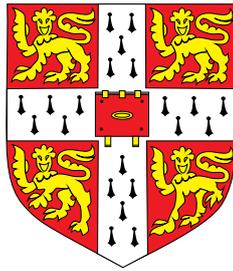


Triadic Resonance Instability in finite-width internal gravity wave beams

Katherine Margaret Grayson
Sidney Sussex College



UNIVERSITY OF CAMBRIDGE

This thesis is submitted for the degree of
Doctor of Philosophy

August 2021

Department of Applied Mathematics and Theoretical Physics,
University of Cambridge.

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the abstract and specified in the text.

Katherine Grayson
November 2021

Abstract

Triadic Resonance Instability in finite-width internal gravity wave beams

Katherine Grayson

Through the use of both experiments and theoretical modelling, this thesis examines the weakly non-linear dynamics of an internal wave beam becoming unstable to Triadic Resonance Instability (TRI). To date, most theoretical work examines the instability in the context of monochromatic plane waves. In the ocean, however, waves seldom take this form, rather they manifest as beams that span the wavenumber spectra. With an aim of developing our understanding of the role of TRI in oceanic settings, this thesis focuses on how the instability evolves in finite-width internal wave beams.

Experiments have been conducted using a new generation of wave maker, featuring a flexible horizontal bottom boundary driven by an array of independently controlled actuators. Using this wavemaker, finite-width internal wave beams of varying amplitude were generated in a linear stratification. Novel experimental results show that when one of these beams becomes unstable via TRI, the approach to a saturated equilibrium state for the triadic waves is not monotonic, rather their amplitudes continue to oscillate without reaching a steady equilibrium. Further diagnostics reveal how the frequencies of the two secondary waves involved in the resonant triad also modulate over time. This behaviour is shown to be a result of the finite spatial extent of the primary beam, which causes cyclic growth and decay of different triadic perturbations. As previously published literature does not consider the triadic energy exchange to be a function of space, it is unable to predict this oscillatory behaviour.

Theoretical modelling is developed to capture the essence of the experimental setup. The model uses numerical solutions of a weakly non-linear system in a two-dimensional framework. A detailed study looks at how different wavenumbers and frequencies of the secondary waves affect the development of TRI in a finite-width beam. The results show how the orientation of the secondary waves has a strong influence on the evolution of the instability, as this determines the duration over which the triadic energy exchange can occur. By including multiple possible resonant waves in the system, the results also capture both the amplitude and frequency oscillations exhibited experimentally. This model not only recapitulates the experimental findings, but provides a tool for the community to dissect the underlying dynamics of the instability.

Both the experimental and theoretical findings presented in this thesis reveal novel insights into how TRI evolves in finite-width internal wave beams. This work thus provides a key to understanding how this instability mechanism may manifest in oceanic scenarios and its potential role in global ocean circulation.

To my parents, Jo and Tina and my sister, Liz.

Acknowledgements

No one achieves anything alone. I don't think I can sum up this PhD with a better phrase and I want to thank everyone who has made it possible. I first want to start by thanking my supervisor, Professor Stuart Dalziel. I could not have been luckier to learn from someone so patient, thoughtful and kind. Stuart has had faith in me from the start, even at times when I didn't and I am eternally indebted for his support and encouragement that has guided me through this PhD. Stuart has gone over and above his role as a supervisor. From spending a weekend building the 'thermal tunnel' with me to delivering jigsaw puzzles in the middle of a pandemic, Stuart has consistently raised the bar as both a supervisor and mentor. He also seems to be an infinite fountain of knowledge that has come to my aid on countless occasions. I cannot begin to thank Stuart enough for getting me to this point, without him this research would not have been possible.

I would also like to thank my co-supervisor, Dr Andrew Lawrie, for the many hours of maths lessons (and chats) over zoom and the wisdom he has provided over these years. His help and advice have been critical in moving my research to new levels. Beyond my supervisor team I am grateful to Neil Stapleton, Duleep Wickramasinghe and Christian Jacobs for helping me to consider my work from a different perspective and the numerous fruitful discussions that have undoubtedly made me a better researcher. Thanks must also go to the Natural Environment Research Council for funding this PhD (grant no. NE/L002507/1).

Next I need to thank the laboratory technicians, without whom none of the experimental work in this thesis would have been possible. Thank you to David Page-Croft for turning my paper scribbles into actual equipment. Thank you to Colin Hitch, the master of the wavemaker, for spending hours lying on the floor underneath my tank fixing the latest leak. Thank you to John Milton for repairing endless broken electrical equipment and answering my urgent pleas for help. Also thank you to Chris Summerfield, Paul Mitton, Andrew Denson, Glen Mitchell and Mark Hallworth who have all provided invaluable assistance for my experiments.

It was not just the laboratory that made my time at DAMTP such a special experience. Firstly, I had the pleasure of having David Parker and Emily Kruger as some of the most fantastic office mates and consequently friends. I was also lucky enough to overlap with Jamie Lee Partridge who, not only made the laboratory such a fun place to work but also taught me the foreign language of DigiFlow. Special thanks must also go to Ben Young, Leo Middleton, Ben Jackson, JP Nijjer, Karol Bacik and the many others in Pav H for the company and community they have all provided during coffee breaks, Christmas dinners and everything in-between.

My life in Cambridge has also been so much more than just my work. I have had the pleasure of making the most wonderful friends who all know how to share the good times and the tough ones. I want to especially thank Mel Jans-Singh, Eléo Fox, Alice Taylor and Rachel Henderson. From the lockdown walks to pub nights and spontaneous adventures to Portugal, I couldn't have asked for better friends. I also want to thank my surf crew, especially Max Jacobs, Barney Salsby and Cristina Jauset Gonzalez who made surfing missions to the North Sea in January some brilliant (and freezing) memories. Throughout my PhD I have also had the pleasure of living with my George Street family, Heather Muir, Joris Witstok and Jeanine Quené. I couldn't have picked better housemates and better people to share lockdown with.

The last few months of writing and finishing experiments have been tough and there is no way I would have made it without my wonderful boyfriend Edoardo Gaude. Not only has Edo borne the brunt of endless moaning and stress, but has cooked every meal for the last four months. Without him I honestly think I would have survived off nothing but a jar of Nutella. Edo has supported me continuously and has played a much bigger role in getting me to the finish line than I think he realises. Finally, to my wonderful family, Mum, Dad, Liz and Granny who have all been there for me since day one and supported me no matter what. In the words of Granny "Keep at it. We are all behind you. Not much use I must say but the thought is there".

Contents

List of Figures	xv
List of Tables	xxxiii
1 Introduction	1
1.1 Internal gravity waves	1
1.2 Thesis outline	4
1.3 Governing equations	5
1.3.1 Introduction	5
1.3.2 Momentum equations	5
1.3.3 Continuity and incompressibility	7
1.3.4 Buoyancy frequency	9
1.3.5 Linear internal waves	10
2 Weakly non-linear theoretical foundations	15
2.1 Introduction	15
2.2 Plane wave solutions	16
2.3 Triadic resonance	17
2.4 Internal wave induced mean flow	18
2.5 Triadic Resonance Instability	19
2.5.1 Overview	19
2.5.2 TRI loci curve	21
2.5.3 Weakly non-linear TRI theory	21
2.5.4 Summary	27
3 Experimental methods and diagnostics	29
3.1 Introduction	29
3.2 Tank	30
3.3 Arbitrary Spectrum Wave Maker	31

3.4	Method of tank filling	34
3.5	Conductivity probe	38
3.6	Vortex ring generation	44
3.7	Synthetic Schlieren	47
3.7.1	Overview	47
3.7.2	Ray tracing	48
3.7.3	Apparatus	51
3.8	Automatic Experiments	52
3.9	Image Filtering	54
3.9.1	Overview	54
3.9.2	Dynamic Mode Decomposition	54
3.9.3	Hilbert Transform	61
3.9.4	Experimental determination of amplitude	62
3.10	Summary	63
4	Experimental Triadic Resonance Instability investigation	65
4.1	Introduction	65
4.2	Experimental wave form	66
4.3	Experimental Results	68
4.3.1	Overview	68
4.3.2	Temporal and spatial modulation	71
4.3.3	Frequency and wavenumber modulation	74
4.4	Amplitude threshold for a finite-width beam	79
4.5	Comparison with previous theory	80
4.5.1	Overview	80
4.5.2	Amplitude development comparison	81
4.5.3	Growth rate comparison	84
4.6	Summary	88
5	A weakly non-linear approach to Triadic Resonance Instability	91
5.1	Introduction	91
5.2	Perturbation expansion	92
5.3	Code development	99
5.3.1	Overview	99
5.3.2	Numerical domains	100
5.3.3	Overall code structure	101
5.4	Linear advection scheme	102

5.4.1	Overview	102
5.4.2	Computational cell	102
5.4.3	Time-step	102
5.4.4	First-order upwind scheme	103
5.4.5	Second-order upwind scheme	104
5.4.6	Initial and boundary conditions	106
5.5	Weakly non-linear interactions	108
5.6	Verification	108
5.6.1	Overview	108
5.6.2	Comparison with analytical solution	109
5.6.3	Energy conservation	111
5.7	Summary	115
6	Weakly non-linear model results	117
6.1	Introduction	117
6.2	Simple triad	118
6.2.1	Overview	118
6.2.2	Outer loci branch	118
6.2.3	Inner loci branch	125
6.3	Multiple secondary waves	129
6.3.1	Overview	129
6.3.2	Temporal and spatial modulation	131
6.3.3	Frequency and wavenumber modulation	133
6.3.4	Growth rate	135
6.4	Higher order expansion	139
6.5	Summary	142
7	Deliberate perturbations	145
7.1	Introduction	145
7.2	Vortex rings in a stratified environment	146
7.2.1	Overview	146
7.2.2	Formation and propagation	146
7.2.3	Internal wave production	147
7.2.4	Collapse distance in a stratified environment	149
7.3	Vortex ring interactions	151
7.3.1	Overview	151
7.3.2	Effect of the wave field on the vortex ring	151

7.3.3	Effect of the vortex ring on the wave field	153
7.4	Potential sub-critical bifurcation	156
7.4.1	Overview	156
7.4.2	Amplitude reduction experiments	157
7.4.3	Secondary wave experiments	162
7.5	Long term vortex ring interactions	167
7.5.1	Overview	167
7.5.2	Long term development of vortex ring interactions	167
7.6	Summary	170
8	Conclusions	173
8.1	Introduction	173
8.2	Thesis summary	174
8.2.1	Finite-width beam	174
8.2.2	Weakly non-linear modelling	175
8.2.3	Deliberate perturbations	176
8.3	Discussion	177
8.3.1	Connection to dynamical systems	177
8.3.2	Three-dimensional considerations	179
8.3.3	Role of beam configuration	181
8.4	Avenues for future experimental and analytical work	181
8.4.1	Wave forcing parameters	182
8.4.2	Weakly non-linear modelling	182
8.4.3	Vortex ring interactions	183
8.5	Further diagnostic developments	184
8.5.1	Experimental control	184
8.5.2	Cameras	185
8.5.3	Particle Image Velocimetry	186
8.6	Concluding remarks	186
	Bibliography	189

List of Figures

1.1	A diagram by MacKinnon (2013). The long purple arrows show the different generation mechanisms of internal waves within the ocean, while the curly arrows indicate the turbulent mixing generated by the breaking of these waves.	3
1.2	Schematic of the oscillating cylinder experiment conducted by Mowbray and Rarity (1967). A cylinder is oscillated vertically at frequency ω in a linear stratification N . The motion produces four wave beams that propagate outwards at angle $\theta = \omega/N$ with respect to the vertical, with group velocity \mathbf{c}_g . The wavenumber vectors \mathbf{k} and phase velocities \mathbf{c}_p are perpendicular to the direction of energy propagation. The four different sign combinations for \mathbf{k} that satisfy linear dispersion (1.20) are indicated for each beam. The quadrants correspond to the quadrants used for the Hilbert Transform discussed in § 3.9.3.	13
2.1	Figure showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector with an origin at (0,0). The black line indicates all the possible locations that satisfy (2.5) corresponding to the TRI condition (2.4). The red dashed lines indicate other possible solutions that are neutrally stable, based on different sign combinations of the two terms on the right hand side of (2.5). The vectors have been normalised by \mathbf{k}_0	22
2.2	A diagram by Bourget et al. (2014), showing their experimental setup in which a primary wave beam is generated from the surface mounted wavemaker. The primary beam has become unstable to TRI and the direction of the resonant waves group velocities is indicated. The control area used to define the interaction region over which energy exchange can occur between the triadic waves is outlined by the black rectangle.	26

3.1	Photo of the main tank in the GK Batchelor Laboratory used for the experiments presented in this thesis, taken from Dobra (2018).	31
3.2	A diagram depicting some of the horizontal rods encased in their Velcro sleeve. These sleeves connect to the underside of the neoprene foam and are able to rotate around rod while still being constrained to the rods vertical motion. This diagram is reproduced from Dobra et al. (2019).	32
3.3	Computer Aided Design drawings of ASWaM. (a) A view from the underside, showing the linear actuators connected to the vertical drive rods via the couplings. The motors are staggered across three vertical planes to enable a tight rod separation of 10 mm. (b) shows a view of ASWaM in-situ with the horizontal support rods exposed. Figures reproduced from Dobra et al. (2019).	33
3.4	Photo of the Arbitrary Spectrum Wavemaker while running a large amplitude sine wave. Due to the large forcing the glycerol located in the cavity underneath has become entrained in the lower region of the stratification causing distortion to the dot pattern located behind the tank. The bearing blocks and motor housing can be seen underneath the tank.	34
3.5	Pictures of the Coleparmer gear pumps used to fill the stratification for the tank. Both pumps were controlled digitally, using the Digiflow software package (Dalziel, 2021).	36
3.6	Gear pump calibration curves. (a) The calibration curve for both pumps across the full range of voltages that the pumps were capable of. (b) High resolution calibration curves for the range of voltages required for filling the stratification.	36
3.7	Timeline of procedure used to fill the stratification in the main tank.	37
3.8	Photograph of the ballcock valve in the fresh water holding tank.	38
3.9	Photograph of the white Delrin tip of the conductivity probe and the outer metal electrode.	39
3.10	Calibration curve for the aspirating conductivity probe. The curve fit was determined using least squares.	40
3.11	The voltage reading from the conductivity probe over 10 seconds. (a) Shows the effect of the electrical interference from the broken traverse motor on the voltage signal of the probe. (b) The probe signal over the same period after the traverse controller was replaced.	41

3.12	The effect of the relationship given in (3.3) on each point in the probe calibration curve given in Figure 3.10 with a calibration temperature C_T of 17.5° C.	42
3.13	Two conductivity profiles taken in June 2019. The black profile was taken soon after the tank was filled, while the blue profile was taken approximately 70 hours later, after which the water temperature had equilibrated to that of the laboratory.	43
3.14	The obtained values for the buoyancy frequency, N , from (3.4) as a function of time. Each graph corresponds to one set of experiments using the same stratification. The green circle of ‘Before’ indicates the profile(s) taken before a single experiment, while the red ‘After’ indicates the profile(s) taken 30 s after the experimental recording finished. The black and blue squares in (c) indicate the two profiles in Figure 3.13.	45
3.15	Schematic showing the view from the camera looking towards the front of the tank. For experiments including the vortex-rings (discussed in Chapter 7), the black ‘L’-shaped tube was mounted as shown before the tank was filled with a linear stratification. Further pictures of the vortex-ring apparatus are shown in Figure 3.16.	46
3.16	Pictures of the vortex ring generator traverse and tube. (a) A picture of the vortex ring generator traverse. Two bike pumps are attached to a linear traverse, mounted on a Bosch profile frame. (b) A photo of the outlet tube which, during experiments, sits partially submerged in the tank. The Bosch profile attachment spans the tank and holds the outlet tube in place when the vortex ring is fired. (c) A connection between the hose running from the bike pumps to the outlet tube shown in (b). The orange bung provides an airtight seal around the connection point ensuring volume conservation.	47
3.17	Side view schematic of the optical arrangement for Synthetic Schlieren. An LED light bank is positioned behind the tank where a random dot pattern is attached. The light rays generated from the light bank refract along a curve through the density stratification, which, for an unperturbed linear stratification, can be modelled by a quadratic. Internal waves generated from the wavemaker then locally distort the stratification and the camera records the apparent displacement of the dots attached to the light bank. A DigiFlow algorithm is then used to determine the gradient of this density perturbation.	48

- 3.18 Pictures of the thermal tunnel. (a) A view from above looking towards the tank, showing the exterior of the tunnel suspended from its frame. (b) A view of the frame around the field of view, which was attached onto the tank. This frame was covered in ‘Velcro’, which connected to the matching ‘Velcro’ on the tunnel, allowing a seal to form around the visualisation window. (c) A view from inside the tunnel where the camera was located, looking towards the tank. 52
- 3.19 Picture of the automatic pinch valve used for the automation of the experiments. When the traverse was moving, the barrel would open, allowing flow through the tube which was connected to the top of the conductivity probe. When the traverse stopped moving, the barrel would pinch close, preventing flow through the tube. 53
- 3.20 Sequence of images showing the difference between the DMD output over frames 300 – 320, for the horizontal and vertical components of the density perturbation. Only the real component of the complex output is shown. The imaginary component of each modal field (not shown) corresponds to the modes complex conjugate. On the left (a) and (c) are the two strongest modes from the vertical gradient of the density perturbation ($\partial\rho/\partial z$), while their corresponding outputs for the horizontal gradient ($\partial\rho/\partial x$) are given in (b) and (d). The frequencies of each image are (a) $\omega = 0.95 \text{ rad s}^{-1}$, (b) $\omega = 0.95 \text{ rad s}^{-1}$, (c) $\omega = 0 \text{ rad s}^{-1}$, (d) $\omega = 1.55 \text{ rad s}^{-1}$. The vertical axis has been normalised by the depth of the tank H , while the horizontal axis has been normalised by the wavelength of the forced wavebeam λ_0 , imposed by the wavemaker. 59
- 3.21 Sequence of images showing the dominant output frequencies from DMD when using the vertical and horizontal components of the gradient of the density perturbation stacked on top of each other. The experiment is the same as that shown in Figure 3.20 and, again, only the real part of the complex output is shown. Stacking the vertical and horizontal components ensures a coupling for the output frequencies. (a) $\omega = 0.95 \text{ rad s}^{-1}$, (b) $\omega = 0 \text{ rad s}^{-1}$ 60

- 3.22 Two images showing the results of the directional filtering from the Hilbert Transform. (a) The RMS amplitude computed from the complex density perturbation field. (b) The RMS amplitude of the density perturbation field produced from the Hilbert Transform, where the wavenumber components in the third quadrant with $\mathbf{k} = (-l, -m)$ are isolated. 62
- 3.23 Two images showing the difference in regions used for averaging the reduced streamfunction after the Hilbert Transform was employed. (a) The domain \mathcal{D}_w is defined by the whole the field of view which is averaged over to obtain $\langle |\tilde{\Psi}_p| \rangle_w$, where $p = 0, 1, 2$ depending on the triadic wave in question. (b) The domain \mathcal{D}_r is shown by the white region, which is averaged over to obtain $\langle |\tilde{\Psi}_0| \rangle_r$. This region is only used for the primary wave. 63
- 4.1 Schematic showing the four wavelength wide waveform generated by ASWaM for the first 60 s of forcing. The blue line shows the actual sinusoidal displacement of the wavemaker while the red line outlines the carrier envelope to the beam. The envelope smooths the wavelengths on either side of the beam with a cosine squared curve while the middle two wavelengths reach maximum height η_0 after 30 s. The x, y and z axes are normalised by the wavelength λ_0 , the period T_0 and the amplitude η_0 of the primary beam, respectively. 68
- 4.2 (a) The vertical gradient of the density perturbation for the full flow field at 549 s ($t/T_0 \approx 83$) into an experiment forced at $\eta_0 = 4$ mm. The white line gives the location for the vertical time series shown in Figure 4.6. (b)–(d) The real part of three of the dominant frequencies produced from the DMD over frames 545 – 565 ($83 \ll t/T_0 \ll 86$). The red arrows overlaid indicate the orientation of the wavenumber vector. Specifically, (b) \mathbb{W}_0 with $\omega_0 = 0.95$ rad s⁻¹ (c) \mathbb{W}_1 with $\omega_1 = 0.358$ rad s⁻¹ (d) \mathbb{W}_2 with $\omega_2 = 0.592$ rad s⁻¹. 70
- 4.3 Figure showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector with an origin at (0,0). The black line indicates all the possible locations that satisfy (2.5) corresponding to the TRI condition (2.4). The blue arrows show the normalised experimentally produced wavenumber vectors for the resonant triad determined from the peak in the one-dimensional power spectrum performed on the snapshots shown in Figure 4.2. 71

- 4.4 Sequence of images at 400 s apart, showing the vertical gradient of the density perturbation for the full flow of the $\eta_0 = 4$ mm experiment. (a) $t/T_0 = 53$, (b) $t/T_0 = 113$, (c) $t/T_0 = 174$, (d) $t/T_0 = 234$, (e) $t/T_0 = 294$, (f) $t/T_0 = 255$, (g) $t/T_0 = 415$, (h) $t/T_0 = 475$ 72
- 4.5 The amplitude of the streamfunction for $\langle |\tilde{\Psi}_0| \rangle_w$ (blue), $\langle |\tilde{\Psi}_1| \rangle_w$ (red), $\langle |\tilde{\Psi}_2| \rangle_w$ (green) for four different experiments. All spatial averaging is performed over the whole domain \mathcal{D}_w (see Figure 3.23 for details) and all amplitudes are normalised against the maximum value of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$. Three of the experiments (a),(b) and (d) were run for 90 minutes, while (c) was run for 180 minutes. The forcing amplitudes of the primary wave beam were (a) $\eta_0 = 3.75$ mm, (b) $\eta_0 = 4$ mm, (c) $\eta_0 = 4$ mm and (d) $\eta_0 = 4.5$ mm. 73
- 4.6 Time frequency spectra computed over the vertical transect shown in white on Figure 4.2(a) for the same four experiments in Figure 4.5. Again (a) $\eta_0 = 3.75$ mm, (b) $\eta_0 = 4$ mm, (c) $\eta_0 = 4$ mm and (d) $\eta_0 = 4.5$ mm. The spectral density is given as the logarithm of the RMS field of both the horizontal and vertical gradient of the density perturbation. The dominant frequencies for each experiment obtained from the DMD are then overlaid in white. 75
- 4.7 (a) Black line showing all the possible locations for the tip of the \mathbf{k}_1 that satisfy (2.5) corresponding to the TRI condition (2.4). The blue arrow show the normalised experimentally produced primary wave vector \mathbf{k}_0 , while the green and red arrow show all of the normalised \mathbf{k}_1 and \mathbf{k}_2 respectively for the full duration of the experiment. Again, the wave vectors correspond to the same experiment shown in Figure 4.5(b) and Figure 4.6(b). (b) The same plot as a function of time. The wave vectors are spaced at intervals of 20 s ($t/T_0 \approx 3$), while the spacing of the loci curve is wider to simply the visualisation. 76

- 4.8 (a) A zoom in of the time frequency spectra for experiment shown in Figure 4.6(b), over the frequencies corresponding to the resonant waves. Again, the dominant modes from the DMD are overlaid in white. (b) A graph showing how well the triadic spatial relationship is satisfied for the same experiment. The height of the line for each time period gives the length between the tip of the \mathbf{k}_1 vector and the start of the \mathbf{k}_2 vector shown in Figure 4.7 normalised by \mathbf{k}_0 . The black dot dashed line shows the time of the snapshots in Figure 4.2, while the dashed line shows the time of the snapshots in Figure 4.9. 78
- 4.9 (a) The vertical gradient of the density perturbation for the full flow field of the same experiment presented in Figure 4.8 at 2109 s ($t/T_0 \approx 319$). (b)–(d) The real part of three of the dominant frequencies produced from the DMD over the frames 2099 – 2119 ($318 \leq t/T_0 \leq 321$). Specifically, (b) $\omega_0 = 0.95 \text{ rad s}^{-1}$ (c) $\omega_1 = 0.323 \text{ rad s}^{-1}$ (d) $\omega_2 = 0.627 \text{ rad s}^{-1}$ 79
- 4.10 Phase diagram showing the amplitude threshold required for the experimental wavebeam to become unstable to TRI. Seven sets of experiments are shown, all with the same \mathbb{W}_0 parameters but with varying forcing amplitudes and in varying strength stratifications. The red and blue circles correspond to instability or no instability respectively. The amplitude has been calculated over a small region $\langle |\tilde{\Psi}_0| \rangle_r$, shown in Figure 3.23, as this most accurately represents the amplitude of the primary wave beam. 80
- 4.11 Amplitude of the primary wave for all 36 experiments given as in terms of both $\langle |\tilde{\Psi}_0| \rangle_w$ and $\langle |\tilde{\Psi}_0| \rangle_r$. As expected, this follows a linear fit given by the black line. The four experiments shown in Figure 4.5 are highlighted for comparison. 81
- 4.12 Results from the \mathcal{M}_{0D} model by Bourget et al. (2014) given in (4.4) for four different forcing amplitudes of the primary wave $|\tilde{\Psi}_{in}|$. The amplitude of the streamfunction for the triadic waves is shown as $|\tilde{\Psi}_0|$ (blue), $|\tilde{\Psi}_1|$ (red), $|\tilde{\Psi}_2|$ (green). The forcing amplitudes of the four plots are (a) $|\tilde{\Psi}_{in}| = 20 \text{ mm}^2 \text{ s}^{-1}$, (b) $|\tilde{\Psi}_{in}| = 30 \text{ mm}^2 \text{ s}^{-1}$, (c) $|\tilde{\Psi}_{in}| = 36 \text{ mm}^2 \text{ s}^{-1}$, (d) $|\tilde{\Psi}_{in}| = 40 \text{ mm}^2 \text{ s}^{-1}$ 83

- 4.13 The same curve as shown in Figure 2.1, showing all the possible unstable solutions for the \mathbf{k}_1 wavenumber vector that satisfy (2.5) corresponding to the TRI condition (2.4), using primary wave input parameters $\omega_0 = 0.95 \text{ rad s}^{-1}$, $(l_0, m_0) = (-0.05, -0.064) \text{ mm}^{-1}$. The curve has now been split into different coloured sections based on the length of l_1/l_0 to correspond to the growth rate curves shown in Figure 4.14. 85
- 4.14 Theoretical growth rates calculated using (4.6) for all the resonant wave parameters that satisfy (2.5), using an input amplitude of $|\tilde{\Psi}_0| = 40 \text{ mm}^2 \text{ s}^{-1}$ and primary wave parameters $\omega_0 = 0.95 \text{ rad s}^{-1}$, $(l_0, m_0) = (-0.05, -0.064) \text{ mm}^{-1}$. The colour and style of each curve corresponds to the different branches indicated in Figure 4.13. (a) Shows the growth rate as a function of κ_1/κ_0 . (b) Shows the growth rate as a function of ω_1/ω_0 . All growth rates have been normalised by the maximum value $\hat{\sigma}_{0D_{\max}}$, where $\hat{\sigma}_{0D}$ corresponds to the growth rates of the solid pink line. 86
- 4.15 Log-linear amplitude plots of the triadic waves for two different experiments with their corresponding σ_{0D} growth rate curves. (a) Shows the same experiment presented in Figure 4.5(b) with $\eta_0 = 4 \text{ mm}$. (c) Shows the same experiment presented in Figure 4.5(d) with $\eta_0 = 4.5 \text{ mm}$. The red dashed line in (a) and (c) shows the linear fit from which the experimental value of the growth rate σ_{exp} was extracted, between the two dashed black vertical lines. (b) and (d) show the corresponding σ_{0D} theoretical growth rates of all the possible wavenumbers given from Figure 4.13 obtained from (4.6). The red star gives the extracted growth rate σ_{exp} from the slope of the linear fit on the left hand log-linear plots. The two grey shaded region indicates the range that κ_1/κ_0 took during the respective experiments on the left. The amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_r$ for (b) is $40 \text{ mm}^2 \text{ s}^{-1}$ and for (d) $\langle |\tilde{\Psi}_0| \rangle_r$ is $43 \text{ mm}^2 \text{ s}^{-1}$ determined by the experimental values shown in Figure 4.11. 87
- 5.1 Nomenclature for the computational cell where $\tilde{\Psi}$ is the stream function and F is the advective flux of the stream function across the face of the control volume. The grid dimensions of the cell are given by Δx and Δz . Compass indexing has been used to avoid confusion with numbering systems. 103
- 5.2 Log-log plot showing the RMS error as a function of grid resolution. The red crosses show the calculated values for six different grid resolutions, while the black line shows a linear best fit. 110

- 5.3 (a) Locations where the flux limiters are acting in the flow to keep the scheme monotonic. The average of the horizontal and vertical λ is given, multiplied by amplitude of the wave beam. (b) The RMS error between the exact solution and the solution produced by the model calculated by (5.47) over the domain with the same resolution. The error occurs across the flanks of the wave where the scheme will not necessarily be second-order due to the flux limiters. Both images are for a grid resolution of $n_x, n_z = 82, 242$ 111
- 5.4 Calculation of ε over the whole domain of size $n_x, n_z = 82, 242$ as a function of time. The blue line gives the change in energy over the domain, while the red line shows the divergence of energy. At steady state, around $t/T_0 \approx 12$, the change of energy is equal to zero, however the divergence is non-zero, indicating a small numerical dissipation. . . 113
- 5.5 The energy dissipation error ε over each cell in a 82×242 domain after the wave has reached steady state. The error across each cell is calculated by (5.50). 113
- 5.6 Log-log plot showing the energy dissipation error ε across the whole domain as a function of grid resolution. The red crosses show the calculated values for six different grid resolutions, while the black line shows a linear best fit. 114
- 6.1 Loci curve showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector based off (2.5) and the input parameters for \mathbb{W}_0 given in (6.1). The approximate locations of the experimental \mathbf{k}_1 vector tips are given by the shaded grey region. The six blue marks correspond to each \mathbb{S}_s used for each simple triad model run, the values of which are shown in Table 6.1. 120
- 6.2 Linear theoretical growth rates σ_{0D} calculated using (4.6) for every possible triad configuration shown in Figure 6.1 with an input amplitude of $|\tilde{\Psi}_{in}| = 35 \text{ mm}^2 \text{ s}^{-1}$. The approximate locations of the experimental parameter range κ_1/κ_0 and ω_1/ω_0 are given by the shaded grey regions. The six blue marks correspond to the wave pairs \mathbb{S}_s used as inputs for each ‘simple triad’ in the \mathcal{M}_{2D} (the values of which are shown in Table 6.1) along with their corresponding $\hat{\sigma}_{0D}$ 121

- 6.3 Six amplitude plots generated using the two-dimensional weakly non-linear model. Each plot corresponds to a simple triad with three domains ($n_d = 3$). The parameters for \mathbb{S}_s corresponding to each plot are given in Table 6.1. (a) \mathbb{S}_a (b) \mathbb{S}_b (c) \mathbb{S}_c (d) \mathbb{S}_d (e) \mathbb{S}_e (f) \mathbb{S}_f . The forcing amplitude for \mathbb{W}_0 in all plots is $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$ and $t_{\text{end}} = 90$ minutes ($t/T_0 \approx 800$). The amplitudes are calculated over the whole visualisation window \mathcal{D}_w and have been normalised against the maximum amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$ 122
- 6.4 Sequence of snapshots at roughly 400 s ($t/T_0 \approx 61$) apart, showing the superposition of the three domains in the model. The parameters for the secondary waves correspond to \mathbb{S}_d . Each domain is multiplied by its respective fast timescales $e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)}$. The forcing amplitude for \mathbb{W}_0 is $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$. The amplitude of each wave is shown in Figure 6.3(d). The timing of each snapshot is (a) $t/T_0 = 121$, (b) $t/T_0 = 151$, (c) $t/T_0 = 272$, (d) $t/T_0 = 333$, (e) $t/T_0 = 424$, (f) $t/T_0 = 484$, (g) $t/T_0 = 545$, (h) $t/T_0 = 606$ 123
- 6.5 Amplitude threshold for instability in the \mathcal{M}_{2D} model as a function of input parameters for \mathbb{W}_1 using the simple triads. The blue crosses represent model runs that do not become unstable while the red crosses indicate the growth of TRI. (a) Amplitude threshold as a function of κ_1/κ_0 (b) Amplitude threshold as a function of ω_1/ω_0 . Again, parameters for \mathbb{W}_0 are given in (6.1). 125
- 6.6 Loci curve showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector based off (2.5) and the input parameters for \mathbb{W}_0 given in (6.1). The approximate locations of the experimental \mathbf{k}_1 vector tips are given by the shaded grey region. The blue mark on the inner branch corresponds to $\mathbb{S}_{\text{inner}}$ used for the \mathcal{M}_{2D} model run. The orientation of the three wave vectors corresponding to $\mathbb{S}_{\text{inner}}$ have also been marked. The wavenumbers and frequencies for $\mathbb{S}_{\text{inner}}$ are given in Table 6.2. . . . 127

- 6.7 Two different simple triad model runs using the \mathcal{M}_{2D} model, both with the same triad configuration corresponding $\mathbb{S}_{\text{inner}}$. Each model run is decomposed to show the individual fields for each wave in the triad. (a) and (b) $|\tilde{\Psi}_0|e^{i\phi_0}$ (c) and (d) $|\tilde{\Psi}_1|e^{i\phi_1}$ (e) and (f) $|\tilde{\Psi}_2|e^{i\phi_2}$. The first model run in (a), (c) and (e) is for a primary beam width of $\Lambda_0 = 4\lambda_0$, while the second model run shown in (b), (d) and (f) is for a beam width of $\Lambda_0 = 6\lambda_0$. The width of the primary beam is overlaid in white for the secondary wave domains. The red lines in (c) and (d) indicate the horizontal transect shown in Figure 6.8. 128
- 6.8 Logarithm of $|\tilde{\Psi}_1|$ across the horizontal cross section marked in red in Figure 6.7(c) and (d). The green line is the cross section for the $\Lambda_0 = 4\lambda_0$ model run shown in (c) and the purple line is the cross section for the $\Lambda_0 = 6\lambda_0$ model run shown in (d). The red lines indicate the boundary of the primary wave beam \mathbb{W}_0 129
- 6.9 Theoretical zero-dimensional growth rates calculated using (4.6) for every possible triad configuration shown in Figure 6.1 using an input amplitude of $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$. The blue marks correspond to the top 60% of both $\hat{\sigma}_{0D}$ and $\hat{\sigma}_{0D}$ and indicate all the \mathbb{W}_1 parameters used for the multiple triad \mathcal{M}_{2D} model run. 130
- 6.10 Sequence of snapshots at approximately 350 s ($t/T_0 \approx 60$) apart, showing the superposition of all of the domains in the model. The input parameters of \mathbb{W}_1 for each domain can be seen by the blue circles in 6.9. Each domain is multiplied by its respective fast timescale $e^{i(\mathbf{k}_{p_j} \cdot \mathbf{x} - \omega_{p_j} t)}$. The forcing amplitude for \mathbb{W}_0 is $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$. The timing of each snapshot is (a) $t/T_0 = 151$, (b) $t/T_0 = 212$, (c) $t/T_0 = 258$, (d) $t/T_0 = 318$, (e) $t/T_0 = 394$, (f) $t/T_0 = 455$, (g) $t/T_0 = 500$, (h) $t/T_0 = 561$. The vertical white line in (a) indicates the location over which the time-series for the spectrogram shown in Figure 6.12 was taken. 131
- 6.11 The magnitude of the reduced streamfunction for the multiple triad \mathcal{M}_{2D} model run outlined in § 6.3.1, with $\langle |\tilde{\Psi}_0| \rangle_w$ (blue) $\sum_{j=1}^g \langle |\tilde{\Psi}_{1_j}| \rangle_w$ (red) and $\sum_{j=1}^g \langle |\tilde{\Psi}_{2_j}| \rangle_w$ (green). The amplitudes were calculated over the whole visualisation window \mathcal{D}_w and have been normalised against the maximum amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$. The model was run for 90 minutes with a forcing amplitude of $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$ 133

- 6.12 Time frequency spectra computed over the vertical transect shown in white on Figure 6.10(a) for the multiple triad \mathcal{M}_{2D} model run outlined in § 6.3.1, with an input amplitude of $|\tilde{\Psi}_{in}| = 28 \text{ mm s}^{-1}$. For this run the amplitude of the secondary waves was suppressed between $800 < t < 1100 \text{ s}$ ($121 < t/T_0 < 167$) so the first 990 s ($t/T_0 = 150$) are removed for clarity. The spectral density is given as the logarithm of the RMS field of both the horizontal and vertical gradient of the reduced streamfunction. The dominant frequencies for each experiment obtained from the DMD algorithm are overlaid in white. 134
- 6.13 Four plots showing the triad with the maximum energy in the system over the first 3300 s ($t/T_0 = 500$) of the model run. The energy of each domain is computed by $E = \kappa^2 \langle |\tilde{\Psi}|_w^2 \rangle$ and the maximum energy is marked with the black dotted line. The two left hand plots (a) and (c), show the secondary waves in wavenumber space, while (b) and (d) are given in terms of frequency. For this run, which included 60 secondary pairs, all of the possible input wavenumbers are shown, marked in blue if they correspond to the outer loci curve, or pink if they correspond to the inner curve. The normalisation given in (6.4) occurs between the red dashed lines. Before the normalisation, the \mathbb{T} with the maximum energy is marked in grey as, for this period, the waves are developing from random background noise. 137

- 6.14 (a) A log-linear plot of the magnitude of the reduced streamfunction for the multiple triad \mathcal{M}_{2D} model run outlined in § 6.3.1 with an initial forcing amplitude of $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm s}^{-1}$. The amplitude of the secondary waves are given as the sum across the domains $\sum_{j=1}^g \langle |\tilde{\Psi}_{1_j}| \rangle_w$ and $\sum_{j=1}^g \langle |\tilde{\Psi}_{2_j}| \rangle_w$ and all the amplitudes are normalised against the maximum amplitude of the primary wave. The amplitudes of the secondary waves have been re-scaled using (6.4) between $(121 < t/T_0 < 167)$, indicated by the black dashed lines. The somewhat hidden red dashed line shows the linear fit between the black dash-dot lines to obtain the \mathcal{M}_{2D} linear growth rate σ_{2D} . The light blue line (lying almost on top of the red line) is the linear fit calculated from the summation of $\hat{\sigma}_{0D}$ given in (6.5). (b) Theoretical growth rates calculated using (4.6) for every possible triad configuration shown in Figure 6.1 using an input amplitude of $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$. The blue circles correspond to the top 60% of both $\hat{\sigma}_{0D}$ and $\hat{\sigma}_{\Sigma_{0D}}$ and indicate all the \mathbb{W}_1 parameters used for the multiple triad \mathcal{M}_{2D} run. The red cross gives the value of the \mathcal{M}_{2D} linear growth rate σ_{2D} obtained from the red-dashed line in (a), while the light blue cross gives the linear growth rate $\hat{\sigma}_{\Sigma_{0D}}$ extracted from the light blue line. 138
- 6.15 Four amplitude plots from the \mathcal{M}_{2D} model including the non-linear interactions at $\mathcal{O}(\epsilon^2 \gamma^1)$. Each plot corresponds to a simple triad with three domains ($n_d = 3$). The parameters for the secondary waves corresponding to each plot are given in Table 6.3. (a) \mathbb{S}_a (b) \mathbb{S}_b (c) \mathbb{S}_c (d) \mathbb{S}_d . The forcing amplitude for \mathbb{W}_0 in all plots is $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$ and the run time is 90 minutes ($t/T_0 \approx 800$). The amplitudes are calculated over the whole visualisation window $\langle |\tilde{\Psi}| \rangle_w$ and have been normalised against the maximum amplitude of $\langle |\tilde{\Psi}_0| \rangle_w$ 140
- 6.16 Snapshots at $t/T_0 = 795$ of the simple triad model run shown in Figure 6.15(c) using the input parameters \mathbb{S}_c . The left hand panels show the magnitude of the reduced streamfunction $|\tilde{\Psi}|$, while the corresponding phase angles θ are given in the right panels. The the fast time scales $e^{i\phi}$ are not included. Each wave in the triad is shown separately. 141

- 7.1 Sequence of images taken 0.25 s apart showing the plan view of the generation of a vortex ring in a linear stratification with $N = 1.85 \pm 0.06 \text{ rad s}^{-1}$ and $Fr \approx 1.4$. The images are visualised using florescence dye and a laser sheet. The outlet diameter of the tube is 40 mm. (a), (b) and (c) show the initial formation of the ring as it exits the tube with the inner vortex core visible. (d), (e) and (f) show the three-dimensional structure begin to collapse as the stratification flattens the ring. (g), (h) and (i) then show a transition to turbulence as vorticity is shed in the wake of the dipole structure. 148
- 7.2 Sequence of images of a vortex ring propagating in a linear stratification of $N = 1.85 \pm 0.06 \text{ rad s}^{-1}$. The vertical gradient of the density perturbation is shown. The outline of the vortex ring is shown in black. (a) $t = 12 \text{ s}$, (b) $t = 30 \text{ s}$, (c) $t = 45 \text{ s}$, (d) $t = 60 \text{ s}$ 149
- 7.3 The non-dimensional collapse distance of the vortex rings in a linear stratification of $N = 1.43 \pm 0.06 \text{ rad s}^{-1}$ as a function of their initial Fr . Experimental data points are marked in red with a linear line of best fit shown in black. All rings used a $F_p = 4.4$, while rate of the piston ranged between 650 and 750 mm s^{-1} 150
- 7.4 The collapse distances of vortex rings in a linear stratification ($N = 1.43 \pm 0.06 \text{ rad s}^{-1}$) as a function of their initial Fr . The red crosses (also shown on Figure 7.3) are for rings in a quiescent environment while the blue crosses show the collapse distance of rings who have interacted with an internal wave field. A linear fit is shown in black for both data sets. Again, all rings used a $F_p = 4.4$, while the rate of the piston ranged between 600 and 750 mm s^{-1} for the blue data set. The grey patch indicates the width of the primary beam. The left flank of the beam is approximately 440 mm away from the outlet tube ($x/D_p = 11$), while the right flank is approximately 940 mm away ($x/D_p = 23.5$). The marker style of each blue data point gives the forcing amplitude η_0 of the wave beam. 152

- 7.5 Sequence of images from a vortex ring interaction experiment showing the vertical gradient of the density perturbation. The forcing amplitude of the primary beam is $\eta_0 = 3.6$ mm, while the vortex ring has a $F_p = 4.4$ and a piston rate of 670 mm s⁻¹. These forcing parameters correspond to an initial velocity of 54.6 mm s⁻¹ ($Fr = 1.54$) and a collapse distance of 655 mm from the outlet ($d/D_p = 16.4$). The quantitative amplitudes of the three triadic-waves are shown in Figure 7.6(b). (a) $t/T_0 = 50.8$, (b) $t/T_0 = 52.2$, (c) $t/T_0 = 57.5$, (d) $t/T_0 = 68.0$, (e) $t/T_0 = 98.3$, (f) $t/T_0 = 134.5$, (g) $t/T_0 = 143.6$, (h) $t/T_0 = 151.2$ 154
- 7.6 The amplitude of the streamfunction for $\langle |\tilde{\Psi}_0| \rangle_w$ (blue), $\langle |\tilde{\Psi}_1| \rangle_w$ (red), $\langle |\tilde{\Psi}_2| \rangle_w$ (green) for four different experiments. All amplitudes were integrated over the whole visualisation window \mathcal{D}_W (see Figure 3.23 for details). As it is necessary to compare the amplitudes between experiments, they have not been normalised. In three of the experiments (a), (b) and (c) a vortex ring of $F_p = 4.4$ was fired into the wave field at $t = 330$ s ($t/T_0 = 50$), indicated by the black dashed line. The forcing amplitudes of the primary wave beam are (a) $\eta_0 = 3.5$ mm, (b) $\eta_0 = 3.6$ mm, (c) $\eta_0 = 3.6$ mm and (d) $\eta_0 = 3.7$ mm. 155
- 7.7 Schematic showing (a) supercritical stability states of the TRI system and (b) showing the sub-critical bifurcation stability states of the primary wave beam for TRI. The solid red and black lines indicate the linearly state equilibrium states while the dashed lines represent the linearly-unstable equilibrium states. The schematic in (b) shows a region whereby the waves could correspond to either the linearly-stable single-wave state or linearly-stable triadic-wave state. 157
- 7.8 The amplitude of the reduced streamfunction $\langle |\tilde{\Psi}| \rangle_w$ for four different experiments. Each experiment had the forcing amplitude reduced after 3630 s ($t/T_0 = 560$), indicated by the black dashed line. The grey patch indicates the time over which the primary wave is averaged for the results in Table 7.1. The forcing amplitudes of each experiment were (a) $\eta_0 = 4$ mm, $\eta_1 = 3.5$ mm. (b) $\eta_0 = 4.5$ mm, $\eta_1 = 3.4$ mm. (c) $\eta_0 = 4.5$ mm, $\eta_1 = 3.3$ mm. (d) $\eta_0 = 4.5$ mm, $\eta_1 = 3.4$ mm. 160

- 7.9 Waveform generated by ASWaM after $t = 1200$ s $t/T_0 = 181$ as the trigger wave is initiated. At this point, the primary beam comprising of four wavelengths is already fully formed. The secondary wave is increased from rest over $1200 < t < 1230$ ($181 < t/T_0 < 186$) and can be seen just to the left of the primary beam. The envelope of both beams is shown in red. The x, y and z axis are normalised by the wavelength λ_0 , period T_0 and amplitude η_0 of the primary beam respectively. . . . 163
- 7.10 Sequence of images from a trigger wave experiment showing the vertical gradient of the density perturbation. The forcing amplitude of the primary beam is $\eta_0 = 3.75$ mm, while the forcing amplitude of the trigger wave is $\eta_2 = 2.5$ mm. The quantitative amplitudes of the waves are shown in Figure 7.11(c). (a) $t/T_0 = 180$, (b) $t/T_0 = 197$, (c) $t/T_0 = 212$, (d) $t/T_0 = 227$, (e) $t/T_0 = 270$, (f) $t/T_0 = 295$, (g) $t/T_0 = 303$, (h) $t/T_0 = 263$ 164
- 7.11 The amplitude of the reduced streamfunction $\langle |\tilde{\Psi}| \rangle_w$ for four different experiments. (a) and (b) both used the constant amplitude forcing described by (4.2), while (c) and (d) both show the results from the trigger wave experiments described by (7.4). The grey patch indicates the time over which the primary wave is averaged for the results in Table 7.2. The forcing amplitudes of each experiment are (a) $\eta_0 = 3.75$ mm. (b) $\eta_0 = 3.75$ mm. (c) $\eta_0 = 3.75$ mm, $\eta_2 = 2.5$ mm. (d) $\eta_0 = 3$ mm, $\eta_2 = 2.5$ mm. 166
- 7.12 The amplitude of the streamfunction for $\langle |\tilde{\Psi}_0| \rangle_w$ (blue), $\langle |\tilde{\Psi}_1| \rangle_w$ (red), $\langle |\tilde{\Psi}_2| \rangle_w$ (green) for four different experiments. All amplitudes were calculated over the whole visualisation window $\langle |\tilde{\Psi}| \rangle_w$ (see Figure 3.23 for details). In three of the experiments (a), (b) and (c) a vortex ring of $F_p = 4.4$ was fired into the wave field at $t = 330$ s ($t/T_0 = 50$), indicated by the black dashed line. The grey patch indicates the time over which the primary wave is averaged for the results in Table 7.3. The forcing amplitudes of the primary wave beam are (a) $\eta_0 = 3.5$ mm, (b) $\eta_0 = 3.6$ mm, (c) $\eta_0 = 3.6$ mm and (d) $\eta_0 = 3.7$ mm. 169

-
- 8.1 Schematic showing the connection between the TRI mechanism in a finite-width beam and a dynamical system. The planes, A, B and C all correspond to different states of the system. In A, the system is only attracted to the linearly-stable single-wave branch. As the amplitude increases, B represents the phase space where the system is can be attracted to either the linearly-stable single-wave branch or to a linearly-stable triadic-wave state, represented by B_1 and B_3 respectively. The state B_2 indicates the unstable equilibrium where the system is repelled from. As the amplitude increases further, plane C corresponds to the solution being continuously attracted to linearly-unstable equilibrium triadic states. Here the solution is attracted towards the saddle point via the stable branch of the manifold, only to be thrown away via the unstable branch. 178
- 8.2 Two photos at 112 s ($t/T_0 = 16.9$) apart of a dye experiment conducted using the same forcing parameters for W_1 outlined in § 4.2. The waveform can be seen on the neoprene cover of ASWaM. The dot pattern for the Synthetic Schlieren has been removed and dye crystals are added to the tank from the free surface. The same quasi-vertical streaks seen in (a) in the middle of the wavebeam have moved considerably by (b) in the direction of the group velocity of the primary wave. 180

List of Tables

5.1	Table showing the RMS error between the model and exact solution calculated from (5.47) as a function of grid resolution.	110
5.2	Table showing the energy dissipation error calculated using (5.50) as a function of grid resolution, where now ε is calculated across the whole domain.	114
6.1	Table giving the input parameters of \mathbb{W}_1 and \mathbb{W}_2 for each \mathbb{S}_s used in the simple triad model. The wave vector locations of each pair can be seen by the blue marks on Figure 6.1.	119
6.2	Table giving the input parameters of \mathbb{W}_1 and \mathbb{W}_2 for $\mathbb{S}_{\text{inner}}$ used in the simple triad \mathcal{M}_{2D} model.	126
6.3	Table giving the input parameters of \mathbb{W}_1 and \mathbb{W}_2 for each \mathbb{S} used in the simple triad \mathcal{M}_{2D} model, this time including the non-linear interaction terms at $\mathcal{O}(\epsilon^2\gamma^1)$	139
7.1	Table showing the results of 18 experiments with comparable amplitudes. The average amplitude of the primary wave $\langle \tilde{\Psi}_0 \rangle_w$ is calculated over a 900 s period. In the case of constant amplitude forcing, this period is either just before the primary beam becomes unstable, or, in the case of no instability, is the last 900 s of the experiment. In the case of variable amplitude forcing, this period is the first 900 s after the amplitude is reduced, as shown by the grey shaded regions in Figure 7.8. The second column indicates if TRI was sustained across the duration of the experiment, while the final column indicates if the amplitude was changed, as shown in (7.2). Figures 7.8(a), (b), (c) and (d) correspond to experiments 1, 12, 14 and 8 respectively. The two experiments highlighted in green are the only two that fit the hypothesis of a hysteresis mechanism to the TRI.	161

- 7.2 Table showing the results of 11 experiments with comparable amplitudes. The average amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$ is calculated over a 900 s period. In the case of constant amplitude forcing, this period is either just before the primary beam becomes unstable, or, in the case of no instability, is the last 900 s of the experiment. In the case of forcing using the trigger wave, this period is the first 900 s after the trigger wave has been removed. All averaging windows are shown by the grey shaded regions in Figure 7.11. The second column indicates if TRI was sustained across the duration of the experiment, while the final column indicates if the trigger wave was fired as given in (7.5). Figures 7.11(a), (b), (c) and (d) correspond to experiments 1, 4, 5 and 11 respectively. . 165
- 7.3 Table showing the results of 17 experiments with comparable amplitudes. The average amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$ is calculated over a 900 s ($t/T_0 = 136$) period. In the case of the simple experiments with no ring, this period is either just before the primary beam becomes unstable, or, in the case of no instability, is the last 900 s of the experiment. In the case of forcing with the vortex ring this period is between $1000 < t < 1900$ s ($150 < t/T_0 < 288$). All averaging windows are shown by the grey shaded regions in Figure 7.12. The second column indicates if TRI was sustained across the duration of the experiment, while the final column indicates if a vortex ring was fired. Figures 7.12(b) and (d) correspond to experiments 17 and 13 respectively, while the amplitude of Figures 7.12(a) and (c) are too small to be listed on this table. 170

Chapter 1

Introduction

1.1 Internal gravity waves

One of the first documented accounts of the effects of internal gravity waves was in the late 19th century by the explorer, Fridtjof Nansen, who on an expedition to the Arctic Ocean, experienced an unprecedented slowing of his ship (Mill, 1897). The vessel had encountered a region of water comprised of two distinct density layers, generated by fresh glacial run off sitting above dense salty water without mixing. As this density interface was approximately at the same depth to the draft of the vessel, the perturbations caused by the ship's motion generated waves between the two layers. The slowing phenomena, named at the time as 'dead water', is now understood to be caused by the drag generated from these interfacial waves (Briscoe, 1975).

Interfacial waves describe a particular type of internal gravity wave, where the motion is confined to the horizontal interface between discrete layers of different density fluid. More generally, internal gravity waves are propagating buoyancy-driven oscillations that move at an angle within a fluid whose density changes continuously with depth. A parcel of fluid perturbed upwards away from its equilibrium position will become heavier than its surroundings, whereby gravity acts as a restoring force, causing it to sink. Inertia will, however, cause the parcel of fluid to overshoot the equilibrium position: at equilibrium, there is no buoyancy force to stop the motion. The parcel of fluid is now lighter than its surroundings, causing it to accelerate upwards again and repeat the motion until damped. The result of this oscillatory motion is internal gravity wave propagation.

A fluid whose density continuously increases with depth is said to be stably stratified. In the ocean, density variations are due to both temperature and salinity, while in the atmosphere they are due to temperature, pressure and moisture content. While

internal waves propagate in any stratified fluid medium, their motions in the atmosphere, due to compressibility, are significantly different to their movement within the ocean (Sutherland, 2010).

Since the reports by Nansen, much work has been conducted on internal gravity waves, often referred to simply as internal waves. Ubiquitous in the ocean, internal waves have been found to play an integral role in the dynamics of oceanic flows, due to their ability to transport energy and momentum and to generate mixing where they break (Wunsch and Ferrari, 2004; Nault and Sutherland, 2007). In the ocean, internal waves arise from four main mechanisms. The first, and arguably the most predominant, is the barotropic tidal currents oscillating over bottom topography on the sea floor, such as ridges, seamounts and canyons (Balmforth et al., 2002; Garrett, 2003). The waves generated from this semi-diurnal flow are often referred to as the internal tide (Bühler and Holmes-Cerfon, 2011). The second mechanism arises from strong tidal currents, such as the Antarctic Circumpolar Current, which flow steadily over a sill or obstacle and cause the generation of ‘lee waves’. Here, the waves propagate in the reference frame of the fluid and appear stationary to the observer. Nikurashin and Ferrari (2013) argue that lee-wave driven mixing is not well represented in current climate models and could account for up to one third of the water-mass transformation driven by internal wave-driven mixing. The third main source of these waves, more prominent at mid-latitudes, is from wind induced surface depressions, which generate waves with frequency near the local inertial frequency (Alford et al., 2007). These waves – with frequencies on the order of days – also feel the effects of Coriolis forces due to the earth’s rotation. If they occur at the surface they are known as ‘inertial waves’ (similar to interfacial waves), whereas if they occur within a stratified fluid they are called ‘inertial gravity waves’. While the dynamics of inertial gravity waves are similar to that of internal gravity waves, in this research only the effects of buoyancy will be considered for wave generation and motion. Finally, mixing within the ocean interior, such as the collapse of turbulent events, or the movements of objects within the ocean, has the potential to generate internal waves. Figure 1.1, reproduced from MacKinnon (2013), depicts the main generation mechanisms of these waves and their role in ocean mixing.

While the frequency of oceanic internal waves is much lower than surface waves, their amplitudes can reach tens to hundreds of meters. Indeed, vertical displacements of up to 500 m with horizontal velocities of 1.5 m s^{-1} , have been reported at locations in the South China Sea (Sarkar and Scotti, 2017). These displacements not only complicate the mapping of the average state of the ocean but also have significant

effects on acoustic transmission, sediment transport and off-shore structures (Garrett and Kunze, 2007).

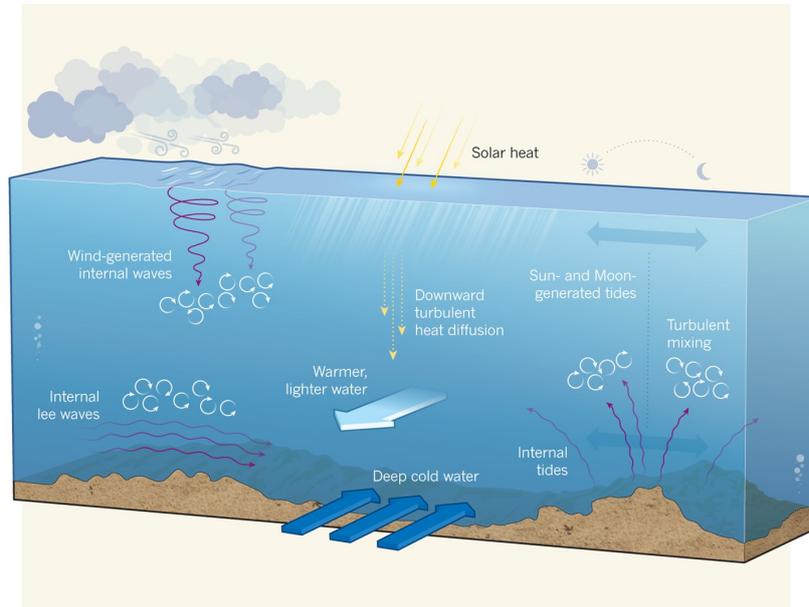


Figure 1.1: A diagram by MacKinnon (2013). The long purple arrows show the different generation mechanisms of internal waves within the ocean, while the curly arrows indicate the turbulent mixing generated by the breaking of these waves.

Despite the generation mechanisms behind internal gravity waves becoming increasingly well explained, their contribution to ocean mixing, and their subsequent role in the maintenance of ocean circulation, is not so well understood. Munk (1966) was amongst the first to suggest that internal waves play a significant part in the deep water vertical mixing of the density stratification within the open ocean. Subsequent work has now shown that the energy flux associated with these waves is not trivial, they are believed to contribute to approximately $0.6 \sim 0.9$ TW of abyssal mixing over the global oceans (Munk and Wunsch, 1998; Egbert and Ray, 2001). This turbulent mixing is crucial in the downward transportation of heat and green house gases, the upward transport of nutrients and supplying the deep ocean with energy to support the global ocean currents (MacKinnon, 2013). Without the presence of deep water mixing, the abyssal ocean would eventually turn into a stagnant, homogeneous pool, with only a sharp interface below the mixed layer (Munk and Wunsch, 1998).

It is, therefore, important to account for the role of internal waves in global climate models, yet it remains impossible to resolve the localised mixing events they generate. While the length scales of internal waves are on the order of 10^1 m to 10^3 m, the length scales involved in their turbulent energy redistribution are on the order of 10^0 m to

10^1 m. Global climate models are unable to resolve these scales and, as a result, the role of internal wave induced mixing is parametrised (Wunsch and Ferrari, 2004). This is often done through the introduction of an eddy diffusivity term that enhances the vertical diffusivity. However, these parametrisations typically fail to account for the spatio-temporal variability inherent in the mixing generated by internal waves (Polzin et al., 1997). A detailed understanding of the mechanisms that transfer internal wave energy to smaller scales is therefore required to capture their variability and effectively include them in global climate models.

One of these energy transfer mechanisms inherent to internal waves is Triadic Resonance Instability (TRI). This weakly non-linear instability mechanism describes a primary wave becoming unstable due to the growth of infinitesimal perturbations in the flow. The growth of these perturbations leads to an energy transfer from the primary wave to two secondary waves, which, together with the primary wave form a resonant triad. Specific details of TRI are given in § 2.5. According to MacKinnon et al. (2017), TRI could play a significant role in internal wave energy dissipation yet our current understanding of the instability is not sufficient to confirm this. As such, unravelling the role of TRI in ocean mixing, especially in the context of more realistic finite-width beams, is of ongoing scientific interest and debate. This thesis therefore seeks to develop understanding of this breakdown mechanism in finite-width internal wave beams through both experiments and weakly non-linear modelling.

1.2 Thesis outline

The remainder of this thesis is outlined as follows. In § 1.3 we present the governing equations in fluid dynamics and the necessary assumptions required to describe the linear theory behind internal wave motion. In Chapter 2 we then examine the weakly non-linear theory of internal waves. We start by looking at monochromatic plane waves, and how they are full solutions to the non-linear equations of motion. With this in mind, we show why it is necessary to consider internal wave interactions in order to understand how the non-linear terms in the equations of motion play a role. The remainder of the Chapter then provides an overview of the published literature for two weakly non-linear mechanisms; internal wave induced mean flow and TRI.

In Chapter 3 we discuss the experimental methods and diagnostic tools used to conduct the experiments presented in this thesis. The first part of the Chapter focuses on the experimental techniques used for the experiments, while the second part outlines how the experiments were quantitatively processed. The first set of experimental results

are then presented in Chapter 4, where we focus on a particular experimental setup of a finite-width internal wave beam becoming unstable to TRI. In this Chapter we see novel experimental results that contrast with how the previous theory (presented in Chapter 2) describes the evolution of TRI in finite-width internal wave beams.

In light of the results from Chapter 4, Chapter 5 introduces the development of the two-dimensional weakly non-linear model, built to encapsulate both the temporal and spatial properties of the instability. We first present the perturbation expansion used, which captures both the two-dimensional wave advection and the weakly non-linear interactions of the instability. The development of the numerical code used to solve the theoretical equations is then given, along with a validation study. Results from this two-dimensional model, along with an experimental comparison, are presented in Chapter 6. This Chapter focuses on how the input parameters for the secondary waves in the triad have a significant impact on the development of the instability due to the finite beam width.

Further original experimental results are then given in Chapter 7, where we look at how deliberate perturbations to the wave field can be used to trigger triadic resonance. We focus on how a vortex-ring interacting with an internal wave field is able to prematurely trigger the growth of the two secondary waves in the flow. Finally, in Chapter 8, key research findings are summarised along with a discussion and potential avenues for further research.

1.3 Governing equations

1.3.1 Introduction

This section introduces the governing equations that are used to describe the dynamics of internal gravity waves. We start with the momentum equations in § 1.3.2, which are known to govern the majority of fluid dynamical problems. In § 1.3.3 some appropriate assumptions about the flow are given, reducing the complexity of the governing equations but remaining valid for internal waves. In § 1.3.4, we introduce the concept of buoyancy frequency for stratified fluids. Finally, § 1.3.5 outlines the unique dynamics of internal waves in the small amplitude, ‘linear’ limit.

1.3.2 Momentum equations

The momentum conservation laws, also known as the Navier-Stokes equations, are an expression of Newton’s second law $F = ma$, where force F is equal to mass m multiplied

by acceleration a . In terms of fluid dynamics, the conservation law prescribes that the rate of change of momentum of a fluid parcel is equal to the sum of the forces acting on it. As is standard practice, because fluid parcels are infinitesimally small, all forces are assumed to be per unit volume. For a non-rotating fluid in a three-dimensional Cartesian co-ordinate system $\mathbf{x} = (x, y, z)$ we have

$$\rho \frac{D\mathbf{u}}{Dt} = -\nabla P - g\rho\hat{\mathbf{z}} + \mu\nabla^2\mathbf{u}, \quad (1.1)$$

where ρ is total density of the fluid, $\mathbf{u} = (u, v, w)$ is the velocity vector and the terms on the right hand side are outlined below. The left hand side of (1.1) expresses the rate of change of momentum of the fluid parcel through the material derivative D/Dt , given as

$$\frac{D}{Dt} = \frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla, \quad (1.2)$$

where ∇ is the gradient operator in three dimensions given by $(\partial_x, \partial_y, \partial_z)$ and t is time. The material derivative captures changes in the moving reference frame of a fluid parcel. The non-linearity of the advection terms, $\mathbf{u} \cdot \nabla$, make it difficult, if not impossible, to construct solutions from the full Navier-Stokes equations except for a very limited number of special cases.

The right hand side of (1.1) then describes the sum of the possible forces acting on the fluid parcel. The first term, $-\nabla P$, is the pressure gradient term, where P is the total pressure. The negative sign indicates that the force acts from regions of high to low pressure. Assuming incompressibility (discussed in § 1.3.3), in quiescent equilibrium the total pressure at a given point is set by the hydrostatic pressure $p_H(z)$, which is the cumulative weight per horizontal area of fluid lying above that point. A perturbation to the fluid also results in the additional dynamic pressure component, denoted p .

The second term $-g\rho\hat{\mathbf{z}}$ in (1.1) is the gravitational restoring force, also known as the buoyancy force, where $g \approx 9.81 \text{ m s}^{-2}$ is the effective acceleration of gravity. The force acts downwards towards earth, indicated by $-\hat{\mathbf{z}}$. In a quiescent state, it is equal and opposite to the hydrostatic pressure gradient meaning the fluid is in hydrostatic balance. In a stratified environment, if a parcel of fluid is displaced from its equilibrium, the buoyancy force acts on the parcel to return it to its original height.

The final term given in (1.1), $\mu\nabla^2\mathbf{u}$ expresses the viscous force, where μ is the assumed constant dynamic viscosity of the fluid and $\nabla^2 = \partial_{xx} + \partial_{yy} + \partial_{zz}$ is the three dimensional Laplacian operator. To understand the relative importance of viscosity in

a fluid dynamics problem, the dimensionless Reynolds number Re is employed

$$Re = \frac{UL}{\nu}, \quad (1.3)$$

where U is the intrinsic velocity of a system, L is the dominant length scale and ν is the kinematic viscosity, given as $\nu = \mu/\rho$. The Reynolds number provides a ratio of the inertial to viscous forces to characterise which plays a more dominant role in the system. As described by Hazewinkel (2010), for internal waves the relevant Reynolds number is

$$Re = \frac{|\mathbf{c}_g|}{|\mathbf{k}|\nu}, \quad (1.4)$$

where $|\mathbf{c}_g|$ is the magnitude of the group velocity of the wave (discussed in § 1.3.5) that provides the velocity scale U and $\mathbf{k} = (l, m)$ is the wavenumber vector that provides the length scale L . For water, the kinematic viscosity at room temperature is approximated at $\nu \approx 10^{-6} \text{ m}^2 \text{ s}^{-1}$ (Haynes et al., 2014). For oceanic scenarios, $|\mathbf{c}_g|$ can be approximated as 10^{-1} m s^{-1} with a corresponding wavenumber magnitude $|\mathbf{k}|$ of 10^{-4} m^{-1} based off the seafloor topography (Dewitt et al., 1986). This results in a $Re \sim 10^9 \gg 1$. With a Re of this magnitude, the effects of viscosity in the ocean can be considered negligible. However, for the experiments presented here, $|\mathbf{c}_g| \sim 10^{-2} \text{ m s}^{-1}$ and $|\mathbf{k}| \sim 10^2 \text{ m}^{-1}$, which gives a $Re \sim 10^2 > 1$. Based on this Re , while it is appropriate to neglect viscous forces in the ocean, it is not always possible to make this approximation for experiments on smaller scales.

1.3.3 Continuity and incompressibility

The law for the conservation of mass states that a parcel of fluid can not gain or lose mass during its motion. In fluid dynamics it is commonly known as the ‘continuity equation’, given here in the general form

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = \kappa \nabla^2 \rho, \quad (1.5)$$

where κ is the molecular mass diffusivity. This form of the continuity equation in (1.5) neglects the contributions from temperature to the density and assumes that ρ is a function of salinity only. As the experiments presented here are isothermal, this is a reasonable approximation. The dimensionless Schmidt number (ν/κ), can then be used to characterise the ratio of viscous diffusion compared to the molecular mass diffusion. For salt-stratified fluids (as in the case of the experiments) the Schmidt

number is approximately 700, meaning the effects of viscosity are almost three orders of magnitude greater than molecular diffusion and we can assume $\kappa = 0$.

For oceanographic flows, the effects of temperature on the density can not be considered negligible. However, due to the large Re and Pe number ¹ in the ocean, it is also possible to neglect the effects of molecular diffusion here. For the remainder of this thesis we neglect molecular diffusion. Under this assumption, using the vector identity

$$\nabla \cdot (\rho \mathbf{u}) = \mathbf{u} \cdot \nabla \rho + \rho \nabla \cdot \mathbf{u}, \quad (1.6)$$

(1.5) can be re-written as

$$\frac{D\rho}{Dt} = -\rho \nabla \cdot \mathbf{u}, \quad (1.7)$$

where again D/Dt is the material derivative given in (1.2). The left hand side of (1.7) describes the total rate of change in density of a fluid parcel. While this expression does not allow the mass of the fluid parcel to change, the right hand side states that a density change can arise if the flow is divergent. For example, if the mass of the parcel was to remain fixed while it was being compressed due to a converging flow field ($\nabla \cdot \mathbf{u} < 0$), the volume of the parcel would decrease. While compressibility is important for some oceanographic processes, for the study of internal waves it is reasonable to consider the flow as incompressible². This results in the right hand side of (1.7) being equal to 0, as we are requiring that a fluids density, not just its mass, must remain constant. Upon re-arranging, we express the incompressibility condition as

$$\nabla \cdot \mathbf{u} = 0, \quad (1.8)$$

which states that the velocity field must be divergence free. This leaves the conservation of mass equation in the form

$$\frac{D\rho}{Dt} = 0. \quad (1.9)$$

Under the assumption of an incompressible fluid, the velocity vector can be defined by a scalar function Ψ , known as the streamfunction, from which the components of \mathbf{u} can be obtained from its derivatives. Specifically,

$$\mathbf{u} = \nabla \times \Psi, \quad (1.10)$$

¹ Pe is the Péclet number, defined as the ratio of the advective transport rate over the diffusive transport rate.

²The assumption of incompressibility effectively states that the speed of sound is infinite compared to the velocity scales of internal waves. Assuming incompressibility, therefore, amounts to removing the physics of sound waves in the equations of motion.

which automatically satisfies the incompressibility condition given in (1.8). As the experiments presented in this thesis examine quasi two-dimensional internal waves (the wavemaker – discussed in § 3.3 – spans the full width the tank) it is appropriate at this point to consider only the $x - z$ plane. For flows restricted in this plane, $\mathbf{u} = \nabla \times (\Psi \hat{\mathbf{y}})$, leaving

$$\mathbf{u} = (u, w) = \left(-\frac{\partial \Psi}{\partial z}, \frac{\partial \Psi}{\partial x} \right). \quad (1.11)$$

For the remainder of this work, we will only focus on two-dimensional flows.

1.3.4 Buoyancy frequency

The density stratification that is present within the ocean, even in the weakly stratified abyss, provides the restoring buoyancy force required for the propagation of internal waves (Sutherland et al., 2000). For stratified fluids it is appropriate to decompose the total density ϱ as

$$\varrho(x, z, t) = \varrho_0 + \bar{\rho}(z) + \rho(x, z, t), \quad (1.12)$$

where ϱ_0 is the reference density, taken as 10^3 kg m^{-3} , $\bar{\rho}(z)$ is the background stratification, which describes the temporally and horizontally averaged density across the depth and ρ is the time dependant perturbation density in two dimensions (x, z) , due to internal wave motion. For a stratified fluid, the equilibrium state is given as the sum of the reference density ϱ_0 , and a correction accounting for the change in density across the depth $\bar{\rho}(z)$.

We then invoke the ‘Boussinesq’ approximation, which assumes that the change in density of the stratification is small compared with its absolute density and so can be ignored except where it gives rise to buoyancy forces (Turner, 1979). In order for the Boussinesq approximation to be valid, the advective accelerations in the system must be much smaller than the acceleration due to gravity $|D\mathbf{u}/Dt| \ll |g|$. In relation to the momentum equation given in (1.1), it amounts to ϱ on the left hand side being equal to ϱ_0 only. It also states that ν is not a function of density and remains constant throughout. According to Sutherland (2010), this is an appropriate approximation for internal waves in the ocean, where there density varies by only a small fraction of ϱ_0 over the depth. For the experiments presented in this thesis, the densities of the stratifications vary by up to 10 %. While this is considerably larger than the oceanic density range, as the accelerations in the experiments remain below the gravitational acceleration, the Boussinesq approximation remains valid.

Under this assumption of a Boussinesq fluid, the buoyancy restoring force is characterised by the ‘buoyancy frequency’ or ‘Brunt–Väisälä frequency’

$$N = \sqrt{-\frac{g}{\varrho_0} \frac{d\bar{\rho}}{dz}}. \quad (1.13)$$

For the rest of this thesis we will refer to ‘buoyancy frequency’ as opposed to ‘Brunt–Väisälä frequency’. This definition of the buoyancy frequency is only appropriate for liquids due to their assumed incompressibility. For air, the anelastic approximation is required (Sutherland, 2010).

1.3.5 Linear internal waves

The term ‘linear’ is used here to describe internal waves which satisfy

$$\epsilon = (|\mathbf{k}|^2 |\Psi|) / N \ll 1, \quad (1.14)$$

where $|\Psi|$ is the magnitude of the wave streamfunction from (1.11), used here as a measure of wave amplitude, N is the buoyancy frequency given in (1.13) and $|\mathbf{k}|$ is the magnitude of the wavenumber vector. In two dimensions, $\mathbf{k} = (l, m) = |\mathbf{k}|(\cos \theta, \sin \theta)$, where θ describes the angle that the lines of constant phase make with the vertical $\hat{\mathbf{z}}$, as shown in Figure 1.2. Due to the lines of constant phase being perpendicular to the wavenumber vector, θ is also the angle between \mathbf{k} and the horizontal $\hat{\mathbf{x}}$. The non-dimensional parameter ϵ , provides a measure of the relative importance of the advection terms, $\mathbf{u} \cdot \nabla$, in the momentum equations (1.1). As these advection terms are quadratic, their contribution is negligible under the assumption of linearity ($\epsilon \ll 1$) and the material derivative given in (1.2) simplifies to $D/Dt \equiv \partial/\partial t$. This is known as the ‘small amplitude’, linear limit.

The mixing events generated by oceanic internal waves, discussed in (1.1), generally occur due to non-linear processes, as wave amplitudes often exceed this linear limit. While it is therefore necessary to understand these non-linear processes, we start here by looking at the properties of linear internal waves.

Applying the Boussinesq approximation, the linearised form of the momentum equation is given as

$$\varrho_0 \frac{\partial \mathbf{u}}{\partial t} = -\nabla p + g\rho \hat{\mathbf{z}} + \mu \nabla^2 \mathbf{u}, \quad (1.15)$$

in which $\nabla^2 = \partial_{xx} + \partial_{zz}$ is now the Laplacian operator in the $x - z$ plane. Only the dynamic pressure p and density fluctuations ρ responsible for buoyancy driven

motion remain on the right hand side of (1.15) due to the assumption of background hydrostatic balance, whereby $dp_H/dz = -g\bar{\rho}$.

Using the density decomposition given in (1.12), the conservation of mass (1.9) can also be written in its linear form

$$\frac{\partial \rho}{\partial t} = -w \frac{d\bar{\rho}}{dz}, \quad (1.16)$$

where w is the vertical component of the two-dimensional velocity vector \mathbf{u} . It is then possible to combine the linearised equations (1.15) and (1.16) into a one equation, using the streamfunction definition for velocity given in (1.11). To do this we take the curl of the linearised momentum equation (1.15) to eliminate pressure and divide through by ρ_0 . Then, taking the partial derivative with respect to time of both sides leaves

$$\frac{\partial^2}{\partial t^2} \nabla^2 \Psi + \frac{g}{\rho_0} \frac{\partial}{\partial x} \frac{\partial \rho}{\partial t} - \nu \frac{\partial}{\partial t} \nabla^2 (\nabla^2 \Psi) = 0, \quad (1.17)$$

where $\nu = \mu/\rho_0$. Eliminating density from the second term of (1.17) using the linear conservation of mass given in (1.16) and substituting the definition for buoyancy frequency in (1.13), results in

$$\frac{\partial^2}{\partial t^2} \nabla^2 \Psi + N^2 \frac{\partial^2 \Psi}{\partial x^2} - \nu \frac{\partial}{\partial t} \nabla^2 (\nabla^2 \Psi) = 0. \quad (1.18)$$

We now have a linearised equation in terms of the single scalar variable Ψ to describe the dynamics of small amplitude, linear internal waves.

It is then possible to seek solutions of the streamfunction in order to solve (1.18). Assuming an inviscid flow ($\nu = 0$) and that Ψ represents an infinite sinusoidal disturbance, known as a ‘plane’ or ‘monochromatic’ wave, Ψ can be mathematically described by

$$\Psi(\mathbf{x}, t) = \tilde{\Psi} e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}, \quad (1.19)$$

where the complex amplitude of the wave is given by the reduced streamfunction $\tilde{\Psi}$ and the periodic phase as $e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}$. The wavenumber vector and frequency of the wave are given as $\mathbf{k} = (l, m)$ and ω respectively. Substituting (1.19) into (1.18) results in the linear dispersion relationship for internal waves,

$$\frac{\omega}{N} = \pm \cos \theta = \pm \frac{|l|}{\sqrt{l^2 + m^2}}. \quad (1.20)$$

This dispersion relationship emphasises a clear difference between internal waves and surface gravity waves, in that the relationship dictates the angle of propagation, as

opposed to the wavelength. Since $0 \leq |\cos \theta| \leq 1$, the dispersion relationship states that the maximum frequency of oscillation is the buoyancy frequency, N . Above this limit, any oscillations will be evanescent disturbances that decay exponentially away from the generation source. From the dispersion relation (1.20) we can define the phase velocity \mathbf{c}_p as

$$\mathbf{c}_p = \frac{\omega}{|\mathbf{k}|^2} \mathbf{k} = \frac{Nl}{|\mathbf{k}|^3} (l, m). \quad (1.21)$$

The group velocity \mathbf{c}_g , which defines the speed at which the energy propagates can then be defined as

$$\mathbf{c}_g = \left(\frac{\partial}{\partial l}, \frac{\partial}{\partial m} \right) \omega = \text{sgn}(l) \frac{Nm}{|\mathbf{k}|^3} (m, -l), \quad (1.22)$$

where $\text{sgn}(l)$ gives direction of the horizontal component of the wavenumber vector.

One of the key distinguishing features of internal waves are that their group velocity \mathbf{c}_g and phase velocity \mathbf{c}_p are perpendicular. They have the same horizontal component ($c_{p_x} = c_{g_x}$) but opposite vertical component ($c_{p_z} = -c_{g_z}$), leaving

$$\mathbf{c}_p \cdot \mathbf{c}_g = 0. \quad (1.23)$$

This means that the summation of the two vectors results in a horizontal distance of

$$\mathbf{c}_p + \mathbf{c}_g = \frac{N}{|\mathbf{k}|}. \quad (1.24)$$

This unique relationship was demonstrated by the oscillating cylinder experiment of Mowbray and Rarity (1967), illustrated in Figure 1.2. The experiment looks at the internal waves produced when oscillating a cylinder vertically in a linearly stratified fluid. Four internal wave beams are produced from this oscillation, propagating away at angle θ from the vertical in both the positive and negative direction. This is an illustration of the anisotropic propagation permitted by the vertical stratification, shown by the unique dispersion relationship (1.20), which permits four possible sign combinations. In each case, the direction of the beams is given by their group velocity, while the perpendicular phase velocities point towards the equilibrium height of the cylinder.

In this chapter we have seen the linear dynamics of internal waves in the context of a linearly stratified, incompressible, Boussinesq fluid. Yet while an understanding of the linear properties of internal waves is fundamental to understanding their motion, it is also essential to understand the non-linear processes of these waves to allow for

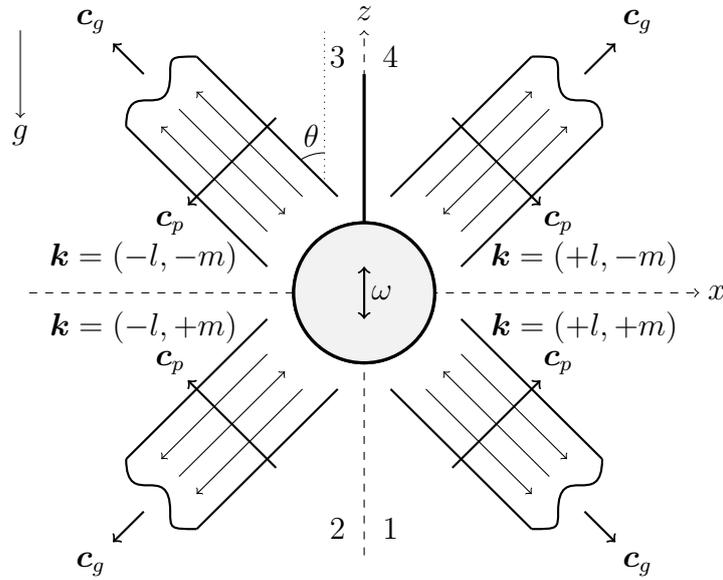


Figure 1.2: Schematic of the oscillating cylinder experiment conducted by Mowbray and Rarity (1967). A cylinder is oscillated vertically at frequency ω in a linear stratification N . The motion produces four wave beams that propagate outwards at angle $\theta = \omega/N$ with respect to the vertical, with group velocity \mathbf{c}_g . The wavenumber vectors \mathbf{k} and phase velocities \mathbf{c}_p are perpendicular to the direction of energy propagation. The four different sign combinations for \mathbf{k} that satisfy linear dispersion (1.20) are indicated for each beam. The quadrants correspond to the quadrants used for the Hilbert Transform discussed in § 3.9.3.

a better comprehension of their energy transfer mechanisms to smaller scales. In the following chapter we introduce the weakly non-linear dynamics of internal waves.

Chapter 2

Weakly non-linear theoretical foundations

2.1 Introduction

The theory presented in Chapter 1 outlines the governing equations of internal waves and their dynamics at leading order. However, truncating solutions in this linear limit provides little insight into the contribution of internal waves to energy transfer and hence global ocean mixing. Indeed, decades of observational records outlined in Garrett and Munk (1972) have shown the presence of an internal wave continuum within the ocean, which cannot be interpreted solely by fundamental tidal modes. This spectrum, known as the ‘energy cascade’, reveals the presence of mechanisms whereby large-scale energy input is converted to smaller spatio-temporal scales (Garrett and Kunze, 2007). These smaller scale internal waves then have the potential to generate localised turbulent mixing, through ‘breaking’ or dissipation (Wunsch and Ferrari, 2004). Yet while this internal wave continuum is known to exist, the mechanisms by which internal waves transfer energy are not so well understood. According to Alford et al. (2015) and MacKinnon et al. (2017) for internal waves, our comprehension of their instability mechanisms, as opposed to their generation, is a much more open question.

This chapter outlines the theory behind two weakly non-linear instability mechanisms of internal waves that are currently believed to play a key role in contributing to the energy cascade. Specifically we focus on internal wave induced mean flow and TRI. The term ‘weakly non-linear’ is used to represent waves that slightly violate the linear limit given in (1.14), so that the role of the advection terms in the momentum equations can no longer be neglected. We first examine the unusual property of monochromatic

plane waves being a solution to the full non-linear equations in § 2.2. We then introduce the concept of triadic resonance in § 2.3 and how non-linear behaviour can be described when considering wave interactions. The first weakly non-linear instability is then introduced in § 2.4, where we examine how triadic interactions play a role in internal wave induced mean flow. Finally, in § 2.5, the Triadic Resonance Instability (TRI) mechanism is examined in more detail. A brief literature review of TRI is presented in § 2.5.1, while § 2.5.3 presents the current weakly non-linear theory for the instability.

2.2 Plane wave solutions

The form of the streamfunction given in (1.19) represents an inviscid plane monochromatic wave. This means that the wave can be modelled by a perfect sinusoidal disturbance with a single frequency and wavenumber for all space and time. Assuming an incompressible, non-rotating, and linearly stratified Boussinesq fluid, this wave form holds the peculiar property of being an exact solution to the full non-linear equations at any amplitude (e.g. Thorpe, 1968; Thorpe and Haines, 1986; Sutherland, 2006). However, while single monochromatic plane waves are convenient mathematically, in nature waves will never take this form. Realistically, they will manifest as beams, confined locally in space and therefore broadly distributed over the wave number spectrum. The focus of analysis using plane wave solutions has been highlighted in the review by Dauxois et al. (2018), who argue that the effects of finite-width and envelope shape have been overlooked when considering the non-linearities of internal waves.

Unexpectedly however, work by Tabaei and Akylas (2003), showed that even the more general solution of time-harmonic uniform plane wave beams in the inviscid limit are also full solutions to the non-linear equations of motion. Here, the superposition of several linear solutions, associated with the same frequency but different wave number, results in the cancellation of the non-linear terms in the equations of motion, leaving only linear solutions. This is due to the uniformity on the along beam direction, which will cause the non-linear terms to vanish even for a wave beam of any general cross beam profile.

Yet despite these wave forms being full solutions to the equations of motion, it does not mean that they are stable ones. This is a fundamental concept in the analysis of both the instabilities discussed in this section, which highlight how interactions between internal waves – even of infinitesimal amplitude – cause the non-linearities in the equations of motion begin to play a role. The remainder of this research is

focused on considering more realistic wave beam configurations along with interactions, in order to capture the weakly non-linear mechanisms of internal waves.

2.3 Triadic resonance

We start by considering the general case of triadic resonance between internal gravity waves. As just discussed in § 2.2, it is necessary to consider the interactions between internal waves in order for the non-linearities to become apparent. This was first explored in the early 1960's by Phillips (1960) and Longuet-Higgins (1962) who examined interactions in the context of surface gravity waves and later extended the theory to internal waves (the review by Phillips (1981) provides a comprehensive account of these relationships and research history). Work by Thorpe (1966) showed that an energy exchange between three internal waves would occur when the frequencies and wavenumbers present in the flow satisfied the condition

$$\mp \phi_0 \pm \phi_1 \pm \phi_2 = 0, \quad (2.1)$$

where ϕ_0 , ϕ_1 and ϕ_2 are the phases of the three triadic waves respectively. The wave phase is defined as

$$\phi = (\mathbf{k} \cdot \mathbf{x} - \omega t), \quad (2.2)$$

where, again, the wavenumber vector is given by $\mathbf{k} = (l, m)$ and the internal wave frequency as ω . The work presented by Thorpe (1966) also showed how two internal waves of fixed frequency in a flow can interact to generate a third wave that satisfies the spatio-temporal resonant condition

$$\pm \phi_2 = \pm \phi_0 \mp \phi_1, \quad (2.3)$$

where ϕ_0 and ϕ_1 are the original wave phases and ϕ_2 is the triadically produced wave phase. We define this interaction in (2.3) as a Two Wave Interaction (TWI). It should be noted that in order for net energy transfer to occur from ϕ_0 and ϕ_1 to generate the resonant ϕ_2 , the frequencies of the three waves must satisfy $|\omega_2| + |\omega_1| = |\omega_0|$. If this temporal condition is not met, while the triad will be resonant, no net exchange of energy will occur. Even in the context of a finite-width beam, a TWI will occur at any amplitude (Davis et al., 2020). Before discussing the more particular case of TRI, we first use this concept of triadic resonance to consider internal wave induced mean flow.

2.4 Internal wave induced mean flow

The generation of three-dimensional background mean flow from internal waves has been cited as one of the potential mechanisms capable of inducing material transport and promoting ocean mixing (Semin et al., 2016). Yet while its presence has been widely considered in atmospheric contexts, its potential role in the oceanic energy cascade has been mainly overlooked (Grisouarda et al., 2013).

The term ‘mean flow’ in the context of internal waves commonly refers to the horizontal mean flow associated with mean vertical vorticity. This type of mean flow, known as ‘streaming’ (Kataoka and Akylas, 2015), is a dissipative mechanism whereby a slowly evolving, non-oscillatory Eulerian flow is generated by the weakly non-linear interactions of a wavebeam with itself. This phenomena has been well documented for other waves, such as ‘acoustic streaming’ for sound waves (Lighthill, 1978) or the currents generated parallel to the coastline due to surface gravity waves (Longuet-Higgins, 1970). In fact, a review by Riley (2001) found that this non-linear mechanism is present in a many homogenous flows.

For internal waves, both viscous attenuation and cross-beam variations are required in order to develop this vortical induced mean flow (Bordes et al., 2012). One of its key features is the irreversible, cumulative transfer of energy from the wave field. While many experimental studies have shown the presence of this three-dimensional flow (Semin et al., 2016; Kataoka et al., 2017), the most complete is probably by Bordes et al. (2012) who reported a strong mean flow generated by a time-harmonic internal wave beam propagating in a tank wider than the beam width. This was important, as it allowed for a recirculating mean flow to develop in the transverse direction, outside of the spatial extent of the beam. As noted by Sutherland (2006), in experiments where waves are confined laterally by tank side walls, horizontal mean flow of this type is unable to develop.

While mean flow can be generated via the interaction of the wave beam with itself, as described above, it is also possible for it occur via the interaction of a wave beam with its reflection by means of a TWI. Due to the dispersion relationship of internal waves, they retain their frequency upon reflection from a boundary, meaning their angle of propagation is also preserved (e.g. Maas and Lam, 1995). A wave reflected from a horizontal boundary or free surface will have a wavenumber vector with an opposite vertical component to the incident beam. Based on the triadic relationship of (2.3) these wave vectors will sum to produce a third ‘wave’ with wavenumber vector aligned with the vertical but with zero frequency and zero horizontal wavenumber (Tabaei and

Akylas, 2003). This non-propagating disturbance can not be classed as a wave but as an oscillatory structure that is confined to the interaction region of the primary beam with its reflection. If considered in a two-dimensional setting, only horizontal vorticity is generated, however this is generally weaker as it is suppressed by the background stratification (Beckebanze et al., 2019).

Both two and three-dimensional numerical simulations were performed by Grisouarda et al. (2013), to better understand the nature of the mean flow generated by a reflected wavebeam. For the two-dimensional case, they found that an Eulerian periodic along-slope current developed, as predicted by the weakly non-linear inviscid two-dimensional analysis of both Thorpe and Haines (1986) and Tabaei et al. (2005). The simulations also agreed with Tabaei et al. (2005) that, in two-dimensions this interaction will be confined solely in the vicinity where the beams interact. In contrast, for their three-dimensional simulations, they witnessed the slow development of the much stronger horizontal mean flow outlined above. Their three-dimensional numerics were also in much closer agreement with their laboratory experiments, which, similar to Bordes et al. (2012), allowed for a horizontal circulating flow to develop in the traverse direction.

2.5 Triadic Resonance Instability

2.5.1 Overview

In addition to wave induced mean flow, we also examine a second weakly non-linear mechanism known as Triadic Resonance Instability (TRI). The theory for TRI is based on the resonant triad interactions given in § 2.3, although unlike a TWI, where two waves interact to produce a third wave, TRI satisfies the more particular spatio-temporal condition

$$\phi_0 = \phi_1 + \phi_2, \quad (2.4)$$

where, again, ϕ defines the wave phase. For the specific TRI case, a single primary wave becomes unstable due to the growth of infinitesimal perturbations in the flow. This growth leads to the formation of a spatio-temporal resonant triad, whereby two secondary waves of lower frequency extract energy from the primary wave (Dauxois et al., 2018). TRI has been cited by Joubaud et al. (2012) as an efficient way to transfer energy from larger to smaller scales, giving it a potential role in the energy cascade within the ocean. Indeed, studies by MacKinnon and Winters (2005) and Alford et al. (2007) suggest that, particularly near a critical latitude of 28.8°N, TRI could play a dominant role in oceanic energy dissipation.

The first reported experimental evidence of TRI for internal and interfacial waves was approximately 50 years ago by Davis and Acrivos (1967) and McEwan (1971). Around the same time, the phenomena was being explored in other fields. Theoretical work by Hasselmann (1967) showed that instability was present in many physical systems as well as for internal waves. Tsytovich (1970) described it as the ‘decay instability’ in plasma physics and McGoldrick (1965) explained the same instability in the context of capillary waves.

In the idealised scenario of an inviscid infinite plane wave, the frequencies of the secondary waves in the triad are equal to half that of the primary wave, motivating the traditional terminology of Parametric Sub-harmonic Instability (PSI) (Fan and Akylas, 2019). In laboratory experiments, however, the role of viscosity cannot be neglected and resonant wave frequencies deviate away from this relationship. For the context of this thesis, it is therefore necessary to use the more appropriate terminology of TRI.

As discussed in § 2.2, in order to adequately represent the form of internal waves in nature, recent focus has been on exploring the dynamics of TRI in spatially localised internal wave beams. Unlike an infinite plane wave that will be unstable to infinitesimal perturbations at any amplitude (Koudella and Staquet, 2006), locally confined beams will only become unstable to TRI above a given amplitude threshold. McEwan et al. (1972) derives one the first analytical expressions for the so-called ‘critical amplitude’, which has since been developed by Bourget et al. (2014) who calculate a growth rate for the instability that includes these finite-width effects. They also show numerically that as the beam width is increased the amplitude threshold decreases. This is due to the perturbations having longer to interact with the underlying primary beam, increasing the region over which energy can be transferred. Recent theoretical work by Karimi and Akylas (2014) has shown how the form of the carrier envelope for a finite-width internal wave-beam has a significant impact on its ability to become unstable.

In addition to theoretical studies, much experimental work has been conducted on the instability (e.g. McEwan and Plumb, 1977; Benielli and Sommeria, 1998). Brouzet et al. (2016) showed how focusing reflections from internal wave attractors leads to steepening of the wave beam which then permits the onset of TRI due to increased amplitude. Experimental studies presented by Joubaud et al. (2012) and Bourget et al. (2013) also provide quantitative measurements of the growth rate for the instability with a comparison against theoretical values.

However, while many studies cite TRI as a mechanism that leads to eventual oceanic mixing, TRI can act to limit the wave steepening that would otherwise lead to breaking, meaning less mixing may occur from the primary wave. Specifically, TRI

transfers energy out of the primary wave field into other wavenumbers and frequencies, limiting the energy available in the primary wave for other phenomena that could lead more directly to wave breaking and mixing. It is clear, therefore, that non-linearities in internal wave beams are fundamental in their dynamics and the pathways towards instability. The review paper by Dauxois et al. (2018) examines the progress on our understanding of the non-linear destabilization of internal wave beams in order to explain their eventual breakdown and their role in energy transfer to dissipative scales.

2.5.2 TRI loci curve

Before examining the weakly non-linear triad energy exchange in § 2.5.3, we briefly look at the form of all the possible secondary waves produced via TRI. Any propagating wave produced by a triadic resonant interaction (e.g. TWI, TRI) must also satisfy the dispersion relationship given in (1.20). Recasting (1.20) in terms of the spatial component of the TRI condition (2.4) provides a relationship for the first resonant wave vector $\mathbf{k}_1 = (l_1, m_1)$ in terms of the primary wave $\mathbf{k}_0 = (l_0, m_0)$

$$\frac{|l_0|}{\sqrt{l_0^2 + m_0^2}} = \frac{|l_1|}{\sqrt{l_1^2 + m_1^2}} + \frac{|l_0 - l_1|}{\sqrt{(l_0 - l_1)^2 + (m_0 - m_1)^2}}. \quad (2.5)$$

The relationship in (2.5) provides all the possible solutions for the of the \mathbf{k}_1 vector produced from the TRI given the \mathbf{k}_0 vector. This is shown the black curve on Figure 2.1, which provides all the possible locations for the tip of the \mathbf{k}_1 vector whose origin lines at (0,0).

All these solutions shown by the black line on Figure 2.1 have a positive growth rate. The other red dashed lines on the figure, correspond to other possible solutions determined by different sign combinations of the two terms on the right hand side of (2.5). These solutions are neutrally stable, and will not be accessed via the TRI mechanism.

2.5.3 Weakly non-linear TRI theory

Under the assumption of a Boussinesq, and hence incompressible fluid, we start with the non-linear equations of motion, which are comprised of the full two-dimensional momentum equations (1.1), the conservation of mass equation (1.9) and the continuity equation (1.8). Again assuming background hydrostatic balance, these non-linear

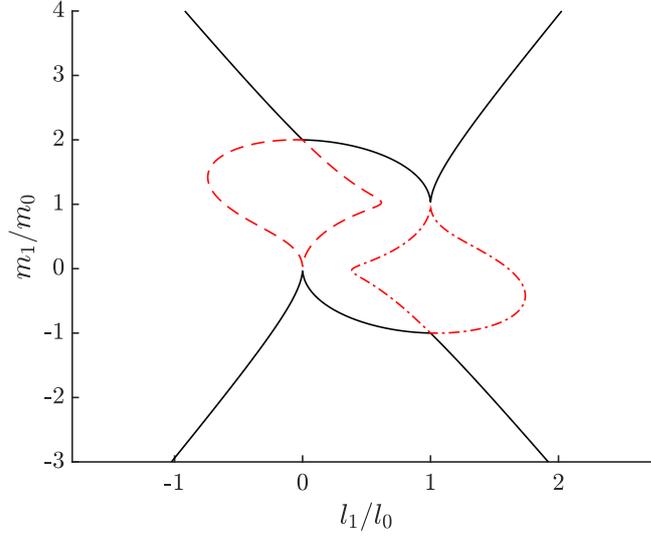


Figure 2.1: Figure showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector with an origin at $(0,0)$. The black line indicates all the possible locations that satisfy (2.5) corresponding to the TRI condition (2.4). The red dashed lines indicate other possible solutions that are neutrally stable, based on different sign combinations of the two terms on the right hand side of (2.5). The vectors have been normalised by \mathbf{k}_0 .

equations are reproduced here as

$$\varrho_0 \frac{Du}{Dt} = -\frac{\partial p}{\partial x} + \mu \nabla^2 u, \quad (2.6)$$

$$\varrho_0 \frac{Dw}{Dt} = -\frac{\partial p}{\partial z} - g\rho + \mu \nabla^2 w, \quad (2.7)$$

$$\frac{D\rho}{Dt} = -w \frac{d\bar{\rho}}{dz}, \quad (2.8)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2.9)$$

where the velocity vector is $\mathbf{u} = (u, w)$, g is the gravitational constant of acceleration, p is the dynamic pressure, μ is the dynamic viscosity and the density decomposition is defined by (1.12). We then wish to reduce this set of coupled partial differential equations given in (2.6), (2.7), (2.8) into a single equation. This is done by taking the curl of the momentum equations (2.6), (2.7), to eliminate pressure, leaving a form of the vorticity equation

$$\frac{\partial}{\partial x} \left(\varrho_0 \frac{Dw}{Dt} + \rho g - \mu \nabla^2 w \right) - \frac{\partial}{\partial z} \left(\varrho_0 \frac{Du}{Dt} - \mu \nabla^2 u \right) = 0. \quad (2.10)$$

From here, we substitute the definition of velocity in (1.11) given in terms of the streamfunction, where $\mathbf{u} = (u, w) = (-\partial\Psi/\partial z, \partial\Psi/\partial x)$. Then dividing through by ϱ_0 and simplifying, (2.10) can be re-written as

$$\frac{D}{Dt}(\nabla^2\Psi) + \frac{g}{\varrho_0} \frac{\partial\rho}{\partial x} - \nu\nabla^2(\nabla^2\Psi) = 0. \quad (2.11)$$

We then take the material derivative of (2.11), which leaves a single equation that is independent of pressure

$$\frac{D^2}{Dt^2}(\nabla^2\Psi) + \frac{g}{\varrho_0} \frac{D}{Dt} \frac{\partial\rho}{\partial x} - \frac{D}{Dt} \nu\nabla^2(\nabla^2\Psi) = 0. \quad (2.12)$$

We then take the partial derivative with respect to x of the density advection (2.8) and multiply through by g/ϱ_0 . Then, by substituting both the streamfunction definition of velocity (1.11) and the definition of the buoyancy frequency (1.13), (2.8) becomes

$$\frac{g}{\varrho_0} \frac{\partial}{\partial x} \frac{D\rho}{Dt} - N^2 \frac{\partial^2\Psi}{\partial x^2} = 0. \quad (2.13)$$

As both (2.12) and (2.13) are equal to zero, we can subtract (2.13) away from (2.12), which results in a single equation in terms of both Ψ and ρ

$$\frac{D^2}{Dt^2}(\nabla^2\Psi) - \frac{D}{Dt} \nu\nabla^2(\nabla^2\Psi) + N^2 \frac{\partial^2\Psi}{\partial x^2} = \frac{g}{\varrho_0} \left(\frac{\partial}{\partial x} \frac{D}{Dt} - \frac{D}{Dt} \frac{\partial}{\partial x} \right) \rho. \quad (2.14)$$

Here we are left with the full non-linear momentum equations as a function of the streamfunction and density, which, in the inviscid, linear limit reduces to (1.18). It is then necessary to determine the form of the streamfunction that will accurately govern the development of TRI. As discussed in § 2.5, if we consider only a single monochromatic plane wave, the non-linear terms in (2.14) would cancel, as plane waves are solutions to the full non-linear equations (Thorpe, 1968). We therefore start by considering the structure of streamfunction first proposed by McEwan and Plumb (1977) and further developed by Koudella and Staquet (2006) and Bourget et al. (2013). Here, their approach was to examine the development of the flow by considering the

summation structure of the streamfunction and perturbation density

$$\Psi = \sum_{p=0}^2 \tilde{\Psi}_p(t) e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)} + \text{c.c.}, \quad (2.15a)$$

$$\rho = \sum_{p=0}^2 \tilde{\rho}_p(t) e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)} + \text{c.c.}, \quad (2.15b)$$

where $\tilde{\Psi}$ and $\tilde{\rho}$ are the reduced forms of the streamfunction and density respectively and c.c. represents the complex conjugate. The summation across p in (2.15a) and (2.15b) represents the assumption that the flow is comprised of three triadic waves that satisfy the relationship given in (2.4).

Before substituting these summations into (2.14), the linear inviscid conservation of mass (1.16) can be used to represent ρ in terms of Ψ . First, (1.16) is expanded using the definition of the streamfunction

$$\frac{\partial \rho}{\partial t} = -w \frac{d\bar{\rho}}{dz} = \frac{\partial \Psi}{\partial x} \frac{d\bar{\rho}}{dz} = \frac{\partial \Psi}{\partial x} \frac{N^2 \varrho_0}{g}. \quad (2.16)$$

We then substitute the forms $\rho = \tilde{\rho} e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}$ and $\Psi = \tilde{\Psi} e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}$. Taking the partial derivatives on both sides and cancelling the $e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}$ associated the fast time components provides the polarisation relationship

$$\tilde{\rho} = -\frac{l}{\omega} \frac{N^2 \varrho_0}{g} \tilde{\Psi}. \quad (2.17)$$

Experimental evidence suggests that the amplitude envelope of the wave beam evolves slowly compared to the temporal phase evolution ($|\partial \tilde{\Psi}_p / \partial t| \ll |\omega_p \tilde{\Psi}_p|$). Under this ‘slow time’ assumption, the above polarisation relationship (2.17) becomes

$$\frac{\partial \tilde{\rho}}{\partial t} \approx -\frac{l}{\omega} \frac{N^2 \varrho_0}{g} \frac{\partial \tilde{\Psi}}{\partial t}. \quad (2.18)$$

Bourget et al. (2013) then neglect the time derivative in Fourier space, seeking the interactions that will govern the slowly evolving amplitude of the wave beam. Substituting (2.15a) and (2.15b) into (2.14), and using the relationship in (2.18) to remove density, the following three ordinary differential equations (ODEs) are obtained

$$\frac{d\tilde{\Psi}_0}{dt} = I_0\tilde{\Psi}_1\tilde{\Psi}_2 - \nu\left(\frac{\kappa_0^2}{2}\right)\tilde{\Psi}_0, \quad (2.19a)$$

$$\frac{d\tilde{\Psi}_1}{dt} = I_1\tilde{\Psi}_0\tilde{\Psi}_2^* - \nu\left(\frac{\kappa_1^2}{2}\right)\tilde{\Psi}_1, \quad (2.19b)$$

$$\frac{d\tilde{\Psi}_2}{dt} = I_2\tilde{\Psi}_0\tilde{\Psi}_1^* - \nu\left(\frac{\kappa_2^2}{2}\right)\tilde{\Psi}_2, \quad (2.19c)$$

where $\kappa = |\mathbf{k}|$ and I_p are the interaction terms defined by

$$I_p = \frac{l_q m_r - m_q l_r}{2\omega_p \kappa_p^2} \left[\omega_p (\kappa_q^2 - \kappa_r^2) + l_p N^2 \left(\frac{l_q}{\omega_q} - \frac{l_r}{\omega_r} \right) \right], \quad (2.20)$$

with subscripts p, q, r corresponding to the circular permutations of the triad waves $(0, 1, 2)$ respectively for each equation. See Bourget et al. (2013) for full details of the derivation of these equations. From here, they conduct a linear stability analysis, assuming that at early times the amplitudes of $\tilde{\Psi}_1$ and $\tilde{\Psi}_2$ are negligible in comparison to $\tilde{\Psi}_0$. Under this assumption, it is possible to combine (2.19b) and (2.19c) to get the second order differential equation

$$\frac{d^2\tilde{\Psi}_1}{dt^2} + \frac{\nu}{2}(\kappa_1^2 + \kappa_2^2) \frac{d\tilde{\Psi}_1}{dt} - \left(I_1 I_2 \tilde{\Psi}_0^2 \tilde{\Psi}_1 - \frac{\nu^2}{4} \kappa_1^2 \kappa_2^2 \right) \tilde{\Psi}_1 = 0. \quad (2.21)$$

The solution to (2.21) leads to the expression $\tilde{\Psi}_{1,2}(t) = A_{1,2}e^{\sigma_+ t} + B_{1,2}e^{\sigma_- t}$ by introducing the growth rate term

$$\sigma_{\pm} = -\frac{\nu}{4}(\kappa_1^2 + \kappa_2^2) \pm \sqrt{\frac{\nu^2}{16}(\kappa_1^2 - \kappa_2^2)^2 + |I_1||I_2||\tilde{\Psi}_0|^2}, \quad (2.22)$$

where A and B are arbitrary constants. Using this expression of growth rate, a vanishingly small amplitude noise will induce linear instability of the primary wave causing the growth of two resonant waves at any amplitude.

While this theory considers the ‘slow-time’ development of the amplitude of the beams, as noted by Sutherland (2013), it is still based upon the assumption that the waves are monochromatic in space and time. Realistically, internal waves will manifest as beams, confined locally in space and therefore broadly distributed over the wave number spectrum. If this beam becomes unstable, the resonant waves have a finite time with which to interact with the underlying primary beam, before leaving the interaction region, at which point the energy transfer will be broken (McEwan and

Plumb, 1977). The role of a finite-width wavebeam on TRI is examined theoretically by Karimi and Akylas (2014) who looked at the effect of the carrier envelope to the primary wave beam. This investigation revealed that only for quasi monochromatic wave beams, comprising a sinusoidal carrier modulated by a locally confined envelope, will short-scale sub-harmonic perturbations be able to evolve. If the width of the locally confined wave beams is too small, then the triadic interaction will not be sustained for sufficient time to allow for the development of TRI.

Noting this shortcoming, a revised model is presented in Bourget et al. (2014), where they attempt to take into account the effect of having a finite-width internal wave beam. Using an energy balance, they define a control area over which the resonant beams can interact with the primary, shown by the reproduced diagram in Figure 2.2.

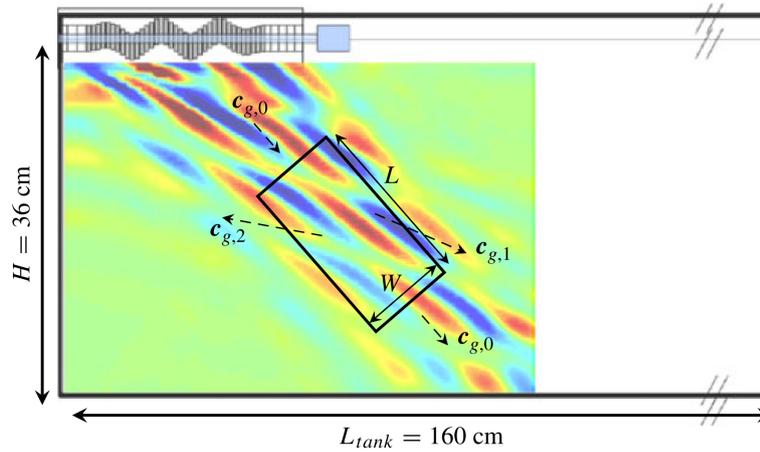


Figure 2.2: A diagram by Bourget et al. (2014), showing their experimental setup in which a primary wave beam is generated from the surface mounted wavemaker. The primary beam has become unstable to TRI and the direction of the resonant waves group velocities is indicated. The control area used to define the interaction region over which energy exchange can occur between the triadic waves is outlined by the black rectangle.

Accounting for the energy flux through the control area via non-linear interactions, viscous attenuation and incoming and outgoing energy flux of the primary beam, the ODEs given in (2.19) become

$$\frac{d\tilde{\Psi}_0}{dt} = I_0 \tilde{\Psi}_1 \tilde{\Psi}_2 - \nu \left(\frac{\kappa_0^2}{2} \right) \tilde{\Psi}_0 + F, \quad (2.23a)$$

$$\frac{d\tilde{\Psi}_1}{dt} = I_1 \tilde{\Psi}_0 \tilde{\Psi}_2^* - \left(\frac{\nu \kappa_1^2}{2} + \frac{|\mathbf{c}_{g1} \cdot \mathbf{e}_{k_0}|}{2W} \right) \tilde{\Psi}_1, \quad (2.23b)$$

$$\frac{d\tilde{\Psi}_2}{dt} = I_2 \tilde{\Psi}_0 \tilde{\Psi}_1^* - \left(\frac{\nu \kappa_2^2}{2} + \frac{|\mathbf{c}_{g2} \cdot \mathbf{e}_{k_0}|}{2W} \right) \tilde{\Psi}_2. \quad (2.23c)$$

The forcing term $F = \mathbf{c}_{g_0}(\tilde{\Psi}_{\text{in}}^* \tilde{\Psi}_{\text{in}} - \tilde{\Psi}_0^* \tilde{\Psi}_0)/(2L\tilde{\Psi}_0^*)$ on the end of (2.23a), represents the energy flux for the primary wave through the control area. The terms on the end of (2.23b) and (2.23c) represent the viscous decay within and flux of energy out of the interaction region of width W . For the remainder of this thesis, we will refer to the above set of zero-dimensional ODEs in (2.23), used to describe the energy exchange in TRI in the context of a finite-width beam, as the zero-dimensional model \mathcal{M}_{0D} .

Using the same linear stability analysis as outlined above, under this new set of ODEs the growth rate term becomes

$$\sigma_{0D\pm} = \frac{1}{4}(\Sigma_1 + \Sigma_2) \pm \sqrt{\frac{1}{16}(\Sigma_1 - \Sigma_2)^2 + |I_1||I_2||\tilde{\Psi}_0|^2}, \quad (2.24)$$

where $\Sigma_p = \nu\kappa_p^2 + |\mathbf{c}_{g_p} \cdot \mathbf{e}_{k_0}|/W$. Here we define σ_{0D} to represent the linear theoretical growth rate for the zero-dimensional model \mathcal{M}_{0D} . The results of this linear stability analysis are thoroughly examined in Chapter 4, where they are compared against experimental results.

2.5.4 Summary

In this chapter we have outlined two weakly non-linear instabilities of internal waves and provided a review of the current theory used to describe them. In particular, we have focused on the theory behind TRI, as this instability forms the main focus of this research. The zero-dimensional ODEs in (2.19) show how published literature describes the energy exchange between the waves in the triad. The revised form of these ODEs, presented in (2.23), then summarises the work of Bourget et al. (2014), who examined the role of the finite-width beam on the interaction via an energy balance over a finite-size interaction region. We have defined these equations as the zero-dimensional model \mathcal{M}_{0D} . These equations provide the starting point for the analysis of the experimental results presented in Chapter 4. Before examining the experimental results of TRI, however, we first outline the experimental methods and diagnostic tools used to obtain them in Chapter 3.

Chapter 3

Experimental methods and diagnostics

3.1 Introduction

This chapter outlines both the experimental methods and diagnostic tools utilised in this thesis to examine internal gravity wave beams. All experiments were undertaken in an 11.40 m long tank (described in § 3.2), situated in the GK Batchelor Laboratory, Cambridge. The tank was filled with a salt stratification, which, for the majority of this project, was done using two gear pumps. The full procedure for filling the tank is detailed in § 3.4. The achieved density profile was then measured using an aspirating conductivity probe mounted to a traverse (outlined in § 3.5) that spanned the full depth of the stratification. Once the tank had been filled, experiments were conducted using a novel wavemaker that was located along a 1 m section of the tank floor. This flexible horizontal boundary generated sinusoidal forcing to produce internal wave beams. In general, these beams propagated upwards at an angle through the fluid, determined by the ratio of the forcing frequency to the buoyancy frequency (1.13). As the beams reached the free surface they were reflected, whereupon their vertical wavenumber component changed sign and they began to propagate downwards. In this fashion they travelled the length of the tank while being continuously damped by viscosity. Full details of the wavemaker are given in § 3.3. For some experimental configurations, vortex rings were fired into the internal wave field using a ‘L’-shaped acyclic tube connected to two bike pumps mounted on a traverse. Details of the vortex ring generation are given in § 3.6.

As internal waves are difficult to discern with the naked eye, the non-intrusive optical technique of Synthetic Schlieren was employed to visualise the flow by determining the

gradient of the density perturbation. Details of Synthetic Schlieren are given in § 3.7 along with the specific apparatus used.

After the experiments were captured, various spatial and temporal diagnostic tools were employed to post-process the results. Using the gradient of the density perturbation, generated from the Synthetic Schlieren, a temporal filtering technique known as Dynamic Mode Decomposition (DMD), was used to decompose the flow into its dominant, non-decaying frequencies. Filtering in Fourier space was then employed using the Hilbert Transform (HT), to isolate the directional components of interest. Details of DMD and the HT are given in § 3.9.2 and § 3.9.3 respectively. Finally, the determination of the wave amplitude, based on the magnitude of the streamfunction, is detailed in § 3.9.4.

3.2 Tank

All experiments were conducted in a 11.40 m long, 0.29 m wide, 0.48 m deep tank in the GK Batchelor Laboratory, a picture of which is shown in Figure 3.1. The tank base and supporting vertical members are constructed out of aluminium, while the walls are constructed out of 0.016 m thick Perspex (acrylic). For the remainder of the thesis, this tank will be referred to as the ‘main’ tank. The clear Perspex allows for visualisation through the tank walls. The tank – which was originally constructed in the late 1970s – was 10 m long, however, in 2015 the 1 m long Arbitrary Spectrum Wave Maker (ASWaM) (see § 3.3) was inserted approximately 2.5 m away from one end, thereby extending the length. Due to the location of the wavemaker, for some of the experiments conducted, internal waves were directed down the ‘long’ end of the tank where the end wall was approximately 7.5 m away. However some experiments also directed the waves along the ‘short’ end, where the end wall was 2.5 m away. This meant that the waves would travel the horizontal distance of either 15 m or 5 m depending on which way they were directed. The exponential viscous decay due to viscosity is given as the energy ratio

$$\frac{E_{\text{end}}}{E_{\text{start}}} = e^{\frac{-\nu\kappa^3 L}{N \sin \theta}} = e^{\frac{-\nu\kappa^2 L}{|cg|}}, \quad (3.1)$$

where L is the distance the wavebeam travels. This expression in (3.1) shows that viscosity causes an exponential decay at a rate of the wavenumber cubed in the direction of the waves group velocity. Using this expression, the wavebeam was found to decay to approximately 2% or 0.01 % of its initial energy before re-interacting with the

generation site. Both directions, therefore, provided sufficient time and distance over which the waves could viscously dissipate to negligible amplitudes.

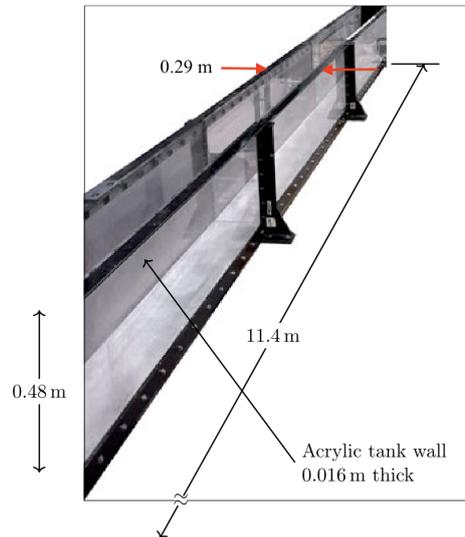


Figure 3.1: Photo of the main tank in the GK Batchelor Laboratory used for the experiments presented in this thesis, taken from Dobra (2018).

3.3 Arbitrary Spectrum Wave Maker

The history of internal wave experiments dates back to 1967, when Mowbray and Rarity (1967) used an oscillating cylinder centred in a stratified fluid to generate four internal wave beams, known as the ‘Saint Andrews Cross’ (a schematic of which was shown in Figure 1.2). In this experiment, the beams produced from the cylinder are one wavelength wide – determined by the diameter of the cylinder – and are comprised of a broad wavenumber spectrum. The oscillating cylinder became the iconic method for studying internal waves for a number of years. More recently, Gostiaux et al. (2007) developed a two-dimensional wavemaker using circular cam-shafts that sit along the side of a tank. For their set-up, the forcing amplitude is determined by the offset from the centre of the cam-shaft to the mounting location, and the wave frequency through the rotation speed. While this presents an original wave generation technique that allows for wave beam generation wider than one wavelength, there are also some drawbacks. Primarily, the forcing parameters can not be altered without the mechanical re-configuration of the drive shafts, meaning this can not be done during the course of an experiment. Another drawback of the wavemaker developed by Gostiaux et al.

(2007), is that it creates a staircase of flat surfaces due to plates of the cam shafts, which introduces additional unwanted signals to the flow. Both of these limitations have been overcome with the new generation of wavemaker presented here.

For all the experiments conducted, internal waves were generated by the Arbitrary Spectrum Wavemaker (ASWaM), which was designed and built in house (Dobra, 2018; Dobra et al., 2019). The ASWaM, also known as the ‘magic carpet’, is a 1 m long, 0.28 m wide wavemaker, which sits along the base of the tank described in § 3.2. It is comprised of a series of 96 computer controlled linear actuators that sit below the tank, each connected to a vertical drive rod via a brass coupling. These vertical rods pass through ‘O-ring’ seals at the base of the tank and connect to 0.28 m long horizontal rods which span the full tank width. A sheet of neoprene foam spans the full length and width of the wavemaker, covering the horizontal rods. Each of these rods is enclosed in a ‘Velcro’ sleeve, which has the ability to rotate around the rod but is constrained to move with the same vertical direction. These sleeves provide the connection with the underside of the neoprene foam, interpolating between the rods to provide smooth, unidirectional forcing. A diagram showing some of these rods and how they connected to the neoprene is depicted in Figure 3.2.

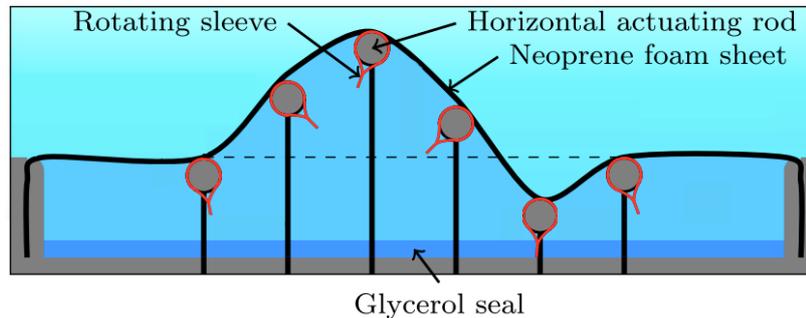


Figure 3.2: A diagram depicting some of the horizontal rods encased in their Velcro sleeve. These sleeves connect to the underside of the neoprene foam and are able to rotate around rod while still being constrained to the rods vertical motion. This diagram is reproduced from Dobra et al. (2019).

When submerged in a stratified fluid, the wavemaker can generate high temporal and spatial resolution, quasi-two dimensional, internal wave beams at amplitudes sufficient to permit wave breaking. Figure 3.3(a) shows a view of ASWaM from the underside of the tank, where the linear actuators connect to the vertical drive rods. Figure 3.3(b) shows ASWaMs positioning within the tank. In both of these figures the neoprene foam sheet covering the horizontal rods has been removed for clarity.

Due to the ASWaM being controlled electronically as opposed to mechanically, programming provides the ability to alter the forcing frequency ω_0 , amplitude η_0

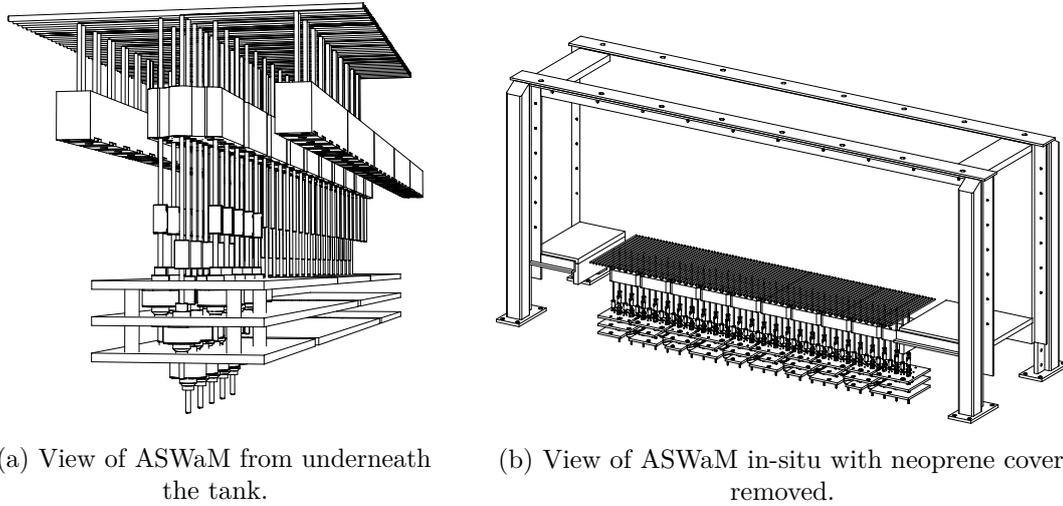


Figure 3.3: Computer Aided Design drawings of ASWaM. (a) A view from the underside, showing the linear actuators connected to the vertical drive rods via the couplings. The motors are staggered across three vertical planes to enable a tight rod separation of 10 mm. (b) shows a view of ASWaM in-situ with the horizontal support rods exposed. Figures reproduced from Dobra et al. (2019).

and wavelength λ_0 in both the temporal and spatial domain during the course of an experiment (the subscript 0 is used throughout to indicate the ‘primary’ or ‘forced’ wave from the wavemaker). This then allows for incremental changes in amplitude from a rest state, avoiding flow separation due to transient startups. The stepper motors are programmed through bespoke software that controls the timings in increments of 30 ns, which allows for fine spatio-temporal resolution. For full details refer to Dobra et al. (2019). Details of the exact waveforms used for the experiments are presented in § 4.2, § 7.4.2 and § 7.4.3.

As the wavemaker is situated along the bottom of the tank, salt water ingress was occurring around the seals of the drive rods and causing corrosion to the motors. To stop this ingress, six months into the PhD, glycerol was gravity fed into the cavity below the wavemaker before the stratification was filled. While the addition of this dense single-component viscous liquid prevented the salt water ingress around the drive rods, its presence provided a nutrient source for microbial growth in the fresher parts of the density stratification. To combat this growth, 0.75 L of sodium hypochlorite solution was added to the fresh water holding tank before filling (see § 3.4). While this prevented the microbial growth, it has accelerated the degradation of the tank, especially the paint work along the bottom and the neoprene sheet covering the wavemaker.



Figure 3.4: Photo of the Arbitrary Spectrum Wavemaker while running a large amplitude sine wave. Due to the large forcing the glycerol located in the cavity underneath has become entrained in the lower region of the stratification causing distortion to the dot pattern located behind the tank. The bearing blocks and motor housing can be seen underneath the tank.

3.4 Method of tank filling

For the first ten months of this research, the stratification within the tank was created using the traditional double bucket technique of Oster and Yamamoto (1963). Two holding tanks of identical volume $V \approx 0.34 \text{ m}^3$ and constant plan form were filled with either fresh water of $\rho_f \approx 0.99 \text{ g cm}^{-3}$ or concentrated salt water at $\rho_s \approx 1.15 \text{ g cm}^{-3}$. An outlet from the ‘fresh’ reservoir than ran to the main tank via a pump, passing through a layer of reticulated foam at the base of the main tank to dissipate momentum. When the filling process started, a valve at the base of the two holding tanks was opened, allowing salt water to slowly feed into the fresh water reservoir due to hydrostatic pressure differences. The fluid then being fed to the main tank from the freshwater reservoir steadily increased density over time, producing a linear stratification. Two main limitations were found with this filling method. Firstly, the holding tanks did not have a large enough combined volume to completely fill the main tank, necessitating additional salt water to be added during the course of the fill, which then impacted the density gradient. Moreover, due to the large density difference between the two reservoirs, the assumption required for the Boussinesq approximation – that the range of density across the depth was small – was no longer valid. The buoyancy frequency of the produced stratification increased by approximately 10 % with increasing depth (Dobra, 2018).

From October 2018 onwards, the tank was filled using two gear pumps, controlled via the DigiFlow software package (Dalziel, 2021). This fill method was implemented to allow for greater control over the stratification: pre-setting the buoyancy frequency

of the stratification (1.13), creating non-linear stratifications and allowing for greater fill depths compared to the double bucket method. Being able to pre-set the value of the buoyancy frequency made it easier to accurately reduce the strength of the stratification, giving greater validity to the Boussinesq approximation.

Two Coleparmer Ismatec BVP-Z Analog Gear Pump Drives and two magnetically driven Coleparmer Micropump L20562 A-Mount Suction Shoe Pump heads were used. Both pumps can be seen in Figure 3.5(a) while a close up of one pump head is shown in Figure 3.5(b). Several calibrations were conducted on the pumps with all combinations of ‘pump head to drive’ tested. It was found that all four combinations produced consistent but different flow rates. The combinations with the most similar flow rate was chosen and arbitrarily set as the ‘salt’ and ‘fresh’ water pumps. They were then re-calibrated *in-situ* – with the pumps being fed at the level of the holding tanks and filling the main tank approximately 1.5 m higher – to account for the change in pressure head due to the different levels. As the calibration was based on mass pumped, they were calibrated using fresh or salt water depending on pump, to ensure that the density of the water was accounted for in the calculated flow rate. The average of three calibrations for both pump configurations can be seen in Figure 3.6. It was found that below a voltage input of 1 V the pumps performed unreliably, with no flow being produced below 0.3 V. During the filling process a minimum voltage input of 1 V was therefore imposed on both pumps. The calibration curves shown in Figure 3.6 were then used to obtain the desired flow rate from each pump during the filling process.

Figure 3.7 provides a timeline for the full filling process. The first step was to fill the reservoirs three days prior to the main tank fill. This allowed for some of dissolved air in the water to be released and for the water to approach room temperature. Reducing the amount of dissolved air meant that fewer bubbles would form on the main tank walls, minimising the interference with the visualisation. On the day of filling, the first step was to fill the main tank with fresh water so that the wavemaker was just submerged. A large amplitude sinusoid was then run on the wavemaker to allow the air bubbles in the neoprene to be removed. After this was run for approximately an hour, the main tank and the cavity underneath the wavemaker, were drained and scraped clean. Using a gravity feed bucket, glycerol was added to the cavity until the level was just below the brass couplings. The glycerol acted to impede salt water leakage onto the motors by preventing evaporation which would lead to salt crystals forcing the seals around the ‘O-rings’ open.

During this time while the cavity was being filled, several other procedures were carried out. First, a piece of coarse reticulated foam was left to soak in fresh water.

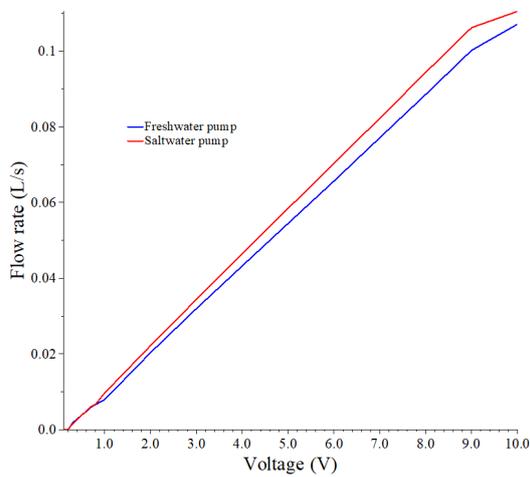


(a) A photo of both of the gear pumps *in-situ*, saltwater pump on the left, freshwater on the right.

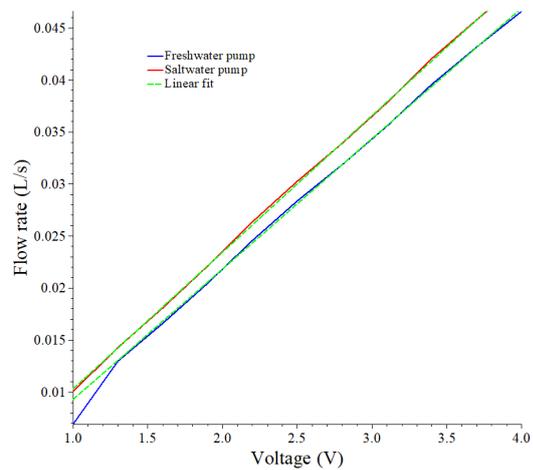


(b) The freshwater pump, showing the inlet and outlet tubes secured with jubilee clips.

Figure 3.5: Pictures of the Coleparmer gear pumps used to fill the stratification for the tank. Both pumps were controlled digitally, using the Digiflow software package (Dalziel, 2021).



(a) Gear pumps calibration across the full range of voltages.



(b) Gear pumps calibration across voltages required for filling the stratification.

Figure 3.6: Gear pump calibration curves. (a) The calibration curve for both pumps across the full range of voltages that the pumps were capable of. (b) High resolution calibration curves for the range of voltages required for filling the stratification.

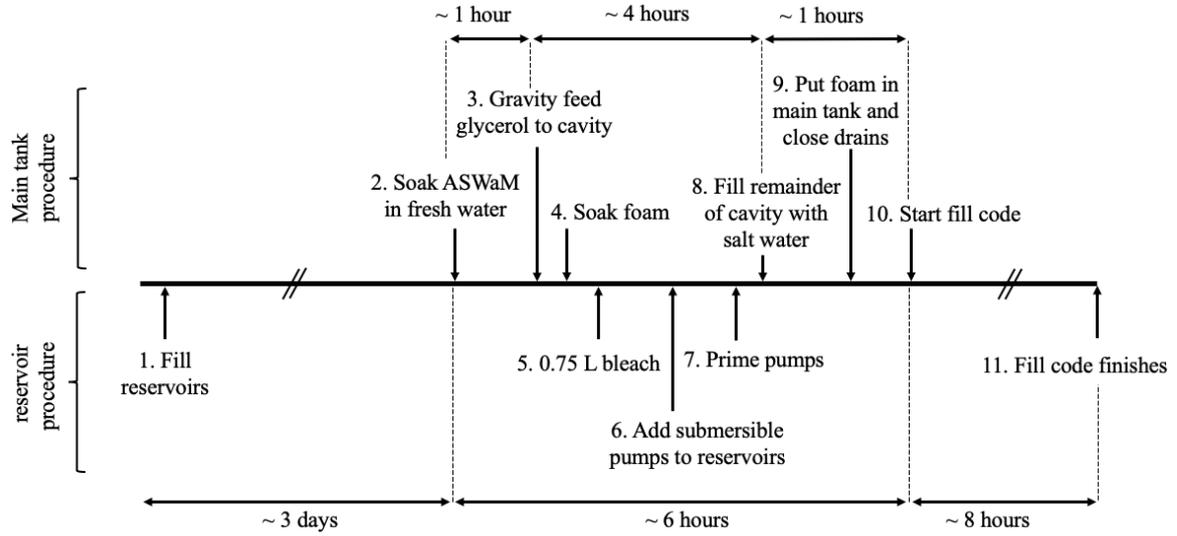


Figure 3.7: Timeline of procedure used to fill the stratification in the main tank.

This foam was later placed underneath the filling pipe that connected the reservoirs to the main tank via the pumps. During filling, it acted to dissipate the momentum of the fluid in order to minimise mixing in the tank. Attention then turned to the reservoirs. First, 0.75 L of sodium hypochlorite solution was added to the fresh water reservoir. As mentioned in § 3.3, this was added to combat microbial growth due the addition of glycerol in the cavity. A submersible pump was then placed into each reservoir and switched on. Care was taken to ensure the outlet of each submersible pump was directed downwards, to avoid unnecessary disturbance of the free surface which would entrain air. The next stage was to prime the gear pumps. The density at the top of each stratification ρ_t , was determined by the total fill rate (based off the time for filling, usually set at 8 hours), the depth of the stratification, normally 0.45 m, and the 1 V minimum voltage requirement for the pumps. For a buoyancy frequency of $N = 1.54 \text{ rad s}^{-1}$, $\rho_t \approx 1.027 \text{ g cm}^{-3}$. The pumps were therefore primed to this density.

Once the glycerol had finished draining from the gravity feed bucket, the remainder of the cavity was slowly filled with salt water until it reached the underside of the neoprene foam. Finally, the main tank was scraped cleaned and the drains were closed. After securing the reticulated foam underneath the outlet pipe, the fill code was started and the gear pumps began to draw water from the reservoirs, filling the tank with the programmed stratification. Due to the preparation time involved, this filling process was usually done over night, as the gear pumps were controlled via DigiFlow. In case

of an error in the filling code, a safety ‘kill’ switch was attached to the main tank that automatically stopped the pumps should the water level get too high.

As the full capacity of the reservoirs was not enough to fill the entire stratification, water was added incrementally during the filling process. The submersible pumps guaranteed that this additional water added to the holding tanks during the fill would be well mixed, ensuring that any changes in temperature or salinity in the reservoirs would be small. This additional water therefore had significantly less impact of the stratification compared to the one produced using the double bucket technique. The water level in the holding tanks was controlled by two ball-cock valves, as shown in Figure 3.8. As the gear pumps removed water, the level inside the holding tank would drop, moving the ball float down. This movement opened the valve, allowing more water to be added until the ball float was raised to a level that would close the valve. This method ensured that the water level in the reservoirs was maintained, while ensuring that additional water was not wasted from the overflow outlet.

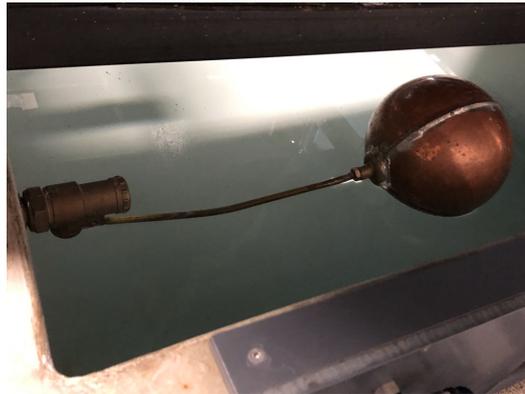


Figure 3.8: Photograph of the ballcock valve in the fresh water holding tank.

3.5 Conductivity probe

In order to measure the achieved density stratification, an aspirating conductivity probe was mounted on a linear traverse that was secured over the tank. The traverse, driven by a stepper motor, moved the probe down through the full depth of the tank at 10 mm s^{-1} , while the probe sampled the passing fluid at 2000 Hz. As the conductivity of a fluid is a function of both temperature and the concentration of salt ions, the voltage output from the probe was used to determine the density stratification.

The probe is comprised of two concentric cylindrical electrodes, electrically insulated from each other except via the fluid that passes through the 0.3 mm diameter opening

in the insulating Delrin tip, a picture of which can be seen in Figure 3.9. When the tip is fully submerged in water, the fluid passing through the narrow hole completes the circuit between the outside of the outer electrode to the inside of the inner one. Due to the small diameter of the hole, the resistivity inside the probe tip is much larger, and the conductivity signal measured by the electrodes is dominated by the fluid passing through the tip. In order to avoid salt crystallisation or the probe tip drying out, the tip was kept submerged in fresh water when not in use.



Figure 3.9: Photograph of the white Delrin tip of the conductivity probe and the outer metal electrode.

The conductivity probe was calibrated using samples of different densities created by mixing fresh and salt water at laboratory temperature. The density of these samples were measured using an Anton Paar DMA 5000 density meter with an accuracy of $\pm 1 \times 10^{-6} \text{ g cm}^{-3}$. The fresh and salt water samples were taken from the tap and left overnight to allow them to reach room temperature. An example calibration curve used for the experiments is shown in Figure 3.10. The density of the stratifications typically ranged between $1.025 \sim 1.140 \text{ g cm}^{-3}$, so the probe was calibrated across the full range of densities with a higher concentration of samples falling in the stratification region. The calibration curve used a least squares fit of a quadratic curve. Generally, as long as the probe was not allowed to dry out, the calibration curve would remain stable over a long period of time.

There were, however, several issues that arose with the conductivity probe, the first of which (for this research project) occurred in November 2018. A new linear traverse was made for the probe to allow for smoother traversing and to stop experiments being delayed due to the traverse being shared by other laboratory users. After some experiments using the new traverse it was clear that the density profiles obtained from the conductivity probe were not accurately representing the stratification. Further

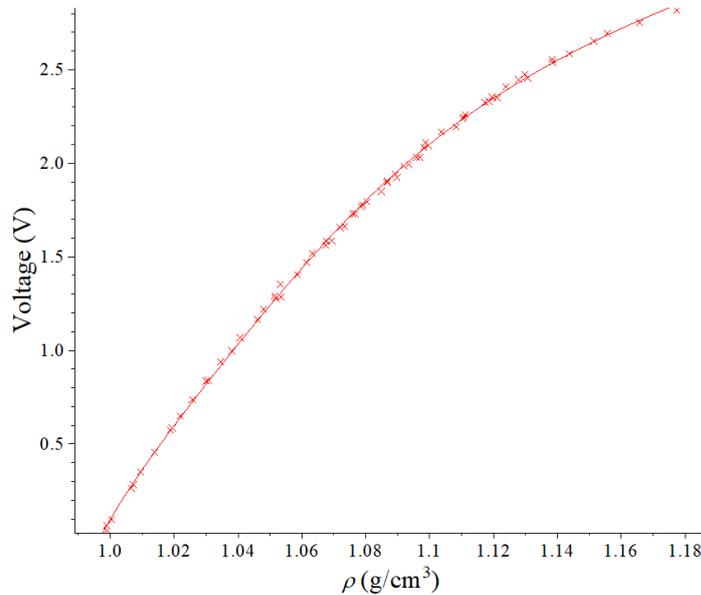


Figure 3.10: Calibration curve for the aspirating conductivity probe. The curve fit was determined using least squares.

diagnostics revealed that there was a significant level of electrical interference in the conductivity measurement coming from the traverse motor controller, regardless of whether the traverse was moving or not. Figure 3.11(a) shows the difference in the conductivity reading from the probe between having the controller box on (blue line) and off (red line). In addition to a significantly greater level of noise, the mean voltage reading was approximately 0.004 V higher with the motor controller box on.

While it was clear that there was a large electrical interference coming from the traverse motor controller box, as this box powered the traverse, it needed to be on during experiments to allow the probe to traverse the density stratification. Note that the interference was not related to the speed of the traverse, simply whether the power was turned on, as it was communicated along the wires from the computer used to set the traverse in motion. Although calibrating with the controller turned on improved the signal through low-pass median filtering, the accuracy remained unacceptable.

A new traverse controller box was therefore commissioned and was received at the end of September 2020, with the delay being primarily due to the shut down of the laboratory between March and August 2020 because of to the COVID-19 pandemic. The signals obtained from the probe with the new controller box are given in Figure 3.11(b), which clearly shows that the mean value of the voltage reading was no longer altered when the new traverse controller was switched on.

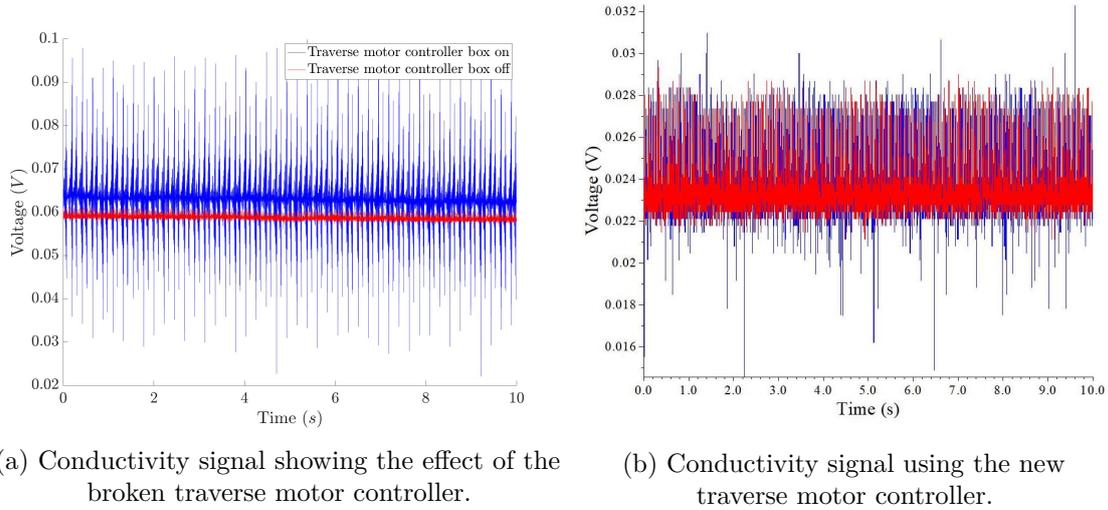


Figure 3.11: The voltage reading from the conductivity probe over 10 seconds. (a) Shows the effect of the electrical interference from the broken traverse motor on the voltage signal of the probe. (b) The probe signal over the same period after the traverse controller was replaced.

The other component that had a significant impact of the conductivity readings was the temperature variations of the water in the stratification. As discussed in § 3.4, the reservoirs were filled approximately three days prior to filling the main tank to allow the water to warm towards room temperature. However, measurements taken in December 2020 showed that due to evaporative cooling, the temperature of the water in the reservoirs equilibrated at approximately 1.5 °C below the laboratory temperature of ~ 20.6 °C. The water used to calibrate the probe experienced the same convective cooling as the water in the reservoirs, meaning it also remained consistently around 1.5 °C lower than the laboratory temperature. Once the main tank was filled, the stratification suppressed the convective heat loss at the free surface. Concomitantly, heat flux from the laboratory gradually warmed the body of the stratification through the tank walls and it asymptotically approached laboratory temperature. The problem with this small increase in temperature was not that it had an impact on the actual density of the stratification, but rather it had an impact on the conductivity.

At constant pressure around room temperature, the equation of state for a water can be approximated by the linear equation

$$\varrho(S, T) \approx \varrho_0(1 - \alpha(T - T_0) + \beta(S - S_0)), \quad (3.2)$$

where S is salinity, T is temperature and α and β are constants. Conductivity, which is used as a measure of the density, is also a function of salinity and temperature,

$C = C(S, T)$. For the temperature changes measured here (± 1.5 °C), it is reasonable to approximate $\varrho \approx \varrho(S)$. However, the same approximation can not be made for the conductivity. Increases in temperature will cause an increase in the conductivity at constant salinity, due to the higher mobility of the salt ions. This effect will be felt more strongly at higher salinities due to the increased concentration of ions. Temperature will have negligible affect on the conductivity of fresh water, as the conductivity is close to zero. To a good approximation, the increase in conductivity with temperature is linear and can be described by

$$C_T = C_{\text{cal}}(1 + \alpha_c(T - T_{\text{cal}})), \quad (3.3)$$

where C_T is the conductivity at a given temperature T , C_{cal} is the conductivity at the calibration temperature T_{cal} and α_c is the fractional increase in conductivity per degree centigrade, which for salt is given as $\sim 2.3\% / \text{°C}$ (Down and Lehr, 2005). Figure 3.12 shows the relationship in (3.3) applied to each point on the calibration curve shown in Figure 3.10. This shows how the effects of temperature on the conductivity are heightened with increased density.

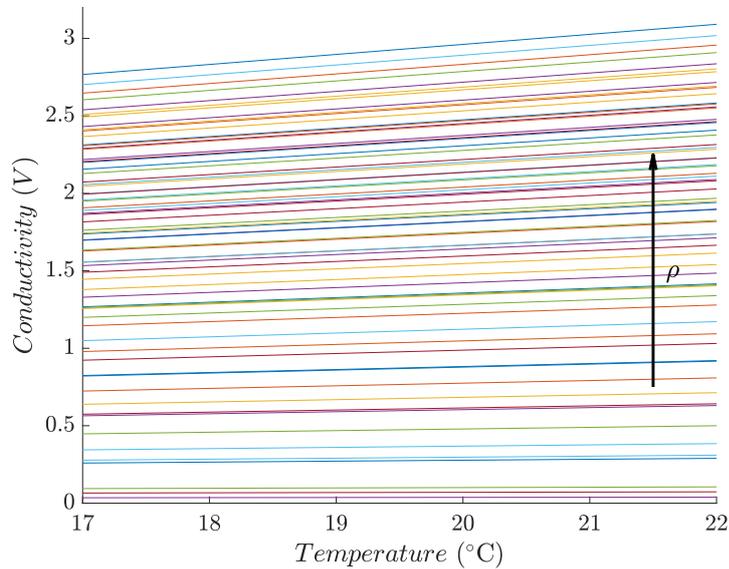


Figure 3.12: The effect of the relationship given in (3.3) on each point in the probe calibration curve given in Figure 3.10 with a calibration temperature C_T of 17.5 °C.

The result of this temperature dependence, means that when the temperature of the stratified water increases in the tank, the saltier, denser water at the bottom will show a greater increase in the conductivity measurement, despite the actual density

difference being negligible. This is shown in Figure 3.13, where two profiles from the same stratification are plotted. The profile in black (with the red linear fit), was taken when the tank was first filled and the water was cooler. The profile in blue (with the green linear fit) was taken approximately 70 hours after, by which point the water column had warmed and equilibrated with the laboratory temperature. The calibration for these profiles was conducted using water that experienced the same convective cooling as the water in the reservoirs, meaning it was around 1.5 °C lower than the laboratory temperature. The earlier black profile, therefore, most accurately represents the stratification, as its temperature was most similar to the temperature that the probe was calibrated at.

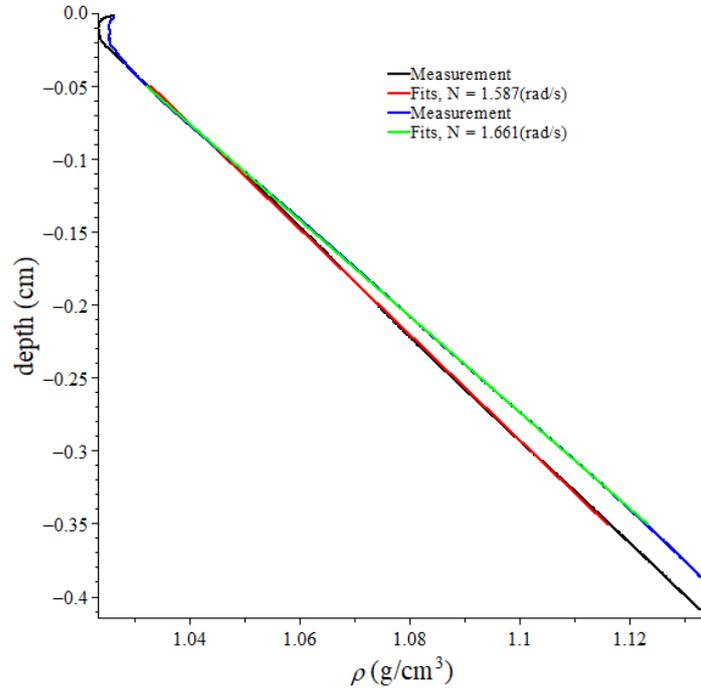


Figure 3.13: Two conductivity profiles taken in June 2019. The black profile was taken soon after the tank was filled, while the blue profile was taken approximately 70 hours later, after which the water temperature had equilibrated to that of the laboratory.

To understand the temperature trend that was affecting the conductivity profiles, the buoyancy frequency, N , was calculated for all the profiles taken over the course of an experimental set (usually four to five days). As given in (1.13), the buoyancy frequency was calculated by

$$N = \sqrt{-\frac{g}{\rho_0} \frac{d\bar{\rho}}{dz}}, \quad (3.4)$$

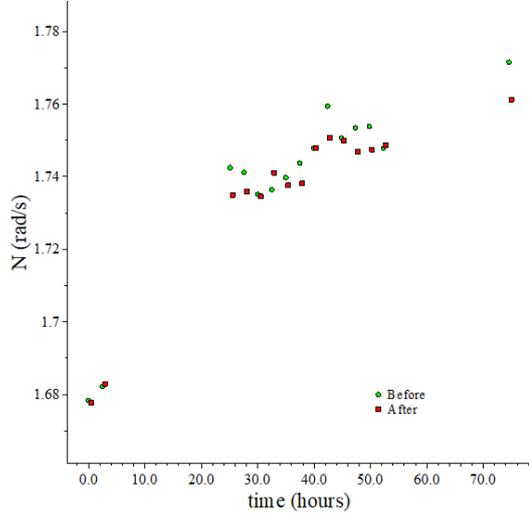
where here the reference density ϱ_0 was the mean density value from the profile and $d\bar{\rho}/dz$ was the density gradient shown by the red and green fits on Figure 3.13. The results for four sets of experiments, conducted over two years, can be seen in Figure 3.14. The two profiles shown in Figure 3.13 are marked on Figure 3.14(c). The black square around the measurement indicates the black profile on Figure 3.13 and the blue square indicates the blue profile. All four graphs in Figure 3.14 show the measured value of the buoyancy frequency increasing over time until they reach an asymptotic value, showing the stratification has equilibrated with the laboratory temperature.

As the water was calibrated at approximately the same temperature of the water in the holding tanks, the first profile taken after the stratification was filled, is the most similar to actual density profile. These temperature dependencies were not recognised until December 2020, due both to the laboratory closure and also the slight mistrust towards the probe readings because of the problems mentioned above. At this point, it was not worth the considerable time that would be required to include a temperature calibration.

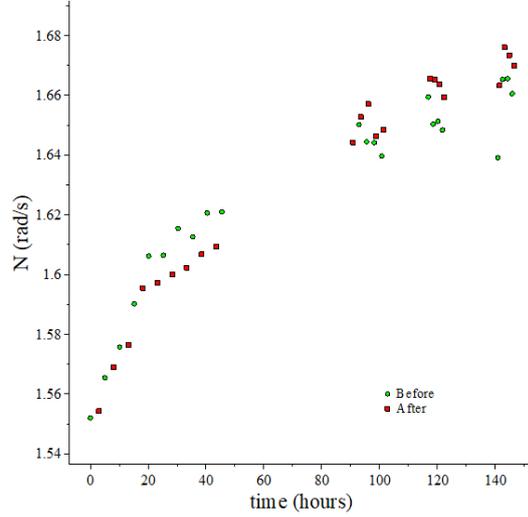
Using the experimental results, a value of N was also obtained from the linear dispersion relation (1.20), based on the angle that the lines of constant phase made with the vertical and the known input forcing frequency. This value was then compared to the values obtained for the buoyancy frequency from the density profiles. As expected, the first profile of each experimental run was always the most similar to the value obtained from the linear dispersion. The difference in these two measurement techniques gave an error to the buoyancy frequency of approximately $\pm 4\%$.

3.6 Vortex ring generation

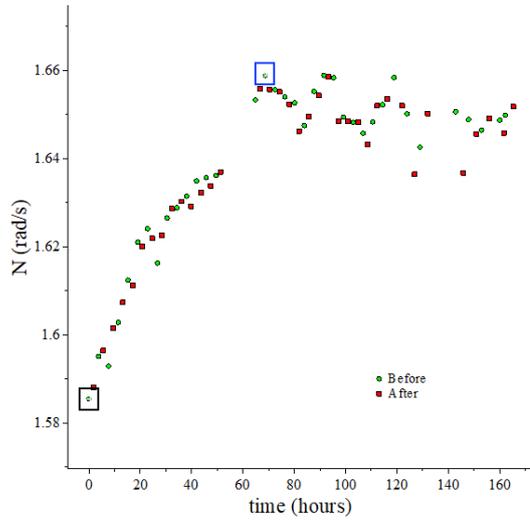
Some of the experiments presented in Chapter 7 examine the interaction between an imposed wave field and a freely evolving turbulent-like structure. Specifically, a vortex ring was taken as the simplified prototype for this turbulent structure and the method used for their generation is described here. Before filling the tank, an ‘L’-shaped tube of constant 0.04 m diameter was inserted and clamped parallel to the length of the tank using a horizontal aluminium profiled bar (Bosch Profile). A schematic of the front view (as seen by the camera) of the full set-up with the vortex-ring tube is shown in Figure 3.15. Care was taken to ensure the tube was centred across the width of the tank. As the stratification in the tank was filled, the tube became partially submerged. The end of the tube exposed to air was then connected via a rubber bung to a plastic hose, which split at the end to form a connection to two bike pumps. The pumps were



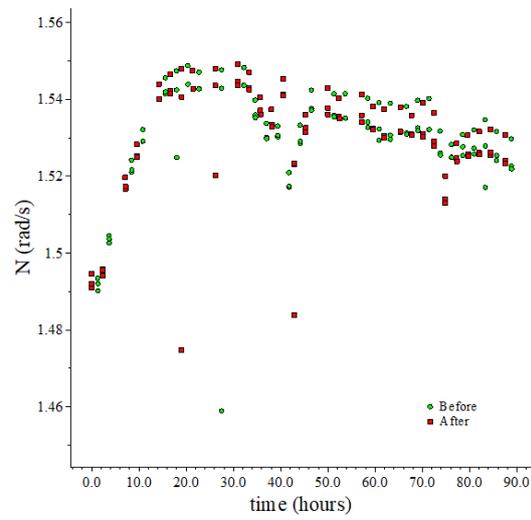
(a) Experiments in January 2019.



(b) Experiments in April 2019.



(c) Experiments in June 2019.



(d) Experiments in December 2020.

Figure 3.14: The obtained values for the buoyancy frequency, N , from (3.4) as a function of time. Each graph corresponds to one set of experiments using the same stratification. The green circle of ‘Before’ indicates the profile(s) taken before a single experiment, while the red ‘After’ indicates the profile(s) taken 30 s after the experimental recording finished. The black and blue squares in (c) indicate the two profiles in Figure 3.13.

mounted on to a linear traverse, which controlled the vertical motion of the pumps handle, meaning air could be forced through the hose and into the outlet tube at a controlled rate. Both bike pumps had a 0.044 m barrel and are shown in Figure 3.16(a). The ‘L’-shaped tube along with its Bosch Profile are shown in Figure 3.16(b) and the rubber bung used to connect the hose to the tube is shown in Figure 3.16(c),

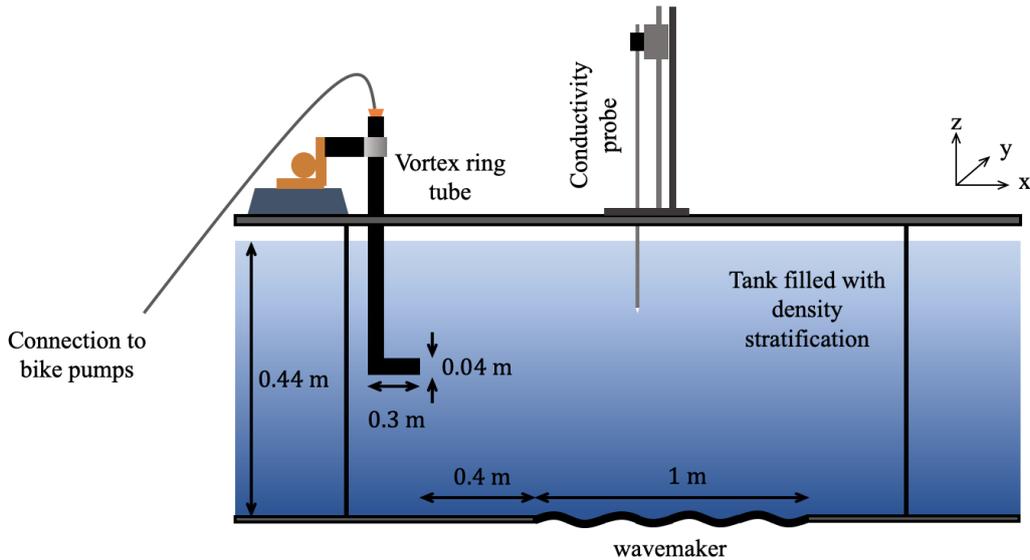


Figure 3.15: Schematic showing the view from the camera looking towards the front of the tank. For experiments including the vortex-rings (discussed in Chapter 7), the black ‘L’-shaped tube was mounted as shown before the tank was filled with a linear stratification. Further pictures of the vortex-ring apparatus are shown in Figure 3.16.

To fire a vortex ring, code from DigiFlow was used to send a hyperbolic tangent profile signal to the traverse motor. The code had the ability to control both the rate of the traverse and the distance over which it travelled. The hyperbolic tangent profile allowed for smooth acceleration of the pumps from rest to high velocities. This motion caused air from both pumps to be ejected through the hose and into the vortex ring tube. Consequently, the water in the tube was forced outwards, rolling up in the process into a vortex ring, which propagated horizontally out of the tube and into the stratification. Approximately 20 seconds after firing the ring, the traverse would slowly move the handles of the bike pumps up again to allow for the motion to be repeated. Details of the properties of the ring generated from this process are presented in Chapter 7.

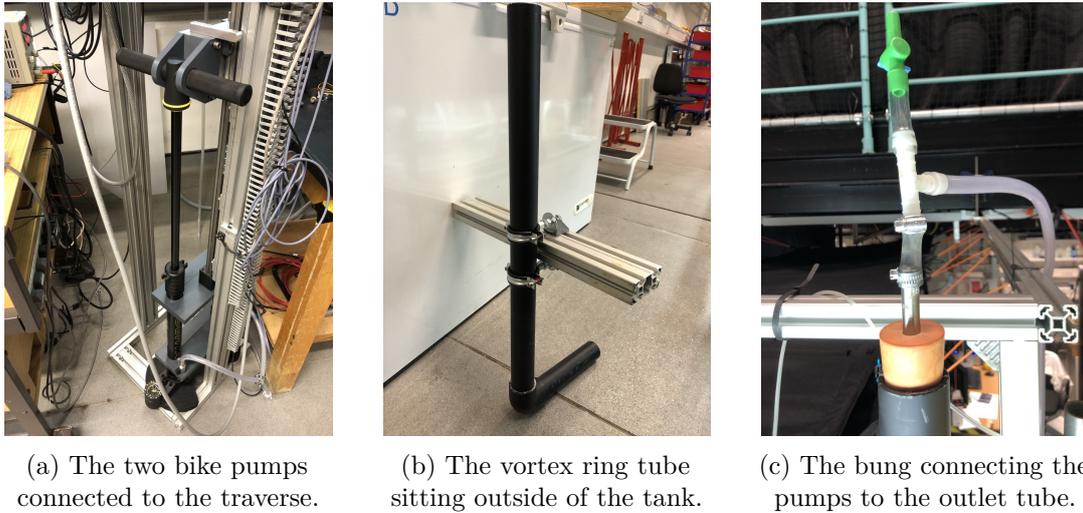


Figure 3.16: Pictures of the vortex ring generator traverse and tube. (a) A picture of the vortex ring generator traverse. Two bike pumps are attached to a linear traverse, mounted on a Bosch profile frame. (b) A photo of the outlet tube which, during experiments, sits partially submerged in the tank. The Bosch profile attachment spans the tank and holds the outlet tube in place when the vortex ring is fired. (c) A connection between the hose running from the bike pumps to the outlet tube shown in (b). The orange bung provides an airtight seal around the connection point ensuring volume conservation.

3.7 Synthetic Schlieren

3.7.1 Overview

As internal waves are difficult to visualise with the naked eye, the non-intrusive technique of Synthetic Schlieren was employed in order to observe and analyse the experiments. Synthetic Schlieren was first proposed by Dalziel et al. (1998) and subsequently developed by Sutherland et al. (1999), Dalziel et al. (2000) and Dalziel et al. (2007). The technique is based on the optical principle that light rays are bent towards regions of higher refractive index, which, in a stratified fluid, correspond to regions of higher density. In the same way that a mirage is generated by heat rising from a hot road, internal waves locally perturb the density stratification, deflecting the light rays as they pass through the tank and distorting the field of view.

To measure this in the laboratory, a dot pattern is illuminated behind the tank and a reference image is captured from the front of the tank, showing the pattern distorted by the quiescent stratification. Localised perturbations from internal wave motion then cause an apparent shift in the light rays, the amplitude of which is measured by the difference of the displaced dot pattern away from the reference image. Using this

displacement amplitude, the gradient of the density perturbation is calculated, which is then used to determine the amplitude of internal wave motion. Further details of the theory behind Synthetic Schlieren is addressed are § 3.7.2.

Unlike the classical Schlieren, first used for internal waves by Mowbray and Rarity (1967), the Synthetic Schlieren uses simple laboratory equipment without the need for large, parabolic mirrors. Figure 3.17 shows the setup used for the experiments presented here. Its sensitivity is such that heat rising from a hand can be detected, especially when located close to the camera. Details on the specific apparatus used in the laboratory, along with the method used to isolate the camera from thermal noise, are given in § 3.7.3.

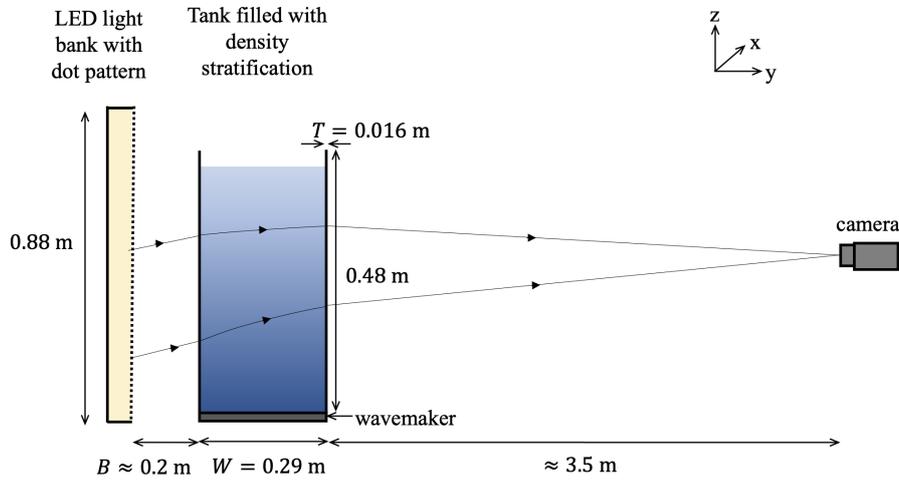


Figure 3.17: Side view schematic of the optical arrangement for Synthetic Schlieren. An LED light bank is positioned behind the tank where a random dot pattern is attached. The light rays generated from the light bank refract along a curve through the density stratification, which, for an unperturbed linear stratification, can be modelled by a quadratic. Internal waves generated from the wavemaker then locally distort the stratification and the camera records the apparent displacement of the dots attached to the light bank. A DigiFlow algorithm is then used to determine the gradient of this density perturbation.

3.7.2 Ray tracing

Synthetic Schlieren is a non-intrusive measurement technique based on optical laws. To understand the theory behind the technique, we start by using Fermat's variational principle, as outlined in Dalziel et al. (2000), which describes the behaviour of light in

an inhomogeneous medium as

$$\delta \int n(x, y, z) ds = 0, \quad (3.5)$$

where n describes the refractive index field and s is the direction orientated along the light ray. The co-ordinate system (x, y, z) , defines x as the component along the length of the tank, y as the component across the width and z as the vertical component, as shown in Figure 3.17. As we are interested in the variation of a light ray that crosses the width of the tank, W , we define $x = \xi(y)$ and $z = \zeta(y)$, where $d\xi/dy$ and $d\zeta/dy$ can be used to describe respectively the horizontal and vertical components of the slope of the light ray at any point across the width. From the variational principle (3.5), the two coupled differential equations are derived

$$\frac{d^2\xi}{dy^2} = \left[1 + \left(\frac{d\xi}{dy} \right)^2 + \left(\frac{d\zeta}{dy} \right)^2 \right] \frac{1}{n} \frac{\partial n}{\partial x}, \quad (3.6a)$$

$$\frac{d^2\zeta}{dy^2} = \left[1 + \left(\frac{d\xi}{dy} \right)^2 + \left(\frac{d\zeta}{dy} \right)^2 \right] \frac{1}{n} \frac{\partial n}{\partial z}, \quad (3.6b)$$

which define the displacements of the light ray due to the gradients of the refractive index field. Therefore, for any ray falling on a particular pixel, (3.6) can be used to back trace where on the original pattern it came from. For the set up used, the camera was located far from the tank to reduce parallax error (Thomas et al., 2009), meaning that the light rays passing through the tank to the camera are approximately horizontal. In addition, we assume the perturbations to the stratification are quasi two-dimensional and the deflections to the light ray are small ($\ll 1$). Under these assumptions, the terms $(d\xi/dy)^2$ and $(d\zeta/dy)^2$ can be considered negligible, allowing (3.6a) and (3.6b) to be decoupled as

$$\frac{d^2\xi}{dy^2} \approx \frac{1}{n} \frac{\partial n}{\partial x}, \quad (3.7a)$$

$$\frac{d^2\zeta}{dy^2} \approx \frac{1}{n} \frac{\partial n}{\partial z}. \quad (3.7b)$$

Integrating (3.7) twice leads to an expression for the path of the light ray, which takes the form of a quadratic

$$\xi = \xi_i + y \tan \phi_\xi + \frac{1}{2} y^2 \frac{1}{n} \frac{\partial n}{\partial x}, \quad (3.8a)$$

$$\zeta = \zeta_i + y \tan \phi_\zeta + \frac{1}{2} y^2 \frac{1}{n} \frac{\partial n}{\partial z}, \quad (3.8b)$$

where the incident co-ordinates of the ray are (ξ_i, ζ_i) and the horizontal and vertical components of angle of the ray as it passes into the tank are given by $\tan \phi_\xi = d\xi/dy$ and $\tan \phi_\zeta = d\zeta/dy$, respectively. We then decompose the refractive index field $n = n_0 + n'$, where n_0 is the nominal refractive index of the field ($n_0 = 1.3332$ for fresh water) and n' is the perturbation to the field, where $n' \ll n_0$. Developments in Dalziel et al. (2007) also highlighted that, in the laboratory setting, there are strong refractive index contrasts between the fluid in the tank n , the tank walls n_t and the surrounding air n_a , which cannot be neglected. Using Snell's law, they showed that the angle of the light changes by a factor of n_0/n_t at the interface with the tank wall and n_0/n_a at the interface with the air. Knowing this, the apparent displacement of the light rays from the source to the camera in both the horizontal and vertical direction $(\Delta\xi, \Delta\zeta)$ is given by

$$\begin{aligned} \Delta\xi &= -\frac{1}{2} W \left(W + 2 \frac{n_0}{n_a} B + 2 \frac{n_0}{n_t} T \right) \frac{1}{n_0} \frac{\partial n'}{\partial x}, \\ \Delta\zeta &= -\frac{1}{2} W \left(W + 2 \frac{n_0}{n_a} B + 2 \frac{n_0}{n_t} T \right) \frac{1}{n_0} \frac{\partial n'}{\partial z}, \end{aligned} \quad (3.9)$$

where B is the distance of the dot pattern behind the tank, W is the width of the tank and T is the thickness of the tank walls, as indicated on Figure 3.17.

It is then necessary to understand the relationship between the refractive index field and density. To a good approximation, the refractive index can be given solely as a linear function of the density (e.g. Haynes et al., 2014) by

$$\nabla n = \beta \frac{n_0}{\varrho_0} \nabla \rho, \quad (3.10)$$

where ϱ_0 is the reference density and the dimensionless constant $\beta \approx 0.184$ (Dalziel et al., 2000). Using this relationship, the measured displacements of the light rays in the horizontal and vertical directions given in (3.9) can be used to express the gradient

of the density perturbation in two-dimensions as

$$\frac{\partial \rho}{\partial x} = -2\varrho_0 \Delta \xi \left(\beta W \left(W + 2 \frac{n_0}{n_a} B + 2 \frac{n_0}{n_t} T \right) \right)^{-1}, \quad (3.11a)$$

$$\frac{\partial \rho}{\partial z} = -2\varrho_0 \Delta \zeta \left(\beta W \left(W + 2 \frac{n_0}{n_a} B + 2 \frac{n_0}{n_t} T \right) \right)^{-1}. \quad (3.11b)$$

All of the inverse ray-tracing algorithms used for the Synthetic Schlieren calculations have been performed using an in-built algorithm in DigiFlow (Dalziel et al., 2007).

3.7.3 Apparatus

A 12 MPixel ISVI IC-X12CXP camera with a Nikkor 35-135 mm zoom lens was used to capture the experiments, at a rate between 1 to 4 frames per second, depending on the overall recording length. The camera was located approximately 3.5 m away from the test section to minimise parallax, as discussed in § 3.7.2. The camera was connected to the computer and all recording was controlled via DigiFlow.

For the first year of the PhD, the dot pattern used for the Synthetic Schlieren was generated by a Dynamic Pattern Generator (DPG), which was a large LCD monitor (a 55" Samsung 4K LCD TV). This presented the advantage of the light source and the mask (dot pattern) being combined into one device. It was noted in 2018, however, that the similarity in resolution of the DPG (8.3 MPixels) and the ISVI camera (12 MPixels) was generating interference patterns in the recorded experiments, so a purpose made LED light bank was used instead. The LED light source, constructed in 2019, is 2.08 m wide and 0.88 m high, allowing for full coverage of the viewing window containing ASWaM. The dot pattern was printed onto overhead transparencies using a laser printer and was then attached to the front of the LED panel. A random black and white dot pattern was used as opposed to a uniform repeating pattern, as aliasing issues arise from repeating patterns if the apparent displacements from the density perturbations are greater than the spacing between the dots (Dalziel et al., 1998).

As Synthetic Schlieren converts all refractive index changes within the field of view into density perturbations, refractive index variations in the surrounding air are incorrectly ascribed to density variations in the fluid. These additional refractive index fluctuations can be generated due to external sources of thermal flux causing density changes within the surrounding air. This ‘thermal noise’ will be particularly apparent when close to the camera. For the experiments presented here, the space between the tank and the camera was a walkway in the laboratory, meaning it was not possible to

remove the thermal fluctuations generated from people walking around the experiment. To help mitigate the effects of this noise, a ‘thermal tunnel’ was constructed, that ran from the camera lens to the perimeter of the visualisation region on the tank, shielding the field of view. The tunnel was constructed using Cordura, a structural nylon with a polyurethane coating on one side to make it impermeable. Pictures of the thermal tunnel are shown in Figure 3.18.

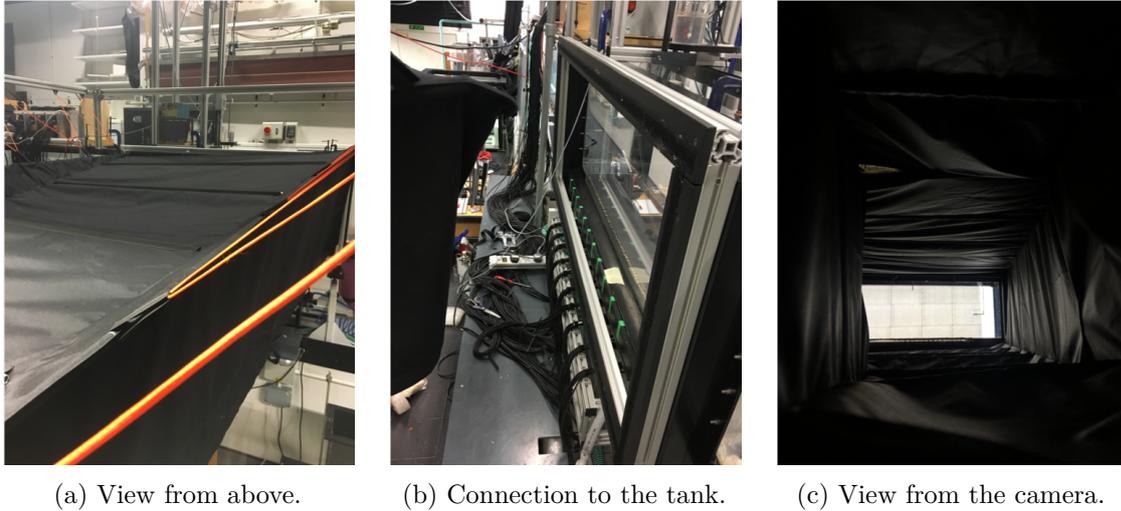


Figure 3.18: Pictures of the thermal tunnel. (a) A view from above looking towards the tank, showing the exterior of the tunnel suspended from its frame. (b) A view of the frame around the field of view, which was attached onto the tank. This frame was covered in ‘Velcro’, which connected to the matching ‘Velcro’ on the tunnel, allowing a seal to form around the visualisation window. (c) A view from inside the tunnel where the camera was located, looking towards the tank.

To hang the tunnel a frame was constructed out of Bosch Profile. A system of pulleys and rope was used to suspend the tunnel and allow for adjustments to align the view from the camera with the tank. To hold its structure, metal rods were inserted along inserts constructed on the exterior of the tunnel, as can be seen in Figure 3.18(a).

3.8 Automatic Experiments

The other way in which the effects of thermal noise were mitigated, was the transition to automation for all of the experimental runs. This was achieved by the start of 2019, meaning all experiments were conducted over night or during the weekends when thermal fluctuations in the laboratory were minimal. Using DigiFlow, a code was created that synchronised the timings of the camera, wavemaker, conductivity probe

and in some cases the vortex ring traverse. Various input parameters were specified, such as the time to record each experiment for, the interval between each experiment, the number of conductivity profiles to take and which experimental waveform to generate from the wavemaker.

Before the automation, the only manual piece of equipment used during the experiments was the pinch valve, which, when opened, allows for water to flow through the tube connected to the conductivity probe. This manual pinch valve was therefore replaced with an automated pinch valve, a picture of which can be seen in Figure 3.19. When the traverse was moving the probe through the stratification, the barrel of the automatic pinch valve would open, allowing water to flow through the tube. As the traverse stopped, the barrel on the pinch valve would close, ceasing the flow of water.

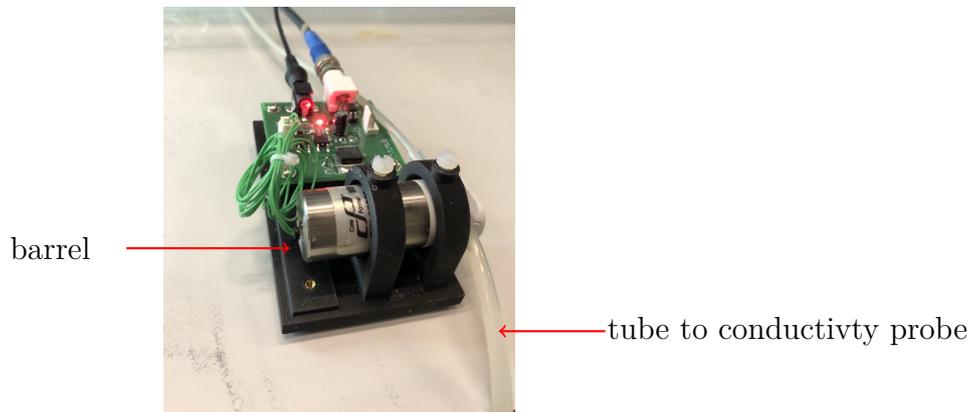


Figure 3.19: Picture of the automatic pinch valve used for the automation of the experiments. When the traverse was moving, the barrel would open, allowing flow through the tube which was connected to the top of the conductivity probe. When the traverse stopped moving, the barrel would pinch close, preventing flow through the tube.

While the automation was originally developed to help minimise thermal noise for the Synthetic Schlieren, there were various other advantages gained from this process. Firstly, timings between experiments could be kept consistent, allowing for a set length of time for all wave disturbances in the tank to dissipate. Running experiments over night or on weekends, meant that external sources of disturbances to the flow were minimal, reducing the chance of unwanted interference. Moreover, it allowed for other activities to be pursued while experiments were running, maximising time efficiency.

3.9 Image Filtering

3.9.1 Overview

The output from the Synthetic Schlieren provides the gradient of the density perturbation of the flow field (see § 3.7). After obtaining this from the experiments, it was necessary to decompose the flow further, into the discrete spatial and temporal components corresponding to separate wave beams. Owing to the unique dispersion relationship of internal waves (1.20) – that links the wavenumber and frequency with the direction of propagation – it is possible to use spatial and temporal filtering techniques to isolate individual waves. The first technique used, outlined in § 3.9.2, was Dynamic Mode Decomposition (DMD). This is a temporal technique that extracts the dominant frequencies and corresponding modal structures of the flow with no *a priori* knowledge of the system. Here we applied DMD to the gradient of the density perturbation fields determined by Synthetic Schlieren. As the dispersion relationship permits four separate directions of propagation for a single frequency, directional filtering in wavenumber space was then required after the DMD analysis. This was done using the Hilbert Transform (HT). Outlined in § 3.9.3, the HT uses Fourier transforms to isolate the desired direction of propagation in two-dimensions, thereby removing unwanted components of the flow associated with a particular frequency. The final processing stage was used to extract the amplitude of these disparate waves at each time step. This process is described in § 3.9.4.

3.9.2 Dynamic Mode Decomposition

Dynamic Mode Decomposition (DMD) was used as the temporal filtering tool for separating out the disparate frequencies of the wave field. This technique provides a convenient method for identifying the dominant frequencies from the Synthetic Schlieren output, without any *a priori* knowledge of the system (Schmid, 2010). It uses an eigen-decomposition of a linearised representation of the underlying temporal operator to establish the ‘modes’ of recurrent structures within an internal wave field. Where DMD excels, is in determining the frequencies and structure of the modes from short time series where there is a discrete spectrum that can reasonably be approximated by a combination of delta functions at slowly evolving frequencies. Fast Fourier Transforms (FFTs) on the other hand, approximate the dominant signal using a polynomial. Over these short sequences the FFT suffers, as, unless the sequence length corresponds to an integer multiple of the signal period, the approximated peak of the signal will be

broader and will not capture the dominant modes so accurately. For this reason DMD was chosen to analyse the slow evolution of the internal wave field becoming unstable to TRI. For this analysis, the DigiFlow macro was used (Dalziel, 2021), which is outlined in the Dynamic Mode Decomposition Algorithm section. Applications of the algorithm, specific to the internal wave experiments, are given in the following section.

Dynamic Mode Decomposition Algorithm

We start with a state vector \mathbf{x}_n which describes the n^{th} frame in a video sequence, where every element in the vector corresponds to a pixel in the frame. A sequence of N discrete frames in an experiment can then be represented by the matrix \mathbf{X}_0 of state vectors

$$\underbrace{\mathbf{X}_0}_{p \times n} = \left[\underbrace{\mathbf{x}_0}_{p \times 1}, \mathbf{x}_1, \dots, \mathbf{x}_{N-1}, \mathbf{x}_N \right], \quad (3.12)$$

where p is the number of scalar values in a given frame and the subscript 0 and N correspond to the first and last frames of the sequence respectively. It is then possible to describe the evolution of the state vector \mathbf{x}_n to \mathbf{x}_{n+1} over some time step δt by a linear matrix operator \mathbf{L}

$$\underbrace{\mathbf{x}_{n+1}}_{p \times 1} = \underbrace{\mathbf{L}_n}_{p \times p} \underbrace{\mathbf{x}_n}_{p \times 1}. \quad (3.13)$$

If the evolution of the system is linear, (3.13) would be exact. However, making the assumption that the non-linear system varies slowly over the time interval δt , then this mapping is equivalent to a linear tangent approximation (Schmid, 2010). For the experiments shown here, this is a valid assumption due to the two discrete timescales whereby the slow time evolution of $\tilde{\Psi}$ is much less than the fast time scales ($|\partial \tilde{\Psi} / \partial t| \ll |\omega_p \tilde{\Psi}|$). The sequence of state vectors \mathbf{X}_0 can then be described as a Krylov sequence (Greenbaum, 1997)

$$\mathbf{X}_0 = [\mathbf{x}_0, \mathbf{L}_0 \mathbf{x}_0, \dots, \mathbf{L}_0^{N-1} \mathbf{x}_0, \mathbf{L}_0^N \mathbf{x}_0]. \quad (3.14)$$

The evolution for the sequence of state vectors can then be written in matrix form as

$$\underbrace{\mathbf{X}_1}_{p \times N} = \underbrace{\mathbf{L}_0}_{p \times p} \underbrace{\mathbf{X}_0}_{p \times N}. \quad (3.15)$$

where \mathbf{X}_1 is the new set of state vectors advanced by time δt

$$\mathbf{X}_1 = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N, \mathbf{x}_{N+1}]. \quad (3.16)$$

However, as the number of pixels p in a frame is approximately $\mathcal{O}(10^7)$, p is generally larger than the number of frames N $\mathcal{O}(10^2)$ and it more computationally efficient to approximate (3.15) as

$$\underbrace{\mathbf{X}_1}_{p \times N} \approx \underbrace{\mathbf{X}_0}_{p \times N} \underbrace{\mathbf{A}_0}_{N \times N}, \quad (3.17)$$

where the evolution of the system can be described by the complex eigenvalues and eigenvectors of \mathbf{A} . If the modulus of these eigenvalues is equal to one, then they represent the steady, non-decaying ‘modes’ of the system and their complex argument sets the frequency.

One way to solve for these eigenvalues is to use a least squares solution to the linear system. However, practically the application of this method is ill-conditioned, as the number of dominant modes extracted can not be more than double the number of waves in the system (as a real valued sinusoid is the sum of two complex conjugate exponentials). The ‘rank’ of the dominant signal would therefore be much smaller than the number of frames N for which the DMD was based, with the rest corresponding to experimental noise. To overcome this, a singular value decomposition (SVD) is employed. The matrix \mathbf{X}_0 is decomposed by

$$\mathbf{X}_0 = \underbrace{\mathbf{U}}_{p \times p} \underbrace{\mathbf{\Sigma}}_{p \times N} \underbrace{\mathbf{T}^\dagger}_{N \times N}, \quad (3.18)$$

where \dagger is the conjugate transpose. The two orthogonal matrices \mathbf{U} and \mathbf{T} contain the normalised eigenvectors of $\mathbf{X}_0 \mathbf{X}_0^\dagger$ and $\mathbf{X}_0^\dagger \mathbf{X}_0$ respectively. The singular values, which are the square root of the eigenvalues of $\mathbf{X}_0 \mathbf{X}_0^\dagger$, are then contained in the diagonal matrix $\mathbf{\Sigma}$. Substituting (3.18) into (3.17), the linear system can be expressed by

$$\underbrace{\mathbf{\Sigma}}_{p \times N} \underbrace{\mathbf{T}^\dagger}_{N \times N} \underbrace{\mathbf{A}}_{N \times N} = \underbrace{\mathbf{U}^\dagger}_{p \times p} \underbrace{\mathbf{X}_1}_{p \times N}. \quad (3.19)$$

In (3.19), the bottom $p - N$ rows of the diagonal matrix $\mathbf{\Sigma}$ are equal to zero and therefore so are the bottom $p - N$ rows of \mathbf{U}^\dagger . Removing these from the expression leads to the form

$$\underbrace{\mathbf{\Sigma}'}_{N \times N} \underbrace{\mathbf{T}^\dagger}_{N \times N} \underbrace{\mathbf{A}}_{N \times N} = \underbrace{\mathbf{U}'^\dagger}_{N \times p} \underbrace{\mathbf{X}_1}_{p \times N}. \quad (3.20)$$

where the superscript $'$ indicates the reduced form. The next step is to further reduce the singular value decomposition. The ‘rank’ r of the flow describes the number of non-decaying modes in the system, given by the number of non-zero singular values in

Σ' , or singular values above a prescribed threshold (Schmid, 2010). For this internal wave system, $r \ll N$, meaning the bottom $N - r$ rows of Σ' and \mathbf{U}'^\dagger can be set to zero without losing information about the flow. Then, as Σ' is diagonal, it follows that the right $N - r$ columns are also set to zero, converting Σ' to a $r \times r$ invertible matrix. Left multiplying both sides by Σ'^{-1} gives

$$\underbrace{\mathbf{T}^\dagger}_{N \times N} \underbrace{\mathbf{A}}_{N \times N} = \underbrace{\Sigma'^{-1}}_{r \times r} \underbrace{\mathbf{U}'^\dagger}_{r \times p} \underbrace{\mathbf{X}_1}_{p \times N}. \quad (3.21)$$

Then, as we are looking for a solution to the linear operator \mathbf{A} , (3.21) is left multiplied by \mathbf{T} . Since, however, the dimensions of matrix Σ'^{-1} are $r \times r$, the right $N - r$ columns of \mathbf{T} are then discarded, leaving the reduced form \mathbf{T}' . This leaves an expression for \mathbf{A} as

$$\underbrace{\mathbf{A}}_{N \times N} = \underbrace{\mathbf{T}'}_{N \times r} \underbrace{\Sigma'^{-1}}_{r \times r} \underbrace{\mathbf{U}'^\dagger}_{r \times p} \underbrace{\mathbf{X}_1}_{p \times N}, \quad (3.22)$$

where \mathbf{A} is of rank r . This means it only has r independent eigenvectors, \mathbf{z} , corresponding to r non-zero eigenvalues, λ , which describe the ‘modes’ of the flow. We then express the eigenproblem as

$$\mathbf{AZ} = \mathbf{Z}\Lambda, \quad (3.23)$$

where \mathbf{Z} is the matrix of eigenvectors \mathbf{z} of \mathbf{A} and Λ is the corresponding diagonal matrix of eigenvalues. Then, right multiplying both sides of (3.23) by \mathbf{Z}^{-1} and substituting into (3.17), gives

$$\mathbf{X}_1 = \mathbf{X}_0 \mathbf{Z} \Lambda \mathbf{Z}^{-1}, \quad (3.24)$$

which, applied n times leaves

$$\mathbf{X}_n = \mathbf{X}_0 \mathbf{Z} \Lambda^n \mathbf{Z}^{-1}. \quad (3.25)$$

Therefore, provided that the linear operator \mathbf{A} varies slowly in comparison with $n\delta t$, the Dynamic Mode Decomposition can be used repeatedly over the experiment to capture the evolution of the dominant modes in the system. This is ideal for the experiments presented here, which, as discussed, have two discrete timescales. Under the assumption that $|\partial\tilde{\Psi}/\partial t| \ll |\omega\tilde{\Psi}|$, the evolution of the non-linear system is much longer than the short time windows used to extract the dominant modes.

Applications of the DMD Algorithm

For this thesis, the DMD algorithm was used to determine the dominant frequencies for the internal waves in the flow field, specifically the frequencies of the two resonant waves produced from the Triadic Resonance Instability (see Chapter 2). The algorithm was performed on the gradient of the density perturbation given from the Synthetic Schlieren (see § 3.7), producing complex modes for both gradient components. The time evolution for the state vectors, δt , was given as the time between frames (normally 1 second) and the number of frames N in each sequence was normally set at 20.

The DMD algorithm performed well at determining the dominant frequencies corresponding to the resonant triad, however the rank of the frequencies produced for the vertical component of the density gradient $\partial\rho/\partial z$ (3.11b), were often different to those produced from the horizontal component $\partial\rho/\partial x$ (3.11a). This is due to the vertical gradient being stronger for lower frequency waves, while the horizontal gradient has a stronger signal for waves with frequency closer to the buoyancy frequency.

An example of the different outputs from the DMD for the horizontal and vertical component is shown in Figure 3.20, which presents snapshots from an experiment conducted in June 2019. The experiment was performed in a linear stratification where $N = 1.55 \pm 0.04 \text{ rad s}^{-1}$ with a simple four wavelength wide wavebeam of forcing amplitude $\eta_0 = 3.5 \text{ mm}$ and a forcing frequency of $\omega_0 = 0.95 \text{ rad s}^{-1}$ ($\omega_0/N \approx 0.62$). Full details of the forcing from the wavemaker for this experiment are given in § 4.2. Due to the forcing amplitude being below the critical value, the primary wavebeam did not become unstable to TRI (the theory behind the amplitude threshold was presented in § 2.5). The two left hand plots in Figure 3.20 show the two dominant frequencies from the vertical component of the Synthetic Schlieren ($\partial\rho/\partial x$) while the right hand plots show the two from the horizontal component ($\partial\rho/\partial z$) over a sequence of 20 frames. Only the real component of the complex output is shown.

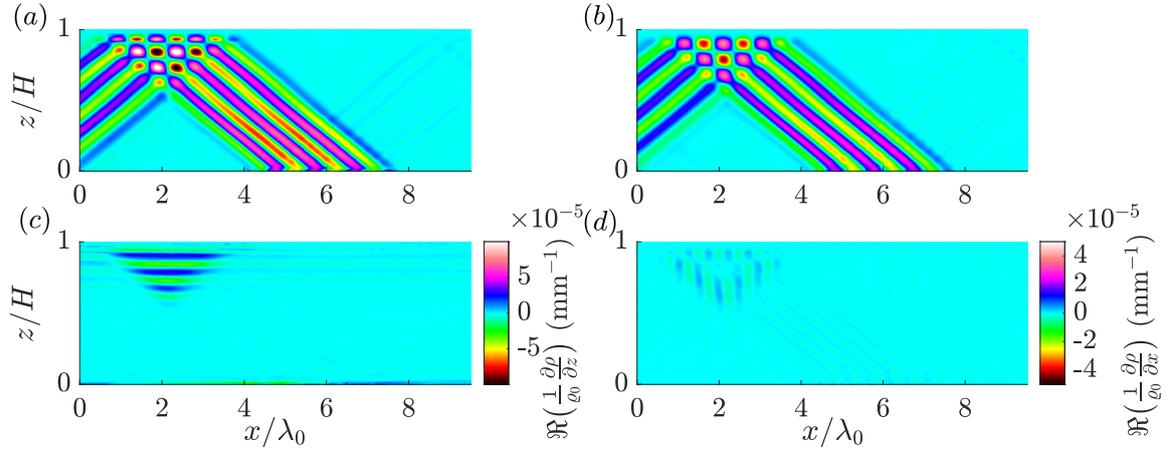


Figure 3.20: Sequence of images showing the difference between the DMD output over frames 300 – 320, for the horizontal and vertical components of the density perturbation. Only the real component of the complex output is shown. The imaginary component of each modal field (not shown) corresponds to the modes complex conjugate. On the left (a) and (c) are the two strongest modes from the vertical gradient of the density perturbation ($\partial\rho/\partial z$), while their corresponding outputs for the horizontal gradient ($\partial\rho/\partial x$) are given in (b) and (d). The frequencies of each image are (a) $\omega = 0.95 \text{ rad s}^{-1}$, (b) $\omega = 0.95 \text{ rad s}^{-1}$, (c) $\omega = 0 \text{ rad s}^{-1}$, (d) $\omega = 1.55 \text{ rad s}^{-1}$. The vertical axis has been normalised by the depth of the tank H , while the horizontal axis has been normalised by the wavelength of the forced wavebeam λ_0 , imposed by the wavemaker.

As expected the most dominant eigenmode for both the vertical and horizontal component was the input frequency of $\omega_0 = 0.95 \text{ rad s}^{-1}$ (shown Figure 3.20(a) and (b) respectively), as this was the prescribed forcing frequency from the wavemaker. Then, for the vertical component, the second dominant eigenmode corresponded to $\omega = 0 \text{ rad s}^{-1}$, as shown in Figure 3.20(c). As discussed in § 2.4, this frequency is due to the interaction of the primary beam with its reflection, which generates zero-frequency, ‘mean’ flow. Due to the strong vertical perturbation generated by this mode (with almost negligible horizontal signal), this frequency was only present for $\partial\rho/\partial z$. In contrast, the second dominant eigenmode for the horizontal component was $\omega = 1.55 \text{ rad s}^{-1}$, shown in Figure 3.20(d). As this component was the same frequency as the buoyancy frequency, its presence was only shown in the horizontal gradient of the density perturbation.

As both the horizontal and vertical component of the Synthetic Schlieren were required for further processing, it was necessary that the DMD modes from $\partial\rho/\partial x$ and $\partial\rho/\partial z$ corresponded to the same frequencies. To ensure this coupling, the horizontal and vertical components from the Synthetic Schlieren were vertically stacked before being analysed by the DMD algorithm. This stacking also removed any bias towards

low or high frequencies for the singular values. As the DMD manipulates the image into a single column vector, this method then meant the dominant eigenvalues for each image were coupled and could be applied to both directional components. Figure 3.21 shows the outputs of the DMD after the images were stacked.

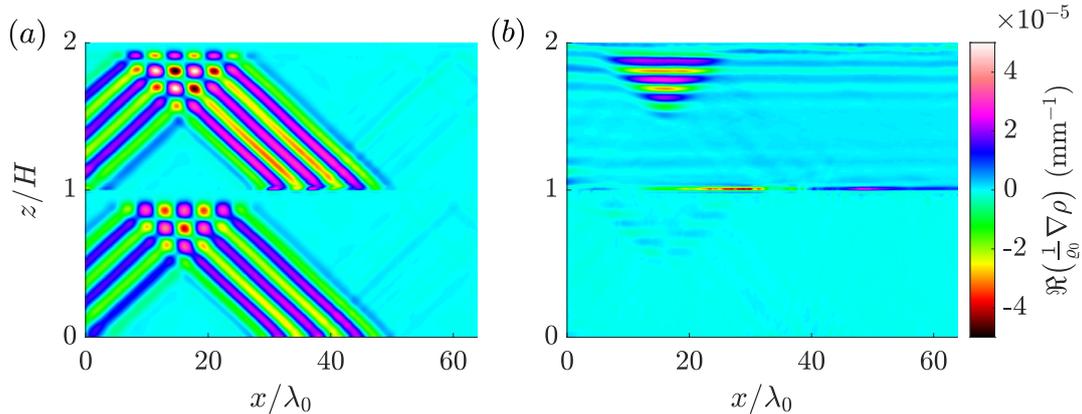


Figure 3.21: Sequence of images showing the dominant output frequencies from DMD when using the vertical and horizontal components of the gradient of the density perturbation stacked on top of each other. The experiment is the same as that shown in Figure 3.20 and, again, only the real part of the complex output is shown. Stacking the vertical and horizontal components ensures a coupling for the output frequencies. (a) $\omega = 0.95 \text{ rad s}^{-1}$, (b) $\omega = 0 \text{ rad s}^{-1}$.

The stacking of the two directional gradient components resulted in a change to the rank of the singular values for the horizontal component of the density perturbation. If a wavebeam became unstable to TRI, the eigenmode with the highest rank for both directional components was always ω_0 . For the vertical component, the subsequent dominant mode was $\omega = 0 \text{ rad s}^{-1}$ before the triadic ω_1 and ω_2 . Yet for the horizontal component, the triadic waves corresponded to the second and third modes. Interestingly, when stacked, the eigenmode with the second strongest signal, still corresponded to $\omega = 0 \text{ rad s}^{-1}$ before the triadic waves.

Once the images had been run through the DMD, they were then separated back into their respective horizontal and vertical components. The next stage of the processing integrated these fields to obtain both the real and imaginary parts of the full density perturbation field $\rho = \tilde{\rho}e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}$ using

$$\rho = \int_x \int_z \nabla \rho \partial z \partial x. \quad (3.26)$$

The full details for this numerical integration process is outlined in Hazewinkel et al. (2011). After this, the full complex density perturbation field ρ was then directionally filtered using the Hilbert Transform outlined below.

3.9.3 Hilbert Transform

Since there are four different sign combinations of the wavenumber vector that satisfy the dispersion relationship for a single frequency, the Hilbert Transform (HT) allows for the directional decomposition of a wave field to determine the propagation directions of separate internal waves. The filtering technique uses Fourier Transforms to isolate waves propagating with different directions. It was first used for internal waves by Mercier et al. (2008), who showed how to isolate quasi two-dimensional internal waves based off the sign of their wavenumber vector. However, their method first uses Fourier transforms in time and frequency windowing to transform a real valued signal into a complex one. This process was not required here as the DMD algorithm had already converted the image sequence into complex modes. The HT was performed using an inbuilt macro in DigiFlow (Dalziel, 2021), briefly outlined below. Its application to the experiments is then detailed.

The algorithm used in DigiFlow takes a two-dimensional Fourier Transform of the complex DMD mode. Once the image is in wavenumber space, it is then possible to filter for the desired sign combinations of the wavenumber vector. The amplitudes of the three quadrants that are not of interest are set to zero, leaving only one sign combination. The wavenumber vector quadrants can be seen in the schematic in Figure 1.2. The inverse Fourier transform is then performed on the remaining signal, to recover the real valued RMS amplitude $|\tilde{\rho}|$ and phase of the original mode in physical space.

For comparison, Figure 3.22(a) shows the root mean square (RMS) of the real and imaginary components of the full density perturbation field from the experiment presented in Figure 3.21. The right panel in (b) shows the RMS amplitude of the density perturbation field produced from the Hilbert Transform algorithm. Here, only the wavenumber vector pointing down and to the left was selected in Fourier space, isolating the primary beam before it was reflected.

Some additional vertical signature is present in Figure 3.22(b) that is not present in (a). This noise is generated from the Fourier transforms used in the HT that are performed over the size of the domain. Zero padding the images before applying the HT algorithm would reduce this vertical artificial signature.

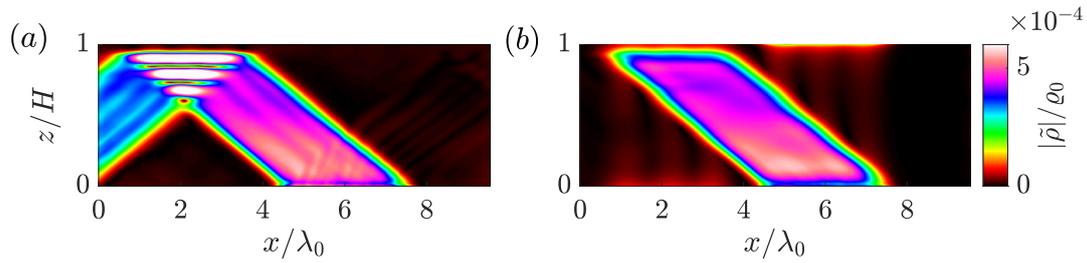


Figure 3.22: Two images showing the results of the directional filtering from the Hilbert Transform. (a) The RMS amplitude computed from the complex density perturbation field. (b) The RMS amplitude of the density perturbation field produced from the Hilbert Transform, where the wavenumber components in the third quadrant with $\mathbf{k} = (-l, -m)$ are isolated.

3.9.4 Experimental determination of amplitude

The final stage of the processing was to extract a quantitative measure of the amplitude of the waves. As noted by Mercier et al. (2010), there is a conversion factor between the amplitude of the wavemaker η_0 and that of the forced beam due to transmission through the stratification, which changes over time. This is further explored with regards to ASWaM by Dobra et al. (2019). For this reason, the amplitudes of the generated waves were measured from the experimental results.

We use the reduced streamfunction $\tilde{\Psi}$ as a measure of the waves amplitude. This was calculated from the linear polarisation relationship given in (2.17), repeated here in terms of each triadic wave $p = 0, 1, 2$

$$\tilde{\Psi}_p = -\frac{\omega_p}{l_p} \frac{g}{N^2 \varrho_0} \tilde{\rho}_p. \quad (3.27)$$

It was then possible to use this relationship to obtain the magnitude of the streamfunction for the experiments. Using the real valued RMS magnitude of the density perturbation field normalised by the reference density $|\tilde{\rho}_p|/\varrho_0$ (after the Hilbert Transform has been employed as shown in Figure 3.22(b)), the real valued magnitude of the reduced streamfunction $\tilde{\Psi}$ was calculated by

$$|\tilde{\Psi}_p| = \left| \frac{\omega_p}{l_p} \frac{g}{N^2} \right| \frac{|\tilde{\rho}_p|}{\varrho_0}, \quad (3.28)$$

where $|\tilde{\Psi}_p| = (\mathbf{x}, t)$. The next step was to define the two different domains \mathcal{D} over which to spatially average the values of $|\tilde{\Psi}_p|$. The first domain \mathcal{D}_w is defined as the

whole image shown in Figure 3.23(a). To extract the measure of $|\tilde{\Psi}_p|$ over this domain

$$\langle |\tilde{\Psi}_p| \rangle_w = \frac{\int_{\mathcal{D}_w} |\tilde{\Psi}_p| d\mathbf{x}}{\int_{\mathcal{D}_w} d\mathbf{x}}. \quad (3.29)$$

The average of $|\tilde{\Psi}_p|$ over this domain contained regions of no wave motion and is not an accurate representation for the actual amplitude of the wave. It is, however, the most appropriate for comparing the amplitude of the primary away against the triadically produced resonant waves, for reasons discussed further in Chapter 4. Moreover, as the HT was employed, averaging over this domain did not include signal from reflections or other waves that are not of interest.

The second domain, \mathcal{D}_r , is given by the small region shown in white on Figure 3.23(b). The value of $|\tilde{\Psi}_0|$ over this domain is given as

$$\langle |\tilde{\Psi}_0| \rangle_r = \frac{\int_{\mathcal{D}_r} |\tilde{\Psi}_0| d\mathbf{x}}{\int_{\mathcal{D}_r} d\mathbf{x}}. \quad (3.30)$$

The average over this region allows for the most direct comparison with theoretical values of $|\tilde{\Psi}_0|$ and for a more accurate representation of the actual amplitude. This region was only ever used for the primary wave.

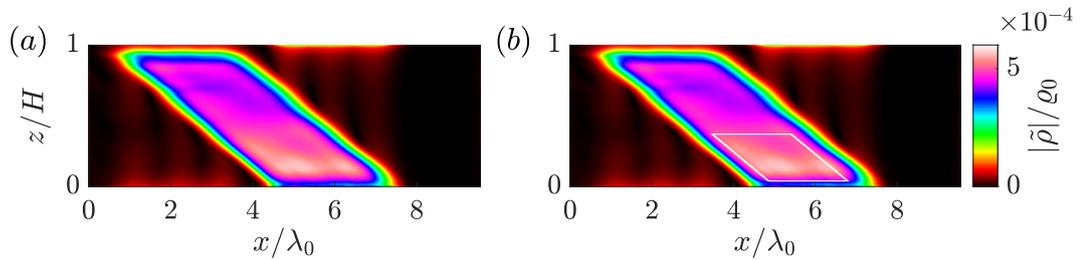


Figure 3.23: Two images showing the difference in regions used for averaging the reduced streamfunction after the Hilbert Transform was employed. (a) The domain \mathcal{D}_w is defined by the whole the field of view which is averaged over to obtain $\langle |\tilde{\Psi}_p| \rangle_w$, where $p = 0, 1, 2$ depending on the triadic wave in question. (b) The domain \mathcal{D}_r is shown by the white region, which is averaged over to obtain $\langle |\tilde{\Psi}_0| \rangle_r$. This region is only used for the primary wave.

3.10 Summary

This chapter has outlined the various experimental techniques used to produce and record the internal wave experiments. We introduced the novel wavemaker used for wave generation and the advances in experimental automation that allowed for

improved quality of the results via noise reduction. The methods for temporal and directional filtering of the experiments have also been outlined. In the next chapter these techniques are utilised to present the first set of experimental results examining a finite-width internal wave beam becoming unstable to TRI.

Chapter 4

Experimental Triadic Resonance Instability investigation

4.1 Introduction

The simplest waveform that can be imposed on the wavemaker is a quasi-monochromatic sinusoid, with a single frequency and dominant wavenumber. Indeed, this was the form of the first experiments that were conducted to understand the propagation of internal waves and to test the capabilities of the wavemaker. Due to the evolution of these experiments, however, this simple waveform became one of the key focuses for the remainder of the research. Rather than the forced wave beam purely propagating down the tank, after sometime, the emergence of a secondary, non-forced wave of lower frequency was seen emanating from the central region of the primary wave. Multiple repeats of the experiment exhibited the same phenomena, which, after some analysis, was revealed to be the result of Triadic Resonance Instability (TRI). Due to the orientation of the wave vectors, only one of the resonant waves could be seen immediately. After performing the DMD algorithm (see § 3.9.2) on the experiments the third wave that completed the triad was found with the same orientation as the primary beam.

After analysing the published literature it was clear that this was not the first time TRI has been witnessed experimentally. In an oscillating cylinder experiment, Clark and Sutherland (2010) attributed the breakdown of a wavebeam due to TRI, meanwhile Joubaud et al. (2012) and Bourget et al. (2013) clearly showed the growth of the instability mechanism for a finite-width beam using their sidewall wavemaker. What was interesting however, was how the experimental results differed from the theory being used to describe the evolution of TRI. The accepted theory for the TRI

interactions predicts that after an initial growth period, the amplitude of all three beams in the triad reach steady state. In contrast, the experiments conducted for this research showed a continuous oscillation, not just in the amplitudes of the three beams, but also of the wavenumbers and frequencies of the two resonant waves.

The structure of this chapter is as follows. In § 4.2 the experimental set-up is outlined, giving details on the wavebeam parameters used in examining the TRI instability. Then § 4.3 presents the experimental results, specifically focusing on how the amplitudes, wavenumbers and frequencies of the triad oscillate over long time-scales (in the order of several hundred buoyancy periods). The amplitude threshold for instability of finite-width beam is then discussed in § 4.4. Finally, § 4.5 provides a comparison of these experimental results with the previous published theory on the development of TRI in finite-width beams. A summary is provided in § 4.6.

4.2 Experimental wave form

Multiple sets of experiments were conducted over the course of two years, all examining a simple wavebeam becoming unstable to TRI. For each set of experiments, the tank was filled with a linear stratification, the buoyancy frequencies of which ranged between $N = 1.46$ to 1.70 ± 0.04 rad s⁻¹. Once the tank was filled, experiments were run for approximately five days in the same stratification, after which the stratification had degraded and it was necessary to drain and clean the tank. Normally filling was done on Wednesday or Thursday to allow for experiments to be run over the weekend using the automated procedure outlined in § 3.8.

Most of the experiments had a duration of 90 minutes and were captured at either 1 or 2 frames per second (fps)¹, however some experiments were run for up to 180 minutes. For all the experiments presented in this chapter, the vertical displacement $z = h(x, t)$ imposed on the neoprene foam was

$$z = h(x, t) = \begin{cases} \Re(f(t) e^{il_0 x} \cos^2(\frac{x-B}{8\pi^2})), & A < x < B, \\ \Re(f(t) e^{il_0 x}), & B < x < C, \\ \Re(f(t) e^{il_0 x} \cos^2(\frac{x-C}{8\pi^2})), & C < x < D, \\ 0, & \text{elsewhere,} \end{cases} \quad (4.1)$$

¹The frame rate of 1 frame per second was chosen to reduce the file size, as each 90 minute experiment is ≈ 45 GB. As the period of the primary wave was $T_0 = 2\pi/\omega_0 \approx 6.6$ s, a frame rate of 1 fps was enough to adequately capture the the period of the forcing.

where the locations A, B, C, D are respectively $7\pi/|l_0|$, $9\pi/|l_0|$, $13\pi/|l_0|$, $15\pi/|l_0|$ and l_0 is the horizontal component of the primary wave vector, here set to -0.05 mm^{-1} . As shown in Figure 1.2, the signs of the components in the wavenumber vector (l, m) provide the direction of propagation. The horizontal separation of the wave crests of the primary wave is $\lambda_0 = 2\pi/|l_0| \approx 125 \text{ mm}$. The spatial structure of the forcing, described by (4.1), takes the form of a beam comprised of four wavelengths enclosed by a cosine-squared envelope on the outer flanks. The temporal forcing $f(t)$ is then described as

$$f(t) = \begin{cases} 0, & \text{for } t \leq 0 \text{ s,} \\ \eta_0 \left(\frac{t}{30}\right) e^{-i\omega_0 t}, & \text{for } 0 \leq t \leq 30 \text{ s,} \\ \eta_0 e^{-i\omega_0 t}, & \text{for } 30 \text{ s} \leq t \leq t_{\text{end}} - 30 \text{ s,} \\ \eta_0 \left(\frac{t_{\text{end}}-t}{30}\right) e^{-i\omega_0 t}, & \text{for } t_{\text{end}} - 30 \text{ s} \leq t \leq t_{\text{end}}, \end{cases} \quad (4.2)$$

where ω_0 and η_0 are the forcing frequency in rad s^{-1} and nominal forcing amplitude in mm of the primary beam, respectively. For the experiments in this chapter, $\omega_0 = 0.95 \text{ rad s}^{-1}$, while η_0 varies depending on the strength of the stratification. Based on this forcing frequency, the period of the primary wave was $T_0 = 2\pi/\omega_0 \approx 6.6 \text{ s}$. The amplitude η_0 specifies the distance from the equilibrium horizontal level of the wavemaker to the crest or trough of the wave. The full crest to trough displacement is, therefore, $2\eta_0$. The end time for the experiment t_{end} , was often 5460 s corresponding to 90 minutes at full amplitude forcing, although in some cases t_{end} was 10800 s for experiments that were run for 180 minutes. The first 60 s of forcing from the wavemaker is shown in Figure 4.1, where the blue line represents the actual sinusoidal displacement of the wavemaker while the red line outlines the carrier envelope to the beam.

The temporal forcing in (4.2) shows that the amplitude of the beam was increased from rest over a 30 s period to minimise the formation of a boundary layer (Dobra, 2018), thereby ensuring maximum transmission of the waves. In the same way, the envelope over the beam was chosen to smooth the sharp edges and prevent discontinuities forming, as observed by Sutherland et al. (1999). While this smoothing was achieved by the cosine squared edges, the envelope also ensured that the middle wavelengths were forced at constant maximum amplitude, η_0 . The fact that this carrier envelope allows the beam to become unstable agrees with the theoretical work by Karimi and Akylas (2014), who showed that TRI is possible for quasi-monochromatic beams comprised of a sinusoidal carrier modulated by a locally confined envelope given that the beam width is large enough to permit the growth of perturbations.

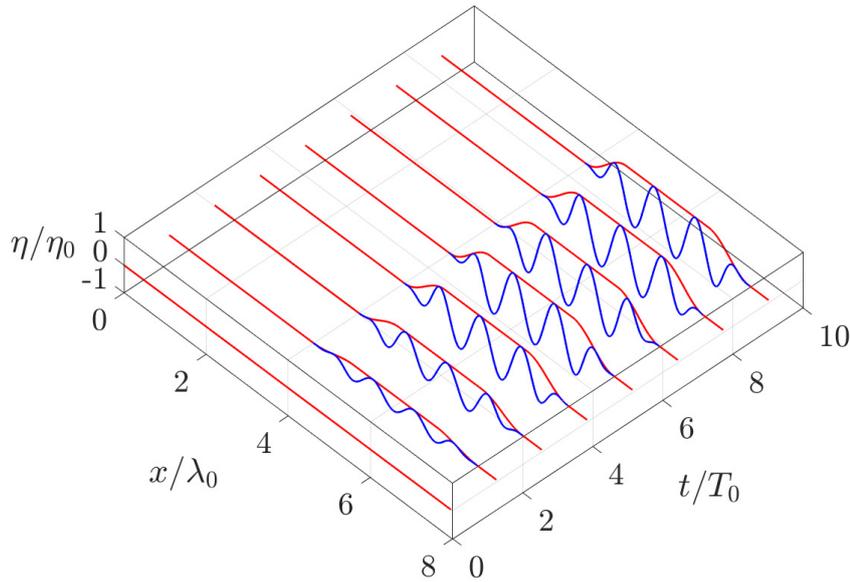


Figure 4.1: Schematic showing the four wavelength wide waveform generated by ASWaM for the first 60 s of forcing. The blue line shows the actual sinusoidal displacement of the wavemaker while the red line outlines the carrier envelope to the beam. The envelope smooths the wavelengths on either side of the beam with a cosine squared curve while the middle two wavelengths reach maximum height η_0 after 30 s. The x, y and z axes are normalised by the wavelength λ_0 , the period T_0 and the amplitude η_0 of the primary beam, respectively.

4.3 Experimental Results

4.3.1 Overview

This section (§ 4.3) examines 36 experiments, captured at one frame per second in the same linear stratification with a buoyancy frequency $N = 1.54 \pm 0.06 \text{ rad s}^{-1}$. The wavebeam forcing used for these experiments is described in § 4.2. The forcing frequency was $\omega_0 = 0.95 \text{ rad s}^{-1}$, providing the ratio $\omega_0/N = 0.61$, while the horizontal component of the primary wave vector was $l_0 = -0.05 \text{ mm}^{-1}$, meaning the beam was directed to the left. The forcing amplitude, η_0 , for the set ranged between 3.25 mm to 4.5 mm. As discussed in § 2.5, for finite-width beams – such as the ones used in these experiments – there is an amplitude threshold that needs to be met before the instability can be triggered. For these experiments, this critical amplitude for instability lay around $\eta_0 \approx 3.9 \text{ mm}$. Further calculations of the amplitude threshold are given in § 4.4. In order to differentiate between the three waves in the triad and

their corresponding parameters, we define

$$\mathbb{W}_p = \{\omega_p, \mathbf{k}_p, \nabla \tilde{\rho}_p, \tilde{\Psi}_p, \dots\}, \quad (4.3)$$

where \mathbb{W} indicates the general wave field and the subscript p corresponds to either 0, 1 or 2 depending on the triadic wave in question.

For the experiments forced above this critical amplitude, the primary wavebeam \mathbb{W}_0 became unstable to TRI. This was apparent by the secondary beam \mathbb{W}_1 that was visible in the flow, emanating from the central region of the primary beam but moving in the opposite direction, with a group velocity down and to the right. This can be observed in Figure 4.2(a), which shows the full flow field at frame 549 ($t/T_0 \approx 83$), from one experiment in the set forced at $\eta_0 = 4$ mm. In order to separate this field into its component frequencies and corresponding modal fields, Dynamic Mode Decomposition (see § 3.9.2) was used. Figure 4.2(b) – (d) shows three of the four non-decaying frequencies produced from the DMD over the frames 545 – 565 ($83 \ll t/T_0 \ll 86$). As expected, Figure 4.2(b) corresponds to \mathbb{W}_0 with frequency $\omega_0 = 0.95$ rad s⁻¹. Figure 4.2(c) and (d) then correspond to the two triadically generated waves \mathbb{W}_1 and \mathbb{W}_2 , with frequencies $\omega_1 = 0.358$ rad s⁻¹ and $\omega_2 = 0.592$ rad s⁻¹ respectively. Summing these frequencies shows that the temporal triadic resonant condition $\omega_0 = \omega_1 + \omega_2$ is satisfied. It should be noted that this triadic condition is not built into the DMD; these frequencies were independently obtained from the algorithm. The fourth non-decaying frequency produced by the DMD corresponded to $\omega = 0$ rad s⁻¹. As the presence of this zero-frequency mode has been discussed in § 2.4 and shown in Figure 3.20, it has not been shown here.

Using the complex modal fields from the DMD, the Hilbert Transform was performed on each triadic waves to determine the direction of its wavenumber vector. The \mathbb{W}_1 wave vector \mathbf{k}_1 was shown to be pointing up and to the right, as opposed to \mathbf{k}_0 that pointed down and to the left, shown by the red arrows on Figure 4.2. As the phase and group velocities of internal waves are perpendicular, the group velocity of \mathbb{W}_1 was therefore down and to the right, as expected from experimental observation. The wave vector \mathbf{k}_2 was also found to be in the same quadrant as \mathbf{k}_0 , down and to the left. Due to its orientation, \mathbb{W}_2 remained initially confined in the spatial extent of the primary beam, hence why its presence is not visible from the snapshot in Figure 4.2(a).

Fast Fourier Transforms (FFTs) were then used to determine the magnitude of the wavenumber vectors corresponding to the secondary waves. The FFT was performed on both the real and imaginary components of the complex DMD output (the real part of which is shown in Figures 4.2(b)–(d)), as here, the fast phase information $e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}$

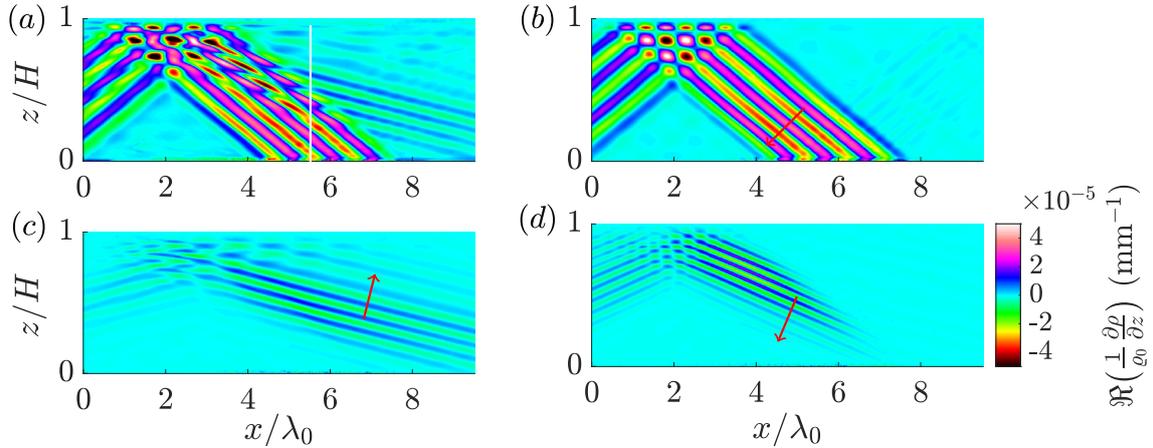


Figure 4.2: (a) The vertical gradient of the density perturbation for the full flow field at 549 s ($t/T_0 \approx 83$) into an experiment forced at $\eta_0 = 4$ mm. The white line gives the location for the vertical time series shown in Figure 4.6. (b)–(d) The real part of three of the dominant frequencies produced from the DMD over frames 545 – 565 ($83 \ll t/T_0 \ll 86$). The red arrows overlaid indicate the orientation of the wavenumber vector. Specifically, (b) \mathbb{W}_0 with $\omega_0 = 0.95$ rad s $^{-1}$ (c) \mathbb{W}_1 with $\omega_1 = 0.358$ rad s $^{-1}$ (d) \mathbb{W}_2 with $\omega_2 = 0.592$ rad s $^{-1}$.

is still present. Specifically, the one-dimensional power spectrum was computed in both the horizontal and vertical directions for the real and imaginary parts of the modal field. Before computing the FFT, the images were embedded in a zero filled matrix to increase the resolution in Fourier space. The wave vector components (l, m) were then obtained by taking the peak value of the combined spectrum from the real and imaginary parts in both the horizontal and vertical directions, respectively.

Using this method, for the snapshots shown in Figure 4.2, $\mathbf{k}_1 = (0.039, 0.156)$ mm $^{-1}$, while $\mathbf{k}_2 = (-0.090, -0.225)$ mm $^{-1}$. Superimposing these wave vectors onto the loci curve given in Figure 2.1, shows that these vectors lie in a closed triangle on top of the black curve, meaning that the spatial triadic resonance condition $\mathbf{k}_0 = \mathbf{k}_1 + \mathbf{k}_2$ is, to a good approximation, also satisfied. This is shown in Figure 4.3.

There is, however, a slight error in how well these vectors satisfy the spatial TRI condition. This is due to the ‘peak’ wavenumber obtained from the FFT spectrum. Unlike the DMD, which approximates the flow by a series of delta functions, the FFT approximates the signal using a polynomial. As these waves beams contain a distribution of wavenumbers, the peak of the polynomial approximation may be skewed by the presence of multiple wavenumbers in the flow and hence not accurately capture the dominant wavenumber. This is discussed further in the following sections.

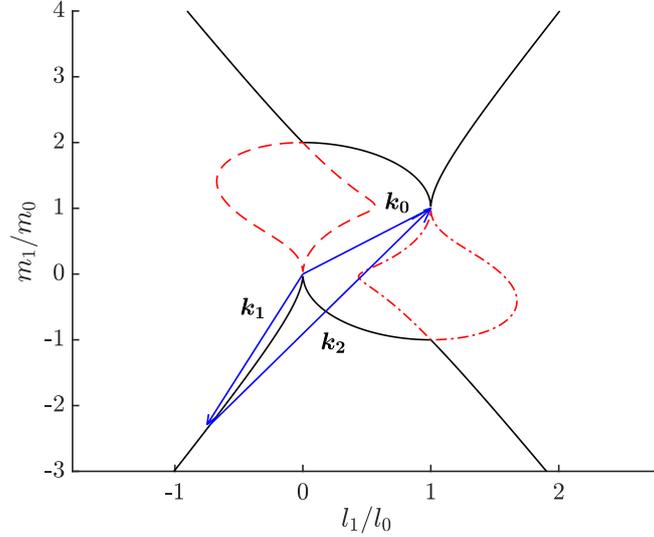


Figure 4.3: Figure showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector with an origin at $(0,0)$. The black line indicates all the possible locations that satisfy (2.5) corresponding to the TRI condition (2.4). The blue arrows show the normalised experimentally produced wavenumber vectors for the resonant triad determined from the peak in the one-dimensional power spectrum performed on the snapshots shown in Figure 4.2.

4.3.2 Temporal and spatial modulation

Initial analysis showed that the triadic resonance conditions of (2.4) were being met and TRI was indeed being witnessed. What was surprising, however, was how the instability developed over the course of the experiment. From qualitatively examining the flow, the amplitudes of the resonant waves appear to grow and then decay over time periods much longer than the wave period T_0 . Figure 4.4 shows snapshots from the same $\eta_0 = 4$ mm forcing experiment presented in Figure 4.2, at roughly 400 s intervals. It is clear from comparing these snapshots that the variability of the resonant \mathbb{W}_1 beam is not only temporal, but spatial. The \mathbb{W}_1 beam seems to gradually move down the primary wave beam while growing in amplitude, then decaying before starting to grow again near the top of the primary beam.

This oscillatory behaviour was demonstrated for the full duration of the experiment, which lasted over 800 forcing periods. The resultant modulation is depicted in Figure 4.5(b), which shows the amplitude of all three triadic beams over the course of the experiment, calculated from the DMD output. The other three graphs correspond to three other experiments in the same set, with different forcing amplitudes. The forcing amplitudes of the experiments in Figure 4.5(a), (c) and (d) were $\eta_0 = 3.75$ mm, $\eta_0 = 4$ mm and $\eta_0 = 4.5$ mm, respectively. The experiment shown in Figure 4.5(c), was run

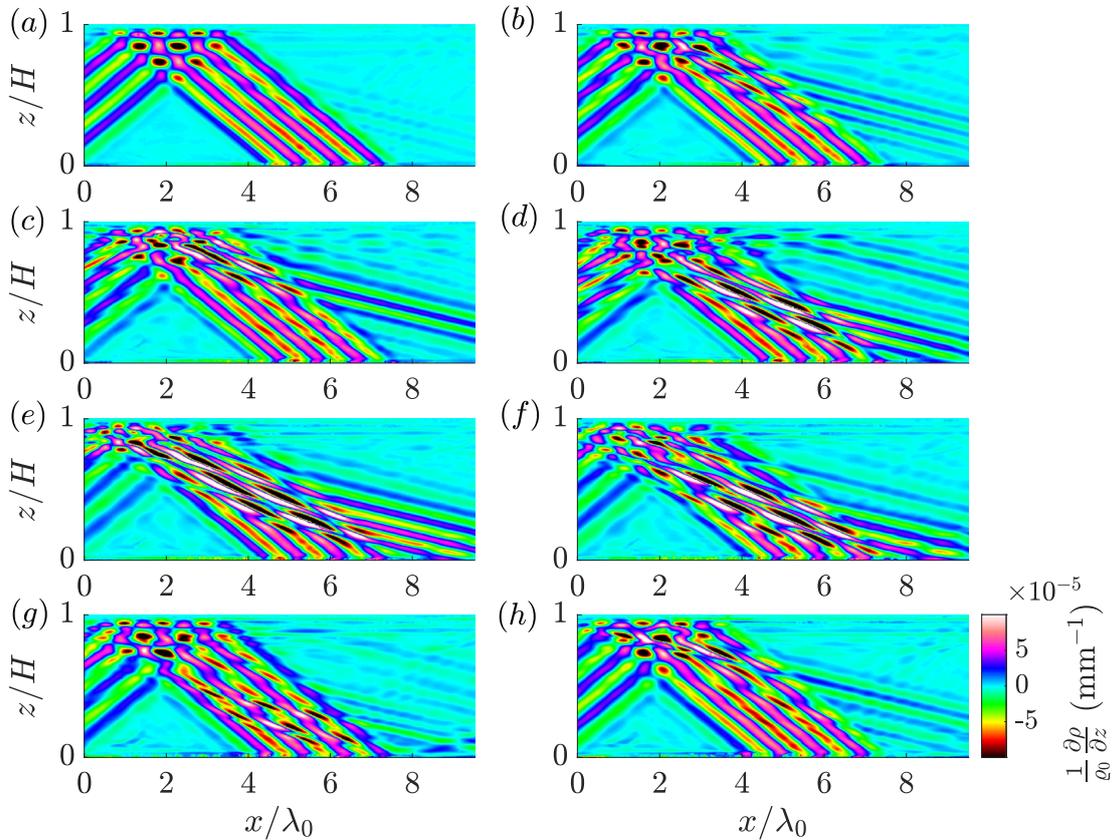


Figure 4.4: Sequence of images at 400 s apart, showing the vertical gradient of the density perturbation for the full flow of the $\eta_0 = 4$ mm experiment. (a) $t/T_0 = 53$, (b) $t/T_0 = 113$, (c) $t/T_0 = 174$, (d) $t/T_0 = 234$, (e) $t/T_0 = 294$, (f) $t/T_0 = 255$, (g) $t/T_0 = 415$, (h) $t/T_0 = 475$.

for 180 minutes, as opposed to the 90 minutes of the other three. All experiments were captured at a frame rate of 1 fps and the DMD interval, N , was kept at 20 frames.

The wave amplitudes shown in Figure 4.5 were determined by taking the arithmetic mean of each triadic wave $|\tilde{\Psi}_p|$ – obtained from (3.28) – over the full domain \mathcal{D}_w using (3.29). As the secondary waves vary spatially over time, using $\langle |\tilde{\Psi}| \rangle_w$ for the measure of amplitude ensures that the fluctuations seen in Figure 4.5 are not due to a confined spatial region over which the secondary waves were being averaged.

One obvious feature in Figure 4.5 is the gradual increase of the mean amplitude of the primary wave over time. In fact, this is also witnessed in lower amplitude forcing experiments that did not become unstable to TRI (not shown here). This increase can not be directly due to the instability, as the TRI mechanism transfers energy from the primary beam to the two secondary waves, as opposed to injecting energy into the primary wave. The amplitude increase might be due to the peristaltic motion of the

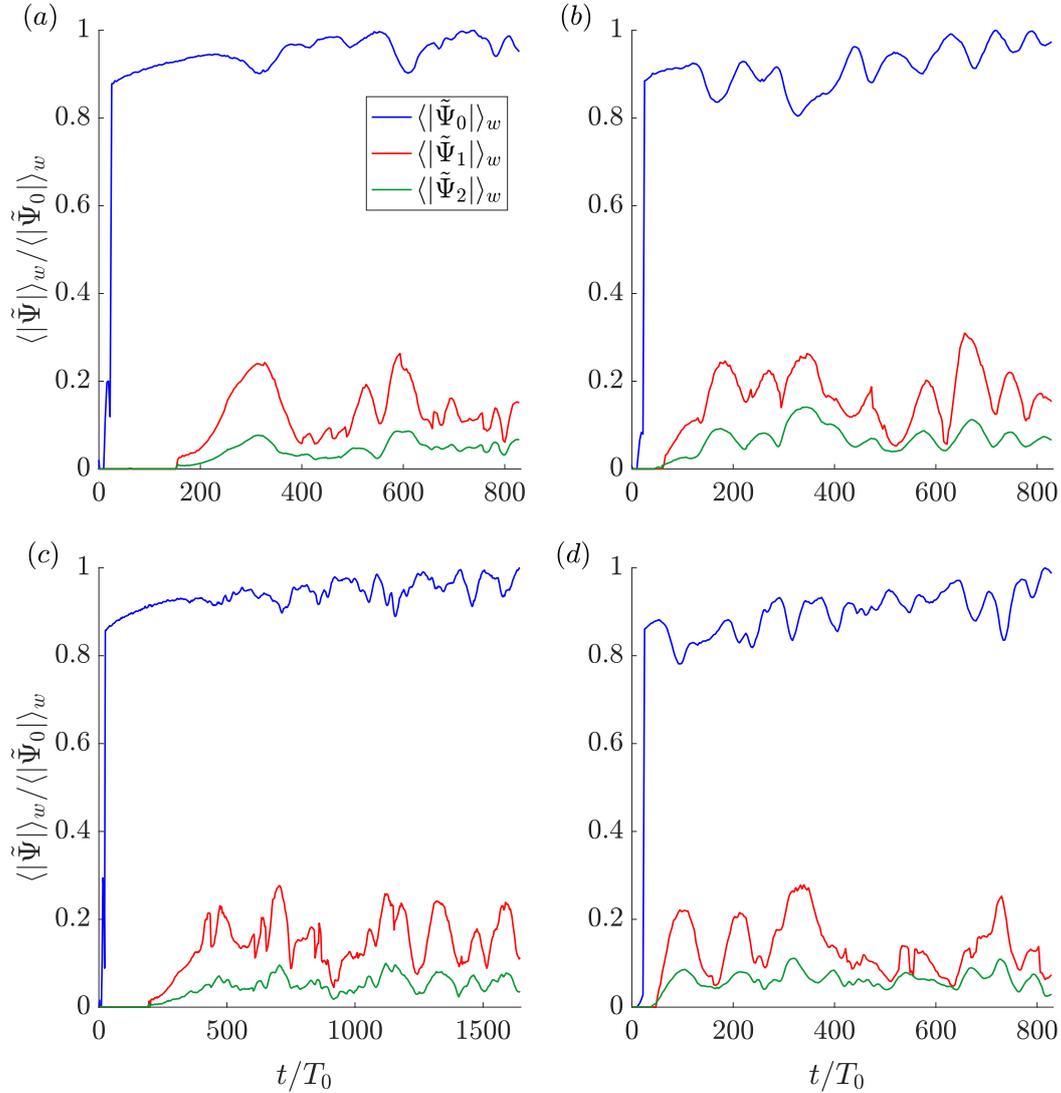


Figure 4.5: The amplitude of the streamfunction for $\langle |\tilde{\Psi}_0| \rangle_w$ (blue), $\langle |\tilde{\Psi}_1| \rangle_w$ (red), $\langle |\tilde{\Psi}_2| \rangle_w$ (green) for four different experiments. All spatial averaging is performed over the whole domain \mathcal{D}_w (see Figure 3.23 for details) and all amplitudes are normalised against the maximum value of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$. Three of the experiments (a),(b) and (d) were run for 90 minutes, while (c) was run for 180 minutes. The forcing amplitudes of the primary wave beam were (a) $\eta_0 = 3.75$ mm, (b) $\eta_0 = 4$ mm, (c) $\eta_0 = 4$ mm and (d) $\eta_0 = 4.5$ mm.

wavemaker leading to a sharpening of the stratification directly above, resulting in an increased coupling efficiency between the energy transfer from the wavemaker to the internal waves.

The coupled modulations in amplitude between the primary wave and the secondary waves, shown in both Figure 4.4 and 4.5, reveal a continued flux in the energy exchange between the waves in the triad. These modulations persisted for the full duration of all of the experiments. Moreover, they seem to be unrelated to the forcing amplitude; when normalised, the amplitude ratios between the triads are similar across all experiments that become unstable. Despite this, Fourier transforms failed to pick up a dominant frequency for these oscillations that was consistent across all the experimental runs.

As we shall see in Chapter 6, we believe part of the oscillation is due to the role of the finite-width beam. As discussed in § 2.5, the energy exchange from the primary beam to the two resonant beams must occur when the secondary waves are in contact with the underlying primary beam. However, the resonant beams are also propagating waves. In the case of W_1 , the orientation of the group velocity (and therefore energy) is in the opposing quadrant to the primary. This means that when a packet of energy in W_1 moves outside of the spatial extent of W_0 , the energy transaction is broken. If the perturbations are not given sufficient time to grow before exiting the primary beam, then the amplitude of the secondary waves decrease, allowing further perturbations to grow in different locations. Figure 4.4 also shows that the region where W_0 , W_1 and W_2 overlap is continuously moving, so that the location across which the transfer of energy occurs is not steady.

4.3.3 Frequency and wavenumber modulation

In order to better understand the oscillatory behaviour of the instability, we turn our attention to the structure of the secondary waves in Fourier space throughout the experiment. Looking at the snapshots in Figure 4.4, W_1 appears to change both its frequency and wavelength gradually over time. From analysing the dominant frequencies of the resonant waves obtained from the DMD across the full experiment, this slow change in frequency is confirmed. The DMD modes show that while the primary wave frequency remains constant, the two secondary wave frequencies modulate slightly over time with a similar period to that of the amplitude oscillations. This is shown in Figure 4.6, which gives the frequency-time spectra computed for the same four experiments presented in Figure 4.5 across the vertical transect marked by the white line in Figure 4.2(a). The frequency-time spectra have been computed using inbuilt Matlab FFT function (Mathworks, 2021) with a signal input length of 150 s,

as indicated by the back time window over Figure 4.2(a). The dominant frequencies obtained from the DMD are then overlaid in white.

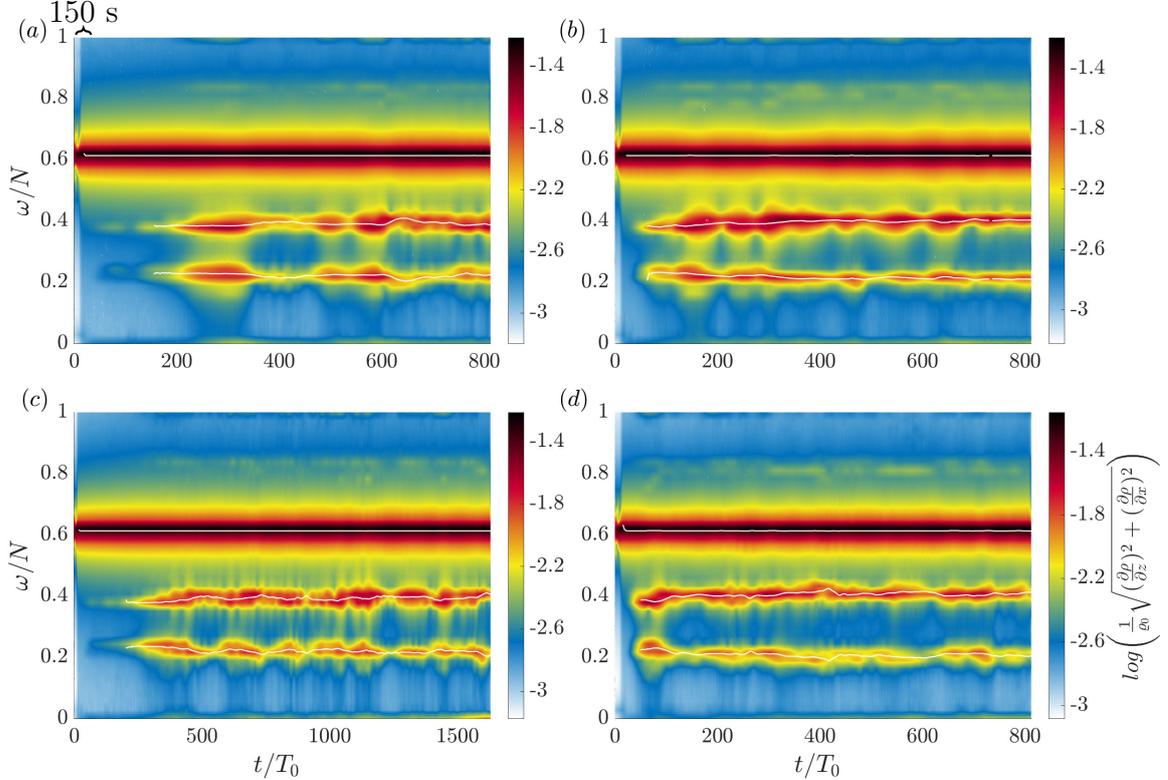


Figure 4.6: Time frequency spectra computed over the vertical transect shown in white on Figure 4.2(a) for the same four experiments in Figure 4.5. Again (a) $\eta_0 = 3.75$ mm, (b) $\eta_0 = 4$ mm, (c) $\eta_0 = 4$ mm and (d) $\eta_0 = 4.5$ mm. The spectral density is given as the logarithm of the RMS field of both the horizontal and vertical gradient of the density perturbation. The dominant frequencies for each experiment obtained from the DMD are then overlaid in white.

Both the DMD and the spectrograms show a modulation to the frequencies of the resonant waves over time. Despite this, when analysing the frequencies obtained from the DMD, the temporal triadic relation $\omega_0 = \omega_1 + \omega_2$ was satisfied at all times (this relationship is not inherent to the DMD algorithm). Other than zero padding, the time series signal used to compute the spectrogram has not been windowed in any way. Qualitatively examining the results of including a Hamming window to the signal showed slightly less variation in the frequency bandwidth although the modulations were still strong.

Interestingly, a similar variation in frequency has been witnessed by both Scoln et al. (2013) and Brouzet et al. (2016) in their studies of non-linear wave attractors. In their experiments, due to increased energy density corresponding to the focusing power

of the attractor, the primary beam breaks down via the TRI mechanism. They also witness slow oscillations to the frequencies and wavenumbers of the two resonant waves after the onset of the instability, however they do not pursue this phenomena further.

As the frequency and wavenumber of internal waves are related via the dispersion relationship, the modulation of the frequency also means a modulation to the wavenumber over time. Using the same method outlined in § 4.3.1, a one-dimensional power spectrum was taken for each complex DMD mode to obtain the wavenumber vector components (l, m) . Again, the peak of the combined spectrum from the real and imaginary components was used to obtain the dominant wavelengths. Figure 4.7(a) shows the same loci curve presented in Figure 4.3 (for the same $\eta_0 = 4$ mm also shown in Figure 4.5(b) and Figure 4.6(b)) but with the wave vectors – computed for each DMD mode – overlaid. Figure 4.7(b), then shows how these wave vectors evolve as a function of time. As each DMD mode used a sequence of 20 frames, the wave vectors are spaced at intervals of 20 s ($t/T_0 \approx 3$) in Figure 4.7(b).

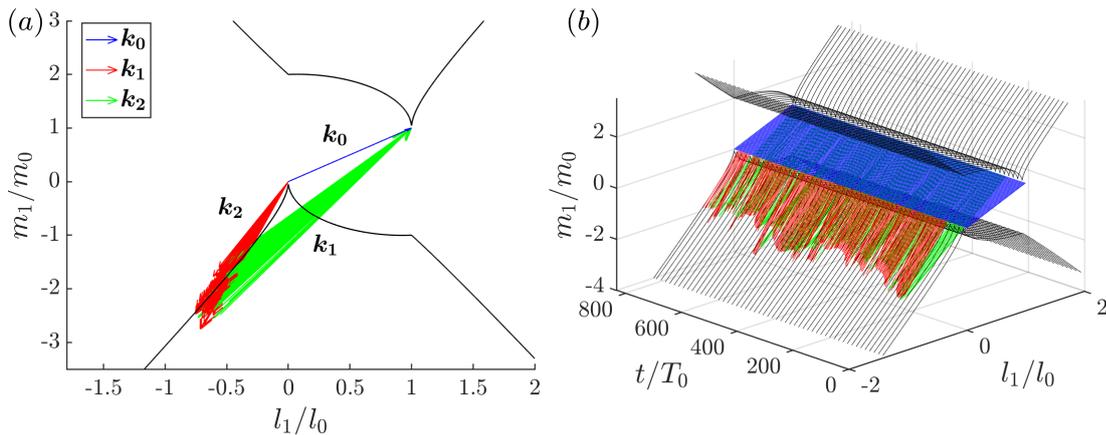


Figure 4.7: (a) Black line showing all the possible locations for the tip of the \mathbf{k}_1 that satisfy (2.5) corresponding to the TRI condition (2.4). The blue arrow show the normalised experimentally produced primary wave vector \mathbf{k}_0 , while the green and red arrow show all of the normalised \mathbf{k}_1 and \mathbf{k}_2 respectively for the full duration of the experiment. Again, the wave vectors correspond to the same experiment shown in Figure 4.5(b) and Figure 4.6(b). (b) The same plot as a function of time. The wave vectors are spaced at intervals of 20 s ($t/T_0 \approx 3$), while the spacing of the loci curve is wider to simply the visualisation.

From Figure 4.7, it can be seen that the spatial triadic resonance condition $\mathbf{k}_0 = \mathbf{k}_1 + \mathbf{k}_2$ is nearly always satisfied, however the agreement for the wavenumber vectors is not as good as it was for the frequencies obtained from the DMD. The reason for this could, in part, be due to the different techniques used to calculate the frequencies as opposed to the wave vectors. As discussed in § 3.9.2, DMD excels when determining the structure of modes over a short time series that can reasonably be

approximated by a combination of delta functions. The dynamic modes capture the dominant flow structures that contribute most to the overall evolution of the measured flow. Using this technique over short time windows (20 s), allowed for the evolution of the resonant frequencies to be accurately captured.

In contrast, the wave vectors in Figure 4.7 were computed using FFTs over each complex DMD output, such as the ones shown in Figure 4.2(b), (c) and (d). Whereas the DMD approximates the underlying energy spectrum by a series of delta functions, the FFT approximates the dominant signal using a polynomial. If multiple wave lengths are present in the flow field, the peak of the FFT approximation will be broader and the corresponding dominant signal may not exactly correspond to the frequency measured from the DMD. In the same way, the underlying spectrograms in Figure 4.6 were computed using FFTs over a time window of 150 s. At certain points in time, the bandwidths of the secondary frequencies are broader, indicating that more than one frequency could be present in the flow.

This is further investigated in Figure 4.8, which shows a comparison between the frequency bandwidth of the spectrogram and the wavenumber vectors. Figure 4.8(a) shows a zoomed view of Figure 4.6(b), looking just at the secondary wave frequencies ω_1 and ω_2 . The lower panel, Figure 4.8(b), then shows how well the spatial resonant condition is satisfied by calculating the distance between the tip of the \mathbf{k}_1 vector and the start of the \mathbf{k}_2 vector for the same experiment.

This analysis was conducted to see if there is a correlation between locations with a broader frequency spectra computed from the FFT frequency spectrogram and an increase in error of the wavenumber vectors, based on how well they form a closed triangle. Figure 4.8 shows a relatively strong correlation between the two, providing evidence that at certain locations, both the frequency signal and the wavenumbers both contain broader signals, indicating the presence of multiple waves in those regions. Based on this analysis, sections of the experiment with a broader frequency spectrum were examined in more detail. Figure 4.9 again shows the real part of the three triadic components from the DMD breakdown of the same experiment presented in Figure 4.8, but now over frames 2099 – 2119 ($318 \leq t/T_0 \leq 321$).

The dominant modes produced from the DMD in Figure 4.9(c) and (d) appear to show two different locations where TRI is occurring over the height of the primary beam. Unlike the modes shown in Figure 4.2, where there was one clear frequency and wavenumber for \mathbb{W}_1 , here the signal is discontinuous across a transition region where the two sections meet. As discussed in § 4.3.2, as a packet of energy in \mathbb{W}_1 exits the underlying primary beam, the energy exchange is broken. If this energy transfer

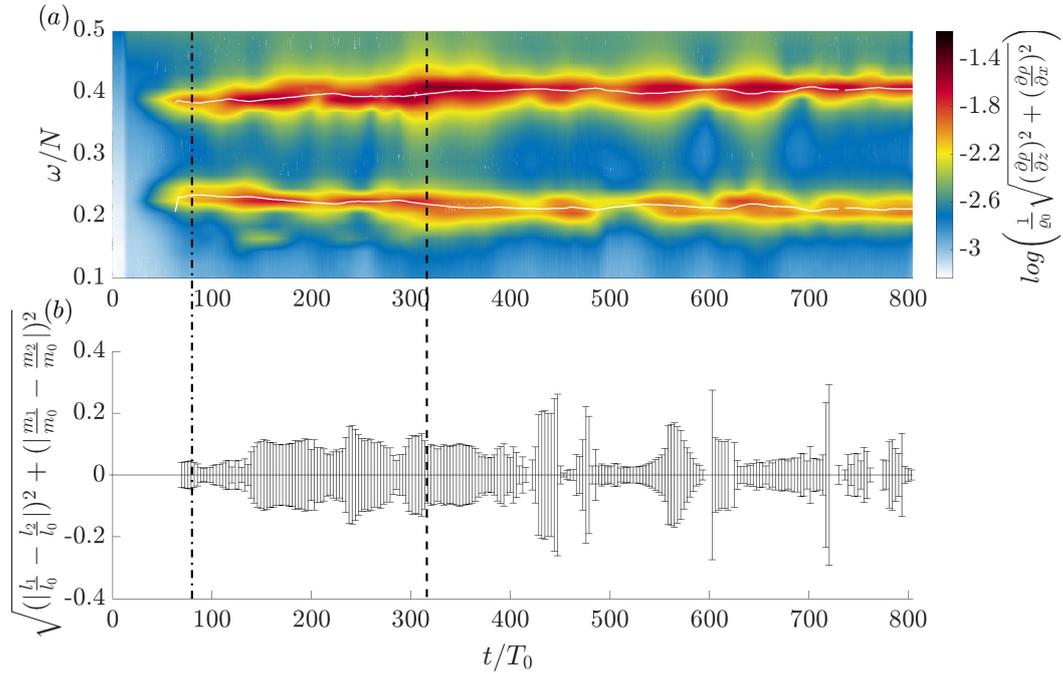


Figure 4.8: (a) A zoom in of the time frequency spectra for experiment shown in Figure 4.6(b), over the frequencies corresponding to the resonant waves. Again, the dominant modes from the DMD are overlaid in white. (b) A graph showing how well the triadic spatial relationship is satisfied for the same experiment. The height of the line for each time period gives the length between the tip of the \mathbf{k}_1 vector and the start of the \mathbf{k}_2 vector shown in Figure 4.7 normalised by \mathbf{k}_0 . The black dot dashed line shows the time of the snapshots in Figure 4.2, while the dashed line shows the time of the snapshots in Figure 4.9.

from the primary wave ceases, then it is able to recover in amplitude and another perturbation is able to grow at a different location. What the analysis here reveals is that new perturbations select slightly different wavenumbers and frequencies of the resonant waves. Indeed, what could be happening, is that as different perturbations grow, the secondary waves with very similar frequencies could interact with each other via the primary wave, generating a slow ‘beating’ effect. This interaction could cause the secondary beams to decay in some locations, while in others it causes a growth in amplitude, amplifying the effects of the modulations.

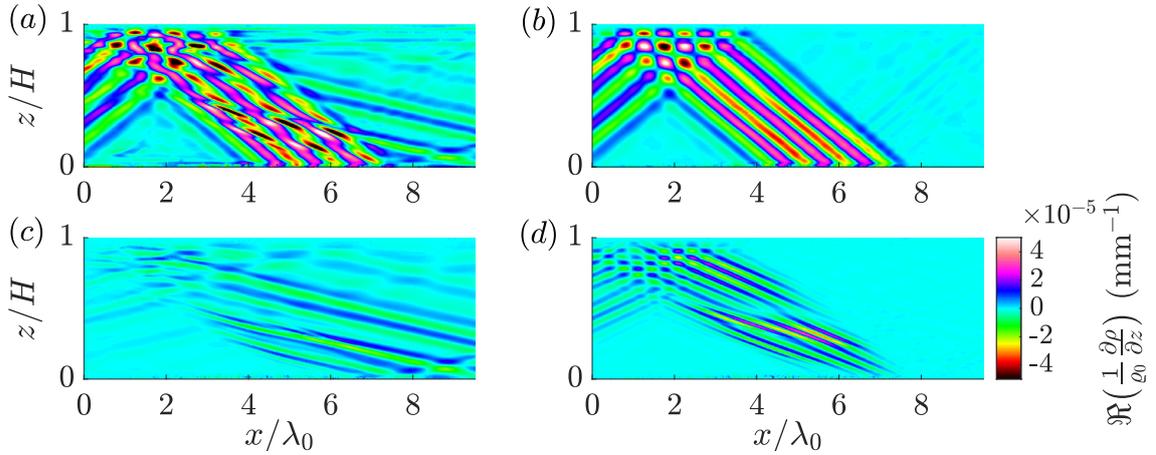


Figure 4.9: (a) The vertical gradient of the density perturbation for the full flow field of the same experiment presented in Figure 4.8 at 2109 s ($t/T_0 \approx 319$). (b)–(d) The real part of three of the dominant frequencies produced from the DMD over the frames 2099 – 2119 ($318 \leq t/T_0 \leq 321$). Specifically, (b) $\omega_0 = 0.95 \text{ rad s}^{-1}$ (c) $\omega_1 = 0.323 \text{ rad s}^{-1}$ (d) $\omega_2 = 0.627 \text{ rad s}^{-1}$.

4.4 Amplitude threshold for a finite-width beam

As discussed in § 2.5, finite-width beams require an amplitude threshold to be met before they will become unstable. The experiments conducted over the course of this research, also showed this threshold to be a function of the strength of the stratification N . Figure 4.10 shows seven sets of experiments, all forced at $\omega_0 = 0.95 \text{ rad s}^{-1}$ and with the wavebeam profile shown in Figure 4.1, but with varying buoyancy frequencies.

Each experiment has been marked with a red or blue circle, corresponding to whether or not it became unstable. A general trend is shown whereby increasing the strength of the stratification (or decreasing the forcing frequency) increases the amplitude threshold. The amplitude is calculated as the arithmetic mean of the region \mathcal{D}_r shown in Figure 3.23, as this most accurately represents the amplitude of the primary wave beam. As seen in Figure 4.5, due to the peristaltic motion of the wavemaker, there was a gradual increase in the mean value of the primary wave over time. Due to this, the value of $\langle |\tilde{\Psi}_0| \rangle_r$ was taken as the mean over the last 200 s of the experiment.

Looking specifically at the set of experiments presented in § 4.3, (indicated by $\omega_0/N = 0.61$ in Figure 4.10), Figure 4.11 shows the measured value of $|\tilde{\Psi}_0|$ for all 36 experiments using the two different regions $\langle |\tilde{\Psi}_0| \rangle_r$ and $\langle |\tilde{\Psi}_0| \rangle_w$. The four experiments presented in Figure 4.5 are also marked. This highlights the importance of using the

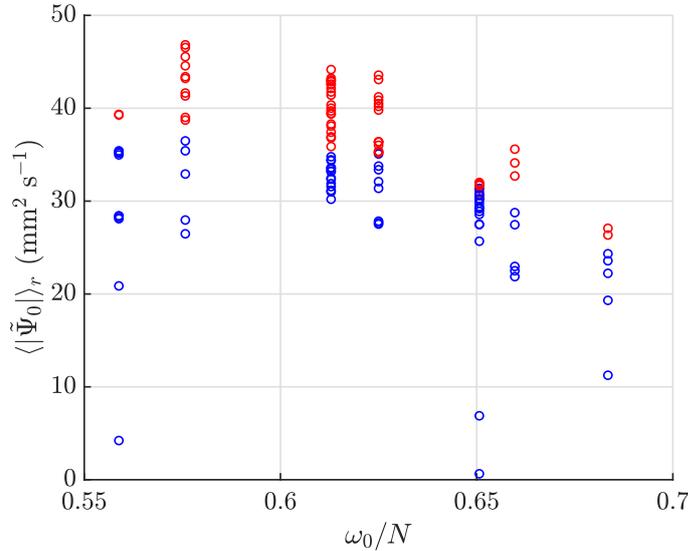


Figure 4.10: Phase diagram showing the amplitude threshold required for the experimental wavebeam to become unstable to TRI. Seven sets of experiments are shown, all with the same \mathbb{W}_0 parameters but with varying forcing amplitudes and in varying strength stratifications. The red and blue circles correspond to instability or no instability respectively. The amplitude has been calculated over a small region $\langle |\tilde{\Psi}_0| \rangle_r$, shown in Figure 3.23, as this most accurately represents the amplitude of the primary wave beam.

experimentally determined measure of the amplitude $|\tilde{\Psi}_0|$ as opposed to the forcing from the wavemaker η_0 , as the two experiments forced at 3.75 mm and 4 mm (marked with a \square and $*$ respectively) result in similar experimental amplitude.

4.5 Comparison with previous theory

4.5.1 Overview

The experimental results presented in this chapter are not the first to explore TRI. The work outlined in § 2.5.3 presents some of the recent experimental and theoretical work by Bourget et al. (2014) looking at the initial growth of TRI for a finite-width beam. This section compares the experimental results presented in § 4.3 with this weakly non-linear theory developed by Bourget et al. (2014) for finite-width beams, which, for clarity, will be referred to as ‘previous theory’ or ‘published literature’.

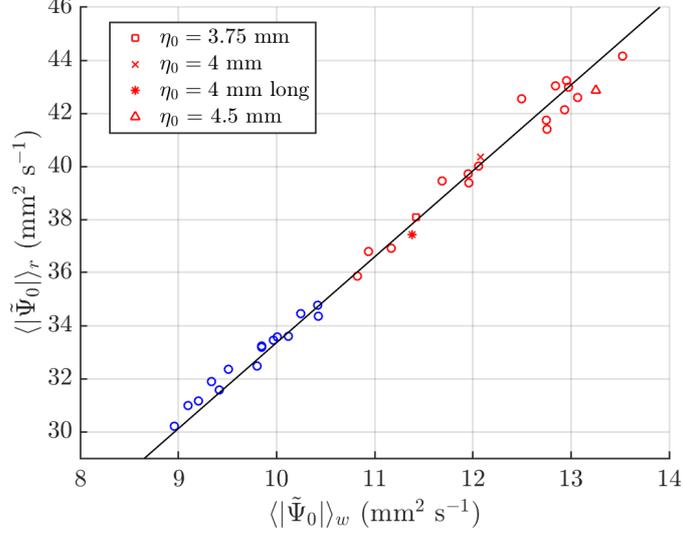


Figure 4.11: Amplitude of the primary wave for all 36 experiments given as in terms of both $\langle |\tilde{\Psi}_0| \rangle_w$ and $\langle |\tilde{\Psi}_0| \rangle_r$. As expected, this follows a linear fit given by the black line. The four experiments shown in Figure 4.5 are highlighted for comparison.

4.5.2 Amplitude development comparison

The first step of comparing the experimental results against the published literature is to understand how the latter describes the development of the instability. This is done by extending the \mathcal{M}_{0D} analysis of Bourget et al. (2014) that examines the weakly non-linear TRI interactions in the context of a finite-width beam. Using Matlab (Mathworks, 2021) the three coupled non-linear ODEs given in (2.23) were numerically integrated to see how the triadic waves develop over time. The ODEs have been repeated here as

$$\frac{d\tilde{\Psi}_0}{dt} = I_0 \tilde{\Psi}_1 \tilde{\Psi}_2 - \nu \left(\frac{\kappa_0^2}{2} \right) \tilde{\Psi}_0 + F, \quad (4.4a)$$

$$\frac{d\tilde{\Psi}_1}{dt} = I_1 \tilde{\Psi}_0 \tilde{\Psi}_2^* - \left(\frac{\nu \kappa_1^2}{2} + \frac{|\mathbf{c}_{g1} \cdot \mathbf{e}_{k_0}|}{2W} \right) \tilde{\Psi}_1, \quad (4.4b)$$

$$\frac{d\tilde{\Psi}_2}{dt} = I_2 \tilde{\Psi}_0 \tilde{\Psi}_1^* - \left(\frac{\nu \kappa_2^2}{2} + \frac{|\mathbf{c}_{g2} \cdot \mathbf{e}_{k_0}|}{2W} \right) \tilde{\Psi}_2, \quad (4.4c)$$

where again, the interaction I term is defined by

$$I_p = \frac{l_q m_r - m_q l_r}{2\omega_p \kappa_p^2} \left[\omega_p (\kappa_q^2 - \kappa_r^2) + l_p N^2 \left(\frac{l_q}{\omega_q} - \frac{l_r}{\omega_r} \right) \right], \quad (4.5)$$

and the forcing term $F = \mathbf{c}_{g_0}(\tilde{\Psi}_{\text{in}}^* \tilde{\Psi}_{\text{in}} - \tilde{\Psi}_0^* \tilde{\Psi}_0)/(2L\tilde{\Psi}_0^*)$. The diagram in Figure 2.2 depicts the interaction region of width W and length L . For the analysis presented here, W was taken as the width over the two middle wavelengths that reached maximum forcing amplitude η_0 (shown in Figure 4.1), making W and L , 250 mm and 532 mm respectively. Although related, the width W given in Figure 2.2 is different from wavebeam width Λ_0 , which, for our experimental setup is ≈ 500 mm. Based on the region defined in Figure 2.2, the width W does not span the full extent of the beam, so the decision was made to only consider the width of the two central wavelengths for this analysis. The wavenumbers and frequencies chosen for the interaction term (4.5) were based off experimentally determined values, specifically the ones described in § 4.3.1. For the primary beam, $\omega_0 = 0.95 \text{ rad s}^{-1}$, $(l_0, m_0) = (-0.05, -0.064) \text{ mm}^{-1}$, while for the resonant beams $\omega_1 = 0.358 \text{ rad s}^{-1}$, $(l_1, m_1) = (0.039, 0.156) \text{ mm}^{-1}$ and $\omega_2 = 0.592 \text{ rad s}^{-1}$, $(l_2, m_2) = (-0.090, -0.225) \text{ mm}^{-1}$. The buoyancy frequency was set at $N = 1.54 \text{ rad s}^{-1}$ and $|\tilde{\Psi}_{\text{in}}|$ ranged between $20 \sim 40 \text{ mm}^2 \text{ s}^{-1}$, comparable to the values for $\langle |\tilde{\Psi}_0| \rangle_r$ shown in Figure 4.11. The results of this analysis can be seen in Figure 4.12.

The first thing to note is that the lowest forcing amplitude of $|\tilde{\Psi}_{\text{in}}| = 20 \text{ mm}^2 \text{ s}^{-1}$, given in Figure 4.12(a), does not show any growth of the secondary waves. For the input parameters to the \mathcal{M}_{0D} model detailed above, the theoretical amplitude threshold was found to be approximately $|\tilde{\Psi}_{\text{in}}| = 25 \text{ mm}^2 \text{ s}^{-1}$. This is lower than the experimental threshold of $\langle |\tilde{\Psi}_0| \rangle_r \approx 35 \text{ mm}^2 \text{ s}^{-1}$, shown in Figure 4.11. This in itself is unsurprising, as various imperfections in the experimental stratification could increase the amplitude threshold required for instability. There are also many spatial aspects of the experiments that are not captured by the \mathcal{M}_{0D} model in (4.4).

The \mathcal{M}_{0D} model results then show that after the amplitude threshold has been surpassed an increase in forcing results in a faster growth of the instability. For the forcing amplitude $|\tilde{\Psi}_{\text{in}}| = 30 \text{ mm}^2 \text{ s}^{-1}$, shown in Figure 4.12(b), the instability occurs around 200 forcing periods, compared to approximately 100 forcing periods for $|\tilde{\Psi}_{\text{in}}| = 40 \text{ mm}^2 \text{ s}^{-1}$, shown in Figure 4.12(d). A similar trend, although less clear, is shown in the experimental results. The lowest amplitude forcing, shown in Figure 4.5(a) (which corresponds to $\langle |\tilde{\Psi}_0| \rangle_r \approx 38 \text{ mm}^2 \text{ s}^{-1}$) takes almost twice as long for the instability to grow compared to the highest amplitude forcing in Figure 4.5(d) (where $\langle |\tilde{\Psi}_0| \rangle_r \approx 43 \text{ mm}^2 \text{ s}^{-1}$). The other similarity between the \mathcal{M}_{0D} model and experimental results is the mean amplitude of the secondary waves. For both, $|\tilde{\Psi}_1|$ sits just below 20% of $|\tilde{\Psi}_0|$ and $|\tilde{\Psi}_2|$ is, to a good approximation, a scaled value of $|\tilde{\Psi}_1|$.

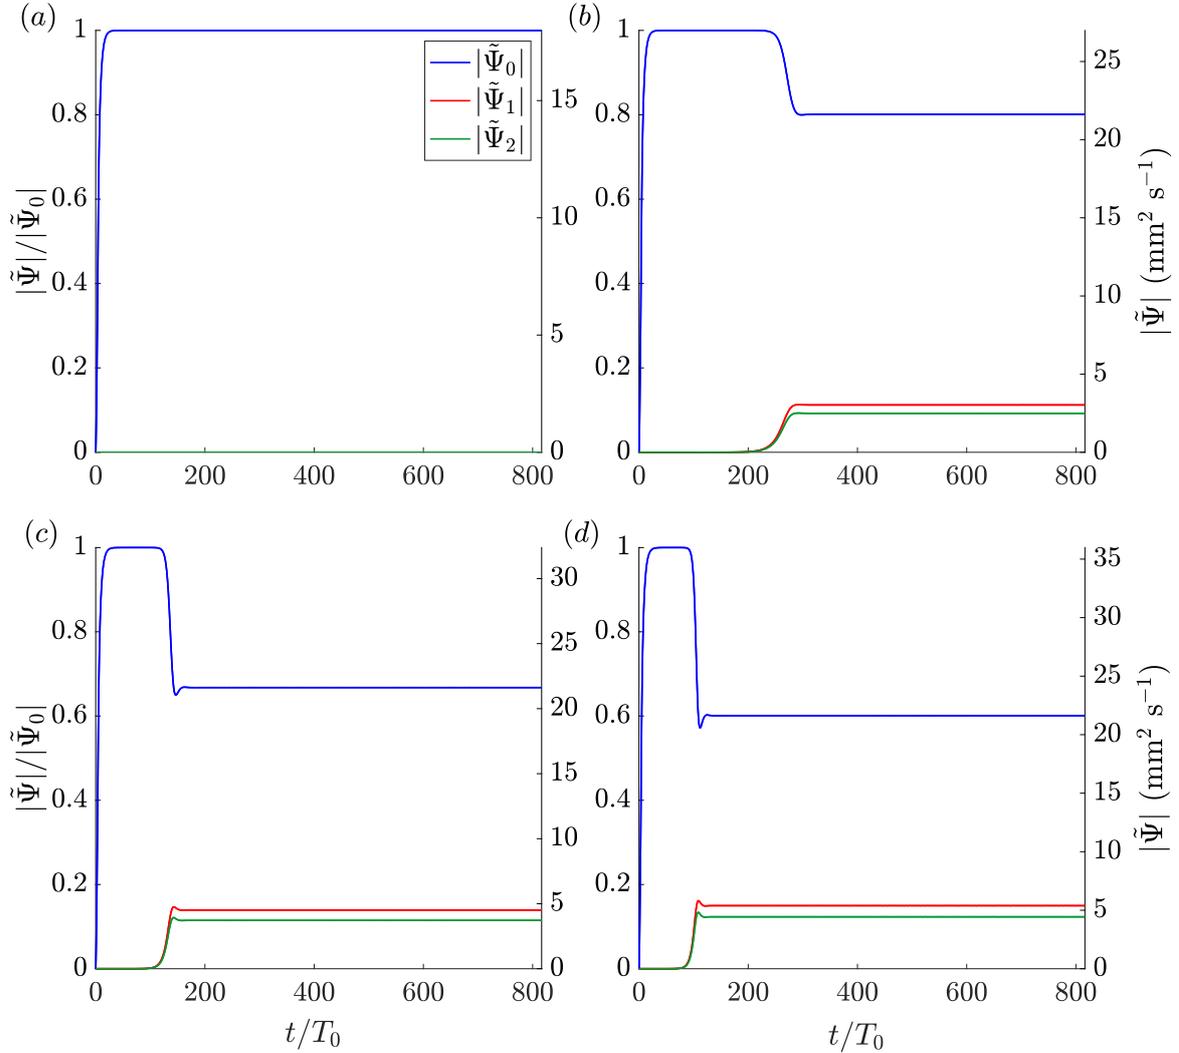


Figure 4.12: Results from the \mathcal{M}_{0D} model by Bourget et al. (2014) given in (4.4) for four different forcing amplitudes of the primary wave $|\tilde{\Psi}_{in}|$. The amplitude of the streamfunction for the triadic waves is shown as $|\tilde{\Psi}_0|$ (blue), $|\tilde{\Psi}_1|$ (red), $|\tilde{\Psi}_2|$ (green). The forcing amplitudes of the four plots are (a) $|\tilde{\Psi}_{in}| = 20 \text{ mm}^2 \text{ s}^{-1}$, (b) $|\tilde{\Psi}_{in}| = 30 \text{ mm}^2 \text{ s}^{-1}$, (c) $|\tilde{\Psi}_{in}| = 36 \text{ mm}^2 \text{ s}^{-1}$, (d) $|\tilde{\Psi}_{in}| = 40 \text{ mm}^2 \text{ s}^{-1}$.

This, however, is where the similarities between Figure 4.5 and Figure 4.12 cease. For the \mathcal{M}_{0D} model results, larger forcing engenders a significant decay in amplitude of the primary wavebeam. Interestingly, Figure 4.12 also shows that the amplitude of the primary wave $|\tilde{\Psi}_0|$ always reaches the same value, regardless of the initial forcing $|\tilde{\Psi}_{in}|$. This raises a question of whether this zero-dimensional scheme conserves energy in the system. This large decay of the primary beam is not witnessed experimentally, where the amplitude of the primary beam oscillates around a mean value given by the forcing from the wavemaker.

The most obvious difference between the two arises in their description of the long term development. The \mathcal{M}_{0D} model results predict that after the initial instability, all waves in the triad reach a steady state whereby the energy exchange from the primary wave to the resonant waves is constant. This is clearly not the case for the experimental results. The slow synchronous amplitude modulations seen experimentally reveal a continuous flux in the energy exchange between the primary wave and the two secondary waves corresponding to different perturbations growing and decaying over the course of the experiment. Interestingly, the theoretical results in Figure 4.12 show that for the three larger amplitudes that cause the primary wave to become unstable, a small oscillation occurs after the initial onset of the instability, however this is very quickly damped and triad returns to its equilibrium state. One thing to note in regards to the previous theory, is that the interaction term (4.5) only permits one set of fixed input parameters for the resonant waves. Already this is limiting the extent to which this analysis can describe the instability, as we have shown in § 4.3.3 that resonant wave parameters are also in a state of continuous flux.

4.5.3 Growth rate comparison

The next stage of the comparison is focused around how well the experimental results match the linear stability analysis given in the previous theory (see § 2.5.3). Again, we focus here on the comparison with the finite-width theory provided by Bourget et al. (2014) as this examines the growth rate for a finite-width beam. Solving the system of coupled ODEs in the \mathcal{M}_{0D} model provides an expression for the growth rate of the resonant wave $\tilde{\Psi}_1$ given in (2.24), repeated here as

$$\sigma_{0D\pm} = \frac{1}{4}(\Sigma_1 + \Sigma_2) \pm \sqrt{\frac{1}{16}(\Sigma_1 - \Sigma_2)^2 + |I_1||I_2||\tilde{\Psi}_0|^2}, \quad (4.6)$$

where $\Sigma_p = \nu\kappa_p^2 + |\mathbf{c}_{g_p} \cdot \mathbf{e}_{k_0}|/W$. From this expression it is possible to plot the linear growth rate σ_{0D} for all resonant waves that satisfy both the triadic resonant conditions

and the dispersion relationship, given by the expression in (2.5). The unstable solution branches to (2.5) are again shown in Figure 4.13 but with each branch coloured to correspond to the relative length l_1/l_0 .

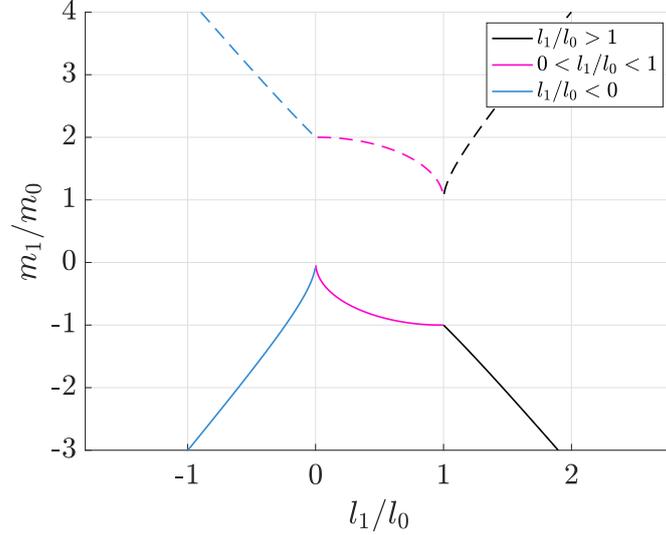


Figure 4.13: The same curve as shown in Figure 2.1, showing all the possible unstable solutions for the \mathbf{k}_1 wavenumber vector that satisfy (2.5) corresponding to the TRI condition (2.4), using primary wave input parameters $\omega_0 = 0.95 \text{ rad s}^{-1}$, $(l_0, m_0) = (-0.05, -0.064) \text{ mm}^{-1}$. The curve has now been split into different coloured sections based on the length of l_1/l_0 to correspond to the growth rate curves shown in Figure 4.14.

All of the possible resonant wave vectors from Figure 4.13 were then substituted into the linear growth rate expression given in (4.6). Figure 4.14 shows these growth rates as a function of both κ_1/κ_0 and ω_1/ω_0 based on $|\tilde{\Psi}_0| = 40 \text{ mm}^2 \text{ s}^{-1}$.

The first thing to note from the growth rate curves in Figure 4.14 is that the maximum value of σ_{0D} corresponds to wave parameters that lie on the central pink branch of Figure 4.13. From here on we will refer to the zero-dimensional linear growth rates of the solid pink branch and the solid blue branch as $\hat{\sigma}_{0D}$ and $\hat{\sigma}_{0D}$ respectively. The maximum values of these branches are therefore denoted $\hat{\sigma}_{0D_{\max}}$ and $\hat{\sigma}_{0D_{\max}}$.

If the pair of secondary wave vectors correspond to the solid pink branch, the wavenumbers will always satisfy $\kappa_1 < \kappa_0$, while $\kappa_2 > \kappa_0$. This means that the energy from the primary wave will be transferred to one larger wavelength and one smaller. Indeed, if the pair of secondary wave vectors correspond to the solid light blue branch there is also a possibility for transfer to both smaller and larger scales if $\kappa_1 < \kappa_0$, meaning $0 > m_1/m_0 > -1$. However, for all the experiments conducted, the tip of the \mathbf{k}_1 wave vector always lay on the outer solid light blue branch and always satisfied

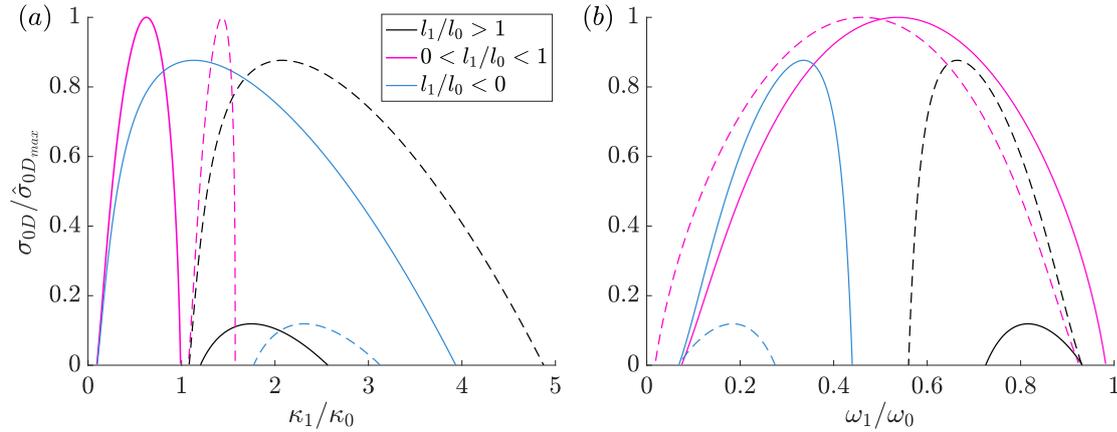


Figure 4.14: Theoretical growth rates calculated using (4.6) for all the resonant wave parameters that satisfy (2.5), using an input amplitude of $|\tilde{\Psi}_0| = 40 \text{ mm}^2 \text{ s}^{-1}$ and primary wave parameters $\omega_0 = 0.95 \text{ rad s}^{-1}$, $(l_0, m_0) = (-0.05, -0.064) \text{ mm}^{-1}$. The colour and style of each curve corresponds to the different branches indicated in Figure 4.13. (a) Shows the growth rate as a function of κ_1/κ_0 . (b) Shows the growth rate as a function of ω_1/ω_0 . All growth rates have been normalised by the maximum value $\hat{\sigma}_{0D_{max}}$, where $\hat{\sigma}_{0D}$ corresponds to the growth rates of the solid pink line.

$m_1/m_0 < -1$ (as shown in Figure 4.7). This means that while wave parameters on the solid light blue branch would also permit a transfer to larger scales, this never occurred experimentally.

To understand how the experimental linear growth rates compare to the theoretical values of σ_{0D} , Figure 4.15(a) and (c) show the log-linear plots of the two experiments shown in Figure 4.5(b) and (d) (with $\eta_0 = 4$ and 4.5 mm , respectively), now with a linear fit for the initial exponential growth of \mathbb{W}_1 shown by the red dashed line. From the slope of this linear fit the experimental growth rate σ_{exp} was extracted. The two growth rate curves shown in Figure 4.15(b) and (d) correspond to the log-linear amplitude plots in (a) and (c) respectively. It is necessary to have separate growth rate curves for the two experiments in Figure 4.15(a) and (c), as they have different forcing amplitudes $|\tilde{\Psi}_0|$, approximated by $\langle |\tilde{\Psi}_0| \rangle_r$. For Figure 4.15(b) $\langle |\tilde{\Psi}_0| \rangle_r = 40 \text{ mm}^2 \text{ s}^{-1}$ and for Figure 4.15(d) $\langle |\tilde{\Psi}_0| \rangle_r = 43 \text{ mm}^2 \text{ s}^{-1}$ determined by the experimental values shown in Figure 4.11. The extracted growth rate σ_{exp} , is marked by a red star and should be compared against the theoretical $\hat{\sigma}_{0D}$ for secondary wave parameters corresponding to solid blue curve (as this was the experimentally accessed branch). The grey patch on Figure 4.15(b) and (d) highlights the experimental range of κ_1/κ_0 .

We first examine, not the extracted experimental growth rates σ_{exp} , but the values of κ_1/κ_0 that they corresponds to. For both experiments, the value of κ_1/κ_0 corresponding to σ_{exp} lie to the right of $\hat{\sigma}_{0D_{max}}$. This was also found to be the case for all the other

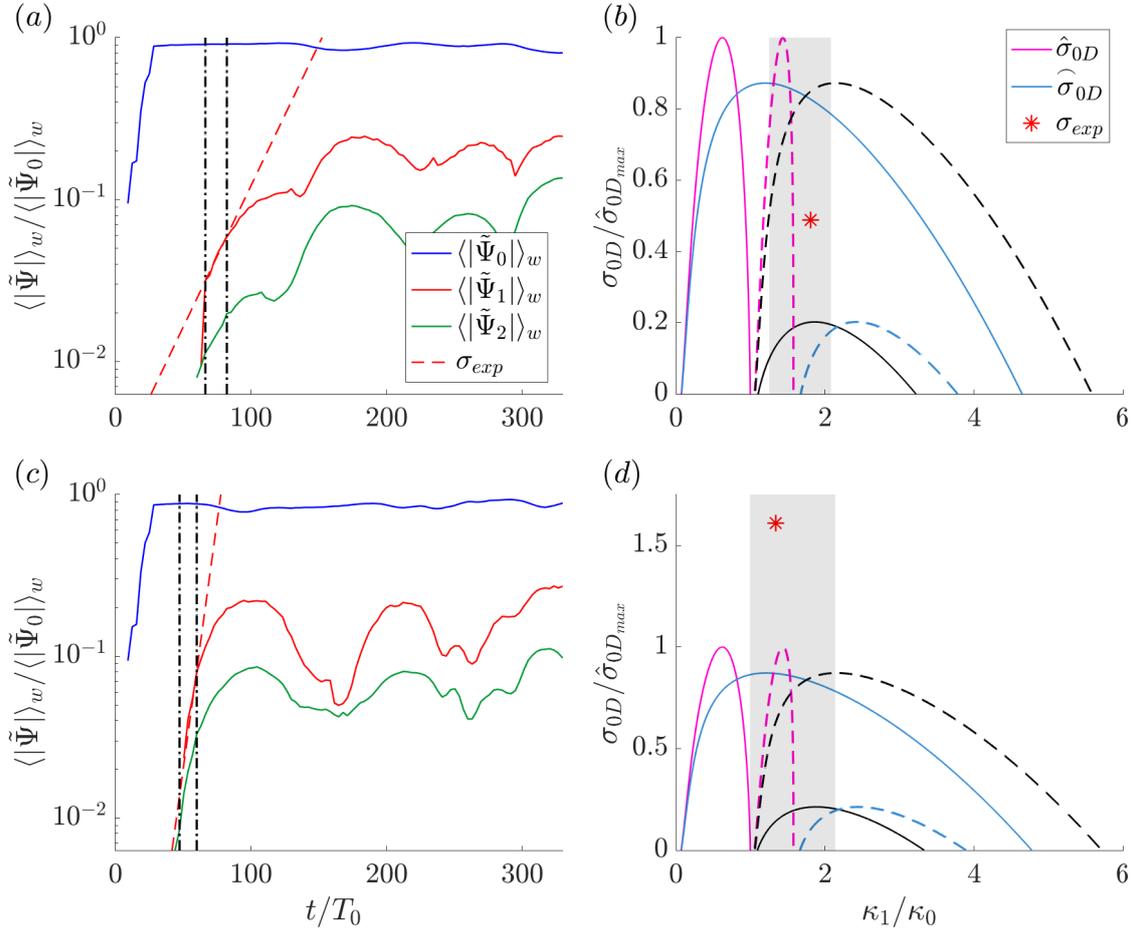


Figure 4.15: Log-linear amplitude plots of the triadic waves for two different experiments with their corresponding σ_{0D} growth rate curves. (a) Shows the same experiment presented in Figure 4.5(b) with $\eta_0 = 4$ mm. (c) Shows the same experiment presented in Figure 4.5(d) with $\eta_0 = 4.5$ mm. The red dashed line in (a) and (c) shows the linear fit from which the experimental value of the growth rate σ_{exp} was extracted, between the two dashed black vertical lines. (b) and (d) show the corresponding σ_{0D} theoretical growth rates of all the possible wavenumbers given from Figure 4.13 obtained from (4.6). The red star gives the extracted growth rate σ_{exp} from the slope of the linear fit on the left hand log-linear plots. The two grey shaded region indicates the range that κ_1/κ_0 took during the respective experiments on the left. The amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_r$ for (b) is $40 \text{ mm}^2 \text{ s}^{-1}$ and for (d) $\langle |\tilde{\Psi}_0| \rangle_r$ is $43 \text{ mm}^2 \text{ s}^{-1}$ determined by the experimental values shown in Figure 4.11.

experimental results examined (not shown here). This is unexpected, as we would predict the experimentally selected value of κ_1 to correspond to $\widehat{\sigma}_{0D_{\max}}$. Again, this off-set could be related to some of the properties of the finite-width beam – not captured in (4.6) – which modify the most unstable wavenumbers. This is further discussed in Chapter 6.

The actual value of the linear growth rate for the two experiments varies significantly. The observed σ_{exp} shown in Figure 4.15(b) lies at approximately half of its corresponding $\widehat{\sigma}_{0D}$, while σ_{exp} in Figure 4.15(d) is nearly twice its corresponding $\widehat{\sigma}_{0D}$. What must be considered here is the accuracy of the experimental results. Experimentally, we are unable to resolve amplitudes below $\langle |\tilde{\Psi}| \rangle_w \sim 10^{-2} \text{ mm}^2 \text{ s}^{-1}$. A question therefore remains as to whether these extracted growth rates actually correspond to the period of exponential linear growth or if we are unable to accurately capture this. Indeed, fitting this curve at earlier times in Figure 4.15(a) would result in a significantly steeper gradient and higher values for σ_{exp} . However, due to the uncertainty in these experimental measurements, this does not seem an appropriate choice. Again, this is discussed further in Chapter 6 when we examine the results from the two-dimensional weakly non-linear modelling.

4.6 Summary

In this chapter we have presented experimental results examining the long term development of TRI for a finite-width beam. The results showed that over the course of 90 minute experiments the triadic waves exhibited a continuous variation in the flux of energy from the primary wave to the secondary waves, resulting in slow coupled amplitude modulations of the triad. Snapshots of the experiments showed that these modulations were not just temporal but also spatial, with the interaction region of the triadic waves continuously varying across the full central region of the primary wave. We believe this oscillatory behaviour to be related to the finite-width of the primary beam. As a packet of energy in the secondary waves exits the underlying primary beam the energy exchange is broken, resulting in an increase in amplitude for the primary beam and decrease for the secondary waves. This increase in amplitude of the primary beam then allows further perturbations to grow, thereby repeating the quasi-periodic cycle. Further analysis revealed that these modulations were also present in Fourier space, as the frequency and wavenumber of the resonant waves were also seen to adjust with a similar period to that of the amplitude. Despite this modulation to the resonant wave parameters, the triadic resonance conditions were

continuously satisfied. These results revealed that the different perturbations selected slightly different secondary wave parameters. Based on these results we proposed that these different frequencies and wavenumbers could then interact with each other, resulting in an off-resonant ‘beating’ via non-linear interactions through the primary wave. These experimental results contrasted with the previous theory, which predicts that after the initial instability the triadic waves reach a steady state amplitude. The following chapter presents the two-dimensional weakly non-linear modelling that was constructed in order to better understand this discrepancy.

Chapter 5

A weakly non-linear approach to Triadic Resonance Instability

5.1 Introduction

In this chapter we focus on the development of the two-dimensional weakly non-linear model that was built to capture the evolution of TRI in finite-width beams. The motivation here lies in the contrast between the previous theory, presented in § 2.5.3, and the experimental results that were presented in Chapter 4. As shown in § 4.5, the published theory does not agree with our experimental observations of TRI, particularly in how the energy exchange between waves in the triad is described.

The published literature that outlines the development of TRI in the context of monochromatic plane waves is given by Koudella and Staquet (2006) and Bourget et al. (2013). Later, Bourget et al. (2014), extends this theory to the context of finite-width beams (see § 2.5.3) however, their analysis is incomplete. As the form of the streamfunction presented in Bourget et al. (2014) is zero-dimensional in space, it does not account for the spatial gradients over which the triadic interactions are occurring. Moreover, how the spatial variability of the primary wave beam will affect the energy exchange between the triad, is not represented. The aim was therefore to expand the model to incorporate two spatial dimensions to better understand how the finite-width of the primary beam affects the evolution of TRI.

The decision was made early on to further explore the weakly non-linear interactions as opposed to implementing a full Computational Fluid Dynamics (CFD) code. The goal here is to dissect the experiments and to isolate the dynamics that are being witnessed, in order to improve the understanding of the system. A CFD code would

have been an inappropriate choice to achieve this, as little would be learnt about the physical mechanisms governing the behaviour witnessed experimentally.

This chapter outlines the development of the two-dimensional model, which is comprised of both the weakly non-linear theoretical equations and the numerical code used to solve the system. To compare against the zero-dimensional \mathcal{M}_{0D} of Bourget et al. (2014), we will refer to this two-dimensional model as \mathcal{M}_{2D} . We start with the perturbation expansion in § 5.2, which presents the non-dimensional parameters and terms associated with consecutive orders in the expansion. The overall structure of the code used to numerically solve these equations is given in § 5.3. The following two sections, § 5.4 and § 5.5, then present how the leading order advection equations and the non-linear interactions were built into the numerical code. Finally, the code is verified in § 5.6.

5.2 Perturbation expansion

We start by assuming two-dimensional continuously stratified Boussinesq fluid with constant buoyancy frequency N . Using the same procedure outlined in § 2.5.3, the full momentum equations of (1.1) and mass conservation of (1.9) are converted to the single non-linear equation (2.14), repeated here as

$$\frac{D^2}{Dt^2}(\nabla^2\Psi) - \nu\frac{D}{Dt}\nabla^2(\nabla^2\Psi) + N^2\frac{\partial^2\Psi}{\partial x^2} = \frac{g}{\varrho_0}\left(\frac{\partial}{\partial x}\frac{D}{Dt} - \frac{D}{Dt}\frac{\partial}{\partial x}\right)\rho. \quad (5.1)$$

It is then necessary to find appropriate forms for the perturbation density ρ and streamfunction Ψ . This is done through the perturbation expansion outlined below, which examines both the spatial and temporal structure of the beam. We start by defining the three non-dimensional parameters

$$\epsilon = (|\tilde{\Psi}_0|\kappa_0^2)N^{-1}, \quad (5.2)$$

$$\gamma = (\kappa_0\Lambda_0)^{-1}, \quad (5.3)$$

$$\chi = (\nu\kappa_0)\mathbf{c}_{g_0}^{-1}, \quad (5.4)$$

where again, $|\tilde{\Psi}_0|$ is the magnitude of the primary wave streamfunction, κ_0 is the magnitude of a primary wavenumber vector \mathbf{k}_0 , \mathbf{c}_{g_0} is its group velocity, N is the buoyancy frequency, Λ_0 is the horizontal width of the primary beam and ν is the kinematic viscosity of water. Here ϵ can be viewed as a non-dimensional measure of the

beam amplitude, which characterises the relative importance of the non-linear $\mathbf{u} \cdot \nabla$ terms in the momentum equations (1.1). The spatial parameter γ , then defines the spatial separation between the wavelengths in the triad compared with the width of the primary wave beam. Finally, χ is the inverse of the Reynolds number given in (1.4). This term quantifies the relative effects of viscosity in the system. We shall restrict our attention to $\epsilon \ll 1$, $\gamma \ll 1$ and $\chi \ll 1$, and introduce re-scaled time and position through the six variables

$$\begin{aligned} \tau_\gamma &= \gamma t, & \xi &= \gamma x, & \zeta &= \gamma z, \\ \tau_\epsilon &= \epsilon t, & \tau_\chi &= \chi t, & \tau_{\epsilon\gamma} &= \epsilon\gamma t. \end{aligned} \quad (5.5)$$

As we shall see, τ_ϵ accounts for the ‘slow non-linear time’ variations to the amplitude while τ_γ governs the ‘slow advection time’ scale. The combination of ϵ and γ in $\tau_{\epsilon\gamma}$ is then introduced to account for higher order derivatives that appear in the expansion. The final timescale given by τ_χ represents the viscous decay of the wave beams. The spatial parameters ξ and ζ account for the gradual spatial variability of the triadic waves in the domain and whether the secondary waves are located within the spatial confines of the primary beam. Utilising the dimensionless amplitude ϵ , we re-write the streamfunction and density as

$$\Psi = \tilde{\Psi}(x, z, t)e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)} = \epsilon\check{\Psi}(\tau_\gamma, \xi, \zeta, \tau_\epsilon, \tau_{\epsilon\gamma}, \tau_\chi)e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}, \quad (5.6)$$

$$\rho = \tilde{\rho}(x, z, t)e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)} = \epsilon\check{\rho}(\tau_\gamma, \xi, \zeta, \tau_\epsilon, \tau_{\epsilon\gamma}, \tau_\chi)e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}, \quad (5.7)$$

where $\check{\Psi} = \epsilon\tilde{\Psi}$ and $\check{\rho} = \epsilon\tilde{\rho}$. We therefore require $|\check{\Psi}| \sim 1$ as $\check{\Psi} = \epsilon\tilde{\Psi} = (|\tilde{\Psi}_0|\kappa_0/N)\check{\Psi}$ and we are interested in small perturbations to the flow. The ‘fast’ phase time scales of the wave are captured by the complex exponential wave form $e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}$. By separating the reduced streamfunction and density using these different space and time scales, the partial derivatives in (5.1) with respect to (x, z, t) can be written as

$$\frac{\partial \Psi}{\partial t} = \epsilon \left(-i\omega + \gamma \frac{\partial}{\partial \tau_\gamma} + \epsilon \frac{\partial}{\partial \tau_\epsilon} + \epsilon\gamma \frac{\partial}{\partial \tau_{\epsilon\gamma}} + \chi \frac{\partial}{\partial \tau_\chi} + \dots \right) \check{\Psi} e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}, \quad (5.8a)$$

$$\frac{\partial \Psi}{\partial x} = \epsilon \left(il + \gamma \frac{\partial}{\partial \xi} + \dots \right) \check{\Psi} e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}, \quad (5.8b)$$

$$\frac{\partial \Psi}{\partial z} = \epsilon \left(im + \gamma \frac{\partial}{\partial \zeta} + \dots \right) \check{\Psi} e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}. \quad (5.8c)$$

$$\frac{\partial \rho}{\partial t} = \epsilon \left(-i\omega + \gamma \frac{\partial}{\partial \tau_\gamma} + \epsilon \frac{\partial}{\partial \tau_\epsilon} + \epsilon \gamma \frac{\partial}{\partial \tau_{\epsilon\gamma}} + \chi \frac{\partial}{\partial \tau_\chi} + \dots \right) \check{\rho} e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}, \quad (5.9a)$$

$$\frac{\partial \rho}{\partial x} = \epsilon \left(il + \gamma \frac{\partial}{\partial \xi} + \dots \right) \check{\rho} e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}, \quad (5.9b)$$

$$\frac{\partial \rho}{\partial z} = \epsilon \left(im + \gamma \frac{\partial}{\partial \zeta} + \dots \right) \check{\rho} e^{i(\mathbf{k} \cdot \mathbf{x} - \omega t)}. \quad (5.9c)$$

As we are also interested in the wave triad, we express the streamfunction and density as the summation

$$\Psi = \sum_{p=0}^2 \epsilon \check{\Psi}_p(\tau_\gamma, \xi, \zeta, \tau_\epsilon, \tau_{\epsilon\gamma}, \tau_\chi) e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)} + \text{c.c.}, \quad (5.10)$$

$$\rho = \sum_{p=0}^2 \epsilon \check{\rho}_p(\tau_\gamma, \xi, \zeta, \tau_\epsilon, \tau_{\epsilon\gamma}, \tau_\chi) e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)} + \text{c.c.}, \quad (5.11)$$

where the subscript p indicates either \mathbb{W}_0 , \mathbb{W}_1 or \mathbb{W}_2 . The perturbation expansions given in (5.8) and (5.9) are then substituted into the summations (5.10) and (5.11) respectively. Using these forms, Ψ and ρ are further substituted into the non-linear momentum equation (5.1). This substitution and following manipulations were performed with the aid of Mathematica (Wolfram Research, 2021) due to the lengthy algebra involved.

We then use the linear relationship (2.17) – repeated here in terms of $\check{\Psi} = \epsilon \check{\Psi}$ and $\check{\rho} = \epsilon \check{\rho}$ as

$$\check{\rho} = -\frac{l}{\omega} \frac{N^2 \rho_0}{g} \check{\Psi}, \quad (5.12)$$

– to remove density from (5.1), as, at leading order in ϵ , the right hand side (RHS) of (5.1) cancels. This means that as the contributions from density on the RHS only appear at second-order in ϵ , using a linear expression to express ρ in terms of Ψ is valid up until this order. The resultant expression obtained provides the Boussinesq viscous equations of motion solely as a function of Ψ that can be used to examine both linear and non-linear dynamics between a triadic set of waves. This was done by collecting around subsequent orders of ϵ , γ and χ . Using the experimental parameters outlined

in § 4.2, the relative amplitude of the non-dimensional parameters are given as

$$1 \gg \underbrace{\epsilon}_{\sim 10^{-1}} \gg \underbrace{\epsilon^2}_{\sim 10^{-2}} \gg \dots \quad (5.13a)$$

$$1 \gg \underbrace{\gamma}_{\sim 10^{-1}} \gg \underbrace{\gamma^2}_{\sim 10^{-2}} \gg \dots \quad (5.13b)$$

$$1 \gg \underbrace{\chi}_{\sim 10^{-2}} \gg \underbrace{\chi^2}_{\sim 10^{-4}} \gg \dots \quad (5.13c)$$

Experimentally, both ϵ and γ vary slightly depending on the amplitude of the primary wave or width of the beam. The inverse Reynolds number χ remains constant as we do not vary κ_0 , ω_0 or N (hence \mathbf{c}_{g_0} is approximately constant). In oceanographic settings, however, χ will be orders of magnitude smaller than ϵ and γ , as the role of viscosity in the ocean can be considered negligible. Yet, as we are looking at the experimental results, we do not neglect the leading order effects of viscosity.

The first set of terms at $\mathcal{O}(\epsilon^0\gamma^0)$ will be equal to zero as this leading order corresponds to a state of rest. We then consider the next order $\mathcal{O}(\epsilon^1\gamma^0)$. For these terms, as ϵ is still at first-order, the non-linearity in (5.1) cancels. Extracting terms around $e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)}$, (where $p = 0, 1, 2$ for each of the waves in the triad), leaves

$$\frac{\omega_p}{N} = \frac{l_p}{\sqrt{(l_p^2 + m_p^2)}}, \quad (5.14)$$

which is the linear dispersion relationship for internal plane waves. As this expression satisfies the full non-linear momentum equations, it is also valid for a single non-linear plane wave.

We then collect around the next order $\mathcal{O}(\epsilon^1\gamma^1)$. Again, as ϵ is still at first-order, the non-linear terms in (5.1) cancel. Looking at the terms around $e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)}$, we obtain

$$\gamma \frac{\partial \check{\Psi}_p}{\partial \tau_\gamma} = \gamma \left(\frac{l_p(\omega_p^2 - N^2)}{\omega_p \kappa_p^2} \frac{\partial \check{\Psi}_p}{\partial \xi} + \frac{\omega_p m_p}{\kappa_p^2} \frac{\partial \check{\Psi}_p}{\partial \zeta} \right), \quad (5.15)$$

which after some re-arranging can be expressed as

$$\gamma \frac{\partial \check{\Psi}_p}{\partial \tau_\gamma} = -\gamma (\mathbf{c}_{g_p} \cdot \nabla_\xi) \check{\Psi}_p, \quad (5.16)$$

where $\nabla_\xi = (\partial/\partial \xi, \partial/\partial \zeta)$. This linear advection equation shows that the streamfunction of each wave in the triad is advected at its respective group velocity. It is already well

known that, for small amplitude internal waves, the group velocity \mathbf{c}_g is the velocity at which energy is transported (e.g. Sutherland, 2010). As energy scales as $\sim \Psi^2$, the advection equation (5.16) is, therefore, not unsurprising.

Before examining the non-linear interaction terms at $\mathcal{O}(\epsilon^2)$, the viscous term at $\mathcal{O}(\epsilon\chi)$ needs to be considered in conjunction with the linear advection. While in the experiments (and ocean settings) these terms are of lower magnitude than the terms at $\mathcal{O}(\epsilon^2)$, they govern the viscous decay of individual waves irrespective of the non-linear interactions, making it appropriate to consider them at the same order as the advection. At $\mathcal{O}(\epsilon\chi)$, again collecting around $e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)}$, we obtain

$$\chi \frac{\partial \check{\Psi}_p}{\partial \tau_\chi} = -\frac{\mathbf{c}_{g0}}{\kappa_0} \frac{\chi}{2} \kappa_p^2 \check{\Psi}_p, \quad (5.17)$$

which shows that the viscous decay of each beam p scales as κ_p^2 . We then combine the two times scales τ_γ and τ_χ by considering the terms at $\mathcal{O}(\epsilon\gamma)$ and $\mathcal{O}(\epsilon\chi)$ together. At first-order in ϵ , the advection equation can be written as

$$\frac{\partial \tilde{\Psi}_p}{\partial t} = -(\mathbf{c}_{gp} \cdot \nabla_{\mathbf{x}}) \tilde{\Psi}_p - \frac{\nu}{2} \kappa_p^2 \tilde{\Psi}_p, \quad (5.18)$$

where $\nabla_{\mathbf{x}} = (\partial/\partial x, \partial/\partial z)$. As the effects of viscosity are now accounted for in the advection equation (5.18), subsequent terms involving χ at higher orders ($\mathcal{O}(\epsilon\gamma\chi)$, $\mathcal{O}(\epsilon^2\chi)$ and $\mathcal{O}(\epsilon^2\gamma\chi)$) were then neglected.

The next stage was to consider terms at order $\mathcal{O}(\epsilon^2\gamma^0)$. Here, terms are no longer only associated with $e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)}$, but are also comprised of cross-terms from the non-linear terms $\mathbf{u} \cdot \nabla$ in (5.1), for example $e^{i(\mathbf{k}_q \cdot \mathbf{x} - \omega_q t) + i(\mathbf{k}_r \cdot \mathbf{x} - \omega_r t)}$. It is, therefore, necessary to consider the terms that sum to one of the triadic waves based on the triadic resonance conditions (2.4). In addition, due to these non-linear cross terms, the RHS of (5.1) no longer cancels so the contribution of these terms must also be included. At $\mathcal{O}(\epsilon^2)$ we recover the same non-linear interactions given in (2.19), given as

$$\epsilon^2 \frac{\partial \check{\Psi}_0}{\partial \tau_\epsilon} = \epsilon^2 I_0 \check{\Psi}_1 \check{\Psi}_2 \rightarrow \frac{\partial \tilde{\Psi}_0}{\partial t} = I_0 \tilde{\Psi}_1 \tilde{\Psi}_2, \quad (5.19a)$$

$$\epsilon^2 \frac{\partial \check{\Psi}_1}{\partial \tau_\epsilon} = \epsilon^2 I_1 \check{\Psi}_0 \check{\Psi}_2^* \rightarrow \frac{\partial \tilde{\Psi}_1}{\partial t} = I_1 \tilde{\Psi}_0 \tilde{\Psi}_2^*, \quad (5.19b)$$

$$\epsilon^2 \frac{\partial \check{\Psi}_2}{\partial \tau_\epsilon} = \epsilon^2 I_2 \check{\Psi}_0 \check{\Psi}_1^* \rightarrow \frac{\partial \tilde{\Psi}_2}{\partial t} = I_2 \tilde{\Psi}_0 \tilde{\Psi}_1^*, \quad (5.19c)$$

where I_p is the interaction term defined by

$$I_p = \frac{l_q m_r - m_q l_r}{2\omega_p \kappa_p^2} \left[\omega_p (\kappa_q^2 - \kappa_r^2) + l_p N^2 \left(\frac{l_q}{\omega_q} - \frac{l_r}{\omega_r} \right) \right]. \quad (5.20)$$

The form of these PDEs in (5.19) is same as those presented in Bourget et al. (2013), although without the viscous term, as this has been accounted for the advection equation at $\mathcal{O}(\epsilon\chi)$. For the majority of the results that are presented in Chapter 6, the solution is truncated at this order and the model simply uses (5.19) to describe the weakly non-linear interactions between the waves.

It is, however, also possible to consider the results of the expansion at the subsequent orders $\mathcal{O}(\epsilon^2\gamma^1)$ and $\mathcal{O}(\epsilon^1\gamma^2)$. Looking first at $\mathcal{O}(\epsilon^1\gamma^2)$, we find that again the non-linear terms cancel, leaving the second-order linear PDE

$$\left(\frac{\partial^2}{\partial \tau_\gamma^2} - \frac{N^2}{\kappa_p^2} \frac{\partial^2}{\partial \zeta^2} - \frac{\omega_p^2}{\kappa_p^2} \nabla_\xi^2 - \frac{4\omega_p}{\kappa_p^2} (\mathbf{k}_p \cdot \nabla_\xi) \frac{\partial}{\partial \tau_\gamma} \right) \check{\Psi}_p = 0, \quad (5.21)$$

which in terms of $\tilde{\Psi}$ becomes

$$\left(\frac{\partial^2}{\partial t^2} - \frac{N^2}{\kappa_p^2} \frac{\partial^2}{\partial z^2} - \frac{\omega_p^2}{\kappa_p^2} \nabla_x^2 - \frac{4\omega_p}{\kappa_p^2} (\mathbf{k}_p \cdot \nabla_x) \frac{\partial}{\partial t} \right) \tilde{\Psi}_p = 0. \quad (5.22)$$

The form of this expression closely resembles the second-order wave equation, although with additional cross-derivative terms (e.g. $\partial^2 \tilde{\Psi} / \partial x \partial t$). Unlike the linear advection in (5.16), which only permits one direction of propagation, this second-order expression would also permit backwards propagation of waves. While this may more accurately capture the dynamics of the experiment, it was considered less significant to understanding the dynamics of TRI and was therefore not built into the model.

The last order to be considered was $\mathcal{O}(\epsilon^2\gamma^1)$. Specifically at this order we are left with the following system of coupled non-linear PDEs

$$\frac{\partial \check{\Psi}_0}{\partial \tau_{\epsilon\gamma}} = \frac{1}{2i\kappa_0^2\omega_0} \left[\check{\Psi}_2(\boldsymbol{\alpha}_{1,2} \cdot \nabla_{\boldsymbol{\xi}} \check{\Psi}_1) + \check{\Psi}_1(\boldsymbol{\alpha}_{2,1} \cdot \nabla_{\boldsymbol{\xi}} \check{\Psi}_2) + \frac{\partial \check{\Psi}_2}{\partial \tau_{\gamma}} \check{\Psi}_1 \beta_{2,1} + \frac{\partial \check{\Psi}_1}{\partial \tau_{\gamma}} \check{\Psi}_2 \beta_{1,2} \right] \quad (5.23a)$$

$$\frac{\partial \check{\Psi}_1}{\partial \tau_{\epsilon\gamma}} = \frac{1}{2i\kappa_1^2\omega_1} \left[\check{\Psi}_0(\boldsymbol{\alpha}_{-2,0} \cdot \nabla_{\boldsymbol{\xi}} \check{\Psi}_2^*) + \check{\Psi}_2^*(\boldsymbol{\alpha}_{0,-2} \cdot \nabla_{\boldsymbol{\xi}} \check{\Psi}_0) + \frac{\partial \check{\Psi}_2^*}{\partial \tau_{\gamma}} \check{\Psi}_0 \beta_{-2,0} + \frac{\partial \check{\Psi}_0}{\partial \tau_{\gamma}} \check{\Psi}_2^* \beta_{0,-2} \right] \quad (5.23b)$$

$$\frac{\partial \check{\Psi}_2}{\partial \tau_{\epsilon\gamma}} = \frac{1}{2i\kappa_2^2\omega_2} \left[\check{\Psi}_0(\boldsymbol{\alpha}_{-1,0} \cdot \nabla_{\boldsymbol{\xi}} \check{\Psi}_1^*) + \check{\Psi}_1^*(\boldsymbol{\alpha}_{0,-1} \cdot \nabla_{\boldsymbol{\xi}} \check{\Psi}_0) + \frac{\partial \check{\Psi}_1^*}{\partial \tau_{\gamma}} \check{\Psi}_0 \beta_{-1,0} + \frac{\partial \check{\Psi}_0}{\partial \tau_{\gamma}} \check{\Psi}_1^* \beta_{0,-1} \right] \quad (5.23c)$$

where the vector $\boldsymbol{\alpha}$ represents a spatial interaction term

$$\boldsymbol{\alpha}_{q,r} = \left[\left((3l_q^2 m_r - 2l_q l_r m_q + m_q^2 m_r)(2\omega_q + \omega_r) - (m_r^3 + l_r^2 m_r)(2\omega_r + \omega_q) + \kappa_q \kappa_r m_r (\omega_r - 2\omega_q) + N l_r m_q \kappa_r \right), \right. \\ \left. \left(- (3l_r m_q^2 - 2l_q m_q m_r + l_q^2 l_r)(2\omega_q + \omega_r) + (l_r^3 + l_r m_r^2)(2\omega_r + \omega_q) + \kappa_r \kappa_q (\omega_q - \omega_r) \right) \right], \quad (5.24)$$

and β represents the scalar co-efficient for the time derivative terms from the linear advection τ_{γ}

$$\beta_{q,r} = -l_r^2 l_q m_r + 2l_q^3 m_r - l_q m_r^3 + l_r^3 m_q - 2l_r l_q^2 m_q + l_r m_r^2 m_q + 2l_q m_r m_q^2 - 2l_r m_q^3. \quad (5.25)$$

For both $\boldsymbol{\alpha}$ and β the subscripts (p, q, r) correspond to the circular permutations of the triad waves $(0, 1, 2)$ respectively for each equation. In the case where either p, q or r is negative – as they are in (5.23b) and (5.23c) – the negative values for both the wavenumber and frequency are required for that wave.

It is possible to substitute the linear advection equation obtained in (5.16) into (5.23). Writing in the form of $\tilde{\Psi}$ leaves

$$\frac{\partial \tilde{\Psi}_0}{\partial t} = \frac{1}{2i\kappa_0^2\omega_0} \left[\tilde{\Psi}_2 \left((\boldsymbol{\alpha}_{1,2} + \beta_{1,2}\mathbf{c}_{g_1}) \cdot \nabla_{\mathbf{x}} \tilde{\Psi}_1 \right) + \tilde{\Psi}_1 \left((\boldsymbol{\alpha}_{2,1} + \beta_{2,1}\mathbf{c}_{g_2}) \cdot \nabla_{\mathbf{x}} \tilde{\Psi}_2 \right) \right] \quad (5.26a)$$

$$\frac{\partial \tilde{\Psi}_1}{\partial t} = \frac{1}{2i\kappa_1^2\omega_1} \left[\tilde{\Psi}_0 \left((\boldsymbol{\alpha}_{-2,0} + \beta_{-2,0}\mathbf{c}_{g_2}) \cdot \nabla_{\mathbf{x}} \tilde{\Psi}_2^* \right) + \tilde{\Psi}_2^* \left((\boldsymbol{\alpha}_{0,-2} + \beta_{0,-2}\mathbf{c}_{g_0}) \cdot \nabla_{\mathbf{x}} \tilde{\Psi}_0 \right) \right] \quad (5.26b)$$

$$\frac{\partial \tilde{\Psi}_2}{\partial t} = \frac{1}{2i\kappa_2^2\omega_2} \left[\tilde{\Psi}_0 \left((\boldsymbol{\alpha}_{-1,0} + \beta_{-1,0}\mathbf{c}_{g_1}) \cdot \nabla_{\mathbf{x}} \tilde{\Psi}_1^* \right) + \tilde{\Psi}_1^* \left((\boldsymbol{\alpha}_{0,-1} + \beta_{0,-1}\mathbf{c}_{g_0}) \cdot \nabla_{\mathbf{x}} \tilde{\Psi}_0 \right) \right] \quad (5.26c)$$

What is left is a set of PDEs that combines the interactions between the triad in both Fourier and physical space. The PDEs are also complex, meaning they would result in a slow phase shift of the complex amplitude of the waves. The other thing to note is that there is also a resemblance to the linear advection equation, where the term $(\boldsymbol{\alpha}_{q,r} + \beta_{q,r}\mathbf{c}_{g_q})$ generates a modified group velocity. However, an investigation into the directionality of this term did not reveal any obvious meaning. The addition of these $\mathcal{O}(\epsilon^2\gamma)$ terms in to the model is briefly discussed in § 6.4.

We have now outlined the linear and weakly non-linear equations obtained from the perturbation expansion given in (5.6) and (5.7) when substituted into the full non-linear equation of motion. The next step is to consider how these equations are solved numerically.

5.3 Code development

5.3.1 Overview

This section outlines the overall structure of the numerical code developed to solve the weakly non-linear theoretical equations obtained in § 5.2 expressed in terms of $\tilde{\Psi}$. The code executes both the linear advection (5.18) and the weakly non-linear interactions (5.19) in one time-step based on the partial time derivative in (5.8a). First, § 5.3.2 outlines the form of the numerical domains, while § 5.3.3 provides details on the overall

code structure. Specific details of the advection scheme and the non-linear interactions are then given in § 5.4 and § 5.5 respectively. The code has been formulated to run in Matlab (Mathworks, 2021).

5.3.2 Numerical domains

The first section of the code sets the input parameters corresponding to each \mathbb{W} and creates the necessary numerical domains. We select the aspect ratio of each domain to be equivalent to the experimental visualisation window of 1210 mm by 410 mm. This corresponds to the number of horizontal n_x and vertical n_z cells being 242 and 82 respectively with spacing $\Delta x = 5.02$ mm and $\Delta z = 5.06$ mm. The result of changing the grid resolution is explored in § 5.6. The number of numerical domains n_d created is equal to the number of waves in the system. Initially, we set $n_d = 3$, corresponding to one set of parameters for each wave \mathbb{W}_0 , \mathbb{W}_1 and \mathbb{W}_2 . For the remainder of this thesis, the case where $n_d = 3$, will be referred to as a ‘simple triad’. As we will see in § 6.3, it is useful to extend this idea to admit multiple different pairs of secondary waves. For this case \mathbb{W}_1 and \mathbb{W}_2 are broken down into subsets

$$\begin{aligned}\mathbb{W}_1^{(g)} &= \{\mathbb{W}_{1_1}, \mathbb{W}_{1_2}, \dots, \mathbb{W}_{1_j}, \dots, \mathbb{W}_{1_g}\}, \\ \mathbb{W}_2^{(g)} &= \{\mathbb{W}_{2_1}, \mathbb{W}_{2_2}, \dots, \mathbb{W}_{2_j}, \dots, \mathbb{W}_{2_g}\},\end{aligned}\tag{5.27}$$

for $j = 1, \dots, g$, where g defines the number of pairs of secondary waves input into the system. The linear superposition of the corresponding elements in both the $\mathbb{W}_1^{(g)}$ and $\mathbb{W}_2^{(g)}$ subsets then forms the overall value for \mathbb{W}_1 and \mathbb{W}_2 respectively. As in (4.3), each wave in the subgroup has its own set of parameters and fields

$$\begin{aligned}\mathbb{W}_{1_j} &= \{\omega_{1_j}, \mathbf{k}_{1_j}, \rho_{1_j}, \tilde{\Psi}_{1_j}, \dots\}, \\ \mathbb{W}_{2_j} &= \{\omega_{2_j}, \mathbf{k}_{2_j}, \rho_{2_j}, \tilde{\Psi}_{2_j}, \dots\}.\end{aligned}\tag{5.28}$$

Each pair of secondary waves (e.g. $\{\mathbb{W}_{1_j}, \mathbb{W}_{2_j}\}$) is then defined as \mathbb{S}_j . As each \mathbb{S}_j always shares the same \mathbb{W}_0 , the triad is defined as

$$\mathbb{T}_j = \{\mathbb{W}_0, \mathbb{W}_{1_j}, \mathbb{W}_{2_j}\} = \{\mathbb{W}_0, \mathbb{S}_j\},\tag{5.29}$$

where \mathbb{W}_0 and the corresponding pair of secondary waves in \mathbb{S}_j , satisfy the triadic resonant condition (2.4). For these cases where multiple secondary waves are considered, $n_d = 1 + 2g$, noting that a single \mathbb{W}_0 is shared between all \mathbb{S}_j . Details of the initial and boundary conditions for each domain are given in § 5.4.6.

5.3.3 Overall code structure

At each time-step in the code, every domain undergoes a linear advection step followed by non-linear interactions. The first step is the linear advection, although the order of these operations is arbitrary. The numerical form of the time derivative in the linear advection equation (5.18) can be written as

$$\frac{\partial \tilde{\Psi}_d}{\partial t} = \frac{\tilde{\Psi}_d^\dagger - \tilde{\Psi}_d^n}{\Delta t} = -\nabla_x \cdot \mathbf{c}_{gd} \tilde{\Psi}_d^n - \frac{\nu}{2} \kappa_d^2 \tilde{\Psi}_d^2, \quad (5.30)$$

where the subscript $d = 1, \dots, n_d$ represents the wave domain (as the number of domains is not necessarily 3, it is not possible to use the subscript p), the values $\tilde{\Psi}_d^n$ and $\tilde{\Psi}_d^\dagger$ respectively represent the initial value of $\tilde{\Psi}_d$ and the new value of $\tilde{\Psi}_d$ after advection and Δt is the time-step over which the code is advanced. Details of how Δt is selected are given in § 5.4.3. Re-arranging (5.30) for $\tilde{\Psi}_d^\dagger$ gives

$$\tilde{\Psi}_d^\dagger = \tilde{\Psi}_d^n + \Delta t \left(-\nabla \cdot \mathbf{c}_{gd} \tilde{\Psi}_d^n - \frac{\nu}{2} \kappa_d^2 \tilde{\Psi}_d^2 \right). \quad (5.31)$$

Details of the numerical advection on the RHS of (5.31) are outlined in § 5.4.

If only linear advection is considered then $\tilde{\Psi}_d^{n+1} = \tilde{\Psi}_d^\dagger$. If, however, the weakly non-linear interactions are considered then

$$\tilde{\Psi}_d^{n+1} = \tilde{\Psi}_d^\dagger + \Delta t \left[\frac{\partial \tilde{\Psi}_d^\dagger}{\partial t} \right], \quad (5.32)$$

where the partial time derivative inside the square brackets is given by the non-linear interactions in (5.19). In § 6.4, we also consider the time derivative to be comprised of a combination of both (5.19) and the complex phase shift given by (5.26). For this case (5.32) becomes

$$\tilde{\Psi}_d^{n+1} = \tilde{\Psi}_d^\dagger + \Delta t \left[\frac{\partial \tilde{\Psi}_d^\dagger}{\partial t} + \frac{\partial \tilde{\Psi}_d^\dagger}{\partial t} \right]. \quad (5.33)$$

The streamfunction domains and the time-step are then updated by

$$\begin{aligned} \tilde{\Psi}_d^n &= \tilde{\Psi}_d^{n+1}, \\ t &= t + \Delta t. \end{aligned} \quad (5.34)$$

where n is the time iteration and process then starts again from (5.31) for the subsequent time-step. If the non-linear interactions are not taken into consideration then code omits the step (5.32) (and also (5.33)) and moves from (5.31) to (5.34).

5.4 Linear advection scheme

5.4.1 Overview

This section provides details on the numerical implementation of the advection scheme given in (5.31). The same advection scheme is used for all the domains.

5.4.2 Computational cell

A finite volume method with a staggered grid was chosen for the development of the advection scheme. The computational cell used is shown in Figure 5.1. As opposed to the finite difference method, which takes the values at the grid nodes, the finite volume method discretizes the governing equations into an arbitrary control volume around each node (Mazumder, 2015), shown by the red dashed box. The advective flux is then evaluated across the faces of the control volume. The estimate of the face-average flux was initially obtained using a first-order scheme, outlined in § 5.4.4, then later using a second-order scheme, discussed in § 5.4.5. In Figure 5.1 we use a compass index co-ordinate system to represent the neighbouring cells in order to avoid confusion with cell numbering in Matlab.

5.4.3 Time-step

An explicit time-stepping scheme was used. Explicit schemes require a sufficiently small time-step to keep the scheme computationally stable. This was achieved via the Courant-Friedrichs-Lewy condition (CFL). The CFL condition requires that the time-step Δt is equal to or smaller than the time required to transport information via the flux across the faces of the volume. The non-dimensional Courant number must, therefore, not exceed unity. The Courant number is given as

$$C = \Delta t \left(\frac{|c_{g_x}|}{\Delta x} + \frac{|c_{g_z}|}{\Delta z} \right) \leq 1, \quad (5.35)$$

where Δx and Δz represent the grid spacing, shown in Figure 5.1 and c_{g_x} and c_{g_z} are respectively the horizontal and vertical components of the group velocity \mathbf{c}_g . Rather than choose the maximum time-step that would just satisfy (5.35), C was arbitrary set at 0.5 and (5.35) was re-arranged to determine the time-step Δt . The value for Δt was then checked to ensure

$$\Delta t < \frac{\nu}{|\Delta x|^2}, \quad (5.36)$$

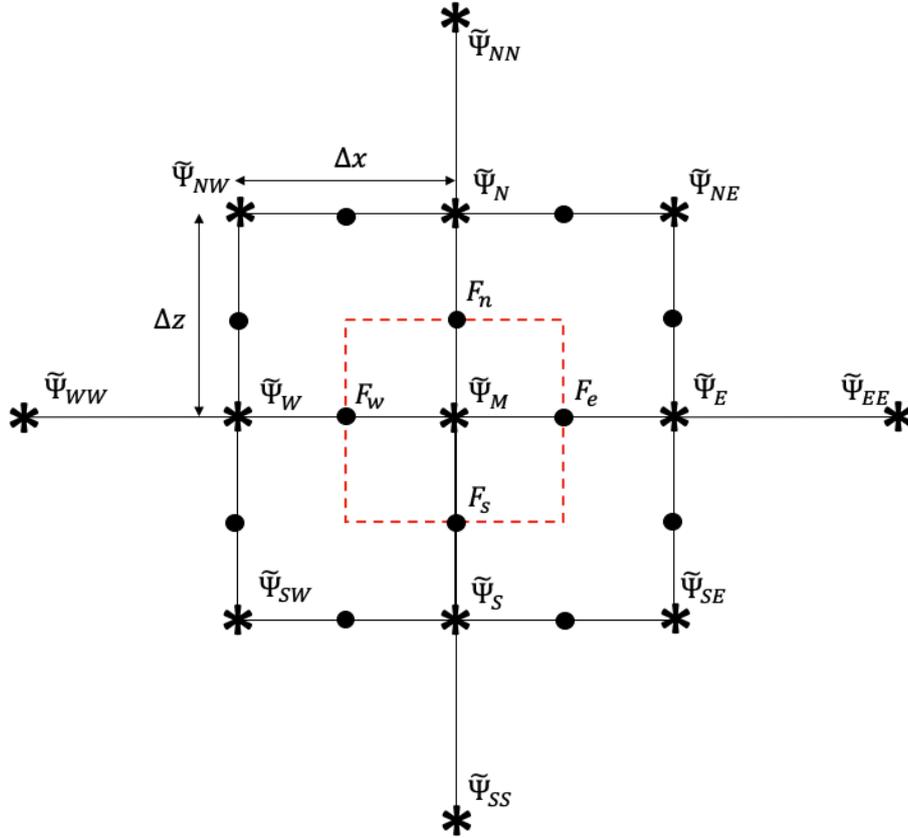


Figure 5.1: Nomenclature for the computational cell where $\tilde{\Psi}$ is the stream function and F is the advective flux of the stream function across the face of the control volume. The grid dimensions of the cell are given by Δx and Δz . Compass indexing has been used to avoid confusion with numbering systems.

although this was rarely a constraint. If the time-step did not satisfy this limit, then (5.36) was used to determine Δt .

The group velocity c_g in (5.35) corresponded to whichever \mathbb{W} in the system had the largest group velocity, to ensure that the code was computationally stable for all waves. Since the group velocities of the waves are not a function of time, the same Δt was used for the entire run.

5.4.4 First-order upwind scheme

The initial implementation of the numerical scheme was a first-order upwind scheme. An upwind scheme uses differencing methods that are biased towards the direction from which the flux is coming. For internal waves this is determined by their relative

group velocity. The fluxes across the control volumes were calculated by

$$\begin{aligned}
F_w = c_{g_x} \tilde{\Psi}_w, \quad \text{where} \quad \tilde{\Psi}_w &= \begin{cases} \tilde{\Psi}_W & \text{for } c_{g_x} > 0, \rightarrow \\ \tilde{\Psi}_M & \text{for } c_{g_x} < 0, \leftarrow \end{cases} \\
F_e = c_{g_x} \tilde{\Psi}_e, \quad \text{where} \quad \tilde{\Psi}_e &= \begin{cases} \tilde{\Psi}_M & \text{for } c_{g_x} > 0, \rightarrow \\ \tilde{\Psi}_E & \text{for } c_{g_x} < 0, \leftarrow \end{cases} \\
F_n = c_{g_z} \tilde{\Psi}_n, \quad \text{where} \quad \tilde{\Psi}_n &= \begin{cases} \tilde{\Psi}_M & \text{for } c_{g_z} > 0, \uparrow \\ \tilde{\Psi}_N & \text{for } c_{g_z} < 0, \downarrow \end{cases} \\
F_s = c_{g_z} \tilde{\Psi}_s, \quad \text{where} \quad \tilde{\Psi}_s &= \begin{cases} \tilde{\Psi}_S & \text{for } c_{g_z} > 0, \uparrow \\ \tilde{\Psi}_M & \text{for } c_{g_z} < 0, \downarrow \end{cases}
\end{aligned} \tag{5.37}$$

The value of the stream function at the central node $\tilde{\Psi}_M$ was then updated by the differences in the fluxes over the control volume using

$$\tilde{\Psi}_M^\dagger = \tilde{\Psi}_M^n + \Delta t (-\nabla \cdot \mathbf{c}_g \tilde{\Psi}_M^n) = \tilde{\Psi}_M^n + \Delta t \left(\frac{F_w^n - F_e^n}{\Delta x} + \frac{F_s^n - F_n^n}{\Delta z} \right). \tag{5.38}$$

The first-order scheme results in a truncation error of order Δx , due to a first-order approximation of the Taylor Series expansion. This error amounts to large numerical diffusion. To minimise this numerical diffusivity the code was upgraded to a second-order scheme.

5.4.5 Second-order upwind scheme

A second-order upwind scheme was employed to reduce the numerical diffusivity inherent in the first-order scheme. For the second-order scheme, the calculation of the

advective fluxes were given by

$$\begin{aligned}
F_w &= c_{g_x} \tilde{\Psi}_w, \\
\text{where } \tilde{\Psi}_w &= \begin{cases} \tilde{\Psi}_W + \frac{\lambda_R}{2} (\tilde{\Psi}_W - \tilde{\Psi}_{WW}) & \text{for } c_{g_x} > 0, \rightarrow \\ \tilde{\Psi}_M + \frac{\lambda_L}{2} (\tilde{\Psi}_M - \tilde{\Psi}_E) & \text{for } c_{g_x} < 0, \leftarrow \end{cases} \\
F_e &= c_{g_x} \tilde{\Psi}_e, \\
\text{where } \tilde{\Psi}_e &= \begin{cases} \tilde{\Psi}_M + \frac{\lambda_R}{2} (\tilde{\Psi}_M - \tilde{\Psi}_W) & \text{for } c_{g_x} > 0, \rightarrow \\ \tilde{\Psi}_E + \frac{\lambda_L}{2} (\tilde{\Psi}_E - \tilde{\Psi}_{EE}) & \text{for } c_{g_x} < 0, \leftarrow \end{cases} \\
F_n &= c_{g_z} \tilde{\Psi}_n, \\
\text{where } \tilde{\Psi}_n &= \begin{cases} \tilde{\Psi}_M + \frac{\lambda_U}{2} (\tilde{\Psi}_M - \tilde{\Psi}_S) & \text{for } c_{g_z} > 0, \uparrow \\ \tilde{\Psi}_N + \frac{\lambda_D}{2} (\tilde{\Psi}_N - \tilde{\Psi}_{NN}) & \text{for } c_{g_z} < 0, \downarrow \end{cases} \\
F_s &= c_{g_z} \tilde{\Psi}_s, \\
\text{where } \tilde{\Psi}_s &= \begin{cases} \tilde{\Psi}_S + \frac{\lambda_U}{2} (\tilde{\Psi}_S - \tilde{\Psi}_{SS}) & \text{for } c_{g_z} > 0, \uparrow \\ \tilde{\Psi}_M + \frac{\lambda_D}{2} (\tilde{\Psi}_M - \tilde{\Psi}_N) & \text{for } c_{g_z} < 0, \downarrow \end{cases}
\end{aligned} \tag{5.39}$$

Again, the value of the stream function $\tilde{\Psi}_M^\ddagger$ was calculated using (5.38). The value of λ , known as the flux limiter, was introduced to ensure monotonicity in the advection scheme. The subscript on the flux limiter R, L, U, D stands for the direction of the flux, as the flux limiter will change depending on the direction of flow. At $\lambda = 1$ the scheme was second-order. However, second-order advection schemes can produce un-physical oscillations. Goudnov's theorem states that numerical schemes for solving partial differential equations (PDEs), can be at most first-order accurate (Godunov, 1954). By including the second-order flux limiter we ensure that the maxima must be non-increasing and the minima non-decreasing, hence preventing the creation of spurious oscillations (Van Leer, 1974). The 'min-mod' function (Versteeg and Malalasekera, 2007) was chosen for the flux limiter

$$\lambda(r) = \max[0, \min(1, r)], \tag{5.40}$$

where r is given as the ratio of the downstream to upstream gradient and is based on the direction of the flow

$$\begin{aligned}
r_R &= \frac{\tilde{\Psi}_M - \tilde{\Psi}_W}{\tilde{\Psi}_W - \tilde{\Psi}_{WW}} && \text{for } c_{g_x} > 0, \rightarrow \\
r_L &= \frac{\tilde{\Psi}_M - \tilde{\Psi}_E}{\tilde{\Psi}_E - \tilde{\Psi}_{EE}} && \text{for } c_{g_x} < 0, \leftarrow \\
r_U &= \frac{\tilde{\Psi}_M - \tilde{\Psi}_N}{\tilde{\Psi}_N - \tilde{\Psi}_{NN}} && \text{for } c_{g_z} > 0, \uparrow \\
r_D &= \frac{\tilde{\Psi}_M - \tilde{\Psi}_S}{\tilde{\Psi}_S - \tilde{\Psi}_{SS}} && \text{for } c_{g_z} < 0, \downarrow
\end{aligned} \tag{5.41}$$

Depending of the gradient ratio in r , $0 \leq \lambda \leq 1$. A monotonic scheme is then achieved by reducing ($0 < \lambda < 1$) or eliminating ($\lambda = 0$) the second-order contributions where necessary.

It should be noted that other second-order flux limiter functions (SUPERBEE, van Leer and van Albada (Versteeg and Malalasekera, 2007)) were tested, however the mid-mod function qualitatively appeared to be the most suitable for the advection scheme.

5.4.6 Initial and boundary conditions

Each domain, \mathbb{W}_d , was initialised with a small-magnitude complex background noise, as without this the instability could not be triggered. At $t = 0$, each cell in the domain contained random complex noise of magnitude $10^{-7} \delta e^{i(2\pi\vartheta)}$, where both δ and ϑ were uniform random numbers spanning the range $[0 \ 1]$. The mean magnitude of this random noise was therefore of the order 5×10^{-8} , while the phase angle was randomly distributed over the range $[0 \ 2\pi]$.

Being a second-order upwind scheme it was necessary to have two outer grid layers at each boundary. The values on these boundaries for each domain were dependant on the direction of the group velocity for that wave domain \mathbb{W} . If $c_{g_x} < 0$ and $c_{g_z} > 0$, as it was for \mathbb{W}_0 and \mathbb{W}_2 , then the outflow boundary conditions were the north and west boundary. For these outflow boundaries the values on the two outer grid layers were set equal to the values at the edge of the inner grid (Versteeg and Malalasekera, 2007) and the waves were simply advected out of the domain. The two inflow boundaries, specified the flux in and, for \mathbb{W}_1 and \mathbb{W}_2 , were set to a small complex value of the form $10^{-7} e^{i(2\pi\varepsilon)}$, where ε represents a constant value over the range $[0 \ 1]$. Simply setting the inflow boundary equal to 0 would mean that the instability would be unable

to develop as the secondary wave domains would be zero for all time. As ε was held constant, the inflow boundary then propagated a constant complex value into the domain with a magnitude comparable to the background noise but with constant phase angle. If the direction of the group velocity was in the other direction (eg. $c_{g_x} > 0$ and $c_{g_z} < 0$), as it was for \mathbb{W}_1 , then the south and east boundary were the outflow boundaries and the north and west boundaries were the inflow boundaries.

The only domain containing a different inflow boundary was the primary wave beam, where it was necessary to impose bottom boundary forcing to simulate the wave maker. As the model only considers the advection of the complex valued reduced streamfunction $\tilde{\Psi}$, the boundary condition is illustrated by the red envelope shown in Figure 4.1, which removes the fast time forcing given by $e^{i(l_0x - \omega_0t)}$. The bottom boundary forcing for the primary wave domain \mathbb{W}_0 was

$$z = h(x, t) = \begin{cases} f(t) \cos^2\left(\frac{x-B}{8\pi^2}\right), & A < x < B, \\ f(t), & B < x < C, \\ f(t) \cos^2\left(\frac{x-C}{8\pi^2}\right), & C < x < D, \\ 0, & \text{elsewhere,} \end{cases} \quad (5.42)$$

where, in the same way as the experimental forcing in (4.1), A, B, C, D are respectively $7\pi/|l_0|$, $9\pi/|l_0|$, $13\pi/|l_0|$, $15\pi/|l_0|$, and l_0 is the horizontal component of the primary wave vector given as -0.05 mm^{-1} . The temporal forcing $f(t)$ was given as

$$f(t) = \begin{cases} 0, & t \leq 0 \text{ s,} \\ \tilde{\Psi}_{\text{in}}\left(\frac{t}{30}\right), & 0 \leq t \leq 30 \text{ s,} \\ \tilde{\Psi}_{\text{in}}, & 30 \text{ s} \leq t \leq t_{\text{end}} - 30 \text{ s,} \\ \tilde{\Psi}_{\text{in}}\left(\frac{t_{\text{end}}-t}{30}\right), & t_{\text{end}} - 30 \text{ s} \leq t \leq t_{\text{end}}. \end{cases} \quad (5.43)$$

Unlike the temporal forcing for the experiments given in (4.2), where the amplitude was prescribed in (mm) by η_0 , here the complex valued $\tilde{\Psi}_{\text{in}}$ was used for the beam amplitude. The structure of $\tilde{\Psi}_{\text{in}}$ was given as $|\tilde{\Psi}_{\text{in}}|e^{i(2\pi\varsigma)}$, where ς was another constant value over the range $[0 \ 1]$. This ensured that the phase angle of the slowly evolving complex amplitude $\tilde{\Psi}_0$ was fixed at the boundary. For comparison with the experimental results, $|\tilde{\Psi}_{\text{in}}|$ was kept similar to $\langle |\tilde{\Psi}_0| \rangle_r$ defined in (3.30). For consistency with the experimental results, t_{end} was normally kept at 5460 s (90 minutes).

5.5 Weakly non-linear interactions

Once the linear advection scheme had been developed it was possible to include in the weakly non-linear interactions into the model. The set of non-linear interaction terms given in (5.19) were converted into the numerical format

$$\tilde{\Psi}_0^{n+1} = \tilde{\Psi}_0^\dagger + \Delta t(I_0 \tilde{\Psi}_1^\dagger \tilde{\Psi}_2^\dagger), \quad (5.44a)$$

$$\tilde{\Psi}_1^{n+1} = \tilde{\Psi}_1^\dagger + \Delta t(I_1 \tilde{\Psi}_0^\dagger \tilde{\Psi}_2^{*\dagger}), \quad (5.44b)$$

$$\tilde{\Psi}_2^{n+1} = \tilde{\Psi}_2^\dagger + \Delta t(I_2 \tilde{\Psi}_0^\dagger \tilde{\Psi}_1^{*\dagger}), \quad (5.44c)$$

where I is the interaction term defined by (5.20), the superscript $\tilde{\Psi}^\dagger$ is the value of $\tilde{\Psi}$ after the advection step and the superscript $\tilde{\Psi}^*$ denotes the complex conjugate. If a simple triad was being considered and $n_d = 3$, then this set of equations would be considered only once. If, however, $n_d > 3$, then the set of equations would be repeated g times (for details on this notation see § 5.3.2). In this case, the input values for I changed on each iteration, stepping through all the parameters corresponding to $\mathbb{W}_1^{(g)} = \{\mathbb{W}_{1_1}, \dots, \mathbb{W}_{1_g}\}$ and $\mathbb{W}_2^{(g)} = \{\mathbb{W}_{2_1}, \dots, \mathbb{W}_{2_g}\}$. At the end of this loop

$$\tilde{\Psi}_0^{n+1} = \tilde{\Psi}_0^\dagger + \Delta t \sum_{j=1}^g (I_{0_j} \tilde{\Psi}_{1_j}^\dagger \tilde{\Psi}_{2_j}^\dagger), \quad (5.45a)$$

$$\tilde{\Psi}_{1_j}^{n+1} = \tilde{\Psi}_{1_j}^\dagger + \Delta t(I_{1_j} \tilde{\Psi}_0^\dagger \tilde{\Psi}_{2_j}^{*\dagger}), \quad (5.45b)$$

$$\tilde{\Psi}_{2_j}^{n+1} = \tilde{\Psi}_{2_j}^\dagger + \Delta t(I_{2_j} \tilde{\Psi}_0^\dagger \tilde{\Psi}_{1_j}^{*\dagger}), \quad (5.45c)$$

where $j = 1, \dots, g$. The same procedure was also followed in the case where the complex non-linear interactions (5.26) were considered.

5.6 Verification

5.6.1 Overview

After the \mathcal{M}_{2D} model was implemented it was necessary to verify that the numerical code was providing the expected result. This verification was done in two ways. Firstly, looking at the linear advection framework, we compared the inviscid solution ($\nu = 0$) to an exact analytical solution derived from the boundary condition. The linear solution was obtained by turning off the non-linear interactions at each time-step, which amounts to removing the step show in (5.32). The exact solution can be compared to \mathbb{W}_0 , as

this is the domain with the boundary forcing, however the results would also valid for domains \mathbb{W}_1 and \mathbb{W}_2 if a source of waves were included into these domains. Results for this comparison are presented in § 5.6.2. Secondly, the conservation of energy was monitored to quantify the numerical dissipation inherent in the scheme. Again, this was done only for the inviscid solution, as including viscosity would also contribute to the dissipation in the system. These results are presented in § 5.6.3.

5.6.2 Comparison with analytical solution

An exact analytical solution of the wave at steady state was generated using

$$\tilde{\Psi}_{\text{exact}} = \begin{cases} \tilde{\Psi}_{\text{in}} \cos^2 \left(\frac{x+z-B}{8\pi^2} \right), & A < x_1 < B, \\ \tilde{\Psi}_{\text{in}}, & B < x_1 < C, \\ \tilde{\Psi}_{\text{in}} \cos^2 \left(\frac{x+z-C}{8\pi^2} \right), & C < x_1 < D, \\ 0, & \text{elsewhere,} \end{cases} \quad (5.46)$$

where, in the same way as the experimental forcing in (4.1), A, B, C, D were taken as $7\pi/|l_0|$, $9\pi/|l_0|$, $13\pi/|l_0|$, $15\pi/|l_0|$ respectively and l_0 was -0.05 mm^{-1} . For this section, $\tilde{\Psi}_{\text{in}}$ was real valued and equal to 1. The parameters x and z correspond to domains giving the horizontal and vertical locations in mm of each grid cell respectively. The value x_1 therefore indicates the bottom row of the domain, where the subscript 1 corresponds to $z = 1$.

To compare the analytical solution with the inviscid solution from the model, the root mean square (RMS) error – also known as the discretisation error – was calculated by

$$\text{error} = \sqrt{\frac{\sum (|\tilde{\Psi}_{\text{exact}}| - |\tilde{\Psi}_0|)^2}{n_x n_z}}, \quad (5.47)$$

where the difference between each cell was taken and squared before summing across all the rows and columns of the domain. According to Versteeg and Malalasekera (2007) any computational fluid dynamics modelling should aim to show a monotonic reduction in the discretisation error over two or three successive levels of mesh refinement. As the source of error in the model comes from the truncation error due to the Taylor Series approximation to the solution, it can be expressed by

$$\text{error}(h) \approx Ch^p, \quad (5.48)$$

where h is the reference size of the mesh, C is some constant and p is the order of the numerical scheme. As the scheme used here is a monotonic second-order scheme, we therefore expect the error to reduce as $(1/n_x)^p$, where $1 < p \leq 2$. The error was computed using (5.47) for six different grid resolutions shown in Table 5.1. The results from Table 5.1 are then shown graphically in Figure 5.2.

n_x	n_z	Δx	Δz	RMS error
61	21	20.17	20.50	0.064
86	29	14.23	14.64	0.045
121	41	10.08	10.25	0.030
173	59	7.03	7.06	0.019
242	82	5.02	5.06	0.013
346	117	2.51	3.53	0.009

Table 5.1: Table showing the RMS error between the model and exact solution calculated from (5.47) as a function of grid resolution.

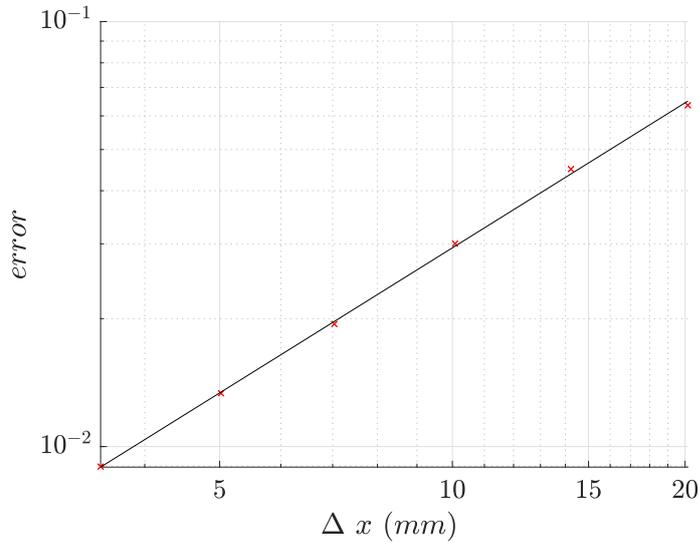


Figure 5.2: Log-log plot showing the RMS error as a function of grid resolution. The red crosses show the calculated values for six different grid resolutions, while the black line shows a linear best fit.

Figure 5.2 shows that the accuracy of the solution increases with smaller grid spacing. The reduction in error follows the form (5.48) where the exponent p is equal to 1.15. This value of error reduction rate is due to the introduction of the flux limiters, which revert the scheme back to first-order in certain locations. This is shown in

Figure 5.3(a), which graphically depicts the locations where the flux limiters are acting for a grid resolution of $n_x, n_z = 82, 242$. An average of the horizontal and vertical flux limiter (in this case $(\lambda_L + \lambda_U)/2$ as \mathbb{W}_0 propagates up and to the left) is shown, multiplied by the primary wave field $|\tilde{\Psi}_0|$. This shows that the estimation of the face-averaged flux lies between first and second-order over a considerable portion of the flow.

Figure 5.3(b) then shows the difference between the model and analytical solution for the same grid resolution of $n_x, n_z = 82, 242$. This Figure demonstrates that the spatial distribution of the dissipation is coincident with the flanks of the wave field where the flux limiters play a role. This is, therefore, the cause of the slightly higher than expected reduction rate of the RMS error over successive grid resolutions. Indeed, Roache (1997) says that it should not be taken for granted that the actual truncation error in a numerical solution will decay exactly in accordance with the formal order p given in (5.48). Despite this, the maximum error per cell is of magnitude 4×10^{-4} and that is relative to $|\tilde{\Psi}_0|$.

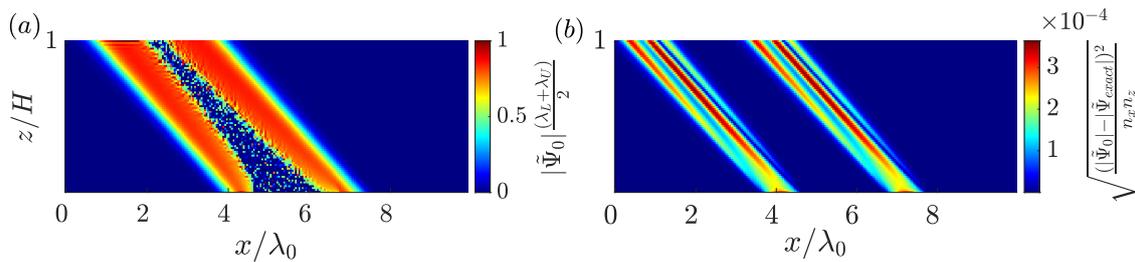


Figure 5.3: (a) Locations where the flux limiters are acting in the flow to keep the scheme monotonic. The average of the horizontal and vertical λ is given, multiplied by amplitude of the wave beam. (b) The RMS error between the exact solution and the solution produced by the model calculated by (5.47) over the domain with the same resolution. The error occurs across the flanks of the wave where the scheme will not necessarily be second-order due to the flux limiters. Both images are for a grid resolution of $n_x, n_z = 82, 242$.

5.6.3 Energy conservation

The other way that the numerical code was verified was through energy conservation. As with the comparison to the exact solution presented in § 5.6.2, we consider energy conservation for the linear advection scheme. Assuming an inviscid flow, through the use of the divergence theorem we can state that

$$\frac{\partial \varphi}{\partial t} + \nabla \cdot (\mathbf{u}\varphi) = 0, \quad (5.49)$$

where φ is some scalar quantity and \mathbf{u} is the velocity vector, which in this case is \mathbf{c}_g . The accuracy of the advection scheme can be tested by seeing if the scalar quantity φ is conserved through (5.49) over a given control volume. As the model advects the streamfunction $\tilde{\Psi}$, for the advection scheme, this quantity was conserved to machine accuracy by construction. We therefore tested how well the energy was conserved across the whole numerical domain. The energy dissipation error ε was obtained by computing

$$\varepsilon = \frac{\partial E}{\partial t} + \nabla \cdot (\mathbf{c}_g E), \quad (5.50)$$

where energy was calculated by $E = \kappa^2 \tilde{\Psi}^2$. This calculation was performed at every time-step. Figure 5.4 shows how the energy changes across the whole domain over time. Specifically, the rate of change of energy $\partial E / \partial t$, is given in blue, while the red line shows the divergence of the energy flux $\nabla \cdot (\mathbf{c}_g E)$, integrated over the whole domain. The beginning of the graph between $0 < t/T_0 < 5$ corresponds to the initial startup of the numerical wavemaker. Here the amplitude of the wave increases from rest over 30 s ($t/T_0 \approx 5$), as given in the boundary forcing (5.42). Over this time, the rate of change of energy is positive, while the net energy flux out of the domain is negative. When the beam hits the top boundary and begins to leave the domain, the rate of energy change begins to decrease as the net flux out of the domain begins to increase towards zero. A steady state is reached after approximately $t = 80$ s ($t/T_0 \approx 12$).

Figure 5.4 shows that at steady state there is a small numerical dissipation in the scheme given by $\varepsilon = \nabla \cdot (\mathbf{c}_g E)$. Figure 5.5 shows the spatial distribution of this numerical dissipation at steady state. Here (5.50) has been calculated on a cell by cell basis. As with Figure 5.3(b), the numerical dissipation is occurring across the flanks of the wave beam, at the locations where the flux limiters act to reduce or remove the second-order contributions. Interestingly, Figure 5.5 also shows larger numerical dissipation at the bottom of the domain over the flanks of the wavemaker. As the cosine squared edges of the envelope smooth the wavebeam in this region, there is a loss in energy directly above the wave maker boundary due to this smoothing.

We then compare how the numerical dissipation changes as function of grid resolution. Again, looking at steady state, the arithmetic mean of ε shown in Figure 5.5 was taken over the whole domain for same six grid resolutions shown in § 5.6.2. The results are shown in Table 5.2 and graphically in Figure 5.6.

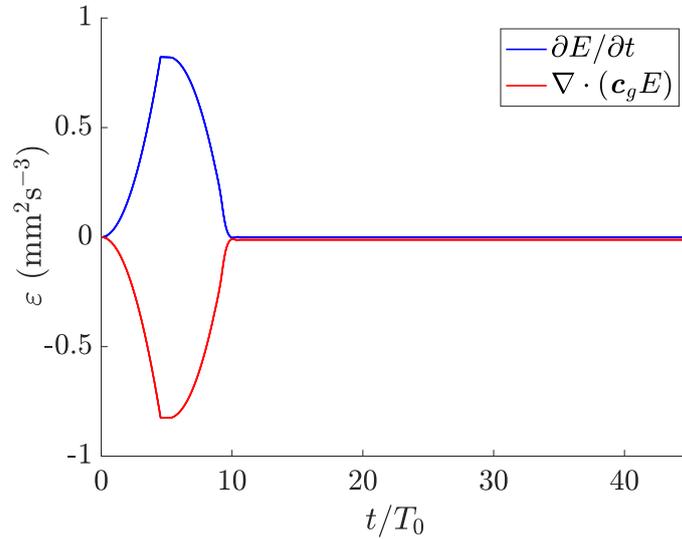


Figure 5.4: Calculation of ε over the whole domain of size $n_x, n_z = 82, 242$ as a function of time. The blue line gives the change in energy over the domain, while the red line shows the divergence of energy. At steady state, around $t/T_0 \approx 12$, the change of energy is equal to zero, however the divergence is non-zero, indicating a small numerical dissipation.

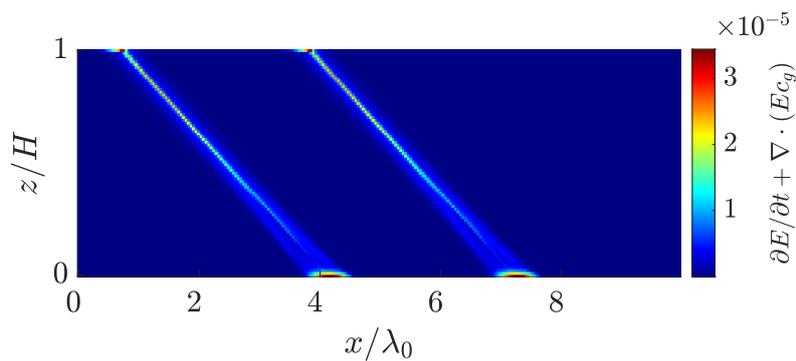


Figure 5.5: The energy dissipation error ε over each cell in a 82×242 domain after the wave has reached steady state. The error across each cell is calculated by (5.50).

n_x	n_z	Δx	Δz	ε
61	21	20.17	20.50	6.06e-06
86	29	14.23	14.64	3.82e-06
121	41	10.08	10.25	2.22e-06
173	59	7.03	7.06	1.19e-06
242	82	5.02	5.06	6.28e-07
346	117	2.51	3.53	2.96e-07

Table 5.2: Table showing the energy dissipation error calculated using (5.50) as a function of grid resolution, where now ε is calculated across the whole domain.

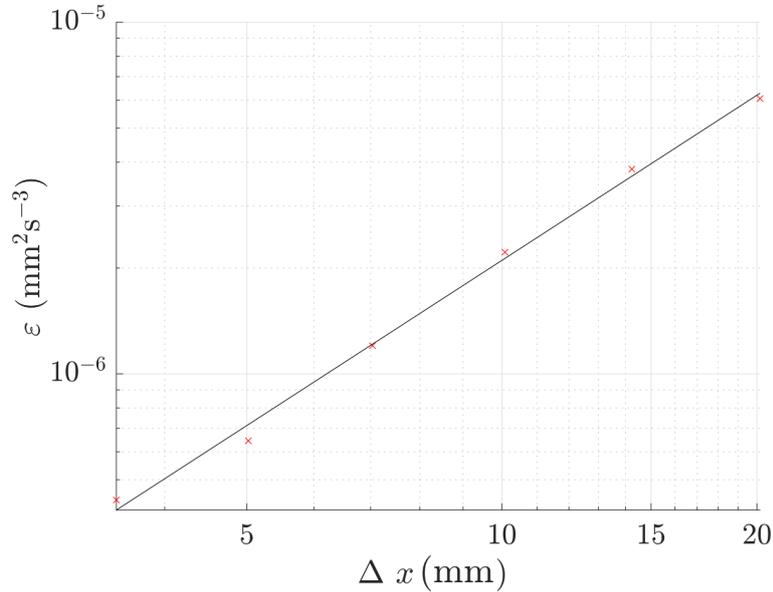


Figure 5.6: Log-log plot showing the energy dissipation error ε across the whole domain as a function of grid resolution. The red crosses show the calculated values for six different grid resolutions, while the black line shows a linear best fit.

Figure 5.6 shows that ε decreases as the resolution increases. Here, the exponent p from (5.48) was found to be 1.52. This is a reasonable decrease in error with resolution, given that the numerical dissipation was coincident with the flux limiters. Overall, the error associated with the numerical dissipation was small and we were confident that the formulation of the scheme – second-order with flux limiters – ensured that any change in energy would cause a decay rather than growth, and that this decay would decrease as the resolution was increased.

5.7 Summary

In this Chapter we have outlined the development of the weakly non-linear \mathcal{M}_{2D} model. Specifically, we have introduced the perturbation expansion used to obtain the two-dimensional theoretical equations and the subsequent numerical code used to solve them. Through the use of the spatial non-dimensional parameter γ we have shown how an advection equation with constant advection velocity \mathbf{c}_g , can be obtained for internal waves at leading order. We then examined the non-linear terms obtained for higher orders of ϵ . We saw that at $\mathcal{O}(\gamma^0\epsilon^2)$, the same non-linear interaction terms between the triad, given in Bourget et al. (2013), were recovered. In § 5.4 we explained how the advection equation was numerically implemented in the code through the use of a monotonic second-order finite volume scheme. The numerical solution of the weakly non-linear equations were then introduced in § 5.5. Finally the numerical code was verified in § 5.6 by comparing the results obtained from the advection scheme scheme to an exact analytical solution and checking that energy was conserved in the system.

Chapter 6

Weakly non-linear model results

6.1 Introduction

This Chapter presents results from the weakly non-linear \mathcal{M}_{2D} model – outlined in Chapter 5 – exploring how well it captures the development of TRI in a finite-width wavebeam. Other than in § 6.2.3, the boundary forcing for \mathbb{W}_0 given in (5.42) is kept the same, to allow for comparisons with the experimental results. The main changes to the model inputs are therefore related to the parameters for \mathbb{W}_1 and \mathbb{W}_2 .

Initially when exploring the model, a simple triad ($n_d = 3$) was examined, corresponding to only one set of input parameters for the \mathbb{W}_1 and \mathbb{W}_2 domains. Results obtained from this set-up showed similar behaviour to the \mathcal{M}_{0D} model presented in Figure 4.12, with decay of the primary beam and all waves reaching a steady equilibrium. As discussed in Chapter 4, however, the frequencies and wavenumbers of \mathbb{W}_1 and \mathbb{W}_2 do not remain fixed throughout the course of an experiment. An investigation was consequently undertaken to see how different input parameters for the secondary waves in the simple triad affected the results. Through this, it was revealed how influential the frequencies and wavenumbers of the secondary waves are in the development of the instability.

In addition to examining a range of simple triads, the model was also modified to include g pairs of secondary waves – as described in (5.27) – allowing for multiple-triads of a similar configuration to be present in the system. For this configuration, $n_d = 2g + 1$, meaning the different pairs of secondary waves are only able to interact with each other via \mathbb{W}_0 . By including multiple possible secondary waves in Fourier space the model makes a significant step in capturing both the spatial and temporal oscillations witnessed experimentally.

This chapter is broken down as follows. In § 6.2 we start by examining results from the model based on simple triads. In § 6.2.2 the triad configuration is located on the blue outer branch of the loci curve shown in Figure 4.13 as, experimentally, the triad always corresponds to this branch. Here we provide a detailed investigation stepping through a range of possible simple triads located along this part of the loci curve. Next, in § 6.2.3, we conduct the same analysis for the simple triad configurations located on the inner pink branch of the loci curve. For the \mathcal{M}_{0D} model, this branch has a higher linear theoretical growth rate. We therefore explore why this branch is not selected experimentally using the \mathcal{M}_{2D} model. The results in § 6.3 then explore the effect of allowing multiple pairs of secondary waves to interact with the primary. The subsequent modulation in both physical and Fourier space is outlined in § 6.3.2 and § 6.3.3 respectively. Finally, the results of including the higher order perturbation expansion in (5.23) are explored in § 6.4 and a summary is presented in § 6.5.

6.2 Simple triad

6.2.1 Overview

We start by looking at results from the model using only a simple triad. As discussed in § 5.3.2, the term ‘simple triad’ refers to the case where only three domains are used in the model ($n_d = 3$), corresponding to one set of input parameters (frequency, wave number etc.) for each wave in the triad. In § 6.2.2 we examine simple triads whose \mathbf{k}_1 vector lies on the outer blue branch of the loci curve shown again in Figure 6.1. Then, in § 6.2.3, we change the triad configuration to lie on the inner pink branch of the loci curve.

6.2.2 Outer loci branch

For the experimental configuration presented in Chapter 4, if the primary beam becomes unstable to TRI, the tip of the \mathbf{k}_1 vector is always located on the outer branch of the loci curve. The approximate location range of these experimental vectors is shown by the grey patch on Figure 6.1. This section investigates the results of the weakly non-linear model using input parameters for \mathbb{W}_1 and \mathbb{W}_2 that correspond to this triad configuration. To allow for comparisons with the experimental results, the parameters for the primary wave domain are the same as the experimental ones. Specifically

$$\mathbb{W}_0 = \{\omega_0 = 0.95 \text{ rad s}^{-1}, l_0 = -0.051 \text{ mm}^{-1}, m_0 = -0.064 \text{ mm}^{-1}\}, \quad (6.1)$$

and $N = 1.54 \text{ rad s}^{-1}$. The boundary forcing for the \mathbb{W}_0 domain is given in (5.42), which describes the four wavelength wide beam with cosine squared smoothing with $\Lambda_0 \approx 500 \text{ mm}$. The input parameters for \mathbb{W}_1 and \mathbb{W}_2 were then selected from the loci curve in Figure 6.1. Using this curve ensures that all secondary pairs satisfy both linear dispersion (1.20) and the triadic resonance condition (2.4).

A range of possible \mathbb{W}_1 were generated by linearly spacing in l_1/l_0 , across the outer blue branch on Figure 6.1. The corresponding \mathbb{W}_2 were determined from the triadic relationship in (2.4). The theoretical linear growth rates $\hat{\sigma}_{0D}$, for this range of \mathbb{W}_1 , were then calculated from (4.6). Based on these values, we examine the 50 pairs with linear growth rates lying within 90% of $\hat{\sigma}_{0D_{\max}}$. The \mathcal{M}_{2D} model was consequently run 50 times, examining each selected triad. Rather than present the results of 50 different runs, six examples are chosen that characterise how the behaviour of the instability changes with the different input parameters for \mathbb{W}_1 and \mathbb{W}_2 . The wavenumbers of the six pairs are shown on the loci curve in Figure 6.1 and their corresponding $\hat{\sigma}_{0D}$ on Figure 6.2. The pairs of secondary waves are labelled \mathbb{S}_s , where s corresponds to either a, b, c, d, e or f . The subscript notation of s is used as opposed to j (described in § 5.3.2), as here we are just considering simple triads; the subscript j is used later when considering multiple \mathbb{T} in the model. The actual parameters corresponding to \mathbb{S}_s are also given in Table 6.1. The amplitude plots showing how the instability develops for these different pairs can be seen in Figure 6.3, using an input amplitude of $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$.

Wave pair \mathbb{S}_s	ω_1 (rad s^{-1})	ω_2 (rad s^{-1})	l_1 (mm^{-1})	l_2 (mm^{-1})	m_1 (mm^{-1})	m_2 (mm^{-1})
\mathbb{S}_a	0.333	0.617	0.024	-0.075	0.106	-0.173
\mathbb{S}_b	0.344	0.606	0.027	-0.079	0.118	-0.184
\mathbb{S}_c	0.348	0.602	0.029	-0.081	0.124	-0.190
\mathbb{S}_d	0.353	0.597	0.031	-0.083	0.130	-0.196
\mathbb{S}_e	0.368	0.582	0.038	-0.090	0.153	-0.220
\mathbb{S}_f	0.400	0.550	0.063	-0.114	0.232	-0.299

Table 6.1: Table giving the input parameters of \mathbb{W}_1 and \mathbb{W}_2 for each \mathbb{S}_s used in the simple triad model. The wave vector locations of each pair can be seen by the blue marks on Figure 6.1.

The six amplitude plots in Figure 6.3 highlight the impact that the input parameters for the secondary waves have on the development of the instability. Despite having the largest theoretical growth rate $\hat{\sigma}_{0D}$ out of all the of pairs, the input parameters in \mathbb{S}_a (shown by the blue circle on Figure 6.2) did not result in any significant growth of the

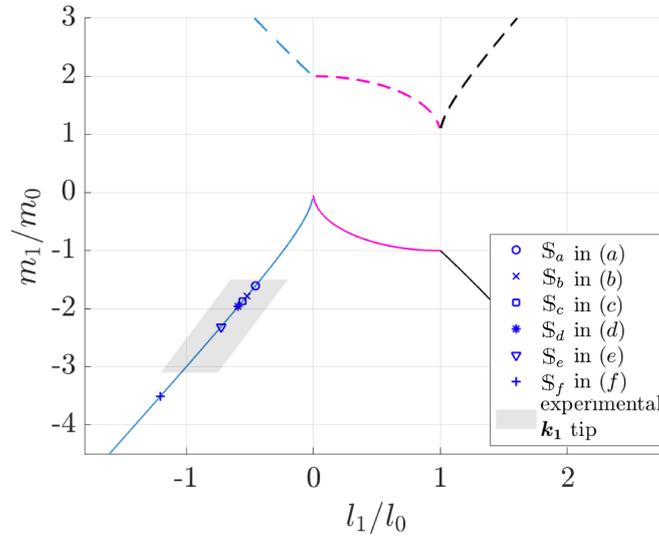


Figure 6.1: Loci curve showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector based off (2.5) and the input parameters for \mathbb{W}_0 given in (6.1). The approximate locations of the experimental \mathbf{k}_1 vector tips are given by the shaded grey region. The six blue marks correspond to each \mathbb{S}_s used for each simple triad model run, the values of which are shown in Table 6.1.

secondary waves when input into the \mathcal{M}_{2D} model. In fact, all pairs with wavenumbers that fell to the left of \mathbb{S}_a on the $\widehat{\sigma}_{0D}$ curve in Figure 6.2(a), when input into \mathcal{M}_{2D} model did not cause the primary beam to become unstable to TRI, despite $\widehat{\sigma}_{0D} > 0$. Even when the forcing amplitude was increased to $|\tilde{\Psi}_{in}| = 45 \text{ mm}^2 \text{ s}^{-1}$ – which is well above the amplitude threshold for instability (see Figure 4.10) – no pairs corresponding to $\kappa_1/\kappa_0 < 1$ caused the primary wave to become unstable when input into the \mathcal{M}_{2D} model.

The grey shaded region on Figure 6.2(a) also shows that $\kappa_1/\kappa_0 < 1$ is the cut-off for instability witnessed experimentally. On this part of the $\widehat{\sigma}_{0D}$ growth rate curve, where $\kappa_1/\kappa_0 < 1$, the length scales involved in the triad would no longer be $\kappa_0 < \kappa_1, \kappa_2$, but rather $\kappa_1 < \kappa_0 < \kappa_2$. This means that the energy transfer from the primary wave would be to one larger length scale (κ_1) and one smaller (κ_2). While these values have a positive growth rate for the zero-dimensional \mathcal{M}_{0D} model, both the experimental results and the \mathcal{M}_{2D} model show that this energy transfer to larger length scales does not occur.

The behaviour of the triad then changes considerably over a small range of input parameters. Looking at Table 6.1, the frequency only changes by $\pm 0.09 \text{ rad s}^{-1}$ between the inputs \mathbb{S}_b , \mathbb{S}_c and \mathbb{S}_d , yet the amplitude of the secondary waves in \mathcal{M}_{2D} becomes much larger as we move down to the right across the $\widehat{\sigma}_{0D}$ curve. Not only

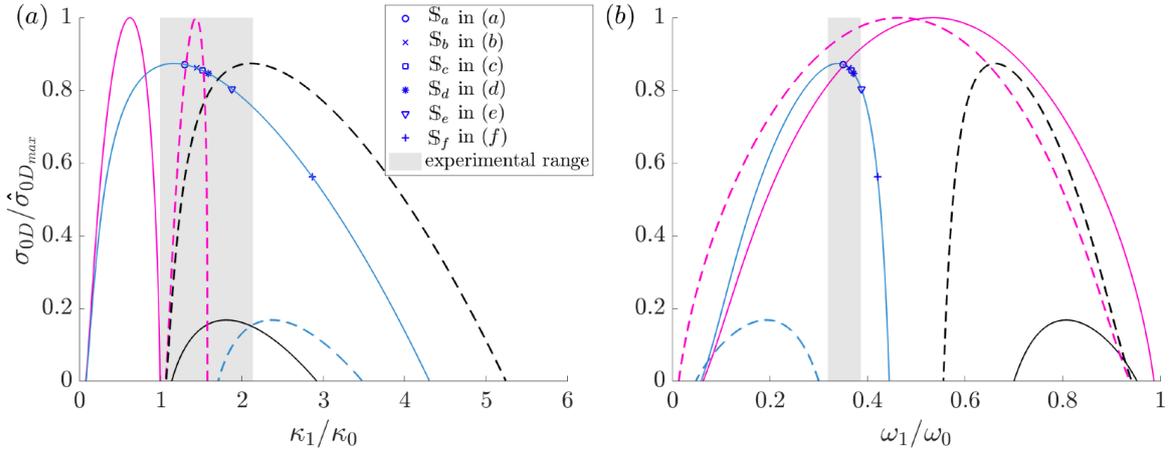


Figure 6.2: Linear theoretical growth rates σ_{0D} calculated using (4.6) for every possible triad configuration shown in Figure 6.1 with an input amplitude of $|\tilde{\Psi}_{in}| = 35 \text{ mm}^2 \text{ s}^{-1}$. The approximate locations of the experimental parameter range κ_1/κ_0 and ω_1/ω_0 are given by the shaded grey regions. The six blue marks correspond to the wave pairs \mathbb{S}_s used as inputs for each ‘simple triad’ in the \mathcal{M}_{2D} (the values of which are shown in Table 6.1) along with their corresponding $\hat{\sigma}_{0D}$.

do these three plots in Figure 6.3(b), (c) and (d) show a significant increase in the amplitude of the secondary waves, they also exhibit similar quasi-periodic modulations to those witnessed experimentally. Yet in the experiments, in addition to the amplitude modulation, the origin of the \mathbb{W}_1 beam migrated significantly over the central region of \mathbb{W}_0 . The snapshots in Figure 6.4 show the results of the \mathcal{M}_{2D} using \mathbb{S}_d , corresponding to the amplitude plot in Figure 6.3(d). What these snapshots show, is that while the amplitude of \mathbb{W}_1 modulates slightly over time, the origin of \mathbb{W}_1 is located towards the top of the primary beam and does not exhibit a large spatial migration over time.

As we then move further down the $\hat{\sigma}_{0D}$ curve to \mathbb{S}_e and \mathbb{S}_f the \mathcal{M}_{2D} amplitude plots in Figure 6.3(e) and (f) lose the oscillatory behaviour and resemble more the steady state results from the \mathcal{M}_{0D} shown in Figure 4.12. However, care should be taken with drawing comparisons between the amplitude plots in Figure 6.3 and Figure 4.12, as in the \mathcal{M}_{0D} model there is no spatial domain to average over, while the spatial averaging in the \mathcal{M}_{2D} model is over an arbitrary visualisation region \mathcal{D}_w . This means that the amplitude ratios of the triadic waves are not comparable between the two models. It is also worth nothing here that the total energy content of the domain in the \mathcal{M}_{2D} model is dependent not only on the amplitude of the waves, but also on the fraction of the domain containing waves. Hence, while these amplitude plots in Figure 6.3(e) and (f) appear to show a significant decrease in the total energy, this is actually due to the wave motion covering less of the domain.

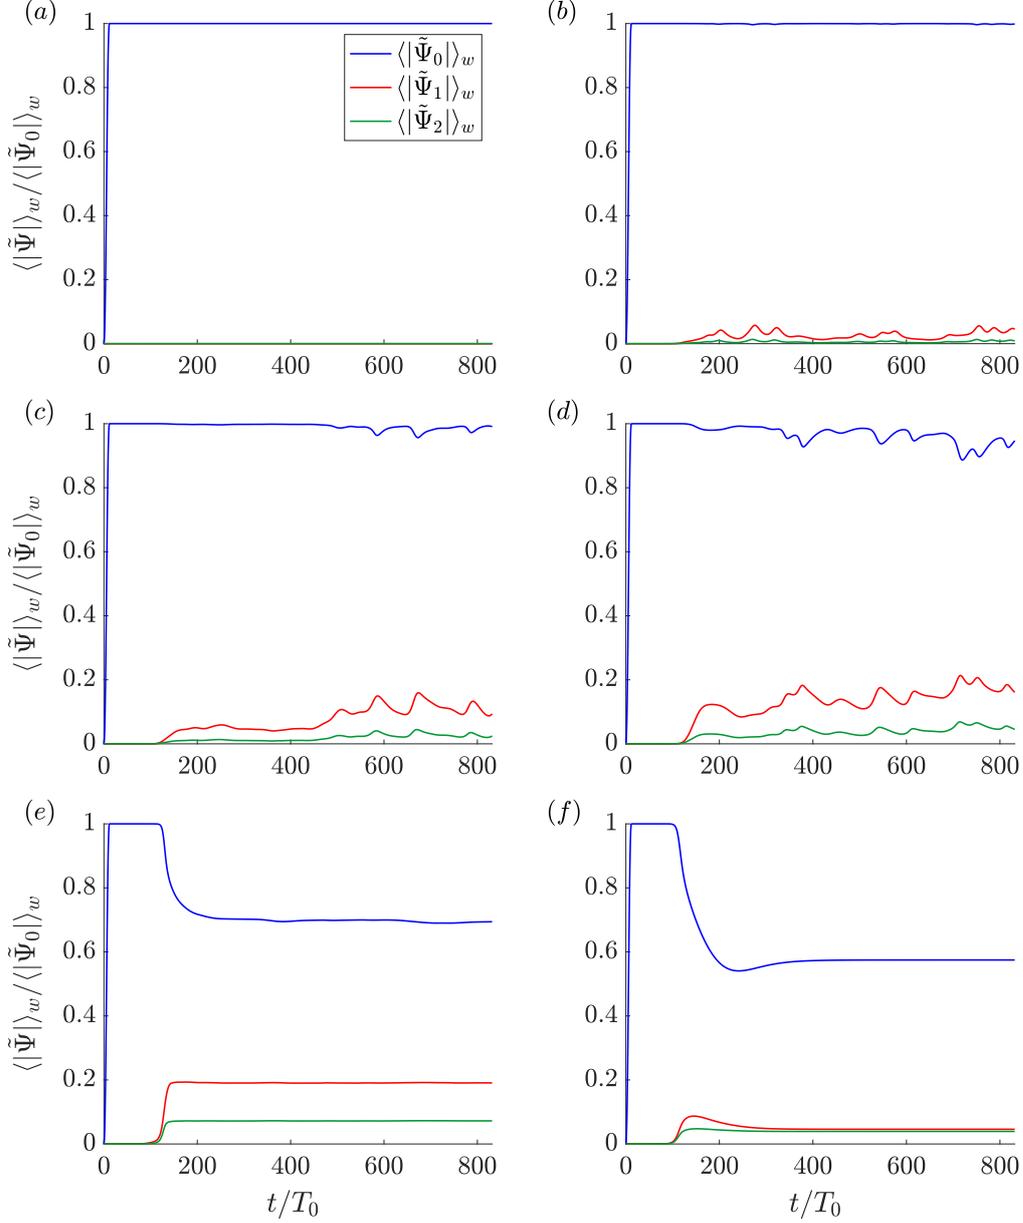


Figure 6.3: Six amplitude plots generated using the two-dimensional weakly non-linear model. Each plot corresponds to a simple triad with three domains ($n_d = 3$). The parameters for \mathcal{S}_s corresponding to each plot are given in Table 6.1. (a) \mathcal{S}_a (b) \mathcal{S}_b (c) \mathcal{S}_c (d) \mathcal{S}_d (e) \mathcal{S}_e (f) \mathcal{S}_f . The forcing amplitude for \mathbb{W}_0 in all plots is $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$ and $t_{\text{end}} = 90$ minutes ($t/T_0 \approx 800$). The amplitudes are calculated over the whole visualisation window \mathcal{D}_w and have been normalised against the maximum amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$.

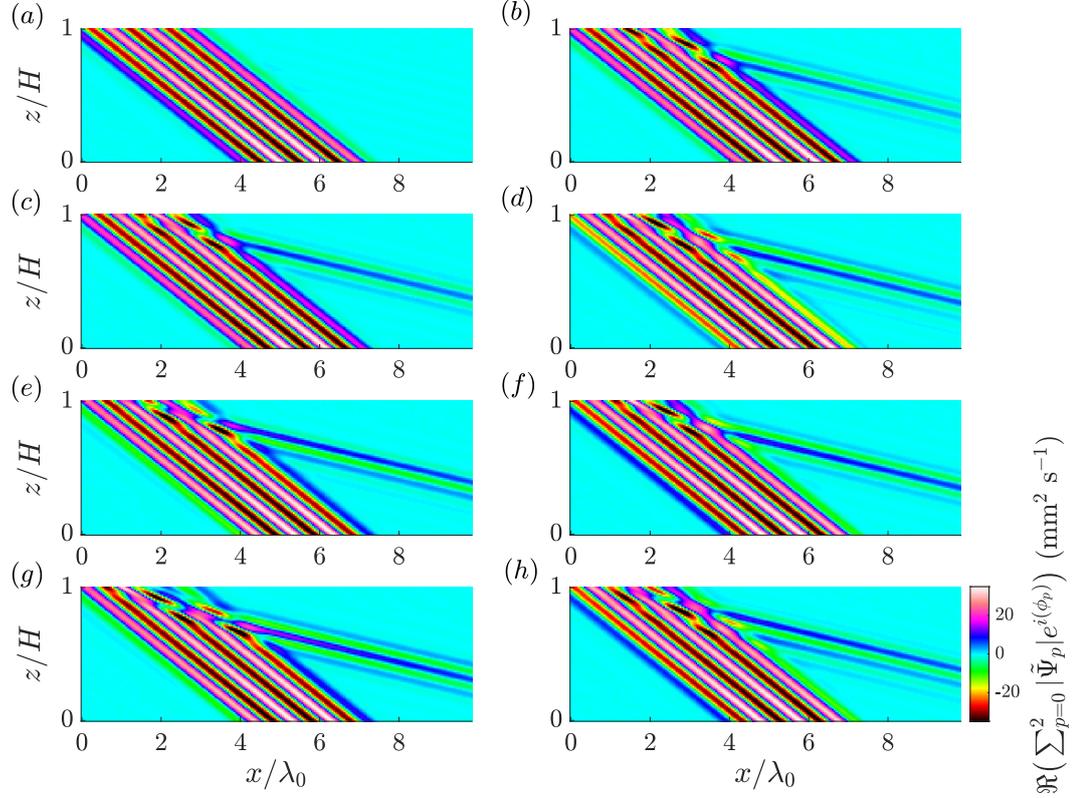


Figure 6.4: Sequence of snapshots at roughly 400 s ($t/T_0 \approx 61$) apart, showing the superposition of the three domains in the model. The parameters for the secondary waves correspond to \mathbb{S}_d . Each domain is multiplied by its respective fast timescales $e^{i(\mathbf{k}_p \cdot \mathbf{x} - \omega_p t)}$. The forcing amplitude for \mathbb{W}_0 is $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$. The amplitude of each wave is shown in Figure 6.3(d). The timing of each snapshot is (a) $t/T_0 = 121$, (b) $t/T_0 = 151$, (c) $t/T_0 = 272$, (d) $t/T_0 = 333$, (e) $t/T_0 = 424$, (f) $t/T_0 = 484$, (g) $t/T_0 = 545$, (h) $t/T_0 = 606$.

Looking back at Table 6.1, the value of ω_1 increases across the pairs from \mathbb{S}_a to \mathbb{S}_f . Based on the dispersion relationship (1.20), an increase in the ratio ω_1/N results in a beam with a steeper angle. In this case, as the frequency of ω_1 for each pair increases, the \mathbb{W}_1 beam becomes steeper and closer to being parallel with \mathbb{W}_0 , thus increasing the spatial region over which the two beams interact. While this increase in ω_1 results in a decrease to ω_2 and therefore a shallower angle for \mathbb{W}_2 , as \mathbb{W}_2 and \mathbb{W}_0 are propagating the same direction, the time in which these waves interact is greater than for \mathbb{W}_1 and \mathbb{W}_0 and it is not the limiting factor in the interaction. This increase in ω_1 means the non-linear interactions are able to occur for longer, leading to this more dramatic decay of the primary beam due to the increased energy transfer to the secondary waves. As discussed in § 4.3.2, the oscillatory behaviour shown in Figures 6.3(b), (c) and (d) is

believed to be a result of this finite-time interaction. When a packet of energy in \mathbb{W}_1 exits the underlying primary beam, the energy exchange between the triad is broken. If the triad has not reached a steady equilibrium, then the triadic interaction decays, thus causing an amplitude growth in \mathbb{W}_0 . The angle of \mathbb{W}_1 in pairs \mathbb{S}_d , \mathbb{S}_c and \mathbb{S}_d is shallower are further away from being parallel with \mathbb{W}_0 in comparison to \mathbb{S}_e and \mathbb{S}_f . This shallower angle of \mathbb{W}_1 means that the triad interaction occurs for less time in the \mathcal{M}_{2D} model. This results in shorter triadic interactions and consequently the witnessed oscillatory behaviour.

If the \mathcal{M}_{2D} was run again using the same inputs, qualitatively the results would look similar, however quantitatively the structure of the oscillations would be different. This is because each domain is initiated with a different random seed which triggers the instability. As with the experiments, this random noise causes different perturbations to grow, hence explaining why the behaviour of these modulations is not exactly repeatable and finding a dominant period of oscillation proved unsuccessful.

Experimentally, the amplitudes of each wave in the triad undergo strong modulations, as shown in Figure 4.5. The shaded regions on Figure 6.2 show the approximate locations of the experimental wavenumbers and frequencies. The frequency band is much narrower compared to the wavenumber band, as it is determined from the DMD which approximates the dominant frequencies of the secondary waves with delta functions generating a narrower range compared to the FFTs used to calculate the spatial vectors. The frequency band shows that the frequencies selected experimentally correspond just to the right of the maximum growth rate $\widehat{\sigma}_{0D_{\max}}$. This experimental band combines the largest values of $\widehat{\sigma}_{0D}$ without permitting transfer to larger scales, which would occur for parameters lying to the left of this peak where $\kappa_1/\kappa_0 < 1$. Figure 6.3 also shows that, for the given set-up, this shaded region also corresponds to the range of parameters that permits this oscillatory behaviour in the \mathcal{M}_{2D} model.

Finally we examine the amplitude threshold for the \mathcal{M}_{2D} model using the simple triad. As shown by Figure 6.3, there is a range of secondary wave parameters in the \mathcal{M}_{2D} model for which the instability will not be triggered. This threshold changes with different forcing amplitudes. Figure 6.5 shows multiple simple triad runs over a range of input amplitudes. Figure 6.5(a) is given as a function of κ_1/κ_0 , while (b) is given as a function of ω_1/ω_0 . The blue crosses indicate that the flow does not become unstable, while the red crosses indicate that the instability is triggered.

Figure 6.5 shows that as the input amplitude is increased, the secondary wave parameters that permit instability, shift. With respect to the $\widehat{\sigma}_{0D}$ curve in Figure 6.2, as the amplitude is increased, the cut-off κ_1/κ_0 for instability moves further to the

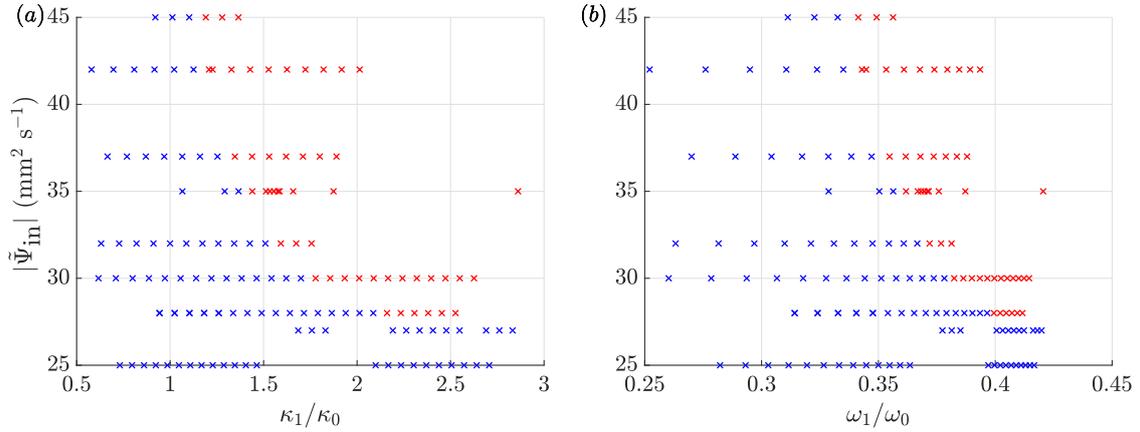


Figure 6.5: Amplitude threshold for instability in the $\mathcal{M}_{2\text{D}}$ model as a function of input parameters for \mathbb{W}_1 using the simple triads. The blue crosses represent model runs that do not become unstable while the red crosses indicate the growth of TRI. (a) Amplitude threshold as a function of κ_1/κ_0 (b) Amplitude threshold as a function of ω_1/ω_0 . Again, parameters for \mathbb{W}_0 are given in (6.1).

left, closer to $\hat{\sigma}_{0\text{D}_{\text{max}}}$ but never crossing the $\kappa_1/\kappa_0 < 1$ boundary. Figure 6.5 also shows that the ‘critical amplitude’ for this \mathbb{W}_0 forcing lies around $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$, as below this amplitude forcing no combination of input parameters causes the primary beam to become unstable. This amplitude threshold is lower than the experimental one of $\langle |\tilde{\Psi}_0| \rangle_r \approx 35 \text{ mm}^2 \text{ s}^{-1}$ shown in Figure 4.11. This in itself is unsurprising as the background noise and external perturbations of the experimental flow will be greater, suppressing instability at lower amplitudes. There is also a difference in how $\langle |\tilde{\Psi}_0| \rangle_r$ is measured compared $|\tilde{\Psi}_{\text{in}}|$, so while they both represent the amplitude of the primary wave near the wavemaker, they are not directly comparable measures.

6.2.3 Inner loci branch

Despite the experiments always selecting the outer branch for the resonant waves, the overall maximum growth from the $\mathcal{M}_{0\text{D}}$ model, $\hat{\sigma}_{0\text{D}_{\text{max}}}$, corresponds to the inner pink curve of the resonant loci diagram, as shown in Figure 6.2. This is true for both the growth rate expression for an infinite plane wave (2.22) and for the finite-width expression (2.24) used to generate the growth rate curves shown in this thesis. As viscosity is considered for these zero-dimensional growth rate curves, this inner pink branch is also the only branch that satisfies secondary wave frequencies corresponding to Parametric Sub-harmonic Instability (PSI), where $\omega_{1,2}/\omega_0 = 0.5$ (in the inviscid limit the outer blue branches also access the $\omega_{1,2}/\omega_0 = 0.5$, see Figure 11 in Bourget et al. (2013) for details). This section therefore examines results from the $\mathcal{M}_{2\text{D}}$ model

using a simple triad looking only at input secondary pairs corresponding to this inner pink branch.

There are two key differences to note regarding \mathbb{W}_1 when the triad corresponds to the inner branch as opposed to the experimentally selected outer branch. The first is that the horizontal orientation of the \mathbf{k}_1 wave vector changes direction. This results in a triad where the horizontal component of all the wave vectors (and group velocities) are in the same direction, as opposed to the outer branch, where l_0 and l_2 are opposing l_1 . The second, is that $\kappa_1 < \kappa_0 < \kappa_2$, meaning the energy transfer would be to one smaller length scale κ_2 and one larger scale κ_1 .

Using the same method outlined in § 6.2.2, 45 pairs of secondary waves were generated that corresponded to the triads on the inner branch of the loci curve, again using the \mathbb{W}_0 parameters given in (6.1). The model was then run 45 times for each simple triad using an input amplitude of $|\tilde{\Psi}_{\text{in}}| = 45 \text{ mm}^2 \text{ s}^{-1}$, again with $\Lambda_0 = 4\lambda_0$ as shown in Figure 4.1. However, for all of the input triads, none caused the primary wave to become unstable. The input amplitude was then increased to $|\tilde{\Psi}_{\text{in}}| = 150 \text{ mm}^2 \text{ s}^{-1}$ – approximately five times the amplitude required for instability in the \mathcal{M}_{2D} model for secondary wave parameters on the outer branch – and still the amplitude of the secondary waves did not surpass $10^{-3} \text{ mm}^2 \text{ s}^{-1}$ (the background noise was of magnitude 5×10^{-8}). As it was clear that the outer branch parameters would trigger TRI well before this inner branch was accessed, we examine the structure of the resonant waves on the inner branch at $|\tilde{\Psi}_{\text{in}}| = 150 \text{ mm}^2 \text{ s}^{-1}$. As no significant behavioural change was noticed with changing input parameters, we examine only one simple triad corresponding to $\hat{\sigma}_{0D_{\text{max}}}$, which we will refer to as $\mathbb{S}_{\text{inner}}$. The values for this input pair can be seen in Table 6.2, while the location and orientation of the wave vectors can be seen on the loci curve in Figure 6.6. The frequencies of the secondary waves in $\mathbb{S}_{\text{inner}}$ are close to $\omega_{1,2}/\omega_0 = 0.5$, corresponding to the specific case of Parametric Sub-harmonic Instability (PSI).

Wave pair	ω_1 (rad s ⁻¹)	ω_2 (rad s ⁻¹)	l_1 (mm ⁻¹)	l_2 (mm ⁻¹)	m_1 (mm ⁻¹)	m_2 (mm ⁻¹)
$\mathbb{S}_{\text{inner}}$	0.492	0.458	-0.016	-0.036	0.048	-0.115

Table 6.2: Table giving the input parameters of \mathbb{W}_1 and \mathbb{W}_2 for $\mathbb{S}_{\text{inner}}$ used in the simple triad \mathcal{M}_{2D} model.

The snapshots in Figure 6.7(a), (c) and (e) show the \mathcal{M}_{2D} model at $t = 2000 \text{ s}$ ($t/T_0 \approx 300$), using the secondary wave pair $\mathbb{S}_{\text{inner}}$. The input amplitude for the primary wave is $|\tilde{\Psi}_{\text{in}}| = 150 \text{ mm}^2 \text{ s}^{-1}$. As the amplitudes of the secondary waves are too small

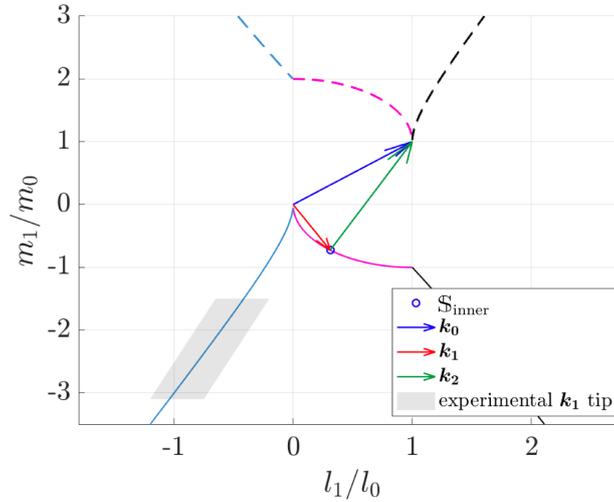


Figure 6.6: Loci curve showing all the possible locations for the tip of the \mathbf{k}_1 wavenumber vector based off (2.5) and the input parameters for \mathbb{W}_0 given in (6.1). The approximate locations of the experimental \mathbf{k}_1 vector tips are given by the shaded grey region. The blue mark on the inner branch corresponds to $\mathbb{S}_{\text{inner}}$ used for the $\mathcal{M}_{2\text{D}}$ model run. The orientation of the three wave vectors corresponding to $\mathbb{S}_{\text{inner}}$ have also been marked. The wavenumbers and frequencies for $\mathbb{S}_{\text{inner}}$ are given in Table 6.2.

to be seen if the triad domains are superposed (as they were in Figure 6.4), each wave in the triad is shown separately. Figure 6.7(a) shows the \mathbb{W}_0 domain, while (c) and (e) show the \mathbb{W}_1 and \mathbb{W}_2 domains respectively. The white dashed lines overlaid on (c) and (e) give the outline of primary wave beam. Interestingly, the detectable amplitude signature of \mathbb{W}_1 does not properly begin until the left outer flank of the primary beam. Between the white lines in (c) and (e), where the largest energy transfer is expected, the amplitude of the secondary waves (especially \mathbb{W}_1) is almost negligible. This could be due to all the group velocities in all three domains being directed towards the left, which may have the effect of pushing the triad interaction out of the primary beam before enough energy can be transferred.

To further explore if this was the case, the $\mathcal{M}_{2\text{D}}$ model was run, again with an input amplitude of $|\tilde{\Psi}_{\text{in}}| = 150 \text{ mm}^2 \text{ s}^{-1}$ and the same \mathbb{W}_0 parameters, but for a beam width of $\Lambda_0 = 6\lambda_0$. Specifically, the values of A, B, C, D are respectively $5\pi/|l_0|, 7\pi/|l_0|, 15\pi/|l_0|, 17\pi/|l_0|$ for the spatial boundary forcing in (5.42). The snapshot in Figure 6.7(b), (d) and (f) correspond to this run.

Despite the wider beam width, the secondary waves shown in Figure 6.7(d) and (f) also do not show a noticeable signature until they reach the left flank of the primary wave. Again, in the middle of the primary beam, where you would expect the largest growth and transfer of energy, the amplitude of the secondary waves appears negligible.

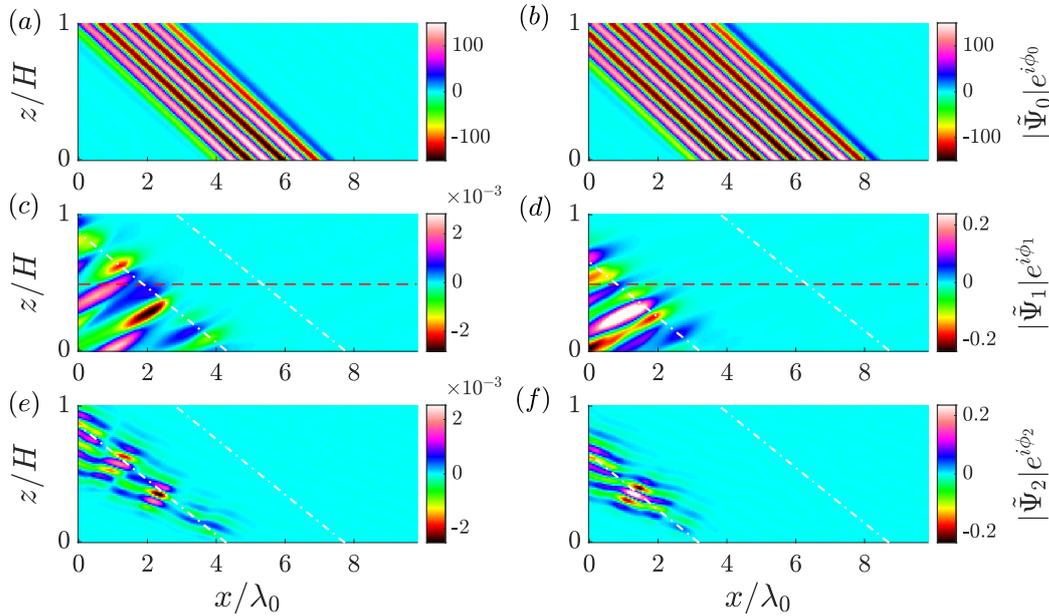


Figure 6.7: Two different simple triad model runs using the \mathcal{M}_{2D} model, both with the same triad configuration corresponding $\mathbb{S}_{\text{inner}}$. Each model run is decomposed to show the individual fields for each wave in the triad. (a) and (b) $|\tilde{\Psi}_0|e^{i\phi_0}$ (c) and (d) $|\tilde{\Psi}_1|e^{i\phi_1}$ (e) and (f) $|\tilde{\Psi}_2|e^{i\phi_2}$. The first model run in (a), (c) and (e) is for a primary beam width of $\Lambda_0 = 4\lambda_0$, while the second model run shown in (b), (d) and (f) is for a beam width of $\Lambda_0 = 6\lambda_0$. The width of the primary beam is overlaid in white for the secondary wave domains. The red lines in (c) and (d) indicate the horizontal transect shown in Figure 6.8.

The amplitude of the secondary waves is, however, two orders of magnitude larger than the corresponding \mathbb{W}_1 and \mathbb{W}_2 for the $\Lambda_0 = 4\lambda_0$ run. This is shown in more detail in Figure 6.8 which shows the logarithm of $\tilde{\Psi}_1$ across the horizontal cross section marked by the red dashed line in Figure 6.7(c) and (d).

In Figure 6.8 we see the logarithmic amplitude of \mathbb{W}_1 for both beam widths across a horizontal transect of the flow (shown by the red dashed line). In the centre of both beams, the background noise begins to show an exponential growth in amplitude. For the smaller beam of $4\lambda_0$, shown in green, the amplitude peaks just outside of the left boundary of \mathbb{W}_0 . The purple line (corresponding to the wider beam of $6\lambda_0$) shows a similar structure to the green line, just shifted to the right by approximately one wavelength, coincident with the difference of one wavelength in width on each flank between the two beams. The purple line has a similar growth rate to the green line but is able to reach a larger amplitude before exiting the spatial extent of the underlying primary beam. The same structure is also seen for lower amplitude forcing, although that has not been shown here.

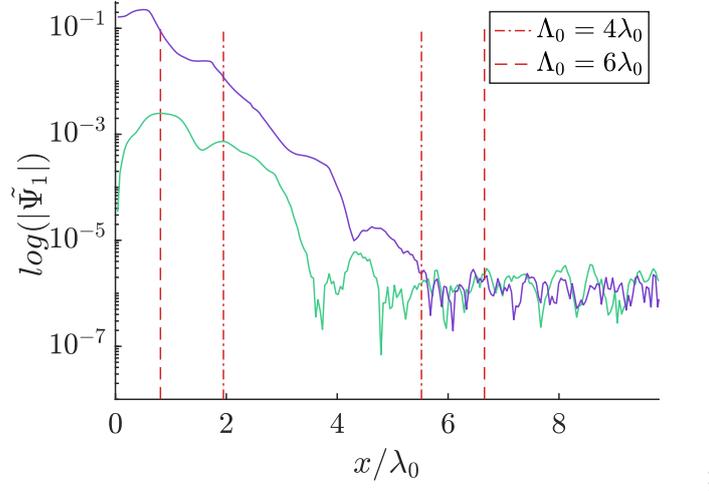


Figure 6.8: Logarithm of $|\tilde{\Psi}_1|$ across the horizontal cross section marked in red in Figure 6.7(c) and (d). The green line is the cross section for the $\Lambda_0 = 4\lambda_0$ model run shown in (c) and the purple line is the cross section for the $\Lambda_0 = 6\lambda_0$ model run shown in (d). The red lines indicate the boundary of the primary wave beam \mathbb{W}_0 .

It appears, therefore, that while energy is transferred from the primary wave and \mathbb{W}_1 shows a linear exponential growth, as all of the horizontal components of the group velocities in the triad are orientated to the same direction, the interaction is advected out of the primary beam before the secondary waves can grow to a significant amplitude. In the limit of an infinite plane wave, this branch may, therefore, dominate the triadic interaction as the triadic energy exchange would be not a function of space.

6.3 Multiple secondary waves

6.3.1 Overview

Despite some of the experimentally observed amplitude modulations being captured in the simple triad runs for a range of input parameters on the outer branch, it is not possible to capture the experimentally observed modulations of \mathbb{W}_1 and \mathbb{W}_2 in Fourier space with only one set of input parameters. In light of this, the model was adapted to be able to include multiple secondary wave pairs, where each pair can interact with the primary beam as shown in (5.45). This model set-up will be referred to as ‘multiple’ triads. This section presents results from the \mathcal{M}_{2D} model using this configuration. Again, the input parameters for \mathbb{W}_0 are given in (6.1) and the boundary forcing in (5.42) with a beam width of $\Lambda_0 = 4\lambda_0$. For the multiple triad runs, the secondary

waves were comprised of subsets

$$\begin{aligned}\mathbb{W}_1^{(g)} &= \{\mathbb{W}_{1_1}, \mathbb{W}_{1_2}, \dots, \mathbb{W}_{1_j}, \dots, \mathbb{W}_{1_g}\}, \\ \mathbb{W}_2^{(g)} &= \{\mathbb{W}_{2_1}, \mathbb{W}_{2_2}, \dots, \mathbb{W}_{2_j}, \dots, \mathbb{W}_{2_g}\}.\end{aligned}\tag{6.2}$$

where, for the results presented in this section $g = 60$. As each \mathbb{S}_j always shares the same \mathbb{W}_0 , each triad \mathbb{T}_j is defined as

$$\mathbb{T}_j = \{\mathbb{W}_0, \mathbb{W}_{1_j}, \mathbb{W}_{2_j}\} = \{\mathbb{W}_0, \mathbb{S}_j\}.\tag{6.3}$$

As $g = 60$, there were 60 \mathbb{S} in the flow. The total number of domains in the model run was, therefore, $n_d = 2g + 1 = 121$. Again, the pairs were selected from the loci curve, this time by evenly spacing in l_1/l_0 across both the outer and inner branch. These pairs were then transposed onto their respective zero-dimensional growth rate curves and the top 60% of both $\hat{\sigma}_{0D}$ and $\hat{\sigma}_{0D}$ were selected. Figure 6.9 shows these zero-dimensional growth rate curves with all the input \mathbb{W}_1 parameters for the multiple triad run shown by the blue circles. The input amplitude was $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$, as this was the amplitude threshold found for the simple triads in the \mathcal{M}_{2D} model shown in Figure 6.5. In § 6.3.2 and § 6.3.3 we look at the results of this multiple triad model run in physical and Fourier space respectively.

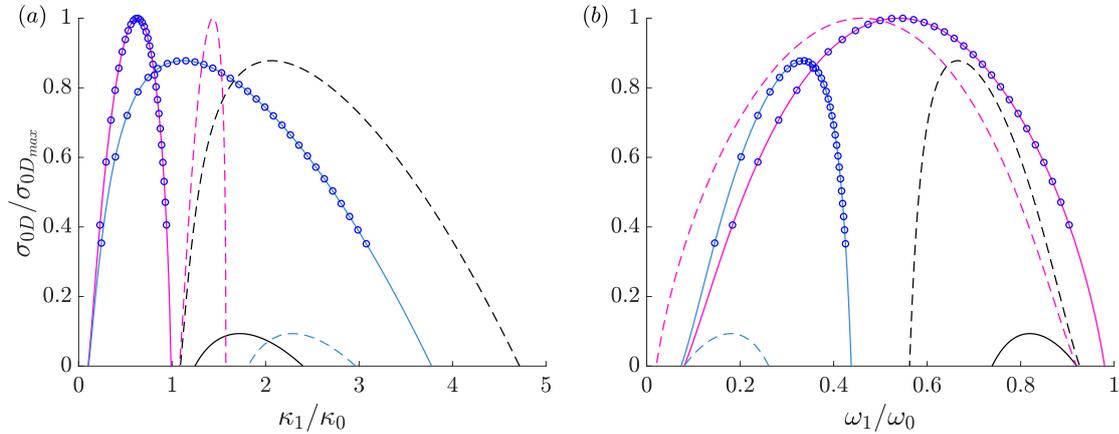


Figure 6.9: Theoretical zero-dimensional growth rates calculated using (4.6) for every possible triad configuration shown in Figure 6.1 using an input amplitude of $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$. The blue marks correspond to the top 60% of both $\hat{\sigma}_{0D}$ and $\hat{\sigma}_{0D}$ and indicate all the \mathbb{W}_1 parameters used for the multiple triad \mathcal{M}_{2D} model run.

6.3.2 Temporal and spatial modulation

The multiple triad model run outlined above in § 6.3.1 was run for 90 minutes ($t/T_0 = 818$) to allow for comparison with the experimental results. Figure 6.10 presents snapshots from the run at approximately 350 s ($t/T_0 \approx 60$) apart.

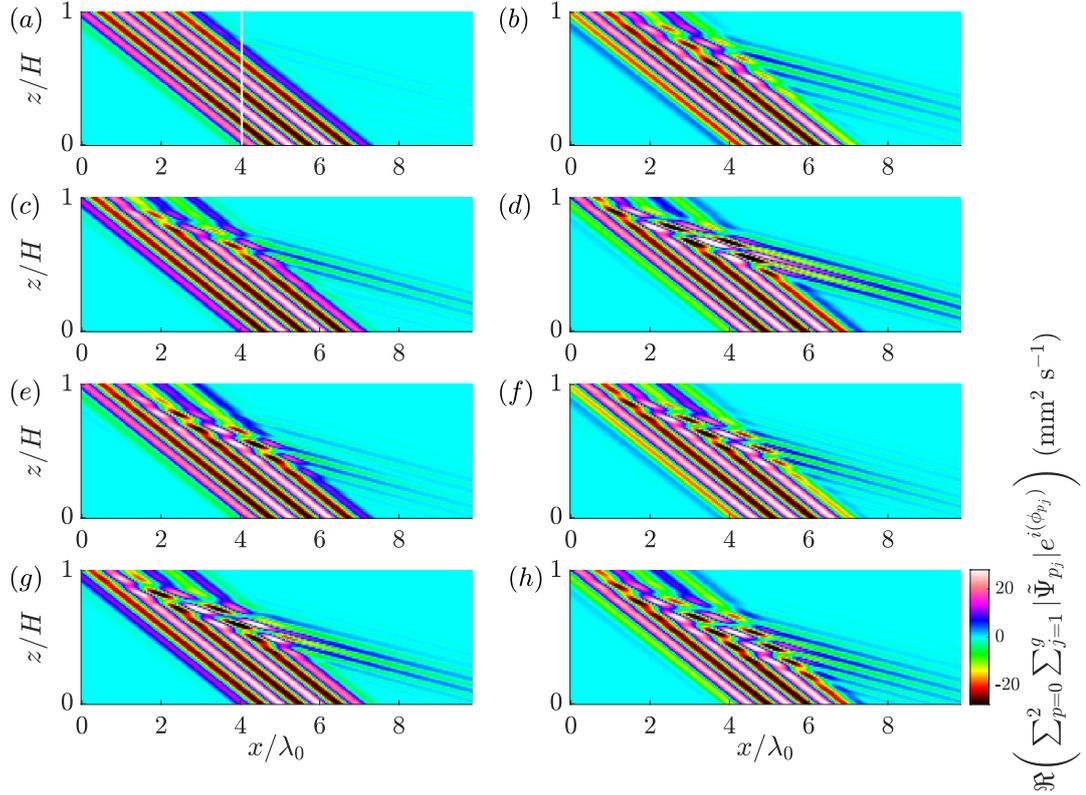


Figure 6.10: Sequence of snapshots at approximately 350 s ($t/T_0 \approx 60$) apart, showing the superposition of all of the domains in the model. The input parameters of \mathbb{W}_1 for each domain can be seen by the blue circles in 6.9. Each domain is multiplied by its respective fast timescale $e^{i(\mathbf{k}_{p_j} \cdot \mathbf{x} - \omega_{p_j} t)}$. The forcing amplitude for \mathbb{W}_0 is $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$. The timing of each snapshot is (a) $t/T_0 = 151$, (b) $t/T_0 = 212$, (c) $t/T_0 = 258$, (d) $t/T_0 = 318$, (e) $t/T_0 = 394$, (f) $t/T_0 = 455$, (g) $t/T_0 = 500$, (h) $t/T_0 = 561$. The vertical white line in (a) indicates the location over which the time-series for the spectrogram shown in Figure 6.12 was taken.

The snapshots in Figure 6.10 show a very different development of the instability compared with the simple triad snapshots in Figure 6.4. When considering the simple triad, the location of the instability did not vary significantly across the height of the primary beam. Moreover, the width of the \mathbb{W}_1 beam remained narrow. This is in contrast to the experimental snapshots shown in Figure 4.4, where the \mathbb{W}_1 beam is wide and its origin can be seen to migrate across the full central region of the primary

beam. Both this increased width and variable spatial location of \mathbb{W}_1 are captured more convincingly in the multiple-frequency snapshots shown here.

In Figure 6.10(a) the structure of the instability can be seen to develop around the top of the primary beam. In the next snapshot 400 s ($t/T_0 = 61$) later, the width of \mathbb{W}_1 increases significantly, also showing more structure in the lower region of the domain. Then in (c) the width of \mathbb{W}_1 decreases again as one perturbation appears to dominate the interaction. The amplitude of \mathbb{W}_1 then increases further in (d) due to a large transfer of energy from the primary wave. Despite this, the \mathbb{W}_1 then decays and Figure 6.10(e) shows two different wavelengths corresponding to the \mathbb{W}_1 beam, highlighting the presence of different triadic perturbations present in the flow. This cyclic behaviour continues and over time the location of the instability migrates further down the primary beam.

What these snapshots also show is that above the triadic interaction, the amplitude of the primary wave decays significantly. This extraction of energy from the primary beam above the region of energy exchange is not witnessed so acutely in the experiments. Because of this, the location of instability migrates more experimentally, as the primary beam still has sufficient amplitude near the free surface to allow growth of a perturbation. In both the simple triad and multiple-frequency triads, perturbations fail to grow again near the upper boundary after the primary beam has decayed significantly. This decay in amplitude of the primary beam can be seen in Figure 6.11 which presents the amplitude of triadic waves during the full multiple triad run. In Figure 6.11 the amplitude of \mathbb{W}_1 and \mathbb{W}_2 are determined by summing across their respective g domains. All the amplitudes have been integrated over the full spatial domain \mathcal{D}_w .

As with the simple triad plots shown in Figure 6.3(b), (c) and (d), the amplitudes of the triadic waves modulate continuously in Figure 6.11 over the full run. What is different, however, is that the typical period of oscillation is longer and the modulations are smoother than in simple triads in Figure 6.3(b), (c) and (d) and the experimental results in Figure 4.5. Experimentally, there is a continuous spectrum of wavenumbers and frequencies present in the background noise, allowing for a much wider range of perturbations to grow and interact with each other. In this multiple triad model run, there is only a finite number of perturbations that are able to grow. This could explain why the typical period of the quasi-periodic modulations is longer, as there is less ‘beating’ between the different triads.

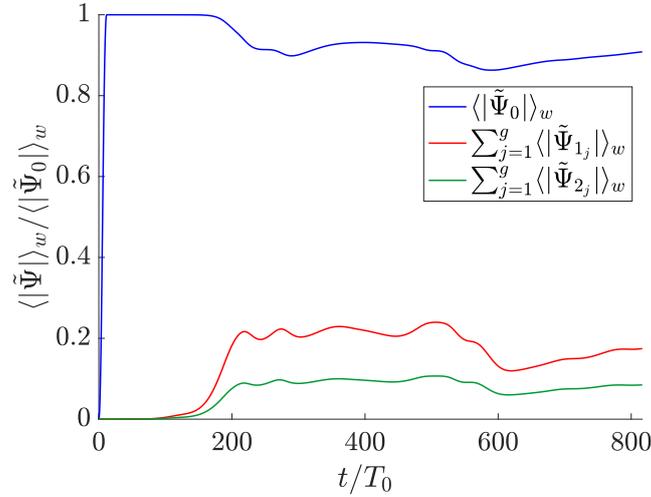


Figure 6.11: The magnitude of the reduced streamfunction for the multiple triad \mathcal{M}_{2D} model run outlined in § 6.3.1, with $\langle |\tilde{\Psi}_0| \rangle_w$ (blue) $\sum_{j=1}^g \langle |\tilde{\Psi}_{1_j}| \rangle_w$ (red) and $\sum_{j=1}^g \langle |\tilde{\Psi}_{2_j}| \rangle_w$ (green). The amplitudes were calculated over the whole visualisation window \mathcal{D}_w and have been normalised against the maximum amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$. The model was run for 90 minutes with a forcing amplitude of $|\tilde{\Psi}_{in}| = 28 \text{ mm s}^{-1}$.

6.3.3 Frequency and wavenumber modulation

The next stage of the analysis is to examine how the multiple triad \mathcal{M}_{2D} run also developed in Fourier space. What is clear experimentally, is how the amplitude modulations are connected to the growth and decay of different triadic perturbations that span a small range of parameters in Fourier space. The snapshot in Figure 6.10(e) reflected this in the \mathcal{M}_{2D} model, as two different \mathbf{k}_1 vectors could be seen simultaneously.

For the results presented in this section (§ 6.3.3), the same parameters for the multiple triad model run presented in § 6.3.1 are used, however – as will be discussed in § 6.3.4 – the amplitude of the secondary waves were suppressed between $800 < t < 1100$ s ($121 < t/T_0 < 167$). This suppression was done to obtain a more accurate measure for the growth rate for the \mathcal{M}_{2D} model.

The same analysis used in § 4.3.3 was performed on the multiple triad run to determine how the parameters of the secondary waves fluctuated in Fourier space. First, the horizontal and vertical gradients of $\tilde{\Psi}$ were taken for the full flow field (obtained from the linear superposition of all n_d domains). A time-series of both the horizontal and vertical gradients was then created across the vertical transect marked by the white line in Figure 6.10(a). Using a time window of 150 s ($t/T_0 \approx 22.7$) with an overlap of 148 s ($t/T_0 \approx 22.4$), FFTs were taken across each vertical position for

both the horizontal and vertical gradients. Finally, the root mean square (RMS) of the gradient components was taken and the results were collapsed across the vertical dimension to provide a two-dimensional frequency spectrogram. The same model run was also analysed using the Dynamic Mode Decomposition (DMD) algorithm (discussed in § 3.9.2) using a time window of 20 s ($t/T_0 \approx 3$). The spectrogram can be seen in Figure 6.12 with the dominant frequencies identified from the DMD algorithm overlaid in white.

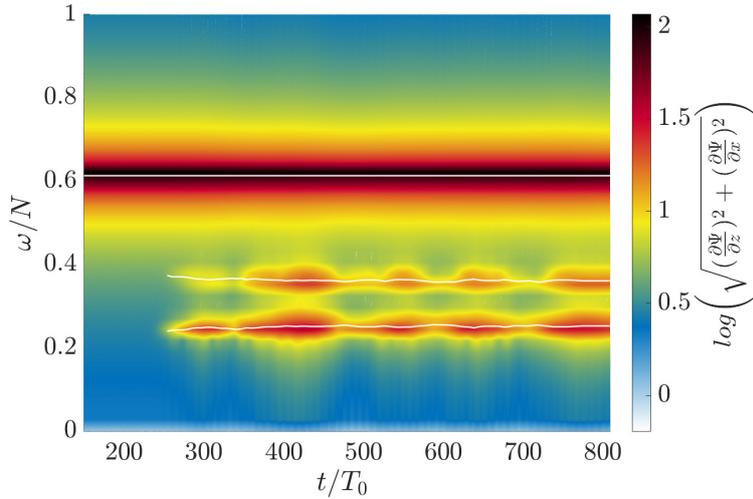


Figure 6.12: Time frequency spectra computed over the vertical transect shown in white on Figure 6.10(a) for the multiple triad \mathcal{M}_{2D} model run outlined in § 6.3.1, with an input amplitude of $|\tilde{\Psi}_{in}| = 28 \text{ mm s}^{-1}$. For this run the amplitude of the secondary waves was suppressed between $800 < t < 1100 \text{ s}$ ($121 < t/T_0 < 167$) so the first 990 s ($t/T_0 = 150$) are removed for clarity. The spectral density is given as the logarithm of the RMS field of both the horizontal and vertical gradient of the reduced streamfunction. The dominant frequencies for each experiment obtained from the DMD algorithm are overlaid in white.

Figure 6.12 shows a strong resemblance to the experimental spectrograms presented in Figure 4.6. Both secondary waves show periods with a strong signal and a wider spectrum of waves combined with periods of relatively weak signal. As with the experimental results, these fluctuations are synchronous between ω_1 and ω_2 . The dominant frequencies determined by the DMD are also seen to fluctuate over time, following the path of the spectrogram, although the modulations are not quite as large as those seen in the experimental results. Despite this modulation, the frequencies obtained from the DMD again satisfy the temporal triadic relationship $\omega_1 + \omega_2 = \omega_0$ to within $\pm 0.002 \text{ rad s}^{-1}$ for all time. As noted previously, this relationship is not an inherent function of the DMD algorithm.

Despite the qualitative similarities between the \mathcal{M}_{2D} model and experimental spectrograms, the range of frequencies for the secondary waves is different. In Figure 6.12, $0.39 < \omega_1/\omega_0 < 0.41$ ($0.24 < \omega_1/N < 0.26$) while experimentally $0.35 < \omega_1/\omega_0 < 0.39$ ($0.20 < \omega_1/N < 0.24$). These values are obtained from the maximum and minimum frequencies from the DMD analysis.

The reason this frequency range shifts is believed to be due to the different forcing amplitude between the multiple triad \mathcal{M}_{2D} model and the experiments. Experimentally, the amplitude required for instability is approximately $\langle |\tilde{\Psi}_0| \rangle_r \approx 35 \text{ mm}^2 \text{ s}^{-1}$. For the multiple triad \mathcal{M}_{2D} model shown here, $|\tilde{\Psi}_{\text{in}}| = 28 \text{ mm}^2 \text{ s}^{-1}$. Figure 6.5(b) shows how the frequency ratio ω_1/ω_0 required for instability in the \mathcal{M}_{2D} model also changes with the forcing amplitude. If the \mathcal{M}_{2D} model was forced at a higher amplitude (similar to that of the experiments) the values of ω_1/ω_0 that would permit instability, decrease to include the experimental range $0.35 < \omega_1/\omega_0 < 0.39$. It is suspected, therefore, that this difference in frequency range for the secondary waves is due to the input amplitude into the model and, based on Figure 6.5(b), a more similar range would be obtained if the multiple triad \mathcal{M}_{2D} model was forced at an amplitude comparable to the experiments. Indeed, this is demonstrated in the simple triad runs forced at $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$, where the ω_1/ω_0 ratio responsible for the amplitude modulations fall within the experimental range.

6.3.4 Growth rate

This section examines the growth rate of the instability for the multiple-frequency \mathcal{M}_{2D} run. As mentioned in § 6.3.3, the flow was allowed to develop beyond linear instability until $t = 800 \text{ s}$ ($t/T_0 \approx 121$), at which point all the domains corresponding to every secondary wave in the flow were re-scaled at every time step using

$$\tilde{\Psi} = 10^{-3} \frac{|\tilde{\Psi}| - |\tilde{\Psi}|_{\min}}{|\tilde{\Psi}|_{\max} - |\tilde{\Psi}|_{\min}} (\cos \theta + i \sin \theta), \quad (6.4)$$

where $|\tilde{\Psi}|(\mathbf{x}, t)$ corresponds to the magnitude of each complex valued cell in a given domain, $|\tilde{\Psi}|_{\min}$ and $|\tilde{\Psi}|_{\max}$ are the smallest and largest magnitudes in the same given domain and $\theta(\mathbf{x}, t)$ is the corresponding phase angle of every cell in the domain given by $\theta = \tan^{-1}(\Im(\tilde{\Psi})/\Re(\tilde{\Psi}))$. The re-scaling given in (6.4) was performed $2g$ times, once for every secondary wave domain $\mathbb{W}_{1_1}, \dots, \mathbb{W}_{1_g}$ and $\mathbb{W}_{2_1}, \dots, \mathbb{W}_{2_g}$ and acted to normalise the magnitude of each domain while retaining the phase angle of the reduced streamfunction in every cell. This re-normalisation was started at $t = 800 \text{ s}$ ($t/T_0 \approx 121$)

into the flow, at which point, each secondary wave had developed an independent structure from the background noise. This meant that when the re-normalisation was stopped at $t = 1100$ s ($t/T_0 \approx 167$), the \mathbb{T} with the maximum growth rate could be determined as the waves already had their developed spatial structure.

Figure 6.13 shows four plots identifying the triad with the maximum energy over the first 3300 s ($t/T_0 \approx 500$) of the model run. The energy of each domain is computed by $E = \kappa^2 \langle |\tilde{\Psi}| \rangle_w^2$. The two left-hand plots (a) and (c) are given in terms of wavenumber ratios κ_1/κ_0 and κ_2/κ_0 respectively, while the right plots (b) and (d) are given in terms of frequency ratios ω_1/ω_0 and ω_2/ω_0 . For this run, which included 60 secondary pairs, all of the input wavenumbers/frequencies are shown, marked in blue if they correspond to the outer loci curve, or pink if they correspond to the inner curve. The black dashed line follows the triad in the flow with the maximum energy. The re-normalisation given in (6.4) occurs between the red dashed lines.

The reason for performing the re-scaling in (6.4) is to understand which \mathbb{T} has the maximum linear growth rate in the multiple triad \mathcal{M}_{2D} model. At initial time ($0 < t/T_0 < 10$) and during the normalisation ($121 < t/T_0 < 167$), the \mathbb{T} with the most energy is simply determined by κ^2 , as all of the amplitudes are comparable. After both of these periods, when the flow is being allowed to develop, the \mathbb{T} with the maximum energy jumps to the smaller wavenumber ratios κ_1/κ_0 . Examining the zero-dimensional theoretical growth rates σ_{0D} on Figure 6.9, shows that $\hat{\sigma}_{0D_{\max}}$ corresponds to $\kappa_1/\kappa_0 \approx 1.2$. Interestingly, what Figure 6.13(a) shows, is that for the \mathcal{M}_{2D} model, after the normalisation, the \mathbb{T} with the maximum energy is close to this maximum $\hat{\sigma}_{0D_{\max}}$, although its value of κ_1/κ_0 is slightly higher at $\kappa_1/\kappa_0 \approx 1.4$. By the time the secondary waves in the \mathcal{M}_{2D} model form visible structure in the flow ($t/T_0 > 250$) the triads with the maximum energy settle at even higher ratios of κ_1/κ_0 . This shows that for the linear exponential growth in the \mathcal{M}_{2D} model, the \mathcal{M}_{0D} model is close to predicting the most unstable \mathbb{T} configuration, however at later times the \mathcal{M}_{2D} model selects triads corresponding to lower values of $\hat{\sigma}_{0D}$. The other point highlighted by Figure 6.13 is that the flow never selects triads on the inner pink branch.

Finally, we explore the actual value of the linear growth rate for the \mathcal{M}_{2D} model, which we will denote σ_{2D} . Figure 6.14(a) shows the amplitudes of the separate waves in the triad on a logarithm scale, with time of the re-scaling indicated by the vertical black dashed lines. As with Figure 6.11, the amplitudes of the \mathbb{W}_1 and \mathbb{W}_2 are given as the summations $\sum_{j=1}^g \langle |\tilde{\Psi}_{1_j}| \rangle_w$ and $\sum_{j=1}^g \langle |\tilde{\Psi}_{2_j}| \rangle_w$, respectively. The red dashed line shows the exponential linear fit for the growth rate from the \mathcal{M}_{2D} model, with the extracted growth, σ_{2D} , rate marked by the red cross on Figure 6.14(b). The light blue

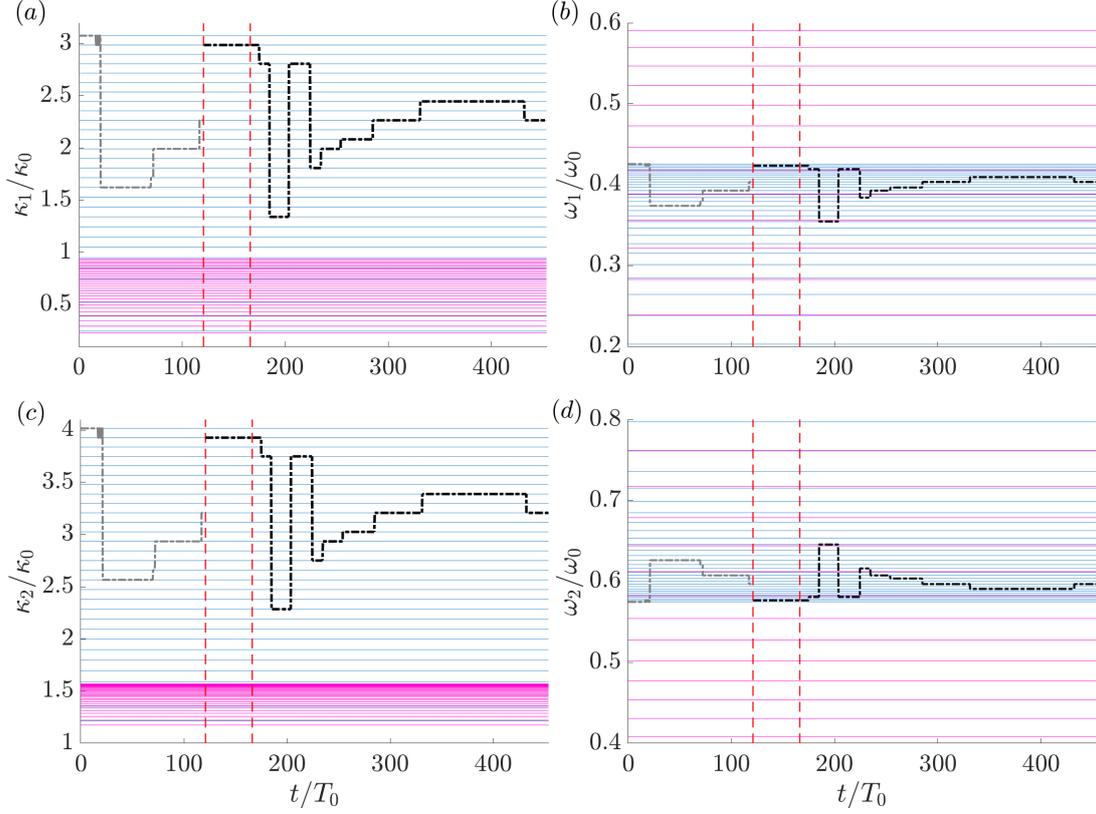


Figure 6.13: Four plots showing the triad with the maximum energy in the system over the first 3300 s ($t/T_0 = 500$) of the model run. The energy of each domain is computed by $E = \kappa^2 \langle |\tilde{\Psi}|_w^2 \rangle$ and the maximum energy is marked with the black dotted line. The two left hand plots (a) and (c), show the secondary waves in wavenumber space, while (b) and (d) are given in terms of frequency. For this run, which included 60 secondary pairs, all of the possible input wavenumbers are shown, marked in blue if they correspond to the outer loci curve, or pink if they correspond to the inner curve. The normalisation given in (6.4) occurs between the red dashed lines. Before the normalisation, the \mathbb{T} with the maximum energy is marked in grey as, for this period, the waves are developing from random background noise.

line (lying almost on top of the red line) is the exponential linear fit for the summation of $\hat{\sigma}_{0D}$ for the 30 input \mathbb{S} (on the outer blue branch) into the \mathcal{M}_{2D} model, calculated by

$$|\tilde{\Psi}|_{\sum \hat{\sigma}_{0D}} = |\tilde{\Psi}_{\text{initial}}| \sum_{j=1}^{(g/2)} e^{\hat{\sigma}_{0D} T}, \quad (6.5)$$

where $T = t_0 - t_1$ is the time window between the two black dash-dotted lines on Figure 6.14(a) that mark t_0 and t_1 and $|\tilde{\Psi}_{\text{initial}}|$ is the amplitude of $\sum_{j=1}^g \langle |\tilde{\Psi}_{1_j}| \rangle_w$ at t_0 . This expression in (6.5) gives the expected value of $|\tilde{\Psi}|$ at time t_1 from the summation of the 30 $\hat{\sigma}_{0D}$ shown by the blue circles on Figure 6.14(b). Only the $\hat{\sigma}_{0D}$ were chosen, as opposed to $\hat{\sigma}_{0D}$ and $\hat{\sigma}_{0D}$, as the triad with the maximum energy in the multiple

triad \mathcal{M}_{2D} model never corresponded to the inner pink branch. This value of $|\tilde{\Psi}|_{\sum \hat{\sigma}_{0D}}$ is then used to calculate the light blue line shown in Figure 6.14(a), from which the growth rate, denoted $\hat{\sigma}_{\sum 0D}$ is extracted and marked on Figure 6.14(b) with the light blue cross.

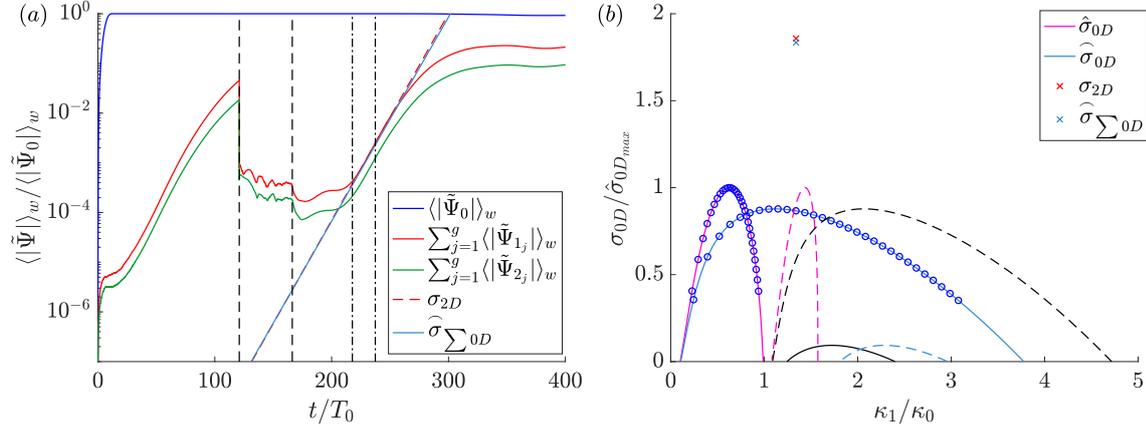


Figure 6.14: (a) A log-linear plot of the magnitude of the reduced streamfunction for the multiple triad \mathcal{M}_{2D} model run outlined in § 6.3.1 with an initial forcing amplitude of $|\tilde{\Psi}_{in}| = 28 \text{ mm s}^{-1}$. The amplitude of the secondary waves are given as the sum across the domains $\sum_{j=1}^g \langle |\tilde{\Psi}_{1_j}| \rangle_w$ and $\sum_{j=1}^g \langle |\tilde{\Psi}_{2_j}| \rangle_w$ and all the amplitudes are normalised against the maximum amplitude of the primary wave. The amplitudes of the secondary waves have been re-scaled using (6.4) between $(121 < t/T_0 < 167)$, indicated by the black dashed lines. The somewhat hidden red dashed line shows the linear fit between the black dash-dot lines to obtain the \mathcal{M}_{2D} linear growth rate σ_{2D} . The light blue line (lying almost on top of the red line) is the linear fit calculated from the summation of $\hat{\sigma}_{0D}$ given in (6.5). (b) Theoretical growth rates calculated using (4.6) for every possible triad configuration shown in Figure 6.1 using an input amplitude of $|\tilde{\Psi}_{in}| = 28 \text{ mm}^2 \text{ s}^{-1}$. The blue circles correspond to the top 60% of both $\hat{\sigma}_{0D}$ and $\hat{\sigma}_{0D_{max}}$ and indicate all the \mathbb{W}_1 parameters used for the multiple triad \mathcal{M}_{2D} run. The red cross gives the value of the \mathcal{M}_{2D} linear growth rate σ_{2D} obtained from the red-dashed line in (a), while the light blue cross gives the linear growth rate $\hat{\sigma}_{\sum 0D}$ extracted from the light blue line.

The σ_{2D} linear growth rate obtained is almost twice the value of $\hat{\sigma}_{0D_{max}}$ calculated from the finite-width expression of Bourget et al. (2014) given in (4.6). However, these two growth rates are not comparable, as one ($\hat{\sigma}_{0D_{max}}$) represents a single \mathbb{S} , while the other (σ_{2D}) represents the summation of multiple \mathbb{S} . The new growth rate, obtained from the linear fit of (6.5), was determined to account for this. Remarkably the values of σ_{2D} and $\hat{\sigma}_{\sum 0D}$ are almost exactly the same. This shows that while the growth rates for the individual \mathbb{S} are not representative, the \mathcal{M}_{0D} model does actually capture the expected growth rate for the instability if multiple secondary waves are accounted for in the flow. In their study of TRI, Bourget et al. (2013) attribute the triad lying on

the outer loci branch to the fact that the $\hat{\sigma}_{0D}$ curve is significantly broader than the $\hat{\sigma}_{0D}$ curve, and therefore shows less wavenumber sensitivity. If this calculation was done in terms of energy, as opposed to amplitude, the broadness of the curve would indeed play a role in the determination of the summation growth rate $\hat{\sigma}_{\Sigma_{0D}}$. The results in Figure 6.14(b) are also in agreement with one of the experiments shown in Figure 4.15(d) where the growth rate obtained was considerably larger than $\hat{\sigma}_{0D_{\max}}$.

6.4 Higher order expansion

For all the of results presented in this chapter so far, we have considered only the non-linear interactions obtained at $\mathcal{O}(\epsilon^2\gamma^0)$ from the perturbation expansion outlined in § 5.2. As can be seen in Figure 6.3 over a certain range of input parameters for the secondary waves, there is a dramatic drop in W_0 across the triad interaction region. This suggests that the non-linear interactions at $\mathcal{O}(\epsilon^2\gamma^0)$ may not be fully capturing the dynamics of the instability, indicating that higher order terms may be important. In this section we present results including the higher order non-linear interactions obtained at $\mathcal{O}(\epsilon^2\gamma^1)$. At this order, we recovered the PDEs shown in (5.26). The way these additional non-linear interactions were integrated into the model was discussed in § 5.5. In order not to complicate the dynamics, we focus only on results from four different simple triads corresponding to the outer loci branch, the parameters for which can be seen in Table 6.3. It should be noted that these are different S_s from those in Table 6.1. Again, the primary wave parameters are given in (6.1) and the input amplitude was $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$.

Wave pair	ω_1 (rad s ⁻¹)	ω_2 (rad s ⁻¹)	l_1 (mm ⁻¹)	l_2 (mm ⁻¹)	m_1 (mm ⁻¹)	m_2 (mm ⁻¹)	κ_1/κ_0
S_a	0.286	0.664	0.013	-0.065	0.070	-0.137	0.849
S_b	0.297	0.651	0.016	-0.067	0.082	-0.160	0.992
S_c	0.309	0.641	0.018	-0.070	0.086	-0.152	1.038
S_d	0.319	0.531	0.020	-0.072	0.093	-0.159	1.130

Table 6.3: Table giving the input parameters of W_1 and W_2 for each S used in the simple triad \mathcal{M}_{2D} model, this time including the non-linear interaction terms at $\mathcal{O}(\epsilon^2\gamma^1)$.

Figure 6.15 shows the amplitude of four different \mathcal{M}_{2D} model runs using the input parameters for the secondary waves given in Table 6.3. All of these runs include the additional non-linear interactions given in (5.26). As with the results for the simple triad presented in Figure 6.3, the instability was not triggered for a range of κ_1/κ_0 , however, the ratio of κ_1/κ_0 that permitted instability, reduced. Specifically, for an

input amplitude of $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$ for the simple triad $\mathcal{M}_{2\text{D}}$ model at $\mathcal{O}(\epsilon^2\gamma^0)$, no instability was triggered for $\kappa_1/\kappa_0 < 1.4$. For the same amplitude, when the higher order terms were included, this limit reduced to $\kappa_1/\kappa_0 < 1$. Although there is a small growth of the secondary waves shown in Figure 6.15(b) where $\kappa_1/\kappa_0 = 0.992$, the growth is small and the wavenumbers are very close to this $\kappa_1/\kappa_0 < 1$ limit.

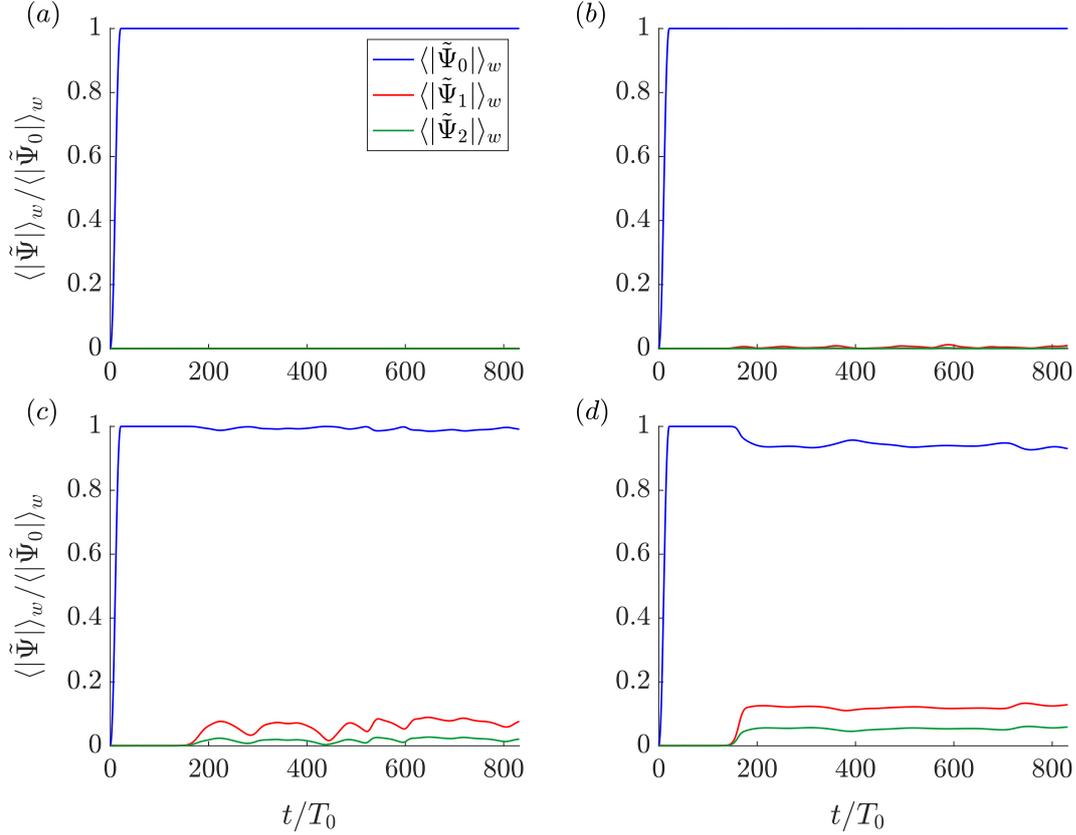


Figure 6.15: Four amplitude plots from the $\mathcal{M}_{2\text{D}}$ model including the non-linear interactions at $\mathcal{O}(\epsilon^2\gamma^1)$. Each plot corresponds to a simple triad with three domains ($n_d = 3$). The parameters for the secondary waves corresponding to each plot are given in Table 6.3. (a) \mathbb{S}_a (b) \mathbb{S}_b (c) \mathbb{S}_c (d) \mathbb{S}_d . The forcing amplitude for \mathbb{W}_0 in all plots is $|\tilde{\Psi}_{\text{in}}| = 35 \text{ mm}^2 \text{ s}^{-1}$ and the run time is 90 minutes ($t/T_0 \approx 800$). The amplitudes are calculated over the whole visualisation window $\langle |\tilde{\Psi}| \rangle_w$ and have been normalised against the maximum amplitude of $\langle |\tilde{\Psi}_0| \rangle_w$.

Considering the pairs \mathbb{S}_c and \mathbb{S}_d , shown respectively in Figure 6.15 (c) and (d), it is clear that the primary wave becomes unstable although the amplitude of the secondary waves remains small. In fact, for all of the input pairs tested, $\langle |\tilde{\Psi}| \rangle_w / \langle |\tilde{\Psi}_0| \rangle_w < 0.18$. Experimentally, and for the multiple triad $\mathcal{M}_{2\text{D}}$ model using only $\mathcal{O}(\epsilon^2\gamma^0)$, this ratio was almost approximately double, sitting around 0.3. In addition, there is only a very

small parameter range over which the amplitude modulations are witnessed. For triads with larger ratios of κ_1/κ_0 than \mathbb{S}_d in Table 6.3, the oscillations disappear and the triad reaches a steady equilibrium. However, despite the amplitude of the secondary waves remaining small, the large decay of the \mathbb{W}_0 beam, witnessed in the $\mathcal{O}(\epsilon^2\gamma^0)$ simple triads in Figure 6.3, is not seen.

Although the amplitude plots in Figure 6.15 do not reveal new insights into the amplitude behaviour of TRI, as the higher order terms included in these model runs contribute to the phase angle of the reduced streamfunction, it is necessary to explore its evolution. Figure 6.16 presents snapshots of the run for \mathbb{S}_c shown in Figure 6.15(c) at $t/T_0 = 795$. Unlike the previous snapshots presented, here the reduced streamfunction is not multiplied by the fast time scales $e^{i\phi}$, rather the magnitude $|\tilde{\Psi}|$ and phase angle θ of $\tilde{\Psi}$ are shown on the left and right panels respectively for each wave in the triad (where $\theta = \tan^{-1} \Im(\tilde{\Psi})/\Re(\tilde{\Psi})$).

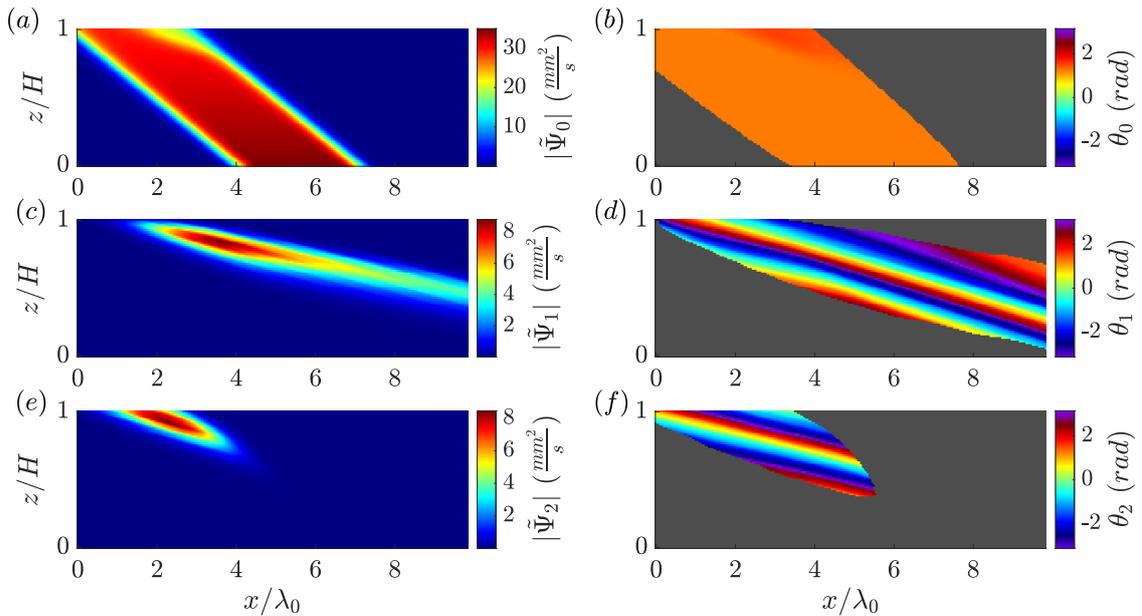


Figure 6.16: Snapshots at $t/T_0 = 795$ of the simple triad model run shown in Figure 6.15(c) using the input parameters \mathbb{S}_c . The left hand panels show the magnitude of the reduced streamfunction $|\tilde{\Psi}|$, while the corresponding phase angles θ are given in the right panels. The fast time scales $e^{i\phi}$ are not included. Each wave in the triad is shown separately.

When only considering the non-linear interactions at $\mathcal{O}(\epsilon^2\gamma^0)$, the phase angle of the reduced streamfunction for all waves varies slowly, similar to the phase angle of the primary wave shown in Figure 6.16(b). As can be seen in Figure 6.16(d) and (f), however, the inclusion of these higher order terms strongly influences the phase of the reduced streamfunction of the secondary waves. The lines of constant phase are seen

to move quickly in the direction of their respective \mathbf{k} , showing that the phase angle is in a state of continuous variability. This is unexpected as, like the amplitude, one would expect the phase angle of the reduced streamfunction to vary slowly in both space and time. However, this continuous variation could be due to the solution trying to move towards the ‘correct’ fast-time phase $\phi = \mathbf{k} \cdot \mathbf{x} - \omega t$. Consider that $\Psi = \tilde{\Psi} e^{i\phi}$, where $\tilde{\Psi} = |\tilde{\Psi}| e^{i\theta}$. This can be expressed as $\Psi = |\tilde{\Psi}| e^{i(\theta+\phi)}$. This variation in θ could be hiding an attempt for the solution to access the preferred ϕ .

There is also a question here as to whether the interactions at this order should be included with the same weighting as the non-linear interactions at $\mathcal{O}(\epsilon^2\gamma^0)$. Moreover, other orders in the expansion also sit at a similar magnitude to $\mathcal{O}(\epsilon^2\gamma^1)$, such as $\mathcal{O}(\epsilon^1\gamma^2)$, the effects of which have not been included in the \mathcal{M}_{2D} model. Overall, further work would need to be conducted to understand how these further non-linear interactions should be built into the model and the role they play in the development of the phase angle. As the results at $\mathcal{O}(\epsilon^2\gamma^0)$ already provided significant insights into the experimental results, this further investigation is considered outside the scope of this thesis.

6.5 Summary

This chapter has provided results based on the weakly non-linear \mathcal{M}_{2D} model outlined in Chapter 5. The model was built with the aim of better understanding the results of Chapter 4, which examined the development of TRI in a finite-width beam. First, the results of the model were investigated looking at a ‘simple triad’, where only three waves were input into the system. This analysis was further split into the wave parameters that correspond to the inner or outer branch of the loci curve. When examining the outer branch – which is also the branch selected experimentally – a study was conducted examining the role of the input parameters for the secondary waves. We saw that over a small range of parameters the amplitude variability witnessed experimentally was captured well, however, as the frequency of the W_1 beam increased, the interaction time of the triad also increased due to W_1 and W_0 becoming closer to being parallel. This resulted in significant decay of the primary beam and the triad reaching a steady state. This decay was similar to the results from the \mathcal{M}_{0D} model presented in § 4.5.

Results from the model were then presented whereby the parameter space for the secondary waves was extended to permit multiple triadic perturbations into the flow. This was in light of the modulations in Fourier space witnessed experimentally, which

would not be possible to capture with the simple triad system. A run was presented that included 60 different pairs of secondary waves all able to interact via their influence on the same primary beam. This run saw the flow develop in a manner more similar to the experiments, with \mathbb{W}_1 showing large variability in both its location and width. The frequency modulation of the resonant waves was also explored. The spectrogram presented in Figure 6.12 qualitatively resembled the experimental spectrograms, with periods of stronger and weaker frequency signal for the secondary waves. A discussion was also provided on why the selected frequencies of the resonant waves did not fall exactly in the experimental range. An analysis of the σ_{2D} also revealed that, when the $\widehat{\sigma}_{0D}$ were summed over the inputs into the multiple triad \mathcal{M}_{2D} model, they predicted the growth rate well. This σ_{2D} also agreed with the experimental results.

The last section § 6.4 examined the effect of including higher order non-linear terms from the perturbation expansion into the model. These terms were seen to strongly affect the phase angle of the reduced streamfunction. Unlike the non-linear terms at $\mathcal{O}(\epsilon^2\gamma^0)$, here the phase angle exhibited a continuous variation in t in the direction of \mathbf{k} . This behaviour did not seem physical as one would expect the phase angle of the reduced streamfunction to vary more gradually over its spatial extent. It was concluded that the addition of these high order terms would need further investigation along with the consideration of other terms that lie at a similar order of magnitude such as $\mathcal{O}(\epsilon^1\gamma^2)$.

Chapter 7

Deliberate perturbations

7.1 Introduction

This chapter moves away from the exploration into the structure of TRI in a finite-width beam and instead focuses on how perturbations to the internal wave field can potentially act as non-linear triggers for the instability. As discussed in § 2.3, there is a distinction between triadic resonance, which encompasses all triadic interactions, and TRI, which refers specifically to the breakdown of a single-wave into a resonant triad. Unlike TRI, where an amplitude threshold for linear instability is present in a finite-width beam, triadic resonance – generated from Two Wave Interactions (TWI) or other mechanisms – can occur at any amplitude, even in a finite-width beam (Varma and Mathur, 2017). In this chapter we explore whether deliberate perturbations to the flow, which can initially give rise to triadic resonance, can provide a non-linear pathway to TRI after their removal.

The motivation here comes from experiments that were conducted to explore the interaction between an imposed wave field and a freely evolving turbulent-like structure. Specifically, a vortex ring was taken as the simplified prototype for a turbulent structure and an exploration was begun into how its propagation is influenced by the internal wave field, and how the vortex ring in turn affects the wave field. Significantly, we demonstrate that a vortex ring propagating through a pre-existing wave field can provide a non-linear trigger for triadic resonance when the amplitude of the imposed wave field is itself too small for the spontaneous growth of TRI.

The initial results also suggested that the triadic resonance generated by the vortex ring perturbation was self-sustaining, indicating that the passage of the vortex ring acted to trigger TRI, despite the amplitude of the beam being lower than critical threshold. This then led to the hypothesis that the occurrence of TRI may follow a non-

linear pathway, described by a sub-critical bifurcation. Experiments were consequently conducted to test this hypothesis without including the complex dynamics of the vortex ring.

This chapter is divided into the following sections. First, in § 7.2, we present the concept of a vortex ring and how its structure evolves in a quiescent stratified environment. Next, in § 7.3, the vortex ring interaction experiments are examined. Two elements are explored here, one looking at how the wave field impacted the propagation of the vortex ring in § 7.3.2, and the second examining how the vortex ring affected the wave field in § 7.3.3. These experiments then provide the starting point for § 7.4, which examines the hypothesis of a sub-critical bifurcation to the TRI mechanism. This was done using two sets of experiments that utilised the unique capabilities of ASWaM. Finally, in light of the results from § 7.4, the longer term evolution of the vortex ring interaction experiments are explored over extended time periods in § 7.5.

7.2 Vortex rings in a stratified environment

7.2.1 Overview

Before considering the effect of the vortex ring on the wave field, we first briefly examine the evolution of a ring in a quiescent environment. We use the term ‘vortex ring’ to describe an initially axisymmetric, toroidal fluid structure with azimuthal vorticity. The apparatus used for their generation has been outlined in § 3.6. First in § 7.2.2 we give details on how these structures form, propagate and eventually breakdown in a stable linear stratification. The internal waves generated by their motion are briefly discussed in § 7.2.3 and finally in § 7.2.4 their collapse distance due to the stratification is discussed.

7.2.2 Formation and propagation

For the experiments considered in this thesis, the vortex rings were produced by partially submerging a cylindrical tube of constant cross sectional area in a linear stratification. The diameter of the tube is 0.04 m. Using a piston-like mechanism, an impulse is delivered to the tube, causing some of the fluid to be ejected from the outlet. As the fluid is forced out of the opening, a cylindrical vortex sheet forms due to the viscous boundary layer on the inside of the tube walls. Upon exiting the tube, this thin boundary layer separates and rolls up into a vortex ring (e.g. Maxworthy, 1972; Didden, 1979). The ring then begins to propagate horizontally through the fluid in the direction

of its symmetric axis, continuously entraining fluid from its surrounds and shedding vorticity in its wake. In a homogenous environment, this entrainment would cause the ring to decelerate and the diameter to grow uniformly in all directions (Jackson, 2021). In a stratification, its motion is somewhat different, as opposing density fluids are entrained on either side of the ring, which weakens the core circulation (Van Atta and Hopfinger, 1988) and causes the ring to flatten.

Initially, if the ring is moving quickly, the density gradient has little impact on its motion. At later times, however, this density gradient causes the ring to collapse from a fully three-dimensional structure to a quasi-two-dimensional one, resembling a horizontally propagating dipole (Scase and Dalziel, 2006b). This continues to propagate, albeit at a slower speed.

The relative effects of velocity compared to the strength of the stratification are characterised by the non-dimensional Froude number

$$Fr = \frac{U}{Na}, \quad (7.1)$$

where N is the buoyancy frequency, U is the initial velocity of the ring and a is taken as half the distance between the cores of the ring as it exits the outlet tube, as shown in Figure 7.1(a).

The photos in Figure 7.1 show the plan view of the generation of a vortex ring in a linearly stratified environment of $N = 1.85 \pm 0.06 \text{ rad s}^{-1}$, visualised using fluorescence dye and a laser sheet. The ring was moving with an initial velocity of approximately $U = 60 \text{ mm s}^{-1}$, with half the distance between the two cores a being 23 mm, giving a Fr of approximately 1.4. The coherent structure of the ring begins to be lost around Figure 7.1(f), where the ring and the wake begin to become turbulent. As the vorticity is shed in the wake of the ring, energy is transferred to smaller scales where it eventually dissipates with some of the energy contributing to mixing of the background stratification.

7.2.3 Internal wave production

Another key difference between the propagation of vortex rings in a stratification compared to a homogenous environment is that their passage generates internal waves. Unlike the ‘permanent’ wave fields generated from ASWaM, with steady frequencies and wavenumbers, vortex rings generate ‘transient’ waves, comprised of a whole spectrum of frequencies and wavenumbers due to their impulsive start and rapidly decaying structure. A detailed study of the internal waves produced from vortex rings is provided

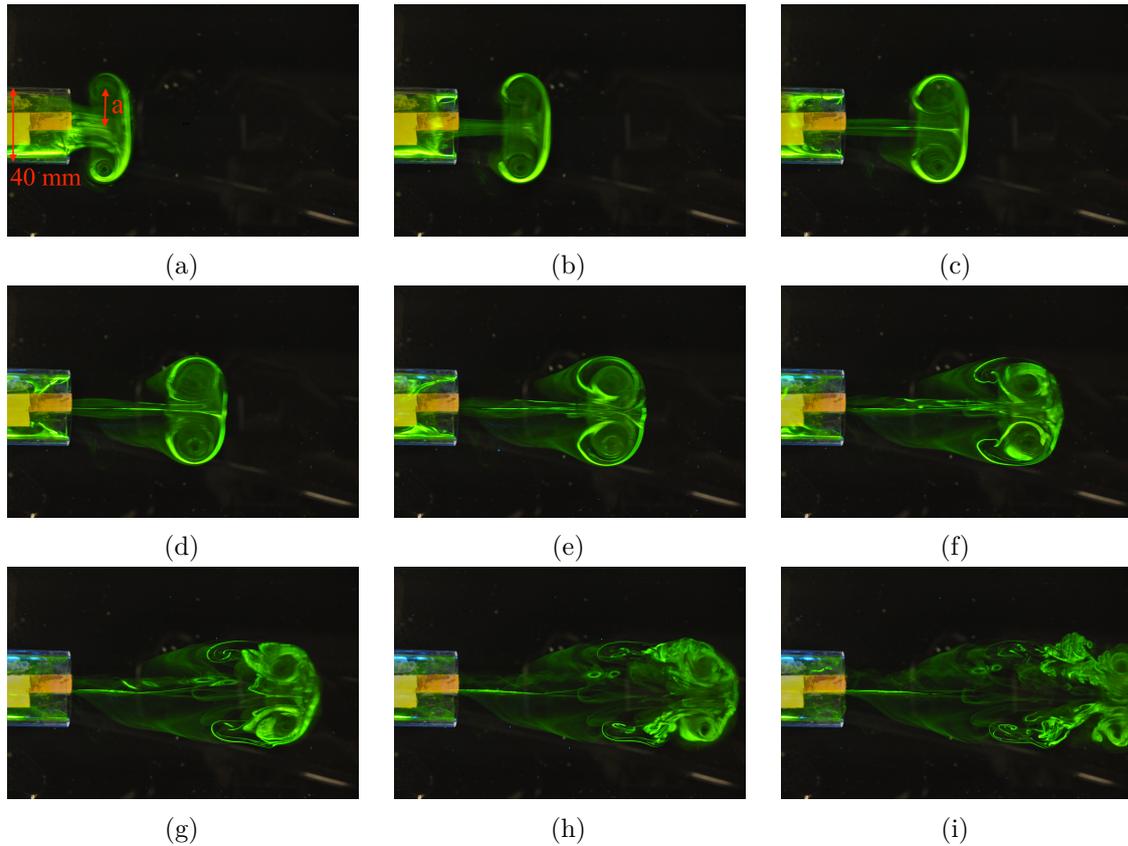


Figure 7.1: Sequence of images taken 0.25 s apart showing the plan view of the generation of a vortex ring in a linear stratification with $N = 1.85 \pm 0.06 \text{ rad s}^{-1}$ and $Fr \approx 1.4$. The images are visualised using fluorescence dye and a laser sheet. The outlet diameter of the tube is 40 mm. (a), (b) and (c) show the initial formation of the ring as it exits the tube with the inner vortex core visible. (d), (e) and (f) show the three-dimensional structure begin to collapse as the stratification flattens the ring. (g), (h) and (i) then show a transition to turbulence as vorticity is shed in the wake of the dipole structure.

by Scase and Dalziel (2006a,b), who show that the emission of these waves also further contributes to the deceleration of the ring.

The snapshots in Figure 7.2 shows a vortex ring with same parameters of that in Figure 7.1, although this time the view from the Synthetic Schlieren is given. The outline of the vortex ring and its turbulent wake (which also produces waves) is marked in black. The structure of the ring itself is not visible however the spectrum of internal waves produced is clearly shown. The outline was produced by tracking the blurred regions of the dot pattern caused by the three-dimensional structure of the ring.

Initially, as the ring begins to propagate in Figure 7.2(b), the flow is dominated by phase-locked lee waves produced from the ring and its turbulent wake. Work by Hopfinger et al. (1991) showed that for rings with $Fr < 4$, lee waves dominate the

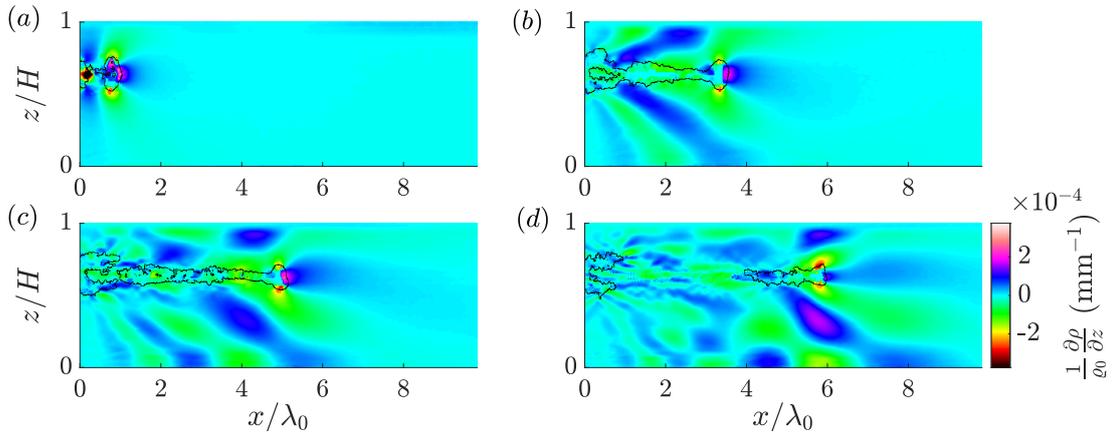


Figure 7.2: Sequence of images of a vortex ring propagating in a linear stratification of $N = 1.85 \pm 0.06 \text{ rad s}^{-1}$. The vertical gradient of the density perturbation is shown. The outline of the vortex ring is shown in black. (a) $t = 12 \text{ s}$, (b) $t = 30 \text{ s}$, (c) $t = 45 \text{ s}$, (d) $t = 60 \text{ s}$.

wave regime, while above this limit shorter wavelength internal waves generated by the turbulent wake dominate. As the ring decelerates its Fr reduces. When $Fr \sim 1$, as is the case around Figure 7.2(d), the system no longer has sufficient kinetic energy to displace the stratified fluid vertically around the diameter of the ring. In the case of lee waves around an isolated obstacle this would lead to a blocking with only part of the flow able to go over the obstacle while the rest goes around (Dalziel et al., 2011). For a vortex ring, the structure itself collapses and the internal waves begin propagate out in front, as shown in Figure 7.2(d). Scase and Dalziel (2006a) showed that the energy imported to the waves increased as Fr decreased towards unity.

7.2.4 Collapse distance in a stratified environment

Two parameters were used to control the properties of the vortex rings used in this thesis. First, the dimensionless Formation number $F_p = L_p/D_p$ defines the length of the piston stroke L_p , over the diameter of outlet tube D_p , which for the experiments shown here is 40 mm. The higher the Formation number, the greater the circulation in the ring. Assuming all vorticity is concentrated in the ring, for small F_p , the vortex ring is comprised of a thin toroidal core described by Lamb's thin core model (Lamb, 1932). With increasing F_p , the thickness of the core increases, eventually reaching the upper limit of Hill's spherical vortex (Hill, 1894). Experimental work by Gharib et al. (1998) found there exists a limiting Formation number (for all combinations of piston velocities and exit diameters) $F_{p_{lim}}$ between 3.6 – 4.5, above which additional circulation manifests as a trailing jet behind the ring.

The other parameter used to control the vortex ring was the speed that the pistons were forced. Increasing the velocity of the pistons creates rings with higher initial Fr number, allowing them to propagate further before being overcome by the stratification. To examine how the piston rate effects the rings development, vortex rings were fired into a linear stratification of $N = 1.43 \pm 0.06 \text{ rad s}^{-1}$. The experiments were captured at one frame per second. The F_p of the rings was kept constant at 4.4, while the rate of the piston ranged between 650 and 750 mm s^{-1} , generating $1 < Fr < 2.4$. As the ring passed through the fluid, its three-dimensional structure distorted the background dot pattern used for the Synthetic Schlieren. A time series was then extracted by considering the evolution of the vertical mean of the defocused pattern. From this the initial velocity was calculated using a least squares linear fit to the initial part of the distance time curve. The collapse location of the nose of the vortex ring d was found from identifying the edge of the curve (which represents the tip of the defocused structure and hence the nose of the collapsed ring). Figure 7.3 shows the non-dimensional collapse distance of the vortex rings (d/D_p) from the outlet tube as a function of their initial Fr .

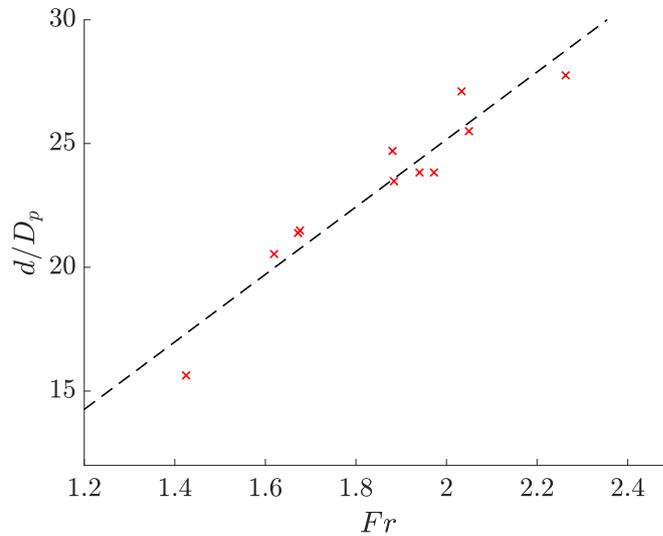


Figure 7.3: The non-dimensional collapse distance of the vortex rings in a linear stratification of $N = 1.43 \pm 0.06 \text{ rad s}^{-1}$ as a function of their initial Fr . Experimental data points are marked in red with a linear line of best fit shown in black. All rings used a $F_p = 4.4$, while rate of the piston ranged between 650 and 750 mm s^{-1} .

Figure 7.3 shows an approximately linear relationship between the collapse location of the ring and its initial velocity. As the experiments were captured at one frame per second, there is some uncertainty on the calculated initial velocity. Despite this, for

the range of Fr considered, these results are in agreement with Johari and Fang (1997) who examined the collapse distance of rings with varying Re and Fr . For rings with $Re = 1350$ and $Fr < 3.2$, they found that the collapse distance increases linearly with increasing velocity. As the rings shown here have a $Re \approx 1300$, this linear relationship is in good agreement. Their actual collapse distances were lower than the ones shown here, however, this is because they considered F_p in the range $1 \sim 1.5$, opposed to $F_p = 4.4$ considered here.

7.3 Vortex ring interactions

7.3.1 Overview

The next stage of the analysis examines how the flow evolves when a vortex ring is fired into an internal wave field. Initially, experiments used a run time of $t_{\text{end}} = 1200$ s (20 minutes). Subsequently, they were repeated with the same forcing parameters but with a run time of $t_{\text{end}} = 3600$ s (1 hour). As these longer experiments are of higher quality, the first 20 minutes of these are presented here to reflect the original results from the shorter experiments conducted initially.

These experiments were undertaken in the same linear stratification used in § 7.2.4 of $N = 1.43 \pm 0.06$ rad s⁻¹, whereby a vortex ring, outlined in § 7.2, was fired into the internal wave field described in § 4.2. There are two main elements that need to be considered when considering this interaction. First, outlined in § 7.3.2, is how the wave field impacts the passage of the vortex ring. Second, discussed in greater detail in § 7.3.3, is how the vortex ring influences the wave field.

7.3.2 Effect of the wave field on the vortex ring

To understand how the presence of the wavebeam impacts the vortex rings, we focus on how the collapse distance of the vortex ring changes when it interacts with the wave field. Vortex rings were fired into a pre-existing wave field, in the same stratification as the experiments shown in § 7.2.4. The same \mathbb{W}_0 forcing parameters and wavebeam envelope outlined in § 4.2 was used, although for the experiments presented here, $l_0 = 0.05$ mm⁻¹, meaning the waves were directed to the right in the tank, away from the outlet of the vortex ring tube. This meant the horizontal component of the group and phase velocities of the internal wave beam and the propagation direction of the vortex ring were in the same direction. For details of the set-up of the vortex-ring tube, see Figure 3.15.

The wave field was allowed to develop for 330 s ($t/T_0 = 50$) before the ring was fired, however as the forcing amplitude of the waves was below the limit for instability, TRI had not occurred during this time. Again, the Formation number of the ring was kept constant at $F_p = 4.4$, while the velocity of the piston ranged between 600 and 750 mm s⁻¹. As the flow was quiescent to the left of the wave field, the initial development of the ring was the same as described in § 7.2, until it reached the left flank of the primary wavebeam after approximately 440 mm ($x/D_p = 11$). Using the same calculation method described in § 7.2.4, the initial velocity of the ring and collapse distance were calculated and are shown in blue on Figure 7.4 as a function of Fr and non-dimensional collapse distance. The collapse distances of a ring in a quiescent stratification, given in Figure 7.3, are also shown in red. The grey patch indicates the width of the primary wave beam.

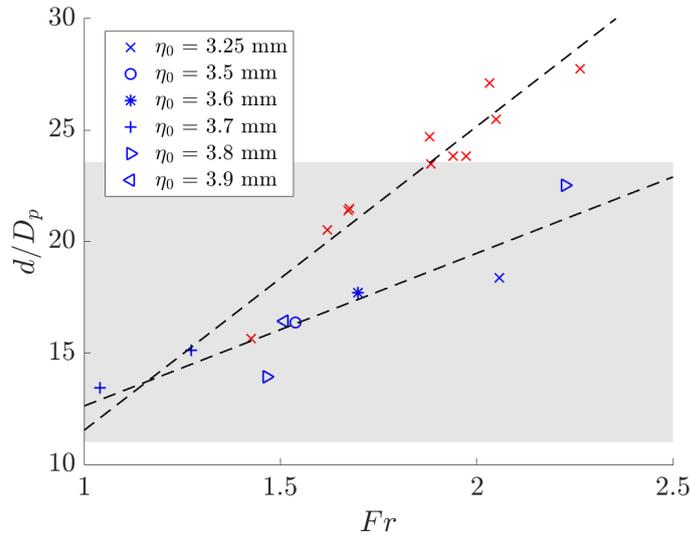


Figure 7.4: The collapse distances of vortex rings in a linear stratification ($N = 1.43 \pm 0.06$ rad s⁻¹) as a function of their initial Fr . The red crosses (also shown on Figure 7.3) are for rings in a quiescent environment while the blue crosses show the collapse distance of rings who have interacted with an internal wave field. A linear fit is shown in black for both data sets. Again, all rings used a $F_p = 4.4$, while the rate of the piston ranged between 600 and 750 mm s⁻¹ for the blue data set. The grey patch indicates the width of the primary beam. The left flank of the beam is approximately 440 mm away from the outlet tube ($x/D_p = 11$), while the right flank is approximately 940 mm away ($x/D_p = 23.5$). The marker style of each blue data point gives the forcing amplitude η_0 of the wave beam.

As with the collapse location of the rings in a quiescent environment, the collapse distance of the rings interacting with the wave beam also follows a linear relationship, although this time with a shallower gradient. This indicates that as the ring interacts

with the wave beam its motion is inhibited, resulting in its premature collapse. Again, these experiments were captured at only one frame per second so there is some uncertainty associated with the initial velocity of the ring, however the results of Figure 7.4 show a clear relationship between the presence of the wavebeam and the premature collapse of the ring. Despite this, the amplitude of the wave field does not seem to have a significant impact on the collapse distance. One would expect that a larger amplitude wave field would cause an earlier collapse of the ring. While this is shown for one of the rings interacting with a wave beam at $\eta_0 = 3.8$ mm, the other point that sits below the line of best fit is for a ring interacting with the lowest amplitude wave. Saying this, the amplitude range is not large and further experiments looking at a larger range of amplitudes using a higher frame rate would provide more insight into this relationship and the mechanisms that promote the collapse.

7.3.3 Effect of the vortex ring on the wave field

The second element to this investigation is how the non-linear perturbation impacted the internal wavebeam. Here we consider the same experiments presented above in § 7.3.2, but this time from the perspective of the wave field. The snapshots in Figure 7.5 show one of the experiments with a vortex ring interacting with the rightward propagating primary wave beam. The forcing amplitude of the wave beam was $\eta_0 = 3.6$ mm; below the amplitude for instability, which lay around 3.8 mm (this threshold changes slightly during the course of the week of an experimental series due to local changes in the stratification, hence the reason it is important to use the measured values for $|\tilde{\Psi}_0|$ rather than solely the forcing amplitude η_0). The first two snapshots, Figure 7.5(a) and (b), clearly show the ring being ejected from the outlet tube and the internal waves produced by its motion. The ring moves with an initial velocity of 54.6 mm s^{-1} ($Fr = 1.54$) and collapses in the middle of the primary wave beam, 655 mm from the outlet ($d/D_p = 16.4$). In Figure 7.5(c), artefacts of the rings passage are still evident, but 70 s later ($t/T_0 = 10.6$), in (d), it is difficult to discern that the wave field was ever perturbed. Despite this, in Figure 7.5(e) and (f), the presence of \mathbb{W}_1 becomes distinguishable, indicating triadic resonance, which does not appear to decay.

What is so intriguing about this experiment, is how, after the collapse of the vortex ring, triadic resonance formed and selected the \mathbb{W}_1 and \mathbb{W}_2 associated with TRI at a slightly stronger forcing, despite the forcing amplitude being below the critical value. Repeating this experiment multiple times (and in different strength stratifications) revealed the same result, that the vortex ring was able to trigger triadic resonance at a range of amplitudes below the stability threshold. The plots in Figure 7.6 show

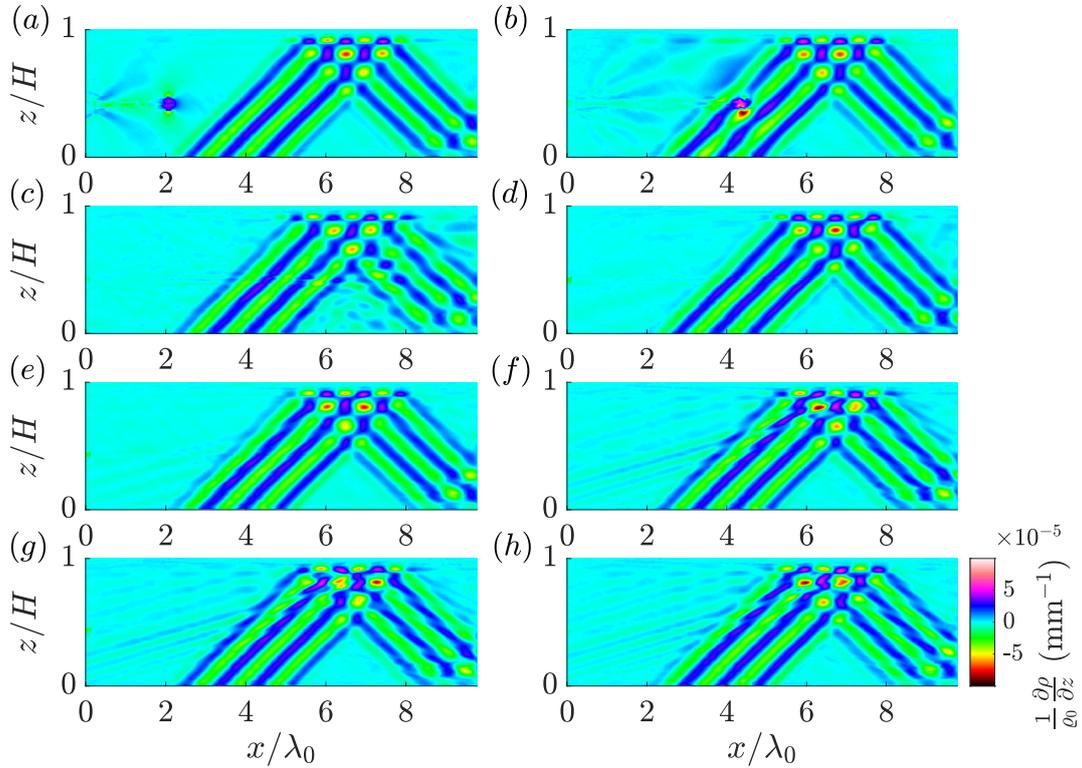


Figure 7.5: Sequence of images from a vortex ring interaction experiment showing the vertical gradient of the density perturbation. The forcing amplitude of the primary beam is $\eta_0 = 3.6$ mm, while the vortex ring has a $F_p = 4.4$ and a piston rate of 670 mm s $^{-1}$. These forcing parameters correspond to an initial velocity of 54.6 mm s $^{-1}$ ($Fr = 1.54$) and a collapse distance of 655 mm from the outlet ($d/D_p = 16.4$). The quantitative amplitudes of the three triadic-waves are shown in Figure 7.6(b). (a) $t/T_0 = 50.8$, (b) $t/T_0 = 52.2$, (c) $t/T_0 = 57.5$, (d) $t/T_0 = 68.0$, (e) $t/T_0 = 98.3$, (f) $t/T_0 = 134.5$, (g) $t/T_0 = 143.6$, (h) $t/T_0 = 151.2$.

this quantitatively. The forcing amplitude η_0 for (a) is 3.5 mm while (b) and (c) are both 3.6 mm. The snapshots shown in Figure 7.5 correspond to the results in Figure 7.6(b). For all three of these experiments, a vortex ring with an F_p of 4.4 was fired at $t = 330$ s ($t/T_0 = 50$), indicated by the black dashed line. The presence of the ring is captured by the jump in amplitude of the primary wave as the two interact. Shortly after, all three of these experiments show a growth in amplitude of the secondary waves in the triad. Interestingly, unlike the growth witnessed in Chapter 4, it is \mathbb{W}_2 (plotted in green) that grows first, before the growth of \mathbb{W}_1 is evident. As \mathbb{W}_1 is the only secondary wave visible qualitatively, this explains why the time lag between the ring and the onset of the interaction appears greater than it is.

As the presence of the secondary waves then persists for 20 minutes (the length of the initial set of experiments), these experiments appear to show that the vortex

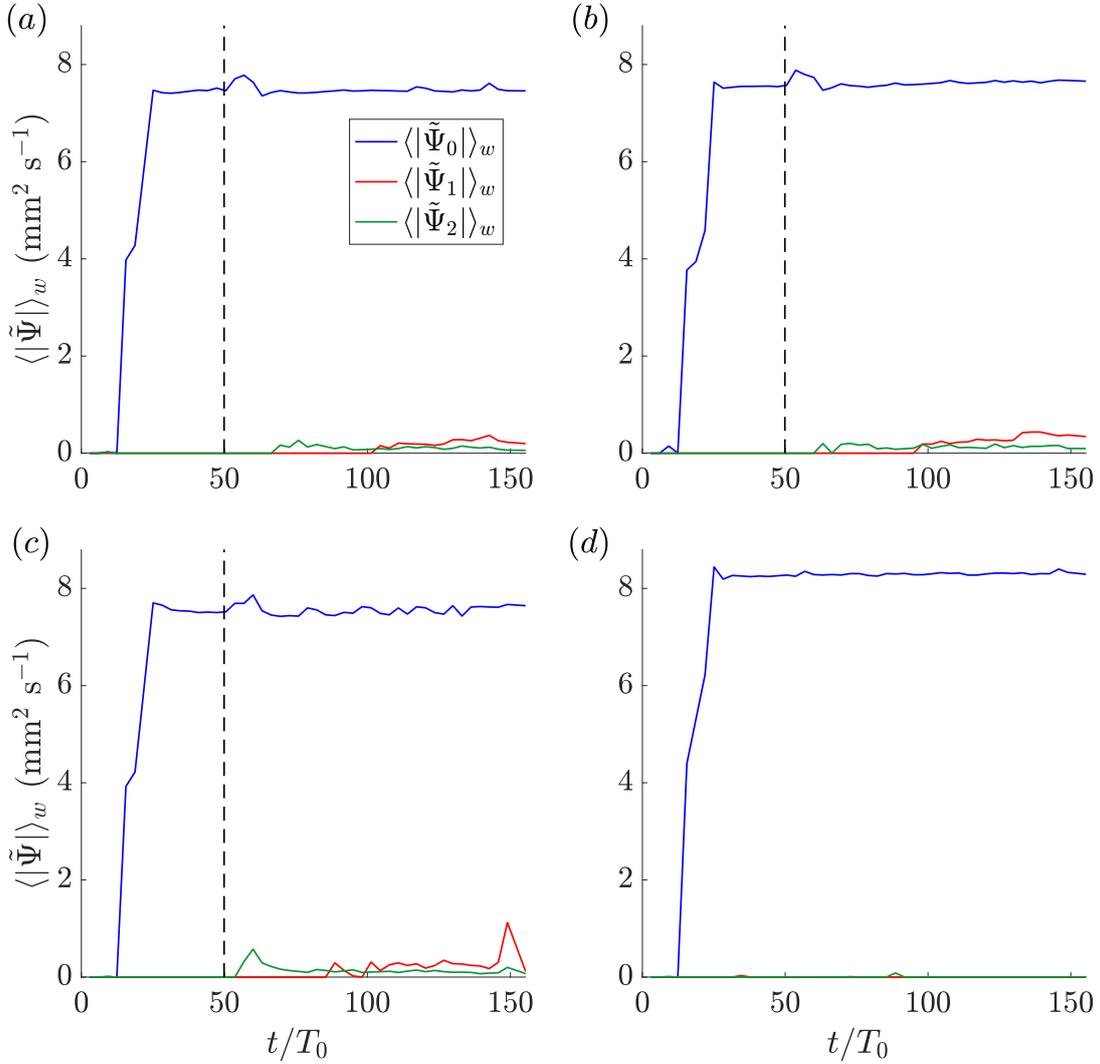


Figure 7.6: The amplitude of the streamfunction for $\langle |\tilde{\Psi}_0| \rangle_w$ (blue), $\langle |\tilde{\Psi}_1| \rangle_w$ (red), $\langle |\tilde{\Psi}_2| \rangle_w$ (green) for four different experiments. All amplitudes were integrated over the whole visualisation window \mathcal{D}_W (see Figure 3.23 for details). As it is necessary to compare the amplitudes between experiments, they have not been normalised. In three of the experiments (a), (b) and (c) a vortex ring of $F_p = 4.4$ was fired into the wave field at $t = 330$ s ($t/T_0 = 50$), indicated by the black dashed line. The forcing amplitudes of the primary wave beam are (a) $\eta_0 = 3.5$ mm, (b) $\eta_0 = 3.6$ mm, (c) $\eta_0 = 3.6$ mm and (d) $\eta_0 = 3.7$ mm.

ring does not just generate a transient triadic interaction, but acts as a non-linear trigger for TRI. Finally, the plot in Figure 7.6(d) shows another experiment from the set, forced at $\eta_0 = 3.7$ mm but without the vortex ring. Even though the amplitude of the primary wave is clearly comparatively larger than the other three experiments, here no TRI is observed.

7.4 Potential sub-critical bifurcation

7.4.1 Overview

Based on the findings from the vortex ring interaction experiments in § 7.3, the internal waves pathway to instability was called into question. The collapse of the vortex-ring inside the beam appeared to provide a non-linear trigger for TRI. This led to the hypothesis that the TRI mechanism may be subject to a sub-critical bifurcation. Previous chapters have shown that a finite-width internal wave beam can have two linearly-stable equilibrium states: one with a single-wave and one with a triad of waves, moving from the former to the later by increasing the amplitude. This indicates that the TRI mechanism follows a supercritical bifurcation. This is illustrated by the diagram in Figure 7.7(a), which shows schematically the equilibrium states of the non-linear system. The dotted black line indicates an equilibrium state that is linearly-unstable, while the solid lines represent equilibrium states that are linearly-stable. As the forcing amplitude is increased, the system transitions away from the horizontal axis where only the linearly-stable primary wave field is present to the linearly-stable triad at location 2.

The vortex ring experiments revealed that this pathway may not be so simple, and in fact, there may exist a region where the system could correspond to either the linearly-stable single-wave branch or the linearly-stable triadic-wave branch depending on other perturbations or forcing history. This sub-critical bifurcation is illustrated by the diagram in Figure 7.7(b). Again, the dotted lines indicate equilibrium states that are linearly-unstable, while the solid lines represent equilibrium states that are linearly-stable. As the amplitude of the beam is increased from location 1 to 2 on Figure 7.7(b), the system transitions from the linearly-unstable single-wave state to the linearly-stable triad at location 3. If, however, the system is perturbed hard enough while at location 1 (as with the vortex rings), it could transition to the linearly-unstable triadic-wave state. As this branch is unstable, it is then thrown away from this point and either

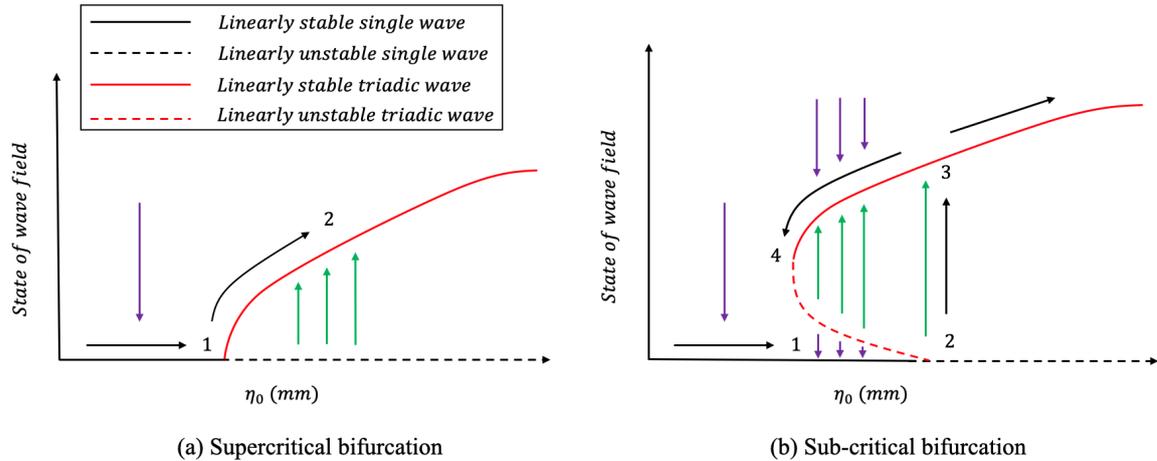


Figure 7.7: Schematic showing (a) supercritical stability states of the TRI system and (b) showing the sub-critical bifurcation stability states of the primary wave beam for TRI. The solid red and black lines indicate the linearly state equilibrium states while the dashed lines represent the linearly-unstable equilibrium states. The schematic in (b) shows a region whereby the waves could correspond to either the linearly-stable single-wave state or linearly-stable triadic-wave state.

jumps back to the linearly-stable single-wave state or moves to the linearly-stable triadic-wave state at 4. These ideas are further discussed in § 8.3.

As the vortex ring appeared to trigger this transition between the stable branches, experiments were designed to explore this pathway to instability without the complex dynamics of the vortex ring. Two different sets of experiments were designed and conducted to test this hypothesis. The first set is presented in § 7.4.2, while the second set is presented in § 7.4.3.

7.4.2 Amplitude reduction experiments

Overview

The first set of experiments are concerned with exploring whether the system exhibits a hysteresis; meaning the stability state of the system depends on its history. Such hysteresis is characteristic of a sub-critical bifurcation. To test this, experiments were designed to quasi-steadily change the amplitude of the primary wave beam over the duration of the run. Initially, the beam is forced at a low enough amplitude so that the primary wave remains at location 1 on the linearly-stable single-wave branch, as shown in Figure 7.7(b). As the forcing amplitude is increased the system becomes linearly-unstable as a single wave and transitions to its new linearly-stable state with a triad of

waves at location 3 (the solid red line). These experiments aimed to access location 4, the region where, due to the previous forcing amplitude, the system remained on the linearly-stable triadic-wave branch. If sufficient energy was given to the system then, by reducing the forcing amplitude, it was hypothesised that it would be possible to maintain the TRI at an amplitude that would not have been sufficient for its generation without the previous energy input.

Experimental procedure

To test the hypothesis of a hysteresis to the TRI mechanism, 36 experiments were conducted in a linear stratification of $N = 1.54 \pm 0.06 \text{ rad s}^{-1}$. Most of the experiments were run for 90 minutes and captured at 1 frame per second (with a small sub-set run for 180 minutes). Out of the 36, 25 used constant amplitude forcing, the results of which have been presented in § 4.3. The other 11 experiments varied the amplitude after 60 minutes of forcing to test the sub-critical bifurcation hypothesis. For these 11 experiments, the vertical displacement imposed on the neoprene foam also had the spatial structure $z = h(x, t)$ given in (4.1), however the temporal forcing $f(t)$ was governed by

$$\frac{f(t)}{e^{-i\omega_0 t}} = \begin{cases} 0 & \text{for } t \leq 0 \text{ s,} \\ \eta_0 \left(\frac{t}{30}\right) & \text{for } 0 < t \leq 30 \text{ s,} \\ \eta_0 & \text{for } 30 < t \leq t_1 \text{ s,} \\ (\eta_0 - a) & \text{for } t_1 < t \leq t_2 \text{ s,} \\ \eta_1 & \text{for } t_2 < t \leq t_{\text{end}} \text{ s,} \\ \eta_1 \left(\frac{1330-t}{30}\right) & \text{for } t_{\text{end}} < t \leq t_{\text{end}} + 30 \text{ s,} \end{cases} \quad (7.2)$$

where η_0 and η_1 are the initial and reduced amplitude respectively, t_1 is the beginning of the amplitude reduction and t_2 is the end time of the reduction. The amplitude reduction parameter a is given by

$$a = \left(\frac{t - t_1}{t_2 - t_1} |\eta_0 - \eta_1| \right). \quad (7.3)$$

For the 11 experiments of this set with variable amplitude, $t_1 = 3630 \text{ s}$ and $t_2 = 3660 \text{ s}$. Another set of experiments was also conducted that looked at the effect of extending the reduction time $t_2 - t_1$ for the amplitude from 30 s to 600 s. This was done to

increase the quasi-steady nature of the amplitude change. As the results are similar to those presented here, they have not been shown.

Experimental results

As discussed in § 4.3.1, for these experiments, the critical forcing amplitude required to trigger TRI lay around $\eta_0 \approx 3.9$ mm. Experiments were consequently run with η_0 in (7.2) slightly higher than this threshold then reduced to an η_1 that was slightly lower, to see if the instability was maintained. Figure 7.8 shows the results of four experiments whose amplitude was reduced after 3630 s ($t/T_0 = 560$), indicated by the black dashed line.

The amplitude plots in Figure 7.8(a) and (b) show that TRI is sustained after the forcing amplitude is reduced, while Figure 7.8(c) and (d) show that the secondary waves decay. What is necessary to determine here, is if the amplitude of the primary wave, after it is reduced, is still larger than the critical amplitude for instability. In order to de-construct the results from this full set of experiments, Table 7.1 quantifies the results of 18 experiments, some of which use constant amplitude forcing and some of which use the variable amplitude given in (7.2). They are ordered by decreasing values of $\langle |\tilde{\Psi}_0| \rangle_w$ and indicate both whether TRI is sustained and if the variable amplitude forcing is used. Only these 18 are shown – as opposed to the full 36 – as the experiments with larger amplitudes always lie well above the amplitude threshold and likewise the experiments with smaller amplitude always remain linearly-stable on the single-wave branch. These 18 are therefore the ones located around the threshold for instability.

The results recorded in Table 7.1 do not show compelling evidence that the TRI mechanism follows a hysteresis curve. Based on the hypothesis, experiments whose amplitude is reduced just below the amplitude threshold should sustain TRI, while those experiments of constant forcing just below the threshold, should not become unstable. Despite this, only experiments 12 and 13 appear to fit this hypothesis (highlighted in green). Both experiments clearly exhibit TRI below the expected amplitude threshold for linear instability, which seems to lie around $\langle |\tilde{\Psi}_0| \rangle_w \approx 10.67 \text{ mm}^2 \text{ s}^{-1}$. The four other variable amplitude experiments shown in Table 7.1, that sit below this amplitude threshold, did not sustain TRI.

The values of $\langle |\tilde{\Psi}_0| \rangle_w$ in Table 7.1 are also quite similar. These values are obtained by averaging over the primary wave beam for 900 s. The location of this averaging period changes depending on the experiment. In the case of constant amplitude forcing, this period is either just before the primary beam becomes unstable, or, in the case of no instability, it is the last 900 s of the experiment. In the case of variable amplitude

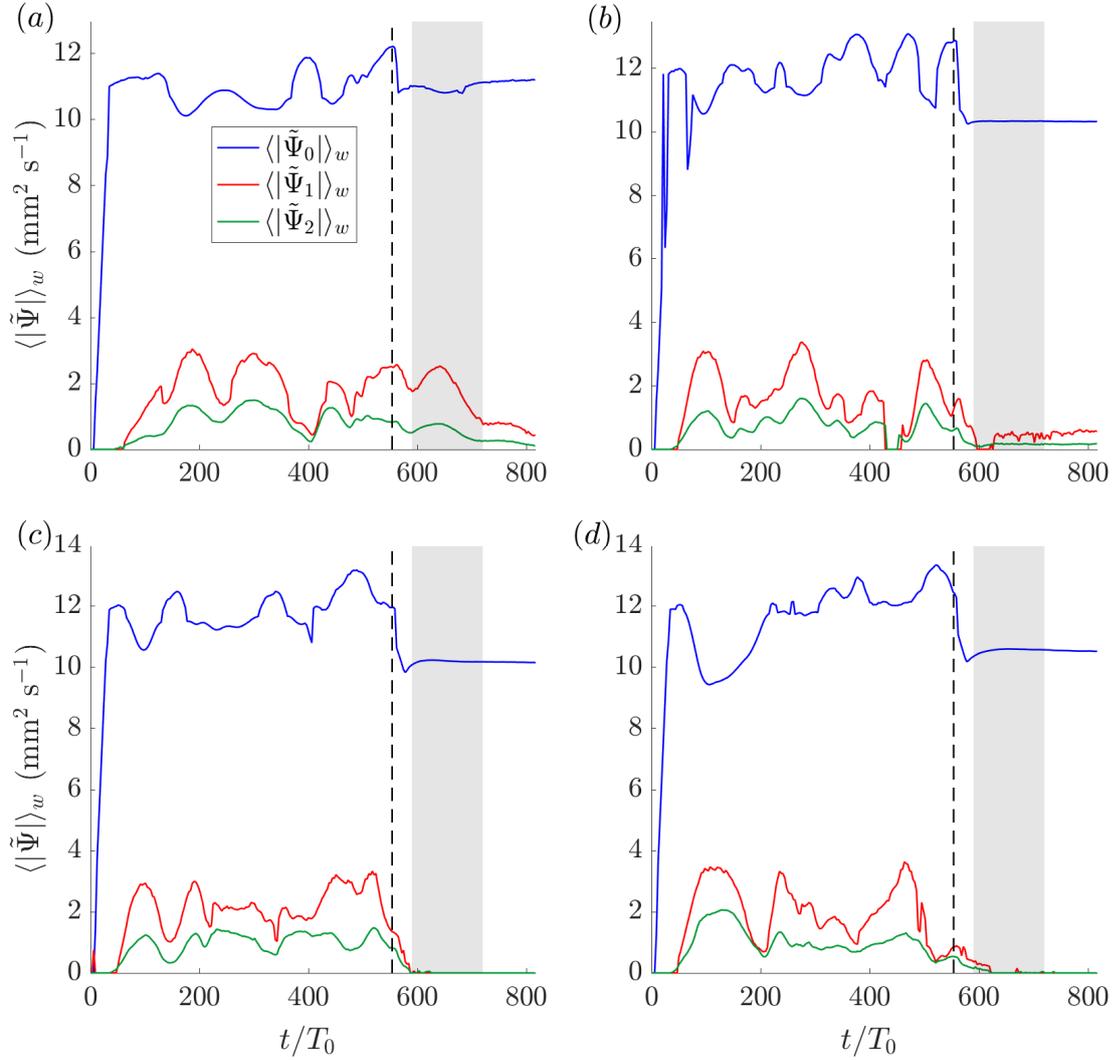


Figure 7.8: The amplitude of the reduced streamfunction $\langle |\tilde{\Psi}| \rangle_w$ for four different experiments. Each experiment had the forcing amplitude reduced after 3630 s ($t/T_0 = 560$), indicated by the black dashed line. The grey patch indicates the time over which the primary wave is averaged for the results in Table 7.1. The forcing amplitudes of each experiment were (a) $\eta_0 = 4$ mm, $\eta_1 = 3.5$ mm. (b) $\eta_0 = 4.5$ mm, $\eta_1 = 3.4$ mm. (c) $\eta_0 = 4.5$ mm, $\eta_1 = 3.3$ mm. (d) $\eta_0 = 4.5$ mm, $\eta_1 = 3.4$ mm.

Experiment no.	$\langle \tilde{\Psi}_0 \rangle_w$ (mm ² s ⁻¹)	TRI	variable forcing
1	10.95	yes	yes
2	10.85	yes	yes
3	10.84	yes	no
4	10.77	yes	no
5	10.74	yes	yes
6	10.67	yes	no
7	10.64	no	yes
8	10.56	no	yes
9	10.40	no	no
10	10.36	no	no
11	10.35	no	yes
12	10.34	yes	yes
13	10.23	yes	yes
14	10.18	no	yes
15	10.10	no	no
16	10.00	no	no
17	9.98	no	no
18	9.94	no	yes

Table 7.1: Table showing the results of 18 experiments with comparable amplitudes. The average amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$ is calculated over a 900 s period. In the case of constant amplitude forcing, this period is either just before the primary beam becomes unstable, or, in the case of no instability, is the last 900 s of the experiment. In the case of variable amplitude forcing, this period is the first 900 s after the amplitude is reduced, as shown by the grey shaded regions in Figure 7.8. The second column indicates if TRI was sustained across the duration of the experiment, while the final column indicates if the amplitude was changed, as shown in (7.2). Figures 7.8(a), (b), (c) and (d) correspond to experiments 1, 12, 14 and 8 respectively. The two experiments highlighted in green are the only two that fit the hypothesis of a hysteresis mechanism to the TRI.

forcing, this period is the first 900 s after the amplitude is reduced, as marked by the grey patch in Figure 7.8. If this averaging window is moved slightly, it is possible to change the amplitude order of the experiments in Table 7.1 as the values of $\langle |\tilde{\Psi}_0| \rangle_w$ are so close together. Given how sensitive this is to small changes in the averaging process, the presence of experiments 12 and 13 do not provide sufficient evidence to validate the hypothesis that the TRI mechanism follows a sub-critical bifurcation.

Overall, there does not seem to be evidence of a hysteresis to the TRI phenomena for the experimental set-up presented here. However, these results focus on only one set of forcing parameters and beam width. It might be possible to access the ‘linearly-unstable triadic-wave equilibrium state’ region, given in Figure 7.7(b), using other experimental forcing or via perturbing the system in another way.

7.4.3 Secondary wave experiments

Overview

The second way that the sub-critical bifurcation hypothesis was tested experimentally, was using a simple perturbation to the wave field that was not the result of a vortex ring. This perturbation came in the form of a second wave generated by ASWaM, that was fired into the primary beam. We will refer to this second wave as the ‘trigger’ wave. Specifically, a beam with parameters that fell in the \mathbb{W}_2 subgroup was used for the trigger wave. By firing a wave that is known to satisfy the conditions for one of the secondary waves in the triad, a two-wave interaction (TWI) occurs. Unlike TRI, where the primary finite-width beam requires sufficient amplitude to become unstable, a TWI interaction will form a resonant triad at any amplitude (see § 2.3 for details). After the triadic interaction was formed, the trigger wave forcing from ASWaM was removed to see if the triad would then be self-sustaining or if it would decay.

Experimental procedure

A set of 15 experiments were conducted in a linear stratification with a buoyancy frequency of $N = 1.47 \pm 0.06$ rad s⁻¹. For these experiments, this critical forcing amplitude required to trigger TRI lay around $\eta_0 \approx 3.8$ mm. Out of these 15 experiments, 7 had the forcing described in (4.1), while the spatial forcing for the remaining 8 was

$$z = h(x, t) = \begin{cases} \Re(f(t) e^{il_0 x} \cos^2(\frac{x-B}{8\pi^2})), & A < x < B, \\ \Re(f(t) e^{il_0 x}), & B < x < C, \\ \Re(f(t) e^{il_0 x} \cos^2(\frac{x-C}{8\pi^2})), & C < x < D, \\ \Re(f(t)_2 e^{il_2 x} \cos^2(\frac{x-B_2}{5\pi^2})), & A_2 < x < B_2, \\ \Re(f(t)_2 e^{il_2 x}), & B_2 < x < C_2, \\ \Re(f(t)_2 e^{il_2 x} \cos^2(\frac{x-C_2}{5\pi^2})), & C_2 < x < D_2, \\ 0, & \text{elsewhere,} \end{cases} \quad (7.4)$$

where the locations $A, B, C, D, A_2, B_2, C_2, D_2$ are $7\pi/|l_0|, 9\pi/|l_0|, 13\pi/|l_0|, 15\pi/|l_0|, 6\pi/|l_2|, 8\pi/|l_2|, 10\pi/|l_2|$ and $12\pi/|l_2|$ respectively. The two horizontal components of the wave vectors l_0 and l_2 are 0.05 mm⁻¹ and 0.08 mm⁻¹ respectively. The two beams described by (7.4) are both enclosed by a cosine squared envelope on the outer wavelengths. The primary beam, as before, is four wavelengths wide while the trigger beam is three wavelengths. The temporal forcing for the primary beam $f(t)$ has been

given in (4.2) while the temporal structure for the secondary beam $f(t)_2$ was given as

$$f(t)_2 = \begin{cases} 0 & \text{for } t \leq 0 \text{ s,} \\ \eta_2 \left(\frac{t}{30}\right) e^{-i\omega_2 t} & \text{for } 1200 \leq t \leq 1230 \text{ s,} \\ \eta_2 e^{-i\omega_2 t} & \text{for } 1230 \text{ s} \leq t \leq 1770 \text{ s,} \\ \eta_2 \left(\frac{t_{\text{end}}-t}{30}\right) e^{-i\omega_2 t} & \text{for } 1770 \text{ s} \leq t \leq 1800 \text{ s,} \\ 0 & \text{for } t \leq t_{\text{end}}, \end{cases} \quad (7.5)$$

where $\omega_2 = 0.61 \text{ rad s}^{-1}$ and η_2 is the amplitude of the secondary wavebeam. For these experiments η_2 is kept fixed at 2.5 mm. This temporal forcing turns the secondary wave beam on for 600 s during the experiment, increasing and decreasing over a 30 s period. The structure of this forcing is shown graphically in Figure 7.9 over $1200 < t < 1230$ ($181 < t/T_0 < 186$), during which the trigger wave is turned on.

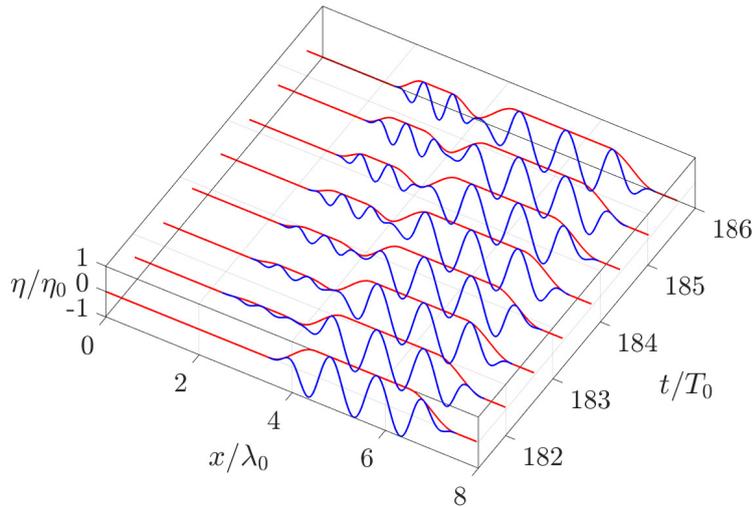


Figure 7.9: Waveform generated by ASWaM after $t = 1200 \text{ s}$ $t/T_0 = 181$ as the trigger wave is initiated. At this point, the primary beam comprising of four wavelengths is already fully formed. The secondary wave is increased from rest over $1200 < t < 1230$ ($181 < t/T_0 < 186$) and can be seen just to the left of the primary beam. The envelope of both beams is shown in red. The x , y and z axis are normalised by the wavelength λ_0 , period T_0 and amplitude η_0 of the primary beam respectively.

Experimental results

The snapshots in Figure 7.10 show a trigger wave experiment with $\eta_0 = 3.75 \text{ mm}$. In Figure 7.10(a) the primary beam has not become unstable. The next snapshot shows

the trigger \mathbb{W}_2 wave start from ASWaM and beginning to interact with the primary wave. By (c) the TWI interaction has emerged and \mathbb{W}_1 can be seen propagating down and to the left. When the trigger wave forcing is removed, initially the \mathbb{W}_1 beam persists, indicating that the TWI has provided sufficient energy for the primary wave to move to the linearly-unstable triadic-wave state (as shown in Figure 7.7)(b). However, by Figure 7.10(h), \mathbb{W}_1 is no longer present in the flow, meaning that the system has been attracted back to the linearly-stable single-wave branch and the trigger wave has left no lasting imprint upon the primary beam.

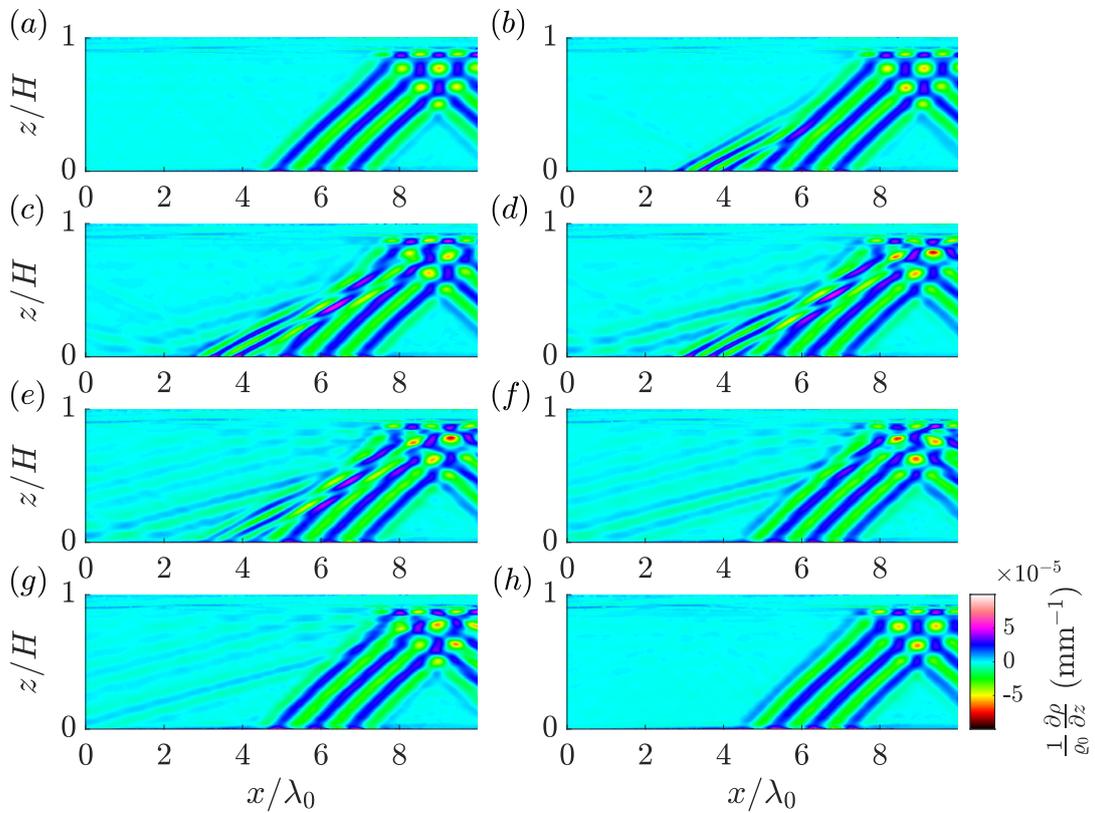


Figure 7.10: Sequence of images from a trigger wave experiment showing the vertical gradient of the density perturbation. The forcing amplitude of the primary beam is $\eta_0 = 3.75$ mm, while the forcing amplitude of the trigger wave is $\eta_2 = 2.5$ mm. The quantitative amplitudes of the waves are shown in Figure 7.11(c). (a) $t/T_0 = 180$, (b) $t/T_0 = 197$, (c) $t/T_0 = 212$, (d) $t/T_0 = 227$, (e) $t/T_0 = 270$, (f) $t/T_0 = 295$, (g) $t/T_0 = 303$, (h) $t/T_0 = 263$.

To see how this interaction looks quantitatively, Figure 7.11 shows the amplitudes of four experiments, two forced at a constant amplitude (a) and (b) and two with the trigger wave forcing, (c) and (d). Figure 7.11(c) corresponds to the snapshots shown in Figure 7.10. The first constant amplitude plot in (a) lies directly on the stability

threshold for TRI as the two secondary waves slowly begin to grow. The second constant amplitude plot in (b) lies below this threshold. For both Figure 7.11(c) and (d), the presence of \mathbb{W}_2 occurs before that of \mathbb{W}_1 , as \mathbb{W}_2 is used as the trigger wave. During the forced triadic interaction, the amplitude of the primary wave decreases, showing that energy is being transferred into the non-forced \mathbb{W}_1 . After the removal of the trigger wave, \mathbb{W}_1 persists for a short time period then also decays. Again, it is necessary to compare the amplitudes of the primary waves to see if any of the trigger wave experiments sustained the triadic interaction below the critical threshold. The results of 11 experiments are given in Table 7.2.

Experiment no.	$\langle \tilde{\Psi}_0 \rangle_w$ ($\text{mm}^2 \text{ s}^{-1}$)	TRI	trigger wave
1	9.02	yes	no
2	8.29	yes	no
3	8.09	no	no
4	7.83	no	no
5	7.30	no	yes
6	6.89	no	no
7	6.47	no	yes
8	6.08	no	yes
9	5.97	no	yes
10	5.93	no	no
11	5.18	no	yes

Table 7.2: Table showing the results of 11 experiments with comparable amplitudes. The average amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$ is calculated over a 900 s period. In the case of constant amplitude forcing, this period is either just before the primary beam becomes unstable, or, in the case of no instability, is the last 900 s of the experiment. In the case of forcing using the trigger wave, this period is the first 900 s after the trigger wave has been removed. All averaging windows are shown by the grey shaded regions in Figure 7.11. The second column indicates if TRI was sustained across the duration of the experiment, while the final column indicates if the trigger wave was fired as given in (7.5). Figures 7.11(a), (b), (c) and (d) correspond to experiments 1, 4, 5 and 11 respectively.

The results of Table 7.2 show that none of the trigger wave experiments caused the TWI interaction to promote the development of TRI. However, part of this could be due to the amplitude range of the primary wave. Again, the amplitude of $\langle |\tilde{\Psi}_0| \rangle_w$ is obtained by averaging over the primary wave beam for 900 s. For these experiments, in the case of constant amplitude forcing, this period is either just before the primary beam becomes unstable, or, in the case of no instability, it is the last 900 s of the experiment. In the case of forcing using the trigger wave, this period is the first 900 s after the trigger wave has been removed. All time averaging windows are shown

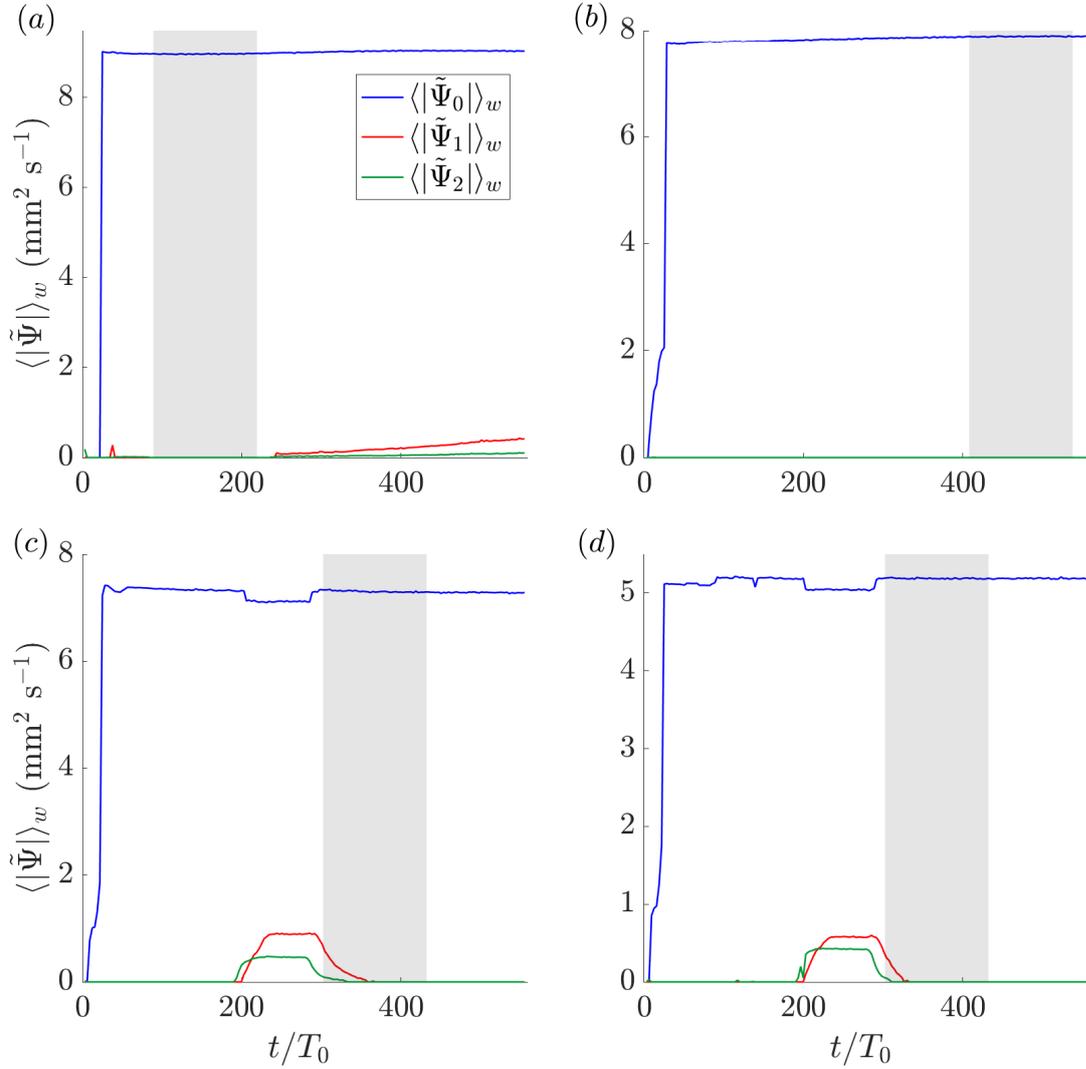


Figure 7.11: The amplitude of the reduced streamfunction $\langle |\tilde{\Psi}| \rangle_w$ for four different experiments. (a) and (b) both used the constant amplitude forcing described by (4.2), while (c) and (d) both show the results from the trigger wave experiments described by (7.4). The grey patch indicates the time over which the primary wave is averaged for the results in Table 7.2. The forcing amplitudes of each experiment are (a) $\eta_0 = 3.75$ mm. (b) $\eta_0 = 3.75$ mm. (c) $\eta_0 = 3.75$ mm, $\eta_2 = 2.5$ mm. (d) $\eta_0 = 3$ mm, $\eta_2 = 2.5$ mm.

by the grey shaded regions in Figure 7.11. For all the trigger wave experiments, the amplitude of the primary wave is considerably less than the critical threshold, which seems to lie around $\langle |\tilde{\Psi}_0| \rangle_w \approx 8.3 \text{ mm}^2 \text{ s}^{-1}$. The next experiment to use a trigger wave was at $\langle |\tilde{\Psi}_0| \rangle_w \approx 7.3 \text{ mm}^2 \text{ s}^{-1}$. At this amplitude the system is unable to access the ‘linearly-unstable triadic-wave equilibrium state’, given in Figure 7.7(b) and remains on the stable single triad branch. As the amplitude values are only translated from η_0 to $\langle |\tilde{\Psi}| \rangle_w$ after processing it is too late to run further experiments in the stratification at the correct amplitude.

Overall, these experiments provided no indication that the TRI mechanism could be triggered non-linearly through a TWI interaction, although this is possibly due to the wrong amplitude range of both W_0 and the trigger W_2 . It is also possible that the amplitude of W_1 and W_2 in the saturated state (linearly-stable triadic-wave branch) may be very small in the sub-critical region. If the trigger wave is of much larger amplitude, it may not be allowing the system to access the unstable equilibrium in the sub-critical region. If this is the case, then the instability in this region would be difficult to measure experimentally. Further work would be needed to explore this, looking at the amplitudes of W_0 closer to the threshold for the TRI or lower amplitudes of the trigger wave. This could also be tested using the \mathcal{M}_{2D} model.

7.5 Long term vortex ring interactions

7.5.1 Overview

In light of the results from § 7.4 it was necessary to re-examine the vortex ring experiments. Neither of the wave-only experiments performed to test the hypothesis of the sub-critical bifurcation indicated that the pathway to instability for the wave field was anything other than linear and could be mapped by a supercritical bifurcation. To further understand the vortex ring interaction experiments, the same experiment was conducted as those presented in § 7.3 although this time running for $t_{\text{end}} = 3600 \text{ s}$, to allow the flow to fully develop. Here we present the full run time of the experiments shown in § 7.3.3 conducted in a linear stratification of $N = 1.43 \pm 0.06 \text{ rad s}^{-1}$.

7.5.2 Long term development of vortex ring interactions

The amplitude plots presented in Figure 7.6 show that after the collapse of the vortex ring within the wave field, the two secondary waves in the resonant triad grow. To mirror the results obtained early on in the research, Figure 7.6 only presented the first 1200 s

of the experiment. In fact, these experiments were run for 3600 s and the full results are presented here. Figure 7.12 shows four of the vortex ring interaction experiments over the full duration. Figure 7.12(a), (b) and (c) are the same as those shown in Figure 7.6(a), (b) and (c). The final amplitude plot in (d) shows an additional vortex ring interaction experiment, again fired into the wave field at $t = 330$ s ($t/T_0 = 50$), indicated by the black dashed line.

Out of the four amplitude plots in Figure 7.12, only (d) appears to show that the passage of the vortex ring acts as a non-linear trigger for TRI. The other three amplitude plots in (a), (b) and (c) exhibit triadic resonance for approximately 990 s ($t/T_0 = 150$) before the resonant waves decay. Again, it is necessary to compare the amplitude of these experiments to see if (d) retained the instability due to the primary wave having sufficient energy to become unstable regardless of the ring. The results in Table 7.3 show 17 experiments from the set of comparable amplitudes, some with the vortex ring and some without.

The results from Table 7.3 show that the experimental stability threshold lies around $\langle |\tilde{\Psi}_0| \rangle_w \approx 8.35 \text{ mm}^2 \text{ s}^{-1}$. Only two experiments below this threshold (experiments 11 and 13 highlighted in green) sustain TRI for the full duration, both of which are vortex ring experiments. While both 11 and 13 exhibit TRI, the other vortex ring experiments, 7 and 9, do not show the instability despite both being at a higher amplitude. It seems unreasonable, therefore, to claim that the vortex ring is able to act as non-linear for the TRI mechanism.

Regardless of whether the TRI was sustained for the full hour, all vortex ring experiments triggered triadic resonance for approximately 990 s. For both sets of experiments presented in § 7.4.2 and § 7.4.3, the instability decayed almost immediately after either the amplitude was reduced or the trigger wave was removed. Some aspect of the vortex ring manages to sustain the triad for a considerable time despite the rings transient nature. It seems unlikely that it is due to the wave numbers and frequencies of the internal waves produced, as the trigger wave experiments also provide the correct input wave parameters for the triad. It could possibly be due to the three-dimensional nature of the ring. The mixing generated by its passage strongly perturbs the wave field and the surrounding stratification. This discontinuity to the background density field will dissipate at a much slower rate than the collapse of the ring, potentially explaining the sustained imprint of the vortex rings passage on the wave field. Although it appears unlikely from the investigation that the pathway to TRI is non-linear, an investigation into the properties of the vortex ring that trigger this behaviour would certainly be worthwhile.

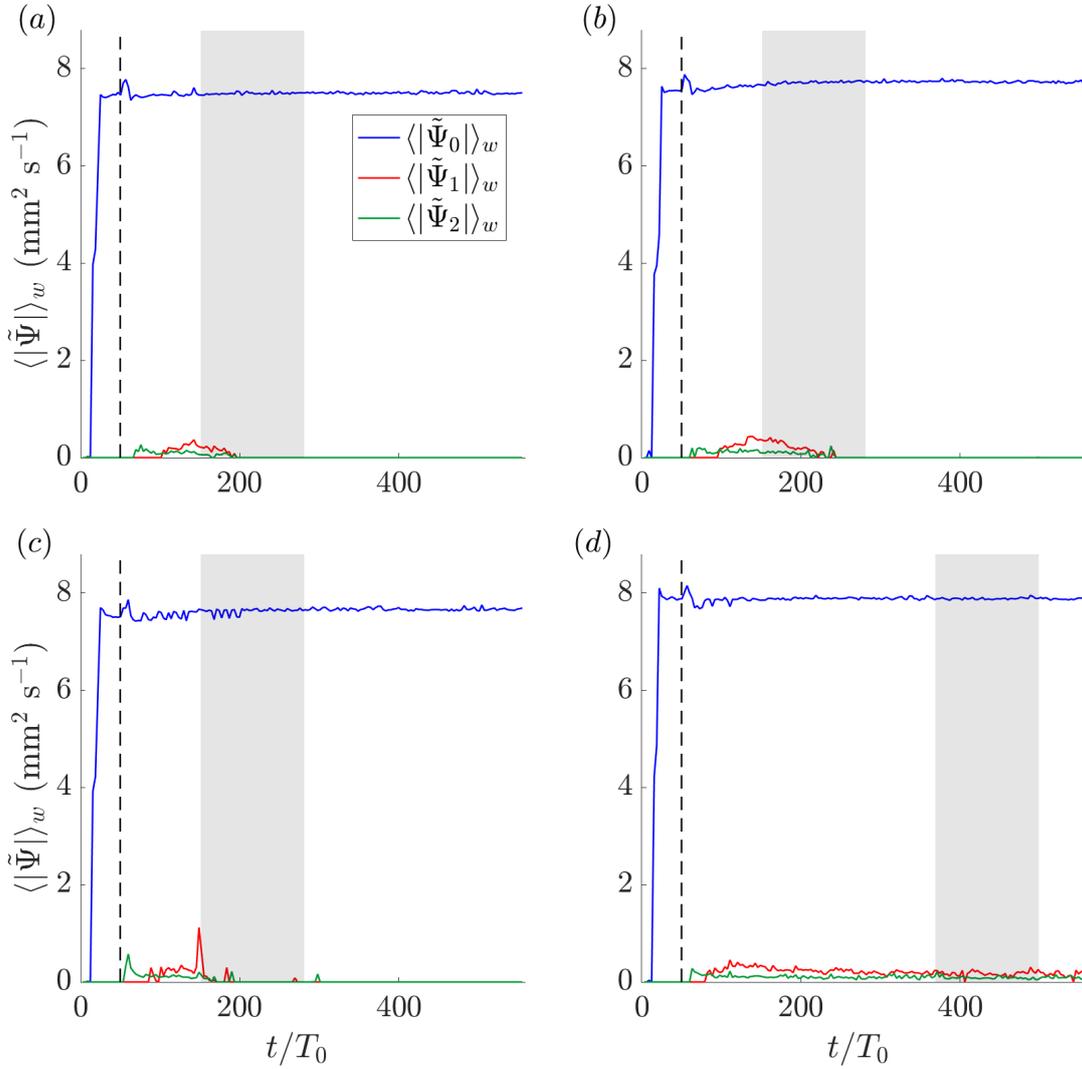


Figure 7.12: The amplitude of the streamfunction for $\langle |\tilde{\Psi}_0| \rangle_w$ (blue), $\langle |\tilde{\Psi}_1| \rangle_w$ (red), $\langle |\tilde{\Psi}_2| \rangle_w$ (green) for four different experiments. All amplitudes were calculated over the whole visualisation window $\langle |\tilde{\Psi}| \rangle_w$ (see Figure 3.23 for details). In three of the experiments (a), (b) and (c) a vortex ring of $F_p = 4.4$ was fired into the wave field at $t = 330$ s ($t/T_0 = 50$), indicated by the black dashed line. The grey patch indicates the time over which the primary wave is averaged for the results in Table 7.3. The forcing amplitudes of the primary wave beam are (a) $\eta_0 = 3.5$ mm, (b) $\eta_0 = 3.6$ mm, (c) $\eta_0 = 3.6$ mm and (d) $\eta_0 = 3.7$ mm.

Experiment no.	$\langle \tilde{\Psi}_0 \rangle_w$ (mm ² s ⁻¹)	TRI	vortex ring
1	8.77	yes	yes
2	8.75	yes	yes
3	8.61	yes	yes
4	8.58	yes	no
5	8.39	yes	no
6	8.32	no	no
7	8.29	no	yes
8	8.08	no	no
9	8.02	no	yes
10	7.96	no	no
11	7.94	yes	yes
12	7.91	no	no
13	7.88	yes	yes
14	7.86	no	no
15	7.81	no	no
16	7.72	no	no
17	7.72	no	yes

Table 7.3: Table showing the results of 17 experiments with comparable amplitudes. The average amplitude of the primary wave $\langle |\tilde{\Psi}_0| \rangle_w$ is calculated over a 900 s ($t/T_0 = 136$) period. In the case of the simple experiments with no ring, this period is either just before the primary beam becomes unstable, or, in the case of no instability, is the last 900 s of the experiment. In the case of forcing with the vortex ring this period is between $1000 < t < 1900$ s ($150 < t/T_0 < 288$). All averaging windows are shown by the grey shaded regions in Figure 7.12. The second column indicates if TRI was sustained across the duration of the experiment, while the final column indicates if a vortex ring was fired. Figures 7.12(b) and (d) correspond to experiments 17 and 13 respectively, while the amplitude of Figures 7.12(a) and (c) are too small to be listed on this table.

7.6 Summary

This chapter was concerned with exploring the effect of deliberately perturbing the internal wave field to understand the pathway to instability. Initial experiments were presented exploring the interaction between a freely evolving turbulent structure and an internal wave field. We examined the impact of both of these structures on the evolution of the other. First we showed how the wave field impeded the propagation of the vortex rings, causing them to undergo a premature collapse. We then showed how the passage of the ring through the wave field also left a imprint on the flow. Several experiments were presented showing how this interaction appeared to provide a non-linear trigger for TRI in the wave field at amplitudes below the threshold for instability for the finite-width beam.

These experiments lead to the hypothesis that the pathway to instability may follow a sub-critical bifurcation. We proposed the idea that this sub-critical bifurcation may contain a hysteresis, meaning that it would be possible for the internal wave system to correspond to either the linearly-stable single-wave branch or the linearly-stable multiple wave branch (shown in Figure 7.7(b)) depending on its history or previous perturbations. This hypothesis was tested through two different experiments that both removed the complex dynamics of the vortex ring's passage. The first attempted to access the linearly-stable multiple wave branch through variable amplitude forcing. Initially, by forcing the system above the amplitude threshold, TRI would be exhibited. Then by decreasing the forcing amplitude we hypothesised that the wave field would remain the linearly-stable triadic-wave branch despite being forced below the critical amplitude required for instability. While somewhat inconclusive, the results of these experiments did not show this to be the case. The second set of experiments conducted to test this hypothesis used a 'trigger' to generate a TWI. Again, it was hypothesised that after the removal of the trigger wave the triad would be self-sustaining, indicating that TRI had occurred. However, the results showed that the none of the TWIs were able to promote TRI, only a long-lived triadic resonance.

The final part of this chapter returned to the vortex ring interaction experiments, this time considering them over a full hour. It was found while the resonant triads actually decayed for most of the experiments, two of them indicated that TRI had been generated from the passage and/or collapse of the ring within the wave beam. In addition, the vortex ring appeared to elicit a stronger non-linear reaction from the primary wave compared with the other 'wave only' sub-critical bifurcation experiments. Potentially the vortex-rings three-dimensional nature and impact on the surrounding stratification might explain why triadic interactions were sustained for a considerable time after its collapse.

Chapter 8

Conclusions

8.1 Introduction

The primary aim of the research presented in this thesis was to understand in detail both the initial growth and the long term development of Triadic Resonance Instability in finite-width internal wave beams. As discussed in Chapter 1, understanding the pathways of internal waves to small scales is fundamental in quantifying their role in global ocean circulation and mixing. While TRI has been cited as a potential non-linear pathway to turbulence in the ocean, there are significant developments to be made regarding our comprehension of this instability in oceanic settings. Indeed, a review by Sutherland (2013) acknowledges this disconnect between theory, which assumes the waves are periodic in space and time, and reality, in which waves are transient and localized.

To the best of our knowledge, very few experimental studies have examined in detail the long term development of TRI in the context of a finite-width beam. The experimental work by Bourget et al. (2014), addressed the importance of finite-width effects, although their theoretical analysis did not fully consider the spatial localisation of the triadic interaction. Moreover, their work focused on the transient start up of the instability and did not examine its evolution over long time periods. Later work by Brouzet et al. (2017) examined the development of TRI over long time periods in internal wave attractors domains, although their focus was on the role of confinement to the development of the secondary waves. How the instability unfolds in unbounded domains has received less attention. These unexplored aspects into the nature of the instability in a more realistic context of a freely propagating finite-width beam formed the basis for the motivation behind the research presented here.

In this final chapter we aim to draw together some of threads running through this work by summarising the results and the key research findings in § 8.2. A discussion of this work with some new perspectives is then given in § 8.3. Based on this discussion and the key results, we will address the research challenges that have been presented and some of the potential avenues for future work that have naturally arisen from this project in § 8.4 and § 8.5. Finally, the overall implications that this research has on the wider community is addressed in the concluding remarks in § 8.6.

8.2 Thesis summary

8.2.1 Finite-width beam

The main focus of this research has been understanding how a finite-width internal wave beam affects the development of TRI. One of the challenges with examining the behaviour of internal waves experimentally is that, in an enclosed tank, there exists only a finite time period before the waves reflect from a boundary and interact with the domain of interest. For the experiments presented here, an 11 m long tank was utilised to study the development of the waves. Using such a long tank ensured that wave reflections had been sufficiently delayed and damped by viscosity so that their impact on the instability – due to interaction with the primary beam – was negligible.

Using the novel wavemaker discussed in § 3.3, experiments were presented in § 4.3 looking at a quasi two-dimensional, four wavelength wide, internal wave beam in a linear stratification with $\omega_0/N \approx 0.62$. Multiple experiments were conducted with a run time of 90 minutes. We saw that for a range of forcing amplitudes, the primary wave beam generated from the wavemaker became unstable via TRI. While in itself not unusual – this instability has been witnessed experimentally for a number of years (e.g. McEwan and Plumb, 1977; Benielli and Sommeria, 1998; Joubaud et al., 2012) – it was the evolution in behaviour of the instability that was of interest. The \mathbb{W}_1 beam, which for the given experimental configuration was the only secondary wave immediately visible in the full flow field, continuously migrated over the full height of the primary beam. During this migration, the structure of the beam also varied, showing periods of large amplitude concentrated over a narrow width and other periods where the beam was wider with lower amplitude.

This state of constant flux was intriguing, especially after examining the published literature on the subject of TRI in finite-width internal wave beams. Numerically integrating the couple of ODEs presented by Bourget et al. (2014) predicted that after

the initial transient start-up, all waves in the triad should reach a steady equilibrium in which the transfer of energy from the primary beam to the secondary beams is constant. In contrast, the experimental results showed an ongoing fluctuation between the amplitude of the primary beam and the two secondary beams. To further understand these interactions, an analysis was undertaken in Fourier space, looking at how the frequency and wavenumbers of the secondary waves evolved over time. Through the use of Dynamic Mode Decomposition and Fast Fourier Transforms we saw how these parameters continuously varied over the course of the experiment with the variation exhibiting a comparable typical period to that of the amplitude. The spectrograms obtained from the FFTs also revealed how the frequency bandwidth for the secondary waves varied over time. Despite these variations, the frequencies obtained from the DMD satisfied the temporal triadic condition at all times.

These results highlighted what we believe to be the quasi-periodic cyclic growth and decay of different triadic perturbations. We suggested that this oscillatory behaviour stems from the finite-nature of the beam, which restricts the triadic interaction in both space and time. As the secondary waves grow and draw energy from the primary beam, the amplitude of the primary wave drops. However, the secondary waves are also propagating and at some point exit the underlying primary beam and the energy transfer is broken. Over time, if this energy exchange can not be maintained then this particular triadic interaction decays, allowing the primary beam to recover in amplitude. At this point, further triadic perturbations – of slightly different wave number and frequency – are able to grow and this aperiodic cycle repeats. In an infinitely wide beam, we believe that the secondary waves of the triad would continue to draw energy until an equilibrium state is reached and this oscillatory behaviour would not be found.

8.2.2 Weakly non-linear modelling

Motivated by these novel experimental results, a weakly non-linear investigation was presented that aimed to capture the spatio-temporal properties of the TRI interaction. This was done through a perturbation expansion that set the slowly evolving amplitude of the waves as a function of both time and space. This expansion was substituted into the full non-linear equations of motion and the model was built by collecting around subsequent orders of the non-dimensional parameters. The solution of the resultant equations required the development of a two-dimensional numerical scheme, \mathcal{M}_{2D} , based on the leading order advection terms of the expansion, onto which the non-linear interactions and viscous effects were incorporated.

Results from the \mathcal{M}_{2D} model were presented in Chapter 6. We first examined the case of a ‘simple triad’, where only three waves were input into the system. Focusing on the triadic parameters that were selected experimentally, we saw how small changes to these secondary waves significantly affected the evolution of the instability. There existed a small range of input parameters over which the amplitude of the triad underwent the same quasi-periodic fluctuations seen experimentally. At one limit of this range, specifically $\kappa_1/\kappa_0 < 1$, the instability never developed. Further analysis revealed that, regardless of the forcing amplitude, the instability was never triggered when $\kappa_1/\kappa_0 < 1$, which would result in the transfer of energy to larger as well as smaller scales. This limit was also seen experimentally. At the other end of the parameter range, as both κ_1/κ_0 and ω_1/ω_0 increased, \mathbb{W}_1 and \mathbb{W}_0 became closer to being parallel. We believe that this increased interaction time for a given packet of energy resulted in significant decay of the primary beam as the triad was able to fully develop.

Results from the model were then presented whereby the parameter space for the secondary waves was extended to permit multiple triadic perturbations into the flow. We examined an \mathcal{M}_{2D} model run that included 60 different pairs of secondary waves all connected through the same primary beam. This run saw the flow develop in a manner more comparable to the experiments, with the \mathbb{W}_1 beam showing large variability in both its origin and width. The modulations in Fourier space of the resonant waves were also explored. The spectrogram presented in Figure 6.12 qualitatively resembled the experimental spectrograms, with periods of wider and narrow frequency bandwidth for the secondary waves, although ω_1 and ω_2 were closer to the PSI limit than those selected experimentally. We believe this to be due to both the lower input amplitude for the multiple frequency simulation and distribution of the frequencies in the system. Overall, while there were differences in detail between the experiments and model, the latter was able to capture the essence of the spatio-temporal modulation of the secondary waves.

8.2.3 Deliberate perturbations

The final part of this thesis focused on a different aspect of the instability. Specifically, we were interested in seeing if deliberate perturbations to the wave field were able to induce instability at amplitudes below the stability threshold. The motivation for these experiments lay in understanding the various pathways to instability, as the ocean is never in a state of quiescent equilibrium.

Overall, it was found that out of the three deliberate perturbations presented, the only perturbation that consistently triggered triadic resonance was the interaction of

a vortex-ring with the wave field. While the triadic resonance did not always persist for the full duration of the experiment, something about the vortex ring gave it the ability to non-linearly trigger the triad in ways that other perturbations could not. The reason for this may be due to the three-dimensional nature of the ring, which not only perturbed the wave field but generated a discontinuity in the stratification from its turbulent propagation. As this discontinuity spread and thinned out across the tank, the stratification became smoother and no longer provided the ‘kick’ for the wave field to become unstable.

Based on these experiments, the nature of the pathway to instability was called into question. An hypothesis was posed suggesting that the instability may be subject to a sub-critical bifurcation, and that there exists a state in phase space where the system could lie on either on a linearly-stable single-wave branch or a linearly-stable triadic-wave branch. These ideas are discussed further in § 8.3.

8.3 Discussion

8.3.1 Connection to dynamical systems

One of the key findings from this research is that once the wavebeam has become unstable via TRI, the state of the triad is continuously varying. What has not been discussed is whether this unsteady behaviour of the system can be represented by a dynamical system. This idea emerged from the vortex ring interaction experiments when it was hypothesised that the triad could correspond to two different linearly-stable equilibrium states and that deliberate forcing could cause the system to ‘jump’ branches due a sub-critical bifurcation.

Current theory suggests that, for a finite-width beam, TRI can be described by two linearly-stable equilibrium states in phase space, one with a single-wave and one with a triad of waves. These states are accessed via changing the forcing amplitude; once a given threshold has been surpassed the system moves from the linearly-stable single-wave branch to the linearly-stable triadic-wave branch through a supercritical bifurcation, shown in Figure 7.7(a). This idea corresponds to the underlying assumption that the triad is comprised of three distinct waves, all with discrete wave numbers and frequencies. Yet this description of the system does not explain the continued unsteadiness to the behaviour being witnessed.

Here, we speculate on the possibility that rather than two discrete stable equilibrium states, there exists a spectrum of points in phase space corresponding to linearly-

unstable equilibrium triadic states. This spectrum could be represented via a plane perpendicular to the amplitude axis in Figure 7.7 as shown in Figure 8.1. As the system is forced higher in amplitude, it becomes attracted to one of these unstable-equilibrium states, which might take the form of a saddle point. This is represented by plane C in Figure 8.1. Here, the solution would be attracted towards the stable branch of the manifold leading to the saddle point. Despite this attraction, the system can never truly access the equilibrium of the saddle point and the solution is thrown away from this node via the unstable branch of the manifold. In this manner the system moves chaotically through the phase space, continuously being attracted to different unstable triadic saddle points but never accessing these singular nodes.

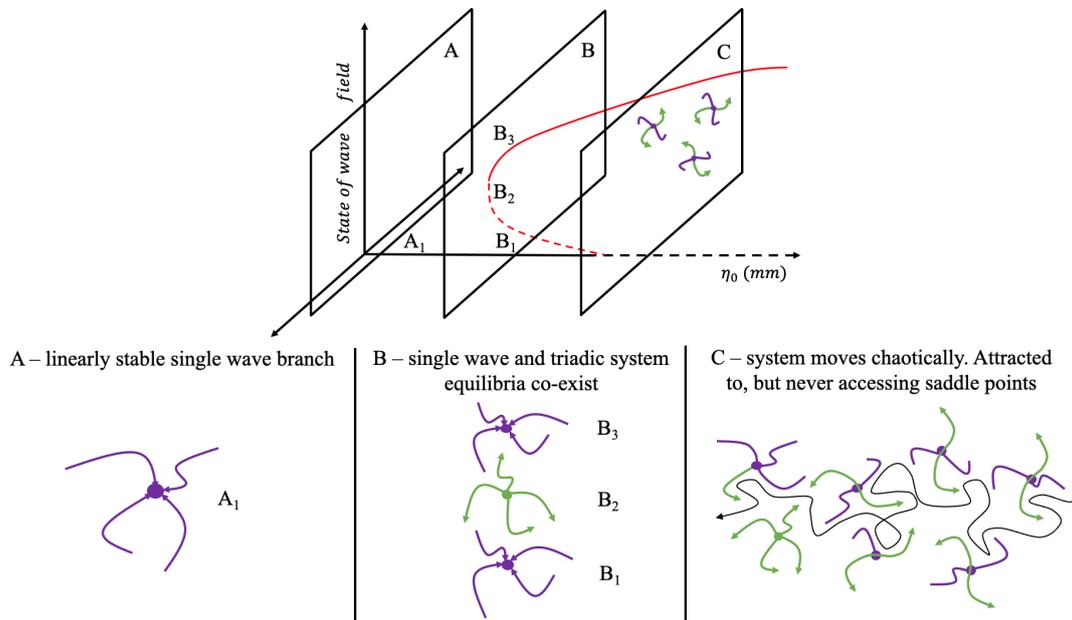


Figure 8.1: Schematic showing the connection between the TRI mechanism in a finite-width beam and a dynamical system. The planes, A, B and C all correspond to different states of the system. In A, the system is only attracted to the linearly-stable single-wave branch. As the amplitude increases, B represents the phase space where the system is can be attracted to either the linearly-stable single-wave branch or to a linearly-stable triadic-wave state, represented by B_1 and B_3 respectively. The state B_2 indicates the unstable equilibrium where the system is repelled from. As the amplitude increases further, plane C corresponds to the solution being continuously attracted to linearly-unstable equilibrium triadic states. Here the solution is attracted towards the saddle point via the stable branch of the manifold, only to be thrown away via the unstable branch.

This idea is embodied in the form of the weakly non-linear model presented in § 6.3, where multiple possible resonant waves were included into the system. In this case, the phase plane was comprised of the discrete set of frequencies given to the model.

By tracking the triad with the maximum energy in the flow, what we saw was that the solution continuously jumped between different perturbations, never settling on one triad. This could relate to the ideas above, that the system moves through this discrete phase space being continuously attracted to the stable branch of each manifold but never reaching a stable equilibrium point. For some of these states that the system is attracted to, it may enter an unstable periodic orbit, where it can spend more time around the saddle node, explaining why the triad with the maximum energy sometimes remained constant for long periods before jumping.

Indeed, even the results exploring the ‘simple triad’ may provide insights into the behaviour of the TRI system in phase space if the equilibrium is (or becomes) a saddle point. This investigation showed that for some input parameters the system underwent unsteady modulations while for other sets of parameters a stable equilibrium was achieved. This could indicate that the former are unstable equilibrium triadic states, while the latter are actually stable equilibrium points that the system is able to access.

Another unexplored element of these experiments is that these triadic resonant interactions, emerging from one primary wave, could be generating wave turbulence. Wave turbulence, which is a still relatively unexplored phenomenon, refers to physical systems with a large number of dispersive and non-linear interacting waves (Davis et al., 2020). The presence of this large amplitude wave turbulence in the ocean, is believed to contribute to the internal wave energy cascade (Beckebanze et al., 2021). The lack of steady equilibrium in the experiments and the seemingly random motion may suggest wave turbulence being generated from a single TRI due to other TWI / TRI interactions from secondary waves.

8.3.2 Three-dimensional considerations

As discussed in the introduction to Chapter 5, a decision was made early on to avoid full Computational Fluid Dynamics (CFD) modelling of the internal waves. This was in order to isolate the dynamics of the TRI that were being witnessed and obtain a more thorough understanding of the flow as opposed to numerically recreating the experiments. While this proved advantageous, many elements of the full flow field are not captured in the weakly non-linear modelling.

One of the elements not captured by the model is the role that the beam envelope plays on the wavenumbers introduced into the flow. Due to the model construction, it is only possible to introduce one domain for the primary beam. This domain, W_0 , comprised a single frequency and wave number. In reality, the finite-width of the beam envelope introduces a whole spectrum of wave numbers in Fourier space, all with

frequency corresponding to the primary beam. Not only do these wave numbers have the ability to undergo TRI themselves, they also propagate at different group velocities. As mentioned in § 2.4, these cross-beam variations generate horizontal mean flow (Bordes et al., 2012). While part of this re-circulating mean flow will be suppressed by the side walls in the experiments presented here, qualitative dye experiments, shown in Figure 8.2, depict that a strong three-dimensional flow – in the direction of the primary waves group velocity (up and to the right) – was still present within the confines of the primary beam.

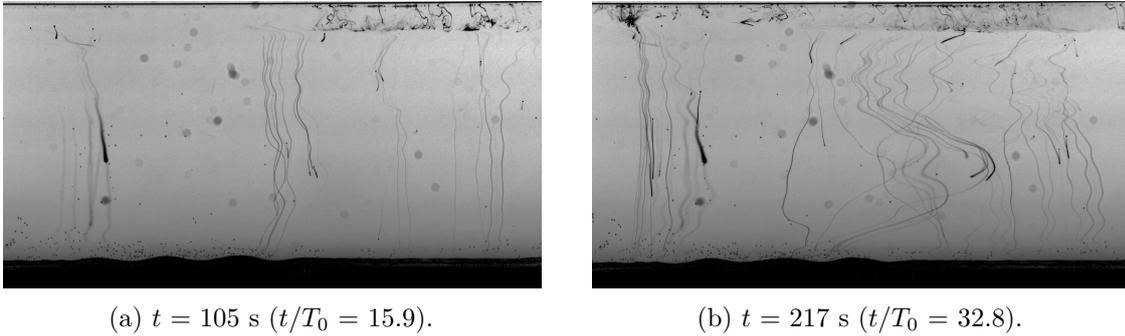


Figure 8.2: Two photos at 112 s ($t/T_0 = 16.9$) apart of a dye experiment conducted using the same forcing parameters for \mathbb{W}_1 outlined in § 4.2. The waveform can be seen on the neoprene cover of ASWaM. The dot pattern for the Synthetic Schlieren has been removed and dye crystals are added to the tank from the free surface. The same quasi-vertical streaks seen in (a) in the middle of the wavebeam have moved considerably by (b) in the direction of the group velocity of the primary wave.

Currently, it is not known how this induced mean flow interacts with the triad. Work by Fan and Akylas (2019), looking at the interaction of a beam with horizontal background mean flow show that the mean flow actually suppresses the triadic interaction over a range of short scale parameters. This is due to the mean flow affecting the group velocity of the perturbations, meaning a packet of energy exits the primary wave beam before being able to extract significant energy via the resonant triad interactions. Their findings were, however, considering the effects of horizontal mean flow. If the mean flow was in the same direction as the group velocity of the primary beam – as we see experimentally in Figure 8.2 – the mean flow could act to advect the perturbations vertically as well as horizontally. This modified group velocity of the triadic perturbations could alter the interaction region of the triad in both space and time. If the triadic perturbations were advected in the same direction as \mathbf{c}_{g_0} , this may explain why experimentally the triad was able to re-ignite near the free surface, yet in the simulations this never occurred. Due to the two-dimensional nature of the model, these three-dimensional mean flow effects cannot be incorporated.

This observation by Fan and Akylas (2019) also agrees with the hypothesis that the triadic interactions never grow on the inner branch of the loci curve, despite the theoretically higher $\hat{\sigma}_{0D}$, due to the horizontal component of all the group velocities in the triad being in the same direction for the inner branch. Here, the horizontal component of every wave in the triad is in the same direction. Having opposing directional components, as seen for the triads on the outer branch of the loci curve, significantly increases the interaction time of the triad, allowing for sufficient energy transfer from the primary wave.

8.3.3 Role of beam configuration

For all of the experiments presented, the width of the beam was always kept at four wavelengths wide, modulated by the same cosine squared envelope. In Chapters 4 and 6, we have hypothesised that part of the observed modulations are related to how long the triadic interactions are contained within the spatial extent of the primary beam. If these modulations are a product of finite-time interactions then they are inherently also a function of the beam width. As we have seen with the weakly non-linear \mathcal{M}_{2D} model, the parameter space over which these oscillations occur is narrow. Wider beams that prolong this interaction may see the selected range of ω_1/ω_0 shift towards smaller values where the interaction time between the beams becomes shorter as the beams get closer to being perpendicular. Conversely, the opposite may be seen in a narrower beam where the interaction time of the triad reduces.

There must also exist a limit on the minimum number of wavelengths required for instability. Work by Karimi and Akylas (2014) found a theoretical limit which was a function of beam amplitude, the angle of both the primary and secondary beams and the specific envelope shape for the beam. Testing this limit would be valuable experimentally, as we saw a significantly lower amplitude threshold required for instability in the experiments as opposed to the model. Further ideas for future work stemming from this research are presented below.

8.4 Avenues for future experimental and analytical work

Inevitably, a wide range of further research avenues are apparent from the work presented in this thesis. This section breaks down these ideas and discusses the work that would be the most valuable to pursue.

8.4.1 Wave forcing parameters

As discussed above, several hypotheses discussed in this thesis are based on how the finite-width beam influences the interaction of the triad. One of the simplest ways to investigate this idea further is to vary the forcing used for the primary wave beam. Specifically, the two parameters that would be worthwhile investigating are the width of the primary and the modulating envelope.

In § 8.3.3 we hypothesised that increasing the beam width would reduce the ratio ω_1/ω_0 , as this would in turn reduce the interaction time of the triad. Due to the unique capabilities of ASWaM, this could be done over the course of one experiment. The beam width could be gradually increased or decreased over a long forcing period. Then the spectrum of frequencies corresponding to the secondary waves could be examined in Fourier space. If over the course of the experiment the range of selected frequencies shifted in accordance with the change in width, experimental evidence could be drawn showing how the secondary wave parameters are a function of the finite-width beam.

The nearly inviscid theoretical predictions by Karimi and Akylas (2014) could also be explored. They propose that only for certain envelope configurations will TRI be permitted. Beams with general locally confined profile would remain stable to short scale sub-harmonic perturbations. Although this analysis is performed in the nearly inviscid limit, how it translates to viscous waves would be of interest. These ideas could be tested through both experiments and using the \mathcal{M}_{2D} model.

From the weakly non-linear model we saw there existed a range of parameters for the ‘simple triad’ for which the instability obtained a steady equilibrium. Experimentally, however, this state was never selected despite a full range of triadic parameters in the flow. Understanding why this unsteady behaviour is selected experimentally may also be explored by changing the input forcing parameters, as different beam widths and configurations may exhibit different oscillatory behaviour.

8.4.2 Weakly non-linear modelling

The development of the weakly non-linear model has provided a high level of insight into the role of the secondary wave parameters on the development of TRI. The next steps with developing this model could be first to examine the effect of changing both the range and the distribution of the secondary wave frequencies in the multiple secondary wave runs, discussed in § 6.3. We have hypothesised that some of the differences in the amplitude modulation between the model and the experiments is due to the distribution of discrete frequencies in the flow. As the secondary waves can

interact non-linearly through the primary wave, they may ‘beat’ with other triads over long time periods. Changing the distribution of the secondary waves across the loci curve may change the period of observed modulations. Also in terms of a dynamical system, changing the distribution would shift the saddle points in phase space. As the number of modes is increased and the system is driven towards a more non-linear state, we might anticipate wave turbulence to arise. To capture this fully, the model would need to be extended to include TWI between \mathcal{S}_{1_j} and \mathcal{S}_{2_j} .

As briefly examined in § 6.4, higher order terms from the perturbation expansion can also be included into the model. There remains a question, however, as to the relative contribution of each higher order and the weighting they should be afforded. Currently, only the $\mathcal{O}(\epsilon^2\gamma^1)$ terms are incorporated into the model while the $\mathcal{O}(\epsilon^1\gamma^2)$ terms of similar magnitude are neglected. Noting that γ is a function of beam width Λ_0 and ϵ is a function of amplitude, the role of these additional non-linear terms should be explored in terms of the parameters of the primary beam, to understand the role of these higher order expansion terms to see how much additional physics are explained within them. Also, as $\chi \sim 1/Re$, this term would be negligible in an oceanographic context and the model could be adapted to more appropriate oceanographic parameters.

The other unexamined aspect of the $\mathcal{O}(\epsilon^2\gamma^1)$ terms shown in (5.26) are the values of $\boldsymbol{\alpha}$ and β . Analysis of the vector $\boldsymbol{\alpha}$ did not reveal any physical significance or understandable adjustment to the \mathbf{c}_g vector. As these higher order terms do not currently play a crucial role in the model, further investigation into the physical meaning behind these coefficients was left to future work, but may provide valuable insights into their higher order contributions.

As the model was built to provide insight into the experimental results, it can also be modified to other experimental set-ups. It would be simple to adapt the model to investigate the role of beam width or wave envelope within spectral limits, as discussed above.

8.4.3 Vortex ring interactions

One element of this work that would be pertinent to explore further is the vortex ring interaction experiments. Based on the data presented in § 7.5, there are still unexplained aspects of this investigation. Firstly, there is some evidence that the triadic resonance generated from the imprint of vortex-ring was self sustaining and the system became unstable to TRI. In light of the potential relationship of the TRI mechanism to a dynamical system – for the parameter space considered – we show that if there is a range in phase space where the single-wave and triadic system equilibria

co-exist, it is small. A worthwhile investigation would be to extend this parameter space to further explore this region of possibly multiple equilibrium solutions and the how the system transitions.

In addition, there are several elements to this experiment that are still not fully understood. All of the vortex rings used in these experiments collapsed within the width of the primary beam, as shown in Figure 7.4. Rings with a higher initial Fr number would propagate further and collapse outside of the primary beam. Conducting experiments with rings of higher Fr , would help to separate the dynamics causing the develop of the triad; the imperfections due to the stratification or the aspects of the rings collapse within the primary beam.

If indeed it is aspects of the vortex ring triggering the generation of the triad, there also remains an open question as to which element of the ring is important. At this point it is unclear whether it is the strength of the waves that are produced while the Fr number of the ring is sufficiently larger than unity, the spectrum of the waves produced as Fr passed through unity during its collapse, or the residence time of the vortex ring inside the wave beam.

Another element highlighted in these vortex-ring experiments was that how \mathbb{W}_2 developed first in the flow. This was in contrast to the experimental TRI results where both secondary waves grew at a similar time. This early growth of the \mathbb{W}_2 beam could be due to the horizontal velocity alignment between the vortex ring and \mathbb{W}_2 . It would be worthwhile exploring this interaction in further detail by changing the orientation of the vortex ring tube so that the ring was fired with an opposite horizontal velocity to that of \mathbb{W}_0 and \mathbb{W}_2 .

8.5 Further diagnostic developments

8.5.1 Experimental control

One large technical development of the experimental procedure was the automation, which significantly improved the quality of the experimental results. Mainly it allowed a uniform protocol across an ensemble of runs. Second, it allowed for larger ensembles to be collected while allowing the researcher to spend more time pursuing the analysis and experimental design, thus improving overall productivity. Perhaps most valuable to the quality of the results, however, was that it allowed the experiments to be conducted over night when the laboratory was devoid of occupants and the ventilation system was turned off. This ensured a quieter thermal and mechanical environment.

There are two components that could now be improved with this set up. Currently the control strategy is lacking the coordination of the timing of the release of the vortex rings with the phase of the internal waves. (This lack of coordination arises due to the separate system used to control the wave maker and the asynchronous protocol that is currently used to communicate between the two systems). One small development would be to provide a hardware link between the two systems to allow the relative phase to be controlled accurately. The other component that requires development is the space available on the BeagleBone Blacks (BBB's) used to control ASWaM. Due to the long experimental runs, maximum storage capacity was often reached with the experimental inputs. This could be developed by incorporating into the automatic code a system that uploads each experimental input to the BBB's before it is run then will delete the experimental input after to free space on the BBB's.

8.5.2 Cameras

For the experiments presented in this thesis a 12 MPixel camera based on CMOS (Complementary Metal Oxide Semiconductor) technology was used. While for most purposes, this is the technology of choice as it offers very good sensitivity, dynamic range and speed, it has the failing that typically only 50% of the sensor's surface is sensitive to light. For Synthetic Schlieren, this effectively introduces (on a per-pixel basis) a lower limit as to how far the pattern must appear to move before there is a response on an individual pixel. While the calculation used to determine the apparent movement of the Synthetic Schlieren relies on information from an ensemble of pixels rather than a single one, the limitation of the CMOS sensors nevertheless degrade the sensitivity and signal-to-noise ratio of the resulting measurements.

The technology used more historically for solid state camera sensors, CCD (Charge Coupled Device), does not have this limitation, as the light sensitivity of the entire sensor is nearly uniformly sensitive to light. However, over the last 15 or so years there has been little development of such sensors, and the performance of the 12 MPixel camera currently being used exceeds that of the 1 MPixel CCD sensors of the older generations of cameras. Recently, however, a new generation of high-resolution medium-speed (sufficient for the present Synthetic Schlieren measurements) CCD-based cameras have started to become available.

If a new camera was purchased, the CCD camera could be used for Synthetic Schlieren measurements, freeing up the existing (much faster) 12 MPixel CMOS camera for Particle Image Velocimetry (PIV) measurements, details of which are outlined below.

8.5.3 Particle Image Velocimetry

All of the experimental data presented has obtained using Synthetic Schlieren. This technique is fundamentally two-dimensional, providing a projection of a three-dimensional density field onto the two-dimensional image plane. While three-dimensional information can be extracted using tomographic approaches from a large number of views (e.g. Hazewinkel et al., 2011), there are aspects of the vortex ring collapse that it will not be able to capture adequately. Consequently, there may be value in obtaining velocity measurements of the collapse using Particle Image Velocimetry (PIV). The manner in which the LED light panel can be turned on and off very rapidly (tests suggest reliable performance up to around 200 Hz) under the control of the computer, gives the scope to interleave high-quality PIV and Synthetic Schlieren measurements in a manner similar to that pioneered by Dalziel et al. (2007).

8.6 Concluding remarks

Overall this thesis presents a non-trivial advancement into our understanding of the behavioural evolution of Triadic Resonance Instability in finite-width internal gravity wave beams. Novel experimental results have shown that the approach to a saturated equilibrium state for the triadic waves is not monotonic, rather their amplitudes continue to oscillate without reaching a steady equilibrium. These results have highlighted the importance of properly accounting for the role of beam width, as this has major implications to the structure of the instability and the secondary wave parameters that complete the triad.

These insights into the structural evolution may help with interpretation of oceanographic data. It is clear that the forcing parameters for the beam will impact the ability of the waves to become unstable so knowledge of the generation mechanisms behind the waves is crucial. In addition, the lack of steady equilibrium highlighted experimentally suggests some of the observed ‘wave turbulence’ in the ocean may not only be due to multiple generations of TWI and TRI, but also to the unsteady evolution of a single TRI. The vortex-ring interactions also show strong relevance to ocean scenarios, as the the ocean is never in a state of quiescent equilibrium. Understanding how perturbations to the flow can impact and potentially trigger the instability may help in detecting the presence of TRI in the ocean. The results of these experiments can also be reversed, as the imprint of various perturbations to the wave field may be used to understand the nature of the original disturbance.

While there are many other mechanisms that need to be considered in oceanographic data, understanding the development of freely evolving finite-width internal wave beams in unbounded domains is fundamental. The results presented in this thesis, therefore, make a substantial contribution to this understanding and have the potential to impact our knowledge of a key instability mechanism and its role in global ocean circulation.

Bibliography

- Alford, M. H., MacKinnon, J. A., Zhao, Z., Pinkel, R., Klymak, J., and Peacock, T. Internal waves across the pacific. *Geophysical Research Letters*, 34:2–7, 2007.
- Alford, M. H., Peacock, T., Mackinnon, J. A., Nash, J. D., Buijsman, M. C., Centuroni, L. R., Chao, S. Y., Chang, M. H., Farmer, D. M., Fringer, O. B., Fu, K. H., Gallacher, P. C., Graber, H. C., Helfrich, K. R., Jachec, S. M., Jackson, C. R., Klymak, J. M., Ko, D. S., Jan, S., Johnston, T. M. S., Legg, S., Lee, I. H., Lien, R. C., Mercier, M. J., Moum, J. N., Musgrave, R., Park, J. H., Pickering, A. I., Pinkel, R., Rainville, L., Ramp, S. R., Rudnick, D. L., Sarkar, S., Scotti, A., Simmons, H. L., St Laurent, L. C., Venayagamoorthy, S. K., Wang, Y. H., Wang, J., Yang, Y. J., Paluszkiwicz, T., and Tang, T. Y. The formation and fate of internal waves in the south china sea. *Nature*, 521(7550):65–69, 2015.
- Balmforth, N. J., Ierley, G. R., and Young, W. R. Tidal conversion by subcritical topography. *Journal of Physical Oceanography*, 32(10):2900–2914, 2002.
- Beckebanze, F., Raja, K. J., and Maas, L. R. M. Mean flow generation by three-dimensional nonlinear internal wave beams. *Journal of Fluid Mechanics*, 864:303–326, 2019.
- Beckebanze, F., Grayson, K. M., Maas, L. R., and Dalziel, S. B. Experimental evidence of internal wave attractor signatures hidden in large-amplitude multi-frequency wave fields. *Journal of Fluid Mechanics*, 915, 2021.
- Benielli, D. and Sommeria, J. Excitation and breaking of internal gravity waves by parametric instability. *Journal of Fluid Mechanics*, 374:117–144, 1998.
- Bordes, G., Venaille, A., Joubaud, S., Odier, P., and Dauxois, T. Experimental observation of a strong mean flow induced by internal gravity waves. *Physics of Fluids*, 24, 2012.

- Bourget, B., Dauxois, T., Joubaud, S., and Odier, P. Experimental study of parametric subharmonic instability for internal plane waves. *Journal of Fluid Mechanics*, 723: 1–20, 2013.
- Bourget, B., Scolan, H., Dauxois, T., Le Bars, M., Odier, P., and Joubaud, S. Finite-size effects in parametric subharmonic instability. *Journal of Fluid Mechanics*, 759: 739–750, 2014.
- Briscoe, M. Introduction to collection of papers on oceanic internal waves. *Journal of Geophysical Research*, 80:289–290, 1975.
- Brouzet, C., Ermanyuk, E. V., Joubaud, S., Sibgatullin, I., and Dauxois, T. Energy cascade in internal-wave attractors. *Epl*, 113(44001), 2016.
- Brouzet, C., Ermanyuk, E., Joubaud, S., Pillet, G., and Dauxois, T. Internal wave attractors: Different scenarios of instability. *Journal of Fluid Mechanics*, 811:544–568, 2017.
- Bühler, O. and Holmes-Cerfon, M. Decay of an internal tide due to random topography in the ocean. *Journal of Fluid Mechanics*, 678:271–293, 2011.
- Clark, H. A. and Sutherland, B. R. Generation, propagation, and breaking of an internal wave beam. *Physics of Fluids*, 22(7):1–16, 2010.
- Dalziel, S. B. *DigiFlow User Guide v3.6.0–4.3.0*. Dalziel Research Partners, 2021.
- Dalziel, S. B., Hughes, G. O., and Sutherland, B. R. Synthetic schlieren. *Proceedings of the 8th International Symposium on Flow Visualization*, 1998.
- Dalziel, S. B., Hughes, G. O., and Sutherland, B. R. Whole field density measurements by ‘synthetic schlieren’. *Experiments in Fluids*, 28:322–335, 2000.
- Dalziel, S. B., Carr, M., Sveen, J. K., and Davies, P. A. Simultaneous synthetic schlieren and piv measurements for internal solitary waves. *Measurement Science and Technology*, 18(3):533, 2007.
- Dalziel, S. B., Patterson, M. D., Caulfield, C. P., and Le Brun, S. The structure of low froude number lee waves over an isolated obstacle. *Journal of Fluid Mechanics*, 689: 3–31, 2011.
- Dauxois, T., Joubaud, S., Odier, P., and Venaille, A. Instabilities of internal gravity wave beams. *Annual Review of Fluid Mechanics*, 50:131–156, 2018.

- Davis, G., Jamin, T., Deleuze, J., Joubaud, S., and Dauxois, T. Succession of resonances to achieve internal wave turbulence. *Physical Review Letters*, 124(20):1–5, 2020.
- Davis, R. E. and Acrivos, A. The stability of oscillatory internal waves. *Journal of Fluid Mechanics*, 30(4):723–736, 1967.
- Dewitt, L. M., Levine, M. D., Paulson, C. A., and Burt, W. V. Semidiurnal internal tide in jasin: observations and simulation. *Journal of Geophysical Research*, 91(C2):2581–2592, 1986.
- Didden, N. On the formation of vortex rings: rolling-up and production of circulation. *Journal of Applied Mathematics and Physics (ZAMP)*, 30, 1979.
- Dobra, T. *Nonlinear interactions of internal gravity waves*. PhD thesis, University of Bristol, 2018.
- Dobra, T., Lawrie, A., and Dalziel, S. The magic carpet: an arbitrary spectrum wave maker for internal waves. *Experiments in Fluids*, 2019.
- Down, R. D. and Lehr, J. *Environmental Instrumentation and Analysis Handbook*. John Wiley and Sons, 2005.
- Egbert, G. D. and Ray, R. D. Estimates of m2 tidal energy dissipation from topex/poseidon altimeter data. *Journal of Geophysical Research*, 106:475–502, 2001.
- Fan, B. and Akylas, T. R. Effect of background mean flow on psi of internal wave beams. *Journal of Fluid Mechanics*, 869:R1, 2019.
- Garrett, C. Internal tides and ocean mixing. *Science*, 301(5641):1858–1859, 2003.
- Garrett, C. and Kunze, E. Internal tide generation in the deep ocean. *Annual Review of Fluid Mechanics*, 39:57–87, 2007.
- Garrett, C. and Munk, W. Space-time scales of internal waves. *Geophysical Fluid Dynamics*, 3:225–264, 1972.
- Gharib, M., Rambod, E., and Shariff, K. A universal time scale for vortex ring formation. *Journal of Fluid Mechanics*, 360:121–140, 1998.
- Godunov, S. K. *Different Methods for Shock Waves*. PhD thesis, Moscow State University, 1954.

- Gostiaux, L., Didelle, H., Mercier, S., and Dauxois, T. A novel internal waves generator. *Experiments in Fluids*, 42:123–130, 2007.
- Greenbaum, A. *Iterative Methods for Solving Linear Systems*. SIAM, 1997.
- Grisouarda, N., Leclair, M., Gostiaux, L., and Staquet, C. Large scale energy transfer from an internal gravity wave reflecting on a simple slope. *Procedia IUTAM*, 8: 119–128, 2013.
- Hasselmann, K. A criterion for nonlinear wave stability. *Journal of Fluid Mechanics*, 30(4):737–739, 1967.
- Haynes, W. M., Lide, D. R., and Bruno, T. J. *CRC handbook of chemistry and physics a ready-reference book of chemical and physical data*. CRC Press, 95th edition, 2014.
- Hazewinkel, J. *Attractors in stratified fluids*. PhD thesis, Universiteit Utrecht, 2010.
- Hazewinkel, J., Maas, L. R. M., and Dalziel, S. B. Tomographic reconstruction of internal wave patterns in a paraboloid. *Experiments in fluids*, 50:247–258, 2011.
- Hill, M. J. M. On a spherical vortex. *Philosophical Transactions of the Royal Society. London.*, pages 213–245, 1894.
- Hopfinger, E. J., Flor, J. B., Chomaz, J. M., and Bonneton, P. Internal waves generated by a moving sphere and its wake in a stratified fluid. *Experiments in Fluids*, 11: 255–261, 1991.
- Jackson, B. M. *The dynamics and mixing properties of vortex rings obliquely impacting a density interface*. PhD thesis, University of Cambridge, 2021.
- Johari, H. and Fang, H. Horizontal vortex ring motion in linearly stratified media. *Physics of Fluids*, 9:2605–2616, 1997. doi: 10.1063/1.869377.
- Joubaud, S., Munroe, J., Odier, P., and Dauxois, T. Experimental parametric subharmonic instability in stratified fluids. *Physics of Fluids*, 24, 2012.
- Karimi, H. H. and Akylas, T. R. Parametric subharmonic instability of internal waves: Locally confined beams versus monochromatic wavetrains. *Journal of Fluid Mechanics*, 757:381–402, 2014.
- Kataoka, T. and Akylas, T. On three-dimensional internal gravity wave beams and induced large-scale mean flows. *Journal of Fluid Mechanics*, 769:621–634, 2015.

- Kataoka, T., Ghaemsaïdi, S. J., Holzenberger, N., Peacock, T., and Akylas, T. R. Tilting at wave beams: A new perspective on the St. Andrew's cross. *Journal of Fluid Mechanics*, 830:660–680, 2017.
- Koudella, C. R. and Staquet, C. Instability mechanisms of a two-dimensional progressive internal gravity wave. *Journal of Fluid Mechanics*, 548:165–196, 2006.
- Lamb, H. *Hydrodynamics*. Cambridge University Press, 1932.
- Lighthill, M. *Waves in Fluids*. Cambridge University Press, 1978.
- Longuet-Higgins, M. S. Resonant interactions between two trains of gravity waves. *Journal of Fluid Mechanics*, 12(3):321–332, 1962.
- Longuet-Higgins, M. Longshore currents generated by obliquely incident sea waves, 1. *Journal of Geophysical Research*, 75:6778–6801, 1970.
- Maas, L. and Lam, F. P. A. Geometric focusing of internal waves. *Journal of Fluid Mechanics*, 300:1–41, 1995.
- MacKinnon, J. Mountain waves in the deep ocean. *Nature*, 501:321, 2013.
- MacKinnon, J. A. and Winters, K. B. Subtropical catastrophe: Significant loss of low-mode tidal energy at 28.9 degrees. *Geophysical Research Letters*, 32:1–5, 2005.
- MacKinnon, J. A., Zhao, Z., Whalen, C. B., Waterhouse, A. F., Trossman, D. S., Sun, O. M., St Laurent, L. C., Simmons, H. L., Polzin, K., Pinkel, R., Pickering, A., Norton, N. J., Nash, J. D., Musgrave, R., Merchant, L. M., Melet, A. V., Mater, B., Legg, S., Large, W. G., Kunze, E., Klymak, J. M., Jochum, M., Jayne, S. R., Hallberg, R. W., Griffies, S. M., Danabasoglu, G., Chassignet, E. P., Buijsman, M. C., Bryan, F. O., Briegleb, B. P., Barna, A., Arbic, B. K., Ansong, J. K., and Alford, M. H. Climate process team on internal wave-driven ocean mixing. *Bulletin of the American Meteorological Society*, 98(11):2429–2454, 2017.
- Mathworks. *MATLAB version 9.10.0 (R2021a)*. The MathWorks Inc., 2021.
- Maxworthy, P. The structure and stability of vortex rings. *Journal of Fluid Mechanics*, 51:15–32, 1972.
- Mazumder, S. *Numerical Methods for Partial Differential Equations: Finite Difference and Finite Volume Methods*. Elsevier Science, 2015.

- McEwan, A. D. Degeneration of resonantly-excited standing internal gravity waves. *Journal of Fluid Mechanics*, 50:431–448, 1971.
- McEwan, A. D. and Plumb, R. A. Off-resonant amplification of finite internal wave packets. *Dynamics of Atmospheres and Oceans*, 2:83–105, 1977.
- McEwan, A. D., Mander, D. W., and Smith, R. K. Forced resonant second-order interaction between damped internal waves. *Journal of Fluid Mechanics*, 55(4): 589–608, 1972.
- McGoldrick, L. F. Resonant interactions among capillary-gravity waves. *Journal of Fluid Mechanics*, 21(2):305–331, 1965.
- Mercier, M. J., Garnier, N. B., and Dauxois, T. Reflection and diffraction of internal waves analyzed with the hilbert transform. *Physics of Fluids*, 20, 2008.
- Mercier, M. J., Martinand, D., Mathur, M., Gostiaux, L., Peacock, T., and Dauxois, T. New wave generation. *Journal of Fluid Mechanics*, 657:308–334, 2010.
- Mill, H. R. Fridtjof nansen’s “farthest north”. *Nature*, 55:393–395, 1897.
- Mowbray, D. and Rarity, B. A theoretical and experimental investigation of the phase configuration of internal waves of small amplitude in a density stratified liquid. *Journal of Fluid Mechanics*, 28:1–16, 1967.
- Munk, W. Abyssal recipes. *Deep-Sea Research and Oceanographic Abstracts*, 13: 707–730, 1966.
- Munk, W. and Wunsch, C. Abyssal recipes ii: Energetics of tidal and wind mixing. *Deep-Sea Research Part I: Oceanographic Research Papers*, 45:1977–2010, 1998.
- Nault, J. T. and Sutherland, B. R. Internal wave transmission in nonuniform flows. *Physics of Fluids*, 19, 2007.
- Nikurashin, M. and Ferrari, R. Overturning circulation driven by breaking internal waves in the deep ocean. *Geophysical Research Letters*, 40:3133–3137, 2013.
- Oster, G. and Yamamoto, M. Density gradient techniques. *Chemical Reviews*, 63: 257–268, 1963.
- Phillips, O. M. On the dynamics of unsteady gravity waves of finite amplitude part 1. the elementary interactions. *Journal of Fluid Mechanics*, 9:193–217, 1960.

- Phillips, O. M. Wave interactions – the evolution of an idea. *Journal of Fluid Mechanics*, 106:215–227, 1981.
- Polzin, K. L., Toole, J. M., Ledwell, J. R., and Schmitt, R. W. Spatial variability of turbulent mixing in the abyssal ocean. *Science*, 276:93–96, 1997.
- Riley, N. Steady streaming. *Annual Review of Fluid Mechanics*, 33:43–65, 2001.
- Roache, P. J. Quantification of uncertainty in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 29:123–160, 1997.
- Sarkar, S. and Scotti, A. From topographic internal gravity waves to turbulence. *Annual Review of Fluid Mechanics*, 49:195–220, 2017.
- Scase, M. M. and Dalziel, S. B. Internal wave fields generated by a translating body in a stratified fluid: an experimental comparison. *Journal of Fluid Mechanics*, 564:305–331, 2006a.
- Scase, M. M. and Dalziel, S. B. An experimental study of the bulk properties of vortex rings translating through a stratified fluid. *European Journal of Mechanics B/Fluids*, 25:302–320, 2006b.
- Schmid, P. J. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010.
- Scolan, H., Ermanyuk, E., and Dauxois, T. Nonlinear fate of internal wave attractors. *Physical Review Letters*, 110:1–5, 2013.
- Semin, B., Facchini, G., Pétrélis, F., and Fauve, S. Generation of a mean flow by an internal wave. *Physics of Fluids*, 28(9), 2016.
- Sutherland, B. R. Internal wave instability: Wave-wave versus wave-induced mean flow interactions. *Physics of Fluids*, 18, 2006.
- Sutherland, B. R. *Internal Gravity Waves*. Cambridge University Press, 2010.
- Sutherland, B. R. The wave instability pathway to turbulence. *Focus on Fluids*, 724:1–4, 2013.
- Sutherland, B. R., Dalziel, S. B., Hughes, G. O., and Linden, P. F. Visualization and measurement of internal waves by ‘synthetic schlieren’. part 1. vertically oscillating cylinder. *Journal of Fluid Mechanics*, 390:93–126, 1999.

- Sutherland, B. R., Hughes, G. O., Dalziel, S. B., and Linden, P. F. Internal waves revisited. *Dynamics of Atmospheres and Oceans*, 31:209–232, 2000.
- Tabaei, A. and Akylas, T. R. Nonlinear internal gravity wave beams. *Journal of Fluid Mechanics*, 482:141–161, 2003.
- Tabaei, A., Akylas, T. R., and Lamb, K. G. Nonlinear effects in reflecting and colliding internal wave beams. *Journal of Fluid Mechanics*, 526:217–243, 2005.
- Thomas, L. P., Marino, B. M., and Dalziel, S. B. Synthetic schlieren: Determination of the density gradient generated by internal waves propagating in a stratified fluid. *Journal of Physics: Conference Series*, 166, 2009.
- Thorpe, S. A. On wave interactions in a stratified fluid. *Journal of Fluid Mechanics*, 24:737–751, 1966.
- Thorpe, S. A. On the shape of progressive internal waves. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 263:563–614, 1968.
- Thorpe, S. A. and Haines, A. P. On the reflection of a train of finite-amplitude internal waves from a uniform slope. *Journal of Fluid Mechanics*, 178:279–302, 1986.
- Tsytovich, V. N. *Nonlinear Effects and Plasma Physics*. Springer US, 1970.
- Turner, J. S. *Buoyancy Effects in Fluids*. Cambridge Monographs on Mechanics. Cambridge University Press, 1979.
- Van Atta, C. W. and Hopfinger, E. J. Vortex ring instability and collapse in a stably stratified fluid. *Experiments in Fluids*, 7(3):97–200, 1988.
- Van Leer, B. Towards the ultimate conservative difference scheme. ii. monotonicity and conservation combined in a second-order scheme. *Journal of Computational Physics*, 14(4):361–370, 1974.
- Varma, D. and Mathur, M. Internal wave resonant triads in finite-depth non-uniform stratifications. *Journal of Fluid Mechanics*, 824:286–311, 2017.
- Versteeg, H. K. and Malalasekera, W. *An Introduction to Computation Fluid Dynamics: The finite volume method*. Pearson Education Limited, 2007.
- Wolfram Research, I. *Mathematica, Version 12.3*. Wolfram Research, Inc., 2021.

Wunsch, C. and Ferrari, R. Vertical mixing, energy, and the general circulation of the oceans. *Annual Review of Fluid Mechanics*, 36:281–314, 2004.