

Multidomain spectral method for the helically reduced wave equation

Stephen R. Lau ^{a,b,*}, Richard H. Price ^c

^a Department of Mathematics and Statistics, University of New Mexico, Albuquerque, NM 87111, USA

^b Division of Applied Mathematics, Brown University, Providence, RI 02912, USA

^c Center for Gravitational Wave Astronomy, Department of Physics and Astronomy, University of Texas at Brownsville, Brownsville, TX 78520, USA

Received 6 February 2007; received in revised form 22 August 2007; accepted 27 August 2007

Available online 14 September 2007

Abstract

We consider the 2+1 and 3+1 scalar wave equations reduced via a helical Killing field, respectively referred to as the 2-dimensional and 3-dimensional helically reduced wave equation (HRWE). The HRWE serves as the fundamental model for the mixed-type PDE arising in the periodic standing wave (PSW) approximation to binary inspiral. We present a method for solving the equation based on domain decomposition and spectral approximation. Beyond describing such a numerical method for solving strictly linear HRWE, we also present results for a nonlinear scalar model of binary inspiral. The PSW approximation has already been theoretically and numerically studied in the context of the post-Minkowskian gravitational field, with numerical simulations carried out via the “eigenspectral method.” Despite its name, the eigenspectral technique does feature a finite-difference component, and is lower-order accurate. We intend to apply the numerical method described here to the theoretically well-developed post-Minkowski PSW formalism with the twin goals of spectral accuracy and the coordinate flexibility afforded by global spectral interpolation.

© 2007 Elsevier Inc. All rights reserved.

MSC: 65M70; 35M10; 5L20; 83C35; 35L05

Keywords: Helical symmetry; Spectral methods; Gravitational waves; Mixed PDE

1. Introduction and preliminaries

1.1. Periodic standing wave (PSW) approximation

The development of gravitational wave detectors has spurred interest in the binary inspiral of compact astrophysical objects, in particular black holes. The challenge of computationally solving Einstein’s equations

* Corresponding author. Address: Division of Applied Mathematics, Brown University, Providence, RI 02912, USA. Tel.: +1 401 863 3058; fax: +1 401 863 1355.

E-mail addresses: lau@dam.brown.edu (S.R. Lau), rprice@phys.utb.edu (R.H. Price).

for such a system has become the focus of many groups throughout the world both for its importance in gravitational wave astronomy and for its role in advancing the understanding of highly dynamical strongly curved spacetimes. Very recently [1–7] a number of techniques have been found to stabilize codes, so that computational evolution of Einstein’s equations can track the last few orbits of binary inspiral.

There is, of course, more to the problem than the last few orbits. For orbiting objects with mass M and separation a , the characteristic measure of the nonlinearity of their gravitational interaction is GM/ac^2 where G , c are the gravitational constant and the speed of light. When this measure is small, the dynamics and the generation of gravitational waves can be found using post-Newtonian methods, an analytic approach in which Einstein’s theory is, in effect, expanded treating GM/ac^2 as a perturbation parameter.

With methods available for the small- GM/ac^2 early stage, and the large- GM/ac^2 last few orbits, what remains needed is an effective method for treating the intermediate epoch, the stage of inspiral in which GM/ac^2 is too large for post-Newtonian methods, but in which many orbits remain. When more than a few orbits still remain, an accurate numerical evolution of Einstein’s equations will be too computationally expensive, at least in the near future.

The periodic standing wave (PSW) approach is an approximation scheme for dealing effectively with this intermediate inspiral epoch. The scheme is based on the fact that during this epoch the orbiting inspiral is very much orbit, and only slightly inspiral. That is, the change in the orbital radius is small for each orbit. (This is, in fact, a criterion for many orbits to remain, and for accurate computational evolution to be daunting.) In the PSW approach the motion of the sources and the fields are assumed to be helically symmetric, that is, all quantities are rigidly rotating in the sense that a change in time by Δt is equivalent to a change in azimuthal angle φ according to $\Delta\varphi \rightarrow -\Omega\Delta t$, where Ω is a constant, the angular velocity with which the fields rigidly rotate.

The imposition of this helical symmetry vastly changes the nature of the mathematical and computational problem. Prior to helical reduction, the problem is that of evolving forward in time a hyperbolic problem (more precisely, a problem that can be cast in hyperbolic form). The imposition of helical symmetry reduces by one the number of independent variables and, more important, changes the problem from a hyperbolic problem to a mixed one, a problem with a region of the manifold (near the rotation axis) in which the equations are elliptic, and an outer region in which the equations are hyperbolic. The boundary conditions for this problem are also unusual. One can have the presence of a source represented by inner Dirichlet boundary conditions on two small topological spheres just outside the location of the sources, the compact astrophysical objects. Alternatively, one can include the objects as explicit inhomogeneities in the equations. The other boundary conditions on the problem represent the radiation in the distant wave zone, and require some discussion.

An important feature of this problem in general relativity is the conservation of the total energy of the system. If energy is leaving in the form of outgoing gravitational radiation, then the orbital motion cannot be helically symmetric; the radius must decrease. For fields other than gravity one could invoke a force, some *deus ex machina*, to keep the orbits unchanging, and have that force not couple to the field being studied. This certainly can be, and has been done in model problems [8–12] but, in principle, cannot be done in general relativity. In Einstein’s gravitation all forces couple to gravity. This is dealt with by computing a standing wave solution of the helically symmetric problem, that is, by imposing standing wave outer boundary conditions. From the standing wave exact solution, an approximate solution is then extracted to the physical problem of outgoing waves.

By “standing waves” in a *linear* theory, we mean here the average of the outgoing wave solution and the ingoing wave solution. No clear meaning exists for standing waves in a nonlinear theory. To define our standing waves, we choose an iterative solution technique for the nonlinear problem in which an average of ingoing and outgoing solutions is taken at each step of iteration.¹ We emphasize here that in each iteration of the PSW problem it is an outgoing (or, trivially different, ingoing) problem that is solved. The present paper will therefore focus on the details of the efficient computation of a problem with outgoing radiation boundary conditions.

¹ Another extension of “standing wave” to nonlinear theories has been discussed in the literature [9], using the minimum amplitude of the multipole in every multipole mode.

Hyperbolic problems with Dirichlet boundary conditions are known not to be well posed. We take a pragmatic approach to whether our mixed-type problems with radiative conditions are well posed. For one thing, the problem arises in what would appear to be a physically well-specified problem; heuristically the problems of nonuniqueness for the Dirichlet problem are removed by the radiation boundary conditions. For another thing, no fundamental instability has been encountered in seeking numerical solutions of our problems. Further indirect evidence that the problem is well posed can be found in the work of Torre, who has shown that a closely related mixed-type linear problem is well-posed if the boundary conditions are of an admissible type that includes Sommerfeld outgoing conditions [13,14].

The strategy of the PSW approach is to solve the helically symmetric binary problem computationally, but otherwise without approximation. From that “exact” helically symmetric solution, an approximation is extracted for the physical problem. The extraction of an outgoing approximation is tantamount to treating the nonlinear standing wave solution as if it were an average of the ingoing and outgoing solutions. (For details of extraction, see [12].) The reason that this is an excellent approximation (as is confirmed by computations with model problems) is that the regions of the physical manifold in which nonlinearities are strong are those near the sources, and in these regions the solution is extremely insensitive to the boundary conditions (ingoing, outgoing, standing wave) in the distant weak field wave region. Since the ingoing and outgoing solutions are nearly identical in this region they can be averaged (or simply replaced by either the ingoing or the outgoing solution). Where the ingoing/outgoing solutions are very different, the theory is approximately linear and hence again the ingoing and outgoing fields can be averaged. This feature, the separation of the region of nonlinearity and the region of waves, is closely related to the argument, given below, that a multidomain spectral method should have advantages for PSW type problems even beyond the advantages it has demonstrated for purely elliptic initial-value problems in the work of Pfeiffer et al. [15,16], or its use in evolution codes by the Caltech–Cornell collaboration [17].

Successful solutions of the PSW approximation will serve a variety of purposes. As discussed above, it will provide the bridge between the post-Newtonian methods and numerical evolution for binary orbits; it will provide near optimal starting points for the numerical evolution of the last few orbits; it may give a useful testbed for studying radiation reaction; solution of the PSW standing wave problem will give a new class of solutions to Einstein’s theory.

Work on PSW computations has already taken several steps using the “eigenspectral method,” an approximate numerical method based on coordinates well adapted to the geometry both close to the sources and in the radiation region. This method, in effect, keeps only the features of the solution that are most important to the structure of the near-source fields and to the radiation, and has been used so far to study nonlinear toy models [8,9,12], linearized general relativity [18], and the post-Minkowski extension of linearized gravity [19]. This method is remarkable for its simplicity in giving approximate solutions, but is ill suited to the high accuracy needed for several purposes. In addition the eigenspectral technique uses a finite-difference component, and provides numerical answers on a specific grid of points in the space of the physical problem. Its results, therefore, require a difficult interpolation from a nonuniform grid if they are to be compared with other results or used as initial conditions for evolution codes. The multidomain spectral method we will describe has neither shortcoming.

To minimize the complexity of issues not directly relevant to the multidomain spectral method, we choose as a specific target problem, more-or-less the problem of Refs. [8,9,12], that of a nonlinear scalar field. The physical problem, of course, is one with three spatial coordinates, but for simplicity both of exposition and of computation, we work here with the two-dimensional helically reduced wave equation, hereafter 2d HRWE. The Cauchy problem for this physical model would involve three independent variables (two spatial, one time), but the helical reduction means we are solving on a manifold of two dimensions. We will also, in Section 4, present some preliminary results for the 3d HRWE, and in the concluding section will discuss application of the method to the 3d HRWE problem. The physical picture is that of two compact scalar charges moving on circular orbits of radius a , with angular velocity Ω , and emitting helically symmetric outgoing waves. The wave speed c will be taken to be unity. The scalar field satisfies a 2d HRWE of the form

$$L\psi + \eta h(\psi) = g, \quad (1)$$

on the region outside the two sources and inside a large radius circle (spherical surface for 3d HRWE) on which outgoing conditions are imposed. Here L is a mixed-type linear HRWE operator (the helically reduced d'Alembertian), h is a nonlinear function of ψ , and the η parameter controls the strength of the nonlinearity. The inhomogeneity g can be used to represent explicit sources. In our 2d HRWE development below we will use inner boundary conditions, not explicit sources, so g will be zero; a nonzero g will be used only in the consideration of the 3d HRWE in the concluding section.

We make one more simplifying choice. We take the solution to have the symmetry $\psi(-x, -y) = -\psi(x, y)$, a choice equivalent to taking the sources to have opposite scalar charge. By setting the total charge equal to zero, we eliminate the bothersome logarithmic monopole that would diverge at large distances. Of course, for gravitation theory there is only charge of a single sign and the monopole cannot be set to zero, but the real physical problem is that for three spatial dimensions, in which the monopole falls off with distance from the sources, and is an acceptable complication.

1.2. Multidomain spectral method

The “two center” domain of the 2d HRWE problem is illustrated in Fig. 1. This figure shows how two annular domains, H and A along with eight rectangular domains, fully cover the physical problem. Note that the symmetry $\psi(-x, -y) = -\psi(x, y)$ means that annulus H represents the information around both the source on the right, and that on the left. Similarly, rectangular domains 1, 2, 3, 7, 8 carry the information about the solution in symmetry related regions on the left side of the physical problem. Using the aforementioned symmetry, we could do away with region 6, which is equivalent to region 4. Nevertheless, we have chosen to keep region 6 so that our code can be tested on elliptic problems which, when posed on the inner region spanned by the rectangles, need not be of definite parity. The annular domain A extends between the two solid circles in Fig. 1. At its outer boundary we impose Sommerfeld-like outgoing radiation boundary conditions to be described below. For most radiation conditions to be applicable, the outer boundary of A must be at least several wavelengths away from the sources. For a typical choice of Ω this means that the radius of the outer boundary must be 10 times or more larger than the separation between the source regions. The dashed circle in Fig. 1 is the “light circle,” the boundary between the inner region in which the operator L is elliptic, and the outer region in which L is hyperbolic. Typically the radius of the the light circle will be at least several times larger than the separation of the centers of the source domains. For a 3d problem, the analogous two-center

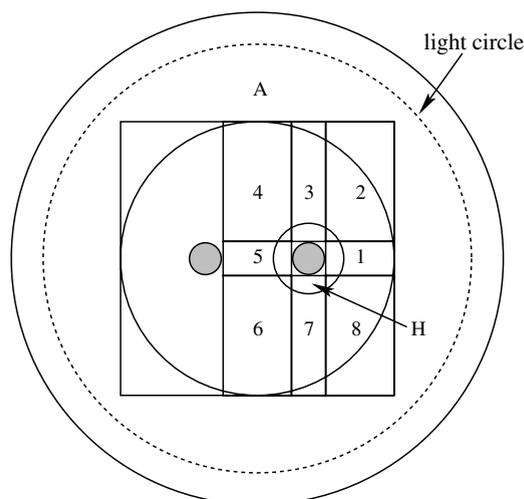


Fig. 1. The 2d HRWE problem and its domain decomposition. The circles that are the outer boundaries of the darkened source regions carry Dirichlet boundary conditions representing the imprint of the sources. Region H , between the two small circles concentric with the source, is the inner annular domain. Domain A between two large solid circles is the outer annular domain. The dashed circle in domain A is the light circle.

domain is actually composed of fewer elements than in the equivalent 2d problem. This is one of the reasons we believe that a multidomain spectral method is well-suited for the computations needed in PSW approximation of binary inspiral. In lieu of a detailed description of the domain decomposition we envision for 3d work, we offer the picture in Fig. 2. Note that cylindrical shell labelled 2, 8, for example, corresponds to both domain 2 and domain 8 in Fig. 1. Indeed, the 2d figure could alternatively be viewed as a cross section of a 3d scenario.

The 10 subdomains H , 1 – 8, A , along with the symmetry $\psi(-x, -y) = -\psi(x, y)$, contain the complete information about ψ . Below we describe a spectral method in which the solution for ψ in each region is considered as an expansion in terms of appropriate basis functions. Moreover, we use the term spectral in the truest sense, since the unknowns we solve for in our numerical approach are in fact the coefficients of the basis functions. (In a later section, while considering the 3d HRWE we will also describe results based in part on a pseudospectral method in which the unknowns are point-values on a spectral grid.) A crucial feature of our numerical method is to separate the physical problem into subdomains so that the type changing nature of the HRWE is then confined within a single subdomain, the outer annulus A . The inner region is then purely elliptical and, despite its two center topology, this inner part of the manifold is amenable to the standard spectral techniques associated with elliptic equations [15,16].

1.3. Brief review of mixed-type problems

In order to place our work in context, we offer some remarks on type-changing PDE in general. Perhaps the most notable examples arise in the mathematical description of transonic flow, and in particular flow over an air foil, a scenario for which subsonic and supersonic regions are respectively described by elliptic and hyperbolic equations. The Frankl–Chaplygin equation [20] $K(y)u_{xx} + u_{yy} = 0$ (the choice $K(y) = y$ determines the familiar Tricomi equation [21]) serves as a fundamental model of this behavior, and Morawetz carried out early and fundamental analytic studies [22] of it and corresponding first order systems (see also [23]). The first truly successful method for numerically calculating transonic flow was put forth in the seminal paper [24] by Murman and Cole. They adopted a relaxation-like method along with type-dependent finite difference stencils. For an account of the impact of the Murman–Cole method in aerodynamics, and a review of more modern CFD approaches towards transonic flow, see the historical perspective of [25]. While the aerodynamic scenario is the most widely known, special examples of mixed-type equations arise in fields as diverse as plasma physics [26] and windshield design [27]. An early work most closely related to our own is Chaohao’s examination [28] of amplifying spiral wave solutions to the 2+1 wave equation. Numerical methods for mixed-type problems tend to be equation specific. The Murman–Cole technique, for example, would seem to have no application for our problem.

Motivated by mixed-type problems, in the late 1950s Friedrichs initiated a program [29] for analyzing a wide class of boundary value problems based on operators whose symmetric part is positive definite and which obey certain admissible boundary conditions. Such *symmetric positive* systems include PDE of hyperbolic,

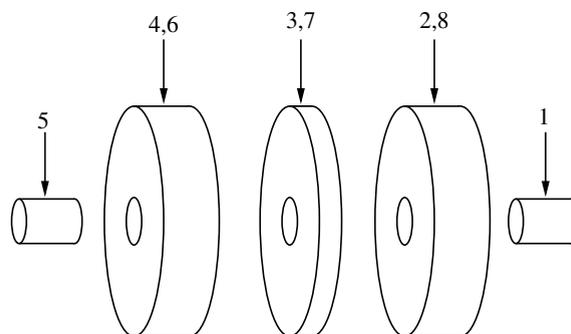


Fig. 2. Domain decomposition for the 3d HRWE Problem. Here the inner spherical shell and outer spherical shell, respectively corresponding to the annuli H and A in Fig. 1, are not shown. The remaining domains (all lying within the elliptic region) are shown in an exploded format for emphasis.

elliptic, and mixed type. A lucid history of the Friedrichs program from both theoretical and numerical perspectives is given by Jensen [30]. Rather early on, Katsanis developed a numerical method for solving Friedrichs systems [31]. Starting with a Friedrichs system, he applied Green’s theorem in a generic cell, and approximated the resulting integral equation. At the discrete level the approach faithfully mimicked both the symmetric positive aspect of the operators as well as admissibility of boundary conditions, and Katsanis went on to numerically examine the Tricomi problem [32]. Although geometrically flexible, the Katsanis method is low-order accurate. A finite-difference method for Friedrichs system was also outlined by Liu [33]. Recent work towards numerically solving Friedrichs systems has drawn on the powerful framework of discontinuous Galerkin methods [30,34]. Although Torre has shown that the 2d HRWE on a disk can be cast into a first-order Friedrichs system (and, indeed, the aforementioned work by Chaohao also made connections with Friedrichs theory), we have nevertheless chosen a classic multidomain spectral method over a discontinuous Galerkin method. This would seem appropriate given the relatively simple geometry of our problem and the expected smoothness of the solutions we seek.

2. Outer annulus and outer boundary conditions

We begin by discussing the 2d HRWE on the outer annulus, labeled A in Fig. 1; with appropriate changes this discussion will also apply to the 3d HRWE on an outer spherical shell. As the relevant PDEs are type-changing on them, these are the most interesting subdomains. In addition, our discussion here will supply some of the details of the HRWE and will exhibit its mathematical features, including its change of type.

In this section, we will describe the radiative outer boundary conditions on A imposed at a large $r = r_{\max} \equiv R$. We shall also speak of an inner boundary condition for A , although when A is considered as a subdomain it is not associated with a true inner boundary condition. The philosophy here is that we must understand the relevant boundary value problem on each subdomain to have all subdomains successfully “glued” together. Inner boundary conditions on A are imposed at a radial value $r = r_{\min} \equiv \epsilon$. Note that ϵ is not small. For the decomposition illustrated in Fig. 1, in fact, ϵ must be greater than the orbital radius.

2.1. 2d HRWE

The linear 2d HRWE of Whelan et al. [8] starts with $-\partial_t^2 \psi + \nabla^2 \psi = g$ and, in terms of polar coordinates r, ϕ , the helical reduction requires that ψ be a function only of r and $\varphi = \phi - \Omega t$. The 2d HRWE then takes the form

$$\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \psi}{\partial r} \right) + \frac{1}{r^2} \zeta(r) \frac{\partial^2 \psi}{\partial \varphi^2} = g(r, \varphi), \tag{2}$$

where $g(r, \varphi)$ is a φ -periodic source term, $\zeta(r) = 1 - r^2 \Omega^2$, and the coordinate ranges are

$$\epsilon \leq r \leq R, \quad 0 \leq \varphi \leq 2\pi. \tag{3}$$

Eq. (2) is clearly elliptic for $r < |\Omega|^{-1}$ and hyperbolic for $r > |\Omega|^{-1}$. One boundary value problem would be to seek solutions to this equation which obey the boundary conditions

$$\psi|_{\epsilon} = \alpha(\varphi), \quad \left(\frac{\partial \psi}{\partial r} - \Omega \frac{\partial \psi}{\partial \varphi} + \frac{\psi}{2r} \right) \Big|_R = \beta(\varphi), \tag{4}$$

here with φ -periodic functions $\alpha(\varphi)$ and $\beta(\varphi)$. This boundary value problem for annulus A is equivalent to the punctured disk examined by Torre [13]. His proof of the well-posedness of this problem assumes that the radial endpoint ϵ lies in the elliptic region, as do we by requiring $\epsilon < \Omega^{-1}$. Torre’s proof places no restriction on R . He allows for R to lie in the hyperbolic region, the elliptic region, or even on the *light cylinder* $r = |\Omega|^{-1}$. In practice of course, R is large and lies in the hyperbolic region, as we will assume.

Fourier transformation of (2) yields

$$\frac{1}{r} \frac{d}{dr} \left(r \frac{\partial \hat{\psi}_n}{\partial r} \right) - \frac{n^2}{r^2} \zeta(r) \hat{\psi}_n = \hat{g}_n(r), \tag{5}$$

and for the corresponding boundary conditions on the mode $\hat{\psi}_n(r)$ we have

$$\hat{\psi}_n(\varepsilon) = \hat{\alpha}_n, \quad \left(\frac{d\hat{\psi}_n}{dr} - in\Omega\hat{\psi}_n + \frac{\hat{\psi}_n}{2r} \right) \Big|_R = \hat{\beta}_n. \tag{6}$$

For now we will assume $n \neq 0$, and consider the zero-mode case separately later on. As an alternative we may also impose exact outgoing-wave boundary conditions, enforced n -by- n via

$$\left(\frac{d\hat{\psi}_n}{dr} - in\Omega\hat{\psi}_n + \frac{\hat{\psi}_n}{2r} \right) \Big|_R = \frac{1}{R} \left[n\Omega R \frac{V'_n(n\Omega R)}{V_n(n\Omega R)} \right] \hat{\psi}_n, \tag{7}$$

where

$$V_\nu(z) = \sqrt{\frac{\pi z}{2}} \exp \left[-i \left(z - \frac{1}{2} \pi \nu - \frac{1}{4} \pi \right) \right] H_\nu^{(+)}(z) \tag{8}$$

is set up to satisfy $V_\nu(z) \sim 1$ as $z \rightarrow \infty$. Here $H_\nu^{(+)}(z)$ is the first cylindrical Hankel function. The “frequency-domain kernel”

$$n\Omega R \frac{V'_n(n\Omega R)}{V_n(n\Omega R)} \tag{9}$$

can be computed as a continued fraction via Steed’s algorithm [35]. Similar kernels appear in studies of radiation boundary conditions for time-domain wave propagation, for example in Refs. [36–39].

2.2. 3d HRWE

For the 3d case we will mostly use the same symbols as for the 2d case. The 3d HRWE of Andrade et al. [11] is

$$\frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial \psi}{\partial r} \right) + \left(\frac{\Delta_{S^2}}{r^2} - \Omega^2 \frac{\partial^2}{\partial \varphi^2} \right) \psi = g(r, \theta, \varphi), \tag{10}$$

where $g(r, \theta, \varphi)$ is again a φ -periodic source term, Δ_{S^2} is the unit two-sphere Laplacian, and the coordinate ranges now are

$$\varepsilon \leq r \leq R, \quad 0 \leq \theta \leq \pi, \quad 0 \leq \varphi \leq 2\pi. \tag{11}$$

Eq. (10) is elliptic for $r \sin \theta < |\Omega|^{-1}$ and hyperbolic for $r \sin \theta > |\Omega|^{-1}$. We seek solutions to this equation which obey the boundary conditions

$$\psi|_\varepsilon = \alpha(\theta, \phi) \quad \left(\frac{\partial \psi}{\partial r} - \Omega \frac{\partial \psi}{\partial \varphi} + \frac{\psi}{r} \right) \Big|_R = \beta(\theta, \phi). \tag{12}$$

We consider the exact outgoing-wave conditions below. As before, we will assume that $\varepsilon < |\Omega|^{-1}$, and explain why in a moment.

Spherical harmonic transformation of (10) yields

$$\frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{d\hat{\psi}_{\ell m}}{dr} \right) - \left[\frac{\ell(\ell+1)}{r^2} - m^2 \Omega^2 \right] \hat{\psi}_{\ell m} = \hat{g}_{\ell m}(r), \tag{13}$$

and we will now take

$$\hat{\psi}_{\ell m}(\varepsilon) = \hat{\alpha}_{\ell m}, \quad \left(\frac{d\hat{\psi}_{\ell m}}{dr} - im\Omega\hat{\psi}_{\ell m} + \frac{\hat{\psi}_{\ell m}}{r} \right) \Big|_R = \hat{\beta}_{\ell m} \tag{14}$$

as the corresponding boundary conditions. We will assume $m \neq 0$, and consider the $m = 0$ case separately later on. As an alternative we may also impose exact outgoing-wave boundary conditions. Enforced mode-by-mode, they have a form similar to Eq. (7),

$$\left(\frac{d\hat{\psi}_{\ell m}}{dr} - im\Omega\hat{\psi}_{\ell m} + \frac{\hat{\psi}_{\ell m}}{r} \right) \Big|_R = \frac{1}{R} \left[m\Omega R \frac{V'_{\ell+1/2}(m\Omega R)}{V_{\ell+1/2}(m\Omega R)} \right] \hat{\psi}_{\ell m}. \tag{15}$$

To cast the radial equation stemming from the 3d HRWE into a form which resembles the radial equation stemming from the 2d HRWE, we substitute $\hat{\psi}_{\ell m} = \hat{\xi}_{\ell m}/\sqrt{r}$ in (13), thereby finding

$$\frac{1}{r} \frac{d}{dr} \left(r \frac{d\hat{\xi}_{\ell m}}{dr} \right) - \frac{(\ell + \frac{1}{2})^2}{r^2} \zeta(r) \hat{\xi}_{\ell m} = \sqrt{r} \hat{g}_{\ell m}(r), \tag{16}$$

where

$$\zeta(r) = 1 - \frac{4m^2\Omega^2 r^2}{(2\ell + 1)^2}. \tag{17}$$

In terms of $\omega_{\ell m} = 2m\Omega/(2\ell + 1) \neq 0$, clearly something special occurs for (16) when $r = |\omega_{\ell m}|^{-1}$, here with $\ell, m \neq 0$. Note that $|\omega_{\ell m}| \leq |\Omega|$, so that $|\Omega|^{-1} \leq |\omega_{\ell m}|^{-1}$. For this reason, we have required $\varepsilon < |\Omega|^{-1}$ above, ensuring that $r = \varepsilon$ lies within each mode’s individual “elliptic region.”

2.3. General form for the outer boundary conditions

From Eq. (5) we pass to a trigonometric rather than exponential representation of the transform $\hat{\psi}_n(r)$ by writing

$$\begin{aligned} \hat{\psi}_n^*(r) \exp(-in\varphi) + \hat{\psi}_n(r) \exp(in\varphi) &= \left[\hat{\psi}_n(r) + \hat{\psi}_n^*(r) \right] \cos(n\varphi) + i \left[\hat{\psi}_n(r) - \hat{\psi}_n^*(r) \right] \sin(n\varphi) \\ &= u_n(r) \cos(n\varphi) + w_n(r) \sin(n\varphi), \end{aligned} \tag{18}$$

with a similar splitting possible for $\hat{\psi}_{\ell m}(r)$, or $\hat{\xi}_{\ell m}(r)$ if preferred.

At $r = R$ all boundary conditions, whether exact or some incarnation of Sommerfeld, may be expressed as follows:

$$Rw'_n(R) + pu_n(R) + qw_n(R) = 0, \quad Ru'_n(R) - pw_n(R) + qu_n(R) = 0. \tag{19}$$

We list the possibilities considered so far in Table 1. Other choices of the form (19) are of course possible. For the exact conditions listed in Table 1, we have made use of the notation

$$v_v(z) = z \frac{V'_v(z)}{V_v(z)}, \tag{20}$$

with the obvious notation for real and imaginary parts.

Turning now to the rather more delicate zero-modes ($n = 0$ for 2d and $m = 0$ for 3d), we note that for $z \rightarrow 0$ the Hankel function $H_v^{(+)}(z)$ is proportional to z^{-v} for $v \neq 0$ and to $\log z$ for $v = 0$, so that in either case

$$zH_v^{(+)\prime}(z)/H_v^{(+)}(z) \rightarrow -v, \tag{21}$$

and hence from the the definition in Eq. (8) we have

$$v_v(0) = \frac{1}{2} - v \tag{22}$$

Table 1
Outer boundary conditions for the 2d and 3d HRWE

Boundary condition	p	q
Sommerfeld on $\hat{\psi}_n$	$n\Omega R$	$\frac{1}{2}$
Sommerfeld on $\hat{\psi}_{\ell m}$	$m\Omega R$	$\frac{1}{2}$
Exact on $\hat{\psi}_n$	$n\Omega R + \text{Im}v_n(n\Omega R)$	$\frac{1}{2} - \text{Re}v_n(n\Omega R)$
Exact on $\hat{\psi}_{\ell m}$	$m\Omega R + \text{Im}v_{\ell+1/2}(m\Omega R)$	$\frac{1}{2} - \text{Re}v_{\ell+1/2}(m\Omega R)$

From this our key result follows. For the $n=0$ homogeneous case, Eq. (5) has solutions $\hat{\psi}_0(r) = c$ or $\hat{\psi}_0(r) \propto \log r$. With $v_0(0) = \frac{1}{2}$, we find $p=0=q$ for the exact outgoing conditions as stated in (19). So this boundary condition rules out the $\log r$ solution. For the $m=0$ homogeneous case, Eq. (13) has solutions $\hat{\psi}_{\ell 0} \propto r^\ell$ and $\hat{\psi}_{\ell 0} \propto r^{-(\ell+1)}$. Now $v_{\ell+1/2}(0) = -\ell$, and $p=0, q=\ell+1$ for the exact outgoing boundary condition. In this case the boundary condition rules out the r^ℓ solution.

3. Sparse spectral approximation of the 2d HRWE

This section describes how we use spectral methods to numerically solve the 2d HRWE on the two center domain shown in Fig. 1. Our method will exploit the spectral *integration preconditioning* (IPC) proposed by Coutsias et al. [40], a general-purpose spectral method for solving ODEs and PDEs. Developing IPC for general orthogonal polynomials, they mostly considered it in the context of ODEs. In this context they presented a detailed theoretical analysis of both conditioning and convergence (see also [41]). Although they also carefully outlined how to apply the method in higher dimensions (with several illuminating two-dimensional examples), they did not explore conditioning issues for PDEs.

We shall follow the IPC approach for the following reasons. Foremost, it is a direct recipe for applying spectral methods in a PDE setting, affords a straightforward way of treating both boundary conditions and the “gluing” of subdomains. Moreover, IPC allows us simultaneously to achieve a sparse banded matrix representation of the 2d HRWE on all subdomains. While we question whether a fully 3d multidomain PSW problem can be treated exclusively with spectral methods based on IPC (further comments to follow), we believe the method is well suited for handling the 2d HRWE on the outer annulus A (or the 3d HRWE on an outer spherical shell). The mixed-type boundary value problem on this subdomain introduces several novel features, and much of our analysis below centers on associated conditioning issues. In short, the IPC approach offers us a quick and easy way to test the fundamental idea of using a multidomain spectral method to solve the 2d HRWE, and has promise for at least one key aspect of true 3d PSW problems.

In following the IPC approach we are able to explicitly form the matrix which represents the 2d HRWE on the two center domain, and we subsequently use Gaussian elimination (as embodied by the `netlib` routine `dgesv`, and sometimes also `dgesvx`) to invert the resulting system. A discussion of the structure of the matrix serves to sharpen our analytic understanding of the linear systems we are dealing with. For 3d problems, however, speed and memory considerations make it impractical to form and solve the full matrix. For the 3d model we consider in the conclusion and for envisioned future 3d work, we have used, and plan to continue with the Krylov-based method GMRES, in which only the specification of a matrix–vector multiply need be implemented. For a Krylov method preconditioning strategies are often necessary to avoid stagnation of the iterative solver.

Despite the name of the method, and despite its intuitive appeal, it is not guaranteed that IPC, especially for rectangular domains, actually improves conditioning (with respect to the problem of matrix inversion). If this is an issue in two dimensions, it is sure to be even more problematic in three dimensions. Since we cannot guarantee that our implementation of IPC will actually improve conditioning, it would be more appropriate to call the method a “sparse formulation.” To refer to the very specific IPC technique it will be convenient, however, for us to continue to use the term “preconditioning.” For simplicity, and to have a uniform treatment, we will use the IPC method for all our subdomains in the 2d model. It should be understood that this uniformity is not necessary. We could treat the elliptic region via a different method (spectral or pseudospectral). In any case, as we verify with numerical experiments, the method yields impressive global accuracy in solving the 2d HRWE on our two center domain.

3.1. Basic formulas for Chebyshev polynomials

Although Ref. [40] considered general orthogonal polynomials, we work solely with Chebyshev polynomials, a classical system of orthogonal polynomials with many useful properties and applications [42,43]. Here we collect only those properties relevant for our discussion, mostly following [40,44,45]. The degree- n Chebyshev polynomial $T_n(\xi)$ is defined by $T_n(\xi) = \cos(n \arccos(\xi))$ for $-1 \leq \xi \leq 1$, showing that we may consistently set $T_{-n}(\xi) = T_n(\xi)$. In our application ξ depends on one of the coordinates, say

$\xi(r) = (2r - r_{\max} - r_{\min}) / (r_{\max} - r_{\min})$ in an outer spherical shell or annulus. The $T_n(\xi)$ are solutions to a singular Sturm–Liouville problem, and therefore particularly suited for approximating solutions to differential equations on $[-1, 1]$ with a wide class of boundary conditions. The $T_n(\xi)$ are orthogonal on the interval $[-1, 1]$ with respect to the weight function $(1 - \xi^2)^{-1/2}$.

We have $T_0(\xi) = 1$, $T_1(\xi) = \xi$, and the following identity:

$$2\xi T_n(\xi) = T_{n+1}(\xi) + T_{n-1}(\xi). \tag{23}$$

If we denote by $\mathbf{T}(\xi)$ the row vector $[T_0(\xi), T_1(\xi), T_2(\xi), \dots]$, then (23) may be expressed as

$$\xi \mathbf{T}(\xi) = \mathbf{T}(\xi) A, \tag{24}$$

where²

$$A = \begin{bmatrix} 0 & \frac{1}{2} & 0 & 0 & 0 & 0 & \dots \\ 1 & 0 & \frac{1}{2} & 0 & 0 & 0 & \dots \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & \dots \\ 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & \dots \\ 0 & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & \dots \\ 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \tag{25}$$

Note that our convention here is to label matrix rows and columns starting from 0 rather than 1. Since A is tridiagonal, the matrix A^m , representing left multiplication by ξ^m through the formula $\xi^m \mathbf{T}(\xi) = \mathbf{T}(\xi) A^m$, has bandwidth $2m + 1$. Moreover, banded matrices $p(A)$ similarly correspond to multiplication by any polynomial $p(\xi)$. The most important special case, is the matrix $A_{T_m} \equiv T_m(A)$ of bandwidth $2m + 1$ corresponding to multiplication by the Chebyshev polynomial $T_m(\xi)$. For $m = 1$, the matrix $A = A_{T_1}$ is given by Eq. (25). More generally the entries of A_{T_m} may be gathered from the identity [40,42]

$$2T_m(\xi) T_n(\xi) = T_{m+n}(\xi) + T_{|m-n|}(\xi). \tag{26}$$

Next we consider the formal expansion

$$u(\xi) = \sum_{n=0}^{\infty} \tilde{u}_n T_n(\xi), \tag{27}$$

where, owing to the orthogonality of the Chebyshev basis, the expansion coefficients have the analytic values [42]

$$\tilde{u}_n = \kappa_n \int_{-1}^1 u(\xi) T_n(\xi) \frac{d\xi}{\sqrt{1 - \xi^2}}. \tag{28}$$

The prefactor κ_n is $1/\pi$ for $n = 0$ and $2/\pi$ for $n > 0$. We denote by $\tilde{\mathbf{u}}$ the column vector of expansion coefficients $[\tilde{u}_0, \tilde{u}_1, \tilde{u}_2, \dots]$. In the vector space of expansion coefficients multiplication by ξ is effected through multiplication by A , that is to say $\tilde{\mathbf{u}}^\xi \equiv A \tilde{\mathbf{u}}$ represents the expansion coefficients $[\tilde{u}_0^\xi, \tilde{u}_1^\xi, \tilde{u}_2^\xi, \tilde{u}_3^\xi, \dots]^t$ for $u^\xi(\xi) \equiv \xi u(\xi)$. Similarly, in the space of expansion coefficients, multiplication by any polynomial $p(\xi)$ may be represented via left multiplication by $p(A)$.

Another well-known formula [40,42,44]

$$T_n(\xi) = \frac{T'_{n+1}(\xi)}{2(n+1)} - \frac{T'_{n-1}(\xi)}{2(n-1)}, \tag{29}$$

yields another matrix identity

² For now we proceed with infinite-size matrices. However, numerical approximation necessarily entails suitable truncations. Our practice will be to use the same symbols for such truncations, and alert the reader whenever the viewpoint shifts.

$$D^2 = \begin{bmatrix} 0 & 0 & 4 & 0 & 32 & 0 & 108 & 0 & 256 & \dots \\ 0 & 0 & 0 & 24 & 0 & 120 & 0 & 336 & 0 & \dots \\ 0 & 0 & 0 & 0 & 48 & 0 & 192 & 0 & 480 & \dots \\ 0 & 0 & 0 & 0 & 0 & 80 & 0 & 280 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & 0 & 120 & 0 & 384 & \dots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 168 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 224 & \dots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots \\ \vdots & \ddots \end{bmatrix}. \tag{36}$$

The matrices D and D^2 obey

$$\mathbf{T}'(\xi) = \mathbf{T}(\xi)D, \quad \mathbf{T}''(\xi) = \mathbf{T}(\xi)D^2. \tag{37}$$

Notice that the first column of D and the first two columns of D^2 are (necessarily) zero. We might accordingly represent these matrices as $D_{\langle 1 \rangle}$ and $D_{\langle 2 \rangle}$ to highlight these facts, but will not employ this notation. Whereas the matrices $B_{[1]}$ and $B_{[1]}^2$ are banded and sparse, D and D^2 are clearly upper triangular and dense. To render the linear systems we encounter in banded form, we will exploit the identities, [40]

$$B_{[1]}D = I_{[1]}, \quad B_{[2]}^2D^2 = I_{[2]}, \quad B_{[2]}^2D = B_{[2]}. \tag{38}$$

Here the notation $I_{[1]}$ means the identity matrix with the first row set to zero; $I_{[2]}$ means the first two rows are zero.

We stress the crucial point that, for the ordering of matrix products in Eq. (38), the identities (38) also hold for $(N + 1)$ -by- $(N + 1)$ truncations of $B_{[1]}$, $B_{[2]}^2$, D , and D^2 .

In order to enforce boundary conditions we introduce both Dirichlet and Neumann vectors (radiation boundary conditions require a bit more thought and will be put off until later). The Dirichlet vectors are

$$\delta^+ = [1, 1, 1, 1, \dots] = [T_0(1), T_1(1), T_2(1), T_3(1), \dots], \tag{39}$$

$$\delta^- = [1, -1, 1, -1, \dots] = [T_0(-1), T_1(-1), T_2(-1), T_3(-1), \dots], \tag{40}$$

while the Neumann vectors are

$$v^+ = [0, 1, 4, 9, \dots] = [T'_0(1), T'_1(1), T'_2(1), T'_3(1), \dots], \tag{41}$$

$$v^- = [1, -1, 4, -9, \dots] = [T'_0(-1), T'_1(-1), T'_2(-1), T'_3(-1), \dots]. \tag{42}$$

The boundary conditions we encounter are easily expressed in terms of these vectors. For example, $\delta^- \cdot \mathbf{\bar{u}} = 0$ is a homogeneous Dirichlet boundary condition at the left endpoint, and $v^+ \cdot \mathbf{\bar{u}} = 0$ is a homogeneous Neumann boundary condition at the right endpoint.

We conclude this preliminary subsection by addressing a few practical matters. First, our discussion so far assumes the standard interval $[-1, 1]$. Modifications, albeit trivial ones, of the matrices A , $B_{[1]}$, $B_{[2]}^2$ and vectors v^\pm are needed for different intervals. For example, if we use the $T_n(\xi(r))$ to approximate functions on $r \in [r_{\min}, r_{\max}]$, then to represent double integration that corresponds to the physical interval we must send $B_{[2]}^2 \rightarrow 0.25(r_{\max} - r_{\min})^2 B_{[2]}^2$. Similar scalings must be made for A , $B_{[1]}$, and v^\pm . However, *we shall retain the same symbols for matrices which are so scaled*. Second, although we may theoretically consider an infinite expansion (27) for a given function u (perhaps a solution to a differential equation), numerical applications entail a suitable truncation $\mathcal{P}_N u(\xi) = \sum_{n=0}^N \bar{u}_n T_n(\xi)$, where \mathcal{P}_N represents the truncation operator. Moreover, most applications do not work directly with the analytic expansion coefficients (28). Rather, one typically approximates the integral in (28) via the quadrature rule stemming from Chebyshev–Gauss–Lobatto nodes (see, for example, [44]). This process introduces aliasing errors, and results in discrete expansion coefficients $\bar{u}_n^{\text{discrete}}$. A more theoretical treatment would, when appropriate, draw a careful distinction between the analytic and discrete expansion coefficients, but this issue is not of primary importance to us.

3.2. Sparse formulation on rectangular domains

In terms of co-moving Cartesian coordinates,

$$x = r \cos \varphi = r \cos(\phi - \Omega t), \quad y = r \sin \varphi = r \sin(\phi - \Omega t), \tag{43}$$

let the scalar field $\psi(x, y)$ obey the inhomogeneous 2d HRWE

$$\left[\partial_x^2 + \partial_y^2 - \Omega^2(x\partial_y - y\partial_x)^2 \right] \psi = g \tag{44}$$

on a rectangle $R = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$. We have also worked with the shifted equation resulting from the mapping $x \rightarrow x - c, y \rightarrow y - d$, where (c, d) is the center of the rectangle. For simplicity, we do not consider a shift here. Via repeated use of the Leibniz rule, we write (44) as

$$\left[\partial_x^2(1 - \Omega^2 y^2) + \partial_y^2(1 - \Omega^2 x^2) - \Omega^2(\partial_x x + \partial_y y - 2\partial_x \partial_y xy) \right] \psi = g, \tag{45}$$

in preparation for our sparse-matrix spectral approximation.

At the theoretical level, we may represent our solution in terms of a double Chebyshev expansion,

$$\psi(x, y) = \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} \tilde{\psi}_{nm} T_n(\zeta(x)) T_m(\eta(y)), \tag{46}$$

where $(\zeta(x), \eta(y))$ is a mapping of our rectangle R onto the unit rectangle $[-1, 1] \times [-1, 1]$. To obtain the system of equations we solve numerically, we consider the truncated series

$$\mathcal{P}_{N,M} \psi(x, y) = \sum_{n=0}^N \sum_{m=0}^M \tilde{\psi}_{nm} T_n(\zeta(x)) T_m(\eta(y)). \tag{47}$$

We represent the finite collection of expansion coefficients as a 1-vector

$$\tilde{\psi} = (\tilde{\psi}_{00}, \tilde{\psi}_{01}, \dots, \tilde{\psi}_{0M}, \tilde{\psi}_{10}, \tilde{\psi}_{11}, \dots, \tilde{\psi}_{1M}, \dots, \tilde{\psi}_{N0}, \tilde{\psi}_{N1}, \dots, \tilde{\psi}_{NM})^t, \tag{48}$$

so that the components $\tilde{\psi}(k) = (\tilde{\psi})_k$ are determined by the direct product representation

$$\tilde{\psi}(n(M + 1) + m) = \tilde{\psi}_{nm}. \tag{49}$$

We then consider the approximation of (45) in terms of $\tilde{\psi}$ and suitable truncations of spectral differentiation matrices,

$$[D_x^2 \otimes (I_y - \Omega^2 A_y^2) + (I_x - \Omega^2 A_x^2) \otimes D_y^2 - \Omega^2(D_x A_x \otimes I_y + I_x \otimes D_y A_y - 2D_x A_x \otimes D_y A_y)] \tilde{\psi} = \tilde{\mathbf{g}}. \tag{50}$$

Here D_x represents the $(N + 1)$ -by- $(N + 1)$ differentiation matrix in the Chebyshev basis $T_n(\zeta(x))$, and A_x represents the $(N + 1)$ -by- $(N + 1)$ matrix corresponding to multiplication by x , with similar statements for D_y and A_y which are both $(M + 1)$ -by- $(M + 1)$. As mentioned at the end of the last subsection, to obtain these matrices appropriate scaling factors must be included with the straightforward truncations of the infinite-size matrices listed above. By the definition of the Kronecker direct product, we have, for example, that

$$(D_x^2 \otimes A_y^2)(n(M + 1) + m, k(M + 1) + p) = (D_x^2)(n, k)(A_y^2)(m, p), \tag{51}$$

in the MATLAB notation $C(i, k)$ for entries of a matrix C_{ik} .

To achieve a sparse and banded representation of the approximation, we multiply (50) by $\mathcal{B} = B_{x[2]}^2 \otimes B_{y[2]}^2$ and exploit the identities (38), thereby reaching

$$\begin{aligned} & [I_{x[2]} \otimes B_{y[2]}^2 (I_y - \Omega^2 A_y^2) + B_{x[2]}^2 (I_x - \Omega^2 A_x^2) \otimes I_{y[2]} - \Omega^2 (B_{x[2]} A_x \otimes B_{y[2]}^2 + B_{x[2]}^2 \otimes B_{y[2]} A_y - 2B_{x[2]} A_x \otimes B_{y[2]} A_y)] \tilde{\psi} \\ & = B_{x[2]}^2 \otimes B_{y[2]}^2 \tilde{\mathbf{g}}. \end{aligned} \tag{52}$$

The coefficient matrix of the system is then

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ \frac{I_{[2]}}{4} & 0 & -\frac{I_{[2]}}{6} & 0 & \frac{I_{[2]}}{24} & 0 & \dots & 0 \\ 0 & \frac{I_{[2]}}{24} & 0 & -\frac{I_{[2]}}{16} & 0 & \frac{I_{[2]}}{48} & \dots & 0 \\ 0 & 0 & \frac{I_{[2]}}{48} & 0 & -\frac{I_{[2]}}{30} & 0 & \dots & 0 \\ 0 & 0 & 0 & \frac{I_{[2]}}{80} & 0 & -\frac{I_{[2]}}{48} & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \frac{I_{[2]}}{4N(N-1)} & 0 & -\frac{I_{[2]}}{2(N^2-1)} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & B_{[2]}^2 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & B_{[2]}^2 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & B_{[2]}^2 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & B_{[2]}^2 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & B_{[2]}^2 \end{bmatrix} - \Omega^2[\dots],$$

where we have only explicitly displayed the Ω -independent contribution to the matrix stemming from the preconditioned Laplacian. This Laplacian contribution exhibits the main features of the overall matrix. Namely, that it is banded, sparse, and contains free rows of zeros. Each entry in either of the matrices above is itself an $(M + 1)$ -by- $(M + 1)$ matrix. Indeed, the $I_{[2]}$ and $B_{[2]}^2$ matrices are $I_{y[2]}$ and $B_{y[2]}^2$ in the notation of Eq. (52), and in the y dimension the truncation is $m = 0, \dots, M$. Overall, the coefficient matrix has $(N + 1)^2$ such blocks. The first two block-rows in the matrices displayed above are empty, giving $2(M + 1)$ free rows of zeros. In each of the remaining $N - 1$ block-rows the first two rows are empty, giving us another $2(N - 1)$ zero rows. Therefore, we have a total of $2(N + M)$ such zero rows at our disposal.

In the free zero rows we insert the “ τ -conditions” [40], that is, the boundary conditions. To illustrate, we here assume Dirichlet conditions $\psi(x_{\min}, y) = f^-(y)$, $\psi(x_{\max}, y) = f^+(y)$, $\psi(x, y_{\min}) = h^-(x)$, $\psi(x, y_{\max}) = h^+(x)$. These boundary conditions can be approximated as

$$\sum_{m=0}^M \tilde{\psi}_{nm} \delta_m^\pm = \tilde{h}_n^\pm, \quad \sum_{n=0}^N \tilde{\psi}_{nm} \delta_n^\pm = \tilde{f}_m^\pm, \tag{53}$$

where one-dimensional Chebyshev projections appear on the right hand sides. There are $2(N + 1) + 2(M + 1)$ boundary conditions in (53), but we now show that they are not all linearly independent due to double counting of corner conditions. The value of $\psi(x_{\max}, y_{\max}) = h^+(x_{\max}) = f^+(y_{\max})$ can be written as a linear combination either of the \tilde{h}_n^+ or of the \tilde{f}_m^+ . This implies a homogeneous linear relationship between the summations in the first and the second set of equations in (53). There are three other such linear relationships that follow from the ways in which $\psi(x_{\max}, y_{\min})$, $\psi(x_{\min}, y_{\max})$, and $\psi(x_{\min}, y_{\min})$ can each be expressed either in terms of sums of \tilde{h}_n^\pm or of \tilde{f}_m^\pm . Therefore, the number of independent equations in (53) is $2(N + 1) + 2(M + 1) - 4 = 2(N + M)$, precisely equal to the number of zero rows.

All methods for eliminating the four linear dependencies in (53) should yield comparable accuracy for a numerical solution. However, dropping the four highest-mode equations ($n = N$ for \pm in the left equation, and $m = M$ for \pm in the right equation) is inconsistent, as confirmed by numerical experiments. Indeed, by throwing out the highest mode on all four edges, one loses information about the corner values. We have chosen to reduce (53) in a consistent way which is at the same time convenient for our direct product representation of the rectangular region. We use all $2N + 2$ of the first set of equations in (53), i.e. those equations involving the \tilde{h}_n^\pm . The Dirichlet vectors δ_m^\pm associated with these conditions are placed within the first two rows of each block in the coefficient matrix (which does not increase the bandwidth of the coefficient matrix beyond the block-diagonal). We must then dispense with four equations from the second set, two for $-$ (left) and two for $+$ (right), and so drop those equations for which $m = M - 1$ and $m = M$. This corresponds to dropping the two highest modes for both left and right Dirichlet conditions, although we have experimented with dropping other modes and, as expected, found little difference. From knowledge of all the \tilde{h}_n^\pm (information determining all four physical corner values) and the \tilde{f}_m^\pm for $m = 0, \dots, M - 2$ we may recover the four values for \tilde{f}_{M-1}^\pm and \tilde{f}_M^\pm , using the four equations for the corner values expressed as linear combinations of the \tilde{f}_m^\pm . The Dirichlet

vectors δ_n^\pm associated with the remaining left/right boundary conditions are then placed in the remaining $2M - 2$ rows of the first two blocks (they reach across block columns and so affect the bandwidth of the coefficient matrix quite a bit). Values appearing on the righthand side of the above boundary conditions, all \tilde{h}_n^\pm and \tilde{f}_m^\pm except for \tilde{f}_{M-1}^\pm and \tilde{f}_M^\pm , replace the appropriate zero entries of the source $B_{x[2]}^2 \otimes B_{y[2]}^2 \tilde{\mathbf{g}}$.

Although, our preconditioning of the rectangular operator has resulted in a sparse matrix (while the original matrix corresponding to the D s is dense), the issue of condition number is more subtle. Clearly, the B s and D s are in some sense inverses of each other. Notice that as an infinite dimensional matrix $B_{[1]}$ has the action: $Q_0^{N-1} \rightarrow Q_1^N$, with Q_k^p denoting the vector subspace of spectral coefficients corresponding to Chebyshev expansion from degree k through degree p polynomials. Suppose we consider the $(N + 1)$ -by- $(N + 1)$ truncation of $B_{[1]}$, and further that we delete the first row and last column of this square matrix, thereby obtaining an N -by- N matrix \bar{B} . Likewise, we take the $(N + 1)$ -by- $(N + 1)$ truncation of D , and delete its first column and last row to obtain \bar{D} . Then we may view $\bar{D} : Q_1^N \rightarrow Q_0^{N-1}$. Taken as square matrices with the same domain and range, \bar{B} and \bar{D} are nonsingular and inverses of each other, whence have the same condition number.³ Such an argument can produce nonsingular matrices \bar{B}^2 and \bar{D}^2 , also inverses of each other with the same condition number. Therefore, although we are ignoring the critical issue of boundary conditions, passing from a coefficient matrix with symbolic form $I \otimes D^2 + D^2 \otimes I$ (corresponding to the Laplacian part of the operator) to one with the form $B^2 \otimes I + I \otimes B^2$ is not clearly advantageous insofar as the conditioning of the resulting linear system is concerned. Nevertheless, one might expect that a better distribution of eigenvalues for the form $B^2 \otimes I + I \otimes B^2$ would lead to faster convergence were we using an iterative solver such as GMRES [46].

3.3. Sparse formulation on annular domains

In our 2d HRWE $L\psi = g$, we now take $x = a + \rho \cos \theta$ and $y = b + \rho \sin \theta$, where the ‘‘hole’’ is located at (a, b) , which is either the center of annulus H or the center of A . In terms of

$$F(\theta) = a \sin \theta - b \cos \theta, \quad G(\theta) = a \cos \theta + b \sin \theta, \tag{54}$$

the HRWE operator is

$$L = \partial_\rho^2 + \rho^{-1} \partial_\rho + \rho^{-2} \partial_\theta^2 - \Omega^2 [F(\theta) \partial_\rho + (1 + \rho^{-1} G(\theta)) \partial_\theta]^2. \tag{55}$$

To prepare for the integration preconditioning, we use $F'(\theta) = G(\theta)$ and $G'(\theta) = -F(\theta)$ along with repeated appeals to the Leibniz rule in order to obtain

$$\rho^2 L\psi = \rho^2 g \tag{56}$$

with

$$\begin{aligned} \rho^2 L = & \partial_\rho^2 \rho^2 (1 - \Omega^2 F^2) + \partial_\rho \rho [-3 + \Omega^2 (2F^2 + G^2 + \rho G)] + \partial_\theta^2 [1 - \Omega^2 (\rho + G)^2] + \Omega^2 \partial_\theta \rho F \\ & - 2\Omega^2 \partial_\theta \partial_\rho \rho F (\rho + G) + 1 - \Omega^2 (G^2 + \rho G). \end{aligned} \tag{57}$$

We now represent the solution on the annular subdomain in terms of a truncated Fourier–Chebyshev expansion,

$$\mathcal{P}_{M,N} \psi(\rho, \theta) = \sum_{n=0}^N \tilde{\psi}_{0n} T_n(\xi(\rho)) + \sum_{k=1}^{\frac{1}{2}M} \sum_{n=0}^N [\tilde{\psi}_{2k-1,n} \cos(k\theta) + \tilde{\psi}_{2k,n} \sin(k\theta)] T_n(\xi(\rho)), \tag{58}$$

where for simplicity we have here chosen M even. For the direct product representation,

$$\tilde{\psi}((N + 1)m + n) = \tilde{\psi}_{mn}, \tag{59}$$

which means each Fourier mode corresponds to its own $(N + 1)$ -by- $(N + 1)$ Chebyshev block. We therefore have the following matrix representation \mathcal{L} of $\rho^2 L$:

³ Recall that the condition number of a matrix A is $\kappa(A) = \|A^{-1}\| \|A\|$.

$$\begin{aligned} \mathcal{L} = & (I_\theta - \Omega^2 F^2) \otimes D_\rho^2 A_\rho^2 + (-3I_\theta + 2\Omega^2 F^2 + \Omega^2 G^2 - 2\Omega^2 D_\theta F G) \otimes D_\rho A_\rho + \Omega^2 (G - 2D_\theta F) \otimes D_\rho A_\rho^2 \\ & - \Omega^2 D_\theta^2 \otimes A_\rho^2 + \Omega^2 [D_\theta F - (I_\theta + 2D_\theta^2)G] \otimes A_\rho + (I_\theta + D_\theta^2)(I_\theta - \Omega^2 G^2) \otimes I_\rho, \end{aligned} \quad (60)$$

where sans serif F and G denote matrices in the Fourier sin/cos basis corresponding to multiplication by F and G . The entries of these matrices are determined by standard trigonometric addition-of-angle formulas, such as $2 \sin \alpha \cos \beta = \sin(\alpha + \beta) + \sin(\alpha - \beta)$. For the scenario of the 3d HRWE, such a spectral approximation for a spherical shell around an inner hole would be rather more problematic (even if the analogous matrix is not explicitly formed). Indeed, for that scenario we would need to contend with Wigner–Clebsch–Gordon coefficients arising from products of spherical harmonics $Y_{\ell m}$.

Since the spectral differentiation matrices D_ρ and D_ρ^2 are dense upper triangular, passage to a sparse-matrix formulation of the problem for an annulus necessarily requires that we apply the integration matrix $B_{\rho[2]}^2$. By contrast, since the matrices D_θ and D_θ^2 are already banded or diagonal, we will achieve a sparse formulation whether or not we precondition with θ integration. *In fact we have employed θ preconditioning*, since it might well yield a better distribution of eigenvalues and better equilibration properties. Nevertheless, for simplicity will here ignore IPC in θ , which, in any case, is logically different from IPC in ρ in that no boundary condition is associated with the periodic direction. The preconditioning integration matrix is then $\mathcal{B} = I_\theta \otimes B_{\rho[2]}^2$, and its application onto (60) yields

$$\begin{aligned} \mathcal{B}\mathcal{L} = & (I_\theta - \Omega^2 F^2) \otimes I_{\rho[2]} A_\rho^2 + (-3I_\theta + 2\Omega^2 F^2 + \Omega^2 G^2 - 2\Omega^2 D_\theta F G) \otimes B_{\rho[2]} A_\rho + \Omega^2 (G - 2D_\theta F) \otimes B_{\rho[2]} A_\rho^2 \\ & - \Omega^2 D_\theta^2 \otimes B_{\rho[2]}^2 A_\rho^2 + \Omega^2 [D_\theta F - (I_\theta + 2D_\theta^2)G] \otimes B_{\rho[2]}^2 A_\rho + (I_\theta + D_\theta^2)(I_\theta - \Omega^2 G^2) \otimes B_{\rho[2]}^2. \end{aligned} \quad (61)$$

The right hand side of $\rho^2 L\psi = \rho^2 g$ is now represented by $\mathcal{B} \cdot (I_\theta \otimes A_\rho^2) \tilde{\mathbf{g}} = (I_\theta \otimes B_{\rho[2]}^2 A_\rho^2) \tilde{\mathbf{g}}$. For the outer annulus we have $(a, b) = (0, 0)$, and so $\rho = r$ and $\theta = \varphi$. Also for this case $F(\theta) = 0 = G(\theta)$, and \mathcal{L} in (60) reduces to

$$\mathcal{L} = I_\varphi \otimes D_r^2 A_r^2 - 3I_\varphi \otimes D_r A_r + D_\varphi^2 \otimes (I_r - \Omega^2 A_r^2) + I_\varphi \otimes I_r. \quad (62)$$

We therefore find

$$\mathcal{B}\mathcal{L} = I_\varphi \otimes I_{r[2]} A_r^2 - 3I_\varphi \otimes B_{r[2]} A_r + D_\varphi^2 \otimes B_{r[2]}^2 (I_r - \Omega^2 A_r^2) + I_\varphi \otimes B_{r[2]}^2 \quad (63)$$

as the preconditioned matrix for the outer annulus.

Turning to the issue of boundary conditions, we first note that for $\mathcal{B}\mathcal{L}$ the first two rows of each block have all zero entries. Into these rows we therefore place the τ -conditions,

$$\sum_{n=0}^N \tilde{\psi}_{mn} \delta_n^\pm = \hat{h}_m^\pm \quad (64)$$

(and similarly for Neumann conditions). Here \hat{h}_m^\pm is the Fourier transform of the boundary conditions $h^\pm(\theta)$, for example with $h^-(\theta) = \psi(\rho_{\min}, \theta)$. Notice that the correct number of zero entries in the preconditioned source correspond to the inhomogeneity \hat{h}_m^\pm . For the radiation “ p, q boundary conditions” (19) we are only concerned with the outer annulus and the simpler operator (63). In this case, the overall $(N + 1)(M + 1)$ -by- $(N + 1)(M + 1)$ matrix $\mathcal{B}\mathcal{L}$ is block diagonal, with $(M + 1)$ blocks. Each block has the structure

$$\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathcal{B}\mathcal{L}^k \end{bmatrix} \quad (65)$$

for some Fourier wave number k . Here the $\mathbf{0}$ represents a row of zeros, and $\mathcal{B}\mathcal{L}^k$ is a nonzero $(N - 1)$ -by- $(N + 1)$ submatrix. As the radiation boundary conditions couple sine and cosine modes of the same wave number, we must consider the two consecutive blocks (one cosine and the other sine) associated with Fourier wave number $k \neq 0$. As depicted in the last equation, each of the two blocks has all zero entries in its first two rows. The p, q boundary conditions are enforced by choosing the overall block neighborhood as follows,

$$\left[\begin{array}{c|c} \delta^- & \mathbf{0} \\ p\delta^+ & R\nu^+ + q\delta^+ \\ \mathcal{B}\mathcal{L}^k & \mathbf{0} \\ \hline \mathbf{0} & \delta^- \\ R\nu^+ + q\delta^+ & -p\delta^+ \\ \mathbf{0} & \mathcal{B}\mathcal{L}^k \end{array} \right], \tag{66}$$

where $\mathbf{0}$ represents either a row or a $(N - 1)$ -by- $(N + 1)$ submatrix of zeros. Boundary conditions for $k = 0$ (the zero mode) are easier to enforce (only a single block need be considered) and are handled similarly.

3.4. Gluing of subdomains

So far we have described individual rectangular and annular subdomains (and their associated τ -conditions) as if these subdomains were decoupled. In fact, we “glue together” all or most of the subdomains shown in Fig. 1. This gluing takes two forms: (i) imposing matching conditions for adjacent rectangles; and (ii) imposing matching conditions for the overlap between an annulus and a set of rectangles. Before describing each case in more detail, we comment on how such gluing is reflected in the overall linear system. Let $\tilde{\psi}^H$ and $\tilde{\psi}^A$ represent the vectors of Fourier–Chebyshev expansion coefficients associated with the spectral representation of the solution on the annuli H and A . Similarly, let $\tilde{\psi}^j$ represent the vector of double-Chebyshev expansion coefficients associated with the spectral representation of the solution on the j th rectangle (with $1 \leq j \leq 8$). The overall set of unknowns is then the concatenation $\tilde{\Psi} = (\tilde{\psi}^H, \tilde{\psi}^1, \tilde{\psi}^2, \tilde{\psi}^3, \tilde{\psi}^4, \tilde{\psi}^5, \tilde{\psi}^6, \tilde{\psi}^7, \tilde{\psi}^8, \tilde{\psi}^A)^t$, which satisfies (for the linear problem) the spectral matrix form of Eq. (1)

$$\mathcal{M}\tilde{\Psi} = \mathcal{B}\tilde{\mathcal{G}}, \tag{67}$$

where $\tilde{\mathcal{G}}$ is a similar concatenation of the sources $\tilde{\mathbf{g}}$ on the individual subdomains and the \mathcal{B} indicates integration preconditioning on all subdomains. Symbolically, the coefficient matrix \mathcal{M} is $\mathcal{B}\mathcal{L}$, here with \mathcal{L} standing for the spectral representation of the HRWE operator L on the whole two center domain.

Each of the 10 subdomains (annuli A and H , as well as rectangles 1–8) in Fig. 1 are represented by one of 10 super-blocks (H – H , 1–1, 2–2, . . . , 8–8, A – A) which sit along the diagonal of the overall super matrix \mathcal{M} representing the PDE on the whole two center domain. We use the term “super-block” here since the matrix corresponding to each subdomain arises, as we have seen, from a direct product structure (and so could be viewed as already in a block form). The supplementary equations needed for gluing are placed within existing zero rows in the same manner as with the τ -conditions. However, the gluing conditions stretch beyond the super-block diagonal, since they are linear relationships between the spectral expansion coefficients on two (or more) separate subdomains. For example, the gluing together of subdomains 1 and 2 (which share a common edge) involves not only filling rows within the 1–1 and 2–2 super-blocks along the diagonal of \mathcal{M} , but also filling rows within the 1–2 and 2–1 off-diagonal super-blocks.

3.4.1. Gluing of rectangles to rectangles

For rectangles which meet at an edge we require both continuity in ψ and its first derivative $\partial\psi/\partial\nu$ (normal to the matching edge). We impose these requirements strongly, that is to say at the level of the numerical solution itself. Consider, for example, rectangles 1 and 2 in Fig. 1, which as indicated share the edge $y = \rho_{\min}$, where ρ_{\min} is the radius of the inner hole (the depicted excised region). We require that values of $\psi(x, \rho_{\min})$ and $\partial\psi/\partial y(x, \rho_{\min})$ agree at the Chebyshev–Gauss–Lobatto collocation points $x(\xi_i)$ whether these values are computed using the spectral coefficients of rectangle 1 or those of rectangle 2. For simplicity, we assume that the number N_1 of spectral elements for rectangle 1 is the same as that N_2 for rectangle 2, so that the matching equations are simply

$$\sum_{m=0}^{M_1} \tilde{\psi}_{nm}^1 \delta_m^+ = \sum_{m=0}^{M_2} \tilde{\psi}_{nm}^2 \delta_m^-, \quad \sum_{m=0}^{M_1} \tilde{\psi}_{nm}^1 \alpha_1 \nu_m^+ = - \sum_{m=0}^{M_2} \tilde{\psi}_{nm}^2 \alpha_2 \nu_m^- \tag{68}$$

for each value of n . The α factors here are the scalings of the Neumann vectors that are necessary since the range of y may not be $[-1, 1]$. (See the discussion following Eq. (42).) These matching conditions are reflected in the overall matrix \mathcal{M} as follows. As the super-block corresponding to each of the subdomains 1 and 2 has been preconditioned in the described fashion, each has a collection of zero rows in which we place the matching conditions. In, say, the zero rows belonging to the super-block 1–1, we insert the first set of conditions given in (68). In the zero rows belonging to the super-block 2–2, we similarly place the Neumann conditions, the second set of conditions given in (68). We note that this filling of zero rows to achieve the required matching does not affect the inhomogeneity, as these are homogeneous conditions (a linear sum of expansion coefficients for one subdomain plus a linear sum of expansion coefficients for another is set equal to zero).

In practice, the gluing procedure is plagued by the same sort of issues that arose in connection with Eq. (53). Ultimately due to redundant counting at corners, the full set of matching conditions responsible for the gluing of all subdomains (including annuli) are not linearly independent. Thus we find that the number of apparent matching (and boundary) conditions is greater than the number of zero rows available in the super matrix \mathcal{M} to hold the such conditions. As we did for a single rectangle, we have chosen to consistently eliminate linear dependencies at left–right edges/interfaces, as illustrated in the next paragraph. The number of zero rows in \mathcal{M} always equal to the number of linearly independent matching (and boundary) conditions, and in principle any two ways of implementing them should be equivalent. While the choice of implementation does of course affect the structure of the super matrix \mathcal{M} , it should not greatly affect the accuracy of our numerical solutions.

Let us turn to our illustrative example, the matching, along the vertical edge $x = x_H + \rho_{\min}$, of rectangles 2 and 3 in Fig. 1. In the matrix for rectangle 2 (super-block 2–2 of the overall super matrix \mathcal{M}) we have reserved $2M_2 - 2$ rows in the first two blocks for enforcing boundary (or matching) conditions at the left and right edges. We will now have $M_2 - 1$ of these available for matching at the left edge of rectangle 2. Likewise, in the matrix for rectangle 3 (super-block 3–3 of the overall super matrix \mathcal{M}) we will have $M_3 - 1$ available rows for the matching along the right edge of rectangle 3. We take $M_1 = M_2 = M$, so that we have a total of $2M - 2$ rows in the overall super matrix \mathcal{M} available to explicitly enforce the matching of rectangles 2 and 3. The full set of equations for gluing along this edge are analogous to (68), but with N s and M s interchanged,

$$\sum_{n=0}^{N_2} \tilde{\psi}_{nm}^2 \delta_n^- = \sum_{n=0}^{N_3} \tilde{\psi}_{nm}^3 \delta_n^+, \quad \sum_{n=0}^{N_2} \tilde{\psi}_{nm}^2 \alpha_2 v_n^- = - \sum_{n=0}^{N_3} \tilde{\psi}_{nm}^3 \alpha_3 v_n^+. \tag{69}$$

We follow a protocol similar to the convenient one outlined for fixing Dirichlet boundary conditions for a single isolated rectangle. Whereas we have chosen to devote a zero row in \mathcal{M} to each of the $2N + 2$ matching conditions (68) for a top-bottom gluing, for a left-right gluing we devote only $M - 1$ zero rows for each set in (69), that is $2M - 2$ rows in all. This means dropping the last two (highest mode) equations from each set. Since we work with the transform of the matching conditions, dropping the highest two modes does not mean we are not enforcing matching at a particular point along the interface. Moreover, although we have chosen not to reflect all of the explicit matching conditions (69) in \mathcal{M} , extra matching of rectangles 2 and 3 is afforded by other matching conditions in the overall linear system. For example, along both the top edge of rectangle 2 and top edge of rectangle 3, we enforce a full set of boundary (or matching) conditions with Dirichlet data coming from annulus A . Since that data will be smooth along the full length of the combined top edge for both rectangles 2 and 3, these top-edge boundary conditions effectively yield extra matching for the 2–3 vertical gluing, in particular enforcing both continuity and differentiability at the physical corner point common to each rectangle.

Another complication in gluing is the implementation of symmetry. Notice that the left edge of rectangle 4, for example, must be *flipped* and then glued to the left edge of rectangle 7. Similar flipped gluing must also be considered for subdomains 5 and 6. We enforce such flipped gluing of edges by demanding continuity in ψ only (since region 7 is already matched to region 6), and filling the available zeros rows of the 4–4 super-block of \mathcal{M} accordingly (as in this case the 6–6 and 7–7 super-blocks would no longer have appropriate free rows available).

Let us provide a brief sketch of rectangle gluing without the “conforming assumption.” Were we gluing, say, rectangle 1 to rectangle 2 (a top-bottom glue) with $N_1 \neq N_2$, we would deal with more complicated equations than (68). The second equation in (68), for example, could be written as

$$\sum_{m=0}^{M_1} \tilde{\psi}_{nm}^1 \alpha_1 v_m^+ = \tilde{h}_n^+, \tag{70}$$

where \tilde{h}_n^+ is simply the explicit sum on the righthand side when $N_1 = N_2$. However, were $N_1 \neq N_2$, it would then arise as the x -Chebyshev transform of the vector

$$h_n^+ = \sum_{p,q=0}^{N_2, M_2} \tilde{\psi}_{pq}^2 T_p(\xi^2(x_n)) \left[\frac{d}{dy} T_q(\eta^2(y)) \right] \Big|_{y=\rho_{\min}}. \tag{71}$$

Here the x_n are the $N_1 + 1$ (scaled) Chebyshev–Gauss–Lobatto points for rectangle 1, and $(\xi^2(x), \eta^2(y))$ are the linear functions mapping rectangle 2 to the square $[-1, 1] \times [-1, 1]$. As there would be $N_1 + 1$ of the equations (70), they would be placed in zero rows corresponding to the 1–1 super-block of \mathcal{M} , and would stretch across the 1–1 and 1–2 super-blocks. Equations for continuity similar to the first set of equations in (68), but using the x -Chebyshev transform for rectangle 2, would determine entries in \mathcal{M} stretching across the 2–1 and 2–2 super-blocks. The roles of 1 and 2 could be interchanged. The situation is nearly the same for a left-right gluing, modulo the issue of dropping the two highest modes.

The rows in the overall super-matrix \mathcal{M} responsible for all of the rectangle-to-rectangle gluing result in \mathcal{M} having a large condition number $\kappa(\mathcal{M})$. Indeed, for the truncations we later consider in our numerical experiments, the matrix for the elliptic region (which includes the H – H super-block as well) has a reciprocal condition number RCOND (a diagnostic determined by `dgesvx`) in the 10^{-6} – 10^{-10} range. As we will see in Section 4.2, such a large condition number leads to an accuracy problem for higher truncations. In an attempt to mitigate this problem, we have experimented with other strategies for rectangle gluing, for example, imposing point-space rather spectral versions of the matching and boundary conditions. Another possibility, suggested by a referee, and for which we had earlier carried out some numerical experiments, is to allow for overlap between adjacent rectangles. In this case the gluing procedure is similar to the described rectangle-annulus gluing, with the key advantage that all Neumann vectors can be done away with (this could lead to improved conditioning). The disadvantage is that significant overlap is typically required for good conditioning, and, therefore, more of the physical space is doubly covered. Beyond an increase in resources, such double cover also further complicates the gluing of the overall concatenation of rectangles to the annuli. Nevertheless, this idea deserves further study. To sum up, we have not exhaustively tried every approach, and our current one is quite possibly not optimal, although we believe it will be hard to do significantly better. Nevertheless, by switching from `dgesv` to `dgesvx`, we are able to achieve the requisite accuracy for our experiments. The routine `dgesvx` includes equilibration (which changes RCOND to the 10^{-4} – 10^{-6} range) and iterative refinement [49]. Moreover, with our approach we have carried out the experiment (“ $L = 5$ with linear mappings”) documented in Section 4.1 of [16], the Laplace equation for the exact solution $\ln(x^2 + y^2)$ on a 10×10 square with an excised 1×1 inner square “hole.” With `dgesv` we attain a best error measure which is about three orders of magnitude larger than the corresponding one reported in Fig. 3 of that reference. However, with `dgesvx` we find nearly identical results, and even a slightly better best error measure.

3.4.2. Gluing of annuli to rectangles

The annuli H and A depicted in Fig. 1 overlap multiple rectangles, and for this overlap the issue of gluing is complicated. Since the issue is essentially the same for the gluing of H to rectangles 1–8 or A to rectangles 1–4 and 6–8, let us here focus on the first case. For example, part of the outer edge of annulus H sits in rectangle 1; we require that the values of ψ along this outer edge of H agree whether they are computed with the spectral representation of ψ on H or with the spectral representation of ψ on rectangle 1. Similarly, the values of ψ at the left edge of rectangle 1 must be the same whether computed using the rectangle 1 or H spectral representation. That is to say, along these boundaries we demand agreement in the numerical solution whether determined by the expansion coefficients $\tilde{\psi}^H$ of the annulus or the coefficients $\tilde{\psi}^1$ for rectangle 1. An essential difference should be noted between this kind of gluing, and the gluing of two rectangles at an edge: Here we use Dirichlet matching along two curves rather than Dirichlet plus Neumann matching along a single shared boundary curve.

For the outer circular boundary $\rho = \rho_{\max}$ of the annulus H , we enforce matching at each of a set of Fourier collocation points $\{\theta_m : m = 1, 2, \dots, M_H\}$ through the following set of equations:

$$\sum_{n=0}^{N_H} \tilde{\psi}_{nm}^H \delta_n^+ = \hat{h}_m^+, \tag{72}$$

where now the \hat{h}_m^+ are not fixed Dirichlet values. They now arise via Fourier transform of

$$h_m^+ = \sum_{p,q=0}^{N_j, M_j} \tilde{\psi}_{pq}^j T_p(\xi^j(x_m)) T_q(\eta^j(y_m)). \tag{73}$$

Here the j th rectangle contains the point (x_m, y_m) corresponding to the collocation point (ρ_{\max}, θ_m) . These equations are placed within the zero rows of the H – H super-block. The explicit matching equations can be captured by expressing the Fourier transform $\hat{h}_m^+ = \sum_{k=1}^{M_H} \mathcal{F}_{mk} h_k^+$ as a matrix relation in terms of the $\tilde{\psi}_{pq}^j$.

The matching along one of the inner edges of a rectangle follows a similar pattern. As a concrete example, take again the inner edge of rectangle 1, and the condition

$$\sum_{n=0}^{N_1} \tilde{\psi}_{nm}^1 \delta_n^- = \tilde{f}_m, \tag{74}$$

where the \tilde{f}_m are now not fixed Dirichlet values. Rather, they arise as the y -Chebyshev transform of the vector

$$f_m = \sum_{p=0}^{N_H} \tilde{\psi}_{0p}^H T_p(\xi(\rho_m)) + \sum_{k=1}^{\frac{1}{2}M_H} \sum_{p=0}^{N_H} \left[\tilde{\psi}_{2k-1,p}^H \cos(k\theta_m) + \tilde{\psi}_{2k,p}^H \sin(k\theta_m) \right] T_p(\xi(\rho_m)). \tag{75}$$

Here the point (ρ_m, θ_m) corresponds to a Chebyshev–Gauss–Lobatto collocation point $(x_{\min}, y(\eta_m))$ along the inner edge, and we have chosen M_H as an even integer for simplicity. (Note that the values of θ_m here, with $0 \leq m \leq M_1$ have no direct relationship to the Fourier collocation points θ_m used in Eqs. (72) and (73).) Again, via a matrix representation $\tilde{f}_m = \sum_{\ell} \mathcal{C}_{m\ell} f_{\ell}$ of the transform, we may express this matching condition more directly. In any case, these conditions are placed within the available zero rows of the 1–1 super-block of \mathcal{M} .

3.5. Nonlinear model

We turn to the 2d nonlinear HRWE discussed in [8], that is

$$L\psi + \eta h(\psi) = g, \quad h(\psi) = \frac{\psi^5}{\psi^4 + \psi_0^4}. \tag{76}$$

Certain aspects of our treatment of the nonlinear term will likely not generalize to 3d. Indeed, we have retained a pure spectral method (coefficients as unknowns), whereas most modern approaches [46] towards solving nonlinear equations via Newton–Krylov spectral methods have centered around the construction of preconditioners within the context of pseudospectral methods (point values as unknowns). Nevertheless, as discussed in the concluding section, we expect that some aspects of our approach will carry over to 3d insofar as an outer spherical shell is concerned.

Again, we let $\tilde{\Psi}$ represent the vector of unknowns, that is the overall concatenation of the spectral coefficients on all subdomains. Using the inverse spectral transformation available on each subdomain, we may produce from $\tilde{\Psi}$ a collection of point-space values $\Psi = \mathcal{F}^{-1} \tilde{\Psi}$ defined on each subdomain’s spectral grid. Here we are using \mathcal{F}^{-1} to formally denote the process of inverse transformation on all subdomains. Our problem now is to solve

$$W[\tilde{\Psi}] = \mathcal{M} \tilde{\Psi} + \eta \mathcal{B} \mathcal{F} h(\mathcal{F}^{-1} \tilde{\Psi}) - \mathcal{B} \tilde{\mathcal{G}} = 0, \tag{77}$$

here with \mathcal{B} representing the application of integration preconditioning on all subdomains, $\tilde{\mathcal{G}}$ the concatenation of all source coefficients $\tilde{\mathbf{g}}$, and $\mathcal{M} = \mathcal{B} \mathcal{L}$ the coefficient matrix for the linear HRWE on the glued-together

two center domain. The term $\mathcal{F}h(\mathcal{F}^{-1}\tilde{\Psi}) = \mathcal{F}h(\Psi)$ is the spectral representation of the nonlinearity, where $h(\Psi)$ is constructed pointwise on each subdomain’s spectral grid.

Solving this nonlinear system of equations requires a Newton–Raphson iteration, and we briefly describe how this has been implemented. First, we note that by using both (26) and addition of angle formulas, such as $2 \sin \alpha \cos \beta = \sin(\alpha + \beta) + \sin(\alpha - \beta)$, we may convert the spectral coefficients $\tilde{\mathbf{u}}$ representing a function u on a particular subdomain into a matrix C_u whose action on the subdomain’s vector space of spectral coefficients corresponds to multiplication by u in physical space. For example, on a rectangular domain, the resulting matrix would arise as

$$\sum_{n=0}^N \sum_{m=0}^M \tilde{u}_{nm} T_n(\xi(x)) T_m(\eta(y)) \rightarrow C_u = \sum_{n=0}^N \sum_{m=0}^M \tilde{u}_{nm} A_{T_n} \otimes A_{T_m}, \tag{78}$$

where $A_{T_n} = T_n(A_x)$ is $(N + 1)$ -by- $(N + 1)$, $A_{T_m} = T_m(A_y)$ is $(M + 1)$ -by- $(M + 1)$, and each matrix has been appropriately scaled (as discussed at the end of Section 3.1). For both annuli and rectangles we have written subroutines for fast computation of such matrices from input coefficients. Then, given a concatenation $\tilde{\mathbf{U}} = (\tilde{\mathbf{u}}^H, \tilde{\mathbf{u}}^1, \tilde{\mathbf{u}}^2, \tilde{\mathbf{u}}^3, \tilde{\mathbf{u}}^4, \tilde{\mathbf{u}}^5, \tilde{\mathbf{u}}^6, \tilde{\mathbf{u}}^7, \tilde{\mathbf{u}}^8, \tilde{\mathbf{u}}^4)^t$ of the spectral coefficients on all subdomains, we may also form a (super-block diagonal) matrix C_u whose action on the vector space of spectral coefficients for the entire two center domain corresponds to multiplication by u in physical space over the entire two center domain.

Let $\tilde{\Psi}^{\text{old}}$ represent the current Newton iterate. Then computing the next iterate $\tilde{\Psi}^{\text{new}} = \tilde{\Psi}^{\text{old}} - \delta\tilde{\Psi}$ requires that we solve the linear equation

$$\left\{ \mathcal{M} + \eta \mathcal{B} C_{h'(\Psi^{\text{old}})} \right\} \delta\tilde{\Psi} = W[\tilde{\Psi}^{\text{old}}], \tag{79}$$

where the derivative function is

$$h'(\psi) = \frac{5\psi^4\psi_0^4 + \psi^8}{[\psi^4 + \psi_0^4]^2}. \tag{80}$$

The matrix $C_{h'(\Psi^{\text{old}})}$ is constructed as follows: (i) obtain the point values $\Psi^{\text{old}} = \mathcal{F}^{-1}\tilde{\Psi}^{\text{old}}$; (ii) construct the set of point values $h'(\Psi^{\text{old}})$; (iii) construct the concatenation of spectral coefficients $\tilde{\mathbf{U}} = \mathcal{F}h'(\Psi^{\text{old}})$; (iv) using the coefficients $\tilde{\mathbf{U}}$, form the matrix $C_{h'(\Psi^{\text{old}})}$ as outlined above. We have taken advantage of the sparse nature of the preconditioner \mathcal{B} to achieve fast matrix–matrix and matrix–vector multiplies, with the former required for quick computation of the matrix–matrix product $\mathcal{B}C_{h'(\Psi^{\text{old}})}$ that is needed for each Newton iterate. Each such linear solve has been performed directly with `dgesv`.

We have found that the described implementation of Newton–Raphson iteration works well for the 2d model, and for a wide range of parameter choices we have not needed to include line searches. Nevertheless, we recognize that such an implementation is not likely a viable option for 3d. Indeed, for spherical shells the scalar spherical harmonics $Y_{\ell m}$ are required for the angular basis functions. Beyond the memory burden of forming matrices for the 3d problem, even representation of multiplication by a nonlinear function in the spectral $Y_{\ell m}$ basis is not as practical, because it would require extensive use of Wigner–Clebsch–Gordan coefficients (the simple multiplicative properties embodied in addition-of-angle formulas are no longer at one’s disposal). However, in a Krylov approach one might use the spherical harmonic transform and inverse transform to the same effect, carrying out all multiplications in point space.

4. Numerical analysis and results

This section considers the conditioning of our numerical problem on the outer annulus A . It also presents our results in solving the HRWE on various subdomains as well as on the whole two center domain. In the first subsection we consider only a single Fourier mode in the outer annulus, and we study the condition number of the matrix that must be inverted. This analysis also pertains to the analogous outer shell problem associated with solving the 3d HRWE. Indeed, as we have shown in the second section, the ODE and boundary conditions governing the 2d and 3d mode equations are essentially the same. The second subsection describes methodology and collects error tables associated with our numerical experiments in solving

the linear and nonlinear 2d HRWE. A third subsection presents preliminary results in solving the 3d HRWE.

4.1. Condition numbers

Consider the radial operator which stems from the linear helical operator via a Fourier-series transform; consider also the radial operator that arises in this way from the Laplace–Poisson operator. With k taken as the wave number, these radial operators are respectively (see left hand side of Eq. (5))

$$L^k = r^2 \frac{d^2}{dr^2} + r \frac{d}{dr} - k^2(1 - \Omega^2 r^2), \quad (\text{Helical operator}) \tag{81}$$

with Dirichlet boundary conditions at the inner radius and exact outgoing radiation boundary conditions at the outer radius (D–R conditions), and

$$P^k = r^2 \frac{d^2}{dr^2} + r \frac{d}{dr} - k^2 \quad (\text{Poisson operator, } \Omega = 0). \tag{82}$$

For the Poisson operator we examine both Dirichlet–Dirichlet (D–D) and Dirichlet–Neumann (D–N) boundary conditions. As the Laplacian is translation invariant, the radial Poisson–Laplace operator P^k can be put into the same Fourier form wherever the annulus lies. The HRWE mixes Fourier modes if the center of the angular coordinates is offset from the center of rotation, as in the case of the hole subdomain H .

Here, however, we will work with the outer annulus A for which the HRWE itself does not mix modes, although, as we have seen, the radiation boundary condition does mix cos and sin modes of the same wave number. Therefore, we essentially have an ODE setting in which the conditioning properties of IPC are documented [40]. We point out that Ref. [40] also refers to the conditioning of a “helical operator.” However, the helical operator of Ref. [40] arises in the context of the Navier–Stokes equations and is of course different from the one we consider.

Before numerically probing the issue of conditioning for the operators above, let us present a heuristic argument as to why IPC achieves improved conditioning in certain ODE settings, including as special cases those of immediate interest to us. Suppose we have the equation

$$\frac{d^2}{dr^2} \xi + \frac{d}{dr} s_1(r) \xi + s_0(r) \xi = g, \tag{83}$$

where the $s_1(r)$ and $s_0(r)$ are polynomials and the equation must be supplemented with appropriate boundary conditions. Here, one should think of $\xi(r)$ as a mode, either $\hat{\psi}_k(r)$ or $\hat{\psi}_{\ell m}(r)$. Performing the IPC technique on the ODE (83), we pass from a spectral approximation of the equation based on differentiation matrices,

$$[D^2 + DS_1(A) + S_0(A)] \xi = g, \tag{84}$$

to the following matrix representation:

$$[I_{[2]} + B_{[2]} S_1(A) + B_{[2]}^2 S_0(A)] \xi = B_{[2]}^2 g, \tag{85}$$

with two available rows of zero in which to put τ -conditions specifying the boundary conditions to be applied. In either representation, $S_1(A)$ and $S_0(A)$ are the matrices which correspond to multiplication by $s_1(r)$ and $s_0(r)$ in the Chebyshev basis. In the passage from (84) to (85), the structure of the coefficient matrix is markedly improved. Indeed, in (84) the coefficient matrix features poorly conditioned dense matrices, save for $S_0(A)$ which is sparse and banded. Moreover, the derivative operators are unbounded as N grows. However, in (85) the coefficient matrix is a perturbation of the identity featuring banded matrices, and the integration matrices remain bounded as N grows.

Let \mathcal{L}^k represent the matrix – built from Chebyshev differentiation matrices – which represents L^k , $\mathcal{B}\mathcal{L}^k$ the corresponding preconditioned matrix, and likewise for \mathcal{P}^k and $\mathcal{B}\mathcal{P}^k$. Boundary conditions must be inserted into these matrices appropriately, and we will return to this point shortly. Our goal is to compute and compare condition numbers for the various operators and boundary conditions above. Given a matrix A , we shall work with the standard condition number $\kappa(A) = \|A\|_2 \|A^{-1}\|_2$ belonging to the matrix 2-norm.

Recall that for the radiation p, q boundary conditions we must consider consecutive blocks, since sin and cos modes of the same wave number are coupled. Suppose we consider the submatrix associated with the two consecutive blocks associated with Fourier wave number k . The unpreconditioned matrix for the helical operator with D–R boundary conditions is then

$$\left[\begin{array}{c|c} \mathcal{L}^k & \mathbf{0} \\ \delta^- & \mathbf{0} \\ p\delta^+ & R\nu^+ + q\delta^+ \\ \hline \mathbf{0} & \mathcal{L}^k \\ \mathbf{0} & \delta^- \\ R\nu^+ + q\delta^+ & -p\delta^+ \end{array} \right], \tag{86}$$

with the understanding that the boundary conditions have been written over the last two rows of \mathcal{L}^k in each block. However, as we have seen in (66), the preconditioned operator is

$$\left[\begin{array}{c|c} \delta^- & \mathbf{0} \\ p\delta^+ & R\nu^+ + q\delta^+ \\ \mathcal{B}\mathcal{L}^k & \mathbf{0} \\ \hline \mathbf{0} & \delta^- \\ R\nu^+ + q\delta^+ & -p\delta^+ \\ \mathbf{0} & \mathcal{B}\mathcal{L}^k \end{array} \right], \tag{87}$$

now with the understanding that the boundary conditions have been written over the first two rows of each block.

As an example of the various other matrices stemming from the Poisson operator, consider the preconditioned matrix with Dirichlet–Neumann boundary conditions,

$$\left[\begin{array}{c|c} \delta^- & \mathbf{0} \\ \nu^+ & \mathbf{0} \\ \mathcal{B}\mathcal{P}^k & \mathbf{0} \\ \hline \mathbf{0} & \delta^- \\ \mathbf{0} & \nu^+ \\ \mathbf{0} & \mathcal{B}\mathcal{P}^k \end{array} \right]. \tag{88}$$

Here we consider this decoupled direct sum of blocks only to make direct comparison with the matrices above. The condition number of the last matrix is of course the same as the condition number of either of its blocks.

For our first investigation of condition numbers we take the Poisson operator associated with D–D boundary conditions. We assume $k = 3$, a domain $[1, 3]$, and consider truncations $N = 8, 16, \dots, 256$, where each decoupled block is $(N + 1)$ -by- $(N + 1)$. We have computed the conditions numbers for these truncations in MATLAB, and recorded our results in Table 2. Owing to the boundary conditions which are inserted into the zeroed rows of the preconditioned operator, we expect growth in the 2-norm condition number for our spectral matrix representation of the preconditioned operator as the truncation size N increases. The issue is significant for boundary conditions involving a Neumann row vector, such as ν^+ , whose entries grow like j^2 with column location j . For $k = 2$ and $\Omega = 0.1$, Table 3 lists preconditioned (top) and unpreconditioned (bottom) condition numbers for the short domain $[5, 15]$. Table 4 lists the same numbers for the larger and more realistic domain $[5, 150]$. For the helical operator the short-domain table shows that our preconditioning is only advantageous initially. As the truncation size N increases, the τ -conditions appear to have a dominant

effect. (Since they mix the two blocks, for the helical operator such conditions would seem especially influential.) We loosely observe that the IPC technique only directly concerns the conditioning of that part of the coefficient matrix which stems from the operator itself, and not the boundary conditions *per se*. Over the whole range of truncations the large-domain table indicates that the condition number of the coefficient matrix is chiefly determined by the operator rather than the τ -conditions. For this large domain we do discern the advantage of preconditioning the helical operator, although not as pronounced as for the Poisson operator. As a way of obtaining a more drastic improvement in conditioning of the helical operator, one might explore either suitable equilibration tailored to the h_r norms discussed in [40] or enforcement of the radiation boundary conditions via a penalty method [44].

In Figs. 3 and 4 we plot the singular values corresponding to the Helical (D–R) columns in Tables 3 and 4, respectively. In Fig. 4, say, the top plot depicts the singular value distributions corresponding to the preconditioned matrices, for which condition numbers have been listed in the top leftmost column of Table 4. The bottom plot shows the same data for unpreconditioned matrices, for which condition numbers have been listed in the bottom leftmost column of Table 4. In either plot, each curve corresponds to one entry of a column, where one should note that the number of plotted singular values is twice the listed truncation (each matrix consists of two modes with one block for each). Therefore, the longest curves corresponds to the 512 truncation, and so forth. Notice that the singular value distributions for preconditioned matrices are more clustered.

4.2. Numerical results for the linear and nonlinear 2d HRWE

Throughout this section we examine the 2d model, and, unless explicitly stated otherwise, all matrix inversions have been carried out with `dgesv`. Those inversions carried out with `dgesvx` will be clearly indicated.

Table 2
2-Norm condition numbers

Truncation	Preconditioned	Unpreconditioned
8	12.5839 (12.6)	3.7740e+03(3.774e+03)
16	20.6559 (20.6)	5.5327e+04(5.533e+04)
32	31.8782 (31.9)	8.4273e+05(8.427e+05)
64	51.1595 (51.2)	1.3140e+07(1.314e+07)
128	86.9027 (90.0)	2.0748e+08(2.075e+08)
256	156.0339 (156.)	3.2976e+09(3.298e+09)

D–D Poisson operator on [1, 3] with $k = 3$. The numbers in parenthesis are those listed in Tables 1 and 4 of [40] for the same problem. That reference's h_0 condition number is the 2-norm condition number.

Table 3
2-Norm condition numbers

Truncation	Poisson (D–D)	Poisson (D–N)	Helical (D–R)
32	157.2452	8.0745e+03	9.3845e+04
64	157.9248	2.0750e+05	2.7642e+06
128	159.5918	6.3311e+06	8.4756e+07
256	164.6046	1.9820e+08	2.6537e+09
512	180.0820	6.2727e+09	8.3984e+10
32	8.8450e+05	6.3270e+06	8.6394e+05
64	1.3791e+07	9.8651e+07	1.3471e+07
128	2.1776e+08	1.5577e+09	2.1270e+08
256	3.4610e+09	2.4757e+10	3.3805e+09
512	5.5191e+10	3.9479e+11	5.3907e+10

Preconditioned (top) and unpreconditioned (bottom) operators on [5, 15] for $k = 2$ and $\Omega = 0.1$.

Table 4
2-Norm condition numbers

Truncation	Poisson (D–D)	Poisson (D–N)	Helical (D–R)
32	1.4634e+04	1.0922e+06	1.9905e+05
64	1.4640e+04	1.0924e+06	2.0635e+05
128	1.4650e+04	1.0949e+06	3.2270e+05
256	1.4673e+04	1.1884e+06	3.0785e+06
512	1.4731e+04	4.2550e+06	9.0043e+07
32	3.9154e+05	2.9145e+07	3.4677e+05
64	6.1040e+06	4.5437e+08	5.4046e+06
128	9.6378e+07	7.1742e+09	8.5335e+07
256	1.5318e+09	1.1402e+11	1.3562e+09
512	2.4426e+10	1.8182e+12	2.1627e+10

Preconditioned (top) and unpreconditioned (bottom) operators on [5, 150] for $k = 2$ and $\Omega = 0.1$.

Consider the linear equation

$$L\psi = Q \frac{\delta(r - x_H)}{x_H} [\delta(\varphi - \pi) - \delta(\varphi)]. \tag{89}$$

where L is the HRWE operator and the inhomogeneity is comprised of two equal but opposite strength δ -functions, one located at $(x, y) = (x_H, 0)$ and the other at $(x, y) = (-x_H, 0)$. Assuming that exact outgoing

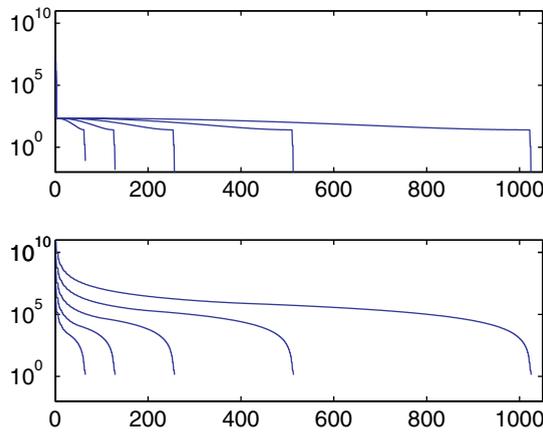


Fig. 3. Singular values for last columns of Table 3. For operators on [5, 15] with $k = 2, \Omega = 0.1$.

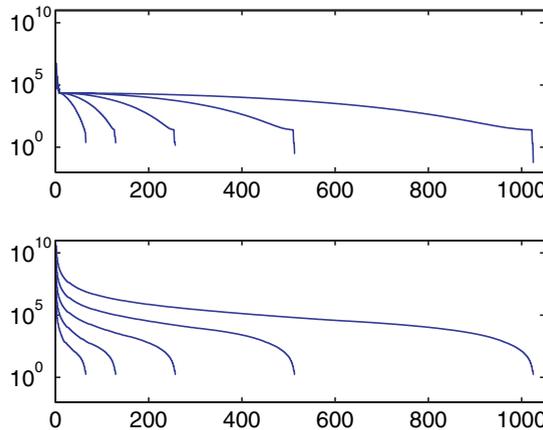


Fig. 4. Singular values for last columns of Table 4. For operators on [5, 150] with $k = 2, \Omega = 0.1$.

radiation conditions are placed at $r = R = r_{\max}$, we may express the exact solution to (89) as the following infinite series [8]:

$$\psi(r, \varphi) = - \sum_m QJ_m(m\Omega r_{<}) [N_m(m\Omega r_{>}) \cos(m\varphi) + J_m(m\Omega r_{>}) \sin(m\varphi)], \quad (90)$$

where the sum is over $m = 1, 3, 5, \dots$ and $r_{>}$ is the larger of x_H and r and similarly for $r_{<}$. In what follows we always take $x_H = 2$, $\Omega = 0.1$, and $Q = 1$. Although we shall not carefully study the pointwise convergence of the series (90), we note that it is poorly convergent whenever r is close to x_H . In using (90) to generate approximations to field values ψ for the experiments below, we have at times used 30 digit precision in MAPLE and included thousands of terms in order to sum the series with double precision accuracy.

In order to make a direct comparison between the series (90) and numerical solutions, we now consider the following three *linear* and *homogeneous* problems.

- (i) *Elliptic problem on inner annulus.* Let ρ represent the radial coordinate as measured from the source point $(x_H, 0)$, so that the inner annulus corresponds to $\rho_{\min} \leq \rho \leq \rho_{\max}$. We consider the homogeneous problem $L\psi = 0$, using the series (90) to place Dirichlet boundary conditions at $\rho = \rho_{\min}$ and $\rho = \rho_{\max}$.
- (ii) *Mixed problem on outer annulus.* We consider $L\psi = 0$ on the outer annulus corresponding to $r_{\min} = \varepsilon \leq r \leq R = r_{\max}$, using the series (90) to place a Dirichlet boundary at $r = \varepsilon$ and adopting exact outgoing conditions at $r = R$. In this setting, we assume $\varepsilon > x_H$ (where ε need not be small). Again, the light circle lies between ε and R .
- (iii) *Mixed problem on two center domain.* In this setting we consider the linear problem $L\psi = 0$ on the two center domain depicted in Fig. 1. We use the series (90) to place Dirichlet boundary conditions at $\rho = \rho_{\min}$ (for both inner “holes”) and adopt exact outgoing conditions at $r = R$.

In each case, we will generate a numerical solution, and then compare it with the series via various choices of error measure. Our numerical experiments for cases (i) and (ii) are meant to facilitate better parameter choices for scenario (iii) experiments.

4.2.1. Linear elliptic problem on inner annulus

For the numerical experiment in (i) above, we use the series to generate Dirichlet data at $\rho_{\min} = 1$ and $\rho_{\max} = 2$. With a double Chebyshev–Fourier expansion and the described integration preconditioning, we numerically solve $L\psi = 0$ on this inner annulus. From the spectral solution we then generate numerical values for $\psi(\rho_{\text{mid}}, \theta)$, where $\rho_{\text{mid}} = 1.5$ and θ is sampled on a fine grid. This ρ_{mid} profile is then compared with the analogous profile given by the exact series (90). All profiles are depicted in Fig. 5, and Table 5 lists the associated errors for the ρ_{mid} profile, where N_H denotes the number of radial Chebyshev elements, and M_H the number of angular Fourier elements.

In this table, as in all remaining tables of the paper, the meaning of each of the error measures is as follows. The absolute sup error is simply the magnitude of the largest error found in the computed solution. For our spectral solutions this means that the solutions must be computed and compared on an evaluation grid. For Table 5 the grid was taken to be 1024 evenly spaced points of θ , at ρ_{mid} . The relative sup error is the absolute sup error divided by the maximum value found on the evaluation grid. The absolute rms error is found by taking the sum of the squares of all absolute errors on the evaluation grid, then dividing that sum by the number of terms that have been added, and finally by taking the square root. The relative rms error is the absolute rms error divided by the rms value of the solution on the evaluation grid. To evaluate the exact series on the evaluation grid, we have first evaluated it to high accuracy with MAPLE on a (uniformly spaced) Fourier grid of 128 points, subsequently using the corresponding Fourier interpolation onto the finer grid. Direct evaluation of the series on the 1024 point grid proved impractical.

4.2.2. Linear mixed problem on outer annulus

For this problem, in the region A of Fig. 1, we choose $\varepsilon = 4.5$ and either $R = 50$ or $R = 150$. Since $\Omega = 0.1$, the light circle always lies within the corresponding annulus. Using (90) to generate Dirichlet data on the circle $r = \varepsilon$ is not problematic, since $4.5 > 2$ and so the series converges reasonably well. We have used the FORTRAN

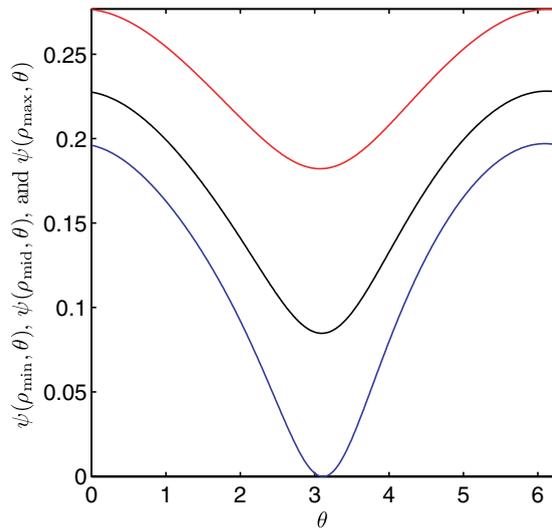


Fig. 5. Constant- ρ sections of $\psi(\rho, \theta)$. The figure depicts the solution (90) expressed in terms of the local polar coordinates (ρ, θ) about the hole. The top curve corresponds to $\rho_{\min} = 1$, the middle to $\rho_{\text{mid}} = 1.5$, and the bottom to $\rho_{\max} = 2$.

Table 5
Errors for the inner annulus problem

N_H, M_H	Absolute sup error	Relative sup error	Absolute rms error	Relative rms error
8, 16	1.091e-05	4.782e-05	5.159e-06	2.964e-05
12, 32	2.551e-09	1.118e-08	1.226e-09	7.040e-09
16, 48	7.646e-13	3.352e-12	3.697e-13	2.124e-12
20, 64	3.192e-16	1.399e-15	1.352e-16	7.766e-16

routines `RJBESL` and `RYBESL` from `netlib` to build up the series (and also `MATLAB`'s Bessel routines for error analysis). In Table 6 we list error measures in the outer annulus for various choices of N_A (number of radial Chebyshev elements) and M_A (number of Fourier elements). In each case, for a fixed choice of M_A we have experimented to (roughly) find the smallest value N_A which achieves the best possible errors. Fig. 6 depicts the numerical solution $\psi(r, \varphi)$ and the error $\Delta\psi(r, \varphi)$, the difference between the numerical and series solutions, corresponding to the bottom table's second to last line. Notice that most of the error is concentrated near the inner boundary $r = \varepsilon$, and this region determines the supremum errors in the tables. For a fixed particular number of angular Fourier elements, these values are, save for the last line, the same for both tables.

Table 6
Errors for the outer annulus problem

N_A, M_A	Absolute sup error	Relative sup error	Absolute rms error	Relative rms error
20, 11	5.433e-04	3.114e-03	3.223e-05	9.171e-01
38, 31	1.045e-07	5.991e-07	3.001e-09	2.142e-04
56, 51	3.028e-11	1.736e-10	6.767e-13	4.831e-08
80, 81	6.661e-16	3.819e-15	8.937e-17	6.380e-12
32, 11	5.433e-04	3.114e-03	3.885e-05	9.979e-01
76, 31	1.045e-07	5.991e-07	5.601e-09	2.223e-03
124, 51	3.028e-11	1.736e-10	8.612e-13	3.419e-07
200, 81	9.853e-16	5.649e-15	2.981e-16	1.183e-10

The top table corresponds to $R = 50$ and the bottom to $R = 150$. These numbers have been extracted from a 1024×1024 uniform grid.

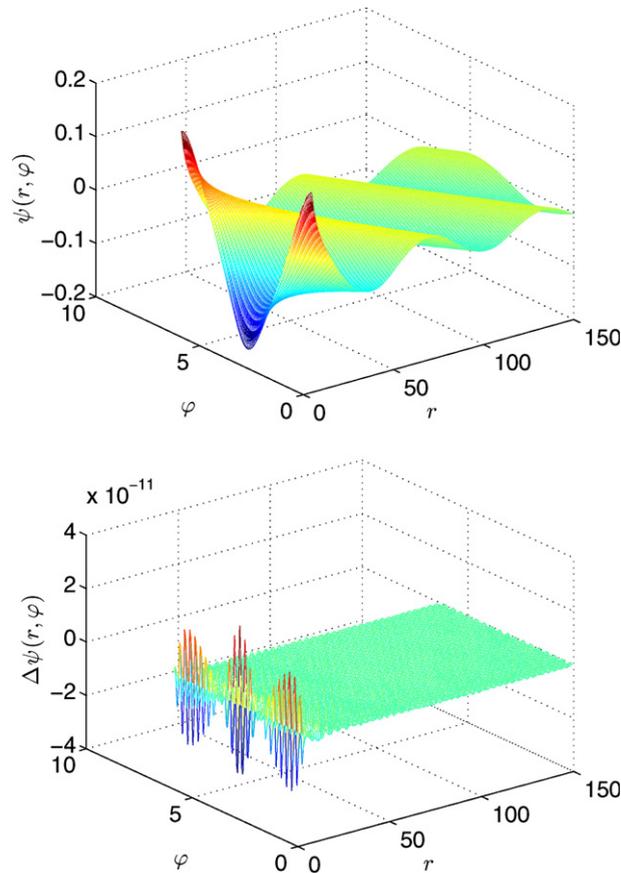


Fig. 6. Outer annulus problem. The plots depict the numerical solution and error for $N_A = 124$ and $M_A = 51$.

We remark that owing to the quadratic growth inherent in the our radiation boundary conditions (based on Neumann vectors v^+), 80 or 200 radial elements is likely excessive. We should prefer to further divide A into sub-annuli all glued together, with fewer Chebyshev elements taken on each sub-annulus. In future 3d work with numerical shells we plan to explore this possibility.

4.2.3. Linear mixed problem on two center domain: iteration between the elliptic region and outer annulus

This numerical experiment tests the full multidomain spectral code. We use the series (90) summed in extended precision with MAPLE to place Dirichlet conditions on the inner holes (in practice, owing to the symmetry of the problem, on one inner hole H). We place radiation conditions at the outer boundary $r = R$ of the outer annulus. We take the hole H and all rectangles 1–8 of Fig. 1 as one glued domain, confined within a large square with side $2L = 10$. The outer annulus A is treated as a separate domain, and we iterate between solves on the hole-plus-rectangles domain and the outer annulus. More precisely, we first solve the hole-plus-rectangles configuration using the series (90) to place Dirichlet conditions on the inner boundary $\rho = \rho_{\min}$ of the hole and adopting Dirichlet zero conditions on the free edges of rectangles which overlap the outer annulus A . This multidomain numerical solution then provides inner Dirichlet conditions at $r = \epsilon = 4.5$ for the solve on the outer annulus. The solution for the outer annulus, in turn, provides a numerical solution from which we obtain corrected Dirichlet conditions for the relevant free edges of rectangles. While this iteration between solves on the hole-plus-rectangles domain and outer annulus is under way, we monitor the rms error between the inner profile $\psi(\epsilon, \varphi)$ from the outer annulus at the current and previous iterate. The iteration is stopped once this inner boundary rms error ceases to decrease, typically after 100 or so iterations. We could of course glue the outer annulus to the hole-plus-rectangles domain, and indeed we do so later when examining the non-

linear model problem, but here we use the described iteration, since we wish to demonstrate its stability. Table 7 lists error measures associated with this experiment with the outer boundary taken as $R = 50$. In the table the choices for N_H , M_H and N_A , M_A stem from our earlier experiments on annulus domains. For simplicity, we have chosen the same numbers N , M for the elements associated with the double Chebyshev expansion on each of the rectangular domains. For the table's last line, we have used `dgesvx` in lieu of `dgesv`, for the hole-plus-rectangles domain only. Without the extra accuracy afforded by `dgesvx` for the elliptic region, the numbers in the last line would be only marginally better than those in the next to last. See the final paragraph in Section 3.4.1 for further germane comments.

4.2.4. Nonlinear mixed problem on two center domain

For our final experiment in 2d, we examine the homogeneous ($g = 0$) nonlinear equation (76). For nonlinear solves we have always chosen to work with the fully glued two center domain ($H +$ all rectangles $1 - 8 + A$). Our numerical experiment with the nonlinear model is to repeat the mixed boundary value problem carried out for the linear model on the full two center domain. In particular, we use the same Dirichlet conditions, from (90), at $\rho = \rho_{\min}$. A major difference is that now we solve simultaneously for the unknowns in all subdomains rather than by back and forth iteration with the outer annulus as we did for the linear problem. Since there is no exact solution available for comparison, we simply compare solutions with different resolutions. As before, the comparison is carried out using output associated with the outer annulus on a 1024×1024 uniform grid. We take $R = 50$, $\eta = -10$, and $\psi_0 = 0.25$. With a numerical solution to the nonlinear problem, we may naively compute its “dipole strength” via Fourier transform of the solution restricted to the outer boundary $r = R$. The nonlinear term corresponding to our parameter choices results in a dipole strength which differs from dipole strength for the linear model ($\eta = 0$) by about 18%, and this difference measures the nonlinear term's strength. Fig. 7 depicts solutions for both the linear and nonlinear models in the vicinity of the holes, showing that even to the eye these solutions differ. Error measures are collected in Table 8.

4.3. Preliminary results for the linear 3d HRWE

We briefly describe two numerical experiments involving the 3d HRWE and meant to suggest that a multidomain spectral approach is viable for 3d PSW approximations, at least insofar as the post-Minkowski scenario is concerned and possibly beyond. The first experiment involves the 3d mixed-problem on an outer spherical shell, analogous to the outer annulus problem studied in Section 4.2.2. The second experiment is a crude two-domain model consisting of an inner cube (lying within the elliptic region and on which we place explicit sources) and an outer spherical shell (intersecting the type-changing cylinder). This single cube serves as a crude substitute for all the subdomains depicted in Fig. 2 (as well as the inner spherical shell not depicted). We would prefer an experiment in which a single cylinder replaced all these subdomains, but as yet have not developed a library of spectral routines for solid cylinders. Our simple two-domain model indicates that iteration between an inner elliptic region and an outer spherical shell is stable. Such an iteration also appears to be quite stable in the 2d setting.

Table 7
Errors for the two center domain problem

N_H, M_H	N, M	N_A, M_A	Absolute sup error	Relative sup error	Absolute rms error	Relative rms error
8, 16	8, 8	20, 11	4.540e-04	2.600e-03	5.362e-05	9.674e-01
12, 32	20, 20	38, 31	9.908e-08	5.680e-07	2.988e-09	2.133e-04
16, 48	30, 30	56, 51	3.023e-11	1.733e-10	7.907e-13	5.645e-08
20, 64	38, 38	80, 81	1.402e-15	8.035e-15	2.291e-16	1.636e-11

These error measures have been extracted from the outer annulus only (again with a 1024×1024 uniformly spaced grid), although the numerical solution is generated on the entire two center domain. The last line of this table has been obtained with `dgesvx` (see text). Compare this table with the top table in Table 6.

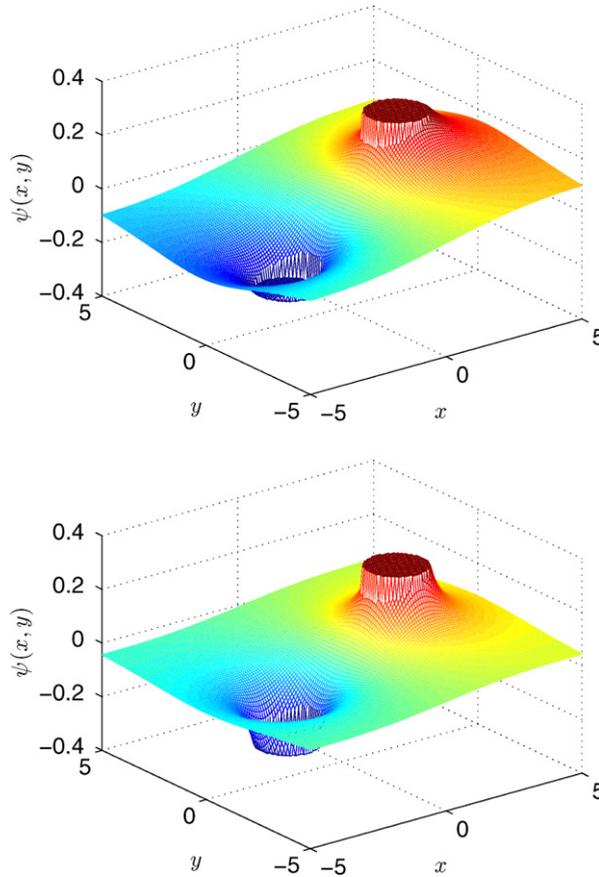


Fig. 7. Comparison of solutions. The plots depicts numerical solutions for the linear (top) and nonlinear (bottom) models. In each case the solution is not known in the immediate neighborhood of the holes, whence on these neighborhoods we have set the solution to a constant value in order to achieve nice plots.

Table 8
Errors for the nonlinear two center domain problem

N_H, M_H	N, M	N_A, M_A	Absolute sup error	Relative sup error	Absolute rms error	Relative rms error
8, 16	8, 8	20, 11	1.096e-03	1.211e-02	1.371e-04	9.976e-01
14, 36	20, 20	42, 39	1.152e-07	1.273e-06	1.140e-08	1.209e-03
18, 54	30, 30	64, 65	3.130e-10	3.459e-09	5.991e-12	6.354e-07

As with the linear model experiment, these error measures are extracted from the outer annulus only. Error measures for the first two lines are taken with respect to the numerical solution corresponding to the next line. Error measures in the last line are taken with respect to the numerical solution which corresponds to incrementing all parameters by 2.

4.3.1. Mixed problem on outer shell

For the first experiment we define the outer spherical shell via $4 = \varepsilon \leq r \leq R = 50$, and consider the following exact series solution to the homogeneous 3d HRWE:

$$\psi(r, \theta, \varphi) = \sum_{\ell=0}^{\ell_{\max}} c_{\ell 0} Y_{\ell 0}(\theta, 0) \frac{a^\ell}{r^{\ell+1}} + \sum_{\ell=1}^{\ell_{\max}} \sum_{m=1}^{\ell} c_{\ell m} Y_{\ell m}(\theta, 0) j_\ell(m\Omega a) [j_\ell(m\Omega r) \sin(m\varphi) + n_\ell(m\Omega r) \cos(m\varphi)], \tag{91}$$

here with $r > a = 1$ and $\Omega = 0.2$. The $c_{\ell m}$ are random expansion coefficients, obeying $|c_{\ell m}| \leq 1$ and drawn from a uniform distribution, $Y_{\ell m}(\theta, \varphi)$ is a scalar spherical harmonic, and $j_\ell(z)$ and $n_\ell(z)$ are spherical Bessel func-

tions. Notice that $\varepsilon < |\Omega|^{-1}$, whence the inner surface of the shell lies within the elliptic region, as discussed in Section 2.2. The form of this series is motivated by the rather more physical series [18]

$$\begin{aligned} \psi(r, \theta, \varphi) = & -2K \sum_{\ell=0}^{\infty} \frac{1}{2\ell+1} Y_{\ell 0}(\pi/2, 0) Y_{\ell 0}(\theta, 0) \frac{a^\ell}{r^{\ell+1}} + 4K\Omega \sum_{\ell=2}^{\infty} \sum_{m=2,4,6,\dots} m Y_{\ell m}(\pi/2, 0) Y_{\ell m}(\theta, 0) \\ & \times j_\ell(m\Omega a) [j_\ell(m\Omega r) \sin(m\varphi) + n_\ell(m\Omega r) \cos(m\varphi)] \end{aligned} \tag{92}$$

corresponding to two equal point charges, that is

$$g(r, \theta, \varphi) = K \frac{\delta(r-a)}{a^2} \delta(\cos \theta) [\delta(\varphi) + \delta(\varphi + \pi)] \tag{93}$$

in Eq. (10). The choice here of like signs for the point sources would make the problem similar to the gravitational problem, with radiation dominated by the quadrupole mode. (In the 2d case we chose equal and opposite signs to avoid the bothersome logarithmic 2d monopole.) We could of course base our test on a truncation of (92), but will instead work with (91) in order to test all modes defined for a given choice of ℓ_{\max} .

We use the series (91) to seed inner boundary conditions at $r = \varepsilon$, which would also be possible for (92), since the locations $r = a, \theta = \pi/2, \varphi = 0, \pi$ do not lie in the shell, and so the series is convergent for $r > a$. This requires a spherical-harmonic decomposition, which we perform with the NCAR routines SPHEREPACK by Adams and Swarztrauber [47]. For this example, we have coded our spectral representation of the 3d HRWE as a matrix–vector multiply. As we do not explicitly form the matrix, we solve the resulting linear system using the iterative method GMRES [48]. Our matrix–vector multiply involves integration preconditioning in the radial variable only. We use the series above with $\ell_{\max} = 7$, which corresponds to $1 + 3 + 5 + 7 + 9 + 11 + 13 + 15 = 64$ modes. The capacity to perform spectral analysis and synthesis (transform and inverse transform) then requires that a corresponding physical grid has an angular resolution of at least $N_\theta \times N_\varphi = 8 \times 15 = 120$ points. In anticipation of using such analysis and synthesis when dealing with nonlinearities, we therefore work with a set of unknowns corresponding to $N_r \times N_\theta \times N_\varphi = 60 \times 8 \times 15 = 7200$, although $60 \times (120 - 64) = 3360$ should be zero. Since only the remaining $60 \times 64 = 3840$ unknowns should be nonzero, as our linear solver we use GMRES, taking the number of iterations $M_{\text{itr}} \leq 3840$. As an example, with $M_{\text{itr}} = 2000$ we produce a numerical solution from which we plot and analyze various cross sections. Over the equatorial plane the numerical solution has a supremum error of 1.194×10^{-9} . For $M_{\text{itr}} = 1600$ the corresponding supremum error is 1.038×10^{-6} .

4.3.2. Mixed problem on two coupled domains

As the second experiment, we consider essentially the same problem described in Ref. [11], that is the linear inhomogeneous 3d HRWE of (10) with two Gaussian sources,

$$g(r, \theta, \varphi) = Qe^{-\kappa[(x-x_H)^2+(y-y_H)^2+z^2]} + \mu Qe^{-\kappa[(x+x_H)^2+(y+y_H)^2+z^2]} \tag{94}$$

Our parameter choices are $Q = 100, \kappa = 3$, with the location of the “hole” (source) at $(x_H, y_H) = (1, 0)$ in the $z = 0$ plane. Furthermore, we choose $\mu = -0.25$ in order to ensure that the stability of our scheme is not predicated on symmetry. We determine the HRWE operator by $\Omega = 0.2$. Our two basic subdomains are an inner cube and an outer shell, respectively determined by $[x_{\min}, x_{\max}] = [y_{\min}, y_{\max}] = [z_{\min}, z_{\max}] = [-3, 3]$ and $[\varepsilon, R] = [2.8, 100]$.

To solve this problem numerically, we work with the two domains as decoupled. On the cube we use a pseudospectral method (point values as unknowns) based on GMRES, with the derivatives in the HRWE operator coded as a matrix–vector multiply. We currently employ no preconditioning whatsoever on the cube in our simple model, and so are limited in what resolutions we can obtain. For the experiment we have taken a $N_x \times N_y \times N_z = 30 \times 30 \times 30$ truncation. For the outer shell, we again use GMRES but with the purely spectral method just described in Section 4.3.1. Therefore, we have adopted a hybrid pseudospectral/spectral approach. Our resolution in the annulus corresponds to a grid size of $N_r \times N_\theta \times N_\varphi = 62 \times 5 \times 9$. Our GMRES-based experiment requires less than 5% of the combined storage required to perform an individual direct solve for each domain, although it does not achieve high accuracy.

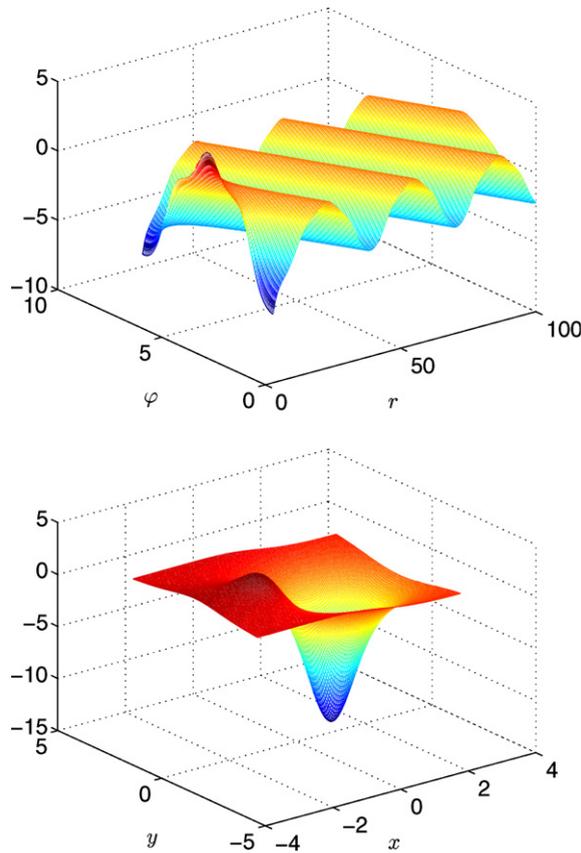


Fig. 8. Cross sections of the 3d numerical solution. The top plot depicts the equatorial cross section on the outer shell, with multiplication by r enhancing $1/r$ fall off. The bottom plot depicts the equatorial cross section on the inner cube.

Starting with the cube, we iterate between the solves. That is to say, we first solve the problem on the cube with Dirichlet-zero boundary conditions. Via the global interpolation then available from the cube solution, we seed inner Dirichlet data $\psi|_e$ for the shell solve. Since we are solving the 3d HRWE on a closed ball (cube plus shell) with no inner boundaries, the solution is only determined up to a free constant. We fix this constant by setting to zero the $\ell = 0$ mode of $\psi|_e$. The solution on the shell then provides corrected Dirichlet data for another cube solve, and so on. The iteration proceeds until the rms error between successive seed data sets $\psi|_e$ ceases to decrease (about 60 iterations). Fig. 8 shows the resulting numerical solution. For this low resolution experiment, the depicted outer shell solution has an absolute supremum error (as measured against a second numerical solution obtained with slightly more resolution on both the cube and shell, but with the same number of angular modes) measuring 1.478×10^{-2} over the equatorial plane. Although we have not considered large values of Ω , we note that the iteration appears to be completely stable and convergent. An instability could set in at higher resolutions. Nevertheless, we find the results from this model 3d code as well as similar iteration in the 2d setting encouraging.

5. Summary, conclusions, and outlook

The motivation for this work arose in problems of gravitational waves and black holes. In a late stage of inspiral the interaction of the holes can be strongly affected by nonlinearities, while the radiation reaction due to gravitational wave emission is weak. It is useful to approximate such a system as having helical symmetry, in which the sources and fields rotate rigidly. Two aspects of this physical problem are then important to the mathematical methods explored in this paper: (i) “Helical reduction,” that is, imposing helical symmetry, converts the PDEs describing the fields into a system of helically reduced wave equations (HRWEs), a mixed sys-

tem that is elliptic inside, and hyperbolic outside a type changing surface. (ii) Nonlinearities are significant only interior to this type changing surface. In this paper we have divided the topologically nontrivial domain of the problem into basic subdomains. Crucially, only a single subdomain contains the type changing surface and the hyperbolic region exterior to it, the outer annulus/shell subdomain (annulus for 2d, shell for 3d). The other subdomain or subdomains span an elliptic problem. We have therefore confined aspect (i) to a topologically simple outer subdomain, and aspect (ii) to the subdomains spanning an elliptic problem.

Sections 2 and 3 of this paper have focused on the model problem of the 2d HRWE for a linear scalar field, though some generalization to 3d and to nonlinear problems has been included. Likewise, while the numerical results in Section 4 have predominately focused on the 2d linear scalar problem, we have also included some results for a nonlinear 2d scalar model and for a linear 3d scalar model. To obtain our numerical results we have necessarily made many choices in the details. This was also case in Section 3; clarity of presentation required some specificity in the adopted methods. It is, however, important to distinguish the full set of choices made from what we believe are robust conclusions that transcend many of the choices. In the present section, we therefore remark on what, in our view, are the central lessons of this work.

The choice to use a spectral (or a pseudospectral) method does not need justification for an elliptic problem. Our mixed problem inherits much of that justification. In particular, we are solving a boundary value problem, so a large set of equations must be solved simultaneously. For such problems spectral methods, which reduce the number of unknowns, are particularly valuable and widely used. A less general issue is the nature of the decomposition of the two center domain of the problem into subdomains. The choice illustrated in Fig. 1 is very much driven by the convenience of rectangular and annular subdomains, for on such regions spectral methods are well developed. For analogous reasons, the similar decomposition in Fig. 2 would be used for the 3d problem.

What we wish to emphasize here is that one should make a distinction between the choices for the outer annulus/shell and those for the inner region, the subdomain or concatenation of subdomains for the elliptical region. (One could also divide the outer shell/annulus into further concentric sub-shells/sub-annuli, but this possibility only adds further technical complexity.) It is quite conceivable that the optimal choice for the outer region is not the same as that for the inner region. One might, for example, treat the inner elliptic problem by whatever method most conveniently or efficiently generates a set of equations for the interior, and then use the spectral method of this paper only for the outer region with its mixed PDEs. Even more freedom is available if the inner and outer regions are solved separately, with an iteration between them, as was done in the second numerical experiment of Section 4.3. In this case, one could use separate iterations to solve within each region, each within the outer loop iteration that feeds values from the inner to the outer region and *vice versa*. We will return to this possibility in a moment.

We have made the choice of using integration preconditioning (IPC) in all subdomains. Whether or not this is an optimal choice remains unresolved. IPC is known to be efficient in reducing the condition number of the coefficient matrix for certain types of ODE boundary value problems. Our study of a novel boundary value problem, one heretofore not considered in the context of spectral methods, has to some extent further elucidated the utility of IPC in such ODE settings. The results in Table 4 demonstrate that IPC reduces the condition number of the matrix stemming from a problem not considered in our primary reference [40]. Moreover, our work demonstrates that IPC affords a relatively direct way to handle a multidomain scenario in two dimensions. (See also [45] for an application of IPC in a 1d multidomain setting.) However, it remains to be seen if its advantages will carry over to 3d and to PDEs significantly different from those of our test problems.

Coming to grips with whether or not IPC is an appropriate method for 3d PSW problems, requires that we ask some broader background questions. In particular, is it important to reduce the condition number of the coefficient matrix? In Section 4, solutions to 2d problems were found by straightforward Gaussian elimination for which conditioning issues are not as critical. For 3d problems, however, we almost certainly expect to use an iterative Krylov method, and GMRES in particular. The efficiency of such iterative solvers is known to depend sensitively on condition number. But such methods are also sensitive to other aspects of the coefficient matrix, in particular to the distribution of eigenvalues [46]. In any case, owing to the relevant domain decomposition (Fig. 2) for a 3d problem, insofar as the elliptic problem is concerned the relevant issue is the effectiveness of IPC for the 3d HRWE on cylindrical and (inner) spherical-shell domains. As noted in [40,50],

special alternate techniques are often necessary for polar and cylindrical subdomains which contain a center or central axis.

We turn now from the open question of whether IPC could be effectively applied to all, or almost all, subdomains of a 3d PSW problem, to the rather clearer question of IPC for the outer region itself. For a scheme in which inner and outer regions are separately solved, we have some enthusiasm for the advantages of IPC. Our enthusiasm stems from the fact that in the outer region the Fourier, or spherical harmonic modes do not mix for a linear problem. This nonmixing originates in the symmetry of the problem, the fact that the annulus or shell is concentric with the coordinate surfaces of constant radius. The situation is closely analogous to the reason that separation of variables applied to this region leads to a set of decoupled radial ODEs. In line with theoretical and numerical arguments presented in [40], our studies indicate that IPC is a promising way of solving the system of ODEs arising in the outer region of our problem.

These simple considerations do not apply, of course, if the problem is nonlinear. The nonlinearity will mix Fourier or spherical harmonic modes. However, the nonlinearities in the underlying physical problem are very weak in the outer region. (Here we are assuming that the inner surface of the outer region is chosen to be near the maximum radius that allows it to be inside the type changing surface and to have the needed overlap with the subdomains of the interior region.) As a result, there would seem to be number of possibilities for handling the nonlinearity. First, it is quite plausible that it is adequate to treat the problem as linear in the outer region. If it is not adequate to treat the nonlinearity as null, then one could exploit the fact that it is weak. The weakly nonlinear problem could be solved by iteration in which each iteration generates a “known” source term, and that source term is projected onto the angular mode basis before the system is solved. Just as would be the case with separation of variables, this procedure would not mix modes (after the projection of the known source is done) so the solution would again be highly efficient. The weakness of the source should guarantee that a few iterations would be sufficient to get the needed resolution. Finally, a Newton–Krylov method might be based upon use of the outer domain’s spectral transform and inverse transform (with all nonlinear operations performed in point space). For a weak nonlinearity the associated Jacobian should not significantly perturb the basic linear operator.

Our next steps with these methods will parallel those of Refs. [8–19]. The method used in those studies is, by its nature, incapable of high precision, whereas our approach is capable of the precision that is needed for some purposes (initial data for evolution codes and studies of radiation reaction for instance). In particular, we will next formulate 3d linear and nonlinear scalar model HRWE problems. We will then use the techniques of Ref. [18] to do linearized gravity, and those of Ref. [19] to do gravity with a post-Minkowski approximation. Finally, we will treat the problem in fully nonlinear general relativity.

Acknowledgments

For helpful conversations and clarifying remarks, we thank E. A. Coutsias, T. Hagstrom, J.S. Hesthaven, and particularly G. von Winckel. This research was supported by NSF Grants PHY 0514282 and PHY 0554367 (to UTB), ARO DAAD19-03-1-0146 (to UNM), and DMS 0554377 and DARPA/AFOSR FA9550-05-1-0108 (to Brown University).

References

- [1] F. Pretorius, Evolution of binary black-hole spacetimes, *Phys. Rev. Lett.* 95 (2005) 121101.
- [2] M. Campanelli, C.O. Lousto, P. Marronetti, Y. Zlochower, Accurate evolutions of orbiting black-hole binaries without excision, *Phys. Rev. Lett.* 96 (2006) 111101.
- [3] M. Campanelli, C.O. Lousto, Y. Zlochower, Spinning-black-hole binaries: the orbital hang-up, *Phys. Rev. D* 74 (2006) 041501.
- [4] M. Campanelli, C.O. Lousto, Y. Zlochower, B. Krishnan, D. Merritt, Spin Flips and Precession in Black-Hole-Binary Mergers, preprint [gr-qc/0612076](http://arxiv.org/abs/gr-qc/0612076), available at <http://lanl.arxiv.org/>.
- [5] J.G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, J. van Meter, Gravitational-wave extraction from an inspiraling configuration of merging black holes, *Phys. Rev. Lett.* 96 (2006) 111102.
- [6] M.A. Scheel, H.P. Pfeiffer, L. Lindblom, L.E. Kidder, O. Rinne, S.A. Teukolsky, Solving Einstein’s equations with dual coordinate frames, *Phys. Rev. D* 74 (2006) 104006.
- [7] B. Brügmann, J.A. González, M. Hannam, S. Husa, U. Sperhake, Calibration of Moving Puncture Simulations, preprint [gr-qc/0610128](http://arxiv.org/abs/gr-qc/0610128), available at <http://lanl.arxiv.org/>.

- [8] J.T. Whelan, W. Krivan, R.H. Price, Quasi-stationary binary inspiral: II. Radiation-balanced boundary conditions, *Class. Quantum Grav.* 17 (2000) 4895.
- [9] J.T. Whelan, C. Beetle, W. Landry, R.H. Price, Radiation-balanced simulations for binary inspiral, *Class. Quantum Grav.* 19 (2002) 1285.
- [10] R.H. Price, Binary inspiral: finding the right approximation, *Class. Quantum Grav.* 21 (2004) S281.
- [11] Z. Andrade, C. Beetle, A. Blinov, B. Bromley, L.M. Burko, I. M. Cranor, R. Owen, R.H. Price, Periodic standing-wave approximation: overview and three-dimensional scalar models, *Phys. Rev. D* 70 (2004) 064001.
- [12] B. Bromley, R. Owen, R.H. Price, Periodic standing-wave approximation: nonlinear scalar fields, adapted coordinates, and the eigenspectral method, *Phys. Rev. D* 71 (2005) 104017.
- [13] C.G. Torre, The helically-reduced wave equation as a symmetric-positive system, *J. Math. Phys.* 44 (2003) 6223.
- [14] C.G. Torre, Uniqueness of solutions to the helically reduced wave equation with Sommerfeld boundary conditions, *J. Math. Phys.* 47 (2006) 073501.
- [15] H. Pfeiffer, Initial data for black hole evolutions, Ph.D. Thesis, Cornell University, 2003.
- [16] H.P. Pfeiffer, L.E. Kidder, M.A. Scheel, S.A. Teukolsky, A multidomain spectral method for solving elliptic equations, *Comput. Phys. Commun.* 152 (2003) 253.
- [17] M. Boyle, L. Lindblom, H. Pfeiffer, M. Scheel, L.E. Kidder, Testing the Accuracy and Stability of Spectral Methods in Numerical Relativity, preprint *gr-qc/0609047*, available at <http://lanl.arxiv.org/>.
- [18] C. Beetle, B. Bromley, R.H. Price, Periodic standing-wave approximation: eigenspectral computations for linear gravity and nonlinear toy models, *Phys. Rev. D* 74 (2006) 024013.
- [19] C. Beetle, B. Bromley, N. Hernández, R.H. Price, The periodic standing-wave approximation: post-Minkowski computations, *Phys. Rev. D* on 8 August 2007, *gr-qc/0708.1141* (submitted for publication).
- [20] F.I. Frankl, On the problems of Chaplygin for mixed subsonic and supersonic flows, *Izv. Akad. Nauk SSSR Ser. Mat.* 9 (1945) 121.
- [21] F.G. Tricomi, Sulle equazioni lineari alle parziali di 2° ordine di tipo misto, *Atti Accad. Naz. Lincei* 14 (1923) 133.
- [22] C.S. Morawetz, A uniqueness theorem for Frankl's problem, *Commun. Pure Appl. Math.* 7 (1954) 697;
C.S. Morawetz, Note on a maximum principle and a uniqueness theorem for an elliptic-hyperbolic equation, *Proc. Roy. Soc. Lond. Ser. A* 236 (1956) 141;
C.S. Morawetz, A weak solution for a system of equations of elliptic-hyperbolic type, *Commun. Pure Appl. Math.* 11 (1958) 315.
- [23] M.H. Protter, Uniqueness theorems for the Tricomi problem, I, II, *J. Rat. Mech. Anal.* 2 (1955) 707;
M.H. Protter, Uniqueness theorems for the Tricomi problem, I, II, *J. Rat. Mech. Anal.* 4 (1955) 721.
- [24] E.M. Murman, J.D. Cole, Calculation of plane steady transonic flows, *AIAA J.* 9 (1) (1971) 114;
E.M. Murman, J.D. Cole, Calculation of plane steady transonic flows, *AIAA J.* 41 (7A) (2003) 301 (Reprinted).
- [25] D.A. Caughey, A. Jameson, Development of computational techniques for transonic flows: an historical perspective, *Symposium Transonicum IV*, vol. 183, Kluwer Academic, Göttingen, 2003.
- [26] T.H. Otway, A boundary-value problem for cold plasma dynamics, *J. Appl. Math.* 3 (2003) 17.
- [27] N. Bilá, Application of symmetry analysis to a PDE arising in the car windshield design, *SIAM J. Appl. Math.* 65 (2004) 113.
- [28] G. Chaoahao, The mixed PDE for amplifying spiral waves, *Lett. Math. Phys.* 16 (1988) 69.
- [29] K.O. Friedrichs, Symmetric positive linear differential equations, *Commun. Pure Appl. Math.* 11 (1958) 333.
- [30] M. Jensen, Discontinuous Galerkin methods for Friedrich systems with irregular solutions, PhD Thesis, University of Oxford, 2004.
- [31] T. Katsanis, Numerical solution of symmetric positive differential equations, *Math. Comput.* 22 (1968) 763.
- [32] T. Katsanis, Numerical solution of tricomi equation using theory of symmetric positive differential equations, *SIAM J. Numer. Anal.* 6 (2) (1969) 236.
- [33] J.-L. Liu, A finite difference method for symmetric positive differential equations, *Math. Comput.* 62 (1994) 105.
- [34] A. Ern, J.-L. Guermond, Discontinuous Galerkin methods for Friedrichs' systems. I. General theory, *SIAM J. Numer. Anal.* 44 (2006) 753;
A. Ern, J.-L. Guermond, Discontinuous Galerkin methods for Friedrichs' systems. Part II. Second-order elliptic PDEs, *SIAM J. Numer. Anal.* 44 (2006) 2363.
- [35] I.J. Thompson, A.R. Barnett, Coulomb and Bessel functions of complex arguments and order, *J. Comput. Phys.* 64 (1986) 490.
- [36] T. Hagstrom, Radiation boundary conditions for the numerical simulation of waves, *Acta Numer.* 8 (1999) 47.
- [37] B. Alpert, L. Greengard, T. Hagstrom, Rapid evaluation of nonreflecting boundary kernels for time-domain wave propagation, *SIAM J. Numer. Anal.* 37 (2000) 1138.
- [38] S.R. Lau, Rapid evaluation of radiation boundary kernels for time-domain wave propagation on blackholes: theory and numerical methods, *J. Comput. Phys.* 199 (2004) 376.
- [39] T. Hagstrom, S. Lau, Radiation boundary conditions for Maxwell's equations: a review of accurate time-domain formulations, *J. Comput. Math.* 25 (2007) 305–336.
- [40] E.A. Coutsias, T. Hagstrom, J.S. Hesthaven, D. Torres, Integration preconditioners for differential operators in spectral τ -methods, *Houston J. Math.* 21 (1996) (Special Issue).
- [41] E.A. Coutsias, T. Hagstrom, D. Torres, An efficient spectral method for ordinary differential equations with rational function coefficients, *Math. Comput.* 65 (1996) 611.
- [42] T.J. Rivlin, Chebyshev Polynomials. From Approximation Theory to Algebra and Number Theory, second ed., Wiley, New York, 1990.
- [43] P. Beckman, Orthogonal Polynomials for Engineers and Physicists, Golem Press, Boulder, 1973.

- [44] J.S. Hesthaven, Integration preconditioning of pseudospectral operators. I. Basic linear operators, *SIAM J. Numer. Anal.* 35 (1998) 1571.
- [45] G. von Winckel, S. Krishna, E.A. Coutsias, Spectral element modeling of semiconductor heterostructures, *Math. Comput. Model.* 43 (2006) 582.
- [46] C.T. Kelley, *Solving Nonlinear Equations with Newton's Method*, SIAM, Philadelphia, 2003.
- [47] J.C. Adams, P.N. Swarztrauber, SPHEREPACK 3.0: a model development facility, *Mon. Weather Rev.* 127 (1999) 1872.
- [48] V. Frayssé, L. Giraud, S. Gratton, J. Langou, A set of GMRES routines for real and complex arithmetics on high performance computers, CERFACS Technical Report TR/PA/03/3, public domain software available on www.cerfacs.fr/algorithm/Softs, 2003.
- [49] J.W. Demmel, *Applied Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [50] J.P. Boyd, *Chebyshev and Fourier Spectral Methods*, second ed., Dover, New York, 2001.