

A Comparison of hp -Adaptive Strategies for Elliptic Partial Differential Equations

WILLIAM F. MITCHELL and MARJORIE A. MCCLAIN, National Institute of Standards and Technology

The hp version of the finite element method (hp -FEM) combined with adaptive mesh refinement is a particularly efficient method for solving PDEs because it can achieve an exponential convergence rate in the number of degrees of freedom. hp -FEM allows for refinement in both the element size, h , and the polynomial degree, p . Like adaptive refinement for the h version of the finite element method, *a posteriori* error estimates can be used to determine where the mesh needs to be refined, but a single error estimate can not simultaneously determine whether it is better to do the refinement by h or p . Several strategies for making this determination have been proposed over the years. These strategies are summarized, and the results of a numerical experiment to study the performance of these strategies is presented. It was found that the reference solution based methods are very effective, but also considerably more expensive, in terms of computation time, than other approaches. The method based on *a priori* knowledge is very effective when there are known point singularities. The method based on the decay rate of the expansion coefficients appears to be the best choice as a general strategy across all categories of problems, whereas many of the other strategies perform well in particular situations and are reasonable in general.

Categories and Subject Descriptors: G.1.8 [Numerical Analysis]: Partial Differential Equations—*elliptic equations; finite element methods*; G.4 [Mathematical Software]: Algorithm design and analysis

General Terms: Algorithms, Experimentation, Performance

Additional Key Words and Phrases: adaptive mesh refinement, hp -adaptive strategy, hp -FEM

ACM Reference Format:

ACM Trans. Math. Softw. V, N, Article A (January YYYY), 38 pages.

DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

1. INTRODUCTION

The numerical solution of partial differential equations (PDEs) is the most computationally intensive part of a wide range of scientific and engineering applications. Consequently the development and application of faster and more accurate methods for solving partial differential equations has received much attention in the past fifty years. Many of the applications at the cutting edge of research are extraordinarily challenging. For these problems it is necessary to allocate computing resources in an optimal way in order to have any chance at solving the problem. Determining the best grid and approximation space on which to compute the solution is a central concern in this regard. Unfortunately, it is rarely possible to determine an optimal grid in advance. Thus, developing self-adaptive techniques which lead to optimal resource allocation is critical for future progress in many fields.

Preprint. Appeared in *ACM Trans. Math. Soft.*, 41 (1), 2014.

Applied and Computational Mathematics Division, National Institute of Standards and Technology. Contribution of NIST, not subject to copyright.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0098-3500/YYYY/01-ARTA \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

Self-adaptive methods have been studied for over 30 years now. They are often cast in the context of finite element methods, where the domain of the PDE is partitioned into a mesh consisting of a number of elements (in two dimensions, usually triangles or rectangles), and the approximate solution is a polynomial over each element. Most of the work has focused on h -adaptive methods. In these methods, the mesh size, h , is adapted locally by means of a local error estimator with the goal of placing the smallest elements in the areas where they will do the most good. In particular, elements that have a large error estimate get refined so that ultimately the error estimates, and presumably the error, are approximately equal over all elements. h -adaptive methods are quite well understood now, and are beginning to be used in practice.

Recently, the research community has begun to focus more attention on hp -adaptive methods. In these methods, one not only locally adapts the size of the mesh, but also the degree of the polynomials, p . The attraction of hp -adaptivity is that the error converges at an exponential rate in the number of degrees of freedom, as opposed to a polynomial rate for fixed p . Much of the theoretical work showing the advantages of hp -adaptive methods was done in the 1980's, but it wasn't until the 1990's that practical implementation began to be studied. The new complication is that the local error estimator is no longer sufficient to guide the adaptivity. It tells you which elements should be refined, but it does not indicate whether it is better to refine the element by h or by p . A method for making that determination is called an hp -adaptive strategy. A number of strategies have been proposed, but it is not clear which ones perform best under different situations, or even if any of the strategies are good enough to be used as a general purpose solver. In this paper we present an experimental comparison of several hp -adaptive strategies.

Any study of this type is necessarily limited in scope. The comparison will be restricted to steady-state linear elliptic partial differential equations on bounded domains in two dimensions with Dirichlet, natural or mixed boundary conditions. The standard Galerkin finite element method will be used with the space of continuous piecewise polynomial functions over triangles that are refined by the newest node bisection method.

The remainder of the paper is organized as follows. In Section 2 we define the equation to be solved, present the finite element method, and give some *a priori* error estimates. In Section 3 we give the details of the hp -adaptive finite element algorithm used in the experiments. Section 4 defines the hp -adaptive strategies to be compared. Section 5 contains the results of the experiments. Finally, we draw our conclusions in Section 6.

2. THE FINITE ELEMENT METHOD

We consider the elliptic partial differential equation

$$\mathcal{L}u = -\operatorname{div}(A\nabla u) + r(x, y)u = f(x, y) \text{ in } \Omega \quad (1)$$

$$u = g_D(x, y) \text{ on } \partial\Omega_D \quad (2)$$

$$\mathcal{B}u = A\nabla u \cdot n + c(x, y)u = g_N(x, y) \text{ on } \partial\Omega_N \quad (3)$$

where Ω is a bounded, connected, polygonal, open region in \mathbb{R}^2 with boundary $\partial\Omega = \partial\Omega_D \cup \partial\Omega_N$, $\partial\Omega_D \cap \partial\Omega_N = \emptyset$, $A = \begin{bmatrix} p(x, y) & 0 \\ 0 & q(x, y) \end{bmatrix}$, and n is the outward unit normal. If $c = 0$, Equation 3 is the natural boundary condition. If, in addition, $p = q = 1$ or $\partial\Omega_N$ consists of line segments that are parallel to the axes, Equation 3 is the Neumann boundary condition. We assume the data in Equations 1-3 satisfy the usual ellipticity and regularity assumptions. In one of the test problems, we extend the equation to a

system of two equations containing a cross derivative term $\partial^2 u / \partial x \partial y$, and in another test problem we include first order derivative terms.

As usual, define the space L^2 by

$$L^2(\Omega) = \{v(x, y) : \iint_{\Omega} v^2 dx dy < \infty\}$$

with inner product

$$\langle u, v \rangle_2 = \iint_{\Omega} uv dx dy$$

and norm

$$\|v\|_2^2 = \langle v, v \rangle_2.$$

We denote by $H^m(\Omega)$ the usual Sobolov spaces

$$H^m(\Omega) = \{v \in L^2(\Omega) : D^\alpha v \in L^2(\Omega), |\alpha| \leq m\}$$

where

$$D^\alpha v = \frac{\partial^{|\alpha|} v}{\partial^{\alpha_1} x \partial^{\alpha_2} y}, \alpha = (\alpha_1, \alpha_2), \alpha_i \in \mathbb{N}, |\alpha| = \alpha_1 + \alpha_2.$$

The Sobolov spaces have inner products

$$\langle u, v \rangle_{H^m(\Omega)} = \iint_{\Omega} \sum_{|\alpha| \leq m} D^\alpha u D^\alpha v dx dy$$

and norms

$$\|v\|_{H^m(\Omega)}^2 = \langle v, v \rangle_{H^m(\Omega)}.$$

We will also refer to the seminorm $|v|_{H^m(\Omega)}$ where the sum is over $|\alpha| = m$.

Let $H_{0, \partial\Omega_D}^1(\Omega) = \{v \in H^1(\Omega) : v = 0 \text{ on } \partial\Omega_D\}$. Let $\tilde{u}_D \in H^1(\Omega)$ be a lift function satisfying the Dirichlet boundary conditions in Equation 2 and define the affine space $\tilde{u}_D + H_{0, \partial\Omega_D}^1(\Omega) = \{\tilde{u}_D + v : v \in H_{0, \partial\Omega_D}^1(\Omega)\}$. Define the bilinear form

$$B(u, v) = \iint_{\Omega} p \frac{\partial u}{\partial x} \frac{\partial v}{\partial x} + q \frac{\partial u}{\partial y} \frac{\partial v}{\partial y} + ruv dx dy + \int_{\partial\Omega_N} cuv ds$$

and the linear form

$$L(v) = \iint_{\Omega} fv dx dy + \int_{\partial\Omega_N} g_N v ds.$$

Then the variational form of the problem is to find the unique $u \in \tilde{u}_D + H_{0, \partial\Omega_D}^1(\Omega)$ that satisfies

$$B(u, v) = L(v) \quad \forall v \in H_{0, \partial\Omega_D}^1(\Omega).$$

The energy norm of $v \in H_{0, \partial\Omega_D}^1$ is defined by $\|v\|_{E(\Omega)}^2 = B(v, v)$.

The finite element space is defined by partitioning Ω into a grid (or mesh), G_{hp} , consisting of a set of N_T triangular elements, $\{T_i\}_{i=1}^{N_T}$ with $\bar{\Omega} = \cup_{i=1}^{N_T} \bar{T}_i$. If a vertex of a triangle is contained in the interior of an edge of another triangle, it is called

a hanging node. We only consider compatible grids with no hanging nodes, i.e. $\bar{T}_i \cap \bar{T}_j$, $i \neq j$, is either empty, a common edge, or a common vertex. The diameter of the element is denoted h_i . With each element we associate an integer degree $p_i \geq 1$. The finite element space V_{hp} is the space of continuous piecewise polynomial functions on Ω such that over element T_i it is a polynomial of degree p_i . The degree of an edge is determined by applying either a minimum rule or a maximum rule over G_{hp} in which the edge is assigned the minimum or maximum of the degrees of the adjacent elements, respectively.

Let $V_D = \{v \in V_{hp} : v = 0 \text{ on all } T_i \text{ for which } \bar{T}_i \cap \partial\bar{\Omega}_D = \emptyset\}$. Let $\hat{u}_D \in V_D$ be a lift function that approximates the Dirichlet boundary conditions via

- (1) $\hat{u}_D = \hat{u}_v + \hat{u}_e$ where \hat{u}_v is piecewise linear,
- (2) $\hat{u}_v = g_D$ at the element vertices on $\partial\Omega_D$, and
- (3) $\int_{\partial\Omega_D} (\hat{u}_e - (g_D - \hat{u}_v))v = 0 \quad \forall v \in V_D$

The finite element solution is the unique function $u_{hp} \in \hat{u}_D + (V_{hp} \setminus V_D)$ that satisfies

$$B(u_{hp}, v_{hp}) = L(v_{hp}) \quad \forall v_{hp} \in V_{hp} \setminus V_D.$$

The error is defined by $e_{hp} = u - u_{hp}$.

The finite element solution is expressed as a linear combination of basis functions $\{\phi_i\}_{i=1}^N$ that span V_{hp} ,

$$u_{hp} = \sum_{i=1}^N \alpha_i \phi_i(x, y).$$

N is called the number of degrees of freedom. For high order elements, there are a number of basis sets used in practice. A number of the hp strategies of Section 4 rely on the basis being a p -hierarchical basis in which the basis functions for a space of degree p are a subset of the basis functions for a space of degree $p + 1$. In the results of Section 5 the p -hierarchical basis of Szabo and Babuška [1991], which is based on Legendre polynomials, is used.

The discrete form of the problem is a linear system of algebraic equations

$$Ax = b \tag{4}$$

where the matrix A is given by $A_{ji} = B(\phi_i, \phi_j)$ and the right hand side is given by $b_j = L(\phi_j)$.

If h and p are uniform over the grid, $u \in H^m(\Omega)$, and the other usual assumptions are met, then the *a priori* error bound is [Babuška and Suri 1987; Babuška and Suri 1990]

$$\|e_{hp}\|_{H^1(\Omega)} \leq C \frac{h^\mu}{p^{m-1}} \|u\|_{H^m(\Omega)} \tag{5}$$

where $\mu = \min(p, m - 1)$ and C is a constant that is independent of h , p and u , but depends on m .

With a suitably chosen hp mesh, and other typical assumptions, the error can be shown [Guo and Babuška 1986] to converge exponentially in the number of degrees of freedom,

$$\|e_{hp}\|_{H^1(\Omega)} \leq C_1 e^{-C_2 N^{1/3}} \tag{6}$$

for some C_1 and $C_2 > 0$ independent of N .

3. HP-ADAPTIVE REFINEMENT ALGORITHM

```

begin with a very coarse grid
form and solve the linear system
repeat
  determine which elements to refine and whether to refine by  $h$  or  $p$ 
  refine elements
  form and solve the linear system
until some termination criterion is met

```

Fig. 1. Basic form of an hp -adaptive algorithm.

There are many variations on the basic form and details of an hp -adaptive algorithm. In this section we describe the particular algorithm used for the results given in Section 5. Note that some of the hp strategies in Section 4 require a different choice for some of the details, or even a modification of the basic algorithm. These exceptions will be noted in Section 4.

The basic form of the hp -adaptive algorithm is given in Figure 1.

For triangle h -refinement, the newest node bisection method [Mitchell 1991] is used. Briefly, a parent triangle is h -refined by connecting its most recently created vertex to the midpoint of the opposite side to form two new child triangles. This may require first refining a chain of neighbor triangles to maintain compatibility of the grid. p -refinement means increasing the degree of the element by one, followed by enforcing the minimum rule for the edges.

Adaptive refinement is guided by a local *a posteriori* error indicator computed for each element. For element T_i it is given by the approximate solution of a local Neumann residual problem:

$$\mathcal{L}e_i = f - \mathcal{L}u_{hp} \quad \text{in } T_i \quad (7)$$

$$e_i = 0 \quad \text{on } \partial T_i \cap \partial \Omega_D \quad (8)$$

$$\mathcal{B}e_i = g_N - \mathcal{B}u_{hp} \quad \text{on } \partial T_i \cap \partial \Omega_N \quad (9)$$

$$\mathcal{B}e_i = -\frac{1}{2} \left[\frac{\partial u_{hp}}{\partial n} \right] \quad \text{on } (\partial T_i \setminus \partial \Omega_D) \setminus \partial \Omega_N \quad (10)$$

where \mathcal{L} , \mathcal{B} , f , g_N , $\partial \Omega_D$, and $\partial \Omega_N$ are defined in Equations 1-3, $\left[\frac{\partial u_{hp}}{\partial n} \right]$ is the jump in the outward normal derivative of u_{hp} across the element boundary, including the coefficients of the natural boundary conditions, and in Equation 10 \mathcal{B} is modified by setting $c(x, y) = 0$. If the degree of T_i is p_i , the approximate solution, $e_{i, hp}$ of Equations 7-10 is computed using the hierarchical bases of exact degree $p_i + 1$, i.e. omitting the bases of degree 0 through p_i . The error indicator for element T_i is then given by

$$\eta_i = \|e_{i, hp}\|_{E(T_i)}$$

The elements that are refined are those for which $\eta_i / \|u_{hp}\|_{E(\Omega)} > \tau / 10 \sqrt{N_T}$ for a prescribed final error tolerance τ .

The primary criterion for program termination is that the relative energy norm of the error be smaller than τ , i.e., $\|e_{hp}\|_{E(\Omega)} / \|u_{hp}\|_{E(\Omega)} < \tau$. An upper bound on the number of degrees of freedom is used as a secondary criterion to avoid run away programs when convergence is slow.

The method for determining whether an element should be refined by h or by p is called an hp -adaptive strategy. Several strategies have been proposed over the years. Many of them will be described in the next section.

4. THE *HP*-ADAPTIVE STRATEGIES

In this section, the *hp*-adaptive strategies that have been proposed in the literature are briefly described. For brevity, many of the details have been omitted. For a detailed description of the strategies, see [Mitchell and McClain 2011a] or [Mitchell and McClain 2011b]. In some cases, these strategies were developed in the context of 1D problems, rectangular elements, or other settings that are not fully compatible with the context of this paper. In those cases, the strategy is appropriately modified for 2D elliptic PDEs and newest node bisection of triangles.

4.1. Use of *a priori* Knowledge of Solution Regularity

It is well known that for smooth solutions *p*-refinement will produce an exponential rate of convergence, but near singularities *p*-refinement is less effective than *h*-refinement. This is a consequence of the *a priori* error bound in Equation 5. For this reason, many of the *hp* strategies use *h*-refinement in areas where the solution is irregular (i.e., locally fails to be in H^m for some finite m , also called nonsmooth) or nearly irregular, and *p*-refinement elsewhere. The simplest strategy is to use any *a priori* knowledge about irregularities. An *hp*-adaptive strategy of this type was presented by Ainsworth and Senior [1999]. In this approach they simply flag vertices in the initial mesh as being possible trouble spots. During refinement an element is refined by *h* if it contains a vertex that is so flagged, and by *p* otherwise. We will refer to this strategy by the name APRIORI.

4.2. Type parameter

Gui and Babuška [1986] presented an *hp*-adaptive strategy using what they call a type parameter, γ . This strategy is also used by Adjerid, Aiffa and Flaherty [1998]. We will refer to this strategy as TYPEPARAM.

Given the error estimates η_{i,p_i} and η_{i,p_i-1} , define

$$R(T_i) = \begin{cases} \frac{\eta_{i,p_i}}{\eta_{i,p_i-1}} & \eta_{i,p_i-1} \neq 0 \\ 0 & \eta_{i,p_i-1} = 0 \end{cases}$$

By convention, $\eta_{i,0} = 0$, which forces *p*-refinement if $p_i = 1$.

R is used to assess the perceived solution smoothness. Given the type parameter, $0 \leq \gamma < \infty$, element T_i is said to be of *h*-type if $R(T_i) > \gamma$, and of *p*-type if $R(T_i) \leq \gamma$. If element T_i is selected for refinement, then refine it by *h*-refinement if it is of *h*-type and *p*-refinement if it is of *p*-type. Note that $\gamma = 0$ gives pure *h*-refinement and $\gamma = \infty$ gives pure *p*-refinement.

For the results of Section 5, we use $\gamma = 0.3$ if the solution has a singularity, and $\gamma = 0.6$ otherwise.¹

4.3. Estimate Regularity Using Larger *p* Estimates

Ainsworth and Senior [1997] presented a strategy based on an estimate of the regularity using three error estimates based on spaces of degree $p_i + 1$, $p_i + 2$ and $p_i + 3$, so we refer to it as NEXT3P.

The error estimate used to approximate the regularity is a variation on the local Neumann residual error estimate given by Equations 7-10 in which Equation 10 is replaced by

$$\mathcal{B}e_i = g_i \text{ on } (\partial T_i \setminus \partial \Omega_D) \setminus \partial \Omega_N$$

¹The value for this parameter, and the parameters of the other strategies, was determined by a preliminary experiment to determine a single value (or possibly two values dependent on singularity) that generally works best, using a subset of the test problems.

where g_i is an approximation of $\mathcal{B}u$ that satisfies an equilibrium condition. This is the equilibrated residual error estimator in [Ainsworth and Oden 2000].

The local problem is solved on element T_i three times using the spaces of degree $p_i + q$, $q = 1, 2, 3$, to obtain error estimates $e_{i,q}$. In contrast to the local Neumann residual error estimate, the whole space over T_i is used, not just the p -hierarchical bases of degree greater than p_i . These approximations to the error converge to the true solution of the residual problem at the same rate the approximate solution converges to the true solution of Equations 1-3, i.e.

$$\|e_i - e_{i,q}\|_{E(T_i)} \approx C(p_i + q)^{-\alpha}$$

where C and α are positive constants that are independent of q but depend on T_i . Using the Galerkin orthogonality

$$\|e_i - e_{i,q}\|_{E(T_i)}^2 = \|e_i\|_{E(T_i)}^2 - \|e_{i,q}\|_{E(T_i)}^2$$

this can be rewritten

$$\|e_i\|_{E(T_i)}^2 - \|e_{i,q}\|_{E(T_i)}^2 \approx C^2(p_i + q)^{-2\alpha}.$$

We can compute $\|e_{i,q}\|_{E(T_i)}^2$ and $p_i + q$ for $q = 1, 2, 3$ from the approximate solutions, so the three constants $\|e_i\|_{E(T_i)}$, C and α can be approximated by fitting the data. Then, using the *a priori* error estimate in Equation 5, the approximation of the local regularity is $m_i = 1 + \alpha$. Use p -refinement if $p_i \leq m_i - 1$ and h -refinement otherwise.

4.4. Estimate Regularity Using Smaller p Estimates

Another approach that estimates the regularity is given by Süli, Houston and Schwab [2000]. This strategy is based on Equation 5 and an estimate of the convergence rate in p using error estimates based on $p_i - 2$ and $p_i - 1$. We will refer to this strategy as PRIOR2P.

Suppose the error estimate in Equation 5 holds on individual elements and that the inequality is an approximate equality. Let η_{i,p_i-2} and η_{i,p_i-1} be *a posteriori* error estimates for partial approximate solutions over triangle T_i using the bases up to degree $p_i - 2$ and $p_i - 1$, respectively. Then

$$\frac{\eta_{i,p_i-1}}{\eta_{i,p_i-2}} \approx \left(\frac{p_i - 1}{p_i - 2}\right)^{-(m_i-1)}$$

and thus the regularity is estimated by

$$m_i \approx 1 - \frac{\log(\eta_{i,p_i-1}/\eta_{i,p_i-2})}{\log((p_i - 1)/(p_i - 2))}$$

Use p -refinement if $p_i \leq m_i - 1$ and h -refinement otherwise.

4.5. Texas 3 Step

The Texas 3 Step strategy [Bey 1994; Oden and Patra 1995; Oden et al. 1992] first performs h -refinement to get an intermediate grid, and follows that with p -refinement to reduce the error to some given error tolerance, τ . We will refer to this strategy as T3S. Note that for this strategy the basic form of the hp -adaptive algorithm is different than that in Figure 1.

The first step is to create an initial mesh with uniform p and nearly uniform h such that the solution is in the asymptotic range of convergence in h . The second step is to perform adaptive h -refinement to reach an intermediate error tolerance $\gamma\tau$ where γ is a given parameter. In the references, γ is in the range 5 – 10, usually 6 in the numerical results. This intermediate grid is created by computing a desired number of

children for each element T_i by a formula that is based on the *a priori* error estimate in Equation 5. The discrete problem is then solved on the intermediate grid. The third step is to perform adaptive p -refinement to reduce the error to the desired tolerance τ . Again, a formula is used to determine the new degree for each element, p -refinement is performed to increase the degree of each element to the desired degree, and the discrete problem is solved on the final grid.

The strategy of performing all the h -refinement in one step and all the p -refinement in one step is adequate for low accuracy solutions (e.g. 1%), but is not likely to work well with high accuracy solution (e.g. 10^{-6} relative error) [Patra 2009]. We extend the Texas 3 Step strategy to high accuracy by cycling through steps 2 and 3 until the final tolerance τ_{final} is met. τ in the algorithm above is now the factor by which one cycle of steps 2 and 3 should reduce the error. Toward this end, before step 2 the error estimate η_0 is computed for the current grid. The final (for this cycle) and intermediate targets are now given by $\eta_T = \tau\eta_0$ and $\eta_I = \gamma\eta_T$. In the results of Section 5 we use $\tau = 0.1$ and $\gamma = 6$.

4.6. Alternate h and p

This strategy, which will be referred to as ALTERNATE, is a variation on T3S that is more like the algorithm of Figure 1. The difference from T3S is that instead of predicting the number of refinements needed to reduce the error to the next target, the usual adaptive refinement is performed until the target is reached. Thus in step 2 all elements with an error indicator larger than $\eta_I/\sqrt{N_0}$ are h -refined. The discrete problem is solved and the new error estimate compared to η_I . This is repeated until the error estimate is smaller than η_I . Step 3 is similar except adaptive p -refinement is performed and the target is η_T . Steps 2 and 3 are repeated until the final error tolerance is achieved.

4.7. Nonlinear Programming

Patra and Gupta [2001] proposed a strategy for hp mesh design using nonlinear programming. We refer to this strategy as NLP. They presented it in the context of quadrilateral elements with one level of hanging nodes, i.e., an element edge is allowed to have at most one hanging node. Here it is modified slightly for newest node bisection of triangles with no hanging nodes. This is another approach that does not strictly follow the algorithm in Figure 1.

Given a current grid with triangles $\{T_i\}$, degrees $\{p_i\}$, h -refinement levels $\{l_i\}$, and error estimates $\{\eta_i\}$, the object is to determine new mesh parameters $\{\hat{p}_i\}$ and $\{\hat{l}_i\}$, $i = 1..N_T$, by solving an optimization problem which can be informally stated as: minimize the number of degrees of freedom subject to the error being less than a given tolerance and other constraints. Computationally, the square of the error is approximated by $\sum_{i=0}^{N_T} \hat{\eta}_i^2$ where $\hat{\eta}_i$ is an estimate of the error in the refined grid over the region covered by T_i , and the number of degrees of freedom over the children of T_i is $2^{\hat{l}_i - l_i} \hat{p}_i^2 / 2$. Thus

the optimization problem is

$$\text{minimize}_{\{\hat{l}_i\}, \{\hat{p}_i\}} \sum_{i=1}^{N_T} 2^{\hat{l}_i - l_i} \frac{\hat{p}_i^2}{2} \quad (11)$$

$$\text{s.t.} \quad \sum_{i=1}^{N_T} \hat{\eta}_i^2 < \hat{\tau}^2 \quad (12)$$

$$\hat{l}_j - 1 \leq \hat{l}_i \leq \hat{l}_j + 1 \quad \forall j \text{ such that } T_j \text{ shares an edge with } T_i \quad (13)$$

$$0 \leq \hat{l}_i \leq l_{max} \quad (14)$$

$$1 \leq \hat{p}_i \leq p_{max} \quad (15)$$

$$l_i - \Delta l_{dec} \leq \hat{l}_i \leq l_i + \Delta l_{inc} \quad (16)$$

$$p_i - \Delta p_{dec} \leq \hat{p}_i \leq p_i + \Delta p_{inc} \quad (17)$$

where $\hat{\tau}$ is the error tolerance for this refinement phase. Equation 13 is a necessary condition for compatibility of the grid (in [Patra and Gupta 2001] it enforces one level of hanging nodes). Equation 14 insures that coarsening does not go beyond the initial grid, and that the refinement level of an element does not exceed a prescribed limit l_{max} . Similarly, Equation 15 insures that element degrees do not go below one or exceed a prescribed limit p_{max} . Also, because many quantities are only approximate, it is wise to limit the amount of change that occurs to any element during one phase of refinement. Equations 16 and 17 restrict the amount of change that can occur at one time.

Since the \hat{l}_i and \hat{p}_i are naturally integers, the optimization problem is a mixed integer nonlinear program, which is known to be NP-hard. To make the problem tractable, the integer requirement is dropped to give a nonlinear program which can be solved by one of several software packages. For the results in Section 5, we used the program ALGENCAN² Version 2.2.1 [Andreani et al. 2007; Birgin 2005]. Following solution of the nonlinear program, the \hat{l}_i and \hat{p}_i are rounded to the nearest integer.

4.8. Predict Error Estimate on Assumption of Smoothness

Melenk and Wohlmuth [2001] proposed a strategy based on a prediction of what the error should be if the solution is smooth. We call this strategy SMOOTH_PRED.

When refining element T_i , assume the solution is locally smooth and that the optimal convergence rate is obtained. If h -refinement is performed and the degree of T_i is p_i , then the expected error on the two children of T_i is reduced by a factor of $\sqrt{2}^{-p_i}$ as indicated by the *a priori* error estimate of Equation 5. Thus if η_i is the error estimate for T_i , predict the error estimate of the children to be $\gamma_h \eta_i / \sqrt{2}^{-p_i}$ where γ_h is a user specified parameter. If p -refinement is performed on T_i , exponential convergence is expected, so the error should be reduced by some constant factor $\gamma_p \in (0, 1)$, i.e., the predicted error estimate of the p -refinement of T_i is $\gamma_p \eta_i$. When the actual error estimate of a child of T_i or p -refinement of T_i becomes available, it is compared to the predicted error estimate. If the error estimate is less than or equal to the predicted error estimate, then p -refinement is indicated for the child. Otherwise, h -refinement is indicated since presumably the assumption of smoothness was wrong. For the results in Section 5 we use $\gamma_h = 2$ and $\gamma_p = \sqrt{0.4}$.

²The mention of specific products, trademarks, or brand names is for purposes of identification only. Such mention is not to be interpreted in any way as an endorsement or certification of such products or brands by the National Institute of Standards and Technology. All trademarks mentioned herein belong to their respective owners.

4.9. Larger of h -Based and p -Based Error Indicators

In 1D, Schmidt and Siebert [2000] proposed a strategy that uses two *a posteriori* error estimates to predict whether h -refinement or p -refinement will reduce the error more. We extend this strategy to bisected triangles and refer to it as H&P_ERREST.

The local Neumann residual error estimate given by Equations 7-10 is actually an estimate of how much the norm of the solution will change if T_i is p -refined. This is because the approximate solution of the local problem is obtained using the p -hierarchical bases that would be added if T_i is p -refined, so it is an estimate of the actual change that would occur. Using the fact that the current space is a subspace of the refined space and Galerkin orthogonality, it can be shown that

$$\|u - \hat{u}_{hp}\|^2 = \|u - u_{hp}\|^2 - \|\hat{u}_{hp} - u_{hp}\|^2$$

where \hat{u}_{hp} is the solution in the refined space. Thus the change in the solution indicates how much the error will be reduced.

A second error estimate for T_i can be computed by solving a local Dirichlet problem

$$\mathcal{L}e_i = f - \mathcal{L}u_{hp} \quad \text{in } T_i \cup T_i^{\text{mate}} \quad (18)$$

$$e_i = g_D - u_{hp} \quad \text{on } \partial(T_i \cup T_i^{\text{mate}}) \cap \partial\Omega_D \quad (19)$$

$$\mathcal{B}e_i = g_N - \mathcal{B}u_{hp} \quad \text{on } \partial(T_i \cup T_i^{\text{mate}}) \cap \partial\Omega_N \quad (20)$$

$$e_i = 0 \quad \text{on } (\partial(T_i \cup T_i^{\text{mate}}) \setminus \partial\Omega_D) \setminus \partial\Omega_N \quad (21)$$

where T_i^{mate} is the element that is refined along with T_i in the newest node bisection method [Mitchell 1991]. The solution to this problem is approximated by an h -refinement of the two elements using only the new basis functions. The error estimate obtained by taking the norm of this approximate solution is actually an estimate of how much the solution will change, or the error will be reduced, if h -refinement is performed.

Schmidt and Siebert divide the two error estimates by the associated increase in the number of degrees of freedom to obtain an approximate error reduction per degree of freedom. In addition or instead, one of the error estimates can be multiplied by a user specified constant to bias the refinement toward h - or p -refinement. In the results of Section 5 the p -based error estimate is multiplied by 2.

The type of refinement that is used is the one that corresponds to the larger of the two modified error estimates.

4.10. Legendre coefficient strategies

We consider two hp -adaptive strategies that are based on the coefficients in an expansion of the solution in Legendre polynomials. In one dimension, the approximate solution in element T_i with degree p_i can be written

$$u_i(x) = \sum_{j=0}^{p_i} a_j P_j(x)$$

where P_j is the j^{th} degree Legendre polynomial scaled to the interval of element T_i .

Mavriplis [1994] estimates the decay rate of the coefficients by a least squares fit of the last four coefficients a_j to $Ce^{-\sigma j}$. Elements are refined by p -refinement where $\sigma > 1$ and by h -refinement where $\sigma \leq 1$. We refer to this strategy as COEF_DECAY.

Houston et al. [2003] present an approach which uses the Legendre coefficients to estimate the regularity of the solution. The regularity is estimated using the root test

yielding

$$m_i = \frac{\log\left(\frac{2p_i+1}{2a_i^{2p_i}}\right)}{2\log p_i}.$$

If $p_i = 1$, use p -refinement. Otherwise, use p -refinement if $p_i \leq m_i - 1$ and h -refinement if $p_i > m_i - 1$. We refer to this strategy as COEF_ROOT.

Both Mavriplis and Houston et al. presented the strategies in the context of one dimension and use the Legendre polynomials as the local basis so the coefficients are readily available. In [Houston et al. 2003] it is extended to 2D for rectangular elements with a tensor product of Legendre polynomials, and the regularity is estimated in each dimension separately, so the coefficients are still readily available. Eibner and Melenik [2007] extended the COEF_DECAY strategy to quadrisedected triangles with an orthogonal polynomial basis. In this study we are using triangular elements which have a basis that is based on Legendre polynomials [Szabo and Babuška 1991]. In this basis there are $3 + \max(j - 2, 0)$ basis functions of exact degree j over an element, so we don't have a unique Legendre polynomial coefficient to use for each degree. Instead, for the coefficients a_j we use the ℓ_1 norm of the coefficients of the degree j basis functions, i.e.

$$a_j = \sum_{\substack{k \text{ s.t. } \deg(\phi_k)=j \\ \text{supp}(\phi_k) \cap T_i \neq \emptyset}} |\alpha_k|$$

4.11. Reference Solution Strategies

Demkowicz and his collaborators developed an hp -adaptive strategy over a number of years, presented in several papers and books, e.g. [Demkowicz 2007; Demkowicz et al. 2002; Rachowicz et al. 1989; Šolín et al. 2004]. We refer to this strategy as REF-SOLN_EDGE because it relies on computing a reference solution and bases the refinement decisions on edge refinements. Note that for this strategy the basic form of the hp -adaptive algorithm is different than that in Figure 1.

The algorithm is first presented for 1D elliptic problems. Given the current existing mesh, $G_{h,p}$, and current solution, $u_{h,p}$, a uniform refinement in both h and p is performed to obtain a fine mesh $G_{h/2,p+1}$. The equation is solved on the fine mesh to obtain a reference solution $u_{h/2,p+1}$.

The next step is to determine the optimal refinement of each element. This is done by considering a p -refinement and all possible (bisection) h -refinements (i.e., all possible assignments of p to the two children of an h -refinement) that give the same increase in the number of degrees of freedom as the p -refinement. In 1D, this means that the sum of the degrees of the two children must be $p + 1$, resulting in a total of p h -refinements and one p -refinement to be examined. For each possibility, the error decrease rate is computed as

$$\frac{|u_{h/2,p+1} - \Pi_{hp,i} u_{h/2,p+1}|_{H^1(T_i)}^2 - |u_{h/2,p+1} - \Pi_{new,i} u_{h/2,p+1}|_{H^1(T_i)}^2}{N_{new} - N_{hp}}$$

where $\Pi_{hp,i} u_{h/2,p+1}$ is the projection-based interpolant of the reference solution in element T_i , and $\Pi_{new,i}$ is the projection onto the resulting elements from any one of the candidate refinements. The refinement with the largest error decrease rate is selected as the optimal refinement. In the case of h -refinement, the degrees may be increased further by a process known as following the biggest subelement error refinement path, which is also used to determine the guaranteed element rate; see [Demkowicz 2007] for details.

Elements that have a guaranteed rate larger than 1/3 the maximum guaranteed rate are selected for refinement; the factor 1/3 is arbitrary.

The 2D algorithm also begins by computing a reference solution on the globally hp -refined grid $G_{h/2,p+1}$. Then for each edge in the grid, the choice between p - and h -refinement, the determination of the guaranteed edge rate, and the selection of edges to refine are done exactly as in 1D, except that a weighted H^1 seminorm is used instead of the more natural $H^{1/2}$ seminorm which is difficult to work with. In the case of bisected triangles, we only consider edges that would be refined by the bisection of an existing triangle.

The h -refinement of edges determines the h -refinement of elements. It remains to determine the degree of each element. As a starting point, element degrees are assigned to satisfy the minimum rule for edge degrees, using the edge degrees determined in the previous step. Then the biggest subelement error refinement path is followed to determine the guaranteed element rate and assignment of element degrees. We again refer to [Demkowicz 2007] for details. Finally, the minimum rule for edge degrees is enforced.

A related, but simpler, approach was developed by Šolín et al. [2008]. We refer to this strategy as REFSOLN_ELEM since it also begins by computing the reference solution, but bases the refinement on elements. The basic form of the hp -adaptive algorithm is different than that in Figure 1 for this strategy, also.

The local error estimate is given by the norm of the difference between the reference solution and the current solution,

$$\eta_i = \|u_{h/2,p+1} - u_{h,p}\|_{H^1(T_i)}$$

and the elements with the largest error estimates are refined. If T_i is selected for refinement, let $p_0 = \lfloor (p_i + 1)/2 \rfloor$ and consider the following options:

- p -refine T_i to degree $p_i + 1$,
- p -refine T_i to degree $p_i + 2$,
- h -refine T_i and consider all combinations of degrees p_0 , $p_0 + 1$ and $p_0 + 2$ in the children.

In all cases the minimum rule is used to determine edge degrees. In [Šolín et al. 2008], quadrisection of triangles is used leading to 83 options to consider. With bisection of triangles, there are only 11 options. Also, since the object of dividing by two to get p_0 is to make the increase in degrees of freedom from h -refinement comparable to that of p -refinement, we use $p_0 = \lfloor (p_i + 1)/\sqrt{2} \rfloor$ since there are only two children instead of four.

For each candidate, the standard H^1 projection $\Pi_{\text{candidate}}^{H^1(T_i)}$ of $u_{h/2,p+1}$ onto the corresponding space is performed, and the projection error in the H^1 norm, $\zeta_{\text{candidate}}$, is computed,

$$\zeta_{\text{candidate}} = \|u_{h/2,p+1} - \Pi_{\text{candidate}}^{H^1(T_i)} u_{h/2,p+1}\|_{H^1(T_i)}$$

as well as the projection error of the projection onto T_i , ζ_i .

Let N_i be the number of degrees of freedom in the space corresponding to T_i , and $N_{\text{candidate}}$ the number of degrees of freedom in the space corresponding to a candidate. After discarding candidates that seem to be outliers, select the candidate that maximizes

$$\frac{\log \zeta_i - \log \zeta_{\text{candidate}}}{N_{\text{candidate}} - N_i}. \quad (22)$$

Following the refinement that is indicated by the selected candidate, the minimum rule for edge degrees is applied.

This algorithm can be modified slightly to bias the refinement towards or away from p refinement to improve the performance. Given a parameter p_{bias} , multiply the value from Equation 22 by it for all the p -refinement candidates. $p_{\text{bias}} > 1$ will bias the refinement toward doing p -refinement, and $p_{\text{bias}} < 1$ will bias the refinement toward doing h -refinement. For the results in Section 5 we use $p_{\text{bias}} = 2$ for most problems, and $p_{\text{bias}} = 4$ for the analytic, mild wave front and both peak problems, which are the easiest problems.

5. NUMERICAL RESULTS

This section contains the results of a numerical experiment to compare the hp -adaptive strategies' performance on a suite of 20 test problems with various difficulties that adaptive refinement should locate. Two metrics are used to compare the strategies: the number of degrees of freedom required to reach a given error tolerance, and the computation time required to reach the tolerance. The test problems and numerical results for each problem are given in Section 5.1, with a summary of the results in Section 5.2.

The full details of the test problems can be found in [Mitchell 2013]. Here we just give a brief description of each problem. Recall that Poisson's equation is $u_{xx} + u_{yy} = f(x, y)$ and Laplace's equation is Poisson's equation with $f = 0$.

Each problem is solved with each hp strategy using the hp -adaptive algorithm of Section 3, except for those strategies that dictate using a variation on that algorithm, as indicated in Section 4. The problems are solved at low accuracy, typically $\tau = 10^{-2}$, and high accuracy, typically $\tau = 10^{-6}$. At the end of each run the number of degrees of freedom and total "wall clock" time to solution are recorded.

The results are given in bar charts in Figures 2–41. The gray bars indicate the number of degrees of freedom required to reach the tolerance, and the black bars indicate the computation time required to reach the tolerance. All results are scaled by the value of the strategy that performed best, so, for example, a value of 1.0 indicates the best strategy, and a value of 0.2 indicates the strategy needed five times as many degrees of freedom or took five times longer than the best strategy. Instances where both of the bars are missing indicate cases where the strategy was unable to achieve the given tolerance within the allotted memory and time resources.

These computations were performed using the adaptive finite element code PHAML Version 1.11.1 [Mitchell 2012] run as a sequential code on an Intel Xeon based computer operating under the CentOS 6.4 release of Linux with kernel 2.6.32-358.6.2.el6.x86_64. PHAML was compiled with the Intel Fortran compiler version 11.1.

5.1. Test Problems and Convergence Graphs

Analytic Solution. The analytic problem in [Mitchell 2013] is Poisson's equation on the unit square with Dirichlet boundary conditions. The solution is the polynomial

$$2^{4p} x^p (1-x)^p y^p (1-y)^p$$

with $p = 10$. 2^{4p} is a normalization factor so that the L^∞ norm is 1.0. For the APRIORI strategy, we choose to always refine by p , i.e., it is just p -adaptive refinement. Results are shown in Figures 2–3.

Reentrant Corner, Nearly Straight. For elliptic partial differential equations, a reentrant (concave) corner in the domain, with interior angle ω , causes a point singularity that behaves like r^α where r is the distance from the corner and $\alpha = \pi/\omega$. The larger ω is, the stronger the singularity. All the reentrant corner problems are Laplace's equation with Dirichlet boundary conditions on $(-1, 1) \times (-1, 1)$ with the




Fig. 2. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the analytic problem.

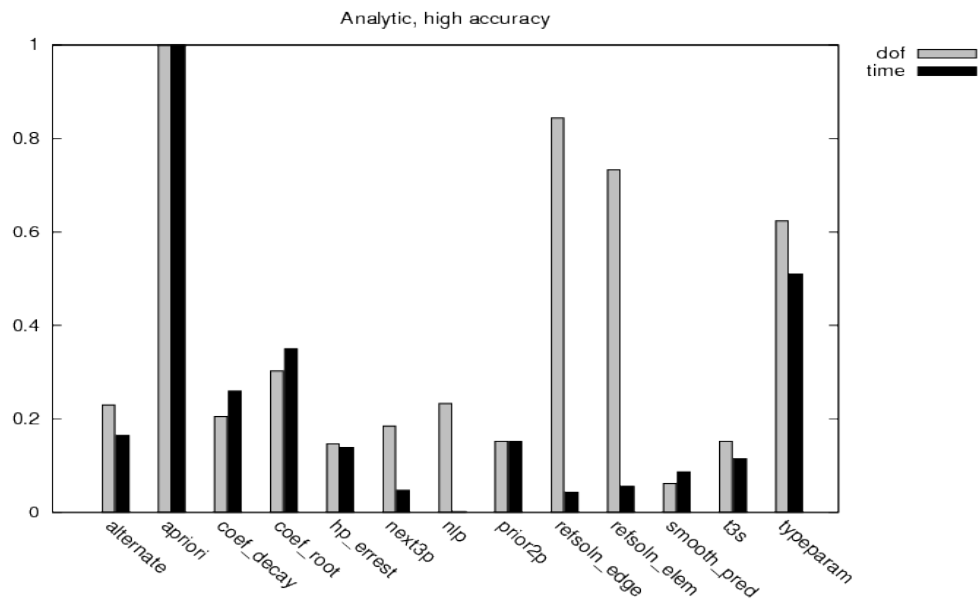


Fig. 3. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the analytic problem.

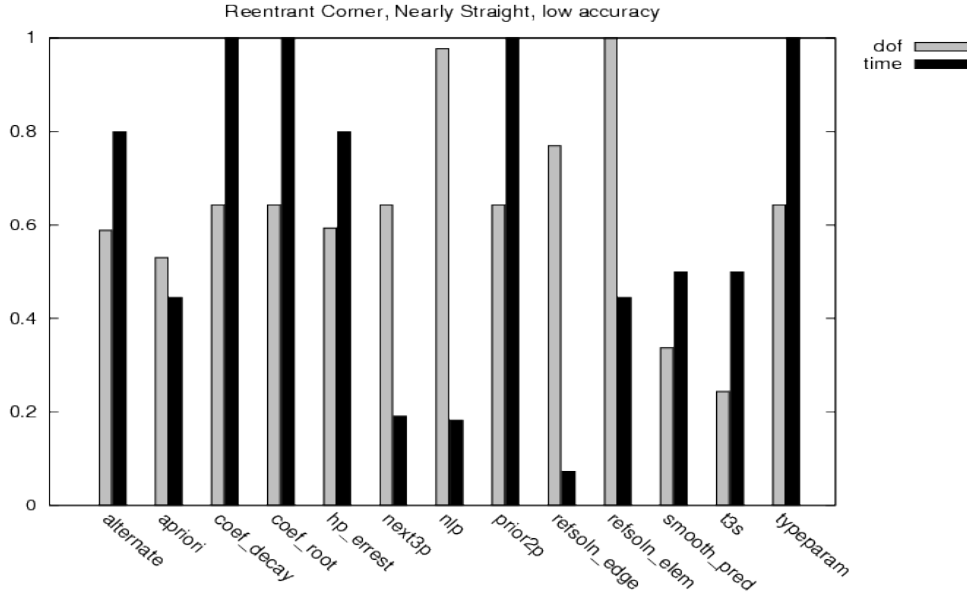


Fig. 4. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-4}$) solution of the nearly straight reentrant corner problem.

section from $2\pi - \omega$ to 2π removed. The solution is

$$r^\alpha \sin(\alpha\theta)$$

where $r = \sqrt{x^2 + y^2}$ and $\theta = \tan^{-1}(y/x)$.

For the nearly straight reentrant corner, $\omega = \pi + .01$. This is a very mild singularity. In all the reentrant corner test problems, the APRIORI strategy refines by h if the element contains the origin and by p otherwise. Results are shown in Figures 4–5.

Reentrant Corner, Wide Angle. This is the reentrant corner problem with $\omega = 5\pi/4$. Results are shown in Figures 6–7.

Reentrant Corner, L-Shaped Domain. The reentrant corner problem with $\omega = 3\pi/2$ is the classic “L domain” problem which is used as an example in many papers on adaptive grid refinement. Results are shown in Figures 8–9.

Reentrant Corner, Narrow Angle. This is the reentrant corner problem with $\omega = 7\pi/4$. Results are shown in Figures 10–11.

Reentrant Corner, Slit. This is the reentrant corner problem with $\omega = 2\pi$. This results in a domain that has a slit along the positive x axis. Results are shown in Figures 12–13.

Linear Elasticity, Mode 1. The linear elasticity problem is a coupled system of two equations with a mixed derivative in the coupling term and different coefficients on the second order x and y terms. The domain is a square with a slit, as in the reentrant corner slit domain problem. The boundary conditions are Dirichlet. For further details, see [Mitchell 2013]. We consider two solutions, referred to as mode 1 and mode 2, by using different boundary conditions. Both solutions have a singularity at the origin, with the mode 1 solution having the stronger singularity. For both problems, the APRIORI strategy refines by h if the element contains the origin and by p otherwise. Results are shown in Figures 14–15.

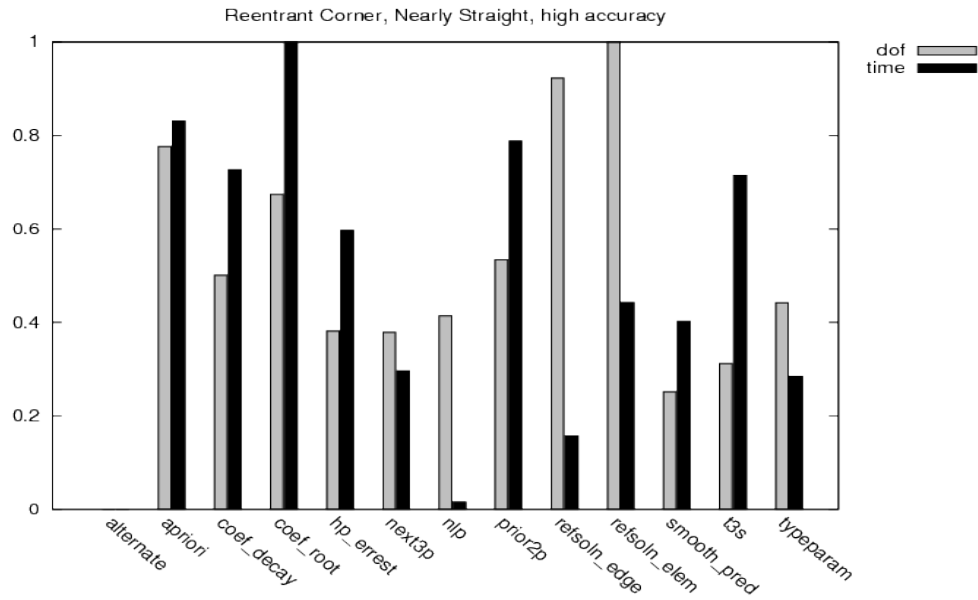


Fig. 5. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-7}$) solution of the nearly straight reentrant corner problem.

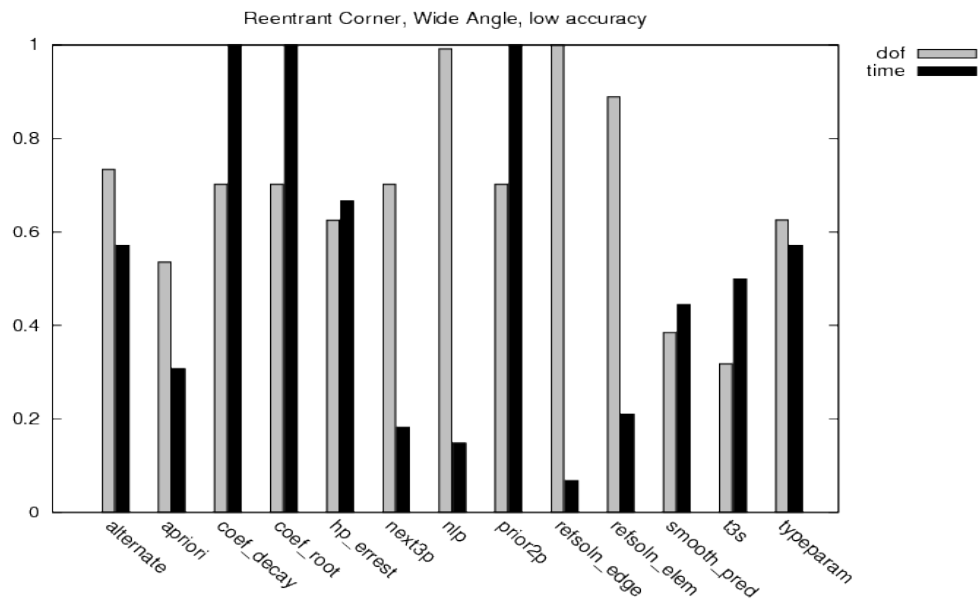


Fig. 6. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the wide angle reentrant corner problem.




Fig. 7. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the wide angle reentrant corner problem.




Fig. 8. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the L-shaped domain problem.




Fig. 9. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the L-shaped domain problem.

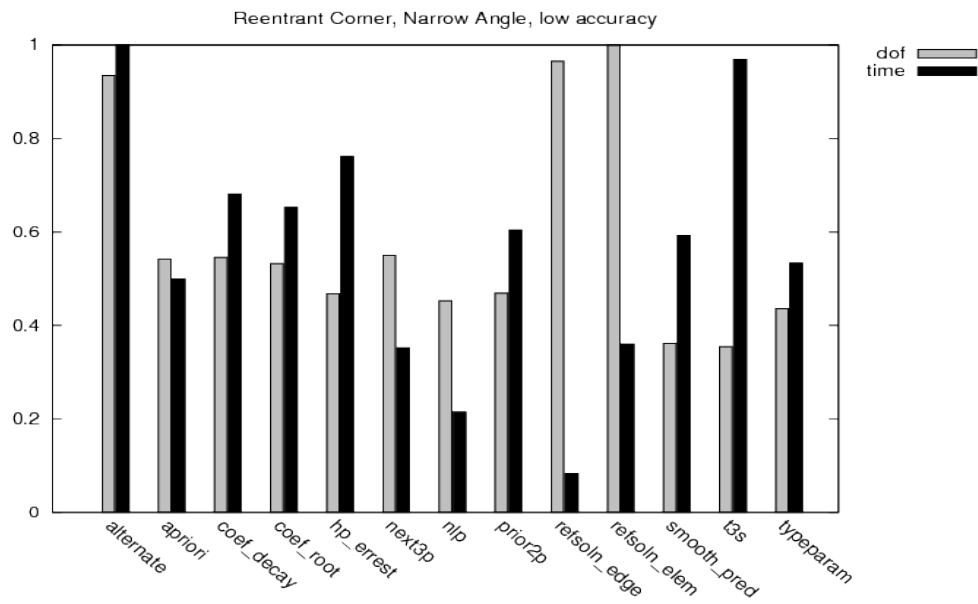


Fig. 10. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the narrow angle reentrant corner problem.

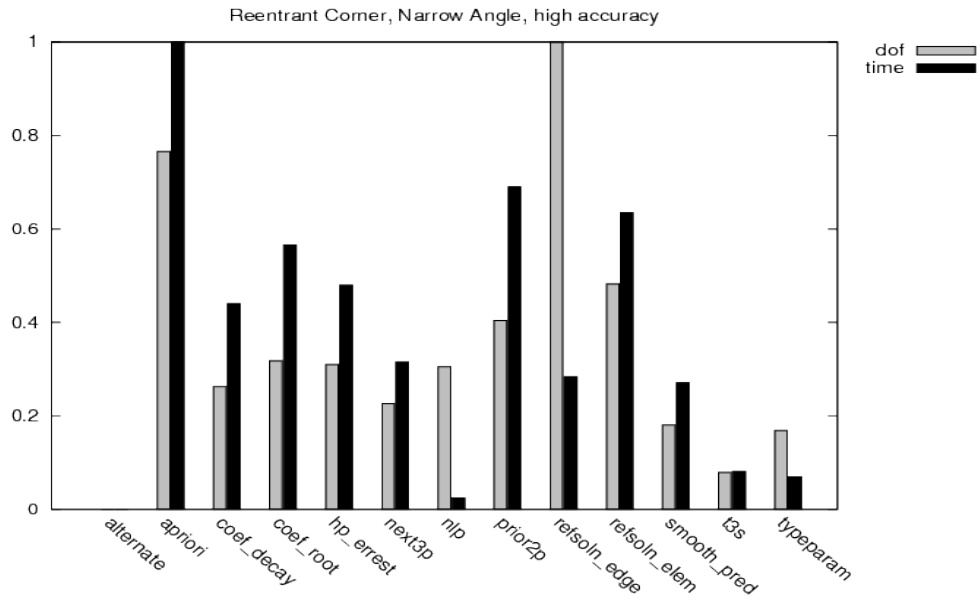


Fig. 11. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the narrow angle reentrant corner problem.

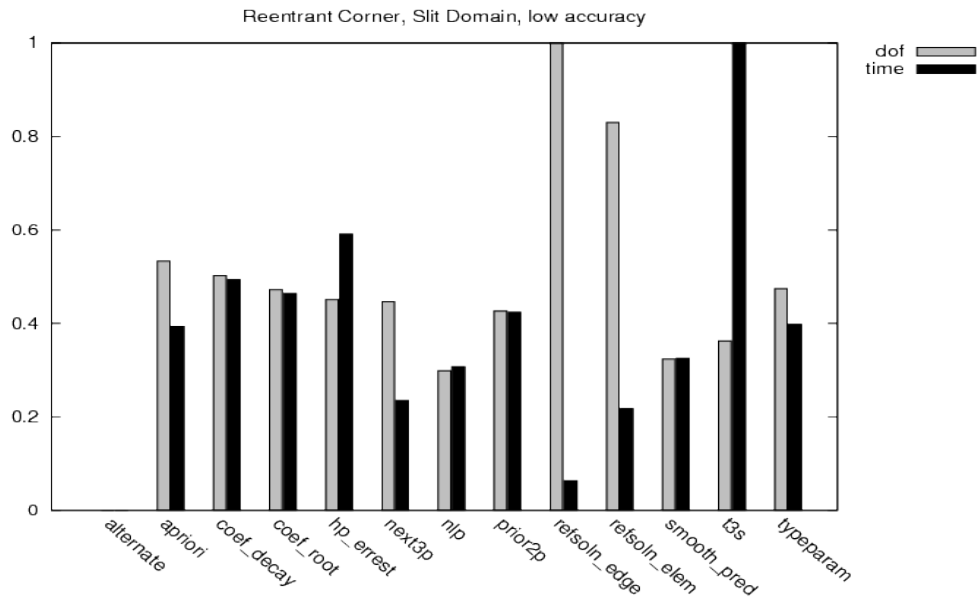


Fig. 12. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the slit domain problem.




Fig. 13. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the slit domain problem.

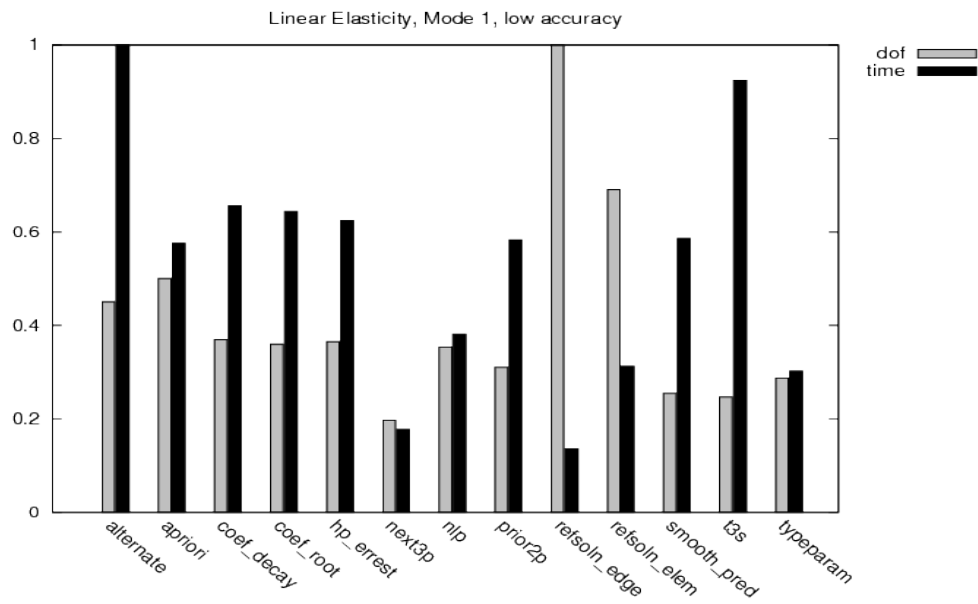


Fig. 14. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the mode 1 linear elasticity problem.




Fig. 15. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the mode 1 linear elasticity problem.

Linear Elasticity, Mode 2. This is the mode 2 solution of the linear elasticity problem. Results are shown in Figures 16–17.

Mild Peak. The peak problem contains a Gaussian peak in the interior of the domain. It is Poisson’s equation on the unit square with Dirichlet boundary conditions. The solution is

$$e^{-\alpha((x-x_c)^2+(y-y_c)^2)}$$

where (x_c, y_c) is the location of the peak, and α determines the strength of the peak. For the easy form of this problem, we use $\alpha = 1000$ and $(x_c, y_c) = (0.5, 0.5)$. The APRIORI strategy refines by h if the element touches the center of the peak and by p otherwise. Results are shown in Figures 18–19.

Sharp Peak. This is the hard version of the peak problem with $\alpha = 100000$ and $(x_c, y_c) = (.51, .117)$. Results are shown in Figures 20–21.

Boundary Layer, Mild. The boundary layer problem is a convection-diffusion equation with first order terms and Dirichlet boundary conditions on $(-1, 1) \times (-1, 1)$. The solution is

$$(1 - e^{-(1-x)/\epsilon})(1 - e^{-(1-y)/\epsilon}) \cos(\pi(x + y))$$

where ϵ controls the strength of the boundary layer. In the easy form of this problem we use $\epsilon = 10^{-1}$. In the APRIORI strategy we refine by h if the element touches either of the boundaries with a boundary layer, and by p otherwise. Results are shown in Figures 22–23.

Boundary Layer, Strong. For the hard version of the boundary layer problem we use $\epsilon = 10^{-3}$. Results are shown in Figures 24–25.

Oscillatory, Mild. The oscillatory problem contains several circular waves which get closer together as you approach the origin. The PDE is a Helmholtz equation with

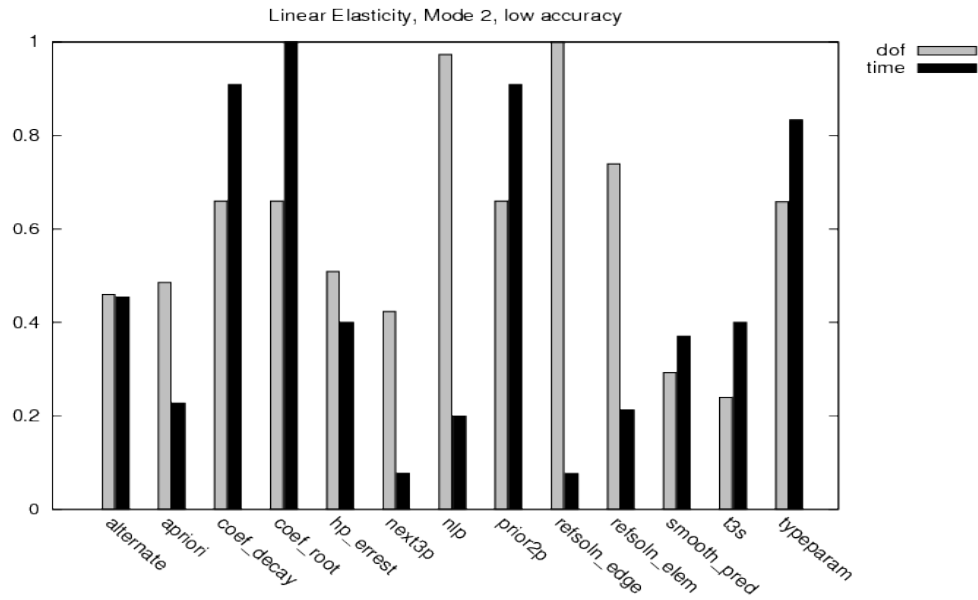


Fig. 16. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the mode 2 linear elasticity problem.

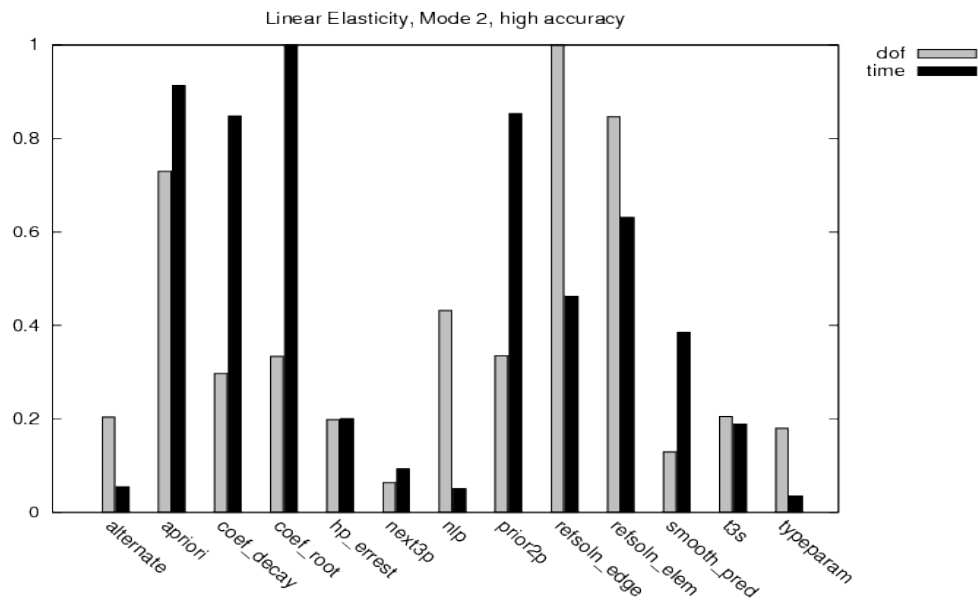


Fig. 17. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the mode 2 linear elasticity problem.

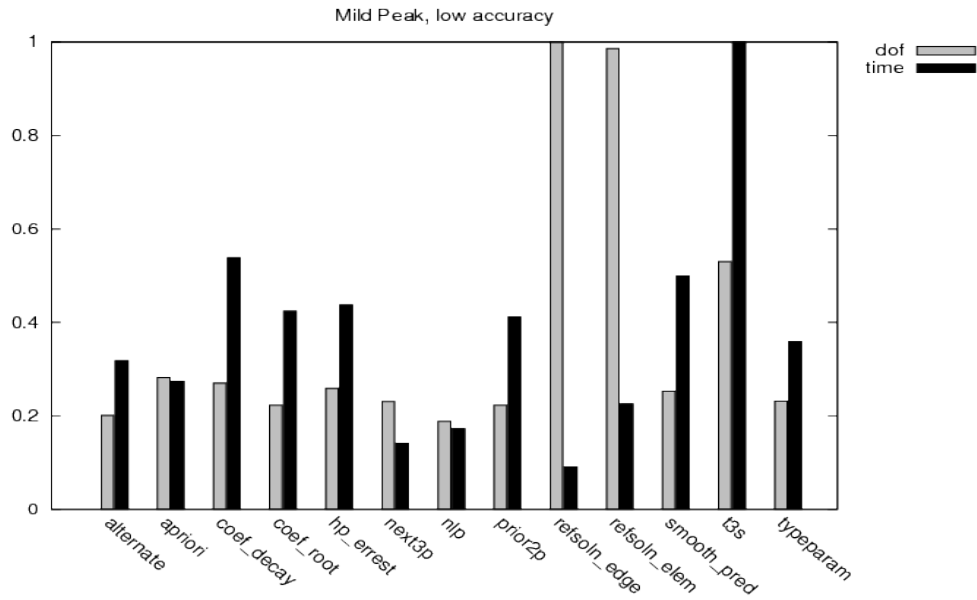


Fig. 18. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the mild peak problem.




Fig. 19. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the mild peak problem.

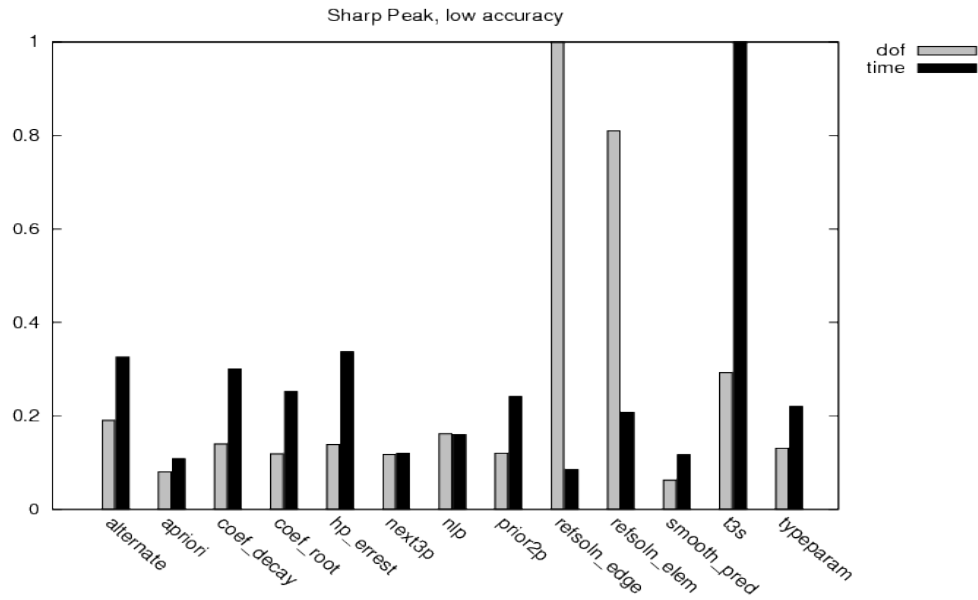


Fig. 20. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the sharp peak problem.

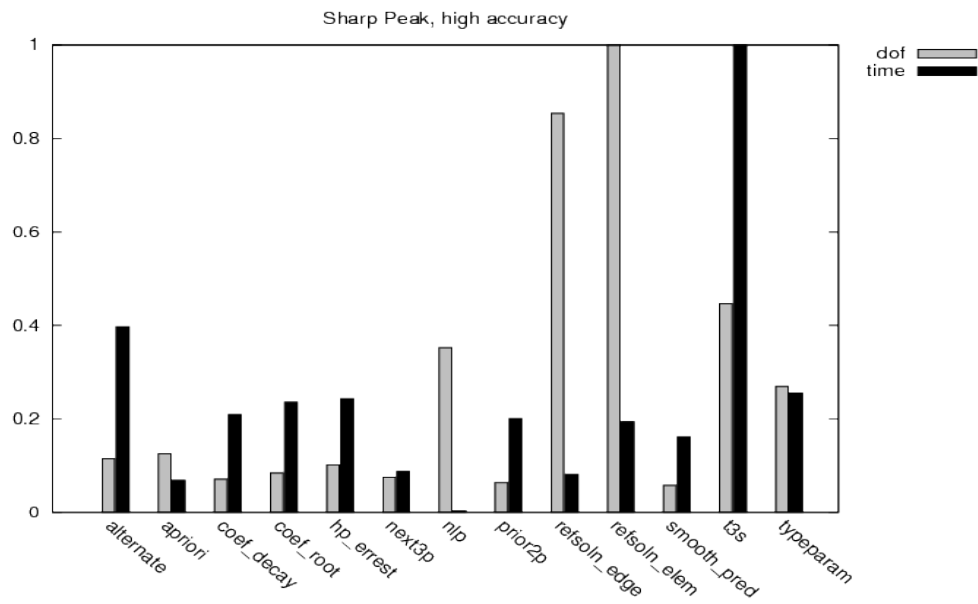


Fig. 21. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the sharp peak problem.




Fig. 22. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the mild boundary layer problem.

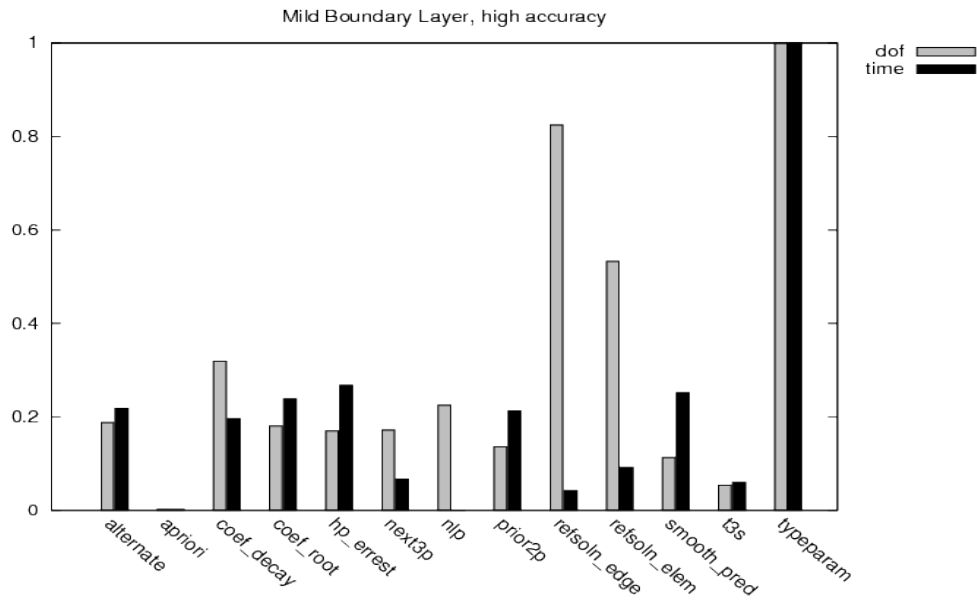


Fig. 23. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the mild boundary layer problem.

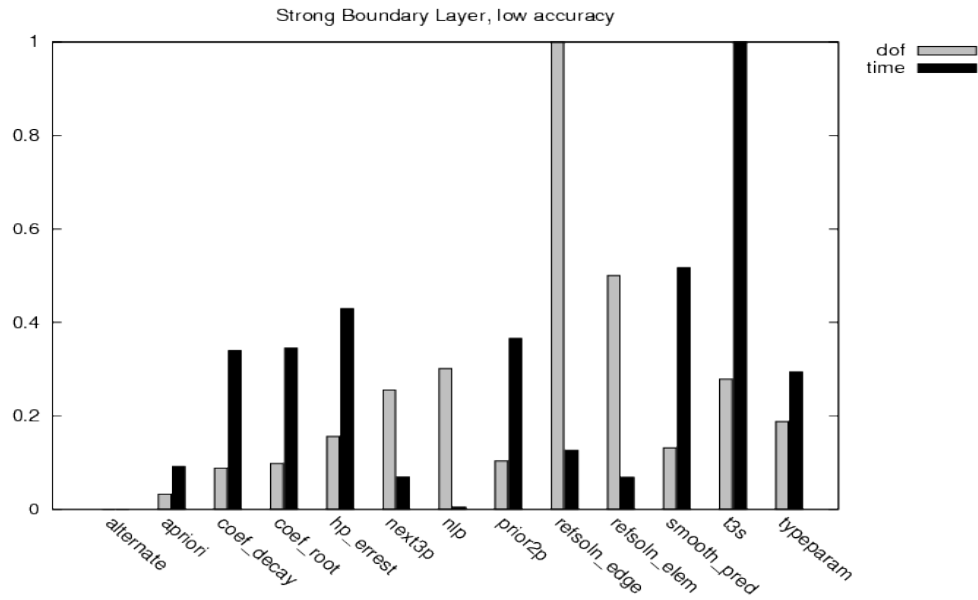


Fig. 24. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the strong boundary layer problem.




Fig. 25. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the strong boundary layer problem.

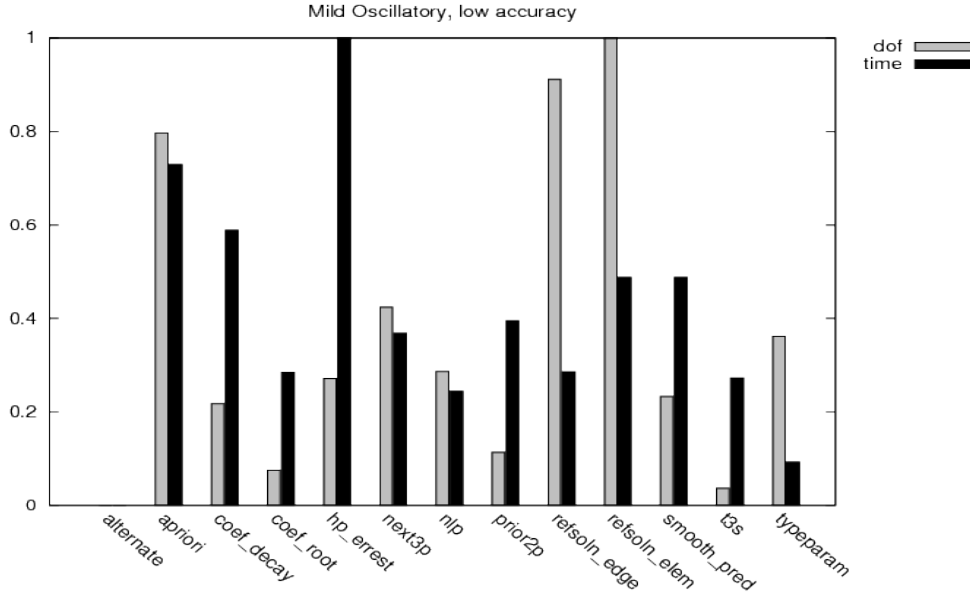


Fig. 26. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the mild oscillatory problem.

Dirichlet boundary conditions on the unit square. The solution is

$$\sin\left(\frac{1}{\alpha + r}\right)$$

where $r = \sqrt{x^2 + y^2}$. The number of oscillations, N , is determined by the parameter $\alpha = \frac{1}{N\pi}$. For the easy form of this problem we use $N = 10.5$. For APRIORI, refine by h if the element touches the origin and by p otherwise. Results are shown in Figures 26–27.

Oscillatory, Strong. For the strong version of the oscillatory problem we use $N = 50.5$. Results are shown in Figures 28–29.

Wave Front, Mild. The circular wave front problem is often used as an example in adaptive grid refinement papers. It is Poisson’s equation with Dirichlet boundary conditions on the unit square. The solution is

$$\tan^{-1}(\alpha(r - r_0))$$

where $r = \sqrt{(x - x_c)^2 + (y - y_c)^2}$. The location of the wave front is defined by a circle with radius r_0 and center (x_c, y_c) . α determines the steepness of the wave front. In addition to the wave front, the solution has a mild singularity at the center of the circle, if the center is located in the closure of the domain. For the easy form of this problem we use $\alpha = 20$, $(x_c, y_c) = (-.05, -.05)$, and $r_0 = 0.7$. The center is chosen outside the domain so that only the wave front is a factor in the adaptivity, not the singularity. With all the wave front problems, for the APRIORI strategy, refine by h if the element touches the circle that defines the location of the wave front and has degree at least 3 (chosen arbitrarily, but works better than degree 1), and by p otherwise. Results are shown in Figures 30–31.




Fig. 27. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-4}$) solution of the mild oscillatory problem.

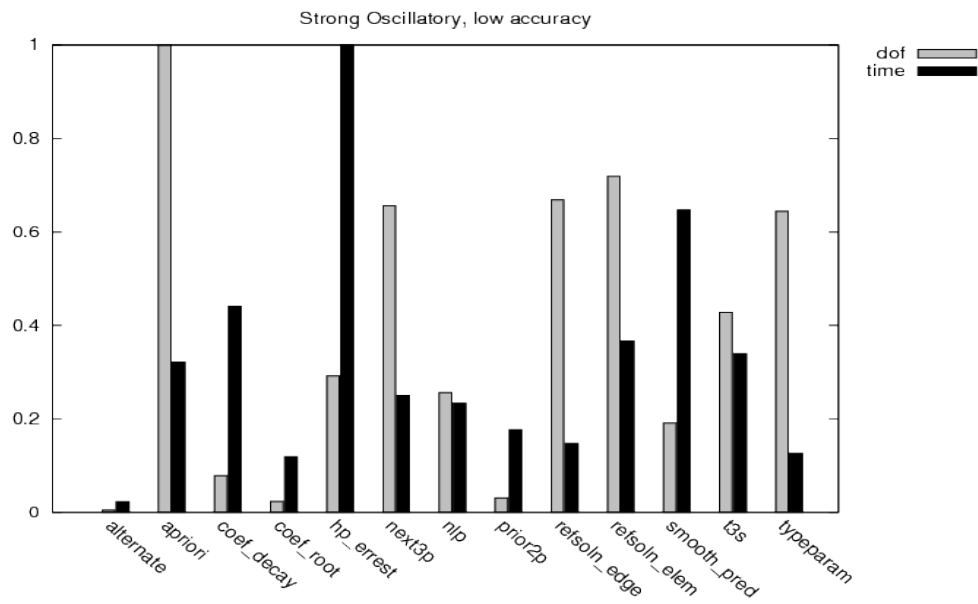


Fig. 28. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the strong oscillatory problem.




Fig. 29. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-4}$) solution of the strong oscillatory problem.

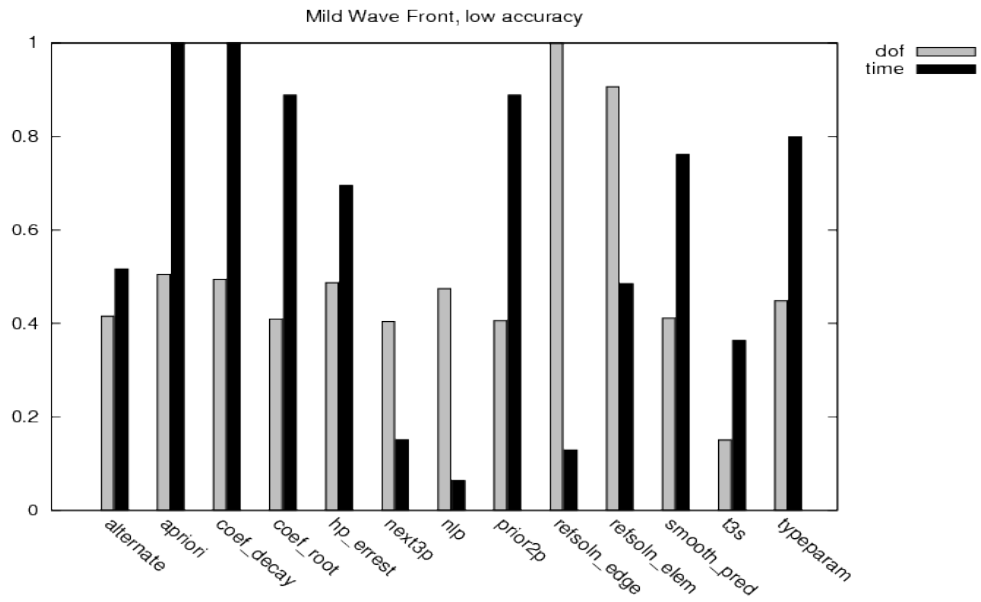


Fig. 30. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the mild wave front problem.

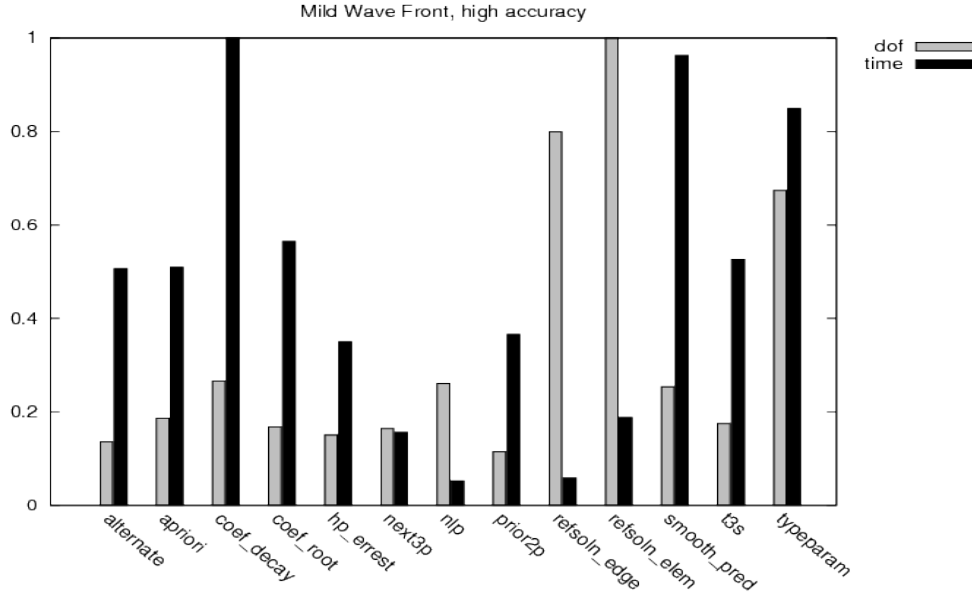


Fig. 31. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the mild wave front problem.

Wave Front, Steep. In the hard version of the wave front problem the location of the wave front is the same, but it is much steeper. The parameters are $\alpha = 1000$, $(x_c, y_c) = (-.05, -.05)$, and $r_0 = 0.7$. Results are shown in Figures 32–33.

Wave Front, Asymmetric. The asymmetric wave front is similar to the steep wave front except the wave front is not symmetric within the domain. The parameters are $\alpha = 1000$, $(x_c, y_c) = (1.5, .25)$, and $r_0 = .92$. Results are shown in Figures 34–35.

Singular Well. This is the wave front problem with the center of the circle placed at the center of the domain and a relatively mild wave front, effectively creating a well with a mild singularity at the center. $\alpha = 50$, $(x_c, y_c) = (.5, .5)$, and $r_0 = .25$. For the APRIORI strategy, refine by h if the element touches the circle that defines the location of the wave front and has degree at least 3, or touches the center of the circle, and by p otherwise. Results are shown in Figures 36–37.

Intersecting Interfaces. The intersecting interfaces problem has piecewise constant coefficients which create a very strong singularity at the center of the domain and discontinuous derivatives along the x and y axes. The boundary conditions are Dirichlet on the domain $(-1, 1) \times (-1, 1)$. For the APRIORI strategy, refine by h if the element touches the origin and by p otherwise. Results are shown in Figures 38–39.

Multiple Difficulties. The multiple difficulties problem combines several of the difficulties of the other problems into a single problem. It contains a reentrant corner point singularity, wave front, peak and boundary layer. For the selected parameters, the peak falls on the wave front, and the wave front intersects the boundary layer and point singularity. The parameters are:

- reentrant corner $\omega = 3\pi/2$
- center of circle for wave front $(0, -3/4)$
- radius of circle for wave front $3/4$
- strength of wave front $\alpha = 200$




Fig. 32. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the steep wave front problem.

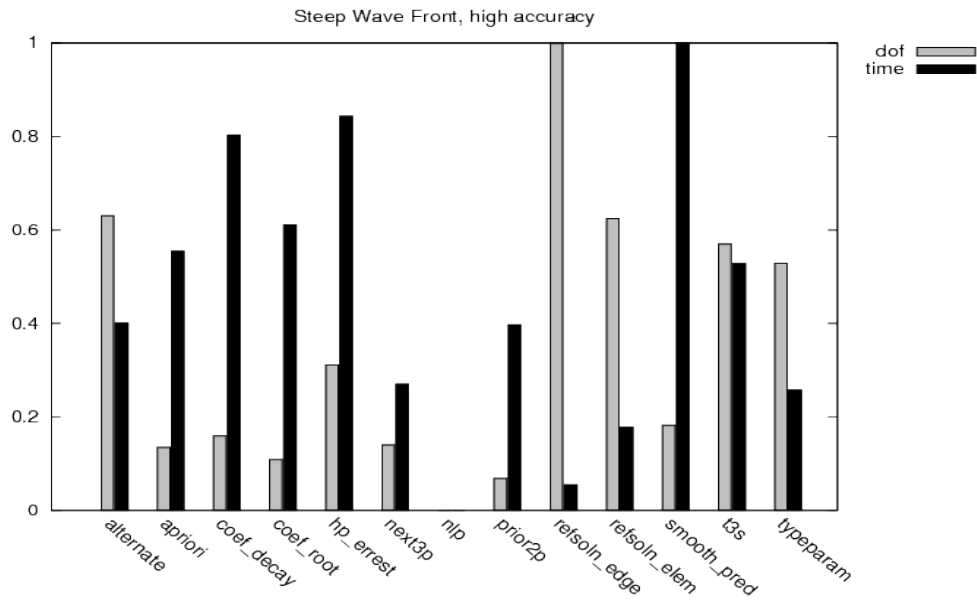


Fig. 33. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the steep wave front problem.

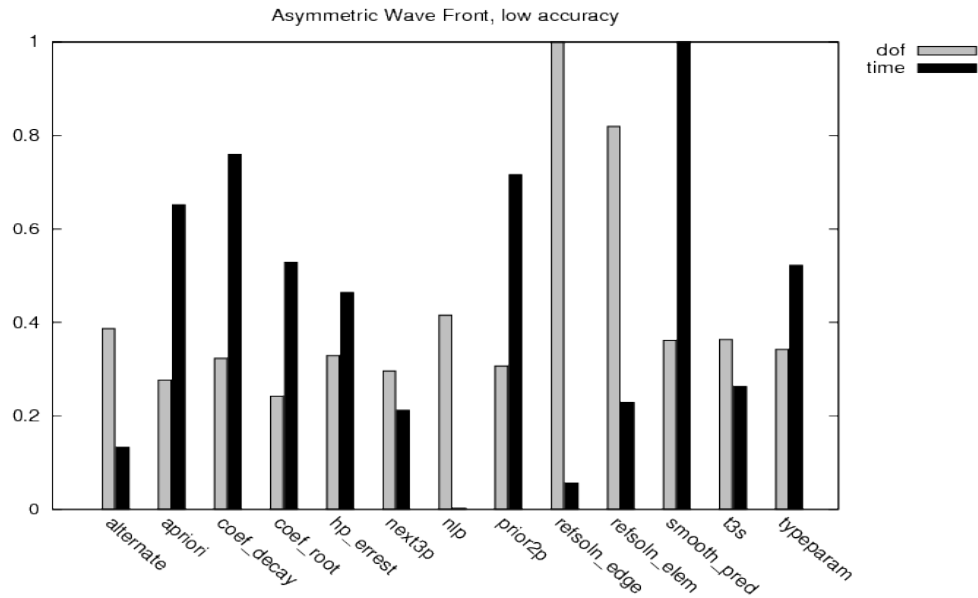


Fig. 34. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the asymmetric wave front problem.

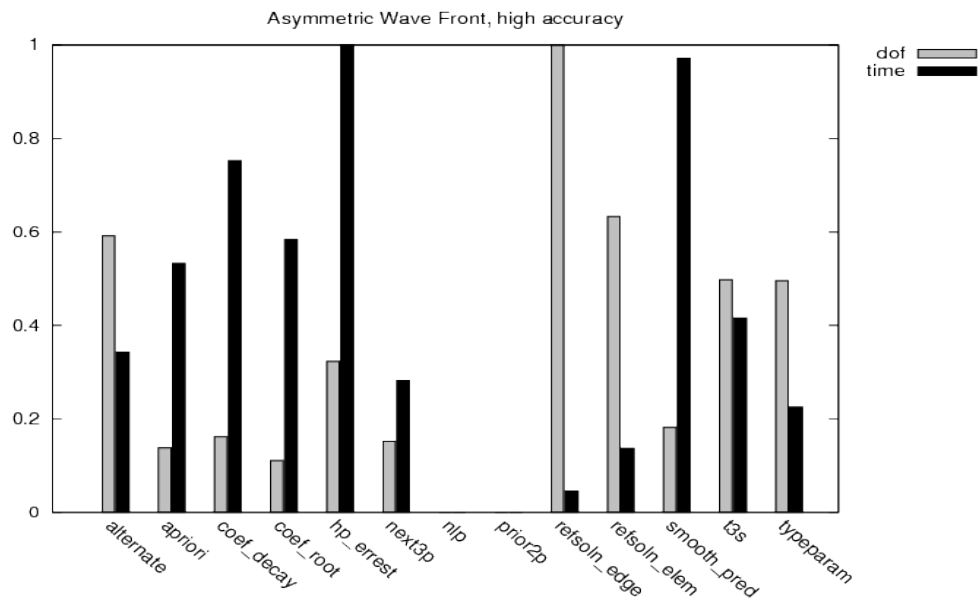


Fig. 35. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the asymmetric wave front problem.

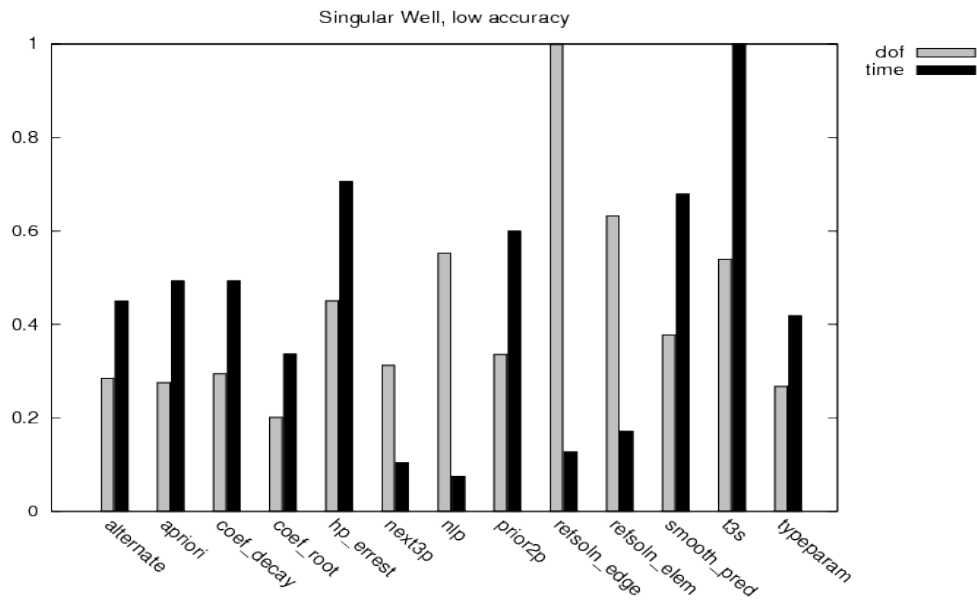


Fig. 36. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the singular well problem.




Fig. 37. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the singular well problem.

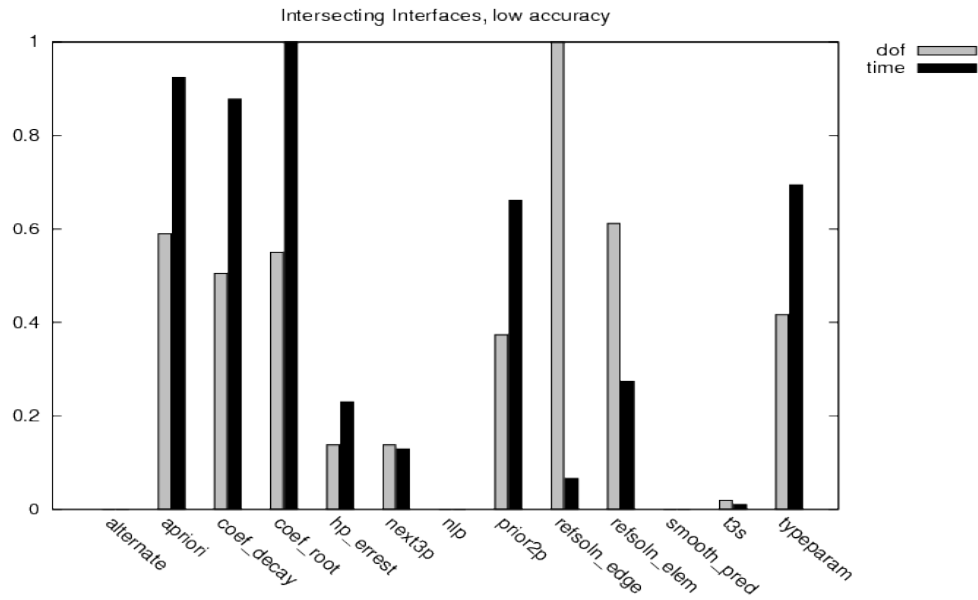


Fig. 38. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-1}$ solution of the intersecting interfaces problem).




Fig. 39. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-2}$) solution of the intersecting interfaces problem.




Fig. 40. Relative performance of the strategies in degrees of freedom and wall clock time for low accuracy ($\tau = 10^{-2}$) solution of the multiple difficulties problem.

Table I. Number of problems for which each strategy required less than twice as many degrees of freedom as the best performing strategy.

strategy	low accuracy			high accuracy		
	easy	hard	singular	easy	hard	singular
ALTERNATE	0	0	5	0	2	1
APRIORI	3	1	7	2	1	8
COEF_DECAY	1	0	7	0	0	2
COEF_ROOT	2	0	6	0	0	2
H&P_ERREST	1	0	4	0	0	0
NEXT3P	2	1	4	0	1	0
NLP	2	0	4	0	0	0
PRIOR2P	2	0	4	0	0	1
REFSOLN_EDGE	5	5	10	5	5	9
REFSOLN_ELEM	5	5	10	5	4	8
SMOOTH_PRED	0	0	0	0	0	1
T3S	1	0	1	0	2	0
TYPEPARAM	2	1	3	3	2	2

- center of peak $(\sqrt{5}/4, -1/4)$
- strength of peak $\alpha = 1000$
- strength of boundary layer $\epsilon = 1/100$

The APRIORI method refines by h in the same cases as it did in the individual problems. Results are shown in Figures 40–41.

5.2. Summary and Observations

In this section, we summarize the results in Section 5.1 to examine the relative performance of the strategies in different situations. The test problems are grouped into six

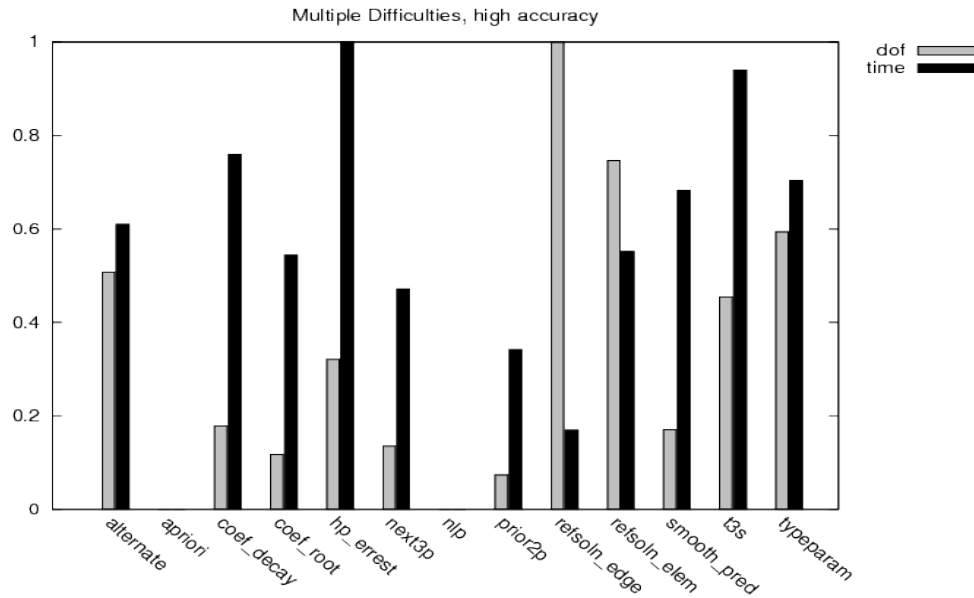


Fig. 41. Relative performance of the strategies in degrees of freedom and wall clock time for high accuracy ($\tau = 10^{-6}$) solution of the multiple difficulties problem.

Table II. Number of problems for which each strategy required less than twice as much computation time as the best performing strategy.

strategy	low accuracy			high accuracy		
	easy	hard	singular	easy	hard	singular
ALTERNATE	1	0	5	1	0	1
APRIORI	3	2	4	3	2	8
COEF_DECAY	5	2	8	3	4	6
COEF_ROOT	3	2	7	2	3	9
H&P_ERREST	3	2	8	2	4	5
NEXT3P	0	0	0	1	0	0
NLP	0	0	0	0	0	0
PRIOR2P	3	2	8	0	1	7
REFSOLN_EDGE	0	0	0	0	0	1
REFSOLN_ELEM	0	0	0	1	1	5
SMOOTH_PRED	2	4	6	2	4	3
T3S	1	3	7	3	4	3
TYPEPARAM	3	2	7	4	0	1

categories: easy problems, hard problems, and singular problems at low accuracy and high accuracy.

While the reader may draw personal conclusions from the data given in Section 5.1, we use the following approach to attempt to summarize those results. We will consider a method to have performed well on a given problem if it is no more than a factor of two worse than the best performing method for that problem, i.e., no more than twice as many degrees of freedom or twice as much computation time. In tables I and II we give the number of problems in each of the six categories for which each strategy performed well in terms of degrees of freedom and computation time, respectively. Bear in mind there are five easy problems, five hard problems, and ten singular problems. Based on these numbers, we make the following observations.

The two reference solution strategies are clearly superior to the other strategies in terms of degrees of freedom, except APRIORI also does well on singular problems. However, the reference solution strategies perform poorly in terms of computation time. Among the other strategies, most perform reasonably well in many categories in terms of computation time, except ALTERNATE, NEXT3P and NLP. COEF_DECAY seems to have the best overall performance across categories. TYPEPARAM does well with easy problems and SMOOTH_PRED does well with hard problems. APRIORI, COEF_ROOT, H&P_ERREST and PRIOR2P do well on singular problems.

6. CONCLUSION AND FUTURE WORK

In this paper we presented the results of a study of strategies for the hp -adaptive finite element method for 2D linear elliptic partial differential equations using newest node bisection of triangles. The hp -strategies are methods for determining how to select between the different possibilities of h - and p -refinement. Thirteen strategies were described and compared in a numerical experiment using 20 test problems. Two metrics for comparison were used: the relative energy norm of the error vs. the number of degrees of freedom and vs. computation time.

We found that the REFSOLN_EDGE and REFSOLN_ELEM strategies performed best in degrees of freedom, and are comparable to each other. However, they are considerably more expensive than the other strategies, except NLP, in computation time. These strategies may be useful for situations where one desires an optimal grid that will be used many times. For problems with known point singularities and no other significant features, APRIORI appears to be the less expensive method of choice. COEF_DECAY appears to be the best choice as a general strategy across all categories of problems, whereas many of the other strategies perform well in particular categories and are reasonable in general.

Since the determination of what strategies to include in this study, other strategies have come to our attention or have come into existence [Bank and Nguyen 2011; Bürg and Dörfler 2011; Wihler 2011]. For future work we will extend the results of this study to include additional strategies as they are discovered. Also, we hope to use the lessons learned from this study to develop a better general purpose hp -strategy. For example, is it possible to get the excellent convergence performance of the reference solution strategies without the expense of computing the reference solution by combining some aspects of the reference solution strategies with some aspects of other strategies? Our conclusion is that, at this time, there is still much opportunity for the development of a general purpose hp -adaptive strategy that is both efficient and effective.

REFERENCES

- ADJERID, S., AIFFA, M., AND FLAHERTY, J. 1998. Computational methods for singularly perturbed systems. In *Singular Perturbation Concepts of Differential Equations*, J. Cronin and R. O'Malley, Eds. AMS, Providence.
- AINSWORTH, M. AND ODEN, J. T. 2000. *a posteriori Error Estimation in Finite Element Analysis*. John Wiley & Sons, New York.
- AINSWORTH, M. AND SENIOR, B. 1997. An adaptive refinement strategy for h-p finite element computations. *Appl. Numer. Math.* 26, 1-2, 165–178.
- AINSWORTH, M. AND SENIOR, B. 1999. hp -finite element procedures on non-uniform geometric meshes: adaptivity and constrained approximation. In *Grid Generation and Adaptive Algorithms*, M. W. Bern, J. E. Flaherty, and M. Luskin, Eds. Vol. 113. IMA Volumes in Mathematics and its Applications, Springer-Verlag, New York, 1–28.
- ANDREANI, R., BIRGIN, E. G., MARTNEZ, J. M., AND SCHUVERDT, M. L. 2007. On augmented Lagrangian methods with general lower-level constraints. *SIAM J. Optim.* 18, 1286–1309.
- BABUŠKA, I. AND SURI, M. 1987. The h - p version of the finite element method with quasiuniform meshes. *RAIRO Modél. Math. Anal. Numér.* 21, 199–238.

- BABUŠKA, I. AND SURI, M. 1990. The p - and h - p versions of the finite element method, an overview. *Comput. Methods Appl. Mech. Engrg.* 80, 5–26.
- BANK, R. E. AND NGUYEN, H. 2011. hp adaptive finite elements based on derivative recovery and super-convergence. *Comput. Visual. Sci.* 14, 287–299.
- BEY, K. S. 1994. An hp adaptive discontinuous Galerkin method for hyperbolic conservation laws. Ph.D. thesis, University of Texas at Austin, Austin, TX.
- BIRGIN, E. G. 2005. TANGO home page. <http://www.ime.usp.br/~egbirgin/tango/>.
- BÜRG, M. AND DÖRFLER, W. 2011. Convergence of an adaptive hp finite element strategy in higher space-dimensions. *Appl. Numer. Math.* 61, 1132–1146.
- DEMKOWICZ, L. 2007. *Computing with hp-adaptive finite elements, Volume 1, One and two dimensional elliptic and Maxwell problems*. Chapman & Hall/CRC, Boca Raton, FL.
- DEMKOWICZ, L., RACHOWICZ, W., AND DEVLOO, P. 2002. A fully automatic hp -adaptivity. *J. Sci. Comput.* 17, 127–155.
- EIBNER, T. AND MELENK, J. M. 2007. An adaptive strategy for hp -FEM based on testing for analyticity. *Comput. Mech.* 39, 5, 575–595.
- GUI, W. AND BABUŠKA, I. 1986. The h , p and h - p versions of the finite element method in 1 dimension. Part 3: The adaptive h - p version. *Numer. Math.* 49, 659–683.
- GUO, B. AND BABUŠKA, I. 1986. The h - p version of the finite element method. Part 1: The basic approximation results. *Comput. Mech.* 1, 21–41.
- HOUSTON, P., SENIOR, B., AND SÜLI, E. 2003. Sobolev regularity estimation for hp -adaptive finite element methods. In *Numerical Mathematics and Advanced Applications*, F. Brezzi, A. Buffa, S. Corsaro, and A. Murli, Eds. Springer-Verlag, Berlin, 619–644.
- MAVRIPLIS, C. 1994. Adaptive mesh strategies for the spectral element method. *Comput. Methods Appl. Mech. Engrg.* 116, 77–86.
- MELENK, J. M. AND WOHLMUTH, B. I. 2001. On residual-based a-posteriori error estimation in hp -FEM. *Adv. Comput. Math.* 15, 311–331.
- MITCHELL, W. F. 1991. Adaptive refinement for arbitrary finite element spaces with hierarchical bases. *J. Comput. Appl. Math.* 36, 65–78.
- MITCHELL, W. F. 2012. PHAML home page. <http://math.nist.gov/phaml>.
- MITCHELL, W. F. 2013. A collection of 2D elliptic problems for testing adaptive grid refinement algorithms. *Appl. Math. Comput.*, to appear.
- MITCHELL, W. F. AND MCCLAIN, M. A. 2011a. A comparison of hp -adaptive strategies for elliptic partial differential equations (long version). NISTIR 7824, National Institute of Standards and Technology.
- MITCHELL, W. F. AND MCCLAIN, M. A. 2011b. A survey of hp -adaptive strategies for elliptic partial differential equations. In *Recent Advances in Computational and Applied Mathematics*, T. E. Simos, Ed. Springer, 227–258.
- ODEN, J. T. AND PATRA, A. 1995. A parallel adaptive strategy for hp finite element computations. *Comput. Methods Appl. Mech. Engrg.* 121, 449–470.
- ODEN, J. T., PATRA, A., AND FENG, Y. 1992. An hp adaptive strategy. In *Adaptive Multilevel and Hierarchical Computational Strategies*, A. K. Noor, Ed. Vol. 157. ASME Publication, 23–46.
- PATRA, A. 2009. private communication.
- PATRA, A. AND GUPTA, A. 2001. A systematic strategy for simultaneous adaptive hp finite element mesh modification using nonlinear programming. *Comput. Methods Appl. Mech. Engrg.* 190, 3797–3818.
- RACHOWICZ, W., ODEN, J. T., AND DEMKOWICZ, L. 1989. Toward a universal h - p adaptive finite element strategy, Part 3. Design of h - p meshes. *Comput. Methods Appl. Mech. Engrg.* 77, 181–212.
- SCHMIDT, A. AND SIEBERT, K. G. 2000. a posteriori estimators for the $h - p$ version of the finite element method in 1D. *Appl. Numer. Math.* 35, 43–66.
- ŠOLÍN, P., ČERVENÝ, J., AND DOLEŽEL, I. 2008. Arbitrary-level hanging nodes and automatic adaptivity in the hp -FEM. *Math. Comput. Simulation* 77, 117–132.
- ŠOLÍN, P., SEGETH, K., AND DOLEŽEL, I. 2004. *Higher-Order Finite Element Methods*. Chapman & Hall/CRC, New York.
- SÜLI, E., HOUSTON, P., AND SCHWAB, C. 2000. hp -finite element methods for hyperbolic problems. In *The Mathematics of Finite Elements and Applications X. MAFELAP*, J. Whiteman, Ed. Elsevier, 143–162.
- SZABO, B. AND BABUŠKA, I. 1991. *Finite Element Analysis*. John Wiley and Sons, New York.
- WIHLER, T. P. 2011. An hp -adaptive strategy based on continuous sobolev embeddings. *J. Comput. Appl. Math.* 235, 8, 2731–2739.

Received Month Year; revised Month Year; accepted Month Year