THE UNIVERSITY OF READING

Department of Meteorology

Displacement Assimilation for Ocean Models

JAMES R. PERCIVAL

A thesis submitted for the degree of Doctor of Philosophy

June 2008

ABSTRACT

Advances in numerical ocean modelling have led to increasing interest in data assimilation for the purpose of ocean forecasting. The majority of current assimilation techniques ignore the integral constraints inherent in ocean dynamics, and instead locally insert information from observations in an Eulerian sense. It has been shown that such methods can generate spurious circulations by affecting the balance of dynamic ocean tracers.

In this thesis techniques are presented which solve a rearrangement problem to perform data assimilation under an integral constraint, effectively performing assimilation in a Lagrangian sense to correct for an assumed initial phase error. A descent algorithm equivalent to three dimensional variational assimilation is developed for the pure rearrangement problem, based on an advective process acting in pseudo-time. This method is shown to assimilate direct observations successfully, even when observations are only available on a limited subdomain. Two novel modifications which improve the rate of convergence are also discussed.

The method is extended to perform in model time, allowing the use of observations valid over a time window, analogous to the behaviour of four-dimensional variational data assimilation. Comparing this technique to a traditional strong formulation 4D-VAR algorithm shows that the displacement method is superior when errors are due to adiabatic forcing or the effects of model resolution, although strongly limited by the coverage and frequency at which observations are available.

Overall this thesis shows that rearrangement methods are a viable and, potentially, beneficial method for ocean data assimilation, provided suitable data sources can be found.

DECLARATION

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

James R. Percival

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION AND MOTIVATION

"A man with one watch knows what time it is.
A man with two watches is never sure."

"Segal's Law"

1.1 MOTIVATION

Ocean data assimilation is a general methodology for estimating values for dynamically evolving oceanic variables by combining the information contained in dynamical ocean models and in observational data. This combination of information generates an estimate of the current state of these variables – an analysis – which may then be used as an initial condition for later numerical forecasting of future ocean conditions. The complex relationship between the various relevant dynamic variables in the ocean produces a chaotic pattern to the growth, with time, of measured differences between forecasts generated from similar initial conditions. Hence the accuracy of analyses is one of the key factors limiting the accuracy of ocean forecasting as a predictive tool.

Many current data assimilation methods are based upon the concept of a 'most-likely' estimate, a representative state that has the greatest probability of having occurred, given previous forecasts and the result of all available observations. The analysis thus represents the modal state of any forecast ensemble and, it is argued, the best viable choice for an initial condition for future deterministic forecast. Analogy can be drawn with the parameterization of small scale mixing processes. There exist myriad possible arrangements of a dynamic variable on scales well below that being modelled. However if an assumption is made as to the probability density of the various arrangements (as is made for example in the mathematical modelling of Brownian motion) and the most likely estimator of the states assumed true (in this case through application of the central limit theorem), then a modelling assumption is generated (Fickian diffusion) which is applicable on the scales of interest.

Such statistical methods require assumptions on the distributions to be represented. In the case of data assimilation this requires the spatial covariances within variable fields and the correlations in the behaviour between different variables. Generally there is not enough data to estimate these accurately. Often these distributions are chosen on an *ad hoc* basis to represent known physical relationships, such as geostrophy.

An alternative approach is to apply these relationships directly, as expressed in the dynamic equations, when creating an assimilation method. This is similar to the philosophy behind the introduction of a variety of 'adiabatic' techniques for sub-grid parameterizations, such as the Gent and McWilliams (1990) method. These adiabatic parameterization methods have been found (Danabasoglu et al., 1994; Böning et al., 1995) to give better fit to observations and more stable results when used for large-scale ocean modelling than schemes which pay less heed to conserved properties. The subject of this thesis is a discussion of a similarly adiabatic approach in the context of ocean data assimilation and the ocean inverse problem.

This chapter begins with a short summary of the history of the ocean inverse problem and the methods that have been applied for its solution. There follows a review of the various sources of data for the Global Ocean Observing system, with particular reference to those currently applied for the ocean data assimilation problem. Some of the current open topics of research in the field of ocean data assimilation are then listed. Finally the key questions to be addressed within this thesis are stated and the format of the thesis as a whole laid out.

1.2   INVERSE PROBLEMS IN THE OCEAN

For a more comprehensive overview of traditional data assimilation methods, with particular reference to the atmosphere, see the textbooks of Daley (1991) or Kalnay (2002). Ghil and Malanotte-Rizzoli (1991) contains a detailed, if slightly dated, discussion of current methods in data assimilation for geophysical fluids. A book by Wunsch (1996) develops the discretized inverse problem in the context of ocean dynamics and discusses the difficulties of the state estimation problem.

1.2.1   DISCOVERING THE OCEAN

The primary interest of meteorologists, dating back to the first synoptic charts drawn up by Admiral Robert Fitzroy in the 1860s, was to accurately determine the transient state of the atmosphere. Climatological data represented primarily a useful 'first guess' for the more important question of the current weather patterns. In

keeping with this philosophy, the goal in atmospheric data assimilation has typically been to use the information contained in meteorological observations to produce an accurate estimate of the current state of the atmosphere - an analysis field.

In contrast with this, work in the oceanographic community was historically directed towards coupling modelling, theory and observations to determine an estimate of the mean state of the ocean, often termed state estimation or inverse modelling. The difference between these two viewpoints is predominantly in terms of timescale. The aim of state estimation is to recover the long-term time average of the ocean state from long, gappy, series of ocean observations, with the available observations often widely separated in space and time. A necessary condition for this method to succeed is the existence of a well-defined stationary mean state, upon which intra-decadal variations appear as small-scale noise. Data assimilation is intimately coupled with the idea of forecasting the magnitude of the transient magnitudes of these small-scale deviations. Early oceanographic work was focused on the area of state estimation for several principle reasons:

(i). Accurate measurements of dynamically important oceanic variables such as temperature and salinity were (and to a certain extent remain, see section 1.4) difficult and expensive to obtain, hard to record and time-consuming to collate. This means that the historical data set available to investigators is small, with poor temporal and spatial coverage, especially in the Southern Ocean.

(ii). It was for a long time believed, following the theory of Ekman (1905), Sverdrup (1947), Munk (1950) and others, that the observed ocean state could be almost entirely explained through a combination of largely static flow in the interior maintaining geostrophic and hydrostatic balance, with primarily wind driven circulations near the surface. This concept of steady ocean currents, "Rivers in the Sea"[1], assumes that the transient signal is small and dynamically unimportant, except possibly for the case of the external wind forcing, which may itself be derived from meteorological observations.

(iii). The ocean has a much smaller Rossby number than the atmosphere in most areas and situations. This non-dimensional number governs the length scale at which geostrophic balance holds. This means that significantly higher spatial resolutions are required to model important mesoscale features, such as

---

[1]The use of this analogy has become something of a cliché in oceanography. Although featured by as illustrious a figure as Fitzroy (1861) in explaining ocean currents to a general scientific audience it was equally criticized by Nanson (1907) as naive. The modern consensus is uniformly negative (Parker, 1998).

mesoscale eddies or transient displacements of the thermocline, which are of interest to end-users such as fisheries and navies. Until recently limitations on available computing resources were such that standard models could not run with grid spacing such that they could resolve these length-scales.

(iv). The observational and model data available at the time supported the idea of a quiescent ocean. In fact the ocean was typically believed laminar in its interior, until this assumption was cast into doubt by accumulated observational data in the 1970s. The unprecedented large-scale and near synoptic observations of the World Ocean Circulation Experiment (WOCE) from 1990 until 1998 exploded this myth, with turbulent structures and great temporal and spatial variability being detected with ubiquity.

1.2.2   A Brief History of State Estimation

Among the earliest attempts at what may be termed ocean state estimation, aiming at determining the long term ocean state, are the discovery of the location and approximate strength of the ocean surface currents by various groups of seafarers, starting with the discovery of the Canary and Guinea Currents by Portuguese sailors in the 15th century (Krauss, 1996). Differences were observed between a ship's apparent positions calculated from their course and speed through the water, 'dead reckoning', and the ship's actual position, as derived from celestial observations. These differences were then ascribed to drift due to surface currents, with repeated observations allowing the flow to be mapped.

This is a specific example of the general class of inverse problems. Both the data assimilation problem for ocean forecasting and state estimation of the ocean fall into this classification. Given a general mathematical operator, the black box of high school mathematics, accepting inputs from one set of data and mapping them in some fashion consistently to another, the outputs, we may define two different classes of problem:

Problem 1.1 (The forward problem) *Given a known input, derive the expected output.*

In the case of the Portuguese sailors this would be the navigational problem of calculating an estimate of the ship's course from the known currents and the observed bearing. Depending on the complexity of the operator in question such problems may be computationally expensive. Nevertheless provided that the operator is known over the whole range of possible inputs, is not pathologically badly

behaved and that the input in question lies in this domain then a unique solution is guaranteed to exist. If our knowledge of the operator is incomplete or approximate then there may exist multiple consistent solutions, but a probabilistic or partial solution may still be found. If the input in question lies outside the domain of the operator then an extension to a larger domain may be chosen and a value still recovered.

Problem 1.2 (The inverse problem) *Given a known output, find the input(s) which generate(s) it.*

This is the problem of calculating the currents from time series of navigational data. For a non-trivial, non-linear operator this is typically much harder, both computationally and philosophically than the forward problem. If, as is typical in physical oceanography, the dimension of the output (the observations) is smaller than the dimension of inputs (the ocean condition) then the problem will be underdetermined and there may exist multiple consistent solutions. If the output lies outside the range of the operator then there may be no consistent solutions at all. In either of these cases the problem must be modified before it can be solved and in general only weak solutions will exist. For example, in the case of linear maps it is possible to find an input which maps to a point nearby to any output in a least squares sense. In this sense any form of data assimilation or state estimation may be viewed as a suitably constrained regression, fitting a curve representing the analysis trajectory through the points provided from observations under some goodness-of fit criterion, which commonly depends on previous analyses.

Among the earliest rigorous inversions in modern oceanography were the box inversions of Wunsch (1977), and others. The equation set inverted here consists of steady state conservation equations for mass, salt and heat, as well as the integrated thermal wind equations, where the vertical integration has been performed assuming a reference level where horizontal density gradients vanish. These equations are applied across pairs of station observations, where stations are situated along survey lines forming a boundary to the 'box' of interest. The aim is to obtain an estimate of the reference level velocity (and thus through the thermal wind relation the full velocity field) as well as the balancing tracer distribution.

A similar method originating at around the same time is the β-spiral method of Stommel and Schott (1977) and others, which uses the same basic equation set, but instead rearranges these into equations for potential density and relative vorticity. These equations are imposed as local, horizontal constraints to give a series of 3D velocity vectors at the reference level. See Olbers and Willbrand (1985) for a review of the various flavours in which the method came.

1.2.3   OCEAN DATA ASSIMILATION

Recently the advances in computer power, coupled with improvements in parameterizations of sub-grid processes and improvements in the physical understanding of ocean dynamics have led to the creation of successively more realistic ocean General Circulation Models (GCMs). There have been a series of attempts at ocean state estimation through the application of adjoint models for the steady part of ocean GCMs (Martel and Wunsch, 1992; Marotzke and Willebrand, 1996; Stammer et al., 2002) using a variety of models and techniques in order to iterate to a steady solution.

In particular, the consortium Estimating the Circulation and Climate of the Ocean have attempted a concerted effort at exploiting new data sets from the WOCE using an advanced Massachusetts Institute of Technology Ocean Model (Marshall et al., 1997) to coupled to automatic differentiation tool to generate the adjoint code. Wunsch and Heimbach (2007) provides an overview of both the mathematical techniques and assumptions involved and the various key results and applications the code has been put to. These range from regional studies (Ayoub, 2006) to estimates of the errors in the model and observations (Stammer, 2005) to estimates of internal mixing in the ocean (Ferreira et al., 2005).

At the same time methods have been developed for data assimilation in ocean forecasting, in analogy with modern weather forecasting through Numerical Weather Prediction (NWP). Interest has grown in the forecasting problem with the realisation that the ocean represents the long-term memory of the climate system and as such the best hope for accurate long-range weather forecasts. It is this interest in deriving the transient state of the ocean system at high resolution from limited data in the near past which we term Ocean Data Assimilation, following the terminology of Wunsch (1996) and by analogy to the nomenclature used in atmospheric science.

In fact, many of the techniques used to provide initial and boundary conditions to these forecasting systems are derived from the Data Assimilation techniques developed for NWP, particularly variations on the Best Linear Unbiased Estimator (BLUE) analysis and three or four dimensional variational methods, based on a statistically optimal minimization of predicted error variances. In the following section we discuss the basis of these techniques.

1.3   TRADITIONAL DATA ASSIMILATION METHODS

The theory behind the derivation of the BLUE equation forms the basis for many modern assimilation techniques, including the Kalman filter, Kalman smoother, 3D-VAR and 4D-VAR techniques. Many operational ocean forecasting centres use variations on these methods. For example, the UK Met Office Forecasting Ocean Assimilation Model (FOAM) uses an Analysis Correction Method (ACM) (Lorenc et al., 1991) to assimilate temperature and salinity profiles. The basis of the ACM is an iterative approximation converging to the BLUE analysis described below. The European Centre for Medium-range Weather Forecasting (ECMWF) uses a 4D-VAR scheme for its operational ocean model, run to assimilate observations over a ten day time window.

No single standard exists for notation in the combined context of data assimilation and numerical modelling. The notation used in this section is based on that prescribed in Ide et al. (1997)[2] and a similar ethos is used when describing assimilation methods in succeeding chapters.

1.3.1   THE BEST LINEAR UNBIASED ESTIMATOR

At a given time, define the (discretized) state of the system of interest, represented by a finite dimension state vector, x_t, over the spatial domain of interest. This is the 'truth' which we wish to find. In the case of ocean forecasting this may be the three dimensional velocity, pressure, density, temperature and salinity given at points corresponding with the model grid of an ocean GCM. Suppose too there exists an initial guess at this state, perhaps from a previous run of the forecast model. Denote this background state by the state vector x_b. Finally let vector, y, contain all available observations assumed valid at the time of interest. Here we assume that the observed variables represent a simple linear combination of the quantities contained in x_b. The background and observation error vectors ϵ_b, ϵ_o, represent

---

[2]The primary difficulties in standardizing notation between the two fields are:

  i. The tendency of dynamicists and modellers to express equations in the continuous notation of partial differential equations, even when it is discretizations of these which are solved, whereas data assimilation is usually presented in the language of linear algebra.

  ii. The tendency to use x as a variable in Cartesian or locally Cartesian space when in the Ide notation this is reserved for the system state.

An entirely consistent and unambiguous notation remains difficult, but the lack of such notation may be limiting the understanding of the data assimilation process in the general scientific community.

differences from the true state,

$$\boldsymbol{x}_b = \boldsymbol{x}_t + \boldsymbol{\epsilon}_b,$$

$$y = \boldsymbol{H}\boldsymbol{x}_t + \boldsymbol{\epsilon}_o,$$

where $\boldsymbol{H} : \mathbb{R}^n \to \mathbb{R}^m$ is the observation operator giving the observational data expected for a given state vector. Under certain conditions Bayes' Theorem allows us to relate a prior estimate of the probability of an event occurring to the posterior conditional probability of the event having occurred, given the result of some set of observations.

Theorem 1.1 (Bayes' Theorem) *For stochastic events X,Y providing the marginal probability of a given observation, $P(Y = y)$, is non-zero, then the posterior probability distribution $P(Y = y|X = x)$ satisfies the relation*

$$P(X = x|Y = y) = \frac{P(Y = y|X = x)P(X = x)}{P(Y = y)},$$

given a prior probability distribution for the state x, $P(X = x)$, and a conditional probability for detecting given observations from a state x, $P(Y = y|X = x)$.

This result may be derived directly from the definition of conditional probability. A full discussion of Bayesian statistics lies outside the scope of this thesis, but for further details the reader is directed to, for example, Lee (1989). Using Bayes theorem it is possible to find a state which represents the maximizer of the posterior probability density. Given unbiased data (so that the expectation of the error for both background and observation vectors is zero) this state minimizes the variance of the analysis error, independently of the data themselves. For suitable distributions, such as Gaussian error distributions, this will be the maximum likelihood estimator for the analysis state.

Introducing the notation $\boldsymbol{B}$, $\boldsymbol{R}$ for the covariance matrices of the background and observation states respectively this analysis state is found (Daley, 1991)[Section 2.2] to be the solution of the BLUE equation,

$$\boldsymbol{x}_a = \boldsymbol{x}_b + \boldsymbol{B}\boldsymbol{H}^T \left(\boldsymbol{H}\boldsymbol{B}\boldsymbol{H}^T + \boldsymbol{R} \right)^{-1} \left(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{x}_b \right). \tag{1.1}$$

Calculating this state thus requires knowledge of the error covariances (as well as expectations, to ensure the data is unbiased). The uniqueness of $\boldsymbol{x}_a$ as a minimizer of the error variance depends on whether the observation operator can distinguish between any two model states. Note that there will not generally be enough in-

formation *a priori* to determine $\boldsymbol{R}$ and $\boldsymbol{B}$ accurately. Instead these are either calculated from model data (especially common in the meteorological community) or are assumed to act only as a smoothing operator to spread the information in limited observational data across some limited local region, weighting towards the location of the observation, that is it acts as a regularization term in a generalized least squares regression. This second view is probably the more common in the oceanographic community, see for example Bennett (1992).

1.3.2 RELATED METHODS

A family of algorithms such as the Successive Correction method (Daley, 1991)[Chapter 3] and the Met Office's ACM reformulate the BLUE analysis as the limiting solution of a set of iterable equations. This allows an incomplete, non-converged, solution to be applied as an initial condition for forecasting, sacrificing accuracy for speed. The Kalman filter method extends the function of the BLUE analysis state into the problem of finding choices for the covariance matrices based on data. These matrices themselves become variables in the system with a new matrix calculated on each assimilation cycle, as new model and observational data becomes available. Sorenson (1970) has a nice discussion of the filter's connection with least squares regression, while Evensen (1994) discusses the basic filter as well as several extensions in an Oceanic context. The Kalman Filter can itself be extended into the Kalman Smoother (Cohn et al., 1994), which sweeps over data available over multiple time levels and thus assimilates observations backwards in time using past information as well as forwards.

The BLUE solution can also be viewed as the minimiser of a function representing the logarithm of the analysis error probability density function,

$$J := \log P(\boldsymbol{\epsilon}_a) = \underbrace{(\boldsymbol{x}_b - \boldsymbol{x})^T \boldsymbol{B}^{-1}(\boldsymbol{x}_b - \boldsymbol{x})}_{\mathcal{J}_b} + \underbrace{(y - H\boldsymbol{x})^T \boldsymbol{R}^{-1}(y - H\boldsymbol{x})}_{\mathcal{J}_o}.$$

The 3D-VAR assimilation methods attempt to find this minimizing state directly using gradient information on $\mathcal{J}[\boldsymbol{x}]$. For systems with linear observation operators the solution found is identical to the BLUE solution and the method is another iterative algorithm for the BLUE analysis; however, the variational approach allows more readily for the extension to cases with non-linear observations than the equivalent extension of the BLUE method, the Extended Kalman Filter. By extending the definition of the observational error density into one including a temporal factor we obtain the 4D-VAR method. Under certain conditions this produces identical

analysis fields over a particular set of data to the Kalman filter. A limit to the utility of this approach is that these variational methods do not produce explicit updates to the covariance or weighting functions, as the Kalman Filter methods do.

1.4   THE OCEAN OBSERVING SYSTEM

A key ingredient in the cookbook of both data assimilation and state estimation, that is estimating the transient and time mean ocean state, is the available observational data set itself. For an assimilation scheme to be effective it is important to understand the source of the data which is to be assimilated and possible sources of error in the data stream. Traditionally observations of the ocean have been far more scarce than those of the atmosphere. This is partly because humans are land-dwelling, with a resultant interest in extremes of weather, while distant from the deep ocean. Moreover the extreme pressures of the ocean deeps and the difficulty and expense of transit over much of the ocean has hampered maritime research. Long-term data sets are limited to tide-gauges and ship-based temperature measurements from naval vessels and voluntary observing ships (VOS).

1.4.1   SHIPBOARD OBSERVATIONS

The greatest proportion of the ocean data available from the VOS consists surface measurements, especially of surface water temperature. This form of data makes up approximately 90% of the historical ship-borne temperature record. Measurement technology has improved markedly from the days when temperatures were collected by a sailor using a bucket and thermometer[3]. Advances in instrumentation and communications technology now allow the automated recording of the temperature of water drawn in to cool the ship's engine, with measurements having a claimed random error of roughly $1.0°$C and systematic biases of $0.5°$C (Kent et al., 1999). This provides a measure of sea-temperature at depths of 3–10m in the shallow, well-mixed layer immediately below the surface.

The development of the inverting thermometer allowed VOS to measure ocean temperature profiles accurately at depth. This technology has improved to the level of the expendable bathythermograph (XBT), a single use probe used in the Ship of Opportunity (SOOP) scheme. As the XBT descends it collects measurements of the vertical temperature profile and pressure which it then transmits back to the launching ship. Approximately 7000 VOS and 140 SOOP currently contribute to

---

[3]Differences in the method of collection and even the type of bucket used have caused considerable difficulties in studying trends in these records (Smith and Reynolds, 2002).

the Global Ocean Observing System (GOOS). The reported positions of XBT deployment in the month of June 2006 is shown in Figure 1.1A. While lying predominantly along popular shipping routes these do also partition many of the oceans, notably with measurements across the North Atlantic and Drake's Passage. This allows for estimates of inter-ocean transports though inversion of the thermal wind equations.

Shipboard, indeed any, salinity measurements are much scarcer, reflecting primarily the unavailability of suitably simple measurement technology. Historically studies have required collection of water samples at depth and then processing through evaporation to determine the total dry salt content. Modern techniques rely on changes to the electrical conductivity of sea-water with ion-content, allowing for in situ salinity measurements. The use of conductivity, temperature and depth (CTD) probes have increased the accuracy of measurement and the speed of processing so that measurements are available in near-real time. These probes are predominantly deployed from research vessels and the number of CTD probes deployed as part of the regular observation network remains small and almost exclusively deployed along prescribed hydrographic sections, in-keeping with the state estimation methodology.

1.4.2   FIXED BUOY, DRIFTER AND FLOAT MEASUREMENTS

The majority of in situ measurements in the GOOS now come from unmanned systems, both arrays of fixed buoys and fleets of mobile Lagrangian floats and drifters (see plots B and C of Figure 1.1). The ATLAS moorings of the TAO/TRITON array provide the basis of the equatorial observing system. They provide accurate temperature measurements at ten levels, from near-surface bulk temperatures at 0.3–1 m to measurements below the thermocline at 500m depth. A few fixed moorings also provide sea-surface salinity measurements.

Fleets of drifting buoys give data on the wind-driven circulation in the mixed layer. The surface-floating buoy is typically attached to a kite-like drogue, so that the path of the buoy is primarily controlled by the near-surface currents, rather than the wind conditions at the surface. The drifter reports near-surface temperature and location via a satellite link. Drifter performance is dependent on the precise type of drogue fitted (Monahan and Monahan, 1973) and a necessary stage of blacklisting is to remove those drifters which appear to have lost their tails.

The ARGO network (Gould and Turton, 2006) will soon be complete and will consist of approximately 3000 free-floating probes. Each float passively tracks with the local flow at a prescribed depth for a preset time period, before returning to

the surface to report its position using the global satellite positioning system. While surfacing the probe records depth profiles of temperature and salinity, available on the timescale of once a fortnight.

Time-integrated information on the flow fluid parcels at the float's prescribed depth of operation is potentially available through tracking the probe's reporting positions. Similarly a variety of sonar tracked passive float systems give information on smaller time and spatial scales. However due to the difficulty of instrumentation there is very little direct information available on real-time Eulerian currents in the GOOS, apart from the five Acoustic Doppler current profilers and the neighbouring point Doppler current meters deployed along the equatorial Pacific in the TAO array. These profilers and meters provide column current information at accuracies of up to ± 5 cm s^{-1}.

1.4.3   REMOTE SENSING OBSERVATIONS

As well as providing communications links with other systems, satellites provide remote sensing information on the ocean state. Due to the high rate of attenuation of the electromagnetic spectrum by liquid water, satellites are limited to surface observations of the ocean. Microwave and Infra Red systems provide complementary methods for detecting the skin sea surface temperature (or more accurately the induced brightness temperatures), with near global coverage and at varying resolutions. The first proper climatological satellite records of SST date back to the launch of the first Advanced Very High Resolution Radiometer (AVHRR) instrument on the TIROS-N satellite in October 1978 (Barton, 1995). This instrument and its successors allowed algorithms based on differential absorption by water vapour to be used to correct for the moist atmosphere. Current claimed accuracy is a root mean square error of ± 0.5°C for the AVHRR/3 instrument on NOAA-15 (Li et al., 2001) or ± 0.3°C for nighttime SSTs on MODIS. Coverage is near global though containing holes in places due to cloud coverage or other processes. The bottom plot of Figure 1.2 shows the number of observations available in a 5 day period on 1° binned gridboxes, which shows the tropics notably better sampled than mid-latitudes. Diurnal processes can cause large differences between the observed skin temperatures and the modelled bulk temperatures, perhaps as large as 1°C (Katsaros, 1980). Retrieval systems and cloud detection algorithms to remove contaminated signals can also lead to large biases in satellite observations (Uddstrom et al., 1999).

Satellite measurements in the visible spectrum and near infra-red provide information on the concentrations of oceanic chlorophyll and thus of ocean primary production. Instruments such as the SeaWiFS sensor (Lavender and Groom, 1999)

provide global coverage over a two-day interval at a resolution of up to 1.1 km. Approximate conservation of chlorophyll then gives data on surface currents and conditions. Scatterometer instruments allow the retrieval of surface wind conditions over the ocean.

Future satellite missions also promise microwave detection of sea surface salinity(SSS). The European Soil Moisture and Ocean Salinity and US Aquarius satellites claim (Koblinsky et al., 2003) a potential accuracy of 0.2 psu at resolutions of 50 km by inverting the relation between the 1.4 GHz brightness temperature, SSTs, surface roughness and SSSs.

Radar altimeters deployed on the ERS, TOPEX and JASON-1 systems provide information on the variation with time of sea surface height and may also allow diagnosis of the steady component, provided the geoid can be estimated with more accuracy by programs such as the CHAMP, GRACE and GOCE projects. The first radar altimeter tests were performed from Skylab in 1973. The basis is simply that the return time of an earthwards directed microwave pulse from space is dependent on the height of the transceiver above the reflecting surface. Later advances provided improvements in orbital tracking and the removal of artifacts due to surface conditions such as wind speed, tidal effects and surface waves. See Wunsch and Stammer (1998) for a detailed discussion of the issues involved. Currently the estimated uncertainty in a single altimeter pass is 4.1 cm for TOPEX or 2.5 cm for JASON-1. The top plot of Figure 1.2 shows 10 days worth of altimetry data from a single instrument, illustrating the near global coverage of the data, particularly in comparison to the in situ data sources.

1.5   CURRENT ISSUES IN OCEAN DATA ASSIMILATION

Any academic thesis must address its topic in relation to the other open questions of research in the field of work in question. To this end this section consists of a survey of current work in some areas of ocean data assimilation and the questions that remain unanswered.

1.5.1   ASSIMILATION OF SATELLITE ALTIMETRY

As can be seen in Table 1.1, remote sensing data from satellites now represents by far the greatest proportion of observational data in the GOOS and, as may be seen in Figures 1.1 and 1.2, has the nearest approximation to global, differentiable, coverage of any data source. Although the opacity of seawater to EM radiation limits observations to the sea-surface, satellite altimetry data contains, through
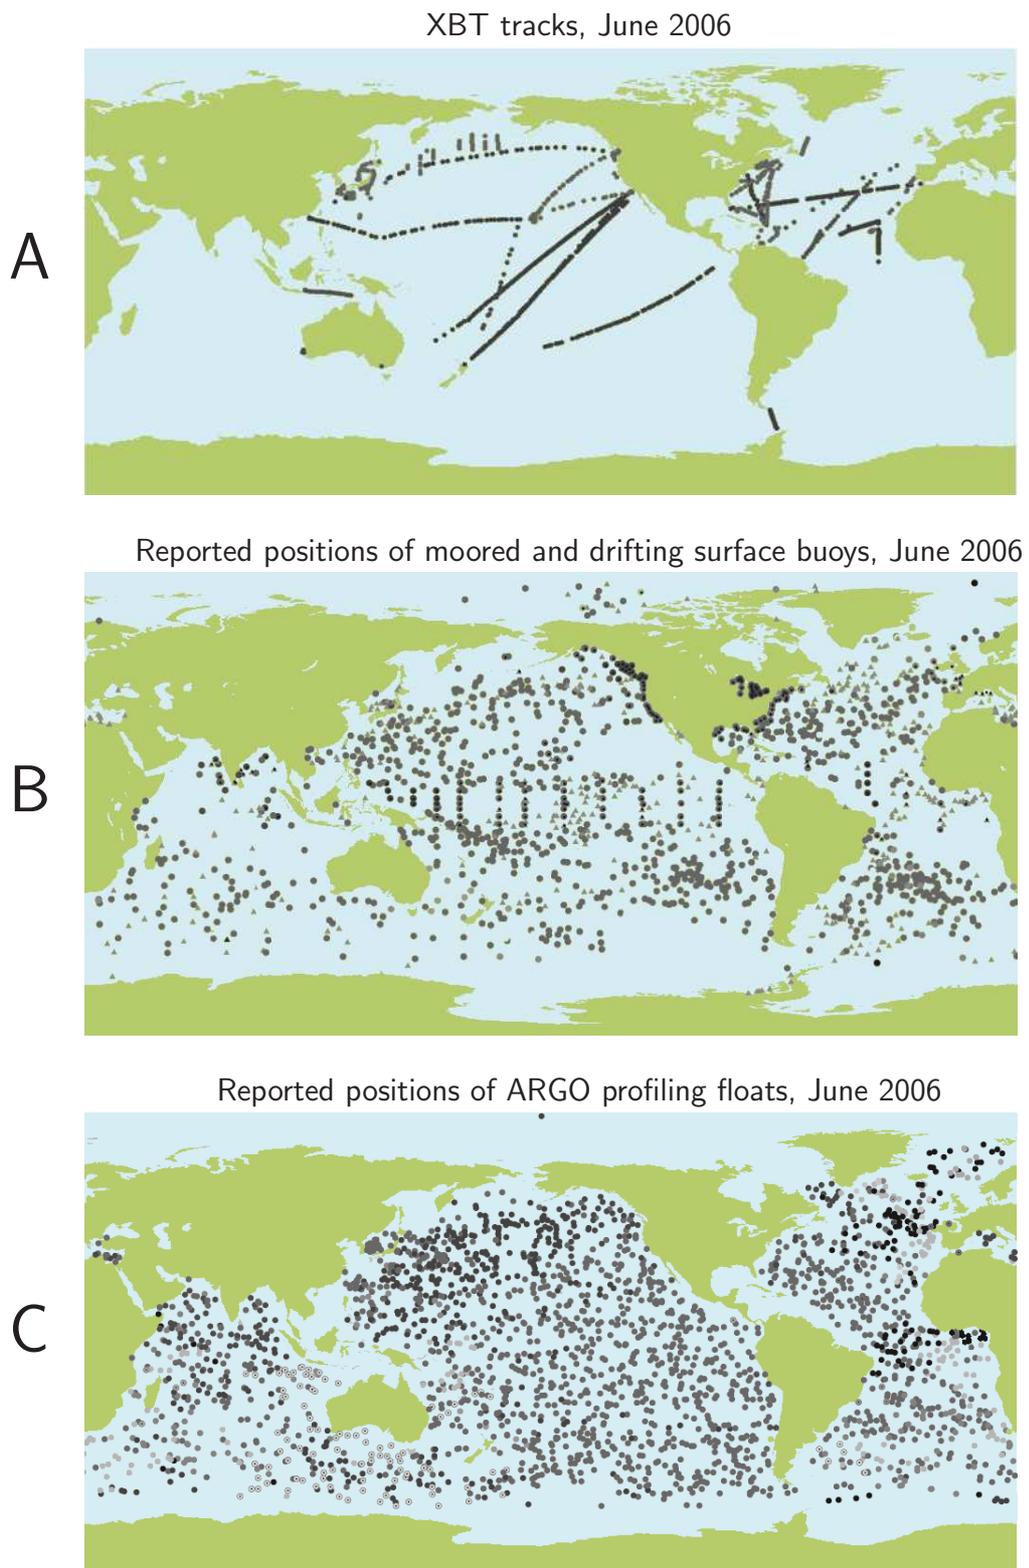
Figure 1.1: Distribution of in situ observations provided to the GOOS
Plots of the locations of individual observations from various sets of instrumentation for one month in 2006.
Data and figures taken from the NOAA Office of Climate Observation
http://www.oco.noaa.gov/
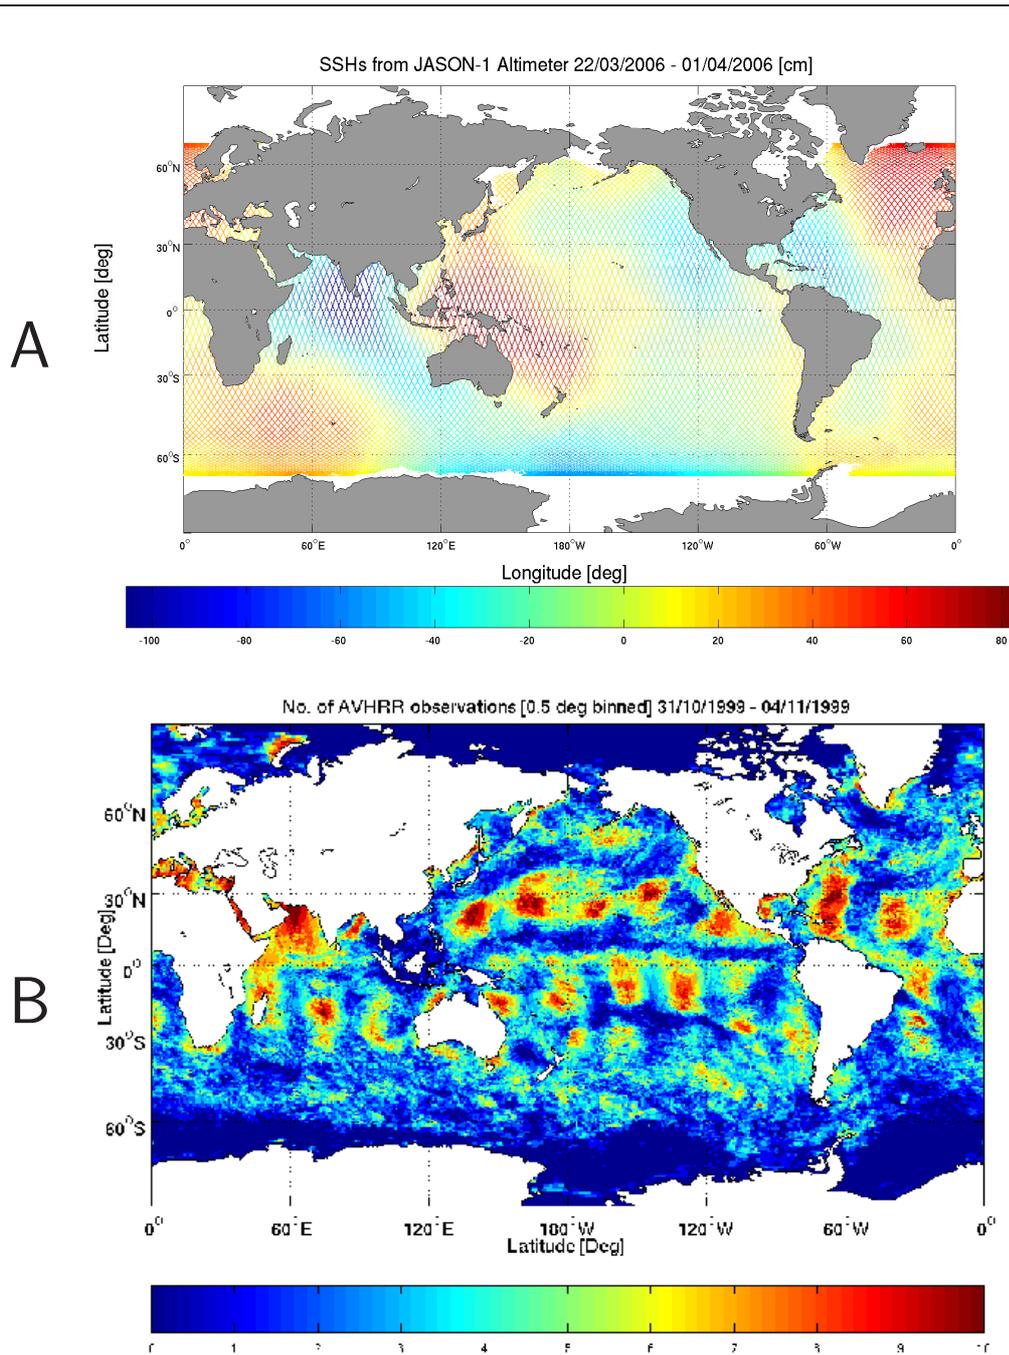
Figure 1.2: Distributions of satellite observations in the GOOS

One weeks worth of satellite data: (A) sea surface altimetry from a single satellite (B) sea surface temperature from multiple satellites. Compare the global coverage with the gappy information from in situ instruments of Figure 1.1.

Figures generated from data taken from the NOAA Satellite and Information Service (JASON data) and National Oceanographic Data Centre (AVHRR data)
http://ibis.grdl.noaa.gov/SAT/SAT.html
http://www.nodc.noaa.gov/woce_V2/disk13/index.htm

Observation Platform	Variables	No. of units	No. of Obs. (Monthly)
DBCP drifting buoys	SST+SLP	1,242	821,413
Moorings - coastal	Temperature profiles	104	
-Tropical	Temperature profiles	78	
+ADCP	current profiles	5	
VOS	SST + meteorology	1,070	52,338
SOOP - Ships		62	
High density lines	T profiles, SSTs	22	
Frequently occupied lines	T profiles, SSTs	17	
Total XBTS Deployed			1088
ARGO Floats	T, S profiles	2,483	
JASON altimeter	SSH	1	1,600,000
AVHRR obs	SSTs	5	2,800,000

Table 1.1: Approximate number of observations from various sources in the GOOS
  This table again shows the vastly greater number of satellite observations available in
  comparison to the in situ data. Note that a single satellite provides twice the number
  of observations of sea surface height as the fleet of drifting buoys provides sea level
  pressure, although obviously at a significantly increased unit cost.
                  Data collated from `http://www.oco.noaa.gov/`,
                      `http://www.bom.gov.au/jcomm/vos/` and
      `http://www.metoffice.com/ravi/networkperf_vosships.html/`

geostrophic and hydrostatic balance, valuable information on both the near-surface
currents and the horizontal gradient of the column density profile. Unfortunately
most modern ocean models are 'primitive equation' level models which generate
SSH (or equivalently surface pressure for those models run in 'rigid-lid' configura-
tion) only as a diagnostic variable. This means the observation operator, as intro-
duced in Section 1.3.1 becomes a nonlinear function of the state variable and the
assumptions of the iterative approximations to the BLUE analysis may fail.

Because of this status as a non-model variable, altimetry data is not usually be
directly assimilated into models in the manner of temperature and salinity obser-
vations. Instead the data must be viewed as a proxy for density errors which are in
turn partitioned into increments in the temperature and salinity states, again based
on an assumed relationship between temperatures and salinities. A number of dif-
ferent approaches have been suggested in the literature. One approach follows
variations of the same statistical optimality argument as that of the BLUE equa-
tion. These methods are based on various local (Oschlies and Willebrand, 1996)
or non-local (Hurlbert et al., 1990) correlation matrices, which relate SSH to the
model prognostic variables of temperature, salinity, pressure and current velocities,
with different studies assuming different relative importance of the various rela-

tionships. Without very large numbers of observations of the prognostic variables in relation to the altimetry measurements these static correlations must themselves be derived from numerical models, with attendant possibility of bias.

An alternative approach is to insert either a pressure or current correction solely on the surface level (at the point of validity of the altimeter measurement) and let the model dynamics themselves transfer the surface anomaly into the ocean interior (Haines et al., 1993). This approach is suited to the aims of state estimation but leads to large spurious signals during the model adjustment which introduce added noise into the system and can even destroy the model circulation pattern. The Cooper and Haines (1996) method, described in detail in section 3.3.1, is an extension of this approach, in which an adjustment is found that generates the full pressure increment required at the top layer, while preserving both column water mass properties and the deep geostrophic structure. The work in this thesis is related to this issue, which is one of assimilation under an integral constraint, in attempting to propagate the observational information available at the surface down into the ocean interior. It is these very ideas which are discussed from Chapter 3 onwards.

1.5.2   ASSIMILATION OF LAGRANGIAN FLOAT DATA

Growing use of, and interest in, Lagrangian floats and drifters such as the Data Buoy Cooperation Panel (DBCP) surface and ARGO subsurface float networks has significantly increased the number, accuracy and spread of SST measurements and of assimilated ocean temperature and salinity profiles. Simply assimilating this data, however, neglects much of the information contained within the reports of probe location returned with the instrument data. Tracking this data provides (time-integrated) information of the Lagrangian currents and transports on the neutral 'parking' density surfaces which the sub-surface probes track.

Such data has historically been used in statistical studies to investigate the mean flow and variance of currents in the deep ocean. However, developments in model resolution and physics now allow the use of this data in correcting subsurface velocity and density fields. Researchers at the University of Miami have directly assimilated information on probe locations through an iterated optimal interpolation technique based on forecast model particle trajectories over a single observation time-step. This has been implemented as a method to increment model thickness in both a quasi-geostrophic layer model (Molcard et al., 2003) and in a more complex primitive equation model (Özgökmen et al., 2003).

In general approaches based on observation linearizations may be expected to

hold when the timescale on which observations of float position are available is small compared to the natural Lagrangian time-scales of the system being observed (Taillandier et al., 2006). For probes such as the ARGO network this ratio is typically of order unity and such linearizations may fail. An alternative approach extends the model equations to include explicit advection equations for the location of each drifter. Under this framework and given the tangent linear model assumption holds (i.e. that locally to any point in time the behaviour of the model trajectory can be accurately described by a linear function) then any number of linear or near-linear assimilation algorithms will hold. This is the basis of the extended Kalman filter technique of Kuznetsov et al. (2003) and a family of techniques based on a 4D-VAR style time dependent variational approach (Kamachi and O'Brien, 1995; Nodet, 2006). Although not directly addressed in the course of this thesis there is an implicit connection between the concepts of assimilating the positional data reported by floats and probes and the rearrangement approach followed in the succeeding chapters. This point is raised again in the future work section of the conclusion.

1.5.3 ASSIMILATION OF SALINITY DATA

Historically salinity observations have only been available from research vessels, due to the complex and expensive instrumentation necessary to obtain uncontaminated water samples and to process them accurately to find their salt content. For this reason updates in salinity values during data assimilation have typically come from either a forced regression towards climatology or from a multivariate assimilation approach, correcting salinity based on the relations with observed errors in other variables, for example the variational approach to correcting salinity from SSH observations and satellite altimetry data of Vossepoel and Behringer (2000) or approaches using SSS and temperature–salinity (T–S) relations based on correlations from climatological data (Vossepoel et al., 1999) or CTD studies (Troccoli and Haines, 1999). This potential under-constraint of a dynamically important variable has been found to lead to the generation of spurious overturning circulations in runs with univariate assimilation, although this effect is reduced under assimilation schemes which preserve the model T–S relationship (Troccoli et al., 2002).

Recent advances in satellite instrumentation and the deployment of the ARGO float fleet have lead to a growing interest in methods for the direct assimilation of SSS and salinity profile data, especially in the presence of temperature data. The surface and near surface observations are particularly valuable, firstly because it is in the mixed layer where salinity can control convection and secondly due to the

current poor representation of mixing and freshwater flux in ocean models. A paper by Haines et al. (2006) describes a method for assimilating salinity profiles onto the salinity fields on isotherms rather than constant depth surfaces. Their argument is that salinity on height surfaces is affected by the fast dynamical processes which move isopycnal surfaces, whereas on isotherms only the slow diffusive processes are important. This method is found to work well when combined with a separate increment from a T–S preserving technique. The conservation approaches discussed in this thesis suggest another approach which we claim should have similar benefits.

1.6   ASSIMILATION OF OBSERVED DISPLACEMENT ERROR

As major motivational factor assimilation of observed displacement error will be discussed in depth. Much of the dynamic activity in the ocean results in coherent structures, for example the separation of warm- and cold-core eddies along western boundary currents such as the Gulf Stream. Due to inaccuracies in both initial conditions and the physics of numerical ocean models these features, when they exist in forecasts, are often displaced spatially from their observed positions. This phase or displacement error is hard to remove using standard Kalman filter techniques (see section 3.1.1), especially in the case of limited observations, so common in physical oceanography. It is the study of this problem and possible solutions which forms the main content of this thesis, although the issues of the previous section, particularly the concept of the T–S relationship from salinity assimilation and the near global coverage of the altimetry observations, help motivate ideas in various sections.

1.6.1   MESOSCALE OCEANOGRAPHY

A wide variety of methods have been created which attempt in various ways to correct errors ascribed to displacement. In oceanography the Cooper & Haines method mentioned in the previous section can be interpreted as reducing the one-dimensional, vertical, displacement error in each water column. The 'contour melding' assimilation routine derived by Mariano (1990) acts by repositioning a finite number of nodal position points which in turn define the shape of contours at various levels of the field being assimilated. The repositioning occurs through a simple weighted sum of the values obtained from a background representation and observations, assumed to cover the entire domain, so in fact likely to be obtained from a retrieval. The 'feature models' of Robinson et al. (1988) act, at the point of initialization of the model only, to insert observed features (primarily jets and eddies)

at the locations found in observational data by inserting objects of a prescribed, analytic, structure with values determined by a limited number of free parameters chosen to match the observations.

1.6.2 THUNDERSTORM MODELLING

Much of the interest in displacement assimilation in the atmosphere has been in models running at what is termed the storm scale, namely at model resolutions of about 1km. On these scales the smoothing effect of traditional data assimilation methods can severely effect the growth and location of extreme features such as convective storms.

The Hoffman et al. (1995) method introduces a separate explicit penalty term governing phase errors due to displacement into the cost-function for data assimilation of SSS/I microwave retrievals of integrated water vapour into a spectral model. This is then minimised with respect to all possible displacements by a variational method similar to 3D-VAR, along with a smoothness term to ensure physical realism of the solution and a 'barrier' term which acts to penalize large scale displacements. Experimentally, (Hoffman and Grassotti, 1996) the method was found to lead to a greater than 45% reduction in the variance of differences against observations when assimilated into ECMWF reanalysis fields.

Ravela et al. (2004) perform a similar displacement assimilation to the Hoffman method, which they term 'Field Alignment'. This again adds specific displacement terms to a 3D-VAR cost function, this time in a spatial formulation, with an explicit smoothness term which penalises the gradient and divergence of the displacements in a Tikhonov (See e.g. Wu (2003)) style formulation. They claim that this arrangement obviates the need for a penalty on the final magnitude of the resultant displacements. In identical twin runs assimilating pressure and velocity into an idealised model the method is found to improve on both the control and runs with 3D-VAR alone in preserving geostrophic balance in the assimilated field and avoiding the spreading of fronts.

Brewster (2002a,b) also attempts to assimilate information on phase through the minimization of a cost-function. Here, however, the minimization is achieved through a 'brute force' shifting method which tests the local effect on the fit with observations of shifting the centre of a control volume of data by a limited amount in all directions. To minimise the shock to the system when the derived optimal displacement field is applied it is introduced as an additional compressible pseudo-wind field which is fixed with respect to time. The method is applied to assimilate Doppler radar readings, radiosonde measurements and surface data into a model of

a severe storm outbreak. Even on this fine scale and considering the compressible nature of the atmosphere it is impressive that the method assimilated information which remained in the system for 3 hours. This suggests that displacement error is a robust concept, even in unlikely geophysical situations.

1.7   KEY THEMES

1.7.1   OVERALL FOCUS

Currently the majority of data assimilation techniques used for ocean forecasting, particularly those in use at operational centres, are Eulerian in nature. That is, they estimate the behaviour of relevant ocean variables at a specific points in space based on the observed behaviour of variables at nearby points and statistical or physical correlations between pairs of locations or variables. This thesis considers the utility of alternative approaches based on Lagrangian arguments regarding water mass properties and the assimilation of information contained in observations of the topology of the flow.

The overall focus of this thesis is:

To understand the benefits and issues involved in using phase error assimilation methods in a general, mathematically rigorous and physically consistent ocean forecasting system.

Here phase error is used to denote, at its most general, errors in the position of features and signals in oceanographic model data, where the magnitude of the feature agrees with observed data.

1.7.2   KEY QUESTIONS CONSIDERED

Within this theme, the key scientific questions which will be addressed in the body of the thesis are:

(i). *To what extent is it feasible to decouple the observed error signal in ocean analyses consistently into signals due to pure phase error and amplitude error?*

We aim for a better understanding of the differences in growth and behaviour of error signals in the phase and amplitude of features present in ocean general circulation models. We hope to diagnose the extent to which the two signals can be quantitatively separated by simple error diagnostics. It is known that many popular Eulerian-based state estimation schemes, including the Kalman filter and BLUE analysis, act to smear extreme features when used to

assimilate data representing phase errors, while most phase correction methods conserve errors in signal amplitude. This separation of signals is thus important since a well-posed method to decouple the two signals allows each form of error to be corrected by an appropriate technique.

(ii). *Given that real oceanographic observations are noisy and limited in scope and number, can a study of novel phase correction algorithms give a quantitative understanding of how observations relate to model data in such methods and what constraints on observations are required?*

We aim to understand how observations combine with the information contained in the background in phase correction methods. This allows comparison with the behaviour of more traditional techniques, which is important in understanding the potential new benefits of applying phase correction techniques in an advanced ocean forecasting system. An understanding of the utility of currently available observational data in phase error correction also tests the feasibility for using phase error correction methods in current operational centres and is important in the planning of future instrumentation and observational networks.

(iii). *Is it possible to develop efficient algorithms which maximise the utility of displacement assimilation techniques, for example in terms of the use of the information contained within observations, in the same way that 4D-VAR methods allow the use of multiple nonlinear observations taken at different times?*

Given that the optimal choice of forecasting technique represents a balance between the cost of forecast production against the value of the information contained within it, it is important to maximise the benefit of any data assimilation method and minimize its cost. For phase error correction methods this requires both the development of efficient algorithms and techniques to maximise the useful information extracted from available observations and a consideration of possible sources of observation.

1.8 THESIS OUTLINE

To consider these questions we investigate the behaviour of various phase correction assimilation methods when applied to idealized problems in a simple Quasi-Geostrophic (QG) ocean model. This simple model allows simulation of mesoscale features and unobserved dynamical variables. We present a variety of new phase error assimilation methods and compare them to traditional techniques. We find that

for suitable assimilation problems these new techniques do indeed provide some benefit. To quantify this improvement new error measures are also developed.

Chapter 2 presents a discussion of the principles, advantages and limitations of the identical twin methodology for investigating data assimilation techniques, reviews the general issues involved in numerical ocean modelling and introduces the particular implementation of a quasi-geostrophic two layer ocean model used for the experimental data presented in this thesis. The equation set is described with reference to its ability to represent real, or realistic, structures and features of the large-scale ocean circulation, thus presenting its suitability as a testbed for oceanographic data assimilation techniques. Different methods for generating model runs with qualitatively differing features are discussed, including both the relevance to current issues in operational oceanography and the relation to the theoretical displacement assimilation techniques discussed in succeeding chapters.

Chapter 3 begins by presenting a discussion, linked to the adiabatic assumption introduced in the previous chapter, for the use of assimilation techniques which conserve the integral material properties of ocean fluid parcels when performing assimilation with limited data. These conservation constraints are then connected to the related problem of assimilating phase errors into a system. Metrics on the relative magnitude of amplitude and phase errors in fully observed fields are developed and the behaviour of the metrics investigated for the QG ocean model under the application of a variety of traditional and novel assimilation schemes. An example of one operationally active adiabatic assimilation method, the Cooper & Haines scheme, is implemented in the context of the QG layer model and results are shown for the effect on observed and unobserved variables of assimilating surface streamfunction. The effect of smoothed observations on small-scale features is discussed, with reference to feedback on large-scale circulation.

Chapter 4 introduces a new direct iterative method for phase correcting assimilation, comparable with 3D-VAR. The method is first derived theoretically, then implemented for a hierarchy of idealized ocean problems. Convergence is discussed, with reference to the theoretical limits available through rearrangement theory and an improved 'multigridding' algorithm developed. Finally the dependence of the method on observations is considered.

Chapter 5 considers the extension of phase correction methods into problems with observations available at multiple times. Current 4-dimensional Eulerian assimilation schemes are discussed, with reference to possible adaptation to phase error assimilation. A method similar to 4D-VAR but optimizing an extra advective term is then developed. This method is implemented into the QG model of chapter 2 and the implementation of the resultant adjoint model described, with results of

a suite of model tests.

Chapter 6 discusses the results of more realistic experiments with the 4D phase correction method. These are then compared with analyses derived from a more traditional 4D-VAR scheme, with specific reference to the differing treatment of phase errors. The analysis dependence on observations and initial conditions is discussed. Final results show the improvement possible in these techniques is highly dependent on the source of the initial error in the forecast to be corrected through assimilation.

Chapter 7 considers the key questions involved in this thesis once more and discusses how far they have been answered. The chapter concludes with a discussion of the remaining open questions in phase error correction and a plan for future work that may come from this thesis.

CHAPTER 2

OCEAN MODELLING

"We were the first that ever burst
  Into that silent sea"

<div align="right">

Samuel Taylor Coleridge
The Rime of the Ancient Mariner, 1798

</div>

To reduce the complex set of equations assumed to govern oceanic flows into a smaller and more easily soluble set, a number of approximations must be made. These approximations inevitably act to filter the possible dynamics of the reduced model and to decrease the number of processes represented compared to the real ocean. One such approximation is the adiabatic condition, which notes that under normal conditions quantities in the ocean are predominantly transported though advection. This leads to key integral constraints on the motion of fluid properties that motivate an interest in phase or displacement error.

This chapter goes on to discuss the identical twin methodology to be used to investigate these methods and then details the approximations made in the quasi-geostrophic model used in the numerical experiments in this thesis. Following the derivation of the model and a description of its behaviour we conclude by justifying the choices and methods to be used to generate the numerical model and initial conditions used to form 'true' and 'observational' data in the data assimilation experiments presented in this thesis.

2.1   THE ADIABATIC CONDITION IN OCEANOGRAPHY

The modelling and study of a time-evolving physical system through partial differential equations requires the choice of an integrable equation set in model space. Since it is not technically feasible to model the various interactions of the approximately 10^{37} molecules which make up the ocean, oceanographers commonly make a series of approximations to produce suitable conceptual and numerical models. For a fuller discussion of the methods involved in general model derivation the

reader is directed towards any number of textbooks in geophysical fluid dynamics, particularly that of Pedlosky (1987). Two of the most important and general of these approximations are :

i. The continuum approximation. Rather than modelling individual molecular motions, the fluid is modelled as a continuous medium and it is assumed that the Navier-Stokes equations apply.

ii. Incompressibility. On the horizontal and length-scales of interest to oceanographers changes in density due to pressure variations are negligible. This means the velocity field in the flow is effectively solenoidal (Batchelor, 1967).

These approximations reduce the dynamics down to what are known as the primitive equations:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u}.\nabla\boldsymbol{u} + 2\boldsymbol{\Omega} \times \boldsymbol{u} = -\frac{\nabla p}{\rho} + \boldsymbol{g}_a + \nabla.(\nu\nabla\boldsymbol{u}) + \boldsymbol{F_u}, \tag{2.1}$$

$$\frac{\partial \Theta}{\partial t} + \boldsymbol{u}.\nabla\Theta = \nabla.(\kappa_\Theta\nabla\Theta) + F_\Theta, \tag{2.2}$$

$$\frac{\partial S}{\partial t} + \boldsymbol{u}.\nabla S = \nabla.(\kappa_s\nabla S) + F_S, \tag{2.3}$$

$$\nabla.\boldsymbol{u} = 0, \tag{2.4}$$

$$\rho = \rho(\Theta, S, p). \tag{2.5}$$

Here the momentum term (2.1) is given in terms of $\boldsymbol{u} = (u, v, w)^T$, the full three-dimensional velocity, $\boldsymbol{\Omega}$ is the planetary angular velocity, p denotes the dynamic pressure field, ρ the local density and $\boldsymbol{g}_a$ the apparent gravitational acceleration. The effective kinematic viscosity is denoted by ν and $\boldsymbol{F_u}$ is a residual term containing all sources and sinks of momentum. The two equations following denote approximate Lagrangian conservation of potential temperature, Θ and salinity, S, with diffusion coefficients κ_Θ, κ_S and forcing terms F_Θ, F_S respectively. Equation (2.5) represents a closure of the equation set, relating density to temperature, salinity and pressure.

The penultimate terms on the right hand side of equations (2.1),(2.2) and (2.3) represent diffusive mixing processes acting on momentum, heat and salt content respectively. These terms are necessary both to parameterize processes which occur on length-scales smaller than those on which our approximations hold and for numerical modelling purposes to maintain numerical stability under the model discretization. Unlike in atmospheric modelling the heat and salinity forcings (more generally all material tracer forcings) and momentum forcing term, F_θ, F_S, and $\boldsymbol{F_u}$ vanish, except on boundaries, primarily the upper boundary, since the ocean is

forced from above. This means that over the majority of the ocean the behaviour of material tracers is governed by a pure advection-diffusion equation. Assuming that the diffusive process is isotropic in the horizontal and that horizontal and vertical diffusion are uncorrelated we may may express this constraint for a general tracer, τ by an evolution equation

$$\frac{\partial \tau}{\partial t} + \boldsymbol{u} \cdot \nabla \tau = \kappa_H \nabla_H^2 \tau + \kappa_V \frac{\partial^2 \tau}{\partial z^2},$$

where ∇_H^2 denotes the horizontal part of the Laplacian operator,

$$\nabla_H^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2},$$

and the isotropic assumption and lack of correlation have been used to replace the diffusivity tensor with the two scalar coefficients. There are two key dimensionless parameters controlling such a flow, the vertical and horizontal Peclet numbers,

$$P_H = \frac{UL}{\kappa_H},$$

$$P_V = \frac{WH}{\kappa_V},$$

where U and W are horizontal and vertical velocity scales, L and H are respectively horizontal and vertical length scales. This form of diffusive process, although a common modelling assumption is physically unrealistic, assuming a truly random exchange of tracer particles between fluid parcels. Nevertheless, it is sufficient for our purposes.

The value of the vertical diffusivity coefficients, κ_V has important consequences for the nature of the observed mean ocean stratifications and the processes which act to maintain it. This value has been a topic of discussion in the literature for many years. The seminal paper, "Abyssal Recipes", of Munk (1966) represented the best estimate of the time based on the average value required to maintain the observed abyssal stratification, with later developments discussed in the sequel (Munk and Wunsch, 1998). The present consensus appears to require a global figure of $\kappa_V = 10^{-4} m^2 s^{-1}$, averaged across the entire abyssal ocean, with a lower diffusivity of $\kappa_V = 10^{-5} m^2 s^{-1}$ in the pelagic ocean. This lower figure is supported by both theoretical studies considering the required steady state heat flux across isopycnal surfaces to maintain equilibrium (Walin, 1982) and practical studies considering the rate of mixing of released passive tracers (Ledwell et al., 1993; Polzin et al., 1997), as well as studies considering the microscale structure of turbulence. The

apparent contradiction between these two values is partially closed by the greatly increased mixing observed to occur near ocean boundaries and in ocean canyons. In these areas diffusivities of up to $10^{-1}m^2s^{-1}$ are estimated to exist, with tracer properties often observed as near uniform. Estimates for horizontal (more correctly along isopycnal) diffusivities are primarily from tracer studies and are typically of order $\kappa_H = 1\mathrm{m}^2\mathrm{s}^{-1}$. Note that these figures are all *eddy* diffusivities, taking finite volume fluid parcels as the unit of diffusion, rather than the far smaller molecular diffusivities governing chemical processes. The observed value thus depends on the length scale assumed in the problem.

Considering the mesoscale flow in an ocean basin such as the North Atlantic, where the other dimensional scales may be chosen as

$$U = 0.1\mathrm{ms}^{-1} \quad W = 2 \times 10^{-3}\mathrm{ms}^{-1}$$
$$L = 5.0 \times 10^4\mathrm{m} \quad H = 1 \times 10^3\mathrm{m},$$

gives estimates for horizontal and vertical Peclet numbers of $P_H = 5 \times 10^4$, $P_V = 2 \times 10^3$. Both of these are significantly greater than unity; in fact these numbers are so large that even considering the hundred-fold bigger diffusivities required by numerical models for stability (Bell, 1998) these ratios are still greater than order one.

The dominant process acting on ocean tracer fields in the ocean interior is thus shown to be advection by the mean velocity, both in reality and in numerical models. This suggests that the concept of water mass properties, that is atomic units of fluid, which carry their tracer values with them as markers, is valid and, provided the surface forcings can be observed or modelled accurately, a leading cause of error will be the phase or displacement errors of section 1.6. If this constraint holds, it is questionable whether traditional data assimilation techniques, based on local modifications of tracer fields (and hence effectively a diabatic forcing, especially in the case of univariate assimilation), remain an optimal use of the limited available information in oceanographic observations. It is this constraint, coupled with the availability of integrated information from satellite altimetry, which justifies our interest in conservative assimilation and the linked issue of displacement and phase error, as will be discussed in Chapter 3.

2.2   THE TWIN EXPERIMENT PHILOSOPHY

We now discuss the methodology to be used to investigate these possible new data assimilation methods. One of the fundamental assumptions made in any data

assimilation method is the existence of the underlying, and dynamically evolving, 'true' state, of which both the background forecast and observational data represent estimates. For the Earth's climate system this true state is of course the unknown transient state of the atmosphere and ocean, defined at some semi-arbitrary resolution determined by the forecast model grid or the density of observations[1]. Since the final analysis state represents a combination of the information contained in the background and the observations (a spatially pointwise linear combination in the case of the BLUE analysis of section 1.3.1, from the functional and mapping portions of the respective polar factorizations in the case of the rearrangement theory approach) neither estimate can be reliably used to verify that the resultant analysis lies close to the true state in the appropriate metric. Validation of techniques and approximations becomes difficult without either independent, unassimilated observations or alternative forecasts, generated using intrinsically different approaches to modelling. In oceanography the scarcity (an large spatial separation) of independent observations limits application of the first technique, while habits among oceanographic modellers, coupled with the nature of the dynamical equations being modelled, hinder the second[2].

One technique which avoids these issues when first investigating a new data assimilation method for use in forecasting the climate system is to begin by assimilating only synthetic data from numerical model runs rather than actual measurements. These are the twin experiments, where 'observations' are actually noisy realizations of the model variables from the 'true' run. If the same model is used for the true state and forecast (the perfect model scenario) then this is an identical twin method. The advantages of such methods follow from complete knowledge of the true state, plus the model producing it, as well as control over the sources of noise in the assimilation system. This makes a study of the use of data assimilation methods practical by giving some notion of falsifiability, in the Popperian sense, while still allowing the methods to be applied to realistic problems. Although success in a twin experiment does not guarantee success with real data, the causes for failure are more amenable to study. This is the approach followed in this thesis.

Criticism for twin experiments concentrates on the level of idealization introduce. Any added noise must be drawn from a distribution with prescribed param-

---

[1]This section ignores the twin questions regarding whether such a resolution dependent true dynamic state exists and, if it does, whether such a state is representative of anything in particular. The first question depends upon the philosophical meaning of 'truth' and the second is easily answered by assessing the perceived value of forecasts of this state to interested parties.

[2]The *reductio ad absurdem* of the verification process is to remove all observations from the data assimilation routine. There will then *always* be exact agreement between forecast, analysis and the zero available observations. This is of course facile.

eters. If these parameters are then used in statistical tests methods may appear more successful in the experiments than they later prove in reality. This particular problem may be alleviated by 'forgetting' the actual noise distribution used in the experiments and approximating it in any other tests. However no twin experiment will have the number of degrees of freedom found in real life and the behaviour of models assimilating simulated data may be quite different from that seen when assimilating observations. For example Tziperman et al. (1992a,b) find greater ill-conditioning of the inverse problem using atlas and hydrographic data for assimilation into a GCM than in twin mode.

Twin experiments are no longer as popular as they once were, with research concentrating on the application of novel methods and techniques to assimilate real data and on assimilating new sources of data using standard techniques. This represents the demands of operational centres for methods which, while sufficiently accurate, are guaranteed not to fail catastrophically, as well as the natural human preference for the novel or unusual. This does not obviate the important role for twin experiments however in identifying methods that will fail. A system which cannot, even with tuning, assimilate synthetic data is highly likely to fail when 'released from captivity' and used in operational models.

2.3   QUASI-GEOSTROPHIC THEORY

Following the twin experiment philosophy all data used in assimilation experiments in this thesis, both forecast and observational, are derived from an idealized model, rather than direct simulation and observations of the real ocean. The ocean model used henceforth solves numerically a discretization of the two layer ocean Quasi-Geostrophic (QG) equations in a parameter and forcing domain suited to modelling a large-scale ocean basin, such as the North Atlantic. Such simple models have been well studied in the past three decades in an effort to understand qualitatively the large-scale behaviour (Bryan, 1963; Holland, 1978) and energetics (Berloff and Meacham, 1998; Berloff and McWilliams, 1999) of wind-forced ocean gyres. Since this model is employed as a proxy for the full ocean dynamics it would be well to understand the underlying approximations and assumptions made in its derivation and how these differ from more complex models of behaviour. As such the derivation of the model equations will be discussed in some depth in the following sections.

2.3.1   THE SHALLOW WATER EQUATIONS

The starting point of this derivation is the Navier-Stokes momentum and volume conservation equations for an incompressible, viscous fluid, given in a frame of reference rotating locally with an angular velocity, $\mathbf{\Omega}$ which may vary with latitude,

$$\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u}.\nabla \boldsymbol{u} + 2\mathbf{\Omega} \times \boldsymbol{u} = -\frac{\nabla p}{\rho} + \boldsymbol{g}_a + \boldsymbol{F_u} + \nu \nabla \cdot (\nabla \boldsymbol{u}), \qquad (2.6)$$

$$\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0. \qquad (2.7)$$

Where $\boldsymbol{u}(\boldsymbol{r}, t)$ denotes the full three-dimensional velocity field given in an Eulerian coordinate system[3] $\boldsymbol{r}$ and at time t, p is the dynamic pressure, ρ the fluid density, ν the kinematic viscosity, $\boldsymbol{g}_a$ the apparent gravity (including the effects of centripetal acceleration) and $\boldsymbol{F_u}$ the sum of all other body forces acting on a fluid parcel.

The first important assumption to be made is the shallow water approximation. In the schematic given in Figure 2.1 the horizontal length scale of interest for the basin, L, is on the order of thousands of kilometres, versus a vertical length scale (D) of thousands of meters. The aspect ratio, $\delta = D/L$, is therefore small compared with unity. A first order balance of terms in the expansion in δ of equation (2.7) requires that

$$W \leq \mathcal{O}(\delta U)$$

where $f = \mathcal{O}(g)$ denotes the usual Landau notation (Marsden and Hoffman, 1999) for functions f, g, which implies that the ratio f/g remains bounded in the limit $g \to 0$. Using this information to get a scale balance on the pressure term in equation (2.6) gives the relation

$$\frac{\partial p}{\partial z} = -\rho g + \mathcal{O}(\delta^2). \qquad (2.8)$$

This result is effectively a restatement of the hydrostatic approximation, but here expressed with an explicit limit on the magnitude of the error term[4]. If the remainder term is assumed to vanish and the system is exactly hydrostatic then quickly propagating, small magnitude, vertical inertia-gravity waves are filtered in the possible dynamics, through a modification to the primitive equation dispersion.

Layer models, as the name suggests, simulate the two dimensional evolution of vector fields defined on fixed surfaces, the layers, rather than solving equations for values at grid-points or for elements on a fixed level of a vertical coordinate system.

---

[3]Although, strictly speaking, these equations should be applied in spherical coordinates, locally these can be approximated by Cartesians. Since this is the frame in which the equations are to be solved numerically in the final model, it is in these coordinates the derivation will be presented.

[4]In fact, for the small Rossby numbers common for geophysical flows, the remainder term may be far smaller than this upper bound suggests.

Due to the extremely stable stratification of the ocean and the fact that density is a good proxy for entropy in the ocean, these layers are normally defined as parallel to isopycnal surfaces, allowing across isopycnal (i.e. co-ordinate vertical) diffusion to be neglected. In Section 2.1 the magnitude of this diffusive effect was discussed and shown to be small. Due to the effects of pressure at depth and the slightly compressible nature of seawater it is strictly speaking the relative potential densities of pairs of water parcels (derived from potential, rather than in situ, temperature) that determines the stability of the water column. This means that the locus of the stable heights of a fluid parcel away from its local region, the neutral surfaces, are a complex function of the global density structure and are impossible to find in the real ocean. Jackett and McDougall (1997) derive a practical neutral density variable suitable for model verification and inverse modelling. Figure 2.1 shows a schematic of the difference between layer and coordinate approaches to modelling the vertical structure of fields in an ocean basin. In layer models data is stored on the solid lines bounding the grey domain, while in level models data is stored on the dotted lines. Layer models thus better preserve isopycnal integral and volume properties, but suffer from variable, and thus possibly large, error in vertical resolution.
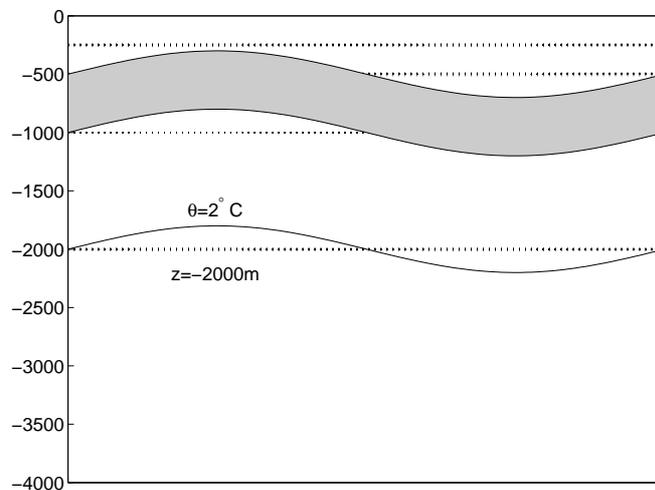
Figure 2.1: Layer systems versus vertical coordinates

A simple schematic showing the difference between layer models, in which the thicknesses of layers, such as the one shaded in grey are stored along with other dynamical properties on a horizontal discretization structure, and level models where properties such as vertical velocity are stored on a vertical discretization divided as by the dotted lines.

The domain of the ocean basin, $\mathcal{D}$, (assumed to be simply connected) is parti-

tioned into N layers, $\mathcal{D}_i$ with $i \in [1, \ldots, N]$, each with boundaries delineated by interfaces $\eta_i(x, y), \eta_{i+1}(x, y)$ and it is assumed that within each layer the potential density is constant, so that

$$\rho(\boldsymbol{r}) = \rho_i \ \forall \ \boldsymbol{r} \in \mathcal{D}_i, \ \text{i.e.} \ \ \forall \boldsymbol{r} := (x, y, z) \in \mathcal{D} \ \text{such that} \ \eta_i < z < \eta_{i+1}.$$

Integrating the relation, (2.8), for pressure change with depth gives, in the limit $\delta \to 0$, an iterative formula for the pressure field at any height within each layer,

$$p_i(x, y, z) = \rho_i g \left(\eta_i - z\right) + p_{(i-1)}[\eta_i], \ \forall \boldsymbol{r} \in \mathcal{D}_i. \tag{2.9}$$

Since this equation is linear in z there is no height dependence in the horizontal pressure gradient terms present in the momentum equations. Velocity fields which are initially depth independent remain depth independent and the equation set within each layer reduces consistently to two z-independent equations for the horizontal components of the velocity in the layer, $\boldsymbol{u}_i = (u_i, v_i, w_i)^T$,

$$\frac{\partial u_i}{\partial t} + u_i \frac{\partial u_i}{\partial x} + v_i \frac{\partial u_i}{\partial y} - fv = -\frac{1}{\rho_0} \frac{\partial p_i}{\partial x} + \nu \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) u, \tag{2.10}$$

$$\frac{\partial v_i}{\partial t} + u_i \frac{\partial v_i}{\partial x} + v_i \frac{\partial v_i}{\partial y} + fu = -\frac{1}{\rho_0} \frac{\partial p_i}{\partial y} + \nu \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) v. \tag{2.11}$$

In generating these equations the Boussinesq approximation has been assumed, that is, variations in density are assumed small compared to some reference density, ρ_0, while the product of density differences and gravitational acceleration remains comparable to unity.

Considering the total mass flux through a small cylinder the height of the layer and of fixed cross-sectional area gives a layer mass conservation equation,

$$\frac{\partial H_i}{\partial t} + \frac{\partial (u_i H_i)}{\partial x} + \frac{\partial (v_i H_i)}{\partial y} = 0, \tag{2.12}$$

where H_i is the layer column thickness, $H_i := \eta_i - \eta_{i+1}$. The three coupled equations (2.10), (2.11) and (2.12) are the layer rotating fluid shallow water equations. Although the original five independent dynamic variables (u, v, w, ρ, p) have been reduced to three (u, v, H) in each layer and the problem set reduced from spanning three spatial dimensions to two[5], this still represents a more complex equation set than can conveniently be solved. Taking the curl of the momentum equations

---

[5]The reduction of dimension within each layer is compensated by a corresponding increase in the total number of equations to match the number of layers.

(2.10), (2.11) removes the gradient term and gives an evolution equation in terms of the vertical component of relative vorticity, $\xi_i := \partial v_i/\partial x - \partial u_i/\partial y$,

$$\frac{\partial \xi_i}{\partial t} + u_i \frac{\partial \xi_i}{\partial x} + v_i \frac{\partial \xi_i}{\partial y} = (\xi + f)\left(\frac{\partial u_i}{\partial x} + \frac{\partial v_i}{\partial y}\right) + \nu \nabla_H \cdot (\nabla \xi_i). \qquad (2.13)$$

Comparing the shallow water vorticity equation (2.13) with the layer mass continuity equation (2.12) shows the existence of a quantity,

$$q_i := \frac{\xi_i + f}{H_i},$$

which is materially conserved in the inviscid limit,

$$\lim_{\nu \to 0} \frac{Dq_i}{Dt} \equiv \frac{\partial q_i}{\partial t} + u\frac{\partial q_i}{\partial x} + v\frac{\partial q_i}{\partial y} = \frac{\nu}{H_i}\nabla_H^2 \xi_i = 0.$$

Although technically accurate only up to a factor of the (constant) layer density, the quantity q_i is usually termed the potential vorticity (PV) of the flow. This conserved quantity replaces the thermodynamic evolution of the primitive equations (2.2–2.3) in the role of a dynamically active tracer of fluid motion.

2.3.2 THE QUASI-GEOSTROPHIC EQUATIONS
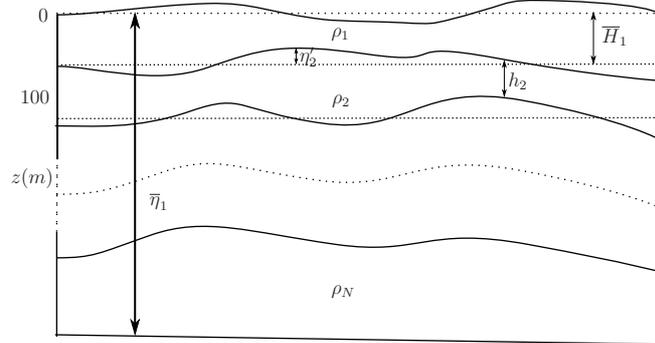
Figure 2.2: Schematic of various height quantities in the QG derivation

The schematic shows the relation of the various height and thickness quantities (H_i, η_i', etc.) in the derivation of the quasi-geostrophic ocean model

Aside from the ocean aspect ratio, δ, the other important non-dimensional pa-

rameter governing the flow is the Rossby number,

$$\epsilon := \frac{U}{fL},$$

where U is a horizontal velocity scale and f is the local value of the Coriolis parameter. This parameter, ϵ, represents a comparison of the magnitude of inertial and Coriolis accelerations in the horizontal momentum equations (2.6). This parameter is small for both atmospheric and oceanic large-scale ($L \sim 10^6$m) flows in middle latitudes, typically:

$$\epsilon_{\text{atmos}} \sim \frac{[10\text{ms}^{-1}]}{[10^{-4}\text{s}^{-1}][10^6\text{m}]} = 10^{-1},$$

$$\epsilon_{\text{ocean}} \sim \frac{[10^{-2}\text{ms}^{-1}]}{[10^{-4}\text{s}^{-1}][10^6\text{m}]} = 10^{-4}.$$

In boundary layer flows such as the Gulf Stream the increased flow rates and reduced spatial scales give a larger value for the relevant dimensionless number, perhaps nearly of order unity, thus it is here that non-linear effects contribute. We therefore expand equations in ϵ and expect from scaling arguments that the leading order behaviour can be described by truncating at first order in ϵ. From $\mathcal{O}(1)$ balance in the momentum equations (2.10-2.11) we obtain geostrophic balance,

$$fv_i = \frac{1}{\rho_0}\frac{\partial p_i}{\partial x} + \mathcal{O}(\epsilon), \ fu_i = -\frac{1}{\rho_0}\frac{\partial p_i}{\partial y} + \mathcal{O}(\epsilon). \tag{2.14}$$

Since the flow is locally two-dimensional it is possible to define a layer streamfunction, $\psi_i(x,y)$, such that

$$u = -\frac{\partial \psi}{\partial y}, \ v = \frac{\partial \psi}{\partial x}. \tag{2.15}$$

The relative vorticity is thus given by

$$\xi_i = \frac{\partial}{\partial x}\left(\frac{\partial \psi}{\partial x}\right) - \frac{\partial}{\partial y}\left(-\frac{\partial \psi}{\partial y}\right) + \mathcal{O}(\epsilon^2) = \nabla^2 \psi + \mathcal{O}(\epsilon^2).$$

From (2.14), (2.9) and (2.15) we find that in the limit of small ϵ

$$\psi_i = \frac{1}{\rho_0 f}\left[p_i + \rho_i gz\right] + C_i,$$

with C_i an arbitrary constant of integration. Dividing the interfaces into their neutral rest and dynamic parts , $\eta_i = \overline{\eta}_i + \eta_i'$ (see Figure 2.2) and defining the streamfunction to vanish at rest gives allows the streamfunction in layer i to be written

as

$$\psi_i = \frac{1}{\rho_0 f}\left(p_{(i-1)}[\overline{\eta}_i + \eta_i'] + \rho_i g \eta_i' + \rho_i g\left(\overline{\eta}_i - z\right)\right) = \psi_{i-1} + \frac{g(\rho_i - \rho_{i-1})}{\rho_0 f}\eta'$$

and hence, on defining the layer reduced gravity, $g_i' = g(\rho_i - \rho_{i-1})/\rho_0$ the deviation in the interface height and the streamfunctions in successive layers are connected through equations,

$$\eta_i' = \frac{f(\psi_i - \psi_{i-1})}{g_i'}.$$

As previously stated the Coriolis parameter is latitude dependent, due to the Earth's spherical geometry. This effect can be simply approximated in Cartesian coordinates by a constant gradient with respect to the Cartesian latitude coordinate y in the f value used,

$$f = f_0 + \beta y.$$

This approximation is good provided f does not vary significantly over the domain, that is provided βL_y is small. Comparing the magnitudes of terms for a typical mid-latitude basin ($\beta = 10^{-11}m^{-1}s^{-1}$) we see that βL is $\mathcal{O}(\epsilon)$ for much of the basin and thus quadratic terms can be assumed to vanish. Provided that the magnitudes of the deviations in interface height, η_i' remain small compared to the resting layer thickness then the shallow water potential vorticity may be expanded in terms of the ratio $\eta_i'/\overline{H}_i$, of interface deviations compared to the average layer depth, $\overline{H}_i$,

$$q_i = \frac{(\nabla^2\psi_i + f_0 + \beta y)}{\overline{H}_i}\left(1 - \frac{f_0(\psi_i - \psi_{i-1})}{g_i'\overline{H}_i} - \frac{f_0(\psi_i - \psi_{i+1})}{g_{i+1}'\overline{H}_i}\right) + \mathcal{O}\left(\frac{\eta_i'^2}{\overline{H}_i^2}, \epsilon^2\right) \quad (2.16)$$

Provided that our assumption that the quadratic terms in ϵ and $\eta'/\overline{H}$ are negligible holds, two terms remain. Firstly, the total column relative vorticity,

$$\nabla^2\psi \sim \frac{\Psi}{L^2},$$

for a length-scale, L and streamfunction anomaly, Ψ, and secondly a term representing the stretching and compression of vortices as the layer interfaces move in geostrophic balance,

$$\frac{f^2\Delta\psi}{g'\overline{H}} \sim \frac{\Psi}{L_D^2},$$

where L_D is the Rossby deformation radius of the system, $L_D := \frac{\sqrt{gH}}{f_0}$ and $g' = g(\rho_2 - \rho_1)/\rho$ the 'reduced gravity'. If the atmospheric pressure is assumed constant, the bottom interface coincides with flat, frictionless, bathymetry then, on ignoring constant terms, the final equation set for a wind forced two layer model (N=2) QG

model is

$$\frac{dq_1}{dt} := \left(\frac{\partial}{\partial t} + \frac{\partial \psi_1}{\partial x} \frac{\partial}{\partial y} - \frac{\partial \psi_1}{\partial y} \frac{\partial}{\partial x} \right) [\xi_1 + \beta y] = \frac{1}{\rho_1 H_1} \left(\frac{\partial \tau^{(y)}}{\partial x} - \frac{\partial \tau^{(x)}}{\partial y} \right) + \nu \nabla^4 \psi_1 \tag{2.17}$$

$$\frac{dq_2}{dt} := \left(\frac{\partial}{\partial t} + \frac{\partial \psi_1}{\partial x} \frac{\partial}{\partial y} - \frac{\partial \psi_1}{\partial y} \frac{\partial}{\partial x} \right) [\xi_2 + \beta y] = \nu \nabla^4 \psi_2 \tag{2.18}$$

The potential vorticity anomalies (ξ_1, ξ_2) are related to the layer stream-functions (ψ_1, ψ_2) through coupled elliptic equations

$$\xi_1 := \nabla^2 \psi_1 - \frac{f_0^2}{g' H_1} (\psi_1 - \psi_2), \tag{2.19}$$

$$\xi_2 := \nabla^2 \psi_2 - \frac{f_0^2}{g' H_2} (\psi_2 - \psi_1). \tag{2.20}$$

The natural condition of no normal flow at material boundaries demands that these form streamlines, i.e that ψ_i is constant on boundaries for all i,

$$\psi_i|_{\delta \mathcal{D}_i} = \Gamma_i(t) \tag{2.21}$$

The lateral boundary condition may be either free slip, expressed as

$$\frac{\partial^2 \psi_i}{\partial n^2} = 0 \ \ \mathbf{n} \perp \delta \mathcal{D}_i, \tag{2.22}$$

for n a unit normal on the boundary, $\delta \mathcal{D}$, or no-slip,

$$\frac{\partial \psi_i}{\partial n} = 0 \ \mathbf{n} \perp \delta \mathcal{D}_i. \tag{2.23}$$

The second condition is the one observed to hold in laboratory experiments, while the first version assumes that boundary processes occur on very fine scales which act to shield the interior flow from frictional effects at the boundary. For low resolution runs modelling flows including western boundary currents then it is possible for the boundary flow to be resolved in as little as two longitudinal grid points. Under these circumstances the free-slip condition (or even one of the many 'hyper-slip' conditions) may be justified. All experiments in this thesis are, however, performed under a non-slip condition.

Regardless of which boundary condition, (2.22) or (2.23), is used, the two equations are both in Von Neumann form and there remains a free system parameter, namely the constant value of the stream function, $\psi_i = \Gamma_i$, equivalent to layer pres-

sure, which is attained on boundaries. This indeterminacy is removed by including a mass conservation equation for each layer,

$$\frac{\partial}{\partial t} \int_{\mathcal{D}_i} , \psi_i dS = 0 \ \ i \in \{1, 2\}, \tag{2.24}$$

which follows directly from the definitions of ψ_i and η_i' chosen in the derivation. The algorithms and techniques employed for the numerical solution of this equation set are detailed and discussed in Appendix A.1.

2.3.3   FORCING FIELDS

The QG layer model formulation assumes constant density within each layer. Moreover distortions of interface height from the neutral position with flat isopycnal surfaces are considered small. This means that the total thickness of any layer cannot vanish and only the top density layer is allowed to outcrop at the surface. It is thus impossible to impose a meridional temperature or salinity gradient across the basin and the model cannot be simply thermodynamically or freshwater forced. This means that the model cannot represent deep density-driven circulations such as the thermohaline circulation in the North Atlantic, or gravity currents such as the Gibraltar overflow but can be forced so as to adequately represent the shallow wind-driven circulations which dominate the balanced flow over much of the North-Atlantic.

The structure of this forcing depends on the applied windstress at the ocean surface. This variable is in reality a complex function of the (time varying) windspeed, the surface currents themselves and the transient condition of the sea surface and is not currently well understood. In this thesis we assume a time-constant wind forcing with variation in the meridional direction only. The forcing has a predominantly symmetric sinusoidal form, with a slight asymmetric component to break the meridional symmetry of the system. This avoids potential problems due to bifurcation of the system behaviour under the symmetry breaking which can be triggered by the accumulation of numerical errors. This in turn reduces the initial spin-up time required by the model from rest towards the point when transfer of vorticity by mesoscale eddies occurs. The precise form of the wind stress, $\boldsymbol{\tau} = (\tau^{(x)}, \tau^{(y)})$ is then

$$\tau^{(x)}(x, y) = \tau_0 \left[\lambda \cos\left(\frac{\pi y}{L_y}\right) - \cos\left(\frac{2\pi y}{L_y}\right) \right], \ \tau^{(y)}(x, y) = 0, \tag{2.25}$$

where the amplitude parameter τ_0 is chosen to be 0.1 Nm^{-2}. This value for τ_0

is close to that typically assumed for both the North Atlantic and North Pacific (Böning et al., 1991), while slightly smaller than the typical value used in most numerical simulations. The value of the asymmetry parameter used here is $\lambda = 0.5$, which produces a moderately small displacement of the zero windstress line along which the jet has a tendency to form. Here as elsewhere in this chapter, y is the (dimensional) distance from the southern boundary and L_y denotes the longitudinal width of the basin. Figure 2.3 shows a plot of this vector field with, for comparison, a plot of the average wind stress data from an ECMWF reanalysis project. The similarity in form between the idealized field and that in effect over the North Atlantic is plain.

In all the experiments that follow $L_y = L_x = 2560$km. This is similar to the configuration used in in Berloff and McWilliams (1999), although in a smaller basin (which offsets the smaller choice of τ_0, maintaining a strongly eddying regime) and with a smaller asymmetric component.
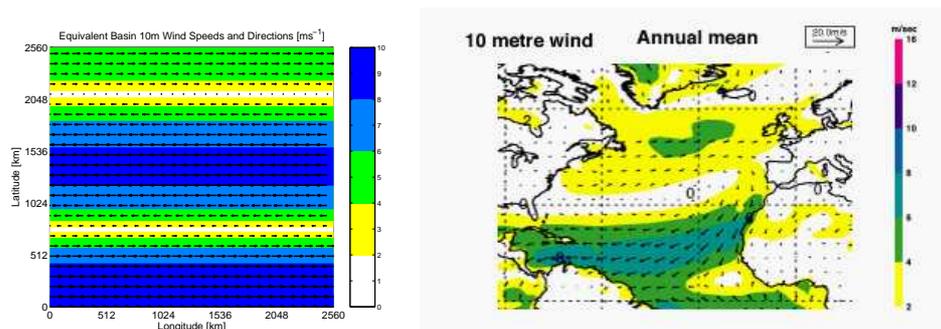
Figure 2.3: Model Wind Stress field and ECMWF ERA40 Climatography

The left hand plot shows the horizontal wind speeds at 10m which, under a typical formula for wind-stress would produce the forcing field, (2.25) used in the QG model. The right hand plot shows an ECMWF ERA40 reanalysis of the 10 m wind climatography on a similar scale. The predominantly zonal pattern shown on the left appears over most ocean basins and the North Atlantic basin has a very similar zonally sinusoidal pattern to it, as well as the slight asymmetry.

ECMWF reanalysis data taken from http://www.ecmwf.int/research/era/

Mean State, $\overline{x}$ Standard Deviation, $S_{N-1}(x)$

Figure 2.4: Long term mean state and standard deviations of q_i, ψ_i

$$\text{mean} : \overline{x} = \sum_{i=1}^{800} \frac{x(t_0 + 12i \text{ hours})}{800} \qquad (S_{N-1}(x))^2 = \sum_{i=1}^{800} \frac{(x(t_0 + 12i \text{ hours}) - x)^2}{799}$$
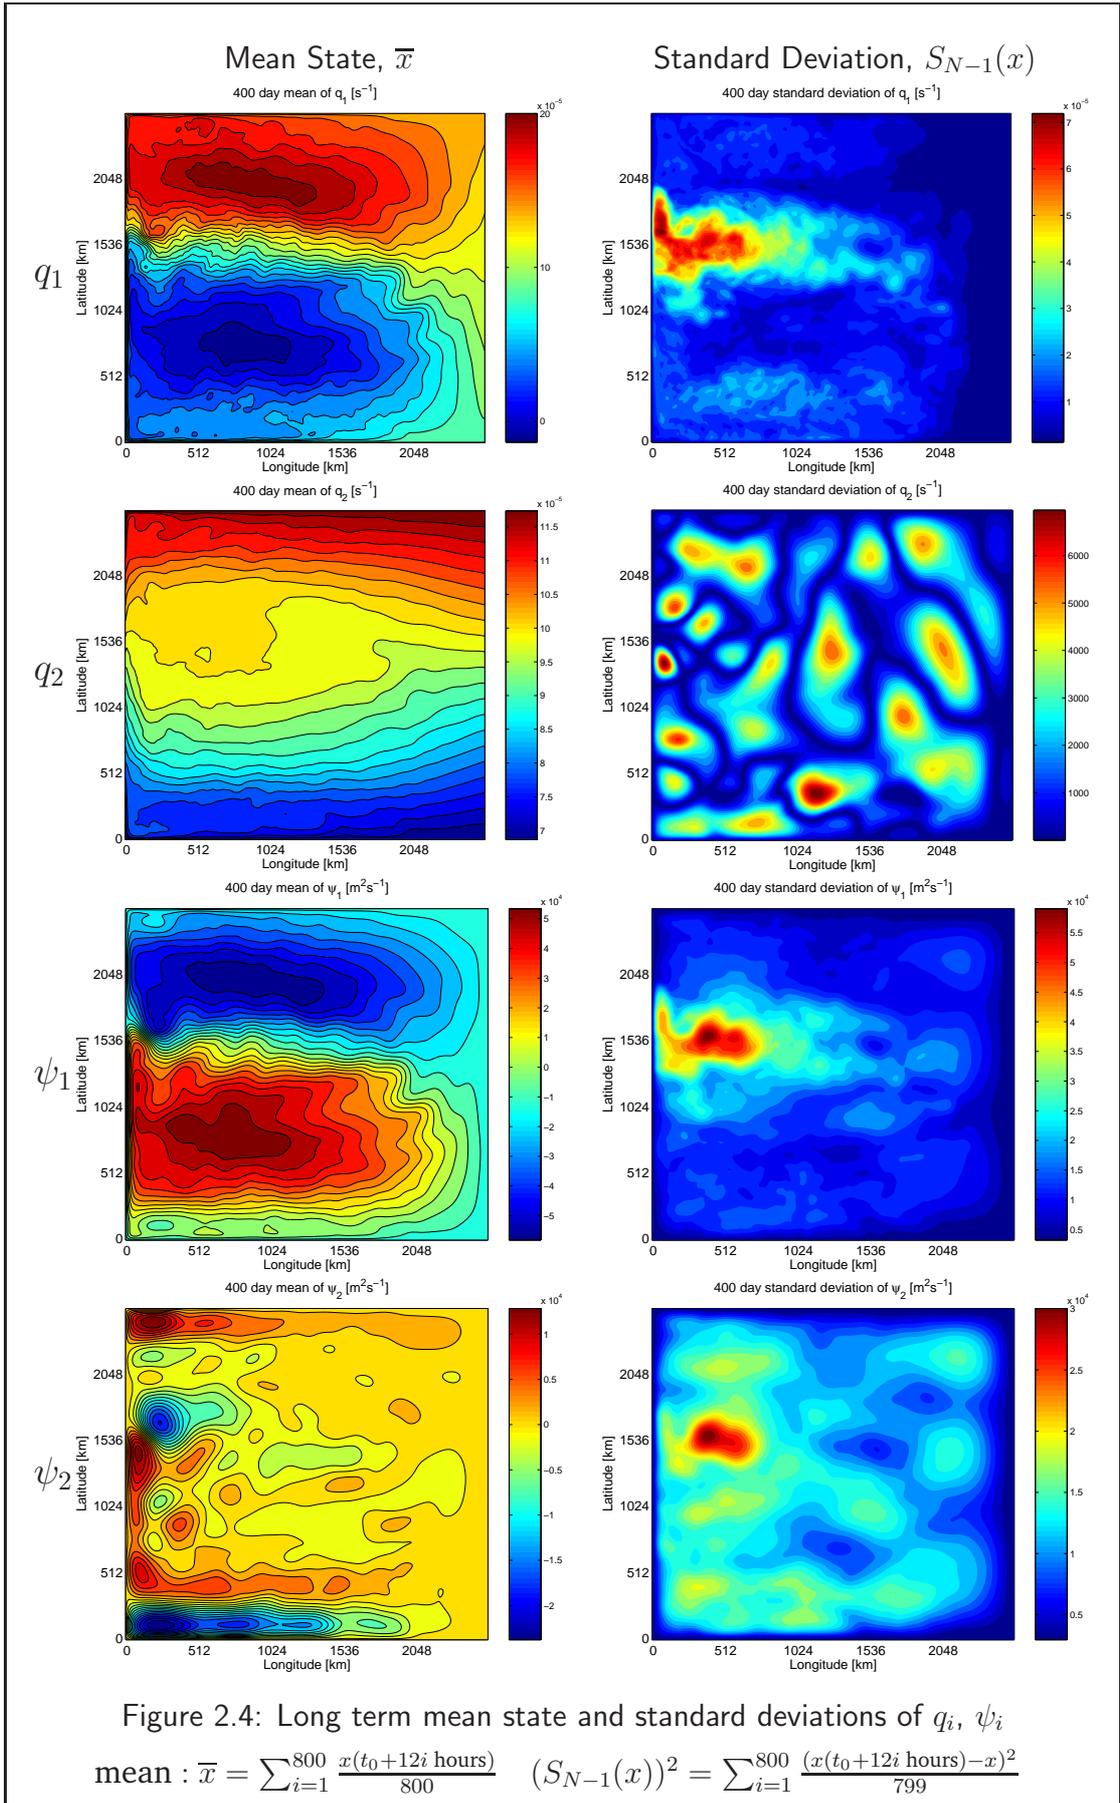
2.4   MODEL BEHAVIOUR

2.4.1   PARAMETER SPACE

The precise model behaviour depends on both the imposed forcing and the various free model parameters. These (dimensional) free parameters can be set by defining the various dimensionless parameters for the system. One possible complete set is determined by giving the various layer aspect ratios, and the Froude, Reynolds, Munk and Rossby numbers of the flow. The aspect ratios are determined entirely by the geometry to be modelled. Here the basin is taken as perfectly square with sides of length $L = 2560$km, the shallow surface layer is assumed to have a depth of 500m and the deep bottom layer a depth of $3,500$ m, for a total basin depth of 4km. The Coriolis and β parameters are chosen to model a typical mid-latitude basin, with $f_0 = 0.7 \times 10^{-4}$ s^{-1}, $\beta = 2 \times 10^{-11}$ m^{-1}s^{-1}. The density difference between the layers is assumed as $\Delta\rho = 1$ kg m^{-3}, equivalent to a vertical temperature gradient of approximately 2.5×10^{-3} K m^{-1}. The only remaining free parameter is the viscosity, η.

For low Reynolds number (highly viscous) flow the equation set has been shown (Bryan, 1963) to rapidly approach a steady state solution with a stable cyclonic subpolar gyre in the north of the basin and counter-rotating anti-cyclonic subtropical gyre in the south (see the mean state images in Figure 2.4) . The precise latitude of the strong eastward jet separating the two gyres is dependent on the magnitude of the asymmetry parameter in the forcing equation. If the Reynolds number of the flow is increased by decreasing viscosity then the steady solution eventually becomes unstable and patterns of mesoscale eddies form along the jet acting to transfer vorticity between the gyres. The location of the main zone where these eddies form clearly shown in the standard deviation field for the top or bottom layer streamfunction in Figure 2.4. Note, however, that the standard deviation in the top layer potential vorticity is dominated by excursions in the separation point of the jet and the bottom layer by variations caused by the slow meanders in the return circulations.

2.4.2   ENERGETICS

We briefly discuss the energetics and conserved and bounded quantities present in the model, to link back to the initial motivations of data assimilation in the presence of conserved quantities and their effects on large scale circulation. The energetics of the model can be bounded by considering the volume integrated be-

Parameter	Value	Notes
$\mathrm{Fr}_i := \dfrac{U}{\sqrt{g'H_i}}$	6×10^{-2} Interior 0.8 Jet	Flow near critical at jet separation point, flow subcritical elsewhere.
$\mathrm{Re} := \dfrac{UL}{\nu}$	100 Interior 2000 Jet	Flow is turbulent everywhere
$\mathrm{Mo} := \dfrac{U^2}{\nu \beta L}$	1×10^8	Boundary Layer is poorly resolved
$\mathrm{Ro} := \dfrac{U}{f_0 L}$	1×10^{-3} Interior 2 Jet	Rossby number is order unity or less everywhere

Table 2.1: Dimensionless parameters used for the QG basin model

haviour of some model quantities, in the manner of Wu (2001). Consider first the volume integral of the first moment of total vorticity,

$$
\begin{aligned}
Q &= \int_{\text{layer 1}} q_1 dV + \int_{\text{layer 2}} q_2 dV, \\
&= H_1 \int_{\mathcal{D}_{\mathcal{H}}} \nabla^2 \psi_1 - \frac{1}{L_1}(\psi_1 - \psi_2) + \beta y dS + H_2 \int_{\mathcal{D}_{\mathcal{H}}} \nabla^2 \psi_2 - \frac{1}{L_2}(\psi_2 - \psi_1) + \beta y dS, \\
&= \beta(H_1 + H_2)\frac{L^2}{2},
\end{aligned}
$$

where the final result follows from mass conservation and the non-slip boundary condition. This states simply that under the non-slip boundary condition the total circulation is conserved. This in turn means that the total PV in the system is conserved, with the forced negative anomaly in the north of the basin balanced by a similar positive anomaly in the south.

We now define inner products for scalars, α, γ and 2-dimensional or layer, vectors, $\boldsymbol{F}, \boldsymbol{G}$, defined over the area of the base of the domain of interest, $\mathcal{D}_H$,

$$
\langle \alpha, \gamma \rangle = \int_{\mathcal{D}_H} \alpha \gamma \, dS,
$$

$$
\langle \boldsymbol{F}, \boldsymbol{G} \rangle = \int_{\mathcal{D}_H} \boldsymbol{F} \cdot \boldsymbol{G} \, dS,
$$

and the resultant metrics on scalars and vectors,

$$
\|\alpha\|^2 = \langle \alpha, \alpha \rangle,
$$

$$\|\boldsymbol{F}\|^2 = \langle \boldsymbol{F}, \boldsymbol{F} \rangle.$$

These are simply the L_2 metrics over $\mathcal{D}_H$ for scalar and vector fields respectively.

This notation allows us to define several physical properties of the model. Firstly we have the total kinetic energy,

$$
\begin{aligned}
\mathcal{K} &:= \rho_0 H_1 \frac{\|\boldsymbol{u}_1\|^2}{2} + \rho_0 H_2 \frac{\|\boldsymbol{u}_2\|^2}{2}, \\
&= \rho_0 H_1 \frac{\|\nabla \psi_1\|^2}{2} + \rho_0 H_2 \frac{\|\nabla \psi_2\|^2}{2}.
\end{aligned}
$$

The next quantity is the available potential energy, $\mathcal{A}$, stored in the displacement of the density interface from its neutral position,

$$\mathcal{A} = g(\rho_2 - \rho_2)\|\eta_2'\|^2 = \frac{\rho_0 f_0^2 \|\psi_1 - \psi_2\|^2}{g'}.$$

Finally there is the total model (relative) enstrophy $\mathcal{Q}$ and $\mathcal{B}$, a measure of the kinetic energy trapped in the baroclinic mode, defined for layer vorticity fields, ζ_i, by

$$
\begin{aligned}
\mathcal{Q} &= \rho_0 H_1 \|\zeta_1\|^2 + \rho_0 H_1 \|\zeta_2\|^2 \\
&= \rho_0 H_1 \|\nabla^2 \psi_1\|^2 + \rho_0 H_2 \|\nabla^2 \psi_2\|^2, \\
\mathcal{B} &= \frac{\rho_0 f_0^2 \|\nabla(\psi_1 - \psi_2)\|^2}{g'}.
\end{aligned}
$$

Using the adjoint nature of the gradient and divergence operators we find evolution equations for $\mathcal{K} + \mathcal{A}$ and for $\mathcal{Q} + \mathcal{B}$,

$$\frac{\partial(\mathcal{K} + \mathcal{A})}{\partial t} = H_1 \left\langle \psi_1, \frac{\partial \tau^{(x)}}{\partial y} \right\rangle - 2\nu \mathcal{Q}, \tag{2.26}$$

$$\frac{\partial(\mathcal{Q} + \mathcal{B})}{\partial t} = \left\langle \nabla^2 \psi_1, \frac{\partial \tau^{(x)}}{\partial y} \right\rangle - 2\nu \mathcal{S}, \tag{2.27}$$

where we define a new functional containing the squares of the third order derivatives in the system,

$$\mathcal{S} = \frac{\rho_0}{2} \left(H_1 \|\nabla \left(\nabla^2 \psi_1 \right)\|^2 + H_2 \|\nabla \left(\nabla^2 \psi_2 \right)\|^2 \right).$$

The first equation states that the total energy in the system is forced by resonance between the surface wind-forcing and pressure fields and reduced through viscous dissipation of the relative enstrophy field. The second equation states that relative enstrophy is forced by baroclinic effects and the resonance between the surface vorticity field and wind forcing and reduced by dissipative effects at higher order.

On application of the Cauchy-Schwartz inequality and then Young's inequality

upper bounds may be placed on the forcing part of the growth rates of (2.26-2.27),

$$H_1 \left\langle \psi_1, \frac{\partial \tau^{(x)}}{\partial y} \right\rangle \leq H_1 \|\nabla^2 \psi_1\| \left\| \nabla^{-2} \frac{\partial \tau^{(x)}}{\partial y} \right\| \leq \frac{\left\| \nabla^{-2} \frac{\partial \tau^{(x)}}{\partial y} \right\|}{2\nu \rho_0} + \nu \mathcal{Q},$$

$$H_1 \left\langle \nabla^2 \psi_1, \frac{\partial \tau^{(x)}}{\partial y} \right\rangle \leq \|\nabla(\nabla^2 \psi_1)\| \left\| \Delta^{-1/2} \frac{\partial \tau^{(x)}}{\partial y} \right\| \leq \frac{\|\Delta^{-1/2} \frac{\partial \tau^{(x)}}{\partial y}\|}{2\nu \rho_0} + \nu \mathcal{R},$$

where $\nabla^{-2} f$ and $\Delta^{-1/2} f$ denote the solutions to

$$\begin{cases} \nabla^2 g = f, & \boldsymbol{r} \in \Omega, \\ g = 0, & \boldsymbol{r} \in \delta\Omega \end{cases}$$

and

$$\begin{cases} \nabla.\boldsymbol{G} = f, & \boldsymbol{r} \in \Omega \\ \boldsymbol{G} = \boldsymbol{0}, & \boldsymbol{r} \in \delta\Omega \end{cases}$$

respectively. Combining these results with equations (2.26-2.27) and assuming the existence of an asymptotic mean allows the derivation of a bound on the asymptotic behaviour of $\mathcal{Q}$ and $\mathcal{S}$,

$$\mathcal{Q} \leq \frac{1}{\nu^2} \left\| \nabla^{-2} \frac{\partial \tau^{(x)}}{\partial y} \right\|,$$

$$\mathcal{S} \leq \frac{1}{\nu^2} \left\| \Delta^{-1/2} \frac{\partial \tau^{(x)}}{\partial y} \right\|.$$

This gives three quantities, one of which is globally conserved by the equation set and two more which have their asymptotic behaviour bounded by the forcing conditions. Similarly there exists a quantity (PV itself) which is a dynamic tracer obeying an advection-diffusion equation. We can identify these with the total ocean heat content and salt content and the local entropy variable. As such the behaviour of PV under assimilation can be expected to be a suitable sand-box test for effects on conservative properties of the various assimilation algorithms and a proxy for the behaviour of variables such as potential temperature and salinity under the application of such methods to more realistic ocean models and to layer GCMs, even though the QG model does not contain any terms representing thermodynamic forcings.

2.5   GENERATING DIVERGENT MODEL TRAJECTORIES

Coupling the conserved properties discussed in the previous section to the ideals of the identical twin experiments philosophy of section 2.2 we conclude this chapter with a discussion of possible methods to generate the separate 'truth' and 'observation' data, noting their effects on conservation and justifying our final selections. Since the applied wind field is time constant and non-random the QG model which has been presented is entirely deterministic. For a given set of initial conditions,

$$\psi(t_o, \boldsymbol{r}) = \psi_0(\boldsymbol{r}),$$

and set of model parameters, $(\boldsymbol{\tau}, \nu)$, the model trajectory,

$$m_{\boldsymbol{\tau}, \nu}(t, \psi_0),$$

is determined for all time intervals, regardless of the initial time and dependent only on the time interval from the start of the run,

$$\Delta t = t - t_0.$$

Generating a differing trajectory to run experiments under the identical twin philosophy described in Section 2.2 thus requires perturbing either the initial conditions or the model itself through the model parameters or equations. Random perturbations to the initial conditions on the model grid-scale added to the vorticity or streamfunction fields introduce unphysical transient motions and may upset prescribed global conditions such as mass balance. Since observations of such transients may adversely affect any assimilation scheme it is necessary to allow a period following the insertion of perturbations for the model to adjust and for dissipation to dampen transients.

One technique to generate physically consistent differences, the singular vector (SV) method (Buizza and Palmer, 1995) is to consider the linearization of the discretized model,

$$M_{ij} = \frac{\partial m(\psi_j)}{\partial \psi_i},$$

where ψ_i represents the i^{th} element of the vector of grid-point streamfunction values. This matrix represents the sensitivity to perturbations in ψ of the discretized model and has a singular value decomposition,

$$\boldsymbol{M} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{V}^T,$$

such that the matrix D is positive definite and diagonal and the columns of U and rows of V^T each form an orthonormal basis,

$$U^T U = I_{\dim(U)}, \; V^T V = I_{\dim(V)}.$$

If D has the natural ordering such that $D_{11} > D_{22} > \ldots D_{nn}$ then the leading rows of V^T represent the spatial structures which grow fastest at the time of linearization and thus the physical modes of the model which will dominate at the end of the run under any random perturbation of the conditions at the initial time. This method is used to populate many Ensemble assimilation and prediction schemes, including that of the ECMWF (Molteni et al., 1996). In the context of adiabatic assimilation however it is not clear that the resultant fields will satisfy the rearrangement condition at any order.

An alternative method, suitable for the generation of phase error is to consider the realization of model trajectories for a single set of initial conditions and parameters values, but valid at two different model times, t_{true} and $t_{\text{background}}$. When these times are close, the background state is also a fair representation of the model state at time t_{true}, but with differences between the states generated by the advective and diffusive processes in the model itself. In the low Peclet number regime this will be dominated by advection, ensuring that differences will be mostly in the position of features, rather than their amplitude.

The experiments presented in this thesis concentrate primarily on differences under variations in the wind forcing and using the advective method. Variations in the diffusivity constant, ν, are avoided due to the strong variability in the transient model dynamics this can generate as well as the possibility of forcing the model towards the steady-state regime. The SV method is not used due to the relatively high cost of generating the perturbations and the question regarding the effect on conservation properties.

2.6   CHAPTER SUMMARY

In this chapter we have built a basis for the numerical and philosophical methods to be presented in the succeeding chapters of this thesis. The chapter began by deriving a conservation principle limiting the motion of material fluid properties across ocean isopycnal surfaces which is found to hold well for large-scale ocean circulations, both in theory and in practice. It is this adiabatic condition which justifies practically our interest in conservative data assimilation, particularly the concepts of rearrangement assimilation to be discussed in the next chapter

We then introduced and defended the identical twin experiment methodology which is to be used later to investigate and quantify the behaviour of the data assimilation schemes under study and the two-layer quasi-geostrophic ocean model which is to form the work-horse for the numerical twin experiments comparing our new and unusual data assimilation techniques with more traditional methods, both conservative and non-conservative, such as are currently used at operational ocean forecasting centres. The behaviour of the model was investigated and shown to present many of the features observed in the real ocean, both qualitatively, in terms of fine scale structure and theoretically in terms of presenting a number of conservation properties, important in light of the previous discussion of the adiabatic condition and our general interest in conservative data assimilation.

The chapter concludes with a discussion of possible choices for a methodology for generating the necessary divergent model trajectories for use in identical twin experiments. We have show that the usual method of insertion of random perturbations is not suitable as a sole choice when a conservative principle is held to exist and have derived several alternative methods. We now move on to the topic of constrained data assimilation in the presence of constraints, particularly integral constraints of the manner of the adiabatic condition.

CHAPTER 3

DISPLACEMENT ASSIMILATION

"... because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns - the ones we don't know we don't know."

Donald Rumsfeld
Department of Defense news briefing
12[th] February 2002

It is necessary to place the original work contained herein within the context of the wider study of adiabatic assimilation. This chapter details the assumptions required for such assimilation methods to be physically consistent and the basic conditions for such methods to be more successful than approaches based on direct modification of Eulerian variables, such as the BLUE equation introduced in Chapter 1. The underlying mathematical theory, known as rearrangement theory, is then introduced together with the related Monge-Kantorovich transport problem. This section is presented in a notation consistent with that used here for the more general data assimilation. Also discussed are the situations in the area of ocean data assimilation in which the formal rearrangement approach is not valid, namely the cases of severely limited or noisy observations.

Following this a new metric on phase/displacement error is presented and discussed with comparison to the standard quadratic cost function, representing an Eulerian measure of distance between points in the function space. The behaviour of the metric is investigated for the quasi-geostrophic ocean model introduced in the previous chapter. We also investigate the quantitative and qualitative behaviour of the model under the assimilation using the Cooper & Haines method under the identical twin philosophy.

3.1   ASSIMILATION AND CONSERVATION PROPERTIES

3.1.1   CONSTRAINTS AND THE BLUE ANALYSIS

The BLUE analysis of section 1.1 is the statistically optimal solution to the naïve data assimilation problem, in the sense of generating an analysis which is the most likely estimator of the true state of the system, only if knowledge of the error statistics of the noise in the system is available, or at least can be accurately estimated. Formally this requires the elements of the background and observation error covariance matrices to determined. In practice these data are still poorly understood and very simple approximations are made in order to speed calculation. Typically, the observations are assumed independent, giving a diagonal observation covariance matrix. The background error covariance matrix is frequently used to spread the information contained in observations into neighbouring grid-points, with the elements of the matrix chosen to reduce in magnitude in proportion to the physical separation of model grid points.

Consider the case of a simple one dimensional model in a periodic domain of length L, discretized into the values at 99 equally spaced grid-points, $\boldsymbol{x} \in \mathbb{R}^{99}$. If the true state of the system is a simple sine wave function of period $L/2$, so that the individual components of the true state vector $\boldsymbol{x}_t$ are given by

$$[\boldsymbol{x}_t]_i = \sin\left(\frac{4\pi i}{100}\right),$$

then two obvious sources of background error may be considered:

i. Amplitude error, where the background is in phase with the truth, but incorrectly scaled by a constant factor, $A \neq 1$, giving components

$$[\boldsymbol{x}_b^{\mathrm{A}}]_i = A \sin\left(\frac{4\pi i}{100}\right) = A[\boldsymbol{x}_t]_i,$$

ii. Phase error, where the magnitude of the signal is correct, but leads or lags the truth through a phase difference, s,

$$[\boldsymbol{x}_b^{\mathrm{phase}}]_i = \sin\left(\left(4\pi\frac{i}{100} + s\right)\right) \approx [\boldsymbol{x}_t]_{i+25s/\pi}.$$

Suppose there are a number, $m < 99$, of direct observations of the true state $\boldsymbol{x}_t$ available on an ordered grid, with larger grid spacing than the model, and that we estimate the observation and model error covariance matrices in a manner consis-

tent with the typical approach, by

$$\boldsymbol{R} = \sigma_o^2 \boldsymbol{I}_m,,$$

$$\boldsymbol{B}_{ij} = \sigma_b^2 \exp\left(-\frac{d(i,j)^2}{\lambda^2}\right).$$

Here σ_o^2, assumed in the BLUE as the variance of the error assumed for the observations is instead simply a weighting function, $d(i,j)^2 = (i-j)^2 \Delta x^2$ is the Eulerian distance between data points and λ is a length-scale weighting the distance over which information may be expected to be spread.

We may now apply the BLUE equation (1.3) to examples of each type of error and generate analyses. The result of two such applications is shown in Figure 3.1. When presented with amplitude error the method proceeds as one might expect, generating an analysis which lies somewhere between the observed values (and therefore in this case the truth) and the background first guess. The initial amplitude error is also reduced, without introducing significant phase error.

The qualitative behaviour of the analysis in the initial presence of phase error is different. The analysis still lies between the observations and background, which is in fact enforced by the BLUE assimilation algorithm and the initial phase error has been reduced, but a significant amplitude error has also been introduced. To capture information on phase error requires many accurate observations to represent the signal without aliasing problems. A standard rule of thumb from numerical modelling suggests 10 observations per wavelength are required to represent a sine wave in a dispersive model (Stephen and Bolmer, 2006) and a similar result is likely to occur here. This problem can be even more serious in the more realistic case when observations are only available on an irregular grid, since this can give a much larger typical spacing for observation grid over much of the domain.

In practice, extra information is available in the form of physical constraints on the relationship between variables, for example the relationship through geostrophic balance between density and velocity fields in the ocean. It is possible to represent some relationships by modifying the choice of the covariance matrices used or by transforming variables before the assimilation step. To apply other constraints it becomes necessary to consider other assimilation techniques.

3.1.2 ASSIMILATION THROUGH REARRANGEMENTS

This section presents the method and notation of rearrangement theory as it applies to the constrained assimilation problem of inserting observational data while
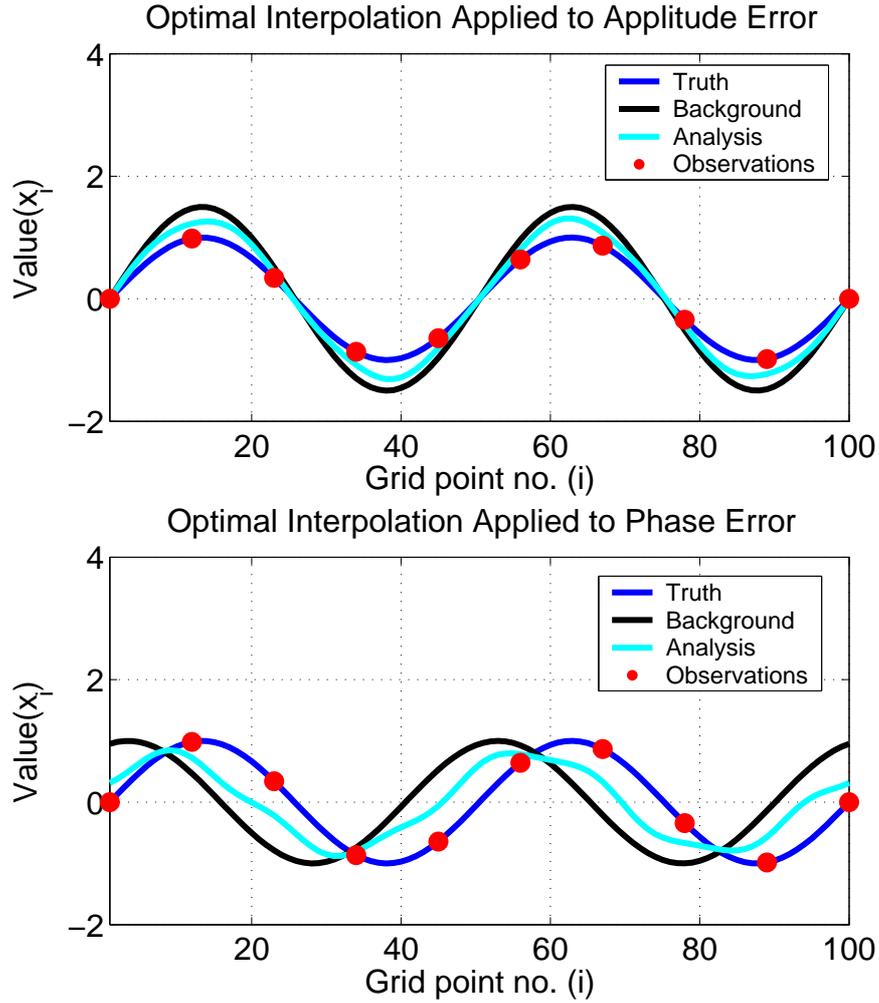
Figure 3.1: A traditional DA method applied to amplitude and phase errors

We apply the optimal interpolation method (Daley, 1991) to generate the BLUE analysis for cases of amplitude error (top plot),

$$[\boldsymbol{x}_b^{\mathrm{A}}]_i = A \sin\left(\frac{4\pi i}{100}\right) = A[\boldsymbol{x}_t]_i,$$

and of phase error (bottom plot),

$$[\boldsymbol{x}_b^{\mathrm{phase}}]_i = \sin\left(\left(\frac{4\pi i}{100} + s\right)\right) \approx [\boldsymbol{x}_t]_{i+100s}.$$

For these plots $A = 1.5$, $s = 0.1$, $\sigma_b = \sigma_o = 1$ and $\lambda = L/25$. Note particularly that there is minimal phase error introduced in the presence of amplitude error, with a maximum lead of about 3 grid-points around grid points 80-90, while in the second case there is an amplitude error of roughly 10-20% introduced across the entire signal, as well as he obvious deformation to the form of the signal.

conserving the original material properties present in the background field. See Douglas (2002) for a much fuller discussion of the mathematics behind rearrangement theory and for a fuller list of references. Cullen (2002) discusses the use of rearrangement theory in meteorology in general, as well as in assimilation specifically.

Consider the problem of tracking the evolution of a passive tracer field in the ocean, for example the concentration of a chemical following a diffuse spillage. Let $C(\boldsymbol{r})$ denote the concentration of the chemical at a point $\boldsymbol{r}$ inside the domain of interest $\mathcal{D}$. Suppose there is no flux of the tracer across the boundary of the domain, $\delta\mathcal{D}$, that there are no new sources or sinks of the chemical and that mixing processes are no longer significant on the time and length scales being modelled[1]. Given a background or hindcast estimate of the initial concentration field, C_0, valid at a time t_0 for all of $\mathcal{D}$ and a forecast for the velocity field, $\boldsymbol{u}_b(t, \boldsymbol{r})$ which applies in $\mathcal{D}$ over a time window $[t_o, t_f]$, then a forecast $C_b(t, \boldsymbol{r})$ for the concentration field in $[t_o, t_f] \times \mathcal{D}$ can obviously be calculated directly by integrating the (forecast) chemical conservation equation for C_b,

$$\frac{\partial C_b}{\partial t} + \boldsymbol{u_b} \cdot \nabla C_b = 0, \tag{3.1}$$

forwards in time starting from the known estimate of initial conditions,

$$C_b(t_0, \boldsymbol{r}) = C_0(\boldsymbol{r}).$$

If the equation (3.1) accurately describes the physics of the true system, that is if the assumptions of no mixing and no forcings are accurate, and the true flow velocity field is incompressible then the real chemical concentration field satisfies the conditions of the following proposition :

Proposition 3.1 (The rearrangement condition) *Let τ be an unforced tracer field on a domain $\mathcal{D}$, which satisfies a simple advection equation,*

$$\frac{\partial \tau}{\partial t} + \boldsymbol{u} \cdot \nabla \tau = 0, \tag{3.2}$$

where $\boldsymbol{u}$ is an incompressible velocity field,

$$\nabla \cdot \boldsymbol{u}(\boldsymbol{r}) = 0, \qquad\qquad \boldsymbol{r} \in \mathcal{D},$$

---

[1]Of course if the original release was from a point source then mixing must dominate over the short time and length scales of the initial release. The case considered here is the one after these initial effects no longer control the flow.

and there is a no normal flow condition on the physical boundaries of $\mathcal{D}$,

$$\boldsymbol{u} \cdot \boldsymbol{n} = 0, \qquad\qquad \boldsymbol{n} \perp \delta\mathcal{D}.$$

Further, let the state of the tracer be defined at some time by the state τ_0,

$$\tau(t_0, \boldsymbol{r}) = \tau_0(\boldsymbol{r}).$$

Then the condition

$$\int_{\mathcal{D}} (\tau - \alpha)\mathcal{H}[\tau - \alpha]dV = \int_{\mathcal{D}} (\tau_0 - \alpha)\mathcal{H}[\tau_0 - \alpha]dV, \qquad (3.3)$$

holds for any choice of t and for any real number, α. Here $\mathcal{H}[x]$ simply denotes the Heaviside function applied pointwise to the field x,

$$\mathcal{H}[x(\boldsymbol{r})] = \begin{cases} 1 & x(\boldsymbol{r}) \geq 0 \\ 0 & x(\boldsymbol{r}) < 0. \end{cases}$$

The proof of this assertion is simple and follows directly from equation (3.2) plus the incompressibility condition on $\boldsymbol{u}$.

Proof :

$$
\begin{aligned}
\frac{d}{dt}\int_{\mathcal{D}} (\tau - \alpha)\mathcal{H}[(\tau - \alpha)]dV =\ & \frac{d}{dt}\int_{\mathcal{D}(\alpha)} (\tau - \alpha)dV, \\
=\ & \int_{\mathcal{D}(\alpha)} \frac{\partial(\tau - \alpha)}{\partial t}dV \quad + \int_{\delta\mathcal{D}(\alpha)} (\tau - \alpha)\boldsymbol{u} \cdot d\boldsymbol{S}, \\
=\ & -\int_{\mathcal{D}(\alpha)} \boldsymbol{u}.\nabla(\tau - \alpha)dV \quad + \int_{\delta\mathcal{D}(\alpha)} (\tau - \alpha)\boldsymbol{u} \cdot d\boldsymbol{S}, \\
=\ & -\int_{\delta\mathcal{D}(\alpha)} (\tau - \alpha)\boldsymbol{u} \cdot d\boldsymbol{S} \quad + \int_{\delta\mathcal{D}(\alpha)} (\tau - \alpha)\boldsymbol{u} \cdot d\boldsymbol{S} = 0.
\end{aligned}
$$

Here $\mathcal{D}(\alpha)$ denotes that subset of $\mathcal{D}$ over which $\mathcal{H}[\tau - \alpha]$ is non-zero, hence the space over which the integral must be evaluated. The domain of the surface integrals, $\delta\mathcal{D}(\alpha)$, is the boundary of this subset. The first term in the expansion comes directly from the time dependence of τ as it appears in the integrand and the second term arises from the application of Leibnitz rule required, since the limit of the integral, $\mathcal{D}(\alpha)$, is itself an implicit function of time.

The rearrangement condition, one form[2] of which is given in equation (3.3), is

---

[2]Other possible forms for the condition include the condition that the integrals of any moment of τ and τ_0 must balance,

$$\int_{\mathcal{D}} \tau^n dV = \int_{\mathcal{D}} \tau_0^n dV \ \forall n \in \mathbb{N}^+,$$

a relation operator on the set of all possible tracer fields on $\mathcal{D}$. The fact that it is satisfied by all fields generated from equation (3.2) can be viewed as stating that the total mass of any material tracer contained within any level set of a tracer field is conserved during advection by an incompressible velocity field. This property follows directly from volume conservation in a Lagrangian interpretation of fluid dynamics. More generally any field on a measurable space which satisfies an equivalent condition in the form of equation (3.3) with respect to a test field is termed (Douglas, 2002) a rearrangement of the test field, that is the equivalence classes of the test field under the relation is the set of its rearrangements. See Figure 3.2 for an example of two functions on the real line which satisfy the condition, as well as a graphical representation of the meaning of the integrals.

Returning to our hypothetical chemical tracer field, if observational data is available to correct the forecast, it is natural (and desirable) to seek to preserve the total mass of the chemical in the system, assuming that the mass contained in the initial estimate was accurate[3]. If it is known that the true velocity field advecting the tracer is incompressible then it follows from assertion 1 and our assumption that the initial concentration pattern is accurate that any analysis field, C_a, over the course of the time window must be a rearrangement of the initial field. Moreover, since the rearrangement property is transitive and the forecast also satisfies the conditions of Proposition 3.1 it is clear that C_a must also be a rearrangement of C_b. Given full, but possibly noisy, observations, C_y, of the true field, C_t, then rearrangement theory provides a method to find an analysis field, C_a, which solves a problem closely related to the standard problem in data assimilation, namely:

Problem 3.1 *Given a background estimate, $C_b(t_f, \boldsymbol{r})$, of a scalar function on $\mathcal{D}$ and full, noisy observations, $C_y(t_f, \boldsymbol{r})$, of the true state of the field, find an analysis state, $C_a(t_f, \boldsymbol{r})$, which minimizes a quadratic cost function,*

$$\mathcal{J}(C) := \int_{\mathcal{D}} (C - C_y)^2 \, dV,$$

subject to the condition that C_a is a rearrangement of C_b.

The solution depends on being able to find what Brenier (1991) has termed a 'polar factorization' of both C_b and C_y. This method of factorization represents an

---

as well as a number of versions given in the formalism of measure theory, see Douglas (2002) and references therein.

[3]This condition that the mass densities are accurately known, even if their locations are poorly understood is a strong, but key assumption in the entire theory that follows. We argue that in general surface fluxes are more easy to ascertain than the subsequent distribution of tracers with depth, justifying our interest in assimilation methods which work when the assumption holds.
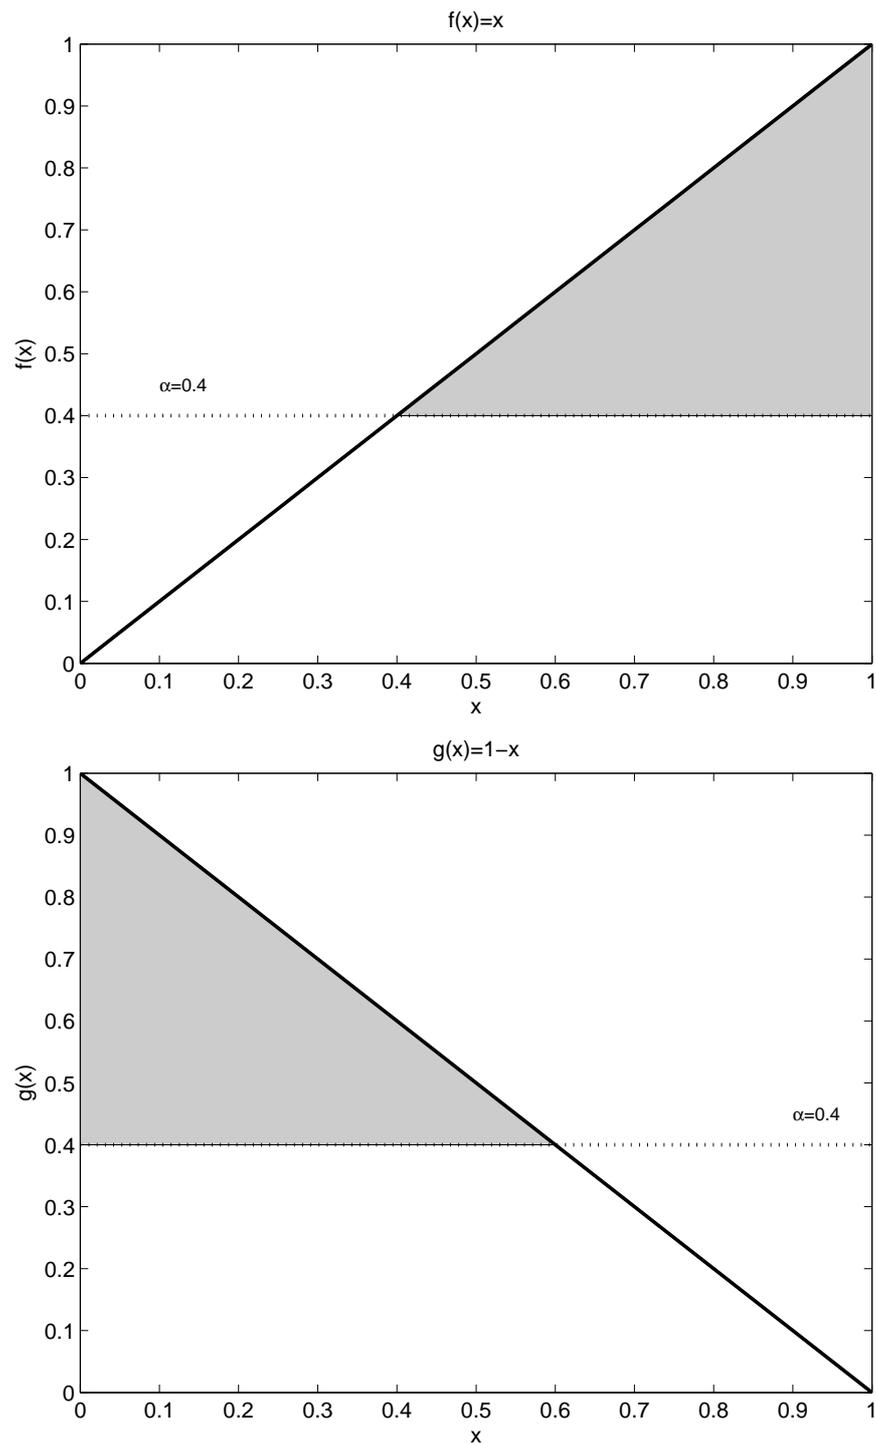
Figure 3.2: Figure denoting a rearrangement of a 1D function

The shaded areas in the two plots are equal. These represent $\int_0^1 (f - \alpha)\mathcal{H}[f - \alpha]dx$ and $\int_0^1 (g - \alpha)\mathcal{H}[g - \alpha]dx$ respectively, for $\alpha = 0.4$. This property remains true regardless of the actual choice of α made. Hence, by our definition in equation (3.3), $g(x)$ is a rearrangement of the function $f(x)$. Figure after Douglas (2002).

arbitrary scalar or vector-valued function by the convolution of a convex function acting on a sphere, $\mathcal{S}$, of the same volume as the domain, $\mathcal{D}$ and a measure preserving mapping from $\mathcal{D}$ to $\mathcal{S}$. For a scalar function, C, the new convex function, f satisfies a modified, spherical rearrangement condition,

$$\int_{\mathcal{S}} (f - \alpha)\mathcal{H}[f - \alpha]dV = \int_{\mathcal{D}} (C - \alpha)\mathcal{H}[C - \alpha]dV \ \forall \alpha \in \mathbb{R},$$

a symmetry condition,

$$f(\boldsymbol{r_1}) = f(\boldsymbol{r_2}), \ \ \forall \|\boldsymbol{r_1}\| = \|\boldsymbol{r_2}\| \, , \ \boldsymbol{r}_1, \boldsymbol{r}_2 \in \mathcal{S},$$

and a radius-wise monotonicity condition,

$$f(\boldsymbol{r_1}) \leq f(\boldsymbol{r_2}), \ \ \forall \|\boldsymbol{r_1}\| < \|\boldsymbol{r_2}\| \, , \ \boldsymbol{r}_1, \boldsymbol{r}_2 \in \mathcal{S}.$$

The convex function produced is unique almost everywhere, up to a change in value at a set of points of zero measure. See Figure 3.3 for an example of the two-dimensional spherically symmetric rearrangement of an SST field. The mapping, s, from $\mathcal{D}$ onto $\mathcal{S}$ is non-unique in that any mapping such that $f(s(\boldsymbol{r})) = C(\boldsymbol{r})$ will suffice. The mapping thus contains all the positional information about the field, while the convex function provides amplitude information about the values the field attains.

The formal solution to Problem 3.1 is that given the polar factorizations for C_y and C_b,

$$C_y(\boldsymbol{r}) = f_y(s_y(\boldsymbol{r})),$$

$$C_b(\boldsymbol{r}) = f_b(s_b(\boldsymbol{r})),$$

we may form a new analysis function,

$$C_a(\boldsymbol{r}) = f_b(s_y(\boldsymbol{r})).$$

This inherits the rearrangement properties of f_b and hence is a rearrangement of C_b, but it is also shown in e.g. Douglas (2002) that this achieves the minimum of $\mathcal{J}(f, g)$ over all rearrangements of f_b. We may thus identify C_a with the desired analysis field.

Note that although advection by an incompressible velocity will always generate a field which is a rearrangement of the initial conditions, not all fields which are rearrangements of each other can be achieved through a period of advection. A simple counter-example considers the functions of Figure 3.2, which cannot be
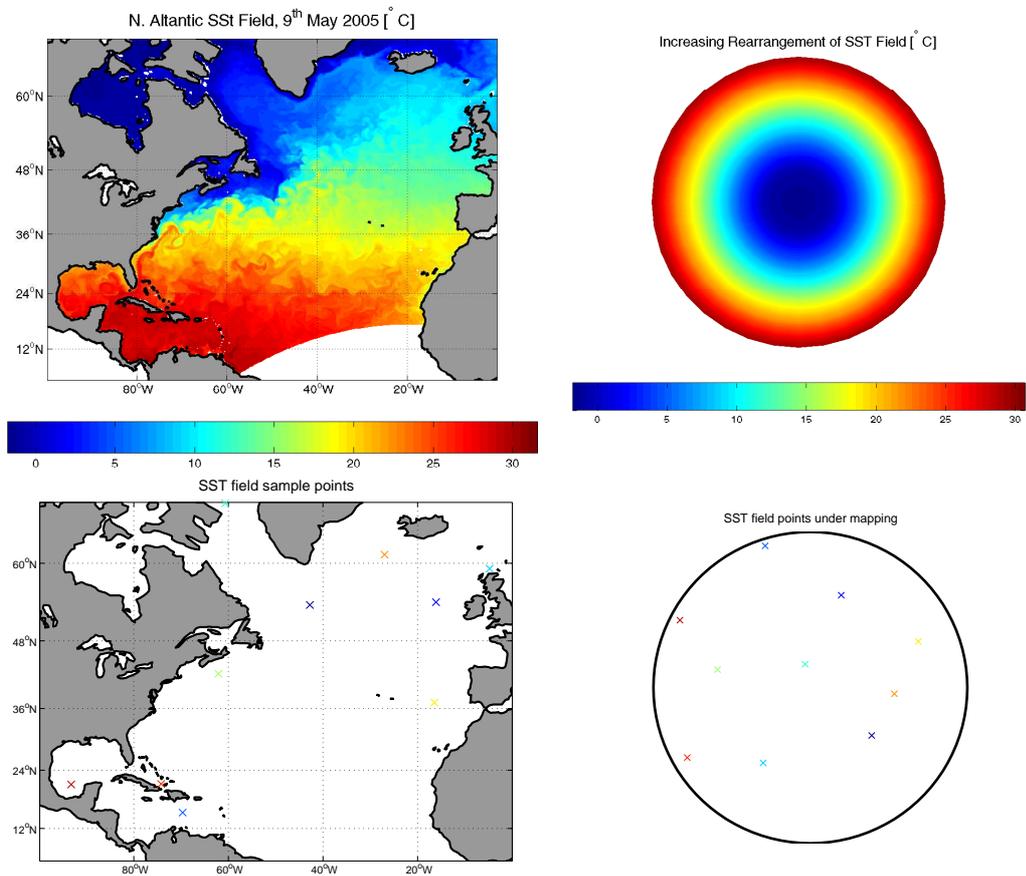
Figure 3.3: Polar factorization of an SST field

The plot on the top left shows an analysis SST field for the North Atlantic from the UK Met Office FOAM model. The plot on the top right shows an increasing rearrangement of the first field, calculated using a simple Matlab algorithm to reorder the data from the first plot under the assumption that the data represent area integrated values. The increasing rearrangement is plotted on a disc of the same surface area as the North Atlantic domain. The complete polar factorization consists of both the increasing rearrangement and the invertible, measure preserving mapping which takes points on the image top left to that on to those top right or vice versa. The plots on the bottom illustrate this; points of equal colour are mapped onto each other under s.

generated by advection by a spatially constant velocity, the only possible choice for an incompressible velocity field in one dimension in a periodic domain. In higher dimensions then we assert that given initial and final conditions which satisfy the rearrangement condition then the transformation of one to the other can weakly be achieved by non-divergent advection, in the sense detailed below. This proposition is important in justifying the use of transport processes for data assimilation in the chapters which follow, particularly Chapter 4, since we show that all rearrangements may be achieved in this way. A key point, however is that the proof does not provide a simple way to find the necessary advection field, at least in finite time. The assimilation algorithms we will later present are all attempts to address this problem efficiently.

Proposition 3.2

Given two square-integrable density functions, f, g on a closed domain $\mathcal{D}$ of dimension greater than one, such that f is a rearrangement of g, (denoted from here onwards by $f \sim g$) it is possible to find a state that is the limit of advection over an infinitely long period (or, equivalently, infinitely energetic advection over a finite period), where the advection maps an initial state f to a state $\hat{g}$, and $\hat{g}$ satisfies the equation

$$\mathcal{J}(g, \hat{g}) := \int_{\mathcal{D}} (g - \hat{g})^2 dV = 0. \tag{3.4}$$

An outline for a pictorial proof is given in Figure 3.4 and the basic steps are given in the following section.

OUTLINE PROOF OF PROPOSITION 3.2

Suppose we have two real-valued, bounded, integrable scalar functions, f, g on a domain ,$\mathcal{D}$, which is a simply connected, closed subset of $\mathbb{R}^d$ with $d > 1$ and that f and g are rearrangement of each other. For any integer N we may partition a subset of the interior of $\mathcal{D}$ into N d-dimensional sub-domains of equal measure, $\delta\mathcal{D}_i$. Define a function on $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$,

$$\mathcal{W}_{[a,b]}(x) = \begin{cases} 1 & a < x < b, \\ 0 & \text{otherwise}, \end{cases}$$

for $a, b, c \in \mathbb{R}$ and consider the domains $U_{[a_j, a_{j+1}]}(f)$, $U_{[a_j, a_{j+1}]}(g)$ defined by,

$$U_{[a_j, a_{j+1}]}(f) = \left\{ \boldsymbol{r} \in \bigcup_{i=1}^{N} \delta\mathcal{D}_i \mid \int_{\delta\mathcal{D}_i} \mathcal{W}_{[a_j, a_{j+1}]}(f) dV = V(\delta\mathcal{D}_i) := \int_{\delta\mathcal{D}_i} dV \right\},$$
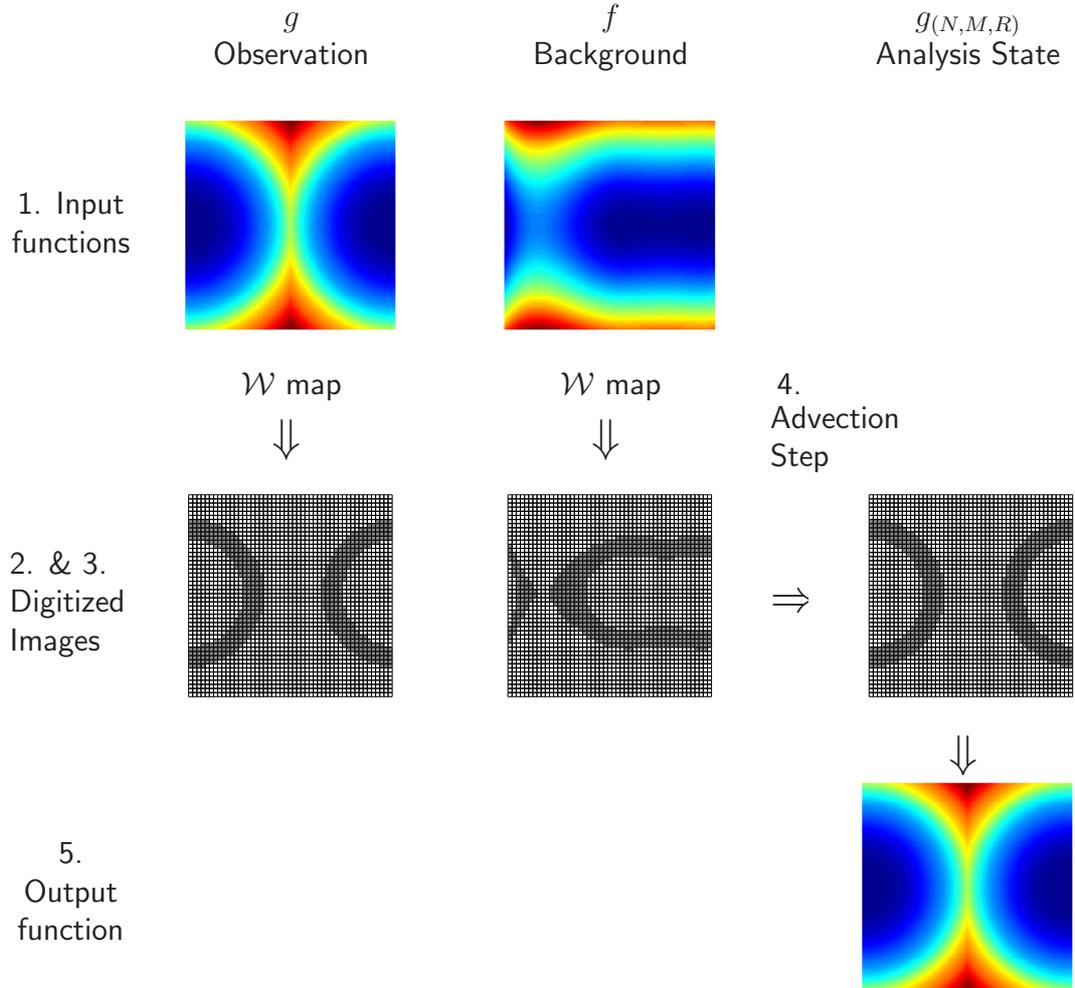
Figure 3.4: Graphical Representation of the algorithm in Proof of Assertion 3.2

1. Two two-dimensional input functions, f and g, to be identified with a model background field and full observational data are provided.

2. The domain of the two functions is quantized into boxes.

3. Digital values are assigned to the boxes, depending on whether the volume of the portion of the box the function maps to a particular interval is equal to the box volume. One particular interval digitization is shown.

4. For each interval of the function co-domain the digitized 'true' areas of the background field are rearranged through the advective process of Brenier and Gangbo (2003) to match the digitized 'true' areas of the observed field.

5. This generates an output analysis field.

As the grid length of the boxes in step (2) and the size of the intervals in step (3) are reduced, then successively better approximations to the observations are produced. For vanishingly small steps the limit of equation (3.4) is achieved.

for some chosen partition, $\mathcal{P} = \{a_1, a_2, \ldots, a_{M+1}\}$, of the bounded co-domain of f and g, with the a_i a set of $M + 1$ monotonically increasing real numbers. From Lemma 1.3 of Brenier and Gangbo (2003) there exists a flow field which in unit time exchanges the contents of neighbouring domains and, by extension, one which attains any arbitrary reordering of the material contained in the $\delta \mathcal{D}_i$, up to a remainder space of arbitrarily small measure, R. It follows that we may produce from f a new state, $g_{(N,M,R)}$, through advection, with maximal intersection of $U_{[a_j, a_{j+1}]}(g)$ and $U_{[a_j, a_{j+1}]}(g_{(N,M,R)})$, for all $j = 1, \ldots, M$. Denote the domain over which the mapping occurs by U^N. Consider the value of the cost function,

$$J(g, g_{(N,M,R)}) := \int_{\mathcal{D}} \left(g - g_{(N,M,R)}\right)^2 dV < \mathcal{J}_N + \mathcal{J}_M + \mathcal{J}_R,$$

where $\mathcal{J}_N$ is an upper bound on the difference between f and g integrated over the space not included in U^N,

$$\mathcal{J}_N = \left[\sup_{x \in \mathcal{D}/U^N} \{f, g\} - \inf_{x \in \mathcal{D}/U^N} \{f, g\} \right]^2 V(\mathcal{D}/U^N),$$

the second term, $\mathcal{J}_M$, is a bound on the discrepancy between g and $g_{(N,M,R)}$ over the space contained in $U^N := \bigcup_{j=1}^{M} U_{[a_j, a_{j+1}]}(g_{(N,M,R)})$, using the definition of the $U_{[a_j, a_{j+1}]}$ to note that the volume of the integrand can be bounded by the a_j,

$$\mathcal{J}_M = \max_{j \in [1, \ldots, M]} |a_j - a_{j+1}|^2 \sum_{\delta D \in U^N} V(\delta \mathcal{D}^i),$$

and $\mathcal{J}_R$ is the remainder term from the space in U^N unchanged by the flow,

$$\mathcal{J}_R = \left[\sup_{\boldsymbol{r} \in U^N} \{f, g\} - \inf_{\boldsymbol{r} \in U^N} \{f, g\} \right]^2 V(R).$$

As $M \to \infty$, $\mathcal{P}$ the partition of the bounded codomain may be chosen so that $\max_{\mathcal{P}} |a_j - a_{j+1}|^2 \to 0$, while the total volume contained in $\sum_{i=1}^{N} \delta \mathcal{D}_i$ remains bounded by volume of the full domain $\mathcal{D}$, hence $\mathcal{J}_M \to 0$. Similarly, for fixed M, integrability ensures that as $N \to \infty$, $U_{\mathcal{P}}(g) \to \mathcal{D}$, while the rearrangement condition ensures $U^N \to U_{\mathcal{P}}(g)$. Hence $V(\mathcal{D}/U^N)$ may be made arbitrarily small through the choice of N, giving $\mathcal{J}_N \to 0$. Finally, the volume of the space R is arbitrary, by the definition of the theorem given above and by taking a suitable long or energetic integration of the flow field we may take $\mathcal{J}_R \to 0$. Combining these separate processes we may chose a sequence of (N, M, R), increasing monotonically in N and M and choosing R of successively smaller volume, such that the limit of

$\mathcal{J}(g, g_{(}N, M, R))$ vanishes.

3.1.3 THE MONGE-KANTOROVICH PROBLEM

The algorithm presented for the polar factorization approach to data assimilation with conservation properties is direct, in the sense of immediately producing an analysis field, given inputs of full observations and a background field. Unfortunately, the method as presented is also univariate, assimilating only the observed field. Although the rearrangement condition (3.3) can be extended into a consistent condition on vector fields and thus to assimilating multiple observed fields over a single domain[4], of itself the method cannot be made truly multivariate, in the sense of generating increments of fields in the null space of the observation operator. Such increments must be generated, either through assumed statistical relations with the observed changes, or derived physical correlations. One physical argument, implicitly assumed in the rearrangement of vector fields, is that a single mapping must be applied concurrently to all material variables, i.e. that the rearrangement is applied as a displacement of fluid parcels, rather than to atoms of the individual variable fields.

If the rearrangement of fluid parcels in question is to be assumed the result of a physical process, i.e. the result of advection by an assimilation velocity field, u_a, then the problem becomes one of determining a suitable value for this u_a. This has an obvious connexion with a class of transport problems known as the Monge-Kantorovich problems. The original physical problem Gaspard Monge considered in his 1781 monograph *"Memoire sur la théorie des déblais et de remblais"* was the rearrangement of a collection of rubble (*reblais*) to fill an equally sized space (*déblais*) in another location, while minimizing the effort (defined as the distance an individual particle of *reblais* must travel) involved in the process. Monge's proof was geometric, based on the existence of a single shortest line tangent to both spaces. Kantorovich formulated a relaxed and more general form of this problem in the more modern language of measure theory.

The Monge-Kantorovich formulation is widely discussed in the literature (see for example Gangbo and McCann (1996)) and has a clear structure for existence arguments, as well as uniqueness for certain choices of distribution. The connection with rearrangement theory comes when these distributions are identified with the background and analysis variable fields respectively. This then shows that, given a

---

[4]Similarly the polar factorization can be extended for functions $f : \mathcal{D} \to \mathbb{R}^n$ to become a convex function ψ on $\mathcal{R}^n$ plus a mapping $s : \mathcal{D} \to \mathbb{R}^n$ such that $f = \nabla_{\mathbb{R}^n} \psi \circ s$ almost everywhere (Brenier, 1991).

suitable choice for the penalty on the possible rearrangements which generate the assimilation field, a unique choice of the mapping, s, may be chosen. This in turn allows for multivariate assimilation in all material model variables.

Algorithm 3.1 (Multivariate rearrangement assimilation)

1. *Calculate (non-unique) polar factorizations of observed and background fields.*

2. *Calculate analysis state for observed variables using formula*

$$C_a(\boldsymbol{r}) = f_b(s_y(\boldsymbol{r})).$$

3. *Calculate mapping $s_K : \mathcal{D} \to \mathcal{D}$ which minimises relevant Kantorovich problem (see below).*

4. *Calculate the new multivariate analysis using formula*

$$x_a(\boldsymbol{r}) = x_b\left(s_K(\boldsymbol{r})\right).$$

The resulting Kantorovich problem may be solved numerically in a number of ways. The technique presented here, from Benamou et al. (2000), is chosen due to its similarity to the 4D-VAR data assimilation method and its relevance to later work from chapter 4 onwards. The Kantorovich problem, when written in the standard form in terms of the marginals of a joint distribution, has no explicit time component. However, one may be introduced by restructuring the problem in the terminology of fluid dynamics.

Problem 3.2 (A fluid-dynamical transport problem) *Given a measure of the work done by an advective process,*

$$\mathcal{E}(\rho, \boldsymbol{u}) = \int_0^1 \int_{\mathcal{D}} F(\boldsymbol{u}, \rho) dV d\tau, \tag{3.5}$$

for F some integrable function of the velocity field and density, find the $(\rho^*, \boldsymbol{u}^*)$ which achieves the infimum of E over the set of all $\rho, \boldsymbol{u}$ which satisfy the boundary value problem,

$$\frac{\partial \rho}{\partial \tau} + \boldsymbol{u} \cdot \nabla \rho = 0, \qquad \text{in } \mathcal{D},$$

with end conditions

$$\rho(0) = \rho_0,$$
$$\rho(1) = \rho_1.$$

The technique operates by iteratively minimizing (a discretization of) the above cost function using an augmented Lagrangian method (Benamou and Brenier, 2001) to impose the constraints. The desired mapping can then be found either by following the Lagrangian path of fluid parcels under this velocity, or by the direct advection of model fields by the u^*.

This type of method has three difficulties as a solution of the oceanographic data assimilation question :

(i). As formulated, the method is univariate, generating an improved analysis field only for the variable for which observations are available. To generate corrections for other variables it is necessary to find a way to generate a unique mapping and apply some physical argument to give a relationship between variables.

(ii). The method works only for variables which are directly observed. Often, particularly with satellite data, the observation obtained is[5] a non-linear function of the physical variable of interest.

(iii). The method requires a full observation field. The typical oceanographic data field is extremely sparse, with many 'data holes' where observations are not available or are not reliable. See the example fields presented in Figure 1.1 in Chapter 1.

3.2   A METRIC ON PHASE AND DISPLACEMENT ERROR

The Monge-Kantorovich techniques discussed above illustrate the dependence of all data assimilation methods on the cost function or metric chosen to measure the disparity with observations and the assumed cost to the integrity of the background solution in modifying it. Comparing the QG layer model developed with the original motivation for the use of adiabatic assimilation methods for ocean assimilation, as related in section 2.1, it is obvious that layer vorticity will play the role of the advected tracer, forced only in the top layer. It is no longer clear however that the rearrangement condition is a useful approximation to the tracer behaviour on short time periods. In this section a metric or index on displacement error is introduced with the metric behaviour discussed in relation to the standard $\mathcal{L}_2$ norm on functions. This metric will then be used to test the approximate conservation of the QG vorticity field under various conditions.

---

[5]At least, is usually assumed to be.

3.2.1   A Rearrangement Condition Metric

As defined previously, two fields f and g on a domain $\mathcal{D}$ are termed rearrangements if the quantity,

$$\mathcal{I}(\alpha) := \int_{\mathcal{D}} (f - \alpha) \, H[f - \alpha] dV - \int_{\mathcal{D}} (g - \alpha) \, H[g - \alpha] dV, \qquad (3.6)$$

vanishes for all real α, with $H[x]$ the Heaviside function as before. This is simply a shuffling of the terms of the rearrangement condition, (3.3). More generally however, provided f and g are bounded, integrable functions and both contain the same mass, that is provided

$$\int_{\mathcal{D}} f dV = \int_{\mathcal{D}} g dV,$$

then $\mathcal{I}(\alpha)$ has bounded support in α and the quantity

$$\mathcal{M} := \int_{-\infty}^{\infty} |\mathcal{I}| d\alpha \qquad (3.7)$$

is finite. The proof of these claims is simple. Let $a = \inf_{\mathcal{D}}\{f, g\}$, $b = \sup_{\mathcal{D}}\{f, g\}$. Then for $\alpha < a$ we have

$$\mathcal{H}(f - \alpha) = \mathcal{H}(g - \alpha) = 1.$$

For $\alpha < a$ we may thus rearrange $\mathcal{I}(\alpha)$ to give

$$\mathcal{I}(\alpha) = \int_{\mathcal{D}} (f - \alpha) - (g - \alpha) dV = \int_{\mathcal{D}}) f - g) dV = 0,$$

where we have used the equal mass condition. Similarly and more trivially, for $\alpha > b$

$$\mathcal{H}(f - \alpha) = \mathcal{H}(g - \alpha) = 0$$

and the integrands in $\mathcal{I}(\alpha)$ vanish. Therefore the support in $\mathcal{I}$ is only in the interval $[a, b]$.

The quantity $\mathcal{M}$ vanishes for the case of $f \sim g$ and is otherwise non-zero for most non-pathological pairs of fields. As such $\mathcal{M}(f, g)$ represents an alternative function for the (non-local) difference between fields of equal mass on a domain, or alternatively a measure of the minimum magnitude of the difference between the two fields, where differences in phase are disregarded. Using the bounded support

of $\mathcal{I}$ the integral in the definition of $\mathcal{M}$ may be defined equivalently as

$$\mathcal{M} := \int_a^b |\mathcal{I}| d\alpha.$$

Using the integral identities

$$\left| \int_{\mathcal{D}} \int_a^b w(\alpha) d\alpha dV \right| \leq \int_{\mathcal{D}} \left| \int_a^b w(\alpha) d\alpha \right| dV \leq \int_{\mathcal{D}} \int_a^b |w(\alpha)| d\alpha dV$$

we may define upper and lower bounds for the magnitude of $\mathcal{M}$. Defining new functions by pointwise comparison of f and g,

$$c(\mathbf{r}) = \min\{f(\mathbf{r}), g(\mathbf{r})\},$$

$$e(\mathbf{r}) = \max\{f(\mathbf{r}), g(\mathbf{r})\},$$

and assuming without loss of generality that $a = 0$, so that f and g are positive functions[6]. The upper bound is then given by

$$\mathcal{M}(f,g) \leq \int_{\mathcal{D}} \left[\int_0^c |f-g| d\alpha + \int_c^e |e-\alpha| d\alpha \right] dV$$
$$\leq \int_{\mathcal{D}} c|f-g| + e(e-c) - \frac{e^2-c^2}{2} dV.$$

It follows that $\mathcal{M}$ is finite. Using the trivial identities

$$|e \pm c| = |f \pm g|,$$

$$e^2 - c^2 = (e+c)(e-c),$$

we see that

$$\mathcal{M}(f,g) \leq \int_{\mathcal{D}} \frac{|f-g||f+g|}{2} dV.$$

Considering the application of this bound to $\mathcal{M}(f, f+\delta f)$ for an otherwise arbitrary perturbation δf which satisfies

$$\int_{\mathcal{D}} \delta f dV = 0$$

we obtain

$$\mathcal{M}(f, f+\delta) \leq \frac{\|\delta f\|^2}{2} + \|f\| \|\delta f\|, \tag{3.8}$$

[6]Generality follows from the fact that $\mathcal{M}(f,g) = \mathcal{M}(f-k, g-k)$ for any constant function k, including the function with k=a.

so that as the standard quadratic cost function,

$$\mathcal{J}(f, f + \delta f) := \frac{1}{2} \int_{\mathcal{D}} (f - [f + \delta f])^2 = \frac{\|\delta f\|^2}{2}, \tag{3.9}$$

vanishes for fixed f as $\|\delta f\| \to 0$, so does $\mathcal{M}$, although possibly more slowly. In other words, any method which reduces the total error present in the system, as measured by the metric $\mathcal{J}$, will naturally reduce the total amplitude error, measured by our new metric, $\mathcal{M}$. The converse of this, that $\mathcal{J}$ cannot vanish for non-zero $\mathcal{M}$, is simply a restatement of the property that rearrangement methods (which by their very nature will preserve $\mathcal{M}$) cannot reproduce perfect fields in the presence of pure amplitude error. The lower bound on $\mathcal{M}$ is derivable by a similar analysis; however, this proves simply that

$$\mathcal{M}(f, g) \geq \left| \|f\|^2 - \|g\|^2 \right|,$$

which is less edifying.

We now have sufficient results to prove the following result :

Theorem 3.1 *$\mathcal{M}$ is a mathematically well defined metric on the set of equivalence classes of the rearrangement operator on continuous functions over bounded subsets of $\mathbb{R}^n$.*

Proof: We will simply show by direct calculation that $\mathcal{M}$ is positive definite, is zero only for identical inputs, is symmetric and satisfies the triangle identity by applying it to fields f, g over some fixed domain.

(i). $\mathcal{M}(f, g) \geq 0$ follows trivially from the modulus sign in the definition of $\mathcal{M}$.

(ii). $\mathcal{M}(f, g) = 0$ if $f \sim g$ follows directly from the definition of $\mathcal{I}(\alpha)$. The converse follows from continuity of $\mathcal{I}$ in α. If $\mathcal{M} = 0$ then $\mathcal{I} = 0$ almost everywhere and from the definition of $\mathcal{M}$ there exists an open neighbourhood round any non-zero point where $\mathcal{I}$ must vanish. So for any point which has $\mathcal{I}(\alpha_0) = K$ we know that for small enough δ

$$|\mathcal{I}(\alpha_0) - \mathcal{I}(\alpha_0 + \delta)| = |\mathcal{I}(\alpha_0) - 0| = K.$$

Hence, on rearranging terms,

$$K = \left| \int_{\mathcal{D}} (f - \alpha_0)\left(\mathcal{H}(f - \alpha_0) - \mathcal{H}(f - \alpha_0 - \delta)\right) + \mathcal{H}(f - \alpha_0 - \delta)\delta dV \right.$$
$$\left. - \int_{\mathcal{D}} (g - \alpha_0)(\mathcal{H}(g - \alpha_0) - \mathcal{H}(g - \alpha_0 - \delta)) + \mathcal{H}(g - \alpha_0 - \delta)\delta dV \right|$$
$$\leq 4|\delta| \int_{\mathcal{D}} dV \to 0 \text{ as } \delta \to 0,$$

thus $\mathcal{I}(\alpha) = 0$ everywhere and $f \sim g$ as claimed. This means the measured distance between two fields is zero if and only if they are in the same equivalence class. That is, identical under rearrangement.

(iii). $\mathcal{M}(f, g) = \mathcal{M}(g, f)$, again trivially from the definition.

(iv). $\mathcal{M}(f, h) \leq \mathcal{M}(f, g) + \mathcal{M}(g, h)$ follows directly from the standard triangle inequality,

$$|\mathcal{I}(\alpha; f, h)| = |\mathcal{I}(\alpha; f, g) - \mathcal{I}(\alpha; g, h)| \leq |\mathcal{I}(\alpha; f, g)| + |\mathcal{I}(\alpha; g, h)|.$$

So $\mathcal{M}$ is a metric as claimed.

3.2.2 METRICS BASED ON MINIMIZATIONS

The quantity $\mathcal{M}$ is conserved over rearrangements of f and g, a property inherited from $\mathcal{I}$. This means that the $\|\delta f\|^2/2 = \mathcal{J}$ terms in the bound (3.8) can be replaced by the tighter limit

$$\mathcal{M}(f, g) \leq \mathcal{J}^* + \sqrt{2}\|f\| \sqrt{\mathcal{J}^*},$$

where $\mathcal{J}^*$ is the limit on the magnitude of $\mathcal{J}$ as defined in equation (3.9) over rearrangements found, for example, by the method of polar factorization,

$$\mathcal{J}^*(f, g) = \min_{\hat{f} \sim f} \mathcal{J}(g, \hat{f}).$$

Therefore the metric defined by $\|f - g\|_*^2 = \mathcal{J}^*(f, g)$ is stricter than $\mathcal{M}$[7]. Calculating $\mathcal{J}^*$ requires the generation of the polar factorization of both f and g however, whereas $\mathcal{M}$ may be found by direct evaluation of the relevant integrals.

We have already met a class of phase error metrics in section 3.1.3. For suitable concave functions, $F(\rho, \boldsymbol{u})$, the minimum of the quantity, $\mathcal{E}$, from (3.5) over the

---

[7]The proof that the norm given by $\mathcal{J}$ also defines a metric over the rearrangement classes is trivial; all results are inherited from $\mathcal{J}$.

relevant constraints defines the square of a metric over the set of rearrangements of f. For $F = \rho|\boldsymbol{u}|$ this metric is the mass weighted sum of the Eulerian distance fluid parcels must travel,

$$d_1(f,g)^2 = \inf_{\gamma \in \Gamma} \int |x - y| d\gamma(x, y),$$

in the measure notation of Section 3.1.3. This quantity is known as the 1st Wasserstein distance. Similarly minimizing $\mathcal{F} = \rho|\boldsymbol{u}|^p/p$ gives the p^{th} Wasserstein distance,

$$d_p(f,g)^2 = \inf_{\gamma \in \Gamma} \int |x - y|^p d\gamma(x, y).$$

Applying any of these distance functions to calculate the distance between f and $\hat{f}$, where $\hat{f}$ is the minimizer of $\mathcal{J}(g, \tilde{f})$ for $\hat{f}$ a rearrangement of f found though the polar factorization method of Section 3.1.2 gives a definition of a unique phase error metric. However, it does not automatically follow that states with large $\mathcal{J}(f,g)$, but small $\mathcal{J}^*(f,g)$, will have large $\mathcal{E}(f,\hat{f})$. Similarly, the two metrics $\mathcal{E}(f,\hat{f})$ and $\mathcal{E}(g,\hat{g})$, may differ, except in the trivial case when f and g satisfy the rearrangement property.

When used to diagnose phase and amplitude errors in practical data assimilation problems, all the metrics discussed here have two main limitations. They all require full knowledge of the functions f and g and they implicitly assume that any error that it is possible to ascribe to phase error *must* be due to phase error. The first issue may be relaxed by only considering retrievals (that is, entire data fields obtained through *post hoc* transformation and extension of observed data through other means) rather than observations in the case of limited observations. Such an approach would on the other hand be likely to contribute to the second problem.

	$\mathcal{J}(f,g)$ low	$\mathcal{J}(f,g)$ high
$\mathcal{M}(f,g)$ low	Errors small	Error probably in phase
$\mathcal{M}(f,g)$ high	Impossible	Difference probably in amplitude

Table 3.1: Behaviour of system for various magnitudes of $\mathcal{M}, \mathcal{J}$.

A simple truth table for identifying sources of error from the respective magnitudes of $\mathcal{J}$ and $\mathcal{M}$. Note that from the bounds found on $\mathcal{M}$ we know one case cannot occur. Properly the two cases in the right hand column test the hypothesis that the discrepancy can be represented as a phase error, rather than determining that it is. If $\mathcal{M}$ is large we know phase error is present, if $\mathcal{M}$ is small then phase error *may* exist.

3.2.3 APPLYING THE NEW METRICS TO THE QG MODEL

The behaviour of the metrics $\mathcal{M}$ and $\mathcal{J}$ differ in the presence of amplitude and phase errors, as described in Table 3.1. Using $\mathcal{M}$ and $\mathcal{J}^*$ as metrics on the differences ascribable purely to amplitude between various pairs of states generated the plots shown in Figure 3.5. The two functions were applied to pairs of top level vorticity fields, $q_1(\boldsymbol{r}, t)$ generated by the QG model of Section 2.3 and differing in various ways. In the first three experiments the difference was in the times at which the states were valid (i.e. the times at which the run was stopped), with a difference of 1, 10 and 100 days respectively in validation time. As may be expected through the presence of both a diabatic forcing term and a mixing term in the evolution equation for q_1 the phase difference under both metrics appears to increase with time. The fourth experiment considers two states valid at the same time and both generated by a 100 day evolution of the same initial conditions but with a 10% difference in the value of τ_0 used in the definition of the wind stress, as per equation (2.25). This experiment shows a significantly larger difference in both $\mathcal{M}$ and $\mathcal{J}^*$ than in any of the experiments based on validation time. This is in contrast to the standard $\mathcal{L}_2$ cost function, $\mathcal{J}$, which is actually smaller in the case of differing forcing than for validation times differing by 100 days.

In all plots the final bar on the histogram shows the result of applying the metrics to two fields chosen to satisfy the rearrangement condition. Under this limit both $\mathcal{M}$ and $\mathcal{J}^*$ vanish, while $\mathcal{J}$ may obviously remain non-zero. Comparing the two distance metrics provided by $\|q_A - q_B\| = \sqrt{\mathcal{J}(q_a, q_b)}$ and $\|q_A - q_B\|^* = \sqrt{\mathcal{J}^*(q_a, q_b)}$ it appears that even in the case of differing forcing, phase error explains 80% of the differences between fields. This supports the previous assertion that phase error correction methods are suitable for assimilation in the QG model.

Note that all results presented are in consideration of the top level vorticity field only. The results for the bottom layer field appear less clear. This seems to be in part due to the generally smaller values of $\mathcal{J}$ leading to greater numerical issues in the calculation of the two amplitude metrics. The routines used to calculate $\mathcal{M}$ and $\mathcal{J}^*$ numerically are discussed in the appendix, section A.3.

3.3 COOPER & HAINES FOR THE QG MODEL

3.3.1 THE COOPER & HAINES ALTIMETRY ASSIMILATION METHOD

The Cooper and Haines (1996) (CH) method is a standard data assimilation technique which attempts to apply the idea of rearrangement of fluid parcels, but
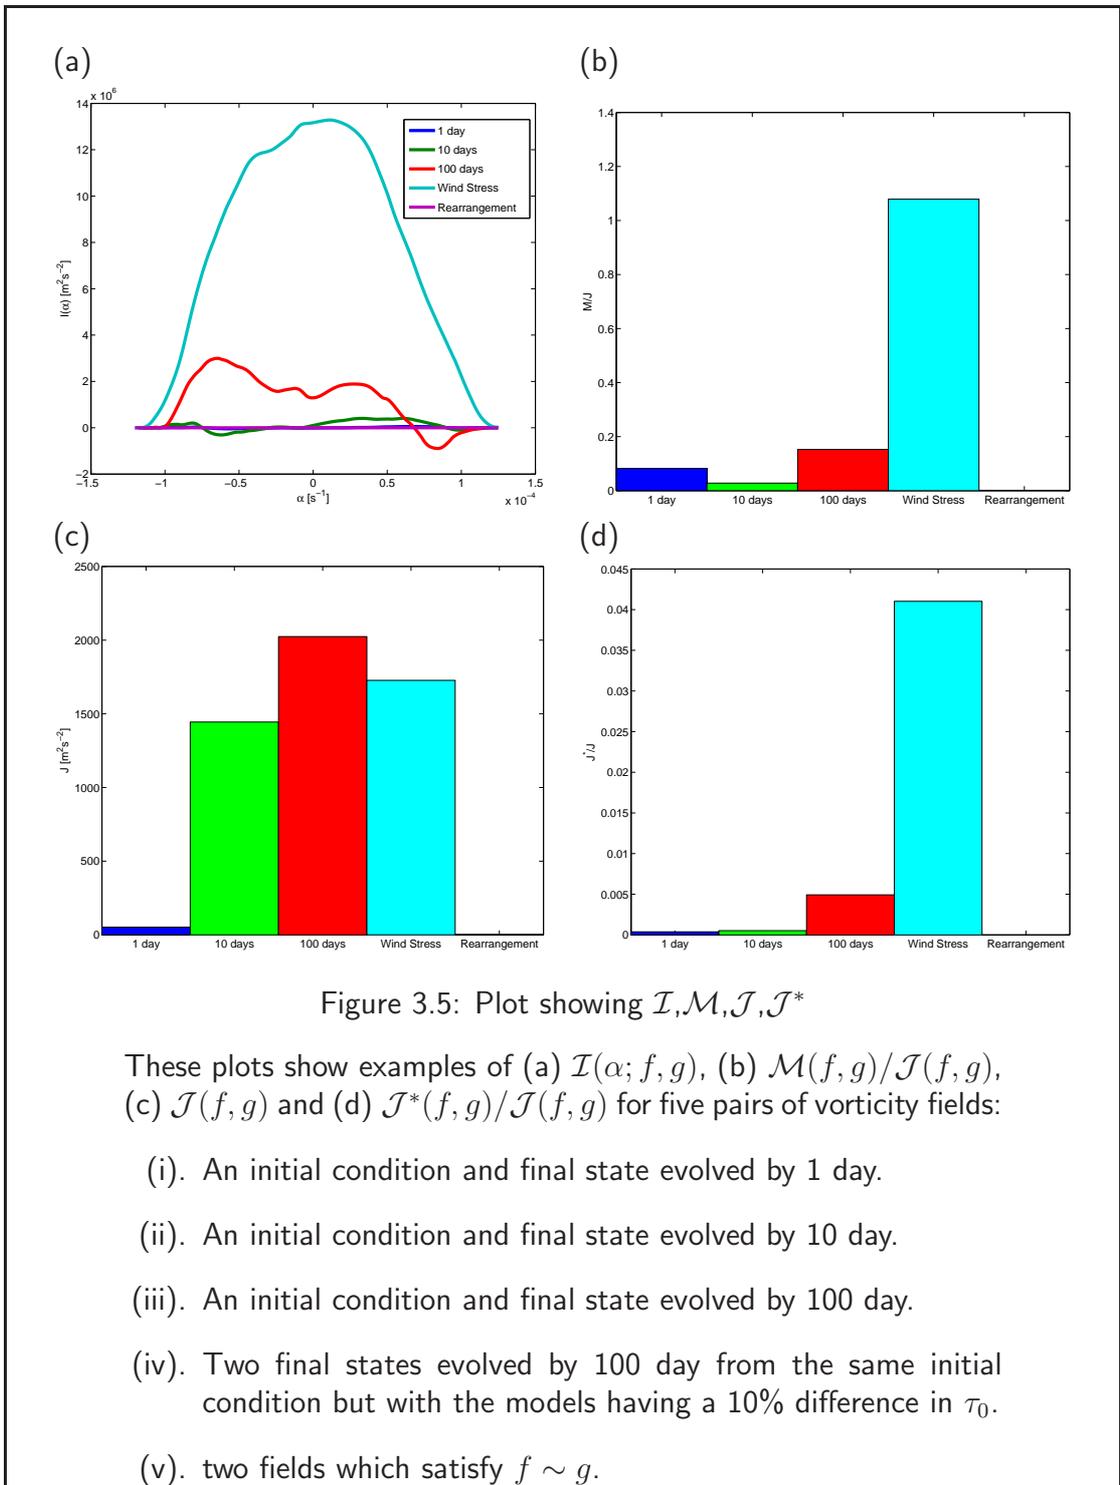
Figure 3.5: Plot showing $\mathcal{I}, \mathcal{M}, \mathcal{J}, \mathcal{J}^*$

These plots show examples of (a) $\mathcal{I}(\alpha; f, g)$, (b) $\mathcal{M}(f, g)/\mathcal{J}(f, g)$, (c) $\mathcal{J}(f, g)$ and (d) $\mathcal{J}^*(f, g)/\mathcal{J}(f, g)$ for five pairs of vorticity fields:

(i). An initial condition and final state evolved by 1 day.

(ii). An initial condition and final state evolved by 10 day.

(iii). An initial condition and final state evolved by 100 day.

(iv). Two final states evolved by 100 day from the same initial condition but with the models having a 10% difference in τ_0.

(v). two fields which satisfy $f \sim g$.

in a simplified one-dimensional sense. We now discuss this method in the context of the QG model presented in chapter 2, before investigating the behaviour of our new metrics under CH assimilation. Given observations of sea surface height (or, equivalently of the ocean pressure field at mean sea-level) from satellite altimetry data, the method attempts to find a vertical shift to the water column below the mixed layer to compensate the observed surface pressure anomaly but such that the pressure, and thus the balanced part of the velocity field, remains unchanged on the lowest dynamic depth level.

Since the CH method is one-dimensional, conservation of water mass properties cannot be maintained on the surfaces where water is entering or leaving the system, (the lowest level above the bottom topography and the first dynamic level below the mixed surface). Instead water properties must be extrapolated from the values and gradients at the ocean floor and mixed layer. It is this, unphysical, change to the column mass which produces the desired correction to pressure at the mixed layer. The CH method can reproduce the observed out-cropping of cold isopycnals in cold core eddies without unduly affecting the (near constant) tracer properties of the deeper layers. Column lowering, however, requires adjustment of the mixed layer properties since the CH method naturally preserves the (too cold) mixed layer temperature and thus presents a cold bias.

3.3.2   SATELLITE ALTIMETRY - OBSERVING THE STREAMFUNCTION

The Cooper and Haines (CH) method was originally derived from considering the behaviour of layer models (Haines, 1991), assimilating information on the behaviour of the top layer streamfunction, ψ_1, using the notation introduced in Section 2.3. This means that the algorithm to apply the CH method to an N-layer quasi-geostrophic (QG) model is far simpler than applying it in a level model with an explicit z component. The condition that bottom-level pressure remains unperturbed is simply

$$\psi_N^{\text{analysis}} = \psi_N^{\text{background}},$$

where the subscript denotes level number. Similarly the constraint that the top level pressure field should be matched to the observed variations in surface height requires simply

$$\psi_1^{\text{analysis}} = \frac{\rho_1 g}{\rho_0 f_0} \eta_1'^{\text{observed}}.$$
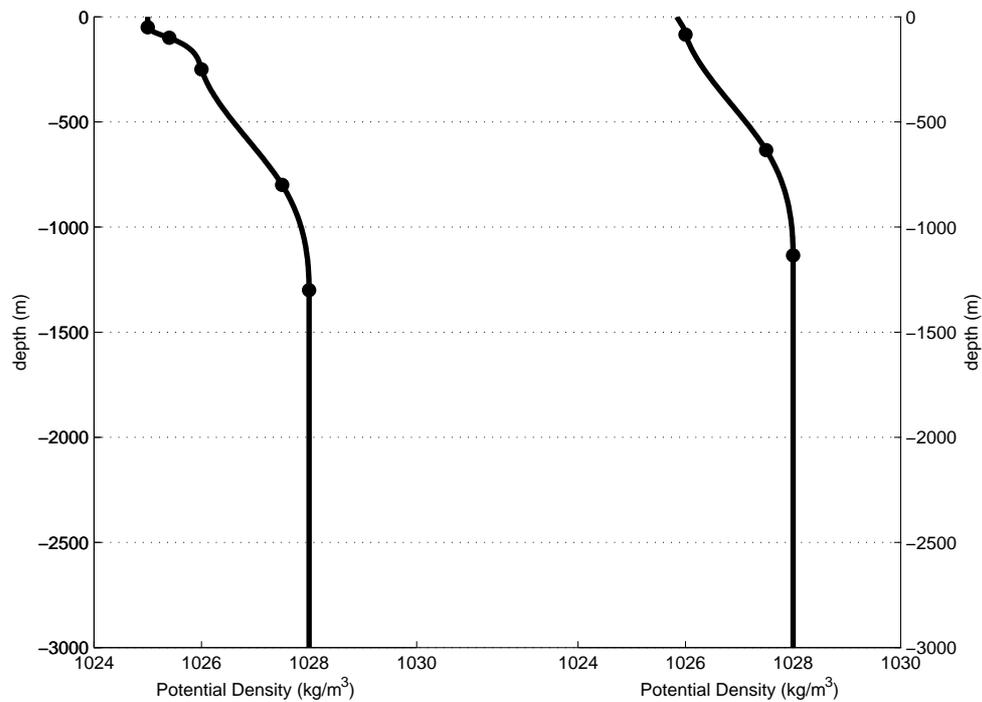
Figure 3.6: Application Of Cooper and Haines Method

A cartoon of the application of the CH method to correct an observed 1000 Pa
low pressure anomaly in sea-surface pressure in a layer model. The profile on
the left represents the unperturbed model profile of potential density. This is
then shifted vertically by 165 m to produce the profile on the right. This profile
matches the observed surface pressure field but leaves the pressure on the lowest
level (z=-3000 m) unchanged. Thus the lowest 165m of this profile represent an
extrapolation of the density at the base of the model. In this case (equivalent to
assimilation of a cold core eddy) the vertical shift is larger than the mixed layer
depth and new cold isopycnals out-crop at the surface. Circles mark the new and
old positions of various fluid parcels; note that the top two parcels are lost.

The changes in other levels follow from the requirement that the shift in interface height be constant at all levels and the equation

$$\Delta\eta_i' = \frac{f\Delta(\psi_i - \psi_{i-1})}{g'} = \Delta\eta_1',$$

holds for each i. Therefore at each level the new analysis streamfunction is given by the relation

$$\psi_i^{\text{analysis}} = \psi_i^{\text{background}} + \frac{\sum_{j=i}^{N} g_j'}{\sum_{j=1}^{N} 1/g_j'} \frac{\rho_1 g}{f_0} \left(\eta_1'^{\text{observed}} - \eta_1'^{\text{background}} \right).$$

For a two layer model this algorithm simplifies further to the one step process

$$\psi_1^{\text{analysis}} = \psi_1^{\text{observed}},$$

$$\psi_2^{\text{analysis}} = \psi_2^{\text{background}}.$$

Hence, in this model, the CH framework is identical to direct insertion of the observed top layer streamfunction into the analysis field, leaving the bottom layer unchanged. This is also related to an algorithm which assumes that the streamfunctions in the top and bottom layers are uncorrelated while applying a 3D assimilation scheme such as optimal interpolation or 3D-VAR to observations of the top layer streamfunction only.

3.3.3 CASE STUDY: ASSIMILATING FULL OBSERVATIONS

The QG CH algorithm was used in twin mode to assimilate the run initialized under a different wind-stress, as used in the metric experiments of Section 3.2.3, into a run starting from the same initial conditions. Both the 'true' and 'background' models were run out for 100 days, to generate a model trajectory discrepancy, as discussed in Section 2.5. The two systems were then run on for a 10 day period, assimilating observations every 24 hours. Three different observation systems were considered, along with a control run with no assimilation:

(i). Full, perfect observations with observations available everywhere,

$$\psi_1^{\text{obs}} = \psi_1^{\text{truth}}.$$

(ii). Full, imperfect retrieval available everywhere,

$$\psi_1^{\text{obs}} = \frac{1}{2}\psi_1^{\text{truth}} + \frac{1}{2}\psi_1^{\text{background}}.$$

Note that we have chosen an unfeasibly small weighting on the observations to make the differences between the true state and analysis more obvious. This is balanced by the assumption of full observations.

(iii). Smoothed retrieval available everywhere,

$$\psi_1^{\text{obs}} = \frac{1}{2}\overline{\psi_1^{\text{truth}}} + \frac{1}{2}\psi_1^{\text{background}},$$

where $\overline{f(\boldsymbol{r})}$ denotes a spatial averaging function on f, with a length scale of approximately 2 grid lengths.

(iv). The control run assumes no observations available,

$$\psi_1^{\text{obs}} = \psi_1^{\text{background}}.$$

The results, in terms of the evolution of the usual quadratic cost function, $\mathcal{J}$, of these experiments are shown in Figure 3.7, with results for the top and bottom layers plotted separately.

As is to be expected, both layers however show eventual convergence towards the true values, as shown by the reduction in $\mathcal{J}$. The saw-toothed pattern in the top-layer cost function represents the points at which the assimilation occurs, when this field is directly forced towards truth. This also explains the rapid convergence towards truth in this layer, even in the cases of non-perfect observations. The behaviour of the bottom layer function is less trivial, showing a lessening in $\mathcal{J}$ in all cases over the course of the assimilation that is more marked than the natural variance shown by the control run[8].

The 'worst case' with only a smoothed retrieval assimilated has almost exactly the same behaviour as the run in which a non-smoothed retrieval is assimilated. This suggests that, at least when such large volumes of data are assimilated, the small scale structures contained in the true field are relatively unimportant compared to the large scale dynamics. Note however that the rate of convergence of the bottom layer streamfunction slows towards the end of the experiment and that

---

[8]That the control run shows also shows a decay in $\mathcal{J}$ is not surprising. This partially represents the decay of the transient mean state signal from the brief perturbation in windstress, used to generate the dissimilar initial states for the assimilation run. This decay occurs because this particular experiment is run under the perfect model assumption.

the final value is higher for both experiments assimilating retrievals. This suggests that the model is settling towards a new dynamic equilibrium in which the assimilation increments act like an additional periodic forcing term. The effects of this on the small scale structure of the model were investigated in a longer term passive tracer experiment.

3.3.4 Effects on Fine Scale Structure

As was previously stated the CH algorithm is not conservative in the QG context. This means that the reduction in $\mathcal{J}$ of the previous section may be achieved through a reduction in amplitude rather than phase error. To illustrate the effect of this on the resulting fine-scale structure, particularly of unassimilated fields and to investigate the relative unimportance of small scales noted in the previous section a passive tracer experiment was performed. Truth and background states were designated as before from a difference in wind stress over a 100 day run. The CH algorithm was then used to assimilate one full smoothed retrieval of top layer streamfunction every 24 hours over a full 30 day period. At the end of this period it was assumed that the transients due to the change in wind forcing would have died out so that any remaining differences represented equilibrium with assimilation. Two new passive tracer fields were then initialized to be spatially identical in both the background and true model states,

$$\tau_x(x,y)_{1,2} = \frac{x}{L_x},$$

$$\tau_y(x,y)_{1,2} = \frac{y}{L_y}.$$

Each τ variable was assumed to satisfy an evolution equation exactly equivalent to the equations for q_1, q_2 so that to a high order, limited only by numerical effects, the $q - \tau$ relations were conserved in each layer. Each run was then continued for a period of 60 days, with continued CH assimilation into the background state every 24 hours, but no assimilation of the τ. The plots in Figure 3.9 shows the result of this experiment in terms of the evolution of a cost function on discrepancy streamfunction and of the tracer fields respectively.

The first point of interest in the graph of $|\mathcal{J}/\mathcal{J}_i|^{1/2}$ (the top plot of Figure 3.9) in the top and bottom layers is that over this longer run there is no obvious convergence in the bottom layer (indeed $\mathcal{J}$ at the end of the experiment is in fact larger than at the beginning), while $|\mathcal{J}|^{1/2}$ converges in the top layer to around a quarter of its initial value. Comparing this with the 50% reduction which applies on each
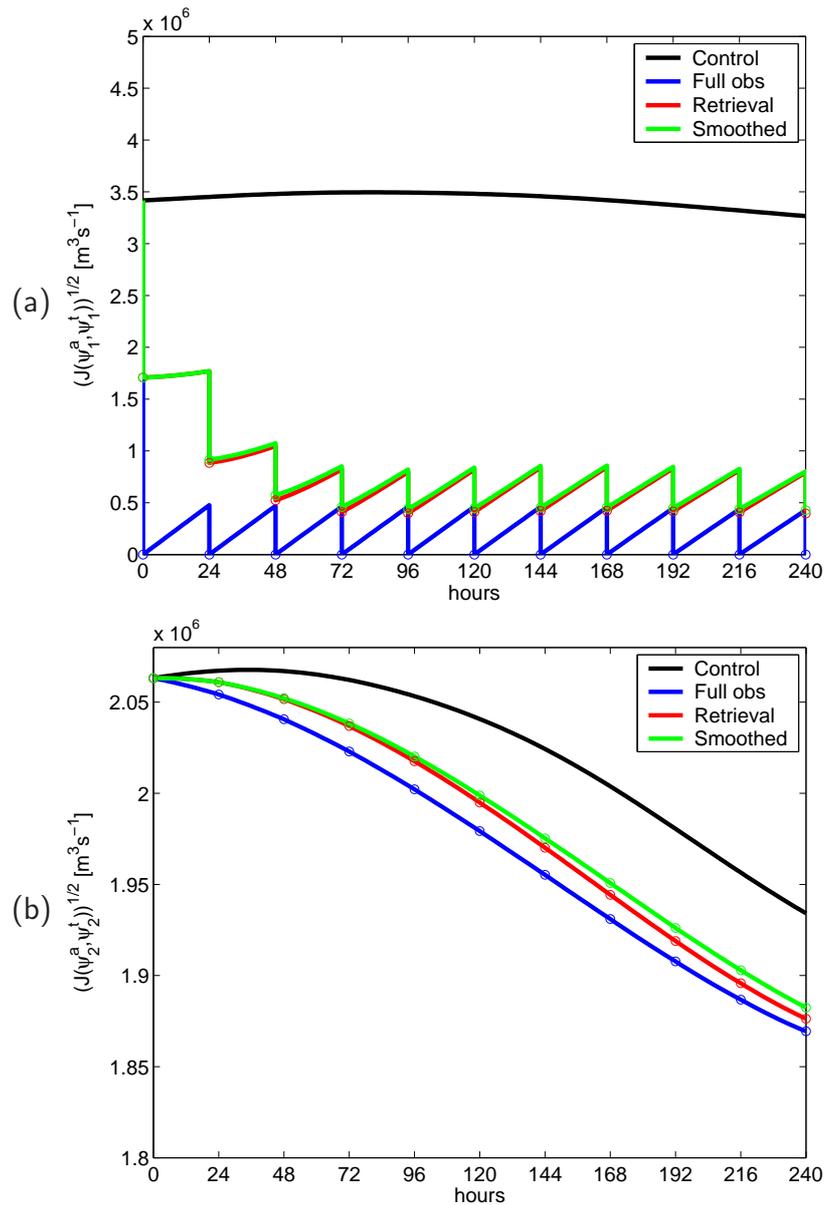
Figure 3.7: Convergence of (a) upper and (b) lower layer $\mathcal{J}$ under CH

The measure of convergence used in each layer is

$$\mathcal{J}(\psi_i^a, \psi_i^t) := \int_{\mathcal{D}_H} (\psi_i^a - \psi_i^t)^2 dS.$$

Circles denote cost function values following CH assimilation steps, when there is an obvious discontinuity in the top layer cost function. The case of full observations is constrained to converge to zero in the top layer, but other assimilation runs show a lack of convergence and a persistent error. The behaviour in the bottom layer is similar in all cases, a slow improvement in the error, levelling off towards the end of the run. Unsurprisingly the case with full observations does best. Note too that the *unassimilated* control run also shows a natural reduction in the lower layer cost function over the course of the run.
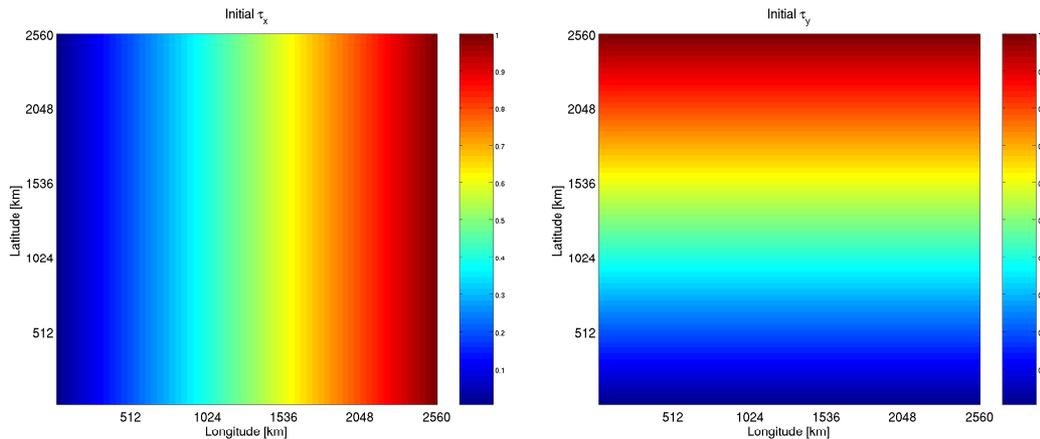
Figure 3.8: Initial τ_x, τ_y

The plots show passive tracer fields, with a zonal and meridional gradient respectively, which are initialized identically in the top and bottom layers of both the true run and the forecast system. This initialization happens after the 100 day spin up period in which trajectories are allowed to diverge and a 30 day initial assimilation period. The final 'truth' and 'analysis' fields are shown in Figure 3.10.

assimilation step gives an error growth rate of $\frac{d|\epsilon|}{dt} = \mathcal{J}_0/8$ day^{-1}, i.e. a doubling time of one day.

Considering the cost functions for the growth of errors in the tracer fields we see that in both cases differences are greater in the bottom layer than in the top, as may be expected since the velocities in the top layer are corrected when the stream function is assimilated. The initial growth of error is faster for the zonal tracer field, τ_y, due to the strong differences in the western boundary current between the true and analysis runs. The errors in each field saturate after 20–50 days, with saturation quicker in the top layer. The final, saturated, value is greatest for the bottom layer of the meridional tracer field, τ_x, again due to the effects of the jet.

Looking at the plots of the true and analysis states at the end of the 60 day experiment we see that the large scale structure of the fields is well represented in all cases. This is in part a result of the relatively small magnitude initial differences between the mean states of the true and background model trajectories. However the small scale structure shows some obvious examples of phase error in the positions of filaments of the tracers in all fields, particularly in the bottom layer, but also notable in the top layer field for τ_y. Amplitude errors are small, a direct result of the way the experiment was designed. This shows that there is a role for dis-

placement assimilation methods in the case of unobserved fields when observations are available on small spatial scales, as is the case in satellite altimetry.

3.4 CHAPTER SUMMARY

In this chapter we have introduced the key theory behind data assimilation in the presence of integral constraints, including presenting the key ideas of rearrangement theory, with regard to assimilating phase errors. These methods were shown to work extremely well in the limited case of pure phase error in fully observed vector fields.

In an attempt to extend these concepts to the more realistic case when only limited observations are available several new metrics have been introduced which quantify the magnitude of the phase, or conservation, error and amplitude, or non-conservative error. Testing these metrics on the two-layer QG model allows limits to be placed on the improvement to be expected under an application of such methods. These are found to be good for small enough space and time-scales (on the order of ten days), but to fail when forcing fields are not known accurately.

For comparison we then presented a data assimilation method for the QG model exactly equivalent to the Cooper and Haines method for assimilation of satellite altimetry. This method was further shown to be derivable from a rearrangement assumption under a constraint on possible movement of fluid parcels. These two derivations show both the useful relative simplicity of this technique and the important limitations that come from the limiting assumption. Finally we presented evidence that such limitations have important damaging effects on small scale fluid structures when the method is applied for cases of noisy or limited data. From these ideas we will now move to developing a similar technique based on instead limiting motions to physical flows along isopycnals.
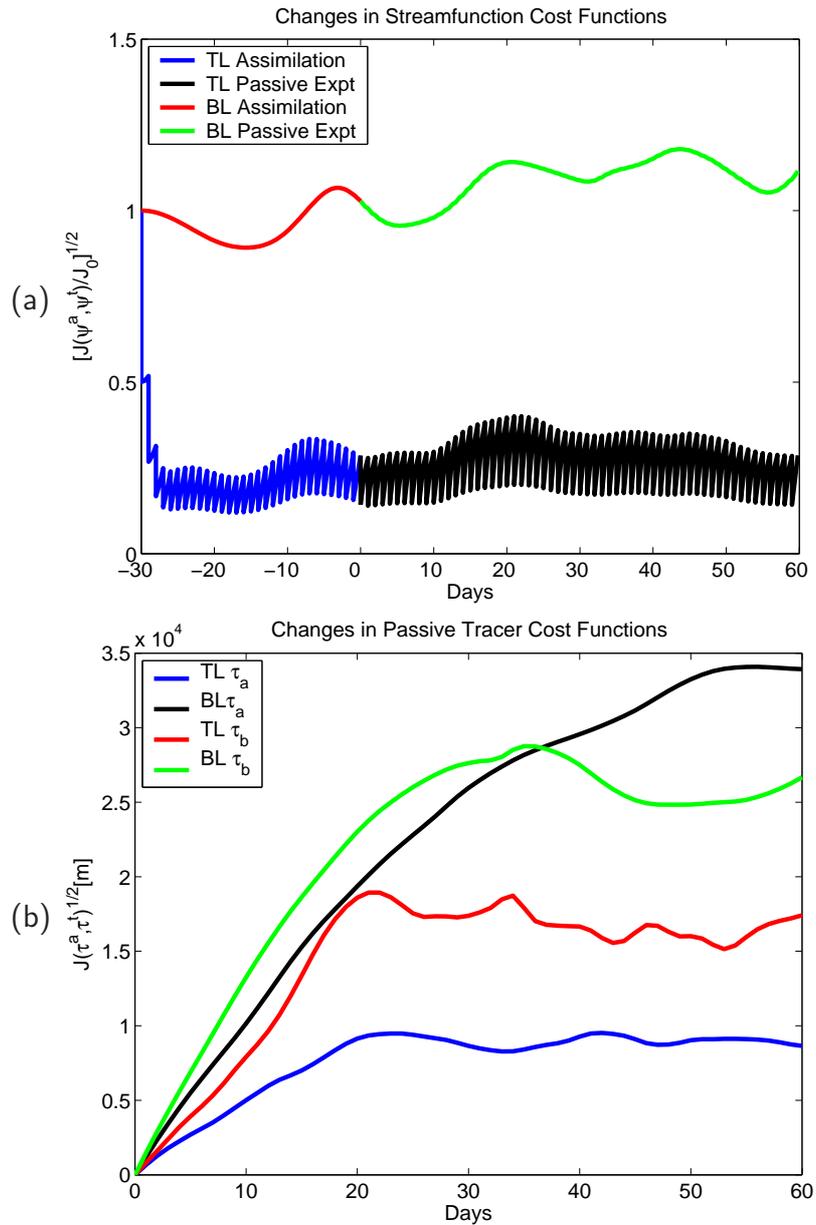
Figure 3.9: Graphs of (a) $\mathcal{J}(\psi_i^a, \psi_i^t)$ and (b) $\mathcal{J}(\tau^a, \tau^t)$

The plots show the variation of the quadratic cost function applied to the evolution of (a) top (TL) and bottom (BL) layer streamfunction and (b) the passive tracer fields from Figure 3.8 in the assimilation experiment. In (a) the time level -30 denotes the start of the assimilation experiment and 0 the time level at which the tracer fields are initialized. These longer runs show the same convergence pattern as the runs in Figure 3.7. Note too the large variation in the cost function in the bottom layer. From (b) we see the saturation of errors in the tracer fields, with the bottom layer saturating more slowly, but with a larger discrepancy.
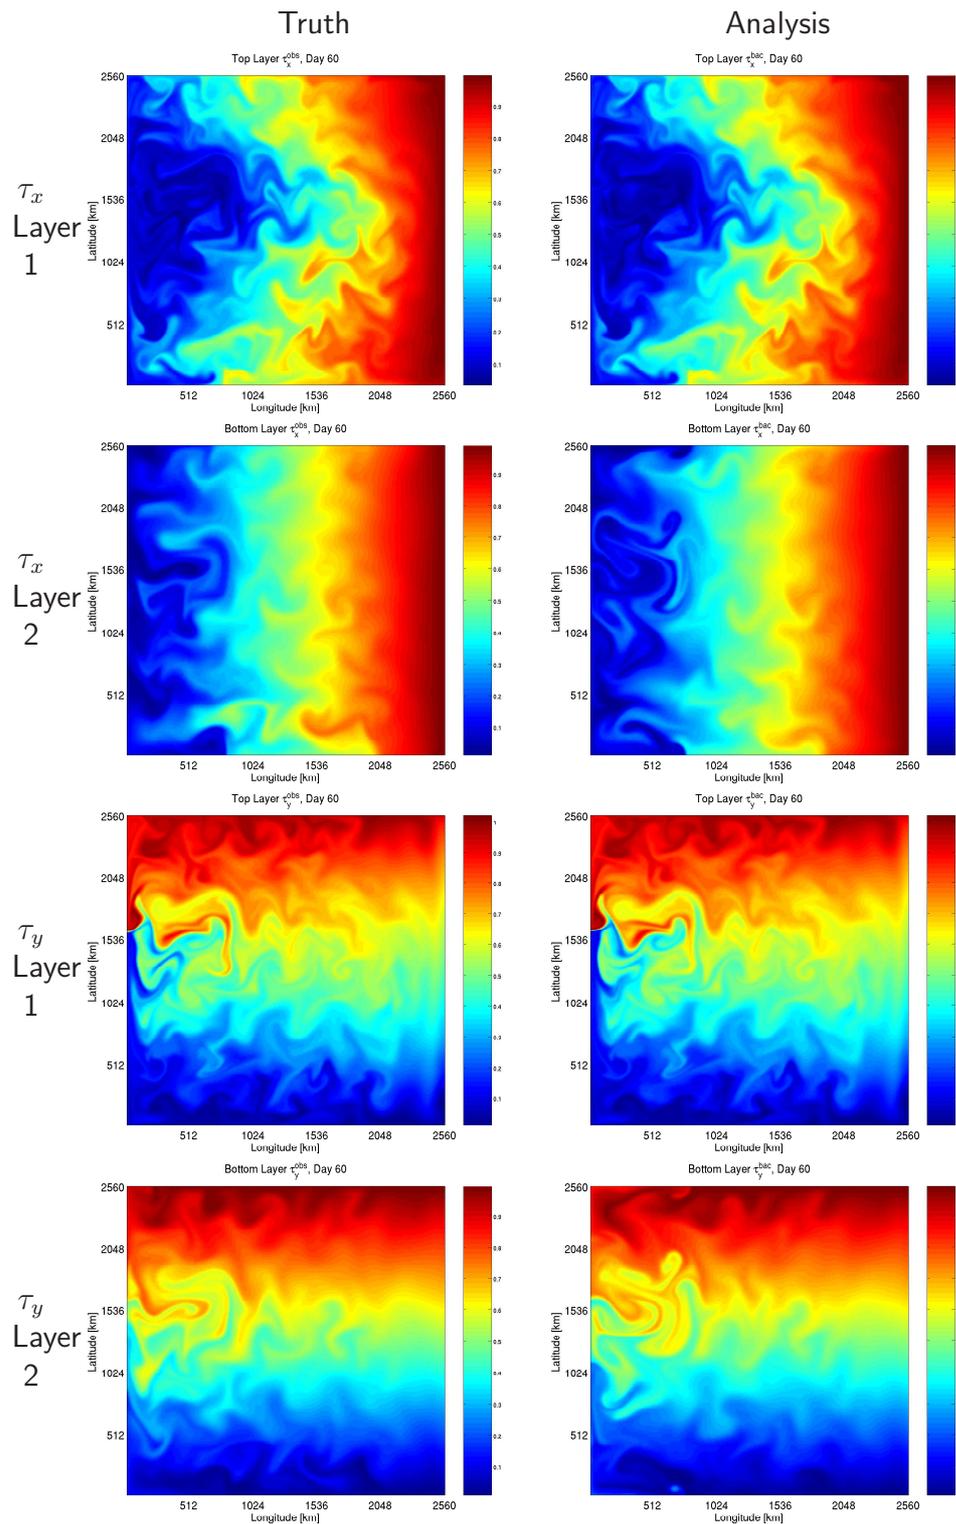
Figure 3.10: Plots of true and analysis tracer fields under CH assimilation

Plots show τ_x and τ_y fields from Figure 3.8 after 60 days, from both the 'true' run and the assimilated analysis field. The top layers appear nearly identical, as may be expected from cost function behaviour. Differences are primarily on the Rossby length of the system. The bottom layer shows more coherent structure to the difference between truth and analysis.

Chapter 4

3D Displacement Assimilation

"Now here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else you must run at least twice as fast as that!"

Lewis Carroll
Alice Through the Looking Glass, 1872

The formulation of the displacement assimilation problem in terms of the rearrangement theory of section 3.1.2, coupled to a Monge-Kantorovich transport problem, as in section 3.1.3 gives an example of an algorithm for an adiabatic displacement scheme. Unfortunately this method is useless for most practical applications in ocean data assimilation. This is due to the twin requirements of full observations and an invertible observation operator. Both these assumptions can fail catastrophically in the real world. This leaves the algorithm suitable only for toy studies or problems where other, non-adiabatic techniques have been previously used to obtain retrievals.

This chapter instead considers a method which retains the adiabatic philosophy of the full rearrangement technique, but which can be consistently applied in the case of limited observations with non-linear, singular observation operators. A family of techniques are developed, based on the use of advection as the adiabatic process which generate paths between rearrangements of a function. This allows the generation of successive rearrangements of an initial field, which tend to minimize an objective cost function based on the difference against observations, similar to the quadratic cost function in standard techniques such as 3D-VAR, or the BLUE equation of Section 1.3.

The methods are investigated for simple toy problems and the QG ocean model of Section 2.3.2. Several techniques to increase the algorithm's rate of convergence are found and the behaviour of the scheme is investigated through use of unassimilated passive tracers and through mapping particle trajectories under assimilation. The scheme is found to be closely related to the Cooper and Haines scheme for stratified data, but behaves markedly differently when assimilating data on isentropic surfaces. Finally the rôle played in the assimilation by the observation field itself is considered, with reference to the cases of limited, noisy and non-model observations.

4.1   Iterative Methods for Displacement Assimilation

This section begins with a review of the notation to be used in the discussion of pure adiabatic data assimilation and with a discussion of the natural connection of this concept with the duality between the Lagrangian and Eulerian perspectives of fluid dynamics. This leads naturally to an iterative data assimilation method exactly equivalent to forced advection of a vector field so as to extremize a penalty functional.

4.1.1   The Minimization Process

Suppose that there exists some first guess prediction for the value of some given (possibly continuous) vector-valued model state variable, $\boldsymbol{x}_b : \mathcal{D} \to \mathbb{R}^n$ of dimension n, which is valid at some chosen time t and defined over the whole domain of interest, $\mathcal{D}$. Since $\mathcal{D}$ will later be taken to represent the ocean we assume it is closed and simply connected[1]. The notation here maintains the use of bold characters for consistency with Section 1.3 when discussing discrete data assimilation and to make obvious the difference from the Cartesian coordinate system of Chapter 2; however, it is important to realise we may still be dealing with infinite dimensional vector spaces rather than limiting ourselves to finite dimensional vectors as in the discrete case. For any particular choice of $\boldsymbol{x}_b \in X$, where X is the set of suitably smooth fields over $\mathcal{D}$, it is possible to define the set $\mathcal{R}(\boldsymbol{x}_b) \subset X$, resulting from the partition of X into equivalence classes of a relation operator, $\sim$,

$$\mathcal{R}(\boldsymbol{x}_b) = \{\boldsymbol{x} \in X : \boldsymbol{x} \sim \boldsymbol{x}_b\},$$

where $\boldsymbol{x} \sim \boldsymbol{y}$ denotes equality in the rearrangement condition of section 3.1.2, namely $\boldsymbol{x} \sim \boldsymbol{x}_b$ if and only if for all $\alpha \in \mathbb{R}^n$ the vector field $\boldsymbol{x} \in X$ satisfies element-wise the condition

$$\int_{\mathcal{D}} (\boldsymbol{x} - \boldsymbol{\alpha}) H[\boldsymbol{x} - \boldsymbol{\alpha}] dV = \int_{\mathcal{D}} (\boldsymbol{x}_b - \alpha)\mathcal{H}[\boldsymbol{x}_b - \boldsymbol{\alpha}] dV. \qquad (4.1)$$

---

[1]Studies of regional oceans introduce the complication of open boundaries. This allows the propagation of information into and out of the domain of interest. This is a valid limitation in what follows, but the issue may be addressed by considering the use of nested models to provide boundary information. If the domain considered is not simply connected it should be partitioned into subdomains which are.

with $\mathcal{H}[\boldsymbol{x}] = \Pi_{i=1}^{n}\mathcal{H}[x_i]$ is a multidimensional extension of the Heaviside function, equal to unity only when all components in the $\boldsymbol{x}$ are positive. Given this notation then the general rearrangement problem to be solved in order to perform adiabatic data assimilation may be expressed succinctly as a simple minimization problem over $\mathcal{R}(\boldsymbol{x}_b)$. If the fit of an arbitrary vector field against known observations can be measured by a quadratic cost function with nonlinear observation operator, as in many standard data assimilation techniques, then we may write this problem in the following form:

Problem 4.1 (The displacement assimilation problem) *Assume we are given a background state, $\boldsymbol{x}_b$, on the domain $\mathcal{D}$ and observations, $\boldsymbol{y} : \mathcal{D}_H \to \mathbb{R}^m$, taken on a domain $\mathcal{D}_H$, together with an associated observation operator, $\boldsymbol{h}[\boldsymbol{x}]$ which maps the state vector $\boldsymbol{x}$ onto a function $\boldsymbol{h} : \mathcal{D}_h \to \mathbb{R}^m$. Define the quadratic cost functional*

$$\mathcal{J}(\boldsymbol{x}; \boldsymbol{y}, \boldsymbol{h}) = \int_{\mathcal{D}_H} \int_{\mathcal{D}_H} (\boldsymbol{y} - \boldsymbol{h}[\boldsymbol{x}])_{\boldsymbol{r}'} w(\boldsymbol{r}', \boldsymbol{r})(\boldsymbol{y} - \boldsymbol{h}[\boldsymbol{x}])_{\boldsymbol{r}} dV \ dV', \qquad (4.2)$$

given by weights, $w : \mathcal{D}_H \times \mathcal{D}_H \to \mathbb{R}$, and suppose that $\mathcal{J}$ defines a measure, statistical or otherwise, of fit of the state $\boldsymbol{x} \in X$ to the observed data. Find the state, $\boldsymbol{x}^*$, which minimizes $\mathcal{J}$ over all possible rearrangements of the background state, $\boldsymbol{x} \in \mathcal{R}(\boldsymbol{x}_b)$.

When the domains or dimensions of the background state and observations are different, $\mathcal{D}_H \neq \mathcal{D}$, $n \neq m$, or the observation operator is not an isomorphic map then $\boldsymbol{h}^{-1}[\mathcal{R}(\boldsymbol{y})]$ and $\mathcal{R}(\boldsymbol{x}_b)$ are not over conformable spaces and the polar factorization technique described in section 3.1.2 must fail. Indeed, in the latter case only weak forms of $\boldsymbol{h}^{-1}[\mathcal{R}(\boldsymbol{y})]$ may exist, as in the case of the pseudo-inverse of non-invertible matrices. Failure here means the inability to generate a unique solution, $\boldsymbol{x} = g(\boldsymbol{x}_b)$, even in the limited sense of the polar factorization solution. There will indeed exist an $\boldsymbol{x}^*$ which minimizes $\mathcal{J}(\boldsymbol{x}; \boldsymbol{y}, \boldsymbol{h})$ over $\mathcal{R}(\boldsymbol{x}_b)$ but, unlike in the case of full observations, there may exist many other $\boldsymbol{x}$ which also achieve the infimum of $\mathcal{J}$, but which are not identical under an $\mathcal{L}_2$ norm, that is which fail to satisfy $\mathcal{J}(\boldsymbol{x}_1^*; \boldsymbol{x}_2^*, \boldsymbol{h}) = 0$. We have lost any sensible definition of uniqueness and the problem is now severely ill-posed.

Rather than regarding the entire question as insoluble and admitting defeat, we instead follow one of the standard techniques in optimization theory as an alternative to the attempt at direct solution of the minimization problem followed by polar factorization; develop an algorithm which generates successive vector fields which are elements in (some approximation to) the set, $\mathcal{R}(\boldsymbol{x}_b)$, which form a sequence which iteratively minimizes the cost functional $\mathcal{J}$. This is the same as the philosophy behind such commonly used techniques in numerical optimization as

the steepest descent method or the conjugate gradients method (Gill et al., 1981). Variations on these two techniques are frequently chosen when assimilating to minimize functions on $\mathbb{R}^n$, as for example in discrete implementations of the 3D-VAR and 4D-VAR data assimilation algorithms.

In general in terms of practical calculation it is neither feasible to generate the full set $\mathcal{R}(\boldsymbol{x}_b)$, nor to calculate $\mathcal{J}$ over it. Nor can we be certain of spanning it to any great degree by any kind of low dimension trial and error, Monte Carlo approach. Instead, by analogy with the steepest descent method, we seek some definition of a search path in $\mathcal{R}(\boldsymbol{x}_b)$ and use this to generate the sequence in $\mathcal{R}(\mathbf{x_b})$ which minimizes $\mathcal{J}$.

Considering the previous implementations of (quasi-)adiabatic rearrangement techniques of Section 1.6 we see that the methods discussed all operate in this fashion. That is, they find an analysis field which achieves a good fit to observations over some relaxed definition of a rearrangement, be it the box shifts (implemented as compressible velocity fields, but viewed as displacements of Brewster) the spectral displacements of Hoffman, or the limit to purely vertical shifts and open bottom boundary of the CH method. One advantage common to this approach and thus possessed by all the above techniques is that it disallows from consideration states generated by unphysical mappings, such as those consisting of a finite number of interchanges of pairs of points in space.

While strictly such bijective mappings do generate adiabatic rearrangements, since they satisfy the rearrangement condition (4.1), they are not one of the forms we are interested in for multivariate data assimilation, since they do not represent physical processes, but instead are mathematical artifacts of the spaces over which we are working. Indeed the existence of such mappings may disappear under aliasing when the problem is discretized for numerical solution.

4.1.2   A QUASI-STEEPEST DESCENT DISPLACEMENT METHOD

To exclude the particle exchange rearrangements discussed above, by physical argument, from this section onward interest is limited to spatially differentiable measure preserving mappings from $\mathcal{D}$ to $\mathcal{D}$ which also have continuous inverses, that is diffeomorphisms of the domain $\mathcal{D}$. Denote the set of these mappings by $G^\dagger(\mathcal{D})$ and suppose there exists some unique $g^\dagger \in G^\dagger$ which generates the state $\boldsymbol{x}^\dagger(\boldsymbol{r}) = \boldsymbol{x}_b(g^\dagger(\boldsymbol{r}))$ which minimises J over all the states generated by the rearrangement

$$\boldsymbol{x}(\boldsymbol{r}) = \boldsymbol{x}(g(\boldsymbol{r})),\ g \in G^\dagger.$$

The concept of path between $\boldsymbol{x}_b$ and $\boldsymbol{x}^\dagger$ may be made concrete by an extension of the mapping function, g into an additional, time-like, dimension. Let g_s be the extended, pseudo-time dependant mapping , $g_s(\boldsymbol{r}, s) : \mathcal{D} \times \mathbb{R} \to \mathcal{D}$, which satisfies the boundary conditions,

$$g_s(\boldsymbol{r}, 0) = \boldsymbol{r},$$

$$g_s(\boldsymbol{r}, 1) = g^\dagger(\boldsymbol{r}),$$

with g_s differentiable in s and $g_s \in G^\dagger$ for all s. This is a smooth curve through the set of diffeomorphisms on $\mathcal{D}$, with endpoints lying respectively on the identity mapping and the generator for the infimum of J. Provided such a curve g_s exists, then it is possible to prove the existence of a couple of conditions on the form that it must take.

Theorem 4.1 (Spatial Boundary Condition) *Points in space which are mapped to the boundary of $\mathcal{D}$, $\delta\mathcal{D}$ under the action of g_s for some value of s must initially lie on the boundary of $\mathcal{D}$ and hence must remain on the boundary for all s.*

Proof of Theorem 4.1 Consider a point $\boldsymbol{r}_1 \in \mathcal{D}$, such that $g_s(\boldsymbol{r}_1, s) \in \delta\mathcal{D}$, the boundary of $\mathcal{D}$, for some s. Suppose that $\boldsymbol{r}_1 \in \mathcal{D} \setminus \delta\mathcal{D}$. Then there exists an open ball, B, centred on $\boldsymbol{r}_1$, with $B \subset \mathcal{D}$, or otherwise $\boldsymbol{r}_1$ would have to lie on $\delta\mathcal{D}$. Now by continuity in s the image $g_s(B, s)$ of B under g_s at time s is an open subset of $\mathcal{D}$. But $\boldsymbol{r}_1 \in B$, so that $g(\boldsymbol{r}_1, s) \in g_s(B, s)$ and there exists an open subset of $\mathcal{D}$ containing $g(\boldsymbol{r}_1, s)$. Hence, $g(\boldsymbol{r}_1, s) \notin \delta\mathcal{D}$. This is a contradiction so therefore $\boldsymbol{r}_1 \in \delta\mathcal{D}$. Applying the same argument to the inverse mapping, g_s^{-1}, shows that all points initially in $\delta\mathcal{D}$ for $s = 0$ must remain in $\delta\mathcal{D}$ and thus proves the final statement of the theorem.

This constraint can be expressed as a condition on the rate of change of g_s at the boundary and $\boldsymbol{n}$, a unit normal to the boundary,

$$\left. \frac{\partial g_s}{\partial s} \right|_{\boldsymbol{r}} . \boldsymbol{n} = 0 \;\; \forall \boldsymbol{r} \in \delta\mathcal{D}.$$

Theorem 4.2 (The incompressibility condition) *Any invertible mapping, $g_s(\boldsymbol{r}, s)$, from a set $\Omega \in \mathbb{R}^n$ onto itself which is differentiable in both $\boldsymbol{r}$ and s and which has the properties that*

  (i). At $s = 0$ g_s is equivalent to the identity mapping,

$$g(\boldsymbol{r}, 0) = \boldsymbol{r}, \tag{4.3}$$

(ii). The mapping g_s is measure preserving on subsets of $\mathcal{D}$, for the chosen measure.

$$\int_V dV = \int_{g_s(V)} dV, \tag{4.4}$$

where V is an arbitrary volume in $\mathcal{D}$ and $g_s(V)$ is the image of that volume under g_s,

must satisfy the incompressibility condition,

$$\nabla \cdot \frac{\partial g_s}{\partial s} = 0 \;\; \forall \boldsymbol{r} \in \mathcal{D}.$$

Proof of Theorem 4.2

From property (4.3) we know that at $s = 0$, $g_s(V) = V$ and hence (4.4) is initially trivially satisfied. So to satisfy the condition for all s from differentiability we obtain the condition that

$$\frac{dV}{ds} := \frac{d}{ds} \int_{g_s(V)} dV = 0$$

for any arbitrary volume $V \subset \mathcal{D}$. But this can be evaluated explicitly using the volume form of Leibniz's Integral Rule to differentiate under the integral,

$$\frac{dV}{ds} = \int_{g(V)} \frac{\partial g_s}{\partial s} \cdot d\boldsymbol{S} = 0.$$

Applying the divergence theorem then gives the condition

$$\frac{dV}{ds} = \int_{g(V)} \nabla \cdot \frac{\partial g_s}{\partial s} dV = 0$$

but since we have assumed g_s invertible this must hold for any volume in $\mathcal{D}$ and hence the condition becomes as required

$$\nabla \cdot \left(\frac{\partial g_s}{\partial s} \right) = 0 \;\; \forall \boldsymbol{r} \in \mathcal{D}.$$

With these results in mind consider again how rearrangement theory was introduced in section 3.1.2. For a field $\rho(\boldsymbol{r}, s) : \mathcal{D} \times \mathbb{R} \rightarrow$ defined by the evolution equation

$$\frac{\partial \rho}{\partial s} + \boldsymbol{u} \cdot \nabla \rho = 0, \tag{4.5}$$

along with an initial condition,

$$\rho_s(\boldsymbol{r}, 0) = \rho_b(\boldsymbol{r}), \tag{4.6}$$

we have $\rho \sim \rho_b$, provided that

$$\nabla \cdot \boldsymbol{u} = 0,$$

$$\boldsymbol{u} \cdot \boldsymbol{n} = 0,$$

that is provided the flow is incompressible and has no flow through material boundaries. The obvious connexion between the viewpoints of mappings g_s and velocity fields $\boldsymbol{u}$ is the definition that

$$\frac{\partial g_s}{\partial s} = \boldsymbol{u}(g_s(\boldsymbol{r}), s),$$

the statement linking the Lagrangian and Eulerian definitions of velocity.

Following these ideas we define a new equivalence relation, similar to the rearrangement condition, relating scalar fields over $\mathcal{D}$. Let us say that ρ_a evolves from ρ_b if there exists a velocity field, $\boldsymbol{u}(\boldsymbol{r}, [0, 1])$, and a scalar field, $\rho(\boldsymbol{r}, [0, T])$, defined over the finite interval $[0, T]$ in the time-like variable s, such that ρ satisfies equations (4.5–4.6) and the boundary condition in s,

$$\rho(\boldsymbol{r}, T) = \rho_a(\boldsymbol{r}). \tag{4.7}$$

Like the rearrangement condition this is again an equivalence relation. Moreover, the equivalence class of that contains the field ρ_b (the set of fields that evolve from ρ_b) is denoted by the notation $\mathcal{W}(\rho_b)$ is, from the results of this section, a subset of the set of rearrangements of ρ_b. From Proposition 2 of Section 3.1.2 we know too that if states which are equal almost everywhere in a Lebesgue measure sense are considered the same, then the set of rearrangements is a subset of the states achievable from evolution and in this definition of path connectedness the two classes are the same. For any ρ_1 which can evolve from ρ_b it is now possible to define a path $\mathcal{P}$ from ρ_b to ρ_1, lying entirely within the set of rearrangements, with the path defined in terms of the various $\rho(\boldsymbol{r})$ generated by any of the possible velocity fields $\boldsymbol{u}(\boldsymbol{r}, s)$ which allow an advected field to attain ρ_1 in finite time.

Since the time-like variable s is arbitrary it may be rescaled at will without affecting the path $\mathcal{P}$, provided only that T and $\boldsymbol{u}$ are rescaled as well. That is, given that $\boldsymbol{u}(\boldsymbol{r}, s)$ generates $\mathcal{P}$ over the interval $[0, T]$ then so does $\boldsymbol{u}'(\boldsymbol{r}, s)$ over the

interval $[0, T']$ where $\boldsymbol{u}', T'$ are defined by

$$\boldsymbol{u}'(\boldsymbol{r}, s') = \frac{\boldsymbol{u}(\boldsymbol{r}, \lambda[s])}{\lambda'}, \quad T' = \lambda T,$$

and $s' = \lambda[s]$ is a new time-like variable. Given this property the definition of path length may be made less ambiguous, though still not unique, by an additional constraint on $\boldsymbol{u}$, namely that

$$\|\boldsymbol{u}\| = 1, \tag{4.8}$$

for all $s \in [0, T]$ and for some chosen norm, $\| \cdot \|$ on vector fields over $\mathcal{D}$. This condition fixes the time scale over which the path is traversed (effectively giving a parametrization of the path length in terms of T). We may thus discount the possibility of paths which traverse the same point more than once in the same direction, by always choosing the shortest path out of all those which traverse the same points.

With a definition of a path over rearrangements it becomes possible to define a methodology for a minimization technique, exactly analogous with the steepest descent technique in regular space. Let $\mathcal{J}(\rho, \rho_o, h)$ be a cost functional on ρ, defined as in equation (4.2). We aim to find a 'local' minimum of $\mathcal{J}(\boldsymbol{x})$ over $\mathcal{R}(\rho_b)$ from the starting state defined by ρ_b. The general algorithm to be used over successive values of the time-like variable s is:

Algorithm 4.1 (Displacement Steepest Descent Method)

1. *Given a current estimate $\rho(s)$ for the field ρ_a minimizing $\mathcal{J}$ over the set of all states which evolve from ρ_b, find the $\boldsymbol{u}^*$ (i.e. the search direction in $\mathcal{R}$) which produces the maximum rate of change in $\mathcal{J}(\rho)$ with respect to s, given that ρ satisfies equation (4.5) and that $\boldsymbol{u}^*$ must satisfy the normalization condition (4.8).*

2. *Set $\boldsymbol{u}(\boldsymbol{r}, s) = -\boldsymbol{u}^*(\boldsymbol{r})$.*

3. *Integrate equation (4.5) forward in time to generate a new test state ρ for the succeeding time level.*

4. *Check for convergence.*

The analogy with the steepest descent method follows from considering the behaviour of a cost function $\mathcal{J}(x)$ under an evolution equation,

$$\frac{\partial x}{\partial s} = d,$$

with d an arbitrary function on $\mathcal{D}$ to be determined. The rate of change of $\mathcal{J}$ with time s is thus given by the equation

$$\frac{d\mathcal{J}}{ds} = \left\langle \nabla_x \mathcal{J}, \frac{\partial x}{\partial s} \right\rangle = \langle \nabla_x \mathcal{J}, d \rangle,$$

which for finite d is minimized when

$$d = -\nabla_x \mathcal{J},$$

that is when the descent direction is parallel to the gradient of $\mathcal{J}$ with respect to s, the steepest descent path.

4.1.3   THE TWO DIMENSIONAL APPROACH

In preparation for the use of synthetic data from the quasi-geostrophic ocean model introduced in Chapter 3 the descent algorithm is here derived in full in the case of two-dimensional observational data and background fields and assuming a streamfunction formulation for the assimilation velocity $\boldsymbol{u}_a$. Let $\mathcal{J}$ be explicitly defined here as a continuous version of the standard quadratic cost function,

$$\mathcal{J}(q) = \frac{1}{2} \int_{\mathcal{D}} (q_o - q)^2 dV,$$

where full observations, q_o, of the conserved variable of potential vorticity, q, are assumed available everywhere and thus $\mathcal{D} \equiv \mathcal{D}_H$ is the (two-dimensional) domain of interest. Since in this context the assimilation velocity $\boldsymbol{u}_a$ is both two-dimensional and incompressible it may be represented by an 'assimilation streamfunction', ψ_a where

$$\boldsymbol{u}_a = \left(-\frac{\partial \psi_a}{\partial y}, \frac{\partial \psi_a}{\partial x} \right)^T.$$

The boundary condition, $\boldsymbol{u} \cdot \boldsymbol{n} = 0$ on $\delta\mathcal{D}$, may be expressed by setting ψ_a to a constant, $\psi_a = 0$ on $\delta\mathcal{D}$. The evolution equation for q expressed in terms of ψ_a is thus

$$\frac{\partial q}{\partial s} + \frac{\partial \psi_a}{\partial x} \frac{\partial q}{\partial y} - \frac{\partial \psi_a}{\partial y} \frac{\partial q}{\partial x} = 0, \tag{4.9}$$

and hence the rate of change of $\mathcal{J}$ with respect to s is given by

$$\mathcal{J}'(s) := \frac{d\mathcal{J}}{ds} = \int_{\mathcal{D}} (q_o - q) \left(\frac{\partial \psi_a}{\partial x} \frac{\partial q}{\partial y} - \frac{\partial \psi_a}{\partial y} \frac{\partial q}{\partial x} \right) dV. \tag{4.10}$$

Choosing a normalization condition of ψ_a,

$$\int_{\mathcal{D}} \|\mathcal{L}(\psi_a)\|^2 dV = 1 \tag{4.11}$$

for some function, $\mathcal{L}$, on scalar fields and applying a Lagrange multiplier approach and the usual variational principle to minimize the extended cost function,

$$\hat{\mathcal{J}}' = \mathcal{J}' + \lambda \left[1 - \int_{\mathcal{D}} \|\mathcal{L}\psi_a\|^2 dV \right],$$

using the boundary condition on ψ_a we see $\mathcal{J}'$ is extremized over permissible ψ_a when

$$\int_{\mathcal{D}} \delta\psi_a \left(\frac{\partial(q_0 - q)}{\partial x} \frac{\partial q}{\partial y} - \frac{\partial(q_0 - q)}{\partial y} \frac{\partial q}{\partial x} + \lambda \mathcal{L}^\dagger \mathcal{L} \psi_a \right) dV = 0$$

where we have defined the adjoint function $\mathcal{L}^\dagger$ given by

$$\langle \alpha, \mathcal{L}\gamma \rangle = \int_{\mathcal{D}} \alpha \mathcal{L}[\gamma] = \int_{\mathcal{D}} \gamma \mathcal{L}^\dagger[\alpha] = \langle \mathcal{L}^\dagger \alpha, \gamma \rangle.$$

For arbitrary $\delta\psi_a$, q this is true if

$$\mathcal{L}^\dagger \mathcal{L} \psi_a = \frac{1}{\lambda} \left(\frac{\partial q}{\partial x} \frac{\partial(q_0 - q)}{\partial y} - \frac{\partial q}{\partial y} \frac{\partial(q_0 - q)}{\partial x} \right), \tag{4.12}$$

where the Lagrange multiplier λ is given by equation (4.11) as

$$\lambda = \sqrt{\int_{\mathcal{D}} \left(\frac{\partial q}{\partial x} \frac{\partial(q_0 - q)}{\partial y} - \frac{\partial q}{\partial y} \frac{\partial(q_0 - q)}{\partial x} \right)^2}.$$

4.2   APPLICATIONS TO TOY PROBLEMS

4.2.1   NUMERICAL IMPLEMENTATION

Following the methodology of the identical twin problem as a test and verification method for data assimilation schemes the displacement descent technique was implemented numerically and applied to a series of toy problems, initially highly idealized and building to the use of data from the quasi-geostrophic ocean model.

For full details of the numerical methods used in implementing the displacement descent algorithm see section B.1 in the Appendix. Briefly, a spatially-centered finite difference approximation to the continuous evolution equation (4.9) was made and used to step the state field q forwards in the time-like variable s, under the action of the advection caused by the test assimilation streamfunction ψ_a. This time

integration step was performed using a fourth order accurate Runga-Kutta method. A fixed time-step was used and the descent of the finite difference approximation to the cost function, $\mathcal{J}$, tested at the end of each time integration step. When the method no longer descended another 'search velocity' was chosen by solving the finite difference approximation of the equation (4.12). In this sense the method is a true steepest descent algorithm.

The evolution equation (4.9) includes no explicit diffusion term. Nevertheless, both the spatial and temporal discretization methods used are known to be slightly dispersive. This effect, coupled with the extreme dependence of ψ_a on the fields q and q_o, plus the requirement of descent in the cost function $\mathcal{J}$ appeared to be sufficient to maintain numerical stability of the algorithm in most cases without explicit diffusion. Such an additional mixing term would indeed mean that over long integrations the method would no longer generate rearrangements of the background field in the continuous limit, as required by the initial adiabatic assumption. No explicit testing for issues such as CFL criteria were found necessary, once the choice of the initial time-step had been tuned to the magnitude of the various problems.

4.2.2   Relative Computational Cost

In order to quantify the relative efficiency of the new method we are introducing we will briefly discuss the expense of the method in terms of computational resources (primarily computer memory, either as RAM or disc space) and CPU time required by the algorithm, in comparison to standard variational assimilation techniques, namely 3D-VAR and an iterated approximation to a BLUE solution.

In terms of disc space, both the new displacement assimilation and the 3D-VAR techniques rely on the explicit formulation of a gradient term, either the assimilation velocity/streamfunction for the new displacement method, or $\nabla_x J$ for 3D-VAR. Since this is to be passed to a minimization routine, it must be store separately from the data on the current analysis state. For the 3D-VAR method this effectively doubles the storage requirements of the original model, since the gradient contains information on the update to every model variable, assuming that the assimilation is multivariate. For our new displacement algorithm we have the lesser cost of an additional model field, either vector (when generating assimilation velocities), or scalar (for cases using a streamfunction). Iterative methods typically have smaller storage requirements than either method, since the new analysis state is found without a minimization step and can overwrite the old one.

In terms of CPU time, the addition of the advection step, effectively requiring the running of a forward model for each calculation of the cost function in the mini-

mization, makes the displacement algorithm much more expensive than traditional 3D methods. Effectively, the cost of a single step becomes similar to that of a single iteration of a 4D-VAR method, even though data at only a single timestep is being assimilated. This additional modelling step will also impact on memory, since the overheads of the advection scheme must be stored, meaning that here both the 3D-VAR and displacement algorithms require space for twice the variables in the original model, with the displacement scheme having the additional over head of the assimilation advection fields, plus the various costs imposed by the numerical advection mode. Note, however, that since the objective function can be tested at each stage there is no requirement to store additional intermediate steps in the advection routine.

4.2.3   ONE-DIMENSIONAL PROBLEMS

We now consider the behaviour of the advective minimization approach on a hierarchy of problems, in an attempt to understand its peculiar advantages and limitations. We commence with a brief discussion of displacement assimilation in a one-dimensional domain. In such domains the only velocity field satisfying the relevant version of the incompressibility condition,

$$\frac{du}{dx} = 0,$$

is one with u constant everywhere. If the domain is also bounded then the no flow condition requires that $u = 0$ everywhere. For our minimization scheme this means the assimilation velocity is static and the method cannot converge. This links back to the lack of a connectivity condition for the rearrangement problem in one dimension. In a periodic domain then u is no longer constrained by boundaries, while under the rescaling argument of the previous section only the sign of the constant u_a is important, determining whether the direction of the shift is positive or negative. Given an accurate line search method the algorithm will thus converge to a local minimum of $\mathcal{J}(\boldsymbol{x})$, with only one calculation of a search direction required. That the algorithm may not find a global minimum and may not choose the direction of shortest path length is obvious from consideration of the behaviour of the functions of Figure 4.1, in which the algorithm is applied for a 1-D function on a periodic domain in six different cases:

(a). Full, perfect observations, with the correct form for the background (error takes the form of a periodic shift in position, the rearrangement condition with respect to the true function is satisfied).
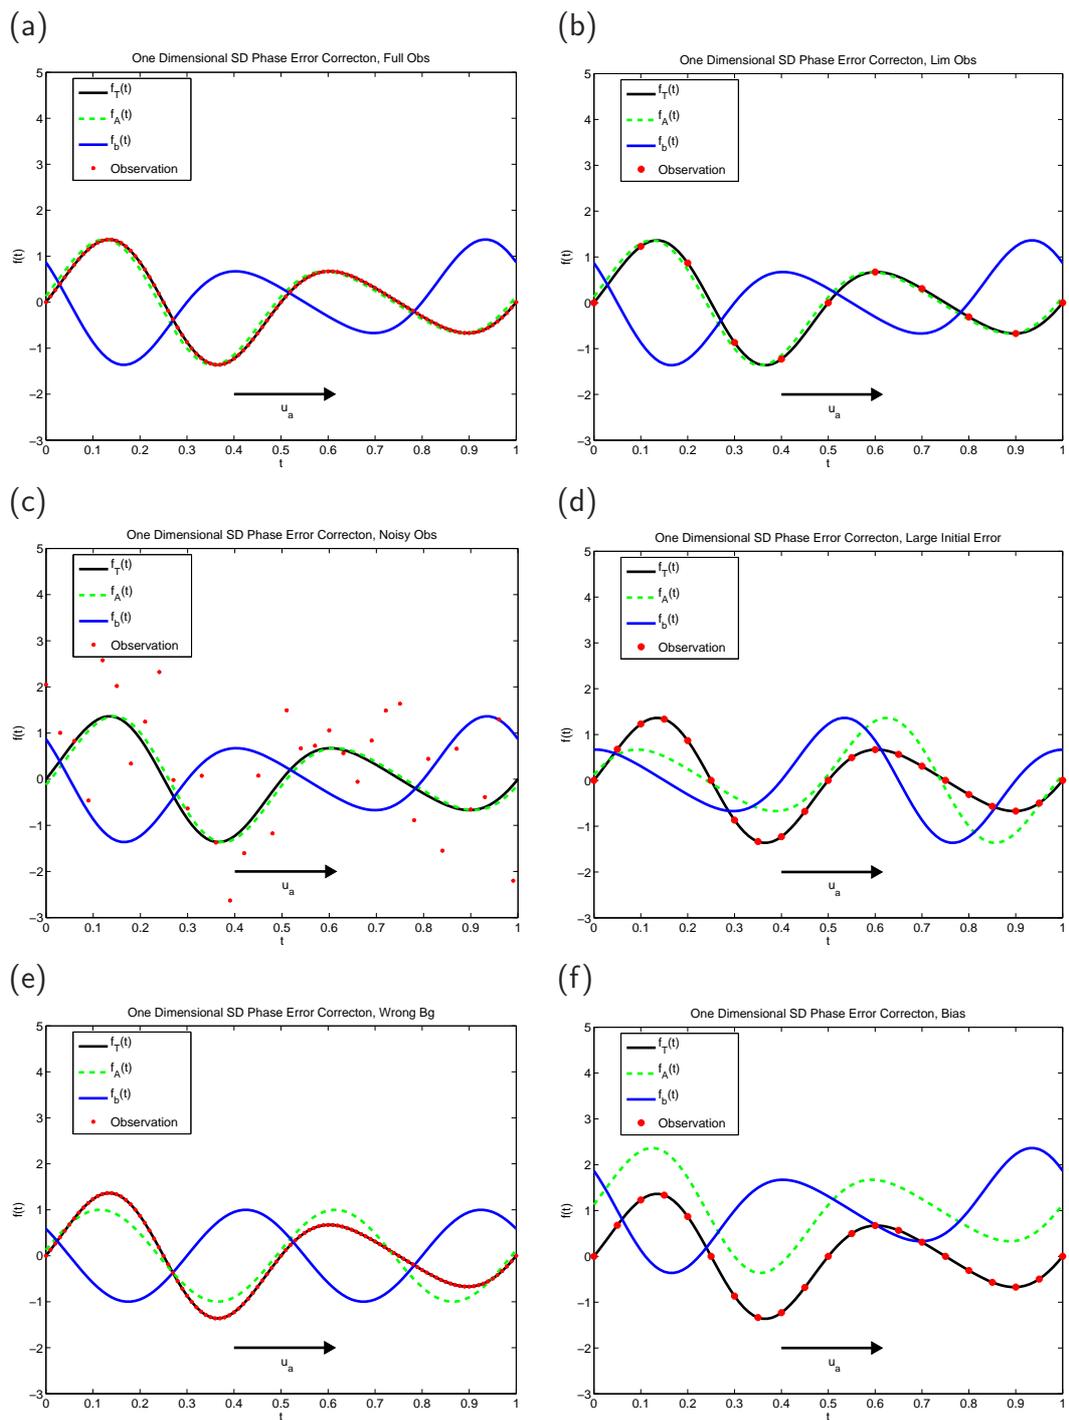
Figure 4.1: The steepest descent algorithm in a 1D periodic domain

In each case a simple iterative algorithm minimizes (green) the square of the residual for observations, y (red), from a background function, f_b (blue). Also plotted is the direction of the assimilation velocity, u_a. (a) Perfect observations everywhere. (b) Limited, perfect observations. (c) Limited, noisy observations. (d) Limited, observations, large initial error, correct. (e) Full observations, incorrect background. (f) Limited perfect observations, fixed bias in background.

(b). Very limited, perfect observations (the number and range of oscillations is sufficient to track the function turning points), correct form for background.

(c). Limited, noisy observations (unbiased Gaussian noise on all observations), correct form for background.

(d). Limited, perfect observations, large initial displacement error, correct form for the background.

(e). Full, perfect observations, incorrect form for background (error includes both phase and amplitude components, the rearrangement condition with respect to the true shape is not satisfied).

(f). Limited perfect observations, fixed bias in background (there is a single valued amplitude error across the whole domain,the background state satisfies a modified form of the rearrangement condition on the true state).

Among these examples only in the case of the large initial error, (d), did the scheme fail to converge to the global minimum of the cost function. In this case the scheme instead found a local minimum with respect to the observations, found by initially minimizing in the opposite direction to that which would have provided the correct analysis solution. This possibility of finding a local minimum is true for any iterative scheme.

Perhaps the most impressive case is that of plot (c), in which the true state is found to with-in a one grid-point lag, even though the form of the solution is not at all obvious from the noisy observations. This case is another example of a problem for which displacement techniques are ideally suited. The final two problems, (e) in particular, show examples of success in situations less appropriate for displacement techniques. In both cases the phase error is corrected for in the final solution, but the initial amplitude error is retained, as is clear from plot (f) where the turning points of the analysis and true functions match one-to-one but are separated by the constant bias. Note that, as was discussed in Section 3.1.2, the one dimensional problem is an extremely special case for the rearrangement technique.

4.2.4   A TWO-DIMENSIONAL PROBLEM FROM AN ALGEBRAIC INITIAL CONDITION

Once again, as in Chapter 3, more complex behaviour develops when considering assimilation problems in domains of two or more dimensions. Consider two

different 'density' fields, ρ^b, ρ^o defined by

$$\rho^b = f(x, z),$$

$$\rho^o = f(x + \theta, z),$$

in the unit square $\mathcal{D} = [0,1] \times [0,1]$ and with f the semi-periodic function

$$f(x, z) = \sin(2\pi x) + kz.$$

This choice is such that the continuous background and observation fields are adiabatic rearrangements of each other[2], $\rho^b \sim \rho^o$. Moreover, for $\theta = k/\Delta x$, the same property holds when discretized, with the quantity inside the volume integral of the rearrangement condition replaced by any quadrature function dependent only on the nearest nodal values. Assuming a nearest-node value approach henceforth the discretized cost function, $\mathcal{J}(\boldsymbol{\rho})$, may be defined as

$$\mathcal{J}(\boldsymbol{\rho}) = \frac{1}{2} \sum_{i \in [1, n_x]} \left[\sum_{j \in [1, n_z]} \left[(\rho_{[n_z \times (i-1)+j]} - \rho^o_{[n_z \times (i-1)+j]})^2 \Delta z_j \right] \Delta x_i \right], \approx \frac{1}{2} \int_{\mathcal{D}} (\rho - \rho^o)^2 dS,$$

where $\boldsymbol{\rho}$ is a vector of nodal values $\rho((i + \frac{1}{2})\Delta x, (j + \frac{1}{2})\Delta z)$ under natural ordering. Here the $\Delta x, \Delta z$ denote the discretization step lengths and the n_x, n_z denote the total number of nodal points in the x and z directions respectively.

From the fact that the discretized rearrangement condition is satisfied it is clear that $\mathcal{J}$ has a unique global minimum of $\mathcal{J} = 0$ over all discretizations of adiabatic rearrangements. This occurs in the simple case of $\boldsymbol{\rho} = \boldsymbol{\rho}^o$. If a monotonically decreasing search path exists in the discretization of $\mathcal{R}(\rho_b)$ which contains the element ρ^o then the algorithm should converge to this state, up to the limits of machine numerical precision and discretization truncation error; otherwise a more local limit will be found.

Running the two-dimensional steepest descent algorithm on this algebraic problem with an $\mathcal{L}_2$ norm on the magnitude of the ψ_a,

$$\int_{\mathcal{D}} \psi_a^2 dS = 1,$$

the scheme was experimentally found to converge, slowly, to a higher limit. The

---

[2]This would of course be true for *any* semi-periodic or periodic function on $\mathcal{D}$. The form given here is chosen as an idealized representation of a problem in which an anomaly in mixed layer depth has become displaced in a model versus observations and must be corrected through assimilation. This is a natural problem, similar to the one for which the CH method was devised.
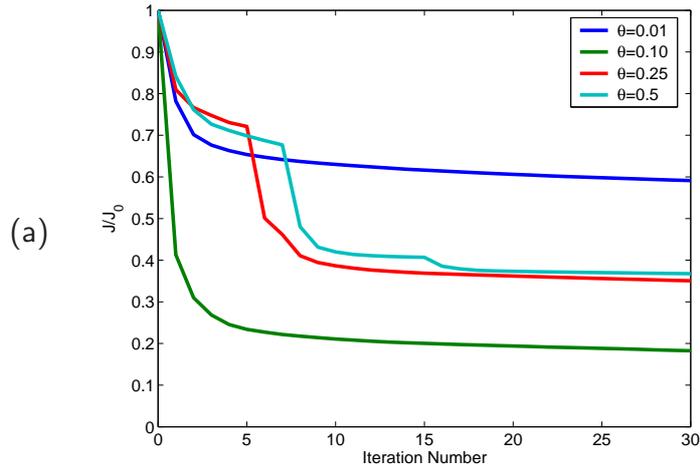
Figure 4.2: Convergence of $\mathcal{J}/\mathcal{J}_0$ for variable θ, $k = 3$, $n_x = 64$, $n_z = 64$
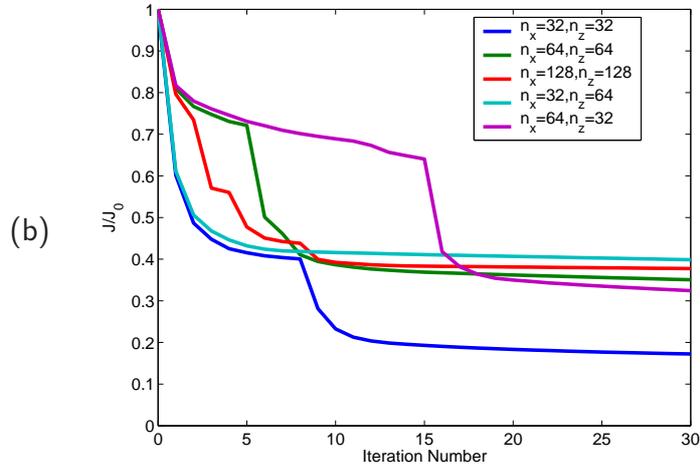
Figure 4.3: Convergence of $\mathcal{J}/\mathcal{J}_0$ for variable n_x, n_z, $k = 3$, $\theta = 0.25$
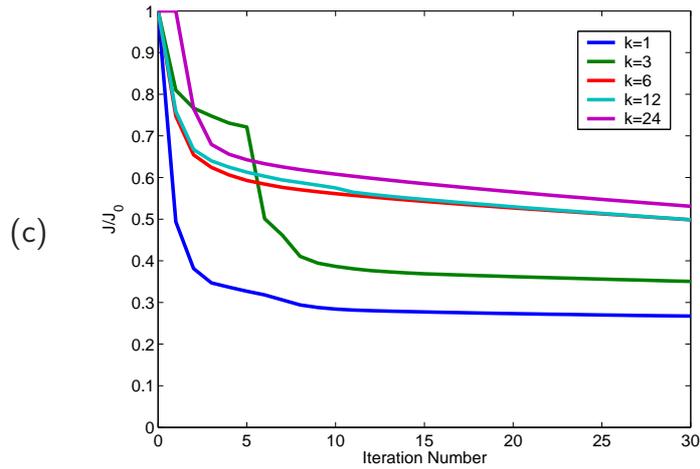
Figure 4.4: Convergence of $\mathcal{J}/\mathcal{J}_0$ for variable k, $n_x = 64$, $n_z = 64$, $\theta = 0.25$

The plots above show the variation in $\mathcal{J}/\mathcal{J}_0 := \int_{\mathcal{D}}(\rho - \rho_o)^2 dV / \int_{\mathcal{D}}(\rho_b - \rho_o)^2 dV$ against iteration of the descent routine with changes in (a) initial phase error, (b) grid step length, (c) density stratification.

precise asymptotic value of this limit and rate of convergence towards it is dependent on the values of the parameters θ, n_x, n_z and k.

As seen in Figure 4.2, the θ dependence expressed itself as a slowing in the rate of convergence and an increase in the value to which the scheme converged as θ was increased over the range $[0, \frac{1}{2}]$. By symmetry values in the range $(\frac{1}{2}, 1)$ above this again reduced the limit and increased the rate of convergence, since they were equivalent to a problem in the range $[0, \frac{1}{2}]$ with the x-axis reversed. This behaviour is easily explained by noting that the value of the cost function at the beginning of the integration, $J(\rho_b)$ is predominantly determined by the θ value used. The problem becomes 'harder', in the sense of containing a greater phase error, which in turn leads to a longer path, a longer integration and thus more accumulation of errors from the numerical schemes used for the advection term. For some values of θ, notably $\theta = 0.5$, the curve for the convergence of $\mathcal{J}/\mathcal{J}_0$ appears to 'step' down towards its final value, with periods of fast convergence, followed by long plateaux where the scheme nearly stalls. It appears these are a function of the poorly conditioned descent problem, where in some sense the method tracks narrow valleys or passes near local minima, before finding steep, straight paths where long steps may be taken. This behaviour is repeated for variations in n_x, n_z and k.

The general form of the dependence upon n_x and n_z was that as the number of points over which the problem was calculated increased the rate of convergence decreased and the limiting value increased. This is plotted for a variety of values in Figure 4.3. This behaviour can easily be explained as an effect of the change upon the CFL limit again causing a longer integration. However this was not symmetric in the two axes, with an increase in n_x (the direction containing the phase error) having more effect than an increase in n_z. This suggests there are genuine effects on the minimization routine due to the difficulty in propagating information across multiple grid points by advection.

The effect of changing k, as shown in Figure 4.4 was initially difficult to explain. Convergence seemed to degrade as k grew, both in the reduction achieved with successive iterations of the routine and the final value which the method appears to attain. This suggested that some process other than the simple transport of fluid in the direction of the phase error was taking place. To investigate the processes the fluid underwent in greater detail two further systems were coupled to the velocity fields generated by the data assimilation scheme.

Specializing to the case $\theta = 0.1$, $n_x = n_z = 64$, $k = 3$ we present the initial discrepancy between the background and observation fields , $\rho^b - \rho^o$ as plot (1) in Figure 4.5, with plot (2) showing the final anomaly between the analysis field after convergence and the observations, $\rho^a - \rho^o$. Note that, while the measure of the total

discrepancy has been shown in Figure 4.2 to reduce to roughly 0.2 of its original value, the local discrepancy in many places around the boundary has increased. It is in these areas that any further advection leads to increases in the cost function and the method stalls. As a test of the magnitude of this effect we introduce a new metric, $\mathcal{J}[\text{box}]$, which calculates the usual quadratic measure of discrepancy, but over a limited rectangular domain with a centre coincident with that of the full domain,

$$\mathcal{J}[\text{box}] = \frac{1}{2} \sum_{i \in [1+\ell_x, n_x-\ell_X]} \left[\sum_{j \in [1+\ell_z, n_z-\ell_z]} (\rho - \rho^o)^2 \Delta z_j \Delta x_i \right],$$

where ℓ_x, ℓ_z denote the 'margins' left unexamined under the metric.

Using this new metric we can investigate two different effects, the proportion of the discrepancy between analysis and observation which lies in a box of varying size, and the extent to which the algorithm has successfully matched the contents of the box, neglecting any data outside it. This gives two new functions, $\mathcal{J}[\text{box}](\rho^a)/\mathcal{J}(\rho^a)$ and $\mathcal{J}[\text{box}](\rho^a)/\mathcal{J}[\text{box}](\rho^b)$, which we plot against the area of the box as plots (3)–(4) in Figure 4.5.

Considering $\mathcal{J}[\text{box}](\rho^a)/\mathcal{J}(\rho^a)$ we see that the proportion of the discrepancy in the box does not decrease linearly with area, but instead initially decreases rapidly, dropping to half its value with a box of 0.8 the original area. This is restating the point above that it is the boundary layers which limit the assimilation method. The graph can in fact be split into three roughly linear sections. For boxes of area 0.8–1.0 the change is fast than unit, since it is over this range we are ignoring the boundary layer. For boxes of area 0.1–0.8 the change is slower, since this is integrating over the area where the assimilation method has been most successful. As the box area drops below 0.1 then numerical artifacts of the integration and assimilation methods dominate and the rate of change again approaches unity. For a square domain with $n_x = n_z = 64$ an area of 0.8 corresponds to a boundary layer of roughly 3 grid points on each boundary. Looking now at $\mathcal{J}[\text{box}](\rho^a)/\mathcal{J}[\text{box}](\rho^b)$ we see that, while over the whole box we have achieved an 80% reduction in cost function, the centre of the box has obtained a better fit, reducing to 0.14 of its original value. While this is still significantly greater than unity, it is still a significant improvement, when viewed as an extra 30% reduction in the cost function.

Studying the form of the final analysis discrepancy we note the obvious grid-scale striations introduced along the vertical boundaries by the numerical advection routine. In some cases these are are locally very large, probably as dominant as the more physical anomalies along the x-axis. To test the role of numerics in this problem we performed a series of tests in the regime $n_x = n_z = 64$, $k = 3$

but allowing θ to tend to zero. The result in terms of reduction in cost function is shown in plot (1) of Figure 4.6, with a log scale for θ. The U shaped curve towards the extreme left of the plot, with oscillations for small θ (the right hand side of the graph) is typical of a minimization problem in the presence of numerical error. Initially as the problem becomes simpler the method improves, but then errors due to the numerics dominate and the code begins to stagnate rapidly, giving the 0.2 reduction in cost function achieved at the end of the window. To illustrate that these are numerical effects, we apply the new rearrangement metric, $\mathcal{M}$, introduced in Section 3.2. We calculate $\mathcal{M}/\mathcal{J}$ for the pairs of fields $[\rho^a, \rho^b]$ and $[\rho^a, \rho^o]$, shown in plots (2)–(3) of Figure 4.6. In the continuum limit with perfect advection these functions should both be zero for all θ, since all functions are rearrangements of each other. Instead, both functions grow rapidly with θ as the noise from the advection routine, shown by the computational mode in the discrepancy field, begins to dominate the minimization problem.

In principal the advection method used could be chosen from those, eg. contour advection, which are strictly conservative in the sense of the rearrangement condition. However this would add additional overhead in reverting to a scalar field form to obtain a value for the cost function, and hence we view this as too expensive. In the more physically realistic assimilation experiments which follow boundary effects appear small, suggesting we lie near the minimum of the U shown in plot (1) of Figure 4.6.

4.2.5 UNOBSERVED TRACERS

The first system was an unobserved tracer field, advected using the same scheme as ρ, but not contributing to the value of ψ. In effect the scheme adds extra equations

$$\frac{\partial \tau_i}{\partial s} + \boldsymbol{u}_a \cdot \nabla \tau_i = 0 \tag{4.13}$$

to the assimilation evolution equations. These tracer equations are discretized and solved numerically using the same approach as in the passive tracer for the QG model of Section 2.3.2. Experiments with this system show several of the advantages of the adiabatic data assimilation philosophy when compared with standard assimilation techniques. There is no spreading of fronts, and no creation of double maxima, even when assimilating data with large gradients. Moreover the scheme is now implicitly, rather than explicitly, multivariate. Rather than requiring statistics on the correlation of variations in τ and ρ, the scheme inverts an assumed $\tau(\rho)$ relation, in the same manner as the salinity assimilation methods of Section 1.5.3.

Figure 4.5: Role of boundaries in the 2D algebraic problem

(1) shows the initial discrepancy between background and observation for the case $n_x = n_z = 64$, $\theta = 0.1$, $k = 3$. (2) Final discrepancy between analysis and observations after application of the displacement assimilation technique. (3) Proportion of total costfunction contained in a square coincident with the full domain and of the indicated area. (4) Relative reduction in costfunction for a square coincident with the full domain and of the indicated area.

Figure 4.6: Numerical error and the 2D algebraic problem

(1) Analysis reduction in cost function for the problem $n_x = n_z = 64$, $k = 3$ against choice of θ. (2) Proportional value of the displacement error metric $\mathcal{M}$ over the discrepancy metric, $\mathcal{J}$ for the analysis state, ρ^a and background state, ρ^b. (3) Proportional value of the displacement error metric $\mathcal{M}$ over the discrepancy metric, $\mathcal{J}$ for the analysis state, ρ^a and observation state, ρ^o.

The effect of the scheme on the discrepancy in the τ field is dependent on how accurately the background field represents the evolution of the true $\tau(\rho)$ relation and how well the scheme can identify the correct destination of water masses from just the ρ field. Figure 4.7 shows the behaviour of two passive tracer fields under one instance of the two-dimensional algebraic problem of Section 4.2.4 (here $k = 3$, $\theta = 0.25$ so that the effective displacements are large). The first tracer, τ_1, has a vertical stratification, and thus a suitable approximate $\tau(\rho)$ relation. This tracer shows a physically sensible analysis field, in which the horizontal displacement between background and observation introduced by the phase shift is observed as a vertical shift in tracers as fluid is moved to balance the density differences. This is precisely the principle behind the CH algorithm of Section 3.3.1. The second tracer field, τ_2, has a horizontal stratification. The resulting analysis field shows the effect of the boundary return currents on the field, producing an image with strong gradients along the top and bottom surfaces. This is the type of effect the CH algorithm could not capture, but which could be argued as unphysical.

As a more general example, consider a fully observed field ρ_1 which contains two features consisting of two water masses V_1, V_2 of equal volume with

$$\int_{V_1} \delta(\rho_1 - \alpha)dV = \int_{V_2} \delta(\rho_2 - \alpha)dV \quad \forall \alpha \in \mathbb{R}.$$

Then the field ρ_2 which has V_1 mapped to V_2 (and vice versa) is an adiabatic rearrangement of ρ_1, in the sense of section 3.1.2. If we suppose there is a further tracer field, τ over the domain, with

$$\int_{V_1} \delta(\tau_1 - \alpha)dV \neq \int_{V_2} \delta(\tau_2 - \alpha)dV$$

for at least one α then the masses can be differentiated by their τ properties. The question is whether incompressibility of the assimilation velocity alone is sufficient additional constraint on the system to force the scheme to find the correct physical choice of rearrangement, mapping V_1 to V_1 and V_2 to V_2, rather than visa-versa. As will be shown in the following section by tracking the paths of individual fluid parcels, it may not always be enough.

4.2.6   TRACING THE PATH OF FLUID PARCEL REARRANGEMENT

To investigate the question posed in the previous section, and to better visualize the path of fluid elements under the scheme, a basic particle scheme was coupled to the velocity fields calculated in the assimilation method. For a particle P at a
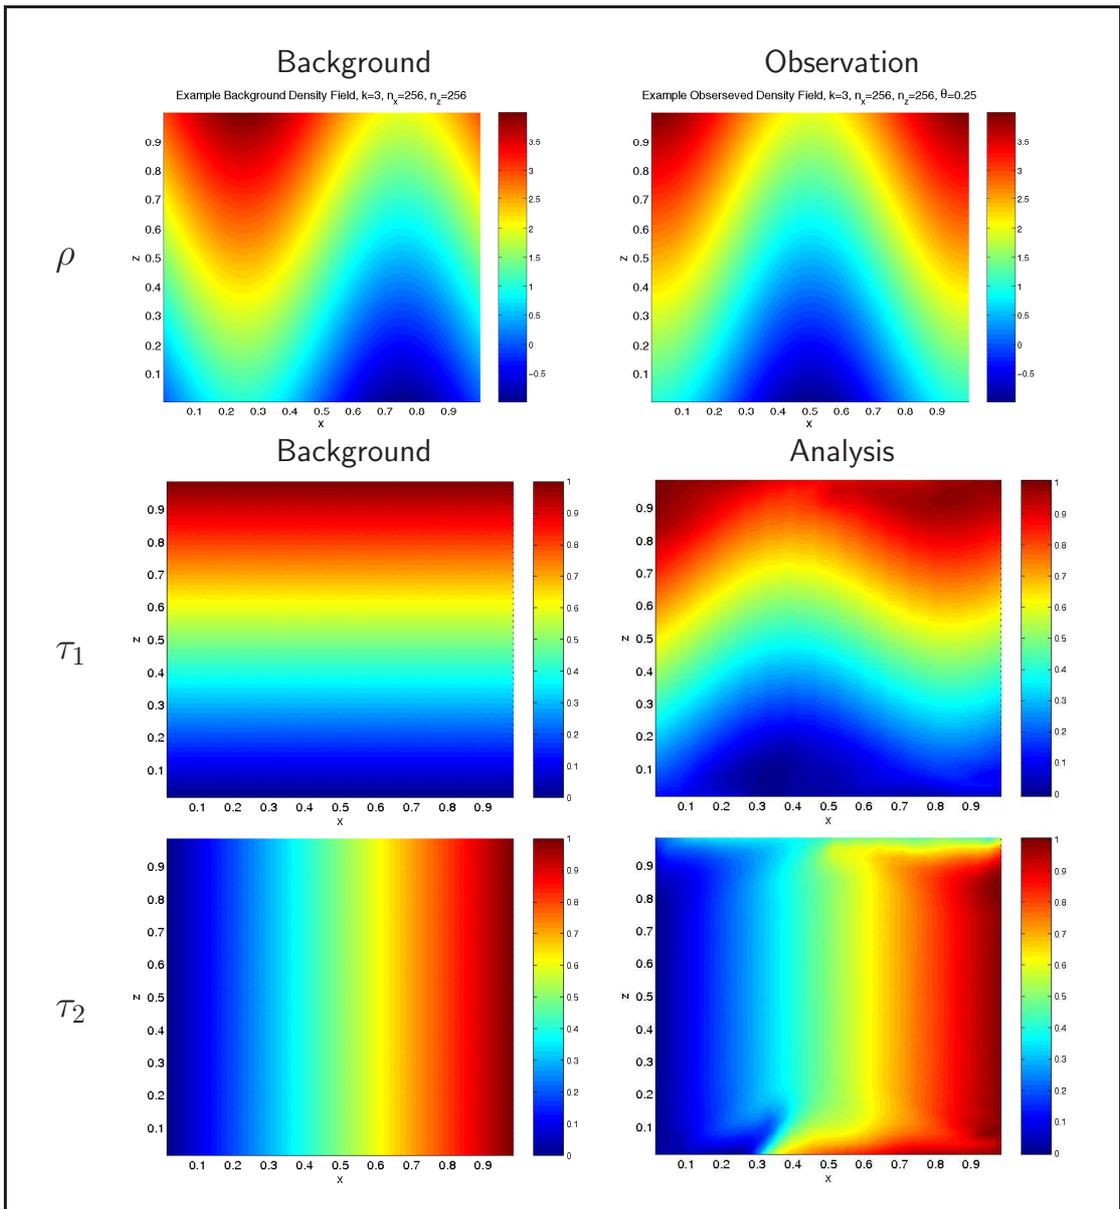
Figure 4.7: Background and Analysis fields for observed and unobserved tracer fields

The top line shows two fields provided as background and observations to the displacement assimilation routine. The other plots show the effect on two unobserved tracer fields, τ_1 with a gradient perpendicular to the direction of the required displacement, τ_2 with a gradient perpendicular to it, advected along with the assimilation . The first field shows a smooth evolution consistent with a dipole-like circulation configuration with lines of no vertical motion at $x = 0.2$ and $x = 0.6$. The second shows more obvious small scale and boundary effects, particularly in the bottom left and top right corners. This shows both the multivariate nature of the assimilation routine as well as providing a visualization of the assimilation process.

point r^k at time step k the local velocity field $u_a(r^k)$ was calculated by linearly interpolating the four nearest velocity points in each direction. The position of the particle is evolved using the simple scheme

$$r^{k+1} = r^k + u_a \Delta t,$$

so that in the continuum limit a particle initially within an arbitrary fluid parcel is transported with it and remains within it for all time. This allows the path of fluid parcels under the assimilation to be tracked. Seeding the system with a regular grid of particles and assimilating one of the 2 dimensional algebraic problems ($n_x = n_z = 32$, $k = 10$, $\theta = 0.25$) produces he particle (and hence fluid parcel) paths shown in Figure 4.8.

This appears to show that the particles follow generally vertical paths, i.e. that the scheme operates mainly by generating overturning circulations similar to the method of Cooper and Haines. This has an obvious connection with the link between the rate of convergence of the method and the k value chosen, since for k large a relatively small (and thus easy to generate) vertical displacement gives a large change in ρ. This means that small vertical displacements fit the model isopycnal layers to the observations well, significantly reducing the cost function, hence giving fast convergence. It does however still leave the scheme open to the potential 'identification' errors as posed in section 4.2.5. This issue can be examined by introducing into the background and observation fields an identifiable 'bulge' of fluid, such that the isopycnal height of the fluid is not a constant in x as in the previous case, while still ensuring that the relationship $\rho_b \sim y$ holds. Since the fluid 'bulge' is a local phenomenon we would expect that the scheme would move particles which started within the bulge close to its final position. The simplest method to achieve this is using images from the QG model of Section 2.3.2, as shown in Figure 4.9.

The most obvious effect of in this more complicated case is the sharp increase in the complexity of particle paths, compared with those in the symmetric case. This is linked to a slower convergence rate, bearing in mind the originally higher value for the cost function due to the extra discrepancy caused by the bulge. Although the paths do show some movement of particles to follow the bulge the total extent of horizontal motion is limited. This is to be expected since most of the fluid to supply the bulge in the analysis will have started in the area the bulge was to be created in the background.

Figure 4.8: The particle system in action for $\theta = 0.25$, $k = 10$, $n_x = n_z = 32$

The main plot has the background state as the coloured plot. The observed and analysis fields are show above. Superimposed are the paths of the particles during the assimilation process. Crosses mark the initial seeding of the particles, with tails marking the subsequent path of the particle under the integration of the assimilation velocity up to its final position, marked by the open circles at the head.

Figure 4.9: The particle flow for a non-symmetric problem

As in Figure 4.8 the main plot shows the background field, with the observations and analysis above. Crosses mark initial particle positions and circles their final location. Note that for this problem, as is typical with the QG model differences between the background state and observations were concentrated around the jet and eddies in the centre left of the basin, exactly where particles have been transported furthest. The curved particle paths at the local the eddies show clearly the rotation the assimilation velocity applies. Note that most of the basin is stagnant under the assimilation however.

4.3   IMPROVING THE RATE OF CONVERGENCE

Although the method successfully reduces the cost function over the course of the iteration, thus implying data assimilation is being performed, and successfully generates an adiabatic rearrangement of the initial background field, the rate of reduction in $\mathcal{J}$ and the tendency to converge to local rather than global minima are both disappointing, compared with the typical behaviour of methods such as 3D-VAR when applied to traditional problems, which can achieve near total reduction in cost function measurements in only a few iterations. To generate an efficient method for adiabatic data assimilation a method must be found to speed convergence and reduce the likelihood of converging to local minima.

Traditionally the steepest descent algorithm is seen as the most basic and little used member of a family of gradient descent, or 'hill-climbing' optimisation techniques such as the family of conjugate gradient method and the Broyden-Fletcher-Glodfarb-Shannon method. Unfortunately these methods are not greatly suitable, due to the extreme non-linearity of the optimization problem in question and the difficulty of mapping to the spaces in question. Alternative, stochastic based methods such as simulated annealing are also inappropriate due to the large dimension of the problem being solved. Instead modifications to the original algorithm are considered, with the aim of finding the most efficient method.

4.3.1   CHOICE OF NORMALIZATION FACTOR

The condition in equation (4.11), with $\mathcal{L}$ the identity operator, represents a restriction on the $\mathcal{L}_2$ norm of the streamfunction. This imposes only a weak constraint on ψ_a while allowing point and line discontinuities which represent local failures in the incompressibility condition. These discontinuities can cause unacceptably large values for the resulting velocity field after the problem is discretized and thus impinge on the CFL limit for the assimilation and on the length in pseudo-time for which the derived search direction for the velocity field $\boldsymbol{u}_a$ actually represents a descent path for the cost function. An obvious alternative choice for $\mathcal{L}$ is to directly constrain the $\mathcal{L}_2$ norm on the assimilation velocity field itself,

$$\int_{\mathcal{D}} \boldsymbol{u}_a \cdot \boldsymbol{u}_a \, dV \equiv \int_{\mathcal{D}} \nabla \psi_a \cdot \nabla \psi_a \, dV = 1,$$

i.e to apply $\mathcal{L}\phi = \nabla\phi$. Under this modification the change to the optimality condition on $\partial J/\partial s$ is such that instead of obtaining

$$\psi_a = \frac{\partial((\rho^o - \rho), \rho)}{\partial(x, y)},$$

with the streamfunction equal to the Jacobian operator applied to the innovation vector and the background field we obtain an elliptic equation for ψ_a,

$$\nabla^2 \psi_a = -\frac{\partial((\rho^o - \rho), \rho)}{\partial(x, y)}.$$

This elliptic problem must be solved under the imposed no-flow boundary condition on ψ_a,

$$\psi_a = 0 \quad \text{on } \delta\Omega.$$

This choice of norm represents a bound on what may be called the 'kinetic energy' of the assimilation step at any pseudo-time level, which is perhaps the most natural choice for a normalization condition, although in principal any normalization condition of the form

$$\int_{\mathcal{D}} \sum_{n=1}^{\infty} m_{2n}(\nabla^{(2n)}\psi_a)^2 + m_{2n+1}(\nabla^{(2n+1)}\psi_a \cdot \nabla^{(2n+1)}\psi_a) \, dV,$$

for a set, m_i, of real, positive constants. In other words any Sobelov norm on ψ_a, may easily be applied.

Investigating the behaviour of the two norms showed that, as might be expected, the numerical method based on the elliptic problem generated significantly smoother streamfunction fields and thus in turn much smoother velocity fields. This was found to lead to significantly longer integration times in s before descent failed for any given search direction and thus a noticeably faster rate of convergence in $\mathcal{J}$. Running the algorithm with norms on ψ_a, $\nabla\psi_a$, $\nabla^2\psi_a$ and $\nabla^3\psi_a$ on the algebraic two dimensional problem produced the results shown in Figure 4.10. As predicted, the $\nabla\psi_a$ shows best convergence properties, although the $\nabla^2\psi$ norm (an enstrophy norm on the assimilation velocity field, $\boldsymbol{u}_a$) showed very similar behaviour. In comparison a Sobelov norm on both $\nabla\psi$ and ψ together, (labelled $\|\nabla\psi\| + \|\psi\|^2/L_D^2$, where $L_D = 30,000$m is the Rossby radius of deformation for the system) produced behaviour only slightly better than for a norm on ψ alone. In this instance the highest order norm tested failed to converge. This may be due to the necessity of specifying three sets of boundary conditions for this problem, since the method requires the integration by parts of the variational term derived

from the normalization function, as well as the usual boundary condition from the underlying transport process.

4.3.2   A MULTISCALING DESCENT ALGORITHM

The implementation of the descent algorithm represents a discretization of the continuous equations presented in Section 4.1.1, which are in turn applied to the discretized source data for observations and an initial background field. There is, however little coupling between the length scales used in the discretization of the source data and those required for numerical implementation of the algorithm itself, apart from the necessity for the final assimilation analysis to lie on the background model grid if it is to be used as an initial condition for further model runs. There is moreover absolutely no coupling between the time discretization used for the model chosen to provide the background field and the discretization chosen for the time-like variable contained in the descent algorithm. Hence although it is natural to choose the same grid structure for the descent algorithm as was used in the original model generating the background field and thus to use an identity mapping to interpolate one field onto the other, this is a choice made of convenience and not a technical requirement.

This suggests it is feasible to apply to the descent method a technique for speeding the convergence of the analysis field equivalent to the use of multigridding methods in the application of minimization techniques such as the conjugate gradient method (Elman, 1998). The underlying idea is to perform the calculation of ψ_a and the accompanying minimisation of $\mathcal{J}(\rho)$ on a coarse, large scale (and thus reduced dimension) grid. In the context of displacement methods this should capture the large-scale shifts required to move water masses into the same vicinity, thus reducing the cost function, without necessitating the small scale shifts needed to capture the fine-scale positioning of filamentary structures. Such methods should also lead to smooth, numerically stable velocity fields allowing longer time-integration steps without failing stability criteria such as the CFL condition.

Evidence of the potential of such a technique can be found by considering a low pass[3] operator, $\overline{f}$, on fields over $\mathcal{D}$, with complementary eddy fields, $f' = f - \overline{f}$. We

---

[3]The nomenclature of the low-pass filter has been inherited into climate science from the field of signal processing, unfortunately without any firm definition of what such a filter actually means, beyond the general assumption that fast, high frequency signals are attenuated, while low, slow frequency signals pass though. For a signal, $f(t)$ and a scale τ, one such operator is

$$\overline{f}(t) = \frac{1}{\tau} \int_{t-\tau/2}^{t+\tau/2} f(t')dt'.$$

assume that the operator $\overline{f}$ satisfies the usual properties,

$$(\overline{f})' = \overline{f'} = 0,$$

$$((\overline{f})) = \overline{f}, \ (f')' = f'.$$

The quadratic cost function, $\mathcal{J}$ for states f and g may then be expressed as

$$\mathcal{J}(f,g) = \frac{1}{2} \int_{\mathcal{D}} (\overline{f} - \overline{g})^2 + 2(\overline{f} - \overline{g})(f' - g') + (f' - g')^2 \ dV,$$

hence, assuming that $\|f' - g'\|$ is initially small (i.e. that displacement errors are initially large scale,)

$$\mathcal{J}(f,g) = \mathcal{J}(\overline{f},\overline{g}) + \mathcal{O}(\|f' - g'\|).$$

Considering the rate of change of $\mathcal{J}$ with s, under the assumptions

$$\frac{\partial \overline{f}}{\partial s} + \boldsymbol{u}_a \cdot \nabla \overline{f} = 0,$$

$$\frac{\partial f'}{\partial s} + \boldsymbol{u}_a \cdot \nabla f' = 0,$$

we obtain the equation

$$\frac{\partial \mathcal{J}}{\partial s} = \int_{\mathcal{D}} \left[(\overline{f} - \overline{g}) + (f' - g') \right] \boldsymbol{u}_a \cdot \nabla \left[(\overline{f} - \overline{g}) + (f' - g') \right] \ dV. \qquad (4.14)$$

Comparing the magnitude of terms in the equation (4.14), changes due to the low pass terms marked with overbars dominate if the two conditions

$$\|\overline{f} - \overline{g}\| \geq \|f' - g'\|, \qquad (4.15)$$

$$\frac{\|\overline{f}\|}{\overline{L}} \gg \frac{\|f'\|}{L'}, \qquad (4.16)$$

are satisfied, where $\overline{L}$ and L' are length scales for the mean and eddy fields respectively, so that for example $\nabla f' \sim \|f'\|/L'$.

Equation (4.15) is assumed satisfied naturally, at least initially, for most geophysical fields for suitably small $\overline{L}$. Plotting a graph of the terms in (4.16) for

---

This attenuates signals on significantly faster than τ, without affecting the amplitude of the long scales and is perhaps the simplest filter in use within the oceanographic community. More complex definitions based on various transform methods are possible however, for examples see Lynch (1997).
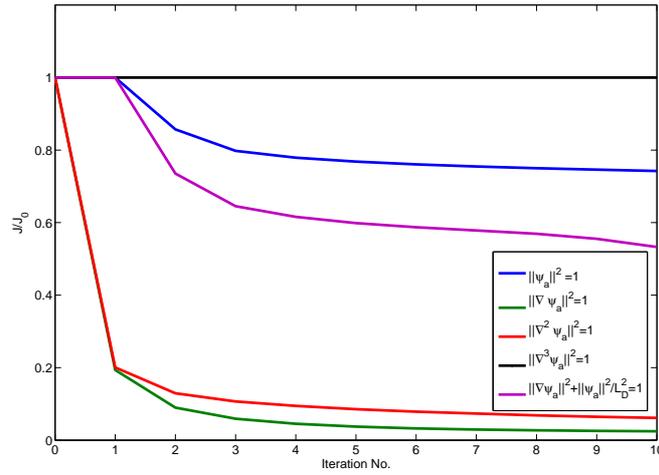
Figure 4.10: Convergence rates for various choices of norm

The plot shows the normalized cost function, $\mathcal{J}/J_0 := \|\rho - \rho_o\|/\|\rho_b - \rho_o\|$, for the shifted sine curve problem of Section 4.2.4. For these runs $n_x = n_z = 128$, $k = 10$ and $\theta = 0.25$. The colours denote the relevant function used for the normalization condition, $\|\mathcal{L}\psi_a\|^2 = 1$. The two functions involving linear terms in ψ_a show no convergence on the first iteration due to problems with the initial choice of step size and are clearly less successful than norms on higher derivatives.
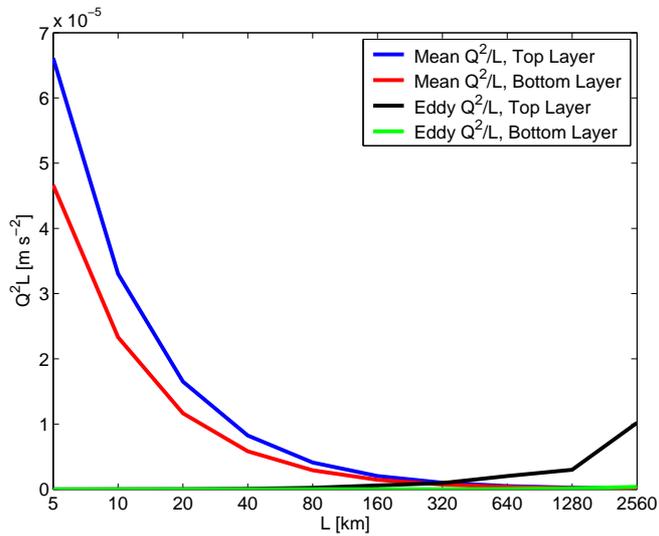
Figure 4.11: Graph of mean and eddy $\|q\|^2/L$ for QG model

The plot shows terms $\|\overline{q_i}\|^2/L$, $\|q_i'\|^2/L$ against length scale, L, for mean and eddy vorticities calculated from $q_i = (\nabla^2 - f_0^2/g'H_i)\psi_i + \beta y$, under the low pass filter of Section 4.3.2. Filtered terms contain more energy for $L < 300$km, in part due to the β term. From equation (4.16) this justifies the multiscale approach.

vorticity data taken from a realization of the QG model produces the results shown in Figure 4.11. This shows that for this case both conditions are initially satisfied for low pass filters acting over the order of eight grid points (80 km) since in both the top and bottom layers the mean function significantly exceeds the eddy function, hence under these conditions the search direction is at first dominated by $\overline{f}$, so that $\boldsymbol{u}_a \approx \nabla \cdot [(\overline{f} - \overline{g})\overline{f}]$.

Consider now the rate of change of $\boldsymbol{u}_a$ itself with s by considering the leading order term in acceptable second order variations of $\mathcal{J}$,

$$\frac{\partial \boldsymbol{u}_a}{\partial s} = \nabla \cdot (\overline{g} - 2\overline{f})\boldsymbol{u}_a \cdot \nabla \overline{f} \quad . \qquad (4.17)$$
$$\sim (\|\overline{\delta}\| + 2\|\overline{f}\|)\|\overline{f}\|/\overline{L}^2$$

This limit will generally decrease with increasing $\overline{L}$, hence the multigrid approach will usually generate more stable search directions than a code based on iterating at only one length scale, again aiding convergence.

Unfortunately a naive approach based purely on mapping the initial background field to a coarse grid alone and performing all calculations there will in general fail to generate an analysis that satisfies the discretized version of the rearrangement condition when re-interpolated back onto the model grid of the background field for outputting analysis states. This is due to the large implicit mixing involved in such an interpolation step. The effect would be to spread fronts and mix water masses defeating the original goal of an adiabatic assimilation method. The obvious compromise is to perform the assimilation using a coarse grid to calculate the ψ_a while retaining information on model variables on a fine-scale grid of the same resolution as the original data. Two approaches are possible, a simple spline fit to a coarse ψ_a field and a scheme based on advecting the mean field through large scale fluxes, while retaining the fine scale structure in situ. The precise numerical details of the implementations of such routines are given in section B.2 in the appendix but a short description is given here.

The first technique involves simply 'super-obing' observations, $\boldsymbol{\rho}^o$, (that is the grouping together of nearby observations into one smoothed or averaged observation) and smoothing the current analysis state, $\boldsymbol{\rho}$, onto a coarse grid using the values from square sets of neighbouring grid boxes. This is performed, for data with observations available on a regular grid, using the operator

$$\overline{f}_{I,J} = \frac{1}{h_{I,J}} \sum_{i=m\times N+1}^{(m+1)*N} \left[\sum_{j=m\times N+1}^{(m+1)*N} [f_{i,j}] \right] ,$$

where $h_{I,J}$ is a term, introduced in anticipation of the case of limited observations, equal to the number of observed elements in the relevant coarse scale grid-box for ρ^o and to the number of background data points (i.e. N^2) in the case of the current iteration of the analysis state, ρ. From these fields a coarse $\overline{\psi_a}$ is calculated using the two-dimensional form of equation (4.17), which is then re-interpolated down to the data grid for the evolution time step, which uses the same approach as before. In this case while the returned velocity field is smooth and large scale, the dimension of problem for the advection step, under which the bulk of the work is done in the method, remains unchanged.

The second technique represents a true multigrid approach to the problem. The same approach as in the first technique is used to generate coarse-scale observations, $\overline{\rho^o}$, analysis, $\overline{\rho}$ and optimal search path stream function $\overline{\psi_a}$. This $\overline{\psi_a}$ is then used to calculate a coarse scale velocity field to be used to evolve the pair of state vectors, $(\overline{\rho}, \rho')$, through two independent evolution equations,

$$\frac{\partial \overline{\rho}}{\partial s} + \overline{\boldsymbol{u}_a} \cdot \nabla \overline{\rho} = 0,$$

$$\frac{\partial \rho'}{\partial s} = 0,$$

where $\rho' = rho - \rho$ is the remaining eddy, or noise field left when the coarse density field is advected. This approach is naturally slightly diffusive on scales below the coarse grid, however the lack of dependence on the eddy terms and the smaller dimension of the advection problem make the code quick both to run and to converge.

Denoting the number of fine scale grid-boxes included in one side of the square coarse scale grid-box by N the method was applied at a number of resolutions to the algebraic problem of Section 4.2.4. Runs at constant $N = 2$ showed improved rates of convergence, but tended to converge to a higher limit than the high resolution $N = 1$ limit. Both the value and rate of convergence of the cost function $\mathcal{J}$ tended to lie somewhere close to that found for the single scale problem of the same dimension, that is discretized with both n_x and n_z halved.

The scheme may also be run with a *variable* value for $N = N^k$, that is as a true multiscale method. Under certain conditions on both the problem and N we may then see not only an improved rate of convergence, but a reduction in the magnitude of the value to which $\mathcal{J}$ converges. In general, the optimal setup appeared to use a saw-tooth or W algorithm acting over the values $N = 1, 2, 4$. Significantly larger choices for N showed little improvement. This suggests that the multigrid method produces a better fit to the global minimum of the system, possibly because
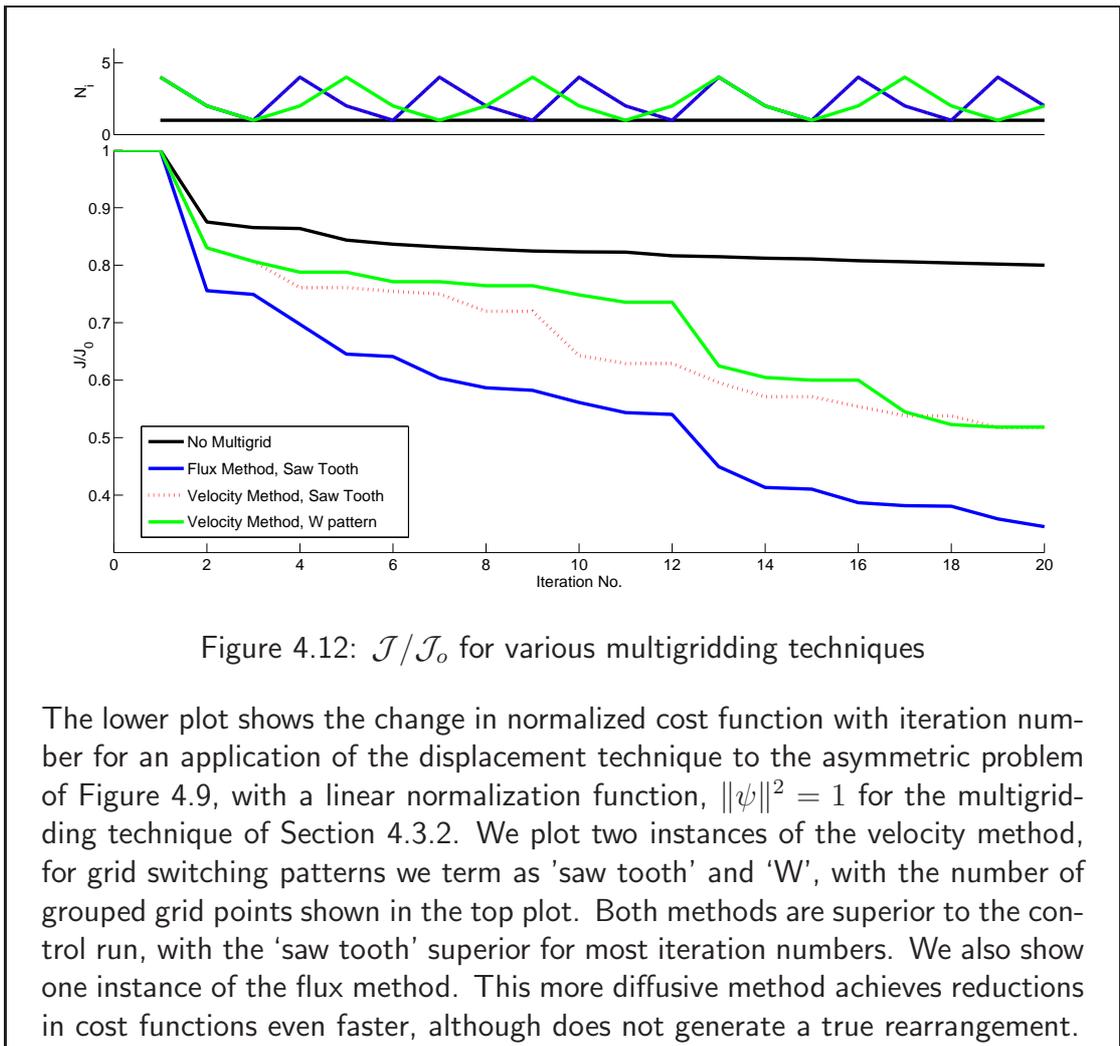
Figure 4.12: $\mathcal{J}/\mathcal{J}_o$ for various multigridding techniques

The lower plot shows the change in normalized cost function with iteration number for an application of the displacement technique to the asymmetric problem of Figure 4.9, with a linear normalization function, $\|\psi\|^2 = 1$ for the multigridding technique of Section 4.3.2. We plot two instances of the velocity method, for grid switching patterns we term as 'saw tooth' and 'W', with the number of grouped grid points shown in the top plot. Both methods are superior to the control run, with the 'saw tooth' superior for most iteration numbers. We also show one instance of the flux method. This more diffusive method achieves reductions in cost functions even faster, although does not generate a true rearrangement.

it allows a better propagation of information across the various length scales on which the scheme operates. Figure 4.12 shows the improvement achieved in the convergence of one particular problem, under both the flux averaged method of super-obing and the velocity averaging method and for both the 'W' and 'saw tooth' algorithm for the grid smoothing/coarsening. For this particular problem, similar to the asymmetric problem illustrated in Figure 4.9, the super-obing flux averaged technique achieves the fastest convergence at nearly three times the rate of the control run.

4.4   DEPENDENCE ON OBSERVATIONS

All previous work has been based on the assumption that values for y were available on all ρ points and at resolution at least equal to the model grid spacing, so that the field may be considered differentiable. In this section we consider the general dependence of adiabatic assimilation methods on observations, particularly in the case where observations are available only at a smaller number of points, which may lie off the model grid or be observations of variables not directly modelled.

4.4.1   ASSIMILATING POINT OBSERVATIONS

The vast majority of real ocean observations are modelled as point observations, that is they represent the value of the observed quantity averaged over length and time scales much smaller than the grid-length and time steps of numerical models. In fact these scales are so small that they can be well approximated as Dirac delta functions in space and time, so that for an observation $y_i[\boldsymbol{r_i}, t]$ of a model variable, valid at at point $\boldsymbol{r_i}$ at time t the observation operator h_i acting on a model state $x(t)$ is

$$h_i[x] = \int_{\Omega} x \delta(\boldsymbol{r} - \boldsymbol{r_i}) dV. \tag{4.18}$$

Combining such discrete data with a continuous constraint is fraught with difficulty so instead in this section we consider the behaviour of a spatially discretized version of the auxiliary equation governing the evolution of the assimilated field,

$$\frac{\partial X}{\partial s} = \boldsymbol{u}_a \cdot \nabla X,$$

when deriving the assimilation cost function.

The model field $x(\boldsymbol{r}, s)$ is replaced with a model vector $\boldsymbol{x}(s)$ constrained element-

wise by the evolution equation

$$\frac{\partial x_i}{\partial s} = -J_{ijk}\psi_j x_k = -J_i[\boldsymbol{\psi}, \boldsymbol{x}],$$

here J_{ijk} is the tensor operator representing the numerical discretization of the Jacobian advection operator, $J(\psi, x)$. The discretized cost function becomes

$$\mathcal{J}(s) = \frac{1}{2}(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{X}(s))^T(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{X}(s))$$

so that the rate of change of $\mathcal{J}$ with respect to s is

$$\frac{\partial \mathcal{J}}{\partial s} = (\boldsymbol{y} - \boldsymbol{H}\boldsymbol{X}(s))^T \boldsymbol{H}\boldsymbol{J}[\boldsymbol{\psi}, \boldsymbol{x}].$$

The normalization condition, equation (4.11), must also be discretized. Allowing $\boldsymbol{V}$ to be a diagonal matrix containing the element volume values of the discretization,

$$V_{ij} = \begin{cases} \Delta V_i & i = j \\ 0 & i \neq j \end{cases},$$

we may write

$$\int_\Omega \boldsymbol{u}.\boldsymbol{u}\,dV \approx \boldsymbol{\psi}^T \boldsymbol{D}^T \boldsymbol{V} \boldsymbol{D} \boldsymbol{\psi} = 1,$$

where $\boldsymbol{D}$ is the discretized matrix operator representing the gradient operator. Using the method of Lagrange multipliers and once again seeking the streamfunction that causes the greatest rate of descent in $\mathcal{J}$ we find

$$\boldsymbol{\psi} = \lambda(\boldsymbol{D}^T \boldsymbol{V} \boldsymbol{D})^{-1}\boldsymbol{J}[\boldsymbol{H}^T(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{x}), \boldsymbol{x}].$$

The major difference between this form and the form with full observations is the introduction of the two observation operator terms, $\boldsymbol{H}^T\boldsymbol{y}$ and $\boldsymbol{H}^T\boldsymbol{H}\boldsymbol{x}$. For point observations on grid nodes the latter term has the extremely simple form

$$\left[\boldsymbol{H}^T\boldsymbol{H}\boldsymbol{x}\right]_i = \begin{cases} x_i & \text{point } \boldsymbol{r}_i \text{ observed} \\ 0 & \text{otherwise.} \end{cases}$$

The precise form of the elliptic problem, $(\boldsymbol{D}^T\boldsymbol{V}\boldsymbol{D})^{-1}$, to be solved is highly dependent on the discretization chosen for the operator $\boldsymbol{J}[\boldsymbol{\psi}, \boldsymbol{x}]$. If the energy and enstrophy preserving Arakawa Jacobian formulation is used (see Appendix A.1), then the operator acts to spread and weight a single observation over neighbouring grid-points, such that the implied direction of flow for such a streamfunction

lies parallel to the spatial gradient of the field x. This agrees well with our natural instinct for how to correct an error at a single point though an advective displacement, by shifting tracer properties down-gradient to match the observation. Although the precise numerics are dependent on our choice of D the effect is to smooth this initial guess. In general the streamfunction produced will have non-zero curl everywhere and thus induce a displacement across the entire basin. This is a consequence of the global nature of the incompressibility and rearrangement conditions. Nevertheless, the majority of the 'assimilation' energy, $\psi^T D^T D \psi$, and the largest displacements both remain close to the observation point.

4.4.2   Assimilating Limited Observations

The question remains as to whether the method can successfully generate rearrangements of a field to match to observations when given limited observational data, without significantly degrading the structures in the analysis field, without large unphysical displacements and without introducing spurious assimilation features. To investigate this the algorithm was run repeatedly using the same states as background and truth each time, but varying the spacing of the assumed observational network, so that we have observations at every n^{th} grid-point, where $n \in 1, 2, 4, 8, 16$. Although this setup is an unrealistic model of the actual GOOS, as shown in Chapter 1, we choose a regular grid structure so that the induced pattern in the analysis field is regular, and thus clearer to the naked eye. Even with a globally irregular grid of random observations there will exist local regions of observations on which an approximately regular grid structure can be imposed. The normalized rate of convergence for these experiments is shown in Figure 4.13, with the final analysis fields plotted in Figure 4.14. Studying these it is obvious that :

 (i). In all cases the method successfully reduces the cost function over the course of the assimilation step. In this sense the algorithm is successfully assimilating the data provided to it.

(ii). The rate of convergence is generally slower as the network becomes more limited, except for the case of the very coarse network, which has fewer observation points than the coarsest model grid used in the multigridding routine.

(iii). As the grain of the observational network coarsens then the magnitude of the spurious displacements introduced increases.

To demonstrate the second point we refer to Figure 4.15. This shows the complex modulus of the result of the two-dimensional Fast Fourier Transform (FFT)
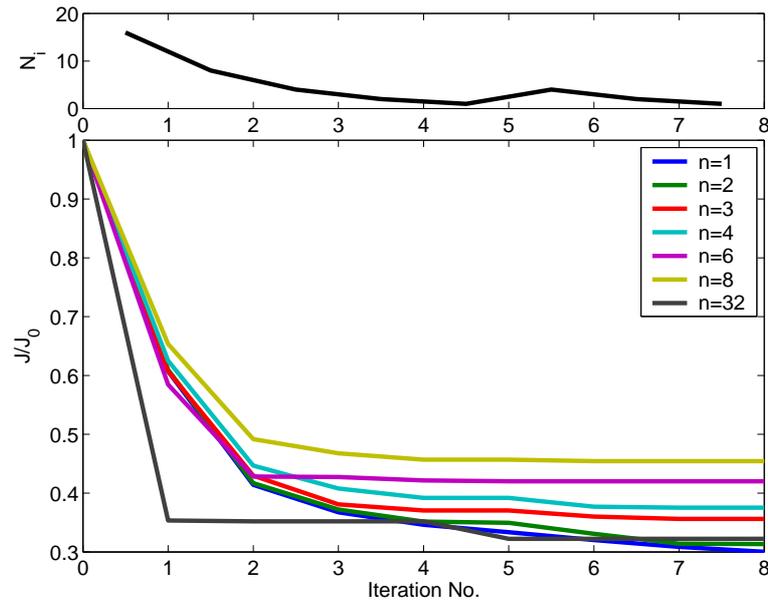
Figure 4.13: $\mathcal{J}/\mathcal{J}_0$ for limited observational networks

The lower plot shows $\|q - q_o\| / \|q_b - q_o\|$ against iteration number for applications of the displacement technique to the problem shown in Figure 4.14 for some choices of limited observations on regular grids. We use a gradient normalization , $\|\nabla \psi_a\|^2 = 1$ and the multigridding scheme shown on the top plot. Here $n\Delta x$ denotes the length of separation between observations, where Δx is the base model grid length. The fine scale and very coarse observations show the same limit of convergence, although the coarse grid initially converges more quickly. The intermediate grids show both slow and poorer convergence under the scheme.
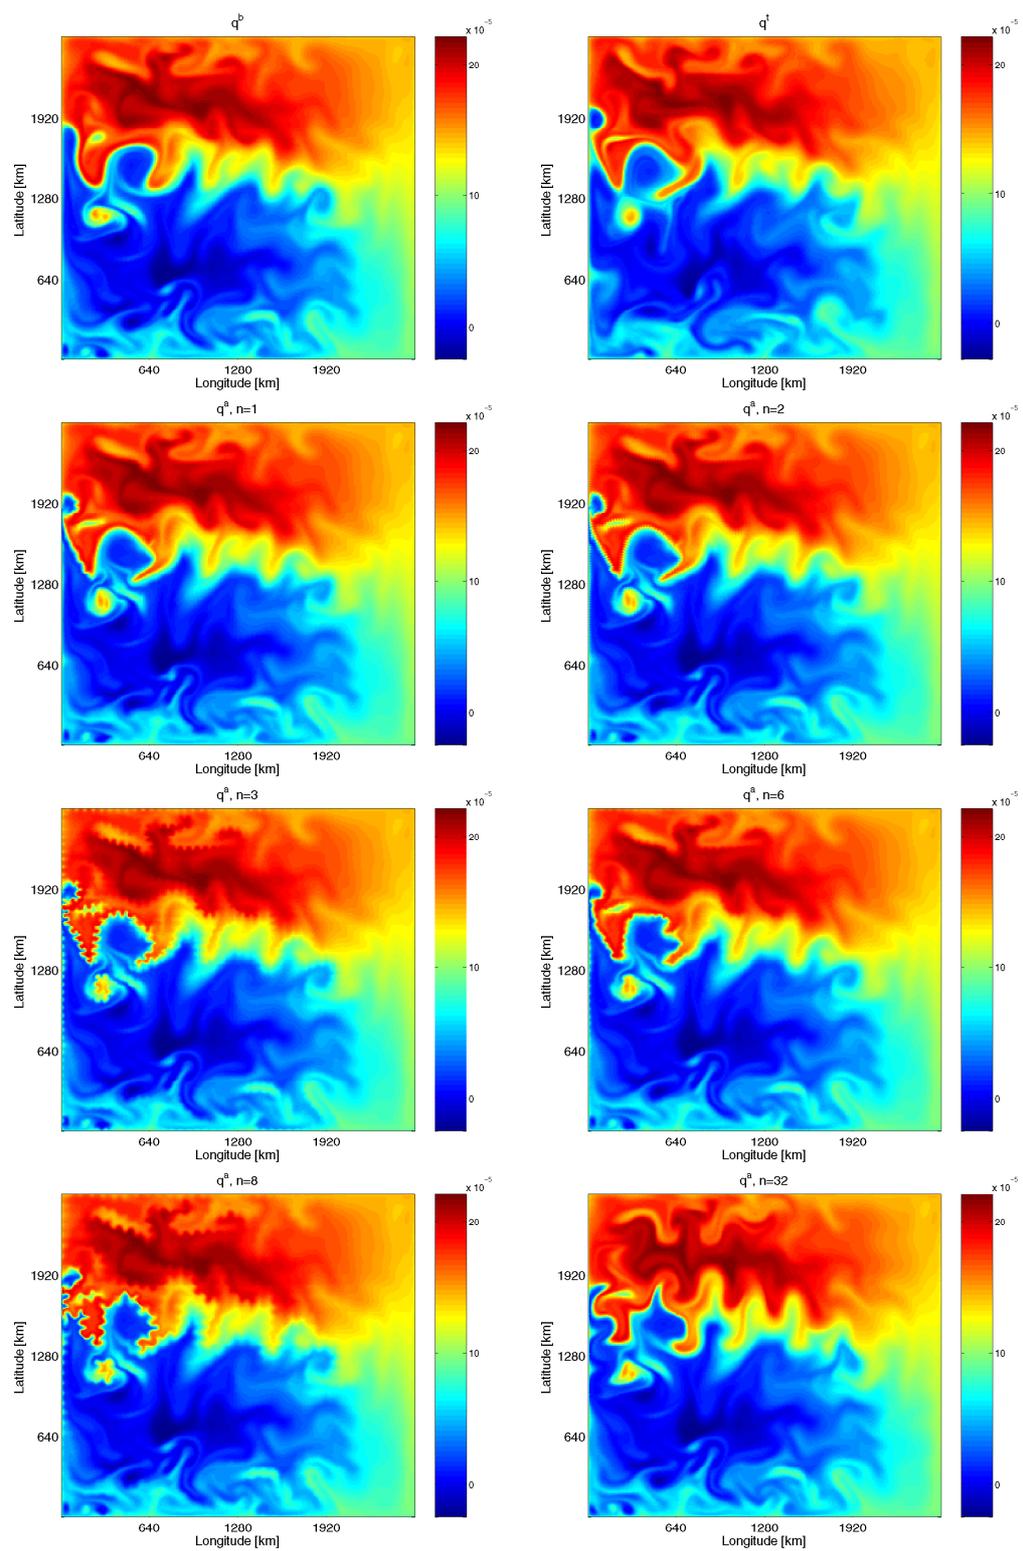
Figure 4.14: Plots of truth, background and analysis for limited observations

Plots showing final analysis fields for the displacement technique under the methods described in Figure 4.13. Note the qualitative successes of the problems with $n = 1, 2, 32$ as well as the 'crinkling' which occurs on the length scale on which observations are available, particularly clear for $n = 3$.

algorithm (Ramirez, 1985) applied to the full final analysis error field , $x_a - x_t$. This algorithm separates the energy contained in the signal into its different (in this case spatial) frequencies. For comparison the FFT of the function

$$g(x, y) = \sin(\frac{\pi x}{L_x}) \sin(\frac{\pi y}{L_x}),$$

is also shown. Initially the field shows the majority of its energy concentrated on mesoscale eddy scales and below. As n increases there is a noticeable shift in energy to scales proportional to $n\Delta x$, visible as the crossed yellow bars in the 2D FFTs. This shows the features of the analysis generated by the assimilation method are directly dependent on the density of the observation network.

4.4.3 THE EFFECT OF NOISY OBSERVATIONS

It is necessary to consider the effect of noise on observations, since in real life observations are prone to error, whether it be truncation of a decimal expansion, the use of a measurement of the temperature of one bucket of water to represent the temperature of an ocean bulk surface temperature or the uncertainty of an altimeter measurement caused by drift in the satellite platform's orbit. In standard assimilation techniques such as 3D-VAR estimates of this error are compared with uncertainties in the forecast of the background state to produce an average state compatible with both. In general such an approach is not compatible with the philosophy of displacement assimilation, where observations are assumed to represent a measure of the Lagrangian, rather than Eulerian, flow pattern, since the observations are assumed to contain a much more accurate representation of true phase than the background field.

In this interpretation noise represents (ideally) small scale false displacements on the back of the large scale true displacement required to assimilate the observations[4]. The errors become important when the magnitude of the displacements required to match to the errors in the observations become larger than those required to correct the difference between the background state and the truth.

It is unfeasible to study the whole gamut of behaviour exhibited by the errors in oceanic observations, however we may show that the issue does not in principle cause all displacement assimilation methods to fail. Considering the highly ide-

---

[4]We assume here that the registered locations of all observations are accurate. If these are not then the observations themselves contain significant displacement errors and any assimilation scheme is likely to fail. This is unlikely with modern data however, due to the very high order accuracy of satellite positioning systems such as GPS, relative to the significant length scales in the ocean. Regular quality control checks at operational centres act quickly to black-list observational sources which regularly mis-report their locations.
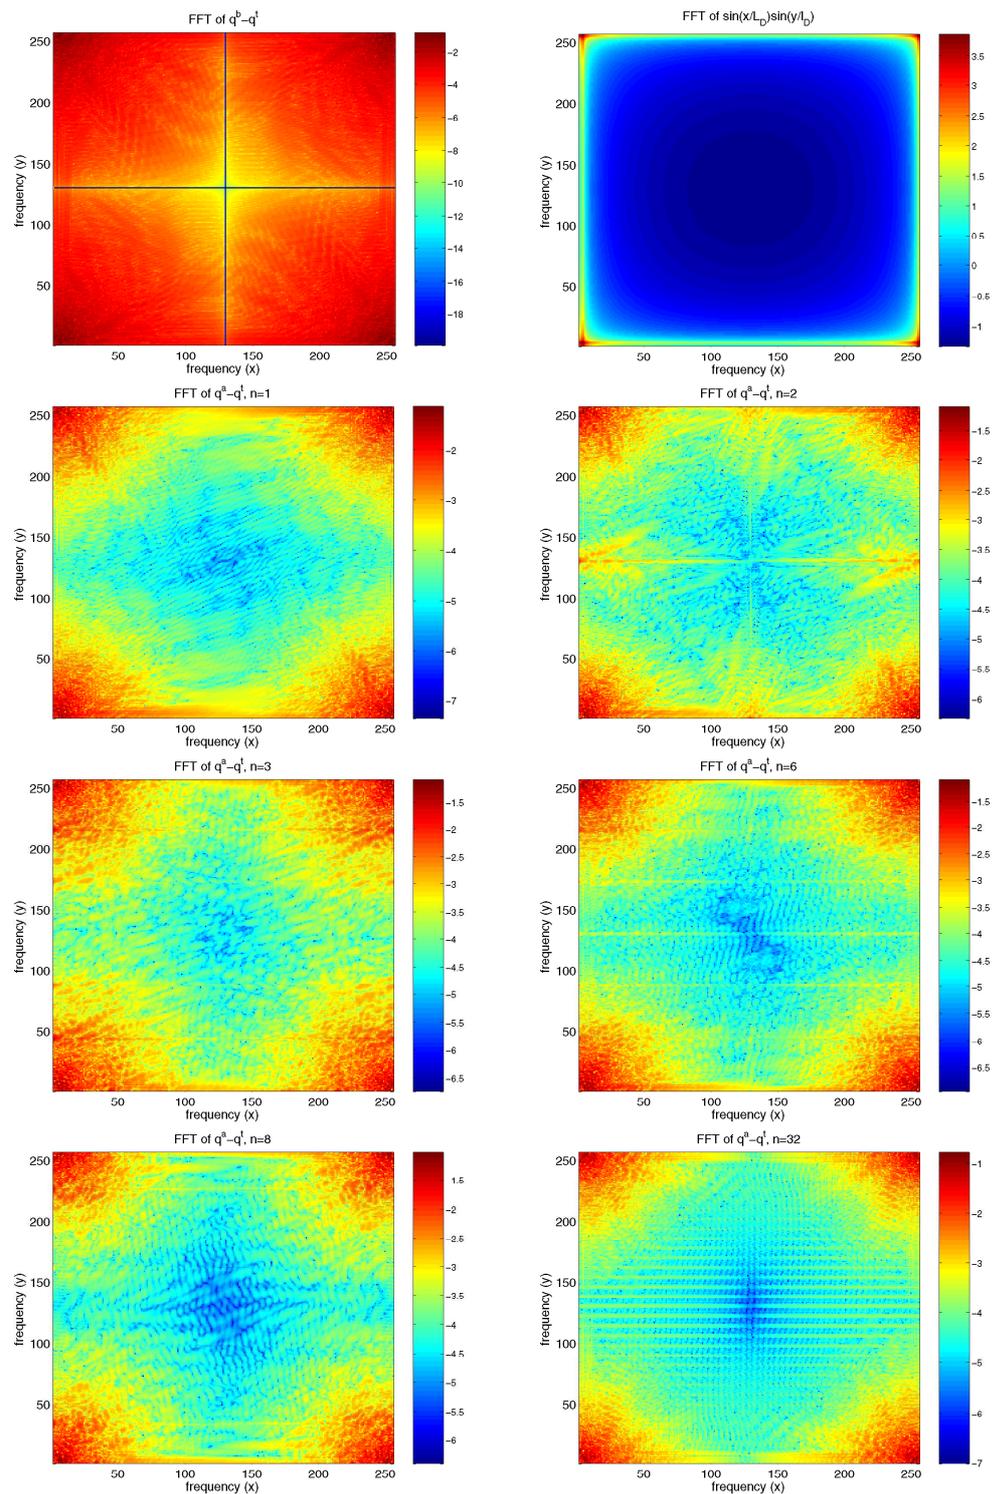
Figure 4.15: FFT analysis of analysis errors

Plots shows 2D FFTs for analysis errors of the plots in Figure 4.14. The obvious horizontal yellow peaks, particularly for $n = 2, 6, 32$ are another example of the energy the assimilation introduces on the wavelength of the observations, the number of peaks increasing in line with n. The asymmetry in the pattern is inherited from the system itself, with strong western boundary currents and weak return flows.

alised situation of the effects of Gaussian noise on the one dimensional problem of Section 4.2.3 produces the results of Figure 4.16, in which noise from a variety of distributions, both biased and unbiased is added to the observations before they are assimilated. In this case the algorithm is particularly insensitive to the noise, since in one dimension the scheme is unable to change the Lagrangian value of the spatial derivative of the background,

$$\frac{d}{ds}\left(\frac{d\rho}{dx}\right) := \left(\frac{\partial}{\partial s} + u_a \frac{\partial}{\partial x}\right)\frac{d\rho}{dx} = 0.$$

This means that the fitting is, in some sense, the matching, through phase shifts, of the two shapes, background and observed, by minimizing some measure of the difference. Large errors are thus only possible when they can be aliased to the shape of the background field, as in the case of large errors, like the second plot of the figure.

In two dimensions the flow field may modify the local gradient of the tracer, since for any scalar field, f, we have

$$\frac{d}{ds}\nabla f := \left(\frac{\partial}{\partial s} + \boldsymbol{u}_a \cdot \nabla\right)\nabla f = \nabla f(\nabla \times \boldsymbol{u}_a) = \nabla f \nabla^2 \psi_a,$$

and the assimilation velocity field is not constrained as irrotational. As such the algorithm may modify the shape of the background field to fit any noise in the observations. This makes it necessary to consider the noise introduced not only into the observed field itself but also into the derived discretized derivative fields. An error of magnitude order ϵ in a discretized field ρ_t can introduce an order $\epsilon/\Delta x$ absolute error into the predicted gradient given by a simple finite differences approach. The relative error is thus $2\epsilon/\|\Delta\rho\|$ which may be very large in areas of small $\|\Delta\rho\|$. Fortunately, provided the background field is close to the truth, these areas of large relative error in observed gradient will be areas of small contribution to ψ_a, since these will be weighted by terms of order $\|\Delta\rho_b\|$.

4.4.4   ASSIMILATING NON-TRACER OBSERVATIONS

The algorithm used is easily modified to consider the question of multiple variables, again following the approach of 3D-VAR algorithms. The matrix $\boldsymbol{H}$ of section 4.4.1 can be replaced by a general non-linear observation operator on $\boldsymbol{X}$, with $\boldsymbol{H}^t$ replaced by the operator adjoint. This approach is described in more detail in the next chapter, in the context of assimilating stream-function observations to correct vorticity fields.
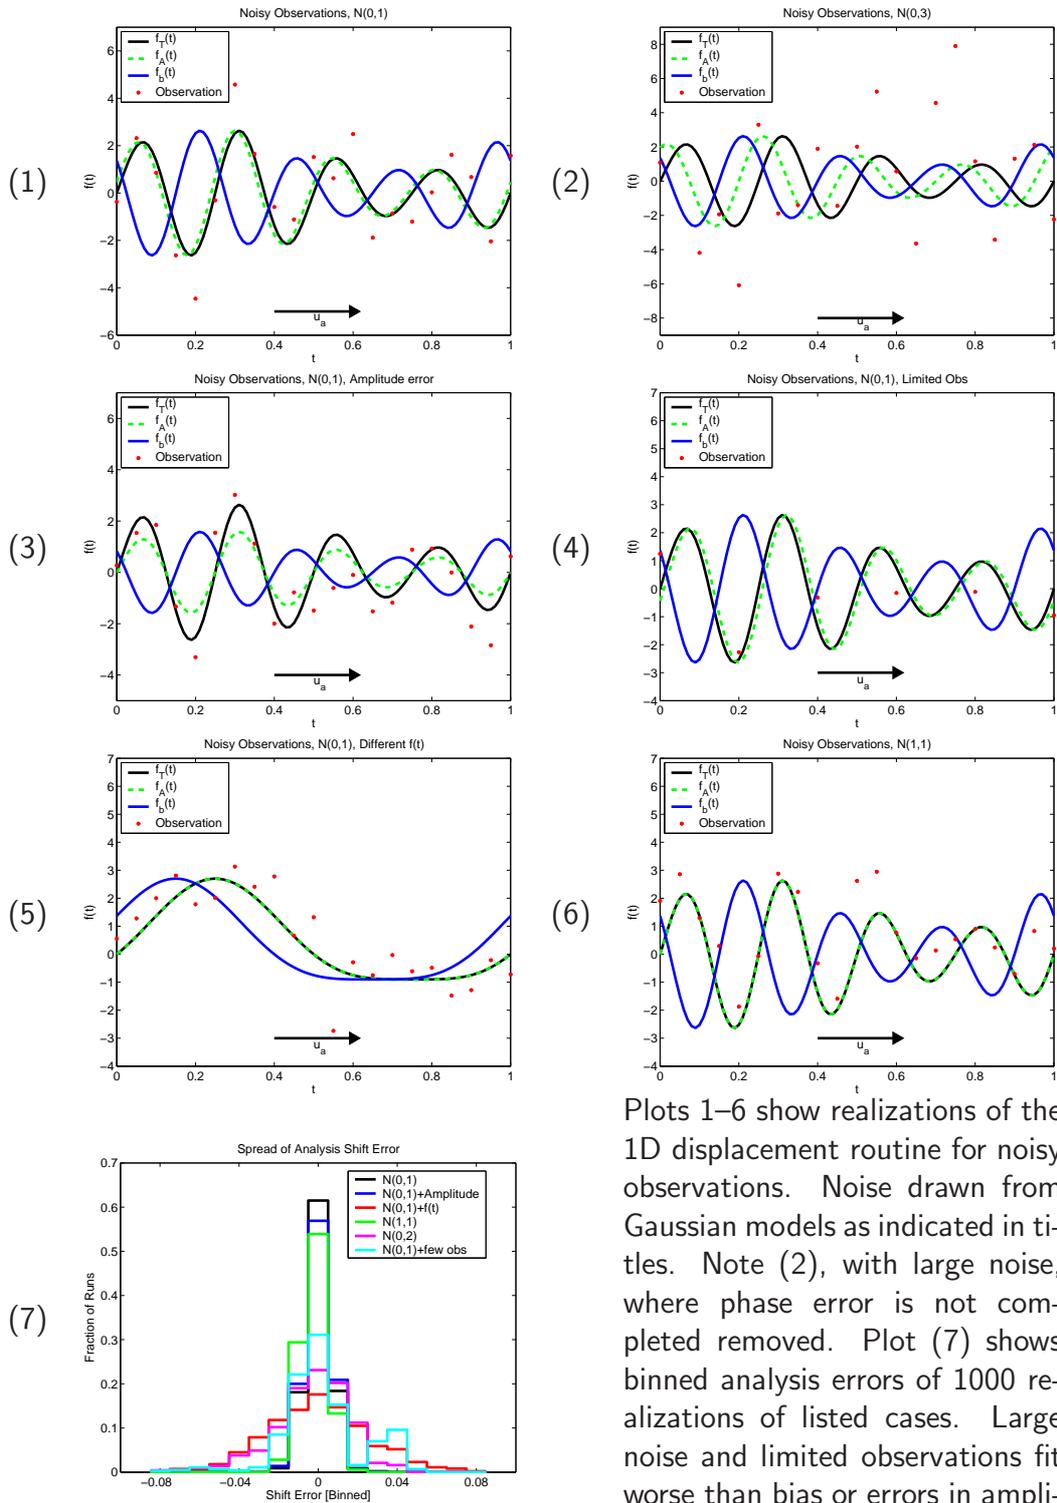
Plots 1–6 show realizations of the 1D displacement routine for noisy observations. Noise drawn from Gaussian models as indicated in titles. Note (2), with large noise, where phase error is not completed removed. Plot (7) shows binned analysis errors of 1000 realizations of listed cases. Large noise and limited observations fit worse than bias or errors in amplitude.

Figure 4.16: The effect of noise on a 1D rearrangement problem

4.5   CHAPTER SUMMARY

This chapter has introduced a family of algorithms which directly minimise a 3D-VAR style cost-function with respect to displacements which can be generated through incompressible advection. This represents an attempt to limit the set of re-arrangements over which the field is minimized to a set of diffeomorphisms of the original field, retaining a physical basis for the displacement. Existence of solutions and convergence properties are discussed. The method does suffer from issues related to non-uniqueness of solutions and the existence of local minima, plus some instability with respect to small changes in the observations provided. The method is applied to a series of progressively more complex problems, culminating in synthetic data representing the distribution of a tracer field in an idealized ocean basin. We find a strong dependence on the form, accuracy and number of the observations supplied to the routine in the ability of the method to generate a successful analysis field.

Chapter 5

4D Displacement Assimilation

"Time, like an ever-rolling stream
  Bears all its sons away;
  They fly forgotten, as a dream
  Dies at the opening day."

Isaac Watts
Man Frail and God Eternal, 1719

Both four dimensional variational assimilation (4D-VAR) and Kalman smoother (though not the Kalman filter, which only carries forward information on the error co-variances) make simultaneous use of observations valid at multiple observation times while calculating an analysis state. In this way they extract the maximum information from the assumed coupling between data in observations and the equations of motion of numerical models. The model equations act to transport information through 'data holes' within observations and thus increase the effective domain, in both space and time, over which observations are valid.

Given the dependence on high-density or full observations of the displacement techniques demonstrated in the previous chapter and the near adiabatic nature of oceanic flows over short time periods it is hypothesised that applying a similar 'time-windowed' approach to the development of a displacement technique may generate an efficient algorithm for displacement data assimilation with operational data. A method is here suggested which achieves this goal through the introduction of an additional 'assimilation velocity' term directly into the tracer equation of a numerical model, which is then used to evolve the analysis state over the time-window within which information contained in observations is assumed relevant to subsequent forecasts.

By choosing this assimilation velocity term to minimise a 4D-VAR style cost-function applied across the whole time window the method is found to generate a final time model analysis state which lies closer to the assumed identical twin true state, as well as generating an augmented model trajectory which converges towards that of the true state over the course of the time-window. In this sense the algorithm may be thought of as correcting for both the initial background error and correcting that fraction of

the model error which may be parameterized through the advective term.

5.1 OBSERVATIONS, INFORMATION AND TIME

Previous implementations of displacement techniques, such as Cooper & Haines or the methods of Brewster and Hoffman presented in Section 1.6, as well as the QG implementation of CH demonstrated in Section 3.3.1, operate only on data available at a single verification time. These methods are generally derived based on an assumption that full observations, or a retrieval thereof, are available over the entire domain of interest at that time, on scales comparable with model resolution. On the other hand, the new generation of data assimilation methods used at operational centres work on the assumption of an interval, or time window, over which spatially separated observations are fitted, using a dynamical model for state evolution. An obvious question is whether this four-dimensional viewpoint can be naturally extended into the philosophy of displacement assimilation while preserving some notion of the rearrangement condition of Section 3.1.2. In this chapter we demonstrate that it can, at the notable disadvantage of linearizing the velocity fields which generate the displacements with respect to the model time step.

5.1.1 AN OBSERVATION MEMORY TIME

Consider a general dynamical system, evolving a state vector, $\boldsymbol{x}$, over a set of discrete times, $t_1, \ldots, t_\infty$. Let $\mathcal{F}$ be the functional such that

$$\boldsymbol{x}[t_n] = \mathcal{F}(t_n, \boldsymbol{x}[t_1], \ldots, \boldsymbol{x}[t_{n-1}])^1.$$

Provided that the evolution functional $\mathcal{F}(t_n)$ is not a totally random process at any time level, it has some memory of all its previous states. That is, denoting the *a priori* probability distribution of state $\boldsymbol{x}[t_n]$ by $P(\boldsymbol{x}[t_n])$ and the resulting conditional probability distribution of state $\boldsymbol{x}[t_n]$ given the state $\boldsymbol{x}$ is known at time t_m by $P(\boldsymbol{x}[t_n]|\,\boldsymbol{x}[t_m])$, there exists a relation

$$P(\boldsymbol{x}[t_n]|\boldsymbol{x}[t_m]) \neq P(x[t_n]).$$

This relation holds, through symmetry, regardless of whether t_m lies to the future or past of t_n. In light of the Bayesian assumption of Section 1.3.1 this means that the

---

[1]We have assumed causality in the development of the system. This does not affect the eventual dependence of state estimates on observations which lie in the relative future however.

posterior estimate of the distribution of the state $x[t_n]$ is therefore dependent, not only on observations taken at time t_n, but also on observations of the states lying in the past and future. In general if the combined system and observations are not entirely deterministic, as has been assumed in any formulation of statistically optimal data assimilation methods, then there will exist a 'memory' timescale, $\mathcal{T}$, over which the utility of observations will fade. This property is illustrated for a simple three state system in Figure 5.1.

The magnitude of this timescale depends on the precise structure of the dynamics of $\mathcal{F}$ and on the predictability of the stochastic part of the system. Typically there will exist markedly different time-scales for different modes in the state, for example the fast and slow modes (labelled inertial gravity and Rossby waves respectively) excited in a geophysical fluid under perturbation.

In a famous paper, which later formed the basis for the development of much of modern chaos theory, Lorenz (1963) posited the existence of a timescale after which two different climate conditions end up statistically as similar to a randomly chosen third state of the system as to each other, regardless of how small the initial anomaly between them was. This timescale, derived in Lorenz's analysis from a numerical implementation of a heavily simplified model of atmospheric convection in a box, represents one method for obtaining an estimate of $\mathcal{T}$ for the climate system. Further estimates for this parameter have since been derived by studying the behaviour of perturbed trajectories of more advanced and realistic GCMs. The precise values vary from the widely quoted figure of two weeks (actually 19 days) from Charney et al. (1966) from a selection of then current atmospheric GCMs to 5 months for the doubling time in the growth of fast mode of errors in a large-scale coupled ocean-atmosphere model (Goswami and Shukla, 1991).

Notwithstanding the variability in such estimates, these figures are, for both the atmosphere and ocean, comparable with the longest scales on which numerical forecasts are currently made. They are moreover far longer than the time-scales at which the models themselves operate and over which new observations become available. Equally importantly in the context of displacement theory, the observation memory scales are shorter than the time-scales over which mixing effects dominate, both in the real ocean (see Section 2.1) and for the QG model used here (see Section 3.2). This means there is information on displacement error contained in both past and present observations of the ocean. This information can be extracted, given a proper understanding of the relevant dynamics.
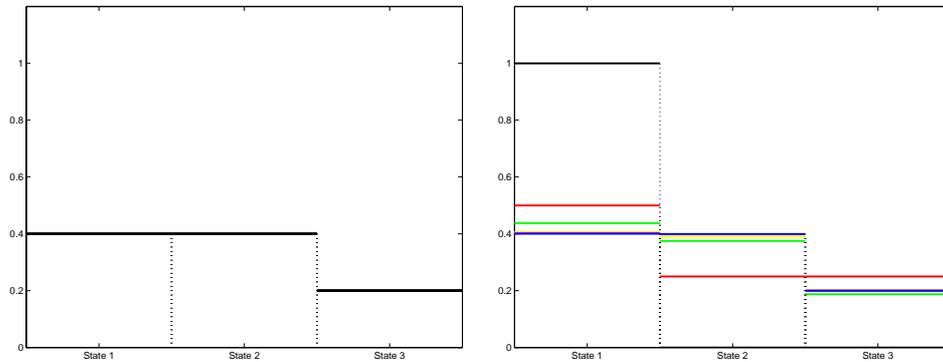
Figure 5.1: Example of observation memory for a simple system

Consider a system with three possible states, controlled by a simple discrete Markov process with an evolution matrix given by

$$A = \begin{pmatrix} \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} & 0 \end{pmatrix}.$$

so that the evolution of the probability distribution $\boldsymbol{x}_{n+1} := P_{n+1}(X = x)$ depends on $P_n(X = x)$ through the relation

$$\boldsymbol{x}_{n+1} = A\boldsymbol{x}_n.$$

The long-term probabilities of such a model existing in each of its states is shown in the left hand plot above. These represent the best probabilities for a guess for the state of the system when no other information is available. The right hand plot shows the probabilistic evolution of the system after a (perfect) observation,

$$\boldsymbol{x}_0 = \begin{pmatrix} 1,0,0 \end{pmatrix}^T$$

This is the PDF given in black. Succeeding PDFs in red, green, yellow, magenta and blue converge back towards the mean state as the system 'forgets' the information contained in the observation.

5.1.2 ASSIMILATION UNDER A TIME WINDOW

Of the assimilation methods which make use of observations at multiple times, including the previously introduced Kalman Smoother (Li and Navon, 2001), as well as the 'Timely' Optimal Interpolation method in use for ocean data assimilation at the UK Met Office[2], perhaps the simplest in form is four dimensional variational data assimilation (4D-VAR). This is also among the most commonly used, being the form of assimilation algorithm used operationally at the ECMWF for ocean data assimilation of most in situ data (Weaver et al., 2003; Vialard et al., 2003), as well as at both the UK Met Office and ECMWF for atmospheric data assimilation. The method is described in some detail here, using the discrete Ide notation of section 1.3.1 and letting $x[t_n]$ denote a system state valid at model time t_n.

Problem 5.1 (Strong Formulation 4D-VAR) *Given vectors of observations $y[t_n]$, $t_n \in [t_0, \ldots, t_f]$ and a background state for x at time t_0, find the initial condition analysis state vector $x_a[t_o]$ which minimizes a cost function, $\mathcal{J}(x[t_0]; y, x_b)$ defined over a time interval $[t_0, t_f]$ by*

$$
\begin{aligned}
\mathcal{J}(x; y, x_b) &= \mathcal{J}_b + \mathcal{J}_o, \\
\mathcal{J}_b &= \frac{1}{2} \left(x[t_o] - x_b \right)^T B^{-1} \left(x[t_o] - x_b \right), \\
\mathcal{J}_o &= \frac{1}{2} \sum_{t_i = t_0}^{t_f} \left(y[t_i] - h(x[t_i], t_i) \right)^T R^{-1}[t_i] \left(y[t_i] - h(x[t_i], t_i) \right),
\end{aligned}
\tag{5.1}
$$

where B and R are, as in Chapter 1, the background and observation error covariance matrices, subject to a strong constraint that the individual states satisfy the prescribed discretized nonlinear model equations exactly,

$$
x[t_{n+1}] = \mathcal{M}(x[t_n], t_n).
\tag{5.2}
$$

In the weak formulation the condition (5.2) is relaxed by allowing for a small discrepancy between the derived analysis state at times greater than t_0 and the state calculated from the model trajectory. This extra parameter is assumed to be able to correct for model errors between the prescribed numerical model, $\mathcal{M}$, and the evolution operator for the true state vector. In both weak and strong formulations the model condition may be enforced using a Lagrange multiplier approach, including

---

[2]the TOI scheme combines the speed of 3D data assimilation with the maximal use of data of 4D methods. Observations are assimilated in a 3D sense, but also evolved with the model and weighted depending on their closeness to the validation time of the relevant analysis (Davey, 2005).

an extra penalty term in the cost function to be minimized;

$$\textbf{Strong Formulation :} \quad \mathcal{J}_{\mathcal{M}} = \sum_{t_i=t_0}^{t_f} \left[\boldsymbol{\theta}[t_i] \cdot (\boldsymbol{x}[t_{i+1}] - \mathcal{M}(\boldsymbol{x}[t_i], t_i)) \right].$$

$$\textbf{Weak Formulation :} \quad \mathcal{J}_{\mathcal{M}} = \sum_{t_i=t_0}^{t_f} \left[\boldsymbol{\theta}[t_i] \cdot (\boldsymbol{x}[t_{i+1}] - \mathcal{M}(\boldsymbol{x}[t_i], t_i) + \boldsymbol{d}[t_i]) + \|\boldsymbol{d}[\boldsymbol{t_i}]\|^2 \right].$$

Extremizing the cost function with respect to the adjoint variable, $\boldsymbol{\theta}$, thus enforces (5.2), up to the minimal discrepancy, $\boldsymbol{d}$, in the weak case.

The minimization problem is typically solved by using an adjoint to the linearized model equations, $\boldsymbol{M}^*$, to calculate the local gradient of the cost function with respect to variations in the initial conditions,

$$\nabla_{\boldsymbol{x}[t_0]}(\mathcal{J} + \mathcal{J}_{\mathcal{M}}) = \boldsymbol{b}^{-1}(\boldsymbol{x}_0 - \boldsymbol{x}_b) + \theta[t_0],$$

where the adjoint variable $\boldsymbol{\theta}$ satisfies the equations

$$\theta[t_i] = \boldsymbol{M}^*(\boldsymbol{\theta}[t_{i+1}], t_{i+1}) + \boldsymbol{H}\boldsymbol{R}^{-1}[t_i](y[t_i] - \boldsymbol{h}(\boldsymbol{x}[t_i], t_i), \quad i \in [t_0, \dots, t_{f-1}],$$

$$\theta[t_f] = \boldsymbol{H}\boldsymbol{R}^{-1}[t_f](y[t_f] - \boldsymbol{h}(\boldsymbol{x}[t_f], t_f),$$

where $\boldsymbol{H}$ is the local linearization of the non-linear observation operator. The calculated gradient is used as a search direction in a suitable descent algorithm.

Under this approach analysis (and of course forecast) states at times greater than t_0 are always calculated directly from the model, $\mathcal{M}$, so that the contribution from the $\mathcal{J}_{\mathcal{M}}$ term is implicitly set to vanish. The values of the θ are calculated by backwards integration of the adjoint model from the terminal condition under the necessary forcings from the observed innovation vectors, $\mathbf{y} - h(\boldsymbol{x})$. This cycling between backwards integration of the adjoint model to obtain a descent path, followed by forward integration of the model equations to find a new approximation to the state trajectory corresponds to the backwards and forwards sweeps of the Kalman Smoother algorithm in that information in the mismatch between observations and background is carried in the positive and negative directions of model time respectively.

When the 4D-VAR scheme is applied operationally approximate or linearized forms of the cost function or adjoint are frequently used to reduce the computational expense of the full algorithm or to speed convergence towards a limiting solution, at the cost of a possibly sub-optimal final analysis state. For example the ECMWF NWP assimilation system operates at half the model resolution of the full

deterministic forecast when assimilating data. The final model trajectory generated can be evolved using the model operator to times greater than the final time of the time window in order to generate a new forecast state. The conceptual difference between using a four dimensional (4D-VAR) and three dimensional (3D-VAR) algorithm to assimilate observations over a short time window in a simple model is illustrated in Figure 5.2. This does not show clearly the advantages of the algorithm though, since in general 4D-VAR algorithms are more suited to problems with limited or noisy observational data. This is because observations valid over a wide range of verification times may be assimilated and the analysis trajectory is constrained to lie on (in the case of strong formulation 4D-VAR) or near (in the case of the weak formulation) a model trajectory.

5.2 A 4D DISPLACEMENT ASSIMILATION TECHNIQUE

5.2.1 TIME WINDOWS AND REARRANGEMENTS

The question remains as to how to combine the advantages of methods based on the 'time window' and 'rearrangement' interpretations of the data assimilation problem consistently. One hypothetical approach would be to couple the increments to the model initial conditions to the three dimensional rearrangement techniques of Chapter 4, constraining the actual applied correction to be in the form of a rearrangement. The method would require both calculation of a 4D-VAR analysis and a rearrangement of the background state that was closest to it, while hoping that the rearrangement state calculated actually represented a better fit to observations than the initial background. Such an approach would be both computationally complicated and inefficient, with no guarantee that the method would actually reduce any objective function.

These issues may be resolved if we remember that the rearrangement condition applies at a single time, whereas the time-dependent nature of algorithms (such as 4D-VAR) is inherited from the form of the cost function to be minimized. By relaxing the rearrangement condition to allow the limit of rearrangements generated by infinitesimal displacements the advective step of the previous chapter may be assumed to occur in model (rather than pseudo) time and the assimilation forcing can thus be represented as an additional advective term in the evolution equation of all material variables of the system.

Under this assumption the evolution equation for an arbitrary model tracer vari-
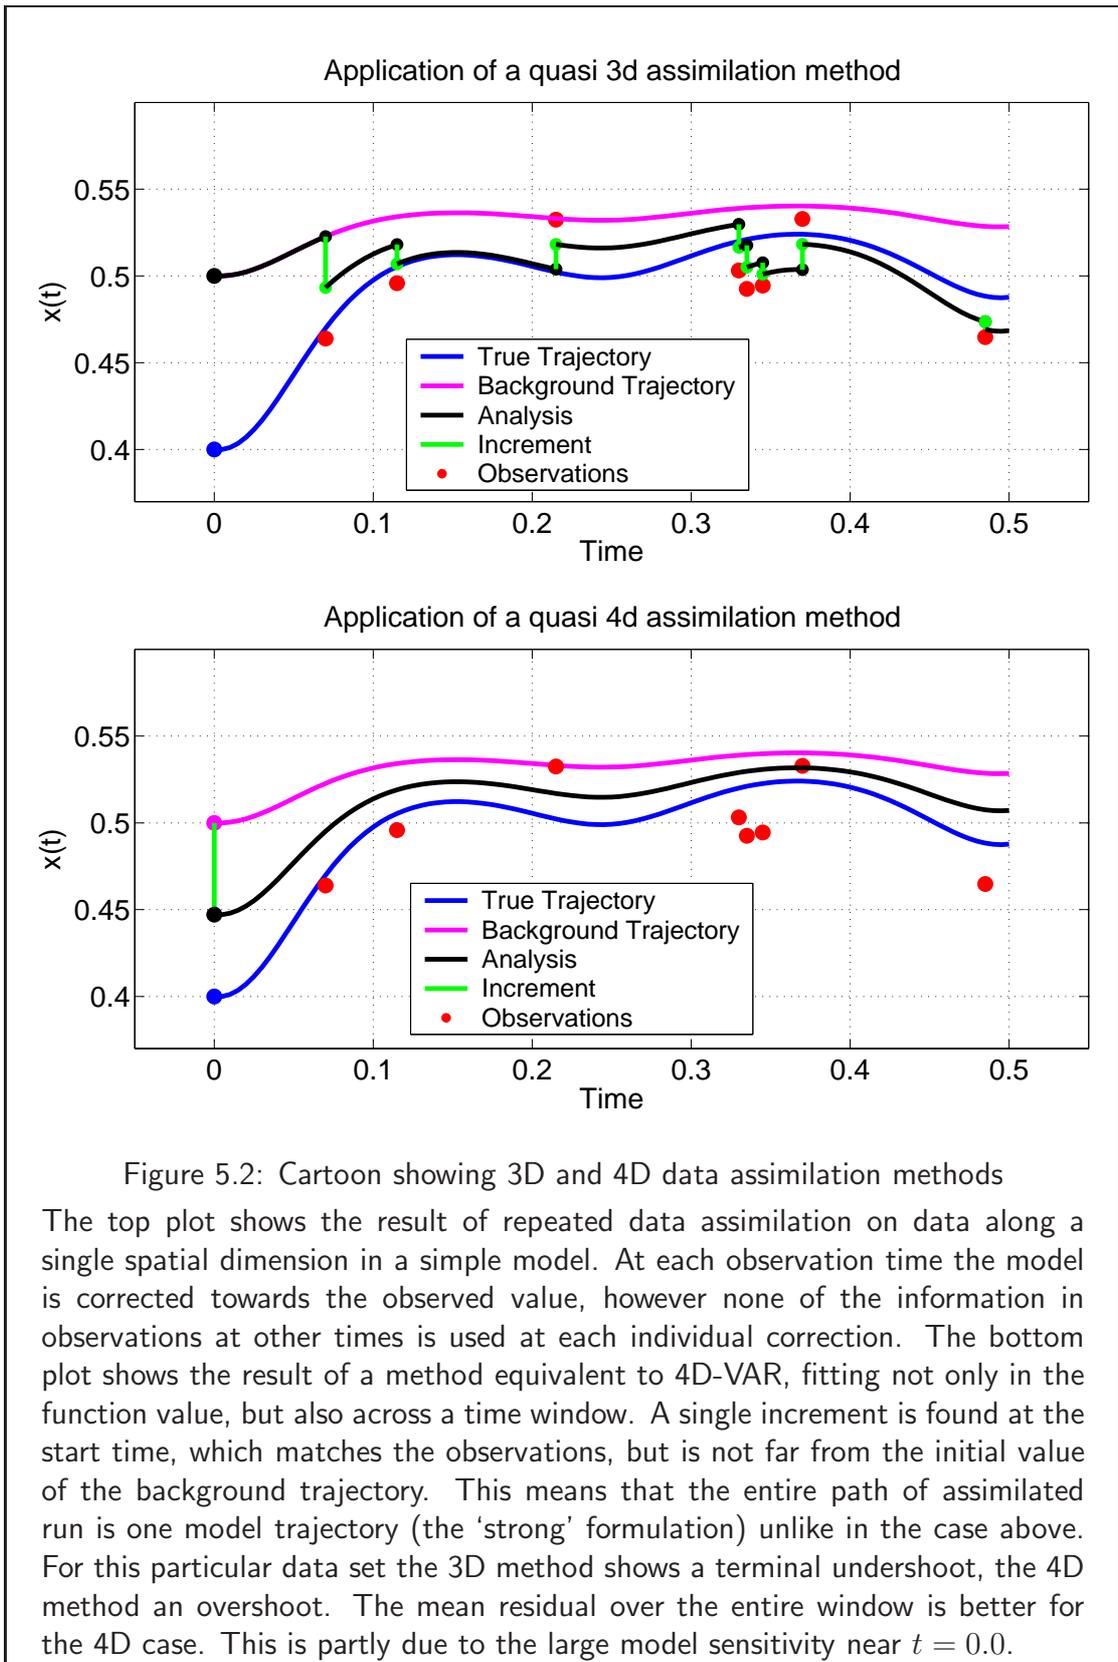
Figure 5.2: Cartoon showing 3D and 4D data assimilation methods

The top plot shows the result of repeated data assimilation on data along a single spatial dimension in a simple model. At each observation time the model is corrected towards the observed value, however none of the information in observations at other times is used at each individual correction. The bottom plot shows the result of a method equivalent to 4D-VAR, fitting not only in the function value, but also across a time window. A single increment is found at the start time, which matches the observations, but is not far from the initial value of the background trajectory. This means that the entire path of assimilated run is one model trajectory (the 'strong' formulation) unlike in the case above. For this particular data set the 3D method shows a terminal undershoot, the 4D method an overshoot. The mean residual over the entire window is better for the 4D case. This is partly due to the large model sensitivity near $t = 0.0$.

able $\tau(x, y, z, t)$ in the domain of interest, $\mathcal{D}$, becomes

$$\frac{\partial \tau}{\partial t} + (\boldsymbol{u} + \boldsymbol{u}_a) \cdot \nabla \tau = F(\tau), \tag{5.3}$$

where $\boldsymbol{u}$ is the standard velocity field derived from the model momentum equation, $\boldsymbol{u}_a$ is the new additional assimilation velocity and $F(\tau)$ denotes the diabatic processes affecting τ, which are assumed small, but non-zero. This is a generalization of the concept of displacement assimilation in what is essentially a nudging of the model, or weak model formulation, while retaining the conception that adiabatic rearrangements through advection dominate the problem.

As suggested previously the new variable, $\boldsymbol{u}_a$, is chosen to minimize a time-windowed observation cost functional,

$$\mathcal{J}(y, \tau) = \int_{t_0}^{t_f} \int_{\mathcal{D}_H} (y - h[\tau])^2 dV dt, \tag{5.4}$$

where the evolution of $\tau(t, \boldsymbol{x})$ is governed exactly by equation (5.3), $y(t, \boldsymbol{x})$ are the observations on the domain $[t_0, t_f] \times \mathcal{D}_H$ to be fitted and $h[\tau, t]$ is as before the observation operator mapping a model state τ to the quantity that it is expected would be observed. The key advantage of this formulation is in its simplicity; the minimization process is for a single function only, guaranteeing that the fit with data, as measured by the cost functional, improves, and removing the "two-step" structure of the hypothetical 4D-VAR plus 3D rearrangements idea suggested above.

5.2.2   Conditions on the Assimilation Velocity

The minimization problem, as presented thus far is seriously ill-posed due to the lack of constraints and the large number of degrees of freedom of even the discretized problem. To alleviate this issue and to reduce the dimension of the space over which the minimization must take place it is possible to list some physically necessary constraints on the chosen form of $\boldsymbol{u}_a$.

(i). The assimilation velocity is assumed incompressible,

$$\nabla \cdot \boldsymbol{u}_a = 0.$$

Firstly this means in the adiabatic limit, where $F(\tau) = 0$, the assimilated model continues to satisfy the rearrangement condition for the entire time-window. Secondly, given that the displacements of interest are slow (by assumption comparable with the timescale on which the model states them-

selves evolve) and of large-scale (again comparable with the model), the assimilation flow itself satisfies the conditions required for the flow velocity field to be assumed solenoidal (Batchelor, 1967, section 3.6). This follows from an assumption that high frequency oscillations, whether in time or space, in the position of a feature can be ascribed to noise in the observations and model.

(ii). The rearrangement assumption requires that the net mass transfer through material boundaries must be zero,

$$\int_{\delta \mathcal{D}} \boldsymbol{u}_a \cdot d\boldsymbol{S} = 0,$$

with $\delta \mathcal{D}$ the boundary of the domain. A consideration of the effects of the displacements at material surfaces (Batchelor, 1967, section 3.3) suggests that this may be achieved through the normal condition of no-normal flow on material boundaries,

$$\boldsymbol{u}_a \cdot \boldsymbol{n} = 0 \text{ on } \delta \mathcal{D},$$

with $\boldsymbol{n}$ any vector normal to $\delta \mathcal{D}$. In principle other choices for the flow boundary condition are possible, for example a sponge condition with inflow along one section of $\delta \mathcal{D}$ balanced by outflow of fluid with identical water mass properties along another section. Whether such a process, for example the removal of excess warm water from the Gulf Stream and its re-introduction near the Mediterranean, has any physical meaning in a data assimilation context for a real ocean model seems unlikely however[3].

(iii). There is no direct physical argument concerning the choice of lateral boundary conditions to apply to $\boldsymbol{u}_a$. Probably the most suitable choice is to apply the same lateral conditions in the data assimilation scheme as are applied to $\boldsymbol{u}$ in the forward ocean model, since it may be argued that the same physics are assumed to apply and this should limit the possible number of extra modes of variability that the assimilation scheme may excite.

There is also the question of what behaviour is to be expected when the observations alone do not contain enough information to make the minimising velocity field unique, as will generally be the case considering the observational data presented in Section 1.4 . In fact, even in the case of full observations of a single tracer field there may exist displacements which leave the cost function (5.4) unchanged,

---

[3]This is certainly obvious for cases with closed boundaries. It could be argued that such a process would represent the required behaviour on open boundaries of a regional model. Nevertheless we would not expect the information present in the outflow region of a regional model to match the information in the inflow region with respect to the existence of eddies, etc.

Figure 5.3: Example of a cost-function preserving rearrangement

The field on the right represents the field on the left rotated anticlockwise about its centre by $90°$, as indicated by the motion of the symbols representing tracer particles. This rearrangement generates no change in any metric on the observed variable, however metrics on particle displacement will show a large signal. It is important that any displacement assimilation method not generate such large scale displacements when not evidenced in data.

for example rearrangements of level sets of the tracer (see Figure 5.3 for an example). This issue cannot be resolved by use of a background term in the manner of the $\mathcal{J}_b$ term of 4D-VAR since this would restrict the ability of the scheme to match with observations without preventing displacements along combined level sets of the background and observations. Similarly, a normalization term, as in Section 4.3.1, would unfairly penalise large displacements in those situations when they are necessary. Instead the problem may be alleviated by initializing the descent routine in a state with zero displacements everywhere and aiming to minimize the excitation of modes orthogonal to the descent direction while gradient sliding.

5.2.3   THE SOLUTION ALGORITHM

By exact analogy with the solution of the 4D-VAR algorithm we couple the evolution constraint, (5.3), to the minimization problem using a Lagrange multiplier approach. Introducing adjoint variables $\boldsymbol{\theta}(t, \boldsymbol{r})$ and an extended cost function,

$$\mathcal{J}(\boldsymbol{u}_a, y) := \int_{t_0}^{t_f} \left[\underbrace{\int_{\mathcal{D}_H} (y - h(q))^2}_{\mathcal{J}_o} + \underbrace{\int_{\mathcal{D}} \boldsymbol{\theta} \cdot \left(\frac{\partial}{\partial t}[q, \boldsymbol{u}] - \mathcal{M}[q, \boldsymbol{u}, \boldsymbol{u}_a] \right)}_{\mathcal{J}_\mathcal{M}} \right] dV \, dt. \quad (5.5)$$

Where here $\mathcal{M}$ denotes the extended non-linear model evolving the state of the system with time, that is to say the standard model together with the extra forced

advective term containing the assimilation velocity. This requires the specification of an initial condition on q and $\boldsymbol{u}$ at some point in the time window,

$$q(t_b, \boldsymbol{x}) = q_0, \quad t_b \in [t_0, t_f],$$

$$\boldsymbol{u} = \boldsymbol{u}_0.$$

The effects of this choice are discussed in section 5.2.5. This method in fact represents a particular form of weak 4D-VAR, in which the form of the discrepancy term is forced to be equal to $\boldsymbol{u}_a \cdot \nabla q$ for some vector field, $\boldsymbol{u}_a$. In this way the method can be thought of as correcting for the adiabatic part of the model error term in the tracer equations, though not transferring any additional momentum to the system.

5.2.4   BASIS FUNCTIONS

The continuous formulation of the four dimensional displacement assimilation (4DDA) algorithm should in theory find instantaneous velocity fields of any magnitude (in the energy norm of Chapter 4) that are sufficiently continuous to satisfy the variational principle of section 5.2.3, so long as they act to extremize the cost function, $\mathcal{J}$. This means that the algorithm is in turn free to generate displacements of multiple different frequencies, including very high ones. This contradicts one of the original assumptions made in the formulation of the incompressible advection term approach, namely that all the assimilation displacements are large-scale and slow, at least in comparison to the standard model dynamics.

In fact in the continuous limit with temporally discrete observations [4] and vanishing diabatic forcings there will exist multiple different choices for the assimilation velocity which will all produce the same final fit to the available observations, but of varying energies. The locally most energetic states correspond in the limit to repeated application of the three dimensional method of Chapter 4, assimilating only at those time levels when observations are available. In this sense the four dimensional algorithm acting in the 'no model' case, when both the forward model velocity and diabatic forcings vanish, and with observations available only at the end of the time window represents a Benemou & Brenier space-time style approach to the case of the three dimensional assimilation problem.

To a certain extent the discontinuity problem vanishes when the problem is discretized for numerical calculation, since the representation then has no dichotomy between a continuous model and discrete data, while the numerical scheme employed makes some sort of assumption about the continuity of the imposed fields.

---

[4]That is the point observations in time, as well as space, as in situ observations often are.

Nevertheless the issues surrounding the time-level of the assimilation velocity do not disappear. Instead a new difficulty arises concerning the dimension of the fields to be stored. Under the time window approach the assimilation velocity field to be applied must be available over the entire time window, but cannot be calculated in the forward model step. Instead the result of the adjoint model must be stored at each intervening time level.

This data storage issue may rapidly exceed available computer capacity, especially in the high resolution, short time-step models for which the adiabatic assimilation routines are most appropriate. In the case of the two layer QG model a single assimilation run at 10km resolution and a one hour time-step to assimilate the data available over a 10 day time-window requires storage for an array of approximately 2^{24} floating point numbers for both the current assimilation streamfunction and its search direction. This represents an over a hundredfold increase on the 2^{16} elements required in storage for the initial value increments generated from the weak 4D-VAR formulation.

Fortunately, a method exists which obviates both the theoretical inconsistency of the continuous model and the practical handicap of the numerical method. Regardless of the physical model chosen the momentum equations in the extended cost function contain no dependence on the rate of change of the assimilation streamfunction with time, i.e. that terms in $\partial \psi^a / \partial t$ never appear in $\mathcal{M}$. If we consider, for example, a Fourier expansion of the assimilation streamfunction,

$$\psi^a(t, \boldsymbol{x}) = \frac{1}{2\pi} \int\limits_{-\infty}^{\infty} \hat{\psi}^a(k, \boldsymbol{x}) e^{ikt} dk. \tag{5.6}$$

It will be noticed that the new definition of the cost function, $\mathcal{J} = \mathcal{J}(\hat{\psi}^a)$, is purely linear in terms of the Fourier modes of the streamfunction, $\hat{\psi}$, and, by considering a variational principle for those terms and reversing the order of the integrations in t and s, we obtain the new search direction,

$$\nabla_{\hat{\psi}^a} \mathcal{J} = \int_{t_0}^{t_f} \left[-\frac{\partial(\theta, \zeta)}{\partial(x, y)} - \beta \frac{\partial \theta}{\partial x} \right] e^{ikt} \, dt.$$

This linear dependence on k implies that the gradients of the different modes can be calculated separately. Moreover, the cost function may be further extended by adding an explicit term penalizing the energy contained in the various frequencies,

$$\mathcal{J}_k = \frac{1}{2} \int_\Omega \int_{-\infty}^{\infty} F(k) \nabla \hat{\psi}^a \cdot \nabla \psi_a \, dk \, dS.$$

The magnitude of the weighting term $F(k)$ then determines the relative magnitude to which modes at different frequencies are allowed to grow. The most severe possible function is a strict cut-off above some critical frequency, K,

$$F(k) = \begin{cases} 0 & |k| \leq K \\ \infty & |k| > K. \end{cases} \tag{5.7}$$

This prohibits the existence of valid solutions with energies at frequencies above K. Alternatively some other monotonic function in $|k|$ could be envisioned, with the growth of some modes unconstrained, some penalised and some forbidden so that

$$F(k) = \begin{cases} 0 & |k| \leq K_1 \\ g & K_1 < |k| < K_2 \\ \infty & |k| > K_2 \end{cases} \tag{5.8}$$

with g an increasing function of $|k|$ which satisfies

$$g(K_1) = 0, \quad g(K_2) = \infty.$$

If g is chosen so that F is smooth in K, then this would be expected to generate solutions which demonstrate less sinusoidal behaviour near high frequency signals.

Under discretization the Fourier integrals become summations. Expressing these for some generalized basis functions, $f_j(t)$, by means of the equation

$$\psi^a(t, \boldsymbol{x}) = \sum_{j=1}^{\infty} f_j(t) \Psi_j^a(\boldsymbol{x}),$$

we see that the spatial gradients of the cost function are calculated by the discretization of

$$\nabla_{\Psi_j^a} \mathcal{J} = \int_{t_0}^{t_f} \left[-\frac{\partial(\theta, \zeta)}{\partial(x, y)} - \beta \frac{\partial \theta}{\partial x} \right] f_j(t) \, dt.$$

The constraint (5.7) is equivalent to truncating the set of basis functions to operate only at certain frequencies,

$$\psi^a(t, \boldsymbol{x}) = \sum_{j=1}^{K} f_j(t) \Psi_j^a(\boldsymbol{x}),$$

and then solving using the general method. Under the alternative constraint (5.8)

then the basis functions are again truncated,

$$\psi^a(t, \boldsymbol{x}) = \sum_{j=1}^{K_1} f_j(t) \Psi_j^a(\boldsymbol{x}) + \sum_{k=K_1}^{K_2} f_j(t) \tilde{\Psi}_j^a(\boldsymbol{x}).$$

The gradient of the cost function with respect to the penalized modes, $\tilde{\Psi}^a$, is then

$$\nabla_{\tilde{\Psi}_j^a} \mathcal{J} = \int_{t_0}^{t_f} \left[-\frac{\partial(\theta, \zeta)}{\partial(x, y)} - \beta \frac{\partial \theta}{\partial x} \right] f(t) \, dt + g_j \nabla^2 \Psi_j^a. \tag{5.9}$$

Since most minimization schemes which use gradient information work on the basis that the new guess Ψ_{n+1} is given by an equation similar to

$$\Psi_{n+1} = \Psi_n - \alpha \nabla_{\Psi_n} \mathcal{J},$$

the extra linear term in (5.9) acts like a linear friction term, damping the energy contained in all those modes which are not excited by the forcing from the Jacobian term between the adjoint variable and vorticity.

5.2.5   TIME PIVOTING AND INITIAL CONDITION ERRORS

As discussed above the solution algorithm requires a choice of initial condition to be defined at some point in the time window, from which the system may be integrated, forwards or backwards. Typically, due to the numerical constraints of the models which are used[5], this point will be the starting time of the window, t_0. Since the assimilation forcing term appears in the evolution equation for the tracer rather than as a step change in its value the variable is constrained locally to remain close to the background trajectory (see Figure 5.4). This means that in cases of large initial condition error the method *cannot* correct the error fully, even if totally adiabatic and will instead tend to create large, numerically unstable, signals around the initial part of the window. This issue can be avoided by coupling the 4DDA method to a 3D method, which is run to correct initial condition errors before the time window is fitted, or by coupling to another 4D method which corrects initial condition errors, such as the 4D-VAR method itself.

Coupling the method directly, with a 4D-VAR algorithm through the use of twin control variables, $\boldsymbol{x}_0$ and $\boldsymbol{u}_a$ represents the approach of Hoffman and Grassotti (1996) and Ravela et al. (2004) in four dimensions. As such, it requires assumptions about the magnitudes of the phase and amplitude errors in the data, in the

---

[5]The forward model is typically only stable integrated forwards, due to diffusion terms. Similarly the adjoint model may be unstable forwards in time, but stable when integrating backwards.

case when both approaches can explain the observed discrepancy. Provided the background conditions are close, in the sense of the rearrangement metrics introduced previously, to the true system state, then this should not be a great issue. When the background and truth differ greatly then the original assumptions used in justifying the use of rearrangement theory are broken and it is justifiable in using methods which only correct for amplitude error.

5.3   ADJOINT MODELS

5.3.1   DISCRETE VERSUS CONTINUOUS ADJOINTS

There currently exist two extremely distinct philosophies among modellers for dealing with the question of numerical calculation of the adjoints to systems based on geophysical fluid flow. The first approach asserts that the continuous equations must be discretized in order to be solved, so that it is the adjoint to the discretization to the tangent linear model, M which must be calculated. This can be achieved by replacing line by line the linearized forward operators by their adjoint equivalent, for example transposing the arrays used in matrix multiplication. This is naturally the approach followed by the various modern automated systems for calculating adjoint codes, such as that of the Tangent linear and Adjoint Model Compiler, or TAMC (Giering and Kaminski).

The major advantage of this method is that the final gradient thus calculated is the best possible representation of the rate of change of the implemented numerical model around the linearization state. Provided the linearization is a good representation of the non-linear dynamics then this maximizes the magnitude of the descent possible on each iteration of the descent step. This approach also allows for the transparent treatment of discontinuous behaviour caused by conditional statements and switches within the forward model, see for example the discussion in Akella and Navon (2006).

The second approach is to consider the implementation of the tangent linear model as an approximation to the genuine continuous tangent linear model which, provided sensible numerics are used, will be re-obtained in the limit of infinitely fine resolution. In this case the discrete adjoint is calculated by first generating the continuous adjoint of the continuous linear model and then discretizing the result using some appropriate technique. For simple model dynamics or small problems this produces easily comprehensible and efficient code and allows the use of differing approximations in the adjoint model from those that are used in the forward
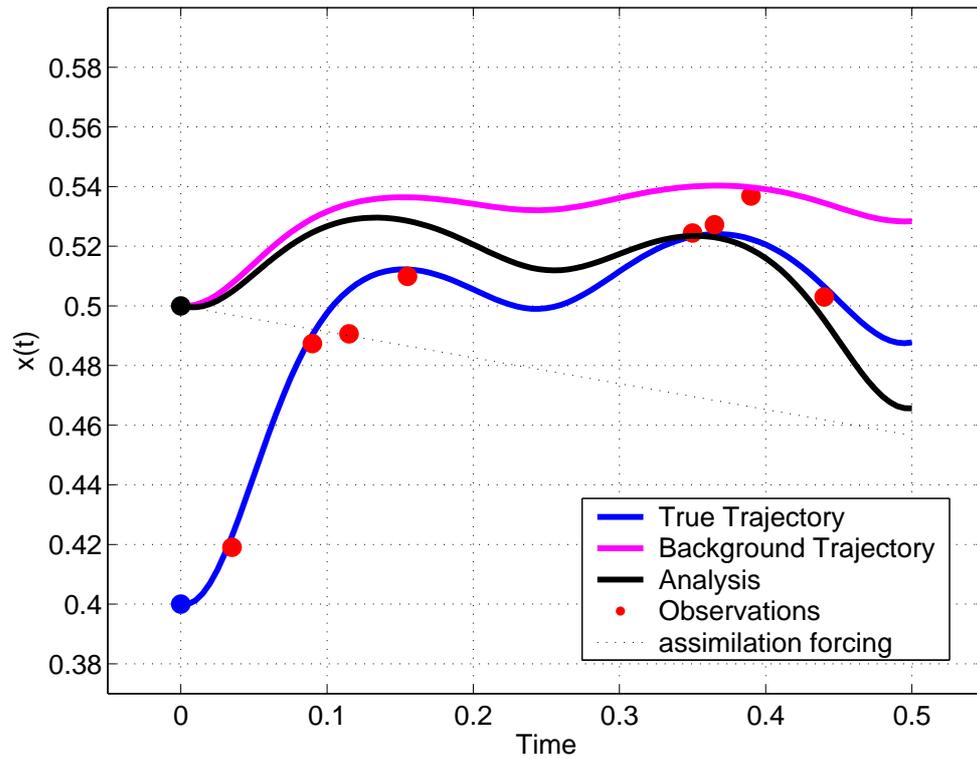
Figure 5.4: Cartoon showing effect of 4DDA on model trajectories

Data is assimilated into a simple one-dimensional model by modelling the evolution equation, in this case through a simple additional linear forcing,

$$\frac{d\boldsymbol{x}}{dt} = \mathcal{M}(\boldsymbol{x}, t) + \alpha,$$

with α chosen to minimize the r.m.s. error compared to observations. It will be noted that the correction due to forcing is most accurate in the centre of the window. The lack of correction near the beginning of the time-window is typical of approaches based on modifying solely the evolution equations.

model when appropriate. More iterations of the chosen descent algorithm may be required than in cases using the first technique to achieve a given reduction in cost function, but with each iteration cheaper to perform the final result may be computationally more efficient. A study by Schiller (1995) found that the accuracy of calculated gradients was relatively unimportant when attempting state estimation in the North Atlantic using the steady state version of an ocean GCM. Sirkes and Tziperman (1997) find that, while both approaches are appropriate for the generation of gradient information, the "adjoint of finite differences" approach shows a large, strongly non-physical computational mode over intermediate timesteps.

In what follows it is thus the "finite difference of adjoint" approach which is taken. The precise numerical implementation of the adjoint models used here are described in detail in section C of the appendix. However for comprehension it should be noted that the code steps four adjoint model variables, $\theta_1, \theta_2, \phi_1, \phi_2$ backwards in time from an initial zero state. These variables correspond to the four Lagrange multipliers applied to maintain the evolution of the forward model quantities, $\zeta_1, \zeta_2, \psi_1, \psi_2$ in the QG two layer equations,

$$
\begin{aligned}
\theta_1 &: \quad \frac{d_1^a}{dt}[\xi_1 + \beta y] - \frac{1}{\rho_0 H_1}\nabla \times \tau - \nu\nabla^4\psi_1 = 0, \\
\theta_2 &: \quad \frac{d_2^a}{dt}[\xi_2 + \beta y] - \nu\nabla^4\psi_2 = 0, \\
\phi_1 &: \quad \xi_1 - \nabla^2\psi_1 + \frac{f_0^2}{g'H_1}(\psi_1 - \psi_2) = 0, \\
\phi_1 &: \quad \xi_2 - \nabla^2\psi_2 + \frac{f_0^2}{g'H_2}(\psi_2 - \psi_1) = 0,
\end{aligned}
$$

as described in Chapter 3. The resulting adjoint optimality equations are

$$
\begin{aligned}
\xi_1 &: \quad -\frac{d_1^a\theta_1}{dt} + \phi_1 + H_1(y - h[\xi_1, \xi_2]) = 0, \\
\xi_2 &: \quad -\frac{d_2^a\theta_2}{dt} + \phi_2 + H_1(y - h[\xi_1, \xi_2]) = 0, \\
\psi_1 &: \quad -\nabla^2\phi_1 + \frac{f_0^2}{g'H_1}\phi_1 - \frac{f_0^2}{g'H_2}\phi_2 + \frac{\partial(\theta_1, \xi_1 + \beta y)}{\partial(x, y)} + \nu\nabla^4\theta_1 - a_1, \\
\psi_2 &: \quad -\nabla^2\phi_2 + \frac{f_o^2}{g'H_2}\phi_2 - \frac{f_0^2}{g'H_1}\phi_1 + \frac{\partial(\theta_2, \xi_2 + \beta y)}{\partial(x, y)} + \nu\nabla^4\theta_2 - a_2,
\end{aligned}
$$

plus boundary conditions, discussed in more detail in the Appendix, Section C. Exactly as with the forward model these form a pair of evolution equations and a pair of elliptic equations and similarly these four equations in fact represent the evolution of two variables, nominally the adjoint vorticity variables, θ_1, θ_2. The solution algorithm is also very similar to that for the forward model. The adjoint vorticities are back-stepped using spatially and temporally centrally differenced fi-

nite difference methods, making use of a co-ordinate transformation to solve the elliptic equations for two uncoupled modes, analogous to the barotropic and baroclinic modes of the two layer forward model. The linearization state in the forward variables at each time level is calculated using linear interpolation between values stored at regular intervals larger than the timestep of either the forward or adjoint models (one hour in each case) to save on storage requirements for data, which can otherwise prove excessive, while avoiding the necessity of recalculating intermediate values using the forward model. This represents another major approximation in the system, particularly when the time interval between storage steps is long. Experimentally the effect on accuracy of this approximation over shorter storage intervals of 4–6 hours seems small however, see Section 6.1.

5.3.2   The Observation Operator

The forcing provided by the observations depends on what variables are observed and on how the measurements are chosen to represent data within the model[6]. The observation method that produces the simplest problem is to assume that observations are available of the vorticity field over a set of continuous connected domains at a scale larger than the discretized model resolution. The observation operator can then be represented in continuous space as the identity operator over the observation domain. This means the observation penalty function of equation(C.4) can be replaced by an identically valued function

$$\mathcal{J}_{\xi \text{obs}} = \frac{1}{2} \sum_{i=1,2} \int_{t_0}^{t_f} \int_{\Omega} H_i[\boldsymbol{r},t] \left(y_i^{\text{ext}} - \xi_i\right)^2 dS dt \tag{5.10}$$

where $H[\boldsymbol{r},t]$ is a step function with

$$H_i[\boldsymbol{r},t] = \begin{cases} 1 & \text{observations of } \xi_i \text{ are available} \\ 0 & \text{no observations of } \xi_i \text{ are available} \end{cases} \tag{5.11}$$

and $y^{\text{ext}}[\boldsymbol{r},t]$ is any extension of the field of observations to the entire domain, i.e. any function that satisfies the condition

$$y^{\text{ext}}[\boldsymbol{r},t] = y[\boldsymbol{r},t] \; \forall \boldsymbol{r} \in \Omega_{\text{obs}} \tag{5.12}$$

---

[6]A measurement in a grid box may be assumed to represent the average value of the variable across the whole box, or the value of a interpolating values at the grid points.

Under these conditions the evolution equations for θ_1, θ_2 become

$$\frac{d_1^a \theta_1}{dt} = -\phi_1 + H_1 \left(y - h[\xi_1, \xi_2] \right),$$

$$\frac{d_2^a \theta_2}{dt} = -\phi_2 + H_2 \left(y - h[\xi_1, \xi_2] \right).$$

An obvious alternative linear function that may be observed is the streamfunction. The top level streamfunction under the model assumptions, specifically the rigid lid assumption, is almost exactly equivalent to the observed sea surface height of the real ocean multiplied by a constant conversion factor. This variable may observed by satellite altimetry as described in Section 1.4.3. Under the same assumptions as for vorticity this means that the observation may be written as

$$\mathcal{J}_{\xi\text{obs}} = \frac{1}{2} \sum_{i=1,2} \int_{t_0}^{t_f} \int_\Omega H_i[\boldsymbol{r}, t] \left(y_i^{\text{ext}} - \psi_i \right)^2 dSdt \qquad (5.13)$$

this gives a revised elliptic set of equations,

$$\nabla^2 \phi_1 - \frac{1}{L_1^2} \phi_1 + \frac{1}{L_2^2} \phi_2 = \frac{\partial(\theta_1, \xi_1)}{\partial(x, y)} + \nu \nabla^4 d_1 - a_1 + H_1 \left(y_i^{\text{ext}} - \psi_1 \right),$$

$$\nabla^2 \phi_2 + \frac{1}{L_1^2} \phi_1 - \frac{1}{L_2^2} \phi_2 = \frac{\partial(\theta_2, \xi_2)}{\partial(x, y)} + \nu \nabla^4 d_2 - a_2 + H_2 \left(y_i^{\text{ext}} - \psi_2 \right).$$

since the observation operator appears only acting on the innovation vector, these may be solved in the same way as previous sections.

5.4   CHAPTER SUMMARY

Due to the limited ocean data available over short validation windows techniques for spreading data in time are required. Previously such techniques have not been combined with the rearrangement type theory due to possible issues with validation times. Here one method has been documented for applying such a scheme in model time by modifying the original model equations to include an assimilation term. A 4D-VAR style adjoint method is used to choose this term optimally so as to minimise an observation cost-function.

CHAPTER 6

TIME WINDOW EXPERIMENTS

"All life is an experiment. The more experiments you make the better."

Ralph Waldo Emerson
(Attributed)

The behaviour of the four dimensional displacement assimilation scheme is investigated by running it over a series of test cases and comparing these with the results of a traditional 4D-VAR algorithm . Firstly the numerical implementation is investigated by using a series of adjoint tests to test the scheme for basic convergence properties and to check that the various approximations made hold.

6.1   ADJOINT MODEL VERIFICATION TESTS

As discussed in Chapter 5, at the heart of both the 4D-VAR method and the new four dimensional displacement assimilation scheme lies a model adjoint code used to calculate a search path for the gradient descent routine used to minimise the target cost function. To test whether the two adjoint models, coded as in the Appendix C, successfully generated (an approximation to) the gradient of the discretized cost function with respect to changes in initial conditions (for the standard 4D-VAR routine) or assimilation streamfunction (for the displacement method technique) a suite of tests were carried out using a simple finite difference gradient approximation.

6.1.1   THE GRADIENT TEST

Consider the general Taylor series expansion of the perturbation of the observation part of the cost functional,

$$\mathcal{J}(\xi) = \frac{1}{2} \int_{\mathcal{D}} (\xi_o - \xi)^2 \, dV,$$

about a state ξ due to a perturbation ϵg, where ϵ is a small parameter and g is an arbitrary (steady) field on $\mathcal{D}$. The leading order terms in the expansion in ϵ are

$$\mathcal{J}(\xi + \epsilon g) = \mathcal{J}(\xi) + \epsilon \int_{\mathcal{D}} \nabla_\xi \mathcal{J} g dV + \mathcal{O}(\epsilon^2),$$

where $\nabla_\xi \mathcal{J}$ denotes the local variation of the functional with respect to variation in ξ, equivalent to the 4D-VAR gradient of the cost functional,

$$\nabla_\xi \mathcal{J} = \frac{\partial}{\partial \xi} \frac{(\xi_o - \xi)^2}{2} = -(\xi_o - \xi).$$

If the arbitrary field is now chosen so that $g = \nabla_\xi \mathcal{J}$ then it follows that the functional, $\mathcal{G}$, defined by

$$
\begin{aligned}
\mathcal{G}(\epsilon) \quad &:= \frac{\mathcal{J}(\xi + \epsilon g) - \mathcal{J}(\xi)}{\epsilon \displaystyle\int_{\mathcal{D}} g^2 dV}, \\
&= \frac{\displaystyle\int_{\mathcal{D}} \nabla_\xi \mathcal{J} g dV}{\displaystyle\int_{\mathcal{D}} g^2 dV} + \mathcal{O}(\epsilon),
\end{aligned}
\tag{6.1}
$$

will converge to unity as ϵ vanishes. More generally any approximation to the first order variation of a cost function with respect to a given parameter may be tested by plotting $\mathcal{G}$ against $\log(\epsilon)$ and checking for convergence. A logarithmic axis is chosen for ϵ since it clearly demonstrates the behaviour over several orders of magnitude. For any numerical implementation of this test there are typically four phases of behaviour:

(i). For very large ϵ the leading order truncation of the Taylor series is a very poor approximation. The behaviour of the function appears essentially random, governed by the coefficients of the higher order terms in the expansion.

(ii). As ϵ vanishes the function rapidly converges towards its limiting value. Unless the adjoint model is perfect and the forward model is linear, discretized over only a single time-step, or the length of the model time step is allowed to tend to zero, then this will not be precisely unity, but should lie close to it, with methods which give a smaller difference preferred.

(iii). There will be an interval of ϵ over which the curve is essentially flat. This is the range of linearization over which the numerical approximation to the gradient of the cost functional has converged towards the analytic solution. Ideally the range of this plateau will be large, allowing for large stepping in the gradient sliding method to be used to minimize $\mathcal{J}$.

(iv). Finally, once ϵ drops below a critical value dependent on the model the curve will produce another series of essentially random oscillations about unity. The critical value depends on the choice of numerics and the floating point representations used in the underlying computer architecture. This fluctuation is due to the accumulation of truncation errors in the forward and adjoint models, as well as in the calculation of $\mathcal{G}$ itself.

Plotting the relevant functions for the 4D-VAR and displacement adjoints to the QG model for various choices of data and parameters produces the graphs shown in Figure 6.2.

6.1.2 RESULTS FOR 4D-VAR METHOD

Application of the gradient test to the 4D-VAR code embedded within the 4DDA method showed accurate generation of the correct gradient with respect to variations in initial condition over almost five orders of magnitude, with typical errors in $\mathcal{G}(\epsilon)$ of order 0.01 relative to unity. The range of ϵ demonstrating convergence was $-\log_{10}\epsilon = 7.5 - -12.5$, with divergence of the gradient, shown by a tendency towards large negative values of $\mathcal{G}(\epsilon)$ for larger values of ϵ and the oscillations typical of numerical truncation errors for smaller values.

Compared with the behaviour found for 4D-VAR codes in operational numerical weather prediction models this represents a relatively small range of convergence, a factor of the use of discretized versions of the continuous adjoint rather than the true discrete adjoint code, coupled with the relatively noisy linear interpolations used to find values for the forward model vorticity in between stored values. Since however the code was developed towards the 4DDA method and was only to be used for generating an initial search direction it was felt that this deficit would be of little relevance.

6.1.3 RESULTS FOR 4DDA METHOD

Compared with the results for the 4D-VAR method the 4DDA gradient tests show a larger domain of convergence, but worse achievement of unity, with convergence for $-\log_{10}\epsilon = 6-13$, but a consistent low bias of 0.1 in the value of $\mathcal{G}(\epsilon)$. Since the 4D-VAR code appears relatively accurate we ascribe this to inaccuracies in the temporal part of the problem, possibly in part the basis function representation of the continuous gradient, which has not been optimised, but is instead chosen for ease of application. It is possible it may also partly represent the larger dimension
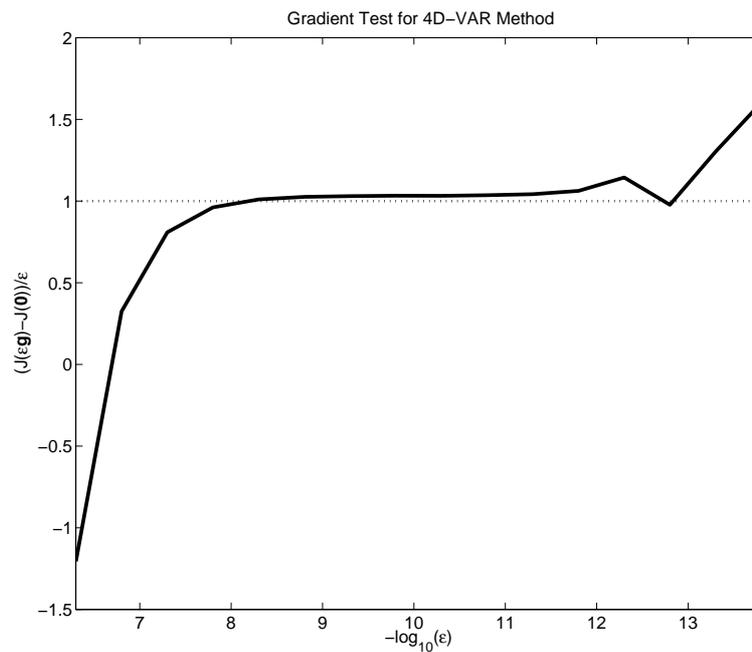
Figure 6.1: Gradient Test for 4D-VAR adjoint code
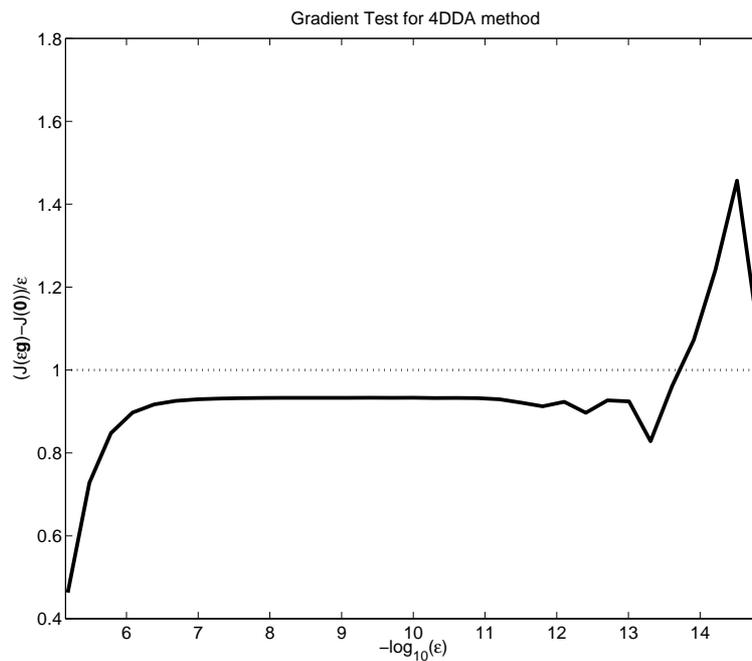
Figure 6.2: Gradient Test for 4DDA adjoint code

The figures show the result of plotting the finite-difference gradient test $\mathcal{G}$ given by (6.1) against step size, ϵ, for the 4D-VAR and 4DDA methods. Both plots show the typical band of convergence bounded for small ϵ by truncation error and for large ϵ by the behaviour of higher order terms. Note that the range over which $\mathcal{G}$ is is constant is actually larger for the 4DDA code, although the method only generates an approximate gradient and thus does not achieve unity.

of the 4DDA problem and thus the increase in the number of small truncation errors in the numerics, which accumulate under the summations implicit in the code integration. Assimilation experiments later in this section show the method is still capable of generating a useful descent direction, and shows not clear bias in the form of the solution produced, hence we accept the code as practicable, while noting that any inefficiency in the 4DDA descent routine as compared to the 4D-VAR method may be partly ascribed to this cause.

6.1.4 RELATIVE COMPUTATIONAL COSTS

As in Section 4.2.2 we briefly discuss the cost of the new 4DDA method in terms of physical computational resources and CPU time. We choose for comparison 4D-VAR, as the obvious choice of another method using data at multiple times. We also briefly discuss the hypothetical method introduced in Section 5.2.1.

As was shown in the development of the 4DDA method in Chapter 5, the 4DDA algorithm generates both the assimilation streamfunction and the 4D-VAR gradient, with only one additional matrix calculation required at each model timestep during the adjoint run, of a cost equivalent to one application of a derivative operator in the forward model. As such, and since there will be many such operations in a single step of the adjoint code, the two methods are effectively equal in computational cost for a single calculation of the descent direction. The hypothetical quasi-3D rearrangement code, however requires one full run of one of the 3D methods of Chapter 4. For a typical situation, in which the number of assimilation steps was of order the number of model timesteps in the assimilation timewindow, this would effectively double the CPU time required in a single iteration of the assimilation code.

In terms of memory usage, the requirement to store the increments from the assimilation streamfunction on multiple timelevels leads to potentially much greater requirements for the 4DDA, than for traditional variational assimilation methods. For a ten day timewindow with a model timestep of an hour, this would be a massive 240-fold increase in memory usage, compared with the requirement for univariate 4D-VAR data assimilation. Following the Fourier series reduction method of Section 5.2.4 we reduce this to a five-fold increase, when storing data on the constant mode and the first two leading order sine and cosine coefficients in the Fourier series representation. This makes the code roughly as memory efficient as 4D-VAR for multivariate assimilation, assuming the number of physical quantities in the model is likely to be of order ten. Since the hypothetical quasi-3D assimilation code requires both the increment retrieved from the adjoint code and the result of the 3D

displacement step, it would represent a tripling of the memory requirements of the forward model.

Since the two methods, 4DDA and 4D-VAR are similar in computing requirements for the process of generating a single line search along a descent direction, we will use the iteration count of these searches in examining the efficiency of the two methods in the test-cases which follow. The codes are run until the length of the normalized line search, drops below a level, 1×10^8, where we expect the numerical truncation errors to dominate.

6.2   TRACER FIELD EXPERIMENTS

Identical twin experiments were performed using both the 4D-VAR and 4DDA methods to attempt to assimilate errors due to

- Adiabatic model errors,

- Adiabatic errors in initial conditions,

- Non-adiabatic (amplitude) errors in initial conditions,

- Different model resolution (model errors).

6.2.1   A TRIVIAL EXAMPLE : PSEUDO-CURRENTS

The most trivial test case for inter-comparison of the 4DDA and 4D-VAR methods is to assimilate full data from a run which has been forced by a similar pseudo-wind field to that which the 4D-VAR method introduces, that is to say we consider as the true state the entire data field from a run of a model, with equations

$$\frac{dq_i}{dt} + \boldsymbol{u}_f \cdot \nabla q_i = F_i.$$

The q_i, $i = 1, 2$ are the relevant layer vorticities, $\boldsymbol{u}_f$ is the applied pseudo wind forcing and d/dt and F_i denote the standard model total advection operator and forcings, so that the standard model is denoted in this shorthand notation by

$$\frac{dq_i}{dt} := \frac{\partial q_i}{\partial t} + \boldsymbol{u}_i \cdot \nabla q_i = F_i.$$

This is the problem to which the adiabatic assimilation method is naturally most suited, and the 4D-VAR method least, in that we are actually attempting to retrieve a model error parameterized as an advection, violating some of the assumptions

for success of the 4D-VAR algorithm. For such problems the 4D-VAR algorithm will produce an analysis state for the initial condition at the start of the time window that differs from the truth, while 4DDA will attempt to converge towards the model bias.

For the results presented here the forcing field was determined by a relation,

$$\boldsymbol{u}_f = \nabla \times \psi_f \boldsymbol{k}, \ \boldsymbol{k} = (0, 0, 1)^T,$$

so that the additional current was given by a forcing stream function defined as

$$\psi_f(x, y)_i = -A \left(\tfrac{r^3}{3} - B\tfrac{r^2}{2} + \tfrac{B^3}{6} \right) \mathcal{H} \left[r^2 - B^2 \right], \quad i = 1$$
$$\psi_f(x, y)_i = 0 \qquad\qquad\qquad i = 2,$$

where $r^2(x, y) = ((x - x_c)^2 + (y - y_c)^2)$, with $\mathcal{H}$ the Heaviside function as before and the parameters chosen as $x_c = y_c = 1280$km, $B = 250$km^2 and $A = 1 \times 10^{-12}$m^{-1}s^{-1}. The resulting 'hump' streamfunction generates rotation around the point (x_c, y_c) in the region $r \le B$ without causing discontinuities at the edge or centre of the rotation.

The streamfunction is plot (1) of Figure 6.3, with the second plot showing the mean top vorticity layer vorticity field in the unforced control field used to provide the background first guess of the model trajectory, and the third and fourth show respectively the difference between the forced 'true' and unforced 'background' top layer vorticity fields at the end of the time window in the experiment, $T = 24$ hours, and the variance in the same field over the entire window. Both these plots show that the effect of the forcing is predominantly concentrated over the area of forcing, $r < B$ and the variance field shows that the dominant effect is in the position of the divide between the high and low vorticity gyres.

Observations at 6 hourly intervals of the resultant vorticity field over the time window of one day were then presented to both the implementations of the 4DDA and 4D-VAR methods. The goal was to achieve the best possible fit to the model trajectory, as measured by the (observation) cost function

$$\mathcal{J}_o = \sum_{t \in T} \int_{\mathcal{D}_H} (q_1^{\text{obs}}[t] - q_1^{\text{back}}[t])^2 + (q_2^{\text{obs}}[t] - q_2^{\text{back}}[t])^2 dS, \ T = [0, 6, 12, 18, 24] \text{ hours.}$$

The growth with time in the individual terms in this summation is shown as plot (1) of Figure 6.4. As may be expected the graph initially shows exponential growth as small errors are amplified. However, the sharp downturn at the end of the window shows that the time window is long enough that small errors have saturated and

are now interacting with the large-scale flow.

As may be expected, given the nature of the error being assimilated (model error directly ascribable to an advective term), the 4DDA method is much the better of the two methods for retrieving the original solution. After only one iteration the method returns an assimilation streamfunction which represents a 70% reduction in $\mathcal{J}_o$. Further gains are however small (see plot (3) in Figure 6.4). In contrast the 4D-VAR Algorithm stagnates immediately, with a neglible reduction in the cost function ($\sim 1 \times 10^{-4}\%$), see plot (4) of Figure 6.4.

The returned assimilation streamfunction (plots (5)–(6) of Figure 6.3) is noticeably different from the forcing applied; there is a tripole rather than mono-pole structure to the function in the top layer, which appears on all four frequencies for which basis functions are prescribed on this run, as well as an induced dipole signal in the bottom layer, which is unforced in the run designated as truth. This demonstrates both the non-linear nature of the induced errors, which induces the bottom layer signal, and the non-uniqueness inherent in the problem, since neither of these two significant differences in streamfunction severely affects the initial reduction in cost function. The induced 4D-VAR increment to the streamfunction also has a multipole structure, but here it appears to contain a sharp wavelike signal (plot (2) in Figure 6.4). The wave structure arises from the inversion of the elliptic operator in the adjoint model. This form of increment is hard to correlate with either the induced forcing or the resultant anomaly. This illustrates how unsuited the strong formulation 4D-VAR system is to problems with an incorrect model, where the assumptions of the assimilation method are invalid.

As stated at the beginning of this section, this test case represents the most trivial problem to which this method may be applied; a well observed system acting under a forcing which is perfectly represented by an advection term. The shortcomings of the method in this situation, failing to retrieve the forcing perfectly and generating signals in locations and on frequencies which are unforced, represent an important limit on the skill of the method in the more realistic cases which follow. In this respect this is a recursive realization of the ideal twin philosophy discussed in Section 2.2. We investigate this further by examining the behaviour of the algorithm over a series of similar test problems in which a forcing streamfunction of a similar form, but of varying amplitude, are used to generate observation states. More precisely we set $A = \epsilon \times 10^{-12}\mathrm{m}^{-1}\mathrm{s}^{-1}$, with ϵ a scaling parameter.

Plotting the final fractional reduction in cost-function the method achieves against the magnitude of the forcing streamfunction produces plot (1) in Figure 6.5. We see that for large amplitude signals the reduction is marginal, while it reaches levels of an 80% reduction, as ϵ is reduced and the effect of the forcing becomes more

linear, in terms of the form of the displacements imparted to features. Once ϵ drops below a critical level, however the method no longer converges and no reduction is possible, as numerical errors once again dominate. This behaviour is familiar from the experiments with very close displacements shown in Section 4.2.4 and Figure 4.6 in particular. Examining the form of the returned assimilation streamfunction at the value of ϵ which maximizes the costfunction reduction (plots (2) – (3) in Figure 6.5) we see that the tripole structure shown in plot (5) of Figure 6.3 is retained in the top layer streamfunction, with only a slight increase in the strength of the centre pole suggesting a better fit to the imposed forcing. In the bottom layer streamfunction however the magnitude of the observed signal is much reduced, and a new 4-pole structure is seen. This is a visible representation of the simplication of the assimilation problem with smaller magnitude forcing. The existance of this window over which the 4DDA method works well to assimilate the trivial forcing problem gives jusification for testing the behaviour in the more complicated assimilation problems to which we will now apply it.

6.2.2 ADIABATIC INITIAL CONDITION ERROR

The second test case considered is the one in which the perfect model assumption is assumed to hold, but the true run differs from the background run in the initial conditions used, which are assumed to differ in such a manner as to approximately satisfy the rearrangement condition on all material tracer variables.

This difference between truth and models is generated using the methods discussed section 2.5, i.e. the (perfect) model itself is used to create data so that truth and background come from the same model trajectory but at different validation times (in fact, strictly speaking over different time windows, since we are now assimilating temporal information). This is a case to which the 4D-VAR algorithm is ideally suited, while it is one for which the 4DDA technique may easily fail completely, for the reasons discussed in section 5.2.5. The correct modification to initial conditions will give a perfect fit with the true model trajectory, but this modification may be poorly approximated by boosts to the model velocity field. This makes the results for the 4DDA method highly dependent on the form of the basis functions to be fitted, particularly the maximum frequency on which the assimilation velocity is allowed to vary, and on the relative magnitude of the phase shift between truth and background compared to the time window over which the observations are to be fitted. Counter-intuitively, small shifts in validation time may lead, over medium-length time windows, to the more difficult assimilation problems than larger shifts. Here difficulty of observation is measured in terms of the fractional reduction in
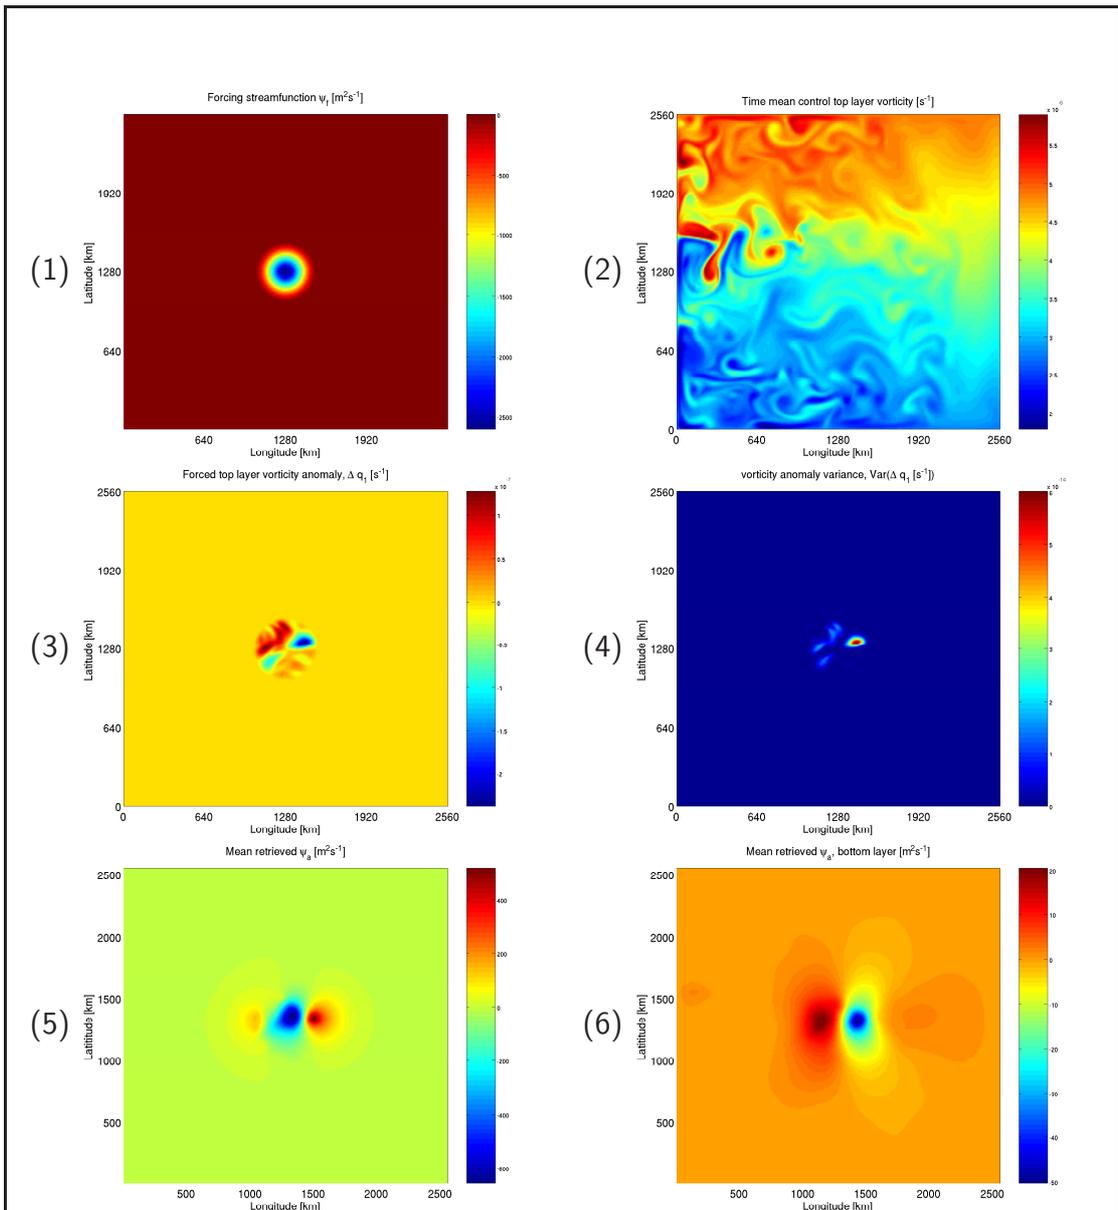
Figure 6.3: Results for trivial forcing problem

(1) Forcing field applied to generate discrepancy. (2) Mean top layer background vorticity. (3) Mean top layer anomaly with observations. (4) Variance of anomaly with observations. (5) Mean top layer assimilation streamfunction. (6) Mean bottom layer assimilation streamfunction.

If the scheme worked perfectly then (1) and (5) would be identical and (6) would be zero everywhere. The actual zonal pattern shown in (5) appears to be inherited from the vorticity field, (2) with the stronger eastern signal appearing in both the mean and variance of the anomaly (3)–(4).

Figure 6.4: Results for trivial forcing problem

(1) Growth with time of the discrepancy between the unforced (background) and forced (observed) fields. (2) Retrieved gradient/sensitivity streamfunction field for the 4D-VAR method. (3) Reduction in cost function under the assimilation 4DDA method, (4) Reduction in cost function under the 4D-VAR assimilation method.

Plot (1) shows that the growth in errors due to the forcing saturates in approximately 22 hours, very near the end of the time window. Comparing (2) to the mean and variance of the anomaly between the fields, plots (3)–(4) of Figure 6.3 shows a wave pattern, inherited from the inversion of the elliptic operator required in the adjoint equations, see Section 5.3. Plot (3) shows both that the 4DDA method rapidly finds a good approximation to the constant forcing field and that the code stagnates at a local minimum after one step. Plot (4) shows simply how unsuited the strong 4D-VAR method is to this type of forcing problem as the cost function refuses to descend.

Figure 6.5: Results for trivial forcing limit tests

(1) Analysis reduction in cost function against magnitude of forcing (2) Mean top layer assimilation streamfunction in most reduced limit. (3) Mean bottom layer assimilation streamfunction in most reduced limit.

(1) shows that there exists a middle-scale of forcing problems for which the 4DDA method can obtain a good fit to observations. For large magnitude and small magnitude problems the method fails for numerical reasons. The form of the returned streamfunction in the top layer, (2) remains fixed across this window (compare with Figure 6.3) while form of streamfunction in the bottom layer, (3) shows an increase in the number of extreme points, coupled to a reduction in the relative magnitude of the signal.

cost function achieved. This appears to be due precisely to the issues of divergence in model trajectories discussed in section 2.5, coupled to the small denominator in the function $\mathcal{J}_0/\mathcal{J}_O^0$ given by the initial difference from truth in this case. Over longer time windows there is less need found for very large initial assimilation velocities, cf. Figure 5.4 and these problems are alleviated.

The results of an example of this form of experiment is summarized in Figure 6.6. The first plot shows an example of the *initial* discrepancy in the vorticity fields, $q_{\text{background}}[t_o] - q_{\text{truth}}[t_0]$. Plot (2) shows the reduction in cost function, defined as in the previous section. In none of our experiments did the method achieve a 1% reduction in streamfunction, despite the presence of full observations. The third and fourth plots show respectively the assimilation streamfunctions obtained for the constant part of the basis functions. Both fields are heavily dominated by the location and direction of the jet separation, with the top-layer setting up an across jet flow and the lower layer a flow increasing the strength of the jet. Considering the effectiveness of the 3D schemes when applied to similar problems it appears the issue is with the non-linear nature of the ideal paths of fluid parcels which lead to the ideal displacements. This is supported by plot (5) in the figure, a histogram showing the integrated kinetic energy equivalent,

$$\mathcal{K} = \sum_{i=1}^{2} \int_{t_0}^{t_f} \int_{\mathcal{D}} \rho_i \nabla \psi_i^a dx dy dt,$$

stored in each of the three modes in question. These modes, a time-constant field, a field varying as a period one sine function over the course of the time window and a field varying like a period one cosine function, are the ones used as temporal basis functions in the assimilation experiment. While over two thirds of the energy lies in the constant mode the sine mode is non-negligible especially compared with the results in the experiments with diabatic error, below. This is an example of failure precisely analogous to Figure 5.4, where attempting to fit to observations across the time window leads to overshoots. There is significantly less energy in the cosine wave field, than in the sine wave even though the periods are the same. This suggests that it is the behaviour in the middle of the time window which is being controlled most closely. This again ties to the impression that the model trajectory has 'overshot' the ideal analysis path at the end of the window and is thus insensitive to variation at the ends of the time window.
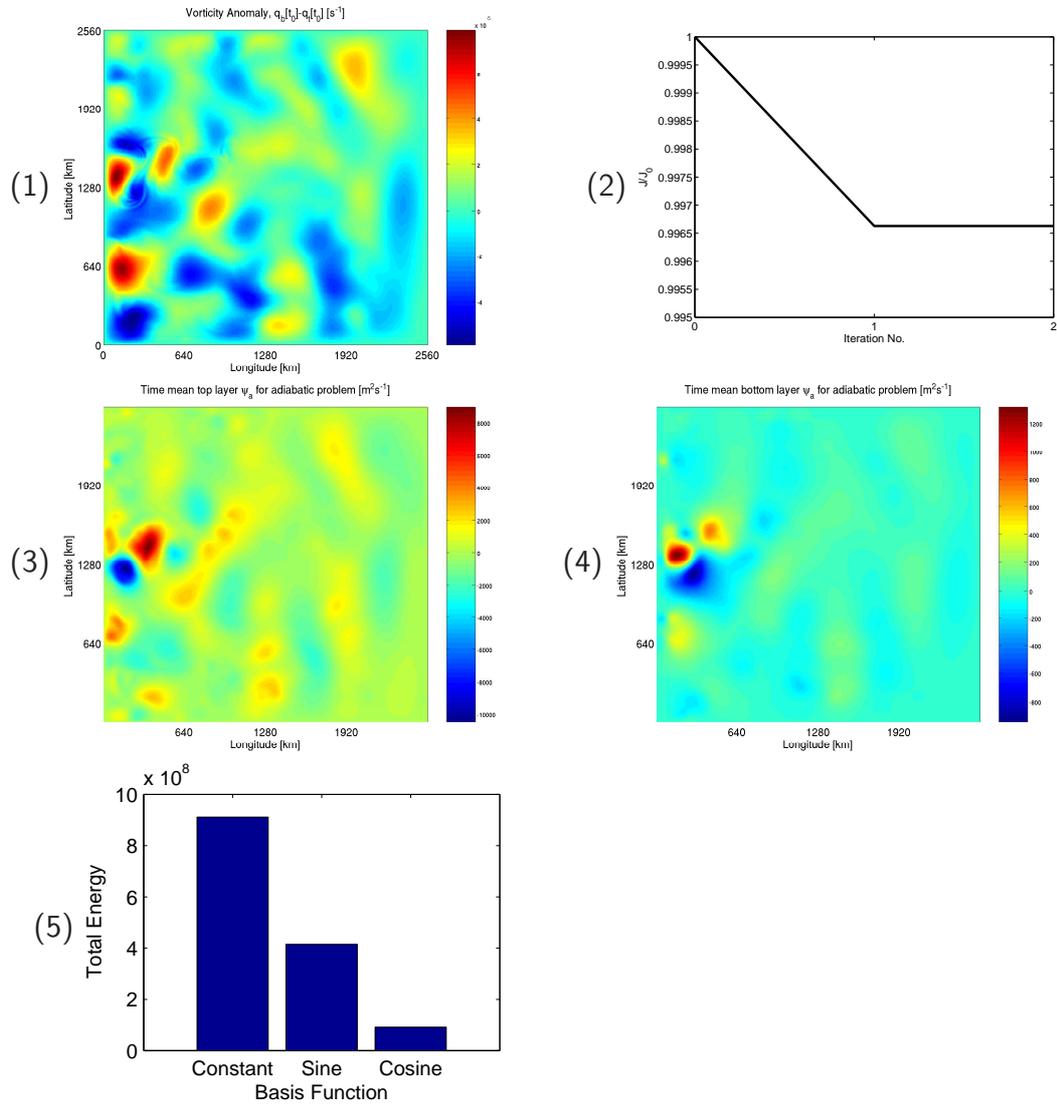
Figure 6.6: Results for adiabatic initial condition error

(1) Initial anomaly in the top layer vorticity field. (2) Reduction in cost function under the 4DDA method, (3) Mean top layer assimilation streamfunction. (4) Mean bottom layer assimilation streamfunction. (5) Energy in the various modes of the assimilation basis functions.

From the time stepping method used to generate the trajectories the anomaly shows the phase error in westward propagating waves. The streamfunctions returned, (3)–(4) are predominantly organized around the position of the jet on the western boundary with the energy almost exclusively in the time-constant mode (see (5)). This sensitivity to the jet means the cost function, (2), does not descend significantly.

6.2.3 DIABATIC INITIAL CONDITION ERROR

The third case considered consists of true and background runs from differing initial conditions with a large diabatic component to the error. This was achieved by using as initial conditions the two differently forced runs of Section 3.2.3. This is the problem to which the 4DDA algorithm is least suited and conversely one which satisfies all the conditions for which the strong-formulation 4D-VAR technique was designed. Errors are large and predominantly in the existence and amplitude of features, rather than in the displacement of their location. The results are shown in Figure 6.7. Plot (1) shows the initial vorticity difference for the experiment. The second shows the reduction in cost function obtained by the displacement assimilation. Although the reduction is greater than in the adiabatic case, it is still small, bearing in mind the quadratic nature of the cost function in question and that for this case the denominator, the initial cost-function, $\mathcal{J}_0 = \mathcal{J}[q_{\text{background}}]$.

Plotting the constant part of the obtained top-layer assimilation stream function gives (3), with contours from the mean of the background top-layer streamfunction superimposed for reference. This shows that the assimilation, or at least the assimilation velocity, is still dominated by features around the jet, although now principally forced by a large meander roughly 400km east of the jet separation point. This feature is also obvious in the vorticity anomaly field. The streamfunction shows a dipole flow, attempting to reduce the meander and move it south.

For this problem the 4D-VAR algorithm rapidly reduces the objective cost function to a level less than 20% of its inital value, as show in plot (5). This rapid reduction is unsurprising. Since the test case is run over a relatively short time window, while the background errors are large and diabatic, the observed difference between background and observations at the points, within the time window, when observations are taken are very similar to the differences between background and truth at the beginning of the window, namely the field the assimilation method is trying to find. This can be observed by studying the form of the gradient search direction supplied by the 4D-VAR algorithm and comparing it with the vorticity anomaly in plot (1). It will be readily observed that many features around the middle left of the box appear prominently in both plots, with only the finer scale structures absent in some places. The excellent behaviour of the 4D-VAR algorithm is to be expected for what is in many respects a trivial problem for this method.

Figure 6.7: Results for diabatic IC error

(1) Initial top layer vorticity anomaly, (2) 4DDA Reduction in cost function, (3) Top layer assimilation streamfunction (colour plot) with mean background top layer streamfunction (contours) (4) 4D-VAR reduction in cost function, (5) Initial 4D-VAR vorticity increment field .
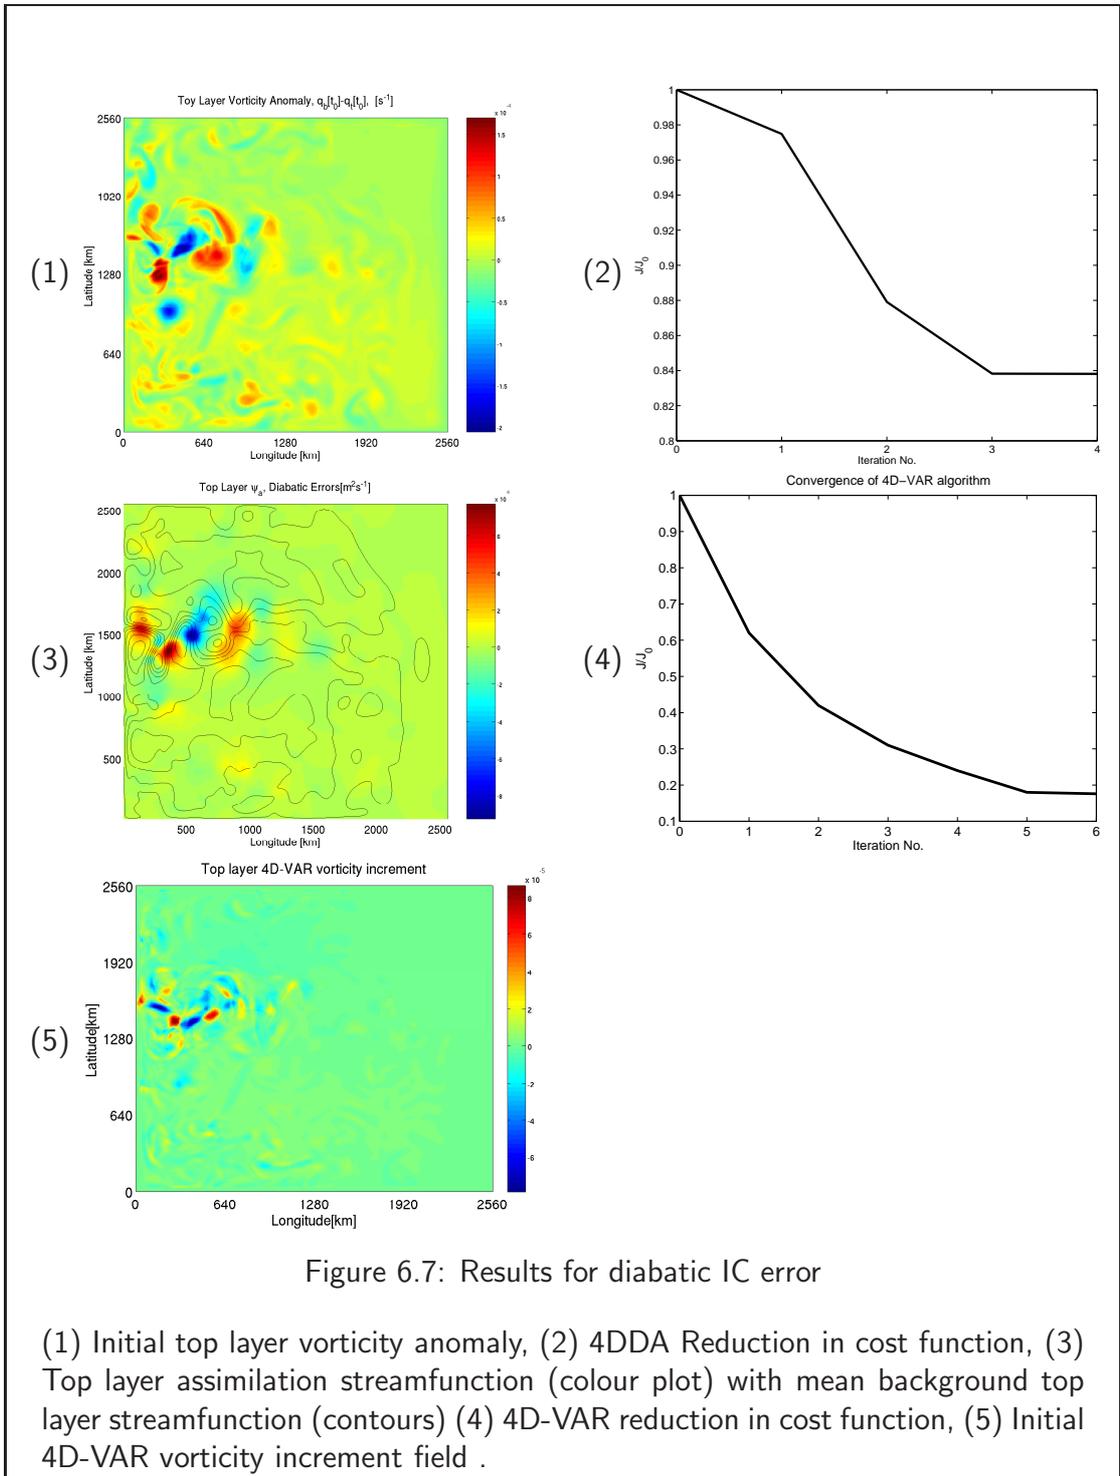
6.2.4 IMPERFECT MODEL EXPERIMENTS

As a final test case dealing with direct observations of tracers themselves we consider the case of more general model errors, which may include an explicit diabatic component. To generate these errors we consider the case of two runs from models based on different discretizations, initialized from the same initial conditions. In this experiment the truth is assumed to be the discrete model based on a 10km grid-point separation distance, Δx, as used in the other sections of this chapter. The background run is from an otherwise identical model with a 20km value for Δx, a value which means that both the westward boundary currents and mesoscale eddies in the system are poorly resolved, but below the Rossby radius of deformation for each layer in the interior of the basin, the scenario commonly termed 'eddy permitting'. The coarse scale model is initialized so that each grid point holds the value of the equivalent fine-scale 'truth' run. This means that gradient information is lost and the interpolation of the coarse-scale grid onto the fine scale no longer satisfies the discretized version of the rearrangement condition. The approach of viewing a fine-scale model as truth in forcing a coarse run has a long history, particularly in studies concentrating on methods for increasing the variability between members of ensemble forecast predictions, in order to better give a larger spread to the model trajectories and cover a wider selection of possible model trajectories near bifurcation paths.

Figure 6.8 shows the results of an experiment assimilating full observations at 6 hourly intervals over a 72 hour time window, fitting in the 4DDA experiment to 7 basis functions, one constant function, 3 sine functions and 3 cosine functions over the time window. Plots (1) and (2) show the initial anomaly, $q_{\text{truth}}[t_f] - q_{\text{background}}[t_f]$, and the final analysis anomaly, $q_{\text{truth}}[t_f] - q_{\text{background}}[t_f]$ following an application of the 4DDA algorithm. The initial anomaly between the two runs represents primarily the difference in the magnitude of the frictional term, $\nu \nabla_H \psi_i$, caused by the smoothing of sharp gradients. This is shown by the large-scale, filamentary, structure of the signal, which is concentrated along the western boundary and, in particular, around the cut-off eddy and meander visible in the top-layer vorticity, plot (8). Plots (3)–(4) show the time-constant mode of returned assimilation streamfunction in colour, with superimposed contour lines from the model streamfunction. In the top layer this shows the assimilation forcing acting to reduce and shift north the strength of the jet, shown by the large positive signal along the jet contour, and to encourage the eddy to move south and roll up, shown by the dipole structure over the centre of the eddy.

Returning to the analysis anomaly, plot (1), we see this approach is successful

in reducing the magnitude of the difference from the fine run, but retains almost exactly the same structure, not introducing any obvious noise. Looking once more at the energy stored in the various modes, plot (6), we see that the time-constant mode contains more than three times the energy of the next largest mode, the period one sine function. This energetic time-constant mode is typical of situations in which the method has converged successfully to an improved analysis solution, as shown by the cost function convergence, (5). Here the cosine modes contain almost as much energy as the sine functions of the same period. This is a signal of a problem for which the code has avoided the overshooting so plain in the previous problem.

For this class of problem the 4D-VAR algorithm fails to descend to an improved analysis. The increment field, plot (7) of Figure 6.8, has a sharp dipole structure centered over the cut off-eddy. Unfortunately, it seems to be too smooth for the gradient descent algorithm to operate successfully, since the maximum step allowed before the increment field is no longer a descent direction is too small.

6.2.5   ASSIMILATING NON-MODEL VARIABLES

We conclude this section of this thesis by returning to the question of the realistic case of assimilating limited, noisy observations. More particularly to the question of assimilation of the satellite altimetry proxy of upper layer sea surface height. As discussed in section 5.3.2, assimilating this form of observation, although linear in the model variables, introduces another pair of elliptic operators into the calculation of the adjoint model. In particular the near singular discrete operator ∇^{-2}, arising from solving for the baroclinic mode, adds appreciably to the run time of the code. In fact we find that generally the code refuses to descend more than minimally in either 4DDA or 4D-VAR formulation. This is disappointing, but suggests there may be insufficient information in such surface observations to perform displacement assimilation without additional regularization terms. This is a possible topic for future study.

6.3   CHAPTER SUMMARY

The validity of the model adjoint codes used in the presented 4D-VAR and 4DDA technique has been tested through comparison with a finite-difference gradient test. The 4DDA code is found to generate a less accurate, but more robust search direction than the same routine in 4D-VAR mode which, it is argued, is due to the rep-
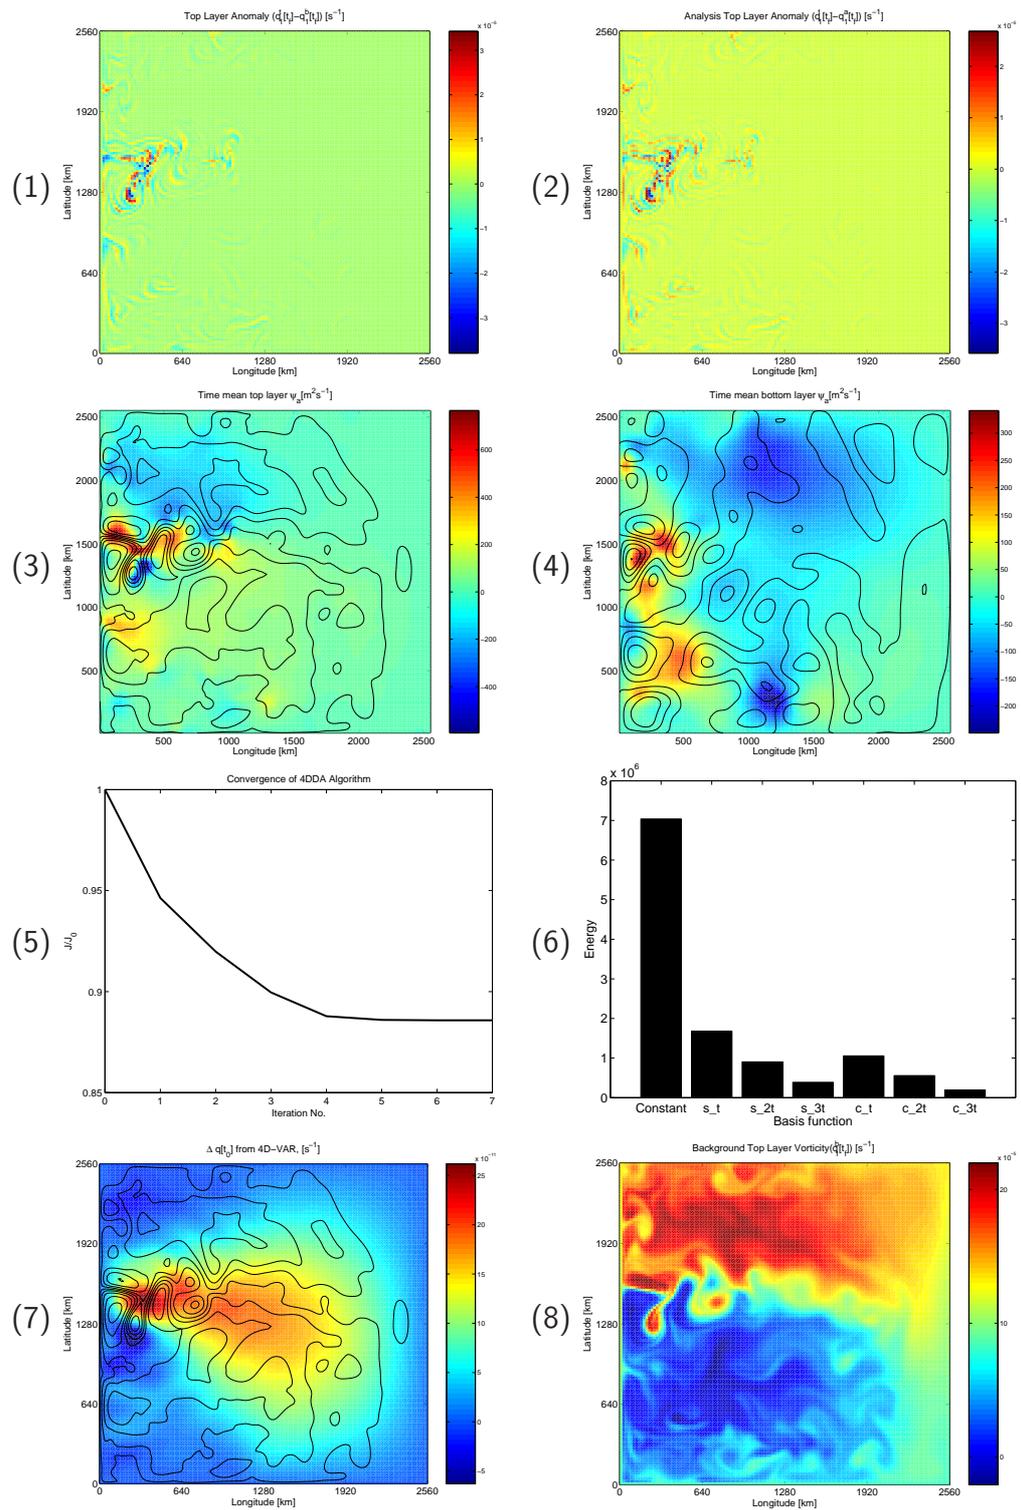
Figure 6.8: Results for model error

(1) Background & (2) Analysis vorticity anomalies. (3) Top layer & (4) Bottom layer assimilation streamfunction (colour) & mean background streamfunction (contours). (5) Reduction in cost function under 4DDA. (6) Energy in basis modes, (7) 4D-VAR gradient. (8) Top layer background vorticity.

resentation of the temporal variation of the assimilation streamfunction allowed through the basis function approach.

Both algorithms were then applied to a series of test problems and the behaviour compared. The 4DDA algorithm was found superior in the cases of model and forcing error, while failing for errors in initial conditions, with broadly the opposite behaviour for the 4D-VAR algorithm. This is understandable given the motivations behind the different formulations, but it is unfortunate that the 4DDA method was relatively poor in the case of adiabatic initial condition error for which it might be considered suitable.

CHAPTER 7

CONCLUSIONS

"Now this is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning."

Winston Churchill
Speech following 2nd Battle of El Alamain
10th November 1942

Conservation properties and phase error are closely related areas in which standard data assimilation algorithms may fail. There are several approaches which can act to alleviate this problem. The aim of this thesis has been to understand the benefits and issues involved in using phase error correction methods in a general, mathematically rigorous and physically consistent ocean forecasting system.

Most previous development in this area has been based in the context of extremely well observed systems. Here the approach has been to extend towards problems with limited observations and to examine destructively under what conditions such approaches fail. This is first performed for an algorithm based on the descent of a spatial cost-function analogous to 3D-VAR.

A novel approach based on a four dimensional algorithm equivalent to a modified version of the weak constraint 4D-VAR algorithm was developed and applied in the context of a quasi-geostrophic two-layer ocean model. This included the development of the forwards and adjoint models as well as error visualisation techniques.

This concluding chapter reviews the content of the thesis as a whole, summarizes the response to the key questions posed in Chapter 1, discusses the potential for the application of these results in operational ocean forecasting and finishes with an enumeration of possible directions for future work.

7.1   Summary of the Thesis Content

In Chapter 2 we have identified a methodology, based on the identical twin philosophy and using trajectories of a simple two layer quasi-geostrophic model, to investigate the behaviour of data assimilation methods which remove phase error from ocean models. In Chapter 3 it was demonstrated that traditional data assimilation techniques can fail in the presence of phase error. We reviewed methods based on the concept of rearrangements and showed that one particular technique, the Cooper & Haines method, which has been implemented operationally, is not of itself sufficient to remove phase error from models once initiated. Chapter 4 instead develops an assimilation approach which can correct for phase error even with limited observations and includes rearrangements excluded from the Cooper & Haines approach. This method is based on adiabatic rearrangements through an advective operator. We find that the technique returns qualitatively representative analyses for a number of semi-idealised problems, provided the number of observations allows suitable features to be identified. Chapter 5 extends the approach into a method we identify with a 4D-VAR type analysis, which allows for information from more observations to be assimilated. This is a novel combination of standard techniques from data assimilation and rearrangement theory. In Chapter 6 we show that the method can be used for assimilation successfully for a number of problems, involving both phase and amplitude errors. However, the technique is limited by the volume of observations available over the time window.

7.2   Review of Key Questions

We now return to considering the questions first posed in section 1.7.2, restate the work presented in this thesis in relation to each, examine the extent to which these key issues have been successfully answered in the course of this thesis and the possible methods of further inquiry which are suggested.

7.2.1   Question 1

To what extent is it feasible to decouple the observed error signal in ocean analyses consistently into signals due to pure phase error and amplitude error?

The rearrangement theory discussed in Chapter 2 and resultant metrics of Chapter 3 present one method for separating the two signals, however this approach is only feasible in the presence of full observations. As such it may be used in compar-

isons of forecast outputs and analyses or retrievals, but cannot be compared directly against data. The iterative approaches of Chapters 4 and 5 represent an attempt to alleviate this requirement, but do not directly generate a measure for the relative phase and amplitude errors in the system. Experiments with both schemes show a general failure to achieve global minima in solving the phase error problem. This suggests that the signals cannot be truly decoupled using the present approach. In this situation it is possible that the simultaneous approach of Hoffman et al. (1995) and Ravela et al. (2004) may be more appropriate.

7.2.2 QUESTION 2

Given that real oceanographic observations are noisy and limited in scope and number, can a study of novel phase correction algorithms give a quantitative understanding of how observations relate to model data in such methods and what constraints on observations?

The experiments of Chapter 4 show that the form and type of observations assimilated have a very large controlling effect on the adiabatic assimilation method used. The method effectively attempts to match the gradient of the background to the gradients of the observation, provided that the observational data was differentiable on the model grid. In this sense with full observations the gradient patterns of observation and background represent the 'shapes' of the fields to be matched. With incomplete observations the behaviour is different. Observations are matched for their values, with the difficulty becoming filling the unobserved region. Even with very limited data surprisingly successful assimilations are possible, in terms of the location and shape of oceanographic features.

7.2.3 QUESTION 3

Is it possible to develop efficient algorithms which maximise the utility of displacement assimilation techniques, for example in terms of the use of the information contained within observations, in the same way that 4D-VAR methods allow the use of multiple nonlinear observations taken at different times?

Which of the presented methods is best used to correct for adiabatic error when assimilating data into a system depends primarily on the form of error which is assumed to be present, as well as on the data themselves. When full observations of model variables are available the calculation of the polar factorization necessary

to minimize phase error directly is trivial and computationally inexpensive[1]. The coupled Monge-Kantorovich problem can be solved approximately with the same computational cost as a 4D-VAR assimilation problem using the methods of Benamou and Brenier (2001). This may however still be excessive considering many operational centres still use sequential or 3D-VAR methods for ocean data assimilation.

When observations are concentrated but not full, or are given in terms of non-model variables with an observation operator whose inversion is non-trivial (at least, one for which the adjoint operator is of less complexity than inversion) then the descent correction methods of Chapter 4 may be used to correct for background errors in states at a single verification time. Experiments have shown that the code is best run using an energy form for the normalization on the assimilation velocity and by choosing a multiscale approach to the assimilation, assimilating the large scale errors in the data faster than those on the smaller scale. For limited observations the multiscale routine approach can be used to fit to the natural scale of the observations themselves and thus reduce 'phantom' oscillation induced by the descent method.

When observations are very limited a four-dimensional approach may be used, so that the information contained in the observations is assimilated at the correct verification time. Introducing an explicit, smooth, assimilation velocity as in Chapter 5 allows for the correction of that portion of the adiabatic error which it is possible to ascribe to errors in the model. This approach is found to be poor at fitting to initial condition errors, for which the 3D approach is more suitable. An ideal adiabatic error correction system will combine the two approaches, both three and four dimensional, using the second approach to correct for model error and to distribute the observational information correctly in time, while using the first to correct for initial condition errors at the beginning of the time window. Such an approach is currently too computationally expensive to be considered generally feasible, but this may change with time.

7.3 IMPLICATIONS FOR OPERATIONAL MODELS

As previously discussed no operational centre currently uses a direct phase error approach for data assimilation purposes, either in NWP or ocean forecasting. The CH algorithm, however, is currently implemented to assimilate satellite altimetry data. The experiments in this thesis are performed using a highly idealized system.

---

[1]The principal costs are the sorting algorithms, plus any costs based on quadrature.

The layer QG model is much simpler than the level based primitive equation models in use for ocean forecasting, and the twin philosophy used in testing ignores many of the degrees of freedom available in the error signal of numerical models when compared to observational data.

Unfortunately most current ocean modes do not include potential vorticity as a prognostic variable. This increases the difficulty of consistently applying the fluid parcel rearrangement idea, since PV represents an important materially conserved quantity which is dynamically active and has a large control over currents through the momentum equation. Regarding other tracer quantities the necessary rearrangements are most sensibly accomplished in 2D on isentropic surfaces. As stated in Section 2.1, these may be approximated by surfaces of constant neutral density, but these density values are themselves dependent on the dynamic tracers of temperature and salinity. One method is thus to calculate increments to tracer values on isopycnal surfaces and iterating with calculation of the change to the density fields.

7.4 FUTURE WORK

Having discussed the extent to which the initial key questions tested during this thesis have been answered we conclude by posing some open questions which remain:

(i). Alternative Cost Functions in the Minimization Problem

Throughout the final four chapters of this thesis, and indeed in the methods discussed in Chapter 3, the fitting to observations is measured by a quadratic cost function. This is in part because such functions are a standard in statistical data assimilation methods, due to the results in the derivation of the BLUE equation of Section 1.3.1. In this case the background term is weighted by the background covariance matrix, usually assumed full, while the observation term is habitually weighted only by the error variances of the instrument in question. In principle the weighting on the observation term for adiabatic methods may be any function of the observed differences with the predictions of the observation operator,

$$\mathcal{J} = \mathcal{F}\left(|\boldsymbol{y} - h[\boldsymbol{x}]|; \boldsymbol{y}, \boldsymbol{x}, h, \boldsymbol{r}, t\right).$$

Terms of order other than a quadratic one give different weighting to the local fitting of the gradient in assimilation velocities to large and small magnitude errors compared to observations in both the three and four dimensional cases. Moreover

the freedom to vary weighting with location, observation type and time allows the method to be selective, depending on background conditions. Such an approach is not implemented here, but may alleviate some of the sensitivity issues both methods have regarding the separation point of the jet. One could, for example use a normalization factor based on the local magnitude of the average background signal, in the same manner the observation error covariance matrix is used in the BLUE equation.

(ii). Best Methods of Combining Phase and Amplitude Assimilation

This thesis has concentrated on the goal of separating signals in data from phase and amplitude error, in order to understand the differences in behaviour of the two sources and thus obtain a clearer idea of the best method of correcting each. In practice though both forms of error pollute model analyses and both must be corrected. Considering the implementations of phase error correction techniques presented in Section 3.3.1, all techniques (with the exception of the Cooper & Haines) also include terms to explicitly correct amplitude error. In each case this requires an *a priori* estimate of the relative magnitudes of the two forms of error, in order to decide the size of the penalization applied to each. To accurately tune such parameters would require a large-scale statistical survey of the presence of each form of error in analysis. This could perhaps be performed by observing the differences between successive analysis valid at a single validification time, as is sometimes used to tune background error covariance matrices. The new metrics introduced in Chapter 2 might be a useful tool in such a survey.

(iii). Assimilating Lagrangian Positional Data

As stated in Chapter 1, there has been a massive effort to reduce the under-observation of the ocean interior through the use of passive floats and drifters. The returns of probe position with time that these probes generate represent information on model phase error, errors in the large scale currents and the effects of sub-gridscale processes. This seems like an appropriate data set for the use of adiabatic assimilation methods. It is possible to model the behaviour of a collection of m probes by considering the evolution of the conditional probability of probe location for a given set of observations,

$$P(\boldsymbol{r}, t) := P(\boldsymbol{x(t)} = \boldsymbol{r} | \boldsymbol{y} = \boldsymbol{r}_y),$$

where $\boldsymbol{r} \in \mathbb{R}^{3m}$ is a vector of probe location co-ordinates, $\boldsymbol{x}$ is the true probe position and $\boldsymbol{y}$ is the usual vector of observations. Consider the following set of as-

sumptions:

i. Each probe is an independent observation (likely to apply when probes are well separated).

ii. The small scale mixing processes which act to spread the probability density are broadly the same as those which mix material properties of fluid parcels. This will obviously fail for processes such as convection.

The individual probability densities,

$$P_i(\boldsymbol{r}, t) := P(x_i(t) = \boldsymbol{r} | \boldsymbol{y} = \boldsymbol{r}_y),$$

will then, in the absence of observations of the ith probe, satisfy the true system tracer evolution equation,

$$\frac{\partial \mathcal{P}_i}{\partial t} + \boldsymbol{u}_t \cdot \nabla \mathcal{P}_i = F(\mathcal{P}_i).$$

The probability densities may thus be treated like any other model tracer and assimilated adiabatically using the same techniques, given a model for observation error,

$$\mathcal{P}_y = Y(r_y).$$

In this manner any of the presented adiabatic assimilation techniques may be used. Note that since, moreover, a single observation of a probe location is enough to give $\mathcal{P}_y$ everywhere, these data are uniquely suited to the polar factorization plus transport problem approach. Using the direct minimization approach of the second half of the thesis would require an extension of the algorithms and codes developed in Chapters 5 and 6 to accept input from the probability functions.

Further research in any of these topics may prove fruitful. Considering the proportion of in situ observations which are now being obtained from floats the third in particular holds interest.

APPENDIX A

MODEL NUMERICS

A.1   NUMERICAL SCHEME

Except for trivial cases, it is impossible to solve the equations governing the evolution of the ocean analytically, even at the level of simplification of the continuous QG equations. In order to compute approximate solutions to the equations numerically the infinite dimensional spatial and temporal domain must be discretized into equations suitable for computers of finite capacity. Possible discretization methods include using truncated spectral functions, truncated integral solutions using finite elements or further approximating the model equations themselves using finite differences. Spectral models are little used in ocean circulation modelling; the complicated geometry of real ocean basins means there is no simple set of orthogonal basis functions to use, unlike compared to the simple spherical harmonics used in atmospheric modelling. Several 'next generation' ocean models, such as the ICOM project (Ford et al., 2004) are now using finite element techniques, primarily due to the relative simplicity of application of asymmetric or adaptive meshes, but the majority of ocean models, including the operational ocean models at both the UK Met Office and the ECMWF use a finite difference modelling approach and it is that methodology which is followed here.

The derivation of the two layer quasi-geostrophic equations is discussed in Section 2.3, so here we will simply present the PDEs to be discretized for easy reference. For a two dimensional domain, $\mathcal{D}(x, y)$, using a choice of principal variables of potential vorticities, q_1, q_2, with equivalent streamfunctions, ψ_1, ψ_2 ,the QG equations are

$$\frac{dq_1}{dt} := \left(\frac{\partial}{\partial t} + \frac{\partial \psi_1}{\partial x} \frac{\partial}{\partial y} - \frac{\partial \psi_1}{\partial y} \frac{\partial}{\partial x} \right) [\xi_1 + \beta y] = \frac{1}{\rho_1 H_1} \left(\frac{\partial \tau^{(y)}}{\partial x} - \frac{\partial \tau^{(x)}}{\partial y} \right) + \nu \nabla^4 \psi_1,$$
$$\text{(A.1)}$$

$$\frac{dq_2}{dt} := \left(\frac{\partial}{\partial t} + \frac{\partial \psi_1}{\partial x} \frac{\partial}{\partial y} - \frac{\partial \psi_1}{\partial y} \frac{\partial}{\partial x} \right) [\xi_2 + \beta y] = \nu \nabla^4 \psi_2, \qquad \text{(A.2)}$$

$$\xi_i := \nabla^2 \psi_1 - \frac{f_0^2}{g' H_1}(\psi_1 - \psi_2), \tag{A.3}$$

$$\xi_2 := \nabla^2 \psi_2 - \frac{f_0^2}{g' H_2}(\psi_2 - \psi_1). \tag{A.4}$$

The other free parameters are the layer densities, ρ_1, ρ_2, average layer depth, H_1, H_2, the wind stress, $(\tau^{(x)}, \tau^{(y)})$, Coriolis parameter, f, β parameter and viscosity, η. These equations are discretized using central finite differencing in both time and space for all derivatives, apart from the friction term, which uses the forward Euler method. This is necessary for stability. This discretization method is applied to variables stored on a regular square grid, denoted the Arakawa A grid in the standard notation of Arakawa and Lamb (1977), see Figure A.1 for a schematic of the stencil. The A grid is chosen, since under central differencing the natural locations off the finite difference approximation to the QG vorticity coincides with the nodal points for streamfunction, which is equivalent to a pressure term. This both obviates the need for extra smoothing operators and allows the use of an energy and enstrophy conserving formulation for the advective Jacobian term in the evolution equations. At nodal points the continuous and discretized forms of variables are assumed to coincide;

$$\psi(x, y, t) = \psi(i\Delta x, j\Delta y, n\Delta t) = \psi_{i,j}^n,$$

$$q(x, y, t) = q(i\Delta x, j\Delta y, n\Delta t) = q_{i,j}^n.$$

Under central differencing the partial derivative with respect to time becomes

$$\left.\frac{\partial \psi}{\partial t}\right|_{(i\Delta x, j\Delta y, n\Delta t)} = \frac{\psi_{i,j}^{n+1} - \psi_{i,j}^{n-1}}{2\Delta t},$$

with the Laplacian operator represented by

$$\left.\nabla^2 \psi\right|_{(i\Delta x, j\Delta y, n\Delta t)} \approx \frac{\psi_{i-1,j}^n - 2\psi_{i,j}^n + \psi_{i+1,j}^n}{(\Delta x)^2} + \frac{\psi_{i,j-1}^n - 2\psi_{i,j}^n + \psi_{i,j+1}^n}{(\Delta y)^2}.$$

The biharmonic operator deriving from isopycnal vorticity diffusion is evaluated assuming $\nabla^4 f = \nabla^2[\nabla^2[f]]$ for any arbitrary scalar field. The Jacobian terms are calculated using the so called Arakawa Jacobian (Arakawa and Lamb, 1981). This is the stencil given by taking equal weightings in an average of the discretizations of the three equivalent continuous forms

A grid
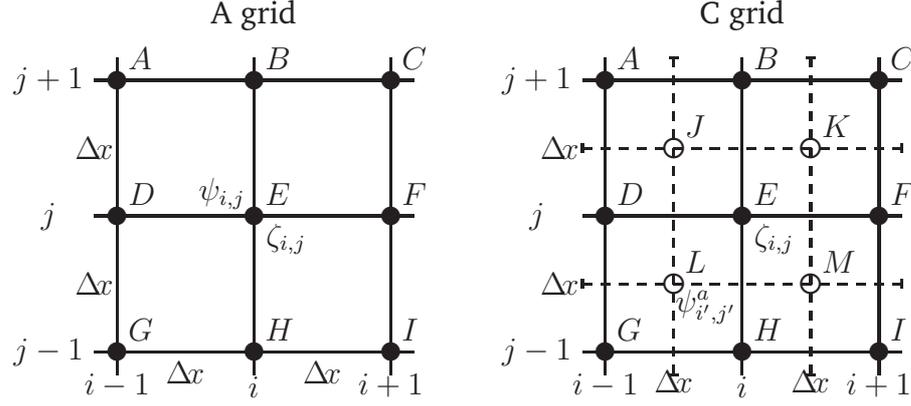
C grid

Figure A.1: The stencil for the Arakawa A and C grids

$$
\begin{aligned}
1. \qquad J(\psi, \xi) &= \frac{\partial \psi}{\partial x}\frac{\partial \zeta}{\partial y} - \frac{\partial \psi}{\partial y}\frac{\partial \zeta}{\partial x}, \\
2. \quad J(\psi, \xi) &= \frac{\partial}{\partial x}\left(\psi\frac{\partial \zeta}{\partial y}\right) - \frac{\partial}{\partial y}\left(\psi\frac{\partial \zeta}{\partial x}\right), \\
3. \quad J(\psi, \xi) &= \frac{\partial}{\partial y}\left(\zeta\frac{\partial \psi}{\partial y}\right) - \frac{\partial}{\partial x}\left(\zeta\frac{\partial \psi}{\partial y}\right).
\end{aligned}
$$

This particular formulation is a common choice in numerical studies since it is possible to show that it conserves the domain integrals both the system kinetic energy, $\rho\|\psi\nabla^2\psi\|_1/2$, and the enstrophy, $\rho\|\nabla^2\psi\|^2$. It is argued that this then prevents growth of nonlinear instabilities. The final discretization for this term is

$$
12\,(\Delta x)^2 \left(\frac{\partial \psi}{\partial x}\frac{\partial \zeta}{\partial y} - \frac{\partial \psi}{\partial y}\frac{\partial \zeta}{\partial x}\right)^n_E =
\begin{aligned}
&\left(\psi^n_H + \psi^n_I - \psi^n_B - \psi^n_C\right)\left(\zeta^n_F + \zeta^n_E\right) \\
&- \left(\psi^n_G + \psi^n_H - \psi^n_A - \psi^n_B\right)\left(\zeta^n_E + \zeta^n_D\right) \\
&+ \left(\psi^n_F + \psi^n_C - \psi^n_D - \psi^n_A\right)\left(\zeta^n_B + \zeta^n_E\right) \\
&- \left(\psi^n_I + \psi^n_F - \psi^n_G - \psi^n_D\right)\left(\zeta^n_E + \zeta^n_H\right) \\
&+ \left(\psi^n_F - \psi^n_B\right)\left(\zeta^n_C + \zeta^n_E\right) \\
&- \left(\psi^n_H - \psi^n_D\right)\left(\zeta^n_E + \zeta^n_G\right) \\
&+ \left(\psi^n_B - \psi^n_D\right)\left(\zeta^n_A + \zeta^n_E\right) \\
&- \left(\psi^n_F - \psi^n_H\right)\left(\zeta^n_E + \zeta^n_C\right),
\end{aligned}
\tag{A.5}
$$

where the subscripts denote the locations on the stencil centering on point E in figure A.1.

The solution algorithm for ξ^{n+1}, ψ^{n+1} proceeds in the following fashion:

i. If ζ_i and ψ_i are known at both current and previous time levels then equations (A.1–A.2) give explicit equations for ζ^{n+1}_i in terms of known variables.

ii. The potential vorticity fields are then inverted to give the streamfunction (and hence column velocities) by transforming variables in equations (A.3–A.4) to

obtain the equations for the (uncoupled) barotropic and baroclinic modes,

$$\nabla^2\left(\psi_1 - \psi_2\right) - \frac{f_0^2(H_1 + H_2)}{g'H_1H_2}(\psi_1 - \psi_2) = \zeta_1 - \zeta_2, \qquad \text{(A.6)}$$

$$\nabla^2\left(H_1\psi_1 + H_2\psi_2\right) = H_1\zeta_1 + H_2\zeta_2. \qquad \text{(A.7)}$$

These two resulting block tridiagonal elliptic equations are then solved to machine precision using the NAG multigrid elliptic 2d solver `nag_pde_ell_mg` (Numerical Algorithms Group Ltd., 1998).

The boundary condition for the baroclinic part is simply that it is spatially constant on the boundary,

$$H_1\psi_1 + H_2\psi_2 = K(t).$$

K may be taken to vanish without modifying the resultant velocity field. The boundary condition for the barotropic mode is solved by decomposing the problem into two parts; Firstly the particular integral problem,

$$\nabla^2\phi - \frac{f_0^2(H_1 + H_2)}{g'H_1H_2}\phi = \chi,$$

is solved with $\phi = 0$ on the boundary, then a scalar multiple of ϕ_0, the solution to the homogeneous equation

$$\nabla^2\phi - \frac{f_0^2(H_1 + H_2)}{g'H_1H_2}\phi = 0,$$

which has unit value on the boundary, is added. The Γ required for

$$\psi_1 - \psi_2 = \phi + \Gamma(t)\phi_0$$

is calculated using the mass conservation equation,

$$\int_{\mathcal{D}} \phi \, dS = 0$$

which implies in turn that

$$\Gamma = \frac{\int \mathcal{D}\phi(t)dS}{\int \mathcal{D}\phi_0 dS}.$$

Note that ϕ_0 need only be calculated once, when the code is initialized.

A.2   THE PASSIVE TRACER ADVECTION SYSTEM

An arbitrary number, N_τ of (column integrated) passive tracer concentration fields, $\tau^j(x,y)_{1,2}$, $j \in \{1,\ldots,N_\tau\}$, $0 \le \tau \le 1$ are assumed to obey the layer tracer transport equations,

$$\frac{D\tau_i^j}{DT} := \left(\frac{\partial}{\partial t} + u_i \frac{\partial}{\partial y} + v_i \frac{\partial}{\partial y}\right)\tau_i^j = \nabla_H \cdot (\kappa_\tau \nabla \tau_i^j), \qquad (A.8)$$

where $i \in 1,2$ as before denotes the layer number and u_i and v_i are the constant geostrophic velocities in each layer given by

$$u_i = -\frac{\partial \psi_i}{\partial y},$$

$$v_i = \frac{\partial \psi_i}{\partial x}.$$

The along-isopyncal diffusivities, κ_τ, of the tracer fields are assumed spatially homogeneous and equal to the along-isopycnal dynamic viscosity of the system, η. This means that bottom layer total vorticity satisfies its own tracer equation and that in the unforced limit the initial q–τ relation is exactly conserved over the course of an integration run. The lateral boundary condition for the diffusion operator, $\nabla \tau . \boldsymbol{n} = 0$, is assumed for $\boldsymbol{n}$ a unit normal to $\delta \mathcal{D}$. This follows naturally from integrating τ across $\mathcal{D}$ and demanding conservation of total τ,

$$\frac{d}{dt}\int_{\mathcal{D}} \tau_i = 0.$$

The equation set (A.8) is discretized on an Arakawa C grid, see Figure A.1, for its conservation properties in flux form and the ease of calculation of edge fluxes for the control volumes. Discretization of the volume flux terms is through the QUICK upwinding form (Leonard, 1979),

$$(u\tau)_e = \begin{cases} u_e\left(\frac{6}{8}\tau_C + \frac{3}{8}\tau_E - \frac{1}{8}\tau_W\right) \text{ for } u_e > 0 \\ u_e\left(\frac{6}{8}\tau_E + \frac{3}{8}\tau_C - \frac{1}{8}\tau_{EE}\right) \text{ for } u_e \le 0 \end{cases} \quad \text{etc}$$

for its third order truncation accuracy and tendency to avoid oscillation and spurious introduction of unphysical negative values for concentration fields. The diffusion term is discretized using a nine point stencil in an attempt to avoid significant computational modes.

Temporal discretization is chosen to be through the standard 4$^{\text{th}}$ order Runga

Kutta algorithm,

$$\tau_1^{n*} = \tau(t_n) + \frac{\Delta t f(t_n, \tau)}{2},$$

$$\tau_2^{n*} = \tau(t_n) + \frac{\Delta t}{2} f\left(\frac{t_n + \Delta t}{2}, \tau_1^*\right),$$

$$\tau_3^{n*} = \tau(t_n) + \frac{\Delta t}{2} f\left(t_n + \frac{\Delta t}{2}, \tau_2^*\right),$$

$$\tau_4^{n*} = \tau(t_n) + \frac{\Delta t}{2} f\left(t_n + \Delta t, \tau_3^*\right),$$

$$\tau^{n+1} = \frac{\left(\tau_1^{n*} + 2\tau_2^{n*} + 2\tau_3^{n*} + \tau_4^{n*}\right)}{6}.$$

This method is chosen primarily for its well-known area of convergence with respect to stability, since it is to be noted that stability criteria are extremely strict around the point of jet separation and it has been shown that the QUICK algorithm is unstable under implementation of a central difference (leapfrog) temporal discretization (Webb et al., 1998).

A.3   NUMERICAL ROUTINES FOR $\mathcal{M}$, $\mathcal{J}^*$ AND $\mathcal{J}$

The integrals in the definitions of the three metrics on functions $f, g : \Omega \to \mathbb{R}$,

$$\mathcal{J}(f, g) = \int_\Omega (f - g)^2 dV,$$

$$\mathcal{J}^*(f, g) = \inf_{\hat{f} \sim f} \int_\Omega (\hat{f} - g)^2 dV,$$

$$\mathcal{M}(f, g) = \int_{-\infty}^{\infty} \int_\Omega |(f - \alpha)\mathcal{H}(f - \alpha) - (g - \alpha)\mathcal{H}(g - \alpha) dV d\alpha$$

must be discretized if these operators are to be applied to numerical data. Of the three the first is the simplest to approximate. Any quadrature may be chosen, with the associated errors easily calculated. For the results presented in Chapter 3 a piecewise flat approach was taken, with f and g assumed equal to their nearest grid-point values. This leads to a discretization conveniently expressed as

$$\mathcal{J} = \sum_{i=1}^{n_x \times n_y} (f_i - g_i)^2 \Delta x \Delta y,$$

where f_i etc. are grid point values under some unique ordering. The truncation error under this approach is of order

$$\epsilon(\mathcal{J}(f,g)) = \int_{\Omega} (\|\nabla f\| + \|\nabla g\|)|f-g|dV,$$

provided ∇f, ∇g are small.

The second function, $\mathcal{J}^*$, is calculated following the polar factorization approach of Section 3.1.2. Once again, an initial nearest grid point assumption is made as to the value of non-nodal points, in keeping with the area averaged value approach of the control volume discretizations of the QG model presented above. This assumption also produces simple, fast code, since under this approach the polar factorization algorithm reduces to a reordering of discrete units, or in terms of equations,

$$f_i = \tilde{f}_{\sigma_f(i)},$$

$$g_i = \tilde{g}_{\sigma_g(i)},$$

where $\hat{f}_i$ is a rearrangement of f with $\tilde{f}_1 \leq \tilde{f}_2 \leq \ldots \leq \tilde{f}_{n_x \times n_y}$. The invertible mapping, $\sigma_f : [1, \ldots, n_x \times n_y] \rightarrow [1, n_x \times n_y]$, is chosen so the equation above holds. The function $\tilde{g}$ and mapping σ_g are defined similarly. The analysis step which minimises $\mathcal{J}(\hat{f}, g)$ over rearrangements of f is then defined as

$$\hat{f} = \tilde{f}_{\sigma_g(i)}.$$

The algorithm thus requires one application of a sorting algorithm to each f_i of g_i, storing the index of the reordering, to generate $\hat{f}_i$. The same $\mathcal{J}$ algorithm as above is then applied to calculate a value for $\mathcal{J}^*$. The local dependence of the error bound given above is replaced by a global one,

$$\epsilon(\mathcal{J}^*(f,g)) = \int_{\Omega} (\sup_f\{\|\nabla f\|\} + \|\nabla g\|)|f-g|dV.$$

To calculate $\mathcal{M}$ two integrals must be evaluated. To reduce the errors in the integrals in α a piecewise linear interpolation assumption must be used, rather than the flat approach used above, since this makes the inner integral more continuous in α, in turn making quadrature more accurate. Calculating the inner integral for fixed α then requires (for two dimensional functions) calculations of the volumes of polygons with regular rectangular or triangular cross section. A routine to calculate this is combined with simple one dimensional numerical integration in α.

APPENDIX B

DISPLACEMENT METHOD NUMERICS

B.1 TWO DIMENSIONAL DISPLACEMENT METHOD NUMERICS

B.1.1 SPATIAL DISCRETIZATIONS

As a proof of concept a discretized form of the continuous method was applied to a full set of point observations of an adiabatically perturbed system. A regular Arakawa C grid (as in the passive tracer advection scheme of the previous section, see figure A.1) was used with equal grid spacings to discretize the interval $[0, 1] \times [0, 1]$, with boundaries passing through ψ points. A system of discretized equations were created with

$$\psi_{i,j}^k = [\overline{\rho_z}^x \overline{y_x}^z - \overline{\rho_x}^z \overline{y_z}^x] , \tag{B.1}$$

$$u_{i,j'}^k = \psi_z^k, \tag{B.2}$$

$$w_{i',j}^k = \psi_x^k, \tag{B.3}$$

$$\tag{B.4}$$

where $\psi_{i,j}^k \approx \psi(i\Delta x, j\Delta z, k\Delta t)$ etc, $i' = i + \frac{1}{2}, j' = j + \frac{1}{2}$ and we have two sets of linear operators,

$$\phi_x = \frac{\phi_{i',j} - \phi_{i'-1,j}}{\Delta x}, \qquad \phi_z = \frac{\phi_{i,j'} - \phi_{i,j'-1}}{\Delta x}, \tag{B.5}$$

$$\overline{\phi}^x = \frac{\phi_{i',j} + \phi_{i'-1,j}}{2}, \qquad \overline{\phi}^x = \frac{\phi_{i,j'} + \phi_{i,j'}}{2}, \tag{B.6}$$

representing discretizations of partial differentiation and smoothing respectively. The lateral boundary condition is applied explicitly by setting $\psi_{i,j} = 0$ on boundary points. The evolution of the ρ field itself through the transport equation,

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{u}_a),$$

was again achieved using the QUICK formulation described in section A.2 for the flux terms in the discretization,

$$\Delta^k \rho_i j = (u^k \rho)_x + (w^k \rho)_z.$$

New values for the ϕ field were calculated whenever the cost function $\mathcal{J}$ ceased to descend. The cost function itself was evaluated as in section A.3 by

$$J(\rho, y) = \sum_{i',j'} \left[w(\rho_{i',j'} - \rho^0_{i',j'})^2 + (\rho_{i',j'} - y_{i',j'})^2 \right] \Delta x_{i',j'} \Delta z_{i',j'}. \tag{B.7}$$

This formulation is consistent with the assumption that ρ values represent volume averaged quantities and as such fits the flux formulation of equation (B.6). Here w represents a weighing towards the background term, i.e. the initial condition for the integration. This was actually found to be unnecessary in the work presented in Chapter 4 where $w = 0$ is assumed throughout.

B.1.2 TEMPORAL DISCRETIZATIONS AND TIME STEPS

Time discretization was by the same Runge-Kutta 4^{th} order method as used in section A.2. The question remains as to how to choose Δt, the length of the integration step. Since this defines both the length of the time integration and the length of the step along the descent path there are conflicting pressures upon it. On the one hand we would like it to be as large as possible to minimise the number of iterations required for convergence. On the other hand we do not want the step size so large that the numerical integration technique for the transport equation becomes unstable. This leads naturally to the decision that the step length will be the minimum of

(i). The step size giving maximum descent in $J(\rho)$

(ii). The maximum step size derived from a CFL condition for numerical stability of the integration,
$$\Delta t < C \frac{\Delta x}{|u|_\infty}, \tag{B.8}$$
where $C < 1$ is some positive constant.

For a linear descent method applied to a quadratic form,

$$K(\boldsymbol{x}) = \frac{1}{2} \boldsymbol{x}^T \boldsymbol{A} \boldsymbol{x} - \boldsymbol{b}^T \boldsymbol{x},$$

an exact form for the maximum descent step size can be found by applying the relation

$$\boldsymbol{x}^{k+1} = \boldsymbol{x}_k - \Delta t \boldsymbol{d}_k$$

and assuming that the point at which $\boldsymbol{d}_k$ is orthogonal to the local gradient of the quadratic form then $\boldsymbol{d}_k$ is no longer a descent direction. Solving for Δt this gives

$$\Delta t = \frac{\boldsymbol{d}_k^T (\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{b})}{\boldsymbol{d}^T \boldsymbol{d}}$$

The minimization in question can be evaluated as a quadratic function in ρ. unfortunately the non-linearity of the transport equation and the Runge-Kutta method mean that the equations generated can no longer be solved directly to obtain an optimal Δt. Linearizations are possible to generate an upper bound for the optimal step size. In practice, however, it is usually easier to calculate the limit from CFL condition, (B.8) and reduce the step size if necessary to obtain descent.

B.2   THE MULTISCALING METHOD

In both methods a new, larger-scale, Arakawa C grid is introduced, with

$$\Delta X = N\Delta x, \qquad \Delta Z = N\Delta z,$$

where capital letters denote properties on the new grid and small letters properties of the old. There are assumed to exist smoothing and coarsening functions mapping fields between the two grids. For most functions the coarsening function represents a volume average of the field value on the fine grid,

$$F_{IJ} = \frac{\sum_{i,j} f_{ij} \Delta x_{ij} \Delta z_{ij}}{\Delta X \Delta Z},$$

where the summation is over all fine cells contained in the new coarse cell. The sole exception is for observational data, where the cell value is the arithmetic mean of all observations in the cell for those containing at least one observation and zero otherwise,

$$Y_{IJ} = \begin{cases} \frac{\sum_{i=1}^{N_o} y_i^{IJ}}{N_o} & N_o > 0, \\ 0 & N_o = 0, \end{cases}$$

where $\boldsymbol{y}^I \boldsymbol{J}$ is the vector of observations in the coarse cell, of length N_o. In both methods a coarse streamfunction is then generated,

$$\Psi_{I,J}^k = \left[\overline{R_z}^X \overline{Y_X}^Z - \overline{R_X}^Z \overline{Y_Z}^x \right]$$

where the smoothing and differential operators of the previous section have been extended to the new grid.

B.2.1   THE STREAMFUNCTION FORMULATION

Under the streamfunction formulation the coarse stream function is immediately smoothed down to the fine grid, using a two-dimensional spline-fitting approach for smoothness. This approach is necessary since a naive approach using linear interpolation would ensure that maxima of the smooth ψ, which represent the centres of local rotation of the fluid, could only lie at the ψ points which coincide on the fine and coarse grids. Having calculated a smoothed streamfunction field the assimilation proceeds exactly as in the single-grid case.

In this case only the streamfunction calculation is carried out on the fine grid, so that the multigridding code represents an additional overhead on each iteration. Advantages arrive through the longer integration times which are possible for a given search direction in the descent routine.

B.2.2   THE FLUX FORMULATION

Under the flux formulation the coarse streamfunction field is used directly to calculate a coarse velocity field through equations

$$U_{I,J'}^k = \Psi_Z^k,$$

$$W_{I',J}^k = \Psi_X^k.$$

These coarse velocities are then used to evolve the coarse density field using the same discretizations as in the single grid case, but now applied on the coarse grid. The total net flux into each cell,

$$\Delta R_{IJ} = R_{IJ}^{k+1} - R_{IJ}^k,$$

is then divided equally across the relevant cells of the fine grid,

$$\rho_{ij}^{k+1} = \rho_{ij}^k + \frac{\Delta x_{ij} \Delta z_{ij}}{\Delta X \Delta Z} \Delta R.$$

This represents an implicit transport between the subcells contained in the coarse cell to achieve this distribution under the given large-scale velocity field. In general this process, while preserving the global integral of the tracer being transported, will be numerically unstable and highly diffusive. Since in this case the velocity field is chosen to smooth the differences between background and observation these problems are obviated to some degree. Since in this case the majority of calculations are carried out on the smooth grid the saving in terms of number of operations is significantly larger. The sole calculation which must be performed on the fine grid is the calculation of the cost function, since this ensures stability by ensuring convergence of the algorithm.

APPENDIX C

ADJOINT MODEL FORMULATIONS

C.1   THE 4D-VAR ADJOINT MODEL

In this formulation the (non-linear) forward model, $\mathcal{M}(\boldsymbol{q}(t_0), t)$, found by integrating the layer model equations,

$$\left(\frac{\partial}{\partial t} + \frac{\partial \psi_1}{\partial x} \frac{\partial}{\partial y} - \frac{\partial \psi_1}{\partial y} \frac{\partial}{\partial x} \right) [q_1 + \beta y] = \frac{1}{\rho_1 H_1} \boldsymbol{\nabla} \times \boldsymbol{\tau} + \nu \nabla^4 \psi_1,$$

$$\left(\frac{\partial}{\partial t} + \frac{\partial \psi_2}{\partial x} \frac{\partial}{\partial y} - \frac{\partial \psi_2}{\partial y} \frac{\partial}{\partial x} \right) [q_2 + \beta y] = \nu \nabla^4 \psi_2,$$

$$q_1 = \nabla^2 \psi_1 - \frac{f_0}{\sqrt{g' H_1}} (\psi_1 - \psi_2) + \beta y,$$

$$q_2 = \nabla^2 \psi_2 - \frac{f_0}{\sqrt{g' H_2}} (\psi_2 - \psi_1) + \beta y,$$

is assumed to evolve a state vector, $\boldsymbol{q}(t_0) = [q_1(t_0), q_2(t_0), \psi_1(t_0), \psi_2(t_0)]$, consisting of the entire fields of vorticity and streamfunction in the two layers, from an initial point in time, t_0 to a new model time, t. The goal is once again to to match observations, $\boldsymbol{y}(t)$, as in section 5.1.2, with the fit measured by a cost function, $\mathcal{J}$, which we will later define explicitly. Introducing the augmented Lagrangian,

$$\mathcal{J}^* = \mathcal{J} + \int_{t_0}^{t_f} \int_{\Omega} \boldsymbol{\theta} \cdot (\boldsymbol{q} - \mathcal{M}(\boldsymbol{q}(t_0), t) dV \, dt,$$

where $\boldsymbol{\theta} = [\theta_1, \theta_2, \phi_1, \phi_2]$ is the vector of Lagrange multipliers which will become our adjoint variables, we obtain formal optimality conditions for the constrained minimization of $\mathcal{J}$ of

$$\boldsymbol{q}(t_0) = \mathcal{M}(q(t_0), t), \text{ Forward model}$$

$$\left.\begin{aligned}
\frac{d_1\theta_1}{dt} &= -\phi_1 + \frac{\partial h}{\partial \zeta_1}(y - h[\zeta_1, \zeta_2]) \\
\frac{d_2\theta_2}{dt} &= -\phi_2 + \frac{\partial h}{\partial \zeta_2}(y - h[\zeta_1, \zeta_2]) \\
\nabla^2\phi_1 - \frac{f_0^2}{g'H_1}\phi_1 + \frac{f_0^2}{g'H_2}\phi_2 &= \frac{\partial(\theta_1, \zeta_1)}{\partial(x, y)} + \nu\nabla^4\theta_1 - a_1 \\
\nabla^2\phi_2 + \frac{f_0^2}{g'H_1}\phi_1 - \frac{f_0^2}{g'H_2}\phi_2 &= \frac{\partial(\theta_2, \zeta_2)}{\partial(x, y)} + \nu\nabla^4\theta_2 - a_2
\end{aligned}\right\} \text{Adjoint model}$$

with boundary conditions $\boldsymbol{\theta}(t_0) = \boldsymbol{\theta}(t_f) = \mathbf{0}$. Since $\nabla\mathcal{J}^* = \theta(t_0)$ the approach of Chapter 4 is followed, the adjoint equations are discretized and stepped backwards from the zero boundary condition at t_f. The calculated $\boldsymbol{\theta}$ is then used in a conjugate gradients descent method until $\mathcal{J}$ is minimized with respect to variations in $q(t_0)$. The numerics are the same as those described below for the 4DDA algorithm.

C.2 THE DISPLACEMENT ADJOINT MODEL

C.2.1 THE ASSIMILATION METHOD

In the examples presented in the main body of this thesis the material quantities being assimilated are the potential vorticity anomalies themselves. Since the flow is being advected in a two dimensional-sense we may introduce an 'assimilation' stream-function, ψ^a, such that

$$u_a = -\frac{\partial \psi^a}{\partial y}, \tag{C.1}$$

$$v_a = \frac{\partial \psi^a}{\partial x}. \tag{C.2}$$

Using (C.1–C.2) and (5.3) to modify (2.17) the continuous evolution equations for the extended system become

$$\left(\frac{\partial}{\partial t} + \frac{\partial \psi_1}{\partial x}\frac{\partial}{\partial y} - \frac{\partial \psi_1}{\partial y}\frac{\partial}{\partial x} + \frac{\partial \psi_1^a}{\partial x}\frac{\partial}{\partial y} - \frac{\partial \psi_1^a}{\partial y}\frac{\partial}{\partial x}\right)[\zeta_1 + \beta y] = \frac{1}{\rho_1 H_1}\boldsymbol{\nabla} \times \boldsymbol{\tau} + \nu\nabla^4\psi_1,$$

$$\left(\frac{\partial}{\partial t} + \frac{\partial \psi_2}{\partial x}\frac{\partial}{\partial y} - \frac{\partial \psi_2}{\partial y}\frac{\partial}{\partial x} + \frac{\partial \psi_2^a}{\partial x}\frac{\partial}{\partial y} - \frac{\partial \psi_2^a}{\partial y}\frac{\partial}{\partial x}\right)[\zeta_2 + \beta y] = \nu\nabla^4\psi_2,$$

Let us introduce the notations

$$\frac{\partial(\psi, \zeta)}{\partial(x, y)} = \frac{\partial \psi}{\partial x}\frac{\partial \zeta}{\partial y} - \frac{\partial \psi}{\partial y}\frac{\partial \zeta}{\partial x},$$

$$\frac{d_i^a(\cdot)}{dt} = \frac{\partial(\cdot)}{\partial t} + \frac{\partial(\psi_i + \psi_i^a, (\cdot))}{\partial(x,y)},$$

$$L_1 = \frac{\sqrt{g'H_1}}{f_0} \ , \ L_2 = \frac{\sqrt{g'H_2}}{f_0},$$

where L_1, L_2 are the Rossby deformation radii within each layer, and including separate Lagrange multipliers for the evolution equation and PV equation we obtain the extended cost function in this case

$$\mathcal{J}^* = \mathcal{J} + \mathcal{J}_{M_\zeta} + \mathcal{J}_{M_\psi} + \mathcal{J}_{\text{bc}} + \mathcal{J}_{\text{ic}} \tag{C.3}$$

$$\mathcal{J} = \frac{1}{2}\int_{t_0}^{t_f}\int_{\Omega_{\text{obs}}} (y - h[q])^2 dSdt, \tag{C.4}$$

$$\mathcal{J}_{M_\zeta} = \int_{t_0}^{t_f}\int_{\mathcal{D}} \theta_1 \left[\frac{d_1^a\zeta_1}{dt} - \frac{1}{\rho_1 H_1}\boldsymbol{\nabla}\times\tau - \nu\nabla^4\psi_1\right] + \theta_2 \left[\frac{d_2^a\zeta_2}{dt} - \nu\nabla^4\psi_1\right] dSdt, \tag{C.5}$$

$$\mathcal{J}_{M_\psi} = \int_{t_0}^{t_f}\int_{\mathcal{D}} \phi_1 \left[\nabla^2\psi_1 - \frac{1}{L_1^2}(\psi_1 - \psi_2) - \zeta_1\right] + \phi_2 \left[\nabla^2\psi_2 - \frac{1}{L_2^2}(\psi_2 - \psi_1) - \zeta_2\right] dSdt, \tag{C.6}$$

$$\mathcal{J}_{\text{bc}} = \sum_i \int_{t_0}^{t_f}\int_{\delta\mathcal{D}} \lambda_i[t,\boldsymbol{x}]\,(\psi_i - k_i) + \mu_i\nabla^2\psi_i\;d\ell dt + a_i\int_{t_0}^{t_f}\int_{\mathcal{D}}(\psi_i - \psi_{i0})\,dS, \tag{C.7}$$

$$\mathcal{J}_{\text{ic}} = \int_{\mathcal{D}} \omega_i[\boldsymbol{x}]\,(\psi_i(t_0,\boldsymbol{x}) - \psi_{i0}[\boldsymbol{x}])\,dS, \tag{C.8}$$

where the $\mathcal{J}_{\text{bc}}$, $\mathcal{J}_{\text{ic}}$ are the explicit form of the penalty terms enforcing the boundary conditions and initial conditions of the model equation and $\mathcal{D}$ is as usual the model domain. The boundary and initial condition penalty terms introduce extra Lagrange multipliers, λ_i, μ_i, a_i and ω_i.

The formal optimality conditions can then be calculated as

$$
\left.\begin{aligned}
\frac{d_1^a q_1}{dt} &= \frac{1}{\rho_1 H_1} \boldsymbol{\nabla} \times \tau + \nu \nabla^4 \psi_1 \\
\frac{d_2^a q_2}{dt} &= \nu \nabla^4 \psi_1 \\
q_1 &= \nabla^2 \psi_1 - \frac{1}{L_1^2}(\psi_1 - \psi_2) \\
q_2 &= \nabla^2 \psi_2 - \frac{1}{L_2^2}(\psi_2 - \psi_1)
\end{aligned}\right\} \quad \text{Forward model} \qquad \text{(C.9)}
$$

$$
\left.\begin{aligned}
\frac{d_1^a \theta_1}{dt} &= -\phi_1 + \frac{\partial h}{\partial q_1}(y - h[q_1, q_2]) \\
\frac{d_2^a \theta_2}{dt} &= -\phi_2 + \frac{\partial h}{\partial q_2}(y - h[q_1, q_2]) \\
\nabla^2 \phi_1 - \frac{1}{L_1^2}\phi_1 + \frac{1}{L_2^2}\phi_2 &= \frac{\partial(\theta_1, q_1)}{\partial(x, y)} + \nu \nabla^4 \theta_1 - a_1 \\
\nabla^2 \phi_2 + \frac{1}{L_1^2}\phi_1 - \frac{1}{L_2^2}\phi_2 &= \frac{\partial(\theta_2, q_2)}{\partial(x, y)} + \nu \nabla^4 \theta_2 - a_2
\end{aligned}\right\} \quad \text{Adjoint model} \qquad \text{(C.10)}
$$

$$
\left.\begin{aligned}
\frac{\partial(\theta_1, q_1)}{\partial(x, y)} &= \beta \frac{\partial \theta_1}{\partial x} \\
\frac{\partial(\theta_2, q_2)}{\partial(x, y)} &= \beta \frac{\partial \theta_2}{\partial x}
\end{aligned}\right\} \quad \text{Gradient w.r.t. } \psi_i^a. \qquad \text{(C.11)}
$$

Thus, given values for q_i, θ_i, it is possible to calculate the gradient of $\mathcal{J}$ with respect to the free variable ψ^a. This can then be used in a gradient search method to minimise the cost function. The complete algorithm is

(i). Integrate (C.9) forward in time over the time window, $[t_o, t_f]$, with an initial condition $\psi(t_0, \boldsymbol{x}) = \psi_0(\boldsymbol{x})$, $q(t_0, \boldsymbol{x}) = q_0(\boldsymbol{x})$ and assimilation forcing, $\psi^a = \psi^{a,i}$.

(ii). Integrate (C.10) backwards in time over the time window, with an initial condition, $\phi(t_f, \boldsymbol{x}) = 0, \theta(t_f, \boldsymbol{x}) = 0$.

(iii). Calculate $\frac{dJ}{d\psi^a}$ using (C.11), use to calculate new $\psi^{a,i+1}$.

(iv). Repeat.

C.2.2   BOUNDARY CONDITIONS

Using a calculus of variations approach and considering a parameterized perturbation $\delta\psi[t_f, \boldsymbol{x}]$, arbitrary beyond a smoothness condition and the spatial boundary conditions, applied to the final state of the system gives an optimality condition

$$
\theta_i[t_f, \boldsymbol{x}] = 0 \;\; \boldsymbol{x} \in \mathcal{D} \qquad \text{(C.12)}
$$

Note that the same process applied at t_0 generates the equation

$$\theta_i[t_0, \boldsymbol{x}] = -\omega_i[\boldsymbol{x}] \ \ \boldsymbol{x} \in \mathcal{D} \tag{C.13}$$

and since optimality in ω requires merely

$$\psi_i(t_0, \boldsymbol{x}) = \psi_{i0}[\boldsymbol{x}] \ \ \boldsymbol{x} \in \mathcal{D} \tag{C.14}$$

we can choose (C.13) as our definition of ω_i and we are free to impose (C.12) as an initial condition at time t_f for the evolution equation for θ_i and then integrate backwards in time.

Collecting together the spatial surface terms in the linearized equation gives the integral in each layer

$$\int_{t_o}^{t_f} \int_{\delta\mathcal{D}} \nu\theta_i\frac{\partial\nabla^2\psi_i'}{\partial n} + \left[\mu - \nu\frac{\partial\theta_i}{\partial n}\right]\nabla^2\psi_i' + \left[\phi_i + \nu\nabla^2\theta_i\right]\frac{\partial\psi_i'}{\partial n}$$
$$- \left[\frac{\partial}{\partial n}\left(\phi_i + \nu\nabla^2\theta_i\right) - \lambda_i + \theta_i\frac{\partial q_i}{\partial s}\right]\psi' - \theta_i\frac{\partial\psi_i + \psi_i^a}{\partial s}q_i' \, d\ell \, dt. \tag{C.15}$$

Since it is possible to define two different functions with, for example,

$$\left.\begin{array}{r}\psi_1 = \psi_2, \\ \dfrac{\partial\psi_1}{\partial n} = \dfrac{\partial\psi_2}{\partial n}, \\ \nabla^2\psi_1 = \nabla^2\psi_2, \\ \dfrac{\partial\nabla^2\psi_1}{\partial n} \neq \dfrac{\partial\nabla^2\psi_2}{\partial n}, \end{array}\right\} \text{ on } \delta\mathcal{D}, \tag{C.16}$$

considering a calculus of variations approach for sufficiently smooth functions gives us four optimality equations

$$\theta_i = 0 \text{ on } \delta\mathcal{D}, \tag{C.17}$$
$$\nu\frac{\partial\theta_i}{\partial n} = \mu_i \text{ on } \delta\mathcal{D}, \tag{C.18}$$
$$\phi_i + \nu\nabla^2\theta_i = 0 \text{ on } \delta\mathcal{D}, \tag{C.19}$$
$$\frac{\partial}{\partial n}\left(\phi_i + \nu\nabla^2\theta_i\right) = \lambda \text{ on } \delta\mathcal{D}, \tag{C.20}$$

where we have simplified using knowledge from the boundary conditions

$$\frac{\partial\psi}{\partial s} = 0 \text{ on } \delta\mathcal{D},$$

$$\frac{\partial q}{\partial s} = 0 \text{ on } \delta\mathcal{D}.$$

Using the free parameter, μ, to solve for equation (C.18) we see that under free slip constraints we have two new boundary conditions on θ, ψ, equations (C.19) and (C.17). The optimality condition on k gives that

$$\int_{\delta\mathcal{D}} \lambda = 0,$$

hence we cannot treat λ as a free parameter to solve equation (C.20), instead we get a third condition,

$$\int_{\delta\mathcal{D}} \frac{\partial}{\partial n}\left(\phi_i + \nu\nabla^2\theta_i\right)d\ell = 0. \tag{C.21}$$

It is this final equation which determines the last free parameters in the system, the constants a_i which appear in the elliptic equations,

$$\nabla^2\phi_1 - \frac{1}{L_1^2}\phi_1 + \frac{1}{L_2^2}\phi_2 = \frac{\partial(\theta_1, q_1)}{\partial(x, y)} + \nu\nabla^4\theta_1 - a_1,$$

$$\nabla^2\phi_2 + \frac{1}{L_1^2}\phi_1 - \frac{1}{L_2^2}\phi_2 = \frac{\partial(\theta_2, q_2)}{\partial(x, y)} + \nu\nabla^4\theta_2 - a_2.$$

C.2.3 NUMERICS

The new forward model is identical to that given in section A.1, except that there is now an additional term,

$$\frac{Dq_i}{Dt} + \boxed{\frac{\partial(\psi_i^a, q_i)}{\partial(x, y)}} = F_i + \nu\nabla^4\psi_i.$$

Since this additional term represents advection of vorticity by a velocity independent of the model velocities there is no reason for it to be stored or calculated on the same model grid. In fact for conservation purposes it is desirable to use a control volumes approach and calculate fluxes on an Arakawa C grid of figure A.1, that is discretizing by central differences the term in the form

$$\frac{\partial(\psi^a, \zeta)}{\partial(x, y)} = \frac{\partial}{\partial y}\left(\frac{\partial\psi^a}{\partial x}\zeta\right) - \frac{\partial}{\partial x}\left(\frac{\partial\psi^a}{\partial y}\zeta\right)$$

$$(\Delta x)^2 \left(\frac{\Delta(\psi^a, \zeta)}{\Delta(x, y)} \right)_E^n = \begin{aligned} &\left(\psi_K^{an} - \psi_M^{an}\right)\left(\zeta_E^n + \zeta_F^n\right)/2 \\ &- \left(\psi_J^{an} - \psi_L^{an}\right)\left(\zeta_D^n + \zeta_E^n\right)/2 \\ &+ \left(\psi_K^{an} - \psi_J^{an}\right)\left(\zeta_E^n + \zeta_B^n\right)/2 \\ &- \left(\psi_M^{an} - \psi_L^{an}\right)\left(\zeta_H^n + \zeta_E^n\right)/2 \end{aligned} \qquad \text{(C.22)}$$

Since the boundary value for the initial forward problem is given by

$$\psi(t_0, \boldsymbol{x}) = \psi_0(\boldsymbol{x})$$

it is impossible to obtain optimality by imposing a boundary condition on $\theta(t_0\boldsymbol{x})$, hence the requirement that the adjoint equation set be integrated backwards from the specified final boundary value

$$\theta(t_f, \boldsymbol{x}) = \boldsymbol{0}.$$

The actual equation set to be solved is very similar to that for the forward model and the solution algorithm is almost the same. The equations are discretized with θ and ϕ values stored on the same grid as ζ and ψ (see figure A.1). Again central differences are used with $\nabla^4 f$ discretized as $\nabla^2[\nabla^2 f]$. Jacobian terms involving properties stored on the A grid are discretized using the Arakawa Jacobian and those involving ψ^a using the form given in (C.22). The gradient information used to update ψ^a requires evaluating

$$\Delta \psi_i^a = -\frac{\partial(\theta_i, \zeta_i)}{\partial(x, y)} - \beta \frac{\partial \theta_i}{\partial x}.$$

Since this information must be stored on the staggered C grid shown in figure A.1 the equation is discretized as

$$(\Delta x)^2 \Delta \psi_L^a = \begin{aligned} &- \left[(\theta_E - \theta_D)/2 + (\theta_G - \theta_H)/2\right] \\ &\times \left[(\zeta_D - \zeta_G)/2 + (\zeta_E - \zeta_H)/2\right] \\ &+ \left[(\theta_D - \theta_G)/2 + (\theta_E - \theta_H)/2\right] \\ &\times \left[(\zeta_E - \zeta_D)/2 + (\zeta_G - \zeta_H)/2\right] \\ &- \beta \left[(\theta_E - \theta_D)/2 + (\theta_G - \theta_H)/2\right] \end{aligned} \qquad \text{(C.23)}$$

This represents a necessary smoothing of the assimilation forcing field, but is flux preserving. Higher order discretizations could be used, but there is unlikely to be major differences to the result. The numerical version of the solution algorithm introduced in the previous section proceeds as follows

(i). Given θ and ϕ at the current and future time levels the evolution equations

$$\frac{\partial \theta_i}{\partial t} = -\frac{\partial(\psi_i, \theta_i)}{\partial(x, y)} - \frac{\partial(\psi_i^a, \theta_i)}{\partial(x, y)} - \phi_i + \frac{\partial h}{\partial \zeta_i}(y^{\text{obs}} - h[\zeta])$$

can be integrated backwards in time to give θ at the previous time level.

(ii). The elliptic equations in ϕ can be rearranged to give

$$\nabla^2(\Phi_1) = \frac{\partial(\theta_1, \zeta_1)}{\partial(x, y)} + \frac{\partial(\theta_2, \zeta_2)}{\partial(x, y)} + \nu \nabla^4(\theta_1 + \theta_2)$$

$$\nabla^2(\Phi_2) - \left(\frac{1}{L_1^2} + \frac{1}{L_2^2}\right)(\Phi_2) = H_2\frac{\partial(\theta_1, \zeta_1)}{\partial(x, y)} - H_1\frac{\partial(\theta_2, \zeta_2)}{\partial(x, y)} + \nu \nabla^4(H_2\theta_1 + H_1\theta_2)$$

where $\Phi_1 = \phi_1 + \phi_2$, $\Phi_2 = H_2\phi_1 - H_1\phi_2$. The two elliptic problems for ϕ at the previous time level are then solved using the NAG multigridded elliptic solver routine as in the forward model code.

(iii). Calculate modifications to the coefficients of the basis functions from the gradient terms using (C.23) and

$$\Delta\psi_f^a = \Delta\psi^a(t)f(t)$$

(iv). Repeat until the beginning of the time window is reached.

APPENDIX D

GLOSSARY OF SYMBOLS AND ACRONYMS

"I am a Bear of Very Little Brain, and long words Bother me."

A.A. Milne
Winnie the Pooh, 1926

TRADITIONAL DATA ASSIMILATION NOTATION

$\boldsymbol{x}_t$, x_t	True state vector/vector field, true scalar field
$\boldsymbol{x}_b$, x_b	Background/forecast state vector/vector field, background sc
$\boldsymbol{x}_a$, x_a	Final analysis state vector/vector field, analysis scalar field
$\boldsymbol{y}$, y	Vector of observations, observation field
$\boldsymbol{H}$, $\boldsymbol{h}$, h	Observation operator matrix, vector, scalar observation field
$\boldsymbol{B}$	Background error covariance matrix
$\boldsymbol{R}$	Observation error covariance matrix

QUASI-GEOSTROPHIC MODEL NOTATION

x	Horizontal distance in the zonal direction
y	Horizontal distance in the meridional direction
H_i	Thickness of the i^{th} layer
ρ_i	Density of the i^{th} layer
f, f_0	Coriolis parameter
β	Beta parameter
ν	Eddy viscosity
$\boldsymbol{\tau} = \left(\tau^{(x)}, \tau^{(y)}\right)$	Wind stress
ψ_i	Geostrophic streamfunction for the i^{th} layer
$u := -\frac{\partial \psi_i}{\partial y}$	Zonal velocity for the i^{th} layer
$v := \frac{\partial \psi_i}{\partial x}$	Meridional velocity for the i^{th} layer
$\xi_i := \frac{\partial v_i}{\partial x} - \frac{\partial u_i}{\partial y}$	Relative vorticity in the i^{th} layer
q_i	Quasi-Geostrophic potential vorticity in the i^{th} layer
$\mathcal{K}$	Total kinetic energy
$\mathcal{A}$	Total available potential energy
$\mathcal{Q}$	Relative enstrophy
$\mathcal{B}$	Baroclinic kinetic energy

DISPLACEMENT ASSIMILATION NOTATION

$\boldsymbol{u}^a, u^a, v^a$	Assimilation velocity
ψ^a	Assimilation streamfunction
s	Timelike assimilation variable (for 3D method)
t	Model time (for 4D method)
$\mathcal{J}$	Discrepancy cost function (3D or 4D)
$\mathcal{M}$	Displacement metric cost function
P	Particle in particle scheme
$\overline{f}$	Low pass filter applied to function f

ACRONYMS

3D-VAR	Three Dimensional Variational assimilation
4D-VAR	Four Dimensional Variational assimilation
4DDA	Four Dimensional Displacement Assimilation
ACM	Analysis Correction Model
BLUE	Best Linear Unbiased Estimator
CH	Cooper and Haines
CTD	Conductivity, Temperature and Depth probe
DA	Data Assimilation
DBCP	Data Buoy Cooperation Panel
ECMWF	The European Centre for Medium Range Weather Forecasting
ERA40	The ECMWF 40 year Reanalysis dataset
FFT	Fast Fourier Transform
FOAM	Forecasting Ocean Assimilation Model
GCM	General Circulation Model
GOOS	Global Ocean Observing System
NWP	Numerical Weather Prediction
PV	Potential Vorticity
QG	Quasi-geostrophic
SOOP	Ship of Opportunity
SSS	Sea Surface Salinity
SST	Sea Surface Temperature
SV	Singular Vector
T–S	Temperature – Salinity
TAO	The Tropical Atmosphere Ocean project
VOS	Voluntary Observing Ships
WOCE	World Ocean Circulation Experiment
XBT	eXpendable BathyThermograph

REFERENCES

"I always have a quotation for everything – it saves original thinking."

Dorothy L. Sayers
Lord Peter Wimsey, Have His Carcase, 1932

S. Akella and I. M. Navon. A comparative study of high resolution advection schemes in the context of data assimilation. *International Journal for Numerical Methods in Fluids*, 51:719–748, 2006.

A. Arakawa and V. R. Lamb. A potential enstrophy and energy conserving scheme for the shallow water equations. *Monthly Weather Review*, 109:18–36, January 1981.

A. Arakawa and V. R. Lamb. Computational design of the basic dynamical processes of the UCLA general circulation model. *Methods in Computational Physics*, 17: 173–265, 1977.

N. Ayoub. Estimation of boundary values in a North Atlantic circulation model using an adjoint method. *Ocean Modelling*, 12(3–4):319–347, 2006.

I. J. Barton. Satellite-derived sea surface temperatures: Current status. *Journal of Geophysical Research*, 100(C5):8777–8790, May 1995.

G. K. Batchelor. *An Introduction to Fluid Dynamics*. Cambridge University Press, 1967.

M. J. Bell. Momentum fluxes in the Bryan/Cox ocean circulation model. *Journal of Atmospheric and Oceanic Technology*, 15(6):1400–1413, December 1998.

J.-D. Benamou and Y. Brenier. A Computational Fluid Mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2001.

J.-D. Benamou, Y. Brenier, and K. Guittet. The Monge-Kantorovitch mass transfer and its Computational Fluid Mechanics formulation. *International Journal of Numerical Methods in Fluids*, 40(1–2):21–30, 2000.

A. F. Bennett. *Inverse Methods in Physical Oceanography*. Cambridge University Press, Cambridge, UK, 1992.

P. Berloff and S. P. Meacham. The dynamics of a simple baroclinic model of the wind-driven circulation. *Journal of Physical Oceanography*, 28(2):361–388, February 1998.

P. S. Berloff and J. C. McWilliams. Large-scale, low-frequency variability in wind-driven ocean gyres. *Journal of Physical Oceanography*, 29:1925–1949, 1999.

C. W. Böning, R. Döscher, and H.-J. Isemer. Monthly mean wind stress and Sverdrup transports in the North Atlantic: A comparison of the Hellerman–Rosenstein and Isemer–Hasse climatologies. *Journal of Physical Oceanography*, 21(2):221–235, February 1991.

C. W. Böning, W. R. Holland, F. O. Bryan, G. Danabasoglu, and J. C. McWilliams. An overlooked problem in model simulations of the thermohaline circulation and heat transport in the Atlantic Ocean. *Journal of Climate*, 8:515–523, 1995.

Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on Pure and Applied Mathematics*, 44(4):375–483, June 1991.

Y. Brenier and W. Gangbo. L^P approximation of maps by diffeomorphisms. *Calculus of Variations*, 16:147–164, 2003.

K. A. Brewster. Phase-correcting data assimilation and application to storm-scale Numerical Weather Prediction. Part i: Method description and simulation testing. *Monthly Weather Review*, 131:480–492, 2002a.

K. A. Brewster. Phase-correcting data assimilation and application to storm-scale Numerical Weather Prediction. Part II: Application to a severe storm outbreak. *Monthly Weather Review*, 131:493–507, 2002b.

K. Bryan. A numerical investigation of a nonlinear model of a wind-driven ocean gyre. *Journal of the Atmospheric Sciences*, 594:594–606, 1963.

R. Buizza and T. N. Palmer. The singular-vector structure of the atmospheric global circulation. *Journal of Atmospheric Science*, 52(9):1434–1456, May 1995.

J. G. Charney, R. G. Feagle, H. Riehl, V. E. Lally, and D. Q. Wark. The feasibility of a global observation and analysis experiment. *Bulletin of the American Meteorological Society*, 47:200–220, 1966.

S. E. Cohn, N. S. Sivakumaran, and R. Todling. A fixed-lag kalman smoother for retrospective data assimilation. *Monthly Weather Review*, 122(12):2838–2867, December 1994.

M. Cooper and K. Haines. Altimetric assimilation with water property conservation. *Journal of Geophysical Research*, 101(C1):1059–1077, 1996.

M. J. P. Cullen. New mathematical developments in atmosphere and ocean dynamics, and their application to computer simulations. In J. Norbury and I. Roulstone, editors, *Large-Scale Atmosphere-Ocean Dynamics I, Analytical Methods and numerical models*, pages 202–287. Cambridge University Press, 2002.

R. Daley. *Atmospheric Data Analysis*. Cambridge University Press, 1991.

G. Danabasoglu, J. C. McWilliams, and P. R. Gent. The role of mesoscale tracer transports in the global ocean circulation. *Science*, 264:1123–1126, 1994.

M. Davey. ENACT enhanced ocean data assimilation and climate prediction, 2005. URL `"www.ecmwf.int/research/EU_projects/ENACT/"`.

R. Douglas. Rearrangements of functions with applications to meteorology and ideal fluid flow. In J. Norbury and I. Roulstone, editors, *Large-Scale Atmosphere-Ocean Dynamics I, Analytical Methods and numerical models*, pages 288–341. Cambridge University Press, 2002.

V. W. Ekman. On the influence of the Earth's rotation on ocean-currents. *Arkiv För Mathematik, Astronomi och Fysik*, 2:1–52, 1905.

H. C. Elman. Multigrid and Krylov subspace methods for the discrete Stokes equations. Technical Report UMIACS-TR-94-76, University of Maryland, Computer Science Department, October 1998.

G. Evensen. Inverse methods and data assimilation in nonlinear ocean models. *Physica D*, 77:108–129, 1994.

D. Ferreira, J. Marshall, and P. Heimbach. Estimating eddy stresses by fitting dynamics to observations using a residual-mean ocean circulation model and its adjoint. *Journal of Physical Oceanography*, 35(10):1891–1910, October 2005.

A. Fitzroy. Outline sketch of the principal varieties and early migrations of the human race. *Transactions of the Ethnological Society of London*, 1:1–11, 1861.

R. Ford, C. C. Pain, M. D. Piggott, A. J. H. Goddard, C. R. E. de Oliveira, and A. P. Umpleby. A nonhydrostatic finite-element model for three-dimensional stratified oceanic flows. part i: Model formulation. *Monthly Weather Review*, 132(12): 2816–2831, 2004.

W. Gangbo and R. J. McCann. The geometry of optimal transportation. *Acta Mathematica*, 177:113–161, 1996.

P. R. Gent and J. R. McWilliams. Isopycnal mixing in ocean circulation models. *Journal of Physical Oceanography*, 20:150–155, 1990.

M. Ghil and P. Malanotte-Rizzoli. Data assimilation in meteorology and oceanography. *Advances in Geophysics*, 33:141–266, 1991.

R. Giering and T. Kaminski. Using TAMC to generate efficient adjoint code: Comparison of automatically generated code for evaluation of first and second order derivatives to hand written code from the Minpack-2 collection. URL `citeseer.ist.psu.edu/496004.html`.

P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic Press, 1981.

B. N. Goswami and J. Shukla. The predictability of a coupled ocean-atmosphere model. *Journal of Climate*, 4(1):3–22, January 1991.

W. J. Gould and J. Turton. Argo - sounding the oceans. *Weather*, 61(1):17–21, January 2006.

K. Haines. A direct method for assimilating sea surface height data into ocean models with adjustments to the deep circulation. *Journal of Physical Oceanography*, 21:843–868, 1991.

K. Haines, R. E. Y. P. Malanotte-Rizzoli, and W. R. Holland. A comparison of two methods for the assimilation of altimeter data into a shallow water model. *Dynamics of Atmospheres and Oceans*, 17:89–133, 1993.

K. Haines, D. Blower, J.-P. Drecourt, C. Liu, A. Vidard, I. Astin, and X. Zhou. Salinity assimilation using $S(T)$: Covariance relationships. *Monthly Weather Review*, 134: 759–771, March 2006.

R. N. Hoffman and C. Grassotti. A technique for assimilating SSM/I observations of marine atmospheric storms: Tests with ECMWF analyses. *Journal of Applied Meteorology*, 35:1177–1188, 1996.

R. N. Hoffman, Z. Liu, J.-F. Louis, and C. Grassotti. Distortion representation of forecast errors. *Monthly Weather Review*, 123:2758–2770, 1995.

W. Holland. The role of mesoscale eddies in the general circulation of the ocean. numerical experiments using a wind-driven quasigeostrophic model. *Journal of Physical Oceanography*, 8:363–392, 1978.

H. E. Hurlbert, D. N. Fox, and E. J. Metzger. Statistical inference of weakly correlated subthermocline fields fields from satellite observations. *Journal of Geophysical Research*, 95:2372–2400, 1990.

K. Ide, P. Courtier, and M. Ghil. Unified notation for data assimilation: Operational, sequential and variational. *Journal of the Meteorological Society of Japan*, 75:189–189, 1997.

D. R. Jackett and T. J. McDougall. A neutral density variable for the world's oceans. *Journal of Physical Oceanography*, 27:237–263, February 1997.

E. Kalnay. *Atmospheric Modelling, Data Assimilation and Predictability*. Cambridge University Press, 2002.

M. Kamachi and J. O'Brien. Continuous data assimilation of drifting buoy trajectories into an equatorial Pacific Ocean model. *Journal of Marine Systems*, 6: 159–178, 1995.

K. B. Katsaros. The aqueous thermal boundary layer. *Boundary-Layer Meteorology*, 18:107–127, 1980.

REFERENCES

E. C. Kent, P. G. Challenor, and P. K. Taylor. A statistical determination of the random observational errors present in voluntary observing ships meteorological reports. *Journal of Atmospheric and Oceanic Technology*, 16(7):905–914, 1999.

C. J. Koblinsky, P. Hildebrand, D. LeVine, F. Pellerano, Y. Chao, W. Wilson, S. Yuch, and G. Lagerloef. Sea surface salinity from space: Science goals and measurement approach. *Radio Science*, 38(4), 2003.

W. Krauss. Comments on the development of our knowledge of the general circulation of the north atlantic ocean. In W. Krauss, editor, *The Warmwatersphere of the North Atlantic Ocean*. Gerbrüder Borntraeger, Berlin, 1996.

L. Kuznetsov, K. Ide, and C. K. R. T. Jones. A method for assimilation of Lagrangian data. *Monthly Weather Review*, 131:2247–2260, October 2003.

S. J. Lavender and S. B. Groom. The SeaWiFS automatic data processing system. *International Journal of Remote Sensing*, 20(6):1051–1056, 1999.

J. R. Ledwell, A. J. Watson, and C. S. Law. Evidence for slow mixing across the pycnocline from an open-ocean tracer release experiment. *Nature*, 364:701–703, 1993.

P. M. Lee. *Bayesian Statistics, An Introduction. Second Edition*. Arnold, 1989.

B. P. Leonard. A stable and accurate convective modelling procedure based on quadratic upstream interpolation. *Computational Methods in Applied Mechanical Engineering*, 19:59–98, 1979.

X. Li, W. Pichel, E. Maturi, P. Clemente-colón, and J. Sapper. Deriving the operational nonlinear multichannel sea surface temperature algorithm coefficients for NOAA-15 AVHRR/3. *International Journal of Remote Sensing*, 22(4):699–704, 2001.

Z. Li and I. Navon. Optimality of variational data assimilation and its relationship with the kalman filter and smoother. *Quarterly Journal of the Royal Meteorological Society*, 127(572):661–683, January 2001.

A. Lorenc, R. Bell, and B. MacPherson. The meteorological office analysis correction data assimilation scheme. *Quarterly Journal of the Royal Meteorological Society*, 117:59–89, 1991.

E. N. Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20:130–141, March 1963.

P. Lynch. The Dolph-Chebyshev window: A simple optimal filter. *Monthly Weather Review*, 125(4):655–660, April 1997.

A. J. Mariano. Contour analysis: A new approach for melding geophysical fields. *Journal of Atmospheric and Oceanic Technology*, 7:285–295, April 1990.

J. Marotzke and J. Willebrand. *The North Atlantic Mean Circulation*. Gerbrüder Borntraeger, Berlin, 1996.

J. E. Marsden and M. J. Hoffman. *Elementary Classical Analysis*. W. H. Freeman, 2nd edition, 1999.

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

F. Martel and C. Wunsch. The North Atlantic circulation in the early 1980s — an estimate from inversion of a finite-difference model. *Journal of Physical Oceanography*, 23:898–924, 1992.

A. Molcard, L. I. Piterbarg, A. Griffa, T. M. Özgökmen, and A. J. Mariano. Assimilation of drifter observations for the reconstruction of the Eulerian circulation field. *Journal of Geophysical Research*, 108(C3):3056–3076, 2003.

F. Molteni, R. Buizza, T. N. Palmer, and T. Petroliagis. The ECMWF ensemble prediction system: Methodology and validation. *Quarterly Journal of the Royal Society for Meteorology*, 122:73–119, 1996.

E. C. Monahan and E. A. Monahan. Trends in drogue design. *Limnology and Oceanography*, 18(6):981–985, November 1973.

W. Munk. Abyssal recipes. *Deep-Sea Research*, 13:707–730, 1966.

W. Munk and C. Wunsch. Abyssal recipes II: energetics of tidal and wind mixing. *Deep-Sea Research. Part 1. Oceanographic research papers*, 45(12):1977–2010, December 1998.

W. H. Munk. On the wind-driven ocean circulation. *Journal of Meteorology*, 7(2):79–93, 1950.

F. Nanson. On North Polar problems. *The Gegraphical Journal*, 30(5):469–487, November 1907.

M. Nodet. Variational assimilation of lagrangian data in oceanography. *Inverse Problems*, 22:245–263, 2006.

Numerical Algorithms Group Ltd. *The NAG F90 Library Manual, Mark 19*, chapter Module 13.2. Oxford, 1998.

D. Olbers and J. Willbrand. The inference of North Atlantic circulation patterns from climatological hydrographic data. *Reviews of Geophysics*, 23:313–356, 1985.

A. Oschlies and J. Willebrand. Assimilation of Geosat altimeter data into an eddy-resolving primitive equation model of the North Atlantic Ocean. *Journal of Geophysical Research*, 101(C6):14175–14190, June 1996.

T. M. Özgökmen, A. Molcard, T. M. Chin, L. I. Piterbarg, and A. Griffa. Assimilation of drifter observations in primitive equation models of midlatitude circulation. *Journal of Geophysical Research*, 108(C7):3238–3254, 2003.

B. B. Parker. Ocean currents: Not really like rivers in the sea. *Mariner's Weather Log*, 42(1):31–36, April 1998.

J. Pedlosky. *Geophysical Fluid Dynamics*. Springer-Verlag, 2nd edition, 1987.

K. L. Polzin, J. M. Toole, J. R. Ledwell, and R. W. Schmitt. Spatial variability of turbulent mixing in the Abyssal Ocean. *Science*, 276:93–96, 1997.

R. W. Ramirez. *The FFT: fundamentals and concepts*. Prentice-Hall, 1985.

S. Ravela, K. Emmanuel, and D. McLaughlin. Multivariate field alignment for data assimilation. *Geophysical Research Abstracts*, 6:6393, 2004.

A. R. Robinson, M. A. Spall, and N. Pinardi. Gulf stream and the dynamics of ring and meander processes. *Journal of Physical Oceanography*, 18:1811–1853, December 1988.

A. Schiller. Journal of marine research. *The mean circulation of the Atlantic Ocean north of 30S determined by the adjoint method*, 53:453–497, 1995.

Z. Sirkes and E. Tziperman. Finite difference of adjoint or adjoint of finite difference? *Monthly Weather Review*, 125:3373–3378, December 1997.

T. M. Smith and R. W. Reynolds. Bias correction for historical sea surface temperatures based on marine mean air temperatures. *Journal of Climate*, 15:73 – 87, January 2002.

H. W. Sorenson. Least-squares estimation from Gauss to kalman. *IEEE Spectrum*, 7: 63–68, July 1970.

D. Stammer. Adjusting internal model errors through ocean state estimation. *Journal of Physical Oceanography*, 35:1143–1153, June 2005.

D. Stammer, C. Wunsch, R. Giering, P. Heimbach, J. Marotzke, A. Adcroft, C. N. Hill, and J. Marshall. Global ocean circulation during 1992–1997, estimated from ocean observations and a general circulation model. *Journal of Geophysical Research*, 107(c9):1–27, 2002.

R. A. Stephen and S. T. Bolmer. Notes for Geoacoustic_TDFD. Technical report, Woods Hole Oceanographic Institution, February 2006.

H. Stommel and F. Schott. The beta spiral and the determination of the absolute velocity field from hydrographic station data. *Deep Sea Research*, 24:325–329, 1977.

H. U. Sverdrup. Wind-driven currents in a baroclinic ocean; with application to the equatorial currents of the eastern Pacific. *Proceedings of the National Academy of Sciences*, 33(11):318–326, 1947.

V. Taillandier, A. Griffa, P.-M. Poulain, and K. Branger. Assimilation of Argo float positions in the north western Mediterranean Sea and impact on ocean circulation simulations. *Geophysical Research Letters*, 33:L11604, 2006.

A. Troccoli and K. Haines. Use of the temperature-salinity relation in a data assimilation context. *Journal of Atmospheric and Oceanic Technology*, 16:2011–2025, 1999.

A. Troccoli, M. Alonso Balmaseda, J. Segschneider, D. L. T. Anderson, K. Haines, T. Stockdale, F. Vitart, and A. D. Fox. Salinity adjustments in the presence of temperature data assimilation. *Monthly Weather Review*, 130:89–102, January 2002.

E. Tziperman, W. C. Thacker, R. B. Long, and S.-M. Hwang. Oceanic data analysis using a general circulation model. Part I: Simulations. *Journal of Physical Oceanography*, 22:1434–1454, December 1992a.

E. Tziperman, W. C. Thacker, R. B. Long, S.-M. Hwang, and S. R. Rintoul. Oceanic data analysis using a general circulation model. Part II: A North Atlantic model. *Journal of Physical Oceanography*, 22:1458–1485, December 1992b.

M. J. Uddstrom, W. R. Gray, R. Murphy, N. A. Oien, and T. Murray. A baysian cloud mask for sea surface temperature retrieval. *Journal of Atmospheric and Oceanic Technology*, 16:117–132, January 1999.

J. Vialard, A. T. Weaver, D. L. T. Anderson, and P. Delecluse. Three- and four- dimensional variational assimilation with a general circulation model of the tropical Pacific Ocean. part ii: Physical validation. *Monthly Weather Review*, 131(7): 1379–1395, July 2003.

F. C. Vossepoel and D. W. Behringer. Impact on sea level assimilation on salinity variability in the western equatorial pacific. *Journal of Physical Oceanography*, 30:1706–1721, 2000.

F. C. Vossepoel, R. W. Reynolds, and L. Miller. Use of sea level observations to estimate salinity variability in the tropical Pacific. *Journal of Atmospheric and Oceanic Technology*, 16:1401–1415, October 1999.

G. Walin. On the relation between sea-surface heat flow and thermal circulation in the ocean. *Tellus*, 34:187–195, 1982.

A. T. Weaver, J. Vialard, and D. L. T. Anderson. Three- and four- dimensional variational assimilation with a general circulation model of the tropical Pacific Ocean. part i: Formulation, internal diagnostics and consistency checks. *Monthly Weather Review*, 131(7):1360–1378, July 2003.

D. J. Webb, B. A. De Cuevas, and C. S. Richmond. Improved advection schemes for ocean models. *Journal of Atmospheric and Oceanic Technology*, pages 1171–1187, October 1998.

J. Wu. Dissipative quasi-geostrophic equations with l^p data. *Electronic Journal of Differential Equations*, 56:1–13, 2001.

L. Wu. A parameter choice method for Tikhonov regularization. *Electronic Transactions on Numerical Analysis*, 16:107–128, 2003.

C. Wunsch. *The Ocean Circulation Inverse Problem*. Cambridge University Press, 1996.

C. Wunsch. Determining the general circulation of the oceans: A preliminary discussion. *Science*, 196(4292):871–875, May 1977.

C. Wunsch and P. Heimbach. Practical global oceanic state estimation. *Physica D*, 230:197–208, 2007.

C. Wunsch and D. Stammer. Global ocean data assimilation and geoid measurements. *Annual Review of Earth and Planetary Science*, 26:219–253, 1998.