Optimal Good-Aspect-Ratio Coarsening for Unstructured Meshes

Gary L. Miller* Dafna Talmor* Shang-Hua Teng†

Abstract

A hierarchical gradient of an unstructured mesh M_0 is a sequence of meshes $M_1, \ldots, M_k$ such that $|M_k|$ is smaller than a given threshold mesh size b. The gradient is well-conditioned if for each i in the range $1 \leq i \leq k$, (1) M_i is well-shaped, namely, elements of M_i have a bounded aspect ratio; and (2) M_i is a coarsened approximation of M_{i-1}. The gradient is node-nested if the set of the nodes of M_i is a subset of that of M_{i-1}. The problem of constructing well-conditioned coarsening gradients is a key step for hierarchical and multi-level numerical calculations. In this paper, we give an algorithm for finding a well-conditioned hierarchical gradient of a two dimensional unstructured mesh. Our algorithm can be used to generate both node-nested and non-nested gradients. The gradient $M_1, \ldots, M_k$ we generate is optimal in the following sense: there exists a constant c such that for any other well-conditioned hierarchical gradient $M'_1, \ldots, M'_k$, $|M_i| \leq c|M'_i|$, that is, the size of the mesh at each level is smaller up to a constant factor.

1 Introduction

The class of hierarchical and multi-level techniques has become one of the most effective and successful classes of numerical techniques for solving partial differential equations (PDEs). These techniques have been used in multigrid methods [4] and multi-level domain decomposition [5].

Numerical methods such as the finite element, finite difference, and finite volume methods apply the following five basic steps to solve a PDE over a domain Ω.

1. Formulate the problem (e.g., in term of PDEs) and describe the geometry and the boundary of a continuous domain D (sometime called geometric modelling).

2. Generate a well-shaped mesh M to approximate the domain.

3. Generate a system of linear or non-linear equations over M for the governing PDEs (e.g., assemble the stiffness matrix and the right hand vector).

4. Solve the system of equations and estimate the error of the solution.

5. Adaptively refine the mesh and return to step (3) if needed.

Once the mesh M is generated, we need to solve a system of linear equations defined over M. A hierarchical method solves this linear system by first constructing a *hierarchical gradient* of meshes $M_0, \ldots, M_k$, where $M_0 = M$ is the finest mesh that discretizes Ω. For each i in the range $1 \leq i \leq k$, the mesh M_i is a geometric coarsening of M_{i-1}. The gist of multigrid methods and other hierarchical numerical methods is the transformation of partial solutions from mesh M_i to mesh M_{i-1} using interpolation, and from mesh M_i to mesh M_{i+1} using restriction. Informally, these hierarchical methods solve a PDE on Ω by first obtaining an initial vector solution either for M_0 or for M_k, and then improving the quality of the vector by transforming it hierarchically up and down the hierarchy while applying some simple and efficient iterative methods at each level.

The simplest form of a *hierarchical gradient* is a series of nested structured meshes (regular grids). Brandt showed, by carefully using restriction and interpolation, that the solution for M_0 can be obtained very efficiently using multigrid methods. See also Bramble, Pasciak and Xu [10]. Nested structured hierarchical gradients are attractive choices in practice because they can be easily generated and because the convergence of the structured multigrid methods is well understood. However, the use of structured regular grids limits the applicability of this simplest class of hierarchical gradients to problems whose domains are simple and whose solution functions have small or constant Hessian [12, 5, 18].

The use of unstructured meshes is inevitable in the solution of complex problems with more intricate domain geometry and solutions. This paper concerns with the problem of generating quality hierarchical gradients for unstructured meshes.

The effectiveness of a hierarchical method that uses

*School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213 {glmiller,tdafna}@cs.cmu.edu.
Supported in part by NSF Grant CCR-9505472

†Department of Computer Science, University of Minnesota, Minneapolis, MN 55455 steng@cs.umn.edu.
Supported in part by an NSF CAREER award (CCR-9502540) and an Alfred P. Sloan Research Fellowship

an unstructured hierarchical gradient $M_0, \ldots, M_k$ depends on the quality of this gradient [10, 5, 6]. In particular, Chan and Zou provided sufficient conditions for multilevel additive Schwarz methods to work on unstructured meshes. Informally, their conditions require, for each i in the range $1 \leq i \leq k$, that (1) M_i is well-shaped, e.g., in two dimensions elements of M_i should have a bounded aspect ratio; and (2) M_i approximates M_{i-1} in the numerical formulation. The coarsening problem can thus be informally defined as: Given a well-shaped mesh M_0 and a threshold size b, construct a gradient $M_1, \ldots M_k$ with $|M_k| \leq b$ that satisfies conditions (1) and (2).

In this paper, we give an algorithm for the coarsening problem. Our algorithm guarantees the quality of the hierarchical gradient (i.e., it produces gradients that satisfies conditions (1) and (2)). It also minimizes the size of the mesh at each level up to a constant factor. A gradient is *node-nested* if the set of the nodes of M_i is a subset of that of M_{i-1}. Our algorithm can be used to generate both node-nested and non-nested coarsening gradients.

A version of this research, for the simpler case of quasi-uniform mesh coarsening, was included in a survey of our work we submitted to the 5th international meshing roundtable [13]. In this paper we address the problem of general unstructured mesh coarsening.

2 The problem of Mesh Coarsening

2.1 Mesh qualities. A two dimensional *domain* Ω is a planar straight-line graph (PSLG), whose boundary is polygonal. The edges in the PSLG can represent boundaries between two materials, points of special interest or holes. In this extended abstract we assume a simple form of this general definition. Our domain is the unit square, and its boundary is the square's four edges.

A *mesh* is a discretization of the domain. The discrete components of the mesh are the mesh *elements.* We mostly discuss triangular meshes, whose elements are triangles. We will often refer to the following three categories of meshes: (1) *grids:* a mesh whose elements are squares of equal size. (2) *quasi-uniform unstructured meshes:* a mesh whose elements' side lengths differ by at most a constant factor. (3) *unstructured meshes:* a general mesh, with no restriction of the elements' size and shape.

Not all meshes perform equally well in numerical computations. Numerical and discretization error depend on the geometric shape and size of the mesh elements [17]. We will use the following definitions for quantifying the geometric shape and mesh element size.

DEFINITION 2.1. *The* **edge-length** *function of a mesh M, el_M, is defined for each $x \in \Omega$ to be the length of the longest edge of all the mesh elements that contain x.*

DEFINITION 2.2. *The* **aspect ratio** *of a triangular mesh element is the smallest angle of the element. The aspect ratio of a mesh is the smallest angle of its elements.*

A mesh is said to be of bounded aspect ratio if its aspect ratio is larger than θ, where θ is a predefined parameter quantifying the mesh quality. We note that there are many definitions for the aspect ratio of a two dimensional mesh which are interchangeable with the above definition [2].

The *el* function is a measure of the mesh element size; the aspect ratio is a measure of the element shape. These two qualities are often at odds: to produce a good aspect ratio mesh conforming to the boundary the mesh often has to be finer. The goal of *mesh generation* is to construct a good aspect ratio mesh whose element sizes are as large as possible (so that the number of mesh elements is as small as possible).

2.2 Mesh coarsenings. A coarsening M' of a mesh M is a mesh whose edge length function $el_{M'}$ is pointwise bigger than el_M but still conforms to the same domain.

The coarsening can be classified as *element-nested, node-nested* or *non-nested.* In general, a triangular mesh does not have any element-nested coarsening, unless it was carefully crafted as such. Furthermore, coarsening, even in the relaxed sense of *node-nested* meshes, can cause a degradation in the aspect ratio of the coarser mesh compared to that of the finer mesh. In this paper, we are interested in generating a sequence of respectively coarser meshes - a gradient of mesh coarsenings. This sequence will be referred to as the *hierarchical coarsening gradient* of M. The objective of this paper is to develop an automatic mesh coarsener which guarantees good aspect ratio of the entire coarsening gradient. We now introduce some definitions to formalize our discussion.

DEFINITION 2.3. *The* **depth** *of an hierarchical gradient is the number of meshes (levels) in the gradient; The* **width** *of level i of the gradient is $|M_i|$, the number of elements in the mesh M_i.*

DEFINITION 2.4. *Let θ be a constant in the range $0 < \theta < \pi/2$; and let $b, \mathcal{I}$ be two positive constants. A hierarchical gradient $M_0, \ldots, M_k$ in two dimensions is a $(\theta, \mathcal{I}, b)$–***Well-Conditioned Hierarchical gradient** *if $|M_k| \leq b$ and for each i in the range $1 \leq i \leq k$ each angle β in M_i satisfies $\theta < \beta < \pi - \theta$. Furthermore, every two adjacent meshes are $\mathcal{I}$-*locally similar, *i.e.* $el_{M_{i+1}} \leq \mathcal{I} el_{M_i}$.

The first condition, good aspect ratio, is motivated by the requirement from iterative methods: The multigrid method and multi-level domain decomposition use an iterative method to "smooth" the residual error at each level. The convergence properties of iterative methods are related to the aspect ratio of the underlying mesh. Chan and Zou [6] showed that bounded aspect ratios at each level are important for both additive Schwarz based multi-level domain decomposition and multigrids. The second condition, the local similarity, is motivated by the restriction and interpolation phases of the multigrid methods, which are used to transform partial solution between meshes in adjacent levels of the hierarchy. To reduce the interpolation and restriction errors, adjacent meshes should approximate each other well. Both local similarity and bounded aspect ratio were used in Chan and Zou's analysis, which showed that they are sufficient for multilevel additive Schwarz methods to work on unstructured meshes.

Given M_0 and $(\theta, \mathcal{I}, b)$, the problem of **hierarchical mesh coarsening** is to find a $(\theta, \mathcal{I}, b)$-well conditioned hierarchical gradient with smallest depth and width.

2.3 Previous approaches to mesh coarsening.

To produce a node nested coarsening, a subset of the fine mesh nodes is picked, and retriangulated to form the coarse mesh. Various approaches to mesh coarsening differ in the methods they use to pick the coarser mesh nodes, and the retriangulation method. The problem of mesh coarsening received much attention, see [9, 19, 11, 15, 1, 7]. However, none of the papers we found address the geometrical issues of mesh coarsening. Properties such as element quality of the mesh or the size of the intersection of the coarse and fine mesh (their local similarity) are not discussed beyond empirical observations.

Picking the node set of the coarser mesh using a maximal independent set (MIS) technique seems to be the most popular approach [5, 9, 19, 11, 15, 7]. The 1D skeleton of the fine mesh is viewed as a graph, and a set of nodes such that no two share an edge is picked (independence). A node not picked for the coarse mesh must neighbour a node that was retained, so the independent set is maximal. An MIS can be constructed to ensure a constant factor reduction in the number of the mesh nodes. However, we now show that the MIS technique can not guarantee the aspect ratio of the coarsened meshes, both for quasi-uniform and for general unstructured meshes.

Quasi-uniform unstructured mesh coarsening: The MIS on the 1D skeleton technique is very successful in reducing the mesh size to a fraction of its original size; however it carries no guarantees for the other qualities of the mesh hierarchy, such as its aspect-ratio. The problem is illustrated in Figure 1: certain choices of an MIS of the original mesh degrade the aspect ratio of the coarser mesh. The aspect ratio degradation compounds with repeated applications. This can be observed even for very uniform meshes, as in the grid-like mesh of the figure.

Figure 1: Repeated applications of MIS can degrade the aspect ratio.

General unstructured mesh coarsening: Above we gave an example of a quasi-uniform mesh for which certain choices of the MIS cause repeated degradation of its aspect ratio. For general unstructured meshes, a much stronger statement is true: there exists an unstructured mesh such that all possible choices of MIS result in a coarsening hierarchy with increasingly worse aspect ratio.

Let M_0 be a one dimensional mesh whose nodes are $P = \{2^{i-1} - 1 : i = 1 \cdots n\}$. The edges of the mesh are between adjacent points, M_0 is a line graph of n nodes. The *aspect ratio* of a one dimensional mesh is the maximum ratio between two adjacent edges, hence the aspect ratio of M_0 is equal to 2.

For this mesh there is no well-conditioned coarsening gradient of depth $\log n$. This statement is a corollary of the results presented in this paper. Since the MIS technique reduces the size of the mesh by a constant factor at each level, it fails to produce good aspect ratio coarsenings.

In particular, the MIS technique defines a hierarchical gradient where the jth mesh $\{M_j\}$ has node set $P_j = \{2^{i2^j-1} - 1 : i = 1 \cdots n/2^j\}$. In other words, P_j is formed by taking every other point of $\mathcal{P}_{j-1}$. The aspect ratio of mesh M_j is therefore at least $2^{2^j} - 1$, and worsens super-exponentially. This one dimensional example can easily be extended to two and three dimensions.

As part of the results in this paper, we will obtain bounds for the shortest (up to a constant) well-conditioned gradient for this mesh and provide simple algorithms for its generation.

3 Function Based Mesh Coarsening

In this section, we present our approach to mesh coarsening. In order to construct a well-conditioned hierarchical gradient we have to overcome the escalating degradation of the mesh quality, demonstrated in the previous section.

Rather than using only the information present in the respectively coarser meshes, we use an intermediate representation which we term "spacing functions". The spacing functions, along with the point set of the initial mesh, capture the information necessary for the generation of the coarser meshes.

At a high level, M_0 defines a spacing function f_0 which describes the typical size and point spacing of M_0. Our idea is to compute a spacing function f_i for each level, and use it to generate M_i. Given the spacing functions our task is then to create a point set that is "spaced" according to that function, and triangulate it (as for mesh generation [12]). We refer this mesh coarsening technique as **function-based coarsening**. It contains four steps: (1) recover the spacing function of the initial mesh; (2) increase the spacing value of the mesh points smoothly to obtain the new spacing functions; (3) delete some mesh nodes so that the remainder nodes are spaced according to the new spacing function; and (4) compute the Delaunay triangulation of the nodes obtained in Step (3).

3.1 Recovering the spacing function: We first formalize our notion of a spacing–function.

DEFINITION 3.1. *Let $\beta > 1$ be a real number. A point set P is β-***spaced** *according to a function f if for any two points $p, q \in P$, $f(p_1) + f(p_2) < \beta\|p - q\|$. The function f is then referred to as the β-spacing function of P.*

The initial spacing function we use is based on the natural spacing of the original mesh:

DEFINITION 3.2. *The* **nearest neighbour** *(NN) function of a point set $P \subset \Omega$ assigns to each point $p \in P$ the distance to the point $q \in P$ nearest to it such that $q \neq p$. It can be extended to the domain Ω by assigning to a point $x \in \Omega$ the radius of the smallest closed ball centered at x and containing at least two points from P.*

3.2 Coarsening the spacing functions:

DEFINITION 3.3. *Let P be a point set in a domain Ω in $\mathbb{R}^2$. Let g be a spacing function over Ω. Let $C > 1$ be a real number. The C-***coarsening** *of f with respect to P is a spacing function over Ω such that for all $x \in \Omega$*

$$f_{g,C,P}(x) = \min_{p \in P} \hat{f}_{g,C,p}(x),$$

where for each point $p \in P$, $\hat{f}_{g,C,p}(x) = C \cdot g(p) + \|p - x\|$. When clear from the context, we omit g from the notation.

We can use a simpler way to generate the coarsening function for quasi-uniform functions g:

Procedure: ONE_LEVEL_COARSEN(M)

Input: M, a mesh over a square Ω.
C, the coarseing factor.
g, the β-spacing function of M.

Output: M_1, a coarser triangular mesh.

Method:

1. Let P_0 be the square's corners; P_1 the mesh nodes located on the square's edges; P_2 the rest of the mesh nodes.
2. Compute $f(p_i) = f_{g,C,P}(p_i)$ for each $p_i \in P$.
3. Let $\hat{P}_2$ be the set of points $\{p\}$ in P_2 whose distance from the boundaries is at least $\delta f(p)$. (δ is a small fixed constant.)
4. Construct a conflict graph with respect to f: $CG(P_0 \cup P_1 \cup \hat{P}_2)$.
5. Let S be a maximal independent set of CG generated by first considering the points P_0, then P_1, and finally $\hat{P}_2$.
6. Return $M_1 = DT(S)$, the Delaunay triangulation of S.

Figure 2: One level function based coarsening.

DEFINITION 3.4. *Let P be a point set in a domain Ω in $\mathbb{R}^2$. Let g be a spacing function over Ω. Let $C > 1$ be a real number. The C-***threshold-coarsening** *function is defined as:*

$$t_{g,C,P}(x) = \max(g(x), C \min_{y \in \Omega} g(y))$$

The set of coarsening functions we suggest is therefore:

- **for general coarsening:** $\{f_{NN,2^i,P}\}$ and β is a constant depending on the aspect ratio (see section 5).
- **for quasi–uniform coarsening:** $\{t_{NN,2^i,P}\}$ and $\beta = 2$.

3.3 Coarsening the meshes: Let M_0 be the initial mesh, C the factor by which the mesh should be coarsened. To coarsen the mesh, we first generate the coarsened spacing function values for each node of the mesh using C, and then pick a subset of the mesh nodes which is β-spaced by this spacing function. One possible method to pick the coarsened mesh nodes is by using a conflict graph.

Definition 3.5. *The* **conflict graph** *of a point set P, CG(P), with respect to a β-spacing function f is a graph CG(P) = (P, E) where*

$$E = \left\{ (p_i, p_j) : \|p_i - p_j\| < \frac{f(p_i) + f(p_j)}{\beta} \right\}$$

Our proposed scheme is outlined in Figure 2. This one-level scheme can be naively extended to a multi-level scheme by repeatedly applying it with coarser and coarser spacing functions f_i, generated using function coarsening and constants of the form C^i for level i.

However this in general produces a gradient that is not node-nested, for it generates the mesh at each level independently of other levels. To construct a node-nested hierarchical gradient, we need a subtler approach. At a high level, we first generate M_k, the coarsest mesh. We then enforce that nodes in M_k will be chosen in M_{k-1}. Repeating this enforcement, we can build a node-nested well-conditioned gradient. The resulting scheme is outlined in Figure 3.

Lemma 3.6. *S_{i+1} is also an independent set of CG_i.*

Proof: Because $C > 1$, $C^{i+1} > C^i$. Therefore, for each point $x \in \Omega$, $f_{C^{i+1},P}(x) \geq f_{C^i,P}(x)$. By definition 3.5, CG_{i+1} is a supergraph of CG_i. Hence any independent set of the former must be independent in the later as well. □

Therefore, Algorithm MULTI_LEVEL_COARSEN correctly generates a node-nested multilevel gradient.

The rest of this paper will show the correctness of this coarsening method, stated in the following two theorems:

Theorem 3.7. *Let M_0 be a mesh whose smallest angle is bounded bellow by θ. The hierarchical gradient $(M_1, ..., M_k)$ produced by the algorithm of Figure 3 has the following properties:*

1. **aspect ratio:** *There is a constant θ_1 depending on θ only such that for $1 \leq i \leq k$, the smallest angle of mesh M_i is bounded bellow by θ_1.*
2. **local similarity:** *There is a constant $\mathcal{I}$ depending on θ only such that for each $1 \leq i \leq k$, $el_{M_i} \leq \mathcal{I} el_{M_{i-1}}$.*

Therefore, the hierarchical gradient generated is $(\theta_1, \mathcal{I}_1, |M_k|)$-well-conditioned.

Theorem 3.8. *Let $M'_1,...,M'_k$ be any $(\theta, \mathcal{I}, b)$-well-conditioned gradient of M_0, for some positive constants $\theta, \mathcal{I}$. Let $M_1,...,M_k$ be the result of the application of the algorithm of Figure 3 on M_0 for k iterations, then there is a constant c such that $\forall i : |M_i| \leq c|M'_i|$. Hence $M_1,...,M_k$ is $(\theta_1, \mathcal{I}_1, c_1 b)$-well-conditioned gradient.*

A simple corollary of the last theorem, is that our algorithm optimizes the number of levels up to an

Procedure: MULTI_LEVEL_COARSEN(M_0)

1. Let k be the length of the required hierarchy.
2. Let M_k = ONE_LEVEL_COARSEN($M_0, f_{C^{k+1},P}$). Let S_k be the point set of M_k.
3. For $i = k-1$ to 1
   - Let S_i be a maximal independent set of CG_i that contains S_{i+1}. Note that by Lemma 3.6 S_{i+1} must be an independent set of CG_i.
   - Let $M_i = DT(S_i)$.
4. Return $(M_1, ..., M_k)$.

Figure 3: Multi level nested function based coarseing.

additive constant factor when coarseing the mesh down to a constant sized mesh:

Corollary 3.9. *If $M'_1,...,M'_k$ is $(\theta, \mathcal{I}, 1)$-well-conditioned gradient, then $M_1,...,M_k$ can be completed to an $(\theta_1, \mathcal{I}_1, 1)$-well-conditioned gradient of depth at most $k + c_1$.*

4 Spacing Functions Qualities and Mesh Qualities

This section focuses on the intimate connection between the mesh qualities and its spacing function qualities. In particular, we show that spacing functions can capture the two most important mesh properties: its elements' shape, and its elements' size and number. This connection is the foundation of our coarsening approach correctness, which will be discussed in the next section.

4.1 From a mesh to a spacing function.

The spacing function we recover from each mesh is the Nearest Neighbour function (*NN*), see Definition 3.2. This function depends only on the node set of the mesh. A generalization of this function, the local feature size function, plays an important role in the analysis of mesh generation algorithms [16, 3, 14]. These functions change slowly spatially, in analogy to the slowly changing element sizes of the well-shaped mesh.

Definition 4.1. *A function f is 1-Lipschitz over a domain Ω if for any two points x,y in Ω, $|f(x) - f(y)| \leq \|x - y\|$.*

The proofs of the following two lemmas are elementary:

Lemma 4.2. (Ruppert [16]) *For each point set P in $\mathbb{R}^d$, NN_P is 1-Lipschitz.*

Lemma 4.3. *NN_P is a 2-spacing function of P.*

For a good aspect ratio mesh M, the *el* function (see Definition 2.1) is equivalent up to a constant factor to *NN*.

THEOREM 4.4. (RUPPERT [16]) *Let M be a mesh with smallest angle bound θ. Let P be the node set of M. There exist two positive constants C_1, C_2 depending on θ only such that $C_1 el_M(x) \leq NN_P(x) \leq C_2 el_M(x)$.*

Since every point set P is 2-spaced according to the 1-Lipschitz function NN_P, these notions by themselves are not powerfull enough to describe good aspect ratio meshes. Intuitively, spacing functions prevent points from clustering, but do not prevent the formation of arbitrarily large gaps. We now formalize the notion of a gap in terms of the spacing function.

DEFINITION 4.5. *Let f be an integrable function over a domain Ω. The f-***area** *of a sub-domain B of Ω is given by $A_f(B) = \int_B 1/(f(x))^2 dA$.*
The following lemma shows triangles of a good aspect ratio mesh are of constant *NN*-area, and relates the number of elements to the area of the domain. The proof is omitted from this extended abstract.

LEMMA 4.6. *Let M be a mesh with smallest angle bound θ.*

1. *For each mesh triangle T:*

$$\frac{\sin^2\theta}{2{C_2}^2} \leq A_{NN_M}(T) \leq \frac{1}{2C_1^2 \sin\theta}$$

2. *Let N be the number of triangles, then:*

$$\frac{\sin^2\theta}{2{C_2}^2} N \leq A_{NN_M}(\Omega) \leq \frac{1}{2C_1^2 \sin\theta} N$$

Where C_1 and C_2 are the constants of Theorem 4.4.
We capture the notion of a gap using empty balls. A ball B is called an *empty* ball with respect to a mesh M if B does not contain any node of M. The following lemma states an empty ball can intersect at most a constant number of triangles. The proof is omitted from this paper.

LEMMA 4.7. *Let M be a mesh with smallest angle bound θ. Any empty ball B intersects at most $C_3 = 4 + 3\frac{4}{\sin^4\theta}$ triangles of M.*
Finally, the following theorem states the relationship between a good aspect ratio mesh and spacing functions. In particular, a good aspect ratio mesh posseses a 2-spacing function and has no large gaps in terms of that spacing function. The next sub-section shows the inverse statement is true as well. Theorem 4.8 is a simple corollary of Lemmas 4.6 and 4.7.

THEOREM 4.8. *Let M be a mesh with smallest angle bound θ.*

1. *Let $C_4 = C_3 \frac{1}{2\sin\theta C_1^2}$. The NN_M-area of each empty ball is at most C_4.*
2. *Let $L = 2 \cdot 8^{C_4}$, and let x be any point of Ω. Any ball containing x whose radius is greater than $L \cdot NN_M(x)$ must contain some node p of the mesh.*

The 1-dimensional analogs of the notion of gaps, and their related theorems, are valid for the function *NN* on the boundaries of Ω (in our case the four edges of the unit square). They were omitted for the sake of brevity.

4.2 From a spacing function to a mesh.

In this section we show that a point set spaced according to a spacing function, where no large gaps in terms of the spacing function occur, is the node set of a good aspect ratio mesh. In particular, we show the Delaunay triangulation of that point set is of good aspect ratio.

The following theorem is the inverse of Theorem 4.8.

THEOREM 4.9. *Let P be a set of points β-spaced by a 1-Lipschitz function f, such that each empty ball B is of constant bounded area, $A_f(B) \leq \alpha$. Then the smallest angle θ of each Delaunay triangle of P is bounded by: $\sin\theta \geq \frac{8^{-\alpha}}{2\beta}$.*

The following theorem shows that the function f β-spacing P is equivalent up to a constant factor to the function NN_P. The connection between f and $el_{\mathrm{DT}(p)}$ then follows by Theorem 4.4.

THEOREM 4.10. *Let P be a set of points β-spaced by a 1-Lipschitz function f, such that for each empty ball B $A_f(B) \leq \alpha$, then:*

$$NN_P(x)/(L^2 + 2L) \leq f(x) \leq (2\beta + 1) NN_P(x)$$

For the sake of brevity, the theorems in this subsection ignored the boundary case. For the analogous results on the boundaries, we further require that the point set is spaced on the boundaries with no large gaps in terms of the spacing function restricted to the boundaries, and that points in the interior of Ω stay away from the boundary: for some constant δ, the distance of each non boundary point p from the boundary must be at least $\delta \cdot f(p)$. The aspect ratio of the Delaunay triangulation then depends on δ as well.

5 Spacing Function Based Coarsening

This section is concerned with the correctness of our coarsening approach, as stated in Theorems 3.8 and 3.7. In particular, we show that given an initial bounded aspect ratio mesh and a coarsening factor C, our algorithm produces a bounded aspect ratio mesh which is (up to a constant factor) the smallest mesh whose *el* function is at most C times larger (point wise) than the *el* function of the original mesh.

5.1 Aspect ratio of the coarser meshes.

We first establish that our coarsening functions are all 1-

Lipschitz. This follows from Definitions 3.3 and 3.4.

LEMMA 5.1. *If* g *is* 1*-Lipschitz, then for any coarsening factor* $C \geq 0$, $f_{g,C,P}$ *and* $t_{g,C,P}$ *are* 1*-Lipschitz.*

The next theorem establishes the main result of this sub-section: that the mesh generated by our one-level coarsening algorithm, see Figure 2, is a bounded aspect ratio mesh.

THEOREM 5.2. *Let* P *be the point set of a mesh* M_0 *with smallest angle bound* θ. *Let* $C > 1$ *be the coarsening factor. Let* M_1 *be the coarser mesh returned by our one-level coarsening algorithm with spacing parameter* $\beta = 12\sqrt{L}$ *and coarsening function* $f_{NN,C,P}$. *Further, assume* $L \geq 4$, *where* L *is the constant of Theorem 4.8. There exists a constant* θ_1 *depending on* β *and* θ *only such that the smallest angle of* M_1 *is larger than* θ_1.

Proof: Let P_1 denote the point set of M_1. P_1 is β–spaced by $f_{NN,C,P}$. Theorem 4.9 implies that if the f-area of any Delaunay ball B is smaller than some positive constant α, then $\sin\theta_1 \geq \frac{8-\alpha}{2\beta}$. Let f stand for $f_{NN,\alpha,P}$. Therefore, we proceed to prove the theorem by showing that the f-area of any empty ball in M_1 is bounded .

We fix some point $p \in P_1$ and a ball B through p. Without loss of generality, let $f(p) = 1$, let $p = (0,0)$, and let the center of B lie on the y-axis, see Figure 4. We show that if the ball B is large, it cannot be empty. Let L be the constant of Theorem 4.8. Let B's radius be $12L$. We show B must contain a point of P_1 in its interior. (Note that B all ready contains p on its boundary).

The center coordinate of B is $(0, 12L)$. Let B_0 be a smaller ball nested in B of radius $3L$ centered at $(0, 2+3L)$. The proof consists of showing that B_0 must contain a point of P, and that this point could not be ruled out by taking a maximal independent set of the conflict graph of P, and hence must belong to P_1 as well.

Since M is a good aspect ratio mesh, Theorem 4.8 implies any ball of radius $LNN_P(x)$ through x must contain a point $q \in P$. Since $f \geq NN_P$, any ball of radius $Lf(x)$ through x must contain a point $q \in P$. Let $x_0 = (0,2)$. x_0 lies on the boundary of B_0, and by the 1-Lipschitz property of f, $f(x_0) \leq f(p) + 2 = 3$. Hence, the ball B_0 which is of radius $3L \geq Lf(x_0)$ must contain a point $q \in P$.

Without loss of generality, let $q = (x, y)$ be on the boundary of B_0. (all the bounds we derive hold if q is internal to B_0 as well). We now show $q \in P_1$ as well. We use the following facts about q:

I. By Lemma 5.3, the distance from q to the boundary of B is larger than $y/2$.

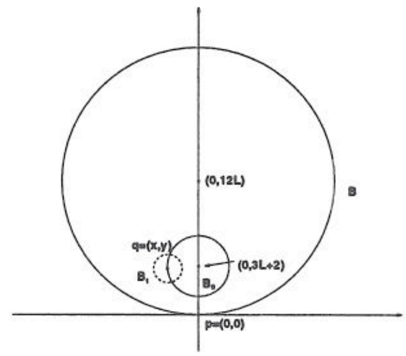

Figure 4: The existence of a mesh point of M_1 in B.

II. By Lemma 5.4, if $L \geq 4$ and $\beta = 12\sqrt{L}$, then $f(q)/\beta \leq y/6$.

Let B_1 be the ball with center q and radius f/β (see the dashed ball in Figure 4.) If point $q \notin P_1$, then there must be some point $w \in P_1$ such that q and w share an edge in the conflict graph used to create P_1. If $w \in B$ then B contains a point of P_1 in its interior, and we are done, therefore w must be outside B and property I implies $d = ||w - q|| \geq y/2$.

Because f is 1-Lipschitz, $f(w) \leq d + f(q)$. q and w share a conflict graph edge, hence $f(q) + f(w) \geq \beta d$. The last two inequalities imply $2f(q) + d \geq \beta d$. By property I, $2f(q) \geq (\beta - 1)d \geq (\beta - 1)y/2$. By property II, $\beta y/3 \geq 2f(q) \geq (\beta - 1)y/2$. Hence $\beta/3 \geq (\beta - 1)/2$. However, the last is a contradiction for $\beta > 3$, as is the case for our β.

This contradiction implies q and w do not share a conflict graph edge, and $q \in P_1$. This in turn implies that if B's radius is greater $12L$, B can not be empty. □

The following two lemmas were necessary for the proof of the theorem:

LEMMA 5.3. *If* $L > 2/3$ *then the distance from any point* (x, y) *on* B_0 *to the boundary of* B *is greater than* $y/2$.

Proof: First notice that $y > 0$. The distance from (x, y) to the boundary of B is equal to $R - \sqrt{x^2 + (y - R)^2}$, where $R = 12L$. Because (x, y) is on the boundary of B_0, we have $x^2 + (y - r - 2)^2 = r^2$, implying $x^2 + y^2 = 2y(2 + r) - 4 - 4r = y(4 + 6L) - 4 - 12L$, for $r = 3L$. This in turn implies, after some manipulations we omit, that $x^2 + (y - r - 2)^2 < (R - y/2)^2$. This last inequality is true when $L > 2/3$, $y > 0$, $L > 0$ and $R = 12L$. Because $y/2 < R$ we have, after taking the square root, $R - \sqrt{x^2 + (y - R)^2} > y/2$, completing the proof. □

LEMMA 5.4. *If $L > 4$ and $\beta = 12\sqrt{L}$, then*

$$\frac{f_{NN,C,P}(q)}{\beta} \le \frac{1 + \|p-q\|}{\beta} \le \frac{y}{6}.$$

Proof:

As shown in the proof of Lemma 5.3 $\|p-q\|^2 = x^2+y^2 = y(4+6L)-4-12L$. Thus, $f(q) \le 1+\|p-q\| \le 1+\sqrt{y(4+6L)-4-12L}$. The Lemma states that $1+\sqrt{y(4+6L)-4-12L} \le \beta y/6$, or equivalently, $\sqrt{y(4+6L)-4-12L} \le \beta y/6 - 1$

Squaring both side, it suffices to show that $y(4+6L)-4-12L < \beta^2 y^2/36 + 1 - \beta y/3$. Hence it is enough to show that $\beta^2 y^2/36 - \beta y/3 - y(4+6L) > 0$. Since $y > 2$ this is true if $\beta^2/18 - \beta/3 - (4+6L) > 0$. The conditions $L > 4$ and $\beta = 12\sqrt(L)$ guarantees this inequality. □

The aspect ratio bound derived in Theorem 5.2 applies to the special case of quasi-uniform meshes as well. However, we can obtain better bounds, using a simpler proof, for quasi-uniform meshes. These bounds are derived for a simpler coarsening spacing function, the threshold spacing function of Definition 3.4.

THEOREM 5.5. *Let P be the point set of a mesh M_0 with smallest angle bound θ. Let $C > 1$ be the coarsening factor. Let M_1 be the coarser mesh returned by our one-level coarsening algorithm with spacing parameter $\beta = 2$ and coarsening function $t_{NN_P,C,P}$. The smallest angle θ_1 of M_1 is bounded bellow by:*

$$\sin\theta_1 \ge \frac{1}{3+4L}$$

where L is the constant from Theorem 4.8, depending on θ only.

5.2 Local similarity. We now show that neighbouring meshes in the hierarchical coarsening gradient generated by our method our localy similar, see Definition 2.4.

The proof of the following lemma is elementary:

LEMMA 5.6. *[linearity and monotonicity] For any $C > 1$ and for any $\gamma \ge 1$. For each point $x \in \Omega$*

$$f_{g,C,P}(x) \le f_{g,\gamma C,P}(x) \le \gamma f_{g,C,P}(x).$$

THEOREM 5.7. *Let M_0 be a mesh whose smallest angle bound is θ. There exists a constant $\mathcal{I}$ depending on θ only such that such that for each mesh M_{i+1} of the hierarchical gradient $(M_1, ..., M_k)$ produced by the Algorithm of Figure 3*

$$el_{M_{i+1}} \le \mathcal{I} el_{M_i}$$

Proof: The coarsening function used to create M_{i+1} is $f_{i+1} = f_{g,2C,P}$, the one used to create M_i is $f_i = f_{g,C,P}$, for some value C. By Lemma 5.6,

$$f_{i+1} \le 2f_i$$

By Theorem 4.10

$$NN_{P_j}/(L^2+2L) \le f_j \le (2\beta+1)NN_{P_j}$$

where P_j is the set of points of mesh M_j. By Theorem 4.4

$$C_1 el_{M_j} \le NN_{P_j} \le C_2 el_{M_j}$$

Hence

$$C_1/(L^2+2L) el_{M_{i+1}} \le f_{i+1} \le 2f_i$$

and

$$2f_i \le 2(2\beta+1)C_2 el_{M_i}$$

and we can take $\mathcal{I}$ to be $2C_2(2\beta+1)(L^2+2L)/C_1$. □

5.3 Size optimality. We now show, up to a constant, that for any $C > 1$ the size of the mesh produced by our one-level coarsening algorithm is the smallest possible. The following Lemma shows that $f_{g,C,P}$ is the largest 1–Lipschitz function that is smaller than $C \cdot g$ at the points of the initial mesh M_0. The proof is elementary, and is omitted from this paper.

LEMMA 5.8. *Let $C > 1$. Let h be a 1-Lipschitz function over the domain Ω such that for all $p \in P$ $h(p) \le C \cdot g(p)$. Then for all $x \in \Omega$, $h(x) \le f_{g,C,P}(x)$.*

We now show that $t_{NN_M,C,P}$ and $f_{NN_M,C,P}$ are equivalent up to a constant factor for a quasi–uniform mesh M. This connection implies it suffices to show that meshes spaced by $f_{NN_M,C,P}$ are of optimal size.

LEMMA 5.9. *Let M be a quasi–uniform mesh, i.e. there exist ratio constant ρ such that:*

$$\min NN_M \le NN_M(x) \le \rho \min NN_M$$

Then

$$f_{NN_M,C,P}/(2L\rho + 2\rho) \le t_{NN_M,C,P} \le f_{NN_M,C,P}.$$

DEFINITION 5.10. *Let M and M' be two well-shaped meshes over a domain Ω. For any positive $C > 1$, we say elements of M' are at most a factor of C larger than those of M if for all $x \in \Omega$, $el_{M'}(x) \le C \cdot el_M(x)$. In other words, for each point $x \in \Omega$, the largest triangle of M' that contains x is no more than a factor C larger than the largest triangle of M that contains x.*

We now state the main result of this sub-section:

THEOREM 5.11. *Let M_0 be mesh with smallest angle bound θ_0. For any $C > 1$, Let M be the mesh obtained by the algorithm of Figure 2 with coarsening factor C. Let M' be a mesh with minimal angle bound θ' whose elements are at most a factor of C larger than those of M_0, then there exists a constant D depending on θ_0 and θ' only such that $|M| \le D|M'|$.*

Proof: By the assumption $el_{M'} \leq Cel_{M_0}$. Since the constant bounds we derived in previous Lemmas depend on the smallest angle, we will refer to them in this proofs as a function of the relevant angle. Let P_0 be the node set of M_0 and let P be the node set of M. By Theorem 4.4,

$$\frac{NN_{M'}}{C_2(\theta')} \leq el_{M'} \leq Cel_{M_0} \leq \frac{C}{C_1(\theta_0)} NN_{M_0}$$

Therefore by Lemma 5.8,

$$NN_{M'} \leq f_{NN_{M_0}, \frac{CC_2(\theta')}{C_1(\theta_0)}, P_0}$$

By Lemma 5.6:

$$NN_{M'} \leq \frac{C_2(\theta')}{C_1(\theta_0)} f_{NN_{M_0}, C, P_0}$$

By Theorem 4.10:

$$NN_{M'} \leq \frac{(2\beta(\theta_0)+1)C_2(\theta')}{C_1(\theta_0)} NN_M$$

where β is the spacing constants used in the coarsening algorithm. Now Lemma 4.6 implies the result. □

A simple application of this Lemma to all the levels of the coarsening gradient proves Theorem 3.8, whereas Theorem 5.2 can be applied to its multi-level version, Theorem 3.7.

6 Practical Concerns

6.1 Mesh quality. The proof outlined in previous sections provided a constant bound on the smallest angle of of the mesh hierarchy. However the bound can be quite small mathematically. Given the practical importance of the coarsening problem, we implemented the algorithm presented in this paper. We now provide some experimental evidence that our coarsening approach indeed produces coarsening meshes of very reasonable quality in practice.

We include the numerical data for the "crack plate" mesh, which was generated by Omar Ghattas and Xiaogang Li of Carnegie Mellon University. The physical problem modeled by the mesh is a plate with a horizontal crack running from the middle of the left edge to the center of the plate [8]. The following Table 1 lists the coarsening hierarchy statistics of our method. See Figure 5 for these meshes.

To offer a comparison, we also implemented a maximal independent set approach. We applied the MIS based program on the same initial mesh and observed a significant and iterative degradation on the smallest angle as shown in Table 2. This smallest angle occurs in the center of the square, so the iterative degradation can not be attributed to edge effects, or our choice of boundary coarsening. The crack mesh is extremely unstructured, and the MIS based program fails to guarantee a bound on the aspect ratio in part because it attempts to reduce the number of mesh elements too much. Note that with our approach, the coarsening factor is doubled in later iterations but not the number of nodes.

coarsening factor	num nodes	num triangles	min angle	max angle
1	5120	9984	41.25	93.88
2.1	2242	4374	18.43	135.01
4	1035	1982	18.23	135.06
8	305	564	18.43	135.01
16	148	272	19.44	123.69
32	119	220	18.01	130.23
64	102	189	18.01	130.23
128	92	170	18.01	130.23

Table 1: Quality of meshes coarsened using a function based approach. The first row describes the original mesh.

coarsening iteration	num nodes	num triangles	min angle	max angle
1	1215	2346	18.43	126.87
2	289	537	14.04	139.40
3	76	131	3.37	135.00

Table 2: Quality of meshes coarsened using a maximal independent set approach.

6.2 Algorithmic efficiency. The algorithm of Figure 2 has a simple $O(n^2)$ time implementation. We an appropriate choice of data structures, we can reduce the algorithmic cost. The two computational expensive steps are the coarsened function construction and the point set selection according these functions.

The spacing function computation as described in Figure 2 takes $O(n^2)$ time. In practice we can use other related spacing function, that still yeild a theoretically correct algorithm. A particulary useful one is the spacing function based on a balanced quad-tree [3] of the fine mesh. The balanced quad-tree can be efficiently coarsened and can be used as an estimation of the spacing function of the underlying mesh. An important next step in our work is to develop such an efficient and practical implementation of our function based coarsening algorithm.

References

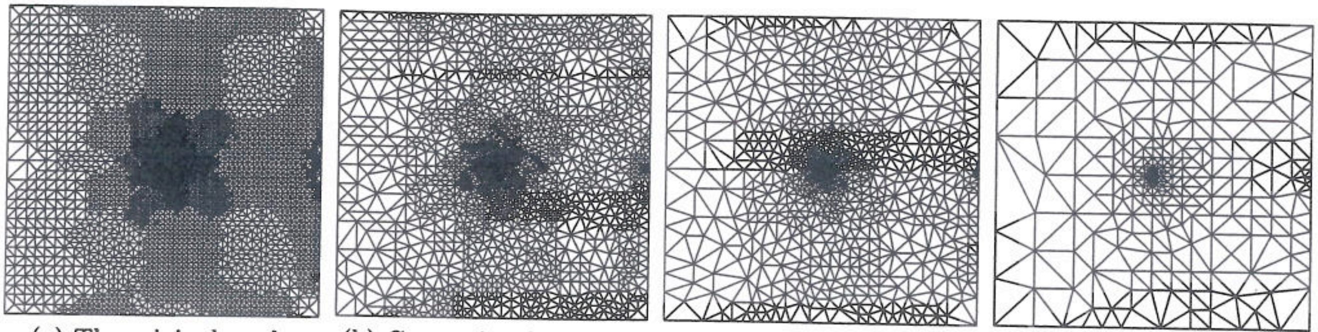

(a) The original mesh. (b) Coarsening factor 2.1. (c) Coarsening factor 4. (d) Coarsening factor 8.

Figure 5: Coarsening the crack plate mesh using a coarsening function approach.

[1] Randolph E. Bank and Jinchao Xu. An algorithm for coarsening unstructured meshes. *To appear in Numer. Math.*

[2] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. Tech. Report CSL-92-1, Xerox PARC, Palo Alto, CA, 1992.

[3] Marshall Bern, David Eppstein, and John Gilbert. Provably good mesh generation. In *FOCS90*, pages 231–241, St. Louis, October 1990. IEEE.

[4] A. Brandt. Multi-level adaptive solutions to boundary value problems. *Mathematics of Computation*, 31:333–390, 1977.

[5] T. F. Chan and B. Smith. Domain decomposition and multigrid algorithms for elliptic problems on unstructured meshes. *Contemporary Mathematics*, pages 1–14, 1993.

[6] Tony F. Chan and Jun Zou. Additive Schwarz domain decomposition methods for elliptic problems on unstructured meshes. CAM Report 95-16, Department of Math, UCLA, March 1995.

[7] S. E. Dorward, L. R. Matheson, and R. E. Tarjan. Unstructured multigrid strategies on massively parallel computers: a case for integrated design. In *Proc. 27th Annu. Hawaii Int. Conf. on System scienes*, pages 169–178, 1994.

[8] Omar Ghattas and Xiaogang Li. Private communications, Department of Civil Engineering, CMU, 1996.

[9] H. Guillard. Node nested multigrid with delaunay coarsening. Tech. report, INRIA, Valbonne, France, 1993.

[10] J. E. Pasciak J. H. Bramble and J. Xu. Convergence estimates for product iterative methods with applications to domain decomposition. *Math. Comput.*, 55:1–22, 1991.

[11] D. J. Mavriplis and D. Venkatakrishnan. Agglomeration multigrid for viscous turbulent flows. *AIAA paper*, 1994.

[12] G. L. Miller, D. Talmor, S.-H. Teng, and N. Walkington. A Delaunay based numerical method for three dimensions: generation, formulation, and partition. In *Proc. 27th Annu. ACM Sympos. Theory Comput.*, pages 683–692, 1995.

[13] Gary L. Miller, Dafna Talmor, Shang-Hua Teng, Noel Walkington, and Han Wang. Control volume meshes using sphere packing: Generation, refinement and coarsening. In *5th international meshing roundtable*, October 1996.

[14] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in three dimensions. In *Proceedings of the ACM Computational Geometry Conference*, pages 212–221, 1992. Also appeared as Cornell C.S. TR 92-1267.

[15] Carl F. Ollivier-Gooch. Upwind acceleration of an upwind euler solver on unstructured meshes. *AIAA Journal*, 33(10):1822–1827, 1995.

[16] Jim Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 83–92, 1992.

[17] G. Strang and G. J. Fix. *An Analysis of the Finite Element Method*. Prentice-Hall, 1973.

[18] Shang-Hua Teng. A geometric approach to parallel hierarchical and adaptive computing on unstructured meshes. In *Fifth SIAM Conference on Applied Linear Algebra*, pages 51–57, June 1994.

[19] D. Venkatakrishnan and D. J. Mavriplis. Agglomeration multigrid for the 3d euler equation. *AIAA paper*, 1994.