

Adaptive spacetime discontinuous Galerkin method for hyperbolic advection–diffusion with a non-negativity constraint

Raj Kumar Pal¹, Reza Abedi^{1,2}, Amit Madhukar¹ and Robert B. Haber^{1,*},[†]

¹*Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, 1206 West Green Street, Urbana, IL 61801, USA*

²*Department of Mechanical, Aerospace & Biomedical Engineering, University of Tennessee Space Institute, 411 B.H. Goethert Parkway, MS 21, Tullahoma, TN 37388, USA*

SUMMARY

Applications where the diffusive and advective time scales are of similar order give rise to advection–diffusion phenomena that are inconsistent with the predictions of parabolic Fickian diffusion models. Non-Fickian diffusion relations can capture these phenomena and remedy the paradox of infinite propagation speeds in Fickian models. In this work, we implement a modified, frame-invariant form of Cattaneo’s hyperbolic diffusion relation within a spacetime discontinuous Galerkin advection–diffusion model. An h -adaptive spacetime meshing procedure supports an asynchronous, patch-by-patch solution procedure with linear computational complexity in the number of spacetime elements. This localized solver enables the selective application of optimization algorithms in only those patches that require inequality constraints to ensure a non-negative concentration solution. In contrast to some previous methods, we do not modify the numerical fluxes to enforce non-negative concentrations. Thus, the element-wise conservation properties that are intrinsic to discontinuous Galerkin models are defined with respect to physically meaningful Riemann fluxes on the element boundaries. We present numerical examples that demonstrate the effectiveness of the proposed model, and we explore the distinct features of hyperbolic advection–diffusion response in subcritical and supercritical flows. Copyright © 2015 John Wiley & Sons, Ltd.

Received 3 December 2014; Revised 16 June 2015; Accepted 22 July 2015

KEY WORDS: Discontinuous Galerkin; partial differential equations; hyperbolic; transport; spacetime; advection–diffusion; non-negativity constraint

1. INTRODUCTION

Advection–diffusion systems arise in numerous science and engineering applications. Examples include transport of diffusing solute species in a fluid flow [1, 2], two-phase fluid flows in heterogeneous porous media [3], fluid flow in oil reservoirs [4], and drift-diffusion in semiconductor devices [5, 6]. Numerical simulation of these systems poses special challenges, especially in the advection-dominated limit where some form of stabilization is required to suppress spurious oscillations around high-gradient features while avoiding excessive numerical dissipation. Among the many stabilization techniques suitable for finite element methods, we mention streamline upwind Petrov–Galerkin (SUPG) methods [7] and Galerkin least squares methods [8]; see [9] for a survey of the extensive literature on this topic. In general, however, these methods do not ensure monotonicity, can be overly diffusive, and may generate solutions that exhibit nonphysical oscillations or violate critical constraints against negative solution values. Thus, research on improved numerical methods for advection–diffusion problems continues.

*Correspondence to: Robert B. Haber, Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, 1206 West Green Street, Urbana, IL 61801, U.S.A.

[†]E-mail: r-haber@illinois.edu

Diffusive processes are most often modeled by Fick's law [10],

$$\mathbf{q} = -\mathbf{K}\nabla\phi \quad (1)$$

in which \mathbf{q} is the diffusive flux and \mathbf{K} is the diffusivity tensor. After accounting for the equation of continuity, the balance law for a species with concentration ϕ is then governed by the evolution equation,

$$\frac{D\phi}{Dt} = -\nabla \cdot \mathbf{q}, \quad (2)$$

in which

$$\frac{D}{Dt} := \frac{\partial}{\partial t} + \mathbf{a} \cdot \nabla \quad (3)$$

is the material derivative written in terms of a material velocity field, \mathbf{a} . Combining (1) and (2) generates a parabolic model. For non-vanishing material velocity, the second term in the material derivative operator (3) generates an advective transport term, resulting in a parabolic advection–diffusion equation.

The parabolic diffusion model (2) suffers the well-known infinite velocity of propagation paradox [11]. Nonetheless, it provides a good match with experimental observations in many applications where the time scales associated with diffusion are much smaller than the time scales associated with advective transport and other processes in the system dynamics. On the other hand, it is critical to model finite propagation velocities in systems where the diffusion time scale is not well separated from the other time scales in the system. Examples include diffusion in crystalline solids [12], dispersion in amorphous materials [13], blood flows [14], and initial dispersion of pollutants into the atmosphere, rivers, and groundwater [15]. This class of problems motivates the numerical method proposed in this work.

Various constitutive relations have been proposed to model finite-velocity diffusive processes [16]. To this end, one of the simplest extensions of Fick's law is the Maxwell–Cattaneo relation [11], given by

$$\left(1 + \lambda_0 \frac{\partial}{\partial t}\right) \mathbf{q} = -\mathbf{K}\nabla\phi, \quad (4)$$

in which the relaxation constant, $\lambda_0 > 0$, is the time lag required to establish steady conditions after a gradient is imposed across a volume element [17]. Combination of (4) and (2) yields a hyperbolic system with finite propagation velocities.

Although (4) resolves the infinite propagation velocity paradox, Christov and Jordan demonstrate in [18] that it violates the Galilean invariance principle in advection–diffusion problems; *i.e.*, when $\mathbf{a} \neq \mathbf{0}$. They propose a revised Maxwell–Cattaneo relation in which a material derivative replaces the temporal partial derivative in (4),

$$\left(1 + \lambda_0 \frac{D}{Dt}\right) \mathbf{q} = -\mathbf{K}\nabla\phi, \quad (5)$$

and show that the resulting hyperbolic system exhibits finite propagation velocities and is Galilean-invariant. We refer to (5) as the modified Maxwell–Cattaneo (mMC) model and use it as our reference hyperbolic model for advection–diffusion from here on.

This paper is concerned with a finite element method for solving hyperbolic advection–diffusion problems, such as (5). A survey of the literature on numerical methods that address hyperbolic diffusion, for both pure diffusion [14, 19–21], and combined advection–diffusion [1, 22, 23], can be found in [24]. For the pure diffusion case, numerical models based on discontinuous Galerkin (DG) finite element methods, [22, 25], are most relevant to the method proposed here. In particular, the adaptive spacetime DG method described in [26] is a direct antecedent to the method advanced in this work. For the mMC advection–diffusion problem, the DG method in [22] is the closest

precedent to the method proposed in this work. It combines a nonadaptive DG spatial discretization with upwind numerical fluxes in a Runge–Kutta DG method [27, 28].

Hyperbolic advection–diffusion models may predict non-smooth wavefronts, particularly for advection-dominated flows. In these cases, the numerical analyst is often left with a choice between overly diffuse resolutions of these wavefronts and sharper renderings accompanied by spurious numerical oscillations and undershoots that can generate negative concentrations in some regions. In the latter case, and depending on the application, violations of the non-negativity condition may result in algorithmic failure, for example, in advection–diffusion–reaction models or in highly anisotropic media; *cf.* [29]. Thus, we are interested in methods that enforce the non-negativity condition without increasing diffusion.

Various schemes have been proposed to generate non-negative, stable solutions in a variety of applications. Shu and coworkers developed flux limiters to ensure non-negative cell averages of density and pressure in finite volume and DG approximate solutions of the Euler equations [30–32]. Berthon and Marche [33], Xing *et al.* [34], Parent [35], and Balsara [36] used similar ideas to enforce non-negative solutions for the shallow water, Euler, and magneto-hydrodynamic equations as well as other conservation laws. Suresh [37] obtained second-order accurate positive solutions by restricting the reconstructions to cell averages of first-order neighbors.

There are some drawbacks to flux-limiting methods. They use expanded, non-compact stencils in their reconstructions and cannot always handle boundary fluxes accurately and consistently. In hyperbolic problems, any modification of numerical fluxes away from their Riemann values compromises the physical significance of the cell-wise conservation properties available in finite volume and discontinuous Galerkin methods. These schemes might not be suitable for use in h -adaptive solvers because average enforcement of non-negativity on a current mesh does not ensure that positive cell averages are preserved under a projection onto a refined mesh. This can lead to catastrophic failure in models where maintaining a non-negative solution is critical to algorithmic stability.

Several recent papers propose optimization techniques as a means to enforce the non-negativity condition as a formal inequality constraint on the solution. Lipnikov *et al.* [38] propose a finite volume method for advection–diffusion problems in which a constrained optimization program determines a minimal nonlinear correction to the advective fluxes to ensure a non-negative solution. Wang and coworkers [39] combine this approach with piecewise-linear reconstructions on polygonal cells to enforce monotonicity. Liska and Shashkov [40] enforce a maximum principle on second-order elliptic problems by solving a constrained optimization problem using piecewise linear finite elements and linear constraints. This approach has been extended to anisotropic diffusion problems in a mixed formulation [29] and to diffusion-controlled reactive systems on unstructured grids [41, 42]. These methods are often implemented with linear finite element bases. This ensures that solution extrema lie at cell vertices, where it suffices to enforce point-wise constraints.

In this paper, we propose an adaptive spacetime DG (SDG) finite element method for hyperbolic advection–diffusion systems, such as the mMC system in (5), that enforces a non-negativity condition on the solution. We use a particular SDG methodology that is fully discontinuous in both space and time and is implemented on unstructured, asynchronous spacetime grids that conform to a so-called causality constraint determined by the characteristic structure of the underlying hyperbolic system; see [26] for an application of this SDG methodology to hyperbolic diffusion; precedents can be traced back to Lowrie *et al.* [43], Falk and Richter [44], and Yin *et al.* [45]. Part of our purpose is to demonstrate the advantages of this computational technology in the context of hyperbolic advection–diffusion. These favorable properties have been demonstrated through a combination of analysis and extensive numerical testing on a variety of applications [26, 46–53]. They include the following:

- Unconditional stability for linear systems, such as (5). Falk and Richter prove stability of SDG solutions for linear, symmetric hyperbolic systems in [44], and Lowrie *et al.* report similar findings in [43] for linear systems of conservation laws, such as the one in this work.
- Compact computational stencils that do not expand with higher-order approximations; *cf.* discussion in first paragraph in Section 4.
- Element-wise conservation with respect to computed Riemann fluxes.

- Linear computational complexity in the number of spacetime elements.
- Computational efficiency increases with polynomial order for suitably smooth solutions—an uncommon result among finite element methods.
- Asynchronous solution structure that does not propagate time-step size limits, in contrast to conventional time-marching schemes and most synchronous, slab-based SDG methods.
- Continual adaptive remeshing in which incremental solutions, error indicators, and remeshing execute locally at a common granularity. Mesh coarsening and connectivity changes are handled geometrically in spacetime to eliminate errors incurred in other adaptive methods when projecting the solution from an old mesh to a new one.
- Embarrassingly parallel structure for high-performance computing.

In addition, we propose a new method that enforces the non-negativity condition locally, in only those elements where the constraint is active. Similar to previous methods, we use constrained optimization methods to enforce non-negative solutions. However, in contrast to several works cited earlier, we do not limit or modify the computed Riemann fluxes so as to preserve the physical and mathematical significance of the SDG method's element-wise conservation property. Instead, we directly constrain the discrete solution to non-negative values. We believe this new approach has applications to other examples of constrained hyperbolic systems beyond the present non-negative hyperbolic advection–diffusion problem.

The following section reviews key aspects of the SDG finite element method and the differential forms notation we use to obtain an objective, coordinate-free formulation of conservation over spacetime control volumes in 3. Section 3 continues with the formulation of the SDG implementation of the mMC advection–diffusion model and details the new method for enforcing non-negative concentration solutions. Section 4 presents numerical examples that demonstrate the performance of the proposed algorithm and explores aspects of the physical response predicted by the mMC advection–diffusion model in subcritical and supercritical flows. Finally, we present conclusions and directions for future research in Section 5.

2. SPACETIME DISCONTINUOUS GALERKIN FINITE ELEMENT METHOD AND DIFFERENTIAL FORMS NOTATION

This section surveys the adaptive spacetime meshing/solution scheme and the differential forms notation we use to develop the frame-independent SDG formulation presented in 3. This material has been published previously in detail. We provide here only a high-level summary to provide the necessary background for understanding the application of the SDG methodology to advection–diffusion systems and the new method for enforcing non-negativity constraints. The interested reader is encouraged to consult the cited publications for detailed coverage of these topics.

2.1. Asynchronous spacetime meshing and solution

Our spacetime discontinuous Galerkin solver works with asynchronous, unstructured spacetime meshes, such as the simple example in $\mathbb{E}^1 \times \mathbb{R}$ in the top part in Figure 1. The inclined arrows indicate characteristic directions (wave trajectories). In this example, waves travel to the right and left at equal, finite speeds. The SDG solution scheme imposes a *causality constraint* on the mesh geometry that requires each element facet to be faster (*i.e.*, closer to horizontal) than the fastest wave trajectory. Thus, all element facets are *space-like* in the terminology of relativity theory, and information flows in only one direction across each element facet. Accordingly, the later facets of a given element are purely outflow surfaces, and the earlier facets are purely inflow. This structure ensures that the solutions in the two shaded elements depend only on the solutions in the earlier elements adjacent to their respective inflow facets. These earlier solution fragments, however, do not depend on the solutions in the shaded elements.

This asymmetric dependency between solutions in adjacent elements generates a partial element ordering by which the global solution can be computed locally, one element at a time. For example, consider the four level-1 elements along the initial-time boundary. The solution in each of these




Figure 1. Spacetime discontinuous Galerkin solution scheme on causal spacetime mesh in $\mathbb{E}^1 \times \mathbb{R}$ (top). Global coupling in non-causal mesh (bottom). Reproduced from [53].

depends only on the initial data along each element’s bottom facet and, in the case of the leftmost and rightmost elements, on prescribed boundary data. The level-1 element solutions can be computed locally and in parallel if multiple computational cores are available. Any level-2 element can be solved as soon as its immediate level-1 neighbors have been solved, even if other level-1 elements remain unsolved. Thus, causal SDG meshes enable asynchronous, element-by-element solutions with linear computational complexity in the number of spacetime elements (assuming the computational cost of solving each element is roughly the same).

The structured mesh at the bottom in Figure 1 covers an individual *slab* in a synchronous SDG scheme, corresponding to a single time step in a conventional time-marching scheme. The characteristic directions indicate symmetric coupling between solutions in adjacent elements. This dependency spreads globally, so that all elements in the slab must be solved simultaneously. This global coupling is entirely an artifact of the synchronous, non-causal discretization; it does not reflect the mathematical structure or physical response of the underlying hyperbolic model.

In practice, we replace the individual elements in Figure 1 with small clusters of simplex elements called *patches* and only subject the exterior patch facets to the causality constraint. Because inter-element boundaries within a patch may be non-causal, we must solve all the elements in each patch simultaneously. We implement this in an advancing-front meshing/solution procedure. In each step, the *Tent Pitcher* algorithm [46, 49] advances a single vertex in time to define a local update of the space-like *front mesh*. The causality constraint limits the vertex’s maximum time increment. After each update of the front mesh, we construct a small mesh of spacetime simplices to cover the region between the old and new fronts; we call this small collection of simplices a *patch*. We immediately solve each new patch as a local finite element problem, store the solution, and then locally update the current front mesh to the patch’s outflow facets. We repeat this patch-by-patch meshing/solution procedure at front vertices that are local minima in time until the entire spacetime analysis domain is solved.

Global coupling significantly increases computational complexity and limits scalability in traditional time-marching schemes and synchronous SDG models relative to our asynchronous causal scheme. Notably, the growth in computational complexity associated with high-order basis functions dominates the improvement in convergence rate in most globally coupled finite element methods, so lower-order elements are favored. In contrast, numerical efficiency increases with polynomial order in the asynchronous SDG solution scheme until the storage required to solve a single patch exceeds

the available CPU cache.[‡] In the context of high-performance computing, the causal SDG solution scheme supports scalable parallel implementations of combined patch-wise meshing and solution steps, while its asynchronous structure is well suited to very large-scale implementations.

2.2. *h*-adaptive meshing

The local spacetime structure of the SDG solver enables a powerful approach to dynamic adaptive meshing [54]. An error indicator is computed for each new patch solution. If the error is acceptable, the patch solution is accepted and stored for use as inflow data for subsequent patches. If the error is too large, the patch is rejected, and the solver passes a demand for mesh refinement to the meshing code. Thus, only a small amount of computational effort is discarded relative to discarding an entire slab or time step in other methods. If the error is too small, the patch is accepted, and a request for subsequent mesh coarsening is issued. The adaptive Tent Pitcher software responds to refinement demands by refining the front mesh before restarting the patch-generation procedure. This, in combination with the causality constraint, produces simultaneous spatial and temporal refinement. For example, two bisection refinements of segments in the old front mesh generate the smaller shaded element in Figure 1. This refined element has both a smaller spatial diameter and a shorter duration, indicating simultaneous adaptive refinement in space and time.

In higher spatial dimensions, we implement common adaptive meshing operations, such as vertex-deletion, edge flips, and vertex motion for mesh smoothing or feature tracking, as spacetime patches with special configurations that conform to the old front mesh on inflow facets and the new front mesh on outflow facets; see [26] for details. This eliminates the need for expensive and error-prone solution projections, as required in conventional adaptive meshing procedures. In particular, projection errors do not compromise the faster convergence rates of higher-order finite elements. We achieve strong, dynamic refinement because adaptive meshing and the patch-wise solutions are implemented as local operations that share a common fine granularity within the SDG algorithm. This contrasts with conventional adaptive methods, where remeshing is typically a global operation that is invoked only after several time steps have been computed.

2.3. *Differential forms and exterior calculus*

In the case of our hyperbolic advection–diffusion model (5), the SDG solution procedure involves solving a system of conservation laws locally on a series of spacetime patches. Rather than work with spatial control volumes, we formulate the conservation relations directly on spacetime control volumes defined by the elements that comprise each patch. In the case of our unstructured causal meshes, this requires a means to express Stokes theorem in a format where surface integrals sum over d -manifolds with arbitrary orientation in spacetime, where d is the spatial dimension. Recalling the absence of a natural metric between space and time in classical mechanics, we observe that no objective definition of an inner product exists for spacetime vectors in $\mathbb{E}^d \times \mathbb{R}$. Accordingly, no frame-invariant definition for spacetime normal vectors is available for describing fluxes across surfaces of spacetime control volumes.

Differential forms and the exterior calculus on manifolds support objective, coordinate-free statements of the Stokes theorem that do not involve normal vectors. In fact, this approach generates elegant conservation statements from which jump conditions, essential to any discontinuous Galerkin model, emerge in a simple and natural manner. A complete exposition of forms and the exterior calculus is beyond the scope of this paper. However, we next present a brief discussion of our representation of spacetime fluxes using forms and recall the elegant statement of the Stokes theorem in exterior calculus. This background should be sufficient for understanding the main ideas in this paper, but the interested reader can find more complete developments of differential forms and exterior calculus in [55–57] and our application of these tools in formulating SDG models in [26].

[‡]Beyond that point, numerical efficiency may drop as the cache-miss rate grows with polynomial order.

Let $\mathcal{D} \subset \mathcal{M} := \mathbb{E}^d \times \mathbb{R}$ be the spacetime manifold under study. Although we use direct notation in our formulation, for the purposes of this discussion, we introduce a basis for vectors, $\{\mathbf{e}_i, \mathbf{e}_t\}_{i=1}^d$, in which the Cartesian spatial basis, \mathbf{e}_i , spans \mathbb{E}^d , and \mathbf{e}_t is the temporal basis for \mathbb{R} . The dual basis for covectors is denoted as $\{\mathbf{e}^i, \mathbf{e}^t\}_{i=1}^d$. We use \wedge and \mathbf{d} to denote the exterior product and exterior derivative operators. Our standard basis for the differentials in 1–forms is $\{dx^i, dt\}_{i=1}^d$. Top $(d + 1)$ -forms, involving spacetime differential volumes, have the singleton basis, $\Omega = dx^1 \wedge \dots \wedge dx^d \wedge dt$. We also require d -forms to represent differentials on surfaces of spacetime volumes. Our basis for the differentials in d -forms is $\{\star dx^i, \star dt\}_{i=1}^d$, in which \star is the Hodge star operator, such that $dx^i \wedge \star dx^j = \delta^i_j \Omega$, $dt \wedge \star dx^j = \mathbf{0}$, $dt \wedge \star dt = \Omega$, and $dx^j \wedge \star dt = \mathbf{0}$ for $i, j = 1, \dots, d$.[§] For spatial dimension, $d = 3$, $\star dx^1 = dx^2 \wedge dx^3 \wedge dt$, $\star dx^2 = -dx^1 \wedge dx^3 \wedge dt$, $\star dx^3 = dx^1 \wedge dx^2 \wedge dt$, and $\star dt = -dx^1 \wedge dx^2 \wedge dx^3$ in which the varying signs reflect the alternating property of exterior products. Note that $\star dt$ is a purely spatial differential volume with dimension, L^3 , while the $\star dx^i$ are mixed differentials in space and time with dimension, L^2T .

In conventional representations, we conceive of scalar, vector, and tensor fields as distinct mathematical entities from the differentials that accompany them in integrals. In forms notation, on the other hand, the fields and differentials are unified entities in which the fields are called the *coefficients* of the differential forms. We use bold, upright symbols to denote vector and tensor fields and bold slanted symbols to denote differential forms. We introduce a useful d -form with covector coefficients, $\star \mathbf{dx} := \mathbf{e}^i \star dx^i$, and note that

$$\mathbf{d}(w \star \mathbf{dx}) = \nabla w \Omega \qquad \mathbf{d}(w \star dt) = \dot{w} \Omega \tag{6a}$$

$$\mathbf{d}(\mathbf{c} \cdot \star \mathbf{dx}) = (\nabla \cdot \mathbf{c}) \Omega \qquad \mathbf{d}(\mathbf{c} \star dt) = \dot{\mathbf{c}} \Omega \tag{6b}$$

$$\mathbf{d}(\mathbf{D} \star \mathbf{dx}) = (\nabla \cdot \mathbf{D}) \Omega \qquad \mathbf{d}(\mathbf{D} \star dt) = \dot{\mathbf{D}} \Omega \tag{6c}$$

in which w and \mathbf{c} are scalar and vector fields, while \mathbf{D} is a tensor field of order 2 or higher.[¶]

Spacetime fluxes and sources of scalar and vector conservation variables take standard forms:

$$\mathbf{F}_w = w \star dt + \mathbf{f}_w \cdot \star \mathbf{dx} \qquad \mathbf{S}_w = s_w \Omega \tag{7a}$$

$$\mathbf{F}_\mathbf{c} = \mathbf{c} \star dt + \mathbf{F}_\mathbf{c} \star \mathbf{dx} \qquad \mathbf{S}_\mathbf{c} = \mathbf{s}_\mathbf{c} \Omega \tag{7b}$$

in which \mathbf{f}_w and $\mathbf{F}_\mathbf{c}$ are spatial fluxes and s_w and $\mathbf{s}_\mathbf{c}$ are sources of vector and tensor conservation variables whose respective density fields are w and \mathbf{c} . The restriction of a spacetime flux to any oriented, spacetime d -manifold delivers the flux of the corresponding conservation variable across the manifold. This restriction, as it applies to the spatial flux component, replaces the contraction of a spatial flux with a normal vector in tensorial descriptions. However, the restriction of a form to a manifold does not involve a spacetime normal vector, so it is not affected by the lack of an inner product on $\mathbb{E}^d \times \mathbb{R}$. Thus, we obtain an objective representation for spacetime fluxes with forms that are not available with tensorial representations.

The forms representation of spacetime fluxes leads to an elegant local statement of conservation. Given a spacetime domain, $\mathcal{D} \subset \mathbb{E}^d \times \mathbb{R}$, consider any open subdomain, $\mathcal{Q} \subset \mathcal{D}$, with a regular boundary, $\partial \mathcal{Q}$. The Stokes theorem in differential forms notation is [56]

$$\int_{\partial \mathcal{Q}} \omega = \int_{\mathcal{Q}} \mathbf{d}\omega, \tag{8}$$

in which ω is a d -form. Conservation requires that the integral on $\partial \mathcal{Q}$ of a spacetime flux, \mathbf{F} , must equal the integral on \mathcal{Q} of the corresponding source, \mathbf{S} . That is, by (8) and the localization theorem,

$$\int_{\partial \mathcal{Q}} \mathbf{F} = \int_{\mathcal{Q}} \mathbf{S} \quad \forall \mathcal{Q} \subset \mathcal{D} \iff \mathbf{d}\mathbf{F} - \mathbf{S} = \mathbf{0} \text{ on } \mathcal{D} \tag{9}$$

[§]For purposes of the summation convention, the \star operator lowers the indices of $\star dx^j$ to subindices.

[¶]Although we have constructed $\star \mathbf{dx}$ with respect to a particular coordinate frame, it is easy to show that it is frame-invariant.

in which the exterior derivatives are in the sense of distributions. Combining (6), (7), and (9) leads directly to familiar local conservation equations for vector and scalar variables:

$$d\mathbf{F}_w - \mathbf{S}_w = \mathbf{0} \text{ on } \mathcal{D} \iff \dot{w} + \nabla \cdot \mathbf{f}_w - s_w = 0 \text{ on } \mathcal{D} \tag{10a}$$

$$d\mathbf{F}_c - \mathbf{S}_c = \mathbf{0} \text{ on } \mathcal{D} \iff \dot{c} + \nabla \cdot \mathbf{F}_c - \mathbf{s}_c = \mathbf{0} \text{ on } \mathcal{D} \tag{10b}$$

Our SDG formulation involves forms with both scalar and vector coefficients, so we extend the usual definition of the exterior product operator to forms with vector coefficients. Let \mathbf{A} and \mathbf{B} be r -forms and s -forms and \mathbf{a} and \mathbf{b} be covector and vector fields, all on \mathcal{D} . Then, $\mathbf{a}\mathbf{A}$ and $\mathbf{b}\mathbf{B}$ are r -forms and s -forms with covector and vector coefficients. Their exterior product is the $(r + s)$ -form with scalar coefficients given by

$$\mathbf{a}\mathbf{A} \wedge \mathbf{b}\mathbf{B} := \mathbf{a}(\mathbf{b})(\mathbf{A} \wedge \mathbf{B}). \tag{11}$$

3. FORMULATION

We introduce the continuum hyperbolic advection–diffusion model as a system of conservation equations and identify the corresponding spacetime fluxes and source terms using differential forms notation. We then introduce *target fluxes* as a unified vehicle for specifying initial, boundary, and Riemann conditions and summarize the Riemann solutions for this problem. Next, following methods described in [26, 47], we formulate weighted residual equations and describe the SDG discretization in which the jump part of the exterior derivative operator generates jump conditions that enforce the target fluxes. We present a new method that imposes non-negativity constraints on concentration solutions and describe the *a posteriori* error indicators that drive our h -adaptive meshing procedure; cf. 2.2.

3.1. Continuum model

Consider the hyperbolic advection–diffusion model of [24] for mass transport of a solute species in the presence of a known, divergence-free (incompressible) background velocity field, $\mathbf{a}(\mathbf{x}, t)$, in which $\mathbf{x} \in \mathbb{E}^d$ denotes spatial position and $t \in \mathbb{R}$ denotes time. Let u denote the solute species concentration and \mathbf{q} its diffusive flux. Then the system of governing equations from [24] can be written as

$$\dot{u} + \nabla u \cdot \mathbf{a} + \nabla \cdot \mathbf{q} - s_u = 0, \tag{12a}$$

$$\tau [\dot{\mathbf{q}} + (\nabla \mathbf{q}) \mathbf{a} + \mathbf{K} \nabla u] + \mathbf{q} = \mathbf{0} \tag{12b}$$

in which \mathbf{K} is a positive-definite diffusivity tensor and τ is a second-order relaxation tensor with dimensions of time. For simplicity, but without loss of generality, we assume isotropic and homogeneous materials from here on. Thus, $\tau = \tau \mathbf{I}$ and $\mathbf{K} = k \mathbf{I}$, where \mathbf{I} is the identity tensor.

Because (12) is a hyperbolic system, the solute species has a finite speed of diffusive transport, given by $c = \sqrt{\kappa/\tau}$ at all locations and in all directions. We define the dimensionless number, $H(\mathbf{x}, t) := |\mathbf{a}(\mathbf{x}, t)|/c$, and classify the flow at position \mathbf{x} and time t as

$$\begin{cases} \text{subcritical} & \text{if } H < 1 \\ \text{critical} & \text{if } H = 1 \\ \text{supercritical} & \text{if } H > 1 \end{cases} \tag{13}$$

In subcritical regions, diffusive transport dominates, and the solute species can propagate upstream against the flow. Upstream solute propagation is impossible in supercritical regions where advection dominates.

We introduce a 2-tuple of conservation fields,

$$\mathbf{u} = \begin{Bmatrix} cu \\ \tau \mathbf{q} \end{Bmatrix}, \tag{14}$$

and assume that the material parameters τ and k are known functions of space and independent of time. Next, we construct the corresponding 2-tuples of spacetime fluxes and spacetime sources:

$$\mathbf{F} = \left\{ \begin{array}{l} u \star dt + (u\mathbf{a} + \mathbf{q}) \star d\mathbf{x} \\ \tau \mathbf{q} \star dt + (\tau \mathbf{q} \otimes \mathbf{a} + u\mathbf{K}) \star d\mathbf{x} \end{array} \right\}, \tag{15a}$$

$$\mathbf{S} = \left\{ \begin{array}{l} s_u \\ -\mathbf{q} \end{array} \right\} \Omega. \tag{15b}$$

Combining (10) and (15), and recalling that \mathbf{a} is divergence-free by assumption, we obtain a forms-based governing system whose coefficients match and that is equivalent to (12):

$$d\mathbf{F} - \mathbf{S} = \left\{ \begin{array}{l} \dot{u} + \nabla u \cdot \mathbf{a} + \nabla \cdot \mathbf{q} - s_u \\ \tau [\dot{\mathbf{q}} + (\nabla \mathbf{q}) \mathbf{a} + \mathbf{K} \nabla u] + \mathbf{q} \end{array} \right\} \Omega = \left\{ \begin{array}{l} 0 \\ \mathbf{0} \end{array} \right\} \Omega \text{ on } \mathcal{D}. \tag{16}$$

3.2. Spacetime discontinuous Galerkin finite element formulation

Let \mathcal{P}_h be a disjoint partition of the spacetime analysis domain \mathcal{D} into N open subdomains (spacetime elements), \mathcal{Q}_α , such that $\mathcal{P}_h = \{\mathcal{Q}_\alpha\}_{\alpha=1}^N$. We use the mesh \mathcal{P}_h to construct a discrete discontinuous Galerkin solution space for paired scalar and vector fields that is piecewise continuous on \mathcal{D} with jumps admissible across all spacetime element boundaries:

$$\mathcal{V}_h := \left\{ \mathbf{w} = \left\{ \begin{array}{l} v \\ \mathbf{p} \end{array} \right\} : v|_{\mathcal{Q}} \in \mathcal{P}_{\mathcal{Q}}^r; p^i|_{\mathcal{Q}} \in \mathcal{P}_{\mathcal{Q}}^s \quad \forall \mathcal{Q} \in \mathcal{P}_h \right\} \tag{17}$$

in which $\mathcal{P}_{\mathcal{Q}}^p$ is the space of polynomial functions of order p on \mathcal{Q} . Although we typically use $r = s$, distinct polynomial orders for the scalar and vector fields are possible. We write $\mathcal{V}_h^{\mathcal{Q}} := \mathcal{V}_h|_{\mathcal{Q}}$ to describe the continuous discrete solution space within each spacetime element, $\mathcal{Q} \in \mathcal{P}_h$.

In view of the discontinuous structure of \mathcal{V}_h , we must address both the diffuse and jump parts of the exterior derivative operator. To this end, we introduce a *target flux function*, denoted as \mathbf{F}^* , on the jump set, $\mathcal{J} := \bigcup_{\mathcal{Q} \in \mathcal{P}_h} \partial \mathcal{Q}$, and expand the governing system (9) as

$$\left\{ \begin{array}{l} (d\mathbf{F} - \mathbf{S})|_{\mathcal{Q}} = \mathbf{0} \\ (\mathbf{F}^* - \mathbf{F})|_{\partial \mathcal{Q}} = \mathbf{0} \end{array} \right\} \quad \forall \mathcal{Q} \in \mathcal{P}_h \tag{18}$$

in which $\mathbf{F}|_{\partial \mathcal{Q}}$ is the trace of \mathbf{F} taken from the interior of \mathcal{Q} . Through suitable selections of \mathbf{F}^* on different parts of \mathcal{J} , we obtain a unified scheme for expressing initial, boundary, and Riemann conditions.

Most discontinuous Galerkin methods enforce a single set of flux jump conditions across inter-element boundaries, where the fluxes on opposing sides of the interface are functions of solution traces from the adjacent elements. Equation (18), on the other hand, generates two sets of jump conditions on each inter-element boundary segment. These express the differences between a common Riemann flux, \mathbf{F}^* , and the traces of fluxes from each adjacent element. In addition to balancing fluxes between elements, the paired jump conditions preserve characteristic structure across inter-element boundaries, thereby improving stability and reducing numerical dissipation. Convenient expressions for the Riemann fluxes are available for the linear conservation laws in this application, so there is no need to introduce approximate numerical fluxes. Section 3.3 next describes the computation of \mathbf{F}^* on the various parts of \mathcal{J} , including Riemann fluxes on inter-element boundaries.

System (18) leads directly to the global weighted residuals statement,

Problem 1 (Global weighted residuals statement)

For each $\mathcal{Q} \in \mathcal{P}_h$, find $\mathbf{u} \in \mathcal{V}_h^{\mathcal{Q}}$ such that

$$\int_{\mathcal{Q}} \mathbf{w} \wedge (d\mathbf{F} - \mathbf{S}) + \int_{\partial \mathcal{Q}} \mathbf{w} \wedge (\mathbf{F}^* - \mathbf{F}) = 0 \quad \forall \mathbf{w} \in (\mathcal{V}_h^{\mathcal{Q}})^* \tag{19}$$

in which $(\mathcal{V}_h^{\mathcal{Q}})^*$ is a discrete vector space whose members take the form, $[v \mathbf{p}]$, where v and \mathbf{p} are scalar and covector fields on \mathcal{Q} such that $v \in \mathcal{P}_{\mathcal{Q}}^r$ and $p_i \in \mathcal{P}_{\mathcal{Q}}^s$.

The Stokes theorem and restriction of (19) to the current patch, Π , generates the weak problem statement used in our implementation; cf. 2.1.

Problem 2 (Weak statement on current patch)

For each $Q \in \Pi$, find $\mathbf{u} \in \mathcal{V}_h^Q$ such that

$$-\int_Q (\mathbf{d}\mathbf{w} \wedge \mathbf{F} + \mathbf{w} \wedge \mathbf{S}) + \int_{\partial Q} \mathbf{w} \wedge \mathbf{F}^* = 0 \quad \forall \mathbf{w} \in (\mathcal{V}_h^Q)^* \tag{20}$$

Equation (20) expands to

$$\begin{aligned} &-\int_Q (\mathbf{d}v \wedge \mathbf{F}_u + v\mathbf{S}_u + \mathbf{d}\mathbf{p} \wedge \mathbf{F}_q + \mathbf{p} \wedge \mathbf{S}_q) \\ &+ \int_{\partial Q} (v\mathbf{F}_u^* + \mathbf{p} \wedge \mathbf{F}_q^*) = 0 \quad \forall [v \ \mathbf{p}] \in (\mathcal{V}_h^Q)^*. \end{aligned} \tag{21}$$

For a given Cartesian coordinate frame, we write (20) in indicial notation as

$$\begin{aligned} &-\int_Q \{vu + v_{,j}(ua^j + q^j) + \dot{p}_j \tau q^j + p_{j,m}(q^j a^m + uK^{jm}) - p_j q^j\} \Omega \\ &+ \int_{\partial Q} \{(vu^* + p_m \tau q^{*m}) \star dt + [v(u^* a^j + q^{*j}) + p_m(q^{*m} a^j + u^* K^{mj})] \star dx^j\} \\ &= 0 \quad \forall [v \ \mathbf{p}] \in (\mathcal{V}_h^Q)^*. \end{aligned} \tag{22}$$

3.3. *Jump conditions and target fluxes*

The jump condition in (18) plays a critical role in our SDG formulation where its weak enforcement in Problems 1 and 2 provides the coupling between solutions on adjacent patches, a means to preserve the characteristic structure of our hyperbolic system across element boundaries, as well as the weak enforcement of initial and boundary conditions. The definition and computation of the target flux function, \mathbf{F}^* , is therefore essential to our method. In general, the value of \mathbf{F}^* depends on location $\mathbf{x} \in \partial Q \subset \mathcal{J}$, the local orientation of ∂Q , the interior and exterior traces of the SDG solution on \mathcal{J} , and initial and boundary data.

The value of \mathbf{F}^* is determined at a location $\mathbf{x} \in \partial Q$ by substituting target values of concentration and diffusive flux, u^* and \mathbf{q}^* , into (15a). We determine u^* and \mathbf{q}^* at \mathbf{x} as follows. We first classify the local orientation of ∂Q at \mathbf{x} as causal or non-causal according to the causality condition described in 2.1. If the orientation is causal, we test whether it is part of the causal–inflow or causal–outflow boundary of Q — if the temporal basis vector, \mathbf{e}_t , points toward the interior of Q it is causal–inflow, and otherwise, causal–outflow. Thus, we have defined a disjoint partition of ∂Q into non-causal, causal–inflow, and causal–outflow parts, denoted respectively as ∂Q^{nc} , ∂Q^{ci} , and ∂Q^{co} , such that $\partial Q = \partial Q^{nc} \cup \partial Q^{ci} \cup \partial Q^{co}$.

If $\mathbf{x} \in \partial Q^{co}$, we set the target values to the corresponding interior traces from Q : $u^* = u|_{\partial Q}$ and $\mathbf{q}^* = \mathbf{q}|_{\partial Q}$ for all $\mathbf{x} \in \partial Q^{co}$. Thus, $\mathbf{F}^* = \mathbf{F}$, and the jump condition is trivially satisfied for any solution. In other words, the solution in Q is left unconstrained along its purely outflow boundary. If $\mathbf{x} \in \partial Q^{ci}$, we check whether \mathbf{x} is on the domain boundary. If $\mathbf{x} \in \partial \mathcal{D}$, then the values of u^* and \mathbf{q}^* are determined by specified initial data. Otherwise, \mathbf{x} lies on an inter-element boundary, and the target values are determined by the trace of the solution in the earlier, adjacent element. In our patch-by-patch solution method, this neighbor element can either be part of the current patch or part of an adjacent, previously solved patch. This assignment of target values guarantees that the causality principle is satisfied in that characteristic information always propagates forward in time across causal element interfaces.

The remaining possibility is that $\mathbf{x} \in \partial Q^{nc}$, either on a non-causal segment of the domain boundary or on a non-causal inter-element interface on the interior of the current patch. In the former case, where $\mathbf{x} \in \partial Q^{nc} \cap \partial \mathcal{D}$, we compute the target values according to prescribed boundary data. The number of prescribed solution components is determined by the local orientation of the spacetime boundary manifold and the characteristic velocities of waves traveling normally to the interface at \mathbf{x} . For simplicity and without loss of generality, we describe boundary conditions only for vertical facets; *i.e.*, space–time facets that are parallel to the time axis. This is sufficient for our present purposes because the boundaries in the numerical examples are all stationary.

The characteristic wave velocities are computed as eigenvalues of the Jacobian tensor containing the partial derivatives of the spatial surface flux vector with respect to the conserved variables [58]. Let \mathbf{n} be the unit normal in space directed outward from Q at \mathbf{x} , and consider a spatial rectangular Cartesian coordinate frame aligned with the normal, such that $\mathbf{e}_1 = \mathbf{n}$. The characteristic variables in this local coordinate system are $(u - q_1/c, q_2, \dots, q_d, u + q_1/c)$, and the corresponding \mathbf{e}_1 components of the wave velocities are $(a_1 - c, a_1$ (repeated $d - 1$ times), $a_1 + c)$. Inflow characteristics are identified by negative normal wave velocities. For a well-posed problem, the number of independent boundary conditions must equal the number of inflow characteristics at \mathbf{x} . This system of boundary relations must determine the inflow characteristic variables as functions of the outflow characteristic variables and prescribed data.

The magnitude and direction of the normal flow velocity, a_1 , influences the number of boundary conditions that must be specified at any $(\mathbf{x}, t) \in \partial Q^{nc} \cap \partial \mathcal{D}$. For example, when $a_1 > 0$ and $|a_1| < c$, only the first characteristic wave velocity, $a_1 - c$, is negative, and accordingly, only the first characteristic variable, $u - q_1/c$, is inflow. Thus, we must prescribe one and only one relation that expresses $u^* - q_1^*/c$, as a function of $\{q_i^*\}_{i=2}^d, u^* + q_1^*/c$ and prescribed data. The remaining characteristic variables are all outflow, and we take the interior trace from Q as the target value for each one so that the jump conditions are trivially satisfied. That is, $q_i^* = q_i|_{\partial Q}$ for $i = 2, \dots, d$ and $u^* + q_1^*/c = (u + q_1/c)|_{\partial Q}$. On the other hand, if $a_1 < 0$ and $|a_1| < c$, all of the characteristic variables except $u + q_1/c$ are inflow, so we must write d relations that express $u^* - q_1^*/c$ and $\{q_i^*\}_{i=2}^d$ as functions of $(u + q_1/c)|_{\partial Q}$ and prescribed data. To complete the specification of target values, we write $u^* + q_1^*/c = (u + q_1/c)|_{\partial Q}$, so the jump condition for the last characteristic variable is trivially satisfied. When $|a_1| > c$, all of the eigenvalues are nonzero and have the same sign. If they are all negative, then we have a pure inflow condition, and we use prescribed boundary data to specify the complete target set, $\{u^*, \mathbf{q}^*\}$. If the eigenvalues are all positive, then the boundary is pure outflow, and all target values are determined by the interior trace of the solution on ∂Q .

The last possibility is that $(\mathbf{x}, t) \in \partial Q^{nc}$ is on an inter-element interface on the interior of the current patch. In this case, we solve a local Riemann problem to determine the target values. The Riemann solution depends on the orientation of the interface, the interior traces of the solutions on both adjacent elements, the local flow velocity, and the local diffusive wave speed. To characterize the local orientation of ∂Q^{nc} , we construct a spatial rectangular Cartesian coordinate frame at $(\mathbf{x}, t) \in \partial Q^{nc}$ in which the basis vector \mathbf{e}_1 is aligned with \mathbf{n} , the unit spatial normal vector at (\mathbf{x}, t) pointing from the ‘–’ to ‘+’ side of the interface. Because the cotangent space to any non-causal manifold must be non-uniform in time, we can parameterize the cotangent space to ∂Q^{nc} at (\mathbf{x}, t) by writing $x_1 = \xi(x_2, \dots, x_d, t)$. Thus, in view of the orthogonal construction of the local spatial coordinate frame, we write $dx^1 = \lambda dt$ in the cotangent space at (\mathbf{x}, t) , in which we call $\lambda = \partial \xi / \partial t|_{\partial Q^{nc}}$ the *normal velocity of ∂Q^{nc} at (\mathbf{x}, t)* .

Now let $\{u^-, q_-^i\}_{i=1}^d$ and $\{u^+, q_+^i\}_{i=1}^d$ denote traces of the solution from the two sides of the interface. We construct the target values at \mathbf{x} in terms of the Riemann solution for $d = 3$ following the procedure described in [58]:

$$u^* = \begin{cases} u^- & \text{if } \lambda \leq a_1 - c \\ u^+ & \text{if } \lambda \geq a_1 + c \\ \frac{u^- + u^+}{2} + \sqrt{\frac{\tau}{\kappa}} \left(\frac{q_-^1 - q_+^1}{2} \right) & \text{otherwise} \end{cases} \quad (23a)$$

$$q_1^* = \begin{cases} q_-^1 & \text{if } \lambda \leq a_1 - c \\ q_+^1 & \text{if } \lambda \geq a_1 + c \\ \frac{q_-^1 + q_+^1}{2} + \sqrt{\frac{\kappa}{\tau}} \left(\frac{u^- - u^+}{2} \right) & \text{otherwise} \end{cases} \quad (23b)$$

$$q_2^* = \begin{cases} q_-^2 & \text{if } \lambda \leq a_1 \\ q_+^2 & \text{otherwise} \end{cases} \quad (23c)$$

$$q_3^* = \begin{cases} q_-^3 & \text{if } \lambda \leq a_1 \\ q_+^3 & \text{otherwise} \end{cases} \quad (23d)$$

The target values for $d = 1$ and $d = 2$ are similar, with one or both of (23c) and (23d) omitted, as appropriate.

3.4. Constrained model for non-negative concentrations

The SDG models with piecewise-constant bases ($p = 0$) are unconditionally stable and produce conservative, non-negative solutions that are free of overshoot and undershoot, even without added stabilization. However, models with $p \geq 1$ are subject to undershoot and nonphysical negative concentrations at sharp fronts that can cause algorithmic failure when combined with certain equations of state. Therefore, a robust solver for hyperbolic advection–diffusion systems must include a means to enforce non-negative concentrations. This section introduces a constrained model that enforces non-negative concentrations with limited impact on the element-wise conservation of the unconstrained SDG method. We do not enforce the constraint by modifying fluxes, because only fluxes derived from Riemann solutions deliver conservation properties that are mathematically and physically consistent with the underlying system. Instead, we enforce non-negativity as a direct inequality constraint on the concentration solution.

The patch-by-patch SDG solution scheme of 2.1 mitigates the computational expense of the constrained model in two ways. In contrast to methods where the constraint must be enforced globally, localizing the constrained optimization problem to individual patches delivers tractable problem sizes. In addition, we only enforce the non-negativity constraint in patches where it is active and incur no additional computational expense in patches where the constraint is inactive. Although we only describe the algorithmic details for linear polynomial bases ($p = 1$) in this work, our method is extensible to higher-order bases at the cost of some additional algorithmic complexity.

Let \mathbf{R} be the discrete residual vector generated by (20) for the current patch Π , let \mathcal{V}_h^Π denote the unconstrained SDG solution space on Π , and let \mathbf{R}_0 be the subvector of \mathbf{R} generated by weighting functions in the piecewise-constant subspace of \mathcal{V}_h^Π . Recalling that $\mathbf{R}_0 = \mathbf{0}$ ensures local conservation for each element in the patch, we introduce

Problem 3 (Constrained patch-level problem)

$$\min_{(u, \mathbf{q}) \in \mathcal{V}_h^\Pi} \frac{\|\mathbf{R}\|^2}{2} \quad (24a)$$

subject to

$$\mathbf{R}_0 = \mathbf{0} \quad (24b)$$

$$\inf u|_\Pi \geq 0. \quad (24c)$$

That is, we minimize the discrete residual vector from Problem 2 in a least-squares sense while enforcing conservation (24b) and non-negative concentrations (24c) on every element in the patch.

The least-squares structure of the minimization problem, (24a), and the linear form of the constraints, (24b) and (24c), guarantee that Problem 3 is convex. Therefore, it has a unique solution if

one exists. Assuming consistent inflow data, it is easy to demonstrate the existence of a piecewise-constant solution to Problem 2 that satisfies both constraints. Therefore, the solution set for Problem 3 is a non-empty set, and there exists a unique solution to the overall program, Problem 3 [59].

For piecewise-linear models, the solution infimum within an element must lie at one or more of the element’s vertices. Thus, ensuring non-negative concentrations at all element vertices ensures non-negative concentrations everywhere in the patch. Accordingly, we replace (24c) with equivalent constraints at a finite set of discrete vertex locations, $\{\mathbf{x}_k\}_{k=1}^{N_p}$:

$$u_k \geq 0; \quad k = 1, 2, \dots, N_p \tag{25}$$

in which \mathbf{x}_k are vertex locations, $u_k := u|_{\mathbf{x}_k}$ and, in view of our use of simplicial spacetime elements, $N_p = (d + 2)N_e$, where N_e is the number of elements in the patch and d is the spatial dimension.

We introduce Lagrange multipliers, $\{\mu_k\}_{k=1}^{N_p}$, corresponding to the discrete non-negativity constraints (25) and a second set of multipliers, $\{\lambda_j\}_{j=1}^{N_c}$, for the components of the equality constraint, (24b). The number of equality constraint components is $N_c = (d + 1)N_e$ in which $(d + 1)$ is the number of conserved field components and N_e is the number of elements in the patch. The constrained minimization problem for the system, (24a), (24b), and (25), is equivalent to

Problem 4 (Min-max problem)

$$\min_{(u, \mathbf{q}) \in \mathcal{V}_h^\Pi} \max_{\lambda_j \in \mathbb{R}, \mu_k \in \mathbb{R}_+} \mathcal{L}_0 := \frac{m^2 \|\mathbf{R}\|^2}{2} - \sum_{k=1}^{N_p} \mu_k u_k + m \sum_{j=1}^{N_c} \lambda_j R_{0j} \tag{26}$$

where $\mathbb{R}_+ := \mathbb{R}^+ \cup \{0\}$ and $m = |\Pi|^{-d/d+1}$ in which $\|\Pi\|$ is the spacetime measure of patch Π .

The scalings by m and m^2 in (26) provide dimensional consistency between the terms in the Lagrangian, \mathcal{L}_0 . This ensures that the relative importance of the individual terms do not vary drastically with mesh refinement and prevents ill-conditioning in the incremental linear system generated from (26); see (30) in the following.

Let α be the vector of coefficients that determine a member of \mathcal{V}_h^Π and ∇_α denote the gradient operator with respect to α . The optimality conditions for a solution to Problem 4 are

$$m^2 (\nabla_\alpha \mathbf{R})^T \mathbf{R} - \sum_{k=1}^{N_p} \mu_k \nabla_\alpha u_k + m \sum_{j=1}^{N_c} \lambda_j \nabla_\alpha R_{0j} = 0 \tag{27a}$$

$$m \mathbf{R}_0 = \mathbf{0} \tag{27b}$$

$$u_k \mu_k = 0; \quad k = 1, 2, \dots, N_p \tag{27c}$$

$$u_k, \mu_k \geq 0; \quad k = 1, 2, \dots, N_p \tag{27d}$$

We use a variant of the nonlinear complementarity (NC) method of Torres and Quintana [60] to replace the complementarity conditions, (27c) and (27d), with a single, smooth equality for each value of k . However, given their simple form, we work directly with the complementarity functions instead of introducing slack variables as in [60]. We adopt the NC function from [61],

$$\varphi(a, b, \varepsilon) = \sqrt{a^2 + b^2 + \varepsilon} - (a + b) \tag{28}$$

in which $0 \leq \varepsilon \ll 1$, and replace (27c) and (27d) with a system of equality constraints,

$$\varphi(u_k, \mu_k, \varepsilon) = 0; \quad k = 1, 2, \dots, N_p, \tag{29}$$

that is equivalent in the limit as $\varepsilon \rightarrow 0$. We used a fixed value, $\varepsilon = 1\text{e-}30$, to generate the numerical results reported in Section 4.

Let $\mathbf{K} = m\nabla_{\alpha}\mathbf{R}$, where $\nabla_{\alpha}\mathbf{R}$ is the stiffness matrix from the unconstrained Problem 2 (similar for $\mathbf{K}_0 = m\nabla_{\alpha}\mathbf{R}_0$), $\boldsymbol{\varphi}$ be the vector of NC constraint functions, and $\mathbf{H} = \nabla_{\alpha}u_k$ be the matrix whose rows are the values of the interpolation functions for u at the vertices. The linearized incremental formulation of the system that comprises (27a), (27b), and (29) is $\mathbf{A}\mathbf{X} = \mathbf{b}$, where

$$\mathbf{A} = \begin{bmatrix} \mathbf{K}^T\mathbf{K} & -\mathbf{H}^T & \mathbf{K}_0^T \\ -\nabla_u\boldsymbol{\varphi}\mathbf{H} & \nabla_{\mu}\boldsymbol{\varphi} & \mathbf{0} \\ \mathbf{K}_0 & \mathbf{0} & \mathbf{0} \end{bmatrix} \quad \mathbf{X} = \begin{Bmatrix} \Delta\boldsymbol{\alpha} \\ \Delta\boldsymbol{\mu} \\ \Delta\mathbf{A} \end{Bmatrix}$$

$$\mathbf{b} = - \begin{Bmatrix} m\mathbf{K}^T\mathbf{R} - \mathbf{H}^T\boldsymbol{\mu} + \mathbf{K}_0^T\mathbf{A} \\ \boldsymbol{\varphi} \\ m\mathbf{R}_0 \end{Bmatrix} \tag{30}$$

in which \mathbf{A} , \mathbf{u} , and $\boldsymbol{\mu}$ are column vectors containing λ_j , u_k , and μ_k . We use a Newton–Raphson procedure to solve (30).

The active condition for the non-negativity constraint at vertex k implies that the constraint, $u_k = 0$, is active in some spacetime neighborhood of k . This condition, in turn, implies that the diffusive flux, \mathbf{q} , should vanish at k . Accordingly, we impose the additional requirement, $\mathbf{q} = \mathbf{0}$, wherever the non-negativity constraint is active. This leads to

Problem 5 (Constrained problem with zero-flux condition)

$$\min \frac{\mathbf{R}^T\mathbf{R}}{2} \tag{31a}$$

subject to

$$\mathbf{R}_0 = \mathbf{0} \tag{31b}$$

$$u_k \geq 0 \quad k = 1, 2, \dots, N_p \tag{31c}$$

$$\mathbf{q}_k = \mathbf{0} \quad \text{if } \mu_k > 0; k = 1, 2, \dots, N_p \tag{31d}$$

Our numerical experience indicates that failure to enforce the constraint (31d) leads to unstable propagation of u into regions where the solution should be zero. Because conditional constraints, such as (31d), are difficult to implement directly [62], we adopt a simplified method instead. We first use (30) to solve Problem 4, and then compute corrected linear distributions for \mathbf{q} in elements with active constraints, consistent with setting $\mathbf{q}_k = \mathbf{0}$ at vertices where $\mu_k > 0$ while retaining unaltered values of \mathbf{q} at the other vertices. Although (31b) in Problem 5 implies element-wise conservation, our simplified implementation of (31d) might compromise that property in elements where the non-negativity constraint is active. In a typical example, the maximum correction to any vertex component of \mathbf{q} was 0.01. However, the maximum conservation error, as measured by $|\mathbf{R}_0|$, was only $4e-7$. Thus, the error in element-wise conservation due to our simplified treatment of (31d) was negligible in this example.

3.5. Adaptive error indicator

We use heuristic, residual-based, a posteriori error indicators to drive the adaptive solution processes in the numerical examples presented in this work. Separate error contributions are computed for the volume residuals and the boundary jump residuals for each element Q , and there are separate error indicators for the scalar and vector conservation relations, denoted respectively by E_1^Q and E_2^Q ; cf. (16). Although the heuristic error indicators appear to work well, we plan to add indicators based on L^1 error norms as a natural alternative in future applications of this work.

User-specified parameters, $\{\underline{E}_{\gamma}, \bar{E}_{\gamma}\}_{\gamma=1,2}$, define target ranges for the element-wise errors: $\underline{E}_{\gamma} \leq E_{\gamma}^Q \leq \bar{E}_{\gamma}$. For each new patch, we first compute the linear SDG solution for Problem 2 and the error indicators, E_{γ}^Q , for each element Q in the patch. We then use these to drive the h -adaptive

SDG solution scheme described in 2.2. If $E_\gamma^Q > \bar{E}_\gamma$ for either value of γ , we flag element Q as needing refinement and reject the patch solution. If $E_\gamma^Q < \underline{E}_\gamma$ for both values of γ , we flag element Q as coarsenable. Otherwise, we flag Q as neither needing refinement nor coarsening. In either of the latter cases, we check the non-negativity condition at the patch's element vertices. If there are no negative values, we accept and store the patch solution. Otherwise, we solve Problem 5 using Newton–Raphson iterations based on (30) and the simplified treatment of the conditional constraint (31d). We globalize the Newton iterations with the non-monotone Armijo line search method of Grippo, Lampariello, and Lucidi [63]. Upon convergence, we compute the error indicators, E_γ^Q , once again and repeat the h -adaptive algorithm. In this way, we only solve the nonlinear constrained problem in patches where it is needed and only after a suitable level of mesh refinement has been completed.

4. NUMERICAL EXAMPLES

The following examples demonstrate the accuracy, robustness, and efficiency of our SDG model, the adaptive solution scheme, and our direct enforcement of the non-negativity constraint. In all cases, we use a broken SDG solution space based on complete, linear spacetime polynomials ($p = 1$) for each solution component within each spacetime element. We do not present convergence studies for our model, but in a previous study of an SDG model for Maxwell–Cattaneo diffusion, Miller *et al.* [26] demonstrate optimal convergence rates of $O(h^{p+1})$ in L^1 -error norms for both solution components, u and \mathbf{q} . Similar results were obtained with a mixed SDG model for linearized elastodynamics in [51].

4.1. Rotating flow

We use the rotating flow problem from [1] to test the accuracy of our SDG solution scheme. The problem setup is illustrated in Figure 2(a). The flow is circular about the center of a square domain, $\Omega = [-1, 1] \times [-1, 1]$, with a prescribed velocity field, $\mathbf{a} = -x_2\mathbf{e}_1 + x_1\mathbf{e}_2$. The material parameters are $k = 10^{-7}$ and $\tau = 1$, consistent with negligible diffusivity. We prescribe homogeneous concentration and flux conditions on inflow boundary segments, as indicated in Figure 2(a), and free conditions on outflow segments. In addition, we enforce $u(x_1, 0) = \sin(\pi x_1)$ on the slit, $x_2 = 0^+$, $x_1 > 0$ at all times.[‡]

As discussed in [7], classical upwind schemes suffer excessive crosswind diffusion on this problem. We solve the problem non-adaptively with a uniform space mesh of 1800 triangular elements. The color plot in Figure 2(b) shows a nearly steady solution at $t = 8$. The absence of a discernible jump between the advected solution and the prescribed values along the slit and the circular structure of the color plot suggest that the solution is accurate and free of spurious crosswind diffusion. Due to the smooth loading, no negative concentrations were generated by the unconstrained SDG model, so the non-negativity constraint was never invoked during the computation of this solution.

4.2. Skew advection

This example, previously studied in [24], demonstrates the accuracy of the SDG model in a problem where the non-negativity constraint is active in some patches within an unstructured, adaptive spacetime mesh. The problem schematic in Figure 3 shows a square domain on which we prescribe the flow velocity, $\mathbf{a} = \cos\alpha\mathbf{e}_1 + \sin\alpha\mathbf{e}_2$, with $\alpha = \pi/4$. The material parameters are $k = 10^{-7}$ and $\tau = 1$. The flow is supercritical everywhere. On the inflow boundary, we enforce $\mathbf{q} = \mathbf{0}$ and prescribe the concentration as shown in the schematic. The solution is free on the outflow boundary. We solve the problem adaptively, starting with a space mesh that comprises only three triangular elements. The adaptive error tolerances were $\underline{\epsilon} = 3.2 \times 10^{-6}$ and $\bar{\epsilon} = 4.8 \times 10^{-6}$.

[‡]Note that this condition is equivalent to prescribing the inflow characteristic variable, $u - q_1/c$, as a function of an outflow characteristic variable, $u + q_1/c$, and problem data. Specifically, $u - q_1/c = 2 \sin(\pi x_1) - (u + q_1/c)$.

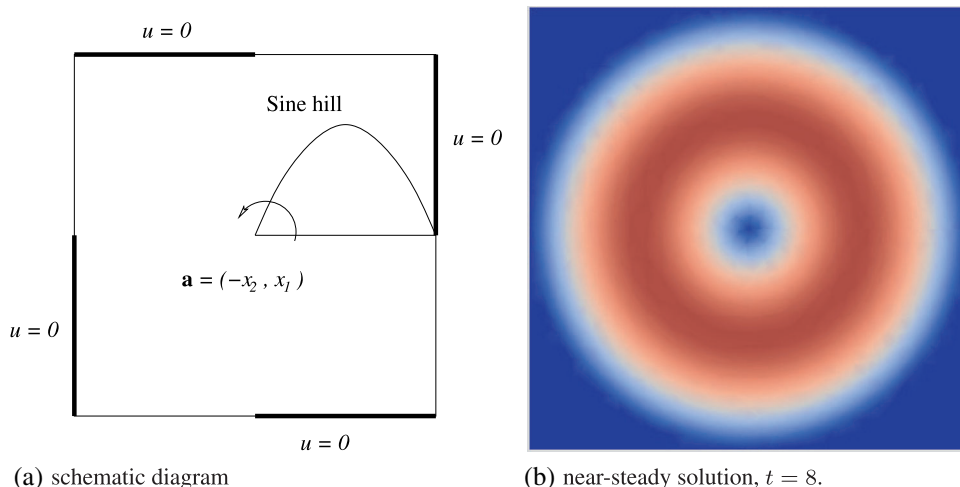


Figure 2. Rotating flow test problem: (a) schematic diagram in which bold lines indicate inflow boundary segments, based on Fig. 15 in [24]; (b) broadband concentration contours for near-steady spacetime discontinuous Galerkin solution with $u \in [0, 1]$, $u = 0 \rightarrow$ blue, $u = 1 \rightarrow$ red.




Figure 3. Schematic of skew advection problem, based on Fig. 13 in [24].

Figure 4(a) displays the concentration solution field at time $t = 0.8$. The grid lines depict the intersection of the $t = 0.8$ plane with the unstructured, tetrahedral spacetime mesh, so the resulting partition of the spatial domain is not representative of the actual element geometry and mesh quality. The sharp concentration front is well resolved by the adaptive scheme, and the non-negativity constraint is duly enforced with no oscillations and monotonic response across the front. Figure 4(b) shows the final steady-state solution and the final space mesh. No spurious oscillations and no excessive crosswind diffusion are observed, even at this late stage of the solution.

Although it is difficult to assess element quality from the tetrahedral mesh intersections in Figure 4(a), the final space mesh in Figure 4(b) confirms that our adaptive spacetime mesh generator maintains a well-graded mesh with good element aspect ratios throughout the simulation. Larger elements appear in regions where the solution is uniform, while the mesh is duly refined to capture the transition across the front. Because the duration of a spacetime element is proportional to its spatial diameter, h , the spacetime volume occupied by an element is $O(h^{d+1})$. Thus, we observe that most, in fact nearly all, of the elements in the adapted spacetime mesh participate directly in resolving the sharp wavefronts in this example. This demonstrates the efficiency of our asynchronous

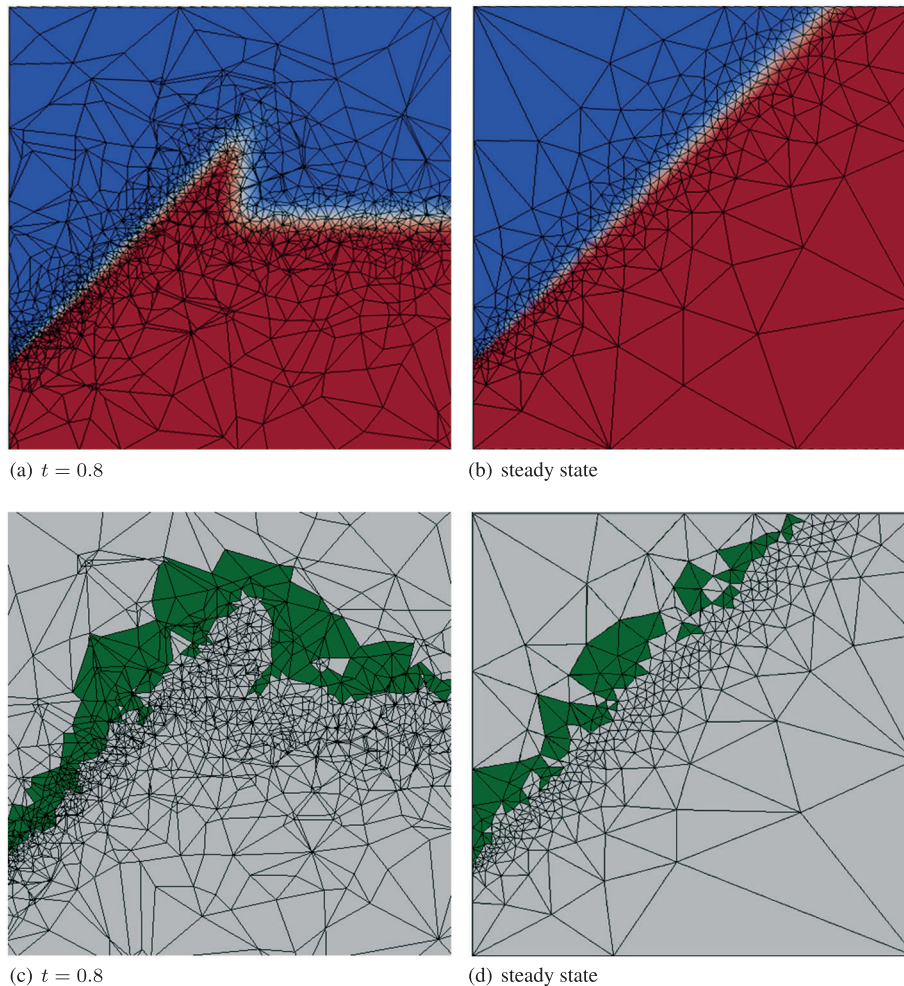


Figure 4. Skew advection: (a), (b) concentration solutions for homogeneous initial data using linear color map for $u \in [0, 1]$ with $u = 0 \rightarrow$ blue, $u = 1 \rightarrow$ red; (c), (d) distributions of patches with active non-negativity constraints for uniform initial concentration, $u_0 = 0.001$.

adaptive algorithm relative to synchronous adaptive methods that must conform to either a global or regional time-step size.

The sudden application of the inflow condition for concentration at $t = 0$ creates a sharp front that slowly diffuses to a smoother profile as the solute advects with the background flow. The remnant of the $t = 0$ front is visible in the upper right quadrant in Figure 4(a), and we see in Figure 4(b) a persistent layer emanating at 45° from the point of discontinuity in the inflow data on the left edge of the domain. Small but finite diffusion causes this layer to thicken with increasing distance from the inflow boundary, even in the continuum weak solution. Notably, the finest mesh refinement is found at the upstream end where the layer thickness is smallest. Given that the continuum solution is discontinuous where the layer meets the inflow boundary, part of the thickening inevitably arises from numerical error that can be controlled, but not eliminated, by using tighter adaptive tolerances. There is also some numerical error associated with the non-negativity constraint. Overall, however, the numerical error is well controlled in this solution.

Our adaptive solution procedure generated 455,465 spacetime patches, of which 59,786 were rejected by the adaptive algorithm. The non-negativity constraint was enforced in 115,523 patches of which 23,676 were rejected, and the cost of solving a constrained patch is about 4.4 times that of solving a patch using the unconstrained SDG procedure. Thus, selective activation of the non-negativity constraint resulted in a 49% savings in computational effort relative to a similar SDG




Figure 5. Schematic of Poiseuille flow problem, based on Figure 6 in [24].

solution procedure in which the constraint algorithm is invoked in all patches. That is, selective enforcement of the constraint cuts execution time roughly in half in this example.

We reran the model with a uniform initial concentration field, $u_0 = 0.001$, over the entire spatial domain, to demonstrate selective enforcement of the non-negativity constraint. These inhomogeneous initial data are consistent with practical problems where the ambient concentration is often nonzero; their effect is to reduce somewhat the number of patches where the constraint is active. Figure 4(c) and (d) highlight the patches where the non-negativity constraint is active. We observe in Figure 4(c) a layer of patches with active constraints along one edge of the band of mesh refinement that indicates the front's current position. The area where the non-negativity constraint is active does not include regions with the most intense mesh refinement, so relatively few patches are affected by the constraint. In the steady condition depicted in Figure 4(d), we observe that the constraint is most active immediately downstream of the inflow point of discontinuity, where the front is sharpest. Diffusion gradually softens the front as it convects downstream, and accordingly, the pattern of patches with active constraints becomes increasingly sparse as the front approaches the outflow boundary. Although our method for selective constraint enforcement appears to work well, we do observe limited propagation of the active constraint condition into regions adjacent to fronts where the solution should be zero. Indeed, the nonzero concentrations in these regions are very small.

The limited problem size on an individual patch enables cost-effective enforcement of the non-negativity constraint in our patch-by-patch solution scheme. With or without selective activation, the cost of a constrained SDG solution grows linearly with the total number of spacetime patches solved, whereas the cost of a constrained implicit time-marching scheme grows geometrically with the sum of the numbers of solution degrees of freedom and Lagrange multipliers in a time-step or spacetime slab. Thus, use of a similar technique for direct enforcement of the non-negativity constraint over a global time step or time slab would be much more expensive and possibly impractical in conventional solvers.

4.3. Poiseuille flow

This example, based on a similar problem studied by Gomez *et al.* in [24], involves Poiseuille flow between two parallel plates with a pollutant injected along a short boundary segment with inward diffusive flux, $q = -\mathbf{q} \cdot \mathbf{n} = 0.02$, as shown in Figure 5.** Zero normal flux conditions are imposed along the top plate and the rest of the lower plate. The steady, incompressible background velocity field is given by $\mathbf{a}(x_1, x_2) = x_2(2 - x_2)\mathbf{e}_1$. The boundary conditions on the left and right sides of the domain for subcritical and supercritical flows are distinct, as described in the following. We interpret the analysis domain as a segment of a larger physical domain with no sources outside the analysis domain. This leads us to adopt boundary conditions in the subcritical case that differ from those used in [24].

We first choose the same material parameters, $\kappa = 2.0$ and $\tau = 1.0$, as in [24]. This results in a subcritical flow everywhere in the domain, with $H_{\max} = 1/\sqrt{2}$. The model advanced in [24] imposes zero concentration and tangential flux on the upstream (left) boundary and zero normal diffusive flux ($\mathbf{q} \cdot \mathbf{n} = 0$) on the downstream (right) boundary. These conditions generate a

**This condition is equivalent to prescribing the inflow characteristic variable, $u - q_1/c$, as a function of an outflow characteristic variable, $u + q_1/c$, and problem data. Specifically, $u - q_1/c = u + q_1/c - 0.04/c$.



Figure 6. Subcritical Poisseuille flow: concentration solutions using linear color map for $u \in [0, 0.01]$ with $u = 0 \rightarrow$ blue, $u \geq 0.01 \rightarrow$ red. (a) $t=2.1$; (b) $t=3.3$; (c) $t=8$; (d) $t=12.5$; (e) Approaching steady state.

mathematically well-posed problem that is suitable for numerical studies. However, they are unphysical according to our interpretation of the analysis domain. For example, in the subcritical case, upstream diffusion should eventually produce positive solute concentrations on the left boundary. Given the proximity of the right boundary to the solute source, we also expect non-vanishing normal diffusion across the downstream boundary at later times. The effects of suppression of positive concentrations along the left boundary and of normal diffusion across the right boundary are evident in the solution displayed in Figure 7 of [24].

To address these issues, we specify vanishing inflow characteristic values and transmitting (free) conditions for the outflow characteristic values to obtain the following relations between the target values (* superscript) and the interior traces (undecorated) of the conserved fields:

$$u^* - q_1^*/c = 0 \quad (\text{upstream and downstream boundaries}) \quad (32a)$$

$$q_2^* = 0 \quad (\text{upstream boundary only}) \quad (32b)$$

$$u^* + q_1^*/c = u + q_1/c \quad (\text{upstream and downstream boundaries}) \quad (32c)$$

Equations (32a) and (32b) are the inflow conditions, and (32c) is the transmitting outflow condition.

We solve the problem adaptively, starting with a space mesh of only four triangular elements and with the adaptive tolerances set to $\underline{\epsilon} = 1.6 \times 10^{-5}$ and $\bar{\epsilon} = 2.4 \times 10^{-5}$. This requires a total of




Figure 7. Supercritical Poiseuille flow with non-negativity constraint: concentration solutions using linear color map for $u \in [0, 0.02]$ with $u = 0 \rightarrow$ blue, $u \geq 0.02 \rightarrow$ red. (a) $t=4$; (b) $t=8$; (c) $t=12$; (d) $t=16$.

2,120,035 spacetime patch solutions during the course of the simulation. We plot the SDG concentration solution at selected times in Figure 6. We cannot compare directly our results with those reported in [24], because the distinct boundary conditions in the two models generate qualitatively different solutions. Nonetheless, the trends in our solution seem reasonable for our intended model.

We observe wave reflections off the plates at early times, $t = 2.1$ and $t = 3.5$, that generate a pattern of interacting waves that is typical of hyperbolic systems. This behavior would not arise in corresponding parabolic systems based on Fickian diffusion models. As the solution proceeds, we observe the effects of both advective and diffusive transport, with most of the pollutant propagating downstream, while a lesser amount diffuses upstream. Consistent with our expectations, we obtain non-vanishing concentration on the left boundary and non-vanishing normal diffusion on the right boundary at later times. We observe no spurious reflections at the upstream and downstream boundaries, so our transmitting conditions for the outflow characteristics appear to work well. Visual inspection of Figures 6(d) and 6(e) suggests that our SDG solution approaches a steady state within the finite analysis domain in which the influx of solute through the lower plate equals the total outflux through the upstream and downstream boundaries.

We also solved the problem with material parameters, $\kappa = 0.01$ and $\tau = 1$, that result in supercritical flow over the solution domain, except for layers along the plates where the speed of the background flow is small. We solved this problem adaptively, starting from the same space mesh of four triangular elements. The transition between the zones with zero and positive concentrations is sharper in this advection-dominated problem. Therefore, we expect undershoot and spurious oscillations at this transition due to Gibbs effects in numerical simulations, unless corrective remedies are applied. Indeed, our SDG solutions exhibit negative concentrations and oscillatory artifacts for $p = 1$ when the non-negativity constraint is deactivated. We emphasize, however, that the amplitudes of the undershoot and oscillations are very small relative to the scale of the overall solution, rendering them invisible unless we use a narrow-band color map to highlight small-amplitude

features. It is therefore unknown whether the solutions displayed with a broadband color map in [24] suffer the same problem.

Although small, the amplitudes of the spurious oscillations generated by our SDG model were too large to be attributed to finite-precision arithmetic; we found them to be insensitive to mesh refinement and to increase with higher-order elements. Because negative concentrations are non-physical and inadmissible in some applications, we implemented the constraint algorithm described in 3.4 to obtain a reliable and robust method for supercritical flows (to date, we have not encountered this problem in solutions for subcritical flows). With this modification in place, our SDG solutions are non-negative, monotone, and free of spurious oscillations.

Figure 7 displays SDG concentration results at selected times for the supercritical case obtained with the non-negativity constraint enforced and adaptive error tolerances, $\underline{\epsilon} = 0.8 \times 10^{-6}$ and $\bar{\epsilon} = 1.2 \times 10^{-6}$. This simulation required a total of 1,582,577 spacetime patches. At early times (e.g., $t = 4$ in Figure 7), most of the pollutant is in the subcritical layer close to the lower plate where diffusive transport dominates and some upstream pollutant propagation is evident. However, advection increases as the solute moves upward, and we observe increased downstream propagation. The flow becomes supercritical a short distance from the wall, and a sharp interface forms because upstream propagation is no longer possible. This effect is more pronounced at later times, $t = 8$ and $t = 12$, where advective effects predominate. A steady state is attained, as shown in Figure 7(d), where the transport is primarily downstream advection with limited diffusion in the vertical direction.

Figure 8(a) and (b) show SDG concentration solutions subject to the non-negativity constraint using a narrow-band color map that emphasizes low-amplitude detail. Even in this stringent display format, there is no evidence of undershoot, oscillations, or spurious upwind transport. On the other hand, spurious oscillations degrade the SDG solution adjacent to and upstream of the region of high




Figure 8. Concentration solutions for supercritical Poiseuille flow using narrow-band color map to accentuate low-amplitude solution features: (a), (b) with and (c), (d) without non-negativity constraint. Color map is linear for $u \in [0, 0.002]$ with $u \leq 0 \rightarrow$ blue, and $u \geq 0.002 \rightarrow$ red.




Figure 9. Schematic of converging–diverging nozzle flow problem. TBC = transmitting boundary condition.

concentration when the non-negativity constraint is not enforced; as in Figures 8(c) and 8(d). Thus, we conclude that the non-negativity constraint eliminates spurious oscillations while preserving sharp fronts between regions of zero and high concentration.

4.4. Flow in a converging–diverging nozzle

As a final example, we model an oscillating flow in a converging–diverging nozzle. The problem schematic in Figure 9 consists of a long channel with a converging–diverging section in its center. Designers use nozzles of this form in turbines, jet engines, etc to accelerate fluid in the narrow region. A chemical solute is injected into the nozzle in its neck region as shown. We assume an unsteady, oscillating background flow, rather than a more typical steady flow, to make the problem more challenging. We enforce transmitting boundary conditions on the left and right sides of the domain, and solve the problem for two sets of fluid properties to explore the distinct character of subcritical and supercritical flows.

Ideally, the assumed background flow should satisfy continuity, incompressibility, and no-slip boundary conditions on the walls of the nozzle. Unfortunately, an analytical solution that satisfies all three of these in our nozzle-shaped domain is not readily available. For this reason, we resort to a flow solution that satisfies continuity and incompressibility, but not the no-slip conditions:

$$\mathbf{a}(x, y) = \begin{cases} 0.5U\mathbf{e}_x & |x| > 2 \\ U\mathbf{e}_x & |x| < 1 \\ U(1/x)\mathbf{e}_x + U(y/x^2)\mathbf{e}_y & x \in (1, 2) \\ -U(1/x)\mathbf{e}_x - U(y/x^2)\mathbf{e}_y & x \in (-2, -1). \end{cases} \quad (33)$$

in which the flow magnitude, U , is a time-dependent function, $U(t) = 0.5 \sin(\pi t)$.

The fluid properties for the subcritical case are $\kappa = 2$, $\tau = 1$, while for the supercritical case, they are $\kappa = 0.01$, $\tau = 1$. The subcritical problem is solved with adaptive error tolerances, $\underline{\epsilon} = 1.6 \times 10^{-5}$ and $\bar{\epsilon} = 2.4 \times 10^{-5}$, and required a total of 1,827,012 spacetime patches. The corresponding values for the supercritical case are $\underline{\epsilon} = 3.2 \times 10^{-6}$ and $\bar{\epsilon} = 4.8 \times 10^{-6}$, with the simulation involving solution of 680,327 patches. Although the error tolerances for the supercritical case demand greater accuracy, the localized nature of the supercritical solution allows the adaptive SDG solver to meet this requirement with fewer patches than in the subcritical case by restricting mesh refinement to regions where the solution gradients change rapidly.

Figure 10 displays a sequence of concentration solutions for the subcritical case. At early times, such as $t = 0.3$, the flow diffuses in all directions, resulting in rapid smoothing of the initial wavefront. We observe additive interference after the wavefront collides with the top wall ($t = 0.5$), resulting in concentrations on the domain interior that can surpass the maximum concentration along the boundary. This behavior contrasts with the response of parabolic advection–diffusion models that obey the maximum principle. Diffusing wavefronts advect along the flow direction at $t = 1.6$, while propagating transversely to form an intricate pattern of interacting waves.

The flow direction reverses at $t = 2$, causing a wavefront to move into the left end of the nozzle, as seen at $t = 2.6$. At long times, diffusion suppresses wave structure in the transverse direction, while a region of increasingly high concentration forms in the central region of the nozzle. The region of highest concentration oscillates to the left and right in response to the magnitude and direction of the background flow and the instantaneous concentration gradients. As the flow in a cycle starts to the left (e.g., at $t = 6$), the higher concentration region advects to the left, while a low-concentration

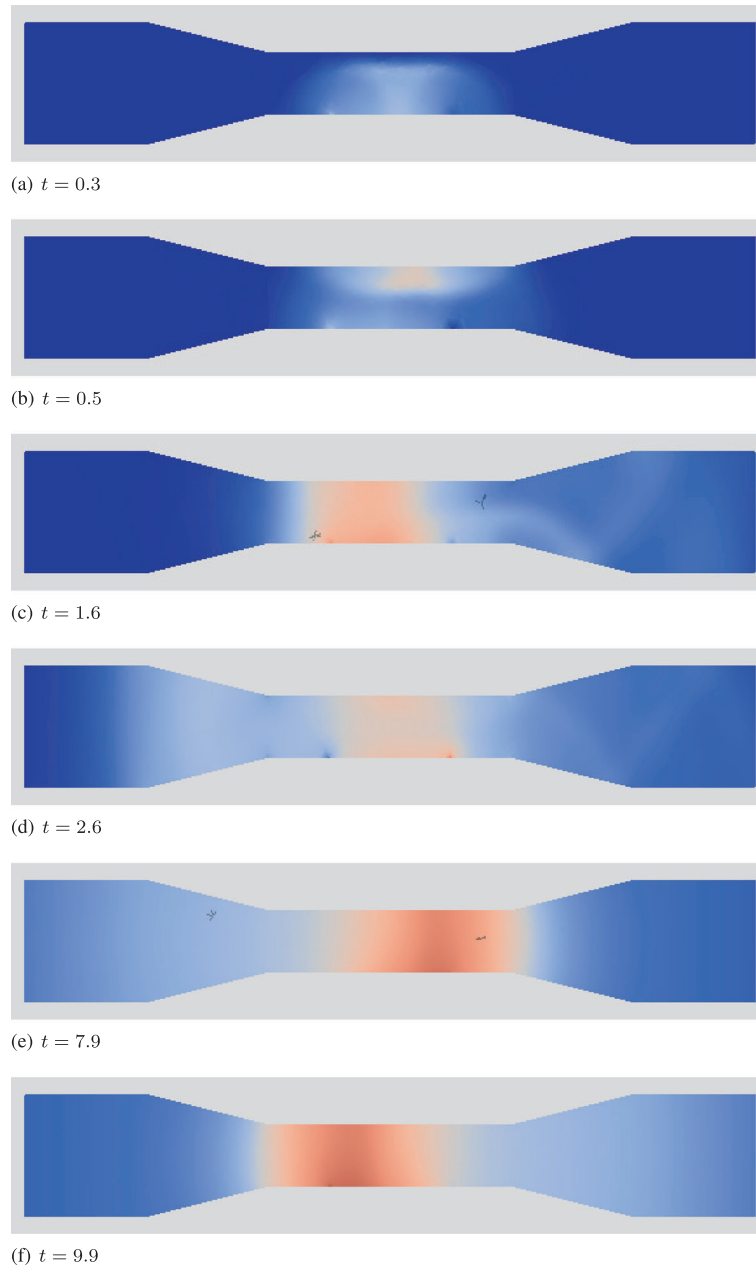


Figure 10. Subcritical flow in a nozzle: concentration solutions using linear color map for $u \in [0, 0.04]$ with $u = 0 \rightarrow$ blue, $u \geq 0.04 \rightarrow$ red. (a) $t=0.3$; (b) $t=0.5$; (c) $t=1.6$ (d) $t=2.6$; (e) $t=7.9$; (f) $t=9.9$.

region forms on the right side of the nozzle. The high-concentration region moves slowly to the right, opposite to the flow direction, at time $t = 7.9$ because of the stronger gradient in that direction and the slowing leftward flow that allows diffusive transport to dominate. Similarly, as the flow direction reverses to the right, the higher concentration region shifts to the left in the confined region toward the end of the cycle at $t = 9.9$. The solution enters a repeating cycle in synchrony with the oscillating background flow after about three cycles. Our use of transmitting boundary conditions at the upstream and downstream boundaries prevents spurious reverse flow when the flow direction changes. Figure 11 displays a sequence of solutions for the supercritical case where the wavefronts are generally sharper. The solute propagates more slowly in the vertical direction than in the flow direction, because now, the primarily longitudinal flow dominates diffusive transport.

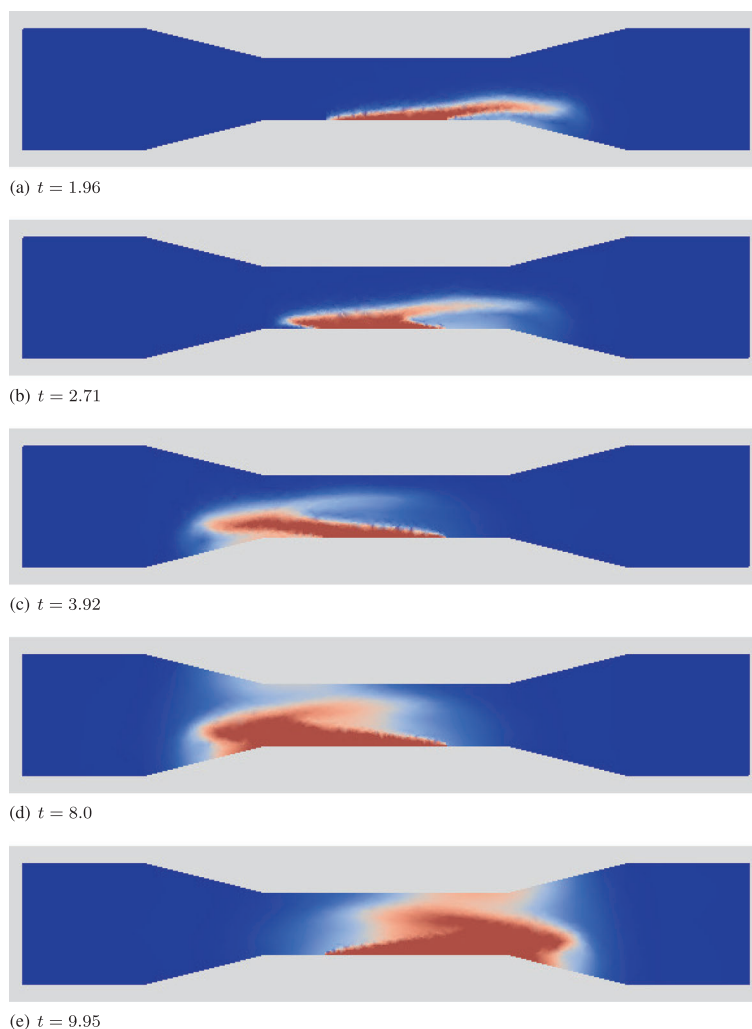


Figure 11. Supercritical flow in a nozzle: concentration solutions using linear color map for $u \in [0, 0.02]$ with $u = 0 \rightarrow$ blue, $u \geq 0.02 \rightarrow$ red.

Advection also dominates when the flow direction reverses, and in contrast to the subcritical flow, the solute species is confined within the nozzle at all times because the flow reverses before the solute can advect to either boundary. Even for long times, such as $t = 8.0$ and $t = 9.95$, the concentration is zero outside the narrowed section of the nozzle. This confinement contrasts with parabolic advection–diffusion models, where a small amount of solute is present everywhere in the domain at all times because of the infinite propagation speed implied by Fickian diffusion. The average solute concentration increases with time as more solute is injected into the domain and none can exit through the boundaries. Thus, unlike the subcritical case, the supercritical solution never enters a repeating cycle.

5. CONCLUSIONS

We have presented a new SDG method for hyperbolic advection–diffusion systems based on a frame-invariant version of the Maxwell–Cattaneo diffusion model. We demonstrated several intrinsic advantages of an asynchronous, patch-wise SDG solution scheme for this application, including element-wise conservation, the use of Riemann fluxes that require no additional upwinding or stabilization, scalability based on linear computational complexity, and an efficient, fine-grained h -adaptive solution scheme. We also proposed and demonstrated a new technique for enforcing

directly a non-negativity constraint on patch-level concentration solutions while preserving the element-wise conservation property.

Several extensions of this work are of practical interest. Given sufficient regularity of the local solution, we know that the efficiency of causal, asynchronous SDG models increases with polynomial order. In contrast to finite volume methods, higher-order SDG stencils remain compact and do not complicate the treatment of boundary conditions. Notably, and in contrast to most other schemes, the optimal convergence and linear computational complexity of SDG solvers deliver better computational efficiency for higher values of p . Thus, higher-order extensions of our model that enable hp -adaptive implementations should further improve the method's performance. The main challenge involves the increasing technical complexity of the implementation of the non-negativity constraint. For $p = 1$, as in this work, the knowledge that element minima must lie at one of the vertices greatly simplifies the problem. The problem is still tractable for $p = 2$, but here, we must consider the possibility that in addition to the vertices, the minimum concentration might occur on the element interior or on any of the element faces or edges. In addition, the location where the minimum value is attained might change as the solution evolves. The continuum constraint becomes more complex for $p > 2$ with the possibility of multiple extrema in each part of the element. Thus, it might become expedient to approximate the continuum constraint with a set of discrete constraints at key locations, such as vertices and quadrature points.

In this work, we adopt a divergence-free velocity field as part of the problem data. However, many practical applications require simultaneous solutions of flow and advection–diffusion problems, with possible relaxation of the incompressibility constraint on the flow field. These extensions are addressed in [23] for a synchronous discontinuous Galerkin model, and we foresee no fundamental difficulties in following a similar approach for the present SDG solver. Calibration and validation of the Maxwell–Cattaneo diffusion model with experimental and observational data are another practical concern. In particular, identification of the relaxation time τ is important to obtain predictive capabilities for problems in which the time scales of advection and diffusion have comparable magnitudes.

Beyond the present application to advection–diffusion systems, the overall approach we propose for the non-negativity constraint shows promise for extension to a more general class of applications that involves constrained hyperbolic partial differential equations. The asynchronous SDG scheme on causality-constrained spacetime meshes replaces a sequence of large-scale constrained optimization problems, defined globally for a single time-step or spacetime slab, with a series of modest-scale, patch-wise problems where the cost of more sophisticated constrained optimization algorithms is not prohibitive. The ability to selectively activate an expensive constraint algorithm on only those patches where it is needed further enhances the efficiency of constrained SDG solution methods. This efficiency could enable direct enforcement of difficult nonlinear and non-smooth constraints without resorting to indirect methods, such as flux limiting or matrix modification, that can introduce unintended side effects. Thus, the general topic of SDG solvers for constrained hyperbolic systems appears to be ripe for continuing research and new applications.

ACKNOWLEDGEMENT

This work was supported in part by United States National Science Foundation grant OCI 09-48393 EAGER ARRA.

REFERENCES

1. Gómez H, Colominas I, Navarrina F, Casteleiro M. A finite element formulation for a convection–diffusion equation based on Cattaneo's law. *Computer Methods in Applied Mechanics and Engineering* 2007; **196**(9):1757–1766.
2. Nielsen DR, van Genuchten MT, Biggar JW. Water flow and solute transport processes in the unsaturated zone. *Water Resources Research* 1986; **22**(9S):89S–108S.
3. Douglas J Jr., Furtado F, Pereira F. On the numerical simulation of waterflooding of heterogeneous petroleum reservoirs. *Computational Geosciences* 1997; **1**(2):155–190.
4. Gerritsen MG, Durlofsky LJ. Modeling fluid flow in oil reservoirs. *Annual Review of Fluid Mechanics* 2005; **37**: 211–238.

5. Chen Z, Cockburn B, Gardner CL, Jerome JW. Quantum hydrodynamic simulation of hysteresis in the resonant tunneling diode. *Journal of Computational Physics* 1995; **117**(2):274–280.
6. Gardner CL. The quantum hydrodynamic model for semiconductor devices. *SIAM Journal on Applied Mathematics* 1994; **54**(2):409–427.
7. Brooks AN, Hughes TJ. Streamline upwind/Petrov–Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier–Stokes equations. *Computer Methods in Applied Mechanics and Engineering* 1982; **32**(1):199–259.
8. Hughes TJ, Franca LP, Hulbert GM. A new finite element formulation for computational fluid dynamics: VIII. the Galerkin/least-squares method for advective–diffusive equations. *Computer Methods in Applied Mechanics and Engineering* 1989; **73**(2):173–189.
9. Hundsdorfer W, Verwer JG. *Numerical Solution of Time-dependent Advection–diffusion–reaction Equations*, Vol. 33. Springer: Berlin, 2003.
10. Crank J. *The Mathematics of Diffusion*. Oxford University Press: New York, 1979.
11. Joseph DD, Preziosi L. Heat waves. *Reviews of Modern Physics* 1989; **61**(1):41.
12. Godoy S, Garcia-Colin L. From the quantum random walk to classical mesoscopic diffusion in crystalline solids. *Physical Review E* 1996; **53**(6):5779.
13. Montroll EW, Scher H. Random walks on lattices. IV. Continuous-time walks and influence of absorbing boundaries. *Journal of Statistical Physics* 1973; **9**(2):101–135.
14. Montecinos GI, Müller LO, Toro EF. Hyperbolic reformulation of a 1D viscoelastic blood flow model and ADER finite volume schemes. *Journal of Computational Physics* 2014; **266**:101–123.
15. Brady J, Koch D. Dispersion in porous media. In *Disorder and Mixing*, vol. 152, Guyon E, Nadal J-P, Pomeau Y (eds)., NATO ASI Series E. Springer: Dordrecht, 1988; 107–122.
16. Compte A, Metzler R. The generalized Cattaneo equation for the description of anomalous transport processes. *Journal of Physics A: Mathematical and General* 1997; **30**(21):7277.
17. Chandrasekharaiah D. Hyperbolic thermoelasticity: a review of recent literature. *Applied Mechanics Reviews* 1998; **51**(12):705–729.
18. Christov C, Jordan P. Heat conduction paradox involving second-sound propagation in moving media. *Physical Review Letters* 2005; **94**(15):154301.
19. Carey G, Tsai M. Hyperbolic heat transfer with reflection. *Numerical Heat Transfer, Part A Applications* 1982; **5**(3):309–327.
20. Yang H. Characteristics-based, high-order accurate and nonoscillatory numerical method for hyperbolic heat conduction. *Numerical Heat Transfer, Part B Fundamentals* 1990; **18**(2):221–241.
21. Shen W, Han S. Numerical solution of two-dimensional axisymmetric hyperbolic heat conduction. *Computational Mechanics* 2002; **29**(2):122–128.
22. Gómez H, Colominas I, Navarrina F, Casteleiro M. A discontinuous Galerkin method for a hyperbolic model for convection–diffusion problems in CFD. *International Journal for Numerical Methods in Engineering* 2007; **71**(11):1342–1364.
23. Gómez H, Colominas I, Navarrina F, Casteleiro M. A mathematical model and a numerical model for hyperbolic mass transport in compressible flows. *Heat and Mass Transfer* 2008; **45**(2):219–226.
24. Gomez H, Colominas I, Navarrina F, París J, Casteleiro M. A hyperbolic theory for advection–diffusion problems: mathematical foundations and numerical modeling. *Archives of Computational Methods in Engineering* 2010; **17**(2):191–211.
25. Wu W, Li X. Application of time discontinuous Galerkin finite element method to heat wave simulation. *International Journal of Heat and Mass Transfer* 2006; **49**:1679–1684.
26. Miller ST, Haber RB. A spacetime discontinuous Galerkin method for hyperbolic heat conduction. *Computer Methods in Applied Mechanics and Engineering* 2008; **198**(2):194–209.
27. Cockburn B, Hou S, Shu CW. The Runge–Kutta local projection discontinuous Galerkin finite element method for conservation laws. iv. the multidimensional case. *Mathematics of Computation* 1990; **54**(190):545–581.
28. Cockburn B, Shu CW. The Runge–Kutta discontinuous Galerkin method for conservation laws V: multidimensional systems. *Journal of Computational Physics* 1998; **141**(2):199–224.
29. Nakshatrala K, Valocchi AJ. Non-negative mixed finite element formulations for a tensorial diffusion equation. *Journal of Computational Physics* 2009; **228**(18):6726–6752.
30. Perthame B, Shu CW. On positivity preserving finite volume schemes for Euler equations. *Numerische Mathematik* 1996; **73**(1):119–130.
31. Zhang X, Shu CW. Positivity-preserving high order finite difference WENO schemes for compressible Euler equations. *Journal of Computational Physics* 2012; **231**(5):2245–2258.
32. Zhang X, Xia Y, Shu CW. Maximum-principle-satisfying and positivity-preserving high order discontinuous Galerkin schemes for conservation laws on triangular meshes. *Journal of Scientific Computing* 2012; **50**(1):29–62.
33. Berthon C, Marche F. A positive preserving high order VFRoe scheme for shallow water equations: a class of relaxation schemes. *SIAM Journal on Scientific Computing* 2008; **30**(5):2587–2612.
34. Xing Y, Zhang X, Shu CW. Positivity-preserving high order well-balanced discontinuous Galerkin methods for the shallow water equations. *Advances in Water Resources* 2010; **33**(12):1476–1493.
35. Parent B. Positivity-preserving high-resolution schemes for systems of conservation laws. *Journal of Computational Physics* 2012; **231**(1):173–189.

36. Balsara DS. Self-adjusting, positivity preserving high order schemes for hydrodynamics and magnetohydrodynamics. *Journal of Computational Physics* 2012; **231**(22):7504–7517.
37. Suresh A. Positivity-preserving schemes in multidimensions. *SIAM Journal on Scientific Computing* 2000; **22**(4):1184–1198.
38. Lipnikov K, Svyatskiy D, Vassilevski Y. A monotone finite volume method for advection–diffusion equations on unstructured polygonal meshes. *Journal of Computational Physics* 2010; **229**(11):4017–4032.
39. Wang S, Yuan G, Li Y, Sheng Z. A monotone finite volume scheme for advection–diffusion equations on distorted meshes. *International Journal for Numerical Methods in Fluids* 2012; **69**(7):1283–1298.
40. Liska R, Shashkov M. Enforcing the discrete maximum principle for linear finite element solutions of second-order elliptic problems. *Communications in Computational Physics* 2008; **3**(4):852–877.
41. Nakshatrala K, Nagarajan H. A numerical methodology for enforcing maximum principles and the non-negative constraint for transient diffusion equations, 2012. arXiv preprint arXiv:1206.0701.
42. Nakshatrala K, Mudunuru M, Valocchi A. A numerical framework for diffusion-controlled bimolecular-reactive systems to enforce maximum principles and the non-negative constraint. *Journal of Computational Physics* 2013; **253**:278–307.
43. Lowrie RB, Roe PL, van Leer B. Space–time methods for hyperbolic conservation laws. In *Barriers and Challenges in Computational Fluid Dynamics*, vol. 6, ICASE/LaRC Interdisciplinary Series in Science and Engineering. Kluwer: Dordrecht, 1998; 79–98.
44. Falk RS, Richter GR. Explicit finite element methods for symmetric hyperbolic equations. *SIAM Journal of Numerical Analysis* 1999; **36**(3):935–952.
45. Yin L, Acharia A, Sobh N, Haber RB, Tortorelli DA. A spacetime discontinuous Galerkin method for elastodynamics analysis. In *Discontinuous Galerkin Methods: Theory, Computation and Applications*, Cockburn B, Karniadakis G, Shu CW (eds). Springer Verlag: Berlin, 2000; 459–464.
46. Abedi R, Chung SH, Erickson J, Fan Y, Garland M, Guoy D, Haber RB., Sullivan J, White S, Zhou Y. Space–time meshing with adaptive refinement and coarsening. *Proceedings 20th Annual ACM Symposium on Computational Geometry*: Association for Computing Machinery, Brooklyn, New York, 2004; 300–309.
47. Palaniappan J, Haber R, Jerrard RL. A spacetime discontinuous Galerkin method for scalar conservation laws. *Computer Methods in Applied Mechanics and Engineering* 2004; **193**(33):3607–3631.
48. Abedi R, Petravic B, Haber RB. A space–time discontinuous Galerkin method for linearized elastodynamics with element-wise momentum balance. *Computer Methods in Applied Mechanics and Engineering* 2006; **195**(25): 3247–3273.
49. Abedi R, Haber B, Thite S, Erickson J. An h-adaptive spacetime-discontinuous Galerkin method for linear elastodynamics. *European Journal of Computational Mechanics/Revue Européenne de Mécanique Numérique* 2006; **15**(6):619–642.
50. Palaniappan J, Miller ST, Haber RB. Sub-cell shock capturing and spacetime discontinuity tracking for nonlinear conservation laws. *International Journal for Numerical Methods in Fluids* 2008; **57**:1115–1135.
51. Miller ST, Kraczek B, Haber RB, Johnson DD. Multi-field spacetime discontinuous Galerkin methods for linearized elastodynamics. *Computer Methods in Applied Mechanics and Engineering* 2009; **199**:34–47.
52. Abedi R, Hawker MA, Haber RB, Matouš K. An adaptive spacetime discontinuous Galerkin method for cohesive models of elastodynamic fracture. *International Journal for Numerical Methods in Engineering* 2009; **1**:1–42.
53. Abedi R, Haber RB. Riemann solutions and spacetime discontinuous Galerkin method for linear elastodynamic contact. *Computer Methods in Applied Mechanics and Engineering* 2014; **270**(0):150–177.
54. Thite S. A unified algorithm for adaptive spacetime meshing with nonlocal cone constraints. *Proceedings of the 21st European Workshop on Computational Geometry*, Eindhoven, Netherlands, 2005; 1–4.
55. Abraham R, Marsden JE, Ra?iu TS, Cushman R. *Foundations of Mechanics*. Benjamin/Cummings Publishing Company Reading: Massachusetts, 1978.
56. Fleming WH. *Functions of Several Variables*, Vol. 14. Springer: New York, 1977.
57. Spivak M. *Calculus on Manifolds*. Perseus Books: Cambridge, Massachusetts, 1965.
58. LeVeque RJ. *Finite Volume Methods for Hyperbolic Problems*, Vol. 31. Cambridge University Press: Cambridge, 2002.
59. Ruszczyński AP. *Nonlinear Optimization*, Vol. 13. Princeton University Press: Princeton, 2006.
60. Torres GL, Quintana VH. Optimal power flow by a nonlinear complementarity method. *Proceedings of the 21st International Conference on Power Industry Computer Applications*, IEEE, Santa Clara, California, 1999; 211–216.
61. Chen B, Harker PT. A non-interior-point continuation method for linear complementarity problems. *SIAM Journal on Matrix Analysis and Applications* 1993; **14**(4):1168–1190.
62. Gelle E, Faltings B. Solving mixed and conditional constraint satisfaction problems. *Constraints* 2003; **8**(2):107–141.
63. Grippo L, Lampariello F, Lucidi S. A nonmonotone line search technique for Newton’s method. *SIAM Journal on Numerical Analysis* 1986; **23**(4):707–716.