

A New Methodology for Anisotropic Mesh Refinement Based Upon Error Gradients

Thomas Apel* Sergei Grosman* Peter K. Jimack† Arnd Meyer*

March 25, 2003

Abstract

We introduce a new strategy for controlling the use of anisotropic mesh refinement based upon the gradients of an *a posteriori* approximation of the error in a computed finite element solution. The efficiency of this strategy is demonstrated using a simple anisotropic mesh adaption algorithm and the quality of a number of potential *a posteriori* error estimates is considered.

Keywords: Anisotropic refinement, *A posteriori* error estimation, Finite element method.

AMS classification: 65N50, 65N30.

1 Introduction

The use of anisotropic mesh refinement in the adaptive finite element solution of partial differential equations (PDEs) with highly anisotropic solutions is widely recognised as having significant potential for improving the efficiency of the solution process, e.g. [3, 5, 12, 20, 26, 33]. Numerous schemes for driving anisotropic mesh adaptivity have been considered in both the engineering, [3, 11, 12, 26, 33], and the numerical analysis literature, [14, 30, 31, 32]. Typically, such schemes are based on *a priori* knowledge of features of the equation and of the nature of the solution (e.g. [6, 23]) or on *a posteriori* knowledge (the numerical solution) to drive the refinement (e.g. [10, 16, 18, 29]). *A priori* knowledge of edge singularities is included in [5, 6] via a special coordinate transformation in the vicinity of that edge leading to the effect that new nodes generated through the adaptive procedure are suitably placed at a location different from the usual midpoints of edges. In [23, Section 6] the authors use *a priori* knowledge of boundary and interior layers to generate a useful anisotropic initial mesh. A different approach is described in [32, 33] where structured anisotropic meshes are used locally near interior layers or shocks, and the approximate solution on a previous mesh is used to determine the position of the layers and to guide the refinement outside the layer. Several authors use the heuristic argument that the local element size parameters should correspond to the ratio of the eigenvalues of the matrix of the (approximated) second order partial derivatives of the solution, with the stretching direction determined by the eigenvector to the largest eigenvalue of this matrix, see [3, 14, 26, 33] and the literature cited there. In this communication we propose an alternative technique for driving anisotropic mesh refinement based upon *a posteriori* error estimation of the finite element error. The technique is related to the method in [27] where local interpolation errors are estimated and equidistributed. The purpose of this work however is to consider whether gradients of *a posteriori* error estimates may be used to control anisotropic refinement in an effective manner.

*Fakultät für Mathematik, Technische Universität Chemnitz, 09107 Chemnitz, Germany

†Computational PDE Unit, School of Computing, University of Leeds, Leeds LS2 9JT, UK

For creating a new mesh, there are three main strategies. The first demands complete *remeshing* on the basis of background information (local mesh sizes, stretching direction); see the overview article [31] and the literature cited therein. Some authors report on anisotropic meshes which have nearly equilateral elements in a local non-Euclidean metric [8]. In this way standard mesh generation techniques are used to solve the meshing problem [13]. Remeshing is quite expensive but one can produce meshes with a gradually changing mesh size and arbitrary stretching directions. The second strategy is based on a subdivision of the existing elements. This approach is inexpensive and fits very well into multi-grid/multi-level strategies for the solution of the corresponding finite element equation system. The subdivision strategy was adapted for anisotropic refinement in [19] and will be investigated also in this paper. The disadvantage is that the initial mesh strongly determines the possible stretching directions of the elements. This can be compensated for by node relocation techniques, sometimes also called adaptive grid orientation [19] or node relaxation techniques [28]. In the third strategy one concentrates on relocating the nodes: this is also called the *r*-version of the finite element method. In order to produce a converging method however, one has to combine this with node insertion or element splitting. In [8, 12] such algorithms are described which allow anisotropic refinement on the basis of a local non-Euclidean metric tensor.

In this paper we investigate the potential of a new strategy for controlling the use of anisotropic mesh refinement. It is based upon the gradients of an *a posteriori* approximation of the error in a computed finite element solution. We focus on a linear, second order, reaction-diffusion test problem

$$-\Delta u + \kappa^2 u = f; \quad u \in \Omega = (0, 1) \times (0, 1), \quad (1)$$

with Dirichlet boundary conditions on $\partial\Omega$. Note that when $f = 0$ equation (1) is satisfied by

$$u = e^{-\kappa x} + e^{-\kappa y} \quad (2)$$

which features highly anisotropic boundary layers when $\kappa \gg 1$. In order to assess the quality of our proposed mechanism for driving the mesh adaptivity described in Sections 2 and 3, we only use a very simple mesh refinement algorithm in this work. This allows us to separate out the issue that we are concerned with here, of how to drive the anisotropic refinement (i.e. provide the information needed to decide where an existing mesh needs to be refined and in which directions) in a robust manner, from the (equally important) issue of how to execute the refinement. The paper concludes with a discussion of a number of important implementation issues and an assessment of the potential of our new approach.

2 An adaptive strategy

The natural norm in which to measure the error of a numerical approximation to the solution of (1) is the energy norm [22] given by

$$\|v\|^2 = \left\| \frac{\partial v}{\partial x} \right\|^2 + \left\| \frac{\partial v}{\partial y} \right\|^2 + \kappa^2 \|v\|^2, \quad (3)$$

where $\|\cdot\|$ represents the usual L^2 norm over Ω . When a set of finite element solutions is obtained using a nested sequence of progressively larger trial spaces (based upon conventional, isotropic, *h*-refinement for example), the final term on the right-hand side of (3) will tend to zero faster than the other two. It may be the case however that, although tending to zero at the same asymptotic rate, the ratio between these two dominant terms is far from one. Our proposal is to drive an anisotropic mesh refinement algorithm based upon the target of equilibrating these two terms, on each element, before reducing them at the same rate using conventional *h*-refinement. Although the asymptotic

convergence rate will not be improved by such a strategy, we expect to see a significant computational gain.

In order to demonstrate the effectiveness of the proposed strategy we begin by considering the exact error to equation (1),

$$e^h = u - u^h, \quad (4)$$

where u is given by (2) and u^h is the Galerkin finite element solution to (1) (with $f = 0$), subject to exact Dirichlet boundary conditions, on a given mesh. Moreover, we also consider only a very straightforward anisotropic refinement algorithm based upon the refinement of rectangles in one of the three ways illustrated in Figure 1 (see also [27]).




Figure 1: The three types of refinement allowed: regular (left), anisotropic in x (middle) and anisotropic in y (right).

When an error on a particular rectangle exceeds some tolerance (typically 20% of the maximum error over all rectangles) it is refined. If the $\kappa^2 \|e^h\|^2$ component of the error is dominant on a rectangle (i.e. $\kappa^2 \|e^h\|^2 > C \| \|e^h\|^2$, where C is typically chosen to equal 0.5), or if

$$\left\| \frac{\partial e^h}{\partial x} \right\| / \left\| \frac{\partial e^h}{\partial y} \right\| \in \left[\frac{1}{2}, 2 \right], \quad (5)$$

then regular refinement takes place on that rectangle. Otherwise, anisotropic refinement takes place: refining in x when $\left\| \frac{\partial e^h}{\partial x} \right\|$ dominates and in y when $\left\| \frac{\partial e^h}{\partial y} \right\|$ dominates. Rectangles are also refined when any edges contain two or more “hanging nodes” (due to successive refinement of a neighbour) so as to prevent excessively large changes in h_x or h_y (the mesh sizes in the x and y directions respectively). This additional refinement is isotropic if the hanging nodes are due to isotropic refinement and anisotropic if they are due to anisotropic refinement, with the former taking precedence in ambiguous cases involving more than one edge of the element having hanging nodes.

Figure 2 illustrates graphs of the error against the number of finite element unknowns when equation (1) is solved, with $\kappa = 10^3$, using the above refinement strategy. In the first (left) case $f = 0$ and the exact solution is (2), whilst in the second (right) case $f \neq 0$ and the Dirichlet boundary conditions are chosen so as to permit the exact solution

$$u = e^{-\kappa x} + e^{-\kappa y} + x^2 + \cos(10y). \quad (6)$$

The main difference between these two cases is that u is essentially zero away from the boundary in the former but is a smooth non-zero function away from the boundary in the latter. The three graphs plotted in each case are all obtained using the same initial isotropic mesh (where each element has an aspect ratio of 1) and correspond to artificially imposing a maximum element aspect ratio after refinement of 1 (i.e. only regular refinement allowed), 16 and 256 respectively. For the purposes of this work we define the aspect ratio (AR) of a rectangular element to be $\max(h_x/h_y, h_y/h_x)$. Furthermore, all finite element calculations are performed using piecewise linear elements obtained by dividing each rectangle into two triangles (with exceptional divisions into three or four triangles

when hanging nodes are present). A discussion of the use of even larger aspect ratios is postponed until Section 4. An illustration of the meshes obtained near to the bottom left corner of the domain in the cases where maximum aspect ratios of 1 and 16 are used is provided in Figure 3: the total error in the computed solutions on these two meshes being almost identical.




Figure 2: Graphs of the error as a function of the number of unknowns for the two test problems (with $\kappa = 10^3$) with a maximum possible aspect ratio of 1 (upper), 16 (middle) and 256 (lower) in each case. For each calculation the adaptivity is based upon the energy norm of the (known) exact error, calculated using a seven point quadrature rule with degree of precision five.




Figure 3: Sections of typical meshes near the bottom left corner of the domain when maximum aspect ratios of 1 and 16 are used.

The results shown in Figure 2 are typical of those obtained for all large values of κ and clearly show a significant advantage from the use of our anisotropic refinement strategy. Note that the asymptotic convergence rates are the same in all cases so the improvement is by a constant factor. Also note that in both examples our adaptive algorithm begins with uniform refinement to drive down the large initial L^2 error before the advantage of using anisotropic refinement is seen.

3 A *posteriori* error estimation

Whilst the numerical results exemplified by those presented in the previous section are extremely encouraging, it is clear that for the proposed methodology to be of any practical value a similar quality of results must also be achievable using *a posteriori* error estimates (as opposed to the exact error). Recently there has been a significant amount of research into the development and analysis of *a posteriori* error estimates that are effective on highly anisotropic grids (see, for example, [15, 20, 21, 30]). Our requirement is even more demanding than this however, since we are restricted to those techniques which yield local estimates of the error as a function: thus enabling us to compute each of the three L^2 norms that appear in (3) on each element.

For the purposes of this investigation we have selected a small number of well-known *a posteriori* error estimation algorithms and contrasted their use with that of the exact error as described in the previous section. The oldest one, BW, is due to Bank and Weiser [9] and consists of locally solving for $\eta_{BW}^h \in V_T$ such that

$$a_T(\eta_{BW}^h, v^h) = \langle f, v^h \rangle_T - a_T(u^h, v^h) \quad \forall v^h \in V_T$$

where $a_T(u, v) = \int_T (\nabla u \cdot \nabla v + \kappa^2 uv)$, $\langle f, v \rangle_T = \int_T f v$, and V_T is a finite dimensional space spanned by the quadratic edge bubble functions belonging to all edges $E \subset \partial T \setminus \partial\Omega$. For the second error estimator considered, AO due to Ainsworth and Oden [2], one has also to solve a local problem, for $\eta_{AO}^h \in \tilde{V}_T$ such that

$$a_T(\eta_{AO}^h, v^h) = \langle f, v^h \rangle_T - a_T(u^h, v^h) + \int_{\partial T \setminus \partial\Omega} g_T v^h \quad \forall v^h \in \tilde{V}_T,$$

where g_T is an approximation to the normal flux $\nabla u|_T \cdot n_T$ on ∂T , and n_T is the outer normal vector. The space \tilde{V}_T is in theory infinite-dimensional and in practice a small finite element space, see Section 4. For the determination of g_T , Ainsworth and Oden introduced the equilibrated residual method, which was modified by Ainsworth and Babuška [1] for the use of large κ ; here this modified error estimator is abbreviated as AB. Another approach is to recover improved gradients $\nabla^R u^h$ by locally averaging ∇u^h and to use $\nabla \eta_{ZZ}^h = \nabla^R u^h - \nabla u^h$. This approximation is due to Zienkiewicz and Zhu [34] and is therefore denoted by ZZ. Note that this estimator provides no estimate of the L^2 -part, $\|e^h\|$, of the error. Recently, Kunert [21] investigated an error estimator, referred to here as GK, which is based on solving local Dirichlet problems for $\eta_{GK}^h \in \tilde{\tilde{V}}_T$,

$$a_{\omega_T}(\eta_{GK}^h, v^h) = \langle f, v^h \rangle_{\omega_T} - a_{\omega_T}(u^h, v^h) \quad \forall v^h \in \tilde{\tilde{V}}_T,$$

where ω_T is the union of T with all elements sharing a face with T , and $\tilde{\tilde{V}}_T$ consists of the element bubble function related to T and certain so-called squeezed edge bubble functions. This error estimator is proved to be reliable and efficient on sensible anisotropic finite element meshes with constants that neither depend on the parameter κ nor on the aspect ratio of the elements. Since all of these error estimators are defined on triangles, piecewise linear triangular finite elements have again been used, as described in the previous section.

In a first test we repeat the calculations of Section 2 on isotropic meshes (i.e. with a maximum aspect ratio of 1). For both examples, and after each finite element solve, we not only compute the exact error function, e^h , but also each of the *a posteriori* estimates (η^h say) described above. Table 1 presents the results of these calculations in the form of effectivity ratios $\|\eta^h\|/\|e^h\|$. The results presented in Table 1 indicate that all estimators can be used on isotropic meshes. The effectivity ratios tend asymptotically to a constant. The ZZ estimate is extremely poor on the very coarse meshes, which is due to the fact that it does not approximate the L^2 component of the energy norm

Problem 1						Problem 2					
Vertices	AB	AO	BW	GK	ZZ	Vertices	AB	AO	BW	GK	ZZ
25	0.83	1.36	0.72	0.69	0.01	25	0.87	2.45	0.82	0.74	0.01
57	0.85	1.42	0.73	0.69	0.01	81	0.83	1.40	0.71	0.67	0.01
121	0.86	1.44	0.74	0.69	0.02	145	0.87	1.55	0.74	0.68	0.02
249	0.90	1.47	0.76	0.70	0.04	278	0.91	1.68	0.78	0.69	0.04
505	0.98	1.47	0.82	0.72	0.09	673	0.98	1.49	0.83	0.72	0.09
1017	1.07	1.36	0.88	0.68	0.19	1185	1.07	1.42	0.88	0.67	0.18
2041	1.12	1.12	0.94	0.60	0.36	2288	1.11	1.39	0.93	0.59	0.35
4089	1.13	1.13	1.00	0.63	0.56	4929	1.13	1.17	0.99	0.62	0.56
8182	1.20	1.20	1.00	0.66	0.66	9181	1.19	1.32	0.99	0.64	0.64
28533	1.28	1.28	1.00	0.72	0.72	31491	1.27	1.33	1.00	0.71	0.71
75598	1.28	1.28	1.00	0.75	0.69	81309	1.28	1.34	0.99	0.74	0.68
277101	1.27	1.27	1.00	0.78	0.65	293067	1.28	1.30	1.00	0.77	0.65

Table 1: The effectivity ratios of a number of error estimates when the adaptive algorithm (based upon the elementwise energy norm of the exact error) is applied to the two test problems with a maximum permitted aspect ratio of 1. The first column for each problem gives the total number of vertices in each mesh.

of the error. Once this component has been driven down the ZZ estimate performs well: consistently yielding an effectivity index of between about 0.6 and 0.7.

In order to assess the suitability of these indicators for our purposes, we return to the calculations that led to the best results shown in Figure 2 (i.e. with a maximum aspect ratio of 256). Table 2 presents the results in a form analogous to the former table. These results show that the GK estimate is the only one that performs uniformly well for the whole range of mesh sizes and aspect ratios (as predicted by the theory [21]). On very coarse meshes, where all elements have an aspect ratio of one, the AB and the BW estimates both appear to perform quite well, and as the meshes are refined uniformly the effectivity index gets closer to one in each case. However, when anisotropic refinement begins to occur the quality of these estimates tends to deteriorate as the maximum aspect ratio grows. This is clearly a potentially undesirable property for our approach. Similar behaviour is observed for the AO algorithm when anisotropic refinement occurs, although this is perhaps not surprising since this estimate yields the same approximation as AB when an element is sufficiently small. For the coarse initial meshes the AO estimate always overestimates the error. (The AO and AB estimates in their theoretical definition, i.e. by solving infinite-dimensional problems, always overestimate the error, but the practical estimates can underestimate it.) As in the previous test the ZZ error estimate underestimates the error on coarse meshes severely. However most significantly, and in contrast to the AO, AB and BW estimates, the ZZ estimate does not appear to be adversely affected by the increasing aspect ratio.

Having assessed the effectiveness of our selected error estimators on sequences of meshes determined from the exact error, we now contrast these meshes with those obtained when the adaptivity is driven by the estimated errors. Results for the same two examples, with the same maximum aspect ratio of 256, are presented in Figure 4. The graphs shown are of the exact error in each case but do not include the graph for the AO estimate. This is because it turns out that, despite providing different numerical values to AB on the coarse initial grids, this estimate leads to very similar sequences of grids to those obtained using AB in both examples (the graphs of error against unknowns are almost indistinguishable).

The graphs shown in Figure 4 are again typical of those obtained for other large values of κ . In

Problem 1						Problem 2					
Vertices (AR)	AB	AO	BW	GK	ZZ	Vertices (AR)	AB	AO	BW	GK	ZZ
25 (1)	0.83	1.36	0.72	0.69	0.01	25 (1)	0.87	2.45	0.82	0.74	0.01
57 (1)	0.85	1.42	0.73	0.69	0.01	81 (1)	0.83	1.40	0.71	0.67	0.01
121 (1)	0.86	1.44	0.74	0.69	0.02	145 (1)	0.87	1.55	0.74	0.68	0.02
249 (1)	0.90	1.47	0.76	0.70	0.04	278 (1)	0.91	1.68	0.78	0.69	0.04
505 (1)	0.98	1.47	0.82	0.72	0.09	673 (1)	0.98	1.49	0.83	0.72	0.09
1017 (1)	1.07	1.36	0.88	0.68	0.19	1185 (1)	1.07	1.42	0.88	0.67	0.18
2041 (1)	1.12	1.19	0.94	0.60	0.36	2288 (1)	1.11	1.39	0.93	0.59	0.35
2567 (2)	1.20	1.21	1.03	0.68	0.56	3407 (2)	1.19	1.25	1.02	0.67	0.55
3086 (4)	1.40	1.41	1.08	0.74	0.66	4085 (4)	1.38	1.50	1.07	0.72	0.63
4643 (8)	1.63	1.63	1.15	0.70	0.72	7351 (8)	1.60	1.70	1.13	0.69	0.70
7256 (16)	1.84	1.84	1.20	0.65	0.70	10533 (16)	1.78	1.91	1.17	0.63	0.66
13074 (32)	2.16	2.16	1.36	0.62	0.65	22090 (32)	2.09	2.15	1.31	0.60	0.62
24352 (64)	2.69	2.69	1.62	0.60	0.62	37921 (64)	2.56	2.59	1.54	0.59	0.60
45350 (128)	3.44	3.44	2.20	0.58	0.61	85765 (128)	3.24	3.25	2.06	0.58	0.59
83765 (256)	4.42	4.42	3.38	0.57	0.60	138298 (256)	3.93	3.93	2.99	0.59	0.58
261030 (256)	3.65	3.65	3.17	0.57	0.59	368598 (256)	3.30	3.30	2.84	0.61	0.58

Table 2: The effectivity ratios of a number of error estimates when the adaptive algorithm (based upon the elementwise energy norm of the exact error) is applied to the two test problems with a maximum permitted aspect ratio of 256. The first column for each problem gives the total number of vertices in each mesh along with the maximum aspect ratio of any rectangle in that mesh.

both examples we see that the AB error estimate proves to be a better driver of the adaptivity than BW, despite there being little to choose between them from the results presented in Table 2. In particular, once the aspect ratio becomes significantly greater than one, the AB estimate appears to provide the better indication of the relative sizes of $\|\frac{\partial e^h}{\partial x}\|$ and $\|\frac{\partial e^h}{\partial y}\|$ on each element.

The ZZ estimate also leads to interesting behaviour in these two tests. In each case it begins anisotropic refinement much sooner than any of the other estimates permit since it has no approximation of the L^2 component of the error. Hence the condition

$$\kappa^2 \|e^h\|^2 \leq C \| \|e^h\| \|^2 \quad (7)$$

is always satisfied. In the first example this proves to be advantageous (thus demonstrating that our choice of $C = 0.5$ in (7) is over-cautious in this case), leading to the maximum aspect ratio of 256 being reached far sooner than when the exact error is used. Since the ZZ estimate is apparently unaffected by these large aspect ratios it continues to do well and ultimately leads to meshes of almost identical quality to those obtained using the exact error to drive the adaptivity. In the second example however, where the solution is a non-zero function away from the boundary, the ZZ estimate performs less well. Again it leads to anisotropic refinement in the boundary layer sooner than the other estimates but, since the L^2 component of the error in the interior has not been eliminated at this stage, the adaptive algorithm runs into difficulties later on. In particular, the actual error in the interior of the domain completely dominates at latter stages of the calculation and is significantly under-estimated by ZZ. This can be seen to result in only small improvements being made between each refinement and solve.

The GK estimate produces a sequence of meshes with an error behaviour closest to the AB estimate. In a first stage both estimators lead to similar meshes, in a second phase the AB estimate




Figure 4: Graphs of the error as a function of the number of unknowns for the two test problems (with $\kappa = 10^3$ and a maximum aspect ratio 256). Each problem is solved five times: the adaptivity being driven by the exact error and the error estimates AB, BW, ZZ, and GK.

is superior while in a third phase this gap is closed by the better performing GK estimate. Clear advantages of the GK estimate are that it yields a better estimate of the norm of the error, see above, and a simpler implementation. In our tests, however, this estimator needs more refinement steps to drive the error down than the exact error or the AB estimate (even though we are careful to ensure that the refinement regimes are the same in all other respects for each of our computations).

Overall therefore, we see that, of the error estimates considered, the AB and GK estimates appear to be the most appropriate to use in the practical situation where the exact error is unavailable, with AB producing slightly better sequences of meshes and GK producing more reliable numerical error estimates.

4 Implementation issues

The results presented in the preceding sections are a small selection from a much larger number of computations, performed with a variety of different parameter choices. In this section we provide a brief overview of some of these parameters, such as the maximum aspect ratio, the interval in (5) and the constant C in (7), and discuss the significance of the particular values selected. We also make some observations on the practical implementation of our chosen error estimates.

Perhaps the most fundamental parameter in our anisotropic algorithm is the maximum permitted aspect ratio, which is taken to be 256 for the examples in Section 3. One advantage of this choice is that, for the sizes of mesh that we consider, the maximum AR is reached before the end of the refinement process and so we are able to observe that the rate of convergence reverts back to approximately one for further refinements. From a theoretical point of view however there is a strong argument against imposing any such upper limit. Instead one could just rely on the refinement selection mechanism (5) to decide whether anisotropic refinement is no longer appropriate since, for any fixed choice of κ , this situation should eventually arise. Practically however this approach leads to a number of difficulties (although, if these can be overcome, permitting larger aspect ratios certainly can deliver superior results in terms of the energy norm of the error versus the number of unknowns).

The growing aspect ratio leads to three types of ill-conditioning which are each addressed here. First, it is necessary to solve the linear systems corresponding to the approximate solutions on the series of meshes. If very anisotropic elements are allowed then these algebraic systems may become

very ill-conditioned. This is clearly a demand on the preconditioner used when solving these systems but is not considered here, see for example [7] and the references therein for work in this direction. In our tests only a simple, sub-optimal, tridiagonal preconditioner is used. Therefore this paper does not consider any time comparisons (although it is reasonable to expect that with an optimal solver the solution time will depend solely on the number of degrees of freedom). Finally we note that, in order to concentrate on the discretization errors alone, these linear systems are always solved more accurately than would normally be the case in this study.

Second, the error estimator itself can be influenced by the anisotropy of the single elements. The estimates AO and AB, for example, require the solution of an infinite-dimensional Neumann boundary value problem on each element, which is usually undertaken with the aid of a small finite element calculation (e.g. using nine or ten cubic basis functions or ten piecewise linear basis functions: we have implemented both but use the latter in this work). The difficulty that occurs is different to that described above since, due to the small number of degrees of freedom, a direct solver is best in this case. The issue here is that the error equation, a partial differential equation with Neumann boundary conditions, must be solved on a very thin domain, i.e. the anisotropic element. This can lead to the occurrence of instabilities due to the shape of the element, and is an inherent property of the AO and AB approaches, thus restricting their applicability (this is discussed in Section 3). Note however that this problem does not produce any errors on the approximate solution u^h itself. Also note that this kind of instability is not necessarily inherent in the local error equation approach: the GK estimate can be proved to be computable in a stable manner [21] independent of the aspect ratio used.

Third, we are able to obtain the correct approximate solution for a mesh within the adaptive series only if the element stiffness matrices of all elements are computed accurately enough. Classical finite element routines cannot guarantee this if the coordinates of adjacent nodes coincide in too many bits (the Jacobian of the mapping from the master element to the real element usually requires differences of coordinates). Here, we restricted the number of succeeding subdivisions to about 20 (usually less), therefore this problem does not occur. A much longer series of subdivisions of some elements requires a special alternative calculation of the Jacobians, other than with differences of nodal coordinates, see for example [24].

A further significant parameter is our choice of $C = 0.5$ in (7). Recall that unless (7) is satisfied, anisotropic refinement is not permitted and so increasing C allows such refinement to occur for a larger relative component of the L^2 norm of the error. It is apparent from the first graph in Figure 4 that in some cases $C = 0.5$ is too cautious a choice (since the ZZ estimate performs better than using the exact error, due to the fact that the former has no L^2 component). Our initial implementation of the adaptive algorithm used $C = 1$ however this generally performs poorly on our second example (and problems similar to it). Other moderate choices for the value of C tend to yield similar performance to $C = 0.5$: altering the point at which the anisotropic refinement begins but generally leading to final meshes of a very similar quality.

The choice of the exact interval on the right-hand side of (5) is also not too critical. Clearly when the quotient on the left-hand side is either very large or very small then we would wish to refine anisotropically (provided that (7) is satisfied). Similarly, when the quotient is very near to one, uniform refinement is appropriate. Calculations using the intervals $[\frac{1}{3}, 3]$ and $[\frac{3}{4}, \frac{4}{3}]$ for example both lead to results for which the graphs of error against unknowns are virtually identical to those obtained when using the interval $[\frac{1}{2}, 2]$ in (5).

Our final remarks concerning implementation issues relate to the complexity of the implementation and the cost of execution of each of the error estimates that have been considered. For the exact error calculations on each triangle we have used a 7-point quadrature rule with algebraic degree of precision 5 throughout this work. A small number of calculations with a more accurate formula (37

points and degree of precision 13) show that the 7-point formula is adequate in almost all cases. The only inaccuracies occur on very coarse grids but these tend to be refined in an identical manner whichever formula is used. Other than the numerical calculation of the energy norm of the exact error, the simplest estimate to compute is ZZ since, unlike the other estimates that we have considered, this does not require any error equations to be solved. The most complex estimate to implement and compute is AB (closely followed by AO) as this requires the careful calculation of Neumann data for the error equations on each triangular element. The BW estimate also requires an error equation to be solved on each triangle however the edge data is far simpler to compute and it is only necessary to solve a 3×3 linear system on each element. The GK estimate does not involve the computation of boundary data for the error equation however the numerical integration over the squeezed bubble ansatz functions does require care. Finally, we remark that the total computational cost for all error estimates is proportional to the number of elements in the mesh, regardless of their anisotropy.

5 Discussion

The numerical calculations reported in Sections 2 and 3 are all based upon $\kappa = 10^3$, which is chosen to be representative of large values of κ . When much smaller values of κ are used the anisotropy in the problem decreases and there is less to choose between the different error estimates in terms of the refinements that they induce. In the limiting case where $\kappa = 0$ the solutions to our test problems are simple smooth functions. However, choosing the right-hand side of (1) appropriately it is possible to manufacture an artificial problem whose solution contains a steep boundary or internal layer. In this situation the energy norm reduces to the H^1 semi-norm and the resulting problem has been considered by a number of authors, including [17, 20, 27]. Our approach works well in this case too however, in practice one would always choose an initial mesh that is able to approximate the data, f in (1), accurately (see, for example, [17, 25]). Hence the initial mesh for this type of problem should always be anisotropic when an anisotropic refinement algorithm is available. There are situations in three dimensions however where an anisotropic solution exists to the Poisson problem with an isotropic right-hand side on certain domains: see [4, 20] for example.

There are a number of additional computational comparisons that could be made in order to obtain further data on the performance of our proposed technique. One particular approach would be to produce a hybrid algorithm based upon combining ZZ with an estimate which is able to provide an approximation to the L^2 norm of the error. This could be achieved in at least two ways. For example, the AB estimate, say, could be used until (7) is satisfied and then the ZZ estimate could be used instead. Alternatively, the two estimates could be combined so as to estimate the energy norm of the error by the sum of the L^2 norm of the AB estimate plus the ZZ estimate. Initial experiments with the first of these strategies are encouraging: generally leading to an improvement over the use of AB alone, as illustrated in Figure 5.

Further comparisons, against existing *ad hoc* refinement strategies, such as those mentioned in the introduction, would also be worth undertaking. The difficulty with this however is that most such strategies implicitly link the criteria for adapting the mesh with the process used for executing the adaptivity. This link makes reliable comparisons quite difficult to achieve. A Hessian approach (e.g. [16]) might be used within the context of our existing refinement algorithm however by only performing anisotropic refinement when eigenvectors of the Hessian are nearly parallel with rectangle edges.

One of the main advantages of the simple mesh refinement algorithm used in this work is its lack of complexity. This therefore allows us concentrate on assessing the quality of the information that we are able to extract from the exact error and the selected error estimates. For practical problems however this Cartesian refinement algorithm is not sufficiently general since there is a need to be able




Figure 5: Graphs of the error as a function of the number of unknowns for the two test problems (with $\kappa = 10^3$ and a maximum aspect ratio 256). Each problem is solved three times: the adaptivity being driven by a hybrid of the AB and ZZ estimates, by the AB estimate alone, and by the exact error.

to align an anisotropic mesh with solution features which may occur in arbitrary directions. It is not the goal of this short paper to consider algorithms for undertaking refinement in this general manner, however we do expect the ideas introduced here still to be applicable in such cases. In particular, provided the adaptivity procedure is able to produce an anisotropic mesh that is well aligned with the anisotropy present in the solution (see, for example, [20] for a discussion of a matching function which is able to quantify this), then it should be possible to drive this adaptivity using the approach described in this work. Developing such an adaptive algorithm, which should be robust and, ideally, maintain the hierarchical data structures required for the fast solution of very poorly conditioned systems of equations, is still a topic of current research however.

Acknowledgements

We are grateful to Gerd Kunert for the valuable discussions that we have had with him whilst undertaking this work which was supported by the DAAD and British Council under the ARC programme. The work of SG was supported by a grant from the state of Saxony, FRG.

References

- [1] M. Ainsworth and I. Babuška. Reliable and robust a posteriori error estimation for singularly perturbed reaction–diffusion problems. *SIAM J. Num. Anal.*, 36(2):331–353, 1999.
- [2] M. Ainsworth and J. T. Oden. A posteriori error estimation in finite element analysis. *Comput. Methods Appl. Mech. Eng.*, 142(1-2):1–88, 1997.
- [3] D. Ait-Ali-Yahia, W. G. Habashi, A. Tam, M.-G. Vallet, and M. Fortin. A directionally adaptive methodology using an edge-based error estimate on quadrilateral grids. *Int. J. Numer. Methods Fluids*, 23:673–690, 1996.
- [4] Th. Apel. *Anisotropic finite elements: Local estimates and applications*. Advances in Numerical Mathematics. Teubner, Stuttgart, 1999. Habilitationsschrift.

- [5] Th. Apel and F. Milde. Comparison of several mesh refinement strategies near edges. *Comm. Numer. Methods Engrg.*, 12:373–381, 1996.
- [6] Th. Apel, R. Mücke, and J. R. Whiteman. Incorporation of a-priori mesh grading into a-posteriori adaptive mesh refinement. In A. Casal, L. Gavete, C. Conde, and J. Herranz, editors, *III Congreso Matemática Aplicada/XIII C.E.D. Y.A. Madrid, 1993*, pages 79–92, Madrid, 1995. Shortened version of Report 93/9, BICOM Institute of Computational Mathematics, 1993.
- [7] Th. Apel and J. Schöberl. Multigrid methods for anisotropic edge refinement. *SIAM J. Numer. Anal.*, 40:1993–2006, 2002.
- [8] R. E. Bank and R. K. Smith. Mesh smoothing using a posteriori error estimates. *SIAM J. Numer. Anal.*, 34:979–997, 1997.
- [9] R. E. Bank and A. Weiser. Some a posteriori error estimates for elliptic partial differential. *Math. Comp.*, 44:283–301, 1985.
- [10] R. Beinert and D. Kröner. Finite volume methods with local mesh alignment in 2-D. In *Adaptive Methods – Algorithms, Theory and Applications*, volume 46 of *Notes on Numerical Fluid Mechanics*, pages 38–53, Braunschweig, 1994. Vieweg.
- [11] J. U. Brackbill. An adaptive grid with directional control. *J. Comp. Phys.*, 108:38–50, 1993.
- [12] G. C. Buscaglia and E. A. Dari. Anisotropic mesh optimization and its application in adaptivity. *Int. J. Numer. Methods Eng.*, 40(22):4119–4136, 1997.
- [13] M. J. Castro-Díaz, F. Hecht, and B. Mohammadi. New progress in anisotropic grid adaption for inviscid and viscous flow simulations. In *Proceedings of the 4th Annual International Meshing Roundtable*, pages 73–85, Albuquerque, NM, 1995. Sandia National Laboratories. Also Report 2671 at INRIA.
- [14] E. F. D’Azevedo and R. B. Simpson. On optimal triangular meshes for minimizing the gradient error. *Numer. Math.*, 59:321–348, 1991.
- [15] M. Dobrowolski, S. Gräf, and C. Pflaum. On a posteriori error estimators in the finite element method on anisotropic meshes. *Electronic Transactions Num. Anal.*, 8:36–45, 1999.
- [16] V. Dolejší. Anisotropic mesh adaptation for finite volume and finite element methods on triangular meshes. *Comput. Vis. Sci.*, 1(3):165–178, 1998.
- [17] W. Dörfler. A convergent adaptive algorithm for Poisson’s equation. *SIAM J. Numer. Anal.*, 33:1106–1124, 1996.
- [18] T. Iliescu. A 3D flow-aligning algorithm for convection-diffusion problems. *App. Math. Letters*, 12(4):67–70, 1999.
- [19] R. Kornhuber and R. Roitzsch. On adaptive grid refinement in the presence of internal and boundary layers. *IMPACT of Computing in Sci. and Engrg.*, 2:40–72, 1990.
- [20] G. Kunert. *A posteriori error estimation for anisotropic tetrahedral and triangular finite element meshes*. PhD thesis, TU Chemnitz, 1999. Logos, Berlin, 1999.
- [21] G. Kunert. Robust local problem error estimation for a singularly perturbed problem on anisotropic finite element meshes. *Math. Model. Numer. Anal.*, 35:1079–1109, 2001.

- [22] G. Kunert. A note on the energy norm for a singularly perturbed model problem. *Computing*, 69:265–272, 2002.
- [23] N. Madden and M. Stynes. Efficient generation of oriented meshes for solving convection-diffusion problems. *Int. J. Num. Meth. Eng.*, 40:565–576, 1997.
- [24] A. Meyer. The adaptive finite element method — can we solve arbitrarily accurate? Preprint SFB393/01-30, TU Chemnitz, 2001. Submitted to *Computing*.
- [25] P. Morin, R. H. Nochetto, and K. Siebert. Data oscillation and convergence of adaptive FEM. Preprint 17/1999, Albert-Ludwigs-Universität Freiburg, Mathematische Fakultät, 1999.
- [26] J. Peraire, M. Vahdati, K. Morgan, and O. C. Zienkiewicz. Adaptive remeshing for compressible flow computation. *J. Comp. Phys.*, 72:449–466, 1987.
- [27] W. Rachowicz. An anisotropic h-type mesh refinement strategy. *Comput. Methods Appl. Mech. Engrg.*, 109:169–181, 1993.
- [28] E. Rank, M. Schweingruber, and M. Sommer. Adaptive mesh generation and transformation of triangular to quadrilateral meshes. *Comm. Numer. Methods Engrg.*, 9:121–129, 1993.
- [29] W. Rick, H. Greza, and W. Koschel. FCT-solution on adapted unstructured meshes for compressible high speed flow computations. In E. H. Hirschel, editor, *Flow simulation with high-performance computers I*, volume 38 of *Notes on Num. Fluid Mechanics*, pages 334–438. Vieweg, 1993.
- [30] K. G. Siebert. An a posteriori error estimator for anisotropic refinement. *Numer. Math.*, 73(3):373–398, 1996.
- [31] R. B. Simpson. Anisotropic mesh transformation and optimal error control. *Appl. Numer. Math.*, 14:183–198, 1994.
- [32] T. Skalický and H.-G. Roos. Anisotropic mesh refinement for problems with internal and boundary layers. *Int. J. Numer. Methods Eng.*, 46:1933–1953, 1999.
- [33] O. C. Zienkiewicz and J. Wu. Automatic directional refinement in adaptive analysis of compressible flows. *Internat. J. Numer. Methods Engrg.*, 37:2189–2210, 1994.
- [34] O. C. Zienkiewicz and J. Z. Zhu. A simple error estimator and adaptive procedure for practical engineering analysis. *Internat. J. Numer. Methods Engrg.*, 24:337–357, 1987.