

An anisotropic hp -adaptation framework for ultraweak discontinuous Petrov–Galerkin formulations

Ankit Chakraborty*, Stefan Henneking, Leszek Demkowicz

Oden Institute, The University of Texas at Austin, United States of America

ARTICLE INFO

Keywords:
DPG
Anisotropy
 hp -adaptivity

ABSTRACT

In this article, we present a three-dimensional anisotropic hp -mesh refinement strategy for ultraweak discontinuous Petrov–Galerkin (DPG) formulations with optimal test functions. The refinement strategy utilizes the built-in residual-based error estimator accompanying the DPG discretization. The refinement strategy is a two-step process: (a) use the built-in error estimator to mark and isotropically hp -refine elements of the (coarse) mesh to generate a finer mesh; (b) use the reference solution on the finer mesh to compute optimal h - and p -refinements of the selected elements in the coarse mesh. The process is repeated with coarse and fine mesh being generated in every adaptation cycle, until a prescribed error tolerance is achieved. We demonstrate the performance of the proposed refinement strategy using several numerical examples on hexahedral meshes.

Acknowledgments: We thank Jacob Badger for fruitful discussions. Ankit Chakraborty, Stefan Henneking and Leszek Demkowicz were supported with NSF award 2103524. Ankit Chakraborty is partially supported with the Peter O'Donnell Jr. Postdoctoral Fellowship. All numerical experiments in this article were performed on Frontera's Intel Cascade Lake (CLX) nodes located at the Texas Advanced Computing Center [1] using DMS22025 allocation.

1. Introduction

Automatic hp -mesh refinement algorithms are powerful tools that aid finite element discretizations in computing solutions of partial differential equations (PDEs) in an efficient and accurate manner. They achieve this efficiency and accuracy by constructing meshes with optimally distributed element size h and polynomial order of approximation p [2,3]. Finite element meshes with optimal element size and polynomial distribution are critical for resolving solution features such as boundary layers in convection-dominated diffusion problems or point and edge singularities in problems with re-entrant corners. In such problems, optimal hp -meshes are indispensable for achieving exponential convergence [4–8]. Designing algorithms capable of generating a sequence of optimal hp -meshes that deliver optimal convergence rates in a problem-agnostic manner has been a significant challenge in finite element research over the past few decades [6,9–17]. Typically,

automatic mesh refinement strategies are driven by computable error estimates. These error estimates are computed using the approximate solution delivered by the discretization scheme. Therefore, the accuracy and stability of the underlying numerical discretization are paramount for the effectiveness of the mesh refinement strategy.

The discontinuous Petrov–Galerkin (DPG) method with optimal test functions, first introduced by Demkowicz and Gopalakrishnan in [18–20], has emerged as a critical technology in terms of robustness and stability over the past decade. Given a stable variational formulation of an underlying PDE and a trial approximation space, the DPG method computes a test space so that the resulting discretization is *inf-sup* stable. The methodology delivers an orthogonal projection in the so-called energy norm. Another significant advantage of the DPG methodology is the presence of a built-in residual-based error estimator, also known as the energy error estimate. This makes the DPG method an ideal candidate for automatic mesh optimization algorithms.

In this article, we focus on the ultraweak DPG finite element formulation with optimal test functions and propose a problem-agnostic anisotropic hp -mesh refinement strategy. It is critical to mention that, for the ultraweak DPG method, the energy norm is *equivalent* to the L^2 -error [21]. Consequently, the method delivers essentially the L^2 -projection of the unknown solution.¹

The proposed refinement strategy consists of the following steps:

* Corresponding author.

E-mail address: ankit.chakraborty@austin.utexas.edu (A. Chakraborty).

¹ See Section 2 for a detailed explanation.

- **Step 1:** Solve the problem on the current *coarse mesh*.
- **Step 2:** Utilize the computed DPG residual to mark coarse mesh elements for refinements.
- **Step 3:** Isotropically *hp*-refine the marked elements to generate a *fine mesh*.
- **Step 4:** Solve the problem on the fine mesh to obtain the fine mesh solution u .
- **Step 5:** Use the fine mesh solution u as a *reference solution* to determine optimal (anisotropic) *hp*-refinements of the selected coarse grid elements.
- **Step 6:** Restore the coarse mesh and execute the optimal *hp*-refinements.

We essentially use the *hp*-algorithm from [2,3]. Optimizing the mesh in the L^2 -space greatly simplifies the original procedure. There is no need for mesh optimization on edges and faces; the *Projection-Based Interpolation* reduces to the L^2 -projection performed on elements only. The optimal refinements of a coarse element K are determined by maximizing the rate (e_{hp}) with which the projection error decreases,

$$e_{hp} := \frac{\|u - P_{\text{coarse}}u\|^2 - \|u - P_{\text{opt}}u\|^2}{N_{\text{opt}} - N_{\text{coarse}}}.$$

Here, P_{coarse} denotes the L^2 -projection onto the coarse mesh, P_{opt} is the projection onto the optimal mesh to be determined, N_{opt} and N_{coarse} denote the number of degrees of freedom (dofs) of the optimal and coarse grid elements, respectively. As the L^2 -projection onto discontinuous polynomial spaces is a purely local operation, the mesh optimization can be trivially performed in parallel.

The main contributions of this paper are: (1) An anisotropic *hp*-mesh refinement strategy that greatly simplifies the algorithm proposed in [2,3] as it *relies solely* on L^2 projections, thereby significantly reducing the implementation complexity of the proposed strategy compared to the approach in [2,3]; (2) Employing the DPG method, and discretization of ultraweak DPG finite element formulations in particular, to enable a robust and reliable L^2 projection-based anisotropic *hp*-mesh refinement strategy that is applicable to *any well-posed* variational problem; (3) Utilizing the DPG residual-based error estimator in Step 3 of the proposed algorithm to compute intermediary fine mesh solutions, which is computationally much more efficient than the global refinement strategy employed in [2,3]; and (4) 3D numerical experiments to demonstrate the efficacy of the proposed anisotropic *hp*-adaptive refinement strategy for problems involving boundary layers and geometric singularities.

The article is organized as follows. Section 2 briefly introduces the ultraweak DPG finite element discretization with optimal test functions. Section 3 provides the details of the mesh optimization algorithm. In Section 4, numerical results show the efficacy of the proposed refinement strategy. Finally, we conclude with a short discussion in Section 5.

2. DPG methodology

The core idea behind the (ideal) DPG method is to automatically generate a stable discretization for a given well-posed variational formulation and an approximate trial space. The method achieves stability by computing an optimal discrete test space [19,22] corresponding to the approximate trial space in such a way that the supremum over the continuous test space in the discrete *inf-sup* [23] is automatically attained over the discrete test space. The optimal test space is obtained by inverting the Riesz map corresponding to the test inner product over a *discontinuous or broken*² test space. Unfortunately, inverting the Riesz operator exactly is impossible due to the infinite-dimensional nature of the continuous test space. Thus, in practical realizations of DPG methods, we approximate the inverse of the Riesz operator by inverting

the Gram matrix induced by the test norm on a larger, but finite-dimensional *enriched* discontinuous test space.³ The use of broken test spaces enables element-wise inversion of the Gram matrix, but it also introduces trace variables defined on the mesh skeleton [24].

We consider a model Poisson problem. Let $\Omega \subset \mathbb{R}^3$ be a bounded Lipschitz domain with boundary Γ split into two disjoint parts: Γ_u and Γ_σ . The first-order formulation of the Poisson problem is given by:

$$\begin{cases} \sigma - \nabla u = \mathbf{0} & \text{in } \Omega, \\ -\nabla \cdot \sigma = f & \text{in } \Omega, \\ u = u_0 & \text{on } \Gamma_u, \\ \sigma \cdot \mathbf{n} = \sigma_0 & \text{on } \Gamma_\sigma, \end{cases} \quad (2.1)$$

where $f \in L^2(\Omega)$ represents the source term and $\mathbf{n}$ denotes the outward normal. Before presenting the ultraweak variational formulation, we briefly introduce the energy spaces used in this article. We define the standard energy spaces as:

$$L^2(\Omega) = \{u : \Omega \rightarrow \mathbb{R} : \|u\| < \infty\},$$

$$H^1(\Omega) = \left\{v : \Omega \rightarrow \mathbb{R} : v \in L^2(\Omega), \nabla v \in (L^2(\Omega))^3\right\}, \quad (2.2)$$

$$\mathbf{H}(\text{div}, \Omega) = \left\{\mathbf{w} : \Omega \rightarrow \mathbb{R}^3 : \mathbf{w} \in (L^2(\Omega))^3, \nabla \cdot \mathbf{w} \in L^2(\Omega)\right\}.$$

In the DPG method, discontinuous energy spaces are used for the test functions. Thus, we must define broken equivalents of $H^1(\Omega)$ and $\mathbf{H}(\text{div}, \Omega)$ spaces for the finite element mesh (Ω_h):

$$H^1(\Omega_h) := \left\{v : \Omega \rightarrow \mathbb{R} : v|_K \in H^1(K) \quad \forall K \in \Omega_h\right\}, \quad (2.3)$$

$$\mathbf{H}(\text{div}, \Omega_h) := \left\{\mathbf{w} : \Omega \rightarrow \mathbb{R}^3 : \mathbf{w}|_K \in \mathbf{H}(\text{div}, K) \quad \forall K \in \Omega_h\right\},$$

where $K \in \Omega_h$ represents an element of the finite element mesh. Use of the broken test spaces [24] leads to the introduction of additional trace unknowns on the mesh skeleton. The traces spaces are defined as:

$$\begin{aligned} H^{1/2}(\Gamma_h) &:= \left\{\hat{u} : \exists u \in H^1(\Omega) \text{ such that } \hat{u} = \gamma^K(u)|_K \text{ on } \partial K \quad \forall K \in \Omega_h\right\}, \\ H^{-1/2}(\Gamma_h) &:= \left\{\hat{\sigma}_n : \exists \sigma \in \mathbf{H}(\text{div}, \Omega) \text{ such that } \hat{\sigma}_n = \gamma_n^K(\sigma)|_K \text{ on } \partial K \quad \forall K \in \Omega_h\right\}, \end{aligned} \quad (2.4)$$

where γ^K and γ_n^K represent continuous and normal trace operators, respectively [25].

Ultraweak formulation. Let $(U, \hat{U})$ be the approximation trial space, V the test space, and V' the dual space of V . Then, the ultraweak DPG formulation of the Poisson problem can be stated as: Given $l \in V'$, find $u \in U$ and $\hat{u} \in \hat{U}$ satisfying:

$$b(u, \mathbf{v}) + \hat{b}(\hat{u}, \mathbf{v}) = l(\mathbf{v}) \quad \forall \mathbf{v} \in V, \quad (2.5)$$

where

$$\begin{aligned} u &= (u, \sigma) \in U := L^2(\Omega) \times (L^2(\Omega))^3, \\ \hat{u} &= (\hat{u}, \hat{\sigma}_n) \in \hat{U} := H^{1/2}(\Gamma_h) \times H^{-1/2}(\Gamma_h), \\ \hat{u} &= u_0 \text{ on } \Gamma_u, \hat{\sigma}_n = \sigma_0 \text{ on } \Gamma_\sigma, \\ \mathbf{v} &= (v, \tau) \in V := H^1(\Omega_h) \times \mathbf{H}(\text{div}, \Omega_h), \end{aligned} \quad (2.6)$$

$$b(u, \mathbf{v}) = (\sigma, \nabla v)_{\Omega_h} + (\sigma, \tau)_{\Omega_h} + (u, \nabla \cdot \tau)_{\Omega_h},$$

$$\hat{b}(\hat{u}, \mathbf{v}) = -\langle \hat{u}, \tau \cdot \mathbf{n} \rangle_{\Gamma_h} - \langle \hat{\sigma}_n, v \rangle_{\Gamma_h},$$

$$l(\mathbf{v}) = (f, v)_{\Omega_h}.$$

In (2.6), $\langle \cdot, \cdot \rangle_{\Gamma_h}$ represents duality pairings defined over mesh skeleton Γ_h ,

² Hence the “D” in the DPG method.

³ We then refer to it as the *practical* DPG method.

$$\begin{aligned}\langle \hat{\mathbf{u}}, \boldsymbol{\tau} \cdot \mathbf{n} \rangle_{\Gamma_h} &:= \sum_{K \in \Omega_h} \langle \hat{\mathbf{u}}, \boldsymbol{\tau} \cdot \mathbf{n}_K \rangle_{\partial K}, \\ \langle \hat{\boldsymbol{\sigma}}_h, \mathbf{v} \rangle_{\Gamma_h} &:= \sum_{K \in \Omega_h} \langle \hat{\boldsymbol{\sigma}}_h, \mathbf{v} \rangle_{\partial K},\end{aligned}\quad (2.7)$$

and

$$(\cdot, \cdot)_{\Omega_h} = \sum_{K \in \Omega_h} (\cdot, \cdot)_{L^2(K)}. \quad (2.8)$$

The broken test space is equipped with the adjoint graph norm [26,27]:

$$\|\mathbf{v}\|_V^2 := \|A_h^* \mathbf{v}\|^2 + \alpha \|\mathbf{v}\|^2 \quad (2.9)$$

where $\alpha > 0$ is a scaling constant, and $A_h^* \mathbf{v} = (\nabla \cdot \boldsymbol{\tau}, \nabla v + \boldsymbol{\tau})_{\Omega_h}$ is the (formal) adjoint operator of $A_h \mathbf{u} = (\boldsymbol{\sigma} - \nabla u, -\nabla \cdot \boldsymbol{\sigma})_{\Omega_h}$ computed element-wise. In this paper, all numerical experiments use $\alpha = 1$.

Another point of view—useful for implementation of the DPG method (see below)—is to reinterpret problem (2.5) with DPG as a mixed method [28,29,22]. Let $V_h \subset V$ denote the *enriched* finite-dimensional test space⁴ and let

$$R_{V,K} : V(K) \rightarrow (V(K))' \quad (2.10)$$

denote the element-wise Riesz operator corresponding to the test inner product. By identifying the residual defined in the (broken) test space (using the fact that the element-wise Riesz map $R_{V,K}$ is an isometry) $\Psi_K := R_{V,K}^{-1}(l_K(\cdot) - b_K(\mathbf{u}_h, \cdot) - \hat{b}_K(\hat{\mathbf{u}}_h, \cdot))$ as a new unknown, we can solve a mixed problem of the form

$$\begin{cases} \psi_h \in V_h, \mathbf{u}_h \in U_h, \hat{\mathbf{u}}_h \in U_h, \\ (\psi_h, \mathbf{v})_V - b(\mathbf{u}_h, \mathbf{v}) - \hat{b}(\hat{\mathbf{u}}_h, \mathbf{v}) = l(\mathbf{v}) & \mathbf{v} \in V_h \\ b(\delta \mathbf{u}_h, \psi_h) = 0 & \delta \mathbf{u}_h \in U_h \\ \hat{b}(\delta \hat{\mathbf{u}}_h, \psi_h) = 0 & \delta \hat{\mathbf{u}}_h \in \hat{U}_h \end{cases} \quad (2.11)$$

where trial spaces U and $\hat{U}$ are discretized with conforming finite-dimensional spaces $U_h \subset U$, $\hat{U}_h \subset \hat{U}$.⁵

Implementation of the DPG method. Next, we briefly discuss the algebraic structure of the resulting linear system and the built-in error estimator. The basis functions for V_h , U_h and $\hat{U}_h$ are denoted by φ_i , ψ_i and $\hat{\psi}_i$ respectively. From (2.6), we construct the following matrices for an element $K \in \Omega_h$,

$$\begin{aligned}G_{K,ij} &= (\varphi_i, \varphi_j)_V, \\ B_{K,ij} &= b_K(\varphi_i, \psi_j), \\ \hat{B}_{K,ij} &= \hat{b}_K(\varphi_i, \hat{\psi}_j), \\ l_{K,i} &= l_K(\varphi_i),\end{aligned}\quad (2.12)$$

where $G_{K,ij}$ represents the element Gram matrix corresponding to the test inner product and approximates the Riesz operator when discretized over $V_h(K)$. $B_{K,ij}$ represents the element stiffness matrix corresponding to the L^2 variables, $\hat{B}_{K,ij}$ represents the element stiffness matrix corresponding to the trace variables, and $l_{K,i}$ is the element load vector. As usual, $b_K(\cdot, \cdot)$, $\hat{b}_K(\cdot, \cdot)$ and $l_K(\cdot)$ denote element K contributions to bilinear forms $b(\mathbf{u}, \mathbf{v})$, $\hat{b}(\hat{\mathbf{u}}, \mathbf{v})$, and linear form $l(\mathbf{v})$, respectively. From (2.12), we obtain a symmetric positive definite linear system for each element K ,

$$\begin{bmatrix} G_K & B_K & \hat{B}_K \\ B_K^T & 0 & 0 \\ \hat{B}_K^T & 0 & 0 \end{bmatrix} \begin{bmatrix} \Psi_K \\ \mathbf{u}_K \\ \hat{\mathbf{u}}_K \end{bmatrix} = \begin{bmatrix} l_K \\ 0 \\ 0 \end{bmatrix}, \quad (2.13)$$

⁴ The enriched test space is discretized with broken (discontinuous) exact sequence spaces of order $p+1$.

⁵ In this article, we employ orientation-embedded exact sequence elements [30] of order p for discretization of the trial space.

where $[\mathbf{u}_K, \hat{\mathbf{u}}_K]^T$ and $[\Psi_K]$ are respectively the solution dofs and the residual dofs associated with element K . In a practical implementation, the mixed system (2.13) is reduced by statically condensing the residual dofs, giving an element-local DPG linear system of the form

$$[B_K \quad \hat{B}_K]^T G_K^{-1} [B_K \quad \hat{B}_K] [\mathbf{u}_K \quad \hat{\mathbf{u}}_K]^T = [B_K \quad \hat{B}_K]^T G_K^{-1} l_K. \quad (2.14)$$

A crucial point for discretization with the DPG method is the fact that the DPG formulations can be implemented within a standard Galerkin finite element code, as long as the code supports conforming discretization of the exact sequence energy spaces [30,31]. Indeed, in terms of implementation, the entire DPG method comes down to setting up the DPG linear system (2.14) for each element in the mesh; both the computation of optimal test functions as well as the DPG residual-based error estimator are local to each element and thus easily parallelizable. After the solution dofs have been obtained by solving the globally assembled system, the element-local DPG residual is easily computed for each element K :

$$\Psi_K = G_K^{-1}(l_K - B_K \mathbf{u}_K - \hat{B}_K \hat{\mathbf{u}}_K). \quad (2.15)$$

An in-depth exposition of the algebraic structure of the linear system induced by DPG formulation for a diffusion problem can be found in [19,32].

The built-in energy error estimate for a mesh element K in the finite element mesh (Ω_h) is given by:

$$\begin{aligned} \|(\mathbf{u}, \hat{\mathbf{u}}) - (\mathbf{u}_h, \hat{\mathbf{u}}_h)\|_{E,K}^2 \\ := \|R_{V,K}^{-1}(l_K(\cdot) - b_K(\mathbf{u}_h, \cdot) - \hat{b}_K(\hat{\mathbf{u}}_h, \cdot))\|_{V(K)}^2. \end{aligned} \quad (2.16)$$

One important aspect of the ultraweak DPG formulation is that the energy norm (2.16), computed with test norm (2.9), is equivalent with the L^2 -norm.⁶ With the element test space $V(K)$ approximated by a finite-dimensional enriched subspace $V_h(K)$, the element error indicators are then computed as⁷:

$$\eta_K := \|\Psi_K\|_{V(K)}^2 = \|G_K^{-1}(l_K - B_K \mathbf{u}_K - \hat{B}_K \hat{\mathbf{u}}_K)\|_{V(K)}^2. \quad (2.17)$$

Equivalence of norms. The ultraweak formulation using DPG methodology involves three groups of variables: the L^2 field variables ($\mathbf{u}$), the trace variables ($\hat{\mathbf{u}}$), and the Riesz representation of the residual (Ψ). Considering an idealized semidiscrete problem where Ψ remains undiscrctized leads to the ideal DPG method. This approach is justified if one discretizes the variables with higher-order elements (enriched spaces). This is indeed the case in the practical DPG method, where the error in Ψ is of higher order, as can be demonstrated by the construction of appropriate Fortin operators. Correspondingly, we can consider an idealized setting where both Ψ and $\hat{\mathbf{u}}$ remain undiscrctized. In this case, the DPG method can be viewed as a practical realization of a weakly conforming discretization of the test space. In one-dimensional problems, the trace variables are scalar values, automatically leading to the best approximation error of traces being zero. In the present work, for three-dimensional problems solved using the practical DPG method, we assume that if we discretize the Riesz representation and the trace variables with higher-order elements, the best approximation error of the traces and the Riesz representation can be neglected compared to the best approximation error of the field variables. In such a case, the energy error corresponding to the adjoint graph norm is indeed equivalent

⁶ The energy norm does not coincide with the L^2 -norm only because the adjoint test norm is not localizable, hence the test norm (2.9) is an adjoint graph norm with an L^2 component scaled with coefficient α ; however, with sufficiently small α , the energy norm is tightly equivalent with the L^2 -norm.

⁷ The energy norm computed on the finite-dimensional enriched test space is equivalent with the L^2 -norm; however, in addition to scaling coefficient α , the equivalence constants depend on the continuity constant of the problem-dependent Fortin operator [33].

to the L^2 error in the field. By restricting ourselves to the ideal DPG method, we can bound the L^2 norm of the field variable with its energy norm as follows:

$$\begin{aligned}
 \|u\| &= \sup_{v \in V(\Omega)} \frac{|(u, A^* v)|}{\|A^* v\|} \\
 &\leq \left(1 + \frac{\alpha}{\gamma^2}\right)^{\frac{1}{2}} \sup_{v \in V(\Omega)} \frac{|(u, A^* v)|}{(\|A^* v\|^2 + \alpha \|v\|^2)^{\frac{1}{2}}} \\
 &= \left(1 + \frac{\alpha}{\gamma^2}\right)^{\frac{1}{2}} \sup_{v \in V(\Omega)} \frac{|(u, A^* v) + \langle \hat{u}, v \rangle|}{(\|A^* v\|^2 + \alpha \|v\|^2)^{\frac{1}{2}}} \\
 &\leq \left(1 + \frac{\alpha}{\gamma^2}\right)^{\frac{1}{2}} \sup_{v \in V(\Omega_h)} \frac{|(u, A_h^* v) + \langle \hat{u}, v \rangle|}{(\|A_h^* v\|^2 + \alpha \|v\|^2)^{\frac{1}{2}}} = \left(1 + \frac{\alpha}{\gamma^2}\right)^{\frac{1}{2}} \|(u, \hat{u})\|_E,
 \end{aligned} \tag{2.18}$$

where γ is the global boundedness below constant, A^* is the continuous adjoint operator, and $V(\Omega)$ denotes the globally conforming (unbroken) test space. Conversely, element K 's contribution to the energy norm is bounded by the sum of the L^2 norm of the field variable u and the minimum energy extension norm of the trace variable $\hat{u}$:

$$\begin{aligned}
 \|(u, \hat{u})\|_{E,K} &= \sup_{v \in V(K)} \frac{|(u, A^* v) + \langle \hat{u}, v \rangle|}{(\|A^* v\|^2 + \alpha \|v\|^2)^{\frac{1}{2}}} \\
 &\leq \sup_{v \in V(K)} \frac{|(u, A^* v)|}{(\|A^* v\|^2 + \alpha \|v\|^2)^{\frac{1}{2}}} + \sup_{v \in V(K)} \frac{|\langle \hat{u}, v \rangle|}{(\|A^* v\|^2 + \alpha \|v\|^2)^{\frac{1}{2}}} \\
 &\leq \|u\| + \|\hat{u}\|_{\hat{U}}
 \end{aligned} \tag{2.19}$$

where $\|\cdot\|_{\hat{U}}$ denotes the minimum energy extension norm. In the present work, we optimize the mesh by minimizing *only* the L^2 error in the fields, assuming that the optimal mesh will also produce an adequate quality approximation of the trace variables. A mesh optimization strategy that also accounts for the trace variables would have significantly higher computational cost and be problem-dependent.

3. Determining optimal hp refinements

The hp -algorithm described in this section is *exactly* the algorithm from [2,3], but specialized to the L^2 -energy space. The corresponding algorithms for the H^1 , $H(\text{curl})$, and $H(\text{div})$ energy spaces, all based on minimizing the *Projection-Based (PB) interpolation error*, are significantly more intricate and consist of several steps reflecting the nature of the particular energy space. For instance, the algorithms for H^1 and $H(\text{curl})$ spaces consist of three stages involving mesh optimization on (interiors of) edges, faces and, finally, elements. The optimal mesh determined in each step serves as a starting point for the optimization in the subsequent step.

In the case of the L^2 -energy space, there are no global conformity requirements; the PB interpolation reduces to just the L^2 -projection, and the mesh optimization takes place over elements only. The implementation of the algorithm is thus much simpler. The second difference between the presented and the original hp -algorithm lies in the involved elements. In the original algorithm, the optimization takes place over *all elements*, whereas here it only does for elements marked for refinement by the DPG residual. The number of elements entering the mesh optimization is thus much smaller.⁸ The *fine mesh* providing the *reference solution* for the mesh optimization is also much smaller than the globally hp -refined mesh used in [2,3]. Fig. 1 illustrates a two-dimensional




Fig. 1. Isotropic hp -refinement of the marked elements: the elements marked for refinement are shaded in red on the coarse mesh.

case of mesh elements being marked by the DPG residual, followed by their isotropic hp -refinement⁹ to generate the *fine mesh*.

The hp -algorithm consists of three steps: the first and third step are purely local (can be done trivially in parallel) while the second step requires a global reduction over the elements preselected for refinement by the DPG residual.

3.1. Step 1: Staging a Competition of Refinements

In the first step of the algorithm, we stage a competition between p - and various anisotropic h -refinements, by computing the so-called guaranteed error reduction rate. The comparison between the various candidate refinements is based on the *error reduction rate* (e_{hp}) defined as:

$$e_{hp} := \frac{\|u - P_{\text{old}} u\|^2 - \|u - P_{\text{new}} u\|^2}{N_{\text{new}} - N_{\text{old}}}, \tag{3.20}$$

where u represents the *reference solution* obtained with the hp -refined mesh generated using the DPG residual, P_{old} is the L^2 -projection onto the original coarse mesh element (space), P_{new} is the L^2 -projection onto a *refined element* (space), N_{new} and N_{old} are the dimensions of the new and old spaces (number of dofs), respectively, and $\|\cdot\|$ denotes the L^2 -norm over the considered element K .

The optimal element refinement is determined by staging a competition among various candidate refinements. For hexahedral elements considered in this paper, there are eight possibilities: no h -refinement (i.e. p -refinement only), three anisotropic h_2 -refinements, three anisotropic h_4 -refinements, and the isotropic h_8 -refinement. Fig. 2 illustrates all possible h -refinement candidates. Each of the eight refinements is accompanied with the determination of the optimal distribution of polynomial degrees. This leads to a catastrophically large number of possible cases for hp -refinement. With $p_x, p_y, p_z \in \{1, \dots, 10\}$, there are “only” 10^3 scenarios for the just p -refined element, but a staggering total of 10^{24} cases for the h_8 -refined element. Clearly, a simple search through all possible cases is not feasible. Instead we rely on the classical p -refinement strategy, see e.g. [35], based on increasing the polynomial order in the subelement with the maximum error. This reduces the discrete search to the so-called *maximum error reduction path* through the vast discrete space of potentially possible refinements.

Maximum error reduction path for a p -refined element. We begin the discussion with the simplest case: p -refinement only. Assuming that the polynomial order can only increase (by one order), there are only a total of $2^3 - 1 = 7$ possible scenarios. The direct search is then possible but can be replaced with a slightly faster dynamic search, as illustrated in Fig. 3. To choose the optimal p -refinement, we traverse from

⁸ Dependent upon the parameter in the Dörfler strategy [34].

⁹ For a three-dimensional hexahedral element, isotropic hp -refinement denotes an isotropic h_8 -refinement followed by an isotropic p -refinement of order 1.




Fig. 2. Various possible h -refinements for a hexahedral element, depicted in (a): (b–d) anisotropic h_2 -refinements; (e–g) anisotropic h_4 -refinements; and (h) isotropic h_8 -refinement.




Fig. 3. Maximum error reduction path for the p -refined element: traversing from (p_x, p_y, p_z) to $(p_x + 1, p_y + 1, p_z + 1)$ for a hexahedral element.

(p_x, p_y, p_z) to $(p_x + 1, p_y + 1, p_z + 1)$ by increasing the order in directions that maximize e_{hp} . For a hexahedral element, the path of traversal has two stages. The first stage has three branches corresponding to p_x, p_y , and p_z . The second stage has two branches corresponding to the remaining directions, with the final configuration being $(p_x + 1, p_y + 1, p_z + 1)$. In Fig. 3, the arrows in red represent the branches corresponding to the highest values of e_{hp} at each stage, and the polynomial order marked in red indicates the polynomial order increased after each stage.

Following the path, we select the p -refinement that delivers the largest error reduction rate. In the case of an affine element, the element Jacobian (jac) is constant, and the L^2 -Piola transform (pullback map) reduces to a scaling with the Jacobian:

$$\phi_j(x) = \frac{1}{\text{jac}} \hat{\phi}_j(\xi), \quad \text{jac} = \left| \frac{\partial x_i}{\partial \xi_j} \right|, \quad (3.21)$$

where ϕ_j is an element L^2 shape function corresponding to a master element shape function $\hat{\phi}_j$. Consequently, the L^2 mass matrix,

$$M_{ij} := \int_K \phi_i \phi_j dx = \frac{1}{\text{jac}} \int_{\hat{K}} \hat{\phi}_i \hat{\phi}_j d\xi, \quad (3.22)$$

is diagonal, and the evaluation of the L^2 projection of a function u onto a subspace spanned by functions $\phi_1, \dots, \phi_N$, reduces to the evaluation of the load vector:

$$P_N u = \sum_{j=1}^N u_j \phi_j, \quad u_j = \frac{1}{M_{jj}} \int_K u \phi_j dx. \quad (3.23)$$

Raising the polynomial order in one direction amounts to adding extra orthogonal shape functions ϕ_{N+l} with $l = 1, \dots, n$. Consequently, evaluation of the error reduction rate reduces to:

$$\frac{\|u - P_N u\|^2 - \|u - P_{N+1} u\|^2}{n} = \frac{\|P_{N+1} u\|^2 - \|P_N u\|^2}{n} \quad (3.24)$$

$$= \frac{1}{n} \sum_{l=1}^n |u_{N+l}|^2 M_{N+l, N+l} = \frac{1}{n} \sum_{l=1}^n \left(\frac{\int_K u \phi_{N+l} dx}{M_{N+l, N+l}} \right)^2 M_{N+l, N+l}$$




Fig. 4. Staging a competition between the p -refined and h -refined element. The maximum error reduction path for the h -refined element traverses from $(p_{K_1}, p_{K_2}) = (1, 1)$ to the winning refinement $(p_{K_1}, p_{K_2}) = (4, 1)$.

$$= \frac{1}{n} \sum_{l=1}^n M_{N+l, N+l}^{-1} \left(\int_K u \phi_{N+l} dx \right)^2.$$

In the case of a general curvilinear element, the L^2 mass matrix is not diagonal, and we use the telescopic solver based on the Cholesky decomposition described in [3, p. 140].

Maximum error reduction path for an h -refined element. Contrary to the pure p -refinement, we always start with a trilinear element where $p_x = p_y = p_z = 1$. The reference solution u is projected onto the subelement mesh and, based on the distribution of the error, subelements are selected for refinement using a *greedy strategy* with a 70% factor. Once the subelements have been identified for p -refinement, the routine described above is employed to determine the optimal p -refinement for each subelement.

Fig. 4 shows the simple case of a 1D element K , starting with polynomial order $p_K = 4$. The subelements of the h -refined element K are denoted K_1 and K_2 , and their respective polynomial orders p_{K_1} and p_{K_2} . The maximum error reduction path for this case (illustrated in Fig. 4c) leads to the winning refinement $(p_{K_1}, p_{K_2}) = (4, 1)$ with the approximate solution shown in Fig. 4b.

The optimal refinement. The selection of the optimal refinement is carried out by comparing the best error reduction rates delivered by the eight differently h -refined meshes. The highest error reduction rate, delivered by the optimal refinement, is called the *guaranteed error reduction rate* and denoted by e_{hp}^* .

3.2. Step 2: Determining Which Elements to Refine

We loop over all considered coarse mesh elements to determine the element with the *best guaranteed error reduction rate* $e_{hp, \max}^*$. In principle, one could then refine only this one element. However, to accelerate the refinements (i.e. reduce the number of refinement steps), a greedy strategy is employed selecting all elements that deliver a rate greater than or equal to 25% of the best guaranteed error reduction rate. Note that this strategy implies that there may be elements initially marked for refinement by the DPG residual which ultimately remain unrefined.

3.3. Step 3: Determining the Final Refinements

For each element selected for refinement in Step 2, we could simply execute the corresponding optimal refinement determined in Step 1; and we do this indeed for the purely p -refined elements. However, when performing h -refinements we typically choose to invest additional dofs by considering the already-performed p -refinements that followed the optimal refinement while investigating error reduction rates in Step 1.

In particular, in Step 1 we recorded the error reduction rates for all subelement meshes following the maximum error reduction path.

On this path, we select the maximum investment (in terms of new dofs) that still delivers 25% of the best guaranteed error reduction rate (meaning it would still satisfy the Step 2 criterion). The rationale for doing so is to reduce the overall number of outer-loop iterations (number of refinement steps) by maximizing the investment in each step as long as it pays off sufficiently (delivering a sufficiently high error reduction rate, as determined by Step 2).

For example, in the 1D case illustrated in Fig. 4, the refinement shown in Fig. 4b won the competition with the p -refinement (Fig. 4a) but, dependent upon the threshold value used in the greedy strategy, we may choose to invest additional dofs in one of the subelements.

Next, we consolidate Steps 1–3 and present the mesh optimization algorithm. In Algorithm 1, tol denotes the user-provided tolerance value for the DPG residual.

Algorithm 1 Mesh Optimization Algorithm.

- 1: Start with an initial trial mesh
- 2: **while** $\eta_{\Omega_h} > \text{tol}$ **do**
- 3: Solve the problem on the current mesh.
- 4: Compute the DPG residual for the current mesh: $\eta_{\Omega_h} = \left(\sum_{K \in \Omega_h} \eta_K \right)^{1/2}$.
- 5: Use the element residuals (η_K) to mark elements for refinements (Dörfler strategy).
- 6: Isotropically hp -refine marked elements to generate the fine mesh.
- 7: Compute the reference solution u using the fine mesh.
- 8: **Step 1:** For each refined element K :
- 9: Determine the best possible p -refinement using the maximum error reduction path.
- 10: Determine the best possible h -refinement using the maximum error reduction path.
- 11: Use error reduction rates to decide between p - and h -refinement.
- 12: Determine the element guaranteed error reduction rate ($e_{hp, K}^*$).
- 13: **Step 2:** Determine the best guaranteed error reduction rate ($e_{hp, \max}^*$).
- 14: Unrefine the mesh.
- 15: **Step 3:** For each element K marked for refinement:
- 16: **if** $e_{hp, K}^* \geq 0.25 e_{hp, \max}^*$ **then**
- 17: Perform the optimal hp -refinement.
- 18: **end if**
- 19: **end while**

3.4. Mesh closure

The hp algorithm is implemented in $hp3D$, a general-purpose finite element code supporting hybrid meshes consisting of elements of all shapes (hexas, tets, prisms, pyramids), conforming discretizations of the exact-sequence spaces (H^1 -, $H(\text{curl})$ -, $H(\text{div})$ -, and L^2 -conforming elements), solution of coupled multiphysics problems, and *anisotropic* hp -refinements [36,37]. $hp3D$ supports MPI/OpenMP parallelism [38,39]

and is available under BSD-3 license.¹⁰ In the code, any h -refinement is executed in two steps. Given a list of elements to refine (along with the requested, possibly anisotropic, h -refinement flags), we proceed as follows.

Closure step 1 (local): Refine the elements from the list in the provided order, enforcing two rules:

- **Compatibility with existing face refinements:** upgrade the requested element refinement flag to accommodate *existing face refinements*.
- **One-irregularity rule for faces:** employ the standard *shelf* or *queue* algorithm ([19, p. 71]) to ensure that no face is refined unless the face¹¹ is unconstrained.

If one of the element faces is constrained, the element is placed on the shelf, and a necessary refinement of the neighbor across the face is executed, to eliminate the constraint. If the one-irregularity rule for faces prohibits the refinement, the corresponding neighbor is placed on the shelf and so on. Once the refinement of the processed element is possible, it is executed and the process resumes with the last element from the shelf. The algorithm proceeds until the shelf is empty. All mesh manipulations (refinements) are supported for meshes that satisfy the one-irregularity rule for faces (not necessary for edges and vertices).

Closure step 2 (global): Loop through all elements and perform additional necessary refinements to eliminate edges and vertices with multiple constraints.

We refer to [39] for a more formal exposition of the algorithms. In the end, in both steps, a number of additional, *unwanted* refinements may be executed. These refinements can be *isotropic* or *anisotropic*, reflecting minimal requirements to eliminate the nodes with multiple constraints. In the ‘global’ hp -refinement driven by the DPG residual, all unwanted refinements are chosen to be isotropic. This is motivated by the fact that an unwillingly refined element (in Step 1) may, in fact, be on the DPG list of wanted refinements. However, once the optimal hp -refinements are determined, all unwanted refinements are executed in a minimal, anisotropic way.

All unwillingly h -refined elements retain their respective polynomial order. In principle, one could attempt to find the corresponding optimal distribution of polynomial orders, but this has been not done in our current implementation. Hence, the presented meshes may be slightly overrefined.

4. Numerical results

4.1. A boundary layer problem

Sharp boundary layers are among the most commonly encountered flow features in computational fluid dynamics. Our first numerical experiment demonstrates the proposed algorithm’s efficacy in resolving such boundary layers. In this test case, we solve a Poisson problem with a manufactured solution containing boundary layers. The manufactured solution is a three-dimensional extension of the solution of the Egger-Schöberl problem [40]. In particular, we solve

$$\begin{aligned} -\nabla^2 u &= f(x, y, z) \quad \text{in } \Omega := (0, 1)^3, \\ u &= 0 \quad \text{on } \Gamma_u, \end{aligned} \quad (4.25)$$

$$\nabla u \cdot \mathbf{n} = g(x, y, z) \quad \text{on } \Gamma_\sigma,$$

where

$$\begin{aligned} \Gamma_u &= ([0, 1] \times [0, 1] \times \{0\}) \cup ([0, 1] \times \{0\} \times [0, 1]) \cup (\{0\} \times [0, 1] \times [0, 1]), \\ \Gamma_\sigma &= ([0, 1] \times [0, 1] \times \{1\}) \cup ([0, 1] \times \{1\} \times [0, 1]) \cup (\{1\} \times [0, 1] \times [0, 1]). \end{aligned} \quad (4.26)$$

In (4.25), $\mathbf{n}$ is the outward normal, and f and g are generated using the exact solution. The exact solution is given by

$$u(x, y, z) = \left(x + \frac{e^{x/\epsilon} - 1}{1 - e^{1/\epsilon}}\right) \left(y + \frac{e^{y/\epsilon} - 1}{1 - e^{1/\epsilon}}\right) \left(z + \frac{e^{z/\epsilon} - 1}{1 - e^{1/\epsilon}}\right). \quad (4.27)$$

The solution exhibits a boundary layer near $x \approx 1$, $y \approx 1$ and $z \approx 1$. The strength of the boundary layer is inversely proportional to ϵ . In this numerical experiment, $\epsilon = 0.005$. The hp -adaptation is initialized with a mesh comprising only eight elements with a constant polynomial order of $(2, 2, 2)$.¹²

Figs. 5a and 5b display the cross-section of an adapted mesh and the corresponding solution contour, respectively. Fig. 6 depicts the polynomial distribution around the boundary layers on an anisotropically adapted hp -mesh. Fig. 7 presents the convergence results, comparing isotropic h -adaptation and the proposed hp -refinement strategy. The Dörfler parameter for both isotropic and hp -refinement is 0.75. In Fig. 7, the depicted error is the combined relative error in all L^2 variables.

Fig. 6 clearly illustrates the strong anisotropy and grading in the element size and the polynomial distribution. The anisotropy and the grading in element size are paramount for resolving strong boundary layers efficiently. The algorithm also prescribes an anisotropic polynomial distribution in the boundary layers instead of an isotropic one. This directional preference of prescribing polynomial orders showcases a significant advantage of the proposed hp -refinement strategy: the ability to complement an anisotropic h -refinement with an anisotropic p -refinement. This approach makes the refinement strategy highly efficient in terms of allocating dofs when the solution exhibits strong anisotropic features. The algorithm does not waste any dofs in directions where the solution variables do not exhibit significant variations.

From Fig. 7, it is evident that anisotropic hp -refinements outperform isotropic h -refinements by orders of magnitude. The convergence plots show the error and the residual against $\sqrt[3]{\text{ndof}}$ (where ndof represents the number of dofs), verifying exponential convergence. In Fig. 7, a reduction in the convergence rate for the hp -refinement can be observed. The slowdown in convergence occurs due to the limiting of the highest polynomial order in the numerical experiments to $p = 6$. The adaptation cycles are initially dominated by h -refinements. Once the boundary layers are resolved, the algorithm starts preferring both p -refinements along with h -refinements. This behavior is expected, since, increasing the polynomial order on coarse meshes while approximating solutions with high gradients can induce spurious oscillations. In Fig. 8, we plot the effectivity index, defined as the ratio of the DPG residual-based built-in error estimator and the composite L^2 error of the field variables u and σ .

4.2. Fichera cube problem

To demonstrate the efficacy of the proposed refinement strategy in the presence of multiple singularities, we solve the well-known Fichera cube problem and perform hp -adaptations using the proposed refinement strategy. The variant of the Fichera cube problem being solved here is given by:

$$\begin{aligned} \nabla^2 u &= 0 \quad \text{in } \Omega := (-1, 1)^3 \setminus [0, 1]^3, \\ u &= 0 \quad \text{on } \Gamma_u, \\ \nabla u \cdot \mathbf{n} &= g(x, y, z) \quad \text{on } \Gamma_\sigma. \end{aligned} \quad (4.28)$$

¹⁰ <https://github.com/Oden-EAG/hp3d>.

¹¹ More precisely, the mid-face node.

¹² In $hp3D$, we employ exact-sequence spaces [31]. Hence, an order of (p_x, p_y, p_z) denotes L^2 shape functions of order $(p_x - 1, p_y - 1, p_z - 1)$.




Fig. 5. Boundary layer problem: (a) cross-section of the mesh showing anisotropic elements required to resolve the boundary layers; and (b) contour plot illustrating the boundary layers on the yz -plane. The boundary layers are along right and top faces of the cross-section.




Fig. 6. Boundary layer problem: an adapted mesh with 855 532 dofs; coloring indicates the polynomial distributions p_x , p_y , p_z in x -, y -, z -direction, respectively. The algorithm prescribes higher-order polynomials anisotropically corresponding to each boundary layer along the x -, y -, and z -axis.




Fig. 7. Boundary layer problem: convergence of relative L^2 error and DPG residual. Even though there is a marginal decrease in the rate of convergence for the hp -refinements (due to limiting $p \leq 6$ in the numerical experiments), both the error and the residual are 2–3 orders of magnitude lower compared to the h -refinements for approximately the same number of dofs.

The domain is created by subdividing a large cube $(-1, 1)^3$ into eight smaller cubes and then removing one of the cubes. The Dirichlet data $u = 0$ is imposed on the three square faces aligned with planes of coordinate axes, i.e.

$$\Gamma_u = ([0, 1] \times [0, 1] \times \{0\}) \cup ([0, 1] \times \{0\} \times [0, 1]) \cup (\{0\} \times [0, 1] \times [0, 1]). \quad (4.29)$$

The volumetric load for the problem is zero. The problem is driven by the Neumann boundary condition on Γ_σ composed of the remaining faces of the cube. The data g correspond to the sum of two-dimensional exact solutions of the L-shaped domain problem on xy -, yz -, and xz -planes. The exact solution of the L-shaped domain problem is given by:

$$u_{\eta, \xi} = r^{\frac{2}{3}} \cos(\theta), \quad r = \sqrt{\eta^2 + \xi^2}, \quad \theta = \tan^{-1} \left(\frac{\xi}{\eta} \right), \quad (4.30)$$

where (η, ξ) denote (x, y) , (y, z) , or (x, z) axes, respectively. These boundary conditions generate a solution with features analogous to an L-shaped domain problem but comprising multiple edge and vertex singularities. While the exact solution for the problem is unknown, the convergence of the DPG residual is shown in Fig. 13.

Figs. 9 and 10 depict the solution contour and the corresponding adapted mesh, respectively. Figs. 11 and 12 illustrate the polynomial distribution associated with the adapted mesh. Fig. 10 shows that the refinement algorithm performs highly anisotropic h -refinements along the edge singularities, generating graded meshes. The anisotropic refinements propagate through the volume to the opposing boundary faces on Γ_σ . The propagation of refinements happens in conjunction to the




Fig. 8. Boundary layer problem: effectivity index, defined as the ratio of the DPG residual-based error estimator and the composite L^2 error of field variables u and σ .




Fig. 9. Fichera cube problem: solution contour. The problem is driven by the Neumann boundary conditions on the L-shaped faces in (a) and the three visible square faces in (b). The faces aligned along the coordinate planes in (a) have the Dirichlet boundary conditions.




Fig. 10. Fichera cube problem: an anisotropically adapted hp -mesh with 1.3M dofs.




Fig. 11. Fichera cube problem: polynomial distribution on the adapted hp -mesh. The algorithm prescribes low-order polynomials anisotropically around each edge singularity along x -, y - and z -axis. Fig. 12 presents a magnified view of the polynomial distribution and anisotropic mesh elements around the singularities.




Fig. 12. Fichera cube problem: magnified view of the mesh and the polynomial distribution near the edge and vertex singularities.

singularities arising from the faces with Neumann boundary conditions. Figs. 11 and 12 clearly indicate that the algorithm chooses lowest order polynomials around the singularities. Moving away from the singularities, the algorithm prescribes higher order polynomials underscoring the smoothness of the solution variables. In Fig. 13, one can observe the exponential convergence of the residual on performing hp -refinements, whereas isotropic h -refinements suffer from a loss of convergence due to the lack of required grading in size and polynomial distribution.

4.3. Eriksson–Johnson problem

We consider a convection-dominated diffusion problem motivated by the Eriksson–Johnson model problem [41]. Here, we extend the exact solution of the two-dimensional problem by multiplying it with a sinusoidal term along z . In particular, we solve

$$\begin{aligned} \frac{\partial u}{\partial x} - \epsilon \nabla^2 u &= f(x, y, z) & \text{in } \Omega := (0, 1)^3, \\ u &= 0 & \text{on } \Gamma_{u_a}, \\ u &= \sin(\pi y) \sin(\pi z) & \text{on } \Gamma_{u_b}, \end{aligned} \quad (4.31)$$

where

$$\Gamma_{u_a} = \partial\Omega \setminus \{0\} \times [0, 1] \times [0, 1] \quad \text{and} \quad \Gamma_{u_b} = \{0\} \times [0, 1] \times [0, 1]. \quad (4.32)$$

The source f and the boundary conditions are computed using the exact solution given by

$$u(x, y, z) = \frac{e^{s_1(x-1)} - e^{s_2(x-1)}}{e^{s_1} - e^{s_2}} \sin(\pi y) \sin(\pi z), \quad (4.33)$$

where

$$s_1 = \frac{1 + \sqrt{1 + 4\pi^2 \epsilon^2}}{2\epsilon} \quad \text{and} \quad s_2 = \frac{1 - \sqrt{1 + 4\pi^2 \epsilon^2}}{2\epsilon}. \quad (4.34)$$




Fig. 13. Fichera cube problem: convergence of the DPG residual.




Fig. 14. Eriksson–Johnson problem: an adapted mesh and solution contour.

In this numerical experiments, $\epsilon = 0.01$. Fig. 14 depicts the cross-section of an adapted mesh and the corresponding solution contour at $z = 0.5$. The solution exhibits a boundary layer along the x -axis with sinusoidal variations along y and z . The variation in the solution is also reflected in the hp -refinements executed by the algorithm. In order to capture the boundary layer, the algorithm generates anisotropic elements parallel to the yz -plane and assigns the highest polynomial order along the x -axis inside the boundary layer. Since the boundary layer is weighted with sinusoidal variations in y and z , the majority of the h -refined elements in the boundary layer are positioned near $y = 0.5$ and $z = 0.5$. Fig. 15 illustrates the adapted mesh with the polynomial distribution along the x -axis. Fig. 16 presents the convergence plots for the relative L^2 error and the residual, demonstrating the efficacy of the proposed hp -refinement strategy for this problem. Finally, Fig. 17 depicts the effectivity index defined as the ratio of the DPG residual-based built-in error estimator and the composite L^2 error of the field variables u and σ .

5. Conclusion

The anisotropic hp -refinement strategy presented in this article utilizes the built-in DPG error-estimator and L^2 projection-based error estimates for the ultraweak variational formulation. The efficacy of the proposed algorithm is demonstrated through numerical experiments containing boundary layers and singularities. The algorithm is able to generate a sequence of meshes that provide exponential convergence. Since the maximum polynomial order in our numerical experiments was capped at $p = 6$ for practical (computational) reasons, we observe a slight loss of optimal convergence rate. Nonetheless, the accuracy of the solutions on the anisotropically refined hp -meshes remains orders of

magnitude better than that on isotropically refined meshes for approximately the same number of dofs. The proposed hp -refinement strategy complements anisotropic h -refinements with anisotropic p -refinements, which allows the algorithm to avoid any superfluous investment (in terms of dofs). In its current form, the proposed adaptation framework is primarily applicable to problems where the location of boundary layers and singularities is known a-priori. It allows one to initiate the adaptation cycles with tensor-product meshes aligned with the layers. For problems where the location of an internal layer or singularity is unknown a-priori, a structured methodological framework is proposed in [42]. In this article, we have avoided the problem of accurately approximating the optimal test functions by performing numerical experiments only with modestly small diffusion constant $\epsilon = 10^{-2}$ and $\epsilon = 10^{-3}$. For alternative strategies to handle smaller values of ϵ , see [26,27,43,44].

Future work. To accelerate the computation of the fine-grid solution and apply the hp -refinement strategy to large-scale multiphysics problems, we intend to integrate the proposed hp -refinement strategy with a scalable DPG multigrid solver [45] used for modeling nonlinear optical waveguides [46,47] and other high-frequency wave applications [48]. Additionally, we aim to extend the proposed refinement strategy to other element types, such as tets, prisms, and pyramids, in order to leverage $hp3D$'s capability to handle hybrid meshes.

Data availability

No data was used for the research described in the article.




Fig. 15. Eriksson–Johnson problem: an adapted mesh with 209 737 dofs; coloring indicates the corresponding polynomial distribution along the x -axis.




Fig. 16. Eriksson–Johnson problem: convergence of relative L^2 error and DPG residual.




Fig. 17. Eriksson–Johnson problem: effectivity index, defined as the ratio of the DPG residual-based error estimator and the composite L^2 error of field variables u and σ .

References

- [1] D. Stanzione, J. West, R.T. Evans, T. Minyard, O. Ghattas, D.K. Panda, Frontera: The Evolution of Leadership Computing at the National Science Foundation, in: Practice and Experience in Advanced Research Computing, Association for Computing Machinery, 2020, pp. 106–111.
- [2] L. Demkowicz, Computing with hp -Adaptive Finite Elements. Vol. I: One and Two Dimensional Elliptic and Maxwell Problems, Chapman and Hall/CRC, 2006.
- [3] L. Demkowicz, J. Kurtz, D. Pardo, M. Paszynski, W. Rachowicz, A. Zdunek, Computing with hp -Adaptive Finite Elements. Vol. II. Frontiers: Three Dimensional Elliptic and Maxwell Problems with Applications, Chapman and Hall/CRC, 2007.
- [4] I. Babuška, B. Guo, The h, p and hp -version of the finite element method; basis theory and applications, Adv. Eng. Softw. 15 (3) (1992) 159–174.
- [5] I. Babuška, W. Gui, The h, p and hp -versions of the finite element method in 1 dimension. Part III. The adaptive hp -version, Numer. Math. 49 (1986) 659–684.
- [6] W. Rachowicz, D. Pardo, L. Demkowicz, Fully automatic hp -adaptivity in three dimensions, Comput. Methods Appl. Mech. Eng. 195 (37) (2006) 4816–4842.
- [7] C. Schwab, p - and hp -Finite Element Methods: Theory and Applications in Solid and Fluid Mechanics, Clarendon Press, 1998.
- [8] I. Babuška, T. Strouboulis, K. Copps, hp optimization of finite element approximations: analysis of the optimal mesh sequences in one dimension, Comput. Methods Appl. Mech. Eng. 150 (1) (1997) 89–108.
- [9] J. Oden, L. Demkowicz, W. Rachowicz, T. Westermann, Toward a universal hp adaptive finite element strategy. Part 2: a posteriori error estimation, Comput. Methods Appl. Mech. Eng. 77 (1) (1989) 113–180.
- [10] L. Demkowicz, J. Oden, W. Rachowicz, O. Hardy, Toward a universal hp adaptive finite element strategy. Part 1: constrained approximation and data structure, Comput. Methods Appl. Mech. Eng. 77 (1) (1989) 79–112.
- [11] W. Rachowicz, J. Oden, L. Demkowicz, Toward a universal hp adaptive finite element strategy. Part 3: design of hp meshes, Comput. Methods Appl. Mech. Eng. 77 (1) (1989) 181–212.
- [12] J.T. Oden, W. Wu, V. Legat, An hp adaptive strategy for finite element approximations of the Navier-Stokes equations, Int. J. Numer. Methods Fluids 20 (8–9) (1995) 831–851.
- [13] P. Devloo, J. Tinsley Oden, P. Pattani, An h - p adaptive finite element method for the numerical simulation of compressible flow, Comput. Methods Appl. Mech. Eng. 70 (2) (1988) 203–235.
- [14] C. Duarte, J. Oden, An h - p adaptive method using clouds, Comput. Methods Appl. Mech. Eng. 139 (1) (1996) 237–262.
- [15] P. Houston, B. Senior, E. Süli, Sobolev regularity estimation for hp -adaptive finite element methods, in: F. Brezzi, A. Buffa, S. Corsaro, A. Murli (Eds.), Numerical Mathematics and Advanced Applications, Springer, Milan, 2003, pp. 631–656.
- [16] P. Houston, E. Süli, hp -adaptive discontinuous Galerkin finite element methods for first-order hyperbolic problems, SIAM J. Sci. Comput. 23 (4) (2001) 1226–1252.
- [17] A. Chakraborty, Optimal approximation spaces for discontinuous Petrov–Galerkin schemes with optimal test functions, Ph.D. thesis, Rheinisch-Westfälische Technische Hochschule Aachen, 2022.

- [18] L. Demkowicz, J. Gopalakrishnan, A class of discontinuous Petrov–Galerkin methods. Part I: The transport equation, *Comput. Methods Appl. Mech. Eng.* 199 (23) (2010) 1558–1572.
- [19] L. Demkowicz, J. Gopalakrishnan, A class of discontinuous Petrov–Galerkin methods. II. Optimal test functions, *Numer. Methods Partial Differ. Equ.* 27 (1) (2011) 70–105.
- [20] L. Demkowicz, J. Gopalakrishnan, A. Niemi, A class of discontinuous Petrov–Galerkin methods. Part III: Adaptivity, *Appl. Numer. Math.* 62 (4) (2012) 396–427.
- [21] J. Zitelli, I. Muga, L. Demkowicz, J. Gopalakrishnan, D. Pardo, V. Calo, A class of discontinuous Petrov–Galerkin methods. Part IV: the optimal test norm and time-harmonic wave propagation in 1D, *J. Comput. Phys.* 230 (7) (2011) 2406–2432.
- [22] L. Demkowicz, J. Gopalakrishnan, Discontinuous Petrov–Galerkin (DPG) method, in: *Encyclopedia of Computational Mechanics Second Edition*, 2017, pp. 1–15.
- [23] I. Babuška, Error-bounds for finite element method, *Numer. Math.* 16 (1971) 322–333.
- [24] C. Carstensen, L. Demkowicz, J. Gopalakrishnan, Breaking spaces and forms for the DPG method and applications including Maxwell equations, *Comput. Math. Appl.* 72 (3) (2016) 494–522.
- [25] L. Demkowicz, *Energy Spaces*, Lecture Notes, The University of Texas at Austin, 2018.
- [26] L. Demkowicz, N. Heuer, Robust DPG method for convection-dominated diffusion problems, *SIAM J. Numer. Anal.* 51 (5) (2013) 2514–2537.
- [27] J. Chan, N. Heuer, T. Bui-Thanh, L. Demkowicz, A robust DPG method for convection-dominated diffusion problems II: adjoint boundary conditions and mesh-dependent test norms, *Comput. Math. Appl.* 67 (4) (2014) 771–795.
- [28] A. Cohen, W. Dahmen, G. Welper, Adaptivity and variational stabilization for convection-diffusion equations, *ESAIM: Math. Model. Numer. Anal.* 46 (5) (2012) 1247–1273.
- [29] W. Dahmen, C. Huang, C. Schwab, G. Welper, Adaptive Petrov–Galerkin methods for first order transport equations, *SIAM J. Numer. Anal.* 50 (5) (2012) 2420–2445.
- [30] F. Fuentes, B. Keith, L. Demkowicz, S. Nagaraj, Orientation embedded high order shape functions for the exact sequence elements of all shapes, *Comput. Math. Appl.* 70 (4) (2015) 353–458.
- [31] L. Demkowicz, *Mathematical theory of finite elements*, *Soc. Ind. Appl. Math.* (2023).
- [32] A. Vaziri Astaneh, F. Fuentes, J. Mora, L. Demkowicz, High-order polygonal discontinuous Petrov–Galerkin (PolyDPG) methods using ultraweak formulations, *Comput. Methods Appl. Mech. Eng.* 332 (2018) 686–711.
- [33] J. Gopalakrishnan, W. Qiu, An analysis of the practical DPG method, *Math. Comput.* 83 (286) (2014) 537–552.
- [34] W. Dörfler, A convergent adaptive algorithm for Poisson’s equation, *SIAM J. Numer. Anal.* 33 (3) (1996) 1106–1124.
- [35] L. Demkowicz, J.T. Oden, T. Strouboulis, Adaptive finite elements for flow problems with moving boundaries. Part 1: Variational principles and a posteriori estimates, *Comput. Methods Appl. Mech. Eng.* 46 (1984) 217–251.
- [36] S. Henneking, L. Demkowicz, *hp3D user manual*, arXiv:2207.12211.
- [37] S. Henneking, S. Petrides, F. Fuentes, J. Badger, L. Demkowicz, *hp3D: a scalable MPI/OpenMP hp-adaptive finite element software library for complex multiphysics applications*, *J. Open Sour. Softw.* 9 (95) (2024) 5946.
- [38] S. Henneking, A scalable hp-adaptive finite element software with applications in fiber optics, Ph.D. thesis, The University of Texas at Austin, 2021.
- [39] S. Henneking, L. Demkowicz, S. Petrides, F. Fuentes, B. Keith, P. Gatto, *Computing with hp Finite Elements. III. Parallel hp3D Code*, 2024, in preparation.
- [40] H. Egger, J. Schöberl, A hybrid mixed discontinuous Galerkin finite element method for convection–diffusion problems, *IMA J. Numer. Anal.* 30 (4) (2009) 1206–1234.
- [41] K. Eriksson, C. Johnson, Adaptive streamline diffusion finite element methods for stationary convection–diffusion problems, *Math. Comput.* 60 (201) (1993) 167–188.
- [42] M. Harris, Flow feature aligned mesh generation and adaptation, Ph.D. thesis, University of Sheffield, 2013.
- [43] L. Demkowicz, T. Führer, N. Heuer, X. Tian, The double adaptivity paradigm: (how to circumvent the discrete inf–sup conditions of Brezzi), *Comput. Math. Appl.* 95 (2021) 41–66.
- [44] J. Salazar, J. Mora, L. Demkowicz, Alternative enriched test spaces in the DPG method for singular perturbation problems, *Comput. Methods Appl. Math.* 19 (3) (2019) 603–630.
- [45] J. Badger, S. Henneking, S. Petrides, L. Demkowicz, Scalable DPG multigrid solver for Helmholtz problems: a study on convergence, *Comput. Math. Appl.* 148 (2023) 81–92.
- [46] S. Henneking, L. Demkowicz, A numerical study of the pollution error and DPG adaptivity for long waveguide simulations, *Comput. Math. Appl.* 95 (2021) 85–100.
- [47] S. Henneking, J. Grosek, L. Demkowicz, Model and computational advancements to full vectorial Maxwell model for studying fiber amplifiers, *Comput. Math. Appl.* 85 (2021) 30–41.
- [48] S. Petrides, L. Demkowicz, An adaptive multigrid solver for DPG methods with applications in linear acoustics and electromagnetics, *Comput. Math. Appl.* 87 (2021) 12–26.