

©Copyright 2012

Matthew O. Williams



# Exploiting Low Dimensionality in Nonlinear Optics and Other Physical Systems

Matthew O. Williams

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

Doctor of Philosophy

University of Washington

2012

Reading Committee:

Jose Nathan Kutz, Chair

Robert O'Malley

Bernard Deconinck

Program Authorized to Offer Degree:  
Applied Mathematics



University of Washington

**Abstract**

Exploiting Low Dimensionality in Nonlinear Optics and Other Physical Systems

Matthew O. Williams

Chair of the Supervisory Committee:  
Professor and Chair Jose Nathan Kutz  
Applied Mathematics

Dimensionality reduction techniques have long been used in a number of fields including nonlinear optics and fluid dynamics. Regardless of the specific technique, the underlying idea is to generate a reduced order model that approximates the dynamics of the system but with fewer degrees of freedom. In the field of nonlinear optics, two of the most popular techniques are coupled mode theory and the variational reduction; both of which are based on a solution ansatz. In other fields however, data driven techniques are more common. These techniques extract a set of basis functions from a training set of data and do not require an explicit solution ansatz. In this thesis, we perform a study of the application of dimensionality reduction techniques, both ansatz based and data driven, to problems in nonlinear optics and other physical disciplines. Specifically, we demonstrate that for pattern forming systems, which are ubiquitous in optics, reduced order models can be used to quickly and accurately compute solution branches, even for bifurcation sequences as complicated as the multi-pulsing transition in a mode-locked laser or the route to chaos in an unstable semiconductor waveguide array laser. We also show that a significant degree of computational savings can be obtained while evolving a system in time by exploiting a low dimensional representation of the system when possible. Although we focus on nonlinear optics in this thesis, the techniques outlined here can be applied to any pattern forming system whose dynamics are essentially low-dimensional. To demonstrate this, we also apply these techniques to surface water waves to obtain periodic solution branches in a physical context other than optics.



## TABLE OF CONTENTS

	Page
List of Figures . . . . .	iv
Chapter 1: Introduction . . . . .	1
1.1 Motivation . . . . .	1
1.2 The Proper Orthogonal Decomposition . . . . .	1
1.3 Pattern Formation and Mode-Locking in Nonlinear Waveguide Arrays . . . . .	3
1.4 Outline of the Thesis . . . . .	5
Chapter 2: Pattern Formation in Nonlinear Optics . . . . .	7
2.1 Nonlinear X-Waves . . . . .	8
2.1.1 Co-Propagating X-Waves . . . . .	8
2.1.2 Counter-Propagating Waves . . . . .	12
2.2 Light-Bullet Formation, Routing, and Control . . . . .	15
2.2.1 System Description . . . . .	15
2.2.2 Governing Equations . . . . .	16
2.2.3 Uniform Gain Dynamics . . . . .	17
2.2.4 Non-Uniform Gain Dynamics . . . . .	20
2.2.5 Gain Mediated Interactions . . . . .	24
2.3 Chapter Summary . . . . .	27
Chapter 3: Ansatz Based Reduction Methods in Nonlinear Optics . . . . .	29
3.1 Experimental Setup . . . . .	30
3.2 Governing Equations . . . . .	31
3.3 Lasing Dynamics . . . . .	32
3.4 Reduced Dimensional Model . . . . .	37
3.5 Chapter Summary . . . . .	43
Chapter 4: The Multi-Pulsing Transition . . . . .	44
4.1 The Multi-Pulsing Transition with Reduced Order Models . . . . .	44

4.1.1	Prototypical Dynamics in the Multi-Pulsing Transition . . . . .	45
4.1.2	The Low Dimensional Model . . . . .	46
4.1.3	Low Dimensional Dynamics . . . . .	52
4.2	Verifying the Multi-Pulsing Transition . . . . .	56
4.2.1	Period-One Breather Solutions . . . . .	57
4.2.2	Torus Bifurcation of Period-One Breathers . . . . .	63
4.2.3	Period-Two Breather Solutions . . . . .	65
4.2.4	Global Bifurcation Structure . . . . .	67
4.3	Chapter Summary . . . . .	70
Chapter 5:	Hybrid Integrators Based on Reduced Order Models . . . . .	73
5.1	The Proper Orthogonal Decomposition and Dynamic Mode Decomposition . . . . .	75
5.1.1	The Proper Orthogonal Decomposition . . . . .	75
5.1.2	Dynamic Mode Decomposition . . . . .	76
5.1.3	Application to the Nonlinear Schrödinger Equation . . . . .	78
5.2	Hybrid Integration with Reduced Order Models . . . . .	80
5.2.1	A Metric for Detecting Low Dimensionality . . . . .	81
5.2.2	Metrics for Confirming Low Dimensionality . . . . .	82
5.3	Cubic Quintic Ginzburg Landau Equation . . . . .	84
5.3.1	Transitions Between Stationary Solutions . . . . .	85
5.3.2	Transitions Between Stationary and Periodic Solutions . . . . .	87
5.3.3	Comparison of Computational Cost . . . . .	87
5.4	Adiabatic Parameter Incrementation in the Waveguide Array Mode Locking Model . . . . .	90
5.5	Chapter Summary and Future Outlook . . . . .	93
Chapter 6:	Dimensionality Reduction in Standing Surface Water Waves . . . . .	95
6.1	Governing Equations . . . . .	96
6.2	Extensions of the Proper Orthogonal Decomposition . . . . .	98
6.3	Infinite Depth . . . . .	101
6.3.1	Problem Formulation . . . . .	101
6.3.2	POD Modes . . . . .	103
6.3.3	POD Dynamics . . . . .	105
6.3.4	Bifurcation Diagram . . . . .	108
6.4	Finite Depth . . . . .	109
6.4.1	Problem Formulation . . . . .	109

6.4.2	Proper Orthogonal Decomposition Modes . . . . .	111
6.4.3	POD Dynamics . . . . .	114
6.4.4	Bifurcation Diagram . . . . .	115
6.5	Chapter Summary . . . . .	117
Chapter 7:	Concluding Remarks . . . . .	120
Appendix A:	Adjoint Continuation Method . . . . .	122
A.1	Description of the Adjoint Continuation Method for the WGAML . . . . .	122
A.2	The Adjoint PDE of the WGAML . . . . .	125
A.3	Computing the stability of periodic solutions . . . . .	126
A.4	Semi-implicit Runge-Kutta Method . . . . .	127
Bibliography	. . . . .	130

## LIST OF FIGURES

Figure Number	Page	
1.1	Two possible laser cavity configurations that include nonlinear mode-coupling from the waveguide array as the mode-locking element. The fiber coupling in and out of the waveguide array occurs at the central waveguide as illustrated. Any electromagnetic field that is propagated into the neighboring waveguides is ejected (attenuated) from the laser cavity. In addition to the basic setup, polarization controllers, isolators, and other stabilization mechanisms may be useful or required for successful operation. . . . .	4
2.1	Schematic of a slab waveguide array theoretically capable of producing light-bullets. Slab waveguides (dark gray) are separated by non-guiding regions (light gray) to produce a weak coupling between adjacent arrays. Gain is applied to the topmost waveguide via current injection. . . . .	8
2.2	Pseudo-color plot of in-phase pulses interacting. The final result is a pair of closely-spaced but distinct X-waves. The white and green dotted lines indicate the center of the X-wave at $Z = 0$ in the 0th and 1st waveguide respectively. The white and green dashed lines represent the center of the X-wave at $Z = 3$ mm. Note that the final separation is slightly larger than the initial separation of the pulses. . . . .	9
2.3	Pseudo-color plot of out of-phase pulses interacting. In this case, the X-waves attract and form a pair of X-waves with a negligible delay in time. The white and green dotted lines denote the center of X-wave in the 0th and 1st waveguides at $Z = 0$ , and the dashed lines represent the center of the X-wave at the present value of $Z$ . . . . .	10
2.4	Plot of the X-wave separation as a function of phase-difference for three different initial separations. The dashed lines show the initial separation and the solid lines of the same color show the final X-wave separation. Regardless of the initial separation, X-waves with small phase-differences repel and X-waves with phase-differences near $\pi$ attract. . . . .	10
2.5	Plot of the final X-wave separation for a pair of in-phase X-waves, shown in black, and $\pi$ out-of-phase X-waves, shown in blue with an initial separation of $\Delta T = 1$ . For sufficiently high initial powers, the in-phase solutions repel and the out-of-phase solutions attract. . . . .	11

2.6	Pseudo-color plot of two unequally sized X-waves with $\Delta\theta = 0$ . The X-wave in waveguide 0 has amplitude 2.0 while the pulse in waveguide 1 has amplitude 1. The larger pulse dominates the dynamics and the attracts the smaller pulse.	11
2.7	Pseudo-color plot of the collision of counter-propagating X-waves. The top series of plots shows the propagation of both the forward and backward X-waves at three snapshots in time. The bottom series of plots follows the propagation of the forward X-wave as it interacts with the backward X-wave.	12
2.8	Plot of the phase shift in radians and time delay of the interacting X-waves. The magnitude of the phase shift is dependent upon the size of the X-wave, but the time-delay remains zero for all initial conditions.	13
2.9	Plot the phase shift of counter-propagating X-waves as a function of initial separation. The different colors correspond to initial conditions with peak pulse amplitudes of 1.50, 1.25, 1.00, 0.75, and 0.50 for the blue, green, red, teal, and purple lines respectively.	14
2.10	Example of a radially symmetric stationary solution with the parameters from Eq. (2.7). The initial condition is noise in the zeroth waveguide and zero in the other two waveguides. The radially symmetric bullet solution forms from noise and appears to be a global attractor of the system. Note that the outer two waveguides also possess rotational symmetry and have qualitatively similar stationary solutions.	18
2.11	The bifurcation diagram and spectrum of selected radially symmetric bullet solutions. In the bifurcation diagram, the blue curve is where the bullet is linearly stable while the red regions are where it is not. In the spectra, red $x$ 's represent eigenvalues with positive real parts. Light-bullet solutions are born out of a saddle node bifurcation near point 1 and go unstable near point 3 due to a Hopf bifurcation.	19
2.12	The amplitude in the zeroth waveguide as a function of time for a time-periodic breather solutions using the parameters in Eq. (2.8). At higher gains, the stationary solution is unstable and the breather solution emerges.	19
2.13	Time evolution of the splitting of a single light-bullet into a pair of light-bullets. Due to the larger gain, the single pulse sheds energy which eventually results in the formation of a second pulse. Note that the resulting bullets are far separated in space. If the second bullet forms in close proximity, it will interact with the original bullet and be absorbed which repeats the splitting process.	20
2.14	Plot of the velocity of the bullet as a function of the slope of the gain profile. The velocity of the bullet is directly and linearly related to the slope of the gain.	21

2.15	Bullet routing using the gain equation in Eq. (2.9). The dotted lines show the location of the bullet center as time progresses. The piecewise linear gain routes the bullet through the junction. . . . .	22
2.16	Bullet control using time dependent gain. The + and o show the $x$ and $y$ coordinates of the bullet center as a function of time. The gray lines show the $x$ and $y$ coordinates of the gain function in Eq. (2.10) as a function of time.	23
2.17	A plot of all possible cases for the NOR gate. Both the clock and input bullets experience a sloped gain which routes all bullets from left to right. The gain is biased so the clock receives less gain than either of the input bullets. The system possesses enough gain to support a single bullet so the clock bullet will be destroyed if either input bullet exists. If neither exist, the clock bullet translates right giving a high output. . . . .	25
2.18	Four possible cases for the NAND gate. Note that each of the four possible bullets experiences a time periodic gain similar to Eq. (2.10) that translates the bullet from left to right. However, the overall gain given to each bullet, from lowest to highest, is: the auxiliary bullet, the clock bullet, the input one bullet, and the input two bullet. The NAND gate has enough gain to support a pair of bullets so if at most one of the inputs bullets exists the output is high. Otherwise, the output is low. The four cases here correspond directly to the four cases in the logic table of the NAND gate in Table 2.1 . . . . .	26
3.1	A schematic figure of the waveguide array laser. The waveguide array contains five waveguides numbered by the proximity to the source of the pump current.	30
3.2	The transverse structure of the waveguide array. This is a five-region waveguide that contains, substrate, AlGaAs cladding, and the active region with the quantum dots. . . . .	30
3.3	Experimental and numerical results measuring the time evolution of the output power for a five-emitter array with a pump current of 400 mA from waveguides 1-5 in (a)-(e) respectively. The power output is asymmetrical, which implies an uneven pump current into each waveguide. The frequency and magnitude of the numerical results (dashed) are similar to experiment (solid), with a slight frequency shift. . . . .	34
3.4	Average output power as a function of pump current in each of the waveguides obtained experimentally (a) and numerically (b). The gain in power is greatest in the first waveguide (+) and is reduced in the second (o), third (●) and fourth ( $\Delta$ ) waveguides. Experimentally, the fifth ( $\triangleright$ ) has a high gain at lower currents and a lower gain at higher currents, but this is not reflected in numerical simulation. Note that despite the current injection ramp, there is a uniform threshold across the waveguides at $\approx 380$ mA. . . . .	36

3.5	Measured phase difference in power between waveguides for pump currents of 395 mA (●), 400 mA (+), 405 mA (×), 410 mA (○), 420 mA (△), 430 mA (▷), and 440 mA (◁). The phase change stays fixed near $\pi$ for all currents except between the fourth and fifth waveguides where there is variation. . . .	37
3.6	Experimental frequency spectrum of the power oscillations vs injection current. The fundamental frequency is directly related to the pump current and appears for all currents larger than 380 mA. . . . .	37
3.7	Frequency spectrum of the power oscillations versus injection current obtained from the simulation. The same pattern as in Fig. 3.6 appears with the fundamental frequency and its harmonics increasing with pump current. Below approximately 380 mA, there is no oscillation and thus no fundamental frequency. . . . .	38
3.8	Plot of the power in waveguides 1-5, shown in (a)-(e) respectively, using the reduced dimensional model. The same parameters and 400 mA injection current as Fig. 3.3 were used except the waveguide separation was taken to be 9 $\mu\text{m}$ . . . . .	40
3.9	Plot of the local maxima and minima of pulse amplitudes for the five-waveguide system with a gain slope of (a) 0 (b) 0.06 (c) 0.12 (d) 0.18 for different waveguide separations. The route to chaos persists, but the location of the Hopf and period-doubling bifurcation changes. . . . .	41
3.10	Plot of the bifurcation locations as a function of waveguide separation for the five waveguide system. The solid line is the location of the Hopf bifurcation while the dashed line is the location of the bifurcation of the limit cycle which is either a period doubling bifurcation, a Neimark-Sacker bifurcation, or a fold of the limit cycle. . . . .	42
4.1	Surface and pseudo-color plots of possible behaviors in the WGAML system. From left to right, a single pulse, a breather, a chaotic, and a two pulse solution are shown. The particular behaviors are determined by the value of $g_0$ which is 2.3, 2.5, 2.6, and 2.7 for the single-pulse, breather, chaotic, and two pulse solutions respectively. . . . .	45
4.2	Plots of the POD mode profile and the error as defined in Eq. (4.5) using the first $N$ POD modes in a low dimensional model. Left: the POD modes of the single-pulse solutions for $g_0 = 2.3$ . Right: the POD modes of the breather solutions for $g_0 = 2.5$ . For the breather solutions, the error was measured at the point of maximum amplitude for both the PDE and ODE solutions due to differences in the period of the limit cycles. . . . .	48

4.3	Plots of POD mode profiles for reconstructing the chaotic and double-pulse solutions of the WGAML. Left: the first six of the twenty-two modes of the chaotic transition needed to capture 99.9% of the energy at $g_0 = 2.6$ ; Right: the first six modes of the double-pulse solution at $g_0 = 2.7$ . Only two of these modes are needed to capture 99.9% of the energy. . . . .	50
4.4	Comparison of the number of POD modes needed to capture 99.9% of the energy for the even-WGAML model and the unrestricted-WGAML model using data taken for $0 \leq Z \leq Z_{\max}$ . The gain used for these results is $g_0 = 2.65$ and $g_0 = 3.1$ respectively. . . . .	51
4.5	POD modes taken from the combined data set for the 0th waveguide that are capable of qualitatively reproducing the dynamics observed in the full WGAML model. . . . .	52
4.6	Top: Surface and pseudo-color plots of the single-pulse, breather, chaotic, and double pulse solutions computed for the finite-dimensional model at $g_0 = 1.5, 2.5, 3.495,$ and $3.5$ respectively. Bottom: The same plots for the even WGAML model taken at $g_0 = 2.3, 2.5, 3.1,$ and $3.2$ respectively. The reduced model accurately reproduces the four behaviors observed. These solutions should be compared to their equivalent behavior in the full WGAML model shown in the same order in Fig. 4.1. . . . .	53
4.7	Top: The bifurcation diagram of energy ( $L^2$ norm) vs $g_0$ for the multi-pulse transition in the POD model. Bottom: The bifurcation diagram of the multi-pulse transition of the WGAML model. The different plots show the same diagram with emphasis on different regions of the transition. For the stationary solutions, emphasized on the left, linearly unstable regions are dashed red lines while linearly stable regions are solid blue lines. Stationary solutions that were not computed explicitly are denoted by dots. For the periodic solutions, the local extrema of the energy are denoted by $x$ 's . . . . .	54
4.8	From top to bottom, plots of the amplitude of the 0th, 1st, and 2nd waveguides over one period at $g_0 = 2.5$ . The solid black line shows the initial condition, and the dashed black line shows the solution after a half-period has elapsed. The green lines show the solutions at evenly-spaced intermediate values of $Z$ . Note that the actual domain is from $T \in [-20, 20)$ and only $T \in [-1, 1]$ was plotted to highlight the region of interest. . . . .	58

- 4.9 (Left) A two-dimensional version of the bifurcation diagram showing the maximum  $L_2$ -norm over a period vs the gain. Stable solutions are plotted in blue and unstable solutions in red. The labeled bifurcation points are shared between the 2D and 3D plots. (Right) The bifurcation diagram of the period-one breather solutions with individual orbits are plotted in the  $(g_0, L_2, H_1)$  coordinate system. The plotted orbits represent one out of every forty periodic orbits computed on the branch. Linearly stable orbits are shown in blue and linearly unstable orbits are shown in red. The solutions at points (1) and (2) are shown in Fig. 4.11 while the solution at (3) is shown in Fig. 4.8. . . . . . 59
- 4.10 On the left, the real and imaginary parts of the bifurcation eigenfunctions of the 0th waveguide are shown in blue and red respectively. These modes are associated with a period doubling bifurcation indicated by a Floquet multiplier crossing the unit circle at  $\mu_{pd} = -1$  in the right figure. Multipliers in or on the unit circle are indicated with a blue dot. Multipliers outside of the unit circle are indicated with a red  $x$ . . . . . 61
- 4.11 Plot of the breather solutions at the points labeled (1) and (2) in Figure 4.9 on the left and right respectively. The solid black line shows the initial condition and the dashed black line shows the solution after a half-period has elapsed. The gray (green) lines show the evolution of the solution in evenly spaced intervals. The left set of plots was taken directly after the Hopf bifurcation, and the right set of plots was taken from the unstable region after the first fold of the limit cycle. Note that the actual domain is from  $T \in [-20, 20)$  and only  $T \in [-1, 1]$  was plotted to highlight the region of interest. . . . . 62
- 4.12 (Top) Bifurcation diagram including the branches of stationary one-pulse, stationary two-pulse, period-one breather, and period-two breather solutions. Branches in blue or green are linearly stable while branches in red are linearly unstable. Branches in solid lines are from stationary (constant amplitude) solutions while branches in dashed lines are  $z$ -periodic solutions. The green dashed lines represent period-two breathers and the blue lines period-one. Hopf, saddle-node, and period-doubling bifurcations are denoted by  $H$ ,  $SN$ , and  $PD$  respectively. A fourth unknown bifurcation is indicated by  $B$ . (Bottom) Examples of the four qualitatively different solution behaviors – stationary one-pulse (single-pulse) solutions, period-one breathers, period-two breathers, and stationary two-pulse (double-pulse) solutions – observed during the multi-pulsing transition. The stationary two-pulse solutions can be treated as two non-interacting stationary one-pulse solutions. . . . . 63

- 4.13 On the left is the plot of the amplitude of the breather solution at  $g_0 = 2.85$  in waveguide 0, and on the right are the Floquet multipliers. Floquet multipliers outside the unit circle are denoted with a grey (red)  $x$ , and multipliers inside or on the unit circle are denoted with a black (blue) dot. There are multipliers outside the unit circle:  $\mu_{\text{pd}} = -5.825$  and  $\mu_{\text{ns}} = 0.1217 \pm 0.9984i$ .  $\mu_{\text{pd}}$  is associated with an odd eigenfunction and has remained outside the unit circle since the period-doubling bifurcation. The additional pair,  $\mu_{\text{ns}}$  corresponds to a torus (Neimark-Sacker) bifurcation that occurs near  $g_0 = 2.8$ . . . . . 64
- 4.14 Example of a stable, period-doubled solution. The symmetry breaking of the period-doubling bifurcation has generated a double-peak structure in all three waveguides. The solid black line shows the initial condition and the dashed black line shows the solution after a half-period has elapsed. The first half-period is shown in blue and the second half-period is shown in green. Note that the actual domain is from  $t \in [-20, 20)$  and only  $t \in [-1, 1]$  was plotted to highlight the region of interest. . . . . 65
- 4.15 (Left) A two-dimensional version of the bifurcation diagram showing the maximum  $L_2$  value over a period versus the gain. Stable solutions are plotted in blue and unstable solutions in red. The labeled bifurcation points are shared between the 2D and 3D plots. (Right) The bifurcation diagram of the period-two solutions to the WGAML. Individual orbits are plotted in  $(g_0, L_2, H_1)$  space and each plotted orbit represents one of every thirty solutions computed on the branch. Linearly stable solutions are plotted with blue edges while unstable orbits are plotted with red edges. . . . . 66
- 4.16 Plots of solutions at the three selected points indicated by (a), (b), and (c) in Fig. 4.15. The solid-black and dashed-black lines show the intensity in the 0th waveguide at  $Z = 0$  and  $z = Z_0/2$ . Intermediate times are plotted in gray (the first half-period is shown in blue and the second half-period is shown in green). Note that the actual domain is from  $T \in [-20, 20)$  and only  $T \in [-1, 1]$  was plotted to highlight the region of interest. . . . . 67
- 4.17 Bifurcation diagram of the stationary and breather solutions of the WGAML. The stationary solutions are indicated by thick lines. Solid blue lines indicate spectrally stable solutions while dotted red lines represent solutions that are spectrally unstable. The period-one limit cycles are shown in green and the period-two in light-blue. The gray plane at  $g_0 = 2.527$  represents the largest value of  $g_0$  for which any of the limit cycles are stable. The gray shaded region indicates the interior of the limit cycle, and it should be noted that the stable two-pulse solution remains within the interior of the limit cycle near the transition value of  $g_0$ . . . . . 69

4.18	Plot of the trajectories of the solutions (in terms of $L_2$ and $H_1$ ) for $g_0 = 2.3, 2.5$ and $2.6$ starting with low-amplitude initial conditions. The long time behavior of the trajectories are shown in blue. The light gray lines show the complete evolution of the solutions from the initial conditions. The light blue lines show the stable single- and double-pulse solutions, while the red lines show the unstable single-pulse solutions. The dark gray lines show other limit cycles at the same value of $g_0$ , such as the period-one breathers (green) and period-two breathers (light-blue). . . . .	70
4.19	Plot of the trajectory of the solution starting from low-amplitude white noise for $g_0 = 2.3, 2.5,$ and $2.6,$ respectively. The solution trajectory is shown in light gray. The green and light-blue lines show the period-one and period-two limit cycles for values of $g_0$ where they exist. The blue points or curves denote the long time behavior of the trajectory. At $g_0 = 2.3,$ the stationary single-pulse solution is the attractor. At $g_0 = 2.5,$ the period-one breather solution is the limiting solution. At $g_0 = 2.6,$ the chaotic translating double-pulse solution gives the long time behavior. . . . .	71
5.1	(a) A solution of the Cubic Quintic Ginzburg Landau equation that, due to time-dependence of the loss, undergoes a Hopf bifurcation from a stationary solution to a time-periodic solution and back. (b) Results from a reduced order POD model that samples data as the solution converges to the stationary solution. Although the static POD model can represent the stationary solution accurately, the ROM is irrelevant after the bifurcation and leads to erroneous results. Regions in red denote when data collection occurs. (c) Results from the hybrid POD/DMD model in this manuscript. By adaptively re-sampling from the PDE, the ROM updates itself to produce accurate results. . . . .	74
5.2	Amplitude plots of the two soliton solution to NLS generated by: (a) numerically solving Eq. (5.9), (b) 3-mode POD representation (Eq. (5.10) with $P = 3$ ), (c) 20-mode DMD representation (Eq. (5.7) with $K = 20$ ). Plots (d), (e), and (f) show the relative error (in the $L^2$ -sense) between the exact solution and the Fourier, POD, and DMD representations of the solution. For reference, the green lines in all three plots show the error of the Fourier representation. The blue lines show the error of the POD and DMD representations in the second and third columns. Because the POD and DMD modes are generated from Fourier data, the accuracy of those methods cannot exceed that of the Fourier representation. In this case, however, both the POD and DMD representations have a relative error within $10^{-5}$ of that of the Fourier basis. Plots (g), (h), and (i) show the ten Fourier modes with largest mean amplitude, the POD mode amplitudes, and the DMD mode amplitudes. . . . .	79

5.3 Results from the PDE and hybrid integrators for the parameter values in Eq. (5.19). As shown by the PDE solution in (a), this system transitions from the initial condition to a triple-pulse solution and then to two different single-pulse solutions. The black lines in (d) and (e) are plots of the energy of the PDE solution. The plots in (b) and (d) show the amplitude and energy of the solution obtained using the hybrid integrator with  $R_D$  as the metric for returning to the full model. In both plots, the red portions denote the times when PDE integrators were used. Note that with the  $R_D$  metric, the hybrid integrator resamples after every change in the system parameters. The plots in (c) and (e) show the same information but using  $R_P$  as the criterion for using the reduced model. Unlike the  $R_D$  metric, the  $R_P$  metric does not always resample when the parameters change. Overall, there is good quantitative agreement between the full PDE and both hybrid integrators. . . . . 86

5.4 Results from the PDE and hybrid integrators for the parameter values in Eq. (5.20). In this case, the system transitions from a hyperbolic secant initial condition, to a non-stiff single-pulse solution, to a time-periodic breather solution, and back to the original non-stiff single-pulse solution. Plots (a), (b), and (c) show the solution amplitude as a function of time and space for the PDE, hybrid integrator with  $R_D$  metric, and hybrid integrator with  $R_P$  metric. Solutions in red denote snapshots where the full PDE integrator was used, and solutions in black were obtained during times the POD integrator was used. Plots (d) and (e) show the energy of the solution for the PDE, shown in black for both plots, and the two hybrid integrators which are shown in blue and green for  $R_D$  and  $R_P$ . With either metric, the hybrid integrator agrees quantitatively with the full PDE, but the choice of metric does impact exactly when the reduced or full integrator is called. . . . . 88

5.5	Plot of the solution amplitude in the 0th waveguide with gain increased quasi-adiabatically (Eq. (5.24)) as computed by (a) <code>ode23</code> (b) the hybrid integrator with the $R_P$ metric, and (c) the hybrid integrator with the $R_D$ metric. The system transitions from the stationary solution, to the breather solution, and finally to the chaotic solution. The time spent in the chaotic solution depends primarily upon the initial condition when the system enters the chaotic regime. Due to this sensitive dependence on the initial condition, the hybrid integrators are unable to accurately reproduce the dynamics in this regime. Plots (d) and (e) show the energy of the solution as a function of time with the energy of the PDE solution given in black. Because of the chaotic nature of the solution, there is a visible difference between the full and hybrid integrators, but there is still good qualitative agreement. As implied by the red background that denotes where the full PDE integrator was used, the majority of the computational speedup comes from the use of the reduced model for the stationary and breather solutions and there is almost no opportunity for a speedup in the chaotic regime. . . . .	92
6.1	A sketch of the water wave problem at time $t$ with mean fluid depth $h$ . The shape of the fluid surface is described by $\eta(x, t)$ . The velocity potential in the fluid is given by $\Phi(x, y, t)$ , and the velocity potential on the surface is $\varphi(x, t) = \Phi(x, \eta(x, t), t)$ . All variables are $2\pi$ -periodic in $x$ . . . . .	96
6.2	(Top) The bifurcation diagram in terms of period and $L^2$ -norm of the standing wave solutions in fluid of infinite depth. Locations marked by an $x$ are where data was used to generate POD modes. (Bottom) Plot of the first ten POD modes for the surface elevation and the dipole density. . . . .	104
6.3	Convergence of the POD modes (blue line) and the Fourier modes (black line) for sharply peaked solution data. The left plot shows the convergence for the fluid surface height, and the right plot shows the convergence for the dipole density in the fluid. Here the POD modes converge much more rapidly than the Fourier modes due to the large wavenumber components of the data. . . . .	105
6.4	Plot of the error in the derivative of a 10-mode Fourier expansion (black-dashed), 10-mode POD expansion (blue-dashed), 22-mode Fourier expansion (black-solid), and 22-mode POD expansion (blue-solid) for the solution plotted on the right. The POD expansions are roughly one to two orders of magnitude more accurate than the equivalently sized Fourier representation at all points along the orbit. . . . .	106

6.5	Plot of the POD mode amplitudes, which are proportional to the amount of $L^2$ -norm each mode captures, for the evolution in Figure 6.4. The oscillatory nature of the mode amplitudes demonstrates that, although the energy in a particular POD mode decays exponentially with the mode number, the amount of $L^2$ -norm each mode captures at a given time does not obey that exponential relationship. As an example, the first POD mode for $\eta$ vanishes at $t = 3.265$ while the fourth mode does not. . . . .	107
6.6	Bifurcation diagram of the time-periodic standing wave solution plotted in terms of the period, $T$ , and the maximum $L^2$ norm squared of the derivative of $\eta$ over a single period. The blue curve is the bifurcation diagram obtained with 10-POD modes. The red curve is the bifurcation diagram obtained with 22-POD modes. The black curve is the bifurcation diagram of the full PDE. With more than 10-POD modes, the branch of solutions is qualitatively reproduced. However, greater quantitative accuracy is obtained by including additional modes. . . . .	108
6.7	The physical domain $\Omega(t) = \{(x, y) : 0 \leq x < 2\pi, 0 < y < \eta(x, t)\}$ is transformed to $\hat{\Omega} = [0, 2\pi) \times [0, 1]$ under the change of variables $\hat{x} = x$ , $\hat{y} = y/\eta(x, t)$ . The problem is reformulated in terms of $\hat{\Phi}(x, y) = \Phi(x, \eta(x, t)y) - \varphi(x)$ , which gives the boundary conditions $\partial_y \hat{\Phi}(x, y) _{y=0} = 0$ and $\hat{\Phi}(x, 1) = 0$ . . . . .	110
6.8	(Top) Bifurcation diagram showing $L^2$ -norm vs. period of the standing wave solutions in fluid of unit depth. The data used to generate POD modes were taken from locations marked by a red +. Near each data point, we plot the surface height evolution over a period. (Bottom) First ten POD modes for the surface elevation and the velocity potential on the surface. Note that there is good agreement between the POD and PDE results even beyond the regimes where data was taken. This demonstrates the use of the POD method in a predictive fashion, i.e. extrapolating beyond the region where data were obtained. . . . .	112
6.9	Plots of the POD modes of $\hat{\Phi}(x, y)$ on the transformed rectangular domain. The modes are ordered from left-to-right and from top-to-bottom. . . . .	113
6.10	Convergence of $L^2$ -norm of the POD representation (blue line) and the Fourier representation (black line) of the rightmost data point (denoted by a red cross, with $T \approx 7.2075$ ) in Fig. 6.8. The POD modes were generated using data from the four solutions indicated in that figure. The tolerance set for the SPOD technique here was $10^{-12}$ . The POD representation is nearly identical to the Fourier representation for the first two modes; as a result, the normalized error is also identical. However, the convergence of the POD representation increases for later modes which are drawn from data sets more similar to the one of interest. . . . .	114

6.11	Plot of the evolution of the fluid surface using (a) 6-POD modes (b) 10-POD modes and (c) 512 Fourier modes. With the 6-mode expansion, the surface evolution can be qualitatively reproduced although there are significant quantitative differences. However, with 10-POD modes there is good quantitative agreement between the low dimensional model and the full 512-Fourier mode PDE solution . . . . .	115
6.12	Plot of the POD mode amplitudes for the evolution in Fig. 6.11, which demonstrates that the maximum value of the mode amplitude decreases as the mode number and frequency of oscillation increase. . . . .	116
6.13	Bifurcation diagrams of the time-periodic standing wave solutions in terms of the period, $T$ , and the maximum $L^2$ -norm over one period. The PDE solution branch is shown in solid black. The 6-mode POD expansion is shown in dashed red lines, and the 10-mode POD expansion is shown in dashed green lines. . . . .	117

## ACKNOWLEDGMENTS

The author wishes to express his sincere appreciation to his advisor J. Nathan Kutz for his support and advice over the past few years. He would also like to acknowledge Steven Cundiff, Edwin Ding, Mingming Feng, Colin McGrath, Peter Schmid, Eli Shlizerman, and Jon Wilkening for their support and advice on all of the projects contained within this thesis.

## Chapter 1

# INTRODUCTION

### **1.1 Motivation**

In many physical disciplines including, but not limited to, nonlinear optics and fluid dynamics, there are many questions that can be answered using techniques from dynamical systems. In particular, the location and stability of different solution branches is often critical when attempting to design or control these types of systems. However, a common difficulty faced by researchers attempting to apply the techniques of dynamical systems to these physical applications is that the underlying governing equations are nonlinear systems of partial differential equations which makes them extremely computationally expensive to study. Therefore, although the standard numerical tools for branch tracking and bifurcation study, such as MATCONT [1] or AUTO [2], can still be used in theory, in practice the time required to execute such codes is too long to be viable.

In this thesis, dimensionality reduction techniques are used to reduce the computational cost of the problem so that standard numerical tools and techniques of applied mathematics such as pseudo-arclength continuation or numerical integrators once again become tractable. The idea of dimensionality reduction is old and well studied. However, the field can be partitioned, roughly, into two groups: dimensionality reduction techniques that assume knowledge about the underlying flow operator and reduction techniques that assume a form of the solution. Over the course of this thesis, both types of techniques will be used to compute bifurcation diagrams and lower the computational cost of evolving systems in time. In contrast to previous works and rather than relying solely upon intuition about the problem, we use data-mining methods to algorithmically extract from a given set of data either an appropriate ansatz for the solution or an effective approximation of the underlying flow operator [3, 4]. Although the many potentially appropriate data-mining methods for dimension reduction differ algorithmically, the underlying goal of each is to identify and parametrize the assumed low dimensional manifold where the dynamics of the system lies.

### **1.2 The Proper Orthogonal Decomposition**

In this section, we provide a short introduction to the most commonly used numerical dimensionality reduction technique in this thesis, the Proper Orthogonal Decomposition (POD). The POD is nearly identical to other techniques such as Principal Component Analysis (PCA), the Hotelling transform, or the Karhunen-Loève Transform (KLT) [4]. Indeed, the difference between these techniques is often a single step such as taking a Fourier Transform of the data or subtracting out the mean. Regardless of how the data is preprocessed, each of these methods is ultimately related to the Singular Value Decomposition (SVD). Although other methods such as the direct method, see Ref. [4], or the method of snapshots, see

Ref. [5], exist and are even commonly used, the SVD is convenient computationally because it is implemented as a built-in routine in many scientific software packages such as MATLAB and NumPy. To generate a complete set of POD modes, a data set, represented here as  $\mathbf{X} \in \mathbb{C}^{N \times M}$ , is compiled. Each column of the matrix is a single snapshot of the PDE solution recorded at a single time with  $N$  spatial degrees of freedom. This is repeated  $M$  times to obtain the complete data set. The SVD factorizes the matrix  $\mathbf{X}$  into the product of three matrices

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*, \quad (1.1)$$

where  $\mathbf{U} \in \mathbb{C}^{N \times N}$ ,  $\mathbf{V} \in \mathbb{C}^{M \times M}$  and  $\mathbf{\Sigma} \in \mathbb{R}^{N \times M}$ . Here, the asterisk denotes the conjugate transpose. In a matrix form, the factorization in Eq. (1.1) is expressed as

$$\mathbf{X} = \begin{bmatrix} \vec{\phi}_1 & \vec{\phi}_2 & \cdots & \vec{\phi}_{N-1} & \vec{\phi}_N \end{bmatrix} \begin{bmatrix} \sigma_1 & & & & \mathbf{O} \\ & \ddots & & & \vdots \\ & & \sigma_j & & \vdots \\ & & & \ddots & \vdots \\ & & & & \sigma_N & \mathbf{O} \end{bmatrix} \begin{bmatrix} \vec{v}_1 \\ \vec{v}_2 \\ \vdots \\ \vec{v}_{M-1} \\ \vec{v}_{M-2} \end{bmatrix}. \quad (1.2)$$

The matrix  $\mathbf{\Sigma}$  is a diagonal matrix with nonnegative elements  $\sigma_j$ . For the case shown in Eq. (1.2),  $M > N$ . Then,  $j = 1, \dots, M$ , and  $\mathbf{O}$  is a matrix of 0s. Otherwise,  $j = 1, \dots, N$ , and the zero-padding occurs on the bottom rows of  $\mathbf{\Sigma}$ . The  $\sigma_j$  are the singular values of  $\mathbf{X}$ , and are ordered such that  $\sigma_1 \geq \sigma_2 \geq \dots \geq 0$ . The matrices  $\mathbf{U}$  and  $\mathbf{V}$  are composed of the eigenvectors  $v_m$  (rows of  $\mathbf{V}^*$ ) and  $\vec{\phi}_n$  (rows of  $\mathbf{U}$ ) of the covariance matrices  $\mathbf{X}^*\mathbf{X}$  and  $\mathbf{X}\mathbf{X}^*$  respectively. As a result, the singular value decomposition allows the decomposition of the  $m$ -th column of  $\mathbf{X}$  into

$$\vec{x}_m = \sum_{n=1}^N \sigma_n v_{mn} \vec{\phi}_n, \quad (1.3)$$

where we again assume that  $M > N$ . With this decomposition, the SVD returns a complete orthonormal set of basis functions for the columns of the data matrix  $\mathbf{X}$ . The elements of this basis are the vectors  $\vec{\phi}_n$  also known as the POD modes.

The POD basis is a complete linear basis for the space of the collected data. Due to this completeness, the direct application of the SVD will provide little to no reduction in the dimension of the state space. One way to reduce the dimensionality of the matrix  $\mathbf{X}$  is to use a subset of the POD basis. The relative importance of the  $n$ -th POD mode  $\vec{\phi}_n$  in the approximation of the matrix  $\mathbf{X}$  is determined by the relative energy  $E_n$  of that mode, defined as [4, 6]

$$E_n = \frac{\sigma_n^2}{\sum_{i=1}^N \sigma_i^2} \quad (1.4)$$

where the total energy is normalized such that  $\sum_{j=1}^N E_j = 1$ . Note that because  $\|\mathbf{X}\|_F^2 = \sum_{i=1}^N \sigma_i^2$ , the energy of each mode is, mathematically, the fraction of the Frobenius norm squared that the rank one approximation of  $\mathbf{X}$  generated by the  $i$ th mode can represent. If the sum of the energies of the retained modes is unity, then these modes can be used

to completely reconstruct  $\mathbf{X}$  as the norm of the residual is zero. Typically, the number of modes required to do this is very large and does not result in a significant amount of dimensionality reduction. In practice, however, it has been found that the matrix can be approximated accurately *and low dimensionally* by simply retaining a set of  $P$ -POD modes whose corresponding energies sum to *almost all*, say for instance 99%, of the total energy. Because  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_N \geq 0$  and as a result  $E_1 \geq E_2 \geq \dots \geq E_N \geq 0$ ,  $P$  can be minimized for any given tolerance by including the modes in a sequential fashion, e.g. a 3-mode expansion contains modes 1, 2, and 3, and halting when the desired fraction of the total energy is obtained. The advantage of using a truncated set of POD modes rather than any other set of modes is that the representation of the data generated by the POD modes is guaranteed to have a smaller least squares error than the representation of the data generated by any other orthonormal set of the same size [4].

Overall, the POD method takes in a single set of data and produces a complete basis for that data. Then, that basis is truncated based on the energy of each mode to generate a low dimensional basis that, if the data was sufficiently low dimensional and representative of the desired dynamics, is capable of accurately representing those dynamics. Once this basis has been obtained, standard spectral methods like the Galerkin projection are used to generate governing equations, in the form of nonlinear systems of ordinary differential equations, for the mode amplitudes [3, 6–10]. In this thesis, we will consider systems of nonlinear PDEs with multiple behaviors requiring data from multiple data sets. In these cases, the logical extensions of the POD method to multiple data sets such as Sequential POD [11] or Trust-Region POD [12] or to systems of equations [13] will be used.

### 1.3 Pattern Formation and Mode-Locking in Nonlinear Waveguide Arrays

The POD and other similar techniques can only decrease the dimension of the system while maintaining accuracy if the observed and desired dynamics both lie on the same low dimensional manifold in the phase space of the full system. In situations such as turbulence where this does not occur, the benefits of a POD approach are marginal at best. Despite this limitation, there are many complex physical systems of current interest where this restriction to a low dimensional manifold does occur, namely pattern forming systems. One such system is mode-locking with nonlinear waveguide arrays.

Nonlinear waveguide arrays, or equivalently arrays of nonlinear fibers, have long been of interest to the photonics community due to the nonlinear mode coupling (NLMC) that occurs in these arrays. NLMC has been both experimentally verified [14–17] and theoretically characterized [15, 18] and is of interest because it can be exploited to produce such effects as: pulse switching and routing [19–21], energy localization [22, 23], light bullet generation [24, 25], X-wave generation [26–28], and optical soliton shaping and control [14, 17, 23, 29]. For this thesis, the critical effect of NLMC is that it acts as the saturable absorber required for passive mode-locking [30–33].

The models derived to characterize the mode-locking consist of two governing equations: one for the fiber cavity and a second for the NLMC element [30, 31] (see Fig. 1.1). Although the two discrete components provide accurate physical models for the laser cavity, characterizing the underlying laser stability and dynamics is often analytically intractable. As a




Figure 1.1: Two possible laser cavity configurations that include nonlinear mode-coupling from the waveguide array as the mode-locking element. The fiber coupling in and out of the waveguide array occurs at the central waveguide as illustrated. Any electromagnetic field that is propagated into the neighboring waveguides is ejected (attenuated) from the laser cavity. In addition to the basic setup, polarization controllers, isolators, and other stabilization mechanisms may be useful or required for successful operation.

result, averaged models are used to better understanding mode-locking behavior. Indeed, the essence of Haus' master mode-locking theory [34] is to approximate discrete elements with a continuous model.

The same approach is used in the Waveguide Array Mode-Locking (WGAML) Model to generate a system of continuous governing equations by averaging the pulse-shaping effects of NLMC and gain provided by the Erbium fiber or pumped semiconductor into the laser cavity dynamics [30]. Numerical simulations have shown that the fundamental behavior in the laser cavity does not change when considering more than five waveguides. Further simplification of the five waveguide model is achieved by making use of the symmetric nature of the coupling and lower intensities in the neighboring waveguides [31]. The resulting approximate evolution dynamics describing WGAML is given by

$$i \frac{\partial A_0}{\partial Z} + \frac{D}{2} \frac{\partial^2 A_0}{\partial T^2} + \beta |A_0|^2 A_0 - ig(Z) \left( 1 + \tau \frac{\partial^2}{\partial T^2} \right) A_0 + i\gamma_0 A_0 + CA_1 = 0, \quad (1.5a)$$

$$i \frac{\partial A_1}{\partial Z} + C(A_0 + A_2) + i\gamma_1 A_1 = 0, \quad (1.5b)$$

$$i \frac{\partial A_2}{\partial Z} + CA_1 + i\gamma_2 A_2 = 0, \quad (1.5c)$$

with

$$g(Z) = \frac{2g_0}{1 + \|A_0\|^2/e_0}. \quad (1.6)$$

This models a system like the one shown in Fig. 1.1 but with only three waveguides labeled zero through two.  $A_0(Z, T)$  is the electric field in the 0th waveguide while the  $A_1(Z, T)$  and  $A_2(Z, T)$  fields model the electromagnetic energy in the neighboring channels of the waveguide array. The normalized parameters in the WGAML are as follows:  $D$  controls the strength and the type (normal or anomalous) of the chromatic dispersion,  $\beta$  is the strength of the Kerr nonlinearity,  $g$  is the saturating gain,  $\tau$  is the bandwidth of the gain,  $\gamma_j$  is a linear loss in the  $j$ th waveguide, and  $C$  controls the strength of the evanescent coupling between adjacent waveguides.

The equations in Eq. (1.5) are posed on the infinite domain so that  $A_0, A_1, A_2 \rightarrow 0$  as  $T \rightarrow \pm\infty$ . Further, since the gain  $g(T)$  saturates with the  $L^2$ -norm of the field, spatially ( $T$ ) periodic solutions (e.g. wavetrains) are not allowed due to physical considerations. For instance, an infinite wavetrain would have infinite energy. The equations governing the neighboring fields are ordinary differential equations. All fiber propagation and gain effects occur in the central waveguide. It is this approximate system that will be the basis for our numerical study. Additionally, the simplified system in Eq. (1.5) provides much analytic insight due to its hyperbolic secant solutions

$$A_0(Z, T) = \eta \operatorname{sech}(\omega T)^{1+iA} e^{i\theta Z}, \quad (1.7)$$

where the solution amplitude  $\eta$ , width  $\omega$ , chirp parameter  $A$ , and phase  $\theta$  satisfy a set of nonlinear equations [30, 31]. This limiting solution forms from low amplitude white-noise initial conditions, appears to have a large basin of attraction, and is the mode-locked pulse so often desired from this sort of experimental setup.

The WGAML model also admits other forms of solutions such as  $Z$ -periodic breather solutions and multiple-pulse solutions. Unlike the single-pulse mode-locked solution, no analytic form for these other cases is known. Nonetheless, each has a significant basin of attraction in the correct parameter regimes, and each is a prime example of a spontaneous pattern forming system where dimensionality reduction techniques are applicable.

#### 1.4 Outline of the Thesis

In the following chapters, we apply the Proper Orthogonal Decomposition and other dimensionality reduction techniques to problems of current research interest in nonlinear optics and other disciplines in the physical sciences. In Chapter 2, additional pattern forming systems that occur in nonlinear optics will be discussed. These include spatially-confined light-bullets that serve as a potential architecture for all-photonics routing and logic as well as nonlinear X-wave formation in waveguide arrays. In Chapter 3, one of the standard ansatz based methods for dimensionality reduction, Coupled Mode Theory, is applied to obtain an understanding of the cause of chaotic behavior observed in an experimental waveguide array. In Chapter 4, we will apply data-driven dimensionality reduction techniques to the multipulsing transition, an ubiquitous effect in mode-locking. Furthermore, the highly accurate

Adjoint Continuation Method (ACM) will also be used to confirm the predictions of the reduced model. In Chapter 5, we exploit reduced models to generate a more efficient hybrid integrator for systems with low dimensional dynamics. In Chapter 6, we demonstrate the application of model reduction techniques to problems outside of nonlinear optics; namely, to the branch of time-periodic standing surface waves in fluids of both finite and infinite depth. Lastly, in Chapter 7, we give some concluding remarks.

## Chapter 2

## PATTERN FORMATION IN NONLINEAR OPTICS

In many optical systems, the dynamics of the electric field in experimental settings such as optical fibers or nonlinear waveguide arrays display clear spatial and temporal patterns. From the perspective of dimensionality reduction, the existence of these structures implies that the observed dynamics are constrained to a low dimensional manifold. Perhaps the most famous and widely used of these coherent mechanisms is temporal mode-locking [34, 35]. However, there are many other patterns that appear as well. In this chapter, we demonstrate the existence and the properties of two additional coherent structures that appear in nonlinear optics: X-waves and light-bullets.

Unlike a mode-locked pulse, an X-wave is a transient pattern that, given enough time, will vanish. Initially, X-waves were considered to be a strictly linear phenomenon. In this context, X-waves were localized solutions of linear wave equations in the diffraction- and dispersion-free limit [36, 37]. Despite their genesis in linear theory, X-waves have also been used to describe *nonlinear* phenomena in settings where the underlying physics is described by a *linear* Schrödinger operator [38]. In nonlinear optics, envelope X-waves were first observed in second-harmonic generation [39] and later extended to explain dynamic filamentation [40] as well as parametric generation in water [41]. They also appear in periodic media such as photonic crystals or Bose condensed gases [42, 43] and, of particular interest here, discrete systems such as nonlinear waveguide arrays (WGAs) [23, 27, 30]. The ubiquitous nature of X-waves also suggests that they have applications outside of conservative models, and as a result, X-wave patterns have been predicted in dissipative systems [44].

Although X-waves appear in many nonlinear settings, they are still considered to be essentially a linear phenomenon, i.e., they can be superimposed without any measurable effects appearing. In Section 2.1, we present numerical experiments which demonstrate that in the context of X-waves in WGAs the interaction dynamics of these waves are nonlinear in nature and more similar to those of solitons than linear solutions.

Light-bullets, on the other hand, are the higher dimensional spatial analog of temporal mode-locked pulses, and are still an emerging photonics technology. The generation, stabilization, and control of light-bullets, which we define here to be continuous wave (CW) nonlinear optical structures confined in all three spatial dimensions, are of interest to a myriad of research groups [45–55]. As a result, there are numerous methods currently proposed for generating light-bullets or fundamentally similar cavity solitons, any of which may eventually lead to the adoption and exploitation of this phenomenon by the broader photonics community. As with all technologies, the successful implementation of light-bullet engineering relies on the ability of the system to reliably and inexpensively produce and control these localized optical structures. To this end, in Section 2.2 we propose an additional method for light-bullet generation and control using a planar semiconductor waveguide array, an example of which is shown in Figure 2.1.




Figure 2.1: Schematic of a slab waveguide array theoretically capable of producing light-bullets. Slab waveguides (dark gray) are separated by non-guiding regions (light gray) to produce a weak coupling between adjacent arrays. Gain is applied to the topmost waveguide via current injection.

## 2.1 Nonlinear X-Waves

In this section, we study the dynamics, formation, and proposed nonlinearity of the X-waves that form during mode-locking with nonlinear waveguide arrays. To do so, we perform a series of numerical experiments that measure the interactions of X-waves in a presumed lossless but otherwise physically realizable nonlinear AlGaAs waveguide array [14]. In particular, we examine the interaction of both co- and counter-propagating X-waves. Although these are but two of the many possible X-wave interactions, the experimental setups required to test these interactions are simpler than any of the alternatives and serve as the most viable means of validating these predictions.

The parameters that will be used are taken from a 3 mm AlGaAs WGA with input pulses generated from a mode-locked laser producing full-width half maximum (FWHM) pulsewidths of 200 fs [23]. The linear coupling coefficient is taken to be  $0.82 \text{ mm}^{-1}$  and the nonlinear self-phase modulation parameter is taken to be  $3.6 \text{ m}^{-1}\text{W}^{-1}$ . In waveguides, chromatic dispersion also is present, but in keeping with previous works it has been neglected in this system [26]. In any experimental fiber-waveguide system, the dispersive effects would be dominated by the fiber which allows it to be neglected here. Additionally, simulations with the inclusion of realistic dispersion values, as measured experimentally [23], show the same dynamics as are demonstrated here. Hence, we are justified in neglecting dispersion in all that follows. Typical peak powers in the WGA are on the order of kilowatts, and the total number of waveguides is 41 [14, 23]. Because of the parameters and initial conditions used, trivial modifications to recent experiments [14, 23, 27] should suffice to confirm the numerical simulations presented here.

### 2.1.1 Co-Propagating X-Waves

The first type of X-wave interaction is between co-propagating X-waves. In this interaction, pulses are injected into adjacent waveguides. The leading-order equations governing the nearest-neighbor coupling (discrete diffraction [14]) of electromagnetic energy in the WGA




Figure 2.2: Pseudo-color plot of in-phase pulses interacting. The final result is a pair of closely-spaced but distinct X-waves. The white and green dotted lines indicate the center of the X-wave at  $Z = 0$  in the 0th and 1st waveguide respectively. The white and green dashed lines represent the center of the X-wave at  $Z = 3$  mm. Note that the final separation is slightly larger than the initial separation of the pulses.

are given by

$$i \frac{\partial A_n}{\partial Z} + \frac{D}{2} \frac{\partial^2 A_n}{\partial T^2} + c(A_{n-1} + A_{n+1}) + \gamma |A_n|^2 A_n = 0, \quad (2.1)$$

where  $A_n$  represents the electric field envelope in the  $n$ th of the  $2N + 1$  waveguides ( $n = -N, \dots, -1, 0, 1, \dots, N$ ) where  $N = 20$  for 41 waveguides. Using the parameter values for the physical waveguide of interest:  $D = 0$ ,  $c = 0.82 \text{ mm}^{-1}$ , and  $\gamma = 3.6 \text{ m}^{-1} \text{ W}^{-1}$ . This set of governing equations has been shown to accurately reproduce experimental findings for pulses with kilowatt peak powers and pulsewidths of hundreds of femtoseconds [23].

To begin, pulses are launched into two adjacent waveguides with

$$A_0(0, T) = \eta_0 \text{sech}(T), \quad \text{and} \quad A_1(0, T) = \eta_1 \text{sech}(T + \Delta T) \times \exp(-i\Delta\theta). \quad (2.2)$$

and  $A_n(0, T) = 0$  for  $n \neq 0, 1$  where  $\eta_0, \eta_1 = 2.0$  and  $\Delta T = 1$ . With these initial conditions, fully nonlinear X-waves are formed during the propagation in the WGA, but the dynamics of the X-waves depend on the relative phase difference between the injected pulses. Figure 2.2 demonstrates a time-history of propagation of Eq. (2.1) with an initial phase difference of  $\Delta\theta = 0$ . Although the majority of the energy in the X-wave remains confined in the initial waveguide, the low amplitude sections interact and perturb the X-waves. Ultimately, the pair of X-waves that results have a slightly larger separation than if the X-waves were input individually, i.e., the X-waves repelled. In contrast, consider the initial conditions in Eq. (2.2) with identical amplitudes but with  $\Delta\theta = \pi$ . A time-history of Eq. (2.1) with these initial conditions is shown in Fig. 2.3. The WGA still generates an X-wave from each of the initial pulses as shown in the second panel in Fig. 2.3. However, the separation in  $T$  between the resulting X-waves is now negligible, i.e., they attract. This demonstrates that, for X-waves with identical amplitudes, the resulting dynamics depend on the initial phase-difference between the pulses.

Furthermore, the dynamics do not, qualitatively, depend on the initial separation between




Figure 2.3: Pseudo-color plot of out of-phase pulses interacting. In this case, the X-waves attract and form a pair of X-waves with a negligible delay in time. The white and green dotted lines denote the center of X-wave in the 0th and 1st waveguides at  $Z = 0$ , and the dashed lines represent the center of the X-wave at the present value of  $Z$ .




Figure 2.4: Plot of the X-wave separation as a function of phase-difference for three different initial separations. The dashed lines show the initial separation and the solid lines of the same color show the final X-wave separation. Regardless of the initial separation, X-waves with small phase-differences repel and X-waves with phase-differences near  $\pi$  attract.

the two X-waves. Figure 2.4 shows the final separation between the X-waves as a function of phase difference for three different initial separations,  $\Delta T = 0.5, 1.0, \text{ and } 1.5$ . In all cases, the X-waves repel for  $\Delta\theta$  near zero, and the X-waves attract for  $\Delta\theta$  near  $\pi$ . As shown in Figure 2.4, the strength of interaction does depend upon the initial separation. As a result, X-waves that have a small initial separation will both repel more strongly and also attract at smaller values of  $\Delta\theta$ . Likewise, X-waves with larger initial separations will repel and attract less strongly. This is necessary for agreement with the limiting case of X-waves that are so far separated that they are non-interacting and will neither attract nor repel.

Similarly, the differences in the maximum pulse amplitude determine the amount of attraction or repulsion but, for sufficiently large initial amplitudes, do not change the qualitative behavior of the X-waves. Figure 2.5 shows the final separation of two co-propagating X-waves that are initially either in-phase or  $\pi$  out-of-phase for a variety of different initial pulse amplitudes. Although the final position depends on the pulse amplitude, the in-phase




Figure 2.5: Plot of the final X-wave separation for a pair of in-phase X-waves, shown in black, and  $\pi$  out-of-phase X-waves, shown in blue with an initial separation of  $\Delta T = 1$ . For sufficiently high initial powers, the in-phase solutions repel and the out-of-phase solutions attract.




Figure 2.6: Pseudo-color plot of two unequally sized X-waves with  $\Delta\theta = 0$ . The X-wave in waveguide 0 has amplitude 2.0 while the pulse in waveguide 1 has amplitude 1. The larger pulse dominates the dynamics and the attracts the smaller pulse.

pulses, shown in black, repel while the out-of-phase pulses, shown in blue, attract. These behaviors persist for a large range of pulse amplitudes, though the amount of attraction or repulsion appears to decrease as the intensity of both pulses increases. The results in Figure 2.5 were obtained using  $\Delta T = 1$ , but the same qualitative results occur even if other separations are used.

In the previous case, the two initial pulses were of equal amplitude. If pulses of different amplitude are launched, the pulse with the largest peak power dominates the interaction. Figure 2.6 shows the interaction dynamics of two X-waves with initial conditions given by Eq. (2.2) where  $\eta_0 = 2.0$  and  $\eta_1 = 1.0$  and  $\Delta T = 1$  and  $\Delta\theta = \pi$ . Unlike the equally sized pulses, the larger pulse dominates the dynamics and incorporates the smaller pulse into a single X-wave structure regardless of the phase difference. Here, it is the amplitude difference and not the phase-difference that determines the resulting dynamics.

To summarize: the interactions between high amplitude co-propagating X-waves are nonlinear in nature and exhibit, in some sense, soliton-like dynamics [56]. That is, whether




Figure 2.7: Pseudo-color plot of the collision of counter-propagating X-waves. The top series of plots shows the propagation of both the forward and backward X-waves at three snapshots in time. The bottom series of plots follows the propagation of the forward X-wave as it interacts with the backward X-wave.

the X-waves attract or repel is dependent upon both the relative phase-difference and the amplitude difference of the two injected pulses. Although there does exist a linear limit in which the X-waves can be superimposed without any additional effects, the conditions for this limit do not occur during standard mode-locking operation of a WGA laser.

### 2.1.2 Counter-Propagating Waves

Another interaction of interest is the interaction between identical but counter-propagating X-waves. This situation could be obtained experimentally by butt-coupling input fibers to opposite ends of a waveguide. The governing equations for counter-propagating waves must include both the forward- and backward-propagating fields. The governing equations, Eq. (2.1), must be modified to account for the second field yielding

$$i\frac{\partial A_n}{\partial Z} + i\sigma\frac{\partial A_n}{\partial T} + \frac{D}{2}\frac{\partial^2 A_n}{\partial T^2} + \gamma(|A_n|^2 + 2|B_n|^2)A_n + c(A_{n-1} + A_{n+1}) = 0 \quad (2.3a)$$

$$i\frac{\partial B_n}{\partial Z} - i\sigma\frac{\partial B_n}{\partial T} + \frac{D}{2}\frac{\partial^2 B_n}{\partial T^2} + \gamma(2|A_n|^2 + |B_n|^2)B_n + c(B_{n-1} + B_{n+1}) = 0 \quad (2.3b)$$

where  $A_n$  is the forward-propagating field of the  $n$ th waveguide, and  $B_n$  is the backward-propagating field of the  $n$ th waveguide. In these equations, the nearest-neighbor coupling and self-phase modulation terms are retained from the co-propagating case, but an additional cross-phase modulation term appears along with a group-velocity term determining the forward ( $+\sigma$ ) and backward ( $-\sigma$ ) directions of propagation.

The collision of counter-propagating X-waves can be accomplished by launching initial pulses on both sides of the waveguide simultaneously. Thus the initial conditions take the

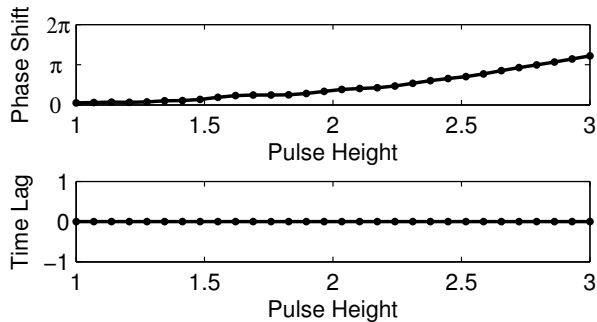


Figure 2.8: Plot of the phase shift in radians and time delay of the interacting X-waves. The magnitude of the phase shift is dependent upon the size of the X-wave, but the time-delay remains zero for all initial conditions.

form

$$A_0(Z, 0) = \eta_+ \operatorname{sech}(Z + \Delta Z), \quad \text{and} \quad B_0(Z, 0) = \eta_- \operatorname{sech}(Z - \Delta Z), \quad (2.4)$$

where  $2\Delta Z$  measures the initial spatial separation between the right moving (forward-propagating) and left moving (backward propagating) X-waves. It must be stressed that unlike the co-propagating interaction the governing equations of this system *occur in a stationary lab frame* and the initial condition is defined for all  $Z$  at  $T = 0$  and not in the usual optical coordinate system (i.e. for all  $T$  at  $Z = 0$ ). With these initial conditions, the X-waves collide at  $Z = 0$ . Figure 2.7 demonstrates the basic collision dynamics. Notice here that the two X-waves pass through each other without visible deformation, and the collision appears, on the surface, to be *linear*.

However, there is an induced nonlinear phase shift due to the collision that depends upon the initial launch intensity of the counter-propagating pulses, which is demonstrated in Figure 2.8. To calculate the phase-shift and time-lag, two simulations were performed. The first simulation includes both the forward- and backward-propagating X-waves. The second simulation only included the forward-propagating X-wave. Both simulations used the same physical and computational parameters as well as the same initial condition for the forward-propagating X-wave. Indeed, the only difference is the presence of the backward-propagating X-wave. After the interaction occurs, the cross-correlation of the forward-propagating X-wave in the 0th waveguide with and without the backwards propagating pulse was computed. As in experiments, the maximum of the convoluted pulses determines both the time-lag of the X-wave and the relative phase difference between them. It is clear from Fig. 2.8 that X-waves act as linear waves for sufficiently low initial intensities. As the injected pulses are made more intense, the nonlinearities begin to exhibit themselves and the phase-shift increases in a monotonic fashion, but as shown in Fig. 2.7, the interaction still appears to be linear to the eye. Indeed, regardless of the pulse height there is no observable spatial lag generated by this interaction.

In order to explore the dependence of the interaction on waveguide length and initial pulse sizes, we consider five different initial values of  $\eta_+$  and  $\eta_-$  for various values of  $\Delta Z$  as




Figure 2.9: Plot the phase shift of counter-propagating X-waves as a function of initial separation. The different colors correspond to initial conditions with peak pulse amplitudes of 1.50, 1.25, 1.00, 0.75, and 0.50 for the blue, green, red, teal, and purple lines respectively.

shown in Fig. 2.9. In all cases, there is no measurable change in the location of the pulse center (pulse delay) due to the interaction, but the amount of phase-shift is dependent upon both the initial separation and the initial pulse height. The results of Fig. 2.9 may be broken down into two separate regimes, low-amplitude initial conditions and high-amplitude initial conditions.

The low amplitude cases are shown in the red, teal, and purple lines in Fig. 2.9 and correspond to maximum pulse amplitudes of 1.00, 0.75, and 0.50. These results are characteristic of lower amplitude solutions as well. In these cases, the amount of phase-shift decreases as the distance between the initial pulses increase. At relatively short initial separations, the small amplitude X-wave has little time to distribute energy to neighboring waveguides through discrete diffraction. Because the phase shift is generated solely by nonlinear effects, the short distances allow the X-waves to be as nonlinear as possible given their initial conditions. Indeed, as the separation is increased the phase-shift decreases because the X-wave distributes energy to the outer waveguides and acts in a linear fashion. With an initial separation of 1.5, there is essentially no noticeable phase-shift. On the other hand, with sufficiently intense initial conditions the X-waves can interact nonlinearly even after propagating for a significant distance because the nonlinear mode-coupling traps energy in the original waveguide.

In this section, we have demonstrated that, in the physical regime associated with mode-locking, X-waves act in a manner more reminiscent of solitons than they do linear structures. Like solitons, when X-waves are placed in a co-propagating arrangement whether they attract or repel depends primarily on the phase difference between the two X-waves. Furthermore when X-waves collide in a counter-propagating situation, both X-waves experience a phase-shift. They do not, however, exhibit the shift in the pulse center that occurs when optical solitons collide. Because X-waves are transient structures, they can in the long  $Z$  limit be treated in a linear fashion. However, we have demonstrated that, with the time scales and optical powers associated with an AlGaAs waveguide array, the assumption of linearity fails to accurately capture the entirety of the interaction dynamics present in the system.

## 2.2 *Light-Bullet Formation, Routing, and Control*

In this section, we demonstrate the feasibility of using nonlinear semiconductor waveguide arrays to generate, control, and route light-bullets. The remainder of the section is structured as follows: in Section 2.2.1, a description of the physical system is given. In Section 2.2.2, the governing equations and model parameters are given. In Section 2.2.3, the creation and stability of light-bullets is discussed for uniform gain profiles. In Section 2.2.4, the routing and control of bullets using non-uniform gain profiles are demonstrated. In Section 2.2.5, the interaction of multiple bullets is used to produce the NAND and NOR gate. Lastly, Section 2.3 contains some concluding remarks and a technological outlook for the WGA device and its applications.

### 2.2.1 *System Description*

The proposed device for light-bullet generation is a slab waveguide array consisting of three planar waveguides, as shown in Fig. 2.1. The waveguides are labeled zero, one, and two from the top of the device to the bottom respectively. Non-guiding regions are placed between the individual waveguides to weaken the coupling between adjacent waveguides. Gain could be provided to the system through a monolithic gold layer attached to the zeroth waveguide [57]. This layer allows for the injection of current into only the zeroth waveguide. One modification useful for applications is to partition the gold layer as in Ref. [25]. By injecting different amounts of current into individual partitions, a non-uniform gain profile is generated with uses that will be discussed later in the section.

This device is envisioned to operate like a vertical cavity surface emitting laser (VCSEL) [58], so the phase velocity of light is in the vertical direction in Fig. 2.1. The light-bullets themselves are stationary in space and confined in the plane of each of the waveguides. This geometrical structure extends the idea of the nonlinear mode coupling (NLMC) [23,30], so critical in temporal mode-locking, from a single dimension to the pair of spatial dimensions in the plane of the waveguide.

The third dimension of confinement is assumed to be created by the structure of the array, which can be accomplished with many current technologies. One particular example is the zero group velocity soliton, which has been obtained theoretically and is the target of ongoing experimental work [59,60]. Another option, particularly suited to the WGA, is the use of defects in Bragg gratings to trap Bragg solitons [61]. By replacing the guiding regions with Bragg gratings and using the insulating regions as defects, the geometrical structure of the WGA itself should be sufficient to generate a trapped Bragg soliton.

Either of the previous two mechanisms would be sufficient to create effectively zero-group velocity bullets. Since the interaction of solitons in the adjacent waveguides is modeled via coupled mode theory [62], the exact mechanism is not important to the overall structure of the governing equations. The slab waveguide mode-locking (SWGAML) model discussed in the next subsection is robust enough to capture both situations. The combination of these two effects, NLMC in the plane of the waveguide and the Bragg or gap-soliton in the normal direction, is sufficient to produce confinement in all three spatial dimensions. As will be shown, these light-bullets act approximately as global attractors in the SWGAML model and will therefore naturally appear from a cold cavity.

### 2.2.2 Governing Equations

Mode-locking in waveguide arrays is created by a competition between the saturable absorption generated by NLMC [32] of the waveguides and the bandwidth limited gain. The WGAML model in Eq. (1.5) describes the temporal mode-locking in traditional ridge waveguide arrays. To model the slab waveguide system in Fig. 2.1, the WGAML model was heuristically extended from one to two spatial dimensions,

$$i\frac{\partial A_0}{\partial t} + \frac{D}{2}\nabla^2 A_0 + \beta|A_0|^2 A_0 + CA_1 + i\gamma_0 A_0 - ig(x, y, t)(1 + \tau\nabla^2)A_0 + i\rho|A_0|^4 A_0 = 0, \quad (2.5a)$$

$$i\frac{\partial A_1}{\partial t} + C(A_0 + A_2) + i\gamma_1 A_1 = 0, \quad (2.5b)$$

$$i\frac{\partial A_2}{\partial t} + CA_1 + i\gamma_2 A_2 = 0, \quad (2.5c)$$

where  $\nabla^2 = \partial_x^2 + \partial_y^2$ . The impact of current injection is modeled as a saturating gain,

$$g(x, t, y) = \frac{2g_0 f(x, y)}{1 + ||A_0||^2/e_0}. \quad (2.6)$$

In Eq. (2.5),  $A_0, A_1$ , and  $A_2$  are the envelopes of the electric fields in the 0th, 1st, and 2nd waveguides respectively. Although mathematically the SWGAML model is nothing more than the WGAML model in Eq. (1.5) with two spacial dimensions, the physical interpretations of the models and the parameter values differ. The SWGAML model is in a stationary (lab) frame and so  $D$  is the diffraction coefficient. Note that for negative refractive index materials  $D$  is negative [63]. This allows the SWGAML model to reproduce both the anomalous and normal dispersion regimes on the WGAML model despite the different physical circumstances.  $\beta$  is the strength of the Kerr nonlinearity,  $\rho$  is proportional to the probability of three photon absorption, the  $\gamma_j$ s are the aggregation of linear losses for each waveguide, and  $C$  is the coupling strength between adjacent waveguides. The saturable gain  $g(x, y, t)$  is of the same form as Eq. (1.6), but the saturation accounts for the depletion of minority charge carriers at high optical intensities rather than the physics of erbium doped fibers. The filtering term,  $g\tau\nabla^2$ , arises from the diffusion of the charge carrier that appears when the charge carriers are modeled explicitly [64].

The function  $f(x, y, t)$  in Eq. (2.6) allows for the possibility of non-uniform gain profiles. In order to uniquely specify the gain, we impose that the mean of  $f(x, y, t)$  is one at all times and  $f(x, y, t) \geq 0$ . Therefore, larger values of  $g_0$  always correspond to larger total injection currents regardless of the exact form of  $f(x, y, t)$ . The addition of a non-uniform gain allows for a variety of additional dynamics not found in the uniform gain case. In particular, a non-uniform gain breaks translational invariance in the system and creates solutions where the bullets translate in space.

### 2.2.3 Uniform Gain Dynamics

The study of uniform gain dynamics is an important first step in describing the SWGAML. With uniform gain, i.e.  $f(x, y, t) = 1$ , there are two distinct asymptotically stable types of mode-locked solutions: stationary solutions and breather solutions. For both types, the system can have an arbitrary number of bullets, e.g. there exists a double bullet stationary solution.

#### *Stationary Light-Bullets*

The first type of mode-locked solution is the stationary (in amplitude) bullet solution which have constant amplitudes and constant phase evolution in time. With uniform gain, the SWGAML is a radially symmetric system of equations. Although non-radially symmetric solutions could exist, the radially symmetric solutions are of lower energy and are therefore the favored solutions when starting from a cold cavity. The radial symmetry can be exploited to remove a spatial dimension from the problem allowing for faster and more accurate simulations. Figure 2.10 shows an example of a radially symmetric solution obtained with

$$(D, C, \gamma_0, \gamma_1, \gamma_2, e_0, \tau, \rho, g_0) = (-1, 10, 0, 0, 10, 1, 0.1, 1, 35). \quad (2.7)$$

These parameter values correspond to an experimentally realizable AlGaAs waveguide array [31]. The system quickly settles into a steady state starting from noisy initial conditions. To obtain this result, a finite difference scheme was used with Neumann boundary conditions at the origin. The computational domain used is larger than that shown in Fig. 2.10 ( $r \in [0, 100]$ ) so the Dirichlet boundary condition used at the far end has a negligible impact on the resulting dynamics.

For the parameters chosen, this solution appears to be the only stable attractor in the system. The software package AUTO [2] was used to track the branch of stationary solutions using a finite difference approximation to compute the derivatives. This coarse approximation to the function was used as a starting point for a more accurate calculation using a Chebyshev collocation method [65]. This collocation approach allows for the eigenvalues of the linearized operator to be computed with spectral accuracy. Errors in the eigenvalues are estimated to be on the order of  $10^{-8}$ . This combined approach was used to generate the bifurcation diagram and sample spectra in Fig. 2.11 for the stationary solutions. In Fig. 2.11, blue regions are linearly stable and red regions are linearly unstable. It is clear that a saddle-node bifurcation is responsible for the creation of the single-pulse solution. The maximum amplitude of the pulse is directly related to the gain provided to the system, but around  $g_0 = 50$  a Hopf bifurcation occurs and the periodic breather solutions become the stable solutions.

#### *Time-Periodic Breather Solutions*

Due to the Hopf bifurcation, it is known that time-periodic breather solutions occur for some range of the bifurcation parameter  $g_0$ . While more complicated than the stationary solutions, breather solutions contain both a larger peak intensity and a larger total energy




Figure 2.10: Example of a radially symmetric stationary solution with the parameters from Eq. (2.7). The initial condition is noise in the zeroth waveguide and zero in the other two waveguides. The radially symmetric bullet solution forms from noise and appears to be a global attractor of the system. Note that the outer two waveguides also possess rotational symmetry and have qualitatively similar stationary solutions.

and may be useful in applications where these traits are desired and a stationary bullet profile not required.

Figure 2.12 shows an example of a time-periodic breather solution using the parameters:

$$(D, C, \gamma_0, \gamma_1, \gamma_2, e_0, \tau, \rho, g_0) = (-1, 10, 0, 0, 10, 1, 0.1, 1, 50). \quad (2.8)$$

All the parameters, except for  $g_0$ , are identical to the single bullet case. In Fig. 2.12, it was assumed that the periodic solutions were also radially symmetric. However, in full two-dimensional simulations it was observed that a low amplitude non-radial background occurs in addition to the central radial bullet [66], which is neglected by the radial solution approximation. It is unclear whether this non-radial background is inherent in the PDE or is a result of the periodic boundary conditions and square domain imposed in the two-dimensional problem. However, in both the radially symmetric case and the full PDE the results are qualitatively similar. If quantitative predictions of bullet dynamics are required, the radially symmetric approximation is no longer valid for the parameters used in simulation, and an analog of transparent boundary conditions for two spatial dimensions is required [67].

### *Multiple-Pulse Solutions*

At larger gains, the breather solution loses stability and a stationary solution with a larger number of bullets dominates. These stationary solutions are equivalent to multiple single-bullet solutions in Sec. 2.2.3. In all the simulations performed, the final solution has been two

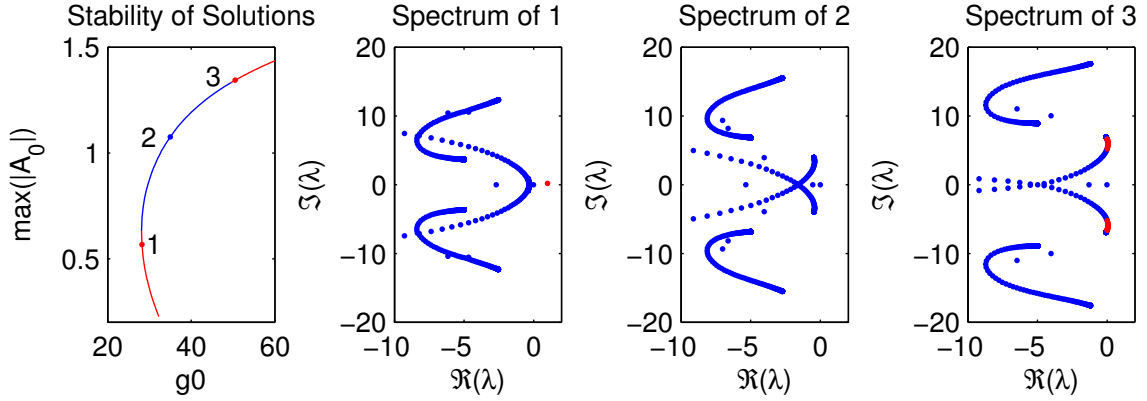


Figure 2.11: The bifurcation diagram and spectrum of selected radially symmetric bullet solutions. In the bifurcation diagram, the blue curve is where the bullet is linearly stable while the red regions are where it is not. In the spectra, red  $x$ 's represent eigenvalues with positive real parts. Light-bullet solutions are born out of a saddle node bifurcation near point 1 and go unstable near point 3 due to a Hopf bifurcation.




Figure 2.12: The amplitude in the zeroth waveguide as a function of time for a time-periodic breather solutions using the parameters in Eq. (2.8). At higher gains, the stationary solution is unstable and the breather solution emerges.

non-interacting stationary bullets, similar to what has been observed in the one-dimensional WGAML [31].

While the final state of the solution is quite simple, the route taken from a single to a double bullet is not, and the transition may persist for long periods of time. Figure 2.13 shows the time evolution of a solution that splits. While the initial condition is that of a single stationary solution bullet, the gain is too large for the stationary or even the breather solutions to be stable. As a result, the bullet radiates energy and eventually a second bullet forms from the energy traveling around the periodic domain. If the second bullet forms in close proximity to the existing bullet, the bullets will recombine into a single unstable bullet, which starts the splitting process over again. Ultimately, this process results in two bullets that are sufficiently separated so that they will not recombine. It should be noted that it is possible to form stationary solutions where the bullets do interact. For identical




Figure 2.13: Time evolution of the splitting of a single light-bullet into a pair of light-bullets. Due to the larger gain, the single pulse sheds energy which eventually results in the formation of a second pulse. Note that the resulting bullets are far separated in space. If the second bullet forms in close proximity, it will interact with the original bullet and be absorbed which repeats the splitting process.

bullets, separations have been found where the relative phase difference between the bullets determines whether or not the pair persists or recombines. However, this situation has only occurred when very specific initial conditions are employed and never when a single bullet was used as the initial conditions. Therefore although stable dipole solutions exist, for simplicity we will assume the multiple bullet solutions are simply a number of non-interacting bullets.

The splitting process in Fig. 2.13 appears to be chaotic in nature and the SWAGML may remain in this chaotic state for protracted periods of time. Note that the 160 time units in Fig. 2.13 is shorter than the typical amount of time requires for bullets to split. For many parameters, this transition process may exist for over a thousand time units and is heavily dependent on the initial condition.

#### 2.2.4 Non-Uniform Gain Dynamics

Using  $f(x, y, t) = 1$ , the governing equations Eq. (2.5) contain a number of symmetries, including translational invariance. The translational invariance eliminates the possibility of creating a single spatial location that attracts bullets in the array. In experimental works

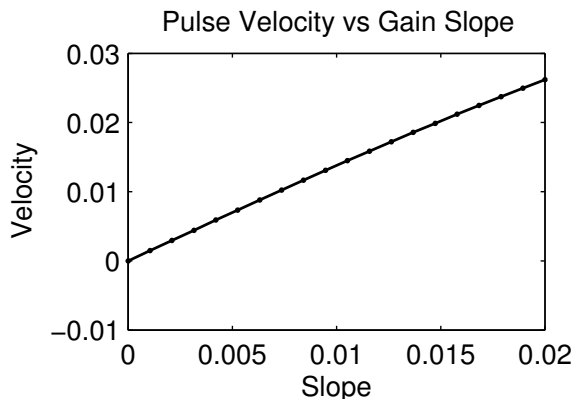


Figure 2.14: Plot of the velocity of the bullet as a function of the slope of the gain profile. The velocity of the bullet is directly and linearly related to the slope of the gain.

on similar systems [68], external lasers were used to pump the system and manipulate the location where bullets formed. A similar effect can be obtained in this system using variations in the current injected to the system.

#### *Constant Gain Slope*

Perhaps the simplest non-uniform gain is a linearly sloped gain profile, where the slope of the gain profile is kept small. This restriction serves two purposes. First, the use of small slopes allows the non-uniformity to be considered a perturbation of the uniform case. In this regime, the branches of solutions obtained in Fig. 2.11 will remain valid to leading order. Therefore, the solutions from the uniform case can be used as initial conditions without a large transient resulting. Second, the inclusion of large gain slopes often creates regions with enough gain that a second bullet simply forms in that region, annihilating the first bullet. Although there is no inherent technological limitation to small gain slopes and this type of motion may be physically realistic, it is more difficult to control and will not be considered here.

Figure 2.14 shows that the velocity of the light bullet is directly related to slope of the gain. For the small slopes used, the bullet is not noticeably deformed from the stationary bullet case, and the bullet trajectory is the same as the gradient of the gain function, i.e. there is no spurious motion in any other direction. The ramped gain case can serve the conceptual starting point for more complex gain profiles, and it provides a simple mechanism for imparting velocity onto the light-bullets. For more complex behaviors, one approach is to use piecewise linear functions to trace out paths in the plane.

#### *Bullet Routing with Sloped Gain*

The use of a sloped gain is a simple method of bullet routing. As shown in Sec. 2.2.4, bullets obtain a velocity in the direction of the gradient of the gain. The simplest types of functions




Figure 2.15: Bullet routing using the gain equation in Eq. (2.9). The dotted lines show the location of the bullet center as time progresses. The piecewise linear gain routes the bullet through the junction.

would be comprised of piecewise linear functions. These functions are capable of robustly routing bullets even through large angles.

As an example, a cross shaped demultiplexer with one input and up to three outputs can be created by generating a single gain ramp and superimposing it with forbidden regions that contain no gain. Here, gain was modeled as

$$f(x, y, t) = \begin{cases} (1 + mx + ny) & |x| < 8 \text{ or } |y| < 8, \\ 0 & \text{otherwise,} \end{cases} \quad (2.9)$$

where  $m$  and  $n$  control the direction in which the bullet moves by changing the gradient of the gain. Figure 2.15 shows the three possible routings with  $m = 0.01$ . The final bullet location is chosen by the value of  $n$ . When  $n = 0.01$  the bullet is routed up, when  $n = 0$  the bullet is routed across, and when  $n = -0.01$  the bullet is routed downward.

In the application of a piecewise linear gain, the regions of zero gain in Fig. 2.15 prevent the bullet from entering the region, which causes the bullet to travel through the junction. The inclusion of these regions of no gain creates great flexibility in the generation and construction of devices used for bullet routing.

### *Time Dependent Gains*

Stationary gain profiles are the simplest to construct but have inherent limitations. As an example, stationary gain profiles are unable to manipulate light bullets into a time-periodic orbit. For small enough gain slopes, light bullets follow the gradient of the gain function. However, no continuous function can have a closed orbit with an always increasing gradient. Additionally as shown in Sec. 2.2.4, the solutions to the SWGAML model have effectively no “momentum” so the discontinuity cannot be overcome.

However, time-periodic orbits can be easily obtained by using a time-periodic gain func-




Figure 2.16: Bullet control using time dependent gain. The + and o show the  $x$  and  $y$  coordinates of the bullet center as a function of time. The gray lines show the  $x$  and  $y$  coordinates of the gain function in Eq. (2.10) as a function of time.

tion. A simple case is a gain profile such as the Gaussian profile:

$$f(x, y, t) = \exp\left(-\alpha\left([x - 10 \cos \omega t]^2 + [y - 10 \sin \omega t]^2\right)\right), \quad (2.10)$$

where  $\alpha = 0.001$  and  $\omega = 2\pi/5000$ . The angular frequency of movement,  $\omega$ , is small relative to the timescale that dynamics typically occur on in the SWGAML. This is consistent with most physical systems where gain manipulations occur on much slower timescales than the evolution of the bullets.

Due to the slow translation of the gain the movement of the bullet can be thought of as an adiabatic process. For any position of the gain, the bullet has sufficient time to follow the gain gradient to the maximum of the gain profile. As the gain translates, this process repeats. The effect is a light bullet that is bound to the top of the translating gain profile and follows a path nearly identical to the moving gain. This process is shown for the Gaussian profile in Eq. (2.10) in Fig. 2.16.

In Fig. 2.16, the  $x$  and  $y$  position of both the bullet and the gain are shown for a variety of times. In order to determine the bullet center, the following formulas were used for both the bullet and the gain:

$$\text{bullet center} = \frac{\iint x |A_0|^2 dx dy}{\|A_0\|^2} \hat{x} + \frac{\iint y |A_0|^2 dx dy}{\|A_0\|^2} \hat{y}. \quad (2.11)$$

As the gain is radially symmetric and the bullet is, to leading order, radially symmetric, this formula accurately describes the centers of both functions. With this measure of position, the bullet and gain overlap to a significant degree. There is a slight lag between the centers of the gain and the bullet as this process is not completely adiabatic. Nonetheless, it is clear that these translating gains can manipulate light bullets in time-periodic orbits and

Table 2.1: Logic table enumerating the input and outputs of the NOR and NAND gates. Zeros in the table denote logical low inputs or outputs, and ones in the table denote logical high inputs or outputs. To be considered a NOR gate or a NAND gate, the device must satisfy the four cases given below.

NOR Gate			NAND Gate		
Input 1	Input 2	Output	Input 1	Input 2	Output
0	0	1	0	0	1
0	1	0	0	1	1
1	0	0	1	0	1
1	1	0	1	1	0

indeed any arbitrary continuous path. Since the current injected into the system is controlled by external electronics, these same external electronics indirectly provide a mechanism for bullet routing and control.

### 2.2.5 Gain Mediated Interactions

In Sec. 2.2.3, it was shown that the SWGAML is capable of supporting multiple-bullet solutions. Furthermore by employing non-uniform gain profiles in Sec. 2.2.4, we demonstrated that it is possible to route both bullets simultaneously. In this section, we exploit the interactions between multiple bullets, which occur through two distinct processes. The first is a direct interaction when two bullets are physically close enough to interact, similar to the interactions observed in the nonlinear Schrödinger equation (NLS) [69]. Similar to NLS, the resulting dynamics of the interaction depends heavily on the separation of the two bullets as well as the relative phase difference between them. The resulting dynamics is therefore primarily influenced by the initial condition of the SWGAML.

In the applications envisioned for this device, it is unlikely one would be able to ensure a particular phase-difference due to the overall phase invariance inherent in all optical systems. The simpler and more robust of the interactions are gain-mediated interactions. These types of interactions occur only through the gain term in Eq. (2.6), where the level of saturation is determined by the  $L^2$  norm. This non-local term allows bullets that are physically separated to influence each other by increasing or reducing the saturation of the gain in the system. While less powerful than direct interactions, this mechanism is still capable of producing both the NOR and NAND logic gates and would therefore be useful in applications. For reference, the logic tables for the NAND and NOR gates are reproduced in Table 2.1.

#### *NOR Gate*

The first of the two master gates implemented with SWGAML is the NOR gate. The NOR gate has up to three inputs. A clock bullet must always exist in the system, but the input bullets will only exist if that particular input is a logical high input. The NOR gate is produced using three stationary gain ramps which route both input bullets as well as the

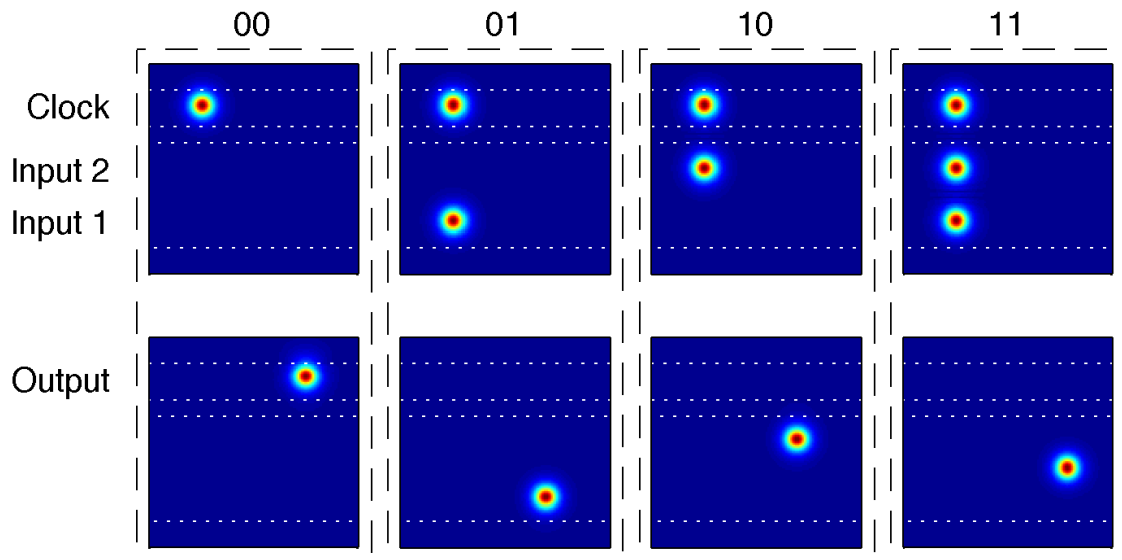


Figure 2.17: A plot of all possible cases for the NOR gate. Both the clock and input bullets experience a sloped gain which routes all bullets from left to right. The gain is biased so the clock receives less gain than either of the input bullets. The system possesses enough gain to support a single bullet so the clock bullet will be destroyed if either input bullet exists. If neither exist, the clock bullet translates right giving a high output.

clock bullet from the input of the NOR gate to the output. If the clock bullet reaches the output of the gate, then the result is logical high. Otherwise, the output is logical low.

What differentiates this from the bullet routing shown in previous sections is that, while there are three bullets, the system is only given sufficient gain to support a single bullet. Therefore, all three of the bullets will decay. To select which of the three bullets survives, each of the gain ramps is given a slightly different amount of gain. Specifically, the clock bullet receives less gain at all points on the ramp than either of the two input bullets. From the stability results of Fig. 2.11, a slightly reduced gain is not of consequence because single bullets are stable for a range of gains.

If the clock bullet is the only bullet in the system, it will have enough gain to persist and translate to the output. If any of the two input bullets are introduced to the system, then all of the bullets will decay. Due to the lower gain given to the clock bullet, it will decay faster than either of the input bullets, and in the end, it will always be destroyed. These results are consistent with a NOR gate.

To verify, Fig. 2.17 shows each of the four possible scenarios for the NOR gate. These results should be compared to the logic table for the NOR gate in Table 2.1. Because these interactions are completely gain mediated, this process does not require any particular phase difference between the input bullets to exist. Furthermore, while both of the input bullets in Fig. 2.17 received the same gain, that need not be the case. As long as the clock

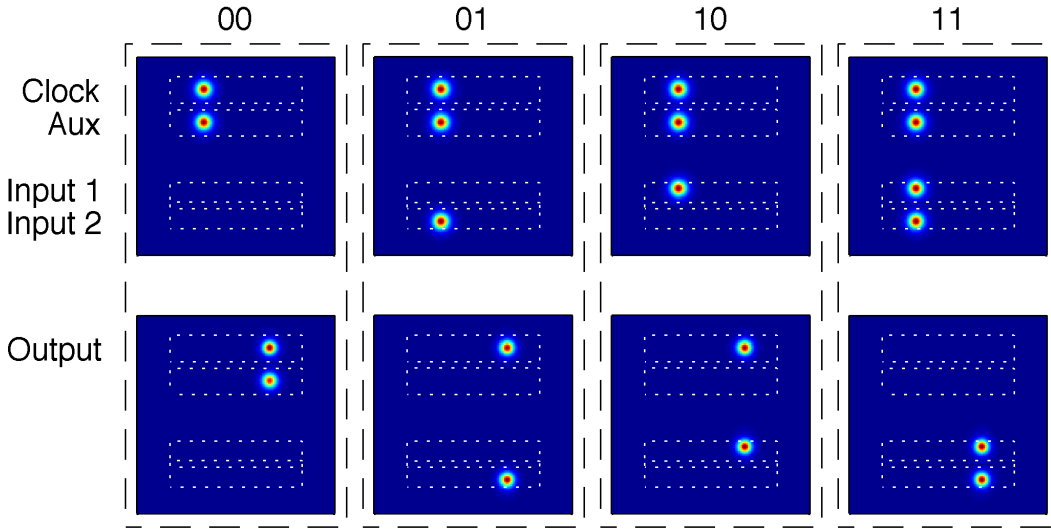


Figure 2.18: Four possible cases for the NAND gate. Note that each of the four possible bullets experiences a time periodic gain similar to Eq. (2.10) that translates the bullet from left to right. However, the overall gain given to each bullet, from lowest to highest, is: the auxiliary bullet, the clock bullet, the input one bullet, and the input two bullet. The NAND gate has enough gain to support a pair of bullets so if at most one of the inputs bullets exists the output is high. Otherwise, the output is low. The four cases here correspond directly to the four cases in the logic table of the NAND gate in Table 2.1

bullet receives less gain than either of the two input bullets, the result will be a NOR gate similar to what is shown in Fig. 2.17.

### *NAND Gate*

SWGAML can also be used to implement the NAND gate. In contrast to the NOR gate, this gate has been produced using time-dependent gains similar to Sec. 2.2.4. For the NAND gate, time dependent translating Gaussian profiles are used to move the clock bullet, the auxiliary bullet, and any of the input bullets. Figure 2.18 enumerates the four possible cases for the NAND gate. The clock bullet and auxiliary bullet must exist regardless of inputs, but there are up to two additional input bullets depending on the inputs to the gate. The gain profiles exist whether or not either of the two inputs exists. Therefore, this NAND gate does not require any prior knowledge about the inputs.

Like the NOR gate, the key to constructing a NAND gate is to provide insufficient gain

for all the bullets and to use differences in the gain to select the order in which bullets are destroyed. For the NAND gate, enough gain for two bullets is provided, and the level of gain provided to each of the bullets is from highest to lowest: input two, input one, the clock bullet, and the auxiliary bullet. If neither of the inputs exists, there is enough gain to support both the clock and auxiliary bullets. Therefore, they will translate to the output. Since the clock bullet reaches the output, the result is logical high. If one input exists, all three bullets will decay. As the auxiliary bullet has the lowest gain, it will be destroyed before the clock or input bullet. At that point, the clock and input bullets translate to the output and again logical high is the output. The final case is if there are two input bullets. In this case, the ordering of gain causes both the clock and auxiliary bullets to be destroyed resulting in the output being logical low. These four cases are consistent with the four possible input cases of the NAND gate given in Table 2.1.

In principal, the NAND gate could be constructed with the same static gain profiles the NOR gate was constructed with, and the NOR gate could be constructed using the same time-dependent gains the NAND gate was constructed with. Furthermore, an auxiliary bullet could be added to the NOR gate using the same gain ordering as the NAND gate but with only enough gain for a single bullet, which makes the physical structure of the NAND and NOR gates identical. This makes it possible to switch the gate from a NAND to a NOR by simply changing the overall level of gain supplied to the system.

Although other gates will not be constructed in this manuscript, the generation of the two master gates shows that fully operational photonic logic devices can be constructed. Since the output of each device is a light-bullet capable of serving as an input to an additional gate and since the NAND and NOR gates are insensitive to small delays in the input bullets, generating larger photonic circuits from these building blocks is trivial, and the other gates can be easily realized through combinations of NAND and NOR. Therefore, it is possible to do with photonics in the SWGAML any computation that could be done digitally.

### **2.3 Chapter Summary**

In this chapter, we have examined two types of patterns that form in systems related to nonlinear waveguide arrays: X-waves and light-bullets. As a result of these studies, we have demonstrated numerically that in the parameter regimes relevant to mode-locking, X-waves are not linear phenomena in the parameter regimes and time-scales that pertain to mode-locking. In the co-propagating case, the relative phase difference between the X-waves determines whether the interaction was attractive or repulsive. For the counter-propagating case, the interactions resulted in a measurable phase shift but no measurable time delay. Taken together, these results demonstrate that the interaction of X-waves in this regime is more akin to the interaction of solitons than it is to linear superposition.

On the other hand, as the higher dimensional analog of mode-locking, light-bullets hold tremendous potential as a critically important technology in the field of photonics and a number of research groups are focused on developing this technology [25,45–55,66]. There are numerous technological methods both proposed and realized for engineering and controlling light bullets, and our approach is certainly not the only viable option for producing light bullets. However, in this chapter we have presented a robust and experimentally realizable

possibility for light bullet formation, stabilization, and control using simple semiconductor waveguide arrays. We have numerically demonstrated that light-bullet solutions form from noise. Furthermore, the governing equations are posed in the lab and not in a traveling frame of reference, so non-uniform gain profiles can be engineered. This nonuniform gain profiles allow for the routing and control of bullets; as a result, the SWGAML architecture can be used to generate all-photonics versions of the master logic gates, and is a strong contender for a future technology for controlling all-optical data streams in a feasible and robust manner.

Although the reasons for studying them differ, both light bullets and X-waves demonstrate the importance of coherent structures in nonlinear optics. Indeed, to obtain the sorts of high optical powers that engineers desire, coherent structures are almost a necessity. However in the examples shown here, analytical solutions for solutions of interest to these problems are unavailable. At best, a numerical description of the solution at a few instances in parameter space can be obtained. In the following chapters, we will outline methods to exploit this fundamental, if sometimes difficult to describe, low-dimensionality in order to compute the bifurcation diagrams of systems similar to the coherent structures sought in this chapter.

## Chapter 3

**ANSATZ BASED REDUCTION METHODS IN NONLINEAR OPTICS**

In this chapter and the following one, we apply dimensionality reduction techniques to allow the study of solution branches and bifurcations that appear in nonlinear optics. There are a number of dimensionality reduction techniques that have historically been used in the realm of optics. Two of the most popular of these techniques are the variational reduction (see e.g., Malomed [70], Anderson [71], or Bale [33] and the references contained therein) and coupled-mode theory (see e.g., Yariv [62]). Although classical and effective, these reduction techniques are restricted by the need for both a fundamental understanding of the underlying physics and the ability to generate an effective solution ansatz. In this chapter, we consider a continuous wave (CW) semiconductor waveguide array. Although complex and even chaotic temporal behaviors have been observed in such devices, a plethora of experimental and theoretical work have demonstrated that the spatial profile remains relatively simple. As a result, coupled mode theory (CMT) can be applied to the system to aid in the study and enhancement of the performance of the semiconductor waveguide array.

Although such devices have shown great *theoretical* promise, experimental effort towards producing high-power waveguide array lasers have yielded mixed results, particularly when compared to competing technologies such as the master-oscillator power amplifier (MOPA). Later theoretical work provided insight into the poor performance of early waveguide array technology. In particular, Botez et al. [72] showed that standard high-index guiding arrays always favor the undesirable out-of-phase normal mode of operation. This led to the so-called leaky-wave-coupled device and a significant improvement in phase-coherent performance [72]. Despite that, waveguide arrays have yet to completely fulfill their promise. The current issue is that high injection levels lead to nonlinear effects that can destroy or degrade the fragile phase-locked state.

In order to characterize the impact of current injection and waveguide geometry, we explore the region of unstable behavior exhibited by a waveguide array quantum dot laser from both a theoretical and experimental standpoint. In what follows, we extend an existing model for the field and carrier dynamics in a waveguide array laser and demonstrate the accuracy of our extended model by comparing with experimental results over a broad range of parameter values. With the underlying model verified, we will then create a reduced order model based on coupled mode theory in order to characterize the route to chaos observed experimentally in this array.

The underlying theoretical description of the semiconductor array dynamics involves a model that couples the optical field propagation to the charge carrier density dynamics. Yariv [62] and others have proposed that the charge density has an influence on the index of refraction of the guiding region. The standard model considered in the literature is derived from a reduction of Maxwell's equations by assuming a transverse electric (TE) wave, the slowly varying envelope approximation, and the effective index method in succession [73].




Figure 3.1: A schematic figure of the waveguide array laser. The waveguide array contains five waveguides numbered by the proximity to the source of the pump current.




Figure 3.2: The transverse structure of the waveguide array. This is a five-region waveguide that contains, substrate, AlGaAs cladding, and the active region with the quantum dots.

Since the conjecture by Yariv, a large body of theoretical and experimental work has been produced in an effort to quantify the dynamics of stripe waveguide lasers [64, 73, 74]. As has been universally observed, the dynamics of even two coupled waveguides can produce complex and chaotic interactions over a large range of parameter space [64, 75–77]. Given the physical phenomena that can be manifest in a larger number of coupled waveguides [14–17], we revisit the coupled array scenario to investigate the nonlinear dynamics exhibited in the semiconductor laser array with a large number of waveguides. Furthermore, we apply a nonuniform injection current to the waveguide array system and investigate its behavior under those conditions.

### 3.1 Experimental Setup

A schematic of the five-emitter laser array is shown in Fig. 3.1 and in cross-section in Fig. 3.2. As shown in Fig. 3.1, the active region consists of a ten-fold stack of InGaAs quantum dot layers embedded in a GaAs waveguide that is sandwiched between p-doped and n-doped AlGaAs cladding layers. The 4.5 mm-long array was fabricated by use of standard photolithography and wet etch processing. Five  $2\ \mu\text{m}$ -wide waveguides were etched in the top cladding layer with a center-to-center spacing of  $8.0\ \mu\text{m}$ . Ti/Au and Ni/AuGe/Ni/Au were used to make the p-type and n-type contacts, respectively. The laser cavity was created by the cleaved facets, which were not coated. The device was mounted p-side up on a copper heat sink that was thermoelectrically cooled to  $12^\circ\text{C}$ .

The pump current was inhomogeneous because the current was injected from one side of the array, as shown in Fig. 3.1. The waveguide, labeled #1, closest to the contact wire had the largest injection current; this current decreases approximately linearly for the successive waveguides, with waveguide #5 receiving the least current. In the Au deposition process the wafer is usually vertical with respect to the deposition direction. However in this case, the wafer is tilted towards the deposition direction. As a result, one side of the waveguide side wall has a smaller angle with the deposition direction which results in a thinner Au film being deposited compared with the Au thickness in other areas. The thinner Au film creates a larger and non-negligible resistance between waveguides. Due to these resistances, when current is injected from one side of the waveguide array, each waveguide has a different current. Therefore, one practical method for controlling the current distribution is to control the tilt of the wafer during the deposition process.

Another more versatile approach is to use multiple contact wires and to attach a single contact wire to each of the waveguides. In that case, the amount of injection current into each of the waveguides can be individually controlled. These contact wires can either have separate current sources, or they can share a single current source using a current-divider-type mechanism to determine the distribution. The key idea is that the distribution of injection current and therefore gain can be accomplished without requiring physical alteration to the waveguide array itself. This method is more complex than the previous method and was not used in the experimental results that follow. However, it is a viable approach for controlling the current distribution that does not require re-fabrication of the array.

In the far-field of the emission from the array, the light from the individual emitters is naturally spatially isolated. The intensity dynamics were recorded with high-speed photodiodes and an oscilloscope.

### 3.2 Governing Equations

The governing equations for this waveguide array are an extension of those derived by Rahman and Winful [64]. Their model used the mean field approach to handle losses at the facets, and included the physical effects of guided mode propagation, coupling between waveguides via the evanescent field, spatial diffraction, gain, and anti-guiding due to the carrier density. Furthermore, they explicitly modeled the carrier dynamics which were subject to both spontaneous and stimulated emission as well as carrier diffusion. In this model, the coupling between the carrier field and the electric field is due solely to stimulated emission, for the electric field equation, and recombination, for the minority carrier density. In non-dimensional form, the governing equations of the optical field and carrier dynamics are given by

$$\frac{\partial U}{\partial \tau} = iC\Delta\eta_{\text{eff}}U + (1 - iR)UV + iL_p^2 \frac{\partial^2 U}{\partial x^2}, \quad (3.1a)$$

$$T \frac{\partial V}{\partial \tau} = p(x) - V - (1 + 2V)|U|^2 + L_e^2 \frac{\partial^2 V}{\partial x^2}, \quad (3.1b)$$

where  $U$  is the envelope of the (transverse) electric field,  $V$  is the charge carrier density,  $\tau$  is the normalized time, and  $x$  is the spatial dimension which, in keeping with previous

works, will remain dimensional. The parameter  $\Delta\eta_{\text{eff}}$  is the difference in the effective index of refraction inside and outside of a waveguide [73].  $R$  models the effect of carrier anti-guiding in the waveguide array.  $T$  and  $L_e$  account for the differences in the time- and length-scales between the field dynamics and the carrier dynamics. Lastly,  $p(x)$  accounts for current injection. The governing equations for this system are derived from a high-frequency asymptotic reduction of Maxwell's equations. In this reduction, a number of key approximations and/or assumptions are made. Specifically, it is assumed that the propagation is dominated by a TE mode. Furthermore, the high-frequency asymptotics result in the slowly varying envelope approximation. This final assumption reduces the propagation from two to one dimension by use of the effective index method [73]. The approximate governing equations are made non-dimensional according to the following scalings:  $U = (\eta_e^2 g \tau_s / 2c \eta_a)^{1/2} \psi$ ,  $V = \frac{1}{2} g \tau_p (N - N_{\text{th}})$ ,  $p(x) = \frac{1}{2} g \tau_p [j(x) \tau_s - N_{\text{th}}]$ ,  $T = \tau_s / \tau_p$ ,  $g = \Gamma a c \eta_a / \eta_e^2$ ,  $1/\tau_p = (c/\eta_e)[(\eta_c/\eta_e)(1-\Gamma)\alpha_c + (2/L)\ln(1/r)]$ ,  $N_{\text{th}} = N_0 + 1/g\tau_p$ ,  $L_p = (c\tau_p/2\eta_e^2 k_0)^{1/2}$ ,  $C = \Gamma c k_0 \eta_a \tau_p / \eta_e^2$ , and  $\tau = t/\tau_p$ .<sup>1</sup> Here,  $\psi$  is the transverse electric field, and  $N$  is the charge carrier density.

The physical constants used can be determined from experiment and the material properties of the waveguide array structure. Table 3.1 gives a complete characterization of the waveguide array considered in the experiments outlined in the previous subsection. The numbers given are in most cases similar to those used by Rahman and Winful [64]. The differences between the experiments presented here and those of Winful [64] are the use of quantum dots in the gain region and current-injection induced gain non-uniformity across the waveguides. In our case, the gain decreases approximately linearly from waveguide #1 to waveguide #5. There are also differences in the optical wavelength, the length of the waveguide arrays, and the index of refraction profile in the transverse direction. The values in Table 3.1 give the following:  $L_p \approx 7.5$ ,  $C \approx 8 \times 10^3$ ,  $\tau_p \approx 22$ , and  $T \approx 87$ .

### 3.3 Lasing Dynamics

To explore the lasing dynamics exhibited in the waveguide array, the governing equations, Eq. (3.1), are simulated with the waveguide parameters in Table 3.1. Following Rahman and Winful [64], these equations were evolved in time by use of a time-splitting procedure and transparent boundary conditions [67]. The split-time equations were then solved by use of the Crank-Nicolson method, which is highly efficient for both diffusion operators (Eq. (3.1b)) and systems related to the Fresnel equation (Eq. (3.1a)) [78].

Of greatest interest is the power output from each waveguide. By use of the experimental methods outlined earlier, the power from each waveguide was measured as a function of time and pump current. For the parameters listed in Table 3.1, numerical simulations were compared to experimental results. The parameter values used are experimentally realistic, though an average center-to-center waveguide separation of  $8.0 \mu\text{m}$  was found to give the best agreement between experiment and simulation. No additional physical effects were added to the system in order to fit the data.

---

<sup>1</sup>The last term in  $1/\tau_p$  is typically  $(1/L)\ln(1/r)$ , so  $N_{\text{th}}$  is the threshold charge carrier density defined by Yariv [62]. The extra factor of two here is for the non-dimensionalization of Eq. (3.1).

Table 3.1: Parameter Values for Waveguide Array Simulation

Parameter Description	Symbol	Units	Value
Wavelength	$\lambda$	$\mu\text{m}$	1.178
Waveguide Ridge width	$W$	$\mu\text{m}$	2.0
Active Layer thickness	$d$	$\mu\text{m}$	0.4
Cavity Length	$L$	$\mu\text{m}$	4500
Injection Current	$I$	mA	400
Cladding Loss	$\alpha_c$	$\mu\text{m}^{-1}$	$1.0 \times 10^{-3}$
Antiguinding Factor	$R$		3
Carrier Density for transparency	$N_0$	$\mu\text{m}^{-3}$	$5 \times 10^4$
Carrier Diffusion length	$L_e$	$\mu\text{m}$	2
Gain Coefficient	$a$		$2 \times 10^{-8}$
Active Layer Index	$\eta_a$		3.443
Cladding Layer Index	$\eta_c$		3.12
Effective Index	$\eta_e$		3.386
Insulator Index	$\eta_i$		1.45
Carrier Lifetime	$\tau_s$	ps	2000
Transverse confinement factor	$\Gamma$		0.75
End Mirror reflectivity	$r$		0.53

To begin, we consider the temporal power fluctuations in each waveguide of the array on the scale of a nanosecond. In Fig. 3.3, we show the oscillatory behavior obtained experimentally where the array does not converge to a steady state. The average power output is higher in the first waveguide than in the fifth waveguide. This is expected as there is a larger injection current in the first waveguide than in the fifth. Thus the first waveguide has a higher average power output than the second waveguide, and so forth. Due to what we believe are imperfections in the fifth waveguide, the trend fails to hold for the fourth and fifth waveguide at low injection currents.

For waveguide arrays, oscillatory behavior like that found experimentally is obtained from simulations of Eq. (3.1). However, the bias towards higher average powers in the first few waveguides cannot be obtained without the assumption that pump current is decaying as a function of waveguide number. To compare simulations with experiment, the largest pump current is assumed to be in the first waveguide, with a linearly decreasing pump strength into successive neighboring waveguides.

Figure 3.3 shows both the numerically computed and experimentally observed nanosecond temporal power fluctuations in each waveguide of the array. The results are qualitatively similar for the first four waveguides, but the final waveguide has a different behavior due to what we believe are fabrication imperfections. Regardless, the agreement is good, given that there are no fitting parameters except the decrease in current.

The decrease in current injection strength as a function of waveguide number is clearly




Figure 3.3: Experimental and numerical results measuring the time evolution of the output power for a five-emitter array with a pump current of 400 mA from waveguides 1-5 in (a)-(e) respectively. The power output is asymmetrical, which implies an uneven pump current into each waveguide. The frequency and magnitude of the numerical results (dashed) are similar to experiment (solid), with a slight frequency shift.

illustrated in Fig. 3.4(a). This figure relates the experimentally determined average output power to the total current injected into the waveguide array structure (see Fig. 3.1). The approximately linear decrease in injection current to the waveguides does not hold for the fourth and fifth waveguide. Regardless, the majority of the power, which is guided along waveguides one, two, and three, is subject to the approximately linear scaling.

When the injection current applied to each waveguide decreases linearly, the numerics differ from the experimental results in two key ways. The first is that the threshold current value of the emitters is not consistent across the array, and the second is the behavior of the fifth waveguide. By heuristically assuming the decrease in the injection currents depends upon both the emitter number and on the level of injection current, the results in Fig. 3.4(b) are obtained. In this plot, the fifth waveguide is receiving the mean amount of current injection in an effort to match results at lower injection current values, and thus emits far too much power at higher current levels. While this approach requires more assumptions than the linearly sloped case, and the behavior of the fifth waveguide is still both qualitatively and quantitatively different, a single threshold current is achieved for all emitters. Additionally, the gain in power is more similar to the experimental results than previously. Thus, by relaxing some of the assumptions made, it is possible to obtain a better match with experiment.

The waveguide spacing and pump current are two critical parameters in determining

the observed dynamics. Specifically, these two in combination impact the threshold current of each individual waveguide. A first observation suggested that if the separation between the fourth and fifth waveguides were less than  $8.0 \mu\text{m}$ , then the coupling between adjacent waveguides would be stronger. A second observation is that for very low pump currents, the difference in threshold currents in each waveguide plays a larger role in the dynamics of the system. The impact of these two observations is that a smaller waveguide separation leads to a stronger coupling between the fourth and fifth waveguide. This enables the fifth waveguide to turn on to a higher output power just above threshold compared with the fourth waveguide, as observed. However, this effect would vanish once injection currents were large enough to render the threshold current irrelevant. This would seem a plausible explanation for the genesis of the fifth waveguide imperfection. In practice, the fabrication of the waveguide array does indeed exhibit imperfections as observed under a scanning electron microscope. The imperfections can include the waveguide spacing, waveguide impurities, and gold deposition for the charge carrier injection. The effects of imperfections and characterization of the waveguide quality will be explored in future experimental work.

As is expected of the oscillatory behavior observed experimentally and theoretically, the phase difference between neighboring waveguides plays a role in the dynamics of the optical field. Experimentally, the phase difference of the output power oscillations of two adjacent waveguides is approximately  $\pi$  over the range of tested pump currents. This suggests the selection of an out-of-phase oscillatory behavior. Figure 3.5 shows the experimentally measured phase difference for pump currents ranging from 395 to 440 mA. The phase difference shifts for the fourth and fifth waveguide. This is consistent with our observation that the fifth waveguide is fundamentally different from the other four. Numerical simulations confirm these experimental findings. Our conjecture is that the out-of-phase oscillatory solutions are stable in relation to the in-phase oscillations for sufficiently well separated waveguides. In contrast, numerical simulations show that if the waveguides are placed too closely, the coupling strength is greatly increased and pairs of *in-phase* waveguides couple in an *out-of-phase* fashion with other pairs of *in-phase* waveguides. In essence, a pair of waveguides acts as a single large waveguide.

The phase difference remains relatively constant over a broad range of pump currents, but the frequency of the output power oscillation increases in direct relation to the injection pump current. Figure 3.6 depicts the experimental results comparing the frequency of the power versus the injection current. Note that all of the waveguides exhibit the same fundamental pattern. Thus, there is a direct relationship between the pump current and the frequency of power oscillation in the output. The dominant frequency increases from approximately 0.75 GHz to 1.25 GHz as current increases from 380 to 500 mA. Below 380 mA, the threshold current has not been reached so the output is simply noise. In addition to the fundamental frequency, the output generates higher harmonics, two of which are clearly seen in Fig. 3.6. These start at approximately 1.6 GHz and 2.3 GHz, respectively, when the injection current is approximately 400 mA. In addition to the higher harmonics, there are streaks of relatively high amplitude for various ranges of currents. These frequencies are generated from nonlinear processes in the waveguide array and the coupling to the carrier dynamics.

Figure 3.7 reproduces numerically the frequency of the power oscillations versus the injection current curves observed experimentally. The numerical simulation computes a




Figure 3.4: Average output power as a function of pump current in each of the waveguides obtained experimentally (a) and numerically (b). The gain in power is greatest in the first waveguide (+) and is reduced in the second (o), third (•) and fourth ( $\Delta$ ) waveguides. Experimentally, the fifth ( $\triangleright$ ) has a high gain at lower currents and a lower gain at higher currents, but this is not reflected in numerical simulation. Note that despite the current injection ramp, there is a uniform threshold across the waveguides at  $\approx 380$  mA.

fundamental frequency of oscillation in addition to the second and third harmonics. In both cases, the increase in the fundamental frequency exhibits roughly a square-root profile as a function of current. This also agrees with numerical simulations previously performed using lower dimensional ODE models [75]. Additionally, streaks of other frequency components are observed just as in the experiment.

The primary difference between the numerics and the experiment lies in the existence of the streaks of high amplitude that lie away from the fundamental frequency and its harmonics. Although these streaks are clearly observed in both cases, the structure appears to be different. Regardless, the theoretical model Eq. (3.1) captures the generation of frequencies due to the nonlinear processes in the waveguide array and carrier response, and so appears to work well in practice.

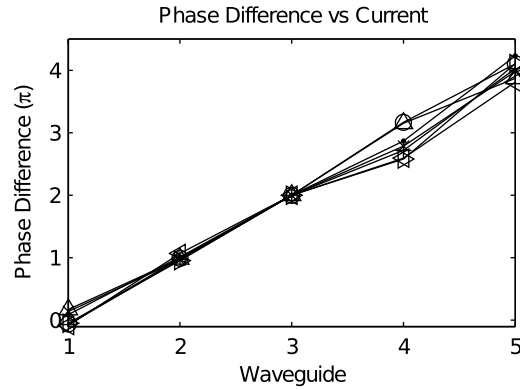


Figure 3.5: Measured phase difference in power between waveguides for pump currents of 395 mA ( $\bullet$ ), 400 mA ( $+$ ), 405 mA ( $\times$ ), 410 mA ( $\circ$ ), 420 mA ( $\triangle$ ), 430 mA ( $\triangleright$ ), and 440 mA ( $\triangleleft$ ). The phase change stays fixed near  $\pi$  for all currents except between the fourth and fifth waveguides where there is variation.




Figure 3.6: Experimental frequency spectrum of the power oscillations vs injection current. The fundamental frequency is directly related to the pump current and appears for all currents larger than 380 mA.

### 3.4 Reduced Dimensional Model

In Sec. 3.2 and 3.3, numerical simulations were performed using the partial differential equations in Eq. (3.1). While these types of simulations are ideal for matching experimental data and making quantitative predictions, the computational cost makes exploring parameter space cumbersome and time consuming. Following the approach of Winful [76], coupled




Figure 3.7: Frequency spectrum of the power oscillations versus injection current obtained from the simulation. The same pattern as in Fig. 3.6 appears with the fundamental frequency and its harmonics increasing with pump current. Below approximately 380 mA, there is no oscillation and thus no fundamental frequency.

mode theory is used to generate a reduced order model. This model, though only capable of yielding qualitative predictions, is useful for studying system dynamics and obtaining an intuitive idea of the response of the waveguide array to a non-uniform injection current.

Unlike in previous works, we do not assume that the gain provided to each of the waveguides is identical. Indeed, the results discussed in Sec. 3.3 suggest that the slope in injection current is a critical component of the observed dynamics. The reduced order model assumes the solution can be represented as a superposition of  $N$  modes:

$$A(x, t) = \sum_{j=1}^N a_j(t) \phi_j(x), \quad (3.2a)$$

$$V(x, t) = V_0(x) + \sum_{j=1}^N v_j(t) (V_j(x) - V_\infty). \quad (3.2b)$$

The individual  $\phi_j$  and  $V_j$  are determined by computing stationary solutions of Eq. (3.1) with a single waveguide. For the single waveguide problem, these solutions are stable and possess a large basin of attraction for physically realistic parameter values. As a result, evolving Eq. (3.1) from white initial conditions yields a good approximation of the solution to the nonlinear eigenvalue problem, which is ultimately solved by `bvp4c` in MATLAB. It should be stressed that the  $\phi_j$  and  $V_j$  are solutions to the single-waveguide problem for the electric field and charge carrier density using the physical parameters and injection current specific to each waveguide. In particular, the  $j$  denotes the waveguide number and not a supermode

(eigenfunction) of the waveguide array. Due to the nonlinearity introduced by the charge carriers, the supermodes of this waveguide array cannot be computed as would be done in the linear case. In this equation,  $V_\infty = -0.5g\tau N_{\text{th}}$  which is the non-dimensionalized charge carrier density arbitrarily far from the waveguides. This value is obtained by assuming that the density of physical charge carriers is zero. Due to the difference in the injection current combined with the nonlinearity of Eq. (3.1), it is necessary to recompute the mode for each of the five waveguides. In this array, the differences in injection current were relatively small, less than 15% difference between the center and outer waveguides. Therefore, continuation methods make computing branches of solutions in each of the individual waveguides numerically tractable.

The modes chosen are the stationary solutions of the single waveguide problem. The electric field is found from a simple linear superposition of the modes. The charge carrier equations are more complex. The solution  $V(x, t)$  is represented as a steady state charge distribution,

$$V_0 = V_\infty + \sum_{j=1}^N \beta_j (V_j - V_\infty), \quad (3.3)$$

and perturbations about that steady state value.

The steady state charge carrier distribution is approximated by assuming that the gain given to any of the modes is zero to leading order. Hence,

$$\int V_0 |\phi_j|^2 dx = 0, \quad (3.4)$$

which allows us to solve for the  $\beta_j$  in Eq. (3.3) and approximate the steady state. Using the basis functions in Eq. (3.2), the following low dimensional model is obtained,

$$\frac{da_j}{dt} = \frac{(1 - iR)}{\|\phi_j\|^2} a_j v_j + \sum_{k=1}^N \kappa_{j,k} a_k, \quad (3.5a)$$

$$\frac{dv_j}{dt} = \frac{1}{T} (P - (\xi_j + 2\gamma_j v_j) |a_j|^2 - (1 + \delta_j) v_j), \quad (3.5b)$$

where  $R$  and  $T$  retain their values from the PDE, and  $\|\phi_j\|$  is the  $L^2$ -norm of the  $j$ th mode.




Figure 3.8: Plot of the power in waveguides 1-5, shown in (a)-(e) respectively, using the reduced dimensional model. The same parameters and 400 mA injection current as Fig. 3.3 were used except the waveguide separation was taken to be  $9 \mu\text{m}$ .

The other parameter values are computed by computing overlap integrals,

$$\kappa_{j,k} = \frac{\int \phi_j^* \phi_k (C(\Delta\eta_0 - \Delta\eta_k) - (i + R)(V_0 - V_k)) dx}{\|\phi_j\|^2}, \quad (3.6a)$$

$$P_j = \int (p_0 - V_0 + L_e^2 \partial_x^2 V_0) |\phi_j|^2 dx, \quad (3.6b)$$

$$\gamma_j = \frac{\int (V_j - V_\infty) |\phi_j|^4 dx}{\int (V_j - V_\infty) |\phi_j|^2 dx}, \quad (3.6c)$$

$$\xi_j = \int (1 + 2V_0) |\phi_j|^4 dx, \quad (3.6d)$$

$$\delta_j = -L_e^2 \frac{\int \partial_x^2 V_j |\phi_j|^2 dx}{\int (V_j - V_\infty) |\phi_j|^2 dx}, \quad (3.6e)$$

where physical parameter values similar to those in Table 3.1 were used. The derivation of these parameter values is identical to that of Ref. [76], but unlike that work it is not possible to normalize the  $\phi_j$  by scaling the modes as the modes in each waveguide are not identical.

The results from the coupled mode theory qualitatively match the results of the PDE. Figure 3.8 shows the results of the coupled mode theory model using the same physical parameter values and 400 mA injection current as in Fig. 3.3 except that the waveguide spacing was taken to be  $9 \mu\text{m}$ . In both the reduced model and the full PDE model, the output of the system is oscillatory. In addition, the frequencies of oscillation match well. The reduced order model is also capable of operating in a steady-state and chaotic fashion so it can recover the three broad types of behaviors observed in the PDE. However, the effect of coupling in the reduced model appears to be stronger than the effect of coupling in the

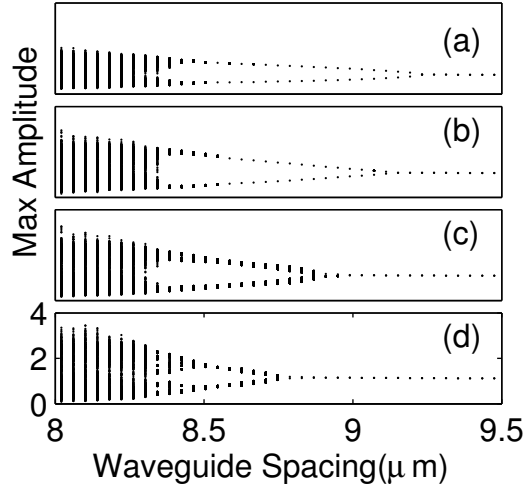


Figure 3.9: Plot of the local maxima and minima of pulse amplitudes for the five-waveguide system with a gain slope of (a) 0 (b) 0.06 (c) 0.12 (d) 0.18 for different waveguide separations. The route to chaos persists, but the location of the Hopf and period-doubling bifurcation changes.

full PDE, which is why the waveguide separation was increased. This is likely due to the assumption that the fields in the waveguide array can be represented as a superposition of modes. In the PDE, the fourth and fifth waveguide have almost no energy and violate this assumption. Therefore, the addition of those modes causes the waveguides to interact more strongly than what occurs in the PDE. Because the reduced order model is able to reproduce the dynamics of the full PDE, it is effective as a qualitative representation of the dynamics in the waveguide array, but because the value of waveguide separation at which these dynamics occur is different, it can not be used in making quantitative predictions about the model.

Using the parameters in Table 3.1, evenly spaced waveguides, and a linear injection current slope, the reduced order model was evolved forward for long periods of time with a total injection current of 450 mA. This gives sufficient time for the system to be drawn to any attractors that may exist at these parameter values. From there, the system was evolved further in time and the output sampled. In order to gain a sense of the system dynamics, local extrema of the amplitude were found. These extrema give a qualitative sense of the dynamics of the system for different gains.

The type of solution is determined by the separation between waveguides. As stated before in the two-waveguide uniform gain case [76], for large separations the interaction between adjacent waveguides is weak and the stationary solutions result. From a stationary solution, MATCONT [1] was used to continue the solution and locate bifurcations using waveguide separation as the bifurcation parameter.

As the separation decreases, the interaction between adjacent waveguides causes the stationary solution to lose stability through a Hopf bifurcation. This Hopf bifurcation creates

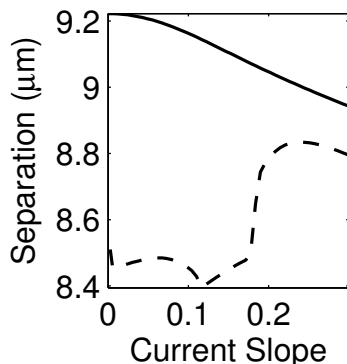


Figure 3.10: Plot of the bifurcation locations as a function of waveguide separation for the five waveguide system. The solid line is the location of the Hopf bifurcation while the dashed line is the location of the bifurcation of the limit cycle which is either a period doubling bifurcation, a Neimark-Sacker bifurcation, or a fold of the limit cycle.

solutions with time-oscillatory amplitudes such as the ones found in the PDE in Sec. 3.3. These periodic solutions persist for a relatively large range of waveguide separations, but they also become unstable when the waveguide separation becomes too small.

The loss of stability of the limit cycle is dependent on the slope of the gain. With two waveguides and no gain slope, a period doubling bifurcation is responsible for the loss of stability [76]. When the gain slope changes, the bifurcation changes from a period-doubling bifurcation to a Neimark-Sacker (torus) bifurcation. For five waveguides, a combination of period-doubling, Neimark-Sacker, and folds of the limit cycle result. After these bifurcations, the system was not continued further, but as shown in Fig. 3.9 the system still continues to chaos. With more degrees of freedom or at certain slopes, the route to chaos is more complicated than the classical period doubling path. Therefore, while the system appears qualitatively similar regardless of the number of waveguides in the system, the dynamics become richer with the introduction of additional waveguides and with non-uniform gain.

From a technological standpoint, non-uniform gain profiles can be used to change the bifurcation point and fundamentally alter the character of the solution. Figure 3.10 shows the waveguide separation where the Hopf bifurcation and the separation where the period doubling bifurcation occurs for the two waveguide problem. The location of the Hopf bifurcation can be shifted by up to  $0.2 \mu\text{m}$  simply by changing the gain slope. Note that the change in the bifurcation location is not monotonic, and for very large gain slopes the Hopf bifurcation again begins to occur at larger waveguide separations. This, however, requires current slopes larger than the ones used in Fig. 3.10.

A similar process occurs for the location of the period-doubling bifurcation. The rapid changes that occur in Fig. 3.10 are not single isolated points. A finer mesh than normal was used to more fully resolve those regions. It is possible some sort of resonance occurs in the system which causes the rather rapid change in location of the bifurcation. It was observed that for two waveguides, these kinks occur where the bifurcation switches from a period

doubling bifurcation to the Neimark-Sacker bifurcation. In the five waveguide system shown in Fig. 3.10, the kinks that occur are where the bifurcation switches from period-doubling to a limit point of the cycle or to a Neimark-Sacker bifurcation. Together, this gives further evidence that some sort of resonance is responsible for the dramatic change in the character of the solution. Nonetheless, it is precisely this region where the system goes to chaos at smaller values of  $S$  than usual.

The result of these observations is that it is possible to create tunable arrays without requiring extra current to be injected into the array. In both Fig. 3.9 and Fig. 3.10, the total current injected into the array remained constant, but the location of the Hopf bifurcation could be shifted by around  $0.2 \mu\text{m}$ . This allows for an array in the stationary regime to be shifted into the oscillatory regime or vice versa. Additionally, the same technique could be used to shift an array operating in the chaotic regime to the oscillatory regime. Furthermore, this requires no physical alterations to the experimental setup as gain can be controlled completely via external electronics.

### 3.5 Chapter Summary

In this chapter, we demonstrated one of the standard methods of dimension reduction in optics, coupled mode theory. Using CMT, we have extended the work of Winful [76] to include larger semiconductor waveguide arrays where each waveguide possesses a unique geometry and level of injection current. Furthermore, we verified the validity of our extended model by comparing the numerical results with experimental results. In both experiment and theory, stable and robust oscillatory operation is produced within the waveguides, with adjacent waveguides having an approximately  $\pi$  phase shift in their output power oscillations. The frequency of oscillation increases proportionally with the level of injection current. Overall by comparing theory and experiment over a wide range of injection currents, we have demonstrated that our extensions of the Winful model match well with physical reality.

With the reduced order model, we first determined the route to chaos that occurs in a larger five-waveguide array when previously only two or three waveguide systems had been considered. Second, by including the effects of a non-uniform gain and waveguide profile, we determined that the non-uniformities of the array serve as a critical operating parameter. In particular, distribution of current injection is useful because it acts as a proxy for waveguide spacing, i.e. it can push the system from stable, to oscillatory, and finally to chaotic behavior. Furthermore, it is one of the few system parameters that can be altered without impacting the physical structure of the array. As a result, it is a key degree of freedom in the design and construction of such arrays that has previously been overlooked and creates the possibility of designing tunable arrays without the need for changing the overall level of injection current or output power.

## Chapter 4

**THE MULTI-PULSING TRANSITION**

The multi-pulsing transition is a mechanism by which the single-pulse solution of a mode-locked laser becomes unstable and, after a sequence of bifurcations, ultimately breaks apart into two, three, and eventually an arbitrarily large number of pulses. The multi-pulsing instability that initiates this transition appears to be ubiquitous in all mode-locked systems, and has been observed both in experiments [79,80] and in numerical simulations [10,79]. The WGAML model and the two dimensional SWGAML model both exhibit this transition, as was demonstrated by Kutz and Sandstede [31] and in Chapter 2. From a physical perspective, the presence of the multi-pulsing instability is critical as it limits the peak power any single mode-locked pulse can have.

Despite the importance of this transition, it remains relatively unexplored. To our knowledge, only a limited theoretical description based on energy arguments [81,82] and a reduced order model for a non-chaotic transition [10] exist. In this chapter, we seek to explore the fundamental nature of a more complex and chaotic multi-pulsing transition via a reduced order model using the traditional numerical methods of dynamical systems and via a specialized method for computing periodic solutions in PDEs. Due to the complex and changing spatial (which, in optical coordinates, is  $T$ ) structures that occur during the multi-pulsing transition, ansatz based methods, such as those in Chapter 3, are not applicable. Instead, methods that are based on data are required. In Sec. 4.1, the multi-pulsing transition is studied using the POD method discussed previously in Sec. 1.2. In order to verify these results, in Sec. 4.2, we use the PDE based adjoint continuation method (ACM) to determine the branches of solutions and their stability without assuming the presence of low dimensional dynamics.

**4.1 The Multi-Pulsing Transition with Reduced Order Models**

In this section, we derive a reduced order model based on the Proper Orthogonal Decomposition (POD), as was discussed in Sec. 1.2, for the WGAML model. We will then apply the standard numerical tools of dynamical systems, such as MATCONT, in order to trace out the branches of solutions and bifurcations that take part in the multi-pulsing transition from one to two pulses. In the following sections, we outline the different qualitative types of dynamics observed in the transition. Then, we derive the framework of a reduced order model for WGAML. Because the validity of the resulting model depends heavily on the data chosen, we demonstrate the dynamics of the reduced order models generated with different data sets. Lastly, we use these models to compute the branches of solutions and bifurcations that occur during the transition from a single-pulse solution to a double-pulse solution.




Figure 4.1: Surface and pseudo-color plots of possible behaviors in the WGAML system. From left to right, a single pulse, a breather, a chaotic, and a two pulse solution are shown. The particular behaviors are determined by the value of  $g_0$  which is 2.3, 2.5, 2.6, and 2.7 for the single-pulse, breather, chaotic, and two pulse solutions respectively.

#### 4.1.1 Prototypical Dynamics in the Multi-Pulsing Transition

There are four prototypical behaviors that occur in the WGAML model in different parameter regimes. Because it is meant to model a physical system, in practice many of the parameter values in the WGAML model are fixed. For consistency with the previous works of Proctor and Kutz [30,31] and to ensure physical relevance, we set

$$(D, \beta, C, \gamma_0, \gamma_1, \gamma_2, \tau, e_0) = (1, 8, 5, 0, 0, 10, 0.1, 1), \quad (4.1)$$

for the remainder of the chapter. Physically, each of these values refer to some inherent material property or capture the waveguide geometry. The only free variable,  $g_0$  in Eq. (1.5), controls the gain supplied to the system. This can be increased experimentally through more vigorous optical pumping for erbium doped fibers or increased levels of carrier injection for semiconductor devices. As a result,  $g_0$  will be the bifurcation parameter in what follows.

Previous works have shown that for the fixed parameters in Eq. (4.1) and  $g_0 = 2.3$ , the WGAML model produces a single-pulse mode-locked solution starting from noise. These were later shown to be linearly stable through the application of the Floquet-Fourier-Hill (FFH) method by Jones and Kutz [83]. The FFH method is spectrally accurate and has been used to determine the spectrum of the linearized operator to a high degree of accuracy. In particular, the pair of zero eigenvalues stemming from the phase invariance and translational invariance have been reproduced numerically with an error less than  $10^{-13}$ . Through direct numerical simulations, breather, chaotic, and double-pulse solutions are also known to exist with larger values of  $g_0$ .

Figure 4.1 summarizes the different possible behaviors in the laser system ordered in terms of increasing gain. The first column on the left shows a single-pulse solution obtained at  $g_0 = 2.3$ . This solution is generated from an initial condition of low-amplitude white-noise. Although the solution is stationary in magnitude, there is an overall phase evolution that occurs. Increasing  $g_0$ , the behavior of the WGAML changes, and the stable solutions

become breather solutions as shown in the second column from the left. The breather solutions are symmetric in  $T$  and the maximum magnitude of their peaks is larger than that of the single-pulse solutions. Jones and Kutz [83] demonstrated that these breather solutions are created by a Hopf bifurcation of the single-pulse solution branch, and that the unstable eigenfunctions are even. When  $g_0$  is increased further, chaotic solutions are observed. An example of such a solution is shown in the third column. For some values of  $Z$ , the chaotic solution resembles the single-pulse solution, but for most values of  $Z$  it resembles a double-pulse solution. We refer to these solutions as chaotic because for certain values of  $Z$  the solution varies between the single-pulse and the double-pulse configurations in an intermittent and chaotic fashion. Further, this intermittency appears to drive various translations of the mode-locked pulse in  $T$ . At higher gain, the double-pulse solution is stabilized, as demonstrated in the right column of Fig. 4.1 with  $g_0 = 2.7$ . Notice that the double-pulse solution is typically not symmetric about the origin due to the translational invariance in  $T$  and the starting initial data. However, the translational invariance of the WGAML model guarantees the existence of a symmetric (in amplitude) double pulse solution. A similar transition from a double-pulse solution to a triple-pulse solution has been observed as well. Such findings suggest that the transition described here is a generic mechanism for the multi-pulse transition from the  $N$ - to  $(N + 1)$ -pulse solution [31].

#### 4.1.2 The Low Dimensional Model

While the phenomenology of the multi-pulse transition can be summarized by the four behaviors described in Fig. 4.1, the bifurcation sequence that connects these behaviors is unknown. In previous works [31, 83], a computational study of the spectrum of the linearized system around the single-pulse solution showed that the transition is initiated by a Hopf bifurcation, where a stable periodic solution (breather) is created. However, the complete sequence and in particular the onset of chaos has not been determined. This is mainly due to the difficulty in studying the stability of non-stationary breather solutions. To overcome the difficulty of studying solutions of an infinite dimensional system and their corresponding infinite dimensional linearized operators, we construct a reduced order model that will represent the dynamics of the full WGAML model.

To construct this model, the POD method in Sec. 1.2 is used, but because the WGAML model is a system of PDEs a slight extension of the method is required. For the WGAML model, the matrix  $\mathbf{X}$  is composed of numerical solutions of the WGAML model taken from *a single waveguide at evenly spaced values of  $Z$  and  $T$* . The SVD is then performed on each data set individually and the largest  $P$  modes, i.e. those that contain the most energy as characterized by their singular values, are retained. Because there are three waveguides in

this system, the solutions to the WGAML model are approximated by

$$A_0(T, Z) = \sum_{j=1}^P a_j(Z) \phi_{0,j}(T) \quad (4.2a)$$

$$A_1(T, Z) = \sum_{j=1}^P b_j(Z) \phi_{1,j}(T) \quad (4.2b)$$

$$A_2(T, Z) = \sum_{j=1}^P c_j(Z) \phi_{2,j}(T) \quad (4.2c)$$

where  $\phi_{i,j}$  is the  $j$ -th POD mode from the  $i$ -th waveguide rather than simply by Eq. (1.3). The Galerkin projection is performed on each of the waveguides individually which yields the system of equations,

$$\begin{aligned} \frac{\partial a_n}{\partial Z} = & \left( i \frac{D}{2} + g\tau \right) \sum_{j=1}^P \langle \phi_{0,n}, \partial_T^2 \phi_{0,j} a_j \rangle + (g - \gamma_0) a_n \\ & + iC \sum_{j=1}^P \langle \phi_{0,n}, \phi_{0,j} b_j \rangle + i\beta \sum_{j=1}^P \sum_{k=1}^P \sum_{m=1}^P \langle \phi_{0,n}, \phi_{0,k} \phi_{0,j} \phi_{0,m}^* \rangle a_j a_k a_m^*, \end{aligned} \quad (4.3a)$$

$$\frac{\partial b_n}{\partial Z} = iC \sum_{j=1}^P (\langle \phi_{1,n}, \phi_{0,j} \rangle a_j + \langle \phi_{1,n}, \phi_{2,j} \rangle c_j) - \gamma_1 b_n, \quad (4.3b)$$

$$\frac{\partial c_n}{\partial Z} = iC \sum_{j=1}^P \langle \phi_{2,n}, \phi_{1,j} \rangle b_j - \gamma_2 c_n, \quad (4.3c)$$

where  $\langle u, v \rangle = \int_{-\infty}^{\infty} u^* v dT$ , and

$$g = \frac{2g_0}{1 + \sum_{j=1}^P |a_j|^2 / e_0}. \quad (4.4)$$

The benefit of treating each waveguide independently is that this allows the POD basis to more easily account for the shifts in the relative phases between the waveguides that occur as  $g_0$  changes.

The form of the system of differential equations obtained in Eq. (4.3) using the Galerkin projection is generic for any orthonormal basis and does not depend on the particular modes generated by the POD. However, the inner products in Eq. (4.3) are critical in determining the dynamics that may be reproduced by this finite dimensional system and whether they will approximate the full WGAML model dynamics. In what follows, we construct the POD modes for each of the four prototypical behaviors observed in Fig. 4.1, and compare the solutions of the resulting reduced model with the full WGAML dynamics.

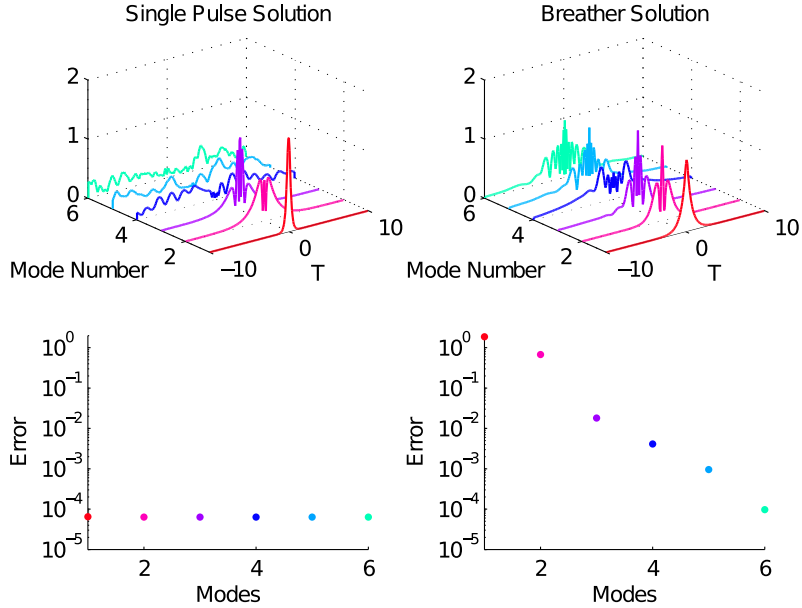


Figure 4.2: Plots of the POD mode profile and the error as defined in Eq. (4.5) using the first  $N$  POD modes in a low dimensional model. Left: the POD modes of the single-pulse solutions for  $g_0 = 2.3$ . Right: the POD modes of the breather solutions for  $g_0 = 2.5$ . For the breather solutions, the error was measured at the point of maximum amplitude for both the PDE and ODE solutions due to differences in the period of the limit cycles.

### *Single-Pulse and Breather Solutions*

For the single-pulse solution, the POD modes are generated by solving the WGAML model for a relatively short  $Z$  ( $Z \in [0, 10]$ ) with  $\Delta Z = 5 \times 10^{-3}$  and  $T \in [-10, 10]$  with  $\Delta T = 20/512$  at  $g_0 = 2.3$  with an initial condition that is a small perturbation away from the single-pulse solution. For each waveguide, the data matrix  $\mathbf{X}$ , which contains a snapshot of the PDE solution in each of its columns, is of size  $512 \times 2000$ . Without the perturbation, all the energy will be contained in the first POD mode that is the single-pulse solution illustrated in Fig. 4.1. Starting from a perturbed solution reveals that the additional POD modes participate in the WGAML dynamics as the system settles to the attractor. The POD modes of the 0th waveguide generated using this methodology are shown in the first column of Fig. 4.2. The first three modes contain over 99.9% of the energy in the system with almost all of this energy in the first mode. The second row in Fig. 4.2 shows the error between the reduced POD system and the WGAML solutions defined as a function of  $P$  where

$$\text{Error}(P) = \|A_{0,\text{PDE}} - A_{0,\text{POD}}(P)\|^2, \quad (4.5)$$

evaluated at  $Z = 500$ . The comparison was performed using similar initial conditions (noise) and identical  $g_0$ . For the single-pulse solution, the accuracy does not measurably improve as

more than a single mode is introduced, and the error remains on the order of  $10^{-5}$ . Thus, a single POD mode is able to accurately capture the mode-locking behavior in this regime. Although the inclusion of additional modes does not increase the accuracy, they do increase the basin of attraction of the single-pulse solution.

The dynamics of the breather solution are more complicated and require additional POD modes to represent. Similar to the single-pulse solution, the six POD modes shown in the second column of Fig. 4.2 capture over 99.9% of the energy; unlike the single-pulse solution, the energy of each of the six modes is non-negligible, and each mode has a clear structure in  $T$ . As  $P$  grows the additional modes are more oscillatory. Comparing the WGAML dynamics to the POD dynamics for  $P = 1$ , the reduced model predicts, at leading order, stationary single-pulse dynamics. For  $P \geq 2$ , the low dimensional model will yield a periodic orbit, in agreement with the WGAML dynamics. With  $P \geq 3$ , there will be a qualitative agreement between the shapes of the orbits. As more modes are added to the system, the results do not change qualitatively, but the reduced model becomes more accurate. To compute the error between the periodic solution of the WGAML dynamics and the POD dynamics, the value of  $Z$  was chosen such that both the ODE and PDE solutions were at their peak amplitude in the 0th waveguide. Due to the difference in the periods between the two solutions, this yields a tractable and reasonable measure of the error between them. Selecting other segments of the period to compare yields different numbers but a similar overall trend in the error. The error is shown in the bottom right plot in Fig. 4.2. For  $P = 1$ , the error is of order one. With three modes the error is reduced to  $10^{-2}$ , and additional modes cause a small but nonzero decrease in the error to  $10^{-5}$ . Therefore we conclude that, in order to obtain qualitative agreement with the PDE, a three-mode reduced order model in either regime should be sufficient.

### *Chaotic POD Modes*

The number of modes required to capture a certain fraction of the total energy is highly dependent on the nature of the underlying attractor. While the reduced order model is able to qualitatively and quantitatively reproduce solutions where the amplitude and phase are periodic in time with only a few POD modes, the same is not true when the WGAML model solutions are translating in  $T$  and when spatio-temporal chaos occurs. In particular, the chaotic solution in the WGAML, shown in the third column of Fig. 4.1, cannot be accurately represented with only a few modes. Figure 4.3 shows the first six POD modes obtained from a solution for the interval  $Z \in [0, 2000]$  with  $\Delta Z = 1$ . The sum of the energy of these six modes does not reach a significant fraction of the total energy. Indeed, we found that twenty-two modes are required to capture 99.9% of the total energy, unlike the one-mode or three-mode equivalent that occurs for the single pulse or breather solutions.

The necessity of using a large number of POD modes could result from any of three primary issues: the double pulse nature of the solution for some intervals in  $Z$ , the chaotic behavior of the solution in  $Z$ , and/or the translation of the solution in  $T$ . The right plot in Fig. 4.3 shows the POD modes of the double-pulse solution. We observe that only a pair of modes is needed to capture 99.9% of the energy in the data, and therefore the double pulse nature of the solution is not the cause of the large number of POD modes observed here.




Figure 4.3: Plots of POD mode profiles for reconstructing the chaotic and double-pulse solutions of the WGAML. Left: the first six of the twenty-two modes of the chaotic transition needed to capture 99.9% of the energy at  $g_0 = 2.6$ ; Right: the first six modes of the double-pulse solution at  $g_0 = 2.7$ . Only two of these modes are needed to capture 99.9% of the energy.

Instead, we must eliminate the chaotic translations.

The necessity of using a large number of POD modes is due to the translations of the pulses in  $T$  that occur in the chaotic regime. The right plot in Fig. 4.3 shows the POD modes of the double pulse solution, which requires only a pair of modes to capture 99.9% of the energy in the data. This demonstrates that there is nothing inherently high dimensional about multiple pulse solutions.

To evaluate the number of modes needed to reproduce the chaotic data without specifically considering the effects of translation, we constrain the solution to even functions (i.e.  $A(Z, T) = A(Z, -T)$ ). Since the WGAML model is invariant with respect to reflections about the origin, it can be restricted to a space of even solutions by choosing the initial condition as an even function. Simulating the WGAML dynamics in the chaotic regime with even initial condition reveals that the system is highly sensitive to odd perturbations. The sensitivity is expressed by the fact that a relatively small numerical error ultimately introduces a significant odd component in the solution. This suggests that the mechanism of chaos in WGAML is instigated by a symmetry breaking bifurcation which drives the translation in  $T$ . Then resulting solution can exhibit complicated behavior in both  $Z$  and  $T$  (as occurs in other systems [84]).

To consider the even WGAML model dynamics in simulations, the sensitivity to odd perturbations can be eliminated by enforcing equality between the  $k$ th and the  $-k$ th Fourier mode of the solution at every  $\Delta Z$  step (with the unpaired mode with the wavenumber  $-\frac{\pi N}{L}$  set to zero). In Fig. 4.4, we compare the number of modes needed to capture 99.9% of the total energy in the even WGAML model and the unrestricted WGAML model as the interval in  $Z$  is increased. The equations were evolved in increments of 100 units in  $Z$  taking 1000 samples per increment. At each increment a new set of POD modes were computed using the complete data from  $Z = 0$  to the current  $Z$  value. Note that this process leaves the previously computed modes unchanged. For the even WGAML model, we observe that




Figure 4.4: Comparison of the number of POD modes needed to capture 99.9% of the energy for the even-WGAML model and the unrestricted-WGAML model using data taken for  $0 \leq Z \leq Z_{\max}$ . The gain used for these results is  $g_0 = 2.65$  and  $g_0 = 3.1$  respectively.

the number of required modes to capture 99.9% of the energy is 15 modes for  $Z \in [0, 1000]$ . This number increases to 16 when the interval of  $Z$  is increased further. Chaos in  $Z$  does not increase the number of modes participating in the dynamics significantly when the even WGAML model is evolved for longer intervals of  $Z$ . However, the number of POD modes involved in the dynamics grows steadily as a function of  $Z$  for the unrestricted WGAML. The unrestricted WGAML required only 12 modes for the first 100 units in  $Z$ . As more data is added, this jumps to 18 and then to 32 modes at  $Z = 2000$ . The jumps occur when the chaotic solution exhibits rapid translations in  $T$ . Because it eliminates these translations, the even-WGAML provides an ideal framework for extracting POD modes from the chaotic regime. Therefore, we will restrict the WGAML model to even solutions. We expect that, although simplified, this model will reproduce the transition from single-pulse to double-pulse solutions and reveal the mechanism that leads to the chaotic behavior in  $Z$  during the transition.

### *The Global Model*

In previous sections, the POD modes were computed for solutions of the WGAML model with a single fixed value of  $g_0$ . Although the reduced order model obtained by a Galerkin projection onto a small number of POD modes accurately reproduced the observed behavior, the validity and accuracy of the resulting model diminishes when  $g_0$  is varied. In addition, the accuracy of the model can be further reduced if a bifurcation in the WGAML model occurs and the truncated set of the POD modes does not include the POD modes of the bifurcating dynamics. To represent the entirety of the multi-pulsing transition, the reduced order model must be valid for a range of  $g_0$  and capture the bifurcations in the transition. Therefore, multiple runs, each with various values of  $g_0$ , need to be combined into a single amalgamated data set. Each individual run consists of a short-time, high-resolution solution,  $Z \in [0, 10]$  and  $\Delta Z = 5 \times 10^{-3}$ . The initial conditions chosen are a small perturbation away




Figure 4.5: POD modes taken from the combined data set for the 0th waveguide that are capable of qualitatively reproducing the dynamics observed in the full WGAML model.

from the attractor at the fixed value of  $g_0$  considered. By combining these individual runs into a single data set, the SVD produces the set of POD modes that best captures the energy of the global dynamics through the transition from one- to two-pulses.

Figure 4.5 shows the first six POD modes of the 0th waveguide for the even-WGAML model. Note that there are two additional sets of modes, one for each of the other two waveguides. The data set includes information from the single-pulse, the breather, the chaotic, and the double-pulse solutions. These six modes capture over 99% of the energy, and 99.9% may be captured by including an additional four modes. It is observed that the resulting modes are not similar to individual modes from Fig. 4.2 or Fig. 4.3. Instead, the POD modes appear to be a nontrivial combination of the modes from all the runs at different gain values. These modes would be impossible to predict *a priori* even with knowledge about the solutions of the system.

#### 4.1.3 Low Dimensional Dynamics

In this subsection, the reduced order model is used to study the multi-pulse transition of the WGAML model. This model is obtained by the Galerkin projection of WGAML model in Eq. (1.5) onto the global POD modes from Fig. 4.5. This produces a system of differential equations of the form shown in Eq. (4.3). To determine the validity of the reduced model, we now compare the full (PDE) WGAML and the POD model dynamics for fixed values of  $g_0$  in the relevant ranges for single-pulse, breather, chaotic, and double-pulse solutions using low amplitude noise as the initial condition for both the full and reduced dynamics. The evolution considered is for a long period of time so that, for each value  $g_0$ , the system approaches the attractor. In this manner, the full and reduced dynamics of the WGAML starting from a cold-cavity configuration can be compared. In addition to comparing the solutions at a single value of  $g_0$ , we also compare the branches of single-pulse and breather solutions in both models.

The first row of Fig. 4.6 shows the single-pulse, breather, chaotic, and double-pulse solutions reconstructed using a six mode POD model. To obtain these solutions, the reduced

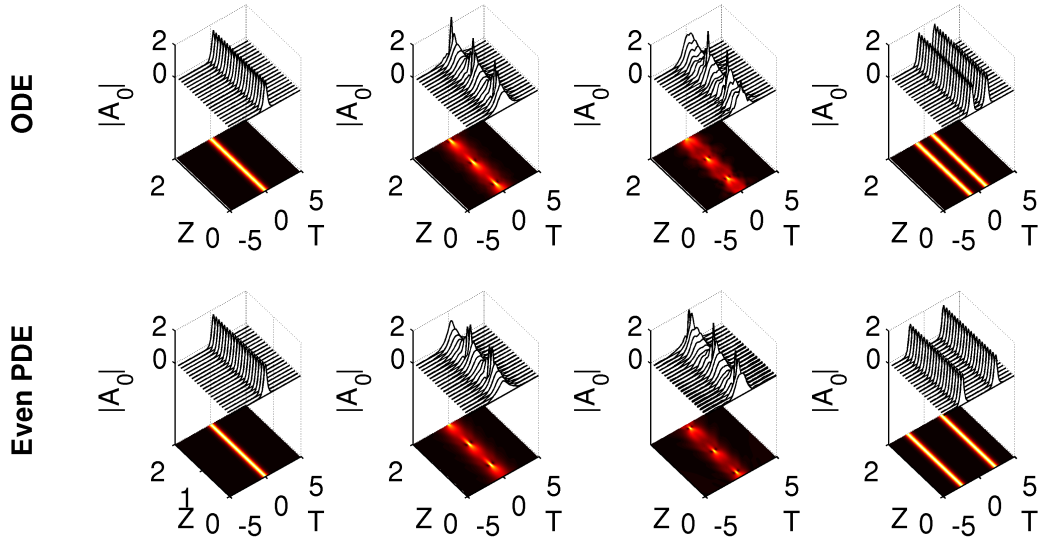


Figure 4.6: Top: Surface and pseudo-color plots of the single-pulse, breather, chaotic, and double pulse solutions computed for the finite-dimensional model at  $g_0 = 1.5, 2.5, 3.495,$  and  $3.5$  respectively. Bottom: The same plots for the even WGAML model taken at  $g_0 = 2.3, 2.5, 3.1,$  and  $3.2$  respectively. The reduced model accurately reproduces the four behaviors observed. These solutions should be compared to their equivalent behavior in the full WGAML model shown in the same order in Fig. 4.1.

model was evolved forward for 1000 units in  $Z$  starting with a low-amplitude white-noise initial condition. These results compare favorably with the same four regimes of the even-WGAML model shown in the second row of Fig. 4.6, whose solutions were obtained by evolving for 200 units in  $Z$  starting from low amplitude but even white-noise. The reduced model qualitatively captures the dynamics and the profile of the solution. The primary difference is the value of  $g_0$  at which these solutions are obtained. The stationary solutions of the POD model lose stability at lower values of  $g_0$  than the stationary solutions of the PDE. Furthermore, the breather solutions in the POD model remain stable for larger values of  $g_0$  than in the PDE. Due to the vastly smaller dimension of the reduced model and the range of  $g_0$  for which it is valid, some disparity between the models is inevitable. However, these results also compare well with the unconstrained PDE solution shown in Fig. 4.1 except for the translations in  $T$  that were disallowed by imposing evenness.

The advantage of the reduced order model over the PDE model is that the bifurcation diagram, including the stability and bifurcations of periodic solutions, can be explicitly calculated and categorized in the reduced model. The bifurcation software MATCONT [1] was used to track and compute the stability of the single-pulse and breather solutions. The chaotic and double-pulse solutions were obtained using direct numerical simulations. While




Figure 4.7: Top: The bifurcation diagram of energy ( $L^2$  norm) vs  $g_0$  for the multi-pulse transition in the POD model. Bottom: The bifurcation diagram of the multi-pulse transition of the WGAML model. The different plots show the same diagram with emphasis on different regions of the transition. For the stationary solutions, emphasized on the left, linearly unstable regions are dashed red lines while linearly stable regions are solid blue lines. Stationary solutions that were not computed explicitly are denoted by dots. For the periodic solutions, the local extrema of the energy are denoted by  $x$ 's .

other solution branches may exist, they do not appear as attractors for white-noise initial conditions and, as such, are not discussed here.

The first row of Fig. 4.7 shows the bifurcation diagram of the reduced order model including solution branches for the single-pulse, breather, chaotic, and double-pulse solutions. The branches marked in blue are linearly stable stationary solutions, and the branches marked in red are linearly unstable stationary solutions. For the periodic solutions, the branches marked by black  $x$ 's denote the extrema of

$$\|\vec{A}\|^2 = \|A_0\|^2 + \|A_1\|^2 + \|A_2\|^2, \quad (4.6)$$

where  $\|A_n\|^2 = \int_{-\infty}^{\infty} |A_n(T)|^2 dT$ . If no extrema exist, the mean value  $\|\vec{A}\|$  is denoted with a black dot. This bifurcation diagram reveals the sequence of bifurcations responsible for the multi-pulsing transition. At the lowest values of  $g_0$ , the only stable solution is the quiescent (trivial) solution as the gain is unable to overcome the cavity losses. The first non-trivial solution is the single-pulse solution, which first appears from the saddle-node bifurcation that is labeled (a) in Fig. 4.7 at  $g_0 = 0.7229$ . The value of  $g_0$  represents the minimum gain needed for a single-pulse solution to exist in the WGAML model. Although only the modes of the 0th waveguide are shown, these single-pulse solutions distribute energy in all

three of the waveguides, and the presence of energy in the other two waveguides is critical for the stabilization of these solutions. For larger values of  $g_0$ , there are two branches of single-pulse solutions that emanate from (a). The first branch is a high-amplitude stable branch of solutions, and the other is a low-amplitude unstable branch. Following the low-amplitude unstable branch of solutions is another saddle-node bifurcation, which is labeled (b) in Fig. 4.7. This bifurcation creates a stable branch of low-amplitude stationary solutions. Because their basin of attraction is small, these solutions are unlikely to appear for white-noise initial conditions. They are, however, known to exist in the full WGAML model [83]. This extra branch of stationary solutions loses stability through a Hopf bifurcation at the point labeled (c) in Fig. 4.7. Due to the low amplitudes and gains associated with this branch, it does not appear to play a role in the multi-pulsing transition.

On the other hand, the branch of larger amplitude stationary solutions is involved in the multi-pulsing transition. This branch of single-pulse solutions can be shown to be stable. Because we have assumed even-solutions, the zero eigenvalue associated with translational invariance no longer exists in the low dimensional system. On the other hand, the zero eigenvalue associated with phase-invariance persists. Although the eigenvalues differ, the stability results agree qualitatively with the full results obtained with the FFH method [83]. For a range of  $g_0$ , all other eigenvalues have negative real parts. As  $g_0$  increases, this branch becomes unstable through a supercritical Hopf bifurcation, which is shown in Fig. 4.7 at (d). As a result, the single-pulse solution becomes unstable but a stable limit-cycle is generated. These limit cycles are the breather solutions of the WGAML model.

As  $g_0$  is increased further, the breather solutions themselves lose stability through two bifurcations in rapid succession. The first is a period-doubling bifurcation. After the period-doubling bifurcation, the solution is still periodic in  $Z$  but with twice the period. In the reduced model, the Floquet multipliers can be computed explicitly, and one of the multipliers was found to be -1. The new, stable, period-doubled limit-cycles almost immediately lose stability through a supercritical Neimark-Sacker bifurcation. The Neimark-Sacker bifurcation, also called a torus bifurcation or secondary Hopf-bifurcation, occurs when a complex conjugate pair of Floquet multipliers leave the unit circle. In this case, the limit cycle loses stability but a torus that encloses the original limit cycle becomes stable. The Neimark-Sacker bifurcation indicates the presence of quasi-periodic solutions in the POD model. The period doubling and Neimark-Sacker bifurcations can be seen in the top right panel of Fig. 4.7 when the pair of extrema splits at the points labeled (e) and (f). MATCONT is unable to track branches of quasi-periodic solutions, but as shown in Fig. 4.7, further bifurcations occur and the ODE eventually exhibits chaos. At the same time, the double-pulse solutions gain stability and the system completes the transition from the single- to the double-pulse solutions.

The bifurcation diagram of the ODE shows remarkable qualitative agreement with the bifurcation diagram of the full WGAML dynamics illustrated in the second row of Fig. 4.7. Unlike quantitatively accurate solutions whose accuracy can be measured by the norm of the difference between the computed and exact solution, there is no good metric for measuring qualitative accuracy. In developing this reduced order model, the criteria we used was to take a representative sample of stable stationary or periodic solutions obtained in the reduced order model and use them as initial conditions for the PDE system which was subsequently

Table 4.1: The values of  $g_0$  associated with each bifurcation in the ODE and PDE. The labels correspond to the bifurcations of the ODE in Fig. 4.7. Values in parentheses are estimated from the same figure.

Label	Bifurcation Type	ODE	even-PDE
(a)	Saddle-Node	0.7229	0.6522
(b)	Saddle-Node	1.424	1.429
(c)	Hopf	1.398	1.369
(d)	Hopf	1.753	2.395
(e)	Period-Doubling	3.3872	(3.0±0.1)
(f)	Neimark-Sacker	3.3873	(3.0±0.1)

evolved in time. At each of these points, the long time dynamics of the PDE system were compared to those of the reduced order model. We found that for most values of  $g_0$ , the same dynamics occurred in both models, e.g., initial conditions that were stationary solutions in the reduced order model became, after an initial transient, stable stationary solutions in the PDE. Due to the differences in  $g_0$  where bifurcations occurred however, there exist regions in parameter space where this does not occur. For instance at  $g_0 = 0.7$ , a stable stationary solution in the reduced order model will transition to a stable periodic solution when used as the initial condition of the PDE. However, in each of these cases, small adjustments to the value of  $g_0$  will eliminate the discrepancy. Because the solutions of the reduced order model satisfied this criteria, we refer to them as qualitatively accurate. By increasing the number of POD modes used, say from 6 modes per waveguide to 12 modes per waveguide, the quantitative accuracy of the reduced order model can be increased.

Additionally, while the linear stability of solutions in the PDE has only been determined for the stationary solutions [83], for those solutions, the types of bifurcations that occur agree completely between the reduced and full models. Additionally, the value of  $g_0$  agrees relatively well for all of the bifurcations as shown in Table 4.1. Using standard methods, explicit stability results have not been obtained for the breather solutions, but the same qualitative bifurcation sequence occurs. The breather solutions lose stability around  $g_0 = 2.9$  which initiates the path to chaos in  $Z$ . Overall, the reduced order model is not accurate enough to quantitatively predict the dynamics of a particular waveguide array, but it is sufficient to determine the overall trends that appear in the multi-pulsing transition.

## 4.2 Verifying the Multi-Pulsing Transition

In the previous section, we proposed a mechanism for the multi-pulsing transition in a waveguide array mode-locked laser using the Proper Orthogonal Decomposition. The results of that study were verified by comparing, qualitatively, the predictions of the reduced model with observations from the full PDE. Although this informal check revealed good agreement between the reduced order model and PDE, we are unable to completely validate the model without directly computing solution branches in the full PDE.

For branches of stationary solutions, the computational cost of standard Newton or

Newton-Krylov methods is small enough that the solution branches can be computed on a standard laptop and to high accuracy in a reasonably short (on the order of an hour) amount of time. However, the same is not true for branches of time-periodic solutions. Orthogonal collocation methods, such as those AUTO and MATCONT use by default, result in overly large systems of equations when applied to nonlinear PDEs. Traditional shooting methods, on the other hand, have fewer degrees of freedom but require repeated and expensive computation of the monodromy matrix. To address this problem, a number of numerical methods have been developed. Among the most popular of these methods are the Recursive Projection Method (RPM) [85] and Newton-Picard methods [86] both of which were designed to obtain solution branches from large black-box systems.

In the WGAML model, the one- and two-pulse solutions have been computed in the work by Kutz and Sandstede [31], and the stability of these solutions has been obtained to high accuracy by Jones and Kutz [83] using the FFH method. As a result, the comparison with the reduced order model is straightforward to perform. To extend this sort of comparison to the periodic solution branches, we use a numerical method called the adjoint continuation method (ACM) to calculate arbitrarily accurate solutions and perform a PDE bifurcation study. For details on the derivation of the ACM as it applies to the WGAML model, see Appendix A. For a general overview of the method and its application to other physical problems, see the work of Wilkening [87–89].

In the remainder of the chapter, we will focus on the individual solution branches that are involved in the multi-pulsing transition. Because the stationary solution branches have been well studied in previous works, we focus solely on the periodic solution branches. In Sec. 4.2.1 we examine the branch of period-one breathers and the bifurcations that impact their stability. In Sec. 4.2.2, we focus on the torus bifurcation of the period-one branch. Although this does not impact stability (as the period-one breather solutions are already unstable), it is relevant to the reduced order model results from Sec. 4.1. In Sec. 4.2.3, the branch of period-two breather solutions is tracked and the loss of stability established. Finally, in Sec. 4.2.4 we discuss potential mechanisms for completing the transition to the double-pulse solution.

#### 4.2.1 *Period-One Breather Solutions*

The first non-trivial  $Z$ -periodic solutions are the period-one breather solutions. These solutions are created by a Hopf bifurcation of a single-pulse stationary solution at  $g_0 = 2.404$  in the full PDE. In order to apply the ACM, as outlined in Appendix A, an initial approximation of the limit cycle is required. From previous numerical experiments in Ref. [31] and Sec. 4.1 with the parameter values in Eq. (4.1), it was found that at  $g_0 = 2.5$  a periodic solution exists and forms from noise [30]. Starting with  $A_0(0, T) = \text{sech}(T)$  and  $A_1(0, T) = A_2(0, T) = 0$ , and evolving until  $Z = 2000$ , the system approaches that periodic solution. From this orbit, the period and phase offset were estimated and an approximation of the limit cycle with an error, as defined in Eq. (A.2), of  $G \sim 10^{-1}$  was obtained. The application of the ACM to this starting point reduced the error to  $G \sim 10^{-25}$  and produced the solution shown in Fig. 4.8.

Figure 4.8 shows the intensity of the fields in all three waveguides as a function of  $Z$




Figure 4.8: From top to bottom, plots of the amplitude of the 0th, 1st, and 2nd waveguides over one period at  $g_0 = 2.5$ . The solid black line shows the initial condition, and the dashed black line shows the solution after a half-period has elapsed. The green lines show the solutions at evenly-spaced intermediate values of  $Z$ . Note that the actual domain is from  $T \in [-20, 20]$  and only  $T \in [-1, 1]$  was plotted to highlight the region of interest.

and  $T$  for  $g_0 = 2.5$ . The black solid and black dashed lines show the solution at the start of the period and after half of the period has elapsed. Fifty-six intermediate values, evenly spaced in  $Z$ , are also plotted in green. This sample solution highlights the non-trivial nature of the breather solutions. Although the 0th waveguide retains a vaguely hyperbolic-secant shape throughout the period, the 1st and 2nd waveguides have at least two local maxima. Despite the more complicated physical representation, *the period-one breather solutions are even functions up to a translation in  $T$* . The periodic oscillation in  $Z$  can be explained by the interchange of pulse energy ( $\|A_0\|^2$ ,  $\|A_1\|^2$ , or  $\|A_2\|^2$ ) between the 0th and the outer two waveguides. For the majority of the period, the 0th waveguide has a low intensity. During this time, the gain remains unsaturated, and the 0th waveguide couples energy into the outer waveguides. At a certain point, this process is reversed and the outer waveguides couple energy back into the 0th waveguide. This causes the 0th waveguide to attain a large peak intensity, but it simultaneously saturates the gain. With the gain saturated, the large amount of energy in the 0th waveguide cannot be maintained. As a result, the energy in the 0th waveguide decreases and the process repeats.

Figure 4.9 shows the branch of period-one breather solutions. To obtain this branch of solutions, we used a simple continuation method. From the solution at  $g_0 = 2.5$ , we change the period,  $Z_0$ , by a small amount and repeat the minimization process as discussed in Appendix A. Linearly stable regions are shown in blue while linearly unstable regions are




Figure 4.9: (Left) A two-dimensional version of the bifurcation diagram showing the maximum  $L_2$ -norm over a period vs the gain. Stable solutions are plotted in blue and unstable solutions in red. The labeled bifurcation points are shared between the 2D and 3D plots. (Right) The bifurcation diagram of the period-one breather solutions with individual orbits are plotted in the  $(g_0, L_2, H_1)$  coordinate system. The plotted orbits represent one out of every forty periodic orbits computed on the branch. Linearly stable orbits are shown in blue and linearly unstable orbits are shown in red. The solutions at points (1) and (2) are shown in Fig. 4.11 while the solution at (3) is shown in Fig. 4.8.

shown in red. To compute the stability of these solutions, a modification of the FFH method, described in Section A.3, has been used. In order to visualize the branch, the branch has been plotted in  $(g_0, L_2, H_1)$  space where

$$L_2(Z) = \int_{-\infty}^{\infty} (|A_0(T, Z)|^2 + |A_1(T, Z)|^2 + |A_2(T, Z)|^2) dT, \quad (4.7a)$$

$$H_1(Z) = \int_{-\infty}^{\infty} (|\partial_T A_0(T, Z)|^2 + |\partial_T A_1(T, Z)|^2 + |\partial_T A_2(T, Z)|^2) dT. \quad (4.7b)$$

Please note that while  $L_2$  is the square of the  $L^2$ -norm,  $H_1$  is not the square of the  $H^1$ -norm because it lacks terms involving  $|A_0(T, Z)|^2$ ,  $|A_1(T, Z)|^2$ , and  $|A_2(T, Z)|^2$ . The bifurcation diagram is composed of roughly 2000 individual computations starting from  $g_0 = 2.5$ , which is labeled as (3). Each PDE computation was performed using 1024 Fourier modes to represent the solution in  $T$  with a domain size of 20 and 500 steps in  $Z$ . With this level of discretization,  $G = \|A_0(T, Z) - A_0(T, 0)\|^2 + \|A_1(T, Z) - A_1(T, 0)\|^2 + \|A_2(T, Z) - A_2(T, 0)\|^2 < 10^{-24}$  for all points on the bifurcation diagram, but most solutions had  $G \sim 10^{-25}$  or better. This corresponds to approximately 12 digits of accuracy.

There are three bifurcations that are critical for describing the period-one breather solutions: the subcritical-Hopf bifurcation that creates the branch, the saddle-node bifurcation of a limit cycle that restores stability, and the period-doubling bifurcation that breaks symmetry and changes the stability.

### *Subcritical Hopf bifurcation*

The subcritical-Hopf bifurcation, shown in Fig. 4.9 at (1), is where the single-pulse stationary solutions bifurcate and create the branch of period-one breather solutions. Because this is a subcritical bifurcation, for  $g_0 < 2.404$  we have both an unstable period-one breather solution and a stable single-pulse solution. Although the existence of the Hopf bifurcation has been initially shown by Kutz and Sandstede [30] and shown with high numerical accuracy by Jones and Kutz [83], it was not known whether the Hopf bifurcation was super- or sub-critical. The reduced order model of the previous section predicted this bifurcation should be super-critical, but by tracking the branch of periodic solutions back to the bifurcation point, we have demonstrated the bifurcation is in fact a sub-critical Hopf bifurcation. This highlights that data-driven reduced order models are heavily dependent on the supplied data, and without a complete and relevant set of data, their dynamics may not be even qualitatively correct.

### *Saddle-node bifurcation*

The next bifurcation that occurs is the saddle-node bifurcation of a limit cycle near  $g_0 = 2.325$ . This bifurcation restores the stability of the period-one branch of solutions as can be seen in Fig. 4.9. The segment of the branch between the saddle-node bifurcation and the period-doubling bifurcation around  $g_0 = 2.523$  is the only region where the period-one breather solutions are stable. This segment of the branch overlaps with the stable regime of the single-pulse solutions. Indeed, for  $g_0 \in (2.325, 2.404)$  there are at least three potential solutions, the stable stationary single-pulse solution, the stable high-amplitude breather solution (shown in blue), and the unstable low-amplitude breather solution (shown in red). In simulations starting from noise, the stationary solution is most likely to appear. This can be argued from an energy perspective. Of the three solutions available, the low amplitude breather solution has the lowest  $L_2$  value (cavity energy) at certain values of  $g_0$  but is unstable. Of the two stable solutions, the stationary single-pulse solution has a lower  $L_2$  value than the period-one breather solution for all  $Z$ . The single-pulse solution is more energetically favorable than the breather solution and therefore more likely to occur. With a similar argument, the stationary two-pulse solutions are also unlikely to form from noise. However because the single-pulse solutions, period-one breathers, and double-pulse solutions are all stable, there is a non-zero probability of any of those solutions appearing when starting from white-noise initial conditions.

### *Period-doubling bifurcation*

The final bifurcation shown in Fig. 4.9 is the period-doubling bifurcation that occurs at  $g_0 = 2.523$ . Figure 4.10 shows the unstable eigenfunction at  $Z = 0$  and the associated Floquet multipliers. The single unstable multiplier at  $\mu_{pd} = -1$  demonstrates that this is a period-doubling bifurcation [90]. Furthermore, the unstable eigenfunction associated with this bifurcation is odd. From Eq. (A.6), it can be shown that for any odd perturbation,  $\dot{A}_0(T, 0)$ ,  $\dot{A}_1(T, 0)$ , and  $\dot{A}_2(T, 0)$ , which is exactly what we have in computing the Floquet multipliers,  $\dot{A}_0(T, Z)$ ,  $\dot{A}_1(T, Z)$  and  $\dot{A}_2(T, Z)$  remain odd for all  $Z$ . Therefore, the unstable




Figure 4.10: On the left, the real and imaginary parts of the bifurcation eigenfunctions of the 0th waveguide are shown in blue and red respectively. These modes are associated with a period doubling bifurcation indicated by a Floquet multiplier crossing the unit circle at  $\mu_{\text{pd}} = -1$  in the right figure. Multipliers in or on the unit circle are indicated with a blue dot. Multipliers outside of the unit circle are indicated with a red  $x$ .

eigenfunction breaks symmetry at all points on the orbit. We believe that the symmetry breaking due to this bifurcation is responsible for the translating behavior that appears when the period-one breather goes unstable.

#### *Deformation of the period-one breathers*

In addition to the bifurcations, the deformations of the period-one breather solutions as we travel along the branch of solutions result in breathers that look qualitatively different at different values of  $g_0$ . Further from the Hopf bifurcation along the branch of solutions, the  $L_2$  values and the  $H_1$  values of solutions increase monotonically, and the solutions deform steadily from the hyperbolic secant solutions of Eq. (1.7) to the nontrivial breather solutions in Fig. 4.8 indicated by (3) in Fig. 4.9.

Figure 4.11 plots the solutions at (1) and (2) in Fig. 4.9 with  $g_0 = 2.398$  and  $2.37$ , respectively, and highlights the deformation of the breathers. The solution at  $g_0 = 2.398$  is a small perturbation away from the single-pulse solution and can be treated as the single pulse solution plus a small  $Z$ -periodic perturbation; thus, all three waveguides have intensities that resemble a hyperbolic secant. Even at this early stage, it is apparent that when  $\|A_0\|$  is maximal, both  $\|A_1\|$  and  $\|A_2\|$  take on their minimal values. Therefore, the dynamic interchange of energy between waveguides is already occurring, if only on a relatively small scale.

Further along the branch at (2) in Fig. 4.9, the magnitude of the oscillation has grown in all three of the waveguides. In addition, the breathing that occurs in the 2nd waveguide now includes the characteristic local minimum at  $T = 0$  that was observed in both the outer waveguides at  $g_0 = 2.5$ . These low-amplitude breather solutions provide the intermediate steps between the Hopf bifurcation and the breather solutions observed in previous

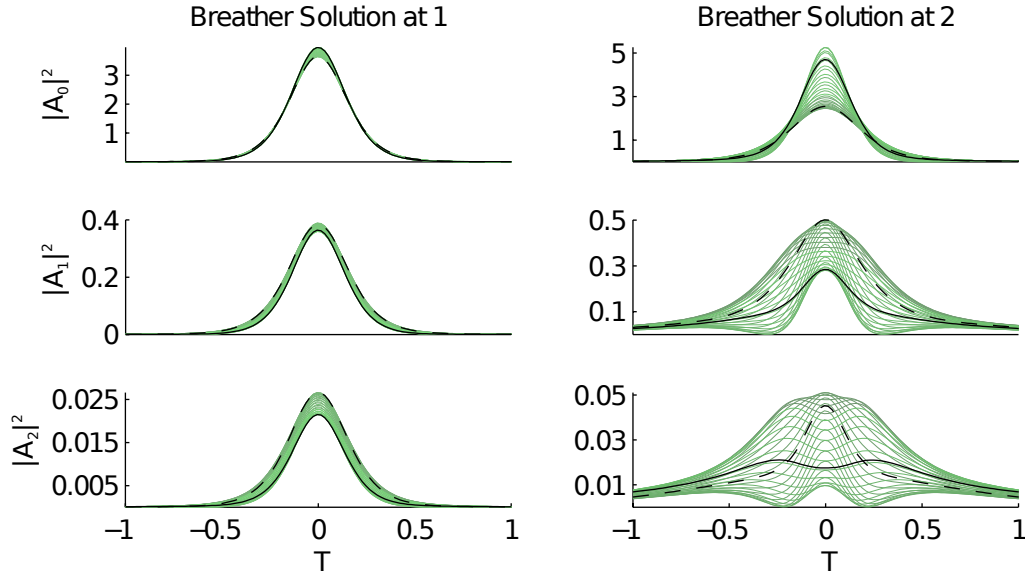


Figure 4.11: Plot of the breather solutions at the points labeled (1) and (2) in Figure 4.9 on the left and right respectively. The solid black line shows the initial condition and the dashed black line shows the solution after a half-period has elapsed. The gray (green) lines show the evolution of the solution in evenly spaced intervals. The left set of plots was taken directly after the Hopf bifurcation, and the right set of plots was taken from the unstable region after the first fold of the limit cycle. Note that the actual domain is from  $T \in [-20, 20)$  and only  $T \in [-1, 1]$  was plotted to highlight the region of interest.

works [30]. In both cases, these solutions have not been previously observed as they are linearly unstable. However, the ACM allows for the continuation of these unstable solutions so the lack of stability is not problematic in this case.

Overall, the period-one breather solutions are critical because they remove the stability of the stationary solutions and bifurcate with an odd mode. The odd mode breaks the even symmetry previously inherent in both the stationary and breather solutions. This symmetry breaking explains how the translations observed in Ref. [31] could occur even when the initial conditions are even functions. Even if the initial condition is truly even, experimental noise or numerical roundoff error will excite the odd mode, which then grows and becomes non-negligible. Furthermore, the structure of the period-one breather branch at low amplitudes, which is initiated by a sub-critical Hopf bifurcation and later gains stability through a saddle-node bifurcation, explains the sudden jump in cavity energy that occurs when  $g_0$  is adiabatically increased through  $g_0 = 2.404$ . Therefore, although the unstable limit cycles revealed by the ACM will never appear in direct numerical simulation from noise, they are necessary to fully explain the dynamics that appear in such simulations.

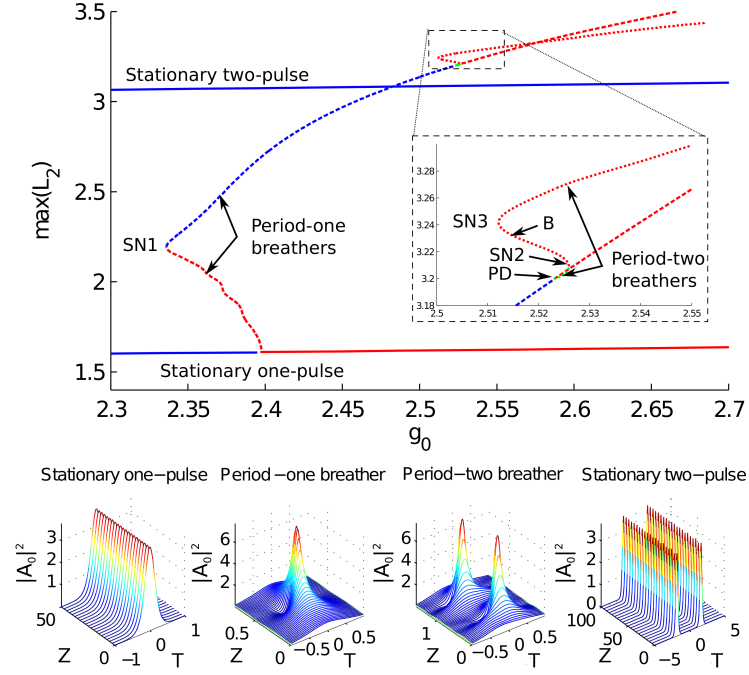


Figure 4.12: (Top) Bifurcation diagram including the branches of stationary one-pulse, stationary two-pulse, period-one breather, and period-two breather solutions. Branches in blue or green are linearly stable while branches in red are linearly unstable. Branches in solid lines are from stationary (constant amplitude) solutions while branches in dashed lines are  $z$ -periodic solutions. The green dashed lines represent period-two breathers and the blue lines period-one. Hopf, saddle-node, and period-doubling bifurcations are denoted by  $H$ ,  $SN$ , and  $PD$  respectively. A fourth unknown bifurcation is indicated by  $B$ . (Bottom) Examples of the four qualitatively different solution behaviors – stationary one-pulse (single-pulse) solutions, period-one breathers, period-two breathers, and stationary two-pulse (double-pulse) solutions – observed during the multi-pulsing transition. The stationary two-pulse solutions can be treated as two non-interacting stationary one-pulse solutions.

#### 4.2.2 Torus Bifurcation of Period-One Breathers

The branch of period-one breather solutions extends beyond the period-doubling bifurcation for values of  $g_0 > 3.0$  and does not appear to terminate. In the full PDE, these solutions are not stable and therefore will not appear in simulations starting from noise. However, as stated in the previous section, the period-doubling bifurcation is caused by an odd bifurcating function and does not appear in systems constrained to be even, such as the one studied in [91]. The solutions that appear at larger values of  $g_0$  are qualitatively similar to those in Sec. 4.2.1; the 0th waveguide resembles an oscillating hyperbolic secant, while the 1st and 2nd waveguides have multiple local maxima. In an even-constrained system, the




Figure 4.13: On the left is the plot of the amplitude of the breather solution at  $g_0 = 2.85$  in waveguide 0, and on the right are the Floquet multipliers. Floquet multipliers outside the unit circle are denoted with a grey (red)  $x$ , and multipliers inside or on the unit circle are denoted with a black (blue) dot. There are multipliers outside the unit circle:  $\mu_{\text{pd}} = -5.825$  and  $\mu_{\text{ns}} = 0.1217 \pm 0.9984i$ .  $\mu_{\text{pd}}$  is associated with an odd eigenfunction and has remained outside the unit circle since the period-doubling bifurcation. The additional pair,  $\mu_{\text{ns}}$  corresponds to a torus (Neimark-Sacker) bifurcation that occurs near  $g_0 = 2.8$ .

branch of solutions remains stable until  $g_0 \approx 2.8$  where a torus (Neimark-Sacker) bifurcation occurs [90].

Figure 4.13 shows the solution at  $g_0 = 2.8$  and the Floquet multipliers of the solution. While the solution itself is still qualitatively similar to the period-one breathers at lower values of  $g_0$ , the difference is that an additional pair of Floquet multipliers has left the unit circle. As shown on the right in Fig. 4.13, there are three unstable multipliers. The eigenfunction now associated with the multiplier at  $\mu_{\text{pd}} = -5.825$  has existed since the period-doubling bifurcation occurred. The new complex-conjugate pair of multipliers,  $\mu_{\text{ns}} = 0.1217 \pm 0.9984i$ , signifies the presence of the torus bifurcation. The eigenfunctions associated with these multipliers are even at  $Z = 0$  and from Eq. (A.6) can be shown to be even for all  $Z$ . Therefore in the even-constrained system, it is the torus bifurcation that is responsible for the loss of stability of the period-one breathers.

In the unconstrained version of the WGAML, this bifurcation does not appear in any meaningful way. Any perturbation will grow primarily in the direction of the eigenfunction with  $\mu = -5.825$  (the one associated with the period doubling bifurcation), and the presence of the torus bifurcation can be neglected. It is only in even-constrained systems, where the period doubling bifurcation cannot occur, that torus bifurcation matters. This case was studied in the previous section and Ref. [91]. In that reduced order model, the torus bifurcation was responsible for the eventual route to chaos in the system. Here we have shown that the same torus bifurcation occurs in this system. Although the adjoint continuation method cannot track the resulting quasi-periodic solutions, this indicates that the reduced order model may be valid.




Figure 4.14: Example of a stable, period-doubled solution. The symmetry breaking of the period-doubling bifurcation has generated a double-peak structure in all three waveguides. The solid black line shows the initial condition and the dashed black line shows the solution after a half-period has elapsed. The first half-period is shown in blue and the second half-period is shown in green. Note that the actual domain is from  $t \in [-20, 20)$  and only  $t \in [-1, 1]$  was plotted to highlight the region of interest.

### 4.2.3 Period-Two Breather Solutions

The period-two branch of solutions bifurcates from the period-one solutions at  $g_0 = 2.523$ . Because the bifurcating eigenfunctions are odd, the period-two breathers are neither even nor odd functions. Figure 4.14 shows a sample solution taken at  $g_0 = 2.527$ . The clearest difference between the period-one breathers and the period-two breathers is that the zeroth waveguide reaches its maximum intensity in two places, near  $T = -0.5$  and  $T = 0.5$ . A similar oscillation occurs in waveguides 1 and 2. This shift in  $T$  is the reason for the doubling of the period. In Fig. 4.14, the solid and dashed lines correspond to  $Z = 0$  and  $Z = Z_0/2$ , respectively. These two lines help show that the solution at  $Z = 0$  and the solution at  $Z = Z_0/2$  are reflections of each other across  $T = 0$ . Due to this symmetry in the problem, the solution spends exactly half the period shifted towards negative values of  $T$  and the other half shifted towards positive values of  $T$ .

From this initial point, we can track the branch of solutions using the period,  $Z_0$ , as the continuation parameter. The bifurcation diagram in Fig. 4.15 shows the period-two solution in  $(g_0, L_2, H_1)$  space where the values of  $L_2$  and  $H_1$  are defined in Eq. (4.7). The stable orbits are shown in blue while the unstable orbits are shown in red. The branch of period-two solutions begins near the point labeled (a) in Fig. 4.15. Immediately after this bifurcation, the period-two solutions are stable until  $g_0 = 2.527$ . Notice that the region for which the period-two solutions are stable is very small with  $\Delta g_0 \approx 0.002$ . For this reason,




Figure 4.15: (Left) A two-dimensional version of the bifurcation diagram showing the maximum  $L_2$  value over a period versus the gain. Stable solutions are plotted in blue and unstable solutions in red. The labeled bifurcation points are shared between the 2D and 3D plots. (Right) The bifurcation diagram of the period-two solutions to the WGAML. Individual orbits are plotted in  $(g_0, L_2, H_1)$  space and each plotted orbit represents one of every thirty solutions computed on the branch. Linearly stable solutions are plotted with blue edges while unstable orbits are plotted with red edges.

the period-two solutions rarely appear in direct numerical simulation.

The stable branch of period-two breathers is terminated by a saddle-node bifurcation around  $g_0 = 2.527$ . At this point, a Floquet multiplier exits the unit circle through  $\mu = 1$ , and there is a fold in the branch of solutions. In a certain sense, this fold is similar to the sub-critical Hopf bifurcation of the period-one branch. The branch of period-two solutions extends below the value of  $g_0$  where the branch began, and the solutions are unstable during this time. The key difference between the period-two branch and the period-one branch is that an additional bifurcation occurs on this low amplitude branch. At this point, an additional multiplier exits the unit circle through  $\mu = 1$ , resulting in two multipliers outside of the unit circle. Although the second saddle-node bifurcation occurs at  $g_0 = 2.251$ , which brings one multiplier back inside the unit circle, the resulting large-amplitude orbits are still unstable. This extra bifurcation is the cause of the vastly different regions of stability despite the qualitative resemblance of the solution branches.

It appears there are no stable branches of either period-one or period-two breathers for larger values of  $g_0$  in the bifurcation diagram. The branch of small-amplitude period-two solutions could potentially be a source of additional branches of periodic solutions. However, the resulting branch of solutions would still be unstable since there remains an additional unstable eigenfunction. Following this branch using the adjoint numerical continuation proved to be unsuccessful, thus the nature of the bifurcation and its co-dimension remains an open question meriting further investigation. To our knowledge, there have been no observations of other periodic solutions that could be associated with a new unknown branch of solutions.

Figure 4.16 in parts (a), (b), and (c) shows sample solutions at varying points along the branch of solutions. Directly after the period-doubling bifurcation at (a), the period-two solution is only a slight perturbation away from the period-one solution, and the peaks are indistinguishable to the eye. Further along the branch at (b), the separation between the peaks increases. This solution is taken just below  $g_0 = 2.527$  and represents the largest




Figure 4.16: Plots of solutions at the three selected points indicated by (a), (b), and (c) in Fig. 4.15. The solid-black and dashed-black lines show the intensity in the 0th waveguide at  $Z = 0$  and  $z = Z_0/2$ . Intermediate times are plotted in gray (the first half-period is shown in blue and the second half-period is shown in green). Note that the actual domain is from  $T \in [-20, 20)$  and only  $T \in [-1, 1]$  was plotted to highlight the region of interest.

separation of pulses that could occur while still keeping the solution stable. At point (c), the pulses are now separated by  $\Delta T = 1$ . This trend continues as  $g_0$  increases beyond point (c). In addition to the increased separation of the points of maximum intensity, the peak intensity decreases as we get further from the period-doubling bifurcation. Despite that, the maximum value of  $L_2$  increases monotonically the further one gets from the period-doubling bifurcation. On the other hand, the maximum value of  $H_1$  decreases monotonically at the same time. Therefore, referring back to Figure 4.15, the period-two solutions are encircled by the period-one solutions in the  $L_2$ - $H_1$  plane. The period-two solutions in turn encircle the two-pulse solutions.

The period-two solutions themselves play a relatively minor role in the transition from the single-pulse to double-pulse solutions. The region of  $g_0$  where the solutions both exist and are stable,  $g_0 \in (2.523, 2.527)$ , is small compared to the regions where the stationary solutions or even the period-one breather solutions are stable. Nonetheless, tracking the branch of period-two breathers suggests that there are no remaining period-one or period-two breathers in the WGAML beyond  $g_0 = 2.527$ .

#### 4.2.4 Global Bifurcation Structure

As described in previous sections, the WGAML model has a variety of solutions and bifurcations when  $g_0$  is increased. In particular, starting on the single-pulse branch, the solutions undergo the following sequence of bifurcations: the single-pulse stationary solution loses sta-

bility and bifurcates to an unstable breather, which undergoes a fold bifurcation and becomes stable. This periodic solution loses stability and bifurcates to a period-two breather, which eventually becomes unstable. After this last bifurcation, additional periodic or stationary solutions on the same branch of solutions were not found. In Fig. 4.17, we summarize our findings in a combined bifurcation diagram in which the stationary and periodic solutions are projected onto the three dimensional  $(g_0, L_2, H_1)$  space. After this projection, the stationary solutions become one-dimensional curves and periodic solutions become two-dimensional closed manifolds. The gray plane indicates the value of  $g_0^u = 2.527$  at which the period-two solution becomes unstable. Our study suggests that from this point on, there are no local bifurcations on this branch that result in stable stationary or periodic solutions. We conjecture that there are no stable stationary or periodic solutions near the one-pulse solutions (other than the one-pulse solutions themselves). However, outside of a neighborhood of this set, breathers and multiple-pulse solutions (such as the two-pulse solution) exist and can be stable.

Indeed the double-pulse solution exists and is stable for values of  $g_0 > g_0^u$ , and one would expect all trajectories to be attracted to it. However, numerical simulations and experiments indicate a more subtle behavior. Solutions initiated in the neighborhood of the two-pulse solution are attracted to it, while solutions initiated in the neighborhood of the one-pulse solution, or in the vicinity of the quiescent solution (low-amplitude white noise), exhibit a typical behavior of intermittent spatio-temporal chaos expressed as translations in  $T$  and non-periodic behavior in  $Z$ . For some trajectories, such a behavior is transient, and eventually the trajectory is attracted to the double-pulse solution. As  $g_0$  is increased, we observe that more initial conditions are attracted to the double-pulse and the transients become shorter. For  $g_0 = g_0^u + \varepsilon$  such transients can be very long and numerical simulations cannot determine whether the attractor for white noise initial data is the double-pulse solution or if there exists a trapping region in which a chaotic attractor reigns. In the latter case the trajectory will hover indefinitely between one-pulse, period-one and period-two unstable solutions.

In Fig. 4.18 and Fig. 4.19, we plot a few typical trajectories, stationary solutions, and relevant periodic orbits in terms of  $(g_0, L_2, H_1)$  for  $g_0 = 2.3, 2.5,$  and  $2.6$ . The trajectories are denoted in light gray and have low-amplitude white noise initial conditions. At  $g_0 = 2.3$ , the solution is rapidly attracted to the stable single-pulse solution denoted by the blue dot. At  $g_0 = 2.5$ , the single-pulse solution is unstable and the trajectory is attracted to the period-one breather solution, which is shown in blue. Lastly for  $g_0 = 2.6 > g_0^u$ , the trajectory eventually is attracted to the stationary two-pulse solution. Prior to that, the solution jumps chaotically between the period-one breather solution, shown in green, and the period-two breather solutions, shown in blue.

By continuing the branch of period-two solutions beyond  $g_0^u$  and projecting individual solutions onto the three dimensional bifurcation diagram we can shed light on this behavior. We observed in the bifurcation diagrams in Figs. 4.15 and 4.17 a change in the shape of the period-two solution at  $g_0 = 2.527$ . Specifically, it is contracted toward the double-pulse solution. Indeed the plot of its evolution over the full period (bottom plot in Fig. 4.16) suggests that it becomes more like the double-pulse solution by having two separated pulses. Combining our findings, we conjecture that there are two possibilities for the source of chaotic




Figure 4.17: Bifurcation diagram of the stationary and breather solutions of the WGAML. The stationary solutions are indicated by thick lines. Solid blue lines indicate spectrally stable solutions while dotted red lines represent solutions that are spectrally unstable. The period-one limit cycles are shown in green and the period-two in light-blue. The gray plane at  $g_0 = 2.527$  represents the largest value of  $g_0$  for which any of the limit cycles are stable. The gray shaded region indicates the interior of the limit cycle, and it should be noted that the stable two-pulse solution remains within the interior of the limit cycle near the transition value of  $g_0$ .

behavior in the multi-pulsing transition. One possibility is that there is a trapping region that includes the neighborhood of the single-pulse solution and extends to the neighborhood of the zero solution. A global bifurcation eventually opens that region and the change in shape of the period-two solution might reveal this bifurcation. Another possibility is that as  $g_0$  is increased, the double-pulse solution and its basin of attraction intersect a neighborhood of the single-pulse solution. As this intersection occurs, the trajectories will asymptotically be attracted to the double-pulse solution. We have observed similar multi-pulsing phenomena in the transition from a general  $N$ -pulse solution to an  $(N + 1)$ -pulse solution. This suggests that the results for the single- to double-pulse transition presented here will be relevant for the more general  $N$ - to  $(N + 1)$ -pulse transition. Although in theory an  $N$ - to  $(N + m)$ -pulse transition for  $m > 1$  may exist, given that the only transitions observed experimentally [34, 80, 81] and theoretically [35, 80, 82] are from  $N$  to  $(N + 1)$  pulses, this alternative transition appears unlikely in practice.




Figure 4.18: Plot of the trajectories of the solutions (in terms of  $L_2$  and  $H_1$ ) for  $g_0 = 2.3$ , 2.5 and 2.6 starting with low-amplitude initial conditions. The long time behavior of the trajectories are shown in blue. The light gray lines show the complete evolution of the solutions from the initial conditions. The light blue lines show the stable single- and double-pulse solutions, while the red lines show the unstable single-pulse solutions. The dark gray lines show other limit cycles at the same value of  $g_0$ , such as the period-one breathers (green) and period-two breathers (light-blue).

### 4.3 Chapter Summary

In this chapter, the solution branches and bifurcations involved in the multi-pulsing transition, as it occurs in a waveguide array mode-locked laser, were studied from two perspectives. The first was a low dimensional POD perspective presented in Sec. 4.1, and the second a highly-accurate fully PDE approach presented in Sec. 4.2.

The reduced model was obtained by projecting the *even* WGAML model onto a characteristic set of six global POD modes per waveguide. This 18 complex (36 real) degree of freedom system, which represents a 64-fold reduction in the state space was able to qualitatively reproduce the dynamics observed in the multi-pulsing transition. In particular, for low values of  $g_0$ , the model completely reproduces the dynamics of the single-pulse solution, including its bifurcations in a region of low amplitude solutions recently discovered in Ref. [83]. With increasing gain  $g_0$ , the single-pulse solution bifurcates to a limit cycle via a Hopf bifurcation, matching to the study of the WGAML model [31]. Our analysis indicates that as that limit cycle grows by increasing  $g_0$  further, it will eventually become chaotic by undergoing a series of bifurcations initiated by a period-doubling bifurcation and then almost immediately followed by a Neimark-Sacker (Torus) bifurcation. The chaos in  $Z$




Figure 4.19: Plot of the trajectory of the solution starting from low-amplitude white noise for  $g_0 = 2.3$ ,  $2.5$ , and  $2.6$ , respectively. The solution trajectory is shown in light gray. The green and light-blue lines show the period-one and period-two limit cycles for values of  $g_0$  where they exist. The blue points or curves denote the long time behavior of the trajectory. At  $g_0 = 2.3$ , the stationary single-pulse solution is the attractor. At  $g_0 = 2.5$ , the period-one breather solution is the limiting solution. At  $g_0 = 2.6$ , the chaotic translating double-pulse solution gives the long time behavior.

is terminated by the double-pulse solution gaining stability and becoming the attractor for low amplitude solutions.

The construction of a reduced model that faithfully reproduces the route from single-pulse to double-pulse solutions has not been previously accomplished using ansatz based reduction methods such as the variational reduction. This is due to the difficulty in formulating a reasonable ansatz for the breather and chaotic solutions. The POD-Galerkin method, applied to the even WGAML model, consists of a nontrivial and low dimensional global basis. Only a few of the modes are needed to represent most of the energy, and therefore it overcomes the difficulty of obtaining a tractable reduced model without a viable ansatz. The optimal global POD basis is generated directly from a combination of short time but high resolution runs that are initiated by a perturbation away from the attractor and combined into a single data set. This methodology captures 99% of the total energy of the entire data in only six modes.

The ACM method applied to the WGAML model reveals a similar pathway. As  $g_0$  is increased, the system undergoes a Hopf bifurcation but, unlike the predictions of the POD model, that bifurcation is sub-critical. The stability of the period-one breather solutions are restored through a saddle-node bifurcation at a lower value of  $g_0$  than the Hopf bifurcation. Stability of the period-one breathers is lost again through a symmetry-breaking period doubling bifurcation. The stability of the resulting period-two breathers is then lost through another saddle-node bifurcation.

The key difference between the two models is that the reduced POD model was con-

structed for the WGAML model restricted to even solutions. This restriction was necessary to bound the number of global modes required to represent the chaos that occurs in the WGAML model. Numerical simulations of the multi-pulsing transition in the even and the unconstrained WGAML model is characterized by a similar route: from single-pulse to chaos and then to a double-pulse. Although the symmetry breaking period doubling bifurcation has an unstable eigenfunction that is odd, further along the branch of period-one breathers is a Neimark-Sacker bifurcation. This is similar to the predictions of the reduced order model. In short, there is qualitative agreement between the solution branches and bifurcations of the reduced POD and the full PDE models.

The POD-Galerkin can be a useful tool for the studies of the dynamics of nonlinear evolution equations describing various optical systems. This is particularly true for cases where the observed dynamics appear to be coherent in nature but non-trivial or novel enough that a canonical ansatz has not yet been established. As long as the reduced model correctly reproduces the dynamics, it can then be studied using the traditional tools and techniques of dynamical systems, such as AUTO [2] or MATCONT [1], in a computationally efficient manner. Furthermore, the POD-Galerkin methodology can be easily modified for changes in the model. Such robustness of the method allows us to study the operating regimes as a function of the parameters in the system. Here the WGAML model was modeled as a three-waveguide array. However, using the same methodology a system of  $N$ -waveguide arrays could be studied. Only minor modifications to Eq. (4.3) are required to model a system with a different number of waveguides. Then parameters for optimal mode-locking in the such a system or the impact of the parameters on the multi-pulsing transition can be studied. In addition, models based on the POD-Galerkin method are sometimes able to predict dynamics in regions outside of the data set where the POD modes were computed, thus the low dimensional model can be used to find interesting regions within the system without having to develop an appropriate ansatz. Once these regions have been identified, by altering the data set used, a more accurate representation of the region can be developed. In summary, the POD reduction gives a new and insightful way to explore the dynamics and bifurcations of a given system.

## Chapter 5

**HYBRID INTEGRATORS BASED ON REDUCED ORDER MODELS**

Modern methods in scientific computing aim to exploit any underlying mathematical structure to make computations more efficient. Adaptive time-stepping schemes for ordinary and partial differential equations are an example of a numerical method that can have a profound impact in a wide range of fields including the mathematical, physical, biological, and engineering sciences. Indeed, some of these schemes, for example those based on embedded Runge-Kutta methods [92], are used so universally that they are now an integral component in most commercial and open source scientific computing software packages. The advantage of these adaptive integrators is particularly profound in time-dependent multi-scale physics problems where large steps can be taken in certain regimes but more refinement is needed to accurately capture fast time scales in other regimes.

Motivated by the increase in performance that can be obtained by adapting the integrator based on the presence (or absence) of certain time dynamics in a problem, in this chapter we propose a hybrid numerical integration scheme that exploits, when possible, the low dimensional nature of a given system. The second plot in Fig. 5.1 demonstrates the key advantage of exploiting low-dimensionality, i.e., the reduced dimensional framework allows for the system to be evolved at a significantly reduced computational cost. For example, the reduced model computes the solution in only 36.18 seconds, and is an order magnitude faster (6% of the computation time) than the full PDE integrator. However, that same figure also demonstrates a key limitation of a static reduced basis like the ones used in Chapters 3 and 4; should the dynamics of the system change, say due to a bifurcation, the predictions of the low dimensional system can become meaningless.

To overcome this limitation, the POD basis must be adapted when dynamics of the system change. This leaves the question of how to inexpensively determine the validity of the reduced order model. A number of researchers have proposed answers to this question including placing bounds on the error of POD-Galerkin methods [93] and introducing entirely new types of reduced order models that allow for *a posteriori* estimates of the error [94]. The approach taken in this manuscript is to compare the results of two model reduction methods: the predictive and equation-free nature of the dynamic mode decomposition (DMD) and the well-studied and potentially accurate POD-Galerkin technique. If the predictions of the two are close, then we will continue to use the reduced model. If not, we revert to the full PDE. The result is a robust yet simple integration scheme that can, for appropriate and non-trivial types of system dynamics, retain many of the computational benefits of reduced order models while avoiding spurious results. Figure 5.1c is an example of the effectiveness of this method. The hybrid integrator is faster than a full PDE integrator and, while not as quick as a static reduced order model, gives valid results. Furthermore, the overhead of this method, even when applied to systems that *do not* exhibit the necessary dynamics, is manageable.




Figure 5.1: (a) A solution of the Cubic Quintic Ginzburg Landau equation that, due to time-dependence of the loss, undergoes a Hopf bifurcation from a stationary solution to a time-periodic solution and back. (b) Results from a reduced order POD model that samples data as the solution converges to the stationary solution. Although the static POD model can represent the stationary solution accurately, the ROM is irrelevant after the bifurcation and leads to erroneous results. Regions in red denote when data collection occurs. (c) Results from the hybrid POD/DMD model in this manuscript. By adaptively re-sampling from the PDE, the ROM updates itself to produce accurate results.

Although there are some exceptions, the generation of the POD modes is typically performed in a static and “offline” fashion where data for the modes are collected prior to generation of the model, and the POD basis itself remains unchanging despite any new incoming information. In order to make the numerical integrator more robust while retaining some of the computational benefits of the reduced model, we adaptively switch between the accurate but expensive PDE simulations, to collect data or evolve the solutions in regimes that are high dimensional, and reduced order simulations that, while inexpensive to integrate, require representative data in order to be constructed and lose accuracy should the dynamics of the system travel too far away from the location in phase space where data was obtained.

The requirement of low dimensional dynamics may appear to be restrictive, but there are many systems of current research interest that satisfy this constraint. In particular, pattern forming systems and systems with spontaneously appearing coherent structures are both interesting for mathematical and scientific reasons [95] and often low dimensional in nature. Examples of such coherent structures in practical applications include mode-locked pulses in nonlinear optics [34], vortices and convective heating in fluid dynamics [4, 13, 96–98] synergies or synchrony in biological systems [5, 99, 100] or social networks [101], as well as many other disciplines in the physical and life sciences. As a result, there is no shortage of potential applications for such an integrator. In this manuscript, we will apply the method to two models for mode-locked lasers: the Cubic Quintic Ginzburg Landau equation (CGLE) and the Waveguide Array Mode-Locking (WGAML) model.

We have divided the chapter into the following sections: Section 5.1 outlines the POD and DMD methods and demonstrates their application on the Nonlinear Schrödinger equa-

tion. Section 5.2 describes the hybrid integrator itself and the metrics it is built upon. Section 5.3 demonstrates the application of the method to the Cubic Quintic Ginzburg Landau equation (CQGLE). Section 5.4 demonstrates the application of the hybrid integrator in a more practical situation, a parameter range sweep through the multi-pulsing transition of a waveguide array mode-locked laser. Lastly, Section 5.5 gives concluding remarks and future outlook.

### 5.1 The Proper Orthogonal Decomposition and Dynamic Mode Decomposition

The Proper Orthogonal Decomposition (POD) and Dynamic Mode Decomposition (DMD), which serve as the basis for the hybrid integrator presented in this manuscript, are two conceptually independent methods for generating reduced order models from a given set of data. In this section, we will first give a short review of the POD and DMD methods as they apply to the hybrid integrator. For a general review of the POD see Sec. 1.2 and for a review of DMD see Refs. [97, 102, 103] as well as the references contained therein.

As both POD and DMD are based on data, the choice of data set is a key factor in the results. In both cases, we require a sequence of  $M$  solution snapshots evenly spaced in time with  $N$  points saved per snapshot, i.e.  $U(x, t_1), U(x, t_2), U(x, t_3), \dots, U(x, t_M)$  where  $t_{m+1} = t_m + \Delta t$  and  $U(x, t_m) \in \mathbb{C}^N$ . These data are arranged in a matrix

$$\mathbf{X} = [U(x, t_1) \quad U(x, t_2) \quad \cdots \quad U(x, t_{M-1}) \quad U(x, t_M)],$$

which results in  $\mathbf{X} \in \mathbb{C}^{N \times M}$ . For the purposes of notation, it will be convenient to define the sub-matrices  $\mathbf{X}_1^{M-1}, \mathbf{X}_2^M \in \mathbb{C}^{N \times M-1}$  as the matrices that contain columns 1 through  $(M-1)$  and 2 through  $M$  of  $\mathbf{X}$  arranged in the same order as in  $\mathbf{X}$ . With this data, we apply the reduced SVD to  $\mathbf{X}_1^{M-1}$  and obtain

$$\mathbf{X}_1^{M-1} = \mathbf{U}\mathbf{\Sigma}\mathbf{W}^*, \quad (5.1)$$

with  $\mathbf{U} \in \mathbb{C}^{N \times K}$ ,  $\mathbf{\Sigma} \in \mathbb{C}^{K \times K}$ , and  $\mathbf{W} \in \mathbb{C}^{M-1 \times K}$ ,  $K$  is the reduced SVD's approximation of the rank of  $\mathbf{X}_1^{M-1}$ , and the asterisk denotes the conjugate transpose. From this point on, the POD and DMD methods differ.

#### 5.1.1 The Proper Orthogonal Decomposition

The POD method has already been outlined in Sec. 1.2. However, in order to use the same singular value decomposition for both the POD and DMD methods, we require a slight modification of the method outlined there. As before, the POD is intrinsically linked to the SVD. Previously the POD method was applied to the entirety of the data set which would be denoted here as  $\mathbf{X}_1^M$ . Here we apply the POD method to a subset of the data to get  $\mathbf{X}_1^{M-1} = \mathbf{U}\mathbf{\Sigma}\mathbf{W}$ , where  $\{\phi_1, \dots, \phi_K\}$  are the columns of  $\mathbf{U}$ .

As usual, we retain only a subset of the POD modes based on the energy content of each

mode. Here the solution is assumed to be a linear superposition of  $P$  POD-modes,

$$U_{\text{POD}}(x, t) = \sum_{p=1}^P a_p(t) \phi_p(x) = \mathbf{\Phi} \mathbf{a}(t), \quad (5.2)$$

where  $\mathbf{\Phi} \in \mathbb{C}^{N \times P}$  is the matrix whose  $p$ -th column is  $\phi_p$ , and  $\mathbf{a}(t)$  is the vector of the  $a_p(t)$ . Then the Galerkin projection is applied to obtain a system of ordinary differential equations for the  $a_p$ . If data is taken while the system is near to an attractor, the energy content in each of the POD modes remains relatively constant, and the resulting reduced order model will be accurate provided the parameter values remain the same. If this is not the case however, then the reduced order model may give spurious results as the neglected modes gain energy and gain importance.

### 5.1.2 Dynamic Mode Decomposition

Dynamic Mode Decomposition (DMD) is a more recent numerical technique that identifies the eigenvalues and eigenvectors of the unknown linear time-independent operator,  $\mathbf{A}$ , that accurately explains the data collected in  $\mathbf{X}$  [97, 102, 104]. It should be noted that DMD can be used even if  $\mathbf{X}$  was generated by a nonlinear operator. In the nonlinear case, the DMD modes are approximations of the modes of the Koopman operator. The Koopman operator is an infinite dimensional and linear operator, but is capable of representing the dynamics of nonlinear, finite dimensional systems [97, 105, 106]. In essence, DMD seeks to approximate this Koopman operator with a finite and sometimes even low-rank linear operator.

For our purposes however, the crux of DMD is to assume that for any of the observed data points,  $\mathbf{X}_1^{M-1} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{M-1}\}$ , we have that

$$\mathbf{x}_{i+1} = \mathbf{A} \mathbf{x}_i, \quad (5.3)$$

which implies that equivalently  $\mathbf{X}_1^{M-1} = \{\mathbf{x}_1, \mathbf{A} \mathbf{x}_1, \mathbf{A}^2 \mathbf{x}_1, \dots, \mathbf{A}^{M-2} \mathbf{x}_1\}$ . As a result, the columns of  $\mathbf{X}_1^{M-1}$  are each elements in a Krylov space. We then attempt to represent the final data point,  $\mathbf{x}_M$ , in terms of this Krylov basis

$$\mathbf{x}_M = \sum_{m=1}^{M-1} b_m \mathbf{x}_m + \mathbf{r}, \quad (5.4)$$

where the  $b_m$  are unknown mode amplitudes, and  $\mathbf{r}$  is the residual that lies outside of the Krylov space. The constraints in Eq. (5.3) and Eq. (5.4) can be expressed as a single algebraic equation

$$\mathbf{A} \mathbf{X}_1^{M-1} = \mathbf{X}_2^M = \mathbf{X}_1^{M-1} \mathbf{S} + \mathbf{r} e_{M-1}^*$$

where  $e_{M-1}$  is the  $(M - 1)$ th unit vector and

$$\mathbf{S} = \begin{bmatrix} 0 & & & b_1 \\ 1 & \ddots & & b_2 \\ 0 & \ddots & \ddots & \vdots \\ & \ddots & \ddots & 0 & b_{M-2} \\ & & & 0 & 1 & b_{M-1} \end{bmatrix},$$

where the  $b_k$  are unknown. This is similar to the Arnoldi method in that the eigenvalues (or Ritz eigenvalues) of  $\mathbf{S}$  approximate some of the eigenvalues of the unknown operator  $\mathbf{A}$ . Rather than compute  $\mathbf{S}$  directly, we instead compute the matrix

$$\tilde{\mathbf{S}} = \mathbf{U}^* \mathbf{X}_2^M \mathbf{W} \mathbf{\Sigma}^{-1}, \quad (5.5)$$

which is related to  $\mathbf{S}$  via a similarity transform [107]. Here  $\mathbf{U}$ ,  $\mathbf{W}$ , and  $\mathbf{\Sigma}$  are from the SVD taken in Eq. (5.1). With this machinery, the DMD modes can be written as

$$\psi_k = \mathbf{U} \mathbf{y}_k, \quad (5.6)$$

where  $\mathbf{y}_k$  is the  $k$ -th eigenvector of  $\tilde{\mathbf{S}}$  associated with the eigenvalue  $\mu_k$ . The eigenvalues,  $\mu_k$ , capture the effects of a discrete step of  $\Delta t$  forward in time. Often, it is beneficial to work with  $\omega_k = \frac{\ln(\mu_k)}{\Delta t}$  to allow the DMD model to be used at intermediate time points. Because we have computed the eigenvectors and eigenvalues of the unknown operator, the solution for all time can be written analytically as

$$D_{\text{DMD}} = \sum_{k=1}^K b_k(0) \psi_k(x) e^{\omega_k t} = \mathbf{\Psi} \text{diag}(e^{\omega t}) \mathbf{b}, \quad (5.7)$$

where  $b_k(0)$  is the initial amplitude of each mode,  $\mathbf{\Psi}$  is the matrix with  $\psi_k$  as the  $k$ -th column,  $\text{diag}(e^{\omega t})$  is the diagonal matrix with  $e^{\omega_k t}$  in the  $k$ th column, and  $\mathbf{b}$  is the vector of  $b_k$ . To compute the mode-amplitudes from a snapshot  $\mathbf{x}$ , we set

$$\mathbf{b} = \mathbf{\Psi}^+ \mathbf{x}, \quad (5.8)$$

where  $^+$  denotes the Moore-Penrose pseudo-inverse [65]. The pseudo-inverse solves the least squares problem for  $\mathbf{b}$  with  $L^2$ -regularization and is consistent with how the DMD modes are originally derived.

Overall, the Dynamic Mode Decomposition generates a set of basis functions that are the eigenvectors of a low dimensional approximation to the full flow operator of the underlying system. These modes, unlike the POD modes, are not orthogonal. However, for each mode the time dynamics are given by the associated  $\omega_k$  without the need to solve a system of ODEs. Due to the inherently approximate nature of the  $\omega_k$ 's, even for data generated by a linear time-independent matrix where they related to the Ritz values, the computed values of  $\omega_k$  often contain a non-negligible amount of error. As a result, the DMD method is best

used to predict the dynamics over a short time-horizon. Furthermore, because the DMD representation is not truncated like the POD representation is, the space spanned by the  $K$  DMD modes is at least as large as the one spanned by the  $P$  dimensional POD basis.

### 5.1.3 Application to the Nonlinear Schrödinger Equation

As an example, we will apply both techniques to the two-soliton solution of the Nonlinear Schrödinger equation (NLSE),

$$\frac{\partial U(x, t)}{\partial t} = i \frac{1}{2} \frac{\partial^2 U(x, t)}{\partial x^2} + i |U(x, t)|^2 U(x, t), \quad (5.9a)$$

$$U(x, 0) = 2\text{sech}(x), \quad (5.9b)$$

in order to compare the methods. This solution of the NLSE has been extensively studied and exhibits non-trivial but also low dimensional time-periodic behavior (see Ref. [108]). This combination makes the NLSE particularly suitable as a pedagogical example.

As described in Sec. 5.1.1, the Galerkin projection is used to obtain equations of motion for the mode amplitudes of the POD modes; this we refer to as the reduced system. In the case of Eq. (5.9), it is

$$\frac{da_i}{dt} = \sum_{p=1}^P \mathbf{L}_{ip} a_p + \left\langle \phi_i, \left| \sum_{p=1}^P a_p \phi_p \right|^2 \left( \sum_{p=1}^P a_p \phi_p \right) \right\rangle, \quad (5.10a)$$

$$\mathbf{L}_{ij} = -i \frac{D}{2} \langle \partial_x \phi_i, \partial_x \phi_j \rangle, \quad (5.10b)$$

$$a_i(0) = \langle \phi_i, U(x, 0) \rangle, \quad (5.10c)$$

where  $\langle u, v \rangle = \int_0^L u^*(x)v(x) dx$  using  $L$  as the width of the computational domain. By numerically solving the low dimensional ODE system in Eq. (5.10), we can then use Eq. (5.2) to rapidly obtain an approximation of  $U(x, t)$  in the PDE. Equation (5.10), as written, is quite general and assumed nothing but the orthogonality of the POD modes. The resulting dynamics depend primarily on the value of the inner products in Eq. (5.10), but in order to compute those we first need to obtain data.

To generate data, Eq. (5.9) was solved with  $L = 40$  using a 256-mode Fourier basis. Starting from a hyperbolic secant initial condition with a maximum amplitude of two, the system was evolved forward in time using an adaptive time-stepping fourth-order Runge-Kutta method (`ode45`) in MATLAB. Figure 5.2 plots (a), (d), and (i) show the amplitude of the solution for  $t \in [0, 5]$ , the relative error between the Fourier and analytic solution, and the amplitude of the largest ten (of 256) Fourier modes as a function of time. In addition to the coherent shape of the resulting solution, the correlation between the Fourier mode amplitudes implies that a significant amount of dimensionality reduction may be achieved by switching to a different basis.

The second column in Fig. 5.2, plots (b), (e), and (h), demonstrate the solution approximation, relative error, and mode-amplitude using a three-mode version of Eq. (5.10). The

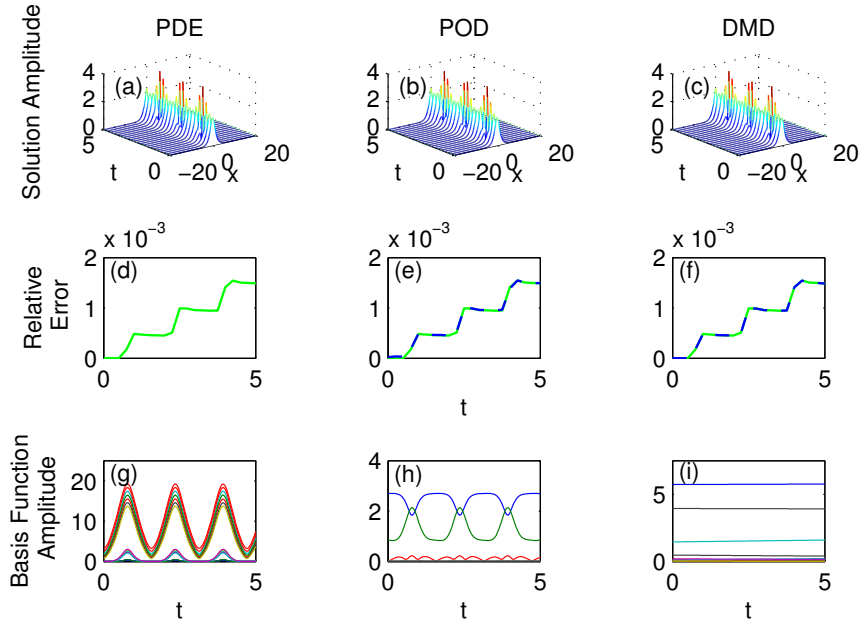


Figure 5.2: Amplitude plots of the two soliton solution to NLS generated by: (a) numerically solving Eq. (5.9), (b) 3-mode POD representation (Eq. (5.10) with  $P = 3$ ), (c) 20-mode DMD representation (Eq. (5.7) with  $K = 20$ ). Plots (d), (e), and (f) show the relative error (in the  $L^2$ -sense) between the exact solution and the Fourier, POD, and DMD representations of the solution. For reference, the green lines in all three plots show the error of the Fourier representation. The blue lines show the error of the POD and DMD representations in the second and third columns. Because the POD and DMD modes are generated from Fourier data, the accuracy of those methods cannot exceed that of the Fourier representation. In this case, however, both the POD and DMD representations have a relative error within  $10^{-5}$  of that of the Fourier basis. Plots (g), (h), and (i) show the ten Fourier modes with largest mean amplitude, the POD mode amplitudes, and the DMD mode amplitudes.

resulting solution approximation is indistinguishable to the eye from the exact PDE solution using many more modes. The time-dynamics of each POD mode, like the Fourier modes, are non-trivial due to the fact that their evolution is governed by a system of nonlinear ODEs. Nonetheless, due to the low dimension of that system the computational time required to evolving the POD model is a factor of ten less than that required to evolve the full PDE model.

The third column in Fig. 5.2 shows the solution reconstruction, relative error, and mode-amplitudes of the DMD representation in Eq. (5.7) using 20 DMD modes. Although the DMD representation has fewer modes than the full PDE, it too can reconstruct the original solution with differences that are indistinguishable to the eye. Unlike the POD or Fourier modes, the time dynamics of the DMD modes are nearly trivial. Indeed because each  $\omega_k$  is almost purely imaginary, mode-amplitudes are similar to pure sinusoids in time. Due

to this limitation on the complexity of the time dynamics, more DMD modes are required to adequately represent the solution, but because their time evolution can be determined analytically, computing the DMD solution at *any point in time* only costs  $\mathcal{O}(NK)$ .

In both cases, the POD and DMD representations of the solution produce errors that are, to the eye, identical to that of the Fourier representation; in fact, the relative error of the POD and DMD methods is larger by at most  $10^{-5}$ . The reason for the similarity between the three methods is that the POD and DMD methods used data taken from the Fourier solution and not from the analytic solution, and as a result, the DMD and POD methods are approximating the *Fourier representation of the solution* and not the analytic solution. Overall this demonstrates that, given the proper data, both reduced order models can reproduce, with quantitative levels of accuracy, the dynamics of a system.

The POD and DMD offer two independent strategies for constructing reduced order models given a set of data. For computational efficiency, both make use of the reduced SVD, but while the POD and SVD are intrinsically linked, DMD is conceptually similar to the Arnoldi and other Krylov methods as it is assumed the data, and therefore the basis, is built from the repeated application of an unknown linear operator [97, 107]. Furthermore, because DMD is approximating the operator itself from the data, no additional work needs to be done to obtain the time dynamics. DMD is truly an example of an equation free technique. The POD, on the other hand, neglects all temporal information. To overcome this, the Galerkin projection must be applied to the system used to obtain the data in order to derive equations for the POD mode amplitudes. The benefit of this equation-aware approach is that the mode-amplitudes can evolve in a nonlinear fashion, which allows POD models to be generally smaller in dimension than an equally accurate DMD model. In the example given here and in what will follow, we assume a case where both approaches are possible, i.e., we are given data spaced equally in time by  $\Delta t$  for DMD, and the governing equations responsible for generating that data are also known for the Galerkin projection.

## 5.2 Hybrid Integration with Reduced Order Models

In this section, we outline the hybrid reduced order model (ROM)/PDE integration scheme, the underlying idea of which is based on adaptively switching between integrating the full PDE model and a reduced order model of the PDE. The PDE integrator is accurate and does not make the critical assumption of low-dimensionality, but due to the large state space and to the potential stiffness issues is computationally expensive. On the other hand, ROMs like the POD method are fast and accurate *if the underlying system is low dimensional*. If not, erroneous results like those given in Fig. 5.1 can be the result. To avoid this, we must do two things. First, we must ensure our hybrid integrator collects useful data for the POD and DMD methods to analyze. Next, we must define two metrics:  $H(\mathbf{X}; \Phi)$ , which determines if the dynamics in the PDE are low dimensional from a given set of PDE data and POD modes, and  $R(\mathbf{a}; \Phi, \Psi)$  which determines if the dynamics of the POD model have remained low dimensional from a single snapshot of the reduced model. Assuming for the moment we have these two metrics, the hybrid integration scheme, at its simplest, consists of four steps:

1. Evolve the PDE for a duration,  $\Delta T$ , and record  $j$  solution snapshots at evenly spaced time intervals. Augment any existing data with these new samples, e.g. if we previously

had  $(m - 1)$  sets of snapshots, the data matrix is now  $\mathbf{X} \in \mathbb{C}^{N \times mj}$ . In practice, we retained a maximum of  $m = 5$  data sets.

2. Compute the POD modes ( $\Phi$ ) and DMD modes ( $\Psi$ ) of  $\mathbf{X}_1^{(m-1)j}$  following Sec. 5.1 and check if  $H(\mathbf{X}; \Phi) > H_{tol}$ . If so, repeat step 1.
3. Evolve the system forward in the reduced model until  $R(\mathbf{a}; \Phi, \Psi) > R_{tol}$ .
4. Purge any saved data in  $\mathbf{X}$ , and repeat from step 1 unless the terminal time is reached.

As shown here, the hybrid integrator is conceptually very simple. By augmenting  $\mathbf{X}$  in step 1, we allow the hybrid integrator to effectively adjust the sampling interval. This is necessary because, to be accurate, the POD method must be performed on a data set that contains a characteristic representation of the low dimensional manifold on which the dynamics are restricted. However, the minimum time-length required to capture this data is attractor dependent. While a conservatively long time interval could be manually selected, the overhead of working with larger data sets and increased granularity of the integration period is greater than the iterative process outlined above. Restricting ourselves to a maximum of 5 updates to the data matrix before restarting is necessary to limit the computational cost of the POD and DMD methods, and is similar to the explicit restarting sometimes used in numerical linear algebra [65]. The key challenge is defining workable metrics as well as the associated tolerances,  $H_{tol}$  and  $R_{tol}$ , for switching between the two models. Currently, we have developed a single metric that measures if a set of observed data is low dimensional and a pair of metrics that reflect whether or not a given snapshot of the reduced order model indicates that the low dimensional behavior is persisting.

### 5.2.1 A Metric for Detecting Low Dimensionality

The metric for low dimensionality operates by directly checking whether or not the reduced  $P$ -mode POD basis can accurately represent new data that it was not developed from. To do this, we need access to a time series of data in  $\mathbf{X}$ , which we subdivide into two parts,  $\mathbf{X}_1^{(m-1)j}$  and  $\mathbf{X}_{(m-1)j+1}^{mj}$ . Applying the POD method to  $\mathbf{X}_1^{(m-1)j}$  yields a set of modes  $\{\phi_1, \dots, \phi_P\}$  as mentioned previously. Let  $\rho : \mathbb{C}^N \rightarrow \mathbb{C}^N$  be a mapping that projects a snapshot of the PDE solution onto the space spanned by the  $P$ -mode POD basis. If Fourier spectral methods are used, then  $\rho = (\Delta x)(\Phi\Phi^*)$  where  $\Delta x$  is the grid spacing. Additionally, let  $\rho_C = I - \rho$  be the complement of that projection operator. With this

$$H(\mathbf{X}) = \frac{\sum_{i=(m-1)j+1}^{mj} \|\rho_C \mathbf{x}_i\|^2}{\sum_{i=(m-1)j+1}^{mj} \|\mathbf{x}_i\|^2} = \frac{\|\rho_C \mathbf{X}_{(m-1)j+1}^{mj}\|_F^2}{\|\mathbf{X}_{(m-1)j+1}^{mj}\|_F^2}. \quad (5.11)$$

Equation (5.11) is the ratio of the Frobenius norm of the data that cannot be represented by the  $P$ -mode POD basis to the total Frobenius norm of the data.

The idea behind this choice of metric is to explicitly determine if a set of  $P$ -POD modes must be augmented to accurately represent additional data, similar to the iterative nature

of sequential POD [11]. If additional modes are required, the POD method has not sampled a significant enough fraction of the attractor, if indeed we are even near to an attractor, to accurately characterize it. Therefore we must continue to collect data. In order to bound the computational costs of this process, we augment the data set at most 5 times before restarting. This sort of explicit restart is meant to limit the maximum size of the data matrix and control the computational cost of the SVD, the most computationally expensive calculation in this algorithm.

When  $H(\mathbf{X}; \Phi) < H_{\text{tol}}$ , the algorithm assumes that the system has both reached an attractor and that the POD basis spans the dynamics on that attractor. However depending on the sampling interval ( $\Delta T$ ), the dynamics of the attractor, or even an un-sampled time-dependence of the parameter values as in Fig. 5.1, this assumption may not hold for all  $t$ . As a result, an additional metric that determines if what is believed to be low dimensional dynamics is, indeed, low dimensional dynamics, is required.

### 5.2.2 Metrics for Confirming Low Dimensionality

In this section, we develop two techniques for determining whether a low dimensional system is exhibiting dynamics that, were they observed in the PDE, would still be considered low dimensional. Due to the on-line nature of the data collection process, the hybrid algorithm is unable to determine whether or not a given set of data is truly a representative sample of the long-time dynamics of the system. Even with the metric outlined in Sec. 5.2.1, the integrator is unable to distinguish a meta-stable state from a stable state or a system with fast dynamics on a slow manifold from a system with just fast dynamics.

The key difficulty is defining these metrics. By design, the POD dynamics are constrained in the  $P$ -dimensional POD subspace and can never exhibit high-dimensional behavior. Error bounds on the POD method have been derived, see [93], but they too require knowledge about the solution dynamics that lay in the complement of the POD subspace. Direct simulation of the PDE itself is the logical source of this information, but in order to obtain a meaningful computational speedup, such checks must be infrequent.

The rationale behind both of the metrics presented in this section is to use DMD as a proxy for the PDE dynamics. The assumptions made by DMD, namely that it has obtained the eigenvalues and eigenfunctions of the underlying Koopman operator, are independent of those made by the POD method. As a result, even if the DMD proxy fails to accurately reproduce the PDE dynamics, there is nothing biasing the DMD and POD methods to fail in the same way, which would produce the types of false negative that must be avoided.

#### *Derivative Comparison Metric*

The first of two metrics we developed directly compares the time derivative of the POD reduced order model with the time derivative of the DMD model with the same state input. This results in the metric

$$R_D(\mathbf{a}) = \frac{\|\partial_t U_{POD} - \partial_t U_{DMD}\|^2}{\|\partial_t U_{POD}\|^2}, \quad (5.12)$$

where  $\partial_t U_{POD}$  is obtained by evaluating the governing equations projected onto the POD modes, e.g., Eq. (5.10).  $\partial_t U_{DMD}$ , on the other hand, is obtained by analytically taking a time derivative of Eq. (5.7). To obtain the DMD mode amplitudes,  $\mathbf{b}$ , from the POD mode amplitudes,  $\mathbf{a}$ , we note that from Eq. (5.8)

$$\mathbf{b} = \Psi^+ U_{POD} = \Psi^+ \Phi \mathbf{a}. \quad (5.13)$$

Once the DMD mode amplitudes have been obtained and a time derivative taken, we can write

$$\partial_t U_{DMD} = (\Psi \Omega \Psi^+ \Phi) \mathbf{a}, \quad (5.14)$$

where  $\Omega$  is the diagonal matrix with the  $\omega_k$  on the diagonal. The four grouped matrices in Eq. (5.14) are constant for a given set of data and can thus be precomputed to form a matrix in  $\mathbb{C}^{N \times P}$ , which is inexpensive to evaluate ( $\mathcal{O}(NP)$ ) in practice. This metric is applied at a frequency equal to that of the sampling frequency in the PDE. Due to the adaptive nature of the step size in the Runge-Kutta methods used, this evaluation occurs approximately every 5 time steps.

For some problems, even the examples shown here, it is reasonable computationally to directly evaluate the time derivative in the PDE. For the case of NLS, the complexity of computing the time derivative directly is  $\mathcal{O}(N \log(N))$ , which is equivalent to the cost of computing the derivative using the reduced order model when  $P \sim \log(N)$ . The results of this metric are essentially equivalent, though slightly more robust, than the version using the DMD proxy. The difference becomes more clear when the right-hand side function becomes more costly to evaluate when other routines such as spectral element methods are used. In short, for this metric the use of either DMD or PDE for the reference time-derivative is workable in practice. The difference is a slight increase in robustness with the PDE version at the cost of a slight increase in computational cost.

### *POD Projection Metric*

The second metric is similar to model-predictive control. The goal is to determine whether or not we would still choose to use a low dimensional model, i.e., if  $H$  is still small, if we were to revert back to the full model and collect data once again. With this metric, the use of DMD is essential as solving the PDE is assumed to be expensive.

Using Eq. (5.13) and Eq. (5.7), we can define an operator  $L_T : \mathbb{C}^N \rightarrow \mathbb{C}^N$  that takes a PDE snapshot, computes the DMD coefficients associated with that snapshot, and evolves the solution forward in time by  $T$ . Similar to Eq. (5.11), we define the projective error to be

$$R_P(\mathbf{a}) = \frac{\|\rho_C L_T U_{POD}\|^2}{\|U_{POD}\|^2} = \frac{\|(\rho_C L_T \Phi) \mathbf{a}\|^2}{\|\mathbf{a}\|^2}. \quad (5.15)$$

Unlike in Eq. (5.11), we only compare the solutions at a single point in time, but this appears to be sufficient as will be demonstrated in the next section. Because  $\rho_C$ ,  $L_T$ , and  $\Phi$  are

fixed, we can combine these operators into a single matrix

$$\rho_C L_T \Phi = (\mathbf{I} - (\Delta x) \Phi \Phi^*) \Psi \text{diag}(e^{\omega T}) \Psi^+ \Phi \quad (5.16)$$

where  $\Delta x$  is the distance between spatial grid points, and  $\text{diag}(e^{\omega T})$  is a diagonal matrix with the  $e^{\omega_k T}$  on the diagonal. The combined operator  $\rho_C L_T \Phi$  first maps from the POD mode amplitudes to physical space ( $\Phi$ ), then maps from physical space to the DMD mode amplitudes,  $\mathbf{b}$ , via the application of the pseudo-inverse. With the initial DMD mode amplitudes, modes, and eigenvalues, the evolution of the amplitudes can be computed analytically. In particular, applying  $e^{\omega_k T}$  is sufficient to evolve the  $k$ th DMD mode amplitude from time  $t$  to  $t+T$ . Left multiplication by the modes again maps back to physical space from the DMD amplitude space. Lastly, the operator  $\rho_C$  removes any components of the DMD solution approximation that lie in the subspace that the POD modes can represent. The result is the part of the solution we could not represent with the current basis. Like  $\mathbf{H}(\mathbf{X}; \Phi)$ , if this fraction exceeds a given tolerance then the dynamics of the system appear to have changed from when data was taken, and resampling is necessary to represent these new dynamics.

Either of the metrics defined here,  $R_P$  or  $R_D$ , appear to work in a range of cases and result in qualitatively similar behaviors. However, details such as exactly when the system transitions do indeed depend on which metric is chosen, and there certainly are cases where one is superior to the other. In the following sections, we use these two metrics on different problems. The data available to the metrics will be collected in batches of twenty, i.e.,  $j = 20$ . Given this data, we will remain in the PDE integrator until  $H < 10^{-6}$  and revert back to the PDE integrator if  $R_P > 10^{-4}$  or  $R_D > 10^{-2}$  depending on which metric is used. For the  $R_P$  metric,  $T = \Delta T \approx 4$ , a single sampling period. These values appear to work on a broad class of problems including a wide variety of the coherent structures that appear in the CQGLE as well as other mode-locking systems like the waveguide array mode-locking (WGAML) model [31] and will likely be reasonable initial estimates for systems with dynamics similar in time-scale to those of the CQGLE.

### 5.3 Cubic Quintic Ginzburg Landau Equation

In order to demonstrate the efficiency of the hybrid method and highlight the differences between the  $R_P$  and  $R_D$  metrics, in this section we apply the method to the Cubic Quintic Ginzburg Landau Equation (CQGLE). CQGLE is a phenomenological model that describes the onset of instabilities in many non-equilibrium physical systems [95]. As a result, it appears in a wide number of fields. Here we focus on the transitions between localized solutions that might arise in the context of mode-locking [34, 35]. As such, CQGLE takes on the form

$$\frac{\partial U}{\partial t} = \left( \frac{i}{2} - \tau \right) \frac{\partial^2 U}{\partial x^2} + \kappa \frac{\partial^4 U}{\partial x^4} + (i + \beta)|U|^2 U + (i\nu + \sigma)|U|^4 U + \gamma U, \quad (5.17)$$

where  $\tau, \kappa, \beta, \nu, \sigma$  and  $\gamma$  are parameter values related to the physical geometry and construction of a laser cavity and  $U$  is the envelope of the optical field [34, 35]. In the optics literature, it is standard to express the space-like variable as  $T$  and the time-like variable

as  $Z$  to denote an optical coordinate frame, but we retain the standard  $t$  and  $x$  coordinate system to avoid confusion.

To solve Eq. (5.17), we employ a 256-mode pseudo-spectral method based on the Fast Fourier Transform with  $x \in [-10, 10)$  for spatial discretization [65]. In time, a 2nd-order adaptive time stepping Runge Kutta method (`ode23` in MATLAB) is used. Unless otherwise stated, the initial condition is a hyperbolic secant. CQGLE has many non-trivial stationary and periodic solutions for a number of different parameter values. However, in all cases this level of spatial discretization and temporal time-stepping shows no appreciable change when the resolution is increased.

The governing equations for the POD approximation of the system are

$$\frac{d\mathbf{a}}{dt} = \left(\frac{i}{2} - \tau\right) \mathbf{L}^{(1)}\mathbf{a} + \kappa \mathbf{L}^{(2)}\mathbf{a} + (\Delta x)\mathbf{\Phi}^* \left( (i + \beta)|\tilde{U}|^2\tilde{U} + (i\nu + \sigma)|\tilde{U}|^4\tilde{U} \right) - \gamma\mathbf{a}, \quad (5.18)$$

where  $\tilde{U} = \mathbf{\Phi}\mathbf{a}$ ,  $\mathbf{L}_{ij}^{(1)} = -\langle \partial_x \phi_i(x), \partial_x \phi_j(x) \rangle$ ,  $\mathbf{L}_{ij}^{(2)} = \langle \partial_x^2 \phi_i(x), \partial_x^2 \phi_j(x) \rangle$ , and  $\Delta x = L/N$  is the spacing between Fourier collocation points. Here  $\mathbf{L}^{(1)}$  and  $\mathbf{L}^{(2)}$  can and should be pre-computed for efficiency. Although the nonlinear terms could, in principle, be pre-computed in a similar fashion, i.e., as an explicit function of  $\mathbf{a}$ , with the quintic term this is sometimes more costly than direct computation, and so is a potential optimization we purposefully choose not to make. Typically Eq. (5.18) is a small system, with  $\mathbf{a}$  having a dimension less than twenty and usually around five. As a result, it is faster and more computationally efficient to evaluate Eq. (5.18) rather than Eq. (5.17).

In order to validate the method, we will examine two prototypical test cases where PDE integration is computationally expensive. Furthermore, for each of these cases the dynamics are visibly low dimensional, but due to a sudden shift in a parameter value through a bifurcation point, a non-adaptive POD model will ultimately fail. Lastly, we will compare the computational cost of PDE integration and hybrid integration in four prototypical cases.

### 5.3.1 Transitions Between Stationary Solutions

The first test case is a rapid transition between stationary solutions of the CQGLE. The parameter values used here are

$$(\tau, \kappa, \beta, \nu, \sigma, \gamma) = \begin{cases} (0.08, -0.05, 1.6, -0.1, -0.1, -\mathbf{0.2}), & \frac{500}{3} < t < \frac{1000}{3}, \\ (0.08, -0.05, 1.6, -0.1, -0.1, -\mathbf{0.1}), & \text{otherwise,} \end{cases} \quad (5.19)$$

where the value in bold changes in a step-like function with time. With these parameters, the dynamics of the system are as follows: the solution transitions from the hyperbolic secant initial condition to a triple-pulse solution. The change in  $\gamma$  destroys the triple-pulse solution and generates a low-amplitude single-pulse solution. The decrease in  $\gamma$  for  $t > \frac{1000}{3}$  allows the pulse to grow in amplitude, but the system remains on the branch of single-pulse solutions. Therefore, this system also demonstrates hysteresis.

Figure 5.3 shows the amplitude of the solution of the PDE system in (a) as well as the




Figure 5.3: Results from the PDE and hybrid integrators for the parameter values in Eq. (5.19). As shown by the PDE solution in (a), this system transitions from the initial condition to a triple-pulse solution and then to two different single-pulse solutions. The black lines in (d) and (e) are plots of the energy of the PDE solution. The plots in (b) and (d) show the amplitude and energy of the solution obtained using the hybrid integrator with  $R_D$  as the metric for returning to the full model. In both plots, the red portions denote the times when PDE integrators were used. Note that with the  $R_D$  metric, the hybrid integrator resamples after every change in the system parameters. The plots in (c) and (e) show the same information but using  $R_P$  as the criterion for using the reduced model. Unlike the  $R_D$  metric, the  $R_P$  metric does not always resample when the parameters change. Overall, there is good quantitative agreement between the full PDE and both hybrid integrators.

solutions computed via the hybrid methods in (b) and (c) using the  $R_D$  and  $R_P$  metrics respectively. The energy ( $\|U\|^2$ ) of the hybrid methods are shown in (d) and (e) as a function of time in blue and green. The black line in both plots is the energy of the PDE solution. In all the plots, regions in red denote time intervals where PDE integration was used and those in black or without background shading denote intervals where the POD/DMD model was used. As expected both models spend the majority of the time running the reduced order model and only revert back to PDE simulations during the transient periods.

The major difference between the  $R_P$  and  $R_D$  hybrid models is that the  $R_D$  model resamples whenever the solution deviates from the state where data was collected. In the present case, it resamples around  $t = \frac{1000}{3}$  despite the fact that the two solutions are on the same solution branch when the  $R_P$  hybrid integrator does not. Indeed a non-adaptive POD basis starting with the data collected at  $t = \frac{500}{3}$  can recover the remainder of the temporal dynamics. The benefit of the  $R_D$  hybrid integrator, however, is that it does capture the transition more accurately because it has reverted to the PDE at those points.

In either case, the hybrid integrators converge very rapidly to the PDE solution at each of the fixed points, although in both cases there are errors that appear during the transition

regions. It should also be noted that, although it does not appear in Fig. 5.3, these transition periods introduce an overall phase difference between the PDE and hybrid models. Due to the phase invariance of CQGLE however, this phase difference has no impact on the resulting dynamics in either model.

### 5.3.2 Transitions Between Stationary and Periodic Solutions

The next test case is the transition between stationary solutions and breather solutions. In this case, we use the following parameters

$$(\tau, \kappa, \beta, \nu, \sigma, \gamma) = \begin{cases} (0.08, 0, 0.66, -0.1, -0.1, -\mathbf{0.1}), & 500 < t < 1500, \\ (0.08, 0, 0.66, -0.1, -0.1, -\mathbf{0.2}), & \text{otherwise.} \end{cases} \quad (5.20)$$

Again all the parameter values, with the exception of  $\gamma$ , are fixed. Due to the changing value of  $\gamma$ , the system will transition from the initial condition to a single-pulse stationary solution, to a breather solution, and back to a stationary solution. Unlike the previous example, for  $\frac{2000}{3} < t < \frac{4000}{3}$  the system has a relatively large transition region between the stationary solution and the breather solution as well as sustained oscillations in amplitude. In essence, despite still being low dimensional the attractor in this case is more complex.

The results of the PDE integrator along with the two hybrid integrators are shown in Figure 5.4. Again, the first row shows a plot of the solution amplitude obtained by the PDE, the hybrid integrator with  $R_D$ , and the hybrid integrator with  $R_P$ . The hybrid integrators accurately capture the dynamics that occur in the full PDE including the switch from stationary to periodic solutions both qualitatively, in terms of the dynamics of the solution amplitude, and quantitatively, in terms of the  $L^2$ -norm. These results should be compared to Figure 5.1b, where the non-adaptive reduced order model fails to reproduce this solution through the first bifurcation point.

The segments in red indicate where PDE integrators were required, and are again the primary difference between the  $R_D$  and  $R_P$  hybrid methods. Similar to the previous section, the  $R_D$  integrator reverts to the PDE during the transition from the breather dynamics back to the stationary solution at  $t = 1500$ . This again highlights that one primary difference between the two integrators is that the  $R_D$  metric becomes large if the solution takes on dynamics different from the ones present when the data was originally sampled. In this case however, the resampling does prove beneficial because the  $R_D$  model for the stationary solution requires fewer POD and DMD modes than the  $R_P$  model, which still retains enough modes to represent breather dynamics. As a result, although both are equally accurate, the  $R_D$  model is faster than the  $R_P$  model for  $t > 1500$ .

### 5.3.3 Comparison of Computational Cost

In this section, we compare the running times of a standard Runge-Kutta based adaptive timestepper for PDEs, the hybrid integrator using  $R_D$  as the metric, and the hybrid integrator using  $R_P$  as the metric. To perform these tests, MATLAB was run in single-core mode and ten iterations of the three numerical integration schemes were run. Four test

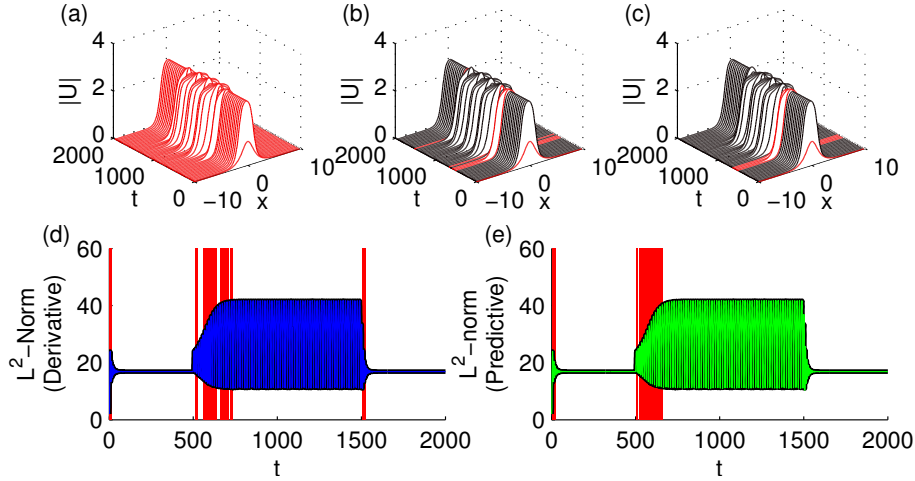


Figure 5.4: Results from the PDE and hybrid integrators for the parameter values in Eq. (5.20). In this case, the system transitions from a hyperbolic secant initial condition, to a non-stiff single-pulse solution, to a time-periodic breather solution, and back to the original non-stiff single-pulse solution. Plots (a), (b), and (c) show the solution amplitude as a function of time and space for the PDE, hybrid integrator with  $R_D$  metric, and hybrid integrator with  $R_P$  metric. Solutions in red denote snapshots where the full PDE integrator was used, and solutions in black were obtained during times the POD integrator was used. Plots (d) and (e) show the energy of the solution for the PDE, shown in black for both plots, and the two hybrid integrators which are shown in blue and green for  $R_D$  and  $R_P$ . With either metric, the hybrid integrator agrees quantitatively with the full PDE, but the choice of metric does impact exactly when the reduced or full integrator is called.

cases were run: the stationary non-stiff case uses the parameter values  $(\tau, \kappa, \beta, \nu, \sigma, \gamma) = (0.08, -0.05, 1.6, -0.1, -0.1, -0.1)$ , the stationary and stiff case uses the parameter values  $(\tau, \kappa, \beta, \nu, \sigma, \gamma) = (0.08, 0, 0.66, -0.1, -0.1, -0.2)$ , and the periodic as well as non-convergent case uses the parameter values  $(\tau, \kappa, \beta, \nu, \sigma, \gamma) = (0.08, 0, 0.66, -0.1, -0.1, -0.2)$ . In the periodic, non-convergent, and non-stiff stationary cases,  $t \in [0, 500]$ , but in the stiff and stationary case,  $t \in [0, 25]$ . To generate the non-convergent case, the relevant metric was artificially set to infinity after it was computed. In this case, the hybrid integrator must constantly check for the existence of low dimensional dynamics (e.g., computing the POD and DMD modes) and frequently map between the low dimensional POD space and the high dimensional Fourier space. Therefore, this serves as an upper bound on the computational overhead associated with the hybrid integrator.

The results of this study are outlined in Table 5.1. The table shows the mean and standard deviation of the runtime for the PDE and both hybrid integrators. In the three cases where a reduced model can be used, the hybrid integrator reduces the integration time by a significant fraction, at least by a factor of two and occasionally up to a factor of ten, using either metric. This includes stationary solutions, both stiff and non-stiff, and

Table 5.1: Comparison of computational costs for a full PDE integrator and the hybrid integrators with the  $R_D$  and  $R_P$  metrics using the methodology outlined in Sec. 5.3.3. The table shows the runtime of each algorithm in seconds as well as the fraction of the problem time (i.e. the  $t$  in the Eq. 5.17) that the hybrid integrator spent running the reduced model. The first three rows demonstrate the speed advantages of the hybrid integrator when the majority of the time is spend in the POD model. The last row is a manufactured worst case scenario where the POD model is never used. This demonstrates that the hybrid integrator adds an overhead of  $\sim 5\%$  due to the repeated checking of the various metrics.

Solution Type	PDE (s)	$R_D$ Time (s)	$R_D$ Frac.	$R_P$ Time (s)	$R_P$ Frac.
Stationary	$81.68 \pm 4.91$	$7.97 \pm 1.10$	0.9672	$11.20 \pm 1.12$	0.9512
Stationary (Stiff)	$614.41 \pm 26.39$	$61.53 \pm 3.46$	0.9032	$293.92 \pm 22.85$	0.5672
Time Periodic	$84.33 \pm 5.67$	$23.54 \pm 1.65$	0.8072	$15.08 \pm 1.34$	0.8712
No Convergence	$83.29 \pm 2.93$	$87.56 \pm 0.73$	0	$87.72 \pm 0.64$	0

time-periodic solutions. That said, the choice of hybrid model has a statistically significant impact on the running time. The performance of the hybrid model is directly related to the percentage of the time the hybrid integrator spends in the reduced dimensional model.

The clearest difference, and a point that warrants closer study, is in the stationary stiff case where the additional time spent by the  $R_P$  integrator in the PDE results in the  $R_P$  metric taking nearly five times as long as the  $R_D$  metric. The reason this occurs is because the  $R_P$  seems to require the solution to reach the attractor before it becomes small, while  $R_D$  will use the reduced order model to compute some of the transient. This is similar to what occurs in Fig. 5.20 during the transition from the stationary to the breather solution. Due to the shorter time-span of the stationary solution, roughly forty percent of the integration time elapses before the solution reaches a good quantitative approximation of the final steady state value. As this case highlights, the  $R_P$  hybrid model does not employ the reduced model when transient behaviors are present, and so in these situations the  $R_D$  hybrid model is demonstrably superior because it will make use of the reduced order model when possible in these regimes. This results in a sufficiently accurate but more quickly computed solution.

The final row demonstrates the penalty incurred due to the repeated low-dimensional checks. Because the  $R_P$  and  $R_D$  models are performing the same computations to calculate  $H(\mathbf{X})$ , the statistically identical nature of the results is actually a desired feature. The performance penalty for the repeated low dimensional checks is relatively minor at just around a 10% increase over the standard PDE solver. In addition, some of this performance penalty comes from the repeated allocation of memory in our implementation in MATLAB, so this premium could likely be reduced further with more aggressive optimizations. Although the SVD is potentially an expensive operation, the reduced SVD used here has an operation count of approximately  $\mathcal{O}(KMN)$  where  $K$  is the rank of the matrix  $\mathbf{X}$ , and  $N, M$  are the number of rows and columns [109, 110]. Recall that  $\mathbf{X}$  is augmented with only 20 samples per sampling interval, so the maximum value that  $M$  and therefore  $K$  can ever obtain is 79 (as we compute the SVD of  $\mathbf{X}_1^{j(m-1)-1}$  with  $m = 5$  and  $j = 20$ ). Furthermore, the previous stages with  $m = 2, 3, 4$  have a smaller data set and thus a reduced cost. As a result, the

overhead of the hybrid integrators is relatively minor compared to the cost of evolving the PDE system.

#### 5.4 Adiabatic Parameter Incrementation in the Waveguide Array Mode Locking Model

In this section, we use the hybrid integrator to evolve the waveguide array mode-locking (WGAML) model through the multi-pulsing transition by quasi-adiabatically increasing the gain supplied to the system [91]. As demonstrated in previous works [31,91], this causes the system to transition from a stationary (in amplitude) solution, to a time-periodic breather solution, to a chaotic solution, and finally to a two-pulse solution. This bifurcation sequence is particularly difficult for low dimensional models due to the high-dimensional spatio-temporal chaos that occurs [31,79,91,111]. The specifics of the WGAML model as well as the relevant parameter values are given in Eq. (1.5) and the relevant segments of Sec. 1.3.

The initial conditions in all that follows is  $A_0(x, 0) = \text{sech}(x)$ , and  $A_1(x, 0), A_2(x, 0) = 0$ . For details and physical justification, see Refs. [30,31]. The parameter values used here are

$$(D, \beta, C, \tau, e_0, \gamma_0, \gamma_1, \gamma_2) = (1, 8, 5, 0.1, 1, 0, 0, 10), \quad (5.21)$$

which corresponds to a physically realistic waveguide. Again in Eq. (1.5) as throughout this manuscript, we do not adopt the optical coordinate system of  $(T, Z)$  but retain the  $(x, t)$  coordinate system.

Similar to the CQGLE, the reduced order model of the WGAML model can be obtained via the straightforward application of the Galerkin projection, resulting in a low dimensional system of ODEs given by

$$\begin{aligned} \frac{d\mathbf{a}}{dt} = & \left( i\frac{D}{2} + g\tau \right) \mathbf{L}^{(1)}\mathbf{a} + iC\mathbf{L}^{(2)}\mathbf{a} + (g - \gamma_0)\mathbf{G}^{(0)}\mathbf{a} \\ & - \gamma_1\mathbf{G}^{(1)}\mathbf{a} - \gamma_2\mathbf{G}^{(2)}\mathbf{a} - i(\Delta x)\beta\Phi_0^*|\tilde{A}_0|^2\tilde{A}_0. \end{aligned} \quad (5.22)$$

To generate this system, we combine the snapshots taken from each of the waveguides into a single vector defined by

$$Q(x, t) = \begin{pmatrix} A_0(x, t) \\ A_1(x, t) \\ A_2(x, t) \end{pmatrix}, \quad (5.23)$$

and generate a set of POD modes for  $Q$ . Therefore, if we have the data sets  $\mathbf{A}_0 \in \mathbb{C}^{N \times M}$ ,  $\mathbf{A}_1 \in \mathbb{C}^{N \times M}$ , and  $\mathbf{A}_2 \in \mathbb{C}^{N \times M}$ , then we acquire the POD modes from  $\mathbf{Q} \in \mathbb{C}^{3N \times M}$ . With this notation,  $\tilde{Q} = \Phi\mathbf{a}$ . For notational simplicity, we also take  $\Phi_0 = [\mathbf{I}, \mathbf{0}, \mathbf{0}]\Phi = [\phi_{0,1}, \phi_{0,2}, \dots]$ ,  $\Phi_1 = [\mathbf{0}, \mathbf{I}, \mathbf{0}]\Phi = [\phi_{1,1}, \phi_{1,2}, \dots]$ ,  $\Phi_2 = [\mathbf{0}, \mathbf{0}, \mathbf{I}]\Phi = [\phi_{2,1}, \phi_{2,2}, \dots]$ , and  $\tilde{A}_0 = \Phi_0\mathbf{a}$ . Here  $\mathbf{I}$  and  $\mathbf{0}$  are the identity matrix and the zero matrix of the appropriate sizes. Then the operators in Eq. (5.22) are:  $\mathbf{L}_{i,j}^{(1)} = \langle \partial_x \phi_{0,i}, \partial_x \phi_{0,j} \rangle$ ,  $\mathbf{L}^{(2)} = 2(\Delta x)\Re(\Phi_0^*\Phi_1 + \Phi_1^*\Phi_2)$ , and  $\mathbf{G}^{(i)} = (\Delta x)\Phi_i^*\Phi_i$ . Here  $\langle u, v \rangle = (\Delta x)u^*v$ , which is spectrally accurate due to the periodic boundary conditions and smooth solution [112]. Similarly, dynamic mode decomposition is performed on  $Q$  and not the individual waveguides.

Note that this is different from the approach taken in Sec. 4.1 where each waveguide was considered separately. The reason the waveguide dynamics were considered separately in that section was to make the reduced order model more flexible to changes in the relative phase between the waveguides as the gain parameter,  $g_0$ , changed. Dynamic mode decomposition, as an equation free method, can only represent the system in its current state, i.e., it is unable to account for parameter values apart from those at which data was taken. As a result, the main benefit of treating the waveguides separately is negated, and it is more computationally efficient to form a reduced model for the system as a whole.

Although the hybrid integrator performs best when the dynamics are constrained to a single attractor, it can also be applied to accelerate parameter space sweeps and exploration. To do so for the WGAML model, we increment the gain,  $g_0$ , in a quasi-adiabatic fashion, i.e. in small discrete steps with the steps well separated in time. For instance, here we use

$$g_0(t) = 2.0 + 0.1 \text{ floor} \left( \frac{t}{1000} \right), \quad (5.24)$$

where, borrowing MATLAB notation,  $\text{floor}(x) = \lfloor x \rfloor$ . For  $t \in [0, 10^4]$ , this could in theory capture all four qualitative behaviors. The 1000 round trips between increments in  $g_0$  provides sufficient time for the stationary and breather solutions to adjust to the shift in the attractor. For the formation of the double-pulsed solution, on the order of 10,000 round trips are often required. As a result, the adiabatic increase used here will not completely capture the complete multi-pulsing transition.

Figure 5.5 demonstrates the application of the hybrid integrators to the WGAML model with the quasi-adiabatic gain profile in Eq. (5.24). The regions outlined in red are those where a full integrator was used, and the regions in black are those where the reduced integrator was used. As before, both hybrid integrators are able to completely represent the dynamics of the stationary solutions and time-periodic solutions. Indeed, for this situation the vast majority of the computational savings are obtained in these two regimes. As before, the  $R_P$  and  $R_D$  metrics force the hybrid integrator back to the full model at different times, but both allow the solutions to qualitatively capture the dynamics observed in the PDE model. This includes capturing the Hopf bifurcation which initiates the transition from the stationary solution branch to the breather solution branch.

The difference in the metrics manifests itself in the instances when the hybrid integrator returns to the full PDE. Along the branch of stationary solutions (approximately  $t < 5000$ ), it is now the  $R_P$  metric that resamples at every change in  $g_0$  while the  $R_D$  metric retains the original set of modes for longer periods of time. Although the stationary solutions are constant in amplitude, the change in  $g_0$  induces both a slight change in the mode amplitude and a slight change in the unplotted phase evolution. The  $R_P$  metric appears to be more sensitive to this phase slip than the  $R_D$  as it shows a larger relative increase than the  $R_D$  metric. It is likely that with proper tuning of the tolerances, both metrics could be made to resample for any change in  $g_0$ , but regardless both hybrid integrators accurately capture the stationary branch of solutions and the switch to the time-periodic breather solution.

Unlike Section 5.3, where the dynamics were nearly ideal for the hybrid integrator, the dynamics here include chaotic solutions where only a small fraction of the time is spent in the reduced order model. Furthermore, it is arguable that even this small fraction is unwar-

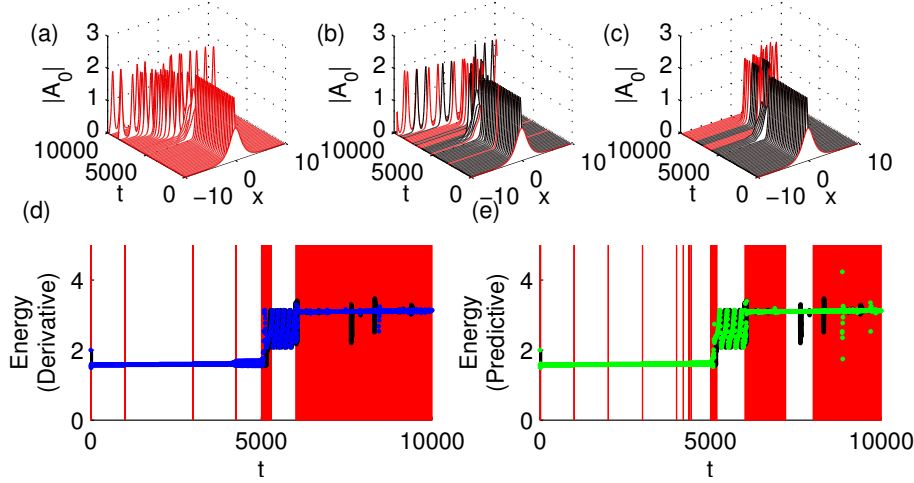


Figure 5.5: Plot of the solution amplitude in the 0th waveguide with gain increased quasi-adiabatically (Eq. (5.24)) as computed by (a) `ode23` (b) the hybrid integrator with the  $R_P$  metric, and (c) the hybrid integrator with the  $R_D$  metric. The system transitions from the stationary solution, to the breather solution, and finally to the chaotic solution. The time spent in the chaotic solution depends primarily upon the initial condition when the system enters the chaotic regime. Due to this sensitive dependence on the initial condition, the hybrid integrators are unable to accurately reproduce the dynamics in this regime. Plots (d) and (e) show the energy of the solution as a function of time with the energy of the PDE solution given in black. Because of the chaotic nature of the solution, there is a visible difference between the full and hybrid integrators, but there is still good qualitative agreement. As implied by the red background that denotes where the full PDE integrator was used, the majority of the computational speedup comes from the use of the reduced model for the stationary and breather solutions and there is almost no opportunity for a speedup in the chaotic regime.

ranted, i.e. the ROM should never be used in these circumstances. Due to the chaotic nature of the dynamics, the three integrators generate visibly different solutions. This limitation cannot be overcome as even switching between a 2nd- and 4th-order Runge-Kutta scheme (`ode23` and `ode45`) will visibly change the results. However, the qualitative dynamics are in agreement with a nearly two-pulse solution oscillating between quasi-stationary behavior and rapid translations in the domain.

In the interest of completeness, it should be remarked that the hybrid integrators shown here have a proclivity to erroneously infer the existence of a stable solution from data that is only meta-stable due to the finite and often short sampling period. For example, suppose there is an unstable eigenfunction with a growth rate that is negligible on the time-scale of the sampling period. This eigenfunction will be neglected by the hybrid integrator and the resulting solution will be artificially stable. In parameter range sweeps like the one performed here, this causes the hybrid integrator to miss the presence of a bifurcation point

until the bifurcation parameter is far enough beyond the bifurcation point that the unstable eigenfunction is retained in the reduced model. This is not an issue of step size or initial value of  $g_0$ , but will occur if, say, the region near the Hopf bifurcation is examined closely. This can be remedied by increasing the sampling period, but this also increases the computational overhead of using the hybrid integrator.

Although it may require some tuning, the hybrid integrator is still an efficient method for performing parameter range sweeps like the one demonstrated here. Using the hybrid integrator reduced the computational cost of the parameter range sweep by approximately a factor of two with either the  $R_D$  or  $R_P$  metric. As a result, although in a realistic case the computational speed up is likely to be limited by legitimately high-dimensional regimes such as the chaotic regime of the WGAML model, the presence of coherent structures is often enough to offset the computational overhead of the method and provide some measure of speedup.

## 5.5 Chapter Summary and Future Outlook

In this chapter, we proposed a hybrid PDE/reduced order model integrator for efficiently evolving pattern forming systems in time. Although the choice of ROM is potentially quite general, we used a combination of the Proper Orthogonal Decomposition and Dynamic Mode Decomposition because both models, which make different assumptions about the nature of the underlying system, can be easily generated from the same data. For both the POD and DMD ROMs, we defined a metric in Eq. (5.11) that indicated when the PDE was low dimensional. Next, we proposed two additional metrics,  $R_D$  and  $R_P$  in Eq. (5.14) and Eq. (5.15) to determine when the reduced order models are no longer valid. These criteria are the crux of whether or not the hybrid integrator is both accurate and computationally efficient.

The choice between the  $R_D$  and  $R_P$  metrics depends on problem specific knowledge and intuition. Both accurately capture the rapid changes in the system that occur as a parameter is varied, and allow the hybrid integrator to reproduce qualitatively different behaviors that a static POD model with a single set of data could not. Furthermore, both significantly outperform the PDE integrator for the test cases discussed earlier.

However, to optimize the performance of the hybrid integrator it is critical to select the correct metric. The  $R_P$  model will use a low dimensional model as long as it can properly capture and represent the future dynamics of the system as approximated by DMD. As a result, during transitions from a higher dimensional attractor to a lower dimensional attractor as in this section or for solutions nearby on the same branch of solutions, the  $R_P$  hybrid model does not need to revert to the PDE for additional data. The drawback of this is that the  $R_P$  model often retains more POD and DMD modes than is strictly necessary to represent the new and lower dimensional solution and will often attempt to capture the dynamics of transient regions with a low dimensional model. The  $R_D$  hybrid model, on the other hand, reverts to the PDE solver whenever the dynamics of the reduced model differs from those that existed when data was taken. As a consequence of the more frequent resampling, the resulting reduced order models tend to be of lower dimension than in the  $R_P$  model. As demonstrated previously, the choice of the correct metric can generate an

additional five-fold speedup in the integration and, therefore, cannot be neglected.

In broader terms, the aim of this chapter was to demonstrate that there are clear benefits, both computational and in terms of robustness, to linking the data collection and POD-mode generation aspects of reduced dimensional models, i.e., generating reduced order models “on the fly”, and the hybrid integrator was just an example of this. This sort of methodology was also used with great success in the trust region POD method [12], and certainly such an idea could be extended to other computational techniques for reduced order models, like the certified reduced basis functions, that provide a more rigorous approximation of the reduced order model error rather than the heuristics used here. Another potential direction is to include some of the nonlinear techniques for reduced order models, such as the one proposed by Sondag et. al. [113], rather than the linear POD method. Overall, the idea of incorporating reduced order models seamlessly with existing full PDE integrators, controllers, etc. is both robust and flexible. Furthermore, it shows great promise as a technique to allow the generation of effective reduced models, which is sometimes more of an art than a science, to be done in a “black box” fashion at the cost of manageable computational overhead.

## Chapter 6

## DIMENSIONALITY REDUCTION IN STANDING SURFACE WATER WAVES

In previous chapters, we applied dimensionality reduction techniques exclusively to problems in mode-locking. There, the low dimensional nature of the mode-locked pulses makes them almost an ideal application for dimensionality reduction techniques. In this chapter, we explore dimensionality reduction in a different application, namely in time-periodic standing surface waves.

Surface water waves have played a central role in shaping applied mathematics. Indeed, major advances in perturbation theory, integrable systems, nonlinear waves, and pattern formation were inspired by the study of water waves [114,115]. From a computational and theoretical point of view, water waves are challenging due to a non-local Dirichlet-to-Neumann operator in the equations of motion [116,117], the nonlinear interaction of dynamics at small and large scales, and the possibility of finite-time singularity formation [118,119]. From a practical point of view, improving our understanding of ocean wave dynamics could help minimize damage due to rogue waves [120] and tsunamis [121,122] and aid in the design of coastal power plants that extract energy from ocean waves [123], to list but a few applications.

In this chapter, we will only consider the dynamics of standing water waves, defined to be spatially and temporally periodic solutions of the free-surface irrotational Euler equations driven by gravity. Such waves were first studied by Lord Rayleigh [124], who incorporated time evolution in the asymptotic techniques developed by Stokes for traveling waves of permanent form [125–127]. Penney and Price [128] extended Rayleigh’s work to 5th order and conjectured that the largest amplitude standing wave would develop sharp 90 degree crest angles each time the fluid comes to rest. Taylor performed wave tank experiments [129] to confirm that large amplitude standing waves form fairly sharp crests close to 90 degrees, but was skeptical about the analysis of Penney and Price. Steering clear of the question of singularity formation, Tadjbakhsh and Keller [130] studied standing waves of finite depth while Concus [131] and Vanden-Broeck [132] considered the effect of surface tension. Small amplitude standing waves were proved to exist by Iooss, Plotnikov and Toland [133] using Nash-Moser theory to overcome a small-divisor problem that arises in the perturbation analysis. Large amplitude standing waves have been computed by Mercer and Roberts [134,135]; Smith and Roberts [136]; Tsai and Jeng [137]; Bryant and Stiassnie [138]; Schultz, Vanden-Broeck, Jiang and Perlin [139]; Okamura [140,141]; Wilkening [89]; and many others. Due to the difficulty of maintaining accuracy in nearly singular free surface flow calculations, many of these studies reached different conclusions regarding the form of the largest amplitude standing wave. Mercer and Roberts [134] predicted interior crest angles as sharp as 60 degree; Schultz et. al. predicted the formation of a cusp rather than a corner; and Okamura [140,141] found that Penney and Price’s conjecture of a 90 degree corner holds.




Figure 6.1: A sketch of the water wave problem at time  $t$  with mean fluid depth  $h$ . The shape of the fluid surface is described by  $\eta(x, t)$ . The velocity potential in the fluid is given by  $\Phi(x, y, t)$ , and the velocity potential on the surface is  $\varphi(x, t) = \Phi(x, \eta(x, t), t)$ . All variables are  $2\pi$ -periodic in  $x$ .

Recently, Wilkening [89] showed that these discrepancies were due to lack of sufficient spatial resolution. Refined calculations reveal that the Penney and Price conjecture is false due to a breakdown of self-similarity [118, 119, 142, 143] at the crests of very large amplitude standing waves [89]. As one progresses into this large amplitude regime, standing waves develop increasingly complex behavior at small scales. Thus, the assumption that a “largest amplitude” standing wave exists, terminating the bifurcation curve through the formation of a geometric singularity, appears to be incorrect. There is also a large literature on pattern formation in Faraday waves, which are standing waves in a container that oscillates vertically with a specified frequency; see for example [144–146].

What we aim to show is that despite their complex spatial structure, the dynamics of the time periodic standing waves are low dimensional even near the breakdown of self-similarity. In particular, only a few spatial modes are needed to reconstruct the waves computed in Refs. [89, 147] and other similar computations in Refs. [134, 135]. Furthermore, the POD modes provide a sufficiently good approximation of these ideal modes that the resulting reduced order models are both accurate and fast. In addition, we demonstrate that such models are equally applicable in either fluids of infinite or finite depth.

The remainder of the chapter is outlined as follows: Section 6.1 describes the governing equations for the standing surface water wave problem. Section 6.2 gives a brief description of the extensions of the POD technique we will require. Section 6.3 describes the application of the POD method to surface waves on water of infinite depth including the reduced dimensional formulation, the set of POD modes used, the analysis of a single periodic solution, and the reproduction of the entire branch. Section 6.4 is similarly structured but for the finite depth case. Lastly, Section 6.5 contains a summary of the chapter and proposed future directions based on the results contained within.

## 6.1 Governing Equations

We consider a two-dimensional spatially-periodic irrotational ideal fluid with a freely evolving surface and a mean fluid depth of  $h$  above a flat bottom. This leads to Laplace’s equation on the domain occupied by the fluid with no-flux boundary conditions on the fluid bottom

and the kinematic condition on the fluid surface. The equations describing the evolution of the fluid surface and velocity potential are:

$$\Delta\Phi = 0, \quad 0 < y < \eta \quad (6.1a)$$

$$\Phi_y = 0, \quad y = 0 \quad (6.1b)$$

$$\Phi = \varphi, \quad y = \eta \quad (6.1c)$$

$$\eta_t + \eta_x \Phi_x = \Phi_y, \quad y = \eta \quad (6.1d)$$

$$\varphi_t = P \left[ -\eta_x \Phi_x \Phi_y - \frac{1}{2} \Phi_x^2 + \frac{1}{2} \Phi_y^2 - g\eta \right], \quad y = \eta. \quad (6.1e)$$

where

$$P[f(x)] = f(x) - \frac{1}{L} \int_0^L f(\alpha) d\alpha. \quad (6.1f)$$

$\Phi(x, y, t)$  is the velocity potential in the fluid,  $\varphi(x, t)$  is the velocity potential on the fluid surface, and  $\eta(x, t)$  is the fluid elevation. A diagram of this problem and the relationship between the various variables is given in Figure 6.1. We further impose that the solution be periodic in space with the period,  $L = 2\pi$ , and we assume the problem has been appropriately scaled such that  $g = 1$ , and we neglect surface tension. The projection operator,  $P$ , used in Eq. 6.1e and defined in Eq. 6.1f, sets the mean velocity potential to be zero at the free surface. This does not alter the dynamics as the velocity potential is only defined up to a constant factor but eliminates a degree of freedom in the computation of time-periodic solutions.

Solving Eq. 6.1 directly is computationally intensive as Laplace's equation must be solved on a domain with a free upper boundary where the evolution of the domain and the boundary conditions involve nonlinear partial differential equations. As a result, a number of methods have been developed to evolve the surface height and velocity potential without directly solving Eq. 6.1a throughout the interior of the fluid. Although other formulations of the problem exist, the choice of surface height and velocity potential are popular because, as observed by Zakharov, they are canonical variables in a Hamiltonian formulation of the problem [116, 148]. Later work by Craig and Sulem introduced a Dirichlet-to-Neumann operator (DNO) that relates the value of the velocity potential on the surface to its normal derivative [149]. In that work, the DNO was approximated via a Taylor expansion [149], but there are other solution methods that rely on different approaches such as expressing the solution spectrally [150] or computing the DNO with boundary integral methods that exploit the Green's function of the Laplace operator [151]. Domain transformations like conformal mapping techniques [152, 153] are another approach for computing the DNO that is extremely effective for water waves that remain smooth (without forming wave crests with regions of high curvature). However, if the wave develops such crests, as is the case in many of the solutions considered here, the collocation points in the conformal mapping spread apart near the crest rather than clustering together there, and the number of Fourier modes needed to represent the solution grows extremely rapidly. In particular for a wave with a moderately sharp crest, 24000 Fourier modes were needed in the conformal mapping technique to achieve the same accuracy as a 512 mode Fourier approach using boundary

integrals. Thus, the latter approach is faster for computing standing waves of extreme form.

To compute the DNO, we adopt the boundary integral approach used in Ref. [89]. For simplicity, this section focuses solely upon the infinite depth case. However, the finite depth case can be handled using a slight modification of the integral operator [147]. In the infinite depth case, the velocity potential in the domain is expressed as

$$\Phi(x, y) = \frac{1}{2\pi} \int_0^{2\pi} K_1(x, y, \alpha) \mu(\alpha) d\alpha, \quad (6.2a)$$

$$K_1(x, y, \alpha) = \text{Im} \left[ \frac{1 + i\eta'(\alpha)}{2} \cot \left( \frac{(x - \alpha) + i(y - \eta(\alpha))}{2} \right) \right], \quad (6.2b)$$

$$\varphi(x) = \frac{\mu(x)}{2} + \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_1(x, \alpha) \mu(\alpha) d\alpha, \quad (6.2c)$$

$$\tilde{K}_1(x, \alpha) = \begin{cases} \text{Im} \left[ \frac{1 + i\eta'(\alpha)}{2} \cot \left( \frac{(x - \alpha) + i(\eta(x) - \eta(\alpha))}{2} \right) - \frac{1}{2} \cot \left( \frac{x - \alpha}{2} \right) \right], & \alpha \neq x, \\ \frac{-\eta''(x)}{2(1 + \eta'(x)^2)}, & \alpha = x. \end{cases} \quad (6.2d)$$

where  $\mu(x)$  is the dipole density in the domain and is initially unknown. Eq. 6.2a expresses  $\Phi$  in the fluid as a double-layer potential over the free surface. This complex form arises when the normal derivative of the Newtonian potential is summed over periodic images to reduce the problem from the whole real line to the interval  $[0, 2\pi)$ . The kernel  $\tilde{K}_1$  is a regularized version of  $K_1$  for points  $y = \eta(x)$  on the surface.  $\tilde{K}_1$  is continuous at  $\alpha = x$ .

To determine  $\mu$ , the limit as  $(x, y) \rightarrow (x, \eta(x))$  is taken, yielding the relationships between  $\varphi$  and  $\mu$  given in Eq. 6.2c. The operator in Eq. 6.2c, which is linear in  $\mu$ , is then expressed in matrix form and inverted to obtain  $\mu$  from a given  $\varphi$ . Once  $\mu(x)$  is known, the partial derivatives of  $\Phi$  are obtained by differentiating and evaluating Eq. 6.2a on the surface. For more details, see Refs. [89, 147] and the references therein. In what follows, we generate different reductions of Eq. 6.1 using the Galerkin projection for infinite depth in Sec.6.3 and for finite depth in Sec.6.4. The reason for the different approaches is to highlight different techniques for approximating the DNO. In the infinite depth case, the relationship between  $\mu$ ,  $\varphi$ , and  $\Phi$  in Eqs. 6.2 is used to compute the DNO. Although this same technique will work in finite depth, there we compute a POD approximation of the fluid interior by solving a suitably transformed version of Eq. 6.1a to obtain the same information.

## 6.2 Extensions of the Proper Orthogonal Decomposition

In this section, we discuss extensions of the Proper Orthogonal Decomposition that will be used in the subsequent sections. For details on the basic POD method, see Sec. 1.2. To review, the POD is a dimensionality reduction technique that takes a set of data and generates the  $M$  dimensional subspace that best represents that data set in the  $L^2$  sense. Normally we define the accuracy of a POD basis in terms of its combined energy. However, one can also directly define the error of a POD representation. Suppose for simplicity that the data are drawn from a single solution of Eq. 6.1 for  $t \in [0, T]$ . Then the error of an

$M$ -mode POD representation can be defined as

$$\text{Error}_{\text{POD}} = \int_0^T \|\psi(x, t) - \hat{\Phi}^{(M)}(x, t)\|^2 dt, \quad (6.3)$$

where  $\psi(x, t)$  is the data,  $\|\cdot\|^2 = \langle \cdot, \cdot \rangle = \int_0^L |\cdot|^2 dx$ ,  $\hat{\Phi}^{(M)}(x, t) = \sum_{m=1}^M \langle \psi(\cdot, t), \phi^{(m)} \rangle \phi^{(m)}(x)$  is the  $M$ -mode reconstruction, and the  $\phi^{(m)}(x)$  are the POD modes that serve as an orthogonal basis for that subspace. The POD basis functions  $\phi^{(m)}(x)$ ,  $m \geq 1$ , may be defined recursively via

$$\phi^{(m)} = \arg \max_{\|\phi\|=1} \int_0^T |\langle \phi(x), \psi(x, t) - \hat{\Phi}^{(m-1)}(x, t) \rangle|^2 dt, \quad (6.4)$$

with  $\hat{\Phi}^{(0)}(x, t) = 0$ . In practice, however, the singular value decomposition (SVD) is used to compute the first  $M$  modes rather than iteratively solving this sequence of nonlinear minimization problems [4, 91].

The data,  $\psi$ , are usually only known at a discrete set of points, say for example  $N$  points in space and  $K$  points in time, and can be represented by a matrix:

$$\mathbf{X} = \begin{bmatrix} \downarrow & \downarrow & \cdots & \downarrow \\ \psi_1 & \psi_2 & \cdots & \psi_K \\ \downarrow & \downarrow & \cdots & \downarrow \end{bmatrix}, \quad \vec{\psi}_k = \begin{bmatrix} \psi(x_1, t_k) \\ \psi(x_2, t_k) \\ \vdots \\ \psi(x_N, t_k) \end{bmatrix}. \quad (6.5)$$

Here  $x_n$  is the  $n$ th sampling point in space and  $t_k$  is the  $k$ th sampling point in time. If  $\psi$  is of dimension larger than one, the sampled points,  $\psi$ , must first be recast as a vector.

For a given set of data, numerical quadrature schemes are chosen based on the location of the sampling points. In particular, for points on a Chebyshev grid, Clenshaw-Curtis integration may be used [112]. For an evenly spaced grid, Newton-Cotes formulas can be used for additional accuracy. In our case, we are integrating periodic functions over their period both spatially and temporally, so the trapezoidal rule is spectrally accurate. Regardless of the numerical scheme chosen, the integral can be represented by the inner product of a weight vector  $\vec{w} = (w_1, w_2, \dots, w_N)$  for integration in space, and  $\vec{\omega} = (\omega_1, \omega_2, \dots, \omega_K)$  for integration in time. Note that different integration schemes can be used to integrate in space and time.

To obtain the POD modes

$$\begin{bmatrix} \left| \right. & \left| \right. & \cdots & \left| \right. \\ \phi^{(1)} & \phi^{(2)} & \cdots & \phi^{(N)} \\ \left| \right. & \left| \right. & \cdots & \left| \right. \end{bmatrix} = \text{diag}(\sqrt{\vec{w}})^{-1} \mathbf{U}, \quad (6.6)$$

we apply the singular value decomposition

$$\text{diag}(\sqrt{\vec{w}}) \mathbf{X} \text{diag}(\sqrt{\vec{\omega}}) = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T. \quad (6.7)$$

Here, for example,  $\text{diag}(\sqrt{\vec{w}})$  is the diagonal matrix with the square root of the elements of  $\vec{w}$  along the diagonal. Note that the resulting POD modes,  $\phi^{(1)}, \dots, \phi^{(N)}$ , are orthogonal with respect to the weighted inner product  $\langle f, g \rangle_w = \sum_{j=1}^N w_j f_j g_j$ . The solution is then approximated as a truncated sum of the POD modes,

$$\hat{\Phi}^{(M)}(x, t) = \sum_{m=1}^M a_m(t) \phi^{(m)}(x) = \sum_{m=1}^M \langle \psi(\cdot, t), \phi^{(m)} \rangle_w \phi^{(m)}(x), \quad (6.8)$$

with  $M \leq \min(N, K)$ . This truncation defines the relative error of an  $M$ -mode expansion,

$$E_M = \frac{\int_0^T \|\psi(x, t) - \hat{\Phi}^{(M)}(x, t)\|^2 dt}{\int_0^T \|\psi(x, t)\|^2 dt}. \quad (6.9)$$

In short, the error is a measure of the difference between the data and the most accurate  $M$ -mode POD reconstruction.

We determine the dimension of the POD basis by setting an error tolerance  $E_{\text{tol}}$ . Then  $M$  is chosen to be the smallest integer such that  $E_M < E_{\text{tol}}$ . Although this sort of error can also be defined for any other  $M$ -dimensional orthogonal basis, for any other linear basis the resulting error  $\tilde{E}_M \geq E_M$ . Because of this property, the POD basis is known as “proper” [4]. For vector-valued functions, each physical variable, e.g. surface height and velocity potential, can be treated either as part of a single coupled set of data [154] or as multiple independent sets of data [91]. In this problem, we chose the latter approach to allow surface height and velocity potential to vary independently. This achieves the same accuracy with slightly (two or three, depending on  $E_{\text{tol}}$ ) fewer modes. To extend the POD to multiple data sets, we used a generalization known as sequential POD (SPOD) [11]. This technique generates a set of  $M$  modes that produces a relative error smaller than  $E_{\text{tol}}$  for each data set by sequentially augmenting an existing set of POD modes.

In order to uniquely specify a periodic solution, a phase condition must be imposed. Following the approach of Wilkening [89], the phase condition is that the surface height at  $t = 0$  must be invariant under translations in space by  $\pi$  (i.e.  $\eta(x, 0) = \eta(x + \pi, 0)$ ), and the surface velocity potential must be invariant when translated in space by  $\pi$  and negated (i.e.  $\varphi(x, 0) = -\varphi(x + \pi, 0)$ ). Equivalently, at  $t = 0$  all the odd wavenumbers of  $\eta$  and all the even wavenumbers of  $\varphi$  must be zero. This phase condition is straightforward to implement in a Fourier representation. Suppose that

$$\eta(x, 0) = \sum_{j=-N/2}^{N/2-1} a_j e^{ijx} \quad \text{and} \quad \varphi(x, 0) = \sum_{j=-N/2}^{N/2-1} b_j e^{ijx}. \quad (6.10)$$

Then  $a_{2j+1} = 0$  and  $b_{2j} = 0$  for the relevant integer values  $j$ . This has the added benefit of allowing symmetries in the problem to be exploited which reduces the computational cost by a factor of four [89].

To make use of these symmetries, we filter the data in each data set and split it into two subsets. In the first subset, the Fourier transform of each snapshot is taken and all of the even wavenumbers are set the zero. Then we apply the SPOD technique to the filtered data

to obtain a basis with  $M_O$  orthonormal functions:  $\phi_O^{(1)}, \phi_O^{(2)}, \dots, \phi_O^{(M_O)}$ . In the second data set, the same Fourier transform is taken but all of the odd wavenumbers are set to zero. Applying the SPOD to this data set generates another orthonormal basis with  $M_E$  elements:  $\phi_E^{(1)}, \phi_E^{(2)}, \dots, \phi_E^{(M_E)}$ . The basis we use is  $M = M_E + M_O$  dimensional, and is the union of these two filtered bases:  $\phi_O^{(1)}, \phi_O^{(2)}, \dots, \phi_O^{(M_O)}, \phi_E^{(1)}, \phi_E^{(2)}, \dots, \phi_E^{(M_E)}$ . The union of these bases is itself an orthonormal basis. By filtering the basis in this manner, we can obtain periodic orbits that also exactly satisfy Eq. 6.10. The elements of the basis are ordered by the fraction of the  $L^2$  norm each individual element captures.

### 6.3 Infinite Depth

In this section, we consider the water wave problem for a fluid of infinite depth. As we are interested in standing waves of extreme form, we adopt the boundary integral method described above for the computation of the DNO. As mentioned previously, conformal mapping methods cause the spacing of collocation points to increase near the wave crests, which decreases the resolution of the grid exactly where it is most needed. Boundary integral methods do not suffer from this problem as the grid spacing can be controlled independently of the evolution equations. Presently, we assume uniform spacing for simplicity; however, many of the underlying solutions of the full PDE were computed with mesh refinement near the wave crest; see Ref. [89].

In order to use a boundary integral approach in a low dimensional setting, it is important to avoid solving large linear systems involving all the collocation nodes of the integral equations. In Sec. 6.3.1, we show that replacing the velocity potential,  $\phi$ , with the dipole density,  $\mu$ , in the evolution equations leads to linear systems of size  $M \times M$  instead of  $N \times N$ , where  $M$  is the number of POD modes used to represent the dynamics and  $N$  is the number of collocation points used in the boundary integral method. This greatly reduces the cost of evolving the interface. In Sec. 6.3.3, we give the evolution of the reduced model on a single periodic orbit and show that the POD basis is more accurate than the Fourier basis. Finally in Sec. 6.3.4, we show the bifurcation diagram of the branch of time-periodic standing wave solutions, and we demonstrate that the accuracy of the representation depends upon the number of modes retained.

#### 6.3.1 Problem Formulation

As written in Eq. 6.1, the two obvious variables to evolve are the surface height and surface velocity potential. Numerical methods based on this approach use a Dirichlet-to-Neumann operator (DNO) to relate the surface velocity potential to the surface fluid velocity [89, 149]. A number of other formulations exist that express the problem in other physical variables. Fenton and Rienecker [150] posed the water wave problem in terms of the surface height and a linear combination of the eigenfunctions of the Laplace equation. Baker, Meiron, and Orszag [155] developed a semi-Lagrangian approach using marker particles to evolve the interface and the vortex sheet strength to represent the fluid velocity. This approach was also used by Mercer and Roberts [134] in their study of large amplitude standing waves. We adopt a hybrid approach in which wave height and dipole density,  $\mu(x, t)$ , are evolved. The

connection to other approaches is that the vortex sheet strength is the derivative of dipole density with respect to arc-length along the curve.

In the POD setting, we represent the state of the system as a linear combination of a set of time-independent modes:

$$\eta(x, t) = \sum_{m=1}^{M_\eta} a_m(t) \eta^{(m)}(x), \quad \mu(x, t) = \sum_{m=1}^{M_\mu} b_m(t) \mu^{(m)}(x), \quad (6.11)$$

where  $a_m, b_m \in \mathbb{R}$  and  $M_\eta$  and  $M_\mu$  are the number of modes in the expansion. Using Eq. 6.2c, the velocity potential on the surface can also be represented as the superposition of  $M_\mu$  modes:

$$\varphi(x, t) = \sum_{m=1}^{M_\mu} b_m(t) \varphi^{(m)}(x, t), \quad (6.12a)$$

$$\varphi^{(m)}(x, t) = \frac{\mu(x)}{2} + \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_1(x, \alpha) \mu^{(m)}(\alpha) d\alpha, \quad (6.12b)$$

where  $\tilde{K}_1$  was given in Eq. 6.2d, and has an implicit time dependence due to  $\eta$ . By taking the time-derivative of Eq. 6.12b, the time-derivative of  $\mu$  can be found by solving

$$\begin{aligned} \frac{\mu_t(x)}{2} + \frac{1}{2\pi} \int \tilde{K}_1(x, \alpha) \mu_t(\alpha) d\alpha &= \varphi_t(x) - \frac{1}{2\pi} \int \tilde{K}_2(x, \alpha) \mu'(\alpha) d\alpha, \\ \tilde{K}_2(x, \alpha) &= \text{Re} \left\{ \frac{\eta_t(x) - \eta_t(\alpha)}{2} \cot \left( \frac{(x - \alpha) + i(\eta(x) - \eta(\alpha))}{2} \right) \right\}, \end{aligned} \quad (6.13)$$

where  $\tilde{K}_2$  becomes  $\frac{\eta_{xt}(x)}{1+\eta_x^2(x)}$  as  $\alpha \rightarrow x$ . Note that this equation is structurally identical to Eq. 6.2c, which is solved to obtain  $\mu$  from  $\varphi$  as an intermediate step in applying the DNO operator in the  $\eta, \varphi$  formulation of [89]. Nevertheless, as explained below, there is a computational advantage to working with  $\mu$  rather than  $\varphi$  in a low dimensional setting.

Using the POD modes of  $\eta$  and  $\mu$  as basis functions, a set of governing equations for the mode amplitudes may be obtained by taking the Galerkin projection of Eq. 6.1,

$$\partial_t a_i = \langle \eta^{(i)}, \Phi_y - \eta_x \Phi_x \rangle, \quad (6.14a)$$

$$\partial_t b_i + \sum_j L_1^{(ij)} (\partial_t b_j) = \left\langle \mu^{(i)}, P \left[ -\eta_x \Phi_x \Phi_y - \frac{1}{2} \Phi_x^2 + \frac{1}{2} \Phi_y^2 - g\eta \right] \right\rangle - \sum_j L_2^{(ij)} b_j, \quad (6.14b)$$

$$L_1^{(ij)} = \left\langle \mu^{(i)}(x), \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_1(x, \alpha) \mu^{(j)}(\alpha) d\alpha \right\rangle, \quad (6.14c)$$

$$L_2^{(ij)} = \left\langle \mu^{(i)}(x), \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_2(x, \alpha) \partial_\alpha \mu^{(j)}(\alpha) d\alpha \right\rangle. \quad (6.14d)$$

This reformulation of the problem is beneficial for low dimensional reproductions because it decreases the cost of computing the DNO. Using the standard boundary integral formulation with a POD representation of  $\eta$  and  $\varphi$ , solving Eq. 6.2c involves inverting a dense  $N \times N$

matrix to compute  $\mu$  from  $\varphi$ . On the other hand, computing  $\varphi$  from  $\mu$  requires a single matrix-vector multiplication. Although the matrix  $L_1$  must be inverted, the elements of  $L_1$  are given by Eq. 6.14c, and  $L_1$  itself is only  $M_\mu \times M_\mu$  because inner products may be taken before the system is inverted. In this work,  $M_\mu \ll N$  and so the resulting system is much lower in dimension. We evolve Eq. 6.14 forward in time using an 8th order explicit Runge-Kutta method by Dormand and Prince, which involves evaluating  $\partial_t a_i$  and  $\partial_t b_i$  twelve times per time step. The most expensive operation in terms of processor time is computing  $\tilde{K}_1$ , mostly due to the number of complex cotangents that arise in the operator, but this step is easily parallelized. The other expensive steps involve level 3 BLAS routines such as matrix-matrix multiplication (e.g. computing  $L_1$  and  $L_2$ ) or matrix factorization (e.g. solving for  $\partial_t b_i$ ). Even running on a single thread, the reduced model is fast in practice, requiring a matter of seconds per period with  $N = 512$ .

### 6.3.2 POD Modes

Figure 6.2 shows a bifurcation diagram of infinite depth standing water waves, computed with the technique developed in Ref. [89]. In that manuscript, numerical continuation was used to accurately compute the relevant branch of standing wave solutions using a Fourier mode amplitude as the continuation parameter. The red crosses in Figure 6.2 represent the four periodic solutions from which the data sets for the POD modes were drawn. This data consists of 801 snapshots of the surface elevation,  $\eta$ , and the dipole density,  $\mu$ , taken at 512 evenly spaced points on the interval  $x \in [0, 2\pi)$  over a single period. A complication that arises is that the POD method is only applicable to a single set of data. To generate modes for multiple data sets like we have here, extensions like the sequential POD (SPOD) technique are required.

The SPOD technique, as derived and implemented in Ref. [11], requires a tolerance and ordering of the data sets to be defined at the outset. Given this tolerance and ordering, the SPOD method begins by computing the POD modes of the first data set and retains only enough modes, say  $m_1$ , to ensure that the error in the representation of the data with the  $m_1$ -dimensional POD basis is less than the predefined tolerance. This POD basis, denoted  $\phi^{(1)}, \phi^{(2)}, \dots, \phi^{(m_1)}$ , is then saved for later use.

For each subsequent data set, the data are first projected onto the orthogonal complement of the pre-existing POD subspace. The POD procedure is performed once again on this projected data, resulting in a new set of POD modes,  $\tilde{\phi}^{(1)}, \tilde{\phi}^{(2)}, \dots, \tilde{\phi}^{(m_2)}, \dots, \tilde{\phi}^{(M)}$ , which span the orthogonal complement of the original POD subspace. The original set of POD modes is then augmented by modes taken from the new set to form the combined subspace spanned by  $\phi^{(1)}, \dots, \phi^{(m_1)}, \tilde{\phi}^{(1)}, \dots, \tilde{\phi}^{(m_2)}$ , where  $m_2$  is the minimum number of additional modes required to satisfy the tolerance of the SPOD method. Because the  $\tilde{\phi}^{(k)}$  are drawn from the orthogonal complement of the preexisting POD subspace, the augmented basis is orthogonal without the need for any additional work. As before, this augmented basis is saved for use in the next iteration. This iterative process can be repeated for any number of data sets. However, the number of modes each data set contributes is dependent upon the tolerance, the order in which the data sets are presented, and the qualitative nature of the dynamics within each data set. Therefore, many nearly identical data sets simply increase

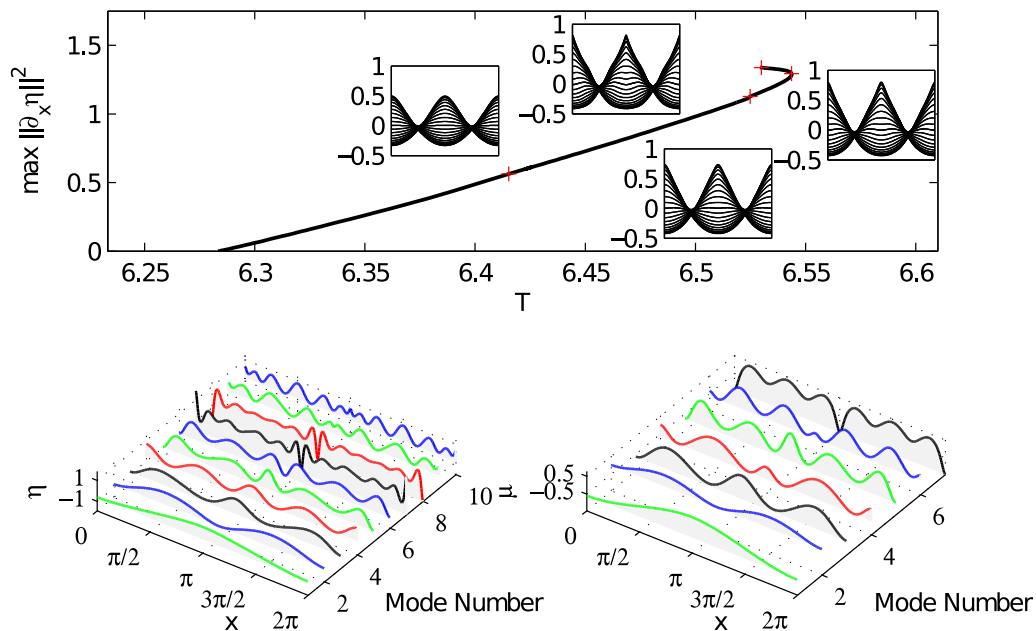


Figure 6.2: (Top) The bifurcation diagram in terms of period and  $L^2$ -norm of the standing wave solutions in fluid of infinite depth. Locations marked by an  $x$  are where data was used to generate POD modes. (Bottom) Plot of the first ten POD modes for the surface elevation and the dipole density.

the computational cost of the method while providing few, if any, additional modes. As a result, the ideal choice of data sets will be relatively small, but contain the prototypical dynamics sampled from different regions in phase space.

In this problem, we used the four data sets indicated in Fig. 6.2. Only a single data set captures low-amplitude solutions; the rest focus on larger-amplitude, sharply crested solutions. More specifically, the first three POD modes of  $\eta$  and  $\mu$  were drawn from the solution with  $T \approx 6.425$ . The remaining ten modes for  $\eta$  and eight modes for  $\mu$  were taken from the large amplitude data sets. The reason for the sparse sampling in the low-amplitude regime is due to the relatively simple dynamics that occur in that region. With the SPOD procedure, the set of POD modes from a single set of data is sufficient to reproduce all the solutions that are qualitatively similar for the tolerance levels we use here.

The benefit of the POD approach is demonstrated in Figure 6.3, which shows the error of the POD reproduction and of a Fourier reproduction for the data set with the maximum value of  $\max \|\partial_x \eta\|^2$  in Fig. 6.2. With fewer than five modes, the accuracy of both methods is similar, but neither basis can accurately reproduce the standing wave solutions of the sharply peaked variety.

Beyond the first five modes, the POD representations of both  $\eta$  and  $\mu$  are significantly more accurate than their Fourier counterparts. For example, the POD expansion is more




Figure 6.3: Convergence of the POD modes (blue line) and the Fourier modes (black line) for sharply peaked solution data. The left plot shows the convergence for the fluid surface height, and the right plot shows the convergence for the dipole density in the fluid. Here the POD modes converge much more rapidly than the Fourier modes due to the large wavenumber components of the data.

than four orders of magnitude more accurate when twenty-five modes are retained. This is due to two reasons. The first is that the sharply peaked nature of the crest requires a large number of Fourier modes to accurately represent which causes the Fourier representation to converge slowly. As discussed previously, the POD representation is not limited in this way. The second is that the later modes are drawn from data sets that are more similar, i.e. closer on the branch of solutions, to the one of interest. As a result, those POD modes are also more relevant to the data set of interest, which helps regain the convergence properties of the POD method for a single data set.

The values of the error as well as the rates of convergence in Fig. 6.3 are specific to the most sharply crested data set. If, for example, the first data set were studied instead, one would initially observe rapid convergence (as the POD modes will be drawn from that set of data) followed by near stagnation once the specified level of error was obtained. Indeed, all that the SPOD approach guarantees is that the resulting basis will represent all the data sets used to a given level of accuracy. However, similar levels of accuracy with the same twenty-five modes are observed in *all of the data points* including those solutions that, although qualitatively similar to the sampled data, did not contribute to the POD basis. Indeed, it is precisely this global increase in accuracy that we will exploit in the upcoming sections.

### 6.3.3 POD Dynamics

In this subsection, we focus on the dynamics of a single periodic solution and compare the dynamics obtained using a POD expansion with that of an equivalently sized Fourier expansion. The dynamics presented here are numerically computed periodic solutions to the reduced system of ODEs in Eq. 6.14. To compute these periodic solutions, we applied the shooting approach in Ref. [89] to the reduced order system and equated solving the




Figure 6.4: Plot of the error in the derivative of a 10-mode Fourier expansion (black-dashed), 10-mode POD expansion (blue-dashed), 22-mode Fourier expansion (black-solid), and 22-mode POD expansion (blue-solid) for the solution plotted on the right. The POD expansions are roughly one to two orders of magnitude more accurate than the equivalently sized Fourier representation at all points along the orbit.

two-point boundary value problem with minimizing the function

$$F(a_0(0), \dots, a_{M_\eta}(0), b_0(0), \dots, b_{M_\mu}(0)) = \frac{1}{2} \sum_{m=1}^{M_\eta} |a_m(T) - a_m(0)|^2 + \frac{1}{2} \sum_{m=1}^{M_\mu} |b_m(T) - b_m(0)|^2 \quad (6.15)$$

where  $a_m(t)$  and  $b_m(t)$  are the mode amplitudes in Eq. 6.11 and  $T$  is the period. This is a nonlinear least-squares problem where the initial conditions are the unknowns, and so the Levenberg-Marquardt algorithm (LMA) [156] was used to numerically obtain a solution. Due to the symmetry properties we enforced in the POD modes, as discussed in Sec.6.2, only roughly half of the modes have a nonzero initial value. In particular, only modes where  $\eta^{(m)}(x) = \eta^{(m)}(x + \pi)$  and  $\mu^{(m)}(x) = -\mu^{(m)}(x + \pi)$  can have non-zero amplitudes at  $t = 0$ . Note that this choice also enforces a phase condition on the orbit. For the orbits presented in this and the following subsections,  $F < 10^{-20}$ . In practice, the dimension of the system is small enough that finite difference methods can be used to compute the Jacobian in reasonable amounts of time, but an analytic Jacobian, which was obtained by linearizing Eq. 6.1 and projecting the linearized system onto a POD basis, was later used to confirm the results. For an initial solution guess, a solution of the full PDE computed in Ref. [89] was used. Due to the slight differences between the full and reduced models, these initial conditions yield values of only  $F \sim 10^{-16}$ , and a few iterations of the LMA were required to obtain the desired level of accuracy. To compute the branch of solutions, this initial result was continued using either the period or a POD mode amplitude at  $t = 0$  as the bifurcation parameter. Results for other points on the branch will be discussed in the following subsection. Here we restrict ourselves to the solution with most extreme form on the branch.

Figure 6.4 shows the error between the full PDE solution and the reduced solution

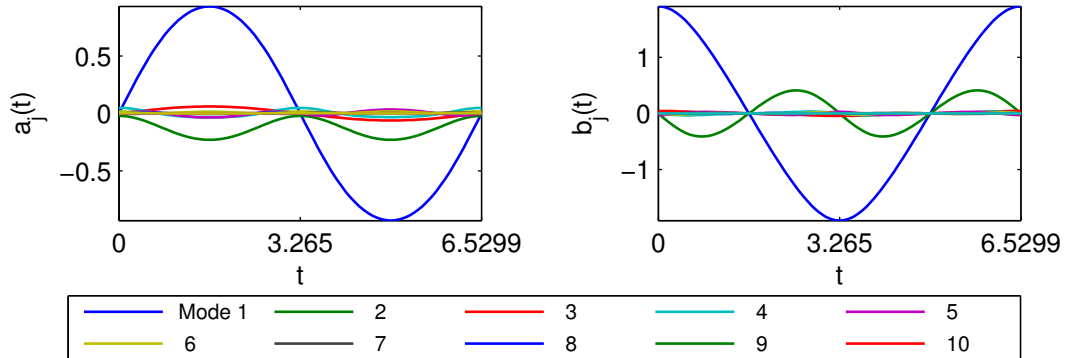


Figure 6.5: Plot of the POD mode amplitudes, which are proportional to the amount of  $L^2$ -norm each mode captures, for the evolution in Figure 6.4. The oscillatory nature of the mode amplitudes demonstrates that, although the energy in a particular POD mode decays exponentially with the mode number, the amount of  $L^2$ -norm each mode captures at a given time does not obey that exponential relationship. As an example, the first POD mode for  $\eta$  vanishes at  $t = 3.265$  while the fourth mode does not.

obtained using a 10-mode Fourier expansion, a 22-mode Fourier expansion, a 10-mode POD expansion, and a 22-mode POD expansion. The error, plotted as a function of time, is given by

$$\partial_x \eta \text{ Error}(t) = \int_0^{2\pi} \|\partial_x \eta(x, t) - \partial_x \tilde{\eta}(x, t)\|^2 dx, \quad (6.16)$$

where  $\tilde{\eta}$  is the low dimensional reproduction of the solution either in terms of Fourier modes or POD modes. It should be noted that the error chosen here is strictly for comparison. By defining error in terms of the derivative, we weight mismatches in the larger wavenumbers more strongly than mismatches in smaller wavenumbers. This makes differences in the reconstruction of the wave crest more obvious to the eye. Other norms, such as the  $H^1$ -norm, could also be used, and give similar results. In general, the POD expansions are one to two orders of magnitude more accurate than the equivalently sized Fourier expansion at all times during the period. Indeed, the 10-mode POD expansion is more accurate than the 22-mode Fourier expansion despite having less than half the degrees of freedom. Although not shown, the POD expansion is also one to two orders of magnitude more accurate in the  $L^2$ -sense as well.

The local maxima of the error in Figure 6.4 corresponds to the times when the wave crests occur. Although more accurate at all times, one advantage of the POD expansion is that several of the POD modes contain high frequency components that are specifically used to represent the wave crests. Recall that because this expansion is drawn from a number of data sets, the modes are not specific to this particular crest shape. Nonetheless at the level of accuracy in Figure 6.4, the 22 POD-mode expansion appears graphically identical to the PDE solution.

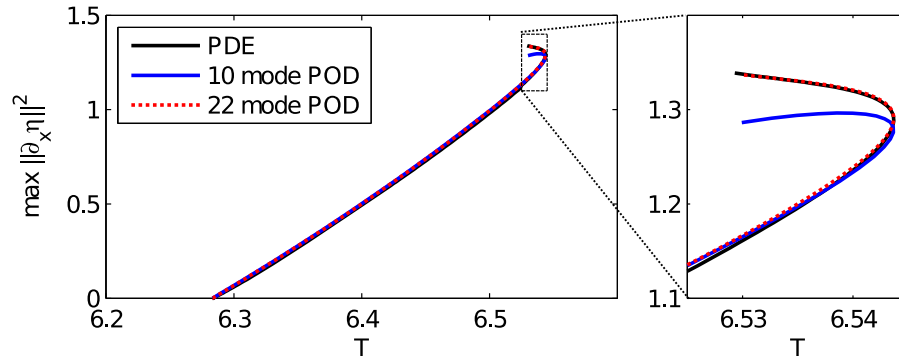


Figure 6.6: Bifurcation diagram of the time-periodic standing wave solution plotted in terms of the period,  $T$ , and the maximum  $L^2$  norm squared of the derivative of  $\eta$  over a single period. The blue curve is the bifurcation diagram obtained with 10-POD modes. The red curve is the bifurcation diagram obtained with 22-POD modes. The black curve is the bifurcation diagram of the full PDE. With more than 10-POD modes, the branch of solutions is qualitatively reproduced. However, greater quantitative accuracy is obtained by including additional modes.

Figure 6.5 shows the mode amplitudes for one period with 10 POD modes for the surface elevation, and the first 7 POD modes for the dipole density. The first few POD modes are nearly sinusoidal both in space and in time, but the higher order POD modes (those that are used to represent the wave crest) tend to have more complex temporal behaviors. The higher order modes tend to obtain their largest amplitudes near  $t = T/4$  and  $3T/4$  in order to reproduce the wave crests that occur. Away from those two times, they undergo small amplitude but rapid oscillations that are essentially sinusoidal in nature. Indeed, the linear combination of these modes enhances the accuracy of the POD expansion at those intermediate times, even if it is unclear how a single mode would contribute on its own. Therefore although the primary purpose of the higher modes is to reconstruct the wave crest, they cannot be neglected at other points.

#### 6.3.4 Bifurcation Diagram

In the previous section, a single solution was reconstructed. We now reconstruct the entire branch of solutions using the data obtained from sampling the four solutions shown in Figure 6.2. This branch was computed using 10 POD modes for the surface elevation, and 7 POD modes for the dipole density. Then the process was repeated using 22 POD modes for the surface and 18 POD modes for the dipole density.

Figure 6.6 shows the branch of solutions obtained with the 10-mode expansion and the 22-mode expansion. As shown in the figure, the 10-mode expansion is able to quantitatively reproduce most of the branch of solutions. In fact, it accurately reproduces both the branch of solutions and individual standing wave solutions with enough accuracy that, for  $T < 6.5$ ,

the differences are negligible to the eye. This is possible because that region does not include sharp wave crests. Therefore, just a few modes can capture most of the dynamics. Beyond this region, more modes are required to accurately reproduce the wave crest.

The right panel of Figure 6.6 focuses on the region of the branch where there exist sharp wave crests. Here the difference between the 10- and 22-mode expansions is clear. Although the 10-mode expansion does qualitatively capture the fold in the branch of solutions, there is a visible difference between the 10-mode solution branch and that of the full PDE. Increasing the number of POD modes to 22 decreases the error as the additional degrees of freedom allow the function to be better represented. Although a number of the additional 12 modes specifically include structures that resemble the sharp wave crest, as shown in Figure 6.5, the POD solution is more accurate at all points on the orbit.

The difference between POD solutions and the full PDE solutions is smaller when measured in the  $L^2$ -norm than in the  $H^1$ -norm. This is partially due to the fact that the  $L^2$ -norm of the error does not weight the missing higher wavenumbers as heavily as the  $L^2$ -norm of the derivative does. Another reason is that the POD modes are chosen to be optimal in the  $L^2$  sense but not in the sense of any other norm. If more accuracy is needed in a different norm, the first step (before increasing the dimension of the basis) should be to recompute the POD modes with respect to the new norm. This will increase the accuracy of the POD representation *as measured by the new norm* for a given number of modes.

## 6.4 Finite Depth

In this section, we generate a reduced-dimensional model for standing wave solutions in a fluid of unit depth ( $h = 1$ ). Although the representation in Sec.6.3 can be adapted for the finite depth case, we instead consider a different formulation that computes POD modes for the fluid interior. For this new formulation, we require access to data from the interior of the fluid in addition to surface data. This alternate formulation is advantageous when techniques, such as the finite element method, are used to solve the problem on the interior of the domain or in modeling three dimensional fluids, where boundary integral methods are difficult to implement and conformal mapping techniques do not apply.

### 6.4.1 Problem Formulation

In the finite depth case, we adopt a direct approach in order to solve Eq. 6.1. Therefore, we wish to solve for the velocity potential on the interior of the domain in order to compute the DNO. The obvious challenge with applying the POD to this problem is compiling a set of data from which to generate the POD modes. Using the unmodified physical domain makes the aggregation of data impossible because the domain itself changes in time. Therefore, the domain with the free surface must be mapped to a fixed, preferably rectangular, domain. There are a number of viable mappings that accomplish this. In two dimensions, conformal maps are a popular choice [153]. However, these approaches require a large number of grid points to represent waves of extreme form because the large negative curvature in the wave profile causes the collocation points to spread apart in physical space in precisely the region where additional mesh refinement is needed. Furthermore, approaches based on conformal maps do not extend to three dimensions.




Figure 6.7: The physical domain  $\Omega(t) = \{(x, y) : 0 \leq x < 2\pi, 0 < y < \eta(x, t)\}$  is transformed to  $\hat{\Omega} = [0, 2\pi] \times [0, 1]$  under the change of variables  $\hat{x} = x$ ,  $\hat{y} = y/\eta(x, t)$ . The problem is reformulated in terms of  $\hat{\Phi}(x, y) = \Phi(x, \eta(x, t)y) - \varphi(x)$ , which gives the boundary conditions  $\partial_y \hat{\Phi}(x, y)|_{y=0} = 0$  and  $\hat{\Phi}(x, 1) = 0$ .

Instead, we map the fluid to a domain with constant depth by adopting the sigma coordinate system in fluid mechanics [157] also known as the c-method in electromagnetics [158]. The new coordinate system, denoted by a hat symbol, is  $\hat{x} = x$  and  $\hat{y} = y/\eta(x, t)$ . This simple transformation, illustrated in Fig. 6.7, produces a time-independent domain that is still easily applicable to large amplitude standing waves and also works in three dimensions. However, a different approach would be needed to model overturning waves, and, unlike a conformal map, this mapping will generate additional terms in Eq. 6.1a.

This transformation alone is not enough to create a consistent reduced model because not all of the transformed boundary conditions are guaranteed to be satisfied. If  $\Phi(x, y)$  is represented as a linear superposition of modes, the Neumann boundary condition at the fluid bottom will be satisfied if each individual mode satisfies the boundary condition. On the other hand, the Dirichlet condition at the surface is  $\Phi(x, 1) = \varphi(x)$ , but if  $\Phi$  is expressed as a linear superposition of modes, this boundary condition cannot be guaranteed to be satisfied since it is not known *a priori* that the basis for  $\Phi$  on the surface span the same space as the basis for  $\varphi$ . Instead, we define a new variable  $\hat{\Phi}(x, \hat{y}) = \Phi(x, \hat{y}) - \varphi(x)$ . The Neumann condition remains essentially unchanged, but the Dirichlet condition on the surface is now  $\hat{\Phi}(x, 1) = 0$ , which again can be satisfied provided that each of the individual modes of  $\hat{\Phi}$  satisfies the boundary condition on the surface. We expand the surface variables as well as this new variable in a superposition of modes:

$$\eta(x, t) = h + \sum_{m=1}^{M_\eta} a_m(t) \eta^{(m)}(x), \quad \varphi(x, t) = \sum_{m=1}^{M_\varphi} b_m(t) \varphi^{(m)}(x), \quad \hat{\Phi}(x, \hat{y}) = \sum_{m=1}^{M_\Phi} B_m \Phi^{(m)}(x, \hat{y}), \quad (6.17)$$

where  $h$  is the average fluid depth and  $\eta^{(m)}$ ,  $\varphi^{(m)}$ , and  $\hat{\Phi}^{(m)}$  are the  $m$ th terms in the modal expansion of the solution.

The methodology used to compute the modes and their dimension will be covered in more detail in Sec.6.4.2. For now, we only require that the individual modes are orthogonal to one another in their respective inner product spaces, e.g.

$$\langle \eta^{(j)}, \eta^{(k)} \rangle = \int_0^{2\pi} \eta^{(j)}(x) \eta^{(k)}(x) dx = \delta_{jk}, \quad \langle \hat{\Phi}^{(j)}, \hat{\Phi}^{(k)} \rangle = \iint_R \hat{\Phi}^{(j)} \hat{\Phi}^{(k)} dA = \delta_{jk},$$

where  $R = [0, 2\pi) \times [0, 1]$ . The Galerkin projection is then used to generate a set of governing equations for the mode amplitudes:

$$\sum_{j=1}^{M_\Phi} L_{mj} B_j(t) = -\langle \hat{\Phi}^{(m)}, \varphi_{xx} \rangle_R, \quad \text{for } m = 1, 2, \dots, M_\Phi, \quad (6.18a)$$

$$\partial_t a_m(t) = \left\langle \eta^{(m)}, \frac{\hat{\Phi}_y}{\eta} - \eta_x \left( \varphi_x - \frac{\eta_x}{\eta} \hat{\Phi}_y \right) \right\rangle, \quad \text{for } m = 1, 2, \dots, M_\eta, \quad (6.18b)$$

$$\partial_t b_m(t) = \left\langle \varphi^{(m)}, -\frac{\eta_x}{\eta} \hat{\Phi}_y \left( \varphi_x - \frac{\eta_x}{\eta} \hat{\Phi}_y \right) - \frac{1}{2} \left( \varphi_x - \frac{\eta_x}{\eta} \hat{\Phi}_y \right)^2 + \frac{1}{2} \left( \frac{\hat{\Phi}_y}{\eta} \right)^2 - g\eta \right\rangle, \quad (6.18c)$$

with

$$L_{mj} = \left\langle \hat{\Phi}^{(m)}, \hat{\Phi}_{xx}^{(j)} + \left( 2y \frac{\eta_x^2}{\eta^2} - y \frac{\eta_{xx}}{\eta} \right) \hat{\Phi}_y^{(j)} + \left( y^2 \frac{\eta_x^2}{\eta^2} + \frac{1}{\eta^2} \right) \hat{\Phi}_{yy}^{(j)} - 2y \frac{\eta_x}{\eta} \hat{\Phi}_{xy}^{(j)} \right\rangle_R, \quad (6.19)$$

where superscripted variables represent individual modes, and variables without superscripts represent the current approximation of the solution in Eq. 6.17. Furthermore, we have dropped the hats from the transformed coordinate system for convenience. These governing equations are a direct consequence of applying the chain rule using the representation of the solution in Eq. 6.17. Equation 6.18a is a transformed version of Laplace's equation in Eq. 6.1a. Because the transformation used is not conformal, Laplace's equation assumes an outwardly more complicated form; however, Eq. 6.18a still enforces incompressibility in the fluid itself. Similarly, Eq. 6.18b is a transformed version of Eq. 6.1d, and Eq. 6.18c is a transformed version of Eq. 6.1e. Because the POD modes taken from a valid solution of the water wave equations must satisfy the boundary conditions (Eq. 6.1c and Eq. 6.1d), which are satisfied identically in the transformed coordinate system and consequently are dropped. In order for this approach to be viable, the entirety of the problem — surface height, surface velocity potential, and the velocity potential in the fluid — must be representable in a low dimensional context.

#### 6.4.2 Proper Orthogonal Decomposition Modes

As stated in the previous subsection, we represent  $\eta$ ,  $\varphi$ , and  $\hat{\Phi}$  in terms of POD modes. Because of the periodic boundary conditions, we collect data on an evenly spaced grid in  $x$ . Due to the Dirichlet and Neumann boundary conditions on the fluid surface and bottom, we evenly extend the domain about  $y = 0$  and use a Chebyshev grid in the  $y$  direction [112]. Therefore, a snapshot of  $\eta$  and  $\varphi$  consists of 512 evenly spaced points in  $x$ , and a snapshot of  $\hat{\Phi}$  consists of 512 evenly spaced points in  $x$  and 41 points, on a Chebyshev grid, in  $y$ . To compute the data for  $\eta$ ,  $\varphi$ , and  $\Phi$  on these grids, Eq. 6.1 was evolved using a suitably modified version of the boundary integral method, similar to Eq. 6.2. While it is not necessary to explicitly form  $\Phi$  to evolve  $\eta$  and  $\varphi$  using the boundary integral formulation, it is easy to evaluate  $\Phi$  on the grid as an extra step, to generate the POD data.

The individual snapshots are collected from periodic orbits taken with four different

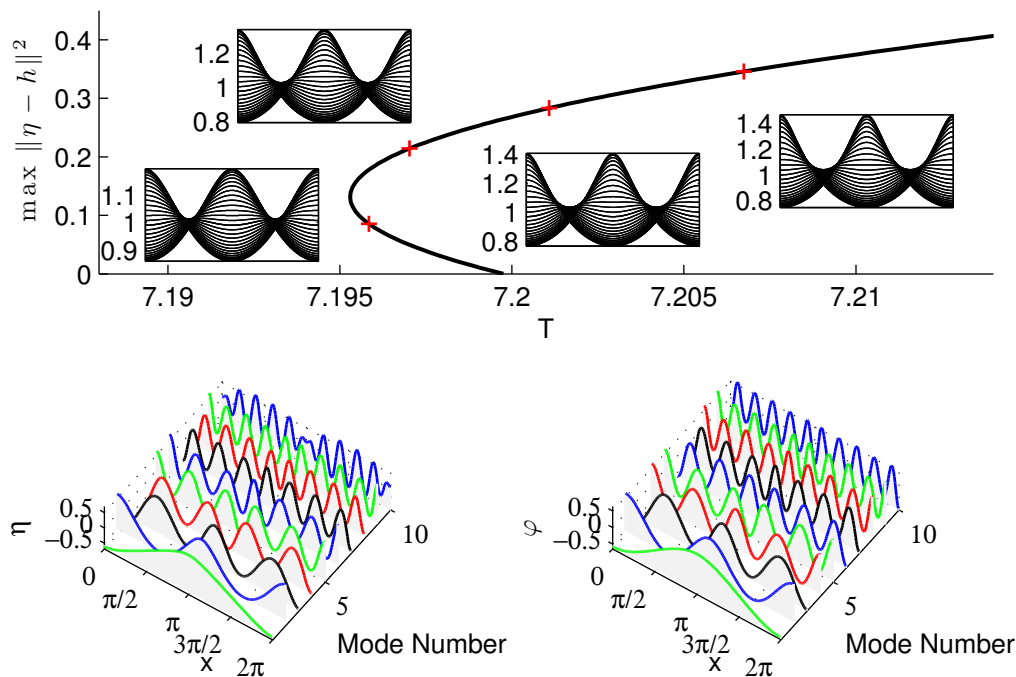


Figure 6.8: (Top) Bifurcation diagram showing  $L^2$ -norm vs. period of the standing wave solutions in fluid of unit depth. The data used to generate POD modes were taken from locations marked by a red +. Near each data point, we plot the surface height evolution over a period. (Bottom) First ten POD modes for the surface elevation and the velocity potential on the surface. Note that there is good agreement between the POD and PDE results even beyond the regimes where data was taken. This demonstrates the use of the POD method in a predictive fashion, i.e. extrapolating beyond the region where data were obtained.

periods. Figure 6.8 shows the branch of finite depth solutions in terms of period,  $T$ , and the maximum  $L^2$ -norm squared of the surface height. The red “+” signs indicate the locations from which data was taken, and the inset plots show the surface height for one period of those solutions. The bottom of Fig. 6.8 shows the ten modes for  $\eta$  and for  $\varphi$  obtained using SPOD with  $E_{\text{tol}} = 10^{-9}$ . The POD modes for the surface height and velocity potential are similar to, but not exactly equal to, the Fourier modes. Due to the relatively smooth profile of the standing wave solutions, it is intuitive that the first few POD modes be similar to the Fourier modes. However, the later POD modes are better described as Fourier modes with an amplitude envelope. This is particularly visible for  $\eta^{(10)}$  but occurs for other POD modes too.

The POD modes of the interior velocity potential,  $\hat{\Phi}$ , are more complicated as shown in Fig. 6.9. The images in the figure show the upper half of a computational domain that is evenly extended to  $y \in [-1, 1]$ . These sixteen modes are computed with the same error threshold of  $E_{\text{tol}} = 10^{-9}$  as the surface modes. The POD modes of the interior do not

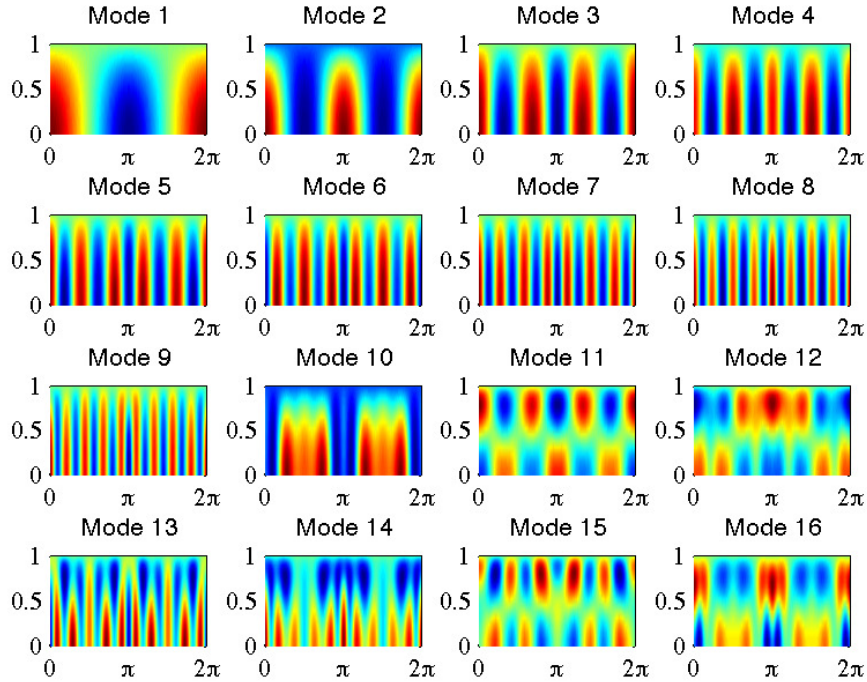


Figure 6.9: Plots of the POD modes of  $\hat{\Phi}(x, y)$  on the transformed rectangular domain. The modes are ordered from left-to-right and from top-to-bottom.

have a natural physical interpretation as they have an indirect dependence on the surface height. However, qualitatively, one can see that like the eigenfunctions  $\cos(kx)\cosh(ky)$  of Laplace's equations, higher POD modes are in general more oscillatory in  $x$  and more rapidly increasing in  $y$ . Demonstrating a quantitative relationship with these eigenfunctions would be difficult because the resulting POD modes (as well as the transformation) depend on the specific combination of  $\varphi$  and  $\eta$  that occurred in the data.

The benefit of using POD modes over Fourier modes is that they better represent the data, in the  $L^2$ -sense, than any other equivalently sized linear set of modes [4]. The limitation of the standard POD technique, however, is that it is applicable to only a single set of data. Using the SPOD method to extend the POD approach to multiple data sets sacrifices that optimality but, in return, yields a more globally applicable basis. In Fig. 6.10 the SPOD method with a tolerance of  $10^{-12}$  was used to generate a set of POD modes using data sets from the indicate points in Fig. 6.8. The data sets were ordered such that data sets with more extreme waveforms were considered later, i.e. from bottom up in Fig. 6.8. For  $\eta$ , the data sets contribute 3, 5, 4, and 2 POD modes respectively. For  $\varphi$ , the data sets contribute 3, 6, 4, and 4 POD modes respectively.

Figure 6.10 shows the convergence of the POD representation to data taken from the




Figure 6.10: Convergence of  $L^2$ -norm of the POD representation (blue line) and the Fourier representation (black line) of the rightmost data point (denoted by a red cross, with  $T \approx 7.2075$ ) in Fig. 6.8. The POD modes were generated using data from the four solutions indicated in that figure. The tolerance set for the SPOD technique here was  $10^{-12}$ . The POD representation is nearly identical to the Fourier representation for the first two modes; as a result, the normalized error is also identical. However, the convergence of the POD representation increases for later modes which are drawn from data sets more similar to the one of interest.

most extreme data set (with  $T \approx 7.2075$ ). For the first two modes, the convergence of the Fourier and POD methods are essentially identical. Indeed, this is because the POD modes are close approximations to the Fourier modes. In some part, this is due to the relatively sinusoidal shape of the lowest-amplitude solution. The later POD modes differ from the Fourier modes and, as a result, reduce the observed error more quickly. This occurs because the later modes are drawn from data sets closer on the branch of solutions than the first data set, and are therefore more applicable to the specific dynamics we want to represent. Overall, despite the loss of provable optimality and the smoother solutions in the finite depth case, the POD modes generated by the SPOD method are still significantly more accurate (one to four orders of magnitude) than the equivalently sized Fourier basis. As a result, there is still the potential for measurable computational gains by switching to a POD basis.

### 6.4.3 POD Dynamics

In this subsection, we consider POD approximations of different dimensions to a single periodic orbit of the PDE. Similar to Sec. 6.3.3, the solution here is a numerically computed periodic solution to the reduced problem with  $F < 10^{-20}$ , as defined in Eq. 6.15. To highlight the differences in the expansions, we present results from the orbit of largest period in Fig. 6.8, but it should be noted that these results are prototypical of other orbits on the branch of solutions.

Figure 6.11 shows the POD reproduction of the solution with largest period in Fig. 6.8 using both a six-mode expansion and a ten-mode expansion. For both these solutions, the period of oscillation is identical (as  $T$  is the bifurcation parameter), but the orbits themselves




Figure 6.11: Plot of the evolution of the fluid surface using (a) 6-POD modes (b) 10-POD modes and (c) 512 Fourier modes. With the 6-mode expansion, the surface evolution can be qualitatively reproduced although there are significant quantitative differences. However, with 10-POD modes there is good quantitative agreement between the low dimensional model and the full 512-Fourier mode PDE solution

are not. With six modes, the POD reduction qualitatively captures the dynamics observed in the PDE. However, without additional modes, the maximum crest height noticeably exceeds that of the full PDE, and the minimum trough depth is smaller than in the PDE as well. In terms of  $L^2$ -norm, the magnitude of the oscillation is too large. On the other hand, the periodic orbit obtained using a ten mode expansion is indistinguishable from the PDE to the eye, i.e. the amplitude of oscillation, in terms of  $L^2$ -norm and the solution at any point on the orbit is visually identical to the same plots from the PDE. There are still quantitative differences between the PDE and POD, but as in standard spectral methods, these can only be corrected by the inclusion of additional modes. Figure 6.12 shows the evolution of the mode-amplitudes of the solution in Fig. 6.11 over one period for both the surface height and surface velocity potential. Note that despite the complex surface shape over a period, the mode amplitudes oscillate in an almost sinusoidal manner, and higher order modes tend to oscillate more rapidly in time, but are also of smaller magnitude.

From Fig. 6.12 it is tempting to conclude that the three easily visible modes,  $a_1$ ,  $a_2$ , and  $a_3$  as well as  $b_1$ ,  $b_2$ , and  $b_3$ , are all the modes that are necessary to completely reproduce the dynamics. Indeed, for a single solution such as the one in Fig. 6.12, a three mode expansion is sufficient to qualitatively reproduce the dynamics observed. However, although the modes with small amplitudes play a relatively minor role in the reconstruction of a single solution, they are critical in reproducing branches of solutions.

#### 6.4.4 Bifurcation Diagram

Individual solutions can be reproduced readily using only a handful of POD modes. However, one of the main goals of this work is to demonstrate that the entire branch of solutions, not just a single element of that set, is low dimensional in nature. In this section, we reproduce

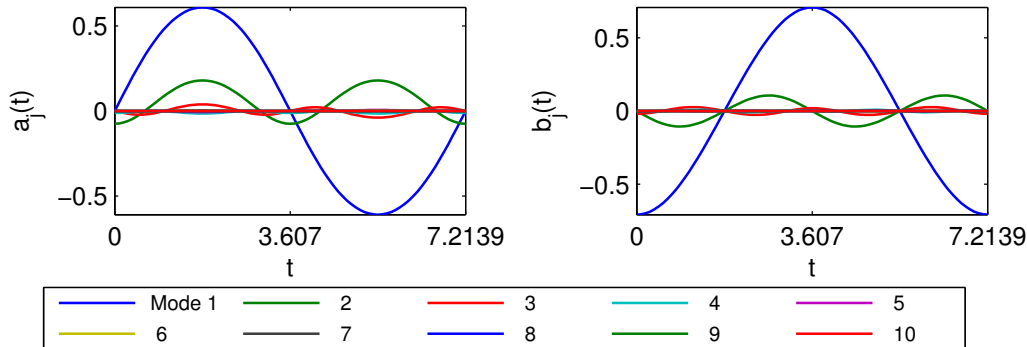


Figure 6.12: Plot of the POD mode amplitudes for the evolution in Fig. 6.11, which demonstrates that the maximum value of the mode amplitude decreases as the mode number and frequency of oscillation increase.

the entire branch of time periodic standing wave solutions using either the first six or all ten of the POD modes in Fig. 6.8 for the surface variables, and the first ten or all sixteen of the interior modes in Fig. 6.9.

Figure 6.13 shows the bifurcation diagram comparing the bifurcation parameter,  $T$ , to the maximum  $L^2$ -norm of the surface height over one period for three cases: the full PDE solution, the six-mode expansion, and the ten-mode expansion. With six POD modes, the bifurcation diagram for the surface height and for the velocity potential can be reproduced in a qualitative fashion. There is a fold in the branch of solutions at approximately the correct period, and as shown in the previous section, the individual solutions are qualitatively similar to the complete PDE solution. On the other hand with ten POD modes, the reduced model is quantitatively accurate, and the two branches of solutions are indistinguishable to the eye.

Furthermore, the POD model accurately predicts the shape of the solution in regions where it did not have access to PDE data. That is, in this problem, the POD model accurately extrapolates beyond the data it could access. This result should not be overstated, and is due primarily to the fact the branch of solutions here is relatively simple with no bifurcations or branch points to complicate the dynamics. Nonetheless, we emphasize that as long as the linear subspace (in phase space) where the dynamics reside does not change too quickly, POD models can be used to accurately predict unobserved results. This implies that the results of this manuscript could have been obtained iteratively — i.e. generating POD modes from a limited set of data, extrapolating beyond that data, solving the PDE based on that extrapolated guess to obtain more data, and repeating — which would eliminate the need for the *a priori* access to complete data relied upon in this manuscript.

Overall, we have reproduced the branch of time-periodic standing wave solutions with unit depth using a six- and ten-mode expansion for both the surface height and velocity potential. Furthermore, a ten and sixteen mode expansion was used for the interior velocity potentials. With the smaller number of modes, the solution branches were obtained qualitatively. With the larger number of modes, the same branch was obtained with quantitative




Figure 6.13: Bifurcation diagrams of the time-periodic standing wave solutions in terms of the period,  $T$ , and the maximum  $L^2$ -norm over one period. The PDE solution branch is shown in solid black. The 6-mode POD expansion is shown in dashed red lines, and the 10-mode POD expansion is shown in dashed green lines.

accuracy. Although the higher order modes play a relatively minor role in reproducing any individual solution, they are critical in reproducing branches of solutions. Hence, although we can use fewer than six modes for individual solutions, six modes appears to be the minimum required to use the same set of modes throughout.

## 6.5 Chapter Summary

In this chapter, we have demonstrated that the branches of time-periodic standing wave solutions in both infinite and finite depth are essentially low dimensional objects by computing a low dimensional subspace from the full PDE dynamics and demonstrating that the dynamics of the system restricted to this subspace are qualitatively similar to the dynamics of the full PDE. For a particular branch of solutions, the sequential POD technique was used to generate a set of basis functions for the subspace that contains the trajectories of the standing wave solutions. Then the Galerkin projection was used to generate a set of evolution equations that have dynamics restricted to this subspace.

The branch of time-periodic solutions was obtained in this reduced model and good agreement with the PDE was observed when a sufficient number of modes were used. For a fluid of infinite depth, a 10-mode expansion in surface height and dipole density was found to be sufficient to reproduce the dynamics on the branch of time-periodic solutions. However to more accurately reproduce the wave crest, a 22-mode expansion was required. For the finite depth case, only 10-surface modes and 10-velocity potential modes were required. However, as few as six of these modes are all that is needed for a qualitatively accurate comparison. In both cases, this is far fewer modes than were required in the full PDE simulation, which demonstrates that the dynamics of time-periodic waves are essentially low dimensional in nature.

The POD technique is general and can readily be applied to different formulations of the water wave problem and with different norms. In this manuscript we use the  $L^2$ -norm and applied the technique to two particular formulations of the problem: the boundary integral method for infinite depth, and a method that tracks  $\Phi$  in the fluid interior for

finite depth. In both cases the data are filtered to produce POD modes that are either symmetric or anti-symmetric with respect to translations in space by  $\pi$ . This allows us to exploit symmetries in the wave water problem to reduce the computational time required by a factor of four. In these two approaches, accurately and quickly computing the Dirichlet to Neumann operator (DNO) is key to generating an efficient low dimensional model — many of the considerations about what variables to represent in a low dimensional framework stem from this issue. Other approaches such as approximating the DNO via Taylor series [149] or using a conformal mapping [153] could readily be adapted to the POD framework and may have computational advantages over what was done here. Furthermore, POD modes obtained with respect to an  $H^1$ - or  $H^2$ -norm may be better suited to reproduce sharp wave crests. This may allow for equally accurate but lower-dimensional solutions in those regimes, as has been observed in other problems in fluid dynamics [159, 160]. Computing the POD modes in the  $H^1$ - or  $H^2$ -norms in this setting is relatively straightforward but has not been undertaken here.

The technique presented in this paper offers an alternative to the approach taken in [89, 147]. In those works, highly optimized and complex scientific codes were used to rapidly compute the same branches of solutions presented in this manuscript with extreme accuracy. In the infinite depth case, the number of unknown Fourier mode coefficients was changed adaptively along the branch, i.e. smooth solutions had fewer degrees of freedom than sharply crested solutions. As a result, the average time required with an 8-core workstation to compute a solution on the branch varies from 2.44 seconds per recorded solution when 20 unknowns were used, to 339 seconds per recorded solution for the sharply crested solutions with 96 unknowns. Meanwhile, the total number of Fourier modes grew from 128 to 1024. The number of time-steps taken by a 5th order Runge-Kutta method was correspondingly increased to maintain accuracy. Overall, the total time required to compute the branch was 5835 seconds on a workstation. We remark that the underlying hardware greatly impacts the execution time. On a laptop, running on a single thread, this execution time increases to 27586 seconds; using GPU hardware on a rackmount server, it improves to 3550 seconds. The benefit of switching from CPUs to GPUs becomes significant for more extreme waves, where several thousand Fourier modes are needed [147].

The POD method shows similar timings using 22 POD modes throughout. Due to the global nature of the modes and the presence of high-wavenumbers throughout, 512 Fourier collocation points were maintained at all points on the branch of solutions. As a result, with a 12-core workstation, the smoother solutions that were quickly computed with the Fourier method took *longer* with the POD method, at an average of 44.46 seconds per recorded solution (compared to 2.44 seconds). For the most sharply crested solutions, on the other hand, the average time required by the POD method actually decreased to 33.07 seconds per recorded solution (compared to 339 seconds). This decrease is primarily due to the closer spacing, in terms of the period of oscillation, of the recorded points for the more sharply crested solutions. As the period was the bifurcation parameter, this allows for a better initial guess for the solution and consequently fewer iterations per solution. Considering the branch as a whole, the POD method required 4496 seconds in total, which is slightly faster than the full model on similar hardware. Further code optimizations and enhancements to the method, such as also adaptively changing the number of POD modes or even simply

neglecting the low amplitude portions of the branch, would reduce this computational time further.

We have also demonstrated that a POD-based approach can be used to produce a viable reduced model even in a non-trivial Hamiltonian system. Because of the Hamiltonian nature of the system, the periodic orbits are not part of limit cycles and have no basins of attraction. Full PDE methods, such as that of Wilkening [89], remain necessary to collect data for the POD method. For small systems, the overhead of the POD method outweighs the benefits. However, the reverse is true for large systems that require many Fourier modes for an accurate representation. In this case, the reduced number of degrees of freedom allows a branch of solutions to be extended quickly. Eventually, reduced order models will lose accuracy, but as demonstrated in this manuscript, extrapolation over non-negligible distances in parameter space is possible. As a result, there is a strong potential for the development of hybrid methods for numerical continuation, similar to the hybrid integrator discussed in Chap. 5. As an example, consider the standard predictor-corrector methods in numerical continuation. The computationally inexpensive extrapolatory power of the POD-method could serve as an almost ideal predictor, allowing larger steps in parameter space to be taken before the more expensive (but potentially heavily optimized) corrector step is called. Although this approach is not required for the standing waves in two-dimensional fluids demonstrated here, the method could be generalized for standing waves in three dimensions, symmetry breaking bifurcations, and stability transitions in Faraday wave patterns to reduce the time required to track solution branches.

In addition to their use as a tool for aiding in the continuation of PDEs, reduced models can also be used to reduce the cost of parameter space exploration. Ultimately it is hoped that the use of reduced models can help answer practical questions fundamental to the understanding of tsunamis, rogue waves, and high-amplitude ocean waves. Given the demonstrated computational complexity of obtaining large amplitude solutions [89, 134, 135, 147], the POD technique offers an intuitive and tractable methodology for characterizing fundamental behaviors, such as tracking bifurcation curves and identifying stability properties, of seemingly high dimensional systems such as the water wave problem considered here.

## Chapter 7

## CONCLUDING REMARKS

In this thesis, we have examined the application of dimensionality reduction techniques to reduce the computational cost required to solve a number of problems in applied mathematics. In nonlinear optics, dimensionality reduction initially took a parametric approach, i.e. a solution ansatz with unknown and time-dependent parameter values was required. These methods included the variational reduction [33, 161] and coupled mode theory (CMT) [62], and the historical impact of these methods alone is sufficient to justify their importance. In particular, CMT is an integral part of the derivation of a number of important inequations in nonlinear optics such as the discrete (nonlinear) Schrödinger equation for the propagation of the optical field in a (nonlinear) fiber or waveguide array, and as such is included in nearly all photonics textbooks. Chapter 3 demonstrates the application of such a technique to a physically valid model for the intensity dynamics in a semiconductor waveguide array. From previous intuition and full PDE simulations we demonstrated that CMT could accurately represent the dynamics of the full system. Then in the reduced dimensional system, we were able to determine the route to chaos observed in that array.

The shortcomings of ansatz based methods appear when a valid ansatz cannot be obtained, either because one does not exist or because the appropriate function has been determined. In Chapter 4, the multi-pulsing transition of a waveguide array laser was such a case. Due to the transitions in the system — from stationary solutions to breather solutions, breather solutions to chaotic solutions, and finally from chaotic solutions to two-pulse solutions — a globally valid ansatz is as yet unknown. By applying data-analysis methods, like the POD, to time-series data from PDE itself, we can allow the dynamics to determine what the optimal basis should be. In this reduced basis, the standard tools for numerical continuation such as MATCONT [1] or AUTO [2] once again becomes viable in practice even when the systems of interest are nonlinear PDEs.

Certainly, there are methods for obtaining the solution branches and bifurcations of nonlinear PDEs directly. Two common black-box techniques are the recursive projection method [85] and Newton-Picard methods [74] which can be applied to dissipative systems with a few unstable eigenvalues. Methods like the Adjoint Continuation Method (ACM), used in Chapter 4, are more generally applicable but require the explicit computation of the adjoint. Applying the ACM to the WGAML model the branch of time periodic solutions involved in the multi-pulsing transition was computed to 10 digits of accuracy. As a result, we finally have the ability to compare the results of the reduced order model to “true” results. This revealed that the reduced order model does qualitatively predict the bifurcation sequence.

The POD and other related methods have long been used in an “offline” fashion. That is, data for the modes is collected first and the reduced model is used from there on out. Indeed, this is the paradigm used in Chapter 4. The static nature of the modes, however,

implies that they cannot adapt to the changes that could occur if, say, a bifurcation occurred. In Chapter 5, we developed a hybrid method that adaptively switches between a full PDE integrator for data collection and accuracy and a reduced order model for speed. We demonstrate the application of this hybrid integrator to two problems from mode-locking: the WGAML model and the Cubic Quintic Ginzburg Landau equation (CQGLE). In both cases, the hybrid integrator was able to make use of the underlying low dimensional nature of the system to reduce the computational effort required for numerical integration. Due to the hybrid nature of the system, it also reverted to the full PDE when more data was needed or, in cases like the chaotic solutions in the PDE WGAML model, no valid low dimensional model existed. Overall although these sorts of hybrid methods sacrifice some of the speed benefits of POD based methods, their added robustness and online nature makes them more widely applicable in practice.

The dimensionality reduction methods discussed in this thesis have primarily been applied to problems in nonlinear optics. This is due, in part, to the presence of mode-locked solutions which are both physically interesting and low dimensional. However, this does not mean that methods like the POD are application specific. Indeed, the POD method, as it is used today, grew out of the study of turbulence and fluid dynamics [4]. It can even be applied to Hamiltonian systems such as the water wave problem as was done in Chapter 6. In the context of generating a low dimensional basis, the presence of an underlying Hamiltonian structure is actually a hindrance as it precludes the existence of the stable stationary solutions and limit-cycles with broad basins of attraction that were so convenient for collecting data. Once the data is collected however, the POD method can be applied just as in the previous cases. Ultimately, we were able to recover the branch of time-periodic standing wave solutions in water of both infinite and finite depth.

Overall, the use of dimensionality reduction techniques can significantly reduce the computational cost of computing bifurcation diagrams, evolving systems, and a number of other related tasks in applied mathematics. Although the restriction of low-dimensionality may seem like a severe one, many systems of interest, particularly those that show the formation of clear patterns, have an underlying low dimensional structure. If an appropriate solution ansatz can be obtained, ansatz based methods are both well studied and powerful. If an appropriate solution ansatz cannot be obtained, then data-driven methods, like the proper orthogonal decomposition, can be used in their place.

In this thesis, we have only scratched the surface of the types and applications of data-driven dimensionality reduction. In the machine learning and image processing communities, linear methods such as the POD have long been in competition with other linear methods such as independent component analysis (ICA) or factor analysis methods and nonlinear methods such as Sammon's mapping, Principal Manifolds, Diffusion Maps, Locally Linear Embedding (LLE), and many others [3]. Despite their widespread use in machine learning and statistics, these methods have only begun to be applied to nonlinear PDEs [113]. As a result, there are still many opportunities to further expand on the ideas presented here and harness the existing work of the machine learning community.

## Appendix A

**ADJOINT CONTINUATION METHOD**

The adjoint continuation method (ACM) is a numerical method developed originally by Wilkening [89] to compute branches of time-periodic solutions for PDE systems. The technique is quite general and applicable to a number of physical and biological systems such as standing surface waves and mode-locked lasers. In this appendix, we discuss the ACM as it applies to the multi-pulsing transition of the WGAML in Chapter 4. In Section A.1, we derive the ACM for the WGAML. In Section A.2 we derive the adjoint equation required for the ACM to be efficient. Lastly in Section A.4, we apply a semi-implicit integration scheme to equations like the WGAML model with gain of the form in Eq. (1.6). For a more general treatment of the ACM and its application to a broader class of problems, see the work of Wilkening in Refs. [87, 88].

**A.1 Description of the Adjoint Continuation Method for the WGAML**

In this section, we describe the ACM for tracking families of time-periodic (or in this case,  $Z$ -periodic) solutions by numerical continuation. Starting with mode-locked solutions of the form of Eq. (1.7), which are trivially  $Z$ -periodic (up to a phase), we will use the method to follow these paths through a sequence of bifurcations leading to more and more complicated, ultimately chaotic, dynamics. Tracking chaotic solutions is beyond the capabilities of the ACM, but it has been observed by Kutz and Standstede [31] that the chaotic solutions, upon further increase of the gain, lead to the formation of the double-pulse solutions. This bifurcation sequence is described in Sec. 4.2.

General methods [162–165] for the solution of nonlinear two-point boundary value problems tend to be geared toward ordinary differential equations, and can be prohibitively expensive for partial differential equations. Recently [87, 166, 167], Wilkening and Ambrose introduced an efficient method of computing time-periodic solutions of nonlinear PDEs. We will refer to this method as the Adjoint Continuation Method (ACM). The idea is to develop a variant of a shooting method [168, 169] in which a nonlinear functional of the initial condition is minimized using adjoint-based optimal control methods [170–172] to obtain a solution of the boundary value problem. Key challenges in adapting this method to the system in Eq. (1.5) include finding solutions that are only periodic up to a phase; incorporating the gain,  $g_0$ , in the adjoint system to allow other variables (such as the period) to be used as bifurcation parameters; and adapting high order, semi-implicit Runge-Kutta methods [173, 174] to handle the case when the terms responsible for stiffness (those involving  $\partial^2 A_0 / \partial T^2$  in Eq. (1.5)) depend non-linearly on  $A_0$  through a gain  $g$  that depends on  $\|A_0\|$ .

Recall that  $z$  is the time-like variable in Eq. (1.5) while  $T$  is the space-like variable. To facilitate the use of spectral methods, we adopt periodic boundary conditions over an interval  $T \in [-L, L)$  where  $L = 20$ .  $L$  is chosen large enough that the solution, which

decays exponentially, is of order  $10^{-10}$  near  $T = \pm L$ . Once  $L$  is large enough, the solution is insensitive to further changes in  $L$ , and may be regarded as a solution over  $\mathbb{R}$  without periodic boundary conditions. In what follows, we take  $L = \pi$  and absorb the appropriate factors from the change of variables into  $D$ ,  $\tau$  and  $e_0$  in Eq. (1.5) and Eq. (1.6). However, for the results in Section 4.2, we transform back to the original domain.

Following the basic approach in [87, 166, 167], we define a nonlinear functional  $G$  of the initial conditions and supposed period that is zero if and only if the solution is  $Z$ -periodic. As we wish to determine the dependence of  $G$  on the parameter  $g_0$  in Eq. (1.6), we add to Eq. (1.5) the equation  $\partial g_0 / \partial Z = 0$ . This will be explained in more detail below. We also wish to find solutions that are only  $z$ -periodic up to a phase. An easy way to do this is to include extra terms on the right hand side of Eq. (1.5) that affect the solution by multiplying  $A_0(Z, T)$ ,  $A_1(Z, T)$  and  $A_2(Z, T)$  by  $e^{-i\theta Z}$ ; we then search for  $\theta$  such that the new problem has fully  $z$ -periodic solutions. The new set of equations is

$$i \frac{\partial A_0}{\partial Z} + \frac{D}{2} \frac{\partial^2 A_0}{\partial T^2} + \beta |A_0|^2 A_0 + C A_1 + i \gamma_0 A_0 - \frac{2i g_0}{1 + \|A_0\|^2 / e_0} \left( 1 + \tau \frac{\partial^2}{\partial T^2} \right) A_0 = \theta A_0, \quad (\text{A.1a})$$

$$i \frac{\partial A_1}{\partial Z} + C (A_0 + A_2) + i \gamma_1 A_1 = \theta A_1, \quad (\text{A.1b})$$

$$i \frac{\partial A_2}{\partial Z} + C A_1 + i \gamma_2 A_2 = \theta A_2, \quad (\text{A.1c})$$

$$\frac{\partial g_0}{\partial Z} = 0. \quad (\text{A.1d})$$

Although  $A_0$ ,  $A_1$  and  $A_2$  in Eq. (1.5) are complex valued, we represent the state of the system (with  $Z$  frozen) as an element of the *real* Hilbert space

$$X = \{(A_0, A_1, A_2, g_0) : A_0, A_1, A_2 \in L^2((-\pi, \pi], \mathbb{C}), g_0 \in \mathbb{R}\}$$

with inner product

$$\langle q_1, q_2 \rangle = \int_{-\pi}^{\pi} \Re \left\{ A_{0,1}(T)^* A_{0,2}(T) + A_{1,1}(T)^* A_{1,2}(T) + A_{2,1}(T)^* A_{2,2}(T) \right\} dT + g_{0,1} g_{0,2},$$

where,  $q_i(T) = (A_{0,i}(T), A_{1,i}(T), A_{2,i}(T), g_{0,i}) \in X$ . Next we define

$$G(q_0, Z, \theta) = \frac{1}{2} \|q(Z_0, \cdot) - q_0(\cdot)\|^2, \quad \|q\|^2 = \langle q, q \rangle, \quad (\text{A.2})$$

where  $q(Z, T)$  solves the initial value problem Eq. (A.1) with initial conditions  $q(0, T) = q_0(T)$ . We note that  $G(q_0, Z, \theta) = 0$  if and only if  $q(Z, T)$  is  $Z$ -periodic, with period  $Z_0$ .

To evaluate  $G$  numerically, we evolve  $q(z, t)$  using a spectral collocation method in  $t$  and a 5th order semi-implicit Runge-Kutta method in  $z$ , that is described in detail in Section A.4.

A prescribed fraction of the Fourier modes (around 40%) are allowed to be non-zero in  $q_0(t)$ . The remaining (high frequency) modes of the initial condition are set to zero to avoid aliasing errors in the computation of  $G$ . To minimize  $G$ , we use the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [156] to vary  $Z_0$ ,  $\theta$ ,  $g_0$  and the non-zero Fourier modes of the initial condition. One of these variables (usually  $g_0$  or  $Z_0$ ) is taken as a bifurcation parameter in the continuation algorithm and removed from the list of variables that BFGS is allowed to vary in search of a minimum. Alternatively, a penalty function [87] can be used to enforce the value of the bifurcation parameter. In either approach, it is useful to be able to use any variable as a bifurcation parameter and switch between them as necessary to traverse turning points and avoid ill-conditioned minimization problems.

BFGS is a quasi-Newton gradient descent method that builds up an approximate inverse Hessian matrix based on the sequence of gradient vectors it encounters in the course of the line searches. Thus, to use the algorithm, we must be able to accurately and efficiently compute the gradient,  $\nabla G$ . Some of the components of  $\nabla G$  can be computed immediately, e.g.

$$\frac{\partial G}{\partial Z} = \left\langle \frac{\partial q}{\partial Z}(Z_0, \cdot), q(Z_0, \cdot) - q_0(\cdot) \right\rangle, \quad \frac{\partial G}{\partial \theta} = \left\langle \begin{pmatrix} iZ_0 A_0(Z_0, \cdot) \\ iZ_0 A_1(Z_0, \cdot) \\ iZ_0 A_2(Z_0, \cdot) \\ 0 \end{pmatrix}, q_0(\cdot) \right\rangle.$$

These inner products are evaluated using the trapezoidal rule at the collocation points of the spectral method. In the formula for  $\frac{\partial G}{\partial \theta}$ , we used the fact that

$$\frac{\partial}{\partial A_0 \theta}(Z, T) = -iZ A_0(Z, T), \quad \int_{-\pi}^{\pi} \Re\{(-iZ_0 A_0(Z, T))^* A_0(Z, T)\} dT = 0,$$

with similar formulas for  $A_1$  and  $A_2$ . It remains to determine  $\frac{\partial G}{\partial g_0}$  and

$$\frac{\partial G}{\partial \Re(\hat{u}_k)}, \quad \frac{\partial G}{\partial \Im(\hat{u}_k)}, \quad \frac{\partial G}{\partial \Re(\hat{v}_k)}, \quad \frac{\partial G}{\partial \Im(\hat{v}_k)}, \quad \frac{\partial G}{\partial \Re(\hat{w}_k)}, \quad \frac{\partial G}{\partial \Im(\hat{w}_k)}, \quad (-k_{\max} \leq k \leq k_{\max}), \quad (\text{A.3})$$

where  $k_{\max}$  is the cutoff beyond which Fourier modes of the initial condition are set to zero. These can all be computed simultaneously by solving a single adjoint PDE as described in Sec. A.2.

In summary, the ACM uses adjoint methods to compute  $G$  and  $\nabla G$  in just two PDE solves: the nonlinear evolution equations Eq. (A.1) are solved to compute  $G$  and two components of the gradient,  $\frac{\partial G}{\partial Z_0}$  and  $\frac{\partial G}{\partial \theta}$ ; then the non-autonomous (but linear) adjoint system Eq.(A.8) is solved to obtain the remaining components of the gradient, namely  $\frac{\partial G}{\partial g_0}$  and those listed in Eq. (A.3). These are the ingredients needed to use the BFGS method to minimize  $G$  and obtain a time-periodic solution. We do this repeatedly for different values of the bifurcation parameter to sweep out families of solutions. When performing numerical continuation, the approximate inverse Hessian matrix from the final iteration of the previous BFGS solve is used to initialize the inverse Hessian matrix for the next solve. As long as the continuation steps are not too large, this dramatically reduces the number of line searches

required to minimize  $G$  due to superlinear convergence in the BFGS algorithm when the Hessian is well approximated.

## A.2 The Adjoint PDE of the WGAML

The goal of the adjoint PDE is to find the variational derivative  $\frac{\delta G}{\delta q_0} = \left( \frac{\delta G}{\delta u_0}, \frac{\delta G}{\delta v_0}, \frac{\delta G}{\delta w_0}, \frac{\partial G}{\partial g_0} \right) \in X$ , which satisfies

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} G(q_0 + \varepsilon \dot{q}_0, Z_0, \theta) = \left\langle \frac{\delta G}{\delta q_0}, \dot{q}_0 \right\rangle \quad (\text{A.4})$$

for every sufficiently smooth  $\dot{q}_0 \in X$ . Here a dot represents a directional derivative with respect to the initial conditions, and we will write  $\dot{G}$  for the left hand side of Eq. (A.4). The quantities in Eq. (A.3) are the real and imaginary parts of the Fourier coefficients of  $\frac{\delta G}{\delta u_0}$ ,  $\frac{\delta G}{\delta v_0}$  and  $\frac{\delta G}{\delta w_0}$ , e.g.

$$\frac{\partial G}{\partial \Re(\hat{u}_k)} = \left\langle \frac{\delta G}{\delta q_0}, \begin{pmatrix} e^{ikT} \\ 0 \\ 0 \\ 0 \end{pmatrix} \right\rangle = 2\pi \Re \left\{ \left( \widehat{\frac{\delta G}{\delta A_{0,0}}} \right)_k \right\}, \quad \frac{\partial G}{\partial \Im(\hat{u}_k)} = \left\langle \frac{\delta G}{\delta q_0}, \begin{pmatrix} ie^{ikT} \\ 0 \\ 0 \\ 0 \end{pmatrix} \right\rangle = 2\pi \Im \left\{ \left( \widehat{\frac{\delta G}{\delta u_0}} \right)_k \right\}.$$

To find a formula for  $\frac{\delta G}{\delta q_0}$ , we evaluate the left hand side of Eq. (A.4) and then manipulate the result to obtain the form on the right of Eq. (A.4). On the left, we have

$$\dot{G} := \left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} G(q_0 + \varepsilon \dot{q}_0, Z_0, \theta) = \left\langle q(Z_0, \cdot) - q_0(\cdot), \dot{q}(Z_0, \cdot) - \dot{q}_0(\cdot) \right\rangle, \quad (\text{A.5})$$

where  $\dot{q}(z, t)$  solves the linearization of Eq. (A.1) about  $q(Z, T)$  with initial condition  $\dot{q}(0, T) = \dot{q}_0(T)$ :

$$\begin{aligned} & i \frac{\partial \dot{A}_0}{\partial z} + \frac{D}{2} \frac{\partial^2 \dot{A}_0}{\partial t^2} + 2\beta |u|^2 \dot{A}_0 + \beta A_0^2 \dot{A}_0^* + C \dot{A}_1 + i\gamma_0 \dot{A}_0 - \frac{2ig_0}{1 + \|u\|^2/e_0} \left( 1 + \tau \frac{\partial^2}{\partial t^2} \right) \dot{A}_0 \\ & + \left[ \frac{4ig_0 \langle A_0, \dot{A}_0 \rangle / e_0}{(1 + \|A_0\|^2/e_0)^2} - \frac{2i\dot{g}_0}{1 + \|A_0\|^2/e_0} \right] \left[ \left( 1 + \tau \frac{\partial^2}{\partial t^2} \right) u \right] = \theta \dot{A}_0, \end{aligned} \quad (\text{A.6a})$$

$$i \frac{\partial \dot{A}_1}{\partial z} + C(\dot{A}_2 + \dot{A}_0) + i\gamma_1 \dot{A}_1 = \theta \dot{A}_1, \quad (\text{A.6b})$$

$$i \frac{\partial \dot{A}_2}{\partial z} + C \dot{A}_1 + i\gamma_2 \dot{A}_2 = \theta \dot{A}_2, \quad (\text{A.6c})$$

$$\frac{\partial \dot{g}_0}{\partial z} = 0. \quad (\text{A.6d})$$

Here  $\langle A_0, \dot{A}_0 \rangle = \left( \int_{-\pi}^{\pi} \Re \{ A_0(z, t)^* \dot{A}_0(z, t) \} dt \right)$  is a real number, and depends on  $z$ . Next we define adjoint variables  $\tilde{q}(\zeta, t) = (\tilde{A}_0(\zeta, t), \tilde{A}_1(\zeta, t), \tilde{A}_2(\zeta, t), \tilde{g}_0)$ , where  $\zeta = Z - z$  is a

reversed time-like variable. We define  $\tilde{q}_0(t) = q(Z, t) - q_0(t)$  so that Eq. (A.5) becomes

$$\dot{G} = \left\langle \tilde{q}_0(\cdot), \dot{q}(Z, \cdot) \right\rangle - \left\langle \tilde{q}_0, \dot{q}_0 \right\rangle.$$

This can be put in the desired form Eq. (A.4) with

$$\frac{\delta G}{\delta q_0}(t) = \tilde{q}(Z, t) - \tilde{q}_0(t) \quad (\text{A.7})$$

provided that  $\left\langle \tilde{q}(Z - z, \cdot), \dot{q}(z, \cdot) \right\rangle$  remains constant for  $0 \leq z \leq Z$ . Differentiation with respect to  $z$  shows that this condition will hold provided that  $\tilde{q}(\zeta, t)$  satisfies

$$\left\langle \frac{\partial \tilde{q}}{\partial \zeta}(\zeta, \cdot), \dot{q}(z, \cdot) \right\rangle = \left\langle \tilde{q}(\zeta, \cdot), \frac{\partial \dot{q}}{\partial z}(z, \cdot) \right\rangle, \quad \tilde{q}(0, t) = \tilde{q}_0(t), \quad \zeta = Z - z.$$

Using Eq. (A.6) to evaluate  $\left\langle \tilde{q}(\zeta, \cdot), \frac{\partial \dot{q}}{\partial z}(z, \cdot) \right\rangle$ , we integrate by parts and collect terms to identify the adjoint system

$$\begin{aligned} \frac{\partial \tilde{A}_0}{\partial \zeta} = & -i \frac{D}{2} \frac{\partial^2 \tilde{A}_0}{\partial t^2} - 2i\beta |u|^2 \tilde{A}_0 + i\beta u^2 \tilde{A}_0^* - \gamma_0 \tilde{A}_0, \\ & + \frac{2g_0}{1 + \|u\|^2/e_0} \left(1 + \tau \frac{\partial^2}{\partial t^2}\right) \tilde{A}_0 - \frac{4g_0 \langle \tilde{A}_0, (1 + \tau \partial_t^2)u \rangle}{e_0(1 + \|u\|^2/e_0)^2} u + i\theta \tilde{A}_0 - iC \tilde{A}_1, \end{aligned} \quad (\text{A.8a})$$

$$\frac{\partial \tilde{A}_1}{\partial \zeta} = -iC(\tilde{A}_0 + \tilde{A}_2) - \gamma_1 \tilde{A}_1 + i\theta \tilde{A}_1, \quad (\text{A.8b})$$

$$\frac{\partial \tilde{A}_2}{\partial \zeta} = -iC \tilde{A}_1 - \gamma_2 \tilde{A}_2 + i\theta \tilde{A}_2, \quad (\text{A.8c})$$

$$\frac{\partial \tilde{g}_0}{\partial \zeta} = \frac{2 \langle \tilde{A}_0, (1 + \tau \partial_t^2)u \rangle}{1 + \|u\|^2/e_0}, \quad \langle \tilde{A}_0, (1 + \tau \partial_t^2)u \rangle = \int_{-\pi}^{\pi} \Re \left\{ \tilde{A}_0^* u - \tau \frac{\partial \tilde{A}_0}{\partial t} \frac{\partial u}{\partial t} \right\} dt. \quad (\text{A.8d})$$

In this equation, as  $\zeta$  runs from 0 to  $Z$ , the variables  $A_0$ ,  $A_1$  and  $A_2$  are evaluated at  $(Z - \zeta, t)$  while  $\tilde{A}_0$ ,  $\tilde{A}_1$  and  $\tilde{A}_2$  are evaluated at  $(\zeta, t)$ . Then using Eq. (A.4) and (A.7) in combination with Eq. (A.8) the remaining components of the gradient can be computed.

### A.3 Computing the stability of periodic solutions

In order to compute the Floquet multipliers of this system, the monodromy matrix [90] was computed for the linearized system in Eq. (A.6) with  $\dot{g}_0 = 0$  and the ODE for  $\dot{g}_0$  dropped. This term is zeroed out because for any given periodic solution  $g_0$  is fixed. The complex solutions of Eq. (A.6) were rewritten as a real system of twice the size by splitting the solutions into real and imaginary parts. If the matrix  $X(z)$  is the fundamental solution

matrix of this system at  $z$ , then the monodromy matrix is

$$M = X(Z_0)X^{-1}(0). \quad (\text{A.9})$$

Because the linearized operator in Eq. (A.6) is non-autonomous,  $X(Z_0)$  must be computed explicitly by solving Eq. (A.6) for the set of initial conditions implicitly described by  $X(0)$ . For the multipliers computed above,  $X(0)$  was represented in the Fourier basis and  $\hat{X}(0) = \text{id}$  where the hat represents the discrete Fourier transform acting on the columns of  $X(0)$ . In order to compute  $A_0(Z, T)$  at intermediate timesteps, the same Hermite interpolation and semi-implicit time-stepper used to solve the adjoint equation were also used to compute the monodromy matrix. Due to the large number of degrees of freedom in the system, this is an expensive computation. However, the majority of the time is spent computing  $X(Z_0)$ , and this portion of the computation is easily parallelizable. With the monodromy matrix assembled, MATLAB was used to compute the eigenvalues.

The linearized equations in Eq. (A.6) has three multipliers where one can show  $\mu_1 = 1$ , one from the phase-condition, one from phase-invariance, and one from translational invariance in  $t$ . This approach recovers these multipliers with  $|\mu_1 - 1| < 10^{-4}$ . Although this approach is too computationally intensive to be used on all solutions, it was used to compute the stability of solutions where a bifurcation was suspected to occur.

#### A.4 Semi-implicit Runge-Kutta Method

The terms involving second derivatives with respect to the space-like variable,  $t$ , cause the systems Eq. (1.5) and (A.8) to be stiff. This means stability requires smaller steps in an explicit scheme than is necessary for accuracy. On the other hand, a fully implicit method would be difficult to implement due to the fact that Eq. (1.5) is nonlinear and (A.8), while linear, is not diagonalized by the Fourier transform. The idea of an implicit-explicit (IMEX) multi-step method [175], or an additive Runge-Kutta (ARK) method [173, 174], is to write the evolution equation as a sum

$$\frac{\partial q}{\partial z} = f(z, q) + g(z, q) \quad (\text{A.10})$$

and devise a scheme in which  $f$  is treated explicitly,  $g$  is treated implicitly, the scheme for  $g$  alone is e.g.  $L$ -stable, and the combined scheme is high order. In the ARK framework, there are two sets of stage derivatives and two Butcher arrays [176], one for  $f$  and one for  $g$ :

$$\begin{aligned} k_i &= f\left(z_n + c_i h, q_n + h \sum_j a_{ij} k_j + h \sum_j \hat{a}_{ij} \ell_j\right), & \frac{c}{b^T} & \Big| \frac{\mathbf{A}}{\hat{\mathbf{A}}} \\ \ell_i &= g\left(z_n + \hat{c}_i h, q_n + h \sum_j a_{ij} k_j + h \sum_j \hat{a}_{ij} \ell_j\right), & \Big| & \Big| \frac{\hat{\mathbf{b}}^T}{\hat{\mathbf{b}}^T} \\ q_{n+1} &= q_n + h \sum_j b_j k_j + h \sum_j \hat{b}_j \ell_j. & \text{for } f & \quad \text{for } g \end{aligned} \quad (\text{A.11})$$

The Butcher arrays satisfy  $a_{ij} = 0$  if  $i \leq j$  and  $\hat{a}_{ij} = 0$  if  $i < j$ . This allows the stage derivatives to be solved for in order,  $l_1, k_1, \dots, l_s, k_s$ , where  $s$  is the number of stages

of the scheme. We used a variant [177] of the 5th order ARK scheme of Kennedy and Carpenter [174], modified so that a fourth order dense output exists. A dense output [176] is a rule for accurately interpolating the solution between timesteps. In our case, we must interpolate the solution of the forward problem Eq. (1.5) to solve the adjoint problem. To achieve fifth order accuracy in the adjoint system, we require a fourth order dense output formula. We use the dense output to generate quartic corrections [178] to cubic Hermite interpolation. Details are given in [177].

For the nonlinear system Eq. (A.1), we set  $q = (A_0, A_1, A_2)$  in Eq. (A.10) and define

$$f(q) = \begin{pmatrix} -i\theta + i\beta|A_0|^2 - \gamma_0 + \frac{2g_0}{1 + \|A_0\|^2/e_0} & iC & 0 \\ iC & -i\theta - \gamma_1 & iC \\ 0 & iC & -i\theta - \gamma_2 \end{pmatrix} \begin{pmatrix} A_0 \\ A_1 \\ A_2 \end{pmatrix}, \quad (\text{A.12})$$

$$g(q) = \begin{pmatrix} \left( i\frac{D}{2} + \frac{2g_0\tau}{1 + \|A_0\|^2/e_0} \right) \frac{\partial^2}{\partial t^2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} A_0 \\ A_1 \\ A_2 \end{pmatrix}.$$

Note that there is no need to include Eq. (1.6) in the evolution equations once the adjoint system is derived;  $g_0$  is simply treated as a parameter in Eq. (A.10) and Eq. (A.12). The main challenge in this procedure is that  $g(q)$  in Eq. (A.12) is non-linear due to the term  $\|A_0\|^2$ . The implicit equation that must be solved has the form

$$(U, V, W) = g[(A_0, A_1, A_2) + \varepsilon(U, V, W)], \quad (\text{A.13})$$

where  $(U, V, W) = \ell_i$  in Eq. (A.11),  $\varepsilon = \hat{a}_{ii}h$ , and  $(A_0, A_1, A_2) = q_n + h \sum_{j=1}^{i-1} [a_{ij}k_j + \hat{a}_{ij}l_j]$ . Due to the form of  $g(q)$  in Eq. (A.12), we have  $V = W = 0$ . We solve for  $U$  in Fourier space:

$$\hat{U}_k + k^2 \left( i\frac{D}{2} + \frac{2g_0\tau}{1 + (2\pi/e_0) \sum_j |\hat{A}_{0,j} + \varepsilon \hat{U}_j|^2} \right) (\hat{A}_{0,k} + \varepsilon \hat{U}_k) = 0, \quad (k \in \mathcal{K}). \quad (\text{A.14})$$

Here  $\mathcal{K} = \{-N/2 + 1, \dots, N/2 - 1\}$ , where  $N$  is the number of Fourier collocation points used to discretize  $(-\pi, \pi]$ , and we set the Nyquist mode to zero,  $\hat{U}_{N/2} = 0$ . In addition, at the end of each time-step, we apply a 36th order filter [179] in which the Fourier modes are multiplied by  $e^{-36(2|k|/N)^{36}}$ . This filter has proved successful in suppressing aliasing errors while allowing high frequency modes to contribute useful information to the solution.

It is convenient to solve for the real and imaginary parts of  $\hat{U}_k$  in Eq. (A.14) separately. Setting

$$\hat{U}_k = a_k + ib_k, \quad \hat{A}_{0,k} = \alpha_k + i\beta_k,$$

the system Eq. (A.14) becomes

$$F_k = \begin{pmatrix} a_k \\ b_k \end{pmatrix} + k^2 \left( \frac{D}{2} J + E(r)I \right) \begin{pmatrix} \alpha_k + \varepsilon a_k \\ \beta_k + \varepsilon b_k \end{pmatrix} = 0, \quad (k \in \mathcal{K}), \quad (\text{A.15})$$

where

$$E(r) = \frac{2g_0\tau}{1 + (2\pi/e_0)r}, \quad r = \sum_j [(\alpha_j + \varepsilon a_j)^2 + (\beta_j + \varepsilon b_j)^2], \quad I = \begin{pmatrix} 1 & \\ & 1 \end{pmatrix}, \quad J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

We solve Eq. (A.15) iteratively, using Newton's method

$$\hat{U}^{(\nu+1)} = \hat{U}^{(\nu)} - [\nabla F(\hat{U}^{(\nu)})]^{-1} F(\hat{U}^{(\nu)}),$$

where  $\hat{U} = (\dots, a_{-1}, b_{-1}, a_0, b_0, a_1, b_1, \dots) \in \mathbb{R}^{2N-2}$ . As a starting guess,  $\hat{U}^{(0)}$ , we drop  $\varepsilon a_j$  and  $\varepsilon b_j$  in the formula for  $r$  and solve Eq. (A.15), which becomes a de-coupled set of  $2 \times 2$  linear systems when  $r$  is frozen in this way. On subsequent iterations of Newton's method, we note that the Jacobian is a rank one perturbation of a  $2 \times 2$  block-diagonal matrix:

$$(\nabla F)_{kj} = \frac{\partial F_k}{\partial (a_j, b_j)} = A_{kj} + B_{kj}, \quad \begin{aligned} \mathbf{A}_{kj} &= \{I + \varepsilon k^2 [(D/2)J + E(r)I]\} \delta_{kj}, \\ \mathbf{B}_{kj} &= 2\varepsilon k^2 E'(r) \begin{pmatrix} \alpha_k + \varepsilon a_k \\ \beta_k + \varepsilon b_k \end{pmatrix} (\alpha_j + \varepsilon a_j, \beta_j + \varepsilon b_j). \end{aligned}$$

Thus, we may use the Sherman-Morrison formula [156] to invert  $\nabla F$ :

$$(\mathbf{A} + uv^T)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1}uv^T\mathbf{A}^{-1}}{1 + v^T\mathbf{A}^{-1}u}.$$

To invert the  $2 \times 2$  diagonal blocks of  $\mathbf{A}$ , we use  $(xI + yJ)^{-1} = \frac{1}{x^2 + y^2}(xI - yJ)$ . Because  $\varepsilon$  is small in practice, the starting guess is close enough that Newton's method converges to machine precision in 3-5 iterations.

For the adjoint system Eq. (A.8), we use the same fifth order additive Runge-Kutta scheme as in the forward problem in Eq. (A.1), but now the equations are non-autonomous and linear

$$\frac{\partial \tilde{q}}{\partial \zeta} = f(\zeta, \tilde{q}) + g(\zeta, \tilde{q}), \quad \tilde{q} = (\tilde{A}_0, \tilde{A}_1, \tilde{A}_2, \tilde{g}_0).$$

We set

$$g(\zeta, \tilde{q}) = \begin{pmatrix} \left( -i\frac{D}{2} + \frac{2g_0\tau}{1 + \|u(Z - \zeta, \cdot)\|^2/e_0} \right) \frac{\partial^2 \tilde{A}_0}{\partial t^2} \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

and define  $f(\zeta, \tilde{q})$  to be the remaining terms of Eq. (A.8). Although  $f(\zeta, \tilde{q})$  and  $g(\zeta, \tilde{q})$  are both linear functions of  $\tilde{q}$ ,  $f(\zeta, \tilde{q})$  contains terms that are difficult to invert in a fully implicit scheme while  $g(\zeta, \tilde{q})$  is easily inverted using the FFT in a similar way to the forward problem above, but without the need for Newton's method. As mentioned at the beginning of this section, we use a dense output formula described in [177] to interpolate the solution  $q(z, t)$  between timesteps to evaluate quantities such as  $\|A_0(Z - \zeta, \cdot)\|$  in the adjoint problem.

## BIBLIOGRAPHY

- [1] A. Dhooge, W. Govaerts, and Y. A. Kuznetsov, “Matcont: A Matlab package for numerical bifurcation analysis of ODEs,” *Association for Computing Machinery Transactions on Mathematical Software*, vol. 29, pp. 141–164, 2003.
- [2] E. Doedel, A. Champneys, F. Dercole, T. Fairgrieve, Y. Kuznetsov, B. Oldeman, R. Paffenroth, B. Sandstede, X. Wang, and C. Zhang, “Auto: Software for continuation and bifurcation problems in ordinary differential equations,” Concordia University, Montreal, Tech. Rep., 2009.
- [3] J. A. Lee and M. Verleysen, *Nonlinear dimensionality reduction*. London: Springer, 2007.
- [4] P. Holmes, J. Lumley, and G. Berkooz, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge: Cambridge University Press, 1996.
- [5] L. Sirovich, “Modeling the functional organization of the visual cortex,” *Physica D: Nonlinear Phenomena*, vol. 96, no. 1-4, pp. 355–366, 1996.
- [6] A. Chatterjee, “An introduction to the proper orthogonal decomposition,” *Current Science*, vol. 78, no. 7, pp. 808–817, 2000.
- [7] M.-L. Rapún and J. M. Vega, “Reduced order models based on local POD plus Galerkin projection,” *J. Comput. Phys.*, vol. 229, no. 8, pp. 3046–3063, 2010.
- [8] D. Alonso, A. Velazquez, and J. M. Vega, “A method to generate computationally efficient reduced order models,” *Computer Methods in Applied Mechanics and Engineering*, vol. 198, no. 33-36, pp. 2683–2691, 2009.
- [9] K. Unisch and S. Volkwein, “Galerkin proper orthogonal decomposition methods for a general equation in fluid dynamics,” *SIAM Journal on Numerical Analysis*, vol. 40, no. 2, pp. 492–515, 2003.
- [10] E. Ding, E. Shlizerman, and J. N. Kutz, “Modeling multipulsing transition in ring cavity lasers with proper orthogonal decomposition,” *Phys. Rev. A*, vol. 82, p. 023823, 2010.
- [11] P. del Sastre and R. Bermejo, “The POD technique for computing bifurcation diagrams: A comparison among different models in fluids,” in *Numerical Mathematics and Advanced Applications*, A. B. de Castro, D. Gómez, P. Quintela, and P. Salgado, Eds. Springer Berlin Heidelberg, 2006, pp. 880–888.

- [12] X. Chen and S. Akella, “A dual-weighted trust-region adaptive POD 4-d var applied to a finite-volume shallow water equations model on the sphere,” *International Journal for Numerical Methods in Fluids*, 2012.
- [13] A. Liakopoulos, P. A. Blythe, and H. Gunes, “A reduced dynamical model of convective flows in tall laterally heated cavities,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 453, no. 1958, pp. 663–672, 1997.
- [14] H. S. Eisenberg, Y. Silberberg, R. Morandotti, A. R. Boyd, and J. S. Aitchison, “Discrete spatial optical solitons in waveguide arrays,” *Phys. Rev. Lett.*, vol. 81, no. 16, pp. 3383–3386, 1998.
- [15] A. B. Aceves, C. De Angelis, T. Peschel, R. Muschall, F. Lederer, S. Trillo, and S. Wabnitz, “Discrete self-trapping, soliton interactions, and beam steering in nonlinear waveguide arrays,” *Phys. Rev. E*, vol. 53, no. 1, pp. 1172–1189, 1996.
- [16] H. S. Eisenberg, R. Morandotti, Y. Silberberg, J. M. Arnold, G. Pennelli, and J. S. Aitchison, “Optical discrete solitons in waveguide arrays. i. soliton formation,” *J. Opt. Soc. Am. B*, vol. 19, no. 12, pp. 2938–2944, 2002.
- [17] U. Peschel, R. Morandotti, J. M. Arnold, J. S. Aitchison, H. S. Eisenberg, Y. Silberberg, T. Pertsch, and F. Lederer, “Optical discrete solitons in waveguide arrays. 2. dynamic properties,” *J. Opt. Soc. Am. B*, vol. 19, no. 11, pp. 2637–2644, 2002.
- [18] S. Jensen, “The nonlinear coherent coupler,” *IEEE Journal of Quantum Electronics*, vol. 18, no. 10, pp. 1580–1583, 1982.
- [19] T. Pertsch, U. Peschel, and F. Lederer, “All-optical switching in quadratically nonlinear waveguide arrays,” *Opt. Lett.*, vol. 28, no. 2, pp. 102–104, 2003.
- [20] R. A. Vicencio, M. I. Molina, and Y. S. Kivshar, “Switching of discrete optical solitons in engineered waveguide arrays,” *Phys. Rev. E*, vol. 70, no. 2, p. 26602, 2004.
- [21] R. A. Vicencio, M. Molina, and Y. S. Kivshar, “All-optical switching and amplification of discrete vector solitons in nonlinear cubic birefringent waveguide arrays,” *Opt. Lett.*, vol. 29, no. 24, pp. 2905–2907, 2004.
- [22] A. B. Aceves, G. G. Luther, C. De Angelis, A. M. Rubenchik, and S. K. Turitsyn, “Energy localization in nonlinear fiber arrays: Collapse-effect compressor,” *Phys. Rev. Lett.*, vol. 75, no. 1, pp. 73–76, 1995.
- [23] D. Hudson, K. Shish, T. Schibli, J. N. Kutz, D. Christodoulides, R. Morandotti, and S. Cundiff, “Nonlinear femtosecond pulse reshaping in waveguide arrays,” *Opt. Lett.*, vol. 33, pp. 1440–1442, 2008.

- [24] M. Matuszewski, I. L. Garanovich, and A. A. Sukhorukov, "Light bullets in nonlinear periodically curved waveguide arrays," *Phys. Rev. A*, vol. 81, no. 4, p. 43833, 2010.
- [25] M. O. Williams, C. W. McGrath, and J. N. Kutz, "Light-bullet routing and control with planar waveguide arrays," *Opt. Express*, vol. 18, no. 11, pp. 11 671–11 682, 2010.
- [26] J. N. Kutz, C. Conti, and S. Trillo, "Mode-locked x-wave lasers," *Opt. Express*, vol. 15, no. 24, pp. 16 022–16 028, 2007.
- [27] Y. Lahini, E. Frumker, Y. Silberberg, S. Droulias, K. Hizanidis, R. Morandotti, and D. N. Christodoulides, "Discrete x-wave formation in nonlinear waveguide arrays," *Phys. Rev. Lett.*, vol. 98, no. 2, p. 23901, 2007.
- [28] P. G. Kevrekidis, J. Gagnon, D. J. Frantzeskakis, and B. A. Malomed, "X , y , and z waves: Extended structures in nonlinear lattices," *Phys. Rev. E*, vol. 75, no. 1, p. 16607, 2007.
- [29] D. N. Christodoulides, F. Lederer, Y. Silberberg, F. Lederer, and Y. Silberberg, "Discretizing light behaviour in linear and nonlinear waveguide lattices," *Nature*, vol. 424, pp. 817–823, 2003.
- [30] J. Proctor and J. N. Kutz, "Theory and simulation of passive mode-locking with waveguide arrays," *Opt. Lett.*, vol. 13, pp. 1015–2013, 2005.
- [31] J. N. Kutz and B. Sandstede, "Theory of passive harmonic mode-locking using waveguide arrays," *Opt. Express*, vol. 16, no. 2, pp. 636–650, 2008.
- [32] B. G. Bale, J. N. Kutz, and B. Sandstede, "Optimizing waveguide array mode-locking for high-power fiber lasers," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 15, pp. 220–231, 2009.
- [33] B. G. Bale and J. N. Kutz, "Variational method for mode-locked lasers," *Journal of the Optical Society of America B*, vol. 25, pp. 1193–1202, 2008.
- [34] H. Haus, "Mode-locking of lasers," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 6, no. 6, pp. 1173–1185, 2000.
- [35] J. N. Kutz, "Mode-locked soliton lasers," *SIAM Review*, vol. 48, pp. 629–678, 2006.
- [36] J.-Y. Lu and J. F. Greenleaf, "Nondiffracting x waves-exact solutions to free-space scalar wave equation and their finite aperture realizations," *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 39, no. 1, pp. 19–31, 1992.

- [37] H. E. Hernandez-Figueroa, M. Zamboni-Rached, E. Recami, and H. E. Hernandez-Figueroa, *Localized Waves*, ser. Wiley Series in Microwave and Optical Engineering, K. Chang, Ed. John Wiley and Sons, 2007.
- [38] C. Conti, S. Trillo, P. Di Trapani, G. Valiulis, A. Piskarskas, O. Jedrkiewicz, and J. Trull, “Nonlinear electromagnetic x waves,” *Phys. Rev. Lett.*, vol. 90, no. 17, p. 170406, 2003.
- [39] P. Di Trapani, G. Valiulis, A. Piskarskas, O. Jedrkiewicz, J. Trull, C. Conti, and S. Trillo, “Spontaneously generated x-shaped light bullets,” *Phys. Rev. Lett.*, vol. 91, no. 9, p. 93904, 2003.
- [40] M. Kolesik, E. M. Wright, and J. V. Moloney, “Dynamic nonlinear x waves for femtosecond pulse propagation in water,” *Phys. Rev. Lett.*, vol. 92, no. 25, p. 253901, 2004.
- [41] D. Faccio, M. A. Porras, A. Dubietis, F. Bragheri, A. Couairon, and P. Di Trapani, “Conical emission, pulse splitting, and x-wave parametric amplification in nonlinear dynamics of ultrashort light pulses,” *Phys. Rev. Lett.*, vol. 96, no. 19, p. 193901, 2006.
- [42] C. Conti and S. Trillo, “Nonspreading wave packets in three dimensions formed by an ultracold Bose gas in an optical lattice,” *Phys. Rev. Lett.*, vol. 92, no. 12, p. 120404, 2004.
- [43] S. Longhi and D. Janner, “X-shaped waves in photonic crystals,” *Phys. Rev. B*, vol. 70, no. 23, p. 235123, 2004.
- [44] K. Staliunas and M. Tlidi, “Hyperbolic transverse patterns in nonlinear optical resonators,” *Phys. Rev. Lett.*, vol. 94, no. 13, p. 133902, 2005.
- [45] Y. Silberberg, “Collapse of optical pulses,” *Opt. Lett.*, vol. 15, pp. 1282–1284, 1990.
- [46] F. Wise and P. D. Trapani, “The hunt for light bullets – spatiotemporal solitons,” *Optics and Photonics News*, vol. 13, pp. 28–32, 2002.
- [47] A. A. Sukhorukov and Y. S. Kivshar, “Slow light bullets in arrays of nonlinear Bragg-grating waveguides,” in *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies*. Optical Society of America, 2006, p. JWB82.
- [48] R. H. Enns and S. S. Rangnekar, “Bistable spheroidal optical solitons,” *Phys. Rev. A*, vol. 45, no. 5, pp. 3354–3357, 1992.

- [49] A. B. Blagoeva, S. G. Dinev, A. A. Dreischuh, and A. Naidenov, "Light bullets formation in a bulk media," *IEEE J. Quant. Elec.*, vol. 27, no. 8, pp. 2060–2065, 1991.
- [50] Y. V. Kartashov, L. Torner, and D. N. Christodoulides, "Soliton dragging by dynamic optical lattices," *Opt. Lett.*, vol. 30, no. 11, pp. 1378–1380, 2005.
- [51] W. Królikowski, U. Trutschel, M. Cronin-Golomb, and C. Schmidt-Hattenberger, "Solitonlike optical switching in a circular fiber array," *Opt. Lett.*, vol. 19, no. 5, pp. 320–322, 1994.
- [52] J. Meier, G. I. Stegeman, D. N. Christodoulides, Y. Silberberg, R. Morandotti, H. Yang, G. Salamo, M. Sorel, and J. S. Aitchison, "Beam interactions with a blocker soliton in one-dimensional arrays," *Opt. Lett.*, vol. 30, no. 9, pp. 1027–1029, 2005.
- [53] Y. Tanguy, T. Ackemann, W. J. Firth, and R. Jäger, "Realization of a semiconductor-based cavity soliton laser," *Phys. Rev. Lett.*, vol. 100, p. 13907, 2008.
- [54] S. Barland, J. R. Tredicce, M. Brambilla, L. A. Lugiato, S. Balle, M. Giudici, T. Maggipinto, L. Spinelli, G. Tissoni, T. Knödl, M. Miller, R. Jäger, T. Knodl, and R. Jager, "Cavity solitons as pixels in semiconductor microcavities," *Nature*, vol. 419, no. 6908, pp. 699–702, 2002.
- [55] V. B. Taranenko and C. O. Weiss, "Incoherent optical switching of semiconductor resonator solitons," *Applied Physics B*, vol. 72, pp. 893–895, 2001.
- [56] L. F. Mollenauer and J. P. Gordon, *Solitons in Optical Fibers: Fundamentals and Applications*. Springer, 2006.
- [57] M. Feng, M. O. Williams, J. N. Kutz, K. L. Silverman, R. P. Mirin, and S. T. Cundiff, "Intensity dynamics in a waveguide array laser," *Optics Communications*, vol. 284, no. 4, pp. 971–978, 2011.
- [58] U. Keller and A. C. Tropper, "Passively modelocked surface-emitting semiconductor lasers," *Physics Reports*, vol. 429, no. 2, pp. 67–120, 2006.
- [59] J. Marangos, "Slow light in cool atoms," *Nature*, vol. 397, pp. 559–560, 1999.
- [60] J. T. Mok, C. M. de Sterke, I. C. M. Liter, and B. J. Eggleton, "Dispersionless slow light using gap solitons," *Nature Physics*, vol. 2, pp. 775–780, 2006.
- [61] P. Y. P. Chen, B. A. Malomed, and P. L. Chu, "Trapping Bragg solitons by a pair of defects," *Phys. Rev. E*, vol. 71, p. 066601, 2005.
- [62] A. Yariv, *Quantum Electronics*. John Wiley and Sons, 1988.

- [63] P. Kockaert, P. Tassin, G. van der Sande, I. Veretennicoff, and M. Tlidi, “Negative diffraction pattern dynamics in nonlinear cavities with left-handed materials,” *Phys. Rev. A*, vol. 74, no. 3, pp. 033 822–+, 2006.
- [64] L. Rahman and H. G. Winful, “Nonlinear dynamics of semiconductor laser arrays: a mean field model,” *IEEE J. Quant. Elec.*, vol. 30, no. 6, pp. 1405–1416, 1994.
- [65] L. N. Trefethen and D. Bau, *Numerical Linear Algebra*. SIAM: Philadelphia, 1997.
- [66] M. O. Williams and J. N. Kutz, “Spatial mode-locking of light bullets in planar waveguide arrays,” *Opt. Express*, vol. 17, no. 20, pp. 18 320–18 329, 2009.
- [67] G. R. Hadley, “Transparent boundary conditions for the beam propagation method,” *IEEE J. Quant. Elec.*, vol. 28, pp. 363–370, 1992.
- [68] *See the FunFacs website (www.funfacs.org)*.
- [69] S. Wabnitz, Y. Kodama, and A. B. Aceves, “Control of optical soliton interactions,” *Optical Fiber Technology*, vol. 1, pp. 187–217, 1995.
- [70] B. A. Malomed, “Chapter 2 variational methods in nonlinear fiber optics and related fields,” ser. Progress in Optics, E. Wolf, Ed., Elsevier, 2002, vol. 43, pp. 71–193.
- [71] D. Anderson, M. Lisak, and A. Berntson, “A variational approach to nonlinear evolution equations in optics,” *Pramana*, vol. 57, no. 5-6, pp. 917–936, 2001.
- [72] B. Botez and G. Peterson, “Modes of phase-locked diode-laser arrays of closely spaced antiguides,” *Elec. Lett.*, vol. 24, pp. 1042–1044, 1988.
- [73] J. Buus, “The effective index method and its application to semiconductor lasers,” *IEEE J. Quant. Elec.*, vol. 18, pp. 1083–1089, 1982.
- [74] G. P. Agrawal, “Lateral analysis of quasi-index-guided injection lasers: transition from gain to index guiding,” *Journal of Lightwave Technology*, vol. 2, pp. 537–543, 1984.
- [75] S. Wang and H. Winful, “Dynamics of phase-locked semiconductor laser arrays,” *Applied Physics Lett.*, vol. 52, pp. 1774–1776, 1988.
- [76] L. Rahman and H. Winful, “Improved coupled-mode theory for the dynamics of semiconductor laser arrays,” *Opt. Lett.*, vol. 18, pp. 128–130, 1993.
- [77] S. Wang and H. Winful, “Propagation model for the dynamics of gain-guided semiconductor laser arrays,” *J. App. Phys.*, vol. 73, pp. 462–464, 1993.

- [78] K. Kawano and T. Kitoh, *Introduction to Optical Waveguide Analysis: Solving Maxwell's Equations and the Schrodinger Equation*. John Wiley and Sons, 2001.
- [79] B. G. Bale, K. Kieu, J. N. Kutz, and F. Wise, "Transition dynamics for multi-pulsing in mode-locked lasers." *Opt. Express*, vol. 17, no. 25, pp. 23 137–23 146, 2009.
- [80] J. M. Soto-Crespo, M. Grapinet, P. Grelu, and N. Akhmediev, "Bifurcations and multiple period soliton pulsations in a passively mode-locked fiber laser," *Phys. Rev. E*, vol. 70, no. 066612, 2004.
- [81] S. Namiki, E. P. Ippen, H. A. Haus, and C. X. Yu, "Energy rate equations for mode-locked lasers," *Journal of the Optical Society of America B*, vol. 14, no. 8, pp. 2099–2111, 1997.
- [82] L. Feng, P. K. A. Wai, and J. N. Kutz, "A geometrical description of the onset of multi-pulsing in mode-locked laser cavities," *JOSA B*, p. (to appear), 2010.
- [83] C. R. Jones and J. N. Kutz, "Stability of mode-locked pulse solutions subject to saturable gain: computing linear stability with the Floquet–Fourier–Hill method," *JOSA B*, vol. 27, no. 6, pp. 1184–1194, 2010.
- [84] E. Shlizerman and V. Rom-Kedar, "Parabolic resonance: A route to Hamiltonian spatiotemporal chaos," *Phys. Rev. Lett.*, vol. 102, no. 3, p. 33901, 2009.
- [85] G. M. Shroff and H. B. Keller, "Stabilization of unstable procedures: The recursive projection method," *SIAM Journal on Numerical Analysis*, vol. 30, pp. 1099–1120, 1993.
- [86] K. Lust, D. Roose, A. Spence, and A. R. Champneys, "an adaptive Newton – Picard algorithm with subspace iteration for computing periodic solutions," *Society*, vol. 19, no. 4, pp. 1188–1209, 1998.
- [87] D. M. Ambrose and J. Wilkening, "Computation of time-periodic solutions of the Benjamin–Ono equation," *Journal of Nonlinear Science*, vol. 20, no. 3, pp. 277–308, 2010.
- [88] J. Wilkening, "Notes on time-stepping the gravity-driven water wave," (contact the author for details).
- [89] ———, "Breakdown of self-similarity at the crests of large-amplitude standing water waves," *Phys. Rev. Lett.*, vol. 107, p. 184501, 2011.
- [90] S. Wiggins, *Introduction to Applied Nonlinear Dynamical Systems and Chaos*, 2nd ed. Springer, New York, 1990.

- [91] M. O. Williams, E. Shlizerman, and J. N. Kutz, “The multi-pulsing transition in mode-locked lasers: a low-dimensional approach using waveguide arrays,” *Journal of the Optical Society of America B*, vol. 27, no. 12, pp. 2471–2481, 2010.
- [92] J. R. Dormand and P. J. Prince, “A family of embedded Runge-Kutta formulae,” *Journal of Computational and Applied Mathematics*, vol. 6, pp. 19–26, 2006.
- [93] M. Rathinam and L. Petzold, “A new look at proper orthogonal decomposition,” *SIAM Journal on Numerical Analysis*, vol. 41, no. 5, pp. 1893–1925, 2004.
- [94] A. Quarteroni, G. Rozza, and A. Manzoni, “Certified reduced basis approximation for parametrized partial differential equations and applications,” *Journal of Mathematics in Industry*, vol. 1, pp. 1–49, 2011.
- [95] M. C. Cross and P. C. Hohenberg, “Pattern formation outside of equilibrium,” *Review of Modern Physics*, vol. 65, no. 3, p. 851, 1993.
- [96] M. Ilak and C. W. Rowley, “Modeling of transitional channel flow using balanced proper orthogonal decomposition,” *Physics of Fluids*, vol. 20, no. 3, p. 34103, 2008.
- [97] K. Chen, J. Tu, and C. Rowley, “Variants of dynamic mode decomposition: connections between Koopman and Fourier analyses,” *Aerospace*, 2011.
- [98] S. L. Brunton and C. W. Rowley, “Low-dimensional state-space representations for classical unsteady aerodynamic models,” *New Horizons*, no. January, pp. 4–7, 2011.
- [99] A. D’Avella, P. Saltiel, and E. Bizzi, “Combinations of muscle synergies in the construction of a natural motor behavior.” *Nature neuroscience*, vol. 6, no. 3, pp. 300–8, 2003.
- [100] E. Shlizerman and K. Schroder, “Neural activity measures and their dynamics,” *SIAM Journal on Applied Mathematics*, (to appear).
- [101] A. Tsoumanis, C. Siettos, I. Kevrekidis, and G. Bafas, “Equation-free multiscale computations in social networks: from agent-based modelling to coarse-grained stability and bifurcation analysis,” *International Journal of Bifurcation and Chaos*, vol. 20, no. 11, p. 3673, 2009.
- [102] P. J. Schmid, “Dynamic mode decomposition of numerical and experimental data,” *J. Fluid. Mech.*, vol. 656, pp. 5–28, 2010.
- [103] P. J. Schmid, L. Li, M. P. Juniper, and O. Pust, “Applications of the dynamic mode decomposition,” *Theoretical and Computational Fluid Dynamics*, vol. 25, no. 1-4, pp. 249–259, 2010.

- [104] P. J. Schmid, “Application of the dynamic mode decomposition to experimental data,” *Exp. Fluids*, vol. 50, no. 4, pp. 1123–1130, 2011.
- [105] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, “Spectral analysis of nonlinear flows,” *J. Fluid. Mech.*, vol. 641, pp. 115–127, 2009.
- [106] I. Mezić, “Spectral properties of dynamical systems, model reduction and decompositions,” *Nonlinear Dynamics*, vol. 41, pp. 309–325, 2005.
- [107] P. J. Schmid, “Dynamic mode decomposition of numerical and experimental data,” *J. Fluid. Mech.*, vol. 656, pp. 5–28, 2010.
- [108] E. Shlizerman, E. Ding, M. O. Williams, and J. N. Kutz, “The proper orthogonal decomposition for dimensionality reduction in mode-locked lasers and optical systems,” *International Journal of Optics*, vol. 2012, no. 11, pp. 1–18, 2011.
- [109] M. Gu and S. C. Eisenstat, “Efficient algorithms for computing a strong rank-revealing QR factorization,” *SIAM J. Sci. Comput.*, vol. 17, no. 4, pp. 848–869, 1996.
- [110] N. Halko, P. G. Martinsson, and J. A. Tropp, “Finding structure with randomness : Probabilistic algorithms for matrix decompositions,” *SIAM Review*, vol. 53, no. 2, pp. 217–288, 2011.
- [111] M. O. Williams, J. Wilkening, E. Shlizerman, and J. N. Kutz, “Continuation of periodic solutions in the waveguide array mode-locked laser,” *Physica D*, 2011.
- [112] L. N. Trefethen, *Spectral methods in MatLab*. SIAM: Philadelphia, 2000.
- [113] B. Sonday, A. Singer, C. W. Gear, and I. G. Kevrekidis, “Manifold learning techniques and model reduction applied to dissipative PDEs,” (submitted to the SIAM Journal on Applied Dynamical Systems).
- [114] R. S. Johnson, *A modern introduction to the mathematical theory of water waves*. Cambridge, UK: Cambridge University Press, 1997.
- [115] H. Lamb, *Hydrodynamics*. Cambridge: University Press, 1895.
- [116] V. E. Zakharov, “Stability of periodic waves on finite amplitude on the surface of a deep fluid,” *Zhurnal Prikladnoi Mekhaniki i Tekhnicheskoi Fiziki*, vol. 9, pp. 86–94, 1968.
- [117] J. W. Miles, “On Hamilton’s principle for surface waves,” *J. Fluid. Mech.*, vol. 83, pp. 153–158, 1977.

- [118] J. Eggers and M. A. Fontelos, “The role of self-similarity in singularities of partial differential equations,” *Nonlinearity*, vol. 22, pp. R1—R44, 2009.
- [119] T. J. Bridges, “Wave breaking and the surface velocity field for three-dimensional water waves,” *Nonlinearity*, vol. 22, pp. 947–953, 2009.
- [120] A. Dyachenko and V. Zakharov, “Modulation instability of Stokes wave\ freak wave,” *Journal of Experimental and Theoretical Physics Letters*, vol. 81, pp. 255–259, 2005.
- [121] J. L. Hammack, “A note on tsunamis: their generation and propagation in an ocean of uniform depth,” *J. Fluid. Mech.*, vol. 60, no. 4, pp. 769–799, 1973.
- [122] Y. Kervella, D. Dutykh, and F. Dias, “Comparison between three-dimensional linear and nonlinear tsunami generation models,” *Theor. Comput. Fluid Dyn.*, vol. 21, pp. 245–269, 2007.
- [123] J. Cruz, *Ocean Wave Energy: Current Status and Future Perspectives*. Springer, 2008.
- [124] J. W. S. Rayleigh, “Deep water waves, progressive or stationary, to the third order approximation,” *Proc. R. Soc. A*, vol. 91, pp. 345–353, 1915.
- [125] G. G. Stokes, “Considerations relative to the greatest height of oscillatory irrotational waves which can be propagated without change of form,” in *Mathematical and physical papers*. Cambridge University Press, 1880, vol. 1, pp. 225–228.
- [126] A. D. D. Craik, “George Gabriel Stokes on water wave theory,” *Ann. Rev. Fluid Mech.*, vol. 37, pp. 23–42, 2005.
- [127] C. J. Amick, L. E. Fraenkel, and J. F. Toland, “On the Stokes conjecture for the wave of extreme form,” *Acta Math.*, vol. 148, no. 1, pp. 192–214, 1982.
- [128] W. G. Penney and A. T. Price, “Finite periodic stationary gravity waves in a perfect liquid, part II,” *Phil. Trans. R. Soc. London A*, vol. 244, pp. 254–284, 1952.
- [129] G. I. Taylor, “An experimental study of standing waves,” *Proc. Roy. Soc. A*, vol. 218, pp. 44–59, 1953.
- [130] I. Tadjbakhsh and J. B. Keller, “Standing surface waves of finite amplitude,” *J. Fluid. Mech.*, vol. 8, pp. 442–451, 1960.
- [131] P. Concus, “Standing capillary-gravity waves of finite amplitude,” *J. Fluid. Mech.*, vol. 14, pp. 568–576, 1962.

- [132] J. M. Vanden-Broeck, “Nonlinear gravity-capillary standing waves in water of arbitrary uniform depth,” *J. Fluid. Mech.*, vol. 139, pp. 97–104, 1984.
- [133] G. Iooss, P. I. Plotnikov, and J. F. Toland, “Standing waves on an infinitely deep perfect fluid under gravity,” *Arch. Rat. Mech. Anal.*, vol. 177, pp. 367–478, 2005.
- [134] G. N. Mercer and A. J. Roberts, “Standing waves in deep water: Their stability and extreme form,” *Phys. Fluids A*, vol. 4, no. 2, pp. 259–269, 1992.
- [135] ———, “The form of standing waves on finite depth water,” *Wave Motion*, vol. 19, pp. 233–244, 1994.
- [136] D. H. Smith and A. J. Roberts, “Branching behavior of standing waves — the signatures of resonance,” *Physics of Fluids*, vol. 11, pp. 1051—1064, 1999.
- [137] C. P. Tsai and D. S. Jeng, “Numerical {Fourier} solutions of standing waves in finite water depth,” *Appl. Ocean Res.*, vol. 16, pp. 185–193, 1994.
- [138] P. J. Bryant and M. Stiassnie, “Different forms for nonlinear standing waves in deep water,” *J. Fluid. Mech.*, vol. 272, pp. 135–156, 1994.
- [139] W. W. Schultz, J. M. Vanden-Broeck, L. Jiang, and M. Perlin, “Highly nonlinear standing water waves with small capillary effect,” *J. Fluid. Mech.*, vol. 369, pp. 253–272, 1998.
- [140] M. Okamura, “Standing gravity waves of large amplitude on deep water,” *Wave Motion*, vol. 37, pp. 173–182, 2003.
- [141] ———, “Almost limiting short-crested gravity waves in deep water,” *J. Fluid. Mech.*, vol. 646, pp. 481–503, 2010.
- [142] M. S. Longuet-Higgins and M. J. H. Fox, “Theory of the almost-highest wave: the inner solution,” *J. Fluid. Mech.*, vol. 80, no. 4, pp. 721–741, 1977.
- [143] K. S. Turitsyn, L. Lai, and W. W. Zhang, “Asymmetric disconnection of an underwater air bubble: persistent neck vibrations evolve into a smooth contact,” *Phys. Rev. Lett.*, vol. 103, p. 124501, 2009.
- [144] L. Jiang, C. Ting, M. Perlin, and W. W. Schultz, “Moderate and steep faraday waves: instabilities, modulation and temporal asymmetries,” *J. Fluid. Mech.*, vol. 329, pp. 275–307, 1996.
- [145] N. Périnet, D. Juric, and L. S. Tuckerman, “Numerical simulation of faraday waves,” *J. Fluid. Mech.*, vol. 635, pp. 1–26, 2009.

- [146] G. Haller, *Chaos Near Resonance*, ser. Applied Mathematical Sciences; 138. Springer-Verlag, NY, 1999.
- [147] J. Wilkening and J. Yu, “Overdetermined shooting methods for high resolution computations of time-periodic water waves.”
- [148] D. P. Nicholls, “Boundary perturbation methods for water waves,” *GAMM-Mitteilungen*, vol. 30, pp. 44—74, 2007.
- [149] W. Craig and C. Sulem, “Numerical simulation of gravity waves,” *J. Comp. Phys.*, vol. 108, pp. 73–83, 1993.
- [150] J. D. Fenton and M. M. Rienecker, “A fourier method for solving nonlinear water-wave problems: application to solitary-wave interactions,” *J. Fluid. Mech.*, vol. 118, pp. 411–443, 1982.
- [151] F. Dias and T. J. Bridges, “The numerical computation of freely propagating time-dependent irrotational water waves,” *Fluid Dynamics Research*, vol. 38, no. 12, p. 803, 2006.
- [152] A. L. Dychenko, V. E. Zakharov, and E. A. Kuznetsov, “Nonlinear dynamics on the free surface of an ideal fluid,” *Plasma Phys. Rep.*, vol. 22, pp. 916—928, 1996.
- [153] P. A. Milewski, J.-M. Vanden-Broeck, and Z. Wang, “Dynamics of steep two-dimensional gravity capillary solitary waves,” *J. Fluid. Mech.*, 2010.
- [154] B. Podvin and P. Le Quéré, “Low-order models for the flow in a differentially heated cavity,” *Physics of Fluids*, vol. 13, no. 11, pp. 3204—3214, 2001.
- [155] G. R. Baker, D. I. Meiron, and S. A. Orszag, “Generalized vortex methods for free-surface flow problems,” *J. Fluid. Mech.*, vol. 123, pp. 477—501, 1982.
- [156] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. Springer-Verlag, 2006.
- [157] N. Phillips, “A coordinate system having some special advantages for numerical forecasting,” *J. Meteor.*, vol. 14, pp. 184—185, 1957.
- [158] J. Chandezon, D. Maystre, and G. Raoult, “A new theoretical method for diffraction gratings and its numerical application,” *J. Optics*, vol. 11, pp. 235—241, 1980.
- [159] M. Hinze and S. Volkwein, “Proper orthogonal decomposition surrogate models for nonlinear dynamical systems: Error estimates and suboptimal control,” in *Dimension Reduction of Large-Scale Systems, Lecture Notes in Computational and Applied Mathematics*, D. Sorensen, P. Benner, and V. Mehrmann, Eds. Springer Berlin Heidelberg, 2005, pp. 261—306.

- [160] K. Kunisch and S. Volkwein, “Galerkin proper orthogonal decomposition methods for parabolic problems,” *Numer. Math.*, vol. 90, pp. 117—148, 2001.
- [161] D. Anderson, “Variational approach to nonlinear pulse propagation in optical fibers,” *Phys. Rev. A*, vol. 27, no. 6, pp. 3135–3145, 1983.
- [162] E. Doedel, H. B. Keller, and J. P. Kernévez, “Numerical analysis and control of bifurcation problems: (i) bifurcation in finite dimension,” *Int. J. Bifurcation and Chaos*, vol. 1, no. 3, pp. 493–520, 1991.
- [163] E. J. Doedel, H. B. Keller, and J. P. Kernévez, “Numerical analysis and control of bifurcation problems: (ii) bifurcation in infinite dimensions.” *Int. J. Bifurcation and Chaos*, vol. 1, no. 4, pp. 745–772, 1991.
- [164] W. J. F. Govaerts, *Numerical methods for bifurcations of dynamical equilibria*. SIAM, Philadelphia, 2000.
- [165] H. B. Keller, *Numerical methods in bifurcation problems*. Springer, New York, 2000.
- [166] D. M. Ambrose and J. Wilkening, “Global paths of time-periodic solutions of the Benjamin–Ono equation connecting pairs of traveling waves.” *Comm. App. Math. and Comp. Sci.*, vol. 4, no. 1, pp. 177–215, 2009.
- [167] ———, “Computation of symmetric, time-periodic solutions of the vortex sheet with surface tension.” *Proceedings of the National Academy of Science*, vol. 107, no. 8, pp. 3361–3366, 2010.
- [168] H. B. Keller, *Numerical Methods for Two-Point Boundary-Value Problems*. Blaisdell, New York, 1968.
- [169] J. Stoer and R. Bulirsch, *Introduction to Numerical Analysis*, 3rd ed. Springer, New York, 2002.
- [170] M. O. Bristeau, O. Pironneau, R. Glowinsky, J. Periaux, and P. Perrier, “On the numerical solution of nonlinear problems in fluid dynamics by least squares and finite element methods. i - least square formulations and conjugate gradient solution of the continuous problems.” *Computer Methods in Applied Mechanics and Engineering*, vol. 17, pp. 619–657, 1979.
- [171] A. Jameson, “Aerodynamic design via control theory.” *J. Sci. Comput.*, vol. 3, no. 3, pp. 233–260, 1988.
- [172] B. Mohammadi and O. Pironneau, *Applied Shape Optimization for Fluids*. Oxford University Press, New York, 2001.

- [173] G. J. Cooper and A. Sayfy, “Additive Runge-Kutta methods for stiff ordinary differential equations.” *Math. Comp.*, vol. 40, pp. 207–218, 1983.
- [174] C. A. Kennedy and M. H. Carpenter, “Additive Runge-Kutta schemes for convection-diffusion-reaction equations,” *Applied Numerical Mathematics*, vol. 44, pp. 139–181, 2003.
- [175] U. M. Ascher, S. J. Ruuth, and B. T. R. Wetton, “Implicit-explicit methods for time-dependent partial differential equations,” *SIAM J. Numer. Anal.*, vol. 32, pp. 797–823, 1995.
- [176] E. Hairer, S. P. Norsett, and G. Wanner., *Solving Ordinary Differential Equations I: Nonstiff Problems*, 2nd ed. Springer, Berlin, 2000.
- [177] J. Wilkening, “A fifth order additive Runge-Kutta method with a fourth order dense output,” *Technical Report, Lawrence Berkeley National Laboratory*, 2010.
- [178] L. F. Shampine, “Some practical Runge-Kutta formulas,” *Mathematics of Computation*, vol. 46, pp. 135–150, 1986.
- [179] T. Y. Hou and R. Li, “Computing nearly singular solutions using pseudo-spectral methods,” *J. Comput. Phys.*, vol. 226, pp. 379–397, 2007.