

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Wave Resonances in Rotating Shear Flows: Weakly-Nonlinear Theories and Numerical Studies

Permalink

<https://escholarship.org/uc/item/4dh5g860>

ISBN

9798189271373

Author

Wang, Jinge

Publication Date

2026-05-01

Peer reviewed|Thesis/dissertation

Wave Resonances in Rotating Shear Flows: Weakly-Nonlinear Theories and Numerical
Studies

by

Jinge Wang

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

Doctor of Philosophy

in

Mechanical Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Philip S. Marcus, Chair
Professor Jon A. Wilkening
Professor Omer Savas

Spring 2026

Wave Resonances in Rotating Shear Flows: Weakly-Nonlinear Theories and Numerical
Studies

Copyright 2026
by
Jinge Wang

Abstract

Wave Resonances in Rotating Shear Flows: Weakly-Nonlinear Theories and Numerical Studies

by

Jinge Wang

Doctor of Philosophy in Mechanical Engineering

University of California, Berkeley

Professor Philip S. Marcus, Chair

Rotating shear flows are fundamental to the dynamics of numerous natural and engineered systems, from protoplanetary accretion disks to aircraft wake vortices. This dissertation investigates the multi-faceted dynamics of hydrodynamic stability in these systems, bridging the fundamental theory of wave-wave resonances in incompressible environments with the global wave-mean flow interactions inherent to rotating, stratified fluids.

Part I investigates the weakly nonlinear stability of incompressible columnar vortices, demonstrating that the triadic resonance of wave modes is governed by a set of hydrodynamic selection rules. Employing a multi-scale perturbation analysis, we prove that resonant interactions between smooth neutral modes are strictly conservative and confined to the Manley—Rowe relations. Using wave pseudoenergy within a large- k WKB framework, we show that the selection rules topologically prohibit intrinsic instability. Consequently, the breakdown of a columnar vortex requires a specific symmetry-breaking mechanism. We identify and analyze two distinct pathways for this destabilization: (1) *Parametric instability*, where external forcing sustains a pumping wave and acts as an energy source. Following the pseudoenergy criterion, we identify the discrete critical layer mode as the sole neutral mode that carries negative pseudoenergy, and thus required for the onset of parametric instability. Using non-degenerate perturbation theory, we generalize classical elliptical instability to arbitrary driving frequencies and geometries, and identify continuous instability configurations involving discrete critical layer modes. (2) *Active critical layers*, where an embedded singularity breaks the Hermitian symmetry of the linear operator, enabling the extraction of mean-flow energy via a wave-mean resonance. Most importantly, Part I establishes that the critical layer is the fundamental engine of instability in columnar vortices, and that the presence of a critical layer is essential to the breakdown of an isolated vortex.

Part II transitions to the global wave-mean resonances of stratified fluids, focusing on the

Zombie Vortex Instability (ZVI). To overcome the numerical stiffness imposed by restorative inertio-gravity waves and rapid advective shear in the three-dimensional cylindrical Boussinesq equations, we develop a parallelized Exponential Time Differencing (ETD) pseudo-spectral solver. By analytically diagonalizing the linear operators, the ETD scheme enables efficient simulations in global cylindrical domains, while natively embedding the physical resonance characteristics into the discrete integration operators. Utilizing these global simulations, we address the geometric limitations of traditional local shearing box approximations, demonstrating that the uniform spacing of zombie vortices observed in prior literature is a result of the uniform background shear assumed within the local models. Instead, we propose the frequency-resetting chain reaction hypothesis: self-replication is driven by localized wave-mean flow resonance, where inertio-gravity wave packets are barotropically emitted, trapped at baroclinic critical layers, and undergo nonlinear roll-up to re-emit new waves at reset frequencies dictated strictly by the local background shear.

Ultimately, this dissertation provides a unified physical framework for wave resonances in rotating shear flows through a synthesis of weakly-nonlinear theories and global numerical studies. Across both incompressible and stratified regimes, the critical layer emerges as the fundamental engine of flow destabilization. Whether acting as the requisite symmetry-breaking mechanism to extract energy in isolated columnar vortices, or functioning dually as the barotropic amplifier and baroclinic spatial trap that drives multi-generational vortex replication, critical layer dynamics dictate the transfer of energy. Together, these theoretical and computational contributions elucidate the complex mechanisms of momentum transport and energy exchange in rotating shear flows, linking the fundamental fluid dynamics of engineered aircraft wakes to the macroscopic behavior of astrophysical accretion disks.

To my parents, Dr. Peijun Wang and Ms. Wenhong Li, for the roots that grounded me.
And to Xiaoyu, for everything, and for every minute.

Contents

Contents	ii
List of Figures	iv
List of Tables	ix
1 Introduction	1
1.1 Rotating Shear Flows	1
1.2 Thesis Scope and Outline	2
I Incompressible Flows: Generic Wave-Wave Resonances and the Critical Layer Imperative	4
2 Stability of a Columnar Vortex	5
2.1 Background	5
2.2 Hydrodynamic Selection Rules and the Optical Analogy	6
2.3 Overview	7
3 General Formulation of Triadic Resonances	9
3.1 Base Flow and Disturbance Equations	9
3.2 Multiple-Scale Perturbation Analysis	11
3.3 Conservative Triadic Resonance	15
4 Selection Rules for Conservative Triadic Resonances	21
4.1 Linear Wave Modes	21
4.2 Conservative Interactions between Smooth Neutral Modes	24
4.3 Pseudoenergy and Selection Rules for Conservative Triads	25
Appendix 4A Linear System in a Constant Rotating Frame	29
5 Generalized Parametric Instability of Columnar Vortices	31
5.1 Historical Context	31
5.2 Resonance Tuning via Non-Degenerate Perturbation Theory	35

5.3	Generalized Parametric Instability of the Lamb–Oseen Vortex	37
	Appendix 5A Spectral Construction of Strain Modes	40
6	Roles of the Critical Layer in the Instability of Columnar Vortices	43
6.1	Critical Layers in Incompressible Flows	43
6.2	The Roles of Stratification in Vortex Stability and Resonance	45
6.3	Conclusion	47
II	Stratified Flows: Global Wave-Mean Resonances and Numerical Simulations	49
7	Stability of Rotating, Stratified, Shear Flows	50
7.1	Motivation and Background	50
7.2	The Local Shearing Box vs. Global Cylindrical Frame	51
7.3	Overview	53
8	Mathematical Formulations	54
8.1	The Boussinesq Equations for Rotating Stratified Shear Flows	55
8.2	Conservation Laws	57
8.3	Linear Stability Equations and Critical Layer Singularities	61
8.4	Poloidal-Toroidal Decomposition	68
	Appendix 8A Inertial Gravity Waves in Static, Stratified Fluids	70
9	A Pseudo-Spectral ETD Scheme for Rotating Stratified Flows	72
9.1	Spectral Expansion and the Pseudo-Spectral Method	73
9.2	Stiffness Associated with the Explicit Treatment of the Background Terms .	76
9.3	The Exponential Time Differencing (ETD) Method	78
9.4	Numerical Implementation	85
9.5	Code Validation	93
	Appendix 9A Calculation of Energy and Momentum with Spectral Decomposition	101
	Appendix 9B Mathematical Identities in Spectral Basis	105
10	Baroclinic Critical Layers and the Zombie Vortex Instability	107
10.1	The Illusion of the Zero-Frequency Mode in Shearing Coordinates	107
10.2	The Frequency-Resetting Chain Reaction Hypothesis	109
10.3	Deploying Global Nonlinear Simulations	110
10.4	Conclusions and Future Work	114
	Appendix 10A Initialization of a Quasi-Equilibrium Gaussian Vortex Perturbation	115
	Bibliography	118

List of Figures

- 3.1 Stability mechanism via mode degeneracy in the inviscid limit. In the absence of degeneracy, discrete neutral modes (represented by the black circles) are topologically constrained to the imaginary axis (neutral stability) and are only permitted to move vertically along it as parameters vary. When a pair of modes becomes degenerate or nearly degenerate (represented by the white concentric circles), a resonant interaction can unfold in two distinct ways: either via detuning, where the modes separate along the imaginary axis, or via instability, where the modes are pulled symmetrically off the axis into the complex plane, resulting into a conjugate pair of growing and decaying modes. 13
- 3.2 Example numerical solutions of the normalized conservative three-wave amplitude equations (3.30a,b,c) for distinct interaction coefficient sign signatures (s_0, s_1, s_2). (a–c) Temporal evolution of the normalized wave intensities $|\hat{A}_j|^2$. (d–f) Corresponding 3D phase space trajectories plotting $|\hat{A}_{1,2}|$ against $|\hat{A}_0|$ and the relative phase factor $\cos(\theta_2 - \theta_1 - \theta_0)$. White and black markers indicate the initial and final states of the evolution, respectively. Left and Right Columns: Bounded resonance regimes, where the system is strictly confined to closed periodic orbits on a compact manifold. Middle Column: The explosive resonance regime ($s_0 = s_1 = -s_2$), where the relative phase locks to a constant value while permitting the system to follow an open trajectory with simultaneous amplitude divergence. See also Craik (1986). 17
- 3.3 Effect of detuning $\Delta\omega$ on parametric instability as given in (3.39). (a) Doppler-shifted eigenfrequencies $\text{Im}(\sigma_{\text{par},j}) - k_j\omega_0/k_0$ for modes $j = 1, 2$. The solid and dashed lines represent the upper and lower branches of the four coupled eigenvalues, respectively. Inside the instability band (shaded region), the upper and lower branches of the same mode locks to the same wave frequency, and the eigenvalues merge toward exact resonance where the modes synchronize to a common Doppler-shifted frequency. (b) Real growth rates $\text{Re}(\sigma)$. In the region where frequencies are locked, the growth rates bifurcate to form a symmetric instability bubble (shaded region indicating positive growth). The two wave modes share identical growth rates, leading to the appearance of a single curve for the unstable branch. Outside this region, the frequency mismatch becomes significant, and the coupled modes are no longer in near resonance. 19

- 4.1 Characteristics of the inviscid linear modes of the Lamb–Oseen vortex obtained numerically by MLEGS (Sangjoon Lee and J. Wang, 2025). (a) Dispersion curves $\omega(k)$ of helical modes ($m = 1$). Solid lines denote Kelvin waves, and dashed lines indicate neutral discrete modes possessing critical layers. The shaded frequency band $\omega_{\text{CL}} = [-m, 0)$ is associated with the critical layer singularity and corresponds to the continuous spectrum. (b–d) Radial velocity component $r\tilde{u}_r(r)$ of selected eigenvectors ($m = 1, k = 3.0$): (b) A regular discrete Kelvin wave, exhibiting a purely real, smooth profile characteristic of global neutral modes; (c) A generic singular critical layer mode from the continuous spectrum. The singularity at the critical radius r_c (vertical dashed line) breaks analytic continuity, resulting in spectral broadening and a distinct phase jump in physical space; (d) A “passive” discrete critical layer mode. Unlike generic critical layer modes, this solution maintains high-order spectral convergence, confirming that the critical layer singularity is effectively suppressed and decoupled from the core dynamics. The eigenvalues associated with the mode profiles shown are marked as circles in the dispersion plot. 23
- 5.1 Graphical representations of resonant triads in the Lamb–Oseen vortex. (a) Resonance sustained by a stationary pumping wave with zero axial wavenumber ($m_p, k_p = 0, \omega_p = 0$). Star symbols mark crossing points between distinct dispersion branches ($k_i = k_j, \omega_i = \omega_j$) in resonance with $|m_p| = |m_i - m_j|$; filled circles indicate intersections with the zero-frequency axis, corresponding to a pair of conjugate modes ($\pm m$) in resonance with a static strain field with symmetry $m_p = 2m$. (b) General triadic resonance satisfying the frequency matching condition $\omega_1 + \omega_2 = \omega_3$ alongside wavenumber matching. The specific examples represent resonance between modes (m, k, ω) : $(1, 5.6, -1.2) + (2, 0.50, -2.0) = (3, 6.0, -3.3)$ (dash-dotted lines with star markers), and $(1, 4.1, -1.2) + (2, 0.98, -2.1) = (3, 5.1, -3.3)$ (dashed lines with circular markers). 34
- 5.2 Iterations of the resonance tuning process for a triad driven by the pump mode $(m_0, k_0, \omega_0) = (0, 2.00, 0.5334)$ of the Lamb–Oseen vortex. Note that iterations 3–6 are overlapping as the frequency mismatch approaches machine precision during the tuning. Solid lines are the dispersion curves of the two free modes, $\omega_0 + \omega_1(k_1)$ and $\omega_2(k_1 + k_0)$. They are provided and shifted here for visual clarity so that their intersection corresponds to perfect resonance. Dashed lines trace the trajectories of the perturbation-based Newton–Raphson algorithm as it projects from one wavenumber approximation $k_1^{(n)}$ to the next $k_1^{(n+1)}$ based on the local group velocity mismatch given by (5.10). The inset (lower right) quantifies the performance, showing the frequency mismatch $|\Delta\omega|$ reducing to a precision of $\mathcal{O}(10^{-1})$ during the first iteration, and to $\mathcal{O}(10^{-14})$ after 6 iterations. 36

- 5.3 (a) The numerically computed profile $g(r) = \tilde{u}_r/r$ for the steady-state strain response of the Lamb–Oseen vortex, converging correctly to 1 in the far field. (b) Overlay of our constructed solution (solid line) against the historical results of C. Eloy and S. Le Dizès (1999) (dashed line), demonstrating perfect mathematical agreement when applying a specific radial scaling factor of $x = 2r$ 37
- 5.4 Resonance loci in the (k_0, k_1) plane for triads involving a helical ($m_1 = 1$) mode and a pumping azimuthal wavenumber of (a) $m_0 = 1$; (b) $m_0 = 2$; (c) $m_0 = 0$. Solid line segments indicate regions capable of triggering parametric instability based on the pseudoenergy criterion, while the gray dashed continuations represent the remaining parts of the resonance loci where parametric instability is suppressed as the discrete critical layer modes transferred to regular Kelvin waves along the dispersion branch. Only loci containing at least one instability segment are plotted. The search is limited to the neutral branches as resolved for each azimuthal wavenumber, except for $m_0 = 0$ where we consider the first 6 radial branches due to the dense spectral spacing. 39
- 7.1 Mapping between the global cylindrical coordinates (r, θ, z) and the local shearing box $(\tilde{x}, \tilde{y}, \tilde{z})$ centered at (r_0, θ_0) . The local coordinates are defined as $\tilde{y} = r - r_0$, $\tilde{x} = -r_0(\theta - \theta_0)$, and $\tilde{z} = z$, such that the local unit vectors correspond to $\tilde{\mathbf{y}} = \hat{\mathbf{r}}$ and $\tilde{\mathbf{x}} = -\hat{\boldsymbol{\theta}}$. Gravity acts in the $-z$ direction (pointing into the page). The local frame co-rotates with the background at an angular velocity $\tilde{\Omega} = \bar{\Omega}(r_0) + \Omega$, where $\bar{\Omega}$ represents the radially dependent angular velocity of the global background flow and Ω is the constant background rotation rate of the global frame. The differential rotation of the background flow is locally approximated as a linear shear profile $\bar{u}_{\tilde{x}} = -\sigma_s \tilde{y}$, where the constant shear rate is defined as $\sigma_s = r_0(d\bar{\Omega}/dr)|_{r_0}$ ($\sigma_s < 0$ in the depicted example). 52
- 9.1 Detailed visualization of the two-stage global transposition encapsulated within the axial transform, converting between the azimuthal (PFP) and fully spectral (PFF) spaces on a $P = 3 \times 2$ MPI grid. **(a)** The first phase operates within `SUBCOMM_2`, exchanging sub-partitions of the complete m -axis to reconstruct complete z -columns in the intermediate (INT) space. **(b)** The second phase operates within `SUBCOMM_1`, exchanging sub-partitions of the complete k -axis to reconstruct complete r -columns in the PFF space. The 2D matrix panels illustrate the exact non-contiguous memory mappings encoded via `MPI_TYPE_CREATE_SUBARRAY`, enabling the `MPI_ALLTOALLW` collective to perform strided transpositions dynamically without explicit buffer packing. 89
- 9.2 Single-node strong scaling performance for a $256 \times 256 \times 128$ physical grid on dual-socket AMD EPYC 7763 processors (128 total cores). Legend labels denote the total step time, the pointwise nonlinear evaluation ($\mathcal{N}$), the ETD advancement (**E**, **NL**), the full space conversions ($\text{PPP} \leftrightarrow \text{FFF}$), the localized radial expansions, and the solenoidal projection ($\mathbb{P}$). The dashed line represents ideal linear scaling. 91

- 9.3 Discrete unstable modes for a stratified Lamb–Oseen vortex at $\bar{N} = 5$. (a) Growth rates λ and (b) wave frequencies ω as functions of the scaled axial wavenumber $k/\bar{N}$. Open circles represent the eigenvalues computed by the present global spectral eigenvalue solver. Three points on the first mode branch ($k/\bar{N} = 1.4, 2.4, 4.8$) are distinctly marked for further structural analysis. Solid black lines denote the benchmark results of S. Le Dizès (2008) obtained via a shooting method with complex contour deformation. The numerical spatial discretization employs $N_r = 252$ radial collocation points, a radial spectral truncation of $N = 244$, and a mapping parameter $L = 6.0$. Only the first three unstable branches are plotted. 95
- 9.4 Frequency maps (left column) and normalized spatial eigenmode profiles $\tilde{u}_\theta$ (right column) tracking the topological evolution of the first mode at the three distinct scaled wavenumbers, as marked in figure 9.3: (a) $k/\bar{N} = 1.4$, (b) $k/\bar{N} = 2.4$, and (c) $k/\bar{N} = 4.8$. The left panels plot the radial profiles of the local epicyclic frequencies (dashed magenta lines) and the co-rotation frequency (solid red line), alongside the mode’s constant wave frequency ω (horizontal green line). The intersections of the horizontal wave frequency line with these respective curves define the radial locations of the turning points and the barotropic critical layer. At $k/\bar{N} = 4.8$, the loss of the inner turning points exposes the mode directly to the barotropic critical layer, triggering the Gibbs-like spectral ringing observed in the spatial profile. 96
- 9.5 Buoyancy field (b') at r - θ slice slightly above the perturbation center plane and r - z slice at $\theta = 245^\circ$ taken at $t = 90$ with $\Delta t = 0.1$ for (a) the ETD scheme that linearizes $\mathbf{u}' \times \bar{\zeta} \hat{\mathbf{z}}$ plus background rotation and stratification, and (b) the full ETD scheme that linearizes both cross terms plus background rotation and stratification. The baseflow is a Lamb–Oseen vortex with a core size of 4, and the Gaussian perturbation is initiated at $(r = 5, \theta = 45^\circ, z = L_z/2)$ with an aspect ratio of 0.26. The flow is subject to $\Omega = 0$ and $\bar{N} = 0.1$. The color intensity corresponds to the magnitude of the field in log-scale. Standard AB-CN scheme develops the same catastrophic high-frequency numerical artifacts observed in (a) but at a much faster pace. 98
- 9.6 Temporal convergence of the ETD scheme. The fractional errors in global kinetic energy (E_K) and available potential energy (E_{AP}) follow the theoretical second-order scaling ($\propto \Delta t^2$) as Δt is reduced, reaching an optimal temporal resolution at $\Delta t \approx 10^{-3}$. Below this threshold, term cancellation in the ETD integration weights causes round-off error to dominate. Errors are evaluated at $t = 20$, corresponding to approximately 3.18 core rotation periods ($T_{\text{core}} = 2\pi$) of the background Lamb–Oseen vortex, ensuring the measurement captures fully developed nonlinear dynamics. 99

- 9.7 Energy balance tracking of the full ETD scheme during the nonlinear integration of the 3D Gaussian perturbation. To maintain inviscid stability and absorb the forward energy cascade, a sixth-order hyperviscosity filter is applied, with its amplitude dynamically adjusted based on the energy magnitude at the spectral tails. The figure comprises four panels: (top left) the temporal evolution of E_K , E_{AP} , and the cumulative shear production; (top right) the internal budget of E_K , explicitly incorporating the buoyancy energy exchange; (bottom left) the internal budget of E_{AP} , matching the buoyancy exchange; and (bottom right) the total mechanical energy balance. Note that the governing equations define the total energy evolution as $d_t(E_K + E_{AP}) = -\mathcal{E}_{\text{shear}}$; here, we denote $\text{PS} = -\mathcal{E}_{\text{shear}}$ 102
- 10.1 Snapshots of zombie vortices in the y - z plane at $x = 0$ marching outward from the initial non-equilibrium Gaussian disturbance field at the origin. Anticyclonic vertical vorticity ω_z is black and cyclonic is white. $N/f = 1$; $\sigma_s/f = -3/4$; and the unit of length in the figure is $4L_x/(3\pi)$, where L_x is the streamwise length of the computational domain, which is periodic in x . Reproduced from Philip S. Marcus, Pei, Jiang, and Hassanzadeh (2013). 108
- 10.2 (a) Mock-up zombie vortex lattices constructed using the frequency-hopping hypothesis. The initial disturbance is assumed to be a point at $y = 0.5$, $z = 0$. The coloring corresponds to WKBJ ray trajectories starting from the previous generation's advection frequency and terminating at the corresponding baroclinic critical layers. Due to mirroring symmetries, each zombie vortex emits exactly four ray paths. (b) The predicted zombie lattice locations. The vertical spacing between neighboring lattices is exactly $\bar{N}$, perfectly mirroring the nonlinear DNS results of Philip S. Marcus, Pei, Jiang, and Hassanzadeh (2013). 111
- 10.3 u'_z -field in r - z slice taken at $t = 20000$ with ray tracing overlays. The background flow is a scaled Lamb–Oseen vortex with an intensity of 4, a core radius of 8, and a linear stratification of $\bar{N} = 0.07$. The initial quasi-equilibrium Gaussian vortex is geometrically centered at $r_0 = 12$. The ray tracings are performed with the wave frequency matching the background flow speed at (a) $r_0 = 12.3$, (b) $r_0 = 11.0$, and (c) $r_0 = 13.0$ for azimuthal wavenumbers up to $m = 10$. Trajectories align strictly with the effective advection frequency at $r_0 = 12.3$ 112

List of Tables

5.1	Summary of resonant triads examined in the literature. A blank axial flow entry indicates a vortex with no axial velocity component. In all listed cases, the perturbation modes have no axial dependency ($\bar{k} = 0$) and the base wave eigenvalues are either zero or assumed zero ($\bar{\sigma} = 0$), except for the Gaussian vortex case. Note that the radial perturbation at the bottom falls into the $(\bar{m}, \bar{k}) = (0, 0)$ self-interaction case that is not considered in this study.	33
5.2	Comparison of the axial wavenumbers (k) and the parametric growth rates computed for the elliptical instability of the Lamb–Oseen vortex between a pair of helical ($m = \pm 1$) wave modes with equal k . The numerical values of the growth rate in this work are calculated based on our generalized interaction coefficients using MLEGS. Note that C. Eloy and S. Le Dizès (1999) only computed the parametric growth rate for the first two degenerate pairs, and their results are scaled by a factor of 2 here due to the radial scaling $x = 2r$ used in their specific pump mode formulation.	38
9.1	Local array shape on each MPI rank for a process grid $N_1 \times N_2$. Distributed dimensions are indicated by the divisor; the undivided dimension is the pencil (fully local). The intermediate (INT) space is a temporary staging area during the axial transform, where the z -pencil is formed before the subsequent radial or azimuthal reorientation.	88
9.2	Key spectral operators, their mathematical forms, and banded-matrix widths. μ is the mapped finite radial coordinate, and c is a constant coefficient.	90

Acknowledgments

This thesis would not exist without the dedication of my silicon-based helpers — an unquantifiable number of CPUs pushed to their limits across the now-retired Stampede2 supercomputer at the Texas Advanced Computing Center (TACC) at The University of Texas at Austin, and the Anvil supercomputer at Purdue University, all generously provided by the NSF ACCESS (formerly XSEDE) program (ID: PHY220056). On a more personal hardware note, I am deeply grateful to the Student Technology Equity Program (STEP) at UC Berkeley, which provided massive support during the most financially stressful days of the global pandemic.

I would like to express my deepest gratitude to my advisor, Prof. Philip S. Marcus. Long before coming to Berkeley, while I was an undergraduate trained as a pure fluid experimentalist, I read about Phil’s classical work on Taylor-Couette flows. Back then, I had heard many anecdotes about Phil, but never in my life did I expect to be advised by him, nor transition into theoretical and numerical work. However, encountering the limitations and the lack of theoretical guidance in my experimental research sparked my deep belief in the necessity of theoretical and numerical studies. That realization prompted a massive shift in my trajectory, and I was very fortunate to have Phil as my advisor during this transition. It feels incredibly fitting that my thesis today focuses on rotating shear flows, echoing the very work of Phil’s that I read about all those years ago. I am incredibly thankful for the level of uncompromising rigor Phil brought out in me. More importantly, he has been a constant inspiration through his immense resilience and his profound love for science — or, more particularly, the Great Red Spot. After all, he quite literally put his new heart into it.

I am also immensely grateful to my committee members for their guidance. Thank you to Prof. Ömer Savaş, who taught me graduate-level fluid mechanics and vortex dynamics, laying the critical foundations of my thesis work. He also fostered in me the invaluable habit of pulling out a classical paper every now and then to perform a complete re-derivation. Thank you to Prof. Jon Wilkening, whose teachings on numerical methods for differential equations deeply benefited the development of my own numerical schemes. During the global pandemic, when everyone was separated into their own isolated bubbles, Prof. Wilkening’s lectures brought a much-needed sense of regularity into a chaotic world. I also want to thank Prof. Alexis Kaminski, who sat on my qualifying committee. Her meticulous care for her students and her lectures, down to every single detail, is truly inspirational. She has been my role model for how to deliver research talks and what it means to be a great scholar.

I must also extend a special thanks to Prof. Douglas Kelley at the University of Rochester. As my undergraduate research advisor, he first led me into the world of fluid dynamics research. His continuous encouragement and support, even throughout my long journey here at Berkeley, have meant a great deal to me.

To my lab mates: Chiyu, Aidi, Haris, Sangjoon (Joon), Haley, and Sungkyu. Too many circumstances prevented us from bonding as much as we otherwise would have, but it is because of all of you that I survived all the hard times. I want to extend a specific thank you to Joon. During the dark days of the pandemic and the incredibly difficult period

that followed, we banded together to tackle our inherited codebase and spent countless hours modernizing a tangled mix of Fortran 77 and 95 into the now-open-sourced MLEGS codebase, a massive undertaking that we navigated entirely on our own. Doing it alongside you made an otherwise arduous process manageable.

I would also like to thank my professional colleagues at Apple: Ta, Jiaqi, Brian, Frank, Qilin, and Huiyang. My time working with you taught me an immense amount about building PDE solvers in an industrial setting. It provided the perfect environment to apply the computing skills I gained during my dissertation, further solidifying my expertise in high-performance computation (where a great number of GPUs on top of the CPUs were drafted into heavy service). You also provided invaluable lessons in productivity, efficiency, and code maintenance. It was deeply inspiring to see how my background in fluid dynamics translated so seamlessly to other physics disciplines, which gave me a much-needed boost of confidence as a researcher.

A massive thank you to all my friends at the Cal Sailing Club. Without you, I would have never completed this long journey, and because of you, this journey took longer than it probably should have. May the wind be with you all.

I owe a profound debt of gratitude to my parents. To my mom, who is now happily enjoying her retirement and traveling around my home country (and the world), something she didn't get to do before I left for college. Thank you for organizing the summer camps that sparked my life-long passion as a sailor and windsurfer, which ultimately led to my fascination with fluid dynamics. To my dad, whose journey from a farming village to a lifelong career in higher education has always inspired me. You have been a profound teacher not only to the countless students you have guided, but to me personally. You have always been endlessly supportive and remain my ultimate role model.

To my partner, Xiaoyu. Thank you for your patience, your trust, and for keeping me fed, caffeinated, and sane, especially while living together in a rented basement for a large portion of my PhD. Despite not drinking coffee yourself, and not knowing how to brew it, you managed to make the most "refreshing" pot of "coffee" to keep me awake on the day of my preliminary exam. Thank you for giving me your unconditional support even after hearing me say that I was "almost done" for the hundredth time, and for miraculously never getting too mad at me when I would disappear to go windsurfing with my CSC folks. I also want to thank you for picking up my habit of staring at aircraft contrails — getting just as excited about them as I do, and never caring if the people around us wonder why two Asian kids are just staring up at the sky for no obvious reason. I could not have done this without you, and I am so excited to start the next phase of our life journey together.

Finally, I want to thank myself for remaining resilient, for never giving up, and for stubbornly insisting on bringing this research to its rightful closure, with the indispensable help, of course, of all the carbon-based life and silicon-based cyber-life mentioned above. And to you, the reader — whether a friend I have known for years or someone I may never meet — thank you for sharing in this journey.

Chapter 1

Introduction

1.1 Rotating Shear Flows

Rotating shear flows are fundamental to the dynamics of countless natural and engineered systems. Their ubiquity spans vast scales, from the astrophysical accretion disks that surround protostars, to atmospheric structures like Jupiter’s Great Red Spot and terrestrial tornadoes, down to the hazardous aircraft wake vortices (AWV) generated by commercial aviation. Understanding the stability mechanisms of these columnar and disk-like structures is paramount across aerospace engineering, fluid mechanics, and astrophysics.

Historically, these flows have been analyzed via linear stability theory, dating back to Lord Kelvin (1880) and his seminal analysis of the Rankine vortex. While this framework successfully classifies the discrete spectrum of regular modes (Kelvin waves) and accurately predicts the initial growth of disturbances in flows subject to centrifugal or Rayleigh–Taylor instabilities, it faces significant limitations. Specifically, in the idealized incompressible and inviscid limit, linear theory strictly dictates neutral stability for monotonic vortex profiles typical of AWV, as well as for Keplerian velocity profiles in accretion disks.

Since linear theory in this incompressible limit implies neutral stability, the destabilization of these coherent structures requires the consideration of finite-amplitude, nonlinear interactions or external forcing. In this context, the classical elliptical instability serves as the archetype (Bayly, 1986; Pierrehumbert, 1986; Waleffe, 1990). Formulated for actual columnar vortices using normal modes by Moore et al. (1975) and Tsai and Widnall (1976), this instability demonstrates that an external strain field triggers a resonant interaction between two neutrally stable Kelvin waves. Subsequent extensive numerical studies by C. Eloy and S. Le Dizès (1999) and others confirmed these analytical growth rates. Yet, this mechanism fundamentally requires a sustained external strain field — a condition typically only satisfied during the late stages of vortex evolution when paired structures, such as trailing aircraft wakes, begin to heavily interact. During its initial, isolated phase, the incompressible vortex remains profoundly robust. Consequently, neither the linear framework nor the externally forced elliptical instability provides a pathway to actively

destabilize an isolated wake vortex.

In addition, the assumption of strict incompressibility is a theoretical idealization. Physical systems, ranging from astrophysical disks to aircraft wakes laden with intense engine exhaust, feature complex thermal conditions and density stratification. The introduction of stratification fundamentally alters the stability landscape, introducing new restorative forces that can circumvent classical incompressible stability bounds. Particularly, the combined action of background rotation, density stratification, and differential azimuthal shear can trigger complex wave-mean flow interactions and profound instabilities. A primary example is the Zombie Vortex Instability (ZVI) (Philip S. Marcus, Pei, Jiang, Barranco, et al., 2015; Philip S. Marcus, Pei, Jiang, and Hassanzadeh, 2013), which is an intrinsic, finite-amplitude instability that thrives in the stably stratified “dead zones” of protoplanetary disks, autonomously generating new generations of “zombie” vortices from a localized initial perturbation.

Understanding the precise physical mechanisms behind the ZVI not only has profound implications for planetary formation but also provides vital insights into the original engineering motivation: the mitigation of AWV. While the ZVI has been extensively documented via numerical simulations utilizing the local Cartesian shearing box approximation, its full topological evolution in a global, unbounded domain remains largely unexplored. Capturing this global evolution is essential to uncovering its true generative mechanism. If the wave-mean resonant interactions induced by stratification drive this intrinsic instability in disks, the implications for aerodynamic wake management are direct. Specifically, it implies that the radial stratification naturally occurring in AWVs is not a negligible thermal artifact, but rather a potential hydrodynamic vulnerability. Actively engineering or adjusting this radial stratification could potentially introduce similar finite-amplitude instabilities, offering a novel theoretical pathway to accelerate vortex breakdown during the robust isolated phase.

1.2 Thesis Scope and Outline

This dissertation investigates the stability of rotating shear flows by bridging fundamental weakly-nonlinear theories in incompressible environments with advanced global numerical simulations in stratified environments. The work is divided into two distinct parts.

Part I focuses on the incompressible limit, generic wave-wave resonances, and the critical layer imperative. We begin by developing a general multiple-scale perturbation theory to isolate resonant triads in columnar vortices. Through this framework, we establish strict hydrodynamic selection rules governed by wave pseudoenergy, proving mathematically that pure wave-wave triads among smooth neutral modes are strictly conservative and incapable of explosive growth. To overcome this fundamental limitation, we identify two distinct symmetry-breaking pathways. The first is a generalized parametric instability driven by sustained external forcing. Expanding upon specific configurations studied in classical literature, we generalize this mechanism to arbitrary frequencies and configurations. We achieve this by introducing non-degenerate perturbation theory to numerically tune for

arbitrary triads. Critically, our pseudoenergy criterion dictates that this forced triad must contain a negative pseudoenergy mode, which for a columnar vortex must take the form of a discrete critical layer mode. The second pathway relies on an *active* critical layer to break the system’s Hermitian symmetry and extract energy directly from the mean flow. For this mechanism, we qualitatively demonstrate the fundamental necessity of active critical layers to achieve intrinsic destabilization.

Part II introduces vertical density stratification, transitioning the focus to global wave-mean resonances and advanced numerical simulations. We expose the geometric limitations of traditional local Cartesian shearing box approximations when studying stratified dynamics, motivating a necessary transition to a global cylindrical domain. To overcome the severe numerical stiffness imposed by restorative inertio-gravity waves and advective shear in the 3D Boussinesq equations, we develop a highly specialized, parallelized pseudo-spectral Exponential Time Differencing (ETD) scheme. Deploying this computational framework for global nonlinear simulations, we culminate with the frequency-resetting chain reaction hypothesis — driven by localized barotropic emission and baroclinic trapping — to comprehensively explain the macroscopic evolution of the Zombie Vortex Instability.

Part I

Incompressible Flows: Generic Wave-Wave Resonances and the Critical Layer Imperative

Chapter 2

Stability of a Columnar Vortex

2.1 Background

To rigorously investigate the intrinsic stability of an isolated columnar vortex, Part I isolates the dynamics within the idealized incompressible, inviscid limit. As outlined in Chapter 1, the eigenmodes of a strictly monotonic vortex profile in this limit are entirely neutral, with eigenvalues confined to the imaginary axis. Transitioning to instability therefore requires a pair of these neutral eigenvalues to collide, creating a geometric spectral *degeneracy* that allows them to bifurcate into growing and decaying modes while preserving the underlying temporal symmetry of the Euler equations.

Unfolding this spectral degeneracy necessitates transcending the linear approximation to consider the effects of nonlinear wave interactions (Schmid et al., 2001). In the context of weakly nonlinear theory, significant attention has historically been focused on vortices subjected to sustained external deformation (see Kerswell, 2002). Classical examples include the elliptical instability driven by an external strain field (C. Eloy and S. Le Dizès, 1999; Feys and S. A. Maslowe, 2016; Kerswell, 1999, 2002; L. Lacaze, Birbaud, et al., 2005; L. Lacaze, Ryan, et al., 2007; McEwan, 1970; Moore et al., 1975; Mora et al., 2021; Tsai and Widnall, 1976), the precessional instability inside a fluid cylinder (Albrecht et al., 2015, 2018; Kobine, 1995; Lagrange, C. Eloy, et al., 2008; Lagrange, Meunier, et al., 2011; Lopez and Marqués, 2018; Mahalov, 1993; Manasseh, 1992, 1994), and the curvature instability arising from the bending of the vortex filament (Blanco-Rodríguez and S. Le Dizès, 2017; Fukumoto and Hattori, 2005). Mathematically, these phenomena constitute a subset of parametric instabilities where two disturbance waves couple with a third, externally imposed pumping mode.

While these studies establish a robust framework for forced vortices, there remains a critical gap in understanding the stability of strictly *isolated* vortices, where no such sustained external deformation exists. This gap is particularly relevant to the mitigation of aircraft wake vortices. These long-lived, counter-rotating vortices are generated at the wingtips and pose a significant hazard to following aircraft (Hallock and Holzäpfel, 2018). The

lifecycle of an AWW system typically involves an initial phase where the two vortices are sufficiently separated to be treated as isolated. Over time, viscous diffusion expands the vortex cores until they begin to impose a mutual strain effect on each other. This mutual interaction triggers mechanisms such as the long-wave Crow instability (Crow, 1970) or the short-wave elliptical instability (Leweke and Williamson, 1998; Moore et al., 1975; Tsai and Widnall, 1976), eventually leading to reconnection and breakdown due to the cancellation of circulation. The engineering challenge lies in the fact that the initial isolated phase is remarkably robust. To accelerate the decay of the wake, it is highly desirable to identify intrinsic weakly-nonlinear mechanisms capable of increasing the growth of disturbances entirely through internal wave coupling during the isolated phase, thereby allowing the vortices to reach sizes sufficient to trigger mutual interactions earlier.

The simplest intrinsic nonlinear mechanism capable of achieving this is *triadic resonance*. This universal phenomenon, occurring in diverse dispersive media including plasma physics, nonlinear optics, and fluid mechanics (Craig, 1986; Weiland and Wilhelmsson, 1977), allows for the secular exchange of energy between three wave modes. In many dynamic systems, these resonant interactions obey a set of conservation laws known as the Manley–Rowe relations (Manley and Rowe, 1956), which constrain the exchange of wave action within each resonant triad. As discussed by Dodin et al. (2008), these relations can be derived deductively to offer a compact form of integrals for dynamical systems with resonances. Crucially, depending on the signs of the interaction coefficients (or pseudoenergies), the solution can yield either a bounded periodic exchange of wave activity or an unbounded simultaneous growth to infinity. The latter solution is known as *explosive instability* (see Becker and Grimshaw, 1993; Craig and Adam, 1979), representing a highly theorized pathway toward accelerating vortex breakdown strictly through internal, unforced wave interactions.

2.2 Hydrodynamic Selection Rules and the Optical Analogy

The physics of wave resonance in vortical flows can be elucidated by drawing a theoretical connection to nonlinear optics. Both fluid vortices and optical media support dispersive waves that can interact nonlinearly. In optical three-wave mixing, energy transfer is efficient only when the phase matching condition ($\Delta \mathbf{k} = 0$) is satisfied (see R. W. Boyd, 2008). Similarly, hydrodynamic resonance requires the geometric matching of spatial wavenumbers (k, m for the axial and azimuthal wavenumbers) and temporal frequencies (ω) among the triad members.

The fundamental distinction between the two systems lies in the energetic status of the background medium. In standard nonlinear optics involving a passive medium (such as a transparent crystal), the total energy of the interacting waves is strictly conserved. Under these conditions, the Manley–Rowe relations constrain the system to a bounded manifold, resulting in an oscillatory exchange between frequencies (see Armstrong et al., 1962). In

contrast, the rotating shear of the background fluid constitutes a vast, active energy reservoir. Superficially, this resembles an optical “active medium” (e.g., a laser gain medium) or a plasma, where the medium can theoretically pump energy into the waves to drive explosive growth.

A central result of Part I, however, is that the availability of this reservoir does not guarantee instability. We establish a framework of hydrodynamic “selection rules” to govern the triadic interactions. Drawing a parallel to quantum mechanics, where selection rules dictate allowable atomic transitions, we demonstrate that for *smooth neutral modes*, the Hamiltonian structure of the incompressible Euler equations imposes strict constraints that confine the interaction completely to the bounded solution branch, thereby prohibiting pure wave-wave triads from achieving sustained growth.

Consequently, the intrinsic breakdown of a columnar vortex requires a symmetry-breaking mechanism to either access the explosive solution branch or violate the conservative constraints entirely. We identify two such pathways. The first is *external forcing*, which functions analogously to an optical parametric oscillator where an external pump wave actively injects energy into the system. This triggers a generalized parametric instability. Unlike classical studies typically limited to resonant configurations dictated by geometric constraints, our framework admits arbitrary driving frequencies and imposes no restrictions on the pumping wavenumbers. The second is the *critical layer*, a flow singularity where the phase velocity of the wave matches the local flow velocity (Fabre et al., 2006; Stéphane Le Dizès and Laurent Lacaze, 2005). The critical layer functions similarly to optical resonant layers found in plasma absorption (Stix, 1992), introducing non-Hermitian effects that allow the wave to extract energy from the background shear. This analogy frames our investigation: intrinsic wave-wave resonance among regular modes is energetically constrained and bounded (passive), whereas instability requires a mechanism to tap into energy reservoirs, either explicitly via external forcing or implicitly via critical-layer singularity.

2.3 Overview

The objective of Part I of this dissertation is to perform a multi-scale perturbation analysis of triadic resonance in generic columnar vortices and to establish the selection rules that distinguish between bounded and explosive solutions. Part I is organized as follows. In Chapter 3, we formulate the disturbance equations using a poloidal-toroidal decomposition and derive the general amplitude equations for triadic resonance. We analyze the specific case of conservative resonance, showing that it obeys the Manley–Rowe relations and theoretically admits either bounded oscillatory exchange or simultaneous explosive growth. In Chapter 4, we apply this framework to columnar vortices. We prove that triadic resonance among smooth neutral modes is always conservative. Using wave pseudoenergy within a large- k WKBJ framework, we then demonstrate that these conservative triads are strictly confined to the bounded solution branch, effectively prohibiting intrinsic explosive instability for smooth modes. Subsequently, in Chapter 5, we relax the isolation constraint and investigate

the generalized parametric instability, employing non-degenerate perturbation theory to tune for arbitrary resonant triads. Finally, in Chapter 6, we discuss the role of critical layer singularities as a necessary pathway to break the selection rules, introducing non-Hermitian effects that allow the wave to access the background shear energy.

Chapter 3

General Formulation of Triadic Resonances

3.1 Base Flow and Disturbance Equations

The base flow of interest is an isolated axisymmetric columnar vortex. In a cylindrical coordinate system (r, θ, z) where the z -axis aligns with the vortex centerline, its velocity profile has the generic form

$$\bar{\mathbf{u}} = \bar{u}_\theta(r)\hat{\boldsymbol{\theta}} + \bar{u}_z(r)\hat{\mathbf{z}}, \quad (3.1)$$

where $\bar{u}_\theta$ and $\bar{u}_z$ are the azimuthal and axial velocity components, and the unit vectors along r , θ and z directions are denoted as $\hat{\mathbf{r}}$, $\hat{\boldsymbol{\theta}}$ and $\hat{\mathbf{z}}$. Accordingly, the angular velocity $\bar{\Omega}(r)$ and axial vorticity $\bar{\zeta}(r)$ are expressed as

$$\bar{\Omega}(r) = \frac{\bar{u}_\theta}{r}, \quad \bar{\zeta}(r) = \frac{1}{r} \frac{d(r\bar{u}_\theta)}{dr}. \quad (3.2a,b)$$

Here, we assume that any provided physical quantities are non-dimensionalized with respect to a velocity scale V_* and a length scale r_* , with the premise that they are well-defined. For instance, in the case of the Rankine vortex, r_* may be taken as the radius of the vorticity patch, and V_* as the maximum azimuthal velocity. Utilizing these characteristic scales, the Reynolds number of the flow is defined as $Re \equiv r_* V_*/\nu$, where ν is the kinematic viscosity of the fluid.

In this study, we will assume that the Reynolds number is sufficiently high ($Re \gg 1$), as is typical for aeronautical applications (L. Lacaze, Ryan, et al., 2007), so viscous diffusion is neglected. The columnar vortex, as defined by (3.1), is an equilibrium solution to the incompressible Euler equations regardless of the profiles of $\bar{u}_\theta(r)$ and $\bar{u}_z(r)$. The multiple-scale perturbation analysis presented in the subsequent sections will be developed for arbitrary $\bar{\mathbf{u}}$ profiles. However, for illustrative purposes and numerical computations, we will refer to the specific case of the Lamb–Oseen vortex, given by

$$\bar{\mathbf{u}} = \frac{1 - e^{-r^2}}{r} \hat{\boldsymbol{\theta}}. \quad (3.3)$$

We consider a small but finite-amplitude disturbance velocity field, $\mathbf{u}$, superposed on the base flow (3.1). The equations of motion for the total velocity field, $\mathbf{v} = \bar{\mathbf{u}} + \mathbf{u}$, are the incompressible Euler equations, here written in the rotation form:

$$\nabla \cdot \mathbf{v} = 0, \quad (3.4a)$$

$$\frac{\partial \mathbf{v}}{\partial t} = \mathbf{v} \times (\nabla \times \mathbf{v}) - \nabla \varphi. \quad (3.4b)$$

where $\varphi = (\mathbf{v} \cdot \mathbf{v})/2 + p/\rho$ with p being the associated pressure field and ρ being the density of the fluid. We require the disturbance $\mathbf{u}$ to be analytic at $r = 0$ and to decay algebraically or exponentially as $r \rightarrow \infty$ to maintain finite kinetic energy. Detailed discussions on the physical and accurate attainment of these conditions can be found in S. Lee and P. S. Marcus (2023, 2025) and Matsushima and P. Marcus (1997).

The disturbance equations are obtained by cancelling out the equilibrium of $\bar{\mathbf{u}}$ from (3.4b), which results into a system of four equations with four primitive variables: three components of $\mathbf{u}$, and φ . Instead of dealing with the primitive form, we note that the incompressibility condition (3.4a) reduces to $\nabla \cdot \mathbf{u} = 0$, so the disturbance velocity field can be decomposed into poloidal and toroidal components:

$$\mathbf{u} = \nabla \times (\psi \hat{\mathbf{z}}) + \nabla \times (\nabla \times (\chi \hat{\mathbf{z}})), \quad (3.5)$$

where ψ and χ are the toroidal and poloidal streamfunctions. To avoid gauge arbitrariness in (3.5), we require

$$\lim_{r \rightarrow \infty} \psi = \lim_{r \rightarrow \infty} \chi = 0, \quad (3.6)$$

and define the following linear, invertible poloidal-toroidal projection operator:

$$\mathbb{P}[\mathbf{u}] = \mathbf{U} \equiv (\psi, \chi)^\top, \quad (3.7)$$

where $\mathbf{U}$ will be referred to as the poloidal-toroidal vector henceforth. The equations of motion for the disturbance can then be expressed in terms of the disturbance vector $\mathbf{U}$:

$$\frac{\partial \mathbf{U}}{\partial t} = \mathbb{N}[\bar{\mathbf{U}}, \mathbf{U}] + \frac{1}{2} \mathbb{N}[\mathbf{U}, \mathbf{U}], \quad (3.8)$$

where

$$\bar{\mathbf{U}} \equiv \mathbb{P}[\bar{\mathbf{u}}], \quad (3.9a)$$

$$\mathbb{N}[\mathbf{U}_1, \mathbf{U}_2] \equiv \mathbb{P}[\mathbf{u}_1 \times (\nabla \times \mathbf{u}_2) - (\nabla \times \mathbf{u}_1) \times \mathbf{u}_2], \quad (3.9b)$$

for arbitrary $\mathbf{U}_j = \mathbb{P}[\mathbf{u}_j]$ ($j = 1, 2$). As the poloidal-toroidal projection of any conservative vector field (i.e., gradient of a scalar potential) is identically zero, (3.8) eliminates the pressure term entirely. Accordingly, the disturbance dynamics are cast in a compact operator form in terms of the poloidal-toroidal vector, while ensuring the velocity field remains solenoidal by construction.

The objective of this study is to investigate the weakly nonlinear stability of the columnar vortex (3.1) with a focus on the triadic interactions between the linear wave modes. Our analysis is based on a multi-scale expansion of the disturbance equations (3.8), which allows us to derive amplitude equations governing the evolution of the wave modes and to identify the conditions under which these interactions lead to instability.

3.2 Multiple-Scale Perturbation Analysis

The analysis commences by constructing an asymptotic expansion of the disturbance vector $\mathbf{U}$. For a small disturbance, since the nonlinearity in (3.8) arises solely from the quadratic term $\mathbb{N}[\mathbf{U}, \mathbf{U}]/2$, it can be expanded in powers of a scale parameter, $\epsilon \ll 1$:

$$\mathbf{U} = \sum_{n=1}^{\infty} \epsilon^n \mathbf{U}^{(n)}(r, \theta, z, t), \quad (3.10)$$

where nonlinear interactions manifest at $\mathcal{O}(\epsilon^2)$ or higher, and the leading-order component $\mathbf{U}^{(1)}$ resembles linear waves. Because the base flow $\bar{\mathbf{u}}$ has only radial variation, the linear solutions of (3.8), denoted $\mathbf{R}$, can be expressed as Fourier series along the azimuthal and axial directions:

$$\mathbf{R}(r, \theta, z, t) = \tilde{\mathbf{R}}(r) e^{i(m\theta + kz) + \sigma t}, \quad (3.11)$$

where m and k are the azimuthal and axial wavenumbers, and σ and $\tilde{\mathbf{R}}$ are the corresponding complex wave frequencies and wave functions respectively. To distinguish between the linear evolution and the inherent oscillatory nature of each linear wave, the complex wave frequency is separated into the real growth rate λ and the oscillation frequency ω :

$$\sigma \equiv \lambda + i\omega. \quad (3.12)$$

Consider a “slow” timescale that is much longer than the oscillation periods of the waves:

$$\tau \equiv \epsilon t, \quad (3.13)$$

$\mathbf{U}^{(1)}$ can be constructed as a superposition of linear wave modes:

$$\mathbf{U}^{(1)} \equiv \sum_j A_j(\tau) \tilde{\mathbf{R}}_j(r) \exp [i(m_j\theta + k_jz + \omega_jt)] + \text{c.c.}, \quad (3.14)$$

where the complex conjugate (of all precedent terms), denoted as c.c., is added to maintain a physical disturbance.

In this multi-scale formulation, the modulations of the wave amplitudes A_j not only reflect the linear growth or decay (λ) of individual modes but also account for their nonlinear interactions at $\mathcal{O}(\epsilon^2)$. In the following analysis, the amplitude equations governing the evolution of A_j ’s will be derived, and their analytical solutions will be discussed.

The linear waves

The linear wave modes are determined by the leading-order problem. Substituting (3.11) into (3.8) and retaining only $\mathcal{O}(\epsilon)$ -terms yields an eigenvalue problem for the linear operator $\mathbb{M}_{mk}$, defined as:

$$\sigma \tilde{\mathbf{R}} = \mathbb{N}_{mk}[\bar{\mathbf{U}}, \tilde{\mathbf{R}}] \equiv \mathbb{M}_{mk} \tilde{\mathbf{R}}, \quad (3.15)$$

where σ and $\tilde{\mathbf{R}}$ are the corresponding eigenvalue and eigenvector, and the subscripts denote the wavenumbers associated with the linear operator. Similarly, the linear matrix operator $\mathbb{M}_{mk}$ admits a left-hand eigenvector $\tilde{\mathbf{L}}_j^{mk}$ that satisfies

$$\sigma_j^{mk} (\tilde{\mathbf{L}}_j^{mk})^H = (\tilde{\mathbf{L}}_j^{mk})^H \mathbb{M}_{mk}, \quad (3.16)$$

which can be normalized to attain biorthogonality with the right-hand eigenvector:

$$\langle \tilde{\mathbf{L}}_j^{mk} | \tilde{\mathbf{R}}_l^{mk} \rangle = \delta_{jl}, \quad (3.17)$$

where the inner product is defined as $\langle \mathbf{X} | \mathbf{Y} \rangle \equiv \int_0^\infty (\mathbf{X} \cdot \mathbf{Y}) r dr$.

It is useful to highlight an important invariance property of the eigenvalue problem (3.15). In particular, while the physical stability properties of the vortex are intrinsic, the wave frequency ω (the imaginary part of the eigenvalue σ) is frame-dependent. In a reference frame translating with constant velocity U_∞ along the z -axis, the transformed linear eigenvalue problem reads:

$$\begin{aligned} \sigma' \tilde{\mathbf{R}}' &= \mathbb{M}_{mk} \tilde{\mathbf{R}}' + \mathbb{P}_{mk} \left[-U_\infty \hat{\mathbf{z}} \times (\nabla \times \mathbb{P}^{-1}[\mathbf{R}']) \right] \\ &= \mathbb{M}_{mk} \tilde{\mathbf{R}}' + ikU_\infty \tilde{\mathbf{R}}', \end{aligned} \quad (3.18)$$

where σ' and $\tilde{\mathbf{R}}'$ are the eigenvalue and eigenvector in the moving frame. This implies a Doppler shift in the wave frequency while the eigenvector structure remains invariant:

$$\omega' = \omega + kU_\infty, \quad \tilde{\mathbf{R}}' = \tilde{\mathbf{R}}. \quad (3.19)$$

While the spectral properties of (3.15) depend on the specific base flow profile, the time-reversibility of the underlying Euler equations implies that any eigenvalue σ with a non-zero real part must have a counterpart with the opposite sign of the real growth rate. Therefore, for an isolated vortex that models aircraft wake vortices or similar long-lived structures, all eigenvalues shall be close to purely imaginary, corresponding to neutral or near-neutral stability. If we restrict our attention to strictly neutral modes with $\lambda = 0$, it is evident that the transition to instability has a specific spectral requirement: a pair of neutral eigenvalues must move towards each other and collide along the imaginary axis, creating a geometric *degeneracy* before they can be “pulled” off the axis into the unstable and decaying half-planes respectively. This mechanism, illustrated conceptually in figure 3.1, acts as a necessary pathway towards instability. In the framework of weakly nonlinear theory, a simplest and most fundamental realization of this mechanism is *triadic resonance*, which provides the nonlinear coupling required to unfold the degeneracy and potentially trigger instabilities.




Figure 3.1: Stability mechanism via mode degeneracy in the inviscid limit. In the absence of degeneracy, discrete neutral modes (represented by the black circles) are topologically constrained to the imaginary axis (neutral stability) and are only permitted to move vertically along it as parameters vary. When a pair of modes becomes degenerate or nearly degenerate (represented by the white concentric circles), a resonant interaction can unfold in two distinct ways: either via detuning, where the modes separate along the imaginary axis, or via instability, where the modes are pulled symmetrically off the axis into the complex plane, resulting into a conjugate pair of growing and decaying modes.

Conditions for triadic resonance

When a disturbance field comprises a superposition of multiple linear eigenmodes, any pair of waves, denoted by indices p and q , interacts via the quadratic nonlinearity $\mathbb{N}[\mathbf{U}, \mathbf{U}]$ of (3.8) and generates a reaction term characterized by the sum of their phases:

$$\exp \left[i \left((m_p + m_q) \theta + (k_p + k_q) z + (\omega_p + \omega_q) t \right) \right].$$

If this nonlinear product effectively projects onto a third eigenmode, a resonant triad is established. This resonant coupling facilitates efficient energy exchange between the modes, potentially destabilizing a flow that is linearly stable to infinitesimal disturbance. Without loss of generality, we designate the triad members as $j = 0, 1$, and 2 , satisfying the resonant conditions described below.

The kinematic conditions required for sustained resonant interaction are:

$$m_0 + m_1 = m_2, \quad k_0 + k_1 = k_2, \quad \omega_2 - (\omega_0 + \omega_1) = \epsilon \Delta \omega, \quad (3.20)$$

where the frequency detuning is assumed to be of $\mathcal{O}(\epsilon)$ or higher to prevent the interaction from averaging out over the timescales of amplitude evolution. Although the linear growth rates λ_j do not explicitly appear in (3.20), modes with significant damping ($|\lambda_j| \gg \epsilon$) effectively decouple from the triad. Meanwhile, modes that are already linearly unstable are

of lesser interest in this weakly nonlinear context. In the rest of this section, we shall assume that all triad members are effectively neutral in the linear regime, with $|\lambda_j| \lesssim \mathcal{O}(\epsilon)$.

Equation (3.20) can be understood as a degeneracy pulling mechanism. Specifically, consider a Galilean transformation to a reference frame moving at the same axial phase velocity as the primary wave (mode 0), $U_\infty = -\omega_0/k_0$. In this frame, the primary wave appears as a stationary deformation ($\omega'_0 = 0$). The frequency matching requirement in (3.20) reduces to *degeneracy* or *nearly-degeneracy* of the other two members:

$$\omega'_2 \approx \omega'_1. \quad (3.21)$$

Physically, the deformation induced by mode 0 breaks the symmetry of the background vortex and couples the previously orthogonal modes 1 and 2. This coupling allows the degenerate pair to interact, potentially causing their eigenvalues to pull off the neutral axis and leading to instability.

Generalized three-wave amplitude equations

To derive the equations governing the temporal evolution of the first-order wave amplitudes, we first approximate the second-order disturbance using the linear modes:

$$\mathbf{U}^{(2)} = \sum_{m,k} \sum_j B_j^{mk}(\tau) \tilde{\mathbf{R}}_j^{mk} \exp [i(m\theta + kz + \omega_j^{mk}t)] + \text{c.c.}, \quad (3.22)$$

where slow modulations of B_j^{mk} 's are introduced following the method of multiple scales.

Substituting the expansions for $\mathbf{U}^{(1)}$ and $\mathbf{U}^{(2)}$ into (3.8) and imposing the solvability condition at $\mathcal{O}(\epsilon^2)$ yields the evolution equations for the triad amplitudes. Assuming $(m_j, k_j) \neq (0, 0)$ to exclude self-interaction cases and retaining terms up to $\mathcal{O}(\epsilon^2)$, we obtain:

$$\frac{d}{d\tau} A_0 \tilde{\mathbf{R}}_0 - \frac{\lambda_0}{\epsilon} A_0 \tilde{\mathbf{R}}_0 - \sum_j \lambda_j^{m_0 k_0} B_j^{m_0 k_0} \tilde{\mathbf{R}}_j^{m_0 k_0} = \mathbb{N}_{m_0 k_0}[\tilde{\mathbf{R}}_1^*, \tilde{\mathbf{R}}_2] A_1^* A_2 e^{i\Delta\omega\tau}, \quad (3.23a)$$

$$\frac{d}{d\tau} A_1 \tilde{\mathbf{R}}_1 - \frac{\lambda_1}{\epsilon} A_1 \tilde{\mathbf{R}}_1 - \sum_j \lambda_j^{m_1 k_1} B_j^{m_1 k_1} \tilde{\mathbf{R}}_j^{m_1 k_1} = \mathbb{N}_{m_1 k_1}[\tilde{\mathbf{R}}_0^*, \tilde{\mathbf{R}}_2] A_0^* A_2 e^{i\Delta\omega\tau}, \quad (3.23b)$$

$$\frac{d}{d\tau} A_2 \tilde{\mathbf{R}}_2 - \frac{\lambda_2}{\epsilon} A_2 \tilde{\mathbf{R}}_2 - \underbrace{\sum_j \lambda_j^{m_2 k_2} B_j^{m_2 k_2} \tilde{\mathbf{R}}_j^{m_2 k_2}}_{\sim \mathcal{O}(\epsilon) \text{ or higher}} = \mathbb{N}_{m_2 k_2}[\tilde{\mathbf{R}}_0, \tilde{\mathbf{R}}_1] A_0 A_1 e^{-i\Delta\omega\tau}. \quad (3.23c)$$

Left-multiplying (3.23) with the adjoint eigenvectors $\tilde{\mathbf{L}}_j$ and applying the biorthogonality condition (3.17), we arrive at the canonical amplitude equations for an arbitrary resonant

triad:

$$\frac{d}{d\tau}A_0 - \frac{\lambda_0}{\epsilon}A_0 = J_0A_1^*A_2e^{i\Delta\omega\tau}, \quad (3.24a)$$

$$\frac{d}{d\tau}A_1 - \frac{\lambda_1}{\epsilon}A_1 = J_1A_0^*A_2e^{i\Delta\omega\tau}, \quad (3.24b)$$

$$\frac{d}{d\tau}A_2 - \frac{\lambda_2}{\epsilon}A_2 = J_2A_0A_1e^{-i\Delta\omega\tau}, \quad (3.24c)$$

with the nonlinear interaction coefficients defined as follows:

$$J_0 \equiv \left\langle \tilde{\mathbf{L}}_0 \left| \mathbb{N}_{m_0k_0}[\tilde{\mathbf{R}}_1^*, \tilde{\mathbf{R}}_2] \right. \right\rangle, \quad (3.25a)$$

$$J_1 \equiv \left\langle \tilde{\mathbf{L}}_1 \left| \mathbb{N}_{m_1k_1}[\tilde{\mathbf{R}}_0^*, \tilde{\mathbf{R}}_2] \right. \right\rangle, \quad (3.25b)$$

$$J_2 \equiv \left\langle \tilde{\mathbf{L}}_2 \left| \mathbb{N}_{m_2k_2}[\tilde{\mathbf{R}}_0, \tilde{\mathbf{R}}_1] \right. \right\rangle. \quad (3.25c)$$

When the wave functions of a resonant triad are known, its interaction coefficients can be calculated from (3.25), and the numerical solutions of (3.24) can be easily computed. Analytical approaches, on the other hand, have thus far yielded only a limited number of qualitative findings for the general three-wave resonance, with no universally acknowledged general analytic solutions (see, for example, Craik, 1986). In the following analysis, we will consider the case of conservative interactions where analytical insights of (3.24) can be obtained.

3.3 Conservative Triadic Resonance

Conditions for conservative interactions

Consider the scenario where a resonant triad satisfies the following hypothesis:

Hypothesis 1 *All three interaction coefficients of the resonant triad are purely imaginary: $J_0, J_1, J_2 \in i\mathbb{R}$.*

In the ideal limit of exact resonance ($\Delta\omega = 0$) between neutral modes ($\lambda = 0$), (3.24) simplify to

$$\frac{d}{d\tau}A_0 = J_0A_1^*A_2, \quad \frac{d}{d\tau}A_1 = J_1A_0^*A_2, \quad \frac{d}{d\tau}A_2 = J_2A_0A_1. \quad (3.26a,b,c)$$

Defining the wave intensity of each wave mode as

$$E_j \equiv |A_j|^2 \in \mathbb{R}^+, \quad (3.27)$$

and utilizing the property $J_j^* = -J_j$ for purely imaginary coefficients, we recover the Manley–Rowe relations:

$$\frac{d_\tau E_0}{J_0} = \frac{d_\tau E_1}{J_1} = -\frac{d_\tau E_2}{J_2} = \Gamma \quad (3.28)$$

where $\Gamma(\tau) \equiv A_0^* A_1^* A_2 - A_0 A_1 A_2^*$. These relations imply the existence of conserved quantities (constants of motion) of the form

$$C_0^+ \equiv \frac{E_0}{J_0} + \frac{E_2}{J_2}, \quad C_1^+ \equiv \frac{E_1}{J_1} + \frac{E_2}{J_2}, \quad C^- \equiv \frac{E_0}{J_0} - \frac{E_1}{J_1}. \quad (3.29a,b,c)$$

Therefore, resonant triads satisfying Hypothesis 1 form a conservative dynamical system, where the exchange of wave intensities is strictly confined to a compact manifold defined by the constants of motion. It is worth noting that there exists a more general condition for conservative resonance (for $p, q = 0, 1, 2$):

$$\arg[J_p] - \arg[J_q] = n\pi, \quad n \in \mathbb{Z},$$

and an equivalent set of constants of motion can be derived.

Analytical solutions and explosive resonance

The analytical solutions of the amplitude equations (3.26a,b,c) are expressible in terms of the elliptic functions (see Weiland and Wilhelmsson, 1977) and consist of two distinct branches: bounded periodic oscillations or simultaneous explosive growth. The selection between these branches is determined solely by the relative signs of the interaction coefficients. To demonstrate this, we normalize (3.26a,b,c) into a universal form:

$$\frac{d}{d\tau} \hat{A}_0 = -s_0 \hat{A}_1^* \hat{A}_2, \quad \frac{d}{d\tau} \hat{A}_1 = -s_1 \hat{A}_0^* \hat{A}_2, \quad \frac{d}{d\tau} \hat{A}_2 = s_2 \hat{A}_0 \hat{A}_1, \quad (3.30a,b,c)$$

where the normalized amplitudes and sign indicators are defined as:

$$\begin{cases} \hat{A}_j = K \frac{A_j}{\sqrt{|J_j|}}, \\ s_j = -iJ_j/|J_j|, \end{cases}$$

with $K \equiv i\sqrt{|J_0||J_1||J_2|}$ and $s_j \in \{+1, -1\}$. The numerical solutions for (3.30a,b,c) can be easily obtained, and a series of example solutions for different combinations of s_j 's are given in figure 3.2. Following the derivation of the Manley–Rowe relations, we obtain the normalized constants of motion:

$$\hat{C}_0^+ \equiv s_0 \hat{E}_0 + s_2 \hat{E}_2, \quad \hat{C}_1^+ \equiv s_1 \hat{E}_1 + s_2 \hat{E}_2, \quad \hat{C}^- \equiv s_0 \hat{E}_0 - s_1 \hat{E}_1, \quad (3.31a,b,c)$$

where $\hat{E}_j = |\hat{A}_j|^2$ represents the normalized positive-definite wave intensity. The dynamics represented by (3.31a,b,c) have two topological regimes:

1. The first regime, *explosive resonance*, occurs strictly when the coefficient of the sum mode (mode 2, defined by $\omega_2 = \omega_0 + \omega_1$) opposes those of both constituent modes:

$$s_0 = s_1 = -s_2. \quad (3.32)$$




Figure 3.2: Example numerical solutions of the normalized conservative three-wave amplitude equations (3.30a,b,c) for distinct interaction coefficient sign signatures (s_0, s_1, s_2) . (a–c) Temporal evolution of the normalized wave intensities $|\hat{A}_j|^2$. (d–f) Corresponding 3D phase space trajectories plotting $|\hat{A}_{1,2}|$ against $|\hat{A}_0|$ and the relative phase factor $\cos(\theta_2 - \theta_1 - \theta_0)$. White and black markers indicate the initial and final states of the evolution, respectively. Left and Right Columns: Bounded resonance regimes, where the system is strictly confined to closed periodic orbits on a compact manifold. Middle Column: The explosive resonance regime ($s_0 = s_1 = -s_2$), where the relative phase locks to a constant value while permitting the system to follow an open trajectory with simultaneous amplitude divergence. See also Craik (1986).

In this specific topology, the invariants $\hat{C}^+$ become differences of wave intensities (e.g., $\hat{E}_2 - \hat{E}_0 = \text{const}$), so the conservation laws place no upper limit on the system, allowing all three wave amplitudes to diverge simultaneously as $(t_c - t)^{-1}$ towards a finite-time singularity t_c , as shown in figure 3.2b.

2. The second regime, *bounded resonance*, encompasses all other sign combinations. Regardless of whether the signs are uniform ($s_0 = s_1 = s_2$) or mixed (e.g., $s_2 = s_0 = -s_1$), at least one of the invariants in (3.31a,b,c) serves as a sum of positive intensities, defining a bounding sphere or ellipsoid in phase space and forcing the system to undergo periodic amplitude oscillations without divergence, as illustrated in figures 3.2a and 3.2c.

As indicated by figure 3.2d-f, the relative phase factor $\cos(\theta_2 - \theta_1 - \theta_0)$ remains constant

for the explosive regime, indicating a phase locking mechanism that sustains the growth; in contrast, the bounded resonance cases exhibit periodic variations in the relative phase, which prevents the system from diverging and confines it to closed orbits in phase space.

Lastly, it is important to address the apparent paradox of explosive growth in a conservative system. The resolution lies in the distinction between the wave intensity E_j tracked by the amplitude equations and the thermodynamic energy of the perturbation. In the presence of a background flow, a mode can extract energy from the mean flow, which further sustains the growth of wave intensities of the triad members while preserving the global energy balance of the fluid system. However, as noted by Weiland and Wilhelmsson (1977) and Craik (1986), this growth eventually outpaces the weakly nonlinear regime. At this threshold, higher-order effects — specifically third-order terms — may introduce frequency detuning or saturation that suppresses the singularity. Nevertheless, explosive triadic resonance remains a critical mechanism because it requires no sustained external forcing to drive the growth; instead, it provides a fast pathway for infinitesimal disturbances to grow rapidly, even if higher-order effects eventually saturate the interaction as the triad members reach large amplitudes.

Parametric instability via external forcing

In addition to the explosive instability, the conservative three-wave resonance can also manifest as a parametric instability when one mode is externally forced and serves as a pump to drive the growth of the other two modes. Without loss of generality, we assume A_0 as a frozen mode with constant amplitude. Under this assumption, the general amplitude evolution equations (3.26a,b,c) reduce to a coupled linear system for the two free modes ($j = 1, 2$):

$$\frac{d}{d\tau}A_0 = 0, \quad \frac{d}{d\tau}A_1 = J_1 A_0^* A_2, \quad \frac{d}{d\tau}A_2 = J_2 A_0 A_1. \quad (3.33a,b,c)$$

Differentiating with respect to τ decouples the system, yielding a second-order ordinary differential equation:

$$\frac{d^2}{d\tau^2}A_j = |A_0|^2 J_1 J_2 A_j, \quad (3.34)$$

which admit solutions of the form:

$$A_j(\tau) = A_j^0 e^{\pm \sqrt{J_1 J_2} |A_0| \tau}. \quad (3.35)$$

Substituting (3.35) back into (3.33a,b,c) reveals the constant coupling between the free modes:

$$\frac{A_2}{A_1} = \pm \sqrt{\frac{J_2}{J_1}} \arg[A_0]. \quad (3.36)$$

Recall that τ is the slow timescale of $\mathcal{O}(\epsilon)$, so the total complex frequency (eigenvalue) for the coupled system is the sum of the original linear complex frequency σ_j and the parametric correction:

$$\sigma_{j,\text{par}} = \sigma_j \pm \epsilon \sqrt{J_1 J_2} |A_0|, \quad (3.37)$$




Figure 3.3: Effect of detuning $\Delta\omega$ on parametric instability as given in (3.39). (a) Doppler-shifted eigenfrequencies $\text{Im}(\sigma_{\text{par},j}) - k_j\omega_0/k_0$ for modes $j = 1, 2$. The solid and dashed lines represent the upper and lower branches of the four coupled eigenvalues, respectively. Inside the instability band (shaded region), the upper and lower branches of the same mode lock to the same wave frequency, and the eigenvalues merge toward exact resonance where the modes synchronize to a common Doppler-shifted frequency. (b) Real growth rates $\text{Re}(\sigma)$. In the region where frequencies are locked, the growth rates bifurcate to form a symmetric instability bubble (shaded region indicating positive growth). The two wave modes share identical growth rates, leading to the appearance of a single curve for the unstable branch. Outside this region, the frequency mismatch becomes significant, and the coupled modes are no longer in near resonance.

which is either periodic when $\text{sgn}[J_1] \cdot \text{sgn}[J_2] > 0$ (J_1 and J_2 are both purely imaginary under Hypothesis 1), or exponential and implying instability when $\text{sgn}[J_1] \cdot \text{sgn}[J_2] < 0$. In the latter case, the conjugate corrections ($\pm$) exactly correspond to the eigenvalues being “pulled” symmetrically off the neutral axis while strictly preserving the time-reversibility of the underlying Hamiltonian system.

In realistic flows, the resonance condition may not be perfectly satisfied. Our general three-wave resonance formulation allows a detuning of $\epsilon\Delta\omega = \omega_2 - (\omega_1 + \omega_0)$. Following the same decoupling procedure as the perfect resonance case, the reduced second-order differential equation for the free modes now includes a first-order term representing the phase drift. For the free mode with amplitude A_2 , this is expressed as

$$\frac{d^2}{d\tau^2} A_2 + i\Delta\omega \frac{d}{d\tau} A_2 - |A_0|^2 J_1 J_2 A_2 = 0. \quad (3.38)$$

Assuming a solution consistent with the slow-time expansion, the corrected complex wave frequencies are now derived as:

$$\sigma_{j,\text{par}} = \sigma_j - i\frac{\epsilon\Delta\omega}{2} \pm \epsilon \sqrt{|A_0|^2 J_1 J_2 - \left(\frac{\Delta\omega}{2}\right)^2}. \quad (3.39)$$

Figure 3.3 illustrates the influence of detuning on the parametric instability based on (3.39). The presence of detuning introduces a threshold for instability, leading to a band of detuning values. Specifically, when the detuning is small enough such that $0 \leq (\Delta\omega/2)^2 < |A_0|^2 J_1 J_2$, the radical in (3.39) becomes real, resulting in a pair of complex conjugate eigenvalues with peak growth rates reached at perfect resonance ($\Delta\omega = 0$). Conversely, if the detuning is too large, the radical remains imaginary, and the instability does not exist.

Finally, we note that the frozen mode assumption is also valid for unforced triads but only in the regime where $|A_0| \gg |A_1|, |A_2|$. As the perturbation amplitudes eventually become comparable, the depletion effects become significant. In the absence of external forcing to replenish A_0 , the system inevitably transitions to the general three-wave conservative interactions confined by the Manley–Rowe relations (3.29a,b,c).

Chapter 4

Selection Rules for Conservative Triadic Resonances

In this section, we apply the general framework of three-wave resonance to the specific context of columnar vortices. We begin by classifying the linear stability properties of columnar vortices, with a particular focus on the characteristics of neutral modes that are relevant for triadic interactions. Building on this classification, we derive a set of selection rules that govern the nonlinear coupling between these modes, which in turn determine the possible resonance scenarios and their associated dynamical behaviors.

4.1 Linear Wave Modes

The linear stability of columnar vortices has been the subject of extensive research (e.g., Fabre et al., 2006; S. Lee and P. S. Marcus, 2023; Mayer and Powell, 1992). Rather than reiterating the full breadth of these studies, our main objective here is to provide a clear classification of the neutral modes ($\sigma = i\omega$), which represent the oscillatory degrees of freedom available for the expansion in (3.14) and subsequent triadic resonance coupling.

To facilitate this classification, we consider the reduction of (3.15) in terms of the primitive variables, which leads to the Howard–Gupta equation (Howard and Gupta, 1962) expressed here in terms of the radial velocity component $\tilde{u}_r(r)$:

$$\frac{d}{dr} \left(S \frac{d(r\tilde{u}_r)}{dr} \right) - \left(\frac{a(r)}{\Phi} + \frac{b(r)}{\Phi^2} + r \right) \frac{\tilde{u}_r}{r} = 0, \quad (4.1)$$

where the auxiliary functions are defined as:

$$\Phi(r) = -i\sigma + m\bar{\Omega}(r) + k\bar{u}_z(r), \quad (4.2a)$$

$$S(r) = \frac{r^2}{k^2 r^2 + m^2}, \quad (4.2b)$$

$$a(r) = r \frac{d}{dr} \left[\frac{S}{r} \left(\frac{d\Phi}{dr} + \frac{2m}{r} \bar{\Omega} \right) \right], \quad b(r) = 2kmS\bar{\Omega} \left[\frac{1}{r} \frac{d\bar{u}_z}{dr} - \frac{k}{m} \bar{\zeta} \right]. \quad (4.2c)$$

For the neutral modes of interest, the coefficients in (4.1) are purely real. Therefore, when in the absence of any singularities, the eigenfunctions $\tilde{u}_r(r)$ possess a global phase coherence. Meanwhile, the condition of mass conservation,

$$\frac{1}{r} \frac{d}{dr}(r\tilde{u}_r) + \frac{im}{r}\tilde{u}_\theta + ik\tilde{u}_z = 0, \quad (4.3)$$

imposes a quadrature relation (90° phase shift) between the phase of $\tilde{u}_r$ and the phases of the azimuthal and axial velocity components. Without loss of generality, we apply a global phase shift to all neutral, non-singular solutions in our subsequent analysis such that the radial velocity is purely imaginary ($\tilde{u}_r \in i\mathbb{R}$), rendering the tangential components real ($\tilde{u}_\theta, \tilde{u}_z \in \mathbb{R}$).

The solutions to (4.1) can carry critical layer singularities, which occur at locations where the Doppler-shifted frequency $\Phi(r)$ vanishes. The nature of the eigenfunctions and their associated dispersion relations are fundamentally influenced by whether such critical layers exist, which leads to two general wave categories:

1. *Regular modes* (or Kelvin waves), which exist when $\Phi(r) \neq 0$ for all r . In this case, the Howard–Gupta equation remains regular, yielding eigenfunctions that are analytic and possess the phase coherence described above. Their dispersion relations, $D(\omega, m, k) = 0$, are discrete and well-approximated by the large- k asymptotic theory (Stéphane Le Dizès and Laurent Lacaze, 2005).
2. *Critical layer modes*, which arise when $\Phi(r_c) = 0$ at some critical radius r_c and form a *continuous spectrum* in the dispersion plot. At the critical radius, the equation is singular, typically resulting in a logarithmic singularity and a phase jump that breaks the global phase coherence (Gally and Smets, 2020).

Notably, within the continuous spectrum, there exists a discrete set of *passive* critical layer modes that maintain structural regularity despite the presence of critical layer singularity. These modes can be interpreted as the continuation of discrete dispersion branches into the continuous spectrum (Fabre et al., 2006), whose dispersion relations can still be approximated within the large- k WKB framework developed by Stéphane Le Dizès and Laurent Lacaze (2005).

Figure 4.1 illustrates the dispersion characteristics of the Lamb–Oseen vortex, highlighting regular Kelvin waves, the continuous spectrum, and the discrete critical layer modes. The numerical results presented here and in subsequent sections are obtained using a spectral collocation method based on the mapped associated Legendre functions (MLEGS), which is capable of accurately resolving both regular and singular modes (S. Lee and P. S. Marcus, 2023; Sangjoon Lee and J. Wang, 2025; Matsushima and P. Marcus, 1997) and is detailed in Appendix 5A. The radial velocity profiles of representative modes are depicted in figure 4.1b–d. It can be seen that, for the discrete critical layer modes (figure 4.1d), the critical layer r_c is located far from the region where the eigenmode is localized, and the singularity is effectively suppressed, resulting in a smooth eigenfunction in the inner region that is




Figure 4.1: Characteristics of the inviscid linear modes of the Lamb–Oseen vortex obtained numerically by MLEGS (Sangjoon Lee and J. Wang, 2025). (a) Dispersion curves $\omega(k)$ of helical modes ($m = 1$). Solid lines denote Kelvin waves, and dashed lines indicate neutral discrete modes possessing critical layers. The shaded frequency band $\omega_{\text{CL}} = [-m, 0)$ is associated with the critical layer singularity and corresponds to the continuous spectrum. (b–d) Radial velocity component $r\tilde{u}_r(r)$ of selected eigenvectors ($m = 1, k = 3.0$): (b) A regular discrete Kelvin wave, exhibiting a purely real, smooth profile characteristic of global neutral modes; (c) A generic singular critical layer mode from the continuous spectrum. The singularity at the critical radius r_c (vertical dashed line) breaks analytic continuity, resulting in spectral broadening and a distinct phase jump in physical space; (d) A “passive” discrete critical layer mode. Unlike generic critical layer modes, this solution maintains high-order spectral convergence, confirming that the critical layer singularity is effectively suppressed and decoupled from the core dynamics. The eigenvalues associated with the mode profiles shown are marked as circles in the dispersion plot.

structurally indistinguishable from a regular mode. In contrast, the surrounding generic modes of the continuous spectrum (figure 4.1c) display distinct phase jumps at the critical radius and lack global phase coherence.

For the purpose of our nonlinear analysis, we categorize the neutral modes based on structural analyticity rather than simply the existence of critical radius. In the subsequent analysis, we shall focus on the set of *smooth neutral modes* who maintain analytic continuity and global phase coherence within their localized region. This group naturally includes all regular Kelvin waves and the discrete subset of critical layer modes.

4.2 Conservative Interactions between Smooth Neutral Modes

We now show that for columnar vortices, the interaction of *smooth neutral modes* (as classified above) is inherently conservative. This ensures that the dynamics are governed strictly by the Manley–Rowe relations derived in the previous section.

Recall that due to mass conservation, the radial velocity component $\tilde{u}_r$ is in quadrature with the azimuthal and axial components. Following the phase convention established earlier, let x_j 's be real functions of r such that:

$$\begin{pmatrix} \tilde{u}_r \\ \tilde{u}_\theta \\ \tilde{u}_z \end{pmatrix} = \begin{pmatrix} ix_1 \\ x_2 \\ x_3 \end{pmatrix}. \quad (4.4)$$

The vorticity field $(\tilde{\omega}_r, \tilde{\omega}_\theta, \tilde{\omega}_z)e^{i(m\theta+kz)+\sigma t}$ inherits this phase structure. Specifically:

$$\begin{pmatrix} \tilde{\omega}_r \\ \tilde{\omega}_\theta \\ \tilde{\omega}_z \end{pmatrix} = \begin{pmatrix} \frac{1}{r}\partial_\theta\tilde{u}_z - \partial_z\tilde{u}_\theta \\ \partial_z\tilde{u}_r - \partial_r\tilde{u}_z \\ \frac{1}{r}\partial_r(r\tilde{u}_\theta) - \frac{1}{r}\partial_\theta\tilde{u}_r \end{pmatrix} = \begin{pmatrix} i(\frac{m}{r}u_z - ku_\theta) \\ -ku_r - \partial_ru_z \\ \frac{1}{r}\partial_r(ru_\theta) + \frac{m}{r}u_r \end{pmatrix} = \begin{pmatrix} ix_4 \\ x_5 \\ x_6 \end{pmatrix}. \quad (4.5)$$

Thus, both velocity and vorticity vectors have the structure (Imag, Real, Real)[†].

The nonlinear interaction term $\mathbf{N}$ involves the cross product $\tilde{\mathbf{u}} \times \tilde{\boldsymbol{\omega}}$. The components of this product follow the pattern:

$$\begin{pmatrix} ix_1 \\ x_2 \\ x_3 \end{pmatrix} \times \begin{pmatrix} ix_4 \\ x_5 \\ x_6 \end{pmatrix} = \begin{pmatrix} x_7 \\ ix_8 \\ ix_9 \end{pmatrix}. \quad (4.6)$$

Therefore, the nonlinear forcing term has the structure (Real, Imag, Imag)[†]. Finally, to compute the interaction coefficient J , we take the inner product with the adjoint $\tilde{\mathbf{L}}$ in the primitive form. Since the adjoint eigenvectors for smooth neutral modes share the same phase structure as the direct modes ($\tilde{L}_r \in i\mathbb{R}$, $\tilde{L}_{\theta,z} \in \mathbb{R}$), the inner product yields:

$$J \sim \int (\tilde{L}_r^* N_r + \tilde{L}_\theta^* N_\theta + \tilde{L}_z^* N_z) r dr \sim (-i)(\text{Re}) + (\text{Re})(i) + (\text{Re})(i) \in i\mathbb{R}. \quad (4.7)$$

This proves that the interaction coefficients are purely imaginary, as assumed in Hypothesis 1. Hence, resonant interactions involving smooth neutral modes in columnar vortices are strictly conservative, governed by the Manley–Rowe relations. This finding sets the stage for the subsequent analysis of whether such conservative triads can exhibit explosive instability or are restricted to bounded oscillations.

4.3 Pseudoenergy and Selection Rules for Conservative Triads

Given the potentially infinite space of possible resonant triads as evidenced by the continuous variation of axial wavenumbers, a brute-force verification of explosive instability through calculations of the nonlinear interaction coefficients is impractical. To address this, we follow Cairns (1979), who leveraged the concept of pseudoenergy from plasma physics and proposed that the onset of hydrodynamic instability can be determined by the relative signs of the triad members’ pseudoenergy. The pseudoenergy can be defined to have the form¹:

$$\mathcal{E}_j \equiv -\frac{1}{4}\omega_j \frac{\partial D}{\partial \omega_j} |A_j|^2, \quad (4.8)$$

where $D(\omega, m, k)$ is the dispersion relation of the wave mode.

For stratified shear flows characterized by piece-wise constant or linear velocity profiles with piece-wise constant or exponential density profiles, studies have identified the existence of explosive resonant triads (Craik and Adam, 1979; Tsutahara, 1984; Tsutahara and Hashimoto, 1986). These results are consistent with the criterion proposed by Cairns (1979), where instability is observed when the highest-frequency mode possesses a pseudoenergy sign opposite to that of the other triad members. In the case of unstratified parallel shear flows with smooth profiles, Becker and Grimshaw (1993) demonstrated that critical layer modes are the sole carriers of negative pseudoenergy, identifying them as prerequisites for explosive resonance.

For vortical flows, Fukumoto and Hattori (2005) employed the pseudoenergy to investigate the parametric instability of the Rankine vortex. They posited that a wave pair in resonance via external strain must possess pseudoenergies of either opposite signs or both zero to trigger instability, a criterion they corroborated via direct eigenvector evaluation. Furthermore, they derived the pseudoenergy for two-dimensional waves directly from the total fluid kinetic energy, providing a partial justification for the application of (4.8) to vortical flows. Subsequently, S. Le Dizès (2008) extended this framework to vortices with

¹The negative sign in definition (4.8) is introduced to ensure that regular Kelvin modes possess positive pseudoenergy. This accounts for the difference in temporal conventions: S. Le Dizès (2008) define the wave phase as $\exp[i(kz + m\theta - \omega t)]$, whereas our formulation assumes $\exp[i(kz + m\theta + \omega t)]$. This sign reversal in the time domain necessitates the adjustment in the energy definition to maintain physical consistency with past literature.

smooth velocity profiles, such as the Lamb–Oseen vortex. By evaluating (4.8) in the inviscid regime via a WKBJ approach (see Stéphane Le Dizès and Laurent Lacaze, 2005), they demonstrated that for unstratified fluids, modes with negative pseudoenergy are confined to specific frequency bands associated with critical layer singularities. Hence, these modes are deemed essential components for the occurrence of the elliptical instability, a finding confirmed by subsequent numerical stability analyzes (Blanco-Rodríguez and S. Le Dizès, 2017; S. Le Dizès, 2008). In the analysis below, we follow the large- k WKBJ framework of Stéphane Le Dizès and Laurent Lacaze (2005) to investigate the pseudoenergy criterion and discuss the stability properties of the general three-wave resonance in columnar vortices.

We first note that, since the nonlinear interaction coefficients (J_0 , J_1 , and J_2) are solely calculated from the linear eigenvectors, they are invariant under the change of the viewing frame due to the Galilean invariance of the linear eigenvalue problem (refer to §3.2). Hence, an explosive triad determined by (3.32) in the original frame remains so in the new frame, and vice versa. A direct application of this invariance is that, instead of evaluating the triad’s pseudoenergy in the original frame, we can choose a specific reference frame in which a triad member appears stationary. In such a frame (denoted by primes), the wave frequencies of the now degenerate modes are exactly the same, and the pseudoenergy criterion used by Fukumoto and Hattori (2005) and S. Le Dizès (2008) can be evaluated by the relative signs of the dispersion slopes:

$$\text{sgn}[J_p \cdot J_q] = \text{sgn}\left[\frac{\partial D'}{\partial \omega'_p} \cdot \frac{\partial D'}{\partial \omega'_q}\right]. \quad (4.9)$$

Regarding the signs of the dispersion slopes, we employ the WKBJ framework, which has been utilized to approximate the linear dispersion relations and eigenvectors of columnar vortices (S. Le Dizès, 2008; Stéphane Le Dizès and Laurent Lacaze, 2005). We explicitly acknowledge that this asymptotic approach is formally derived for the large- k limit, and its accuracy varies across the parameter space. As demonstrated by Stéphane Le Dizès and Laurent Lacaze (2005), the WKBJ predictions align closely with exact numerical calculations for small azimuthal wavenumbers ($m \leq 2$), but exhibit noticeable deviations at higher m and very small k . Because our pseudoenergy analysis fundamentally relies on evaluating the relative signs of the dispersion slopes, any quantitative deviations in the predicted topology of the dispersion branches will inherently impact the accuracy of the sign-based criterion in these specific regimes. Consequently, these deviations establish a formal boundary for our approach. While the WKBJ formulation provides a highly tractable analytical framework for investigating the general properties of three-wave resonance, the subsequent stability conclusions are drawn strictly within this asymptotic context, and these limitations must be kept in mind when extending the predictions to non-asymptotic regimes. Specifically, in the large- k limit, the k^2 -proportional terms in the Howard–Gupta equation (4.1) dominate, and a local radial wavenumber can be defined as:

$$k_r(r) = k \sqrt{\frac{\Delta(r)}{\Phi(r)^2}}, \quad (4.10)$$

which is real in regions where the generalized Rayleigh discriminant $\Delta(r)$ is positive:

$$I \equiv \{r : \Delta(r) \equiv 2\bar{\zeta}(r)\bar{\Omega}(r) - \Phi(r)^2 > 0\}. \quad (4.11)$$

The eigenvectors are thus oscillatory in these regions, while evanescent elsewhere. Stéphane Le Dizès and Laurent Lacaze (2005) demonstrated that the neutral modes are localized in I and selected by a discretization condition that leads to quantized dispersion relations, whose derivative can be approximated as an integral over the localized region:

$$\begin{aligned} \frac{\partial D}{\partial \omega} &= \frac{\partial}{\partial \omega} \int_I \sqrt{\frac{\Delta}{\Phi^2}} dr \\ &= - \int_I \frac{\Delta + \Phi^2}{\Delta^{\frac{1}{2}} \Phi^2} \text{sgn}[\Phi] dr. \end{aligned} \quad (4.12)$$

The integrand in (4.12) shows that the sign of $\partial_\omega D$ depends solely on the sign of Φ in the oscillatory region. For the smooth neutral modes considered here, $\Phi(r)$ retains a uniform sign throughout I : for the regular modes, this is guaranteed by the absence of critical layers; for the discrete critical layer modes, the zero-crossing of Φ occurs outside the region of integration. Furthermore, since Φ represents the Doppler-shifted frequency relative to the local flow, which is frame-independent, the sign of $\partial_\omega D$ remains invariant regardless of the observing frame.

We now demonstrate that regular Kelvin waves cannot form explosive triads by using (4.12) and exploiting its frame invariance. Because this proof fundamentally relies on the dispersion slopes derived from the WKBJ approach, the mathematical guarantee inherits the asymptotic boundaries of the method.

Corollary 1 *In the large- k limit, regular modes (Kelvin waves) of a columnar vortex alone cannot form explosive triads.*

Proof 1 *By setting proper azimuthal rotation speed and z -directional translation speed, we can always position the continuous spectrum of the base flow such that it straddles zero:*

$$\min[\bar{\Omega}(r)] \leq 0 \leq \max[\bar{\Omega}(r)] \quad \text{and} \quad \min[\bar{u}_z(r)] \leq 0 \leq \max[\bar{u}_z(r)],$$

which implies that for any set of wavenumbers (m and k), the local flow frequency range contains zero:

$$\min[m\bar{\Omega} + k\bar{u}_z] \leq 0 \leq \max[m\bar{\Omega} + k\bar{u}_z]. \quad (4.13)$$

Regular modes, by definition, possess eigenfrequencies ω that lie strictly outside the continuous spectrum (the range of the background fluid frequencies) to avoid critical layers. Therefore, in this frame, the Doppler-shifted frequency $\Phi(r)$ must satisfy:

$$\Phi(r) < 0 \quad (\implies \omega < 0) \quad \text{or} \quad \Phi(r) > 0 \quad (\implies \omega > 0)$$

uniformly across the entire domain. This implies that the frame frequency ω and the intrinsic frequency Φ share the same sign for every regular mode.

Recall that the pseudoenergy is proportional to $\omega \cdot \partial_\omega D$. From (4.12), we know that $\text{sgn}[\partial_\omega D] = -\text{sgn}[\Phi]$. Thus, all regular Kelvin waves possess positive pseudoenergy in this reference frame. Since the nature of an instability (bounded or explosive) must be frame-invariant, and we have found a frame where all pseudoenergies are positive, regular modes cannot form explosive triads. We formally restrict this conclusion to the parameter regimes where the WKBJ framework accurately captures the dispersion topology (i.e., the large- k limit and small azimuthal wavenumbers, $m \leq 2$).

The prohibition of explosive instability extends to the broader class of smooth neutral modes. This generalization follows directly from the resonant structure of the Doppler-shifted frequencies.

Corollary 2 *In the large- k limit, explosive triadic resonance between smooth neutral modes is prohibited, except for the singular case where all three modes are static ($\omega_0 = \omega_1 = \omega_2 = 0$).*

Proof 2 *Explosive instability requires a triad whose interaction coefficient signs follow the relation (3.32). Assume, for the sake of contradiction, that a smooth triad exists with the first two modes sharing the same sign, $\text{sgn}[J_0] \cdot \text{sgn}[J_1] > 0$. By (4.9) and the property that $\text{sgn}[\partial_\omega D] = -\text{sgn}[\Phi]$ in the WKBJ approximation, this implies:*

$$\text{sgn}[\omega_0 + m_0\bar{\Omega} + k_0\bar{u}_z] = \text{sgn}[\omega_1 + m_1\bar{\Omega} + k_1\bar{u}_z]. \quad (4.14)$$

According to the resonance condition in (3.20), the frequency of the third mode is the sum of the first two. Since Φ is linear in ω, m, k , we obtain:

$$\begin{aligned} \text{sgn}[\omega_2 + m_2\bar{\Omega} + k_2\bar{u}_z] &= \text{sgn}[(\omega_0 + \omega_1) + (m_0 + m_1)\bar{\Omega} + (k_0 + k_1)\bar{u}_z] \\ &= \text{sgn}[(\omega_0 + m_0\bar{\Omega} + k_0\bar{u}_z) + (\omega_1 + m_1\bar{\Omega} + k_1\bar{u}_z)]. \end{aligned} \quad (4.15)$$

Thus, $\text{sgn}[\Phi_2] = \text{sgn}[\Phi_0] = \text{sgn}[\Phi_1]$. This indicates $\text{sgn}[J_0] = \text{sgn}[J_1] = \text{sgn}[J_2]$, which corresponds strictly to bounded resonance.

Overall, the prohibition of explosive instability between smooth neutral modes is strictly established here for the asymptotic regime where the WKBJ approximation remains valid. As discussed previously, this formally restricts our analytical guarantee to the large- k limit and small azimuthal wavenumbers ($m \leq 2$) (see Stéphane Le Dizès and Laurent Lacaze, 2005). Within this asymptotic domain, the only theoretical exception to this prohibition is the singular scenario identified by Moore et al. (1975), where all three interacting modes have zero frequency ($\omega = 0$). However, while the formal proof is constrained by these asymptotic limits, its physical implications are considerably more general. S. Le Dizès (2008) explicitly demonstrated that for monotonic profiles like the Lamb–Oseen vortex, the WKBJ theory provides a good approximation of the dispersion relations and their essential topology

extending all the way down to $k = 0$ for low- m modes. Although triadic interactions strictly dominated by high azimuthal wavenumbers fall outside the proven scope of these selection rules, the accuracy of the WKB topological predictions for these low-order azimuthal modes across the axial wavenumber spectrum is highly relevant. Consequently, we conclude that isolated monotonic vortices, such as those used in aeronautics applications for modelling aircraft wake vortices, are broadly resistant to explosive triadic resonance between smooth neutral modes across a wide and practical range of wave parameters, extending beyond the strict large- k limit.

The selection rules, however, do not preclude instability if external or internal symmetry-breaking is introduced. Examples include parametric instability driven by external forcing, or the excitation of critical layer singularities where the phase uniformity required for the selection rules no longer hold.

Appendix 4A Linear System in a Constant Rotating Frame

Consider an observing frame rotating with a constant angular velocity $\bar{\Omega}\hat{\mathbf{z}}$, a velocity field will appear to have a superposed velocity $-r\bar{\Omega}\hat{\boldsymbol{\theta}}$, and the differential operators in the rotating frame and the original frame are related as follows:

$$\boldsymbol{\nabla} = \boldsymbol{\nabla}' \quad \text{and} \quad \frac{\partial}{\partial t} = \frac{\partial}{\partial t'} - \bar{\Omega} \frac{\partial}{\partial \theta'}, \quad (4.16)$$

where $\{\cdot\}'$ denotes measurements and operators made with respect to the rotating frame. The Euler equations give

$$\begin{aligned} \frac{\partial \mathbf{v}'}{\partial t} - \bar{\Omega} \frac{\partial \mathbf{v}'}{\partial \theta} &= (\mathbf{v}' + r\bar{\Omega}\hat{\boldsymbol{\theta}}) \times (\boldsymbol{\nabla} \times (\mathbf{v}' + r\bar{\Omega}\hat{\boldsymbol{\theta}})) - \boldsymbol{\nabla} \varphi \\ &= \mathbf{v}' \times (\boldsymbol{\nabla} \times \mathbf{v}') + (r\bar{\Omega}\hat{\boldsymbol{\theta}}) \times (\boldsymbol{\nabla} \times \mathbf{v}') \\ &\quad - (2\bar{\Omega}\hat{\mathbf{z}} \times \mathbf{v}' - \boldsymbol{\nabla} \bar{\Omega}^2 r^2) - \boldsymbol{\nabla} \varphi. \end{aligned} \quad (4.17)$$

It is useful to note that

$$\hat{\boldsymbol{\theta}} \times (\boldsymbol{\nabla} \times \mathbf{v}') = \frac{1}{r} \boldsymbol{\nabla} (r v'_{\theta}) - \frac{1}{r} \frac{\partial \mathbf{v}'}{\partial \theta}, \quad (4.18)$$

so (4.17) reduces to

$$\frac{\partial \mathbf{v}'}{\partial t} = \mathbf{v}' \times (\boldsymbol{\nabla} \times \mathbf{v}') + \boldsymbol{\nabla} (r\bar{\Omega} v'_{\theta} + \bar{\Omega}^2 r^2 - \varphi) - 2\bar{\Omega}\hat{\mathbf{z}} \times \mathbf{v}'. \quad (4.19)$$

Taking poloidal-toroidal projection of (4.19) gives

$$\frac{\partial \mathbf{v}'}{\partial t} = \mathbb{P} [\mathbf{v}' \times (\boldsymbol{\nabla} \times \mathbf{v}') - 2\bar{\Omega}\hat{\mathbf{z}} \times \mathbf{v}'], \quad (4.20)$$

so the transformed linear eigenvalue problem reads:

$$\begin{aligned}
 \sigma' \tilde{\mathbf{R}}' &= \mathbb{M}_{mk} \tilde{\mathbf{R}}' + \mathbb{P}_{mk} \left[-r \bar{\Omega} \hat{\boldsymbol{\theta}} \times (\boldsymbol{\nabla} \times \mathbb{P}^{-1}[\mathbf{R}']) \right. \\
 &\quad \left. + \mathbb{P}^{-1}[\mathbf{R}'] \times (\boldsymbol{\nabla} \times (-r \bar{\Omega} \hat{\boldsymbol{\theta}})) - 2 \bar{\Omega} \hat{\mathbf{z}} \times \mathbb{P}^{-1}[\mathbf{R}'] \right] \\
 &= \mathbb{M}_{mk} \tilde{\mathbf{R}}' + \mathbb{P}_{mk} \left[-r \bar{\Omega} \hat{\boldsymbol{\theta}} \times (\boldsymbol{\nabla} \times \mathbb{P}^{-1}[\mathbf{R}']) \right] \\
 &= \mathbb{M}_{mk} \tilde{\mathbf{R}}' + im \bar{\Omega} \tilde{\mathbf{R}}',
 \end{aligned} \tag{4.21}$$

which implies a shift in the wave frequency while the growth rate and the eigenvector are unchanged:

$$\omega' = \omega + im \bar{\Omega} \quad \text{and} \quad \tilde{\mathbf{R}}' = \tilde{\mathbf{R}}. \tag{4.22}$$

Chapter 5

Generalized Parametric Instability of Columnar Vortices

In the preceding chapter, we established through the principle of pseudoenergy that explosive triadic resonance is strictly forbidden among smooth, neutral modes in isolated columnar vortices, provided the interacting waves fall within the valid large- k asymptotic regime. The resulting selection rules dictate that breaking this conservative internal constraint inherently requires a carrier of negative pseudoenergy, such as a critical layer, or a departure from the proven bounds of the WKBJ approximation. However, when the system is no longer isolated, but rather subjected to external energy injection via a pumping wave, these conservative constraints can be bypassed. This chapter explores this mechanism, known as parametric instability. We begin by reviewing the historical context of the classical elliptical instability, tracing its origins from pure analytical flows to the normal mode formulations for columnar vortices. We then generalize this framework to admit arbitrary pumping waves, developing a systematic tuning method based on non-degenerate perturbation theory to locate exact resonant triads across the continuous spectrum. Finally, we apply this generalized methodology to the Lamb–Oseen vortex, numerically validating our framework against classical results and systematically uncovering new unstable resonant loci.

5.1 Historical Context

Specific instances of parametric instability have been studied extensively in the past, most notably the elliptical instability driven by an external strain field (Kerswell, 2002; L. Lacaze, Birbaud, et al., 2005; L. Lacaze, Ryan, et al., 2007; Waleffe, 1990). The historical foundation of these studies begins with the analytical investigation of pure elliptical flows. Instead of considering the complex correction of a vortex, these early studies constructed a base flow consisting of a pure rotation of uniform vorticity 2γ superposed with an irrotational stretching field of strain rate e . This combination generates elliptical streamlines described

by the streamfunction

$$\psi = -\frac{1}{2}[(\gamma - e)x^2 + (\gamma + e)y^2], \quad (5.1)$$

characterized by an aspect ratio of $\sqrt{(\gamma + e)/(\gamma - e)}$. Numerical studies of this flow by Pierrehumbert (1986) demonstrated that the most unstable modes involve the bending and stretching of originally straight vortex lines into planar sinusoidal curves. A formal analytical theory was subsequently presented by Bayly (1986) using plane-wave solutions, which Waleffe (1990) expanded upon to derive a maximum instability growth rate of

$$\sigma_{\max} = \frac{9}{16} \frac{e}{\gamma} \quad (5.2)$$

for inertial waves in the limit of small strain.

This analytical limit for pure elliptical flows served as a critical benchmark, which C. Eloy and S. Le Dizès (1999) later adapted to estimate the maximum growth rates for actual localized columnar vortices as

$$\sigma = \frac{9}{16} S_0, \quad (5.3)$$

where S_0 is the strain rate precisely at the vortex center. The theory was subsequently broadened by Stéphane Le Dizès and Christophe Eloy (1999) to encompass multipolar strain fields. By defining an m -fold symmetric streamline as

$$\psi(r, \theta; p) = \frac{r^2}{2} - p \frac{r^m}{m} \cos(m\theta) = C, \quad (5.4)$$

where p is a real positive parameter controlling the strength of the strain field, the associated asymmetry can be quantified by the strain amplitude

$$e_m \equiv p \left(\frac{2mC}{m-1} \right)^{(m-2)/2}. \quad (5.5)$$

Based on this formulation, they derived the analytical maximum growth rate limits for quadrupolar ($m = 2$), triangular ($m = 3$), and square ($m = 4$) strain fields to be

$$\sigma_2 = \frac{9}{16} e_2, \quad \sigma_3 = \frac{49}{32\sqrt{3}} e_3, \quad \text{and} \quad \sigma_4 = \frac{3}{2} e_4, \quad (5.6)$$

respectively.

The mathematical transition from pure, unbounded elliptical flows to the stability analysis of actual localized columnar vortices was formalized through the normal mode framework established by Moore et al. (1975) and Tsai and Widnall (1976). They developed a standard “recipe” for what is often termed the Tsai-Widnall-Moore-Saffman (TSWM) instability. This approach begins by asymptotically expanding the unperturbed axisymmetric base flow to $\mathcal{O}(e)$ to obtain a steady-state ($\bar{\sigma} = 0$) velocity correction field characterized by an azimuthal

Instability	Vortex model	Axial flow	$\bar{m}$	(m_1, m_2)	Reference
Elliptical instability	Circular vortex	-	2	$(-1, 1)$	Moore et al. (1975)
	Rankine vortex	-	2	$(-1, 1)$	Tsai and Widnall (1976)
	Burger's vortex	-	2	$(-1, 1)$	C. Eloy and S. Le Dizès (1999)
	Lamb–Oseen vortex	-	2	$(-1, 1)$	
	Gaussian vortex	-	2	$(-1, 1)$	S. Le Dizès and Laporte (2002)
	Rankine vortex	constant axial flow inside vortex core	2	$m_2 - m_1 = 2$	L. Lacaze, Birbaud, et al. (2005)
	Batchelor vortex	jet-like axial flow	2	$m_2 - m_1 = 2$	L. Lacaze, Ryan, et al. (2007)
	Moore-Saffman vortex	jet-like axial flow	2	$m_2 - m_1 = 2$	Feys and S. A. Maslowe (2016)
Multi-polar strain	Rankine vortex	-	2, 3, 4	$m_2 - m_1 = \bar{m}$	C. Eloy and S. Le Dizès (2001)
Curvature instability	Batchelor vortex	jet-like axial flow	1	$m_2 - m_1 = 1$	Blanco-Rodríguez and S. Le Dizès (2017)
Radial perturbation	Lamb–Oseen vortex	-	0	-	S. A. Maslowe (2022)

Table 5.1: Summary of resonant triads examined in the literature. A blank axial flow entry indicates a vortex with no axial velocity component. In all listed cases, the perturbation modes have no axial dependency ($\bar{k} = 0$) and the base wave eigenvalues are either zero or assumed zero ($\bar{\sigma} = 0$), except for the Gaussian vortex case. Note that the radial perturbation at the bottom falls into the $(\bar{m}, \bar{k}) = (0, 0)$ self-interaction case that is not considered in this study.

symmetry of $\bar{m} = 2$ and no axial variation ($\bar{k} = 0$). The instability inherently arises from the convective triadic interaction between this externally imposed, stationary strain mode and two resonant linear normal modes (Kelvin waves) of the original unperturbed vortex. In this classical framework, the interacting wave pair is strictly constrained to satisfy the discrete geometric and temporal resonance conditions:

$$m_2 - m_1 = \pm \bar{m}, \quad k_1 = k_2, \quad \text{and} \quad \sigma_2 - \sigma_1 = 0. \quad (5.7)$$

Because this destabilization mechanism relies on resonant short-wavelength perturbations localized tightly within the vortex core, it is fundamentally categorized as a short-wave instability.

Following this foundational recipe, subsequent research systematically expanded the analysis of short-wave instabilities across a variety of vortex models. As summarized in Table 5.1, C. Eloy and S. Le Dizès (1999) computationally validated the analytical growth rates for the Burger and Lamb–Oseen vortices. The framework was later extended to more complex rotating shear flows by adding a uniform axial flow inside the Rankine vortex core (L. Lacaze, Birbaud, et al., 2005), and by analyzing the jet-like axial flows characteristic of the Batchelor and Moore-Saffman vortices (Feys and S. A. Maslowe, 2016; L. Lacaze, Ryan, et al., 2007). Additionally, the resonant configurations themselves were diversified to include multipolar strain instabilities where $m_2 - m_1 = 2, 3, 4$ (C. Eloy and S. Le Dizès, 2001), and curvature instabilities driven by the bending of the vortex filament corresponding to an azimuthal symmetry of $\bar{m} = 1$ (Blanco-Rodríguez and S. Le Dizès, 2017).




Figure 5.1: Graphical representations of resonant triads in the Lamb–Oseen vortex. (a) Resonance sustained by a stationary pumping wave with zero axial wavenumber ($m_p, k_p = 0, \omega_p = 0$). Star symbols mark crossing points between distinct dispersion branches ($k_i = k_j$, $\omega_i = \omega_j$) in resonance with $|m_p| = |m_i - m_j|$; filled circles indicate intersections with the zero-frequency axis, corresponding to a pair of conjugate modes ($\pm m$) in resonance with a static strain field with symmetry $m_p = 2m$. (b) General triadic resonance satisfying the frequency matching condition $\omega_1 + \omega_2 = \omega_3$ alongside wavenumber matching. The specific examples represent resonance between modes (m, k, ω) : $(1, 5.6, -1.2) + (2, 0.50, -2.0) = (3, 6.0, -3.3)$ (dash-dotted lines with star markers), and $(1, 4.1, -1.2) + (2, 0.98, -2.1) = (3, 5.1, -3.3)$ (dashed lines with circular markers).

However, despite the diversity of the base flows investigated and the variations in the azimuthal wavenumbers, all these extended classical phenomena constitute a highly constrained subclass of the generalized parametric instability. Specifically, they universally rely on a specific restricted configuration: a pumping wave that is strictly stationary ($\bar{\sigma} = 0$) and geometrically two-dimensional ($\bar{k} = 0$). Our objective in this chapter is to lift these rigid geometric and temporal constraints. By generalizing the resonance methodology, we establish a unified framework capable of admitting arbitrary pumping waves of any frequency and spatial geometry, allowing us to locate and compute resonant triads without relying on the underlying conjugate symmetries that restricted the classical studies.

5.2 Resonance Tuning via Non-Degenerate Perturbation Theory

To move beyond the constraints of classical elliptical instabilities, we consider the general triadic resonance framework where the parameters of the pump mode (m_0, k_0, ω_0) are treated as free variables. In classical configurations driven by a stationary pump wave, resonant triads are located strictly at the discrete intersections of the dispersion curves, which were initially identified by Moore et al. (1975) at the zero-crossings of the dispersion curves via the conjugate symmetry of the linear eigenvalues, as highlighted by solid round marks in figure 5.1a for the Lamb–Oseen vortex, and later extended to the general cross points (C. Eloy and S. Le Dizès, 1999; Pradeep and Hussain, 2001), as indicated by solid stars in the same figure.

On contrary, when the pumping frequency and axial wavenumber are non-zero, identifying resonant triads becomes significantly more complex. Schematically, generalized resonant triads form a closed parallelogram in the fully continuous (k, ω) dispersion space, as shown in figure 5.1b: the vectors (k_0, ω_0) and (k_1, ω_1) define the adjacent sides, while the sum mode (k_2, ω_2) completes the vertex. Because these vertices can translate continuously along their respective dispersion branches, the number of theoretically permitted resonant triads is effectively infinite.

For the general case, if the full dispersion relations are known *a priori*, identifying these generalized resonant triads is geometrically straightforward: one can simply shift the origin of the dispersion curve of a constituent mode along the dispersion vector of the pump mode until an intersection with the third mode’s dispersion curve is found. This shifting is equivalent to adjusting the relative axial wavenumbers of the participating modes. However, mapping the global dispersion landscape at high resolution is computationally prohibitive for complex base flow profiles. Furthermore, numerical simulations of weakly-nonlinear amplitude equations require a strict tolerance for perfect resonance ($\Delta\omega \rightarrow 0$); otherwise, phase drifting will suppress the parametric instability.

To overcome this without requiring exhaustive global dispersion mapping, we introduce an iterative tuning method based on non-degenerate perturbation theory. Because the dispersion relation varies continuously with the axial wavenumber k , we can perform a linear expansion of the frequency mismatch with respect to a small wavenumber correction Δk . Assuming the pump mode k_0 is fixed, the frequency mismatch expansion is given by:

$$\Delta\omega(k + \Delta k) \approx \Delta\omega(k) + \left(\frac{\partial\omega_2}{\partial k} - \frac{\partial\omega_1}{\partial k} \right) \Delta k = 0. \quad (5.8)$$

In our numerical implementation, we determine the local group velocities by applying a finite test shift δk to the axial wavenumbers of the free modes ($k'_j = k_j + \delta k$). We then compute the resulting eigenvalue shift $\delta\sigma_j$ via the first-order correction derived from non-degenerate perturbation theory:

$$\delta\sigma_j = \left\langle \tilde{\mathbf{L}}_j \left| \left(\mathbb{M}_{m_j k'_j} - \mathbb{M}_{m_j k_j} \right) \tilde{\mathbf{R}}_j \right. \right\rangle. \quad (5.9)$$




Figure 5.2: Iterations of the resonance tuning process for a triad driven by the pump mode $(m_0, k_0, \omega_0) = (0, 2.00, 0.5334)$ of the Lamb–Oseen vortex. Note that iterations 3–6 are overlapping as the frequency mismatch approaches machine precision during the tuning. Solid lines are the dispersion curves of the two free modes, $\omega_0 + \omega_1(k_1)$ and $\omega_2(k_1 + k_0)$. They are provided and shifted here for visual clarity so that their intersection corresponds to perfect resonance. Dashed lines trace the trajectories of the perturbation-based Newton–Raphson algorithm as it projects from one wavenumber approximation $k_1^{(n)}$ to the next $k_1^{(n+1)}$ based on the local group velocity mismatch given by (5.10). The inset (lower right) quantifies the performance, showing the frequency mismatch $|\Delta\omega|$ reducing to a precision of $\mathcal{O}(10^{-1})$ during the first iteration, and to $\mathcal{O}(10^{-14})$ after 6 iterations.

By projecting the local group velocity mismatch, the wavenumber correction required to achieve perfect resonance is iteratively estimated as:

$$\Delta k \approx \delta k \cdot \frac{\sigma_1 + \sigma_0 - \sigma_2}{\delta\sigma_2 - \delta\sigma_1}. \quad (5.10)$$

This algorithm operates as a highly efficient Newton–Raphson root-finding scheme. As demonstrated in Figure 5.2, the frequency mismatch $|\Delta\omega|$ typically drops by orders of magnitude within a single iteration, rapidly reaching machine precision ($\mathcal{O}(10^{-14})$). This tuning capability successfully generalizes our stability analysis, allowing us to systematically locate exact parametric instability configurations across an arbitrary spectrum of pumping frequencies and geometries.

It is crucial to recognize that while our algorithm can locate an effectively countless number of kinematic resonances, resonance does not inherently guarantee instability. The

pseudoenergy selection rules established in the previous chapter must still be satisfied. For a generalized parametric instability to physically manifest and overcome the bounded Manley–Rowe constraints, the mathematically tuned triad must possess a negative energy mode (specifically, a critical layer mode). The presence of this active singularity is what ultimately couples the externally forced energy of the pump wave into the intrinsic modal dynamics, driving the exponential growth.

5.3 Generalized Parametric Instability of the Lamb–Oseen Vortex

To rigorously validate our generalized framework, we first reconstruct the classical elliptical instability for the Lamb–Oseen vortex. The steady-state ($\bar{\sigma} = 0$) strain mode generated by the external forcing requires special mathematical treatment because its boundary conditions ($u_\theta \rightarrow r$ as $r \rightarrow \infty$) differ fundamentally from the regular wave modes that naturally decay in the far field. We resolve this by formulating a specific boundary value problem for the strain field, the complete spectral construction of which is detailed in Appendix 5A. The resulting numerical solution for the core response profile $g(r)$ and its exact comparison with historical literature are shown in Figure 5.3. It is perfectly consistent with the results derived by C. Eloy and S. Le Dizès (1999), noting that their formulation utilized a radial variable transformation $x \equiv 2r$ and a solution function $F(x) = r^2 g(r)$, which effectively scales the boundary limits by a factor of 2.




Figure 5.3: (a) The numerically computed profile $g(r) = \tilde{u}_r/r$ for the steady-state strain response of the Lamb–Oseen vortex, converging correctly to 1 in the far field. (b) Overlay of our constructed solution (solid line) against the historical results of C. Eloy and S. Le Dizès (1999) (dashed line), demonstrating perfect mathematical agreement when applying a specific radial scaling factor of $x = 2r$.

C. Eloy and S. Le Dizès (1999)		Present study	
k	$\text{Re}[\sigma_{\text{par}}]/\epsilon$	k	$\text{Re}[\sigma_{\text{par}}]/\epsilon$
2.26	1.3790	2.260814	1.379248
3.94	1.3888	3.957585	1.388915
-	-	5.611789	1.391302
-	-	7.252515	1.392252

Table 5.2: Comparison of the axial wavenumbers (k) and the parametric growth rates computed for the elliptical instability of the Lamb–Oseen vortex between a pair of helical ($m = \pm 1$) wave modes with equal k . The numerical values of the growth rate in this work are calculated based on our generalized interaction coefficients using MLEGS. Note that C. Eloy and S. Le Dizès (1999) only computed the parametric growth rate for the first two degenerate pairs, and their results are scaled by a factor of 2 here due to the radial scaling $x = 2r$ used in their specific pump mode formulation.

In the classical configuration defined by Moore et al. (1975), the free modes interacting with this strain field are chosen via the azimuthal symmetry of the linear dispersion relations ($m_2 = -m_1 = 1$, with $\omega_1 = \omega_2 = 0$), which correspond to the zero-crossings of the dispersion curves. By incorporating our numerically constructed steady-state strain field directly into our generalized amplitude equations and locate the zero-crossings numerically following our perturbation approach, we calculate the resulting parametric growth rates for these pairs. As detailed in Table 5.2, our computed growth rates are in excellent agreement with the numerical results generated by C. Eloy and S. Le Dizès (1999) to four significant figures (again, accounting for the scaling factor of 2 inherent in their pump mode formulation). Furthermore, leveraging our spectral solver, we naturally extend this classical calculation to higher wavenumbers ($k > 4$), demonstrating that the parametric growth rate continues to monotonically increase, albeit at a progressively diminishing rate.

Having validated the method, we now move beyond the classical elliptical constraints to explore the fully generalized parametric instability of the Lamb–Oseen vortex. A unique and critical advantage of our flexible resonance tuning method is the ability to easily perform “mode swapping”. If a localized triad yields only bounded oscillations because the free modes share the same pseudoenergy sign as the pump mode (e.g., a critical layer pump driving two regular positive-energy free modes), we can mathematically swap the pump mode m_0 with one of the free modes m_1 . This dynamical reconfiguration forces the new free modes to possess opposite pseudoenergies (e.g., one regular and one critical layer mode), thereby satisfying the pseudoenergy criterion and unlocking the parametric instability. Consequently, as long as the initial pump mode and the first free mode possess opposite pseudoenergy signs, a parametric instability configuration is mathematically guaranteed to exist regardless of the specific physical nature of the resulting sum mode.

Applying this mode-swapping logic, we perform a systematic search for generalized




Figure 5.4: Resonance loci in the (k_0, k_1) plane for triads involving a helical ($m_1 = 1$) mode and a pumping azimuthal wavenumber of (a) $m_0 = 1$; (b) $m_0 = 2$; (c) $m_0 = 0$. Solid line segments indicate regions capable of triggering parametric instability based on the pseudoenergy criterion, while the gray dashed continuations represent the remaining parts of the resonance loci where parametric instability is suppressed as the discrete critical layer modes transferred to regular Kelvin waves along the dispersion branch. Only loci containing at least one instability segment are plotted. The search is limited to the neutral branches as resolved for each azimuthal wavenumber, except for $m_0 = 0$ where we consider the first 6 radial branches due to the dense spectral spacing.

parametric instabilities by actively isolating triads where the first two modes maintain opposite pseudoenergy signs. Relying on the energy signs established by the WKBJ framework, we effectively reduce the search to scanning the continuous parameter space of the arbitrary pump mode m_0 while fixing the first free mode m_1 as a discrete critical layer mode. Restricting our numerical search to the discrete neutral branches within the range $m \in [-3, 3]$ and $0 \leq |k| \leq 10$, we successfully locate numerous, continuous parametric instability configurations driven by pumping wavenumbers of $m_0 = 0, 1, 2$ interacting with discrete critical layer modes of $m_1 = 1$. These configurations are mapped comprehensively in Figure 5.4, highlighting the specific solid line segments in the (k_0, k_1) resonance loci where the theoretical pseudoenergy constraints are successfully satisfied. We computed the parametric growth rates for these generalized configurations directly from their nonlinear interaction coefficients, finding them all to be real and positive. This fundamentally confirms the predictive power of the pseudoenergy criterion and demonstrates that, with arbitrary resonance tuning, parametric instability in columnar vortices is both pervasive and effectively countless.

It is worth pointing out that, because of the conjugate and azimuthal symmetries of the underlying eigenvalue problem, the existence of parametric configurations involving $m_1 = 1$ implies the existence of a symmetric set of $m_1 = -1$, and the corresponding configurations are not presented here to avoid redundancy. Additionally, although we are not able to locate any parametric configurations involving a critical layer mode of $m = 2$ within the specified search range, we expect such configurations to exist at higher wavenumbers due to the limited span of their neutral branches at lower wavenumbers, which is evident in figure 5.1. We also want to note that our numerical search is strictly restricted to smooth neutral modes, which are verified by monitoring their spectral expansion coefficients based on the spectral convergence guaranteed by the MLEGS basis functions for smooth profiles. This distinguishes our work from Fabre et al. (2006) and Stéphane Le Dizès and Laurent Lacaze (2005), where dispersion curves of the discrete critical layer modes are extended to lower wavenumbers as the critical layer singularity moves closer to the vortex core and becomes more prominent. Nevertheless, while we do not claim to have exhausted the entire parameter space, the results, as indicated by the continuous resonance loci in figure 5.4, suggest that parametric instability is relatively common and effectively countless. Most importantly, we have demonstrated that, with the pseudoenergy criterion and the perturbation-based tuning method, we can efficiently locate parametric configurations for arbitrary flow profiles, pumping frequencies, and geometries.

Appendix 5A Spectral Construction of Strain Modes

To strictly validate our generalized parametric instability framework against the classical elliptical instability, we must accurately construct the steady-state velocity correction field (the strain mode) generated by an external quadrupolar ($m = 2$) strain. We utilize a highly efficient spectral projection method based on Mapped Associated Legendre Polynomials (MLEGS).

Assuming a time-independent external strain field of the form

$$\mathbf{u}_{\text{strain}} = e \cdot r \sin(2\theta) \hat{r} + e \cdot r \cos(2\theta) \hat{\theta}, \quad (5.11)$$

we expect the velocity components of the strain mode to take the form:

$$u_r(r, \theta) = \tilde{u}_r(r) \sin(2\theta), \quad u_\theta(r, \theta) = \tilde{u}_\theta(r) \cos(2\theta), \quad u_z(r, \theta) = \tilde{u}_z(r) \cos(2\theta). \quad (5.12)$$

The boundary conditions demand $\tilde{u}_r, \tilde{u}_\theta \rightarrow r$ and $\tilde{u}_z \rightarrow 0$ as $r \rightarrow \infty$, while remaining completely regular at the origin. Linearizing the incompressible Euler equations about the base flow $\bar{\mathbf{u}}$ and applying the continuity constraint allows us to express the tangential and axial components purely in terms of the radial component. Let $\bar{\zeta}(r) = \frac{1}{r} \frac{d}{dr} (r \bar{u}_\theta)$ denote the background axial vorticity. The entire system reduces to a single boundary value problem for the function $g(r) \equiv \frac{1}{r^2} (r \tilde{u}_r)$:

$$r^2 g''(r) + 5r g'(r) - \frac{r^2 \frac{d}{dr} \bar{\zeta}}{\bar{u}_\theta} g(r) = 0, \quad (5.13)$$

subject to the refined boundary conditions:

$$g(r) \rightarrow 1 \text{ as } r \rightarrow \infty, \quad \text{and} \quad g(r) = \mathcal{O}(1) \text{ as } r \rightarrow 0. \quad (5.14)$$

To solve (5.13), we construct the solution using a linear combination of Mapped Associated Legendre Polynomials (Matsushima and P. Marcus, 1997). MLEGS is uniquely suited for this semi-infinite domain because the polynomials explicitly possess the mathematical pole properties

$$P_{L_n}^m(r) \rightarrow \mathcal{O}(r^{|m|}) \text{ as } r \rightarrow 0, \quad (5.15a)$$

$$P_{L_n}^m(r) \rightarrow \mathcal{O}(r^{-|m|}) \text{ as } r \rightarrow \infty. \quad (5.15b)$$

Consequently, only the 0th-order polynomials ($m = 0$) naturally satisfy the constant $\mathcal{O}(1)$ boundary behavior required at both the origin and infinity, and we represent the solution as

$$g(r) = \sum_{n=0}^{\infty} g_n^0 P_{L_n}^0. \quad (5.16)$$

Crucially, for this spectral representation to be complete and mathematically valid, the background vorticity term must be compatible with the pole behaviors of the MLEGS basis. Specifically, the expression $r^2 \frac{d}{dr} \bar{\zeta} / \bar{u}_\theta$ must be expandable in Taylor series near the origin and Laurent series near infinity. For the Lamb–Oseen vortex where $\bar{u}_\theta(r) = q(1 - e^{-r^2})/r$, the axial vorticity derivative explicitly yields $-r^2 \frac{d}{dr} \bar{\zeta} / \bar{u}_\theta = 4r^2/(e^{r^2} - 1)$. This term is strictly analytic at the origin and decays exponentially at infinity, perfectly matching the required pole behaviors and guaranteeing spectral convergence.

We translate the continuous differential equation into a discrete matrix operator $\mathbf{G}$ operating on the spectral coefficients $\mathbf{g}^0$. Recognizing that

$$r^2 g'' + 5r g' = r \frac{d}{dr} \left(r \frac{dg}{dr} \right) + 4r \frac{dg}{dr}, \quad (5.17)$$

we utilize the established spectral operator $\mathbf{D}$ representing $r \frac{d}{dr}$, alongside a diagonal physical-space multiplication matrix $\mathbf{H}$ for the background vorticity term. The null space of the resulting operator:

$$\mathbf{G} = \mathbf{D}^2 + 4\mathbf{D} - \mathbf{T}^{-1}\mathbf{H}\mathbf{T}, \quad (5.18)$$

contains the exact solution to the boundary value problem. Extracting the zero-frequency eigenvector and transforming it via the physical space projection operator $\mathbf{T}$ instantly yields the accurate profile $g(r)$.

To reintegrate this solution into the triadic resonance framework established in Chapter 3, we reconstruct the velocity field. The radial and azimuthal components are cleanly expressed as the curl of a plane streamfunction $\psi(r, \theta) = -\frac{1}{2}r^2 g(r) \cos(2\theta)$. Applying Fourier expansion in the azimuthal direction, the $\bar{m} = 2$ streamfunction reduces to the purely real component $\psi_2(r) = -\frac{1}{4}r^2 g(r)$.

However, a critical numerical complication arises: the mapped associated Legendre basis requires perturbations to decay strictly faster than $r^{-|m|}$ at infinity for spectral convergence. Because $\psi_2(r) \propto r^2 g(r) \sim \mathcal{O}(r^2)$, the raw strain mode diverges mathematically at far field. To correct this without corrupting the delicate inner-core dynamics where the wave interactions actually occur, we apply a smooth hyperbolic tangent damping function:

$$G(r) = \frac{1}{2} \left[1 - \tanh \left(\frac{r - r_c}{\Delta r} \right) \right]. \quad (5.19)$$

Multiplying $\psi_2(r)$ by $G(r)$ forces an aggressive, smooth decay beyond a chosen cutoff radius r_c (typically placed far outside the vortex core, e.g., $r_c = 20$). This avoids the severe Gibbs phenomena caused by sharp piecewise cutoffs, while guaranteeing that the perturbation meets the strict pole decay conditions. Finally, the spectral coefficients for the damped forcing field f_n^2 are efficiently calculated via Gauss-Legendre quadrature integration:

$$f_n^2 \approx \frac{1}{(N_n^2)^2} \sum_i \psi_2(r(\mu_i)) G(r(\mu_i)) P_n^2(\mu_i) W(\mu_i), \quad (5.20)$$

where N_n^m are the normalization factors, and μ_i and $W(\mu_i)$ are the Gauss-Legendre abscissas and weights, respectively, for the mapped associated Legendre polynomials. This procedure yields the spectral coefficients of the strain mode, which can then be directly incorporated into the nonlinear interaction coefficients for the parametric instability analysis.

Chapter 6

Roles of the Critical Layer in the Instability of Columnar Vortices

6.1 Critical Layers in Incompressible Flows

The selection rules derived in Chapter 4 establish that, within the proven asymptotic bounds, isolated columnar vortices are theoretically resistant to explosive resonance when the interacting waves remain smooth and neutral. This constraint applies not only to regular Kelvin waves but also to the discrete subset of critical layer modes, provided their critical layers are located sufficiently far from the vortex core where the background vorticity gradient is negligible. In such cases, the singularity is effectively passive, and the modes continue to adhere strictly to the conservative constraints. Consequently, the onset of intrinsic instability requires a physical mechanism capable of violating the strict conservation laws imposed by the Manley–Rowe relations. While the preceding chapter quantified how external forcing can actively pump energy into the system to trigger parametric instability, we now turn to a qualitative discussion of the second fundamental pathway: the active critical layer. In contrast to the rigorous perturbation analysis employed earlier, the objective of this chapter is to outline the physical principles by which critical layer singularities break the Hermitian symmetry of the linear operator. This internal symmetry-breaking allows resonant disturbances to extract energy directly from the background shear, transforming the structurally conservative system into an open, non-conservative one capable of intrinsic instability.

To understand the physical mechanism of the critical layer, we start by introducing the energy budget of the disturbance field. The time rate of change of the total disturbance kinetic energy, E_K , is governed solely by the shear production term:

$$\begin{aligned} \frac{d}{dt}E_K &= - \int_V \mathbf{u} \cdot (\mathbf{u} \cdot \nabla \bar{\mathbf{u}}) dV \\ &= - \int_V \left(u_r u_\theta \cdot r \frac{d}{dr} \bar{\Omega} + u_r u_z \cdot \frac{d}{dr} \bar{u}_z \right) dV, \end{aligned} \tag{6.1}$$

where the two terms in the integral represent the work done by the disturbance against the rotational shear ($r \frac{d}{dr} \bar{\Omega}$) and the axial shear ($\frac{d}{dr} \bar{u}_z$), respectively. This relation indicates that, without external forcing (which is the mechanism for the parametric instability), a disturbance must be able to tap into the energy of the background shear via wave-mean interactions in order to sustain its growth.

In this regard, the critical layer is an efficient pathway to the background shear. The nature of the critical layer is a local resonance between the wave mode and the background flow at the critical radius r_c , where the wave phase velocity matches the local flow velocity, suggesting a zero Doppler-shifted frequency: $\Phi(r_c) = 0$. In the linear theory governed by the Howard–Gupta equation (4.1), this wave-mean resonance manifests as a mathematical singularity, leading to a logarithmic divergence in the vorticity fluctuation and essentially separating the eigenmodes into disjointed pieces in the inviscid limit. In fact, numerical calculations of the Lamb–Oseen vortex reveal pairs of degenerate modes sharing the same critical radius, where one mode is localized inside the critical radius and the other outside (S. Lee and P. S. Marcus, 2023). As the linear eigenvalue problem is merely a mathematical idealization, the singularity is not expected to persist in the physical system. And the manner in which these piecewise linear solutions are connected to form a global solution depends entirely on the chosen regularization method that reintroduces the physics neglected by the leading-order inviscid approximation. In the general study of critical layers, both viscosity and nonlinearity are recognized as mechanisms capable of smoothing the singularity (Stewartson, 1977; T. Warn and H. Warn, 1976, 1978). Their relative importance is dictated by the competition between viscous diffusion and nonlinear effects; while often studied as distinct processes, they are not mutually exclusive and may exert comparable influence during the transition from a viscously dominated regime to a highly nonlinear state (Benney and Bergeron, 1969).

In our analysis, a diminishing-viscosity regularization is implicitly adopted through the large- k asymptotic framework of Stéphane Le Dizès and Laurent Lacaze (2005). This approach assumes that an inviscid solution of the linear problem should be the limit of its viscous counterpart as the Reynolds number tends to infinity. Mathematically, this justifies deforming the integration contour into the complex plane to avoid the singularity (Lin, 1955), leading to a finite, complex residue that ensures causality and provides the necessary connection between the piecewise solutions. Notably, this process introduces a discrete phase jump across r_c (typically $-\pi$) that persists even in the high- Re limit and allows the critical layer mode to exert a torque and extract kinetic energy from the background shear. Within the concept of pseudoenergy, this extraction identifies the critical layer mode as a negative pseudoenergy mode. In a physical medium, the sign of a mode’s pseudoenergy determines its role in the disturbance’s total energy budget: a positive energy mode requires a net energy investment to grow, while a negative energy mode grows by effectively withdrawing energy from the supporting flow. Therefore, a disturbance consisting solely of positive energy modes cannot undergo simultaneous growth, as the amplification of one member would necessitate the decay of others to satisfy the total budget. The negative energy mode removes this constraint and can act as a source that enables simultaneous amplification, which explains

why critical layer modes are essential for both the explosive resonance and the parametric instability that are analyzed in Chapter 4.

As the critical radius r_c moves deeper into the mode’s localized region, its role in the triadic exchange becomes increasingly dominant. Although the individual modes may remain neutral in an inviscid linear sense, the active critical layer facilitates a continuous energy exchange with the background shear that renders the triadic resonance *non-conservative*. This is reflected in the interaction coefficients becoming complex ($J \in \mathbb{C}$), a result of the complex residues associated with the regularized critical layers. Because the system is no longer conservative, it is liberated from the Manley–Rowe relations and the strict selection rules that govern smooth neutral modes. Consequently, the theoretical prohibition against simultaneous amplification is lifted. However, the introduction of an open energy pathway implies that the classical notion of “explosive resonance”, a concept inherently rooted in closed, conservative energy exchanges, becomes physically ambiguous. Without these overarching analytical constraints, general stability conclusions can no longer be broadly drawn. Instead, the stability characteristics of the triadic system shall be assessed on a case-by-case basis by explicitly computing the specific complex interaction coefficients for each individual triad.

In addition, at higher amplitudes, nonlinearity often becomes the dominant regularization effect. This is particularly evident in the emergence of “stacked” resonances. Due to the linear relation $\Phi_2(r) = \Phi_0(r) + \Phi_1(r)$ governing a resonant triad, if any two members share a critical layer, the third mode is geometrically constrained to possess a critical layer at the precise same location. Such a configuration is likely to trigger a nonlinear cascade, as these primary modes can spawn higher-order harmonics or secondary interactions that remain spatially locked to the *same* critical radius. In these instances, the interaction transcends the simple triadic resonance assumed in our weakly-nonlinear formulation. The presence of strong, localized nonlinearities violates standard asymptotic ordering, thereby necessitating a dedicated regularization framework (S. Maslowe, 1986). While such nonlinear critical layer theory is well-established for planar geometries (Benney and Bergeron, 1969; Stewartson, 1977; T. Warn and H. Warn, 1976, 1978), the specific case of three-dimensional critical layers in cylindrical columnar vortices remains comparatively unexplored and warrants further investigation.

6.2 The Roles of Stratification in Vortex Stability and Resonance

Throughout Part I, we established the stability of isolated columnar vortices within the strict limit of homogeneous, unstratified flow. While this incompressible limit provides a foundational baseline, introducing background density variations fundamentally alters the stability landscape of the vortex. Crucially, the mathematical architecture of our generalized triadic resonance formulation, including the formal definition of pseudoenergy as introduced

in Cairns (1979) and Craik and Adam (1979), remains entirely applicable. However, because density variations alter the underlying dispersion relations and wave structures, the strict prohibition of explosive resonance derived previously for unstratified smooth, neutral modes no longer applies. This additional physical mechanism influences wave interactions on two distinct levels: first, by globally modifying the underlying linear modes, and second, by actively altering the dynamics near the critical layers.

At the first level, macroscopic density variations alter the underlying normal modes that serve as the leading-order basis for wave-wave interactions. While foundational studies have established that variable-density vortices are susceptible to intrinsic instabilities (Joly et al., 2005; Saunders, 1973; Sipp et al., 2005), the primary consequence for our perturbation framework is the modification of the resonant participants themselves. For instance, Chang and Llewellyn Smith (2021) investigated the linear stability of a strained vortex featuring a density jump at its interface. Their formulation reveals that the underlying mechanism perfectly aligns with our unified triadic resonance configurations, specifically manifesting as a special subcase of the elliptical instability (or, the broader parametric instability). They found that the introduction of the density ratio fundamentally modifies the inertial properties of the interacting waves, which in turn alters the growth rates of the elliptical instability. Because their study remains anchored to the specific Moore–Saffman–Tsai–Widnall configuration (see §5), evaluating these two-layer, stratified constructions through the broader lens of generalized triadic resonance and the pseudoenergy criterion presents a highly compelling avenue for future research. Thus, even without considering internal singularities, stratification actively tunes the efficacy of established wave-wave resonances.

At the second level, introducing a density gradient exerts a profound effect on the wave dynamics near the critical layers. In the context of aircraft wake vortices, phenomena such as exhaust entrainment or wingtip heating can introduce *radial* density variations. When a radial density step is introduced, the critical layer ($\Phi(r_c) = 0$) serves as a crucial locus for instability. As explicitly demonstrated by Dixit and Govindarajan (2011), positioning a radial density interface (at a radius r_j) triggers a powerful two-wave resonance. To isolate the physical mechanism, they plotted the dispersion curves with respect to r_j by uncoupling the effects, separating a homogeneous Rankine core from an isolated circular density interface. They showed that an instability emerges precisely where two otherwise neutral branches intersect: a density wave propagating on the interface, and a Kelvin mode of opposite pseudoenergy propagating on the vortex core. When the density step r_j is placed near the critical radius, these two dispersion curves cross, causing the modes to phase-lock and form an instability. While Dixit and Govindarajan (2011)’s formulation elegantly justifies the pseudoenergy criterion for a two-wave resonance, it naturally motivates a more comprehensive inquiry. A direction for future AWV research is to systematically investigate generalized *triadic* resonance under the fully coupled effects of background shear and *continuously* varying radial density profiles. Exploring how these complex dispersion curves interact to generate new resonant pathways could provide the theoretical blueprint for degrading the AWV structure significantly faster than natural viscous diffusion.

While radial stratification offers a promising theoretical pathway for aeronautical flow

control, the geometric alignment of the stratification dictates the specific mathematical nature of the critical-layer resonance. In geophysical and astrophysical flows, background stratification is typically aligned along the axis of rotation rather than the radial direction. This geometric distinction fundamentally changes the singularity: axial stratification supports the propagation of three-dimensional internal gravity waves, giving rise to distinct *baroclinic critical layers* located exactly where the Doppler-shifted frequency matches the axial buoyancy frequency. At these locations, resonant disturbances tap directly into the potential energy of the axial stratification. These specific baroclinic critical layers serve as the fundamental engine driving the Zombie Vortex Instability (ZVI) (Barranco, Pei, et al., 2018; Philip S. Marcus, Pei, Jiang, and Barranco, 2016; Philip S. Marcus, Pei, Jiang, Barranco, et al., 2015; Philip S. Marcus, Pei, Jiang, and Hassanzadeh, 2013), with relevant nonlinear critical layer theory examined by C. Wang and Balmforth (2020, 2021). Having established the rigorous constraints of unstratified resonance in Part I, Part II of this dissertation will pivot to this stratified astrophysical regime, exploring the properties of these axial baroclinic critical layers and their profound role in the stability of rotating shear flows.

6.3 Conclusion

In this part of the dissertation, we have unified the theory of vortex instability under the framework of hydrodynamic “selection rules”. By performing a multi-scale perturbation analysis, we derived the general amplitude equations for triadic resonance of columnar vortices in incompressible flows. We established that the resonant interaction between smooth neutral wave modes, specifically regular Kelvin waves and discrete critical layer modes with passive singularities, is inherently conservative. Consequently, the dynamics of these triads are governed by the Manley–Rowe relations, which constrain their energy exchange to compact manifolds.

To determine the stability of these conservative interactions, we employed the concept of wave pseudoenergy within a large- k WKB framework by Stéphane Le Dizès and Laurent Lacaze (2005). We demonstrated that the sign of pseudoenergy is invariant under a translating frame, and critical layer modes are the sole carriers of negative pseudoenergy. This leads to a set of selection rules that topologically prohibit the formation of explosive resonant triads. This finding provides a theoretical explanation for the robustness of isolated columnar vortices, showing that intrinsic instability is not a generic property of resonance but is suppressed by the underlying Hamiltonian structure of the flow.

Despite the prohibition of explosive resonance, we showed that the conservative framework admits *parametric instability* when the system is driven by external forcing. In this scenario, the external pump acts as an energy reservoir that bypasses the Manley–Rowe constraints. To explore this, we developed a robust and efficient tuning method based on non-degenerate perturbation theory, capable of locating resonant configurations for arbitrary pumping frequencies and wavenumbers. We successfully validated this method on the Lamb–Oseen vortex, reproducing classical results for elliptical instability and identifying a broad

class of new parametric instabilities involving critical layer modes.

We further identified *active critical layers* as an alternative pathway to bypass the selection rules. Unlike smooth neutral modes, active critical layers introduce non-Hermitian effects that allow the disturbance to extract energy directly from the background shear, potentially transforming the conservative system into an explosive one. These findings suggest that the mitigation of aircraft wake vortices requires specific symmetry-breaking mechanisms: either tuned external forcing to trigger parametric instability or the engineering of critical layers (e.g., via thermal stratification) to activate these forbidden transitions. Part II of this dissertation will explore the baroclinic critical layers.

Finally, while our analysis has focused on the inviscid regime, the broader framework represented by (3.24) can incorporate the influence of viscosity, provided that the leading-order mode structures align with their inviscid counterparts. Extensive discussions on conservative three-wave resonance involving linearly damped modes in plasma physics and hydrodynamics are available in Weiland and Wilhelmsson (1977) and Craik (1986), and a separate investigation into the case of columnar vortices is expected in the future. Moreover, the possibility of explosive triads persists in non-conservative scenarios. As critical layer modes can be regularized by both viscosity (S. Lee and P. S. Marcus, 2023) and nonlinearity, the latter of which naturally exists in triadic resonance, we anticipate that they will hold significance in non-conservative resonance. A comprehensive numerical analysis of resonance involving regularized critical layer modes in incompressible flows is currently being conducted at the time of this writing.

Part II

Stratified Flows: Global Wave-Mean Resonances and Numerical Simulations

Chapter 7

Stability of Rotating, Stratified, Shear Flows

7.1 Motivation and Background

In geophysical and astrophysical systems, ranging from planetary atmospheres and terrestrial oceans to stellar interiors and protoplanetary accretion disks, fluid dynamics are governed by the simultaneous coexistence of background rotation, stable density stratification, and differential shear (Armitage, 2011; Barranco and Philip S. Marcus, 2005). Understanding the stability of these environments is a critical open problem in fluid mechanics, primarily because it dictates energy dissipation and mass transport. In protoplanetary disks, for example, gas and dust must continuously lose angular momentum to de-orbit and accrete onto the central protostar. Molecular viscosity is orders of magnitude too weak to account for the observed accretion rates. While the magnetorotational instability (MRI) can drive vigorous turbulence in highly ionized regions, disks harbor extensive “dead zones” where ionization is too low for the MRI to operate (Gammie, 1996). Consequently, the necessary de-orbiting and accretion in these regions must rely on purely hydrodynamic instabilities.

Historically, the existence of such purely hydrodynamic instabilities presented an apparent paradox, as the flow appears rigidly stabilized in both its radial and vertical dimensions. Radially, the primary balance between the central protostar’s gravity and the fluid’s centrifugal force dictates a Keplerian angular velocity profile ($\bar{\Omega} \propto r^{-3/2}$). Because the specific angular momentum of this profile increases outward, the orbits are classically deemed stable against axisymmetric perturbations according to Rayleigh’s centrifugal criterion. Vertically, stable density stratification, quantified by the axial Brunt-Väisälä (or buoyancy) frequency, $\bar{N}$, which represents the natural oscillation frequency of a displaced fluid parcel, suppresses simple overturning motions. As long as the fluid is stably stratified ($\bar{N}^2 > 0$), any vertical perturbation is countered by a strong restorative buoyancy force. Together, these two independent restoring forces seemingly lock the fluid into a deeply stable state.

However, this classical stability paradox has been debunked because it relies on a

fundamental oversimplification: criteria like Rayleigh’s evaluate rotation and stratification as isolated, one-dimensional stabilizing forces, often assuming constant density or non-stratified parallel shear. When the full thermodynamic interplay of variable density, rotation, and shear is properly coupled, the very stratification that suppresses vertical overturning introduces new baroclinic wave pathways that circumvent classical unstratified limits. For instance, the stratorotational instability (SRI) can destabilize a Rayleigh-stable flow, though it requires the presence of physical boundaries (Dubrulle, B. et al., 2005; Molemaker et al., 2001; Yavneh et al., 2001). In unbounded domains, the radiative instability can develop, though its analytical framework builds specifically on the assumption of strong stratification (Stéphane Le Dizès and Billant, 2009). Operating in a distinctly different parameter regime, the zombie vortex instability (ZVI) (Philip S. Marcus, Pei, Jiang, and Barranco, 2016; Philip S. Marcus, Pei, Jiang, Barranco, et al., 2015; Philip S. Marcus, Pei, Jiang, and Hassanzadeh, 2013; Wohlever et al., 2025) is a finite-amplitude, self-replicating nonlinear phenomenon driven by baroclinic critical layers that uniquely requires stratification, rotation, and shear to all be of the same order of magnitude.

To theoretically and numerically capture these complex wave-mean flow interactions in rotating, stratified shear flows, the vast majority of past studies, with few exceptions, have relied extensively on the local shearing box approximation. While this local model provides significant analytical and computational simplifications, it inherently alters the global geometry and curvature of the system. Therefore, to rigorously investigate the stability of these flows, it is critical to evaluate the limitations of the shearing box approximation and compare it directly against a global cylindrical framework, which is the focus of the following section.

7.2 The Local Shearing Box vs. Global Cylindrical Frame

The shearing box approximation serves as a powerful idealized model designed to isolate the local dynamics of a differentially rotating disk or vortex. It begins with the full equations of motion in a global rotating frame and considers a small Cartesian patch $(\tilde{x}, \tilde{y}, \tilde{z})$ co-rotating with the fluid at a specific fiducial radius r_0 . To construct the local flow field, the global cylindrical velocity profile is Taylor-expanded around this central point. The zeroth-order uniform advection term is removed by transitioning to the local co-rotating frame, and only the first-order term is retained. This mathematical truncation results in a strictly linear shear flow relative to the center of the box: $\bar{u}_{\tilde{x}}(\tilde{y}) \approx -\sigma_s \tilde{y}$, where $\tilde{y} = r - r_0$ is the local radial coordinate, and the constant shear rate σ_s is defined by the velocity gradient of the global flow evaluated at r_0 .

The defining feature of this Cartesian idealization is that the vorticity of the background flow, $\bar{\omega}_{z,\text{box}} = -d\bar{u}_{\tilde{x}}/d\tilde{y} = \sigma_s$, is strictly constant throughout the entire computational domain. The shearing box is fundamentally a local asymptotic model, mathematically valid only in the limit of large radii where the global curvature effects (terms of order $\mathcal{O}(1/r)$)




Figure 7.1: Mapping between the global cylindrical coordinates (r, θ, z) and the local shearing box $(\tilde{x}, \tilde{y}, \tilde{z})$ centered at (r_0, θ_0) . The local coordinates are defined as $\tilde{y} = r - r_0$, $\tilde{x} = -r_0(\theta - \theta_0)$, and $\tilde{z} = z$, such that the local unit vectors correspond to $\tilde{\mathbf{y}} = \hat{\mathbf{r}}$ and $\tilde{\mathbf{x}} = -\hat{\boldsymbol{\theta}}$. Gravity acts in the $-z$ direction (pointing into the page). The local frame co-rotates with the background at an angular velocity $\bar{\Omega} = \bar{\Omega}(r_0) + \Omega$, where $\bar{\Omega}$ represents the radially dependent angular velocity of the global background flow and Ω is the constant background rotation rate of the global frame. The differential rotation of the background flow is locally approximated as a linear shear profile $\bar{u}_{\tilde{x}} = -\sigma_s \tilde{y}$, where the constant shear rate is defined as $\sigma_s = r_0(d\bar{\Omega}/dr)|_{r_0}$ ($\sigma_s < 0$ in the depicted example).

approach zero. It is within this specific, highly constrained environment that phenomena like the ZVI were originally modeled, leveraging the simplicity of periodic boundary conditions and uniform shear to track their temporal evolution (Barranco and Philip S. Marcus, 2006).

However, this Cartesian idealization possesses inherent mathematical and physical limitations that obscure the true global nature of wave-mean flow interactions (Regev and Umurhan, 2008). Specifically, the shearing box artificially decouples the local shear rate from the local background rotation, strictly enforces non-physical spatial periodicity, and imposes a linear background velocity profile that artificially restricts the macroscopic radial evolution of large-scale modes. Consequently, while instabilities like the ZVI clearly manifest within the shearing box, attributing their geometric characteristics, such as the uniform spatial spacing between vortex generations, to the artificial periodicity of the box is physically problematic.

To enforce proper geometric constraints and investigate whether such instabilities can be sustained physically outside idealized numerical setups, it is an absolute necessity to transition the study to a global cylindrical domain (Armitage, 1998; Ebrahimi and Blackman, 2016; Machida et al., 2000; Zhu et al., 2020). In a true global flow, the instability mechanisms of the ZVI are fundamentally linked to the existence of spatially localized critical layers. For instance, the condition for a baroclinic critical layer at a radius r_c dictates that the wave's intrinsic (Doppler-shifted) frequency must match the background buoyancy frequency: $\omega - m\bar{\Omega}(r_c) = \pm N$. Because the background angular velocity $\bar{\Omega}(r)$ varies continuously with radius in a global fluid, these critical layers naturally localize at specific, discrete radii without

requiring arbitrary domain boundaries to define their spacing.

By moving to a global cylindrical framework, we restore the natural coupling between shear, rotation, and curvature. This allows us to track the propagation of inertio-gravity waves across varying shear gradients and observe their localized trapping at critical layers. To maintain physical relevance and avoid the unphysical boundaries or divergent far-field velocities associated with unbounded power-law models, our global investigations will utilize generic, spatially bounded flow profiles, such as the Lamb–Oseen vortex. This transition guarantees that any identified resonance mechanisms or instability spacings are intrinsic hydrodynamic properties of the fluid, rather than geometric artifacts of the computational domain.

7.3 Overview

The remainder of Part II is dedicated to establishing the mathematical formulations, numerical schemes, and physical analyses necessary to study rotating, stratified shear flows globally. In Chapter 8, we lay out the 3D Boussinesq equations in cylindrical coordinates, define the exact energy conservation laws, and derive the linear stability equations that govern critical layer resonance. In Chapter 9, we analyze the severe numerical stiffness inherent to explicit time integration of these equations and formulate a highly-parallelized, semi-analytical pseudo-spectral Exponential Time Differencing (ETD) scheme to solve them efficiently. Finally, in Chapter 10, we deploy this advanced numerical framework to perform global nonlinear simulations, analyze baroclinic wave trapping, and formally present our frequency-resetting chain reaction hypothesis to explain the global evolution of the Zombie Vortex Instability.

Chapter 8

Mathematical Formulations

While Part I of this thesis focused on wave-wave interactions and triadic resonance within incompressible flows, Part II investigates wave-mean flow interactions in stratified environments. To explore how infinitesimal perturbations extract energy from a background state, this chapter establishes the mathematical framework governing the dynamics of rotating, stratified shear flows. The objective is to construct a theoretical bedrock that bridges the continuous three-dimensional physical equations to both the localized stability analysis and the direct numerical simulations developed in the subsequent chapters.

The chapter is structured to systematically analyze the governing equations in their primitive form before reformulating them for computational integration. We begin in §8.1 by formulating the nonlinear Boussinesq equations in an unbounded global cylindrical domain. Building directly upon these primitive equations, §8.2 derives the global conservation laws for mechanical energy and axial angular momentum. These integral balances are essential for identifying the Reynolds stress and buoyancy flux mechanisms that sustain hydrodynamic instabilities, and they provide the analytical constraints needed to validate the numerical schemes developed in Chapter 9.

While the conservation laws identify the macroscopic sources of instability, they do not explain the localized mechanics of energy extraction. Therefore, in §8.3, the linearized governing equations are reduced to a one-dimensional eigenvalue problem to conduct a normal mode stability analysis. This serves multiple purposes. First, it introduces the linear wave modes that characterize the disturbance away from singularities such as the critical layers, and utilizes a short-wavelength WKBJ approximation to provide insight into their spatial structure. Second, it formally defines critical layers as localized regions of wave-mean flow resonance and establishes their roles in wave dynamics. By demonstrating how these critical layers trap propagating waves and actively extract energy from the background shear, this section outlines the physical mechanisms relevant to the Zombie Vortex Instability (ZVI).

Finally, in §8.4, we transition from theoretical analysis to computational preparation. Because direct numerical integration of the primitive equations is complicated by the implicit pressure gradient and the divergence-free constraint, we introduce the poloidal-toroidal de-

composition. This mathematical transformation analytically eliminates the pressure variable and reformulates the solenoidal vector field into scalar potentials, which is utilized by the numerical schemes developed in Chapter 9 as an efficient projection filter and pressure solver during numerical integration.

8.1 The Boussinesq Equations for Rotating Stratified Shear Flows

We consider the dynamics of a viscous, stratified fluid within an unbounded global cylindrical domain defined by coordinates (r, θ, z) with unit vectors $(\hat{\mathbf{r}}, \hat{\boldsymbol{\theta}}, \hat{\mathbf{z}})$, where $0 \leq r < \infty$. The fluid is subject to a constant gravitational acceleration $\mathbf{g} = -g\hat{\mathbf{z}}$, and the system is observed in a frame of reference rotating about the vertical axis with constant angular velocity: $\boldsymbol{\Omega} = \Omega\hat{\mathbf{z}}$, as depicted in figure 7.1. The primitive state variables of the fluid are its velocity $\mathbf{u}$, gauge pressure p , and rescaled buoyancy $b \equiv \rho g / \rho_0$, where ρ is the fluid density. The latter two are defined relative to a hydrostatic reference state of constant density ρ_0 .

As the primary focus of Part II is the interaction between wave perturbations and the mean flow, we decompose the total flow state into a steady, axisymmetric background component, denoted by overbars, and a time-dependent disturbance component, denoted by primes (e.g., $p = \bar{p} + p'$). The background state is considered to be an azimuthal shear flow with radial dependence,

$$\bar{\mathbf{u}} \equiv \bar{u}_\theta(r)\hat{\boldsymbol{\theta}}, \quad (8.1)$$

whose angular velocity $\bar{\Omega}$ and axial vorticity $\bar{\zeta}$ can be expressed as

$$\bar{\Omega}(r) = \frac{\bar{u}_\theta}{r}, \quad (8.2a)$$

$$\bar{\zeta}(r) = \frac{1}{r} \frac{d(r\bar{u}_\theta)}{dr} = 2\bar{\Omega} + r \frac{d\bar{\Omega}}{dr}. \quad (8.2b)$$

The background buoyancy is assumed to depend solely on the vertical coordinate and follow a linear profile, characterized by a constant Brunt-Väisälä frequency

$$\bar{N} \equiv \sqrt{-\frac{d}{dz}\bar{b}(z)}. \quad (8.3)$$

We further require the disturbance component to remain analytic at the origin ($r = 0$) and to decay sufficiently rapidly to zero in the far field ($r \rightarrow \infty$), ensuring that it is physically confined and that both its total mass anomaly and energy remain finite. A more detailed discussion on these pole conditions is presented in §8.4.

We adopt the Boussinesq approximation to govern the flow dynamics, assuming that density variations are sufficiently small relative to ρ_0 that they are negligible in the inertial terms but retained in the gravitational force. Under this approximation, the specified

background state constitutes an equilibrium solution in the inviscid limit, provided the background pressure field $\bar{p}(r, z)$ satisfies the hydrostatic and cyclogeostrophic balances:

$$\partial_r \bar{p} = \rho_0(\bar{\Omega} + 2\Omega)\bar{u}_\theta, \quad (8.4a)$$

$$\partial_\theta \bar{p} = 0, \quad (8.4b)$$

$$\partial_z \bar{p} = -\bar{\rho}g, \quad (8.4c)$$

where the radial pressure gradient provides the necessary centripetal force to balance the Coriolis and inertial effects of the background rotation, while the vertical gradient balances the static stratification. To maintain the steady-state assumption inherent in our flow state decomposition, we neglect the slow viscous diffusion of the background flow, assuming it occurs on a timescale much longer than the transient wave dynamics of interest. This physical justification ensures that the background profile remains fixed, preventing non-physical drift, and implies that the viscous and diffusive operators in the governing equations act exclusively on the perturbation fields.

The governing equations of the system are obtained by removing the background equilibrium from the full Boussinesq equations. Noting that the advection of the linear background stratification simplifies to $(\mathbf{u} \cdot \nabla)\bar{b} = -u'_z \bar{N}^2$, the governing equations for the disturbance are written in the vorticity form:

$$\nabla \cdot \mathbf{u}' = 0, \quad (8.5a)$$

$$\partial_t \mathbf{u}' = \mathbf{u}' \times \boldsymbol{\omega}' - 2\Omega \hat{\mathbf{z}} \times \mathbf{u}' - b' \hat{\mathbf{z}} - (\mathbf{u}' \cdot \nabla) \bar{\mathbf{u}} - (\bar{\mathbf{u}} \cdot \nabla) \mathbf{u}' + \nu \nabla^2 \mathbf{u}' - \nabla \Pi, \quad (8.5b)$$

$$\partial_t b' = -(\mathbf{u}' \cdot \nabla) b' + u'_z \bar{N}^2 - (\bar{\mathbf{u}} \cdot \nabla) b' + \kappa \nabla^2 b', \quad (8.5c)$$

where $\boldsymbol{\omega}' = \nabla \times \mathbf{u}'$, ν and κ are the kinematic viscosity and thermal diffusivity of the fluid, respectively, and $2\Omega \hat{\mathbf{z}} \times \mathbf{u}'$ in (8.5b) represents the Coriolis force acting on the disturbance velocity. Here, Π is an effective pressure potential that absorbs the kinematic pressure (p'/ρ_0), the centrifugal force, and the scalar potential associated with the vector identity $(\mathbf{u}' \cdot \nabla) \mathbf{u}' = \nabla(|\mathbf{u}'|^2/2) - \mathbf{u}' \times \boldsymbol{\omega}'$.

It is important to emphasize that to maintain mathematical consistency, the two advective cross terms between the mean flow and the disturbance must share the same advection or vorticity form:

$$\begin{aligned} & -(\mathbf{u}' \cdot \nabla) \bar{\mathbf{u}} - (\bar{\mathbf{u}} \cdot \nabla) \mathbf{u}' \\ &= \frac{1}{2} [\nabla \times (\bar{\mathbf{u}} \times \mathbf{u}') + \mathbf{u}' \times \bar{\zeta} \hat{\mathbf{z}} + \bar{\mathbf{u}} \times \boldsymbol{\omega}'] + \frac{1}{2} [\nabla \times (\mathbf{u}' \times \bar{\mathbf{u}}) + \bar{\mathbf{u}} \times \boldsymbol{\omega}' + \mathbf{u}' \times \bar{\zeta} \hat{\mathbf{z}}] \quad (8.6) \\ &= \mathbf{u}' \times \bar{\zeta} \hat{\mathbf{z}} + \bar{\mathbf{u}} \times \boldsymbol{\omega}', \end{aligned}$$

which requires the cancellation of the two curl operations in each cross term of the intermediate step in (8.6) to achieve equality.

The Boussinesq equations (8.5) can be non-dimensionalized using the reference density ρ_0 , a characteristic velocity U_0 , and a length scale L_0 , where the choices of the latter

two characteristic scales depend on the particular flows of interest. This gives a set of dimensionless parameters, which are the Rossby number (Ro), the Reynolds number (Re), and the horizontal Froude number (Fr):

$$Ro \equiv \frac{U_0}{2\Omega \cdot L_0}, \quad Re \equiv \frac{U_0 \cdot L_0}{\nu}, \quad Fr \equiv \frac{U_0}{L_0 \cdot \bar{N}}, \quad (8.7a,b,c)$$

where Fr is the ratio between the advective timescale, L_0/U_0 , and the characteristic timescale of the internal gravity waves, $1/\bar{N}$, reflecting the relative strength between inertial forces and the restoring buoyancy force.

8.2 Conservation Laws

Having established the governing equations in primitive variables, we now examine their global energetic properties. To establish a framework for physical interpretation as well as the validation of our numerical schemes, we derive the conservation laws for the Boussinesq equations (8.5).

We define the total kinetic energy ($E_{K,tot}$), disturbance kinetic energy (E_K), and available potential energy (E_{AP}) as:

$$E_{K,tot} \equiv \frac{1}{2} \int_V |\mathbf{u}|^2 dV, \quad (8.8a)$$

$$E_K \equiv \frac{1}{2} \int_V |\mathbf{u}'|^2 dV, \quad (8.8b)$$

$$E_{AP} \equiv \frac{1}{2\bar{N}^2} \int_V (b')^2 dV. \quad (8.8c)$$

The volume integration extends over the unbounded cylindrical domain ($0 \leq r < \infty$) that is periodic in the azimuthal (θ) and axial (z) directions. The assumption of axial periodicity is physically justified by considering a domain length sufficient to isolate the transient dynamics of interest from vertical boundary effects. We note that while the background buoyancy $\bar{b}(z)$ is linear and non-periodic, it is structurally decoupled from the disturbance equations (8.5) and is unaffected by this assumption. Furthermore, we enforce the boundary condition that all localized disturbance fields, such as the perturbation velocity $\mathbf{u}'$ and buoyancy b' , decay to zero as $r \rightarrow \infty$. While theoretical background shear profiles analogous to the local shearing box (e.g., constant vorticity profiles) may possess infinite total energy when integrated over an unbounded radial domain, finite and rigorous conservation laws can still be constructed by focusing exclusively on the bounded disturbance fields.

To establish the theoretical framework, we first consider a generalized, fully bounded velocity field $\mathbf{u}$ and buoyancy b that decay at infinity. For clarity, our derivations here omit the viscous and diffusive terms, which will be reintroduced in the final energy budgets as dissipative sinks. Taking the inner product of the inviscid momentum equation with $\mathbf{u}$ yields

the evolution of the kinetic energy density. Noting that the Coriolis force does no mechanical work ($\mathbf{u} \cdot (2\Omega \hat{\mathbf{z}} \times \mathbf{u}) = 0$) and the nonlinear advection of the total kinetic energy can be expressed as a divergence, we obtain:

$$\begin{aligned} \partial_t \left(\frac{1}{2} \mathbf{u}^2 \right) &= -\mathbf{u} \cdot \nabla \left(\frac{p}{\rho_0} + \frac{1}{2} \mathbf{u}^2 \right) - bu_z \\ &= -\nabla \cdot \left[\left(\frac{p}{\rho_0} + \frac{1}{2} \mathbf{u}^2 \right) \mathbf{u} \right] - bu_z, \end{aligned} \quad (8.9)$$

where we have utilized the divergence-free condition $\nabla \cdot \mathbf{u} = 0$. Integrating over the domain volume V , the divergence term evaluates to zero by Gauss's theorem due to the periodic and radially decaying boundary conditions.

For a linearly stratified fluid, the total buoyancy is $b = \bar{b}(z) + b'$, where the background profile is $\bar{b}(z) = -\bar{N}^2 z + C$. The work done by this background stratification can be rewritten as a divergence:

$$\bar{b}u_z = (-\bar{N}^2 z + C)u_z = \nabla \cdot \left[\left(-\frac{1}{2} \bar{N}^2 z^2 + Cz \right) \mathbf{u} \right]. \quad (8.10)$$

Integrating this divergence over the domain volume reduces it to boundary fluxes via Gauss's theorem. Because the background shear flow is azimuthal, the total radial and vertical velocities are identical to the localized disturbance velocities ($u_r = u'_r$, $u_z = u'_z$), which decay to zero in the far field. Consequently, the radial boundary flux vanishes. Furthermore, incompressibility and radial decay dictate that the net vertical mass flux across any horizontal plane is zero ($\iint u'_z r dr d\theta = 0$), meaning the vertical boundary fluxes also evaluate to zero. The volume integral vanishes despite the explicit non-periodic presence of z^2 . This leaves the time rate of change of the generalized kinetic energy dependent solely on the bounded buoyancy perturbation:

$$\frac{d}{dt} E_{K,\text{tot}} = \partial_t \int_V \frac{1}{2} \mathbf{u}^2 dV = - \int_V (b' u_z) dV. \quad (8.11)$$

Similarly, the time rate of change of the available potential energy (E_{AP}) associated with the buoyancy perturbation is calculated via the thermodynamic equation:

$$\begin{aligned} \frac{d}{dt} E_{AP} &= \frac{1}{\bar{N}^2} \int_V b' \partial_t b' dV \\ &= -\frac{1}{\bar{N}^2} \int_V \nabla \cdot \left(\frac{1}{2} b'^2 \mathbf{u} \right) dV + \int_V (b' u_z) dV \\ &= \int_V (b' u_z) dV. \end{aligned} \quad (8.12)$$

Summing these two energy components reveals the exact conservation of total available energy,

$$\frac{d}{dt} (E_K + E_{AP}) = 0, \quad (8.13)$$

for an isolated wave or generalized decaying field in the diffusion-free limit.

When analyzing the actual disturbance fields $(\mathbf{u}', b')$ superimposed on a steady background shear $\bar{\mathbf{u}}$, the total disturbance energy is no longer conserved due to the continuous exchange of momentum with the base flow. The evolution of the disturbance kinetic energy evaluates to:

$$\begin{aligned} \frac{d}{dt} E_K &= \int_V (-u'_z b' - \mathbf{u}' \cdot ((\bar{\mathbf{u}} \cdot \nabla) \mathbf{u}' + (\mathbf{u}' \cdot \nabla) \bar{\mathbf{u}})) dV \\ &= \int_V \left(-u'_z b' - \left(r \frac{d\bar{\Omega}}{dr} \right) u'_r u'_\theta \right) dV. \end{aligned} \quad (8.14)$$

The final term represents the shear production, which dictates the transfer of kinetic energy between the background shear and the disturbance. Since the background field is purely two-dimensional, $u_z = u'_z$, and the evolution of the disturbance available potential energy remains identical to (8.12).

Re-incorporating the viscous and diffusive terms from the Boussinesq equations, the resulting complete energy budget for the disturbance field is governed by:

$$\frac{d}{dt} E_K = -\mathcal{E}_{\text{exc}} - \mathcal{E}_{\text{shear}} - \mathcal{E}_{\text{visc}}, \quad (8.15a)$$

$$\frac{d}{dt} E_{\text{AP}} = \mathcal{E}_{\text{exc}} - \mathcal{E}_{\text{diff}}, \quad (8.15b)$$

$$\frac{d}{dt} (E_K + E_{\text{AP}}) = -\mathcal{E}_{\text{shear}} - \mathcal{E}_{\text{visc}} - \mathcal{E}_{\text{diff}}, \quad (8.15c)$$

where the energy exchange, shear production, and dissipation terms are defined as:

$$\mathcal{E}_{\text{exc}} = \int_V (b' \cdot u'_z) dV, \quad (8.16a)$$

$$\mathcal{E}_{\text{shear}} = \int_V \left(r \frac{d}{dr} \bar{\Omega} \right) u'_r u'_\theta dV, \quad (8.16b)$$

$$\mathcal{E}_{\text{visc}} = \nu \int_V |\nabla \mathbf{u}'|^2 dV, \quad (8.16c)$$

$$\mathcal{E}_{\text{diff}} = \frac{\kappa}{N^2} \int_V |\nabla b'|^2 dV. \quad (8.16d)$$

In the absence of diffusion, the total disturbance energy $E_K + E_{\text{AP}}$ evolves solely due to the shear production term $\mathcal{E}_{\text{shear}}$. Note that $\mathcal{E}_{\text{shear}}$ only acts on E_K because the azimuthal shear is orthogonal to the direction of stratification. In contrast, the buoyancy flux $\mathcal{E}_{\text{exc}}$ appears with opposite signs in the E_K and E_{AP} equations. It serves as an internal transfer mechanism, driving the reversible conversion between kinetic and potential energy via vertical fluid motion. If the energetic scope is expanded to consider the total flow field, encompassing both the steady background state and the transient disturbance, both of these interaction terms become internal to the system and do not appear as external sources or sinks. This guarantees that the total energy of the full system is conserved in the diffusion-free limit.

Lastly, for an ideal inviscid fluid, the total axial angular momentum of the disturbance,

$$\mathcal{L}_z = \int_V r u'_\theta dV, \quad (8.17)$$

is a strictly conserved quantity. Integrating the azimuthal component of the disturbance momentum equation multiplied by r gives:

$$\frac{d}{dt} \mathcal{L}_z = - \int_V r [(\mathbf{u}' \cdot \nabla) \mathbf{u}' + (\bar{\mathbf{u}} \cdot \nabla) \mathbf{u}' + (\mathbf{u}' \cdot \nabla) \bar{\mathbf{u}}]_\theta dV - \int_V 2\Omega r u'_r dV - \int_V \frac{\partial p'}{\partial \theta} dV. \quad (8.18)$$

Each term on the right-hand side integrates identically to zero. The pressure gradient vanishes because it is the azimuthal derivative of a single-valued function evaluated over a 2π -periodic domain. By utilizing the divergence-free condition, the nonlinear self-advection and linear background advection terms can be expanded into a combination of pure divergences and azimuthal derivatives: $\nabla \cdot (r u'_\theta \mathbf{u}') + \partial_\theta (\bar{u}_\theta u'_\theta) + r u'_r \bar{\zeta}(r)$, which evaluate to zero over the radially decaying and azimuthally and axially periodic boundaries. The remaining shear interaction term ($r u'_r \bar{\zeta}$) and the Coriolis term ($2\Omega r u'_r$) both vanish as a mathematical consequence of mass conservation. Because the disturbance is solenoidal, integrating the continuity equation over any cylindrical volume of arbitrary radius R demands that the net radial mass flux across that cylindrical surface must be zero: $\int_0^{2\pi} \int_0^{L_z} u'_r(R, \theta, z) R d\theta dz = 0$. Because the horizontal average of u'_r must be identically zero at every radius, any integral of the form $\int_V f(r) u'_r dV$ (including $2\Omega r u'_r$ and $r \bar{\zeta}(r) u'_r$) evaluates to zero. Therefore, the global conservation of the disturbance's total axial angular momentum is guaranteed.

Ill-posedness of the standard potential energy

For completeness, we note why the standard potential energy definition,

$$E_P \equiv - \int_V z b' dV, \quad (8.19)$$

is mathematically ill-formed and not utilized for the disturbance energy balance. If we attempt to derive the evolution of E_P by multiplying the inviscid disturbance buoyancy equation by $-z$ and integrating over the volume, it yields:

$$\begin{aligned} \frac{dE_P}{dt} &= \int_V z (\mathbf{u}' \cdot \nabla b') dV + \int_V z (\bar{\mathbf{u}} \cdot \nabla b') dV - \bar{N}^2 \int_V z u'_z dV \\ &= \int_V \nabla \cdot (z b' \mathbf{u}') dV - \int_V b' u'_z dV + \int_V \nabla \cdot (z b' \bar{\mathbf{u}}) dV - \bar{N}^2 \int_V z u'_z dV. \end{aligned} \quad (8.20)$$

By applying physical boundary conditions and mass conservation, three of these four terms evaluate cleanly. First, because the background flow $\bar{\mathbf{u}}$ is azimuthal, it generates zero flux across both the vertical and radial boundaries, meaning $\int_V \nabla \cdot (z b' \bar{\mathbf{u}}) dV = 0$. Second,

because incompressibility and radial decay dictate that the net vertical mass flux across any infinite horizontal plane must be zero ($\iint u'_z r dr d\theta = 0$), the fourth integral vanishes identically ($\int_V z u'_z dV = 0$). Third, the middle integral returns the reversible energy exchange term, $-\mathcal{E}_{\text{exc}}$.

However, the first divergence term does not vanish. While the disturbance fields are periodic, the advective flux vector $zb'\mathbf{u}'$ explicitly contains the non-periodic axial coordinate z . Evaluating its volume integral via Gauss's divergence theorem yields a net vertical boundary flux of $L_z \iint b'(r, \theta, L_z) u'_z(r, \theta, L_z) r dr d\theta$, which evaluates to a non-zero value because the disturbance buoyancy and vertical velocity are dynamically correlated variables. The presence of this unbalanced boundary flux causes the energy evolution to scale explicitly with the arbitrary periodic domain height L_z , rendering the standard potential energy mathematically ill-formed in an axially periodic domain¹. To maintain strict mathematical consistency, we must therefore mandate the use of the available potential energy (E_{AP}), which relies exclusively on bounded, periodic perturbation variables.

8.3 Linear Stability Equations and Critical Layer Singularities

While the energy budgets derived in §8.2 successfully quantify the macroscopic mechanisms by which a disturbance extracts energy from the mean flow (e.g., the shear production $\mathcal{E}_{\text{shear}}$), they do not explain the localized mechanics of how a specific propagating wave taps into these energetic reservoirs. To identify the physical phenomena driving wave-mean energy transfer, we now transition from global integral properties to localized wave dynamics via a formal linear stability analysis.

Linear stability equations

The linear stability problem based on the inviscid Boussinesq disturbance equations can be studied by normal mode analysis. The linearized equations are homogeneous in terms of variables t , θ , and z , so we can express the disturbance as wave modes of the following form:

$$\mathbf{u}' = \tilde{\mathbf{u}}(r) \exp\{i(m\theta + kz) + \sigma t\}, \quad (8.21a)$$

$$b' = \tilde{b}(r) \exp\{i(m\theta + kz) + \sigma t\}, \quad (8.21b)$$

$$p'/\rho_0 = \tilde{p}(r) \exp\{i(m\theta + kz) + \sigma t\}, \quad (8.21c)$$

¹For a truly unbounded domain in z with decaying boundary conditions that are sufficiently fast, the surface flux vanishes, and corresponding conservation laws can be constructed.

where m and k are the azimuthal and vertical wavenumbers, and $\sigma \equiv \lambda + i\omega$ is the complex growth rate. The linearized equations now become, term-by-term:

$$\hat{\mathbf{r}} : \quad \varphi \tilde{u}_r = 2(\bar{\Omega} + \Omega) \tilde{u}_\theta - \mathrm{d}_r \tilde{p}, \quad (8.22a)$$

$$\hat{\boldsymbol{\theta}} : \quad \varphi \tilde{u}_\theta = -(\bar{\zeta} + 2\Omega) \tilde{u}_r - \frac{im}{r} \tilde{p}, \quad (8.22b)$$

$$\hat{\mathbf{z}} : \quad \varphi \tilde{u}_z = -\tilde{b} - ik\tilde{p}, \quad (8.22c)$$

$$\tilde{b} : \quad \varphi \tilde{b} = \bar{N}^2 \tilde{u}_z, \quad (8.22d)$$

where $\bar{\Omega}$ and $\bar{\zeta}$ are the angular velocity and vertical vorticity of the baseflow, respectively, and $\varphi(r)$ is the local Doppler-shifted complex eigenvalue², defined as:

$$\varphi(r) \equiv \sigma + im\bar{\Omega}(r). \quad (8.23)$$

For a static background ($\bar{\Omega} = 0$), (8.22) reduces to standard inertial gravity wave equations, which are detailed in Appendix 8A.

The linearized disturbance equations can be further reduced to a single second-order ordinary differential equation in terms of the radial velocity component $\tilde{u}_r$. This reduction is achieved by systematically eliminating the other variables ($\tilde{u}_\theta$, $\tilde{u}_z$, $\tilde{b}$, and $\tilde{p}$) through algebraic manipulation of the linearized equations. In particular, we note that the disturbance density purely depends on $\tilde{u}_z$:

$$\tilde{b} = \frac{\bar{N}^2}{\varphi} \tilde{u}_z, \quad (8.24)$$

so (8.22c) can be rearranged to express $\tilde{u}_z$ in terms of $\tilde{p}$:

$$\tilde{u}_z = -ik \frac{\varphi}{\varphi^2 + \bar{N}^2} \tilde{p}. \quad (8.25)$$

We also note that mass conservation states:

$$\frac{im}{r} \tilde{u}_\theta + ik \tilde{u}_z + \frac{1}{r} \mathrm{d}_r [r \tilde{u}_r] = 0, \quad (8.26)$$

which gives

$$\tilde{u}_\theta = i \cdot \frac{k^2}{m} \frac{r\varphi}{\varphi^2 + \bar{N}^2} \tilde{p} + i \frac{1}{m} \mathrm{d}_r [r \tilde{u}_r]. \quad (8.27)$$

Using (8.22b), we express $\tilde{p}$ in terms of $\tilde{u}_r$:

$$\tilde{p} = \frac{r(\varphi^2 + \bar{N}^2)m}{(k^2 r^2 + m^2)\varphi^2 + m^2 \bar{N}^2} \left(-\frac{\varphi}{m} \mathrm{d}_r [r \tilde{u}_r] + i(\bar{\zeta} + 2\Omega) \tilde{u}_r \right). \quad (8.28)$$

²For a purely neutral wave, the local Doppler-shifted eigenvalue becomes purely imaginary: $\varphi(r) = i\Phi(r)$, where Φ is the real Doppler-shifted wave frequency as defined in (4.2).

Substituting this back yields $\tilde{u}_\theta$ strictly in terms of $\tilde{u}_r$. Finally, from (8.22a), we obtain the final differential equation entirely in terms of the radial velocity component. For notational compactness, we define the auxiliary variable

$$g(r) \equiv r\tilde{u}_r, \quad (8.29)$$

so that the final second-order linear ordinary differential equation governing the linear wave modes is³:

$$0 = g'' + \left\{ \frac{2\varphi\varphi'}{\varphi^2 + \bar{N}^2} + \frac{1}{r} - \frac{S'}{S} \right\} g' - \left\{ \frac{k^2 \cdot (\varphi^2 + \Delta(r))}{\varphi^2 + \bar{N}^2} + \frac{m^2}{r^2} + \frac{2im(\bar{\zeta} + 2\Omega)\varphi'}{r(\varphi^2 + \bar{N}^2)} + \frac{im[S\bar{\zeta}' - (\bar{\zeta} + 2\Omega)S']}{r\varphi S} \right\} g, \quad (8.30)$$

where

$$S(r) \equiv (k^2 r^2 + m^2)\varphi^2 + m^2 \bar{N}^2, \quad (8.31a)$$

$$\Delta(r) \equiv 2(\bar{\zeta} + 2\Omega)(\bar{\Omega} + \Omega). \quad (8.31b)$$

Here, $\Delta(r)$ is the generalized Rayleigh discriminant, which reduces to the classical Rayleigh discriminant in the unstratified limit ($\bar{N}^2 = 0$) and serves as a key parameter that affects the stability of the shear flow.

WKBJ approximation and spatial wave structure

To bridge the global cylindrical framework with local dynamics and to understand the spatial localization of the wave modes, we employ a short-wavelength WKBJ approximation. This method is mathematically justified when the radial wavelength of the wave is much shorter than the length scale over which the background flow properties vary, allowing the approximation $\frac{d}{dr} \rightarrow ik_r$. Applying this directly to the linearized primitive equations of motion yields an algebraic system, and the requirement for a non-trivial solution defines the local dispersion relation:

$$k_r^2 \varphi - k_r \left(m \frac{d\bar{\Omega}}{dr} \right) + \varphi \left[\frac{m^2}{r^2} + \frac{\varphi^2 + \Delta(r)}{\varphi^2 + \bar{N}^2} k^2 \right] = 0. \quad (8.32)$$

This local dispersion relation exhibits structural correspondence with the dispersion relation derived within the local shearing box approximation (Wohlever et al., 2025) with the

³The first-order coefficients originally contain terms: $im(2\bar{\Omega} - \bar{\zeta})/(r\varphi) + \varphi'/\varphi$. The derivative of the Doppler-shifted eigenvalue is exactly $d_r \varphi = im d_r \bar{\Omega}$, so the vorticity term reduces to a logarithmic derivative:

$$im(2\bar{\Omega} - \bar{\zeta})/(r\varphi) = -im d_r \bar{\Omega} / \varphi = -\varphi' / \varphi,$$

which perfectly cancels the $+\varphi'/\varphi$ term.

following parameter mapping:

$$f \Longleftrightarrow 2(\bar{\Omega}(r) + \Omega), \quad (8.33a)$$

$$k_x \Longleftrightarrow -\frac{m}{r}, \quad (8.33b)$$

$$\sigma_s \Longleftrightarrow r \frac{d\bar{\Omega}}{dr}, \quad (8.33c)$$

where f is the Coriolis parameter, k_x is the streamwise wavenumber, and σ_s is the constant shear rate in the shearing box approximation.

This generalized relation contains two crucial physical limits. First, in the axisymmetric limit ($m = 0$), the shear interaction term vanishes, and (8.32) simplifies to

$$k_r^2 + k^2 \frac{\sigma^2 + \Delta(r)}{\sigma^2 + \bar{N}^2} = 0, \quad (8.34)$$

yielding:

$$\sigma^2 = -\frac{k_r^2 \bar{N}^2 + k^2 \cdot \Delta(r)}{k_r^2 + k^2}. \quad (8.35)$$

For a locally-propagating mode ($k_r \in \mathbb{R}$), the numerator must be negative for a temporal instability to exist ($\sigma^2 > 0$). This analytically recovers the generalized Rayleigh criterion for centrifugal instability in an unstratified flow ($\bar{N} = 0$):

$$\Delta(r) < 0. \quad (8.36)$$

This confirms that centrifugal instability is natively captured by the formulation and remains a necessary but not sufficient condition for instability under stable stratification ($\bar{N}^2 > 0$).

Second, for non-axisymmetric, nearly neutral waves ($m \neq 0, \lambda = 0$), the complex Doppler-shifted frequency becomes purely imaginary: $\varphi(r) = i(m\bar{\Omega}(r) + \omega)$. Assuming the shear interaction term $m \frac{d}{dr} \bar{\Omega}$ is $\mathcal{O}(\delta)$ justifies an expansion of the radial wavenumber:

$$k_r = k_r^{(0)} + \delta k_r^{(1)}. \quad (8.37)$$

Separating the local dispersion relation by orders of δ yields the leading-order spatial structure:

$$\beta^2(r) \equiv (k_r^{(0)})^2 = -\frac{m^2}{r^2} - \frac{\varphi^2 + \Delta(r)}{\varphi^2 + \bar{N}^2} k^2, \quad (8.38)$$

where the real quantity β^2 dictates the WKB spatial envelope of the wave. Assuming strong stratification ($\bar{N}^2 \gg |\varphi^2|$), the wave exhibits oscillatory, freely propagating behavior in WKB regions where $\beta^2 > 0$ and decays exponentially in evanescent regions where $\beta^2 < 0$. The WKB turning points, r_t , which separate these regions, occur at the roots

$$\beta^2(r_t) = 0. \quad (8.39)$$

S. Le Dizès (2008) successfully utilized these specific turning points to construct approximated global discrete modes via asymptotic matching, demonstrating how turning points act as potential barriers bounding the wave structure.

Furthermore, the first-order $\mathcal{O}(\delta)$ correction from the expansion provides the small imaginary part of k_r :

$$\delta k_r^{(1)} = \frac{m}{2\varphi} \frac{d\bar{\Omega}}{dr}. \quad (8.40)$$

Because φ is purely imaginary, this term represents the local spatial exponential growth or decay of the wave's amplitude. The term $\delta k_r^{(1)}$ serves as the mathematical signature of the local energy exchange mechanism, demonstrating that amplitude modification is driven exclusively by the local shear gradient, $\frac{d}{dr}\bar{\Omega}$, and inherently vanishes for axisymmetric modes.

Critical layers

While the WKBJ analysis successfully maps the spatial envelopes of the linear waves, it naturally partitions the domain through two distinct geometric features: turning points and critical layers. At turning points, the local radial wavenumber vanishes ($k_r = 0$). These locations act as kinematic boundaries where waves reflect, separating regions of propagating and evanescent behavior. The governing differential equation remains perfectly regular at these locations; the wave solutions are well-established on either side of the turning point and merely require proper asymptotic connection formulas to continuously link the distinct wave regimes.

In contrast, critical layers represent localized regions of physical resonance and active energy exchange between the disturbance and the background state. In the idealized limit of inviscid, linear, neutral modes ($\lambda = 0$), this underlying physical resonance manifests mathematically as the roots of the leading-order coefficients of (8.30) falling precisely on the real spatial axis. The vanishing of the highest-order derivative coefficient at these resonant locations creates a formal singularity in the exact governing differential equation. However, it is crucial to recognize that this singularity is a result of the mathematical formulation in the linear limit rather than a physical infinite state. Various mechanisms naturally regularize the equations at these locations, such as incorporating viscosity, accounting for nonlinear effects, or recognizing the finite growth or decay of the modes ($\lambda \neq 0$), which shifts the mathematical poles into the complex plane. Regardless of the specific regularization applied to resolve the mathematics, the underlying physical resonance persists unaltered as the primary driver of wave dynamics within the critical layer.

Evaluating the roots of the leading order coefficients of (8.30) under the neutral wave condition reveals three distinct singularities:

1. The Baroclinic Critical Layer:

$$\varphi^2(r_c) + \bar{N}^2 = 0, \quad (8.41)$$

which occurs when the Doppler shifted wave frequency resonates with the natural buoyancy frequency ($\omega = -m\bar{\Omega}(r_c) \pm \bar{N}$). For stationary waves ($\omega = 0$), this critical layer locates at the radius where $\bar{\Omega}(r_c) = \pm \bar{N}/m$.

2. The Barotropic Critical Layer:

$$\varphi(r_c) = 0, \quad (8.42)$$

which represents co-rotation resonance, where the wave phase speed matches the local angular velocity of the background mean flow.

3. The Apparent Singularity:

$$S(r_s) = 0, \quad (8.43)$$

which arises from the root of the highest order derivative coefficient.

In contrast to the active critical layers, the apparent singularity at $S(r_s) = 0$ is an algebraic artifact inherent to the reduction of the governing equations. Rearranging its definition yields the local dispersion relation for a two-dimensional internal gravity wave:

$$\varphi^2(r) = -\bar{N}^2 \frac{m^2/r^2}{k^2 + m^2/r^2}. \quad (8.44)$$

This condition identifies a kinematic resonance where the frequency of the three-dimensional disturbance matches a localized two-dimensional shell mode. At this radius, the radial velocity of the perturbation vanishes ($\tilde{u}_r \rightarrow 0$), and the fluid motion becomes purely azimuthal and vertical. In the unreduced, three-dimensional linearized equations, this transition is perfectly analytic; the wave simply lacks radial momentum at r_s , but all perturbation variables remain bounded and smooth. The mathematical singularity emerges exclusively when attempting to condense the system into the single one-dimensional equation (8.30). As the location possesses no physical trapping mechanism or true wave-mean flow energy extraction, it cannot drive intrinsic instability. If an eigenvalue solver improperly evaluates this algebraic artifact and computes a fictitious temporal growth rate, it drives a spurious instability that carries no physical meaning and must be rigorously filtered out of the linear stability results⁴.

The critical layers, on the other hand, dictate the active energy dynamics of the instability. Specifically, as established by the WKBJ analysis, the baroclinic critical layer serves as a highly localized wave trapper. As $\varphi^2 \rightarrow -\bar{N}^2$, the local radial wavenumber $k_r \rightarrow \infty$, causing the radial group velocity of the internal gravity wave to stall. Energetically, this resonance forces the horizontal kinetic energy of the propagating wave to transfer into the vertical velocity component $\tilde{u}_z$, generating accumulations of available potential energy. This baroclinic trapping is the essential mechanism that arrests wave propagation and leads to the

⁴The eigenvalue solvers employed by this dissertation work are all constructed from full, unreduced primitive linearized equations, so this algebraic division by zero is fully bypassed.

nonlinear defect roll-up required by the Zombie Vortex Instability (C. Wang and Balmforth, 2020, 2021). Concurrently, the barotropic critical layer remains equally relevant to the global energy balance, acting as the active engine that extracts kinetic energy from the background shear, as discussed in Part I of this thesis.

To quantify how these localized resonances dictate the energy dynamics of the system, a spatial conservation law can be established. This is achieved by casting the linear ordinary differential equation into a self-adjoint form, a procedure that formally isolates the conserved wave action flux.

Because an integrating factor is defined up to an arbitrary multiplicative constant, we construct it directly as

$$P(r) = \exp \left(\int C_1 dr \right), \quad (8.45)$$

where C_1 is the coefficient of g' in (8.30). For constant $\bar{N}$, all terms in C_1 form exact logarithmic differentials, allowing the integrating factor to simplify to:

$$P(r) = \frac{r(\varphi^2 + \bar{N}^2)}{S(r)}. \quad (8.46)$$

Because φ is purely imaginary for neutral modes, φ^2 is strictly real. Consequently, $P(r)$ is a purely real function. Multiplying the original differential equation (8.30) by $P(r)$ yields the self-adjoint equation:

$$(P(r)g')' + Q(r)g = 0, \quad (8.47)$$

where $Q(r) \equiv P(r)C_2$ is the effective zeroth-derivative coefficient.

To evaluate the true nature of the critical layer singularities, we explicitly construct $Q(r)$. By substituting the geometric identity

$$\frac{\bar{\zeta}'}{r} \equiv \frac{d}{dr} \left(\frac{\zeta_a}{r} \right) + \frac{\zeta_a}{r^2}, \quad (8.48)$$

where $\zeta_a = \bar{\zeta} + 2\Omega$ is the absolute background vorticity, and applying the logarithmic derivative of the integrating factor P'/P , the final highly complex terms of C_2 mathematically condense into a single exact spatial derivative. Multiplying by $P(r)$, the effective coefficient reduces to:

$$Q(r) = -P(r) \left[\frac{k^2(\varphi^2 + \Delta(r))}{\varphi^2 + \bar{N}^2} + \frac{m^2}{r^2} \right] - \frac{im}{\varphi} \frac{d}{dr} \left(P(r) \frac{\bar{\zeta} + 2\Omega}{r} \right). \quad (8.49)$$

Because φ is purely imaginary, this condensed $Q(r)$ is strictly real. By multiplying (8.47) by the complex conjugate g^* , subtracting the complex conjugate of the resulting equation, and applying Abel's identity, the real $Q(r)$ terms perfectly cancel. The expression condenses to an invariant derivative:

$$\frac{d}{dr} \left[P(r)(g^*g' - gg^*) \right] = 0. \quad (8.50)$$

This allows us to formally define the conserved quantity $F(r)$ by isolating the phase shift required for radial wave propagation:

$$F(r) = \text{Im} (P(r)g^*g'). \quad (8.51)$$

The invariant $F(r)$ is not merely a mathematical abstraction; it represents the radial wave action flux of the perturbation. As discussed by Cheng (1975) and Kalashnik and Chkhetiani (2017), for inertial gravity waves in stratified shear flows, the imaginary part of the self-adjoint Wronskian is directly proportional to the physical Eliassen-Palm flux normalized by the intrinsic Doppler-shifted frequency. In regions free of singularities, $dF/dr = 0$, mathematically proving that the radial wave action flux is strictly conserved along the continuous medium. A smooth wave profile cannot intrinsically amplify or decay; it merely propagates or reflects. The only mechanism capable of breaking this conservative bound, and thereby serving as an active engine for instability, is the presence of a mathematical singularity, which introduces a logarithmic branch cut that yields a discontinuous jump in the wave action flux ($\Delta F \neq 0$).

The Frobenius analysis of the self-adjoint equation distinguishes the two active critical layers by their respective impacts on the wave action flux. At the barotropic critical layer ($\varphi = 0$), the integrating factor $P(r)$ is smooth and finite, while $Q(r)$ dictates a simple pole. This yields indicial roots of 0 and 1, generating a fundamental wave solution with a linear logarithmic branch cut ($g \sim A + B(r - r_c) \ln(r - r_c)$). Because the spatial derivative g' diverges logarithmically, the imaginary part of g^*g' is abruptly altered, yielding a mathematically discontinuous jump in the wave action flux ($\Delta F \propto \pm\pi|A|^2$) that signifies a direct exchange of kinetic energy with the background shear via over-reflection. At the baroclinic critical layer ($\varphi^2 + \bar{N}^2 = 0$), the mathematical structure is distinct. Here, the vanishing of the coefficient $P(r)$ corresponds to the $k_r \rightarrow \infty$ behavior established in the earlier local analysis. While the barotropic resonance acts as an active energy engine, this mathematical feature reinforces the role of the baroclinic critical layer as a localized wave trapper, where the kinematic stalling of the radial group velocity facilitates the secular accumulation of energy.

8.4 Poloidal-Toroidal Decomposition

The governing Boussinesq equations (8.5) are formulated in terms of the primitive perturbation variables ($p', b', \mathbf{u}'$). Because the velocity field is strictly solenoidal according to (8.5a), we can leverage a poloidal-toroidal decomposition (Chandrasekhar, 1981; Sangjoon Lee and Philip S. Marcus, 2023; Matsushima and P. Marcus, 1997; J. Wang et al., 2026) to enforce the divergence-free constraint and analytically eliminate the pressure variable.

By selecting the axial unit vector $\hat{\mathbf{z}}$ as the reference vector, any generic solenoidal vector field $\mathbf{A}(r, \theta, z)$ can be expressed in terms of two scalar potentials:

$$\mathbf{A} = \nabla \times (\psi \hat{\mathbf{z}}) + \nabla \times [\nabla \times (\chi \hat{\mathbf{z}})], \quad (8.52)$$

where $\psi(r, \theta, z)$ and $\chi(r, \theta, z)$ are the toroidal and poloidal streamfunctions, respectively. Taking the axial component and the axial curl of this decomposition yields a pair of two-dimensional Poisson equations:

$$-\nabla_{\perp}^2 \chi = \hat{\mathbf{z}} \cdot \mathbf{A}, \quad (8.53a)$$

$$-\nabla_{\perp}^2 \psi = \hat{\mathbf{z}} \cdot (\nabla \times \mathbf{A}), \quad (8.53b)$$

where the transverse Laplacian operator $\nabla_{\perp}^2$ is defined as

$$\nabla_{\perp}^2 \equiv \frac{1}{r} \frac{\partial}{\partial r} r \frac{\partial}{\partial r} - \frac{1}{r^2} \frac{\partial^2}{\partial \theta^2}. \quad (8.54)$$

To uniquely determine ψ and χ from a known vector field, one must invert these Poisson equations, which can be obtained by a two-dimensional convolution with the free-space Green's function, $G(r, \theta) = \ln r / (2\pi)$. As discussed by Sangjoon Lee and Philip S. Marcus (2023), this convolution is mathematically well-defined as long as the vector components decay algebraically or faster in the far field ($r \rightarrow \infty$), bounded by:

$$A_r \sim \mathcal{O}(r^{-1-p_1}), \quad A_{\theta} \sim \mathcal{O}(r^{-1-p_2}), \quad A_z \sim \mathcal{O}(r^{-2-p_3}), \quad (8.55a,b,c)$$

for constants $p_1, p_2, p_3 \geq 0$. If $\mathbf{A}$ represents a velocity field, these bounds guarantee that its kinetic energy density yields a finite total energy over the unbounded domain, which naturally satisfies our physical assumptions for the disturbance field. Furthermore, because the vector field is constructed purely from spatial derivatives of the poloidal-toroidal potentials, any constant offset is physically inconsequential. Hence, the remaining gauge freedom associated with the Laplacian inversion can be eliminated by imposing the following boundary conditions:

$$\lim_{r \rightarrow \infty} \psi(r, \theta, z) = 0, \quad (8.56a)$$

$$\lim_{r \rightarrow \infty} \chi(r, \theta, z) = 0. \quad (8.56b)$$

This establishes a linear and invertible projection operator, $\mathbb{P}_{\text{sol}}$, mapping the space of compliant solenoidal vector fields to their poloidal-toroidal components:

$$\mathbb{P}_{\text{sol}}[\mathbf{A}] \equiv \begin{bmatrix} \psi \\ \chi \end{bmatrix}. \quad (8.57)$$

In cylindrical coordinates, the spatial curl operators commute with both the full three-dimensional Laplacian (∇^2) and the transverse Laplacian ($\nabla_{\perp}^2$). This property implies that the linear viscous diffusion operator in the governing Boussinesq equations can be applied directly to the scalar potentials rather than operating on the primitive velocity vectors. Moreover, it allows us to extend our projection operator to a general projection filter, $\mathbb{P}$, by invoking the Helmholtz decomposition. The Helmholtz theorem states that any arbitrary

vector field $\mathbf{F}$ can be written as the sum of a solenoidal field and an irrotational scalar gradient ($\mathbf{F} = \mathbf{A} + \nabla \Pi$). The extended projection operator isolates the poloidal-toroidal components of only the solenoidal part:

$$\mathbb{P}[\mathbf{F}] \equiv \mathbb{P}_{\text{sol}}[\mathbf{A}]. \quad (8.58)$$

Note that while the poloidal potential χ for a strictly solenoidal field can be found directly from the axial component A_z , isolating the solenoidal part of a general, non-solenoidal field requires taking both the axial curl and the axial curl of the curl of $\mathbf{F}$ to annihilate the irrotational scalar field ($\nabla \times \nabla \Pi \equiv \mathbf{0}$). Exploiting the operator commutativity, this defines the generalized Poisson equations for the extended projection operator:

$$-\nabla_{\perp}^2 \psi = \hat{\mathbf{z}} \cdot (\nabla \times \mathbf{F}), \quad (8.59a)$$

$$-\nabla_{\perp}^2 (-\nabla^2 \chi) = \hat{\mathbf{z}} \cdot (\nabla \times \nabla \times \mathbf{F}). \quad (8.59b)$$

By solving (8.59), the extended projection operator effectively filters out any scalar gradient ($\mathbb{P}[\nabla \Pi] \equiv \mathbf{0}$) and extracts only the solenoidal component of $\mathbf{F}$.

Applying the extended projection operator $\mathbb{P}$ to the momentum equations (8.5b) annihilates the effective pressure gradient $\nabla \Pi$. This definition of $\mathbb{P}$ serves a vital computational purpose. By reformulating the system into a coupled set of scalar equations, it bypasses the need for an explicit pressure Poisson solver. For traditional implicit-explicit methods, such as a baseline AB2-CN scheme, this allows the system to be integrated entirely in the reduced scalar space, though this approach will inevitably encounter the severe numerical stiffness of the governing equations. Conversely, in the semi-analytical ETD schemes formulated in Chapter 9, diagonalizing the linear operators is significantly more efficient in the primitive velocity space. In that context, $\mathbb{P}$ functions as a direct, algebraic solenoidal filter. Rather than computationally marching a separate pressure field, intermediate velocity fields are mapped to the poloidal-toroidal space and inverted ($\mathbf{F}_{\text{sol}} = \mathbb{P}_{\text{sol}}^{-1} [\mathbb{P}[\mathbf{F}]]$), satisfying mass conservation with spectral accuracy.

Appendix 8A Inertial Gravity Waves in Static, Stratified Fluids

For a static background, the linearized equations (8.22) reduce to the standard governing equations for inertial gravity waves:

$$\hat{\mathbf{r}} : \quad \sigma \tilde{u}_r = 2\Omega \tilde{u}_{\theta} - \text{d}_r \tilde{p} \quad (8.60a)$$

$$\hat{\boldsymbol{\theta}} : \quad \sigma \tilde{u}_{\theta} = -2\Omega \tilde{u}_r - \frac{\text{i}m}{r} \tilde{p} \quad (8.60b)$$

$$\hat{\mathbf{z}} : \quad \sigma \tilde{u}_z = -\tilde{b} - \text{i}k \tilde{p} \quad (8.60c)$$

$$\tilde{b} : \quad \sigma \tilde{b} = \bar{N}^2 \tilde{u}_z \quad (8.60d)$$

By expressing all the variables in terms of $\tilde{p}$, the governing equations simplify to the standard wave equation:

$$\tilde{p}'' + \frac{1}{r}\tilde{p}' + \left(k_h^2 - \frac{m^2}{r^2}\right)\tilde{p} = 0, \quad (8.61)$$

where

$$k_h^2 = -k^2 \frac{(2\Omega)^2 + \sigma^2}{\bar{N}^2 + \sigma^2}. \quad (8.62)$$

Note that two fundamental singularities, $\sigma = \pm i\bar{N}$ and $\sigma = \pm 2\Omega$, exist in the parameter space. They represent the natural resonant frequencies of the fluid, at which the wave motion simplifies to purely buoyancy-driven or purely inertial oscillations.

By enforcing regularity at origin, the globally continuous solutions to (8.61) are the Bessel functions of the first kind:

$$\tilde{p} = C \cdot J_m(k_h \cdot r). \quad (8.63)$$

Considering neutral modes ($\sigma = i\omega$), the dispersion relation for valid propagating modes ($k_h > 0$) is:

$$\omega^2 = \frac{\bar{N}^2 k_h^2 + (2\Omega)^2 k^2}{k_h^2 + k^2}, \quad (8.64)$$

which implies that the frequency of inertial gravity waves is physically bounded between the two natural resonant frequencies of the system:

$$\min(|\bar{N}|, 2|\Omega|) < |\omega| < \max(|\bar{N}|, 2|\Omega|). \quad (8.65)$$

Chapter 9

A Pseudo-Spectral ETD Scheme for Rotating Stratified Flows

The mathematical formulations established in Chapter 8 provide the theoretical basis for analyzing wave-mean flow interactions. However, the direct numerical simulation of these governing equations is complicated by the presence of disparate physical timescales. In rotating, stratified shear flows, the high-frequency restorative forces of inertio-gravity waves and the rapid advective transport of the background flow introduce severe numerical stiffness, which heavily restricts the allowable time step in standard integration schemes. The objective of this chapter is to develop a robust pseudo-spectral numerical framework capable of bypassing these explicit stability limits, enabling the long-term, high-resolution simulation of global cylindrical flows.

To facilitate the development of the numerical schemes, we first cast the continuous Boussinesq disturbance equations into a compact vector-operator notation. Defining the state vector in the primitive velocity basis as

$$\mathbf{v} = [u'_r, u'_\theta, u'_z, b']^\top, \quad (9.1)$$

the governing equations (8.5) can be written abstractly as:

$$\partial_t \mathbf{v} = \mathcal{L}_{\text{wave}} \mathbf{v} + \mathcal{L}_{\text{diff}} \mathbf{v} + \mathcal{L}_{\text{cross}} \mathbf{v} + \mathcal{N}(\mathbf{v}) - \mathbf{\Pi}, \quad (9.2)$$

subject to the solenoidal constraint $\nabla \cdot \mathbf{u}' = 0$. Here, the column vector $\mathbf{\Pi} = [\nabla \Pi, 0]^\top$ absorbs all effective irrotational potentials, including the kinematic pressure perturbation, the centrifugal potential, and the scalar gradient derived from the vorticity form of the nonlinear terms. The constant-coefficient linear matrix operators governing the background rotation and stratification ($\mathcal{L}_{\text{wave}}$) and the viscous and thermal diffusion ($\mathcal{L}_{\text{diff}}$) are explicitly given by:

$$\mathcal{L}_{\text{wave}} = \begin{bmatrix} 0 & 2\Omega & 0 & 0 \\ -2\Omega & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 0 & \bar{N}^2 & 0 \end{bmatrix}, \quad \mathcal{L}_{\text{diff}} = \begin{bmatrix} \nu \nabla^2 & 0 & 0 & 0 \\ 0 & \nu \nabla^2 & 0 & 0 \\ 0 & 0 & \nu \nabla^2 & 0 \\ 0 & 0 & 0 & \kappa \nabla^2 \end{bmatrix}. \quad (9.3a,b)$$

The remaining operators, representing the linear interactions with the steady background shear flow ($\mathcal{L}_{\text{cross}}$) and the nonlinear self-advection ($\mathcal{N}$), act on the state vector to produce the following column vectors:

$$\mathcal{L}_{\text{cross}}\mathbf{v} = \begin{bmatrix} -(\bar{\mathbf{u}} \cdot \nabla)\mathbf{u}' - (\mathbf{u}' \cdot \nabla)\bar{\mathbf{u}} \\ -(\bar{\mathbf{u}} \cdot \nabla)b' \end{bmatrix}, \quad \mathcal{N}(\mathbf{v}) = \begin{bmatrix} \mathbf{u}' \times \boldsymbol{\omega}' \\ -(\mathbf{u}' \cdot \nabla)b' \end{bmatrix}. \quad (9.4\text{a,b})$$

To integrate the system (9.2) in time, we rely on the method of fractional steps. Throughout the numerical recipes presented in this chapter, superscript notations such as $t_{N+1/2}$ are used to denote intermediate computational stages within a single time step, rather than evaluations of the flow field at physical half-time intervals. The integration time step Δt is kept uniform, and active time-step adjustment mechanisms are not considered in this work. Across all numerical schemes evaluated in this study, the numerical treatment of the diffusion and nonlinear terms remains consistent: the viscous and thermal diffusion operator ($\mathcal{L}_{\text{diff}}$) is integrated implicitly using the unconditionally stable Crank-Nicolson (CN) method, while the nonlinear self-advection ($\mathcal{N}$) is integrated explicitly using the second-order Adams-Bashforth (AB2) method. The fundamental difference between the schemes presented herein lies entirely in the specific numerical treatment of the linear wave ($\mathcal{L}_{\text{wave}}$) and cross-term ($\mathcal{L}_{\text{cross}}$) operators.

The chapter is structured to detail the spatial discretization, diagnose the inherent numerical stiffness, and formulate the stabilizing integration strategy. §9.1 outlines the spatial discretization using mapped associated Legendre polynomials and Fourier series. §9.2 discusses the numerical stiffness that arises when standard implicit-explicit (IMEX) methods are applied to rotating stratified shear flows. To overcome these constraints, §9.3 formulates a semi-analytical Exponential Time Differencing (ETD) framework, detailing the analytical diagonalization of the linear operators and demonstrating how physical resonance characteristics, such as critical layers, are natively embedded into the discrete operators. §9.4 provides the numerical implementation details, including numerical stabilization techniques and the parallelization architecture. Finally, §9.5 provides comprehensive code validation, including linear eigenmode verification against theoretical benchmarks and nonlinear energy conservation tracking.

9.1 Spectral Expansion and the Pseudo-Spectral Method

The unbounded cylindrical domain is expanded using Fourier series in the periodic azimuthal and axial directions, and mapped associated Legendre polynomials in the radial direction (Matsushima and P. Marcus, 1997). The semi-infinite radial domain ($0 \leq r < \infty$) is mapped onto a finite domain $\mu \in [-1, 1]$ via the algebraic transformation:

$$\mu(r) = \frac{r^2 - L^2}{r^2 + L^2}, \quad (9.5)$$

where $L > 0$ is a tunable mapping parameter. The radial basis is constructed from the standard associated Legendre polynomials, $P_n^m(\mu)$:

$$P_{Ln}^m(r) \equiv P_n^m(\mu(r)), \quad (9.6)$$

and an arbitrary scalar field ϕ (representing any component of $\mathbf{v}$) can thus be expanded as:

$$\phi(r, \theta, z, t) = \sum_{m,k=-\infty}^{\infty} \sum_{n=|m|}^{\infty} \tilde{\phi}_{mnk}(t) P_{Ln}^m(r) e^{i(m\theta+kz)}, \quad (9.7)$$

where m and k are the azimuthal and axial wavenumbers, and n is the degree of the associated Legendre polynomials. For any real scalar field $\phi \in \mathbb{R}^3$, the reality condition requires complex conjugate relations between (m, k) and $(-m, -k)$ expansions, allowing the solver to compute only half of the wavenumbers, e.g., $m \geq 0$.

The calculation of the spectral coefficients $\tilde{\phi}_{mnk}(t)$ is based on the orthogonality condition:

$$\int_0^\infty P_{Ln}^m(r) P_{Ln'}^m(r) w(r) dr = \delta_{nn'} (N_n^m)^2, \quad (9.8)$$

where the normalization constant N_n^m and the mapped weight function $w(r)$ are defined as:

$$(N_n^m)^2 = 2(n + |m|)! / ((2n + 1)(n - |m|)!), \quad (9.9a)$$

$$w(r) = \frac{4L^2 r}{(r^2 + L^2)^2}. \quad (9.9b)$$

By exploiting this orthogonality relation, the spectral coefficients are extracted by integrating the product of the physical field $\hat{\phi}_{mk}(r, t)$, the basis function $P_{Ln}^m(r)$, and the weight function $w(r)$ over the radial domain. Because this continuous integral over $[0, \infty)$ maps to the bounded interval $\mu \in [-1, 1]$, the numerical integration can be evaluated efficiently using Gauss-Legendre quadratures defined on the μ grid.

The mapped associated Legendre basis is advantageous for unbounded cylindrical domains for several theoretical and computational reasons. The analytical mapping allocates half of the radial collocation points to the inner region $0 \leq r \leq L$, yielding high numerical resolution where vortex dynamics and shear gradients are typically most prominent. This localized resolution tuning was exploited by Sangjoon Lee and Philip S. Marcus (2023) in their linear stability analysis of incompressible columnar vortices. Additionally, the basis functions behave as C^∞ at $r = 0$, satisfying the analyticity conditions at the coordinate singularity. In the far field ($r \rightarrow \infty$, corresponding to $\mu \rightarrow 1$), the mapped polynomials decay as $\mathcal{O}(r^{-|m|})$. Therefore, for any C^∞ function that decays algebraically or exponentially as $r \rightarrow \infty$, the mapped associated Legendre polynomials form a Galerkin basis set that achieves exponential spectral convergence (Sangjoon Lee and Philip S. Marcus, 2023; Matsushima and P. Marcus, 1997).

It is important to mention that, when performing the poloidal-toroidal decomposition, expanding the scalar streamfunctions (ψ and χ) directly in the mapped associated Legendre basis requires the inclusion of a supplemental logarithmic term,

$$P_l(r) \equiv \ln \frac{L^2 + r^2}{2L^2}, \quad (9.10)$$

which accounts for the $\mathcal{O}(r^{-1})$ far-field asymptotic behavior of the azimuthal velocity field (u_θ) and the mean axial components of the velocity and vorticity fields (Matsushima and P. Marcus, 1997). Taking the derivative of this logarithmic term maps it back into the standard polynomial space:

$$r \frac{d}{dr} P_l = P_{L0}^0 + P_{L1}^0. \quad (9.11)$$

As a result, the primitive velocity ($\mathbf{u}'$) and vorticity ($\boldsymbol{\omega}'$) fields are expanded using the standard $P_{Ln}^m(r)$ basis without logarithmic corrections. As detailed in Appendix 9B, the projected logarithmic coefficients ψ_l for the nonlinear term $\mathbf{u}' \times \boldsymbol{\omega}'$ and the linear wave term $-(2\Omega \hat{\mathbf{z}} \times \mathbf{u}' + b' \hat{\mathbf{z}})$ evaluate to zero, ensuring they do not dynamically complicate the time evolution equations.

We shall also note that the density perturbation b' is subjected to a specialized physical constraint. Because the prescribed background stratification $\bar{b}(z)$ is horizontally uniform and dynamically rigid, it cannot redistribute or absorb a net horizontal mass deficit. Any physically realizable localized perturbation must restrict its own total mass anomaly to be finite within the unbounded domain, implying that its axisymmetric ($m = 0$) modes must decay at least as fast as $\mathcal{O}(r^{-4})$ as $r \rightarrow \infty$. In the mapped domain, this suggests that the axisymmetric modes must be divisible by $(1 - \mu)^2$, which mathematically translates to:

$$\sum_{n=0}^{\infty} \tilde{b}'_{0nk} = 0 \quad \text{and} \quad \sum_{n=0}^{\infty} \tilde{b}'_{0nk} \cdot \frac{n(n+1)}{2} = 0 \quad \text{for all } k. \quad (9.12)$$

Since the governing equations and the spectral spatial operators conserve the domain-integrated mass anomaly, these spectral constraints need only be satisfied by the initial conditions, and the numerical scheme will preserve this asymptotic decay without explicit filtering during time integration.

With the complete spectral basis, the governing equations (9.2) are integrated following a pseudo-spectral approach, which requires transitioning among three distinct computational spaces:

- **FFF (Function-Function-Function) space:** The fully spectral domain where spatial variables are expanded by (9.7). Linear operators with constant coefficients (such as diffusion) are evaluated here as algebraic operations.
- **PPP (Physical-Physical-Physical) space:** The fully physical domain evaluated at discrete spatial grid points.

- **PFF (Physical-Function-Function) space:** A hybrid space where the azimuthal and axial dimensions remain in Fourier space, but the radial dimension is evaluated at physical collocation points to accommodate the radial dependencies of $\mathcal{L}_{\text{cross}}$.

To evaluate the nonlinear advective product $\mathcal{N}(\mathbf{v})$, we bypass computationally expensive convolution integrals in the spectral space by transforming the state vector to the PPP space, where direct point-by-point multiplication is performed. The result is subsequently transformed back to the FFF space before the linear operators are applied.

Beyond evaluating nonlinearities, the spectral representation offers computational advantages for the linear physics through the sparsity of the resulting matrices. Applying the transverse Laplacian $\nabla_{\perp}^2$ to the basis function P_{Ln}^m yields the relation

$$\nabla_{\perp}^2|_{m,n} = \frac{4n(n+1)L^2}{(L^2 + r^2)^2}. \quad (9.13)$$

This property, along with the operational recurrence relations documented by Matsushima and P. Marcus (1997), ensures that radial differential operators map natively to narrow-banded matrices, reducing the computational complexity and memory footprint of the numerical solver.

Furthermore, because the azimuthal and axial dimensions are expanded in the orthogonal Fourier basis $e^{i(m\theta + kz)}$, the three-dimensional linear operators ($\mathcal{L}_{\text{wave}}$, $\mathcal{L}_{\text{cross}}$, and $\mathcal{L}_{\text{diff}}$) decouple with respect to m and k . They decompose into independent, banded block matrices for each (m, k) pair, which is suited for massive parallelization. As will be detailed in §9.4, the spectral (m, k) modes are distributed across independent computing processors via a 2D pencil decomposition using the Message Passing Interface (MPI) (Dalcin et al., 2019; Sangjoon Lee and J. Wang, 2025). This strategy allows linear operators to be evaluated simultaneously without inter-process communication. And cross-processor communication is incurred exclusively during the full spatial transformations required for the nonlinear products.

9.2 Stiffness Associated with the Explicit Treatment of the Background Terms

A standard approach to time integration relies on treating the linear operators $\mathcal{L}_{\text{wave}}$ and $\mathcal{L}_{\text{cross}}$ explicitly alongside the nonlinear terms. The motivation for this formulation is its ability to operate primarily within the reduced poloidal-toroidal state space, $\mathbf{V} = [\psi, \chi, b]^T$. By converting the primitive velocity components to scalar potentials via the poloidal-toroidal projection $\mathbb{P}$ introduced in §8.4, the number of evolving variables is reduced from four to three.

Because the curl operators defining the poloidal-toroidal decomposition commute with the three-dimensional Laplacian ($\mathbb{P}\nabla^2\mathbf{u}' = \nabla^2\mathbb{P}[\mathbf{u}']$), the implicit diffusion operator $\mathcal{L}_{\text{diff}}$ can be evaluated directly on the poloidal-toroidal basis. Applying the forward projection ($\mathbb{P}$)

extracts the solenoidal component of any vector field, thereby annihilating the irrotational potential field ($\mathbb{P}[\nabla \Pi] \equiv \mathbf{0}$) without requiring a dedicated pressure Poisson solver. The numerical scheme proceeds as follows:

Numerical recipe 1 (Explicit AB2-CN scheme in poloidal-toroidal space)

1. **Explicit Forcing and Projection:** Given the spectral poloidal-toroidal state $\mathbf{V}^{t_N}$, reconstruct the primitive variables:

$$\mathbf{v}^{t_N} = \mathbb{P}_{\text{sol}}^{-1}[\mathbf{V}^{t_N}]. \quad (9.14)$$

Evaluate the explicit forcing vector in the physical (PPP) space and project to poloidal-toroidal components in spectral (FFF) space:

$$\mathbf{F}^{t_N} = \mathbb{P} [\mathcal{L}_{\text{wave}} \mathbf{v}^{t_N} + \mathcal{L}_{\text{cross}} \mathbf{v}^{t_N} + \mathcal{N}(\mathbf{v}^{t_N})]. \quad (9.15)$$

2. **AB2 Extrapolation:** Advance the projected explicit forcing using the historical Adams-Bashforth weights to form an intermediate, undiffused state:

$$\mathbf{V}^{t_{N+1/2}} = \mathbf{V}^{t_N} + \Delta t \left(\frac{3}{2} \mathbf{F}^{t_N} - \frac{1}{2} \mathbf{F}^{t_{N-1}} \right). \quad (9.16)$$

3. **Implicit Diffusion with CN:** Solve the uncoupled, implicit Helmholtz equations for the diffused poloidal-toroidal state to finalize the time step:

$$\mathbf{V}^{t_{N+1}} = \left(\text{diag}^{-1}(\nu, \nu, \kappa) \frac{2}{\Delta t} - \nabla^2 \right)^{-1} \left(\text{diag}^{-1}(\nu, \nu, \kappa) \frac{2}{\Delta t} \mathbf{V}^{t_{N+1/2}} + \nabla^2 \mathbf{V}^{t_N} \right). \quad (9.17)$$

Because this scheme operates natively in the poloidal-toroidal space, it features a low memory footprint. The commutativity of the diffusion operator ensures that the implicit solve preserves the divergence-free condition, isolating the pressure treatment to the explicit projection step. This scheme is robust for physical scenarios characterized by a hydrostatic background state under low rotation and stratification.

However, the explicit treatment of $\mathcal{L}_{\text{wave}}$ and $\mathcal{L}_{\text{cross}}$ introduces a severe operational limitation for rapidly rotating or strongly stratified flows. The matrix operator $\mathcal{L}_{\text{wave}}$ as given in (9.3a) exhibits a block-diagonal structure with eigenvalues $\lambda_{\text{wave}} = \pm i 2\Omega$ and $\pm i \bar{N}$. These purely imaginary eigenvalues correspond to the physical frequencies of the epicyclic oscillations and internal gravity waves. Treating such restorative, high-frequency physics explicitly using the Adams-Bashforth method strictly bounds the allowable time step to maintain absolute numerical stability. Specifically, the integration time step is bottlenecked by the fastest linear waves¹:

$$\Delta t_{\text{AB2-CN}} \lesssim \mathcal{O}(1/\max(2\Omega, \bar{N})). \quad (9.18)$$

¹We assume a stably stratified background ($\bar{N}^2 > 0$) in this stiffness analysis, which yields the purely imaginary eigenvalues responsible for the severe oscillatory constraint. It is worth noting, however, that the semi-analytical exponential time differencing scheme developed in subsequent sections remains mathematically valid and strictly accurate for unstable stratifications ($\bar{N}^2 < 0$), as it naturally and exactly integrates the resulting real exponential growth factors.

Furthermore, when a background flow possesses strong azimuthal shear, advection dominates the cross-term operator $\mathcal{L}_{\text{cross}}$, yielding eigenvalues that scale with the local angular velocity and imposing an additional advective constraint on the time step. Consequently, for flows where the background rotation, stratification, or shear is significantly faster than the timescale of the nonlinear disturbances, this explicit formulation forces the numerical scheme to take restrictively small time steps, rendering long-term integration impractical.

We must emphasize here that, while it may seem intuitive to bypass these explicit stability limits by absorbing $\mathcal{L}_{\text{wave}}$ and $\mathcal{L}_{\text{cross}}$ into the implicit Crank-Nicolson solve alongside diffusion, doing so introduces severe numerical artifacts. The Crank-Nicolson method approximates the exact exponential time evolution using a $(1, 1)$ Padé approximant. For purely imaginary eigenvalues, this approximation is unconditionally stable, yielding an amplification factor with a modulus of 1. However, at large Courant numbers, the Padé approximant suffers from phase-speed retardation. This numerical dispersion causes high-frequency waves to artificially slow down and steepen, generating spurious ringing oscillations. In rotating and stratified shear flows, these non-physical dispersive artifacts can phenomenologically masquerade as genuine vortex instabilities. Therefore, an implicit CN treatment of the wave and cross terms cannot be trusted, which motivates our formulation of a semi-analytical integration method.

9.3 The Exponential Time Differencing (ETD) Method

General formulation and the basic wave operator

To resolve the numerical stiffness without artificially retarding phase speeds or inducing numerical dispersion, we deploy an Exponential Time Differencing (ETD) scheme, which exponentiates the stiff linear operators and integrates the linear wave and advective physics analytically over the time step. ETD frameworks have proven highly effective in stabilizing pseudo-spectral simulations across a variety of stiff systems, from canonical nonlinear wave equations governing interacting solitons (Whalen et al., 2015) to fully coupled fluid boundary problems, such as dispersive surface tension in quasi-periodic water waves (Wilkening and Zhao, 2021).

Unlike the AB2-CN scheme, the ETD formulation operates primarily in the primitive velocity space ($\mathbf{v}$). Exponentiating the fully coupled linear operators in the primitive form allows for a clean diagonalization, whereas the poloidal-toroidal form is not easily tractable. To establish the mathematical foundation of our approach, we first derive a basic ETD scheme that isolates and exponentiates solely the constant-coefficient wave operator. For a hydrostatic background state, this basic ETD scheme is sufficient to bypass the stiffness associated with the high-frequency inertio-gravity waves. In subsequent sections, we will extend this formulation to include the cross-term operator within the exponential, which is necessary for strongly sheared flows.

To maintain notational consistency for the matrix exponentials and primitive-form vectors, we define:

$$\mathbf{L} \equiv \mathcal{L}_{\text{wave}}, \quad \mathbf{f}(\mathbf{v}) = \mathcal{L}_{\text{cross}} \mathbf{v} + \mathcal{N}(\mathbf{v}), \quad \mathbf{g}(\mathbf{v}) = \mathcal{L}_{\text{diff}} \mathbf{v}(t) - \mathbf{\Pi}(t). \quad (9.19\text{a,b,c})$$

Multiplying the continuous governing equation by the matrix integrating factor $e^{-\mathbf{L}t}$ gives:

$$\partial_t (e^{-\mathbf{L}t} \mathbf{v}) = e^{-\mathbf{L}t} (\mathbf{f}(t) + \mathbf{g}(t)), \quad (9.20)$$

which can be integrated exactly over a single time step from t_N to t_{N+1} :

$$\begin{aligned} \mathbf{v}^{t_{N+1}} &= e^{\mathbf{L}\Delta t} \mathbf{v}^{t_N} + \int_0^{\Delta t} e^{\mathbf{L}(\Delta t - \tau)} (\mathbf{f}(t_N + \tau) + \mathbf{g}(t_N + \tau)) d\tau \\ &= \mathbf{E} \mathbf{v}^{t_N} + \mathbf{NL} \left(\frac{3}{2} \mathbf{f}^{t_N} - \frac{1}{2} \mathbf{f}^{t_{N-1}} + \frac{1}{2} \mathbf{g}^{t_N} + \frac{1}{2} \mathbf{g}^{t_{N+1}} \right) + \mathcal{O}(\Delta t^3), \end{aligned} \quad (9.21)$$

where the discrete integration operators are defined as:

$$\mathbf{E} \equiv e^{\mathbf{L}\Delta t}, \quad \mathbf{NL} \equiv \int_0^{\Delta t} e^{\mathbf{L}(\Delta t - \tau)} d\tau = \mathbf{L}^{-1} (e^{\mathbf{L}\Delta t} - \mathbf{I}). \quad (9.22\text{a,b})$$

Note here that the explicit forcing, the diffusive terms, and the irrotational potential all continuously evolve over the integration interval $\tau \in [0, \Delta t]$, and their integrals are approximated as constants to be factored out of the integration operator. Specifically, the explicit term $\mathbf{f}$ is extrapolated using the historical second-order Adams-Bashforth weights, while the implicit diffusion and potential fields $\mathbf{g}$ are approximated using a Crank-Nicolson temporal average.

The right-hand side of (9.21) contains the unknown future state $\mathbf{g}^{t_{N+1}}$. To isolate it without corrupting the second-order accuracy, we utilize the Taylor series expansions of the two ETD operators:

$$\mathbf{E} = \mathbf{I} + \Delta t \mathbf{L} + \frac{\Delta t^2}{2} \mathbf{L}^2 + \mathcal{O}(\Delta t^3), \quad (9.23\text{a})$$

$$\mathbf{NL} = \Delta t \mathbf{I} + \frac{\Delta t^2}{2} \mathbf{L} + \mathcal{O}(\Delta t^3). \quad (9.23\text{b})$$

By expanding the grouped implicit terms from (9.21), one can show that applying the $\mathbf{NL}$ operator to the CN average is mathematically equivalent, up to $\mathcal{O}(\Delta t^3)$ local truncation error, to applying the exponential operator directly to the historical implicit state:

$$\mathbf{NL} \left(\frac{1}{2} \mathbf{g}^{t_N} + \frac{1}{2} \mathbf{g}^{t_{N+1}} \right) = \mathbf{E} \tilde{\mathbf{g}}^{t_N} + \tilde{\mathbf{g}}^{t_{N+1}} + \mathcal{O}(\Delta t^3), \quad (9.24)$$

where $\tilde{\mathbf{g}} \equiv (\Delta t/2) \mathbf{g}$ and we use the expansion $\tilde{\mathbf{g}}^{t_{N+1}} = \tilde{\mathbf{g}}^{t_N} + \mathcal{O}(\Delta t^2)$. Substituting this approximation back into our numerical formulation and expanding $\tilde{\mathbf{g}}$ separates the known

and unknown states:

$$\left(\mathbf{I} - \frac{\Delta t}{2} \mathcal{L}_{\text{diff}}\right) \mathbf{v}^{t_{N+1}} + \tilde{\mathbf{I}}^{t_{N+1}} = \mathbf{E} \left(\mathbf{v}^{t_N} + \frac{\Delta t}{2} \mathcal{L}_{\text{diff}} \mathbf{v}^{t_N} - \tilde{\mathbf{I}}^{t_N} \right) + \mathbf{NL} \left(\frac{3}{2} \mathbf{f}^{t_N} - \frac{1}{2} \mathbf{f}^{t_{N-1}} \right), \quad (9.25)$$

where $\tilde{\mathbf{I}} \equiv (\Delta t/2) \mathbf{II}$ is the scaled irrotational potential vector. The right-hand side of equation (9.25) depends strictly on known quantities, forming an intermediate state $\mathbf{v}^{t_{N+1/2}}$. Applying the continuous poloidal-toroidal projection $\mathbb{P}$ to both sides annihilates the updated potential $\tilde{\mathbf{I}}^{t_{N+1}}$. Because $\mathbb{P}$ and $\mathcal{L}_{\text{diff}}$ commute, this algebraic reduction directly returns the uncoupled, implicit Helmholtz equations for the diffused poloidal-toroidal state.

To compute the exact matrix exponentials, the constant-coefficient operator $\mathbf{L} = \mathcal{L}_{\text{wave}}$ is diagonalized as

$$\mathbf{L} = \mathbf{SJS}^{-1}, \quad (9.26)$$

where $\mathbf{J}$ is the diagonal matrix of eigenvalues, representing the inertial and internal gravity wave frequencies, and $\mathbf{S}$ and $\mathbf{S}^{-1}$ are the corresponding eigenvector matrix and its inverse:

$$\mathbf{J} = \text{diag} (i2\Omega, -i2\Omega, i\bar{N}, -i\bar{N}), \quad (9.27a)$$

$$\mathbf{S} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ i & -i & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & -i\bar{N} & i\bar{N} \end{bmatrix}, \quad \mathbf{S}^{-1} = \frac{1}{2} \begin{bmatrix} 1 & -i & 0 & 0 \\ 1 & i & 0 & 0 \\ 0 & 0 & 1 & i/\bar{N} \\ 0 & 0 & 1 & -i/\bar{N} \end{bmatrix}. \quad (9.27b)$$

The discrete integration operators can be evaluated as:

$$\mathbf{E} = \mathbf{S} e^{\mathbf{J}\Delta t} \mathbf{S}^{-1}, \quad \mathbf{NL} = \mathbf{SJ}^{-1} (e^{\mathbf{J}\Delta t} - \mathbf{I}) \mathbf{S}^{-1}. \quad (9.28a,b)$$

An important technical detail associated with the numerical evaluation of the $\mathbf{NL}$ operator is that, whenever $\mathbf{J}$ contains zero eigenvalues (e.g., when the background rotation is absent, $\Omega = 0$, causing $\mathbf{J}^{-1}$ to become undefined), a singularity arises and must be addressed. This indeterminate form is strictly mathematical, not physical, and is resolved by applying L'Hôpital's rule for the corresponding diagonal components, denoted by $\lambda_{\mathbf{L}}$:

$$\lim_{\lambda_{\mathbf{L}} \rightarrow 0} \frac{e^{\lambda_{\mathbf{L}} \Delta t} - 1}{\lambda_{\mathbf{L}}} = \Delta t. \quad (9.29)$$

Consequently, any zero-eigenvalue of the diagonalized linear operator asymptotes to Δt , effectively reducing the ETD integration to a standard AB2 explicit update for that specific component (e.g., for the non-rotating case, the horizontal velocity components are reduced to AB2 for time integration). Furthermore, for non-zero but infinitesimally small arguments ($0 < |\lambda_{\mathbf{L}} \Delta t| \ll 1$), direct floating-point evaluation of the exponential term suffers from catastrophic cancellation. To construct a mathematically airtight architecture across all dynamical scales, this numerical limit can be formally bypassed by substituting a Taylor series expansion for the $\mathbf{NL}$ operator, a robust technique established in generalized ETD

frameworks to prevent precision loss near the zero-mode limits (Whalen et al., 2015; Wilkening and Zhao, 2021). While this treatment is necessary for fully generalized applications, the target phenomenological regime of the present study (e.g., the zombie vortex instability) relies on dynamically balanced background states lacking extreme temporal scale separation. This allows the solver to operate efficiently at large time steps safely above the precision floor, rendering the runtime series expansion unnecessary for the simulations presented herein.

Likewise, a critical physical singularity arises for the z -invariant modes. Consider the $k = 0$ wavenumber of the disturbance field: $\mathbf{u}'_{m0} \equiv \mathbf{u}'_{m0}(r, \theta)$. The solenoidal condition (8.5a) is linear and suggests that it can be constructed from a set of poloidal-toroidal potential functions: $\mathbf{u}'_{m0} = (\partial_\theta \psi'_{m0}/r)\hat{\mathbf{r}} - \partial_r \psi'_{m0}\hat{\boldsymbol{\theta}} - \nabla_\perp^2 \chi'_{m0}\hat{\mathbf{z}}$. The result of $\mathcal{L}_{\text{wave}}$ applied to $\mathbf{u}'_{m0}$ is then evaluated to be $2\Omega\partial_r \psi'_{m0}\hat{\mathbf{r}} + (2\Omega\partial_\theta \psi'_{m0}/r)\hat{\boldsymbol{\theta}} - b'_{m0}\hat{\mathbf{z}}$, whose divergence is given by $2\Omega\nabla^2 \psi_{m0}$. Hence, it is instantaneously balanced by the two-dimensional pressure field ($\Pi_{m0} = 2\Omega\psi_{m0}$), leaving only the axial component, so the Coriolis force exerts exactly zero net dynamic effect on the physical evolution of the $k = 0$ modes. If the standard $\mathbf{E}$ and $\mathbf{NL}$ operators are blindly applied to these modes, the scheme will unphysically rotate them, corrupting the kinetic energy of the flow. To enforce the correct physics, the discrete horizontal ETD operators for all $k = 0$ modes must be manually overridden to match their non-rotating limits: $\mathbf{E}_{k=0} \rightarrow \mathbf{E}_{\Omega=0}$ and $\mathbf{NL}_{k=0} \rightarrow \mathbf{NL}_{\Omega=0}$.

With the operators explicitly defined and the constraints satisfied, the basic ETD time-stepping algorithm proceeds as follows²:

Numerical recipe 2 (ETD scheme in primitive space)

1. **Explicit Forcing and Filtering:** Given the primitive state $\mathbf{v}^{t_N}$, compute the explicit forcing $\mathbf{f}^{t_N}$ in the physical (PPP) space, and project the explicit terms to poloidal-toroidal space and back to enforce strict incompressibility prior to time integration for numerical stability:

$$\mathbf{f}_{\text{sol}}^{t_N} = \mathbb{P}_{\text{sol}}^{-1} [\mathbb{P}[\mathcal{L}_{\text{cross}}\mathbf{v}^{t_N} + \mathcal{N}(\mathbf{v}^{t_N})]] . \quad (9.30)$$

2. **ETD Advancement:** Apply the exact matrix exponentials to the historical primitive state and the AB2-extrapolated explicit forcing to obtain an intermediate, undiffused state:

$$\mathbf{v}^{t_{N+1/2}} = \mathbf{E} \left(\mathbf{v}^{t_N} + \frac{\Delta t}{2} \mathcal{L}_{\text{diff}} \mathbf{v}^{t_N} - \tilde{\Pi}^{t_N} \right) + \mathbf{NL} \left(\frac{3}{2} \mathbf{f}_{\text{sol}}^{t_N} - \frac{1}{2} \mathbf{f}_{\text{sol}}^{t_{N-1}} \right) . \quad (9.31)$$

3. **Implicit Diffusion:** Apply the forward poloidal-toroidal projection $\mathbb{P}$ to algebraically drop the potential gradient $\tilde{\Pi}^{t_{N+1}}$. Solve the resulting uncoupled, implicit Helmholtz

²Richardson extrapolation can be applied to initiate the time-stepping, which consists of two half steps and one full step, each following a first-order exponential time differencing scheme that uses forward Euler method to extract the nonlinear and cross terms, and backward Euler method for the pressure term and the diffusion term.

equations in the spectral (FFF) space for the diffused poloidal-toroidal state to finalize the explicit physics:

$$\mathbf{V}^{t_{N+1}} = \left(\mathbf{I} - \frac{\Delta t}{2} \mathcal{L}_{\text{diff}} \right)^{-1} \mathbb{P}[\mathbf{v}^{t_{N+1/2}}]. \quad (9.32)$$

The final primitive velocity is subsequently reconstructed via

$$\mathbf{v}^{t_{N+1}} = \mathbb{P}_{\text{sol}}^{-1}[\mathbf{V}^{t_{N+1}}]. \quad (9.33)$$

4. **Pressure Extraction:** Calculate the updated irrotational potential using the residual of the intermediate state, to be stored for the subsequent time step:

$$\tilde{\Pi}^{t_{N+1}} = \mathbf{v}^{t_{N+1/2}} - \left(\mathbf{I} - \frac{\Delta t}{2} \mathcal{L}_{\text{diff}} \right) \mathbf{v}^{t_{N+1}}. \quad (9.34)$$

It is worth mentioning that the poloidal-toroidal projection in the explicit forcing evaluation (9.30) is an optional treatment to improve the numerical stability. In essence, it decomposes the pressure term into two distinct components: the first instantaneously addresses velocity changes caused by nonlinear interactions, while the second accounts for changes arising from the explicit time integration of the linear term. While this additional step neither increases nor reduces the order of the local truncation error of the numerical scheme, it ensures that any field that is fed into the time integration operators of the ETD schemes is divergence-free, which improves the overall stability of the code.

Diagonalization of the background shear

While the basic ETD scheme formulated above stabilizes the restorative wave forces, evaluating the background shear cross terms explicitly via the Adams-Bashforth method remains impractically stiff for strongly sheared flows. To achieve robust simulation speeds, the analytical integration must be expanded to encompass these shear interactions alongside the wave operator.

To incorporate the cross terms into the exact matrix exponential, we must first select their mathematical formulation, as it directly dictates the computational tractability of the resulting scheme. If $\mathcal{L}_{\text{cross}}$ is expressed in the vorticity form,

$$\mathcal{L}_{\text{cross}}(\mathbf{v}) = \bar{\mathbf{u}} \times \boldsymbol{\omega}' + \mathbf{u}' \times \bar{\zeta}(r) \hat{\mathbf{z}}, \quad (9.35)$$

it isolates two distinct physical mechanisms: the macroscopic sweeping and stretching of the perturbation vortices ($\bar{\mathbf{u}} \times \boldsymbol{\omega}'$) and the local Coriolis-like restoring force driven by the background vorticity ($\mathbf{u}' \times \bar{\zeta} \hat{\mathbf{z}}$). However, the curl operator required to obtain $\boldsymbol{\omega}'$ inextricably couples the different velocity components within the vorticity stretching term. Exponentiating this differential operator would require the inversion of a massive, globally coupled matrix system, rendering exact integration computationally prohibitive.

Alternatively, one might consider a hybrid scheme where the algebraic term $-\bar{\zeta}\hat{\mathbf{z}} \times \mathbf{u}'$ is analytically diagonalized alongside the wave operator, and the differential term $\bar{\mathbf{u}} \times \boldsymbol{\omega}'$ is absorbed into the explicitly evaluated forcing vector. Nonetheless, our numerical experiments showed this partial ETD implementation inevitably suffers from numerical blow-up at moderate time steps, indicating that the scheme is still critically stiff due to the explicit treatment of the remaining cross term.

Casting the cross terms in their primitive advective form resolves this structural limitation:

$$\mathcal{L}_{\text{cross}} = -(\bar{\mathbf{u}} \cdot \nabla) \mathbf{u}' - (\mathbf{u}' \cdot \nabla) \bar{\mathbf{u}}. \quad (9.36)$$

Because the background flow is strictly azimuthal ($\bar{\mathbf{u}} = r\bar{\Omega}(r)\hat{\boldsymbol{\theta}}$), evaluating these directional derivatives generates geometric off-diagonal terms alongside the advective diagonal $-im\bar{\Omega}(r)$, all of which do not contain radial derivatives. The fully coupled linear operator $\mathbf{L} = \mathcal{L}_{\text{wave}} + \mathcal{L}_{\text{cross}}$ is thus algebraic in r . This guarantees that in the hybrid PFF space, $\mathbf{L}$ forms an independent, 4×4 block-diagonal matrix at each radial collocation point and each azimuthal wavenumber:

$$\begin{aligned} \mathbf{L}(r, m) &= \begin{bmatrix} \mathbf{L}_H & \mathbf{0} \\ \mathbf{0} & \mathbf{L}_V \end{bmatrix} \\ &= \begin{bmatrix} -im\bar{\Omega}(r) & 2\Omega + 2\bar{\Omega}(r) & 0 & 0 \\ -2\Omega - \bar{\zeta}(r) & -im\bar{\Omega}(r) & 0 & 0 \\ 0 & 0 & -im\bar{\Omega}(r) & -1 \\ 0 & 0 & \bar{N}^2 & -im\bar{\Omega}(r) \end{bmatrix}. \end{aligned} \quad (9.37)$$

We can therefore diagonalize and exponentiate this system at each (r, m) -pair independently without coupling the global domain.

Because the horizontal ($\mathbf{L}_H$) and vertical ($\mathbf{L}_V$) blocks cleanly decouple, they are diagonalized separately. For the horizontal block, the eigenvalues mathematically capture the Doppler-shifted epicyclic frequencies, yielding the following diagonal and eigenvector matrices:

$$\mathbf{J}_H = \begin{bmatrix} -i(m\bar{\Omega} + \sqrt{\Delta}) & 0 \\ 0 & -i(m\bar{\Omega} - \sqrt{\Delta}) \end{bmatrix}, \quad (9.38a)$$

$$\mathbf{S}_H = \begin{bmatrix} \frac{i\sqrt{\Delta}}{2\Omega + \bar{\zeta}} & -\frac{i\sqrt{\Delta}}{2\Omega + \bar{\zeta}} \\ 1 & 1 \end{bmatrix}, \quad \mathbf{S}_H^{-1} = \begin{bmatrix} -i\frac{2\Omega + \bar{\zeta}}{2\sqrt{\Delta}} & \frac{1}{2} \\ i\frac{2\Omega + \bar{\zeta}}{2\sqrt{\Delta}} & \frac{1}{2} \end{bmatrix}, \quad (9.38b)$$

where the generalized Rayleigh criterion $\Delta(r)$ follows the previous definition given in (4.11). For the vertical block, the eigenvalues correspond to the Doppler-shifted buoyancy frequencies:

$$\mathbf{J}_V = \begin{bmatrix} -i(m\bar{\Omega} + \bar{N}) & 0 \\ 0 & -i(m\bar{\Omega} - \bar{N}) \end{bmatrix}, \quad (9.39a)$$

$$\mathbf{S}_V = \begin{bmatrix} -\frac{i}{\bar{N}} & \frac{i}{\bar{N}} \\ 1 & 1 \end{bmatrix}, \quad \mathbf{S}_V^{-1} = \begin{bmatrix} \frac{i\bar{N}}{2} & \frac{1}{2} \\ -\frac{i\bar{N}}{2} & \frac{1}{2} \end{bmatrix}. \quad (9.39b)$$

Notably, the vertical eigenvectors remain identical to those of the basic unsheared ETD scheme, as the advective cross term physically applies a Doppler shift to the phase accumulation rate without altering the fundamental vertical eigenstructure.

This complete analytical diagonalization reveals a profound physical elegance behind the ETD framework: the fundamental stability limits and resonance characteristics of the fluid are intrinsically encoded within the mathematical structure of the integration operators. The matrix $\mathbf{J}_H$ yields purely imaginary (oscillatory) eigenvalues if and only if $\Delta(r) > 0$ universally. This guarantees that the numerical scheme rigorously respects the physical requirement for centrifugal stability of the background flow, ensuring the analytical eigenvectors remain well-defined and avoiding the mathematical singularity at $\Delta(r) = 0$. Similarly, the vertical formulation becomes singular when $\bar{N} = 0$; however, this unstratified limit falls outside the scope of the stratified regimes considered in the present study.

Furthermore, the formulation naturally accommodates critical layer phenomena. As evidenced by (9.38a) and (9.39a), zero eigenvalues will occur at critical radii, r_c , where the local Doppler-shifted frequencies vanish: $\bar{\Omega}(r_c) = \pm \bar{N}/m$ (baroclinic critical layers), or at the turning points, r_t , where $\bar{\Omega}(r_t) = \pm \sqrt{\Delta(r_t)}/m$ (surface turning points). In the classic wave dynamics of rotating shear flows, the two surface turning points act as reflective boundaries that spatially sandwich the corotation resonance located at $\bar{\Omega}(r_c) = 0$ for a stationary mode (barotropic critical layer), forming a resonant wave cavity. Physically, the static (zero-frequency) critical layers correspond to regions of local wave resonance, absorption, and momentum transfer (Booker and Bretherton, 1967; S. Maslowe, 1986), and they have been shown to incite instability dynamics such as the zombie vortex instability (Barranco, Pei, et al., 2018; Philip S. Marcus, Pei, Jiang, Barranco, et al., 2015; Philip S. Marcus, Pei, Jiang, and Hassanzadeh, 2013; C. Wang and Balmforth, 2021; Wohlever et al., 2025). Because the zero eigenvalues at these dynamically vital regions are handled by the analytical L'Hôpital limit (9.29) in our ETD framework, the numerical scheme inherently captures the critical layer dynamics without requiring *ad hoc* numerical intervention.

Operationally, this advanced ETD scheme follows the same algorithm detailed in §9.3, with three implementation distinctions. First, because the linear cross terms are absorbed analytically into $\mathbf{L}(r, m)$, the explicit forcing vector is reduced to the nonlinear self-advection: $\mathbf{f}^{t_N} = \mathcal{N}(\mathbf{v}^{t_N})$. Second, because the matrices $\mathbf{E}$ and $\mathbf{NL}$ are radially dependent, the ETD advancement must be computed in the hybrid PFF space, which requires applying updates at each radial point before transforming back to the spectral FFF space. Third, by fully diagonalizing the advective cross terms, the instantaneous pressure balance of the z -invariant ($k = 0$) modes is assimilated into the eigenvalue structure, eliminating the need for the manual non-rotating overrides required by the basic scheme.

This advanced ETD scheme removes the background flow advective stability constraint. Because both the wave forces and shear interactions are integrated analytically, the allowable explicit time step is dictated by the temporal truncation error of the nonlinear perturbation advection rather than the fast background kinematics.

9.4 Numerical Implementation

Numerical Stabilization, Aliasing, and Outgoing Waves

For pseudo-spectral methods, evaluating nonlinear advective products in the physical collocation space intrinsically generates high-frequency harmonics. In high-Reynolds-number or inviscid simulations, the physical diffusion operator ($\mathcal{L}_{\text{diff}}$) is insufficient to dissipate this forward energy cascade. To prevent artificial energy pile-up at the truncation limits and the subsequent back-scattering of aliasing errors, the numerical framework utilizes two stabilization paradigms.

For the periodic azimuthal and axial directions, discretized using uniform Fourier series, one can apply the standard 2/3 de-aliasing rule to geometrically truncate the corrupted high-frequency harmonics. Alternatively, and mandated for the semi-infinite radial direction where the non-uniform Gauss-Legendre grid renders geometric truncation less effective, the framework relies on a dynamically adjusted hyperviscosity and hyperdiffusivity filter to exponentially suppress the forward cascade (Barranco and Philip S. Marcus, 2006; J. P. Boyd, 1998; Canuto et al., 2007).

When employing the hyperviscosity operator, rather than applying computationally rigid, high-order differential operators (e.g., ∇^p), the hyperviscosity operator is implemented as an exact integrating factor applied to the spectral modal amplitudes at the conclusion of each time step. For a given scalar variable ϕ , the generalized spectral filter takes the form:

$$\phi_{\text{filtered}}^{t_{N+1}} = \phi^{t_{N+1}} \exp \left(-\Delta t \sum_{j \in \{r, \theta, z\}} \nu_{\text{hyp}, j} \mathcal{M}_j(w_{\text{op}, j}) \left(\frac{w_{\text{op}, j}}{w_{\text{op}, j}^{\text{max}}} \right)^{p/2} \right), \quad (9.40)$$

where $\nu_{\text{hyp}, j}$ are the directional dissipation coefficients, and $w_{\text{op}, j}$ represents the specific directional eigenvalue of the underlying spectral basis operators (i.e., $w_{\text{op}, \theta} = m^2$, $w_{\text{op}, z} = k^2$, and $w_{\text{op}, r} = n(n+1)$, where n is the degree of the associated Legendre polynomial following the spectral relation (9.13)). The masking function, $\mathcal{M}_j(w_{\text{op}, j})$, is a smooth hyperbolic tangent step profile designed to evaluate to zero for the lower wavenumbers, ensuring large-scale inertial physics remain undamped.

To ensure the dissipation rate perfectly matches the dynamically varying nonlinear cascade without over-damping the flow, the coefficients $\nu_{\text{hyp}, j}$ can be adjusted dynamically every hundreds of time steps via a Proportional-Derivative (PD) feedback control loop. Using Parseval's theorem (detailed in Appendix 9A), the 1D kinetic and available potential energy spectra are evaluated along the periodic azimuthal and axial directions. The spectral shape of the high-wavenumber tail is then extracted by projecting the logarithmic energy spectra, $\log_{10} E(x)$, onto the first and second Legendre polynomials (P_1 and P_2), where $x \in [-1, 1]$ maps across the tail wavenumbers. These projections quantify the local spectral slope and curvature, respectively. A positive curvature indicates an upward concavity (an energy bottleneck), while the slope indicates the bulk dissipation rate. The dissipation coefficients can then be actively updated based on the deviation from a theoretical inertial-range target

state:

$$\nu_{\text{hyp},j}^{t+\Delta t} = \nu_{\text{hyp},j}^t \left[1 + K_p (C^t - C^{t-\Delta t}) + K_d (C^t - C_{\text{target}}) \right], \quad (9.41)$$

where C is the measured spectral metric (curvature or slope), and K_p and K_d are the proportional and derivative gains. For the azimuthal and axial directions, the corresponding metrics and damping coefficients can be directly evaluated from their respective 1D spectra. Because the non-periodic mapped radial basis precludes a direct 1D Parseval spectral density, the radial metric C_r is formulated as a weighted average of the azimuthal and axial metrics, biased strictly towards the direction exhibiting the more severe spectral pile-up. This control architecture dynamically ramps up the hyperviscosity and hyperdiffusivity exactly when spectral blocking is detected, and smoothly relaxes them once the turbulent cascade is properly resolved.

Finally, we must manage outgoing waves radiating toward the domain boundaries. For the unbounded radial direction, the mapped Legendre basis natively provides an algebraic decay as $r \rightarrow \infty$. A radial sponge layer is avoided since applying a spatially varying damping profile in the radial direction of a strongly sheared background flow injects localized spurious vorticity, which can trigger non-physical instabilities. On the other hand, to safely absorb waves reaching the periodic boundaries in the axial direction, a physical 1D axial sponge layer is adapted from (Mahdinia et al., 2017). A damping term $-D(z)\phi$ is appended to the evolution equations, utilizing the spatial profile:

$$D(z) = \frac{1 - T(z)}{\tau}, \quad (9.42)$$

where $T(z)$ is a hyperbolic tangent top-hat function maintaining $T = 1$ throughout the interior domain, and τ dictates the damping timescale near the boundaries. To preserve the global second-order temporal accuracy of the ETD scheme without violating incompressibility, this damping is integrated in the hybrid PFP space at the end of each fractional step:

$$\phi_{\text{sponged}}^{t_{N+1}} = \phi^{t_{N+1}} \exp(-D(z)\Delta t). \quad (9.43)$$

Most importantly, it is applied exclusively to the scalar potentials (ψ, χ) and the buoyancy perturbation (b') , which guarantees that the reconstructed 3D velocity field remains strictly divergence-free.

Parallelization and Performance

Parallelization architecture

Our numerical solver employs a hybrid MPI/OpenMP parallelization paradigm, specifically tailored to the memory and communication requirements of three-dimensional pseudo-spectral methods.

The global computational domain is distributed across compute nodes using a two-dimensional pencil decomposition, while shared-memory concurrency via OpenMP is utilized

for all on-node, thread-level operations. The MPI process topology is configured as a two-dimensional Cartesian grid of $P = N_1 \times N_2$ ranks, maintaining a near-square aspect ratio. Unlike finite-difference methods, where a square topology is chosen to minimize surface-to-volume ghost cell exchanges, the global data volume in spectral transpositions is independent of the process grid aspect ratio. Instead, a near-square grid ($N_1 \approx N_2$) minimizes the sum of the sub-communicator sizes ($N_1 + N_2$), thereby minimizing the total number of point-to-point messages during the `MPI_ALLTOALLW` collectives and mitigating network latency overhead. From this topology, two orthogonal sub-communicators are instantiated: `SUBCOMM_1` (of size N_1) and `SUBCOMM_2` (of size N_2). To ensure optimal load balancing across this grid, the domain decomposition employs a near-uniform partition strategy: an arbitrary dimension of size N distributed across p ranks assigns $\lfloor N/p \rfloor + 1$ elements to the first $N \bmod p$ ranks and $\lfloor N/p \rfloor$ elements to the remainder. This rigorously bounds the maximum memory and workload imbalance to a single element per dimension.

In the pencil decomposition, the 3D data array remains contiguous and undivided along one spatial or spectral dimension (the “pencil”), while the remaining two dimensions are distributed across the $N_1 \times N_2$ process grid. Let N_r , N_θ , and N_z denote the number of grid points in the physical radial, azimuthal, and axial directions, respectively, and let N_n , N_m , and N_k denote their corresponding spectral modes. In the physical space (PPP) and the azimuthal-Fourier space (PFP), data is structured as a θ -pencil: the azimuthal direction (θ or m) is fully local to every rank, while the radial and axial dimensions are partitioned by `SUBCOMM_1` and `SUBCOMM_2`, respectively. Because the contiguous m -axis is locally owned, the real-to-complex azimuthal discrete Fourier transforms (DFTs) can execute without any inter-process communication. Conversely, the fully spectral operations require an r -pencil configuration (the PFF and FFF spaces), where the radial column is fully local, and the spectral axes m and k are distributed across `SUBCOMM_2` and `SUBCOMM_1`. Since the majority of the numerical operators are applied in PFF and FFF spaces, N_1 and N_2 are explicitly assigned such that the larger sub-communicator aligns with the larger of the two distributed spectral dimensions (N_m or N_k). This targeted alignment guarantees that the workload remains optimally distributed across the process grid during the dominant phases of the time integration. Future versions of the solver might consider dynamically tuning the ratio of N_1 and N_2 to precisely match the ratio of N_k and N_m ; while this would yield perfectly isotropic local spectral blocks that might further optimize OpenMP thread scheduling, overall improvements could be limited and highly case-dependent, as departing from a near-square process grid risks exacerbating network latency during the collective transpositions. The evolution of the local array shapes defining these spaces is summarized in Table 9.1.

Because our formulation relies on a pseudo-spectral method, the nonlinear convective products are evaluated pointwise on the physical collocation grid, whereas the spatial derivatives, linear operators, and the semi-analytical ETD advancement are resolved in their corresponding spectral domains. Consequently, the core parallelization challenge lies in efficiently transforming data between the physical and spectral representations.

To transition between the PFP and PFF spaces, the pencil orientation must be systematically rotated. This global two-stage transposition is strictly encapsulated within the axial

Space	Local shape	Pencil
PPP	$[N_r/N_1, N_\theta, N_z/N_2]$	θ
PFP	$[N_r/N_1, N_m, N_z/N_2]$	m
INT	$[N_r/N_1, N_m/N_2, N_z]$	z
PFF	$[N_r, N_m/N_2, N_k/N_1]$	r
FFF	$[N_n, N_m/N_2, N_k/N_1]$	n

Table 9.1: Local array shape on each MPI rank for a process grid $N_1 \times N_2$. Distributed dimensions are indicated by the divisor; the undivided dimension is the pencil (fully local). The intermediate (INT) space is a temporary staging area during the axial transform, where the z -pencil is formed before the subsequent radial or azimuthal reorientation.

FFT subroutines, ensuring that the high-level ETD time-stepping logic remains completely free of explicit message-passing operations. Following the forward direction from the physical to the spectral domain, a process that is exactly reversible during the inverse transform, the internal sequence proceeds as follows (illustrated in figure 9.1):

1. **First Transposition Phase (PFP $\rightarrow$ INT):** The z -dimension must be made contiguous to evaluate the axial DFT. As shown in figure 9.1a, the local m -axis is sub-partitioned into N_2 independent blocks. A collective swap is executed strictly within `SUBCOMM_2`. During this phase, each rank sends its m -subdivisions to the other ranks in the sub-communicator and receives the requisite z -segments in return. Upon completion, the z -dimension is fully reassembled locally, yielding the intermediate (INT) space where the independent axial DFTs are performed, transforming the z -axis into the spectral k -axis.
2. **Second Transposition Phase (INT $\rightarrow$ PFF):** To prepare for the radial transforms and the ETD advancement, a contiguous r -dimension is required. As depicted in figure 9.1b, the local k -dimension is sub-partitioned into N_1 blocks. A second collective swap is executed across `SUBCOMM_1`, where ranks exchange k -subdivisions to gather their complementary r -segments. This reassembles the complete radial dimension locally, mapping the data natively into the PFF space.

Executing these multi-dimensional block exchanges typically incurs severe computational overhead from explicitly packing and unpacking non-contiguous send and receive buffers. To completely circumvent this, the transpositions utilize zero-copy memory routing via the `MPI_TYPE_CREATE_SUBARRAY` interface. During initialization, specific derived datatypes are constructed to encode the multi-dimensional extents and memory strides of the sub-partitions. This instructs the MPI library on how to traverse the buffers directly, enabling the `MPI_ALLTOALLW` collectives to scatter from and gather into the primary simulation arrays natively, which drastically reduces memory bandwidth pressure.




Figure 9.1: Detailed visualization of the two-stage global transposition encapsulated within the axial transform, converting between the azimuthal (PFP) and fully spectral (PFF) spaces on a $P = 3 \times 2$ MPI grid. **(a)** The first phase operates within `SUBCOMM_2`, exchanging sub-partitions of the complete m -axis to reconstruct complete z -columns in the intermediate (INT) space. **(b)** The second phase operates within `SUBCOMM_1`, exchanging sub-partitions of the complete k -axis to reconstruct complete r -columns in the PFF space. The 2D matrix panels illustrate the exact non-contiguous memory mappings encoded via `MPI_TYPE_CREATE_SUBARRAY`, enabling the `MPI_ALLTOALLW` collective to perform strided transpositions dynamically without explicit buffer packing.

Operator	Form	Structure
DELSQH	$(1 - \mu)^{-2} \nabla_{\perp}^2$	Diagonal
DEL2	∇^2	Penta-diagonal
IHELM	$(\nabla^2 - cI)^{-1}$	Penta-diagonal LU
IHELMP	$(\nabla^{2p} + cI)^{-1}$	$(2p + 1)$ -diagonal LU
MULXM	$(1 - \mu) \cdot$	Tri-diagonal mat-vec
XXDX	$(1 - \mu^2) \partial_{\mu} = r \partial_r$	Tri-diagonal mat-vec

Table 9.2: Key spectral operators, their mathematical forms, and banded-matrix widths. μ is the mapped finite radial coordinate, and c is a constant coefficient.

Once the data is mapped into the PFF space, the r -pencils are fully contiguous, allowing the radial spectral expansion (the mapped associated Legendre expansion) to proceed purely locally. This transformation projects the data into the FFF space and exploits the parity of the associated Legendre polynomials via a symmetric/antisymmetric decomposition, effectively halving the floating-point operation count while utilizing highly optimized BLAS matrix-vector products for the Gauss quadrature. Furthermore, a triangular truncation is applied in the FFF space, activating only the lower $N_r - m$ radial modes for each azimuthal wavenumber m .

A defining mathematical advantage of this pseudo-spectral discretization is that the spatial differential operators decouple completely across the spectral azimuthal and axial modes. This produces tightly banded matrix operators for each specific (m, k) pair, whose structures are summarized in table 9.2 for the primary spectral operators. Because these modes are independent, the semi-analytical ETD advancement is applied directly in the PFF and FFF spaces without any MPI overhead. OpenMP threads are dispatched to concurrently process the local pool of decoupled radial columns within each MPI rank. The threads navigate the banded-matrix operations, such as the tri-diagonal evaluations and the penta-diagonal Helmholtz LU solver (IHELM), in an embarrassingly parallel fashion, requiring no inter-thread synchronization beyond the implicit barriers at the end of the time-stepping phases.

Parallel performance and scaling

The hybrid architecture isolates all inter-process communication to the transpositions bridging the nonlinear physical evaluations and the linear spectral advancements. Specifically, each evaluation of the nonlinear term requires one forward and one backward transformation between the PPP and PFF spaces for each state variable, thereby invoking exactly four MPI_ALLTOALLW collectives per variable during each time step. The per-rank data volume exchanged during each collective scales as $V_{\text{comm}} \sim \mathcal{O}(N_r N_m N_k / P)$, satisfying ideal weak and strong scaling expectations for all-to-all patterns.




Figure 9.2: Single-node strong scaling performance for a $256 \times 256 \times 128$ physical grid on dual-socket AMD EPYC 7763 processors (128 total cores). Legend labels denote the total step time, the pointwise nonlinear evaluation ($\mathcal{N}$), the ETD advancement ($\mathbf{E}, \mathbf{NL}$), the full space conversions (PPP $\leftrightarrow$ FFF), the localized radial expansions, and the solenoidal projection ($\mathcal{P}$). The dashed line represents ideal linear scaling.

We assess the parallel performance through a series of strong scaling tests. To isolate algorithmic scaling properties from the hardware-induced latency variances typical of multi-node network interconnects, the performance profiling here is strictly confined to a single high-performance computing node on the Purdue Anvil supercomputer (Song et al., 2022). The compute node is equipped with dual-socket AMD EPYC 7763 processors clocked at 2.45GHz, providing a total of 128 physical cores with a layered cache architecture including 32KB of L1 data (L1d) and instruction (L1i) caches, 512KB of L2, and 32MB of L3 cache.

The test problem uses a physical collocation grid of $256 \times 256 \times 128$. Because the azimuthal transform utilizes a real-to-complex FFT, the 256 physical azimuthal points are stored as 128 complex Fourier modes. Furthermore, to reserve spectral capacity for high-order radial differentiation, the effective spectral resolution is internally truncated to $250 \times 128 \times 127$ modes. This resolution ensures a memory footprint large enough to test performance under memory bandwidth saturation, yet it remains well within the total physical RAM of the compute node to prevent performance degradation from disk swapping.

The actual hardware-software interaction is revealed in the strong scaling results depicted in figure 9.2. For the explicit nonlinear computation ($\mathcal{N}$), one would theoretically expect ideal linear scaling since the pointwise multiplications require zero MPI communication. However, because its computation traverses the entire physical grid and involves heavy, continuous memory access across the full θ -pencil arrays, it rapidly saturates the system’s memory controllers at low core counts. This is reflected in the scaling curve, where the nonlinear evaluation speedup stalls severely between 2 and 32 processors. As a result, the

speedup of the physical-space operations initially stalls before recovering efficiently between 64 and 128 cores. This super-linear recovery is a direct cache-fitting consequence of the pencil decomposition: as P increases, the local sub-arrays shrink simultaneously in both partitioned dimensions until they reside comfortably within the low-latency L3 caches.

In contrast, the ETD advancement (**E, NL**) and projection (**P**) operators involve denser algebraic operations on smaller localized blocks. Consequently, they maintain near-ideal scaling from 2 to 8 cores before hitting a delayed efficiency plateau. This mid-range degradation is attributed to a combination of memory bandwidth saturation and minor computational load imbalances, arising from the indivisible number of spectral modes (e.g., the odd truncation limit of 127 axial modes) distributed unevenly across the MPI processor grid. Both operators similarly recover at higher core counts. The localized radial transforms (i.e., radial spectral expansion and reverse) maintain consistent scaling throughout. Meanwhile, the full space conversion ($\text{PPP} \leftrightarrow \text{FFF}$), which dictates the primary MPI communication cost due to the algorithm's pseudo-spectral nature, experiences an initial latency penalty from 2 to 4 cores, likely due to crossing internal NUMA or socket boundaries, before scaling efficiently at larger process counts. These scaling results demonstrate that, during the active time-stepping phase, the linear matrix advancements from ETD introduce no significant performance penalty. Because these operators act via localized block-diagonal multiplications in the spectral space, they scale highly efficiently and does not suffer from the communication overhead associated with global data transposes.

To provide a complete assessment of the ETD framework's computational footprint, it is also necessary to mention its initialization overhead. While the scheme requires pre-computing the integration operators during simulation startup, this associated cost is mathematically negligible. The full ETD framework requires diagonalization for each radial collocation point and azimuthal wavenumber; however, utilizing the analytical symbolic expressions (9.38) and (9.39) ensures the computational complexity of this initialization phase remains strictly $\mathcal{O}(N_r \cdot N_m)$, where N_r and N_m denote the number of radial collocation points and azimuthal wavenumbers, respectively.

Ultimately, this combination of limited initialization overhead and highly efficient localized time-stepping underscores the computational viability of the ETD framework. While the per-step computational and memory requirements remain comparable to conventional mixed IMEX schemes, such as AB2-CN, the framework is designed to provide a significant advantage through the expected allowance of larger time steps due to the analytical treatment of the stiff linear operators. This section's scaling results further demonstrate that the parallelization strategy is highly successful, with the robust recovery and near-ideal scaling achieved at higher process counts confirming the efficacy of the 2D pencil decomposition. By effectively distributing the domain such that the local sub-arrays fit entirely within the high-speed cache architecture of the processors, the numerical framework achieves excellent parallel efficiency and is well-optimized for production deployment on large-scale computing clusters.

9.5 Code Validation

To validate the numerical framework and evaluate the performance of the exponential time differencing (ETD) scheme, we use a stably stratified Lamb–Oseen vortex as the background shear flow. Without loss of generality, we consider a generic vortex with core radius $a = 1$ and circulation $\Gamma = 2\pi$, such that its azimuthal velocity profile is given by:

$$\bar{u}_\theta(r) = \frac{1 - e^{-r^2}}{r}. \quad (9.44)$$

For the numerical experiments detailed below, the Brunt–Väisälä frequency ranges from $\bar{N} = 0.1$ to 5.0. We further consider the vortex in a non-rotating environment ($\Omega = 0$). While the inclusion of a background rotation will fundamentally alter the stability landscape by modifying the Coriolis restoring forces, shifting the generalized Rayleigh discriminant $\Delta(r)$, and potentially introducing rotation-dependent mechanisms such as the stratorotational instability, setting $\Omega = 0$ serves two specific validation purposes. First, it isolates the baroclinic shear-stratification interaction, aligning our numerical setup precisely with the documented benchmark of S. Le Dizès (2008) for a clean physical cross-reference. Second, from an algorithm perspective, the background rotation enters the linearized governing equations exclusively as an additive constant to the background axial vorticity within the horizontal transformation matrix ($\mathbf{S}_H$ and $\mathbf{S}_H^{-1}$). Therefore, setting a zero background rotation does not simplify the algebraic structure, the cross-variable coupling, or the numerical stiffness of the linear operators. The full functional machinery of the ETD diagonalization is therefore fully exercised and validated even in the absence of system rotation.

Linear stability and spatial validation

We validate the linear spatial operators ($\mathcal{L}_{\text{wave}} + \mathcal{L}_{\text{cross}}$) and the global accuracy of the mapped Legendre discretization by computing the discrete global eigenmodes of the stratified Lamb–Oseen vortex. We emphasize that as the nonlinear operator $\mathbf{u}' \times \boldsymbol{\omega}'$ has been thoroughly validated in preceding studies utilizing the mapped associated Legendre basis (Matsushima and P. Marcus, 1997), the linear operators are the primary focus here, as they constitute the functional foundation of the ETD scheme.

Assuming normal mode disturbances of form (8.21), as given in §8.3, the linearized, inviscid Boussinesq equations form a generalized eigenvalue problem:

$$\sigma \tilde{\mathbf{v}} = \mathbf{L}(r, m) \tilde{\mathbf{v}} - \tilde{\boldsymbol{\Pi}}. \quad (9.45)$$

The velocity eigenvector, $\tilde{\mathbf{u}} = [\tilde{u}_r, \tilde{u}_\theta, \tilde{u}_z]^\top$, is subject to the linear solenoidal constraint:

$$\frac{1}{r} \frac{d}{dr}(r \tilde{u}_r) + \frac{im}{r} \tilde{u}_\theta + ik \tilde{u}_z = 0. \quad (9.46)$$

To enforce this divergence-free condition and simultaneously eliminate the pressure perturbation $\tilde{\Pi}$, the global eigenvalue matrix is constructed using the poloidal-toroidal formulation. Specifically, we define spectral test functions in the reduced scalar basis $(\tilde{\psi}, \tilde{\chi}, \tilde{b})$ and expand these potentials in the mapped Legendre space to reconstruct the corresponding primitive variables. The linear operator is then applied in the PFF space at each radial collocation point. Finally, the resulting state is projected back to the poloidal-toroidal form in the FFF space, yielding the corresponding columns of the global eigenvalue matrix.

At each collocation point, the local linear dynamics are governed by the purely algebraic operator $\mathbf{L}(r, m)$. Exploiting the analytical diagonalization derived for our ETD scheme in §9.3, this operator is evaluated via the similarity transformation:

$$\mathbf{L}(r, m) = \mathbf{S}(r) \mathbf{J}(r, m) \mathbf{S}^{-1}(r). \quad (9.47)$$

The diagonal matrix $\mathbf{J}(r, m)$ explicitly separates into a restorative component, which remains constant across all spatial wavenumbers (m, k) , and an advective Doppler shift that scales solely with the azimuthal wavenumber m :

$$\mathbf{J}(r, m) = \text{diag}(\mathrm{i}\omega_H^\pm(r), \pm \mathrm{i}\bar{N}) - \mathrm{i}m\bar{\Omega}(r)\mathbf{I}, \quad (9.48)$$

where the horizontal epicyclic frequencies are defined as $\omega_H^\pm(r) = \pm\sqrt{\Delta(r)}$ with $\Delta(r)$ being the generalized Rayleigh discriminant given by (4.11), and $\mathbf{I}$ is the 4×4 identity matrix. These epicyclic frequencies dictate the natural, inertial restoring oscillations of a fluid parcel displaced radially within the rotating background shear, and are governed entirely by the generalized Rayleigh discriminant. Because the constituent matrix $\mathbf{L}(r, m)$ is completely independent of the axial wavenumber k , our global eigenvalue solver can efficiently scan across a large parameter space without the need to re-calculate $\mathbf{L}$ every single time.

Before discussing the eigenvalue results, it is highly instructive to re-visit the large- k WKBJ framework (S. Le Dizès, 2008) detailed in §8.3, which helps understand the spatial structures of the calculated eigenmodes. Assuming small azimuthal wavenumber m , the local radial wavenumber $\beta(r)$, defined by (8.38), further reduces to:

$$\beta(r) = k \sqrt{\frac{\varphi^2(r) + \Delta(r)}{\varphi^2(r) + \bar{N}^2}}, \quad (9.49)$$

which indicates the local oscillatory or evanescent behavior of the wave and exhibit kinematic turning points r_t , defined by the roots of the numerator: $\varphi^2(r_t) + \Delta(r_t) = 0$, or equivalently $\varphi(r_t) = \mathrm{i}\omega_H^\pm(r_t)$, where the local Doppler-shifted wave frequency matches the local epicyclic frequencies encoded in our ETD operator. Furthermore, comparing the definition of the barotropic critical layer (8.42) with the definition of the turning points reveals that the frequency curve of the barotropic critical layer will always sit between the two epicyclic frequencies ($\omega_H^\pm$). As studied in S. Le Dizès (2008), for a typical ring mode of the Lamb–Oseen vortex under strong stratification, there are generally three such turning points, and the barotropic critical layer is located exactly inside the evanescent (negative β^2) region




Figure 9.3: Discrete unstable modes for a stratified Lamb–Oseen vortex at $\bar{N} = 5$. (a) Growth rates λ and (b) wave frequencies ω as functions of the scaled axial wavenumber $k/\bar{N}$. Open circles represent the eigenvalues computed by the present global spectral eigenvalue solver. Three points on the first mode branch ($k/\bar{N} = 1.4, 2.4, 4.8$) are distinctly marked for further structural analysis. Solid black lines denote the benchmark results of S. Le Dizès (2008) obtained via a shooting method with complex contour deformation. The numerical spatial discretization employs $N_r = 252$ radial collocation points, a radial spectral truncation of $N = 244$, and a mapping parameter $L = 6.0$. Only the first three unstable branches are plotted.

bounded by two of these turning points, implying that the turning points shield the active wave from the singular barotropic critical layer.

We benchmark our eigenvalue solver against the numerical results of S. Le Dizès (2008), who utilized a shooting method with complex contour integration to obtain the dispersion curves of the helical ($m = 1$) modes at $\bar{N} = 5$. As shown in figure 9.3, our global spectral eigenvalue solver accurately captures the three unstable branches. An error of less than 1% is achieved for all wavenumbers of the second ring mode, and for wavenumbers up to $k/\bar{N} = 25$ for the first ring mode. At higher wavenumbers, the wave frequencies of the first ring mode obtained by our real-axis solver begin to show minor spectral oscillations (with errors remaining below 2.5%). These minor deviations are a direct consequence of evaluating the linear critical layer singularity on the purely real collocation axis. As analyzed in detail below, this real-axis evaluation also governs the spectral behavior of the primary first mode, which undergoes a unique structural evolution that fully exposes this mathematical singularity.

As pointed out by S. Le Dizès (2008), the first mode (the most unstable branch) exhibits unique physical characteristics that separate it from the standard ring mode family. Figure 9.4 tracks the structural evolution of this mode across three distinct stages: before




Figure 9.4: Frequency maps (left column) and normalized spatial eigenmode profiles $\tilde{u}_\theta$ (right column) tracking the topological evolution of the first mode at the three distinct scaled wavenumbers, as marked in figure 9.3: (a) $k/\bar{N} = 1.4$, (b) $k/\bar{N} = 2.4$, and (c) $k/\bar{N} = 4.8$. The left panels plot the radial profiles of the local epicyclic frequencies (dashed magenta lines) and the co-rotation frequency (solid red line), alongside the mode's constant wave frequency ω (horizontal green line). The intersections of the horizontal wave frequency line with these respective curves define the radial locations of the turning points and the barotropic critical layer. At $k/\bar{N} = 4.8$, the loss of the inner turning points exposes the mode directly to the barotropic critical layer, triggering the Gibbs-like spectral ringing observed in the spatial profile.

its peak growth rate ($k/\bar{N} = 1.4$), at its peak ($k/\bar{N} = 2.4$), and after it visibly detaches from the theoretical curve ($k/\bar{N} = 4.8$).

Initially (figure 9.4a), the mode behaves similarly to a standard ring mode; its wave frequency intersects the ω_H^- curve at two distinct inner turning points and the ω_H^+ curve at an outer turning point. Because the barotropic critical layer ($\varphi = 0$) resides between the ω_H^+ and ω_H^- curves, it is located within the resulting evanescent region ($\beta^2 < 0$). Correspondingly, its mode profiles are smooth and well-resolved. In contrast, as the wavenumber increases and the mode reaches its peak growth rate (figure 9.4b), its wave frequency ω drops to the local minimum (valley) of the ω_H^- curve. At this juncture, the two inner turning points collide. Eventually, as the frequency drops further (figure 9.4c), these inner turning points completely vanish. Without this inner evanescent potential barrier, the oscillatory region extends entirely to the center of the vortex ($r = 0$), transforming the disturbance into a core mode. More critically, the removal of the turning points strips away the spatial shielding. The barotropic critical layer ($\varphi = 0$) is no longer hidden in the evanescent zone; it emerges directly inside the active oscillatory region as a naked mathematical singularity.

While S. Le Dizès (2008) and Riedinger et al. (2010) bypass the critical layer singularity using complex contours or mappings, and Sangjoon Lee and Philip S. Marcus (2023) optimizes the spatial mapping parameter L to smoothly resolve it, our spectral discretization is evaluated strictly on the real radial axis to serve the physical nonlinear ETD scheme. Attempting to fit the resulting singular, logarithmic phase jump associated with this unshielded critical layer using globally smooth polynomials naturally triggers the Gibbs-like spectral ringing observed in the $\tilde{u}_\theta$ profiles in figure 9.4c. This ringing directly accounts for the deviations from the numerical eigenvalues of S. Le Dizès (2008) at high wavenumbers. However, while critical layers are physical resonance phenomena, the mathematical singularity introduced at these locations exists only in the linear, inviscid limit. Because finite-amplitude nonlinear effects and viscosity inherently smoothen and regularize such resonances in full temporal simulations, complex mapping is physically unwarranted. As our eigenvalue solver is constructed strictly to validate our real-axis spatial operators, the observed spectral behaviors are entirely expected, and the real-axis resolution remains highly satisfactory for the target nonlinear physics.

Temporal convergence and energy conservation

To confirm the temporal accuracy of the ETD scheme, we conduct a fractional error convergence test. The combination of AB2 for nonlinear forcing extrapolation and CN for diffusion approximation, as given by (9.21), suggests a global temporal error of $\mathcal{O}(\Delta t^2)$ for the overall ETD scheme.

To strictly isolate the temporal truncation error without contamination from spatial aliasing or boundary discontinuities, the convergence test is performed on a band-limited flow field. We discretize the domain using a spatial resolution of $(N_r, N_\theta, N_z) = (160, 256, 128)$ with a radial mapping parameter of $L = 10.0$ and an axial domain length of $L_z = 10\pi$. To prevent nonlinear spectral blocking from polluting the temporal measurement,




Figure 9.5: Buoyancy field (b') at r - θ slice slightly above the perturbation center plane and r - z slice at $\theta = 245^\circ$ taken at $t = 90$ with $\Delta t = 0.1$ for (a) the ETD scheme that linearizes $\mathbf{u}' \times \bar{\zeta} \hat{\mathbf{z}}$ plus background rotation and stratification, and (b) the full ETD scheme that linearizes both cross terms plus background rotation and stratification. The baseflow is a Lamb–Oseen vortex with a core size of 4, and the Gaussian perturbation is initiated at $(r = 5, \theta = 45^\circ, z = L_z/2)$ with an aspect ratio of 0.26. The flow is subject to $\Omega = 0$ and $\bar{N} = 0.1$. The color intensity corresponds to the magnitude of the field in log-scale. Standard AB-CN scheme develops the same catastrophic high-frequency numerical artifacts observed in (a) but at a much faster pace.

the spectral coefficients are de-aliased using a truncation of $(N_n, N_m, N_k) = (152, 84, 42)$, and a hyperviscosity of order 6 is turned on to manage any energy pile-up at higher-order radial modes, as discussed in §9.4.

The flow is initialized with a low-amplitude, 3D localized off-axis Gaussian perturbation superposed on the columnar background flow. While the perturbation’s structure can be derived using the cyclo-geostrophic and hydrostatic balance framework to achieve a quasi-equilibrium (Mahdinia et al., 2017), we deliberately widen the axial length scale of the Gaussian envelope relative to a quasi-equilibrium. This artificial broadening ensures the perturbation is smoothly and fully resolved by the axial grid, circumventing the Gibbs ringing and spurious energy injections associated with under-resolved vertical gradients. Finally, a hyperviscosity filter is temporarily applied to the initial condition to scrub any residual spectral noise.

Before evaluating the formal accuracy of the solver, we first assess the stabilizing properties of the proposed analytical diagonalization. We compare the integration stability of the baseline AB2-CN method against two variants of the ETD scheme detailed in §9.3: a partial implementation that diagonalizes only the Coriolis-like restoring force ($\mathbf{u}' \times \bar{\zeta} \hat{\mathbf{z}}$) alongside the wave operator, and the full ETD formulation (9.37) that analytically integrates both advective cross terms. The simulations are advanced using a moderate time step of $\Delta t = 0.1$. To ensure that the numerical stability limit is strictly dominated by the advective shear rather than high-frequency restorative waves, the background state is configured with weak stratification ($\bar{N} = 0.1$) and zero background rotation ($\Omega = 0$) for this preliminary investigation.




Figure 9.6: Temporal convergence of the ETD scheme. The fractional errors in global kinetic energy (E_K) and available potential energy (E_{AP}) follow the theoretical second-order scaling ($\propto \Delta t^2$) as Δt is reduced, reaching an optimal temporal resolution at $\Delta t \approx 10^{-3}$. Below this threshold, term cancellation in the ETD integration weights causes round-off error to dominate. Errors are evaluated at $t = 20$, corresponding to approximately 3.18 core rotation periods ($T_{\text{core}} = 2\pi$) of the background Lamb–Oseen vortex, ensuring the measurement captures fully developed nonlinear dynamics.

As shown in figure 9.5a, explicitly evaluating the macroscopic stretching term ($\bar{\mathbf{u}} \times \boldsymbol{\omega}'$) while treating only the local restoring force analytically fails to resolve the numerical stiffness of the system. The partial ETD scheme generates a spurious, high-frequency ring band near the rotation axis, where the local angular velocity is highest, that rapidly corrupts the spatial domain and forces a numerical blow-up. The standard AB2-CN scheme exhibits identical structural artifacts but crashes significantly faster. This confirms that both components of the cross term fundamentally contribute to the advective stiffness of the governing equations. Conversely, as depicted in figure 9.5b, only the full ETD scheme that analytically integrates both cross terms maintains absolute stability without generating any trace of unphysical artifacts or dispersion.

Having established the necessity of the full analytical diagonalization for numerical stability, we monitor the kinetic energy (E_K) and the available potential energy (E_{AP}) to evaluate the temporal truncation error. For this test, the background stratification is increased to $\bar{N} = 1.0$ to ensure the presence of significant buoyancy dynamics and energy exchange, providing a more stringent assessment of the temporal accuracy across both the

momentum and buoyancy fields. The numerical evaluation of these energy terms is detailed in Appendix 9A. As depicted in figure 9.6, after integrating for more than three vortex turn-around times, the fractional errors for both E_K and E_{AP} decay with a slope of 2 on a log-log scale for time steps down to $\Delta t \approx 10^{-3}$. This confirms that the exact matrix exponentiation of the linear operators preserves the global second-order accuracy of the underlying AB2 and CN approximations without introducing stiff transients.

As Δt is reduced below 10^{-3} , the fractional error deviates from the theoretical second-order scaling and begins to increase, entering a round-off dominated regime. While reaching a precision floor at very small time steps is the expected, correct mathematical behavior for any discrete numerical recipe, the ETD scheme encounters this limit earlier than standard IMEX methods. This elevated floor arises during the direct floating-point evaluation of the nonlinear integration operator (**NL**). When evaluating terms proportional to $(e^z - 1)/z$ evaluated at $z = \lambda_{\mathbf{L}} \Delta t$ (where $\lambda_{\mathbf{L}}$ represents the eigenvalues of the diagonalized linear operator), an extremely small time step forces the exponential to approach unity, leading to catastrophic cancellation and a loss of significant digits in standard double-precision arithmetic.

To construct a mathematically airtight numerical architecture across all dynamical scales, the evaluation of the **NL** operator's diagonal components can be formalized as a three-tier strategy. First, exact stationary resonances ($\lambda_{\mathbf{L}} = 0$, representing critical layers and turning points) are resolved analytically via the L'Hôpital limit defined previously in (9.29). Second, for non-zero but small arguments ($0 < |z| \ll 1$), the catastrophic cancellation is formally resolved by substituting a Taylor series expansion,

$$\frac{e^z - 1}{z} = \sum_{l=1}^{\infty} \frac{z^{l-1}}{l!}, \quad (9.50)$$

a standard robust technique utilized in quasi-periodic ETD formulations to preserve precision near the zero-mode limits (Wilkening and Zhao, 2021). Finally, for $|z| \gtrsim 10^{-3}$, the exact discrete exponential formula is applied.

While incorporating this series expansion fully eliminates the premature precision floor, such treatment is omitted in the current convergence tests as it is unnecessary for our target applications. The fundamental purpose of the ETD framework here is to bypass the severe explicit stability restrictions that force traditional IMEX schemes into taking impractically small time steps. Crucially, the target phenomena of this solver, such as critical layer resonances and the zombie vortex instability, require the background shear, stratification, and rotation to be dynamically balanced and on the same order of magnitude. This phenomenological regime inherently lacks the extreme temporal scale separation and parasitic high-frequency waves that would otherwise dictate a microscopic temporal resolution. Therefore, to efficiently capture the slow, long-term nonlinear evolution of these instabilities, the ETD scheme is explicitly designed to operate safely and accurately at large time steps ($\Delta t \gg 10^{-3}$). Consequently, the early round-off flooring is effectively isolated from our operational regime and does not impact the practical application of the solver,

but the aforementioned treatment is necessary for general applications. Finally, we note that the absolute magnitude of the fractional error is higher for E_{AP} than for E_K . This is physically expected, as the explicit nonlinear advection of the strongly stratified buoyancy field ($\mathbf{u}' \cdot \nabla b'$) possesses a larger absolute truncation error prefactor than the corresponding momentum advection.

Beyond formal temporal accuracy, the nonlinear robustness of the ETD scheme is verified against the exact energy conservation laws given by (8.15). We track the energy balance of the Gaussian perturbation by monitoring the evolution of the kinetic and available potential energy relative to the cumulative shear production and the internal buoyancy energy exchange. These source and exchange terms are evaluated at each diagnostic step and subsequently integrated over time using the trapezoidal rule.

As shown in figure 9.7, the numerical scheme exhibits exceptional kinematic consistency. The bottom-right panel demonstrates that the net change in the total mechanical energy precisely matches the cumulative background shear production (PS), strictly adhering to (8.15c). Furthermore, the individual energy budgets are properly maintained: the change in the disturbance kinetic energy balances the sum of the shear production and the buoyancy energy exchange (top right, consistent with (8.15a)), while the change in the available potential energy mirrors the buoyancy exchange (bottom left, consistent with (8.15b)). Ultimately, this strict adherence to the fundamental energy balances confirms that the fully diagonalized spectral ETD formulation successfully suppresses spurious numerical artifacts, enabling stable, physically faithful nonlinear integration at large time steps.

Appendix 9A Calculation of Energy and Momentum with Spectral Decomposition

We established the conservation laws for the disturbance fields governed by (8.5) in §8.2. To validate the numerical recipes detailed in this chapter, we will need to evaluate the corresponding energy terms in their discrete spectral representations.

We begin by considering the available potential energy, which requires integrating the square of the buoyancy perturbation, $b'(r, \theta, z)$. In our spectral framework, this field is represented as a double Fourier series:

$$b'(r, \theta, z) = \sum_{m=-\infty}^{\infty} \sum_{\tilde{k}=-\infty}^{\infty} b'_{m\tilde{k}}(r) \exp \left\{ \left[i \left(m\theta + \frac{2\pi}{L_z} \tilde{k}z \right) \right] \right\}, \quad (9.51)$$

where m and $\tilde{k}$ are integers representing the azimuthal and axial wavenumbers. The integer $\tilde{k}$ is utilized explicitly to enforce the periodic boundary condition along the axial direction of length L_z . The total available potential energy is defined as the volume integral of $(b')^2$ scaled by the background stratification. Applying orthogonality in the periodic azimuthal and axial directions reduces the three-dimensional integral to a sum of one-dimensional radial




Figure 9.7: Energy balance tracking of the full ETD scheme during the nonlinear integration of the 3D Gaussian perturbation. To maintain inviscid stability and absorb the forward energy cascade, a sixth-order hyperviscosity filter is applied, with its amplitude dynamically adjusted based on the energy magnitude at the spectral tails. The figure comprises four panels: (top left) the temporal evolution of E_K , E_{AP} , and the cumulative shear production; (top right) the internal budget of E_K , explicitly incorporating the buoyancy energy exchange; (bottom left) the internal budget of E_{AP} , matching the buoyancy exchange; and (bottom right) the total mechanical energy balance. Note that the governing equations define the total energy evolution as $d_t(E_K + E_{AP}) = -\mathcal{E}_{\text{shear}}$; here, we denote $PS = -\mathcal{E}_{\text{shear}}$.

integrals:

$$\begin{aligned} E_{\text{AP}} &\equiv \frac{1}{2\bar{N}^2} \int_0^{L_z} \int_0^{2\pi} \int_0^\infty (b')^2 r dr d\theta dz \\ &= \frac{2\pi L_z}{\bar{N}^2} \sum_{\tilde{k}} \left\{ \sum_{m>0} \int_0^\infty |b'_{m\tilde{k}}(r)|^2 r dr + \frac{1}{2} \int_0^\infty |b'_{0\tilde{k}}(r)|^2 r dr \right\}, \end{aligned} \quad (9.52)$$

where we have utilized the conjugate symmetry for real-valued fields, $b'_{-m, -\tilde{k}}(r) = b'^*_{m\tilde{k}}(r)$.

To evaluate the continuous integral over r strictly within the spectral space, we apply the algebraic mapping (9.5) of the mapped associated Legendre polynomials, which can be rewritten as:

$$r = L\sqrt{(1+\mu)/(1-\mu)}, \quad (9.53)$$

where $\mu \in [-1, 1)$ and L is the mapping scale parameter. The differential area element transforms as

$$r dr = \frac{L^2}{(1-\mu)^2} d\mu, \quad (9.54)$$

allowing us to evaluate the radial integrals exactly utilizing Gauss-Legendre quadrature. By sampling at the collocation points μ_n with corresponding quadrature weights W_n , the total available potential energy is computed numerically as

$$E_{\text{AP}} = \frac{2\pi L_z L^2}{\bar{N}^2} \sum_{\tilde{k}} \sum_n \left\{ \sum_{m>0} \frac{|b'_{m\tilde{k}}(\mu_n)|^2}{(1-\mu_n)^2} + \frac{1}{2} \frac{|b'_{0\tilde{k}}(\mu_n)|^2}{(1-\mu_n)^2} \right\} W_n. \quad (9.55)$$

The above quadrature approach applies identically to the volumetric energy exchange between the disturbance buoyancy and the vertical velocity. Representing the vertical velocity u_z as a double Fourier series analogous to the buoyancy, the integrated total energy exchange evaluates to

$$\int_V (b' u_z) dV = 4\pi L_z L^2 \sum_{\tilde{k}} \sum_n \left\{ \sum_{m>0} \text{Re} \left[\frac{b'_{m\tilde{k}}(\mu_n) u_{z, m\tilde{k}}^*(\mu_n)}{(1-\mu_n)^2} \right] + \frac{1}{2} \text{Re} \left[\frac{b'_{0\tilde{k}}(\mu_n) u_{z, 0\tilde{k}}^*(\mu_n)}{(1-\mu_n)^2} \right] \right\} W_n. \quad (9.56)$$

While the available potential energy and the energy exchange rely on straightforward scalar integrals, evaluating the total kinetic energy requires handling the full perturbation velocity vector. This is achieved efficiently by exploiting the poloidal-toroidal decomposition (8.52). Let

$$\mathbf{u}' = \mathbf{u}_\psi + \mathbf{u}_\chi. \quad (9.57)$$

The total kinetic energy expands into the sum of the poloidal energy, the toroidal energy, and their cross interaction:

$$E_K = \frac{1}{2} \int_V (\mathbf{u}_\psi + \mathbf{u}_\chi)^2 dV = \frac{1}{2} \int_V (\mathbf{u}_\psi^2 + \mathbf{u}_\chi^2 + 2\mathbf{u}_\psi \cdot \mathbf{u}_\chi) dV. \quad (9.58)$$

The cross interaction integrates strictly to zero for decaying and periodic fields. Using standard vector identities and the property $\nabla \times \hat{z} = 0$, the integrand of the cross interaction expands as $\mathbf{u}_\chi \cdot \mathbf{u}_\psi = \nabla \cdot [(\partial_z \psi)(\nabla \chi \times \hat{z})]$. Because this interaction is a pure divergence, its volume integral evaluates to zero by Gauss's theorem. The contributions from the strictly poloidal and toroidal fields simplify independently via integration by parts. The toroidal contribution reduces to

$$\int_V \mathbf{u}_\chi^2 dV = - \int_V \chi \nabla_\perp^2 \chi dV, \quad (9.59)$$

and the poloidal contribution reduces to

$$\int_V \mathbf{u}_\psi^2 dV = \int_V \nabla_\perp^2 \psi \nabla^2 \psi dV. \quad (9.60)$$

Combining these, the baseline kinetic energy in terms of the scalar potentials is

$$E_K = \frac{1}{2} \int_V (-\chi \nabla_\perp^2 \chi + \nabla_\perp^2 \psi \nabla^2 \psi) dV. \quad (9.61)$$

As mentioned in §8.4 and discussed in Matsushima and P. Marcus (1997), specific global flows require logarithmic components in its poloidal-toroidal decompositions for completeness. In particular, the streamfunction decomposes into a localized component and a structurally imposed logarithmic tail: $\chi \equiv \chi_o + \chi_l P_l(r)$, where $P_l(r) \sim \ln(r)$. Inserting this decomposition into the toroidal energy integral produces a self-energy term for the macroscopic tail, given by $-\chi_l^2 \int_V P_l(r^{-1} \partial_r (r \partial_r P_l)) dV$. Integrating this term over an unbounded radial domain yields a logarithmic divergence. To obtain a physically meaningful measure of the perturbation energy that characterizes the instability, we evaluate the regularized kinetic energy. This regularization discards the divergent self-energy of the unperturbed $1/r$ tail while explicitly retaining the finite cross interaction that dictates the physical momentum exchange between the localized disturbance χ_o and the asymptotic boundary structure $\chi_l P_l$:

$$\widetilde{E}_K = \frac{1}{2} \int_V \{-\chi_o \nabla_\perp^2 (\chi_o + 2\chi_l P_l) + \nabla_\perp^2 \psi \nabla^2 \psi\} dV. \quad (9.62)$$

Beyond the energy budgets, the discrete spectral framework also facilitates the exact evaluation of the domain-integrated momenta. The total vertical momentum of the disturbance, $\mathcal{P}_z$, evaluates entirely to the axisymmetric and vertically invariant $m = 0, \tilde{k} = 0$ Fourier mode:

$$\mathcal{P}_z = \int_V u_z dV = 2\pi L_z \int_0^\infty u_{z,00}(r) r dr = 2\pi L_z L^2 \sum_n u_{z,00}(\mu_n) \frac{W_n}{(1 - \mu_n)^2}. \quad (9.63)$$

Because the vertical velocity is related to the poloidal streamfunction via $u_z = \nabla_\perp^2 \psi$, we can define a mapped spectral function

$$h(\mu) \equiv \nabla_\perp^2 \psi_{00} / (1 - \mu)^2. \quad (9.64)$$

Expanding h in the standard Legendre basis, $h(\mu) = \sum_n h_n P_n(\mu)$, the momentum integral simplifies exactly to the zeroth Legendre coefficient, yielding

$$\mathcal{P}_z = 4\pi L_z L^2 h_0. \quad (9.65)$$

Similarly, the total axial angular momentum, $\mathcal{L}_z$, isolates strictly to the zeroth mode of the radial derivative of the toroidal streamfunction:

$$\mathcal{L}_z = \int_V r u_\theta dV = \int_V r \left(-\partial_r \chi + \frac{1}{r} \partial_\theta \partial_z \psi \right) dV = -2\pi L_z \int_0^\infty r^2 d_r \chi_{00}(r) dr. \quad (9.66)$$

Finally, the spectral formulation explicitly demonstrates the non-conservative nature of the standard potential energy in a periodic domain, as anticipated theoretically. If evaluated directly, the standard potential energy (K_P) retains an explicit dependence on the domain boundaries:

$$K_P = \int_0^{L_z} \int_0^{2\pi} \int_0^\infty z b'(r, \theta, z) r dr d\theta dz = 2\pi \sum_{\tilde{k}} \left(\int_0^\infty b'_{0\tilde{k}}(r) r dr \right) \int_0^{L_z} z \exp\left(i \frac{2\pi}{L_z} \tilde{k} z\right) dz. \quad (9.67)$$

Integration by parts on the axial coordinate reveals the non-zero boundary evaluation:

$$\int_0^{L_z} z \exp\left(i \frac{2\pi}{L_z} \tilde{k} z\right) dz = \left[\frac{L_z}{i2\pi\tilde{k}} z \exp\left(i \frac{2\pi}{L_z} \tilde{k} z\right) \right]_0^{L_z} - \int_0^{L_z} \frac{L_z}{i2\pi\tilde{k}} \exp\left(i \frac{2\pi}{L_z} \tilde{k} z\right) dz = \frac{L_z^2}{i2\pi\tilde{k}}. \quad (9.68)$$

Because this term is strictly non-zero for periodic modes where $\tilde{k} \neq 0$, the standard potential energy inherently depends on the artificial vertical domain scale rather than strictly the internal physics. This reinforces the necessity of utilizing the available potential energy when quantifying the physical evolution of the stratified disturbance.

Appendix 9B Mathematical Identities in Spectral Basis

We establish the mathematical identities utilized for evaluating the spectral volume integrations here. Specifically, we demonstrate that for generalized vector fields satisfying appropriate boundary conditions, the volume integrals of certain rotational operators evaluate exactly to zero.

First, consider the volume integral of $\hat{\mathbf{z}} \cdot \nabla \times (\mathbf{A} \times \nabla \times \mathbf{B})$ for two arbitrary vector fields $\mathbf{A}$ and $\mathbf{B}$. Utilizing standard vector calculus identities, the axial component of this operator can be expressed exactly as the difference of two pure divergence terms:

$$\hat{\mathbf{z}} \cdot \nabla \times (\mathbf{A} \times \nabla \times \mathbf{B}) = \nabla \cdot (A_z \nabla \times \mathbf{B}) - \nabla \cdot ((\nabla \times \mathbf{B})_z \mathbf{A}). \quad (9.69)$$

Integrating this quantity over the domain volume V and applying Gauss's theorem reduces the evaluation strictly to surface boundary fluxes. As noted by Matsushima and P. Marcus

(1997), provided that the vector fields are regular at the origin and decay to zero as $r \rightarrow \infty$, the net boundary flux vanishes completely, yielding:

$$\int_V \hat{\mathbf{z}} \cdot \nabla \times (\mathbf{A} \times \nabla \times \mathbf{B}) dV = 0. \quad (9.70)$$

Furthermore, in the specific case where $\mathbf{A}$ represents a strictly azimuthal background flow ($\mathbf{A} = A_\theta(r)\hat{\boldsymbol{\theta}}$), the axial component A_z is zero, and the remaining divergence term simplifies to $-r^{-1}\partial_\theta((\nabla \times \mathbf{B})_z A_\theta(r))$. When integrated over the full cylindrical volume, the expression (9.70) also evaluates exactly to zero due to azimuthal periodicity, irrespective of whether $A_\theta(r)$ decays in the far field.

Similarly, we evaluate the domain integral involving the background rotation, shear, and buoyancy. For a solenoidal perturbation velocity $\mathbf{u}'$ and buoyancy perturbation b' , we define the composite vector field

$$\mathbf{V} = (2\Omega + f(r)) \hat{\mathbf{z}} \times \mathbf{u}' + b' \hat{\mathbf{z}}. \quad (9.71)$$

Utilizing the identity $\hat{\mathbf{z}} \cdot \nabla \times \mathbf{V} = \nabla \cdot (\mathbf{V} \times \hat{\mathbf{z}})$ and applying the vector triple product, the axial curl reduces analytically to a strictly horizontal divergence:

$$\hat{\mathbf{z}} \cdot \nabla \times \mathbf{V} = \nabla \cdot [(2\Omega + f(r)) \mathbf{u}'_\perp], \quad (9.72)$$

where $\mathbf{u}'_\perp = u'_r \hat{\mathbf{r}} + u'_\theta \hat{\boldsymbol{\theta}}$ represents the horizontal velocity. Integrating this exact divergence over the domain volume yields:

$$\int_V \hat{\mathbf{z}} \cdot \nabla \times ((2\Omega + f(r)) \hat{\mathbf{z}} \times \mathbf{u}' + b' \hat{\mathbf{z}}) dV = \int_V \nabla \cdot [(2\Omega + f(r)) \mathbf{u}'_\perp] dV. \quad (9.73)$$

By the divergence theorem, this volume integral reduces entirely to horizontal boundary fluxes evaluated at $r = 0$ and $r \rightarrow \infty$. Because the physical velocity field must be regular at the origin and decay to zero in the far field, the radial boundary fluxes are identically zero. Consequently, the integral evaluates to zero independent of the axial boundary conditions.

Chapter 10

Baroclinic Critical Layers and the Zombie Vortex Instability

The preceding chapters established the mathematical formulations and developed the numerical infrastructure required to simulate wave-mean flow interactions in rotating, stratified environments. With the global conservation laws defined and the ETD pseudo-spectral scheme validated, we now apply this framework to investigate the physical mechanisms driving the ZVI.

While substantial analytical and numerical progress has been made using the local shearing box approximation (Barranco and Philip S. Marcus, 2006; Philip S. Marcus, Pei, Jiang, and Barranco, 2016; Philip S. Marcus, Pei, Jiang, and Hassanzadeh, 2013; Wohlever et al., 2025), explanations for the regular vortex spacing and the dominance of the zero-frequency wave mode have remained heavily tied to local Cartesian coordinates. To confirm whether the ZVI is a robust physical phenomenon capable of existing in laboratory experiments, we must transition to a global perspective. This chapter analyzes the limitations of previous Cartesian interpretations and introduces the frequency-resetting chain reaction hypothesis to explain the self-replication of zombie vortices. We construct a multi-generational mock-up to bridge local and global dynamics, detail the deployment of our global cylindrical simulations, employ computational ray tracing to analyze the resulting wave dynamics, and discuss the physical shear requirements for sustaining multi-generational vortex lattices.

10.1 The Illusion of the Zero-Frequency Mode in Shearing Coordinates

A prominent feature of the ZVI observed in previous local Cartesian studies is the formation of a uniform lattice of vortices (see figure 10.1), indicating a regular, distinct spatial gap between successive generations. In addition to this uniform spatial spacing, these studies identified a dominant instability mode characterized by a temporal frequency of exactly zero




Figure 10.1: Snapshots of zombie vortices in the y - z plane at $x = 0$ marching outward from the initial non-equilibrium Gaussian disturbance field at the origin. Anticyclonic vertical vorticity ω_z is black and cyclonic is white. $N/f = 1$; $\sigma_s/f = -3/4$; and the unit of length in the figure is $4L_x/(3\pi)$, where L_x is the streamwise length of the computational domain, which is periodic in x . Reproduced from Philip S. Marcus, Pei, Jiang, and Hassanzadeh (2013).

($\omega = 0$). Both the geometric regularity of the lattice and the stationarity of the primary mode have been central to Cartesian interpretations of the instability.

However, this zero-frequency selection is not an intrinsic dynamical property of the instability in an inertial frame. To simulate these dynamics, Barranco and Philip S. Marcus (2006) and Philip S. Marcus, Pei, Jiang, and Hassanzadeh (2013) solved the equations of motion in a time-dependent shearing coordinate system (x', y', z', t') , where $x' = x - \sigma_s y t$ and σ_s is the constant background shear rate. This mathematical transformation effectively absorbs the background advection term into the definition of the coordinate itself. As a consequence, the condition $\omega = 0$ in the shearing coordinate is mathematically equivalent to the statement that the perturbation's phase speed exactly matches the local background velocity. When transformed back to the inertial laboratory frame, this corresponds to a non-zero, advective wave frequency $\omega = -k_x \bar{u}_y(y_c)$ or $\omega = -m \bar{\Omega}(r_c)$. Thus, the 'stationary' modes of the shearing box are a kinematic artifact of the coordinate transformation, and are physically identical to advected wave packets in a global system.

This coordinate artifact initially contributed to the formulation of the ‘periodic realignment’ hypothesis (Wohlever et al., 2025), which posited that wave-mean flow resonance occurs when the differential shear shifts the wave phase by exactly one domain length L_x during one buoyancy period. While this geometric derivation successfully recovers the correct uniform spatial gap between vortex generations,

$$|\Delta y_c| = \frac{\bar{N}}{\sigma_s k_x}, \quad (10.1)$$

the explicit invocation of the domain length L_x is a geometric artifact of the shearing box. When the definition of the streamwise wavenumber $k_x = 2\pi m/L_x$ is substituted into the resonance condition, the domain length L_x mathematically cancels out.

This cancellation is profound: it reveals that the instability’s length scale is entirely independent of the domain size, implying that global periodicity is not the dynamical driver of the spacing. The mechanism must be inherently local, governed strictly by the local shear profile and stratification, rather than the global domain topology.

10.2 The Frequency-Resetting Chain Reaction Hypothesis

To resolve the discrepancy between local periodic models and global non-uniform flows, we propose that the ZVI is driven by a ‘frequency-resetting’ chain reaction. In this framework, the self-replication of zombie vortices is a discrete, iterative process reliant on local barotropic amplification and baroclinic trapping via critical layers, where the wave frequency ω is reset at each generation. The cycle for a single generation proceeds as follows:

1. Barotropic Emission: A coherent vortex structure at location r_n (or y_n in Cartesian coordinates) is advected by the local mean flow. By perfectly matching the local advection frequency, the perturbation satisfies the condition for a barotropic critical layer at its source:

$$\omega_n = -m\bar{\Omega}(r_n) \quad \text{or} \quad \omega_n = -k_x \sigma_s y_n. \quad (10.2)$$

As established in our linear stability analysis in §8.3, this resonance isolated via the self-adjoint wave action flux allows the perturbation to actively extract kinetic energy from the background shear via over-reflection mechanisms, resulting in the continuous emission of an amplified inertio-gravity wave packet.

2. Baroclinic Trapping: The emitted wave packet propagates radially until it encounters a baroclinic critical layer at the neighboring location r_{n+1} . This spatial location is defined where the intrinsic (Doppler-shifted) wave frequency matches the buoyancy frequency $\bar{N}$:

$$\omega_n + m\bar{\Omega}(r_{n+1}) = \pm\bar{N}. \quad (10.3)$$

Substituting the barotropic emission frequency (10.2) into this relation yields the exact spatial gap condition:

$$m(\bar{\Omega}(r_n) - \bar{\Omega}(r_{n+1})) = \pm\bar{N}. \quad (10.4)$$

At this radius, the radial group velocity of the wave stalls. The horizontal kinetic energy of the propagating wave funnels entirely into the vertical velocity component, leading to the localized, secular accumulation of available potential energy.

3. Nonlinear Roll-up and Frequency Reset: As the trapped energy accumulates at r_{n+1} , it modifies the mean flow to create a jet-like defect. As analyzed by C. Wang and Balmforth (2021), this defect eventually becomes susceptible to a secondary shear instability, causing the layer to roll up into a new coherent vortical structure. Crucially, this new structure is a material object entrained by the local flow. It effectively abandons the original wave frequency ω_n of its parent and emits a *new* wave packet at a frequency determined strictly by its new local environment: $\omega_{n+1} = -m\bar{\Omega}(r_{n+1})$.

This frequency reset implies that the spacing of ZVI bands is determined by a constant step in the baseflow’s angular velocity for the global cylindrical coordinate:

$$\Delta\bar{\Omega}(r) = \pm \frac{\bar{N}}{m}, \quad (10.5)$$

or a constant spatial step in the uniform background shear flow for the local shearing box ($\sigma_s \Delta y_c = \pm \bar{N}/k_x$). Consequently, in global cylindrical flows where the shear varies with radius, this mechanism predicts a non-uniform spatial banding, contrasting directly with Cartesian uniform lattices.

To demonstrate this hypothesis, we construct a multi-generational ray-tracing mock-up using the local Cartesian parameters of Philip S. Marcus, Pei, Jiang, and Hassanzadeh (2013). In this conceptual model, we terminate a WKBJ ray when it reaches a baroclinic critical layer, and restart a new ray from that exact location using a reset frequency matched to the local mean flow. As shown in figure 10.2, the resulting recursive lattice flawlessly reconstructs the spatial distribution of zombie vortices observed in the original direct numerical simulations. This kinematic equivalence serves as a strong evidence that the ZVI is driven by localized wave-mean flow interactions rather than global periodic boundary conditions.

10.3 Deploying Global Nonlinear Simulations

To test the frequency-hopping hypothesis via global numerical simulation, the flow must be initialized with a robust, long-lasting perturbation capable of serving as a continuous wave source. Injecting an arbitrary, unbalanced velocity or density field results in the rapid emission of transient acoustic and gravity waves. This rapid dispersion severely depletes the localized energy before any sustained wave-mean flow resonance can establish.

Therefore, we initialize the simulation with a localized Gaussian disturbance $(p', \rho', \mathbf{u}')$ mathematically derived to be in a state of cyclo-geostrophic and hydrostatic (CG-H) balance relative to the columnar background flow. Adapting the theoretical framework of Hassanzadeh et al. (2012) to a sheared cylindrical background (the full derivation of which is detailed




Figure 10.2: (a) Mock-up zombie vortex lattices constructed using the frequency-hopping hypothesis. The initial disturbance is assumed to be a point at $y = 0.5$, $z = 0$. The coloring corresponds to WKBJ ray trajectories starting from the previous generation's advection frequency and terminating at the corresponding baroclinic critical layers. Due to mirroring symmetries, each zombie vortex emits exactly four ray paths. (b) The predicted zombie lattice locations. The vertical spacing between neighboring lattices is exactly $\bar{N}$, perfectly mirroring the nonlinear DNS results of Philip S. Marcus, Pei, Jiang, and Hassanzadeh (2013).

in Appendix 10A), the aspect ratio of the initial vortex $\alpha' = H'/L'$ must satisfy:

$$(\alpha')^2 = \frac{Ro' (1 + \sqrt{2} Ro') f_{\text{eff}}^2}{(N'_c)^2 - \bar{N}^2}, \quad (10.6)$$

where Ro' is the local Rossby number, f_{eff} is the effective Coriolis parameter at the central radius, and N'_c is the perturbation buoyancy frequency. Initializing the global velocity components from this balance ensures that the injected disturbance survives long enough to dynamically interact with the background flow. Because of its finite spatial extent, this coherent vortex naturally emits a broad, continuous spectrum of inertio-gravity waves into the domain. This provides the necessary seed perturbations, allowing the flow itself to naturally select and amplify the resonant frequencies that trigger the ZVI.

Using the ETD scheme, the nonlinear simulation of a Lamb–Oseen vortex was executed with this balanced disturbance centered at $r_0 = 12$. The simulation successfully captured one of the primary mechanics of the chain reaction. As the injected perturbation advected with the local background flow and emitted its broad spectrum of inertio-gravity waves, our analysis revealed the crucial selection dynamic: the instability mechanism actively selected specific waves from this continuum. The emitted waves whose frequencies perfectly matched




Figure 10.3: u'_z -field in r - z slice taken at $t = 20000$ with ray tracing overlays. The background flow is a scaled Lamb–Oseen vortex with an intensity of 4, a core radius of 8, and a linear stratification of $\bar{N} = 0.07$. The initial quasi-equilibrium Gaussian vortex is geometrically centered at $r_0 = 12$. The ray tracings are performed with the wave frequency matching the background flow speed at (a) $r_0 = 12.3$, (b) $r_0 = 11.0$, and (c) $r_0 = 13.0$ for azimuthal wavenumbers up to $m = 10$. Trajectories align strictly with the effective advection frequency at $r_0 = 12.3$.

the local advection frequency near the vortex core satisfied the condition for a barotropic critical layer. Consequently, the local flow environment effectively turned the disturbance into a barotropic amplifier, continuously extracting energy from the background shear to fuel the emission of those specific, resonant inertio-gravity wave packets. These amplified wave packets propagated radially and stalled precisely at the spatial locations dictated by the baroclinic critical layer resonance condition ($\Delta\bar{\Omega}(r) = \pm\bar{N}/m$). At these radii, secular accumulation of vertical kinetic and available potential energy was visually documented, validating the local trapping mechanism defined in step 2 of our hypothesis.

To verify the advective frequency selection driving this trap, we overlay WKBJ ray tracing directly onto the numerical data extracted from the global simulation. The trajectory of a wave packet in the meridional plane is determined by the group velocity ratio:

$$\frac{dz}{dr} = \Re \left[\frac{\partial_k D}{\partial_{k_r} D} \right] = \Re \left[\frac{2k\varphi \frac{\varphi^2 + \Delta}{\varphi^2 + \bar{N}^2}}{2k_r\varphi - m \frac{d}{dr} \bar{\Omega}} \right]. \quad (10.7)$$

As shown in figure 10.3, the ray trajectories align precisely with the wave propagation patterns in the simulated vertical velocity (u'_z) field only when calculated using a frequency locked to an effective advection radius of $r_{\text{eff}} \approx 12.3$. This slight radial shift from the geometric center ($r_0 = 12$) occurs because the finite radial width of the initial Gaussian perturbation excites a continuum of frequencies. Within this bandwidth, the system pref-

entially amplifies the mode corresponding to $r_{\text{eff}} \approx 12.3$ likely due to more favorable local shear gradients for barotropic over-reflection.

To test the sensitivity of this selection, trajectories were computed based on frequencies near the periphery of the disturbance ($r = 11.0$ and $r = 13.0$). These alternative trajectories diverge significantly from the numerical data. Furthermore, computing trajectories for a hypothetical quasi-steady wave ($s = 0$) relative to the global inertial frame yields negative curvature in the upper half of the r - z plane, fundamentally contradicting the positive curvature observed in the simulation. This definitively confirms that the wave packet frequency is not an arbitrary global eigenvalue, but is strictly selected by the local flow conditions at the source.

Despite successfully capturing the barotropic emission and baroclinic trapping, the secondary phase of the chain reaction, the generation of a second-generation zombie vortex, did not manifest within the simulated timeframe. Instead, the wave energy remained trapped without additional roll-up. It is crucial to distinguish this localized wave emission and trapping process from the global linear radiative instabilities analyzed in past literature (e.g., S. Le Dizès, 2008; Riedinger et al., 2010). The radiative instability framework relies on the assumption of strong stratification and fundamentally does not account for the existence of a baroclinic critical layer. In our system, while the initial wave propagation adheres to linear wave dynamics outside the critical layers, the baroclinic critical layer acts as a perfect wave trap. This trapping creates a singularity in the inviscid linear equations, causing global linear normal mode analysis to break down. Therefore, the accumulation of energy at the baroclinic trap and its subsequent evolution represents a complex wave-mean flow interaction that cannot be described by the radiative instability formulations.

Because linear theory cannot resolve the ultimate fate of this singular trapped wave, the ZVI framework hypothesizes a strictly nonlinear mechanism, where the local background shear drives a secondary roll-up of the defect to continue the chain reaction. The absence of this secondary roll-up in our global simulation can be directly attributed to the specific kinematic structure of the Lamb–Oseen vortex. Since its azimuthal shear decays rapidly ($\sim r^{-2}$), the local shear gradient at the outer radii, where the baroclinic critical layers reside, drops below the energetic threshold required to trigger the secondary shear instability evaluated by C. Wang and Balmforth (2021). Consequently, the chain reaction is effectively arrested after the primary wave trapping. To incite higher-generation zombie replication in future numerical studies or laboratory experiments, it may be necessary to utilize background flow profiles that maintain a sustained shear region over a broader radial extent. Such profiles could potentially provide the secondary baroclinic traps with sufficient local kinetic energy to overcome stratification, offering a rigorous pathway to test whether this mechanism truly drives the complete, multi-generational nonlinear roll-up of the ZVI in a global domain.

10.4 Conclusions and Future Work

This dissertation investigated the multi-faceted dynamics of hydrodynamic stability, bridging the fundamental theory of wave-wave resonances in incompressible environments to the complex wave-mean flow interactions driving instability in rotating, stratified systems.

Part I performed a multi-scale perturbation analysis of triadic resonance in generic columnar vortices. By formulating the disturbance equations via a poloidal-toroidal decomposition and deriving the general amplitude equations, we established the strict selection rules that distinguish between bounded oscillatory energy exchange and explosive instability. We proved that for smooth neutral modes, triadic resonance is always conservative and strictly obeys the Manley–Rowe relations. Utilizing wave pseudoenergy within a large- k WKBJ framework, we demonstrated that these conservative triads are confined exclusively to the bounded solution branch, effectively prohibiting intrinsic explosive instability among smooth modes. By relaxing the isolation constraint to investigate generalized parametric instability, we ultimately established that critical layer singularities are an absolute necessity to break these selection rules. These critical layers introduce essential non-Hermitian effects, providing the necessary pathway for waves to access the background shear energy and trigger instability.

Part II transitioned to the global wave-mean flow resonances inherent to stratified fluids, focusing specifically on the mechanics of the Zombie Vortex Instability (ZVI) and the introduction of the three-dimensional baroclinic critical layer. To overcome the severe numerical stiffness imposed by the three-dimensional Boussinesq equations in cylindrical coordinates, we developed a highly specialized, parallelized Exponential Time Differencing (ETD) pseudo-spectral scheme. By analytically diagonalizing the coupled linear operators in the primitive velocity space, the ETD framework natively embeds the physical resonance characteristics of the flow directly into the discrete integration operators. This approach successfully decoupled the numerical stability limit from the fast background kinematics, thereby bypassing the geometric constraints of traditional local shearing box approximations and enabling robust global simulations.

Applying this computational framework, we demonstrated that the uniform spacing of zombie vortices observed in prior literature is not a geometric artifact of Cartesian periodicity, but rather the consequence of a localized wave-mean flow resonance. We proposed the frequency-resetting chain reaction hypothesis: a coherent disturbance acts as a barotropic amplifier, continuously emitting inertio-gravity waves that propagate radially until they are trapped at a baroclinic critical layer. The secular accumulation of energy at this trap is hypothesized to drive a secondary nonlinear roll-up, creating a new vortex that resets the emission frequency based on its local advective environment. Our global cylindrical simulations successfully confirmed the initial barotropic emission, the highly specific advective frequency selection, and the subsequent spatial wave trapping at the baroclinic critical layer. While the instability arrests after this primary trap in our numerical tests, this outcome does not contradict the proposed framework; rather, it remains strictly consistent with the speculation that full multi-generational replication requires sustained background shear to fuel the secondary roll-up.

Future work will extend the developed ETD numerical framework to explore background flow profiles featuring sustained shear gradients. This configuration will provide the necessary local kinetic energy at secondary trapping locations to rigorously test whether our proposed mechanism can drive the complete multi-generational roll-up of the ZVI in a global domain. Additionally, subsequent investigations will incorporate radial density stratification into the physical model. The inclusion of radial density gradients is highly relevant to the dynamics of wake vortices, where the interplay between centrifugal forces and radial buoyancy fundamentally modifies the energy exchange at the barotropic critical layer, critically altering wave propagation and the overall destabilization mechanics. Ultimately, validating these localized critical layer interactions in cylindrical geometries provides the foundational theory required to bridge the gap toward experimental verification, extending these instability mechanisms to complex astrophysical and engineering flows.

Appendix 10A Initialization of a Quasi-Equilibrium Gaussian Vortex Perturbation

To initialize a long-lasting, non-axisymmetric perturbation that does not rapidly dissipate into transient gravity waves, we adapt the theoretical framework of Hassanzadeh et al. (2012) to derive a state of cyclo-geostrophic and hydrostatic (CG-H) balance embedded within a rotating background shear flow under constant vertical stratification ($\bar{N}$). While the original framework derived a universal aspect ratio for an isolated, quasi-equilibrium vortex in a purely hydrostatic background state, we extend this formulation to obtain the properties of a localized disturbance in CG-H balance relative to a sheared background flow in global cylindrical coordinates. This balance ensures the perturbation remains coherent long enough to dynamically interact with the flow and trigger the desired instability.

We formulate the pressure perturbation p' as a function strictly of the local coordinates originating from the vortex center, $(r_{\text{local}}, z_{\text{local}})$. Applying the quasi-equilibrium assumption ($\partial_t \approx 0$) and assuming the radial and axial velocity components are negligible ($u'_r \approx u'_z \approx 0$), the perturbation Boussinesq momentum equations simplify to a scalar force balance in the local radial direction and a standard hydrostatic balance vertically:

$$\frac{\partial p'}{\partial r_{\text{local}}} = \rho_0 u'_{\theta, \text{local}} \left[f_{\text{eff}} + \frac{u'_{\theta, \text{local}}}{r_{\text{local}}} \right], \quad (10.8)$$

$$\frac{\partial p'}{\partial z_{\text{local}}} = -\rho' g, \quad (10.9)$$

where f_{eff} is the effective Coriolis parameter at the perturbation's central radius r_0 :

$$f_{\text{eff}} = 2 \left(\Omega + \bar{\Omega}(r_0) \right). \quad (10.10)$$

Integrating the hydrostatic balance (10.9) yields an expression for the central pressure anomaly:

$$p'_c = \rho_0 H'^2 \left(\bar{N}^2 - (N'_c)^2 \right), \quad (10.11)$$

where H' is the characteristic vertical scale of the perturbation and N'_c is the perturbation buoyancy frequency at the vortex center. Conversely, performing an approximate integration of the radial balance equation (10.8) by evaluating the terms at their characteristic scales relates the central pressure anomaly to the characteristic azimuthal velocity, u'_{char} , occurring at the radius of maximum velocity, R_v :

$$-\frac{p'_c}{L'} \approx \rho_0 u'_{\text{char}} \left(f_{\text{eff}} + \frac{u'_{\text{char}}}{R_v} \right), \quad (10.12)$$

where L' is the characteristic horizontal length scale of the pressure anomaly.

Defining the aspect ratio $\alpha' \equiv H'/L'$ and introducing the local Rossby number as

$$Ro' \equiv \frac{u'_{\text{char}}}{L' \cdot f_{\text{eff}}}, \quad (10.13)$$

we equate the two expressions for p'_c to arrive at the general aspect ratio formula:

$$(\alpha')^2 \approx \frac{Ro' \left(1 + Ro' \cdot \frac{L'}{R_v} \right) f_{\text{eff}}^2}{(N'_c)^2 - \bar{N}^2}. \quad (10.14)$$

Evaluating (10.14) requires a relationship between the pressure scale L' and the velocity scale R_v , which is dictated by the specific structural model chosen for the vortex. For our numerical experiments, we impose a Gaussian pressure anomaly:

$$p'(r_{\text{local}}, z_{\text{local}}) = p'_c \exp \left(-\frac{r_{\text{local}}^2}{L'^2} - \frac{z_{\text{local}}^2}{H'^2} \right). \quad (10.15)$$

The local azimuthal velocity is obtained by linearizing the radial balance equation for a weak perturbation:

$$u'_{\theta, \text{local}} \approx \frac{1}{\rho_0 f_{\text{eff}}} \frac{\partial p'}{\partial r_{\text{local}}}, \quad (10.16)$$

which yields the velocity profile:

$$u'_{\theta, \text{local}} \propto r_{\text{local}} \exp(-r_{\text{local}}^2/L'^2). \quad (10.17)$$

Differentiating this profile to find the extremum determines the radius of maximum velocity: $R_v = L'/\sqrt{2}$. Substituting this geometric relation into (10.14) yields the final aspect ratio formula for the initialized Gaussian vortex:

$$(\alpha')^2 = \frac{Ro' (1 + \sqrt{2} Ro') f_{\text{eff}}^2}{(N'_c)^2 - \bar{N}^2}. \quad (10.18)$$

Consequently, the buoyancy perturbation $b' = -(1/\rho_0) \partial p' / \partial z_{\text{local}}$ and the local azimuthal velocity component are explicitly defined in the local coordinate frame as:

$$b'(r_{\text{local}}, z_{\text{local}}) = 2z_{\text{local}}(\bar{N}^2 - (N'_c)^2) \exp\left(-\frac{r_{\text{local}}^2}{L'^2} - \frac{z_{\text{local}}^2}{H'^2}\right), \quad (10.19a)$$

$$u'_{\theta, \text{local}}(r_{\text{local}}, z_{\text{local}}) = \frac{H'^2(\bar{N}^2 - (N'_c)^2)}{f_{\text{eff}}} \left(-\frac{2r_{\text{local}}}{L'^2}\right) \exp\left(-\frac{r_{\text{local}}^2}{L'^2} - \frac{z_{\text{local}}^2}{H'^2}\right). \quad (10.19b)$$

Finally, projecting (10.19) onto the global basis $(\hat{\mathbf{r}}, \hat{\boldsymbol{\theta}})$ provides the disturbance velocity components in the global cylindrical coordinates required for simulation initialization:

$$u'_r(r, \theta) = u'_{\theta, \text{local}} \left(\frac{r_0}{r_{\text{local}}} \sin(\theta_0 - \theta) \right), \quad (10.20)$$

$$u'_\theta(r, \theta) = u'_{\theta, \text{local}} \left(\frac{r - r_0 \cos(\theta - \theta_0)}{r_{\text{local}}} \right). \quad (10.21)$$

Bibliography

- Albrecht, T. et al. (July 2015). “Triadic resonances in precessing rapidly rotating cylinder flows”. In: *Journal of Fluid Mechanics* 778, R1.
- (Feb. 2018). “On triadic resonance as an instability mechanism in precessing cylinder flow”. In: *Journal of Fluid Mechanics* 841, R3.
- Armitage, Philip J. (June 1998). “Turbulence and Angular Momentum Transport in a Global Accretion Disk Simulation”. In: *The Astrophysical Journal* 501.2, p. L189.
- (2011). “Dynamics of Protoplanetary Disks”. In: *Annual Review of Astronomy and Astrophysics* 49. Volume 49, 2011, pp. 195–236.
- Armstrong, J. A. et al. (Sept. 1962). “Interactions between Light Waves in a Nonlinear Dielectric”. In: *Phys. Rev.* 127 (6), pp. 1918–1939.
- Barranco, Joseph A. and Philip S. Marcus (Apr. 2005). “Three-dimensional Vortices in Stratified Protoplanetary Disks”. In: *The Astrophysical Journal* 623.2, p. 1157.
- (2006). “A 3D spectral anelastic hydrodynamic code for shearing, stratified flows”. In: *Journal of Computational Physics* 219.1, pp. 21–46.
- Barranco, Joseph A., Suyang Pei, and Philip S. Marcus (Dec. 2018). “Zombie Vortex Instability. III. Persistence with Nonuniform Stratification and Radiative Damping”. In: *The Astrophysical Journal* 869.2, p. 127.
- Bayly, B. J. (Oct. 1986). “Three-Dimensional Instability of Elliptical Flow”. In: *Physical Review Letters* 57.17, pp. 2160–2163.
- Becker, J. M. and R. Grimshaw (Dec. 1993). “Explosive resonant triads in a continuously stratified shear flow”. In: *Journal of Fluid Mechanics* 257, pp. 219–228.
- Benney, D J and R F Bergeron (Sept. 1969). “A New Class of Nonlinear Waves in Parallel Flows”. In: *Studies in Applied Mathematics* 48.3, pp. 181–204.
- Blanco-Rodríguez, F. J. and S. Le Dizès (2017). “Curvature instability of a curved Batchelor vortex”. In: *Journal of Fluid Mechanics* 814, pp. 397–415.
- Booker, John R. and Francis P. Bretherton (1967). “The critical layer for internal gravity waves in a shear flow”. In: *Journal of Fluid Mechanics* 27.3, pp. 513–539.
- Boyd, John P. (1998). “Two Comments on Filtering Artificial Viscosity for Chebyshev and Legendre Spectral and Spectral Element Methods: Preserving Boundary Conditions and Interpretation of the Filter as a Diffusion”. In: *Journal of Computational Physics* 143.1, pp. 283–288.
- Boyd, Robert W (May 2008). *Nonlinear Optics*. Elsevier.

- Cairns, R. A. (1979). “The role of negative energy waves in some instabilities of parallel flows”. In: *Journal of Fluid Mechanics* 92.1, pp. 1–14.
- Canuto, Claudio et al. (Sept. 2007). *Spectral Methods*. Springer Science & Business Media.
- Chandrasekhar, S. (1981). *Hydrodynamic and Hydromagnetic Stability*. Dover Books on Physics Series. Dover Publications.
- Chang, Ching and Stefan G. Llewellyn Smith (2021). “Density and surface tension effects on vortex stability. Part 2. Moore–Saffman–Tsai–Widnall instability”. In: *Journal of Fluid Mechanics* 913, A15.
- Cheng, Hsi (1975). “An invariant for inertial gravity waves in wind shear”. In: *Geophysical Fluid Dynamics* 7.1, pp. 165–173. eprint: <https://doi.org/10.1080/03091927508242619>.
- Craik, A. D. D. (Jan. 1986). *Wave interactions and fluid flows*. Cambridge, UK: Cambridge University Press.
- Craik, A. D. D. and J. A. Adam (1979). “‘Explosive’ resonant wave interactions in a three-layer fluid flow”. In: *Journal of Fluid Mechanics* 92.1, pp. 15–33.
- Crow, S. C. (1970). “Stability theory for a pair of trailing vortices”. In: *AIAA Journal* 8.12, pp. 2172–2179. eprint: <https://doi.org/10.2514/3.6083>.
- Dalcin, Lisandro, Mikael Mortensen, and David E. Keyes (2019). “Fast parallel multidimensional FFT using advanced MPI”. In: *Journal of Parallel and Distributed Computing* 128, pp. 137–150.
- Dixit, Harish N. and Rama Govindarajan (2011). “Stability of a vortex in radial density stratification: role of wave interactions”. In: *Journal of Fluid Mechanics* 679, pp. 582–615.
- Dodin, I.Y., A.I. Zhmoginov, and N.J. Fisch (2008). “Manley–Rowe relations for an arbitrary discrete system”. In: *Physics Letters A* 372.39, pp. 6094–6096.
- Dubrulle, B. et al. (2005). “A hydrodynamic shear instability in stratified disks”. In: *A&A* 429.1, pp. 1–13.
- Ebrahimi, F. and E. G. Blackman (Mar. 2016). “Radially dependent large-scale dynamos in global cylindrical shear flows and the local cartesian limit”. In: *Monthly Notices of the Royal Astronomical Society* 459.2, pp. 1422–1431. eprint: <https://academic.oup.com/mnras/article-pdf/459/2/1422/8040856/stw724.pdf>.
- Eloy, C. and S. Le Dizès (1999). “Three-dimensional instability of Burgers and Lamb–Oseen vortices in a strain field”. In: *Journal of Fluid Mechanics* 378, pp. 145–166.
- (Mar. 2001). “Stability of the Rankine vortex in a multipolar strain field”. In: *Physics of Fluids* 13.3, pp. 660–676.
- Fabre, D., D. Sipp, and L. Jacquin (2006). “Kelvin waves and the singular modes of the Lamb–Oseen vortex”. In: *Journal of Fluid Mechanics* 551, pp. 235–274.
- Feys, J. and S. A. Maslowe (2016). “Elliptical instability of the Moore–Saffman model for a trailing wingtip vortex”. In: *Journal of Fluid Mechanics* 803, pp. 556–590.
- Fukumoto, Y. and Y. Hattori (Mar. 2005). “Curvature instability of a vortex ring”. In: *Journal of Fluid Mechanics* 526, pp. 77–115.
- Gallay, T. and D. Smets (Sept. 2020). “Spectral stability of inviscid columnar vortices”. In: *Analysis & PDE* 13.6, pp. 1777–1832.

- Gammie, Charles F. (Jan. 1996). “Layered Accretion in T Tauri Disks”. In: *Astrophysical Journal* 457, p. 355.
- Hallock, J. N. and F. Holzäpfel (Apr. 2018). “A review of recent wake vortex research for increasing airport capacity”. In: *Progress in Aerospace Sciences* 98, pp. 27–36.
- Hassanzadeh, Pedram, Philip S. Marcus, and Patrice Le Gal (2012). “The universal aspect ratio of vortices in rotating stratified flows: theory and simulation”. In: *Journal of Fluid Mechanics* 706, pp. 46–57.
- Howard, L. N. and A. S. Gupta (1962). “On the hydrodynamic and hydromagnetic stability of swirling flows”. In: *Journal of Fluid Mechanics* 14.3, pp. 463–476.
- Joly, L., J. Fontane, and P. Chassaing (2005). “The Rayleigh–Taylor instability of two-dimensional high-density vortices”. In: *Journal of Fluid Mechanics* 537, pp. 415–431.
- Kalashnik, Maxim V. and Otto Chkhetiani (2017). “Generation of Gravity Waves by Singular Potential Vorticity Disturbances in Shear Flows”. In: *Journal of the Atmospheric Sciences* 74.1, pp. 293–308.
- Kelvin, L. (Sept. 1880). “Vibrations of a columnar vortex”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 10.61, pp. 155–168.
- Kerswell, R. R. (Mar. 1999). “Secondary instabilities in rapidly rotating fluids: inertial wave breakdown”. In: *Journal of Fluid Mechanics* 382, pp. 283–306.
- (Jan. 2002). “Elliptical Instability”. In: *Annual Review of Fluid Mechanics* 34.1, pp. 83–113.
- Kobine, J. Jonathan (1995). “Inertial wave dynamics in a rotating and precessing cylinder”. In: *Journal of Fluid Mechanics* 303, pp. 233–252.
- Lacaze, L., A.-L. Birbaud, and S. Le Dizès (Jan. 2005). “Elliptic instability in a Rankine vortex with axial flow”. In: *Physics of Fluids* 17.1, p. 017101.
- Lacaze, L., K. Ryan, and S. Le Dizès (2007). “Elliptic instability in a strained Batchelor vortex”. In: *Journal of Fluid Mechanics* 577, pp. 341–361.
- Lagrange, R., C. Eloy, et al. (Aug. 2008). “Instability of a fluid inside a precessing cylinder”. In: *Physics of Fluids* 20.8, p. 081701.
- Lagrange, R., P. Meunier, et al. (Jan. 2011). “Precessional instability of a fluid cylinder”. In: *Journal of Fluid Mechanics* 666, pp. 104–145.
- Le Dizès, S. (2008). “Inviscid waves on a Lamb–Oseen vortex in a rotating stratified fluid: consequences for the elliptic instability”. In: *Journal of Fluid Mechanics* 597, pp. 283–303.
- Le Dizès, S. and F. Laporte (2002). “Theoretical predictions for the elliptical instability in a two-vortex flow”. In: *Journal of Fluid Mechanics* 471, pp. 169–201.
- Le Dizès, Stéphane and Paul Billant (Sept. 2009). “Radiative instability in stratified vortices”. In: *Physics of Fluids* 21.9, p. 096602. eprint: https://pubs.aip.org/aip/pof/article-pdf/doi/10.1063/1.3241995/14737870/096602_1_online.pdf.
- Le Dizès, Stéphane and Christophe Eloy (Feb. 1999). “Short-wavelength instability of a vortex in a multipolar strain field”. In: *Physics of Fluids* 11.2, pp. 500–502.
- Le Dizès, Stéphane and Laurent Lacaze (Oct. 2005). “An asymptotic description of vortex Kelvin modes”. In: *Journal of Fluid Mechanics* 542, pp. 69–96.

- Lee, S. and P. S. Marcus (July 2023). “Linear stability analysis of wake vortices by a spectral method using mapped Legendre functions”. In: *Journal of Fluid Mechanics* 967, A2.
- (July 2025). “Transient growth of a wake vortex and its initiation via inertial particles”. In: *Journal of Fluid Mechanics* 1014, A16.
- Lee, Sangjoon and Philip S. Marcus (2023). “Linear stability analysis of wake vortices by a spectral method using mapped Legendre functions”. In: *Journal of Fluid Mechanics* 967, A2.
- Lee, Sangjoon and Jinge Wang (Sept. 2025). *MLegS: Modernized and Parallelized Mapped Legendre Spectral Method Code*. Version v1.0.2.
- Leweke, T. and C. H. K. Williamson (1998). “Cooperative elliptic instability of a vortex pair”. In: *Journal of Fluid Mechanics* 360, pp. 85–119.
- Lin, C.C (1955). *The Theory of Hydrodynamic Stability*. Cambridge University Press.
- Lopez, J. M. and F. Marqués (Jan. 2018). “Rapidly rotating precessing cylinder flows: forced triadic resonances”. In: *Journal of Fluid Mechanics* 839, pp. 239–270.
- Machida, M., Mitsuru R. Hayashi, and R. Matsumoto (Feb. 2000). “Global Simulations of Differentially Rotating Magnetized Disks: Formation of Low- β Filaments and Structured Coronae”. In: *The Astrophysical Journal* 532.1, p. L67.
- Mahalov, A. (Apr. 1993). “The instability of rotating fluid columns subjected to a weak external Coriolis force”. In: *Physics of Fluids A: Fluid Dynamics* 5.4, pp. 891–900.
- Mahdinia, Mani et al. (2017). “Stability of three-dimensional Gaussian vortices in an unbounded, rotating, vertically stratified, Boussinesq flow: linear analysis”. In: *Journal of Fluid Mechanics* 824, pp. 97–134.
- Manasseh, Richard (1992). “Breakdown regimes of inertia waves in a precessing cylinder”. In: *Journal of Fluid Mechanics* 243, pp. 261–296.
- (1994). “Distortions of inertia waves in a rotating fluid cylinder forced near its fundamental mode resonance”. In: *Journal of Fluid Mechanics* 265, pp. 345–370.
- Manley, J. M. and H. E. Rowe (1956). “Some General Properties of Nonlinear Elements-Part I. General Energy Relations”. In: *Proceedings of the IRE* 44.7, pp. 904–913.
- Marcus, Philip S., Suyang Pei, Chung-Hsiang Jiang, and Joseph A. Barranco (Dec. 2016). “ZOMBIE VORTEX INSTABILITY. II. THRESHOLDS TO TRIGGER INSTABILITY AND THE PROPERTIES OF ZOMBIE TURBULENCE IN THE DEAD ZONES OF PROTOPLANETARY DISKS”. In: *The Astrophysical Journal* 833.2, p. 148.
- Marcus, Philip S., Suyang Pei, Chung-Hsiang Jiang, Joseph A. Barranco, et al. (July 2015). “ZOMBIE VORTEX INSTABILITY. I. A PURELY HYDRODYNAMIC INSTABILITY TO RESURRECT THE DEAD ZONES OF PROTOPLANETARY DISKS”. In: *The Astrophysical Journal* 808.1, p. 87.
- Marcus, Philip S., Suyang Pei, Chung-Hsiang Jiang, and Pedram Hassanzadeh (Aug. 2013). “Three-Dimensional Vortices Generated by Self-Replication in Stably Stratified Rotating Shear Flows”. In: *Phys. Rev. Lett.* 111 (8), p. 084501.
- Maslowe, S. (Jan. 1986). “Critical Layers in Shear Flows”. In: *Annual Review of Fluid Mechanics* 18.1, pp. 405–432.

- Maslowe, S. A. (Feb. 2022). “Linear instability of a perturbed Lamb–Oseen vortex”. In: *Fluid Dynamics Research* 54.1, p. 015513.
- Matsushima, T. and P.S. Marcus (1997). “A Spectral Method for Unbounded Domains”. In: *Journal of Computational Physics* 137.2, pp. 321–345.
- Mayer, E. W. and K. G. Powell (Dec. 1992). “Viscous and inviscid instabilities of a trailing vortex”. In: *Journal of Fluid Mechanics* 245, pp. 91–114.
- McEwan, A. D. (Feb. 1970). “Inertial oscillations in a rotating fluid cylinder”. In: *Journal of Fluid Mechanics* 40.3, pp. 603–640.
- Molemaker, M. Jeroen, James C. McWilliams, and Irad Yavneh (June 2001). “Instability and Equilibration of Centrifugally Stable Stratified Taylor-Couette Flow”. In: *Phys. Rev. Lett.* 86 (23), pp. 5270–5273.
- Moore, D. W., P. G. Saffman, and J. T. Stuart (1975). “The instability of a straight vortex filament in a strain field”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 346.1646, pp. 413–425.
- Mora, D. O. et al. (July 2021). “Three-dimensionality of the triadic resonance instability of a plane inertial wave”. In: *Physical Review Fluids* 6 (7), p. 074801.
- Pierrehumbert, R. T. (Oct. 1986). “Universal Short-Wave Instability of Two-Dimensional Eddies in an Inviscid Fluid”. In: *Physical Review Letters* 57.17, pp. 2157–2159.
- Pradeep, D. S. and F. Hussain (2001). “Core dynamics of a strained vortex: instability and transition”. In: *Journal of Fluid Mechanics* 447, pp. 247–285.
- Regev, O. and O. M. Umrhan (Jan. 2008). “On the viability of the shearing box approximation for numerical studies of MHD turbulence in accretion disks”. In: *Astronomy & Astrophysics* 481.1, pp. 21–32.
- Riedinger, Xavier, Stéphane Le Dizès, and Patrice Meunier (2010). “Viscous stability properties of a Lamb–Oseen vortex in a stratified fluid”. In: *Journal of Fluid Mechanics* 645, pp. 255–278.
- Saunders, Peter M (1973). “The instability of a baroclinic vortex”. In: *Journal of Physical Oceanography* 3.1, pp. 61–65.
- Schmid, Peter J., J. E. Marsden, and L. Sirovich (2001). *Stability and transition in shear flows*. Springer New York.
- Sipp, Denis et al. (2005). “Stability of a vortex with a heavy core”. In: *Journal of Fluid Mechanics* 526, pp. 67–76.
- Song, X. Carol et al. (2022). “Anvil - System Architecture and Experiences from Deployment and Early User Operations”. In: *Practice and Experience in Advanced Research Computing 2022: Revolutionary: Computing, Connections, You*. PEARC '22. Boston, MA, USA: Association for Computing Machinery.
- Stewartson, K. (1977). “The evolution of the critical layer of a Rossby wave”. In: *Geophysical & Astrophysical Fluid Dynamics* 9.1, pp. 185–200. eprint: <https://doi.org/10.1080/03091927708242326>.
- Stix, Thomas Howard (1992). *Waves in plasmas*. New York: American Institute Of Physics.

- Tsai, C.-Y. and S. E. Widnall (1976). “The stability of short waves on a straight vortex filament in a weak externally imposed strain field”. In: *Journal of Fluid Mechanics* 73.4, pp. 721–733.
- Tsutahara, M. (Aug. 1984). “Resonant interaction of internal waves in a stratified shear flow”. In: *Physics of Fluids* 27.8, pp. 1942–1947.
- Tsutahara, M. and K. Hashimoto (Sept. 1986). “Three-wave resonant interactions and multiple resonances in two-layer and three-layer flows”. In: *Physics of Fluids* 29.9, pp. 2812–2818.
- Waleffe, F. (Jan. 1990). “On the three-dimensional instability of strained vortices”. In: *Physics of Fluids A: Fluid Dynamics* 2.1, pp. 76–80.
- Wang, Chen and Neil J. Balmforth (2020). “Nonlinear dynamics of forced baroclinic critical layers”. In: *Journal of Fluid Mechanics* 883, A12.
- (2021). “Nonlinear dynamics of forced baroclinic critical layers II”. In: *Journal of Fluid Mechanics* 917, A48.
- Wang, Jinge, Sangjoon Lee, and Philip S. Marcus (2026). “Perturbation analysis of triadic resonance in columnar vortices: selection rules and the roles of external forcing and critical layers”. arXiv: 2402.05287 [physics.flu-dyn].
- Warn, T. and H. Warn (1976). “On the Development of a Rossby Wave Critical Level”. In: *Journal of Atmospheric Sciences* 33.10, pp. 2021–2024.
- (1978). “The Evolution of a Nonlinear Critical Level”. In: *Studies in Applied Mathematics* 59.1, pp. 37–71. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/sapm197859137>.
- Weiland, J. and H. Wilhelmsson (1977). *Coherent Non-linear Interaction of Waves in Plasmas*. International series in natural philosophy 88. Oxford, UK: Pergamon Press.
- Whalen, P., M. Brio, and J.V. Moloney (2015). “Exponential time-differencing with embedded Runge–Kutta adaptive step control”. In: *Journal of Computational Physics* 280, pp. 579–601.
- Wilkening, Jon and Xinyu Zhao (2021). “Spatially Quasi-Periodic Water Waves of Infinite Depth”. In: *Journal of Nonlinear Science* 31.3, p. 52.
- Wohlever, Haley Y., Chung-Hsiang Jiang, and Philip S. Marcus (2025). “Zombie Vortex Instability in rotating, stratified, shear flows. I. Inviscid flows.” Manuscript in preparation.
- Yavneh, Irad, James C. McWilliams, and M. Jeroen Molemaker (2001). “Non-axisymmetric instability of centrifugally stable stratified Taylor-Couette flow”. In: *Journal of Fluid Mechanics* 448, pp. 1–21.
- Zhu, Zhaohuan, Yan-Fei Jiang, and James M Stone (Apr. 2020). “Global 3D radiation magnetohydrodynamic simulations for FU Ori’s accretion disc and observational signatures of magnetic fields”. In: *Monthly Notices of the Royal Astronomical Society* 495.3, pp. 3494–3514. eprint: <https://academic.oup.com/mnras/article-pdf/495/3/3494/33344771/staa952.pdf>.