GPU-accelerated discontinuous Galerkin methods on hybrid meshes

Jesse Chan^{a,*}, Zheng Wang^a, Axel Modave^a, Jean-Francois Remacle^b, T. Warburton^a

^a*Department of Computational and Applied Mathematics, Rice University, 6100 Main St, Houston, TX, 77005*

^b*School of Engineering, Universite catholique de Louvain, Avenue Georges Lemaitre 4, B-1348 Louvain-la-Neuve, Belgium*

Abstract

We present a time-explicit discontinuous Galerkin (DG) solver for the time-domain acoustic wave equation on hybrid meshes containing vertex-mapped hexahedral, wedge, pyramidal and tetrahedral elements. Discretely energy-stable formulations are presented for both Gauss-Legendre and Gauss-Legendre-Lobatto (Spectral Element) nodal bases for the hexahedron. Stable timestep restrictions for hybrid meshes are derived by bounding the spectral radius of the DG operator using order-dependent constants in trace and Markov inequalities. Computational efficiency is achieved under a combination of element-specific kernels (including new quadrature-free operators for the pyramid), multi-rate timestepping, and acceleration using Graphics Processing Units.

1. Introduction

Highly efficient solution techniques exist for high order finite element and Galerkin discretizations on hexahedral meshes. Presently, producing high quality hexahedral meshes for complex domains is a difficult and non-robust procedure. Hybrid meshes, which consist of wedge and pyramidal elements in addition to hexahedra and tetrahedra, have been proposed to leverage the efficiency of hexahedral elements for more general geometries. Such meshes are commonly produced by adding a transitional layer of wedges and pyramids to a structured hexahedral mesh and constructing an unstructured tetrahedral mesh to fill in the remaining space [40].

Recently, techniques have been developed to automatically generate conforming hex-dominant hybrid meshes with a high percentage of hexahedral elements [2], which we aim to leverage to develop efficient solvers on more general geometries. We note that the utility of high order time-explicit DG using hybrid (including polygonal) meshes has been explored for polynomial physical frame discretizations in [18]. We consider in this work approximation spaces defined under a reference-to-physical mapping. [Additionally, we restrict ourselves to meshes consisting of vertex-mapped elements, whose physical mappings are defined only in terms of vertex positions and lowest-order vertex functions on the reference element.](#)

Spectral and *hp*-finite element methods on hybrid meshes were explored early on by Sherwin, Karniadakis, Warburton, and Kirby [41, 43, 25] using orthogonal polynomial basis functions constructed for hexahedral, wedge, pyramidal, and tetrahedral elements. While approximation spaces for the hexahedron, wedge, and tetrahedron have remained relatively unchanged since their introduction, more recent work has focused on the construction of alternative basis functions and approximation spaces for the pyramid. Bedrosian introduced a set of low-order vertex functions for the pyramid which yielded polynomial traces, allowing for conformity with low order finite elements for other shapes [3]. Unlike other elements, however, his pyramidal shape functions were rational in nature. Bergot, Cohen, and Duruffe extended this construction in [4, 7] to produce a high order basis for the pyramid (which we refer to as the *rational basis*). It was additionally shown that, under non-affine mappings of the reference pyramid, the mapped rational basis contains a

*Principal Corresponding author

Email addresses: Jesse.Chan@caam.rice.edu (Jesse Chan), zw14@caam.rice.edu (Zheng Wang), modave@caam.rice.edu (Axel Modave), Jean-Francois.Remacle@uclouvain.be (Jean-Francois Remacle), timwar@caam.rice.edu (T. Warburton)

complete polynomial space. As a result, the rational basis provides optimal high order convergence rates on vertex-mapped elements, which is not true of the polynomial pyramid basis. Rational pyramidal bases were also adapted to the $H(\text{curl}), H(\text{div})$ setting and exact sequence spaces in [6, 5, 35]. The construction of exact sequence high order bases for elements of all types (including pyramids) is addressed in recent work by Fuentes, Keith, Demkowicz, and Nagaraj [16].

Our focus is on high order discontinuous Galerkin (DG) methods on hybrid meshes. We pay particular attention to the efficient implementation of solvers on hardware accelerators. The computation-intensive nature of time-explicit methods lends itself well to modern accelerator architectures such as Graphics Processing Units (GPUs), and was exploited by Klöckner, Warburton, Bridge and Hesthaven to construct an efficient high order DG solver on tetrahedral meshes using a single GPU [26]. Exposing both coarse and fine-grained parallelism resulted in several times speedup [relative to serial CPU implementations](#) for meshes with around a hundred thousand elements at high orders of approximation. Increases in on-board memory in modern GPUs have allowed for the solution of even larger problems on a single unit. Additional gains in efficiency and problem size may be achieved via multi-rate timestepping and distributed-memory parallelism [21, 20]. However, these technologies have been developed on all-hex or all-tetrahedra meshes, and while heterogeneous computing on hybrid meshes has been explored in the context of fluid dynamics and the closely related Flux Reconstruction method [48], GPU-accelerated DG methods on hybrid meshes have not yet been analyzed in detail.

Crucial to the efficiency of GPU-accelerated DG methods is limiting the memory required. Since each GPU has a relatively small amount of on-device storage, maximum problem sizes are typically [determined by available memory](#). Classical techniques for hexahedral and tetrahedral elements limit the amount of data that must be stored, and recent developments in the construction of basis functions [45, 46, 11] make it possible to develop low-storage DG methods on pyramids and wedges as well. This work leverages each of these technologies to construct a low-memory GPU-accelerated high order DG solver on hybrid meshes for the acoustic wave equation. Additionally, order-dependent global and local timestep restrictions for general hybrid meshes are derived using constants in discrete trace inequalities for each element.

[The aim of this work is to present a unified discussion of theoretical, practical, and computational aspects of DG on hybrid meshes, including implementation and performance on GPU architectures. These include issues of accuracy \(high order approximation\), stability \(energy stability of the variational formulation, stable timestep restrictions\), and efficiency \(computational performance, storage requirements\).](#) The structure of this paper is as follows: Section 2 introduces a low-storage DG method for hybrid meshes. The low-storage treatment of mass matrices is addressed through a careful choice of basis functions for each specific type of element. The variational formulation and element-specific operations are given, and a computational implementation on many-core architectures is described. Section 3 derives order-dependent timestep restrictions for each type of element based on the constants in discrete trace inequalities. Finally, Section 4 reports numerical and computational results.

2. A Low-storage DG method on hybrid meshes

Typical DG methods result in a system of ordinary differential equations of the form

$$\frac{du}{d\tau} = M^{-1}Au,$$

where τ is time, A is a linear operator defined by the discretization, and M is the global mass matrix, which is block diagonal in nature. Efficient implementations of DG must account for the inverse of M in a fast and memory-efficient manner. One common strategy is to store the factorization or explicit inverse of each block of M^{-1} ; however, on GPU and other accelerator architectures, on-device memory is typically limited to a few gigabytes. Explicit storage of factorizations can rapidly exhaust available memory and limit the maximum problem size on a single GPU at high orders of approximation. For example, assuming a mesh of 1 million planar tetrahedra in single precision, at order $N = 4$, storage of M^{-1} already exceeds the 1 GB of memory available on early GPUs used to accelerate DG methods [26], without accounting for the storage cost of other relevant data.

Figure 1: Types of elements present in hex-dominant meshes. From left to right: reference [hexahedron](#), wedge, pyramid, and [tetrahedron](#).

We consider meshes consisting of hexahedra, wedges, pyramids, and tetrahedra. Given an order $N = 1$ basis on each element, we may then define low-order vertex functions. A map from the reference element to a physical element may then be defined by interpolating physical vertex positions with these vertex functions. We refer to such elements as *vertex-mapped* elements, and restrict ourselves to meshes consisting of such elements in this work. Furthermore, we define J to be the determinant of the mapping Jacobian, such that

$$\int_K u = \int_{\hat{K}} u J,$$

where $K, \hat{K}$ are the physical and reference elements, respectively.

2.1. Basis functions

We will sidestep the cost of storage of M^{-1} by tailoring our choice of basis functions for each element. Assuming that all elements are vertex-mapped, we will construct basis functions for the hexahedron, wedge, and pyramid which yield diagonal mass matrices. For vertex-mapped tetrahedra, the mass matrix will not be diagonal, but will be a scalar multiple of the reference mass matrix.

2.1.1. Hexahedron

We take the reference hexahedron to be the bi-unit cube, with coordinates $-1 \leq a, b, c \leq 1$. The approximation space of order N on the hexahedron is simply the tensor product space of polynomials of order N in each direction. For a vertex-mapped hexahedra $\mathcal{H}$, this space and its dimension N_p are given as follows

$$P_N(\mathcal{H}) = \{x^i y^j z^k, \quad 0 \leq i, j, k \leq N\}, \quad N_p = (N + 1)^3.$$

We also have that $J \in P_1(\hat{\mathcal{H}})$, and an appropriate quadrature rule can be constructed using a tensor product of standard 1D Gauss-Legendre rules. If we take the orthogonal basis

$$\phi_{ijk}(a, b, c) = \ell_i(a) \ell_j(b) \ell_k(c) \quad 0 \leq i, j, k \leq N,$$

where ℓ_i is the i th Lagrange polynomial at the i th Gauss-Legendre node, the mass matrix defined by

$$M_{ijk,lmn} = \int_{\mathcal{H}} \phi_{ijk} \phi_{lmn} \, dx = \int_{\hat{\mathcal{H}}} \phi_{ijk} \phi_{lmn} J \, d\hat{x}$$

is diagonal for $J \in P_1(\hat{\mathcal{H}})$. To apply M^{-1} , we store only of the inverse of the diagonal of the full mass matrix, which requires the storage of $N_p K$ values, instead of storing the $O(N_p^2 K)$ entries required for the factorization of the global block-diagonal mass matrix.

We may also achieve a diagonal mass matrix by taking Lagrange polynomials at tensor product Gauss-Legendre-Lobatto nodes, which includes points at both -1 and 1 . Applying Gauss-Legendre-Lobatto quadrature results in a diagonal mass matrix, and is the basis behind the Spectral Element Method [37, 24]. We refer to the choice of basis using Gauss-Legendre nodes as the GL basis, and the basis using Gauss-Legendre-Lobatto nodes as the SEM basis. We will compare both of these choices in this work.

Figure 2: SEM and GL nodes on the quadrilateral. For GL nodes, additional surface nodes must be specified (shown as squares).

2.1.2. Tetrahedron

For the reference tetrahedron, an orthogonal basis may be given as the product of Jacobi polynomials on the bi-unit cube [38, 27, 14]

$$\phi_{ijk}(a, b, c) = P_i^{0,0}(a)(1-b)^i P_j^{2i+1,0}(b)(1-c)^{i+j} P_k^{2(i+j+1),0}(c).$$

The tetrahedral basis is then given by the Duffy mapping from the bi-unit cube to the bi-unit right tetrahedron with coordinates (r, s, t)

$$r = (1+a) \left(\frac{1-b}{2} \right) \left(\frac{1-c}{2} \right) - 1, \quad s = (1+b) \left(\frac{1-c}{2} \right) - 1, \quad t = c.$$

The approximation space for a vertex-mapped tetrahedra is the space of polynomials of total order N with dimension N_p

$$P(\mathcal{T}) = \{x^i y^j z^k, \quad i+j+k \leq N\}, \quad N_p = \frac{(N+1)(N+2)(N+3)}{6}.$$

A convenient fact about vertex-mapped (planar) tetrahedra is that J is constant on each element. Thus, each local mass matrix is simply a constant scaling of the reference mass matrix, and only the reference mass matrix is stored. In our implementation, we choose a nodal Lagrange basis (defined implicitly using the above orthogonal basis) where the nodes are placed at optimized interpolation points on the simplex [44].

2.1.3. Wedge

An orthogonal basis for the reference wedge may be given in terms of Jacobi polynomials on the bi-unit cube

$$\phi_{ijk}(a, b, c) = P_i^{0,0}(a) P_j^{0,0}(b) \left(\frac{1-c}{2} \right)^i P_k^{2i+1,0}(c), \quad 0 \leq i, j \leq N, \quad 0 \leq k \leq N-i.$$

Applying a Duffy-type transform gives a basis on the reference bi-unit wedge with coordinates (r, s, t)

$$r = (1+a) \left(\frac{1-c}{2} \right) - 1, \quad s = b, \quad t = c.$$

For a vertex-mapped wedge $\mathcal{W}$, $J \in P_1(\widehat{\mathcal{W}})$, and the approximation space of order N on the wedge and its dimension are

$$P_N(\mathcal{W}) = \{x^i y^j z^k, \quad 0 \leq i, j \leq N, \quad 0 \leq k \leq N-i\}, \quad N_p = \frac{(N+1)^2(N+2)}{2}.$$

We may construct an appropriate quadrature rule on the reference wedge based on the tensor product of Gauss-Legendre rules in the r, s coordinate, and a Gauss-Legendre rule with weights $(1, 0)$ in the t coordinate. Alternatively, a wedge cubature may be constructed by taking the tensor product of an optimized cubature

on the triangle with Gauss-Legendre cubature in the s coordinate. We take the latter approach in this work, using rules from Xiao and Gimbutas [49]. The triangle quadrature is chosen such that the rule is exact for polynomials of total order $2N + 1$, similar to Gauss-Legendre quadrature.

Decreasing storage costs for the wedge mass matrix is less straightforward. One common approach is to construct a basis using the tensor product of triangular basis functions with Lagrange polynomials at $N + 1$ Gauss-Legendre points. The resulting mapped basis is orthogonal in the b direction, and the local mass matrix becomes block diagonal, with each block corresponding to a triangular slice of the wedge. Unfortunately, for general vertex-mapped wedges, J is bilinear on each triangular slice, and each block of the local mass matrix corresponds to a J -weighted L^2 inner product on a triangle. Since there is **no clear polynomial basis** which is orthogonal in this inner product for an arbitrary bilinear J , the blocks of each local mass matrix are typically dense and distinct from element to element. Factorizations of each block are required in order to efficiently apply the mass matrix inverse, resulting in increased storage costs.

We may address this storage cost by using a Low-Storage Curvilinear DG (LSC-DG) approach [45, 46], where we define the H^1 non-conforming basis $\tilde{\phi}_{ijk}$

$$\tilde{\phi}_{ijk} = \frac{\phi_{ijk}}{\sqrt{J}},$$

such that

$$M_{ijk,lmn} = \int_{\widehat{\mathcal{W}}} \tilde{\phi}_{ijk} \tilde{\phi}_{lmn} J d\widehat{x} \int_{\widehat{\mathcal{W}}} \frac{\phi_{ijk}}{\sqrt{J}} \frac{\phi_{lmn}}{\sqrt{J}} J d\widehat{x} = \widehat{M}_{ijk,lmn}.$$

As a result, each mass matrix is identical to the reference mass matrix, allowing us to apply M^{-1} while accounting for only a single mass matrix inverse over all elements.

The use of LSC-DG results in a rational basis; consequentially, standard quadratures are not exact and aliasing errors may result from underintegration. To sidestep these issues, we restrict ourselves to variational formulations that may be written in a skew-symmetric fashion, such that DG is energy stable **regardless** of the choice of quadrature [46]. Additionally, since our approximation space now contains rational functions, standard polynomial approximation results do not hold for LSC-DG wedges. Under a quasi-regular scaling assumption on elements in our mesh, we recover instead an L^2 approximation bound of the form

$$\|u - \Pi_w u\|_{L^2(K)} \leq Ch^{N+1} \left\| \frac{1}{\sqrt{J}} \right\|_{L^\infty(K)} \left\| \sqrt{J} \right\|_{W^{N+1,\infty}(K)} \|u\|_{W^{N+1,2}(K)},$$

where the weighted projection $\Pi_w u$ is the projection onto the LSC-DG space, h is the size of the element K , and $\|\cdot\|_{W^{N+1,\infty}(K)}$ denotes the L^∞ Sobolev norm of order $N + 1$ over K . For comparison, the projection error bound for curvilinear mappings using standard mapped bases replaces the Sobolev in the above bound with an L^∞ norm. In other words, while the approximation error for standard curvilinear elements depends only on the extremal values of J , it depends additionally on the smoothness of J when using rational LSC-DG basis functions.

2.1.4. Pyramid

Unlike hexahedral, wedge, and tetrahedral elements, the basis functions for the pyramid are rational in nature. For a vertex-mapped pyramid $\mathcal{P}$, the approximation space and dimension N_p are

$$B_N(\mathcal{P}) = P^N(\mathcal{P}) \oplus \sum_{k=0}^{N-1} \left\{ \left(\frac{xy}{1-z} \right)^{N-k} x^i y^j, 0 \leq i + j \leq k \right\}, \quad N_p = \frac{(N+1)(N+2)(2N+3)}{6},$$

where $P_N(\mathcal{P})$ is the space of polynomials of total order N on the physical pyramid. For a vertex-mapped pyramid, we have that the change of variables factor $J \in B_1(\widehat{\mathcal{P}})$. An appropriate quadrature rule is defined on the reference pyramid using a tensor product of Gauss-Legendre rules in the r, s coordinate, and a Gauss-Legendre rule with weights $(2, 0)$ in the t coordinate.

Since J is again non-constant, each local mass matrix is distinct and dense, increasing storage costs for DG methods on pyramids. However, it was shown in [11] that there exists an orthogonal basis ϕ_{ijk} on vertex-mapped pyramids which spans $B_N(\mathcal{P})$. This basis is defined on the bi-unit cube as

$$\phi_{ijk}(a, b, c) = \sqrt{\frac{C_{N-k}^N}{w_i w_j}} \ell_i^k(a) \ell_j^k(b) \left(\frac{1-c}{2}\right)^k P_{N-k}^{2k+3}(c), \quad 0 \leq i, j \leq k, \quad 0 \leq k \leq N,$$

where $\ell_i^k(a), \ell_j^k(b)$ are order k Lagrange polynomials at $(k+1)$ Gauss-Legendre nodes, and $P_{N-k}^{2k+3}(c)$ is the Jacobi polynomial of degree k with order-dependent weight $2k+3$. w_i is the i th Gauss-Legendre quadrature weight, and the normalization constant $C_{N-k}^N = (N+2)/(2^{2k+2}(2k+3))$. Due to the use of both orthogonal Jacobi polynomials and nodal polynomials at quadrature points, we refer to ϕ_{ijk} as the “semi-nodal” basis.

The pyramid basis is then defined under the Duffy-type mapping from the bi-unit cube to the bi-unit right pyramid

$$r = (1+a) \left(\frac{1-c}{2}\right) - 1, \quad s = (1+b) \left(\frac{1-c}{2}\right) - 1, \quad t = c.$$

The mass matrix is diagonal for vertex-mapped pyramids under such a basis, and the entries of the mass matrix are the evaluation of J at the quadrature points a_i^k, b_j^k . As with hexahedral elements, the application of M^{-1} requires only storage of the diagonals of each local mass matrix.

2.2. Variational formulation

We consider the acoustic wave equation on domain Ω . We assume zero forcing and a free surface boundary condition $p = 0$ on the boundary of the domain $\partial\Omega$. This may be written in first order form

$$\begin{aligned} \frac{1}{\kappa} \frac{\partial p}{\partial \tau} + \nabla \cdot \mathbf{u} &= 0 \\ \rho \frac{\partial \mathbf{u}}{\partial \tau} + \nabla p &= 0, \end{aligned}$$

where p is acoustic pressure, $\mathbf{u}$ is velocity, and ρ and κ are density and bulk modulus, respectively. We assume also a triangulation of the domain Ω_h consisting of elements K with faces f , and that ρ and κ are piecewise constant on each element. We will denote the union of the faces f as Γ_h . Let $(p^-, \mathbf{u}^-)$ denote the solution fields on the face of an element K , and let $(p^+, \mathbf{u}^+)$ denote the solution on the neighboring element adjacent to that face. We may then define the jump of p and the normal jump and average of the vector velocity $\mathbf{u}$ componentwise

$$[p] = p^+ - p^-, \quad [\mathbf{u}] = \mathbf{u}^+ - \mathbf{u}^-, \quad \{\{\mathbf{u}\}\} = \frac{\mathbf{u}^+ + \mathbf{u}^-}{2}.$$

Defining $\mathbf{n}^-$ as the outward unit normal on a given face, the (local) variational formulation for the discontinuous Galerkin method is then

$$\int_K \frac{1}{\kappa} \frac{\partial p}{\partial \tau} \phi^- \, dx = \int_K \mathbf{u} \cdot \nabla \phi^- \, dx + \int_{\partial K} \left(\frac{1}{2} \tau_p [p] - \mathbf{n}^- \cdot \{\{\mathbf{u}\}\} \right) \phi^- \, dx \quad (1)$$

$$\int_K \rho \frac{\partial \mathbf{u}}{\partial \tau} \psi^- \, dx = - \int_K \nabla p \cdot \psi^- \, dx + \int_{\partial K} \frac{1}{2} (\tau_u [\mathbf{u}] \cdot \mathbf{n}^- - [p]) \psi^- \cdot \mathbf{n}^- \, dx, \quad (2)$$

where $\tau_p = 1/\{\{\rho c\}\}$, $\tau_u = \{\{\rho c\}\}$, and $c^2 = \kappa/\rho$ is the speed of sound. Only one equation has been integrated by parts, resulting in a skew-symmetric and energy-stable formulation (the choice of which equation to integrate by parts is arbitrary; both produce skew-symmetric formulations). Proofs of energy stability typically rely on the exact evaluation of integrals [32]. Due to the use of LSC-DG on wedges, physical basis functions become rational in nature and are not exactly integrable using standard quadrature rules. The skew-symmetry of the variational form guarantees energy stability regardless of quadrature rule [46].

Figure 3: The choice of nodal basis for the hexahedron determines the quadrature rule on quadrilateral faces of both pyramids (shown above) and wedges.

While we adopt a skew-symmetric variational formulation when necessary for energy stability, we do not always compute with it. Instead, we use the fact that the skew-symmetric form (2) is sometimes equivalent to the twice integrated-by-parts “strong form” [22]

$$\begin{aligned} \int_K \frac{1}{\kappa} \frac{\partial p}{\partial \tau} \phi^- dx &= - \int_K \nabla \cdot \mathbf{u} \phi^- dx + \int_{\partial K} \frac{1}{2} (\tau_p \llbracket p \rrbracket - \mathbf{n} \cdot \llbracket \mathbf{u} \rrbracket) \phi^- dx \\ \int_K \rho \frac{\partial \mathbf{u}}{\partial \tau} \psi^- dx &= - \int_K \nabla p \cdot \psi^- dx + \int_{\partial K} \frac{1}{2} (\tau_u \llbracket \mathbf{u} \rrbracket \cdot \mathbf{n}^- - \llbracket p \rrbracket) \psi^- \cdot \mathbf{n}^- dx. \end{aligned}$$

This holds, for example, if integration by parts holds when volume and surface integrals are replaced with quadrature approximations, and is related to the “summation-by-parts” property in finite differences.

The choice between the strong/skew formulations is governed by the choice between GL and SEM nodal bases for the hexahedron, which determines the quadrature rule on quadrilateral faces (see Figure 3). When SEM quadrature is used, surface integrals over the quadrilateral face are computed inexactly; as a result, the strong formulation may not be equivalent to the weak formulation, and may not be energy stable.

2.2.1. *Stable formulations for hexahedra*

For hexahedral elements, the strong and skew-symmetric forms are equivalent under both GL and SEM bases. This was shown by Kopriva and Gassner in [28] and for SEM, relies on the exact cancellation of Gauss-Legendre-Lobatto quadrature errors in volume and surface integrals. We may thus exploit the improved computational efficiency of the strong form for hexahedra for both SEM or GL nodal bases. For tetrahedra, since the mapping is constant for planar simplices, volume integrals are computed exactly using quadrature free techniques, and surface integrals are computed exactly using quadrature. Thus, the skew-symmetric and strong forms are equivalent for vertex-mapped tetrahedra, and we adopt the strong form for improved computational efficiency.

2.2.2. *Stable formulations for wedges and pyramids*

For the wedge, we are restricted to the skew symmetric form, since integrals over the rational LSC-DG basis are inexact for any polynomial cubature rule. For the pyramid, the energy stability of the variational form depends on the choice of SEM or GL quadrature for the quadrilateral face. Under Gauss-Legendre quadrature, surface integrals for the pyramid are exact, and both the strong and skew-symmetric form are equivalent. However, if SEM quadrature is used for the quadrilateral face, the surface integral is inexact and we must use the skew-symmetric form for stability.

We may now choose freely between a SEM or GL nodal basis for the hexahedron, which also determines the choice of quadrature on quadrilateral faces. SEM nodes tends to be more efficient than Gauss-Legendre nodes — with Gauss-Legendre quadrature, since the nodal points lie in the interior of the hexahedron, an additional evaluation step is necessary to compute the solution at points on the surface. Since the SEM nodal basis contains both volume and surface quadrature points as degrees of freedom, evaluation of the solution on the surface requires only retrieval of a single degree of freedom, and this extra step is avoided.

	SEM Formulation	GL Formulation
Tetrahedra	Nodal, strong form	Nodal, strong form
Pyramid	Semi-nodal, skew-symmetric form	Semi-nodal, strong form
Wedge	Modal LSC-DG, skew-symmetric form	Modal LSC-DG, skew-symmetric form
Hexahedra	Nodal SEM, strong form	Nodal Gauss-Legendre, strong form

Table 1: Summary of stable bases and formulations for element type-specific variational forms.

Additionally, due to the Lagrange property of the SEM nodal basis functions, the surface contributions are sparser for the SEM formulation than for the GL formulation.

However, it is also known that the inexact integration in SEM quadrature reduces the accuracy of the resulting solution [28], though the inner products generated by both underintegrated SEM quadrature and full Gauss-Legendre quadrature are known to be equivalent (with constants that do not grow in N). We will also show in Section 4.1 that SEM underintegration can reduce the observed order of convergence for smooth solutions. Additionally, the use of SEM on hybrid meshes also restricts the pyramid to the skew-symmetric form, which is typically less computationally efficient than the strong form, while the use of Gauss-Legendre nodes on the hex allows for the use of the strong form for the pyramid.

2.2.3. Summary of energy stable formulations

We naturally arrive at two energy stable formulations based on the two nodal bases for the hexahedron, which are summarized for each element type in Table 1. We refer to the formulation using SEM hexahedra as the “SEM formulation”, and the formulation using Gauss-Legendre hexahedra as the “GL formulation”. We will compare the accuracy and efficiency of each formulation in Section 4.

2.3. Element-specific operations

Our numerical implementation tailors both the local variational formulation and computational operations to each specific type of element. For each element type, we construct three kernels

1. **Volume kernel:** compute the volume contribution (integrals over the interior of the element K).
2. **Surface kernel:** compute the surface contribution (integrals over the surface of the element ∂K).
3. **Update kernel:** apply the inverse of the mass matrix where necessary, execute a step of Adams-Bashforth, and evaluate/store the solution at cubature points on the surface.

Algorithms describing the implementation of relevant kernels for each element type are given in the following sections. For the wedge, pyramid, and tetrahedron, surface and update kernels are very similar to those given in [11] for evaluation of the solution at surface cubature points. Their implementation differs only by the specific surface cubature, which is constructed by combining appropriate cubatures for triangular and quadrilateral faces. For triangular faces, we use quadrature rules for the triangle computed by Xiao and Gimbutas [49] that are exact for polynomials of degree $2N$. For quadrilateral faces, we use a tensor product Gauss-Legendre or Gauss-Legendre-Lobatto (SEM) cubature rule with $(N+1)^2$ points, depending on which hexahedral nodal basis is chosen.

2.3.1. Hexahedral kernels

The volume contribution for the hexahedron may be computed efficiently by exploiting both the Lagrange property of the nodal basis and the tensor-product form of the basis functions. Basis functions for the hexahedron are constructed

$$\phi_{ijk}(r, s, t) = \ell_i(r)\ell_j(s)\ell_k(t),$$

where $\ell_i(r)$ is the Lagrange polynomial at the i th node in the r direction, and similarly for $\ell_j(s), \ell_k(t)$. This implies that, for trilinear J and polynomial u of degree $\leq N$, the evaluation of volume integrals using quadrature reduces to

$$\int_{\hat{K}} \phi_{lmn} \frac{\partial u}{\partial x} J = \sum_{l', m', n'=1}^{N+1} w_{l', m', n'} J_{l', m', n'} \left. \frac{\partial u}{\partial x} \right|_{l', m', n'} \phi_{lmn}|_{l', m', n'} = w_{lmn} J_{lmn} \left. \frac{\partial u}{\partial x} \right|_{l, m, n}$$

due to the Lagrange property of ϕ_{lmn} at quadrature nodes. The values of derivatives $\frac{\partial u}{\partial x}$ require computation of $\frac{\partial u}{\partial r}$, $\frac{\partial u}{\partial s}$, $\frac{\partial u}{\partial t}$ at nodal points, which reduces to the computation of 1D derivatives

$$\left. \frac{\partial u}{\partial r} \right|_{lmn} = \sum_{i=1}^{N+1} \frac{\partial \ell_i(r_l)}{\partial r} \sum_{j=1}^{N+1} \ell_j(s_m) \sum_{k=1}^{N+1} u_{ijk} \ell_k(t_n) = \sum_{i=1}^{N+1} \frac{\partial \ell_i(r_l)}{\partial r} u_{imn} = \sum_{i=1}^{N+1} D_{li}^{1D} u_{imn}. \quad (3)$$

The integral is then scaled by $M_{lmn,lmn}^{-1} = (w_{lmn} J_{lmn})^{-1}$, removing the weight and geometric factor multiplying the derivative. This implementation is described in detail in Algorithm 1.

Algorithm 1 Algorithm for the hexahedron volume kernel (both SEM and GL).

- 1: **procedure** HEXAHEDRON VOLUME KERNEL
 - 2: Load 1D derivative operator D_{ij}^{1D} into shared memory.
 - 3: **for** each element **do**
 - 4: **for** each volume node **do**
 - 5: Load nodal values into shared memory.
 - 6: Prefetch geometric factors into thread local memory.
 - 7: Synchronize threads.
 - 8: Apply D^{1D} to each slice of u_{ijk} to compute reference derivatives (3).
 - 9: Scale by geometric factors to compute physical derivatives.
 - 10: Write the RHS contribution of node $\mathbf{x}_{ijk}$ to global memory.
-

The computation of numerical fluxes requires the evaluation of the solution on the surface of a hexahedron. This is done in the update kernel for the hexahedron, and is described in Algorithm 2. For a 1D nodal basis at GL points, we may evaluate the solution at endpoints ± 1 via

$$u(-1) = \sum_{k=1}^{N+1} \ell_k(-1) u_k, \quad u(1) = \sum_{k=1}^{N+1} \ell_k(1) u_k = \sum_{k=1}^{N+1} \ell_{N-k}(-1) u_k, \quad (4)$$

where we have used symmetry of the GL points across 0. We store values of $\ell_k(-1)$ in an array $V_k^f = \ell_k(-1)$. For a SEM basis, $\ell_k(-1) = \delta_{k1}$, and the evaluation of the basis on the boundary reduces to the loading of the proper nodal degrees of freedom on each face.

Algorithm 2 Algorithm for computing surface values for the GL hexahedron (used in the update kernel).

- 1: **procedure** COMPUTE GL SURFACE VALUES
 - 2: Load array V^f into shared memory.
 - 3: **for** each element **do**
 - 4: **for** each volume node **do**
 - 5: Load nodal values into shared memory.
 - 6: Synchronize threads.
 - 7: **for** each surface node **do**
 - 8: Determine the line of volume nodes required to compute a face node.
 - 9: **for** each line volume node **do**
 - 10: Accumulate contributions to face values (4) into local memory.
 - 11: Store the face nodal values into global memory.
-

The computation of the hex surface contribution differs slightly between SEM and Gauss-Legendre nodal bases. We compute contributions face by face, mapping surface integrals to the reference quadrilateral. For example, on the face corresponding to $t = -1$, this gives

$$\int_r \int_s \ell_{ijk} u J^s = \int_r \int_s \ell_i(r) \ell_j(s) \ell_k(-1) u J^s|_{r,s,-1} = w_i w_j \ell_k(-1) u J^s|_{r_i,s_j,-1} = w_i w_j V_k^f u J^s|_{r_i,s_j,-1}. \quad (5)$$

where w_i, w_j are 1D quadrature weights and J^s is the determinant of the Jacobian mapping from the physical to reference quadrilateral face. For a Gauss-Legendre basis, $\ell_k(-1)$ must be evaluated explicitly. For the SEM formulation, the evaluation of $\ell_k(\pm 1)$ may be skipped, and further optimizations may be done by noting that by the inverse mass matrix may be premultiplied into the geometric factors, and that due to the Lagrange property of nodal polynomials, surface RHS contributions for interior nodes are zero. This quantity is then scaled by the inverse mass matrix to compute the surface contribution to the RHS, and the full procedure is outlined in Algorithm 3.

Algorithm 3 Algorithm for the hexahedron GL surface kernel.

```

1: procedure HEXAHEDRON SURFACE KERNEL
2:   Load array  $V^f$  into shared memory.
3:   for each element do
4:     for each surface node do
5:       Load solution surface values from global memory.
6:       Compute numerical flux using solution surface values and geometric factors.
7:       Store numerical flux values in shared memory.
8:     for each volume node do
9:       For faces  $r, s, t \pm 1$ , compute the surface RHS contribution using (5).
10:      Scale the surface RHS contribution by diagonal of  $M^{-1}$ , accumulate into global memory.

```

2.3.2. Tetrahedral kernels

Assuming Lagrange polynomials $\ell_j(r, s, t)$ defined by N_p nodal points (r_i, s_i, t_i) on a tetrahedron, we may define operators which evaluate derivatives at those same nodal points

$$D_{ij}^r = \frac{\partial \ell_j(r_i, s_i, t_i)}{\partial r}, \quad D_{ij}^s = \frac{\partial \ell_j(r_i, s_i, t_i)}{\partial s}, \quad D_{ij}^t = \frac{\partial \ell_j(r_i, s_i, t_i)}{\partial t}.$$

As shown in [22], RHS contributions may be computed using only the above derivative operators

$$S_{ij}^x = \int_K \ell_i \frac{\partial \ell_j}{\partial x}, \quad M^{-1} S^x \mathbf{u} = \frac{\partial u(\mathbf{x})}{\partial x},$$

where $\mathbf{x}, \mathbf{u}$ are vectors of the nodal positions $\mathbf{x}_i$ and nodal degrees of freedom u_j . The implementation of the tetrahedron volume kernel closely follows the strategy in [26].

In the update kernel, we evaluate solution fields at cubature points on the surface. For nodal tetrahedra, this evaluation may also be done face-by-face to reduce cost at high orders, since traces of solution fields depend only on face nodal values.

2.3.3. Wedge elements

Evaluation of the skew-symmetric form involves the computation of two types of volume integrals

$$\int_K \tilde{\phi}_i \frac{\partial u}{\partial x}, \quad \int_K \frac{\partial \tilde{\phi}_i}{\partial x} u, \quad \tilde{\phi}_i = \frac{\phi_i}{\sqrt{J}}.$$

Since derivatives in x (and similarly for y, z) of the LSC-DG basis are given by

$$\frac{\partial \tilde{\phi}}{\partial x} = \frac{\partial (\phi / \sqrt{J})}{\partial x} = \frac{1}{\sqrt{J}} \left(\frac{\partial \phi}{\partial x} - \frac{\phi}{2J} \frac{\partial J}{\partial x} \right),$$

we precompute and store the values of the physical gradient $\nabla_{xyz} J$ at each cubature point. Derivatives of u and the first volume integral $\int_K \tilde{\phi}_i \frac{\partial u}{\partial x}$ reduce to integrals over the reference element

$$u = \sum_{j=1}^{N_p} u_j \frac{\phi_j}{\sqrt{J}}, \quad \frac{\partial u}{\partial x} = \sum_{j=1}^{N_p} u_j \left(\frac{\partial \phi_j}{\partial x} - \frac{\phi_j}{2J} \frac{\partial J}{\partial x} \right), \quad \int_K \tilde{\phi}_i \frac{\partial u}{\partial x} = \int_{\hat{K}} \frac{\phi_i}{\sqrt{J}} \frac{\partial u}{\partial x} J = \int_{\hat{K}} \phi_i \frac{\partial u}{\partial x}.$$

Similarly, the second volume integral $\int_K \frac{\partial \tilde{\phi}_i}{\partial x} u$ reduces to

$$\begin{aligned} \int_K \frac{\partial \tilde{\phi}_i}{\partial x} u &= \int_{\hat{K}} \left(\frac{\partial \phi_i}{\partial x} - \frac{\phi_i}{2J} \frac{\partial J}{\partial x} \right) u = \int_{\hat{K}} \left(\frac{\partial \phi_i}{\partial r} \frac{\partial r}{\partial x} + \frac{\partial \phi_i}{\partial s} \frac{\partial s}{\partial x} + \frac{\partial \phi_i}{\partial t} \frac{\partial t}{\partial x} - \phi_i \frac{1}{2J} \frac{\partial J}{\partial x} \right) u \\ &= \int_{\hat{K}} \left[\frac{\partial \phi_i}{\partial r}, \frac{\partial \phi_i}{\partial s}, \frac{\partial \phi_i}{\partial t}, \phi_i \right] \cdot \left[u \frac{\partial r}{\partial x}, u \frac{\partial s}{\partial x}, u \frac{\partial t}{\partial x}, -u \frac{1}{2J} \frac{\partial J}{\partial x} \right]. \end{aligned}$$

For LSC-DG, we split the evaluation of the skew-symmetric formulation into the computation of the trial integrand (involving the solution) and the test integrand (involving test functions ϕ_i). Intermediate values used in the test integrand are stored in global memory. In the first step, derivatives of the solution are computed, and values of the solution at cubature points are premultiplied by appropriate change of variable factors, as indicated in the computation of the second volume integral above. In the second step, test functions and their derivatives in the reference r, s, t coordinates are evaluated at cubature points and summed up to compute all integrals.

The splitting of the wedge volume kernel into two steps allows us to load geometric factors and derivatives of J only in the first step, while maintaining fast coalesced memory access patterns. Additionally, splitting into two kernels and writing intermediate values to global memory avoids the overuse of shared memory, which can reduce workgroup occupancy.

Finally, since the wedge is a tensor product of triangle and line elements, we may additionally decompose interpolation and derivative operators into 2D triangle and 1D operators. Exploiting the tensor product nature of these operators on the wedge leads to a reduction in the cost of quadrature. We assume that wedge basis functions may be decomposed as $\phi_k^\Delta(r, t) \phi_l^{1D}(s)$, where $\phi_k^\Delta(r, t)$ are orthogonal basis functions over the triangle [38, 27] and $\phi_l^{1D}(s)$ are 1D Legendre polynomials. Our cubature is defined also as a tensor product of a triangle and 1D cubature rule with N_c^{tri} and $(N+1)$ points, respectively. For cubature points (r_i, s_j, t_i) and weights w_{ij} , where i is the index for a triangle cubature point and j is the index for a 1D cubature rule, we precompute the values of $\phi_k(r, t), \phi_l(s)$ (as well as their derivatives)

$$\begin{aligned} V_{ik} &= \phi_k^\Delta(r_i, t_i), & V_{ik}^r &= \frac{\partial \phi_k^\Delta(r_i, t_i)}{\partial r}, & V_{ik}^t &= \frac{\partial \phi_k^\Delta(r_i, t_i)}{\partial t} & 0 \leq i \leq N_c^{\text{tri}}, \\ V_{jl}^{1D} &= \phi_l^{1D}(s_j), & D_{jl}^{1D} &= \frac{\partial \phi_l^{1D}(s_j)}{\partial s} & 0 \leq j \leq N+1. \end{aligned}$$

Then, the solution may be evaluated at cubature points by the following

$$u(r_i, s_j, t_i) = \sum_{k=1}^{N_p^{\text{tri}}} \sum_{l=1}^{N+1} u_{kl} \phi_k^\Delta(r, t) \phi_l^{1D}(s) = \sum_{k=1}^{N_p^{\text{tri}}} \sum_{l=1}^{N+1} u_{kl} V_{ik} V_{jl}^{1D} \quad (6)$$

where u_{ij} are coefficients for the wedge, decomposed into triangle and 1D indices i and j , respectively. The reference derivatives may be computed in a similar fashion.

Algorithm 4 Algorithm for part 1 of the wedge volume kernel.

- 1: **procedure** WEDGE VOLUME KERNEL PART 1
 - 2: Load 1D interpolation/derivative operators V^{1D}, D^{1D} into shared memory.
 - 3: **for** each element **do**
 - 4: **for** each cubature point **do**
 - 5: Load solution coefficients into shared memory.
 - 6: Synchronize threads.
 - 7: Compute (using the tensor product form) the solution and reference coordinate derivatives (6)
 - 8: Write out solution values at cubature points to global memory.
-

Given solution values and derivatives at cubature points, integrals in the weak form are computed through scaling by cubature weights and multiplication by transposes of matrices used in evaluating the solution

$$\begin{aligned} \int_{\hat{K}} \phi_{ij} \frac{\partial u}{\partial x} &= \sum_{k=1}^{N_c^{\text{tri}}} (V_{ki})^T \sum_{l=1}^{N+1} (V^{1D})_{li}^T \frac{\partial u(\mathbf{x}_{kl})}{\partial x} w_{kl}, \\ \int_{\hat{K}} \frac{\partial \phi_{ij}}{\partial x} u &= \sum_{k=1}^{N_c^{\text{tri}}} \sum_{l=1}^{N+1} w_{kl} [D_{ki}^r V_{lj}^{1D}, V_{ki} D_{lj}^{1D}, D_{ki}^t V_{lj}^{1D}, V_{ki} V_{lj}^{1D}] \cdot \left[u \frac{\partial r}{\partial x}, u \frac{\partial s}{\partial x}, u \frac{\partial t}{\partial x}, -u \frac{1}{2J} \frac{\partial J}{\partial x} \right]. \end{aligned} \quad (7)$$

Scaling by the diagonal mass matrix inverse completes the computation of wedge volume contributions.

Algorithm 5 Algorithm for part 2 of the wedge volume kernel.

- 1: **procedure** WEDGE VOLUME KERNEL PART 2
 - 2: Load transposed 1D interpolation/derivative operators $(V^{1D})^T, (D^{1D})^T$ into shared memory.
 - 3: **for** each element **do**
 - 4: Load solution values at cubature points from global memory into shared memory.
 - 5: **for** each basis function **do**
 - 6: Compute integrals via (7).
 - 7: Scale the volume contribution by the diagonal of M^{-1} and write to global memory.
-

The procedure for the wedge volume kernel is summarized in Algorithms 4 and 5. For Algorithm 5, we again reuse loaded triangle operators, which are constant for each 1D summation. Additionally, we may load premultiplied information and execute dot products in the summation using `float4` data and operations, which are pipelined for faster execution on certain GPUs. Compared to a serialized sum-factorization approach, the tensor product sum in Algorithms 4, 5 does not decrease the total number of computations. We found that parallel implementations of sum factorization required larger amounts of either shared or register memory, both of which decrease occupancy. Instead, we exploit the fact that triangle operators are constant for each 1D summation, increasing data reuse.

2.3.4. Pyramidal elements

The form of the pyramid mapping and the semi-nodal pyramid basis lend themselves to simpler evaluation of integrals. We introduce a quadrature-free method of computing volume contributions, which improves upon the quadrature-based method for the pyramid basis described in [11]. Let u be a function in the span of the pyramid basis on the bi-unit reference pyramid $\hat{\mathcal{P}}$; then,

$$\begin{aligned} \int_{\hat{\mathcal{P}}} u(r, s, t) \phi_{lmn}(r, s, t) &= \int_a \frac{\ell_l^n(a)}{\sqrt{w_l^n}} \int_b \frac{\ell_m^n(b)}{\sqrt{w_m^n}} \int_c u(a, b, c) \frac{P_{N-n}^{2n+3}(c)}{\sqrt{C_{N-n}^N}} \left(\frac{1-c}{2} \right)^{2+n} \\ &= \sqrt{\frac{w_l^n w_m^n}{C_{N-n}^N}} \int_c u(a_l^n, b_m^n, c) P_{N-n}^{2n+3}(c) \left(\frac{1-c}{2} \right)^{2+n}. \end{aligned} \quad (8)$$

Since we assume u is in the pyramidal space, it may be represented with degrees of freedom u_{ijk} . Evaluating $\frac{\partial u}{\partial r}$ using the chain rule gives

$$\frac{\partial u}{\partial a} \frac{\partial a}{\partial r} = \left(\frac{2}{1-c} \right) \sum_{k=0}^N \sum_{i=0}^k \sum_{j=0}^k \frac{1}{\sqrt{w_i^k w_j^k}} \frac{\partial \ell_i^k(a)}{\partial a} \ell_j^k(b) \frac{P_{N-k}^{2k+3}(c) \left(\frac{1-c}{2} \right)^k}{\sqrt{C_{N-k}^N}} u_{ijk}.$$

Putting this together with (8), we have

$$\int_{\hat{\mathcal{P}}} \phi_{lmn} \frac{\partial u}{\partial r} = \sum_{k=0}^N \int_c \frac{P_{N-k}^{2k+3} P_{N-n}^{2n+3} \left(\frac{1-c}{2} \right)^{1+k+n}}{\sqrt{C_{N-n}^N C_{N-k}^N}} \sum_{i=0}^k \sum_{j=0}^k \sqrt{\frac{w_l^n w_m^n}{w_i^k w_j^k}} \frac{\partial \ell_i^k(a_l^n)}{\partial a} \ell_j^k(b_m^n) u_{ijk}.$$

Denoting

$$D_{lmn,ijk}^r = \sqrt{\frac{w_l^n w_m^n}{w_i^k w_j^k}} \frac{\partial \ell_i^k(a_l^n)}{\partial a} \ell_j^k(b_m^n) M_{nk}^c, \quad M_{nk}^c = \int_c \frac{P_{N-k}^{2k+3} P_{N-n}^{2n+3} \left(\frac{1-c}{2}\right)^{1+k+n}}{\sqrt{C_{N-n}^N C_{N-k}^N}},$$

we have that D^r behaves as a weak derivative operator over the reference element

$$\int_{\hat{\mathcal{P}}} \phi_{lmn} \frac{\partial u}{\partial r} = \sum_{k=0}^N \sum_{i=0}^k \sum_{j=0}^k D_{lmn,ijk}^r u_{ijk}. \quad (9)$$

We may similarly define D^s such that

$$\int_{\hat{\mathcal{P}}} \phi_{lmn} \frac{\partial u}{\partial s} = D_{lmn,ijk}^s = \sqrt{\frac{w_l^n w_m^n}{w_i^k w_j^k}} \ell_i^k(a_l^n) \frac{\partial \ell_j^k(b_m^n)}{\partial b} M_{nk}^c.$$

To take derivatives with respect to the t coordinate, we use the chain rule for $\frac{\partial u}{\partial t}$

$$\begin{aligned} \frac{\partial u}{\partial t} &= \frac{\partial u}{\partial a} \frac{\partial a}{\partial t} + \frac{\partial u}{\partial b} \frac{\partial b}{\partial t} + \frac{\partial u}{\partial c} = \frac{\partial u}{\partial a} \left(\frac{1+a}{2}\right) \left(\frac{2}{1-c}\right) + \frac{\partial u}{\partial b} \left(\frac{1+b}{2}\right) \left(\frac{2}{1-c}\right) + \frac{\partial u}{\partial c} \\ &= \left(\frac{1+a}{2}\right) \frac{\partial u}{\partial r} + \left(\frac{1+b}{2}\right) \frac{\partial u}{\partial s} + \frac{\partial u}{\partial c} \end{aligned}$$

An operator for the t derivative is constructed similarly. We first define a derivative operator in c

$$D_{lmn,ijk}^c = \sqrt{\frac{w_l^n w_m^n}{w_i^k w_j^k}} \ell_i^k(a_l^n) \ell_j^k(b_m^n) \int_c \frac{\partial \left[P_{N-k}^{2k+3}(c) \left(\frac{1-c}{2}\right)^k \right]}{\partial c} \frac{P_{N-n}^{2n+3}(c) \left(\frac{1-c}{2}\right)^{2+n}}{\sqrt{C_{N-n}^N C_{N-k}^N}}.$$

Using this, D^t may then be defined as

$$D_{lmn,ijk}^t = \left(\frac{1+a_i^k}{2}\right) D_{lmn,ijk}^r + \left(\frac{1+b_j^k}{2}\right) D_{lmn,ijk}^s + D_{lmn,ijk}^c.$$

We now exploit the fact that both geometric factors $\frac{\partial rst}{\partial xyz}$ and the determinant of the Jacobian J are constant in the c (and thus t) direction for all vertex-mapped pyramids [6, 11]. As a result, geometric factors are constant in the index n . This implies that the integral of a derivative in the x coordinate over a physical pyramid may be given as

$$\begin{aligned} \int_{\mathcal{P}} \phi_{lmn} \frac{\partial u}{\partial x} &= \int_{\hat{\mathcal{P}}} \phi_{lmn} \left(\frac{\partial u}{\partial r} \frac{\partial r}{\partial x} + \frac{\partial u}{\partial s} \frac{\partial s}{\partial x} + \frac{\partial u}{\partial t} \frac{\partial t}{\partial x} \right) J \\ &= J_{lm} \left\{ \left(\frac{\partial r}{\partial x} \right)_{lm} \sum_{ijk} D_{lmn,ijk}^r u_{ijk} + \left(\frac{\partial s}{\partial x} \right)_{lm} \sum_{ijk} D_{lmn,ijk}^s u_{ijk} + \left(\frac{\partial t}{\partial x} \right)_{lm} \sum_{ijk} D_{lmn,ijk}^t u_{ijk} \right\}. \end{aligned}$$

The above expression is simply multiplication by (square) semi-nodal operators D^r, D^s, D^t and entry-wise scalings by geometric factors, similarly to nodal methods for hexahedral or tetrahedral elements. Furthermore, the scaling by J_{lm} is removed when applying the inverse of the diagonal mass matrix

$$M_{lmn,lmn}^{-1} = \frac{1}{J_{lm}}.$$

When computing integrals in the skew-symmetric form, we may simply apply the transpose of each operator

$$\int_{\hat{\mathcal{P}}} \frac{\partial \phi_{lmn}}{\partial r} u J = \sum_{ijk} (D^r)_{lmn,ijk}^T J_{ij} u_{ijk}. \quad (10)$$

and similarly for $(D^s)^T, (D^t)^T$. In this case, the factor of J is not cancelled out after multiplying by M^{-1} . The pyramid kernel is described in detail in Algorithm 6. Similarly to the wedge, the algorithm is made up of two parts (computation of the trial and test portion). However, unlike the wedge, the semi-nodal nature of the pyramid basis allows us to store intermediate values efficiently in shared memory.

Algorithm 6 Algorithm for the (skew-symmetric) pyramid volume kernel.

```

1: procedure PYRAMID VOLUME KERNEL
2:   for each element do
3:     Load geometric factors and solution coefficients from global memory.
4:     Premultiply solution coefficients by geometric factors, load into shared memory.
5:     for each basis function do
6:       Compute weak derivatives using (9).
7:       Compute weak derivatives (integrated by parts) using (10).
8:       Scale the volume contribution by the diagonal of  $M^{-1}$  and write to global memory.
```

2.4. Many-core implementation

The low-storage DG method described above has been implemented in **hybridg**, a portable, GPU accelerated solver for the acoustic wave equation on hybrid meshes. Portability between OpenMP, OpenCL and CUDA platforms is achieved through the OCCA programming model [33]. The solution is evolved in time using a multi-rate Adams-Bashforth (MRAB) scheme. Support for multiple GPUs is facilitated through MPI, though load-balancing strategies for scalability remain to be investigated. The kernels for each element type are optimized according to the operations described above. Additional optimizations may be done by tuning the number of elements processed per workgroup (batching), following strategies outlined in [26]. GPUs consist of several “workgroups”, each of which typically contains 32 synchronized threads. In this implementation, one or more elements are assigned to a workgroup, and each thread processes work corresponding to a single degree of freedom. Since the number of degrees of freedom is typically not a perfect multiple of 32, assigning only a single element to a workgroup results in idle threads. Automatically tuning batch sizes for each element brings the number of active threads closer to a multiple of 32, while controlling the increased memory costs and decreased occupancy which may result from processing multiple elements per workgroup. In Section 4.3, we give results where optimal batch sizes are determined experimentally for each element type and approximation order N .

Due to large variations in element sizes, as well as variations in element type-dependent trace constants, local timestep restrictions may vary significantly over the mesh. To sidestep overly restrictive global timesteps, we use multi-rate timestepping in the form of a 3rd order Adams-Bashforth scheme. The local timestep $d\tau_K$ for each element is derived in Section 3, and the global timestep $d\tau_{\min}$ is taken as the minimum over all local timesteps. For the MRAB scheme, N_{levels} timestep levels are determined via

$$d\tau_{\text{lev}} = 2^{N_{\text{levels}} - \text{lev}} d\tau_{\min}, \quad 1 \leq \text{lev} \leq N_{\text{levels}}.$$

Elements are then binned into each timestep level — a given element K is assigned a timestep of $d\tau_{\text{lev}}$ if $d\tau_{\text{lev}-1} \leq d\tau_K \leq d\tau_{\text{lev}}$. A second sweep through the mesh moves elements from coarser timesteps to a finer timesteps in order to guarantee that neighboring elements differ only by one level at most (or that the timesteps of neighboring elements differ only by a factor of 1 or 2). The solution is then updated according to the coarsest timestep, and each timestep level is evolved $2^{\text{lev}-1}$ times for every coarse timestep.

Since the details of the implementation are independent of element type, we refer the reader to [21, 17, 34] for a description of the multi-rate scheme on triangular and tetrahedral meshes.

3. Timestep restrictions

Stable local and global timesteps are necessary for multi-rate schemes. We derive here timestep restrictions for hybrid meshes by bounding the DG operator norm. These bounds are given in terms of order-dependent constants and physical/geometric quantities for each element type.

To simplify notation, we define group trial and test variables U, V

$$U = \begin{pmatrix} p \\ \mathbf{u} \end{pmatrix}, \quad V = \begin{pmatrix} v \\ \boldsymbol{\tau} \end{pmatrix}.$$

Under time-explicit DG methods, the discretized acoustic wave equation results in a system of ODEs

$$\frac{dU}{d\tau} = M^{-1}A(U), \quad V^T M U = \sum_{K \in \Omega_h} \left(\int_K \frac{1}{\kappa} p v + \int_K \rho \mathbf{u} \boldsymbol{\tau} \right)$$

where M is the scaled mass matrix. The global matrix A is given as the sum of local matrices $A = \sum_{K \in \Omega_h} A_K$, where A_K is given by the local spatial discretization over an element $K \in \Omega_h$

$$\begin{aligned} V^T A_K U &= \int_K \mathbf{u} \cdot \nabla \phi^- dx - \int_K \nabla p \cdot \boldsymbol{\psi}^- dx \\ &+ \int_{\partial K} \left(\frac{1}{2} \tau_p \llbracket p \rrbracket - \mathbf{n}^- \cdot \{ \{ \mathbf{u} \} \} \right) \phi^- dx + \int_{\partial K} \mathbf{n}^- \frac{1}{2} (\tau_u \llbracket \mathbf{u} \rrbracket \cdot \mathbf{n}^- - \llbracket p \rrbracket) \boldsymbol{\psi}^- dx. \end{aligned} \quad (11)$$

We are interested in deriving bounds on the real and imaginary parts of the spectra of the DG operator in order to accurately determine a stable timestep restriction. For example, we may bound the timestep by $d\tau < 1/\rho(M^{-1}A)$, the inverse of the spectral radius of $M^{-1}A$.

Previous work has been done in estimating stable timesteps for structured grids [31, 42], and explicit information about the spectra of the discrete DG operator has been derived by Krivodonova and Qin in 1D [29, 30]. We use an approach similar to Cohen, Ferrieres, and Pernet in [13] to derive bounds for $\rho(M^{-1}A)$ that depend on explicit expressions for the constants in discrete trace and Markov inequalities. The first bounds the L^2 norm of a function on the boundary ∂K of an element K by the L^2 norm of the function over the element

$$\|u\|_{L^2(\partial K)}^2 \leq C_T(N, K) \|u\|_{L^2(K)}^2$$

while the latter bounds the L^2 norm of the gradient of a function by the L^2 norm of the function in the interior

$$\|\nabla u\|_{L^2(K)}^2 \leq C_M(N, K) \|u\|_{L^2(K)}^2.$$

For SEM, the L^2 norm is replaced by the equivalent discrete norm computed using GLL quadrature. The mesh and order-dependent constants $C_T(N, K), C_M(N, K)$ determine the timestep restriction required for stability. Since we wish to take the timestep as large as possible, we require sharp expressions for these constants in order to accurately estimate $\rho(M^{-1}A)$.

We note that the bounds we derive assume polynomial basis functions. This assumption is violated for non-affine mappings of the wedge, where we use rational LSC-DG basis functions. However, for such bases, we may use weighted Markov and trace inequalities from [45] in place of standard polynomial inequalities.

3.1. Bounds on the spectra

We derive here bounds on the spectra of the DG operator based on the constants in trace and Markov inequalities over reference elements. These bounds apply for both the GL formulation with full mass matrix M (or true L^2 inner product) and the SEM lumped mass matrix M_{SEM} (or discrete inner product based on GLL quadrature); the difference between the two is reflected in the constants in each inequality.

The eigenvalues λ and corresponding eigenvectors v of $M^{-1}A$ are given by the generalized eigenvalue problem $Av = \lambda Mv$. We may derive a more detailed bound on the spectra of $M^{-1}A$ by decomposing the matrix into symmetric and skew-symmetric parts

$$A = (A^s + A^k), \quad A^s = \frac{1}{2}(A + A^T), \quad A^k = \frac{1}{2}(A - A^T).$$

From [1], we have that

$$\begin{aligned}\lambda_{\min}(M^{-1}A^s) &\leq \operatorname{Re}(\lambda(M^{-1}A)) \leq \lambda_{\max}(M^{-1}A^s) \\ \lambda_{\min}(M^{-1}A^k) &\leq \operatorname{Im}(\lambda(M^{-1}A)) \leq \lambda_{\max}(M^{-1}A^k),\end{aligned}$$

In other words, the magnitude of the real and imaginary parts of the spectra are bounded by $\rho(M^{-1}A^s)$ and $\rho(M^{-1}A^k)$. Since A^s is symmetric and real, its spectral radius is given by the generalized Rayleigh quotient

$$\rho(M^{-1}A^s) = \max \left| \frac{U^T A^s U}{U^T M U} \right|.$$

From the definition of A_K in (11), we can show that

$$|U^T A^s U| = \left| -\frac{1}{2} \sum_{f \in \Gamma_h} \int_f \left\{ \tau_p \llbracket p \rrbracket^2 + \tau_u \llbracket \mathbf{u}_n \rrbracket^2 \right\} \right| \leq \frac{1}{2} \sum_{f \in \Gamma_h} \left\{ \tau_p \|\llbracket p \rrbracket\|_{L^2(f)}^2 + \tau_u \|\llbracket \mathbf{u}_n \rrbracket\|_{L^2(f)}^2 \right\}, \quad (12)$$

where we have used the normal jump $\llbracket \mathbf{u}_n \rrbracket = \mathbf{u}^+ \mathbf{n}^+ + \mathbf{u}^- \mathbf{n}^-$. We may bound the norm of the jumps

$$\frac{1}{2} \|\llbracket p \rrbracket\|_{L^2(f)}^2 \leq \|p^+\|_{L^2(f)}^2 + \|p^-\|_{L^2(f)}^2, \quad \frac{1}{2} \|\llbracket \mathbf{u}_n \rrbracket\|_{L^2(f)}^2 \leq \|\mathbf{u}^+ \mathbf{n}^+\|_{L^2(f)}^2 + \|\mathbf{u}^- \mathbf{n}^-\|_{L^2(f)}^2,$$

and bound the sum of jumps and averages over faces by boundary values

$$\sum_{f \in \Gamma_h} \frac{1}{2} \|\llbracket p \rrbracket\|_{L^2(f)}^2 \leq \sum_{K \in \Omega_h} \|p\|_{L^2(\partial K)}^2, \quad \sum_{f \in \Gamma_h} \|\llbracket \mathbf{u}_n \rrbracket\|_{L^2(f)}^2 \leq \sum_{K \in \Omega_h} \|\mathbf{u}\|_{L^2(\partial K)}^2.$$

Combined with (12) and changing to a sum over elements, we can apply trace inequalities to arrive at

$$\begin{aligned}|U^T A^s U| &\leq \sum_{K \in \Omega_h} \left\{ \tau_{p,K} \|p\|_{L^2(\partial K)}^2 + \tau_{u,K} \|\mathbf{u} \cdot \mathbf{n}\|_{L^2(\partial K)}^2 \right\} \\ &\leq \sum_{K \in \Omega_h} C_T(N, K) \left\{ \tau_{p,K} \|p\|_{L^2(K)}^2 + \tau_{u,K} \|\mathbf{u}\|_{L^2(K)}^2 \right\},\end{aligned}$$

where $C_T(N, K)$ is the constant in the trace inequality over the boundary ∂K of the element K and $\tau_{p,K}, \tau_{u,K}$ are the maximum values of the penalties over the faces of an element. Combining these bounds, we have

$$\begin{aligned}\max \left| \frac{U^T A^s U}{U^T M U} \right| &\leq \max \frac{\sum_{K \in \Omega_h} C_T(N, K) \left\{ \tau_{p,K} \|p\|_{L^2(K)}^2 + \tau_{u,K} \|\mathbf{u}\|_{L^2(K)}^2 \right\}}{\sum_{K \in \Omega_h} \left\{ \frac{1}{\kappa_K} \|p\|_{L^2(K)}^2 + \rho_K \|\mathbf{u}\|_{L^2(K)}^2 \right\}} \\ &\leq \max_K C_T(N, K) \max \left(\tau_{p,K} \kappa_K, \frac{\tau_{u,K}}{\rho_K} \right)\end{aligned}$$

The spectral radius of $M^{-1}A^k$ may be bounded by noting that skew-symmetric matrices are normal, which implies that $\rho(M^{-1}A^k) = \|M^{-1}A^k\|$. A characterization of this norm is given as follows: for a skew-symmetric matrix $Z \in \mathcal{R}^{n \times n}$,

$$\|Z\| = \max_{f, g \in \mathcal{R}^n} \frac{2g^T Z f}{\|f\|^2 + \|g\|^2}.$$

This characterization has been proven in other settings (see, for example, [12]). A straightforward modification of this theorem to use the norm $\|f\|_M = \sqrt{f^T M f}$ gives that

$$\|M^{-1}A^k\| = \max_{U, V \in \mathcal{R}^n} \frac{2V^T A^k U}{\|U\|_M^2 + \|V\|_M^2}$$

where the denominator $\|U\|_M^2 + \|V\|_M^2$ is

$$\|U\|_M^2 + \|V\|_M^2 = \sum_{K \in \Omega_h} \left\{ \frac{1}{\kappa_K} \|p\|_{L^2(K)}^2 + \rho_K \|\mathbf{u}\|_{L^2(K)}^2 + \frac{1}{\kappa_K} \|v\|_{L^2(K)}^2 + \rho_K \|\boldsymbol{\tau}\|_{L^2(K)}^2 \right\}$$

From (11), we may compute the skew-symmetric form

$$V^T A^k U = \sum_{K \in \Omega_h} \left\{ \int_K \mathbf{u} \cdot \nabla v - \int_K \nabla p \cdot \boldsymbol{\tau} \right\} + \sum_{f \in \Gamma_h} \int_f (-\{\{\mathbf{u}\}\} \llbracket v_n \rrbracket + \{\{\boldsymbol{\tau}\}\} \llbracket p_n \rrbracket)$$

where we have used the vector-valued jumps $\llbracket p_n \rrbracket = p^+ \mathbf{n}^+ + p^- \mathbf{n}^-$ and $\llbracket v_n \rrbracket = v^+ \mathbf{n}^+ + v^- \mathbf{n}^-$.

We may bound the volume terms using Markov and Young's inequality

$$\begin{aligned} \sum_{K \in \Omega_h} \left| \int_K \mathbf{u} \cdot \nabla v - \int_K \nabla p \cdot \boldsymbol{\tau} \right| &\leq \sum_{K \in \Omega_h} \|\mathbf{u}\|_{L^2(K)} \|\nabla v\|_{L^2(K)} + \|\nabla p\|_{L^2(K)} \|\boldsymbol{\tau}\|_{L^2(K)} \\ &\leq \sum_{K \in \Omega_h} \frac{1}{2} \sqrt{C_M(N, K)} \left(\|p\|_{L^2(K)}^2 + \|\mathbf{u}\|_{L^2(K)}^2 + \|v\|_{L^2(K)}^2 + \|\boldsymbol{\tau}\|_{L^2(K)}^2 \right) \\ &\leq \frac{1}{2} \sqrt{C_M(N, K)} \max \left(\kappa_K, \frac{1}{\rho_K} \right) \left(\|U\|_M^2 + \|V\|_M^2 \right) \end{aligned}$$

where $C_M(N, K)$ is the constant in the Markov inequality for the element K .

The surface terms may also be bounded using Young's inequality

$$\begin{aligned} \left| \sum_{f \in \Gamma_h} \int_f (-\{\{\mathbf{u}\}\} \llbracket v_n \rrbracket + \{\{\boldsymbol{\tau}\}\} \llbracket p_n \rrbracket) \right| &\leq \sum_{f \in \Gamma_h} \left(\|\{\{\mathbf{u}\}\}\|_{L^2(f)} \|\llbracket v_n \rrbracket\|_{L^2(f)} + \|\{\{\boldsymbol{\tau}\}\}\|_{L^2(f)} \|\llbracket p_n \rrbracket\|_{L^2(f)} \right) \\ &\leq \frac{1}{2} \sum_{K \in \Omega_h} \left(\|p\|_{L^2(\partial K)}^2 + \|\mathbf{u}\|_{L^2(\partial K)}^2 + \|v\|_{L^2(\partial K)}^2 + \|\boldsymbol{\tau}\|_{L^2(\partial K)}^2 \right) \\ &\leq \frac{1}{2} \sum_{K \in \Omega_h} C_T(N, K) \left(\|p\|_{L^2(K)}^2 + \|\mathbf{u}\|_{L^2(K)}^2 + \|v\|_{L^2(K)}^2 + \|\boldsymbol{\tau}\|_{L^2(K)}^2 \right) \\ &\leq \frac{1}{2} \max_K C_T(N, K) \max \left(\kappa_K, \frac{1}{\rho_K} \right) \left(\|U\|_M^2 + \|V\|_M^2 \right) \end{aligned}$$

Combining these two bounds gives

$$\|M^{-1} A^k\| = \max_{U, V \in \mathcal{R}^n} \frac{2V^T A^k U}{\|U\|_M^2 + \|V\|_M^2} \leq \max_K \max \left(\kappa_K, \frac{1}{\rho_K} \right) \left(\sqrt{C_M(N, K)} + C_T(N, K) \right).$$

3.2. Constants in trace inequalities

An explicit expression for the trace inequality constant for a general d -simplex was given by Warburton and Hesthaven in [47], and was extended by Hillewaert in his thesis to hexahedra and wedges in [23]. While the trace constant for the pyramid does not appear to be available in an explicit form, Hillewaert also proposed an empirical fit of the trace inequality constant for pyramids, which was based on a lower bound for the trace constant that could be derived in closed form. An improved asymptotically optimal constant was given by Chan and Warburton in [10]. The trace inequality is as follows: for u in the approximation space, the norm of u on a face $\hat{f}$ of the reference element $\hat{K}$,

$$\|u\|_{L^2(\hat{f})}^2 \leq C_f(N) \|u\|_{L^2(\hat{K})}^2.$$

These constants were derived assuming triangular faces with area 2 and quadrilateral faces of area 4 on bi-unit reference elements.

Element type	Triangular face	Quadrilateral face	$C_T(N)$
Hexahedra		$(N+1)^2/2$	$4(N+1)^2$
Wedge	$(N+1)^2/2$	$(N+1)(N+2)$	$(\sqrt{2}+3)(N+1)(N+2)$
Pyramid	$(N+1)(N+2)/2$	$(N+1)(N+3)$	$(\sqrt{2}+2)(N+1)(N+3)$
Tetrahedra	$(N+1)(N+3)/2$		$(\sqrt{3}+3)(N+1)(N+3)/2$

Table 2: Summary of the trace constant $C_f(N)$ in the discrete trace inequalities for different faces of various reference elements, as well as the analytic trace constant $C_T(N) = \max_f (C_f(N)/|f|) |\partial\hat{K}|$ over the surface $\partial\hat{K}$.

Figure 4: Comparison of face-based bound $C_T(N)$ to numerically computed trace inequality constants on each reference element.

The constants in the trace inequality over the surface $\partial\hat{K}$ of a reference element may be given in terms of the constants over the different types of faces. Assume the trace inequality is derived for face f . Then, if we have N_f faces of a reference element $\hat{K}$,

$$\|u\|_{L^2(\partial\hat{K})}^2 = \sum_{j=1}^{N_f} \|u\|_{L^2(f_j)}^2 = \sum_{j=1}^{N_f} \frac{|f_j|}{|f|} \|u\|_{L^2(f)}^2 \leq \sum_{j=1}^{N_f} \frac{C_f(N)}{|f|} |f_j| \|u\|_{\hat{K}}^2 \leq \max \frac{C_f(N)}{|f|} |\partial\hat{K}| \|u\|_{\hat{K}}^2.$$

where we have used Hölder's inequality with $p = 1$ and $q = \infty$ and the fact that, under an affine transform (i.e. rotation), $\|u\|_{f_j}^2 = |f_j|/|f| \|u\|_f^2$. We will refer to $C_T(N)$ as the constant in the trace inequality over the full surface ∂K , such that

$$C_T(N) = \max_f \frac{C_f(N)}{|f|} |\partial\hat{K}|$$

We summarize these trace constants $C_f(N)$ and $C_T(N)$ in Table 2. We remark that the constant $C_T(N)$ for tetrahedra may be equivalently derived using the trace inequality from [47]

$$\|u\|_{\partial K}^2 \leq \frac{(N+1)(N+3)}{3} \frac{|\partial K|}{|K|} \|u\|_K$$

when the element K is taken to be the bi-unit right tetrahedron $\hat{K}$.

The constants in Table 2 have the benefit of being fully explicit in N . However, since the values of $C_T(N)$ are based on face-by-face estimates (instead of considering the whole surface ∂K), surface trace inequalities under these constants are not tight. Figure 4 compares numerically computed trace inequality constants to the derived analytic constants $C_T(N)$ in Table 2. For each element up to $N \leq 7$, the estimated analytic constant is a factor of 2-3 larger than the computed constant.

For the reference wedge, pyramid, and tetrahedra, we have not found explicit expressions for numerically computed constants in surface trace inequalities. However, for the hexahedron, we are able to derive that the constant in the bounds

$$\|u\|_{L^2(\partial\hat{K})}^2 \leq C_T(N) \|u\|_{L^2(\hat{K})}^2, \quad \|u\|_{\text{SEM}(\partial\hat{K})}^2 \leq C_{\text{SEM}}(N) \|u\|_{\text{SEM}(\hat{K})}^2$$

are given exactly by

$$C_T(N) = \frac{3(N+1)(N+2)}{2}, \quad C_{\text{SEM}}(N) = \frac{3N(N+1)}{2}$$

where the bound for the SEM inequality is valid for $N \geq 1$. These constants are a refinement of explicit bounds given by Ern and Burman [9] and Evans and Hughes [15], and proofs are included in Appendix A.1.

Under a mapping from reference element $\hat{K}$ to physical element K , we may derive a bound in terms of the determinant of the Jacobian J and the determinant of the surface Jacobian J^s

$$\|u\|_{L^2(\partial K)}^2 \leq C_T(N) \|J^s\|_{L^\infty(\partial \hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})} \|u\|_{L^2(K)}^2. \quad (13)$$

If the LSC-DG basis is used for the wedge, the polynomial trace inequalities used here may be replaced by trace inequalities for weighted polynomial spaces [46], resulting in a bound of the form

$$\|u\|_{L^2(\partial K)}^2 \leq C_T(N) \left\| \frac{J^s}{J} \right\|_{L^\infty(\partial \hat{K})} \|u\|_{L^2(K)}^2.$$

For affine mappings of faces where J^s, J are constant, the ratio of their norms reduces to the ratio of the physical surface area to reference surface area divided by the ratio of physical element volume to reference element volume

$$\|J^s\|_{L^\infty(\partial \hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})} = \frac{|\partial K| / |\partial \hat{K}|}{|K| / |\hat{K}|}.$$

3.3. Constants in Markov inequalities

Expressions for the constants in Markov inequalities over simplices are given in [36]; however, these bounds have not yet been extended to hexahedra, tetrahedra, and pyramids. We take a more heuristic approach here and compute numerically the constant in Markov inequalities over the reference element. We may bound

$$\|\nabla u\|_{L^2(K)}^2 = \sum_{k=1}^3 \left\| \frac{\partial u}{\partial x_k} \right\|_{L^2(K)}^2 \leq C_{rst}^2 \|J\|_{L^\infty(\hat{K})} \|\nabla_{rst} u\|_{L^2(\hat{K})}^2,$$

where $C_{rst} = \max_{rst,xyz} \left\| \frac{\partial r,s,t}{\partial x,y,z} \right\|_{L^\infty(\hat{K})}$ is the max norm of the Jacobian matrix for the element K . The constants in the reference Markov inequality

$$\|\nabla_{rst} u\|_{L^2(\hat{K})}^2 \leq C_M(N) \|u\|_{L^2(\hat{K})}^2$$

may then be computed numerically over the reference element through the solution of an eigenvalue problem. The bound is completed by bounding reference quantities by physical quantities using J^{-1}

$$\|\nabla u\|_{L^2(K)}^2 \leq C_M(N) C_{rst}^2 \|J\|_{L^\infty(\hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})} \|u\|_{L^2(K)}^2.$$

3.4. N -dependent bounds for $\rho(M^{-1}A)$

To summarize, the trace constant over a mapped element K is given by

$$C_T(N, K) = C_T(N) \|J^s\|_{L^\infty(\partial \hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})},$$

while the Markov constant $C_M(N, K)$ is given by

$$C_M(N, K) = C_M(N) C_{rst}^2 \|J\|_{L^\infty(\hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})}.$$

The real part of the spectra of $M^{-1}A$ is bounded by the trace inequality constant and physical parameters

$$|\operatorname{Re}(\lambda)| \leq \max_K \max \left(\tau_{p,K} \kappa_K, \frac{\tau_{u,K}}{\rho_K} \right) C_T(N) \|J^s\|_{L^\infty(\partial\hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})}. \quad (14)$$

Likewise, the imaginary part of the spectra of $M^{-1}A$ depends on the quantity

$$|\operatorname{Im}(\lambda)| \leq \max_K \max \left(\kappa_K, \frac{1}{\rho_K} \right) \left(\sqrt{C_M(N, K)} + C_T(N, K) \right). \quad (15)$$

The latter term expands to

$$C_{rst} \sqrt{C_M(N) \|J\|_{L^\infty(\hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})} + C_T(N) \|J^s\|_{L^\infty(\partial\hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})}.$$

We analyze the bounds derived above for the acoustic wave equation with $\kappa = \rho = 1$ and compare them to numerically computed spectra of $M^{-1}A$. Figure 5 shows the computed spectra of $M^{-1}A$ for a hybrid mesh using the GL formulation. Bounds on the real and imaginary parts of the spectra derived by computing $\rho(M^{-1}A^s)$, $\rho(M^{-1}A^k)$ are also included as dotted lines.

Figure 5: Spectra for a hybrid mesh using Gauss-Legendre quadrature, along with bounds (dotted lines) computed from $\rho(A^s)$ and $\rho(A^k)$, and analytic bounds (squared lines). The bounds for $N = 2, 3$ are truncated for visualization purposes.

We note that our estimates for $\rho(M^{-1}A^k)$ provide a loose bound on the imaginary part of the spectra. However, we observe that the spectral radius of the symmetric part $\rho(M^{-1}A^s)$ provides a relatively tight bound on $\rho(M^{-1}A)$. Motivated by this observation, we compare $\rho(M^{-1}A)$ to four different bounds based on $\rho(M^{-1}A^s)$:

1. The true spectral radius of $M^{-1}A^s$.
2. The bound (14), where the trace inequality constant $C_T(N, K)$ is computed numerically as the maximum trace constant over each mapped element K .
3. The bound (14), where the trace inequality constant $C_T(N)$ is computed numerically over the reference element $\hat{K}$ (Appendix B).
4. The bound (14), where the trace inequality constant $C_T(N)$ is given by the estimate (13).

In the latter two bounds, the L^∞ norm of J, J^{-1} , and J^s is approximated by taking the maximum value over quadrature points. Table 3 shows each bound compared to $\rho(M^{-1}A)$ for various N . The bound on $\rho(M^{-1}A)$ using computed trace inequality constants over the reference element is less than a factor of 2 away from the true spectral radius, and we use this to estimate the largest stable timestep for a given mesh and discretization. The bound on $\rho(M^{-1}A)$ using analytically derived trace constants over faces provides a looser bound (factor of 4-5 away from the true spectral radius), but has the advantage of being fully explicit in N .

N	1	2	3	4
$\rho(M^{-1}A)$	10.89	18.66	28.18	41.09
$\rho(M^{-1}A^s)$	11.26	19.30	29.05	42.86
Computed $C_T(N, K)$	12.22	20.84	32.76	47.47
Computed $C_T(N)$	17.24	29.47	46.32	67.13
Analytic $C_T(N)$	38.53	74.73	124.54	186.81

Table 3: Spectral radius of $\rho(M^{-1}A)$ and various bounds for a hybrid mesh and $N = 1, \dots, 4$.

Bounds for the spectral radius under the SEM formulation may be derived by substituting in trace inequality constants using SEM quadrature over quadrilateral faces. Computed trace inequality constants $C_T(N)$ for both the GL and SEM formulations are given in Appendix B. Analytic expressions for constants in trace inequalities may also be derived using norm equivalences between SEM and L^2 norms for polynomials of order N .

Numerical experiments confirm that $\rho(M^{-1}A)$ is bounded by the maximum trace inequality constant $\max_K C_T(N, K)$. Since a stable timestep is given by $d\tau \leq 1/\rho(M^{-1}A)$, we estimate the global timestep by

$$d\tau = \frac{C}{\max_K (C_{\rho, \kappa}^K C_T(N, K))} = \frac{C}{\max_K (C_{\rho, \kappa}^K C_T(N) C_J^K)},$$

where $C_{\rho, \kappa}^K = \max\left(\tau_{p, K} \kappa_K, \frac{\tau_{u, K}}{\rho_K}\right)$, C is a tunable global CFL constant, and C_J^K is

$$\begin{aligned} C_J^K &= \|J^s\|_{L^\infty(\partial\hat{K})} \|J^{-1}\|_{L^\infty(\hat{K})} \quad \text{if the element is a hex, pyramid, or tet, or} \\ C_J^K &= \left\| \frac{J^s}{J} \right\|_{L^\infty(\partial\hat{K})} \quad \text{if the element is an LSC-DG wedge.} \end{aligned}$$

Since we have estimated the global timestep by taking the maximum trace constant over all elements, we may also use local trace constants to estimate stable local timesteps $d\tau_K$ using the constants in local trace inequalities

$$d\tau_K = \frac{C}{C_{\rho, \kappa}^K C_T(N) C_J^K}. \quad (16)$$

4. Numerical experiments

In this section, we present results from **hybridg** for meshes of individual element types and hybrid meshes at various orders of approximation.

4.1. Verification: individual element types

We begin by comparing the SEM and GL formulations for a series of uniformly refined meshes. Hex meshes are constructed by subdividing the unit cube into hexahedra of size h , and we construct wedge and pyramid meshes by further subdividing each hexahedra into either 2 wedges or 6 pyramids, as shown in Figure 6. Figure 7 shows numerical convergence rates obtained using the resonant cavity solution

$$p(x, y, z, \tau) = \sin(\pi x) \sin(\pi y) \sin(\pi z) \cos(\sqrt{3}\pi\tau)$$

over the unit cube $[0, 1]^3$. Since the solution is smooth, the best approximation in L^2 converges with a rate of h^{N+1} , though DG is guaranteed only a convergence rate of $h^{N+1/2}$ for general meshes. We observe optimal h^{N+1} convergence rates for the GL formulation, while the SEM formulation yields computed convergence rates somewhere in between $h^{N+1/2}$ and h^{N+1} (though for the hex, these rates are less informative since SEM

Figure 6: Subdivisions of hexahedra used to construct uniform meshes of wedges and pyramids.

Figure 7: Convergence rates for both the SEM and GL formulations for meshes of hexes, wedges, and pyramids.

errors do not follow a constant rate as closely). For all orders and meshes, the error for the GL formulation is lower than that of the SEM formulation, though this difference is less pronounced at higher N .

We also compare standard wedges to LSC-DG wedges in Figure 8 by performing convergence tests on two sequences of meshes. For one sequence, we refine a uniform mesh of affine wedges with the GL formulation (referred to as “GL Wedge”). For the other, we begin with a mesh of randomly warped wedges, and refine the mesh by bisection to produce a sequence of meshes (referred to as “GL LSC wedge”). Optimal rates of convergence are observed for both sequences (for the SEM formulation, similar behavior to standard SEM wedges is also observed when using LSC SEM wedges). We note that the observed optimal convergence rates for LSC-DG are likely dependent on the fact that refinement by bisection produces asymptotically affine elements. If we perturb vertex positions at each level of mesh refinement, we do not in general observe optimal convergence rates. This issue is observed also for standard mapped approximation spaces [8], though it is likely exacerbated by the LSC-DG approximation.

4.2. Verification: hybrid cube meshes

We check also numerical convergence rates for unstructured hybrid meshes. L^2 errors are computed on a sequence of hybrid meshes containing elements of all types. Each element is refined by a self-similar splitting, except for pyramids, which are refined into both pyramids and tetrahedra.¹ These meshes contain 201, 1704, 14016, and 113664 elements, respectively. The breakdown of the number of each individual elements of each type is given in Table 4. 5 MRAB levels are used, with a CFL constant of 0.5. Results are computed using double precision; single precision behaves similarly, though the error stalls at around 5×10^{-7} .

¹Each element is refined according to the pattern specified by the “refine by splitting” option in GMSH [19].

Figure 8: Convergence of LSC-DG (right) on a mesh of warped wedges (left). The vertex positions are perturbed by 10%, and the mesh is then refined by splitting. The convergence of LSC-DG is plotted using a black line, while red lines compare the convergence of DG on an un-warped mesh of wedges.

	Hexahedra	Wedge	Pyramids	Tetrahedra	Total elements
Mesh 1	84	10	24	83	201
Mesh 2	672	80	96	856	1704
Mesh 3	5376	640	384	7616	14016
Mesh 4	43008	5120	1536	64000	113664

Table 4: Number of elements of each type for each hybrid mesh used in convergence tests.

Under the assumption that the mesh size h is halved at each time, we estimate the asymptotic convergence rate in Figure 9 for $N = 1, 2, 3$. For both the SEM and GL formulations, we observe orders of convergence (in between $h^{N+1/2}$ and h^{N+1}) similar to those reported for single element-type meshes in Section 4.1.

4.3. Computational results

In Section 4.3.1 and 4.3.2, we quantify the computational cost of DG solvers on hybrid meshes in terms of runtime per degree of freedom. While the approximation power per degree of freedom varies from element to element, the cost per dof gives a rough estimate of the efficiency of each kernel. In Section 4.3.3, we quantify the computational efficiency of the solver in terms of estimated bandwidth and [floating point operations per](#)

Figure 9: An exploded view of the hybrid mesh with 201 elements (left) and L^2 errors for both the SEM and GL formulation (right).

second (GFLOPS) for each element kernel. All computations were run on a single Nvidia GTX 980 GPU in single precision.

4.3.1. Cost of SEM vs GL formulations

Since the computational structure of the volume kernel is identical between the SEM and GL formulations, any additional cost associated with the GL formulation lies in the surface and update kernels. We implemented two separate surface and update kernels and report the relative per-dof speedup of SEM compared to GL for orders $N = 1, \dots, 5$ in Table 5. For all orders, the cost of the SEM update is lower or equal to the cost of the GL update. For low orders, the SEM-tailored surface kernel performs slightly worse than that of the GL kernel². However, as the order N increases, the SEM kernel becomes cheaper due to the fact that data from interior nodes does not need to be accessed. When comparing total runtime at $N > 1$, the SEM formulation becomes roughly 5 – 10% cheaper per-dof than the GL formulation, similar to the 10 – 15% cost savings using SEM over GL as reported by Kopriva and Gassner on CPU architectures [28].

N	1	2	3	4	5
Ratio of surface cost of SEM vs GL	1.1243	1.0597	0.9350	0.8986	0.8329
Ratio of update cost of SEM vs GL	0.9884	0.7437	1.0000	0.8408	0.9000
Ratio of total cost of SEM vs GL	1.0432	0.9286	0.9798	0.9163	0.9134

Table 5: Total time-per-dof cost for SEM vs GL hexahedra.

4.3.2. Cost per element type

In this section, we compare the computational cost of hexahedra, wedges, and pyramids relative to the computational cost of tetrahedra. The results reported are for tuned computational kernels, where the number of elements processed per workgroup has been chosen in order to minimize the runtimes of the volume, surface, and update kernels for each element [34, 17]. As suggested in [26], automation of this process is crucial for portable performance across various architectures, especially for hybrid meshes where parameters must be tuned for 12 separate kernels.

Table 6 shows the time-per-dof cost for $K \approx 100000$ hexahedral, wedge, or pyramidal elements relative to the time-per-dof cost of $K \approx 100000$ tetrahedral elements. While hexahedra are observed to be faster per-dof than tetrahedra at higher orders, wedge and pyramidal volume kernels are observed to be 1-2 times slower per-dof than tetrahedra kernels. This may be due in part to the fact that, for planar tetrahedra, geometric factors are constant over an element. As a result, only a few values must be loaded per element independently of N . This is in contrast to hexahedra, wedges and pyramids, for which the number of geometric factors to be loaded increases along with the total number of nodal or cubature points.

We note that these numbers give optimistic estimates on per-dof costs for two reasons. First of all, for meshes containing only planar tetrahedra, the cost of the surface and update kernels may be further reduced by replacing surface quadratures with lift operators for nodal bases [22, 26]. Secondly, on hex-dominant meshes, the number of wedges and pyramids is expected to be much smaller than the number of hexahedra and tetrahedra. The runtimes reported previously are for asymptotically “large” numbers of elements, and may increase for a small number of elements.

To illustrate the dependence of runtime on number of elements, we timed the cost of the volume kernel over 500 timesteps on meshes containing approximately 100, 1000, 10000, and 30000 elements. For ease of presentation, we show results only for the volume kernel, though the behavior is similar for the surface and update kernels. The order of approximation is varied from $N = 1, \dots, 5$ for each mesh, and the kernel runtime is reported in Figure 10. Since the computational structure of each volume kernel is the same

²The slower runtime may be due to additional conditional statements and the structure of memory accesses in the SEM-tailored kernel. We note that it is always possible to run the SEM formulation using GL kernels, and thus the cost of the SEM surface kernel can always be made in practice to be less than or equal to the cost of the GL kernel

N	1	2	3	4	5
Hex SEM kernels					
Volume	1.2238	1.0373	1.1147	0.9538	0.7772
Surface	0.9631	0.8735	0.5666	0.6038	0.5808
Update	1.0254	0.8427	0.9161	0.8762	0.6667
Total (all kernels)	1.0492	0.9108	0.8222	0.7987	0.6738
Hex GL kernels					
Volume	1.2238	1.0373	1.1147	0.9538	0.7772
Surface	0.8567	0.8243	0.6060	0.6719	0.6974
Update	1.0374	1.1332	0.9161	1.0422	0.7407
Total (all kernels)	1.0058	0.9808	0.8392	0.8717	0.7377
Wedge kernels					
Volume	2.965	1.800	2.999	2.596	2.975
Surface	0.912	0.778	0.99	2.086	1.566
Update	1.414	1.02	1.111	1.321	1.485
Total (all kernels)	1.5937	1.1450	1.5972	2.0407	2.0242
Pyramid kernels					
Volume	1.912	1.653	2.418	3.228	2.842
Surface	0.959	0.855	0.634	1.657	1.157
Update	1.51	1.033	1.112	1.169	1.804
Total (all kernels)	1.3705	1.1386	1.2788	2.0470	1.9280

Table 6: Time-per-dof cost (relative to tetrahedra) of hexahedra, wedge and pyramid kernels.

Figure 10: Time-per-dof cost of hexahedral, wedge, pyramidal, and tetrahedral volume kernels on meshes with $K \approx 100, 1000, 10000$ and 30000 elements at various orders N .

between the SEM and GL formulations (with the exception of the pyramid³) we show timings for only four volume kernels.

Since the time-per-dof decreases as more elements are processed, if a mesh contains a small number of wedges or pyramids, the actual cost of a wedge or pyramid relative to a tetrahedron may be slightly greater than the reported values of Table 6, though this effect is less pronounced at higher orders of approximation. For example, for $N = 3$ on the largest hybrid mesh (consisting of 113664 elements, with 5120 prisms, 1536 pyramids, and 64000 tetrahedra), prisms and pyramids were respectively 1.7563 and 1.8849 times more expensive per dof than tetrahedra, instead of the estimated 1.6 and 1.28 times suggested by Table 6.

4.3.3. Estimated GFLOPS and bandwidth

³Under the GL formulation, the pyramid volume kernel may use the more efficient strong formulation. We show timings, GFLOPS, and estimated bandwidth for the skew-symmetric pyramid kernel only.

Finally, we present estimated GFLOPS and effective bandwidth counts for each kernel in Table 7. Experiments are run on a Nvidia GTX 980 card, with a peak performance of 4612 GFLOPS/s and peak bandwidth of 224 GB/s. Comparing only the volume kernels, we see efficiency resulting from both effective hardware utilization and reduced computational load. For example, the wedge volume kernel shows low estimated bandwidth with high estimated GFLOPS, indicating that its performance is compute-bound. We note that the observed GFLOPS are relatively high, which may be due to the heavy use of `float4` operations within the wedge volume kernel. On the other hand, the bandwidth and GFLOPS for the tetrahedron volume kernel are more balanced. Finally, though it is similar in structure to the tetrahedral kernel, the pyramid volume kernel shows high estimated bandwidth and low GFLOPS, implying that it is IO bound. This is likely due to the fact that the skew-symmetric form is used, and requires the loading of twice as many operators (derivative matrices and their transposes) as the tetrahedral kernel.

For hexahedral kernels, the operators are explicitly stored in shared memory, removing caching effects. As a result, the reported bandwidth and GFLOPS are both significantly lower than for other elements, but near-peak bandwidth numbers are achieved for $N \geq 3$, implying the kernel is IO bound. This near-peak performance is likely due to the loading of multiple geometric change-of-variables factors per node per element and the low arithmetic intensity of tensor-product operations, which are effectively hidden during memory retrievals. It has been noted in [39] that, for vertex-mapped hexahedra, the IO bound nature of the hex volume kernel may be addressed by computing these geometric factors on the fly inside the kernel, loading only vertex positions (regardless of order).

For wedges, pyramids, and tetrahedra, bandwidth is reported both with and without counting operator loads per element. The reported bandwidth numbers with operator loads are often higher than peak rates, which is likely due to the fact that compiler optimizations and caching effects are not taken into account during estimates. Removing operator loads from bandwidth counts reveals that the bandwidth for prism, pyramid, and tetrahedron kernels decreases as N increases, verifying that operations for these three element types are compute-bound at high orders.

5. Conclusions

We have presented in this paper a high order discontinuous Galerkin solver on hybrid, hex-dominant meshes. A careful selection of basis functions controls the storage cost of the scheme, making it suitable for accelerators and GPUs. Two energy-stable formulations are given, based on either SEM or Gauss-Legendre nodal bases for the hexahedron. Explicit bounds on the spectra are computed in terms of the order-dependent constants in the trace inequalities for each type of element, and are used to determine local timestep restrictions for multi-rate timestepping. Convergence rates are reported for each formulation, and computational cost is assessed by estimating the GFLOPS, bandwidth, and cost-per-dof of each element.

The focus of this work has been mainly based presenting a DG solver for hybrid meshes on a single GPU. Future work will focus on the following areas:

- **Wedge basis:** while the LSC-DG wedge basis achieves optimal rates of convergence, there is additional complexity and computational cost compared to polynomial bases (necessity of quadrature, storage of derivatives of J). A low-storage polynomial wedge basis could reduce costs for wedge volume kernels and surface/update kernels for wedges, pyramids, and tetrahedra.
- **Hybrid meshes compared to tetrahedral meshes:** while cost-per-dof comparisons have been made between tetrahedra and other element types, DG on hex-dominant meshes should be compared to optimized nodal DG on fully tetrahedral meshes. The effect on efficiency of the ratio between hexes and elements of other types should also be carefully investigated.
- **Parallelization and scalability:** multiple GPUs are necessary to handle larger problem sizes. Performance of multiple GPU-accelerated DG will require load balancing and communication strategies for multi-rate solvers on hybrid meshes.

N	1	2	3	4	5	1	2	3	4	5
Hex SEM										
Volume	100	134	173	199	232	152	160	169	165	167
Surface	68	51	57	47	35	88	68	81	70	54
Update	46	34	49	40	44	159	112	155	125	133
Hex GL										
Surface	60	44	52	43	32	77	59	74	64	49
Update	45	46	49	48	48	160	152	155	149	148
Wedge										
Volume	242	552	1032	1551	1956	217/149	271/156	294/122	316/107	300/79
Surface	197	418	566	452	771	166/103	273/98	328/70	247/31	408/33
Update	166	403	736	904	960	177/116	309/132	472/128	533/99	536/69
Pyramid										
Volume	77	161	238	259	392	220/102	444/106	646/83	699/52	1054/49
Surface	187	344	850	482	922	163/107	236/100	510/131	271/46	497/55
Update	136	333	699	898	670	155/103	270/123	464/137	546/116	397/58
Tet										
Volume	123	200	401	550	699	238/145	322/125	595/139	784/116	974/95
Surface	149	228	423	596	791	122/81	147/65	244/64	327/58	421/50
Update	161	260	599	779	931	205/147	238/129	434/160	509/144	559/112

(a) GFLOPS

(b) Est. bandwidth (GB/s)

Table 7: GFLOPS and estimated effective bandwidth counts for each kernel of each element type (after optimization of the number of elements processed per workgroup). For wedges, pyramids, and tetrahedra, bandwidth is reported (with operator loads)/(without operator loads).

6. Acknowledgements

Jesse Chan is supported by the Rice University CAAM Department Pfeffer Postdoctoral Fellowship and by NSF (award number DMS-1216674). He would additionally like to acknowledge and thank the Jennifer J. Young Memorial Fund. Axel Modave is an Honorary Fellow of the Belgian American Educational Foundation (BAEF), and acknowledges Wallonie-Bruxelles International (WBI) and Shell for their support. T. Warburton is supported partially by ANL (award number 1F-32301, subcontract on DOE DE-AC02-06CH11357).

References

- [1] Ali R Amir-Moéz and Arnold Lionel Fass. *Elements of Linear Spaces*, volume 26. Pergamon Press, Oxford, 1962.
- [2] Tristan Carrier Baudouin, Jean-François Remacle, Emilie Marchandise, François Henrotte, and Christophe Geuzaine. A frontal approach to hex-dominant mesh generation. *Advanced Modeling and Simulation in Engineering Sciences*, 1(1):1–30, 2014.
- [3] G Bedrosian. Shape functions and integration formulas for three-dimensional finite element analysis. *International journal for numerical methods in engineering*, 35(1):95–108, 1992.
- [4] Morgane Bergot, Gary Cohen, and Marc Duruflé. Higher-order finite elements for hybrid meshes using new nodal pyramidal elements. *Journal of Scientific Computing*, 42(3):345–381, 2010.
- [5] Morgane Bergot and Marc Duruflé. Approximation of $H(\text{div})$ with high-order optimal finite elements for pyramids, prisms and hexahedra. *Communications in Computational Physics*, 14(5):1372–1414, 2013.

- [6] Morgane Bergot and Marc Duruflé. High-order optimal edge elements for pyramids, prisms and hexahedra. *Journal of Computational Physics*, 232(1):189–213, 2013.
- [7] Morgane Bergot and Marc Duruflé. Higher-order discontinuous Galerkin method for pyramidal elements using orthogonal bases. *Numerical Methods for Partial Differential Equations*, 29(1):144–169, 2013.
- [8] Lorenzo Botti. Influence of reference-to-physical frame mappings on approximation properties of discontinuous piecewise polynomial spaces. *Journal of Scientific Computing*, 52(3):675–703, 2012.
- [9] Erik Burman and Alexandre Ern. Continuous interior penalty *hp*-finite element methods for advection and advection-diffusion equations. *Mathematics of computation*, 76(259):1119–1140, 2007.
- [10] Jesse Chan and T Warburton. *hp*-finite element trace inequalities for the pyramid. *Computers & Mathematics with Applications*, 2015.
- [11] Jesse Chan and T Warburton. Orthogonal bases for vertex-mapped pyramids. *arXiv preprint arXiv:1502.07703*, 2015. Accepted to SISC.
- [12] Fan Chung and Franklin Kenter. Discrepancy inequalities for directed graphs. *Discrete Applied Mathematics*, 176:30–42, 2014.
- [13] Gary Cohen, Xavier Ferrieres, and Sébastien Pernet. A spatial high-order hexahedral discontinuous Galerkin method to solve Maxwell’s equations in time domain. *Journal of Computational Physics*, 217(2):340–363, 2006.
- [14] Moshe Dubiner. Spectral methods on triangles and other domains. *Journal of Scientific Computing*, 6(4):345–390, 1991.
- [15] John A Evans and Thomas JR Hughes. Explicit trace inequalities for isogeometric analysis and parametric hexahedral finite elements. *Numerische Mathematik*, 123(2):259–290, 2013.
- [16] Federico Fuentes, Brendan Keith, Leszek Demkowicz, and Sriram Nagaraj. Orientation embedded high order shape functions for the exact sequence elements of all shapes. Technical report, ICES, 2015.
- [17] Rajesh Gandham, DS Medina, and T. Warburton. GPU accelerated discontinuous Galerkin methods for shallow water equations. *arXiv preprint arXiv:1403.1661*, 2014.
- [18] Gregor J Gassner, Frieder Lörcher, Claus-Dieter Munz, and Jan S Hesthaven. Polymorphic nodal elements and their application in discontinuous Galerkin methods. *Journal of Computational Physics*, 228(5):1573–1590, 2009.
- [19] Christophe Geuzaine and Jean-François 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):1309–1331, 2009.
- [20] Nico Gödel, Nigel Nunn, Tim Warburton, and Markus Clemens. Scalability of higher-order discontinuous Galerkin FEM computations for solving electromagnetic wave propagation problems on GPU clusters. *Magnetics, IEEE Transactions on*, 46(8):3469–3472, 2010.
- [21] Nico Gödel, Steffen Schomann, Tim Warburton, and Markus Clemens. GPU accelerated Adams–Bashforth multirate discontinuous Galerkin FEM simulation of high-frequency electromagnetic fields. *Magnetics, IEEE Transactions on*, 46(8):2735–2738, 2010.
- [22] Jan S Hesthaven and Tim Warburton. *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*, volume 54. Springer, 2007.
- [23] Koen Hillewaert. *Development of the Discontinuous Galerkin Method for High-resolution, Large Scale CFD and Acoustics in Industrial Geometries*. PhD thesis, University de Louvain, 2013.

- [24] George Karniadakis and Spencer Sherwin. *Spectral/hp element methods for computational fluid dynamics*. Oxford University Press, 2013.
- [25] RM Kirby, TC Warburton, I Lomtev, and GE Karniadakis. A discontinuous Galerkin spectral/hp method on hybrid grids. *Applied numerical mathematics*, 33(1):393–405, 2000.
- [26] Andreas Klöckner, Tim Warburton, Jeff Bridge, and Jan S Hesthaven. Nodal discontinuous Galerkin methods on graphics processors. *Journal of Computational Physics*, 228(21):7863–7882, 2009.
- [27] Tom Koornwinder. Two-variable analogues of the classical orthogonal polynomials. In *Theory and application of special functions (Proc. Advanced Sem., Math. Res. Center, Univ. Wisconsin, Madison, Wis., 1975)*, pages 435–495. Academic Press New York, 1975.
- [28] David A Kopriva and Gregor Gassner. On the quadrature and weak form choices in collocation type discontinuous Galerkin spectral element methods. *Journal of Scientific Computing*, 44(2):136–155, 2010.
- [29] Lilia Krivodonova and Ruibin Qin. An analysis of the spectrum of the discontinuous Galerkin method. *Applied Numerical Mathematics*, 64:1–18, 2013.
- [30] Lilia Krivodonova and Ruibin Qin. An analysis of the spectrum of the discontinuous Galerkin method II: Nonuniform grids. *Applied Numerical Mathematics*, 71:41–62, 2013.
- [31] Ethan J Kubatko, Clint Dawson, and Joannes J Westerink. Time step restrictions for Runge-Kutta discontinuous Galerkin methods on triangular grids. *Journal of Computational Physics*, 227(23):9697–9710, 2008.
- [32] Johan Malm, Philipp Schlatter, Paul F Fischer, and Dan S Henningson. Stabilization of the spectral element method in convection dominated flows by recovery of skew-symmetry. *Journal of Scientific Computing*, 57(2):254–277, 2013.
- [33] David S Medina, Amik St-Cyr, and T Warburton. OCCA: A unified approach to multi-threading languages. *arXiv preprint arXiv:1403.0968*, 2014.
- [34] A Modave, A St-Cyr, WA Mulder, and T Warburton. A nodal discontinuous Galerkin method for reverse-time migration on GPU clusters. *Geophysical Journal International*, 203(2):1419–1435, 2015.
- [35] Nilima Nigam and Joel Phillips. Numerical integration for high order pyramidal finite elements. *ESAIM: Mathematical Modelling and Numerical Analysis*, 46(02):239–263, 2012.
- [36] Sevtap Ozisik, Beatrice Riviere, and Tim Warburton. On the constants in inverse inequalities in L^2 . Technical report, Rice University, 2010.
- [37] Anthony T Patera. A spectral element method for fluid dynamics: laminar flow in a channel expansion. *Journal of computational Physics*, 54(3):468–488, 1984.
- [38] Joseph Proriol. Sur une famille de polynomes à deux variables orthogonaux dans un triangle. *Comptes Rendus Hebdomadaires des Séances de L’Académie des Sciences*, 245(26):2459–2461, 1957.
- [39] J.F. Remacle, R. Gandham, and T. Warburton. GPU accelerated spectral finite elements on all-hex meshes. *arXiv preprint arXiv:1506.05996*, 2015.
- [40] Joachim Schöberl, H Gerstmayr, and R Gaisbauer. NETGEN-automatic mesh generator, 2012.
- [41] Spencer J Sherwin, T. Warburton, and George Em Karniadakis. Spectral/hp methods for elliptic problems on hybrid grids. *Contemporary Mathematics*, 218:191–216, 1998.
- [42] Thomas Toulorge and Wim Desmet. CFL conditions for Runge-Kutta discontinuous Galerkin methods on triangular grids. *Journal of Computational Physics*, 230(12):4657–4678, 2011.

- [43] T. Warburton. *Spectral/hp methods on polymorphic multidomains: Algorithms and applications*. PhD thesis, Brown University, 1999.
- [44] T. Warburton. An explicit construction of interpolation nodes on the simplex. *Journal of engineering mathematics*, 56(3):247–262, 2006.
- [45] T. Warburton. A low storage curvilinear discontinuous Galerkin time-domain method for electromagnetics. In *Electromagnetic Theory (EMTS), 2010 URSI International Symposium on*, pages 996–999. IEEE, 2010.
- [46] T. Warburton. A low-storage curvilinear discontinuous Galerkin method for wave problems. *SIAM Journal on Scientific Computing*, 35(4):A1987–A2012, 2013.
- [47] T. Warburton and Jan S Hesthaven. On the constants in hp -finite element trace inverse inequalities. *Computer methods in applied mechanics and engineering*, 192(25):2765–2773, 2003.
- [48] Freddie D Witherden, Antony M Farrington, and Peter E Vincent. PyFR: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach. *Computer Physics Communications*, 185(11):3028–3040, 2014.
- [49] H Xiao and Zydrunas Gimbutas. A numerical algorithm for the construction of efficient quadrature rules in two and higher dimensions. *Comput. Math. Appl.*, 59:663–676, 2010.

Appendix A. Explicit trace inequalities for tensor product elements

The constant in the surface trace inequality for tensor product elements may be computed explicitly. Throughout these proofs, we define ϕ_j be the normalized Legendre polynomials, orthogonal with respect to the standard $L^2([-1, 1])$ inner product. We begin by proving a result in 1D.

Lemma Appendix A.1. *Let $u \in P^N([-1, 1])$. Then,*

$$|u(-1)| + |u(1)| \leq \frac{(N+1)(N+2)}{2} \|u\|_{L^2([-1, 1])}^2$$

Proof. In 1D, we wish to derive an expression for the spectral radius $\rho(M_s)$, where

$$(M_s)_{ij} = \phi_j(-1)\phi_i(-1) + \phi_j(1)\phi_i(1).$$

and $\phi_j(x)$ are normalized Legendre polynomials. Since Legendre polynomials are symmetric across $x = 0$, $\phi_j(-1) = \phi_j(1)$ for j even, and $\phi_j(-1) = -\phi_j(1)$ for j odd. Then, assuming i and j do not share the same parity, we have

$$(M_s)_{ij} = \phi_j(-1)\phi_i(-1) + \phi_j(1)\phi_i(1) = 0.$$

This implies M_s is block diagonal, with each block corresponding to even/odd powers of the Legendre polynomial. Both the even block M^e and odd block M^o reduce down to

$$(M^e)_{ij} = (M^o)_{ij} = 2\phi_j(-1)\phi_i(-1) = 2\phi_j(1)\phi_i(1).$$

These are both rank 1 matrices, such that

$$\rho(M^e) = 2 \sum_{j=\text{even}} \phi_j(-1)^2, \quad \rho(M^o) = 2 \sum_{j=\text{odd}} \phi_j(-1)^2.$$

For both even and odd i, j , $\phi_j(-1)^2 = (2i+1)/2$, implying that

$$\rho(M_s) = 2 \max(\rho(M^e), \rho(M^o)) = \frac{(N+1)(N+2)}{2}.$$

□

The extremal polynomial for the above inequality may also be characterized explicitly.

Lemma Appendix A.2. *Define r_i to be the quadrature points of the $(N+2)$ -point Gauss-Legendre-Lobatto (GLL) rule. Let u be the order N Lagrange polynomial that is zero at N interior GLL nodes and unity at either $r = 1$ or $r = -1$. Then,*

$$|u(-1)| + |u(1)| = \frac{(N+1)(N+2)}{2} \|u\|_{L^2([-1,1])}^2$$

Proof. We note that, while the $(N+1)$ -point GLL quadrature does not integrate polynomials of order $2N$ exactly, polynomials of order $2N$ are integrated exactly by the $(N+2)$ -point GLL quadrature.

Since $u(r)$ is symmetric or antisymmetric across $r = 0$ (this may be shown by explicitly writing out the interpolating Lagrange polynomial), $u(1) = \pm u(-1)$ depending on whether N is even or odd. Then, $\|u\|_{L^2}^2$ is given as

$$\|u\|_{L^2}^2 = \int_{-1}^1 u^2 = \sum_{j=1}^{N+2} w_j u(r_j) = w_1 u(-1)^2 + w_{N+2} u(1)^2 = u(1)^2 (w_1 + w_{N+2}).$$

The surface integral of u^2 is just $u(1)^2 + u(-1)^2 = 2u(1)^2$. As a result,

$$\frac{2u(1)^2}{u(1)^2(w_1 + w_{N+2})} = \frac{1}{w_1} = \frac{(N+1)(N+2)}{2}.$$

by the fact that the GLL weights for an $(N+2)$ -point rule are each $\frac{2}{(N+1)(N+2)}$ at the endpoints. $\square$

A curious coincidence is that, since the $N+1$ -point Gauss-Radau-Jacobi (GRJ) rule contains all but one endpoint of the $(N+2)$ -point GLL rule, the extremal polynomial is identical to the extremal polynomial for triangle faces, which is unity at the -1 GRJ point and zero at all others [47].

We may prove a similar result, replacing the mass matrix with the underintegrated Gauss-Legendre-Lobatto quadrature mass matrix.

Lemma Appendix A.3. *Let $u \in P^N([-1, 1])$. Then, for $N \geq 1$, the following bound is tight*

$$|u(-1)| + |u(1)| \leq \frac{N(N+1)}{2} \|u\|_{\text{SEM}}^2.$$

Proof. We take the basis $\phi_i(r) = \ell_i(r)/\sqrt{w_i}$, where $\ell_i(r)$ is the Lagrange polynomial at the i th GLL node and w_i are the quadrature weights of the $(N+1)$ -point GLL rule. Since this basis is orthogonal with respect to the SEM inner product, the constant in the trace inequality is the spectral radius of the surface mass matrix under this basis, which has entries $(M_s)_{ij} = \phi_i(-1)\phi_j(-1) + \phi_i(1)\phi_j(1)$. Applying the Lagrange property, we see that M_s is diagonal with only two nonzero diagonal entries

$$(M_s)_{11} = \frac{1}{w_1}, \quad (M_s)_{N+1, N+1} = \frac{1}{w_{N+1}}.$$

$\rho(M_s)$ is the maximum of these two entries, and the bound is proven by noting $w_1 = w_{N+1} = \frac{2}{N(N+1)}$. The extremal polynomial for this bound is any convex combination of $\phi_1(r)$ or $\phi_{N+1}(r)$. $\square$

Appendix A.1. Tensor product elements

For a d -dimensional hypercube $[-1, 1]^d$, we may define an orthonormal tensor product basis as

$$\phi_{i_1, \dots, i_d} = \prod_{k=1}^d \phi_{i_k}(r_k), \quad 0 \leq i_k \leq N,$$

where r_k is the coordinate in the k th direction. Using the above basis, we may prove the following theorem:

Theorem Appendix A.4. Let $u \in \text{span} \{\phi_{i_1, \dots, i_d}\}$ for $0 \leq i_1, \dots, i_d \leq N$. Then,

$$\|u\|_{L^2(\partial\hat{K})}^2 \leq d \frac{(N+1)(N+2)}{2} \|u\|_{L^2(\hat{K})}^2$$

Furthermore, this bound is tight.

Proof. The surface mass matrix over the d -dimensional hypercube is

$$(M_s)_{i_1, \dots, i_d, j_1, \dots, j_d} = \sum_{k=1}^d (\phi_{i_k}(-1)\phi_{j_k}(-1) + \phi_{i_k}(1)\phi_{j_k}(1)) \prod_{l \neq k} \int_{-1}^1 \phi_{i_l} \phi_{j_l} dr_l.$$

By the orthogonality of Legendre polynomials, this reduces to

$$\begin{aligned} (M_s)_{i_1, \dots, i_d, j_1, \dots, j_d} &= \sum_{k=1}^d (M_s^k)_{i_1, \dots, i_d, j_1, \dots, j_d}, \\ (M_s^k)_{i_1, \dots, i_d, j_1, \dots, j_d} &= (\phi_{i_k}(-1)\phi_{j_k}(-1) + \phi_{i_k}(1)\phi_{j_k}(1)) \prod_{l \neq k} \delta_{i_l j_l}. \end{aligned}$$

This results in a block diagonal matrix in all indices i_l, j_l for $l \neq k$. Also, by the same argument as in the 1D case, if pairs of indices i_k, j_k do not share the same parity,

$$(\phi_{i_k}(-1)\phi_{j_k}(-1) + \phi_{i_k}(1)\phi_{j_k}(1)) = 0,$$

implying that each matrix M_s^k consists of two blocks of even and odd indices. Additionally, each matrix M_s^k contains the same entries as the 1D surface mass matrix and are identical up to a permutation. An application of Weyl's inequality then gives

$$\rho(M_s) \leq \sum_{k=1}^d \rho(M_s^k) = d \frac{(N+1)(N+2)}{2}.$$

We may show that this bound is tight by constructing the extremal polynomial in d dimensions as the tensor product of the 1D polynomial. Take $u_{1D}(r_k)$ to be the 1D extremal polynomial in the k th coordinate and define

$$u(r_1, \dots, r_d) = \prod_{k=1}^d u_{1D}(r_k).$$

The $(N+2)^d$ tensor product GLL rule integrates this exactly, such that for the reference element $\hat{K} = [-1, 1]^d$,

$$\|u\|_{L^2(\hat{K})}^2 = \sum_{i_1=1}^{N+2} \sum_{i_2=1}^{N+2} \dots \sum_{i_d=1}^{N+2} w_{i_1} w_{i_2} \dots w_{i_d} u_{1D}^2(r_{1,i_1}) u_{1D}^2(r_{2,i_2}) \dots u_{1D}^2(r_{d,i_d}) = 2^d w_1^d.$$

where 2^d is the number of vertices on the d -dimensional hypercube.

The surface norm of u for a d -dimensional hypercube is the sum of the integrals over the $2d$ faces. Since each face itself is a hypercube of dimension $d-1$, the integral of u over the face is $2^{d-1} w_1$, and the surface norm of u reduces to

$$\|u\|_{L^2(\hat{K})}^2 = \sum_{f=1}^{2d} 2^{d-1} w_1^{d-1} = d 2^d w_1^{d-1},$$

and we may conclude that

$$\frac{\|u\|_{L^2(\partial\hat{K})}^2}{\|u\|_{L^2(\hat{K})}^2} = \frac{d}{w_1} = d \frac{(N+1)(N+2)}{2}.$$

□

N	1	2	3	4	5	6	7	8	9
Hexahedron	9	18	30	45	63	84	108	135	165
Wedge	9.93	18.56	29.03	42.99	58.80	78.01	99.27	123.76	150.48
Pyramid	11.68	20.89	32.84	47.59	65.17	85.60	108.90	135.07	164.11
Tetrahedron	12.22	20.46	29.18	41.65	54.45	71.10	88.32	109.04	130.67

Table B.8: Computed trace constants $C_T(N)$ using the full mass matrix for various reference elements up to $N = 9$.

N	1	2	3	4	5	6	7	8	9
Hexahedron	3	9	18	30	45	63	84	108	135
Wedge	46.46	51.70	63.94	83.32	104.76	132.20	161.22	195.95	232.52
Pyramid	31.58	37.12	47.59	60.95	77.14	96.31	118.52	143.80	172.13

Table B.9: Computed trace constants $C_{\text{SEM}}(N)$ using SEM quadrature for various reference elements up to $N = 9$. SEM quadrature is used to compute the mass matrix and surface mass matrices when quadrilateral faces are present. The constants for tetrahedra are identical in both cases.

For $N > 1$, similar steps may be followed to prove the equivalent trace inequality under SEM quadrature in d dimensions

$$\|u\|_{\text{SEM}(\partial\hat{K})}^2 \leq d \frac{N(N+1)}{2} \|u\|_{\text{SEM}(\hat{K})}.$$

We note that Ern and Burman also utilized GLL weights to prove a similar trace inequality in d dimensions

$$\|u\|_{L^2(\partial\hat{K})}^2 \leq d \frac{(N+1)(N+2)}{2} \left(2 + \frac{1}{N}\right)^d \|u\|_{L^2(\hat{K})}^2.$$

The additional factor of $(2 + 1/N)^d$ results from the equivalence constant between the L^2 norm and discrete L^2 norm resulting from underintegration using GLL quadrature.

Appendix B. Computed Markov and trace inequality constants

For wedge, pyramid, and tetrahedral elements, we have not been able to determine explicit expressions for the constants $C_T(N)$ in trace inequalities over reference elements $\hat{K}$

$$\|u\|_{L^2(\partial\hat{K})}^2 \leq C_T(N) \|u\|_{L^2(\hat{K})}^2.$$

However, $C_T(N)$ may be computed as the maximum eigenvalue of a generalized eigenvalue problem

$$M_s v = \lambda M v,$$

where M_s is the surface mass matrix and M is the volume mass matrix. We may similarly compute the trace inequality constants $C_{\text{SEM}}(N)$ for the case when SEM quadrature is used for hexahedra and quadrilateral faces. For the Markov inequality on the reference element

$$\|\nabla u\|_{L^2(\hat{K})}^2 \leq C_M(N) \|u\|_{L^2(\hat{K})}^2$$

we may similarly compute the constant $C_M(N)$ as the maximum eigenvalue of

$$K v = \lambda M v,$$

where $K_{ij} = \int_{\hat{K}} \nabla \phi_j \cdot \nabla \phi_i$ is the stiffness matrix over the reference element. We assemble the surface mass matrix and stiffness matrix using surface and volume quadratures which are exact (to numerical roundoff) for each element type at each order of approximation. Constants in the trace and Markov inequalities are then computed from the solution of generalized eigenvalue problems using the `eig` routine in MATLAB. We summarize computed values of $C_T(N)$, $C_{\text{SEM}}(N)$ and $C_M(N)$ up to $N = 9$ in Tables B.8, B.9 and B.10. For a SEM hex, we may determine the constant in the Markov inequality using the full mass matrix value of $C_M(N)$ and equivalence constants between discrete SEM and L^2 inner products.

N	1	2	3	4	5	6	7	8	9
Hexahedron	3.00	18.00	55.73	137.51	293.97	562.17	985.92	1616.24	2511.47
Wedge	12.00	54.27	142.63	308.34	585.89	1021.64	1663.85	2574.06	3814.56
Pyramid	12.92	60.05	175.51	405.43	809.95	1460.93	2442.26	3849.94	5792.11
Tetrahedron	20.00	78.62	195.58	403.91	744.85	1265.54	2021.09	3073.26	4491.62

Table B.10: Computed Markov constants $C_M(N)$ for various reference elements up to $N = 9$.