

A Multilevel Approach for Obtaining Locally Optimal Finite Element Meshes

Peter K. Jimack, Rashid Mahmood, Mark A. Walkley and Martin Berzins

Computational PDE Unit, School of Computing

University of Leeds, Leeds LS2 9JT, UK

Abstract

In this paper we consider the adaptive finite element solution of a general class of variational problems using a combination of node insertion, node movement and edge swapping. The adaptive strategy that is proposed is based upon the construction of a hierarchy of locally optimal meshes starting with a coarse grid for which the location and connectivity of the nodes is optimized. This grid is then locally refined and the new mesh is optimized in the same manner. Results presented indicate that this approach is able to produce better meshes than those possible by more conventional adaptive strategies and in a relatively efficient manner.

1 Introduction

Automatic mesh generation is an important computational tool for the finite element analysis of a wide variety of engineering problems ranging from structural analysis through to computational fluid dynamics for example. For many of these problems the use of *unstructured* grids offers many advantages over structured grids, such as permitting the straightforward triangulation of geometrically complex domains or allowing the mesh density to be adapted according to the behaviour of the solution. In this paper we are concerned mainly with the latter of these properties of unstructured grids: the natural manner in which numerous forms of mesh adaptivity are permitted.

Broadly speaking mesh adaptivity algorithms may be categorised as belonging to one of two classes. The first of these, generally referred to as *h*-refinement, involves adding vertices and elements to

the mesh in some manner. This may be through local refinement (e.g. [10]) or through more global remeshing (e.g. [15]) but has the general aim of increasing the number of vertices and elements in those regions of the domain where some measure of the error is unacceptably high. The second class of approach, often referred to as r -refinement, adapts the mesh in such a way that the number of vertices and elements remains essentially unchanged. This is typically achieved through the use of node movement (e.g. [7, 8]), where the mesh is continuously deformed so as to increase the density of vertices in those regions of the domain with the highest errors, or through the use of edge swapping (e.g. [11]), where the number and position of the vertices is held fixed but the way in which they are connected together is allowed to change. (There is also a third class of adaptive algorithm, known as p -refinement, which involves increasing the degree of the finite element approximation space on a fixed mesh, however we do not consider this approach here. See, for example, [1] or [16] for further details.)

In this paper we present a new hybrid algorithm that combines the local insertion and movement of vertices with the local swapping of edges in order to attempt to obtain optimal finite element meshes for a general class of problem. These are variational problems which may be posed in the following form (or similar, according to the precise nature of the boundary conditions):

$$\min_{u: \Omega(\subset \mathcal{R}^m) \rightarrow \mathcal{R}^n} \int_{\Omega} F(\underline{x}, u, \nabla u) d\underline{x} \quad (1)$$

for some energy density function $F : \mathcal{R}^m \times \mathcal{R}^n \times \mathcal{R}^{m \times n} \rightarrow \mathcal{R}$. Physically this variational form may be used to model problems in linear and nonlinear elasticity, heat and electrical conduction, motion by mean curvature and many more (see, for example, [2],[7],[8]). Throughout the majority of this paper we restrict our attention to the two-dimensional case where $m = 2$, however generalizations to three dimensions ($m = 3$) are considered towards the end.

For variational problems of the form (1), the fact that the exact solution minimises the energy functional provides a natural optimality criterion for the design of computational grids using r -refinement. Indeed, the idea of locally minimising the energy with respect to the location of the vertices of a mesh of fixed topology has been considered by a number of authors (e.g. [5],[14]), as has the approach of locally minimising the energy with respect to the connectivity of a mesh with fixed vertices (e.g. [11]). Accordingly, the specific algorithms that we use for node movement are generalizations of those used in [6] and [14], and the edge swapping is based upon [11]. Full details of these algorithms and how

they are combined with local h -refinement are presented in the following section.

In section 3 it is then demonstrated that combining the above r -refinement and h -refinement approaches in an appropriate manner allows locally optimal grids to be obtained which are better (in terms of energy minimisation) than using either strategy alone. The approach taken is to start with a very coarse mesh which is optimised using r -refinement. This is then refined locally to create a new mesh with a greater number of elements and vertices which can itself be optimised. By repeating this process a number of times a hierarchy of locally optimal meshes is obtained. Since the initial mesh at each level of the hierarchy is produced by local refinement of an optimal mesh at the previous level it follows that this typically provides a reliable starting point when optimising the new mesh. The results presented demonstrate that the proposed hybrid algorithm is able to provide a mesh which allows the solution of (1) to be approximated to an arbitrary error tolerance using substantially fewer vertices and elements than through h -refinement alone. Furthermore, it is also demonstrated that, for a fixed size of mesh, this multilevel approach invariably finds a better locally optimal solution than is obtained by applying r -refinement directly to a regular starting mesh of the same fixed size.

The paper concludes by addressing possible generalizations of the technique to three-dimensional problems and discussing the strengths and weaknesses of the proposed hybrid algorithm for obtaining locally optimal meshes in two dimensions.

2 Multilevel adaptivity

In this section we consider the adaptive finite element solution of problems of the form (1), first using r -refinement and then adding h -refinement to obtain a sequence of optimal meshes. The r -refinement approach is described first and consists of a combination of node movement and edge swapping in order to minimize the energy functional for a given size of mesh. At this stage the mesh may be any triangulation of the domain Ω , which is assumed to be a subset of \mathcal{R}^2 (i.e. $m = 2$ in equation (1)), and we only consider piecewise linear finite element trial functions.

2.1 Locally optimal meshes

We define a locally optimal mesh for the finite element solution of (1) to be a mesh with the following properties.

1. There exists some number $\varepsilon > 0$ such that if any node is displaced by any distance $\delta \leq \varepsilon$ in any direction (subject to the constraint that a boundary node remains on the boundary and the domain is not altered) the finite element solution on the new mesh has an energy which is no less than the energy of the finite element solution on the locally optimal mesh.
2. By noting that each internal edge of a triangulation is shared by exactly two triangles then, if the union of these two triangles is a convex quadrilateral, we may obtain a modified triangulation by swapping the diagonal of this quadrilateral, as shown in Figure 1. The finite element solution on any such modified triangulation has an energy which is no less than the energy of the finite element solution on the locally optimal mesh.

In order to obtain such a mesh from a given starting mesh we use an approach which is based heavily on that of [14]. This approach combines node movement and edge swapping in a manner which only requires the solution of local problems in order to converge to a global solution of the full problem, (1), on a locally optimal grid. For clarity we describe the node movement and the edge swapping algorithms separately and then discuss how they may be combined.

A necessary condition for the position of each node of the triangulation to be optimal is that the derivative of the energy functional with respect to each nodal position is zero. Like the approach of [14] our algorithm seeks to reduce the energy functional monotonically by moving each node in turn until the derivative with respect to the position of each node is zero. Whilst this does not guarantee with absolute certainty that we reach a local minimum (as opposed to a saddle point or even a local maximum), the presence of rounding errors combined with the downhill nature of the technique ensures that in practice any other outcome is almost impossible.

The algorithm proposed consists of a number of sweeps through each of the nodes in turn until convergence is achieved. At the beginning of each sweep the gradient, with respect to the position of

each node, of the energy functional

$$E = \int_{\Omega} F(\underline{x}, u^h, \underline{\nabla} u^h) d\underline{x} \quad (2)$$

is found (where u^h is the latest piecewise linear finite element solution). This is done using a slightly different approach to that described in [14], based upon [8] instead. In [8] it is proved that if \underline{s}_i is the position vector of node i then

$$\frac{\partial E}{\partial s_{id}} = \int_{\Omega} \left\{ \left[F \delta_{dj} - \frac{\partial u_k^h}{\partial x_d} F_{,3kj} \right] \frac{\partial \alpha_i}{\partial x_j} + F_{,1d} \alpha_i \right\} dx, \quad (3)$$

where α_i is the usual local piecewise linear basis function at node i , s_{id} is the d^{th} component of \underline{s}_i ($d = 1$ to m), $F_{,p}$ represents the derivative of F with respect to its p^{th} argument, other suffices represent components of tensors, δ_{dj} is the Kronecker delta and repeated suffices imply summation ($j = 1$ to m and $k = 1$ to n). Note that using (3) the gradients with respect to all of the vertices in the mesh may be assembled in a single pass of the elements. These gradients are then sorted by magnitude and the nodes visited one at a time, starting with the largest gradient.

When each node is visited the direction of steepest descent,

$$\underline{s} = - \frac{\partial E}{\partial \underline{s}_i}, \quad (4)$$

is used in order to determine along which line the node should be moved. The distance that the node is moved along this line is computed using a one-dimensional minimization of the energy subject to the constraint that the node should not move more than a proportion w ($0 < w < 1$) of the distance from its initial position to its opposite edge (see Figure 2). Once a new position for the node has been found the value of the solution, u_i say, at that node must be updated by solving the local problem

$$\min_{u_i \in \mathcal{R}^n} \int_{\Omega_i} F(\underline{x}, u^h, \underline{\nabla} u^h) d\underline{x}. \quad (5)$$

Here Ω_i is the union of all elements which have node i as a vertex and Dirichlet conditions are imposed on $\partial\Omega_i$ using the latest values for u^h . Once this update is complete the same process is undertaken for the next node in the sorted list and when each node has been visited the sweep is complete. Provided convergence has not been achieved the next sweep may then begin.

Using the above approach the interior nodes could move in any direction however a slight modification is required for nodes on the boundary of Ω . These nodes may only be moved tangentially along

the boundary and even then this is subject to the constraint that the domain remains unaltered. Where this constraint is not violated the downhill direction of motion along the boundary is easily computed by projecting \underline{s} from (4) onto the local tangent of the boundary. The one-dimensional minimization in this direction is then completed as for any other node. On Dirichlet boundaries the updated value of u is of course prescribed however on any other type of boundary it must be computed by solving a local problem of the same form as (5).

Once convergence with respect to the position of each node has been achieved a further reduction in the energy of the solution is sought by the use of edge swapping. Following [11] a loop through each of the internal edges of the mesh is completed and, for each edge, the local energy on the two triangles on either side is computed. The edge is then swapped in the manner illustrated in Figure 1 and the new local energy over the two triangles on either side is computed. If this energy is less than the original local energy on the quadrilateral then the edge swap is kept; otherwise it is rejected. Once the loop through each of the edges has been completed it is repeated until there is an entire pass for which no edges are swapped.

Of course the grid is unlikely to be locally optimal at this point since the edge swapping will generally cause the node locations to become sub-optimal. Hence it is necessary to alternate between the node movement and the edge swapping algorithms until the whole process has converged. The downhill nature of each step in the process guarantees that this will eventually occur. Despite this guarantee however, for pragmatic reasons it is useful to be able to impose a small number of additional constraints on the allowable meshes. For example, in our implementation of the edge swapping algorithm parameters are provided for the maximum number of edges that may be connected to a single node and for the smallest angle allowed in any triangle. Similarly, for the node movement algorithm a parameter is provided to specify the smallest area allowed for any element (and any triangle which shrinks to a size below this threshold is removed from the mesh by a simple element/node deletion algorithm). Numerous other such parameters could also be included.

2.2 Local mesh refinement

The main difficulty with the r -refinement strategy introduced in the previous subsection is that it is impossible to know *a priori* how many nodes or elements will be required in order to get a sufficiently accurate finite element solution to any given variational problem. Even an optimal mesh with a given number of nodes may not be adequate for obtaining a solution of a desired accuracy. For this reason some form of mesh refinement is essential.

In [14] global mesh refinement is combined with r -refinement and it is demonstrated that this provides better solutions than the use of uniform global h -refinement alone. In this work we extend these results in a number of ways. Firstly, by generalizing to systems of equations (i.e. $n > 1$ in (1)) and secondly, by using local (rather than uniform) h -refinement in conjunction with r -refinement. This, we demonstrate, leads to further efficiency gains above and beyond those observed in [14]. In addition, we also demonstrate that the hierarchical approach of starting with a coarse grid and then optimizing, refining, optimizing, refining, etc., provides a far more robust adaptive algorithm than simply refining first and then optimizing the node positions and the mesh topology at the end.

For the purposes of this two-dimensional work two different local refinement algorithms have been considered. The first of these divides all triangles which are to be refined into four children (as used in [10] for example and illustrated in the top half of Figure 3) whilst the other divides all triangles which are to be refined into just two children (as used in [4] for example and shown in the bottom half of Figure 3). In each case any “hanging nodes” left at the end (again see Figure 3) are removed by bisecting the neighbouring elements and then performing local edge swapping.

There are many possible ways in which the h -refinement might be combined with the r -refinement to produce a hybrid algorithm. Our experience suggests that a robust approach is to always refine an optimized mesh and then to interpolate the coarser solution onto the refined mesh as a starting point for the next level of optimization. It also appears to be advantageous to approximately optimize the nodal solution values first before attempting to optimize the positions. The local refinement itself attempts to subdivide all elements for which the local energy is greater than $X\%$ of the largest energy value calculated on any single element. Typically X is chosen to be somewhere between 40 and 80.

In the following section the performance of this hybrid algorithm is assessed using a number of

test problems. In each case comparisons are made with the approach of [14], in which r -refinement is combined with global h -refinement, and with the more conventional approach of using local h -refinement on its own.

3 Numerical results

In this section we study three representative test problems in order to assess the quality of the adaptive technique that has been described. The first of these is an artificial test case however the second problem is taken from [8] and [13] whilst the third appears in [3], [9] and [14].

3.1 Problem one

We begin by considering the simple two-dimensional reaction-diffusion equation

$$-\Delta u + \frac{1}{\varepsilon^2} u = 0, \quad \underline{x} \in \Omega = (0, 1) \times (0, 1), \quad (6)$$

subject to the Dirichlet boundary conditions

$$u = e^{-x_1/\varepsilon} \quad (7)$$

throughout $\partial\Omega$. This boundary condition is chosen so that (7) is also the exact solution of (6) over the whole domain Ω . Furthermore, solving equation (6) corresponds to minimizing the energy functional

$$E = \frac{1}{2} \int_{\Omega} \left[\frac{\partial u}{\partial x_i} \frac{\partial u}{\partial x_i} + \frac{u^2}{\varepsilon^2} \right] d\underline{x}, \quad (8)$$

and so this clearly falls into the general class defined by problem (1). Indeed, substituting (7) into (8) shows that the energy of the exact solution is given by

$$E = \frac{1}{2\varepsilon} [1 - e^{-2/\varepsilon}]. \quad (9)$$

For the purposes of these experiments however we restrict our consideration to the single case $\varepsilon = 0.01$, for which $E = 50.0000$.

Initially the problem is solved on a uniform coarse grid containing just 32 elements. This grid is then optimized using the r -refinement approach of Subsection 2.1 and the total energy reduces from

374.473 to 50.8937, reflecting the fact that before optimization there were no degrees of freedom in the boundary layer near $x_1 = 0$. Following [14] this optimal grid may now be uniformly refined to produce 128 elements which may themselves be optimized. This leads to a solution with a total stored energy of 50.1137. A further global refinement and optimization then leads to a solution with a total stored energy of 50.0158 on a mesh of 512 elements: this sequence of locally optimal meshes is shown in Figure 4.

Figure 5 illustrates two further meshes of 512 elements: the first obtained by global refinement of the initial uniform mesh and the second by optimizing this grid directly. The energies of the solutions on these meshes are 103.630 and 50.2311 respectively thus illustrating, for this example at least, the superiority of the hierarchical approach when r -refinement is combined with global h -refinement. It is clear from these meshes that although the second grid is locally optimal it suffers from the problem that too many of the degrees of freedom, inherited from the first grid, lie in a part of the domain that is far from the boundary layer.

The purpose of this paper however is to propose that the hybrid algorithm should combine r -refinement with local h -refinement and Figure 6 shows a sequence of meshes obtained in this manner. The first mesh is the same one, containing 32 elements, that appears in Figure 4, whilst the second, third and fourth meshes contain 42, 94 and 323 elements respectively and were obtained by refining into 2 children only those elements whose local energy exceeded 60% of the maximum local energy on any element. The total energies of the solutions on these four meshes are 50.8937, 50.3408, 50.1010 and 50.0085 respectively: clearly illustrating the superiority of the use of local rather than global h -refinement within the hybrid algorithm.

To conclude our discussion of this example we illustrate the advantage of applying the hybrid approach hierarchically by contrasting it with the use of local h -refinement alone, possibly followed by a single application of r -refinement. Figure 7 shows two grids of 1048 elements and two grids of 784 elements that were obtained in this manner (again using a threshold of $X = 60$ for the local refinement). The total energies of the solutions on the 1048 element meshes (obtained by local one-to-four refinement alone and then a single application of the mesh optimization at the end) are 54.8553 and 50.0536 respectively, whilst the total energies of the solutions on the 784 element meshes (obtained by

local one-to-two refinement plus a final optimization) are 51.4939 and 50.0714 respectively. We see that in both cases, despite the fact that the second of each pair of meshes is locally optimal, the quality of these local optima are not as good as that obtained using the hierarchical approach. A summary of all of the computations made for this test problem is provided in Table 1.

Elements	Energy	Description
32	374.473	Figure 4 (top left)
32	50.8937	Figure 4 (top right)
128	50.1137	Figure 4 (bottom left)
512	50.0158	Figure 4 (bottom right)
512	103.630	Figure 5 (left)
512	50.2311	Figure 5 (right)
32	50.8937	Figure 6 (top left)
42	50.3408	Figure 6 (top right)
94	50.1010	Figure 6 (bottom left)
323	50.0085	Figure 6 (bottom right)
1048	54.8553	Figure 7 (top left)
1048	50.0536	Figure 7 (top right)
784	51.4939	Figure 7 (bottom left)
784	50.0714	Figure 7 (bottom right)

Table 1: Summary of the results obtained for Problem one (the global energy minimum is 50.0000).

3.2 Problem two

We now consider the more challenging problem of calculating the displacement field for a two-dimensional linear elastic model of an overhanging cantilever beam supporting a vertical point load at the end of the cantilever, as illustrated in Figure 8. For this problem $m = n = 2$ and the energy functional is given by

$$E = \frac{1}{2} \int_{\Omega} \frac{\partial u_i}{\partial x_j} C_{ijkl} \frac{\partial u_k}{\partial x_\ell} d\mathbf{x} - \int_{\Omega} \rho b_i u_i d\mathbf{x} - \int_{\partial_\theta} \theta_i u_i ds. \quad (10)$$

Here, all repeated suffices are summed from 1 to 2, \mathbf{C} is the usual fourth order elasticity tensor (in this case corresponding to a Young's modulus $E = 100$ and a Poisson ratio $\nu = 0.001$), $\rho \underline{b}$ provides the external body forces due to gravity and \underline{t} represents the traction boundary condition (in this case a point load as illustrated in Figure 8). The left half of the lower boundary is fixed whilst the rest of the boundary, ∂_{θ} say, is free. Unlike for the first example we do not know an exact solution for this problem and so the optimal value for the stored energy is not known *a priori*.

As before we begin by solving the problem on a uniform coarse grid, this time containing 64 elements. This grid is then optimized using the r -refinement algorithm to reduce the total energy from -0.201352 to -0.253210 . This optimal grid may now be uniformly refined to produce 256 elements which are also optimized, leading to a solution with a total stored energy of -0.302353 . One further global refinement and optimization then leads to a solution with a total stored energy of -0.338964 on a mesh of 1024 elements. This sequence of locally optimal meshes is shown in Figure 9.

Figure 10 illustrates two further meshes of 1024 elements. The first of these is obtained by global refinement of the initial uniform mesh and the second by optimizing this grid directly. The energies of the solutions on these meshes are -0.306791 and -0.329249 respectively and so we again observe the superiority of the hierarchical approach when r -refinement is combined with global h -refinement.

As for the previous example, our goal is to assess the hybrid algorithm that combines r -refinement with *local* h -refinement hence Figure 11 shows a sequence of meshes obtained in this manner. The first mesh is the same one, containing 64 elements, that appears in Figure 9, whilst the second and third meshes contain 224 and 455 elements respectively and were obtained by refining into 2 children only those elements whose local energy exceeded 60% of the maximum local energy on any element. The total energies of the solutions on these three meshes are -0.253210 , -0.308351 and -0.363313 respectively: again illustrating the superiority of the use of local rather than global h -refinement within the hybrid algorithm.

We again conclude our example by illustrating the advantage of applying the hybrid approach hierarchically by contrasting it with the use of local h -refinement alone, possibly followed by a single application of r -refinement. Figure 12 shows two grids of 674 elements and two grids of 462 elements that were obtained in this manner (again using a threshold of $X = 60$ for the local refinement). The

total energies of the solutions on the 674 element meshes (obtained by local one-to-four refinement alone and then a single application of the mesh optimization at the end) are -0.325679 and -0.342525 respectively, whilst the total energies of the solutions on the 462 element meshes (obtained by local one-to-two refinement plus a final optimization) are -0.325879 and -0.342355 respectively. As before it is clear that the quality of the locally optimal meshes obtained in this manner is inferior to that of meshes obtained using the hierarchical approach. A summary of all of the computations made for this test problem is provided in Table 2.

Elements	Energy	Description
64	-0.201352	Figure 9 (top)
64	-0.253210	Figure 9 (second)
256	-0.302353	Figure 9 (third)
1024	-0.338964	Figure 9 (bottom)
1024	-0.306791	Figure 10 (top)
1024	-0.329249	Figure 10 (bottom)
64	-0.253210	Figure 11 (top)
224	-0.308351	Figure 11 (middle)
455	-0.363313	Figure 11 (bottom)
674	-0.325679	Figure 12 (top)
674	-0.342525	Figure 12 (second)
462	-0.325879	Figure 12 (third)
462	-0.342355	Figure 12 (bottom)

Table 2: Summary of the results obtained for Problem two (the global energy minimum is unknown).

3.3 Problem three

The final two-dimensional problem that we consider also involves just one dependent variable (i.e. $n = 1$ in (1)) however it features a solution which is singular at the origin. The energy functional

corresponds to the Laplacian operator and is given by

$$E = \frac{1}{2} \int_{\Omega} \frac{\partial u}{\partial x_i} \frac{\partial u}{\partial x_i} d\mathbf{x}, \quad (11)$$

where the presence of repeated suffices again implies summation from 1 to 2. The domain, Ω , is the unit disc with a 45° sector removed, as illustrated in Figure 13, and Dirichlet boundary conditions consistent with the exact solution $u = r^{2/7} \sin \frac{2\theta}{7}$ are applied throughout $\partial\Omega$. Since the exact solution is known in this case so is the true value of the global minimum of E in (11): 0.392699.

As with the previous examples the problem is first solved on a coarse initial mesh, in this case with just 28 elements, which is then optimized. This locally optimal mesh is then refined globally and optimized to three further levels, giving meshes of 112, 448 and 1792 elements respectively. These meshes are shown in Figure 14 and their corresponding solutions have energies of 0.549242, 0.434828, 0.404352 and 0.396215.

Once again, it may be observed that the approach of optimizing the mesh at each level after global refinement is superior to applying global h -refinement alone and then optimizing the resulting mesh. Figure 15 shows two meshes, each containing 1792 elements, that were obtained by this method. The energies of the solutions on these meshes are 0.438164 (uniform h -refinement only) and 0.405547 (after optimization), which are significantly worse than for the final mesh of Figure 14.

To conclude this example, we now consider the application of local h -refinement in our hybrid algorithm. Figure 16 shows a sequence of four meshes of 28, 107, 255 and 1275 elements respectively. In order to contrast the solutions on these meshes with those obtained on the meshes shown in Figure 14 we have forced refinement of each of the edges on the circular boundary so that the domains correspond to the four domains in Figure 14. Further refinement (one element to two children) has then been permitted locally for any elements whose energy is greater than 60% of the maximum energy on any single element. This local refinement is executed repeatedly on each domain until it is necessary to refine the boundary elements again. The total energies of the solutions on the four meshes shown in Figure 16 are 0.549242 (the same mesh as in Figure 14), 0.431777, 0.402413 and 0.395183 respectively.

Again we have seen the advantage of using the hierarchical mesh optimization approach with local, rather than global, refinement. Furthermore, when local h -refinement is used on its own, even if this is followed by mesh optimization, the resulting grids are not as good. Two pairs of such grids, containing

1437 (one-to-four refinement) and 1413 (one-to-two refinement) elements respectively, are illustrated in Figure 17. For these examples the corresponding finite element solutions have total energies of 0.407613 and 0.398523 (1437 elements before and after optimization) and 0.402199 and 0.398123 (1413 elements before and after optimization) respectively. (For the purposes of comparison, we have artificially refined those edges on the circular boundary so as to ensure that the domains are identical to the final domains in Figures 14 to 16.) A summary of all of the computations made for this test problem is provided in Table 3.

Elements	Energy	Description
28	0.549242	Figure 14 (top left)
112	0.434828	Figure 14 (top right)
448	0.404352	Figure 14 (bottom left)
1792	0.396215	Figure 14 (bottom right)
1792	0.438164	Figure 15 (left)
1792	0.405547	Figure 15 (right)
28	0.549242	Figure 16 (top left)
107	0.431777	Figure 16 (top right)
255	0.402413	Figure 16 (bottom left)
1275	0.395183	Figure 16 (bottom right)
1437	0.407613	Figure 17 (top left)
1437	0.398523	Figure 17 (top right)
1413	0.402199	Figure 17 (bottom left)
1413	0.398123	Figure 17 (bottom right)

Table 3: Summary of the results obtained for Problem three (the global energy minimum is 0.392699).

4 Discussion

4.1 Two dimensions

The three examples of the previous section have clearly illustrated that the quality of the final mesh produced when using the proposed hybrid algorithm is better, in the sense that the finite element solution has a lower energy, than that obtained by using either h -refinement or r -refinement alone. Furthermore it is demonstrated that combining the mesh optimization with local h -refinement is superior to the global refinement approach used in [14]. Finally, the advantage of using the hierarchical approach, whereby the intermediate level meshes are optimized, is also apparent: an excellent combination of small mesh sizes and low energies for the corresponding finite element solutions being achieved.

When discussing the merits of our proposed algorithm it is important to note that there are some problems for which the benefits may not be quite so substantial as those observed in the three examples above. A common feature to each of these examples is the desirability of clustering the majority of the mesh elements in a relatively small subset of the domain. When a problem is such that the optimal mesh is more uniformly distributed across the domain the local refinement algorithm will show little or no advantage over the global approach of [14] since almost all elements of the mesh will need to be refined when moving from one level to the next. This is a phenomenon that we have observed in at least one example that we have considered (the nonlinear problem used as the second test problem in [14]). Nevertheless, even in this case, our variant of the algorithm performed no worse than that used in [14].

4.2 Three dimensions

Up to this point our discussion has been restricted mainly to the solution of two-dimensional problems (i.e. $m = 2$ in (1)). We now conclude the paper by considering how the multilevel approach may also be applied to obtain optimal tetrahedral meshes when solving problems in three dimensions.

The definition of a locally optimal mesh in Subsection 2.1 contains two components. One is that the position of each vertex of the mesh should be locally optimal, whilst the other is that the connectivity of these vertices should also be locally optimal. The first of these is quite straightforward to generalize to three dimensions however the edge-swapping part of the definition is more complex. This is discussed

in some detail in [6] for example, where a number of different stencils are used to modify the local topology of the mesh depending upon how many elements share an edge. As a general rule, if there are e elements sharing a particular edge ($e \geq 3$) then the union of these elements may be reconnected in a way that replaces them with $2e - 4$ different elements. This is made even more complicated by the fact that there are numerous alternative ways to reconnect the region in this manner, all of which need to be considered when seeking a local optimum. In [6] the objective is just to improve the geometric quality of the mesh and so it is not always necessary to consider all possible edge swaps (they never locally reconnect the mesh when $e > 7$ for example). Moreover, because these local reconnections of the mesh allow the possibility of introducing new elements and edges it is not entirely straightforward to guarantee the termination of an energy minimization algorithm based upon this approach.

Due to these difficulties associated with edge swapping, we restrict this initial discussion on producing locally optimal tetrahedral meshes to the problem of optimizing the node locations only. This means that we will consider a mesh to be optimal if it satisfies the first of the two conditions enumerated in Subsection 2.1.

The node movement part of our algorithm then generalizes simply to three dimensions. The derivatives of the energy with respect to the nodal positions may still be computed using (3) with a single loop over the elements of the mesh. This list may then be sorted and, beginning with the largest values of $|\frac{\partial E}{\partial \mathbf{x}_i}|$, the nodes may be moved in turn. In each case the movement requires an (approximate) one-dimensional minimization in the direction of steepest descent (given by (4)). As in two dimensions we may also introduce artificial constraints on this minimization to prevent the possibility of the mesh from becoming too degenerate. Once the updated location of node i has been found it is a simple matter to modify the corresponding solution value through a local solve on the patch of elements, Ω_i , surrounding that node.

The local refinement part of the algorithm could either be implemented by dividing each tetrahedron into two children (as in [4] for example) or into eight children (as in [12] for example). It is the latter approach that we use here, and this is illustrated in Figure 18. The removal of “hanging nodes”, which appear when neighbouring elements are at different refinement levels, is achieved through the use of a transitional refinement layer, as illustrated in Figure 19.

For a simple test problem we consider the following generalization of the first equation solved in the previous section.

$$-\Delta u + \frac{1}{\varepsilon^2} u = 0, \quad \underline{x} \in \Omega = (0, 1) \times (0, 1) \times (0, 1), \quad (12)$$

subject to the Dirichlet boundary conditions

$$u = e^{-x_1/\varepsilon} \quad (13)$$

throughout $\partial\Omega$. As with the two-dimensional example (13) is the true solution of (12) over all of Ω , and the corresponding energy functional ((8) but with this modified Ω and summation of the repeated suffices from 1 to 3) has a minimum value given by (9). We again choose $\varepsilon = 0.01$ to yield a thin boundary layer near $x_1 = 0$ and an optimal energy $E = 50.0000$.

Following the approach used for testing the two-dimensional algorithm in Section 3, we begin by assessing the performance of three-dimensional multilevel mesh optimization when combined with global h -refinement. Initially the test problem is solved on a regular coarse grid of 384 tetrahedral elements, as illustrated in Figure 20. This mesh is then optimized and the total energy of the solution reduces from 378.628 to 104.857. Three levels of uniform refinement, each followed by optimization, then yield solutions with energies of 59.9077, 52.3988 and 50.7552 on meshes of 3072, 24576 and 196608 elements respectively. To see that this final mesh is superior to one obtained without multilevel optimization the problem is then solved on a three level uniform refinement of the initial mesh shown in Figure 20 (with 196608 elements therefore), to yield a solution with energy 67.2790. When this mesh is locally optimized however the energy only decreases to a value of 52.4342.

We now demonstrate that the potential advantages of using local refinement with the multilevel optimization also appear to apply in three dimensions. Starting with the locally optimal 384 element grid, a sequence of three further meshes is obtained through local h -refinement (again using a threshold of $X = 60$) followed by local optimization. These meshes contain 2655, 16933 and 100866 tetrahedral elements and the corresponding solutions have energies of 59.9024, 52.3814 and 50.7460 respectively. Finally, we demonstrate the superiority of this final mesh over one obtained using only local h -refinement followed by local optimization at the end. This comes from the observation that a grid of 573834 elements obtained using only local h -refinement yields a solution energy of 54.8852

and when this is optimized the solution energy only reduces to 51.3324. A summary of all of these computational results is provided in Table 4.

Elements	Energy	Description
384	104.857	
3072	59.9077	Multilevel optimization and global h -refinement.
24576	52.3988	
196608	50.7552	
196608	67.2790	Global h -refinement followed by optimization.
196608	52.4342	
384	104.857	
2655	59.9024	Multilevel optimization and local h -refinement.
16933	52.3812	
100866	50.7460	
573834	54.8852	Local h -refinement followed by optimization.
573834	51.3324	

Table 4: Summary of the results obtained for the three-dimensional test problem (the global energy minimum is 50.0000).

Because of the difficulties in visualizing very large unstructured tetrahedral meshes we do not include pictures of all of the grids described above. Nevertheless, it is perhaps informative to include a couple of representative examples. Figure 21 therefore shows a mesh of 10687 elements created as part of the above sequence of local h -refinements. The solution on this mesh has an energy of 103.461. In contrast to this, Figure 22 shows a locally optimal mesh of 2655 elements, created as part of the sequence of multilevel optimizations with local h -refinement. Although containing many fewer elements than the mesh in Figure 21 the solution on this mesh also has a lower energy of 59.9024. Despite not appearing to be particularly smooth, the mesh in Figure 22 certainly seems to possess the excellent qualities of being both fine in the direction perpendicular to the boundary layer (near the face $x_1 = 0$) and quite coarse in the directions parallel with the layer. It is anticipated that the addition of a suitable

edge-swapping strategy will, as in two dimensions, further improve the quality of these meshes.

5 Conclusion

In this paper we have presented a technique for producing finite element solutions to variational problems on locally optimal meshes. The major contribution is to propose a multilevel approach which is shown to lead to better quality meshes, with fewer elements, than those obtained by using alternative techniques. Furthermore, based on the ideas presented in [14], there is no need to solve any global problems other than on an initial coarse grid. Extensive numerical results have been presented for a variety of problems in two dimensions and more provisional results have been described for a three-dimensional example. All of these numerical experiments have proved to be extremely encouraging.

Some additional work is still required however to turn this promising technique into efficient, reliable and robust general-purpose software. For example, the use of edge-swapping has proven to be highly beneficial in two dimensions and an approach similar to that used (in a different context) in [6] is therefore also likely to be well worth including. In addition, although global solves are not strictly necessary, there may well be efficiency gains to be made through the use of approximate global solves at appropriate points in the algorithm (immediately after h -refinement for example): these should be investigated carefully. As another example, it is still an open question as to how accurately the mesh needs to be optimized at each level of the hierarchy before local refinement takes place. Related to this, it also is unclear how accurately it is necessary to solve each of the one-dimensional minimization problems that are encountered at each node within each sweep of the node movement algorithm. Other issues that should be considered further concern the importance of the order in which nodes and edges are visited during local optimization sweeps and the possibility of making more aggressive use of the element/node deletion algorithm that is currently only employed when elements shrink to zero.

Acknowledgements

RM gratefully acknowledges the Government of Pakistan for financial support in the form of a merit scholarship. The work of MAW was funded through EPSRC research grant GR/M00077.

References

- [1] I. Babuska, B.A. Szabo and I.N. Katz, *The p-Version of the Finite Element Method*, SIAM Journal on Numerical Analysis, vol.18, pp.515–545, 1981.
- [2] J.M. Ball, P.K. Jimack and T. Qi, *Elastostatics in the presence of a temperature distribution or inhomogeneity*, Zeitschrift Fur Angewandte Mathematic Und Physik, vol.43, pp.943–973, 1992.
- [3] R.E. Bank, *PLTMG Users' Guide 7.0*, SIAM, Philadelphia, 1994.
- [4] E. Bänsch, *An adaptive finite element strategy for the 3-dimensional time-dependent Navier-Stokes equations*, Journal of Computational and Applied Mathematics, vol.36, pp.3–28, 1991.
- [5] M. Delfour, G. Payre and J.-P. Zolésio, *An optimal triangulation for second-order elliptic problems*, Computer Methods in Applied Mechanics and Engineering, vol.50, pp.231–261, 1985.
- [6] L.A. Freitag and C. Ollivier Gooch, *Tetrahedral mesh improvement using swapping and smoothing*, International Journal for Numerical Methods in Engineering, vol.40, pp.3979–4002, 1997.
- [7] P.K. Jimack, *A best approximation property of the moving finite element method*, SIAM Journal on Numerical Analysis, vol.33, pp.2206–2232, 1996.
- [8] P.K. Jimack, *An optimal finite element mesh for elastostatic structural analysis problems*, Computers and Structures, vol.64, pp.197–208, 1997.
- [9] C. Johnson, *Numerical solution of partial differential equations by the finite element method*, Cambridge University Press, Cambridge, 1987.
- [10] R. Lohner, *An adaptive finite element scheme for transient problems in CFD*, Computer Methods in Applied Mechanics and Engineering, vol.61, pp.267–281, 1987.
- [11] S. Rippa and B. Schiff, *Minimum energy triangulations for elliptic problems*, Computer Methods in Applied Mechanics and Engineering, vol 84, pp.257–274, 1990.

- [12] W. Speares and M. Berzins, *A 3-d unstructured mesh adaptation algorithm for time-dependent shock dominated problems*, International Journal for Numerical Methods in Fluids, vol 25, pp.81–104, 1997.
- [13] B.H.V. Topping and A.I. Khan, *Parallel Finite Element Computations*, Saxe-Coburg Publications, Edinburgh, 1996.
- [14] Y. Tourigny and F. Hulsemann, *A new moving mesh algorithm for the finite element solution of variational problems*, SIAM Journal on Numerical Analysis, vol. 35, pp.1416–1438, 1998.
- [15] N.P. Weatherill, *Grid adaptation using a distribution of sources applied to inviscid compressible flow simulations*, International Journal for Numerical Methods in Fluids, vol 19, pp.739–764, 1994.
- [16] O.C. Zienkiewicz, D.W. Kelly and J.P. Gago, *The Hierarchical Concept in Finite Element Analysis*, International Journal of Computers and Structures, vol.16, pp.53–65, 1982.

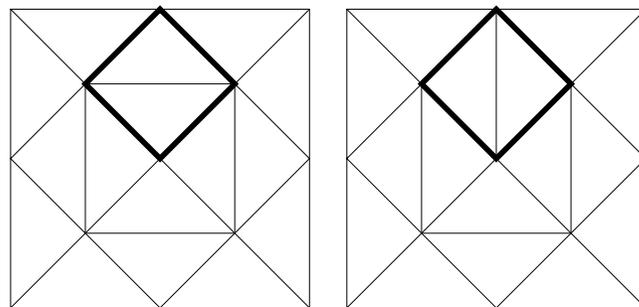


Figure 1: An illustration of the modification of a mesh by the swapping of a single edge.

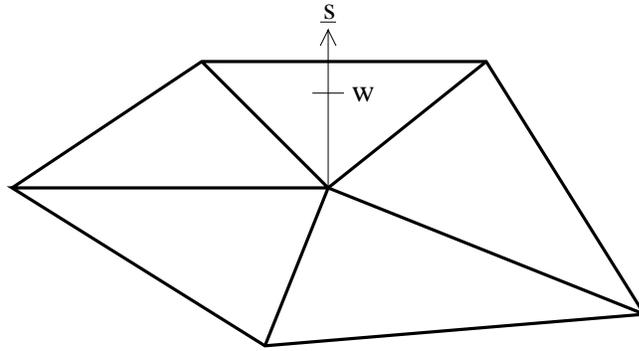


Figure 2: An illustration of local node movement: \underline{s} is the direction of steepest descent for the node motion and w represents the maximum distance that the node may move in this direction.

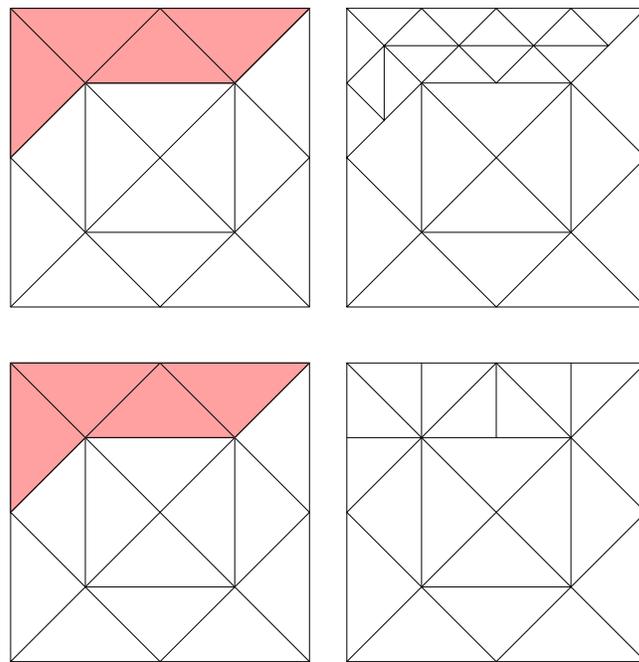


Figure 3: An illustration of the refinement of certain (shaded) elements of a mesh using one-to-four subdivision (top) and one-to-two subdivision (bottom).

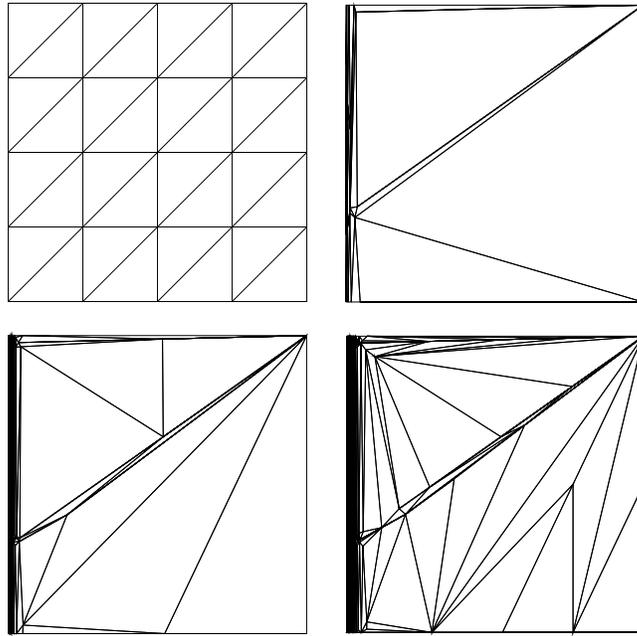


Figure 4: An initial mesh (top left) followed by a sequence of meshes obtained by r -refinement and then combinations of global h -refinement with r -refinement.

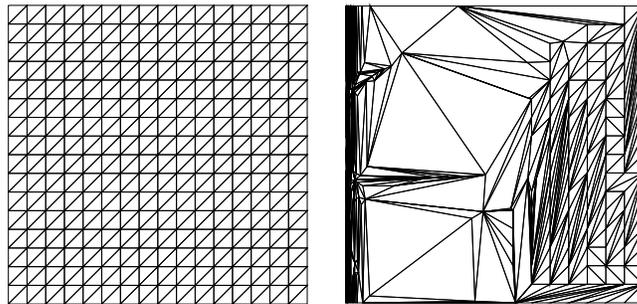


Figure 5: A globally refined mesh of 512 elements and the corresponding locally optimized mesh.

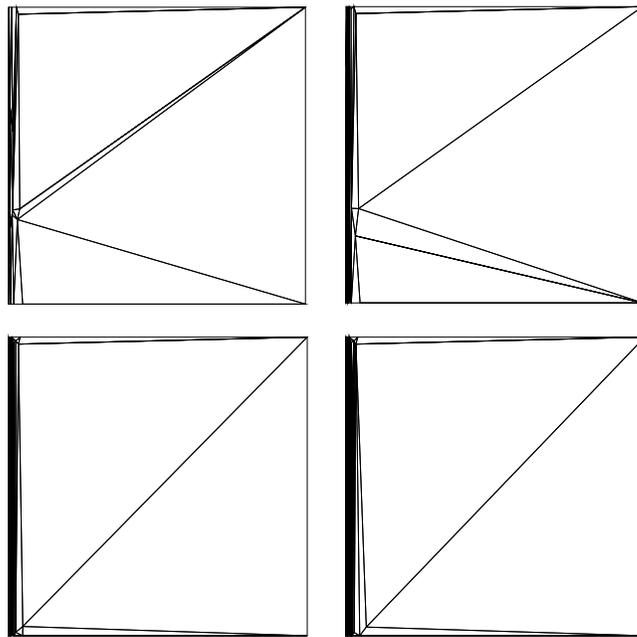


Figure 6: A sequence of meshes obtained by r -refinement of an initial coarse grid (top left) and then combinations of local h -refinement followed by r -refinement.

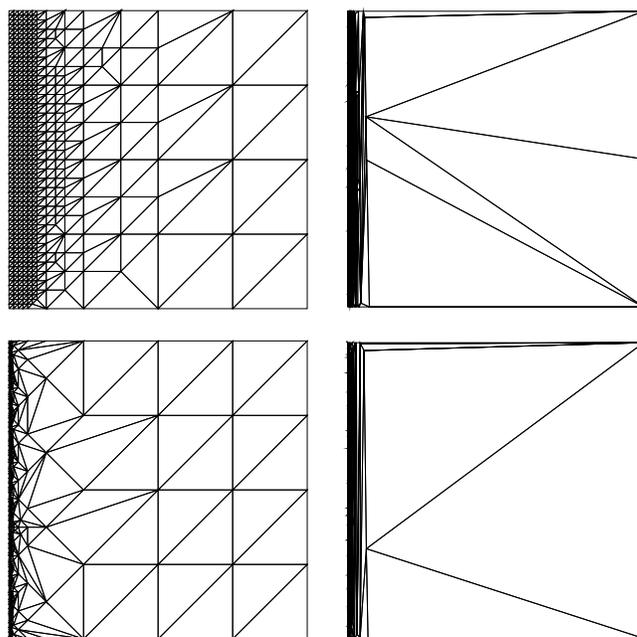


Figure 7: A pair of meshes of 1048 elements obtained using local one-to-four h -refinement (top left) followed by optimization and a pair of meshes of 784 elements obtained using local one-to-two h -refinement (bottom left) followed by optimization.

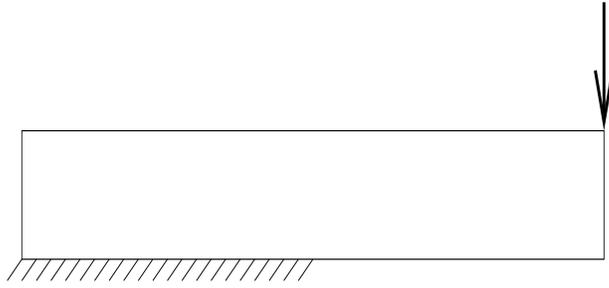


Figure 8: An illustration of the overhanging cantilever beam with a vertical point load at the end of the cantilever.

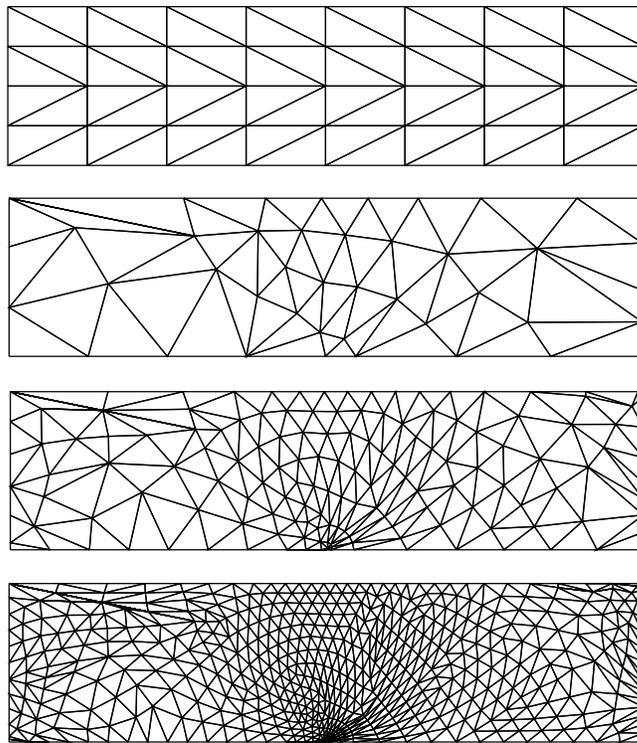


Figure 9: An initial mesh followed by a sequence of meshes obtained by r -refinement and then combinations of global h -refinement with r -refinement.

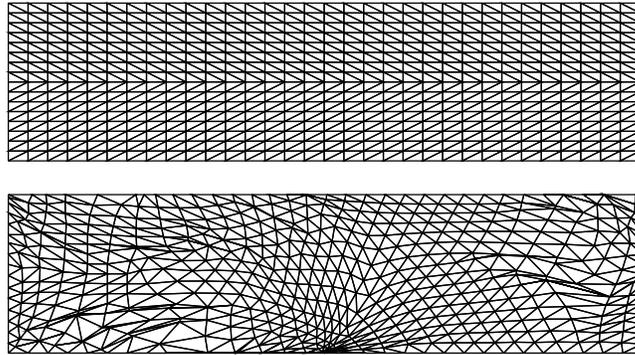


Figure 10: A globally refined mesh of 1024 elements and the corresponding locally optimized mesh.

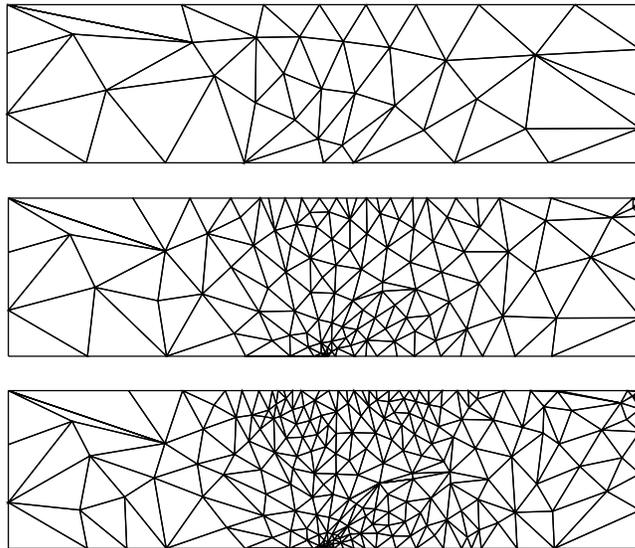


Figure 11: A sequence of meshes obtained by r -refinement of an initial coarse grid and then combinations of local h -refinement followed by r -refinement.

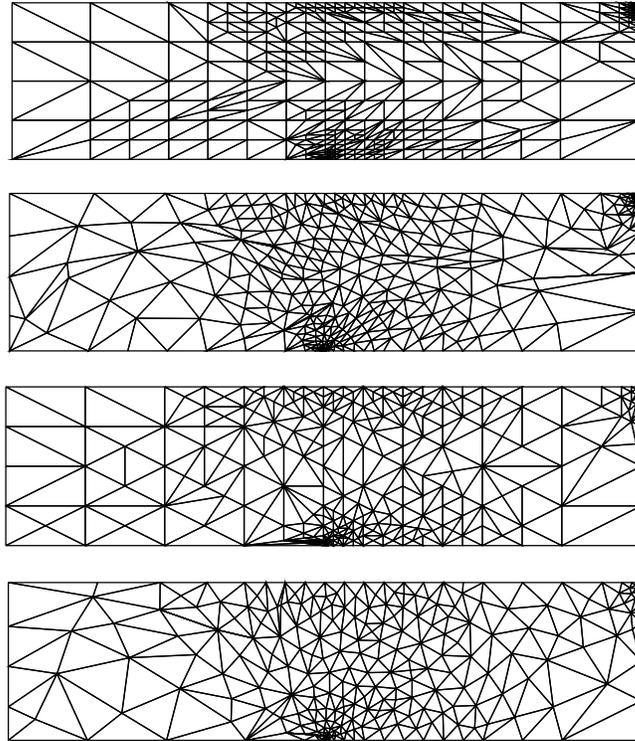


Figure 12: A pair of meshes of 674 elements obtained using local one-to-four h -refinement (top) followed by optimization (second) and a pair of meshes of 462 elements obtained using local one-to-two h -refinement (third) followed by optimization (bottom).

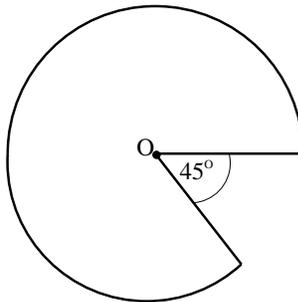


Figure 13: An illustration of the domain for the singular problem.

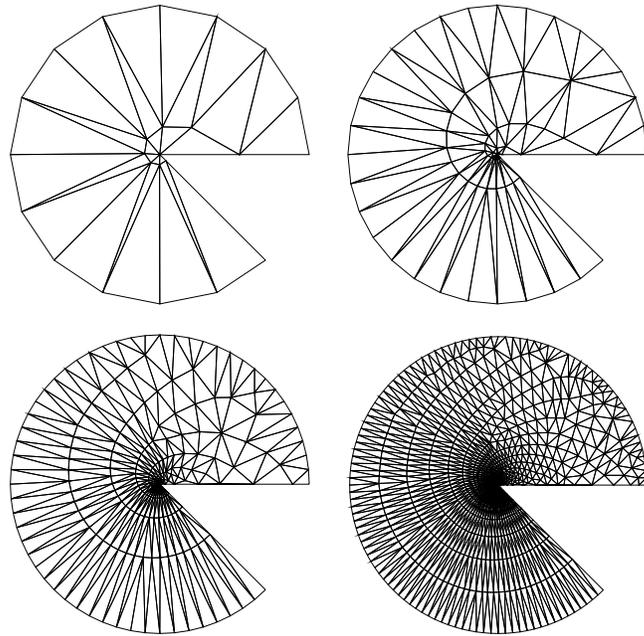


Figure 14: A sequence of meshes obtained by r -refinement of an initial coarse grid (top left) and then combinations of global h -refinement followed by r -refinement.

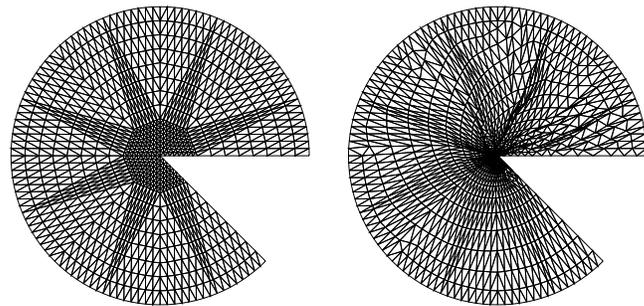


Figure 15: A globally refined mesh of 1792 elements and the corresponding locally optimized mesh.

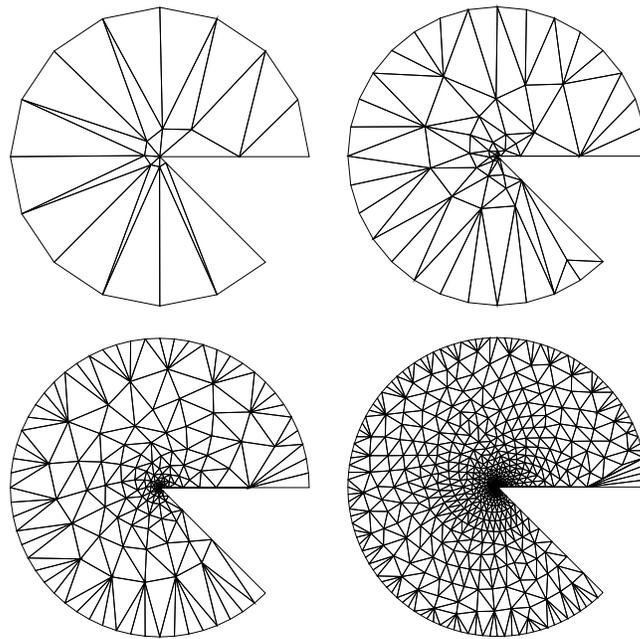


Figure 16: A sequence of meshes obtained by r -refinement of an initial coarse grid (top left) and then combinations of local h -refinement followed by r -refinement.

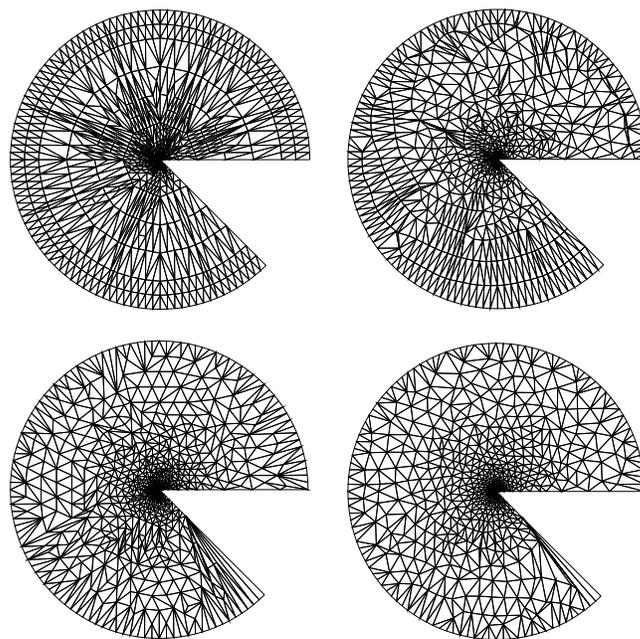


Figure 17: A pair of meshes of 1437 elements obtained using local one-to-four h -refinement (top left) followed by optimization and a pair of meshes of 1413 elements obtained using local one-to-two h -refinement (bottom left) followed by optimization.

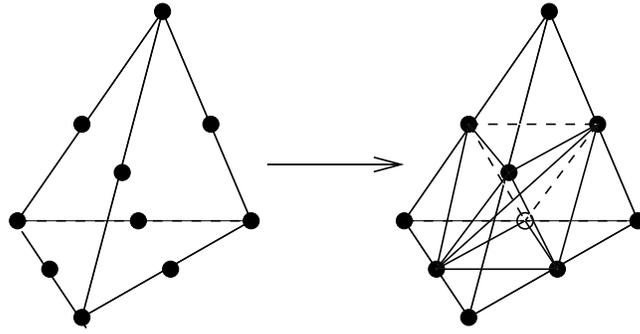


Figure 18: An illustration of the regular refinement of a tetrahedron into eight children by bisecting each edge.

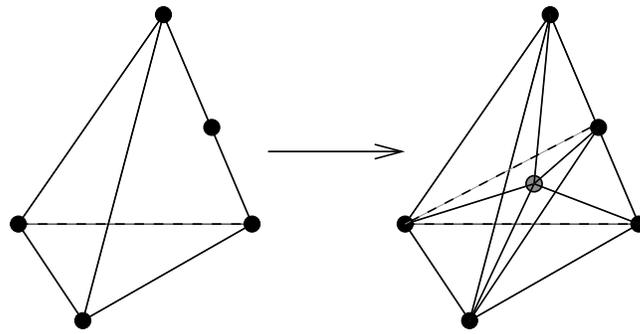


Figure 19: An illustration of the transitional refinement of a tetrahedron when it has a single “hanging node”.

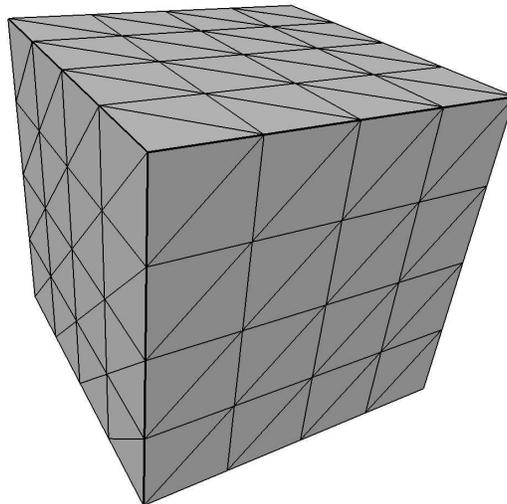


Figure 20: An illustration of an initial uniform mesh containing 384 tetrahedral elements.

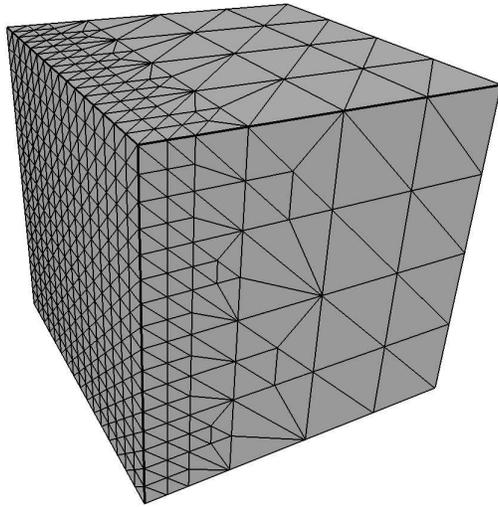


Figure 21: An illustration of a 10687 element mesh for the solution of equation (12) using local h -refinement alone.

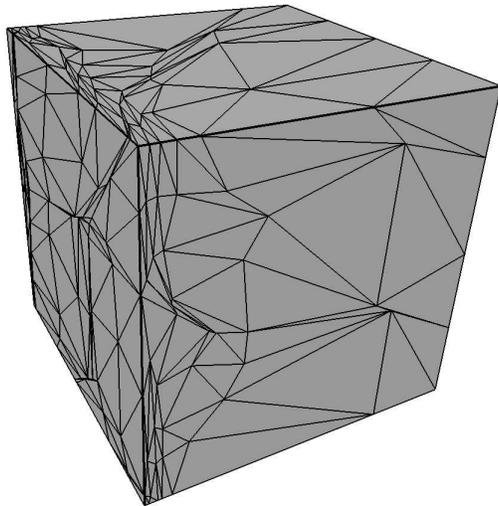


Figure 22: An illustration of a 2655 element mesh for the solution of equation (12) using the hybrid algorithm.