GYRE: an open-source stellar oscillation code based on a new Magnus Multiple Shooting scheme

R. H. D. Townsend⋆ and S. A. Teitler

Department of Astronomy, University of Wisconsin-Madison, 2535 Sterling Hall, 475 N. Charter Street, Madison, WI 53706, USA

ABSTRACT

We present a new oscillation code, GYRE, which solves the stellar pulsation equations (both adiabatic and non-adiabatic) using a novel Magnus Multiple Shooting numerical scheme devised to overcome certain weaknesses of the usual relaxation and shooting schemes appearing in the literature. The code is accurate (up to sixth-order in the number of grid points), robust, efficiently makes use of multiple processor cores and/or nodes and is freely available in source form for use and distribution. We verify the code against analytic solutions and results from other oscillation codes, in all cases finding good agreement. Then, we use the code to explore how the asteroseismic observables of a $1.5\,M_\odot$ star change as it evolves through the red-giant bump.

Key words: methods: numerical – stars: evolution – stars: interiors – stars: oscillations – stars: variables: general.

1 INTRODUCTION

The field of asteroseismology has been reinvigorated in recent years thanks to the wealth of new observational data provided by space-based instruments. Over the past decade there have been three satellite missions with specific asteroseismic objectives: *MOST* (Walker et al. 2003; Matthews 2007) launched in 2003, *CoRoT* (Michel et al. 2008; Baglin et al. 2009) launched in 2006 and *Kepler* (Borucki et al. 2009; Gilliland et al. 2010; Kjeldsen et al. 2010) launched in 2009. Exciting results from these missions include the discovery that nearly all γ Doradus and δ Scuti stars are hybrid pulsators (Grigahcène et al. 2010), ensemble asteroseismic analysis of solar-like oscillations in hundreds of solar-type stars (Chaplin et al. 2011) and the detection of solar-like oscillations in a large sample of red giants (De Ridder et al. 2009). See also the reviews by Christensen-Dalsgaard & Thompson (2011) and Chaplin & Miglio (2013) for further highlights.

Interpreting these new observations requires the seismologist's analogue to the telescope: a stellar oscillation code which calculates the eigenfrequency spectrum of an arbitrary input stellar model. Comparing a calculated spectrum against a measured one provides a concrete metric for evaluating a model, and therefore constitutes the bread and butter of quantitative asteroseismology. Although the task of iteratively improving model parameters has been quite cumbersome in the past, there are now tools available that largely automate this process. The Asteroseismic Modelling Portal (AMP; Metcalfe, Creevey & Christensen-Dalsgaard 2009) provides a web-based front end for asteroseismic analysis and model optimization of solar-like stars, using the Aarhus Stellar Evolution code (ASTEC; Christensen-Dalsgaard 2008a) to build models and the Aarhus Pulsation code (ADIPLS; Christensen-Dalsgaard 2008b) to calculate their eigenfrequencies. Likewise, the widely adopted MESA stellar evolution code (Paxton et al. 2011, 2013) includes an asteroseismology module also based on ADIPLS, and offering similar optimization capabilities to AMP.

Such tools place ever-increasing demands on the oscillation codes that underpin them. A code will typically be executed hundreds or thousands of times during an optimization run, and must therefore make efficient use of available computational resources (e.g. multiple processor cores and/or cluster nodes). The code must be robust, running and producing sensible output without manual intervention like hand-tuning. The code must have an accuracy that matches or exceeds the frequency precision now achievable by satellite missions. Finally, it is preferable that the code addresses the various physical processes that inevitably complicate calculations, such as non-adiabaticity, rotation and magnetic fields.

These desiderata motivated us to develop a new oscillation code, GYRE, which we describe in the present paper. The code is based on a 'Magnus Multiple Shooting' (MMS) scheme for solving the linearized pulsation equations, devised by us to address various pitfalls encountered with the standard relaxation and shooting schemes appearing in the literature. The following section reviews these schemes and the existing oscillation codes which use them. Section 3 then describes the MMS scheme in detail, and Section 4 discusses how the scheme is implemented in GYRE. We present example calculations in Section 5, and discuss and summarize the paper in Section 6.

2 BACKGROUND

The differential equations and algebraic boundary conditions governing small-amplitude non-radial oscillations of a star about an

⋆ E-mail: townsend@astro.wisc.edu

equilibrium background state – the so-called linearized stellar pulsation equations, presented in the adiabatic case in Appendix A – constitute a two-point boundary value problem (BVP) in which the oscillation frequency ω serves as an eigenvalue (for a comprehensive review, see the monographs by Cox 1980; Unno et al. 1989; Aerts, Christensen-Dalsgaard & Kurtz 2010; Smeyers & van Hoolst 2010). Although there exist special cases where analytic solutions exist (e.g. Pekeris 1938), in general this BVP must be solved numerically. Oscillation codes specializing in this task were first described a half century ago, and since then many different numerical schemes have been proposed in the literature. The following sections review the two most prevalent, and Section 2.3 then briefly discusses other approaches which have been adopted.

2.1 Relaxation schemes

Relaxation schemes for BVPs replace the derivatives in the differential equations with finite-difference approximations specified on a grid. Applied to pulsation problems, the finite-difference relations together with boundary conditions and a normalization condition establish a (typically large) system of algebraic equations in which the unknowns are the dependent variables y at the grid points plus the dimensionless oscillation frequency ω. Because these equations are non-linear (specifically, bi-linear in y and ω^2 for adiabatic pulsation), the simultaneous determination of all unknowns requires iterative improvement of a trial solution – for instance using the procedure developed initially for stellar evolution calculations by Henyey, Forbes & Gould (1964), which can be regarded as a multi-dimensional Newton–Raphson algorithm. (Unno et al. 1989 present a detailed implementation of this procedure specifically tailored to the pulsation equations).

The convergence of the Henyey scheme depends on how close a trial solution is to a true solution. Castor (1971) proposed an elegant approach to finding good trial solutions for radial pulsation problems, which Osaki & Hansen (1973) subsequently adapted to the non-radial case. One of the boundary conditions is set aside, allowing the system of algebraic equations to be solved at any ω using a standard *linear* algorithm (e.g. Gaussian elimination). The overlooked boundary condition is then used to construct a discriminant function $\mathcal{D}(\omega)$ which vanishes when the boundary condition is satisfied. Clearly, the roots of $\mathcal{D}(\omega)$ correspond to the eigenfrequencies of the full BVP; thus, good trial solutions can be obtained by isolating and refining these roots.

Relaxation using the Castor (1971) approach has proven very popular, forming the basis for many oscillation codes including the BOOJUM code (Townsend 2005), the Nice Oscillation code (NOC; Provost 2008), the Granada Oscillation code (GRACO; Moya & Garrido 2008) and the LNAWENR code (Suran 2008). It is generally robust, but can run into difficulty when the discriminant function exhibits singularities. These arise when the dependent variable used for normalization naturally exhibits a zero at the point where the normalization is applied. Unno et al. (1989) propose addressing this problem by dividing the discriminant by one of the dependent variables evaluated at the opposite boundary to the overlooked boundary condition. This approach works well for the adiabatic pulsation equations within the Cowling (1941) approximation (where perturbations to the gravitational potential are neglected), because neither of the dependent variables in this second-order BVP is ever zero at the boundaries. However, in more general cases no such guarantees can be made, and the division itself can make the singularities recrudesce.

Attempts at more-sophisticated fixes to the singular discriminant problem seem similarly doomed to failure. Because the singularities ultimately stem from the imposition of an inappropriate normalization, it is better to avoid normalization altogether when searching for eigenfrequencies; this is the approach taken by the MMS scheme (Section 3).

2.2 Shooting schemes

Shooting schemes treat BVPs as a set of initial value problems (IVPs), with matching conditions applied where pairs of these IVPs meet. In the stellar oscillation literature 'double shooting' (also termed 'shooting to a fitting point') is most commonly encountered: IVPs are integrated from each boundary towards an internal fitting point, with initial values determined from the boundary conditions. The mismatch between solutions at the fitting point is quantified by a discriminant function $\mathcal{D}(\omega)$ which vanishes when the integrations match. As before, the roots of $\mathcal{D}(\omega)$ correspond to the eigenfrequencies of the BVP.

Hurley, Roberts & Wright (1966) and Smeyers (1966, 1967) were among the first to apply double shooting to the pulsation equations. Scuflaire (1974) adopted a simplified version of the scheme, where the fitting point is placed at a boundary and only one IVP integration is performed (so-called 'single shooting' or 'simple shooting'); however, the integration can become unstable when approaching a boundary where the differential equations become singular (i.e. the inner boundary, and in polytropic models the outer boundary too). Modern oscillation codes based on double shooting include ADIPLS (which can use either shooting or relaxation), the Porto Oscillation Code (POSC; Monteiro 2008) and the OSCROX code (Roxburgh 2008).

Christensen-Dalsgaard (1980) discusses a complication that arises when using double shooting to solve the adiabatic pulsation equations without the Cowling (1941) approximation. This fourth-order BVP requires integrating two linearly independent solutions from each boundary. In evanescent regions these solutions are dominated by an exponentially growing component, and they can easily become numerically linearly dependent. The problem cannot be fixed by switching the direction of integration (as one might do with an IVP), because the BVP has an 'exponential dichotomy' – components that grow and decay exponentially in both directions.[1] This is a well-established weakness of single/double shooting schemes in general and has been extensively analysed in the BVP literature (see e.g. the excellent monograph by Ascher, Mattheij & Russell 1995). Happily, the same literature provides a number of strategies for avoiding this weakness. One of them, the Riccati method, has already been used by Gautschy & Glatzel (1990) to explore highly non-adiabatic oscillations (and see also Valsecchi et al. 2013). Another, multiple shooting, forms the basis of the MMS scheme.

2.3 Other approaches

Although shooting and relaxation dominate in the stellar oscillation literature, they are by no means the only schemes used. The groundbreaking paper by Hurley et al. (1966), already mentioned above as an early instance of shooting, also describes a collocation method (and the authors allude to the possibility of a third approach, which

---

[1] In fact this is a good thing; as Hoog & Mattheij (1987) demonstrate, an exponential dichotomy is a necessary condition for a BVP to be well conditioned.

can be recognized as relaxation!). Collocation methods approximate BVP solutions as a superposition of basis functions (e.g. Chebyshev polynomials) which satisfy the differential equations exactly at a set of nodes. Recently, Reese, Lignières & Rieutord (2006) again used collocation to explore oscillations of polytropes, but this time incorporating the effects of rapid rotation.

The finite element method (FEM) shares some similarities with collocation methods, also using superpositions of basis functions. However, the functions are chosen to minimize certain integrals representative of the solution error. Two examples of FEM-based oscillation codes are FILOU (Suárez & Goupil 2008) and PULSE (Brassard & Charpinet 2008).

One other approach garnering some interest is inverse iteration. As with relaxation, the differential equations are approximated with finite differences. However, the resulting algebraic equations are explicitly structured as a generalized linear eigenvalue problem, which is then solved using the well-established technique of inverse iteration (e.g. Golub & van Loan 1996). With a good trial solution convergence is rapid. This approach is used by the MAD code (Dupret 2001) and the Liège oscillation code (OSC; Scuflaire et al. 2008).

3 THE MAGNUS MULTIPLE-SHOOTING SCHEME

3.1 Problem statement

The MMS scheme solves BVPs defined by a system of linear, homogeneous, first-order ordinary differential equations

$$\frac{d\boldsymbol{y}}{dx} = \mathbf{A}(x)\boldsymbol{y} \tag{1}$$

defined on the interval $x^{a} \leq x \leq x^{b}$, together with boundary conditions applied at each end of the interval,

$$\mathbf{B}^{a}\boldsymbol{y}(x^{a}) = \mathbf{0}, \qquad \mathbf{B}^{b}\boldsymbol{y}(x^{b}) = \mathbf{0}. \tag{2}$$

For n equations, $\boldsymbol{y} \in \mathbb{C}^{n}$ is the vector of dependent variables and $\mathbf{A} \in \mathbb{C}^{n \times n}$ is the Jacobian matrix. If n_{a} of the boundary conditions are applied at the inner point x^{a} and the remaining $n_{b} \equiv n - n_{a}$ are applied at the outer point x^{b}, then $\mathbf{B}^{a} \in \mathbb{C}^{n_{a} \times n}$ and $\mathbf{B}^{b} \in \mathbb{C}^{n_{b} \times n}$.

3.2 Multiple shooting

Multiple shooting is a natural extension of the single/double shooting schemes discussed in Section 2.2 which avoids the numerical difficulties encountered when the system of equations exhibits an exponential dichotomy. Ascher et al. (1995) discuss it in considerable depth; here, we highlight the important aspects. The interval is divided up into a grid of N points

$$x^{a} \equiv x^{1} < x^{2} < \cdots < x^{N-1}, \quad x^{N} \equiv x^{b}. \tag{3}$$

The solution to the BVP at any point in the kth subinterval $x^{k} \leq x \leq x^{k+1}$ ($k = 1, 2, \ldots, N-1$) is written as

$$\boldsymbol{y}(x) = \mathbf{Y}(x; x^{k})\boldsymbol{y}^{k}, \tag{4}$$

where $\boldsymbol{y}^{k} \equiv \boldsymbol{y}(x^{k})$ and the fundamental solution $\mathbf{Y}(x; x') \in \mathbb{C}^{n \times n}$ is the matrix function satisfying the IVP

$$\frac{d\mathbf{Y}}{dx} = \mathbf{A}(x)\mathbf{Y}, \qquad \mathbf{Y}(x'; x') = \mathbf{I}. \tag{5}$$

Here, $\mathbf{I}$ is the rank-n identity matrix.

The requirement that $\boldsymbol{y}$ be continuous at subinterval edges imposes the matching condition

$$\boldsymbol{y}^{k+1} = \mathbf{Y}^{k+1;k}\boldsymbol{y}^{k}, \tag{6}$$

where we use the shorthand

$$\mathbf{Y}^{k+1;k} \equiv \mathbf{Y}(x^{k+1}; x^{k}) \tag{7}$$

for the fundamental solution matrix spanning the kth subinterval. There are $N - 1$ such matching conditions, and in combination with the boundary conditions (2) they lead to the system of algebraic equations

$$\mathbf{S}\boldsymbol{u} = \mathbf{0}. \tag{8}$$

The vector of unknowns $\boldsymbol{u} \in \mathbb{C}^{Nn}$ packs together the dependent variables at the grid points,

$$\boldsymbol{u} = \begin{pmatrix} \boldsymbol{y}^{1} \\ \boldsymbol{y}^{2} \\ \vdots \\ \boldsymbol{y}^{N-1} \\ \boldsymbol{y}^{N} \end{pmatrix}, \tag{9}$$

and the system matrix $\mathbf{S} \in \mathbb{C}^{Nn \times Nn}$ has a staircase structure (e.g. Fourer 1984) given by

$$\mathbf{S} = \begin{pmatrix} \mathbf{B}^{a} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ -\mathbf{Y}^{2;1} & \mathbf{I} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -\mathbf{Y}^{3;2} & \mathbf{I} & \cdots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & -\mathbf{Y}^{N;N-1} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{B}^{b} \end{pmatrix}. \tag{10}$$

As a linear homogeneous system, equation (8) admits non-trivial solutions only when the determinant of the system matrix vanishes,

$$\det(\mathbf{S}) = 0. \tag{11}$$

This can be recognized as the characteristic equation of the BVP. In the case of the pulsation equations $\mathbf{S}$ depends implicitly on ω; thus, the stellar eigenfrequencies are the roots of the discriminant function

$$\mathcal{D}(\omega) = \det[\mathbf{S}(\omega)], \tag{12}$$

and can be determined using a suitable root-finding algorithm. Setting ω equal to one specific eigenfrequency, the corresponding eigenfunctions are first constructed on the shooting grid $\{x^{k}\}$ by finding the non-trivial vector $\boldsymbol{u}$ satisfying equation (8). Then, the eigenfunctions at any point in any subinterval follow from applying equation (4).

Inspecting the form of $\mathbf{S}$ suggests that the system of equations can be greatly simplified to

$$\mathbf{S}^{c}\boldsymbol{u}^{c} = \mathbf{0}, \tag{13}$$

where

$$\boldsymbol{u}^{c} = \begin{pmatrix} \boldsymbol{y}^{1} \\ \boldsymbol{y}^{N} \end{pmatrix} \tag{14}$$

and

$$\mathbf{S}^{c} = \begin{pmatrix} & \mathbf{B}^{a} & & \mathbf{0} \\ -\mathbf{Y}^{N;N-1}\mathbf{Y}^{N-1;N-2}\dots\mathbf{Y}^{3;2}\mathbf{Y}^{2;1} & & \mathbf{I} \\ & \mathbf{0} & & \mathbf{B}^{b} \end{pmatrix}. \tag{15}$$

Unfortunately, this approach, known as compactification, suffers from a similar issue to single/double shooting: when evaluating the product of fundamental solution matrices in the above expression, the columns become numerically linearly dependent (see Ascher et al. 1995, for a more-detailed discussion).

3.3 Magnus integrators

To evaluate the fundamental solution matrices $\mathbf{Y}^{k+1;k}$ in each of the $N-1$ subintervals, the MMS scheme builds on an approach proposed by Gabriel & Noels (1976). These authors approximated the Jacobian matrix of the adiabatic pulsation equations as piecewise-constant in each subinterval ('shell' in their terminology). In the present context this leads to a fundamental solution

$$\mathbf{Y}^{k+1;k} = \exp(\mathbf{A}\Delta x^{k}), \tag{16}$$

where $\Delta x^{k} \equiv x^{k+1} - x^{k}$ (a derivation of this result appears below). This expression involves matrix exponentiation – a topic discussed at length by Moler & Van Loan (2003), who survey the strengths and weaknesses of nineteen different methods. Here we focus on eigen-decomposition (their method 14), both for pedagogic purposes and because it is adopted in the GYRE code (Section 4.3). The Jacobian matrix is written as

$$\mathbf{A} = \mathbf{M_A}\mathbf{\Lambda_A}\mathbf{M_A}^{-1}, \tag{17}$$

where $\mathbf{\Lambda_A} \in \mathbb{C}^{n\times n}$ is a diagonal matrix whose non-zero elements are the eigenvalues $\{\lambda_\mathbf{A}^i\}$ $(i = 1, \dots, n)$ of $\mathbf{A}$, and the columns of the matrix $\mathbf{M_A} \in \mathbb{C}^{n\times n}$ comprise the corresponding eigenvectors. With this decomposition, the fundamental solution (16) becomes

$$\mathbf{Y}^{k+1;k} = \mathbf{M_A}\exp(\mathbf{\Lambda_A}\Delta x^{k})\mathbf{M_A}^{-1}, \tag{18}$$

where the non-zero elements of the diagonal matrix $\exp(\mathbf{\Lambda_A}\Delta x^{k})$ are

$$[\exp(\mathbf{\Lambda_A}\Delta x^{k})]_{ii} = \exp(\lambda_\mathbf{A}^i\Delta x^{k}). \tag{19}$$

An instructive physical narrative for these equations can be obtained by substituting equation (18) into (6) to yield

$$\mathbf{y}^{k+1} = \mathbf{M_A}\exp(\mathbf{\Lambda_A}\Delta x^{k})\mathbf{M_A}^{-1}\mathbf{y}^{k}. \tag{20}$$

The matrices on the right-hand side of equation (20) correspond to a sequence of operations which advance $\mathbf{y}$ from the kth grid point to the $k+1$ point. First, $\mathbf{y}^{k}$ is projected on to a set of basis vectors given by the rows of $\mathbf{M_A}^{-1}$. This projection amounts to decomposing $\mathbf{y}^{k}$ into contributions from n independent waves. Then, the amplitudes and phases of the waves are evolved across the subinterval by applying the diagonal matrix $\exp(\mathbf{\Lambda_A}\Delta x^{k})$. In evanescent zones all eigenvalues are real and only the wave amplitudes change, whereas in propagation zones one or more eigenvalues are complex and the wave phases also change. Finally, the waves are projected back into physical space by the matrix $\mathbf{M_A}$.

The Gabriel & Noels (1976) approach can be generalized by recognizing it as an application of a simple yet powerful theorem proposed by Magnus (1954). Subject to certain convergence criteria, the solution to the IVP (5) can be written as the matrix exponential

$$\mathbf{Y}(x;x') = \exp[\mathbf{\Omega}(x;x')], \tag{21}$$

where the Magnus matrix $\mathbf{\Omega} \in \mathbb{C}^{n\times n}$ has a series expansion whose leading terms are

$$\mathbf{\Omega}(x;x') = \int_{x'}^{x} \mathbf{A}(x_1)\,\mathrm{d}x_1$$
$$- \frac{1}{2}\int_{x'}^{x}\left[\int_{x'}^{x_1}\mathbf{A}(x_2)\,\mathrm{d}x_2, \mathbf{A}(x_1)\right]\,\mathrm{d}x_1 + \cdots \tag{22}$$

(here, $[\dots,\dots]$ denotes the matrix commutator). Blanes et al. (2009) present a detailed review of Magnus's theorem, covering both its mathematical underpinnings and its practical application to solving systems of differential equations.

In the context of the MMS scheme, Magnus's theorem gives the fundamental solution matrix within each subinterval as

$$\mathbf{Y}^{k+1;k} = \exp[\mathbf{\Omega}(x^{k+1};x^{k})]. \tag{23}$$

If the Jacobian matrix $\mathbf{A}$ is independent of x, then all terms but the first in the expansion (22) vanish and the Magnus matrix is simply

$$\mathbf{\Omega}(x^{k+1};x^{k}) = \mathbf{A}\,\Delta x^{k}. \tag{24}$$

By combining equations (23) and (24) we recover the fundamental solution (16) obtained using the Gabriel & Noels (1976) approach. This constant-Jacobian case is the only one having a closed-form expression for the Magnus matrix; however, as discussed by Blanes et al. (2009) it is relatively straightforward to construct approximations to equation (22) which are correct to some specified order in the subinterval width Δx^{k}. Specifically, if second-order Gauss–Legendre quadrature is used to evaluate the integrals in the expansion (22), then the Magnus matrix becomes

$$\mathbf{\Omega}(x^{k+1};x^{k}) \approx \mathbf{A}\left(x^{k} + \frac{\Delta x^{k}}{2}\right)\Delta x^{k} + \mathcal{O}[(\Delta x^{k})^3]. \tag{25}$$

Higher-order quadratures lead to correspondingly higher-order approximations; in Appendix (B) we quote fourth- and sixth-order accurate expressions for the Magnus matrix, also based on Gauss–Legendre quadrature. We refer to the fundamental solutions (23) using these expressions as the Magnus GLo integrators, where GL for stands for Gauss–Legendre and $o = 2, 4, 6$ indicates the order of accuracy. These integrators all share the useful property of not requiring Jacobian evaluations at the subinterval endpoints, which allows them to gracefully handle singularities at the interval boundaries.

3.4 Commentary

Compared to single/double shooting schemes, multiple shooting has the advantage that the subinterval width Δx^{k} can always be chosen sufficiently small that the columns of the associated fundamental solution matrix $\mathbf{Y}^{k+1,k}$ remain linearly independent, even in the presence of an exponential dichotomy. This is because the matrix approaches the identity matrix in the limit $\Delta x^{k} \to 0$. In practice, this choice rarely needs to be made explicitly; small Δx^{k} is already desirable in the interests of accuracy.

Despite being based on shooting, certain parts of the MMS scheme bear a strong resemblance to relaxation. In particular, the staircase structure of the system matrix (10) also arises in finite-difference approximations to the BVP differential equations (see Section 2.1). This is no coincidence: it is straightforward to demonstrate that any relaxation scheme can be built from a multiple shooting scheme (and vice versa) by choosing a suitable numerical method for the IVP (4).

In principle, equation (8) could be solved using Castor's method: by replacing one of the boundary conditions in the first or last block rows with a normalization condition, the system of equations becomes inhomogeneous and can be solved for any ω. The replaced boundary condition then serves as the discriminant function. However, this would be an obtuse way to solve a homogeneous linear problem, and it is little wonder that problems arise (e.g. the singularities discussed in Section 2.1). The determinant-based discriminant we propose in equation (12) is the natural approach, and given that the elements of **S** are finite it is guaranteed to be well behaved. Dupret (2001) and Scuflaire et al. (2008) successfully use a similar method to find trial solutions for their inverse iteration schemes.

The MMS scheme is the first explicit application of Magnus's theorem to stellar oscillations. However, as we demonstrate in the preceding section the Gabriel & Noels (1976) method is equivalent to shooting using a Magnus integrator. More recently, Christensen-Dalsgaard (2008b) mentions that the ADIPLS code can optionally use a similar approach (second-order Magnus based on Newton–Cotes quadrature) to integrate the adiabatic pulsation equations within the Cowling (1941) approximation. In both papers, the authors recognize the schemes' key strength that they can resolve solutions which vary on arbitrarily small spatial scales – something that fixed-stepsize IVP solvers cannot do.

4 THE GYRE CODE

GYRE is a new oscillation code which uses the MMS scheme described above to calculate the eigenfrequencies and eigenfunctions of an input stellar model. Although GYRE can address both adiabatic and non-adiabatic pulsation problems, in this paper we focus on the adiabatic case (documented in Appendix A) because our primary goal is to introduce the MMS scheme and the code.

GYRE is written in FORTRAN 2008 with a modular architecture that allows straightforward extension to handle more-complicated problems. To take advantage of multiple processor cores and/or cluster nodes it is parallelized using a combination of OpenMP (Dagum & Menon 1998) and Message Passing Interface (MPI; Dongarra et al. 1995). In brief, a typical GYRE run involves the following steps: first, a stellar model is either read from file or built analytically (Section 4.1), and the calculation grids are constructed (Section 4.2). A scan through frequency space then searches for sign changes in the discriminant $\mathcal{D}(\omega)$, which are used as initial guesses for the discriminant roots (Section 4.5). After these roots are found, the corresponding eigenfunctions are reconstructed (Section 4.6). The following sections further discuss these steps and provide other salient implementation details.

4.1 Stellar model

GYRE supports three classes of stellar model, each providing the dimensionless structure coefficients V, A^*, U, c_1 and Γ_1 appearing in the pulsation equations (see Appendix A). Evolutionary models are generated by a stellar evolution code, polytropic models are based on solutions to the Lane–Emden equation, and analytic models rely on explicit expressions for the structure coefficients. Both evolutionary and polytropic models are specified on a discrete radial grid, with cubic spline interpolation used to evaluate the structure coefficients between grid points. Different options exist for constructing the splines, with the monotonicity-preserving algorithm by Steffen (1990) being the default.

4.2 Grid construction

GYRE offers a number of strategies for establishing the grids used for multiple shooting (see equation 3) and eigenfunction reconstruction (discussed below in Section 4.6). For evolutionary and polytropic models, the grid can be cloned from the corresponding model grid, with the option of oversampling certain subintervals. GYRE can also create an ab initio grid following a variety of recipes. The simplest of these is the 'double geometric' grid with subinterval widths given by

$$\Delta x^k = (1 + g)\Delta x^{k-1} \qquad k \le M,$$
$$\Delta x^k = (1 + g)\Delta x^{k+1} \qquad k > M. \tag{26}$$

Here, $M = N/2$ for even N and $M = (N-1)/2$ for odd N, and the growth factor g is determined from the requirement that

$$\sum_{k=1}^{N-1} \Delta x^k = 1. \tag{27}$$

The subinterval sizes at the boundaries are fixed by a user-specified stretching parameter s representing the ratio between the average subinterval size and the boundary size; thus,

$$\Delta x_1 = \Delta x_{N-1} = \frac{1}{s(N-1)}. \tag{28}$$

The double geometric grid has greatest resolution near the inner and outer boundaries – a useful property because the components of the Jacobian matrix typically vary fastest near these boundaries.

4.3 Fundamental solution calculation

GYRE calculates the fundamental solutions with one of the Magnus GLo integrators (the order o is configurable at run time). The matrix exponential in equation (23) is evaluated using eigendecomposition, and so the fundamental solution in the kth subinterval is obtained as

$$\mathbf{Y}^{k+1;k} = \mathbf{M}_\Omega \exp(\mathbf{\Lambda}_\Omega)\mathbf{M}_\Omega^{-1}. \tag{29}$$

Here, $\mathbf{\Lambda}_\Omega$ and $\mathbf{M}_\Omega$ are the eigenvalue and eigenvector matrices of the Magnus matrix $\mathbf{\Omega}(x^{k+1}; x^k)$, itself taken from one of equations (25), (B1) or (B4) for the GL2, GL4 or GL6 integrators, respectively. The eigendecomposition is implemented with calls to the LAPACK linear algebra library (Anderson et al. 1999); OpenMP directives are used to distribute the work for the $N - 1$ subintervals across multiple processor cores.

The eigendecomposition can fail if $\mathbf{\Omega}(x^{k+1}; x^k)$ lacks a complete set of linearly independent eigenvectors (i.e. the matrix is defective; see Golub & van Loan 1996). In such cases one of the alternative matrix exponentiation methods discussed by Moler & Van Loan (2003) must be used. In practice, we have never encountered this situation; nevertheless, as a precaution GYRE is configured to abort with an error when it detects a defective or near-defective Magnus matrix.

4.4 Determinant evaluation

To evaluate the determinant of the system matrix **S**, GYRE first constructs the LU factorization

$$\mathbf{S} = \mathbf{LU}, \tag{30}$$

where **L** is a lower-triangular matrix with unit diagonal elements and **U** is an upper triangular matrix. The determinant of a triangular

matrix is simply the product of its diagonal elements, and it therefore follows that

$$\det(\mathbf{S}) = \prod_{j=1}^{Nn} \mathbf{U}_{jj}. \qquad (31)$$

GYRE undertakes the factorization (30) using the structured algorithm described by Wright (1992, 1994), which is specifically targeted at matrices arising in multiple shooting schemes. The cyclic reduction version of the algorithm is implemented because it produces the same results whether run serially or in parallel. Block-row pairs are distributed across multiple processor cores using OpenMP directives, and then factorized with calls to the LAPACK and BLAS (Blackford et al. 2002) libraries. We explored an MPI version of the algorithm for use on clusters, but found that inter-node communication overhead produces poor performance on systems larger than a few tens of nodes. As the dimension of $\mathbf{S}$ becomes large the determinant calculation risks overflowing the computer floating-point range; therefore, GYRE evaluates equation (31) with extended-range floating point arithmetic, build on the object-oriented capabilities of recent FORTRAN dialects.

4.5 Eigenfrequency searching

GYRE searches for eigenfrequencies within a user-specified frequency interval $\omega_a \leq \omega \leq \omega_b$ by first evaluating the discriminant $\mathcal{D}(\omega)$ at n_ω points distributed within this interval. A change in the sign of $\mathcal{D}(\omega)$ between an adjacent pair of points signals that a root is bracketed, and the pair is passed as starting guesses to a root-finding routine based on the algorithm described by Brent (1973).

To leverage multi-node clusters GYRE parallelizes the initial discriminant evaluations and subsequent root searches with calls to the MPI library. This is in addition to the OpenMP parallelization described above for the eigendecompositions and LU factorization.

4.6 Eigenfunction reconstruction

For each eigenfrequency found as a root of $\mathcal{D}(\omega)$, GYRE reconstructs the corresponding eigenfunctions on the shooting grid $\{x^k\}$ by solving the algebraic system (8). The LU factorization (30) reduces this system to

$$\mathbf{U}\,\mathbf{u} = \mathbf{0}. \qquad (32)$$

Because the upper triangular matrix $\mathbf{U}$ is singular when ω is an eigenfrequency, one of its diagonal elements – say, $\mathbf{U}_{j'j'}$ – must be zero to within numerical uncertainties. The elements of the solution vector can then be written as

$$\mathbf{u}_j = \begin{cases} -[\tilde{\mathbf{U}}^{-1}\tilde{\mathbf{u}}]_j & j < j', \\ 1 & j = j', \\ 0 & j > j', \end{cases} \qquad (33)$$

where the square matrix $\tilde{\mathbf{U}}$ is formed from the first $j' - 1$ rows and columns of $\mathbf{U}$, and the vector $\tilde{\mathbf{u}}$ is formed from the first $j' - 1$ elements of the j' column of $\mathbf{U}$. Because $\tilde{\mathbf{U}}$ is upper triangular, the product $\tilde{\mathbf{U}}^{-1}\tilde{\mathbf{u}}$ is evaluated trivially by back-substitution (e.g. Golub & van Loan 1996). With $\mathbf{u}$ determined in this way, $\{\mathbf{y}^k\}$ can be unpacked using equation (9).

Non-trivial solutions to the pulsation equations cannot completely vanish at any point (otherwise, they would vanish everywhere). This means that the zero element $\mathbf{U}_{j'j'}$ should be located in the bottom-right corner of the matrix, such that $j' > (N - 1)n$. GYRE

explicitly tests whether this condition is met, and flags violations to indicate a problem with the solution. Our experience has been that these violations arise when the grid spacing is too large in one or more subintervals, preventing the Magnus expansion (22) from converging. The fix is invariably to reduce Δx^k by increasing N.

Once the eigenfunctions are obtained on the shooting grid, GYRE evaluates them on the separate reconstruction grid using a secant-line approximation to the Magnus matrix within each subinterval. In the kth subinterval this is

$$\boldsymbol{\Omega}(x; x^k) \approx w_k(x)\boldsymbol{\Omega}(x^{k+1}; x^k), \qquad (34)$$

where the weight function

$$w_k(x) \equiv \frac{x - x^k}{\Delta x^k} \qquad (35)$$

varies between 0 ($x = x^k$) and 1 ($x = x^{k+1}$). The solution at any point in the subinterval is then efficiently calculated as

$$\mathbf{y}(x) = \mathbf{M}_\Omega \exp[\boldsymbol{\Lambda}_\Omega\, w_k(x)]\mathbf{M}_\Omega^{-1}\,\mathbf{y}^k, \qquad (36)$$

where $\boldsymbol{\Lambda}_\Omega$ and $\mathbf{M}_\Omega$ are the same eigenvalue and eigenvector matrices obtained during the fundamental solution construction (cf. Section 4.3); no further eigendecomposition is required.

Eigenfunctions resulting from this procedure are C^∞ continuous within subintervals and are C^0 continuous at the edges. As final steps GYRE normalizes the eigenfunctions to have a mode inertia $\mathcal{E} = MR^2$ (see Aerts et al. 2010, their equation 3.139), and then classifies them in the standard Eckart (1960)–Scuflaire (1974)–Osaki (1975) (ESO) scheme by enumerating the acoustic- and gravity-wave winding numbers n_p and n_g. For dipole modes the ESO scheme can fail in certain circumstances (see e.g. Christensen-Dalsgaard & Mullan 1994), and so GYRE instead uses the modified scheme developed by Takata (2006).

5 CALCULATIONS

5.1 Eigenfrequencies of the $n = 0$ polytrope

As an initial test of GYRE and the underlying MMS scheme, we calculate radial and non-radial eigenfrequencies of the $n = 0$ polytrope model with $\Gamma_1 = 5/3$ (the so-called homogeneous compressible sphere). Because exact expressions exist for these frequencies (Pekeris 1938) this exercise allows an assessment of how the code and the scheme perform as the resolution of the shooting grid is varied.

Fig. 1 illustrates typical results, plotting the absolute error $\epsilon(\omega) \equiv |\omega - \omega_{ex}|$ in the eigenfrequency (with ω_{ex} being the exact value) as a function of the number of grid points N, for the $\ell = 1$ mode with $(n_p, n_g) = (1, 0)$ (traditionally labelled the dipole p_1 mode). The three curves show data from runs using the GL2, GL4 and GL6 Magnus integrators; in all cases the shooting grid is the double geometric grid described in Section 4.2 with stretching parameter $s = 10^3$.

The figure clearly reveals that the eigenfrequency error follows an asymptotic scaling $\propto N^{-2}$, N^{-4} and N^{-6} for the GL2, GL4 and GL6 integrators, respectively. This is the expected behaviour: the GLo integrator is o'th order accurate, leading to fundamental solutions (cf. equation 23) with an error scaling as Δx^{o+1}. Accumulated over the $N - 1$ subintervals the global error of the shooting scheme is then $\epsilon \sim (N - 1)\Delta x^{o+1} \sim N^{-o}$ (where we have used $\Delta x \sim N^{-1}$), which is the scaling seen in the figure. (For $N \gtrsim 2 \times 10^3$ the GL6 integrator departs from the asymptotic behaviour described, because
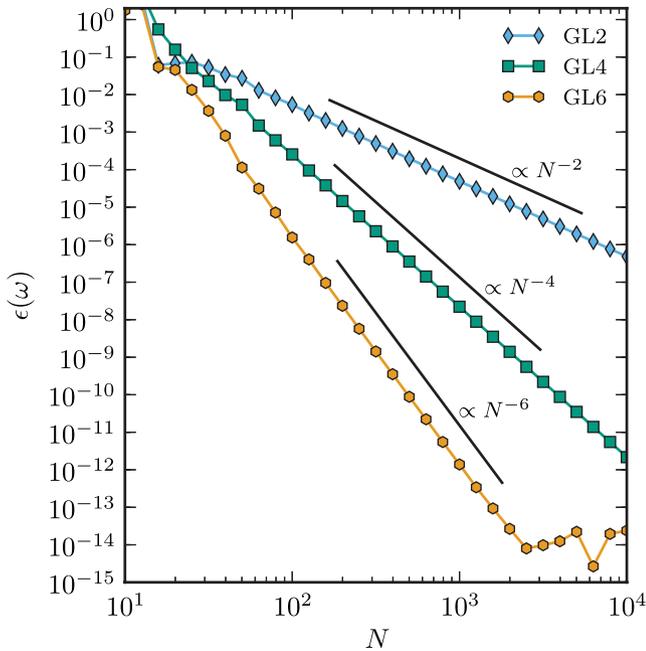
Figure 1. The absolute error in the dimensionless eigenfrequency of the dipole p_1 mode of the $n = 0$ polytrope, plotted as a function of the number of grid points. The three curves correspond to GYRE's GL2, GL4 and GL6 Magnus integrators, while the thick lines show the corresponding asymptotic scalings $\epsilon(\omega) \propto N^{-2}$, N^{-4} and N^{-6}, respectively.

numerical rounding becomes the dominant contributor towards the error).

Results for other radial and non-radial modes are comparable to those shown in the figure for the dipole p_1 mode. This confirms that the MMS scheme with the Magnus GL2, GL4 and GL6 integrators yields eigenfrequencies whose departures from exact values scale as the inverse second, fourth and sixth power of the grid size.

5.2 Inter-code comparison with ESTA model M4k

As a second verification exercise, we use GYRE to calculate eigenfrequencies of the 'M4k' model described by Moya et al. (2008). This model was produced with the ASTEC stellar evolution code (Christensen-Dalsgaard 2008a) and represents a 1.5 M$_\odot$ star at an age 1.35 Gyr, about half-way through its main-sequence evolution; it was adopted by Moya et al. (2008) as the basis for their comprehensive comparison of oscillation codes[2] from nine different research groups participating in the *CoRoT* Evolution and Seismic Tools activity (ESTA; see Lebreton et al. 2008). GYRE's shooting grid is cloned from the model grid without any oversampling (see Section 4.2), and the gravitational constant is set to the same value $G = 6.6716823 \times 10^{-8}\,\mathrm{cm^3\,g^{-1}\,s^{-2}}$ adopted by Moya et al. (2008). As in the preceding section, we perform separate runs using the GL2, GL4 and GL6 Magnus integrators.

Table 1 compares the GYRE linear eigenfrequencies against those obtained with GRACO, which was used as the reference code in the Moya et al. (2008) study. Across the range $20 \leq \nu \leq 2\,500\,\mu\mathrm{Hz}$ considered by these authors, the absolute error between the GYRE and GRACO frequencies is $\lesssim 4\,\mathrm{nHz}$ for radial modes, rising to $\lesssim 14\,\mathrm{nHz}$ in the $\ell = 3$ case. These values are comparable to the frequency differences found by Moya et al. (2008) between GRACO and the other

[2] All mentioned in Section 2.

Table 1. Maximum absolute differences, in nHz, between the GRACO and GYRE linear frequencies for $l = 0, \ldots, 3$ modes of the M4k model. Values are tabulated for each of GYRE's integrators and for the same low ($20 \leq \nu \leq 80\,\mu\mathrm{Hz}$), medium ($80 \leq \nu \leq 500\,\mu\mathrm{Hz}$) and high ($500 \leq \nu \leq 2\,500\,\mu\mathrm{Hz}$) frequency regions adopted for discussion purposes by Moya et al. (2008).

ℓ	GL2 Low	GL2 Med.	GL2 High	GL4 Low	GL4 Med.	GL4 High	GL6 Low	GL6 Med.	GL6 High
0	–	0.44	3.39	–	0.39	0.49	–	0.39	0.48
1	0.67	3.61	4.39	0.64	3.39	0.64	0.64	3.39	0.64
2	1.15	7.83	3.90	1.10	7.72	0.38	1.10	7.72	0.38
3	1.63	13.11	3.55	1.57	13.12	0.34	1.57	13.12	0.34

oscillation codes. The GL4 and GL6 integrators produce almost identical results, indicating that the frequency differences between them are much smaller than their differences with GRACO.

To further illustrate the comparison between GYRE and GRACO, Fig. 2 plots the differences between the $\ell = 2$ large separations $\Delta\nu_2$ and the $\ell = 1, 3$ small separations $\delta\nu_{13}$, both as a function of frequency (see e.g. Aerts et al. 2010, for a definition and discussion of these asteroseismic parameters). These two plots are intended for direct comparison against the middle panels of figs 12 and 15, respectively, of Moya et al. (2008). They confirm that GYRE is in good agreement with the other oscillation codes. (The zig-zag feature at $\nu \approx 350\,\mu\mathrm{Hz}$ in the left-hand panel of Fig. 2 is also seen when comparing GRACO against other codes, and therefore is not due to GYRE).

5.3 Asteroseismology through the RGB Bump

As a 'first science' experiment using GYRE we now explore how the seismic properties of a 1.5 M$_\odot$ star change as it passes through the so-called red giant branch (RGB) bump. During this evolutionary phase the hydrogen-burning shell encounters the abundance discontinuity left by the convective envelope during first dredge-up, causing a temporary reversal in the star's luminosity growth as it ascends the RGB. This reversal shows up in the Hertzsprung–Russell diagram (HRD) as a narrow zig-zag in a single star's evolutionary track, and it causes a bump in the luminosity function of cluster members on the RGB, hence the name (see Salaris, Cassisi & Weiss 2002, for a more detailed discussion). Our decision to focus on the red bump has two motivations: on the one hand RGB stars are an area of especial recent interest (see Section 1) and on the other they are a challenge for any oscillation code to model due to the extremely short spatial scale of eigenfunctions in their radiative cores.

Models for the 1.5 M$_\odot$ star spanning the bump phase are obtained by running the `1.5M_with_diffusion` test-suite calculation of the MESA STAR stellar evolution code, revision 4930 (see Paxton et al. 2011, 2013). The resulting track in the HRD is plotted in Fig. 3, with the inset magnifying the RGB bump phase. For each of the ~300 models spanning this phase we use GYRE to find $\ell = 1$ modes in the frequency range $30 \leq \nu \leq 40\,\mu\mathrm{Hz}$, chosen to loosely correspond to the frequency of maximum power $\nu_{\max}$ predicted by the standard scaling relation for solar-like oscillations (e.g. Brown et al. 1991; Kjeldsen & Bedding 1995). Fig. 4 illustrates results from this exercise, plotting linear frequencies as a function of model age. The figure also shows the frequency dependence of the normalized mode inertia E (as defined by Aerts et al. 2010, their equation 3.140; not to be confused with the un-normalized inertia $\mathcal{E}$) for a single model with an age $\approx 2.705\,\mathrm{Gyr}$ which places it near the minimum luminosity encountered during the bump phase.
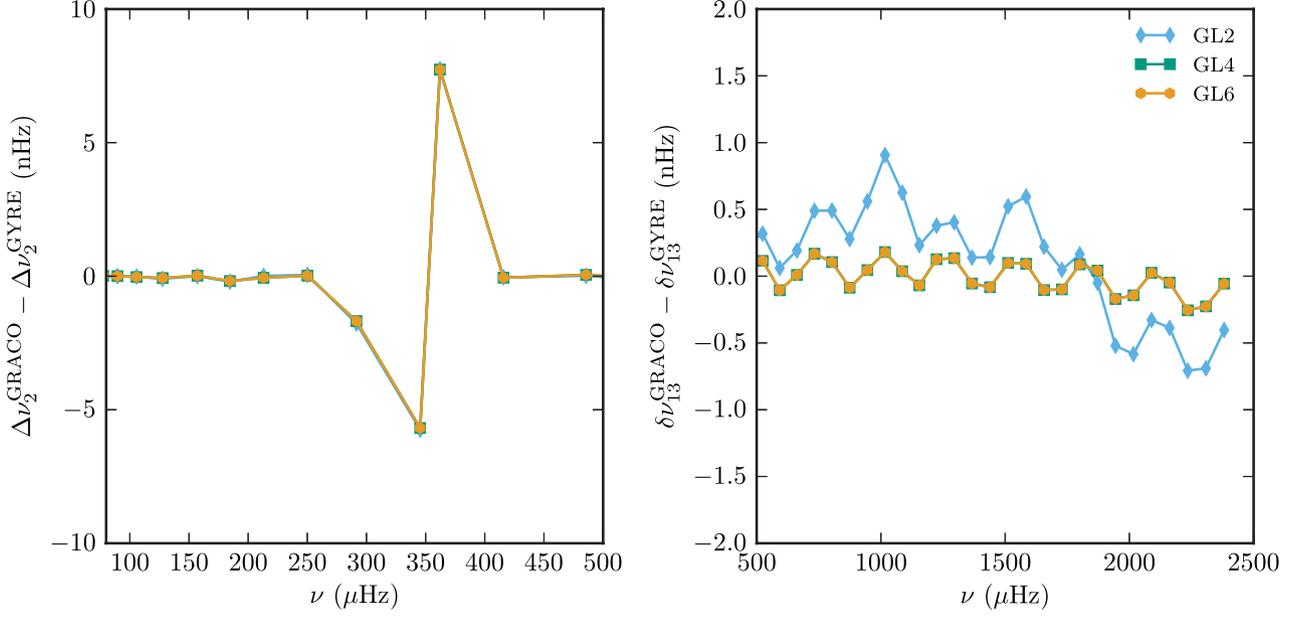
Figure 2. Differences between the GRACO and GYRE $\ell = 2$ large frequency separations (left) and $\ell = 1, 3$ small-frequency separations (right) of the M4k model, plotted as a function of linear frequency. The three curves correspond to GYRE's GL2, GL4 and GL6 Magnus integrators; in the left-hand panel, all three curves sit atop each other, while in the right-hand panel the GL4 and GL6 curves overlap. The panels should be compared against figs 12 and 15, respectively, of Moya et al. (2008).
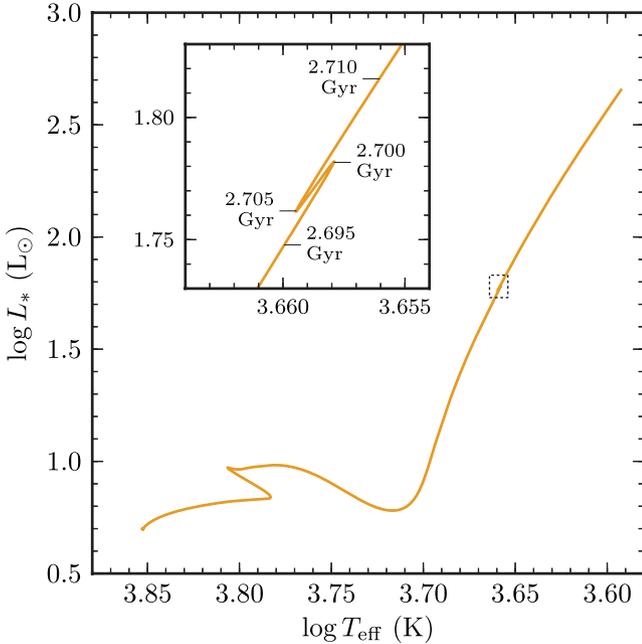

Figure 3. The evolutionary track of the 1.5 M$_\odot$ star plotted in the HRD. The inset magnifies the RGB bump phase (shown in the main diagram by the dotted rectangle), where the star's luminosity growth undergoes a temporary reversal; tick marks indicate the stellar age.

The figure depicts many of the features characteristic to red-giant oscillations. The frequency spectrum is dominated by a dense forest of g modes with radial orders $n \equiv n_{\rm p} - n_{\rm g}$ in the range $-577 \leq n \leq -401$. These modes are trapped in the radiative interior of the star where the Brunt–Väisälä frequency is large. A small subset of the modes have frequencies close to those of envelope p modes, and coupling between the two leads to the distinctive

pattern of avoided crossings (Aizenman, Smeyers & Weigert 1977) displayed in the left-hand panel of the figure. During an avoided crossing, the greatly enhanced mode amplitude in the low-density stellar envelope leads to a much-reduced normalized inertia, as can be seen in the right-hand panel. It is these low-inertia modes which dominate the observed frequency spectra of RGB stars, as they are easiest to excite to measurable amplitudes by stochastic processes (see e.g. Chaplin & Miglio 2013).

To further illustrate the change in mode properties during an avoided crossing, Fig. 5 plots the differential inertia $\mathrm{d}\mathcal{E}/\mathrm{d}x$ (which is proportional to the kinetic energy density) as a function of fractional radius x for the modes labelled 'A' and 'B' in the right-hand panel of Fig. 4. Mode A is involved in an avoided crossing and has an appreciable amplitude in both core and envelope. In contrast, Mode B is confined to the radiative core and has a negligible amplitude at the surface, accounting for its much enhanced normalized inertia compared to mode A. For both modes the radial wavelength in the radiative interior is very short due to the large Brunt–Väisälä frequency there, leading to the highly oscillatory behaviour (well resolved by GYRE) seen in the inset of the figure.

Returning to Fig. 4, the evolution through the RGB bump phase reveals itself by a temporary increase in the otherwise-decreasing frequencies of the avoided crossings. This is a direct consequence of contraction of the star's envelope during the bump luminosity reversal, which shortens the sound crossing time and therefore elevates p-mode frequencies. The frequencies of g modes are largely unaffected, as the radiative interior of the star changes only slowly during bump passage.

The same general behaviour can also be seen in Fig. 6, which plots asteroseismic observables – the p-mode large frequency separation $\Delta\nu_1$ and the g-mode period separation $\Delta\Pi_1$ – as a function of stellar age for the same $\ell = 1$ modes. The frequency separations closely follow the scaling $\Delta\nu \propto \tau^{-1}$ predicted by asymptotic relations (e.g. Aerts et al. 2010), where $\tau \equiv \sqrt{R^3/GM}$ is the star's dynamical time-scale. The gradual decrease in the period separation arises due
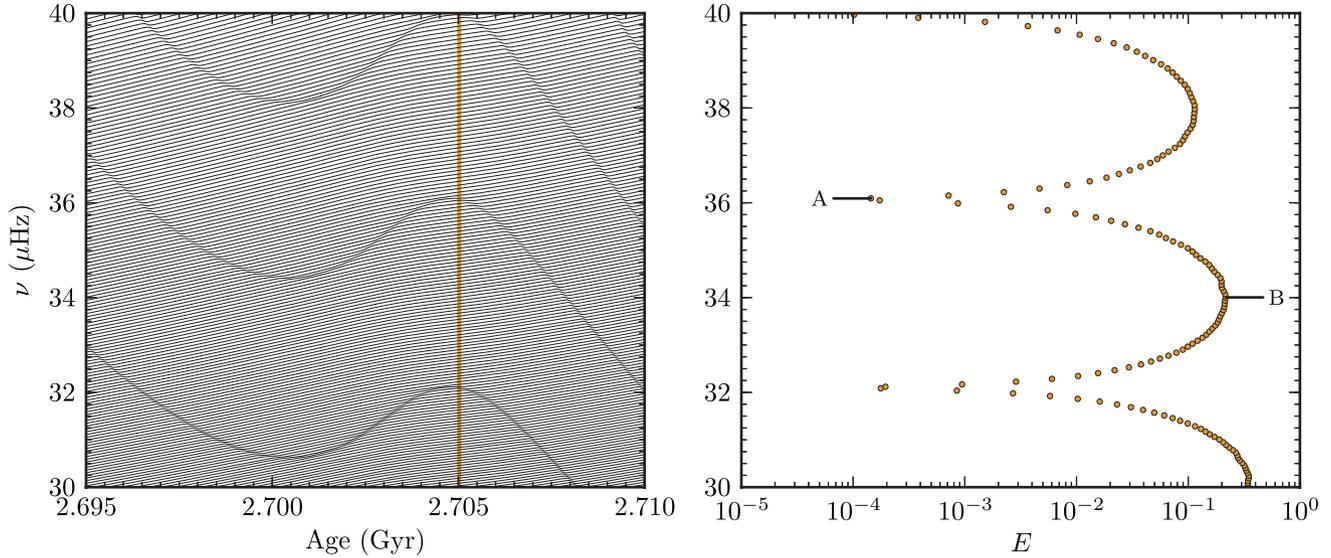
Figure 4. Linear eigenfrequencies of dipole modes of the 1.5 M$_\odot$ model plotted as a function of stellar age (left), together with the normalized mode inertia of the 2.705 Gyr model (right; also highlighted in the left-hand panel by the vertical line). The differential inertia of the modes labelled A and B is shown in Fig. 5.
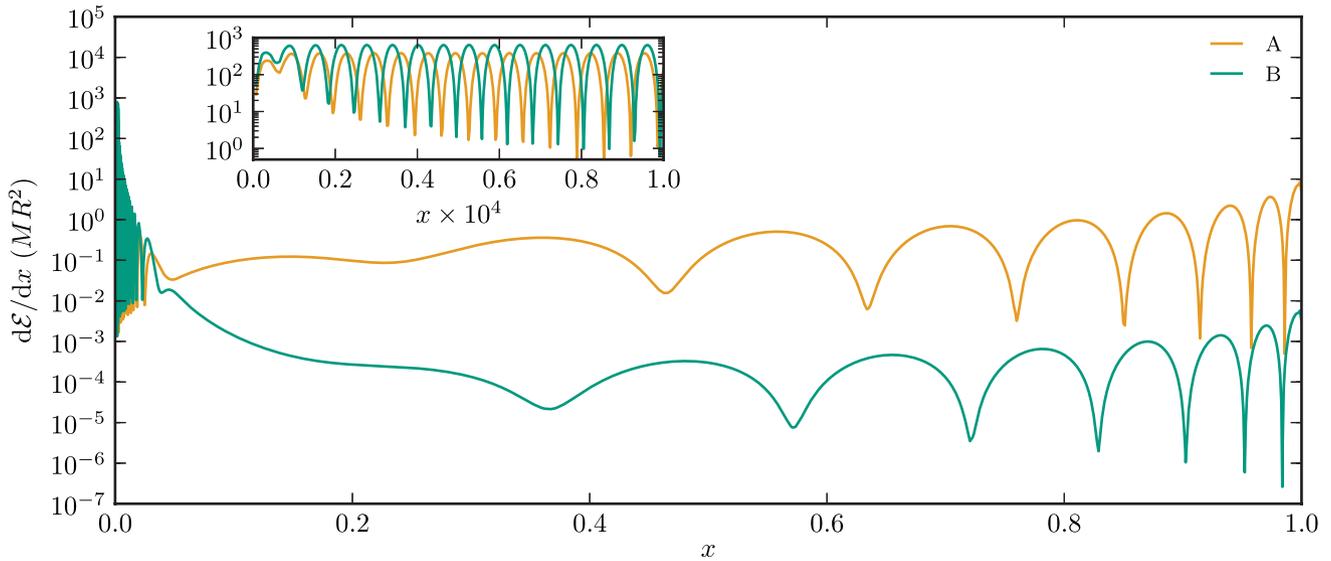

Figure 5. The differential inertia of the two modes labelled in the right-hand panel of Fig. 4, plotted as a function of fractional radius. The formal classifications are $(n_p, n_g) = (8, 476)$ for mode A and $(n_p, n_g) = (7, 505)$ for mode B. The inset magnifies the centremost region, illustrating the very small spatial scale of the modes there.

to the growing mass and shrinking radius of the degenerate helium core, which together raise the gravitational acceleration and hence Brunt–Väisälä frequency there.

Motivated by this analysis, we can speculate whether the RGB bump manifests itself in asteroseismic observables. In their presentation of initial *Kepler* observations, Kallinger et al. (2010) discuss a distinct subpopulation of RGB stars which they identify as bump stars. However, it might be argued that these stars (seen e.g. as the 'B' feature in their fig. 9b) are simply an extension of the core-helium-burning red clump stars to lower effective temperatures. To explore this issue further, it is necessary first to disentangle the RGB and the red clump. As demonstrated by Bedding et al. (2011), this separation can be achieved on the basis of measured period

separations, which are now becoming available for large numbers of stars (Stello et al. 2013).

5.4 Parallel scaling

To explore how the performance of GYRE scales on parallel architectures, we measure the execution time T of the initial root bracketing calculations (see Section 4.5) for $\ell = 1$ modes of the M4k model introduced in Section 5.2. These calculations are undertaken on a cluster of 24 nodes, each containing two four-core AMD Opteron processors, networked on an Infiniband switched fabric. Fig. 7 illustrates results from calculations using different numbers of nodes n_{node} and cores per node n_{core}, plotting the speedup and
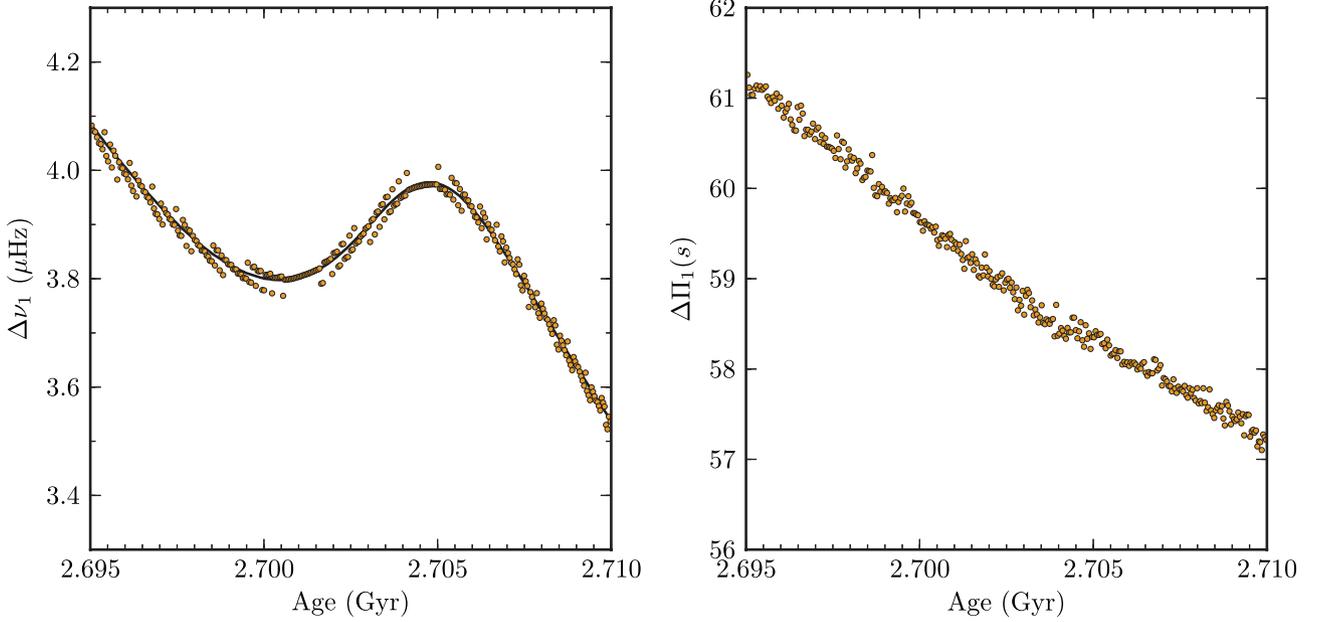
Figure 6. The p-mode large frequency separation (left) and g-mode period separation (right) of dipole modes of the 1.5 $M_\odot$ model, plotted as a function of stellar age. The large separation is measured between the minima of the normalized inertia arising from avoided crossings with the $n_p = 7$ and 8 envelope p modes, while the period separation is taken as the maximum period spacing of modes situated between these two minima. The solid curve in the left-hand panel follows the scaling $\Delta\nu \propto \tau^{-1}$ and is normalized to pass through the initial data point.
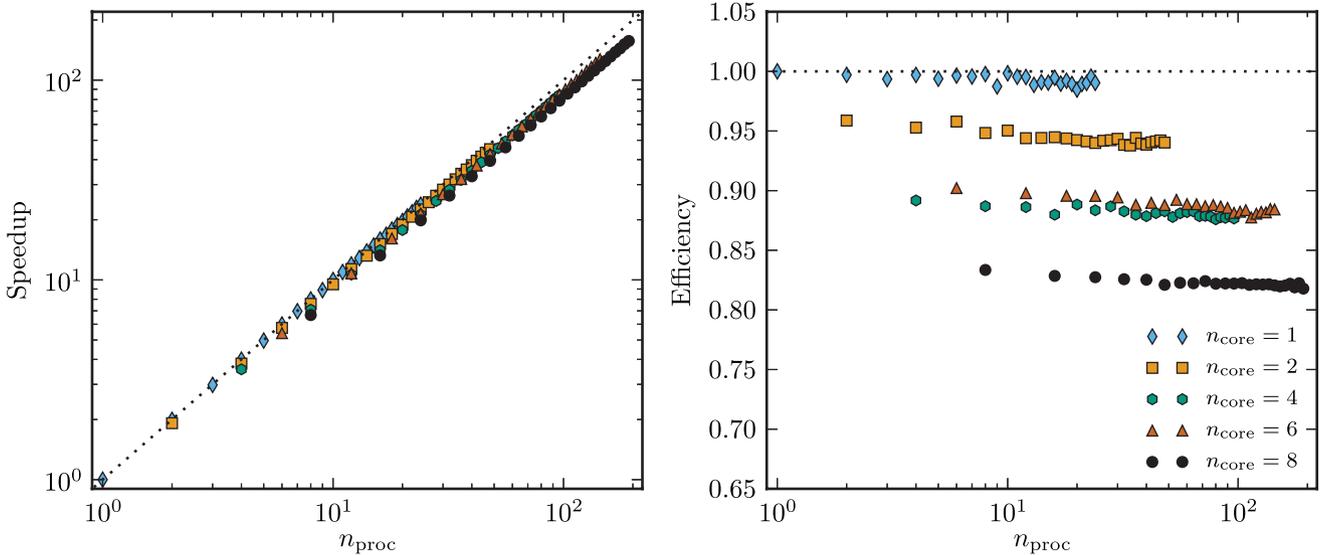
Figure 7. The speedup (left) and efficiency (right) of the root bracketing calculations, plotted as a function of the total number of processors for different combinations of cores and nodes. Each point is based on an average over five separate runs. The dotted lines show the ideal linear speedup and unity efficiency cases.

efficiency against the total number of processors $n_{proc} \equiv n_{node} \times n_{core}$. The speedup $T(1)/T(n_{proc})$ measures the overall performance of the code relative to the single-processor case, while the efficiency $T(1)/[n_{proc}T(n_{proc})]$ indicates how effectively individual processors are utilized, again relative to the single-processor case.

As discussed in Section 4, GYRE implements a hybrid approach to parallelization: OpenMP directives allow multiple cores to participate in the construction and subsequent LU factorization of a single system matrix, while MPI calls allow cluster nodes to evaluate multiple discriminants concurrently. The figure confirms that this approach is largely successful, with the speedup increasing monotonically with n_{proc}. While a modest decline in efficiency can be seen

towards larger n_{core}, this is not much cause for concern when running GYRE on today's commonly available multi-core architectures. That said, to take full advantage of next-generation architectures such as Intel's many-core Xeon Phi co-processor (used in the NSF Stampede cluster at the Texas Advanced Computing Center) it will be necessary to further improve the code's efficiency, presumably via increased OpenMP parallelization.

6 DISCUSSION AND SUMMARY

In the preceding sections we introduce a new MMS scheme for solving linear homogeneous BVPs (Section 3), together with an

oscillation code GYRE that implements this scheme to calculate eigenfrequencies and eigenfunctions of stellar models (Section 4). Initial test calculations indicate that the code is accurate, robust and makes efficient use of computational resources (Section 5).

GYRE debuts in an arena which is already well populated with oscillation codes (cf. Section 2). However, of all these codes only ADIPLS is freely available, and it is restricted to adiabatic pulsation – clearly not an optimal arrangement given the data analysis challenges now facing the field (Section 1). We are therefore pleased to make the GYRE source code open.[3] for use and distribution under the GNU General Public License. Our hope is that a community of practice (e.g. Turk 2013) will arise around the code, bringing together users and developers to shape the code's future evolution in ways that best serve the field and its participants.

As we mention in Section 4, although the present paper focuses primarily on adiabatic pulsation, GYRE can also address non-adiabatic problems. This capability requires a few minor adjustments to the MMS scheme, already implemented in GYRE, which we plan to describe in detail in a forthcoming science-oriented paper. Looking further into the future, we intend to extend GYRE to include the effects of stellar rotation – first within the traditional approximation (e.g. Townsend 2005) and then using the spherical harmonic expansion approach pioneered by Durney & Skumanich (1968) and recently adopted by various groups (e.g. Lee 2001; Reese et al. 2006; Ouazzani, Dupret & Reese 2012). The expansion approach results in BVPs with large numbers of unknowns ($4\,h$ in the adiabatic approximation when h spherical harmonics are used), and will therefore be a particularly appropriate target for testing the robustness and performance scalability of the MMS scheme and GYRE.

Alongside these code development activities, we plan to interface GYRE with the MESA STAR evolution code. GYRE can already natively read models produced by MESA STAR (cf. Section 5.3); the next step is to wrap GYRE in a callable interface and integrate it into MESA STAR's asteroseismic module. This will open up the possibility of large-scale, automated adiabatic and non-adiabatic asteroseismic analyses, in turn facilitating investigation of issues such as core rotation in RGB stars (e.g. Mosser et al. 2012), instability strips in white dwarfs (e.g. Fontaine & Brassard 2008) and the surprising incidence of opacity-driven oscillations in low-metallicity environments (e.g. Salmon et al. 2012).

ACKNOWLEDGEMENTS

---

[3] See http://www.astro.wisc.edu/~townsend/gyre/

REFERENCES

Aerts C., Christensen-Dalsgaard J., Kurtz D. W., 2010, Astronomy and Astrophysics Library: Asteroseismology. Springer, Dordrecht, p. 866

Aizenman M., Smeyers P., Weigert A., 1977, A&A, 58, 41

Anderson E. et al., 1999, LAPACK Users' Guide, 3rd edn. SIAM, Philadelphia

Ascher U. M., Mattheij R. M. M., Russell R. D., 1995, Numerical Solution of Boundary Value Problems for Ordinary Differential Equations. SIAM, Philadelphia

Baglin A., Auvergne M., Barge P., Deleuil M., Michel E. CoRoT Exoplanet Science Team, 2009, in Pont F., Sasselov D., Holman M. J., eds, Proc. IAU Symp. 253, Transiting Planets. Cambridge Univ. Press, Cambridge, p. 71

Bedding T. R. et al., 2011, Nat, 471, 608

Blackford L. S. et al., 2002, ACM Trans. Math. Softw., 28, 135

Blanes S., Casas F., Oteo J. A., Ros J., 2009, Phys. Rep., 470, 151

Borucki W. et al., 2009, in Pont F., Sasselov D., Holman M. J., eds, Proc. IAU Symp. 253, Transiting Planets. Cambridge Univ. Press, Cambridge, p. 289

Brassard P., Charpinet S., 2008, Ap&SS, 316, 107

Brent R. P., 1973, Algorithms for Minimization without Derivatives. Prentice-Hall, Englewood Cliffs

Brown T. M., Gilliland R. L., Noyes R. W., Ramsey L. W., 1991, ApJ, 368, 599

Castor J. I., 1971, ApJ, 166, 109

Chaplin W. J., Miglio A., 2013, preprint (arXiv:1303.1957)

Chaplin W. J. et al., 2011, Sci, 332, 213

Christensen-Dalsgaard J., 1980, MNRAS, 190, 765

Christensen-Dalsgaard J., 2008a, Ap&SS, 316, 13

Christensen-Dalsgaard J., 2008b, Ap&SS, 316, 113

Christensen-Dalsgaard J., Mullan D. J., 1994, MNRAS, 270, 921

Christensen-Dalsgaard J., Thompson M. J., 2011, in Brummell N. H., Brun A. S., Miesch M. S., Ponty Y., eds, Proc. IAU Symp. 271, Astrophysical Dynamics: from Stars to Galaxies. Cambridge Univ. Press, Cambridge, p. 32

Cowling T. G., 1941, MNRAS, 101, 367

Cox J. P., 1980, Theory of Stellar Pulsation. Princeton Univ. Press, Princeton

Dagum L., Menon R., 1998, IEEE Comput. Sci. Eng., 5, 46

De Ridder J. et al., 2009, Nat, 459, 398

Dongarra J. J., Otto S. W., Snir M., Walker D., 1995, Commun. ACM, 39, 84

Dupret M. A., 2001, A&A, 366, 166

Durney B., Skumanich A., 1968, ApJ, 152, 255

Dziembowski W. A., 1971, Acta Astron., 21, 289

Eckart C., 1960, Hydrodynamics of Oceans and Atmospheres. Pergammon Press, Oxford

Fontaine G., Brassard P., 2008, PASP, 120, 1043

Fourer R., 1984, SIAM Rev., 26, 1

Gabriel M., Noels A., 1976, A&A, 53, 149

Gautschy A., Glatzel W., 1990, MNRAS, 245, 154

Gilliland R. L. et al., 2010, PASP, 122, 131

Golub G. H., van Loan C. F., 1996, Matrix Computations. Johns Hopkins Univ. Press, Baltimore

Grigahcène A. et al., 2010, ApJ, 713, L192

Henyey L. G., Forbes J. E., Gould N. L., 1964, ApJ, 139, 306

Hoog F. R. D., Mattheij R. M. M., 1987, SIAM J. Numer. Anal., 24, 89

Hurley M., Roberts P. H., Wright K., 1966, ApJ, 143, 535

Kallinger T. et al., 2010, A&A, 522, A1

Kjeldsen H., Bedding T. R., 1995, A&A, 293, 87

Kjeldsen H., Christensen-Dalsgaard J., Handberg R., Brown T. M., Gilliland R. L., Borucki W. J., Koch D., 2010, Astron. Nachr., 331, 966

Lebreton Y. et al., 2008, Ap&SS, 316, 1

Lee U., 2001, ApJ, 557, 311

Magnus W., 1954, Commun. Pure Appl. Math., 7, 649

Matthews J. M., 2007, Commun. Astroseismol., 150, 333

Metcalfe T. S., Creevey O. L., Christensen-Dalsgaard J., 2009, ApJ, 699, 373

Michel E. et al., 2008, Sci, 322, 558

Moler C., Van Loan C., 2003, SIAM Rev., 45, 3

Monteiro M. J. P. F. G., 2008, Ap&SS, 316, 121

Mosser B. et al., 2012, A&A, 548, A10

Moya A., Garrido R., 2008, Ap&SS, 316, 129

Moya A. et al., 2008, Ap&SS, 316, 231

Osaki J., 1975, PASJ, 27, 237

Osaki Y., Hansen C. J., 1973, ApJ, 185, 277

Ouazzani R.-M., Dupret M.-A., Reese D. R., 2012, A&A, 547, A75

Paxton B., Bildsten L., Dotter A., Herwig F., Lesaffre P., Timmes F., 2011, ApJS, 192, 3

Paxton B. et al., 2013, ApJS, 208, 4

Pekeris C. L., 1938, ApJ, 88, 189

Provost J., 2008, Ap&SS, 316, 135

Reese D., Lignières F., Rieutord M., 2006, A&A, 455, 621

Roxburgh I. W., 2008, Ap&SS, 316, 141

Salaris M., Cassisi S., Weiss A., 2002, PASP, 114, 375

Salmon S., Montalbán J., Morel T., Miglio A., Dupret M.-A., Noels A., 2012, MNRAS, 422, 3460

Scuflaire R., 1974, A&A, 36, 107

Scuflaire R., Montalbán J., Théado S., Bourge P.-O., Miglio A., Godart M., Thoul A., Noels A., 2008, Ap&SS, 316, 149

Smeyers P., 1966, Ann. Astrophys., 29, 539

Smeyers P., 1967, Bull. Soc. R. Sci. Liège, 36, 357

Smeyers P., van Hoolst T., 2010, Linear Isentropic Oscillations of Stars: Theoretical Foundations. Springer, Dordrecht

Steffen M., 1990, A&A, 239, 443

Stello D. et al., 2013, ApJ, 765, L41

Suárez J. C., Goupil M. J., 2008, Ap&SS, 316, 155

Suran M. D., 2008, Ap&SS, 316, 163

Takata M., 2006, PASJ, 58, 893

Townsend R. H. D., 2005, MNRAS, 360, 465

Turk M. J., 2013, preprint (arXiv:1301.7064)

Unno W., Osaki Y., Ando H., Saio H., Shibahashi H., 1989, Nonradial Oscillations of Stars, 2nd edn. Univ. Tokyo Press, Tokyo

Valsecchi F., Farr W. M., Willems B., Kalogera V., 2013, ApJ, 773, 39

Walker G. et al., 2003, PASP, 115, 1023

Wright S. J., 1992, SIAM J. Sci. Stat. Comput., 13, 742

Wright S. J., 1994, Numer. Math., 67, 521

APPENDIX A: PULSATION EQUATIONS

This appendix briefly summarizes the pulsation BVP solved by GYRE in the adiabatic case. The independent variable is the fractional radius $x = r/R$, with r being the radial coordinate and R the stellar radius, while the components of the dependent variable vector $\mathbf{y}$ are

$$y_1 = \frac{\xi_r}{r} x^{2-\ell}, \qquad y_2 = \frac{1}{gr}\left(\frac{p'}{\rho} + \Phi'\right) x^{2-\ell},$$

$$y_3 = \frac{1}{gr}\Phi' x^{2-\ell}, \qquad y_4 = \frac{1}{g}\frac{d\Phi'}{dr} x^{2-\ell}. \tag{A1}$$

Here, the symbols have the same meaning as in Unno et al. (1989); specifically, ξ_r is the radial displacement perturbation and p' and Φ' are the Eulerian perturbations to the pressure and gravitational potential, respectively. These definitions mirror the dimensionless variables introduced by Dziembowski (1971), except that we introduce the scaling $x^{2-\ell}$ to ensure that the variables approach constant values at the origin $x = 0$ – desirable behaviour from a numerical perspective.

Given the definitions above, the differential equations governing linear, adiabatic non-radial oscillations can be written in the canonical form (1) with a Jacobian matrix

$$\mathbf{A} = x^{-1}\tilde{\mathbf{A}}, \tag{A2}$$

where

$$\tilde{\mathbf{A}} = \begin{pmatrix} \frac{V}{\Gamma_1} - 1 - \ell & \frac{\ell(\ell+1)}{c_1\omega^2} - \frac{V}{\Gamma_1} & \frac{V}{\Gamma_1} & 0 \\ c_1\omega^2 - A^* & A^* - U + 3 - \ell & -A^* & 0 \\ 0 & 0 & 3 - U - \ell & 1 \\ UA^* & \frac{UV}{\Gamma_1} & \ell(\ell+1) - \frac{UV}{\Gamma_1} & -U + 2 - \ell \end{pmatrix}. \tag{A3}$$

The dimensionless oscillation frequency ω is related to the linear frequency ν via

$$\omega = 2\pi\nu\sqrt{\frac{R^3}{GM}}, \tag{A4}$$

and the other variables again have the same meaning as in Unno et al. (1989).

The requirement that solutions remain regular at the centre leads to the inner boundary conditions

$$\mathbf{B}^{\mathrm{a}} = \begin{pmatrix} c_1\omega^2 & -\ell & 0 & 0 \\ 0 & 0 & \ell & -1 \end{pmatrix}, \tag{A5}$$

evaluated at $x = x^{\mathrm{a}} = 0$. Likewise, the requirement that the Lagrangian pressure perturbation δp vanishes at the stellar surface, and that Φ' vanishes at infinity, leads to the outer boundary conditions

$$\mathbf{B}^{\mathrm{b}} = \begin{pmatrix} 1 & -1 & 1 & 0 \\ U & 0 & \ell+1 & 1 \end{pmatrix} \tag{A6}$$

evaluated at $x = x^{\mathrm{b}} = 1$. GYRE offers the option of outer boundary conditions based on more-sophisticated treatments of the stellar atmosphere; these include the prescriptions by Dziembowski (1971) and Unno et al. (1989). However, for all calculations presented in Section 5 the zero-δp condition incorporated in equation (A6) is adopted.

APPENDIX B: MAGNUS MATRICES

For convenience, this section presents expressions for the Magnus matrices (taken from Blanes et al. 2009) which are used by GYRE's GL4 and GL6 Magnus integrators.

B1 GL4 Magnus integrator

Using a fourth-order Gauss–Legendre quadrature, the Magnus matrix in the kth subinterval is approximated as

$$\Omega(x^{k+1}; x^k) \approx \alpha_{4,1} - \frac{1}{12}[\alpha_{4,1}, \alpha_{4,2}] + \mathcal{O}[(\Delta x^k)^5]. \tag{B1}$$

Here,

$$\alpha_{4,1} = \frac{\Delta x^k}{2}(\mathbf{A}_1 + \mathbf{A}_2), \qquad \alpha_{4,2} = \Delta x^k \sqrt{3}(\mathbf{A}_2 - \mathbf{A}_1) \tag{B2}$$

and $\mathbf{A}_i \equiv \mathbf{A}(x_i^k)$ ($i = 1, 2$) are the Jacobian matrices evaluated at the two Gauss–Legendre nodes within the subinterval,

$$x_1^k = x^k + \left(\frac{1}{2} - \frac{\sqrt{3}}{6}\right)\Delta x^k, \qquad x_2^k = x^k + \left(\frac{1}{2} + \frac{\sqrt{3}}{6}\right)\Delta x^k. \tag{B3}$$

Note that the above expression for $\alpha_{4,2}$ corrects an error in equation (253) of Blanes et al. (2009).

B2 GL6 Magnus integrator

Using a sixth-order Gauss–Legendre quadrature, the Magnus matrix in the kth subinterval is likewise approximated as

$$\mathbf{\Omega}(x^{k+1}; x^k) \approx \boldsymbol{\alpha}_{6,1} + \frac{1}{12}\boldsymbol{\alpha}_{6,3}$$

$$+ \frac{1}{240}[-20\boldsymbol{\alpha}_{6,1} - \boldsymbol{\alpha}_{6,3} + \mathbf{C}_1, \boldsymbol{\alpha}_{6,2} + \mathbf{C}_2] + \mathcal{O}[(\Delta x^k)^7]. \qquad \text{(B4)}$$

Here,

$$\boldsymbol{\alpha}_{6,1} = \Delta x^k \mathbf{A}_2, \qquad \boldsymbol{\alpha}_{6,2} = \frac{\Delta x^k \sqrt{15}}{3}(\mathbf{A}_3 - \mathbf{A}_1),$$

$$\boldsymbol{\alpha}_{6,3} = \frac{10\Delta x^k}{3}(\mathbf{A}_3 - 2\mathbf{A}_2 + \mathbf{A}_1), \qquad \text{(B5)}$$

while

$$\mathbf{C}_1 = [\boldsymbol{\alpha}_{6,1}, \boldsymbol{\alpha}_{6,2}], \qquad \mathbf{C}_2 = -\frac{1}{60}[\boldsymbol{\alpha}_{6,1}, 2\boldsymbol{\alpha}_{6,3} + \mathbf{C}_1] \qquad \text{(B6)}$$

and $\mathbf{A}_i \equiv \mathbf{A}(x_i^k)$ ($i = 1, 2, 3$) are the Jacobian matrices evaluated at the three Gauss–Legendre nodes within the subinterval,

$$x_1^k = x^k + \left(\frac{1}{2} - \frac{\sqrt{15}}{10}\right)\Delta x^k, \qquad x_2^k = x^k + \frac{\Delta x^k}{2},$$

$$x_3^k = x^k + \left(\frac{1}{2} + \frac{\sqrt{15}}{10}\right)\Delta x^k. \qquad \text{(B7)}$$