The Implications of Parallel Processing on *h-p* Adaptive Finite Element Analysis for Electromagnetics

Steve McFee and Dennis Giannacopoulos
Electrical Engineering Department, McGill University
3480 University Street, Montreal, PQ, H3A 2A7, Canada

Abstract — **The primary implications of parallel processing on *h-p* adaptive finite element methods for electromagnetic analysis are investigated. Aside from the conventional benefits and costs associated with the parallelization of the essentially non-adaptive finite element modules, significant fundamental advantages that are unique to the adaptive process itself are explored. First, the overall speedup potential of local error estimator evaluation and *h-p* discretization refinement is superior to that of finite element solver execution in parallel environments, and therefore justifies the use of more complex and computationally intensive adaption control strategies. Second, the availability of parallel processing motivates the comparative assessment of different discretization strategies at each *h-p* refinement step to help guide the evolution of the adaption. Practical results representing a range of parallel configurations are computed to illustrate the concepts.**

Index terms — **Finite element methods, adaptive systems, error analysis, parallel algorithms, electromagnetic analysis.**

I. INTRODUCTION

Today, the study of finite element methods for the analysis of electromagnetic systems in electrical engineering is a well-established and mature research area [1]. Further, a number of *h*-type and *p*-type adaptive systems have been developed, and are now in reasonably widespread use [2], [3]. However, *h-p* adaptive finite element analysis for electromagnetic systems is still a relatively new and underdeveloped area of research [4]; and conclusive contributions towards an efficient and effective coupling of parallel processing methods and *h-p* adaption for electromagnetics are virtually unheard of. To date, no results for *h-p* adaptive electromagnetics applications are reported in the mainstream literature. Even outside the electromagnetics area, the only parallelization results available for *h-p* adaption are limited to conventional translations of previously reported algorithms developed for sequential environments [5].

The objective of this contribution is to lay the ground work and establish the primary requirements for the development of efficient and effective parallel processing formulations for *h-p* adaptive finite element methods for electromagnetic analysis. Unlike previous approaches, the emphasis is on developing and tuning *h-p* adaptive subsystems and strategies to better exploit the inherent strengths of a parallel environment, as opposed to building more efficient parallelizations of existing algorithms which were not developed with parallel processing in mind. It is firmly believed that only a fundamental approach that takes the nature of the parallel environment into consideration from the ground up can yield an *h-p* adaptive system that is capable of exploiting the full potential of that parallel environment.

The *h-p* adaptive process for finite element analysis may be viewed as a simple system comprised of complex subsystems. While the individual subsystems can be sophisticated and very large, they are also quite well-understood and essentially non-adaptive. On the other hand, while the main feedback system consists of only a handful of steps, they can be quite sensitive and fairly subtle in their interaction. The *h-p* adaptive feedback loop for finite element analysis is defined as follows [4]:

A. Generate initial finite element discretization.

Repeat:

- B. Solve finite element problem.
- C. Evaluate solution accuracy; if adequate then **STOP**.
- D. Identify regions of inadequate discretization.
- E. Determine required *h-p* discretization refinements.
- F. Update finite element discretization.

The present work is focused on the steps which are particular to the adaptive process itself, i.e., C, D, E and F; the development of optimized parallelizations for the finite element subsystems associated with steps A and B is left to later contributions.

Manuscript received November 3, 1997. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada.

II. PARALLEL *h-p* ADAPTIVE PROCESSES

According to current theory and practice, it is important to keep the net computational cost (measured as elapsed runtime) of the adaptive control and discretization refinement processes small compared to that of the finite element solver, in order to achieve the full potential of the *h-p* adaptive approach [3], [5]. Based on typical *h*-type and *p*-type adaptive implementations, an elapsed runtime cost ratio of 10:1 or more is often realized. This distribution of computational effort indicates an empirical balance between adaptive refinement and solution calculations that seems to be effective for most practical adaptive systems running on sequential machines. However, for parallel implementations, this balance can be much less efficient, because the full-scale speedup potential of local error estimator evaluation and *h-p* discretization refinement is far superior to that of finite element solver execution.

For dense matrices in practical parallel environments, linear system solutions cost $O(n^2)$ in elapsed time, (compared to $O(n^3)$ for serial processing) [6]; while, for the sparse matrices typical of finite element analysis, the best execution rates are limited to $O(n \log n) \rightarrow O(n^{1.5})$ for the most efficient parallel formulations, due to data communication and indirect addressing costs [7]. Unfortunately, this performance is not all that much better than what can be achieved with efficient serial processing [6], [8]. However, the speedup potential for error estimator evaluation under parallel implementations is far more substantial. These calculations are strongly localized, and usually only rely on the data associated with the element under evaluation. In the ideal case for SIMD implementations, the runtime cost is essentially independent of the number of elemental evaluations computed! Therefore, relative to serial implementations, the use of more complex and computationally intensive adaption strategies is indicated. Further, a parallel environment permits the comparative assessment of competing discretization schemes at each *h-p* refinement step, to help guide the evolution of the adaption.

0018-9464/98$10.00

III. Investigative Studies

At each refinement step, the main concerns of *h-p* adaption are: *where* should extra degrees of freedom (DOF) be inserted; *what type* of DOF should be used; and *how many* DOF should be added. The following investigations have been designed to explore the possibilities and potential advantages of addressing these points within a parallel processing environment. The first study examines the value of using pairs of complementary error estimators to determine where to add DOF to a discretization. In this case, an unbiased average of two complementary errors is used to assess each element, at each adaptive step, within a practical *h*- followed by *p*-adaption system. The second study investigates the potential benefits of constructing and solving both an *h*- and a *p*-refined discretization at each adaptive step, in order to determine what type of DOF should be added to the discretization at each step. In this case, both a single step depth search (2 refinement scenarios: pure *h* and pure *p*) and a double step depth search (4 refinement scenarios: pure *h* followed by pure *h*; pure *h* followed by pure *p*; pure *p* followed by pure *p*; and pure *p* followed by pure *h*) are examined. The third study also addresses the "what type of DOF" question. In this case, the potential value of adding a mixture of *h*- and *p*-type DOF at each adaptive step is tested: 50% of the prescribed DOF update are inserted as *p*-type, guided by a *p*-type error estimator; and the remaining 50% are added as *h*-type, according to an *h*-type error estimator. The fourth study examines the advantages of constructing and solving a range of discretizations that differ only in the number of new DOF added, at each adaptive step, in order to determine how many DOF should be added to the discretization at each step. In this case, four different %DOF refinement levels, ranging from a 25% DOF update, to refining every element in the discretization, are investigated and compared within a practical *h*-followed by *p*-adaption system. The final study also addresses the "how many DOF" question. In this case, two straightforward schemes which are based on the distribution and relative strengths of the errors over the discretization are used to determine how many DOF should be added at each adaptive step. The first scheme simply directs that all elements with above average error levels should be refined; the second scheme scans the error level list, sorted by descending magnitude, for the first statistically significant abrupt jump in error level, and then directs that all elements with errors above that level should be refined.

IV. Results

The five investigative studies described above were carried out using two basic test systems: the standard "L" benchmark problem defined in [9]; and a high-frequency variation based on the same 2D geometry and initial mesh. Briefly, Fig.1 represents 1/4 of the cross-section of an infinitely long, translationally symmetric, air-filled, coaxial line – for test system 1. The objective is to resolve the electrostatic scalar potential field in the air between the conductors, when a unit voltage difference is maintained across the conductors. For test problem 2, Fig. 1 represents a sharply truncated 90° corner in a planar microstrip circuit. In this case, the goal is to resolve the variation of E_N in the substrate between the strip and the ground plane, assuming that: one port carries unit excitation; the other port is shorted; the boundaries are modelled as perfect magnetic walls; and the system operates at a normalized frequency = 1/3 port width.

The results for the first three investigative studies are based on test problem 1 calculations, using an initial discretization of eight first-order triangles. The results for studies four and five are based on test problem 2 and an initial discretization of eight second-order triangles. In all five cases, comparative adaptive performance results are described in terms of normalized functional error versus cumulative computational cost. Further, all reported results are derived from and representative of test data spanning a minimum 1000-fold reduction in functional error. Finally, 50% DOF updates (per step) were used exclusively in the first three studies; and only *h*- followed by *p*-adaption was employed in studies one, four and five.

The results of the first investigative study are reported in Fig. 2. The standard field-discontinuity and the recently developed functional-gradient based error estimators [9] were used as a complementary pair for *h*-adaption; while the PDE residual and hierarchal coefficient estimators [3] were used for *p*-adaption. In addition to the performance obtained by using the averages of the two sets of estimators, three related curves are plotted to gauge adaption efficiency. The two standard adaption results, corresponding to the best (most efficient) and worst of the four possible pairings of the two *h*-type and two *p*-type estimators, together with the optimal uniform refinement *h*- followed by *p*-adaption result, are provided.

The results of the second study are reported in Fig. 3. To fix the focus of the investigation, the PDE residual error estimator was used throughout for both *h*- and *p*-adaption. In this study, optimal *h*- followed by *p*-adaption results for both 50% DOF (Ideal h-p) and uniform (Uniform h-p) updates are provided for comparison. The single step depth search results are denoted Level 1 h-p; the double step are labelled Level 2 h-p.

The results of investigative study three are reported in Fig. 4. As in the previous study, the PDE residual error estimator was used throughout. In this case, the performance results for pure *h*-adaption, pure *p*-adaption, and optimal uniform refinement *h*-followed by *p*-adaption have been provided for comparison.

The results of the fourth study are reported in Fig. 5. In this case, the PDE residual estimator was used for *h*-adaption, and the hierarchal coefficient estimator was applied for *p*-adaption.

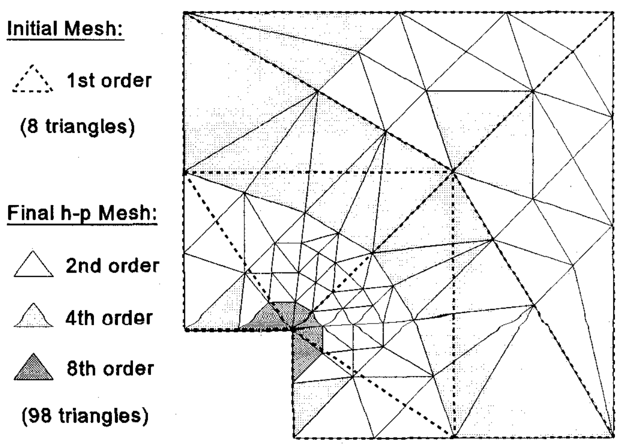

Fig. 1. 2D geometry and initial mesh for the two test problems; and final *h-p* discretization for first investigative study, based on averaged error estimators.

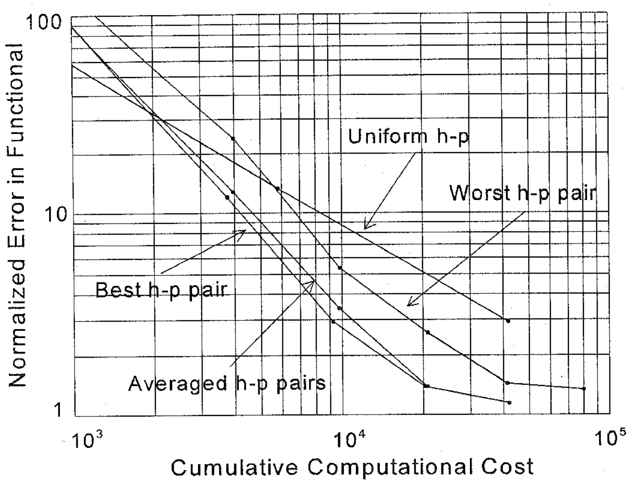

Fig. 2. Comparative *h-p* adaption performance results for the first investigative study, on using averages of complementary pairs of error estimators.

The four DOF refinement levels examined were: 25% DOF; 50% DOF; 100% DOF; and uniform refinement. These four refinement updates were constructed, solved and compared at each adaptive step, to determine how many DOF to add to the discretization at each step. For comparative purposes, the performance of this search-based result (labelled Mixed) is plotted together with the four fixed %DOF update adaption results.

The results of the fifth investigative study are reported in Fig. 6. As in the fourth study, the PDE residual error estimator was used for *h*-adaption, and the hierarchal coefficient error estimator was used for *p*-adaption. The "above average" DOF result is labelled Average; the "abrupt jump" DOF result is labelled Variable. The four fixed %DOF update adaption results given in Fig. 5 are also plotted in Fig. 6 for comparison.

V. Analysis of Results

Each of the investigative studies described in section III, and implemented in section IV, were designed to explore the possibilities and potential advantages of using parallel processing in *h-p* adaptive finite element analysis for electromagnetics. The

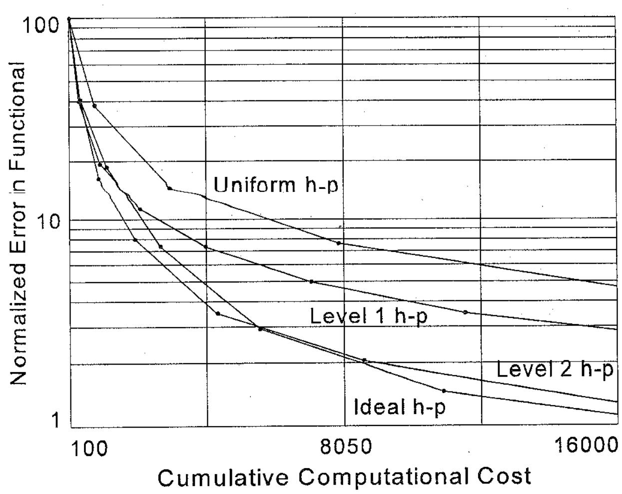

Fig. 3. Comparative *h-p* adaption performance results for the second investigative study, on monitoring both *h*- and *p*-type updates simultaneously.

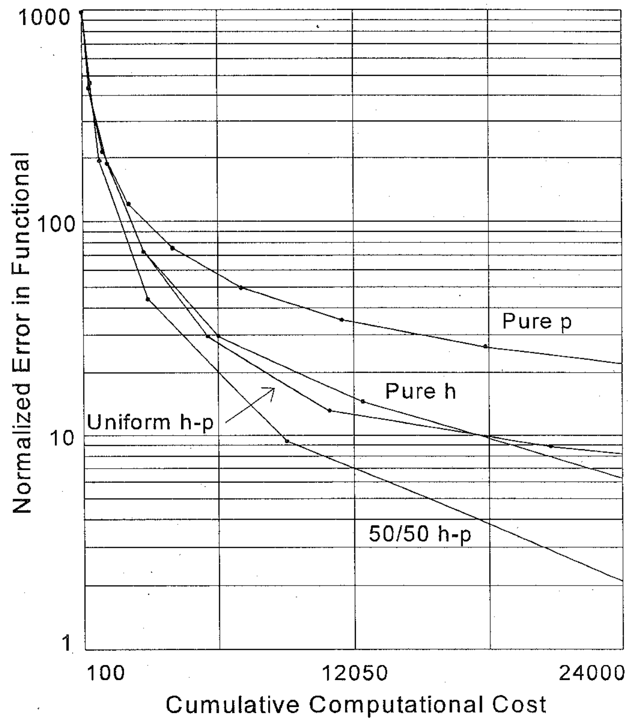

Fig. 4. Comparative adaption performance results for the third investigative study, on using hybrid *h*- and *p*-type refinements in the same adaptive step.

results of these computational experiments provide supportive evidence for a number of basic hypotheses on the considerable advantages associated with a parallel processing environment. The most significant of these findings are summarized below.

A. Facts: Placement of new DOF is critical to adaption performance; two or more error estimators can be evaluated nearly as inexpensively as one in a parallel environment. Hypothesis: Complementary estimators should be able to render more detailed and decisive error distributions with less refinement-model dependent bias and distortion, and therefore yield increased adaption efficiency and stability

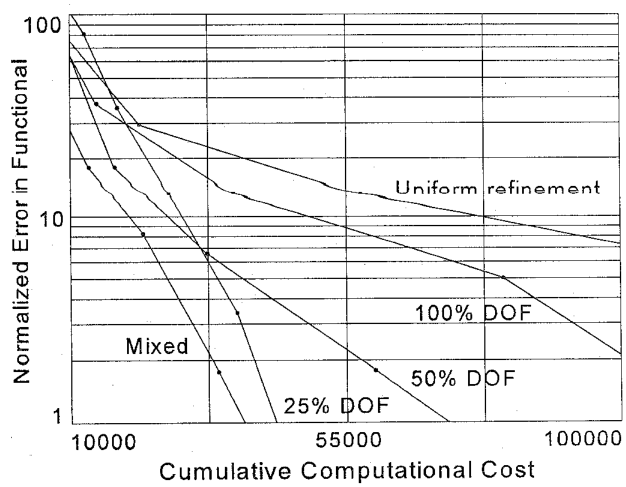

Fig. 5. Comparative *h-p* adaption performance results for the fourth investigative study, on monitoring multiple %DOF updates simultaneously.

at little added cost in a parallel processing environment. Support: (study one) Averaged estimator techniques can perform substantially better than methods that apply the same estimators singly; in this test, the averaged approach was comparable to the best individual strategy available.

B. Facts: The types of DOF added to an evolving discretization are very important to *h-p* adaption performance; two or more refinement scenarios can be constructed and evaluated almost as cheaply as one in a parallel environment. Hypothesis: Comparative evaluations of potential *h*- and *p*-refinement update scenarios at each step should be able to deliver more appropriately focused discretizations, and thereby lead to better optimized *h-p* adaption trajectories. Support: (study two) Locally optimized refinement type selections can yield substantial improvements in adaption efficiency over standard *h* followed by *p* strategies; and, the benefit seems to increase with the depth of the search and number of refinement update scenarios considered.

C. Facts: The amount of DOF added to an evolving discretization is also important to adaption performance; two or more mesh refinements can be constructed and evaluated nearly as inexpensively as one in a parallel environment. Hypothesis: Comparative evaluations of different %DOF refinements should be able to identify more effective updates, and yield better optimized *h-p* adaption trajectories. Support: (study four) Locally optimized %DOF updates can yield striking improvements in overall adaption efficiency over standard fixed size %DOF refinement updates; in this test, the optimized strategy strongly outperformed every one of the standard updates used to determine it!

D. Facts: The performance of *h-p* adaptive formulations for sequential environments can be improved when the types and amount of DOF may be tuned to each adaptive step. Hypothesis: The parallel processing strategies discussed above should be able to be adapted to sequential adaption systems, to yield related benefits in a serial environment. Support: (study three) The enhanced performance of the 50/50 mixture of *h*- and *p*-refinements (per adaptive step) indicates that non-trivial *h-p* adaption efficiency improvements can be realized at relatively small additional cost in a serial environment. Further, (study five) the remarkable efficiency achieved by the "abrupt jump" (Variable) DOF updates shows that excellent serial adaption performance improvements are possible at a very reasonable cost.

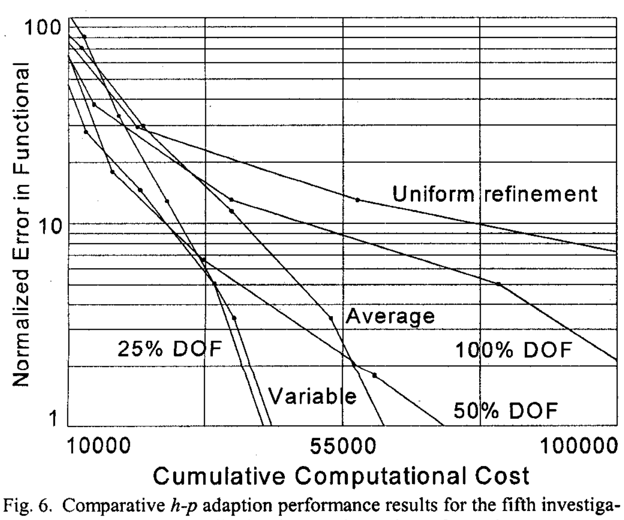

Fig. 6. Comparative *h-p* adaption performance results for the fifth investigative study, on using error distributions to determine DOF update amounts.

VI. CONCLUSIONS

This preliminary research has demonstrated that there can be substantial advantages associated with the development of *h-p* adaptive finite element methods for electromagnetic analysis in parallel environments. The fundamental investigative approach adopted in this work has revealed the possibilities and potential benefits of designing parallel processing adaptive strategies and modules from the ground up, as opposed to simply constructing parallelizations of existing sequential algorithms. Finally, new *h-p* adaption error estimation and local refinement strategies for practical parallel processing have been developed and tested.

A number of related issues have been purposefully excluded from this presentation in order to balance its scope, specificity and clarity. The authors believe that many of these interesting and important matters justify the focus of future research, e.g.

A. Which adaptive strategies and formulations are best suited to which types of parallel environments; different parallel computing facilities possess different strengths and weaknesses, e.g. data communications costs and overhead.

B. There are indications that other adaption modules can also benefit from a parallel processing redesign, e.g. *h-p* mesh generation [10]. What are the potential gains and costs?

C. The parallel adaption strategies presented in this contribution are essentially independent; how can they be used in combination effectively, and what are the implications?

REFERENCES

[1] P.P. Silvester and R.L. Ferrari, *Finite Elements for Electrical Engineers*, Third Edition, Cambridge University Press, Cambridge, 1996.

[2] P. Fernandes, P. Girdinio, M. Repetto and G. Secondo, "Refinement strategies in adaptive meshing", *IEEE Trans. Magn.*, Vol. 28(2), pp. 1739-1742, March, 1992.

[3] S. McFee and J.P. Webb, "Adaptive finite element analysis of microwave and optical devices using hierarchal triangles", *IEEE Trans. Magn.*, Vol. 28(2), pp. 1708-1711, March, 1992.

[4] D. Giannacopoulos and S. McFee, "Towards optimal *h-p* adaption near singularities in finite element electromagnetics", *IEEE Trans. Magn.*, Vol. 30(5), pp. 3523-3526, September, 1994.

[5] J.T. Oden and A. Patra, "Parallel adaptive strategy for *hp* finite element computations", *Comp. Meth. Appl. Mech. Eng.*, Vol. 121(2), pp. 449-470, March 1995.

[6] G.H. Golub and C.F. Van Loan, *Matrix Computations*, Second Edition, Chapter 6, Johns Hopkins University Press, Baltimore, Maryland, 1989.

[7] A. Benaini, D. Laiymani and G.R. Perrin, "A reconfigurable parallel algorithm for sparse cholesky factorization", *Lecture Notes in Computer Science: Parallel Algorithms for Irregularly Structured Problems*, Vol. 980, pp. 261-274, Springer, New York, 1995.

[8] I.S. Duff, A.M. Erisman and J.K. Reid, *Direct Methods for Sparse Matrices*, Oxford Univesity Press, New York, 1989.

[9] S. McFee and D. Giannacopoulos, "Optimal discretization based refinement criteria for finite element adaption", *IEEE Trans. Magn.*, Vol. 32, pp. 1357-1360, May, 1996.

[10] S. McFee and J.P. Webb, "Automatic mesh generation for *h-p* adaption", *IEEE Trans. Magn.*, Vol. 29(2), pp. 1894-1897, March, 1993.