




University of Reading

Department of Mathematics and Statistics

Assimilation of future ocean current measurements using satellites

Laura Risley

Thesis submitted for the degree of Doctor of Philosophy

July 2025

Declaration

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

Laura Risley

Abstract

Accurate forecasting of the ocean is crucial for numerous applications, including enhancing numerical weather prediction through coupled ocean-atmosphere systems. Upcoming satellite missions will provide measurements of ocean surface currents from space, offering the potential for significant improvements in forecasting accuracy. Data assimilation is the technique that integrates observations with numerical models to produce ocean forecasts. Existing ocean data assimilation systems rely on assumptions about velocity fields that are sufficient in the absence of ocean current observations. However, the future availability of satellite current measurements necessitates a reassessment of these assumptions and advances in the treatment of velocity fields within the existing systems.

In this thesis, we propose a modification to NEMOVAR, the variational data assimilation system used by the NEMO (Nucleus for European Modelling of the Ocean) model, introducing alternative velocity variables to those currently used. The proposed variables are ageostrophic streamfunction and velocity potential. We design and test a novel transformation to these variables within both a simplified shallow water model and an idealised configuration of the NEMO model. This transformation is analysed both mathematically and numerically in the shallow water model. In implementing this transformation, we conduct an in-depth exploration of the boundary conditions associated with these new variables and investigate a numerical artifact in the solution.

The transformation is further examined in the NEMO configuration, with a focus on integrating these variables into the NEMOVAR assimilation system. Assimilation experiments are performed to evaluate the impact of this new variable formulation on the assimilation of future ocean surface current observations. Results demonstrate that the proposed approach improves the accuracy of the analysis produced by the assimilation of such measurements.

This thesis provides a comprehensive assessment of the theoretical and practical implications of adopting these alternative velocity variables and establishes a strong foundation for their future application within NEMOVAR.

Acknowledgments

I would first like to thank my supervisors, Prof. Amos Lawless, Dr. Matt Martin, and Dr. Anthony Weaver. I feel incredibly fortunate to have had such supportive supervisors. Thank you, Amos, for our weekly meetings and your unwavering encouragement throughout this rollercoaster of a PhD — you always reminded me I was on the right track when I needed it most. I will miss working with all three of you. Thank you to SCENARIO, the Met Office, NCEO, and CERFACS for providing the funding that enabled me to complete this PhD and attend valuable conferences and workshops. Without this funding, the research completed would not have been possible.

I would also like to thank the Data Assimilation Research Centre (DARC). I am truly grateful to have been part of a research group filled with such wonderful and inspiring people. In particular, thank you to Ieuan for being my DARC buddy from day one. We've gone all the way from DA babies to where we are now, and the journey's been way better with a good friend. I want to thank my office, 3L67 (including Jake). You are truly the best office, thank you for all the cake! Thank you to Sarah, Rosie, and Ben. A PhD can be lonely, but having lovely, dog-obsessed friends like you has made it anything but. Speaking of which, a huge thank you to Baxter and Snoop. There is nothing better than a furry cuddle at the end of a tough PhD day, and you two do it best.

Thank you to my siblings — not only for being my best friends, but for always being there. Thank you, Amy, for checking in on me with our daily calls and rants. You've kept me sane. Thank you, Tom, for listening to every presentation I gave throughout the PhD. I wonder if you might now be the most DA-knowledgeable musician out there? Finally, thank you to Scott for feeding me, being there every single day, always making me smile, and patiently listening to everything on my mind. You've often told me I could have done this PhD without you — but I would never have wanted to.

I dedicate this thesis to my momma. You've done so much for me, more than I could ever fully express here (this thesis is already really long!). You've always believed in me, especially when I didn't believe in myself. You inspired me to reach for my dreams. Thank you for everything. I love you more than words can say.

Contents

1	Introduction	19
1.1	Outline	23
2	Mathematical background	26
2.1	Helmholtz theorem	26
2.1.1	Uniqueness and boundary conditions	27
2.2	The matrix problem	28
2.2.1	Tikhonov's regularisation	29
2.2.2	Numerical minimisation algorithms	30
2.3	Shapiro filter	30
2.4	Fast-Fourier transform	31
2.5	Laplacian with Neumann boundary conditions	32
2.6	Error metrics	33
2.7	Summary	34
3	Ocean data assimilation and the control variable transform	36
3.1	Ocean concepts	37
3.1.1	Geostrophic balance	39
3.2	Variational data assimilation	40
3.2.1	Incremental 3D VAR	42
3.2.2	Incremental 3D FGAT	43
3.3	Control variable transform	44
3.3.1	Implementing the CVT within incremental 3D VAR	45
3.4	Operational data assimilation systems	47
3.4.1	Atmospheric data assimilation	47

3.4.2	Oceanic data assimilation	48
3.4.3	NEMOVAR	52
3.5	Proposed control variables	55
3.5.1	Treatment of boundary conditions	56
3.6	Summary	58
4	Alternative velocity control variables using the shallow water equations	60
4.1	Shallow water model	60
4.1.1	Discretisation	62
4.2	CVT using the shallow water equations	65
4.2.1	U - Transform	66
4.2.2	T - Transform	66
4.2.3	The discrete transformation to streamfunction and velocity potential	67
4.2.4	Boundary conditions for streamfunction	68
4.3	Numerical implication	72
4.3.1	Method	72
4.3.2	Checkerboard pattern	74
4.3.3	Spatial correlations	76
4.3.4	The Shapiro filter	79
4.4	Summary and conclusions	82
5	Alternative velocity control variables in NEMOVAR	85
5.1	Nucleus for European Modelling of the Ocean	86
5.1.1	Coordinate system in NEMO	86
5.1.2	GYRE configuration	86
5.2	Variable transformations	89
5.2.1	Balance operator	89
5.2.2	The discrete transformation to streamfunction and velocity potential	91
5.3	Performing the discrete transformation	93
5.3.1	Method	93
5.3.2	Transformation experiments	95
5.3.3	Summary	99
5.4	Checkerboard pattern	100

5.4.1	The Shapiro filter	102
5.5	Summary and conclusions	111
6	Assimilation experiments in NEMOVAR	114
6.1	Experiment overview	115
6.1.1	Spin-up procedure	115
6.1.2	Experiment details	116
6.1.3	Experiment aims	116
6.2	Model behaviour	118
6.3	Data assimilation set-up	121
6.3.1	Observations	123
6.3.2	Background error covariance settings	125
6.4	Single observation experiments	126
6.5	Assimilation experiments	131
6.5.1	Methods of comparison	132
6.5.2	One day assimilation	132
6.6	Year long assimilation	134
6.6.1	Errors averaged over the domain	134
6.6.2	RMSE averaged over time	148
6.6.3	Summary	150
6.7	Vertical velocities	151
6.8	Month long assimilation	157
6.8.1	Sensitivity to standard deviations	157
6.8.2	Summary	161
6.9	Summary and conclusions	161
7	Conclusion	163
7.1	Research questions	164
7.2	Main conclusions	168
7.3	Future work	168
	Appendices	171

A	Cross-correlations of control variables using the SWEs	171
A.1	Method	171
A.2	Results and discussion	173
B	Geostrophic balance	174
C	Approaches to implementing the CVT for the shallow water model	175
C.1	One system approach	175
C.1.1	Adding the Shapiro filter	175
C.1.2	Alternative filtering approach	176
C.2	Two system approach	177
D	A Lagrangian approach	184
D.1	Continuous equations	184
D.1.1	Tikhonov's regularisation	187
D.2	Discrete equations	187
E	Testing of the adjoint and gradient	194
E.1	The adjoint test	194
E.2	The gradient test	194
F	Non-divergence of the velocity generated by streamfunction	196
F.1	Away from the boundaries	198
F.2	Near the boundaries	198
G	FFT boundary conditions	201
H	No velocity observations	203
I	Single observation experiments	206

List of Figures

3.1	Variational data assimilation schemes	41
3.2	Configurations of the Arakawa C-grid for the placement of streamfunction	57
4.1	2D shallow water equations	61
4.2	Arakawa C-grid	63
4.3	Arakawa C-grid with the averaged values of streamfunction located at the corner	70
4.4	Arakawa C-grid configuration with extended land	71
4.5	Model output of the shallow water model	74
4.6	Ageostrophic streamfunction and velocity potential increments for the SWEs	75
4.7	Reconstructed ageostrophic velocity increments for the SWEs	76
4.8	Locations of the points chosen for the dominant signal analysis	77
4.9	Dominant signal analysis for zonal velocity	78
4.10	Spatial correlations for streamfunction	79
4.11	Workflow to finding the ageostrophic streamfunction and velocity potential increments with the Shapiro filter after the transformation	81
4.12	Reconstructed ageostrophic velocity increments for the SWEs with the Shapiro filter	82
5.1	NEMO coordinate system	87
5.2	NEMO horizontal grid	89
5.3	NEMO lateral boundary conditions	89
5.4	Model increment workflow for the GYRE configuration	94
5.5	Surface horizontal velocity increments from the GYRE configuration	94
5.6	RMSE of the reconstructed horizontal velocity errors at the surface	97

5.7	Average RMSE of the reconstructed horizontal velocity across all depths . . .	98
5.8	Norm of the residual at the surface and number of iterations until convergence across all depths	98
5.9	Surface streamfunction, velocity potential and reconstructed velocities . . .	100
5.10	Horizontal lines for the FFT	101
5.11	FFT spectrums for no filter	103
5.12	FFT spectrums for one iteration of the Shapiro filter	105
5.13	FFT spectrums for five iteration of the Shapiro filter	106
5.14	Surface streamfunction and reconstructed velocities with one iteration of the Shapiro filter	107
5.15	Vertical cross-section of streamfunction and reconstructed velocities	108
5.16	RMSE of the reconstructed velocities at all the depths, including filter experiments	109
5.17	2D FFT spectrums of streamfunction at the surface	110
5.18	Surface velocity reconstruction errors for the high frequency test	111
6.1	Spin-up workflow for the GYRE configuration	116
6.2	Nature run and free-run model fields on the first day	119
6.3	Mean of the nature run and free-run model fields over a year	120
6.4	Standard deviation of the nature run and free-run model fields over a year	122
6.5	Coverage of the observations used in the assimilation experiments	124
6.6	Single observation experiment locations	127
6.7	Velocity increments of the single observation experiments	129
6.8	Speed increments of the single observation experiments	130
6.9	Surface velocity increments for control and initial assimilation experiments	133
6.10	RMSE calculated over the whole domain at the surface	135
6.11	Temperature difference between model runs.	137
6.12	MAE calculated over the whole domain at the surface	138
6.13	Standard deviation calculated over the whole domain at the surface	140
6.14	Error and observations in August	142
6.15	Time averaged RMSE for zonal velocity in July and September.	143
6.16	RMSE across the depths over one year	145

6.17	MAE across the depths over one year	146
6.18	Standard deviation across the depths over one year	147
6.19	Difference in RMSE across the domain for the control experiments	148
6.20	Difference in RMSE across the domain	149
6.21	RMSE of the vertical velocities	152
6.22	Spatial mean the vertical velocities	153
6.23	Spatial mean the diagnosed vertical velocity increments	154
6.24	Spatial mean the vertical velocity increments	155
6.25	Spatial mean the horizontal velocity increments	156
6.26	RMSE for the month long experiments	158
6.27	MAE for the month long experiments	159
6.28	RMSE and MAE for the month long experiments across the depths	160
A.1	Cross-correlations of control variables	172
C.1	One system approach to finding the ageostrophic streamfunction and velocity potential increments	180
C.2	One system approach to finding the ageostrophic streamfunction and velocity potential increments with the Shapiro filter after the transformation	181
C.3	One system approach to finding the ageostrophic streamfunction and velocity potential increments with the Shapiro filter within the transformation	182
C.4	Two system approach to finding the ageostrophic streamfunction and velocity potential increments	183
D.1	A visual demonstration of the discrete Laplacian operator used to solve the vorticity equation for streamfunction	193
G.1	FFT of periodic and non-periodic signals.	202
I.1	Single observation velocity increments - inactive region	207
I.2	Single observation velocity increments - near the boundary	208
I.3	Single observation speed increments - inactive region	209
I.4	Single observation speed increments - near the boundary	210

List of Tables

3.1	Ocean data assimilation systems	49
4.1	Parameter values for the SWEs	73
5.1	Convergence and regularisation experiments for the GYRE configuration .	96
5.2	Shapiro filter experiments for the GYRE configuration	107
6.1	Description of the assimilation experiment settings	117
6.2	Observation error standard deviations	124
6.3	Background error covariance settings	126

Notation

Data assimilation

$\tilde{z}$	Control variable (assumed completely uncorrelated background errors)
δx	Increment
Λ	Spatial correlation matrix
$\mathbf{B}$	Background error covariance matrix
$\mathbf{d}$	Residual vector
$\mathbf{H}$	Linearised observation operator
$\mathbf{R}$	Observation error covariance matrix
$\mathbf{y}$	Vector of observations
$\mathcal{T}$	T-Transform
$\mathcal{U}$	U-Transform
h	Observation operator
x	Model variable
x^a	Analysis
x^b	Background state
z	Control variable (assumed uncorrelated variables, but with spatial correlations remaining)

Mathematics

α Time weighting parameter for the SISL scheme

Δ_x Discretisation of $\partial/\partial x$

Δ_x^2 Discretisation of $\partial^2/\partial x^2$

Δ_y Discretisation of $\partial/\partial y$

Δ_y^2 Discretisation of $\partial^2/\partial y^2$

$\mathbb{R}$ Real numbers

$\mathbf{L}$ Discrete Laplacian operator

$\mathcal{F}$ Shapiro filter

$\mathcal{J}(\cdot)$ Cost function

$\mathcal{L}$ Lagrange function

μ Regularisation parameter

∇ Gradient

∇^2 Laplacian

∂ Partial derivative

Ocean data assimilation

Σ Matrix of standard deviations

$\mathbf{C}$ Correlation matrix

$\mathbf{K}$ Linearised balance operator

$\mathbf{W}$ Weighting matrix

K Balance operator

Ocean model

$\cdot_B$	Balanced component
$\cdot_U$	Unbalanced component
χ	Velocity potential
Δt	Time step length
$\Delta x, \Delta y$	Spatial step lengths
η	Sea surface height
γ	Linear drag coefficient
$\hat{\mathcal{D}}$	Discretisation of divergence
$\hat{\xi}$	Discretisation of vorticity
λ	Longitude
$\mathbf{F}$	Forcing
$\mathbf{v}$	Velocity vector
$\mathbf{v}_\chi$	Velocity generated by velocity potential
$\mathbf{v}_\psi$	Velocity generated by streamfunction
$\mathbf{x}$	Spatial coordinate, $\mathbf{x} = (x, y)$
$\mathcal{D}$	Divergence
$\nabla \cdot \mathbf{v}$	Horizontal divergence
$\nabla \eta$	Pressure gradient
Ω	Angular speed of the earth
ψ	Streamfunction
ρ	Density
τ	Wind stress

τ_i	Mean wind intensity
τ_s	Seasonal oscillation intensity
φ	Latitude
ξ	Vorticity
a	Radius of the earth
f	Coriolis force
$f_0 + \beta \mathbf{y}$	Coriolis parameter on a β -plane
g	Gravitational acceleration
H	Resting depth of shallow water model
L	Shallow water domain length
p	Pressure
R	Rossby number
S	Salinity
T	Temperature
t	Time
u	Zonal velocity
v	Meridional velocity
w	Vertical velocity

Acronyms and abbreviations

3D FGAT 3D First guess at appropriate time

3D VAR Three-dimensional variational data assimilation

4D VAR Four-dimensional variational data assimilation

ADCP Acoustic Doppler Current Profilers

BODAS Bluelink Ocean Data Assimilation System

CERFACS Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique

CFL Courant-Friedrichs-Lewy

CG Conjugate gradient

CVs Control variables

CVT Control variable transform

DA Data assimilation

DFT Discrete Fourier transform

ECCC Environment and Climate Change Canada

ECMWF European Centre for Medium-range Weather Forecasts

EnKF Ensemble Kalman filter

EOFs Empirical orthogonal functions

ESA European Space Agency

FFT Fast-Fourier transform

FNMO Fleet Numerical Meteorology and Oceanography Center

FOAM Forecast Ocean Assimilation Model

HF High frequency

IAU Incremental Analysis Update

INRIA National Institute of Research in Digital Science and Technology

JMA Japan Meteorological Agency

JPL Jet Propulsion Laboratory

MAE Mean absolute error

MOVE Multivariate Ocean Variational Estimation

MSE Mean square error

NAVOCEANO Naval Oceanographic Office

NCEO National Centre of Earth Observation

NCODA Navy Coupled Ocean Data Assimilation

NEMO Nucleus for European Modelling of the Ocean

NEMOVAR NEMO Variational Data Assimilation system

NWP Numerical weather prediction

ODYSEA Ocean Dynamics and Surface Exchange with the Atmosphere

OSSEs Observing system simulation experiments

QCM Quick covs method

RMS Root mean square

RMSE Root mean square error

ROMS Regional Ocean Modelling System

SAM Système d'Assimilation Mercator

SEEK Singular evolutive extended Kalman filter

SISL Semi-implicit semi-Lagrangian scheme

SKIM Sea surface Kinematics Multiscale monitoring

SLA Sea level anomaly

SSH Sea surface height

SST Sea surface temperature

SVD Singular value decomposition

SWEs Shallow water equations

TSCV Total surface current velocities

Chapter 1

Introduction

Ocean forecasts play a vital role in a wide range of applications, including coastal management (Dré villon et al., 2008), the understanding of marine animal species (Dickey, 2003) and numerical weather prediction (NWP) (Lea et al., 2015; de Rosnay et al., 2022). The accuracy of these forecasts relies heavily on the data assimilation process, which combines observational data with numerical ocean models, weighted by their uncertainties, to give optimal initial conditions (Kalnay, 2003). Variational data assimilation is commonly used in ocean applications (Anderson et al., 1996) and treats the data assimilation problem using a least-squares approach, where a cost function is minimised to estimate the most likely ocean state. Within variational data assimilation the model variables are often transformed into alternative variables to simplify the assumptions associated with the data assimilation problem. In this thesis we look at an alternative transformation to be used in variational ocean data assimilation in preparation for future ocean current measurements.

In this work, we are particularly interested in the treatment of ocean currents. Accurate forecasts of ocean currents are critical for marine safety and offshore operations (Waters et al., 2024b). However, there is a limited availability of surface current velocity measurements on a global scale. Some observations are available near coastlines, primarily from high frequency (HF) radar systems. These ocean currents derived from HF radar allow some coastal data assimilation systems to assimilate HF radar currents (Hoteit et al., 2009; Ngodock et al., 2015). Near-surface currents can also be derived from the changing positions of surface drifters, (Helber et al., 2023; Smith et al., 2023), but these

measurements are sparse compared to the scales of variability in ocean currents. Globally, this limited availability of ocean current measurements means they are generally not assimilated operationally (Waters et al., 2024a). The ocean currents are instead estimated through their balance relationships with other ocean variables which are observed.

Recently proposed satellite missions, such as Harmony, the next European Space Agency (ESA) Earth Explorer 10 mission (Stoffelen, 2022), aim to accurately measure ocean surface current velocities from space, presenting an opportunity to assimilate velocity current data. To fully leverage these potential future measurements, it is imperative that we improve current ocean data assimilation systems to effectively integrate these observations. Idealised experiments in Waters et al. (2024b) demonstrate the potential impact of satellite surface current data, showing significant improvements in the quality of ocean analyses and forecasts. The aim of this work is to further improve the assimilation of such data.

To solve the variational data assimilation we often perform a transformation from the model variables to control variables, that we assume to be uncorrelated. This is known as the control variable transform (CVT) (Lorenc et al., 2000; Parrish and Derber, 1992). In this thesis, we focus on the velocity control variables used in NEMOVAR. NEMOVAR is a variational data assimilation software used with the ‘Nucleus for European Modelling of the Ocean’ (NEMO) model. It is used operationally at the UK Met Office (Barbosa Aguiar et al., 2024) and the European Centre for Medium-Range Weather Forecasts (ECMWF) (Balmaseda et al., 2013).

Due to the limited observations of horizontal velocities in the ocean, the velocity variables selected in the control space currently have no impact on the assimilation of observed variables (e.g. temperature, salinity and sea surface height (SSH), Lea et al. (2014)). The velocities are updated based on balance relationships with observed variables. When actively assimilating these future surface current measurements, the transformed velocity variables will have a much greater impact on the assimilation, as the velocity control variables will directly update the velocities. In NEMOVAR, the control variable transform involves using the inverse of a balance operator to decompose state variables into balanced and unbalanced components. The transformed velocity variables are the ageostrophic components of the horizontal velocity vector (Weaver et al., 2005). However, since the ageostrophic velocities are components of the same vector, we expect

them to be highly correlated, which undermines the assumption of uncorrelated variables. Therefore, we seek alternative velocity control variables to be used in variational ocean data assimilation.

Commonly in atmospheric data assimilation the velocities are split into their non-divergent (streamfunction) and irrotational (velocity potential) components, using Helmholtz Theorem. The velocity control variables are then based on these components. The uncorrelated assumption is more suitable for these control variables than the velocities themselves (Daley, 1993). The application of the atmospheric methods for the velocity CVT is challenging in the ocean due to the presence of coastlines. These introduce complications when solving elliptic equations, as they require handling complex boundary conditions. The transformation from the velocities to their non-divergent and irrotational parts has been discussed in previous studies such as Lynch (1989) and Watterson (2001). They propose various boundary conditions to impose on streamfunction and velocity potential. Li et al. (2006) notably address the difficulty in specifying explicit boundary conditions to streamfunction at coastal boundaries. They suggest a method that involves adjusting the discretised location of the streamfunction on the computational grid to avoid specifying explicit boundary conditions.

Beyond the issue of boundary conditions, a CVT based on Helmholtz decomposition may help reduce spurious vertical motions produced by ocean data assimilation. Such motions are a known issue in ocean models (Ford and Barciela, 2017). By separating the velocity field into non-divergent and irrotational components, the decomposition has the potential to better control these vertical motions introduced by data assimilation.

In this work, we demonstrate using ageostrophic streamfunction and velocity potential as velocity control variables in the ocean, by adopting the method proposed by Li et al. (2006). This thesis investigates how we implement these alternative control variables in NEMOVAR and whether they have the potential to improve the analyses produced. We aim to answer the following research questions:

1. **What numerical challenges and implications arise from the discrete transformation of the velocities to their irrotational and non-divergent parts?**

- We implement the transformation to ageostrophic streamfunction and velocity

potential within an idealised shallow water model framework.

- We identify two distinct approaches for discretising this transformation and provide a detailed analysis of their implications on the boundary conditions, an aspect that has been largely overlooked in previous studies.
- We uncover a previously unreported issue of grid-scale noise arising from the transformation. We assess its impact on the data assimilation performance and discuss methods to mitigate it.

2. How can the discretisation of this transformation be extended from an idealised framework to application within a realistic ocean model?

- We implement the transformed velocity variables within an operational variational data assimilation system (NEMOVAR) and integrate them into a complex multivariate balance operator.
- We conduct a detailed sensitivity analysis of key transformation parameters, providing guidance on their optimal settings and assessing their influence on the transformation.
- We discover that the grid-scale noise identified in the idealised model also emerges in this more complex, realistic setting. We offer practical solutions and implementation strategies to address this issue in operational contexts.

3. What are the implications of using these alternative velocity variables in variational data assimilation, and how do these variables respond to the assimilation of velocity observations?

- We carry out the inaugural assimilation experiments using the transformed velocity variables within NEMOVAR, using idealised ocean current observations.
- We compare the performance of the transformed variables against the current operational CVT and demonstrate their potential to improve the quality of the resulting analysis.
- We perform a novel investigation into the influence of these variables on vertical velocity fields, offering new insights into their role.

These research questions lead us to three key results:

1. We **design** a novel control variable transformation using alternative velocity variables for use in NEMOVAR.
2. We **discover** new implications of implementing this transformation that emerge from a detailed mathematical analysis.
3. We **demonstrate** the potential for improving the assimilation of future ocean current measurements using the new control variable transformation.

1.1 Outline

This thesis is structured as follows.

Chapter 2 will introduce some of the key mathematical concepts that are essential to this work. We begin with Helmholtz theorem which underpins the variable transformations used throughout the thesis. We then present the general matrix inversion problem, highlighting issues that lead to ill-posed problems. This leads onto the discussion of regularisation, which can be used to deal with ill-posed problems. Following this, we explore other important mathematical tools, including filtering, Fourier transforms, the Laplacian operator and error metrics — all of which play an important role in this work.

Chapter 3 will introduce ocean data assimilation. This starts with a brief overview of key oceanographic concepts, including geostrophic balance. We then explain the importance of ocean forecasts and introduce data assimilation. Among the various data assimilation techniques available, this work focuses on variational data assimilation, particularly the incremental formulation. We represent the error statistics using control variable transformations, which we introduce. We then specify the current control variables that are used in atmospheric and ocean data assimilation, providing a discussion of the current limitations of the velocity control variables in ocean data assimilation. Finally, we present the alternative control variable transform that we propose to improve ocean variational data assimilation.

Chapter 4 introduces the shallow water equations, which are used as a simplified model to investigate the proposed velocity control variables. We describe the model equations and detail their discretisation using a semi-implicit semi-Lagrangian scheme.

We then outline the process of applying the proposed control variable transformation within the shallow water model, highlighting the complexities introduced by boundary conditions. During the computation of the alternative velocity control variables, we encounter a numerical issue, known as the checkerboard pattern. This chapter presents the issue both analytically and numerically using the shallow water model. We propose a potential solution to this numerical issue — while acknowledging its own limitations.

Chapter 5 extends the investigation of chapter 4 to NEMO. Here, we implement these alternative velocity control variables in NEMOVAR and evaluate the transformation using a configuration of the NEMO model known as the GYRE configuration. We firstly introduce the NEMO model, the NEMOVAR data assimilation system, and the GYRE configuration. Next, we describe the implementation of the control variables within the NEMOVAR balance operator. We investigate the sensitivity of this transformation to various parameters and identify the settings that are most suitable for this study. The transformation to the alternative velocity control variables is then performed, and we demonstrate that the checkerboard issue observed in chapter 4 also arises in NEMOVAR. This is shown numerically using the GYRE configuration. Finally, we discuss the proposed solution to this issue and evaluate its sensitivity and limitations.

Chapter 6 builds on the work of chapter 5 and uses the GYRE configuration to perform idealised assimilation experiments in NEMOVAR with the proposed velocity control variables. We describe the experiment aims and set-up, detailing the observations, background, nature run, control and assimilation experiments. We also discuss the methods of comparison used to evaluate the analyses produced by the experiments. We perform assimilation experiments with the previous and alternative control variables, these include single observation experiments. The aim is to demonstrate that the alternative control variables have the potential to improve the analysis. We evaluate the effect of the proposed control variables on the magnitudes of the vertical velocities, which is a potential benefit of these new variables. Additionally, we investigate the sensitivity of the analysis results to the choice of error standard deviations.

Chapter 7 contains the main conclusions of this work. We summarise the key findings for the analytical and numerical experiments performed using the shallow water model and GYRE configuration. From these we draw conclusions regarding the appropriateness of the alternative velocity control variables proposed for NEMOVAR. Finally, we offer

suggestions on further work to be done to progress the implementation of these variables operationally.

We now introduce the mathematical concepts integral to this thesis.

Chapter 2

Mathematical background

In this chapter, we introduce the key mathematical concepts that underpin the methods and analyses presented throughout this work. This chapter aims to equip the reader with the foundational tools necessary to fully engage with the mathematical framework of the project.

2.1 Helmholtz theorem

A core principle used throughout this work is the decomposition of a vector field into its rotational (non-divergent) and divergent (irrotational) components, a result formalised by the Helmholtz Theorem. The Helmholtz decomposition was introduced by Helmholtz (1858). This foundational theorem in fluid dynamics is particularly valuable, as it enables the simplification of vector fields by isolating properties such as incompressibility and vorticity (Bhatia et al., 2012). Beyond fluid dynamics, Helmholtz theorem finds wide-ranging applications in fields including astrophysics, robotics, computer graphics, and imaging (Bhatia et al., 2012). Here, we present the generic Helmholtz Theorem for a field and then specify how this can be applied to a velocity field.

Theorem 2.1.1 (Helmholtz Theorem) (*Denaro, 2003*) *Any smooth vector field $\mathbf{f} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ can be expressed as the sum of the gradient of a scalar potential, D , and the curl of a vector potential, $\vec{R}$,*

$$\mathbf{f} = \nabla D + \nabla \times \vec{R} \tag{2.1}$$

where ∇D is irrotational ($\nabla \times \nabla D = 0$) and $\nabla \times \vec{R}$ is non-divergent ($\nabla \cdot \nabla \times \vec{R} = 0$).

2.1.1 Uniqueness and boundary conditions

It is important to consider whether Helmholtz Theorem can be applied uniquely to a particular field. In domains where the decomposition is not unique, a component of the flow exists that is both non-divergent and irrotational, this is called the harmonic component (Bhatia et al., 2012; Schoder et al., 2020).

The uniqueness of the Helmholtz decomposition was proven in unbounded domains by Blumenthal (1905) when a vanishing condition at infinity is imposed. This states that the field vanishes at infinity and as such the irrotational, non-divergent and harmonic component must also do the same. Blumenthal (1905) shows that the only harmonic component to vanish at infinity is the zero function. This follows from the maximum principle. As such the decomposition is unique.

In a bounded domain, a non-zero harmonic function may be present. However, certain boundary conditions can enforce a unique decomposition. Bhatia et al. (2012) provide an in-depth analysis of the uniqueness of Helmholtz Theorem. In particular they state that a unique decomposition can be obtained if either of these boundary conditions are satisfied:

- the irrotational component is normal to the boundary, i.e. $\nabla D \times \vec{n} = 0$;
- the non-divergent component is parallel to the boundary, i.e. $\nabla \times \vec{R} \cdot \vec{n} = 0$;

where $\vec{n}$ is the outward normal to the boundary. These boundary conditions can be simplified and uniqueness is also satisfied if there is no flow on the boundary. This theorem can be extended to a velocity field. Therefore, uniqueness would require zero velocity at the boundary.

Theorem 2.1.2 (Helmholtz Theorem for a horizontal velocity field)

The horizontal velocity field, $\mathbf{v} \equiv (u, v)$ - where u is the zonal velocity and v is the meridional velocity - can be decomposed into rotational and divergent parts,

$$\mathbf{v} = \mathbf{v}_\psi + \mathbf{v}_\chi \quad (2.2)$$

where

$$\mathbf{v}_\psi = [\mathbf{k} \times \nabla \psi]_h, \quad \mathbf{v}_\chi = [\nabla \chi]_h. \quad (2.3)$$

Here ψ is the streamfunction, χ is the velocity potential and $\mathbf{k}$ is the unit vector in the vertical. As ∇ is a 3D gradient operator, we take the horizontal components of each of these terms, denoted by ${}_h$.

This is equivalent to,

$$u = -\frac{\partial\psi}{\partial y} + \frac{\partial\chi}{\partial x} \quad (2.4)$$

and

$$v = \frac{\partial\psi}{\partial x} + \frac{\partial\chi}{\partial y} \quad (2.5)$$

where (x, y) is the 2D spatial coordinate. From Helmholtz Theorem we can deduce the Poisson equations for divergence, $\mathcal{D}$, and vorticity, ξ ,

$$\mathcal{D} = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = \nabla_h^2 \chi, \quad (2.6)$$

$$\xi = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} = \nabla_h^2 \psi. \quad (2.7)$$

In this work, we apply the Helmholtz decomposition to partition the oceanic velocity field into its irrotational and non-divergent components. We further examine the conditions under which this decomposition is unique within the ocean domains considered.

2.2 The matrix problem

Throughout this work, we encounter discrete linear systems that arise in the form of matrix equations. Consider the generic matrix problem,

$$\mathbf{A}\mathbf{x} = \mathbf{b} \quad (2.8)$$

where $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^m$. The direct problem of computing $\mathbf{b}$ from $\mathbf{x}$ is usually well-defined. However, the inverse problem, solving (2.8) for $\mathbf{x}$ given $\mathbf{b}$, requires more careful consideration. If $\mathbf{A}$ is a square matrix, i.e. $m = n$, a unique solution to the inverse problem exists if and only if,

- $\mathbf{A}$ is non-singular,
- the determinant of $\mathbf{A}$ is non-zero, i.e. $\det(\mathbf{A}) \neq 0$, $\mathbf{A}$ is invertible,

- $\mathbf{A}$ is full rank, i.e. $\text{rank}(\mathbf{A}) = n$,
- the kernel of $\mathbf{A} = 0$,
- $\mathbf{A}$ has linearly independent rows and columns.

These statements are all equivalent. If $m \neq n$, we can redefine the matrix problem, (2.8), as a least squares problem (Golub et al., 1999),

$$\min_x \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2. \quad (2.9)$$

Direct problems are typically well-posed whereas inverse problems are often ill-posed (Groetsch, 1993), whereby small changes in $\mathbf{b}$ will lead to large changes in the solution, $\mathbf{x}$. Hansen (1997) discusses different methods used to deal with ill-posed problems and in this work we use regularisation. This replaces an ill-posed linear system with a nearby system that is less sensitive to perturbations. In the next section we discuss Tikhonov's regularisation.

2.2.1 Tikhonov's regularisation

Regularisation methods are used to deal with non-unique or underdetermined problems. Tikhonov (1963) and Phillips (1962) independently introduced Tikhonov's regularisation, which involves adding a term to the linear system (2.9) to give

$$\min_x \left\{ \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \mu \|\mathbf{x}\|_2^2 \right\} \quad (2.10)$$

where μ is a positive and finite regularisation parameter (Hanke and Hansen, 1993). This parameter controls the weight of the regularisation term relative to the residual norm (Hansen, 1997), determining how sensitive the regularised solution is to the error in $\mathbf{b}$ (Calvetti and Reichel, 2003). If μ is chosen too small, the regularised equation remains too close to the original ill-posed problem, retaining its instability. Conversely, if μ is too large, the solution becomes overly smooth and may no longer accurately reflect the original problem. Determining an optimal regularisation parameter is a challenging task (Groetsch, 1993, 1984).

There are multiple methods discussed in the literature to determine the value of the regularisation parameter and this determination is problem dependent. A particular

method, called the L-curve, is concluded as the most robust method according to Hansen (1997). The L-curve involves plotting the logarithm of the regularised solution against the logarithm of the residual norm for multiple values of μ . The result is an L-shaped curve, where the vertex represents the optimal regularisation parameter, which is a compromise between the perturbation and regularisation error. While this thesis does not make use of this method, sensitivity tests are applied instead. The method is recommended for consideration in future studies.

2.2.2 Numerical minimisation algorithms

As discussed, we look to minimise linear systems such as (2.10). This requires numerical minimisation algorithms. One particular algorithm is known as the conjugate gradient (CG) method.

The CG method is a well-known algorithm for solving symmetric, positive definite systems of linear equations in the form of (2.8) and for minimising unconstrained non-linear functions, combining simplicity with computational efficiency (Nocedal and Wright, 2006; Nazareth, 2009; Shewchuk et al., 1994). The CG method improves upon the steepest descent approach by generating conjugate directions, leading to faster convergence. It is especially effective for large, sparse systems where direct methods like Cholesky decomposition are impractical due to memory constraints, making it a popular and practical approach in computational linear algebra and non-linear optimisation (Nazareth, 2009; Shewchuk et al., 1994).

In this work, we use the CG method to minimize the linear systems that arise. This choice is motivated by the large, sparse linear systems typical in data assimilation, for which CG offers both efficiency and scalability compared to alternative methods.

2.3 Shapiro filter

Numerical solutions to non-linear systems in meteorological and oceanic applications are often prone to the spurious growth of short-wavelength components, which necessitates the use of smoothing or filtering techniques to maintain stability and accuracy (Phillips, 1959; Shapiro, 1970).

Spatial filters are widely employed in computational physics to smooth data fields and control numerical noise (Falissard, 2017). In particular, low-pass filters are commonly used in applications such as weather forecasting and computational fluid dynamics to selectively damp high-frequency components while preserving physically relevant structures (Falissard, 2015, 2017). Shapiro filters are standard in various high-resolution models across geophysical and engineering disciplines (Falissard, 2013, 2017), these are high-order explicit filters capable of removing spurious grid-scale oscillations without affecting the large-scale features of the solution (Shapiro, 1970, 1971, 1975).

Shapiro filters operate as linear low-pass filters that target and eliminate 2Δ -scale waves—oscillations that span two grid intervals—while minimising the damping of long-wavelength components (Shapiro, 1975; Falissard, 2013). They are constructed with $2n$ -order accuracy and a $(2n + 1)$ -point stencil and they ensure strong attenuation of high-frequency noise while maintaining numerical fidelity (Falissard, 2013, 2015). Although widely adopted in interior domains, their application near boundaries remains more challenging, and a few adaptations have been proposed to address this limitation (Falissard, 2015).

This work uses the 2D second-order Shapiro filter. Let us consider a 2-dimensional field, $\mathbf{X} \in \mathbb{R}^{n \times p}$. The output $\mathbf{Y} \in \mathbb{R}^{n \times p}$, of applying the 2D second-order Shapiro filter is as follows (Shapiro, 1970):

$$\mathbf{Y}_{i,j} = \frac{1}{16}(4\mathbf{X}_{i,j} + 2(\mathbf{X}_{i-1,j} + \mathbf{X}_{i+1,j} + \mathbf{X}_{i,j-1} + \mathbf{X}_{i,j+1}) + \mathbf{X}_{i+1,j+1} + \mathbf{X}_{i+1,j-1} + \mathbf{X}_{i-1,j+1} + \mathbf{X}_{i-1,j-1}), \quad (2.11)$$

$\forall (i, j) \in [1, n] \times [1, p]$. The relevant boundary conditions must be applied to $\mathbf{X}_{i,j}$ for $i = 0, n + 1$ and $j = 0, p + 1$. Here, we formulate the Shapiro filter using finite differences on a rectilinear grid, for reasons that will be discussed later in this work.

The Shapiro filter is applied to address a key numerical artifact revealed in this work, which we discuss in chapters 4 and 5. In these chapters, we also describe how boundary conditions are enforced when applying the Shapiro filter.

2.4 Fast-Fourier transform

The Fourier transform is a mathematical tool we exploit during this work. It is a type of integral transform, in which any waveform can be re-written as the sum of sinusoidal

functions. The Fourier transform of an integrable function $f : \mathbb{R} \rightarrow \mathbb{C}$ is defined as

$$\mathcal{F}\{f(t)\} = \hat{f}(k) = \int_{-\infty}^{\infty} e^{-2\pi i k t} f(t) dt \quad (2.12)$$

and the inverse, if it exists, is defined as

$$\mathcal{F}^{-1}\{\hat{f}(k)\} = f(t) = \int_{-\infty}^{\infty} e^{2\pi i k t} \hat{f}(k) dk. \quad (2.13)$$

The discrete Fourier transform (DFT) of a function f_j , where $j = 0, \dots, N$, is given by

$$F_n = \sum_{j=0}^{N-1} f_j e^{-2\pi i n j / N} \quad (2.14)$$

where k is the wavenumber. The inverse discrete Fourier transform is given by

$$f_j = \sum_{n=0}^{N-1} F_n e^{2\pi i n j / N}. \quad (2.15)$$

Given a 2D array with N_1 rows and N_2 columns and entries $f_{j,k}$ we define the 2D DFT as

$$F_{m,n} = \sum_{k=0}^{N_2-1} e^{-2\pi i k n / N_2} \left(\sum_{j=0}^{N_1-1} e^{-2\pi i j m / N_1} f_{j,k} \right). \quad (2.16)$$

This is found by taking the DFT of each row and then column of the original matrix.

The Fast-Fourier transform (FFT) is a highly efficient algorithm for computing the DFT, reducing computational complexity from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$ by exploiting the iterative structure of DFT coefficients (Cooley and Tukey, 1965; Cochran et al., 1967; Cooley et al., 1969). In addition to significantly lowering computation time, the FFT also minimises round-off errors. A detailed historical account of the FFT is provided by Cooley et al. (1967). Later in this work, we use the FFT to analyse field signals suspected of containing numerical artifacts.

2.5 Laplacian with Neumann boundary conditions

We now introduce the Laplacian matrix with certain boundary conditions. We previously discussed the Poisson equations that can be derived from the Helmholtz decomposition in section 2.1. In this work we discuss solving these equations numerically for streamfunction and velocity potential. This involves solving Poisson equations with Neumann boundary conditions. Consider a 1D Laplacian operator,

$$\nabla^2 \psi = \frac{\partial^2 \psi}{\partial x^2}. \quad (2.17)$$

The discrete Laplacian matrix, with Neumann boundary conditions, is

$$\frac{1}{(\Delta x)^2} \begin{pmatrix} -1 & 1 & 0 & . & . & . & . & . & 0 & 0 \\ 1 & -2 & 1 & 0 & . & . & . & . & 0 & 0 \\ 0 & 1 & -2 & 1 & 0 & . & . & . & 0 & 0 \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . & . \\ 0 & 0 & . & . & . & 0 & 1 & -2 & 1 & 0 \\ 0 & 0 & . & . & . & 0 & 0 & 1 & -2 & 1 \\ 0 & 0 & . & . & . & . & . & 0 & 1 & -1 \end{pmatrix}$$

where Δx is the grid-spacing. Now, consider a 2D Laplacian operator,

$$\nabla^2 \psi = \frac{\partial^2 \psi}{\partial x^2} + \frac{\partial^2 \psi}{\partial y^2}. \quad (2.18)$$

The discrete Laplacian matrix, with Neumann boundary conditions, is

$$\begin{pmatrix} -((\Delta x)^{-2} + (\Delta y)^{-2}) & (\Delta x)^{-2} & 0 & . & 0 & (\Delta y)^{-2} & 0 & . & . & 0 \\ (\Delta x)^{-2} & -((\Delta x)^{-2} + 2(\Delta y)^{-2}) & (\Delta x)^{-2} & 0 & . & 0 & (\Delta y)^{-2} & 0 & . & 0 \\ 0 & . & 0 & (\Delta y)^{-2} & 0 & . & 0 & (\Delta x)^{-2} & -((\Delta x)^{-2} + 2(\Delta y)^{-2}) & (\Delta x)^{-2} \\ 0 & . & . & 0 & (\Delta y)^{-2} & 0 & . & 0 & (\Delta x)^{-2} & -((\Delta x)^{-2} + (\Delta y)^{-2}) \end{pmatrix}.$$

It is evident that both matrices are tridiagonal and symmetric. They are also singular, as they have non-empty null spaces, spanned by the vector $(1, \dots, 1)^T$. These matrices include both negative and zero value eigenvalues. To solve these Laplacian equations with Neumann boundary conditions, it is necessary to reformulate the problem into an equivalent form involving a modified matrix with positive eigenvalues, ensuring the existence of a unique and stable solution.

2.6 Error metrics

In the final chapters of this thesis we use various error metrics. In chapter 5 these are used to assess reconstruction errors and in chapter 6 these are used to assess the analyses produced by various assimilation experiments. Here we define these metrics.

Suppose we have an analysis generated by an experiment, x^a , and some form of truth, x^t . A method that can be used to assess the error is the root mean square error (RMSE).

The RMSE is given by,

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{k=1}^n (x_k^a - x_k^t)^2} \quad (2.20)$$

where k indicates the sampling over time or space. The RMSE is the square root of the mean squared error (MSE). The MSE can be decomposed into variance and bias components, so the RMSE reflects both the random (variance) and systematic (bias) aspects of the analysis error. However, unlike the MSE, the RMSE retains the units of the analysis variable being more easily interpretable as a typical error magnitude (Wilks, 2011). When the analysis is perfect the RMSE is zero. The larger the RMSE the worse the accuracy of the analysis.

An assessment of the mean can be demonstrated using the mean absolute error (MAE), defined as follows

$$\text{MAE} = \frac{1}{n} \sum_{k=1}^n |x_k^a - x_k^t|. \quad (2.21)$$

The MAE is the arithmetic average of the absolute values of the differences between the members of each pair of analysis and truth. The MAE is zero if the analysis is perfect and increases as the discrepancies between the pair become larger (Wilks, 2011).

RMSE is commonly used as the standard statistical tool when evaluating model performance in meteorological and climate science (Chai and Draxler, 2014). However MAE is also a useful tool and Williamson et al. (1992) discusses the advantages of using MAE over RMSE. In contrast, Chai and Draxler (2014) highlight the benefit of using RMSE in certain circumstances. Hodson (2022) notes that statistically neither metric is inherently better, this is dependent on the distribution of the errors.

The final metric is the standard deviation of the error, which reveals how the variability of the truth has been captured in the analysis. This is given by,

$$\sigma = \sqrt{\frac{\sum_{k=1}^n |\epsilon_k - \bar{\epsilon}|^2}{n}} \quad (2.22)$$

where $\epsilon_k = x_k^a - x_k^t$ and $\bar{\epsilon} = \frac{1}{n} \sum_{k=1}^n \epsilon_k$ is the mean error.

2.7 Summary

In this chapter, we have presented the fundamental mathematical tools that will be utilised in the subsequent chapters. In the next chapter, we introduce ocean data assimilation,

outlining the specific methodologies and providing context for the work presented in this thesis.

Chapter 3

Ocean data assimilation and the control variable transform

This chapter introduces ocean data assimilation (DA) and the variable transformations within DA. The goal is to provide the background and motivation necessary for the chapters and work that follow.

We first introduce fundamental ocean concepts before delving into data assimilation techniques. We present the concept of variational data assimilation, in particular, three-dimensional variational data assimilation (3D VAR). We describe the incremental formulation of 3D VAR, and introduce 3D VAR first guess at appropriate time (FGAT). These variational methods are commonly used in the ocean and will be the main focus of this work. We then establish the control variable transform (CVT), a variable transformation commonly used within the incremental formulation of variational data assimilation to deal with the complexity of the background error covariance matrix. We demonstrate the process of the CVT within incremental 3D VAR. We finally discuss the use of the CVT in both the atmosphere and ocean operationally, providing context for the interest in defining new velocity control variables in the ocean. We outline the control variables we propose to be used in the ocean and a numerical approach for implementing them.

3.1 Ocean concepts

The oceans, spanning approximately 71% of the Earth’s surface (Lalli and Parsons, 1997; Stewart, 2008), represent a vast and crucial component of our planet. The ocean plays a critical role in global systems. Given its vastness, it is essential that we accurately understand the processes occurring within it and make reliable forecasts about its future state.

Ocean data assimilation is key to enhancing the accuracy of ocean model forecasts and reanalysis, bringing them closer to real-world conditions. Ocean forecasts are important for a variety of applications, such as monitoring water quality (Novellino et al., 2024); coastal management (Dré villon et al., 2008); mitigating storm damage and flooding of coastal areas (Kourafalou et al., 2015); Navy operations (Burnett et al., 2014); understanding and prediction of marine animal species and ship navigation (Dickey, 2003). In particular, the forecasts of ocean currents are essential for marine safety and offshore operations (Waters et al., 2024b). By advancing the current ocean data assimilation systems, we have the capability to improve numerical weather predictions and future climate forecasts (Sluka et al., 2016; Zhang et al., 2020). The ocean and atmosphere drive each other (Stewart, 2008) and coupled ocean-atmosphere systems are now run operationally (Guiavarc’h et al., 2019).

Ocean data assimilation is reliant on observations of ocean variables. Commonly, observations of sea surface temperature (SST), sea surface height (SSH), temperature and salinity, are assimilated in operational systems (Lea et al., 2014). These observations consist of satellite SST; in situ SST data from drifting buoys, ships and moored buoys; temperature and salinity profile data from various sources such as Argo floats (Atkinson et al., 2014); and sea level anomaly (SLA) data from satellite altimetry.

However, there is limited availability of surface current velocity measurements globally. *In situ* measurements of ocean currents are available around some coastlines from high-frequency (HF) radars. These ocean currents derived from HF radar allow some coastal data assimilation systems to assimilate HF radar currents (Hoteit et al., 2009; Ngodock et al., 2015). Near-surface currents can also be derived from the changing positions of surface drifters. Helber et al. (2023) and Smith et al. (2023) describe their two-step method for assimilating ocean drifter measurements, but these measurements are

sparse compared to the scales of variability in ocean currents. Additionally, the open ocean ADCPs (Acoustic Doppler Current Profilers) provide some isolated measurements of ocean currents. Isern-Fontanet et al. (2017) provide an extensive review of ocean surface current measurements and what is assimilated. Globally, this limited availability of ocean current measurements means they are generally not assimilated operationally (Waters et al., 2024a). The corrections to the ocean currents are estimated through their balance relationships with other ocean variables which are observed. This lack of velocity observations leads to them being poorly constrained in ocean data assimilation.

This gap highlights the need for improvements in the treatment of surface currents within ocean data assimilation frameworks. Several recently proposed satellite missions aim to address this limitation by measuring ocean surface current velocities from space. These include Harmony, the next European Space Agency (ESA) Earth Explorer 10 mission (Stoffelen, 2022); SKIM, a previous ESA Earth explorer 9 candidate (Ardhuin et al., 2019); SEASTAR, a previous ESA Earth Explorer 11 candidate (Gommenginger et al., 2019) and the NASA Earth System Explorers Mission concept ‘Ocean Dynamics and Surface Exchange with the Atmosphere’ (ODYSEA, 2023). These satellites employ different measurement techniques: some infer ocean surface currents using Doppler-based methods, others use radar. The Harmony mission is set to launch in 2029 and uses synthetic aperture radar. To fully leverage these potential future measurements, it is imperative that we improve current ocean data assimilation systems to effectively integrate these observations. Idealised experiments have shown the potential impact of satellite surface current data with Waters et al. (2024b) having shown large improvements in the quality of ocean analyses and forecasts. The aim of this work is to further improve methods for assimilating such data. This thesis develops an assimilation approach, which is necessary for determining how to best extract information from these proposed satellites.

Efficiently utilising future ocean current measurements hinges on leveraging the balance relationships present in the oceans. Some fundamental relationships are described in detail by Weaver et al. (2005). Given our focus on surface currents, the geostrophic balance relationship is of primary interest.

3.1.1 Geostrophic balance

The horizontal momentum equations are derived from the continuity equations and Newton's 2nd law. They can be used to describe the motion of a 2D stratified, rotating, incompressible fluid with uniform density, as follows

$$\frac{D\mathbf{v}}{Dt} + f\mathbf{k} \times \mathbf{v} + \frac{1}{\rho_0}\nabla p = \mathbf{F} \quad (3.1)$$

where t is time, $\mathbf{x}=(x,y)$ is the 2D spatial coordinate, $\mathbf{v}(\mathbf{x},t) = (u,v)$ is the horizontal velocity vector, p is pressure and $\mathbf{k}$ is the unit vector in the vertical. Any additional forcing is represented by $\mathbf{F}$, ρ_0 is the uniform density and $f = 2\Omega \sin \varphi$ is the Coriolis parameter where Ω is the Earth's rotation rate and φ is latitude (Gill, 2016). The Earth's rotation produces the Coriolis force (Stewart, 2008). The smallness of the Rossby number is a measure of the significance of rotation for a particular phenomenon, and as such is a key measure of the importance of the Coriolis force (Friedlander and Serre, 2002). The Rossby number is given by

$$R = \frac{U}{fL} \quad (3.2)$$

where L is the characteristic length scale of the motion and U is the horizontal velocity scale characteristic of the motion. Large-scale flows are defined as those with sufficiently large L and as such a small Rossby number. These flows experience a strong Coriolis force (Pedlosky, 2013).

Geostrophic balance is the approximation of most large-scale flows in the ocean at mid-latitudes where the Coriolis force balances the horizontal pressure gradient. We can derive the geostrophic balance equations from the momentum equations (Stewart, 2008),

$$u_s = -\frac{1}{f\rho_0}\frac{\partial p}{\partial y} \quad \text{and} \quad v_s = \frac{1}{f\rho_0}\frac{\partial p}{\partial x}. \quad (3.3)$$

Near the surface of the ocean, the shallow pressure is related hydrostatically to the elevation of the sea surface, η , by $p = g\rho_0\eta$, where g is gravitational acceleration (Stewart, 2008). Therefore the near-surface geostrophic flow can be calculated from sea surface height as

$$u_s = -\frac{g}{f}\frac{\partial \eta}{\partial y} \quad \text{and} \quad v_s = \frac{g}{f}\frac{\partial \eta}{\partial x}. \quad (3.4)$$

Having presented some key concepts of oceanography, we now discuss data assimilation more generally.

3.2 Variational data assimilation

Data assimilation (DA) combines observations and a background estimate (from a numerical model), weighted by their uncertainties, to give optimal initial conditions for a numerical forecast (Kalnay, 2003). DA is used in various fields such as ocean modelling (Anderson et al., 1996) and numerical weather prediction (Lorenc, 1986). It is a particularly innovative and useful process as it takes into consideration our knowledge of the errors of the observed data and numerical model.

Variational data assimilation formulates the data assimilation problem as a least squares problem, which involves minimising a cost function. This cost function combines the observed and model (background) data, weighted by their uncertainties (Lorenc, 1986; Talagrand and Courtier, 1987). These uncertainties are represented as error covariance matrices. In variational DA we assume that the background and observation uncertainties can be described by Gaussian distributions. Gaussian distributions are given by their mean and covariance only. Therefore, we represent the errors in the background and observations through their covariances. A benefit of variational data assimilation is its adaptability to various observations. Observed data comes from many sources, satellites, for example, and these observations can be indirectly related to the quantities we are trying to forecast. Additionally, observations may not be at the same locations as our model grid points. Variational DA is equipped to adapt to these situations through the non-linear observation operator, which maps the model state to the observations to enable an appropriate comparison.

In variational data assimilation, numerical minimisation methods are used to iteratively minimise cost functions (Lorenc, 1986). The minimiser of the cost function is known as the analysis, $\mathbf{x}^a$. There are two main branches of variational data assimilation: three-dimensional variational assimilation (3D VAR) and four-dimensional variational assimilation (4D VAR). 3D VAR assimilates observations at a given time, as shown in figure 3.1a. Observations are only distributed in space. Whereas 4D VAR allows for observations to be distributed in time. As seen in figure 3.1b, observations are compared to the evolving model state, which is fully determined by the initial conditions. Thus, the method focuses on estimating the initial state that leads the model to best fit observations throughout the assimilation window.




Figure 3.1: Variational data assimilation schemes.

We focus our attention on 3D VAR. The cost function, $\mathcal{J}(\mathbf{x})$ (with $\mathbf{x} \in \mathbb{R}^n$), to be minimised in 3D VAR is formulated as a non-linear least squares problem, where information from the observations and background are combined. This gives,

$$\mathcal{J}(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}^b)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) + \frac{1}{2}(\mathbf{y} - h(\mathbf{x}))^T \mathbf{R}^{-1}(\mathbf{y} - h(\mathbf{x})) \quad (3.5)$$

where $\mathbf{x}^b \in \mathbb{R}^n$ is the background state, $\mathbf{B} \in \mathbb{R}^{n \times n}$ is the background error covariance matrix, $\mathbf{y} \in \mathbb{R}^p$ is a vector of observations, $h : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is the (potentially) non-linear observation operator, and $\mathbf{R} \in \mathbb{R}^{p \times p}$ is the observation error covariance matrix. We now explain each of these terms.

The background state, sometimes referred to as the *a priori* estimate, describes an estimate of our current state before information from the latest observations is accounted for. This is often generated from a previous run of an assimilation cycle. In variational data assimilation the background distribution is assumed to be a Gaussian distribution with mean $\mathbf{x}^b$ and error covariance $\mathbf{B}$. $\mathbf{B}$ specifies the covariance of the background forecast errors, and largely determines how information is spread from observed variable locations to unobserved points in space and also other unobserved variables (Bannister et al., 2008). The observation operator maps the state vector to observation space. This enables a comparison between the model and observations. Finally, the likelihood function is assumed to be a Gaussian with mean $\mathbf{y}$ and error covariance $\mathbf{R}$.

We need to minimise (3.5) in order to find the analysis, $\mathbf{x}^a$. This usually involves using a gradient descent algorithm such as the conjugate gradient method (see section 2.2.2). These algorithms require the gradient of the cost function at each iteration of the minimisation. The gradient is given by

$$\nabla \mathcal{J}(\mathbf{x}) = \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^b) - \mathbf{H}^T \mathbf{R}^{-1}(\mathbf{y} - h(\mathbf{x})), \quad (3.6)$$

which requires the adjoint operator $\mathbf{H}^T \in \mathbb{R}^{p \times n}$. The tangent linear operator, $\mathbf{H}$, is the observation operator linearised about the current state (Waller et al., 2017). We now introduce the incremental formulation of variational data assimilation.

3.2.1 Incremental 3D VAR

In practice, the minimisation of non-linear functions in variational data assimilation can be computationally expensive, especially in 4D VAR due to the need to integrate a non-linear model and an adjoint repeatedly (Lawless et al., 2005). To reduce computational cost, an incremental approach was proposed (Courtier et al., 1994) whereby we solve the DA problem in terms of increments to the background. Here, the cost function with non-linear constraints (3.5) is replaced by a series of linear least-squares cost functions with linear constraints (Courtier et al., 1994; Lawless et al., 2005). These cost functions are quadratic and can be minimised iteratively.

We define $\mathbf{x}^{(k)}$ to be the k^{th} estimate of the solution of (3.5). The next estimate will be

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \delta \mathbf{x} \quad (3.7)$$

where $\delta \mathbf{x} \in \mathbb{R}^n$ is the increment and a solution of the following linearised cost function (Lawless et al., 2005),

$$\mathcal{J}(\delta \mathbf{x}) = \frac{1}{2} \left(\delta \mathbf{x} - (\mathbf{x}^b - \mathbf{x}^{(k)}) \right)^T \mathbf{B}^{-1} \left(\delta \mathbf{x} - (\mathbf{x}^b - \mathbf{x}^{(k)}) \right) + \frac{1}{2} \left(\mathbf{H} \delta \mathbf{x} - \mathbf{d}^{(k)} \right)^T \mathbf{R}^{-1} \left(\mathbf{H} \delta \mathbf{x} - \mathbf{d}^{(k)} \right) \quad (3.8)$$

where $\mathbf{d}^{(k)} = \mathbf{y} - h(\mathbf{x}^{(k)})$ is the residual vector, called the innovation vector when $k = 0$ (Kang et al., 2014). The gradient is given by

$$\nabla \mathcal{J}(\delta \mathbf{x}) = \mathbf{B}^{-1} \left(\delta \mathbf{x} - (\mathbf{x}^b - \mathbf{x}^{(k)}) \right) + \mathbf{H}^T \mathbf{R}^{-1} \left(\mathbf{H} \delta \mathbf{x} - \mathbf{d}^{(k)} \right). \quad (3.9)$$

The minimisation of (3.8) is known as the *inner loop*. This step is located within the *outer loop*, where we update the previous estimate, using (3.7) (Lawless, 2013). The process of implementing incremental 3D VAR, as described in Lawless et al. (2005), is as follows,

1. For the first iteration, $k = 0$, we choose the initial guess to be the background state, $\mathbf{x}^{(0)} = \mathbf{x}^b$.
2. For $k = 0, 1, \dots$ we calculate the residual vector, $\mathbf{d}^{(k)} = \mathbf{y} - h(\mathbf{x}^{(k)})$.
3. Minimise the quadratic cost function (3.8) for perturbation $\delta\mathbf{x}$. The initial guess will be $\delta\mathbf{x} = 0$.
4. Update the estimate, $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \delta\mathbf{x}$.
5. Repeat steps 2 - 4 until desired convergence criterion is satisfied.

Here, step 3 is considered the *inner loop* and steps 2,4 and 5 the *outer loop*. Incremental variational data assimilation significantly reduces the computational cost of the minimisation in variational data assimilation as the inner loop can be performed at a lower resolution to the outer loop and a simpler observation operator is used (Courtier et al., 1994).

3.2.2 Incremental 3D FGAT

3D VAR is a suitable data assimilation method for dynamical systems where the change is slow over the assimilation window. It is common for operational centres to compromise between 4D VAR and 3D VAR by using 3D FGAT (first guess at appropriate time). 3D FGAT differs from the 3D VAR we introduced in section 3.2.1 as it is a simplification of 4D VAR; the time dependency is retained in the residual vector but the increments are assumed constant in time (Lee et al., 2004; Weaver et al., 2003; Lorenc and Rawlins, 2005). This reduces the computational cost associated with 4D VAR. 3D FGAT is used in some operational ocean DA systems which we discuss in section 3.4.2. Using incremental 3D FGAT, we now consider observations over a time window $t = [0, J]$. The incremental 3D FGAT cost function is given by,

$$\mathcal{J}(\delta\mathbf{x}_0) = \frac{1}{2} \left(\delta\mathbf{x}_0 - (\mathbf{x}_0^b - \mathbf{x}_0^{(k)}) \right)^T \mathbf{B}^{-1} \left(\delta\mathbf{x}_0 - (\mathbf{x}_0^b - \mathbf{x}_0^{(k)}) \right) + \frac{1}{2} \sum_{j=0}^J \left(\mathbf{H}_j \delta\mathbf{x}_0 - \mathbf{d}_j^{(k)} \right)^T \mathbf{R}_j^{-1} \left(\mathbf{H}_j \delta\mathbf{x}_0 - \mathbf{d}_j^{(k)} \right) \quad (3.10)$$

where $\mathbf{d}_j^{(k)} = \mathbf{y}_j - h_j(\mathbf{x}_j)$ and $\mathbf{x}_j$ is calculated from $\mathbf{x}_0$ by evolving the non-linear model. The gradient is given by

$$\nabla \mathcal{J}(\delta(\mathbf{x}_0)) = \mathbf{B}^{-1} \left(\delta \mathbf{x}_0 - (\mathbf{x}_0^b - \mathbf{x}_0^{(k)}) \right) + \sum_{j=0}^J \mathbf{H}_j^T \mathbf{R}_j^{-1} \left(\mathbf{H}_j \delta \mathbf{x}_0 - \mathbf{d}_j^{(k)} \right). \quad (3.11)$$

3.3 Control variable transform

The background error covariance matrix, $\mathbf{B}$, is a particularly important component of variational data assimilation, but can be difficult to represent. In practice, we are unable to explicitly formulate $\mathbf{B}$ as a matrix due to its size. For example, a state vector can be of length $n = 10^9$ as in the 1/12 degree global ocean NEMO model used at the Met Office. As $\mathbf{B}$ is of size $n \times n$, it is extremely large. To fully represent and manipulate a matrix of this size is not feasible. Furthermore, finding the inverse of a matrix of this size is impossible. A method used to avoid the need to calculate the inverse of $\mathbf{B}$ is the control variable transform.

The control variable transform (CVT) is a well known concept in data assimilation and was first introduced by Parrish and Derber (1992). It is commonly used in variational data assimilation systems across the world. In the CVT, the assimilation is performed with a set of variables, called the control variables, that differ from the original model variables (Lorenc et al., 2000; Parrish and Derber, 1992). After the analysis is found in this new control space, it is transformed back to the original model space. The control variables are chosen such that their errors can be considered to be approximately uncorrelated, thereby simplifying the structure of $\mathbf{B}$ in the control space.

The control variable transform is frequently used within the incremental formulation of variational data assimilation (Lorenc et al., 2000), where we solve the data assimilation problem in terms of increments to the background, as described in section 3.2.1. We define a control variable $\delta \mathbf{z}$ such that there is a transformation $\mathcal{U}$ to the model variables. This is known as the U-transform:

$$\delta \mathbf{x} = \mathcal{U} \delta \mathbf{z}. \quad (3.12)$$

The T-transform is the inverse, and allows a transformation from model to control space:

$$\delta \mathbf{z} = \mathcal{T} \delta \mathbf{x}. \quad (3.13)$$

Now the inner loop minimisation of the incremental 3D VAR cost function (3.8) is re-written in terms of the control variables, $\delta\check{\mathbf{z}}$,

$$\mathcal{J}(\delta\check{\mathbf{z}}) = \frac{1}{2} \left(\delta\check{\mathbf{z}} - \delta\check{\mathbf{z}}_b^{(k)} \right)^T \mathcal{U}^T \mathbf{B}^{-1} \mathcal{U} \left(\delta\check{\mathbf{z}} - \delta\check{\mathbf{z}}_b^{(k)} \right) + \frac{1}{2} \left(\mathbf{H} \mathcal{U} \delta\check{\mathbf{z}} - \mathbf{d}^{(k)} \right)^T \mathbf{R}^{-1} \left(\mathbf{H} \mathcal{U} \delta\check{\mathbf{z}} - \mathbf{d}^{(k)} \right) \quad (3.14)$$

where $\delta\check{\mathbf{z}}_b^{(k)} = \mathcal{T}(\mathbf{x}^b - \mathbf{x}^{(k)})$. As stated previously, the control variables, $\delta\check{\mathbf{z}}$, are chosen such that their errors can be assumed to be approximately uncorrelated. As such, in the case where $\mathcal{U}$ is a square matrix, the background error covariance matrix of the control variables is the identity matrix by definition. The U-transform replaces the calculations involving the background error covariance matrix. As a result of this transform, we implicitly have $\mathbf{B} = \mathcal{U} \mathcal{U}^T$. The cost function in control space no longer contains the original background error covariance matrix and the calculation of $\mathbf{B}^{-1}$ is no longer necessary. This statement is also valid when $\mathcal{U}$ is a rectangular matrix due to the minimisation procedure which constrains $\delta\check{\mathbf{z}} - \delta\check{\mathbf{z}}_b^{(k)}$ to the subspace spanned by $\mathcal{U}^T$ (Ménétrier and Auligné, 2015). Therefore, we can effectively re-write (3.14) as,

$$\mathcal{J}(\delta\check{\mathbf{z}}) = \frac{1}{2} \left(\delta\check{\mathbf{z}} - \delta\check{\mathbf{z}}_b^{(k)} \right)^T \left(\delta\check{\mathbf{z}} - \delta\check{\mathbf{z}}_b^{(k)} \right) + \frac{1}{2} \left(\mathbf{H} \mathcal{U} \delta\check{\mathbf{z}} - \mathbf{d}^{(k)} \right)^T \mathbf{R}^{-1} \left(\mathbf{H} \mathcal{U} \delta\check{\mathbf{z}} - \mathbf{d}^{(k)} \right). \quad (3.15)$$

In fact, both the U- and T-transforms contain two transformations within them. The first transforms between physical variables (parameter transform) and the second transforms spatially (spatial transform),

$$\mathcal{U} = \mathcal{U}_p \mathcal{U}_s \quad \text{and} \quad \mathcal{T} = \mathcal{T}_s \mathcal{T}_p.$$

This gives $\mathbf{B} = \mathcal{U}_p \mathcal{U}_s \mathcal{U}_s^T \mathcal{U}_p^T$. In the T-transform, the parameter transform removes the correlations between variables and the spatial transform removes the spatial correlations. The work in this thesis will be solely concerned with the parameter transform and we do not go into detail regarding the spatial transform. Hereafter we use $\mathcal{U}$ to denote the parameter transform.

3.3.1 Implementing the CVT within incremental 3D VAR

The following steps are performed to implement the CVT within incremental 3D VAR (detailed in section 3.2.1),

1. For the first iteration of the outer loop, $k = 0$, define the initial guess to be the background state,

$$\mathbf{x}^{(0)} = \mathbf{x}^b.$$

2. Calculate the residual vector,

$$\mathbf{d}^{(k)} = \mathbf{y} - h(\mathbf{x}^{(k)}).$$

3. We transform the background increment using the T-transform for iteration k ,

$$\delta\check{\mathbf{z}}_b^{(k)} = \mathcal{T}^{(k)}(\delta\mathbf{x}_b^{(k)}) = \mathcal{T}^{(k)}(\mathbf{x}^{(k)} - \mathbf{x}^b),$$

which is not necessary for the first iteration since $\delta\check{\mathbf{z}}_b^{(k)} = 0$.

4. Find the value of $\delta\check{\mathbf{z}}$ that minimises

$$\mathcal{J}(\delta\check{\mathbf{z}}) = \frac{1}{2} \left(\delta\check{\mathbf{z}} - \delta\check{\mathbf{z}}_b^{(k)} \right)^T \left(\delta\check{\mathbf{z}} - \delta\check{\mathbf{z}}_b^{(k)} \right) + \frac{1}{2} \left(\mathbf{H}\mathcal{U}\delta\check{\mathbf{z}} - \mathbf{d}^{(k)} \right)^T \mathbf{R}^{-1} \left(\mathbf{H}\mathcal{U}\delta\check{\mathbf{z}} - \mathbf{d}^{(k)} \right).$$

5. Transform the solution using the U-transform for iteration k :

$$\delta\mathbf{x} = \mathcal{U}^{(k)}\delta\check{\mathbf{z}}.$$

6. Update the current guess:

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \delta\mathbf{x}.$$

7. Repeat the outer loop (steps 2 - 6) until desired convergence is reached.

The T-transform, used in step 3, can be avoided altogether if both the U- and T-transforms remain static (Katz, 2007). In this case the background increment calculated in step 3 becomes

$$\delta\check{\mathbf{z}}_b^{(k+1)} = \delta\check{\mathbf{z}}_b^{(k)} - \delta\check{\mathbf{z}}^{(k)}. \quad (3.16)$$

Although we have shown that the T-transform is not directly used to calculate the analysis, the transformation is necessary for computing covariances of the control variables, which are used as the background error statistics for the assimilation (Katz et al., 2011).

3.4 Operational data assimilation systems

In this section we discuss some of the data assimilation systems used operationally, both in the atmosphere and ocean. We emphasise the use of the control variable transform in these systems. This will lead to the discussion of the alternative velocity control variables that we propose for the ocean.

Commonly, the CVT in the ocean is based on physically-based balance relationships derived from the equations governing the ocean circulation (Derber and Bouttier, 1999; Weaver et al., 2005). Some state variables are split into their balanced and unbalanced components, and the control variables are chosen from these. In particular, geostrophic balance is used to find the balanced components of the velocities, as described in section 3.1.1. However, we firstly discuss atmospheric data assimilation.

3.4.1 Atmospheric data assimilation

Before exploring ocean data assimilation, we first briefly discuss atmospheric data assimilation. A variety of data assimilation methods are used in atmospheric data assimilation. An in-depth review of these different methods is given by Bannister (2017). Variational data assimilation is commonly used and, in particular, the velocity control variables used in the atmosphere are commonly based on Helmholtz Theorem. Recall from section 2.1 that Helmholtz Theorem decomposes the horizontal velocities into their irrotational (divergence) and non-divergent (vorticity) parts. It is often assumed in the atmosphere that the rotational component of the flow is fully balanced, and that this balance is entirely represented by the vorticity (Parrish and Derber, 1992; Katz et al., 2011). In the Met Office atmospheric DA system it is assumed that streamfunction is totally balanced, as streamfunction is the rotational component of the velocity, given by the Poisson equation for vorticity (2.7). The Met Office control variables are streamfunction, velocity potential and unbalanced pressure (Bannister, 2008; Lorenc et al., 2000; Lorenc, 2003) where the balanced component of pressure is related to streamfunction (vorticity) through the linear balance equation. Similarly, ECMWF assumes vorticity to be totally balanced and uses vorticity, unbalanced divergence, unbalanced temperature and surface pressure as their control variables (Derber and Bouttier, 1999). This has recently been updated to also include humidity (Semane and Bonavita, 2025).

Cullen (2003) and Katz et al. (2011) describe the benefit of using a CVT based on potential vorticity rather than vorticity in the atmosphere. This allows for vorticity to have both balanced and unbalanced components. However, this requires the solution of a complex 3D elliptic problem, and as such vorticity is more commonly used.

We propose using a CVT based on Helmholtz Theorem in the ocean, but with differing assumptions for the totally balanced variable, which we describe in section 3.5. Oceanic data assimilation is quite different from atmospheric data assimilation for a number of reasons:

- The presence of coastal boundaries.
- The temporal and spatial scales differ in the ocean.
- The observations in the ocean are more sparse than in the atmosphere, particularly in the subsurface ocean.

For these reasons, the way data assimilation schemes are set up and applied in the ocean can differ from their use in the atmosphere, as we will now discuss.

3.4.2 Oceanic data assimilation

There are various data assimilation schemes used in ocean data assimilation. In this section, we describe some of the notable schemes used for both global ocean and regional ocean applications, with a summary provided in table 3.1. As this table does not include all ocean DA systems, a more thorough overview of operational ocean data assimilation schemes is provided in Martin et al. (2024). This section then focuses on how the current variational schemes in particular can be improved for assimilating future ocean current measurements.

Global ocean data assimilation systems use a variety of sequential and variational algorithms. Regional ocean models face different difficulties to the global ocean. The first and most distinct is due to the presence of open boundaries that define a regional ocean. Regional oceans are driven by more local features and are usually run at a much higher resolution than global models (Edwards et al., 2015). For example, the Met Office global ocean model runs at both $1/4^\circ$ and $1/12^\circ$ resolution (Mignac et al., 2025) whereas

DA System	Developing/operating institutions	References	DA Method	Domain
NEMOVAR (Nucleus for European Modelling of the Ocean Variational DA system)	Met Office, ECMWF (European Centre for Medium-range Weather Forecasts), CERFACS, INRIA	Weaver et al. (2005), Waters et al. (2015), Zuo et al. (2019)	4D VAR, hybrid schemes. Operationally uses 3D FGAT with control variables temperature, unbalanced salinity, unbalanced SSH and unbalanced velocities.	Global and regional
SAM2 (Système d'Assimilation Mercator)	Mercator, ECCC (Environment and Climate Change Canada)	Lellouche et al. (2013), Pham et al. (1998), Lellouche et al. (2018), Smith et al. (2016)	SEEK-FGAT (singular evolutive extended Kalman filter). Background error covariances modelled by an ensemble of multivariate 3D anomalies from a multi-year hindcast simulation. 3D VAR scheme provides correction for slowly evolving large-scale biases in temperature and salinity.	Global and regional
EnKF-C	The Bluelink Partnership	Brassington et al. (2023)	Deterministic ensemble Kalman Filter.	Global
NCODA 3D VAR (Navy Coupled Ocean Data Assimilation)	FNMO (Fleet Numerical Meteorology and Oceanography Center), NAVOCEANO (Naval Oceanographic Office)	Smith et al. (2017), Cummings et al. (2010), Cummings and Smedstad (2013)	3D VAR	Global
MOVE (Multivariate Ocean Variational Estimation)	JMA (Japan Meteorological Agency)	Usui et al. (2015), Hirose et al. (2019), Fujii et al. (2023)	MOVE/MRI.COM-G3 is a global scheme is composed of two schemes of differing resolution. MOVE-4DVAR is regional and uses 4D VAR with control variables composed of amplitudes for the vertical coupled temperature-salinity EOF modes	Global and western north pacific region
ROMS3DVAR (Regional Ocean Modeling System 3D VAR)	JPL (Jet Propulsion Laboratory)	Li et al. (2015)	Multiscale 3D VAR with control variables temperature, unbalanced salinity, non-steric height and unbalanced velocities streamfunction and velocity potential.	California coastal region
ROMS 4D-Var	University of California, Santa Cruz	Moore et al. (2011), Moore et al. (2023)	4D VAR with control variables temperature, unbalanced salinity, unbalanced SSH and unbalanced velocities.	Regional
OceanVar2	CMCC	Dobricic and Pinardi (2008), Storto et al. (2016), Coppini et al. (2023)	3D VAR with background errors split into horizontal and vertical mode. Horizontal covariances are taken as Gaussian functions, vertical covariances are represented by EOFs (Empirical Orthogonal Functions).	Global and Mediter-ranean sea

Table 3.1: Ocean data assimilation systems from around the globe. Details of which institutions develop/run these, which DA scheme is used and the domain.

their regional shelf seas model runs at $7km$ and $1.5km$ resolution (Tonani et al., 2019). Regional data assimilation schemes are often more sophisticated, partly because the smaller domain sizes make the use of advanced methods more computationally feasible. Additionally, regional models must represent more complex processes, such as tides, which play a significantly greater role in shallow coastal waters. Here we discuss the schemes used in both global and regional domains.

Although variational methods are common in the ocean, Kalman filters are used by some. The Système d’Assimilation (SAM2) of Mercator and Environment and Climate Change Canada (ECCC), use SEEK-FGAT. This is a reduced-order Kalman filter based on the singular evolutive extended Kalman filter (SEEK) formulation with a fixed basis. The background error covariances are modelled by an ensemble of multivariate three-dimensional anomalies derived from a multi-year hindcast simulation and a 3D VAR scheme provides a correction for the slowly evolving large-scale biases in temperature and salinity (Lellouche et al., 2013; Pham et al., 1998; Lellouche et al., 2018; Smith et al., 2016) both globally and regionally. The Bluelink Partnership, part of the Australian government, uses a deterministic ensemble Kalman filter in their global DA system (Brassington et al., 2023).

Both 3D and 4D VAR are regularly used in regional ocean data assimilation systems. For the Mediterranean sea, Dobricic and Pinardi (2008) proposed a 3D VAR scheme that aims to reduce the complexity of the background error covariance matrix and allow for regional variability and complex coastlines. Unlike the conventional transformations to control space in meteorological applications (as described in section 3.3), in Dobricic and Pinardi (2008) the covariances are split into their horizontal and vertical modes where the horizontal covariances are assumed Gaussian and vertical covariances are represented by multivariate Empirical orthogonal functions (EOFs). This was extended for use in the global domain by Storto et al. (2016).

Globally, Fleet Numerical Meteorology and Oceanography Centre (FNMOC) and the Naval Oceanographic Office (NAVOCEANO) have a DA system called NCODA 3D VAR, that uses a 3D VAR scheme (Cummings et al., 2010; Cummings and Smedstad, 2013). A 4D VAR scheme, MOVE-4DVAR, is used for the Western North Pacific region at Japan Meteorological Agency (JMA). The scheme estimates temperature and salinity fields above 1500m using a multivariate variational method with vertically coupled temperature

salinity (T-S). EOFs model the decomposition of the background error covariance matrix and the control variables are composed of amplitudes for the vertical coupled T-S EOF modes (Usui et al., 2015; Hirose et al., 2019). JMA also run a global DA scheme called MOVE/MRI.COM-G3. Observations are directly assimilated into MOVE-G3A through a 4DVAR method, and the temperature and salinity analysis fields are downscaled into MOVE-G3F through incremental analysis updates (Fujii et al., 2023).

Moreover, an operational variational regional ocean data assimilation scheme, used in various regions such as the California coast, based on 4D VAR has been developed by Moore et al. (2023). The control variables are taken as temperature and the unbalanced components of salinity, SSH and the horizontal velocities. This is largely based on the balance operator used NEMOVAR (Weaver et al., 2005), which we describe in section 3.4.3. The spatial correlations are modelled using an explicit diffusion operator as detailed in Weaver and Courtier (2001).

Li et al. (2008a) and Li et al. (2008b) describe the implementation of ROMS3DVAR, another variational regional ocean data assimilation scheme used in the California coastal region by the Jet Propulsion Laboratory (JPL) at the California Institute of Technology. The control variables used in ROMS3DVAR are temperature (T), salinity (S), non-steric height (η'), and unbalanced streamfunction (ψ_U) and velocity potential χ_U . The balance relationships used to find these control variables are based on two dynamical constraints; hydrostatic and geostrophic balance. Using hydrostatic balance, η can be split into steric and nonsteric components. Steric height is the component of height which occurs when fluctuations in density do not change the total mass of the water column (Li et al., 2008a). Unlike larger scale models, the nonsteric part is considered non-negligible for coastal data assimilation. Geostrophic balance is applied to both steric and nonsteric height to find the geostrophic velocity. Li et al. (2008a) note that streamfunction and velocity potential can be considered more suitable control variables than the velocity components in data assimilation. As such the ageostrophic (unbalanced) velocities are transformed to the ageostrophic ψ and χ . Li et al. (2008a) organise the control variables into three groups η' ; ψ_U and χ_U ; T and S . The errors within these groups are assumed uncorrelated. As such, their incremental 3D VAR cost function is broken into three parts, with three separate background error covariance matrices and observation error covariance matrices. The main focus of their work is treating the spatial correlations. Li et al. (2008a) propose

using a Kronecker product to define the correlation matrices. This has the capability of expressing the anisotropic and inhomogeneous features of coastal oceans and guarantees positive definiteness. A multi-scale data assimilation framework was then introduced to ROMS3DVAR by Li et al. (2015).

In this thesis we are particularly interested in NEMOVAR, a variational data assimilation software for the Nucleus for European Modelling of the Ocean (NEMO) model, used both in global and regional domains. NEMOVAR is used operationally at the UK Met Office (Waters et al., 2015) and European Centre for Medium-Range Weather Forecasts (ECMWF) (Mogensen and Balmaseda, 2012; Balmaseda et al., 2013; Zuo et al., 2019). NEMOVAR supports multiple DA schemes, including both variational and hybrid approaches, but uses 3D FGAT operationally. In the next section, we provide further details on the configuration of the NEMOVAR system used in this work. Other variables, such as sea-ice concentration, could also be included, but these are not considered here.

3.4.3 NEMOVAR

In NEMOVAR the background error covariance matrix, $\mathbf{B}$, can be written as

$$\mathbf{B} = \mathbf{K}\mathbf{\Sigma}\mathbf{C}\mathbf{\Sigma}\mathbf{K}^T \quad (3.17)$$

where $\mathbf{C}$ is the univariate correlation matrix, $\mathbf{C} = \mathbf{C}^{1/2}\mathbf{C}^{1/2^T}$, $\mathbf{\Sigma}$ is the diagonal matrix containing the standard deviations and $\mathbf{K}$ is the balance operator. We can compare 3.17 to $\mathbf{B} = \mathcal{U}_p\mathcal{U}_s\mathcal{U}_s^T\mathcal{U}_p^T$ in section 3.3, where $\mathcal{U} = \mathcal{U}_p\mathcal{U}_s$. Here, $\mathbf{K}$ takes the role of $\mathcal{U}_p$ and $\mathcal{U}_s = \mathbf{\Sigma}\mathbf{C}^{1/2}$. As discussed previously, for the rest of this work, when we refer to the U-transform we will solely be concerned with the parameter transform, which is dictated by the balance operator, $\mathbf{K}$.

In section 3.3 the $\mathcal{U}$ operator acts on $\delta\mathbf{z}$, i.e. $\delta\mathbf{x} = \mathcal{U}\delta\mathbf{z} = \mathbf{K}\mathbf{\Sigma}\mathbf{C}^{1/2}\delta\mathbf{z}$. Since we are concerned only with $\mathbf{K}$, it will be convenient to redefine what we mean by a control variable from now on, and so we define $\delta\mathbf{z} = \mathbf{\Sigma}\mathbf{C}^{1/2}\delta\mathbf{z}$. Hence, the relationship between model variables, $\delta\mathbf{x}$ and $\delta\mathbf{z}$ is $\delta\mathbf{x} = \mathbf{K}\delta\mathbf{z}$. The univariate correlations are defined using a diffusion operator (Weaver et al., 2016), which can easily account for coastlines, keeping increments confined to specific basins. As discussed previously, for the rest of this work, when we refer to the U-transform we will solely be concerned with the parameter transform, which is dictated by the balance operator, $\mathbf{K}$.

In NEMO, the state vector increment contains the following physical variables: potential temperature δT , salinity δS , sea surface height (SSH) $\delta \eta$ and the components of the horizontal velocity vector $\delta \mathbf{v} = (\delta u, \delta v)^T$ to give $\mathbf{x} = (\delta T, \delta S, \delta \eta, \delta \mathbf{v})^T$. The control vector contains the unbalanced components of the state variables, denoted $\cdot_U$, apart from potential temperature which is taken in its totality. The control space in NEMOVAR is $\delta \mathbf{z} = (\delta T, \delta S_U, \delta \eta_U, \delta \mathbf{v}_U)^T$. We denote the balanced components as $\cdot_B$.

As mentioned above, in NEMOVAR, a balance operator, K , is used to describe the balance relationship between variables (Weaver et al., 2005) and transform from control to state space. The inverse of the balance operator is designed to decorrelate the model variables. A linearised version of the balance operator, $\mathbf{K}$, is used in the incremental formulation. This balance operator can be summarised by the following sequence of equations,

$$\begin{aligned}\delta T &= \delta T, \\ \delta S &= \mathbf{K}_{ST}\delta T + \delta S_U := \delta S_B + \delta S_U, \\ \delta \eta &= \mathbf{K}_{\eta\rho}\delta\rho + \delta\eta_U := \delta\eta_B + \delta\eta_U, \\ \delta \mathbf{v} &= \mathbf{K}_{vp}\delta p + \delta \mathbf{v}_U := \delta \mathbf{v}_B + \delta \mathbf{v}_U,\end{aligned}\tag{3.18}$$

where

$$\begin{aligned}\delta\rho &= \mathbf{K}_{\rho T}\delta T + \mathbf{K}_{\rho S}\delta S, \\ \delta p &= \mathbf{K}_{p\rho}\delta\rho + \mathbf{K}_{p\eta}\delta\eta,\end{aligned}$$

are diagnostic quantities corresponding to density, ρ , and pressure, p . The transformation from the variable(s) n to m are represented by $\mathbf{K}_{mn}$. These are the balance relationships (Weaver et al., 2005) and they are contained within the balance operator $\mathbf{K}$.

The balance operator determines the cross-variable covariances which make up the multivariate component of the background error covariance matrix. To find the control variables we remove the balanced component using the inverse of the balance operator,

$\mathbf{K}^{-1}$. This is simply obtained by performing the inverse of (3.18):

$$\delta T = \delta T,$$

$$\delta S_U = \delta S - \delta S_B,$$

$$\delta \eta_U = \delta \eta - \delta \eta_B,$$

$$\delta \mathbf{v}_U = \delta \mathbf{v} - \delta \mathbf{v}_B.$$

This inverse sequence of balance relations leaves temperature unaltered but removes the parts of salinity that could be related to temperature, the parts of SSH that could be related to temperature and salinity, and the parts of velocity that could be related to temperature, salinity and SSH (Weaver et al., 2005).

The balance relationships are discussed in detail in Weaver et al. (2005). Temperature is taken in totality and, as described, all the balance relations stem from there. Balanced salinity increments are determined from the vertical derivatives of the temperature and salinity background state. Balanced density is found by linearising the equation of state around a reference state. SSH is computed diagnostically as a function of temperature and salinity using the dynamic height relation (Weaver et al., 2005). In this work, we are mainly concerned with the treatment of the velocity variables in both global and regional domains, but specifically in NEMOVAR. The balance relation, $\mathbf{K}_{v,p}(\delta \mathbf{p})$, is geostrophic balance (3.3) and so the balanced velocities are given by their geostrophic components. Geostrophic balance breaks down near the equator where the Coriolis force approaches zero. A separate method described in Weaver et al. (2005) is used to deal with this area, whereby the zonal velocity is taken to be in second order geostrophic balance and the meridional velocity is reduced to zero at the equator.

The ageostrophic velocities are used as control variables in NEMOVAR. Whilst these are currently suitable velocity control variables, the velocity components are highly correlated and the assumption that their error covariances are uncorrelated is incorrect. As such, alternative control variables are highly desirable in preparation for future surface current measurements to be assimilated. Addressing this need is a central focus of this thesis. In this work, we propose new velocity control variables for NEMOVAR and discuss their implementation and potential implications.

Building on the previous discussion of data assimilation schemes and control variables used in both global and regional ocean applications, we now introduce the proposed

velocity control variables for use in NEMOVAR. These are inspired by approaches commonly used in atmospheric data assimilation and bear some similarities to those in Li et al. (2008a), but differ overall due to the distinct framework and methodology employed in NEMOVAR.

3.5 Proposed control variables

In the NEMOVAR ocean DA system, the ageostrophic velocities are used as the velocity control variables. This assumes errors in the ageostrophic velocities are uncorrelated, so that an ocean-current observation updates each velocity component independently. However, the velocity components are intricately dependent on each other, as they are coupled through the momentum equations, the incompressibility constraint, and advective transport. Independent updates therefore fail to capture this mutual influence. To improve the assimilation of ocean-current measurements, we seek alternative velocity control variables that allow the components to update each other. We propose the use of new control variables in NEMOVAR, namely a CVT that is based on Helmholtz Theorem, which will allow for the ageostrophic velocity components to be correlated. The use of streamfunction and velocity potential as control variables has been used in other variational data assimilation schemes both in the atmosphere and ocean, as described in sections 3.4.1 and 3.4.2. As we are proposing an alternative approach to the velocity variables in NEMOVAR, the totally balanced variable is taken as temperature, and not vorticity (or potential vorticity) as done in the atmosphere (Katz et al., 2011).

It is often stated that control variables based on streamfunction and velocity potential are more appropriate than the velocity vector components since the assumption that they are uncorrelated is more suitable (Daley, 1993, Section 5.2). Daley (1993) assumes there is no correlation between the velocity potential and streamfunction. The justification for this comes from the observational work of Hollingsworth and Lönnberg (1986), who conducted experiments to compute the cross-correlations between streamfunction and velocity potential errors. Their results indicate that these cross-correlation terms are small in the background error forecast, suggesting a weak correlation. Panchev (1971) discusses this from a more theoretical perspective. This assumption is discussed briefly in appendix A. However, this work does not focus on testing the uncorrelated assumption.

An additional benefit of using velocity control variables derived from the Helmholtz theorem is that they give better control over horizontal divergence, and as such the vertical motions. The coupling of physical and ocean biogeochemical data assimilation systems is hindered due to the extremely high sensitivity of ocean biogeochemistry to spurious vertical velocities (Park et al., 2018). These arise due to imbalances induced during the physical ocean data assimilation where the analysis increments produced are not exact solutions to the equations of the model (Pilo et al., 2018). Physical ocean data assimilation has been widely found to spuriously increase the vertical mixing and vertical velocities, degrading the biogeochemical model simulations (Ford and Barciela, 2017). Park et al. (2018) state that even small spurious vertical transports across the ocean can have disastrous impacts on the biogeochemical states. By utilising velocity control variables based on their non-divergent and irrotational components, we gain direct control over the divergence of the velocity analysis increment. Consequently, through the continuity equation, this allows us to effectively constrain vertical motions.


Furthermore, according to Li et al. (2008*b*) the auto-error covariances of velocity potential and streamfunction are expected to be more isotropic and homogeneous than the velocity covariances. We do not investigate this statement as it is beyond the scope of this work.

As stated, we propose using a CVT in NEMOVAR based on Helmholtz Theorem. Specifically, we suggest using ageostrophic streamfunction and velocity potential as the velocity control variables. These are used in a regional ocean DA system by Li et al. (2008*a,b*), discussed in section 3.4.2. We now discuss a difficulty in using these alternative control variables.

3.5.1 Treatment of boundary conditions

In order to use alternative velocity control variables based on Helmholtz Theorem in NEMOVAR, we must transform from the ageostrophic velocities to their irrotational and non-divergent parts - velocity potential and streamfunction. There is a difficulty in performing this transformation due to the need to explicitly specify boundary conditions for streamfunction in an ocean domain. This issue is caused by the discretisation of the model variables on the Arakawa C-grid (Arakawa and Lamb, 1977). The Arakawa C-grid

is used in the NEMO model (Madec et al., 2023) and commonly in ocean modelling. The traditional Arakawa C-grid used in Lynch (1989) and Watterson (2001), and shown in figure 3.2a, places the discretised streamfunction, ψ , in the corners of the grid.



(a) Arakawa C-grid as seen in Lynch (1989) and Watterson (2001) whereby streamfunction is located in the corner of each grid cell. (b) Alternative Arakawa C-grid presented by Li et al. (2006) whereby streamfunction is located in the centre of the grid cell.

Figure 3.2: Configurations of the Arakawa C-grid for the placement of streamfunction, ψ . The green shaded region represents land.

The placement of streamfunction in the corner of the grid implies that explicit boundary conditions must be imposed on streamfunction. However, specifying a value for the streamfunction at the boundary is challenging, as its exact value at the boundary is not known. The choice of boundary conditions for streamfunction has been discussed in many studies. Lynch (1989) tested 8 different combinations of boundary conditions on a limited domain using Sangster’s method for the partitioning problem (Sangster, 1960). However, when reconstructing the velocity from the computed streamfunction and velocity potential, the reconstruction error was poor (Li et al., 2006).

Alternatively, De Verdière and Ollitrault (2016) observe that the physical boundary condition of zero normal velocity on the solid boundaries can be imposed by forcing the streamfunction to be constant there. They use the velocity to calculate the vorticity and, with that, calculate the streamfunction by solving the Poisson equation (2.7). They note that in a simply connected domain, such as the North Atlantic and Pacific oceans,

streamfunction can be set to zero arbitrarily on the boundary. However, particularly in the southern hemisphere, the large islands mean that the value of streamfunction is undetermined. Webb (1996) highlights that island boundary conditions require integration around each island’s perimeter. Similarly, Han and Huang (2020) and Watterson (2001) impose constant streamfunction along solid boundaries in limited domains. Yet, as discussed by Li et al. (2006), enforcing constant streamfunction at the boundary implies non-divergent flow, which in turn assumes geostrophic balance—a condition not valid at coastlines. These considerations further underscore the need to avoid prescribing explicit boundary conditions for the streamfunction.

As the value of streamfunction can not be constant along the boundary, neither can the ageostrophic component. To circumvent having to specify explicit boundary conditions Li et al. (2006) proposed a solution - they moved streamfunction to the centre of the grid, as shown in figure 3.2b. We propose using this configuration of the Arakawa C-grid, presented by Li et al. (2006), to implement alternative velocity control variables in NEMOVAR - ageostrophic streamfunction and velocity potential. In this work we investigate implications of these control variables that have not been discussed before. The following chapters will answer the questions proposed in chapter 1 that are yet to be discussed in the literature.

3.6 Summary

In this chapter we introduced ocean data assimilation, with a focus on the variational data assimilation schemes used in the ocean. Variational DA formulates the data assimilation problem as a cost function to be minimised. This cost function combines the background and observations, weighted by their uncertainties, which are represented by error covariance matrices. By minimising the cost function, we determine the analysis. Variational data assimilation uses numerical minimisation techniques to iteratively minimise this cost function. We then presented the incremental formulation of variational data assimilation, where we solve the problem in terms of increments to the background, involving an inner and outer loop. This provides flexibility in reducing the computational cost of variational data assimilation as the inner loop can be run at a lower resolution to the outer loop. We then described the control variable transform, a transformation to control

space from model space where the control variables are considered less correlated than the model variables. These transformations are used within the incremental formulation of the data assimilation problem, and so we detailed the process of performing the CVT. Finally, we identified the CVTs currently used in both the atmosphere and ocean, specifically looking at the treatment of the velocity variables. We exposed the issues with the current velocity control variables and proposed new control variables for NEMOVAR. We now present results of investigating this new CVT using a shallow water model.

Chapter 4

Alternative velocity control variables using the shallow water equations

In this chapter we identify and examine a numerical issue with implementing ageostrophic streamfunction and velocity potential as velocity control variables. In particular, using a shallow-water model we focus on the numerical implication of transforming to these variables based on the discretisation proposed by Li et al. (2006).

We first introduce the shallow water equations (SWEs). These are partial differential equations that describe the flow of an incompressible fluid and are commonly used as a simplified model to examine ocean dynamics. We describe the discretisation of the shallow water model using a semi-implicit semi-Lagrangian scheme. We then detail how the CVT using the SWEs can be formulated to include the alternative velocity variables, ageostrophic streamfunction and ageostrophic velocity potential. We explain the discrete process of transforming to these variables from the ageostrophic velocities and the issues surrounding the boundary conditions. We conclude this section with a discussion of a numerical artefact discovered when performing this transformation - a checkerboard pattern.

4.1 Shallow water model

The ocean is a fluid that can be described to a good approximation by the primitive equations. Lions et al. (1992) describes how the primitive equations are fundamental

equations of the ocean. However, we consider an approximation simpler than the primitive equations. The SWEs can be obtained by integrating the primitive equations over a fluid layer depth and linearising about a basic state at rest. They describe a shallow layer of fluid and are very useful for modelling geophysical fluid dynamics (Pedlosky, 2013; Ghil and Malanotte-Rizzoli, 1991).




Figure 4.1: Schematic illustration of the 2D shallow water model.

In this study, we use a 2D shallow water model on a β -plane in mid-latitudes, based on the model of Stommel (1948). The model simulates the wind-driven circulation in a closed ocean basin, including a western boundary current. The model describes the horizontal flow of an incompressible fluid with a free surface and solid lower boundary. The vertical pressure gradient is assumed to be in hydrostatic balance. The Stommel model is described by the SWEs with the addition of linear drag and wind stress;

$$\frac{D\eta}{Dt} + H \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right) = 0, \quad (4.1)$$

$$\frac{Du}{Dt} - (f_0 + \beta y)v = -g \frac{\partial \eta}{\partial x} - \gamma u + \frac{\tau_x}{\rho_0 H}, \quad (4.2)$$

$$\frac{Dv}{Dt} + (f_0 + \beta y)u = -g \frac{\partial \eta}{\partial y} - \gamma v + \frac{\tau_y}{\rho_0 H}, \quad (4.3)$$

where

$$\frac{D}{Dt} \equiv \frac{\partial}{\partial t} + u \frac{\partial}{\partial x} + v \frac{\partial}{\partial y} \quad (4.4)$$

and η is the surface elevation, (x, y) is the 2D spatial coordinate and H is the ocean depth which is taken to be constant (see figure 4.1). We have (u, v) , the horizontal velocity vector, g is the gravitational acceleration, γ is a linear drag coefficient and ρ_0 is a uniform density. The approximation to the Coriolis parameter for a β -plane, is given by $f_0 + \beta y$. Here, $f_0 = 2\Omega \sin \theta_0$ and $\beta_0 = \frac{2\Omega}{a} \cos \varphi_0$ where Ω is the angular speed of the Earth, θ_0 is a reference latitude and a is the radius of the Earth. The SWEs are defined on the closed domain $[0, L_x] \times [0, L_y]$ with the following boundary conditions. On the eastern and western boundary we have

$$u(0, y, t) = u(L_x, y, t) = \frac{\partial v}{\partial x}(0, y, t) = \frac{\partial v}{\partial x}(L_x, y, t) = 0, \quad y \in [0, L_y] \quad \forall t.$$

On the northern and southern boundary we have

$$v(x, 0, t) = v(x, L_y, t) = \frac{\partial u}{\partial y}(x, 0, t) = \frac{\partial u}{\partial y}(x, L_y, t) = 0, \quad x \in [0, L_x] \quad \forall t.$$

The wind stress, $\boldsymbol{\tau}$, acts on the surface and is defined by (4.5), where the seasonality of the wind forcing is based loosely on Lévy et al. (2010):

$$\boldsymbol{\tau} = \begin{pmatrix} \tau_x \\ \tau_y \end{pmatrix} = \tau_n \begin{pmatrix} -\cos\left(\frac{2\pi y}{L_x}\right) \\ 0 \end{pmatrix} \quad \text{with} \quad \tau_n = \tau_i - \tau_s \cos\left(\frac{(t - t_{max})\pi}{t_{min} - t_{max}}\right), \quad (4.5)$$

where τ_i is the mean wind intensity, τ_s is the seasonal oscillation intensity, t is the current time in hours from initial time, t_{max} is the number of hours to reach the end of a chosen period of time and t_{min} is the number of hours to reach the mid-point of the chosen period (half of t_{max}).

4.1.1 Discretisation

Throughout this work, we use finite differences on a rectilinear grid, consistent with the discretisation available in NEMO/NEMOVAR and the grids they employ. We discretise the SWEs presented above using a semi-implicit semi-Lagrangian scheme (SISL). Two alternative schemes were considered but deemed not appropriate. The first was the Forward-Backward scheme (Matsuno, 1966; Beckers and Deleersnijder, 1993). This Eulerian scheme enforced a severe restriction on the time step, due to the propagation of inertia gravity waves and the Courant-Friedrichs-Lewy (CFL) criterion. To be able to increase the time step, we then implemented a semi-implicit Euler scheme. However,

this approach does not fully address instabilities arising from the advection terms, which are discretised using centred differences. In order to allow the solution of the SWEs to be unconditionally numerically stable, we use a semi-implicit semi-Lagrangian scheme. The SISL removes the stability constraint imposed by rapid gravity wave propagation and eliminates the non-linear instability caused by the non-linear advection terms that appear in the Eulerian framework (Durrant, 2013). The SISL scheme has less stringent stability conditions on the time step and more accurate space discretisation (Kalnay, 2003).




Figure 4.2: Arakawa C-grid.

Furthermore, as discussed in section 3.5.1, we use the Arakawa C-grid (Arakawa and Lamb, 1977), which staggers the variables such that velocity components, u and v , are at the mid-points on the eastern and northern boundaries of the grid cell, respectively, and elevation is at the centre of the grid cell, as shown in figure 4.2. The discrete boundary conditions follow Madec et al. (2023) whereby we enforce the physical boundary condition of zero normal velocity at the boundary. The tangential velocities are chosen to have zero normal gradient at the boundary, consistent with the tangential boundary condition adopted in the configuration of NEMO used in chapters 5 and 6. The tangential boundary conditions are discussed in section 5.1.2. The flow is initialised at rest with zero elevation.

Semi-implicit semi-Lagrangian Scheme

In this section, we describe how we implement the SISL scheme. The velocity divergence ($\nabla \cdot \mathbf{v}$) and pressure gradient ($\nabla \eta$) are terms essential to gravity wave propagation (Durran, 2013). Therefore, these are the terms of the SWEs we treat implicitly. We treat the rest of the terms explicitly at departure points. Firstly, we consider the equation for surface elevation (4.1) and discretise the derivatives. We indicate the arrival and departure points with the subscripts a and d , respectively:

$$\frac{\eta_a^{n+1} - \eta_d^n}{\Delta t} + H \left(\alpha_0 \left[\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right]_a^{n+1} + (1 - \alpha_0) \left[\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right]_d^n \right) = 0$$

where superscripts indicate the time-level. The coefficient α_0 is a time-weighting parameter that lies in the interval $[0.5, 1]$ (Lawless et al., 2003), which is multiplied to the velocity divergence. The departure points are interpolated to the arrival points using cubic interpolation. We then rearrange the above to have arrival points on the left hand side and known departure points on the right hand side:

$$\eta_a^{n+1} + H \Delta t \alpha_0 \left[\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right]_a^{n+1} = \eta_d^n - H \Delta t (1 - \alpha_0) \left[\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right]_d^n. \quad (4.6)$$

If we discretise (4.6) at a specific point on the grid, we have

$$\eta_{a_{i,j}}^{n+1} + H \Delta t \alpha_0 \left[\left(\frac{u_{a_{i+\frac{1}{2},j}}^{n+1} - u_{a_{i-\frac{1}{2},j}}^{n+1}}{\Delta x} \right) + \left(\frac{v_{a_{i,j+\frac{1}{2}}}^{n+1} - v_{a_{i,j-\frac{1}{2}}}^{n+1}}{\Delta y} \right) \right] = E_{i,j}^n$$

where E

$$E^n = \eta_d^n - H \Delta t (1 - \alpha_0) \left[\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right]_d^n.$$

Similar to the elevation, we discretise the momentum equations for the velocities, (4.2) and (4.3), by introducing weighting parameters α_1 and α_2 . We apply these to the pressure gradient terms:

$$u_a^{n+1} + g \Delta t \alpha_1 \left[\frac{\partial \eta}{\partial x} \right]_a^{n+1} = u_d^n + \Delta t (f_0 + \beta y) v_d^n - \Delta t \gamma u_d^n + \frac{\tau_x}{\rho H} - g \Delta t (1 - \alpha_1) \left[\frac{\partial \eta}{\partial x} \right]_d^n \quad (4.7)$$

and

$$v_a^{n+1} + g \Delta t \alpha_2 \left[\frac{\partial \eta}{\partial y} \right]_a^{n+1} = v_d^n - \Delta t (f_0 + \beta y) u_d^n - \Delta t \gamma v_d^n - g \Delta t (1 - \alpha_2) \left[\frac{\partial \eta}{\partial y} \right]_d^n. \quad (4.8)$$

Recall from (4.5) that $\tau_y = 0$. If we specify (4.7) and (4.8) at specific points on the grid we have,

$$u_{a_{i-\frac{1}{2},j}}^{n+1} + g \Delta t \alpha_1 \frac{\eta_{a_{i,j}}^{n+1} - \eta_{a_{i-1,j}}^{n+1}}{\Delta x} = Z_{i-\frac{1}{2},j}^n$$

where Z is the right hand side of (4.7):

$$Z^n = u_d^n + \Delta t(f_0 + \beta y)v_d^n - \Delta t\gamma u_d^n + \frac{\tau_x}{\rho H} - g\Delta t(1 - \alpha_1)\left[\frac{\partial\eta}{\partial x}\right]_d^n$$

and

$$v_{a_{i,j-\frac{1}{2}}}^{n+1} + g\Delta t\alpha_2 \frac{\eta_{a_{i,j}}^{n+1} - \eta_{a_{i,j-1}}^{n+1}}{\Delta y} = M_{i,j-\frac{1}{2}}^n$$

where M is the right hand side of (4.8):

$$M^n = v_d^n - \Delta t(f_0 + \beta y)u_d^n - \Delta t\gamma v_d^n - g\Delta t(1 - \alpha_2)\left[\frac{\partial\eta}{\partial y}\right]_d^n.$$

In order to solve this system of equations, we substitute the arrival points for the velocities into (4.6) to have an equation that only involves η implicit terms. This gives a final equation for the elevation,

$$(1 + 2C_1 + 2C_2)\eta_{a_{i,j}}^{n+1} - C_1(\eta_{a_{i+1,j}}^{n+1} + \eta_{a_{i-1,j}}^{n+1}) - C_2(\eta_{a_{i,j+1}}^{n+1} + \eta_{a_{i,j-1}}^{n+1}) = R_{i,j}^n \quad (4.9)$$

where

$$C_1 = \frac{H(\Delta t)^2 g \alpha_0 \alpha_1}{(\Delta x)^2}, \quad C_2 = \frac{H(\Delta t)^2 g \alpha_0 \alpha_2}{(\Delta y)^2}$$

and

$$R_{i,j}^n = E_{i,j}^n - \frac{H\Delta t\alpha_0}{\Delta x} \left(Z_{i+\frac{1}{2},j}^n - Z_{i-\frac{1}{2},j}^n \right) - \frac{H\Delta t\alpha_0}{\Delta y} \left(M_{i,j+\frac{1}{2}}^n - M_{i,j-\frac{1}{2}}^n \right).$$

Equation (4.9) is in the form $\mathbf{Ax} = \mathbf{b}$, as discussed in section 2.2, where the left hand side of (4.9) is the matrix-vector product $\mathbf{Ax}$ and $\mathbf{b} = R$. We solve this using the conjugate gradient (CG) method (see section 2.2.2).

We now describe how to implement ageostrophic streamfunction and ageostrophic velocity potential as velocity control variables in the shallow water model.

4.2 CVT using the shallow water equations

The chosen velocity control variables are ageostrophic streamfunction, $\delta\psi_U$ and ageostrophic velocity potential, $\delta\chi_U$. Here, the subscript U refers to the unbalanced (ageostrophic) part of the variable. We now describe the process of performing both the U-Transform and T-transform outlined in section 3.3, using the shallow water model.

4.2.1 U - Transform

To perform the U-transform, (3.12), to find the velocity model variables from the control variables,

$$\delta \mathbf{x} = \mathcal{U} \delta \mathbf{z}$$

where

$$\delta \mathbf{z} = \begin{pmatrix} \delta \psi_U \\ \delta \chi_U \end{pmatrix} \quad \text{and} \quad \delta \mathbf{x} = \begin{pmatrix} \delta u \\ \delta v \end{pmatrix},$$

we apply the following steps:

1. Calculate δu_U and δv_U from $\delta \psi_U$ and $\delta \chi_U$ using the Helmholtz equations:

$$\delta u_U = -\frac{\partial \delta \psi_U}{\partial y} + \frac{\partial \delta \chi_U}{\partial x}, \quad (4.10)$$

$$\delta v_U = \frac{\partial \delta \psi_U}{\partial x} + \frac{\partial \delta \chi_U}{\partial y}. \quad (4.11)$$

2. Compute the balanced velocity increments, δu_B and δv_B , using the geostrophic balance relation,

$$\delta u_B = -\frac{g}{f} \frac{\partial \delta \eta}{\partial y} \quad \text{and} \quad \delta v_B = \frac{g}{f} \frac{\partial \delta \eta}{\partial x}. \quad (4.12)$$

3. Compute the full velocity increments, δu and δv , from

$$\delta u = \delta u_B + \delta u_U \quad \text{and} \quad \delta v = \delta v_B + \delta v_U.$$

4.2.2 T - Transform

The T-transform is required to transform from the horizontal velocities to ageostrophic streamfunction and ageostrophic velocity potential. Although the T-transform is not directly used in (3.14) to calculate the analysis, the transformation is necessary for computing background error covariances of the control variables (Katz et al., 2011). The T-transform is given by

$$\delta \mathbf{z} = \mathcal{T} \delta \mathbf{x}$$

where $\delta \mathbf{z}$ and $\delta \mathbf{x}$ are as defined in the U-transform. We formulate this transformation as follows,

1. Compute the balanced velocity increments, δu_B and δv_B , from $\delta \eta$ using geostrophic balance (4.12).

2. Compute the ageostrophic velocity increments using,

$$\delta u_U = \delta u - \delta u_B \quad \text{and} \quad \delta v_U = \delta v - \delta v_B.$$

3. Compute the ageostrophic streamfunction, $\delta\psi_U$, and ageostrophic velocity potential, $\delta\chi_U$, from δu_U and δv_U by inverting the Helmholtz relations, (4.10) and (4.11).

As discussed in section 2.1, in a bounded domain, the Helmholtz decomposition is not necessarily unique as a non-zero harmonic function may be present. The harmonic component is zero if flow on the boundary is zero and the domain is simply connected (Bhatia et al., 2012). As described in section 4.1, our shallow water domain is a closed basin with zero normal velocity at the boundary but the tangential velocity is non-zero. Therefore, the Helmholtz decomposition of the velocities into their irrotational and non-divergent parts is not unique in this domain. We now describe how to perform this decomposition to ageostrophic velocity potential and ageostrophic streamfunction numerically and how to resolve the non-uniqueness of the decomposition.

4.2.3 The discrete transformation to streamfunction and velocity potential

In order to solve (4.10) and (4.11) to find ageostrophic streamfunction and velocity potential in the T-transform, we discretise the Helmholtz equations to form a linear system involving the horizontal velocities and their rotational and divergent parts (Li et al., 2006),

$$\mathbf{A}\delta\mathbf{z} = \delta\mathbf{v}_U \tag{4.13}$$

where

$$\delta\mathbf{z} = \begin{pmatrix} \delta\psi_U \\ \delta\chi_U \end{pmatrix}, \quad \delta\mathbf{v}_U = \begin{pmatrix} \delta u_U \\ \delta v_U \end{pmatrix}, \quad \mathbf{A} = \begin{pmatrix} -\Delta_y & \Delta_x \\ \Delta_x & \Delta_y \end{pmatrix}$$

with Δ_x and Δ_y representing the discretisation of $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$, respectively. For our application, $\mathbf{A}$ is not usually a square matrix because $\delta\mathbf{z}$ and $\delta\mathbf{v}_U$ are not required to be the same length. We discuss this further later in this section. Therefore, the linear system given by (4.13) may not be solved directly to find $\delta\mathbf{z}$. Instead, we find the least squares solution in the 2-norm by minimising the following cost function with respect to

$\delta \mathbf{z}$ (Li et al., 2006):

$$\mathcal{J}(\delta \mathbf{z}) = \frac{1}{2}(\delta \mathbf{v}_U - \mathbf{A}\delta \mathbf{z})^T(\delta \mathbf{v}_U - \mathbf{A}\delta \mathbf{z}). \quad (4.14)$$

An alternative approach to solving the above is detailed in appendix C.2, whereby we separate the calculation of $\delta \psi_U$ and $\delta \chi_U$, by solving the Poisson equations for divergence, (2.6), and vorticity, (2.7). In appendix D, we use Lagrangian multipliers to demonstrate how from Helmholtz theorem we can derive the Poisson equations for the continuous equations, putting the Li et al. (2006) approach in a continuous setting. The Lagrangian approach also reveals the boundary conditions we impose on the velocity, which we discuss below. In the discrete case we also derive similar equations, showing there are two approaches to solving this linear system. The first is obtained by minimising (4.14) with respect to $\delta \mathbf{z}$. The second involves the discretised Poisson equations, (2.6) and (2.7). Both these approaches are illustrated in appendix C. These approaches are equivalent and we use the former approach in this chapter.

In section 3.5.1, we discussed the placement of streamfunction on the Arakawa C-grid. In this work we adopt the arrangement proposed by Li et al. (2006), whereby streamfunction is located at the centre of grid cell, as shown in figure 3.2b. However, this grid arrangement confirms that $\mathbf{A}$ is not a square matrix. There is also the issue of the non-uniqueness of the Helmholtz decomposition. Li et al. (2006) suggest using Tikhonov's regularisation to overcome these problems. Tikhonov's regularisation method, discussed in section 2.2.1, involves adding a regularisation term to a cost function. Therefore, we add a regularisation term to (4.14) to give

$$\mathcal{J}_\mu(\delta \mathbf{z}) = \frac{1}{2}(\delta \mathbf{v}_U - \mathbf{A}\delta \mathbf{z})^T(\delta \mathbf{v}_U - \mathbf{A}\delta \mathbf{z}) + \frac{1}{2}\mu \delta \mathbf{z}^T \delta \mathbf{z} \quad (4.15)$$

where μ is the positive regularisation parameter and $\mathcal{J}_\mu(\delta \mathbf{z})$ is called Tikhonov's functional. Due to the regularisation term the cost function $\mathcal{J}_\mu(\delta \mathbf{z})$ has a unique minimum.

4.2.4 Boundary conditions for streamfunction

The method proposed by Li et al. (2006) is designed to avoid specifying explicit boundary conditions for streamfunction. By placing streamfunction at the centre of the grid cell, it is no longer necessary to impose a value of streamfunction at the boundary. Here, we demonstrate that boundary conditions still exist and how these appear in the

discretisation. Li et al. (2006) mention that the solution to the Helmholtz equation relates to Neumann boundary conditions, but do not discuss how these conditions emerge from the discretisation. Here, we provide a detailed explanation of these boundary conditions. Recall from section 4.2.3, the discretised Helmholtz decomposition is given by,

$$\delta u_U = -\Delta_y \delta \psi_U + \Delta_x \delta \chi_U \quad (4.16)$$

and

$$\delta v_U = \Delta_x \delta \psi_U + \Delta_y \delta \chi_U. \quad (4.17)$$

The discrete derivatives are defined in (4.13). The centred derivatives of velocity potential are located at the correct velocity points; however this is not the case for the centred derivatives of streamfunction. We need to calculate $\Delta_y \delta \psi_U$ and $\Delta_x \delta \psi_U$ at u and v points respectively. Their locations are shown on the Arakawa C-grid in figure 3.2. However, using a centred difference representation of the derivatives, these are not naturally located at these points. Instead, $\Delta_y \delta \psi_U$ is at a v -point and $\Delta_x \delta \psi_U$ at a u -point. Positioning the derivatives at the necessary velocity points requires averaging them. There are two possible approaches to achieve this,

Approach A:

1. Calculate $\Delta_y \delta \psi_U$ at the v -point by calculating the y -derivative of $\delta \psi_U$ using centred differencing.
2. Average these values to the u -points.
3. Calculate $\Delta_x \delta \psi_U$ at the u -point by calculating the x -derivative of $\delta \psi_U$ using centred differencing.
4. Average these values to the v -points.

Approach B:

1. Average $\delta \psi_U$ to the corner of the grid cells; $\delta \bar{\psi}_U$ represents the averaged values of $\delta \psi_U$ at the corner point (see figure 4.3).
2. Find the value of $\Delta_y \delta \psi_U$ at the u -point by calculating the y -derivative of $\delta \bar{\psi}_U$ using centred differencing.

3. Find the value of $\Delta_x \delta \psi_U$ at the v -point by calculating the x -derivative of $\delta \bar{\psi}_U$ using centred differencing.




Figure 4.3: Arakawa C-grid with the averaged values of streamfunction located at the corner of the grid cell.

Away from the boundaries, on a uniform grid, these two approaches are identical. We also find these approaches lead to an issue of grid-scale noise we call the checkerboard pattern. This will be discussed in detail in section 4.3.2. Near the boundaries, these approaches should also be equivalent but care must be taken in the treatment of the streamfunction at the boundary.

Suppose we want to calculate $u_{i-\frac{1}{2},j-1}$ using (4.16). Assume this point is near the south boundary as shown in figure 4.4. We have,

$$\delta u_{i-\frac{1}{2},j-1} = -[\Delta_y \delta \psi_U]_{i-\frac{1}{2},j-1} + \underbrace{[\Delta_x \delta \chi_U]_{i-\frac{1}{2},j-1}}_{=\frac{1}{\Delta x}(\delta \chi_{U,i,j-1} - \delta \chi_{U,i-1,j-1})} \quad (4.18)$$

Below we present both approaches to calculating the streamfunction derivative.

Approach A

For Approach A, we must first calculate $\Delta_y \delta \psi_U$ at the surrounding v -points. We have two points located away from the boundary which we have the following expressions for:

- $[\Delta_y \delta \psi_U]_{i,j-\frac{1}{2}} = \frac{1}{\Delta y}(\delta \psi_{U,i,j} - \delta \psi_{U,i,j-1})$ and
- $[\Delta_y \delta \psi_U]_{i-1,j-\frac{1}{2}} = \frac{1}{\Delta y}(\delta \psi_{U,i-1,j} - \delta \psi_{U,i-1,j-1}).$




Figure 4.4: The same grid as figure 3.2 with the land extended into the domain. The blue circle highlights point $u_{i-\frac{1}{2}, j-1}$ which we calculate using the discrete Helmholtz equations in this section.

However, the remaining two derivatives are located on the southern boundary, $[\Delta_y \delta \psi_U]_{i, j-\frac{3}{2}}$ and $[\Delta_y \delta \psi_U]_{i-1, j-\frac{3}{2}}$, and so we must impose a boundary condition on the normal derivatives of streamfunction. By setting them equal to zero we have Neumann boundary conditions. Therefore, when we average these derivatives to the u -point we have,

$$\begin{aligned}
 [\Delta_y \delta \psi_U]_{i-\frac{1}{2}, j-1} &= \frac{1}{4} \left([\Delta_y \delta \psi_U]_{i, j-\frac{1}{2}} + [\Delta_y \delta \psi_U]_{i-1, j-\frac{1}{2}} + [\Delta_y \delta \psi_U]_{i, j-\frac{3}{2}} + [\Delta_y \delta \psi_U]_{i-1, j-\frac{3}{2}} \right) \\
 &= \frac{1}{4\Delta y} \left((\delta \psi_{U_{i,j}} - \delta \psi_{U_{i,j-1}}) + (\delta \psi_{U_{i-1,j}} - \delta \psi_{U_{i-1,j-1}}) + 0 + 0 \right) \\
 &= \frac{1}{4\Delta y} \left(\delta \psi_{U_{i,j}} - \delta \psi_{U_{i,j-1}} + \delta \psi_{U_{i-1,j}} - \delta \psi_{U_{i-1,j-1}} \right).
 \end{aligned}$$

Approach B

For Approach B, we first average $\delta \psi_U$ to the corners of the grid cells. The first corner, which is located away from the boundary, can be easily expressed. However, the second corner lies near the boundary and is more challenging to define:

- $\delta \bar{\psi}_{U_{i-\frac{1}{2}, j-\frac{1}{2}}} = \frac{1}{4} \left(\delta \psi_{U_{i,j}} + \delta \psi_{U_{i-1,j}} + \delta \psi_{U_{i,j-1}} + \delta \psi_{U_{i-1,j-1}} \right)$ and
- $\delta \bar{\psi}_{U_{i-\frac{1}{2}, j-\frac{3}{2}}}$.

Specifying a value of $\delta\bar{\psi}_{U_{i-\frac{1}{2},j-\frac{3}{2}}}$ would be an explicit boundary condition which we aim to avoid due to the reasons discussed and because the method in Li et al. (2006) is specifically designed to circumvent this. Instead, we extrapolate the information from the nearby ocean points, assuming streamfunction is zero at land points. This is equivalent to having streamfunction with zero gradient across the boundary, i.e. Neumann boundary conditions. This gives,

$$\delta\bar{\psi}_{U_{i-\frac{1}{2},j-\frac{3}{2}}} = \frac{1}{2} \left(\delta\psi_{U_{i-1,j-1}} + \delta\psi_{U_{i,j-1}} \right). \quad (4.19)$$

Therefore, the derivative becomes,

$$\begin{aligned} [\Delta_y \delta\psi_U]_{i-\frac{1}{2},j-1} &= \frac{1}{\Delta y} \left(\delta\bar{\psi}_{U_{i-\frac{1}{2},j-\frac{1}{2}}} - \delta\bar{\psi}_{U_{i-\frac{1}{2},j-\frac{3}{2}}} \right) \\ &= \frac{1}{\Delta y} \left(\frac{1}{4} \left(\delta\psi_{U_{i,j}} + \delta\psi_{U_{i-1,j}} + \delta\psi_{U_{i,j-1}} + \delta\psi_{U_{i-1,j-1}} \right) - \frac{1}{2} \left(\delta\psi_{U_{i-1,j-1}} + \delta\psi_{U_{i,j-1}} \right) \right) \\ &= \frac{1}{4\Delta y} \left(\delta\psi_{U_{i,j}} - \delta\psi_{U_{i,j-1}} + \delta\psi_{U_{i-1,j}} - \delta\psi_{U_{i-1,j-1}} \right). \end{aligned}$$

In conclusion, Approach A and B are equivalent at the boundaries when we impose Neumann boundary conditions to streamfunction; i.e. the normal derivatives of streamfunction are zero at the boundary. Therefore, although this method avoids the need for explicit boundary conditions, implicit boundary conditions exist.

Having presented the theory behind the transformation to the alternative velocity control variables, in the next section we implement this numerically using the SWEs and discuss a computational issue that arises.

4.3 Numerical implication

In this section, we discuss the numerical implementation of the CVT proposed and outline a numerical problem we have discovered - a checkerboard pattern. To illustrate the transformation numerically we use the SWEs, described in section 4.1.

4.3.1 Method

We run the discretised model using the SWEs, presented in section 4.1, with a model time step $\Delta t = 1800s$. We chose a domain of size $L_x, L_y = 10^6m$ with grid spacing

Parameter	Symbol	Value
Coriolis parameter	f_0	$10^{-4}s^{-1}$
β -coefficient	β	$10^{-11}m^{-1}s^{-1}$
Gravitational acceleration	g	$10ms^{-1}$
Linear drag coefficient	γ	$10^{-6}s^{-1}$
Constant density	ρ_0	$1000kgm^{-3}$
Resting depth	H	$1000m$
Mean wind intensity	τ_i	$0.02kgm^{-1}s^{-2}$
Seasonal wind intensity	τ_s	$0.01kgm^{-1}s^{-2}$
Wind stress max	τ_{max}	720hrs (30 days)
Wind stress min	τ_{min}	360hrs (15 days)

Table 4.1: Physical parameter values for the SWEs.

$\Delta x = \Delta y = 10^4m$, similar to Stommel (1948). The values of the model parameters chosen for the SWEs are detailed in table 4.1. To allow the development of dynamical instabilities in the shallow water model, the linear drag coefficient is set to a small value compared to the magnitude of the velocities. While linear drag acts to dampen velocity and suppress instabilities, applying only weak drag allows instability to emerge. We choose a wind intensity that varies with a period of one month. The time weighting parameters in the semi-Lagrangian scheme are chosen to be 0.6. This enables us to be close to second order accuracy while still ensuring stability (Lawless et al., 2003). We run the discrete model for 331 days, with the fields at the final time step shown in figure 4.5. We can see the model has spun-up with a cyclonic/anticyclonic pair of eddies in the west side of the domain.

In variational data assimilation, the Helmholtz relation is applied to assimilation increments, $\delta\psi$ and $\delta\chi$, in the CVT. Therefore, as a proxy for an assimilation increment, we calculate the difference between the velocities over one day and transform to streamfunction and velocity potential increments using the process described in section 4.2.2. We add an additional component to step 2 of this process, whereby we apply the Shapiro filter to the ageostrophic velocity increments (see section 2.3 for the filter definition) to remove any noise caused by the geostrophic balance calculation, appendix




Figure 4.5: Output from running the discretised shallow water model for 331 days. The left is elevation (m) and the right is speed (ms^{-1}) with velocity vectors superimposed.

B discusses why the filter is necessary here. Moreover, we minimise (4.14) iteratively using the CG method. Appendix E details the tests performed to ensure the adjoint and gradient implemented are correct. For the CG method, we chose the following convergence criterion:

$$\|\nabla \mathcal{J}_\mu(\delta \mathbf{z}_k)\|_2 < 10^{-5} \|\nabla \mathcal{J}_\mu(\delta \mathbf{z}_0)\|_2, \quad (4.20)$$

where k is an iteration of the CG method and $\delta \mathbf{z}_0$ contains the initial guess for $\delta \psi$ and $\delta \chi$, which we set to be zero. We choose the regularisation parameter as $\mu = 10^{-12} m^{-2}$. A series of sensitivity experiments show that this value allows for efficient convergence whilst still providing an accurate solution. However, more sophisticated techniques, such as the L-curve, may be used and are described in Hansen (1997).

After minimising (4.15) for $\delta \psi_U$ and $\delta \chi_U$ (fields shown in figure 4.6) we reconstruct the ageostrophic velocity increments using (4.10) and (4.11). We now discuss a grid-scale noise issue associated with this decomposition.

4.3.2 Checkerboard pattern

Figure 4.7 presents the velocity reconstruction errors, which are relatively small in magnitude compared to the velocities themselves. This demonstrates that the transformation presented in section 4.2 is successful. However, there is the presence of grid-scale noise which we now discuss.




Figure 4.6: The ageostrophic streamfunction and velocity potential increments, calculated using the SWEs. These are calculated using Helmholtz equations, from the ageostrophic velocity increments. On the left we have the streamfunction increment and on the right the velocity potential increment (m^2s^{-1}).

A computational mode is a solution or pattern caused by a discretisation that does not appear in the continuous solution. In this section we demonstrate a computational mode that arises during the computation of streamfunction derivatives. It is caused by a loss of information due to averaging. For example, at $\delta u_{U, i+\frac{1}{2}, j}$ (using the grid shown in figure 3.2b, this point is away from the boundary) we have,

$$\delta u_{U, i+\frac{1}{2}, j} = -\left(\Delta_y \delta \psi_U\right)_{i+\frac{1}{2}, j} + \left(\Delta_x \delta \chi_U\right)_{i+\frac{1}{2}, j}. \quad (4.21)$$

When calculating $\Delta_y \delta \psi_U$ through either approach A or B described in section 4.2.4, this gives

$$\left(\Delta_y \delta \psi_U\right)_{i+\frac{1}{2}, j} = \frac{1}{4\Delta y} \left(\delta \psi_{U, i, j+1} - \delta \psi_{U, i, j-1} + \delta \psi_{U, i+1, j+1} - \delta \psi_{U, i+1, j-1} \right). \quad (4.22)$$

We obtain a similar result for $\Delta_x \delta \psi_U$ at $v_{i, j+\frac{1}{2}}$. The resulting derivatives do not use any information from the adjacent values of streamfunction, i.e. points (i, j) and $(i+1, j)$. We can clearly see this numerical artefact using the SWEs. The checkerboard pattern is apparent in the reconstructed velocities, as seen in figure 4.7. Figure 5 in Li et al. (2006) shows the same problem, although it was not discussed in that article.

This issue of grid-scale noise is discussed throughout literature due to various discretisation choices. For example, the unstaggered Arakawa-A grid can cause horizontal uncoupling due to calculating over distances twice the grid space length (Kalnay, 2003;




Figure 4.7: The difference between the original ageostrophic velocity increments, calculated using the SWEs, and the reconstructed ageostrophic velocity increments. These are reconstructed using Helmholtz equations, from the ageostrophic streamfunction and velocity potential increments. On the left we have the zonal velocity increment error and on the right the meridional velocity increment error (ms^{-1}).

Skamarock and Klemp, 2008). The Arakawa-C grid requires spatial averaging for the Coriolis terms which also causes this computational mode (Marshall et al., 1997; Weller et al., 2012).

The checkerboard pattern caused by streamfunction is particularly problematic when estimating background error covariances for this variable in data assimilation applications. Typically, this would involve computing error samples of ψ (and χ) by minimising (4.15) given error samples of u and v (e.g., from ensembles). Some background error correlation models, such as those based on a diffusion operator, require an estimate of the correlation aspect tensor, which can be obtained using an approach that requires numerically approximating the first-order directional derivatives of the error field (Weaver et al., 2021; Chrust et al., 2024). Such an estimate would be severely contaminated by grid-scale noise. We now demonstrate how the checkerboard pattern can appear in the statistics of the control variables.

4.3.3 Spatial correlations

As discussed in the previous section, we need to calculate statistics of the control variables as part of the CVT to provide the background error covariances for assimilation. Here we calculate the spatial correlations of the ageostrophic streamfunction increments field to highlight the presence of the checkerboard pattern.

To examine the spatial correlations of the ageostrophic streamfunction increments, we pick one particular point in the domain, and then calculate the correlation of each grid point with that chosen point using the ‘quick covs method’ (Polavarapu et al., 2005) as follows:

1. Run the shallow water model described in section 4.1 for 200 days after spinning the model up for 200 days.
2. Calculate the day-to-day model differences as a proxy for the background error, $\delta\eta$, δu and δv .
3. Calculate the current velocity control variables, δu_U and δv_U , using geostrophic balance (4.12), for each daily difference field and then calculate the new velocity control variables, $\delta\psi_U$ and $\delta\chi_U$, using the T-transform, for each daily increment.
4. Calculate the correlation between the error at the chosen point in the domain and the error in each other grid point for the streamfunction field.




Figure 4.8: Locations of the points chosen for the dominant signal analysis performed using the shallow water model. The underlying fields are generated from running the discretised shallow water model for 331 days.

The Met Office FOAM system currently utilises operational forecasts with a 24-hour interval (Blockley et al., 2014). We adopt this interval length in our correlation experiment. To assess whether this interval is suitable for our shallow water model, we investigate the sensitivity of the model to the daily forecast interval by evaluating whether there is a dominant signal present in the forecasts. Our objective is to ensure that the computed statistics reflect the characteristics of the forecast errors rather than those of the model’s inherent variability. A similar approach is adopted by (Katz et al., 2011),

who identify and remove the diurnal cycle to prevent it from dominating the statistical analysis.

We do so by generating a forecast and plotting the values of the variables at fixed points in space against time, these fixed points are shown in figure 4.8. For each variable, we select different points across the domain with varying levels of activity, ensuring a diverse sampling of locations. We perform this analysis for elevation and the velocities. In figure 4.9 we show the values of the zonal velocity at these points for a range of time frames. The absence of periodic waves, oscillations, or repeating patterns suggests that there are no signals present in u that would dominate the statistics. Elevation and meridional velocity show very similar result with no dominant signals. As such, the choice of a daily time interval is deemed appropriate for this work.




Figure 4.9: Zonal velocity values at fixed spatial points (with coloured lines matching the dots in second panel of figure 4.8) are plotted against time for different intervals.

Using the daily model differences as a proxy for the errors, we calculate the proposed velocity control variables, $\delta\psi_U$ and $\delta\chi_U$, using the process described in section 4.2.2, for each daily error field. Finally, we calculate the correlation between the chosen point in the domain and each other grid point for the ageostrophic streamfunction error field.

In figure 4.10a we show the resulting correlation field generated by the method described, where our chosen point was (400km, 400km). This plot shows that the spatial correlations of streamfunction for this particular point exhibit the checkerboard pattern. We see the large-scale correlation fields with the grid-scale noise present. This demonstrates that the checkerboard pattern in the streamfunction field is being carried through to the statistics calculated of it. As discussed previously this would have a direct impact on the assimilation if such estimates of the errors were used to specify the background error covariances. To mitigate this issue, we apply a Shapiro filter.




Figure 4.10: Spatial correlation of point (400km, 400km) of the ageostrophic streamfunction increments. (a) No filtering . (b) Shapiro filter is applied to ageostrophic streamfunction increments as a post-processing.

4.3.4 The Shapiro filter

One way to mitigate this checkerboard issue is to apply a filter. A potential filter is the Shapiro filter, described in section 2.3. An approach to using the Shapiro filter would be to apply the filter within the Helmholtz transformation. However this is unable to remove the checkerboard pattern. We demonstrate this by applying the filter to streamfunction directly within the Helmholtz transformation. This involves averaging streamfunction

values to the corner of the grid cells and then calculating the necessary derivatives, this follows Approach B in section 4.2.4. We can apply the Shapiro filter to the averaged values of streamfunction at the corner of the grid cells. If we aim to calculate $\Delta_y \delta \psi_U$ at $\delta u_{U_{i+\frac{1}{2},j}}$ as in section, 4.2.4 we have:

$$\left(\Delta_y \delta \psi_U \right)_{i+\frac{1}{2},j} = \frac{1}{\Delta y} (\delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{1}{2}}} - \delta \bar{\psi}_{U_{i+\frac{1}{2},j-\frac{1}{2}}}) \quad (4.23)$$

where the averaged values of streamfunction are located at the corner of the grid cells (figure 4.3). We apply the Shapiro filter to these corner points to give

$$\mathcal{F} \delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{1}{2}}} = \frac{1}{16} \left(4\delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{1}{2}}} + 2(\delta \bar{\psi}_{U_{i-\frac{1}{2},j+\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{3}{2},j+\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{1}{2},j-\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{3}{2}}}) + \right. \\ \left. \delta \bar{\psi}_{U_{i+\frac{3}{2},j+\frac{3}{2}}} + \delta \bar{\psi}_{U_{i+\frac{3}{2},j-\frac{1}{2}}} + \delta \bar{\psi}_{U_{i-\frac{1}{2},j+\frac{3}{2}}} + \delta \bar{\psi}_{U_{i-\frac{1}{2},j-\frac{1}{2}}} \right)$$

and

$$\mathcal{F} \delta \bar{\psi}_{U_{i+\frac{1}{2},j-\frac{1}{2}}} = \frac{1}{16} \left(4\delta \bar{\psi}_{U_{i+\frac{1}{2},j-\frac{1}{2}}} + 2(\delta \bar{\psi}_{U_{i-\frac{1}{2},j-\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{3}{2},j-\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{1}{2},j-\frac{3}{2}}} + \delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{1}{2}}}) + \right. \\ \left. \delta \bar{\psi}_{U_{i+\frac{3}{2},j+\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{3}{2},j-\frac{3}{2}}} + \delta \bar{\psi}_{U_{i-\frac{1}{2},j+\frac{1}{2}}} + \delta \bar{\psi}_{U_{i-\frac{1}{2},j-\frac{3}{2}}} \right).$$

These combine in (4.23) to give

$$\left(\Delta_y \delta \psi_U \right)_{i+\frac{1}{2},j} = \frac{1}{16\Delta y} \left(2(\delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{3}{2}}} - \delta \bar{\psi}_{U_{i+\frac{1}{2},j-\frac{1}{2}}} - \delta \bar{\psi}_{U_{i+\frac{1}{2},j-\frac{3}{2}}}) + \right. \\ \left. \delta \bar{\psi}_{U_{i-\frac{1}{2},j+\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{3}{2},j+\frac{1}{2}}} + \delta \bar{\psi}_{U_{i+\frac{3}{2},j+\frac{3}{2}}} + \delta \bar{\psi}_{U_{i-\frac{1}{2},j+\frac{3}{2}}} - \right. \\ \left. \delta \bar{\psi}_{U_{i+\frac{3}{2},j-\frac{1}{2}}} - \delta \bar{\psi}_{U_{i-\frac{1}{2},j-\frac{1}{2}}} - \delta \bar{\psi}_{U_{i+\frac{3}{2},j-\frac{3}{2}}} - \delta \bar{\psi}_{U_{i-\frac{1}{2},j-\frac{3}{2}}} \right). \quad (4.24)$$

As discussed, the streamfunction values have been averaged from the centre of the grid cell to the corner, for example

$$\delta \bar{\psi}_{U_{i+\frac{1}{2},j+\frac{1}{2}}} = \frac{1}{4} (\delta \psi_{U_{i,j+1}} + \delta \psi_{U_{i+1,j+1}} + \delta \psi_{U_{i,j}} + \delta \psi_{U_{i+1,j}}). \quad (4.25)$$

By substituting the un-averaged values of streamfunction, as given by (4.25), into (4.24) we get

$$\left(\Delta_y \delta \psi_U \right)_{i+\frac{1}{2},j} = \frac{1}{64\Delta y} \left[6 \left(\delta \psi_{U_{i,j+1}} - \delta \psi_{U_{i,j-1}} + \delta \psi_{U_{i+1,j+1}} - \delta \psi_{U_{i+1,j-1}} \right) + \right. \\ \left. 3 \left(\delta \psi_{U_{i,j+2}} - \delta \psi_{U_{i,j-2}} + \delta \psi_{U_{i+1,j-2}} - \delta \psi_{U_{i+1,j-1}} \right) + 2 \left(\delta \psi_{U_{i-1,j+1}} - \delta \psi_{U_{i-1,j-1}} + \delta \psi_{U_{i+2,j+1}} - \delta \psi_{U_{i+2,j-1}} \right) \right. \\ \left. + \left(\delta \psi_{U_{i-1,j+2}} - \delta \psi_{U_{i-1,j-2}} + \delta \psi_{U_{i+2,j+2}} - \delta \psi_{U_{i+2,j-2}} \right) \right]. \quad (4.26)$$

The resulting derivative still does not incorporate information from adjacent streamfunction values, $\delta\psi_{U_{i,j}}$ and $\delta\psi_{U_{i+1,j}}$, and as such, the checkerboard pattern has not been eliminated. Hence, we consider an alternative approach in which we apply one iteration of the Shapiro filter to $\delta\psi_U$ after the transformation from the ageostrophic velocity increments. The full workflow of transforming to the alternative velocity control variables and applying the Shapiro filter to the ageostrophic streamfunction is illustrated in figure 4.11 and discussed in appendix C. We apply Neumann boundary conditions to streamfunction when applying the Shapiro filter.




Figure 4.11: Workflow to finding the ageostrophic streamfunction and velocity potential increments with the Shapiro filter after the transformation. From the model output we find the daily increments, then use geostrophic balance to find the ageostrophic velocity increments and apply the Shapiro filter. We then find the the ageostrophic streamfunction and velocity potential increments through the inverse of the Helmholtz decomposition. We then again apply the Shapiro filter to the ageostrophic streamfunction increment.

The errors in the reconstructed velocities, when we apply the Shapiro filter to the ageostrophic streamfunction increment as a post-process, no longer display the




Figure 4.12: The difference between the original ageostrophic velocity increments, calculated using the SWEs, and the reconstructed ageostrophic velocity increments. These are reconstructed using Helmholtz equations, from the ageostrophic streamfunction and velocity potential increments where the Shapiro filter has been applied to streamfunction. On the left we have the zonal velocity increment error and on the right the meridional velocity increment error (ms^{-1}).

checkerboard pattern and neither do the resulting spatial correlations (figure 4.10b). The Shapiro filter also has little effect on the magnitude of the spatial correlations, with an average absolute magnitude of 0.269 with no filter and 0.275 after the filter has been applied. However, the magnitudes of the reconstruction errors, as seen in figure 4.12, increase significantly compared to the original errors shown in figure 4.7. This is to be expected as we are altering the streamfunction field before reconstructing the velocities. However, it is a significantly large increase that requires further investigation, which we conduct in the section 5.4.1.

4.4 Summary and conclusions

In this chapter we have detailed a new CVT using the SWEs, discussed the implications of this CVT on the boundary conditions of the control variables and discovered a major numerical issue with this transformation from the ageostrophic velocities to ageostrophic streamfunction and velocity potential. We provide both analytical and numerical discussions for all.

Initially, we set-up the shallow water model and discretised it using a semi-implicit Semi-Lagrangian scheme. We then detailed the transformation from the velocities to the

proposed velocity variables, using this shallow water model. A method proposed in the literature, to transform velocities into their irrotational and non-divergent components, eliminates the need for imposing explicit boundary conditions on the streamfunction. However, we first demonstrate that implicit boundary conditions are still present for streamfunction. Furthermore, we have shown theoretically and then numerically using the SWEs on a β -plane, that there is a numerical artefact arising from this transformation. The calculation of streamfunction derivatives within the Helmholtz equations results in a checkerboard pattern, stemming from a loss of information during the decomposition. When performing the T-transform, from the ageostrophic velocities to ageostrophic streamfunction and ageostrophic velocity potential, we have proposed applying the Shapiro filter to the ageostrophic streamfunction increment after the transformation as a practical remedy. We investigated a consequence of the checkerboard pattern on the spatial correlations of streamfunction. These clearly showed the presence of the checkerboard pattern, illustrating how the grid-noise could pollute any statistics calculated.

In appendix A, we present preliminary work on calculating cross-correlations between the velocity control variables, providing a brief examination of the uncorrelated assumption used in the CVT for the proposed control variables, geostrophic streamfunction and the velocity potential. The appendix highlights the presence of a checkerboard pattern in the cross-correlation fields, similar to that observed in the auto-correlation fields. Our findings also indicate that the proposed variables are in fact more strongly cross-correlated than the ageostrophic velocity components. Possible contributors to these correlations are the use of Tikhonov’s regularisation in the modified Helmholtz transformation or the choice of method to calculate these correlations. Despite of this and our discovery of a numerical issue with this transformation to alternative velocity variables, the overall potential benefits discussed in section 3.5 are still worth pursuing in NEMOVAR. We do not pursue the correlation analysis further in this work. However, a more detailed study is needed to draw definitive conclusions about the correlations between these variables, which would be extremely valuable to employing these variables in NEMOVAR.

In the next chapter we implement these alternative velocity variables in the balance operator used in NEMOVAR and demonstrate the transformation using a configuration

of NEMO. We identify the numerical issue discovered here and provide a more extensive investigation into the solution.

Chapter 5

Alternative velocity control variables in NEMOVAR

In this chapter, we integrate the proposed velocity control variables — ageostrophic streamfunction and velocity potential — into the NEMOVAR balance operator. We assess the performance of the transformation in a realistic data assimilation system by testing the implementation using the GYRE configuration of NEMO.

Firstly, we introduce the GYRE configuration, and the configuration set up. We then describe how the balance operator currently used in NEMOVAR can be updated to use ageostrophic streamfunction and velocity potential as the velocity control variables. The process of transforming from the velocities to their non-divergent and irrotational parts differs from that used in the simpler shallow water domain, and here we detail these differences. Next, we demonstrate the transformation from the velocity increments to streamfunction and velocity potential numerically using the GYRE configuration. This involves sensitivity experiments to determine certain parameter settings. Finally, we discuss a numerical implication of this transformation that appeared using the shallow water model and a potential solution. This chapter extends the work from chapter 4 to the NEMOVAR data assimilation code for application with NEMO.

5.1 Nucleus for European Modelling of the Ocean

Nucleus for European Modelling of the Ocean (NEMO) is a framework for ocean modelling (Madec et al., 2023). In this and the following chapter we use NEMO 4.0.4. NEMOVAR (Mogensen et al., 2009) is the variational data assimilation software for the NEMO model introduced in section 3.4.3. It is used operationally at the UK Met Office (Waters et al., 2015; Mignac et al., 2025) and ECMWF (Balmaseda et al., 2013; Chrust et al., 2024). In this chapter, we update the current balance operator used in NEMOVAR to include alternative velocity control variables. In order to demonstrate the transformation from the velocities to their non-divergent and irrotational parts, we use model output generated by the so-called GYRE configuration of NEMO.

5.1.1 Coordinate system in NEMO

All configurations in NEMO use a deformed curvilinear coordinate system on a sphere. As described in Madec et al. (2023) and shown in figure 5.1, NEMO has a set of orthogonal curvilinear coordinates (i, j, k) . The geographical coordinate system given by (λ, φ, z) describes the longitude, $\lambda(i, j)$, latitude, $\varphi(i, j)$ and depth below reference sea level, $z(k)$, respectively. Spherical coordinate systems create a singularity at the North Pole, which causes challenges in ocean models. To mitigate this, NEMO modifies the coordinate system to place the singularity over land. In the ORCA family of grids used by NEMO global configurations, a tripolar grid is employed whereby the two northern hemisphere poles are placed over land while the south pole remains over land in Antarctica. The transformation of the coordinate system over the pole introduces scaling functions into the model equations. These are described in section 1.3.1 of Madec et al. (2023). Throughout this section we will present the equations in spherical coordinates, without this transformation.

5.1.2 GYRE configuration

The GYRE configuration of the NEMO ocean model uses the primitive and mass continuity equations to simulate the seasonal cycle of a double-GYRE box model (Lévy et al., 2010). We have the following variables: $\mathbf{v} = (\mathbf{v}_h, w)$, the velocity vector, where




Figure 5.1: The geographical coordinate system (λ, φ, z) and the curvilinear coordinate system (i, j, k) used in NEMO. Figure 1.2 from Madec et al. (2023).

$\mathbf{v}_h = (u, v)$ is the horizontal velocity vector and w is the vertical velocity, T , temperature; S , salinity and ρ is the in-situ density given by $\rho = \rho(T, S, p)$ where p is the pressure. The primitive equations in NEMO are defined as follows,

$$\left[\frac{D\mathbf{v}_h}{Dt} \right]_h = -f\mathbf{k} \times \mathbf{v}_h - \frac{1}{\rho_0} \nabla_h p + D^v + F^v, \quad (5.1)$$

$$\frac{\partial T}{\partial t} = -\nabla \cdot (T\mathbf{v}) + D^T + F^T, \quad (5.2)$$

$$\frac{\partial S}{\partial t} = -\nabla \cdot (S\mathbf{v}) + D^S + F^S, \quad (5.3)$$

where

$$\left[\frac{D}{Dt} \right]_h \equiv \frac{\partial}{\partial t} + \frac{u}{a \cos \varphi} \frac{\partial}{\partial \lambda} + \frac{v}{a} \frac{\partial}{\partial \varphi} \quad (5.4)$$

and t is time; ρ_0 is the reference density; $f = 2\Omega \cdot \mathbf{k}$ is the Coriolis acceleration (where Ω is the Earth's angular velocity vector); D^v, D^T, D^S are the parametrisations of small-scale physics; F^v, F^T, F^S are surface forcing terms and a is the radius of the earth. Additional assumptions are made when formulating these equations. These are described in section 1.1.1 of Madec et al. (2023). The GYRE configuration uses a free surface formulation described in section 1.2.2 of Madec et al. (2023). This introduces an additional model variable: sea surface height (SSH), η , which describes the shape of the air-sea interface.

NEMO is discretised on an Arakawa C-grid, similar to the shallow water model presented in section 4.1. The centre of each grid cell is referred to as the T-point, and the corner as the f-point, as seen in figure 5.2.

Configuration set-up

The primitive equations for the GYRE configuration are solved in an idealised 3D domain and are forced by an analytical function of seasonal forcing. The configuration allows the investigation of the spontaneous generation of a large number of interacting transient mesoscale eddies. The domain is a closed rectangular basin on the β -plane, bounded by vertical walls and a flat bottom. The GYRE configuration is used to represent an idealised north Atlantic basin, eventually developing a western boundary current reminiscent of the Gulf Stream.

The circulation is forced by wind and buoyancy fluxes which vary seasonally in a sinusoidal manner (with winter and summer extrema). The wind stress is zonal. The model is initiated from rest with vertical profiles of temperature and salinity, which vary with depth but are horizontally uniform, across the whole domain. The model takes a few years to spin up, the procedure we use to spin up the model is described in section 5.3.1. The model simulates sea surface height (elevation), horizontal and vertical velocities, salinity and temperature on a 3D grid. The GYRE configuration has 31 vertical levels with the first layers located at $5m$, $15m$, $26m$, $36m$, $48m$, $60m$, and $75m$ depth. The horizontal resolution can be set to any value between 1° to $1/54^\circ$.

We must prescribe boundary conditions to both the normal and tangential velocities. There is zero flow across the boundary, i.e. the velocity normal to the boundary is zero. This is a physical boundary condition at coastlines. There are several options for the tangential boundary conditions in NEMO. These are described fully in Madec et al. (2023) and illustrated in figure 5.3,

- (a) Free-slip: normal derivative of the tangential velocity is zero at the boundary;
- (b) No-slip: tangential velocity is zero at the boundary;
- (c) Partial free-slip: the tangential velocity is smaller than the offshore velocity;
- (d) Strong no-slip: no-slip but with higher friction.

The GYRE configuration uses free-slip conditions at all boundaries. This is commonly used in the global ocean. However, Guiavarc’h et al. (2025) describe the use of partial slip boundary conditions in the Southern Ocean in the Met Office global configurations.




Figure 5.2: The horizontal grid used in the NEMO model. Figure 3.3 from Madec et al. (2023).




Figure 5.3: Lateral boundary conditions, (a) free-slip, (b) no-slip, (c) partial free-slip (d) strong no-slip. Figure 7.2 from Madec et al. (2023).

In the following section we describe how to implement the alternative velocity control variables in the NEMOVAR balance operator.

5.2 Variable transformations

As described in section 3.4.2, at present the ageostrophic velocities are used as the velocity control variables in NEMOVAR and we hypothesise that these are highly correlated. In this section, we describe how to implement the alternative velocity control variables, ageostrophic streamfunction and velocity potential increments, in NEMOVAR.

5.2.1 Balance operator

We propose a modification to the linearised balance operator (3.18) used in NEMOVAR. The balance operator, $\mathbf{K}$, is used to describe the balance relationship between variables and the inverse is designed to decorrelate the variables (Weaver et al., 2005). We use the same balance operator given by (3.18) but add an additional transformation to the final

equation - splitting the ageostrophic horizontal velocities into to their non-divergent and irrotational parts given by streamfunction and velocity potential. This gives

$$\begin{aligned}
\delta T &= \delta T, \\
\delta S &= \mathbf{K}_{ST} \delta T + \delta S_U := \delta S_B + \delta S_U, \\
\delta \eta &= \mathbf{K}_{\eta\rho} \delta \rho + \delta \eta_U := \delta \eta_B + \delta \eta_U, \\
\delta \mathbf{v}_h &= \mathbf{K}_{vp} \delta p + \delta \mathbf{v}_{h_U} := \delta \mathbf{v}_{h_B} + \delta \mathbf{v}_{h_U} = \delta \mathbf{v}_{h_B} + \delta \mathbf{v}_{\psi_U} + \delta \mathbf{v}_{\chi_U}
\end{aligned} \tag{5.5}$$

where $\delta \rho$ and δp are as described in section 3.4.3, the horizontal components of the velocities are denoted by $_h$, and $\delta \mathbf{v}_{\psi_U}$ and $\delta \mathbf{v}_{\chi_U}$ are the non-divergent and irrotational components of the velocity increments, respectively. The ageostrophic velocity increments are decomposed using the Helmholtz decomposition, which we introduced in section 2.1. The Helmholtz equations are as such:

$$\delta u = -\frac{1}{a} \frac{\partial \delta \psi}{\partial \varphi} + \frac{1}{a \cos \varphi} \frac{\partial \delta \chi}{\partial \lambda} \quad \text{and} \quad \delta v = \frac{1}{a \cos \varphi} \frac{\partial \delta \psi}{\partial \lambda} + \frac{1}{a} \frac{\partial \delta \chi}{\partial \varphi}. \tag{5.6}$$

For simplicity, we formulate the Helmholtz equations (5.6) in spherical coordinates. As discussed in section 5.1.1, NEMO uses a curvilinear coordinate system (i, j, k) which can be transformed to from the spherical coordinates (λ, φ, z) . This transformation is detailed in section 1.3.1 of Madec et al. (2023).

The new control vector is given by $\delta \mathbf{z} = (\delta T, \delta S_U, \delta \eta_U, \delta \psi_U, \delta \chi_U)^T$. The inverse of the balance operator, $\mathbf{K}^{-1}$, is used to compute $\delta \mathbf{z}$ from $\delta \mathbf{x}$:

$$\begin{aligned}
\delta T &= \delta T, \\
\delta S_U &= \delta S - \delta S_B, \\
\delta \eta_U &= \delta \eta - \delta \eta_B, \\
\delta \psi_U &= \mathcal{H}_\psi (\delta \mathbf{v}_h - \delta \mathbf{v}_{h_B}), \\
\delta \chi_U &= \mathcal{H}_\chi (\delta \mathbf{v}_h - \delta \mathbf{v}_{h_B})
\end{aligned}$$

where $\mathcal{H}_\psi$ and $\mathcal{H}_\chi$ represent the inverse Helmholtz decomposition operators, which reconstruct the velocity field from its rotational and divergent components. These should not be confused with the observation operators: the non-linear operator h or its linearised form $\mathbf{H}$, which instead map model states to the observation space. The inverse Helmholtz decomposition, solving for streamfunction and velocity potential from the horizontal

velocities, is a part of the T-transform and is essential for computing statistics of the control variables. However, the T-transform is not required during the assimilation process itself. Instead, the forward transformation, which derives horizontal velocities from the streamfunction and velocity potential, is used as part of the U-transform, which is employed during the data assimilation procedure. While this forward transformation is straightforward to apply, performing the inverse is considerably more complex. In the following section, we describe the formulation of this inverse in NEMOVAR. The rest of this chapter is focused on the inverse transformation, we test the forward transformation in chapter 6.

5.2.2 The discrete transformation to streamfunction and velocity potential

We aim to solve the Helmholtz equations (5.6) discretised as follows,

$$\delta u_U = -\frac{1}{s_\varphi}\Delta_\varphi\delta\psi_U + \frac{1}{s_\lambda}\Delta_\lambda\chi \quad \text{and} \quad \delta v_U = \frac{1}{s_\lambda}\Delta_\lambda\psi + \frac{1}{s_\varphi}\Delta_\varphi\chi \quad (5.7)$$

to find the ageostrophic streamfunction and velocity potential increments, where we have the scale factors $s_\lambda = a \cos \varphi$, $s_\varphi = a$ and Δ_λ and Δ_φ are the discretisation of $\frac{\partial}{\partial \lambda}$ and $\frac{\partial}{\partial \varphi}$. For the shallow water model, we used the conjugate gradient (CG) method (see section 2.2.2) to minimise a cost function (4.14). In NEMOVAR we form a similar cost function. The NEMO model uses a non-uniform grid for its configurations, including the GYRE configuration. The use of a non-uniform grid requires weighting for each area element. This is implemented using $\mathbf{W}_V$, which is a diagonal matrix of area elements at the velocity points. We formulate the following cost function,

$$\mathcal{J}(\Phi) = \frac{1}{2}(\delta \mathbf{v}_{h_U} - \mathbf{A}\Phi)^T \mathbf{W}_V (\delta \mathbf{v}_{h_U} - \mathbf{A}\Phi) \quad (5.8)$$

where

$$\Phi = \begin{pmatrix} \delta\psi_U \\ \delta\chi_U \end{pmatrix} \quad \text{and} \quad \delta \mathbf{v}_{h_U} = \begin{pmatrix} \delta u_U \\ \delta v_U \end{pmatrix}$$

and here

$$\mathbf{A} = \begin{pmatrix} -\frac{1}{s_\varphi}\Delta_\varphi & \frac{1}{s_\lambda}\Delta_\lambda \\ \frac{1}{s_\lambda}\Delta_\lambda & \frac{1}{s_\varphi}\Delta_\varphi \end{pmatrix}. \quad (5.9)$$

Setting the gradient of the cost function (5.8),

$$\nabla \mathcal{J}(\Phi) = -\mathbf{A}^T \mathbf{W}_V (\delta \mathbf{v}_{h_U} - \mathbf{A} \Phi), \quad (5.10)$$

to zero gives the following expression for Φ

$$\mathbf{A}^T \mathbf{W}_V \mathbf{A} \Phi = \mathbf{A}^T \mathbf{W}_V \delta \mathbf{v}_{h_U} \quad (5.11)$$

which we solve using the CG method. Similar to the shallow water model in section 4.1, the GYRE configuration is a bounded simply connected domain and the horizontal velocity vector is non-zero on the boundary. Therefore, the Helmholtz decomposition and the solution to (5.11) is not unique.

Non-uniqueness of Helmholtz Theorem

In the shallow water model, we apply Tikhonov's regularisation to address the non-uniqueness of the Helmholtz decomposition (see section 4.2.3). For NEMOVAR, we also use Tikhonov's regularisation. The inclusion of regularisation updates (5.11) to give

$$(\mathbf{A}^T \mathbf{W}_V \mathbf{A} + \mu \mathbf{W}_T) \Phi = \mathbf{A}^T \mathbf{W}_V \delta \mathbf{v}_U \quad (5.12)$$

where $\mathbf{W}_T$ is a diagonal matrix of area elements at T-points, and μ is a regularisation parameter defined as

$$\mu = \hat{\mu} \cdot \max(\text{diag}(\mathbf{G})) \quad (5.13)$$

where $\hat{\mu}$ is a non-dimensional parameter and

$$\mathbf{G} = \mathbf{W}_T^{-1} \mathbf{A}^T \mathbf{W}_V \mathbf{A}. \quad (5.14)$$

This is chosen as such to tune the relative magnitude of two terms with a non-dimensional coefficient. This particular choice of parameter removes the area weighting, thereby ensuring equal weighting across all grid cells. As a result, only a single regularisation parameter is required for the entire domain. The strength of the regularisation is controlled by $\hat{\mu}$.

Once again we implement the method proposed by Li et al. (2006), whereby streamfunction is located at the centre of the grid cell (see figure 3.2b). As discussed in section 4.2.4, there are two approaches to calculating the streamfunction derivatives. In

NEMOVAR we implement approach B (described in section 4.2.4), whereby we average the streamfunction derivatives to the corners of the grid cells before calculating the derivatives. A note about implementing this approach is presented in appendix F. We now discuss implementing this transformation to the alternative control variables.

5.3 Performing the discrete transformation

In this section we implement and test the discrete transformation from the horizontal velocities to their non-divergent and irrotational parts, in the context of NEMOVAR. To accomplish this, we first run the GYRE configuration to generate model increment fields. From these fields, we perform the inverse of the Helmholtz transformation to compute the non-divergent and irrotational components of the velocity increments. We then test this transformation by reconstructing the velocities and analysing the reconstruction errors.

5.3.1 Method

We run the GYRE configuration presented in section 5.1.2. The GYRE configuration is initialised from rest; therefore, to obtain model fields with significant and realistic ocean structures, the model must be spun up for some time. Since running high-resolution models is computationally expensive, we begin with a low-resolution setup and gradually increase the resolution. We generate the model output by performing the following steps, also illustrated in figure 5.4.

1. Run the GYRE configuration at a 1° resolution (GYRE1) for 1000 years, starting from rest.
2. Run the GYRE configuration at a $1/12^\circ$ resolution (GYRE12) for 10 years starting from an interpolated file at the end of the 1000 year GYRE1 run.
3. Run the GYRE configuration at a $1/36^\circ$ resolution (GYRE36) for 1 month starting from an interpolated file at the end of the 10 year GYRE12 run.

The highest resolution run allows small scale structures to form. However, we interpolate the file at the end of the 1 month GYRE36 run back to the GYRE12

configuration to be consistent with chapter 6 where we perform data assimilation experiments using the GYRE12 configuration.




Figure 5.4: Process to generate the one day increment from the GYRE configuration used in the transformation experiments in this chapter.

We run the GYRE configuration for one day at the $1/12^\circ$ resolution and take the difference between this day and the previous day. This generates a 1 day difference in the model variables which we use as a proxy for an assimilation increment for demonstration purposes here. The horizontal velocity increments at the surface that are produced using this procedure are shown in figure 5.5. In the following sections we decompose the velocity increments into their non-divergent and irrotational components, rather than the ageostrophic components, as done in chapter 4. We then reconstruct the velocities to investigate the errors associated with the transformation.




Figure 5.5: Surface horizontal velocity increments from running the GYRE configuration for 1010 years and 1 month. The zonal velocity increment is on the left and the meridional velocity increment on the right (ms^{-1}).

To transform the horizontal velocity increments to streamfunction and velocity potential increments we solve (5.12) using the CG method by performing a 2D minimisation at each vertical level. For the CG method, the maximum number of iterations is set to 10,000. These inverse experiments we perform are inexpensive to execute given the scale of the GYRE configuration domain. Setting a high maximum number of iterations allows us to observe convergence clearly, though in practice, convergence was achieved well before reaching this maximum. It is also worth noting that no preconditioning was applied in this setup; with an appropriate preconditioner, the number of required iterations would likely decrease significantly. Furthermore, we define the convergence criterion based on the residual reduction relative to its initial value, which must be less than a specified tolerance.

In order to perform the inverse transformation we must first answer the following questions

- What value to use as the regularisation parameter, $\hat{\mu}$?
- What value to use as the convergence (accuracy) tolerance?

5.3.2 Transformation experiments

In this section, we conduct experiments to choose the value of the convergence tolerance and regularisation parameter for our GYRE configuration set-up and show the sensitivity of the transformation to a variety of settings. We experiment with four values for both the convergence tolerance and regularisation parameter. For each experiment we transform from the horizontal velocity increments, generated using the method described in section 5.3.1, to streamfunction and velocity potential. We then reconstruct the velocities from these fields using the Helmholtz equations (5.7). For each experiment, we calculate the root mean square (RMS) of the zonal and meridional velocity error, defined by (2.20). The experiment results are shown in table 5.1. The entries include the number of iterations the CG method took to converge and the RMSE of the surface velocities. The table confirms some expected results: as the regularisation parameter and convergence tolerance decrease, the number of iterations required for convergence increases, but the RMSE of the velocities becomes smaller.

Experiment number	Convergence tolerance	Regularisation parameter ($\hat{\mu}$)	RMSE zonal velocity	RMSE meridional velocity	Iterations
1	10^{-3}	10^{-3}	0.0048567	0.0043384	111
2	10^{-3}	10^{-4}	0.0011231	0.000991	225
3	10^{-3}	10^{-5}	0.0004596	0.0003853	354
4	10^{-3}	10^{-6}	0.0004176	0.0003455	377
5	10^{-4}	10^{-3}	0.004849	0.0043435	161
6	10^{-4}	10^{-4}	0.0011162	0.0009634	371
7	10^{-4}	10^{-5}	0.0004094	0.0002951	827
8	10^{-4}	10^{-6}	0.0002719	0.0002023	1519
9	10^{-5}	10^{-3}	0.0048491	0.0043423	212
10	10^{-5}	10^{-4}	0.0011158	0.0009631	530
11	10^{-5}	10^{-5}	0.0004017	0.0002997	1255
12	10^{-5}	10^{-6}	0.0002615	0.0001945	3246
13	10^{-6}	10^{-3}	0.0048492	0.0043423	262
14	10^{-6}	10^{-4}	0.0011159	0.000963	688
15	10^{-6}	10^{-5}	0.0004018	0.0002994	1751
16	10^{-6}	10^{-6}	0.0002613	0.0001943	4754

Table 5.1: Convergence and regularisation parameter experiments for the GYRE configuration. The RMSE values are for the surface velocities and the number of iterations are for the convergence of the minimisation at the surface. The highlighted row is our chosen settings for the final experiment.

To visually examine and understand the results from table 5.1, we plot the RMSE values for the three smallest regularisation parameters at the surface in figure 5.6. We exclude experiments 1, 5, 9 and 13, with $\hat{\mu} = 10^{-3}$, as their RMSE values are significantly larger than the smaller regularisation parameters. When comparing the RMSE values of the velocities to the magnitude of the original velocity increments, a value of $\hat{\mu} = 10^{-3}$ results in approximately a 5% error. However, reducing $\hat{\mu}$ to $\leq 10^{-4}$ decreases the error to 1% or even less. It is evident from figure 5.6 that tightening the convergence tolerance has progressively less impact, with further reductions yielding diminishing returns. In particular, the difference between 10^{-6} and 10^{-5} is almost negligible, suggesting that




Figure 5.6: RMSE of the reconstructed horizontal velocity at the surface for various regularisation parameters and convergence tolerance values. The left is the zonal velocity and the right the meridional velocity. The blue line (convergence tolerance = 10^{-6}) is masked by the orange line.

beyond a certain point, a stricter tolerance offers little to no practical benefit. Figure 5.7 shows the RMSE values averaged across all vertical levels. We see that the RMSE at the surface is representative of the mean values across the depths. We conclude from figures 5.6 and 5.7, that reducing the convergence tolerance below 10^{-5} is unnecessary.

Consider the results from table 5.1 where the tolerance is equal to 10^{-5} . Figure 5.8a shows the norm of the residual in (5.12) (the left hand side minus the right hand side) at the surface, as the iterations proceed. The rate as which the residual falls is inversely related to the computational time taken to converge to the solution. We see that the number of iterations before convergence increases drastically the smaller the regularisation parameter. To illustrate the impact of reducing this parameter to a more extreme value and demonstrate the negative effect on convergence we add a new experiment with $\hat{\mu} = 10^{-7}$, which does not feature in table 5.1. Figure 5.8b shows the number of iterations for each regularisation parameter value at each vertical level. We see that the number of iterations to converge is reasonably consistent through the vertical levels, except for this additional regularisation parameter of 10^{-7} , which sees a sharp rise in iterations with depth. This suggests that vectorising the computation — computing the transformations for all levels simultaneously — could be a viable strategy to reduce




Figure 5.7: Average RMSE of the reconstructed horizontal velocity across all depths for various regularisation parameters and convergence tolerance values. The left is the zonal velocity and the right the meridional velocity. The blue line (convergence tolerance = 10^{-6}) is masked by the orange line.




Figure 5.8: The norm of the residual at the surface and the number of iterations to convergence across all depths with an accuracy tolerance of 10^{-5} and various values of regularisation parameter.

communication steps and improve efficiency.

From these results, we choose to set the accuracy tolerance and regularisation parameter, $\hat{\mu}$, to both be 10^{-5} . This convergence value is chosen as such since smaller values beyond this threshold give minimal improvement and result in significantly more iterations to converge. We conclude that $\hat{\mu} = 10^{-5}$ is also an appropriate value for our particular system from the above analysis. Future work would look to use a more precise

method of deciding this regularisation parameter. There are multiple methods discussed in literature to determine an optimal value of Tikhonov’s regularisation parameter. A particular method, called the L-curve, as discussed in section 2.2.1, is concluded as the most robust method according to Hansen (1997).

Figures 5.9a and 5.9b show the streamfunction and velocity potential increments obtained after performing the inverse Helmholtz decomposition with these parameters (experiment 11 in table 5.1). The errors in the velocities after reconstruction are given in figures 5.9c and 5.9d. We see that the error is relatively small compared to the magnitude of the velocity increments, approximately 1% of increment fields (seen in figure 5.5). In figures 5.9c and 5.9d a checkerboard pattern is apparent. This is caused by the averaging involved in calculating streamfunction derivatives, as discussed in section 4.3.2. We discuss this issue further in section 5.4.

5.3.3 Summary

In this section we performed the transformation from the velocity increments to their irrotational and non-divergent parts, velocity potential and streamfunction. We performed experiments to test the sensitivity of the transformation to the regularisation parameter and convergence tolerance. Our results showed that reducing the tolerance and regularisation parameter below a certain threshold has little impact on solution sensitivity but causes a substantial increase in the number of iterations required for convergence. We propose the following parameter settings when performing the transformation from velocity increments to their non-divergent and irrotational parts, using the GYRE configuration:

Accuracy tolerance	10^{-5}
$\hat{\mu}$	10^{-5}

In the next section we perform the transformation using these settings and discuss the checkerboard pattern. As with this section, we do not use the full balance operator discussed in section 5.2.1. Notably, the checkerboard pattern would still arise when the transformation is applied in the full balance operator, to the ageostrophic velocity increments.




Figure 5.9: Top row: The resulting surface streamfunction and velocity potential increments from performing the inverse Helmholtz decomposition with accuracy tolerance of 10^{-5} and $\hat{\mu} = 10^{-5}$ (experiment 11 in table 5.1). Bottom row: The errors obtained when reconstructing the horizontal velocities.

5.4 Checkerboard pattern

In this section, we examine the checkerboard pattern that emerges when the horizontal velocities are decomposed into their non-divergent and irrotational components. We then assess the sensitivity of this transformation to a proposed solution involving the application of a Shapiro filter, and discuss additional issues that arise from this approach.

We introduced the checkerboard pattern in section 4.3.2, when performing the inverse of the Helmholtz decomposition to find streamfunction and velocity potential, using the SWEs. As discussed in section 4.3.2, the checkerboard pattern introduces noise into the

control variable fields, which can distort calculated statistics. This leads to an inaccurate background error covariance matrix and poses challenges when estimating length scales.

In the previous section, we performed the transformation from the horizontal velocities to streamfunction and velocity potential using the GYRE configuration of NEMO. From an initial visual inspection of the streamfunction field in figure 5.9a, the checkerboard pattern is not obvious; however it is clearly visible in figures 5.9c and 5.9d, the velocity reconstruction errors at the surface. The checkerboard computational mode differs on a non-uniform grid from that produced using the SWEs on a uniform grid. Recall, the checkerboard pattern occurs when information from adjacent streamfunction points is lost when calculating the streamfunction derivatives. On a uniform grid, adjacent streamfunction values are completely eliminated. However, on a non-uniform grid, such as the one used in NEMO, information from nearby points is still removed, but not entirely.




Figure 5.10: The horizontal lines we perform the Fast-Fourier transform on in order to detect the checkerboard pattern.

To demonstrate the presence of the checkerboard pattern, we use a Fast-Fourier

transform (FFT) (see section 2.4). We take multiple horizontal slices of the streamfunction field, shown in figure 5.10. For each line we calculate the discrete Fourier transform (DFT) sample frequencies using the FFT, as seen in the left column of figure 5.11. In appendix G, we discuss the impact of applying the FFT to a non-periodic wave. This non-periodicity introduces distortions in the high-frequency components. To address this issue, we adopt a mirroring approach: the original lines are extended by appending their mirror image, effectively doubling their length. As a result, the FFT produces wavenumbers that are twice as dense as those obtained from the original (non-mirrored) data. However, only the even wavenumbers correspond directly to the wavenumbers of the standard FFT without mirroring, and it is these values that we present in this and the following section.

We also apply this Fourier analysis to the meridional velocity reconstruction error field, as seen in the right column of figure 5.11. In these figures we show the spectra calculated from the Fourier analysis at the surface, given by the red line. The analysis is repeated at all vertical levels and the blue line shows the averaged amplitude across the levels. The grey shaded region provides the range of values. The amplitude shows a distinct peak at the lowest frequencies (low wavenumbers) as expected, indicating a wave the size of the domain across all sample lines. Line 4, which intersects an area of intensified dynamical activity (e.g., oceanic eddies), displays elevated spectral energy at these low frequencies — an expected result given the enhanced variability and multiscale structures present. For all sampled lines, we observe a small but distinct high-frequency (high wavenumber) peak in streamfunction, corresponding to a wavelength of two grid lengths — the shortest wavelength resolvable on the grid. This is characteristic of a checkerboard pattern, which is even more prominently visible in the meridional wind reconstruction error.

Having confirmed the existence of the checkerboard pattern, we adopt the same approach as in section 4.3.4, applying the Shapiro filter (see section 2.3) to attempt to remove the associated noise.

5.4.1 The Shapiro filter

We apply the second order Shapiro filter to the streamfunction output produced by experiment 11 at each depth; the surface field before filtering is shown in Figure 5.9a. The filtering is performed as a post-processing step, with the justification detailed in




Figure 5.11: The spectra generated from applying the FFT to the lines shown in 5.10, the top row for line 1 and the bottom row for line 4. Streamfunction (left) and meridional velocity reconstruction error (right). No filter is applied to the streamfunction field. The red line is the spectrum at the surface, blue is the amplitude averaged across the 31 vertical levels and the grey shaded region is the range of values across the depths.

section 4.3.4. We begin by applying a single iteration of the Shapiro filter and assessing its effect on the checkerboard pattern through the Fourier transform. First, we transform the velocity increments to streamfunction and velocity potential, as per experiment 11. Next, we apply the Shapiro filter to the streamfunction field. Finally, we reconstruct the velocities using the filtered streamfunction field along with the original velocity potential field. The FFT is then applied to the horizontal lines shown in figure 5.10 for both streamfunction and the meridional velocity reconstruction error using the process described previously.

Figure 5.12 shows that the pronounced peaks at high frequencies, which we saw in figure 5.11, have been removed by applying one iteration of the Shapiro filter to streamfunction. This is the case for both the streamfunction and the meridional velocity error fields. However, the structure of the meridional velocity error signal has become more uniform, suggesting that the Shapiro filter has eliminated dominant spatial patterns even at large scales. As a result, the remaining signal appears increasingly noise-dominated. Since this is an error field, the presence of a noisy pattern is considered acceptable as we expect no particular signal. We also observe slight alterations in the low frequencies, which are discussed in more detail later in the section.

Figure 5.12 suggests that a single iteration of the Shapiro filter is sufficient to remove the checkerboard pattern. To verify this, we repeat the previous experiment, increasing the number of filter iterations to five. As shown in figure 5.13, the streamfunction spectrum remains essentially unchanged, indicating that further iterations are unnecessary as the FFT shows no additional removal of the checkerboard pattern. The meridional velocity error spectrum again remains fairly uniform. Therefore, we conclude that only one iteration of the Shapiro filter is necessary when performing this transformation using the GYRE configuration.

In figure 5.14a we show the streamfunction field after applying the single iteration of the Shapiro filter and figures 5.14b and 5.14c are the errors in the reconstructed velocities, all at the surface. The reconstruction errors appear significantly larger than those observed without the filter (see figures 5.9c and 5.9d). In figure 5.15 we have a vertical cross-section of streamfunction and the velocity errors along line 4 in figure 5.10. From these we clearly see that the magnitudes are larger at the surface and in the areas of high eddy activity.

Similar to the experiments detailed in table 5.1 we calculate the RMSE of the




Figure 5.12: As figure 5.11 with the addition of one iteration of the Shapiro filter applied to the streamfunction field.




Figure 5.13: As figure 5.11 with the addition of five iterations of the Shapiro filter applied to the streamfunction field.




Figure 5.14: (a) The resulting surface streamfunction increment from performing the inverse Helmholtz decomposition with accuracy tolerance of 10^{-5} , $\hat{\mu} = 10^{-5}$ and applying one iteration of the Shapiro filter as a post-process. (b) and (c) The reconstructed surface velocity errors.

Experiment number	Convergence tolerance	Regularisation parameter ($\hat{\mu}$)	RMSE zonal velocity	RMSE meridional velocity	Shapiro filter iterations
17	10^{-5}	10^{-5}	0.0102558	0.0094623	1
18	10^{-5}	10^{-5}	0.0227913	0.0211791	5

Table 5.2: Shapiro filter experiments for the GYRE configuration. The RMSE values are for the surface velocities and the number of iterations are for Shapiro filter applied to the streamfunction increment.

reconstructed velocities, as shown in table 5.2. The RMSE of the zonal and meridional reconstructed velocities are 0.0103 and $0.0094ms^{-1}$, respectively for one iteration of the Shapiro filter. Both these values are larger than all the RMSE values given in table 5.1.




Figure 5.15: Vertical profiles taken along line 4 in figure 5.10. (a) The resulting streamfunction increment from performing the inverse Helmholtz decomposition with accuracy tolerance of 10^{-5} , $\hat{\mu} = 10^{-5}$ and applying 1 iteration of the Shapiro filter as a post-process. (b) and (c) The reconstructed surface velocity errors.

Comparing these RMSE to the magnitude of the original velocity increments (seen in figure 5.5) gives approximately a 10% error. Figure 5.16 shows the RMSE at all the depths for the experiments in table 5.1 (excluding $\hat{\mu} = 10^{-3}$) and experiment 17 from table 5.2. The majority of the lines are hidden beneath line 15, and none beneath line 17. We observe that at all depths, applying the Shapiro filter to streamfunction results in higher RMSEs in the reconstructed velocities compared to all experiments with varying tolerance and regularisation values — rendering the process of tuning these parameters almost redundant.

A larger reconstruction error compared with the original transformation without the




Figure 5.16: RMSE of the reconstructed velocities at all the depths for the experiments in table 5.1 (excluding $\hat{\mu} = 10^{-3}$) and experiment 17 from table 5.2.

filter is expected. The application of the Shapiro filter is supposed to remove small-scale, high-frequency features from the streamfunction field. By eliminating this information, we reduce the accuracy of streamfunction in representing the true field. Consequently, the reconstruction is based on a distorted version of the streamfunction, leading to larger errors. We can see in figures 5.14b and 5.14c that these errors are particularly large at the boundaries. The Shapiro filter, like most spatial filters, works by smoothing the signal based on neighbouring values. Near the boundaries, the filter has fewer neighbouring points to consider, which can result in less accurate filtering near the edges. This effect can lead to larger discrepancies in the reconstructed fields at the boundaries. However, while some increase in error is expected with the use of the Shapiro filter, the fact that it has increased the reconstruction error by more than an order of magnitude and the error is far worse than all the sensitivity experiments is an unacceptably poor outcome. We now investigate whether this specific filter choice is responsible for the large errors.

Filter test

As discussed, applying the Shapiro filter to streamfunction dramatically increases errors in the horizontal velocity reconstruction. To determine whether these errors are unacceptably large due to the suppression of low-frequency components — rather than

the high frequencies, which we expect to cause some error but not to this extent — we perform the following test using the FFT:

1. Apply the 2D FFT to the streamfunction field produced by experiment 11.
2. Set the highest frequencies of the FFT spectra to zero (wavenumbers 170 : 176 in the x -direction and 110 : 116 in the y -direction).
3. Reconstruct the streamfunction field by reversing the 2D FFT.
4. Apply Helmholtz equations (5.7) to reconstruct the velocities and calculate the reconstruction errors.




Figure 5.17: 2D FFT spectrums of streamfunction at the surface. Left: The 2D FFT of the streamfunction field shown in figure 5.9a produced by experiment 11. Middle: The 2D FFT of the streamfunction field with the Shapiro filter applied, experiment 17. Right: The 2D FFT of the streamfunction field produced by experiment 11 with the highest frequencies set to zero.

Here we do not use the mirroring process discussed for the 1D FFT in section 5.4 but simply apply the 2D FFT to the whole domain. The 2D FFT spectrums are shown in figure 5.17 and the reconstruction errors can be seen in figure 5.18. We observe that the reconstruction errors are half the size of those produced using the Shapiro filter, as shown in Figures 5.14b and 5.14c, except near the boundaries. This suggests that while the second order Shapiro filter is a low-pass filter that should primarily affect the smallest wavelengths, as discussed in section 2.3, it appears to negatively impact the larger wavelengths, leading to significant reconstruction errors. However, the high-frequency filter we have applied in this test has removed some of the checkerboard pattern but not in its entirety, as we see this pattern near the north-west and south-east boundaries

in figure 5.18a and towards the eastern and western boundaries in figure 5.18b. This test demonstrates that the velocity reconstruction error is highly sensitive to the choice of filter. It highlights the need for a filter that more effectively targets high-frequency components and improves boundary treatment — potentially through the use of a higher-order Shapiro filter.




Figure 5.18: Surface velocity reconstruction errors using a FFT. The velocities are reconstructed from the streamfunction field with the highest wavenumbers set to zero. The left is the zonal velocity error and the right is the meridional velocity error (ms^{-1}).

5.5 Summary and conclusions

In this chapter we have described how the NEMOVAR balance operator can be updated to use ageostrophic streamfunction and velocity potential as the velocity control variables instead of the ageostrophic velocities, which are currently used. To demonstrate the additional transformation added to the balance operator, we used the GYRE configuration of NEMO.

Firstly, we introduced the GYRE configuration and the discretisation used in the NEMO model. Next, we detailed how to revise the NEMOVAR balance operator by incorporating an additional transformation that transforms the ageostrophic velocity increments into ageostrophic streamfunction and velocity potential increments — the inverse of the Helmholtz decomposition. We then formulated this inverse transformation

using NEMOVAR notation, accounting for NEMO’s spherical coordinates on a non-uniform grid. This naturally led to a discussion of the non-uniqueness of this decomposition and using Tikhonov’s regularisation, as previously applied in the shallow water framework outlined in section 4.2.3.

Furthermore, we performed the transformation of velocity increments into their non-divergent and irrotational components. Importantly, rather than applying the full balance operator discussed in section 5.2.1, we focused solely on decomposing the velocity increments into their irrotational and non-divergent parts. This section aimed to illustrate the transformation and assess its sensitivity to specific choices, as these are independent of the transformation’s input. This approach can later be extended to the ageostrophic components when incorporating the full balance operator.

To implement this transformation, we must address the non-uniqueness of the Helmholtz decomposition and establish appropriate convergence tolerance. We conducted multiple experiments, initially testing different values for both the convergence tolerance and the regularisation parameter. From these experiments, we determined that a tolerance of 10^{-5} is appropriate. Below this threshold, the magnitude of reconstructed velocity errors became significantly less sensitive, whereas the number of iterations required for convergence increased substantially. Finally, we established that $\mu = 10^{-5} \cdot \max(\text{diag}(\mathbf{W}_T^{-1} \mathbf{A}^T \mathbf{W}_V \mathbf{A}))$ is a suitable regularisation parameter for this study.

After confirming the parameter settings, we analysed the streamfunction output from this transformation. Visually, it is unclear whether the checkerboard pattern observed in the shallow water domain is present here. However, it is clearly visible in the velocity reconstruction errors. To verify this numerically, we applied a Fourier transform, which revealed the presence of grid-scale waves — indicative of the checkerboard pattern in both the streamfunction and in the reconstructed velocities. To mitigate this issue, we applied the second order Shapiro filter to the streamfunction used to reconstruct the velocities and found that a single iteration is sufficient to eliminate the grid-scale noise. However, this comes at the cost of significantly increasing the error in the reconstructed velocities. Diagnostics revealed that a filter specifically targeting the highest frequencies would be more effective.

In conclusion, we have extended the work presented in chapter 4 by demonstrating how the newly proposed velocity control variables can be implemented in the balance

operator of NEMOVAR. These variables have not previously been used with this operator and represent a novel development within NEMOVAR. We showed through detailed tests the sensitivity of the transformation (to these new variables) to the associated parameters and identified the emergence of checkerboard patterns - results which should be addressed when applying these control variables operationally.

In chapters 4 and 5 we have investigated the inverse transformation - decomposing the velocities into their non-divergent and irrotational components. This transformation is part of the T-transform and is necessary for computing statistics of the control variables. However, the T-transform is not required during the assimilation process. In the next chapter, we investigate the U-transform, by conducting assimilation experiments using the proposed control variables in NEMOVAR.

Chapter 6

Assimilation experiments in NEMOVAR

In this final chapter we perform assimilation experiments in NEMO and NEMOVAR using the proposed velocity control variables, ageostrophic streamfunction and velocity potential. In chapter 5 we discussed the process of transforming to the newly proposed control variables from the velocity increments as part of the T-transform in NEMOVAR (this is the inverse of the Helmholtz transformation). We also presented the challenges faced when performing this transformation, and the measures that must be taken to combat these. Here, we test the forward transformation. We run assimilation experiments to determine whether these new control variables have the potential to improve the analysis produced compared to the current control variables, the ageostrophic velocities. To do so we use the idealised GYRE configuration of NEMO and the NEMOVAR assimilation system.

Firstly, we discuss in detail the aims and design of these idealised twin assimilation experiments. We then describe the model behaviour and the generation of a high-resolution model run that we consider to be the ‘true’ state of the ocean. Next, we discuss the observations and background used in the data assimilation experiments. Initially, we conduct experiments with single observations to analyse the increments produced by our new control variables. We then present the results of the assimilation experiments using the full set of observations. We conclude that the proposed control variables demonstrate the ability to improve the assimilation of velocity observations. More fine tuning of error

covariance parameters should lead to further improvements in performance.

6.1 Experiment overview

In this chapter, we perform cycling twin experiments to assess the impact of assimilating velocity observations using the new NEMOVAR balance operator, which incorporates the proposed velocity control variables (described in section 5.2.1). These are observing system simulation experiments (OSSEs; Masutani et al. (2010)), which are designed to evaluate the potential impact of an observing system using synthetic data. Specifically, we test the assimilation of velocity observations using the proposed control variables. We generate a high-resolution nature run that we consider to be our ‘truth’. Outputs from this nature run are used to generate idealised observations with artificial noise included, which are then used in multiple assimilation experiments at lower resolution. The ability of the assimilation to recover the ‘true’ state of the ocean is assessed by comparing the outputs of the assimilation runs to the nature run. We also generate a model free-run which is at a lower resolution to the nature run, matching that used in the assimilation experiments. The purpose of this lower-resolution free-run is to demonstrate the impact of not using data assimilation and observations. We use the GYRE configuration of NEMO as our model, which is described in section 5.1.

6.1.1 Spin-up procedure

The nature run in these twin experiments is a high-resolution free-running model which we consider to be the ‘true’ state of the ocean for the purposes of these idealised experiments. The free-run is a low-resolution free-running model. As illustrated in figure 6.1, we perform the following steps to spin-up the GYRE configuration:

1. Run the GYRE configuration at a 1° resolution (GYRE1) for 1000 years, starting from rest.
2. Run the GYRE configuration at a $1/12^\circ$ resolution (GYRE12) for 10 years starting from an interpolated file at the end of the 1000 year GYRE1 run.
3. Run the GYRE configuration at a $1/12^\circ$ resolution (GYRE12_DA) for a further year starting from the end of the 10 year GYRE12 run.

4. Run the GYRE configuration at a $1/36^\circ$ resolution (GYRE36) for 1 year starting from an interpolated file at the end of the 10 year GYRE12 run.

Using this spin up we generate the nature run, free-run and initial conditions for the assimilation experiments:

1. Run the GYRE configuration at a $1/36^\circ$ resolution (GYRE36_NR) for 1 year starting the end of the 1 year GYRE36 run. This is taken as the **nature run**. We use the nature-run to generate the synthetic observations used in this these experiments, see section 6.3.1.
2. Run the GYRE configuration at a $1/12^\circ$ resolution (GYRE12_FR) for 1 year starting the end of the 1 year GYRE12_DA run. This is taken as the **free-run**.
3. The assimilation runs start from the end of the GYRE12_DA spin up.




Figure 6.1: Process of spinning up the GYRE configuration for the nature run, assimilation experiments and free-run of the model, including the resolution of the model.

6.1.2 Experiment details

In table 6.1 we list the assimilation experiments to be performed in this chapter. The details of each experiment will be discussed in the following sections.

6.1.3 Experiment aims

These experiments aim to address the third research question of this thesis, described in chapter 1. This question examines the implications of using the proposed velocity

Experiment	Velocity control variables	Observations	Background error standard deviations	Duration
Free-run	-	-	-	1 year
Control 1	-	Standard	$\sigma_u, \sigma_v = 0.2ms^{-1}$	1 year
Control 2	$\delta u_U, \delta v_U$	Standard and velocity	$\sigma_u, \sigma_v = 0.2ms^{-1}$	1 year
New CVs 1	$\delta\psi_U, \delta\chi_U$	Standard and velocity	$\sigma_\chi, \sigma_\psi = 10^3m^2s^{-1}$	1 year
New CVs 2	$\delta\psi_U, \delta\chi_U$	Standard and velocity	$\sigma_\chi = 0m^2s^{-1},$ $\sigma_\psi = 10^3m^2s^{-1}$	1 year
New CVs 3 - sensitivity	$\delta\psi_U, \delta\chi_U$	Standard and velocity	$\sigma_\chi = 0m^2s^{-1},$ $\sigma_\psi = 10^4m^2s^{-1}$	1 month
New CVs 4 - sensitivity	$\delta\psi_U, \delta\chi_U$	Standard and velocity	$\sigma_\chi = 0m^2s^{-1},$ $\sigma_\psi = 10^2m^2s^{-1}$	1 month
New CVs 5 - sensitivity	$\delta\psi_U, \delta\chi_U$	Standard and velocity	$\sigma_\chi = 0m^2s^{-1},$ $\sigma_\psi = 10^1m^2s^{-1}$	1 month

Table 6.1: Description of the assimilation experiment settings.

control variables in NEMOVAR and their response to assimilating future ocean current measurements. The main aims of these experiments are:

1. To demonstrate the use of the proposed control variables in NEMOVAR.
2. To show that the proposed control variables are able to improve the assimilation of velocity data in NEMOVAR.
3. To evaluate the sensitivity of the assimilation to choices in the prescribed error covariance statistics of the proposed control variables.
4. To assess the impact of the assimilation of surface velocity data on the model's vertical velocities using the current and proposed control variables.

Previous studies, Li et al. (2008*a,b*) and Li et al. (2015), also perform assimilation experiments in a regional domain using these proposed control variables, ageostrophic

streamfunction and velocity potential. However our experiments differ from these in multiple ways,

- We are using the NEMOVAR balance operator which uses the diffusion operator to model the spatial correlations in the background errors rather than the Kronecker product, which is used by Li et al. (2008a).
- Our experiments are focused on the effect of the proposed control variables and comparing these to the previous control variables used in NEMOVAR - the ageostrophic velocities. Previous studies have used these control variables but not investigated the impact compared to other velocity control variables.
- We use idealised satellite velocity observations in our experiments. Li et al. (2008b) conduct both single-observation and multi-observation assimilation experiments. In the latter they do not assimilate any velocity observations; the velocity analysis increment arises solely from the imposed geostrophic balance. Li et al. (2015) perform OSSEs that include velocity data, but the observations are obtained from HF radar measurements, which are limited to coastal regions, whereas we use satellite-derived velocity observations that provide coverage across the open ocean.

6.2 Model behaviour

In this section we compare the fields generated by the nature run to a free-run of the model. By comparing these two fields in this section we illustrate some of the differences expected in the assimilation experiments due to the lower resolution.

Figure 6.2 presents the full fields at the surface — SSH, temperature, zonal, and meridional velocity — on the first day after the spin-up. For both the nature run and free-run, we observe a western boundary current, indicated by higher velocities. Additionally, eddies of different scales are clearly seen in the SSH fields. Notably, the nature run exhibits more pronounced small-scale features across all fields compared to the free-run which we expect due to the higher resolution. This is observed similarly in the salinity field (not shown).

We also compare the mean and variability (represented by the standard deviation) of the fields over the course of the year, as shown in figures 6.3 and 6.4, respectively.




Figure 6.2: GYRE configuration model fields on the first day of the runs after the spin up at the surface. The left column is the nature run, interpolated to $1/12^\circ$ resolution and the right column is the free-run. The top row shows elevation, second row is temperature, third row is zonal velocity and the final row is meridional velocity.




Figure 6.3: As figure 6.2 but the mean fields.

The mean values for all variables exhibit similar patterns to the first day of the run. The nature run shows greater variability in the SSH field compared to the free-run, whereas the temperature field exhibits the opposite. This may be explained by the higher resolution of the eddies, which enhance the transport of temperature. When time-averaged, this results in increased variability in the temperature field, as the eddies redistribute heat more effectively. The nature run, being higher resolution, may better constrain these features, resulting in a smoother temporal evolution. In contrast, the free-run may generate spurious variability due to unresolved sub-grid processes or weaker constraints. Salinity exhibits a similar pattern to temperature (not shown). The velocities also have a higher variability in the nature run compared to the free-run. Overall, variability shown in figure 6.4 is highest near the western boundary current.

In the next section we describe the data assimilation set-up for the assimilation experiments.

6.3 Data assimilation set-up

In this section we detail the data assimilation set-up for the assimilation experiments. These settings are relevant to both the single observation experiments in section 6.4 and the assimilation experiments in sections 6.5-6.8. We run the assimilation experiments using NEMO and NEMOVAR. NEMOVAR uses 3D FGAT as the operational assimilation method (Mogensen et al., 2009; Mogensen and Balmaseda, 2012; Waters et al., 2015), discussed in section 3.2.2. We run the assimilation experiments on a daily cycle. On each cycle a one-day run of NEMO reads in the observations and interpolates the model values to the observation locations at the nearest model time-step and outputs these innovations. NEMOVAR reads in these innovations together with a gridded model background file (which is used in the multivariate balance calculations) and error covariance information, and produces the analysis increments on the model grid for the model variables. These analysis increments are read into a second run of NEMO during which they are applied slowly using Incremental Analysis Updates (IAU). NEMO expects increments in temperature, salinity, sea-surface height and the horizontal velocities.




Figure 6.4: As figure 6.2 but the standard deviation fields.

6.3.1 Observations

Generating the observations

The observations used in this chapter are simulated from the nature run described in section 6.1.1. For the standard observations (excluding those of velocity), we use the times and day of year as well as locations of real observations used in the Met Office FOAM system for the year 2016. Similar to the experiments in Waters et al. (2024a). We assume that the observation errors are uncorrelated Gaussian noise, therefore we add Gaussian noise with spatially constant standard deviations to the interpolated nature run values. These standard deviations are given in table 6.2 and are also used in the specification of the observation error covariance matrix. Table 6.2 also indicates which observation types are standard, as referred to in table 6.1.

Standard observation types

We use the standard observations types for all the assimilation experiments, these are: in-situ temperature and salinity profiles; in-situ and satellite sea surface temperature (SST) observations and altimeter observations of sea level anomaly (SLA). The coverage of these observations on the first day of the assimilation experiments is shown in figure 6.5.

Velocity observations

Some of our experiments also use observations which could be obtained from a SKIM-like satellite mission (Ardhuin et al., 2019). This is velocity data that has been simulated using the open-source SKIMulator tool (Gaultier and Ubelmann, 2024). This simulates measurements of total surface current velocities (TSCV) across a wide satellite swath as show in figures 6.5e and 6.5f. We assimilate 2D TSCVs which give both the zonal and meridional component of the velocities along the SKIM swath. This swath moves across the domain over time. The SKIMulator tool provides several levels of output. While it can generate radial velocities in the satellite’s reference frame, in our case we use the next level of output, where these radial velocities have been mapped—via an optimal interpolation procedure — onto a grid within the satellite swath, from which zonal and meridional velocities are estimated. A recent study by the European Space Agency Assimilation of TSCV (ESA A-TSCV) project (Waters et al., 2024a; Mirouze




Figure 6.5: Coverage of the observations used in the assimilation experiments on the first day of the assimilation.

et al., 2024; Waters et al., 2024b) performed OSSEs to demonstrate the potential positive impact TSCV observations could have on ocean forecasting. The simulated TSCV data used here has the same sampling characteristics as was used in that project.

Observation	Error standard deviation	Standard
Temperature profile	$1^{\circ}C$	✓
Salinity profile	$0.18psu$	✓
SST	$0.3^{\circ}C$	✓
SLA	$0.05m$	✓
TSCV	$0.1ms^{-1}$	

Table 6.2: Observation error standard deviations.

Assimilating the observations

We have a daily assimilation cycle and use IAU to incorporate the observations. The IAU is the process of adding analysis increments into a model forecast in a gradual manner (Bloom et al., 1991, 1996). In NEMOVAR, for each day a model forecast is generated from the previous analysis state. This does not produce any data assimilation increments but instead generates the model equivalent of the observations. The daily analysis increments are then calculated and then gradually applied throughout the day using the IAU. Using the IAU helps to reduce the shock of adding the increments to the model fields (Waters et al., 2015) and reduce any adverse impact the increments may have on sensitive model fields (Bloom et al., 1996).

The observations used in each assimilation experiment are given in table 6.1. Having discussed the details of the observations used in the assimilation experiments we now briefly discuss the background error covariance settings.

6.3.2 Background error covariance settings

The background error covariance matrix and balance operator, K , in NEMOVAR are discussed in section 3.4.3 with the updated balance operator that includes the new velocity control variables described in section 5.2.1. The standard deviations and lengthscales that we specify at the surface of the GYRE configuration are given in the table 6.3. The vertical lengthscales are three times the vertical grid resolution, in table 6.3 we use the grid resolution near the surface.

The values for the error standard deviations for the model variables are chosen based on the magnitude of the full fields, representing roughly a 10% error. We see that streamfunction is of magnitude $10^3 - 10^4$ (from chapter 5) and choose the standard deviation to be 10^3 . This is also approximately a 10% error. We investigate the sensitivity of the assimilation to this value in section 6.8. The standard deviation of velocity potential is chosen to have an equivalent weighting. Future work would use a more objective method to measure these standard deviations.

Variable	Error standard deviation	Horizontal lengthscale	Vertical lengthscale
η_U	$0.01m$	4°	-
T	$0.5^\circ C$	$35km$	$30m$
S_U	$0.25psu$	$35km$	$30m$
u_U	$0.2ms^{-1}$	$35km$	$30m$
v_U	$0.2ms^{-1}$	$35km$	$30m$
ψ_U	$10^3 m^2 s^{-1}$	$35km$	$30m$
χ_U	$10^3 m^2 s^{-1}$	$35km$	$30m$

Table 6.3: Background error covariance settings at the surface.

6.4 Single observation experiments

Before presenting the full assimilation results, we first evaluate the impact of introducing the proposed control variables, ageostrophic streamfunction and ageostrophic velocity potential, through a set of idealised single observation experiments. These experiments act as a tool for assessing how the new variables influence the structure and spatial spread of the analysis increments, as governed by the background error covariance matrix.

The aims of these experiments are to:

- Demonstrate the implementation of the updated balance operator (5.5) in NEMOVAR (section 5.2.1);
- Determine whether the proposed control variables lead to physically sensible changes in the analysis;
- Illustrate the forward transformation from control to model space of the new CVT.

Single observation experiments are widely used to investigate the background error covariances of an assimilation system. The resulting analysis increments are proportional to the forecast error covariances between the observed quantity and all other variables and locations in the domain (Thépaut et al., 1996). They provide valuable insight into how observational information spreads via the background error covariance matrix (Xiao and Sun, 2007).

To this end, we use two types of idealised observations: a single idealised TSCV observation with zonal and meridional components of the innovation equal to 0.1ms^{-1} ; and an altimeter observation of SLA with innovation of -0.1m . The SLA observation illustrates how the velocities are currently updated in NEMOVAR through the geostrophic balance relation. The TSCV observation reveals how directly updating the unbalanced component alters the velocity increments.

Each observation is assimilated independently on the first day of the assimilation run, following the GYRE12_DA spin-up (section 6.1.1, figure 6.1). The observation and background error covariance settings are described in section 6.3. The experiment with the altimeter observation uses the current NEMOVAR balance operator (3.18). We run two experiments with the TSCV observation, the first (Control) using the current balance operator, with the ageostrophic velocities as control variables. The second (New CVs) using the updated balance operator (5.5), which uses the proposed control variables, ageostrophic streamfunction and ageostrophic velocity potential. For this second experiment we use the background error covariance settings of the new CVs 1 experiment, described in table 6.1. The resulting analysis increments allow us to visualise the effect of the control variables.




Figure 6.6: Locations of the single observation experiments.

We conduct these experiments at three different locations in the domain (shown in figure 6.6):

- An active region near the western boundary current at 30° N latitude and 78° W

longitude;

- An inactive region at 30° N latitude and 65° W longitude;
- A north-eastern boundary region at 39° N latitude and 55° W longitude.

The increments produced at these three locations are largely comparable and therefore in the rest of this section we show results from the active location, given by the black cross in figure 6.6. The results from the other two locations are presented in appendix I.

In figure 6.7, we present the velocity increments generated by the single observation experiments. These figures show a $6^\circ \times 6^\circ$ area centred on the observation location, which is indicated by the black cross in figure 6.6. We first note that the increments produced by the TSCV observation using the proposed control variables appear reasonable when compared to those generated using the current balance operator. This is a positive outcome, achieved without extensive parameter tuning.

A key motivation for introducing the new control variables is that they allow for correlated ageostrophic (unbalanced) velocity components; in contrast the current NEMOVAR balance operator (3.18) assumes these components are uncorrelated and updates them independently. As shown in figure 6.7, the unbalanced velocity increments associated with the current control variables produce an isotropic update centred around the observation location, with a relatively uniform magnitude and consistent sign. In contrast, the increments resulting from the new control variables exhibit a similarly sized spatial footprint but a more complex structure, including variations in magnitude and a clear change in sign across the region.

We also observe that the full velocity increments are significantly larger in magnitude than the unbalanced components. The full speed increments are also notably larger than the unbalanced speed increments, as seen in figure 6.8. This is due to the magnitude of the unbalanced velocity background error variances, which were not calculated based on any outputs of the gyre configuration.

Therefore, the full velocity increments are not substantially influenced by the unbalanced component in this configuration. However, in more realistic ocean models with stronger ageostrophic flows, the new control variables could have a more pronounced impact.




Figure 6.7: The velocity increments at the surface for the single observation experiments. The first column are velocity increments generated from the SLA experiment. The second column are the increments from the first experiment using a TSCV observation with current control variables and the final column is the second TSCV experiment with the new control variables. The first row is the unbalanced zonal velocity increments, second row is the full zonal velocity increments, the third row is the unbalanced meridional velocity increments and the final row is the full meridional velocity increments.

The full velocity increments from the SLA observation experiment exhibit a circulation pattern around the observation location, as shown by the direction of the speed vectors in




Figure 6.8: The speed increments at the surface with velocity vectors superimposed for the single observation experiments. The first column is the speed increment from the SLA observation experiment. The second column are the speed increments from the first experiment using a TSCV observation with current control variables and the final column is the second TSCV experiment with the new control variables. The first row is the unbalanced speed increments, and the final row is the full speed increments.

figure 6.8. This pattern primarily arises from the geostrophic balance used to update the velocity field. In particular, a velocity increment at the location of the SLA innovation requires an associated SSH gradient to maintain geostrophic balance, which results in dipolar SSH features flanking the innovation location. These features, in turn, generate a rotational velocity pattern consistent with the observed circulation. In contrast, the velocity increments from the TSCV experiments display a structurally different pattern. While these increments also arise mainly from the balanced (geostrophic) component, the velocity field is further modified by contributions from the unbalanced component and spatial spreading via the diffusion operator in NEMOVAR, which is used to model spatial background error correlations (Weaver et al., 2016). Consequently, the increment pattern differs significantly from that in the SLA case, reflecting the distinct characteristics of the assimilated observations rather than a fundamental change in the balance constraints.

In summary, the single observation experiments demonstrate that the inclusion of

velocity data, along with the proposed control variables, produces sensible and physically consistent increments. These differ from the increments produced by the SLA observation as expected. These results provide confidence to proceed with assimilation experiments using a full suite of observations, as described in section 6.3.1.

6.5 Assimilation experiments

We run multiple assimilation experiments in the following sections, the details of which can be found in table 6.1. Firstly, we have a free-run of the model, as described in section 6.1. Secondly, we run two control experiments. The control experiments are the baseline experiments whereby we use the current balance operator used in the NEMOVAR system. For the first control experiment we assimilate the standard observation types. As the velocities are unobserved they are only updated through the geostrophic balance relations. For the second control experiment we also assimilate the simulated TSCV observations described in section 6.3.1 and the ageostrophic velocity increments are the velocity control variables.

Finally, we run assimilation experiments with the proposed velocity control variables - ageostrophic streamfunction and velocity potential increments. The experiments using the new control variables all assimilate the synthetic velocity observations. Without velocity observations the change in velocity control variables has no effect on the analysis, see appendix H. The first experiment with the new control variables (new CVs 1 in table 6.1) uses $\sigma_\psi, \sigma_\chi = 10^3 m^2 s^{-1}$ for the standard deviation of streamfunction and velocity potential. See section 6.3.2 for a discussion of this choice. This is varied in subsequent experiments, as described in section 6.8. The second experiment (new CVs 2) with the new control variables sets $\sigma_\psi = 10^3 m^2 s^{-1}$ and $\sigma_\chi = 0 m^2 s^{-1}$. By setting the background error standard deviation of ageostrophic velocity potential to zero, we are effectively assuming that the background is perfect for this component. As a result, the analysis increment for that variable is forced to be zero, meaning this control variable is not used during the assimilation. In practical terms, this removes the ageostrophic divergent component of the flow from the analysis update. As discussed in section 3.5, a theoretical benefit of using these new control variables is that we can control the horizontal divergence introduced by the assimilation of velocity data, thereby reducing any spurious vertical motions generated

by the data assimilation.

6.5.1 Methods of comparison

Before presenting the results of the assimilation experiments, we first detail the methods used to evaluate the analyses generated by each experiment, x^a , compared to the nature-run (the ‘truth’), x^t . We use three methods of comparison in this work all discussed and compared in section 2.6: root mean square error (RMSE), mean absolute error (MAE) and standard deviation of the error.

When assessing the analysis produced by the experiments, we are evaluating three key aspects:

1. Has the mean improved?
2. Is the variability better captured?
3. How accurate is the analysis?

A method that can be used to assess all of the above is the RMSE, defined in section by (2.20). An assessment of the mean can be demonstrated using the MAE, defined by (2.21). Finally, we calculate the standard deviation of the error, which reveals how the variability of the truth has been captured in the analysis, defined by (2.22). In the next sections we present the results of the assimilation experiments using these assessment metrics.

6.5.2 One day assimilation

Given the computational expense of running long assimilation experiments, we initially test that our data assimilation system behaves as expected on the first day. Initially, we run the free-run, control 1, control 2, new CVs 1 and new CVs 2 experiments for a one day period. We aim to test whether the increments produced by the new control variables, new CVs 1, are sensible and whether new CVs 2 can reduce the magnitude of the vertical velocities diagnosed from the increments. Throughout this chapter we will refer to the middle of the top model layer (top 10m of the model) as the surface.

Figure 6.9 shows the velocity increments on the first day of the second control experiment and the two assimilation experiments with the new control variables just




Figure 6.9: Surface velocity increments. The first column is the second control experiment which includes the synthetic velocity observations. The second column is the first experiment with the new control variables and the final column is the second experiment with the new control variables, where the divergent part of the flow is set to zero. The first row is the zonal velocity increments, second row is the meridional velocity increments and the final row is the vertical velocities diagnosed from the horizontal velocity increments.

described, at the surface. For the horizontal velocities, we clearly see that the increments for the new control variables are sensible due to the similar magnitude and pattern as the control experiment, but there are some small scale differences, which we would expect. We also examine the magnitude of the vertical velocities that are diagnosed from the horizontal velocity increments. This is calculated in NEMOVAR by integrating the continuity equation, which depends on the horizontal divergence. Figure 6.9 shows that the introduction of the new control variables has reduced the magnitude of the vertical

velocity increments diagnosed from the horizontal velocity increments. When the $\delta\chi_U$ component is removed, the increments decrease even further. The small-scale vertical velocity structures are likely associated with the ageostrophic component of the flow, whereas the larger-scale features, that we see align with the altimeter tracks, come from the balanced component. Notably, removing the ageostrophic divergent part of the flow eliminates most of the small-scale vertical structures. This is an encouraging result, as it suggests that the assimilation process is less likely to introduce spurious vertical velocities when horizontal velocity increments are applied. In the following section, we extend these experiments to a full year and compare the analyses to the ‘truth’.

6.6 Year long assimilation

In this section, we conduct year long experiments. Having established that the velocity increments using our new control variables are satisfactory after one assimilation cycle, we now investigate their impact over multiple assimilation cycles. We aim to compare their long-term performance with the control experiments and the free-run. We run the following experiments for one year; the free-run, control 1, control 2, new CVs 1 and new CVs 2; details are given in table 6.1. For each experiment, we compare the analysis to the ‘truth’ given by the nature run. As the nature run is at a higher resolution, $1/36^\circ$, we interpolate this to $1/12^\circ$, to be able to directly compare the experiments’ results with the nature run.

6.6.1 Errors averaged over the domain

For each experiment we compare the analysis to the nature run by calculating the RMSE, MAE and standard deviation over time at each depth, as described in section 6.5.1, for each model variable. Figures 6.10, 6.12, 6.13 present the RMSE, MAE and standard deviation, respectively, at the surface for each experiment over the course of the year. As expected, the free-run (blue line) yields the highest RMSE for the model variables, shown in figure 6.10, indicating the poorest performance. However, this trend is less pronounced for the velocity components, where the first control experiment (orange line) also performs poorly, as shown in figures 6.10d and 6.10e. This is due to the absence of

velocity observations in the first control experiment. Although the assimilation updates the velocities through balance relationships, the lack of direct velocity data limits the improvement in these variables compared to SSH, temperature, and salinity, which are directly observed.



Figure 6.10: RMSE calculated over the whole domain at the surface.






Figure 6.10 (continued).




Figure 6.11: Nature run minus the free-run for temperature at the surface, on the first day of the assimilation run, 1st February, (left) and 1st May (right).

A notable feature in the free-run is the decline in surface temperature RMSE, shown in figure 6.10b. This reduction is also evident in the MAE in figure 6.12b, but not in the standard deviation (figure 6.13b). To further investigate this behaviour, figure 6.11 compares surface temperature differences between the nature run and the free-run on day 1 (1st February) and on May 1st. Initially, the free-run generally underestimates surface temperature across most of the domain, with the largest discrepancies along the western boundary current. By May 1st, the temperature differences in areas away from the western boundary current decrease significantly, showing that the free-run values become closer to those of the nature run.

The control 2 experiment, given by the green line in figures 6.10, 6.12 and 6.13, incorporates synthetic velocity observations as described in section 6.3.1. This experiment shows an overall improvement in the analysis across all model variables for all error metrics. This result highlights the value of future ocean current measurements for data assimilation and aligns with the findings of Waters et al. (2024b). We run two additional year long experiments to test the new velocity control variables; one using $\sigma_\psi, \sigma_\chi = 10^3 m^2 s^{-1}$ (new CVs 1), and another with the divergent component of the ageostrophic flow removed by setting $\sigma_\chi = 0 m^2 s^{-1}$ (new CVs 2). Overall, the RMSE, MAE, and standard deviation results from the experiments using these new control variables are promising.



Figure 6.12: MAE calculated over the whole domain at the surface.




Figure 6.12 (continued).

Between May and August, the new control variables show improvements in RMSE, MAE, and standard deviation — given by the red and purple lines in figures 6.10, 6.12, and 6.13 — compared to the control experiments. This improvement is especially pronounced in the RMSE of horizontal velocities (figures 6.10d and 6.10e) where the new control variables substantially outperform the control 2 experiment. In the first half of the year, excluding the divergent component of the ageostrophic flow (purple line) has little effect on the error metrics compared to the experiment that includes it (red line). This indicates

that, during this period, the assimilation is relatively insensitive to the removal of the divergent component. This result is promising, as it suggests that the settings used in the new CVs 2 experiment may provide a realistic approach to constraining vertical velocity magnitude. This is explored further later in this chapter.




Figure 6.13: Standard deviation calculated over the whole domain at the surface.






Figure 6.13 (continued).

In contrast, during the second half of the year, particularly from August onward, the experiment that excludes the divergent component (purple line) begins to show a worse performance for horizontal velocities. This is most evident in figures 6.12d and 6.12e, where the MAE of new CVs 2 generally exceeds that of new CVs 1 (red line), and in some cases even performs worse than control 2 (green line). A similar trend is observed for RMSE and standard deviation metrics during this period. In August, both new control variable experiments show a decline in performance, with a marked increase in RMSE. Nevertheless, the experiment that retains the divergent component (red line) recovers after the August–September dip and achieves performance comparable to control 2 for most variables, as seen in figures 6.10, 6.12, and 6.13.




Figure 6.14: The top row is the difference between the zonal velocity of new CVs 1 from the nature run and the bottom row is the zonal velocity observations. These are for two dates in August, the left 19th August and the right 23rd August.

The increase in RMSE in August for both new CVs experiments, without a

corresponding rise in MAE, suggests the presence of large errors or outliers, to which RMSE is more sensitive. Further analysis indicates that the degradation begins around 21st August. To investigate this further, we examine the raw error from the new CVs 1 experiment relative to the nature run (figure 6.14, top row). The panels show the error on 19th and 23rd August, during which the RMSE spike occurs. A pronounced error increase is visible on 23rd August in the far western region of the domain. As shown in the bottom row of figure 6.14, there are no observations in this area on 23rd August, whereas observations are present on 19th August. This pattern is consistent for both zonal and meridional velocity components. The lack of observational constraints in this region appears to have led to this RMSE spike. However, the RMSE subsequently returns to improved levels after this brief degradation for new CVs 1.

For the new CVs 2 experiment, a similar increase in error is observed in the western region of the domain. Figure 6.15 shows the time-averaged RMSE for zonal velocity in July (left panel) and September (right panel). The error is clearly higher in September, particularly in the western part of the domain. The degraded performance of this experiment in the second half of the year is probably due to the absence of the divergent component in the control variable formulation as this degradation is not observed in new CVs 1. Without this component being updated by the data assimilation, the model may be less capable of dynamically adjusting the velocity fields in regions where observations are temporarily absent. This hypothesis warrants further investigation to better understand the role of the divergent flow in maintaining dynamical consistency under sparse velocity observational coverage.




Figure 6.15: Time averaged RMSE for zonal velocity in July and September at the surface for new CVs 2.

Overall, figures 6.10, 6.12 and 6.13 suggest that the inclusion of the new velocity variables leads to a more substantial reduction in the mean error than in the variability. These results indicate that the new velocity control variables have potential to enhance the assimilation of future ocean current data, which is a key objective of this work. All error metrics show reasonable performance, and with further tuning of certain parameters, we anticipate that these improvements can be made more significant.

We also examine the RMSE, MAE, and standard deviation as functions of depth, averaged over time and space (figures 6.16, 6.17, and 6.18). In the deeper ocean without direct ocean measurements, data assimilation updates the variables here indirectly through mechanisms such as through the barotropic component of the increments, which apply depth-uniform corrections derived from surface observations like SSH. For temperature and salinity, the free-run exhibits the largest errors near the surface, with control 1 consistently performing second worst across all three metrics. Below approximately 1000m, the free-run improves sharply, and its performance surpasses the assimilation experiments, which show increased errors at these depths. As this work focuses on surface velocities and no velocity observations are available at depth, we do not explore this behaviour further; however, it presents an interesting avenue for future research. The horizontal velocity fields show a similar trend, though the assimilation experiments begin to degrade at deeper depths. Interestingly, the RMSE, MAE and standard deviation see an improvement from assimilating velocity data below 1500m, even though the velocity observations are only at the surface. Overall, new CVs 1 delivers the best performance across all variables and depths, followed closely by control 2 and new CVs 2. Notably, new CVs 2 outperforms control 2 at the surface, which is an encouraging result for our new control variables.




Figure 6.16: The RMSE over the depths of the assimilation results over one year.




Figure 6.17: The MAE over the depths of the assimilation results over one year.




Figure 6.18: The standard deviations over the depths of the assimilation results over one year.

6.6.2 RMSE averaged over time

Here, we present RMSE results averaged over time to identify where in the domain the largest errors occur and how they are spatially distributed. This analysis also highlights regional differences between experiments. We first compare the second control experiment (which includes velocity observations) with the first control experiment. For SSH, temperature, and salinity, the improvements are relatively minor (not shown). However, for horizontal velocities, the addition of velocity data leads to a more substantial improvement, as shown in figure 6.19. In this figure, red indicates a reduction in error relative to the previous experiment, blue indicates an increase, and white indicates no change. The most notable improvement occurs along the western boundary current, a region characterised by strong eddy activity, suggesting that assimilating velocity data is particularly beneficial in dynamically active regions.




Figure 6.19: Difference in the RMSE averaged over time at the surface. This is the difference between horizontal velocity errors in the two control experiments. Red indicates a reduction in error from assimilating the velocity data using the old control variables compared to not assimilating velocity data.

For the experiments using the new control variables, we compare both configurations to the second control experiment. The spatial differences in RMSE, averaged over time, are shown in figure 6.20. The first column displays the difference between the new CVs 1 experiment (which includes the divergent component of the ageostrophic flow) and the second control. Once again, red indicates a reduction in error relative to the control experiment, blue indicates an increase, and white indicates no change. While the differences are not substantial across the domain, the most noticeable changes occur in eddy-rich regions, where there is a mix of localised improvements and degradations. Overall, the errors remain comparable in much of the domain for all variables. The second




Figure 6.20: Differences in the RMSE averaged over time at the surface. The first column is the control 2 minus new CVs 1 and the second column is the control 2 minus new CVs 2. The first row is SSH, and then temperature, salinity, zonal velocity and meridional velocity respectively. Red indicates a reduction in error of the new CVs experiments compared to control 2.

column shows the difference between the new CVs 2 experiment (excluding the divergent component) and the second control. In this case, we observe more widespread degradation, particularly in the velocity fields and temperature. Although the increase in temperature RMSE appears significant compared to the other variables, it is important to note that the RMSE for temperature is larger, as we see in figure 6.10b. Therefore, the relative increase in error is not as severe as it may initially seem. The notable increase in velocity RMSE is likely linked to the degradation observed after August, discussed previously in the domain-averaged RMSE plots for the zonal and meridional velocities (figures 6.10d and 6.10e). Spatially, the degradation is most pronounced in the northern part of the domain, although localised improvements appear in the southwestern corner, coinciding with high eddy activity and the western boundary current. This pattern contrasts with the elevated RMSE seen in September in figure 6.15b, which highlights the spatial and temporal variability of the error growth.

Overall, these results suggest that setting the standard deviation of velocity potential to zero has a more complex and spatially variable impact on RMSE than the globally averaged results in figure 6.10 initially suggest. While the new CVs 2 experiment leads to performance degradation in some regions — particularly in the north — it also yields improvements in dynamically active areas. This highlights a trade-off in performance and underscores the importance of further investigating the spatial sensitivity of the control variable formulation. By setting $\sigma_\chi = 0$, the background errors in the velocities are reduced, effectively giving less weight to the velocity observations. This adjustment may contribute to the observed differences between the new CVs 1 and new CVs 2 experiments.

6.6.3 Summary

In these year-long experiments, we identify that the new control variables show considerable promise. The surface-averaged errors indicate that the new control variables lead to improved performance compared to the original variables when assimilating velocity data. Notably, the performance during the first half of the year is significantly better. However, after August, where we identify a gap in the velocity observations near the western boundary current and regions of intense eddy activity, the performance of the new control variables deteriorates. In particular, the experiment in which the divergent

component is removed fails to recover and continues to underperform relative to the other experiments.

When examining performance with depth, we observe a decline in assimilation quality at deeper levels. However, since the surface layers are of primary interest, this degradation is not a major concern for this configuration; performance within the upper $2000m$ remains reasonable. In the context of an operational system, where deeper layers are likely to be of greater relevance, this warrants more detailed investigation. Time-averaged RMSE analyses show that the largest improvements and degradations in performance occur in regions of strong eddy activity. Additionally, the removal of the ageostrophic velocity potential introduces more pronounced errors across the domain. Despite the overall poorer spatial performance, there are several periods in the time series plots where this experiment performs comparably to the second control experiment and the experiment that includes the divergent component. For this reason, we continue to explore this experimental setup further and test the hypothesis that it may help reduce the magnitude of the vertical velocity.

6.7 Vertical velocities

As discussed, by setting the standard deviation of the ageostrophic velocity potential increment to zero (in experiment new CVs 2), we are removing the ageostrophic divergent component of the velocity from the analysis increments. In section 6.5 we illustrated that the magnitude of the vertical velocity increment diagnosed from the horizontal velocity increments, after one day of assimilation, appeared to decrease with this setting. Here we assess the impact on the vertical velocities over a longer period of assimilation to assess whether using the new control variables and the additional removal of the ageostrophic divergent flow helps to reduce the spurious generation of vertical velocities by the assimilation. It is important to note that the vertical velocity increments shown in figure 6.9 are diagnosed from the horizontal velocity increments. However, in NEMO vertical increments are not directly applied during the IAU; instead, the vertical velocities are adjusted by the model during the assimilation process.

In the previous section, our analysis focused primarily on the surface of the domain. However, when examining vertical velocities, we consider subsurface layers, as vertical

velocity is not defined at the surface itself due to the use of a Lorenz staggered grid (Holdaway et al., 2013). Therefore, in this section, we extend our focus to include near-surface depths.

Firstly, we evaluate the RMSE, MAE, and standard deviation of the vertical velocities, consistent with the previous section. As discussed in section 3.5, the control experiments are expected to perform poorly, since ocean data assimilation tends to artificially inflate the magnitude of vertical velocities (Ford and Barciela, 2017). Figure 6.21 shows the RMSE, averaged over the horizontal domain and a year, across different depths. As expected, the free-run performs significantly better than the control experiments. Among the assimilation runs, control 2 performs best, while both the new control variables and control 1 experiments perform similarly poorly across depths. At greater depths, the experiment removing the ageostrophic velocity potential increment shows the worst performance. These patterns are temporally consistent across all depths and are similarly evident in both the MAE and standard deviation.






Figure 6.21: RMSE of the vertical velocities across the depths over a year.

Although the performance of our new control variables with respect to vertical velocity errors is disappointing, our primary interest in this work lies in the magnitude of the vertical velocities; therefore, we examine the spatial mean of their absolute values, see figure 6.22. As expected, the free-run gives the smallest vertical velocities. However, all assimilation experiments see a sharp increase in the mean at the beginning of the assimilation period. Interestingly, during the first half of the year, the new CVs 1 experiment shows slightly improved performance; however, this advantage deteriorates

after August. This decline may be related to the broader degradation in performance of the new control variable experiments from August onward, as discussed in section 6.6. Overall, it seems as though the introduction of the new control variables does not lead to any improvement, and notably, the removal of the ageostrophic divergent flow appears to further increase the magnitude of the vertical velocities. This outcome is contrary to our initial hypothesis, which proposed that eliminating the ageostrophic divergent component would reduce vertical velocity magnitudes. The results, however, clearly suggest otherwise. It is important to note that this analysis is based on the spatial mean of the full vertical velocity field produced by the assimilation.



(a) 10m depth




(b) Averaged across the depths


Figure 6.22: Spatial mean of the absolute values of the vertical velocities averaged over the domain.

To further assess whether the new control variables offer any benefit in reducing

vertical velocities, we examine the increments diagnosed from the horizontal velocities, as shown in figure 6.9. Specifically, we calculate the spatial mean of the absolute values of these increments at various depths over the course of a year. However, the resulting time series plots are highly noisy, making it difficult to extract meaningful trends. Consequently, we focus our analysis on the month of March.



(a) 10m depth



(b) Average across the depths

Figure 6.23: Spatial mean of the absolute values of the vertical velocity increments diagnosed from the horizontal velocity increments averaged over the domain for March.

Figure 6.23 shows the spatial mean of the absolute values of the diagnosed vertical velocity increment at 10m depth, as well as the values calculated at each depth over the month. Comparing the new CVs 2 experiment to control 2, we observe that the removal of the ageostrophic divergent flow consistently reduces the magnitude of the diagnosed vertical velocity increment throughout the month. Interestingly, control 1 has small

vertical velocities at some times when examining the 10m depth (figure 6.23a). However, when averaged across all depths (figure 6.23b), it is evident that this experiment has the largest mean vertical velocity increments overall. The observed reduction in diagnosed vertical velocity increments when setting $\sigma_\chi = 0m^2s^{-1}$ confirms that we are able to reduce the magnitude of the vertical velocities diagnosed from the horizontal velocities. However, this raises an important question: if the diagnosed increments are reduced, why do the magnitudes of the vertical velocity fields from the assimilation remain so large?





Figure 6.24: Spatial mean of the absolute values of the diagnosed vertical velocity increments and difference between the vertical velocities at the end of the assimilation period and the start, averaged over the domain for the new CVs 2 experiment for the year.


To address this question, we examine the difference in vertical velocities produced by the model before and after the assimilation period. This difference demonstrates the impact of the increments applied to the model variables through the IAU, as discussed in section 6.3.1. Figure 6.24 shows the magnitude of this difference for the new CVs 2 experiment, where $\sigma_\chi = 0m^2s^{-1}$, averaged over the entire year. The difference in vertical velocity before and after assimilation is significantly larger than the diagnosed increments, indicating a strong response of the model to the assimilation process. This large discrepancy suggests that the model continues to generate spurious vertical motions. In particular, it implies that the balanced component of the vertical velocity is divergent and therefore plays a dominant role in determining the total vertical velocity magnitude. It should be noted that the diagnosed increments are defined at the same vertical levels as the horizontal velocities, and thus include values at the surface. In contrast, the model's

vertical velocity is prescribed to be zero at the surface.

To further illustrate this point, figure 6.25 presents the same analysis but for the horizontal velocity components. In this figure, the solid purple line represents the analysis increments, while the dashed line shows the difference between model fields before and after the application of the IAU (i.e., the increments after the IAU). Here, we see that the diagnosed and post-IAU increments match closely for horizontal velocities, highlighting the contrast with the vertical velocity case and making the discrepancy in vertical velocity even more striking.



(a) Zonal velocity



(b) Meridional velocity

Figure 6.25: Spatial mean of the absolute values of the horizontal velocity increments from the DA and the increments from the IAU (difference between the horizontal velocities before and after the assimilation period), averaged over the domain for the new CVs 2 experiment for the year.

Overall, we conclude that using the proposed control variables and setting $\sigma_\chi =$

$0m^2s^{-1}$ does reduce the magnitude of the diagnosed vertical velocity increments. However, the model still introduces spurious vertical velocities during the assimilation process, indicating that the balanced component of the vertical velocity is divergent and thus plays a dominant role in determining the total vertical velocity magnitude. This suggests that further work is needed to improve the NEMOVAR balance operator so that it can better account for vertical velocities.


6.8 Month long assimilation

In this section we present results from multiple month long experiments to test the sensitivity of the assimilation analysis to the specified background error standard deviation of streamfunction. These are experiments new CVs 2-5 described in table 6.1.


6.8.1 Sensitivity to standard deviations

Having examined the errors from the year-long experiments, we now investigate the sensitivity of the analysis to variations in the background error standard deviation of the streamfunction over a one-month period, while keeping the standard deviation of the velocity potential fixed (and all other aspects of the assimilation). This investigation is not intended to prescribe an optimal configuration for NEMOVAR in future experiments, but rather to offer insight into the sensitivity of the analysis to these parameter choices. The experiments are conducted for the month of March, selected because the analysis errors observed over the annual period showed a marked decrease in February and reached a stable state by March. We present only the RMSE and MAE metrics, as the year-long results demonstrated that the standard deviation patterns closely mirrored those of RMSE for the velocities. Figures 6.26 and 6.27 show the horizontal velocity RMSE and MAE results, respectively, at the surface.

In the year-long experiments, we set $\sigma_\psi = 10^3 m^2 s^{-1}$, a value chosen to be consistent with the typical magnitude of the streamfunction field, as discussed in section 6.3.2. As shown in figures 6.26 and 6.27, increasing this value by an order of magnitude to $\sigma_\psi = 10^4 m^2 s^{-1}$ (pink line) leads to a substantial degradation in performance, with RMSE and MAE consistently worse than those of the second control experiment that incorporates



(a) Zonal velocity




(b) Meridional velocity


Figure 6.26: RMSE for the month long experiments.

velocity data. This is likely due to overfitting the velocity data and indicates a high sensitivity of the assimilation to large increases in σ_ψ .

Conversely, reducing σ_ψ by one or two orders of magnitude has a much smaller impact. For $\sigma_\psi = 10^2 m^2 s^{-1}$ and $10^1 m^2 s^{-1}$ (black and gold lines), the RMSE remains comparable throughout the month-long period, with $\sigma_\psi = 10^3 m^2 s^{-1}$ generally yielding the lowest zonal velocity RMSE and occasionally outperforming in the meridional component as well. While the RMSE exhibits considerable variability — consistent with the year-long findings — the lower values of σ_ψ show similar average performance. The MAE is even less sensitive to reductions in σ_ψ , with all tested values producing largely comparable results across both velocity components. We also investigate the impact of setting $\sigma_\psi = 0$ and



(a) Zonal velocity



(b) Meridional velocity

Figure 6.27: MAE for the month long experiments.

find that the resulting errors are comparable to those obtained using the lowest standard deviations presented in this section. The lack of sensitivity to the reduction in σ_ψ may indicate that the balanced component of the flow contributes most significantly to the analysis, and that reducing the ageostrophic component has a relatively minor effect. In a more realistic configuration with spatially and temporally varying wind forcing, the ageostrophic velocity component would likely be larger than in the GYRE setup, and therefore more sensitive to reductions in the standard deviations.

We extend the RMSE and MAE analysis across depth levels in figure 6.28 and observe consistent patterns; the largest standard deviation ($\sigma_\psi = 10^4 m^2 s^{-1}$, pink line) results in significantly higher errors, while the smaller values produce nearly indistinguishable




Figure 6.28: RMSE and MAE for the month long experiments across the depths for the horizontal velocities. The top and second row are RMSE. The third and bottom row are MAE.

results. Similar trends are observed in the other model variables (SSH and salinity) whereas temperature seems to have a negligible sensitivity to the changing of the standard deviations. These findings indicate that increasing the standard deviation of the streamfunction, while keeping the velocity potential unchanged, can substantially degrade the quality of the analysis. In contrast, reducing σ_ψ by one or more orders of magnitude has minimal impact, with $\sigma_\psi = 10^3 m^2 s^{-1}$ typically yielding the most favourable results on average.

6.8.2 Summary

In this set of experiments, we investigate the sensitivity of the analysis to the new control variables by varying the specified background error standard deviation of the streamfunction field. We find that increasing the standard deviation to a magnitude comparable to the largest values of the field significantly worsens the results. In contrast, reducing the standard deviations by multiple orders of magnitude leads to an improved performance. These results offer guidance for the use of these variables within the GYRE configuration and an insight into the sensitivity to the specified background error standard deviation of the streamfunction field. However, the results would likely differ in a more realistic setup with more complex wind forcing. Ideally, these parameters should be estimated using statistical methods that provide spatially varying variances, thereby capturing regions of greater uncertainty within the domain.

6.9 Summary and conclusions

In this chapter we conducted assimilation experiments to test whether the proposed velocity control variables - ageostrophic streamfunction and velocity potential - can provide an improvement to the analysis produced by NEMOVAR, compared to the previous control variables.

After outlining the experimental setup, observations, and comparison methodology, we presented the results of multiple assimilation experiments. First, we investigated the increments generated by the single-observation experiments. These experiments demonstrated that the inclusion of velocity data, along with the proposed velocity

control variables, produced sensible increments. Following this, we conducted assimilation experiments using the full suite of available observations. Two control experiments were conducted. The first employed the standard observations currently used operationally, while the second incorporated synthetic velocity observations using the previous control variables, the ageostrophic velocities. Subsequently, two assimilation experiments were performed using ageostrophic streamfunction and velocity potential as control variables. In the second of these experiments, we removed $\delta\chi_U$, effectively eliminating increments to the divergent component of the ageostrophic flow. As anticipated, this removal led to a reduction in the magnitude of the vertical velocity increment, after one assimilation cycle.

Each experiment was then run over a one-year period. Overall, the error metrics; RMSE, MAE, and standard deviation, showed either further reduction or comparable behaviour when velocity observations were introduced and when the new velocity variables were used, marking a promising outcome. However, despite these improvements, the overall magnitude of the vertical velocities was not significantly reduced, suggesting that a substantial portion of the increase is attributable to divergence in the balanced part of the increments. Additionally, the method by which increments are introduced into the model through the IAU may also be generating spurious vertical velocities. We finally conducted month-long experiments to assess the sensitivity of the analysis to the background error standard deviation set for the streamfunction, while continuing to exclude the divergent flow component. The results showed that increasing the standard deviation of the streamfunction by an order of magnitude degraded the quality of the analysis, whereas smaller standard deviation values generally produced comparable results.

Overall, we have demonstrated that the proposed velocity control variables are able to improve the analysis produced by NEMOVAR when assimilating velocity observations using the GYRE configuration. Their implementation in more realistic models — with stronger ageostrophic signals — could yield an even greater impact. However, spurious vertical motions in the model, generated by the data assimilation, were not reduced, indicating that further investigation into the treatment of the balanced component of the flow is required.

Chapter 7

Conclusion

In this thesis we designed a new velocity CVT for use in NEMOVAR, assessing its implications and demonstrating its potential for assimilating future ocean current measurements.

Ocean forecasts are important to many applications such as numerical weather prediction and marine safety. These forecasts heavily rely on data assimilation. This is the process of combining observations with numerical models, weighting each by their associated errors. In the ocean, we commonly use variational data assimilation, in which a least-squares problem is solved. Within variational data assimilation, a transformation is performed from model variables to control variables, which simplifies the problem by assuming that the control variables are uncorrelated. In ocean data assimilation, the velocity control variables have required less attention due to the limited assimilation of velocity observations. However, with the prospect of future ocean current measurements, these variables require more careful consideration.

In chapter 2, we introduced the key mathematical concepts for this thesis, such as Helmholtz Theorem and the Shapiro filter. In chapter 3, we presented an overview of data assimilation, with a focus on incremental variational data assimilation and the control variable transform (CVT). We reviewed the existing CVTs used in variational ocean data assimilation and discussed the limitations of the current velocity variables in NEMOVAR. We also presented the method proposed by Li et al. (2006), proposing its use in NEMOVAR.

In chapter 4, we developed a CVT using these new variables within a shallow water

model, exploring the numerical challenges involved in implementing the transformation. In chapter 5, we extended this work to NEMOVAR, highlighting the intricacies of using these variables in an operational framework, as well as their sensitivity to parameters and the persistent numerical issues within the CVT. Finally, in chapter 6, we conducted assimilation experiments to demonstrate the potential positive impact these new variables could have within NEMOVAR.

7.1 Research questions

In chapter 1, we identified three key questions that the research in this thesis set out to answer:

1. What numerical challenges and implications arise from the discrete transformation of the velocities to their irrotational and non-divergent parts?
2. How can the discretisation of this transformation be extended from an idealised framework to application within a realistic ocean model?
3. What are the implications of using these alternative velocity variables in variational data assimilation, and how do these variables respond to the assimilation of velocity observations?

We now address each of these based on the work presented in chapters 4 - 6.

1. What numerical challenges and implications arise from the discrete transformation of the velocities to their irrotational and non-divergent parts?

In chapter 4, we performed the discrete transformation from the ageostrophic velocities into their irrotational and non-divergent parts using a shallow water model. To achieve this, we applied the method proposed by Li et al. (2006), which formulates the discrete Helmholtz decomposition as a linear system. We solved this system by minimising a cost function to obtain a least-squares solution. Due to the non-uniqueness of the Helmholtz decomposition in our domain, the method uses Tikhonov regularisation. The Li et al.

(2006) method additionally moves the discrete location of streamfunction on the Arakawa-C grid to circumvent the need for explicit boundary conditions. In this work, we identified an alternative approach to perform the discrete Helmholtz decomposition, by solving the Poisson equations, (2.6) and (2.7), directly, as discussed in detail in appendix D. We found this method to be equivalent to the cost function approach, but more intricate to implement due to placement of streamfunction on the grid. Consequently, we opted to proceed with the cost function method.

We formulated a novel control variable transformation based on the shallow water model and defined the U- and T-transforms. In examining how to compute the streamfunction derivatives in (4.16) and (4.17), we found two equivalent methods on a uniform grid. However, both imply the existence of implicit (Neumann) boundary conditions on streamfunction — an implication not thoroughly discussed in the literature. This arose due to averaging streamfunction to the boundary when calculating the derivatives. We provided a detailed mathematical explanation of how these Neumann boundary conditions arise.

In implementing the Li et al. (2006) method, we encountered an issue of grid-scale noise, which appeared as a checkerboard pattern in the velocity reconstructions. This problem stems from the averaging of streamfunction derivatives and was evident both analytically and numerically in the velocity error fields. Through spatial correlation analysis using the SWEs, we showed how this pattern can contaminate the statistics of the control variables. To mitigate this, we proposed using a second order Shapiro filter. A single iteration proved effective in removing the checkerboard noise. We examined the placement of the filter within the transformation process and determined that post-processing was necessary to suppress the noise effectively. However, we also observed that applying the filter unfortunately introduced larger reconstruction errors in the velocities.

Overall, we developed a new control variable transformation with ageostrophic streamfunction and velocity potential as the velocity control variables. This was implemented in the shallow water model and led to the discovery of several critical challenges associated with applying the Li et al. (2006) method for decomposing velocities into their non-divergent and irrotational components.

2. How can the discretisation of this transformation be extended from an idealised framework to application within a realistic ocean model?

In chapter 5, we integrated the proposed velocity control variables — ageostrophic streamfunction and velocity potential — into NEMOVAR. Building on the work from chapter 4, we extended the approach from the shallow water model to the more realistic ocean modelling framework, NEMO. To achieve this, we updated the existing balance operator in NEMOVAR to include a transformation from the ageostrophic velocity components to ageostrophic streamfunction and velocity potential. This required formulating the Helmholtz decomposition in discrete form using spherical coordinates. We described implementing these control variables in a more complex system compared to the SWEs, including the need for spatial weighting and a more sophisticated approach to regularisation.

After defining the new balance operator and the discretised Helmholtz decomposition within it, we demonstrated the transformation of velocity fields into their non-divergent and irrotational components using the GYRE configuration of NEMO. We examined the impact of varying the convergence tolerance and regularisation parameter by reconstructing the velocities after the transformation. We found that, beyond a certain threshold, further reductions in these parameters had negligible impact on reconstruction errors but incurred substantial computational cost.

Using a Fourier analysis of both the streamfunction field and the reconstruction errors, we observed that the checkerboard pattern, previously identified in the shallow water model, also appeared in this more complex domain and coordinate system. To address this, we proposed the use of the Shapiro filter, a method already used in NEMOVAR, to suppress this numerical artifact. To determine the number of iterations required, we again performed a Fourier analysis and found that a single iteration was sufficient to remove the checkerboard noise. However, the Shapiro filter led to an increase in the reconstruction errors. We revealed that this was due to the filter not only removing the small-scale noise associated with the checkerboard pattern but also modifying large-scale wave components.

Overall, we successfully incorporated the new velocity control variables into the complex balance operator used in NEMOVAR, along with the discretised Helmholtz decomposition

adapted to the NEMO coordinate system. We demonstrated the transformations sensitivity to parameter choices and confirmed the presence of the checkerboard pattern. We concluded that, to effectively overcome this issue, a filtering approach that more selectively targets small-scale waves is necessary.

3. What are the implications of using these alternative velocity variables in variational data assimilation, and how do these variables respond to the assimilation of velocity observations?

In chapter 6, we performed new assimilation experiments comparing the current velocity control variables — ageostrophic velocities — with the newly proposed ageostrophic streamfunction and velocity potential, in NEMOVAR. These represent the first experiments to use the new velocity control variables alongside idealised SKIM-like satellite velocity observations, within the GYRE configuration of NEMO.

Single-observation experiments demonstrated that the inclusion of these new velocity control variables produced sensible analysis increments compared to the current system. Over the course of year-long experiments, we demonstrated that the inclusion of these new variables in NEMOVAR has the potential to improve the quality of the analysis. Notably, this improvement was observed even without extensive tuning of the background error statistics. We also conducted experiments in which the divergent component of the ageostrophic velocities was removed by setting the ageostrophic velocity potential increment to zero. This led to a reduction in the magnitude of the ageostrophic vertical velocity increments, a significant benefit of the new control variables. However, we also found that this approach did not reduce the overall magnitude of the full vertical velocity increments. Instead, the balanced component of the vertical velocity was responsible for the increase. This highlights a limitation and suggests that additional strategies are required to effectively constrain vertical velocities in the system. We investigated the sensitivity of the analysis to the standard deviation of the streamfunction while keeping the divergent part of the flow removed. We found that the analysis was more sensitive to larger values of the standard deviation, whereas smaller values produced relatively similar results.

Overall, we have demonstrated the positive impact new velocity control variables can

have on the analysis produced by NEMOVAR when assimilating future ocean current measurements. We have shown that the assimilation results are sensitive to the choice of background error statistics and that constraining the vertical velocities requires more than simply introducing new unbalanced control variables.

7.2 Main conclusions

The three key conclusions of this thesis are:

1. We **designed** a novel control variable transformation using alternative velocity variables for use in NEMOVAR.
2. We **discovered** new implications of implementing this transformation that emerged from a detailed mathematical analysis.
3. We **demonstrated** the potential for improving the assimilation of future ocean current measurements using the new control variable transformation.

7.3 Future work

This thesis set out to provide an insight into how ageostrophic streamfunction and velocity potential can be used as the velocity control variables in NEMOVAR. We discussed some of the challenges and impacts of using these. We now consider how this work may be extended in order to provide further insight. In the long term, the main aim of future research is to implement these assimilation experiments within a full global model, incorporating a realistic observing system and background error statistics derived from more robust methodologies. In the remainder of this section, we outline several steps that can be taken to move toward this goal.

Firstly, our experiments focused on a single type of domain, a closed ocean basin. Since the control variables were designed to avoid the need for explicit boundary conditions, it would be valuable to investigate their use in other domain types, such as limited-area domains, as done in Li et al. (2006), or global domains with islands and narrow straits. This would provide further insight into how the issues we observed, such as

the checkerboard pattern, filtering effects, and boundary-related challenges, manifest themselves in more varied domains.

Another task for further research is to investigate the cross-correlations of the proposed control variables. In section 3.5, we discussed that control variables based on streamfunction and velocity potential are more suitable than those based directly on velocity components. A statement frequently agreed upon in literature based on the conclusions of Daley (1993). This was briefly explored in appendix A, where we examined the cross-correlations of the proposed control variables. These preliminary experiments indicated that ageostrophic streamfunction and velocity potential were actually more correlated than the ageostrophic velocity components, regardless of whether the Shapiro filter was applied. One possible contributing factor is the modification of the Helmholtz decomposition through the use of Tikhonov regularisation, which may have affected the independence of the transformed variables. However, these findings are not conclusive. Further investigation is required to assess the correlation between ageostrophic streamfunction and velocity potential increments. This should ideally be carried out in a more realistic setting, such as the GYRE configuration or a global configuration, to provide a more definitive understanding.

On the topic of statistics, the assimilation experiments presented in chapter 6 were conducted using the new control variables. As discussed in section 6.3.2, the standard deviation values for the control variables were assumed to be constant throughout the entire domain and were chosen based on the magnitude of the model fields. We subsequently tested the sensitivity of the assimilation to changes in the standard deviation of the streamfunction variable. However, these assimilation experiments should be repeated using a more robust approach, such as the NMC method (Parrish and Derber, 1992), to estimate the background error statistics. In appendix A we found strong cross correlations between the proposed velocity control variables, using a rudimentary method to estimate forecast errors (namely the "quick covs" method). A better forecast error estimation method, such as the NMC method, may provide a more accurate estimation of these cross correlations, which may be found to be smaller for these proposed variables. If this were found, it would be a further validation in favour of the use of ageostrophic streamfunction and velocity potential as velocity control variables.

Another valuable next step in extending this work would be to consider the filtering

process more analytically. As discussed in section 5.4.1, there is a need to replace the Shapiro filter with an alternative that more selectively targets high-frequency components of the FFT spectrum. Future research should aim to identify and evaluate such a filter. Importantly, this investigation should go beyond numerical experiments and include a thorough analytical study to better understand the filtering properties and their impact on the control variables and overall assimilation performance.

The assimilation experiments, conducted in chapter 6, could also be extended by incorporating in situ data, such as HF radar observations, to evaluate how these new variables might impact NEMOVAR in current operational settings, rather than only in future scenarios. Additional assessment methods could be introduced, such as using Lagrangian techniques to track particle trajectories; testing uncertainty by perturbing the initial conditions of the assimilation experiments; and evaluating forecast skill in addition to analysis accuracy.

As concluded in section 6.7, the introduction of the new control variables did not succeed in constraining the vertical velocities. However, this issue appeared to be primarily related to the balanced component of the velocity field. Further work is needed to understand why the balanced component of the vertical velocities increases and how this behaviour might be better constrained within the assimilation system.

While we have highlighted the unresolved issues associated with using these alternative velocity variables in NEMOVAR, this thesis has demonstrated that they hold the potential for improving the current ocean data assimilation system. As such, this approach presents an exciting and worthwhile opportunity for further research.

Appendix A

Cross-correlations of control variables using the SWEs

In section 4.3 we discussed the numerical implications of transforming to ageostrophic streamfunction and velocity potential from the horizontal velocities, using the shallow water model. Here, we investigate the cross-correlations of the proposed control variables and how using the Shapiro filter affects these correlations. The correlation analysis described in this section will test the validity of the assumption that the control variable errors are uncorrelated. In particular, we discuss whether the newly proposed control variables are less correlated than the ageostrophic velocities; taking into consideration the numerical implications and solutions discussed in section 4.3. We aim to address these points using the numerical model described in section 4.1, by calculating the correlations between the background errors of the control variables. To begin, we detail the method to calculate these correlations.

A.1 Method

To examine the correlations between control variables, we do so by calculating the cross-variable correlations at each grid point using the ‘quick covs method’ (Polavarapu et al., 2005), similar to section 4.3.3. Using this method, we take differences between model forecasts at regular time intervals. These forecast differences act as a proxy for background errors. We take differences until we have a larger enough data set. Using these differences,

we look at the correlations of the difference fields. We investigate the correlation between control variables by transforming the forecast time-differences of the state variables using the T-transform, described in section 4.2.2. The following steps are taken to perform the correlation analysis using the shallow water model:

1. Run the shallow water model described in section 4.1 for 200 days after spinning the model up for 200 days.
2. Calculate the day-to-day model differences as a proxy for the background error: $\delta\eta$, δu and δv .
3. Calculate the current velocity control variables, δu_U and δv_U , using geostrophic balance (4.12), for each daily difference field.
4. Calculate the new velocity control variables, $\delta\psi_U$ and $\delta\chi_U$, using the T-transform, for each daily increment.
5. Calculate the correlations between control variables increments at each grid point.

As discussed in section 4.3.3 we use a 24-hour interval for our forecasts. We are solely concerned with the correlations between the velocity control variables, $\text{corr}(\delta u_U, \delta v_U)$ and $\text{corr}(\delta\psi_U, \delta\chi_U)$.




Figure A.1: The cross-correlations of the control variables across the domain. Titles include the control variables and the percentage of the field that could be considered approximately uncorrelated i.e. the % of the field below a certain threshold, chosen to be 0.2. $\mathcal{F}$ represents the Shapiro filter.

A.2 Results and discussion

First, we calculate the correlations between the ageostrophic velocities, the control variables currently used in NEMOVAR. The correlation field, $\text{corr}(\delta u_U, \delta v_U)$, is shown in figure A.1a. To further understand how correlated these two variables are we calculate the percentage of the field that could be considered approximately uncorrelated, by finding the % of the field below a certain threshold, chosen to be 0.2. We find that only a third of the correlation field could be considered approximately uncorrelated.

We then consider the correlations of the proposed velocity control variables, ageostrophic streamfunction and velocity potential. Using the method described in section A.1 we calculate the correlations. Figure A.1b shows the correlation field, with 27.8% of the field considered uncorrelated. This field clearly displays the checkerboard pattern discussed in section 4.3.2, therefore we apply the Shapiro filter, $\mathcal{F}$, to the ageostrophic streamfunction increment, after performing the transformation from the horizontal velocities. We then calculate $\text{corr}(\mathcal{F}\delta\psi_U, \delta\chi_U)$. The resulting correlation field is shown in figure A.1c, with only 28.4% of the field uncorrelated.

It is often considered that control variables based on streamfunction and velocity potential are more appropriate than the velocity vector components since the assumption that they are uncorrelated is more suitable (Daley, 1993, Section 5.2). Daley (1993) assumes there is no correlation between velocity potential and streamfunction. The justification for this comes from the observational and theoretical work of Hollingsworth and Lönnberg (1986) and Panchev (1971) respectively. However, through our correlation analysis we find that ageostrophic streamfunction and velocity potential are more correlated than the ageostrophic velocities, regardless of the use of the Shapiro filter. We must consider that altering the Helmholtz transformation by using Tikhonov's regularisation could have contributed to these higher correlations. Alternatively, these results could be improved by using a better background error proxy method. However, further study is needed to provide a concrete conclusion regarding the correlation between ageostrophic streamfunction and velocity potential increments.

Appendix B

Geostrophic balance

When calculating the balanced component of the velocities using geostrophic balance, we also introduce a computational mode that leads to noisy structures. This can be demonstrated on a uniform grid by calculating $u_{B_{i+\frac{1}{2},j}}$ using (4.12),

$$\begin{aligned} u_{B_{i+\frac{1}{2},j}} &= \frac{1}{4} \left(\frac{\eta_{i,j+1} - \eta_{i,j}}{\Delta y} + \frac{\eta_{i,j} - \eta_{i,j-1}}{\Delta y} + \frac{\eta_{i+1,j+1} - \eta_{i+1,j}}{\Delta y} + \frac{\eta_{i+1,j} - \eta_{i+1,j-1}}{\Delta y} \right) \\ &= \frac{1}{4\Delta y} \left(\eta_{i,j+1} - \eta_{i,j-1} + \eta_{i+1,j-1} - \eta_{i+1,j+1} \right). \end{aligned}$$

A similar expression can be found for v_B . The balanced component of the velocity has no dependence on the adjacent elevation values, i.e. $\eta_{i,j}$ and $\eta_{i+1,j}$. This leads to the checkerboard pattern discussed in section 4.3.2. This is a disadvantage of using the Arakawa-C grid, as discussed by Randall (1994). Marshall et al. (1997) and Weller et al. (2012) discuss a similar Coriolis mode on the Arakawa-C grid. To resolve this issue when performing the transformation in section 4.3.1 with the SWEs, we apply the Shapiro filter to the ageostrophic components of the velocities, see section 4.3.4. This removes any noise that would negatively affect the Helmholtz decomposition.

Appendix C

Approaches to implementing the CVT for the shallow water model

In this section we describe and illustrate the potential CVT processes to consider when implementing the alternative velocity control variables, ageostrophic streamfunction and velocity potential. We describe how we transform from the model variables to these control variables and how there are two approaches to solving the Helmholtz decomposition for streamfunction and velocity potential.

C.1 One system approach

The one system approach is as follows: from the model output we find the daily increments. We apply geostrophic balance to these increments and calculate the ageostrophic velocity increments. We transform from the ageostrophic velocity increments to ageostrophic streamfunction and velocity potential increments using the inverse of the Helmholtz decomposition. This process is detailed in section 4.2 and illustrated in figure C.1.

C.1.1 Adding the Shapiro filter

In section 4.3.4 and appendix B we discuss applying the Shapiro filter to the ageostrophic velocity increments and the ageostrophic streamfunction increment. We update the discretisation of the linear system in section 4.2.3 to include this filtering as follows:

1. Find δu_U and δv_U using geostrophic balance (4.12).
2. Apply the Shapiro filter to the ageostrophic velocities to give, $\mathcal{F}\delta u_U$ and $\mathcal{F}\delta v_U$.
3. We find the least squares solution to the following linear system,

$$\mathbf{A}\delta\mathbf{z} = \delta\mathbf{v}_{U_f}$$

where

$$\delta\mathbf{z} = \begin{pmatrix} \delta\psi_U \\ \delta\chi_U \end{pmatrix} \quad \text{and} \quad \delta\mathbf{v}_{U_f} = \mathcal{F}\delta\mathbf{v}_U = \begin{pmatrix} \mathcal{F}\delta u_U \\ \mathcal{F}\delta v_U \end{pmatrix}$$

and

$$\mathbf{A} = \begin{pmatrix} -\Delta_y & \Delta_x \\ \Delta_x & \Delta_y \end{pmatrix} \tag{C.1}$$

by minimising the following cost function with respect to $\delta\mathbf{z}$,

$$\mathcal{J}_f(\delta\mathbf{z}) = \frac{1}{2}(\mathcal{F}\delta\mathbf{v}_U - \mathbf{A}\delta\mathbf{z})^T(\mathcal{F}\delta\mathbf{v}_U - \mathbf{A}\delta\mathbf{z}). \tag{C.2}$$

This involves setting the gradient of (C.2) equal to zero,

$$\nabla \mathcal{J}_f(\delta\mathbf{z}) = -\mathbf{A}^T(\mathcal{F}\delta\mathbf{v}_U - \mathbf{A}\delta\mathbf{z}) = 0,$$

$$\mathbf{A}^T\mathbf{A}\delta\mathbf{z} = \mathbf{A}^T\mathcal{F}\delta\mathbf{v}_U$$

$$\delta\mathbf{z} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathcal{F}\delta\mathbf{v}_U.$$

4. Apply the Shapiro filter to $\delta\psi_U$ to give $\mathcal{F}\delta\psi_U$.
5. Use $\mathcal{F}\delta\psi_U$ and $\delta\chi_U$ as the velocity control variables.

This updated process is illustrated in figure C.2.

C.1.2 Alternative filtering approach

In section 4.3.4 we discuss applying the Shapiro directly to the streamfunction derivatives within the Helmholtz transformation. The discretisation of the linear system in section 4.2.3 including this filtering is given by:

1. Find δu_U and δv_U using geostrophic balance (4.12).

2. We find the least squares solution to the following linear system,

$$\mathbf{A}^f \delta \mathbf{z} = \delta \mathbf{v}_U$$

where

$$\delta \mathbf{z} = \begin{pmatrix} \delta \psi_U \\ \delta \chi_U \end{pmatrix} \quad \text{and} \quad \delta \mathbf{v}_U = \begin{pmatrix} \delta u_U \\ \delta v_U \end{pmatrix}$$

and

$$\mathbf{A}^f = \mathbf{A}\mathbf{F} = \begin{pmatrix} -\Delta_y & \Delta_x \\ \Delta_x & \Delta_y \end{pmatrix} \begin{pmatrix} \mathcal{F} & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} -\mathcal{F}\Delta_y & \Delta_x \\ \mathcal{F}\Delta_x & \Delta_y \end{pmatrix} \quad (\text{C.3})$$

by minimising the following cost function with respect to $\delta \mathbf{z}$,

$$\mathcal{J}_f(\delta \mathbf{z}) = \frac{1}{2}(\delta \mathbf{v}_U - \mathbf{A}\mathbf{F}\delta \mathbf{z})^T(\delta \mathbf{v}_U - \mathbf{A}\mathbf{F}\delta \mathbf{z}). \quad (\text{C.4})$$

This involves setting the gradient equal to zero,

$$\nabla \mathcal{J}_f(\delta \mathbf{z}) = -\mathbf{F}^T \mathbf{A}^T (\delta \mathbf{v}_U - \mathbf{A}\mathbf{F}\delta \mathbf{z}) = 0,$$

$$\mathbf{F}^T \mathbf{A}^T \mathbf{A}\mathbf{F}\delta \mathbf{z} = \mathbf{F}^T \mathbf{A}^T \delta \mathbf{v}_U$$

$$\delta \mathbf{z} = (\mathbf{F}^T \mathbf{A}^T \mathbf{A}\mathbf{F})^{-1} \mathbf{F}^T \mathbf{A}^T \delta \mathbf{v}_U.$$

3. Use $\delta \psi_U$ and $\delta \chi_U$ as the velocity control variables.

This process is shown in figure C.3. This approach was not adopted as the CVT process using the shallow water model. We showed in section 4.3.4 that this placement of the filter failed to adequately remove the checkerboard pattern.

C.2 Two system approach

We refer to the discrete linear system described in section 4.2.3. There we solved for ageostrophic streamfunction and velocity potential by minimising (4.15) directly with respect to $\delta \mathbf{z}$ using the conjugate gradient (CG) method (see section 2.2.2). We refer to this as the one system approach, and the workflow is illustrated as a flow diagram in figure C.1. An alternative approach involves solving the Helmholtz equations for streamfunction and velocity potential is by separating the calculation of $\delta \psi_U$ and $\delta \chi_U$, i.e. solving the divergence and vorticity equations,

$$\mathcal{D}_U = \frac{\partial \delta u_U}{\partial x} + \frac{\partial \delta v_U}{\partial y} = \nabla^2 \delta \chi_U, \quad (\text{C.5})$$

$$\xi_U = \frac{\partial \delta v_U}{\partial x} - \frac{\partial \delta u_U}{\partial y} = \nabla^2 \delta \psi_U. \quad (\text{C.6})$$

We discretise these equations to give

$$\mathbf{L} \delta \chi_U = \hat{\mathcal{D}}_U \quad (\text{C.7})$$

and

$$\mathbf{L} \delta \psi_U = \hat{\xi}_U, \quad (\text{C.8})$$

where $\mathbf{L}$ is the discrete Laplacian operator.

Similar to the one system approach, we use Tikhonov's regularisation to address the non-uniqueness of the Helmholtz decomposition (see section 4.2.3). To find the minimum of the cost function, (4.15) in the one system approach, we set the gradient equal to zero,

$$\begin{aligned} \Delta \mathcal{J}_\mu(\delta \mathbf{z}) &= -\mathbf{A}^T(\delta \mathbf{v}_U - \mathbf{A} \delta \mathbf{z}) + \mu \delta \mathbf{z} = 0 \\ \Rightarrow \mathbf{A}^T \mathbf{A} \delta \mathbf{z} + \mu \delta \mathbf{z} &= \mathbf{A}^T \delta \mathbf{v}_U \\ \Rightarrow (\mathbf{A}^T \mathbf{A} + \mu \mathbf{I}) \delta \mathbf{z} &= \mathbf{A}^T \delta \mathbf{v}_U \end{aligned} \quad (\text{C.9})$$

where $\mathbf{I}$ is the identify matrix and $\mathbf{A} = \begin{pmatrix} -\Delta_y & \Delta_x \\ \Delta_x & \Delta_y \end{pmatrix}$. We note that $\mathbf{A} = \mathbf{A}^T$ and therefore

$$\mathbf{A}^T \mathbf{A} = \begin{pmatrix} \Delta_x^2 + \Delta_y^2 & 0 \\ 0 & \Delta_x^2 + \Delta_y^2 \end{pmatrix} = \begin{pmatrix} \mathbf{L} & 0 \\ 0 & \mathbf{L} \end{pmatrix} \quad (\text{C.10})$$

where Δ_x^2 and Δ_y^2 are the discretisation of $\frac{\partial^2}{\partial x^2}$ and $\frac{\partial^2}{\partial y^2}$. We also have

$$\mathbf{A}^T \delta \mathbf{v}_U = \begin{pmatrix} -\Delta_y \delta u_U + \Delta_x \delta v_U \\ \Delta_x \delta u_U + \Delta_y \delta v_U \end{pmatrix} = \begin{pmatrix} \hat{\xi}_U \\ \hat{\mathcal{D}}_U \end{pmatrix}$$

and so (C.9) can separate into two separate Poisson equations, each with their own regularisation parameter,

$$(\mathbf{L} + \mu_\chi \mathbf{I}) \delta \chi_U = \hat{\mathcal{D}}_U \quad (\text{C.11})$$

and

$$(\mathbf{L} + \mu_\psi \mathbf{I}) \delta \psi_U = \hat{\xi}_U. \quad (\text{C.12})$$

To solve (C.11) and (C.12) to find streamfunction and velocity potential we must address the fact that the 2D Laplacian matrix is singular with both negative and zero value

eigenvalues due to the Neumann boundary conditions imposed on streamfunction, see sections 2.5 and 2.5. To transform to a solvable matrix problem, we must first have positive eigenvalues, and so we instead solve,

$$-(\mathbf{L} + \mu_\chi \mathbf{I})\delta\chi_U = -\hat{\mathcal{D}}_U \quad (\text{C.13})$$

and

$$-(\mathbf{L} + \mu_\psi \mathbf{I})\delta\psi_U = -\hat{\xi}_U. \quad (\text{C.14})$$

We refer to this alternative approach as the two system approach, and the flow diagram is shown in figure C.4. Through the Lagrangian approach in appendix D, we are able to demonstrate that the Laplacian operator used in (C.14) must be carefully discretised due to the location of streamfunction on the grid.




Figure C.1: One system approach to finding the ageostrophic streamfunction and velocity potential increments. From the model output we find the daily increments, then use geostrophic balance to find the ageostrophic velocity increments. We then find the ageostrophic streamfunction and velocity potential increments through the inverse of the Helmholtz decomposition.




Figure C.2: One system approach to finding the ageostrophic streamfunction and velocity potential increments with the Shapiro filter after the transformation. From the model output we find the daily increments, then use geostrophic balance to find the ageostrophic velocity increments and apply the Shapiro filter. We then find the the ageostrophic streamfunction and velocity potential increments through the inverse of the Helmholtz decomposition. We then again apply the Shapiro filter to the ageostrophic streamfunction increment.




Figure C.3: One system approach to finding the ageostrophic streamfunction and velocity potential increments within the transformation. From the model output we find the daily increments, then use geostrophic balance to find the ageostrophic velocity increments. We then find the the ageostrophic streamfunction and velocity potential increments through the inverse of the Helmholtz decomposition. The Shapiro filter is applied to the ageostrophic streamfunction derivatives within the transformation.




Figure C.4: Two system approach to finding the ageostrophic streamfunction and velocity potential increments. From the model output we find the daily increments, then use geostrophic balance to find the ageostrophic velocity increments. We then find the ageostrophic divergence and vorticity. Finally, we find the ageostrophic streamfunction and velocity potential increments through the Poisson equations.

Appendix D

A Lagrangian approach

In this section we describe how Lagrangian multipliers can be used to derive the Poisson equations (2.6) and (2.7) from Helmholtz Theorem. The Lagrangian approach also reveals the boundary conditions we impose on the velocities. We focus on the use of Helmholtz Theorem on a uniform grid as described in chapter 4. We first address the Helmholtz equations in the continuous setting. We then adapt this to investigate the effect of introducing Tikhonov’s regularisation, as used in the Li et al. (2006) method. In the discrete case we also derive similar equations. However, we clearly show that the Laplacian operator applied to ψ is different from the standard Laplacian operator. This is due to placing ψ at the centre of the grid and the averaging involved in finding the derivatives, as discussed in section 4.3.2. The following two sections are self-contained notation-wise, and we discuss the general Helmholtz decomposition for the full model fields, rather than the ageostrophic increments. The results are equivalent analytically.

D.1 Continuous equations

In this section, we formulate an energy minimisation problem to find the velocity fields u and v , constrained by the Helmholtz decomposition, that best approximate the given fields $\tilde{u}$ and $\tilde{v}$. We want to minimise the following,

$$\mathcal{J}(\mathbf{x}) = \frac{1}{2} \int_Y \int_X (u - \tilde{u})^2 + (v - \tilde{v})^2 \, dxdy. \quad (\text{D.1})$$

subject to the constraint $\mathbf{x} = \mathcal{H}(\mathbf{z})$. Here, $\mathbf{x} = \begin{pmatrix} u \\ v \end{pmatrix}$, $\mathbf{z} = \begin{pmatrix} \psi \\ \chi \end{pmatrix}$ and $\mathcal{H}$ is the Helmholtz relation detailed in theorem 2.1.2. To do so, in a continuous setting, we define a Lagrangian,

$$\mathcal{L} = \mathcal{J}(\mathbf{x}) + \int_S \boldsymbol{\lambda}^T (\mathbf{x} - \mathcal{H}(\mathbf{z})) dS \quad (\text{D.2})$$

where $\boldsymbol{\lambda} = \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix}$ are the Lagrangian multipliers. This can be re-written as

$$\mathcal{L} = \mathcal{J}(\mathbf{x}) + \int_Y \int_X \left[\lambda_1 \left(u - \left(-\frac{\partial \psi}{\partial y} + \frac{\partial \chi}{\partial x} \right) \right) + \lambda_2 \left(v - \left(\frac{\partial \psi}{\partial x} + \frac{\partial \chi}{\partial y} \right) \right) \right] dx dy \quad (\text{D.3})$$

where X and Y define the spatial domain of integration. We find the stationary values of $\mathcal{L}$. These are found by computing the first order variation of $\mathcal{L}$ with respect to the perturbations $\delta\lambda_1, \delta\lambda_2, \delta u, \delta v, \delta\psi$ and $\delta\chi$, as follows,

$$\begin{aligned} \delta\mathcal{L} = \delta\mathcal{J} + \int_Y \int_X & \left[\delta\lambda_1 \left(u - \left(-\frac{\partial \psi}{\partial y} + \frac{\partial \chi}{\partial x} \right) \right) + \delta\lambda_2 \left(v - \left(\frac{\partial \psi}{\partial x} + \frac{\partial \chi}{\partial y} \right) \right) \right. \\ & \left. + \lambda_1 \left(\delta u - \left(-\frac{\partial \delta\psi}{\partial y} + \frac{\partial \delta\chi}{\partial x} \right) \right) + \lambda_2 \left(\delta v - \left(\frac{\partial \delta\psi}{\partial x} + \frac{\partial \delta\chi}{\partial y} \right) \right) \right] dx dy \end{aligned} \quad (\text{D.4})$$

with

$$\delta\mathcal{J} = \int_Y \int_X (u - \tilde{u})\delta u + (v - \tilde{v})\delta v dx dy. \quad (\text{D.5})$$

To find the stationary values of $\mathcal{L}$ we require $\delta\mathcal{L}$ to vanish for any given perturbation. Therefore, to establish these conditions (D.4) must be reduced to an expression that isolates the perturbations. As the first order perturbations are arbitrary, the stationarity condition is satisfied if and only the coefficients of the perturbations, $\delta\lambda_1, \delta\lambda_2, \delta u, \delta v, \delta\psi$ and $\delta\chi$, are equal to zero. First, we consider coefficients of the δu and δv from (D.4) and (D.5). By setting these coefficients equal to zero, we obtain

$$u - \tilde{u} + \lambda_1 = 0 \Rightarrow \lambda_1 = \tilde{u} - u, \quad (\text{D.6})$$

$$v - \tilde{v} + \lambda_2 = 0 \Rightarrow \lambda_2 = \tilde{v} - v. \quad (\text{D.7})$$

Next, consider the first-order derivatives,

$$\int_Y \lambda_1 \frac{\partial \delta\psi}{\partial y} dy = \left[\lambda_1 \delta\psi \right]_Y - \int_Y \frac{\partial \lambda_1}{\partial y} \delta\psi dy, \quad (\text{D.8})$$

$$\int_X \lambda_2 \frac{\partial \delta\psi}{\partial x} dx = \left[\lambda_2 \delta\psi \right]_X - \int_X \frac{\partial \lambda_2}{\partial x} \delta\psi dx, \quad (\text{D.9})$$

$$\int_X \lambda_1 \frac{\partial \delta \chi}{\partial x} dx = \left[\lambda_1 \delta \chi \right]_X - \int_X \frac{\partial \lambda_1}{\partial x} \delta \chi dx, \quad (\text{D.10})$$

$$\int_Y \lambda_2 \frac{\partial \delta \chi}{\partial y} dy = \left[\lambda_2 \delta \chi \right]_Y - \int_Y \frac{\partial \lambda_2}{\partial y} \delta \chi dy. \quad (\text{D.11})$$

We substitute (D.6) and (D.7) into the boundary terms of (D.8) - (D.11),

$$\left[\lambda_1 \delta \psi \right]_Y = \left[(\tilde{u} - u) \delta \psi \right]_Y = 0,$$

$$\left[\lambda_2 \delta \psi \right]_X = \left[(\tilde{v} - v) \delta \psi \right]_X = 0,$$

$$\left[\lambda_1 \delta \chi \right]_X = \left[(\tilde{u} - u) \delta \chi \right]_X = 0,$$

$$\left[\lambda_2 \delta \chi \right]_Y = \left[(\tilde{v} - v) \delta \chi \right]_Y = 0,$$

due to the no-flow through condition, $u = \tilde{u} = 0$ and $v = \tilde{v} = 0$ on the boundary. This reduces (D.4) to,

$$\begin{aligned} \delta \mathcal{L} = \int_Y \int_X & \left[\delta \lambda_1 \left(u - \left(-\frac{\partial \psi}{\partial y} + \frac{\partial \chi}{\partial x} \right) \right) + \delta \lambda_2 \left(v - \left(\frac{\partial \psi}{\partial x} + \frac{\partial \chi}{\partial y} \right) \right) \right. \\ & + \delta u \left((u - \tilde{u} + \lambda_1) \right) + \delta v \left((v - \tilde{v} + \lambda_2) \right) \\ & \left. + \delta \psi \left(\frac{\partial \lambda_2}{\partial x} - \frac{\partial \lambda_1}{\partial y} \right) + \delta \chi \left(\frac{\partial \lambda_1}{\partial x} + \frac{\partial \lambda_2}{\partial y} \right) \right] dx dy. \quad (\text{D.12}) \end{aligned}$$

Moreover, by setting the coefficients of $\delta \lambda_1$ and $\delta \lambda_2$ equal to zero, we retrieve the Helmholtz constraint. Finally, setting the coefficients of $\delta \psi$ and $\delta \chi$ equal to zero, and using (D.6) and (D.7), we obtain

$$\begin{aligned} \frac{\partial \lambda_2}{\partial x} - \frac{\partial \lambda_1}{\partial y} &= 0 & \frac{\partial \lambda_1}{\partial x} + \frac{\partial \lambda_2}{\partial y} &= 0 \\ \Rightarrow \frac{\partial(\tilde{v} - v)}{\partial x} - \frac{\partial(\tilde{u} - u)}{\partial y} &= 0 & \Rightarrow \frac{\partial(\tilde{u} - u)}{\partial x} + \frac{\partial(\tilde{v} - v)}{\partial y} &= 0 \\ \Rightarrow -\frac{\partial \tilde{u}}{\partial y} + \frac{\partial \tilde{v}}{\partial x} &= -\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} & \Rightarrow \frac{\partial \tilde{u}}{\partial x} + \frac{\partial \tilde{v}}{\partial y} &= \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}. \end{aligned}$$

We observe that the terms on the left-hand side involving $\tilde{u}$ and $\tilde{v}$ correspond to the definitions of vorticity, $\tilde{\xi}$ (2.7), and divergence, $\tilde{D}$ (2.6). Since u and v satisfy the Helmholtz equations given by (2.4) and (2.5), we can also rewrite the right-hand side accordingly, to give

$$\tilde{\xi} = -\frac{\partial}{\partial y} \left(-\frac{\partial \psi}{\partial y} + \frac{\partial \chi}{\partial x} \right) + \frac{\partial}{\partial x} \left(\frac{\partial \psi}{\partial x} + \frac{\partial \chi}{\partial y} \right) \quad \tilde{D} = \frac{\partial}{\partial x} \left(-\frac{\partial \psi}{\partial y} + \frac{\partial \chi}{\partial x} \right) + \frac{\partial}{\partial y} \left(\frac{\partial \psi}{\partial x} + \frac{\partial \chi}{\partial y} \right).$$

We have derived the discrete Poisson equations, $\tilde{\xi} = \nabla^2 \psi$ and $\tilde{D} = \nabla^2 \chi$ and therefore must solve these to find u and v .

D.1.1 Tikhonov's regularisation

As discussed in section 2.2.1, Tikhonov's regularisation can be used to improve the convergence and conditioning of an inverse problem, detailed in section 4.2.3. We can use the Lagrangian approach to demonstrate how this constraint is implemented in our system. Tikhonov's regularisation adds an additional constraint to the cost function (D.1),

$$\mathcal{J}_T = \mathcal{J} + \frac{1}{2} \int_Y \int_X \mu_1 (\psi - \psi_b)^2 + \mu_2 (\chi - \chi_b)^2 \, dxdy. \quad (\text{D.13})$$

This form penalises deviations of the control variables, ψ and χ , from their respective background (or prior) values, ψ_b and χ_b . In the specific case discussed by Li et al. (2006), the regularisation term takes the form $\frac{1}{2} \alpha \mathbf{z}^T \mathbf{z}$, which penalises the squared amplitude of the control variable. This corresponds to choosing $\mathbf{z}_b = 0$, i.e., a prior value of zero. Following this approach, we set $\psi_b, \chi_b = 0$ which simplifies the regularised cost function to $\mathcal{J}_T = \mathcal{J} + \frac{1}{2} \int_Y \int_X \mu_1 \psi^2 + \mu_2 \chi^2 \, dxdy$. We then take the first order variation to give

$$\delta \mathcal{J}_T = \delta \mathcal{J} + \int_Y \int_X \mu_1 \psi \delta \psi + \mu_2 \chi \delta \chi \, dxdy.$$

We want to minimise this cost function with the constraints given previously. This leads to the following Poisson equations

$$\tilde{\xi} = \nabla^2 \psi - \mu_1 \psi \quad \text{and} \quad \tilde{D} = \nabla^2 \chi - \mu_2 \chi. \quad (\text{D.14})$$

D.2 Discrete equations

We now use the Lagrangian approach to derive the discrete Poisson equations. Recall figure 3.2b, we discretise on the Arakawa C-grid with streamfunction located at the centre of the grid cells, as proposed by Li et al. (2006). When discretising, we use masks to determine what are land and ocean points (this mask is defined at the centre of the grid cells) and to enforce boundary conditions on the velocities (this mask is defined on the boundaries of the grid cells). These are given by,

$$m = \begin{cases} 1 & \text{at ocean points} \\ 0 & \text{at land points} \end{cases} \quad (\text{D.15})$$

and

$$\epsilon_u, \epsilon_v = \begin{cases} 0 & \text{on the boundary} \\ 1 & \text{not on the boundary.} \end{cases} \quad (\text{D.16})$$

The boundary condition masks ϵ_u, ϵ_v enforce the physical boundary condition of zero normal velocity at the boundary. To simplify the notation, we have $m_{i,j}\psi_{i,j} = \psi'_{i,j}$ and $m_{i,j}\chi_{i,j} = \chi'_{i,j}$. The Helmholtz equations (2.4) and (2.5) can be discretised to give,

$$\begin{aligned} u_{i+\frac{1}{2},j} &= \epsilon_{u_{i+\frac{1}{2},j}}(-\Delta_y \psi'_{i+\frac{1}{2},j} + \Delta_x \chi'_{i+\frac{1}{2},j}) \\ &= \epsilon_{u_{i+\frac{1}{2},j}} \left[-\frac{1}{4\Delta y} (\psi'_{i+1,j+1} - \psi'_{i+1,j-1} + \psi'_{i,j+1} - \psi'_{i,j-1}) + \frac{1}{\Delta x} (\chi'_{i+1,j} - \chi'_{i,j}) \right] \end{aligned} \quad (\text{D.17})$$

and

$$\begin{aligned} v_{i,j+\frac{1}{2}} &= \epsilon_{v_{i,j+\frac{1}{2}}}(\Delta_x \psi'_{i,j+\frac{1}{2}} + \Delta_y \chi'_{i,j+\frac{1}{2}}) \\ &= \epsilon_{v_{i,j+\frac{1}{2}}} \left[\frac{1}{4\Delta x} (\psi'_{i+1,j+1} - \psi'_{i-1,j+1} + \psi'_{i+1,j} - \psi'_{i-1,j}) + \frac{1}{\Delta y} (\chi'_{i,j+1} - \chi'_{i,j}) \right]. \end{aligned} \quad (\text{D.18})$$

To derive the Poisson equations we define a Lagrange function similar to section D.1, as follows,

$$\begin{aligned} \hat{\mathcal{L}} &= \mathcal{J}(\mathbf{x}) + \sum_i \sum_j \left[\hat{\lambda}_{1_{i+\frac{1}{2},j}} \left(\epsilon_{u_{i+\frac{1}{2},j}} u_{i+\frac{1}{2},j} - \frac{\epsilon_{u_{i+\frac{1}{2},j}}}{4\Delta y} (\psi'_{i+1,j+1} - \psi'_{i+1,j-1} + \psi'_{i,j+1} - \psi'_{i,j-1}) \right. \right. \\ &\quad \left. \left. + \frac{\epsilon_{u_{i+\frac{1}{2},j}}}{\Delta x} (\chi'_{i+1,j} - \chi'_{i,j}) \right) \right. \\ &\quad \left. + \hat{\lambda}_{2_{i,j+\frac{1}{2}}} \left(\epsilon_{v_{i,j+\frac{1}{2}}} v_{i,j+\frac{1}{2}} - \frac{\epsilon_{v_{i,j+\frac{1}{2}}}}{4\Delta x} (\psi'_{i+1,j+1} - \psi'_{i-1,j+1} + \psi'_{i+1,j} - \psi'_{i-1,j}) + \frac{\epsilon_{v_{i,j+\frac{1}{2}}}}{\Delta y} (\chi'_{i,j+1} - \chi'_{i,j}) \right) \right] \end{aligned} \quad (\text{D.19})$$

where

$$\hat{\mathcal{J}} = \sum_i \sum_j \frac{1}{2} \left[\epsilon_{u_{i+\frac{1}{2},j}} \left(u_{i+\frac{1}{2},j} - \tilde{u}_{i+\frac{1}{2},j} \right)^2 + \epsilon_{v_{i,j+\frac{1}{2}}} \left(v_{i,j+\frac{1}{2}} - \tilde{v}_{i,j+\frac{1}{2}} \right)^2 \right] \quad (\text{D.20})$$

and $\hat{\lambda}_1, \hat{\lambda}_2$ are Lagrangian multipliers. Similar to the continuous case, we can obtain the Poisson equations by taking the derivative of the Lagrange function, $\hat{\mathcal{L}}$, with respect to all the variables and setting them equal to zero to find the stationary values. As such we have

$$\delta \hat{\mathcal{J}} = \sum_i \sum_j \left[\epsilon_{u_{i+\frac{1}{2},j}} (u_{i+\frac{1}{2},j} - \tilde{u}_{i+\frac{1}{2},j}) + \epsilon_{v_{i,j+\frac{1}{2}}} (v_{i,j+\frac{1}{2}} - \tilde{v}_{i,j+\frac{1}{2}}) \right]. \quad (\text{D.21})$$

Firstly, setting the derivatives with respect to the Lagrangian multipliers, $\partial \hat{\mathcal{L}} / \partial \hat{\lambda}_1$ and $\partial \hat{\mathcal{L}} / \partial \hat{\lambda}_2$, equal to zero gives the Helmholtz constraint. We now set $\epsilon_{u_{i+\frac{1}{2},j}} \hat{\lambda}_{1_{i+\frac{1}{2},j}} = \hat{\lambda}'_{1_{i+\frac{1}{2},j}}$,

$\epsilon_{v_{i,j+\frac{1}{2}}} \hat{\lambda}_{2_{i,j+\frac{1}{2}}} = \hat{\lambda}'_{2_{i,j+\frac{1}{2}}}$, $\epsilon_{u_{i+\frac{1}{2},j}} u_{i+\frac{1}{2},j} = u'_{i+\frac{1}{2},j}$, $\epsilon_{v_{i,j+\frac{1}{2}}} v_{i,j+\frac{1}{2}} = v'_{i,j+\frac{1}{2}}$ for simplicity and calculate derivatives with respect to the velocities,

$$\frac{\partial \hat{\mathcal{L}}}{\partial u_{i+\frac{1}{2},j}} = \frac{\partial \hat{\mathcal{J}}}{\partial u_{i+\frac{1}{2},j}} + \hat{\lambda}'_{1_{i+\frac{1}{2},j}} = u'_{i+\frac{1}{2},j} - \tilde{u}'_{i+\frac{1}{2},j} + \hat{\lambda}'_{1_{i+\frac{1}{2},j}}$$

$$\frac{\partial \hat{\mathcal{L}}}{\partial v_{i,j+\frac{1}{2}}} = \frac{\partial \hat{\mathcal{J}}}{\partial v_{i,j+\frac{1}{2}}} + \hat{\lambda}'_{2_{i,j+\frac{1}{2}}} = v'_{i,j+\frac{1}{2}} - \tilde{v}'_{i,j+\frac{1}{2}} + \hat{\lambda}'_{2_{i,j+\frac{1}{2}}}.$$

We equate these to zero to give the following,

$$\frac{\partial \hat{\mathcal{L}}}{\partial \delta u_{i+\frac{1}{2},j}} = 0 \Rightarrow \hat{\lambda}'_{1_{i+\frac{1}{2},j}} = \tilde{u}'_{i+\frac{1}{2},j} - u'_{i+\frac{1}{2},j} \quad (\text{D.22})$$

and

$$\frac{\partial \hat{\mathcal{L}}}{\partial \delta v_{i,j+\frac{1}{2}}} = 0 \Rightarrow \hat{\lambda}'_{2_{i,j+\frac{1}{2}}} = \tilde{v}'_{i,j+\frac{1}{2}} - v'_{i,j+\frac{1}{2}}. \quad (\text{D.23})$$

We finally calculate the derivatives of $\hat{\mathcal{L}}$ with respect to streamfunction and velocity potential,

$$\begin{aligned} \frac{\partial \hat{\mathcal{L}}}{\partial \psi_{i,j}} = \frac{1}{4\Delta y} & \left(\hat{\lambda}'_{1_{i-\frac{1}{2},j-1}} - \hat{\lambda}'_{1_{i-\frac{1}{2},j+1}} + \hat{\lambda}'_{1_{i+\frac{1}{2},j-1}} - \hat{\lambda}'_{1_{i+\frac{1}{2},j+1}} \right) \\ & - \frac{1}{4\Delta x} \left(\hat{\lambda}'_{2_{i-1,j-\frac{1}{2}}} - \hat{\lambda}'_{2_{i+1,j-\frac{1}{2}}} + \hat{\lambda}'_{2_{i-1,j+\frac{1}{2}}} - \hat{\lambda}'_{2_{i+1,j+\frac{1}{2}}} \right) \end{aligned} \quad (\text{D.24})$$

$$\frac{\partial \hat{\mathcal{L}}}{\partial \chi_{i,j}} = -\frac{1}{\Delta x} \left(\hat{\lambda}'_{1_{i-\frac{1}{2},j}} - \hat{\lambda}'_{1_{i+\frac{1}{2},j}} \right) - \frac{1}{\Delta y} \left(\hat{\lambda}'_{2_{i,j-\frac{1}{2}}} - \hat{\lambda}'_{2_{i,j+\frac{1}{2}}} \right). \quad (\text{D.25})$$

By equating (D.24) and (D.25) to zero and substituting in the expressions for the Lagrangian multipliers, (D.22) and (D.23), we obtain

$$\begin{aligned} & \frac{1}{4\Delta y} \left[\tilde{u}'_{i-\frac{1}{2},j-1} - \tilde{u}'_{i-\frac{1}{2},j+1} + \tilde{u}'_{i+\frac{1}{2},j-1} - \tilde{u}'_{i+\frac{1}{2},j+1} - \left(u'_{i-\frac{1}{2},j-1} - u'_{i-\frac{1}{2},j+1} + u'_{i+\frac{1}{2},j-1} - u'_{i+\frac{1}{2},j+1} \right) \right] \\ & = \frac{1}{4\Delta x} \left[\tilde{v}'_{i-1,j-\frac{1}{2}} - \tilde{v}'_{i+1,j-\frac{1}{2}} + \tilde{v}'_{i-1,j+\frac{1}{2}} - \tilde{v}'_{i+1,j+\frac{1}{2}} - \left(v'_{i-1,j-\frac{1}{2}} - v'_{i+1,j-\frac{1}{2}} + v'_{i-1,j+\frac{1}{2}} - v'_{i+1,j+\frac{1}{2}} \right) \right] \end{aligned} \quad (\text{D.26})$$

and

$$\frac{1}{\Delta y} \left[\tilde{v}'_{i,j-\frac{1}{2}} - \tilde{v}'_{i,j+\frac{1}{2}} - \left(v'_{i,j-\frac{1}{2}} - v'_{i,j+\frac{1}{2}} \right) \right] = -\frac{1}{\Delta x} \left[\tilde{u}'_{i-\frac{1}{2},j} - \tilde{u}'_{i+\frac{1}{2},j} - \left(u'_{i-\frac{1}{2},j} - u'_{i+\frac{1}{2},j} \right) \right]. \quad (\text{D.27})$$

From (D.27) we obtain

$$\underbrace{\frac{\tilde{u}'_{i+\frac{1}{2},j} - \tilde{u}'_{i-\frac{1}{2},j}}{\Delta x} + \frac{\tilde{v}'_{i,j+\frac{1}{2}} - \tilde{v}'_{i,j-\frac{1}{2}}}{\Delta y}}_{\hat{D}_{i,j}} = \frac{u'_{i+\frac{1}{2},j} - u'_{i-\frac{1}{2},j}}{\Delta x} + \frac{v'_{i,j+\frac{1}{2}} - v'_{i,j-\frac{1}{2}}}{\Delta y} \quad (\text{D.28})$$

where $\hat{D}_{i,j}$ is the discrete divergence. From (D.26) we obtain,

$$\begin{aligned}
& \underbrace{-\frac{1}{4\Delta y}(\tilde{u}'_{i-\frac{1}{2},j+1} - \tilde{u}'_{i-\frac{1}{2},j-1} + \tilde{u}'_{i+\frac{1}{2},j+1} - \tilde{u}'_{i+\frac{1}{2},j-1}) + \frac{1}{4\Delta x}(\tilde{v}'_{i+1,j-\frac{1}{2}} - \tilde{v}'_{i-1,j-\frac{1}{2}} + \tilde{v}'_{i+1,j+\frac{1}{2}} - \tilde{v}'_{i-1,j+\frac{1}{2}})}_{\hat{\xi}_{i,j}} \\
&= -\frac{1}{4\Delta y}(u'_{i-\frac{1}{2},j+1} - u'_{i-\frac{1}{2},j-1} + u'_{i+\frac{1}{2},j+1} - u'_{i+\frac{1}{2},j-1}) + \frac{1}{4\Delta x}(v'_{i+1,j-\frac{1}{2}} - v'_{i-1,j-\frac{1}{2}} + v'_{i+1,j+\frac{1}{2}} - v'_{i-1,j+\frac{1}{2}})
\end{aligned} \tag{D.29}$$

where $\hat{\xi}_{i,j}$ is the discrete vorticity. We firstly examine the equation for divergence. From (D.28) and using the discrete Helmholtz equations ((D.17) and (D.18)) we have

$$\begin{aligned}
\hat{D}_{i,j} &= \frac{u'_{i+\frac{1}{2},j} - u'_{i-\frac{1}{2},j}}{\Delta x} + \frac{v'_{i,j+\frac{1}{2}} - v'_{i,j-\frac{1}{2}}}{\Delta y} \\
&= \frac{1}{\Delta x} \left[-\frac{1}{4\Delta y}(\psi'_{i+1,j+1} - \psi'_{i+1,j-1} + \psi'_{i,j+1} - \psi'_{i,j-1}) + \frac{1}{\Delta x}(\chi'_{i+1,j} - \chi'_{i,j}) \right. \\
&\quad \left. + \frac{1}{4\Delta y}(\psi'_{i,j+1} - \psi'_{i,j-1} + \psi'_{i-1,j+1} - \psi'_{i-1,j-1}) + \frac{1}{\Delta x}(\chi'_{i,j} - \chi'_{i-1,j}) \right] \\
&\quad + \frac{1}{\Delta y} \left[-\frac{1}{4\Delta x}(\psi'_{i+1,j+1} - \psi'_{i-1,j+1} + \psi'_{i+1,j} - \psi'_{i-1,j}) + \frac{1}{\Delta x}(\chi'_{i,j+1} - \chi'_{i,j}) \right. \\
&\quad \left. + \frac{1}{4\Delta y}(\psi'_{i+1,j} - \psi'_{i-1,j} + \psi'_{i+1,j-1} - \psi'_{i-1,j-1}) + \frac{1}{\Delta x}(\chi'_{i,j} - \chi'_{i,j-1}) \right].
\end{aligned}$$

We combine the ψ' and χ' terms above. As such the ψ' terms fully cancel out, and we are left with,

$$\hat{D}_{i,j} = \frac{1}{\Delta x^2}(\chi'_{i+1,j} - 2\chi'_{i,j} + \chi'_{i-1,j}) + \frac{1}{\Delta y^2}(\chi'_{i,j+1} - 2\chi'_{i,j} + \chi'_{i,j-1}) = \Delta_x^2 \chi'_{i,j} + \Delta_y^2 \chi'_{i,j}$$

where Δ_x^2 and Δ_y^2 are the discretisation of $\frac{\partial^2}{\partial x^2}$ and $\frac{\partial^2}{\partial y^2}$. This gives,

$$\hat{D}_{i,j} = \mathbf{L}\chi'_{i,j} \tag{D.30}$$

where $\mathbf{L}$ is the discrete Laplacian operator. This is the standard discrete Laplacian operator using centred differencing. We now show that the same cannot be said for

vorticity. From (D.29) we have,

$$\begin{aligned}
\hat{\xi}_{i,j} = & -\frac{1}{4\Delta y} \left[-\frac{1}{4\Delta y} \left(\psi'_{i,j+2} - \psi'_{i,j} + \psi'_{i-1,j+2} - \psi'_{i-1,j} - \psi'_{i,j} + \psi'_{i,j-2} - \psi'_{i-1,j} + \psi'_{i-1,j-2} \right. \right. \\
& + \psi'_{i+1,j+2} - \psi'_{i+1,j} + \psi'_{i,j+2} - \psi'_{i,j} - \psi'_{i+1,j} + \psi'_{i+1,j-2} - \psi'_{i,j} + \psi'_{i,j-2} \left. \right) \\
& + \frac{1}{\Delta x} \left(\chi'_{i,j+1} - \chi'_{i-1,j+1} - \chi'_{i,j-1} + \chi'_{i+1,j+1} - \chi'_{i,j+1} - \chi'_{i+1,j-1} + \chi'_{i,j-1} \right) \left. \right] \\
& + \frac{1}{4\Delta x} \left[\frac{1}{4\Delta x} \left(\psi'_{i+2,j} - \psi'_{i,j} + \psi'_{i+2,j-1} - \psi'_{i,j-1} - \psi'_{i,j} + \psi'_{i-2,j} - \psi'_{i,j-1} + \psi'_{i-2,j-1} \right. \right. \\
& + \psi'_{i+2,j+1} - \psi'_{i,j+1} + \psi'_{i+2,j} - \psi'_{i,j} - \psi'_{i+1,j} + \psi'_{i-2,j+1} - \psi'_{i,j} + \psi'_{i-2,j-1} \left. \right) \\
& + \frac{1}{\Delta y} \left(\chi'_{i+1,j} - \chi'_{i+1,j-1} - \chi'_{i-1,j} + \chi'_{i-1,j-1} + \chi'_{i+1,j+1} - \chi'_{i+1,j} - \chi'_{i-1,j+1} + \chi'_{i-1,j} \right) \left. \right].
\end{aligned}$$

The χ' subsequently cancel, leaving the ψ' terms only. Upon further examination, it is revealed that the discretised streamfunction terms, $\Delta_y \psi$ and $\Delta_x \psi$, from the discrete Helmholtz decomposition (D.17) and (D.18) are present, this reduces the above to

$$\begin{aligned}
\hat{\xi}_{i,j} = & \frac{1}{4\Delta y} \left(\Delta_y \psi'_{i-\frac{1}{2},j+1} - \Delta_y \psi'_{i-\frac{1}{2},j-1} + \Delta_y \psi'_{i+\frac{1}{2},j+1} - \Delta_y \psi'_{i+\frac{1}{2},j-1} \right) \\
& + \frac{1}{4\Delta x} \left(\Delta_x \psi'_{i+1,j-\frac{1}{2}} - \Delta_x \psi'_{i-1,j-\frac{1}{2}} + \Delta_x \psi'_{i+1,j+\frac{1}{2}} - \Delta_x \psi'_{i-1,j+\frac{1}{2}} \right) = \bar{\Delta}_x^2 \psi'_{i,j} + \bar{\Delta}_y^2 \psi'_{i,j} \\
\Rightarrow \hat{\xi}_{i,j} = & \bar{\mathbf{L}} \psi'_{i,j}
\end{aligned} \tag{D.31}$$

where $\bar{\Delta}_x^2$ are $\bar{\Delta}_y^2$ and $\bar{\mathbf{L}}$ are non-standard discrete operators caused by the placement of streamfunction at the centre of the grid. Therefore, when using the the two-process method described in section C.2, care must be taken when implementing these non-standard operators. Additionally, it is important to note that the Lagrangian formulation yields an expression for the vorticity evaluated at the centre of the grid cell — a consequence that is consistent with the placement of the streamfunction at the same location. However, on an Arakawa C-grid vorticity is usually located at the corner of the grid-cell (Madec et al., 2023), the more common location of streamfunction in literature (Lynch, 1989; Watterson, 2001). Accordingly, the vorticity should be computed as follows,

$$\xi_{i+\frac{1}{2},j+\frac{1}{2}} = \Delta_y^2 \psi'_{i+\frac{1}{2},j+\frac{1}{2}} + \Delta_x^2 \psi'_{i+\frac{1}{2},j+\frac{1}{2}} \tag{D.32}$$

where

$$\Delta_y^2 \psi'_{i+\frac{1}{2},j+\frac{1}{2}} = \frac{1}{\Delta y} \left(\Delta_y \psi'_{i+\frac{1}{2},j+1} - \Delta_y \psi'_{i+\frac{1}{2},j} \right) = \frac{1}{\Delta y} \left[\frac{1}{4\Delta y} \left(\psi'_{i,j+2} - \psi'_{i,j} + \psi'_{i+1,j+2} - \psi'_{i+1,j} \right) - \frac{1}{4\Delta y} \left(\psi'_{i,j+1} - \psi'_{i,j-1} + \psi'_{i+1,j+1} - \psi'_{i+1,j-1} \right) \right]$$

and

$$\Delta_x^2 \psi'_{i+\frac{1}{2},j+\frac{1}{2}} = \frac{1}{\Delta x} \left(\Delta_x \psi'_{i+1,j+\frac{1}{2}} - \Delta_x \psi'_{i,j+\frac{1}{2}} \right) = \frac{1}{\Delta x} \left[\frac{1}{4\Delta x} \left(\psi'_{i+2,j+1} + \psi'_{i+2,j} - \psi'_{i,j+1} - \psi'_{i,j} \right) - \frac{1}{4\Delta x} \left(\psi'_{i+1,j+1} + \psi'_{i+1,j} - \psi'_{i-1,j+1} - \psi'_{i-1,j} \right) \right].$$

This approach can be most clearly demonstrated using figure D.1. We firstly calculate the streamfunction derivatives at the necessary locations, given by the red and yellow dots, as described in section 4.3.2 (this involves finding the derivatives and averaging them to the correct velocity locations). The second derivatives are then calculated at the corner points of the grid, given by the green dots, where vorticity is located.

Hence, when applying this two-system approach of solving the Poisson equations, to find ageostrophic streamfunction and velocity potential, special care must be taken to use the correct Laplacian operator. The intuitive or "natural" method of implementing the Laplacian operator is not numerically equivalent and should therefore be avoided.




Figure D.1: A visual demonstration of the Laplacian operator used to solve (C.8) using the streamfunction location proposed by Li et al. (2006) given by the blue dot. The streamfunction derivatives are located at red and yellow dots, these are calculated as described in section 4.3.2 (this involved finding the derivatives and averaging them to the correct velocity locations). The second derivatives are then calculated at the corner points of the grid, given by the green dots, where vorticity is located.

Appendix E

Testing of the adjoint and gradient

When performing the transformation to ageostrophic streamfunction and velocity potential, in chapter 4, there are occasions where we calculate the adjoint and the gradient of an operator. This section describes the validation tests for these calculations.

E.1 The adjoint test

The adjoint test is used to verify whether an adjoint operator has been implemented correctly (Li et al., 1994; Lawless et al., 2003). Let $\mathbf{M}$ be a linear operator and $\mathbf{M}^T$ its adjoint. To test this, we generate a random perturbation $\delta\mathbf{x}$ and apply the operator $\mathbf{M}$ to obtain $\mathbf{M}\delta\mathbf{x}$. We apply the adjoint to this result to give, $\mathbf{M}^T\mathbf{M}\delta\mathbf{x}$. To check the validity of the adjoint we verify the following identity holds,

$$\langle \mathbf{M}\delta\mathbf{x}, \mathbf{M}\delta\mathbf{x} \rangle = \langle \delta\mathbf{x}, \mathbf{M}^T\mathbf{M}\delta\mathbf{x} \rangle. \quad (\text{E.1})$$

This can be verified by examining the difference between the dot products and checking if it is sufficiently close to zero, within the limits of machine precision.

E.2 The gradient test

The gradient test verifies the accuracy of a function's gradient calculation (Navon et al., 1992; Li et al., 1994). Suppose we have a function, $J(\mathbf{x})$ and

$$J(\mathbf{x} + \alpha\mathbf{h}) = J(\mathbf{x}) + \alpha\mathbf{h}^T\nabla J(\mathbf{x}) + O(\alpha^2) \quad (\text{E.2})$$

where α is a scalar and $\mathbf{h}$ is of unit length, i.e.

$$\mathbf{h} = \frac{\nabla J(\mathbf{x})}{\|\nabla J(\mathbf{x})\|}.$$

We define,

$$\Phi(\alpha) = \frac{|J(\mathbf{x} + \alpha\mathbf{h}) - J(\mathbf{x})|}{\alpha\mathbf{h}^T\nabla J(\mathbf{x})} = 1 + O(\alpha). \quad (\text{E.3})$$

The gradient test involves plotting the value of $\Phi(\alpha)$ as α approaches zero. If the values are consistent, i.e. they are close to 1 as expected from (E.3), then the gradient calculation is correct. To ensure that the plot values do not deviate significantly from 1, we also plot $1 - \Phi(\alpha)$ for verification.

Appendix F

Non-divergence of the velocity generated by streamfunction

As discussed in section 2.1, Helmholtz theorem states that a velocity field can be split into its non-divergent and irrotational parts. Notably, the velocity generated by streamfunction should be non-divergent i.e. $\nabla \cdot \mathbf{v}_\psi = 0$. Here we investigate this statement and expose the intricacies of coding these equations. Note we deal with the full model fields in this chapter, rather than the ageostrophic increments. The result would be the same analytically.

We discretise the divergence of the velocity generated by streamfunction,

$$\mathcal{D}_\psi = \nabla \cdot \mathbf{v}_\psi = \frac{\partial}{\partial x} \left(\frac{-\partial \psi}{\partial y} \right) + \frac{\partial}{\partial y} \left(\frac{\partial \psi}{\partial x} \right) \quad (\text{F.1})$$

on a uniform grid to give

$$\hat{\mathcal{D}}_\psi = \Delta_x \left(-\Delta_y \psi \right) + \Delta_y \left(\Delta_x \psi \right). \quad (\text{F.2})$$

Recall figure 3.2b, we discretise on the Arakawa C-grid, with streamfunction located at the centre of the grid cell, as proposed by Li et al. (2006). When discretising we use masks as described in section D.2. The land mask m is given by (D.15) and the boundary condition masks ϵ_u, ϵ_v (given by D.16) enforce the physical boundary condition of zero normal velocity at the boundary. Therefore at a given location on the grid we have,

$$\hat{\mathcal{D}}_{\psi_{i,j}} = m_{i,j} \left(\Delta_x \left(-\Delta_y \psi \right) \right)_{i,j} + m_{i,j} \left(\Delta_y \left(\Delta_x \psi \right) \right)_{i,j}. \quad (\text{F.3})$$

In section 4.2.4 we discussed two different approaches for calculating the streamfunction derivatives and how they are equivalent on a uniform grid. However, these approaches are not equivalent when on a non-uniform grid. This is also linked to the masking used. Therefore, in NEMOVAR we use approach B. In this section we use approach B to investigate the divergence of the velocity generated by streamfunction. This involves averaging the values of streamfunction before calculating derivatives, as follows,

1. Average ψ to the corner of the grid cells, as given by $\bar{\psi}$ in figure 4.3. The land mask is applied to the values of ψ before doing so.
2. To calculate $\left(\Delta_x\left(-\Delta_y\psi\right)\right)_{i,j}$,
 - (a) Find the value of $(\Delta_y\psi)$ at the u-point by calculating the y-derivative of $\bar{\psi}$ using centred differencing.
 - (b) Apply the boundary mask to these u-points.
 - (c) Find the value of $\Delta_x\left(-\Delta_y\psi\right)$ at the centre of the grid cell by calculating the x-derivative of $(\Delta_y\psi)$ using centred differencing.
3. To calculate $\left(\Delta_y\left(\Delta_x\psi\right)\right)_{i,j}$,
 - (a) Find the value of $(\Delta_x\psi)$ at the v-point by calculating the x-derivative of $\bar{\psi}$ using centred differencing.
 - (b) Apply the boundary mask to these v-points.
 - (c) Find the value of $\Delta_y\left(\Delta_x\psi\right)$ at the centre of the grid cell by calculating the y-derivative of $(\Delta_x\psi)$ using centred differencing.

This is the process we implement in NEMOVAR to calculate this divergence. The corner value of streamfunction is given as

$$\bar{\psi}_{i+\frac{1}{2},j+\frac{1}{2}} = \frac{1}{M} \left(m_{i,j+1}\psi_{i,j+1} + m_{i,j}\psi_{i,j} + m_{i+1,j+1}\psi_{i+1,j+1} + m_{i+1,j}\psi_{i+1,j} \right) \quad (\text{F.4})$$

with $M = \max(1, m_{i,j+1} + m_{i,j} + m_{i+1,j+1} + m_{i+1,j})$ is the number of surrounding ocean

points. From the process described above we have,

$$\begin{aligned} \left(\Delta_x (-\Delta_y \psi) \right)_{i,j} &= -\frac{1}{\Delta x} \left[\epsilon_{u_{i+\frac{1}{2},j}} (\Delta_y \psi)_{i+\frac{1}{2},j} - \epsilon_{u_{i-\frac{1}{2},j}} (\Delta_y \psi)_{i-\frac{1}{2},j} \right] \\ &= -\frac{1}{\Delta x} \left[\frac{\epsilon_{u_{i+\frac{1}{2},j}}}{\Delta y} \left(\psi'_{i+\frac{1}{2},j+\frac{1}{2}} - \psi'_{i+\frac{1}{2},j-\frac{1}{2}} \right) - \frac{\epsilon_{u_{i-\frac{1}{2},j}}}{\Delta y} \left(\bar{\psi}_{i-\frac{1}{2},j+\frac{1}{2}} - \bar{\psi}_{i-\frac{1}{2},j-\frac{1}{2}} \right) \right] \end{aligned} \quad (\text{F.5})$$

and

$$\begin{aligned} \left(\Delta_y (\Delta_x \psi) \right)_{i,j} &= \frac{1}{\Delta y} \left[\epsilon_{v_{i,j+\frac{1}{2}}} (\Delta_x \psi)_{i,j+\frac{1}{2}} - \epsilon_{v_{i,j-\frac{1}{2}}} (\Delta_x \psi)_{i,j-\frac{1}{2}} \right] \\ &= \frac{1}{\Delta y} \left[\frac{\epsilon_{v_{i,j+\frac{1}{2}}}}{\Delta x} \left(\bar{\psi}_{i+\frac{1}{2},j+\frac{1}{2}} - \bar{\psi}_{i-\frac{1}{2},j+\frac{1}{2}} \right) - \frac{\epsilon_{v_{i,j-\frac{1}{2}}}}{\Delta x} \left(\bar{\psi}_{i+\frac{1}{2},j-\frac{1}{2}} - \bar{\psi}_{i-\frac{1}{2},j-\frac{1}{2}} \right) \right]. \end{aligned} \quad (\text{F.6})$$

F.1 Away from the boundaries

Away from the boundaries, the masks will be set to $\epsilon_u, \epsilon_v = 1$ and $m = 1$. This gives

$$\bar{\psi}_{i+\frac{1}{2},j+\frac{1}{2}} = \frac{1}{4} \left(\psi_{i,j+1} + \psi_{i,j} + \psi_{i+1,j+1} + \psi_{i+1,j} \right).$$

We obtain

$$\left(\Delta_x (-\Delta_y \psi) \right)_{i,j} = -\frac{1}{4\Delta x \Delta y} \left((\Delta_y \psi)_{i+1,j+\frac{1}{2}} + (\Delta_y \psi)_{i+1,j-\frac{1}{2}} + (\Delta_y \psi)_{i-1,j+\frac{1}{2}} + (\Delta_y \psi)_{i-1,j-\frac{1}{2}} \right)$$

and

$$\left(\Delta_y (\Delta_x \psi) \right)_{i,j} = \frac{1}{4\Delta x \Delta y} \left((\Delta_x \psi)_{i-\frac{1}{2},j+1} + (\Delta_x \psi)_{i+\frac{1}{2},j+1} + (\Delta_x \psi)_{i-\frac{1}{2},j-1} + (\Delta_x \psi)_{i+\frac{1}{2},j-1} \right).$$

Adding these two terms we get $\hat{\mathcal{D}}_{\psi_{i,j}} = 0$, meaning non-divergence is preserved away from the boundary.

F.2 Near the boundaries

We now calculate the value of $\hat{\mathcal{D}}_{\psi_{i,j}}$ for the case shown in figure 4.4, at a point located near the eastern boundary. We have,

$$\begin{aligned} \left(\Delta_x (-\Delta_y \psi) \right)_{i,j} &= -\frac{1}{\Delta x} \left[0 \times (\Delta_y \psi)_{i+\frac{1}{2},j} - 1 \times (\Delta_y \psi)_{i-\frac{1}{2},j} \right] \\ &= \frac{1}{\Delta x \Delta y} \left[\left(\bar{\psi}_{i-\frac{1}{2},j+\frac{1}{2}} - \bar{\psi}_{i-\frac{1}{2},j-\frac{1}{2}} \right) \right] \end{aligned} \quad (\text{F.7})$$

and

$$\begin{aligned} \left(\Delta_y \left(\Delta_x \psi \right) \right)_{i,j} &= \frac{1}{\Delta y} \left[1 \times (\Delta_x \psi)_{i,j+\frac{1}{2}} - 1 \times (\Delta_x \psi)_{i,j-\frac{1}{2}} \right] \\ &= \frac{1}{\Delta x \Delta y} \left[\left(\bar{\psi}_{i+\frac{1}{2},j+\frac{1}{2}} - \bar{\psi}_{i-\frac{1}{2},j+\frac{1}{2}} \right) - \left(\bar{\psi}_{i+\frac{1}{2},j-\frac{1}{2}} - \bar{\psi}_{i-\frac{1}{2},j-\frac{1}{2}} \right) \right] \end{aligned} \quad (\text{F.8})$$

where

$$\begin{aligned} \bar{\psi}_{i-\frac{1}{2},j+\frac{1}{2}} &= \frac{1}{4} \left(\psi_{i-1,j+1} + \psi_{i-1,j} + \psi_{i,j+1} + \psi_{i,j} \right) \\ \bar{\psi}_{i-\frac{1}{2},j-\frac{1}{2}} &= \frac{1}{4} \left(\psi_{i-1,j} + \psi_{i,j} + \psi_{i-1,j-1} + \psi_{i,j-1} \right) \\ \bar{\psi}_{i+\frac{1}{2},j+\frac{1}{2}} &= \frac{1}{2} \left(\psi_{i,j+1} + \psi_{i,j} \right) \\ \bar{\psi}_{i+\frac{1}{2},j-\frac{1}{2}} &= \frac{1}{2} \left(\psi_{i,j} + \psi_{i,j-1} \right). \end{aligned}$$

Therefore,

$$\left(\Delta_x (-\Delta_y \psi) \right)_{i,j} = \frac{1}{4\Delta x \Delta y} \left[\psi_{i-1,j+1} - \psi_{i-1,j-1} + \psi_{i,j+1} - \psi_{i,j-1} \right]$$

and

$$\left(\Delta_y \left(\Delta_x \psi \right) \right)_{i,j} = \frac{1}{4\Delta x \Delta y} \left(\psi_{i,j+1} - \psi_{i-1,j+1} - \psi_{i,j-1} + \psi_{i-1,j-1} \right).$$

These combine to give

$$\begin{aligned} \hat{\mathcal{D}}_{\psi_{i,j}} &= 1 \times \left(\Delta_x (-\Delta_y \psi) \right)_{i,j} + 1 \times \left(\Delta_y \left(\Delta_x \psi \right) \right)_{i,j} \\ &= \frac{1}{4\Delta x \Delta y} \left(\psi_{i-1,j+1} - \psi_{i-1,j-1} + \psi_{i,j+1} - \psi_{i,j-1} \right) \\ &\quad + \frac{1}{4\Delta x \Delta y} \left(\psi_{i,j+1} - \psi_{i-1,j+1} - \psi_{i,j-1} + \psi_{i-1,j-1} \right) \\ &= \frac{1}{2\Delta x \Delta y} \left(\psi_{i,j+1} - \psi_{i,j-1} \right) \\ &\neq 0. \end{aligned}$$

We have shown that the streamfunction induced velocity is divergent near the boundaries with how the code is currently setup in NEMOVAR. However, it is important to point out a conflict we have made in this method of masking. In (F.5) we have applied the mask ϵ_u , which enforces the value of zero at the boundary, to the y-derivatives of streamfunction. The logic behind this is that we must apply some form of masking in

NEMOVAR to a value at the boundary. However, by doing so we have assumed that the y-derivative of streamfunction has the same boundary condition as the normal velocity. Essentially enforcing $\Delta_y \psi = 0$ on the eastern and western boundaries. This implies that streamfunction is constant along these boundaries, which would further imply that we should re-arrange (F.8) to give

$$\frac{1}{\Delta x \Delta y} \left[\underbrace{\left(\bar{\psi}_{i+\frac{1}{2}, j+\frac{1}{2}} - \bar{\psi}_{i+\frac{1}{2}, j-\frac{1}{2}} \right)}_{=0} - \left(\bar{\psi}_{i-\frac{1}{2}, j+\frac{1}{2}} - \bar{\psi}_{i-\frac{1}{2}, j-\frac{1}{2}} \right) \right]. \quad (\text{F.9})$$

The first term above is equal to zero as we state that ψ is constant along the boundary. As such, when we combine (F.7) and (F.9) we obtain $\hat{\mathcal{D}}_{\psi, j} = 0$. Therefore, we do in fact preserve non-divergence with the Li et al. (2006) approach, the issue of divergence at the boundary arises due to the discretisation and masking applied in NEMOVAR.

This setting of the boundary conditions — where the streamfunction is held constant along the boundary — is undesirable. As noted by Li et al. (2006), Watterson (2001) assume a constant streamfunction along the coastline, which implies a non-divergent flow—an assumption that breaks down in regions with significant Ekman transport. This discretisation implemented in NEMOVAR does introduce some complications as mentioned. However, it is important to emphasise that in NEMOVAR, the streamfunction derivative at the boundary is not currently used independently. Once it is combined with the velocity potential derivative, the resulting velocity field is subjected to a boundary condition that enforces zero velocity, rendering the masking of derivatives effectively redundant. The Li et al. (2006) approach still ensures we do not enforce any explicit values of streamfunction on the boundary. We have demonstrated the above using a uniform grid for simplicity of notation; however, the remarks extend to a non-uniform grid, as used in NEMOVAR.


In summary, while this issue of masking the streamfunction derivatives should be carefully considered in future developments and applications of this discretisation and masking approach in NEMOVAR, it does not presently impact the assimilation process with the new control variables.

Appendix G


FFT boundary conditions

In chapter 5, we utilise the Fast-Fourier Transform (FFT) to examine the impact of the checkerboard pattern. The Fourier transform assumes periodic boundary conditions, but in section 5.4, we apply the FFT to 1D and 2D fields that are non-periodic. To investigate the implications of this, we analyse two simple waveforms. The first, illustrated in figure G.1a, is a full sine wave (over 2π), which is periodic. The second, shown in figure G.1b, is a partial sine wave that is non-periodic. Fourier spectra are displayed on the right of both waves.

Interestingly, in figure G.1b, despite the non-periodic nature, the assumption of periodic boundary conditions does not significantly distort the general shape and magnitudes of the frequencies compared to figure G.1a. However, there is a small amount of low-frequency power present that should not be there. Although this result may depend on the input field, a similar result is also discussed in Baxter et al. (2011). Therefore, in the Fourier analysis presented in section 5.4, it is important to be mindful of the subtle distortion introduced in the high-frequency components due to the non-periodic nature of the input fields.



(a) Left: Sine wave. Right: Spectrum of the wave calculated using the FFT.



(b) Left: Partial sine wave. Right: Spectrum of the wave calculated using the FFT.

Figure G.1: FFT of periodic and non-periodic signals.

Appendix H

No velocity observations

This thesis aims to prepare for the assimilation of future ocean current measurements, by improving the velocity variables that are used in ocean data assimilation. However, does changing the velocity variables have any effect on the analysis if there are no velocity observations? To answer this question, we investigate the effect of changing the velocity variables with only a single height observation in our simple shallow water model, described in section 4.1.

Recall the variational data assimilation cost function with the control variable transform,

$$\mathcal{J}(\mathbf{z}) = \frac{1}{2}\mathbf{z}^T\mathbf{z} + \frac{1}{2}\left(\mathbf{y} - \mathbf{h}(\mathcal{U}\mathbf{z})\right)^T \mathbf{R}^{-1}\left(\mathbf{y} - \mathbf{h}(\mathcal{U}\mathbf{z})\right) \quad (\text{H.1})$$

with $\mathbf{x} = \mathcal{U}\mathbf{z}$ where $\mathbf{x}$ is the model variable, $\mathcal{U}$ is the U-transform and $\mathbf{z}$ is the control variable, as described in section 3.3. The gradient of the cost function is given by

$$\nabla \mathcal{J}(\mathbf{z}) = \mathbf{z} + \mathcal{U}^T \mathbf{H}^T \mathbf{R}^{-1} \left(\mathbf{y} - \mathbf{h}(\mathcal{U}\mathbf{z}) \right). \quad (\text{H.2})$$

Here, we investigate whether the gradient of the cost function changes with respect to the new control variables, given no velocity observations. In the shallow water model we have $\mathbf{x} = (\eta, u, v)^T$. Suppose we have only a single observation of height, η_o , with $\mathbf{R} = \sigma_\eta$ and $\mathbf{h} = \mathbf{H} = (1, 0, 0)$.

To begin, we examine the gradient of the cost function when using the velocity control variables currently used in NEMOVAR, $\mathbf{z} = (\eta, u_U, v_U)^T$, the ageostrophic velocities.

Substituting our single observation into (H.2) we have,

$$\begin{aligned} \begin{pmatrix} \frac{\partial \mathcal{J}}{\partial \eta} \\ \frac{\partial \mathcal{J}}{\partial u_U} \\ \frac{\partial \mathcal{J}}{\partial v_U} \end{pmatrix} &= \begin{pmatrix} \eta \\ u_U \\ v_U \end{pmatrix} + \begin{pmatrix} 1 & -\frac{g}{f} \left(\frac{\partial}{\partial y} \right)^T & \frac{g}{f} \left(\frac{\partial}{\partial x} \right)^T \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \frac{1}{\sigma_\eta} \left(\eta_o - (1, 0, 0) \begin{pmatrix} \eta \\ u \\ v \end{pmatrix} \right) \\ &= \begin{pmatrix} \eta + \frac{1}{\sigma_\eta} (\eta_o - \eta) \\ u_U \\ v_U \end{pmatrix}. \end{aligned}$$

Now suppose we use the proposed velocity control variables $\mathbf{z} = (\eta, \psi_U, \chi_U)^T$, ageostrophic streamfunction and velocity potential. The gradient becomes,

$$\begin{aligned} \begin{pmatrix} \frac{\partial \mathcal{J}}{\partial \eta} \\ \frac{\partial \mathcal{J}}{\partial \psi_U} \\ \frac{\partial \mathcal{J}}{\partial \chi_U} \end{pmatrix} &= \begin{pmatrix} \eta \\ \psi_U \\ \chi_U \end{pmatrix} + \begin{pmatrix} 1 & -\frac{g}{f} \left(\frac{\partial}{\partial y} \right)^T & \frac{g}{f} \left(\frac{\partial}{\partial x} \right)^T \\ 0 & -\left(\frac{\partial}{\partial y} \right)^T & \left(\frac{\partial}{\partial x} \right)^T \\ 0 & \left(\frac{\partial}{\partial x} \right)^T & \left(\frac{\partial}{\partial y} \right)^T \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \frac{1}{\sigma_\eta} \left(\eta_o - (1, 0, 0) \begin{pmatrix} \eta \\ u \\ v \end{pmatrix} \right) \\ &= \begin{pmatrix} \eta + \frac{1}{\sigma_\eta} (\eta_o - \eta) \\ \psi_U \\ \chi_U \end{pmatrix}. \end{aligned}$$

As seen above, the absence of velocity observations results in no update to the velocity control variables. Since we have a linear system, the gradient can be set to zero to find the analysis, i.e. $\nabla \mathcal{J}(\mathbf{z}^a) = 0$. We have here in both cases the height analysis is equal to

$$\eta^a = \frac{\eta_o}{1 - \sigma_\eta}.$$

When using the ageostrophic velocities as control variables, we have

$$\begin{aligned} u_U^a = 0 &\Rightarrow u^a = -\frac{g}{f} \left(\frac{\partial \eta}{\partial y} \right) + 0 \\ v_U^a = 0 &\Rightarrow v^a = \frac{g}{f} \left(\frac{\partial \eta}{\partial x} \right) + 0 \end{aligned}$$

and using ageostrophic streamfunction and velocity potential as control variables, we have

$$\begin{aligned} \psi_U^a = 0 &\Rightarrow u^a = -\frac{g}{f} \left(\frac{\partial \eta}{\partial y} \right) + 0 \\ \chi_U^a = 0 &\Rightarrow v^a = \frac{g}{f} \left(\frac{\partial \eta}{\partial x} \right) + 0. \end{aligned}$$

Therefore, without velocity observations, we have shown in a simplified system that modifying the velocity control variables does not influence the resulting analysis. This conclusion also applies to more complex CVTs, such as the balance operator used in NEMOVAR. As described in section 3.4.3, the NEMOVAR balance operator takes temperature as the lead variable, and then separates salinity, sea surface height (SSH), and horizontal velocity into balanced and unbalanced components. This separation is performed in this specific order meaning the velocity variables do not influence the preceding variables — namely temperature, salinity, and SSH. As a result, we observe no impact on the analysis when modifying the velocity variables, if there are no direct velocity observations.

Appendix I

Single observation experiments

This appendix provides supplementary plots for the single observation experiments described in section 6.4, which were designed to evaluate the impact of introducing the proposed control variables — ageostrophic streamfunction and ageostrophic velocity potential — on the structure of the analysis increments. We perform these experiments at three different locations in the domain (shown in figure 6.6):

- An active region near the western boundary current at 30° N latitude and 78° W longitude;
- An inactive region at 30° N latitude and 65° W longitude;
- A north-eastern boundary region at 39° N latitude and 55° W longitude.

With section 6.4 presenting the results for the active region. The figures in this section are all centred around the observation locations.

In figures I.1 and I.3, we present the results from the inactive location, indicated by the yellow cross in figure 6.6. The structure of the increments is very similar to those observed in the active region (figures 6.7 and 6.8), though the magnitudes of the full increments are smaller, as expected in a region of reduced dynamical activity. An even greater reduction is seen near the boundary, marked by the red cross in figure 6.6. The corresponding results are shown in figures I.2 and I.4, where the TSCV speed increments are an order of magnitude smaller. Despite the reduced magnitudes, the spatial structures remain consistent with those observed in the active region.




Figure I.1: The velocity increments at the surface for the single observation experiments. The first column is the SLA observation experiment. The second column is the first experiment using a TSCV observation with current control variables and the final column is the second TSCV experiment with the new control variables. The first row is the unbalanced zonal velocity increments, second row is the full zonal velocity increments, the third row is the unbalanced meridional velocity increments and the final row is the full meridional velocity increments. These are the increments produced in the inactive region, indicated by the yellow cross in figure 6.6.




Figure I.2: The velocity increments at the surface for the single observation experiments. The first column is the SLA observation experiment. The second column is the first experiment using a TSCV observation with current control variables and the final column is the second TSCV experiment with the new control variables. The first row is the unbalanced zonal velocity increments, second row is the full zonal velocity increments, the third row is the unbalanced meridional velocity increments and the final row is the full meridional velocity increments. These are the increments produced near the boundary, indicated by the red cross in figure 6.6.




Figure I.3: The speed increments with velocity vectors superimposed at the surface for the single observation experiments. The first column is the SLA observation experiment. The second column is the first experiment using a TSCV observation with current control variables and the final column is the second TSCV experiment with the new control variables. The first row is the unbalanced speed increments, and the final row is the full speed increments. These are the increments produced in the inactive region, indicated by the yellow cross in figure 6.6.




Figure I.4: The speed increments at the surface with velocity vectors superimposed for the single observation experiments. The first column is the SLA observation experiment. The second column is the first experiment using a TSCV observation with current control variables and the final column is the second TSCV experiment with the new control variables. The first row is the unbalanced speed increments, and the final row is the full speed increments. These are the increments produced near the boundary, indicated by the red cross in figure 6.6.

Bibliography

- Anderson, D., Sheinbaum, J. and Haines, K. (1996), ‘Data assimilation in ocean models’, *Reports on Progress in Physics* **59**(10), 1209.
- Arakawa, A. and Lamb, V. R. (1977), ‘Computational design of the basic dynamical processes of the UCLA general circulation model’, *General circulation models of the atmosphere* **17**(Supplement C), 173–265.
- Ardhuin, F., Brandt, P., Gaultier, L., Donlon, C., Battaglia, A., Boy, F., Casal, T., Chapron, B., Collard, F., Cravatte, S. et al. (2019), ‘SKIM, a candidate satellite mission exploring global ocean currents and waves’, *Frontiers in Marine Science* **6**, 209.
- Atkinson, C. P., Rayner, N. A., Kennedy, J. J. and Good, S. A. (2014), ‘An integrated database of ocean temperature and salinity observations’, *Journal of Geophysical Research: Oceans* **119**(10), 7139–7163.
- Balmaseda, M. A., Mogensen, K. and Weaver, A. T. (2013), ‘Evaluation of the ECMWF ocean reanalysis system ORAS4’, *Quarterly journal of the royal meteorological society* **139**(674), 1132–1161.
- Bannister, R. N. (2008), ‘A review of forecast error covariance statistics in atmospheric variational data assimilation. II: Modelling the forecast error covariance statistics’, *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* **134**(637), 1971–1996.
- Bannister, R. N. (2017), ‘A review of operational methods of variational and ensemble-variational data assimilation’, *Quarterly Journal of the Royal Meteorological Society* **143**(703), 607–633.

- Bannister, R. N., Katz, D., Cullen, M. J. P., Lawless, A. S. and Nichols, N. K. (2008), ‘Modelling of forecast errors in geophysical fluid flows’, *International journal for numerical methods in fluids* **56**(8), 1147–1153.
- Barbosa Aguiar, A., Bell, M. J., Blockley, E., Calvert, D., Crocker, R., Inverarity, G., King, R., Lea, D. J., Maksymczuk, J., Martin, M. J. et al. (2024), ‘The Met Office Forecast Ocean Assimilation Model (FOAM) using a 1/12-degree grid for global forecasts’, *Quarterly Journal of the Royal Meteorological Society* **150**(763), 3827–3852.
- Baxter, G., Dance, S., Lawless, A. and Nichols, N. (2011), ‘Four-dimensional variational data assimilation for high resolution nested models’, *Computers & fluids* **46**(1), 137–141.
- Beckers, J.-M. and Deleersnijder, E. (1993), ‘Stability of a FBTCs scheme applied to the propagation of shallow-water inertia-gravity waves on various space grids’, *Journal of computational physics* **108**(1), 95–104.
- Bhatia, H., Norgard, G., Pascucci, V. and Bremer, P.-T. (2012), ‘The Helmholtz-Hodge decomposition — A survey’, *IEEE Transactions on visualization and computer graphics* **19**(8), 1386–1404.
- Blockley, E., Martin, M., McLaren, A., Ryan, A., Waters, J., Lea, D., Mirouze, I., Peterson, K., Sellar, A. and Storkey, D. (2014), ‘Recent development of the Met Office operational ocean forecasting system: an overview and assessment of the new global FOAM forecasts’, *Geoscientific Model Development* **7**(6), 2613–2638.
- Bloom, S. C., Takacs, L. L. and Brin, E. (1991), A scheme to incorporate analysis increments gradually in the GLA assimilation system, in ‘Ninth Conf. on Numerical Weather Prediction, Denver, CO, Amer. Meteor. Soc’, pp. 110–112.
- Bloom, S., Takacs, L., Da Silva, A. and Ledvina, D. (1996), ‘Data assimilation using incremental analysis updates’, *Monthly Weather Review* **124**(6), 1256–1271.
- Blumenthal, O. (1905), ‘Über die zerlegung unendlicher vektorfelder’, *Mathematische Annalen* **61**(2), 235–250.

- Brassington, G. B., Sakov, P., Divakaran, P., Aijaz, S., Sweeney-Van Kinderen, J., Huang, X. and Allen, S. (2023), OceanMAPS v4. 0i: a global eddy resolving EnKF ocean forecasting system, *in* ‘OCEANS 2023-Limerick’, IEEE, pp. 1–8.
- Burnett, W., Harper, S., Preller, R., Jacobs, G. and LaCroix, K. (2014), ‘Overview of operational ocean forecasting in the US Navy: Past, present, and future’, *Oceanography* **27**(3), 24–31.
- Calvetti, D. and Reichel, L. (2003), ‘Tikhonov regularization of large linear problems’, *BIT Numerical Mathematics* **43**, 263–283.
- Chai, T. and Draxler, R. R. (2014), ‘Root mean square error (RMSE) or mean absolute error (MAE)? - arguments against avoiding RMSE in the literature’, *Geoscientific model development* **7**(3), 1247–1250.
- Chrust, M., Weaver, A. T., Browne, P., Zuo, H. and Balmaseda, M. A. (2024), ‘Impact of ensemble-based hybrid background-error covariances in ECMWF’s next-generation ocean reanalysis system’, *Quarterly Journal of the Royal Meteorological Society* .
- Cochran, W. T., Cooley, J. W., Favin, D. L., Helms, H. D., Kaenel, R. A., Lang, W. W., Maling, G. C., Nelson, D. E., Rader, C. M. and Welch, P. D. (1967), ‘What is the fast Fourier transform?’, *Proceedings of the IEEE* **55**(10), 1664–1674.
- Cooley, J. W., Lewis, P. A. and Welch, P. D. (1967), ‘Historical notes on the fast Fourier transform’, *Proceedings of the IEEE* **55**(10), 1675–1677.
- Cooley, J. W., Lewis, P. A. and Welch, P. D. (1969), ‘The fast Fourier transform and its applications’, *IEEE Transactions on Education* **12**(1), 27–34.
- Cooley, J. W. and Tukey, J. W. (1965), ‘An algorithm for the machine calculation of complex Fourier series’, *Mathematics of computation* **19**(90), 297–301.
- Coppini, G., Clementi, E., Cossarini, G., Salon, S., Korres, G., Ravdas, M., Lecci, R., Pistoia, J., Goglio, A. C., Drudi, M. et al. (2023), ‘The Mediterranean Forecasting System–Part 1: Evolution and performance’, *Ocean Science* **19**(5), 1483–1516.

- Courtier, P., Thépaut, J.-N. and Hollingsworth, A. (1994), ‘A strategy for operational implementation of 4D-Var, using an incremental approach’, *Quarterly Journal of the Royal Meteorological Society* **120**(519), 1367–1387.
- Cullen, M. J. P. (2003), ‘Four-dimensional variational data assimilation: A new formulation of the background-error covariance matrix based on a potential-vorticity representation’, *Quarterly Journal of the Royal Meteorological Society* **129**(593), 2777–2796.
- Cummings, J. A. and Smedstad, O. M. (2013), Variational data assimilation for the global ocean, in ‘Data assimilation for atmospheric, oceanic and hydrologic applications (Vol. II)’, Springer, pp. 303–343.
- Cummings, J. A., Smith, S., Carroll, S. N. and Spence, P. (2010), User’s manual for the Navy coupled ocean data assimilation 3D variational analysis (NCODA 3DVar) system, version 3.0, Technical report.
- Daley, R. (1993), *Atmospheric data analysis*, number 2, Cambridge university press.
- de Rosnay, P., Browne, P., de Boissésou, E., Fairbairn, D., Hirahara, Y., Ochi, K., Schepers, D., Weston, P., Zuo, H., Alonso-Balmaseda, M. et al. (2022), ‘Coupled data assimilation at ECMWF: Current status, challenges and future developments’, *Quarterly Journal of the Royal Meteorological Society* **148**(747), 2672–2702.
- De Verdière, A. C. and Ollitrault, M. (2016), ‘A direct determination of the world ocean barotropic circulation’, *Journal of Physical Oceanography* **46**(1), 255–273.
- Denaro, F. M. (2003), ‘On the application of the Helmholtz–Hodge decomposition in projection methods for incompressible flows with general boundary conditions’, *International Journal for Numerical Methods in Fluids* **43**(1), 43–69.
- Derber, J. and Bouttier, F. (1999), ‘A reformulation of the background error covariance in the ECMWF global data assimilation system’, *Tellus A: Dynamic Meteorology and Oceanography* **51**(2), 195–221.
- Dickey, T. D. (2003), ‘Emerging ocean observations for interdisciplinary data assimilation systems’, *Journal of Marine Systems* **40**, 5–48.

- Dobricic, S. and Pinardi, N. (2008), ‘An oceanographic three-dimensional variational data assimilation scheme’, *Ocean modelling* **22**(3-4), 89–105.
- Dréville, M., Bourdallé-Badie, R., Derval, C., Lellouche, J., Rémy, E., Tranchant, B., Benkiran, M., Greiner, E., Guinehut, S., Verbrugge, N. et al. (2008), ‘The GODAE/Mercator-ocean global ocean forecasting system: results, applications and prospects’, *Journal of Operational Oceanography* **1**(1), 51–57.
- Durran, D. R. (2013), *Numerical methods for wave equations in geophysical fluid dynamics*, Vol. 32, Springer Science & Business Media.
- Edwards, C. A., Moore, A. M., Hoteit, I. and Cornuelle, B. D. (2015), ‘Regional ocean data assimilation’, *Annual review of marine science* **7**(1), 21–42.
- Falissard, F. (2013), ‘Genuinely multi-dimensional explicit and implicit generalized Shapiro filters for weather forecasting, computational fluid dynamics and aeroacoustics’, *Journal of Computational Physics* **253**, 344–367.
- Falissard, F. (2015), ‘Uneven-order decentered Shapiro filters for boundary filtering’, *Journal of Computational Physics* **292**, 168–175.
- Falissard, F. (2017), ‘Chebyshev-like generalized Shapiro filters for high-accuracy flow computations’, *Journal of Computational Physics* **336**, 595–607.
- Ford, D. and Barciela, R. (2017), ‘Global marine biogeochemical reanalyses assimilating two different sets of merged ocean colour products’, *Remote Sensing of Environment* **203**, 40–54.
- Friedlander, S. and Serre, D. (2002), *Handbook of mathematical fluid dynamics*, Elsevier.
- Fujii, Y., Yoshida, T., Sugimoto, H., Ishikawa, I. and Urakawa, S. (2023), ‘Evaluation of a global ocean reanalysis generated by a global ocean data assimilation system based on a four-dimensional variational (4DVAR) method’, *Frontiers in Climate* **4**, 1019673.
- Gaultier, L. and Ubelmann, C. (2024), ‘SKIM-like data simulation: System description, configuration and simulations’, *preparation for Remote Sensing*.

- Ghil, M. and Malanotte-Rizzoli, P. (1991), Data assimilation in meteorology and oceanography, *in* ‘Advances in geophysics’, Vol. 33, Elsevier, pp. 141–266.
- Gill, A. E. (2016), *Atmosphere—ocean dynamics*, Elsevier.
- Golub, G. H., Hansen, P. C. and O’Leary, D. P. (1999), ‘Tikhonov regularization and total least squares’, *SIAM journal on matrix analysis and applications* **21**(1), 185–194.
- Gommenginger, C., Chapron, B., Hogg, A., Buckingham, C. and Fox-Kemper, B. (2019), ‘SEASTAR: A mission to study ocean submesoscale dynamics and small-scale atmosphere-ocean processes in coastal’, *Shelf and Polar Seas, Frontiers in Marine Science* **6**.
- Groetsch, C. (1984), ‘The theory of Tikhonov regularization for Fredholm equations’, *Boston Pitman Publication* **104**.
- Groetsch, C. W. (1993), *Inverse problems in the mathematical sciences*, Vol. 52, Springer.
- Guiavarc’h, C., Roberts-Jones, J., Harris, C., Lea, D. J., Ryan, A. and Ascione, I. (2019), ‘Assessment of ocean analysis and forecast from an atmosphere–ocean coupled data assimilation operational system’, *Ocean Science* **15**(5), 1307–1326.
- Guiavarc’h, C., Storkey, D., Blaker, A. T., Blockley, E., Megann, A., Hewitt, H., Bell, M. J., Calvert, D., Copsey, D., Sinha, B. et al. (2025), ‘GOSI9: UK global ocean and sea ice configurations’, *Geoscientific Model Development* **18**(2), 377–403.
- Han, L. and Huang, R. X. (2020), ‘Using the Helmholtz decomposition to define the Indian Ocean meridional overturning streamfunction’, *Journal of Physical Oceanography* **50**(3), 679–694.
- Hanke, M. and Hansen, P. C. (1993), ‘Regularization methods for large-scale problems’, *Surv. Math. Ind* **3**(4), 253–315.
- Hansen, P. (1997), *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, Monographs on Mathematical Modeling and Computation, Society for Industrial and Applied Mathematics.
- URL:** <https://books.google.co.uk/books?id=3vxh3p03p6AC>

- Helber, R. W., Smith, S. R., Jacobs, G. A., Barron, C. N., Carrier, M. J., Rowley, C. D., Ngodock, H. E., Yaremchuk, M., Pasmans, I., Bartels, B. P. et al. (2023), ‘Ocean drifter velocity data assimilation, part 1: Formulation and diagnostic results’, *Ocean Modelling* **183**, 102195.
- Helmholtz, H. v. (1858), ‘Über integrale der hydrodynamischen gleichungen, welche den wirbelbewegungen entsprechen.’
- Hirose, N., Usui, N., Sakamoto, K., Tsujino, H., Yamanaka, G., Nakano, H., Urakawa, S., Toyoda, T., Fujii, Y. and Kohno, N. (2019), ‘Development of a new operational system for monitoring and forecasting coastal and open-ocean states around Japan’, *Ocean Dynamics* **69**, 1333–1357.
- Hodson, T. O. (2022), ‘Root mean square error (RMSE) or mean absolute error (MAE): When to use them or not’, *Geoscientific Model Development Discussions* **2022**, 1–10.
- Holdaway, D., Thuburn, J. and Wood, N. (2013), ‘Comparison of lorenz and Charney–Phillips vertical discretisations for dynamics–boundary layer coupling. part I: Steady states’, *Quarterly Journal of the Royal Meteorological Society* **139**(673), 1073–1086.
- Hollingsworth, A. and Lönnberg, P. (1986), ‘The statistical structure of short-range forecast errors as determined from radiosonde data. part I: the wind field’, *Tellus A* **38**(2), 111–136.
- Hoteit, I., Cornuelle, B., Kim, S. Y., Forget, G., Köhl, A. and Terrill, E. (2009), ‘Assessing 4D-VAR for dynamical mapping of coastal high-frequency radar in San Diego’, *Dynamics of Atmospheres and Oceans* **48**(1-3), 175–197.
- Isern-Fontanet, J., Ballabrera-Poy, J., Turiel, A. and García-Ladona, E. (2017), ‘Remote sensing of ocean surface currents: A review of what is being observed and what is being assimilated’, *Nonlinear Processes in Geophysics* **24**(4), 613–643.
- Kalnay, E. (2003), *Atmospheric modeling, data assimilation and predictability*, Cambridge university press.
- Kang, Y.-H., Kwak, D. Y. and Park, K. (2014), ‘Multigrid methods for improving the

- variational data assimilation in numerical weather prediction', *Tellus A: Dynamic Meteorology and Oceanography* **66**(1), 20217.
- Katz, D. (2007), The Application of PV-based Control Variable Transformations in Variational Data Assimilation, PhD thesis, University of Reading.
- Katz, D., Lawless, A. S., Nichols, N. K., Cullen, M. J. P. and Bannister, R. N. (2011), 'Correlations of control variables in variational data assimilation', *Quarterly Journal of the Royal Meteorological Society* **137**(656), 620–630.
- Kourafalou, V., De Mey, P., Le Hénaff, M., Charria, G., Edwards, C., He, R., Herzfeld, M., Pascual, A., Stanev, E., Tintoré, J. et al. (2015), 'Coastal ocean forecasting: system integration and evaluation', *Journal of Operational Oceanography* **8**(sup1), s127–s146.
- Lalli, C. and Parsons, T. R. (1997), *Biological oceanography: an introduction*, Elsevier.
- Lawless, A. S. (2013), 'Variational data assimilation for very large environmental problems, vol. 13', *Large scale inverse problems: computational methods and applications in the Earth Sciences* **2**, 55–90.
- Lawless, A. S., Gratton, S. and Nichols, N. K. (2005), 'An investigation of incremental 4D-Var using non-tangent linear models', *Quarterly Journal of the Royal Meteorological Society* **131**(606), 459–476.
- Lawless, A. S., Nichols, N. K. and Ballard, S. (2003), 'A comparison of two methods for developing the linearization of a shallow-water model', *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* **129**(589), 1237–1254.
- Lea, D., Martin, M. and Oke, P. (2014), 'Demonstrating the complementarity of observations in an operational ocean forecasting system', *Quarterly Journal of the Royal Meteorological Society* **140**(683), 2037–2049.
- Lea, D., Mirouze, I., Martin, M., King, R., Hines, A., Walters, D. and Thurlow, M. (2015), 'Assessing a new coupled data assimilation system based on the Met Office coupled atmosphere–land–ocean–sea ice model', *Monthly Weather Review* **143**(11), 4678–4694.

- Lee, M.-S., Barker, D., Huang, W. and Kuo, Y.-H. (2004), First guess at appropriate time (FGAT) with WRF 3DVAR, *in* ‘WRF/MM5 Users Workshop’, pp. 22–25.
- Lellouche, J.-M., Greiner, E., Le Galloudec, O., Garric, G., Regnier, C., Drevillon, M., Benkiran, M., Testut, C.-E., Bourdalle-Badie, R., Gasparin, F. et al. (2018), ‘Recent updates to the Copernicus marine service global ocean monitoring and forecasting real-time 1/12 high-resolution system’, *Ocean Science* **14**(5), 1093–1126.
- Lellouche, J.-M., Le Galloudec, O., Drévillon, M., Régnier, C., Greiner, E., Garric, G., Ferry, N., Desportes, C., Testut, C.-E., Bricaud, C. et al. (2013), ‘Evaluation of global monitoring and forecasting systems at Mercator Océan’, *Ocean Science* **9**(1), 57–81.
- Lévy, M., Klein, P., Tréguier, A.-M., Iovino, D., Madec, G., Masson, S. and Takahashi, K. (2010), ‘Modifications of gyre circulation by sub-mesoscale physics’, *Ocean Modelling* **34**(1-2), 1–15.
- Li, Y., Navon, I. M., Yang, W., Zou, X., Bates, J., Moorthi, S. and Higgins, R. (1994), ‘Four-dimensional variational data assimilation experiments with a multilevel semi-lagrangian semi-implicit general circulation model’, *Monthly weather review* **122**(5), 966–983.
- Li, Z., Chao, Y. and McWilliams, J. C. (2006), ‘Computation of the streamfunction and velocity potential for limited and irregular domains’, *Monthly weather review* **134**(11), 3384–3394.
- Li, Z., Chao, Y., McWilliams, J. C. and Ide, K. (2008a), ‘A three-dimensional variational data assimilation scheme for the regional ocean modeling system’, *Journal of Atmospheric and Oceanic Technology* **25**(11), 2074–2090.
- Li, Z., Chao, Y., McWilliams, J. C. and Ide, K. (2008b), ‘A three-dimensional variational data assimilation scheme for the regional ocean modeling system: Implementation and basic experiments’, *Journal of Geophysical Research: Oceans* **113**(C5).
- Li, Z., McWilliams, J. C., Ide, K. and Farrara, J. D. (2015), ‘Coastal ocean data assimilation using a multi-scale three-dimensional variational scheme’, *Ocean Dynamics* **65**, 1001–1015.

- Lions, J.-L., Temam, R. and Wang, S. (1992), ‘On the equations of the large-scale ocean’, *Nonlinearity* **5**(5), 1007.
- Lorenc, A., Ballard, S., Bell, R., Ingleby, N., Andrews, P., Barker, D., Bray, J., Clayton, A., Dalby, T., Li, D. et al. (2000), ‘The Met Office global three-dimensional variational data assimilation scheme’, *Quarterly Journal of the Royal Meteorological Society* **126**(570), 2991–3012.
- Lorenc, A. C. (1986), ‘Analysis methods for numerical weather prediction’, *Quarterly Journal of the Royal Meteorological Society* **112**(474), 1177–1194.
- Lorenc, A. C. (2003), ‘Modelling of error covariances by 4D-Var data assimilation’, *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* **129**(595), 3167–3182.
- Lorenc, A. C. and Rawlins, F. (2005), ‘Why does 4D-Var beat 3D-Var?’, *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* **131**(613), 3247–3257.
- Lynch, P. (1989), ‘Partitioning the wind in a limited domain’, *Monthly weather review* **117**(7), 1492–1500.
- Madec, G., Bourdallé-Badie, R., Chanut, J., Clementi, E., Coward, A., Ethé, C., Iovino, D., Lea, D., Lévy, C., Lovato, T., Martin, N., Masson, S., Mocavero, S., Rousset, C., Storkey, D., Müeller, S., Nurser, G., Bell, M., Samson, G., Mathiot, P., Mele, F. and Moulin, A. (2023), ‘NEMO ocean engine’.
URL: <https://doi.org/10.5281/zenodo.6334656>
- Marshall, J., Adcroft, A., Hill, C., Perelman, L. and Heisey, C. (1997), ‘A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers’, *Journal of Geophysical Research: Oceans* **102**(C3), 5753–5766.
- Martin, M. J., Hoteit, I., Bertino, L. and Moore, A. M. (2024), ‘Data assimilation schemes for ocean forecasting: state of the art’, *State of the Planet Discussions* **2024**, 1–16.
- Masutani, M., Schlatter, T. W., Errico, R. M., Stoffelen, A., Andersson, E., Lahoz, W., Woollen, J. S., Emmitt, G. D., Riishøjgaard, L.-P. and Lord, S. J. (2010), Observing

- system simulation experiments, in ‘Data Assimilation: Making sense of observations’, Springer, pp. 647–679.
- Matsuno, T. (1966), ‘Numerical simulation of the primitive equations by a simulated backward difference method’, *J. Meteorol. Soc. Japan* **44**, 76–84.
- Ménétrier, B. and Auligné, T. (2015), ‘An overlooked issue of variational data assimilation’, *Monthly Weather Review* **143**(10), 3925–3930.
- Mignac, D., Waters, J., Lea, D. J., Martin, M. J., While, J., Weaver, A. T., Vidard, A., Guiavarc’h, C., Storkey, D., Ford, D. et al. (2025), ‘Improvements to the Met Office’s global ocean–sea ice forecasting system including model and data assimilation changes’, *Geoscientific Model Development* **18**(11), 3405–3425.
- Mirouze, I., Rémy, E., Lellouche, J.-M., Martin, M. J. and Donlon, C. J. (2024), ‘Impact of assimilating satellite surface velocity observations in the Mercator Ocean international analysis and forecasting global 1/4° system’, *Frontiers in Marine Science* **11**, 1376999.
- Mogensen, K., Balmaseda, M., Weaver, A., Martin, M. and Vidard, A. (2009), ‘NEMOVAR: A variational data assimilation system for the NEMO ocean model’, *ECMWF newsletter* **120**, 17–22.
- Mogensen, K. and Balmaseda, W. (2012), ‘The NEMOVAR ocean data assimilation system as implemented in the ECMWF ocean analysis for system 4’.
- URL:** <https://www.ecmwf.int/node/11174>
- Moore, A. M., Arango, H. G., Broquet, G., Powell, B. S., Weaver, A. T. and Zavala-Garay, J. (2011), ‘The regional ocean modeling system (ROMS) 4-dimensional variational data assimilation systems: Part I - system overview and formulation’, *Progress in Oceanography* **91**(1), 34–49.
- Moore, A. M., Arango, H. G., Wilkin, J. and Edwards, C. A. (2023), ‘Weak constraint 4D-Var data assimilation in the regional ocean modeling system (ROMS) using a saddle-point algorithm: Application to the california current circulation’, *Ocean Modelling* p. 102262.

- Navon, I. M., Zou, X., Derber, J. and Sela, J. (1992), ‘Variational data assimilation with an adiabatic version of the NMC spectral model’, *Monthly weather review* **120**(7), 1433–1446.
- Nazareth, J. L. (2009), ‘Conjugate gradient method’, *Wiley Interdisciplinary Reviews: Computational Statistics* **1**(3), 348–353.
- Ngodock, H., Muscarella, P., Carrier, M., Souopgui, I. and Smith, S. (2015), Assimilation of HF radar observations in the Chesapeake-Delaware Bay region using the Navy Coastal Ocean Model (NCOM) and the Four-Dimensional Variational (4DVAR) method, *in* ‘Coastal Ocean Observing Systems’, Elsevier, pp. 373–391.
- Nocedal, J. and Wright, S. J. (2006), ‘Conjugate gradient methods’, *Numerical optimization* pp. 101–134.
- Novellino, A., Arnaud, A., Schiller, A. and Wan, L. (2024), ‘End user applications for ocean forecasting: present status description’, *State of the Planet Discussions* **2024**, 1–6.
- ODYSEA (2023).
URL: <https://odysea.ucsd.edu/>
- Panchev, S. (1971), *Random Functions and Turbulence*, Pergamon Press 444pp.
- Park, J.-Y., Stock, C. A., Yang, X., Dunne, J. P., Rosati, A., John, J. and Zhang, S. (2018), ‘Modeling global ocean biogeochemistry with physical data assimilation: A pragmatic solution to the equatorial instability’, *Journal of Advances in modeling earth systems* **10**(3), 891–906.
- Parrish, D. F. and Derber, J. C. (1992), ‘The National Meteorological Center’s spectral statistical-interpolation analysis system’, *Monthly Weather Review* **120**(8), 1747–1763.
- Pedlosky, J. (2013), *Geophysical fluid dynamics*, Springer Science & Business Media.
- Pham, D. T., Verron, J. and Roubaud, M. C. (1998), ‘A singular evolutive extended kalman filter for data assimilation in oceanography’, *Journal of Marine systems* **16**(3–4), 323–340.

- Phillips, D. L. (1962), ‘A technique for the numerical solution of certain integral equations of the first kind’, *Journal of the ACM (JACM)* **9**(1), 84–97.
- Phillips, N. A. (1959), ‘An example of non-linear computational instability’, *The atmosphere and the sea in motion* **501**, 504.
- Pilo, G. S., Oke, P. R., Coleman, R., Rykova, T. and Ridgway, K. (2018), ‘Impact of data assimilation on vertical velocities in an eddy resolving ocean model’, *Ocean Modelling* **131**, 71–85.
- Polavarapu, S., Ren, S., Rochon, Y., Sankey, D., Ek, N., Koshyk, J. and Tarasick, D. (2005), ‘Data assimilation with the Canadian middle atmosphere model’, *Atmosphere-ocean* **43**(1), 77–100.
- Randall, D. A. (1994), ‘Geostrophic adjustment and the finite-difference shallow-water equations’, *Monthly Weather Review* **122**(6), 1371–1377.
- Sangster, W. E. (1960), ‘A method of representing the horizontal pressure force without reduction of station pressures to sea level’, *Journal of Atmospheric Sciences* **17**(2), 166–176.
- Schoder, S., Roppert, K. and Kaltenbacher, M. (2020), ‘Helmholtz’s decomposition for compressible flows and its application to computational aeroacoustics’, *SN Partial Differential Equations and Applications* **1**, 1–20.
- Semane, N. and Bonavita, M. (2025), ‘Reintroducing the analysis of humidity in the stratosphere’.
- URL:** *nbsp;*
- Shapiro, R. (1970), ‘Smoothing, filtering, and boundary effects’, *Reviews of geophysics* **8**(2), 359–387.
- Shapiro, R. (1971), ‘The use of linear filtering as a parameterization of atmospheric diffusion’, *Journal of Atmospheric Sciences* **28**(4), 523–531.
- Shapiro, R. (1975), ‘Linear filtering’, *Mathematics of computation* **29**(132), 1094–1097.

- Shewchuk, J. R. et al. (1994), ‘An introduction to the conjugate gradient method without the agonizing pain’.
- Skamarock, W. C. and Klemp, J. B. (2008), ‘A time-split nonhydrostatic atmospheric model for weather research and forecasting applications’, *Journal of computational physics* **227**(7), 3465–3485.
- Sluka, T. C., Penny, S. G., Kalnay, E. and Miyoshi, T. (2016), ‘Assimilating atmospheric observations into the ocean using strongly coupled ensemble data assimilation’, *Geophysical Research Letters* **43**(2), 752–759.
- Smith, G. C., Roy, F., Reszka, M., Surcel Colan, D., He, Z., Deacu, D., Belanger, J.-M., Skachko, S., Liu, Y., Dupont, F. et al. (2016), ‘Sea ice forecast verification in the canadian global ice ocean prediction system’, *Quarterly Journal of the Royal Meteorological Society* **142**(695), 659–671.
- Smith, S., Ngodock, H., Carrier, M., Shriver, J., Muscarella, P. and Souopgui, I. (2017), ‘Validation and operational implementation of the navy coastal ocean model four dimensional variational data assimilation system (NCOM 4DVAR) in the okinawa trough’, *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications (Vol. III)* pp. 405–427.
- Smith, S. R., Helber, R. W., Jacobs, G. A., Barron, C. N., Carrier, M., Rowley, C., Ngodock, H., Pasmans, I., Bartels, B., DeHaan, C. et al. (2023), ‘Ocean drifter velocity data assimilation part 2: Forecast validation’, *Ocean Modelling* **185**, 102260.
- Stewart, R. H. (2008), *Introduction to physical oceanography*, Robert H. Stewart.
- Stoffelen, A. (2022), The ESA Earth Explorer 10 Harmony mission for ocean-atmosphere interaction, in ‘AGU Fall Meeting Abstracts’, Vol. 2022, pp. OS11B–02.
- Stommel, H. (1948), ‘The westward intensification of wind-driven ocean currents’, *Eos, Transactions American Geophysical Union* **29**(2), 202–206.
- Storto, A., Masina, S. and Navarra, A. (2016), ‘Evaluation of the CMCC eddy-permitting global ocean physical reanalysis system (C-GLORS, 1982–2012) and its assimilation components’, *Quarterly Journal of the Royal Meteorological Society* **142**(695), 738–758.

- Talagrand, O. and Courtier, P. (1987), ‘Variational assimilation of meteorological observations with the adjoint vorticity equation. I: Theory’, *Quarterly Journal of the Royal Meteorological Society* **113**(478), 1311–1328.
- Thépaut, J.-N., Courtier, P., Belaud, G. and Lemaître, G. (1996), ‘Dynamical structure functions in a four-dimensional variational assimilation: A case study’, *Quarterly Journal of the Royal Meteorological Society* **122**(530), 535–561.
- Tikhonov, A. N. (1963), On the solution of ill-posed problems and the method of regularization, in ‘Doklady akademii nauk’, Vol. 151, Russian Academy of Sciences, pp. 501–504.
- Tonani, M., Sykes, P., King, R. R., McConnell, N., Péquignet, A.-C., O’Dea, E., Graham, J. A., Polton, J. and Siddorn, J. (2019), ‘The impact of a new high-resolution ocean model on the Met Office North-West European Shelf forecasting system’, *Ocean Science* **15**(4), 1133–1158.
- Usui, N., Fujii, Y., Sakamoto, K. and Kamachi, M. (2015), ‘Development of a four-dimensional variational assimilation system for coastal data assimilation around Japan’, *Monthly Weather Review* **143**(10), 3874–3892.
- Waller, J. A., Dance, S. L. and Nichols, N. K. (2017), ‘On diagnosing observation-error statistics with local ensemble data assimilation’, *Quarterly Journal of the Royal Meteorological Society* **143**(708), 2677–2686.
- Waters, J., Lea, D. J., Martin, M. J., Mirouze, I., Weaver, A. and While, J. (2015), ‘Implementing a variational data assimilation system in an operational 1/4 degree global ocean model’, *Quarterly Journal of the Royal Meteorological Society* **141**(687), 333–349.
- Waters, J., Martin, M. J., Bell, M. J., King, R. R., Gaultier, L., Ubelmann, C., Donlon, C. and Van Gennip, S. (2024a), ‘Assessing the potential impact of assimilating total surface current velocities in the Met Office’s global ocean forecasting system’, *Frontiers in Marine Science* **11**, 1383522.
- Waters, J., Martin, M. J., Bell, M. J., King, R. R., Gaultier, L., Ubelmann, C., Donlon, C. and Van Gennip, S. (2024b), ‘The impact of simulated total surface current velocity

- observations on operational ocean forecasting and requirements for future satellite missions’, *Frontiers in Marine Science* **11**, 1408495.
- Watterson, I. (2001), ‘Decomposition of global ocean currents using a simple iterative method’, *Journal of Atmospheric and Oceanic Technology* **18**(4), 691–703.
- Weaver, A. and Courtier, P. (2001), ‘Correlation modelling on the sphere using a generalized diffusion equation’, *Quarterly Journal of the Royal Meteorological Society* **127**(575), 1815–1846.
- Weaver, A. T., Chrust, M., Ménétrier, B. and Piacentini, A. (2021), ‘An evaluation of methods for normalizing diffusion-based covariance operators in variational data assimilation’, *Quarterly Journal of the Royal Meteorological Society* **147**(734), 289–320.
- Weaver, A. T., Deltel, C., Machu, É., Ricci, S. and Daget, N. (2005), ‘A multivariate balance operator for variational ocean data assimilation’, *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography* **131**(613), 3605–3625.
- Weaver, A. T., Tshimanga, J. and Piacentini, A. (2016), ‘Correlation operators based on an implicitly formulated diffusion equation solved with the Chebyshev iteration’, *Quarterly Journal of the Royal Meteorological Society* **142**(694), 455–471.
- Weaver, A., Vialard, J. and Anderson, D. (2003), ‘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.
- Webb, D. (1996), ‘An ocean model code for array processor computers’, *Computers & Geosciences* **22**(5), 569–578.
- Weller, H., Thuburn, J. and Cotter, C. J. (2012), ‘Computational modes and grid imprinting on five quasi-uniform spherical C grids’, *Monthly weather review* **140**(8), 2734–2755.
- Wilks, D. S. (2011), *Statistical methods in the atmospheric sciences*, Vol. 100, Academic press.

- Williamson, D. L., Drake, J. B., Hack, J. J., Jakob, R. and Swarztrauber, P. N. (1992), ‘A standard test set for numerical approximations to the shallow water equations in spherical geometry’, *Journal of computational physics* **102**(1), 211–224.
- Xiao, Q. and Sun, J. (2007), ‘Multiple-radar data assimilation and short-range quantitative precipitation forecasting of a squall line observed during IHOP_2002’, *Monthly Weather Review* **135**(10), 3381–3404.
- Zhang, S., Liu, Z., Zhang, X., Wu, X., Han, G., Zhao, Y., Yu, X., Liu, C., Liu, Y., Wu, S. et al. (2020), ‘Coupled data assimilation and parameter estimation in coupled ocean-atmosphere models: a review’, *Climate Dynamics* **54**, 5127–5144.
- Zuo, H., Balmaseda, M. A., Tietsche, S., Mogensen, K. and Mayer, M. (2019), ‘The ECMWF operational ensemble reanalysis–analysis system for ocean and sea ice: a description of the system and assessment’, *Ocean science* **15**(3), 779–808.