An energy-based discontinuous Galerkin discretization of the elastic wave equation in second order form

Daniel Appelö[a,1,*], Thomas Hagstrom[b,2]

[a] *Department of Applied Mathematics, University of Colorado, Boulder, University of Colorado, 526 UCB Boulder, CO 80309.*
[b] *Department of Mathematics, Southern Methodist University, PO Box 750156, Dallas, TX 75275.*

Abstract

We present an application of our general formulation [1] to construct energy based, arbitrary order accurate, discontinuous Galerkin spatial discretizations of the linear elastic wave equation. The resulting methods are stable and, depending on the choice of numerical flux, conserve or dissipate the elastic energy. The performance of the method is demonstrated for problems with manufactured and exact solutions. Applications to more realistic problems are also presented. Implementations of the methods are freely available at [2].

Keywords:

1. Introduction

Accurate propagation of linear elastic waves in heterogeneous material and complex geometry is important in many fields such as design of mechanical structures, non-destructive testing, civil engineering and seismic exploration. In general, the challenge in numerically propagating waves is the long distance of propagation, requiring high order methods able to control the growth of dispersion errors.

Discontinuous Galerkin (dG) methods have excellent dispersion properties. The research devoted to dG has been substantial over the last decade and a half and we will not attempt to review the whole literature here (the textbook by Hesthaven and Warburton, [3], is an excellent reference) but limit our discussion to methods pertaining to the linear time dependent elastic wave equation.

Discontinuous Galerkin methods for elastic waves either discretize the governing equations in some first order formulation (in space and time), e.g. velocity-strain or velocity-stress, or directly work with the equations in second order form. Examples of methods based on first order systems which construct the numerical fluxes based on the Riemann problem are [4, 5, 6]. Another method that works with a first order system is the staggered dG method in [7]. Methods that work with the equations in second order form are the interior penalty method [8] and the symmetric interior penalty method [9], as well as the local discontinuous Galerkin methods; see [10] for a stability analysis applicable to a wide range of methods of this type. Another class of methods is the so-called space-time discontinuous Galerkin methods; see for example the early application in [11] and more recently [12, 13]. We postpone to §4 comparisons with other discretization techniques after our method, which differs from all of these, has been fully described.

In our proposed approach we approximate the equations in second order form in space but introduce the velocity to reduce the order to one in time. As the velocity is naturally connected to the kinetic

*Corresponding author
Email addresses: daniel.appelo@colorado.edu (Daniel Appelö), hagstrom@smu.edu (Thomas Hagstrom)

[1] Supported in part by NSF Grant DMS-1319054. Any conclusions or recommendations expressed in this paper are those of the author and do not necessarily reflect the views NSF.

[2] Supported in part by ARO Contract W911NF-09-1-0344 and NSF Grants OCI-0904773, DMS-1418871. Any conclusions or recommendations expressed in this paper are those of the author and do not necessarily reflect the views of ARO or NSF.

energy and the displacement, through the stress, is connected to the potential energy, this form allows for a dG formulation that mimics the dynamics of the energy of the system. The proposed method is a direct application of our general formulation for wave equations in second order form [1]. The crucial (and to our knowledge novel) step in our formulation is to test the equation governing the time derivatives of the displacements, not directly against a test function but against a quantity related to the potential energy. This immediately leads to an energy identity for the dG formulation under very general assumptions on the approximation spaces. This test quantity is often invariant to certain transformations and therefore does not fully specify the time derivatives of the displacements. However we show that it is straightforward to add independent equations to close the system without affecting the energy identity.

Finite difference methods are also often used for wave propagation problems. Until recently an obstacle to constructing high order and curvilinear solvers for the second order formulation was the stable enforcement of traction boundary conditions, but there are now multiple options, [14, 15, 16], all relying on summation by parts techniques to prove energy stability. Although these methods typically have a smaller spectral radius than dG and spectral element methods, they may become less efficient when solving problems where free surface or interface waves are important.

The classic parameter used to measure resolution requirements for finite difference methods, points per wavelength (PPW) introduced in the seminal paper [17] by Kreiss and Oliger, has recently been found to be an inadequate measure of the resolution requirements when surface waves or interfaces are present [18, 14]. The analysis of Kreiss and Petersson, [18], uses a modified equation approach to show that the number of points required to reach a fixed error at a fixed time scales with $(\lambda/\mu)^{1/p}h^p$ if a method that approximates the surface waves to pth order of accuracy is used. The results presented in the experiments section indicates that *this effect is not important for the proposed discretization* (and probably not for other dG methods either). As the analysis in [18] is based on a modified equation it does apply to dG as well, however it is well known that the wave speeds in dG discretizations are approximated with about twice the order of the method [19, 9]. Now the dispersion error in the surface wave for dG is, say, $(\lambda/\mu)^{\frac{1}{2p}}h^{2p}$ which is still very small compared to h^p unless λ/μ is (unphysically) large. We believe that this is the reason why we do not observe any degradation when $\lambda/\mu \gg 1$[3].

As the basic theoretical analysis and framework of the general method has already been presented in [1] we focus here on the formulation for the elastic wave equation and the method's performance for a sequence of numerical examples. Through these experiments, whose computer implementations are freely available from [2], we hope to convince practitioners that our method is worth considering. Here we work in two dimensions but note that the method can easily be generalized to three dimensions.

The rest of the paper is organized as follows. In Section 2 we present the governing equations and discuss the dynamics of the energy of a solid body subject to boundary conditions. In Section 3 we present the variational formulation and discuss how to add independent equations to compensate for the invariants of the elastic energy. We also discuss how to choose the numerical fluxes at inter-element boundaries and at physical boundaries. Section 4 briefly outlines the discretization and the implementation of the method. In Section 5 we present a comprehensive collection of numerical experiments illustrating the method's performance for manufactured solutions, for classical problems such as the propagation of Rayleigh, Lamb and Stoneley waves, as well as mode conversion at an interface between two solids. We conclude the numerical experiments with two applications: wave propagation in a slab with a stiff inclusion and an application in uncertainty quantification for non-destructive testing. The last section summarizes the paper and discusses possible extensions and improvements.

Finally, we again note that the computer codes used in the examples below are freely available from [2].

2. The Elastic Wave Equation

Let $\rho = \rho(x_1, x_2)$ be the density and $u_i = u_i(x_1, x_2, t)$, $i = 1, 2$ be the displacements in the x_1 and x_2 direction. Assuming small deformations the the linear isotropic elastic wave equation, governing the time

[3] As far as we know, the dG dispersion relation for elastic surface or interface waves has not been studied yet but we expect it to behave similarly as for other waves.

evolution of the displacements, can then be written in second order form

$$\rho\frac{\partial^2 u_1}{\partial t^2} = \nabla\cdot\vec{F_1}, \quad (x_1,x_2)\in S\subset\mathbb{R}^2,\, t>0, \tag{1}$$

$$\rho\frac{\partial^2 u_2}{\partial t^2} = \nabla\cdot\vec{F_2}, \quad (x_1,x_2)\in S\subset\mathbb{R}^2,\, t>0. \tag{2}$$

Here the functions $\vec{F_i}, i=1,2$ are composed of the elements of the stress tensor

$$\vec{F_1} = \begin{pmatrix} (2\mu+\lambda)\frac{\partial u_1}{\partial x_1} + \lambda\frac{\partial u_2}{\partial x_2} \\ \mu(\frac{\partial u_1}{\partial x_2} + \frac{\partial u_2}{\partial x_1}) \end{pmatrix}, \quad \vec{F_2} = \begin{pmatrix} \mu(\frac{\partial u_1}{\partial x_2} + \frac{\partial u_2}{\partial x_1}) \\ \lambda\frac{\partial u_1}{\partial x_1} + (2\mu+\lambda)\frac{\partial u_2}{\partial x_2} \end{pmatrix},$$

expressed here using the Lamé parameters $\lambda=\lambda(x_1,x_2)$ and $\mu=\mu(x_1,x_2)$.

To close the system we prescribe initial conditions for the displacements and velocities

$$u_i(x_1,x_2,0) = g_i(x_1,x_2), \quad \frac{\partial u_i}{\partial t}(x_1,x_2,0) = h_i(x_1,x_2), \quad i=1,2, \tag{3}$$

and boundary conditions

$$\gamma_i\frac{\partial u_i}{\partial t} + \kappa_i R(\vec{n}\cdot\vec{F_i}) = 0, \quad (x_1,x_2)\in\partial S, \quad i=1,2. \tag{4}$$

Here R is a constant to make the equations dimensionally consistent, $\vec{n}$ is the outward pointing normal, $\kappa_i\geq 0$, $\gamma_i\geq 0$, and we normalize $\gamma_i^2+\kappa_i^2=1$. The choice $\kappa_i=0$ corresponds to a homogenous Dirichlet boundary condition on $\frac{\partial u_i}{\partial t}$; $\gamma_i=0$ corresponds to a traction free boundary, and the choice $\gamma_i\kappa_i>0$ will dissipate the energy of the system and can be thought of as a low order non-reflecting boundary condition.[4]

Remark 1. *In the above equations we have chosen, for the brevity of the presentation, not to include body forces but we note that it is straightforward to include such forces; see [1] . Indeed, the example presented in §5.2.2 includes body and boundary forcing.*

The equations (1)-(2) can be derived by taking the variational derivative of the potential energy density:

$$G = \frac{\lambda}{2}\left(\frac{\partial u_1}{\partial x_1} + \frac{\partial u_2}{\partial x_2}\right)^2 + \frac{2\mu}{2}\left(\left(\frac{\partial u_1}{\partial x_1}\right)^2 + \left(\frac{\partial u_2}{\partial x_2}\right)^2 + \frac{1}{2}\left(\frac{\partial u_1}{\partial x_2} + \frac{\partial u_2}{\partial x_1}\right)^2\right). \tag{5}$$

Equations (1)-(2) then take the form

$$\rho\frac{\partial^2 u_i}{\partial t^2} = \sum_{k=1}^{2}\frac{\partial}{\partial x_k}\left(\frac{\partial G}{\partial u_{i,k}}\right), \quad i=1,2, \tag{6}$$

where $u_{i,k}$ denotes differentiation of u_i in the kth space coordinate. The equations (1)-(2) and (6) are, of course, the same. In particular note that

$$\vec{F_i} = \begin{bmatrix} \frac{\partial G}{\partial u_{i,1}} \\ \frac{\partial G}{\partial u_{i,2}} \end{bmatrix}, \quad i=1,2.$$

[4]For the general theory of well-posed boundary conditions for second order wave equations see [20, 21]. The theory is applied to the traction-free boundary condition in elastodynamics in [18]. Generally speaking, in addition to the energy estimates following from (4), one must analyze so-called generalized eigenvalues, which is beyond the scope of this paper.

2.1. The Energy and Energy Derivative

Our variational formulation will be motivated by the continuous energy and its change in time and we therefore now recall the basic facts about the energy of an elastic system. The total energy in a domain, S, with sufficiently smooth boundary is obtained by integrating the kinetic and potential energy density

$$E(t) = \int_S \sum_{i=1}^{2} \frac{\rho}{2} \left(\frac{\partial u_i}{\partial t} \right)^2 + G \, dS, \tag{7}$$

and the rate of change of that energy is

$$E'(t) = \int_S \sum_{i=1}^{2} \left[\frac{\partial u_i}{\partial t} \rho \frac{\partial^2 u_i}{\partial t^2} + \sum_{k=1}^{2} \frac{\partial^2 u_i}{\partial t \partial x_k} \left(\frac{\partial G}{\partial u_{i,k}} \right) \right] dS.$$

Assuming classical solutions, integrating by parts, and using (6) we find that $E'(t)$ is completely determined by the contribution on the boundary

$$E'(t) = \int_{\partial S} \sum_{i=1}^{2} \frac{\partial u_i}{\partial t} \left[\sum_{k=1}^{2} n_k \frac{\partial G}{\partial u_{i,k}} \right] ds, \tag{8}$$

or equivalently using $\vec{F}_1$ and $\vec{F}_2$

$$E'(t) = \int_{\partial S} \sum_{i=1}^{2} \frac{\partial u_i}{\partial t} \left[\vec{n} \cdot \vec{F}_i \right] ds. \tag{9}$$

Now, using the boundary conditions (with $1 = \gamma_i^2 + \kappa_i^2$) we find that the rate of change of the energy is

$$E'(t) = - \int_{\partial S} \sum_{i=1}^{2} \gamma_i \kappa_i \left[\frac{1}{R} \left(\frac{\partial u_i}{\partial t} \right)^2 + R \left(\vec{n} \cdot \vec{F}_i \right)^2 \right] ds. \tag{10}$$

Thus, as mentioned above, the rate of change of the energy will depend on the parameters in the boundary conditions; it is conserved when traction free or homogenous Dirichlet boundary conditions are used and dissipated when $\gamma_i \kappa_i > 0$.

We note that the energy balance holds under weaker conditions, and in particular it will be conserved across material interfaces and for solutions with less than C^2 regularity in the linear case. Here we will apply the method to problems with material discontinuities and singular sources, but restrict our analysis to problems with smooth solutions.

Our goal is to introduce a discontinuous Galerkin method whose (discrete) solution obeys the same energy estimate as the continuous solution. To do this we first choose an appropriate variational formulation on a single element which after integration by parts (on each element) satisfies an energy estimate of the type (8) with boundary terms from interior and boundary faces. We then choose numerical fluxes, that couple interior elements in a stable fashion, and specify boundary "ghost states" that reproduce the right hand side of (10). Note that our approximations are polynomials within each element, but will be discontinuous across element boundaries. As in [8], we can rigorously obtain a weak formulation of the elastic wave equation for displacements in $(H^m)^d$ for $d \leq 3$ and $m > 3/2$, so that their traces are elements of $(H^1)^d$. The discretizations will then be defined in broken Sobolev spaces, as is typical for dG methods.

Our method is based on approximating the displacements, u_i, and the velocities, v_i. We therefore write the governing equations as a first order system in time

$$\frac{\partial u_i}{\partial t} - v_i = 0, \qquad i = 1, 2, \tag{11}$$

$$\rho \frac{\partial v_i}{\partial t} = \sum_k \frac{\partial}{\partial x_k} \left(\frac{\partial G}{\partial u_{i,k}} \right), \qquad i = 1, 2. \tag{12}$$

3. An Energy Based dG Method for the Elastic Wave Equation

Let the finite element mesh

$$\bar{\Omega} = \bigcup_j \Omega_j,$$

be a discretization of S consisting of geometry-conforming and nonoverlapping mesh elements Ω_j with piecewise smooth boundaries. Here we will consider the elements to be quadrilaterals, possibly with curved sides, but note that the method can equally well be implemented using simplices. We also note that in our implementation we use conforming grids but that this is not required.

We denote the piecewise tensor product polynomial approximations to the displacements $\vec{u} = [u_1, u_2]^T$ and the velocities $\vec{v} = [v_1, v_2]^T$ by $\vec{u}^h = [u_1^h, u_2^h]^T$ and $\vec{v}^h = [v_1^h, v_2^h]^T$.

On a single element Ω_j the approximation $(\vec{u}^h, \vec{v}^h)$ will be a tensor product polynomial in the space $(\mathbb{Q}^{q_u}(\Omega_j))^2 \times (\mathbb{Q}^{q_v}(\Omega_j))^2$, where $\mathbb{Q}^q$ denotes vector valued tensor product polynomials of degree q in each coordinate on the reference element. Let $P_i(z)$ be a hierarchical polynomial basis, then we take

$$\phi_{i,j,1} = \begin{pmatrix} P_i(r)P_j(s) \\ 0 \end{pmatrix}, \quad \phi_{i,j,2} = \begin{pmatrix} 0 \\ P_i(r)P_j(s) \end{pmatrix}, \quad i,j = 0,\ldots,q,$$

to be the basis for $(\mathbb{Q}^q)^2$.

3.1. A Variational Formulation Based on the Energy

We are now ready to formulate our Galerkin variational problem based on testing equations (11) and (12). Equation (12) is tested against the test functions as usual but in order to obtain an energy estimate mimicking the continuous case we test equation (11) in a non-standard way.

Problem 1. *On each element, require that for all test functions*

$$(\vec{\phi}_u, \vec{\phi}_v) \in (\mathbb{Q}^{q_u}(\Omega_j))^2 \times (\mathbb{Q}^{q_v}(\Omega_j))^2,$$

the following variational formulation holds:

$$\int_{\Omega_j} \sum_{i=1}^{2}\sum_{k=1}^{2} \left(\frac{\partial}{\partial x_k} \left(\frac{\partial G}{\partial \phi_{u_i,k}} (\nabla \vec{\phi}_u) \right) \right) \left(\frac{\partial u_i^h}{\partial t} - v_i^h \right) = \int_{\partial\Omega_j} \sum_{i=1}^{2}\sum_{k=1}^{2} \left(\frac{\partial G}{\partial \phi_{u_i,k}} (\nabla \vec{\phi}_u)\, n_k \right) \left(\frac{\partial u_i^h}{\partial t} - v_i^* \right), \tag{13}$$

$$\int_{\Omega_j} \sum_{i=1}^{2} \rho \phi_{v_i} \frac{\partial v_i^h}{\partial t} + \sum_{k=1}^{2} \frac{\partial \phi_{v_i}}{\partial x_k} \frac{\partial G}{\partial u_{i,k}} (\nabla \vec{u}^h) = \int_{\partial\Omega_j} \sum_{i=1}^{2} \phi_{v_i} (\vec{w}_i^* \cdot \vec{n}). \tag{14}$$

Integrating the first equation by parts yields an alternative variational formulation

$$\int_{\Omega_j} \sum_{i=1}^{2}\sum_{k=1}^{2} \left(\frac{\partial G}{\partial \phi_{u_i,k}} (\nabla \vec{\phi}_u) \right) \frac{\partial}{\partial x_k} \left(\frac{\partial u_i^h}{\partial t} - v_i^h \right) = \int_{\partial\Omega_j} \sum_{i=1}^{2}\sum_{k=1}^{2} \left(\frac{\partial G}{\partial \phi_{u_i,k}} (\nabla \vec{\phi}_u)\, n_k \right) \left(v_i^* - v_i^h \right). \tag{15}$$

The formulations (13) and (15) are sometimes referred to as the strong and weak formulation, [3]. The notation refers to the requirement of smooth test functions in the first formulation and the possibility of non-smooth test functions in the latter. As already mentioned, here we use smooth test and trial functions and the formulations are thus interchangeable.

Now, for a standard discontinuous Galerkin formulation the next step would be to choose the numerical fluxes $v_1^*, v_2^*, \vec{w}_1^*, \vec{w}_2^*$ so that the method is stable and consistent. Here however as we have tested equation (11) against the right hand side of (6) with the *derivatives* of the test function $\vec{\phi}_u$ as arguments we will obviously "lose" equations for the constant elements of $(\mathbb{Q}^{q_u}(\Omega_j))^2$. Perhaps less obvious is that we will lose yet another equation. To see this more clearly we return to the starting point of our method, the potential energy density.

3.2. Invariants of the Elastic Energy

The elastic energy (7) has three invariants. These invariants correspond to scalar multiples of the time independent solutions $(u_1, u_2) = (1, 0)$, $(u_1, u_2) = (0, 1)$, and $(u_1, u_2) = (x_2, -x_1)$, having in common that they do not contribute to the potential energy density G.

To see how this affects the variational formulation note that on an element Ω_j the local mass matrix corresponding to (15) will be obtained from

$$\int_{\Omega_j} \left((2\mu+\lambda)\frac{\partial\phi_1}{\partial x_1} + \lambda\frac{\partial\phi_2}{\partial x_2}\right)\frac{\partial^2 u_1}{\partial x_1 \partial t} + \left(\mu\left(\frac{\partial\phi_1}{\partial x_2} + \frac{\partial\phi_2}{\partial x_1}\right)\right)\frac{\partial^2 u_1}{\partial x_2 \partial t} + \left(\mu\left(\frac{\partial\phi_1}{\partial x_2} + \frac{\partial\phi_2}{\partial x_1}\right)\right)\frac{\partial^2 u_2}{\partial x_1 \partial t} + \left(\lambda\frac{\partial\phi_1}{\partial x_1} + (2\mu+\lambda)\frac{\partial\phi_2}{\partial x_2}\right)\frac{\partial^2 u_2}{\partial x_2 \partial t} d\Omega, \tag{16}$$

where we have suppressed the subscript u on the test function. As we are considering hierarchical bases the two first vectors[5] of the basis can be taken as

$$P_0(z) = 1, \quad P_1(z) = z, \quad z \in [-1, 1].$$

Considering first the constant vectors $\phi_{0,0,1} = (1, 0)^T$ and $\phi_{0,0,2} = (0, 1)^T$ in the basis we see that the expression (16) vanishes. If we tried to use these equations the mass matrix would become singular. Therefore we replace the variational formulation, when tested against $\phi_{0,0,1}$ and $\phi_{0,0,2}$, with the independent equations

$$\int_\Omega (\frac{\partial u_i}{\partial t} - v_i) = 0, \quad i = 1, 2. \tag{17}$$

Next consider the linear vectors in the basis

$$\phi_{1,0,1} = [r, 0]^T, \quad \phi_{0,1,2} = [0, s]^T, \quad \phi_{0,1,1} = [s, 0]^T, \quad \phi_{1,0,2} = [0, r]^T.$$

Using the first two vectors (16) reduces to

$$\int_{\Omega_j} \begin{bmatrix} (2\mu+\lambda)r_1 & \lambda r_1 & \mu r_2 & \mu r_2 \\ \lambda s_2 & (2\mu+\lambda)s_2 & \mu s_1 & \mu s_1 \end{bmatrix} \begin{bmatrix} \frac{\partial^2 u_1}{\partial x_1 \partial t} \\ \frac{\partial^2 u_2}{\partial x_2 \partial t} \\ \frac{\partial^2 u_1}{\partial x_2 \partial t} \\ \frac{\partial^2 u_2}{\partial x_1 \partial t} \end{bmatrix} d\Omega.$$

Here we have used the chain rule to express $\partial_{x_1} = r_1\partial_r + s_1\partial_s$, etc. For the last two vectors (16) becomes

$$\int_{\Omega_j} \begin{bmatrix} (2\mu+\lambda)s_1 & \lambda s_1 & \mu s_2 & \mu s_2 \\ \lambda r_2 & (2\mu+\lambda)r_2 & \mu r_1 & \mu r_1 \end{bmatrix} \begin{bmatrix} \frac{\partial^2 u_1}{\partial x_1 \partial t} \\ \frac{\partial^2 u_2}{\partial x_2 \partial t} \\ \frac{\partial^2 u_1}{\partial x_2 \partial t} \\ \frac{\partial^2 u_2}{\partial x_1 \partial t} \end{bmatrix}.$$

[5]To avoid confusion between the finite element and the element of a basis we use the term vector for the elements in the basis in this section.

Now, assuming that $r_1 s_2 \neq 0$, we can use the two former "rows in the mass matrix" to reduce the latter two to

$$\int_{\Omega_j} \begin{bmatrix} 0 & 0 & \mu & \mu \\ 0 & 0 & \mu & \mu \end{bmatrix} \begin{bmatrix} \frac{\partial^2 u_1}{\partial x_1 \partial t} \\ \frac{\partial^2 u_2}{\partial x_2 \partial t} \\ \frac{\partial^2 u_1}{\partial x_2 \partial t} \\ \frac{\partial^2 u_2}{\partial x_1 \partial t} \end{bmatrix}.$$

These are obviously linearly dependent and we therefore replace the variational formulation for $\phi_{0,1,1} = [s, 0]^T$ with

$$\int_{\Omega_j} \left(\frac{\partial^2 u_1}{\partial x_2 \partial t} - \frac{\partial^2 u_2}{\partial x_1 \partial t} \right) - \left(\frac{\partial v_1}{\partial x_2} - \frac{\partial v_2}{\partial x_1} \right) = 0. \tag{18}$$

If $r_1 s_2 = 0$ we can instead replace the variational formulation for $\phi_{1,0,1} = [r, 0]^T$ with equation (18).

Practically speaking, in the assembly process used in our computer code we simply compare $m_{12} = \int_{\Omega_j} |r_1 s_2| d\Omega$ with $m_{21} = \int_{\Omega_j} |r_2 s_1| d\Omega$ and replace the variational formulation for $\phi_{0,1,1}$ with (18) if $m_{12} > m_{21}$. If $m_{12} < m_{21}$ we instead replace the variational formulation for $\phi_{1,0,1}$ with (18). We note that combining all equations containing $\frac{\partial u_i}{\partial t}$ and inverting the resulting mass matrix we obtain a simpler form which might prove more useful for general unstructured grids:

$$\frac{\partial u_i}{\partial t} - v_i|_{\Omega_j} = \text{boundary terms}. \tag{19}$$

To summarize, a full set of equations uniquely defining the time derivatives of our approximation are given by (14), (13) or (15) augmented by (17) and (18).

Remark 2. *At a first glance the equation relating velocity to the time derivative of the displacement, which in a single dimension can be written as*

$$M \hat{\mathbf{u}}_t = S \hat{\mathbf{v}} + \mathbf{L}_{\mathrm{R}} (\hat{\mathbf{v}}^*_{\mathrm{R}} - \hat{\mathbf{v}}_{\mathrm{R}}) - \mathbf{L}_{\mathrm{L}} (\hat{\mathbf{v}}^*_{\mathrm{L}} - \hat{\mathbf{v}}_{\mathrm{L}}), \tag{20}$$

appears to be less efficient than traditional formulations with diagonal mass matrices. Indeed if use the Legendre polynomials $P_0(z) = 1, P_1(z) = z, \ldots$, as our basis, the $N \times N$ mass matrix M (on the reference element) is a dense symmetric matrix with elements

$$M_{i,j} = \begin{cases} \frac{2}{3} & i = j = 1, \\ 3i(i-1) & i = 2, 4, 6 \ldots, \quad j = i, i+2, i+4, \ldots, \\ 3j(j-1) & j = 2, 4, 6 \ldots, \quad i = j, j+2, j+4, \ldots, \\ 0 & \textit{otherwise}. \end{cases} \tag{21}$$

However, since $M = S$ we always have that $M^{-1} S \hat{\mathbf{v}} = \hat{\mathbf{v}}$. Further, a direct computation shows that the symmetric pentadiagonal matrix A, with elements

$$A_{i,j} = \begin{cases} \frac{3}{2} & i = j = 1, \\ \frac{(2i-1)}{3((2i-1)^2 - 2^2)} & i = j = 2, 3, \ldots, N-2 \\ \frac{1}{6(2i-3)} & i = j = N-1, N, \\ -\frac{1}{6(2i+1)} & i = 2, 3, \ldots, j = i+2, \\ -\frac{1}{6(2j+1)} & j = 2, 3, \ldots, i = j+2, \\ 0 & \textit{otherwise}, \end{cases} \tag{22}$$

is the inverse of M. Now, as the elements of the vectors (lift operators) $\mathbf{L}_{\mathrm{R}}$ and $\mathbf{L}_{\mathrm{L}}$ are $(\mathbf{L}_{\mathrm{R}})_i = 3/2(i-1)i$ and $(\mathbf{L}_{\mathrm{L}})_i = (-1)^{i+1} (\mathbf{L}_{\mathrm{R}})_i$ the vectors $A\mathbf{L}_{\mathrm{R}}$ and $A\mathbf{L}_{\mathrm{L}}$ take the simple form

$$A\mathbf{L}_{\mathrm{R}} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1/2 \\ 1/2 \end{bmatrix}, \qquad A\mathbf{L}_{\mathrm{L}} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ -1/2 \\ 1/2 \end{bmatrix}. \tag{23}$$

Thus, the flux correction only affects the time evolution of the last two expansion coefficients and the computation of $\hat{\mathbf{u}}_t$ in simply an assignment plus four fused multiply adds independent of the order of approximation. Similarly in higher dimensions the cost of computing $\hat{\mathbf{u}}_t$ is dominated by the cost of computing $\hat{\mathbf{v}}_t$.

3.3. Numerical Fluxes

Let the discrete energy be defined as the sum over the energy contributions from each element

$$E^h(t) = \sum_j \int_{\Omega_j} \sum_{i=1}^{2} \frac{\rho}{2}(v_i^h)^2 + G^h(\nabla \vec{u}^h). \tag{24}$$

Computing the time derivative of the energy on a single element Ω_j we find that

$$\frac{dE_j^h(t)}{dt} = \int_{\Omega_j} \sum_{i=1}^{2} v_i^h \rho \frac{\partial v_i^h}{\partial t} + \sum_{k=1}^{2} \frac{\partial^2 u_i^h}{\partial t \partial x_k} \left(\frac{\partial G^h}{\partial u_{i,k}^h} \right). \tag{25}$$

This integral can be transformed into an integral over the boundary of the element by using the variational equations (14) and (15). Precisely, choosing $\vec{\phi}_u = \vec{u}^h$ and $\vec{\phi}_v = \vec{v}^h$ in (14) and (15) and adding those two equations together we find

$$\frac{dE_j^h(t)}{dt} = \int_{\partial\Omega_j} \sum_{i=1}^{2} v_i^h (\vec{w}_i^* \cdot \vec{n}) + \sum_{k=1}^{2} \left(\frac{\partial G^h}{\partial u_{i,k}} (\nabla \vec{u}^h)\, n_k \right) \left(v_i^* - v_i^h \right). \tag{26}$$

Focusing on the contribution from a face F_l between two elements we follow standard notation and label the current element by "$-$" and the adjacent element by "$+$". Further, to make room for this notation we suppress the superscript h. Then the contribution takes the form

$$\begin{aligned} &\int_{F_l} \sum_{i=1}^{2} \left[v_i^- (\vec{w}_i^* \cdot \vec{n}^-) + \sum_{k=1}^{2} \left(\frac{\partial G^-}{\partial u_{i,k}} (\nabla \vec{u}^-)\, n_k^- \right) \left(v_i^* - v_i^- \right) \right] \\ &\quad - \left[v_i^+ (\vec{w}_i^* \cdot \vec{n}^-) + \sum_{k=1}^{2} \left(\frac{\partial G^+}{\partial u_{i,k}} (\nabla \vec{u}^+)\, n_k^- \right) \left(v_i^* - v_i^+ \right) \right], \end{aligned} \tag{27}$$

where we have used that $\vec{n}^+ = -\vec{n}^-$.

We consider the parametrized flux

$$\vec{w}_i^* \cdot \vec{n}^- = \sum_{k=1}^{2} \left(n_k^- \left(\alpha_i \frac{\partial G^+}{\partial u_{i,k}} + (1-\alpha_i) \frac{\partial G^-}{\partial u_{i,k}} \right) \right) - \frac{\beta_i}{R}(v_i^- - v_i^+), \tag{28}$$

$$v_i^* = \left[(1-\alpha_i) v_i^+ + \alpha_i v_i^- \right] - \tau_i R \sum_{k=1}^{2} \left(n_k^- \left(\frac{\partial G^-}{\partial u_{i,k}} - \frac{\partial G^+}{\partial u_{i,k}} \right) \right), \tag{29}$$

which yields the energy contribution

$$-\int_{F_l} \sum_{i=1}^{2} \frac{\beta_i}{R}(v_i^- - v_i^+)^2 + \tau_i R \sum_{k=1}^{2} \left(n_k^- \left(\frac{\partial G^-}{\partial u_{i,k}} - \frac{\partial G^+}{\partial u_{i,k}} \right) \right)^2. \tag{30}$$

The method is thus dissipative when $\beta_i > 0, \tau_i > 0$ and conservative when $\beta_i = \tau_i = 0$. The choice of α_i is arbitrary; $\alpha_i = 1/2$ recovers the standard central flux and $\alpha_i = 0$ or $\alpha_i = 1$ recovers alternating fluxes. An upwind like flux is obtained by the choice $(\alpha_i, \beta_i, \tau_i) = (1/2, \beta_i > 0, \tau_i > 0)$.

We note that many authors use Riemann/Godunov fluxes [4, 6, 12, 13]. The forms above can only reproduce such fluxes in the case where element normals coincide with coordinate directions. However, Riemann/Godunov fluxes can be easily incorporated into our formulation by slightly generalizing (28)-(29) to replace $\mathrm{diag}(\beta_i)$, $\mathrm{diag}(\tau_i)$ by nondiagonal matrices. However in our experiments here we stick to the simpler forms.

3.4. Boundary Conditions

In the case of homogenous boundary conditions we follow [1] and set the outside states on a physical boundary as

$$v_i^+ = v_i^- - 2\gamma_i \left(\gamma_i v_i^- + \kappa_i R \left(\sum_{k=1}^{2} n_k^- \frac{\partial G^-}{\partial u_{i,k}} \right) \right),$$
$$\sum_{k=1}^{2} n_k^+ \frac{\partial G^+}{\partial u_{i,k}} = -\sum_{k=1}^{2} n_k^- \frac{\partial G^-}{\partial u_{i,k}} + 2\kappa_i \left(\gamma_i v_i^- + \kappa_i R \left(\sum_{k=1}^{2} n_k^- \frac{\partial G^-}{\partial u_{i,k}} \right) \right).$$

For the case of non homogenous Dirichlet boundary conditions, that is when $\gamma_i = 1, \kappa_i = 0$, and

$$\frac{\partial u_i}{\partial t} = d_i(x, y, t), \quad (x, y) \in \partial S, \quad i = 1, 2,$$

we simply set the outside state so that the average equals the boundary condition

$$v_i^+ = d_i(x, y, t) - v_i^-,$$
$$\sum_{k=1}^{2} n_k^+ \frac{\partial G^+}{\partial u_{i,k}} = -\sum_{k=1}^{2} n_k^- \frac{\partial G^-}{\partial u_{i,k}}.$$

If we want to impose a traction boundary condition, $\gamma_i = 0, \kappa_i = 1$, and

$$\vec{n} \cdot \vec{F}_i = T_i(x, y, t),$$

we set

$$v_i^+ = v_i^-,$$
$$\sum_{k=1}^{2} n_k^+ \frac{\partial G^+}{\partial u_{i,k}} = -T_i(x, y, t) + \sum_{k=1}^{2} n_k^- \frac{\partial G^-}{\partial u_{i,k}}.$$

4. Details of the Implementation

In this section we outline the implementation of the above method. The implementation is also available as open source software at [2].

To approximate the solution on an element Ω_j we first map the element to the reference element $\hat{\Omega} = \{(r, s) \in [-1, 1]^2\}$ and expand the solution components in a tensor product Chebyshev polynomial expansion. For example we approximate

$$u_1(x_1(r, s), x_2(r, s)) \approx u_1^h(x_1(r, s), x_2(r, s)) =$$
$$\sum_{l_1=0}^{q_u} \sum_{l_2=0}^{q_u} (\hat{u}_1^h(t))_{l_1,l_2} T_{l_1}(r) T_{l_2}(s), \quad x_1, x_2 \in \Omega_j.$$

Here $T_n(x) = \cos(n \arccos(x))$.

We assume that the elements are either straight-sided quadrilaterals or curved-sided quadrilaterals. Typically we use curved elements at the physical boundaries and straight-sided elements in the interior. We assume we know the curves bounding the elements so that we can use a simple transfinite interpolation to map the elements to the reference element. See Chapter 6 and Algorithm 98 in [22] for a comprehensive description of the transfinite interpolation mapping.

To compute the element and face integrals we use a Legendre-Gauss-Lobatto rule of sufficiently high degree. Typically we use $q_u + 6$ quadrature points. The metric coefficients required to express derivatives

with respect to the physical coordinates x_1, x_2 (via the chain rule) on the reference element are computed using Fornberg's finite difference formula algorithm `weights.f` at the quadrature points; see [23].

The resulting mass matrices are factored as a pre-computation using standard dense LAPACK routines. The application of the stiffness matrices to the solutions vectors are computed using BLAS routines and the element-wise linear systems of equations are solved using LAPACK as well.

Here the time evolution of the problem is done by the classic fourth order Runge-Kutta formulas but we note that it would also be possible to evolve the equations using a Taylor series method of arbitrary order.

4.1. Comparison to Existing Methods

It is difficult to compare discretization techniques as there are many different aspects that influence what is a "good method". If the final goal is that the method should be used by practitioners, it is important that it is stable and easy to understand and use. If on the other hand the goal is to maximize the efficiency then it is important that the method comes with a very high computation to communication ratio and minimal data transfer so that the method maximizes the utilization of current and future hardware.

As mentioned in the introduction, a wide variety of discontinuous Galerkin methods have been proposed for simulating elastic waves. Most focus on the dG discretization in space, with time handled either by a standard ODE solver, as in our case, or by the so-called ADER approach where a temporal Taylor series is approximated via the transformation of time to spatial derivatives [6]. An exception are the space-time dG methods [11, 12, 13]. Methods differ by using, as we do, displacements as basic variables [8, 9, 10, 12, 13], or transforming to first-order form with velocity-stress [5, 6, 7] or velocity-strain [4] as basic variables. They also differ by either conserving energy as in SIPG [9] and first order formulations with central fluxes [5] or staggered meshes [7], or dissipating energy using upwind fluxes [4, 6, 8, 12, 13]. Our method allows both energy-conserving and energy-dissipating fluxes, and we present experiments with both. However our upwind flux differs from the Riemann/Godunov fluxes used in the references. Essentially all methods, including ours, demonstrate optimal convergence of order $q+1$ if degree q polynomials are used to discretize the basic variables, at least for some numerical experiments.

Comparing with other methods which use displacement variables, a strength of the proposed method in comparision with interior penalty methods is that its stability does not require any grid dependent penalty terms. Although an IPDG method typically will be stable if the penalty parameters are chosen large enough, too large a penalty parameter will result in numerical stiffness and a consequent loss of efficiency. Comparisons with the space-time methods are more complex as there are obvious differences in the grid generation and time-stepping process. However it does seem that the weak formulations in [12] share some similarities with ours.

The upwind methods that work with first order forms of the equations are also stable and have been shown to be highly efficient when implemented on accelerators on both straight-sided and curved simplices [24, 25]. In particular the low-storage curvilinear discontinuous Galerkin [25] by Warburton eliminates the need to store a custom mass matrix on each curvilinear element by building element specific approximation spaces based on a scaling by the square root of the element Jacobian. A potential drawback of our method is that, due to the non-standard variational form we use, it may not be possible to construct such approximation spaces, necessitating the storage of a custom lift matrix. (Our current implementation uses a custom mass matrix, but as mentioned above this can be easily avoided.) This drawback can be partially offset by the fewer variables that need to be stored when working with the equations in second order rather than first order form.

5. Numerical Examples

In this section we assume that suitable scalings have been performed so that all quantities are dimensionless. Throughout we also set the dimensionality constant R to have a numerical value of unity.

5.1. Time-step Constraints and Scaling of the Upwind Flux for Large c_p/c_s

To experimentally investigate the time-step stability properties of the method we consider the unit square discretized by square elements of size $h \times h$. We enforce stress free boundary conditions (the results are similar for Dirichlet boundary conditions) and explicitly construct the matrix, M, which maps the degrees of freedom to their time derivatives. We restrict the presentation to the specific case $q_u = q_v$ and central or upwind fluxes. In addition the computational results we present in this example use $\mu = 1$, $\rho = 1$ but as the scaling below appears (based on observations from other numerical experiments not reported here) to be valid for general μ and ρ we keep them as parameters in the expressions.

We estimate, $\Lambda_{\max}$, the eigenvalue of M with largest amplitude, by computing the largest eigenvalue of M for a sequence of grids and then choosing the largest (scaled) value obtained in the sequence. In general, we find that the change of the largest eigenvalue for the different grids is small.

Starting with the central flux we aim to find a scaling, $S(q_u, \lambda)$, of $\Lambda_{\max}$ such that

$$\frac{h\Lambda_{\max}}{S(q_u, \lambda)} \approx \text{constant}. \tag{31}$$

In Figure 1 we display plots of $h\Lambda_{\max}$ versus $S(q_u, \lambda)$ and $\frac{h\Lambda_{\max}}{S(q_u,\lambda)}$ as a function of λ for $q_u = 2, \ldots, 9$ using the empirically determined scaling

$$S(q_u, \lambda) = \sqrt{\frac{2\mu + \lambda}{\rho}}(q_u + 3/2)^2. \tag{32}$$

We note that this is different than the typical scaling used for finite difference methods

$$S(q_u, \lambda) = \sqrt{\frac{3\mu + \lambda}{\rho}},$$

or for the dG method by Wilcox et al.

$$S(q, \lambda) = \sqrt{\frac{2\mu + \lambda}{\rho}} \max(1, q^2).$$

To leading order in q_u^2 and λ the scalings are the same but (empirically) we have found that (32) is a good model for our method.

Now for the upwind flux we have the freedom to choose the penalty parameters and the choice will have an impact on the allowable time-step for large λ (or equivalently large c_p/c_s). In Figure 2 we plot the full spectrum for a single grid keeping $q_u = 4$ fixed (results for other choices of q_u are similar) and varied $\lambda = 1, 2, 4, \ldots, 1024$. We also use either the central flux or the upwind flux with $\tau_i = 1, \sqrt{\rho}/\sqrt{2\mu + \lambda}, \rho/(2\mu + \lambda)$. As can be seen the last choice for τ_i gives a method whose spectrum resembles the central flux method and with a timestep that is limited by the eigenvalues along the imaginary axis. The first two choices for τ_i result in spectra with real negative eigenvalues that limit the timestep.

Choosing $\tau_i = \frac{\rho}{2\mu+\lambda}$ we repeat the scaling experiment. The results, displayed in Figure 3, are similar to those obtained with the central flux. Throughout the rest of the paper, when we refer to the upwind flux, this choice with $\tau_i = \frac{\rho}{2\mu+\lambda}$ is what we mean.

In conclusion we have that the timestep stability limit appears to be well described by

$$\Delta t < \frac{\text{CFL} \cdot h}{(q_u + \frac{3}{2})^2 \sqrt{\frac{2\mu+\lambda}{\rho}}}, \tag{33}$$

where h refers, for example, to the smallest diameter of the elements and CFL depends on the time stepper. We again note that the experiments are carried out for the specific case $\rho = \mu = 1$ but that the above expression, where we have left ρ,μ to be parameters, appears to hold for general ρ and μ. Roughly one should require $\text{CFL} \cdot h\Lambda_{\max}/S(q_u, \lambda)$ to be inside the stability region of the time stepper. For Runge-Kutta we typically choose CFL ~ 2 for meshes of reasonable quality.

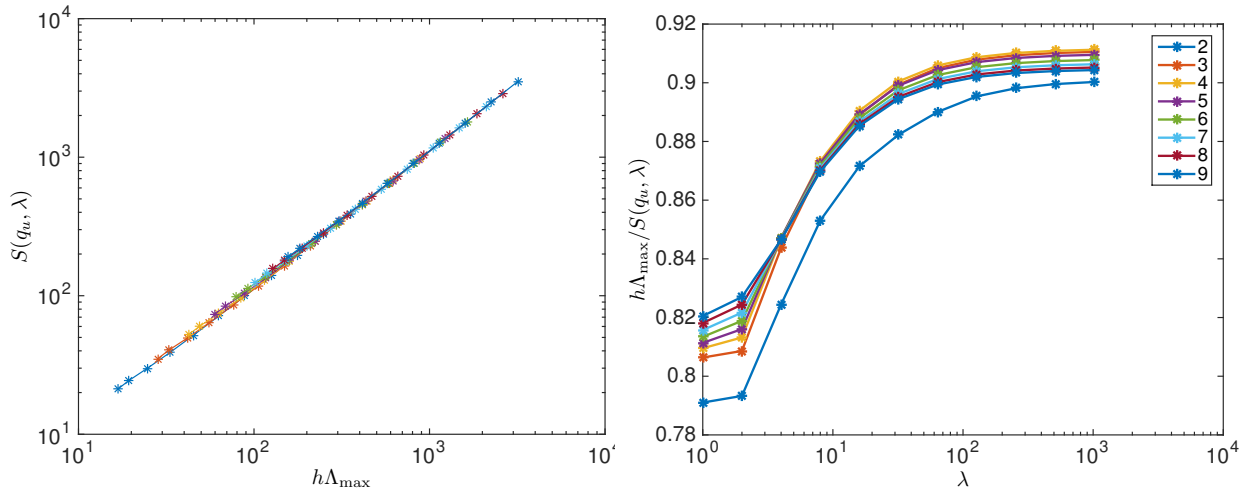

Figure 1: To the left we plot the quantity (32) as a function of the largest eigenvalue scaled by h for many different q_u and λ. As can be seen the curves for the different methods and material parameters collapse to a single curve with an almost constant slope. To the right we plot the slope of that curve as a function of λ. The results are for the central flux.

5.2. Choice of Approximation Spaces and Flux

In this section we consider the impact of the choice of the degrees of the approximations spaces for the displacement and velocity. We either choose $q_u = q_v + 1$ or $q_u = q_v$. We also investigate how the accuracy of the method is affected by the choice of flux.

5.2.1. Harmonic Vibrations of the Unit Square

We first consider the harmonic vibration of the unit square with traction free boundary conditions. Then a solution to the elastic wave equation is

$$u_1(x_1, x_2, t) = \cos(\omega t)\cos(\pi x_1)\sin(\pi x_2), \tag{34}$$
$$u_2(x_1, x_2, t) = -\cos(\omega t)\sin(\pi x_1)\cos(\pi x_2), \tag{35}$$

with $\omega = \sqrt{2\mu}$. Our mesh is a uniform discretization of the unit square split into smaller squares with side-length $h = 1/n_{\text{sq}}, n_{\text{sq}} = 2, 3, \ldots, 20$. We set $\mu = 1, \lambda = 1$ and simulate until time $t = \frac{\pi}{2\omega}$ when we measure the error. Here CFL is 0.1 to eliminate any temporal error effects. Note that we have reduced CFL for accuracy rather than stability. The results of the computations are presented in Figure 4.

The reported error, $e_{2,u}$, is the sum of the errors in the displacements made relative with the L_2-norm of the displacement at the initial time. Precisely we define

$$e_{2,u}(t) = \left(\frac{\sum_{i=1}^{2} \int_S \left(u_i - u_i^h\right)^2}{\sum_{i=1}^{2} \int_S \left(u_i(\cdot, \cdot, 0)\right)^2} \right)^{\frac{1}{2}}. \tag{36}$$

Measured rates of convergence are reported in Tables 1 and 2. Here we use least squares and the data points from the error graphs in Figure 4 to find the rates. The least squares fit is performed using errors from the four finest grids. Note that for the highest order methods we excluded the portions of the curves where the error saturated, i.e. for $h \lesssim 1/10$ for this problem.

The results displayed in the figures show relatively clean convergence with all the error curves bending slightly downwards as the errors approach the asymptotic regime. For the case $q_u = q_v + 1$ the results in Table 1 indicate that the rates of convergence are q_u for both the central and upwind flux. The main outlier is the case scheme C65 which has a higher rate. The left sub-figure also shows that the slopes are similar for

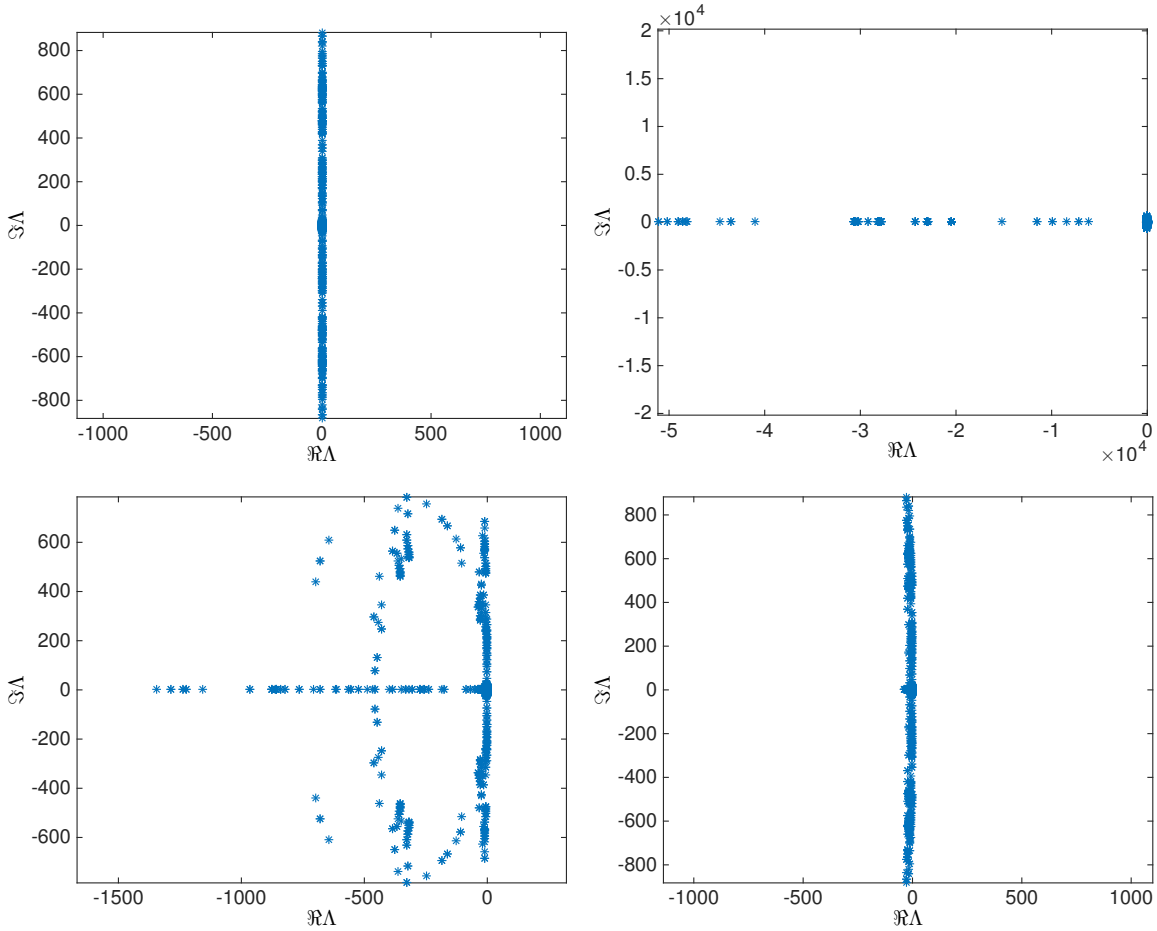

Figure 2: The upper left plot displays the spectrum of the matrix M when using the central flux. The upper right, lower left and lower right plots display the spectra of the matrix M when using and upwind flux with $\tau_i = 1, \sqrt{\rho}/\sqrt{2\mu+\lambda}, \rho/(2\mu+\lambda)$ respectively. Note that the range of the axis are different in the different plots. Also note that the last choice of τ_i yields a spectrum which is similar to the spectrum when the central flux is used.

both fluxes and that the error levels are almost identical when q_u is odd and that the upwind flux is more accurate when q_u is even.

If on the other hand the orders are chosen to be the same we find that the upwind methods with odd $q_u = q_v$ e.g. 3,5,7 appear to converge at a rate close to the optimal rate q_u+1 but only at a rate approximately q_u when $q_u = q_v$ is even; see Table 2. The schemes with a central flux on the other hand do not appear to benefit at all from the extra degree of freedom for the velocities. In fact, the rates are only q_u and $q_u - 1$ when $q_u = q_v$ is odd and even respectively. For all orders the upwind flux provides a more accurate answer.

Errors in the Velocity and Stress

For the upwind methods with odd $q_u = q_v = 3, 5, 7$ we also compute the errors in the velocity and in the stress components $\frac{\partial G}{\partial u_{1,1}}$ and $\frac{\partial G}{\partial u_{1,2}}$. The results, which can be found in Figure 5, indicate that the rates of convergence in these quantities are sub-optimal appearing to asymptotically approach q_u.

5.2.2. Forced Manufactured Method Solution on the Unit Square

Next we investigate the rates of convergence for non-homogenous problems. Again we consider the unit square $S = [0, 1]^2$ and impose Dirichlet conditions on the boundary. The initial data, boundary conditions

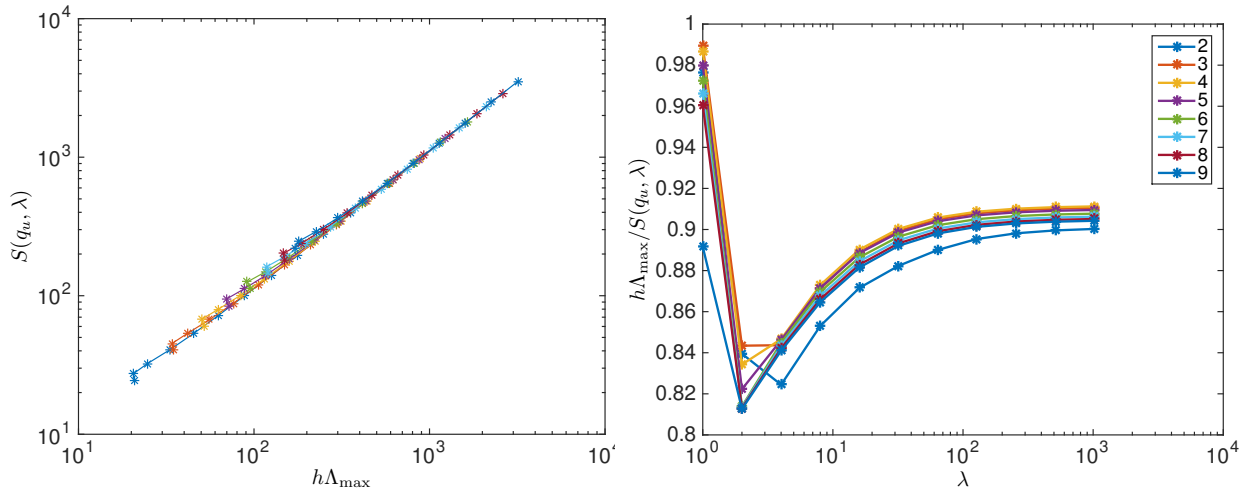

Figure 3: To the left we plot the largest eigenvalue scaled by h for the time-stepping matrix against the quantity (32) for different q_u and λ. As can be seen the curves for the different methods and material parameters collapse to a single curve with an almost constant slope. To the right we plot the slope of that curve as a function of λ. The results are for the upwind flux with $\tau_i = \frac{\rho}{2\mu+\lambda}$.

Table 1: Harmonic Vibrations of the Unit Square: Estimated rates of convergence for central and upwind fluxes using $q_u = q_v+1$.

q_u	2	3	4	5	6	7	8
C	2.23	3.05	4.29	4.99	7.12	6.81	7.46
U	2.14	2.93	4.28	4.98	6.20	6.76	8.66

Table 2: Harmonic Vibrations of the Unit Square. Estimated rates of convergence for central and upwind flux using $q_u = q_v$.

q_u	2	3	4	5	6	7	8
C	1.01	2.95	2.99	4.95	4.97	6.96	7.61
U	1.99	4.00	4.00	5.80	5.98	7.90	8.49

and forcing of the governing equations are chosen so that the solution is

$$u_1 = A_1 \cos(\omega t)\sin(k_x x + x_0)\sin(k_y y + y_0), \tag{37}$$
$$u_2 = A_2 \cos(\omega t)\sin(k_x x + x_0)\sin(k_y y + y_0), \tag{38}$$

with $A_1 = 1$, $A_2 = -1$, $k_x = 2.5\pi$, $k_y = 2\pi$, $x_0 = 5$, $y_0 = -10$. Our mesh is a uniform discretization of the unit square split into smaller squares with side-length $h = 1/n_{\text{sq}}, n_{\text{sq}} = 2, 3, \ldots, 20$.

In our first experiment we set $\omega = 2\pi$ and simulate until time $t = 1$ with CFL $= 0.1$. The material parameters are chosen to be constants, $\mu = 1, \lambda = 1$.

The errors plotted as a function of the grid size h are displayed in Figure 6 and the rates of convergence estimated using linear least squares are reported in Table 3 and 4. Again, the least squares fit is performed using errors from the the four finest grids.

From the figures we see that the errors for the upwind flux behave similarly to the problem above but that the central flux exhibits a longer pre-asymptotic regime where the errors behave in an unpredictable manner. The tables again display that the rates for $q_u = q_v = 3, 5, 7$ are higher than q_u but this time only $q_u = q_v = 3$ appears to reach the optimal rate $q_u + 1$. The methods $q_u = q_v = 5, 7$ appear to give a rate of

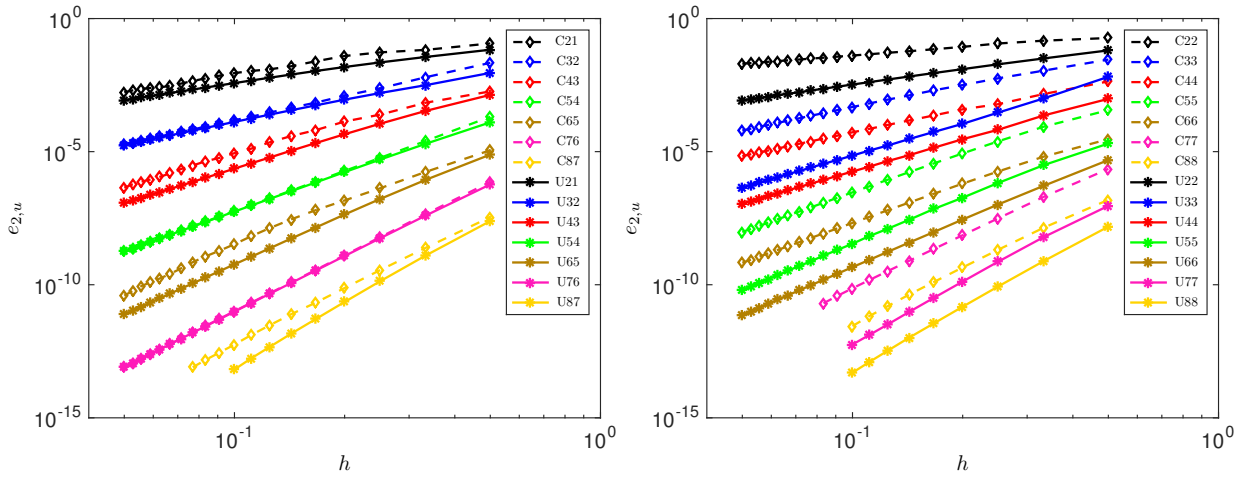

Figure 4: Harmonic Vibrations of the unit square. Errors at the final time for the harmonic vibrations of the traction free unit square. To the left $q_u = q_v + 1 = 2, \ldots, 8$ with central and upwind flux. To the right $q_u = q_v = 2, \ldots, 8$ with central and upwind flux.

Table 3: Estimated rates of convergence using central or upwind flux for the manufactured solution example. Here we are using $q_u = q_v + 1 = 2, \ldots, 8$.

q_u	2	3	4	5	6	7	8
C	2.02	3.57	4.84	6.93	6.42	8.27	5.92
U	2.69	4.09	4.35	5.27	5.99	7.51	8.27

Table 4: Estimated rates of convergence using central or upwind flux for the manufactured solution example. Here we are using $q_u = q_v = 2, \ldots, 8$.

q_u	2	3	4	5	6	7	8
C	1.20	2.82	3.44	4.02	4.96	6.16	7.29
U	2.45	4.07	4.49	5.44	6.11	7.52	8.67

convergence closer to $q_u + 1/2$. For this forced problem we also find that the upwind schemes with $q_u = q_v$ and $q_u = q_v + 1$ behave more similarly with respect to rates of convergence as well as with respect to actual error levels.

Although some methods appear to have optimal, or close to optimal rates the results are somewhat surprising as both experiments and analysis for the scalar wave equation show optimal convergence for all $q_u = q_v + 1$ and upwind fluxes [1].

5.2.3. A Comparison with a Summation by Parts Discretization

To see how our method compares to a high order finite difference scheme we repeat example 5.2.1 from [15]. The initial data, forcing and boundary conditions are chosen as in §5.2.2 with all the parameters the same except for λ, which is now taken to be 100. The computational domain is again the unit square $S = [0, 1]^2$ and we impose (non-homogenous) traction boundary conditions on all sides.

As in [15] we solve until time 2 on a sequence of four uniform grids refined by a factor of two between each level. We present results for the upwind flux and with $q_u = q_v = 3$ and 5 and with CFL=2.5. In [15] the coarsest grid has 25 elements in each direction so in order to match the number of degrees of freedom we use 4 elements for $q_u = 5$ and 6 elements for $q_u = 3$ (note that $4 \times (5 + 1) = 6 \times (3 + 1) = 24 \approx 25$.)

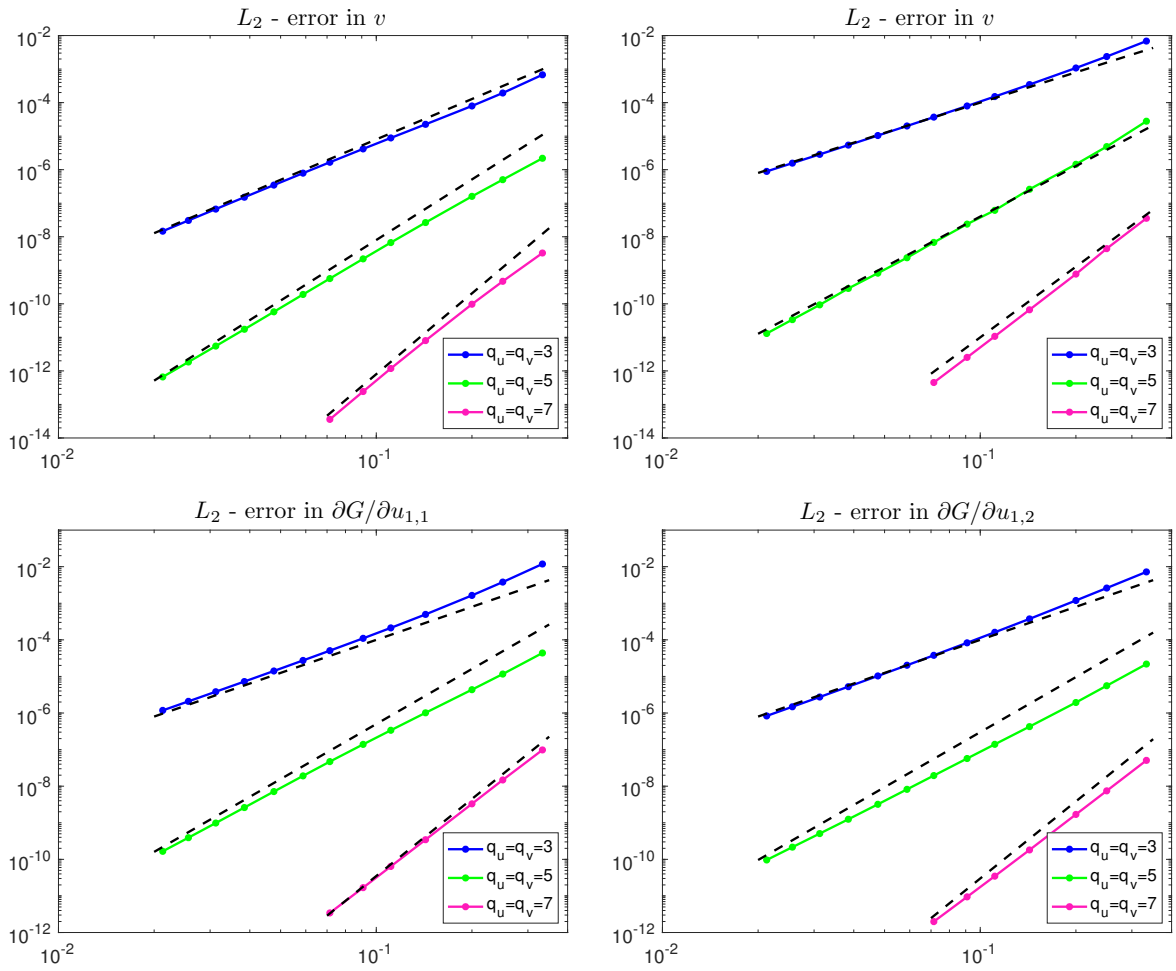

Figure 5: From the top right the sum of the L_2 errors in the displacements, the sum of the L_2 errors in the velocities, L_2 errors in $\frac{\partial G}{\partial u_{1,1}}$ and L_2 errors in $\frac{\partial G}{\partial u_{1,2}}$. The results are for $q_u = q_v = 3, 5, 7$ and the dashed line have slopes 4, 6 and 8 in the top right figure and 3, 5 and 7 in the other figures.

The errors as a function of time are displayed in Figure 7 for the four different grids and for the two choices for q_u. Comparing with Figure 1 in [15] we note that our method yields about one order of magnitude better results at fourth order. Our method with $q_u = 5$ is almost two orders of magnitude better than the "sixth order method" of [15], probably because their global order of accuracy is, in fact, only 5 (as is also noted in [15]). It is also worth noting that the error grows very slowly with time, significantly slower than the fourth order method of [15].

In Table 5 we present the errors at the final time and computed rates of convergence. Note that we now observe the optimal rate of convergence $q_u + 1$. This is surprising as the only difference between this experiment is that we now impose the traction and not the velocity on the boundary. Comparing the results in the table with Table 1 in [15] we see again that for the same number of degrees of freedom we observe smaller errors than for the SBP-SAT discretization of [15].

Above we compared the results based on the number of degrees of freedom and it should be pointed out that our methods take ~ 1.6 and 2.2 times smaller time steps compared to the methods in example 5.2.1 in [15].

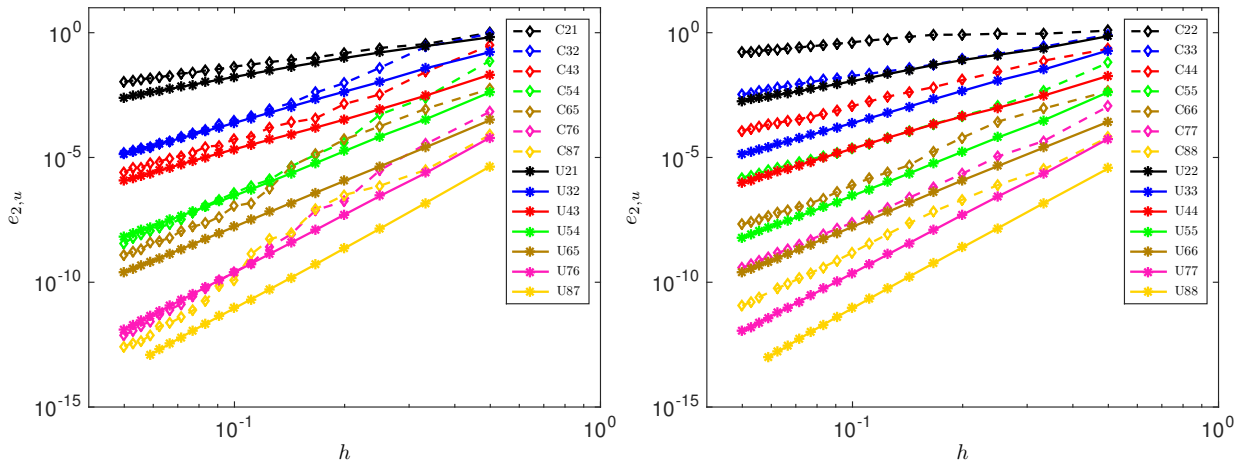

Figure 6: Manufactured Solution: To the left $q_u = q_v+1 = 2, \ldots, 8$ with central and upwind flux. To the right $q_u = q_v = 2, \ldots, 8$ with central and upwind flux.

Table 5: Example from [15]: Errors at time 2 for different grids and orders of approximation

h	error $q_u = 3$	rate	h	error $q_u = 5$	rate
1/6	3.55(-3)		1/4	8.92(-4)	
1/12	1.77(-4)	4.32	1/8	8.12(-6)	6.78
1/24	1.09(-5)	4.02	1/16	1.11(-7)	6.19
1/48	6.79(-7)	4.01	1/32	1.78(-9)	5.97

5.2.4. Harmonic Vibrations of an Annulus

This example is inspired by example 6.4 in [26]. We consider the harmonic oscillations of the radially symmetric mode

$$u_r(r) = \cos(\omega_k t)\left[J_1(\alpha_k t) + B_k Y_1(\alpha_k r)\right], \tag{39}$$

in an annulus with inner radius $r_i = 1/2$ and outer radius $r_o = 1$. The constants α_k, B_k are determined by the boundary conditions, which we take to be homogenous Dirichlet boundary condition on the inner boundary and homogenous traction on the outer boundary. We consider two cases $\lambda = 100$, $k = 1$ and $\lambda = 1$, $k = 5$. These correspond to constants $\alpha_1 = 3.578028273880645$ and $B_1 = 2.524281875355463$ for the first case and $\alpha_5 = 28.28892419267416$, $B_5 = 1.069885069032360$ for the second case.

We solve on a sequence of structured grids defined by the grid points

$$\begin{aligned} x_{kl} &= r_i \cos(k\frac{2\pi}{n_\Theta}) + \frac{l}{n_r}(r_o - r_i)\cos(k\frac{2\pi}{n_\Theta}), \\ y_{kl} &= r_i \sin(k\frac{2\pi}{n_\Theta}) + \frac{l}{n_r}(r_o - r_i)\sin(k\frac{2\pi}{n_\Theta}), \\ & k = 0, \ldots, n_\Theta, \quad l = 0, \ldots, n_r, n_r = n_\Theta/4. \end{aligned}$$

To make the element next to the boundary conform with it we use a transfinite interpolation and over-integrate using tensor product LGL quadrature. The metric terms at the quadrature points are computed using a high order finite difference approximation; see [23] (see also the description in §4). In this experiment we take CFL to be equal to 1 and choose $n_\Theta = 32, 36, \ldots, 60$ when $\lambda = 1$ and $n_\Theta = 20, 24, \ldots, 40$ when $\lambda = 100$.

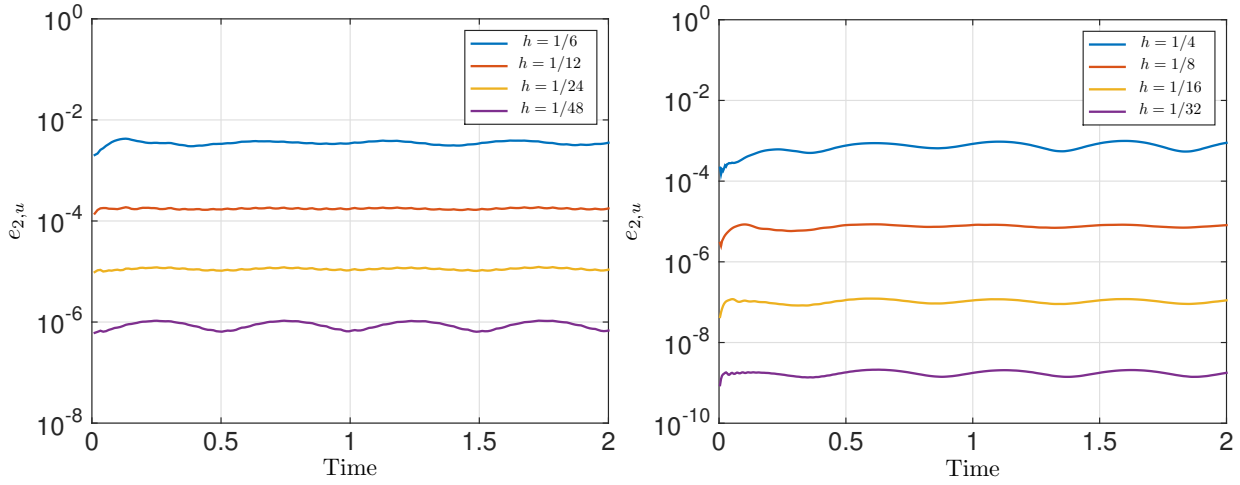

Figure 7: Errors as a function of time for the example from [15]. To the left $q_u = 3$, to the right $q_u = 5$. Note that the axis are adjusted to coincide with Figure 1 in [15].

The solution is a pure pressure wave so $\alpha_k^2 = \rho\omega_k^2/(2\mu + \lambda)$. To measure the error in the solution we solve until time $t = \frac{5\pi}{2\omega_k}$ (2.5 periods) at which time the exact solution is identically zero. The errors are reported as in the previous examples, i.e. as L_2 errors made relative with the solution at the initial time, see equation (36).

Table 6: Harmonic oscillations of an annulus with $\lambda = 1$, $\mu = 1$, $\rho = 1$. Estimated rates of convergence for central and upwind flux using $q_u = q_v$ and $q_u = q_v + 1$. Here h is taken to be the square root of the area of the element with the smallest area.

$q_u = q_v$	2	3	4	5	6	7
C	2.34	3.00	3.62	5.37	8.61	6.79
U	2.91	4.44	4.45	6.70	7.00	8.62
$q_u = q_v + 1$	2	3	4	5	6	7
C	1.96	6.32	4.83	7.21	6.69	9.86
U	2.86	4.30	3.97	6.42	6.65	8.55

The results, displayed in Figure 8 and Table 6 and 7, again give somewhat varying estimates for the rates of convergence (the rates of convergence are obtained from a least squares fit of the four last data-points). Focusing on the figure it is however clear that the upwind flux gives much steadier rates of convergence. It also appears as the rates are better for higher order methods and that the upwind flux gives rates of convergence greater than q_u for almost all cases.

Based on these and the above results we conclude that the upwind flux with $q_u = q_v$ and odd results in methods with robust and observed optimal or close to optimal convergence rates. To this end we will present results below for an upwind flux and $q_u = q_v$.

5.3. Exact Solutions

In this section we perform due diligence and demonstrate the performance of our method on known analytic and semi-analytic solutions to the elastic wave equation in simple geometries.

5.3.1. Rayleigh Wave

Perhaps the most well known wave supported by the elastic wave equation is the Rayleigh surface wave. This wave travels along a traction-free surface and decays exponentially in the normal direction from the

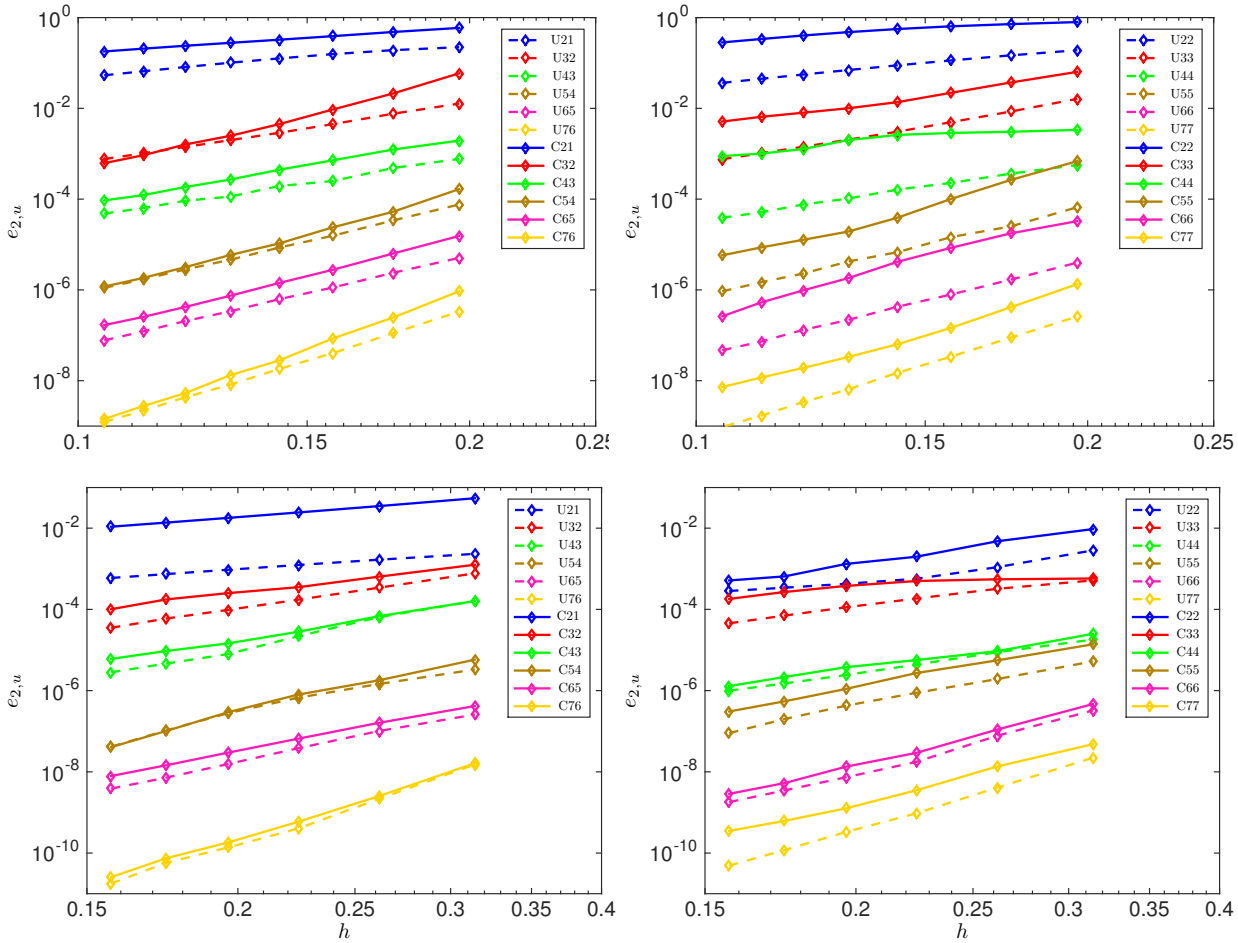

Figure 8: Harmonic oscillations of an annulus with $\lambda = 1$, $\mu = 1$, $\rho = 1$ (top two figures) and $\lambda = 100$, $\mu = 1$, $\rho = 1$ (bottom two). To the left $q_u = q_v + 1 = 1, \ldots, 7$ with central and upwind flux. To the right $q_u = q_v = 1, \ldots, 7$ with central and upwind flux. The notation in the legend refers to the flux used and the degree of the elements, e.g. U21 refers to upwind flux with $q_u = 2, q_v = 1$.

surface. The decay of the wave increases as the ratio $\frac{\lambda}{\mu}$ increases and this has been observed to lead to difficulties for some numerical approximations; see [14]. Following the example in Section 5.4 of [14] we consider an isentropic solid with density $\rho = 1$ in $x_1 > 0$ and with a traction free surface at $x_1 = 0$. Then the Rayleigh wave can be written

$$\begin{aligned}\vec{u} &= e^{-\omega x_1\sqrt{1-\xi^2}}\begin{pmatrix}\cos(\omega(x_2+c_r t))\\ \sqrt{1-\xi^2}\sin(\omega(x_2+c_r t))\end{pmatrix}\\ &\quad+\left(\frac{\xi^2}{2}-1\right)e^{-\omega x_1\sqrt{1-\frac{\xi^2\mu}{2\mu+\lambda}}}\begin{pmatrix}\cos(\omega(x_2+c_r t))\\ \sin(\omega(x_2+c_r t)/\sqrt{1-\frac{\xi^2\mu}{2\mu+\lambda}})\end{pmatrix}.\end{aligned}$$

Here $c_r = \xi\sqrt{\mu}$ is the Rayleigh phase velocity and ξ is a solution of the dispersion equation

$$\sqrt{1-\xi^2}\sqrt{1-\frac{\xi^2\mu}{2\mu+\lambda}}-\left(\frac{\xi^2}{2}-1\right)^2=0. \tag{40}$$

As in [14] we fix $\lambda = 1$ and vary μ. We compute solutions in the rectangular domain $(x_1, x_2) \in$

Table 7: Harmonic oscillations of an annulus with $\lambda = 100$, $\mu = 1$, $\rho = 1$. Estimated rates of convergence for central and upwind flux using $q_u = q_v$ and $q_u = q_v + 1$.

$q_u = q_v$	2	3	4	5	6	7
C	4.00	2.84	4.17	6.11	6.64	6.44
U	1.94	3.94	4.15	6.38	6.39	8.33
$q_u = q_v + 1$	2	3	4	5	6	7
C	2.26	3.42	4.26	8.40	5.95	8.66
U	2.07	4.44	5.65	7.82	6.44	8.59

$[0, 10] \times [0, 1]$ and choose $\omega = 2\pi$ so that the domain contains exactly one wavelength in the direction along the surface. The computational domain is discretized into square elements of size $1/N$. To compare to the results in [14] we introduce the degrees of freedom per wavelength as

$$P_r = N(q_u + 1). \tag{41}$$

We impose free surface conditions at the surface $x_1 = 0$ and use the exact solution as forcing to set Dirichlet conditions (for the velocity) on all other boundaries.

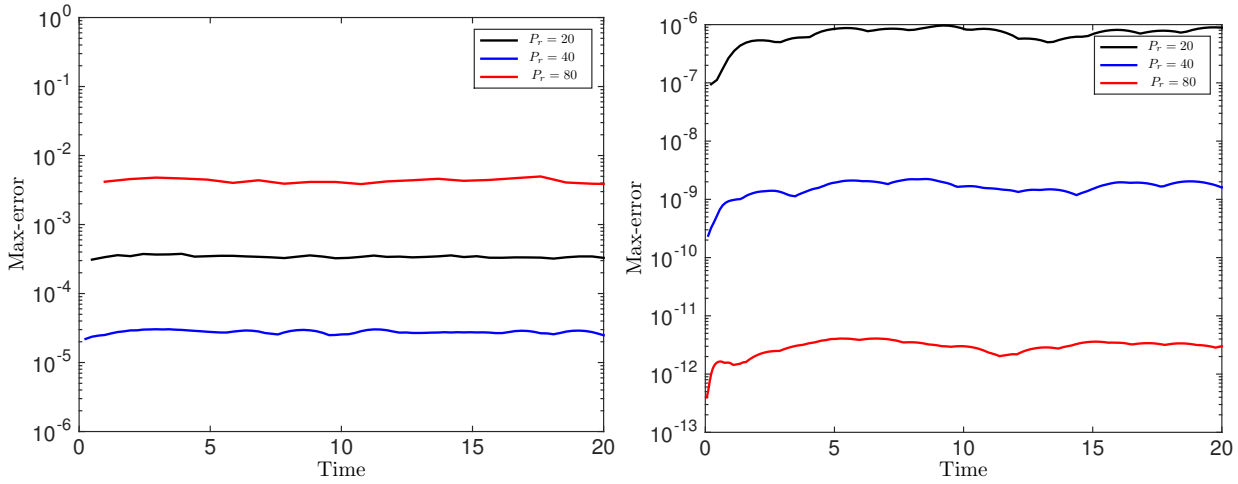

Figure 9: Rayleigh Wave: Max errors as a function of time when $\lambda = 1$, $\mu = 0.01$ for 20, 40, and 80 degrees of freedom per wavelength. To the left $q_u = 3$. To the right $q_u = 9$.

Following [14] we first take $\mu = 0.01$ and solve until time 20. We use the upwind flux, $q_u = q_v$ and CFL $= 0.5$ for $q_u = 9$ and CFL $= 1.0$ for $q_u = 3$. In Figure 9 we present results displaying the maximum error as a function of time. To the left in Figure 9 the results with $q_u = 3$ are displayed and to the right the results with $q_u = 9$. The number of elements in the x_2-direction, 5, 10 and 20 and 2, 4 and 8 respectively, were chosen to match the number of points per wavelength used in [14]. Comparing the results to Figure 4 in [14] we note that unlike [14] we do not see any significant growth in time of the error. Also, the error levels are a bit lower for $q_u = 3$ compared to [14] while the results with $q_u = 9$ are significantly better.

Kreiss and Petersson [18] studied the Rayleigh wave in the incompressible limit $\mu/\lambda \to 0$ using normal modes and modified equation techniques. The main result is that the classical Kreiss-Oliger points per wavelength theory for finite difference methods [17] is inadequate for predicting the number of points per wavelength required when propagating surface waves. Kreiss and Petersson found that to meet a predetermined tolerance in the error of the Rayleigh phase velocity the number of points per wavelength for

a pth order finite difference method must scale as $\sim (\lambda/\mu)^{1/p}$. These results have also been confirmed by finite difference simulations in [14, 15].

We repeat the second experiment in [14] where the influence of the ratio of λ/μ is studied. The problem setup is the same as above with the exception that μ is set to $1, 0.1, 0.01, 0.001, 0.0001$ (again $\lambda = 1$ is fixed) and that the final time is one period (which depends on μ).

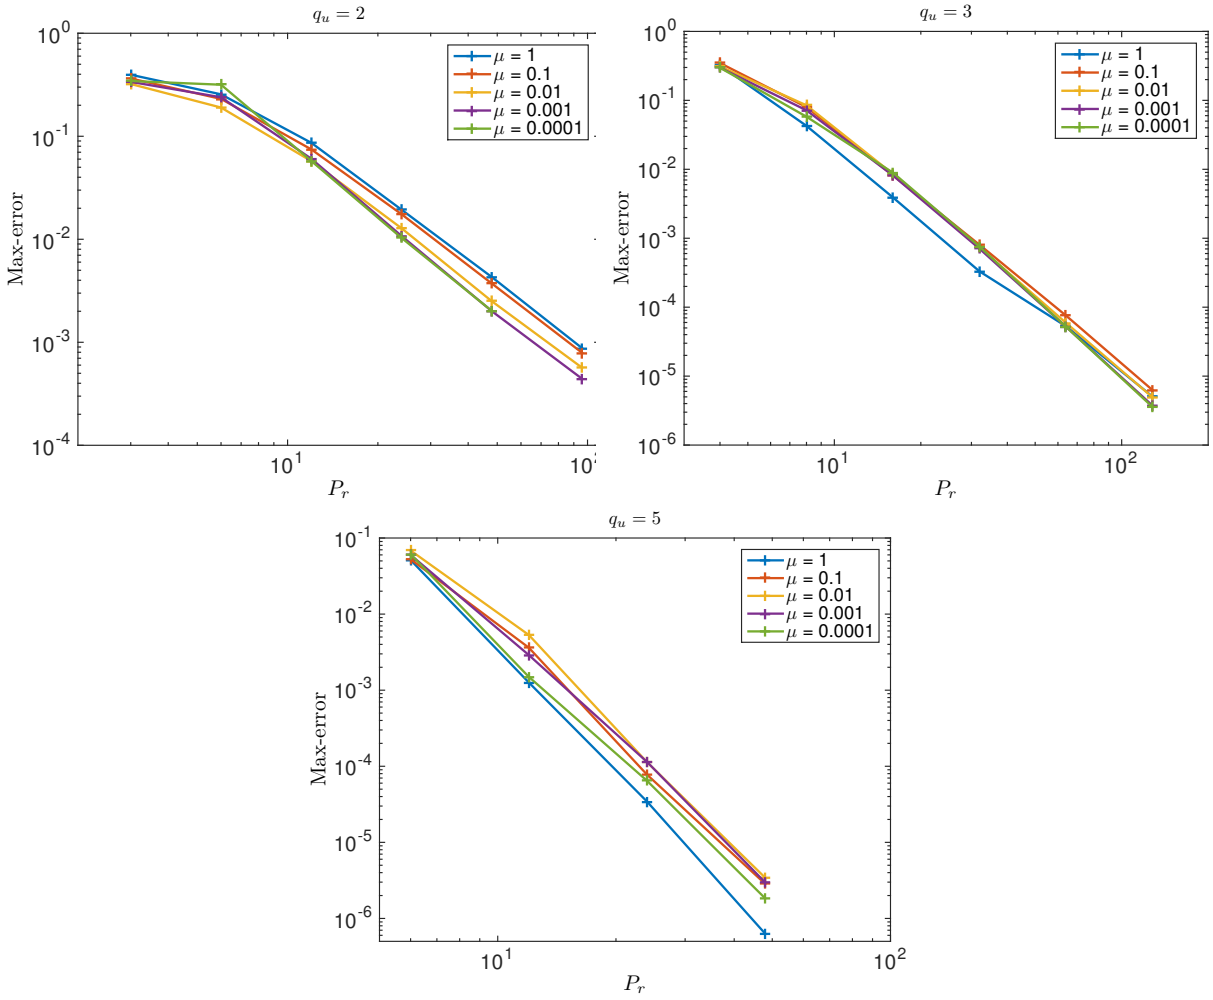

Figure 10: Rayleigh Wave: Max-error as a function of number of degrees of freedom, P_r, for different values of μ and different $q_u = q_v$. Unlike the results reported in [14, 15] we do not experience any significant degratation in the performance of the method when $\lambda/\mu \to \infty$.

We again perform simulations with the upwind flux and with $q_u = q_v = 2, 3, 5$. The constant CFL is taken to be 0.5. The results are displayed in Figure 10. As in [14] we report the maximum error at the final time as a function of number of degrees of freedom, P_r, for different values of μ. As can be seen in the figure our method does not appear to suffer from any significant degradation in performance as $\lambda/\mu \to \infty$.

We have not yet carried out a theoretical dispersion analysis of our method but we suspect that this excellent propagation of the surface waves can be explained by the superconvergence of the dispersion errors. Precisely we suspect, as for other dG methods (see for example [19]), that the dispersion errors converge at a rate about twice the degree of the polynomial approximation space.

5.3.2. Lamb Wave

As another test of the free surface conditions we consider a rectangular waveguide occupying $(x_1, x_2) \in [-\infty, \infty] \times [-d/2, d/2]$ with free surface conditions at $x_2 = \pm d/2$. Such a waveguide supports Lamb waves. Following Wilcox et al. [4] we consider the symmetric Lamb wave

$$u_1(x_1, x_2, t) = (-kB_1 \cos(px_2) - qB_2 \cos(qx_2)) \sin(kx_1 - \omega t),$$
$$u_2(x_1, x_2, t) = (-pB_1 \sin(px_2) + kB_2 \sin(qx_2)) \cos(kx_1 - \omega t),$$

where the parameters p and q are defined as

$$p^2 = \frac{\omega^2}{2\mu + \lambda} - k^2, \quad q^2 = \frac{\omega^2}{\mu} - k^2,$$

and the wavenumber k and frequency ω are related through the dispersion relation

$$\frac{\tan(qd)}{\tan(pd)} + \frac{4k^2pq}{(q^2 - k^2)^2} = 0.$$

The amplitudes B_1 and B_2 are unique up to a scaling constant and their ratio satisfies:

$$\frac{B_1}{B_2} = \frac{2\mu kq \cos(qd)}{(\lambda k^2 + (2\mu + \lambda)p^2) \cos(pd)}.$$

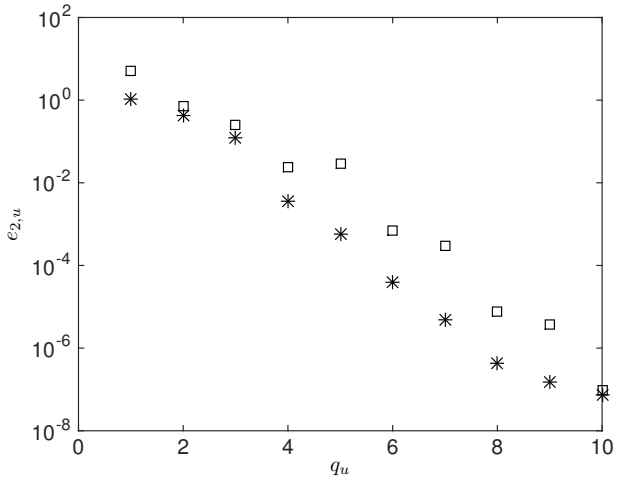

Figure 11: The Lamb wave problem. Squares represent central flux and stars the upwind flux. Both methods use $q_u = q_v$.

As in [4] we consider the computational domain $(x_1, x_2) \in [-1, 1] \times [-1/2, 1/2]$, i.e. $d = 1/2$, and a wavenumber $k = 2\pi$. We set the material parameters $(\lambda, \mu, \rho) = (2, 1, 1)$, which leads to $\omega = 13.137063197233$. The amplitudes are set to $B_1 = 126.1992721468$, $B_2 = 53.88807700007$.

We discretize the domain with a uniform grid with 8×4 square elements, set CFL $= 2.5$, and impose Dirichlet conditions on the left and right ends of the domain $x_1 = \pm 1$. We solve until time $t = 1$ with $q_u = q_v$ and either the upwind or central flux. The results, presented in Figure 11, display the spectral accuracy of the methods. Note that as above the upwind flux gives a monotone and robust convergence while the central flux appears to converge in steps of 2-increments with q_u.

5.3.3. Mode Conversion at an Interface

Next we consider an example with piecewise constant material properties consisting of the mode conversion of an incoming P-wave impinging on the interface $x_2 = 0$. Let the interface x_2 separate two domains with material properties $(\lambda_a, \mu_a, \rho_a)$ and $(\lambda_b, \mu_b, \rho_b)$. Then the angles of the incoming (plane) P-wave, α_{ip}, the reflected P and S-wave, α_{rp}, α_{rs}, and the transmitted P and S-wave, α_{tp}, α_{ts} (defined as in Figure 12) obey Snell's law:

$$\frac{\sin\alpha_{ip}}{c_{ap}} = \frac{\sin\alpha_{rp}}{c_{ap}} = \frac{\sin\alpha_{rs}}{c_{as}} = \frac{\sin\alpha_{tp}}{c_{bp}} = \frac{\sin\alpha_{ts}}{c_{bs}},$$

Here $c_{ap}^2 = (2\mu_a + \lambda_a)/\rho_a$, $c_{bp}^2 = (2\mu_b + \lambda_b)/\rho_b$, $c_{as}^2 = \mu_a/\rho_a$, $c_{bs}^2 = \mu_b/\rho_b$.

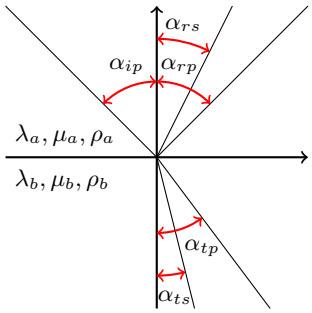

Figure 12: Snell's law for an incoming P-wave.

We may separate the solutions in the upper and lower part of the domain as

$$\vec{u}_a = \vec{u}_{ip} + \vec{u}_{rp} + \vec{u}_{rs},$$
$$\vec{u}_b = \vec{u}_{tp} + \vec{u}_{ts},$$

where

$$\begin{aligned}
\vec{u}_{ip} &= A_{ip} \begin{bmatrix} \sin(\alpha_{ip}) \\ -\cos(\alpha_{ip}) \end{bmatrix} \cos\left(k_{ap}\left[x_1\sin(\alpha_{ip}) - x_2\cos(\alpha_{ip})\right] - \omega t\right), \\
\vec{u}_{rp} &= A_{rp} \begin{bmatrix} \sin(\alpha_{rp}) \\ \cos(\alpha_{rp}) \end{bmatrix} \cos\left(k_{ap}\left[x_1\sin(\alpha_{rp}) + x_2\cos(\alpha_{rp})\right] - \omega t\right), \\
\vec{u}_{rs} &= A_{rs} \begin{bmatrix} \cos(\alpha_{rs}) \\ -\sin(\alpha_{rs}) \end{bmatrix} \cos\left(k_{as}\left[x_1\sin(\alpha_{rs}) + x_2\cos(\alpha_{rs})\right] - \omega t\right), \\
\vec{u}_{tp} &= A_{tp} \begin{bmatrix} \sin(\alpha_{tp}) \\ -\cos(\alpha_{tp}) \end{bmatrix} \cos\left(k_{bp}\left[x_1\sin(\alpha_{tp}) - x_2\cos(\alpha_{tp})\right] - \omega t\right), \\
\vec{u}_{ts} &= A_{ts} \begin{bmatrix} \cos(\alpha_{ts}) \\ \sin(\alpha_{ts}) \end{bmatrix} \cos\left(k_{bs}\left[x_1\sin(\alpha_{ts}) - x_2\cos(\alpha_{ts})\right] - \omega t\right).
\end{aligned}$$

The wave numbers $k_{ap}, k_{as}, k_{bp}, k_{bs}$ are given by

$$k_{ap} = \frac{\omega}{c_{ap}}, \quad k_{as} = \frac{\omega}{c_{as}}, \quad k_{bp} = \frac{\omega}{c_{bp}}, \quad k_{bs} = \frac{\omega}{c_{bs}}.$$

Now, if we prescribe the material parameters, the incoming angle, and the amplitude and angular frequency of the incoming P-wave, we can obtain the amplitudes of the reflected and transmitted waves by requiring that the total wave fields in domains a and b satisfy the interface conditions:

$$\vec{u}_a = \vec{u}_b, \quad x_2 = 0, \tag{42}$$

$$(2\mu_a + \lambda_a)\frac{\partial u_{2a}}{\partial x_2} + \lambda_a \frac{\partial u_{1a}}{\partial x_1} = (2\mu_b + \lambda_b)\frac{\partial u_{2b}}{\partial x_2} + \lambda_b \frac{\partial u_{1b}}{\partial x_1}, \quad x_2 = 0, \tag{43}$$

$$\mu_a \left(\frac{\partial u_{1a}}{\partial x_2} + \frac{\partial u_{2a}}{\partial x_1} \right) = \mu_b \left(\frac{\partial u_{1b}}{\partial x_2} + \frac{\partial u_{2b}}{\partial x_1} \right), \quad x_2 = 0. \tag{44}$$

To demonstrate the performance of our method for problems with discontinuous media we consider the problem defined by the parameters: $\alpha_{ip} = \pi/3$, $\omega = 2\pi$, $A_{ip} = 1$, $(\lambda_a, \mu_a, \rho_a) = (3, 2, 1)$ and $(\lambda_b, \mu_b, \rho_b) = (1.1.1)$.

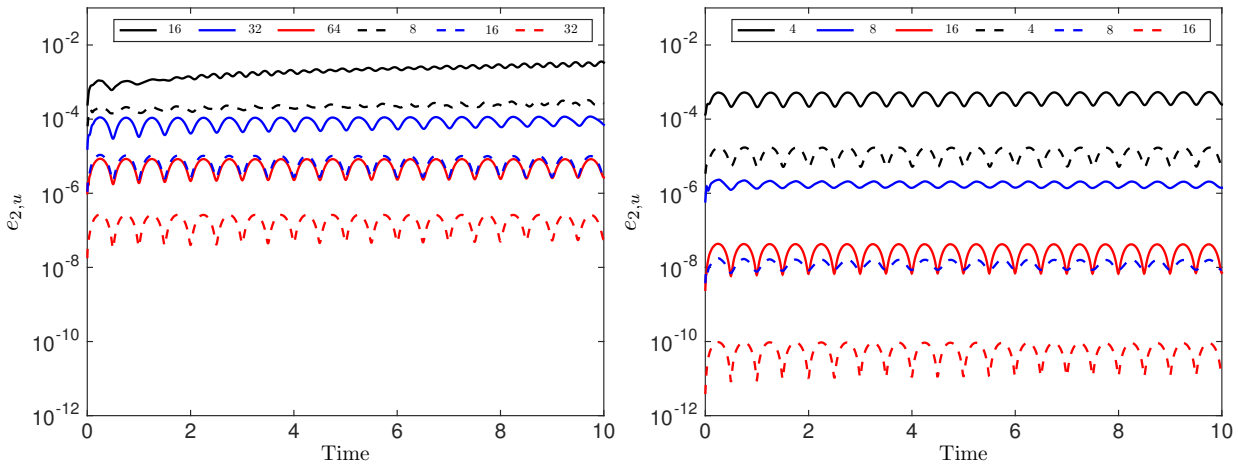

Figure 13: Errors as a function of time for the mode conversion example. To the left the results with $q_u = 3$ correspond to solid lines and $q_u = 5$ to dashed lines. To the right the results with $q_u = 7$ correspond to solid lines and $q_u = 9$ to dashed lines. The labels refer to the number of elements used in the x_1 direction.

We solve on the domain $(x_1, x_2) \in [-1, 1] \times [-2, 2]$ using the upwind flux and with $q_u = q_v = 3, 5, 7, 9$. We set CFL $= 0.5$ and record the error (36) as a function of time. The mesh consists of squares of size $2/N$.

In Figure 13 we display the errors as a function of time for the different choices of q_u and for different N. We note that the high order methods give very accurate results even for a very coarse discretization.

In Table 8 we report the errors at the final time and rates of convergence for the different choices of q_u and the different grids. The rates of convergence appear to be consistently greater than q_u.

Table 8: Errors and convergence rates for the mode conversion example.

$q_u = 3, N$	$e_{2,u}(t = 10)$	rate	$q_u = 5, N$	$e_{2,u}(t = 10)$	rate
16	3.35e-03	-	8	2.63e-04	-
32	7.02e-05	5.58	16	2.84e-06	6.54
64	2.66e-06	4.72	32	5.57e-08	5.67
$q_u = 7, N$	$e_{2,u}(t = 10)$	rate	$q_u = 9, N$	$e_{2,u}(t = 10)$	rate
4	2.48e-04	-	4	5.30e-06	-
8	1.41e-06	7.46	8	8.64e-09	9.26
16	7.12e-09	7.63	16	1.25e-11	9.43

5.3.4. Stoneley Wave

The Stoneley wave, [27], exists at an interface between two solids and is similar to the Rayleigh wave in that it decays exponentially away from the interface. To demonstrate the spectral accuracy (with increasing

q_u) of our solver we evolve the Stonely wave used in the example from Section 6.6 in [4] at an interface $x_2 = 0$. The material parameters for $x_2 > 0$ are $(\lambda_1, \mu_1, \rho_1) = (3, 3, 10)$ and for $x_2 < 0$ they are $(\lambda_1, \mu_1, \rho_1) = (1, 1, 1)$. With this setup the Stoneley wave takes the form

$$
\begin{aligned}
u_1(x_1, x_2, t) &= \Re\left[\left(ikB_1 e^{-kb_{1p}x_2} + kb_{1s}B_2 e^{-kb_{1s}x_2}\right) e^{i(kx_1-\omega t)}\right], \quad x_2 > 0,\\
u_2(x_1, x_2, t) &= \Re\left[\left(-kb_{1p}B_1 e^{-kb_{1p}x_2} + ikB_2 e^{-kb_{1s}x_2}\right) e^{i(kx_1-\omega t)}\right], \quad x_2 > 0,\\
u_1(x_1, x_2, t) &= \Re\left[\left(ikB_3 e^{kb_{2p}x_2} - kb_{2s}B_4 e^{kb_{2s}x_2}\right) e^{i(kx_1-\omega t)}\right], \quad x_2 < 0,\\
u_2(x_1, x_2, t) &= \Re\left[\left(kb_{2p}B_3 e^{kb_{2p}x_2} + ikB_4 e^{kb_{2s}x_2}\right) e^{i(kx_1-\omega t)}\right], \quad x_2 < 0.
\end{aligned}
$$

Here

$$
k = \frac{\omega}{c_{\text{st}}}, \quad b_{lp} = \sqrt{1 - \frac{c_{\text{st}}^2}{(2\mu_l + \lambda_l)/\rho_l}}, \quad b_{ls} = \sqrt{1 - \frac{c_{\text{st}}^2}{\mu_l/\rho_l}}, \; l = 1, 2,
$$

and c_{st} is the Stoneley wave speed. It is determined by an algebraic solvability condition derived from the interface conditions (see [4, 27]) and depending on the material parameters. For the material parameters above $c_{\text{st}} = 0.546981324213884$. Once c_{st} has been determined the amplitudes B_1, B_2, B_3, B_4 follow (up to a scaling with a constant.) As in [4] we set them to be $B_1 = -i0.2952173626624$, $B_2 = -0.6798795208473$, $B_3 = i0.5220044931212$, and $B_4 = -0.9339639688697$. Finally, we set $\omega = c_{\text{st}}$.

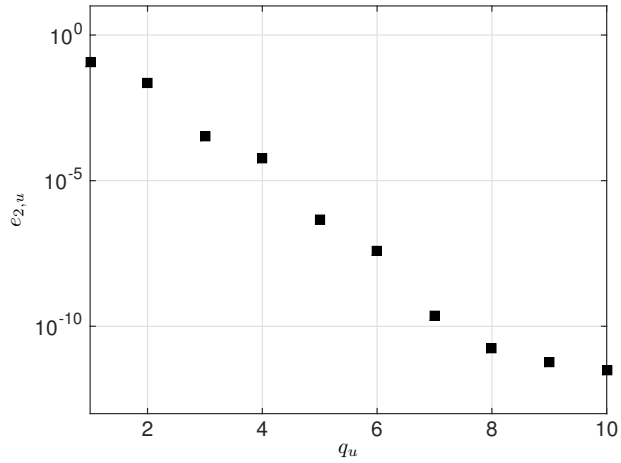

Figure 14: This plot of error versus element order illustrates the spectral convergence of them method when applied to the Stoneley wave also considered in [4].

The computational domain is chosen as in [4], $(x_1, x_2) \in [-1, 1] \times [-20, 20]$ and we impose the exact solution via Dirichlet conditions on all boundaries. We use an upwind flux with $q_u = q_v$ and set CFL = 2.5 and solve until time 1. The results, found in Figure 14, clearly demonstrate the spectral convergence of the method with increasing q_u.

5.3.5. A Singular Source Problem

As an example with a singular source term we solve a variation of Lamb's problem. We consider a half space $x_2 \leq 0$ with a free surface at $x_2 = 0$ and a point force at $x_1 = 0$ and directed downward. The time dependence of the forcing is

$$
s(t) = -2(t-2)\pi^2 e^{-\pi^2(t-2)^2}.
$$

Using methods with $q_u = q_v = 3, 5$ and 7 and the upwind flux we evolve the solution until time 10 with CFL = 2. The discretization of the delta function follows directly from the properties of the delta function.

We record the solution on the surface at $x_1 = 5$ and compare the results to the semi-exact solution obtained by EX2DDIR by Berg and If. In Figure 15 we display the L_2 error in time at the receiver for a sequence of grids with square elements. As can be seen in the figure the errors asymptote to the rates of convergence 4, 6, and 8.

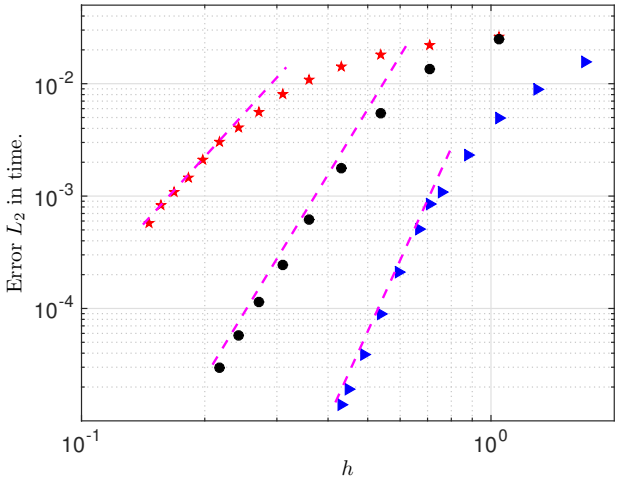

Figure 15: L_2-errors in time for the points source problem as a function of the grid size. The dashed asymptotes correspond to rates of convergence 4,6 and 8. The red stars, black circles and blue triangles correspond to methods with $q_u = q_v = 3, 5$ and 7 and upwind flux.

Remark 3. *Before moving on to the applications we note that the convergence rates we observe above does vary depending on the choice of numerical flux as well as the degree of the polynomial. Comparing the examples we also note that optimal order of accuracy at a rate $q_u + 1$ is achieved when $q_u = q_v$ and odd together with the upwind flux and homogenous traction or Dirichlet boundary conditions or non-homogenous traction boundary conditions. With non-homogenous Dirichlet conditions we observe a decrease in the rate of convergence by one half to $q_u + 1/2$. For best efficiency we thus recommend using the upwind flux and odd and equal degrees for the velocities and displacements.*

5.4. Applications

In this section we present two more practical applications of our method. The first problem is taken from the textbook by LeVeque, [28]. In the second application we consider how the uncertainty in the geometry of a pipe composed of two slightly non-concentric circles is propagated to quantities of interest, QOI.

5.4.1. Le Veque's Stiff Inclusion Problem

As a first application example we repeat Example 22.3 from [28]. A rectangular elastic slab $(x_1, x_2) \in [0, 2] \times [0, 1]$ with $\lambda = 2$ and $\mu = 1$ has a stiff inclusion, $(x_1, x_2) \in [0.5, 1.5] \times [0.4, 0.6]$, with $\lambda = 200$ and $\mu = 100$. The density is $\rho = 1$ throughout. The slab is initially at rest and the problem is forced through the left boundary, $x_1 = 0$, where we impose the Dirichlet boundary condition for the velocities

$$v_1(0, x_2, t) = \begin{cases} \sin(\pi t/0.025) & \text{if } t < 0.025, \\ 0 & \text{if } t \geq 0.025, \end{cases} \qquad v_2(0, x_1, t).$$

The three remaining sides of the slab are taken to be free of traction.

The solution is evolved until time $t = 0.4$ using $q_u = q_v = 15$ and the upwind flux. The grid is Cartesian with elements of size $1/20 \times 1/20$. In Figure 16 we display $|(\vec{F}_1)_2|$ and $|(\vec{F}_1)_1 + (\vec{F}_2)_2|$ at times 0.324 (left) and 0.4 (right). As in [28] we use a nonlinear color scale (most likely not the same) to emphasize the features

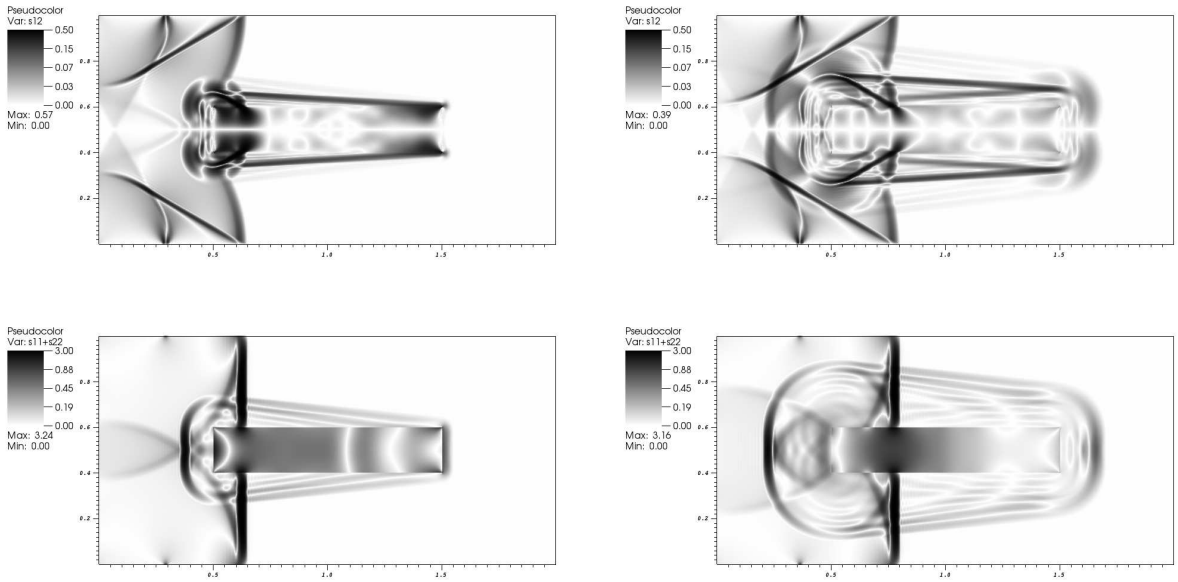

Figure 16: Example from LeVeque. Top: $|(\vec{F}_1)_2|$ at times 0.324 (left) and 0.4 (right). Bottom: $|(\vec{F}_1)_1 + (\vec{F}_2)_2|$ at times 0.324 (left) and 0.4 (right). Note that the color scale is nonlinear.

of the solution. The results appear to agree well with Figure 22.5 in [28]. To estimate the error we repeat the computation with $q_u = 13$ and compare the solutions at the final time. The errors in u_1 and u_2 made relative with their maximum amplitudes at the final time are displayed in Figure 17. The largest point-wise errors are on the order of 0.1% and a bit larger than expected. It may be that this is can be attributed to the lack of smoothness in the time dependent forcing, which is continuous but not differentiable at the time instances when it turns on and off.

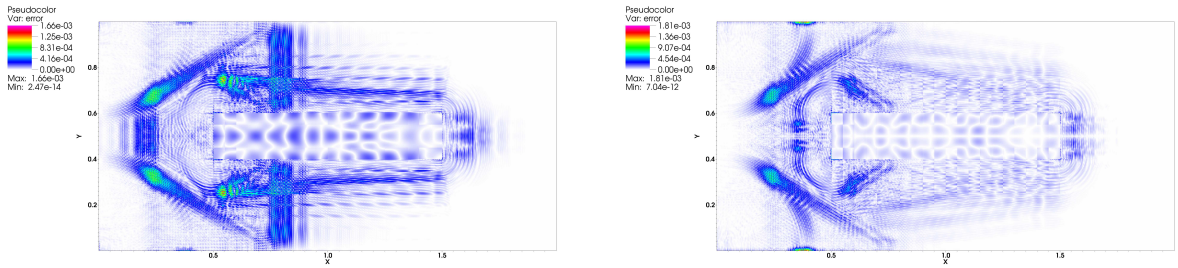

Figure 17: Pointwise errors relative to a reference solution in u_1 to the left and u_2 to the right.

5.4.2. Non Destructive Testing and Uncertainty Quantification

In this example we consider a pipe whose outer wall is a circle of radius 1 and centered at $(x_1, x_2) = (0, 0)$ and with an inner wall being a circle of radius 1/2. The center of the inner circle is $(x_1, x_2) = (-0.05 + 0.1\mathcal{X}_c, 0)$ where $\mathcal{X}_c$ is a uniformly distributed random variable, $\mathcal{X}_c \in \mathcal{U}(0, 1)$.

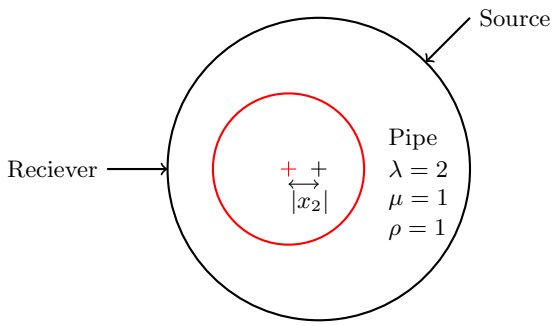

Figure 18: Sketch of the non-destructive testing of a pipe.

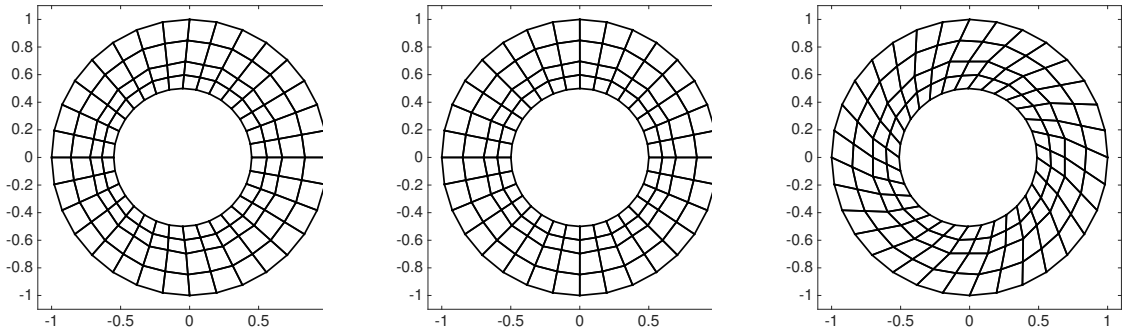

Figure 19: Typical grids used in the UQ example.

The pipe is initially at rest and we force the problem by prescribing a boundary forcing centered at the angle $\pi/4$; see Figure 18. Precisely we set:

$$\begin{aligned} \vec{n} \cdot \vec{F}_1 &= \cos(\theta) s_{\mathrm{NDT}}(x_1, x_2, t), \\ \vec{n} \cdot \vec{F}_2 &= \sin(\theta) s_{\mathrm{NDT}}(x_1, x_2, t), \end{aligned}$$

where

$$s_{\mathrm{NDT}}(x_1, x_2, t) = e^{-\left(\theta - \frac{\pi}{4}\right)/0.25)^2} \left(1 - \left(\frac{t-\tau}{\sigma}\right)^2\right) e^{-\frac{1}{2}\left(\frac{t-\tau}{\sigma}\right)^2}.$$

Here we choose $\sigma = 1/10$, $\tau = 0.75$. The pipe is clamped at the inner boundary, i.e. the displacement is zero.

Let $u_n(\theta, t)$ be the outward normal component of the displacement at the outer boundary and at an angle θ. We consider four (normalized) QOIs: $Q_i(\mathcal{X}_c) = \hat{Q}_i(\mathcal{X}_c)/\hat{Q}_i(0)$, $i = 1, \ldots, 4$,

$$\hat{Q}_1(\mathcal{X}_c) = \int_0^T \int_0^{2\pi} u_n(\theta, t, \mathcal{X}_c) d\theta dt, \tag{45}$$

$$\hat{Q}_2(\mathcal{X}_c) = \int_0^T u_n(\frac{\pi}{4}, t, \mathcal{X}_c) d\theta, \tag{46}$$

$$\hat{Q}_3(\mathcal{X}_c) = u_n(\frac{\pi}{4}, T, \mathcal{X}_c), \tag{47}$$

$$\hat{Q}_4(\mathcal{X}_c) = \int_0^{2\pi} u_n(\theta, T, \mathcal{X}_c) d\theta. \tag{48}$$

To compute approximations to the expected values of the QOIs we use stochastic collocation; see e.g. [29]. Here we expand $u_n(\theta, t, \mathcal{X}_c)$ in a basis consisting of Lagrange interpolating polynomials defined by the Clenshaw-Curtis nodes $z_l = (1 + \cos(\frac{l\pi}{N_{\rm CC}}))/2$, $l = 0, \ldots, N_{\rm CC}$. This allows us to approximate, e.g.,

$$E(\hat{Q}_1(\mathcal{X}_c)) \equiv \int_0^1 \hat{Q}_1(\mathcal{X}_c) d\mathcal{X}_c \approx \sum_{l=0}^{N_{\rm CC}} w_l \hat{Q}_1(z_l). \tag{49}$$

Here w_l are the Clenshaw-Curtis weights associated with the nodes z_l. The integrals in the above formulas are computed using the trapezoidal rule.

As the Clenshaw-Curtis nodes are nested we can use the the finest grid as a reference solution and then inspect how the convergence of the expected values depends on the order of the numerical method as well as the number of samples in the stochastic collocation. As described in [29] the rate of convergence of a stochastic collocation method depends on the quadrature rule, the numerical method and the smoothness of the solution with respect to the probability space. The paper [29] presents a rather complete characterization of how the smoothness of the solution to the scalar wave equation depends on uncertainties in the material coefficients for different initial data.

The examples and theory developed in [29] are motivated by applications in seismology and exploration where the material parameters in the ground are uncertain. The example we present here can be thought of as an example arising in non-destructive testing where the material parameters are typically well known but the internal geometry could perhaps be known only with certain precision (for example due to manufacturing tolerances). We note that in non-destructive testing the end goal is often to find out if there are imperfections or defects inside the object. This type of problem could also be treated with our method but is beyond the scope of this paper.

To assess the convergence of the expected values we perform $N_{\rm CC} + 1$ computations for the different locations of the interior circle. Each of the $N_{\rm CC} + 1$ realizations requires its own grid which we generate automatically using the grid generator GMSH [30]. Typically the grids are of high quality, see the two leftmost grids in Figure 19, but occasionally the elements are somewhat deformed, as in the grid to the right in Figure 19. The handling of the curved elements is described in §4.

We choose $N_{\rm CC} = 2^8$ and compute solutions at time $T = 5$ using $q_u = q_v = 3, 5, 7, 9$ and using the upwind flux with CFL $= 0.5$. The computed expected values of the three different QOI using $N_{\rm CC} = 2^8$ can be found in Table 9. As can be seen the QOIs converge with increasing order.

Using the computation with $q_u = 9$ as a reference value we then compute errors for all the nested levels of the Clenshaw-Curtis nodes. The results are displayed in Figure 20. It is clear that the convergence of the expected values depends on the number of quadrature points used as well as on the discretization. The results for all the QOIs are improved with increased order but the jumps in the error levels are somewhat different. The principal difference between the QOIs is whether or not they are integrated in space and/or in time. A careful study of the smoothing properties of the integrals and the numerical errors associated with their approximation, following the ideas in [29], is a topic for future study.

Table 9: Expected values for the four different QOI for different $q_u = q_v$ using the full set of Clenshaw-Curtis nodes.

q_u	$E(\hat{Q}_1)$	$E(\hat{Q}_2)$	$E(\hat{Q}_3)$	$E(\hat{Q}_4)$
3	0.7593239	0.9148653	0.76725	-0.21405469
5	0.7577141	0.9183187	0.77871	-0.23945853
7	0.7577082	0.9185063	0.77758	-0.23951380
9	0.7577076	0.9185057	0.77761	-0.23951366

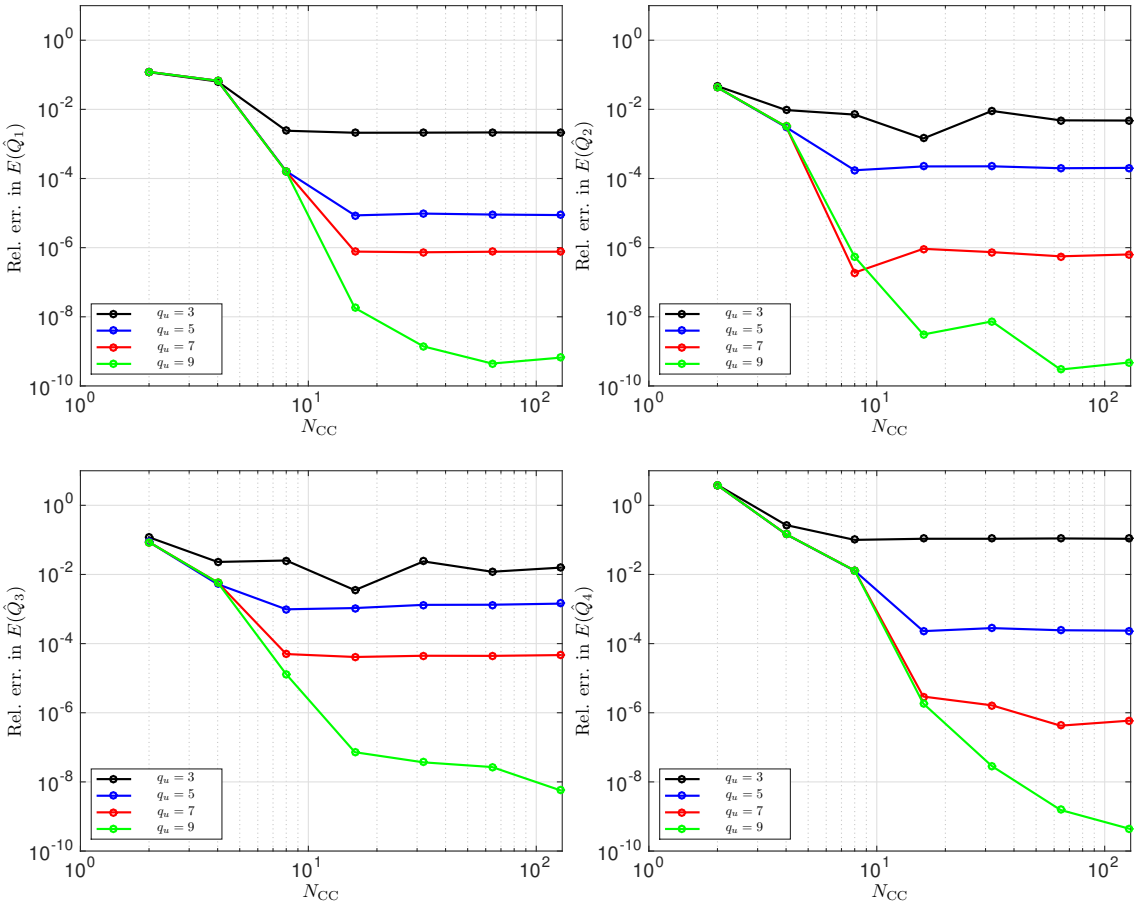

Figure 20: Relative errors in the expected values of the different QOIs as a function of the number of quadrature points and for different q_u.

6. Summary

In summary, we have presented an application of our general formulation [1] to the elastic wave equation. The resulting semi-discrete method is of arbitrary order of accuracy and either conserves or dissipates the elastic energy, depending on the choice of flux.

We investigated how the choice of approximation spaces and fluxes impacts the observed order of accuracy. The best (and most robust) order of convergence is obtained when the upwind flux is used and when the degrees of the polynomials used to approximate the velocity and displacement are the same.

The accuracy and spectral properties of the spatial discretization were demonstrated on a sequence of problems containing both manufactured and exact solutions. Two more realistic applications of the method were also presented.

Finally, our current implementation of the method is freely available at [2]; we encourage the reader to try it out.

7. Acknowledgement

We would like to thank the anonymous reviewers for their constructive comments and suggestions.

References

[1] D. Appelö, T. Hagstrom, A new discontinuous Galerkin formulation for wave equations in second order form, SIAM Journal On Numerical Analysis 53 (6) (2015) 2705–2726.

[2] D. Appelö, T. Hagstrom, `dg_dath`, `bitbucket.org/appelo/dg_dath_elastic_v1.0` (2015).

[3] J. S. Hesthaven, T. Warburton, Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications, Vol. 54, Springer, New York, 2008.

[4] L. C. Wilcox, G. Stadler, C. Burstedde, O. Ghattas, A high-order discontinuous Galerkin method for wave propagation through coupled elastic–acoustic media, Journal of Computational Physics 229 (24) (2010) 9373 – 9396.

[5] S. Delcourte, L. Fezoui, N. Glinsky-Olivier, A high-order Discontinuous Galerkin method for the seismic wave propagation, ESAIM: Proc. 27 (2009) 70–89.

[6] M. Käser, M. Dumbser, An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes - i. the two-dimensional isotropic case with external source terms, Geophysical Journal International 166 (2) (2006) 855–877.

[7] E. T. Chung, C. Y. Lam, J. Qian, A staggered discontinuous Galerkin method for the simulation of seismic waves with surface topography, Geophysics 80 (4) (2015) T119–T135.

[8] B. Riviere, M. Wheeler, Discontinuous finite element methods for acoustic and elastic wave problems. part i: semidiscrete error estimates, Contemporary Mathematics 329 (2003) 271–282.

[9] J. D. De Basabe, M. K. Sen, M. F. Wheeler, The interior penalty discontinuous Galerkin method for elastic wave propagation: grid dispersion, Geophysical Journal International 175 (1) (2008) 83–93.

[10] P. F. Antonietti, B. Ayuso de Dios, I. Mazzieri, A. Quarteroni, Stability analysis of discontinuous Galerkin approximations to the elastodynamics problem, Journal of Scientific Computing 68 (1) (2016) 143–170.

[11] T. J. R. Hughes, G. M. Hulbert, Space-time finite element methods for elastodynamics: Formulations and error estimates, Comput. Methods Appl. Mech. Eng. 66 (3) (1988) 339–363.

[12] R. Abedi, B. Petracovici, R. B. Haber, A space–time discontinuous Galerkin method for linearized elastodynamics with element-wise momentum balance, Computer Methods in Applied Mechanics and Engineering 195 (25) (2006) 3247–3273.

[13] S. Miller, B. Kraczek, R. Haber, D. Johnson, Multi-field spacetime discontinuous Galerkin methods for linearized elastodynamics, Computer Methods in Applied Mechanics and Engineering 199 (1–4) (2009) 34 – 47.

[14] B. Sjögreen, N. A. Petersson, A fourth order accurate finite difference scheme for the elastic wave equation in second order formulation, Journal of Scientific Computing 52 (1) (2012) 17–48.

[15] K. Duru, G. Kreiss, K. Mattsson, Stable and high-order accurate boundary treatments for the elastic wave equation on second-order form, SIAM Journal on Scientific Computing 36 (6) (2014) A2787–A2818.

[16] D. Appelö, N. A. Petersson, A stable finite difference method for the elastic wave equation on complex geometries with free surfaces, Communications in Computational Physics 5 (1) (2009) 84–107.

[17] H.-O. Keiss, J. Oliger, Comparison of accurate methods for the integration of hyperbolic equations, Tellus 24 (1972) 199–215.

[18] H.-O. Kreiss, N. A. Petersson, Boundary estimates for the elastic wave equation in almost incompressible materials, SIAM Journal on Numerical Analysis 50 (3) (2012) 1556–1580.

[19] M. Ainsworth, Dispersive and dissipative behavior of high-order discontinuous Galerkin finite element methods, J. Comput. Phys. 198 (2004) 106–130.

[20] R. Sakamoto, Hyperbolic Boundary Value Problems Cambridge, Cambridge University Press, 1978.

[21] H. O. Kreiss, O. E. Ortiz, N. A. Petersson, Initial-boundary value problems for second order systems of partial differential equations, ESAIM M2AN (46) (2012) 559–593.

[22] D. A. Kopriva, Implementing spectral methods for partial differential equations: Algorithms for scientists and engineers, Scientific computation, Springer, 2009.

[23] B. Fornberg, A Practical Guide to Pseudospectral Methods, Vol. 1, Cambridge University Press, Cambridge, 1996.

[24] A. Klöckner, T. Warburton, J. Bridge, J. S. Hesthaven, Nodal discontinuous Galerkin methods on graphics processors, Journal of Computational Physics 228 (21) (2009) 7863–7882.

[25] T. Warburton, A low-storage curvilinear discontinuous Galerkin method for wave problems, SIAM Journal on Scientific Computing 35 (4) (2013) A1987–A2012.

[26] D. Appelö, J. W. Banks, W. D. Henshaw, D. W. Schwendeman, Numerical methods for solid mechanics on overlapping grids: Linear elasticity, Journal of Computational Physics 231 (18) (2012) 6012–6050.

[27] R. Stoneley, Elastic waves at the surface of separation of two solids, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences 106 (738) (1924) 416–428.

[28] R. J. LeVeque, Finite volume methods for hyperbolic problems, Cambridge University Press, Cambridge, 2002.

[29] M. Motamed, F. Nobile, R. Tempone, A stochastic collocation method for the second order wave equation with a discontinuous random speed, Numerische Mathematik 123 (3) (2013) 493–536.

[30] C. Geuzaine, J.-F. Remacle, Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities, International Journal for Numerical Methods in Engineering 79 (11) (2009) 1309–1331.