

A Rationale for p -refinement with the Vector Helmholtz Equation and Two Dimensional Vector Finite Elements

R. Stephen Preissig
Texas Instruments Incorporated
5707 Peachtree Parkway
Norcross, GA 30092

Andrew F. Peterson
School of ECE
Georgia Institute of Technology
Atlanta, GA 30332-0250
peterson@ee.gatech.edu

Abstract: A preliminary study of p -refinement with vector finite elements is reported. Results suggest that improved accuracy can be obtained from representations employing a mixture of polynomial orders instead of a uniform polynomial order. Results also suggest that it might be possible to jump directly from the local error in a $p=0$ expansion to a final representation employing 5 or more polynomial orders. In addition, a new set of hierarchical curl-conforming vector basis functions is proposed.

Introduction

Over the years, there have been many extensions and variations on the classical scalar finite element method. The recent introduction of vector finite elements (edge elements) has created the opportunity for an analogous development of the vector finite element method. One aspect of finite elements is the possibility of adaptive refinement of the finite element mesh, such as the h -refinement process where portions of the mesh are refined to achieve smaller cells and higher accuracy where required, and the p -refinement strategy, where the polynomial order of the representation is selectively increased throughout portions of the mesh [1-3]. Based on work in scalar finite element analysis, it is generally thought that improved convergence can be obtained with one of these refinement schemes or a mixture of h -refinement and p -refinement strategies.

This article considers the benefits of a p -refinement approach for vector finite elements. Texts such as those by Akin [1] and Zienkiewicz

and Taylor [2] discuss p -refinement for scalar finite element applications. The use of vector finite elements, however, is relatively new, and very little has been done to study adaptive methods. Salazar-Palma et al. have explored h -refinement with vector finite elements [3], and several commercial packages (Ansoft's High Frequency Structure Simulator, for one [4]) employ h -refinement with vector elements. Although p -refinement has not been widely considered for the vector case, most aspects of p -refinement for the vector formulation are similar to the scalar formulation. One major difference between the scalar and vector formulations is the basis set, which is correspondingly scalar or vector.

Basis functions

The present investigation considers two-dimensional formulations based on the curl-curl form of the vector Helmholtz equation. A space of vector basis functions suitable for use with the vector Helmholtz equation was introduced by Nedelec in 1980 [5]. These functions maintain tangential cell-to-cell continuity and are known as *curl conforming*. Nedelec's mixed-order basis reduces the number of null valued eigensolutions of the curl-curl operator, which are physically meaningless in a source-free region, and therefore improves computational efficiency [6]. Vector basis functions consistent with Nedelec's spaces have been developed [3, 7]. A goal of the present study was to investigate p -refinement approaches using functions from Nedelec's curl-conforming space.

In practice, p -refinement techniques often incorporate hierarchical basis sets, where the functions comprising order p contain all functions of lower polynomial order. In this manner, p -refinement can be carried out by simply adding a few additional functions during each pass rather than changing the entire set. Several hierarchical vector basis sets have been proposed. Webb and Forghani [8] proposed vector functions for tetrahedral cells that can easily be adapted to triangular cells. They actually only presented functions for the lowest two orders ($p=0$ and $p=1$ in our notation). Wang [9] proposed different vector basis functions based on orthogonal polynomials which may offer improved linear independence. Carrie and Webb [10] proposed a third variety of vector functions for triangles based on Jacobi polynomials. The preceding functions do not appear consistent with Nedelec's spaces. More recently, several authors proposed alternative sets for tetrahedral cells that do appear to be consistent with Nedelec's spaces [11-14]. For completeness, Appendix A summarizes a set of hierarchical functions for triangular cells ranging up to polynomial order 4.5.

Although a practical p -refinement implementation would employ hierarchical bases, for the present investigation we used the interpolatory vector bases of [7] rather than a hierarchical set. The triangular-cell curl-conforming bases of [7], at a given order, provide a representation that is equivalent to that of the hierarchical functions of Appendix A. (In infinite precision arithmetic, the results would be identical.) It is also apparent that a useful implementation of p -refinement would necessarily employ a large dynamic range of basis orders. In the following, results are presented based on a mixture of up to 5 different polynomial orders within the same finite element mesh.

An additional feature of a p -refinement technique is the definition of transition elements, to smooth transition between two regions of different polynomial order. The basis functions employed in this study [7] are each associated with an edge or a patch within the mesh. The patch functions are entirely local to one triangular cell, and contribute no tangential component to any of the cell edges. The tangential continuity of the expansion is ensured by using the same order basis for those functions interpolatory on both sides of a given edge and assigning the same

coefficients to these functions. Thus, these two edge functions may be considered a single edge function which spans both patches common to that edge. Consequently, a transition element is formed when the polynomial order of the functions associated with different edges of the same triangular patch differ. To transition from order p to order q , the basis functions associated with one edge of a triangle may be of order p , while those of another edge are of order q . (We generally constrain the method to no more than two different orders per patch.)

At any stage of the refinement process, a polynomial order is assigned to each edge within the mesh. (A patch order may also be defined as the average of the orders of each of the three edges associated with that patch, rounded up to the next integer.) The goal of the optimization scheme is therefore to determine the order of each edge for the optimal distribution of unknowns throughout the mesh. A more detailed presentation and derivation of the formulation used in this study is presented in reference [15].

Formulation

To investigate the potential advantages of p -refinement, we first wish to determine whether it is possible to obtain a better accuracy/efficiency trade-off using a mixture of basis orders throughout a mesh than it is with a single order. As a canonical problem of interest, we consider the two-dimensional resonant cavity application for the transverse electric (TE) polarization. The vector Helmholtz equation for this situation is

$$\nabla \times \nabla \times \mathbf{E} = k^2 \mathbf{E} \quad (1)$$

For a cavity with perfect electric (PEC) walls, the tangential component of the electric field must vanish on the walls. A weak formulation of the problem is constructed as delineated in Chapter 9 of [6], and we refer the readers to that text for the details. In the following, we consider two specific cavities to illustrate our investigation. The first geometry is a square cavity partially loaded with a dielectric slab with $\epsilon_r=10$. The second structure is a circular cavity with a septum to the center. These examples both possess an exact analytical solution against which to measure the error in a numerical result for the dominant resonant frequency. In addition, the modal

solutions of both of these problems have regions of high variation as well as regions of low variation (Figure 1). The modes of the dielectric-loaded square cavity have relatively high variation within the dielectric, while the circular cavity has a singularity at the tip of the septum, and therefore has a large variation in the vicinity of the tip.

The dielectric-loaded square cavity was studied with three meshes, labeled A, B, and C, having 37, 74 and 158 patches, respectively. The circular cavity was studied with three meshes, labeled D, E and F, having 36, 66 and 128 patches, respectively. In each case, the patch size was relatively uniform throughout the mesh, but the meshes are unstructured. Meshes A and D are shown in Figure 2. For the circular cavity, parabolic curvilinear cells were used along the boundary.

First phase of the study

In the first part of our investigation, we attempted to determine whether a mixture of polynomial orders throughout the mesh offered better accuracy for a given number of unknowns than a representation with a uniform polynomial order. To study this issue, we developed an iterative optimization algorithm that locally adjusted the polynomial order (both up and down) while holding the total number of unknowns fixed. The iterative optimization method used in this study was based on three assumptions:

- (1) The normal discontinuity in the \mathbf{D} -field at cell edges is proportional to the local error. (The normal discontinuity is zero in the true solution.)
- (2) It is optimal for the error to be uniformly distributed throughout the mesh as opposed to localized.
- (3) A localized increase in the number of unknowns will improve the localized solution, decreasing the localized error.

The optimization routine attempts to minimize the standard deviation of the normal discontinuities throughout the mesh for a given number of unknowns. The program first calculates the statistical quantities of interest for the initial distribution of unknowns and attempts to improve the uniformity by increasing the number of

unknowns in highly discontinuous regions and decreasing the number of unknowns in regions with relatively low discontinuities. It was found that the actual limits by which the “high” and “low” discontinuities are defined affect the convergence rate of the optimized solution but not the final solution (if they are picked within reason, of course). These quantities were picked somewhat arbitrarily and will not be discussed here.




Figure 1. The dominant modes in the square and circular cavities.

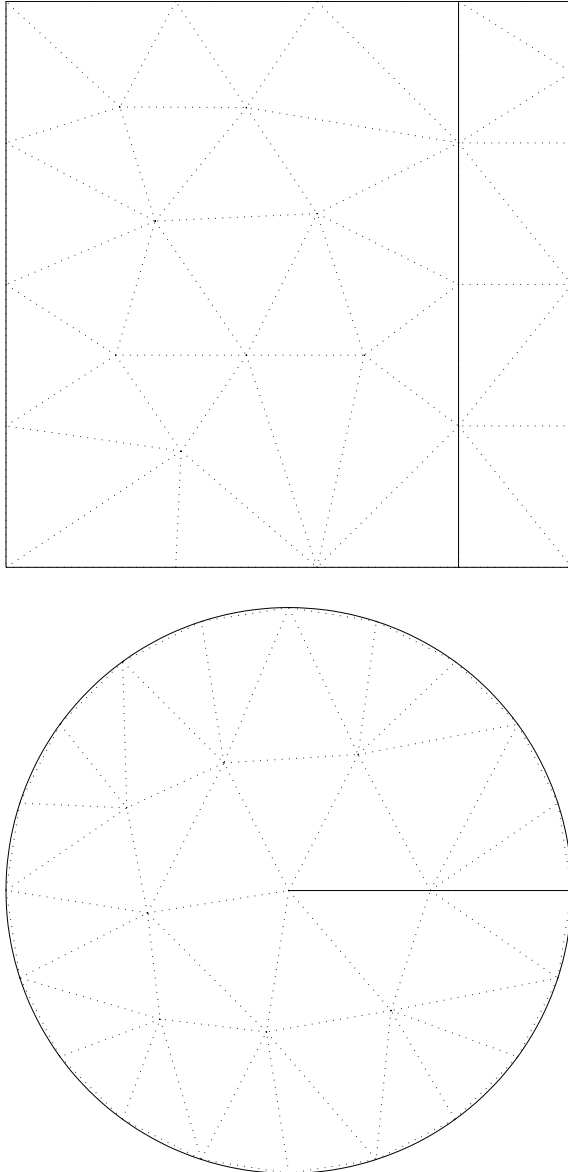


Figure 2. The coarsest meshes for the square and circular cavities.

The reallocation of unknowns is repeated until the program can no longer decrease the standard deviation of the discontinuities in the mesh. This final distribution of unknowns is then regarded as the optimal distribution of unknowns and the error associated with this distribution of unknowns (the error in the dominant resonant wavenumber) is used in a plot of error versus number of unknowns. There is, of course, no guarantee that the result is actually optimal.

Second phase of the study

The first phase of the study investigated whether a more accurate solution was possible with a fixed number of unknowns. In the second phase of the study, we attempted to determine if one could efficiently realize a nearly optimal distribution of unknowns, while avoiding the cumbersome optimization process used in phase 1.

As a first step in the approach, we solve the problem using a uniform zero-order representation throughout the mesh. We then determine the normal discontinuity (in the \mathbf{D} -field) produced at each edge in the mesh by that zero-order representation. The second step is to use that error distribution to immediately assign a “final” polynomial order to each edge in the mesh.

The ambiguity of such an approach is in how many orders and how many edges of each order to assign, since the process is constrained by the total number of unknowns desired. For this study, we cheated — we used the results of the earlier iterative optimization procedure to determine the number of orders and percentage of each order to use. In other words, if for a particular mesh the iterative method determined (after many passes) that 15% of the edges were assigned polynomial order 4, then 15% of the edges were assigned that order in the single step algorithm. However, in the single step procedure the edges chosen for refinement were selected based solely upon the extent of the zero-order normal discontinuity at that edge. Thus, the procedure is not a complete p-refinement algorithm at this point, and our specific approach is not practical for implementation. Data presented below suggest, however, that an efficient algorithm for distributing the unknowns among several orders is possible. The second phase of the study, while not a self contained single step p-refinement capability, does provide insight into the accuracy possible with relatively little computational effort.

Results

Figures 3-8 show the percent error in the dominant cavity resonant frequency determined by various methods versus number of unknowns. Each figure depicts the “homogeneous order solution,” the error obtained using a $p=0$

representation, a $p=1$ representation, etc., throughout the entire mesh.

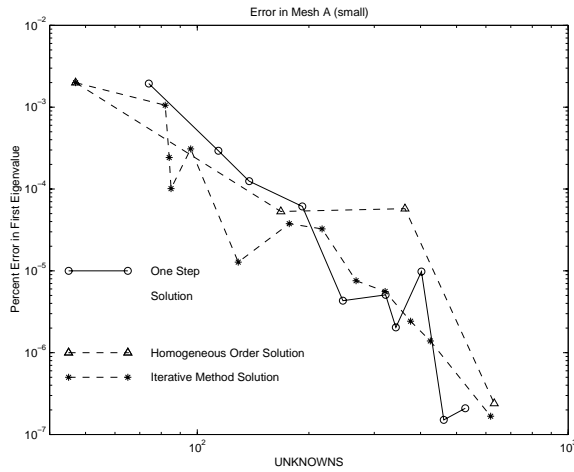


Figure 3. Error Versus Unknowns for Square Cavity, Mesh A, Iterative and One Step Methods.

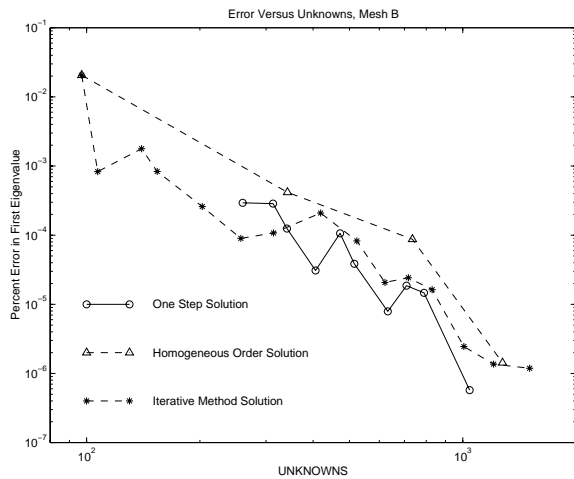


Figure 4. Error Versus Unknowns for Square Cavity, Mesh B, Iterative and One Step Methods.

Each figure also shows the error produced by the iterative optimization process, which employed a mixture of various polynomial orders throughout the mesh. The iterative process attempted to minimize the standard deviation of the normal discontinuity within the dominant mode. For most of the data, the iterative process produced better accuracy for a given number of unknowns (using a mixture of polynomial orders) than the homogeneous solution (a single polynomial order).

Figures 3-8 also show the “one step” solution obtained by jumping from the zero-order result directly to the final distribution of unknowns. The one step solution is usually as good (and occasionally better) than that produced by the gradual iterative process. This suggests that the local error associated with the zero-order solution is a meaningful predictor for the distribution of unknowns in the final result.

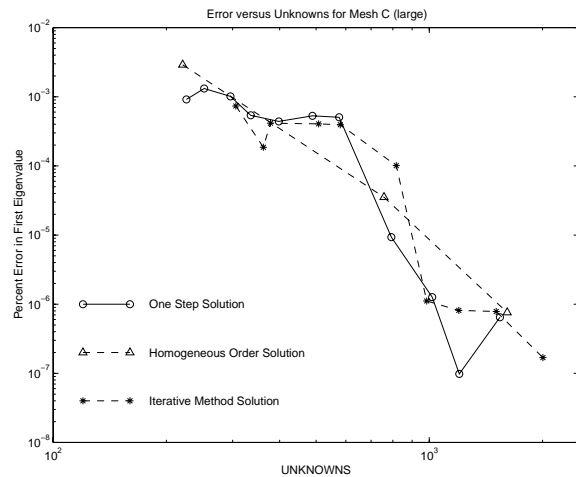


Figure 5. Error Versus Unknowns for Square Cavity, Mesh C, Iterative and One Step Methods.

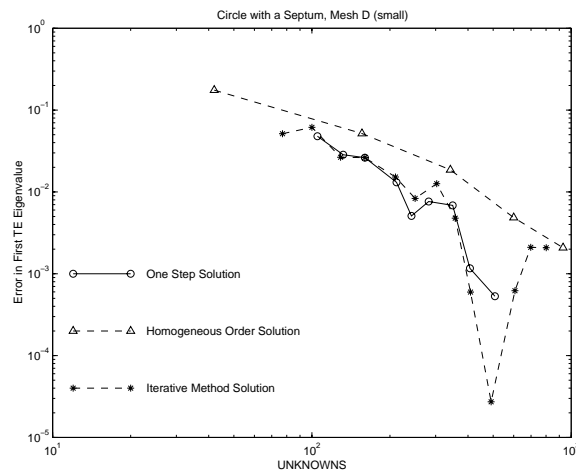


Figure 6. Error Versus Unknowns for Circular Cavity, Mesh D, Iterative and One Step Methods.

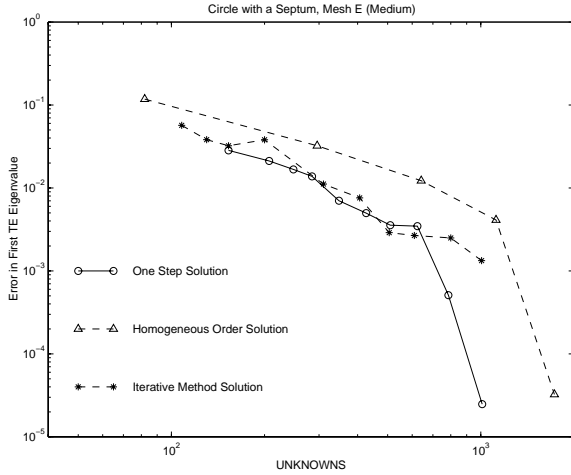


Figure 7. Error Versus Unknowns for Circular Cavity, Mesh E, Iterative and One Step Methods.

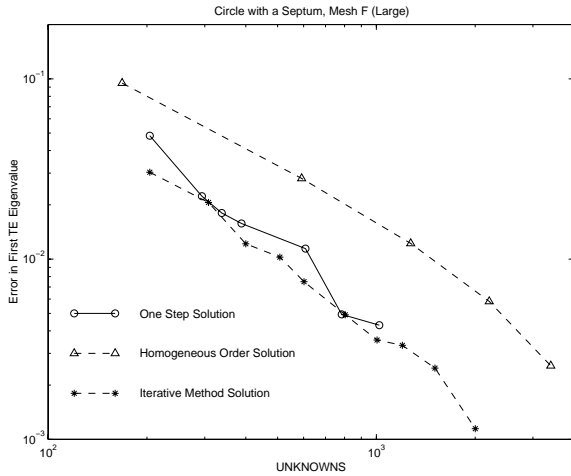


Figure 8. Error Versus Unknowns for Circular Cavity, Mesh F, Iterative and One Step Methods.

For Figures 3-5, representing the square cavity, the optimal distribution of unknowns tended to be that with a single order throughout the air-filled part of the cavity and a higher order throughout the dielectric part. The error curves tend to zigzag up and down as the number of unknowns is increased, due to the fact that a single order in the air filled part of the cavity and a higher order throughout the dielectric part can only be achieved for specific numbers of unknowns. When the algorithm is forced to optimize to a number of unknowns for which this distribution is not possible, it is less efficient. In contrast, the error curves in Figures 6-8 for the circular cavity

show a more uniform behavior. Figure 9 illustrates the distribution of degrees of freedom for mesh F, when optimized for 1000 unknowns.




Figure 9. Distribution of degrees of freedom for mesh F when optimized at 1000 unknowns.

As an initial step in developing a control algorithm for p-refinement, we present several plots showing the distribution of polynomial order associated with a given number of unknowns. Figures 10 and 11 show the transition profiles for the square dielectric-loaded cavity (Mesh B) and the circular cavity with a baffle (Mesh E), respectively. The transition profiles show the percentage of the total number of unknowns assigned to each polynomial order for a given number of unknowns. These data are produced using the iterative optimization process. The number of unknowns is normalized to the number of unknowns in the homogeneous zero order case. For example, a normalized number of unknowns of “3” corresponds to three times as many unknowns as in the zero order homogeneous case.

Figures 12 and 13 show the percentage of edges of each polynomial order in terms of a “transition point.” The local error (discontinuity in \mathbf{D} -field) level is organized into a list by edges; the transition point is the position in that list where the transition between orders is assigned. The unknowns in the one step method are assigned

directly by the ordered list of discontinuities. For example, mesh A has 37 edges. If it is determined that there are to be 30 zero order edges and 7 first order edges after refinement, the 0/1 normalized transition point would be $30 / 37 = 0.81$.




Figure 10. Transition Profile for Square Cavity, Mesh B.




Figure 11. Transition Profile for Circular Cavity, Mesh E.

Because the unknowns in the one step method are assigned directly by the ordered list of discontinuities, the normalized transition point also corresponds to the percentage of edges to which each order is assigned. The 1/2 transition point corresponds to the percentage of edges with

order less than two in the optimized distribution of unknowns. Thus, the difference between the 1/2 transition point and the 0/1 transition point is the percentage of edges with order one in the optimized distribution. Thus, the bar graphs show the percentage of edges assigned to each order.

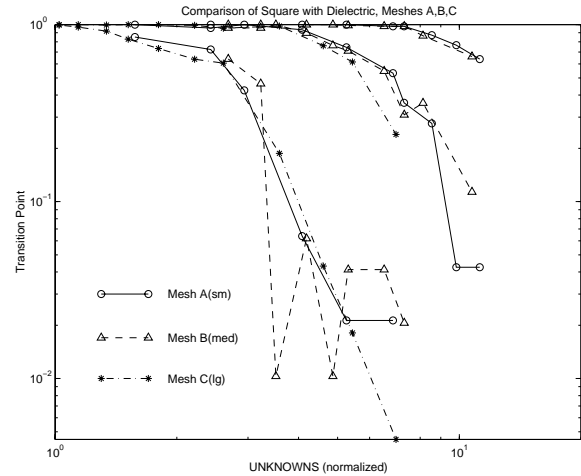


Figure 12. Comparison of Transition Profiles for the Square Cavity, Meshes A, B and C.




Figure 13. Comparison of Transition Profiles for the Circular Cavity, Meshes D, E and F.

It appears from Figure 12 that the normalized transition points of the different meshes for the square dielectric-loaded cavity occur at roughly the same point. Figure 13 suggests that the normalized transition points of the different meshes of the circular cavity also occur at roughly

the same point. This would indicate that the optimal normalized transition points are not strongly dependent on the number of patches or their arrangement. The transition points of the circular cavity do not, however, occur at the same points as those for the square cavity. Thus, there are other variables that do change these profiles. The similarity of Figures 12 and 13 suggest that it might be possible to develop an algorithm to assign transition points, leading to an efficient implementation of the one step p-refinement procedure, applicable to a wider range of problems than those considered here. Additional research is required to determine an efficient strategy for assigning transition points.

Conclusions

Aspects of p-refinement for vector finite elements have been investigated for two-dimensional cavity applications. Data suggest that representations with mixed polynomial orders offer better accuracy than those with a uniform polynomial order throughout the mesh. A one step p-refinement strategy is studied where the final polynomial order is assigned based on the normal field discontinuity in the zero-order solution. Results suggest that the zero-order result gives meaningful information about the regions to refine, and implies that such a strategy can provide a better accuracy-versus-efficiency trade off than methods based on a uniform polynomial order.

Two further studies would prove of immediate benefit. First, a study detailing the sensitivity of the error versus unknown curves to perturbations of the transition points would give an indication of to what accuracy the transition points must be found in order to reap the gains of a one step p-refinement method. The fact that the error trends for the one step solutions did not differ appreciably from those of the iterative method suggests that there may be significant freedom in choosing these transition points.

A second route of further study would be to attempt to correlate these optimized transition points to factors in the order zero results, such as the standard deviation of the discontinuities in the

order zero solution. If these other factors could be identified, it would be possible to implement an efficient one step p-refinement scheme.

Appendix A: A set of hierarchical vector basis functions for triangles

A set of proposed hierarchical functions for triangular cells is presented in Table 1. These are defined in terms of simplex coordinates (L_1, L_2, L_3) [6]. The lower-order members of this set are similar to the Webb and Forghani functions [8], but have been modified to satisfy the Nedelec conditions [5]. Functions have been included up to order 4/5. For source-free regions, the intent is to use the entire set up to mixed order 0/1, mixed order 1/2, mixed order 2/3, and so on. In other words, a refinement would involve increasing the order from 1/2 to 2/3, not from mixed 1/2 to complete 2. In each case the highest-order members of a given mixed group satisfy the Nedelec conditions.

In Table 1, functions of degrees 3 and 4 are constructed using polynomial products such as $(3L_1 - L_2)(L_1 - L_2)(L_1 - 3L_2)$ in order to enhance the linear independence of the functions. There are a number of ways of constructing such functions, and it can be argued that $(L_1 - L_2)$ is an equally valid way of expressing an equivalent degree of freedom, although possibly not as linearly independent — consider, for instance, the similarity in the shapes of $(L_1 - L_2)$ and $(L_1 - L_2)$. Other specific products could be used.

The vector basis set reported in Table 1 appears to be equivalent to the sets of functions recently reported by Savage [11], Andersen and Volakis [12-13], and Webb [14], if those sets are converted in a fairly obvious manner to triangular cells. This equivalence implies that an appropriate combination of any of these sets of functions conform to Nedelec's curl-conforming spaces. However, the specific elements of each set are different and parameters such as the associated matrix condition numbers may be different as a consequence.

| Table 1 Proposed Hierarchal Vector Bases | | |
|---|--|--|
| mixed order 0/1: | $L_2 \nabla L_3 - L_3 \nabla L_2$ $L_3 \nabla L_1 - L_1 \nabla L_3$ $L_1 \nabla L_2 - L_2 \nabla L_1$ | 3 edge-based functions total degrees of freedom = 3 |
| complete order 1: | $\nabla(L_1 L_2)$ $\nabla(L_1 L_3)$ $\nabla(L_2 L_3)$ | 3 edge-based functions total degrees of freedom = 6 |
| mixed order 1/2: | $L_1 (L_2 \nabla L_3 - L_3 \nabla L_2)$ $L_2 (L_3 \nabla L_1 - L_1 \nabla L_3)$ | 2 cell-based functions total degrees of freedom = 8 |
| complete order 2: | $\nabla\{L_1 L_2 (L_1 - L_2)\}$ $\nabla\{L_1 L_3 (L_1 - L_3)\}$ $\nabla\{L_2 L_3 (L_2 - L_3)\}$ $\nabla\{L_1 L_2 L_3\}$ | 3 edge-based functions 1 cell-based function total degrees of freedom = 12 |
| mixed-order 2/3: | $L_1 L_2 (L_2 \nabla L_3 - L_3 \nabla L_2)$ $L_1 L_3 (L_2 \nabla L_3 - L_3 \nabla L_2)$ $L_2 L_3 (L_3 \nabla L_1 - L_1 \nabla L_3)$ | 3 cell-based functions total degrees of freedom = 15 |
| complete order 3: | $\nabla\{L_1 L_2 (2L_1 - L_2) (L_1 - 2L_2)\}$ $\nabla\{L_1 L_3 (2L_1 - L_3) (L_1 - 2L_3)\}$ $\nabla\{L_2 L_3 (2L_2 - L_3) (L_2 - 2L_3)\}$ $\nabla\{L_1 L_2 L_3 (L_1 - L_2)\}$ $\nabla\{L_1 L_2 L_3 (L_1 - L_3)\}$ | 3 edge-based functions, 2 cell-based functions total degrees of freedom = 20 |
| mixed-order 3/4: | $L_1 L_2 (L_1 - L_2) (L_2 \nabla L_3 - L_3 \nabla L_2)$ $L_1 L_3 (L_1 - L_3) (L_2 \nabla L_3 - L_3 \nabla L_2)$ $L_2 L_3 (L_2 - L_3) (L_3 \nabla L_1 - L_1 \nabla L_3)$ $L_1 L_2 L_3 (L_3 \nabla L_1 - L_1 \nabla L_3)$ | 4 cell-based functions total degrees of freedom = 24 |

| Table 1, cont. | | |
|---|---|--|
| Proposed Hierarchal Vector Bases | | |
| complete order 4: | $\nabla\{L_1 L_2 (3L_1 - L_2) (L_1 - L_2) (L_1 - 3L_2)\}$ $\nabla\{L_1 L_3 (3L_1 - L_3) (L_1 - L_3) (L_1 - 3L_3)\}$ $\nabla\{L_2 L_3 (3L_2 - L_3) (L_2 - L_3) (L_2 - 3L_3)\}$ $\nabla\{L_1 L_2 L_3 (2L_1 - L_2) (L_1 - 2L_2)\}$ $\nabla\{L_1 L_2 L_3 (2L_1 - L_3) (L_1 - 2L_3)\}$ $\nabla\{L_1 L_2 L_3 (2L_2 - L_3) (L_2 - 2L_3)\}$ | 3 edge-based functions, 3 cell-based functions total degrees of freedom = 30 |
| mixed-order 4/5: | $L_1 L_2 (2L_1 - L_2) (L_1 - 2L_2) (L_2 \nabla L_3 - L_3 \nabla L_2)$ $L_1 L_3 (2L_1 - L_3) (L_1 - 2L_3) (L_2 \nabla L_3 - L_3 \nabla L_2)$ $L_2 L_3 (2L_2 - L_3) (L_2 - 2L_3) (L_3 \nabla L_1 - L_1 \nabla L_3)$ $L_1 L_2 L_3 (L_1 - L_2) (L_3 \nabla L_1 - L_1 \nabla L_3)$ $L_1 L_2 L_3 (L_1 - L_3) (L_3 \nabla L_1 - L_1 \nabla L_3)$ | 5 cell-based functions total degrees of freedom = 35 |

References

- [1] J. E. Akin, *Finite Elements for Analysis and Design*, Academic Press, San Diego, 1994.
- [2] O. C. Zienkiewicz and R. L. Taylor, *The Finite Element Method*. McGraw-Hill, New York, 1989.
- [3] M. Salazar-Palma et al., *Iterative and Self-Adaptive Finite Elements in Electromagnetic Modeling*. Boston: Artech House, 1998.
- [4] *Ansoft HFSS Users Manual*, October 1999.
- [5] J. C. Nedelec, "Mixed finite elements in R3," *Num. Math.*, vol. 35, pp. 315-341, 1980.
- [6] A. F. Peterson, S. L. Ray and R. Mittra, *Computational Methods For Electromagnetics*, IEEE Press, New York, 1998.
- [7] R. D. Graglia, D. R. Wilton and A. F. Peterson, "Higher Order Interpolatory Vector Bases for Computational Electromagnetics," *IEEE Transactions on Antennas and Propagation*, vol. 45, no. 3, pp. 329-342, March 1997.
- [8] J. P. Webb and B. Forghani, "Hierarchal Scalar and Vector Tetrahedra," *IEEE Trans. Magnetics*, vol. 29, pp. 1495-1498, March 1993.
- [9] J. S. Wang, "Hierarchic 'Edge' Elements for High-Frequency Problems," *IEEE Transactions on Magnetics*, vol. 33, no. 2, pp. 1536-1539, March 1997.
- [10] C. Carrie and J. P. Webb, "Hierarchal Triangular Edge Elements Using Orthogonal Polynomials," *Digest of the*

- 1997 *IEEE Antennas and Propagation International Symposium*, vol. 2, pp. 1310-1313, July 1997.
- [11] J. S. Savage, "Comparing high order vector basis functions," *Proceedings of the 14th Annual Review of Progress in Applied Computational Electromagnetics*, Monterey, CA, pp. 742-749, 1998.
- [12] L. S. Andersen and J. L. Volakis, "Hierarchical Tangential Vector Finite Elements for Tetrahedra," *IEEE Microwave and Guided Wave Letters*, vol. 8, pp.127-129, March 1998.
- [13] L. S. Andersen and J. L. Volakis, "Development and application of a novel class of hierarchical tangential vector finite elements for electromagnetics," *IEEE Transactions on Antennas and Propagation*, vol. 47, pp. 112-120, January 1999.
- [14] J. P. Webb, "Hierarchical vector basis functions of arbitrary order for triangular and tetrahedral finite elements," *IEEE Transactions on Antennas and Propagation*, vol. 47, pp. 1244-1253, August 1999.
- [15] R. S. Preissig, *Local P Refinement in Two Dimensional Vector Finite Elements*, Master's Thesis, Georgia Institute of Technology, Atlanta, GA, 1998.



that time he has worked for Texas Instruments in a field support role and is currently located in Atlanta, Georgia.



Andrew F. Peterson received the B.S., M.S., and Ph.D. degrees in Electrical Engineering from the University of Illinois, Urbana-Champaign in 1982, 1983, and 1986 respectively. Since 1989, he has been a member of the faculty of the School of Electrical and Computer Engineering at the Georgia Institute of Technology, where he is now Professor and Associate Chair for Faculty Development. Within ACES, he served for six years (1991-1997) as a member of the Board of Directors, and has been the Finance Committee Chair and the Publications Committee Chair.