The Low Dimensionality of Time-Periodic Standing Waves in Water of Finite and Infinite Depth

Matthew O. Williams[†][*], Eli Shlizerman[†], Jon Wilkening[‡], and J. Nathan Kutz[†]

Abstract. Time periodic standing waves are demonstrated to be low-dimensional by use of the proper orthogonal decomposition (POD). Moreover, the nonlinear dynamics of the system restricted to this low-dimensional linear subspace are shown to accurately recover the spatio-temporal full PDE dynamics. A global set of modes, generated with sequential POD, are then used to produce time-periodic standing wave branches as a function of the period. This representation quantitatively reproduces the entire branch, including both large and small amplitude solutions, using only a few POD modes. This technique gives the potential for a new direction of exploration in this challenging problem, including an efficient way to characterize the bifurcation structure and stability of these solutions.

Key words.

AMS subject classifications.

1. Introduction. 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 [29, 32]. 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 [38, 64], the nonlinear interaction of dynamics at small and large scales, and the possibility of finite-time singularity formation [3, 19]. From a practical point of view, improving our understanding of ocean wave dynamics could help minimize damage due to rogue waves [17] and tsunamis [22, 30] and aid in the design of coastal power plants that extract energy from ocean waves [12], to list but a few applications.

In this manuscript, we 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 [48], who incorporated time evolution in the asymptotic techniques developed by Stokes for traveling waves of permanent form [1, 10, 54]. Penney and Price [44] 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 [56] 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 [55] studied standing waves of finite depth while Concus [8] and Vanden-Broeck [60] considered the effect of surface tension. Small amplitude standing waves were proved to exist by Iooss, Plotnikov and Toland [27] 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 [36, 37]; Smith and Roberts [53]; Tsai

[*][mowill@uw.edu]
[†]Department of Applied Mathematics, University of Washington, Seattle, WA 98195-2420 USA
[‡]Department of Mathematics, University of California, Berkeley, CA 94720-3840 USA

and Jeng [58]; Bryant and Stiassnie [5]; Schultz, Vanden-Broeck, Jiang and Perlin [50]; Okamura [42, 43]; Wilkening [61]; 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 [36] predicted interior crest angles as sharp as 60 degree; Schultz et. al. predicted the formation of a cusp rather than a corner; and Okamura [42, 43] found that Penney and Price's conjecture of a 90 degree corner holds true. Recently, Wilkening [61] 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 [3, 19, 35, 59] at the crests of very large amplitude standing waves [61]. 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 [21, 28, 45].

The purpose of this manuscript is to demonstrate 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, we show that only a few spatial modes are needed to reconstruct the waves computed in Refs. [61, 62] and other similar computations in Refs. [36, 37]. To do so, we use the proper orthogonal decomposition (POD) to demonstrate that the dynamics of the full PDE are essentially confined, in the L^2 sense, to a low dimensional linear subspace. Then using the basis for this subspace as modes, we generate reduced models by defining suitable projection operators, based on the Galerkin projection, for two discretization approaches for the PDE that describes the water wave system. The reduced models are shown to qualitatively reproduce the dynamics and bifurcation structure of the full PDE in the parameter regimes where the POD modes were obtained.

Although many other dimensionality reduction techniques exist [33], the POD method — which is essentially identical to principal component analysis (PCA), the Karhunen-Loéve (KL) expansion, the Hoteling transform, and metric multidimensional scaling (MDM) — is particularly advantageous because it is easily incorporated in numerical methods and is fundamentally related to the singular value decomposition. The POD method can be applied to either numerically generated or experimentally obtained data [16, 49], and is most often used in systems from physical, biological, or engineering applications [4, 7, 13, 15, 24–26, 34, 51, 52, 63]. With this basis in hand, the POD modes are often paired with the Galerkin projection technique to generate a system of ODEs for the mode amplitudes to study systems with spatial and temporal patterns [24].

In the dimensionality reduction context, the choice of specialized spatial basis functions plays a critical role. Typical spectral methods use basis functions, such as Fourier modes, that are not specific to any particular system. Rather, these methods are selected because they are spectrally accuracy or harness the speed of the fast Fourier transform. Therefore, many modes are often required to obtain accurate solutions. By contrast, the POD method takes a problem-specific approach. Sets of data containing information about the evolution of the system are used to generate POD modes specific to a particular flow operator. These modes are chosen so that the space they span captures more of the solution, in terms of the L^2–norm,

than any other linear basis of the same dimension [24]. As a result, even a low-dimensional POD basis can be accurate. Figure 1.1 illustrates the effectiveness of the POD method in revealing the low dimensionality of time-periodic standing waves in contrast to the Fourier method. As shown there, time-periodic solutions with relatively sharp peaks within the period require many Fourier modes to represent the peaks. Due to the number of modes required, these solutions have only recently been computed through highly sophisticated computational schemes [36, 37, 61]. In contrast, the POD approach represents these solutions with only a few modes. This suggests that using the POD basis in combination with the Galerkin projection to represent the dynamics can improve efficiency of standing wave computations.

The paper is outlined as follows: Section 2 demonstrates the accuracy of a low-dimensional POD representation of three prototypical types of standing wave solutions. Section 3 describes the governing equations for the water wave problem. Section 4 gives a brief practical description of the POD technique. Section 5 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 is similarly structured but for the finite depth case. Lastly, Section 7 contains some concluding remarks and proposed future directions.

2. Dimensionality of Periodic Solutions. The fundamental idea behind dimensionality reduction in partial differential equations is to generate a low-dimensional basis. Commonly used bases, such as Fourier modes or Chebyshev polynomials [57], are selected for their numerical properties and not due to their structure or the dynamics of the problem at hand. As a result, there may be pairs of basis functions whose time dynamics are correlated. Dimensionality reduction techniques, like the Proper Orthogonal Decomposition, seek to identify and eliminate this redundant information, and as a result, reduce the dimension of the system without a loss of accuracy.

Figure 1.1 highlights the potential of these techniques in the water wave problem by applying the POD approach to three prototypical time-periodic solutions and comparing the POD representation to a Fourier representation of equal size. The three representative solutions are: a low amplitude and nearly sinusoidal solution (left); a larger amplitude and smooth solution (center); and a large amplitude solution of extreme form (right). The first row depicts the surface height of the solution over one period. The black curve is the slice in time where the wave crest has formed. Notice that it is precisely at this wave crest (and an equivalent solution that forms a half-period later) where the solution becomes most cusp-like.

The second row shows a detailed comparison between the different reconstructions at the wave crest. The *black* curve is the solution of the full PDE. The *green* curves are POD reconstructions, and the *red* curves are the Fourier reconstructions. Dashed curves are 3-mode expansions, dotted curves are 5-mode expansions, and curves with crosses are 40-mode expansions. For the low-amplitude sinusoidal solution on the left, both Fourier and POD representations accurately represent the wave crest with only three modes. Furthermore, the POD modes for this solution are nearly identical to the Fourier cosine modes. The differences are most evident when the solution contains a sharp crest that requires high Fourier wavenumbers to accurately represent. Examples of these two cases are shown in the center and rightmost columns. In these columns, notice that both the three- and five-mode Fourier
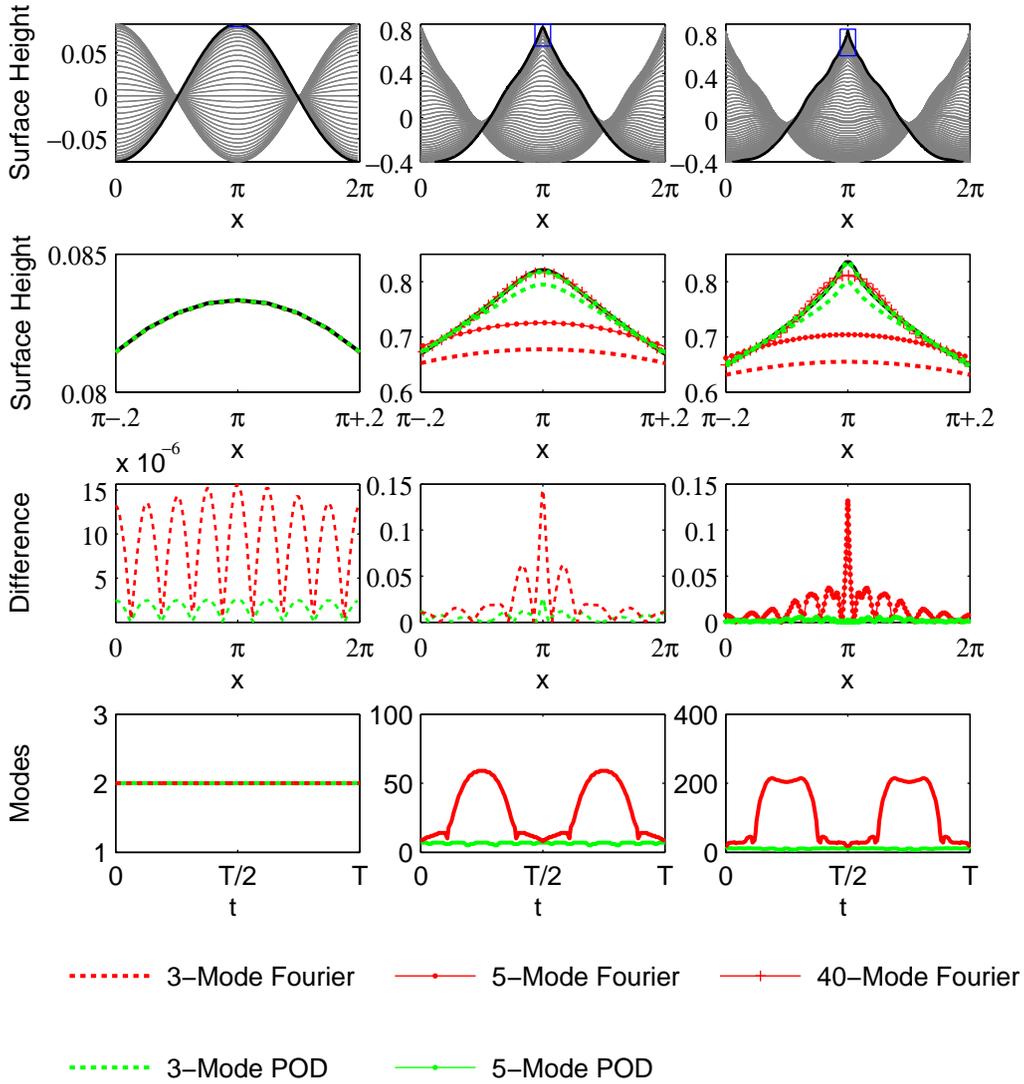
Figure 1.1.   *Comparison of POD and PDE simulations of standing waves at three different points on the branch of solutions. (Top row) Plots of the evolution over a single period. (Remaining rows) Red curves show Fourier cosine reproductions while green curves show POD reproductions. Dashed lines represent 3-mode expansions, dotted lines represent 5-mode expansions, and crosses represent 40-mode expansions. The second row gives a detailed look at the wave crest for the solution in bold. The third row plots the absolute value of the difference between the low-order reproduction and the full PDE. The final row gives the number of modes required to achieve an L^∞ error of $\leq 10^{-3}$.*

representations cannot capture the sharp wave crest. Ultimately, it takes a forty-mode Fourier expansion to accurately reproduce the solution in the center column and approximately 120-modes to accurately reproduce the right-most column. Because the POD representation is specific to a particular set of data, it can reproduce both sharply crested or sinusoidal curves using a limited number of modes. Notice that with only five-POD modes the wave crest can
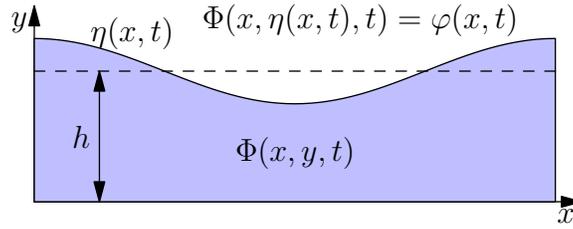
Figure 3.1. *A sketch of the water wave problem at time t with mean fluid depth h. The shape of the fluid surface is described by $\eta(x,t)$. The velocity potential in the fluid is given by $\Phi(x,y,t)$, and the velocity potential on the surface is $\varphi(x,t) = \Phi(x,\eta(x,t),t)$. All variables are 2π-periodic in x.*

be qualitatively reproduced in all three cases. Furthermore, even lower-dimensional solutions such as the three-mode expansions can capture critical details such as a sharply peaked crest; although, additional modes are required for quantitative accuracy.

The third row in Figure 1.1 shows the absolute value of the difference between the PDE and the low-dimensional representation, i.e. the absolute value of the superposition of all neglected modes, plotted on a linear scale. The error of the Fourier representation has a maximum at $x = \pi$ primarily caused by the formation of the wave crest. There is also an oscillatory error throughout the domain that roughly corresponds to the neglected mode of largest amplitude. The POD representation, overall, has a lower error than the Fourier representation. In some regimes, such as the left-column in Figure 1.1, the POD modes and Fourier modes are essentially the same. It should be noted, however, that despite the similarities, the POD error is indeed smaller. In the center and right columns the difference is clear: the error at the wave crest is visibly smaller than in the Fourier case because the POD representation has modes that can accurately represent the wave crest. The fourth row shows the number of POD or Fourier modes required to represent the solution with an L^∞ error of at most 10^{-3} for all points on the orbit, and highlights that the single set of POD modes can be as accurate as a far larger Fourier basis at all points in time. In particular, the Fourier representation requires many modes near $t = T/4$ and $t = 3T/4$ due to the sharpness of the peak when the wave crests occur, while the POD representation requires no additional modes during these times. Therefore, the POD basis is expected to be optimal for the whole period, representing the sharp peaks as well as the sinusoidal spatial structures. Because the differences are most visible near the wave crest, we have chosen to stress the comparison between the two representations there. However, the POD representation is the best linear representation, in the L^2-sense and therefore for all x, for these dynamics for the entire period.

The implications of these test cases are clear: a carefully chosen linear basis can accurately capture the dynamics that occur in the water wave problem at all points in space and in time. Determining this new optimal basis is the goal of many dimensionality reduction techniques [24, 33] and many options exist. However as demonstrated by Figure 1.1, the POD method is extremely effective in the regions of interests when computing time-periodic solutions.

3. 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:

$$\Delta\Phi = 0, \qquad\qquad 0 < y < \eta \qquad (3.1a)$$

$$\Phi_y = 0, \qquad\qquad y = 0 \qquad (3.1b)$$

$$\Phi = \varphi, \qquad\qquad y = \eta \qquad (3.1c)$$

$$\eta_t + \eta_x\Phi_x = \Phi_y, \qquad\qquad y = \eta \qquad (3.1d)$$

$$\varphi_t = P\left[-\eta_x\Phi_x\Phi_y - \frac{1}{2}\Phi_x^2 + \frac{1}{2}\Phi_y^2 - g\eta\right], \qquad\qquad y = \eta. \qquad (3.1e)$$

where

$$P[f(x)] = f(x) - \frac{1}{L}\int_0^L f(\alpha)\,d\alpha. \qquad (3.1f)$$

$\Phi(x, y, t)$ 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 are given in Figure 3.1. 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. 3.1e and defined in Eq. 3.1f, 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. 3.1 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. 3.1a 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 Zhakarov, they are canonical variables in a Hamiltonian formulation of the problem [40, 64]. 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 [9]. In that work, the DNO was approximated via a Taylor expansion [9], but there are other solution methods that rely on different approaches such as expressing the solution spectrally [20] or computing the DNO with boundary integral methods that exploit the Green's function of the Laplace operator [14]. Domain transformations like conformal mapping techniques [18, 39] 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 the solution in the right column of Figure 1.1, 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. [61]. 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 [62]. In the infinite depth case, the velocity potential in the domain is expressed as

$$\Phi(x, y) = \frac{1}{2\pi} \int_0^{2\pi} K_1(x, y, \alpha)\mu(\alpha)\, d\alpha, \tag{3.2a}$$

$$K_1(x, y, \alpha) = \text{Im}\left[\frac{1 + i\eta'(\alpha)}{2} \cot\left(\frac{(x - \alpha) + i(y - \eta(\alpha))}{2}\right)\right], \tag{3.2b}$$

$$\varphi(x) = \frac{\mu(x)}{2} + \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_1(x, \alpha)\mu(\alpha)\, d\alpha, \tag{3.2c}$$

$$\tilde{K}_1(x, \alpha) = \begin{cases} \text{Im}\left[\frac{1 + i\eta'(\alpha)}{2} \cot\left(\frac{(x - \alpha) + i(\eta(x) - \eta(\alpha))}{2}\right) - \frac{1}{2} \cot\left(\frac{x - \alpha}{2}\right)\right], & \alpha \neq x, \\ \frac{-\eta''(x)}{2(1 + \eta'(x)^2)}, & \alpha = x. \end{cases} \tag{3.2d}$$

where $\mu(x)$ is the dipole density in the domain, and is initially unknown. Eq. 3.2a expresses Φ 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 μ, the limit as $(x, y) \to (x, \eta(x))$ is taken, yielding the relationships between φ and μ given in Eq. 3.2c. The operator in Eq. 3.2c, which is linear in μ, is then expressed in matrix form and inverted to obtain μ from a given φ. Once $\mu(x)$ is known, the partial derivatives of Φ are obtained by differentiating and evaluating Eq. 3.2a on the surface. For more details, see Refs. [61, 62] and the references therein. In what follows, we generate different reductions of Eq. 3.1 using the Galerkin projection for infinite depth in §5 and for finite depth in §6. The reason for the different approaches is to highlight different techniques for approximating the DNO. In the infinite depth case, the relationship between μ, φ, and Φ in Eqs. 3.2 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. 3.1a to obtain the same information.

4. The Proper Orthogonal Decomposition. The Proper Orthogonal Decomposition (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. Suppose for simplicity that the data are drawn from a single solution of Eq. 3.1 for $t \in [0, T]$. Then the error of an M-mode POD representation can be defined as

$$\text{Error}_{\text{POD}} = \int_0^T \|\psi(x, t) - \hat{\Phi}^{(M)}(x, t)\|^2\, dt, \tag{4.1}$$

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

$$\phi^{(m)} = \arg\max_{\|\phi\|=1} \int_0^T |\langle \phi(x), \psi(x,t) - \hat{\Phi}^{(m-1)}(x,t)\rangle|^2 \, dt, \tag{4.2}$$

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 [24, 63].

The data, ψ, 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:

$$\mathbf{X} = \begin{bmatrix} | & | & & | \\ \vec{\psi}_1 & \vec{\psi}_2 & \cdots & \vec{\psi}_K \\ | & | & & | \end{bmatrix}, \qquad \vec{\psi}_k = \begin{bmatrix} \psi(x_1, t_k) \\ \psi(x_2, t_k) \\ \vdots \\ \psi(x_N, t_k) \end{bmatrix}. \tag{4.3}$$

Here x_n is the nth sampling point in space and t_k is the kth sampling point in time. If ψ is of dimension larger than one, the sampled points, ψ, 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 [57]. 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 $\vec{w} = (w_1, w_2, \ldots, w_N)$ for integration in space, and $\vec{\omega} = (\omega_1, \omega_2, \ldots, \omega_K)$ for integration in time. Note that different integration schemes can be used to integrate in space and time.

To obtain the POD modes

$$\begin{bmatrix} | & | & & | \\ \phi^{(1)} & \phi^{(2)} & \cdots & \phi^{(N)} \\ | & | & & | \end{bmatrix} = \mathrm{diag}\left(\sqrt{\vec{w}}\right)^{-1} \mathbf{U}, \tag{4.4}$$

we apply the singular value decomposition

$$\mathrm{diag}\left(\sqrt{\vec{w}}\right) \mathbf{X} \, \mathrm{diag}\left(\sqrt{\vec{\omega}}\right) = \mathbf{U\Sigma V}^T. \tag{4.5}$$

Here, for example, $\mathrm{diag}(\sqrt{\vec{w}})$ is the diagonal matrix with the square root of the elements of $\vec{w}$ along the diagonal. Note that the resulting POD modes, $\phi^{(1)}, \ldots, \phi^{(N)}$, 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,

$$\hat{\Phi}^{(M)}(x,t) = \sum_{m=1}^M a_m(t)\phi^{(m)}(x) = \sum_{m=1}^M \langle \psi(\cdot, t), \phi^{(m)}\rangle_w \phi^{(m)}(x), \tag{4.6}$$

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

$$E_M = \frac{\int_0^T \|\psi(x,t) - \hat{\Phi}^{(M)}(x,t)\|^2 \, dt}{\int_0^T \|\psi(x,t)\|^2 \, dt}. \tag{4.7}$$

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_{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" [24]. 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 [47] or as multiple independent sets of data [63]. 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_{tol}) fewer modes. To extend the POD to multiple data sets, we used a generalization known as sequential POD (SPOD) [13]. This technique generates a set of M modes that produces a relative error smaller than E_{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 [61], the phase condition is that the surface height at $t = 0$ must be invariant under translations in space by π (i.e. $\eta(x, 0) = \eta(x + \pi, 0)$), and the surface velocity potential must be invariant when translated in space by π and negated (i.e. $\varphi(x, 0) = -\varphi(x + \pi, 0)$). Equivalently, at $t = 0$ all the odd wavenumbers of η and all the even wavenumbers of φ must be zero. This phase condition is straightforward to implement in a Fourier representation. Suppose that

$$\eta(x, 0) = \sum_{j=-N/2}^{N/2-1} a_j e^{ijx} \quad \text{and} \quad \varphi(x, 0) = \sum_{j=-N/2}^{N/2-1} b_j e^{ijx}. \tag{4.8}$$

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 that reduce the computational cost by a factor of four [61].

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: $\phi_O^{(1)}, \phi_O^{(2)}, \ldots, \phi_O^{(M_O)}$. 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: $\phi_E^{(1)}, \phi_E^{(2)}, \ldots, \phi_E^{(M_E)}$. The basis we use is $M = M_E + M_O$ dimensional, and is the union of these two filtered bases: $\phi_O^{(1)}, \phi_O^{(2)}, \ldots, \phi_O^{(M_O)}, \phi_E^{(1)}, \phi_E^{(2)}, \ldots, \phi_E^{(M_E)}$. 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. 4.8. The elements of the basis are ordered by the fraction of the L^2 norm each individual element captures.

5. 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. [61].

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 §5.1, we show that replacing the velocity potential, ϕ, with the dipole density, μ, 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 §5.3, 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 §5.4, 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.

5.1. Problem Formulation. As written in Eq. 3.1, 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 [9, 61]. A number of other formulations exist that express the problem in other physical variables. Fenton and Rienecker [20] 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 [2] 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 [36] 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:

$$\eta(x, t) = \sum_{m=1}^{M_\eta} a_m(t)\eta^{(m)}(x), \qquad \mu(x, t) = \sum_{m=1}^{M_\mu} b_m(t)\mu^{(m)}(x), \tag{5.1}$$

where $a_m, b_m \in \mathbb{R}$ and M_η and M_μ are the number of modes in the expansion. Using Eq. 3.2c, the velocity potential on the surface can also be represented as the superposition of M_μ modes:

$$\varphi(x, t) = \sum_{m=1}^{M_\mu} b_m(t)\varphi^{(m)}(x, t), \tag{5.2a}$$

$$\varphi^{(m)}(x, t) = \frac{\mu(x)}{2} + \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_1(x, \alpha)\mu^{(m)}(\alpha) \, d\alpha, \tag{5.2b}$$

where $\tilde{K}_1$ was given in Eq. 3.2d, and has an implicit time dependence due to η. By taking

the time-derivative of Eq. 5.2b, the time-derivative of μ can be found by solving

$$\frac{\mu_t(x)}{2} + \frac{1}{2\pi} \int \tilde{K}_1(x,\alpha)\,\mu_t(\alpha)\,d\alpha = \varphi_t(x) - \frac{1}{2\pi} \int \tilde{K}_2(x,\alpha)\mu'(\alpha)\,d\alpha, \qquad (5.3)$$

$$\tilde{K}_2(x,\alpha) = \text{Re}\left\{ \frac{\eta_t(x) - \eta_t(\alpha)}{2} \cot\left(\frac{(x-\alpha) + i(\eta(x) - \eta(\alpha))}{2} \right) \right\},$$

where $\tilde{K}_2$ becomes $\frac{\eta_{xt}(x)}{1+\eta_x^2(x)}$ as $\alpha \to x$. Note that this equation is structurally identical to Eq. 3.2c, which is solved to obtain μ from φ as an intermediate step in applying the DNO operator in the η, φ formulation of [61]. Nevertheless, as explained below, there is a computational advantage to working with μ rather than φ in a low-dimensional setting.

Using the POD modes of η and μ as basis functions, a set of governing equations for the mode amplitudes may be obtained by taking the Galerkin projection of Eq. 3.1,

$$\partial_t a_i = \langle \eta^{(i)}, \Phi_y - \eta_x \Phi_x \rangle, \qquad (5.4a)$$

$$\partial_t b_i + \sum_j L_1^{(ij)}(\partial_t b_j) = \left\langle \mu^{(i)}, P\left[-\eta_x \Phi_x \Phi_y - \frac{1}{2}\Phi_x^2 + \frac{1}{2}\Phi_y^2 - g\eta \right] \right\rangle - \sum_j L_2^{(ij)} b_j, \quad (5.4b)$$

$$L_1^{(ij)} = \left\langle \mu^{(i)}(x), \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_1(x,\alpha)\mu^{(j)}(\alpha)\,d\alpha \right\rangle, \qquad (5.4c)$$

$$L_2^{(ij)} = \left\langle \mu^{(i)}(x), \frac{1}{2\pi} \int_0^{2\pi} \tilde{K}_2(x,\alpha)\partial_\alpha \mu^{(j)}(\alpha)\,d\alpha \right\rangle. \qquad (5.4d)$$

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 η and φ, solving Eq. 3.2c involves inverting a dense $N \times N$ matrix to compute μ from φ. On the other hand, computing φ from μ requires a single matrix-vector multiplication. Although the matrix L_1 must be inverted, the elements of L_1 are given by Eq. 5.4c, 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. 5.4 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 $N = 512$.

5.2. POD Modes. Figure 5.1 shows a bifurcation diagram of infinite depth standing water waves, computed with the technique developed in Ref. [61]. 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 5.1 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, η, and the dipole density, μ, taken

Figure 5.1. *(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.*

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. [13], 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 $\phi^{(1)}, \phi^{(2)}, \ldots, \phi^{(m_1)}$, 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, $\tilde{\phi}^{(1)}, \tilde{\phi}^{(2)}, \ldots, \tilde{\phi}^{(m_2)}, \ldots, \tilde{\phi}^{(M)}$, 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 $\phi^{(1)}, \ldots, \phi^{(m_1)}, \tilde{\phi}^{(1)}, \ldots \tilde{\phi}^{(m_2)}$, 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
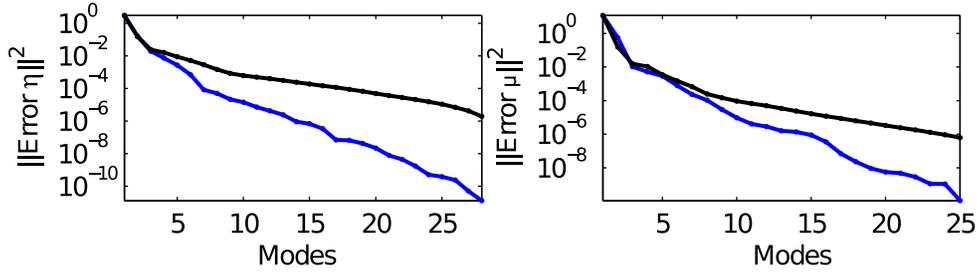
Figure 5.2. *Convergence of the POD modes (blue line) and the Fourier modes (black line) for sharply peaked solution data. The left plot shows the convergence for the fluid surface height, and the right plot shows the convergence for the dipole density in the fluid. Here the POD modes converge much more rapidly than the Fourier modes due to the large wavenumber components of the data.*

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 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. 5.1. 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 η and μ were drawn from the solution with $T \approx 6.425$. The remaining ten modes for η and eight modes for μ 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 5.2, 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. 5.1. 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 η and μ are significantly more accurate than their Fourier counterparts. For example, the POD expansion is more 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. 5.2 are specific to the

most sharply crested data set. If the first data set were studied instead, one would observe rapid convergence initially (as the POD modes will be drawn from that set of data) followed by near stagnation once the specified level of tolerance 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.

5.3. 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. 5.4. To compute these periodic solutions, we applied the shooting approach in Ref. [61] to the reduced order system and equated solving the two-point boundary value problem with minimizing the function

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

(5.5)

where $a_m(t)$ and $b_m(t)$ are the mode amplitudes in Eq. 5.1 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) [41] was used to numerically obtain a solution. Due to the symmetry properties we enforced in the POD modes, as discussed in §4, only roughly half of the modes need to vary. 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. 3.1 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. [61] 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 5.3 shows the error between the full PDE solution and the reduced solution 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

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

(5.6)

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
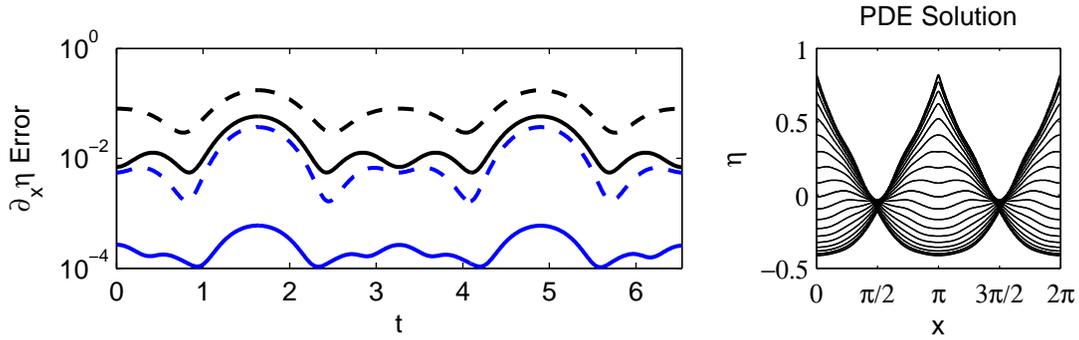
Figure 5.3. *Plot of the error in the derivative of a 10-mode Fourier expansion (black-dashed), 10-mode POD expansion (blue-dashed), 22-mode Fourier expansion (black-solid), and 22-mode POD expansion (blue-solid) for the solution plotted on the right. The POD expansions are roughly one to two orders of magnitude more accurate than the equivalently sized Fourier representation at all points along the orbit.*

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 5.3 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 5.3, the 22 POD-mode expansion appears graphically identical to the PDE solution.

Figure 5.4 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.
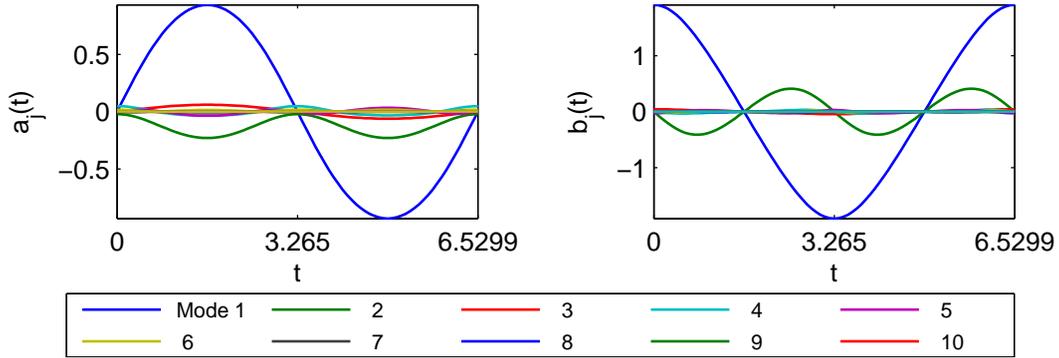
Figure 5.4. *Plot of the POD mode amplitudes, which are proportional to the amount of L^2-norm each mode captures, for the evolution in Figure 5.3. The oscillatory nature of the mode amplitudes demonstrates that, although the energy in a particular POD mode decays exponentially with the mode number, 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 η vanishes at $t = 3.265$ while the fourth mode does not.*
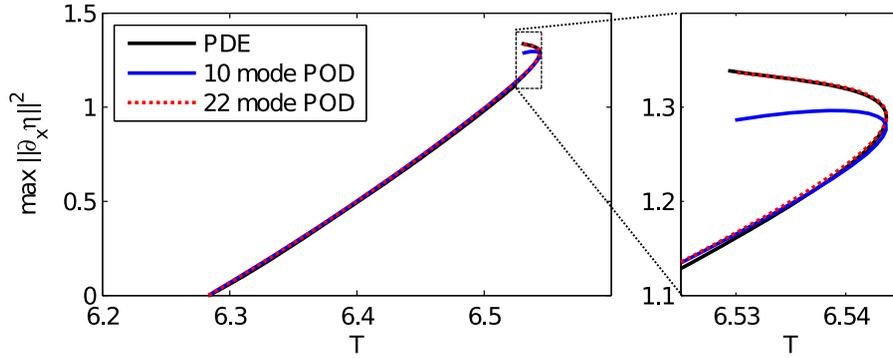
Figure 5.5. *Bifurcation diagram of the time-periodic standing wave solution plotted in terms of the period, T, and the maximum L^2 norm squared of the derivative of η over a single period. The blue curve is the bifurcation diagram obtained with 10-POD modes. The red curve is the bifurcation diagram 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.*

5.4. 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 5.1. 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 5.5 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 $T < 6.5$,
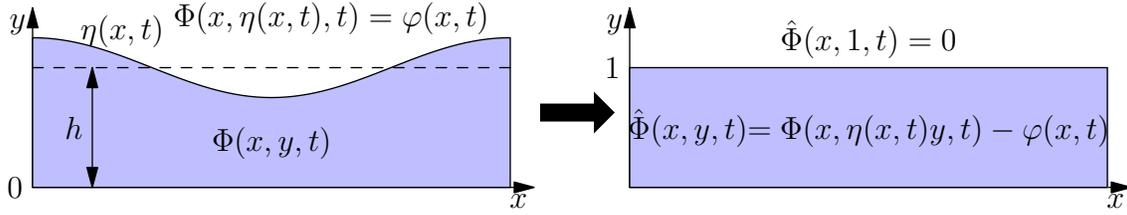
Figure 6.1. *The physical domain $\Omega(t) = \{(x, y) : 0 \le x < 2\pi, 0 < y < \eta(x, t)\}$ is transformed to $\hat{\Omega} = [0, 2\pi] \times [0, 1]$ under the change of variables $\hat{x} = x$, $\hat{y} = y/\eta(x, t)$. The problem is reformulated in terms of $\hat{\Phi}(x, y) = \Phi(x, \eta(x, t)y) - \varphi(x)$, which gives the boundary conditions $\partial_y \hat{\Phi}(x, y)|_{y=0} = 0$ and $\hat{\Phi}(x, 1) = 0$.*

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 5.5 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 5.4, 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. 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. 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 §5 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.

6.1. Problem Formulation. In the finite depth case, we adopt a direct approach in order to solve Eq. 3.1. 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 [39]. 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.

Instead, we map the fluid to a domain with constant depth by adopting the sigma coordinate system in fluid mechanics [46] also known as the c-method in electromagnetics [6]. 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 Figure 6.1, 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. 3.1a.

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 Φ 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 Φ on the surface span the same space as the basis for φ. 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:

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

(6.1)

where h is the average fluid depth and $\eta^{(m)}$, $\varphi^{(m)}$, and $\hat{\Phi}^{(m)}$ are the mth 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 §6.2. For now, we only require that the individual modes are orthogonal to one another in their respective inner product spaces, e.g.

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

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:

$$\sum_{j=1}^{M_\Phi} L_{mj} B_j(t) = -\langle \hat{\Phi}^{(m)}, \varphi_{xx} \rangle_R, \qquad \text{for } m = 1, 2, \ldots, M_\Phi, \tag{6.2a}$$

$$\partial_t a_m(t) = \left\langle \eta^{(m)}, \frac{\hat{\Phi}_y}{\eta} - \eta_x \left(\varphi_x - \frac{\eta_x}{\eta} \hat{\Phi}_y \right) \right\rangle, \qquad \text{for } m = 1, 2, \ldots, M_\eta, \tag{6.2b}$$

$$\partial_t b_m(t) = \left\langle \varphi^{(m)}, -\frac{\eta_x}{\eta} \hat{\Phi}_y \left(\varphi_x - \frac{\eta_x}{\eta} \hat{\Phi}_y \right) - \frac{1}{2} \left(\varphi_x - \frac{\eta_x}{\eta} \hat{\Phi}_y \right)^2 + \frac{1}{2} \left(\frac{\hat{\Phi}_y}{\eta} \right)^2 - g\eta \right\rangle, \tag{6.2c}$$

with

$$L_{mj} = \left\langle \hat{\Phi}^{(m)}, \hat{\Phi}_{xx}^{(j)} + \left(2y \frac{\eta_x^2}{\eta^2} - y \frac{\eta_{xx}}{\eta} \right) \hat{\Phi}_y^{(j)} + \left(y^2 \frac{\eta_x^2}{\eta^2} + \frac{1}{\eta^2} \right) \hat{\Phi}_{yy}^{(j)} - 2y \frac{\eta_x}{\eta} \hat{\Phi}_{xy}^{(j)} \right\rangle_R, \tag{6.3}$$

where superscripted variables represent individual modes, and variables without superscripts represent the current approximation of the solution in Eq. 6.1. 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. 6.1. Equation 6.2a is a transformed version of Laplace's equation in Eq. 3.1a. Because the transformation used is not conformal, Laplace's equation assumes an outwardly more complicated form; however, Eq. 6.2a still enforces incompressibility in the fluid itself. Similarly, Eq. 6.2b is a transformed version of Eq. 3.1d, and Eq. 6.2c is a transformed version of Eq. 3.1e. Because the POD modes taken from a valid solution of the water wave equations must satisfy the boundary conditions (Eq. 3.1c and Eq. 3.1d), those equations 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.

6.2. Proper Orthogonal Decomposition Modes. As stated in the previous subsection, we represent η, φ, 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 [57]. Therefore, a snapshot of η and φ 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 η, φ, and Φ on these grids, Eq. 3.1 was evolved using a suitably modified version of the boundary integral method, similar to Eq. 3.2. While it is not necessary to explicitly form Φ to evolve η and φ using the boundary integral formulation, it is easy to evaluate Φ on the grid as an extra step, to generate the POD data.

The individual snapshots are collected from periodic orbits taken with four different periods. Figure 6.2 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 Figure 6.2 shows the ten modes for η and for φ obtained using
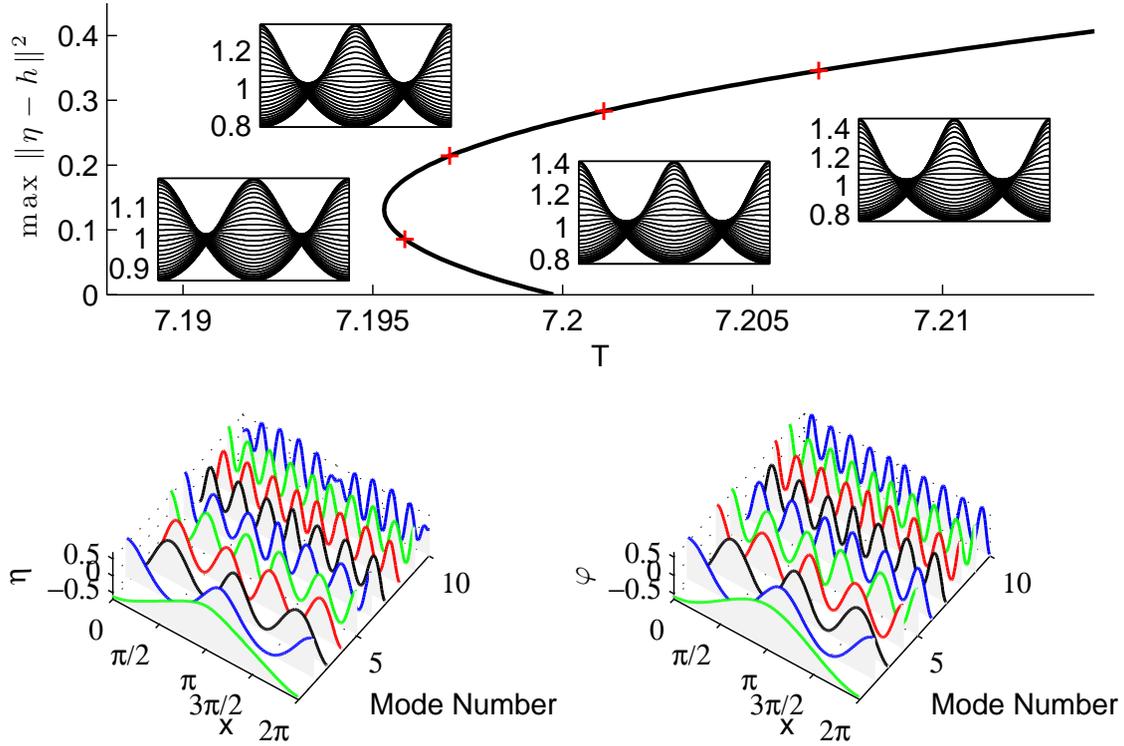
Figure 6.2. *(Top) Bifurcation diagram showing L^2-norm vs. period of the standing wave solutions in fluid of unit depth. The data used to generate POD modes were taken from locations marked by a red +. Near each data point, we plot the surface height evolution over a period. (Bottom) First ten POD modes 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.*

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 Figure 6.3. 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 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
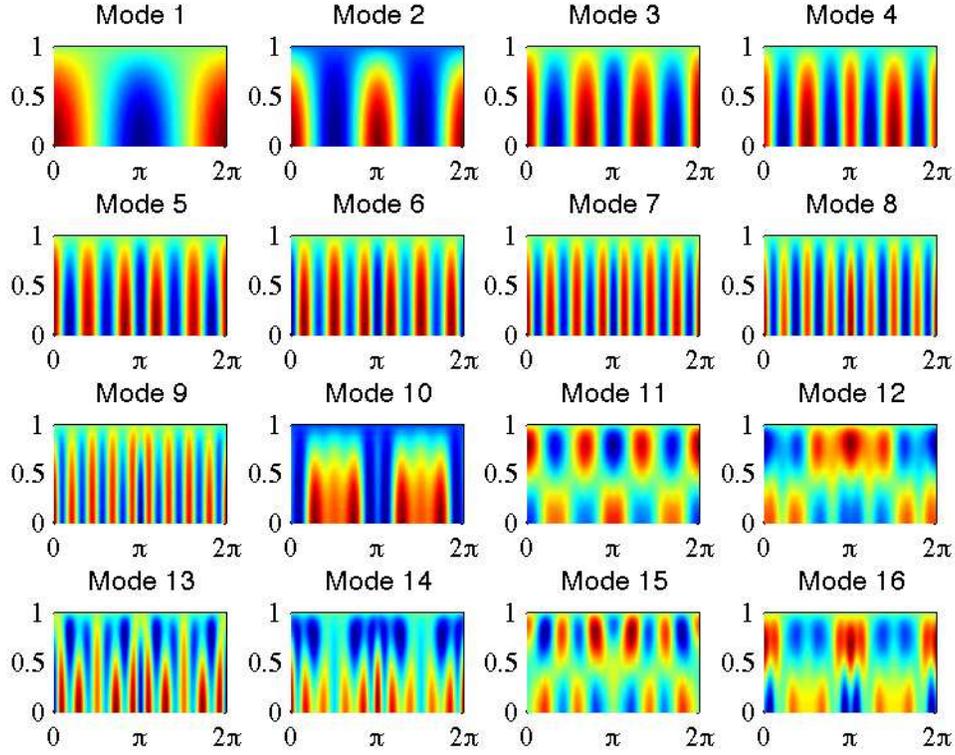
Figure 6.3. *Plots of the POD modes of $\hat{\Phi}(x, y)$ on the transformed rectangular domain. The modes are ordered from left-to-right and from top-to-bottom.*

combination of φ and η 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 [24]. 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 returns, yields a more globally applicable basis. In Fig. 6.4 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. 6.2. The data sets were ordered such that data sets with more extreme waveforms were considered later, i.e. from bottom up in Fig. 6.2. For η, the data sets contribute 3, 5, 4, and 2 POD modes respectively. For φ, the data sets contribute 3, 6, 4, and 4 POD modes respectively.

Figure 6.4 shows the convergence of the POD representation to data taken from the 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
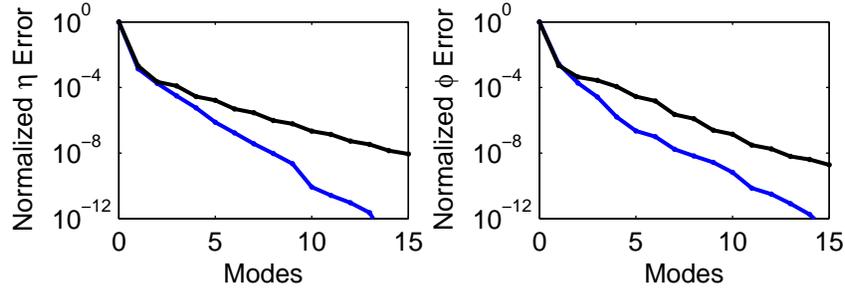
Figure 6.4. *Convergence of L^2-norm of the POD representation (blue line) and the Fourier representation (black line) of the rightmost data point (denoted by a red cross, with $T \approx 7.2075$) in Fig. 6.2. The POD modes were generated using data from the four solutions indicated in that figure. The tolerance set for the SPOD technique here a 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.*
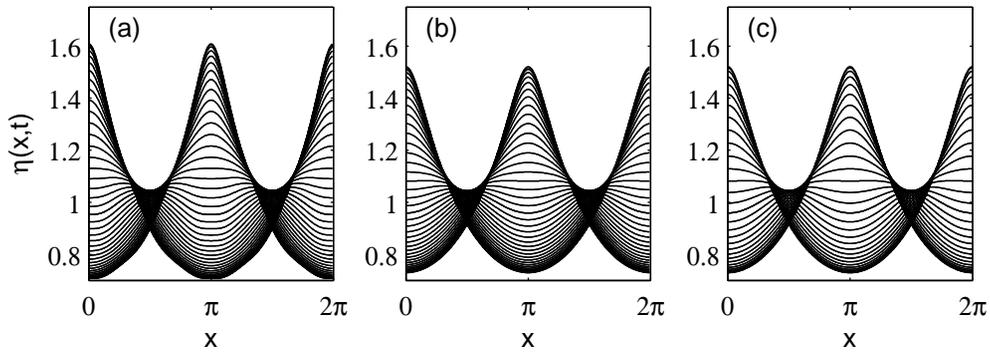
Figure 6.5. *Plot of the evolution of the fluid surface using (a) 6-POD modes (b) 10-POD modes and (c) 512 Fourier modes. With the 6-mode expansion, the surface evolution can be qualitatively reproduced although there are significant quantitative differences. However, with 10-POD modes there is good quantitative agreement between the low-dimensional model and the full 512-Fourier mode PDE solution*

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.

6.3. POD Dynamics. In this subsection, we consider POD approximations of different dimensions to a single periodic orbit of the PDE. Similar to §5.3, the solution here is a numerically computed periodic solution to the reduced problem with $F < 10^{-20}$, as defined in Eq. 5.5. To highlight the differences in the expansions, we present results from the orbit of largest period in Figure 6.2, but it should be noted that these results are prototypical of other orbits on the branch of solutions.
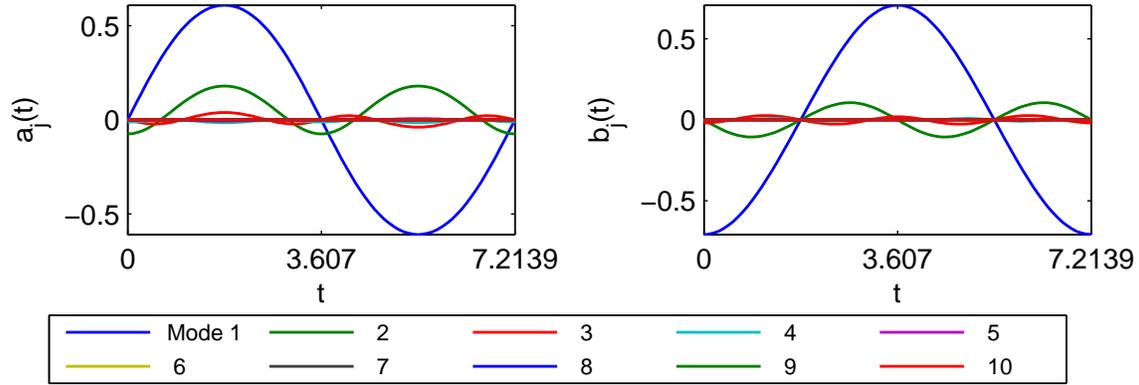
Figure 6.6. *Plot of the POD mode amplitudes for the evolution in Figure 6.5, which demonstrates that the maximum value of the mode amplitude decreases as the mode number and frequency of oscillation increase.*

Figure 6.5 shows the POD reproduction of the solution with largest period in Figure 6.2 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 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 6.6 shows the evolution of the mode-amplitudes of the solution in Figure 6.5 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 Figure 6.6 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 Figure 6.6, 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. 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 the entire branch of time periodic standing wave solutions using either the first six or all ten of the POD modes in Figure 6.2 for the surface variables, and the first ten or all sixteen of the interior modes in Figure 6.3.
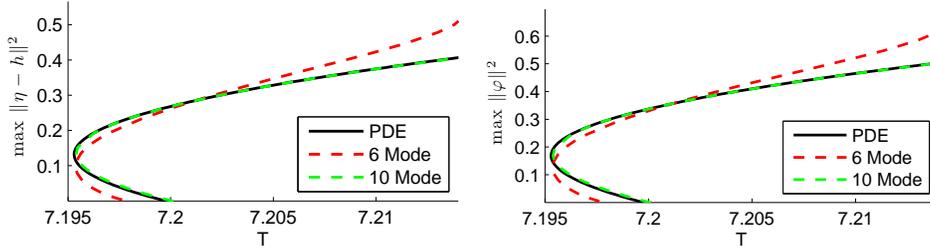
Figure 6.7. *Bifurcation diagrams of the time-periodic standing wave solutions in terms of the period, T, and the maximum L^2-norm over one period. The PDE solution branch is shown in solid black. The 6-mode POD expansion is shown in dashed red lines, and the 10-mode POD expansion is shown in dashed green lines.*

Figure 6.7 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 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.

7. Conclusions. In this manuscript, 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 demon-

strating 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 Φ 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 π. 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 [9] or using a conformal mapping [39] 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 [23, 31]. 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 [61, 62]. In those papers, 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 [62].

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.

In this manuscript, we have 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 [61], 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. 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 [36, 37, 61, 62], 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.

REFERENCES.

[1] C. J. AMICK, L. E. FRAENKEL, AND J. F. TOLAND, *On the Stokes conjecture for the wave of extreme form*, Acta Math., 148 (1982), pp. 192–214.

[2] G. R. BAKER, D. I. MEIRON, AND S. A. ORSZAG, *Generalized vortex methods for free-surface flow problems*, J. Fluid Mech., 123 (1982), pp. 477–501.

[3] T. J. BRIDGES, *Wave breaking and the surface velocity field for three-dimensional water waves*, Nonlinearity, 22 (2009), pp. 947–953.

[4] D. S. BROOMHEAD AND G. P. KING, *Extracting Qualitative Dynamics from Experimental Data*, Physica D, 20 (1986), pp. 217–236.

[5] P. J. BRYANT AND M. STIASSNIE, *Different forms for nonlinear standing waves in deep water*, J. Fluid Mech., 272 (1994), pp. 135–156.

[6] J. CHANDEZON, D. MAYSTRE, AND G. RAOULT, *A new theoretical method for diffraction gratings and its numerical application*, J. Optics, 11 (1980), pp. 235–241.

[7] E. A. CHRISTENSEN, M. BRONS, AND J. N. SORENSEN, *Evaluation of Proper Orthogonal Decomposition-based decomposition techniques applied to parameter dependent non-turbulent flows*, SIAM J. Sci. Comput., 21 (2000), pp. 1419–1434.

[8] P. CONCUS, *Standing capillary-gravity waves of finite amplitude*, J. Fluid Mech., 14 (1962), pp. 568–576.

[9] W. CRAIG AND C. SULEM, *Numerical simulation of gravity waves*, J. Comp. Phys., 108 (1993), pp. 73–83.

[10] A. D. D. CRAIK, *George Gabriel Stokes on water wave theory*, Ann. Rev. Fluid Mech., 37 (2005), pp. 23–42.

[11] M. C. CROSS AND P. C. HOHENBERG, *Pattern formation outside of equilibrium*, Rev. Mod. Phys., 65 (1993), p. 851.

[12] J. CRUZ, *Ocean Wave Energy: Current Status and Future Perspectives*, (2008).

[13] 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.

[14] F. DIAS AND T. J. BRIDGES, *The numerical computation of freely propagating time-dependent irrotational water waves*, Fluid Dynamics Research, 38 (2006), p. 803.

[15] E. DING, E. SHLIZERMAN, AND J. N. KUTZ, *Modeling multipulsing transition in ring cavity lasers with proper orthogonal decomposition*, Physical Review A, 82 (2010), p. 23823.

[16] P. DRUAULT AND C. CHAILLOU, *Use of Proper Orthogonal Decomposition for reconstructing the 3D in-cylinder mean-flow field from PIV data*, Comptes Rendus Mécanique, 335 (2007), pp. 42–47.

[17] A. DYACHENKO AND V. ZAKHAROV, *Modulation instability of Stokes wave → freak wave*, JETP Letters, 81 (2005), pp. 255–259.

[18] A. L. DYCHENKO, V. E. ZAKHAROV, AND E. A. KUZNETSOV, *Nonlinear dynamics on the free surface of an ideal fluid*, Plasma Phys. Rep., 22 (1996), pp. 916–928.

[19] J. EGGERS AND M. A. FONTELOS, *The role of self-similarity in singularities of partial differential equations*, Nonlinearity, 22 (2009), pp. R1—-R44.

[20] J. D. FENTON AND M. M. RIENECKER, *A Fourier method for solving nonlinear water-*

wave problems: application to solitary-wave interactions, Journal of Fluid Mechanics, 118 (1982), pp. 411–443.

[21] G. HALLER, *Chaos Near Resonance*, Applied Mathematical Sciences; 138, Springer-Verlag, NY, 1999.

[22] J. L. HAMMACK, *A note on tsunamis: their generation and propagation in an ocean of uniform depth*, J. Fluid Mech., 60 (1973), pp. 769–799.

[23] 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.

[24] P. HOLMES, J. LUMLEY, AND G. BERKOOZ, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, 1996.

[25] M. ILAK AND C. W. ROWLEY, *Reduced-order modeling of channel flow using traveling POD and balanced POD*, in 3rd AIAA Flow Control Conference, 2006.

[26] M. ILAK AND C. W. ROWLEY, *Modeling of transitional channel flow using balanced proper orthogonal decomposition*, Physics of Fluids, 20 (2008), p. 034103.

[27] G. IOOSS, P. I. PLOTNIKOV, AND J. F. TOLAND, *Standing waves on an infinitely deep perfect fluid under gravity*, Arch. Rat. Mech. Anal., 177 (2005), pp. 367–478.

[28] L. JIANG, C. TING, M. PERLIN, AND W. W. SCHULTZ, *Moderate and steep Faraday waves: instabilities, modulation and temporal asymmetries*, J. Fluid Mech., 329 (1996), pp. 275–307.

[29] R. S. JOHNSON, *A modern introduction to the mathematical theory of water waves*, Cambridge University Press, Cambridge, UK, 1997.

[30] Y. KERVELLA, D. DUTYKH, AND F. DIAS, *Comparison between three-dimensional linear and nonlinear tsunami generation models*, Theor. Comput. Fluid Dyn., 21 (2007), pp. 245–269.

[31] K. KUNISCH AND S. VOLKWEIN, *Galerkin proper orthogonal decomposition methods for parabolic problems*, Numer. Math., 90, (2001), pp. 117–148.

[32] H. LAMB, *Hydrodynamics*, University Press, Cambridge, 1895.

[33] J. A. LEE AND M. VERLEYSEN, *Nonlinear Dimensionality Reduction*, Springer, 2007.

[34] 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, 453 (1997), pp. 663–672.

[35] M. S. LONGUET-HIGGINS AND M. J. H. FOX, *Theory of the almost-highest wave: the inner solution*, J. Fluid Mech., 80 (1977), pp. 721–741.

[36] G. N. MERCER AND A. J. ROBERTS, *Standing waves in deep water: Their stability and extreme form*, Phys. Fluids A, 4 (1992), pp. 259–269.

[37] G. N. MERCER AND A. J. ROBERTS, *The form of standing waves on finite depth water*, Wave Motion, 19 (1994), pp. 233–244.

[38] J. W. MILES, *On Hamilton's principle for surface waves*, Journal of Fluid Mechanics, 83 (1977), pp. 153–158.

[39] P. A. MILEWSKI, J.-M. VANDEN-BROECK, AND Z. WANG, *Dynamics of steep two-dimensional gravity capillary solitary waves*, Journal of Fluid Mechanics, 664 (2010),

pp. 466–477.

[40] D. NICHOLLS, *Boundary Perturbation Methods for Water Waves*, GAMM-Mitteilungen, 30 (2007), pp. 44–74.

[41] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer, New York, 1999.

[42] M. OKAMURA, *Standing gravity waves of large amplitude on deep water*, Wave Motion, 37 (2003), pp. 173–182.

[43] M. OKAMURA, *Almost limiting short-crested gravity waves in deep water*, J. Fluid Mech., 646 (2010), pp. 481–503.

[44] W. G. PENNEY AND A. T. PRICE, *Finite periodic stationary gravity waves in a perfect liquid, Part {II}*, Phil. Trans. R. Soc. London A, 244 (1952), pp. 254–284.

[45] N. PÉRINET, D. JURIC, AND L. S. TUCKERMAN, *Numerical simulation of Faraday waves*, J. Fluid Mech., 635 (2009), pp. 1–26.

[46] N. PHILLIPS, *A coordinate system having some special advantages for numerical forecasting*, J. Meteor., 14 (1957), pp. 184–185.

[47] B. PODVIN AND P. LE QUÉRÉ, *Low-order models for flow in a differentially heated cavity*, Physics of Fluids, 13 (2011), pp. 3204–3214.

[48] J. RAYLEIGH, *Deep water waves, progressive or stationary, to the third order approximation*, Proc. R. Soc. A, 91 (1915), pp. 345–353.

[49] P. J. SCHMID, K. .E. MEYE, AND O. PUST, *Dynamic mode decomposition and proper orthogonal decomposition of flow in a lid-driven cylindrical cavity*, in Proc. of the 8th Intl. Symp. on Particle Image Velocimetry, 2009.

[50] W. SCHULTZ, J. M. VANDEN-BROECK, L. JIANG, AND M. PERLIN, *Highly nonlinear standing water waves with small capillary effect*, J. Fluid Mech., 369 (1998), pp. 253–272.

[51] L. SIROVICH, *Turbulence and the dynamics of coherent structures part I: coherent structures**, Quart. App. Math., 45 (1987), pp. 561–571.

[52] L. SIROVICH, *Modeling the functional organization of the visual cortex*, Physica D: Nonlinear Phenomena, 96 (1996), pp. 355–366.

[53] D. H. SMITH AND A. J. ROBERTS, *Branching behavior of standing waves — the signatures of resonance*, Phys. Fluids, 11 (1999), pp. 1051–1064.

[54] 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, vol. 1, Cambridge University Press, 1880, pp. 225–228.

[55] I. TADJBAKHSH AND J. B. KELLER, *Standing surface waves of finite amplitude*, J. Fluid Mech., 8 (1960), pp. 442–451.

[56] G. I. TAYLOR, *An experimental study of standing waves*, Proc. Roy. Soc. A, 218 (1953), pp. 44–59.

[57] L. TREFETHEN, *Spectral methods in MatLab*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.

[58] C. P. TSAI AND D. S. JENG, *Numerical Fourier solutions of standing waves in finite water depth*, Appl. Ocean Res., 16 (1994), pp. 185–193.

[59] 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, 103 (2009), 124501.

[60] J. M. VANDEN-BROECK, *Nonlinear gravity-capillary standing waves in water of arbitrary*

uniform depth, J. Fluid Mech., 139 (1984), pp. 97–104.

[61] J. Wilkening, *Breakdown of self-similarity at the crests of large amplitude standing water waves*, Phys. Rev. Lett, 107 (2011), 184501.

[62] J. Wilkening and J. Yu, *Overdetermined shooting methods for computing time-periodic water waves with spectral accuracy.*

[63] M. O. Williams, E. Shlizerman, and J. N. Kutz, *The multi-pulsing transition in mode-locked lasers: a low-dimensional approach using waveguide arrays*, J. Opt. Soc. Am. B, 27 (2010), pp. 2471–2481.

[64] V. E. Zakharov, *Stability of periodic waves on finite amplitude on the surface of a deep fluid*, Zhurnal Prildadnoi Mekhaniki i Tekhnicheskoi Fiziki, 9 (1968), pp. 86–94.