An optimization-based approach for high-order accurate discretization of conservation laws with discontinuous solutions

M. J. Zahr[a,b,1,*], P.-O. Persson[a,b,2]

[a]*Mathematics Group, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, United States*
[b]*Department of Mathematics, University of California, Berkeley, Berkeley, CA 94720, United States*

Abstract

This work introduces a novel discontinuity-tracking framework for resolving discontinuous solutions of conservation laws with high-order numerical discretizations that support inter-element solution discontinuities, such as discontinuous Galerkin or finite volume methods. The proposed method aims to align inter-element boundaries with discontinuities in the solution by deforming the computational mesh. A discontinuity-aligned mesh ensures the discontinuity is represented through inter-element jumps while smooth basis functions interior to elements are only used to approximate smooth regions of the solution, thereby avoiding Gibbs' phenomena that create well-known stability issues. Therefore, very coarse high-order discretizations accurately resolve the piecewise smooth solution throughout the domain, provided the discontinuity is tracked. Central to the proposed discontinuity-tracking framework is a discrete PDE-constrained optimization formulation that simultaneously aligns the computational mesh with discontinuities in the solution and solves the discretized conservation law on this mesh. The optimization objective is taken as a combination of the the deviation of the finite-dimensional solution from its element-wise average and a mesh distortion metric to simultaneously penalize Gibbs' phenomena and distorted meshes. It will be shown that our objective function satisfies two critical properties that are required for this discontinuity-tracking framework to be practical: (1) possesses a local minima at a discontinuity-aligned mesh and (2) decreases monotonically to this minimum in a neighborhood of approximately $\mathcal{O}(h)$, whereas other popular discontinuity indicators fail to satisfy the latter. Another important contribution of this work is the observation that traditional reduced space PDE-constrained optimization solvers that repeatedly solve the conservation law at various mesh configurations are not viable in this context since the stability issues caused by Gibbs' phenomena may make it impossible to solve the discrete conservation law on non-aligned meshes. Therefore, we advocate a gradient-based, full space solver where the mesh and conservation law solution converge to their optimal values simultaneously and therefore never require the solution of the discrete conservation law on a non-aligned mesh. The merit of the proposed method is demonstrated on a number of one- and two-dimensional model problems including the L^2 projection of discontinuous functions, Burgers' equation with a discontinuous source term, transonic flow through a nozzle, and supersonic flow around a bluff body. We demonstrate optimal $\mathcal{O}(h^{p+1})$ convergence rates in the L^1 norm for up to polynomial order $p = 6$ and show that accurate solutions can be obtained on extremely coarse meshes.

Keywords: r-adaptivity, shock tracking, high-order methods, discontinuous Galerkin, full space PDE-constrained optimization, transonic and supersonic flow

1. Introduction

Even with the continued advancement of computational hardware and algorithms, it is apparent that important developments are still required to improve the predictive capability of the computational simu-

---

*Corresponding author
Email addresses: `mjzahr@lbl.gov` (M. J. Zahr), `persson@berkeley.edu` (P.-O. Persson)
[1]Luis W. Alvarez Postdoctoral Fellow, Computational Research Division, Lawrence Berkeley National Laboratory
[2]Associate Professor, Department of Mathematics, University of California, Berkeley

lation tools that are used for design. For example, it is widely believed that higher fidelity is required for problems with propagating waves, turbulent fluid flow, nonlinear interactions, and multiple scales [51]. This has resulted in a significant interest in high-order accurate methods, such as discontinuous Galerkin (DG) methods [15, 26], which have the potential to produce fundamentally more reliable solutions than traditional numerical schemes.

One of the remaining problems with the DG methods is their sensitivity to under-resolved features, in particular for non-linear problems where the spurious oscillations often cause a break-down of the numerical solvers. This is exacerbated for problems with shocks, where the natural dissipative mechanisms introduced by DG methods through jump terms is insufficient to stabilize the solution for high-order approximations. Since shocks are present in many important problems in fields such as aerospace, astrophysics, and combustion, the lack of efficient ways to handle them is a fundamental challenge for the wide adoption of these new numerical methods.

One of the most straight-forward approaches for stabilizing high-order methods in the presence of shocks is to identify the elements close to the shocks using a sensor, and reduce the corresponding approximation orders [7, 12]. This increases the numerical dissipation which helps stabilize the solution, and in the extreme case of piecewise constant solutions will completely avoid any oscillations. An obvious drawback with this strategy is that it will typically result in a first-order scheme in the affected elements. This can partially be addressed by using h-adaptivity and remeshing [16, 28]; however, generating efficient anisotropic meshes in 3D is still a difficult task and the resulting meshes typically have a very large number of degrees of freedom around the shocks. Further issues with these methods include moving shocks, which require constant remeshing and solution transfer, and a difficulty for implicit solvers with large timesteps since order reduction strategies typically cannot be differentiated for a nonlinear Newton solver. A related class of methods use limiters based on the ENO/WENO schemes [22, 35, 30] to limit the high-order representation of the solution inside the elements [13, 42, 57, 33, 55]. While these schemes can be remarkably robust and stable for very strong shocks, they are suffer from spurious oscillations, have difficulties with fully unstructured elements, particularly in 3D, and are typically only used for explicit timestepping.

Another popular way to stabilize problems with shocks is to explicitly add artificial viscosity to the governing equations [49]. In [6, 7, 24], this was applied to the DG method using element-wise viscosity based on the residual of the Euler equations. To address issues with the consistency of the resulting scheme, a proper discretization of the artificial viscosity was used in [41] together with a resolution-based indicator based on the decay-rate of the high-order terms. The resulting scheme could smoothly resolve shocks using subcell resolution, and the continuous nature of both the sensor and the viscosity made it possible to get fully converged steady-state solutions and accurate time-stepping using implicit methods. Further developments of this approach include some refinements of the parameter selection [31], the solution of an auxiliary PDE for the viscosity field [5], and the extension to implicit time-stepping for transient problems with smooth viscosity fields [39]. An alternative approach for the shock sensor is to somehow incorporate the physics of the problem [29, 9, 36]. Overall, these artificial viscosity approaches are widely used, but like limiting they require anisotropic h-adaptivity to be competitive [2, 53], which again is difficult in 3D, particularly for moving shocks, and produces large meshes.

In this work, we propose a fundamentally different approach based on the idea of shock tracking or shock fitting [45, 46, 8, 23, 48, 56, 47, 4, 43, 20, 38] using r-adaptivity. For generality, we will use the term *discontinuity tracking* in this work, although our motivating applications involve shock waves. Since the solution basis in DG methods (as well as the finite volume methods) naturally support discontinuities between each element, it is in principle possible to keep an unmodified high-order DG discretization for problems with strong discontinuities, provided the element faces can be aligned with the discontinuities of the problem. However, this is a very challenging problem since the unknown discontinuity surfaces can have complex geometries and if the elements are even just slightly misaligned, the discrete equations cannot be solved due to spurious oscillations. Instead, we employ a new optimization formulation that aims to align faces of the computational mesh with discontinuities and solve the discrete conservation law. The optimization formulation is classified as fully discrete r-adaptive PDE-constrained optimization since the discretized conservation law defines nonlinear equality constraints and the state vector and nodal positions

of the computational mesh are optimization variables. In this framework, the objective function must possess a local minima at a discontinuity-aligned mesh and must monotonically approach this minima in a neighborhood of radius approximately $h/2$ for a solver to successfully locate such a minima. Here, we propose a simple robustness-based discontinuity indicator that will be experimentally shown to satisfy both criteria, while a number of intuitive physics- and error-based discontinuity indicators fail to satisfy the monotonicity property. Finally, we use a full space approach to solve the PDE-constrained optimization problem whereby the discrete solution of the conservation law and the nodal positions of the mesh *simultaneously* converge to their optimal values without ever requiring the solution of the discrete conservation law on a non-aligned mesh. The popular alternative, the reduced space approach, is not a viable solver as it requires the solution of the discrete conservation law on non-aligned meshes, which may be impossible due to Gibbs' instabilities. We demonstrate our method can produce high-order accurate solutions on remarkably coarse meshes, for a range of problems in 1D and 2D.

The remainder of the paper is organized as follows. Section 2 introduces the governing system of steady, inviscid conservation laws on a parametrized domain and its discretization using a discontinuous Galerkin method. Since deformation of the computational domain is a key ingredient in this work, a parametrized domain deformation is introduced at the continuous level and the conservation laws are recast on a fixed reference domain. In the end, Section 2 reduces the conservation law to a discrete nonlinear system of equations that depend on the discrete state vector and computational mesh, which is the point of departure for the proposed discontinuity-tracking framework, introduced in Section 3. The proposed PDE-constrained, *r*-adaptive formulation is presented in Section 3.1. The objective function used in this work, introduced in Section 3.2, is a combination of new discontinuity indicator and a standard mesh distortion indicator. In this section, we show the indicator possesses a minima at a discontinuity-aligned mesh and decreases monotonically to this minima, whereas the latter condition fails for some existing indicators. Section 3.3 discusses the full space solver required for the proposed discontinuity-tracking and a detailed discussion on why a traditional reduced space PDE-constrained optimization approach is not sufficient. Section 3.4 discusses a number of practical details including implementation, initialization of the nonlinear optimization solver, and robustness and efficiency of the method. Finally, Section 4 presents a number of one- and two-dimensional test problems that demonstrate optimal $\mathcal{O}(h^{p+1})$ convergence up to $p = 6$ and accurate solutions are obtained on extremely coarse meshes.

2. Governing equations and high-order numerical discretization

Consider a general system of N_c conservation laws, defined on the physical domain $\Omega \subset \mathbb{R}^d$,

$$\nabla \cdot \mathcal{F}(U) = 0 \quad \text{in} \ \ \Omega, \tag{1}$$

where $U(x) \in \mathbb{R}^{N_c}$ is the solution of the system of conservation laws at $x \in \Omega$ and $\mathcal{F}(U) \in \mathbb{R}^{N_c \times d}$ is the physical flux. While we solely consider steady conservation laws, i.e., Ω is a d-dimensional spatial domain and the solution is independent of time, the conservation law in (1) and the proposed numerical method encapsulate the unsteady case where Ω is the d-dimensional space-time domain and the spatial domain is $(d-1)$-dimensional. Furthermore, we assume the solution U contains discontinuities, in which case the conservation law (1) holds away from these discontinuities.

The proposed optimization-based method for tracking discontinuities, introduced in Section 3, is built upon existing numerical discretizations that possess the following properties: 1) represents a stable and convergent discretization of the conservation law in (1), 2) allows for deformation of the computational domain, and 3) employs a solution basis that supports discontinuities between computational cells or elements. While the proposed method will be valid for any numerical discretization that satisfies these requirements, such as finite volume or discontinuous Galerkin methods, we focus on high-order DG methods given their potential to deliver accurate solutions on very coarse discretizations, provided the discontinuities are tracked.

The remainder of this section will detail the discretization of the conservation law (1) using DG such that it reduces to the discrete form

$$\boldsymbol{r}(\boldsymbol{u}, \boldsymbol{x}) = \boldsymbol{0} \tag{2}$$
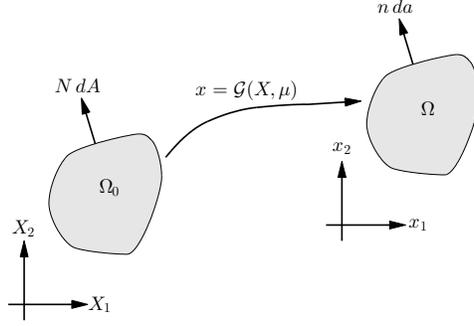
Figure 1: Mapping between reference and physical domains.

where $\boldsymbol{u} \in \mathbb{R}^{N_u}$ is the discrete representation of the conservation law state U, $\boldsymbol{x} \in \mathbb{R}^{N_x}$ is the discrete representation of the conservation law domain Ω, and $\boldsymbol{r}$ is the discretized conservation law. A standard nodal, isoparametric DG method is used for the discretization, with special attention given to treatment of the domain deformation and the numerical fluxes, which must be selected carefully in a tracking method since the inter-element jumps do not tend to zero on the discontinuity surface. The discrete form of the conservation law (16) will be the point of departure for discussion of the proposed optimization-based tracking method and the primary reason a different underlying discretization, e.g., finite volume, could be used in place of DG with only minor required modifications. Section 3 will mention the modifications required to use a finite volume method as the underlying discretization.

2.1. Transformed conservation law from deformation of physical domain

Before introducing a discretization of (1) it is convenient to explicitly treat deformations to the domain of the conservation law Ω. Given the divergence-form of the conservation law in (1), i.e., no explicit time derivative term, the domain deformation is handled directly. That is, the physical domain can be taken as the result of a parametrized diffeomorphism applied to a reference domain (Figure 1)

$$\Omega = \mathcal{G}(\Omega_0, \mu), \tag{3}$$

where $\Omega_0 \subset \mathbb{R}^d$ is a fixed reference domain, $\mu \in \mathbb{R}^{N_\mu}$ is a vector of parameters, and $\mathcal{G} : \mathbb{R}^d \times \mathbb{R}^{N_\mu} \to \mathbb{R}^d$ is the parametrized diffeomorphism defining the domain mapping. Under the domain mapping (3), the conservation law becomes

$$\nabla \cdot \mathcal{F}(U) = 0 \quad \text{in} \ \ \mathcal{G}(\Omega_0, \mu). \tag{4}$$

For convenience, the conservation law on the physical domain Ω is transformed to a conservation law on the reference domain Ω_0 using the procedure in [40] to yield

$$\nabla_X \cdot F(u, \mu) = 0 \quad \text{in} \ \ \Omega_0 \tag{5}$$

where ∇_X denotes spatial derivatives with respect to the reference domain Ω_0 with coordinates X. The transformed state vector, u, and flux, F, take the form

$$u = g_\mu U, \qquad F(u, \mu) = g_\mu \mathcal{F}(g_\mu^{-1} u) G_\mu^{-T} \tag{6}$$

where $G_\mu(X) = \frac{\partial}{\partial X} \mathcal{G}(X, \mu)$ is the deformation gradient of the domain mapping and $g_\mu(X) = \det G_\mu(X)$ is the Jacobian. The transformed conservation law in (5)-(6) is, by construction, identical to the Arbitrary Lagrangian-Eulerian formulation of the steady conservation law in (4). For details regarding the derivation of the transformed equations, the reader is referred to [40].

While a numerical scheme and implementation can be designed around (4) or (5), we chose the formulation on the reference domain in (5) as it provides a convenient framework to handle domain deformation. In this case, the governing equation depends on the domain deformation *solely* through the transformed flux function and definition of the state vector. This detail leads to a clean implementation, particularly when computing required derivative quantities, as discussed in Section 3.4.

4

2.2. Discontinuous Galerkin discretization of transformed conservation law

In this section, a standard nodal discontinuous Galerkin method [15, 26] is used to discretize the transformed conservation law (5) with special attention given to the numerical fluxes. Let $\mathcal{E}_{h,p}$ represent a discretization of the reference domain Ω_0 into non-overlapping, potentially curved, computational elements, where h is a mesh element size parameter and p is the polynomial order associated with the curved elements. The DG construction begins by considering the element-wise weak form of the governing equation in (5) that results from multiplication by a test function ψ, integration over a single element $K \in \mathcal{E}_{h,p}$, and application of the divergence theorem

$$\int_K \nabla_X \cdot F(u, \mu) \cdot \psi \, dV = \int_{\partial K} \psi \cdot F(u, \mu) N \, dA - \int_K F(u, \mu) : \nabla_X \psi \, dV = 0, \tag{7}$$

where N is the outward normal to the surface ∂K. The global weak form, upon which DG methods are built, arises from the summation of the local residuals over all elements in $\mathcal{E}_{h,p}$

$$\sum_{K \in \mathcal{E}_{h,p}} \int_{\partial K} \psi \cdot F(u, \mu) N \, dA - \int_{\Omega_0} F(u, \mu) : \nabla_X \psi \, dV = 0. \tag{8}$$

At this point, a numerical flux, F^*, is introduced in the first term of (8)

$$\sum_{K \in \mathcal{E}_{h,p}} \int_{\partial K} \psi \cdot F^*(u, \mu, N) \, dA - \int_{\Omega_0} F(u, \mu) : \nabla_X \psi \, dV = 0 \tag{9}$$

to ensure the flux is single-valued along ∂K where u is multi-valued. In DG methods, the numerical flux must be chosen such that it is consistent to ensure local conservation [14], i.e., $F^*(\bar{u}, \mu, N) = F(\bar{u}, \mu)N$ where $\bar{u}(X)$ is single-valued, and leads to a stable discretization. Given that inter-element jumps tend to zero under refinement, the particular choice of numerical flux is less important, provided it is stable and consistent [15, 52]. However, recall the goal of this work is to *track* discontinuities with element faces, which implies inter-element jumps at the discontinuity will not tend to zero under refinement and care must be taken when choosing a numerical flux such that it is consistent with the governing equation at discontinuities. In this work, Roe's method [44] with the Harten-Hyman entropy fix [23] is used for the numerical fluxes at interior faces and the appropriate boundary conditions determine the numerical fluxes on faces that intersect $\partial \Omega$. The Harten-Hyman entropy fix is used to ensure non-physical rarefaction shocks do not appear.

To establish the finite-dimensional form of (9), we introduce the isoparametric finite element space of piecewise polynomial functions associated with the mesh $\mathcal{E}_{h,p}$:

$$\mathcal{V}_{h,p} = \left\{ v \in [L^2(\Omega_0)]^{N_c} \mid v|_K \circ \mathcal{T}_K \in [\mathcal{P}(K_0)]^{N_c} \; \forall K \in \mathcal{E}_{h,p} \right\}$$

$$\hat{\mathcal{V}}_{h,p}(\mu) = \left\{ v \in [L^2(\Omega)]^{N_c} \mid v|_{\mathcal{G}(K,\mu)} \circ \mathcal{T}_{\mathcal{G}(K,\mu)} \in [\mathcal{P}(K_0)]^{N_c} \; \forall K \in \mathcal{E}_{h,p} \right\}$$

where $\mathcal{P}_p(K_0)$ is the space of polynomial functions of degree at most $p \geq 1$ on the parent element K_0 and $K = \mathcal{T}_K(K_0)$ defines a mapping from the parent element to element $K \in \mathcal{E}_{h,p}$. For notational brevity, we assume all elements map from a single parent element. The parametrized space $\hat{\mathcal{V}}_{h,p}(\mu)$ will be used in later sections to define an L^2 projection onto a piecewise polynomial space attached to the physical domain $\Omega = \mathcal{G}(\Omega_0, \mu)$. The Petrov-Galerkin weak form in (9) becomes: find $u_{h,p} \in \mathcal{V}_{h,p}$ such that for all $\psi_{h',p'} \in \mathcal{V}_{h',p'}$

$$r_{h,p}^{h',p'}(u_{h,p}, \mu) := \sum_{K \in \mathcal{E}_{h,p}} \int_{\partial K} \psi_{h',p'} \cdot F^*(u_{h,p}, \mu, N) \, dA - \int_{\Omega_0} F(u_{h,p}, \mu) : \nabla_X \psi_{h',p'} \, dV = 0, \tag{10}$$

where $r_{h,p}^{h',p'}$ is the finite-dimensional residual of (9) corresponding to the trial space $\mathcal{V}_{h,p}$ and test space $\mathcal{V}_{h',p'}$. In this work, these spaces are taken to be the same and we have the Galerkin weak form: find $u_{h,p} \in \mathcal{V}_{h,p}$

such that for all $\psi_{h,p} \in \mathcal{V}_{h,p}$

$$\sum_{K \in \mathcal{E}_{h,p}} \int_{\partial K} \psi_{h,p} \cdot F^*(u_{h,p},\, \mu,\, N)\, dA - \int_{\Omega_0} F(u_{h,p},\, \mu) : \nabla_X \psi_{h,p}\, dV = 0. \tag{11}$$

The Petrov-Galerkin residual in (10) will be used in Section 3.2.2 to define an accuracy-based discontinuity indicator. To obtain the discrete form of (11), we introduce a basis $\{\varphi_i\}_{i=1}^{N_p}$ for $\mathcal{P}_p(K_0)$, where N_p is the number of basis functions. While this can be any valid basis, a flexibility afforded by the DG framework, it is convenient for the proposed shock tracking method to work with a *nodal basis*. Define a set of nodes $\{\xi_j\}_{j=1}^{N_p}$ within the parent element K_0 and let φ_i be the Lagrange polynomial associated with nodes ξ_i over element K_0 with the property $\varphi_i(\xi_j) = \delta_{ij}$. For brevity, we introduce notation for the Lagrange basis defined on the reference element $K \in \mathcal{E}_{h,p}$: $\varphi_i^K(X) = \varphi_i(\mathcal{T}_K^{-1}(X))$. The finite-dimensional solution, $u_{h,p}$, in each element $K \in \mathcal{E}_{h,p}$ is written in terms of its discrete expansion coefficients as

$$u_{h,p}(X)|_K = \sum_{i=1}^{N_p} \boldsymbol{u}_i^K \varphi_i^K(X), \tag{12}$$

where $\boldsymbol{u}_i^K \in \mathbb{R}^{N_c}$ is the solution at node i of element K. The finite-dimensional test functions, $\psi_{h,p}$ are expanded similarly

$$\psi_{h,p}(X)|_K = \sum_{i=1}^{N_p} \boldsymbol{\psi}_i^K \varphi_i^K(X), \tag{13}$$

where $\boldsymbol{\psi}_i^K \in \mathbb{R}^{N_c}$ is the test function at node i of element K. Under the isoparametric assumption, the physical coordinates are expanded in the nodal basis as

$$x_{h,p}(X)|_K = \sum_{i=1}^{N_p} \boldsymbol{x}_i^K \varphi_i^K(X), \tag{14}$$

where $\boldsymbol{x}_i^K \in \mathbb{R}^d$ is the coordinate of node i of element K. It will also be convenient in later sections to introduce the corresponding nodal coordinates in the reference domain, $\boldsymbol{X}_i^K \in \mathbb{R}^d$. From this expansion, the deformation gradient, G, and Jacobian, g, required to define the transformed state vector and fluxes are

$$G_{h,p}(X)|_K = \sum_{i=1}^{N_p} \boldsymbol{x}_i^K \frac{\partial \varphi_i^K}{\partial X}(X)$$
$$g_{h,p}(X)|_K = \det G_{h,p}(X)|_K. \tag{15}$$

Therefore, in the isoparametric setting, the mapping to the physical domain is completely determined from the nodal positions of each element $\boldsymbol{x}_i^K$. To ensure the domain mapping is continuous, nodal positions co-located in the reference domain are required to be co-located in the physical domain. This is accomplished by using a single parameter to define the position of all nodes co-located in the reference domain, which is equivalent to using the nodal positions of the *continuous* high-order mesh corresponding to $\mathcal{E}_{h,p}$ as the parameter set. Let $\boldsymbol{x} \in \mathbb{R}^{N_x}$ denote these positions, then $\boldsymbol{x}_i^K$ is reconstructed from $\boldsymbol{x}$ by selecting the appropriate entry based on the position of the corresponding node in the reference domain. The corresponding nodal coordinates, denoted $\boldsymbol{X} \in \mathbb{R}^{N_x}$, in the reference domain are defined similarly. With this notation, the continuous domain mapping is parametrized by $\mu = \boldsymbol{x}$, i.e., $\Omega = \mathcal{G}(\Omega_0,\, \boldsymbol{x})$. The remainder of the document will use $\boldsymbol{x}$ to parametrize the domain deformation instead of μ. Finally, the integrals in (11) are evaluated using high-order Gaussian quadrature rules to yield the discrete form of the governing equations

$$\boldsymbol{r}(\boldsymbol{u},\, \boldsymbol{x}) = 0, \tag{16}$$

where $\boldsymbol{u} \in \mathbb{R}^{N_u}$ ($N_u = N_c N_p |\mathcal{E}_{h,p}|$) is the solution vector comprised of the coefficients $\boldsymbol{u}_i^K$ for all elements. The discrete form of the governing equations in (16) will be the point of departure for the remainder of

the document. The proposed shock tracking method presented in the next section will equally apply to an underlying discretization that takes the form of (16) and satisfies the requirements outlined in the beginning of this section, such as finite volume methods.

3. High-order discontinuity tracking via optimization-based r-adaptivity

In this section, we present a novel method designed to align discontinuous features in a finite-dimensional solution basis with features in the solution itself, with a driving application being transonic and supersonic compressible flow where shock waves are present. In the discretization setting outlined in Section 2, e.g., discontinuous Galerkin or finite volumes, this amounts to aligning element faces, where discontinuities are supported, with discontinuities in the solution. With the discontinuous features tracked with element faces, very coarse high-order discretizations effectively resolve the solution throughout the domain. However, in general, the location of the discontinuities is not known *a-priori* and therefore the face alignment cannot be done explicitly. To circumvent this difficulty that has plagued previous attempts at discontinuity tracking, the proposed method is built around three critical contributions:

1) a PDE-constrained optimization formulation with optimization variables taken as the nodal positions of the *high-order* mesh, i.e., an optimization-based r-adaptive framework, that allows elements to curve to accurately track the location of the discontinuity,

2) an objective function that attains a (local) minimum at a discontinuity-aligned mesh and monotonically decreases to this minimum in a neighborhood of radius approximately $h/2$,

3) a PDE-constrained optimization solver that does not require solution of the discrete PDE (16) on meshes not aligned with the discontinuity as this cannot be done reliably in a numerical scheme due to Gibbs' phenomena around the discontinuity (see Section 1 for the complete discussion).

The remainder of this section will detail these contributions and summarize the proposed high-order discontinuity-tracking algorithm in its entirety.

3.1. Optimization formulation for r-adaptivity

The proposed method for high-order resolution of discontinuous solutions of conservation laws reformulates the discrete nonlinear system in (16) that searches for a discrete solution $\boldsymbol{u}$ on a given mesh $\boldsymbol{x}$ to an optimization problem that searches for the discrete solution and mesh that minimize some objective function, $f : \mathbb{R}^{N_{\boldsymbol{u}}} \times \mathbb{R}^{N_{\boldsymbol{x}}} \to \mathbb{R}$, and satisfy the discrete PDE

$$
\begin{aligned}
\underset{\boldsymbol{u}\in\mathbb{R}^{N_{\boldsymbol{u}}},\, \boldsymbol{x}\in\mathbb{R}^{N_{\boldsymbol{x}}}}{\text{minimize}} \quad & f(\boldsymbol{u},\, \boldsymbol{x}) \\
\text{subject to} \quad & \boldsymbol{r}(\boldsymbol{u},\, \boldsymbol{x}) = \boldsymbol{0}.
\end{aligned}
\tag{17}
$$

The objective function, discussed in detail in Section 3.2, is constructed such that it is locally minimized at a discontinuity-aligned mesh. The proposed method falls into the category of r-adaptive methods given that it moves nodes to improve approximation quality; however, it has stronger requirements than existing r-adaptive methods as we aim to align mesh *faces* with the discontinuity rather than just nodes. See Figure 4 (top left) for a case where several mesh nodes are aligned with a discontinuity without any faces being aligned. Such situations are avoided through the construction of the objective function in Section 3.2 that is minimized when faces are aligned with discontinuities.

While the general formulation of the optimization problem in (17) is desirable in that it does not require any *a-priori* knowledge of the location of the discontinuity, it can be cumbersome due to the large number of optimization variables and difficult to avoid ill-conditioning and non-uniqueness in the domain mapping, i.e., severe mesh distortion or inversion. Ill-conditioning and non-uniqueness of the domain mapping will be addressed in the construction of the objective function in Section 3.2. To reduce the number of optimization variables, we re-parametrize the optimization problem in (17). Let $\boldsymbol{\phi} \in \mathbb{R}^{N_{\phi}}$ be a vector of parameters such

that $N_{\phi} \leq N_{\boldsymbol{x}}$ and $\mathcal{A} : \mathbb{R}^{N_{\phi}} \to \mathbb{R}^{N_{\boldsymbol{x}}}$ be an injective mapping that maps $\phi \mapsto \boldsymbol{x}$, i.e., $\boldsymbol{x} = \mathcal{A}(\phi)$. For clarity, we use $\boldsymbol{x}(\phi)$ in place of $\mathcal{A}(\phi)$. Then the optimization problem in (17) becomes

$$
\begin{aligned}
&\underset{\boldsymbol{u} \in \mathbb{R}^{N_{\boldsymbol{u}}}, \, \phi \in \mathbb{R}^{N_{\phi}}}{\text{minimize}} \quad f(\boldsymbol{u}, \, \boldsymbol{x}(\phi)) \\
&\text{subject to} \qquad \boldsymbol{r}(\boldsymbol{u}, \, \boldsymbol{x}(\phi)) = \boldsymbol{0}.
\end{aligned}
\tag{18}
$$

In addition to reducing the number of optimization parameters, the mapping $\mathcal{A}$ can incorporate mesh smoothing operations to improve mesh quality. However, despite the advantages of re-parametrization, the fact that $\mathcal{A}$ is not necessarily bijective limits the reachable mesh configurations and care must be taken to ensure the parametrization is capable of aligning with discontinuities. In this document, three different mesh parametrizations will be considered: the entire mesh, a single node parametrization with piecewise linear mesh smoothing, and a n-node parametrization with linear elastic smoothing in d dimensions. In the most general case where the entire mesh is used as optimization variables, the re-parametrization map is the identity map, $\mathcal{A}(\phi) = \phi$, or any other bijection, and (18) is equivalent to the original optimization problem. The other mesh parametrizations used in this work will be discussed in Sections 3.2 and 4.4.

3.2. Shock tracking objective function

A critical contribution of this work is the specific form of the objective function in the PDE-constrained optimization problem in (18). For the proposed framework to successfully align faces of the computational mesh with discontinuities in the solution the objective function must 1) attain a local minimum at some discontinuity-aligned mesh and 2) monotonically decrease to such a minima is a neighborhood of radius approximately $h/2$. The first requirement ensures a local minima of (18) leads to a desired mesh and the second requirement gives gradient-based numerical optimization methods hope to locate it even in the worst-case scenario where element faces are initially $h/2$ away from discontinuities. Thus, we seek an objective function that is an effective discontinuity indicator that monotonically decreases from large values at non-aligned meshes to smaller values at aligned meshes.

A host of discontinuity indicators exist in the literature [4, 41] as sensing discontinuous features is a backbone of many numerical methods for either capturing or tracking discontinuities. Two common classes of discontinuity indicators are physics-based and error-based. Physics-based indicators typically involve the residual associated with the integral form of the conservation law, i.e., the Rankine-Hugoniot conditions [4]. Error-based indicators use the magnitude of an approximate error measure such as adjoint-based error estimators [17] or the residual of the solution in a richer test space [1]. In this work, we propose a departure from these traditional discontinuity indicators and consider an indicator driven purely by the *robustness* of the numerics [41]. Recall from Section 1 that high-order discretizations of conservation laws fail when attempting to resolve discontinuous solutions due to severe spurious oscillations. Therefore we propose an objective function that penalizes these oscillations by considering the *element-wise* distance of the numerical solution from its mean value, i.e.,

$$
f_{shk}(\boldsymbol{u}, \, \boldsymbol{x}) = h_0^{-2} \sum_{K \in \mathcal{E}_{h,p}} \int_{\mathcal{G}(K, \, \boldsymbol{x})} \left\| u_{h,p} - \bar{u}_{h,p}^K \right\|_{\boldsymbol{W}}^2 dV
\tag{19}
$$

where the dependence of the finite dimensional solution $u_{h,p}$ on the discrete representation is implied, $\bar{u}_{h,p}^K$ is the mean value of $u_{h,p}$ over element K, $\boldsymbol{W} \in \mathbb{R}^{N_c \times N_c}$ is the symmetric positive semi-definite matrix that defines the local semi-norm, and h_0 is the length scale of the reference mesh $\mathcal{E}_{h,p}$

$$
\bar{u}_{h,p}^K = \frac{1}{|\mathcal{G}(K, \, \boldsymbol{x})|} \int_{\mathcal{G}(K, \, \boldsymbol{x})} u_{h,p} \, dV, \qquad |\mathcal{G}(K, \, \boldsymbol{x})| = \int_{\mathcal{G}(K, \, \boldsymbol{x})} dV, \qquad h_0 = \left(\frac{1}{|\mathcal{E}_{h,p}|} \int_{\Omega_0} dV \right)^{1/d}.
\tag{20}
$$

The h_0^{-2} factor ensures the indicator scales as $\mathcal{O}(1)$ instead of $\mathcal{O}(h^2)$. This objective function does not explicitly depend on the governing equation, only the approximation space $\mathcal{V}_{h,p}$, nor does it take into account the accuracy of the finite-dimensional solution, which marks a departure from physics-based and error-based discontinuity indicators.

8

Remark. *The discontinuity indicator in (19) will be identically zero, and therefore a useless objective function, if a piecewise constant polynomial approximation space is used ($p = 0$), which is equivalent to a finite volume scheme. In this case, the pointwise values of the finite dimensional solution $u_{h,p}$ in (19) can be replaced by its non-limited reconstruction using the solution in neighboring elements [34].*

The discontinuity indicator in (19) is not well-suited as the objective function in the discontinuity-tracking optimization setting (18) in its current form because it is agnostic to a poor quality or inverted mesh that may arise from certain choices of $\boldsymbol{x}$. Therefore we construct the objective function as a weighted combination of (19) and a function $f_{msh} : \mathbb{R}^{N_{\boldsymbol{x}}} \to \mathbb{R}$ that penalizes mesh distortion

$$f(\boldsymbol{u},\, \boldsymbol{x};\, \alpha) = f_{shk}(\boldsymbol{u},\, \boldsymbol{x}) + \alpha f_{msh}(\boldsymbol{x}). \tag{21}$$

In the remainder of the document the dependence on α will be dropped unless explicitly required. In this work, we use the following mesh distortion measures

$$f_{msh}(\boldsymbol{x}) = \begin{cases} h_0 \displaystyle\sum_{K \in \mathcal{E}_{h,p}} \left| \dfrac{h_0}{|\mathcal{G}(K,\, \boldsymbol{x})|} - 1 \right| & d = 1 \\[2em] h_0^d \displaystyle\sum_{K \in \mathcal{E}_{h,p}} \dfrac{1}{|\mathcal{G}(K,\, \boldsymbol{x})|} \displaystyle\int_{\mathcal{G}(K, \boldsymbol{x})} \left(\dfrac{\|G_{h,p}\|_F^2}{(\det G_{h,p})_+^{2/d}} \right)^r & \text{otherwise}, \end{cases} \tag{22}$$

where $r = 2$ is used in this work. For the general case of $d > 1$, the quantity in (22) is widely accepted in the high-order meshing community as reasonable metric to penalize mesh distortion and is well-suited for optimization [32, 18]. However, this mesh distortion measure is not useful for the case $d = 1$. Therefore we treat the case $d = 1$ separately and define mesh distortion as the deviation from a uniform mesh, modified to heavily penalize zero-volume elements. The h_0^d factors are again included to ensure the f_{msh} scales independently of h. Since both f_{shk} and f_{msh} are independent of h and p, the value of α that appropriately balances discontinuity tracking and mesh regularization should be mesh independent, i.e., independent of h and p.

3.2.1. Behavior of objective function: L^2 projection of discontinuous function

In this section, we use a simple one- and two-dimensional model problem to demonstrate the behavior of the proposed discontinuity indicator (19) near a discontinuity-aligned mesh and verify it possesses the desired properties advocated at the beginning of Section 3.2, namely attains a local minimum at a discontinuity-aligned mesh and monotonically approaches such a minima. Additionally, we show the impact of the mesh regularization term (22) on the combined objective function (21) as a function of the weighting parameter α. Unlike the physics- and error-based indicator, the proposed indicator is independent of the conservation law, which provides the opportunity to study the indicator *without* the complication of a partial differential equation. Instead, we consider the L^2 projection of a discontinuous function onto the piecewise polynomial space $\hat{\mathcal{V}}_{h,p}(\boldsymbol{x})$.

A complete description of the L^2 projection in the finite-dimensional setting of Section 2.2 is provided in Section 4.1. Here, we let $\boldsymbol{u}_H(\boldsymbol{x})$ denote the discrete representation of the L^2 projection of the discontinuous function $H : \Omega \to \mathbb{R}$ onto the space $\hat{\mathcal{V}}_{h,p}(\boldsymbol{x})$ and investigate the behavior of

$$j_\alpha(\boldsymbol{\phi}) = f(\boldsymbol{u}_H(\boldsymbol{x}(\boldsymbol{\phi})),\, \boldsymbol{x}(\boldsymbol{\phi});\, \alpha), \tag{23}$$

where $\boldsymbol{x}(\boldsymbol{\phi})$ denotes a parametrization of the computational mesh. In order to graphically study the objective function, it is important to use a single degree of freedom to parametrize the mesh. This is accomplished in one spatial dimension ($d = 1$) by using $\boldsymbol{\phi}$ to define the mapping of a single point $\bar{x}$ in the reference domain Ω_0 and mapping the remaining points piecewise linearly between the fixed boundaries and $\boldsymbol{\phi}$

$$x_{h,p}(X;\, \boldsymbol{\phi}) = \begin{cases} x_l + (X - x_l)\dfrac{\boldsymbol{\phi} - x_l}{\bar{x} - x_l} & X \leq \bar{x} \\[1.5em] x_r - (x_r - X)\dfrac{x_r - \boldsymbol{\phi}}{x_r - \bar{x}} & X > \bar{x} \end{cases} \tag{24}$$

9
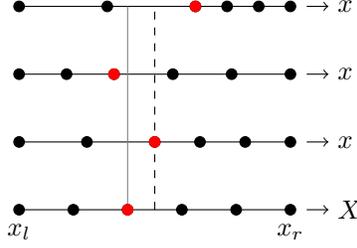
Figure 2: Single degree of freedom $1d$ mesh parametrization in (24) corresponding to a function with a discontinuity at (- - -). The choice of $\bar{x}$ (——) is chosen to align with a mesh node (●) near the position of the discontinuity and the parameter ϕ reduces to the position of this node. The remaining nodes (●) are determined by uniformly distributing them in the intervals $(x_l, \bar{x})$ and $(\bar{x}, x_r)$. The first configuration (*bottom*) shows the reference domain with a uniform nodal distribution and the remaining configurations corresponds to the deformation with $\phi = x^*$, $x^* - 0.75h$, $x^* + 0.75h$, respectively. This range of ϕ is the same range used to perform the sweep of the objective function (23) in Section 3.2.1 and Figure 3.

where $\Omega_0 = (x_l, x_r) \subset \mathbb{R}$. In the discrete setting, the mapping in (24) is applied to the nodal positions of the reference domain $\boldsymbol{X}$ to define the nodal positions of the physical domain $\boldsymbol{x}$ as

$$\boldsymbol{x}_i(\boldsymbol{\phi}) = \begin{cases} x_l + (\boldsymbol{X}_i - x_l)\dfrac{\phi - x_l}{\bar{x} - x_l} & \boldsymbol{X}_i \leq \bar{x} \\ x_r - (x_r - \boldsymbol{X}_i)\dfrac{x_r - \phi}{x_r - \bar{x}} & \boldsymbol{X}_i > \bar{x} \end{cases} \tag{25}$$

for $j \in \{1, \ldots, N_{\boldsymbol{x}}\}$. This single degree of freedom mesh parametrization possesses the obvious property $\boldsymbol{x}_i(\bar{x}) = \boldsymbol{X}_i$, which implies that the mesh quality will be perfect in terms of the distortion metric f_{msh} in (22) at $\phi = \bar{x}$ if the reference nodes are uniformly distributed. Additionally, if the reference nodes are uniformly distributed, the parametrization in (25) inherently incorporates ideal smoothing since the resulting nodes will be uniformly spaced in the intervals $[x_l, \bar{x}]$ and $(\bar{x}, x_r)$. A depiction of the single degree of freedom mesh parametrization is provided in Figure 2. This figure also shows the range that ϕ will be swept to study the behavior of $j_\alpha(\phi)$ next.

With this single degree of freedom parametrization, we consider the graph of $j_\alpha(\phi)$ for various values of α in Figure 3, where $\boldsymbol{u}_H$ corresponds to the L^2 projection of the piecewise sinusoidal function in (51) onto a piecewise polynomial basis with 17 elements and polynomial orders $p \in \{1, 2, 3, 4\}$. The reference nodes are uniformly distributed and therefore f_{msh} possesses a minimum at $\phi = \bar{x}$. For the case with no mesh regularization, the proposed objective function satisfies the two desired properties, i.e., possesses a local minima when the mesh is aligned with the discontinuity and monotonically approaches this minima in a neighborhood of radius $h/2$. As the penalty parameter is increased to $\alpha = 1$, the shape of the graph changes in that values near $\bar{x}$ decrease, but monotonicity is not broken. As α increases to 10, the monotonicity property breaks down and a local minima is introduced at $\bar{x}$, which becomes stronger as α is increased to 20. Once α is sufficiently large ($\alpha = 1000$), the mesh regularization term dominates causing the local minima at the shock location to vanish. Similar trends were observed when using coarser and finer discretization into piecewise polynomials. This simple test shows that the discontinuity indicator proposed in (21) is well-suited as the objective function in our discontinuity-tracking framework, provided the penalty parameter is chosen to *balance* the scale of the discontinuity indicator and mesh regularization terms.

Since the mesh distortion metric is defined separately for $d = 1$ and $d > 1$, it is important to repeat the investigation into $j_\alpha(\phi)$ in $d = 2$ dimensions. The same mesh parametrization is applied to the *first* coordinate in the two-dimensional setting with the second coordinate held fix

$$\boldsymbol{x}_i(\boldsymbol{\phi}) \cdot \boldsymbol{e}_1 = \begin{cases} x_l + (\boldsymbol{X}_i \cdot \boldsymbol{e}_1 - x_l)\dfrac{\phi - x_l}{\bar{x} - x_l} & \boldsymbol{X}_i \leq \bar{x} \\ x_r - (x_r - \boldsymbol{X}_i \cdot \boldsymbol{e}_1)\dfrac{x_r - \phi}{x_r - \bar{x}} & \boldsymbol{X}_i > \bar{x} \end{cases} \tag{26}$$

$$\boldsymbol{x}_i(\boldsymbol{\phi}) \cdot \boldsymbol{e}_2 = \boldsymbol{X}_i \cdot \boldsymbol{e}_2$$

10
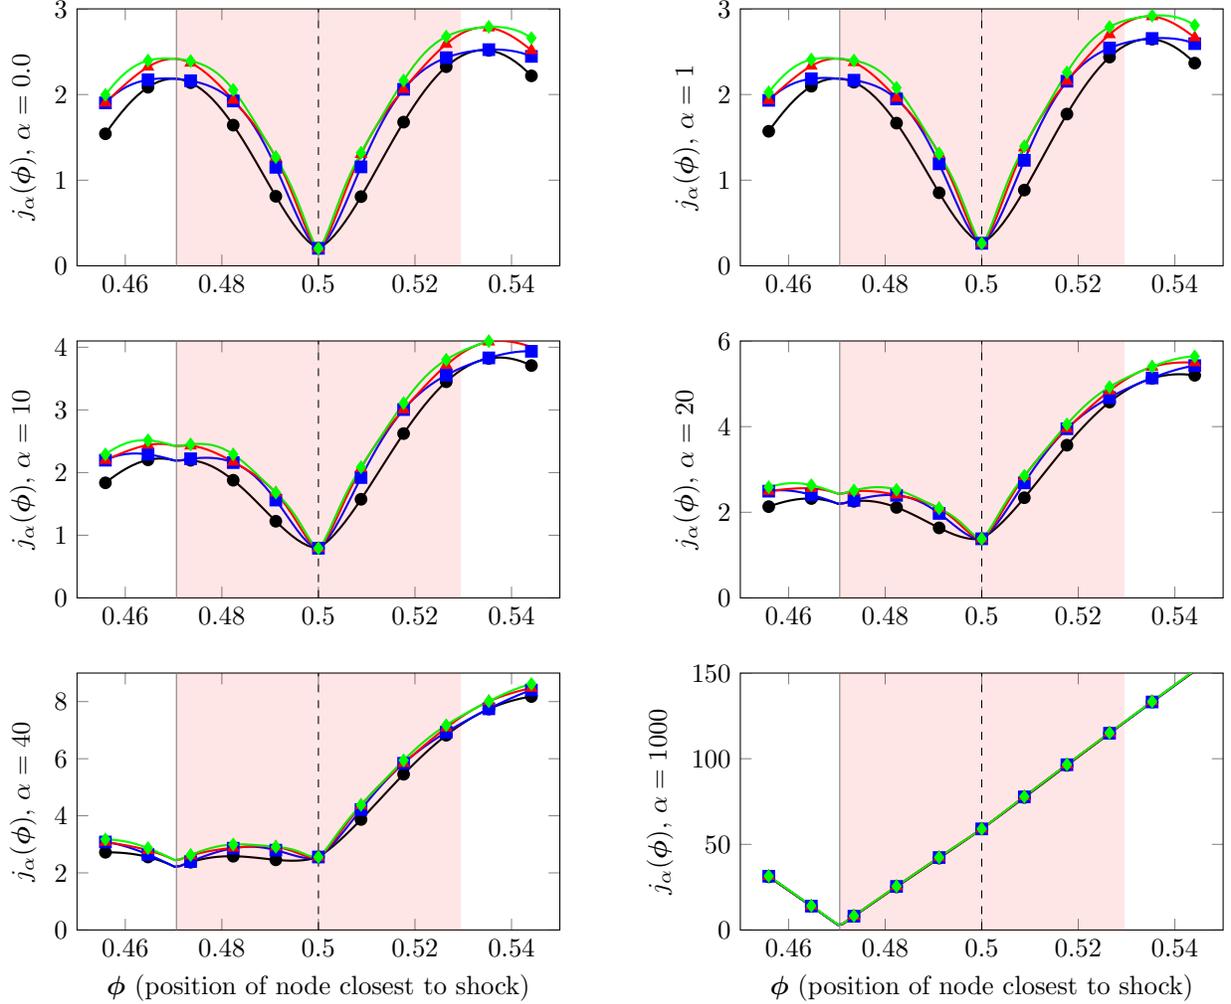
Figure 3: The objective function $j_\alpha(\phi)$ corresponding to the L^2 projection of the discontinuous function in (51) onto a piecewise polynomial basis of degree $p = 1$ (—●—), $p = 2$ (—■—), $p = 3$ (—▲—), $p = 4$ (—◆—) with $|\mathcal{E}_{h,p}| = 17$ elements for various values of the mesh regularization parameter α under the mesh parameterization in (24) and Figure 2. The parameter ϕ is swept over the range $[x^* - 0.75h, \, x^* + 0.75h]$, where x^* is the location of the discontinuity (- - -). The position of $\bar{x}$ in (24) is indicated with (——); when $\phi = \bar{x}$ nodes are uniformly distributed and f_{msh} takes its minimum values. The red shaded region identifies the neighborhood of radius $h/2$ about x^*. The objective function with no mesh regularization ($\alpha = 0$) in the *top left* plot shows the two desired properties discussed in Section 3.2, namely, a minimum at x^* and monotonically decreases to this minimum in a neighborhood centered at x^* with radius $h/2$. The remaining plots show increasing values of the regularization parameter α. The minimum does not noticeably move until α is sufficiently large ($\alpha = 40$) at which point is moves from x^* to $\bar{x}$, i.e., the objective function transitions from prioritizing discontinuity tracking to mesh regularization.

11
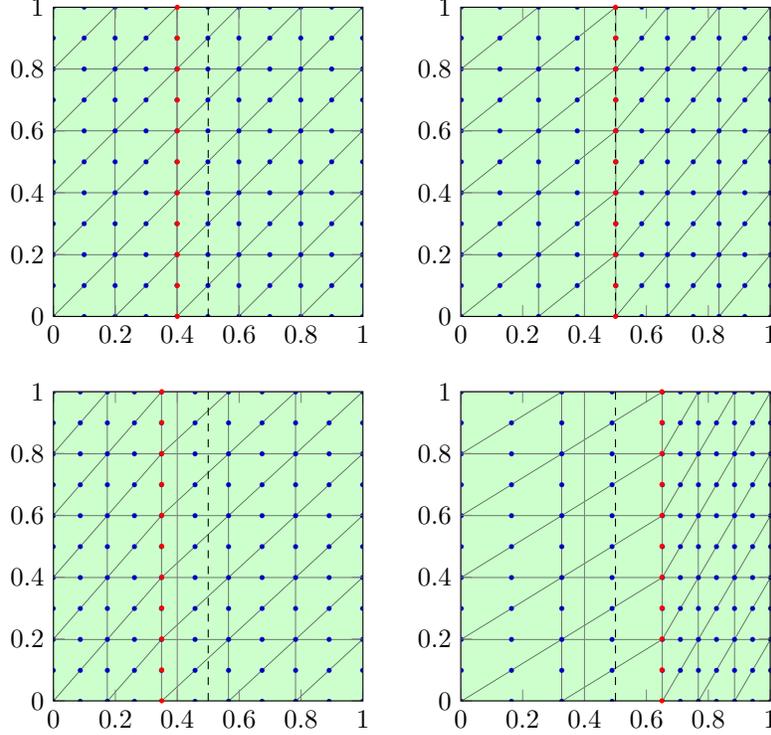
Figure 4: Single degree of freedom $2d$ mesh parametrization in (24) and problem setup for investigation into the proposed objective function in (23) as applied to the L^2 projection of the discontinuous function in (27). The choice of $\bar{x}$ (——) is chosen to align with a line of mesh nodes (•) near the position of the discontinuity (- - -) and the parameter ϕ reduces to the position of this node. The remaining nodes (x-coordinates only; y-coordinates are frozen) are determined by uniformly distributing them in the intervals $(x_l, \bar{x})$ and $(\bar{x}, x_r)$. The first configuration (*top left*) shows the reference domain with a uniform nodal distribution and the remaining configurations corresponds to the deformation with $\phi = x^*$, $x^* - 0.75h$, $x^* + 0.75h$. This range of ϕ is the same range used to perform the sweep of the objective function (23) in Section 3.2.1 and Figure 5.

A depiction of the single degree of freedom mesh parametrization in $d = 2$ dimensions is provided in Figure 4. This figure also shows the range that ϕ will be swept to study the behavior of $j_\alpha(\phi)$ next.

Figure 5 shows the graph of $j_\alpha(\phi)$ for various values of α corresponding to this single degree of freedom parametrization, where $\boldsymbol{u}_H$ corresponds to the L^2 projection of the piecewise sinusoidal function

$$\eta(x,\,y) = \begin{cases} 0.25\sin(4\pi x) - 0.5, & x < 0 \\ 0.25\sin(4\pi x) + 0.5, & x > 0, \end{cases} \tag{27}$$

$(x,\,y) \in \Omega = (-1,\,1) \times (-1,\,1)$, onto a piecewise polynomial basis with 50 elements and polynomial orders $p \in \{1,\,2,\,3\}$. Similar to the one-dimensional case, the proposed objective function satisfies the two desired properties i.e., possesses a local minima when the mesh is aligned with the discontinuity and monotonically approaches this minima in a neighborhood of radius about $h/2$. Also, as the penalty parameter is increased the mesh regularization is preferred and the local minima at the discontinuity disappears. Similar trends were observed when using coarser and finer discretization into piecewise polynomials.

3.2.2. Comparison with physics-based and error-based indicators: inviscid Burgers' equation

With the merit of the proposed discontinuity indicator verified in the previous section, we compare it with other metrics that have been used as discontinuity or refinement indicators. The first is an *error-based* indicator that measures the extent that a Galerkin solution corresponding to a test and trial space $\mathcal{V}_{h,p}$
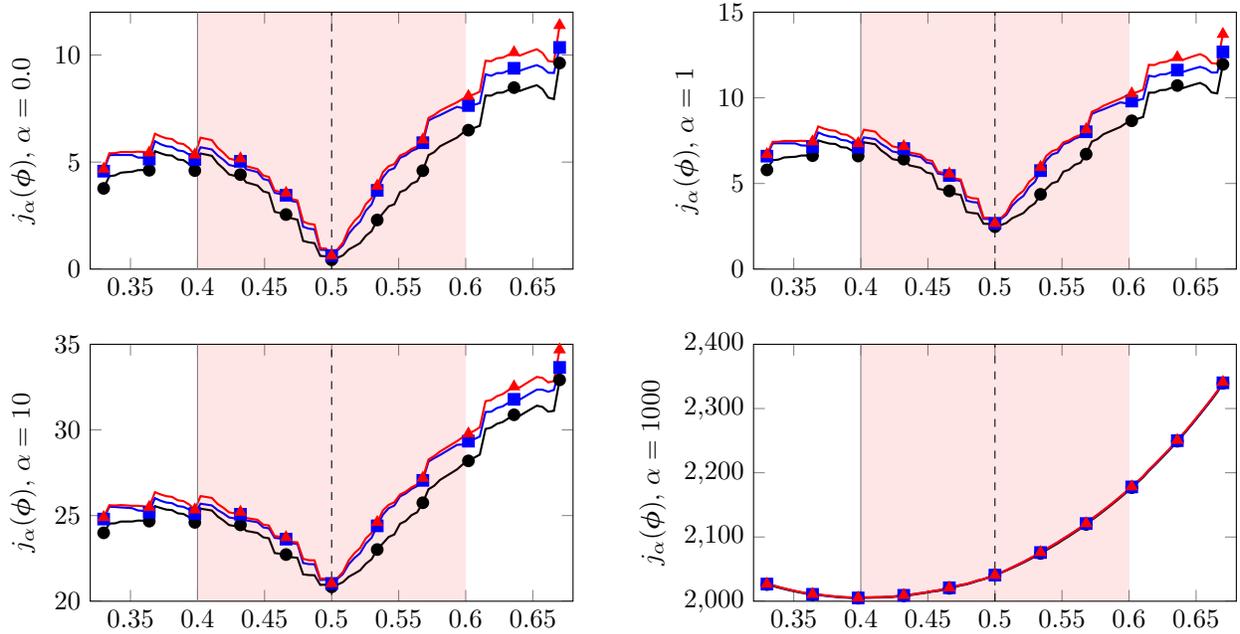
Figure 5: The objective function $j_\alpha(\phi)$ corresponding to the L^2 projection of the discontinuous function in (27) onto a piecewise polynomial basis of degree $p = 1$ (—●—), $p = 2$ (—■—), $p = 3$ (—▲—) with $|\mathcal{E}_{h,p}| = 50$ elements for various values of the mesh regularization parameter α under the mesh parameterization in (24) and Figure 4. The parameter ϕ is swept over the range $[x^* - 0.75h, x^* + 0.75h]$, where x^* is the location of the discontinuity (- - -). The position of $\bar{x}$ in (24) is indicated with (——); when $\phi = \bar{x}$ nodes are uniformly distributed and f_{msh} takes its minimum values. The red shaded region identifies the neighborhood of radius $h/2$ about x^*. The objective function with no mesh regularization ($\alpha = 0$) in the *top left* plot shows the two desired properties discussed in Section 3.2, namely, a minimum at x^* and monotonically decreases to this minimum in a neighborhood centered at x^* with radius $h/2$. The remaining plots show increasing values of the regularization parameter α. The minimum does not noticeably move until α is sufficiently large at which point is moves from x^* to $\bar{x}$, i.e., the objective function transitions from prioritizing discontinuity tracking to mesh regularization.

13

satisfies the residual in an enriched test space $\mathcal{V}_{h'}^{p'}$

$$f_{h,p}^{h',p'}(\boldsymbol{u}, \boldsymbol{x}) = r_{h,p}^{h',p'}(u_{h,p}, \boldsymbol{x}) \tag{28}$$

where $r_{h,p}^{h',p'}$ is the finite-dimensional Petrov-Galerkin residual corresponding to the trial space $\mathcal{V}_{h,p}$ and test space $\mathcal{V}_{h'}^{p'}$ defined in (10). The rationale behind this type of indicator is there will be no contribution from regions of the domain away from the discontinuity that are well-resolved and near the discontinuity the oscillations will be penalized by the additional test functions. We will consider two instances of this class of indicators: h-refinement where $h' = h/2$, $p' = p$ and p-refinement where $h' = h$, $p' = p + 1$. We also consider a *physics-based* indicator based on the Rankine-Hugoniot conditions [4]. Recall the Rankine-Hugoniot conditions for a steady state conservation law reduce to

$$F(u^+) = F(u^-) \quad \text{on } \Gamma_s, \tag{29}$$

where Γ_s is any surface in the domain and u^+, u^- are the values of the solution u on either side of Γ_s. This will hold in smooth regions of the domain or in the case Γ_s is the shock surface; however, it will be violated in regions in the domain with non-physical jumps that arise from non-aligned meshes.

The condition in (29) is converted to a discontinuity indicator by integrating the residual over all inter-element faces

$$f_{rh}(\boldsymbol{u}, \boldsymbol{x}) = \sum_{\partial K \in \partial \mathcal{E}_{h,p}} \int_{\partial K} \left\| F(u_{h,p}^-) - F(u_{h,p}^+) \right\|_{\boldsymbol{W}}^2 \, ds, \tag{30}$$

where $\boldsymbol{W} \in \mathbb{R}^{N_c \times N_c}$ is a symmetric positive semi-definite matrix that defines the local semi-norm. In this work, we take $\boldsymbol{W} = \boldsymbol{e}_1 \boldsymbol{e}_1^T$, where $\boldsymbol{e}_1 \in \mathbb{R}^d$ is the first canonical unit vector.

To study the indicators in (28)-(30) we consider the modified inviscid Burgers' equation with a discontinuous source term (55) whose solution is given by (57). A complete description of this problem is deferred to Section 4.2. Here, we let $\boldsymbol{u}$ denote the discrete representation of the numerical solution of the inviscid Burgers' equation in (55) using the test and trial space $\mathcal{V}_{h,p}$ and investigate the behavior of various indicators: 1) the proposed discontinuity tracking indicator in (19), 2) a residual-based indicator under p-refinement ($p' = p + 1$), 3) a residual-based indicator under h-refinement ($h = h/2$), 4) and the Rankine-Hugoniot indicator

$$j_{shk}(\boldsymbol{\phi}) = f_{shk}(\boldsymbol{u}(\boldsymbol{x}(\boldsymbol{\phi})), \boldsymbol{x}(\boldsymbol{\phi})) \tag{31}$$

$$j_{h,p}^{h,p+1}(\boldsymbol{\phi}) = f_{h,p}^{h,p+1}(\boldsymbol{u}(\boldsymbol{x}(\boldsymbol{\phi})), \boldsymbol{x}(\boldsymbol{\phi})) \tag{32}$$

$$j_{h,p}^{h/2,p}(\boldsymbol{\phi}) = f_{h,p}^{h/2,p}(\boldsymbol{u}(\boldsymbol{x}(\boldsymbol{\phi})), \boldsymbol{x}(\boldsymbol{\phi})) \tag{33}$$

$$j_{rh}(\boldsymbol{\phi}) = f_{rh}(\boldsymbol{u}(\boldsymbol{x}(\boldsymbol{\phi})), \boldsymbol{x}(\boldsymbol{\phi})). \tag{34}$$

To study these indicators in isolation, the mesh regularization term is excluded. We consider the same single degree of freedom parametrization from the previous section, i.e., $\boldsymbol{x}(\boldsymbol{\phi})$ is given by (25). The graphs of these functions over the range $\boldsymbol{\phi} \in [x^* - 0.75h, \, x^* + 0.75h]$ where $x^* = 0.0$ is the discontinuity location, is shown in Figure 6 for a discretization with 40 elements of orders $p = 1, 2, 3$. It is clear from this figure that *all* of the functions considered possess a local minimum when the mesh is perfectly aligned with the discontinuity ($\boldsymbol{\phi} = x^*$). However, the proposed objective function is the only one that *monotonically* approaches this minimum in a neighborhood of radius $h/2$. The other indicators are extremely oscillatory and in many cases are only monotonic in a neighborhood of radius $h/10$ about the minimum. It is interesting to note that as the polynomial order of the discretization increases, the physics- and error-based indicators becomes more oscillatory, while the proposed indicator appears to be agnostic to the polynomial order. This implies the physics- and error-based indicators are not suitable candidates for the objective function in the discontinuity-tracking framework since the mesh faces would have to be initialized very close to the unknown discontinuity for any hope of convergence. On the other hand, the monotonicity of the proposed indicator around the discontinuity suggests that a gradient-based optimization solver will be capable of aligning the mesh with

14
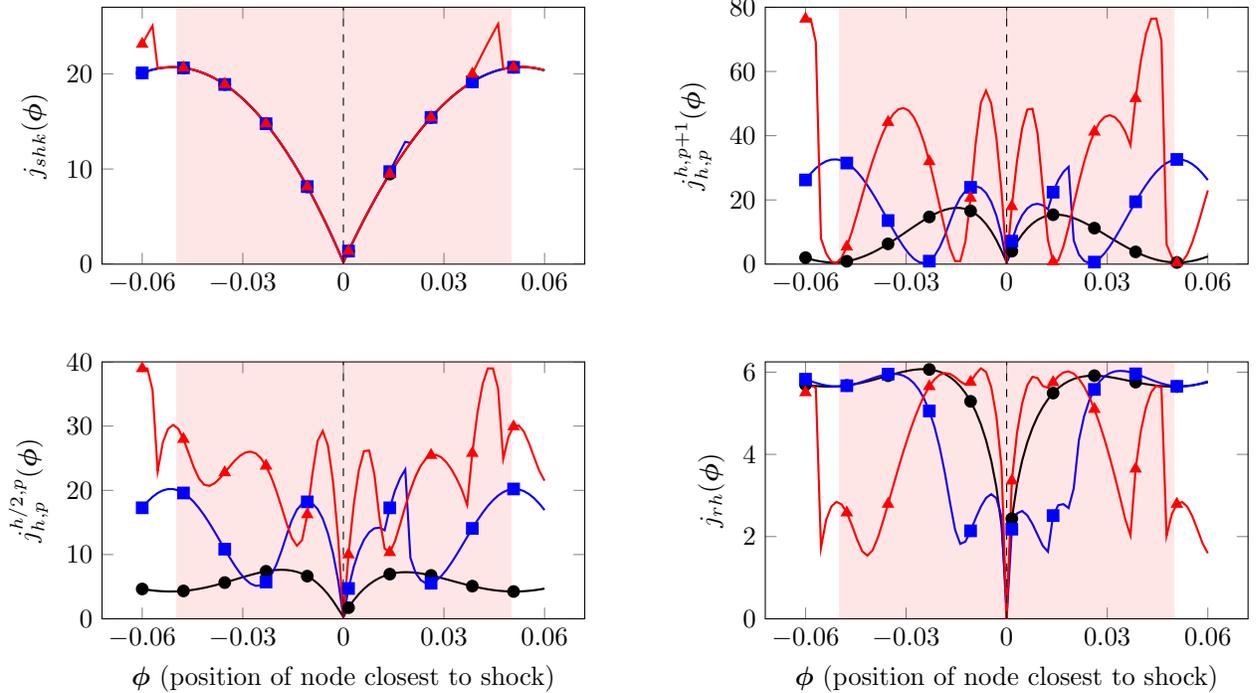
Figure 6: A side-by-side comparison of the behavior of the proposed objective function with standard physics- and error-based objective functions (*top left*: deviation from mean (31), *top right*: residual-based p-refinement (32), *bottom left*: residual-based h-refinement (33), *bottom right*: Rankine-Hugoniot conditions (34)) corresponding to the solution of the inviscid Burgers' equation with a discontinuous source term in (55) using a piecewise polynomial basis of degree $p = 1$ (—●—), $p = 2$ (—■—), $p = 3$ (—▲—) with $|\mathcal{E}_{h,p}| = 40$ elements without mesh regularization $\alpha = 0$. The single degree mesh parameterization from (24) and Figure 2 is used and the parameter ϕ is swept over the range $[x^* - 0.75h, x^* + 0.75h]$, where x^* is the location of the discontinuity (- - -). The red shaded region identifies the neighborhood of radius $h/2$ about x^*. Even with the added complexity of the conservation law, the proposed objective function attains a minimum when the mesh is aligned with the discontinuity and approaches this monotonically in a neighborhood of radius $h/2$. The other discontinuity indicators possess the correct minimum, but clearly do not satisfy monotonicity.

discontinuities even in the worst case where a face is $\mathcal{O}(h/2)$ away from the discontinuity. In the remainder of this document, we will solely consider the proposed discontinuity indicator with mesh regularization term (21).

3.3. Full space optimization solver

The final ingredient to complete the description of the proposed discontinuity tracking method is the numerical solver for the PDE-constrained optimization problem in (18) with the objective function in (21) and constraints in (16). The first, and most common, PDE-constrained optimization solver is known as the reduced space approach [27] where the PDE constraint is explicitly enforced by defining the implicit function $\boldsymbol{u}(\boldsymbol{x})$ as the solution of $\boldsymbol{r}(\boldsymbol{u}, \boldsymbol{x}) = \boldsymbol{0}$. The PDE-constrained optimization problem reduces to the unconstrained optimization problem

$$\underset{\boldsymbol{\phi} \in \mathbb{R}^{N_\phi}}{\text{minimize}} \quad f(\boldsymbol{u}(\boldsymbol{\phi}), \boldsymbol{x}(\boldsymbol{\phi})) \tag{35}$$

where the state variable $\boldsymbol{u}$ is no longer an optimization variable as it can be determined given $\boldsymbol{\phi}$. Standard unconstrained optimization solvers such as quasi-Newton methods [37] can be applied to solve (35). The reduced space approach is popular because the number of parameters is greatly reduced compared to the full space formulation in (18) due to the elimination of $\boldsymbol{u}$ and sophisticated PDE solvers can be employed as a black box to query the implicit function $\boldsymbol{u}(\boldsymbol{x})$ and its gradient using the sensitivity or adjoint method [27].

15

Despite these advantages, the reduced space approach cannot, in general, be used to solve the discontinuity-tracking PDE-constrained optimization problem in (18). The reduced space formulation implicitly assumes the PDE constraint (16) can be solved given $\boldsymbol{x}(\boldsymbol{\phi})$ *for any $\boldsymbol{\phi}$ in the feasible set* (taken as $\mathbb{R}^{N_\phi}$). Recall from the discussion in Section 1 that the PDE constraint cannot be solved unless $\boldsymbol{x} = \boldsymbol{x}^*$, where $\boldsymbol{x}^*$ is the discrete representation of any discontinuity-aligned mesh, or enough artificial viscosity is added to stabilize the solution procedure. Since the mesh will not align with discontinuities until convergence, the reduced space approach requires repeated evaluation of the PDE solution at meshes not aligned with the discontinuity and can only be used if sufficient viscosity, or Laplacian-based diffusion, is added during intermediate iterations. Such an approach would lead to a heuristic homotopy method and certainly impact, if not prevent, convergence of the optimization procedure.

To overcome the numerous issues associated with the reduced space approach to PDE-constrained optimization for the discontinuity-tracking problem in (18), we propose using the *full space approach* [21, 10]. In this setting, the PDE state vector $\boldsymbol{u}$ and mesh parameters $\boldsymbol{\phi}$ are taken as *independent* optimization variables, as suggested in the formulation in (18), and converged to their optimal values simultaneously. This implies the solution of the PDE is *never required away from a discontinuity-aligned mesh $\boldsymbol{x}^*$* and overcomes the fundamental difficulty with the reduced space formulation. Since $\boldsymbol{u}$ and $\boldsymbol{\phi}$ are varying independently, the PDE constraint cannot be eliminated and a nonlinearly constrained optimization solver is required. In this work, we use the Sequential Quadratic Programming (SQP) method [11] implemented in the well-known optimization software SNOPT [19] as our constrained optimization solver. Any SQP method will reduce the nonlinear optimization problem in (18) to a sequence of quadratic programs of the form

$$\begin{aligned}
&\underset{\Delta\boldsymbol{u}\in\mathbb{R}^{N_{\boldsymbol{u}}},\,\Delta\boldsymbol{\phi}\in\mathbb{R}^{N_\phi}}{\text{minimize}} \quad q(\Delta\boldsymbol{u},\,\Delta\boldsymbol{\phi}) \\
&\text{subject to} \quad \boldsymbol{r}(\boldsymbol{u}_k,\,\boldsymbol{x}(\boldsymbol{\phi}_k)) + \frac{\partial\boldsymbol{r}}{\partial\boldsymbol{u}}(\boldsymbol{u}_k,\,\boldsymbol{x}(\boldsymbol{\phi}_k))\Delta\boldsymbol{u} + \frac{\partial\boldsymbol{r}}{\partial\boldsymbol{\phi}}(\boldsymbol{u}_k,\,\boldsymbol{x}(\boldsymbol{\phi}_k))\Delta\boldsymbol{\phi} = \boldsymbol{0}
\end{aligned} \tag{36}$$

where $\boldsymbol{u}_k$ and $\boldsymbol{\phi}_k$ correspond to the kth iterate of the optimization variables and $q : \mathbb{R}^{N_{\boldsymbol{u}}} \times \mathbb{R}^{N_\phi} \to \mathbb{R}$ is the quadratic Taylor approximation of the objective function where the Hessian is replaced by the Hessian of the Lagrangian (or a symmetric positive definite approximation)

$$\mathcal{L}(\boldsymbol{u},\,\boldsymbol{\phi},\,\boldsymbol{\lambda}) = f(\boldsymbol{u},\,\boldsymbol{\phi}) - \boldsymbol{\lambda}^T \boldsymbol{r}(\boldsymbol{u},\,\boldsymbol{x}(\boldsymbol{\phi}))$$

with Lagrange multiplier $\boldsymbol{\lambda} \in \mathbb{R}^{N_{\boldsymbol{u}}}$. The solution of this quadratic program provides search directions that are used to update the optimization variables as

$$\begin{aligned}
\boldsymbol{u}_{k+1} &= \boldsymbol{u}_k + \tau\Delta\boldsymbol{u} \\
\boldsymbol{\phi}_{k+1} &= \boldsymbol{\phi}_k + \tau\Delta\boldsymbol{\phi},
\end{aligned} \tag{37}$$

where τ is determined via a linesearch[3] on a merit function [37]. It can be shown that the KKT conditions of (36) are equivalent to the linearized KKT conditions of (18) [11, 37] and therefore SQP methods are equivalent to applying damped quasi-Newton methods to the KKT system of the nonlinear optimization problem (18).

An optimization setting such as the linesearch-based SQP method in (36)-(37) demonstrates the superiority of the full space approach compared to the reduced space approach *for the discontinuity-tracking problem*. In the reduced space formulation, the complete nonlinear solution of $\boldsymbol{r}(\boldsymbol{u},\,\boldsymbol{x}(\boldsymbol{\phi})) = \boldsymbol{0}$ for a given $\boldsymbol{\phi}$ is required, which causes severe robustness issues if $\boldsymbol{x}(\boldsymbol{\phi})$ is not aligned with discontinuities and will likely not be possible since oscillations will cause the solver to terminate at non-physical values. However, the full space approach only requires the solution of the quadratic program in (36) that only involves the *linearized governing equation*

$$\boldsymbol{r}(\boldsymbol{u}_k,\,\boldsymbol{x}(\boldsymbol{\phi}_k)) + \frac{\partial\boldsymbol{r}}{\partial\boldsymbol{u}}(\boldsymbol{u}_k,\,\boldsymbol{x}(\boldsymbol{\phi}_k))\Delta\boldsymbol{u} + \frac{\partial\boldsymbol{r}}{\partial\boldsymbol{\phi}}(\boldsymbol{u}_k,\,\boldsymbol{x}(\boldsymbol{\phi}_k))\Delta\boldsymbol{\phi} = \boldsymbol{0}.$$

---

[3]Trust-region SQP methods could be used in place of a linesearch-based SQP method. In this case, the trust-region radius is determined at the beginning of the iteration and the quadratic program is restricted to the trust region.

The solution $\boldsymbol{u}_k$ and the mesh $\boldsymbol{x}(\boldsymbol{\phi}_k)$ are updated simultaneously based on the result of the linesearch. Therefore, at every optimization iteration the mesh is updated alongside the solution and the linesearch (or trust-region) prevents steps that cause non-physical values as they would adversely affect the merit function.

While we use the general nonlinear optimization software SNOPT in this work to demonstrate the merit of the proposed discontinuity-tracking framework, there is a considerable opportunity for efficiency to develop a specialized constrained optimization solver that takes advantage of the considerable structure present in (18), namely, nonlinear equality constraints with an invertible partial Jacobian $\dfrac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}$. Furthermore, a specialized optimizer will be able to re-use data structures and operations from the PDE implementation and incorporate homotopy measures (viscosity or temporal relaxation) into the discrete PDE for particularly difficult problems. Such a solver is the subject of ongoing work.

3.4. Practical considerations: implementation, initialization, robustness, and efficiency

In this section we discuss a number of practical considerations important in the application of the proposed method to interesting problems. These include relevant implementation details, initialization of the optimization problem in (18), and the expected efficiency compared to standard steady state solvers or a reduced space optimization formulation.

Implementation

The implementation of the proposed method consists of three main components: the residual $\boldsymbol{r}(\boldsymbol{u},\, \boldsymbol{x})$ and objective $f(\boldsymbol{u},\, \boldsymbol{x})$ functions, the definition of the mapping from optimization parameters to nodal coordinates of the high-order mesh $\boldsymbol{x} = \mathcal{A}(\boldsymbol{\phi})$, and the nonlinear optimization solver for (18),

The high-dimensional nature of the optimization problem in (18) with $N_{\boldsymbol{u}} + N_{\boldsymbol{\phi}}$ optimization variables, requires a gradient-based optimization solver as derivative-free methods are impractical in many dimensions. Any such solver will require objective function and gradient evaluations as well as the constraint function and gradient evaluations, i.e.,

$$
\begin{aligned}
&f(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})), && \frac{\partial f}{\partial \boldsymbol{u}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})), && \frac{\partial f}{\partial \boldsymbol{\phi}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})), \\
&\boldsymbol{r}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})), && \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})), && \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{\phi}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})).
\end{aligned}
\tag{38}
$$

These are the same terms that would be required if a reduced space optimization approach was used; however, in that setting, the derivative terms would have to be combined using the sensitivity or adjoint method to compute the *total derivative* of the objective function with respect to $\boldsymbol{\phi}$. The partial derivatives with respect to $\boldsymbol{\phi}$ are computed via a simple application of the chain rule

$$
\begin{aligned}
\frac{\partial f}{\partial \boldsymbol{\phi}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})) &= \frac{\partial f}{\partial \boldsymbol{x}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi}))\frac{\partial \boldsymbol{x}}{\partial \boldsymbol{\phi}}(\boldsymbol{\phi}) \\
\frac{\partial \boldsymbol{r}}{\partial \boldsymbol{\phi}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi})) &= \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{x}}(\boldsymbol{u},\, \boldsymbol{x}(\boldsymbol{\phi}))\frac{\partial \boldsymbol{x}}{\partial \boldsymbol{\phi}}(\boldsymbol{\phi})
\end{aligned}
\tag{39}
$$

where $\dfrac{\partial \boldsymbol{x}}{\partial \boldsymbol{\phi}} = \dfrac{\partial \mathcal{A}}{\partial \boldsymbol{\phi}}$.

The terms involving the objective and residual function, i.e.,

$$
\begin{aligned}
&f(\boldsymbol{u},\, \boldsymbol{x}), && \frac{\partial f}{\partial \boldsymbol{u}}(\boldsymbol{u},\, \boldsymbol{x}), && \frac{\partial f}{\partial \boldsymbol{x}}(\boldsymbol{u},\, \boldsymbol{x}), \\
&\boldsymbol{r}(\boldsymbol{u},\, \boldsymbol{x}), && \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}(\boldsymbol{u},\, \boldsymbol{x}), && \frac{\partial \boldsymbol{r}}{\partial \boldsymbol{x}}(\boldsymbol{u},\, \boldsymbol{x})
\end{aligned}
\tag{40}
$$

are solely dependent on the discretization of the conservation law and can be completely separated from the mesh parametrization and optimization solver. The implementation of the objective function and its

derivatives are straightforward due to the element-wise integral form of (19) that can be expressed in terms of element mass matrices. The residual function and its derivative with respect to the state variable $\boldsymbol{u}$ are standard terms required by an implicit steady state solver. The derivative of the residual function with respect to the nodal coordinates $\boldsymbol{x}$ is not standard and usually only required for the purposes of gradient-based shape optimization [54]. However, the formulation of the conservation law on a *reference domain* in (9) implies the elements over which the integrals are performed $K \in \mathcal{E}_{h,p}$ are independent of $\boldsymbol{x}$ and the dependence on $\boldsymbol{x}$ is restricted to the *pointwise* definition of the transformed fluxes. This implies the implementation and data structures used to compute $\dfrac{\partial \boldsymbol{r}}{\partial \boldsymbol{x}}$ will mimic those for $\dfrac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}$.

The mesh parametrization can be treated independently of the PDE discretization and must implement the mapping from $\boldsymbol{\phi}$ to the mesh coordinates $\boldsymbol{x} = \mathcal{A}(\boldsymbol{\phi})$ as well as the Jacobian of this mapping, i.e.,

$$
\begin{aligned}
\boldsymbol{x}(\boldsymbol{\phi}) &= \mathcal{A}(\boldsymbol{\phi}) \\
\frac{\partial \boldsymbol{x}}{\partial \boldsymbol{\phi}}(\boldsymbol{\phi}) &= \frac{\partial \mathcal{A}}{\partial \boldsymbol{\phi}}(\boldsymbol{\phi}).
\end{aligned}
\tag{41}
$$

The most convenient and general mesh parametrization uses the entire mesh as the optimization variables, i.e.,

$$
\begin{aligned}
\boldsymbol{\phi} &= \boldsymbol{x} \\
\mathcal{A}(\boldsymbol{\phi}) &= \boldsymbol{\phi} \\
\frac{\partial \mathcal{A}}{\partial \boldsymbol{\phi}}(\boldsymbol{\phi}) &= \boldsymbol{I}.
\end{aligned}
\tag{42}
$$

However, this burdens the optimization solver due to the large number of optimization variables and likely introduces local minima that do not correspond to discontinuity-aligned meshes. A general single-node mesh parametrization for one-dimensional problems is presented in Section 3.2.1 and used in Sections 3.2.1-3.2.2, 4.1-4.3. A general approach in arbitrary dimensions that effectively reduces the number of optimization variables selects $\boldsymbol{\phi}$ to represent coordinates of nodes in the vicinity of discontinuities and smoothing based on linear elasticity is used to determine the positions of the remaining nodes. This assumes *a-priori* knowledge of the shock region, which can easily be found by introducing sufficient viscosity and finding the corresponding steady state using traditional methods. Regardless of the choice of mesh parametrization, the terms in (40) and (41) can easily be combined to yield the required derivative information for gradient-based optimization solvers in (38).

The full space optimization solver is the final component required to realize the proposed discontinuity-tracking method in practice. A number of extremely robust optimization packages exist that could be used for this purpose such as SNOPT [19], IPOPT [50], and the Trilinos package ROL [25] with ROL being one of the only parallel solvers available. However, the problem in (18) possesses considerable structure that most black box optimizers cannot leverage such as an invertible partial constraint Jacobian $\dfrac{\partial \boldsymbol{r}}{\partial \boldsymbol{u}}$, only equality constraints, and specialized, parallel, sparse data structures for the terms in (40) from the PDE discretization. Therefore, it is desirable to use solvers that can leverage such information or develop a specialized solver as discussed in Section 3.3.

Initialization

The initial guess for the optimization variables $\boldsymbol{u}$, $\boldsymbol{\phi}$ is a particularly important consideration in gradient-based optimization with nonlinear constraints due to the presence of local minima at meshes not aligned with discontinuities. The objective function in (21) was requirement to mostly eliminate local minima due to the *monotonicity* constraint, but cannot be guaranteed. In this work we use an initial guess for the mesh parameters $\boldsymbol{\phi}_0$ such that the resulting mesh is identical to the reference mesh $\mathcal{E}_{h,p}$. This simple mesh initialization strategy will be shown to be quite robust in the numerical examples in Section 4. Unlike the reduced space approach to optimization that only requires an initial guess for $\boldsymbol{\phi}$, the full space approach requires initialization of $\boldsymbol{u}$. Naturally, the solution of $\boldsymbol{r}(\boldsymbol{u}, \boldsymbol{x}(\boldsymbol{\phi}_0)) = \boldsymbol{0}$ is a desirable initial guess since $\boldsymbol{u}_0$,

ϕ_0 would satisfy the PDE constraint. However, this is not possible since $\boldsymbol{x}(\phi_0)$ corresponds to a mesh that is not aligned with discontinuities. Instead, define

$$\boldsymbol{r}_\nu(\boldsymbol{u},\,\boldsymbol{x}) = \boldsymbol{0} \tag{43}$$

as a discontinuous Galerkin discretization of a *viscous* conservation law with viscosity parameter ν

$$\nabla \cdot \mathcal{F}(U) + \nabla \cdot \mathcal{F}_\nu(U,\,\nabla U) = 0 \quad \text{in } \mathcal{G}(\Omega_0,\,\mu), \tag{44}$$

where $\mathcal{F}_\nu$ becomes the zero map as $\nu \to 0$. The reference [3] provides an overview of various methods available for the treatment of viscous terms in a DG setting. Since we are only looking for an initial guess for the optimizer, the viscous terms in (44) can either correspond to physical viscosity, e.g., the viscous terms in the Navier-Stokes equations, or non-physical Laplacian smoothing. Let $\bar{\nu}$ be a value of the viscosity parameter that is as close to zero as possible but sufficiently large that the solution of

$$\boldsymbol{r}_{\bar{\nu}}(\boldsymbol{u},\,\boldsymbol{x}(\phi_0)) = \boldsymbol{0} \tag{45}$$

can be found reliably. Then the solution of (45) is taken as the initial guess for the state variable.

These heuristics provide high quality initial guesses, particularly at low polynomial orders such as $p = 1$, but issues arise for high-order elements due to the introduction of local minima not necessarily sensed by the proposed indicator (19). Therefore, we propose a homotopy method on the polynomial order since even the $p = 1$ solution finds a quality approximation of the shock location and solution. The transfer of the PDE solution $\boldsymbol{u}$ and mesh $\boldsymbol{x}$ to a higher order mesh is a trivial injection operation and the corresponding optimization parameters ϕ are determined to exactly, if possible, reconstruct the high-order mesh.

Robustness

Recall from Section 1 and 3.2 that robustness is a key consideration in the design of the proposed method. In particular, the objective function aligns the faces of the mesh with the discontinuities to avoid severe oscillations in the solution and a full space optimization approach is used to ensure the PDE solution is never required at non-aligned meshes. These measures effectively solve the robustness issue of traditional steady-state solvers provided the optimization problem is properly initialized as discussed in the previous section.

Efficiency

A natural question regarding efficiency of the proposed method arises given the shift from a traditional nonlinear system of equations in $N_{\boldsymbol{u}}$ variables to a nonlinearly constrained optimization problem in $N_{\boldsymbol{u}} + N_\phi$ variables and $N_{\boldsymbol{u}}$ constraints. In fact, the proposed discontinuity-tracking method is expected to perform similarly to a traditional steady state solver *on a given mesh*. This comes from the observation that *per iteration* an implicit steady state solver will require evaluation of the residual and its Jacobian with respect to $\boldsymbol{u}$ and a linear solve with this Jacobian. A full space solver will also require these operations as well as the additional terms in (38), i.e., the objective function and its derivatives and the Jacobian of the residual with respect to $\boldsymbol{x}$. The evaluation of the objective function and its derivatives are inexpensive operations: they require a single pass over the mesh, have obvious data structures, and a parallel implementation is straightforward. The evaluation of the Jacobian of the residual with respect to $\boldsymbol{x}$ will incur a cost similar to the evaluation of the Jacobian with respect to $\boldsymbol{u}$.

Despite a similar per iteration performance on a given mesh, the proposed method that tracks discontinuities in the solution is expected to achieve *orders of magnitude* reduction in degrees of freedom compared to traditional shock capturing methods such as limiting and artificial viscosity, even when combined with adaptive mesh refinement. This reduction is expected because the discontinuity is perfectly represented by the solution basis given there is sufficient resolution for the isoparametric elements to track the *geometry* of the shock. Thus mesh resolution is driven by the geometry of the shock and the smooth features in the solution away from the shock, which allows for very coarse discretizations on high-order meshes as will be seen in Section 4.

3.5. Summary of proposed discontinuity-tracking method

Sections 3.1-3.4 provide a complete description of the proposed optimization-based discontinuity-tracking method. The key ingredients of the proposed method are:

- a high-order discretization of steady conservation laws in (1) that supports discontinuities along element faces and is parametrized by the positions of the high-order mesh nodes; in this work we use a high-order discontinuous Galerkin method with isoparametric elements

- the PDE-constrained optimization formulation in (17)

- a parametrization of the mesh as $\boldsymbol{x} = \mathcal{A}(\boldsymbol{\phi})$ where $\boldsymbol{\phi}$ is a reduced set of optimization parameters and $\mathcal{A}$ maps them to the nodes of the high-order mesh, possibly including mesh smoothing; the most general choice is the identity map, with other options presented in Section 3.2.1 and 4.4

- the objective function in (21) that penalizes spurious oscillations and mesh distortion; numerical experiments on one-dimensional test cases in Section 3.2.1- 3.2.2 show this indicators has a local minima at discontinuity-aligned mesh and decreases monotonically to this minima in a neighborhood of radius approximately $h/2$

- the full space PDE-constrained optimization solver that simultaneously converges the solution and mesh to their optimal values without ever requiring the PDE solution on a non-aligned mesh and

- the practical issue of the initialization of the non-convex optimization problem in (17) with *viscosity* solutions of the conservation law or the result of the discontinuity tracking method at a lower polynomial order.

4. Numerical experiments

4.1. L^2 projection of discontinuous function onto piecewise polynomial basis

In this section, the proposed framework is investigated without the complication of a conservation law. Instead, an L^2 projection of a given discontinuous function onto a piecewise polynomial space $\hat{\mathcal{V}}_{h,p}(\boldsymbol{x})$ induced by the discretization $\mathcal{E}_{h,p}$ of the reference domain Ω_0 is considered. The purpose of this section is three-fold: 1) demonstrate the merit of the proposed framework across a range of mesh sizes and polynomial orders, 2) study the convergence behavior of the full space approach, and 3) verify optimal $p+1$ convergence rates for the proposed discontinuity-tracking framework and compare to alternative methods.

The L^2 projection of a function $\eta \in L^2(\Omega)$ onto $\hat{\mathcal{V}}_{h,p}(\boldsymbol{x})$ is defined as the function $u \in \hat{\mathcal{V}}_{h,p}(\boldsymbol{x})$ such that

$$\int_\Omega \psi u \, dV = \int_\Omega \psi \eta \, dV, \qquad \forall \psi \in \hat{\mathcal{V}}_{h,p}(\boldsymbol{x}). \tag{46}$$

Expansion of the solution and test function in the basis $\{\varphi_i\}$ according to (12) and (13) and restriction to element $K \in \mathcal{E}_{h,p}$ leads to the discrete problem

$$\left(\int_{\mathcal{G}(K,\boldsymbol{x})} \varphi_i^K \varphi_j^K \, dV \right) \boldsymbol{u}_i^K = \int_{\mathcal{G}(K,\boldsymbol{x})} \varphi_i^K \eta \, dV, \qquad i \in \{1, \ldots, N_p\} \tag{47}$$

that can be written compactly as

$$\boldsymbol{M}^K(\boldsymbol{x}) \cdot \boldsymbol{u}^K = \boldsymbol{\eta}^K(\boldsymbol{x}), \tag{48}$$

where the element mass matrix and contribution of the function η are

$$\boldsymbol{M}_{ij}^K(\boldsymbol{x}) = \int_{\mathcal{G}(K,\boldsymbol{x})} \varphi_i^K \varphi_j^K \, dV, \qquad \boldsymbol{\eta}_i^K(\boldsymbol{x}) = \int_{\mathcal{G}(K,\boldsymbol{x})} \varphi_i^K \eta \, dV, \qquad i,j \in \{1, \ldots, N_p\}. \tag{49}$$
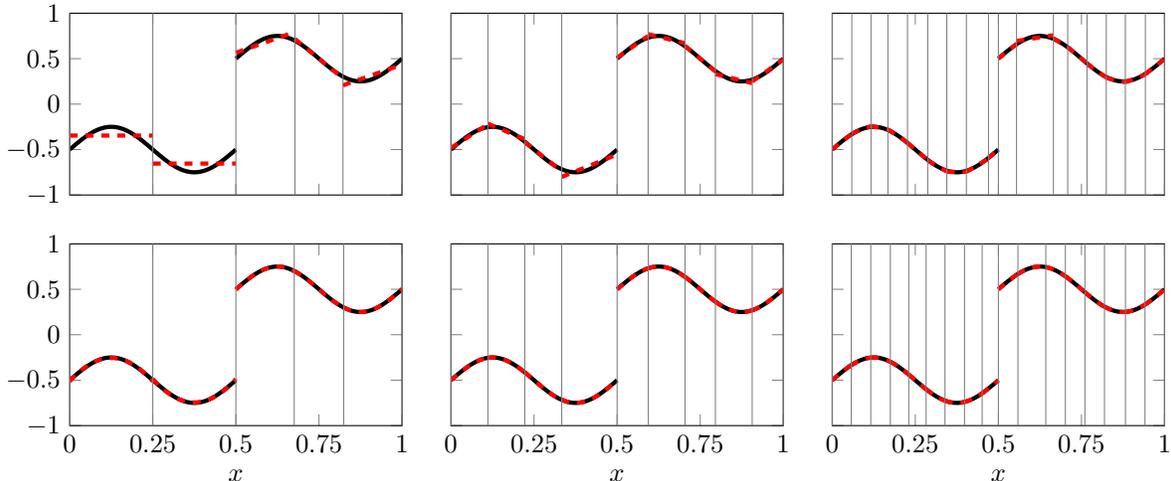
Figure 7: Results of discontinuity-tracking framework using the objective function defined in (21) with $\alpha = 1$ applied to the L^2 projection of the discontinuous function in (51) onto discretizations with $|\mathcal{E}_{h,p}| = 5$ (*left column*), $|\mathcal{E}_{h,p}| = 9$ (*middle column*), $|\mathcal{E}_{h,p}| = 17$ (*right column*) elements and polynomial order $p = 1$ (*top row*) and $p = 3$ (*bottom row*). The $p = 2$ results are excluded for brevity as they are visually identical to the $p = 3$ results. Legend: piecewise sinusoidal function in (51) (——), mesh (——) and solution (- - -) obtained from discontinuity-tracking framework.

Therefore, the discrete residual for the case of an L^2 projection takes the form

$$\boldsymbol{r}(\boldsymbol{u}, \boldsymbol{x}) = \boldsymbol{M}(\boldsymbol{x}) \cdot \boldsymbol{u} - \boldsymbol{\eta}(\boldsymbol{x}) = \boldsymbol{0}. \tag{50}$$

In general, discontinuities present in η will not align with element faces and care must be taken to ensure the right-hand side of (47) is integrated accurately using quadrature. In this work, for any element $K \in \mathcal{E}_{h,p}$ such that $\mathcal{G}(K, \boldsymbol{x})$ contains discontinuities of η, we compute special quadrature rules that do not traverse these discontinuities. This special case of a L^2 projection does not require the full space solver advocated in Section 3.3 since the governing equation in (50) can be solved given any $\boldsymbol{x}$ that does not invert the mesh, i.e., the implicit function $\boldsymbol{u}(\boldsymbol{x})$ is well-defined, and the reduced space approach is a valid option. However, we use the full space approach to test the proposed method in this simple scenario. The mesh $\boldsymbol{x}$ is initialized as a uniform mesh and $\boldsymbol{u}$ is initialized as the solution of (50) on the uniform mesh, regardless of polynomial order.

We begin by considering the discontinuous, piecewise sinusoidal function in one spatial dimension:

$$\eta(x) = \begin{cases} 0.25 \sin(4\pi x) - 0.5, & x < 0 \\ 0.25 \sin(4\pi x) + 0.5, & x > 0 \end{cases} \tag{51}$$

for all $x \in \Omega = (-1, 1)$. The one-dimensional mesh $\mathcal{E}_{h,p}$ is parametrized with one degree of freedom corresponding to each *non-boundary* node in the *continuous, high-order* mesh. The results of the discontinuity-tracking method, applied to the L^2 projection of (51), are shown in Figures 7-8 for a range of mesh sizes and polynomial orders. For an appropriate choice of α ($\alpha = 1$ in this case), the discontinuity is tracked and the solution is well-resolved with as few as 5 elements of degree $p = 3$. If α is too small, mesh regularization is ignored and the element containing the discontinuity is crushed. If α is too large, discontinuity-tracking is ignored and the mesh remains uniform.

The convergence history of the full space optimization solver corresponding to the mesh with 17 $p = 1$ elements for $\alpha = 1$ is provided in Figure 9. It shows the solution of the discretized PDE is only required *at convergence*, i.e., once the objective is minimized and the grid is aligned with the discontinuities. In fact, the PDE constraint violation is relatively large until the objective function has been reduced to nearly its final value and the discontinuity is closely tracked. This is a feature of the full space approach and
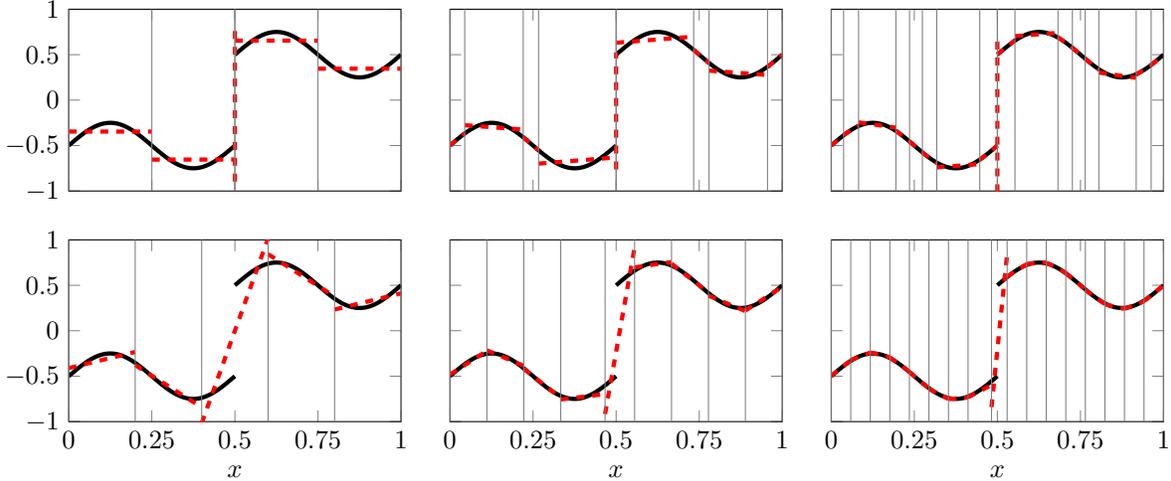
Figure 8: Results of discontinuity-tracking framework using the objective function defined in (21) with $\alpha = 0$ (*top row*) and $\alpha = 10$ (*bottom row*) applied to the L^2 projection of the discontinuous function in (51) onto discretizations with $|\mathcal{E}_{h,p}| = 5$ (*left column*), $|\mathcal{E}_{h,p}| = 9$ (*middle column*), $|\mathcal{E}_{h,p}| = 17$ (*right column*) elements and polynomial order $p = 1$. Legend: piecewise sinuosity function in (51) (——), mesh (——) and solution (- - -) obtained from discontinuity-tracking framework.

provides robustness not possible with the reduced space approach since a PDE solution is never required on non-aligned meshes.

The convergence of the L^2 projection of (51) onto a sequences of increasingly refined meshes with polynomial orders $p = 1, 2, 3, 4, 5, 6$ to the exact function is provided in Figure 10. To ensure the mesh regularization term does not impede convergence, the regularization parameter is set to zero ($\alpha = 0$) and the single degree of freedom parametrization of (2) is used. The L^1 error

$$e_{L^1}(u, u_{h,p}) = \int_\Omega |u - u_{h,p}| \, dV \tag{52}$$

is used to measure convergence of the finite-dimensional solution $u_{h,p}$ to the analytical solution u given by (51). Despite the L^1 norm being the weakest norm, it is appropriate to study problems with discontinuities as it encapsulates the error in the smooth region of the solution and discontinuity location and is expected to converge at the appropriate rate of $\mathcal{O}(h^{p+1})$ if the discontinuity location converges at this rate. Figure 10 shows the expected convergence rate for polynomial orders up to $p = 6$ and a highly accurate solution is possible with as few as 4 $p = 6$ elements.

With the convergence rates of the proposed method established, we compare its performance to alternative methods: 1) uniform refinement and 2) adaptive mesh refinement. The mesh adaptation algorithm uses a straight-forward refinement strategy where elements with large values of the residual in an enriched (h-refined) test space are split. The comparison is provided in Figure 11 for $p = 1, 3, 5$. As expected, uniform refinement is only first-order $\mathcal{O}(h)$ accurate and the adaptive mesh refinement achieves the optimal order of accuracy $\mathcal{O}(h^{p+1})$ eventually. Our discontinuity-tracking provides better accuracy for a given number of elements across a range of mesh sizes and polynomial orders since the discontinuity is accurately tracked and the basis functions are optimally utilized in regions of the domain where the solution is smooth.

To close this section, we consider the L^2 projection of the discontinuous, piecewise constant function in two spatial dimensions where the discontinuity surface is a circle

$$\eta(x) = \begin{cases} 2, & x^2 + y^2 < r^2 \\ 1, & x^2 + y^2 > r^2 \end{cases} \tag{53}$$

for all $(x, y) \in \Omega = (0, 1) \times (0, 1)$. Due to the difficulty associated with producing quadrature rules that do not cross discontinuities in two-dimensions on curved meshes, the discontinuous function in (53) is replaced

22
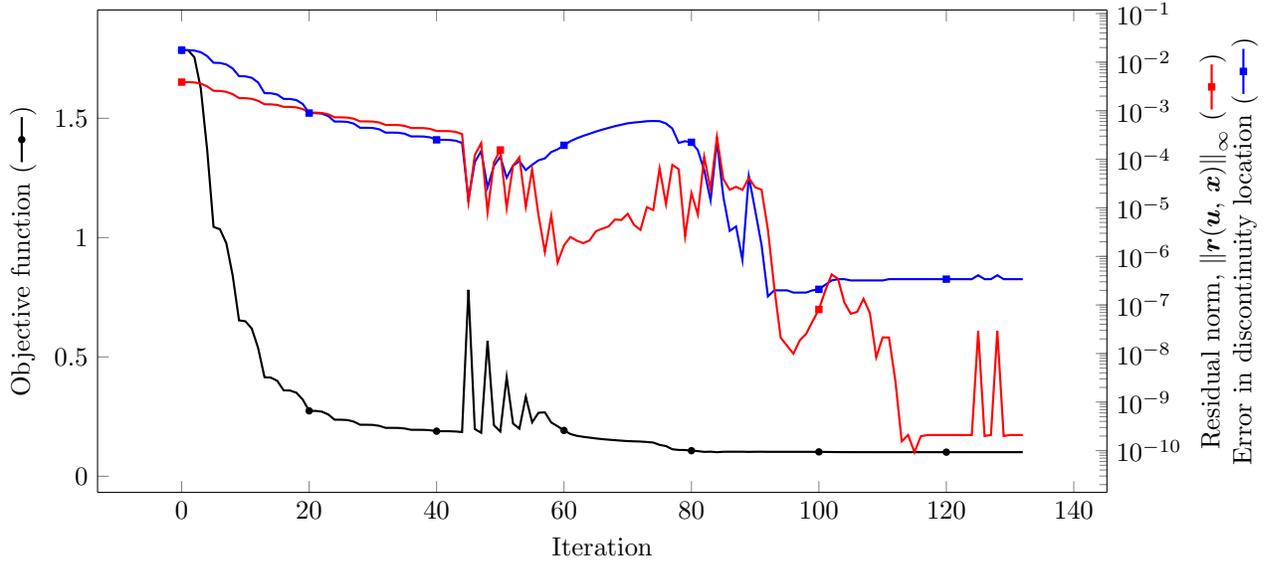
Figure 9: Convergence history of full space optimization solver (SNOPT [19]) applied to discontinuity-tracking optimization problem in (18) with objective function in (21) ($\alpha = 1$) corresponding to the L^2 projection of the discontinuous function in (51) onto a mesh with $|\mathcal{E}_{h,p}| = 17$, $p = 1$ elements. Legend: Value of the objective function (——•——, *left axis*), the residual norm $\|\mathbf{r}(\mathbf{u}, \mathbf{x})\|_\infty$ (——■——, *right axis*), and error in the location of the discontinuity (——■——, *right axis*).
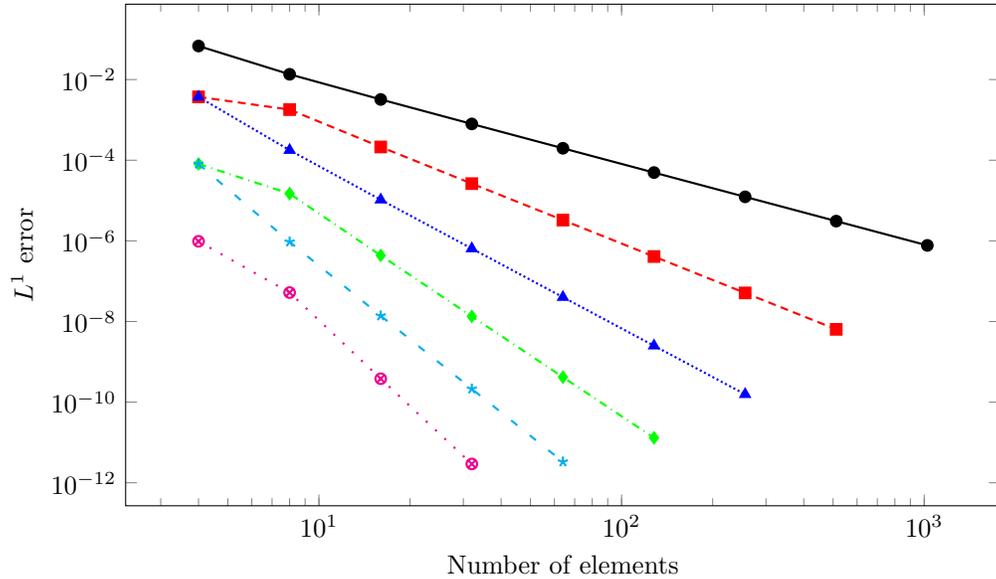
Figure 10: Convergence of discontinuity-tracking method using the single degree of freedom parametrization in (25) and no mesh regularization ($\alpha = 0$) applied to the L^2 projection of the discontinuous function in (51) for polynomial orders $p = 1$ (——•——), $p = 2$ (– –■– –), $p = 3$ (·····▲·····), $p = 4$ (–·–◆–·–), $p = 5$ (– ∗ –), $p = 6$ (·⊗·). The expected convergence rates of $p + 1$ are obtained in all cases.

23
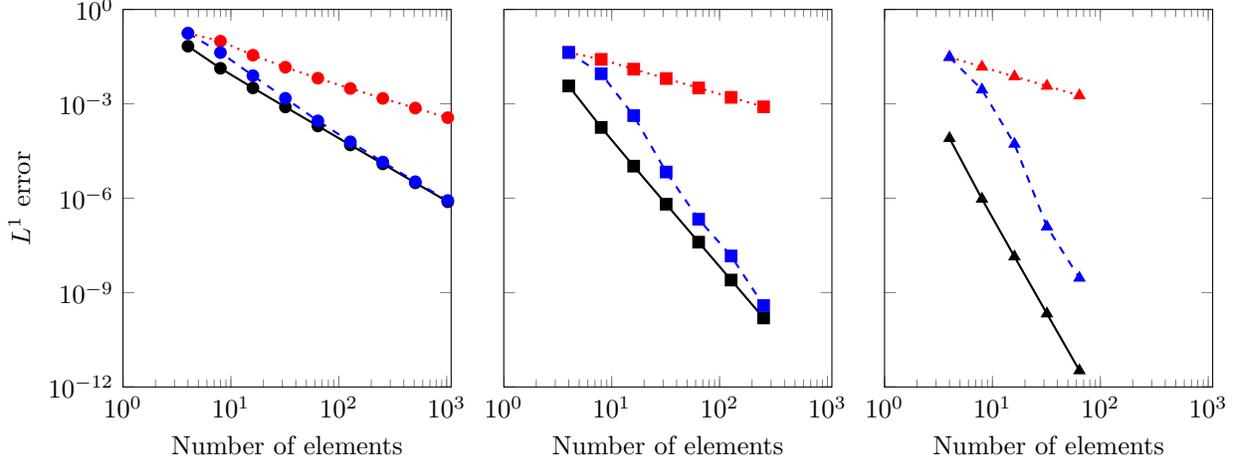
Figure 11: Comparison of convergence of discontinuity-tracking method (single degree of freedom parametrization in (25) and no mesh regularization, $\alpha = 0$) to uniform and adaptive mesh refinement at fixed polynomial orders for the L^2 projection of the discontinuous function in (51). The discontinuity-tracking method and adaptive mesh refinement achieve the expected $p + 1$ convergence rate in the asymptotic regime while uniform refinement is limited to first order due to poor shock resolution. Legend: discontinuity-tracking with $p = 1$ (—●—), $p = 3$ (—■—), $p = 5$ (—▲—); uniform refinement with $p = 1$ (··●··), $p = 3$ (··■··), $p = 5$ (··▲··); and adaptive mesh refinement with $p = 1$ (-●-), $p = 3$ (-■-), $p = 5$ (-▲-).

with a smooth approximation

$$\eta_\nu(x) = 1 + \frac{1}{2}\left[1 - \tanh\left(\nu^{-1}\left(\sqrt{(x - c_x)^2 + (y - c_y)^2} - r\right)\right)\right] \tag{54}$$

where ν is a viscosity-like parameter that controls the steepness of the smoothed discontinuity. In this work, $\nu = 10^{-3}$ is used. To ensure the term on the right-hand-side of (47) is integrated accurately, a Gaussian quadrature rule with 20 quadrature points is used. The results of the discontinuity-tracking framework applied to the L^2 projection of (54) onto a coarse mesh of only 50 elements with polynomial orders $p = 1, 2, 3$ is shown in Figure 12. The mesh is parametrized using all continuous, high-order, non-boundary mesh nodes for the most general parametrization. The L^2 projection onto a uniform, non-aligned mesh is clearly a poor approximation to the piecewise constant function where the discontinuity surface is a circle. Despite being a poor approximation to the function in (53), this mesh and solution are used as the initial guess for the discontinuity-tracking method with the $p = 1$ mesh. Subsequently, the mesh and solution produced by our discontinuity-tracking method using polynomial order $p = \bar{p}$ is used to initialize the method with polynomial order $p = \bar{p} + 1$. The coarse mesh with $p = 2$ and $p = 3$ track the discontinuity very well, sufficiently regularize the mesh, and accurately approximate the piecewise constant function in (53).

4.2. Inviscid Burgers' equation with discontinuous source term

In this section, the proposed discontinuity-tracking framework is applied to the inviscid, modified one-dimensional Burgers' equation with a discontinuous source term from [5]

$$\frac{\partial}{\partial x}\left(\frac{1}{2}u^2\right) = \beta u + f(x), \quad \text{for } x \in \Omega \subset \mathbb{R}, \tag{55}$$

where $\beta = -0.1$ and

$$f(x) = \begin{cases} (2 + \sin(\frac{\pi x}{2}))(\frac{\pi}{2}\cos(\frac{\pi x}{2}) - \beta), & x < 0 \\ (2 + \sin(\frac{\pi x}{2}))(\frac{\pi}{2}\cos(\frac{\pi x}{2}) + \beta), & x > 0. \end{cases} \tag{56}$$
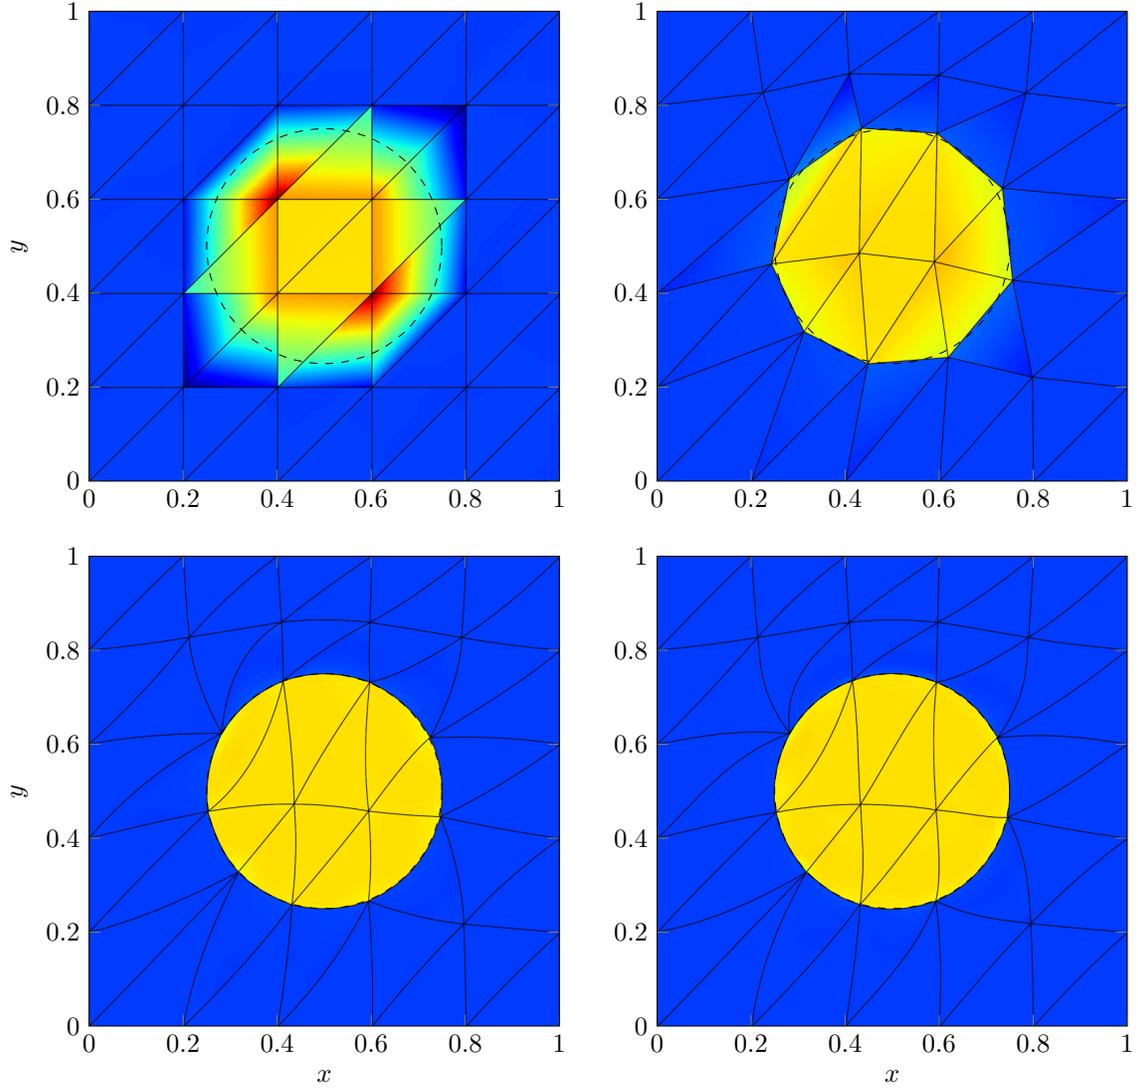
Figure 12: Results of discontinuity-tracking framework using the objective function defined in (21) with $\alpha = 1$ applied to the L^2 projection of the function in (54) with steep gradients onto discretizations with $|\mathcal{E}_{h,p}| = 50$ and polynomial order $p = 1$ (*top row*), $p = 2$ (*bottom left*), and $p = 3$ (*bottom right*). *Top left*: L^2 projection of (54) onto uniform non-aligned mesh with 50 $p = 1$ elements. *Top right*: result of the discontinuity-tracking framework with $p = 1$ elements using the non-aligned mesh and solution in the *top left* as the initial guess. *Bottom left*: result of the discontinuity-tracking framework with $p = 2$ elements using the $p = 1$ aligned mesh and solution in the *top right* as the initial guess. *Bottom right*: result of discontinuity-tracking framework with $p = 3$ elements using the $p = 2$ aligned mesh and solution in the *bottom left* as the initial guess.
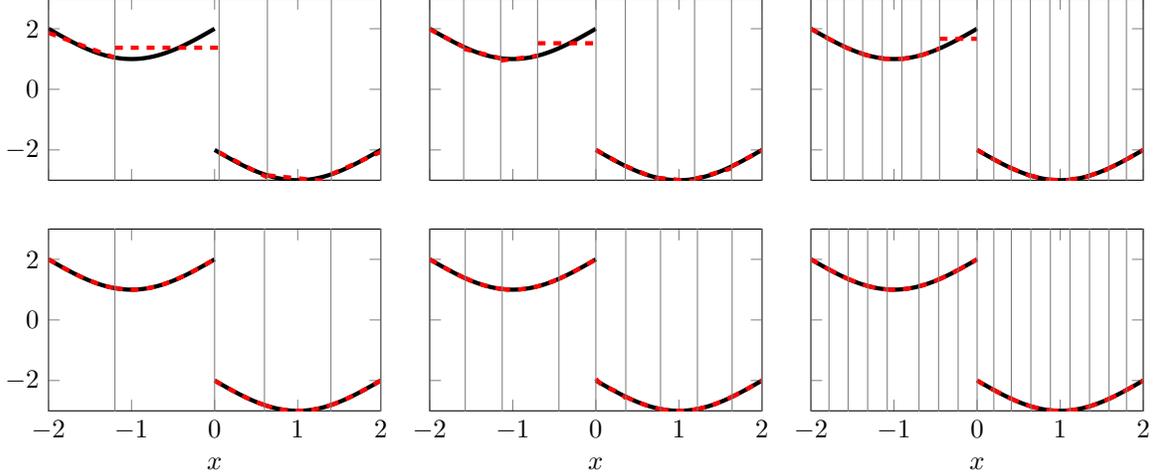
Figure 13: Results of discontinuity-tracking framework (objective function defined in (21) with $\alpha = 0.1$) applied to the solution of the inviscid Burgers' equation (55) onto discretizations with $|\mathcal{E}_{h,p}| = 5$ (*left column*), $|\mathcal{E}_{h,p}| = 9$ (*middle column*), $|\mathcal{E}_{h,p}| = 17$ (*right column*) elements and polynomial order $p = 1$ (*top row*) and $p = 3$ (*bottom row*). The $p = 2$ results are excluded for brevity as they are visually identical to the $p = 3$ results. Legend: exact solution (——) of the modified inviscid Burgers' equation in (55), mesh (——) and solution (- - -) obtained from discontinuity-tracking framework.

It can be shown that the solution of (55) is

$$
u(x) = \begin{cases} 2 + \sin(\frac{\pi x}{2}), & x < 0 \\ -2 - \sin(\frac{\pi x}{2}), & x > 0. \end{cases} \tag{57}
$$

The domain is taken as $\Omega = (-2, 2)$ with Dirichlet boundary conditions specified at both ends corresponding to the exact solution, i.e., $u(-2) = 2$ and $u(2) = -2$. The discretization of (55) proceeds according to the formulation in Section 2 and the numerical flux is taken as the solution of the exact Riemann problem corresponding to the time-dependent version of (55) without source terms ($\beta = 0$ and $f(x) = 0$).

The one-dimensional mesh $\mathcal{E}_{h,p}$ is parametrized with one degree of freedom corresponding to each *non-boundary* node in the *continuous, high-order* mesh. The results of the discontinuity-tracking method, applied to the solution of the modified, inviscid Burgers' equation are shown in Figure 13 for a range of mesh sizes and polynomial orders. For an appropriate choice of α ($\alpha = 0.1$ in this case), the discontinuity is tracked and the solution is well-resolved with as few as 5 $p = 3$ elements.

Convergence of the proposed discontinuity-tracking method to the exact solution in (57) using a sequences of increasingly refined meshes with polynomial orders $p = 1, 2, 3, 4, 5, 6$ is provided in Figure 14. Similarly to the previous section, the L^1 error is used to quantify the error in the solution and the single degree of freedom mesh parametrization (2) without regularization ($\alpha = 0$) is used to ensure mesh regularization does not impede convergence. Figure 14 shows the expected convergence rate for most polynomial orders up to $p = 6$ and a highly accurate solution is possible with as few as 4 $p = 6$ elements. Figure 15 compares the performance of the discontinuity-tracking method to uniform and adaptive mesh refinement. As expected, uniform refinement is only first-order $\mathcal{O}(h)$ accurate and adaptive mesh refinement nearly achieves the optimal order of accuracy $\mathcal{O}(h^{p+1})$. For this test case involving a conservation law, the improved accuracy provided by tracking the discontinuity is more pronounced than the previous case of an L^2 projection of an analytical function, e.g., for the $p = 3$ case, an error of 10^{-4} is possible with only 10 elements when the discontinuity is tracked while adaptive mesh refinement requires over 700 elements.

4.3. Transonic, inviscid flow through nozzle

Our final one-dimensional example considers the relevant situation of transonic, inviscid flow through a converging-diverging nozzle. The quasi-one-dimensional Euler equations are used to model the inviscid,

26
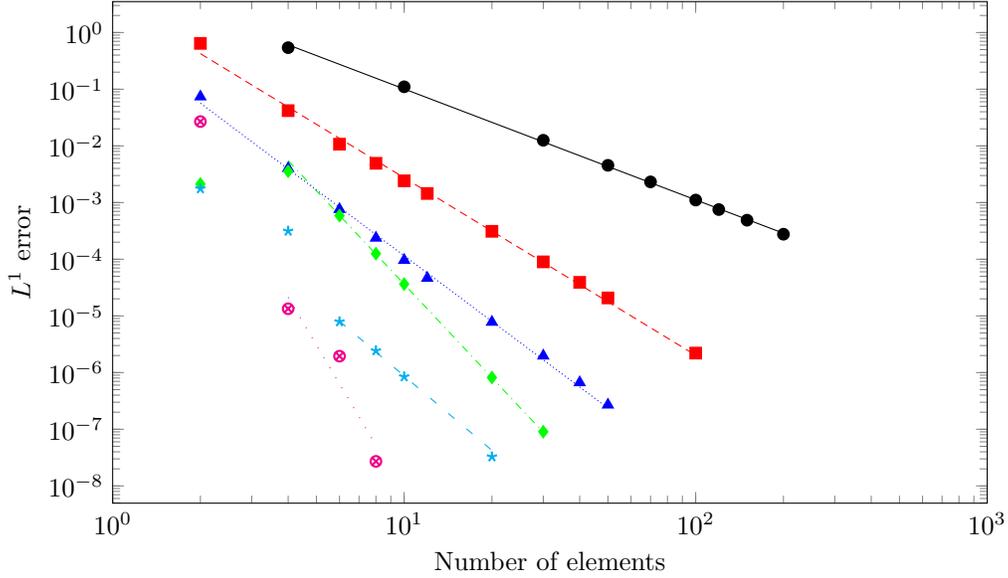
Figure 14: Convergence of discontinuity-tracking method using the single degree of freedom parametrization in (25) and no mesh regularization ($\alpha = 0$) applied to the modified inviscid Burgers' equation in (55) for polynomial orders $p = 1$ (●), $p = 2$ (■), $p = 3$ (▲), $p = 4$ (◆), $p = 5$ (★), $p = 6$ (⊗). The expected convergence rates of $p + 1$ are obtained in most cases. The slopes of the best-fit lines to the data points in the asymptotic regime are: $\angle -1.95$ (——), $\angle -3.13$ (-- --), $\angle -3.85$ (········), $\angle -5.47$ (– · – ·), $\angle -4.36$ (– –), $\angle -8.67$ (· · · ·).
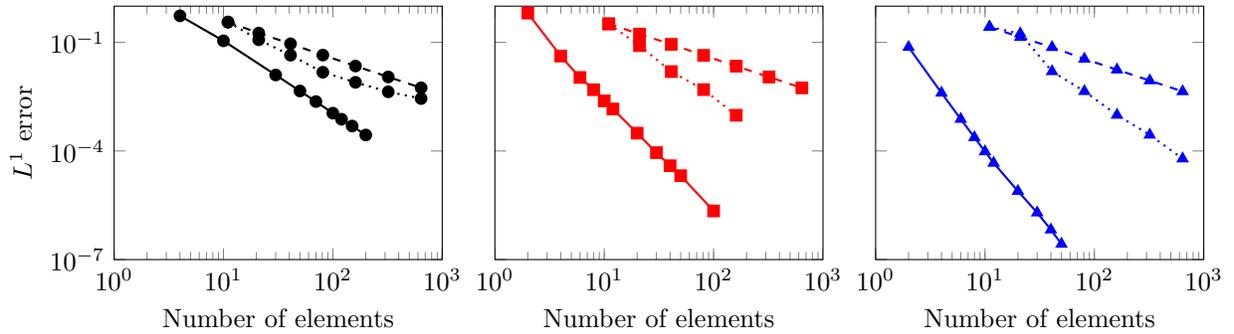
Figure 15: Comparison of convergence of discontinuity-tracking method (single degree of freedom parametrization in (25) and no mesh regularization, $\alpha = 0$) to uniform and adaptive mesh refinement at fixed polynomial orders for modified inviscid Burgers' equation. The discontinuity-tracking method achieves the expected $p + 1$ convergence rate in the asymptotic regime while uniform refinement is limited to first order and adaptive mesh refinement falls short of optimal convergence rates due to poor shock resolution. Legend: discontinuity-tracking with $p = 1$ (—●—), $p = 2$ (—■—), $p = 3$ (—▲—); uniform refinement with $p = 1$ (- ● -), $p = 2$ (- ■ -), $p = 3$ (- ▲ -); and adaptive mesh refinement with $p = 1$ (·· ● ··), $p = 2$ (·· ■ ··), $p = 3$ (·· ▲ ··).

27
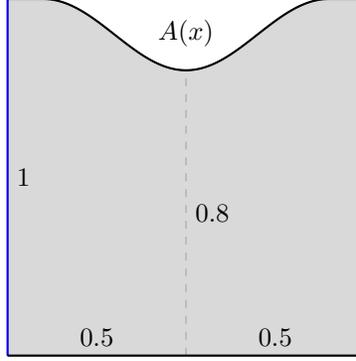
Figure 16: Geometry and boundary conditions for nozzle flow. Boundary conditions: inviscid wall (——), inflow (——), outflow (——).

compressible flow in a variable-area stream tube $A(x)$

$$\frac{\partial}{\partial x}(A\rho u) = 0,$$
$$\frac{\partial}{\partial x}(A[\rho u^2 + p]) = \frac{p}{A}\frac{\partial A}{\partial x},$$
$$\frac{\partial}{\partial x}(A[\rho E + p]u) = 0 \tag{58}$$

where ρ is the fluid density, u is the fluid velocity, p is the thermodynamic pressure, and

$$\rho E = \rho e + \frac{1}{2}\rho u^2 \tag{59}$$

is the total energy. The pressure is related to ρE by the equation of state

$$p = (\gamma - 1)\left(\rho E - \frac{1}{2}\rho u^2\right) \tag{60}$$

for a perfect gas with ratio of specific heats $\gamma = 1.4$. The domain is taken as $\Omega = (0, 1)$ and the nozzle profile takes the form

$$A(x) = \begin{cases} 1 - (1 - T)\cos(\pi(x - 0.5)/0.8)^2 & x \in [0.1,\, 0.9] \\ 1 & \text{otherwise} \end{cases} \tag{61}$$

where $T = 0.8$ is the height of the nozzle throat. The boundary conditions weakly impose the farfield conditions $\rho_i = 1.0$, $u_i = 1.0$, $M_i = 0.40$ at the inflow and $\rho_o = 1.0$, $u_o = 1.0$, $M_o = 0.45$ at the outflow using Roe's approximate Riemann solver. A schematic of the domain, including boundary conditions, is provided in Figure 16. The discretization of the governing equations proceed according to the formulation in Section 2 and Roe's approximate Riemann solver corresponding to the time-dependent version of (58) is used for the numerical flux.

The one-dimensional mesh $\mathcal{E}_{h,p}$ is parametrized with one degree of freedom corresponding to each *non-boundary* node in the *continuous, high-order* mesh. The results of the discontinuity-tracking method, applied to the solution of the quasi-1d Euler equations are shown in Figure 17 for a range of mesh sizes and polynomial orders. For an appropriate choice of α ($\alpha = 0.1$ in this case), the discontinuity is tracked and the solution is well-resolved with as few as 8 $p = 2$ elements.

Convergence of the proposed discontinuity-tracking method to a reference solution using a sequence of increasingly refined meshes with polynomial orders $p = 1, 2$ is provided in Figure 18. The reference solution
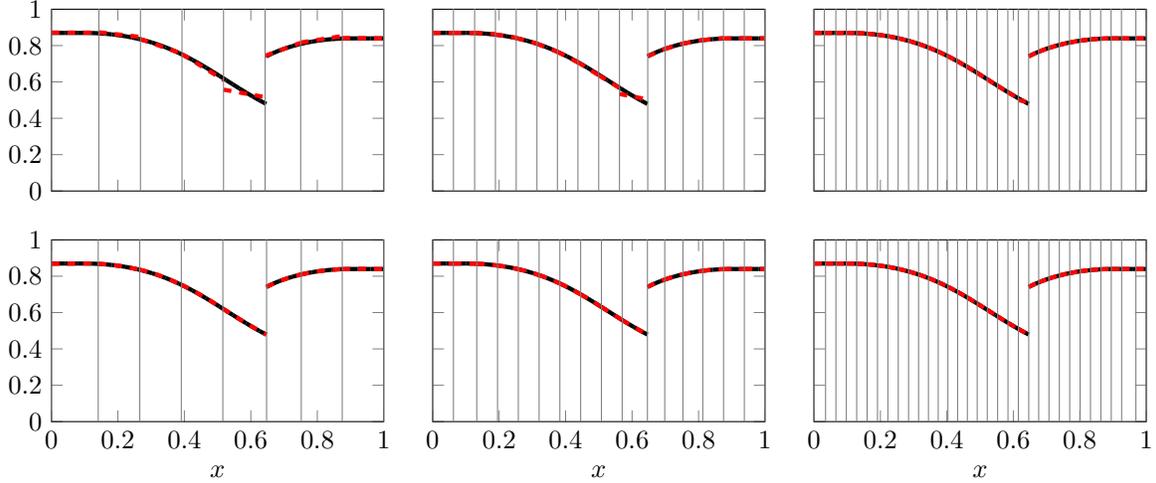
Figure 17: Results of discontinuity-tracking framework (objective function defined in (21) with $\alpha = 0.1$) applied to the quasi-1d Euler equations in (58) with $|\mathcal{E}_{h,p}| = 8$ (*left column*), $|\mathcal{E}_{h,p}| = 16$ (*middle column*), $|\mathcal{E}_{h,p}| = 32$ (*right column*) elements and polynomial order $p = 1$ (*top row*) and $p = 2$ (*bottom row*). Legend: reference solution to the quasi-1d Euler equations determined from aggressive adaptive mesh refinement with over 80000 $p = 1$ elements (——), mesh (——) and solution (- - -) obtained from discontinuity-tracking framework applied to quasi-1d Euler equations.

is computed using DG on a highly adapted mesh near the discontinuity consisting of over 80000 $p = 1$ elements. Similar to the previous sections, the L^1 error is used to quantify the error in the solution. To ensure the mesh regularization term does not impede convergence, the regularization parameter is set to zero ($\alpha = 0$) and the single degree of freedom parametrization of (2) is used. Figure 18 shows the expected convergence rate for polynomial orders $p = 1$, 2 and a highly accurate solution is possible with as few as 20 $p = 2$ elements. The discontinuity-tracking method is shown to outperform the adaptive mesh refinement method with $p = 1$, at least in terms of number of elements for a required accuracy, which is limited to second-order accuracy.

4.4. Supersonic, inviscid flow around cylinder

With the merit of the proposed discontinuity tracking framework established in one spatial dimension, we turn our focus to its performance on two-dimensional conservation laws, specifically the two-dimensional steady Euler equations

$$
\frac{\partial}{\partial x_i}(\rho u_i) = 0,
$$
$$
\frac{\partial}{\partial x_i}(\rho u_i u_j + p) = 0 \quad \text{for } i = 1, 2, \tag{62}
$$
$$
\frac{\partial}{\partial x_i}(u_j(\rho E + p)) = 0,
$$

where ρ is the fluid density, u_1, u_2 are the velocity components, and E is the total energy. For an ideal gas, the pressure p has the form

$$
p = (\gamma - 1)\rho \left(E - \frac{1}{2}u_k u_k \right), \tag{63}
$$

where γ is the adiabatic gas constant. The discretization of the governing equations proceed according to the formulation in Section 2 and Roe's approximate Riemann solver corresponding to the time-dependent version of (62) is used for the numerical flux. All farfield and inviscid wall boundary conditions are weakly imposed using the Roe solver.

The two-dimensional problem considered is supersonic flow around a circle at Mach 2. Due to symmetry, only one quarter of the domain is modeled; see Figure 19 for a schematic of the domain and boundary

29
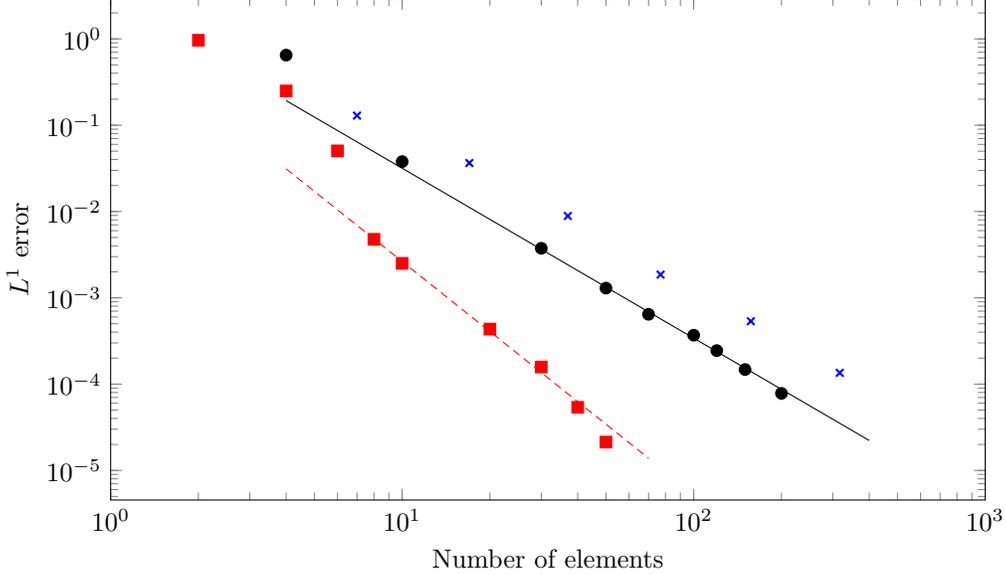
Figure 18: Convergence of discontinuity-tracking method using the single degree of freedom parametrization in (25) and no mesh regularization ($\alpha = 0$) applied to the quasi-1d Euler equations (58) for polynomial orders $p = 1$ (●) and $p = 2$ (■). The expected convergence rates of $p + 1$ are obtained in both cases. The slopes of the best-fit lines to the data points in the asymptotic regime are: $\angle - 1.97$ (——) and $\angle - 2.70$ (‑ ‑ ‑ ‑). A reference second-order method that uses adaptive mesh refinement with $p = 1$ elements (✖).

conditions. For simplicity, we use a simple mesh parametrization that does not include the position of all the nodes of the continuous high-order mesh, but rather a well-chosen subset of these nodes. The remainder of the nodes are determined through a linear operator that incorporates mesh smoothing. In this work, we use linear elasticity with prescribed displacements at the parametrized nodes and in the normal direction along domain boundaries. For this problem, we use a mesh with only $|\mathcal{E}_{h,p}| = 48$ elements and polynomial degree up to $p = 4$. Figure 20 shows these meshes and identifies the parametrized nodes whose displacements compose ϕ and correspond to $N_\phi = 12, 24, 36, 48$ for the $p = 1, 2, 3, 4$ meshes, respectively.

As a non-convex optimization problem underlies the discontinuity-tracking framework, the performance of the full space solver relies on a quality initial guess. For the case of $p = 1$, the initial guess for the mesh is taken as the reference mesh, shown in Figure 20, while for the $p > 1$ case, the mesh is initialized with the solution of the discontinuity-tracking problem on the $p - 1$ mesh. The initial state vector, regardless of
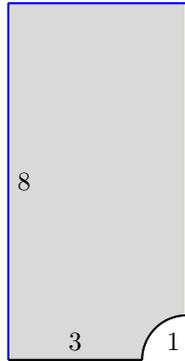
Figure 19: Geometry and boundary conditions for supersonic flow around cylinder. Boundary conditions: inviscid wall/symmetry condition (——) and farfield (——).
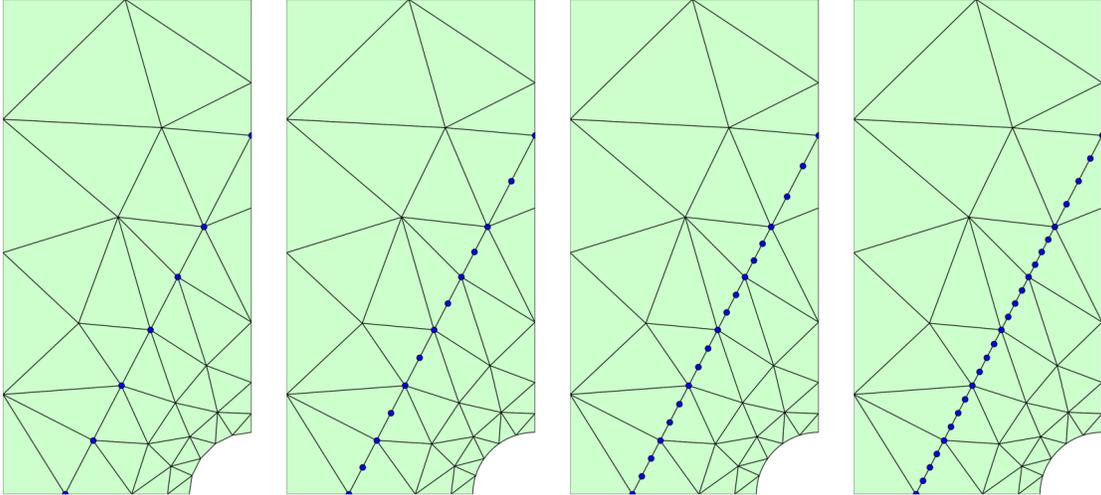
Figure 20: Reference domain for and mesh for supersonic flow around cylinder with $|\mathcal{E}_{h,p}| = 48$ elements and polynomial orders $p = 1$ (*left*), $p = 2$ (*middle left*), $p = 3$ (*middle right*), and $p = 4$ (*right*). The blue circles identify parametrized nodes, i.e., nodal positions that whose displacements compose the optimization variables ϕ. Only the displacement normal to the boundary are taken as optimization variables for the two parametrized nodes that lie on boundaries.

polynomial degree, is taken as a viscosity solution of (62) with a global viscosity parameter $\nu = 0.05$, i.e., a solution of the Naiver-Stokes equations, on the initial mesh, as discussed in Section 3.4.

The result of the discontinuity-tracking method using the above initialization strategy and the reference domain in Figure 20 is provided in Figures 21-22. Figure 21 shows the high-order computational mesh superimposed on the fluid density, element-wise discontinuity indicator (f_{shk}), and element-wise mesh distortion indicator (f_{msh}) for the viscosity solution of the $p = 1$, non-aligned mesh (e.g., initial guess for the $p = 1$ discontinuity-tracking problem) and the inviscid solution using the discontinuity-tracking framework for the $p = 1, 2, 3, 4$ meshes (mesh regularization parameter is set to $\alpha = 0.1$). The non-aligned mesh is clearly underresolved and the shock is smeared out. Several elements contain non-trivial under- and over-shoot, which is identified by the discontinuity indicator that takes large values in elements near the discontinuity. The mesh distortion indicator is identically zero since the reference mesh is used. It is clear that the discontinuity-tracking framework, even for $p = 1$, considerably reduces the overshoot in the solution near the shock at the price of mesh distortion. The $p = 1$ solution is still clearly underresolved and the shock surface is faceted since the element cannot bend in the piecewise linear setting. However, the $p = 2, 3, 4$ reference meshes deliver quite accurate and smooth solutions away from the shock and track the smooth, curved shock surface very well. Since the search space is enriched in the cases relative to the $p = 1$ case, the discontinuity indicator and mesh distortion indicator also improve in most elements. Figure 22 shows the fluid density corresponding to the non-aligned, viscosity solver and the inviscid, tracking framework without the mesh edges to emphasize the smooth, high-quality solution obtained away from the shock.

To quantify the quality of the solution obtained from the discontinuity-tracking framework, we consider two metrics from inviscid fluid mechanics: the total enthalpy $H = (\rho E + p)/\rho$, and the stagnation pressure. In a steady, inviscid flow the total enthalpy is constant and the stagnation pressure, p_0, is given by

$$p_0 = \bar{p} \frac{1 - \gamma + 2\gamma \bar{M}^2}{\gamma + 1} \left(\frac{(\gamma + 1)^2 \bar{M}^2}{4\gamma \bar{M}^2 - 2(\gamma - 1)} \right)^{\frac{\gamma}{\gamma - 1}}, \tag{64}$$

where $\bar{p}$ and $\bar{M}$ are the freestream pressure and Mach number, respectively. Therefore we use the following error metrics to quantify the performance of the discontinuity tracking framework without requiring a
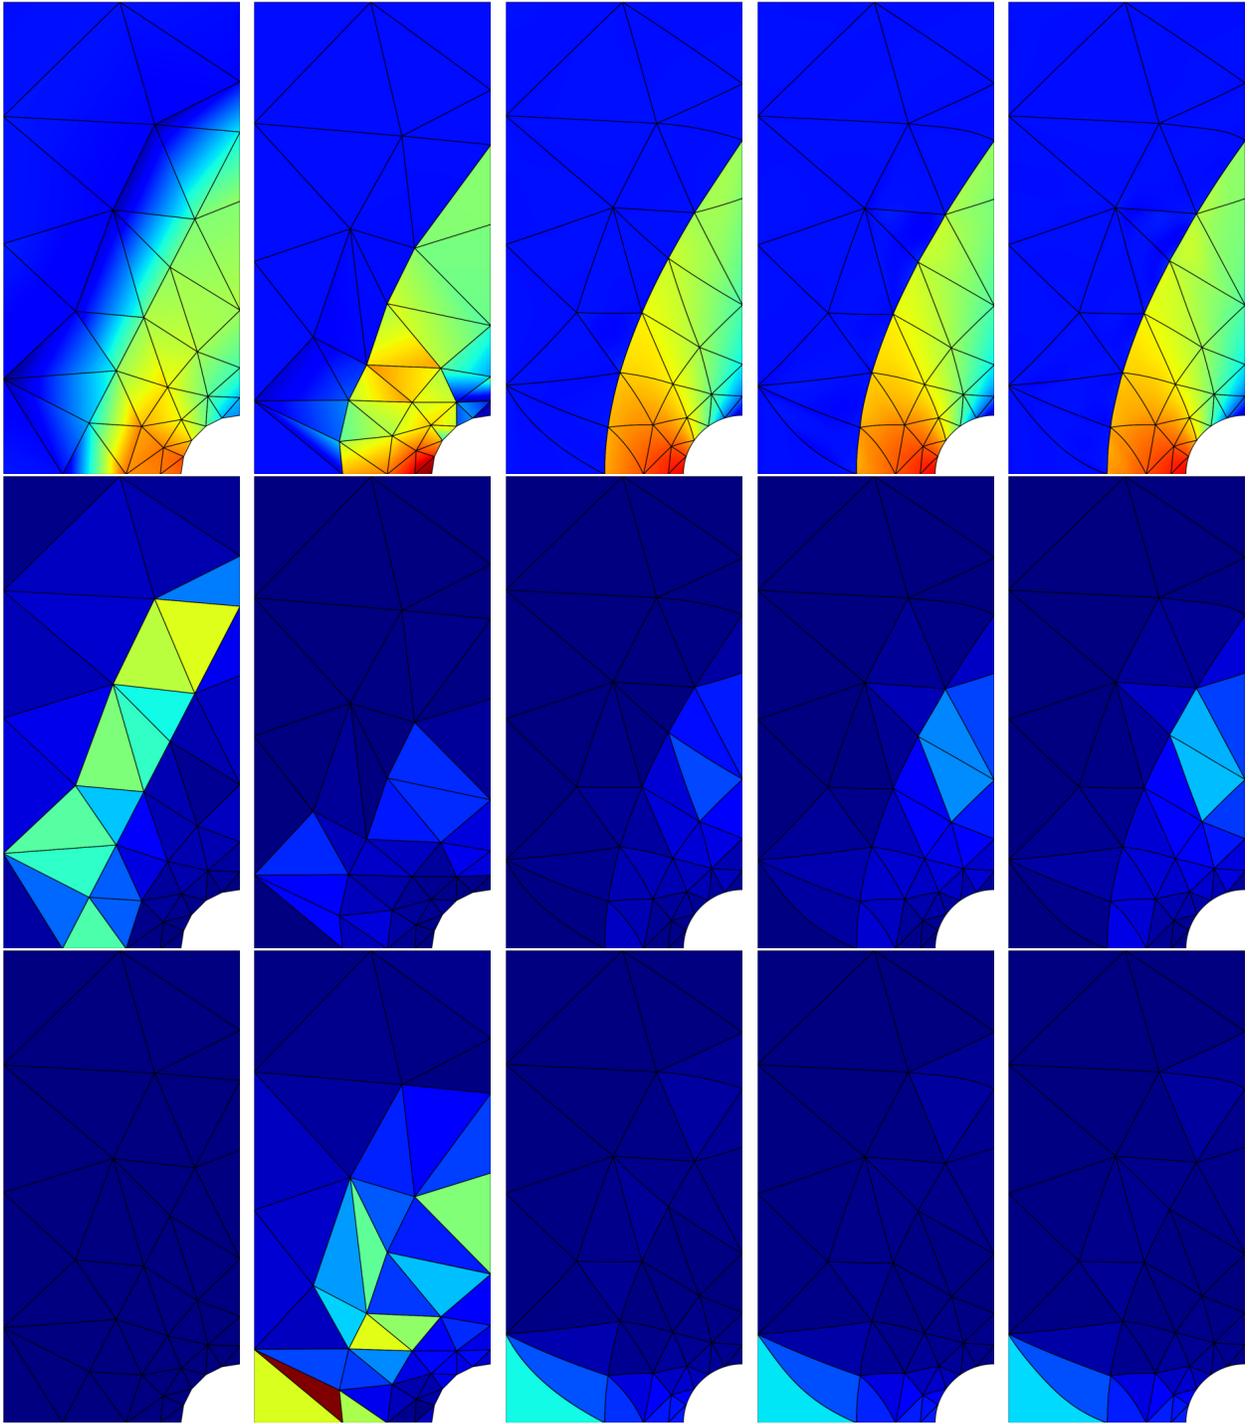
Figure 21: Discontinuity-tracking results corresponding to the reference meshes in Figure 20 with mesh edges superimposed. *Rows*: fluid density (*top*), element-wise discontinuity indicator (*middle*), element-wise mesh distortion indicator (*bottom*). *Columns*: viscosity solution on the non-aligned $p = 1$ mesh (*left*) and the inviscid, discontinuity-tracking solution on the $p = 1$ (*middle left*), $p = 2$ (*middle*), $p = 3$ (*middle right*), $p = 4$ (*right*) meshes with mesh regularization $\alpha = 0.1$.
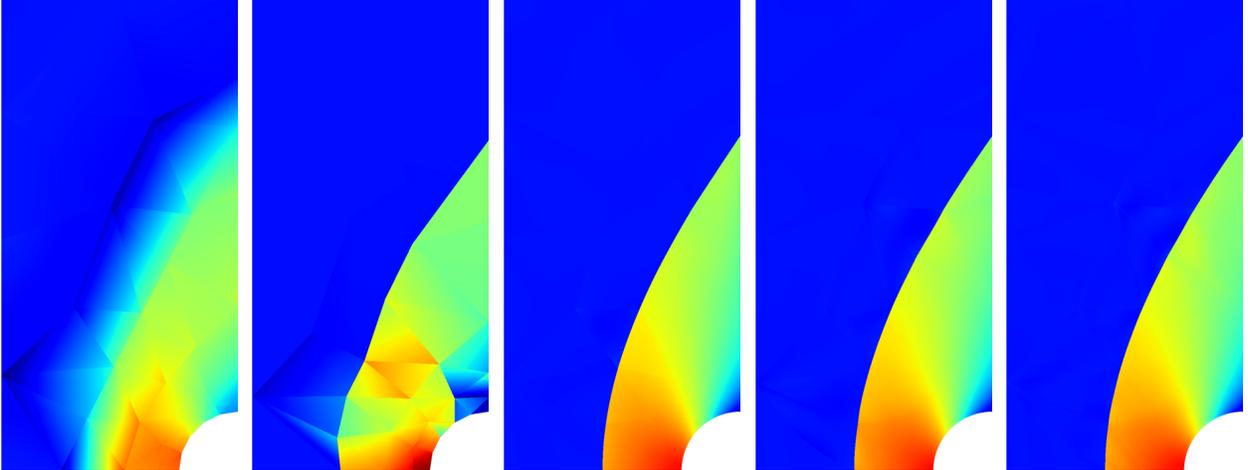
Figure 22: Fluid density corresponding to viscosity solution on the non-aligned $p = 1$ mesh (*left*) and the inviscid, discontinuity-tracking solution on the $p = 1$ (*middle left*), $p = 2$ (*middle*), $p = 3$ (*middle right*), $p = 4$ (*right*) meshes with mesh regularization $\alpha = 0.1$.

reference solution

$$e_p(\boldsymbol{u}, \boldsymbol{x}) = |p_s(\boldsymbol{u}, \boldsymbol{x}) - p_0|$$

$$e_H(\boldsymbol{u}, \boldsymbol{x}) = \sqrt{\frac{\int_\Omega (H(\boldsymbol{u}, \boldsymbol{x}) - H_0)^2 \, dv}{\int_\Omega dv}}, \tag{65}$$

where $p_s(\boldsymbol{u}, \boldsymbol{x})$ is the pressure at the stagnation point, p_0 is the analytical stagnation pressure (64), $H(\boldsymbol{u}, \boldsymbol{x})$ is the total enthalpy, and H_0 is the farfield total enthalpy. These quantities are summarized in Table 1 for the discontinuity-tracking framework at all polynomial orders considered. The table shows the discontinuity-

Polynomial order (p)	1	2	3	4
Degrees of freedom $(N_{\boldsymbol{u}})$	576	1152	1920	2880
Enthalpy error (e_H)	0.0106	0.000462	0.00151	0.000885
Stagnation pressure error (e_p)	0.0711	0.00479	0.0112	0.000616

Table 1: Discontinuity-tracking performance summary, including the number of degrees of freedom and enthalpy and stagnation pressure error for the reference mesh in Figure 20 with 48 elements and polynomials orders $p = 1$, $p = 2$, $p = 3$, $p = 4$ with mesh regularization $\alpha = 0.1$.

tracking framework provides accurate solutions on discretizations with very few degrees of freedom, e.g., $N_{\boldsymbol{u}} \sim \mathcal{O}(10^3)$ with errors on the order of $e_H \sim \mathcal{O}(10^{-3})$.

5. Conclusion

This paper introduces a high-order accurate, nonlinearly stable discontinuity-tracking framework for solving steady conservation laws with discontinuous solution features such as shock waves. The method leverages the discontinuities between computational elements present in the finite-dimensional solution basis in the context of a discontinuous Galerkin or finite volume discretizations to track discontinuities in the underlying solution. The approximate Riemann solver ensures the numerical flux is consistent with the governing equations and provides appropriate stabilization through upwinding. Central to the proposed tracking framework is a PDE-constrained optimization formulation of the discrete conservation law whose objective drives the computational mesh to align with discontinuities and constraints ensure the discrete conservation law is satisfied. The proposed objective function not only attains its minimum when the mesh

33

is aligned with discontinuities, but also approaches the minimum monotonically in a neighborhood of radius $h/2$, which is critical for a gradient-based optimizer to locate it. A mesh regularization term is included to avoid poor-quality meshes and is shown to not destroy the minima or monotonicity of the indicator, for reasonable values of the mesh quality parameter α. The optimization problem is solved using a full space optimization solver whereby the PDE solution and mesh simultaneously converge to their optimal values, ensuring the discrete PDE solution is never required on non-aligned meshes and avoiding non-robustness issues that arise due to Gibbs' phenomena in these situations. This simultaneous convergence behavior of the full space approach is demonstrated in Figure 9 where the discrete conservation law residual is quite large until the objective is sufficiently reduced, i.e., the discontinuity is tracked.

A host of numerical experiments in one and two spatial dimensions are provided to demonstrate the merit of proposed framework. Simple one-dimensional tests problems are provided in Section 3.2 to show the objective function in (19) possesses the appropriate minima and monotonicity property, whereas other popular shock indicators have the appropriate minima, but are highly oscillatory even when the mesh is nearly aligned with the discontinuity. The entire tracking framework is studied on a slew of one- and two-dimensional test problems, including the L^2 projection of discontinuous functions onto a piecewise polynomial basis, the modified steady inviscid Burgers' equation with a discontinuous source term, the quasi-one-dimensional Euler equations (transonic, inviscid flow through a nozzle), and the supersonic two-dimensional Euler equations. We showed the framework successfully tracks discontinuities in the solution even when the nodal positions of the entire mesh are taken as optimization variables, provided the mesh regularization parameter is set appropriately. In the two-dimensional test problems with $p > 1$, the high-order mesh *curves* to provide a high-order alignment with the shock surface. We also showed optimal $\mathcal{O}(h^{p+1})$ convergence rates in the L^1 norm up to $p = 6$ for the L^2 projection problem and inviscid Burgers' equation and showed the discontinuity tracking framework can significantly outperform popular alternatives such as uniform or adaptive mesh refinement. Finally, the discontinuity-tracking framework proved extremely effective in resolving supersonic (Mach 2) flow around a cylinder.

While the proposed discontinuity-tracking framework shows considerable promise as an efficient, high-order, nonlinearly stable method for resolving discontinuous solutions of conservation laws, a number of important research issues must be investigated for this to be a truly viable approach. First, as discussed in Section 3.3, there is considerable structure in the optimization problem (18), e.g., only equality constraints, partition of optimization variables into $\boldsymbol{u}$ and $\boldsymbol{\phi}$, invertible constraint Jacobian with respect to $\boldsymbol{u}$, efficient sparse parallel solvers available for operations involving $\boldsymbol{r}$ and its derivatives, that black-box optimizers cannot leverage. Therefore, there is significant opportunity for efficiency and robustness by developing a specialized solver that can leverage this structure and incorporate well-known homotopy strategies into the constraint since $\boldsymbol{r}(\boldsymbol{u}, \boldsymbol{x}) = 0$ is difficult to solve for high-speed flows that contain shocks. Additionally, it is often the case that, during the full space iterations, the discontinuity indicator and mesh quality can be significantly improved by performing *local topology changes* between elements, for example using the so-called edge flip and face swap operations. Future work will focus on incorporating these discrete mesh updates into the optimization solution procedure since it has the potential to drastically improve the quality of the converged mesh and reduce the number of optimization iterations. We also intend to extend the proposed framework to the *time-dependent* setting, where discrete mesh operations will be crucial. Further research into the choice of numerical flux is also warranted since the inter-element jumps do not tend to zero under refinement along the shock surface, an unfamiliar situation for high-order discontinuous Galerkin methods. While the method was only introduced and tested for inviscid conservation laws, where perfect discontinuities arise, we expect the framework to also be useful for viscous problems with smooth, high-gradient solution features, but further investigation is required. Finally, this paper only considers relatively simple model problems in one and two dimensions. Further testing is required for more complex flows in two and three dimensions and will be the subject of future work.

Acknowledgments

References

[1] Mark Ainsworth and J Tinsley Oden. A posteriori error estimation in finite element analysis. *Computer methods in applied mechanics and engineering*, 142(1-2):1–88, 1997.

[2] Frédéric Alauzet and Adrien Loseille. A decade of progress on anisotropic mesh adaptation for computational fluid dynamics. *Comput.-Aided Des.*, 72:13–39, 2016.

[3] Douglas N Arnold, Franco Brezzi, Bernardo Cockburn, and L Donatella Marini. Unified analysis of discontinuous Galerkin methods for elliptic problems. *SIAM Journal on Numerical Analysis*, 39(5):1749–1779, 2002.

[4] M. J. Baines, S. J. Leary, and M. E. Hubbard. Multidimensional least squares fluctuation distribution schemes with adaptive mesh movement for steady hyperbolic equations. *SIAM Journal on Scientific Computing*, 23(5):1485–1502, 2002.

[5] G. E. Barter. *Shock capturing with PDE-based artificial viscosity for an adaptive, higher-order discontinuous Galerkin finite element method*. PhD thesis, M.I.T., June 2008.

[6] F Bassi and S Rebay. Accurate 2D euler computations by means of a high order discontinuous finite element method. In *Fourteenth International Conference on Numerical Methods in Fluid Dynamics*, pages 234–240. Springer, 1995.

[7] C. E. Baumann and J. T. Oden. A discontinuous *hp* finite element method for the Euler and Navier-Stokes equations. *Int. J. Numer. Methods Fluids*, 31(1):79–95, 1999. Tenth International Conference on Finite Elements in Fluids (Tucson, AZ, 1998).

[8] JB Bell, GR Shubin, and JM Solomon. Fully implicit shock tracking. *Journal of Computational Physics*, 48(2):223–245, 1982.

[9] Ankit Bhagatwala and Sanjiva K Lele. A modified artificial viscosity approach for compressible turbulence simulations. *J. Comput. Phys.*, 228(14):4965–4969, 2009.

[10] George Biros and Omar Ghattas. Parallel lagrange–Newton–Krylov–Schur methods for PDE-constrained optimization. part i: The Krylov–Schur solver. *SIAM Journal on Scientific Computing*, 27(2):687–713, 2005.

[11] Paul T Boggs and Jon W Tolle. Sequential quadratic programming for large-scale nonlinear optimization. *Journal of Computational and Applied Mathematics*, 124(1):123–137, 2000.

[12] A. Burbeau, P. Sagaut, and Ch.-H. Bruneau. A problem-independent limiter for high-order Runge-Kutta discontinuous Galerkin methods. *J. Comput. Phys.*, 169(1):111–150, 2001.

[13] Bernardo Cockburn and Chi-Wang Shu. The Runge-Kutta discontinuous Galerkin method for conservation laws. V. Multidimensional systems. *J. Comput. Phys.*, 141(2):199–224, 1998.

[14] Bernardo Cockburn and Chi-Wang Shu. Runge–Kutta discontinuous Galerkin methods for convection-dominated problems. *Journal of Scientific Computing*, 16(3):173–261, Sep 2001.

[15] Bernardo Cockburn and Chi-Wang Shu. Runge-Kutta discontinuous Galerkin methods for convection-dominated problems. *J. Sci. Comput.*, 16(3):173–261, 2001.

[16] Alain Dervieux, David Leservoisier, Paul-Louis George, and Yves Coudière. About theoretical and practical impact of mesh adaptation on approximation of functions and PDE solutions. *Internat. J. Numer. Methods Fluids*, 43(5):507–516, 2003. ECCOMAS Computational Fluid Dynamics Conference, Part I (Swansea, 2001).

[17] Krzysztof J Fidkowski and David L Darmofal. Review of output-based error estimation and mesh adaptation in computational fluid dynamics. *AIAA Journal*, 49(4):673–694, 2011.

[18] A. Gargallo-Peiró, X. Roca, J. Peraire, and J. Sarrate. A distortion measure to validate and generate curved high-order meshes on CAD surfaces with independence of parameterization. *Internat. J. Numer. Methods Engrg.*, 106(13):1100–1130, 2016.

[19] Philip E Gill, Walter Murray, and Michael A Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Journal on Optimization*, 12(4):979–1006, 2002.

[20] James Glimm, Xiaolin Li, Yingjie Liu, Zhiliang Xu, and Ning Zhao. Conservative front tracking with improved accuracy. *SIAM Journal on Numerical Analysis*, 41(5):1926–1947, 2003.

[21] Max D Gunzburger. *Perspectives in flow control and optimization.* SIAM, 2002.

[22] Ami Harten, Björn Engquist, Stanley Osher, and Sukumar R. Chakravarthy. Uniformly high-order accurate essentially nonoscillatory schemes. III. *J. Comput. Phys.*, 71(2):231–303, 1987.

[23] Ami Harten and James M Hyman. Self adjusting grid methods for one-dimensional hyperbolic conservation laws. *Journal of computational Physics*, 50(2):235–269, 1983.

[24] Ralf Hartmann and Paul Houston. Adaptive discontinuous Galerkin finite element methods for the compressible Euler equations. *J. Comput. Phys.*, 183(2):508–532, 2002.

[25] Michael A Heroux, Roscoe A Bartlett, Vicki E Howle, Robert J Hoekstra, Jonathan J Hu, Tamara G Kolda, Richard B Lehoucq, Kevin R Long, Roger P Pawlowski, Eric T Phipps, et al. An overview of the Trilinos project. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):397–423, 2005.

[26] J. S. Hesthaven and T. Warburton. *Nodal discontinuous Galerkin methods*, volume 54 of *Texts in Applied Mathematics*. Springer, New York, 2008. Algorithms, analysis, and applications.

[27] Michael Hinze, René Pinnau, Michael Ulbrich, and Stefan Ulbrich. *Optimization with PDE constraints*, volume 23. Springer Science & Business Media, 2008.

[28] A. Huerta, E. Casoni, and J. Peraire. A simple shock-capturing technique for high-order discontinuous Galerkin methods. *Internat. J. Numer. Methods Fluids*, 69(10):1614–1632, 2012.

[29] A. Jameson, W. Schmidt, and E. Turkel. Numerical solution of the Euler equations by finite volume methods using Runge Kutta time stepping schemes. In *14th Fluid and Plasma Dynamics Conference*, 1981.

[30] Guang-Shan Jiang and Chi-Wang Shu. Efficient implementation of weighted ENO schemes. *J. Comput. Phys.*, 126(1):202–228, 1996.

[31] Andreas Klöckner, Tim Warburton, and Jan S. Hesthaven. Viscous shock capturing in a time-explicit discontinuous Galerkin method. *Mathematical Modelling of Natural Phenomena*, 6(3):57–83, 2011. arXiv:1102.3190.

[32] Patrick M. Knupp. Algebraic mesh quality metrics. *SIAM Journal on Scientific Computing*, 23(1):193–218, 2001.

[33] L. Krivodonova, J. Xin, J.-F. Remacle, N. Chevaugeon, and J. E. Flaherty. Shock detection and limiting with discontinuous galerkin methods for hyperbolic conservation laws. *Applied Numerical Mathematics*, 48(3):323–338, 2004.

[34] Randall J. LeVeque. *Finite volume methods for hyperbolic problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge, 2002.

[35] Xu-Dong Liu, Stanley Osher, and Tony Chan. Weighted essentially non-oscillatory schemes. *J. Comput. Phys.*, 115(1):200–212, 1994.

[36] D. Moro. *An adaptive high order Reynolds-averaged Navier-Stokes solver with transition prediction*. PhD thesis, M.I.T., December 2014.

[37] Jorge Nocedal and S Wright. *Numerical optimization, series in operations research and financial engineering*. Springer, 2006.

[38] Jayandran Palaniappan, Scott T Miller, and Robert B Haber. Sub-cell shock capturing and spacetime discontinuity tracking for nonlinear conservation laws. *International Journal for Numerical Methods in Fluids*, 57(9):1115–1135, 2008.

[39] P.-O. Persson. Shock capturing for high-order discontinuous Galerkin simulation of transient flow problems. In *21st AIAA Computational Fluid Dynamics Conference, San Diego, CA*, Jun 2013. AIAA-2013-3061.

[40] P-O Persson, J Bonet, and J Peraire. Discontinuous Galerkin solution of the navier–stokes equations on deformable domains. *Computer Methods in Applied Mechanics and Engineering*, 198(17):1585–1595, 2009.

[41] P.-O. Persson and J. Peraire. Sub-cell shock capturing for discontinuous Galerkin methods. In *44th AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada*, 2006. AIAA-2006-0112.

[42] Jianxian Qiu and Chi-Wang Shu. Runge-Kutta discontinuous Galerkin method using WENO limiters. *SIAM J. Sci. Comput.*, 26(3):907–929, 2005.

[43] Philip Roe and Hiroaki Nishikawa. Adaptive grid generation by minimizing residuals. *International Journal for Numerical Methods in Fluids*, 40(1-2):121–136, 2002.

[44] Philip L Roe. Approximate Riemann solvers, parameter vectors, and difference schemes. *Journal of computational physics*, 43(2):357–372, 1981.

[45] GR Shubin, AB Stephens, and HM Glaz. Steady shock tracking and Newton's method applied to one-dimensional duct flow. *Journal of Computational Physics*, 39(2):364–374, 1981.

[46] Greg R Shubin, AB Stephens, HM Glaz, AB Wardlaw, and LB Hackerman. Steady shock tracking, Newton's method, and the supersonic blunt body problem. *SIAM Journal on Scientific and Statistical Computing*, 3(2):127–144, 1982.

[47] F Taghaddosi, WG Habashi, G Guevremont, and D Ait-Ali-Yahia. An adaptive least-squares method for the compressible euler equations. *International Journal for Numerical Methods in Fluids*, 31(7):1121–1139, 1999.

[48] J.-Y. Trepanier, M. Paraschivoiu, M. Reggio, and R. Camarero. A conservative shock fitting method on unstructured grids. *Journal of Computational Physics*, 126(2):421 – 433, 1996.

[49] J. Von Neumann and R. D. Richtmyer. A method for the numerical calculation of hydrodynamic shocks. *J. Appl. Phys.*, 21:232–237, 1950.

[50] Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.

[51] Z. J. Wang, Krzysztof Fidkowski, Rémi Abgrall, and et al. High-order CFD methods: current status and perspective. *Internat. J. Numer. Methods Fluids*, 72(8):811–845, 2013.

[52] V. Wheatley, H. Kumar, and P. Huguenot. On the role of Riemann solvers in discontinuous Galerkin methods for magnetohydrodynamics. *Journal of Computational Physics*, 229(3):660 – 680, 2010.

[53] Masayuki Yano and David L. Darmofal. An optimization-based framework for anisotropic simplex mesh adaptation. *J. Comput. Phys.*, 231(22):7626–7649, 2012.

[54] Matthew J. Zahr and Per-Olof Persson. An adjoint method for a high-order discretization of deforming domain conservation laws for optimization of flow problems. *Journal of Computational Physics*, 326(Supplement C):516 – 543, 2016.

[55] Xiangxiong Zhang and Chi-Wang Shu. On maximum-principle-satisfying high order schemes for scalar conservation laws. *J. Comput. Phys.*, 229(9):3091–3120, 2010.

[56] Xiaolin Zhong. High-order finite-difference schemes for numerical simulation of hypersonic boundary-layer transition. *J. Comput. Phys.*, 144(2):662–709, 1998.

[57] Jun Zhu, Jianxian Qiu, Chi-Wang Shu, and Michael Dumbser. Runge-Kutta discontinuous Galerkin method using WENO limiters. II. Unstructured meshes. *J. Comput. Phys.*, 227(9):4330–4353, 2008.