

A Refinement-by-Superposition Approach to Fully Anisotropic hp-Refinement for Improved Efficiency in CEM

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

29-09-2021 / 05-10-2021

CITATION

Corrado, Jeremiah; Harmon, Jake; Notaros, Branislav (2021): A Refinement-by-Superposition Approach to Fully Anisotropic hp-Refinement for Improved Efficiency in CEM. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.16695163.v1>

DOI

[10.36227/techrxiv.16695163.v1](https://doi.org/10.36227/techrxiv.16695163.v1)

A Refinement-by-Superposition Approach to Fully Anisotropic hp -Refinement for Improved Efficiency in CEM

Jeremiah Corrado, Jake J. Harmon, *Graduate Student Member, IEEE*, and Branislav M. Notaroš, *Fellow, IEEE*

Abstract—We present an application of fully anisotropic hp -adaptivity over quadrilateral meshes for $H(\text{curl})$ -conforming discretizations in Computational Electromagnetics (CEM). Traditionally, anisotropic h -adaptivity has been difficult to implement under the constraints of the Continuous Galerkin Formulation; however, Refinement-by-Superposition (RBS) facilitates anisotropic mesh adaptivity with great ease. We present a general discussion of the theoretical considerations involved with implementing fully anisotropic hp -refinement, as well as an in-depth discussion of the practical considerations for 2-D FEM. Moreover, to demonstrate the benefits of both anisotropic h - and p -refinement, we study the 2-D Maxwell eigenvalue problem as a test case. The numerical results indicate that fully anisotropic refinement can provide significant gains in efficiency, even in the presence of singular behavior, substantially reducing the number of degrees of freedom required for the same accuracy with isotropic hp -refinement. This serves to bolster the relevance of RBS and full hp -adaptivity to a wide array of academic and industrial applications in CEM.

Index Terms—computational electromagnetics, continuous Galerkin, finite element method, higher order methods, hp -refinement, anisotropic h -refinement, anisotropic p -refinement, multi-level, refinement-by-superposition.

I. INTRODUCTION

FULLY anisotropic hp -refinement over quadrilateral and hexahedral discretizations, an under-explored paradigm in Computational Electromagnetics (CEM) and Finite Element Methods (FEM), facilitates a significant enhancement in the tuning of discretizations for accurate *and* efficient simulations. However, given the difficulty of implementing isotropic h -adaptivity over quadrilateral or hexahedral cells under the constraints of a Continuous Galerkin Formulation, little work has been done to extend popular methodologies to support anisotropic h -adaptivity. That is not to say that the necessary theoretical foundations are absent [1], [2], rather, the requisite implementation complexity has impeded adoption of these techniques. The implementation difficulties can be avoided entirely by employing triangular or tetrahedral discretizations [3] or a discontinuous Galerkin Formulation [4]; however, non-rectangular cells are sub optimal for the linear independence of unit vectors in many CEM applications, and discontinuous Galerkin formulations do not maximize per-DoF efficiency. Therefore, we explore a far less burdensome approach to

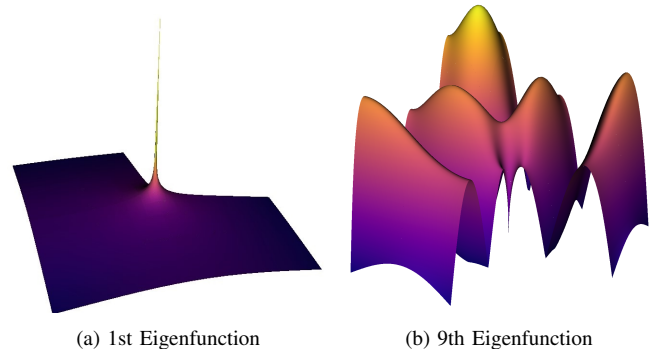


Fig. 1. Two solutions to the Maxwell Eigenvalue Problem on a unit size L-Shaped waveguide. Both solutions are shown as the magnitude of the Electric Field and exhibit sharp behavior around the re-entrant corner.

anisotropic hp -adaptivity with quadrilateral cells under a Continuous Galerkin formulation by leveraging a Refinement-by-Superposition (RBS) method.

Previous work has demonstrated the RBS method's ability to produce state-of-the-art exponential convergence on the 2-D Maxwell eigenvalue problem, even in the presence of singular or non-smooth behavior [5], of the form depicted in Fig.1. We note that without isolation of the irregular solution behavior through suitable h -refinements, p -refinement alone is insufficient to achieve exponential convergence. Thus, any computational method that aims to converge quickly regardless of non-idealities in the solution behavior must also implement some form of local h -adaptivity in concert with p -refinement [6]–[8].

The distinguishing feature of the RBS method as an approach to h -refinement is its simple formulation and ease of implementation. This stands in stark contrast to more mainstream h -adaptivity methods, such as the constrained nodes approach, which necessitate sophisticated algorithms to treat hanging-nodes. As described in [9], the RBS approach enforces continuity requirements by construction and is also able to handle arbitrary degrees of mesh irregularity (edges or faces can have any number of hanging nodes) without any special treatment.

Additionally, RBS not only supports anisotropic h -refinement without significant theoretical alterations, but there are also very few practical difficulties associated with its implementation. This is not the case for other, more traditional approaches to h -refinements over quadrilateral or hexahedral cells. Although excellent results can be achieved for applica-

Jeremiah Corrado, Jake J. Harmon, and Branislav M. Notaroš are with the department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523-1373 USA (e-mail: jccorrado@rams.colostate.edu, jake.harmon@ieee.org, branislav.notaros@colostate.edu).

tions in CEM with isotropic RBS hp -adaptivity [5], we aim to show the further benefits of introducing anisotropy in both h and p .

The remainder of this paper is organized as follows. Section II gives a short overview of the fully isotropic formulation of the RBS method. This is only intended to set the stage for the introduction of anisotropic adaptivity. More detailed descriptions of RBS can be found in [5], [9], [10]. Section III explores some practical considerations and implementation details for introducing anisotropy to the RBS method. A specific framework is discussed for the 2-D case; however, it generalizes trivially to 3-D. Finally, Section IV presents an experimental analysis of the method using the 2-D Maxwell Eigenvalue problem as a benchmark. The previously mentioned mesh with singular and non-smooth eigenpairs is used to show the efficacy of the method in situations where hp -refinement is required to achieve exponential convergence. The results show that anisotropic hp -refinements can achieve the same accuracy as fully isotropic refinements with significantly fewer Degrees of Freedom (DoFs), illustrating the usefulness of the method to other challenging problems in CEM.

II. OVERVIEW OF THE REFINEMENT-BY-SUPERPOSITION METHOD

The RBS method is an alternative approach to h -refinement that is well suited for a wide variety of computational methods. We focus here on discretizations that employ $H(\text{curl})$ - or $H(\text{div})$ -conforming hierarchical basis over quadrilateral or hexahedral cells (such as those in [11]). We also reference the shape-function classification and coordinate system described in [12].

In contrast to the more typical Refinement-by-Replacement (RBR) method where h -refinements decompose a cell into a set of smaller cells, RBS superimposes a set of “child cells” over a “parent cell” without removing the parent cell from the mesh. This subtle change obviates the need for special treatment of hanging nodes as local h -refinements introduce new nodes on a separate refinement layer and leave the parent edges completely intact.

Alternatively, the direct treatment of hanging nodes calls for more complex enforcement of continuity. A common RBR-based strategy is to introduce constrained nodes into the mesh. These nodes add DoFs to the system which are not actually free (they do not contribute entropy to the system) but are carefully constructed such that boundary conditions are satisfied. As such, local h -refinements are allowed, but per-DoF efficiency is reduced and implementation complexity is high.

The RBS method makes a different trade-off between implementation difficulty and resource requirements. By leaving the parent cell in the system, and allowing its DoFs to remain active as necessary, continuity is enforced by construction. The nature of this approach therefore reduces the implementation complexity at the expense of reduced matrix sparsity due to the increased interactions between refinement layers. These denser systems can require more memory and compute to solve; however, the lower implementation complexity often outweighs such a cost.

The following algorithm enforces continuity between neighboring cells and linear independence between refinement layers (additional details and illustrations related to this process in the context of isotropic refinements may be found in [5]):

- 1) Iterate over each cell and enumerate all possible shape functions based on the cell’s expansion orders, then associate them with the relevant geometric components (cell, face, node, or edge)
 - For $H(\text{curl})$, associate the shape functions based on the non-zero tangential components
 - For $H(\text{div})$, associate the shape functions based on the non-zero normal components
- 2) Iterate over each non-cell geometric component (faces, nodes, and edges)
 - For each shape function associated with that component, search for matching shape functions on neighboring cell(s)
 - If a match is found, designate it as a new DoF in the connected system
- 3) Iterate over each non-cell geometric component a second time
 - If the component has descendants with active DoFs, deactivate its DoFs (ex: if an edge has two direct descendant edges which possess sets of matched shape functions, deactivate the edges DoFs)
- 4) Iterate over each cell
 - If the cell has descendants, leave its cell-type shape functions inactive; otherwise, add them to the connected system as new DoFs

The above procedure is summarized more succinctly by the following remarks:

- Whenever possible, only the DoFs associated with the most- h -refined geometric components within an ancestry tree should be active
- If a node, edge, or face does not have an equally h -refined neighbor with which to match shape functions, then the responsibility falls on the nearest ancestor component (with a valid neighbor) to enforce continuity

Fig. 2 shows an example of an RBS mesh with the active geometric components annotated. This serves as visual representation of the above algorithm. Some anisotropic h - and p -refinements are also included to demonstrate that they are fully supported by the RBS framework. (For this figure and others, the vertical spacing simply represents the addition of refinement layers and has no physical significance to the geometry of the mesh.)

III. INTRODUCTION OF ANISOTROPIC h -REFINEMENTS TO RBS: PRACTICAL CONSIDERATIONS AND IMPLEMENTATION DETAILS

A. Implementation Goals

An anisotropic h -refinement, i.e., a directional h -refinement, is advantageous where the intensity of non-smooth or singular behavior is also directionally dependent or is confined to one small region of a cell. In such cases, isotropic h -refinement




Fig. 2. Example of a mixed-order 2-D Mesh constructed using the RBS framework. Nodes and Edges with active shape-functions are highlighted in Purple (those shared across refinement layers are highlighted on both layers). Cells that have active Edge-Type or Node-Type shape functions with inactive Cell-Type functions are colored around the edges and transparent in the center. Cells without active shape functions are completely transparent.

may introduce more DoFs than necessary to improve solution accuracy. Anisotropic h -adaptivity can therefore contribute to improved efficiency or even faster convergence rates by introducing new unknowns in a more frugal manner [4]. Likewise, anisotropic p -adaptivity permits increasing a cell’s polynomial expansion order in only one direction to drive more efficient improvements to the solution accuracy.

The methodology described in section II imposes few limitations on the shape or size of the superimposed cells. Previous works (such as [5]) have limited h -refinements to a simple 4-cell isotropic superposition; however, any set of child cells is permissible so long as the following conditions are met:

- 1) The child cells cover the entirety of the parent cell
- 2) The child cells do not extend into neighboring cells or beyond the boundary of the mesh
- 3) No internal hanging nodes are introduced among the child cells S.T. internal continuity is enforced naturally

For example, a 9-Cell isotropic superposition is equally as valid as a 4-Cell superposition, as the parent cell uses the same mechanism to enforce continuity in either case.

One can imagine taking advantage of the wide variety of anisotropic h -refinements allowed by the RBS method; however, it is useful to impose a few practical limitations to generate a simple implementation with maximum usability. In other words, given so few limitations on the shape, size, number, and orientation of the child cells, it is difficult to construct a simple h -refinement scheme that is easily targeted by adaptive methods (such as those described in [13]) or by human users while capturing all possible h -refinements. A practical goal for any anisotropic h -refinement implementation is to strike a balance between feature-richness and simplicity.

B. Implementation Details for 2-D FEM

For a specific demonstration of the procedure, we study an anisotropic RBS implementation designed for 2-D FEM which aims to illustrate the benefits of anisotropy while minimizing




Fig. 3. The three types of h -refinement used in our 2-D anisotropic h -refinement implementation; all illustrated as a refinement of a single square cell

implementation overhead. This framework generalizes to 3-D trivially. The three types of h -refinements considered here (named T-, U-, and V-Type) are shown in Fig. 3. T-Type refinement is identical to the isotropic h -refinement used in [5]. It is also equivalent to the successive application of a U-Type refinement and two V-Type refinements (or the inverse); however it is implemented directly for the sake of simplicity.

The two anisotropic h -refinements, U-Type and V-Type, superimpose only two new cells over the parent cell improving the u -directed or v -directed resolutions respectively, where the coordinates u and v correspond to those of the reference cell, while leaving the opposite direction unaffected. We also designate here that horizontal edges (those superimposed with a new node during a U-Type refinement) are called u -directed edges, and vertical edges (those superimposed with a new node during a V-Type refinement) are called v -directed edges.

Within this framework, it is useful for each cell to keep track of its “ h -refinement level” in the u and v directions. In other words, each cell must know how many U-Type and V-Type refinements were applied over the history of its construction from the base layer. When a cell is h -refined, the refinement levels of the child cells are a function of the parent’s refinement levels and the refinement type:

- T-type:

$$(u_{child}, v_{child}) = (u_{parent} + 1, v_{parent} + 1)$$

- U-type:

$$(u_{child}, v_{child}) = (u_{parent} + 1, v_{parent})$$

- V-type:

$$(u_{child}, v_{child}) = (u_{parent}, v_{parent} + 1)$$

With this information linked to each cell, it becomes straightforward to determine whether two adjacent cells can match edge-type shape functions or if four adjacent cells can match node-type shape functions. This is explored first by example, using the mesh configurations in Fig. 4, then a more explicit formulation is given.

The mesh in Fig.4(a) has no edges shared between refinement layers, so shape function matching is relatively simple and resembles the isotropic case. Here, the children of Cell 0 and Cell 1 have h -refinement levels of (1,1) and (0,1) respectively. Cells 3 and 6 share a v -directed edge, and their v -directed h -refinement levels match, therefore they can match edge-type shape functions. The same is true for Cells 5 and 7. As explained in Section II, these edges are more h -refined than

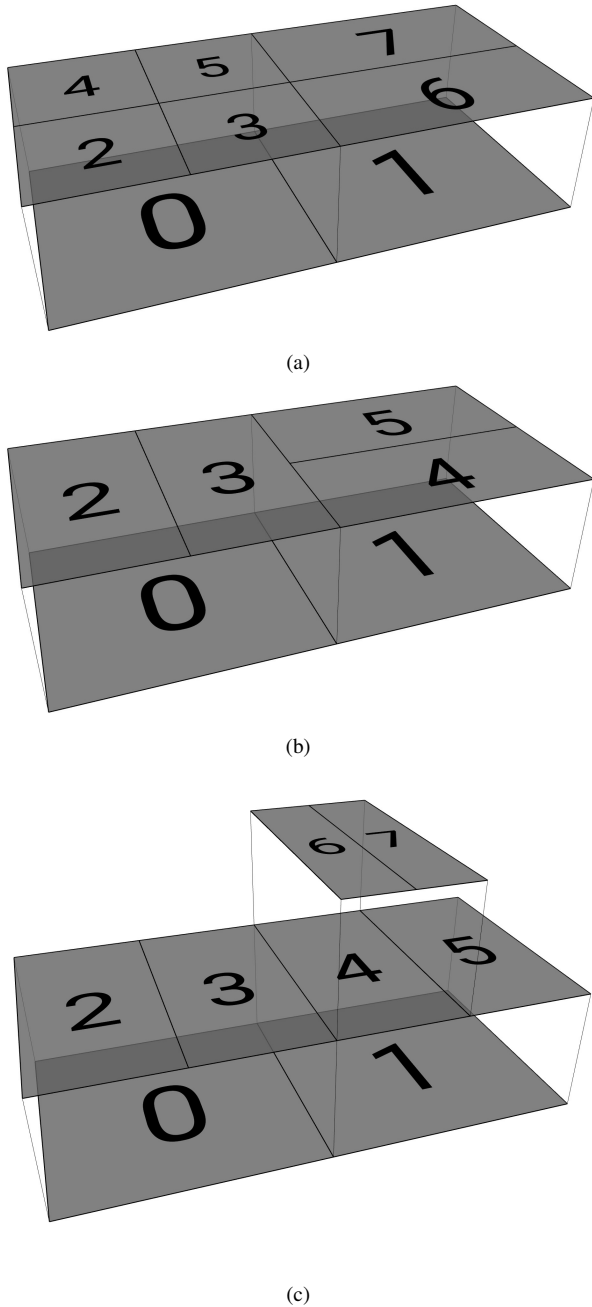


Fig. 4. Examples of Edge-Type shape function matching cases introduced with anisotropic h -refinement. These cases all focus on continuity enforcement over the v -directed edge shared by Cells 0 and 1. (a) A mesh with a single isotropic h -refinement and one V-Type refinement. (b) A Mesh with one U-Type and one V-Type h -refinement. Outside the context of RBS, this would introduce a hanging-node along the central edge. (c) A 3-Layer Mesh with only U-Type h -refinements illustrating the necessity for a cell-ranking system.

their parent edge, and they can both support shape-functions, therefore edge-type functions on Cells 0 and 1 are unnecessary to enforce continuity and are left inactive.

An important feature of the U- and V-Type h -refinements is that only 4 new edges are constructed along the border of the parent cell (as opposed to the 8 generated by a T-Type refinement). The other two edges are shared with the parent cell, which can introduce some ambiguity with shape function matching. In Fig.4(b), the central edge is shared by Cells 0, 1, and 3. As such, the edge-type shape functions on Cell 1 can match with those on Cells 0 or 3 which have h -refinement levels of (0, 0) and (1, 0) respectively. In keeping with the RBS concept that the most h -refined cells should be responsible for maintaining continuity whenever possible, Cells 3 and 1 will match edge-type shape functions and those on Cell 0 will remain inactive.

This logic extends further when multiple cells of the same relevant h -refinement level share a single edge. For example in Fig.4(c), Cells 0 and 3 share many possible matching combinations with Cells 1, 4, and 6, which all have v -directed h -refinement levels of zero. To select a single matching pair, the choices on both sides of the edge are ranked by their u -directed refinement level, and the highest-ranked cells on either side are chosen. As such, shape-functions on Cells 3 and 6 are used to enforce continuity over the central edge.

The logic in the above examples is captured by the following remarks:

- All edges in the mesh must maintain two lists of adjacent cells—one for each side—which are ranked first by the cells refinement-level associated with the edges own direction, then by its refinement level associated with the opposite direction
- The highest-ranked cell on each side of the edge is used to construct a valid pair if and only if the edge itself is needed to enforce continuity (which is determined by the broader RBS procedure given in Section II)

An extension of this logic to include four cells is used to match node-type shape functions. In such cases, cells are ranked by the four relevant edges in the same manner, and each edge must "agree" with the neighboring two edges on which cells to select.

No major changes to the isotropic RBS method, described in Section II, are required. The algorithm used to match shape functions on neighboring cells and enforce continuity requirements remains essentially the same, except for the implementation of Step 2, which is amended to include the slightly more complex matching procedure described above.

IV. NUMERICAL RESULTS

A. Benchmark Problem Description

We now establish the advantages of introducing hp -anisotropy to the RBS method by solving the Maxwell Eigenvalue problem on an L-shaped waveguide terminated by Dirichlet boundary conditions proposed by [14]. Additionally, we only consider TE propagation modes by asserting that the solution is purely transversal. The Maxwell Eigenvalue

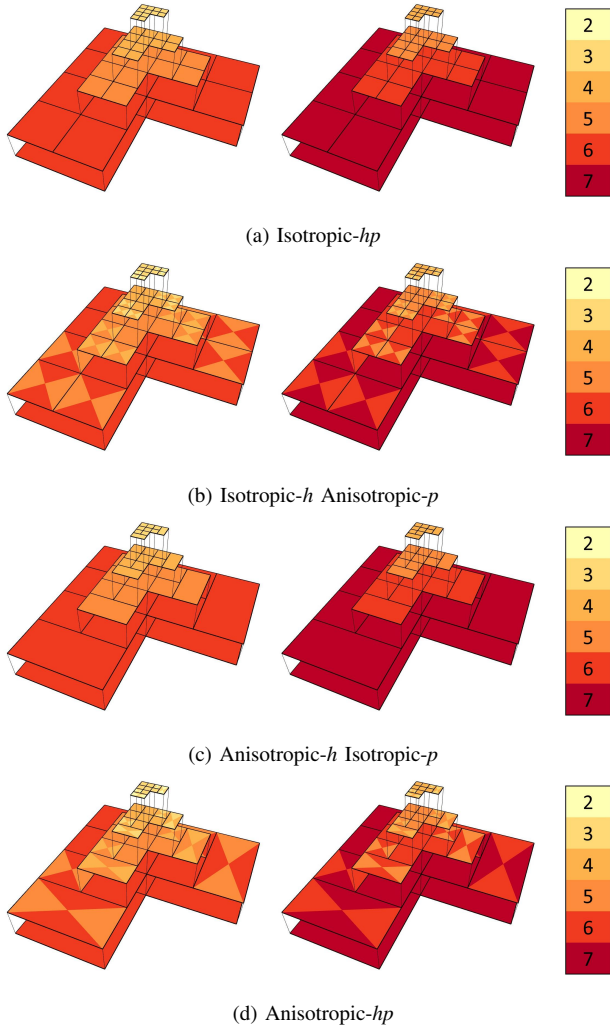


Fig. 5. Illustration of the *a priori* refinement strategies on a five-layer mesh ($k = 4$). Figures in the left column show the initial mesh and Figures in the right column show the same mesh after one refinement iteration. In all cases, the refinement (from left to right) increments u - and v -directed expansion orders by 1 on all cells. The color scales on the right designate expansion orders. (p -anisotropy is shown on the dual-colored cells S.T. the expansion order associated with a given edge is represented by the color touching that edge.)

problem is formulated as follows:

Find $U = \{\mathbf{u}_{hp}, \lambda_{hp}\} \in B_{hp} \times \mathbb{R}$ such that

$$a(\mathbf{u}_{hp}, \phi_{hp}) = \lambda_{hp} m(\mathbf{u}_{hp}, \phi_{hp}) \quad \forall \phi_{hp} \in B_{hp}, \quad (1)$$

for $B_{hp} \subset H(\text{curl}; \Omega)$, where $m(\mathbf{u}_{hp}, \phi_{hp}) = \langle \mathbf{u}_{hp}, \phi_{hp} \rangle$, and $a(\mathbf{u}_{hp}, \phi_{hp}) = \langle \nabla_t \times \mathbf{u}_{hp}, \nabla_t \times \phi_{hp} \rangle$.

The convergence behavior is evaluated for the four combinations of (An)Isotropic- h and (An)Isotropic- p refinement. The solutions express singular and non-smooth behavior making it an ideal benchmark for evaluating the effectiveness of hp -methods. The two eigenfunctions in question are the first—shown in Fig. 1(a)—which contains a singularity along the re-entrant corner, and the ninth—shown in Fig. 1(b)—which has non-smooth behavior on the re-entrant corner and more complex behavior elsewhere.

B. Refinement Strategies

A set of four *a priori* refinement strategies are used to illustrate the difference in convergence behavior with and without anisotropy. An example of each is shown in Fig. 5. Although the first and ninth eigenfunctions are quite different, they share a key feature of sharp behavior around the re-entrant corner, and thus the same strategies are sufficient to produce exponential convergence on both eigenpairs.

The first strategy, shown in Fig. 5(a), is fully isotropic and is identical to the strategy used in [5] to illustrate the method's capacity for exponential convergence. Here, a given number of layers: k , describes the number of refinement layers that were added to the base discretization. Each new layer is generated with a T-Type refinement of the three cells surrounding the re-entrant corner. To construct the starting discretization, the cells on the first layer are assigned an expansion order of $k + 2$ (ensuring a minimum order of 3 on the top layer). Subsequent layers have an expansion order 1 less than the previous layer.

Anisotropic p -refinements are introduced by reducing the expansion order of a few select cells by 1 in only one direction. The specific pattern of anisotropic expansion orders is shown in Fig. 5(b). The targeted cells and expansion directions were chosen based on the areas of the mesh where the u -directed electric field is much more intense than the v -directed electric field or vice versa.

Anisotropic h -refinements are introduced by replacing the outer two T-Type refinements with a U/V-Type refinement followed by a V/U-Type refinement of the resultant cell closest to the re-entrant corner. As shown in Fig. 5(c), the mesh remains coarser over the regions that are farther from the re-entrant corner. As such, new DoFs and smaller-scale cells are only introduced around the re-entrant corner where they are most needed to capture the sharp solution behavior.

Finally, a strategy with anisotropic hp -refinements is shown in Fig. 5(d). This is simply a combination of the previous two strategies and constitutes the DoF savings from both.

C. Results and Discussion

Convergence behavior is evaluated as follows: for each of the strategies shown in Fig. 5, an initial mesh is generated based on a given value of k . Then, accuracy data is collected by repeatedly p -refining each cell in the mesh and solving the Maxwell eigenvalue problem described above. The relative error at each refinement iteration is computed as the absolute difference between the solution eigenvalue and an accurate numerical benchmark [13]. Results for the first eigenvalue are shown in Fig. 6 and results for the ninth are shown in Fig. 7. Both figures include results for $k=4$ and $k=8$. A maximum expansion order of 14 is employed across experiments, meaning that the $k=4$ plots contain more refinement iterations (as those experiments start with a lower maximum expansion order).

In all cases, the relative error is nearly identical across the four refinement strategies for any given iteration; however, the number of DoFs varies significantly. Universally, the fully anisotropic refinement pattern performs the most efficiently (requiring fewest NDoFs to achieve a given accuracy), then the Anisotropic- h Isotropic- p pattern, then the Isotropic- h

TABLE I
Efficiency Gains of anisotropic strategies at each refinement iteration for $k=4$ on the 1st eigenvalue

Refinement Iteration	1	2	3	4	5	6	7	8	9
<i>Isotropic-h Anisotropic-p</i>	1.1463	1.1251	1.1062	1.0906	1.0793	1.0704	1.0633	1.0575	1.0526
<i>Anisotropic-h Isotropic-p</i>	1.2931	1.2866	1.2825	1.2794	1.2769	1.2750	1.2733	1.2720	1.2709
<i>Anisotropic-hp</i>	1.4462	1.4195	1.3951	1.3757	1.3611	1.3497	1.3405	1.3330	1.3267

TABLE II
Efficiency Gains of anisotropic strategies at each refinement iteration for $k=8$ on the 1st eigenvalue

Refinement Iteration	1	2	3	4	5
<i>Isotropic-h Anisotropic-p</i>	1.0881	1.0855	1.0774	1.0697	1.0633
<i>Anisotropic-h Isotropic-p</i>	1.2760	1.2901	1.2903	1.2886	1.2872
<i>Anisotropic-hp</i>	1.3673	1.3824	1.3743	1.3643	1.3559

Anisotropic- p pattern, and finally the fully isotropic pattern is the least efficient. Additionally, it is clear that the anisotropic h -refinements constitute a larger improvement in efficiency than the anisotropic p -refinements.

The efficiency improvements obtained by each of the three anisotropic strategies are given numerically in Tables 1-4, with the most significant improvements from each iteration highlighted in bold. Results are given for $k=4$ and $k=8$ on both eigenpairs. The values shown are the ratio of the anisotropic strategies per-DoF efficiency to the isotropic per-DoF efficiency for the same iteration. In other words, these values express how many times more efficient the anisotropic solution was than the isotropic solution for some refinement iteration. Or more explicitly, the efficiency gains are expressed by the following equation, where η denotes per-DoF efficiency, ε denotes the accuracy (or inverse of the relative error), and N denotes the NDoFs:

$$\frac{\eta}{\eta_{iso}} = \left(\frac{\varepsilon}{\varepsilon_{iso}} \right) * \left(\frac{N_{iso}}{N} \right) \quad (2)$$

For most iterations, the achieved accuracy is approximately equal across strategies, meaning that the above equation is well approximated by the ratio of the number of DoFs. Thus, for the experiments shown here, efficiency gains are equivalent to the DoF savings introduced by the anisotropic strategy.

Across all experiments, anisotropic h -refinements present the most significant gains in efficiency, while anisotropic p -refinements present a smaller, but still useful, boost in efficiency. Notably, the efficiency gains associated with the first refinement iterations for the 9th eigenvalue exhibit some erratic behavior (which can also be deduced from Fig. 7). In all other cases, the fully anisotropic refinement strategy yielded the largest gains in efficiency, closely followed by the anisotropic- h isotropic- p strategy, then followed by the isotropic- h anisotropic- p strategy. These insights align with the data shown in Figures 6 and 7, and indicate that the combined application of anisotropic h - and p -refinements is a highly effective strategy to increase the per-DoF efficiency of an FEM simulation.

TABLE III
Efficiency Gains of anisotropic strategies at each refinement iteration for $k=4$ on the 9th eigenvalue

Refinement Iteration	1	2	3	4	5	6	7	8	9
<i>Isotropic-h Anisotropic-p</i>	1.3281	1.1547	1.1061	1.0897	1.0788	1.0701	1.0631	1.0573	1.0525
<i>Anisotropic-h Isotropic-p</i>	1.2006	1.2612	1.2840	1.2781	1.2768	1.2748	1.2733	1.2719	1.2708
<i>Anisotropic-hp</i>	2.1333	1.4272	1.3971	1.3732	1.3604	1.3492	1.3402	1.3327	1.3265

TABLE IV
Efficiency Gains of anisotropic strategies at each refinement iteration for $k=8$ on the 9th eigenvalue

Refinement Iteration	1	2	3	4	5
<i>Isotropic-h Anisotropic-p</i>	0.6199	1.0788	1.0752	1.0683	1.0627
<i>Anisotropic-h Isotropic-p</i>	0.8081	1.1713	1.2875	1.2862	1.2854
<i>Anisotropic-hp</i>	0.6122	1.2517	1.3698	1.3600	1.3526

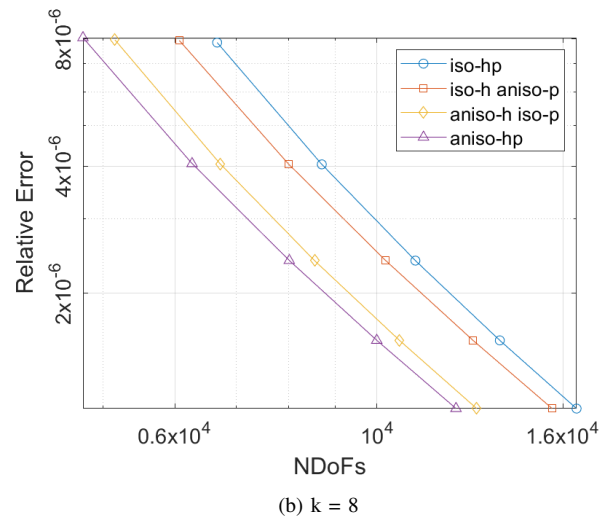



Fig. 6. Convergence behavior of four refinement strategies shown on a log-log scale for the 1st eigenvalue. Results are given for $k=4$ and $k=8$.




Fig. 7. Convergence behavior of four refinement strategies shown on a log-log scale for the 9th eigenvalue. Results are given for $k=4$ and $k=8$.

V. CONCLUSION

We have demonstrated an extension of the capabilities of the Refinement-by-Superposition approach to full hp -adaptivity by including anisotropic h - and p -refinements. Using the 2-D Maxwell eigenvalue problem over an L-shaped waveguide as a benchmark; we showed that anisotropic refinements, particularly in h , present a significant advantage in computational efficiency. This is consistent with a theoretical understanding, as anisotropic h -refinements permit the construction of mesh configurations that capture small-scale behavior without introducing redundant DoFs into the system. Due to the challenging nature of the benchmark problem, we are confident that these efficiency improvements will be broadly applicable to other difficult problems in CEM.

This work also serves to further illustrate the benefits of the RBS method's low implementation complexity. Other approaches to h -refinement on quadrilateral or hexahedral discretizations involve complex frameworks for boundary condition enforcement which do not easily lend themselves to the implementation of anisotropy. The RBS method, by contrast, imposes few limitations on the configuration of the

superimposed cells, and thus supports h -anisotropy essentially for free. A conservative implementation of h -anisotropy for the 2-D case was described in detail and can readily be expanded to include more refinement types and can also be generalized for 3D FEM.

REFERENCES

- [1] W. Rachowicz, "An anisotropic h-type mesh-refinement strategy," *Computer Methods in Applied Mechanics and Engineering*, vol. 109, no. 1, pp. 169–181, 1993. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/004578259390231L>
- [2] D. Schötzau, C. Schwab, and R. Stenberg, "Mixed hp-fem on anisotropic meshes ii: Hanging nodes and tensor products of boundary layer meshes," *Numerische Mathematik*, vol. 83, no. 4, pp. 667–697, 1999.
- [3] V. Dolejší, G. May, F. Roskovec, and P. Solin, "Anisotropic hp-mesh optimization technique based on the continuous mesh and error models," *Computers and mathematics with applications (1987)*, vol. 74, no. 1, pp. 45–63, 2017.
- [4] S. Giani and M. Seaid, "Multi-hp adaptive discontinuous galerkin methods for simplified pn approximations of 3d radiative transfer in non-gray media," *Applied Numerical Mathematics*, vol. 150, pp. 252–273, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168927419302570>
- [5] J. J. Harmon, J. Corrado, and B. M. Notaroš, "A refinement-by-superposition hp-method for $h(\text{curl})$ - and $h(\text{div})$ -conforming discretizations," *TechRxiv*, Jun 2021.
- [6] W. Gui and I. Babuška, "The h, p and h-p versions of the finite element method in 1 dimension, part i," *Numerische Mathematik*, vol. 49, no. 6, pp. 577–612, Nov 1986.
- [7] W. Gui and I. Babuška, "The h, p and h-p versions of the finite element method in 1 dimension, part ii," *Numerische Mathematik*, vol. 49, no. 6, pp. 613–657, Nov 1986.
- [8] W. Gui and I. Babuška, "The h, p and h-p versions of the finite element method in 1 dimension, part iii," *Numerische Mathematik*, vol. 49, no. 6, pp. 659–683, Nov 1986.
- [9] N. Zander, T. Bog, S. Kollmannsberger, D. Schillinger, and E. Rank, "Multi-level hp-adaptivity: High-order mesh adaptivity without the difficulties of constraining hanging nodes," *Computational Mechanics*, vol. 55, no. 3, p. 499–517, 2015.
- [10] N. Zander, T. Bog, M. Elhaddad, F. F. S. Kollmannsberger, and E. Rank, "The multi-level hp-method for three-dimensional problems: Dynamically changing high-order mesh refinement with arbitrary hanging nodes," *Computer Methods in Applied Mechanics and Engineering*, vol. 310, p. 252–277, 2016.
- [11] M. M. Kostić and B. M. Kolundžija, "Maximally orthogonalized higher order bases over generalized wires, quadrilaterals, and hexahedra," *IEEE Transactions on Antennas and Propagation*, vol. 61, no. 6, pp. 3135–3148, 2013.
- [12] M. M. Ilic, A. Z. Ilic, and B. M. Notaros, "Efficient large-domain 2-D FEM solution of arbitrary waveguides using p-refinement on generalized quadrilaterals," *IEEE Transactions on Microwave Theory and Techniques*, vol. 53, no. 4, pp. 1377–1383, 2005.
- [13] J. J. Harmon and B. M. Notaroš, "Adaptive hp-refinement for time-harmonic Maxwell eigenvalue problems: Method and benchmarks," *In Review*, 2021.
- [14] M. Dauge, "Benchmark computations for Maxwell equations for the approximation of highly singular solutions," <https://perso.univ-rennes1.fr/monique.dauge/benchmax.html>.