Zonal Jets and Turbulence

William Andrew Jackman

A dissertation submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Mathematics University College London

November 16, 2022

I, William Andrew Jackman, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Disclaimer

The author has benefited from work by his supervisor, Prof. J. G. Esler. The calculation in Appendix (A.1) is solely Prof. Esler's.

Elements of Chapter 4 and Chapter 5 form part of a manuscript which has been recently submitted for publication in the *Journal of Fluid Mechanics*. All the calculations and numerical results are the authors', but much of the presentation of sections 4.1–4.3 and Chapter 5 should be considered as joint work.

Abstract

From the atmospheres of giant planets to terrestrial oceans, self-organisation of turbulence into zonal jets is a ubiquitous feature of geophysical fluid dynamics. The relationship between jets and turbulence is rich, with many unanswered mathematical questions. The aim of this thesis is to extend the existing literature and understanding of jets using numerical and theoretical approaches.

It's well known the number of jets that form in fluid dynamical models can spontaneously change. Such transitions can occur with increasing rarity as the model parameters vary. The first interesting problem in this thesis studies the use of a 'rare event algorithm' to calculate transition probabilities when it is too expensive to do so by direct numerical simulation. We verify the effectiveness of the algorithm in this context and successfully adapt the algorithm for application to a deterministic two-layer (baroclinic) model.

The second problem focuses on the stochastically forced single-layer (barotropic) model, with the motivating question 'Can eddy momentum fluxes be parameterised?'. If so, the mean jet profile can be solved directly. The key idea is if the turbulent scale is small compared to the background flow, eddy motions only 'see' local properties of the flow. For the problem to be mathematically tractable we adopt a second order statistical truncation of the model (known variously as CE2 or SSST). Within this framework, we simplify an existing expression for eddy momentum flux in terms of the local background shear. The new expression can be readily computed and is shown to agree with the statistical model when the background flow is stable and steady. The expression is tested as a method of parameterising small scale turbulence in various scenarios with great success. An

Abstract 5

important discovery of the work is the role that emergent secondary (barotropic) instabilities play in the mean momentum balance for jets.

Impact Statement

The expertise, knowledge, discoveries, and insights presented within this thesis engage with, and contribute to, the broad academic research in the field of rapidly rotating planetary turbulence. Novel theoretical insights, methodologies and calculations are put forward in this work, some of which have been presented at conferences and seminars to the international academic community. Parts of the thesis have been submitted for publication in a leading journal for the subject material. In various sections of the thesis there is an explicit expression for the potential of future scholarly avenues in the form of open questions and natural extensions to the research.

As with all geophysical fluid dynamics research, the subject territory naturally has a broad impact outside of academia given that we all experience, and are ultimately at the mercy of, the climate. This specific work places an emphasis on methods to calculate the likelihood of rare events in turbulence models. Advances in probability estimates of rare weather events (heat waves, flooding, hurricanes, etc...) provide a public health and environmental benefit, particularly in a world where rare events are increasingly extreme.

Acknowledgements

I'm greatly indebted to Prof. Esler. For the patient support, the supervision, and for the personal encouragement, I express my sincere gratitude.

I would like to thank the UCL Mathematics department as a whole for providing the entirety of my university education. Many of the staff, past and present, academic or otherwise, directly and indirectly, knowingly or unknowingly, have contributed toward my mathematical development.

Most of all, the past four years would have been inconceivable without the unconditional support of my family.

Contents

1	Intr	oductor	y Material	21
	1.1	Barotr	opic, Quasi-Geostrophic Equations	24
		1.1.1	Derivation	24
		1.1.2	Energy and Enstrophy	28
		1.1.3	Fourier-Space and Quadratic Cascades	29
		1.1.4	Stochastic Forcing	31
		1.1.5	Non-dimensional Parameters for Jet Dynamics	34
		1.1.6	Model Integrations	39
	1.2	Two-L	ayer Baroclinic, Quasi-Geostrophic Equations	43
		1.2.1	Derivation	43
		1.2.2	Boundary Conditions and Parameter Specification	46
		1.2.3	Energy	47
		1.2.4	Parameter Space and Results	49
	1.3	Numer	rical methods	54
		1.3.1	Numerical Methods For Barotropic Flow	54
		1.3.2	Numerical Methods For Two-Layer Flow	56
2	Jet '	Fransiti	ons and Rare Events	59
	2.1	Introdu	uction	59
	2.2	Adapti	ive Multi-Level Splitting Algorithm (AMS)	63
		2.2.1	The Algorithm	63
		2.2.2	Initialisation	65
		2.2.3	Scoring Function	65

Contents 9

		2.2.4	Probability Estimators	67
		2.2.5	Return Times	70
		2.2.6	Parallelising	71
		2.2.7	Results	75
	2.3	Deterr	ninistic Adaptive Multilevel Splitting Algorithm (DAMS)	81
		2.3.1	The Algorithm	81
		2.3.2	Splitting Perturbation: Idea and Theory	82
		2.3.3	Splitting Perturbation: Practical Implementation	84
		2.3.4	Scoring Function	86
		2.3.5	Results	87
3	Seco	ond Ord	ler Closure Models	98
	3.1	Introd	uction	98
	3.2	Quasi-	Linear Approximation	101
	3.3	CE2/S	3T: Theory, Derivation and Equations	103
		3.3.1	Derivation	103
		3.3.2	The Link Between CE2 and QL	105
		3.3.3	CE2 in Zonal Fourier-Space	105
	3.4	Result	s	107
4	A Lo	ocal Clo	sure Theory for Momentum Flux in the Limit of Small Scale	e
	Ford	cing		109
	4.1	Introd	uction	109
	4.2	The S	Y14 local theory	112
		4.2.1	The SY14 result for the momentum flux in a linear shear flow	v112
		4.2.2	General Properties of the Momentum Flux Formula	117
		4.2.3	The SY14 Momentum Flux Closure Applied to a Jet Flow .	123
	4.3	Mome	ntum Balance in Equilibrated Jets	126
		4.3.1	Experimental Design: Separation of Momentum Flux Con-	
			tributions	127
		4.3.2	Momentum Flux Decomposition in an Equilibrated let Flow	129

В	Colo	phon		185
	A.3	Kernel	Properties	180
		A.2.2	$\langle u'u' \rangle$ and $\langle v'v' \rangle$	179
		A.2.1	Reynolds Stress $\langle u'v' \rangle$	178
	A.2		s for Reynolds Stress Tensor on Steady Background Flow	
	A.1	Sheare	d Disturbance Solution	176
A	SY1	4 Local	Closure Theory	176
Aŗ	pend	ices		176
A =-	mond	ioos		174
6	Gen	eral Co	nclusions	173
		5.3.2	Addition Numerical Results	168
		5.3.1	Set-Up and Main Results	164
	5.3	A Typi	cal Equilibrated Jet Flow in β -plane Turbulence	164
		5.2.2	Additional Numerical Results	159
		5.2.1	Set-Up and Main Results	158
	5.2	A Stab	le Zonal Flow With Alternating Jets	158
	5.1	A Stab	le Monotonic Zonal Flow	156
5	App	lication	s of the Local Closure Theory	154
		4.5.3	Scattering Experiments Over Time-Variable Shear	146
		4.5.2	Time-variable Momentum Flux Formulae	
		4.5.1	Why Time-Variability Matters	
	4.5		Variability of the Local Closure	
		4.4.3	Scattering Experiment Numerical Results	139
		4.4.2	Formulae and Eddy Kinetic Energy Interpretation	138
		4.4.1	Reynolds Stress Tensor	137
	4.4	The Re	eynolds Stress Tensor	137
		4.3.3	Linear Stability Analysis of the Equilibrium Jet Flow	134

List of Figures

1.1	Snapshots of vorticity in freely decaying 2D turbulence in a doubly-	
	periodic domain. Time increases clockwise from top-left, along with	
	the energy containing length scales	23
1.2	Integration of (1.1.21) from spin-up. Top: Hovmöller and final zonal	
	velocity. Middle: final vorticity. Bottom: final energy enstrophy.	
	Parameters: $\beta = 5.26$, $\mu = 0.0012$, $\varepsilon = 0.0024$, $k_f = 14.5$ with an	
	annulus half width of 0.6. Numerical resolution is 256×256 and	
	$v_4 = 2.5 \times 10^{-8}$	41
1.3	Hovmöller contour plots for zonal velocity against time in QG2D	
	simulations. Top to bottom: $\alpha = (2, 4, 8, 12) \times 10^{-3}$. Left to	
	right $\beta = [2^2, 3^2, 4^2, 5^2]/2$ Zonostrophy values according to (1.1.66)	
	given in each case. Even these low resolution 128×128 , $k_f = 10$	
	give the expected scalings. All simulations are run for $50\alpha^{-1}$ non-	
	dimensional time units	42
1.4	Two-layer model results with $(W_c, \kappa_M) = (3.25\pi, 0.07)$. Top: con-	
	toured is the Hovmöller plot for the depth averaged zonal mean	
	flow and plotted are the mean velocity profiles as labelled. Bottom:	
	final snapshot of the depth averaged (barotropic) potential vorticity	
	$Q^+ + \beta y$	50
1.5	Same as in Figure 1.4 but with a narrower equilibrium profile and a	
	reduced friction, $(W_c, \kappa_M) = (2.4\pi, 0.03)$	51
1.6	Same as in Figure 1.4 but with parameters setting for two jets.	
	$(W_c, \kappa_M) = (3\pi, 0.02).$	52

1.7	Same as in Figure 1.4 but with parameters setting for two jets.	
	$(W_c, \kappa_M) = (3.25\pi, 0.05)$	53
2.1	Schematic representation of the AMS algorithm with an ensemble	
	size of three. The first iteration, on the left, shows S_1 being split	
	from one of the other two trajectories (in this case S_2) at score Q_1 .	
	The purple trajectory is the new trajectory and the yellow star marks	
	the split. The next iteration of the algorithm, on the right, shows the	
	adjusted $\{S_1, S_2, S_3\}$ ensemble, re-ordered by ascending maximum	
	scores. A new split is shown following the same procedure but this	
	time branching off S_3	63
2.2	Histogram of return times for the 2 jet state from 10000 simulations.	
	Reset times greater than 400 are not shown. This is for an improved	
	display of the histogram at lower τ values	71
2.3	Transitions between two and three jets in QG2D found by DNS. Top:	
	Hövmoller plot of zonal mean velocity, bottom: corresponding line	
	plot of $ \overline{\zeta}_2 $ (blue) and $ \overline{\zeta}_3 $ (orange) showing abrupt transitions	75
2.4	A selection of two-to-three jet transitions for QG2D found using	
	AMS, corresponding to the transitions found from ensembles 11-18	
	in Table 2.1	79
2.5	Zonal vorticity Fourier component amplitudes for the middle-middle	
	(left) and bottom-left (right) transitions in Figure 2.4. Note the right	
	hand side has a longer time axis	79
2.6	The potential vorticity at the beginning and end of the bottom-left	
	transition in Figure 2.4	80
2.7	Transitions in the reduced zonal Fourier component amplitude space	
	$(\overline{\zeta}_2 , \overline{\zeta}_3 , \overline{\zeta}_4)$. The boxes are the sets <i>A</i> and <i>B</i> and the transition	
	from red to blue is from two to three jets. The left panel is for the	
	eight transitions from ensembles 11-18. The right panel plots a full	
	ensemble of 100 members where all trajectories were guided towards	
	set <i>B</i>	80

2.8	The initial perturbation magnitudes are $\varepsilon_s = (1, 0.25, 10^{-4}, 10^{-12})$	
	with corresponding colours (blue, orange, yellow, purple). i.e. de-	
	creasing ε_s corresponds to a delayed perturbation growth and trajec-	
	tory decorrelation	85
2.9	Top: Hovmöller of $U^+(y,t)$ in a two-layer flow with parameters	
	$(W_c, \kappa_M) = (3.25\pi, 0.05)$. Below: $ U_1^+ $ (blue) and $ U_2^+ $ (orange) as	
	defined in (2.3.9). The clear flipping of magnitudes indicates that	
	the low wavenumbers will make a good scoring function	88
2.10	Average jet number in the (W_c, κ_M) parameter space found through	
	DNS. The approximate 1.5 jet number locations are given by the	
	black circled crosses	89
2.11	Results from DNS varying κ_M with $W_c = 2.6\pi$ fixed (corresponding	
	to the vertical sweep in Figure 2.10). The upward pointing black	
	triangles indicate the mean waiting time to go from a single to double	
	jet. Downwards pointing arrows are the mean waiting time to go	
	from a double to a single jet. Blue dots mark the number of observed	
	transitions per 10^5 time units, indicated by the right hand side axis	89
2.12	Histogram of 'reset times' from 2000 simulations. $(W_c, \kappa_M) =$	
	$(2.6\pi, 0.028)$	91
2.13	Transitions from a single to double jet for friction values $\kappa_M = 0.028$	
	(left) and 0.01 (right) found using DAMS. In each plot there are 200	
	transitions taken randomly from an algorithm implementation. The	
	top row is in (U_1^+ , U_2^+) -space, and the bottom row is an Argand	
	diagram with U_1^+ plotted in red and U_2^+ plotted in blue. The grey	
	boxes/coloured circles give the boundaries of sets A and B	93
2.14	Example of a transition with $\kappa_M = 0.028$ where a second jet forms	
	'south' of the single jet	94
2.15	Example of a transition with $\kappa_M = 0.028$ where a second jet forms	
	'north' of the single jet	95
2.16	Example of a transition with $\kappa_M = 0.01$	96

2.17	Example of a transition with $\kappa_M = 0.01$ where the jet oscillates
	before a second jet gains sufficient strength to reach set B 97
3.1	Diagram of triad interactions. From left to right: wave-mean interaction where a wave interacts directly with the zonal flow, eddy-eddy interaction of equal and opposite signed waves resulting in Reynolds stress on the mean flow, and lastly all other eddy-eddy interactions not involving mean zonal flow. The first two interactions are retained for QL theory while the final interaction on the right is dropped 103
3.2	Top to bottom: Nonlinear, quasi-linear and CE2 Hovmöller contour plots of $U(y,t)$
3.3	Line plots of U(y) at $t=5/\mu$ for the nonlinear (nlue), quasi-linear (orange) and CE2 (yellow) simulations. These are the same simulations as plotted in Figure 3.2
4.1	Left: the function $K(\phi,m)$ against $\phi \in [-\pi/2,\pi/2]$ for various values of m . When $m\gg 1$, K is sinusoidal (blue). As $m\to 0$, $K\to 1$ for all ϕ except an $O(m)$ region at $\phi=\pi/2$, where an $O(m^{-1})$ minimum is obtained. Right: the same plot as the left with rescaled axes illustrates the self-similar behaviour of $K(\phi,m)$ about its minimum. Dotted black plots the $m\to 0$ solution $mK(\phi,m)=-\pi(m/\theta)^2\exp(-m/\theta)$ found in the appendix
4.2	A plot illustrating the range of $K(\phi, m)$. Plotted are the ϕ -supremum and ϕ -infimum, $K_+(m)$ (red) and $K(m)$ (blue) respectively. The shaded grey region therefore is the range of $K(\phi, m)$. The dashed black line corresponds to the upper bound for Reynolds stress established in SY14.

- 4.4 Results from scattering experiments with $U(y) = 2\sin y$ and Q = 1.22, Z = 1.94 and three different values of F (corresponding to forcing wavenumber $k_f = (8,32,128)$ respectively, or F = (0.677,0.169,0.042) in the upper row, and F = (0.844,0.211,0.053) in the lower row). The top and bottom rows show the results for the wave forcings WF2 and WF3 respectively (see Figure 4.3 caption). The middle panels compare calculated and predicted $\langle u'v' \rangle$ across the full domain, and the right panels show a close-up of the situation near the east and west jet cores. The theoretical results (4.1.2, WB19, dashed purple), (4.2.15, SY14, red dotted line) and (4.2.18, $m \to \infty$, dashed blue) are also plotted. . . . 126

4.8	Linear stability analysis for the CE2 profile in Figure 4.6. The
	contour plot shows the maximum eigenmode growth rate kc_i as
	δeta and k are varied. Here, $\delta eta = \tilde{eta} - eta$ measures the deviation of
	the Coriolis parameter \tilde{eta} used for the stability analysis compared
	to the actual β value used in the CE2 simulation. Two distinct
	unstable regions are visible at $k = 3,4$ and $k = 6$, corresponding
	to instabilities at the westward and eastward jets as labelled. The
	unstable wavenumbers agree with those identified in Figure 4.6.
	The right-hand side line plots show the normalised momentum flux
	divergence (solid blue) of the most unstable mode at the points
	indicated by the green and blue dots. For reference the PV gradient
	$\tilde{\beta} - U_{yy}$ is also shown (dotted red)
4.9	As in Figure 4.4, but plotting $\langle u'u' \rangle$ instead of $\langle u'v' \rangle$
4.10	As in Figure 4.4, but plotting $\langle v'v' \rangle$ instead of $\langle u'v' \rangle$ 142
4.11	All Reynolds stress components from the SY14 local closure theory.
	The background flow is $U(y) = 2 \sin y$ and the forcing structure is
	WF3
4.12	Hovmöller plots for: (top panel) specified mean shear profile
	$\partial_y U(y,t) = \sin(\mu t)\cos(yt)$, and (bottom panels) momentum flux
	$\langle u'v' \rangle$ from CE2 simulation, time-dependent SY14 formula (4.5.4),
	and the instantaneous SY14 formula (4.2.10)
4.13	The results from Figure 4.12 at the latutude $y = 0$. The plots are
	local shear $U_y(0,t)$ (dotted red), and momentum flux $\langle u'v' \rangle$ from
	CE2 (blue), time-dependent SY14 (red) and instantaneous SY14
	(yellow)
4.14	As in Figure 4.12, but with the mean flow profile $U(y,t) = (2 +$
	$\sin(\mu t)\sin(yt)$. Hovmöller plots for: (top panel) specified mean
	shear profile $\partial U(y,t)$, and (bottom panels) momentum flux $\langle u'v' \rangle$
	from CE2 simulation, time-dependent SY14 formula (4.5.4), and the
	instantaneous SY14 formula (4.2.10)

4.15	The results from Figure 4.14 at the latutude $y = 0$. The plots are
	local shear $U_y(0,t)$ (dotted red), and momentum flux $\langle u'v' \rangle$ from
	CE2 (blue), time-dependent SY14 (red) and instantaneous SY14
	(yellow)
5.1	Left: $U(y)$ from CE2 (blue) and the SY14 parameterised model
	(purple marked) for the a channel flow linearly relaxed to $U_0 = \operatorname{erf}(y)$
	(red). Middle: similar to left, but plots the difference $U-U_0$. Right:
	the forcing profile $f(y)$. The forcing profile is also indicated on
	the two left-most panels with a blue gradient indicating the forcing
	magnitude
5.2	CE2 and SY14 σ quantities from the radiatively damped experiment.
	The top panel compares the mean wind profile with U_0 for CE2
	and SY14 σ (with the optimal $\sigma^* = 0.237/k_f$). The bottom panel
	investigates results in more detail by plotting the profile deviation
	from the radiatively relaxed profile, $U - U_0$. SY14 σ results are given
	for $\sigma = [0.125/k_f, \sigma^*, 0.5/k_f]$. In all cases $k_f = 16$
5.3	CE2 and SY14 σ quantities from the radiatively damped experiments
	A and D. The top panel shows the solution U from CE2 and SY14 σ
	when $\sigma = \sigma^*$. The bottom panels give a closer inspection of the
	solutions by plotting the difference with the equilibrium radiative
	solution, $U - U_0$. In these panels solutions with additional values of
	σ are also plotted ($\sigma = [0.125/k_f, \sigma^*, 0.5/k_f]$, in all cases $k_f = 16$). 16.
5.4	CE2 and SY14 σ quantities from the radiatively damped experiment
	E. The left panel plots U from CE2 and SY14 σ when $\sigma = \sigma^*$, and
	the right panel shows the difference with the equilibrium radiative
	solution, $U - U_0$. In the right panel solutions are presented for
	$\sigma = [0.125/k_f, \sigma^*, 0.5/k_f]$ (in all cases $k_f = 16$)

5.5	Left: snapshots of potential vorticity $\zeta + \beta y$ and relative vorticity
	ζ from the NL simulation reported in Figures 4.5 and 4.6. Right:
	snapshots of the same quantities from an SY14 σ simulation with
	the same parameters. In the SY14 σ model the stochastic forcing is
	replaced with the deterministic forcing term described in the main
	text ($\sigma = 0.237/k_f, k_f = 16$)
5.6	Hovmöller plot of the SY14 σ model for fully developed barotropic
	β -plane jets with $\sigma = 0.23/k_f$. Plotted is energy against time (to-
	tal energy (blue), mean flow energy (orange), eddy kinetic energy
	(yellow))
5.7	Wavenumber decomposition for momentum flux in the $SY14\sigma$
	model with $\sigma = 0.237/k_f$
5.8	Contour plots showing the spin-up of the $SY14\sigma$ simulation
	reported in Figure 5.5. Vorticity is contoured at times $t =$
	$[0.02, 0.04, 0.08, 10]/\mu$ (left to right, top to bottom) 169
5.9	As Figure 5.5, but with $\sigma = 0.5/k_f$ for the right hand panels 170
5.10	As Figure 5.6, but with $\sigma = 0.5/k_f$. Plotted is energy against time
	(total energy (blue), mean flow energy (orange), eddy kinetic energy
	(yellow))
5.11	As Figure 5.5, but with $\sigma = 0.125/k_f$ for the right hand panels 171
5.12	As Figure 5.6, but with $\sigma = 0.125/k_f$

List of Tables

1.1	Numerical parameters for the barotropic model	56
2.1	AMS results for two-to-three jet transitions in QG2D	76
2.2	DAMS results for one to two jet transitions in the two-layer model.	
	For each implementation the algorithm stopped when all members	
	transitioned $(r = N)$	90
5.1	Parameter setting references for the radiatively damped model	160

Chapter 1

Introductory Material

Arguably turbulence is the oldest and least satisfactorily resolved problem in fluid dynamics. Despite ubiquity over a wide-range of fluid scenarios from volcanic plumes to blood flow in the cardiovascular system, analytic theories of turbulence are incomplete or over-idealised. Exponential advances in computational power over the past few decades have provided an alternative line of attack, but numerical simulation encompassing enough scales of interest to fully account for turbulent dynamics is often out of reach (at least within reasonable computation time). Modelling the oceans and atmosphere, with scales of motion ranging from metres to hundreds of kilometres, is particularly susceptible to such shortcomings.

That said, significant contributions towards a turbulent theory were developed throughout the 20th century, predominantly beginning with Kolmogorov [1]. His inertial range theory is concerned with an idealised 3D turbulence that is statistically isotropic (the same in all directions), statistically homogenous (the same at all points) and has an external energy input at some length scale to sustain the turbulence. Supposing the Reynolds number is large enough so that the length scale of viscous dissipation is much shorter than the length scale of the energy injection scale, then there is an 'inertial range' of length scales between the two. For a length scale l_I in the inertial range, inertial range theory hypothesises that the statistics at l_I depend only on l_I and the energy injection rate. It is an assumption of locality in wavenumber space. Dimensional analysis yields that the energy spectrum function, in wavenumber space, is proportional to wavenumber to the power -5/3 in the

inertial range. Energy is fluxed towards small scales, producing a 'whorls become smaller whorls, and so on to viscosity' scenario [2].

Progress was made with 2D turbulence in the 1960s by Kraichnan [3] and Batchelor [4] predicting, and numerically verifying, a -5/3 energy spectra in 2D also. However, in 2D, energy is directed towards *large* length scales (referred to as the inverse cascade), with the -5/3 relationship holding on length scales larger than the energy injection scale. Energy is typically removed by a linear friction in 2D models at large length scales, rather than viscous dissipation. Enstrophy (half vorticity squared) is also conserved in 2D turbulence. It's predicted that a k^{-3} profile transfers enstrophy to smaller length scales than the injection scale, where it's removed by viscosity.

The upscale energy transfer of freely decaying 2D turbulence is shown in Figure 1.1 which gives snapshots of vorticity from a frictionless 2D simulation on a doubly periodic domain. Energy is conserved, but clearly occupies larger length scales as the simulation progresses.

The inertial range theories of turbulence are not a comprehensive theory of turbulence, however. In turbulence with a mean background flow, there may be a strong interaction between the turbulent scales and the mean flow directly, implying a non-local energy transfer and breaking the hypothesis of inertial range theories. In flows with a differential rotation, anisotropy is induced which can culminate in large scale zonal jets. A theory for anisotropic turbulence and mean-turbulent interactions are of key interest in planetary fluid dynamics, in particular for the gas giants Jupiter and Saturn [see e.g. 5, and refs. therein]. The photogenic jets tracing lines of longitude on Jupiter have been known for centuries (Saturn is known to have similar jets, though less visible). Despite its age, however, mathematical study of this phenomenon is relatively recent and is currently ongoing. Understanding the mechanisms at play in the simple models capturing the same essential physics is helpful, and perhaps necessary, as a means to understand various real-life observations or predictions from more advanced models [for instance 6, the recent unearthing of the surprising properties Jupiter's internal vertical structure].

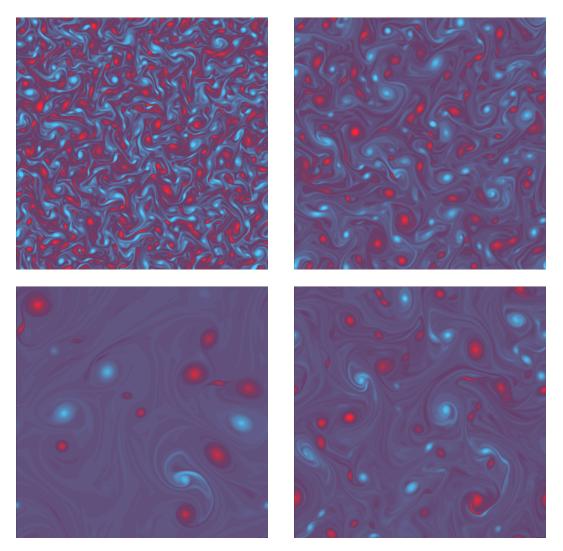


Figure 1.1: Snapshots of vorticity in freely decaying 2D turbulence in a doubly-periodic domain. Time increases clockwise from top-left, along with the energy containing length scales.

Perhaps the first directly relevant theoretical breakthrough for jets in rotating turbulence is the contribution from Rhines [7], where jet spacing is predicted to scale like $\sqrt{U/\beta}$ for a typical velocity U and differential rotation β . The scaling has been confirmed broadly across the across the subject literature since.

Vallis [8] extended the 2D inverse cascade idea to the beta-plane, finding a theoretical energy spectrum surface in 2D wavenumber space. Briefly, due to a mismatch of (anisotropic) Rossby wave period and turbulent eddy turnover time, particular low meridional wavenumbers aren't efficiently excited so the inverse energy cascade is diverted around an anisotropic shape on its journey to small

wavenumbers. The consequence is a disproportionate energy density on the lowest zonal wavenumbers.

In this Chapter two models to study zonal jets are introduced. First is the quasi-geostrophic 2D (QG2D) turbulence model. The model is an excellent test-bed for theories and is cheaper numerically than more complex models. The second model is a two-layer model, similar to the Phillips model [9], where the layers are relaxed towards an unstable vertical shear. Both models, as we shall see, are capable of producing turbulent-jet dynamics.

Starting with QG2D, both models will be derived and the important nondimensional parameters associated with jets will be discussed. Numerical results are provided after each model description. The last section details the methods and considerations of the numerical integrations for the models.

Both models will be relevant for Chapter 2, where rare event algorithms are investigated. In Chapters 3, 4 and 5, a theoretical approach to turbulence-mean interaction problems is studied, with attention being limited to the QG2D model.

1.1 Barotropic, Quasi-Geostrophic Equations

1.1.1 Derivation

The governing equation for the QG2D model is the two-dimensional barotropic vorticity equation on a β -plane, subject to stochastic forcing and a linear drag (1.1.21). QG2D has long been considered a key 'toy' model for understanding the physics of jet formation, in particular the long-lived zonal jets of Jupiter and Saturn. This section derives the equation from a more general perspective to introduce geophysical fluid dynamics principles and place the validity of the vorticity equation in a wider context.

To begin, assume an incompressible fluid in a thin layer with aspect ratio $\alpha \ll 1$. It is a standard exercise to derive the shallow water equations as an approximation to the Navier-Stokes equations in this scenario [29, p. 59-62]. The key steps are to show that the hydrostatic approximation holds at leading order in the vertical direction. This implies that horizontal accelerations are independent of the vertical

coordinate z, and the vertical momentum equation can be depth integrated for a mass conservation equation. The rotating shallow water equations over a flat topography are found as

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \mathbf{f} \times \mathbf{u} = -g \nabla \delta, \tag{1.1.1}$$

$$\frac{\partial \delta}{\partial t} + \nabla \cdot (h\mathbf{u}) = 0. \tag{1.1.2}$$

Here \boldsymbol{u} is the horizontal velocity vector, g is gravity, H is the mean fluid depth, δ is the variation of the fluid surface from H (so $h = H + \delta$ is total depth), and \boldsymbol{f} is a background rotation.

Assuming length scales of motion are smaller than the radius of the planet allows the β -plane approximation

$$\mathbf{f} = f\hat{\mathbf{k}} = (f_0 + \beta y)\hat{\mathbf{k}},$$

where $f_0 = 2\Omega \sin \theta$ and $\beta = 2(\Omega/a)\cos \theta$ for Ω the angular velocity of the planet.

To derive the single layer quasi-geostrophic equations we need to highlight some significant scales, non-dimensionalise accordingly and apply an asymptotic approach.

If U is a typical velocity and L is a typical length scale, one of the most fundamental quantities in geophysical fluid dynamics is the Rossby number

$$Ro = \frac{U}{fL},\tag{1.1.3}$$

which compares the strengths of the advective and Coriolis terms. On Earth, $Ro \ll 1$ in the ocean and atmosphere. The most severe approximation from small Ro is to ignore the advection terms entirely which gives a relation between the Coriolis and pressure terms called *geostrophic balance*. Geostrophic balance indicates the scaling for the pressure term is

$$\delta \sim \frac{fUL}{g},$$
 (1.1.4)

which may also be written as

$$\delta \sim RoH \frac{L^2}{L_D^2},\tag{1.1.5}$$

where the Rossby radius of deformation,

$$L_D = \frac{\sqrt{gH}}{f},\tag{1.1.6}$$

has been introduced. It's a known result that \sqrt{gH} is a scaling for the speed of surface waves in the shallow water equations, so L_D can be thought of as a ratio of wave speed to planetary rotation.

Letting a prime denote non-dimensional variables, we non-dimensionalise using (x,y)=L(x',y'), $\mathbf{u}=U\mathbf{u}'$, t=L/Ut', $\delta=Rof^2L^2/g\delta'$, $\beta=U/L^2\beta'$ and $f=f_0(f_0'+Ro\beta'y')$. The momentum and mass equations (1.1.1)-(1.1.2) become

$$Ro\left\{\frac{\partial \mathbf{u}'}{\partial t'} + \left(\mathbf{u}' \cdot \nabla'\right) \mathbf{u}'\right\} - \mathbf{f}' \times \mathbf{u}' = -\nabla' \delta', \tag{1.1.7}$$

$$Ro\frac{L}{L_D} \left\{ \frac{\partial \delta'}{\partial t'} + \boldsymbol{u}' \cdot \nabla' \delta' \right\} + \nabla' \cdot \boldsymbol{u}' = 0, \tag{1.1.8}$$

Now looking for solutions as an expansion in Rossby number

$$\mathbf{u}' = \mathbf{u}_0' + Ro\mathbf{u}_1' + Ro^2\mathbf{u}_2' + \dots, \tag{1.1.9}$$

$$\delta' = \delta_0' + Ro\delta_1' + Ro^2\delta_2' + \dots, \tag{1.1.10}$$

the momentum equation at leading order simply gives geostrophic balance. The two equations can be cross differentiated to cancel δ'_0 getting $\nabla' \cdot \boldsymbol{u}'_0 = 0$, which is also the conclusion of the mass equation at leading order. The mass equation at the next order introduces \boldsymbol{u}_1 , viz.

$$\frac{L}{L_D} \left\{ \frac{\partial \delta_0'}{\partial t'} + \boldsymbol{u}_0' \cdot \nabla' \delta_0' \right\} + \nabla' \boldsymbol{u}_1' = 0, \tag{1.1.11}$$

as does the same order from the momentum balance

$$\frac{\partial \mathbf{u}_0'}{\partial t'} + (\mathbf{u}_0' \cdot \nabla') \mathbf{u}_0' - \beta' y' \mathbf{k} \times \mathbf{u}_0' - f_0' \mathbf{k} \times \mathbf{u}_1' = -\nabla' \delta_1'. \tag{1.1.12}$$

The δ_1 term is eliminated by taking the curl here, getting the vorticity equation

$$\frac{\partial q_0'}{\partial t'} + \left(\mathbf{u}_0' \cdot \nabla'\right) \left(\zeta_0' + \beta' y'\right) = -f_0' \nabla' \cdot \mathbf{u}_1' \tag{1.1.13}$$

where $\zeta' = v_x' - u_y'$ is two-dimensional vorticity (and its subscripts here denote its orders as expected). The associated streamfunction ψ_0' satisfies $\nabla^2 \psi_0' = \zeta_0'$ and $(u_0', v_0') = (-\partial_y, \partial_x) \psi_0'$. Also geostrophic balance implies $\psi_0' = \delta_0' / f_0'$. Eliminating u_1' using (1.1.11) in the above finds

$$\frac{\partial}{\partial t} \left\{ \nabla^2 \psi_0' - \frac{L^2}{L_D^2} f_0'^2 \psi_0' \right\} + \left(\mathbf{u}_0' \cdot \nabla' \right) \left\{ \nabla^2 \psi_0' + \beta' y' - f_0'^2 \frac{L^2}{L_D^2} \psi_0' \right\} = 0. \quad (1.1.14)$$

Dimensions may be reintroduced for the shallow-water quasi-geostrophic equation

$$\frac{D}{Dt}\left(\nabla^2 \psi + \beta y - \frac{1}{L_D^2}\psi\right) = 0, \qquad (1.1.15)$$

where $D/Dt = \partial_t + \boldsymbol{u} \cdot \nabla$. The term in the brackets is the quasi-geostrophic shallow water *potential vorticity*. The ψ/L_D^2 term captures the effects of 3D vortex stretching in a 2D setting. As a final approximation, this term is neglected as a result of considering motions on scales smaller than the deformation radius. For 2D vorticity $\zeta = v_x - u_y$, then, we have

$$\zeta_t - \psi_y \zeta_x + \psi_x \zeta_y + \beta \psi_x = 0, \qquad (1.1.16)$$

with the streamfunction relation

$$\zeta = \nabla^2 \psi. \tag{1.1.17}$$

Velocity components are

$$u = -\psi_y, \quad v = \psi_x, \tag{1.1.18}$$

which satisfy mass conservation

$$\nabla \cdot \boldsymbol{u} = u_x + v_y = 0. \tag{1.1.19}$$

The 2D potential vorticity, q, is simply the absolute vorticity (composed of the relative and background planetary vorticity)

$$q = \zeta + \beta y. \tag{1.1.20}$$

The QG2D model is a modified version of equation (1.1.16) that includes forcing, drag and diffusion terms. The full equation, which has long been considered a toy-model for atmospheric jets, is

$$\zeta_t + J(\psi, \zeta) + \beta \psi_x = -\mu \zeta + (-1)^{n+1} \nu_{2n} \nabla^{2n} \zeta + \sqrt{\varepsilon} \eta. \tag{1.1.21}$$

The notation for the Jacobian operator $J(\cdot,\cdot)$ has been introduced here, defined by

$$J(\rho, \sigma) := \rho_x \sigma_y - \rho_y \sigma_x, \tag{1.1.22}$$

i.e. $J(\psi, \zeta) = \nabla \cdot (u\zeta)$. This term contains the nonlinearity, and consequently much of the difficulty involved in understanding QG2D. The final $\sqrt{\varepsilon}\eta$ term represents a stochastic forcing on the flow which is detailed in section 1.1.4. For now it's enough to state the forcing term supplies energy to the system and ε is the energy injection rate. The $-\mu\zeta$ term appearing is referred to as a drag or friction term which is frequently used in this model. Besides the convenience, the linearity is also a theoretical parameterisation of the damping imposed on a flow from the basal Ekman layer. Lastly, a hyper-diffusion term is included as a means of viscous dissipation.

For simplicity QG2D uses a square domain \mathcal{D} with dimensions $2\pi L_d \times 2\pi L_d$ and the boundary conditions are doubly-periodic.

1.1.2 Energy and Enstrophy

The spatially averaged kinetic energy over the entire domain has various representations using standard integral theorems, e.g.

$$\mathcal{E} = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} \frac{1}{2} \left(u^2 + v^2 \right) = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} \frac{1}{2} |\nabla \psi|^2 = \frac{-1}{|\mathcal{D}|} \int_{\mathcal{D}} \frac{1}{2} \psi \zeta. \tag{1.1.23}$$

Employing the rightmost expression, an energy budget is found by multiplying (1.1.21) by $-\psi$ and taking the domain mean. In addition, taking the expectation over forcing realisations finds the equilibrium energy budget

$$\frac{\mathrm{d}\mathcal{E}}{\mathrm{d}t} = -2\mu\mathcal{E} + \varepsilon - D. \tag{1.1.24}$$

The final term is the energy sink from hyper-diffusion

$$D = \frac{v_{2n}}{|\mathcal{D}|} \int_{\mathcal{D}} \mathbb{E}\left[|\nabla \psi|^n\right], \tag{1.1.25}$$

which clearly always removes energy. Supposing v_{2n} is such that $D \ll \varepsilon$, neglecting D in the energy balance gets the equilibrium expected energy as

$$\mathcal{E} = \frac{\varepsilon}{2\mu}.\tag{1.1.26}$$

Similar to \mathcal{E} , another quadratic of interest in QG2D is the domain averaged enstrophy, \mathcal{Z} , defined as

$$\mathcal{Z} = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} \frac{1}{2} \zeta^2. \tag{1.1.27}$$

Like energy, enstrophy is provided by η . However, in contrast, enstrophy is cascaded *downscale*. The discretisation of numerical simulations sets a smallest resolved scale of motion. Without a strong enough diffusion at small scales simulations are prone to enstrophy 'pile-up' at the smallest resolved scale. The hyper-diffusion must be calibrated carefully to remove enstrophy before pile-up, but have a weak effect on the larger scales of motion.

1.1.3 Fourier-Space and Quadratic Cascades

To continue discussion it is necessary to the introduce the Fourier transform of a function f(x) defined on domain $[0, 2\pi L_d)$ to be

$$f_k = \frac{1}{2\pi L_d} \int_0^{2\pi L_d} f(x)e^{-ikx} dx, \qquad f(x) = \sum_k f_k e^{ikx}.$$
 (1.1.28)

Note that k isn't necessarily integer here as it depends on the domain size. Precisely, k takes values n/L_d for $n \in \mathbb{Z}$. The n-dimensional Fourier transform on some domain \mathcal{D} is simply an application of the above transform on each dimension,

$$f_{\mathbf{k}} = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} f(x, y) e^{-i\mathbf{k} \cdot \mathbf{x}} d\mathbf{x}, \qquad f(\mathbf{x}) = \sum_{\mathbf{k}} f_{\mathbf{k}} e^{i\mathbf{k} \cdot \mathbf{x}}.$$
 (1.1.29)

This specification of the Fourier transform is designed so that domain averaged quantities are found simply as the zeroth mode, i.e.

$$\frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} f(\mathbf{x}) d\mathbf{x} = \frac{1}{|\mathcal{D}|} \int_{\mathcal{D}} f_{\mathbf{0}} d\mathbf{x} = f_{\mathbf{0}}.$$
 (1.1.30)

Supposing $f(\mathbf{x})$ is a real function, so that $f = f^*$, where an asterisk (*) denotes the complex conjugate, we must have

$$\sum_{k} f_{k} e^{ik \cdot x} = \sum_{k} f_{k}^{*} e^{-ik \cdot x} = \sum_{k} f_{-k}^{*} e^{ik \cdot x}.$$
 (1.1.31)

This yields the *conjugate symmetry* condition

$$f_{k} = f_{-k}^{*}. (1.1.32)$$

Fourier space provides the natural environment to consider the turbulent energy/enstrophy cascades discussed in the introduction. Now we briefly give an argument for an inverse energy cascade in Fourier space [e.g. Vallis 10, pgs. 370-371]. Defining a characteristic wavelength for the energy as

$$k_{\mathcal{E}} = \frac{\int |\mathbf{k}| \mathcal{E}_{\mathbf{k}} d\mathbf{k}}{\int \mathcal{E}_{\mathbf{k}}}, \tag{1.1.33}$$

a measure of the energy distribution's width can be found using

$$I = \int (|\mathbf{k}| - k_{\mathcal{E}})^2 \mathcal{E}_{\mathbf{k}} d\mathbf{k}, \qquad (1.1.34)$$

$$= \int |\mathbf{k}|^2 \mathcal{E}_{\mathbf{k}} d\mathbf{k} - 2k_{\mathcal{E}} \int |\mathbf{k}| \mathcal{E}_{\mathbf{k}} d\mathbf{k} + k_{\mathcal{E}}^2 \int \mathcal{E}_{\mathbf{k}} d\mathbf{k}, \qquad (1.1.35)$$

$$= \mathcal{Z} - k_{\mathcal{E}}^2 \mathcal{E}. \tag{1.1.36}$$

The energy budget (1.1.24) in the absence of forcing, friction and viscous dissipation implies conservation of energy (similarly, enstrophy is also conserved). Under the natural assumption that the energy distribution spreads out from some initial distribution,

$$0 < \frac{\mathrm{d}I}{\mathrm{d}t} = -2k_{\mathcal{E}} \frac{\mathrm{d}k_{\mathcal{E}}}{\mathrm{d}t} \mathcal{E},\tag{1.1.37}$$

implying $\mathbf{k}_{\mathcal{E}}$ decreases; energy inverse cascades towards smaller wavenumbers (larger physical length scales). By the vorticity-streamfunction relation (1.1.17), it follows that energy and enstrophy are related via $\mathcal{Z}_{\mathbf{k}} = \mathbf{k}^2 \mathcal{E}_{\mathbf{k}}$, and a similar argument as for the energy cascade is easily applicable to enstrophy, finding a forward cascade.

In QG2D, due to the random forcing, there isn't an exact energy conservation, but the same idea applies – energy is injected at a characteristic scale and transferred upscale before significant energy loss to friction.

1.1.4 Stochastic Forcing

1.1.4.1 Overview

The final term in (1.1.21) is a means of supplying energy. Without such a forcing, the energy budget (1.1.24) implies that any energy in the system from initial conditions will tend to zero due to the drag and viscous terms. The stochastic forcing term may be viewed as an attempt to model the energy injection from baroclinic instability on the barotropic mode (see section 1.2), or as localised injections of vorticity modelling moist convection penetrating into the modelled layer of fluid. Without explicitly solving any internal or external processes, it is typical to turn to a stochastic guess. The focus in the bulk of the literature adopting the QG2D model isn't on the details of the stochastic forcing as it's not considered to have a significant effect on the larger

scales of 2D turbulence such as jets [11] (this is a large topic of debate in Chapters 4 and 5 of this thesis). The only real requirement is usually that the characteristic forcing length scale, L_f , is taken to be relatively small.

The stochastic forcing term is defined through spatial and temporal correlations. The most common-place stochastic forcing for QG2D is delta-correlated in time. Essentially the forcing has no 'memory', or is *white*. The spatial correlation for η is taken to be described fully by its two-point correlation; the forcing is *Gaussian*. The two-point correlation is usually chosen to be an annulus in Fourier space since this provides a homogenous isotropic forcing centred on a wavelength of choice.

Taking the Fourier transform of the vorticity equation (1.1.21) yields

$$\frac{\mathrm{d}\zeta_{k}}{\mathrm{d}t} = \left\{ \text{deterministic terms} \right\}_{k} + \sqrt{\varepsilon}\eta_{k} \tag{1.1.38}$$

so we have a set of complex SPDEs

$$d\zeta_{\mathbf{k}} = \{\text{deterministic terms}\}_{\mathbf{k}} dt + \sqrt{\varepsilon} \gamma_{\mathbf{k}} dW_{\mathbf{k}}. \tag{1.1.39}$$

where η_k (and hence η) is defined through γ_k of our choosing. The dW_k are complex, i.e.

$$dW_{\mathbf{k}} = dW_{\mathbf{k}}^{R} + idW_{\mathbf{k}}^{I}, \tag{1.1.40}$$

with both W_k^R and W_k^I being independent standard Brownian processes. Conjugate symmetry to keep $\eta(\mathbf{x})$ real requires

$$\eta_{\mathbf{k}} = \eta_{-\mathbf{k}}^*, \tag{1.1.41}$$

thus we take

$$dW_{\boldsymbol{k}} = (dW_{-\boldsymbol{k}})^*, \qquad (1.1.42)$$

and

$$\gamma_{k} = \gamma_{-k}^{*}. \tag{1.1.43}$$

One consequence of the complex variables to be aware of is the unusual result that

$$\mathbb{E}\left[dW_{\mathbf{k}}dW_{\mathbf{j}}\right] = \begin{cases} 2dt, & \text{if } \mathbf{j} = -\mathbf{k}, \\ 0, & \text{otherwise.} \end{cases}$$
 (1.1.44)

Defining η in this manner implies it is a white-in-time stochastic noise defined *entirely* by it's second order statistic

$$\Pi(\mathbf{x}, \mathbf{x}') = \mathbb{E}\left[\eta(\mathbf{x})\eta(\mathbf{x}')\right]. \tag{1.1.45}$$

Here the expectation is over realisations of η . The covariance Π fully defines the statistics of η .

Note that

$$\Pi(\mathbf{x}, \mathbf{x}') = \mathbb{E} \left[\eta \left(\mathbf{x}, t \right) \eta \left(\mathbf{x}', t' \right) \right],$$

$$= \sum_{\mathbf{k}, \mathbf{j}} 2 \delta_{\mathbf{k}, -\mathbf{j}} \gamma_{\mathbf{k}} \gamma_{\mathbf{j}} e^{\mathrm{i}(\mathbf{k} \cdot \mathbf{x} + \mathbf{j} \cdot \mathbf{x}')},$$

$$= \sum_{\mathbf{k}} 2 \gamma_{\mathbf{k}} \gamma_{-\mathbf{k}} e^{\mathrm{i}(\mathbf{k} \cdot (\mathbf{x} - \mathbf{x}'))},$$

$$= \sum_{\mathbf{k}} 2 |\gamma_{\mathbf{k}}|^2 e^{\mathrm{i}(\mathbf{k} \cdot (\mathbf{x} - \mathbf{x}'))},$$

$$= \Pi(\mathbf{x} - \mathbf{x}').$$
(1.1.46)

implying that Π a function of only the vector difference between the two points, *regardless* of where they are situated. This means the forcing is homogenous.

Moreover, it is clearly seen that the two-dimensional Fourier coefficients of Π are

$$\Pi_{\boldsymbol{k}} = 2|\gamma_{\boldsymbol{k}}|^2. \tag{1.1.47}$$

1.1.4.2 Isotropic Forcing

For isotropic forcing one must go beyond homogeneity and also restrict Π to a function of only the magnitude of x - x'. For this it is required that γ_k is dependent only on |k| (i.e γ is radially symmetric in wavenumber space).

Then the expression above – in it's continuous limit where the wave vector discretisation effects are ignored – can be manipulated using a change of coordinates to a polar system (r, θ) with $\theta = 0$ corresponding to the x - x' direction as

$$\Pi(|\mathbf{x} - \mathbf{x}'|) = 2 \int_0^\infty \int_0^{2\pi} |\gamma_r|^2 e^{ir|\mathbf{x} - \mathbf{x}'|\cos\theta} r dr d\theta.$$
 (1.1.48)

Using the fact that $\int_0^{2\pi} \exp(iz\cos\theta) d\theta = 2\pi J_0(z)$ for J_0 the zeroth order Bessel function of the first kind, this can be written

$$\Pi(|\mathbf{x} - \mathbf{x}'|) = 4\pi \int_0^\infty r |\gamma_r|^2 J_0(r|\mathbf{x} - \mathbf{x}'|) dr.$$
 (1.1.49)

1.1.4.3 Energy Normalisation Constraint

A calculation for the expected domain averaged energy injection rate from the forcing, is required in order to normalise η appropriately. The constraint is

$$\varepsilon = \varepsilon \sum_{\mathbf{k}} \frac{|\gamma_{\mathbf{k}}|^2}{|\mathbf{k}|^2}.$$
 (1.1.50)

This means that η is independent of ε , with the normalisation contraint

$$\sum_{\mathbf{k}} \frac{|\gamma_{\mathbf{k}}|^2}{|\mathbf{k}|^2} = 1. \tag{1.1.51}$$

Equivalently,

$$\sum_{k} \frac{\Pi_k}{2|k|^2} = 1. \tag{1.1.52}$$

1.1.5 Non-dimensional Parameters for Jet Dynamics

Rather than investigating the change of dynamics over the $(\beta, \mu, \varepsilon, L_d, L_f)$ parameter space, it's more sensible to consider dimensionless parameters which control important features of jets: strength, stability, number of jets etc. The reason is that similar jet dynamics may occur across some lower-dimensional surface within the parameter space.

Key Dimensional Parameters

It is not possible to discuss non-dimensional parameters without introducing two important dimensional quantities first. First, the Rhines length scale [7], defined as

$$L_{\mathrm{Rh}} \sim \sqrt{\frac{U}{\beta}},$$
 (1.1.53)

has been widely found to predict the spacing of the barotropic jets (up to a constant prefactor). Here, U is a typical root mean square velocity. Essentially L_{Rh} this is the cross-over length scale balancing beta and advection terms in the vorticity equation.

A second notable dimensional quantitity is the estimated length scale at which freely decaying turbulence starts to 'feel' planetary rotation. By balancing the frequency of freely decaying turbulence ($\varepsilon^{1/3}k^{2/3}$) from Kolmogorov's dimensional analysis with Rossby wave frequency (β/k), the 'anisotropic length scale' is [10]

$$L_{\beta} = \left(\frac{\varepsilon}{\beta^3}\right)^{1/5}.\tag{1.1.54}$$

Non-dimensional Parameters

A quantisation parameter, Q, which gives an indicator of the expected number of jets in the domain is simply defined for the doubly periodic domain as the meridional domain length divided by the Rhines scale

$$Q = \frac{L_d}{L_{\rm Ph}}.\tag{1.1.55}$$

The ratio of the Rhines scale, $L_{\rm Rh}$, to the 'anisotropic length scale', L_{β} , gives the *zonostrophy* of the flow

$$Z = \frac{L_{\rm Rh}}{L_{\beta}} = \frac{\beta^{1/10} \varepsilon^{1/20}}{\mu^{1/4}}.$$
 (1.1.56)

Zonostrophy was first introduced by [12], and subsequently refined to the above form in [13] and [14]. The idea is that if Z is large $(L_{\rm Rh} \gg L_{\beta})$, turbulence will organise itself anisotropically at length scales smaller than the jet scale. Conversely, if Z is small then it is not clear that jets will form since the turbulence will still be approximately isotropic at the Rhines scale. Consequently Z is considered an

indicator of jet strength and stability – something our numerical experiments are in keeping with as well as the literature [for a beautiful demonstration see 15].

Notice that when formulating L_{β} above, the Rossby wave frequency was taken as k/β which is an *isotropic* approximation to *anisotropic* Rossby wave frequency $-\beta k_x/|\mathbf{k}|^2$. An illuminating theory for zonal jets is put forward by [8] using the anisotropic Rossby frequency to derive an anisotropic wave-turbulence boundary. Setting the anisotropic Rossby wave and turbulence frequencies as equal (as in the derivation of L_{β}), gives an anisotropic boundary $\varepsilon^{1/3} |\mathbf{k}|^{2/3} = -\beta k_x/|\mathbf{k}|^2$. This boundary is a dumbbell shape in Fourier space, oriented so that the two 'centres' are on the k_x axis. The theory is that turbulence cannot cross this boundary easily due the frequency mis-match, so the boundary acts to funnel the inverse energy cascade onto the k_v axis creating (zonally oriented) anisotropy. Whilst this inverse cascade type theory of spontaneous jet formation is attractive and is backed up numerically for fairly low zonostrophy flows, the past decade or two have made it clear that, in high zonostrophy, non-local interactions between the turbulence and the zonal mean field become more important than the classic local cascade interactions [see for instance 16]. A key point to realise here is that the Buckingham π theorem reveals that there is actually only one non-dimensional parameter which can be formulated from the parameters ε , μ and β . So regardless if the assumptions are valid or not for some particular derivation of Z, it inevitably arrises as an important non-dimensional parameter.

Simulating a large zonostrophy is perhaps best achieved through reducing μ and ε by equal proportion. However, enstrophy needs to be removed at the small scales whilst keeping the viscous term negligible with respect to the forcing, hence increasing zonostrophy will typically require an increase in resolution and computational cost. An O(10) zonostrophy is achieved in [15] creating extremely robust jets. Flows with $Z\gg 1$ are called 'zonostrophic'.

Another non-dimensional parameter which we expect to be important in QG2D is the forcing number

$$F = \frac{L_f}{L_{\rm Rh}},\tag{1.1.57}$$

which is the ratio of the characteristic forcing length scale, L_f , to the Rhines scale $L_{\rm Rh}$ (jet length scale). It can be thought of loosely as a measure of how many stochastically forced eddies fit between adjacent zonal jet tips.

In summary, the important non-dimensional parameters are zonostrophy, Z, jet quantisation Q, and the forcing number F. The relevant regime for strong and steady jets is $Z \gg 1$ and $F \ll 1$, and naturally we are interested in a Q value which will produce at least a single jet on the domain.

Time scale separation

A final dimensionless quantity that is of interest and is prominent in the literature, is a timescale separation between the eddy turnover time and the zonal timescale

$$\alpha = \frac{\tau_{\text{adv}}}{\tau_{\text{drag}}}.$$
 (1.1.58)

Following the non-dimensionalisation procedure in [17], QG2D is considered on the dimensionless domain $[0,2\pi)\times[0,2\pi)$ with a theoretical mean total equilibrium energy on the domain of 1. This is achieved by taking the advective timescale as

$$\tau_{\text{adv}} = \frac{L_d}{2\pi\sqrt{\mathcal{E}}} = \frac{L_d}{2\pi}\sqrt{\frac{2\mu}{\varepsilon}},\tag{1.1.59}$$

where $\mathcal{E}=\varepsilon/2\mu$ is the (dimensional) expected energy per unit area. The drag timescale is μ^{-1} so it follows that

$$\alpha = \mu \tau_{\text{adv}} = \frac{L_d}{2\pi} \sqrt{\frac{2\mu^3}{\varepsilon}}.$$
 (1.1.60)

Non-dimensionalising the vorticity equation (1.1.21) with length and time scales L_d and τ_{adv} respectively gets

$$\zeta'_{t'} + J(\psi', \zeta') = -\beta' \psi'_{x'} - \alpha \zeta' + \sqrt{2\alpha} \eta' - v'_{2n} (-\nabla^2)^n \zeta', \tag{1.1.61}$$

where $\beta' = L_d \tau_{\text{adv}} \beta$ and $\eta' = 2\pi L_d \tau_{\text{adv}}^{1/2} \eta$. The 2π factor here (and in τ_{adv} above) is to maintain consistency with Bouchet et al. ([17], [18]). Essentially the forcing term η'

is re-normalised to have unit energy injection rate on the *entire* non-dimensionalised domain, rather than per unit area. It is then observed from (1.1.61) that the non-dimensional expected energy over the entire domain is one (c.f. $\mathcal{E} = \varepsilon/2\mu$). Writing the equation in this form, particularly with the $\sqrt{2}$ factor, links QG2D to Langevin dynamics.

Within the timescale separation framework, the model is viewed as having two dimensionless parameters (α, β') . These are found from the dimensional quantities $(L_d, \beta, \varepsilon, \mu)$ via

$$\alpha = \frac{L_d}{2\pi} \left(\frac{2\mu^3}{\varepsilon}\right)^{1/2},\tag{1.1.62}$$

$$\beta' = \frac{L_d^2 \beta}{2\pi} \left(\frac{2\mu}{\varepsilon}\right)^{1/2}.$$
 (1.1.63)

Notice we can write the non-dimensionalised beta parameter as proportional to the square of the ratio of domain length to Rhines scale

$$\beta' = \frac{L_d^2}{2\pi \left(\mathcal{E}^{1/2}/\beta\right)} = \frac{\sqrt{2}}{2\pi} \left(\frac{L_d}{L_{\rm Rh}}\right)^2$$
 (1.1.64)

predicting that the quantisation parameter Q scales according to

$$Q \propto \beta'^{1/2},\tag{1.1.65}$$

hence only one of Q or β' needs to be specified. We can also write the zonostrophy (1.1.56) in the non-dimensional variables

$$Z = \left(\frac{\sqrt{2}}{2\pi} \frac{\beta'}{\alpha^2}\right)^{1/10}.\tag{1.1.66}$$

This gives the peculiar relationship $Z \propto (Q/\alpha)^{1/5}$, which is not overly insightful but it makes the point that only two of Z, Q and α actually *need* to be specified. In the remainder of this Chapter, and in Chapter 2, the non-dimensionalised form (1.1.61) of QG2D is used, and the prime is dropped from β' with the understanding that β is

no the non-dimensionalised Coriolis parameter defined above.

1.1.6 Model Integrations

In this section some example results from numerical simulations that capture the essence of the barotropic model are reported. The specifics of the numerical implementation are detailed in section 1.3.

Firstly, simulations with the same parameters and initial conditions are not identical since the model is determined by a realisation of the stochastic forcing which is different for each simulation. However, over all repeated simulations where the only difference is the forcing realisations, the jet existence (or lack thereof) has been consistent - no flipping between a dipole and jet state is found as in [19], for instance. This is due to the dynamics being controlled by the non-dimensional parameters described in 1.1.5, in particular the zonostrophy parameter (of which [19] does not have, since $\beta = 0$ in that problem).

An example of typical jet formation from spin-up (zero initial energy) is shown in Figure 1.2. The top panel shows a space-time (or Hovmöller) plot of zonal velocity. The remaining panels are a zonal velocity plot and the vorticity, energy and enstrophy fields at the end of the simulation.

We can see that as the energy increases from zero to equilibrium, the number of jets decreases. The Rhines scale for jets should really be considered as an equilibrium theory, but the decreasing jet number can be loosely interpreted as the jet scaling being related to energy [20] (c.f. $L_{\rm Rh} \sim U^{1/2} \sim \mathcal{E}^{1/4}$). The simulation has been run beyond the time when equilibrium energy has been reached and the three jets appear reasonably well established. Changes in the number of jets are possible even once equilibrium energy is obtained. These are referred to as *transitions* and they are the subject of Chapter 2.

The zonal velocity panel shows that the east and west jets are asymmetric. East jets are sharp, while the west jets are more rounded with a typical maximum curvature β . The asymmetry is discussed in more depth in Chapter 4, but may be viewed several ways. For instance, the Rayleigh-Kuo necessary condition for instability in 2D flow is that the potential vorticity gradient $U_{yy} - \beta$ changes sign

in the domain [21]. Since the potential vorticity gradient is more liable to change sign at westward jets, the idea that the growth of westward jets is halted due to a counteracting instability mechanism has been suggested [22, 23]. Additionally, a simple but powerful idea which captures the profile asymmetry is the staircase limit ([24], [15]). The idea is that the potential vorticity field mixes within each 'layer' between the cores of the east jets. The conclusion from a perfect mixing is a staircase potential vorticity profile; equivalently, a parabolic profile for zonal velocity, i.e. $U_{yy} = \beta$, is predicted from perfect potential vorticity mixing. It is observed that the combination of a staircase profile and the assumption that the zonal flow contains the majority of energy provides a simple constraint on the jet width – it is a simple exercise to show this width scales like $L_{\rm Rh}$.

The contour plots of energy and enstrophy illustrate their tendency to the large and the small scales respectively. Unsurprisingly, the bulk of energy is located at the jet cores (both east and west) where the fluid velocity is typically largest. Enstrophy isn't restricted to the jet cores and large contributions are due to small eddies and filaments which can be anywhere in the flow.

Results from a parameter sweep over (α, β) are shown in Figure 1.3. The results clearly show coherent jet formation and maintenance over a large time scale. A large time scale in this context is taken as multiple α^{-1} non-dimensional time units. The values for β increase such that $\sqrt{\beta}$ is proportional to (2,3,4,5); it is observed that the jet number roughly follows these values, confirming the Rhines scale is inversely proportional to $\sqrt{\beta}$.

The jets appear stronger with larger zonostrophy in most cases. This agrees with the discussion in section 1.1.5. The simulations in the right-most column with $\beta = 12.5$ are the obvious exceptions. All of these simulations have an approximate jet number of five, highlighting that the forcing scale to jet width, F, does become important to the dynamics when it is not sufficiently small. We do not worry about this since the simulations provided in the parameter sweep were run at a relatively low resolution, and hence a relatively small k_f (large forcing scale) was chosen. In hindsight, the choice of $k_f = 10$ was smaller than necessary, however it's

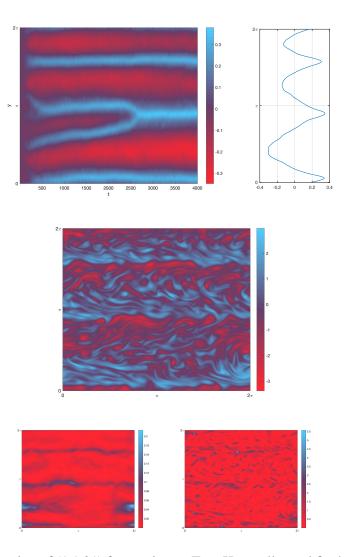


Figure 1.2: Integration of (1.1.21) from spin-up. Top: Hovmöller and final zonal velocity. Middle: final vorticity. Bottom: final energy enstrophy. Parameters: $\beta = 5.26$, $\mu = 0.0012$, $\varepsilon = 0.0024$, $k_f = 14.5$ with an annulus half width of 0.6. Numerical resolution is 256×256 and $v_4 = 2.5 \times 10^{-8}$.

encouraging to show that the expected behaviour is followed for the two to four jet state simulations in spite of this. All our interest will be in a jet number of three or less in subsequent chapters and k_f will also be larger.

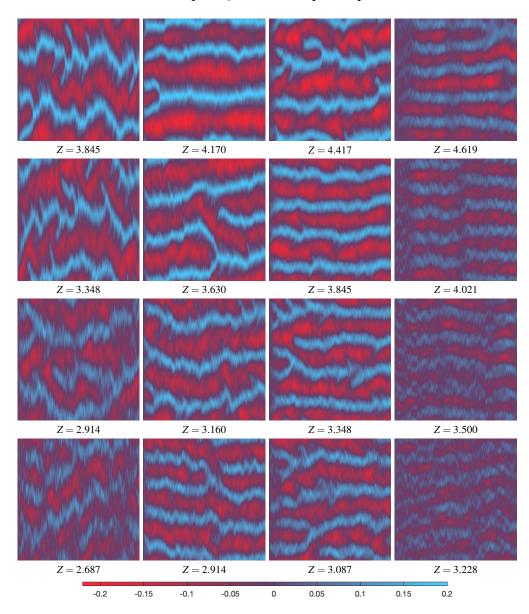


Figure 1.3: Hovmöller contour plots for zonal velocity against time in QG2D simulations. Top to bottom: $\alpha = (2, 4, 8, 12) \times 10^{-3}$. Left to right $\beta = [2^2, 3^2, 4^2, 5^2]/2$ Zonostrophy values according to (1.1.66) given in each case. Even these low resolution 128×128 , $k_f = 10$ give the expected scalings. All simulations are run for $50\alpha^{-1}$ non-dimensional time units.

1.2 Two-Layer Baroclinic, Quasi-Geostrophic Equations

1.2.1 Derivation

In this section we derive the two-layer model. The governing equations are given in (1.2.15)-(1.2.16) and the specifics of the model set-up are in the following section.

For a two-layer model we assume a set-up of two layers of fluid in which each have a constant density and mean depth H. Motions are three dimensional and the interface between the layers can move. All variables are indexed with a subscript j=1,2 to denote the upper layer (1) the lower layer (2). i.e $\mathbf{u}_1=(u_1,v_1,w_1)^T$ is the velocity in the upper layer. The total depth of each layer is h_j and δ_j is the displacement of each layer about its mean height. Density and pressure are denoted ρ_j and p_j respectively. As in QG2D, we make the β -plane approximation for the Coriolis parameter $\mathbf{f}=(f_0+\beta y)\hat{\mathbf{k}}$.

The Rossby number is identical to the shallow water definition, and as before we let

$$Ro = \frac{U}{f_0 L} \ll 1,\tag{1.2.1}$$

where U and L are a typical velocity flow speed and length.

For a stable stratification under gravity, $\rho_1 < \rho_2$, and we define

$$\rho^+ = \frac{\rho_1 + \rho_2}{2},\tag{1.2.2}$$

$$\Delta = \frac{\rho_2 - \rho_1}{\rho^+}.\tag{1.2.3}$$

The Boussinesq approximation, $\Delta \ll 1$, implies a reduced gravity

$$g' = \Delta g, \tag{1.2.4}$$

which acts on the interface between the layers.

Also analogous to the shallow water equations previously, we have that wave

speed scales like $\sqrt{g'H}$ which provides the *internal* Rossby deformation radius

$$L_D = \frac{\sqrt{g'H}}{\sqrt{2}f_0},\tag{1.2.5}$$

and also the Froude number

$$Fr = \frac{\sqrt{2}U}{\sqrt{g'H}}. (1.2.6)$$

These are related via $Fr/Ro = L/L_D$.

We will use the internal deformation radius as a horizontal length scale to non-dimensionalise, hence Ro/Fr = 1, implying $Fr \ll 1$ in the same way as Ro.

The incompressible, inviscid Euler equations in each layer are

$$\frac{\partial \boldsymbol{u}_{j}}{\partial t} + (\boldsymbol{u}_{j} \cdot \nabla) \boldsymbol{u}_{j} + \boldsymbol{f} \times \boldsymbol{u}_{j} = -\frac{1}{\rho_{j}} \nabla p_{j} - g\boldsymbol{k}, \qquad (1.2.7)$$

$$\nabla \cdot \boldsymbol{u}_i = 0. \tag{1.2.8}$$

We non-dimensionalise using L_D and U. We also redefine $\mathbf{u}_j = (u_j, v_j)^T$ as a *horizontal* velocity and refer to the vertical velocity as w_j . Using a prime for dimensionless quantities, the non-dimensionalised equations are

$$\left(1 + (-1)^{j} \frac{\Delta}{2}\right) \left(\frac{D \mathbf{u}'_{j}}{D t'} + (1 + Ro\beta' y') \hat{\mathbf{k}} \times \mathbf{u}'_{j}\right) = -\frac{1}{\rho_{j}} \nabla' p'_{j}, \tag{1.2.9}$$

$$\left(1 + (-1)^{j} \frac{\Delta}{2}\right) \left(\frac{\Delta F r^{2} \alpha^{2}}{2}\right) \frac{D w'_{j}}{D t'} = \frac{\Delta F r^{2}}{2 Ro} \frac{\partial \rho'}{\partial z'} - \left(1 + (-1)^{j} \frac{\Delta}{2}\right), \tag{1.2.10}$$

$$\nabla \cdot \mathbf{u}'_{j} = -\frac{\partial w'_{j}}{\partial z'}, \tag{1.2.11}$$

where $\alpha = H/L$ and we have used

$$\rho_j = \left(1 + (-1)^j \frac{\Delta}{2}\right). \tag{1.2.12}$$

We make the shallow water approximation, $\alpha \ll 1$ and the Boussinesq approximation, $\Delta \ll 1$ (and also $\Delta \ll Ro$). The vertical momentum equation can be used to show

that the upper layer displacement is $O(\Delta)$ while the interfacial displacement is O(1). Hence the Boussinesq approximation implies the free surface displacement is negligible when we consider expansions in Ro, and we may consider the surface as a frictionless rigid lid. This is equivalent to taking the *external* deformation radius to infinity, drawing a parallel to QG2D.

Expanding in powers of Ro with

$$u'_{j} = u_{j}^{(0)} + Rou_{j}^{(1)} + Ro^{2}u_{j}^{(2)} + \dots$$
 (1.2.13)

and similar for other quantities, we drop the Δ and α terms in the non-dimensionalised momentum and mass equations and equate orders of Ro. Leading order gives geostrophic balance (in each layer). The next order is then manipulated, similarly to QG2D, to arrive at (primes dropped for brevity)

$$\frac{\partial Q_j}{\partial t} + J(\psi_j, Q_j) + \beta \frac{\partial \psi_j}{\partial x} = 0, \qquad (1.2.14)$$

where

$$Q_j = \nabla^2 \psi_j + (-1)^j \left(\frac{\psi_1 - \psi_2}{2} \right). \tag{1.2.15}$$

Each layer, then, is similar to the equation for ζ in QG2D (before including dissipation and stochastic terms), except that Q_j is vorticity with a modification that couples the layer's dynamics.

Following Lee & Held (1991) [25] we include extra terms on the right hand side

$$\frac{\partial Q_j}{\partial t} + J(\psi_j, Q_j) + \beta \frac{\partial \psi_j}{\partial x} = -\delta_{j2} \kappa_M \nabla^2 \psi_2 + v_{2n} (-\nabla^2)^{n+1} \psi_j \qquad (1.2.16)$$

$$- (-1)^j \kappa_T \left(\frac{\psi_1 - \psi_2}{2} - \tau_e \right).$$

As in QG2D, n and v_{2n} are the degree and diffusivity of hyper-diffusion and κ_M is an Ekman drag coefficient acting on the lower layer. A new feature appearing is a radiative damping towards the equilibrium temperature field τ_e ; the parameter κ_T controls the strength. The exact profile $\tau_e(y)$ will be specified in section 1.2.2, but it

should relax the baroclinic velocity to an order one field.

Finally, it's useful to introduce the barotropic and baroclinic form of a layer variable, denoted with '+' and '-' respectively, as

$$f^{\pm} = \frac{f_1 \pm f_2}{2}.\tag{1.2.17}$$

These can easily be converted back for each layer as

$$f_i = f^+ + (-1)^{j-1} f^-.$$
 (1.2.18)

Consequently we have the simple potential vorticity-streamfunction relations

$$Q^{+} = \nabla^{2} \psi^{+}, \quad Q^{-} = (\nabla^{2} - 1) \psi^{-}.$$
 (1.2.19)

1.2.2 Boundary Conditions and Parameter Specification

Similar to Lee (1997) [26], we consider the two-layer model to have fixed channel walls at $y = \pm L_y \pi$ and let the x domain be periodic over $[0, 2\pi L_x]$. Boundary conditions at the walls in y read (e.g. Phillips (1954) [27])

$$\overline{\psi}_{1_y}=0$$
 (Phillips: upper layer),
$$(\partial_t + \kappa_M) \, \overline{\psi}_{2_y}=0 \text{ (Phillips: lower layer)},$$

$$\psi_{i_x}=0 \text{ (impermeable walls, } i=1,2), \tag{1.2.20}$$

where the over-bar denotes a zonal (x) average. For correct evaluation of the hyperdiffusion at the boundaries we specify the higher order conditions at the walls

$$\frac{\partial}{\partial y} \left(\nabla^2 \psi_i' \right) = \frac{\partial^3}{\partial y^3} \left(\nabla^2 \psi_i' \right) = \frac{\partial}{\partial y} \left(\frac{\partial^2 \overline{\psi_i}}{\partial y^2} \right) = \frac{\partial^3}{\partial y^3} \left(\frac{\partial^2 \overline{\psi_i}}{\partial y^2} \right) = 0, \tag{1.2.21}$$

where prime denotes the perturbation field from the zonal mean field [28].

We also note that since $\psi'_{i_x} = 0$, we simply have $\psi_i = \overline{\psi_i}$ on the boundary. i.e at

the channel walls

$$\psi_i' = 0. (1.2.22)$$

Care must be taken to satisfy conditions (1.2.20) and (1.2.21) numerically. This is discussed in section 1.3.

The baroclinic velocity associated with τ_e is

$$U_e^- = -\frac{\partial \tau_e}{\partial y}. ag{1.2.23}$$

The physical motivation behind the zonal wind profile is to model an equatorially centred radiative heating, e.g. from the sun. In the context of an idealised model there is some flexibility over the choice of U_e , and in this thesis it is taken to be (as in Lee (1997) [26]) a 'flat top' shape with a Gaussian decay

$$U_{e_1} - U_{e_2} = -2\frac{\partial \tau_e}{\partial y} = \begin{cases} 1, & \text{if } |y| < W_c, \\ \exp\left[-(|y| - W_c)^2 / \sigma^2\right], & \text{if } |y| > W_c. \end{cases}$$
(1.2.24)

The profile U_e is plotted in Figure 1.5 (dotted). For reference, the corresponding τ_e is

$$\tau_e = \begin{cases} -y/2, & \text{if } |y| < W_c, \\ -\frac{1}{2} \text{sgn}(y) \left(\frac{\sigma \sqrt{\pi}}{2} \text{erf} \left(\frac{|y| - W_c}{\sigma} \right) + W_c \right), & \text{if } |y| > W_c, \end{cases}$$
(1.2.25)

and has a typical value of order unity in the baroclinically unstable region keeping consistency with the non-dimensionalisation.

1.2.3 Energy

An energy equation is easily found from (1.2.16) by taking the product of the upper (lower) equation with the upper (lower) streamfunction and integrating over the

domain (e.g. [29]). This finds

$$\frac{\mathrm{d}E}{\mathrm{dt}} = \frac{\mathrm{d}}{\mathrm{dt}} \int \frac{1}{2} \left(|\nabla \psi_1|^2 + |\nabla \psi_2|^2 \right) + \left(\frac{\psi_1 - \psi_2}{2} \right)^2 \mathrm{d}\mathbf{x}, \tag{1.2.26}$$

$$= \frac{-1}{2} \int \frac{dQ_1}{dt} \psi_1 + \frac{dQ_2}{dt} \psi_2 dx, \qquad (1.2.27)$$

$$= -\kappa_{M} \int |\nabla \psi_{2}|^{2} d\mathbf{x} + \frac{\kappa_{T}}{2} \int \frac{\psi_{1} - \psi_{2}}{2} \left(2\tau_{e} - \frac{\psi_{1} - \psi_{2}}{2} \right) d\mathbf{x}, \qquad (1.2.28)$$
$$+ \nu_{2n} \int \psi_{1} (-\nabla^{2})^{n+1} \psi_{1} + \psi_{2} (-\nabla^{2})^{n+1} \psi_{2} d\mathbf{x}.$$

The first integral is composed of two terms, the first is the usual kinetic energy and the second is the potential energy

$$PE = \int \left(\frac{\psi_1 - \psi_2}{2}\right)^2 d\mathbf{x}. \tag{1.2.29}$$

A fundamental property of the two-layer model is the exchange between kinetic and potential energy. The imposed temperature profile supplies potential energy to an out of equilibrium flow, the potential energy grows disturbances through baroclinic instability, thereby transitioning to kinetic energy.

In the final expression, the baroclinic relaxation towards the imposed temperature gradient τ_e has an impact on the energy that depends on the state of the flow. The Ekman friction term is negative definite and so only serves to remove energy. The energy and enstrophy cascades are still relevant to the barotropic mode, hence the viscous term is included to remove enstrophy pile-up at the smallest scales as in QG2D. Integration by parts on the final term finds that hyper-diffusion always acts to remove energy. At a high enough numerical resolution (to allow a small v_{2n} parameter) this should be a negligible loss of energy.

It should be emphasised that since the temperature profile replenishes energy lost to friction and viscous dissipation, an essential difference between the two-layer model and QG2D is that the two-layer model doesn't require an artificial stochastic forcing to inject energetic eddies. The two-layer flow, albeit turbulent, is *deterministic*.

1.2.4 Parameter Space and Results

A horizontal strip with periodicity 20π and a meridional length 90 is used in [26]. In this work a square domain with dimensions $20\pi \times 20\pi$ is taken. This has the same aspect ratio as the QG2D model, but the main benefit is that the single and double jet regimes that are relevant for Chapter 2 don't require the unstable region to be as large as some of the experiments in [26]. A meridional length of 90 is therefore an unnecessary use of numerical power. We do however follow [26] fixing $\beta = 0.25$, $\sigma = \pi$, $\kappa_T = \kappa_M/3$ and study the (W_c, κ_M) parameter space. The flow is initialised so that $U^- = U_e^-$ (with $U_2 = 0$) with the addition of a small random eddy perturbation field allowing a unique realisation of turbulence to unfold.

The results provided in Figures 1.4 - 1.7 indicate that the jet number increases with an increased W_c or a decreasing κ_M . This is a result that has been observed over many more (unreported) simulations. It is also found that transitions between a single and double jet state can occur in some parameter regimes (e.g. Figure 1.7). Even the simulation in Figure 1.5, which looks like a very steady single jet state, can be seen to have a distinct second jet form briefly. Exploring larger W_c , it is possible to generate three or more jets, but the focus here on the two-layer model will concern the single and double jet state only.

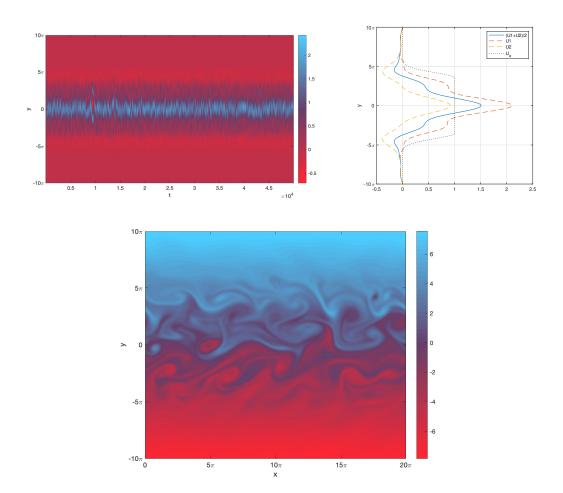


Figure 1.4: Two-layer model results with $(W_c, \kappa_M) = (3.25\pi, 0.07)$. Top: contoured is the Hovmöller plot for the depth averaged zonal mean flow and plotted are the mean velocity profiles as labelled. Bottom: final snapshot of the depth averaged (barotropic) potential vorticity $Q^+ + \beta y$.

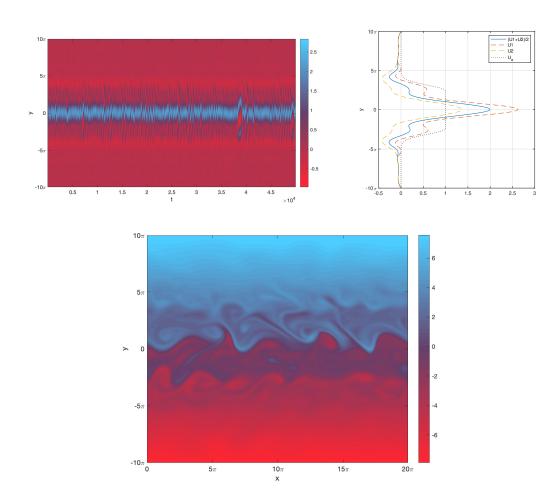


Figure 1.5: Same as in Figure 1.4 but with a narrower equilibrium profile and a reduced friction, $(W_c, \kappa_M) = (2.4\pi, 0.03)$.

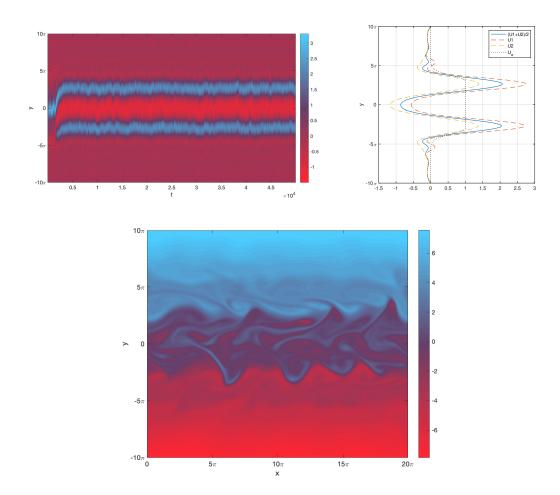


Figure 1.6: Same as in Figure 1.4 but with parameters setting for two jets. $(W_c, \kappa_M) = (3\pi, 0.02)$.

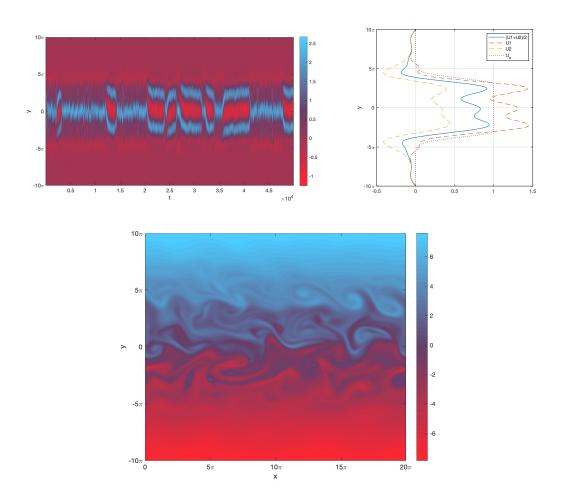


Figure 1.7: Same as in Figure 1.4 but with parameters setting for two jets. $(W_c, \kappa_M) = (3.25\pi, 0.05)$.

1.3 Numerical methods

1.3.1 Numerical Methods For Barotropic Flow

To integrate the vorticity equation (1.1.61) a second order Runge-Kutta scheme is used. On a doubly periodic domain the boundary conditions are periodic, making it sensible to work spectrally. Taking the Fourier transform (1.1.29) of (1.1.61)

$$\frac{\partial}{\partial t} \zeta_{\mathbf{k}} = -J(\psi, \zeta)_{\mathbf{k}} - i\beta k_x \psi_{\mathbf{k}} - \alpha \zeta - v_{2n} |\mathbf{k}|^{2n+2} \zeta_{\mathbf{k}} + \sqrt{2\alpha} \eta_{\mathbf{k}}. \tag{1.3.1}$$

The potential vorticity-streamfunction relation is simply

$$\zeta_{\mathbf{k}} = -|\mathbf{k}|^2 \psi_{\mathbf{k}},\tag{1.3.2}$$

which crucially means the streamfunction is easily inverted as $\psi_{\mathbf{k}} = -\zeta_{\mathbf{k}}/|\mathbf{k}|^2$. This is very fast numerically since it only requires element-wise matrix division. Velocity components are related by

$$u_{\mathbf{k}} = -\mathrm{i}k_{y}\psi_{\mathbf{k}} = \frac{\mathrm{i}k_{y}}{|\mathbf{k}|^{2}}\zeta_{\mathbf{k}},\tag{1.3.3}$$

$$v_{\mathbf{k}} = \mathrm{i}k_x \psi_{\mathbf{k}} = -\frac{\mathrm{i}k_x}{|\mathbf{k}|^2} \zeta_{\mathbf{k}},\tag{1.3.4}$$

which are just as easy to calculate.

The bulk of numerical work is contained in the nonlinearity in the Jacobian term. We have

$$J(\psi,\zeta) = \psi_x \zeta_y - \psi_x \zeta_x = u\zeta_x + v\zeta_y = (u\zeta)_x + (v\zeta)_y - (u_x + v_y)\zeta_{xy}$$
(1.3.5)
= $(u\zeta)_x + (v\zeta)_y$

and the k^{th} component is found as

$$\{J(\boldsymbol{\psi},\boldsymbol{\zeta})\}_{\boldsymbol{k}} = \left\{ (u\boldsymbol{\zeta})_x + (v\boldsymbol{\zeta})_y \right\}_{\boldsymbol{k}},\tag{1.3.6}$$

$$= ik_x (u\zeta)_{\mathbf{k}} + ik_y (v\zeta)_{\mathbf{k}}, \qquad (1.3.7)$$

but to compute this one must compute the physical u, v and ζ from ζ_k and then Fourier transform $u\zeta$ and $v\zeta$ back into Fourier space. This involves three inverse and two forward 2D Fourier transforms. A regular (1D) discrete Fourier transform with MATLAB's Fast Fourier Transform (FFT) function is $O(N \ln N)$ for N the number of discrete points or wavenumbers. The 'divide and conquer' algorithm used for FFT is optimised if N has small prime-factors. For this reason we always use a power of 2 for the resolution.

Supposing the (single dimensional) resolution of N (even), then in general N wavenumbers are needed. Generally these can be thought of in the wavenumber order

$$[0,1...N/2,-N/2+1,...,-2,-1].$$

Conjugate symmetry can be utilised to save memory as it means only knowledge of 'non-negative' Fourier components is required. The choice of which waves are 'non-negative' is actually somewhat arbitrary, but for two dimensions on an N^2 grid the choice is made to divide the wave vectors into 'positive' waves K^+ (which are solved) and 'negative' waves K^- (which can be derived), according to

$$K^{+} = \left\{ (k_x, k_y) : \left(\left(0 < k_x \le \frac{N}{2} \right) \cap \left(-\frac{N}{2} \le k_y \le \frac{N}{2} \right) \right) \right.$$

$$\left. \cup \left(\left(k_x = 0 \right) \cap \left(0 < k_y \le \frac{N}{2} \right) \right) \right. \right\}, \tag{1.3.8}$$

and

$$K^{-} = \{ \mathbf{k} : \mathbf{k} + \mathbf{j} = \mathbf{0} \text{ for } \mathbf{j} \in K^{+} \}.$$
 (1.3.9)

This divide can be thought of as following the sign of k_x and when $k_x = 0$ follows the sign of k_y .

To integrate numerically, a time-step δt must be selected. A small time-step provides numerical stability but takes more computation and more time. The aim is to choose the largest time-step such that the simulation is numerically stable. The

Courant-Friedrichs-Lewy condition for numerical stability is

$$\frac{u\delta t}{\delta x} \ll C. \tag{1.3.10}$$

Here u is the magnitude of fluid velocity, $\delta x = 2\pi/N$ is the grid scale and C depends on the integration scheme. For Runge-Kutta schemes, C is typically unity.

The condition gives an inverse relationship between resolution and time-step. Hence if the resolution is doubled, the integration time not only increases from the calculation of terms in the vorticity equation (predominantly the $O((N \ln N)^2)$ 2D Fourier transforms), but will also double from the time-step reduction. Typically it's found that doubling the resolution increases the integration time by a factor of ~ 10 .

A difficult aspect of numerical integration is choosing the hyper-diffusion coefficient. On the one hand, this must be large enough to remove enstrophy before pile-up at the small scales which is unphysical and potentially leads to numerical instability. On the other hand the coefficient must be small enough to have minimal interference on the energy balance. Investigating the enstrophy in spectral space by looking at plots such as $\ln |(\zeta^2)_{\bf k}|$ (logarithm of the 2D transform of enstrophy), we can tune the viscous coefficient so that enstrophy decays (at least) exponentially before pile-up on the shortest waves.

In all cases a fourth order hyper diffusion (n = 2) is chosen, and the numerical parameters for the various simulation resolutions used in this thesis are given in Table 1.1.

Grid points, <i>N</i>	δt	v_4
128	0.02	1×10^{-7}
256	0.01	2.5×10^{-8}
512	0.005	4×10^{-9}

Table 1.1: Numerical parameters for the barotropic model.

1.3.2 Numerical Methods For Two-Layer Flow

Since there are boundary conditions at the channel walls in the two-layer set-up, we use finite-difference methods in y. Spectral methods are still used in x as the model

is zonally periodic. The grid points in y are placed in the centre of the N meridional divisions. i.e.

$$y_n = 2\pi L_y \left(\frac{n-0.5}{N} - \frac{1}{2}\right), n = 1, \dots, N.$$
 (1.3.11)

The spacing between grid points is simply

$$\delta y = \frac{2\pi L_y}{N}.\tag{1.3.12}$$

A key difference for implementation is we evaluate y derivatives using differentiation matrices. Derivatives on the interior of the domain can be found using adjacent points fairly intuitively (or more formally using Taylor expansions at the grid scales about the the grid point of interest). Recalling that the variables are discretised at the 'half-way points', differentiation matrices satisfying Neumann (N) and Dirichlet (D) conditions can be found as

Similarly, second derivatives adopting Neumann and Dirichlet conditions at the walls are

$$\mathbf{N}_{2} = \frac{1}{\delta y^{2}} \begin{bmatrix} -1 & 1 & 0 & 0 & \ddots & 0 \\ 1 & -2 & 1 & 0 & \ddots & \ddots \\ 0 & 1 & -2 & 1 & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \\ 0 & \ddots & \ddots & 1 & -2 & 1 \\ 0 & \ddots & \ddots & 0 & 1 & -1 \end{bmatrix}, \quad \mathbf{D}_{2} = \frac{1}{\delta y^{2}} \begin{bmatrix} -3 & 1 & 0 & 0 & \ddots & 0 \\ 1 & -2 & 1 & 0 & \ddots & \ddots \\ 0 & 1 & -2 & 1 & 0 & \ddots & \ddots \\ 0 & 1 & -2 & 1 & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ 0 & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ 0 & 0 & 1 & -2 & 1 & 0 & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ 0 & 0 & 0 & \ddots & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots \\ 0 & 0 & 0 & 0 & 0 & 1 & -3 \end{bmatrix}.$$

For the high order derivatives for hyper-diffusion, it is first observed that

$$\nabla^2 \psi_j = q_j - (-1)^j (\psi_1 - \psi_2) / 2, \tag{1.3.13}$$

so the first laplacian is easily found using the right hand side. One can then use

$$\nabla^2 = -k^2 \mathbf{I} + \mathbf{N}_2, \tag{1.3.14}$$

to premultiplying each $\left[\nabla^2 \psi_j\right]_k$ twice through each zonal wavenumber, k. This satisfies boundary conditions (1.2.21).

Chapter 2

Jet Transitions and Rare Events

2.1 Introduction

We have seen from the integrations carried out in Chapter 1 that a zonal flow profile may sporadically transition back and forth between a different number of jets. This indicates the mean flow has multiple quasi-equilibrium states that it can abruptly transition between. This is a widely observed phenomenon in turbulent flow. For instance, the Kuroshio current [30], magnetic field reversals in turbulent dynamos [31, 32], various two-dimensional scenarios [33, 19, 34], three-dimensional turbulence experiments [35], convection turbulence [36, 37, 38]) and the bimodal behaviour in the turbulent wake from flow past a cylinder [39] are all examples known to exhibit transitional behaviour between turbulent attractors. There are even known bistability phenomena between laminar and turbulent flow states [40].

Developing tools to efficiently calculate the likelihood of the transitions in the models introduced in Chapter 1 may serve as a stepping stone to calculate the probabilities of rare weather events in more realistic, numerically expensive models. The layer models may also still be able to capture the essential physics of jet transitions, going some way to understand why, for example, between 1939 and 1940 Jupiter remarkably lost a jet to be replaced with three white anticyclones [41].

The most straightforward approach for finding a rare event and estimating its statistics is direct numerical simulation (DNS) (i.e. integrate a simulation, or run many trials, until the rare event is observed multiple times). This is also referred to

as naïve Monte-Carlo or crude Monte-Carlo. The approach suffers as events become increasingly rare. To see why, suppose a random variable achieves its rare event when $X \in B$, with probability $0 . The Naïve Monte-Carlo method draws an i.i.d. sample <math>(X_1, X_2, \dots, X_N)$ from X and gets a probability estimator

$$\hat{p} = \frac{1}{N} \sum_{i=1}^{N} \mathbf{1}_{B}(X_{i}), \tag{2.1.1}$$

where $\mathbf{1}_B(X)$ is the indicator function which takes the value of 1 if $X \in B$ and 0 otherwise. The i.i.d. variables $\mathbf{1}_B(X_i)$ are Bernoulli with success probability p, so the normalised variance is

$$\mathbb{V}\left(\frac{\hat{p}}{p}\right) = \frac{1-p}{Np} \approx \frac{1}{Np},\tag{2.1.2}$$

for small p. The normalised variance should be order one or less, which demands $N \gtrsim 1/p$ for an acceptable variance. For a rare event, with $p \ll 1$, this demands a large sample which can quickly become impractical. To achieve the same variance with a reduced computational effort in sampling we can use a *rare event algorithm*. As an aside, events may even be so rare that the limitations of a practical sample size render a single rare event observation as unlikely using DNS (i.e. if the practical sample size is much less than 1/p). In this case a rare event algorithm may be needed just to show the existence of an event.

There are two main algorithmic approaches to rare event sampling: importance sampling and splitting methods. Importance sampling draws samples from a biased distribution to make a rare event more probable. Consequently importance sampling requires enough knowledge of the system to choose a good importance distribution and is not well suited to high-dimensional turbulence.

Splitting algorithms were first introduced in the 50s by Kahn & Harris [42] and Rosenbluth & Rosenbluth [43] and there are now various modifications. The applications have been wide-spread, from molecular dynamics [44] to air traffic management [45], but the underlying idea is consistent. The algorithms hinge on simultaneously integrating an *ensemble* of trajectories and guiding the members

towards rare event realisations through a systematic pruning and cloning. Simulations which are deemed the most unlikely to achieve the rare event are removed and replaced with a clone of another, more promising, ensemble member. After cloning, stochasticity in the system *splits* the cloned members (avoiding the redundancy of an ensemble with identical members). If correctly implemented, the splitting algorithm will drive all the ensemble members to achieve the rare event. The final ensemble can be thought of as a selected subset of the effective ensemble that also comprises of all the deleted simulations and their potential clones; probability estimates are easily obtained from such considerations and are generally known to be unbiased for most splitting algorithms.

The likeliness of achieving a rare event is monitored by values of a scoring function. In the original versions of splitting algorithms, the entire ensemble is integrated and splitting occurs at regular intervals. Members which don't meet a specific value of the scoring function (or score, or splitting level) are discarded and replaced by a clone of another member. The scores are decided beforehand. Such algorithms are referred to as 'multi-level splitting algorithms'. Relatively recently an *adaptive* multi-level splitting algorithm (AMS) has been put forward, which, as the name suggests, adapts the splitting levels on-the-fly [46]. The advantage is no decision needs to be made for the scoring level values (one of the most difficult aspects of multi-level splitting algorithms). It can be proved [47] that if optimally implemented it's known that the variance of estimators is reduced for multi-level splitting methods over DNS. Optimal implementation, however, depends on an optimal choice of scoring function which is usually out of reach. The challenge of implementing AMS is therefore to choose a *good enough* scoring function; this is specifically discussed in sections 2.2.3 and 2.2.4.

Much of the research into rare event algorithms has been in low dimensions. For example, applications to Ornstein-Ulhenbeck processes and low dimension potential-well problems [46, 48, 49]. These simple SDEs offer analytic solutions for rare event probabilities, allowing the algorithm estimators to be checked. They are also numerically fast so various algorithms may be compared and verified to

be more efficient than DNS (or not). An overview of many key results and notable contributions can be found in the historical perspective [50].

The extension to a variable such as potential vorticity on an $N \times N$ resolution grid is relatively straightforward, however the computational memory and power needed for an ensemble of simulations can be immense. Especially when one considers that accurate statistics require multiple algorithm implementations. That said, the technological advances of recent decades have allowed an application of rare event algorithms to high-dimensional rare events. Statistics of extreme heatwaves over western Europe in a climate model which are out of reach of DNS have been accumulated using a splitting algorithm [51, 52] and AMS has been applied directly to the QG2D model to study the transitions between two and three jets [18, 53]. Notably, [18] is the first use of a rare event algorithm in a model of turbulence. Similar work is repeated and discussed at depth in section 2.2 as an appropriate introduction to AMS and a test of algorithm implementation. We then adapt AMS to find transitions in the two-layer model in section 2.3. The achievement here is the adaptation from a stochastic model to a deterministic model (since AMS requires stochasticity to split trajectories).

2.2 Adaptive Multi-Level Splitting Algorithm (AMS)

2.2.1 The Algorithm

Let \mathbb{R}^d be the space of trajectories of some stochastic dynamics with two attractive states which are contained within the open subsets $A, B \in \mathbb{R}^d$. Also suppose $A \cap B = \emptyset$. It should be likely that if a trajectory leaves either set it returns to the same set. We are interested in the rare event that a trajectory moves from set A to set B, without going back to set A.

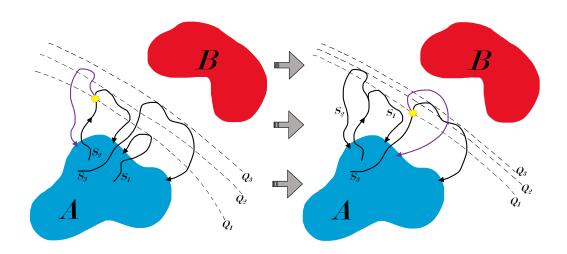


Figure 2.1: Schematic representation of the AMS algorithm with an ensemble size of three. The first iteration, on the left, shows S_1 being split from one of the other two trajectories (in this case S_2) at score Q_1 . The purple trajectory is the new trajectory and the yellow star marks the split. The next iteration of the algorithm, on the right, shows the adjusted $\{S_1, S_2, S_3\}$ ensemble, re-ordered by ascending maximum scores. A new split is shown following the same procedure but this time branching off S_3 .

The first step in setting up AMS is to define a scoring function, $Q: \mathbb{R}^d \to \mathbb{R}$, with

$$Q(X) = \begin{cases} 0, & \text{if } X \in A, \\ 1, & \text{if } X \in B. \end{cases}$$
 (2.2.1)

The efficacy of the algorithm will depend on the relationship between the values of

Q(X) for X outside of A and B and the *actual* probability of reaching B before A (see section 2.2.4 for more precise details).

We let N be the size of the ensemble, label the simulations in the ensemble S_1, S_2, \ldots, S_N , denote by r the number of transitioned simulations and denote by n the number of splits.

To begin the algorithm we initialise trajectories with initial conditions which are distributed according to the probability density function for a flow given it's in set A (practical implementation of this is discussed in section 2.2.2). These simulations are then run from these initial conditions until they reach some small escape score Q_0 near the boundary of A and then they are continued until they (probably) return to A. The maximum score of each trajectory is tracked and referred to as its score. When all trajectories return to set A, we are in a position to begin the algorithm's iterative process. The simulations are re-indexed by ascending score, i.e. S_1, S_2, \dots, S_N with respective scores $Q_1 \leq Q_2 \leq \cdots \leq Q_N$. The simulation S_1 is discarded and one of the remaining simulations is randomly nominated to be cloned. The cloned version, S_1^* , is only cloned until it reaches the score Q_1 . The clone S_1^* is then integrated from this point – the new forcing realisation splitting it from its sibling member – until reaching set A or B. The simulations are then re-indexed and the process is repeated again on whichever simulation now has the smallest score. When a particle reaches B it should still be considered part of the ensemble and available for cloning selection. The algorithm can be chosen to stop when r reaches some value, typically r = N. The algorithm procedure is described in Algorithm 1 and represented schematically in Figure 2.1.

The algorithm described above is known as 'last particle AMS', although we will simply refer to it as AMS. In full generality, the algorithm may discard a variable number of trajectories, k_j , at the j^{th} iteration. The algorithm then picks a clones randomly from the remaining trajectories for each discarded trajectory and the clones are all split from the score of the k_j^{th} last particle. It's observed that k = 1 yields the best numerical results [48], so we opt for the last particle method and don't clutter the algorithm with unnecessary parameters.

Algorithm 1: AMS

```
1 r = 0
n = 0
3 for i = 1 : N do
       initialise S_i in set A
       run S_i until Q_i \ge Q_0 > 0 and Q_0 small.
5
       continue until either returning to A or reaching B
7 arrange simulations in ascending order of high score and re-index
8 update r if needed (unlikely)
   while r < N do
       discard S_1
10
       clone S_j as S_1 up until Q_1 for random j \in 2,...,N
11
       n = n + 1
12
       run S_1 until returning to A or reaching B
13
       arrange simulations in ascending order of high score and re-index
14
       update r if needed
15
16 \mathbb{P}(\{Q=Q_0\} \rightarrow B) = \left(\frac{N-1}{N}\right)^n
```

2.2.2 Initialisation

Ideally, we should draw initial conditions from set A according to the invariant measure of the dynamical system conditional on being in set A and then run the flow until first exit. This is practically achieved by drawing initial conditions from a large data file of a long run. The idea is that the ergodicity of the system implies minimal bias here. To reduce any potential bias from sampling from the same specific long run, the initial conditions are also simulated with their own stochastic forcing realisations for a time before being allowed to reach the boundary of A. This is at least a step beyond using a single initial condition for the entire ensemble as in $[53]^1$.

The simulations are then continued until they return to set *A* and their highest scores are recorded.

2.2.3 Scoring Function

A crucial component for the efficacy of AMS is choosing a sensible scoring function to determine how likely a given trajectory is to reach the rare event set. Ideally,

¹ It strikes the author that this is a bad practice. Moreover, by default this is presumably also the method used in the original study [18], as it is not mentioned and [53] is essentially an extension and detailed explanation of the same work.

the theoretical scoring function is simply the *committer function*: the probability of reaching set *B* before set *A* from any given point. It actually turns out that any committer-like function which has isosurfaces which align with the isosurfaces of the event probability (see section 2.2.4) is an optimal choice. The dilemma for rare event algorithms is that their optimisation requires knowledge of the isosurfaces of the committer function, but if we already knew the committer function then the rare event problem is solved.

For the QG2D implementation, we take the same approach as [18], where the scoring function depends on the absolute value of the second, third and fourth Fourier modes of zonal vorticity, i.e. $|\overline{\zeta}_2|$, $|\overline{\zeta}_3|$ and $|\overline{\zeta}_4|$. For given parameters (α,β) which can exhibit two and three jet states, we can investigate the (empirical) probability density function of the lowest wavenumbers and find a clear separation in Fourier space. This allows a sensible choice of lower and upper bounds for each $|\overline{\zeta}_k|$ to define each set A and B, i.e. A and B are distinct cuboids in the $(|\overline{\zeta}_2|,|\overline{\zeta}_3|,|\overline{\zeta}_4|)$ space. If P is a point in this space corresponding to time t of trajectory X_t , let |P-A| denote the shortest distance of P from the boundary of P (in the P-P) and likewise for P-P. The scoring function used is

$$Q(X_t) = \begin{cases} \frac{|P-A|}{2|P-B|}, & \text{if } |P-A| < |P-B|, \\ 1 - \frac{|P-B|}{2|P-A|}, & \text{if } |P-A| > |P-B|. \end{cases}$$
 (2.2.2)

This is a crude guess at the isosurfaces of the committer function. In fact, it's pointed out [18, 53] that the typical path between the sets A and B is asymmetric in the fourth wave. The two jet to three jet transition first sees a large growth in fourth-wave amplitude for instance. This surely implies the level surfaces of a committer function cannot be aligned so symmetrically as (2.2.2). It will be shown below the crude scoring function (2.2.2) works nonetheless. This is an encouraging sign for the rare event algorithms' ease of application.

With the parameters $(\alpha, \beta, k_f, d_k) = (0.0012, 5.26, 14.5, 0.6)$, where k_f and d_k and the radius and half width of the forcing annulus, we define jet state boundaries

from DNS evidence. The two jet state, set A, is defined by

$$|\overline{\zeta}_2| \in [0.22, 0.231],$$
 (2.2.3)

$$|\overline{\zeta}_3| \in [0, 0.022],$$
 (2.2.4)

$$|\overline{\zeta}_4| \in [0, 0.06],$$
 (2.2.5)

and the three jet state, set B, by

$$|\overline{\zeta}_2| \in [0.1, 0.14],$$
 (2.2.6)

$$|\overline{\zeta}_3| \in [0.2, 0.3],$$
 (2.2.7)

$$|\overline{\zeta}_4| \in [0.12, 0.175].$$
 (2.2.8)

2.2.4 Probability Estimators

Suppose we have a trajectory X_t , where $0 \le t \le s$, for s the hitting time for set A or B. The event probability we are interested in is

$$p = \mathbb{P}(X_s \in B | Q(X_0) = Q_0), \tag{2.2.9}$$

and it is intuitive that its probability estimator may be found at any point during AMS as

$$\hat{p} = \frac{r}{N} \left(\frac{N-1}{N} \right)^n. \tag{2.2.10}$$

This is because r/N is the fraction of transitions in the ensemble and (N-1)/N is the 'inverse growth' of the effective ensemble at each of the n splits. This estimator is known to be unbiased [54]. When AMS is implemented to the completed situation where the entire ensemble consists of transitions, the r/N factor is unity and the expression simplifies to $\hat{p} = (1 - 1/N)^n$.

Useful insights can be derived from the AMS probability estimator (2.2.10). For example, for a given ensemble size the number of algorithm iterations, n, is expected to scale like

$$n \propto -\ln(p). \tag{2.2.11}$$

So as p becomes small the number of algorithm iterations to find some number of rare events has a better scaling than the 1/p scaling for DNS to find a fixed number of rare events (an 'iteration' is a different process in AMS and DNS of course, but the scaling is promising for AMS as events become rare).

Another property which can be seen from (2.2.10) is the nonlinear relationship between n and r. For instance, it is a simple exercise to show from (2.2.10) that the required number of iterations for AMS to find N transitions is approximately only double the number of iterations to find the first transition if $p \sim 1/N^2$. For smaller p, this idea of increased returns becomes more extreme, hinting that it is beneficial to run the algorithm to completion (r = N). However, it is unknown how the variance of the estimator \hat{p} depends on the algorithm-stopping r value. The picture is also distorted since not all algorithm iterations will have the same expected computational time – simulations split at different scores will typically have different expected hitting times. It will be seen in the results section that we run AMS implementations with r = 1 and r = N for a direct inspection of the variance of \hat{p} in the most extreme cases.

An optimal implementation of AMS could be considered as minimising the variance of the probability estimator \hat{p} against the computational time [50]. Before we can discuss results related to the variance of the probability estimator \hat{p} , though, it is necessary to introduce some notation. Denoting s_l the hitting time of the open threshold with score l as

$$s_l = \inf_t \{Q(X_t) > l\}.$$
 (2.2.12)

The committer function

$$Q^*(X) = \mathbb{E}[\mathbf{1}_B(X_s)|X_0 = X], \tag{2.2.13}$$

is simply the probability of reaching the rare event set *B* from some position *X*. We also define the conditional distribution

$$\eta_l(\varphi) = \mathbb{E}[\varphi(X_{s_l})|s_l < s], \tag{2.2.14}$$

where φ is any test function, and the level probabilities

$$p_l = \mathbb{P}(s_l \le s),\tag{2.2.15}$$

for $l \in [0, 1]$. i.e. $p_0 = 1$ and $0 < p_1 \ll 1$.

Now, there is a central limit theorem

$$\sqrt{N}\hat{p} = \mathcal{N}\left(p, \sigma^2\right),\tag{2.2.16}$$

associated with AMS [e.g. 47, and references therein], where

$$\sigma^{2} = -p^{2} \ln p - \int_{Q_{0}}^{1} \mathbb{V}_{\eta_{l}}(Q^{*}) p_{l} dp_{l}.$$
 (2.2.17)

and

$$\mathbb{V}_{n_l}(Q^*) = \mathbb{V}\left[\mathbb{E}(\mathbf{1}_B(X_s)|X_{s_l})\right]. \tag{2.2.18}$$

The integration in (2.2.17) is over the level sets of the scoring function Q and the integrand is the variance (across the given level set) of the *real* probability of reaching set B from such a position. If we know and use the committer function Q^* for the scoring function, such variance vanishes since $\mathbb{E}(\mathbf{1}_B(X_s)|X_{s_l})$ is constant for any given scoring level l. Any *committer-like* function [50], which has a gradient everywhere in the same direction (though not necessarily the same magnitude) as Q^* , or equivalently shares the same isosurfaces as Q^* , will also be an optimal choice.

Insightful as (2.2.17) may be from a theoretical vantage, we cannot calculate the integral without the knowledge of a comitter-like function, rendering it practically useless.

It is worth mentioning that in the hypothetical case that the committer function is used, the number of iterations is Poisson distributed [55] according to

$$n \sim \text{Pois}(-N \ln p),$$
 (2.2.19)

giving direct access to the distribution of \hat{p} . The unwanted effects of a sub-optimal

scoring function, though, are not insignificant [also 55].

Finally, it is put forward that the upper bound on AMS variance is "at worst twice as bad as DNS" [47] for the same computational effort. Although disappointing, at least AMS variance is theoretically bounded (even with a worst permissible scoring function).

2.2.5 Return Times

The probability estimator \hat{p} is a fairly arbitrary statistic in of itself. It simply gives the probability of transitioning from the threshold with score Q_0 to set B without first entering set A. A more fundamental statistic of interest is: given we are in set A, how long can we expect to wait until we are in set B? This is particularly important for risk assessment, for example, where one needs to know the expected frequency of a catastrophe.

The expected waiting time may be estimated from \hat{p} if we know the expected time taken to *reset* a simulation (e.g. [48]). The reset time, τ , is the total time taken to integrate from the initialisation at $Q = Q_0$ until returning to set A and passing back out of A to the $Q = Q_0$ threshold again. The reset time is not defined if the simulation reaches set B before returning to set A after initialisation.

To access the waiting time, consider a single simulation initialised in set A and integrated forward in time. Typically we'll expect the score to pass back-and-forth many times through $Q = Q_0$ before a transition to set B (since p is small). If we imagine that when we reach $Q = Q_0$ (having just exited set A) that we count returning to set A as failure and transitioning to set B as a success. Now we essentially have a Bernoulli trial with success probability estimated by \hat{p} . Using the reset time, it's easy to conclude the expected waiting time to transition from A to B is

$$\mathbb{E}[T_{A\to B}] \approx \frac{\tau}{p}.\tag{2.2.20}$$

The transition time itself, τ_{trans} (i.e. the expected time to go *directly* from $Q = Q_0$ to B), and the small time, τ_1 , to initially reach $Q = Q_0$ from a random initial condition in A could be added for precision. But presuming events are rare and abrupt it will

be a small adjustment. The condition is $p \ll \tau/(\tau_{\text{trans}} + \tau_1)$.

Data from 10000 QG2D simulations are shown in the Figure 2.2. Each simulation was initialised using the same initialisation process as AMS and the scoring function is defined as above. All other parameters are kept the same, too. The mean reset time is $\tau = 55.1957...$ with a standard deviation of 111.80896.... Assuming a central limit theorem, the 95% confidence interval is [53.004..., 57.387...].

The 10000 simulations took roughly 30 hours to integrate directly, which in the context of rare event algorithms is not a very long time. It is therefore feasible to reduce the confidence interval if desired.

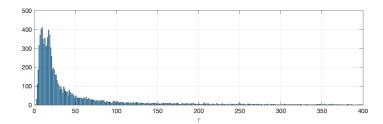


Figure 2.2: Histogram of return times for the 2 jet state from 10000 simulations. Reset times greater than 400 are not shown. This is for an improved display of the histogram at lower τ values.

2.2.6 Parallelising

It is possible, generally, to clone/split k trajectories at each splitting implementation. It is found the algorithm's variance is minimised with k = 1 [48], but also noted that there is a potential to run in parallel if k > 1 which might reverse the numerical costs associated with the increased variance. For instance, [53] parallelised with k = 100. Another parallelising option is to simply run many algorithm implementations in parallel with smaller ensembles. This is equivalent to dividing the pool of potential clones to subsets of a 'global' ensemble. It seems inevitable this will produce variance but a detailed study would be interesting. Here, we give a method to parallelise the 'last particle' version of AMS.

For simplicity consider parallelising AMS on two cores only. We can't immediately run the two worst scoring simulations in parallel for two reasons: the worst scorer may never exceed the second worst scorer (in which case the second

worst scorer may be called on as the clone), and also the second worst scorer may be assigned the worst scorer as a clone when it needs to be split which can't be done until the worst scorer 'overtakes'. The solution to this is simple, we run the worst scoring simulation and wait until it achieves the second worst simulation's high-score. Then (and only then) we can confidently randomly select a clone for the second worst scorer and run the two in parallel until completion. If the worst scorer never achieves the second worst's high-score then we don't parallelise and we run the worst scorer again.

To parallelise with multiple cores, consider labelling the simulations $S_1, S_2, ..., S_N$ in ascending order according to their respective high-scores $Q_1, Q_2, ..., Q_N$. We then clone and run for S_1 until it achieves Q_2 , then clone for S_2 and run S_1 and S_2 in parallel until they both achieve Q_3 , then clone for S_3 and run S_1 , S_2 and S_3 in parallel until they achieve Q_4 etc... At some level n not all $S_1,...,S_n$ will reach score Q_{n+1} . At this point, let all the simulations run until completion, re-order and begin again. Psuedo-code is given in Algorithm 2.

Although the parallelised AMS doesn't consistently use parallelisation, it is expected that with a large enough number of simulations, the worst scoring simulations are tightly bunched. It's a tough competition to be last, typically a quick dethroning is expected from the next-best neighbour when cloning randomly from the larger (better performing) ensemble.

Note however that with each round of splitting, before starting with the losing simulation again, there are potentially long waiting times for the final simulation(s) to finish. It is an annoying consequence of AMS (compared to, say, multi-level splitting with fixed splitting times) and a waste of parallel capabilities. An anticipated step further in sophistication for a parallelised last particle AMS might monitor which simulations have finished and if they are guaranteed losers, prematurely begin the next round of splitting. This could entirely bypass the bottleneck that comes with waiting for the longest simulations to finish. Lastly, we note that although this parallelisation procedure is slightly slower and more convoluted than simply running multiple simultaneous implementations of AMS each on a single core, it is far more

memory efficient (and memory constraints become important when saving thousands of turbulent flows at relatively high temporal resolution).

Algorithm 2: AMS - parallel

```
1 r = 0
 n = 0
 3 for i = 1 : N do
       initialise S_i in set A
       run S_i until Q_i \ge Q_0 > 0 and Q_0 small.
       continue until either returning to A or reaching B
 7 arrange simulations in ascending order of high score and re-index
 8 update r if needed (unlikely)
   while r < N do
       P_{max} = N - r
10
11
       for p = 1 : P_{max} do
           clone S_i onto S_p for random j \in \{p+1, p+2, ..., N\} up to Q_p
12
13
           for q = 1 : p do
14
              run S_q until reaching Q_{p+1} or returning to set A
15
           if none reach Q_{p+1} then
16
              break
17
           if only some reach Q_{p+1} (but not all) then
18
               clone the worst scoring sim that didn't reach Q_{p+1}. Run this
19
                along with those that reached Q_{p+1} until completion.
               n=n+1
20
               if any sims reach set B then
21
                   update r appropriately
22
               break
23
       re-order
25 \mathbb{P}(\{Q=Q_0\}	o B)=\left(rac{N-1}{N}
ight)^n
```

2.2.7 Results

DNS may be used to find transitions with parameter values $(\alpha, \beta) = (0.0012, 5.26)$. A long integration of the QG2D model is given in Figure 2.3. Three transitions from two to three jets are observed over the entire integration time of 3.67×10^6 time units. The time spent in the two jet state is approximately 8.075×10^5 time units per transition. This isn't out of keeping with the estimate in [18] (supplementary material) which finds five transitions from DNS with an average time of 5.3×10^5 per two-to-three jet transition. The difference is easily accounted for in the error sampling only a handful of transitions. There may also be differences in underlying statistics due to alternative numerical integration schemes.

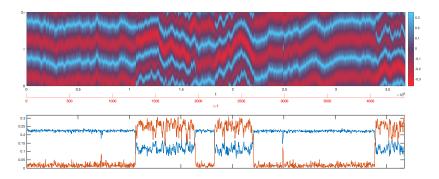


Figure 2.3: Transitions between two and three jets in QG2D found by DNS. Top: Hövmoller plot of zonal mean velocity, bottom: corresponding line plot of $|\overline{\zeta}_2|$ (blue) and $|\overline{\zeta}_3|$ (orange) showing abrupt transitions.

Results for the AMS algorithm applied to QG2D with the same parameter values are given in Table 2.1. In all cases, the rare event is a transition from two to three jets.

There are three groups of algorithm implementations, each consisting of 10 ensembles:

- 1. Ensembles 1-10 with N = 50 and r = 1.
- 2. Ensembles 11-20 with N = 1000 and r = 1.
- 3. Ensembles 21-30 with N = 1000 and r = 1000.

where N is the ensemble size and r is the desired number of transitions from the algorithm.

Ens.	N	r	n	$\hat{p} (\times 10^{-6})$
1	50	1	427	3.58562
2	50	1	592	0.12790
3	50	1	269	87.26586
4	50	1	357	14.74820
5	50	1	427	3.58562
6	50	1	742	0.00618
7	50	1	444	2.54335
8	50	1	707	0.01253
9	50	1	928	0.00014
10	50	1	96	2875.63170
11	1000	1	5513	4.02288
12	1000	1	2989	50.26252
13	1000	1	4726	8.84092
14	1000	1	5550	3.87668
15	1000	1	2662	69.71555
16	1000	1	4308	13.43146
17	1000	1	6042	2.36963
18	1000	1	9033	0.11887
19	1000	1	7961	0.34742
20	1000	1	7962	0.34707
21	1000	1000	8726	161.60331
22	1000	1000	9811	54.57641
23	1000	1000	8843	143.75150
24	1000	1000	9550	70.86186
25	1000	1000	9132	107.65606
26	1000	1000	8920	133.09293
27	1000	1000	9049	116.97767
28	1000	1000	8971	126.47214
29	1000	1000	9619	66.13498
30	1000	1000	9367	85.09974

Table 2.1: AMS results for two-to-three jet transitions in QG2D.

The first success is that all ensembles manage to successfully run until completion. This at least shows that the algorithm reliably runs in a finite and feasible time.

We now compare the statistics of the probability estimator \hat{p} for the three groups. The mean probability estimator for the first group is 2.99×10^{-4} with a normalised variance of 9.19. The same quantities for the second group are 1.53×10^{-5} and 2.52, and for the third group they are 1.07×10^{-4} and 0.11. The results in the first two groups are not very statistically useful since their normalised variance is greater than one. In these groups it's clear that 10 ensembles is not a large enough sample (ensemble number 10 alters the estimator from the first group by an order of magnitude). The high variance associated with the r=1 AMS implementation in the second group is not unexpected as each ensemble is waiting for only one rare event. It is encouraging that the variance is markedly reduced from group two to group three. Moreover, the reduction in normalised variance outweighed the extra

computational time taken to find the remaining N-1 transitions, indicating the extra computational effort provides an increased benefit.

To compare DNS with AMS we will use group three. AMS's group three took 15 times more computer time than DNS (approximately 30 weeks to DNS's two). Despite this large factor, AMS still outperformed DNS in the task of *finding* rare events. AMS found 100,000 events to DNS's three, meaning AMS found rare events over two-hundred times faster than DNS. Some of these rare events share ancestry, however even eliminating all such events, AMS finds *distinct* rare events at a faster rate than DNS.

To obtain a confidence interval on the transition waiting time from AMS we assume a central limit theorem for \hat{p} . The 95% confidence interval for the transition waiting time can then be found (e.g. using the reset time τ found in section 2.2.5). The estimated transition waiting time is 5.83×10^5 with the confidence interval $[4.28, 6.55] \times 10^5$. We note that the estimated transition time is similar to the (very approximate) DNS values given above.

To compare with DNS, we assume that DNS rare event observations are Poisson distributed. For a 95% confidence interval to have the same (relative) range magnitude as the AMS interval, we need to observe approximately 80 events (giving the confidence interval for the number of events as [63.4, 99.6]). Using the confidence interval from AMS as a guide, 80 events would require between 18 and 29 weeks of computer time by DNS. This does not include the DNS time taken to transition back to a two jet state after the rare event like Figure 2.3, since if we were looking for two to three jet transitions exclusively then we would reset the DNS simulation after each transition. The conclusion is that DNS would provide a slightly more precise statistic over the same computational time used as for the AMS group three experiment. With a smaller friction parameter so that events are more rare, it is expected that AMS will outperform DNS (due to the $n \propto -\ln p$ relationship).

The single transitions from each of the ensembles 11-18 are plotted in Figure 2.4 and consolidate the previously observed quality that nucleations start in the westward jet [18, 53]. It is most usual for a nucleated jet to form in a westward jet

and grow in zonal velocity until the flow reaches set *B*, but it can be observed that trajectories may take alternative routes to set *B*. The first two simulations of the top row of Figure 2.4 are noteworthy as they clearly experience *two* distinct nucleations before fully transitioning. Note as well that the time taken to transition has large variability. Figure 2.5 plots the second and third wave components of the mean vorticity profile for two of the transitions, highlighting that in some cases there is a relatively direct path between sets *A* and *B*, while in other cases the majority of the progress towards set *B* is made early on, with some time being taken to gain the required jet strength and spacing to actually be in set *B*. Before and after snapshots of potential vorticity are given in Figure 2.6.

Figure 2.7 shows trajectories in the three dimensional $(|\overline{\zeta}_2|, |\overline{\zeta}_3|, |\overline{\zeta}_4|)$ space relevant for the scoring function. The sets A and B are the grey cuboids and transitions are coloured red to blue as they transition from two to three jets. The first panel shows the eight transitions from ensembles 11-18. Of particular interest is the particularly wayward trajectory near set B. This trajectory corresponds to the top left transition in Figure 2.4 and its distinction from other trajectories is capturing the peculiar second nucleation towards the end of the transition. The second panel in Figure 2.7 has been taken from a test ensemble of 100 members (not included in Table 2.1) where all the members transitioned. Similar to previous results [18, 53], the transitions take a non-direct route from A to B, following a curve rather than a line, which at first grows in the direction of increasing the fourth Fourier component magnitude. The ensemble of 100 transitions is seen to have relatively few ancestors at set A; the 100 trajectories are independent realisations once they reach set B, but many of them share the same dynamics for some initial time prior to splitting (seen here as a branching of the trajectories exiting set A).

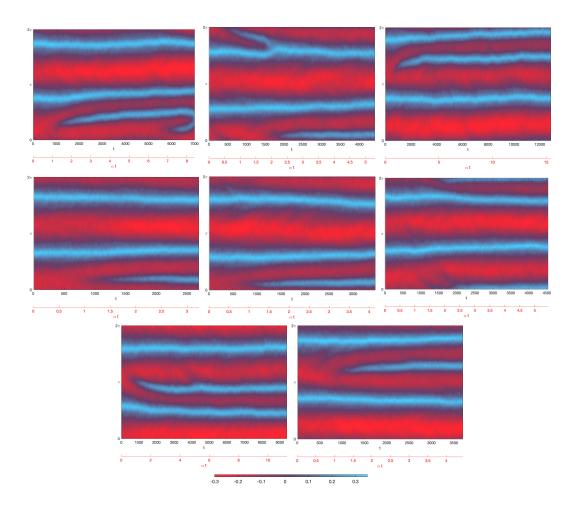


Figure 2.4: A selection of two-to-three jet transitions for QG2D found using AMS, corresponding to the transitions found from ensembles 11-18 in Table 2.1.

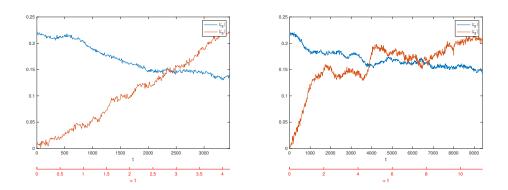


Figure 2.5: Zonal vorticity Fourier component amplitudes for the middle-middle (left) and bottom-left (right) transitions in Figure 2.4. Note the right hand side has a longer time axis.

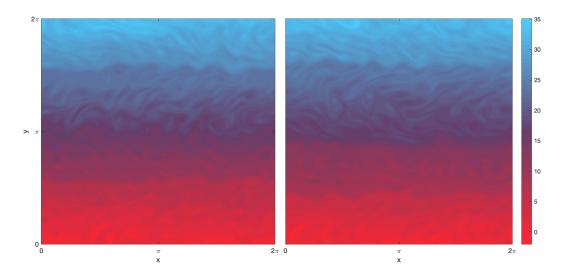


Figure 2.6: The potential vorticity at the beginning and end of the bottom-left transition in Figure 2.4.

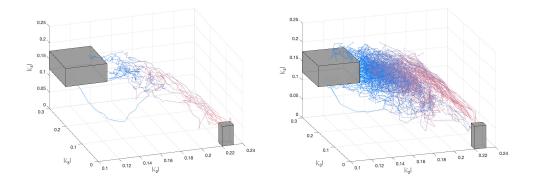


Figure 2.7: Transitions in the reduced zonal Fourier component amplitude space $(|\overline{\zeta}_2|, |\overline{\zeta}_3|, |\overline{\zeta}_4|)$. The boxes are the sets *A* and *B* and the transition from red to blue is from two to three jets. The left panel is for the eight transitions from ensembles 11-18. The right panel plots a full ensemble of 100 members where all trajectories were guided towards set *B*.

2.3 Deterministic Adaptive Multilevel Splitting Algorithm (DAMS)

Many physical models are deterministic, not stochastic. That is to say, once initial conditions are prescribed, the future dynamics are fully determined. Rare events are not confined to stochastic systems, so can a rare event algorithm be devised to find rare events in deterministic models? In this section we describe a deterministic adaptive multilevel splitting algorithm (DAMS) and provide results of the algorithm in the two-layer model. To our knowledge it is the first time the idea has been implemented and the first time any rare event algorithm has been applied to a baroclinic model of turbulence.

2.3.1 The Algorithm

It is clear with a moments thought that AMS cannot be applied to a deterministic model. Without the random forcing, cloned trajectories do not develop distinct dynamics since simulations with any shared ancestry must be identical for all time. This can be overcome by adding a small perturbation at the splitting time. If a system is sufficiently chaotic, the perturbation will grow exponentially with time, leading to cloned trajectories developing distinguished dynamics.

The idea to add a perturbation at splitting times in AMS was first used to study rare energy fluctuations in the deterministic Lorenz '96 model [56] (a low dimensional model of chaos). The perturbation added was a random Gaussian noise at the time of splitting. The justification is that it's a non-physical model of chaos and the emphasis is showing AMS can be modified to successfully find rare events in deterministic chaos for the first time. Similarly, the deterministic collapse of turbulence in plane Couette flow has recently been successfully studied by adding a small random noise at splitting times [57].

We question whether a random perturbation is the most relevant and propose an alternative with some advantages. The perturbation is fully detailed in the following sections 2.3.2–2.3.3.

The rare event algorithm we term DAMS is identical to AMS except line 11 in

Algorithm 1 where it should read "clone S_j as S_1 and perturb", i.e. DAMS perturbs using our new perturbation method.

2.3.2 Splitting Perturbation: Idea and Theory

We propose a more natural method than the random perturbations used previously to split trajectories in deterministic applications of AMS [56, 57]. The method will make use of some general results of chaotic dynamical systems which we introduce here.

Supposing we have an ergodic dynamical system which has been (effectively) run to equilibrium from $t = -\infty$,

$$q_t = \mathcal{N}q \tag{2.3.1}$$

where \mathcal{N} is a time-dependent nonlinear operator acting on q. We may linearise (2.3.1) about a given solution trajectory q(t) to find a *tangent linear model*

$$\tilde{q}_t = \mathcal{L}\tilde{q}. \tag{2.3.2}$$

The time-dependent linear operator \mathcal{L} is dependent on the original trajectory q(t).

As an example, if we consider the material conservation equation

$$q_t = -uq_x - vq_y, \tag{2.3.3}$$

looking for solutions of the form $q+\tilde{q}$ (for some specific solution q) and linearising finds the equation

$$\tilde{q}_t = -u\tilde{q}_x - v\tilde{q}_y - \tilde{u}q_x - \tilde{v}q_y = \left\{-u\partial_x - v\partial_y + q_x\partial_y\nabla^{-2} - q_y\partial_x\nabla^{-2}\right\}\tilde{q}. \quad (2.3.4)$$

The operator in the curly brackets gives the linear operator in this simple example.

In general, given an initial condition \tilde{q}_0 at $t=t_0$, the tangent linear model has the *formal* solution

$$\tilde{q}(t,t_0) = \exp\left(\int_{t_0}^t \mathcal{L} dt\right) \tilde{q}_0, \qquad (2.3.5)$$

where the exponential can be interpreted in the power series sense.

We are now in a position to give a working definition of the maximal Lyapunov exponent

$$\lambda = \sup_{\tilde{q}_0 \in \mathcal{Q}_0} \lim_{t_0 \to -\infty} \frac{1}{t - t_0} \ln \left(\frac{\|\tilde{q}(t, t_0)\|}{\|\tilde{q}_0\|} \right), \tag{2.3.6}$$

where here and throughout $\|\cdot\|$ is the L^2 -norm

$$\|\phi\| = \int_{\mathcal{D}} \phi^2 d\mathbf{x}. \tag{2.3.7}$$

The set Q_0 is the set of all suitable randomly generated initial conditions. In fact, almost all (in the probabilistic sense) $\tilde{q}_0 \in Q_0$ qualify as a supremum in the above, with only sets of measure zero in Q_0 recovering other Lyapunov exponents. Let us denote this 'almost everywhere' subset of Q_0 by \overline{Q}_0 . This allows for a consistent definition of the *maximal Lyapunov vector* (MLV)

$$\tilde{q}_{\text{MLV}}(t) = \pm \lim_{t_0 \to -\infty} \frac{\tilde{q}(t, t_0)}{\|\tilde{q}(t, t_0)\|} \bigg|_{\tilde{q}_0 \in \overline{Q}_0}, \tag{2.3.8}$$

where the sign can be chosen by some criterion to confer uniqueness. A crucial property of the MLV is it is independent of the particular \tilde{q}_0 , as all initial conditions in $\overline{\mathcal{Q}}_0$ develop into the MLV structure. The MLV depends on the underlying nonlinear trajectory and is hence a unique time-evolving property of q(t). In contrast, the maximal Lyapunov exponent, λ , is a property of the dynamical system.

The new splitting procedure for DAMS perturbs a simulation in the direction of the MLV at the time of the split with some relatively small specified magnitude ε_s , i.e. the perturbation at splitting times is $\varepsilon_s \tilde{q}_{\text{MLV}}$.

A perturbation at t=0 is expected to grow like $\exp(\lambda t)$ so an approximate convergence to a MLV is expected for $t\gg \lambda-\lambda_2$ (where λ_2 is the second largest Lyapunov exponent). Provided the splitting time satisfies the inequality, *almost any* change in initial conditions with magnitude $\varepsilon_s \exp(-\lambda t)$ approximately converges to $\varepsilon_s \tilde{q}_{\text{MLV}}$.

The crucial conceptual leap from AMS to DAMS, then, is that we are approxi-

mating a natural exploration of initial conditions. Consequently, the trajectories from DAMS correspond (at least to close approximation) to a realisable deterministic flow. The 'splitting' can be considered as selecting a new initial condition in set A, such that the dynamics are close to the clone up to the splitting time.

2.3.3 Splitting Perturbation: Practical Implementation

To approximate the MLV at time T, we consider the flow history back to $T-\Delta T$ for some $\Delta T>0$. We begin by adding a noise at time $T-\Delta T$ with a small magnitude ε_p . The model can then be integrated and the perturbation monitored. However, since nonlinear effects on the perturbation are unwanted, we must not let the perturbation grow large relative to the flow. To avoid this we employ a simple trick: the perturbation is rescaled back to magnitude ε_p several times throughout the integration over $[T-\Delta T,T]$. Provided ΔT is large enough $(\Delta T\gg \lambda-\lambda_2)$, the final perturbation at T gives an approximation to the MLV direction.

The two-layer model exhibits a single jet state when the parameters are $(W_c, \kappa_M) = (2.2\pi, 0.01)$. We use this parameter setting to test different ε_p and ΔT for the calculation of the linear tangent space. It's determined that $\varepsilon_p = 10^{-12}$ (relatively very small, see discussion below on ε_s for context) and $\Delta T = 25$ with a perturbation rescaling every time unit is sufficient to approximate the MLV. At this value of ΔT the MLV approximation is found to have negligible dependence on the initial perturbation (i.e. the MLV approximations have a small standard deviation magnitude relative to the MLV magnitude).

A difficult component of DAMS is deciding the final splitting perturbation magnitude, ε_s . Since we are using a linear approximation for the perturbation, we should not take ε_s to be large compared to the magnitude of the original flow. However, choosing ε_s too small results in the perturbed clone decorrelating from the original flow at an impractically late time (later than the typical lifetime of a trajectory in the rare event algorithm for example).

To determine a sensible value for ε_s in the two-layer model, we run some experiments with parameters $(W_c, \kappa_M) = (2.2\pi, 0.01)$. For various values of ε_s , we add $\varepsilon_s \tilde{q}_{\text{MLV}}$ at t = 0 and monitor the growth of the perturbation.

Plotted in Figure 2.8 is the magnitude of the perturbation after splitting. Exponential growth is observed until trajectory decorrelation halts the linear theory described above. The perturbation in the two-layer model is taken to be applied to the 'vector' $Q = (Q_1, Q_2)^T$. The typical magnitude of Q in the two-layer model with the above parameters is approximately 3×10^5 . Much of this is contained in the zonal profile due to the imposed temperature gradient. The magnitude of Q without the zonal contribution is approximately 2.5×10^3 . Normalisation of the MLV means the splitting perturbation has magnitude ε_s , hence $\varepsilon_s \ll 10^3$ is the condition for relatively small perturbations. However, this magnitude comparison misses that the MLV has a tendency to be skewed to specific regions of the domain. With $\varepsilon_s = 1$, sometimes the MLV has been found to exceed 10^{-2} times the magnitude of Q on some small part of the domain. It is expected that $\varepsilon_s = 0.25$ will keep the perturbation to less than 10^{-2} times the original trajectory everywhere in the domain. We decide to take $\varepsilon_s = 0.25$. This value also implies a total perturbation magnitude $O(10^{-4})$ of the original trajectory, as well as a decorrelation time within the typical trajectory lifetime.

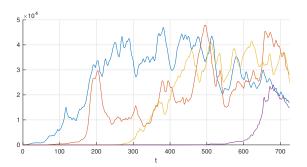


Figure 2.8: The initial perturbation magnitudes are $\varepsilon_s = (1, 0.25, 10^{-4}, 10^{-12})$ with corresponding colours (blue, orange, yellow, purple). i.e. decreasing ε_s corresponds to a delayed perturbation growth and trajectory decorrelation.

Are statistical estimators influenced by the parameter ε_s ? If $\varepsilon_s > 0$ but small enough not to alter trajectories significantly on the lifetimes of trajectories, we cannot expect to find transitions easily through algorithm iterations. For a given computational effort, then, there doesn't appear to be any certainty that the variance of the probability estimator from DAMS will be superior to DNS unless care is taken with the choice of ε_s . As for estimator bias, it's supposed that the exploration close

to some initial condition is enough to sample rare events in an unbiased manner, but unbiasedness is yet to be proven rigorously for DAMS.

At first glance, it's not immediately clear that the decorrelation considerations above aren't relevant to high dimensional AMS problems. For instance, it is not clear that scores will decorrelate over trajectory lifetimes in QG2D if we have weak forcing so that the system is dominated by advective terms. However, it is known that AMS is unbiased regardless and has an upper limit on variance.

2.3.4 Scoring Function

As in QG2D, we may also look at Fourier coefficients of zonal flow in the two-layer model to determine a scoring function. Since the jets are confined to an internal region of the domain with a width related to the equilibrium profile width, we look at coefficients U_k^+ defined as

$$U_k^+ = \frac{1}{2(W_c + 2\sigma)} \int_{-W_c - 2\sigma}^{W_c + 2\sigma} U^+(y) \exp\left(-ik\frac{y + W_c + 2\sigma}{W_c + 2\sigma}\right) dy, \qquad (2.3.9)$$

recalling that W_c is the half width of the flat top region of the equilibrium profile and $\sigma = \pi$ is its exponential decay away from the flat top region. The U_k^+ are essentially the modes of $U^+(y)$ on the 'jet domain' $[W_c - 2\sigma, W_c + 2\sigma]$.

It is found that the values of $|U_1^+|$ and $|U_2^+|$ distinguish the single and double jet states well across all the range of W_c tested. Figure 2.9 shows the Hovmöller contour plot for a simulation with transitions, along with a graph for the corresponding values of $|U_1^+|$ and $|U_2^+|$. The abrupt transitions and their representation by modes share resemblance to Figure 2.3 for QG2D.

It has been found that in the two-layer model the values of $|U_3^+|$ are not very distinct in each state, and we continue with a scoring function which only considers the magnitude of the first two Fourier components. The fact that DAMS, as we shall see, is found to work with a scoring function of only two variables is very reassuring for the potential of future improvement.

The scoring function we will use is essentially the same as the scoring function for the single layer model, i.e. (2.2.2), but in two dimensions. The sets are sensitive

to the model parameters. When $(W_c, \kappa_M) = (2.2\pi, 0.01)$, the sets the single jet state, set A, is defined by

$$|U_1^+| \in [0.5, \infty), \tag{2.3.10}$$

$$|U_2^+| \in [0, 0.25],$$
 (2.3.11)

and the double jet state, set B, is defined by

$$|U_1^+| \in [0, 0.2], \tag{2.3.12}$$

$$|U_2^+| \in [0.55, \infty).$$
 (2.3.13)

The sets for the parameter setting $(W_c, \kappa_M) = (2.6\pi, 0.028)$ are: set A;

$$|U_1^+| \in [0.3, \infty), \tag{2.3.14}$$

$$|U_2^+| \in [0, 0.33],$$
 (2.3.15)

and set B;

$$|U_1^+| \in [0, 0.15],$$
 (2.3.16)

$$|U_2^+| \in [0.5, \infty). \tag{2.3.17}$$

Starting with a primitive scoring function, it is possible to improve with little extra knowledge of the system. As we shall see, a simple improvement on the scoring function in this model could be devised by including the phase of the Fourier components as well as the magnitude. This is because jet states are symmetric about y = 0 which gives them a preferred location in y. This is a quality that is missing in QG2D where the target states are translationally invariant in y from meridional periodicity.

2.3.5 Results

We will begin by looking at results from DNS for the two-layer model. Dozens of long integrations at the 128×128 resolution have been undertaken across the

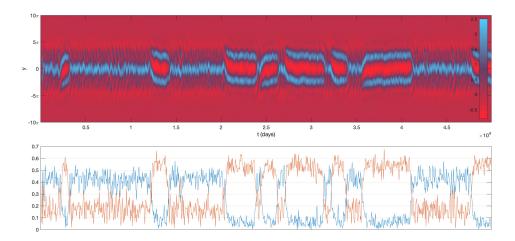


Figure 2.9: Top: Hovmöller of $U^+(y,t)$ in a two-layer flow with parameters $(W_c, \kappa_M) = (3.25\pi, 0.05)$. Below: $|U_1^+|$ (blue) and $|U_2^+|$ (orange) as defined in (2.3.9). The clear flipping of magnitudes indicates that the low wavenumbers will make a good scoring function.

 (W_c, κ_M) parameter space, focusing mostly on regions of the parameter space near to where transitions occur. Figure 2.10 shows the average jet number observed at points in the parameter space.

A curve of interest in the parameter space is the curve where the average jet number is 1.5, indicating the jet is equally likely to be a single or double jet state.

Figure 2.11 provides DNS data for the number of observed transitions and the mean waiting times for various κ_M values with $W_c = 2.6\pi$ fixed. The transition times for each transition direction is plotted on a logarithmic axis with 95% confidence intervals. The transitions times appear nearly linear with κ_M on these axes. This indicates there may be an Arrhenius type law $T_{A\to B} \propto e^{\nabla V/\kappa_M}$. The deviation from this relationship might be expected though when varying κ_M only, since if we don't also adjust W_c accordingly, the expected jet number changes as we change κ_M (and hence the expected underlying transition times).

Following the data from DNS we apply DAMS to two parameter settings: $(W_c, \kappa_M) = (2.6\pi, 0.028)$ and $(W_c, \kappa_M) = (2.2\pi, 0.01)$. The larger friction value has frequent transitions, with ~ 1.5 jets, and we've found accurate statistics from DNS. The aim, since this is the first time DAMS has been used, is to check that DAMS gives comparable results. The smaller friction setting is such that no transitions have

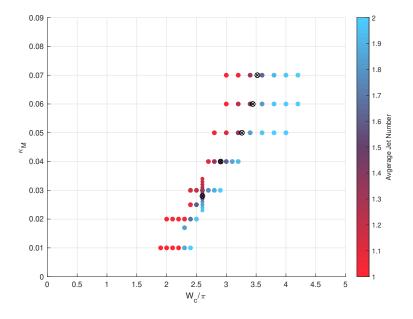


Figure 2.10: Average jet number in the (W_c, κ_M) parameter space found through DNS. The approximate 1.5 jet number locations are given by the black circled crosses.

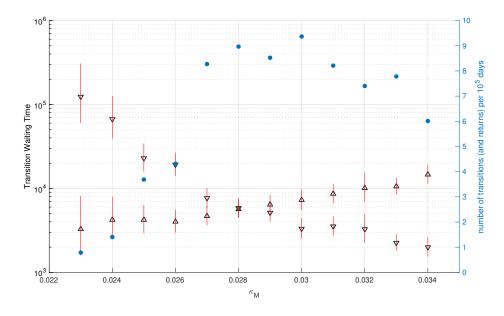


Figure 2.11: Results from DNS varying κ_M with $W_c = 2.6\pi$ fixed (corresponding to the vertical sweep in Figure 2.10). The upward pointing black triangles indicate the mean waiting time to go from a single to double jet. Downwards pointing arrows are the mean waiting time to go from a double to a single jet. Blue dots mark the number of observed transitions per 10^5 time units, indicated by the right hand side axis.

been observed through DNS. However, both a steady single and double jet state have been found with the small friction setting if we use suggestive initial conditions. This has allowed us to define the sets *A* and *B* to apply DAMS.

We perform the algorithm on ten ensembles in each parameter setting with results given in Table 2.2. The first group of ten contains the $(W_c, \kappa_M) = (2.6\pi, 0.028)$ results and the second contains the $(W_c, \kappa_M) = (2.2\pi, 0.01)$ results. The first group uses 500 ensemble members, the latter uses 1000 and the algorithm is run until the entire ensemble reaches set B.

Ens.	W_c	κ_{M}	N	n	p̂
1			500	1100	1.10559×10^{-1}
2			500	1215	0.87822×10^{-1}
3			500	1213	0.88175×10^{-1}
4			500	1236	0.84207×10^{-1}
5	2.6π	0.028	500	1148	1.00430×10^{-1}
6			500	1140	1.02051×10^{-1}
7			500	1226	0.85910×10^{-1}
8			500	1266	0.79299×10^{-1}
9			500	1081	1.14846×10^{-1}
10			500	1057	1.20499×10^{-1}
11			1000	9283	9.25608×10^{-5}
12			1000	11344	1.17733×10^{-5}
13			1000	12711	0.29986×10^{-5}
14			1000	13018	0.22056×10^{-5}
15	2.2π	0.01	1000	11332	1.19155×10^{-5}
16			1000	9766	5.70897×10^{-5}
17			1000	11591	0.91955×10^{-5}
18			1000	11050	1.57996×10^{-5}
19			1000	11277	1.25896×10^{-5}
20			1000	12498	0.37108×10^{-5}

Table 2.2: DAMS results for one to two jet transitions in the two-layer model. For each implementation the algorithm stopped when all members transitioned (r = N).

For the large friction, the reset time is found, from 2000 simulations, as $\tau \sim 450$ with a confidence interval approximately 450 ± 25 (figure 2.12). The large reset time was an oversight in implementation. Set A is larger than ideal, meaning τ (and p) are larger than optimal. Much of the work in this DAMS example is therefore is shifted towards the initialisation steps. To get an accurate waiting time we need thousands of 'reset observations' which are expensive at ~ 450 time units each. Moreover, since the desired sample size of reset times is larger than 1/p for this example, calculating τ provides a reasonable DNS estimate on p. It is found that 208 of the 2000 simulations transitioned to a two jet state.

Continuing with the value $\tau = 450$ anyhow gives a wait time as 4.62×10^3 .

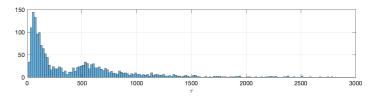


Figure 2.12: Histogram of 'reset times' from 2000 simulations. $(W_c, \kappa_M) = (2.6\pi, 0.028)$.

Assuming a central limit theorem for \hat{p} , the 95% confidence interval for the waiting time (again, with $\tau = 450$ assumed precise) is $[4.23, 5.09] \times 10^3$. This is in agreement with the DNS result in Figure 2.11. The confidence intervals have an overlap region and cover a similar range in each case. The DNS result in Figure 2.11 took approximately one week of computer time to find 66 transitions while DAMS took 8 weeks to find 5000 transitions. Further, about 10 days of DAMS was calculating τ , which is equivalent to a DNS sampling of 208 rare events which provides a smaller confidence interval for \hat{p} from DNS than DAMS (about half the size), despite using eight times less computer time.

In conclusion we can say that DAMS is more efficient at finding transitions at the high friction parameter setting, whereas DNS provides a better estimate for the transition waiting time.

For the low friction setting a direct comparison is not entirely possible since DNS did not find any transitions (the simulation was stopped after 8 weeks)! The mean \hat{p} value is 2.94×10^{-5} with relative variance 1.79. The mean reset time from 1000 simulations has been calculated as $\tau = 262$. The implied waiting time from τ and the mean \hat{p} value is of the order 1×10^7 , which corresponds to around 10 weeks of computer time for DNS. DAMS took a total of ~ 15 weeks, but if we increased the algorithm iterations, the statistics so far indicate that DAMS will perform better than DNS at this friction value. While caution should be exercised reading too much into these (high variance) statistics, it is clear that DAMS is orders of magnitude more efficient at finding rare events than DNS. In the same amount of time that DNS found zero transitions, DAMS found approximately 5000.

Figures 2.13–2.17 below show data from DAMS and some DAMS discovered transitions. Figure 2.13 plots the variables U_1^+ (red) and U_2^+ (blue) on an Argand

diagram, giving a clear indication of how to improve the scoring function. The two distinct routes from set A to set B in the Argand diagrams also reveals the approximately even split in the probability that the second jet will nucleate on the north or south of the existing single jet. Such a split is not visible in the plots in the $(|U_1|, |U_2|)$ space. Figure 2.17 shows how the lower friction case allows for a jet oscillation before reaching the two jet state (this behaviour is actually visible near the figure '8' shape made by the blue line plots in the right-most Argand diagram in Figure 2.13). Most transitions however follow the more direct approach and don't oscillate, e.g. Figure 2.16. It is clear that DAMS has found a variety of transition behaviours and jet dynamics far more efficiently than DNS would be able to in the small friction regime.

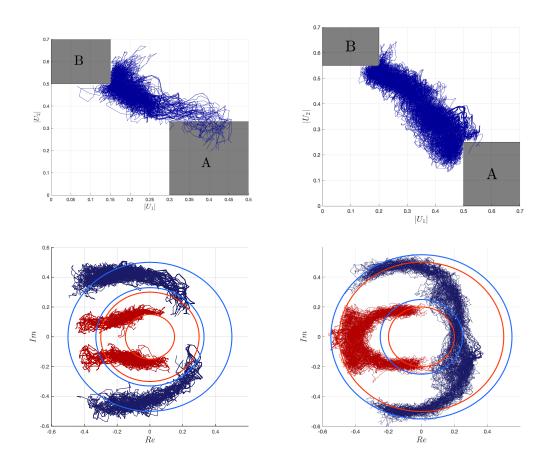


Figure 2.13: Transitions from a single to double jet for friction values $\kappa_M = 0.028$ (left) and 0.01 (right) found using DAMS. In each plot there are 200 transitions taken randomly from an algorithm implementation. The top row is in $(|U_1^+|, |U_2^+|)$ -space, and the bottom row is an Argand diagram with U_1^+ plotted in red and U_2^+ plotted in blue. The grey boxes/coloured circles give the boundaries of sets A and B.

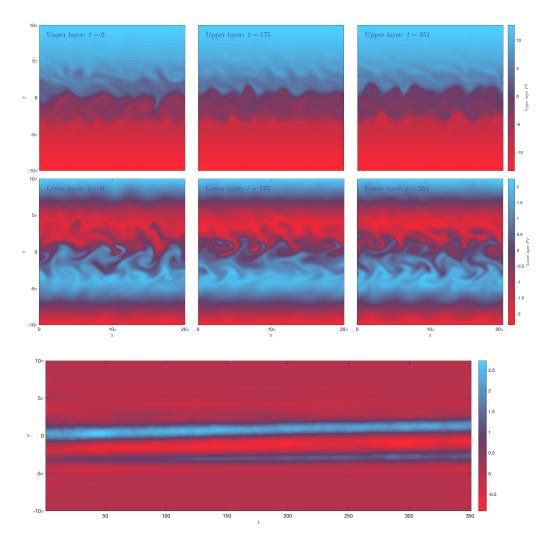


Figure 2.14: Example of a transition with $\kappa_M = 0.028$ where a second jet forms 'south' of the single jet.

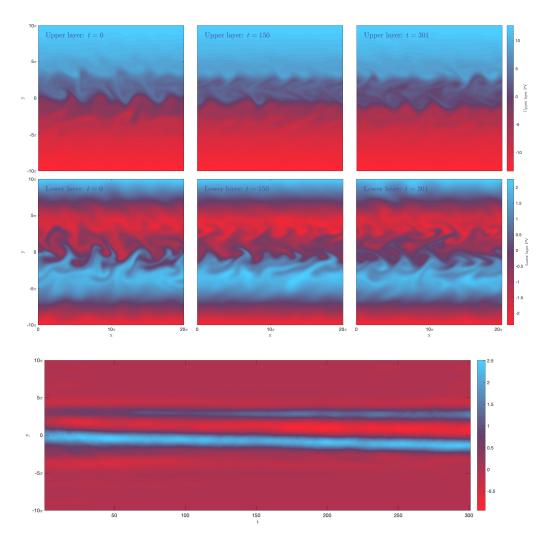


Figure 2.15: Example of a transition with $\kappa_M = 0.028$ where a second jet forms 'north' of the single jet.

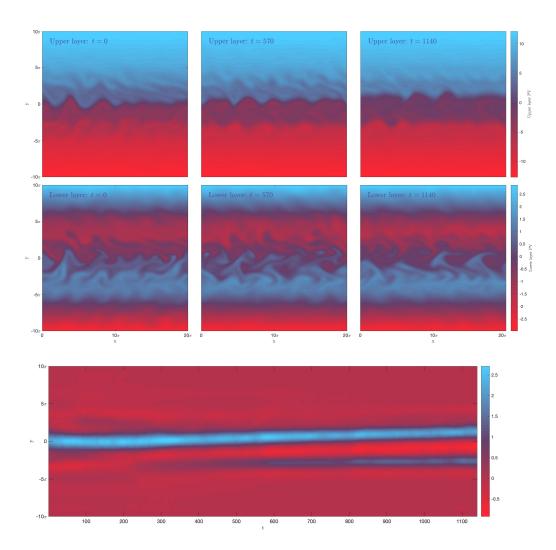


Figure 2.16: Example of a transition with $\kappa_M = 0.01$.

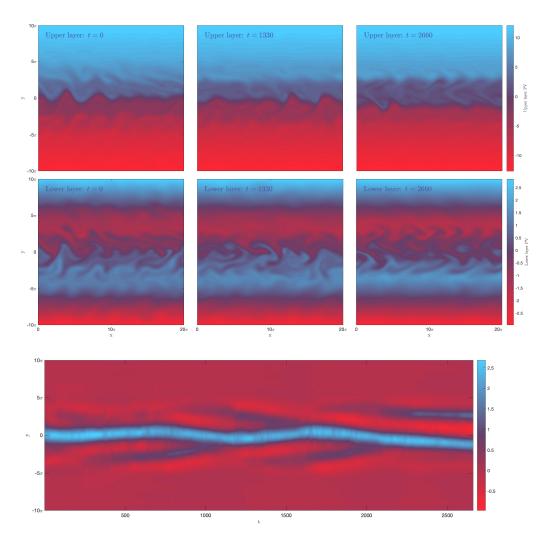


Figure 2.17: Example of a transition with $\kappa_M = 0.01$ where the jet oscillates before a second jet gains sufficient strength to reach set B.

Chapter 3

Second Order Closure Models

This Chapter provides an overview of the motivating idea of statistical truncations for turbulence. A well-known statistical closure model of stochastically forced, quasi-geostrophic β -plane turbulence is reviewed and introduced mathematically, ready for theoretical analysis in Chapter 4.

3.1 Introduction

Previous Chapters have emphasised the difficulties of integrating turbulence numerically over multiple eddy turnover times. The computational expense not only limits the investigation into rare, transitional, weather events, but in general any long-time study of the jet dynamics. For instance, gathering the statistics of turbulent fluctuations about a fixed state requires integration over multiple eddy turnovers [58]. Recent decades, however, have seen the emergence of a statistical approach which aims to access a turbulent flow's statistics *directly*, rather than empirically averaging a particular realisation of the flow. The resulting model (section 3.3) forms the basis for the remaining Chapters of this thesis.

A key idea is that the statistics of the QG2D model are fully described by the cumulants of its vorticity field, and the aim is to find equations for the cumulants.

The first few cumulants are

$$C_{1}(\mathbf{x}) = \mathbb{E}\left[\zeta\left(\mathbf{x}\right)\right],$$

$$C_{2}(\mathbf{x}_{1}, \mathbf{x}_{2}) = \mathbb{E}\left[\left(\zeta\left(\mathbf{x}_{1}\right) - \mathbb{E}\left[\zeta\left(\mathbf{x}_{1}\right)\right]\right)\left(\zeta\left(\mathbf{x}_{2}\right) - \mathbb{E}\left[\zeta\left(\mathbf{x}_{2}\right)\right]\right)\right],$$

$$C_{3}(\mathbf{x}_{1}, \mathbf{x}_{2}, \mathbf{x}_{3}) = \mathbb{E}\left[\left(\zeta\left(\mathbf{x}_{1}\right) - \mathbb{E}\left[\zeta\left(\mathbf{x}_{1}\right)\right]\right)\left(\zeta\left(\mathbf{x}_{2}\right) - \mathbb{E}\left[\zeta\left(\mathbf{x}_{2}\right)\right]\right)\left(\zeta\left(\mathbf{x}_{3}\right) - \mathbb{E}\left[\zeta\left(\mathbf{x}_{3}\right)\right]\right)\right],$$

where the expectation is specific to the problem at hand (Note that while the first three cumulants are just centred moments, this is not the case from the fourth cumulant onwards ¹.

An equation for the first cumulant is easily found by taking the expectation of the evolution equation for ζ (for example equation 1.1.21). The crux of turbulence formation lays in the quadratic nonlinearity of the Navier-Stokes equations, and generally the quadratic nonlinear term couples the first and second cumulant. Continuing on to higher cumulants, the equation for the n^{th} cumulant can be derived through manipulation of the evolution equation and the previous cumulants. The nonlinear quadrature introduces the $(n+1)^{th}$ cumulant into the n^{th} cumulant's equation, meaning we are left with an infinite hierarchy of cumulant equations. This presents an immediate hurdle from a practical standpoint.

There are two main approaches to deal with the infinite hierarchy. The first is to use 'functional' or 'field-theoretic' methods, pioneered by [59], which aim to solve the entire probabilistic landscape. Whilst the functional approach has a theoretical appeal it has received limited attention in the geophysical fluid dynamics literature. It is, frankly, beyond the scope of this work, and the description of the methods as being "mostly useless in understanding turbulence" [60], whether correct of not, led the author to favour an alternative route. The second approach - which is the focus for this Chapter and underpins the second half of this thesis - truncates, or closes, the cumulant hierarchy. A brief overview of the literature aiming to close the hierarchy of idealised isotropic, homogenous turbulence is presented in [60]. The impression is that it is not clear that closure theories hold a promising

This follows from the definition of the cumulant: the n^{th} cumulant is the n^{th} derivative of $\ln \mathbb{E}(e^{tX})$ with respect to t, evaluated at t = 0)

route forward for *idealised* turbulence, as non-linearity is such a crucial feature that truncating the hierarchy at low order is insufficient, requiring higher order cumulants to be artificially parameterised. It turns out however, mercifully, that for strong turbulence-mean interaction problems (considered 'less nonlinear' than idealised turbulence [29]), relatively simple closures have been found that offer, at the minimum, qualitatively good results. It follows that a key motivator for closure models is that many atmospheric flows exhibit strong wave-mean interaction, in particular the upper atmospheres of giant planets Jupiter and Saturn [e.g. 5, and references therein].

Retaining only the first and second cumulants, referred to as closing the hierarchy at the second order, has been studied extensively. The successful framework is referred to as stochastic structural stability theory, or simply SSST/S3T (following [61]), or CE2 (following [62]). CE2 was developed after S3T, and it is subtly different in the rationale, but the two theories result in a mathematically identical set of equations when applied to the QG2D model. The CE2 approach is not limited to QG2D, for example CE2 can be applied to two-layer baroclinic flow [63, 64, 65], the Lorenz63 system [66], jets and vortices [67], and magnetohydrodynamics [68, 69].

Statistical truncation at the second order is closely related to a *quasi-linear ap- proximation* (QL) – which is described in section 3.2 – with CE2 being interpreted as modelling QL statistics. It has been proposed [17, 70] that the QL model is formally asymptotic to the fully nonlinear QG2D model in the limit of large zonostrophy. This claim is actually questioned in section 4.5, however, there is a large amount of evidence supporting the practical use of the QL approximation and the CE2 model for predicting seemingly accurate jet structures at large zonostrophy [see for instance 62, 71, 72, 73]. This is exciting as the QL/CE2 framework is particularly amenable to theoretical analysis.

In this thesis we use CE2 in its original form for the specific problem of β plane turbulence (as in, say, [71]). Variations of CE2 that capture higher order
correlations have been formulated, taking names such as CE2*, CE2.5 and CE3
[72, 74]. Typically, these suffer from requiring extra model parameters or numerical

intervention to avoid 'negative probability densities' [75]. In fact one of the brilliant properties of CE2 is that quadratic quantities are conserved, providing a physical realisability (non-negative probability densities). Another drawback with extensions to the CE2 model is that they are prone to the 'curse of dimensionality' (they occupy an increasing number of dimensions depending on complexity). An impressive extension of CE2, which doesn't suffer the aforementioned problems, is the promise of a generalised CE2 model (e.g. the natural extension of the generalised QL model [75]). Here, the closure permits the quasi-linear theory to include *some* of the nonlinear eddy-eddy interactions which are discarded in QL theory. The model needn't be much more numerically demanding, but can achieve far greater accuracy in simulations compared to CE2 in a variety of applications [68, 69, 76]. This is specifically true in QG2D with zonostrophy values below which QL simulations cope [75]. Similar to the generalised CE2 is a version of CE2 that uses a proper orthogonal decomposition, whereby only the most relevant structures in the second cumulant are retained [77].

3.2 Quasi-Linear Approximation

The quasi-linear (QL) approximation to the QG2D vorticity equation (1.1.21), has as its starting point the Reynolds decomposition

$$\varphi = \bar{\varphi} + \varphi'$$

where $\bar{\varphi}$ denotes the zonal mean of a quantity φ , and φ' its non-zonal perturbation. Decomposing the vorticity equation (1.1.21), and neglecting the non-zonal terms given by $J(\psi', \zeta') - \overline{J(\psi', \zeta')}$ in the perturbation equation, results in the QL equations

$$\bar{\zeta}_t = -(\bar{v}'\bar{\zeta}')_y - \mu\bar{\zeta} + (-1)^{n+1}v_{2n}\partial_v^{2n}\bar{\zeta} + \sqrt{\varepsilon}\bar{\eta}$$
(3.2.1)

$$\zeta_t' = -\bar{u}\zeta_x' - v'\bar{\zeta}_y - \beta v' - \mu \zeta' + (-1)^{n+1}v_{2n}\nabla^{2n}\zeta' + \sqrt{\varepsilon}\eta'.$$
 (3.2.2)

Note that we have used the dimensional QG2D model (1.1.21), rather the non-dimensionalised form (1.1.61). This keeps consistency with the literature, and all subsequent QL/CE2 equations and results can easily be interpreted under the non-dimensionalisation used in Chapter 1.

The QL approximation is best understood in terms of wave mode interactions in the zonal direction. Considering the zonal Fourier transform of the QG2D equation (1.1.21)

$$\frac{\partial}{\partial t}\zeta_k = -\{J(\psi,\zeta)\}_k - \mathrm{i}\beta k\psi_k - \mu\zeta_k - \nu_{2n}(k^2 - \partial_{yy}^2)^{2n+2}\zeta_k + \sqrt{\varepsilon}\eta_k, \quad (3.2.3)$$

where a subscript k denotes the kth zonal Fourier component. Zonal eddy-eddy interactions occur only in the quadratic nonlinearity. This can be decomposed into triad interactions

$$\{J(\boldsymbol{\psi},\boldsymbol{\zeta})\}_{k} = \left\{ (u\boldsymbol{\zeta})_{x} + (v\boldsymbol{\zeta})_{y} \right\}_{k} = ik(u\boldsymbol{\zeta})_{k} + \partial_{y}(v\boldsymbol{\zeta})_{k} = \sum_{j+l=k} \left(iku_{j}\boldsymbol{\zeta}_{l} + \partial_{y}v_{j}\boldsymbol{\zeta}_{l} \right).$$

$$(3.2.4)$$

By construction, the QL approximation is equivalent to restricting the interactions to only those with j, l or k equal to 0, i.e.

$$\{J(\psi(y),\zeta(y))\}_{k} = \begin{cases} ik\overline{u}\zeta_{k} + \partial_{y}\overline{\zeta}v_{k}, & \text{if } k \neq 0, \\ \partial_{y}\sum_{j}v_{j}\zeta_{-j}, & \text{if } k = 0. \end{cases}$$
(3.2.5)

This is the same as saying the QL equations allow for nonlinear interactions only between pairs of wave modes and the zonal mean flow, and not between triads of wave modes (these restrictions are shown schematically in Figure 3.1). This excludes the possibility of a spectrally local inverse cascade of energy, as in the classical Kraichnan-Batchelor 2D turbulence theory, [3, 78, 4].

Aside from the advantages of the QL model's theoretical accessibility compared to its nonlinear counterpart, the QL model is also numerically much faster than the nonlinear model. Using spectral methods the QL model only requires a one-dimensional meridional Fourier transform to solve the nonlinear terms, whereas the

fully nonlinear model must use two-dimensional Fourier transforms.

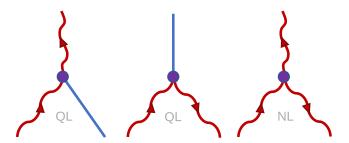


Figure 3.1: Diagram of triad interactions. From left to right: wave-mean interaction where a wave interacts directly with the zonal flow, eddy-eddy interaction of equal and opposite signed waves resulting in Reynolds stress on the mean flow, and lastly all other eddy-eddy interactions not involving mean zonal flow. The first two interactions are retained for QL theory while the final interaction on the right is dropped.

3.3 CE2/S3T: Theory, Derivation and Equations

3.3.1 Derivation

Following e.g. [62], the CE2 equations are obtained by introducing an averaging operator $\langle \cdot \rangle$, which in different interpretations is an ensemble average over the noise process η , or a zonal average [71], or a temporal average [17], or a more sophisticated average over persistent coherent structures [77]. In any case, $\langle \cdot \rangle$ satisfies the Reynolds decomposition property $\langle \varphi_1 \varphi_2 \rangle = \langle \varphi_1 \rangle \langle \varphi_2 \rangle - \langle \varphi_1' \varphi_2' \rangle$, where $\varphi_i' = \varphi_i - \langle \varphi_i \rangle$. The CE2 equations are most conveniently formulated for the averaged zonal velocity $U = \langle \overline{u} \rangle$ and the second cumulant of vorticity

$$\mathcal{Z}(\mathbf{x}_1, \mathbf{x}_2, t) = \langle \zeta'(\mathbf{x}_1, t) \zeta'(\mathbf{x}_2, t) \rangle. \tag{3.3.1}$$

(Note that although its related, \mathcal{Z} here is not the same as the enstrophy which uses the same symbol in Chapter 1). To derive the CE2 equations, note that the time derivative of \mathcal{Z} is

$$\partial_t \mathcal{Z}(\mathbf{x}_1, \mathbf{x}_2, t) = \langle \partial_t \zeta'(\mathbf{x}_1) \zeta'(\mathbf{x}_2, t) + \zeta'(\mathbf{x}_1, t) \partial_t \zeta'(\mathbf{x}_2) \rangle. \tag{3.3.2}$$

Now considering the non-zonal part of the QG2D equation (1.1.21)

$$\zeta_t' = -\overline{u}\zeta_x' - v'\overline{\zeta}_y - (u'\zeta')_x - (v'\zeta')_y + (\overline{v'\zeta'})_y$$
$$-\beta\psi_x' - \mu\zeta' - v_{2n}(-\nabla^2)^n\zeta' + \sqrt{\varepsilon}\eta, \qquad (3.3.3)$$

an equation for the second cumulant is found by:

- multiplying (3.3.3) at position \mathbf{x}_1 by $\zeta'(\mathbf{x}_2)$,
- taking the $\langle \cdot \rangle$ average, and
- adding the same result again with x_1 and x_2 interchanged.

The result is

$$\frac{\partial}{\partial t} \mathcal{Z}(\boldsymbol{x}_{1}, \boldsymbol{x}_{2}) = \left[-U_{1} \partial_{x_{1}} - (\beta - U_{1}'') \partial_{x_{1}} \mathcal{L}_{1} - \mu - v_{2n} (\nabla_{1}^{2})^{n} \right] \mathcal{Z}(\boldsymbol{x}_{1}, \boldsymbol{x}_{2}) + \left[-U_{2} \partial_{x_{2}} - (\beta - U_{2}'') \partial_{x_{2}} \mathcal{L}_{2} - \mu - v_{2n} (\nabla_{2}^{2})^{n} \right] \mathcal{Z}(\boldsymbol{x}_{1}, \boldsymbol{x}_{2}) + \left\{ 3\text{rd order terms} \right\} + \varepsilon \Pi(\boldsymbol{x}_{1}, \boldsymbol{x}_{2}), \tag{3.3.4}$$

where subscripts n = 1, 2 denote the reference variable, $\mathcal{L}_n = \nabla_n^{-2}$ is the inverse Laplacian acting on the n^{th} variable (so that $\psi(\mathbf{x}_n) = \mathcal{L}_n q(\mathbf{x}_n)$, and $\Pi = \langle \eta'(\mathbf{x}_1, t) \eta'(\mathbf{x}_2, t) \rangle$ is the correlation matrix defining the stochastic forcing η . Also U_n'' denotes the second derivative of U with respect to y at latitude y_n .

A key point is that due to taking the stochastic average $\langle \cdot \rangle$ the cumulant equations are deterministic. The CE2 equations are found by dropping the third order terms (closing the hierarchy at the second order), and recognising that we can write $\mathcal{Z} = \mathcal{Z}(x,y_1,y_2)$ for $x = x_1 - x_2$. The full CE2 equations are

$$U_{t} = -\langle u'v' \rangle_{v} - \mu U - (-1)^{n} v_{2n} \partial_{v}^{2n} U, \qquad (3.3.5)$$

$$\partial_{t} \mathcal{Z} + (U_{1} - U_{2}) \,\partial_{x} \mathcal{Z} = -\left(\left(\beta - U_{1}^{"}\right) \nabla_{1}^{-2} - \left(\beta - U_{2}^{"}\right) \nabla_{2}^{-2}\right) \partial_{x} \mathcal{Z}$$

$$-2\mu \mathcal{Z} - (-1)^{n} v_{2n} \left(\nabla_{1}^{2n} + \nabla_{2}^{2n}\right) \mathcal{Z} + \varepsilon \Pi.$$
(3.3.6)

To close the equations it is necessary to evaluate the Reynolds stress term $\langle u'v' \rangle$

appearing in (3.3.5) in terms of \mathcal{Z} . From the relationship between relative vorticity ζ' and velocity, the required result is

$$\langle u'v'\rangle(y) = -\left[\partial_{y_1}\partial_x\nabla_1^{-2}\nabla_2^{-2}\mathcal{Z}(x,y_1,y_2)\right]_{x\to 0,\ y_1\to y_2=y}.$$
 (3.3.7)

3.3.2 The Link Between CE2 and QL

We note that dropping the third order terms neglects terms which are cubic in primed quantities, hence the CE2 model is often described as providing a model for the QL statistics. It is difficult to discern what exactly this means. After all, a *literal* average of QL realisations will be zero since there is no meridional preference in the system.

To understand how to interpret CE2, we point out that there are actually more subtle third order terms which are neglected, e.g. for the first term on the right-hand side of (3.3.4), we must make the approximation that

$$\langle \overline{u'}v'\zeta'\rangle \simeq \langle \overline{u'}\rangle\langle v'\zeta'\rangle.$$
 (3.3.8)

This approximation neglects interactions between perturbations against the $\langle \cdot \rangle$ mean. The physical interpretation here really depends the definition of $\langle \cdot \rangle$. One interpretation is that the $\langle \cdot \rangle$ average captures the collective average action on the mean flow from all possible turbulent realisations, *without* taking into account the individual development of the mean flow from any particular individual realisation. To say CE2 is 'a model for the QL statistics', means that CE2 averages the QL processes over a background flow profile that is (a) universal (to all the processes) and (b) develops according to the averaged Reynolds stress from all QL processes.

3.3.3 CE2 in Zonal Fourier-Space

A useful practical tool is that the second cumulant equation (3.3.6) may be Fourier transformed in x for

$$\partial_{t}\tilde{\mathcal{Z}}_{k}+\mathrm{i}k\left(U_{1}-U_{2}\right)\tilde{\mathcal{Z}}_{k}=-\mathrm{i}k\left\{ \left(\beta-U_{1}''\right)\nabla_{1}^{-2}-\left(\beta-U_{2}''\right)\nabla_{2}^{-2}\right\}\tilde{\mathcal{Z}}_{k}-2\mu\tilde{\mathcal{Z}}_{k}+\varepsilon\tilde{\Pi}_{k},\tag{3.3.9}$$

and may be written in matrix form

$$\partial_t \tilde{Z}_k + \mathcal{A}_k \tilde{Z}_k + \tilde{Z}_k \mathcal{A}_k^{\dagger} = \varepsilon \tilde{\Pi}_k, \tag{3.3.10}$$

where \dagger is the complex transpose and \mathcal{A}_k is an operator which depends on U(y). For clarity the tilde ($\tilde{\cdot}$) has been introduced here to denote the zonal Fourier transform (it has been unnecessary until now, but it provides helpful distinction in the following Chapter). It is worth noting here that

$$\tilde{Z}_k(y_1, y_2) = \langle \tilde{\zeta}_k(y_1) \tilde{\zeta}_{-k}(y_2) \rangle, \tag{3.3.11}$$

where k > 0. In general the CE2 system is solved by integrating (3.3.5) and (3.3.10) forward in time until an approximate equilibrium is reached. When CE2 is solved in this thesis, we use MATLAB's inbuilt ODE45 solver to integrate the first and second cumulant equations (3.3.10) and (3.3.10). Spectral methods are used in the zonal direction for numerical simulations, and standard finite difference methods in the meridional. The QL approximation means the number of active zonal modes is limited to the largest forced wavenumber. With caution that crucial mechanisms at shorter zonal waves are not being missed (such as mean flow instabilities), there is opportunity to save computation time by dropping the inactive modes from the QL/CE2 calculations. The issues around mean flow instabilities are attended to in detail in the following chapter.

We end by noting that the equilibrium equation

$$\mathcal{A}_{k}\tilde{Z}_{k} + \tilde{Z}_{k}\mathcal{A}_{k}^{\dagger} = \varepsilon\tilde{\Pi}_{k}, \tag{3.3.12}$$

is a Lyapunov equation which can be solved readily numerically at each wave k with MATLAB's inbuilt Lyapunov solver. The solution provides the exact statistics of the QL system over a fixed background flow U(y).

3.4 Results

The CE2 equations have been shown to capture jet formation and maintenance, predicting at least qualitatively the spacing and structure of jets [71, 62, 72, 73]. For a brief review, we report a single result from NL, QL and CE2 simulations in Figure 3.2. In each set of results, the NL, QL and CE2 simulations predict the same number of jets and feature east-west asymmetry. The zonal mean velocity U at $t=5/\mu$ is plotted in Figure 3.3, where the CE2, QL and NL results all agree qualitatively well with each other. CE2 has been shown to break down for this problem as zonostrophy is reduced and the physical justification of the closure is broken [71]. The tendency seen in the Hovmöller plots for QL jets to be more latitudinally steady than NL jets has been observed previously [79]. The physical parameter settings are $(\mu, \varepsilon, \beta, L_d, L_f) = (0.0002, 0.0002/(2\pi)^2, 3, 1, 1/16)$, which gives $(Z, Q, F, \alpha, \beta') = (5.1, 4.3, 1.7, 0.00028, 3)$. The forcing is an annulus centred on the wavenumber $k_f = 16$, and a 256×256 grid with $v_4 = 2.5^{-8}$ was used.

Having introduced the CE2/QL framework, we are now in a position to move on to Chapter 4 where we approach its solution theoretically.

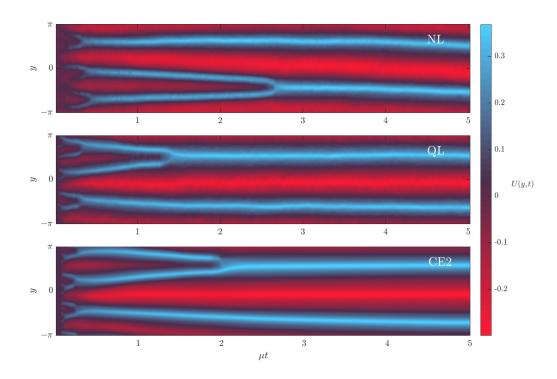


Figure 3.2: Top to bottom: Nonlinear, quasi-linear and CE2 Hovmöller contour plots of U(y,t).

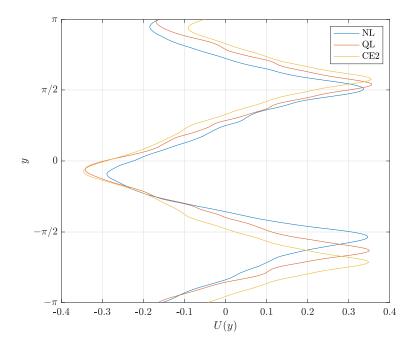


Figure 3.3: Line plots of U(y) at $t = 5/\mu$ for the nonlinear (nlue), quasi-linear (orange) and CE2 (yellow) simulations. These are the same simulations as plotted in Figure 3.2.

Chapter 4

A Local Closure Theory for Momentum Flux in the Limit of Small Scale Forcing

In this Chapter zonal jet formation in β -plane turbulence is investigated with the focus on whether an accurate closure can be developed for the eddy momentum fluxes due to small-scale random forcing. The approach of Srinivasan and Young [80] is developed to give a relatively simple expression for the local Reynolds stress due to the random forcing in the QG2D model when the characteristic length scale is much less than the jet spacing. In typical jet flows, however, it is demonstrated that the Srinivasan-Young flux is not the full story because momentum fluxes due to jet-scale waves, present as a result of distinct barotropic instabilities, also play a key role in the momentum balance. The current Chapter provides a full theoretical analysis of the Srinivasan-Young closure, setting the scene for its practical use as a tool for parameterising small scale turbulence in Chapter 5.

4.1 Introduction

Recall that the QG2D takes the doubly-periodic domain of dimension $2\pi L_d \times 2\pi L_d$, and the most physically significant nondimensional parameters characterising QG2D

are (e.g. section 1.1.5)

$$Z = \frac{L_{Rh}}{L_{\beta}}, \quad Q = \frac{L_d}{L_{Rh}}, \quad F = \frac{L_f}{L_{Rh}},$$
 (4.1.1)

where Z is zonostrophy, Q is a jet quantisation parameter, and F is a forcing number (the alternative parameters α and β' used in Chapters 1 and 2 can be derived from Z and Q as described in section 1.1.5). The focus is on the regime of most relevance to the giant planets, $Z \gg 1$ and $F \ll 1$. In this regime, it has been suggested that F has a subtle effect on the jet dynamics [see e.g. 15, where a range of different forcings are considered], with an apparent stronger tendency to 'potential vorticity staircase' [24] formation at larger F. A more detailed examination of the role of F is provided in [20] to the same conclusion. In the small F limit, the excited waves are short compared to the scale of the jets, and will therefore only 'see' a linear shear flow locally. It turns out that the exact solution for the second-order moments in the CE2 system, in the special case where the mean flow is a linear shear flow, is known [80, SY14 hereafter]. Solutions following the SY14 approach appear to hold the promise of a local closure for the ensemble average wave momentum flux $\langle u'v' \rangle$. Such a closure – a long-time goal of researchers – would allow the cumulant hierarchy to be closed at the first order. The result would be a single equation describing the evolution of the ensemble mean zonal wind U(y,t).

Recently, using a different but essentially equivalent approach to SY14, [23, WB19 hereafter] have shown that in the joint limits $Z\gg 1$, $F\ll 1$, under some mild restrictions on the forcing η to be discussed below, the momentum flux closure becomes independent of the details of η and is given by

$$\langle u'v'\rangle = \varepsilon/U_{\rm y}.$$
 (4.1.2)

A similar closure (modified slightly for polar geometry) has been postulated for the related problem of vortex condensate formation in forced 2D Navier-Stokes turbulence [81], and can be used to predict the radial structure of the vortex condensate, thus demonstrating that a local closure theory can be successful in the right setting.

However, in the β -plane jets problem – as WB19 recognised – it is clear that the closure (4.1.2) is inadequate. First, (4.1.2) fails because the equilibrated flow must include jet minima and maxima where $U_y = 0$, and where the predicted momentum fluxes of (4.1.2) are therefore singular. Second, (4.1.2) does not depend on β , and therefore it cannot possibly lead to equilibrated jets with spacing on the Rhines scale $L_{\rm Rh}$. These observations, however, do not necessarily preclude a local theory. An extended local theory can be hypothesised, which could allow for higher order terms in F to introduce dependency on β and the flow curvature U_{yy} , and which would require a particular focus on boundary layer regions where $|U_y| = O(F)$. In principle such a theory might resolve both of the above issues and lead to a closure that is valid everywhere in the flow.

Using a series of carefully designed CE2 and fully nonlinear numerical simulations, the aim is to demonstrate the following:

- That in the right parameter regime ($F \ll 1, Z \gg 1$) the SY14 approach *does* lead to accurate predictions for the Reynolds stresses $\langle u'v' \rangle$ due to the waves induced by the small-scale forcing, including at jet maxima and minima where (4.1.2) breaks down.
- That in a typical equilibrated state, both the eastward and westward jets fluctuate around a state which is marginally unstable to barotropic instability. Jet-scale barotropic waves due to these instabilities are present on both jets, and these have a distinct characteristic structure on each jet. Momentum fluxes due to these waves which are here disentangled from momentum fluxes due to the small-scale forcing have a key role in the momentum balance of the jets.
- Because the barotropic waves emerge on the jet scale, i.e. they are global, a
 purely local closure for the jet profile of the type suggested above is doomed
 to failure.

The plan of the current Chapter is to review the results of SY14 and reconcile them with those of WB19, clarifying the limit in which (4.1.2) is valid. Then

in section 4.3, CE2 and QL simulations designed to demonstrate the main points above are described and the results are analysed. A careful study of the linear stability properties of both the eastward and westward jets is performed to clarify the role of the waves generated by barotropic instability in the momentum balance of the equilibrated flow. The Chapter ends with sections 4.4 and 4.5 where further extensions of the SY14 local theory are discussed such as introducing time variability and considering the momentum flux components $\langle u'u' \rangle$ and $\langle v'v' \rangle$. The latter sections highlight the key ideas the author would recommend for extending this work in the future, while the former sections lay the ground-work for the final Chapter of this thesis.

4.2 The SY14 local theory

4.2.1 The SY14 result for the momentum flux in a linear shear flow

A re-working of the SY14 solution to the second cumulant equation (3.3.6), for the special case in which $U(y) = \gamma y$ is a steady, constant shear flow, is presented. The aim is to obtain a relatively simple and easy to evaluate formula for the steady momentum flux $\langle u'v' \rangle$, obtained from the solution to (3.3.6) as $t \to \infty$, which will be used to understand the outcome of the numerical calculations below.

Recall that the focus is restricted to a time-independent homogeneous forcing covariance $\Pi(\mathbf{x})$, which takes the argument $\mathbf{x} = (x, y)^T = \mathbf{x}_1 - \mathbf{x}_2$, and its Fourier transform $\hat{\Pi}(\mathbf{k})$ defined by

$$\hat{\Pi}(\mathbf{k}) = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} \Pi(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}} d\mathbf{x}, \quad \Pi(\mathbf{x}) = \int_{\mathbb{R}^2} \hat{\Pi}(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}} d\mathbf{k}.$$
 (4.2.1)

(Here an infinite domain has been introduced for this set-up, and in a slight abuse of notation a hat is used for both continuous and discrete Fourier transforms, with the distinction being obvious depending on the set-up). The key points related to the forcing (as described in Chapter 1) are that $\hat{\Pi}(\mathbf{k}) = \langle \hat{\eta}(\mathbf{k}) \hat{\eta}(-\mathbf{k}) \rangle$, where $\hat{\eta}$ is the Fourier transform of the stochastic forcing η in (1.1.21), and since η is real it

has conjugate symmetry $\hat{\boldsymbol{\eta}}(-\boldsymbol{k}) = \hat{\boldsymbol{\eta}}^*(\boldsymbol{k})$, so it follows that $\hat{\boldsymbol{\Pi}} = \langle |\hat{\boldsymbol{\eta}}(\boldsymbol{k})|^2 \rangle$ is real and has the symmetry $\hat{\boldsymbol{\Pi}}(\boldsymbol{k}) = \hat{\boldsymbol{\Pi}}(-\boldsymbol{k})$. Further, $\boldsymbol{\eta}$ must be normalised so that $\boldsymbol{\varepsilon}$ is the energy injection rate in (1.1.21), which in Fourier space corresponds to the constraint (1.1.52), which in this context is

$$\int_{\mathbb{R}^2} \frac{\hat{\mathbf{\Pi}}(\mathbf{k})}{2|\mathbf{k}|^2} \, \mathrm{d}\mathbf{k} = 1. \tag{4.2.2}$$

Appendix (A.1) provides an important result from SY14: the 'sheared disturbance' solution for \hat{Z} at late times $(t \gg \mu^{-1})$. This is found in the appendix as

$$\hat{\mathcal{Z}}_{\mathbf{k}}(t) = \int_0^\infty \hat{\Pi}(k, l + k\Delta(w, t)), \qquad (4.2.3)$$

where

$$\Delta(w,t) = \int_{t-w}^{t} \gamma(\tau) \, d\tau, \qquad (4.2.4)$$

is a time-history integral of the background shear. The solution (4.2.3) is a slight generalisation of the SY14 approach by introducing time-dependence for the background shear (we shall return to the time-independent case shortly, but keep the generality for now as it is needed for section 4.5). SY14 realised that from the solution (4.2.3), it is possible to calculate the Reynolds stress $\langle u'v' \rangle$ by exploiting the linearity of (3.3.6) to separate the contributions to $\langle u'v' \rangle$ from forcing at each wavevector \mathbf{k} . The total Reynolds stress is found as

$$\langle u'v'\rangle(t) = -\int_{\mathbb{R}^{2}} \frac{kl\hat{\mathcal{Z}}}{(k^{2}+l^{2})^{2}} \,\mathrm{d}\boldsymbol{k},$$

$$= -\varepsilon \int_{\mathbb{R}^{2}} \int_{0}^{\infty} \frac{kle^{-2\mu w}}{(k^{2}+l^{2})^{2}} \tilde{\Pi}(k,l+k\Delta(w,t)) \,\mathrm{d}w \,\mathrm{d}\boldsymbol{k},$$

$$= -\varepsilon \int_{\mathbb{R}^{2}} \int_{0}^{\infty} \frac{k(l-k\Delta(w,t))e^{-2\mu w}}{(k^{2}+(l-k\Delta(w,t))^{2})^{2}} \tilde{\Pi}(k,l) \,\mathrm{d}w \,\mathrm{d}\boldsymbol{k},$$

$$= \varepsilon \int_{\mathbb{R}^{2}} \left\{ \frac{\tilde{\Pi}_{\boldsymbol{k}}}{2\gamma(t)|\boldsymbol{k}|^{2}} + \frac{1}{2} \int_{0}^{\infty} \frac{\mathrm{d}}{\mathrm{d}w} \left(\frac{e^{-2\mu w}}{\gamma(t-w)} \right) \frac{\tilde{\Pi}(k,l)}{k^{2}+(l-k\Delta(w,t))^{2}} \,\mathrm{d}w \right\} \,\mathrm{d}\boldsymbol{k}.$$

$$(4.2.5)$$

where the last line is found by parts and (k, l) are the components of **k**. Now let's

assume the background flow is time-independent, i.e.

$$\Delta(t, w) = w\gamma, \tag{4.2.6}$$

so that (4.2.5) becomes

$$\langle u'v'\rangle = \varepsilon \int_{\mathbb{R}^2} \left\{ \frac{\tilde{\Pi}_{\mathbf{k}}}{2\gamma |\mathbf{k}|^2} - \frac{\mu}{\gamma} \int_0^\infty \frac{\tilde{\Pi}(k,l)e^{-2\mu w}}{k^2 + (l - k\gamma)^2} dw \right\} d\mathbf{k}. \tag{4.2.7}$$

One of the original realisations that motivated the latter half of this thesis was that expression (4.2.7) simplifies dramatically, with a full explanation detailed in Appendix A.2.1. A fundamental point is that it can be shown that the *amplitude* of the wavevector at which energy is injected is irrelevant in the SY14 analysis. Rather, for forcing at each k, the SY14 result depends only on the phase angle $\phi = \tan^{-1}(l/k)$. Physically this is due to the fact that there is no intrinsic length scale associated with a constant shear flow. Consequently it is possible to express the general result in terms of an energy input density $\rho^{\varepsilon}(\phi)$, which is defined to be

$$\rho^{\varepsilon}(\phi) = \int_0^\infty \frac{\hat{\Pi}(k\boldsymbol{e}_k(\phi))}{k^2} k \, dk, \text{ where } \boldsymbol{e}_k(\phi) = (\cos\phi, \sin\phi)^T, \quad -\frac{\pi}{2} < \phi \le \frac{\pi}{2}.$$
(4.2.8)

The symmetry $\hat{\Pi}(\pmb{k})=\hat{\Pi}(-\pmb{k})$ means that $\rho^{\varepsilon}(\phi)=\rho^{\varepsilon}(\phi+\pi)$, which explains the restriction to $-\frac{\pi}{2}<\phi\leq\frac{\pi}{2}$. The energy input constraint means that

$$\int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \, \mathrm{d}\phi = 1, \tag{4.2.9}$$

so $\rho^{\varepsilon}(\phi)$ can be thought of as a density function describing the distribution of the energy input with respect to the phase angle (c.f. a probability density). Notice that the case of isotropic forcing, i.e. forcing on an annulus in wavenumber space in the idealised limit of an infinite domain, corresponds to $\rho^{\varepsilon}(\phi) = 1/\pi$ (constant).

The general result is that the momentum flux can be expressed as

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) K(\phi, m) d\phi,$$
 (4.2.10)

where $m = 2\mu/\gamma$ is a parameter governed by the ratio of the Rayleigh friction rate μ to the shear rate γ . The kernel function $K(\phi, m)$ appearing in (4.2.10) quantifies the contribution of a wave forced at phase angle ϕ to the momentum flux. It is given by

$$K(\phi, m) = 1 + \frac{1}{m} |z|^2 \text{Im} \{e^z E_1(z)\},$$
 (4.2.11)

where,

$$z(\phi, m) = m(-\tan\phi + i), \ (-\frac{\pi}{2} < \phi \le \frac{\pi}{2}).$$
 (4.2.12)

Here $E_1(z)$ is the complex exponential integral with the branch cut on the negative real axis

$$E_1(z) := \int_z^\infty \frac{e^{-t}}{t} dt, \quad |Arg(z)| < \pi.$$
 (4.2.13)

(Note that for fixed m the branch cut of $E_1(z)$ is not crossed as ϕ varies, and that the restriction $\phi < |\pi/2|$ ensures that $\tan \phi$ is continuous.)

The result (4.2.10) is an alternative, more accessible presentation of the results in equations B6-B9 of SY14 – the SY14 result B9 involves calculating an infinite time integral of an infinite sum of Chebyshev polynomials. (Initially the result (4.2.10) was actually derived from B9 using properties of Chebyshev polynomials and various tricks in mathematical methods. This was a delightful, yet confusing surprise, until the realisation that the integral in (4.2.7) can be manipulated directly).

Expression (4.2.10) has several advantages over those of SY14. First $E_1(z)$ is a tabulated function which is implemented in most mathematical software packages (e.g. as ExpIntegralE[z,1] in Mathematica, expint(z) in matlab, etc.), meaning that (4.2.10) is relatively straightforward to evaluate. Also $E_1(z)$ has well-known series expansions for both $|z| \ll 1$ and $|z| \gg 1$, which can be exploited to understand more about the behaviour of $\langle u'v' \rangle$ for any forcing Π , allowing results obtained by SY14 and [82] for rather specific forcings to be generalised to all $\rho^{\varepsilon}(\phi)$. Finally, the form (4.2.10) is particularly amenable to the case of forcing applied at discrete wavenumbers, e.g. on the doubly-periodic domain typical for numerical simulations.

In that case

$$\hat{\Pi}(\mathbf{k}) = \sum_{\mathbf{k}_j} \rho_j^{\varepsilon} |\mathbf{k}_j|^2 \delta(\mathbf{k} - \mathbf{k}_j),$$

where the sum is over all of the discretised wavenumber vectors \mathbf{k}_j generated by the domain, and ρ_j^{ε} is the fraction of energy injected at wavenumber \mathbf{k}_j (with k > 0). By construction $\sum_{\mathbf{k}_j, k > 0} \rho_j^{\varepsilon} = 1$. The result (4.2.10) then simplifies to

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} \sum_{\mathbf{k}_j, k>0} \rho_j^{\varepsilon} K(\phi_j, m),$$
 (4.2.14)

which is a convenient form for accurate comparison with the numerical results below. (A technical point related to the discretised domain is that forcing on the zonal mean (i.e. with k=0, or at $\phi=\pm\pi/2$) self-evidently does not contribute to the wave-driven momentum flux $\langle u'v'\rangle$, and must be treated separately, which explains why the sum over wavevectors with k>0. In this case ε and ρ_j^ε must be adjusted in (4.2.14) to include only the total energy injected into the non-zonal modes, and the fraction of this total associated with each k_j , respectively, so that $\sum_{k_j,k>0} \rho_j^\varepsilon = 1$.)

Given a specific energy input density $\rho^{\varepsilon}(\phi)$, a useful alternative form of (4.2.10) is

$$\langle u'v' \rangle = \frac{\varepsilon}{2\mu} G\left(\frac{2\mu}{\gamma}\right) \text{ where } G(m) := m \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) K(\phi, m) d\phi.$$
 (4.2.15)

One reason (4.2.15) is useful is that it is natural to scale the momentum flux $\langle u'v'\rangle$ in terms of the expected equilibrium energy $\mathcal{E} = \mathcal{E}/2\mu$ in the system. Also the function G(m) is a non-dimensional function of m which is unique to each energy input density $\rho^{\mathcal{E}}(\phi)$. In general, considering the problem in terms of \mathcal{E} and μ (e.g. rather than \mathcal{E} explicitly) is the natural way to interpret the problem with regards to the nondimensionalisation put forward in Chapter 1 (where essentially the parameters are scaled to normalise the non-dimensional expected equilibrium energy). It is not a difficult task to interpret the results going forward with the nondimensional parameters from Chapter 1, but for the purpose of generality dimensional parameters are retained.

4.2.2 General Properties of the Momentum Flux Formula

In order to understand the general momentum flux formulae (4.2.10) and (4.2.15), and to extend and place in context previous results by SY14, WB19 and others, it is useful to understand the behaviour of the kernel function $K(\phi,m)$. Figure 4.1 shows $K(\phi,m)$ for several values of m, with the left panel showing the full range $\left(-\frac{\pi}{2} < \phi \le \frac{\pi}{2}\right)$ and the right panel a scaled region (around $\phi = \frac{\pi}{2}$) to illustrate some of the self-similar behaviours as $m \to 0$. Some key properties of $K(\phi,m)$, established in Appendix A.3 and presented here for the case m > 0, are:

- (i) $K(\phi, m) < 1$ for all $\phi \in (-\frac{\pi}{2}, \frac{\pi}{2})$ and all $m \in \mathbb{R} \setminus \{0\}$.
- (ii) The limiting values at the boundaries of the interval are $K(\pm \frac{\pi}{2}, m) = 0$.
- (iii) $K(\phi,m)$ has a single minimum $K_-(m)$ and in the limit $m \to 0$, the location is asymptotically close to $\phi = \pi/2 m/2$ and $K_-(m) \sim -4\pi e^{-2}/m$.
- (iv) In the limit $m \to 0$, $K(\phi,m) = 1 m(\phi + \frac{\pi}{2})\sec^2\phi + O(m^2\log m)$. Note that this formula applies to fixed ϕ , and therefore cannot be used to approximate $K(\phi,m)$ when $|\phi \pm \pi/2| \sim O(m)$, meaning that it cannot describe the minima of $K(\phi,m)$ seen in Figure 4.1.
- (v) As $m \to \infty$, $K(\phi, m) = -m^{-1} \sin 2\phi + 2m^{-2} \cos \phi \cos 3\phi + O(m^{-3})$.
- (vi) $\int_{-\pi/2}^{\pi/2} K(\phi, m) d\phi = 0.$

Much of the structure of $K(\phi,m)$ can be understood by considering the momentum fluxes generated by plane waves with different ϕ , and how those plane waves evolve in a linear shear flow, under the so-called Orr mechanism [83]. When friction is large $(m \to \infty)$, plane waves are dissipated before they evolve under the action of the shear flow, and $K(\phi,m)$ is an odd function of ϕ determined by the momentum flux of waves at the local phase angle ϕ at which they are generated. Hence $K(\phi,m)$ is negative for positive ϕ and positive for negative ϕ [see explanation on pgs. 516-517 of 10]. When friction is low $(m \to 0)$, by contrast, the waves generated at angle ϕ will be long-lived and as time evolves their phase angle will decrease monotonically,

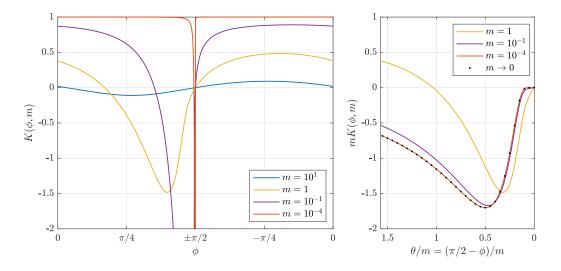


Figure 4.1: Left: the function $K(\phi,m)$ against $\phi \in [-\pi/2,\pi/2]$ for various values of m. When $m \gg 1$, K is sinusoidal (blue). As $m \to 0$, $K \to 1$ for all ϕ except an O(m) region at $\phi = \pi/2$, where an $O(m^{-1})$ minimum is obtained. Right: the same plot as the left with rescaled axes illustrates the self-similar behaviour of $K(\phi,m)$ about its minimum. Dotted black plots the $m \to 0$ solution $mK(\phi,m) = -\pi(m/\theta)^2 \exp(-m/\theta)$ found in the appendix.

as the wave is advected by the shear flow. Waves generated at an initial angle ϕ close to $\frac{\pi}{2}$, i.e. with an extreme tilt against the shear, undergo considerable transient growth due to the Orr mechanism as their phase angle approaches zero. This transient growth occurs on a timescale $1/\gamma\theta$, where $\theta=\frac{\pi}{2}-\phi$. As they evolve, these waves are also damped on a timescale μ^{-1} . The waves which generate the most extreme negative momentum fluxes are therefore those with $\theta\approx\mu/\gamma$, a value which allows for strong transient growth prior to the wave being dissipated, but with sufficient dissipation occurring that the wave is attenuated significantly by the time its phase angle becomes negative, so that the positive momentum fluxes during this stage of its life cycle are insufficient to cancel the earlier negative momentum fluxes.

Properties (i-vi) are useful as they make it relatively straightforward to prove and extend previous results regarding $\langle u'v' \rangle$ in a simple setting. The most important of these are:

1. **Bounds on the momentum flux**: Since $K(\phi, m)$ is bounded on $[-\frac{\pi}{2}, \frac{\pi}{2}]$ exact bounds on $\langle u'v' \rangle$ follow naturally, which apply to any choice of the energy

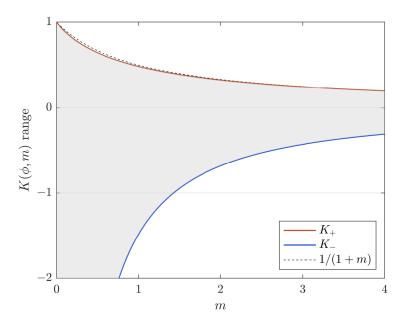


Figure 4.2: A plot illustrating the range of $K(\phi,m)$. Plotted are the ϕ -supremum and ϕ -infimum, $K_+(m)$ (red) and $K_-(m)$ (blue) respectively. The shaded grey region therefore is the range of $K(\phi,m)$. The dashed black line corresponds to the upper bound for Reynolds stress established in SY14.

input density $\rho^{\varepsilon}(\phi)$, For positive shear $(\gamma > 0)$ these are

$$\frac{\varepsilon}{\gamma}K_{-}(m) \leq \langle u'v' \rangle \leq \frac{\varepsilon}{\gamma}K_{+}(m), \text{ where } K_{+}(m) = \sup_{\phi \in [-\frac{\pi}{2}, \frac{\pi}{2}]} \{K(\phi, m)\},$$
(4.2.16)

and K_- is the corresponding infimum. Note that property (i) ensures that $K_+ < 1$. Each bound can be attained exactly by setting $\rho^{\varepsilon}(\phi) = \delta(\phi - \phi_{\pm}(m))$ where $\phi = \phi_{\pm}(m)$ denotes the locations of the supremum and infimum respectively. Interestingly, the upper bound improves only very slightly on the bound $\langle u'v' \rangle \leq \varepsilon/\gamma(1+m)$ found by SY14 from the energy power integral. The lower bound, which satisfies $K_-(m) \sim 4\pi e^{-2}m^{-1}$ when $m \to 0$, is only useful when m is order unity. Equation (4.2.11) allows K_\pm to be evaluated numerically to high accuracy, and the results are shown in Figure 4.2 along with the SY14 bound.

2. The low friction limit (WB19 result): A key question raised in the Introduc-

tion is under what circumstances is the low friction approximation (4.1.2), $\langle u'v'\rangle \approx \varepsilon/\gamma$ valid? It turns out that a simple necessary condition for (4.1.2) to be valid in the limit $m \to 0$, is that there is a wave angle cut-off in the forcing density $\rho^{\varepsilon}(\phi)$. That is to say, it is necessary that there exists a cut-off angle ϕ_m such that $\rho^{\varepsilon}(\phi) = 0$ for all ϕ satisfying $\phi_m < |\phi| < \pi/2$. For example, such a cut-off occurs naturally for simulations in a finite domain when the forcing is on an annulus in wavenumber space, because in this case k is bounded below and l is bounded above. In this case (4.1.2) can be refined by inserting the expansion in (iv) above into (4.2.10), to give (for m > 0)

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} \left(1 - m \int_{-\phi_m}^{\phi_m} \left(\phi + \frac{\pi}{2} \right) \sec^2 \phi \ \rho^{\varepsilon}(\phi) \ d\phi \right) + O(m^2 \log m). \quad (4.2.17)$$

$$\to \frac{\varepsilon}{\gamma} \quad \text{as } m \to 0.$$

The phase angle cut-off is evidently necessary in (4.2.17) in order for the leading order correction term to avoid the singularities in the integrand at $\phi = \pm \pi/2$ and consequently remain bounded.

In practice, while formally valid for $m \ll 1$, (4.2.17) will be an accurate approximation to (4.2.10) only if $|\frac{\pi}{2} - \phi_m| \gg m$, in order that the leading correction term in (4.2.17) remains O(m). While this condition will always be formally satisfied if ϕ_m is fixed and the limit $m \to 0$ is taken, long domains which allow for forcing at wavenumbers with $|k| \ll |l|$ will have $|\frac{\pi}{2} - \phi_m| \ll 1$, and will therefore require m to be very small before (4.2.17) becomes accurate. A further practical point is that, when applied to slowly varying jets, (4.2.17) cannot hold everywhere because no matter how small the friction, m will not remain small in boundary layer regions around the jet extrema, where $|U_y|/\mu \lesssim 1$.

3. The high friction / low shear limit: The behaviour in the limit $m \to \infty$, corresponding to high friction or weak shear, is always important in flows with jets close to the jet extrema. The behaviour of $\langle u'v' \rangle$ in this limit can be

investigated by inserting the expansion in (v) into (4.2.10). A cancellation of leading terms follows and the result is

$$\langle u'v' \rangle = -\frac{\varepsilon}{2\mu} \left(\int_{-\pi/2}^{\pi/2} \sin 2\phi \, \rho^{\varepsilon}(\phi) \, d\phi - \frac{2}{m} \int_{-\pi/2}^{\pi/2} \cos \phi \, \cos 3\phi \, \rho^{\varepsilon}(\phi) \, d\phi \right) + O(m^{-2}). \tag{4.2.18}$$

Reassuringly, (4.2.18) shows explicitly that $\langle u'v' \rangle$ is bounded as $m \to \infty$ for all possible $\rho^{\varepsilon}(\phi)$, consistent with the findings of SY14. The leading term can be recognised as the momentum flux arising from the high friction solution of (3.3.6), in which all terms involving the shear are neglected to give $\mathcal{Z} = \varepsilon \Pi/2\mu$. If $\rho^{\varepsilon}(\phi)$ is taken to be an even function of ϕ , as is typical in simulations, the first term vanishes and the second term is dominant. In this case $\langle u'v' \rangle$ is linear in the shear $\gamma \equiv U_{\gamma}$ near to the jet extrema.

4. **The isotropic forcing paradox**: The property (vi) ensures that if the forcing in (4.2.10) is isotropic, i.e. $\rho^{\varepsilon}(\phi) = 1/\pi$, then

$$\langle u'v'\rangle = 0$$
, for all m . (4.2.19)

This remarkable result, that an isotropic forcing leads to zero momentum flux, was first discovered by [84]¹ and was strongly emphasised by SY14. At first glance, considering the limit $m \to 0$, (4.2.19) appears to be in contradiction to the WB19 result (4.1.2). However, isotropic forcing with $\rho^{\varepsilon}(\phi) = 1/\pi$ does not satisfy the wave angle cut-off property which is required for (4.1.2) to hold. The subtlety here is the remarkable structure of $K(\phi,m)$ which allows properties (i) and (vi) to hold simultaneously as $m \to 0$. Essentially, an increasingly thin boundary layer close to $\phi = \pi/2$, illustrated in the right panel of Figure 4.1, allows the integral property (vi) $(\int_{-\pi/2}^{\pi/2} K(\phi,m) d\phi = 0)$

¹ The publication [84] builds on ideas from [85], both studying the evolution of small perturbations in plane Couette flow. The set-up in these manuscripts do not involve friction, but essential insights related to the Orr mechanism are discovered. In particular that the Reynolds stress of an (initial) isotropic perturbation is identically zero for all time. Linear damping does not effect this property.

to hold even as $K(\phi, m) \to 1$ for all $\phi \in (-\frac{\pi}{2}, \frac{\pi}{2})$. Physically, what is occurring as $m \to 0$ is that only those waves with the most extreme tilt against the shear $(k \ll l)$ can contribute a negative momentum flux, but the momentum flux for these waves becomes increasingly large in magnitude because they become very long-lived when friction is low. In any finite domain, with forcing confined to an annulus in wavenumber space, waves with such extreme tilts will not exist and (4.1.2) holds as $m \to 0$. The above reasoning explains why the many numerical experiments of β -plane turbulence [e.g. 82, 22, 16, 15] which use a discretised 'isotropic' forcing in a finite domain, result in jet flows with decidedly non-zero $\langle u'v' \rangle$. An alternative perspective on the same issue is given in [86]. There, the focus is on the length of time required for the negative momentum fluxes of the 'extreme tilt' waves to become established, which becomes increasingly long as $m \to 0$. As a result, switching the order of the limits $m \to 0$ and $t \to \infty$ in the solution of (3.3.6) results in isotropic forcing giving (4.1.2) instead of (4.2.19). Here the focus is on steady equilibria and consequently the limit $t \to \infty$ is always taken first.

Given a specific energy input density $\rho^{\varepsilon}(\phi)$, (4.2.15) shows that momentum flux is determined by the function G(m). Figure 4.3 shows G(m) for some simple examples of different wave forcings (WF1, WF2 and WF3 hereafter, see caption), together with the corresponding limiting forms obtained from (4.1.2) and (4.2.17) (valid as $m \to 0$) and (4.2.18) (valid as $m \to \infty$). The asymptotic expressions are seen to do a good job of approximating G(m) for $m < 10^{-1}$ and $m > 10^{1}$, with (4.2.17) representing a significant improvement on (4.1.2). However, both the small m and large m expressions are inaccurate in the range $10^{-1} \lesssim m \lesssim 10^{1}$. In Figure 4.3 there is a notable contrast between G(m) for forcings for which $\rho^{\varepsilon}(\phi)$ is an even function of ϕ (middle and lower panels) compared to otherwise (top panel). In the latter case, G(m) tends towards a constant value as $m \to \infty$, consistent with the fact that the forcing in this case generates a non-zero momentum flux even in the absence of a shear flow (see the leading term in 4.2.18).

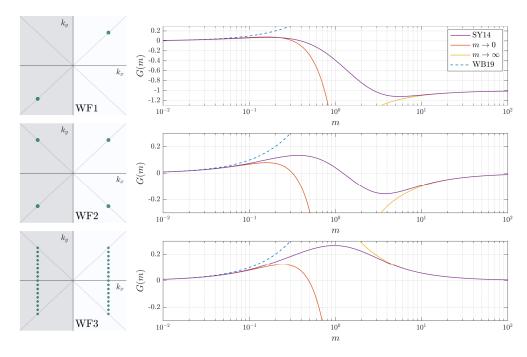


Figure 4.3: Examples of the SY14 momentum flux function G(m) defined in (4.2.15) for the wave forcings WF1, WF2 and WF3. The right panels show G(m) (purple), the limiting forms (4.2.17) (red curves) and (4.2.18) (orange curves) and the WB19 result (4.1.2, G(m) = m) (dashed blue). The left panels show the patterns of the wave forcings WF1, WF2 and WF3 in wavenumber space. WF1: $\phi_1 = \pi/4$, WF2: $\{\phi_1, \phi_2\} = \{-\pi/4, \pi/4\}$ and $\rho_j^{\varepsilon} = 1/2$, WF3: $\phi_j = \tan^{-1}(j/8)$ for j = -8, ..., 0, ..., 8 and $\rho_j^{\varepsilon} \propto \cos^2 \phi_j$.

4.2.3 The SY14 Momentum Flux Closure Applied to a Jet Flow

Up to this point the focus has been on constant shear flows with $U(y) = \gamma y$. How relevant is the SY14 expression for the momentum flux when U(y) is instead a smooth jet flow? Intuitively when the parameter F, which determines the ratio of the forcing scale to the jet scale, satisfies $F \ll 1$ one would expect that the eddies 'see' only a linear shear locally and SY14 will be accurate. A key question is how small does F have to be in practice for the SY14 formula to hold.

The reason the above question is important is that if a regime exists in which SY14 holds, the question of whether the equilibrium jet profile U(y) is ever determined by a purely zonal closure can be addressed. The SY14 zonal closure equation is obtained by inserting (4.2.15) into (3.3.5) to give

$$\mu U = -\frac{\varepsilon}{2\mu} \partial_{y} G\left(\frac{2\mu}{U_{y}}\right) - v_{2n} (-\partial_{yy})^{n} U, \qquad (4.2.20)$$

where the function G is defined in (4.2.15). In principal (4.2.20), which becomes a second-order ordinary differential equation for U when hyperdiffusion is neglected, can be solved numerically to obtain an equilibrium jet profile U(y), for any G. A special case is G(m) = m, which corresponds to the WB19 situation in which $\langle u'v' \rangle = \varepsilon/U_y$ (see 4.1.2). For this case the exact solution, neglecting hyperdiffusion, can be given in terms of the inverse of the error function [86]

$$U(y) = 2\sqrt{\mathcal{E}}\operatorname{erf}^{-1}(y/\sqrt{A^2\pi\mathcal{E}}) \text{ for } -A\sqrt{\pi\mathcal{E}} < y < A\sqrt{\pi\mathcal{E}},$$
 (4.2.21)

where $\mathcal{E} = \varepsilon/2\mu$. Here the existence of a latitude y=0 where U=0 is assumed, and $A=U_y(0)>0$ is the local (undetermined) shear there. The profile (4.2.21), which becomes singular at $y=\pm A\sqrt{\pi\mathcal{E}}$ has nevertheless been argued by WB19 to be a reasonable fit to the 'between jet' zonal wind profiles seen in simulations. Obviously, however, the singularities in (4.2.21) mean that it cannot be accurate everywhere. Replacing G(m)=m with the full SY14 expression in (4.2.20) avoids singularities, but it is an important finding of this work that the SY14 expression does not generally result in more accurate jet profiles. Much of what follows will be dedicated to explaining why this is the case.

A good starting point for investigating the SY14 closure is to perform a scattering experiment. That is, the CE2 equation (3.3.6) is solved numerically for a fixed steady jet-like wind profile U(y), and the resulting momentum flux $\langle u'v'\rangle$ is compared to the theoretical prediction (eqn. 4.2.15 with γ replaced with U_y). Note that the CE2 solution is exactly equivalent to that obtained by solving (3.2.2) with $\bar{u} = U(y)$ held fixed and obtaining $\langle u'v'\rangle$ by statistical averaging, however CE2 is clearly more efficient since statistical error is eliminated. Additionally, when U(y) is held fixed, (3.3.6) becomes a Lyapunov equation (see 3.3.12), which is much cheaper to solve numerically compared to the full time-dependent CE2 equations. High resolution solutions of (3.3.6) can therefore be obtained, allowing for a numerical investigation of the small F limit.

In a 2π -periodic domain, the scattering experiments are performed on a fixed flow $U(y) = 2\sin y$, with quantisation and zonostrophy parameters Q = 1.22 and

Z=1.94 respectively. The value of Q is consistent with a single jet in the domain, and the modest value of Z allows boundary layers² near jet extrema to be resolved, which becomes prohibitively expensive at higher Z. By varying the forcing wavenumber k_f across experiments, the effect of varying F in the range 0.04 to 0.84 is investigated. Figure 4.4 shows the resulting momentum flux $\langle u'v' \rangle$ for the wave forcing patterns WF2 and WF3, illustrated in the left panels, and described further in the Figure 4.3 caption. The middle panels show the main comparison and demonstrate that, away from the jet extrema, in both cases there is good agreement between the CE2 scattering experiments at all three values of F and the SY14 asymptotic result (4.2.15), and to a lesser extent the simpler expression (4.1.2).

The right panels show a blow-up of the situation close to the jet extrema, where $\langle u'v' \rangle$ varies rapidly in thin boundary layers, in both the CE2 calculations and in the theory (4.2.15). The theoretical expression (4.2.15) depends only on U_y and not β , and therefore is symmetric at both east and west jets. The CE2 solutions include the effect of β and therefore the convergence to (4.2.15) as $F \to 0$ is seen to be rather different for each jet. In the upper panel, in which the forcing is concentrated on just a few waves, even the sign of the momentum flux is opposite to the (4.2.15) prediction for the larger values of F, showing that convergence is slow. When the forcing is distributed over a wider range of wavenumbers, as seen in the lower panel, convergence is more uniform. In both sets of calculations convergence is significantly slower at the east jet, where the PV gradient $\beta - U_{yy}$ is large and positive, compared to the west jet where the magnitude of $\beta - U_{yy}$ is smaller.

In summary, the scattering experiments show clearly that, for sufficiently small but nevertheless physically reasonable (and numerically accessible) values of F, the SY14 expression (4.2.15) can do a good job of predicting the momentum fluxes

$$rac{2\mu}{U_{\scriptscriptstyle V}} \sim rac{\mu L_{Rh}}{\mathcal{E}^{1/2}} \sim rac{\mu^{5/4}}{eta^{1/2} ar{\epsilon}^{1/4}} = Z^{-1/5},$$

relates zonostrophy inversely to the (local) m as $Z = m^{-5}$. Increasing zonostrophy means the (local) m at any given latitude is reduced, and the regions of non-zero shear where the WB19 solution is accurate expand, developing a thin boundary layer near to jet extrema where WB19 is inaccurate and the SY14 formula regularises.

² The 'boundary layers' in question refer to the small latitudinal regions around extrema where the momentum flux expression (4.2.10) regularises. This occurs for large (local) $m = 2\mu/U_y$. The scaling argument

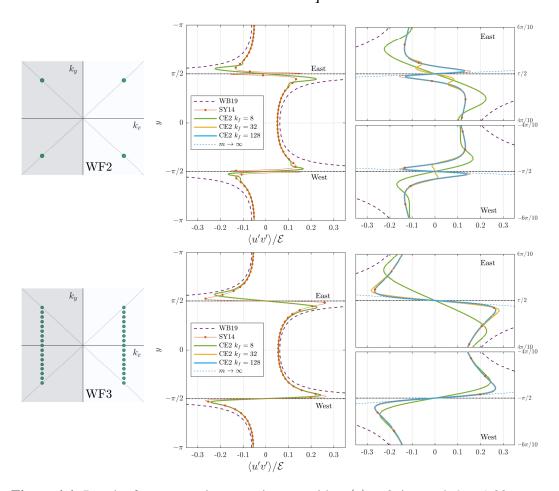


Figure 4.4: Results from scattering experiments with $U(y)=2\sin y$ and Q=1.22, Z=1.94 and three different values of F (corresponding to forcing wavenumber $k_f=(8,32,128)$ respectively, or F=(0.677,0.169,0.042) in the upper row, and F=(0.844,0.211,0.053) in the lower row). The top and bottom rows show the results for the wave forcings WF2 and WF3 respectively (see Figure 4.3 caption). The middle panels compare calculated and predicted $\langle u'v' \rangle$ across the full domain, and the right panels show a close-up of the situation near the east and west jet cores. The theoretical results (4.1.2, WB19, dashed purple), (4.2.15, SY14, red dotted line) and (4.2.18, $m \to \infty$, dashed blue) are also plotted.

everywhere, for a steady, smooth and, crucially, *stable* zonal flow U(y). Why, then, does the closure (4.2.20) not describe the equilibrium jets in β -plane turbulence calculations? This question will be addressed next.

4.3 Momentum Balance in Equilibrated Jets

In this section the aim is to explain why purely zonal closures such as (4.2.20) cannot by themselves describe the equilibrium jet profiles in β -plane turbulence. Evidence from QL simulations, CE2 calculations and fully nonlinear simulations will be used

to show that momentum fluxes due to waves arising from jet-scale instabilities also have an essential role in the momentum balance. This is not a new observation (see [87]), and it has been previously suggested that the westward jets in β -plane turbulence are marginally stable to barotropic instability [e.g 22]. The aim here is to demonstrate that the same is true at the eastward jets, and to quantify the role of these secondary instabilities in the momentum balance. A specific choice of wave forcing is used in the simulations to facilitate the analysis, which is described next.

4.3.1 Experimental Design: Separation of Momentum Flux Contributions

A standard approach in β -plane simulations of jets is to use a (near) isotropic wave forcing η , in which the forcing is applied only at wavevectors k which lie within an annulus $(k_- < |k| < k_+)$. However, for fixed Q and Z, the equilibrium jet structure in simulations is known to be largely insensitive to the details of the forcing mechanism [e.g. 15]. Here, the freedom to choose the forcing structure is exploited by instead using the forcing WF3, which is concentrated on a single zonal wavenumber $k = k_f$ (see discussion above). In all but one simulation, WF3 is augmented (slightly) by also including an extremely weak forcing applied to waves with $k \neq k_f$ to allow for the excitation of instabilities. The reason for this is twofold:

• To allow the momentum fluxes associated with waves directly forced by η and those generated by secondary instabilities to be clearly distinguished. Exploiting the properties of the Fourier transform in x, it is helpful to first decompose the momentum fluxes into contributions from each zonal wavenumber k,

$$\langle u'v'\rangle_k(y) = 2\operatorname{Re}\left(\tilde{u}_k(y)\,\tilde{v}_k^*(y)\right),\tag{4.3.1}$$

where \tilde{u}_k and \tilde{v}_k denote the coefficient of the k term in the x-Fourier transform of u' and v' respectively. This allows for the decomposition of the momentum flux $\langle u'v' \rangle = \langle u'v' \rangle_D + \langle u'v' \rangle_S$ into directly forced and secondary components

defined to be

$$\langle u'v'\rangle_D = \langle u'v'\rangle_{k_f} \tag{4.3.2}$$

$$\langle u'v'\rangle_S = \sum_{k \neq k_f} \langle u'v'\rangle_k.$$
 (4.3.3)

The above decomposition pre-supposes that no secondary instabilities will occur at $k=k_f$, which is the case in the simulations, because $F\ll 1$ and there is a clear scale separation between the jet-scale secondary instabilities and the forcing. Forcing on wavenumbers in an annulus does not allow for a clean decomposition, because the zonal wavenumbers associated with the secondary instabilities will also be forced directly. Below, it will also prove useful to further decompose $\langle u'v'\rangle_S$ to help distinguish between instabilities on the east and west jets.

• The second motivation for choice of forcing is to introduce a long wave cut-off for the excitation of zonal wave modes. It is now widely accepted that baroclinic instability is the primary mechanism driving the extratropical jets on the giant planets [e.g. 88, 89, 90] and, if η is to represent this process, it should have a long wave cut-off in the zonal direction [i.e. there should be no stochastic excitation of waves with $|k_x| < k_{\min}$ for some $k_{\min} > 0$, c.f. the Eady or Phillips model for baroclinic instability, e.g, 10]. As discussed above, the existence of a long-wave cut-off is quantitatively important because very long zonal waves can make a disproportionately large (and, if not treated carefully, an unphysical) contribution to the momentum flux in the local theory [see also 86].

With the forcing described above, a reference simulation (REF) is performed at parameter settings (Z,Q,F)=(5.05,2.91,1.01). (This value of F corresponds to forcing at $k_f=16$). For the purposes of comparison and model validation, simulation REF is repeated in the full nonlinear equations (1.1.21), the QL equations (3.2.1-3.2.2) and the CE2 equations (3.3.5-3.3.6). The resolution in the nonlinear simulation is 256^2 Fourier modes, while for QL and CE2 it is 16×256 since wavenumbers

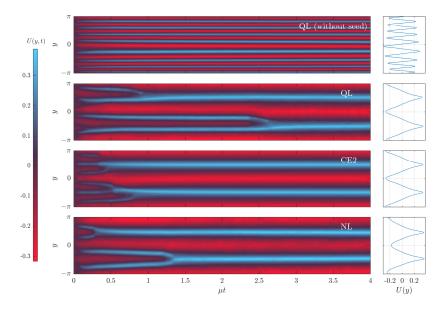


Figure 4.5: Left panels: Hovmöller plots of zonal mean flow U(y,t). (Top) QL simulation with wave forcing WF3 with the 'seed forcing' omitted. (Second) QL simulation with wave forcing WF3 including the seed forcing. (Third) CE2 simulation with WF3 including the seed forcing. (Bottom) Fully nonlinear simulation with WF3. Right panels: time average of U(y,t) over the last $0.5\mu^{-1}$ time period. In all simulations Z=5.05, Q=2.91 and F=1.01

 $|k| > k_f$ are not necessary. Hyper-diffusivity in all cases is $v_4 = 2.5 \times 10^{-8}$, which is sufficient to remove enstrophy at small scales in the nonlinear simulation, but remains insignificant in the energy balance. The results are reported next.

4.3.2 Momentum Flux Decomposition in an Equilibrated Jet Flow

Figure 4.5 shows the spin-up and equilibration of jets in the QL, CE2 and nonlinear (NL) simulations. The left panels are Hovmöller plots of the zonal mean wind U(y,t) and the right panels show a time-mean taken over the final $0.5\mu^{-1}$ time units in each simulation. The QL and CE2 simulations result in very similar jet profiles, as expected since the latter describes the statistics of the former. The jets in the NL simulation are less symmetric, with a more rounded westward jet, reflecting the fact that at $Z\approx 5$ the NL simulation only approaches the high zonostrophy ($Z\gg 1$) regime in which QL is expected to be a good approximation to (1.1.21). Nevertheless Figure 4.5 confirms the usefulness of QL and CE2 as simplified models of the full nonlinear behaviour.

The top panel of Figure 4.5 shows a QL simulation which gives a simple illustration of the importance of secondary instabilities for the equilibrated jet profiles. In this simulation the 'seed forcing' in WF3 is switched off, meaning that waves with zonal wavenumber $k \neq k_f$ cannot be excited. Despite the fact that, energetically speaking, the forcing is effectively unchanged, since the fraction of energy input into the seed forcing into waves with $k \neq k_f$ is less than 10^{-4} of the total, the outcome of the simulation is radically different. Thin jets are formed with widths far less than the Rhines scale seen in the other simulations. Further QL simulations (not shown) demonstrate that the equilibrated jets in the QL simulation (second row) are independent of the amplitude used for the seed forcing, providing further evidence that the role of the seed forcing is to excite instability.

To understand the mechanisms at play it is informative to look first at the momentum flux decomposition for the CE2 and QL simulations, shown in the top two rows of Figure 4.6, because these are somewhat cleaner than their NL counterpart and give near identical results. The left panels show the time-mean zonal wind U(y)for reference, and the remaining panels show the contributions to the momentum flux convergence - i.e. the wave induced force on the zonal flow - for different zonal wavenumbers. Results for QL are obtained using a 'jet-following' averaging procedure to compensate for any gradual evolution in jet position which would otherwise smear out the results. The jet-following procedure resets in the origin in the y-direction prior to averaging, by phase-shifting all quantities in Fourier space to the phase of the jets, as determined by the second Fourier coefficient of U(y,t). Using this method, averages are calculated over a long period $(50\mu^{-1})$ of equilibrium, to obtain good statistical convergence. CE2 adopts an alternative averaging procedure, since equilibrium solutions for CE2 undergo small oscillations [74] about a steady jet configuration. In the solution presented the oscillations are relatively small, but a temporal average is taken anyway once a steady state is reached (from $t = 49\mu^{-1}$ to $50\mu^{-1}$).

The directly forced contribution to the momentum flux convergence $-\partial_y \langle u'v' \rangle_D$ (for $k=k_f=16$) is shown in the rightmost panels. The secondary instability

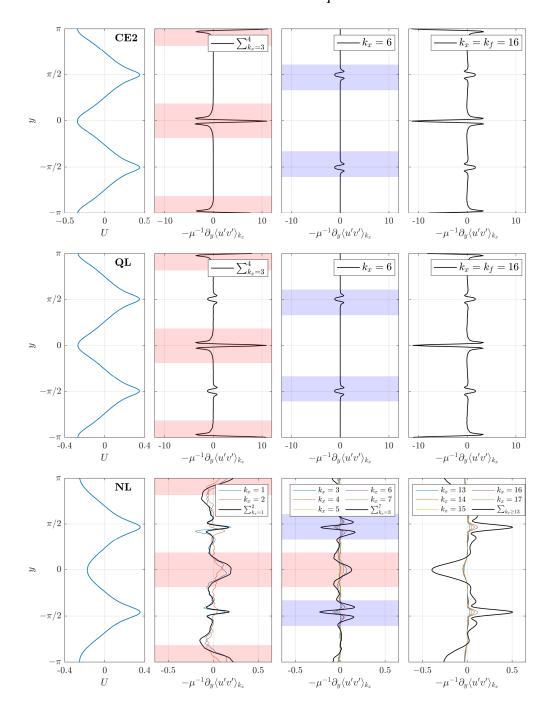


Figure 4.6: Long-time equilibrium mean quantities from the CE2 (top), QL (middle) and NL (bottom) simulations reported in Figure 4.5. The left panels show the average mean wind profile U(y). The other panels plot $-\mu^{-1}\partial \langle u'v'\rangle_k$ for significant modes k. The CE2 and QL results are similar, clearly identifying distinct instabilities at the westward (shaded red) and eastward (shaded blue) jets which counteract the jet-sharpening contribution from the forcing wave $k=k_f$. Analysis of the NL simulation reveals wave ranges performing similar roles as the QL and CE2 simulations: e.g. k=1,2 damps westward jet growth, k=3-7 all have a similar structure counteracting the eastward jet sharpening, and the waves k>13 perform the jet sharpening role of the forcing wave k_f .

contribution $-\partial_y \langle u'v' \rangle_S$ is found to be entirely dominated by zonal wavenumbers k=3,4 and 6 and these contributions are plotted in the middle panels. A striking feature of Figure 4.6 is that at the jet cores there is a cancellation between two large terms: the direct wave induced force which acts to accelerate both the westward and eastward jets, and the contributions from the secondary instabilities which are decelerating. At the westward jets in particular, the direct and secondary momentum flux convergences are about an order of magnitude larger than the frictional force in the momentum balance in equation (3.3.5). The main contributions at the eastward and westward jets have somewhat different scales (k=6 and k=3,4 respectively), which will be shown below is characteristic of the scale of barotropic instability on each jet. The secondary instabilities are strictly meridionally non-local, however their influence is not felt across the full domain, and there are wide regions in the jet flanks where $\langle u'v' \rangle_S \ll \langle u'v' \rangle_D$. Figure 4.7 shows that, within these wide regions, SY14 formula (4.2.10) is an excellent prediction for the Reynolds stress.

The momentum flux decomposition for the NL simulation is shown in the bottom row of Figure 4.6. There are striking similarities between NL and QL/CE2 in the *patterns* of the momentum flux convergences at each jet, but also significant differences in both the magnitudes and latitudinal scales, as well as the breakdown by zonal wavenumber. The differences can be accounted for by the fact that *Z* is finite in the NL simulation, and agreement between NL and QL is expected to improve at higher *Z*. Finite *Z* affects the NL flow statistics in two important ways:

- First, there is significant wave-wave interaction in the nonlinear simulation, which leads to wave energy being scattered in wavenumber space. This affects the direct wave induced force $-\partial_y \langle u'v'\rangle_D$ by spreading its contribution over a range of wavenumbers centred on $k=k_f$, and similarly the spectrum of $-\partial_y \langle u'v'\rangle_S$ is broadened to a larger range of k. This is clear from the bottom panel of figure 4.6.
- Second, the average variance of the mean wind profile about its temporal
 mean is roughly five times greater in the NL case compared to QL (both cases
 are calculated using the 'jet following' procedure). The increased variance is

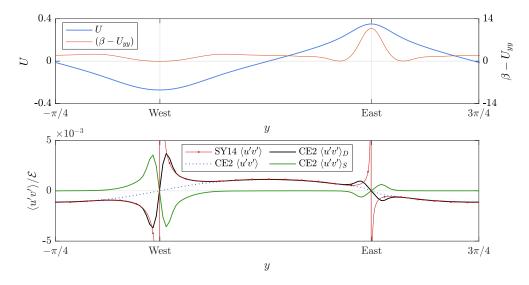


Figure 4.7: A further analysis of the CE2 simulation in Figure 4.6. The half domain $[-\pi/4, 3\pi/4]$ is reported to view a single eastward and westward jet. The top panel shows the mean zonal wind profile (blue) and the potential vorticity gradient (orange). The bottom panel shows momentum flux quantities scaled by \mathcal{E} . These are the CE2 results for $\langle u'v' \rangle$ (dotted blue), $\langle u'v' \rangle_D$ (black) and $\langle u'v' \rangle_S$ (green) and the SY14 $\langle u'v' \rangle$ (red dotted line). Sufficiently distant from the points where $\beta - U_{yy} \approx 0$ it is observed that $\langle u'v' \rangle_D \approx 0$ and the SY14 solution agrees with the CE2.

exhibited as (a) less steady relative jet positions, which smears the statistics to give weaker and broader patterns for the momentum flux convergence, and (b) increased variance about some current jet state, which is conducive to a larger likelihood that the RK condition for instability is met at any given latitude, broadening the range of potential instabilities.

The above effects account for the difference between the CE2 /QL and NL momentum flux decompositions in Figure 4.6, since the latter can be viewed as a 'smeared out' version of the former, but is otherwise qualitatively similar. In summary, momentum fluxes from the directly forced waves act to accelerate both the eastward and westward jets, and opposing these are momentum fluxes from relatively long (k = 1 to 7) jet-scale waves which appear to derive from secondary instabilities. To investigate the origin of these long waves more thoroughly, a linear instability analysis of the time-mean flow is conducted next.

4.3.3 Linear Stability Analysis of the Equilibrium Jet Flow

To investigate the origin of the emergent long waves that drive the secondary momentum fluxes $\langle u'v'\rangle_S$, a linear stability analysis of the time-mean flow in the CE2 simulation is presented next. The CE2 simulation is chosen for analysis because the time-variability of U(y,t) is significantly lower than for the stochastic simulations, meaning that its time average is an excellent statistical representation of the actual state at any fixed time. One of the pitfalls of a single numerical linear stability calculation of a time-mean state of a system near marginal stability is that the results can return large number of wave modes with near-zero growth rates, making the results hard to interpret. The goal of the linear stability analysis, therefore, is to investigate whether there are states *nearby* to the time-mean state (in a sense to be described) in which there are specific wave modes with *significant* growth rates. The idea is that, as U(y,t) evolves it will spend a proportion of its time in these more unstable states, exciting the strongly unstable wave modes, which will then persist because they are relatively weakly damped when the system is outside the unstable regime.

The linear stability problem is formulated by seeking a solution to (3.2.2) of the form

$$\psi'(x, y, t) = \text{Re} \{ \Psi(y) \exp (ik(x - ct)) \},$$
 (4.3.4)

resulting in the generalised eigenvalue problem

$$\mathcal{L}\Psi = c\mathcal{M}\Psi,\tag{4.3.5}$$

for linear operators

$$\mathcal{L} = \left(ikU + \mu - (-1)^{n+1}v_{2n}\left(\partial_{yy} - k^2\right)^n\right)\left(\partial_{yy} - k^2\right) + ik\left(\tilde{\beta} - U_{yy}\right),$$

$$\mathcal{M} = ik\left(\partial_{yy} - k^2\right).$$

Here $\tilde{\beta} = \beta + \delta \beta$, where $\delta \beta$ is a perturbation to the value of β used in the CE2 simulations, which has been introduced as a device to investigate the stability of

'nearby' states to the CE2 time-mean flow. The assumption here is that the stability properties of nearby states generated by varying $\delta\beta$ as a control parameter are representative of those of the nearby states generated by fluctuations in U(y,t). It's believed this is reasonable as growth rates are largely determined by the width and magnitude of the PV gradient reversal region in which $\beta - U_{yy}$ changes sign, as suggested by the Rayleigh-Kuo necessary criterion for instability [21], which states that a sign change must be present for instability in the inviscid system. Note that the reference time mean profile, with $\delta\beta = 0$, just satisfies the Rayleigh-Kuo criterion, with small regions of opposite sign PV gradient located at the westward jet extrema and in flanks of the eastward jet.

The generalised eigenvalue problem (4.3.5) is discretised on a grid of 1024 points by replacing \mathcal{L} and \mathcal{M} with the matrices obtained when standard centred-difference approximations replace y-derivatives. The calculated growth rate kc_i of the fastest growing mode, obtained by taking the imaginary part of the computed eigenvalues c of (4.3.5), is plotted in Figure 4.8 as a function of $(k, \delta\beta)$. As $\delta\beta$ is reduced and the system becomes more unstable, growing waves are seen to emerge at around $k \approx 3.7$ and $k \approx 6.1$, which correspond to instabilities at the westward and eastward jets respectively. Notably, there is a significant asymmetry between strong positive growth rates at negative $\delta\beta$ and weak decay rates at positive $\delta\beta$, supporting the idea that the system need spend only a relatively small fraction of time in the unstable regime to support the emergence of these waves.

The latitudinal structure of the momentum flux convergence associated with each unstable wave can be calculated from the corresponding eigenvector Ψ . The right panels show the calculations for a typical unstable wave on the westward jet (green dot) and eastward jet (blue dot) respectively. The pattern of the latitudinal structure in each case is seen to be close to those calculated in the equilibrium QL / CE2 and NL simulations, shown in Figure 4.7. This correspondence in the momentum flux structures at each jet, together with the close matches in the respective emergent zonal wavenumbers, (k = 3, 4 for the westward jet and k = 6 for the eastward jet), provides conclusive evidence that each jet is independently marginally

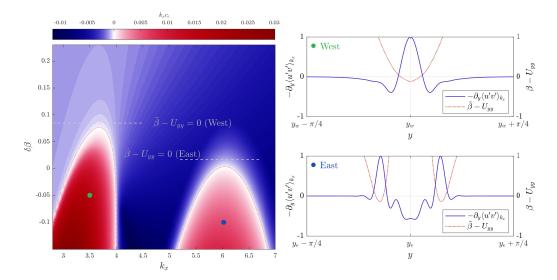


Figure 4.8: Linear stability analysis for the CE2 profile in Figure 4.6. The contour plot shows the maximum eigenmode growth rate kc_i as $\delta\beta$ and k are varied. Here, $\delta\beta = \tilde{\beta} - \beta$ measures the deviation of the Coriolis parameter $\tilde{\beta}$ used for the stability analysis compared to the actual β value used in the CE2 simulation. Two distinct unstable regions are visible at k = 3,4 and k = 6, corresponding to instabilities at the westward and eastward jets as labelled. The unstable wavenumbers agree with those identified in Figure 4.6. The right-hand side line plots show the normalised momentum flux divergence (solid blue) of the most unstable mode at the points indicated by the green and blue dots. For reference the PV gradient $\tilde{\beta} - U_{yy}$ is also shown (dotted red).

stable to barotropic instability, and that these barotropic instabilities are the source of the secondary instabilities which contribute to $\langle u'v'\rangle_S$ in the equilibrium simulations reported above.

The marginally stable nature of the barotropic jets may explain the small oscillatory behaviour of CE2 solutions (observed in the author's private calculations, and in the literature [74]). It is presumably the case that near to CE2 equilibria, cycles exist where the mean flow is 'over-sharpened' by $\langle u'v'\rangle_D$, meaning instabilities develop and $\langle u'v'\rangle_S$ 'over-damps' the mean flow, and so on.

The results and discussion presented over the course of this Chapter so far brings an end to the analysis of $\langle u'v' \rangle$ from the point of view of an equilibrium local closure theory. The results are utilised in Chapter 5, where practical applications using the SY14 formula (4.2.10) as a parameterisation small-scale forcing are considered. The remaining sections of this Chapter present additional results from the local closure theory.

4.4 The Reynolds Stress Tensor

4.4.1 Reynolds Stress Tensor

The Reynolds stress $\langle u'v' \rangle$ appears in the mean flow equation (3.3.5), so has been the natural mean eddy quadratic quantity to study. However, $\langle u'v' \rangle$ is just one component of the more general Reynolds stress *tensor*, which is

$$\mathbf{R} = \begin{bmatrix} \langle u'u' \rangle & \langle u'v' \rangle \\ \langle v'u' \rangle & \langle v'v' \rangle \end{bmatrix}. \tag{4.4.1}$$

In this section it's shown that the local theory can be used to calculate all the components of \mathbf{R} . One motivation for this is to notice that over a fixed background flow, the statistics of ζ are Gaussian. At any given point \mathbf{x} in the domain the velocity distribution is

$$\begin{pmatrix} u(\mathbf{x}) \\ v(\mathbf{x}) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} U(y) \\ 0 \end{pmatrix}, \mathbf{R}(\mathbf{x}) \right)$$
(4.4.2)

In other words, \mathbf{R} encodes a complete picture of the single-point flow statistics over the background flow. The extent to which this holds for a dynamic background flow as in the full CE2 system is an interesting avenue of exploration [58]. From here, the properties of the bivariate normal distribution can be used to calculate the probability distribution function of $u'(\mathbf{x})v'(\mathbf{x})$. Obtaining a closed form for the distribution of the product of two correlated normal random variables is a problem dating back to 1936, with the closed form solution only being found recently [91].

In principle it is possible to calculate two-point flow statistics, such as the distribution of $u'(\mathbf{x},t)v'(\mathbf{x}',t)$ with the local closure theory. To do this, the SY14 formulation of Reynolds stress in (4.2.5) would contain a physical space exponential term $e^{\mathrm{i}l(y'-y)}$ in the integrand to deal with the meridional separation (this term drops out in the single-point correlation calculations presented since y'-y=0). Any zonal separation between two points gives a correlation which can be recovered from the knowledge of the zonal wave mode decomposition, i.e. the components $\langle u'v'\rangle_k$, from the SY14 calculation as is.

Explicit calculations for the two-point correlations would provide knowledge of

the spatial structure and distribution of eddy fluctuations. This would be the natural direction the author would take this work next for various reasons. For instance it is known that CE2 tends to reach steady states [69] (although transitions in some scenarios have been found [76]). It would be interesting to see if the higher order moments of \mathcal{Z} could provide a natural method to 're-introduced' noise into the CE2 system that recovers the possibility of transitions.

The probability density function of $\overline{u'v'}$ may even provide access to transitional probabilities between turbulent attractors, like the rare events studied in Chapter 2, through a large deviations principle [58]. A theoretical approach to rare events of this sort does not rely on producing any individual realisation of the rare event, so it will almost certainly be a more direct route than any rare event algorithm. This is extremely exciting, but presumably goes beyond a local theory as the role of secondary instabilities in jet transitions is unknown.

4.4.2 Formulae and Eddy Kinetic Energy Interpretation

Formulae for $\langle u'u' \rangle$ and $\langle v'v' \rangle$ are derived in Appendix A.2.2 and read

$$\langle u'u'\rangle = \frac{\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon} \frac{|z|^2}{m} \left(-\frac{1}{m} \operatorname{Im} \left\{ e^z E_1(z) \right\} - \operatorname{Re} \left\{ e^z E_1(z) + \frac{1}{z^*} \right\} \right) d\phi, \quad (4.4.3)$$

$$\langle v'v'\rangle = \frac{\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon} \frac{|z|^2}{m} \left(-\frac{1}{m} \operatorname{Im} \left\{ e^z E_1(z) \right\} + \operatorname{Re} \left\{ e^z E_1(z) + \frac{1}{z^*} \right\} \right) d\phi. \quad (4.4.4)$$

The interrelationship between the components of \mathbf{R} is highlighted with the solutions written in the above form, e.g. as

$$\mu\left(\langle u'u'\rangle + \langle v'v'\rangle\right) = \varepsilon \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \frac{|z|^{2}}{m} \operatorname{Im}\left\{e^{z}E_{1}(z)\right\} d\phi,$$

$$= \varepsilon - \gamma \left(\frac{\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \left(1 - \frac{|z|^{2}}{m} \operatorname{Im}\left\{e^{z}E_{1}(z)\right\}\right)\right) d\phi,$$

$$= \varepsilon - \gamma \langle u'v'\rangle. \tag{4.4.5}$$

This result can be derived directly by multiplying the vorticity equation by ψ' and taking the average (e.g. SY14 [80] eqn 26). Note that equation (4.4.5) rearranges for

the nondimensional energy equation

$$\mathcal{E} = \frac{\langle u'v'\rangle}{m} + EKE, \tag{4.4.6}$$

where EKE is eddy kinetic energy. It is clear from the local theory, then, that

$$EKE = \mathcal{E} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \left(1 - K(\phi, m)\right), \tag{4.4.7}$$

with $K(\phi, m)$ the kernel function defined by (4.2.11) and $\rho^{\varepsilon}(\phi)$ the energy injection rate distribution. The same analysis for $\langle u'v' \rangle$ can be neatly converted to an equivalent expression or statement about EKE. For example, the isotropic forcing paradox, where $\langle u'v' \rangle = 0$ is equivalent to *all* the expected equilibrium energy being made up of eddy kinetic energy. The reason for relationship (4.4.6) can also be interpreted from the equilibrium zonal energy equation (see [23], eqn 3.17),

$$\partial_{y} \left(U \langle u'v' \rangle \right) = \gamma \langle u'v' \rangle - 2\mu \frac{U^{2}}{2}. \tag{4.4.8}$$

Since the first term is a divergence, $\gamma\langle u'v'\rangle$ can be interpreted as the energy injection term into the zonal mean flow. The local theory implies that the second and third terms balance so that $\langle u'v'\rangle/m$ is the local zonal energy, i.e. recovering (4.4.6). It is recognised in [23] that letting $\gamma\langle u'v'\rangle$ equal ε (interpreted as all the small scale eddy energy injection being transferred locally to the mean flow before being dissipated) recovers the WB19 inverse shear result. An alternative view of the work in the early sections of this Chapter, therefore, is that it has quantified the small reservoir of eddy kinetic energy which is not transferred locally to the mean flow, but is in fact dissipated.

4.4.3 Scattering Experiment Numerical Results

The expressions (4.4.3) and (4.4.4) above for $\langle u'u' \rangle$ and $\langle v'v' \rangle$ are confirmed in the limit of small scale forcing by repeating the scattering experiment presented in section 4.2.3 (Figure 4.4). To recall, a steady background flow profile $U(y) = 2\sin(y)$ is specified and the second cumulant equation is solved for various forcing structures.

Since the solutions are steady, the second cumulant equation is a Lyapunov equation which can be solved at high resolution, gaining access to small F.

Figures 4.9 and 4.10 are analogous to Figure 4.4, but plot $\langle u'u' \rangle$ and $\langle v'v' \rangle$ respectively (rather than $\langle u'v' \rangle$). The conclusions are much the same as for $\langle u'v' \rangle$; all \boldsymbol{R} components:

- (i) are accurately predicted by the local theory away from jet cores for all *F* in the test range,
- (ii) converge at the jet cores as $F \rightarrow 0$, and
- (iii) converge faster at the eastward jet core than the westward jet core.

Figure 4.11 shows the relationship between all \mathbf{R} components (in this case the forcing is WF3). Also plotted is EKE, showing the tendency for the small amount of eddy energy in the model to concentrate at the jet extrema.

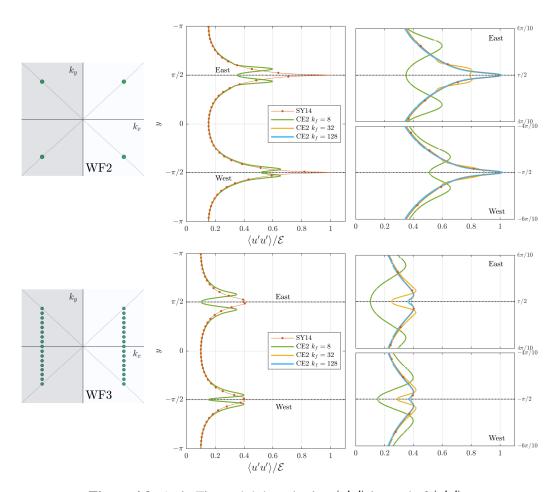


Figure 4.9: As in Figure 4.4, but plotting $\langle u'u' \rangle$ instead of $\langle u'v' \rangle$

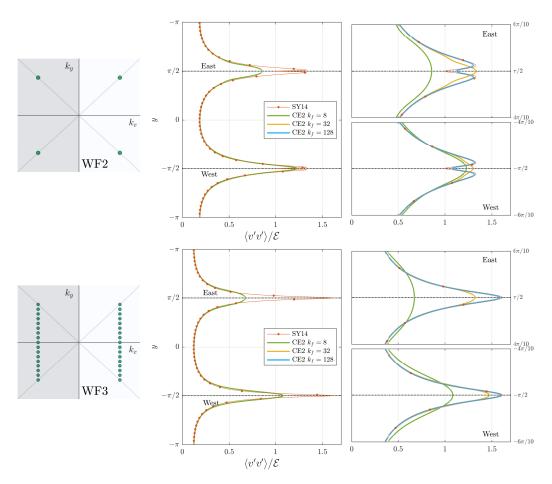


Figure 4.10: As in Figure 4.4, but plotting $\langle v'v' \rangle$ instead of $\langle u'v' \rangle$

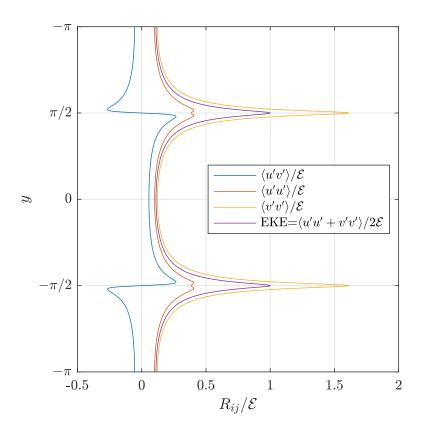


Figure 4.11: All Reynolds stress components from the SY14 local closure theory. The background flow is $U(y) = 2 \sin y$ and the forcing structure is WF3.

4.5 Time-Variability of the Local Closure

The previous sections of this Chapter have convincingly shown that small scale forcing over an equilibrium mean flow profile can be parameterised accurately using a local theory. It is not obvious, however, why taking the background shear flow as steady is valid.

4.5.1 Why Time-Variability Matters

Recall from the local theory that the time-dependent sheared-disturbance solution for $\mathcal Z$ is

$$\hat{\mathcal{Z}}_{\mathbf{k}}(t) = \int_{-\infty}^{t} e^{-2\mu w} \tilde{\Pi}(k, l + k\Delta(w, t)) \, \mathrm{d}w. \tag{4.5.1}$$

It is quite plain to see that this solution, which is derived from the CE2 framework, will generally vary on the $O(1/\mu)$ frictional timescale, owing to the exponential term. Since $\langle u'v' \rangle$ is encoded in \mathcal{Z} , ignoring time-history in a dynamical zonal mean equation could be disastrous. For instance, from (3.3.5), a zonal mean model for a jet profile could be proposed as

$$\frac{\partial U}{\partial t} = F(U) - \mu U - (-1)v_4 \partial_y^4 U, \tag{4.5.2}$$

where the functional F parameterises the Reynolds stress term. By the above analysis, an ideal F will vary on the frictional time scale. An essential observation is that the background flow is dynamic on the frictional time-scale also (from the balance of $\partial_t U$ and μU), so the time-variability of F is generally important in (4.5.2). This is a point which has been entirely missed in the 'kinetic theory' of jet dynamics put forward in [17, 70], where, effectively the same equation as (4.5.2) appears in [17] (their eqn. 88), with no recognition of time-variability. In fairness, an equation such as (4.5.2) is justifiable as a practical tool to find equilibrium states, as we will see in Chapter 5, where the background flow has limited dynamics. However, at the core of the proof in [17] that the QL model is an asymptotic approximation to the fully nonlinear model is the assumption (p. 586):

"...the fast non-zonal degrees of freedom instantaneously relax to their

stationary Gaussian distribution."

The 'stationary Gaussian distribution' here refers to the eddy distribution over a steady background flow (i.e. the solution of the steady Lyapunov equation 3.3.12). Clearly it cannot be expected that this statement is correct in light of the sheared disturbance solution (4.5.1). Whilst some hand-waving can be made for the practicality of such a statement to aid in approximating equilibria, it does not meet the rigour required for a formal asymptotic proof of convergence between QL and the nonlinear models.

Whilst on the topic of QL and nonlinear convergence, this thesis has convincingly shown the existence of large scale non-zonal structures as crucial for jet development and maintenance, and it is non-obvious where developing instabilities fit within the time-scale separation analysis. It may well be that the time scale associated with the non-zonal structures themselves can be considered on the slow time-scale (in which case the non-zonal structures could be included in the mean component, c.f. [74]), but the exponential nature of developing instabilities over a dynamic profile suggests the overall picture is more complicated. Even if it is the case that the non-zonal structures formally fit into a time-scale separation analysis, why are their nonlinear eddy-eddy interactions not significant? If, as results from [15] suggest, there is no decline in the O(1) gradients of the emergent non-zonal waves as zonostrophy is increased, then it stands to reason that nonlinear processes will remain indispensable to the overall jet dynamics.

4.5.2 Time-variable Momentum Flux Formulae

The general time-dependent solution for Reynolds stress from the local closure theory is (from the penultimate line in 4.2.5)

$$\langle u'v'\rangle(t) = -\varepsilon \int_{\mathbb{R}^2} \int_0^\infty \frac{k(l-k\Delta(w,t))e^{-2\mu w}}{(k^2 + (l-k\Delta(w,t))^2)^2} \tilde{\Pi}(k,l) \, \mathrm{d}w \mathrm{d}\boldsymbol{k}. \tag{4.5.3}$$

Following many of the same details as in Appendix A.2.1 for the steady case, it's a simple exercise to convert the above expression into polar coordinates for

$$\langle u'v'\rangle(t) = -2\varepsilon \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \sec^2 \phi \int_0^{\infty} \frac{(\tan \phi - \Delta(w,t))e^{-2\mu w}}{(1 + (\tan \phi - \Delta(w,t))^2)^2} dw d\phi. \quad (4.5.4)$$

This formula is relatively straightforward to calculate numerically for some given forcing structure and flow history U(y,t). The numerical calculation of (4.5.4) is simplified further by noting that the time-integral can be computed recursively at each time-step; if $\langle u'v'\rangle(t_0)$ is known then the calculation at the next time step, $\langle u'v'\rangle(t_0+\delta t)$, can utilise the result that $\langle u'v'\rangle(t_0)\exp(-\delta t)$ provides the relevant integral for $w \le t_0$. With this trick, formula (4.5.4) is almost no more computationally expensive to calculate on the fly than the instantaneous momentum flux formula (4.2.10). The diagonal Reynolds stress tensor components are easily found under the similar considerations. Starting with

$$\langle u'u'\rangle(t) = \varepsilon \int_{\mathbb{R}^2} \frac{l^2 \hat{\mathcal{Z}}_{\boldsymbol{k}}(t)}{(k^2 + l^2)^2} d\boldsymbol{k},$$
 (4.5.5)

$$\langle v'v'\rangle(t) = \varepsilon \int_{\mathbb{R}^2} \frac{k^2 \hat{\mathcal{Z}}_{\boldsymbol{k}}(t)}{(k^2 + l^2)^2} d\boldsymbol{k},$$
 (4.5.6)

and plugging in the sheared disturbance solution and converting to polars finds

$$\langle u'u'\rangle(t) = 2\varepsilon \int_{-\pi/2}^{\pi/2} \int_0^\infty \rho^{\varepsilon}(\phi) \sec^2 \phi \frac{(\tan \phi - \Delta(t, w))^2 e^{-2\mu w}}{(1 + (\tan \phi - \Delta(t, w))^2)^2} dw d\phi, \qquad (4.5.7)$$

$$\langle v'v'\rangle(t) = 2\varepsilon \int_{-\pi/2}^{\pi/2} \int_0^\infty \rho^{\varepsilon}(\phi) \sec^2 \phi \frac{e^{-2\mu w}}{(1 + (\tan \phi - \Delta(t, w))^2)^2} dw d\phi. \tag{4.5.8}$$

4.5.3 Scattering Experiments Over Time-Variable Shear

To investigate the quantitative extent to which time-variability is important in the local closure theory momentum flux predictions, similar experiments to that of the scattering experiments in section 4.2.3 are performed. This time, however, a time-varying background flow profile is used and a comparison is made between: the time-dependent momentum flux formula (4.5.4), the momentum flux from direct calculation of the second cumulant equation (3.3.6), and the 'instantaneous' SY14

expression (4.2.10).

It happens that the results and conclusions for $\langle u'u' \rangle$ and $\langle v'v' \rangle$ are much the same as those drawn for $\langle u'v' \rangle$, so only the results for the most relevant component $\langle u'v' \rangle$ are reported. Section 4.4 provides enough evidence that the local theory is equally valid for all components of **R**, without labouring the point.

4.5.3.1 First Experiment

In the first experiment, the time variable profile is taken as $U(y,t)=2\sin(\mu t)\sin(y)$. Figure 4.12 reports results significantly after spin up, once the initial conditions in CE2 are irrelevant. The Hovmöller plots show that the CE2 result agrees far better with the time-dependent formula than the the time-independent formula. In particular, the time-dependent solution is accurate near to times where the profile amplitude, $\sin(\mu t)$, is small or passes through zero. At these times, the time-independent solution is clearly inaccurate. In particular, the effects of the profile amplitude passing through zero have a markedly delayed impact on the momentum flux from CE2 which is not 'seen' by the time-independent solution. Interestingly, all three solutions are in good agreement (away from cusps) at times when the profile amplitude becomes relatively large. Any comparison is useless near to the profile cusps $y=\pm\pi/2$, for all the reasons previously discussed relating to finite F (these regions have also been intentionally made large, by choosing Z small, to exaggerate deviations from the WB19 solution - see discussion surrounding Figure 4.4).

A different perspective from the same experiment is provided in Figure 4.12. Here, attention is restricted is to the latitude y = 0 where at any given time the background shear attains a local maximum/minumum. This filters out the distractions of the jet extrema and allows closer inspection of the quantities provided in the Hovmöller plots. First of all, it is immediate that the time-independent SY14 formula (yellow) completely misrepresents the situation whenever the local background shear (dotted red) is small. More surprising is that although the time-dependent SY14 solution offers an improvement, there is still a discrepancy compared to the CE2 result. It is believed that this is due to the hyper-viscous effects which are not modelled in the SY14 solution. By design, hyper-diffusion acts disproportionately on

larger wavenumbers, meaning as eddies are sheared the effects become more or less pronounced. As a result the effect of hyper-diffusion modulates with the shear, giving an explanation why the time-dependent SY14 momentum flux plotted is similar to the CE2 plot, but slightly distorted and time-lagged. In the case of constant shear, the degree to which hyper-diffusion is noticeable is limited, since its overall effect is not being changed. The small role of hyper-diffusion in equilibrium states is partly hidden by calculating \mathcal{E} empirically for the CE2 experiments. It would be interesting to attempt an inclusion of hyper-diffusion in the SY14 framework (i.e. starting from modifying the sheared-disturbance solution 4.2.3), but from a theoretical point of view the interest is the limit of small F and negligible hyper-diffusion. As it is known that CE2 is definitely accurate when a large enough resolution can be obtained to accommodate small F and small V_4 for the case of a steady background flow (Figure 4.4), it is expected that the time-dependent local theory would agree with the CE2 results if the computational power was invested to replicate the experiment at an increased resolution that can accommodate an accordingly small hyper-diffusion and forcing wavelength.

4.5.3.2 Second Experiment

The observation that the time-independent solution is reasonably accurate when the background shear profile has a relatively large amplitude, motivates studying the profile $U(y,t)=(2+\sin(\mu t))\sin(y)$. At a given latitude, this profile has a shear that undergoes significant variation on the frictional time-scale, but does not change sign. This situation is arguably more typical for a jet system modelling, say, seasonal amplification of an existing jet structure.

Hovmöller results from this experiment, which are provided in Figure 4.14, offer a stark contrast to the previous example. The momentum flux predictions from the time-dependent formula, time-independent formula, and the CE2 solutions all agree relatively well. In fact, the momentum flux comparison at latitude y=0, reported in Figure 4.15, shows the difference between CE2 and the SY14 theory is mostly due to issues unrelated to time-dependence (i.e. hyper-diffusion), since the time-dependent and time-independent formulae agree better than they do with the

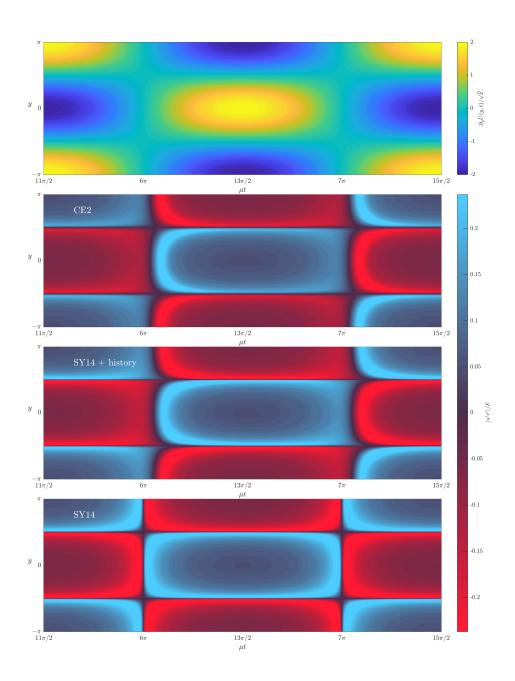


Figure 4.12: Hovmöller plots for: (top panel) specified mean shear profile $\partial_y U(y,t) = \sin(\mu t)\cos(yt)$, and (bottom panels) momentum flux $\langle u'v'\rangle$ from CE2 simulation, time-dependent SY14 formula (4.5.4), and the instantaneous SY14 formula (4.2.10).

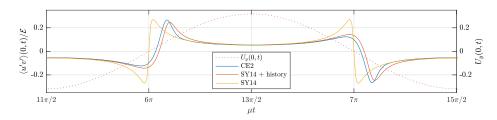


Figure 4.13: The results from Figure 4.12 at the latutude y = 0. The plots are local shear $U_y(0,t)$ (dotted red), and momentum flux $\langle u'v' \rangle$ from CE2 (blue), time-dependent SY14 (red) and instantaneous SY14 (yellow).

CE2 solution. The small difference is that the time-dependent SY14 solution has a small time-lag compared to the instantaneous SY14 solution.

4.5.3.3 Further Discussion

Why is the agreement in momentum flux results improved so dramatically in Figures 4.14 and 4.15 compared to the first experiment? The following provides an explanation which argues taking the agreement between the time dependent/independent results with, at the least, a pinch of salt.

Supposing, as is the case in the second experiment, that the profile shear does not change sign and remains $O(\gamma_0)$, for γ_0 a characteristic shear value. The time history integral $\Delta(w,t)$ will grow monotonically and indefinitely the further back it's taken, i.e.

$$\Delta(w,t) = \int_{t-w}^{t} \gamma(\tau) d\tau = O(\gamma_0 w). \tag{4.5.9}$$

The vector $(k, l+k\Delta(w,t))^T$ features in the sheared disturbance solution (4.2.3), but it is noticeable given (4.5.9) that at some time w_0 , the vector $(k, l+k\Delta(w,t))^T$ will escape the range of forcing vectors for all relevant (k,l), i.e., $\tilde{\Pi}(k,l+k\Delta(w,t))=0$ for all $w \geq w_0$. In fact, $w_0 = O(l_{\text{max}}/\gamma_0 k_{\text{min}})$ where k_{min} is the smallest zonal wave being forced and l_{max} is the largest physically relevant meridional wavenumber. In the simulations, for instance, l_{max} can be no greater than the wavelength at which hyper-diffusion becomes important. Now solution (4.5.1) implies that second order statistics vary significantly on a time scale of order min $\{O(l_{\text{max}}/\gamma_0 k_{\text{min}}), O(1/\mu)\}$. So provided $l_{\text{max}}/k_{\text{min}} \ll \gamma_0/\mu$, the integrand in equation (4.5.1) contributes only for $w < w_0 \ll 1/\mu$, and as the profile varies on an $O(1/\mu)$ time scale we can write

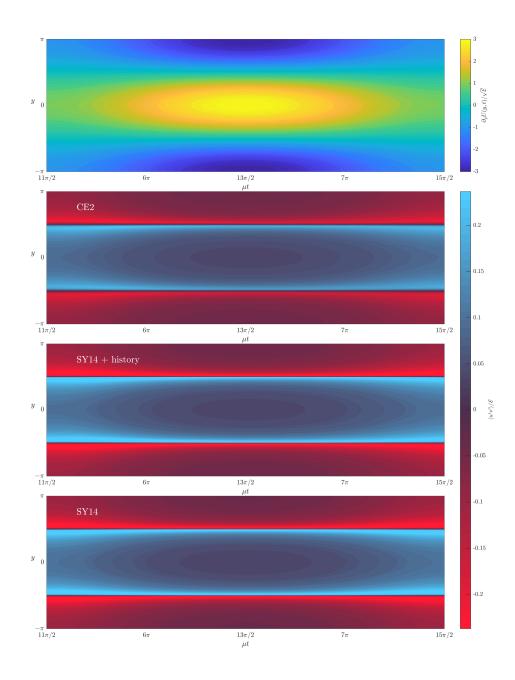


Figure 4.14: As in Figure 4.12, but with the mean flow profile $U(y,t) = (2+\sin(\mu t))\sin(yt)$. Hovmöller plots for: (top panel) specified mean shear profile $\partial U(y,t)$, and (bottom panels) momentum flux $\langle u'v' \rangle$ from CE2 simulation, time-dependent SY14 formula (4.5.4), and the instantaneous SY14 formula (4.2.10)

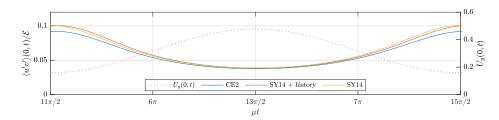


Figure 4.15: The results from Figure 4.14 at the latutude y = 0. The plots are local shear $U_y(0,t)$ (dotted red), and momentum flux $\langle u'v' \rangle$ from CE2 (blue), time-dependent SY14 (red) and instantaneous SY14 (yellow).

for $w < w_0$

$$\Delta(w,t) = \int_{t-w}^{t} \gamma(\tau) d\tau = \int_{t-w}^{t} \gamma(t) + \gamma'(t)(\tau - t) + \dots d\tau, \qquad (4.5.10)$$

$$= \gamma(t)w + O(1/\mu)d\tau. \tag{4.5.11}$$

In this situation the second order cumulant convergence is on a faster time scale than the mean flow dynamics and the steady shear approximation $\Delta(w,t) = \gamma w$ is the first order solution.

Note that this is only an explanation for the experiments presented. The above argument does not hold if the typical profile shear does not satisfy $\gamma_0 \gg (\mu l_{\rm max}/k_{\rm min})$, so if:

- $F \to 0$ (since $F \to 0$ requires $l_{max} \to \infty$), meaning time dependency must be considered in a detailed theoretical study of the local closure theory.
- The profile shear passes through, or close to, zero. The case of shear passing through zero is particularly interesting, as it means $\Delta(w,t)$ is non-monotonic which allows a significant, delayed, contribution from the integrand in (4.2.3).

These considerations explain very clearly why in the first experiment the time-independent solution is different from the CE2 solution near to the times when the background shear flow passes through zero and why the difference is skewed towards the times after the profile passes through zero. It is also put forward that the maximum effective meridional wavenumber associated with numerical limitations means that the solutions can agree well with each other once the (albeit dynamic)

background profile attains a sufficient shear.

In conclusion, the time-variability of second-order cumulants is typically on the same time-scale as the rate of change of the background zonal flow. This must be true as the zonal flow controls the evolution of second cumulants. This is true even in the limit of large zonostrophy. Notwithstanding, if the mean flow is in a quasi-equilibrium state, where the background flow only undergoes small fluctuations about it's mean state, it might be that the local closure theory without time dependence provides a good approximation to the QL dynamics in the zonostrophic limit. It would be particularly interesting to investigate if the errors in the proof of convergence between the QL and NL models [17, 70] can be adjusted in light of the findings put forward here. This is left as a potential for future work.

Chapter 5

Applications of the Local Closure Theory

In the previous Chapter it was established that the SY14 closure for the momentum flux (4.2.15) can accurately predict the direct contribution $\langle u'v'\rangle_D$ due to stochastic forcing, at least in the limit $F\to 0$. In the equilibrated jet flows, however, an equally important contribution $\langle u'v'\rangle_S$ is present due to barotropic instabilities at both the eastward and westward jets. This Chapter address the question of whether, when and how exactly the SY14 closure can practically be used to model stochastically forced zonal flows, by exploring three settings of increasing complexity:

- 1. **A stable monotonic zonal flow**: For this class of flows the SY14 closure and, because there are no turning points, the simple approximation (4.1.2) can be used to predict $\langle u'v'\rangle_D$. Provided the stochastic forcing is not too strong, the equilibrium flow will remain in the stable regime and there will be no secondary instabilities, i.e. $\langle u'v'\rangle_S = 0$. This scenario is closely analogous to the 'vortex condensate' situation, in which it has been shown [e.g. 81] that zonal closures can be successful.
- 2. **A stable zonal flow with alternating jets**: In this scenario a strong 'radiative' damping is used to maintain the flow in a stable state consisting of an eastward and westward jet. This is a tougher test of SY14 than the monotonic flow above, since the flow must adjust to rapid changes in $\langle u'v'\rangle_D$ near the jet cusps.

However, it remains simpler than the β -plane jet flows of set-up 3, because by design the flow is stable and $\langle u'v'\rangle_S = 0$. This regime is key to understanding the mechanism by which the directly forced momentum fluxes $\langle u'v'\rangle_D$ lead to the onset of secondary instability.

3. A typical equilibrated jet flow in β -plane turbulence: In this flow we know from the previous Chapter that it is essential to include some representation of barotropic instability so that $\langle u'v'\rangle_S$ is captured. Our approach below is simply to replace the stochastic forcing with a deterministic term based on SY14 to model $\langle u'v'\rangle_D$, and to permit barotropically unstable waves to emerge spontaneously in the flow, by including non-zonal perturbations in the initial conditions. Therefore, our parametrisation in this case is not purely zonal, but can still be useful when $F \ll 1$, i.e. there is a large scale separation between the forcing scale and the jet scale, since the latter determines the scale of the barotropic waves.

A practical consideration when implementing the SY14 closure concerns the large gradients in the calculated $\langle u'v'\rangle_D$ which invariably emerge in thin boundary layers near jet cusps. Numerically these boundary layers can always be resolved with a sufficiently high resolution grid. However, it is often computationally more practical to apply a smoother to the momentum flux $\langle u'v'\rangle_D$ calculated from (4.2.15). Here we use a kernel smoother which acts on a function f(y) according to

$$f_{\sigma}(y) = \int_{-\infty}^{\infty} f(y') K_{\sigma}(y - y') \, \mathrm{d}y', \qquad (5.0.1)$$

where $K_{\sigma}(y)$ denotes a smoothing kernel with characteristic length scale σ . In practice a Gaussian with variance σ^2 is used. It turns out that, for a significant range of σ , this modified closure (denoted SY14 σ hereafter) actually improves the comparison with our numerical results, for the simple reason that the smoothing in (5.0.1) can replicate the effect of finite F (see Figure 4.4).

Results from the three different flow scenarios are reported next.

5.1 A Stable Monotonic Zonal Flow

To generate a stable monotonic zonal flow, the Rayleigh friction term in the original governing equation (1.1.21) is modified so that the flow is relaxed towards a prescribed profile $U_0(y) = \operatorname{erf}(y)$, modelling e.g. a wind stress term in an oceanic flow. In the y-direction sidewall boundaries at $y = \pm \pi$ replace periodicity, and to minimise wall effects, the stochastic excitation is confined to the shear zone, so the governing equation becomes

$$\zeta_t + \mathbf{u} \cdot \nabla \zeta + \beta v = -\mu(\zeta - \zeta_0) - v_4 \nabla^4 \zeta + \sqrt{\varepsilon f(y)} \eta, \qquad (5.1.1)$$

with $\zeta_0 = -2\exp(-y^2)/\sqrt{\pi}$. The local energy injection rate is now $\varepsilon f(y)$ with $f(y) = \exp\left(-y^2/2\sigma_n^2\right)/\sqrt{2\pi}\sigma_n$ and $\sigma_n = \pi/8$. Since the flow is monotonic it is not necessary to use the smoothed closure (SY14 σ described above), therefore results from (5.1.1) can be compared with those from the unsmoothed SY14 closure equation which is (c.f. 4.2.20)

$$\partial_t U = -\mu (U - U_0) - \frac{\varepsilon}{2\mu} \partial_y \left(f(y) G\left(\frac{2\mu}{U'(y)}\right) \right) + v_4 \partial_y^4 U, \tag{5.1.2}$$

with G(m) given by (4.2.15).

Figure 5.1 shows a comparison between the equilibrated states of the CE2 equations obtained from (5.1.1) with the equilibrated solution of the SY14 closure equation (5.1.2). The parameters $(\mu, (2\pi)^2 \varepsilon/\mu, \beta) = (0.002, 4, 2)$ have here been chosen in order that the stochastic forcing is rather weak, in the sense that the equilibrated flow remains close to the relaxation profile $U_0(y)$. It is clear (see middle panel) that SY14 does an excellent job of describing the deviation from $U_0(y)$ induced by the stochastic forcing in the CE2 model (the small differences here can be attributed to finite F effects).

Evidently, a class of channel flows exist in which zonal local closure theories are entirely successful. This explains the success of local theories in vortex condensate flows [81]. Here, the set-up is identical to the stochastically forced barotropic vorticity dynamics governed by (1.1.21), but with $\beta = 0$. Without planetary rotation

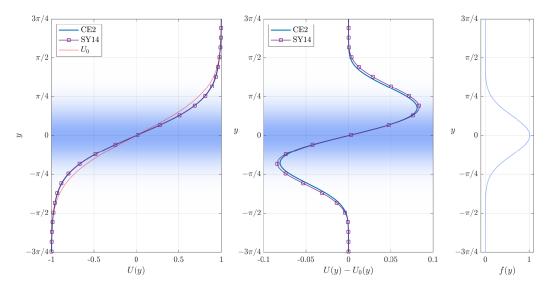


Figure 5.1: Left: U(y) from CE2 (blue) and the SY14 parameterised model (purple marked) for the a channel flow linearly relaxed to $U_0 = \operatorname{erf}(y)$ (red). Middle: similar to left, but plots the difference $U - U_0$. Right: the forcing profile f(y). The forcing profile is also indicated on the two left-most panels with a blue gradient indicating the forcing magnitude.

the large scale mean flow structures are domain sized vortex condensates. The WB19 inverse shear result (modified for polar geometry) has been shown to be accurate near to the vortex core [81, 92] where the shear is largest. However, it is noticeable that the solution is only accurate very close to the condensate core. It is expected that the SY14 local closure theory can be modified to give an improvement over the WB19 solution in this context as well, thus providing theoretical access to the profile of the vortex condensate up to a larger radius.

We conclude the discussion on monotonic zonal flows by mentioning that further experiments (not shown) reveal that, even in this simple setting, the regime in which the SY14 local zonal closure theory remains accurate is rather restricted. Increasing the forcing (higher ε) or broadening the forcing region (higher σ_n) tends to lead to the formation of local extrema with $U_y = 0$, which lead to secondary instabilities by a mechanism to be described in the next subsection. For the reasons described in Chapter 4, SY14 is sensitive to finite F effects at local extrema, and consequently local extrema deserve special attention, in the following flow scenario.

5.2 A Stable Zonal Flow With Alternating Jets

5.2.1 Set-Up and Main Results

To generate a stable alternating jet flow, equation (1.1.21) is modified to include a 'radiative relaxation' term, as has been used by e.g. [93] to model the effect of large scale radiative damping on the jets of the giant planets. The governing equation, solved on the doubly-periodic domain, becomes

$$\zeta_t + \mathbf{u} \cdot \nabla \zeta + \beta v = r(\bar{\psi} - \Psi_{\text{rad}}) - \mu \zeta - v_4 \nabla^4 \zeta + \sqrt{\varepsilon} \eta. \tag{5.2.1}$$

Here the streamfunction $\psi = \nabla^{-2}\zeta$, and $\Psi_{\text{rad}} = \cos(y)$ so that the zonal flow is effectively relaxed towards a sinusoidal 'deep jet' profile [c.f. 94], given here by

$$U_0(y) = \frac{1}{1 + (\mu/r)}\sin(y). \tag{5.2.2}$$

The reason for introducing the radiative relaxation term is that it selectively damps large scale eddies, and to maintain the stability of the equilibrium flow, it was found to be necessary to use a relatively strong relaxation and damping. This set-up thus allows exploration of strongly relaxed jets without necessarily causing excessive damping of eddy activity at the forcing scale.

The parameter settings for our main calculation, using the CE2 truncation of (5.2.1) are $(\beta, \mu, r, \varepsilon) = (22, 0.05, 0.05, 0.05/(2\pi)^2)$, solving numerically on a grid of 256 points. The large value of β was found to be necessary to suppress secondary instabilities. Useful insight into the expected behaviour at the jet cusps is provided by the scattering experiments of section 4.2.3, where typical momentum fluxes for a fixed sinusoidal jet were plotted for different values of the forcing scale parameter F. Strong sensitivity to F is therefore expected, and since it is not computationally feasible to explore the $F \rightarrow 0$ limit in these interactive CE2 calculations, we take $k_f = 16$ and aim to compare our results with the SY14 σ closure (i.e. equation (5.1.2) modified for this system, with the smoother (5.0.1) applied), which can capture the qualitative effects of finite F.

Figure 5.2 shows the comparison between the CE2 calculation and solutions obtained from the SY14 σ closure equation (solved numerically on a 2048 grid). The SY14 σ results show sensitivity to the smoothing parameter σ , reflecting the sensitivity to F in the CE2 momentum flux patterns seen in Figure 4.4. For the present calculation, an optimal value of $\sigma = \sigma_* = 0.237/k_f$ results in good agreement between CE2 and SY14 σ throughout the domain. Admittedly this good agreement requires a flow specific parameter fit for σ , and more research is required to determine how best finite F effects can be captured in a modified or extended SY14.

The eddy induced changes to U seen in Figure 5.2 serve to illustrate why the regime in which $\langle u'v'\rangle_S=0$, and the purely zonal closure based on SY14 σ is applicable, is restricted to a narrow region of parameter space in which the stochastic forcing is weak and the relaxation of the jet is strong. Because $-\partial_y \langle u'v'\rangle_D$ peaks strongly near the jet cores, even rather modest eddy forcing, which has little impact on the flow elsewhere, has a strong effect at the jet tips, acting to sharpen them. The jet curvature $|U_{yy}|$ at the jet tip and flanks increases rapidly with forcing strength, leading to $\beta - U_{yy}$ changing sign and thus to secondary instability. The effect is more pronounced as F is reduced, or σ is reduced in SY14 σ , as seen in Figure 5.2. Through this sharpening mechanism, stochastic forcing at a jet cusp has a much stronger tendency to lead to secondary instability than elsewhere, explaining why the stable alternating jet regime discovered here occupies such a narrow region of parameter space, especially so at low F.

5.2.2 Additional Numerical Results

We briefly report that additional parameters setting for the radiatively damped experiment were explored, labelled A-E with values for μ , r, ε , β and σ^* as given in Table 5.1.

We first report the results from simulations A-D. In general, all the results from A-D are similar, and the SY14 σ and CE2 solutions agree well. Quantities from the two most extreme experiments in this set, A and D, are plotted in Figure 5.3. Note that experiment D is the 'main' parameter setting reported in the previous section. In all cases the profiles become sharpened at both the eastward and westward jet

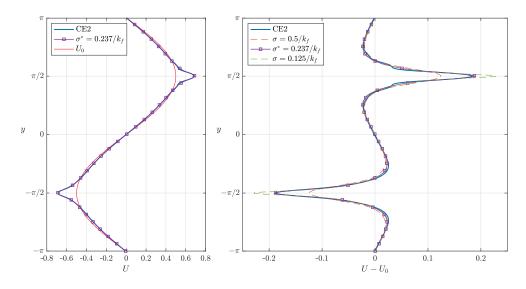


Figure 5.2: CE2 and SY14 σ quantities from the radiatively damped experiment. The top panel compares the mean wind profile with U_0 for CE2 and SY14 σ (with the optimal $\sigma^* = 0.237/k_f$). The bottom panel investigates results in more detail by plotting the profile deviation from the radiatively relaxed profile, $U - U_0$. SY14 σ results are given for $\sigma = [0.125/k_f, \sigma^*, 0.5/k_f]$. In all cases $k_f = 16$.

Reference	μ	r	$(2\pi)^2 \varepsilon/\mu$	β	$k_f \sigma^*$ (3.s.f.)
A	0.05	0.05	0.05	4	0.207
В	0.05	0.05	0.20	8	0.215
C	0.05	0.05	0.40	12	0.226
D	0.05	0.05	1.00	18	0.237
E	0.01	0.01	0.05	6	0.223

Table 5.1: Parameter setting references for the radiatively damped model.

cores. As ε/μ is increased (A to D), the force from the parameterised term takes a stronger role in the momentum balance and sharpens the profile to a greater degree. In simulation A, the profile only differs no more than 5% from U_0 at any latitude and the small β value means there is very little difference between CE2 and SY14 σ results. The largest disparities are found in the regions surrounding the cusps. The disagreement is east-west symmetric, implying it is not due to next order β effects, but rather finite k_f and the manner in which the smoothing function non-locally distributes the large Reynolds stress divergence at the jet core to the surrounding region in SY14 σ . Moving from A through to D, the profile sharpening must be balanced by increasing β , hence the large value for β reported for the main study result. Surprisingly, the 4.5 times increase in β doesn't change the picture that much

between experiments A and D.

Experiment E has similar parameters to experiment A except with a reduced friction. The results from experiment E are plotted in Figure 5.4 and offer a similar analysis to the results from experiment A. That is, CE2 and SY14 σ agree well except with the same small disagreement in the flanks. Reducing μ means the profile from experiment E has slightly sharper jets than experiment A as the boundary layer regions around the jet extrema in which $\langle u'v'\rangle_D$ regularises the WB19 type solution is narrowed. Consequently the required β is larger for experiment E than experiment A. It is not expected that more experiments with $\mu=0.01$ will reveal qualitatively different behaviour to the simulations A-D. Reducing ε/μ (as in experiments A-D) would simply follow a similar pattern, except that β will need to be slightly larger in the CE2 experiments to ensure stability.

Reassuringly the optimal smoothing parameter σ^* (which is calculated by minimising the l_2 -norm of the difference between the SY14 σ solution and the CE2 result) remains relatively consistent and O(1) throughout the experiments. Practically, this meant that finding sensible β values was quite simple, since an estimate for a minimum β could be made from examining an SY14 σ solution with a ball-park σ value. Going forwards, now that we have a method for dealing with cusps, we are now ready to apply the SY14 σ parameterisation in the more realistic equilibrated jet flow to be studied next.

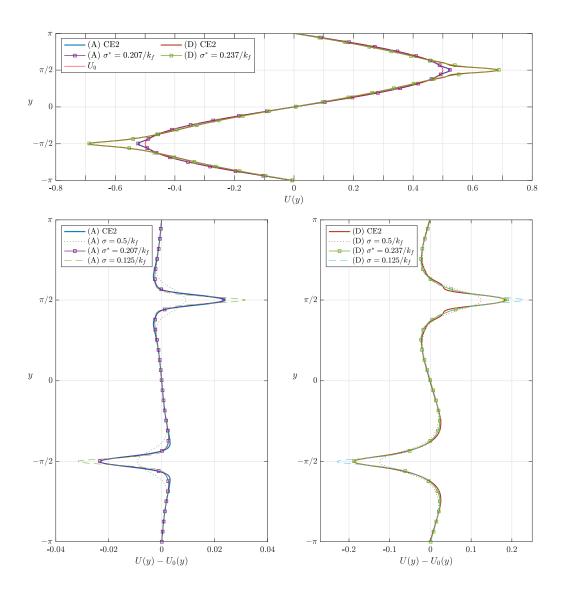


Figure 5.3: CE2 and SY14 σ quantities from the radiatively damped experiments A and D. The top panel shows the solution U from CE2 and SY14 σ when $\sigma = \sigma^*$. The bottom panels give a closer inspection of the solutions by plotting the difference with the equilibrium radiative solution, $U - U_0$. In these panels solutions with additional values of σ are also plotted ($\sigma = [0.125/k_f, \sigma^*, 0.5/k_f]$, in all cases $k_f = 16$).

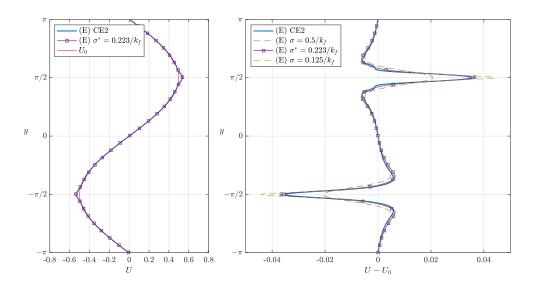


Figure 5.4: CE2 and SY14 σ quantities from the radiatively damped experiment E. The left panel plots U from CE2 and SY14 σ when $\sigma = \sigma^*$, and the right panel shows the difference with the equilibrium radiative solution, $U - U_0$. In the right panel solutions are presented for $\sigma = [0.125/k_f, \sigma^*, 0.5/k_f]$ (in all cases $k_f = 16$).

5.3 A Typical Equilibrated Jet Flow in β -plane

Turbulence

5.3.1 Set-Up and Main Results

Finally, we are ready to address the question of whether the SY14 σ closure can be adapted to parameterise the stochastic term in the canonical nonlinear β -plane turbulence simulations. Recall that SY14 σ can capture only the directly forced momentum flux $\langle u'v'\rangle_D$, and not the secondary flux $\langle u'v'\rangle_S$ due to barotropic instability, which we know from section 4.3 is equally important in the equilibrated momentum balance. In the absence of a (known) means to parameterise $\langle u'v'\rangle_S$, our approach is simply to allow unstable waves to develop spontaneously in the parameterised flow, by solving the deterministic nonlinear equation

$$\zeta_t + \mathbf{u} \cdot \nabla \zeta + \beta v = -\mu \zeta + (-1)^{n+1} v_{2n} \nabla^{2n} \zeta - \frac{\varepsilon}{2\mu} \left[\partial_{yy} G \left(\frac{2\mu}{\overline{\zeta}} \right) \right]_{\sigma}, \quad (5.3.1)$$

where the square brackets $[\cdot]_{\sigma}$ denote that the smoother (5.0.1) is applied. Equation (5.3.1) is simply (1.1.21) with the stochastic term replaced by the deterministic SY14 σ parameterisation.

Figure 5.5 compares PV and vorticity snapshots in the NL simulation of section 4.3.2 with those from a corresponding integration of the parameterised equation (5.3.1). The parameter settings and numerical configurations are identical, and both simulations are integrated for the same length of time until an equilibrated state is reached. The value $\sigma = 0.237/k_f$, found in the alternating jet experiment above, is used for the smoothing parameter in SY14 σ . The same WF3 forcing with $k_f = 16$ is being used, and the smoothing is performed on 2048 grid points.

There are several striking similarities between the two simulations apparent in Figure 5.5:

- Excellent correspondence between the structure of the equilibrated jets in the two simulations.
- Relatively short waves $(k_x = 4-7)$ propagating on the PV barriers at the core

of the eastward jets. These correspond to the eastward jet instability identified in the linear analysis of Figure 4.8.

• Longer waves ($k_x = 1-3$) propagating on the PV barriers located at the flanks of the westward jets. Note that, while these waves are only clearly visible in snapshots of the parameterised run, their presence in the NL simulation is evident from the momentum flux decomposition of Figure 4.6. They correspond to the waves generated by the westward jet instability identified in Figure 4.8.

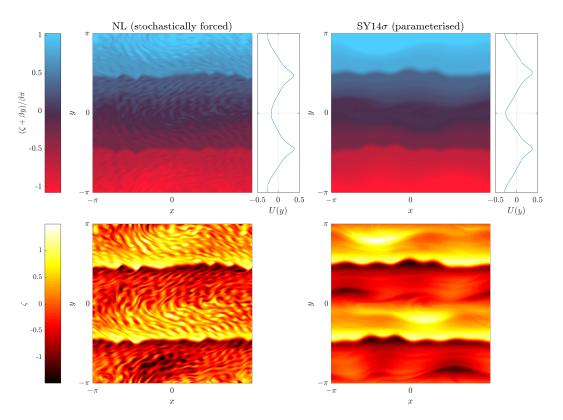


Figure 5.5: Left: snapshots of potential vorticity $\zeta + \beta y$ and relative vorticity ζ from the NL simulation reported in Figures 4.5 and 4.6. Right: snapshots of the same quantities from an SY14 σ simulation with the same parameters. In the SY14 σ model the stochastic forcing is replaced with the deterministic forcing term described in the main text ($\sigma = 0.237/k_f$, $k_f = 16$).

From the evidence compiled above, it is clear that SY14 σ is successful because it captures the direct momentum flux $\langle u'v'\rangle_D$ sufficiently accurately to permit the flow to reach a marginally stable equilibrium in which emergent waves generate the correct secondary flux $\langle u'v'\rangle_S$. Further analysis of the momentum budget is

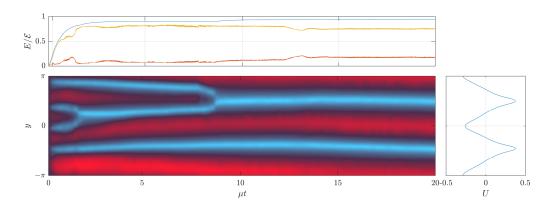


Figure 5.6: Hovmöller plot of the SY14 σ model for fully developed barotropic β -plane jets with $\sigma = 0.23/k_f$. Plotted is energy against time (total energy (blue), mean flow energy (orange), eddy kinetic energy (yellow)).

proveided in Figure 5.7, supporting this interpretation. Here, a time average, over the final μ^{-1} time period of the simulation, for the components $-\partial_y \langle u'v' \rangle_{k_x}$ is reported in the panels of Figure 5.7. In the left panel are waves $k_x = 1$ and 2, and in the right panel the waves 3–8. It is evident that the longest waves are working in symbiotic manner similar to the NL simulation (e.g. Figure 4.6) so that their overall force on the mean flow is to damp the westward jet extrema. Similarly the shorter waves in the right hand panel concentrate their overall effect on the eastward jets. Waves greater than 8 do not give a significant Reynolds stress contribution. Note that the magnitude of the components is different in each of panels (so that the right-hand panel details are visible), and the magnitudes typically lay between the equivalent quantities in the NL and QL simulations (Figure 4.6).

The Hovmöller plot in Figure 5.6 reveals that the spin-up time for the SY14 σ simulation is on a significantly longer time-scale than the NL simulation (see Figure 4.5). This is remedied in the next subsection considering alternative values for σ . Also plotted in Figure 5.6 is the energy against time during spin-up for SY14 σ . From an energetics point of view, the parameterised term is the energy source in this model, injecting energy into the mean flow, and the instabilities emerge by drawing energy from the mean flow. The smoothing process described reduces the energy that the parameterisation supplies the SY14 σ simulation. If the smoothing is set up ideally, the 'missing' energy in the SY14 σ model simply corresponds to the eddy

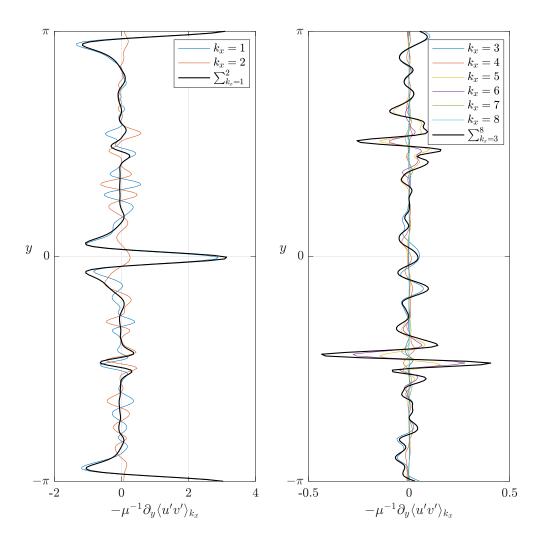


Figure 5.7: Wavenumber decomposition for momentum flux in the SY14 σ model with $\sigma = 0.237/k_f$.

kinetic energy associated with the (now parameterised) small scale perturbations. It noticeable that there are small jumps in the total energy when jets merge, this is because the smoothing process acts most severely at jet extrema, hence reducing the number of jet extrema increases the energy injection from the parameterised forcing term.

The main drawback of SY14 σ is that it doesn't give a computational saving compared to the NL simulation, in fact the computational cost is almost identical as both runs are on the same grid over the same time period. However, this need not be true in general, particularly as F is reduced. This is because, as inspection

of Figure 5.5 confirms, the length scale of waves in the parameterised simulation is set by the jet spacing, not by the forcing scale. Therefore a scenario with $F \ll 1$ in which the parameterised run is many orders of magnitude cheaper than NL would not be difficult to set up.

5.3.2 Addition Numerical Results

Figure 5.8 shows vorticity contour plots of vorticity during spin-up from the same SY14 σ simulation in the previous section. Besides their pretty looks, these contour plots reveal much about the spin-up and equilibration of the SY14 σ model. It is observed that in the first panel, when $t = 0.02/\mu$, the flow is essentially a parallel shear flow. At this point, the parameterised term has been providing energy to the mean flow via the small initial zonal perturbation. The effect of the parameterisation is to create, and sharpen, 'jets' at the smoothing scale (i.e. the forcing scale k_f by design). The initial, small, random non-zonal perturbations are negligible to the dynamics until the second panel $(t = 0.04/\mu)$, where the background flow has become sufficiently unstable, and enough time has elapsed that the perturbations have excited secondary instabilities. By the third panel $(t = 0.08/\mu)$, it appears the spin-up has entered a phase of nonlinear turbulent equilibration. Being careful not to read too much into the spin-up and equilibration procedure, it is nevertheless interesting to note that as time progresses, and the jets become wider, the typical unstable zonal wave numbers decrease. This agrees with the interpretation that the length scale of the emergent secondary instabilities scale in proportion to the jet scale, and therefore that there is a length scale separation between the secondary instabilities and the small forcing length scale.

We now report results from identical parameterised simulations as the above SY14 σ simulation except with different smoothing parameter values σ . The values $\sigma = 0.5/k_f$ and $\sigma = 0.125/k_f$ are tested (in contrast with $\sigma = \sigma* = 0.237/k_f$ in the 'main' experiment), the same as the test σ values reported in the radiatively damped model of the previous section.

Figures 5.9 and 5.10 show the results from the simulation with $\sigma = 0.5/k_f$. In this situation, the parameterised term is being over-smoothed. The Hovmöller

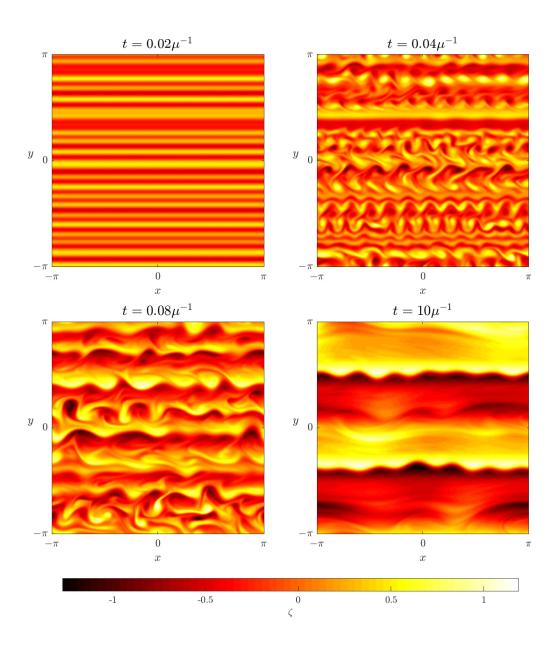


Figure 5.8: Contour plots showing the spin-up of the SY14 σ simulation reported in Figure 5.5. Vorticity is contoured at times $t = [0.02, 0.04, 0.08, 10]/\mu$ (left to right, top to bottom).

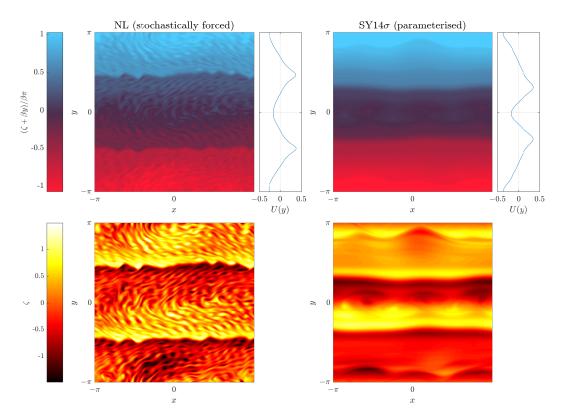


Figure 5.9: As Figure 5.5, but with $\sigma = 0.5/k_f$ for the right hand panels.

plot in Figure 5.10 reveals the haphazard and peculiar route the mean flow takes to find the final structure shown in Figure 5.9. Without the sharp eastward jets the eastward instabilities are at smaller wavenumbers ($k_x = 1,2$). A westward jet instability at wavenumber $k_x = 3$ is particularly pronounced in Figure 5.9. Despite the over-smoothing generating a relatively unrecognisable profile, it is clear that qualitatively the same marginal stability processes are occurring.

Figures 5.11 and ?? are results from the simulation with $\sigma = 0.125/k_f$ which, given the radiatively damped model's results, we might have expected to be an under-smoothing scenario. It happens however that dynamics from this simulation are very similar to the $\sigma = \sigma*$ simulation, implying that there's quite a wide region of viable smoothing parameters which may predict realistic equilibrium jet structures. A significant difference is that the $\sigma = 0.125/k_f$ simulation reached the two jet state on the same timescale as the NL simulation (c.f. Figure 4.5). This is due to the jets being sharpened more readily by the parameterised term than in the simulation with $\sigma = \sigma^*$.

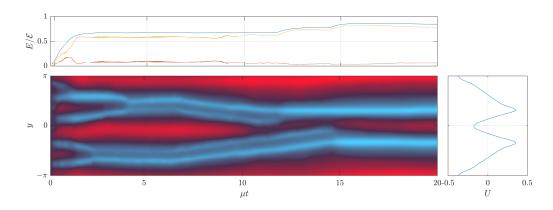


Figure 5.10: As Figure 5.6, but with $\sigma = 0.5/k_f$. Plotted is energy against time (total energy (blue), mean flow energy (orange), eddy kinetic energy (yellow)).

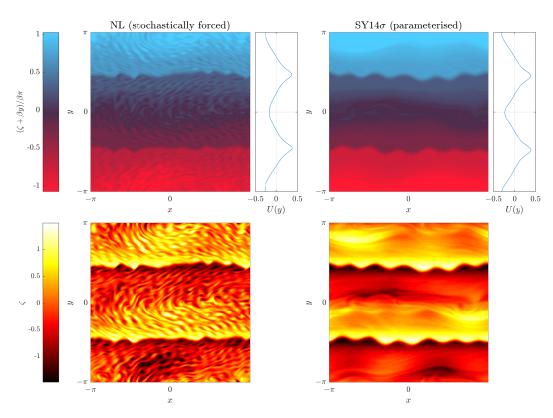


Figure 5.11: As Figure 5.5, but with $\sigma = 0.125/k_f$ for the right hand panels.

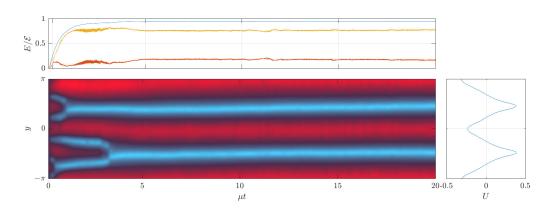


Figure 5.12: As Figure 5.6, but with $\sigma = 0.125/k_f$.

Chapter 6

General Conclusions

The initial, broad, objective of this work was to research planetary jets in β -plane turbulence models. Over the course of the work, the aim has been refined into two distinct and well-known topics related to the QG2D model: spontaneous jet transitions, and statistical closure theories. We first give an overview of the former, before discussing the latter and drawing the final conclusions.

An investigation into rare jet transitions is the subject of Chapter 2. This was inspired by, and gives an in depth account of the work by Bouchet et. al. [18]. The adaptive multilevel splitting (AMS) algorithm is introduced and verified to work in the barotropic setting. Noting that similar rare jet transitions occur in a two-layer (baroclinic) model, the original research of Chapter 2 begins with the design of a new algorithm which can be applied to the two-layer set-up. The chief difference here is that as the two-layer model is deterministic, and as the AMS algorithm relies on stochasticity to 'split' flow trajectories, the AMS algorithm cannot systematically find rare events in the deterministic setting. Considering the two-layer model from a dynamical systems perspective, a modification of AMS to a deterministic adaptive multilevel splitting (DAMS) algorithm is postulated. The idea is that the two-layer model is chaotic, so a small random change in initial conditions provides a natural 'perturbation' to be added to make a 'cloned' trajectory 'split'. The chaotic nature of baroclinic turbulence is found to decorrelate the 'clones' rapidly enough that DAMS successfully guides ensembles towards rare events. In fact, spontaneous transitions between attractors are found using DAMS that are so rare that they were not found with comparable computational power by direct numerical simulation.

The application of AMS or DAMS comes with daunting numerical difficulties such as large computation time and required memory. Both of these hurdles are overcome, primarily using savvy storage techniques and realistic parameter regimes, rather than state-of-the-art resources. It is therefore a wonderful prospect that these ideas, with more adept tools, should be able to quantify rare events dramatically more efficiently than direct numerical simulation.

The second main topic of the thesis starts with Chapter 3, where the QL approximation and its related closure theory, CE2, are introduced and discussed. Chapter 4, in which the solution to the CE2 equation is considered in the limit of small scale forcing is arguably the theoretical pinnacle of the thesis. The idea that a small scale forcing only 'sees' the local properties of the mean flow (borrowing heavily from the work in Srinivasan & Young [80] and Woillez & Bouchet [23]), allows a 'local closure theory' for momentum flux to be developed. Exact formulae for momentum flux components are derived based on the local closure theory, with the physical insight that the most important aspect of the forcing structure is the phase density of its energy injection rate. The formulae are shown to be accurate in the limit of small scale forcing through a series of numerical experiments. It is shown in this work that the structure of the forcing in high zonostrophy flows is, typically, unimportant.

A fundamental mechanism in jet development and maintenance is made apparent in Chapter 4 when it is shown that the QG2D model forced on a single (relatively large) zonal wave mode generates symmetric, tightly packed, jets. The fact that these jets break the Rayleigh-Kuo condition for instability, and that when small perturbations on longer wave modes are introduced the asymmetry and classical Rhines jet spacing return, implies that emergent barotropic instabilities are important. This is not a new idea [i.e. 22, 23], but section 4.3 isolates the exact roles of the emergent instabilities through careful experimental design. In particular, it is found that there are instabilities associated with the westward *and* eastward jet extrema.

A successful model for the large-scale β -plane jet dynamics is presented in Chapter 5 where the small scale turbulence in the QG2D model is accurately captured

using a parameterisation based on the local closure theory. The mean flow in the model is 'driven' towards sharp jets by the parameterised forcing, while emergent instabilities are allowed to develop and regulate the growth of jets in a similar manner to the full model. It is remarkable that the parameterised model - which is derived assuming a quasi-linear approximation and allows no interaction between the emergent instabilities and the small scale forcing - is in many respects comparable to the equivalent fully nonlinear stochastically forced simulation at the same parameters. The nonlinear barotropic dynamics, with persistent non-zonal coherent structures, is a stringent test for the efficacy of a local closure parameterisation. The success implies that there may be great benefit for the local closure theory parameterisation in the expanding list of turbulence-mean interaction problems that the CE2 framework is proving useful.

A critical message of this thesis is that the time-variability of eddy statistics in closure models has been overlooked in the literature. It has previously been believed that β -plane turbulence can be viewed as a time-scale separation problem between the (slow) zonal mean flow and the (fast) turbulent fluctuations (Bouchet et. al. [17]), with the claim that there is a formal asymptotic convergence of the QL and NL models in the limit of time-scale separation (high zonostrophy). The basis of the proof is found to be false in section 4.5 when it is shown theoretically, and experimentally, that eddy statistics do not 'instantaneously' converge. Rather, eddy statistics converge on the slow timescale and therefore the history of the mean flow profile is important within the closure frameworks. This must be considered in future studies. In particular, in works that attempt to approach rare events problems in turbulence using a statistical closure approach. From the evidence presented, it is almost certain that the rare events studied in the first half of the thesis exhibit mean flow dynamics that will provide a significant time-variability of second order eddy statistics.

Appendix A

SY14 Local Closure Theory

A.1 Sheared Disturbance Solution

Using collective coordinates (as suggested in [16])

$$y = y_1 - y_2,$$
 (A.1.1)

$$\bar{y} = \frac{y_1 + y_2}{2},\tag{A.1.2}$$

Laplacian operators are transformed for n = 1, 2 by

$$\nabla_n^2 = \nabla^2 - (-1)^n \partial_y \partial_{\bar{y}} + \frac{1}{4} \partial_{\bar{y}}^2. \tag{A.1.3}$$

If $U_y = \gamma$, $\tilde{Z}_k(y, \bar{y})$ in 3.3.12 must be independent of \bar{y} . This is quite remarkable given the β terms, but without curvature the Coriolis parameter is unimportant. The second cumulant equation (3.3.9) in the collective coordinates reads simply

$$\partial_t \tilde{\mathcal{Z}}_k + iky\gamma \tilde{\mathcal{Z}}_k = -2\mu \tilde{\mathcal{Z}}_k + \varepsilon \tilde{\Pi}_k. \tag{A.1.4}$$

With a linear shear, $U = \gamma y$, and the second cumulant equation above with a linear shear is

$$\partial_t \hat{\mathcal{Z}}_{k} + \gamma(t)k\partial_x \hat{\mathcal{Z}}_{k} = -2\mu \hat{\mathcal{Z}}_{k} + \varepsilon \hat{\Pi}_{k}. \tag{A.1.5}$$

This is a first order PDE subject to an initial condition

$$\mathcal{Z}_{\boldsymbol{k}}(0) = \mathcal{Z}_0(\boldsymbol{k}),\tag{A.1.6}$$

which can be solved by the method of characteristics. The characteristics satisfy

$$\frac{\mathrm{d}t}{\mathrm{d}s} = 1, \quad \frac{\mathrm{d}k}{\mathrm{d}s} = 0, \quad \frac{\mathrm{d}l}{\mathrm{d}s} = -\gamma(t)k, \quad \frac{\mathrm{d}\hat{Z}_k}{\mathrm{d}s} = -2\mu\mathcal{Z} + \varepsilon\hat{\Pi}.$$
 (A.1.7)

so characteristic projections are defined by

$$(t(s), k(s), l(s)) = (s, k_0, l_0 - \Gamma(s)k_0), \tag{A.1.8}$$

where $\Gamma(s) = \int_0^s \gamma(u) du$. Integrating along the projections

$$\hat{\mathcal{Z}}_{k}(t) = \int_{0}^{t} \varepsilon \hat{\Pi}(k(q), l(q)) e^{-2\mu(t-q)} dq + \hat{\mathcal{Z}}_{0}(k_{0}, l_{0})$$
(A.1.9)

$$= \int_0^t \varepsilon \hat{\Pi}(k, l + (\Gamma(t) - \Gamma(t - w))k) e^{-2\mu w} dw + \hat{Z}_0(k_0, l_0)e^{-2\mu t}, \quad (A.1.10)$$

where the change of variable w = t - q has been made for the final line. At late times $(t \gg \mu^{-1})$

$$\hat{\mathcal{Z}}_{k}(t) = \varepsilon \int_{0}^{\infty} \hat{\Pi}(k, l + k\Delta(w, t)), \qquad (A.1.11)$$

where we have introduced the time history integral

$$\Delta(w,t) = \int_{t-w}^{t} \gamma(\tau) d\tau. \tag{A.1.12}$$

A.2 Results for Reynolds Stress Tensor on Steady Background Flow

A.2.1 Reynolds Stress $\langle u'v' \rangle$

Assuming a steady background flow $\Delta = \gamma w$, the Reynolds stress found in the main text (equation 4.2.7) is

$$\langle u'v'\rangle = -\varepsilon \int_{\mathbb{R}^2} \left\{ -\frac{\tilde{\Pi}_{\boldsymbol{k}}}{2\gamma |\boldsymbol{k}|^2} + \frac{\mu}{\gamma} \int_0^\infty \frac{e^{-2\mu w}}{k^2 + (l - kw\gamma)^2)} \tilde{\Pi}(k, l) \mathrm{d}w \right\} \mathrm{d}\boldsymbol{k}.$$

By (4.2.2) the first term is ε/γ (i.e. the inverse shear result).

For the latter term we proceed initially as in SY14, transforming the k integral into polar form with $(k,l) = (\kappa \cos \phi, \kappa \sin \phi)$ for

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} - \frac{\varepsilon\mu}{\gamma} \int_{-\pi}^{\pi} \int_{0}^{\infty} \int_{0}^{\infty} \frac{\hat{\Pi}(\kappa\cos\phi, \kappa\sin\phi)e^{-2\mu w}}{\cos^{2}\phi + (\sin\phi - \gamma w\cos\phi)^{2}} \frac{1}{\kappa} dw d\kappa d\phi. \quad (A.2.1)$$

Instead of specifying a radial structure of $\hat{\Pi}$, we recognise that the κ integral serves only to determine the energy injection at a given wave angle ϕ . The symmetry $\hat{\Pi}(\mathbf{k}) = \hat{\Pi}(-\mathbf{k})$ implies the wave angles ϕ and $\phi + \pi$ have equal energy input. The fact that the denominator of the integrand is also invariant under $\phi \to \phi + \pi$ means we can let $\rho^{\varepsilon}(\phi)$ represent the total energy injection of the wave angles ϕ and $\phi + \pi$ (as given by 4.2.8) and consider the integral on the half domain $\phi \in (-\pi/2, \pi/2]$, i.e.

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} + \frac{2\varepsilon\mu}{\gamma} \int_{-\pi/2}^{\pi/2} \int_{0}^{\infty} \rho^{\varepsilon}(\phi) \sec^{2}\phi \frac{e^{-2\mu w}}{1 + (\tan\phi - \gamma w)^{2}} dw d\phi, \qquad (A.2.2)$$

The change of variable $s = \gamma w - \tan \phi$ finds

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} \left(1 - \frac{2\mu}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \sec^2 \phi e^{-2\mu \tan \phi/\gamma} \left(\int_{-\tan \phi}^{\infty} \frac{e^{-2\mu s/\gamma}}{1 + s^2} ds \right) d\phi \right). \tag{A.2.3}$$

The inside s integral can be performed with the general result

$$\int_{-x}^{\infty} \frac{e^{-ms}}{1+s^2} ds = -\operatorname{Im} \left\{ e^{mi} E\left(m(-x+i)\right) \right\}, \tag{A.2.4}$$

allowing us to write

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} \left(1 + \frac{1}{m} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) m^{2} \sec^{2}(\phi) \operatorname{Im} \left\{ e^{m(-\tan\phi + i)} E\left(m(-\tan\phi + i)\right) \right\} d\phi \right), \tag{A.2.5}$$

where $m = 2\mu/\gamma$. Equivalently, this can be written in the form in the main text

$$\langle u'v'\rangle = \frac{\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) K(\phi, m) d\phi, \quad K(\phi, m) = 1 + \frac{1}{m} |z|^{2} \operatorname{Im} \left\{ e^{z} E(z) \right\}. \quad (A.2.6)$$

where $z(\phi, m) = m(-\tan \phi + i)$.

A.2.2 $\langle u'u' \rangle$ and $\langle v'v' \rangle$

To derive formulae for $\langle u'u' \rangle$ and $\langle v'v' \rangle$, we can use analogous calculations to those above for $\langle u'v' \rangle$. It is straightforward to show that the time-dependent solutions are

$$\langle u'u'\rangle = \frac{2\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \sec^2 \phi e^{-m\tan\phi} \left(\int_{-\tan\phi}^{\infty} \frac{s^2 e^{-ms}}{(1+s^2)^2} ds \right) d\phi \tag{A.2.7}$$

$$\langle v'v'\rangle = \frac{2\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon}(\phi) \sec^2 \phi e^{-m\tan \phi} \left(\int_{-\tan \phi}^{\infty} \frac{e^{-ms}}{(1+s^2)^2} ds \right) d\phi$$
 (A.2.8)

To analyse further, the following results are useful

$$\begin{split} 4\int_{-t}^{\infty} \frac{s^2 e^{-ms}}{(1+s^2)^2} \mathrm{d}s &= \left[(m+\mathrm{i}) e^{-\mathrm{i}m} E_1(-m(s-\mathrm{i})) + (m-\mathrm{i}) e^{\mathrm{i}m} E_1(-m(s+\mathrm{i})) + \frac{2s e^{-ms}}{1+s^2} \right]_{-\infty}^{-x}, \\ &= -(m+\mathrm{i}) e^{-\mathrm{i}m} E_1(-m(t+\mathrm{i})) - (m-\mathrm{i}) e^{\mathrm{i}m} E_1(m(-t+\mathrm{i})) - \frac{2t e^{mt}}{1+t^2}, \\ 4\int_{-t}^{\infty} \frac{e^{-ms}}{(1+s^2)^2} \mathrm{d}s &= \left[(m-\mathrm{i}) e^{-\mathrm{i}m} E_1(m(s-\mathrm{i})) + (m+\mathrm{i}) e^{\mathrm{i}m} E_1(m(s+\mathrm{i})) - \frac{2s e^{-ms}}{1+s^2} \right]_{-\infty}^{-x}, \\ &= (m-\mathrm{i}) e^{-\mathrm{i}m} E_1(-m(t+\mathrm{i})) + (m+\mathrm{i}) e^{\mathrm{i}m} E_1(m(-t+\mathrm{i})) + \frac{2t e^{mt}}{1+t^2}. \end{split}$$

(These are obtained by writing the denominator as partial fractions.)

Again, letting $z = m(-\tan \phi + i)$, we have

$$\begin{split} 2e^{-m\tan\phi} \int_{-\tan\phi}^{\infty} \frac{s^2 e^{-ms}}{(1+s^2)^2} \mathrm{d}s &= \frac{-(m+\mathrm{i})e^{z^*} E_1(z^*) - (m-\mathrm{i})e^z E_1(z)}{2} - \frac{\tan\phi}{\sec^2\phi}, \\ &= -m \mathrm{Re} \left\{ e^z E_1(z) + \frac{1}{z^*} \right\} - \mathrm{Im} \left\{ e^z E_1(z) \right\}, \\ 2e^{-m\tan\phi} \int_{-\tan\phi}^{\infty} \frac{e^{-ms}}{(1+s^2)^2} \mathrm{d}s &= \frac{(m-\mathrm{i})e^{z^*} E_1(z^*) + (m+\mathrm{i})e^z E_1(z)}{2} + \frac{\tan\phi}{\sec^2\phi}, \\ &= m \mathrm{Re} \left\{ e^z E_1(z) + \frac{1}{z^*} \right\} - \mathrm{Im} \left\{ e^z E_1(z) \right\}, \end{split}$$

so the solutions for $\langle u'u' \rangle$ and $\langle v'v' \rangle$ are

$$\begin{split} \langle u'u' \rangle &= \frac{\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon} \frac{|z|^2}{m} \left(-\frac{1}{m} \mathrm{Im} \left\{ e^z E_1(z) \right\} - \mathrm{Re} \left\{ e^z E_1(z) + \frac{1}{z^*} \right\} \right) \mathrm{d}\phi, \\ \langle v'v' \rangle &= \frac{\varepsilon}{\gamma} \int_{-\pi/2}^{\pi/2} \rho^{\varepsilon} \frac{|z|^2}{m} \left(-\frac{1}{m} \mathrm{Im} \left\{ e^z E_1(z) \right\} + \mathrm{Re} \left\{ e^z E_1(z) + \frac{1}{z^*} \right\} \right) \mathrm{d}\phi. \end{split}$$

A.3 Kernel Properties

Two useful expansions of the exponential integral are:

$$E_1(z) = \gamma_1 - \ln(z) - \sum_{n=1}^{\infty} \frac{(-z)^n}{nn!},$$
(A.3.1)

where γ_1 is Euler's constant [95, p. 229], and

$$e^{z}E_{1}(z) = \sum_{n=0}^{\infty} \frac{(-1)^{n}n!}{z^{n+1}},$$
 (A.3.2)

for large values of Re(z) [96, p. 3].

Proofs of the six properties of $K(\phi, m)$ are as follows:

1. $K(\phi, m) < 1$ for all $\phi \in (-\pi/2, \pi/2)$ and all $m \in \mathbb{R} \setminus \{0\}$.

The definition of $K(\phi, m)$ (equation 4.2.11) prompts the observation

$$\operatorname{sgn}\left(\frac{1}{m}|z|^{2}\operatorname{Im}\left\{e^{z}E_{1}(z)\right\}\right) = \operatorname{sgn}\left(\frac{1}{m}\operatorname{Im}\left\{e^{\operatorname{i}m}\int_{m(-\tan\phi+\mathrm{i})}^{\infty}\frac{e^{-t}}{t}\mathrm{d}t\right\}\right),$$

$$= \operatorname{sgn}\left(\frac{1}{m}\operatorname{Im}\left\{\int_{-m\tan\phi}^{\infty}\frac{s-\mathrm{i}m}{s^{2}+m^{2}}e^{-s}\mathrm{d}s\right\}\right),$$

$$= \operatorname{sgn}\left(-\int_{-m\tan\phi}^{\infty}\frac{e^{-s}}{s^{2}+m^{2}}\mathrm{d}s\right),$$

$$= -1,$$

proving that $K(\phi, m) < 1$.

- 2. The limiting values at the boundaries of the interval are $K(\pm \frac{\pi}{2}, m) = 0$. Corollary of property (v), see below.
- 3. $K(\phi,m)$ has a single minimum $K_{-}(m)$. In the limit $m \to 0$, $K_{-}(m) \sim -4\pi e^{-2}/m$, and the location of the minimum is asymptotically close to $\phi = \pi/2 m/2$.

It can be shown $K(\phi, m)$ solves the differential equation

$$\cos^2 \phi \frac{\partial}{\partial \phi} K(\phi, m) = (\sin 2\phi - m) K(\phi, m) - \sin 2\phi, \tag{A.3.3}$$

so any minimum must solve

$$K_{-}(m) = \frac{\sin(2\phi_{-})}{\sin(2\phi_{-}) - m}.$$
(A.3.4)

Seeking a minimum point ϕ_- such that $c=m/\theta_-=O(1)$, where $\theta_-=\pi/2-\phi_-$, the leading order behaviour for $K(\phi_-,m)$ as $m\to 0$ is found from (A.3.1) as

$$K(\phi_{-}, m) = -\frac{1}{|m|} \left(\frac{m}{\theta_{-}}\right)^{2} \pi e^{-\frac{m}{\theta_{-}}} + O(1).$$
 (A.3.5)

If the minimum point exists, the leading order terms of (A.3.4) and (A.3.5) balance, i.e.

$$-\frac{c^2}{|m|}\pi \exp(-c) = \frac{2}{c-2} + O(m^2), \tag{A.3.6}$$

which rearranges for

$$c = 2\left(1 - |m|\left(\frac{1}{c^2\pi e^{-c}}\right) + O(m^2)\right) = 2 + O(m).$$
 (A.3.7)

Since c = O(1), the solution is self-consistent and in the small m limit

$$K_{-}(m) = -\frac{4\pi e^{-2}}{m} + O(1)$$
, with $\phi_{-} = \pi/2 - m/2 + O(m^2)$. (A.3.8)

4. As $m \to 0$, $K(\phi, m) = 1 - m(\frac{\pi}{2} + \phi) \sec^2 \phi + O(m^2 \log m)$.

For any fixed ϕ , the complex number $z = m(-\tan \phi + i)$ has infinitesimal magnitude in the limit $m \to 0$. Thus it follows from the expansion (A.3.1), and the exponential series $e^z = \sum_{n=0}^{\infty} \frac{z^n}{n!}$, that we have

$$\operatorname{Im} \{e^{z} E_{1}(z)\} = -\operatorname{Arg}(z) + \operatorname{Im} \{z \ln(z)\} + \dots$$
 (A.3.9)

The argument of z is $Arg(z) = \phi + \frac{\pi}{2}$, so using the above in the definition of $K(\phi, m)$, and noting $|z^2|/m = \sec^2 \phi$, the approximation is found.

5. As $m \to \infty$, $K(\phi, m) = -m^{-1} \sin 2\phi + 2m^{-2} \cos \phi \cos 3\phi + O(m^{-3})$. Plugging into the series expansion stated above for large Re(z), i.e. large $m \tan(\phi)$, we have

$$K(\phi, m) = 1 + \frac{1}{m} |z|^2 \operatorname{Im} \left\{ e^z E_1(z) \right\} = 1 + \frac{1}{m} \operatorname{Im} \left\{ \sum_{n=0}^{\infty} \frac{(-1)^n n! (z^*)^{n+1}}{|z|^{2n}} \right\}$$

$$= \sum_{n=1}^{\infty} \frac{(-1)^n n! \cos^{2n} \phi}{m^n} \operatorname{Im} \left\{ (\tan \phi + i)^{n+1} \right\}$$
(A.3.10)

For $\phi \in (-\pi/2, \pi/2)$ we can use that $\tan(\phi) + i = \sec \phi e^{i\theta}$ to show

$$\cos^{n+1} \phi \operatorname{Im} \left\{ (\tan \phi + i)^{n+1} \right\} = \sin ((n+1)\pi/2) \cos ((n+1)\phi) \quad (A.3.11)$$
$$-\sin ((n+1)\phi) \cos ((n+1)\pi/2),$$

and depending on n odd or even one of the right-hand side terms vanishes. Substituting (A.3.11) into (A.3.10), a little work finds that $K(m, \phi)$ can be written

$$K(\phi, m) = \sum_{n=1}^{\infty} a_n m^{-n}, \quad a_n = \begin{cases} (-1)^{\frac{n+1}{2}} n! \sin((n+1)\phi) \cos^{n-1}(\phi), & \text{if } n \text{ odd,} \\ (-1)^{\frac{n}{2}+1} n! \cos((n+1)\phi) \cos^{n-1}(\phi), & \text{if } n \text{ even.} \end{cases}$$
(A.3.12)

The first few terms are

$$K(\phi, m) = -\frac{1}{m}\sin(2\phi) + \frac{2}{m^2}\cos(\phi)\cos(3\phi) + \frac{6}{m^3}\cos^2(\phi)\sin(4\phi) + \dots$$
(A.3.13)

Hence,

$$\langle u'v'\rangle = -\frac{\varepsilon}{2\mu} \int_0^{\pi} \rho^{\varepsilon}(\phi) \left(\sin(2\phi) - \frac{2}{m} \cos(\phi) \cos(3\phi) + O(m^{-2}) \right) d\phi. \tag{A.3.14}$$

The series is valid for $m \tan(\phi) \to \infty$. This covers the case $m \to \infty$ with ϕ fixed, but also the limit $\phi \to \pm \pi/2$ for any $m \neq 0$, concluding

$$\lim_{\phi \to \pm \pi/2} K(\phi, m) = 0. \tag{A.3.15}$$

i.e. property (ii).

6.
$$\int_{-\pi/2}^{\pi/2} K(\phi, m) d\phi = 0$$
.

By definition of $K(\phi, m)$ (equation 4.2.11),

$$\int_{-\pi/2}^{\pi/2} K(\phi, m) \, d\phi = \pi + \frac{1}{m} \int_{-\pi/2}^{\pi/2} |z|^2 \operatorname{Im} \left\{ e^z E_1(z) \right\} d\phi, \tag{A.3.16}$$

Changing the variable of integration to z (using $dz/d\phi = -|z|^2/m$) and using

parts

$$\int_{-\pi/2}^{\pi/2} K(\phi, m) d\phi = \pi - \lim_{t \to \infty} \operatorname{Im} \left\{ \int_{mt+im}^{-mt+im} e^{z} E_{1}(z) dz \right\},$$

$$= \pi - \lim_{t \to \infty} \operatorname{Im} \left\{ \left[e^{z} E_{1}(z) + \ln(z) \right] \Big|_{mt+im}^{-mt+im} \right\},$$
(A.3.17)

$$= \pi - \lim_{t \to \infty} \text{Im} \left\{ \left[e^z E_1(z) + \ln(z) \right] \Big|_{mt + im}^{-mt + im} \right\}, \quad (A.3.18)$$

$$=0, (A.3.19)$$

where the $e^z E_1(z)$ terms in the final evaluation are vanishingly small (by expansion A.3.2) and the logarithms cancel the π .

Appendix B

Colophon

This document was set in the Times Roman typeface using LATEX and BibTEX, composed with TEXShop. All numerical experiments were performed in MATLAB, as well as almost all of the figures.

Bibliography

- [1] A. N. Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Proc. USSR Acad. Sci.*, 30:301–305, 1941.
- [2] L. F. Richardson. *Weather prediction by numerical process*. Cambridge university press, 2007.
- [3] R. H. Kraichnan. Inertial ranges in two-dimensional turbulence. *The Physics of Fluids*, 10(7):1417–1423, 1967.
- [4] G. K. Batchelor. Computation of the energy spectrum in homogeneous two-dimensional turbulence. *The Physics of Fluids*, 12(12):II–233, 1969.
- [5] B. Galperin and P. L. Read. Zonal Jets: Phenomenology, Genesis, and Physics. Cambridge University Press, 2019.
- [6] Y. Kaspi, E. Galanti, W. B. Hubbard, D. J. Stevenson, S. J. Bolton, L. Less, T. Guillot, J. Bloxham, J. E. P. Connerney, H. Cao, et al. Jupiter's atmospheric jet streams extend thousands of kilometres deep. *Nature*, 555(7695):223–226, 2018.
- [7] P. Rhines. Waves and turbulence on the beta plane. *J. Fluid Mech.*, 69:417–443, 1975.
- [8] G. K. Vallis and M. E. Maltrud. Generation of mean flows and jets on a beta plane and over topography. *J. Phys. Oceanogr.*, 23(7):1346–1362, 1993.

- [9] N. A. Phillips. A simple three-dimensional model for the study of large-scale extratropical flow patterns. *J. Meteorol.*, 8:381–394, 1951.
- [10] G. K. Vallis. *Atmospheric and Oceanic Fluid Dynamics*. Cambridge University Press, Cambridge, 2006.
- [11] T DelSole. Can quasigeostrophic turbulence be modeled stochastically? *J. Atmos. Sci.*, 53(11):1617–1633, 1996.
- [12] B. Galperin, S. Sukoriansky, N. Dikovskaya, P. L. Read, Y. H. Yamazaki, and R. Wordsworth. Anisotropic turbulence and zonal jets in rotating flows with a beta-effect. *Nonlinear Process. Geophys.*, 13(1):83–98, 2006.
- [13] S. Sukoriansky, N. Dikovskaya, and B. Galperin. On the arrest of inverse energy cascade and the rhines scale. *J. Atmos. Sci.*, 64:3312–3327, 2007.
- [14] B. Galperin, S. Sukoriansky, and N. Dikovskaya. Zonostrophic turbulence. *Physica Scripta*, 2008(T132):014034, 2008.
- [15] R. K. Scott and D. G. Dritschel. The structure of zonal jets in geostrophic turbulence. *J. Fluid Mech.*, 711:576–598, 2012.
- [16] K. Srinivasan and W. R. Young. Zonostrophic Instability. *J. Atmos. Sci.*, 69(5):1633–1656, 2012.
- [17] F. Bouchet, C. Nardini, and T. Tangarife. Kinetic theory of jet dynamics in the stochastic barotropic and 2D Navier-Stokes equations. *J. Statist. Phys.*, 153(4):572–625, 2013.
- [18] F. Bouchet, J. Rolland, and E. Simonnet. Rare event algorithm links transitions in turbulent flows with activated nucleations. *Phys. Rev. Lett.*, 122(7):074502, 2019.
- [19] F. Bouchet and E. Simonnet. Random changes of flow topology in two-dimensional and geophysical turbulence. *Phys. Rev. Lett.*, 102(9):094504, 2009.

- [20] R. K. Scott and A.-S. Tissier. The generation of zonal jets by large-scale mixing. *Physics of Fluids*, 24(12):126601, 2012.
- [21] H. Kuo. Dynamic instability of two-dimensional nondivergent flow in a barotropic atmosphere. *J. Atmos. Sci.*, 6(2):105–122, 1949.
- [22] N. C. Constantinou, B. F. Farrell, and P. J. Ioannou. Emergence and equilibration of jets in beta-plane turbulence: applications of stochastic structural stability theory. *Journal of the Atmospheric Sciences*, 71(5):1818–1842, 2014.
- [23] E. Woillez and F. Bouchet. Barotropic theory for the velocity profile of jupiter's turbulent jets: an example for an exact turbulent closure. *J. Fluid Mech.*, 860:577–607, 2019.
- [24] D. G. Dritschel and M. E. McIntyre. Multiple jets as PV staircases: the Phillips effect and the resilience of eddy transport barriers. *J. Atmos. Sci.*, 65:855–874, 2008.
- [25] S. Lee and I. M. Held. Subcritical instability and hysteresis in a two-layer model. *J. Atmos. Sci.*, 48(8):1071–1077, 1991.
- [26] S. Lee. Maintenance of Multiple Jets in a Baroclinic Flow. *J. Atmos. Sci.*, 54(13):1726–1738, 1997.
- [27] N. A. Phillips. Energy transformations and meridional circulations associated with simple baroclinic waves in a two-level quasi-geostrophic model. *Tellus*, 6:273–286, 1954.
- [28] S. B. Feldstein and I. M. Held. Barotropic decay of baroclinic waves in a two-layer beta-plane model. *J. Atmos. Sci.*, 46:3416–3430, 1989.
- [29] J. Pedlosky. Geophysical Fluid Dynamics. Springer-Verlag, New York, 1987.
- [30] M. J. Schmeits and H. A. Dijkstra. Bimodal behavior of the Kuroshio and the Gulf Stream. *Journal of physical oceanography*, 31(12):3435–3456, 2001.

- [31] M. Berhanu, R. Monchaux, S. Fauve, N. Mordant, F. Pétrélis, A. Chiffaudel, F. Daviaud, B. Dubrulle, L. Marié, F. Ravelet, et al. Magnetic field reversals in an experimental turbulent dynamo. *EPL (Europhysics Letters)*, 77(5):59001, 2007.
- [32] M. Berhanu, B. Gallet, R. Monchaux, M. Bourgoin, P. Odier, J. Pinton, N. Plihon, R. Volk, S. Fauve, N. Mordant, et al. Bistability between a stationary and an oscillatory dynamo in a turbulent flow of liquid sodium. *J. Fluid Mech.*, 641:217–226, 2009.
- [33] V. Dallas, K. Seshasayanan, and S. Fauve. Transitions between turbulent states in a two-dimensional shear flow. *Physical Review Fluids*, 5(8):084610, 2020.
- [34] P. N. Loxley and B. T. Nadiga. Bistability and hysteresis of maximum-entropy states in decaying two-dimensional turbulence. *Physics of Fluids*, 25(1):015113, 2013.
- [35] F. Ravelet, L. Marié, A. Chiffaudel, and F. Daviaud. Multistability and memory effect in a highly turbulent flow: Experimental evidence for a global bifurcation. *Phys. Rev. Lett.*, 93(16):164501, 2004.
- [36] D. S. Zimmerman, S. A. Triana, and D. P. Lathrop. Bi-stability in turbulent, rotating spherical Couette flow. *Physics of Fluids*, 23(6):065104, 2011.
- [37] G. S. Huisman, R. C. Van Der Veen Ca, C. Sun, and D. Lohse. Multiple states in highly turbulent Taylor–Couette flow. *Nature communications*, 5(1):1–5, 2014.
- [38] P. Wei, S. Weiss, and G. Ahlers. Multiple transitions in rotating turbulent Rayleigh-Bénard convection. *Phys. Rev. Lett.*, 114(11):114506, 2015.
- [39] O. Cadot, A. Evrard, and L. Pastur. Imperfect supercritical bifurcation in a three-dimensional turbulent wake. *Physical Review E*, 91(6):063005, 2015.
- [40] D. Barkley, B. Song, V. Mukund, G. Lemoult, M. Avila, and B. Hof. The rise of fully turbulent flow. *Nature*, 526(7574):550–553, 2015.

- [41] J. H. Rogers. *The giant planet Jupiter*, volume 6. Cambridge University Press, 1995.
- [42] H. Kahn and T. E. Harris. Estimation of particle transmission by random sampling. *National Bureau of Standards applied mathematics series*, 12:27–30, 1951.
- [43] M. N. Rosenbluth and A. W. Rosenbluth. Monte Carlo calculation of the average extension of molecular chains. *The Journal of Chemical Physics*, 23(2):356–359, 1955.
- [44] I. Teo, c. g. Mayne, K. Schulten, and T. Lelièvre. Adaptive multilevel splitting method for molecular dynamics calculation of benzamidine-trypsin dissociation time. *Journal of chemical theory and computation*, 12(6):2983–2989, 2016.
- [45] H. A. P. Blom, G. J. Bakker, and J. Krystul. Rare Event Estimation for a Large-Scale Stochastic Hybrid System with Air Traffic Application., 2009.
- [46] F. Cérou and A. Guyader. Adaptive multilevel splitting for rare event analysis. *Stochastic Analysis and Applications*, 25(2):417–443, 2007.
- [47] F. Cérou, B. Delyon, A. Guyader, and M. Rousset. On the asymptotic normality of adaptive multilevel splitting. *SIAM/ASA Journal on Uncertainty Quantification*, 7(1):1–30, 2019.
- [48] F. Cérou, A. Guyader, T. Lelievre, and D. Pommier. A multiple replica approach to simulate reactive trajectories. *The Journal of chemical physics*, 134(5):054108, 2011.
- [49] J. Rolland and E. Simonnet. Statistical behaviour of adaptive multilevel splitting algorithms in simple models. *Journal of Computational Physics*, 283:541–558, 2015.
- [50] F. Cérou, A. Guyader, and M. Rousset. Adaptive multilevel splitting: Historical perspective and recent results. *Chaos*, 29(4):043108, 2019.

- [51] F. Ragone, J. Wouters, and F. Bouchet. Computation of extreme heat waves in climate models using a large deviation algorithm. *Proceedings of the National Academy of Sciences*, 115(1):24–29, 2018.
- [52] F. Ragone and F. Bouchet. Rare event algorithm study of extreme warm summers and heatwaves over Europe. *Geophysical Research Letters*, 48(12):e2020GL091197, 2021.
- [53] E. Simonnet, J. Rolland, and F. Bouchet. Multistability and rare spontaneous transitions in barotropic β -plane turbulence. *J. Atmos. Sci.*, 78(6):1889–1911, 2021.
- [54] C. Bréhier, M. Gazeau, L. Goudenège, T. Lelièvre, M. Rousset, et al. Unbiasedness of some generalized adaptive multilevel splitting algorithms. *The Annals of Applied Probability*, 26(6):3559–3601, 2016.
- [55] E. Simonnet. Combinatorial analysis of the adaptive last particle method. *Statistics and Computing*, 26(1-2):211–230, 2016.
- [56] J. Wouters and F. Bouchet. Rare event computation in deterministic chaotic systems using genealogical particle analysis. *J. Phys. A*, 49(37):374002, 2016.
- [57] J. Rolland. Collapse of transitional wall turbulence captured using a rare events algorithm. *J. Fluid Mech.*, 931, 2022.
- [58] F. Bouchet, J. B. Marston, and T. Tangarife. Fluctuations and large deviations of Reynolds stresses in zonal jet dynamics. *Phys. Fluids*, 30(1):015110, 2018.
- [59] E. Hopf. Statistical hydromechanics and functional calculus. *Arch. Ration. Mech. Anal.*, 1:87–123, 1952.
- [60] U. Frisch. *Turbulence: the legacy of AN Kolmogorov*. Cambridge university press, 1995.
- [61] B. F. Farrell and P. J. Ioannou. Structural stability of turbulent jets. *J. Atmos. Sci.*, 60(17):2101–2118, 2003.

- [62] J. B. Marston, E. Conover, and T. Schneider. Statistics of an unstable barotropic jet from a cumulant expansion. *J. Atmos. Sci.*, 65(6):1955–1966, 2008.
- [63] B. F. Farrell and P. J. Ioannou. A Theory of Baroclinic Turbulence. *J. Atmos. Sci.*, 66:2445–2454, 2009.
- [64] J. B. Marston. Statistics of the general circulation from cumulant expansions. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(4):041107, 2010.
- [65] N. A. Bakas and P. J. Ioannou. Is spontaneous generation of coherent baroclinic flows possible? *J. Fluid Mech.*, 862:889–923, 2019.
- [66] K. Li, J. B. Marston, S. Saxena, and S. M. Tobias. Direct statistical simulation of the Lorenz63 system. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 32(4):043111, 2022.
- [67] S. M. Tobias and J. B. Marston. Direct statistical simulation of jets and vortices in 2D flows. *Physics of Fluids*, 29(11):111111, 2017.
- [68] A. Child, R. Hollerbach, B. Marston, and S. Tobias. Generalised quasilinear approximation of the helical magnetorotational instability. *Journal of Plasma Physics*, 82(3):905820302, 2016.
- [69] A. Plummer, J. B. Marston, and S. M. Tobias. Joint instability and abrupt nonlinear transitions in a differentially rotating plasma. *Journal of Plasma Physics*, 85(1), 2019.
- [70] F. Bouchet, C. Nardini, and T. Tangarife. Stochastic averaging, large deviations and random transitions for the dynamics of 2D and geostrophic turbulent vortices. *Fluid Dyn. Res.*, 46(6):061416, 2014.
- [71] S. M. Tobias and J. B. Marston. Direct statistical simulation of out-of-equilibrium jets. *Phys. Rev. Lett.*, 110(10):104502, 2013.
- [72] F. Ait-Chaalal, T. Schneider, B. Meyer, and J. B. Marston. Cumulant expansions for atmospheric flows. *New Journal of Physics*, 18(2):025019, 2016.

- [73] B. F. Farrell and P. J. Ioannou. Structure and spacing of jets in barotropic turbulence. *J. Atmos. Sci.*, 64(10):3652–3665, 2007.
- [74] J.B. Marston, W. Qi, and S. M. Tobias. *Direct Statistical Simulation of a Jet*, page 332–346. Cambridge University Press, 2019.
- [75] J. B. Marston, G. P. Chini, and S. M. Tobias. Generalized quasilinear approximation: application to zonal jets. *Phys. Rev. Lett.*, 116(21):214501, 2016.
- [76] S. M. Tobias, J. S. Oishi, and J. B. Marston. Generalized quasilinear approximation of the interaction of convection and mean flows in a thermal annulus. *Proceedings of the Royal Society A*, 474(2219):20180422, 2018.
- [77] A. Allawala, S. M. Tobias, and J. B. Marston. Dimensional reduction of direct statistical simulation. *J. Fluid Mech.*, 898, 2020.
- [78] G. K. Batchelor. *Introduction to fluid dynamics*. Cambridge University Press, Cambridge, 1967.
- [79] L. Cope. *The Dynamics of Geophysical and Astrophysical Turbulence*. PhD thesis, University of Cambridge, 2020.
- [80] K. Srinivasan and W. R. Young. Reynolds stress and eddy diffusivity of β -plane shear flows. *J. Atmos. Sci.*, 71(6):2169–2185, 2014.
- [81] J. Laurie, G. Boffetta, G. Falkovich, I. Kolokolov, and V. Lebedev. Universal profile of the vortex condensate in two-dimensional turbulence. *Phys. Rev. Lett.*, 113(25):254503, 2014.
- [82] N. A. Bakas and J. P. Ioannou. A theory for the emergence of coherent structures in beta-plane turbulence. *J. Fluid Mech.*, 740:312–341, 2014.
- [83] W. M'F. Orr. The stability or instability of the steady motions of a perfect liquid and of a viscous liquid. Part II: A viscous liquid. In *Proceedings of the Royal Irish Academy. Section A: Mathematical and Physical Sciences*, volume 27, pages 69–138. JSTOR, 1907.

- [84] B. Farrell. Developing disturbances in shear. *J. Atmos. Sci.*, 44(16):2191–2199, 1987.
- [85] T. G. Shepherd. Time development of small disturbances to plane Couette flow. *J. Atmos. Sci.*, 42(17):1868–1872, 1985.
- [86] A. Frishman. The culmination of an inverse cascade: mean flow and fluctuations. *Physics of Fluids*, 29(12):125102, 2017.
- [87] S. M. Tobias, K. Dagon, and J. B. Marston. Astrophysical fluid dynamics via direct statistical simulation. *The Astrophysical Journal*, 727(2):127, 2011.
- [88] J. Liu and T. Schneider. Mechanisms of jet formation on the giant planets. *J. Atmos. Sci.*, 67:3652 –3672, 2010.
- [89] R. M. B. Young, P. L. Read, and Y. Wang. Simulating Jupiter's weather layer. Part I: Jet spin-up in a dry atmosphere. *Icarus*, 326:225–252, 2019.
- [90] P. L. Read, D. Kennedy, N. Lewis, H. Scolan, F. Tabataba-Vakili, Y. Wang, S. Wright, and R. M. B. Young. Baroclinic and barotropic instabilities in planetary atmospheres: energetics, equilibration and adjustment. *Nonlin. Proc. Geophys.*, 27(2):147–173, 2020.
- [91] S. Nadarajah and T. K. Pogány. On the distribution of the product of correlated normal random variables. *Comptes Rendus Mathematique*, 354(2):201–204, 2016.
- [92] A. Frishman and C. Herbert. Turbulence statistics in a two-dimensional vortex condensate. *Physical review letters*, 120(20):204505, 2018.
- [93] R. K. Scott and L. M. Polvani. Equatorial superrotation in shallow atmospheres. *Geophys. Res. Lett.*, 35:L24202, 2008.
- [94] S. I. Thomson and M. E. McIntyre. Jupiter's unearthly jets: a new turbulent model exhibiting statistical steadiness without large-scale dissipation. *J. Atmos. Sci.*, 73(3):1119–1141, 2016.

- [95] M. Abramowitz and I. Stegun. *Handbook of mathematical functions*. Dover, New York, 1972.
- [96] N. Bleistein and R. A. Handelsman. *Asymptotic expansions of integrals*. Ardent Media, 1975.