

Oden Institute REPORT 19-20

December 2019

A numerical study of the pollution error and DPG adaptivity for long waveguide simulations

by

Stefan Henneking and Leszek Demkowicz



Oden Institute for Computational Engineering and Sciences
The University of Texas at Austin
Austin, Texas 78712

Reference: Stefan Henneking and Leszek Demkowicz, "A numerical study of the pollution error and DPG adaptivity for long waveguide simulations," Oden Institute REPORT 19-20, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, December 2019.

A numerical study of the pollution error and DPG adaptivity for long waveguide simulations

Stefan Henneking and Leszek Demkowicz

Oden Institute for Computational Engineering and Sciences,
The University of Texas at Austin, Austin, TX 78712, USA,
E-mail: stefan@oden.utexas.edu, leszek@oden.utexas.edu

December 11, 2019

Abstract

High-frequency wave propagation has many important applications in acoustics, elastodynamics, and electromagnetics. Unfortunately, the finite element discretization for these problems suffer from significant numerical pollution errors that increase with the wavenumber. It is critical to control these errors to obtain a stable and accurate method. We study the effect of pollution for very long waveguide problems in the context of robust discontinuous Petrov-Galerkin (DPG) finite element discretizations. Our numerical experiments show that the pollution primarily has a diffusive effect causing energy loss in the DPG method while phase errors appear less significant. We report results for 3D vectorial time-harmonic Maxwell problems in waveguides with more than 8000 wavelengths. Our results corroborate previous analysis for the Galerkin discretization of the Helmholtz operator by Melenk and Sauter (“Wavenumber explicit convergence analysis for Galerkin discretizations of the Helmholtz equation”. In: *SIAM J. Numer. Anal.* 49.3 (2011), pp. 1210–1243). Additionally, we discuss adaptive refinement strategies for multi-mode fiber waveguides where the propagating transverse modes must be resolved sufficiently. Our study shows the applicability of the DPG error indicator to this class of problems. Finally, we illustrate the importance of load balancing in these simulations for distributed-memory parallel computing.

1 Introduction

Motivation. It is well-known that an accurate numerical solution in high-frequency wave problems is difficult to obtain. Acoustic or electromagnetic wave propagation problems with high frequency in the time-harmonic setting lead to partial differential equations (PDEs) with an indefinite Helmholtz or Maxwell operator, respectively.

Therefore, many advanced solver techniques for Hermitian positive definite discrete systems are not directly applicable to these problems. Additionally, finite element discretizations are typically dependent on satisfying the Nyquist stability criterion, implying that all propagating wave frequencies must be “resolved” to a certain extent in order to have a stable discretization. Even though advanced finite element methods such as the discontinuous Petrov-Galerkin (DPG) method can circumvent the stability problem and deliver a robust discretization for any wave number [10, 26, 27], they do not eliminate the numerical pollution error in multiple dimensions. Pollution can manifest itself in different forms: commonly, we observe a diffusive effect causing wave attenuation and/or a dispersive effect resulting in a phase shift. It is critical to control the pollution error for obtaining accurate results.

Background and literature. Numerical pollution in wave propagation has been studied extensively for the Bubnov-Galerkin finite element method, as well as discontinuous Galerkin methods, least-squares methods, and various other approaches (see [3, 19, 11, 1, 22, 10] and references therein); however, to the best of our knowledge numerical studies have mostly been limited to acoustic wave problems with a moderate number of wavelengths. In this paper, we are reporting numerical results for the DPG method applied to the 3D vectorial time-harmonic Maxwell problem in long waveguides with more than 8000 wavelengths and high order of approximation. We attempt to corroborate theoretical results by Melenk and Sauter [22] and provide guidance as to how to best discretize wave problems with high frequency.

Outline. First, we introduce the DPG methodology and we discuss how the pollution error enters the DPG formulation as a perturbation parameter in the best approximation, in contrast to the classical Bubnov-Galerkin method where the perturbation parameter is present in the stability constant. We briefly go over the mathematical setting for the time-harmonic Maxwell equations in a linear waveguide in the context of the broken ultraweak variational formulation. For the numerical pollution study, we are reporting the experiments for a simple case with one propagating mode in a rectangular waveguide with the intent of eliminating other effects unrelated to pollution. We discuss our findings in the light of theoretical pollution error estimates from the literature. For waveguides with multiple propagating modes, we discuss different adaptive refinement strategies; in particular, we show numerical results for multi-mode fiber simulations illustrating the need for mesh adaptivity in the numerical solution to such problems. In the case of a distributed-memory parallel simulation, we emphasize the importance of dynamic load balancing for this particular problem.

2 The DPG methodology

The DPG method of Demkowicz and Gopalakrishnan [7, 9, 8] has been used to solve applications in viscoelasticity [15], acoustic and electromagnetic wave propagation [27,

25, 18], compressible fluid dynamics [6], and linear elasticity [21]. Its stability properties make it particularly applicable to high-frequency wave propagation problems, where pre-asymptotic stability is essential for driving efficient *hp*-adaptivity [26]. Through on-the-fly computation of problem-dependent optimal test functions, the DPG method guarantees a stable discretization for any well-posed linear variational problem.

2.1 The ideal DPG method

Consider an abstract variational problem of the form,

$$\begin{cases} u \in \mathcal{U}, \\ b(u, v) = l(v), \quad v \in \mathcal{V}, \end{cases} \quad (2.1)$$

where \mathcal{U} (trial space) and \mathcal{V} (test space) are Hilbert spaces, and $b(\cdot, \cdot)$ is a continuous bilinear (sesquilinear) form on $\mathcal{U} \times \mathcal{V}$ (with continuity constant M),

$$|b(u, v)| \leq M \|u\|_{\mathcal{U}} \|v\|_{\mathcal{V}}, \quad (2.2)$$

that satisfies the continuous inf-sup condition (with inf-sup constant γ),

$$\inf_{\|u\|_{\mathcal{U}}=1} \sup_{\|v\|_{\mathcal{V}}=1} |b(u, v)| =: \gamma > 0. \quad (2.3)$$

And, the continuous linear (antilinear) form $l(\cdot)$ satisfies the compatibility condition,

$$l(v) = 0 \quad \forall v \in \mathcal{V}_0, \quad \text{where } \mathcal{V}_0 := \{v \in \mathcal{V} : b(u, v) = 0 \quad \forall u \in \mathcal{U}\}. \quad (2.4)$$

Let \mathcal{U}' and \mathcal{V}' denote the space of continuous linear (antilinear) functionals on \mathcal{U} and \mathcal{V} , respectively. By the Babuška-Nečas theorem, the variational problem (2.1) is well-posed, i.e., there exists a unique solution $u \in \mathcal{U}$ that depends continuously upon the data,

$$\|u\|_{\mathcal{U}} \leq \frac{1}{\gamma} \|l\|_{\mathcal{V}'}. \quad (2.5)$$

Consider finite-dimensional subspaces $\mathcal{U}_h \subset \mathcal{U}$ and $\mathcal{V}_h \subset \mathcal{V}$, where $\dim(\mathcal{U}_h) = \dim(\mathcal{V}_h)$, and the corresponding discrete abstract variational problem,

$$\begin{cases} u_h \in \mathcal{U}_h \subset \mathcal{U}, \\ b(u_h, v_h) = l(v_h), \quad v_h \in \mathcal{V}_h \subset \mathcal{V}. \end{cases} \quad (2.6)$$

If the discrete inf-sup condition is satisfied, i.e.,

$$\inf_{\|u_h\|_{\mathcal{U}}=1} \sup_{\|v_h\|_{\mathcal{V}}=1} |b(u_h, v_h)| =: \gamma_h > 0, \quad (2.7)$$

then the discrete problem (2.6) is well-posed; and by Babuška's theorem [2],

$$\|u - u_h\|_{\mathcal{U}} \leq \frac{M}{\gamma_h} \inf_{\omega_h \in \mathcal{U}_h} \|u - \omega_h\|_{\mathcal{U}}, \quad (2.8)$$

where u is the exact solution of (2.1), M/γ_h is the stability constant, and the best approximation error $\inf_{\omega_h \in \mathcal{U}_h} \|u - \omega_h\|_{\mathcal{U}}$ is measured in the trial norm $\|\cdot\|_{\mathcal{U}}$. The continuous inf-sup condition (2.3) does *not* in general imply the discrete inf-sup condition (2.7).

In the DPG method, the issue of discrete stability is solved by finding a unique test space, called the *optimal test space* \mathcal{V}^{opt} . Given any trial space $\mathcal{U}_h \subset \mathcal{U}$, define its optimal test space by,

$$\mathcal{V}^{opt} := T(\mathcal{U}_h), \quad (2.9)$$

where the *trial-to-test* operator $T : \mathcal{U} \rightarrow \mathcal{V}$ is defined by,

$$(Tu_h, v)_{\mathcal{V}} = b(u_h, v) \quad \forall u_h \in \mathcal{U}_h, v \in \mathcal{V}. \quad (2.10)$$

For any $u_h \in \mathcal{U}_h$, equation (2.10) uniquely defines Tu_h by the Riesz representation theorem. Let $\mathcal{B} : \mathcal{U} \rightarrow \mathcal{V}'$ denote the linear operator induced by the form $b(\cdot, \cdot)$,

$$\langle \mathcal{B}u, v \rangle_{\mathcal{V}' \times \mathcal{V}} = b(u, v), \quad v \in \mathcal{V}, \quad (2.11)$$

where $\langle \cdot, \cdot \rangle_{\mathcal{V}' \times \mathcal{V}}$ denotes the duality pairing on $\mathcal{V}' \times \mathcal{V}$. Then, $T = \mathcal{R}_{\mathcal{V}}^{-1} \mathcal{B}$, where $\mathcal{R}_{\mathcal{V}} : \mathcal{V} \rightarrow \mathcal{V}'$ is the Riesz map. In other words, for every trial function u_h , the trial-to-test operator defines a unique optimal test function, $v_h = \mathcal{R}_{\mathcal{V}}^{-1} \mathcal{B}u_h$. The optimal test functions realize the supremum in the inf-sup condition. Indeed,

$$\sup_{0 \neq v \in \mathcal{V}} \frac{|b(u_h, v)|}{\|v\|_{\mathcal{V}}} = \sup_{0 \neq v \in \mathcal{V}} \frac{|(Tu_h, v)|}{\|v\|_{\mathcal{V}}} \leq \|Tu_h\|_{\mathcal{V}} = \frac{|b(u_h, v_h)|}{\|v_h\|_{\mathcal{V}}}. \quad (2.12)$$

Therefore, $\gamma_h \geq \gamma$, i.e., discrete stability is guaranteed by construction.

2.2 Breaking the test space

In the discussion so far, we have neglected computational aspects of the DPG method. One question that arises immediately in the context of practicality is the cost of the inversion of the global Riesz map $\mathcal{R}_{\mathcal{V}}$. Let Ω denote the global domain of interest, with Lipschitz boundary $\Gamma \equiv \partial\Omega$, and let Ω_h denote a suitable finite element triangulation of Ω , with mesh skeleton Γ_h . By “breaking” the test space, i.e., employing a larger discontinuous test space, $\mathcal{V}(\Omega_h) \supset \mathcal{V}(\Omega)$, the inversion of the Riesz map on Ω is localized and can be done independently element-wise. The element-local computational costs are still significant but they can be parallelized efficiently and fast integration techniques can be implemented to accelerate computation by more than one order of magnitude [23, 4]. By reducing the regularity of the test space, new interface terms arise on the mesh skeleton with interface unknowns \hat{u} . The resulting variational problem is,

$$\begin{cases} u \in \mathcal{U}, \hat{u} \in \hat{\mathcal{U}}, \\ b(u, v) + \langle \hat{u}, v \rangle_{\Gamma_h} = l(v), \quad v \in \mathcal{V}(\Omega_h), \end{cases} \quad (2.13)$$

where $\langle \cdot, \cdot \rangle_{\Gamma_h}$ denotes an appropriate duality pairing on the mesh skeleton. The new interface unknowns may be interpreted as Lagrange multipliers [8].

The stability of the formulation with broken test spaces is inherited from the continuous problem. In particular, the broken formulation (2.13) is well-posed with a mesh-independent stability constant of the same order as the stability constant for the continuous problem [5].

2.3 The practical DPG method

Until now, the trial-to-test operator T has only been defined in the infinite-dimensional setting (2.10). To compute optimal test functions in practice, the inversion of the Riesz map must be approximated on a *truncated* finite-dimensional test space $\mathcal{V}^r \subset \mathcal{V}$, $\dim(\mathcal{V}^r) \gg \dim(\mathcal{U}_h)$, also called the *enriched test space*.

With the Riesz map defined on the truncated test space, $\mathcal{R}_{\mathcal{V}^r} : \mathcal{V}^r \rightarrow (\mathcal{V}^r)'$, the approximate trial-to-test operator $T^r : \mathcal{U}_h \rightarrow \mathcal{V}^r$ is defined by,

$$T^r := \mathcal{R}_{\mathcal{V}^r}^{-1} \iota^T \mathcal{B}, \quad (2.14)$$

where $\iota : \mathcal{V}^r \rightarrow \mathcal{V}$ is the inclusion map.

Consequently, the practical DPG method with optimal test functions solves,

$$\begin{cases} u_h \in \mathcal{U}_h \subset \mathcal{U}, \\ b(u_h, v_h) = l(v_h), \quad v_h \in \mathcal{V}^r = T^r \mathcal{U}_h, \end{cases} \quad (2.15)$$

with the additional interface term from (2.13) when breaking the test space \mathcal{V}^r .

Because the optimal test functions are approximated, some stability loss is inevitable. Several papers have addressed the issue of controlling and quantifying the stability loss in the DPG method [16, 24].

3 Pollution study

We first introduce the functional setting for the broken ultraweak DPG formulation of the time-harmonic Maxwell equations in a linear waveguide. Then, we discuss how the frequency ω enters the DPG error analysis, and how an *hp*-refinement strategy can effectively control the pollution error. We report numerical results for the propagation of the fundamental mode in a rectangular waveguide.

3.1 Function spaces

For our time-harmonic Maxwell model problem, we define the following standard energy spaces on a bounded domain $\Omega \subset \mathbb{R}^3$ with Lipschitz boundary $\Gamma \equiv \partial\Omega$:

$$\begin{aligned} L^2(\Omega) &:= \{y : \Omega \rightarrow \mathbb{C} : \|y\| < \infty\}, \\ H(\text{curl}, \Omega) &:= \{\mathbf{q} : \Omega \rightarrow \mathbb{C}^3 : \mathbf{q} \in (L^2(\Omega))^3, \nabla \times \mathbf{q} \in (L^2(\Omega))^3\}, \end{aligned} \quad (3.1)$$

where $\|\cdot\|$ is the $L^2(\Omega)$ norm.

Suppose Ω is partitioned into a set Ω_h of open disjoint elements $\{K\}_{K \in \Omega_h}$ with Lipschitz element boundaries $\{\partial K\}_{K \in \Omega_h}$. The broken energy spaces defined on the finite element mesh Ω_h are:

$$\begin{aligned} L^2(\Omega_h) &:= \{y \in L^2(\Omega) : y|_K \in L^2(K) \ \forall K \in \Omega_h\} = L^2(\Omega), \\ H(\text{curl}, \Omega_h) &:= \{\mathbf{q} \in (L^2(\Omega))^3 : \mathbf{q}|_K \in H(\text{curl}, K) \ \forall K \in \Omega_h\} \supset H(\text{curl}, \Omega). \end{aligned} \quad (3.2)$$

Additionally, we must define energy spaces for the trace unknowns that arise from breaking the test space. These spaces are defined on the mesh skeleton Γ_h . In the model problem, we need the following trace space:

$$H^{-1/2}(\text{curl}, \Gamma_h) := \{\hat{\mathbf{q}} \in \prod_{K \in \Omega_h} H^{-1/2}(\text{curl}, \partial K) : \exists \mathbf{q} \in H(\text{curl}, \Omega) : \gamma_t(\mathbf{q}|_K) = \hat{\mathbf{q}}\}, \quad (3.3)$$

where the continuous and surjective tangential trace operator is defined element-wise [5]:

$$\gamma_t : H(\text{curl}, \Omega_h) \rightarrow \prod_{K \in \Omega_h} H^{-1/2}(\text{curl}, \partial K). \quad (3.4)$$

Finally, we introduce the minimum energy extension norm for the trace unknowns:

$$\|\hat{\mathbf{q}}\|_{H^{-1/2}(\text{curl}, \Gamma_h)} := \inf_{\substack{\mathbf{q} \in H(\text{curl}, \Omega) \\ \gamma_t(\mathbf{q}|_K) = \hat{\mathbf{q}}}} \|\mathbf{q}\|_{H(\text{curl}, \Omega)}. \quad (3.5)$$

3.2 Problem formulation

We consider the linear, isotropic time-harmonic Maxwell problem:

$$\begin{cases} \nabla \times \mathbf{E} = -i\omega\mu\mathbf{H} & \text{in } \Omega, \\ \nabla \times \mathbf{H} = i\omega\varepsilon\mathbf{E} & \text{in } \Omega, \\ \mathbf{n} \times \mathbf{E} = \mathbf{n} \times \mathbf{E}_0 & \text{on } \Gamma, \end{cases} \quad (3.6)$$

where \mathbf{E} and \mathbf{H} are the complex vector-valued time-harmonic electric and magnetic field, respectively; ω is the angular wave frequency; \mathbf{n} is the outward unit normal; and ε and μ are the scalar-valued electric permittivity and magnetic permeability, respectively. And we assume sufficiently regular boundary data.

3.2.1 Ultraweak formulation

The ultraweak variational formulation is obtained by testing (3.6) with test functions (\mathbf{F}, \mathbf{G}) , integrating over Ω , and relaxing both equations:

$$\begin{cases} \mathbf{E}, \mathbf{H} \in (L^2(\Omega))^3, \\ (\mathbf{E}, \nabla \times \mathbf{F}) + (i\omega\mu\mathbf{H}, \mathbf{F}) = -\langle \mathbf{n} \times \mathbf{E}_0, \mathbf{F} \rangle_\Gamma, \quad \mathbf{F} \in H(\text{curl}, \Omega), \\ (\mathbf{H}, \nabla \times \mathbf{G}) - (i\omega\varepsilon\mathbf{E}, \mathbf{G}) = 0, \quad \mathbf{G} \in H(\text{curl}, \Omega) : \mathbf{n} \times \mathbf{G} = 0 \text{ on } \Gamma. \end{cases} \quad (3.7)$$

3.2.2 Broken ultraweak formulation

By breaking the test space, we introduce new trace unknowns on the mesh skeleton Γ_h . The broken ultraweak variational formulation is:

$$\begin{cases} \mathbf{E}, \mathbf{H} \in (L^2(\Omega))^3, \\ \hat{\mathbf{E}} \in \hat{\mathcal{U}}_1, \hat{\mathbf{H}} \in \hat{\mathcal{U}}_2, \\ (\mathbf{E}, \nabla_h \times \mathbf{F}) + \langle \mathbf{n} \times \hat{\mathbf{E}}, \mathbf{F} \rangle_{\Gamma_h} + (i\omega\mu\mathbf{H}, \mathbf{F}) = 0, \quad \mathbf{F} \in H(\text{curl}, \Omega_h), \\ (\mathbf{H}, \nabla_h \times \mathbf{G}) + \langle \mathbf{n} \times \hat{\mathbf{H}}, \mathbf{G} \rangle_{\Gamma_h} - (i\omega\varepsilon\mathbf{E}, \mathbf{G}) = 0, \quad \mathbf{G} \in H(\text{curl}, \Omega_h), \end{cases} \quad (3.8)$$

where h denotes element-wise operations, and,

$$\hat{\mathcal{U}}_1 := \{ \hat{\mathbf{q}} \in H^{-1/2}(\text{curl}, \Gamma_h) : \mathbf{n} \times \hat{\mathbf{q}} = \mathbf{n} \times \mathbf{E}_0 \text{ on } \Gamma \}, \quad (3.9)$$

$$\hat{\mathcal{U}}_2 := H^{-1/2}(\text{curl}, \Gamma_h). \quad (3.10)$$

3.3 Pollution estimates

The mathematical setting for different variational formulations of the time-harmonic Maxwell equations in context of DPG is analyzed in much detail in [5]. We recap a few points that are relevant for our discussion regarding the ultraweak formulation.

Define the following group variables:

$$\mathbf{u} = (\mathbf{E}, \mathbf{H}), \quad \mathbf{v} = (\mathbf{F}, \mathbf{G}). \quad (3.11)$$

The Maxwell operator from (3.6) can be written as:

$$A\mathbf{u} = (\nabla \times \mathbf{E} + i\omega\mu\mathbf{H}, \nabla \times \mathbf{H} - i\omega\varepsilon\mathbf{E}); \quad (3.12)$$

then, the bilinear form corresponding to the ultraweak formulation (3.7) is:

$$b(\mathbf{u}, \mathbf{v}) = (\mathbf{u}, A^*\mathbf{v}), \quad (3.13)$$

with the formal Adjoint operator defined by:

$$A^*\mathbf{v} = (\nabla \times \mathbf{F} + i\omega\varepsilon\mathbf{G}, \nabla \times \mathbf{G} - i\omega\mu\mathbf{F}). \quad (3.14)$$

In the ideal ultraweak DPG method with unbroken test spaces, the Adjoint graph norm, $\|\mathbf{v}\|_{\mathcal{V}} = \|A^*\mathbf{v}\|$, for the test space yields the *optimal test norm*; with this norm, the method delivers the L^2 projection. For the broken formulation (3.8) we define an additional trace group variable, $\hat{\mathbf{u}} = (\hat{\mathbf{E}}, \hat{\mathbf{H}})$, and obtain the bilinear form:

$$b(\mathbf{u}, \mathbf{v}) + \hat{b}(\hat{\mathbf{u}}, \mathbf{v}) = (\mathbf{u}, A^*\mathbf{v})_{\Omega_h} + \langle \mathbf{n} \times \hat{\mathbf{E}}, \mathbf{F} \rangle_{\Gamma_h} + \langle \mathbf{n} \times \hat{\mathbf{H}}, \mathbf{G} \rangle_{\Gamma_h}. \quad (3.15)$$

The broken ultraweak variational problem (3.8) can then be written as:

$$\begin{cases} \mathbf{u} \in \mathcal{U}, \hat{\mathbf{u}} \in \hat{\mathcal{U}}, \\ b(\mathbf{u}, \mathbf{v}) + \hat{b}(\hat{\mathbf{u}}, \mathbf{v}) = l(\mathbf{v}), \quad \mathbf{v} \in \mathcal{V}. \end{cases} \quad (3.16)$$

The optimal test norm is not localizable, but we can augment it with an additional term to obtain a quasi-optimal test norm: $\|\mathbf{v}\|_{\mathcal{V}}^2 = \|A^*\mathbf{v}\|^2 + \alpha\|v\|^2$, with scaling parameter $\alpha \in \mathcal{O}(1)$. The quasi-optimal test norm is robustly equivalent with the optimal test norm, i.e. independent of the frequency ω , and the robust stability constant is maintained in the broken formulation [5]. This implies that the approximation error is bounded by the best approximation error (BAE) uniformly in ω :

$$\|\mathbf{u} - \mathbf{u}_h\|^2 + \|\hat{\mathbf{u}} - \hat{\mathbf{u}}_h\|_{\hat{q}}^2 \leq C \left[\inf_{\mathbf{w}_h} \|\mathbf{u} - \mathbf{w}_h\| + \inf_{\hat{\mathbf{w}}_h} \|\hat{\mathbf{u}} - \hat{\mathbf{w}}_h\|_{\hat{q}} \right], \quad (3.17)$$

where constant C is independent of the mesh and frequency ω , and $\|\cdot\|_{\hat{q}}$ refers to the minimum energy extension norm defined in (3.5). The estimate also implies that the L^2 best approximation is pollution free because it is independent of ω . In one dimension, the BAE for the traces is zero, thus the method is in fact pollution free [10, 27]. In multiple dimensions, however, the BAE for the traces, measured in $\|\cdot\|_{\hat{q}}$, does depend on the frequency ω and the method exhibits numerical pollution.

The estimate for the standard Galerkin method on the other hand has a stability constant that is not ω -independent, thus the Galerkin discretization is not robustly stable. The DPG method hides the perturbation parameter ω in the best approximation and by doing so yields a stable discretization for any wavenumber. This can practically be exploited by starting computation on a coarse mesh where the pollution error is high, and driving hp -adaptivity with the DPG error indicator. This approach yields superior meshes for resolving localized waves [27, 26].

A wavenumber explicit analysis for the Helmholtz equation is presented for the DPG method in [10] and for the Galerkin method in [22]. For the Galerkin discretization, Melenk and Sauter show that quasi-optimality is obtained under the conditions that $\omega h/p$ is sufficiently small and p is at least $\mathcal{O}(\log \omega)$, where h is the mesh size and p the polynomial order of approximation [22]. Based on these estimates, the best approach to dealing with the pollution error may be an hp -strategy that preferably increases the polynomial order p while keeping ωh constant for increasing frequency.

In the next section, we study the pollution error with numerical experiments for many wavelengths and discuss the observations with regard to the suggested hp -strategy and its applicability to the DPG method for the time-harmonic Maxwell problem.

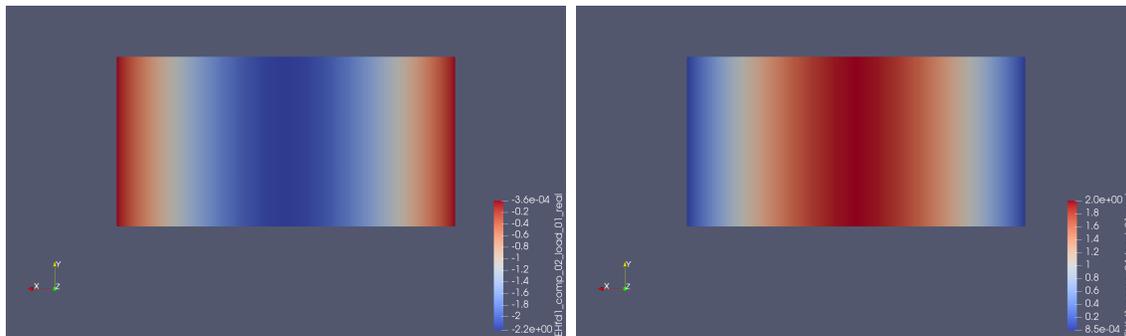
3.4 Numerical results

The propagation of an electromagnetic field in a waveguide is governed by the Maxwell equations. We assume that the time-harmonic setting is justified and that the waveguide medium is nonmagnetic, dielectric, and no free charges are present. While non-linear effects and anisotropic, inhomogeneous material properties play important roles in research on fiber optics, for the purpose of this pollution study we will assume the waveguide medium is linear, isotropic and homogeneous. If we prescribe idealized PEC

boundary conditions, then under these assumptions the Maxwell equations reduce to (3.6). Note that in fact we do prescribe non-zero Dirichlet boundary conditions (BCs) at the input to excite the waveguide and impedance BCs at the output but PEC BCs everywhere else. For this simplified setting, the propagating field in a waveguide can be described as a superposition of guided modes. These modes are eigenfunction solutions to the reduced scalar Helmholtz problem which can be solved analytically for simple domains (e.g, see [17, 20]). Consider a rectangular waveguide with the domain Ω , in Cartesian coordinates:

$$\Omega = (0.0, 1.0) \times (0.0, 0.5) \times (0, L),$$

where L is the length of the waveguide. The fundamental mode in this waveguide is the transverse electric TE_{10} mode, depicted in Fig. 1. The fundamental mode is not very oscillatory in the transverse direction. At the waveguide end, we employ an absorbing impedance boundary condition that matches the wave impedance for the fundamental mode. In the rectangular waveguide experiments, the cross-section is modeled with two hexahedral elements which is justified by the simple transverse mode profile (cf. Fig. 1).



(a) Electric field component E_y

(b) Magnetic field component H_x

Figure 1: TE_{10} transverse fields in rectangular waveguide in a plane normal to the z -axis

In our first experiment, we analyze the relative field error, measured in the L^2 norm, for the propagating fundamental mode in waveguides of different length L . The smallest waveguide has a length equivalent to one wavelength of the fundamental mode, and the longest one has 8192 wavelengths. As we increase the length L , we keep the number of elements per wavelength (i.e., degrees of freedom (DOFs) per wavelength) constant. In particular, we choose a discretization with four elements per wavelength. Fig. 2 shows the relative field error for these waveguides for uniform order of approximation, ranging from $p = 4$ to $p = 8$. In all numerical experiments, we are using the enrichment order $\Delta p = 1$ for the test space to approximate optimal test functions.

We make several observations: for a fixed number of wavelengths, higher polynomial order yields significantly smaller (more than one order of magnitude) errors, as

expected; for every order of approximation, the field error starts to increase if the waveguide is long enough despite keeping the DOFs per wavelength constant; and for higher p , this pollution effect is “kicking in” at a later point, i.e., more wavelengths can be computed with higher order before the pollution error is measurable. Furthermore, to maintain some desired accuracy, one needs to increase the polynomial order in nearly regular intervals. For example, to achieve 1% accuracy for 4 wavelengths, it is sufficient to use $p = 4$; at 64 wavelengths, $p = 5$ is needed; with $p = 6$, computing up to 1024 wavelengths is feasible with this error margin; and $p = 7$ would most likely be sufficient for 16384 wavelengths. At a closer look, these intervals resemble a logarithmic dependency on the polynomial order ($4 * 2^4 = 64$, $64 * 2^4 = 1024$, $1024 * 2^4 = 16384$). In other words, these results corroborate theoretical estimates by Melenk and Sauter predicting that control of the pollution error would require increasing p logarithmically with the wavenumber.

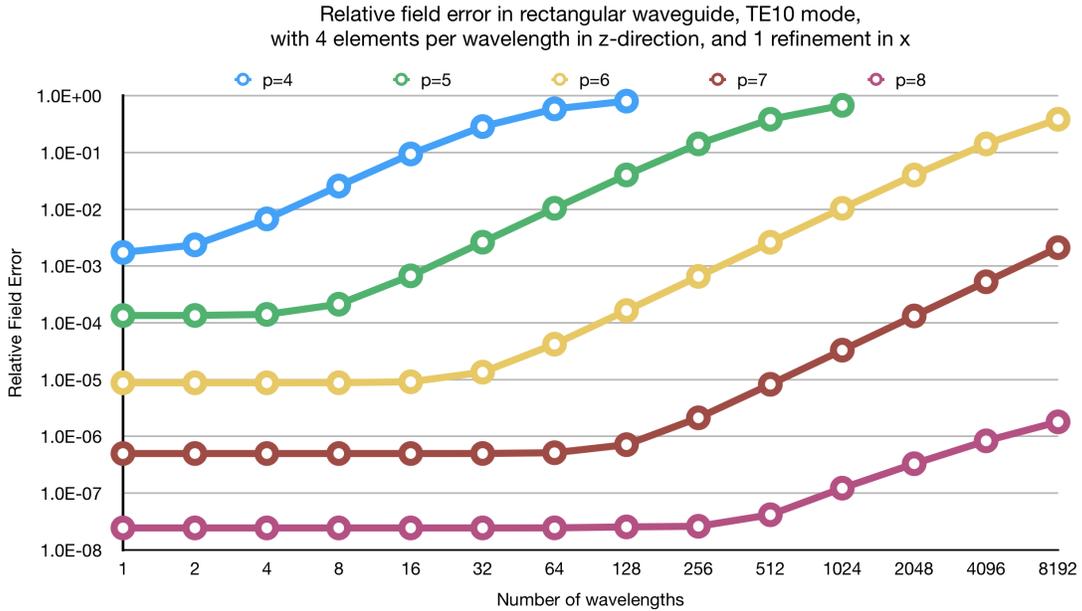


Figure 2: Relative field error for uniform order of approximation

In our experiments, the pollution was primarily a diffusive error causing wave attenuation. This is in agreement with previous observations for the DPG method [10]. A practical way of measuring this diffusivity in waveguide applications is to compute power flux through the cross-section of the waveguide at different points in z . In a linear, dielectric waveguide with PEC boundary conditions the fundamental mode should be carried without loss of power. Fig. 3 shows the measured power loss between the waveguide input ($z = 0$) and output ($z = L$) for different polynomial orders. Note that $p = 8$ has less than 0.005% loss of power in all tested waveguides. The pollution error is clearly visible in terms of power loss. We also observe the same logarithmic dependency for increasing polynomial order, illustrated by the near-equidistant parallel

character of the lines.

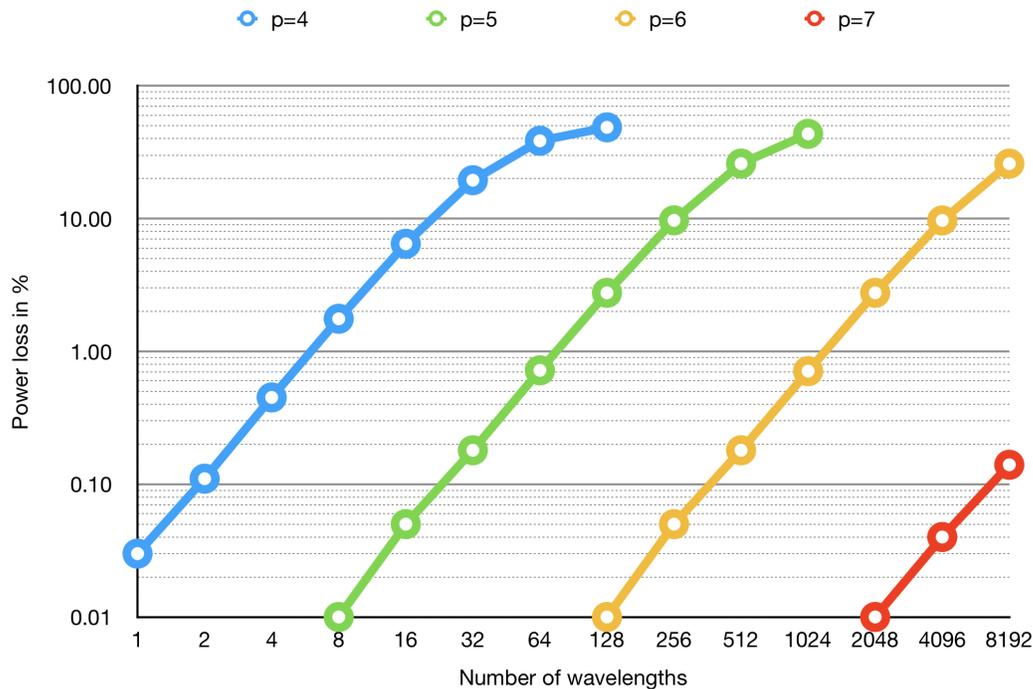


Figure 3: Power loss for uniform order of approximation

Moving on to additional experiments, we keep our focus on the same rectangular waveguide but with different potential approaches of dealing with the pollution error. It may be reasonable to assume that since the wave is propagating in one direction (along z), it will be sufficient to increase the order of approximation anisotropically or to increase the number of elements through anisotropic h -refinements in z . Exploring both of these options (cf. Fig. 4), we find that neither one of these approaches yields satisfactory results. First, in Fig. 4a, we use fifth-order polynomials in the radial discretization ($px = py = 5$) of the waveguide and increase the anisotropic order from $pz = 4$ to $pz = 7$. While the error decreases initially, it begins stagnating at $pz = 6$ (note that the $pz = 7$ error coincides almost exactly with $pz = 6$). The same observation is made for uniform order $p = 5$ with varying number of elements per wavelength (ranging from 2 elements to 16 elements). Our findings indicate that the pollution error comes from an interplay between the mode resolution (radial discretization) and the wave resolution in the direction of propagation. In other words, *increasing the number of DOFs anisotropically does not suffice asymptotically to control the pollution error*.

Finally, we measure the loss of power for both anisotropic refinement cases, plotted in Fig. 5. Expectably, we observe the same stagnation in the diffusive pollution, consistent with the errors measured in the previous plot.

We have conducted these experiments on different waveguides (rectangular waveguides, circular waveguides, and step-index fibers) with various propagating modes,

and the observations are all consistent with the observations presented up to here; we therefore omit showing additional numerical results for those cases.

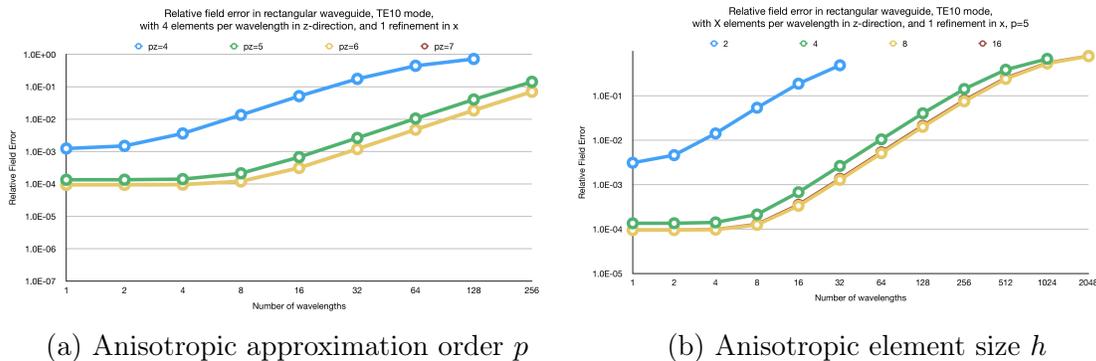


Figure 4: Relative field error for anisotropic refinements

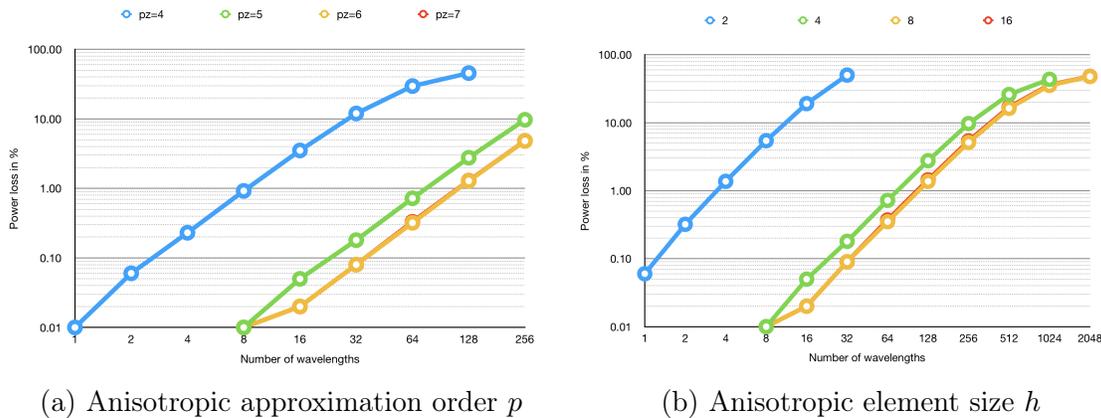


Figure 5: Power loss for anisotropic refinements

4 Adaptivity study

In our adaptivity study, we focus on a different aspect of resolving the propagating wave. We have shown that the interplay between resolving the wave along the direction of propagation and resolving the transverse mode profile is important in controlling the pollution error. For that reason, the finite element mesh should be sensitive to different mode profiles and adapt to resolve them appropriately. This is especially important in waveguide applications where significant transfer of power occurs between different guided modes. Our tool for adapting the mesh “on-the-fly” is the DPG residual that serves as an error estimator in the energy norm. We briefly recap the important points in the derivation of this indicator and proceed with numerical experiments in multi-mode step-index fibers. Lastly, we will look at the load imbalance that results from adapting the mesh to different propagating modes.

4.1 The DPG error indicator

The DPG method can be reformulated as a minimum residual method [8],

$$u_h = \arg \min_{w_h \in \mathcal{U}_h} \|l - \mathcal{B}w_h\|_{\mathcal{V}'}^2 = \arg \min_{w_h \in \mathcal{U}_h} \|\mathcal{R}_{\mathcal{V}}^{-1}(l - \mathcal{B}w_h)\|_{\mathcal{V}}^2, \quad (4.1)$$

where the residual is minimized in the dual test norm $\|\cdot\|_{\mathcal{V}'}$. Taking the Gâteaux derivative, we obtain a minimum residual formulation,

$$\begin{cases} u_h \in \mathcal{U}_h, \\ (\mathcal{R}_{\mathcal{V}}^{-1}(l - \mathcal{B}u_h), \mathcal{R}_{\mathcal{V}}^{-1}\mathcal{B}w_h)_{\mathcal{V}} = 0, \quad w_h \in \mathcal{U}_h. \end{cases} \quad (4.2)$$

Furthermore, we define the energy norm $\|\cdot\|_E$ on the trial space \mathcal{U} by,

$$\|u\|_E := \|\mathcal{B}u\|_{\mathcal{V}'} = \|\mathcal{R}_{\mathcal{V}}^{-1}\mathcal{B}u\|_{\mathcal{V}}. \quad (4.3)$$

We define ψ as the Riesz representation of the residual,

$$\psi := \mathcal{R}_{\mathcal{V}}^{-1}(l - \mathcal{B}u_h). \quad (4.4)$$

Notice that when u_h minimizes the residual (4.1), then,

$$(\psi, \mathcal{R}_{\mathcal{V}}^{-1}\mathcal{B}w_h)_{\mathcal{V}} = 0, \quad w_h \in \mathcal{U}_h. \quad (4.5)$$

We arrive at a mixed Galerkin formulation,

$$\begin{cases} u_h \in \mathcal{U}_h, \psi \in \mathcal{V}, \\ (\psi, v)_{\mathcal{V}} + b(u_h, v) = l(v), \quad v \in \mathcal{V}, \\ b(w_h, \psi) = 0, \quad w_h \in \mathcal{U}_h. \end{cases} \quad (4.6)$$

The error measured in the energy norm can be computed explicitly,

$$\|u - u_h\|_E = \|\mathcal{B}(u - u_h)\|_{\mathcal{V}'} = \|l - \mathcal{B}u_h\|_{\mathcal{V}'} = \|\mathcal{R}_{\mathcal{V}}^{-1}(l - \mathcal{B}u_h)\|_{\mathcal{V}} = \|\psi\|_{\mathcal{V}}, \quad (4.7)$$

hence $\|\psi\|$ offers a built-in a-posteriori error indicator. Finally, note that the choice of the test norm $\|\cdot\|_{\mathcal{V}}$ is critical, as it dictates the norm in which the method converges. The natural choice for the test norm in the ultraweak formulation is the adjoint graph norm [8].

4.2 Multi-mode step-index fibers

We consider a dielectric step-index waveguide. More precisely, consider a weakly-guiding large mode area (LMA) step-index fiber made of silica glass. See Tab. 1 for a description of the model parameters for the fiber. For weakly-guiding fibers, $(n_{core} - n_{clad})/n_{core} \ll 1$, and the guided modes are linearly polarized (LP) modes. The V -number is a “normalized frequency” that determines how many guided modes

are supported by the particular fiber. For example, if $V < 2.405$, then the fiber is single-mode. LMA fibers have a relatively large core radius and support multiple modes; the fiber we are using has $V \approx 4.43$, and it supports four guided modes: $\{LP_{01}, LP_{11}, LP_{21}, LP_{02}\}$. The fiber axis is assumed to be aligned with the z -axis, and the length of the fiber is L . Fig. 6 illustrates the guided modes for this particular fiber, showing the magnitude of the electric field in the center of the fiber cross-section.

For multi-mode propagation we are using a perfectly matched layer (PML) absorbing boundary condition. We refer to [28] for details on the derivation and implementation of a PML for the DPG method.

Table 1: Step-index fiber: parameters

Description	Symbol	Value
Wavelength	λ	1064 nm
Core radius	r_{core}	12.7 μm
Cladding radius	r_{clad}	127 μm
Core refractive index	n_{core}	1.4512
Cladding refractive index	n_{clad}	1.45
Numerical aperture	NA	0.059
V-number	V	4.43

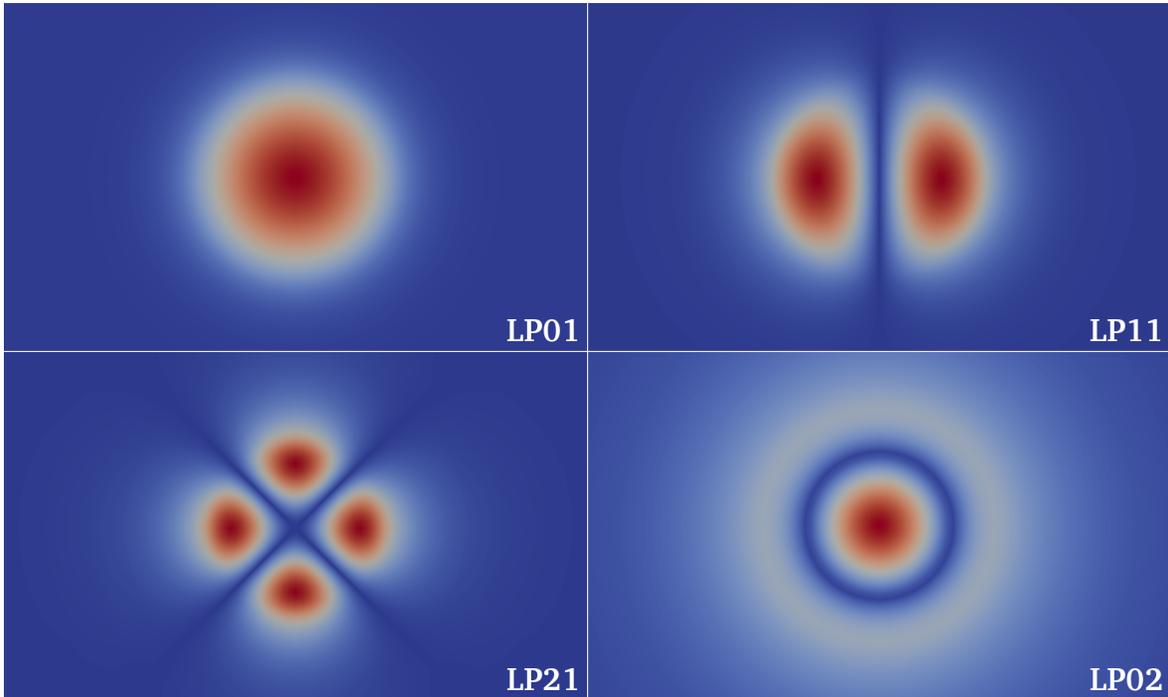


Figure 6: Guided modes in LMA fiber (magnitude of the electric field)

For this particular fiber, the power confinement (amount of energy confined to the

core region) of each mode is,

$$\{\Gamma_{01}, \Gamma_{11}, \Gamma_{21}, \Gamma_{02}\} \approx \{96.11, 88.77, 74.79, 59.58\}\%. \quad (4.8)$$

Clearly, the optimal discretization of the fiber cross-section is different for each of these modes. That is, to capture the oscillations of the higher-order modes near the core-cladding interface, a finer discretization is needed than for the fundamental mode LP_{01} . In particular, higher-order modes demand more refinements (or degrees of freedom) outside of the fiber core, when compared to the fundamental mode that is mostly confined to the core region. Therefore, one geometry cannot be optimal for capturing any propagating mode.

Suppose we are interested in simulating the transverse mode instability (TMI) phenomenon [14] in active gain fiber amplifiers. The TMI is characterized by the chaotic transfer of energy between the fundamental mode and the higher-order modes. One challenge in computing a numerical solution to the resulting nonlinear Maxwell problem is to capture modes accurately when they occur. With mode instabilities, it is not known a-priori which modes will be propagating in which parts of the fiber. Refining the initial geometry globally to better resolve higher-order modes increases the computational cost dramatically and may render the computation infeasible for large problem instances. Adaptivity, on the other hand, can be used to refine the mesh where it is needed for capturing these modes locally, and the overall computational cost will be kept significantly lower.

4.3 Numerical experiments

In the following experiments, we are aiming to establish the efficacy of adaptivity based on the DPG residual for resolving higher-order modes. In the broken DPG setting, the residual is computed through element-wise contributions, i.e.,

$$\|\psi\|_{\mathcal{V}}^2 = \sum_{j=1}^n \|\psi|_{K_j}\|_{\mathcal{V}(K_j)}^2, \quad (4.9)$$

where $K_j, j = 1, \dots, n$, denotes the j -th element. After each solve, elements are marked for refinement if they satisfy a certain criterion. We use a strategy for marking elements that is based on Dörfler's marking [13]:

1. Sort the element residuals $\|\psi|_{K_j}\|_{\mathcal{V}(K_j)}^2$ in descending order, i.e.,

$$\|\psi|_{K_1}\|_{\mathcal{V}(K_1)} \geq \|\psi|_{K_2}\|_{\mathcal{V}(K_2)} \geq \dots \geq \|\psi|_{K_n}\|_{\mathcal{V}(K_n)}. \quad (4.10)$$

2. Mark elements $K_j, j = 1, \dots, J$, where $J \leq n$ is the smallest integer for which the following is true:

$$\sum_{j=1}^J \|\psi|_{K_j}\|_{\mathcal{V}(K_j)}^2 \geq \kappa \|\psi\|_{\mathcal{V}}^2, \quad (4.11)$$

where $\kappa \in (0, 1)$.

At this point, with some choice of κ , elements have been marked for refinement. What is not as clear is how to optimally refine each marked element when the hp mesh supports anisotropic adaptive refinements in both element size h and polynomial order p . The choice will ultimately be problem-dependent.

In our experiments, we choose an initial mesh with uniform polynomial order $p = 5$, two elements per wavelength in z -direction (direction of propagation), and a radial (transverse) hybrid discretization using curvilinear hexahedral and prismatic elements.

Fig. 7 illustrates the initial geometry discretization in the fiber cross-section (not drawn to scale): four prisms are used to model the center of the fiber core, and they are surrounded by three layers of four hexahedra each. We refer to these different layers as “domains” and enumerate them from 1 to 4 moving radially outward from the center of the fiber to the cladding boundary. The choice of the initial discretization was informed by the fiber geometry, the fact that all guided modes decay exponentially in the cladding region, and by conducting numerical tests primarily with the fundamental mode. For a relatively short fiber of 16 wavelengths, this initial geometry captures the fundamental mode very well with regard to several physical quantities of interest (e.g., conservation of power, mode confinement). Higher-order modes are not captured quite as well, and for fibers with many wavelengths (i.e., several hundred or a few thousand wavelengths) we observe more significant pollution errors in the propagation of these modes. For example, the errors can be observed in small oscillations of the mode powers along the fiber, diffusive pollution effects, or an unsteady power confinement ratio.

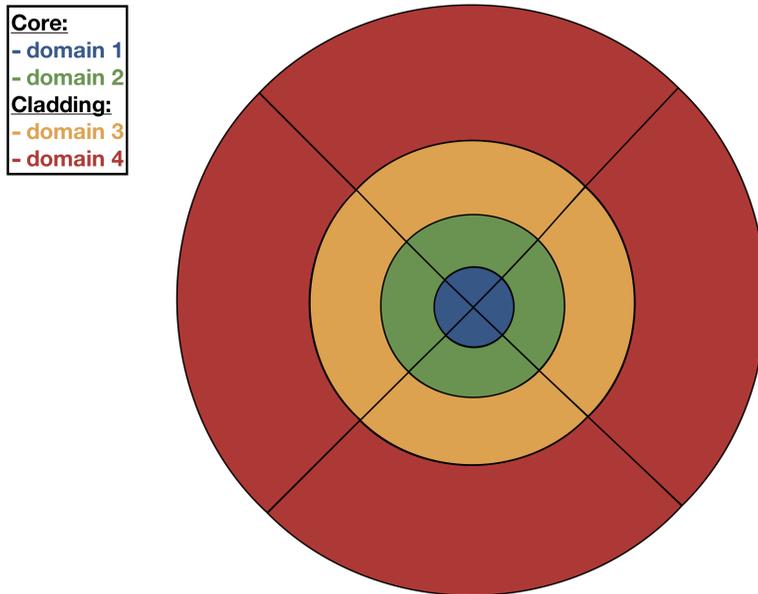


Figure 7: Initial geometry discretization in fiber cross-section (not drawn to scale)

We use DPG to perform multiple adaptive mesh refinements, each based on the respective previous solution and residual, to test the residual error indicator for cap-

turing different modes. As a test case, we will look at the 16 wavelengths fiber. The goal is to observe the sensitivity of the adaptive refinements toward specific modes that are propagating. We choose the parameter $\kappa = 0.5$ in (4.11) and proceed with four adaptive refinement steps after the initial solution. Note that mesh regularity requirements may perform some additional refinements to “close the mesh”, i.e., to obtain a mesh that is 1-irregular.

First, we apply isotropic h -refinements for marked elements. Fig. 8 shows the domains of refinement in the fiber. Each plot illustrates how the mesh is successively refined for one particular guided mode propagating in the fiber. For the fundamental mode, the error indicator is marking elements for refinement primarily in the fiber core, where most of the energy is located. The first three refinement steps exclusively refine in the outer and inner core region. None of the adaptive refinements for higher-order modes refine inside the inner core region. It is notable how sensitive the error indicator is to these different modes. In the case of the LP_{02} mode, the code primarily refines in the outer cladding region; this is likely to be the case because the initial cladding geometry discretization is too coarse to handle the exponential decay of the remaining energy in the transverse field.

Next, we repeat the experiment with anisotropic (radial) h -refinements. Radial refinements are of interest because higher-order modes are only more oscillatory in the transverse field, but they are not more oscillatory in the direction of propagation. In other words, the guided modes have very similar propagation constants (in fact, higher-order modes oscillate slightly slower than the fundamental mode). Therefore, if the numerical pollution is low for the fundamental mode, we may assume that the resolution in the direction of propagation is “good enough” for approximating any higher-order guided modes. Then, radial refinements (in h or p) are the more economical way of capturing these modes. For anisotropic h -adaptive refinements, depicted in Fig. 9, a similar pattern emerges for the higher-order modes but the picture is quite different for the LP_{01} mode. The fundamental mode repeatedly refines elements in the same domain, indicating that the anisotropic refinements do not decrease the local residuals in a way the isotropic ones did. This indicates that the fundamental mode is already well approximated in the transverse field and the residual demands refinements in z -direction for better accuracy of the numerical solution. This interplay between the resolution in different directions is critical when studying the pollution error for guided modes.

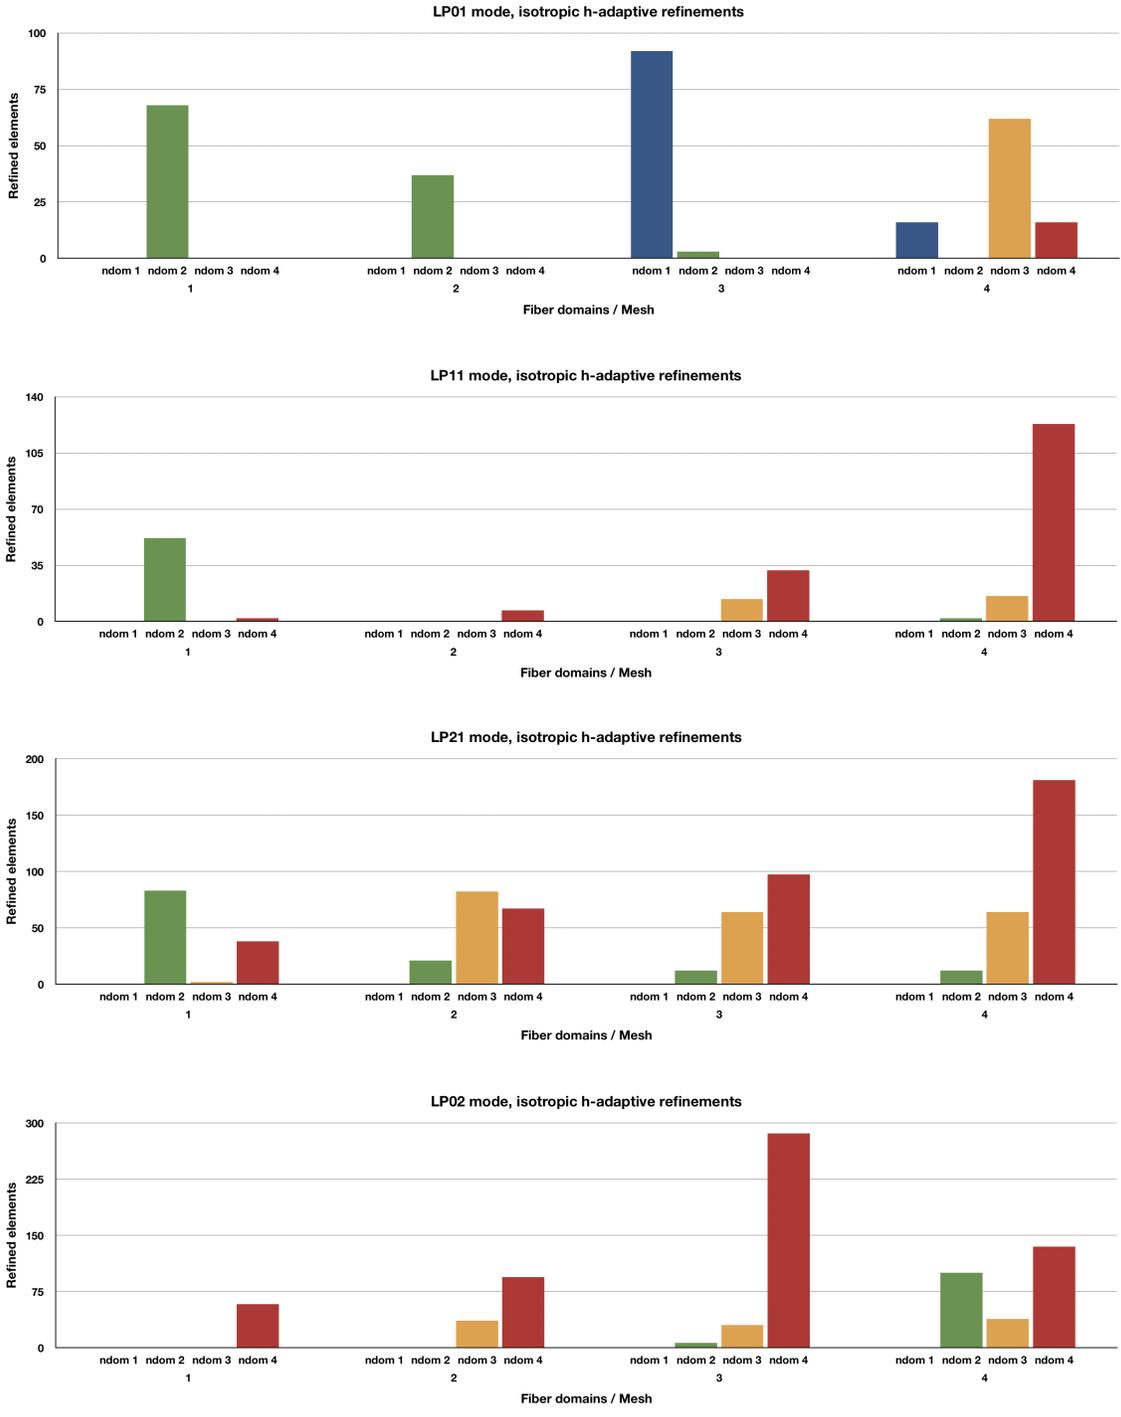


Figure 8: Isotropic h -adaptive refinements: fiber domains

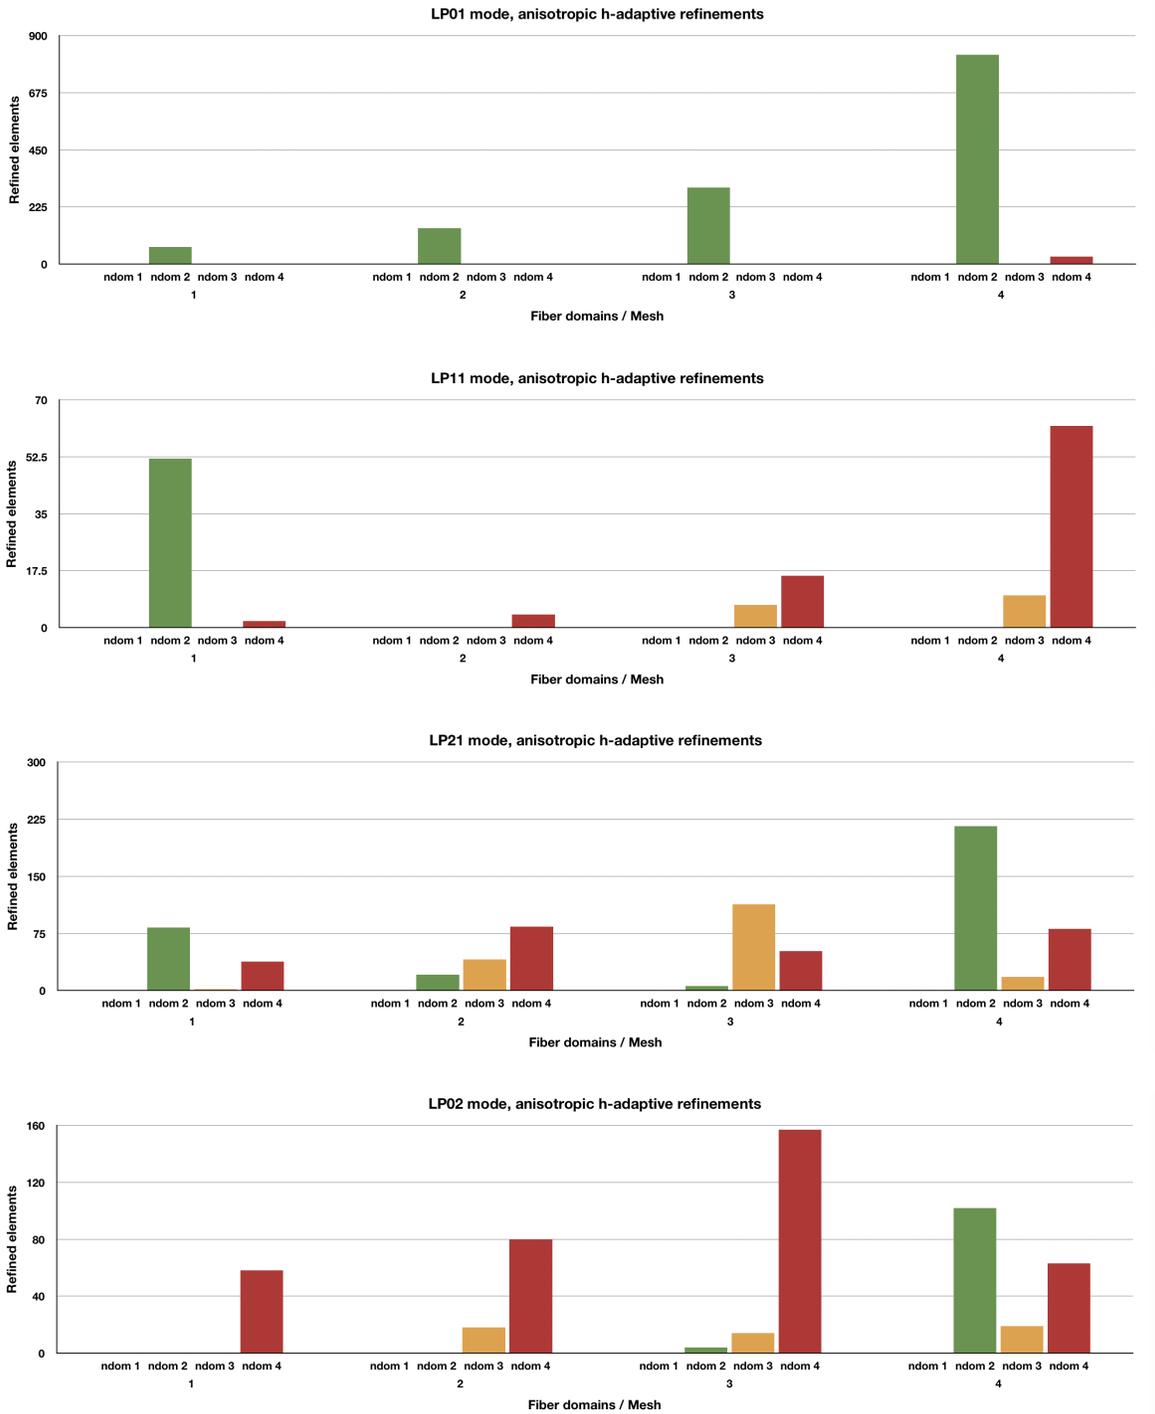


Figure 9: Radial (anisotropic) h -adaptive refinements: fiber domains

Fig. 10 shows how the total residual evolves in both scenarios: we observe that the higher-order modes benefit much from anisotropic refinements, making this the preferred choice for improving the numerical solution with fewer degrees of freedom. For the fundamental mode, we find that the residual does not further decrease through anisotropic refinements indicating the mode is captured quite well by the initial current geometry.

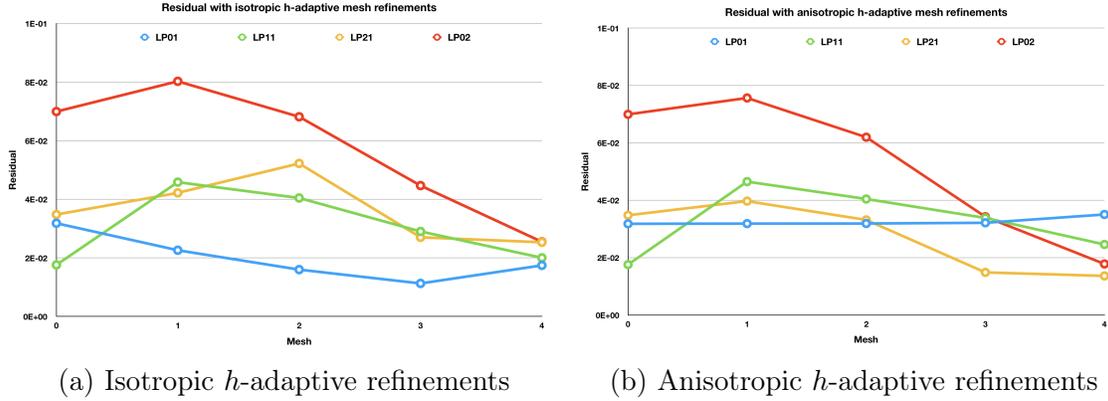


Figure 10: DPG residual in adaptive mesh refinements

4.4 Load balancing

In the parallel computation of the fiber problem, we partition the geometry into subdomains, each owned by one distinct MPI process. The rank of each MPI process is the ID of the subdomain it owns. Initially, we partition the fiber directly based on geometric cuts orthogonal to the fiber axis. Fig. 11 illustrates what the partitioning looks like for four subdomains.

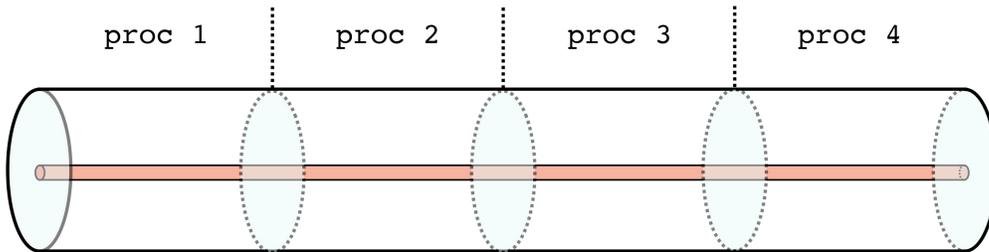


Figure 11: Initial (static) load distribution in the step-index fiber

This initial static partitioning is a good choice because it keeps the interfaces between subdomains small, making it possible to compute large fibers in parallel with a nested dissection solve and obtain good weak scaling. As the adaptive mesh refinements proceed, we must dynamically repartition the domain to retain load balance. A

large number of different repartitioners are available in open source software packages, such as Zoltan [12]. We believe that in this instance, graph partitioning that strives for minimum cuts is a good choice because it keeps the subdomain interfaces relatively small. In the broken ultraweak Maxwell formulation, we are using element interior dofs (electromagnetic fields: L^2) as weights for graph vertices, and trace DOFs on faces (electromagnetic fluxes: $H(\text{curl})$ trace) as weights for the graph's edges. We are omitting connectivities from edge degrees of freedom to provide a sparser graph and accelerate partitioning. ParMetis or PT-Scotch can be used for approximating the partitioning problem. As an alternative, we use a custom dynamic fiber repartitioner that forces orthogonal cuts through the domain while trying to maximize load balance and minimize data migration, similar to recursive coordinate bisection partitioners. This custom repartitioner can perform orders of magnitude faster than graph partitioning because it relies primarily on geometry information.

We are studying how the workload in different subdomains changes without repartitioning. Fig. 12 and Fig. 13 show the workload per MPI rank in a fiber of 16 wavelengths, partitioned into 8 subdomains, with h -adaptive isotropic and anisotropic refinements, respectively. Both plots show the results for the two higher-order modes LP_{21} and LP_{02} . Here, the workload is simply shown as the number of subdomain interior DOFs, i.e., all solution DOFs that are part of a subdomain excluding the trace DOFs on the subdomain interfaces.

Isotropic refinements lead expectedly to higher load imbalance because each isotropic h -refinement increases the local DOFs by another factor of two compared to the anisotropic h -refinement. Two observations stand out: firstly, the subdomains closer to the fiber input appear to exhibit higher residuals hence more refinements are observed in that region; secondly, towards the end of the fiber many refinements are picked up in the sixth subdomain, and almost none in the very last one. The latter observation is an effect from the PML boundary layer at the fiber end. In this short fiber, the PML is active in the last two subdomains (i.e., the layer encompasses about four wavelengths). When the wave enters the boundary layer, it exhibits exponential decay due to the coordinate stretching. This initial decay must be captured accurately by the numerical solution. We see that the DPG residual recognizes the need for more refinements in this region and marks elements in the sixth subdomain. By the time the wave enters the last subdomain, owned by rank seven, it has decayed so far that the residual remains fairly small and almost none of the elements are marked in the adaptive procedure.

It is evident that dynamic load balancing is critical in the simulation of the TMI phenomenon or other applications with transfer energy between guided modes. Through repartitioning of the fiber domain, we obtain significant speed up in the solve, especially for our parallel nested dissection solver.

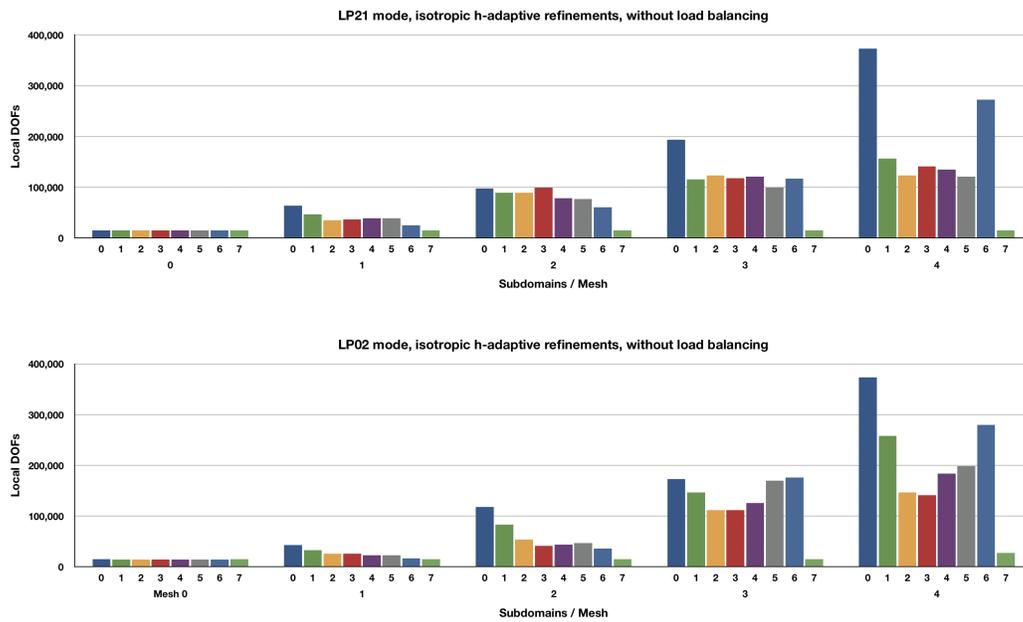


Figure 12: Isotropic adaptive h -refinements: load imbalance

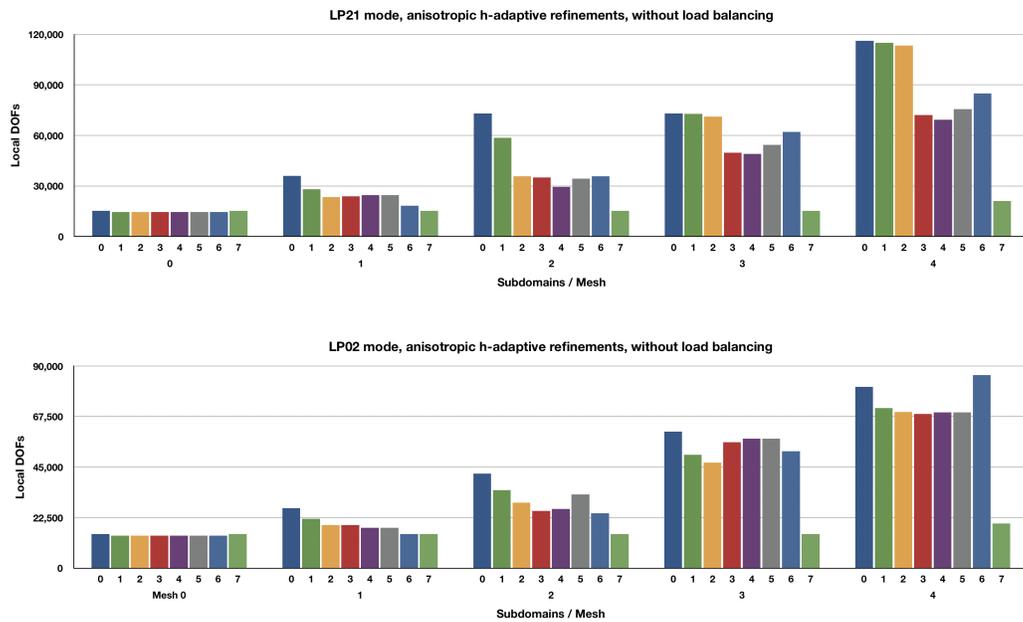


Figure 13: Radial (anisotropic) adaptive h -refinements: load imbalance

5 Summary

We have studied the effect of numerical pollution for guided wave problems with many wavelengths. We found that the pollution error in the ultraweak DPG setting is primarily present in the form of a diffusive effect that causes attenuation of the propagating wave. This is consistent with previous observations for the DPG method where the phase error appears to be relatively small. Furthermore, we were able to corroborate theoretical estimates by Melenk and Sauter, indicating a logarithmic dependency of the pollution error on the polynomial order of approximation. Based on our numerical results, we agree that the best strategy for resolving many wavelengths is one based on hp -refinements that first “capture” the wave (discretizing each wavelength by a few elements with moderate p so that $\omega h/p$ is sufficiently small) and subsequently increase the polynomial order p with $\mathcal{O}(\log \omega)$ if needed while keeping ωh constant. For the 3D waveguide problem where the wave is truly a superposition of transverse modes propagating in one direction, we emphasize that, perhaps counterintuitively, controlling the pollution error asymptotically required additional degrees of freedom in both the transverse direction and the direction of propagation.

For controlling the error in a multi-mode waveguide, we described suitable adaptive refinement strategies. The error indicator we used is based on the Riesz representation of the residual in the ultraweak DPG formulation. With broken test spaces this indicator is computed locally and in parallel. We have shown the importance of adaptive refinements in capturing different propagating modes in a fiber waveguide. For such problems, the efficacy of the DPG residual is remarkable as it sharply recognizes where the “energy” of the solution is located. The numerical results in the weakly-guiding optical fiber waveguide indicate that anisotropic h -adaptive refinements can be a much more efficient choice than isotropic h -adaptive refinements for these problems.

In a distributed-memory computation the repartitioning of the multi-mode fiber becomes essential to maintain load balance. We observed that the PML boundary layer requires additional refinements especially as the wave enters the PML region. Through dynamic repartitioning the load imbalance can be nearly eliminated and the time to solution decreases significantly. This is especially important in problems with localized features in the solution such as the transverse mode instability in fiber amplifiers.

Acknowledgement

This work was partially supported by AFOSR FA9550-17-1-0090 and AFOSR 18RD-COR018.

References

- [1] M. Ainsworth, P. Monk, and W. Muniz. “Dispersive and dissipative properties of discontinuous Galerkin finite element methods for the second-order wave equation”. In: *J. Sci. Comput.* 27.1-3 (2006), pp. 5–40.
- [2] I. Babuška. “Error-bounds for finite element method”. In: *Numer. Math.* 16.4 (1971), pp. 322–333.
- [3] I. Babuška and S. Sauter. “Is the pollution effect of the FEM avoidable for the Helmholtz equation considering high wave numbers?” In: *SIAM J. Numer. Anal.* 34.6 (1997), pp. 2392–2423.
- [4] J. Badger, S. Henneking, and L. Demkowicz. “Fast integration of DPG matrices based on tensorization (Part II - Prismatic elements)”. In: *Oden Institute Report 19-15* (2019).
- [5] C. Carstensen, L. Demkowicz, and J. Gopalakrishnan. “Breaking spaces and forms for the DPG method and applications including Maxwell equations”. In: *Comput. Math. Appl.* 72.3 (2016), pp. 494–522.
- [6] J. Chan, L. Demkowicz, and R. Moser. “A DPG method for steady viscous compressible flow”. In: *Comput. & Fluids* 98 (2014), pp. 69–90.
- [7] L. Demkowicz and J. Gopalakrishnan. “A class of discontinuous Petrov-Galerkin methods. II. Optimal test functions”. In: *Numer. Methods Partial Differential Equations* 27.1 (2011), pp. 70–105.
- [8] L. Demkowicz and J. Gopalakrishnan. “Discontinuous Petrov-Galerkin (DPG) Method”. In: *Encyclopedia of Computational Mechanics Second Edition* (2017), pp. 1–15.
- [9] L. Demkowicz, J. Gopalakrishnan, and A. H. Niemi. “A class of discontinuous Petrov-Galerkin methods. Part III: Adaptivity”. In: *Appl. Numer. Math.* 62.4 (2012), pp. 396–427.
- [10] L. Demkowicz, J. Gopalakrishnan, I. Muga, and J. Zitelli. “Wavenumber explicit analysis of a DPG method for the multidimensional Helmholtz equation”. In: *Comput. Methods Appl. Mech. Engrg.* 213 (2012), pp. 126–138.
- [11] A. Deraemaeker, I. Babuška, and P. Bouillard. “Dispersion and pollution of the FEM solution for the Helmholtz equation in one, two and three dimensions”. In: *Internat. J. Numer. Methods Engrg.* 46.4 (1999), pp. 471–499.
- [12] K. Devine et al. “Zoltan Data Management Services for Parallel Dynamic Applications”. In: *Comput Sci Eng* 4.2 (2002), pp. 90–97.
- [13] W. Dörfler. “A convergent adaptive algorithm for Poisson’s equation”. In: *SIAM J. Numer. Anal.* 33.3 (1996), pp. 1106–1124.

- [14] T. Eidam et al. “Experimental observations of the threshold-like onset of mode instabilities in high power fiber amplifiers”. In: *Opt. Express* 19.14 (2011), pp. 13218–13224.
- [15] F. Fuentes, L. Demkowicz, and A. Wilder. “Using a DPG method to validate DMA experimental calibration of viscoelastic materials”. In: *Comput. Math. Appl.* 325 (2017), pp. 748–765.
- [16] J. Gopalakrishnan and W. Qiu. “An analysis of the practical DPG method”. In: *Math. Comp.* 83.286 (2014), pp. 537–552.
- [17] D. Griffiths. *Introduction to electrodynamics*. AAPT, 2005.
- [18] S. Henneking, J. Grosek, and L. Demkowicz. “A DPG Maxwell approach for studying nonlinear thermal effects in active gain fiber amplifiers”. In: *arXiv preprint arXiv:1912.01185* (2019).
- [19] F. Ihlenburg and I. Babuška. “Finite element solution of the Helmholtz equation with high wave number part II: the hp version of the FEM”. In: *SIAM J. Numer. Anal.* 34.1 (1997), pp. 315–358.
- [20] J. Jackson. *Classical electrodynamics*. John Wiley & Sons, 2012.
- [21] B. Keith, F. Fuentes, and L. Demkowicz. “The DPG methodology applied to different variational formulations of linear elasticity”. In: *Comput. Methods Appl. Mech. Engrg.* 309 (2016), pp. 579–609.
- [22] J.M. Melenk and S. Sauter. “Wavenumber explicit convergence analysis for Galerkin discretizations of the Helmholtz equation”. In: *SIAM J. Numer. Anal.* 49.3 (2011), pp. 1210–1243.
- [23] J. Mora and L. Demkowicz. “Fast integration of DPG matrices based on sum factorization for all the energy spaces”. In: *Comput. Methods Appl. Math.* 19.3 (2019), pp. 523–555.
- [24] S. Nagaraj, S. Petrides, and L. Demkowicz. “Construction of DPG Fortin operators for second order problems”. In: *Comput. Math. Appl.* 74.8 (2017), pp. 1964–1980.
- [25] S. Nagaraj et al. “A 3D DPG Maxwell approach to nonlinear Raman gain in fiber laser amplifiers”. In: *J. Comput. Phys.* 2 (2019), p. 100002.
- [26] S. Petrides. “Adaptive multilevel solvers for the discontinuous Petrov–Galerkin method with an emphasis on high-frequency wave propagation problems”. PhD thesis. The University of Texas at Austin, 2019.
- [27] S. Petrides and L. Demkowicz. “An adaptive DPG method for high frequency time-harmonic wave propagation problems”. In: *Comput. Math. Appl.* 74.8 (2017), pp. 1999–2017.
- [28] A. Vaziri, B. Keith, and L. Demkowicz. “On perfectly matched layers for discontinuous Petrov-Galerkin methods”. In: *Comput. Mech.* 63.6 (2019), pp. 1131–1145.