

Optimal Discretization-Based Load Balancing for Parallel Adaptive Finite-Element Electromagnetic Analysis

Dennis D. Giannacopoulos, *Member, IEEE*

Abstract—The potential benefits of employing optimal discretization-based refinement criteria to achieve load balancing in parallel adaptive finite-element electromagnetic analysis are considered. Specifically, the ability of this class of adaption refinement criteria to resolve an effective domain decomposition based on initial discretizations with only relatively few degrees of freedom are examined. The new load-balancing method is evaluated with adaptively refined finite-element solutions for benchmark systems using practical adaption models.

Index Terms—Adaptive systems, electromagnetic analysis, finite-element methods, parallel processing.

I. INTRODUCTION

ADAPTIVE finite-element analysis (AFEA) for electromagnetic simulations is now a well-established research area [1], [2]. Adaptive solution strategies are capable of intelligently evolving efficient distributions of degrees of freedom (DOF) over the problem domain, by establishing solution error distributions and selectively adding DOF to the discretization to correct them. Hence, the primary benefit of AFEA is the efficient and accurate computational analysis of large continuum problems, for only a relatively small fraction of the cost of nonadaptive finite-element methods [1].

Despite the advantages of AFEA, the electromagnetic simulation of very complex structures can still be intractable within conventional sequential programming models [3]. The main difficulty is that a very large number of DOF are necessary for computing accurate solutions. Currently, one promising approach to overcome this computational barrier is to combine AFEA with high-performance computing (HPC) methods, such as parallel and distributed simulations [4], [5]. However, parallel and distributed AFEA simulations introduce challenging implementation issues that do not arise with simpler solution strategies [5]. In particular, adaptive algorithms evolve nonuniform discretizations that can make the task of balancing processor computational workload more difficult than with uniform structures. For example, load balancing can be especially problematic for distributed programming models when using domain decomposition based approaches to partition finite-element discretizations into subdomains, which are then assigned to specific processors. Essentially, some subdomains require more refinement

than others as a discretization is adapted. Thus, an initial discretization that is evenly distributed among processors in terms of workload, can lead to a severe workload imbalance as the adaption progresses toward more highly enriched and nonuniform discretizations [5].

Various types of load-balancing approaches for AFEA have been proposed [5]. However, one common drawback is the high communication costs of assessing the severity of the load imbalance, computing a more balanced work load distribution, and redistributing the improved work load among the available processors at a given iteration of the adaptive process. The purpose of this contribution is to investigate the practical value of an alternative approach of addressing the load-balancing problem for AFEA, which requires substantially less communication overall. Specifically, the ability of a class of adaption refinement criteria [2] to resolve an effective domain decomposition based on initial discretizations with only relatively few DOF are examined.

II. OPTIMAL DISCRETIZATION-BASED LOAD BALANCING

Typically, remapping or redistribution of DOF to available processors is performed repeatedly by load-balancing algorithms, at various stages of the adaption. In the case of remapping, a domain decomposition is recomputed to repartition the entire discretization into subdomains that, ideally, represent equal amounts of computation. Alternatively, DOF can be redistributed among existing subregions in order to balance the processor workload. Relatively sophisticated algorithms have been developed based on both approaches over the past several years [5]. However, the communication-to-computation cost ratios involved can still be quite high and reduce the overall efficiency of the AFEA process.

An alternative load-balancing approach for AFEA, which requires considerably less communication among processors, is to predict an efficient domain decomposition for the entire adaption process based on a single initial discretization containing relatively few DOF. The main advantages of this approach are twofold: it avoids having to solve the load-balancing problem repeatedly on large discretizations, and interprocessor communication, for the purpose of load balancing, is required only once. Clearly, the effectiveness of this type of approach is strongly dependent on its ability to predict how a coarse initial discretization will evolve throughout the adaptive process.

In this work, optimal discretization-based refinement criteria [2] are used to assess the relative solution error over an initial finite-element discretization during the early stages of the

Manuscript received July 1, 2003. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada.

The author is with the Department of Electrical and Computer Engineering, McGill University, Montreal, QC, H3A 2A7 Canada (e-mail: dennis.giannacopoulos@mcgill.ca).

Digital Object Identifier 10.1109/TMAG.2004.824713

adaptation. Subsequently, subregions are defined by equally distributing the total estimated solution error over each subregion, as predicted by the refinement criteria. Thus, if the number of subregions created is an integer multiple of the number of available processors, each processor can be assigned a partition of the initial discretization with approximately equal estimated solution error. Each processor may then apply AFEA to its corresponding partition until solution convergence is achieved. The hypothesis that must hold for this approach to yield a balanced work load, is that having approximately equal error in each partition of the initial mesh will result in equal work for the processors in the ensuing computations. The validity of this hypothesis is examined with fundamental benchmark systems in the following section.

A. Optimal Discretization-Based Refinement Criteria

Optimal discretization-based refinement criteria originate from an explicit formulation for computing optimal finite-element discretizations directly [2]. Namely, in [2] a set of optimization equations are developed for the geometric discretization parameters (i.e., the element vertex positions), which can be used to define AFEA refinement criteria. Thus, the relative discretization errors over a finite-element mesh can be estimated in terms of how well these equations are satisfied. Hence, the optimal discretization-based refinement criteria are defined implicitly as measures of the residuals of the geometric optimization equations [2], [7], [8]. Therefore, the newly proposed approach is based on equally distributing these residuals over each partition in a discretization. The essential components of the underlying derivation are discussed next.

Consider the scalar Helmholtz equation

$$\nabla \cdot (p\nabla u) + k^2 u = g \quad (1)$$

in the enclosed region Ω bounded by the surface S , with boundary conditions $u = u_d$ on S_d and $\partial u/\partial n = 0$ on S_n , where u is the electromagnetic field unknown to be solved for, p is a material-related parameter equal to ε, μ^{-1} , or 1, k is the wave number of the system, and g is a source function. If the surface S bounding the entire problem region Ω is comprised of the union of the Dirichlet and Neumann surfaces, i.e., $S = S_d + S_n$, then a corresponding functional may be written as

$$F(u) = \frac{1}{2} \int_{\Omega} \{p\nabla u \cdot \nabla u - k^2 u^2 + 2gu\} d\Omega \quad (2)$$

for which the true solution to (1) is the admissible function u that renders the above functional F stationary. For example, in two-dimensional (2-D) Cartesian problems, where $u = u(x, y)$, the optimization equations for the scalar triangular element with vertex positions $(x_l, y_l), l = 1, 2, 3$, corresponding to x_l and y_l , respectively, are given by the matrix forms

$$\frac{1}{2} \mathbf{u}^T \mathbf{V} \mathbf{u} - b_l \frac{k^2}{2} \mathbf{u}^T \mathbf{B} \mathbf{u} + b_l \mathbf{u}^T \mathbf{B} \mathbf{g} = 0 \quad (3)$$

and

$$\frac{1}{2} \mathbf{u}^T \mathbf{W} \mathbf{u} - c_l \frac{k^2}{2} \mathbf{u}^T \mathbf{B} \mathbf{u} + c_l \mathbf{u}^T \mathbf{B} \mathbf{g} = 0 \quad (4)$$

evaluated over the elements that share the vertex in question. Here, \mathbf{u} and \mathbf{g} are the field solution and source term vectors, respectively. The square matrices \mathbf{V} and \mathbf{W} are related to the

Laplacian part of the functional (2), and their entries are defined by

$$V_{ij} = \frac{1}{4A^2} \sum_{m=1}^3 \sum_{n=1}^3 \left[2A \frac{\partial(c_m c_n)}{\partial x_l} - (b_m b_n + c_m c_n) b_l \right] I_{ijmn} \quad (5)$$

and

$$W_{ij} = \frac{1}{4A^2} \sum_{m=1}^3 \sum_{n=1}^3 \left[2A \frac{\partial(b_m b_n)}{\partial y_l} - (b_m b_n + c_m c_n) c_l \right] I_{ijmn} \quad (6)$$

where A is the elemental area, and b_i and c_i are geometric parameters related to an element's vertex positions and can be defined as follows, with the subscripts progressing modulo 3

$$b_i = y_{i+1} - y_{i-1}, \quad \text{and} \quad c_i = x_{i-1} - x_{i+1}. \quad (7)$$

I_{ijmn} is the elemental integral (in homogeneous coordinates) of the product of the derivatives of the i th and j th basis functions, with respect to the m th and n th simplex coordinates. The second and third terms in (3) and (4) are related to the wave and source functional terms, respectively, where B_{ij} is the elemental integral (in homogeneous coordinates) of the product of the i th and j th basis functions. These functional-based formulae are valid for any choice of legitimate finite-element basis functions and may be computed for uniform- or mixed-order meshes, as may be required by specific refinement models such as h -, p -, or hp -adaptive methods. Analogous one-dimensional (1-D) and three-dimensional (3-D) formulas are given in [2]. Once the optimization equations for the vertex positions have been evaluated numerically for a given mesh, they may be used in various ways as refinement criteria [2], [7], [8]. One simple approach, referred to as "Type-A" in [7], is to assess a weighted sum of the residuals of the vertex-based optimization equations for each element and use these values to rank the elements for refinement. The practical benefits of using these types of refinement criteria for load balancing in parallel AFEA are examined in the next section.

B. Mesh Partitioning Algorithms for Load Balancing

To focus ideas, three mesh partitioning strategies for achieving load balancing are considered in this work. The first approach, labeled Uniform/Nonadaptive, is based on a uniform partitioning of an initial mesh into subregions containing equal numbers of DOF. The second strategy, Uniform/Adaptive, is based on uniformly partitioning the discretization into subregions containing equal numbers of DOF after several adaptive refinements of the initial mesh. Finally, the new approach is based on using optimal discretization-based refinement criteria to partition the same mesh as for the second method above into subregions of equal estimated solution error, as explained in the previous section.

III. RESULTS

A simple 1-D free-space example is presented first in order to examine the potential significance of the new load-balancing approach for AFEA. In addition, the effectiveness of the newly

TABLE I
PERCENT OF TOTAL WORKLOAD VERSUS PROCESSOR NUMBER

| Method | Processor | | | |
|----------------------|-----------|-------|-------|-------|
| | 1 | 2 | 3 | 4 |
| Uniform/Non-adaptive | 77.3% | 9.09% | 9.09% | 4.54% |
| Uniform/Adaptive | 36.4% | 27.3% | 22.7% | 13.6% |
| New Approach | 22.7% | 27.3% | 27.3% | 22.7% |

proposed load-balancing approach is examined with practical AFEA models for a 2-D Laplace system. For both cases, the three mesh partitioning strategies described above were considered and their relative performance assessed based on h -adaption refinements.

A. 1-D: Free Space Point Charge Test System

This static benchmark system was first studied in [6] and is based on the classical point singularity model in free space. The objective for this benchmark system is to compute the functional value for a radial neighborhood close to the point charge and spanning a 100-fold decay in electric scalar potential: the point charge, of magnitude $10^{-9}/9$ C, is located at the origin, and the two boundaries of the problem domain are set at radial distances of 0.1 m and 10 m away from the charge. The results for the 1-D analysis are given in Table I. All results were computed using standard first-order Lagrangian elements. Evidently, the new approach is able to achieve a better workload balance than the other two approaches considered.

Although a 1-D analysis for this system is possible, the formulation used properly accounts for the entire 3-D volume associated with the problem, in order to correctly apply the underlying variational principle. The primary feature of this system is the rapid field solution variation close to the singularity. This feature is common to many practical devices that contain sharp material corners and has been shown to drastically reduce finite-element solution convergence rate. In addition, this test system was considered, primarily, to examine the potential of the new approach more readily and with fewer obstacles than 2-D or 3-D test systems permit.

B. 2-D: Laplace Test System

The 2-D Laplace benchmark system is described by Fig. 1. It is one quarter of a square coaxial line in cross-section—the standard “L” problem. The conductor boundary conditions are 1 V and 0 V, as indicated, and the symmetry planes are labeled N . All results reported for this system were computed using standard second-order Lagrangian elements. Performance results for the load-balancing studies are presented in Fig. 2. Note, the new approach is able to achieve an almost ideal workload balance among the processors compared to the other two approaches, which yield significant load imbalances.

For reference, sample partitioned meshes are shown in Figs. 3–7 for the various load-balancing algorithms considered in this work. Specifically, the initial eight element mesh used for each of the three mesh partitioning strategies is shown in Fig. 1, as well as the four partitions (labeled A, B, C, and D) used for the Uniform/Nonadaptive load-balancing method. Fig. 3 shows a sample mesh for the Uniform/Nonadaptive

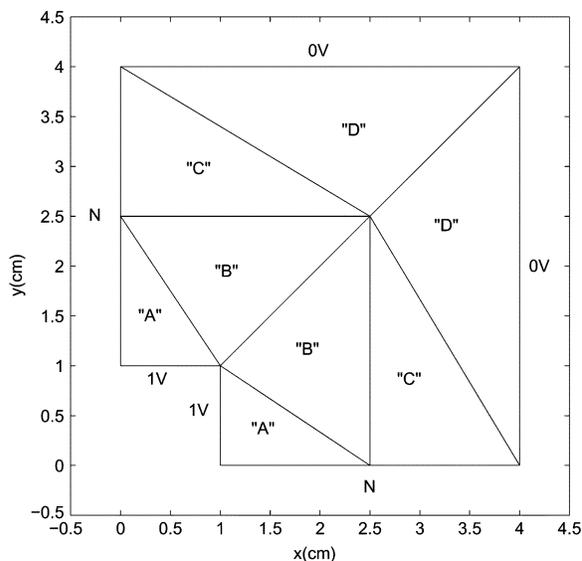


Fig. 1. Laplace 2-D benchmark system: initial mesh (eight triangles) and partitions for Uniform/Nonadaptive method.

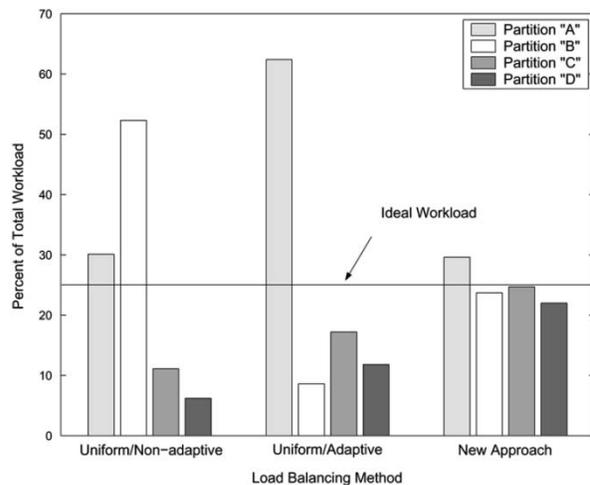


Fig. 2. Laplace 2-D test system results.

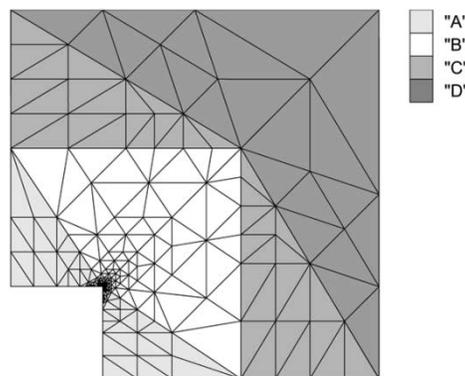


Fig. 3. Sample mesh for Uniform/Nonadaptive method (352 triangles).

method applied to the initial mesh with the partitions shown in Fig. 1. The partitions for the Uniform/Adaptive strategy and the new approach are shown in Figs. 4 and 6, respectively. Note that both of these meshes are identical, by definition, but the subdomains assigned to the various processors are quite

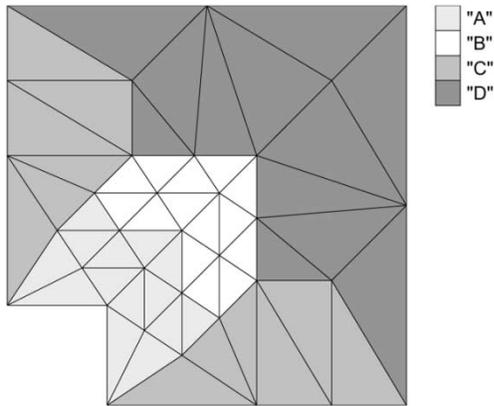


Fig. 4. Partitions for Uniform/Adaptive method (52 triangles).

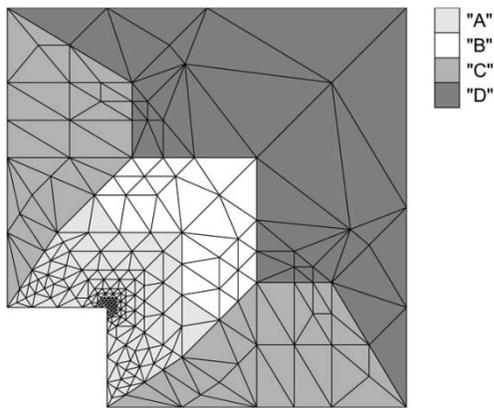


Fig. 5. Sample mesh for Uniform/Adaptive method (372 triangles).

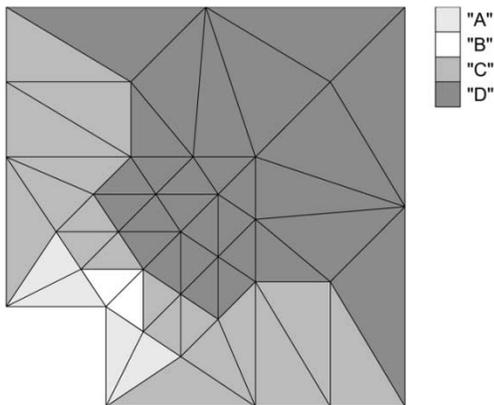


Fig. 6. Partitions for the new approach (52 triangles).

different. For the Uniform/Adaptive strategy, the partitioning is according to the criterion of equal numbers of DOF per subdomain; for the new approach, equal error per subdomain was the objective, as explained in the previous section. The estimated solution errors for each of the partitions shown in Fig. 6 were within $\pm 2.5 \times 10^{-4}$ of each other. Sample meshes resulting from the two approaches are shown in Figs. 5 and 6, respectively. It may be noted that for each approach considered, the adaptive process was terminated when the global solution

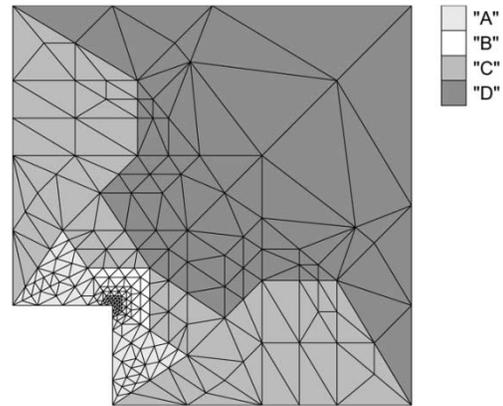


Fig. 7. Sample mesh for the new approach (372 triangles).

per cent error estimate was less than 0.1%, and the results reported in Fig. 2 are based on these terminal meshes.

IV. CONCLUSION

A new load-balancing approach for parallel AFEA that uses optimal discretization-based refinement criteria has been proposed and evaluated for scalar electromagnetic systems. The results for the benchmark systems investigated demonstrate that the new approach is able to achieve effective domain decompositions based on initial discretizations with only relatively few DOF. Compared with existing techniques, the new approach requires considerably less communication among processors. These preliminary findings suggest that further algorithmic sophistications and experimental studies in 3-D may be merited. For instance, future research should focus on possible extensions of the method for application to transient problems involving eddy currents or spatial motions.

REFERENCES

- [1] P. Fernandes, P. Girdinio, P. Molfino, and M. Repetto, "Local error estimates for adaptive mesh refinement," *IEEE Trans. Magn.*, vol. 24, pp. 299–302, 1988.
- [2] D. Giannacopoulos, "Optimal Discretization-Based Adaptive Finite Element Analysis for Electromagnetics," Ph.D. Thesis, McGill University, 1998.
- [3] —, "Toward optimal error distributions in adaptive finite element electromagnetic analysis for microelectronic interconnection structures," *IEEE Trans. Magn.*, vol. 38, pp. 401–402, 2002.
- [4] S. McFee and D. Giannacopoulos, "The implications of parallel processing on h - p adaptive finite element analysis for electromagnetics," *IEEE Trans. Magn.*, vol. 34, pp. 3284–3287, 1998.
- [5] C. J. Liao and Y. C. Chung, "Tree-based parallel load-balancing methods for solution-adaptive finite element graphs on distributed memory multicomputers," *IEEE Trans. Parallel Distrib. Syst.*, vol. 10, pp. 360–370, 1999.
- [6] D. Giannacopoulos and S. McFee, "Toward optimal h - p adaption near singularities in finite element electromagnetics," *IEEE Trans. Magn.*, vol. 30, pp. 3523–3526, 1994.
- [7] S. McFee and D. Giannacopoulos, "Optimal discretization based refinement criteria for finite element adaption," *IEEE Trans. Magn.*, vol. 32, pp. 1357–1360, 1996.
- [8] D. Giannacopoulos and S. McFee, "Functional derivatives and optimal discretization based refinement criteria for adaptive finite element analysis with scalar tetrahedra," *IEEE Trans. Magn.*, vol. 35, pp. 1326–1329, 1999.