

Optimal Coarsening of Unstructured Meshes

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

Abstract

A bounded aspect-ratio coarsening sequence of an unstructured mesh \mathcal{M}_0 is a sequence of meshes $\mathcal{M}_1, \dots, \mathcal{M}_k$ such that:

- \mathcal{M}_i is a bounded aspect-ratio mesh, and
- \mathcal{M}_i is an approximation of \mathcal{M}_{i-1} that has fewer elements,

where a mesh is called a bounded aspect-ratio mesh if all its elements are of bounded aspect-ratio. The sequence is node-nested if the set of the nodes of \mathcal{M}_i is a subset of that of \mathcal{M}_{i-1} . The problem of constructing good quality coarsening sequences is a key step for hierarchical and multi-level numerical calculations. In this paper, we give an algorithm for finding a bounded aspect-ratio, node-nested, coarsening sequence that is of optimal size: that is, the number of meshes in the sequence, as well as the number of elements in each mesh, are within a constant factor of the smallest possible.

1 Introduction

Numerical methods such as the finite element, finite difference, and finite volume methods apply the following basic steps to solve a partial differential equation (PDE) over a continuous domain Ω :

1. **Geometric modeling:** the continuous domain is approximated using a simpler, discrete description.
2. **Mesh generation:** the interior of the domain is decomposed into a mesh M of simple and well-shaped elements.
3. **Approximation:** a system of linear or non-linear equations is formed over M for the governing PDEs (e.g., the stiffness matrix and the right hand vector are assembled).
4. **Solution:** the system of equations is solved, and the error of the solution is estimated.
5. **Adaptive refinement:** if necessary, the mesh is refined and steps 4 and 5 are repeated over the refined mesh.

*School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, {glmiller,tdafna}@cs.cmu.edu. Supported in part by NSF Grants CCR-9505472 and CCR-9706572.

†Department of Computer Science, University of Illinois at Urbana-Champaign, 1304 W. Springfield, Urbana, IL 61801-2987, steng@cs.uiuc.edu. Supported in part by the DOE Accelerated Strategic Computing Initiative (ASCI) Center for Simulation of Advanced Rockets, an NSF CAREER award (CCR-9502540), an Alfred P. Sloan Research Fellowship.. Part of the work will Department of Computer Science, University of Minnesota, Minneapolis, MN 55455.

Once the mesh M is generated, a system of equations is formed over M and has to be solved. The class of multi-level techniques has become one of the most effective and successful classes of numerical techniques for solving PDEs. These techniques have been used in multigrid methods [6] and multi-level domain decomposition [7]. A multi-level method solves a system of equations by first constructing a *coarsening sequence* of meshes M_0, \dots, 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} . At the heart of multigrid methods and other multi-level 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 methods solve a PDE on Ω by first obtaining an initial solution either for M_0 or for M_k , and then improving the quality of the solution by transforming it up and down the coarsening sequence hierarchy while applying some simple and efficient iterative methods at each level.

The simplest form of a coarsening sequence is a series of nested structured meshes (regular grids). For this class of meshes, Brandt [6] showed, by carefully using restriction and interpolation, that the solution for M_0 can be efficiently obtained using multigrid methods. (See also Bramble, Pasciak and Xu [5].) Nested structured coarsening sequences are attractive choices in practice because they can be easily generated and because the convergence behavior of the structured multigrid methods on such sequences is well understood. However, the use of structured regular grids limits the applicability of the method to problems whose domains are simple and whose solution functions have a small or constant Hessian [7, 16, 21].

The use of unstructured meshes is inevitable in the solution of complex problems with more intricate domain geometry and solutions. The theory behind applying multigrid and multi-level methods to unstructured meshes is not as well understood as it is for the structured case. However, it is an area of active research, with increasing theoretical progress [5, 7, 8]. Also, in practice, multigrid methods on unstructured meshes are popular despite this current lack of full theoretical foundations [3, 18].

The algorithmic and computational geometry aspects of generating coarsening sequences are important issues to address in view of the theoretical and experimental interest in applying multigrid and multi-level methods to unstructured meshes. The effectiveness of a multi-level method that uses an unstructured coarsening sequence M_0, \dots, M_k depends on the quality of this sequence [5, 7, 8]. In particular, Chan and Zou [8] provided sufficient conditions for multi-level 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 sequence M_1, \dots, M_k with $|M_k| = O(b)$ that satisfies conditions (1) and (2).

The problem of mesh coarsening can be viewed as a generalization of the problem of mesh generation. In particular, a sequence of meshes has to be constructed rather than a single mesh, and furthermore, the coarsener is restricted to sample only from the discrete set of the initial mesh nodes instead of having the freedom to sample arbitrarily from the domain. Thus, the challenges of mesh coarsening, as contrasted with mesh generation, are twofold: first, that of understanding the shape and size of each mesh in the sequence that satisfies conditions (1) and (2) specified above, and second, the challenge of constructing such a sequence over a discrete node set. However, without the foundations laid by recent advances in the theoretical understanding of mesh generation techniques, as in the work of Bern, Eppstein and Gilbert [4], Chew [9], Ruppert [19], Mitchell and Vavasis [17], and Miller, Talmor, Teng, Walkington [16], this work would not be possible.

In this paper we describe a new algorithm for the coarsening problem in two dimensions. This algorithm guarantees that the coarsening sequence satisfies conditions (1) and (2). It also minimizes the size of the mesh at each level up to a constant factor. A sequence is *node-nested* if the set of the nodes of M_i is a subset of that of M_{i-1} . The approach presented here can be used to generate both node-nested and non-nested coarsening sequences.

The paper is organized as follows: Section 2 gives basic definitions and more formal discussion of the problem of mesh coarsening. Section 3 reviews previous work on mesh coarsening. Section 4 presents the coarsening algorithms. Section 5 explains spacing functions and their properties — spacing functions are the main technique behind our coarsening algorithm. Sections 6 and 7 show the correctness of our approach for one-level and multi-level schemes respectively. Section 8 discusses the run-time complexity of our scheme, and efficient variations of the approach. Section 9 gives some experimental evidence that the approach generates high-quality coarsening sequences in practice. Section 10 concludes the paper.

2 The problem of mesh coarsening

This section provides the background for the rest of the paper. We first introduce the basic definitions and the notation used in this paper. We then formally define the problem of mesh coarsening.

2.1 Basic mesh qualities

A **Planar Straight Line Graph (PSLG)** is a collection of line segments and points in the plane, closed under intersection. A **domain** Ω is a planar straight-line graph (PSLG), whose boundary is a collection of polygons (i.e., holes are allowed). The PSLG is a linear modeling of some continuous domain in the plane, possibly containing internal edges and points representing boundaries between two materials, and points of special interest.

A two-dimensional **mesh** is a discretization of the geometric domain, described by a PSLG, into simple polygonal **elements**. The intersection of two neighboring elements must either be an edge of both elements, or a vertex of both. The **mesh nodes** are the vertices of its elements. The following notation is used to describe a mesh: $\mathcal{M} = (\mathcal{P}, \mathcal{T}, \mathcal{B})$, where \mathcal{P} is the set of mesh nodes, E is the edge set, and \mathcal{B} is the PSLG boundary description. Nodes from \mathcal{P} that reside on boundary segments or points are called **boundary nodes**. To distinguish between boundary nodes and points of \mathcal{B} , we refer to the latter as boundary vertices. Edges from E that reside on **boundary segments** of \mathcal{B} are called **boundary edges**.

The mesh can be **structured** or **unstructured**. For a structured mesh, the local geometrical and combinatorial structure of the mesh at each of its nodes is identical, with the possible exception of the mesh’s boundary nodes. Unstructured meshes allow for full flexibility, and the local geometrical and combinatorial structure can differ from one node to another. Structured meshes often use quadrilateral elements, whereas unstructured meshes typically use triangular elements.

The following mesh categories play an important part in this paper: **grids** are uniform structured meshes composed of square elements; **quasi-uniform unstructured meshes** are unstructured triangular meshes, whose elements have edge-lengths that are within a constant factor of each other; **general unstructured meshes** are unstructured triangular meshes, with no special restrictions.

A mesh is generated as an intermediate step of a numerical method to compute or simulate physical quantities over the original domain. Not all meshes perform equally well in the subsequent

numerical computations. Numerical and discretization errors depend on the geometric shape and size of the mesh elements. The following measures are defined to quantify these mesh qualities.

Definition 2.1 (edge-length function (el_M)) *Let M be a mesh over a two-dimensional domain Ω . For each $x \in \Omega$, $\text{el}_M(x)$ is defined to be the length of the longest edge of all the mesh elements that contain x .*

Note that only the mesh nodes and the mesh edges are contained in more than one element. This definition has been used in several mesh generation papers, such as [4, 17, 19].

The geometric quality of a triangular element is measured using its aspect-ratio. In general, long and skinny triangles are undesirable; for example see [1].

Definition 2.2 (aspect-ratio) *The **aspect-ratio** of a mesh element is R/r , where R is the radius of the smallest ball containing the element and r is the radius of the largest ball contained in the element. The aspect-ratio of a mesh is the maximal aspect-ratio among its elements.*

If a mesh has a bounded aspect-ratio, then its smallest angle is bounded from below by a constant and vice versa. Therefore, by “a mesh has a bounded aspect-ratio”, we mean that its aspect-ratio is smaller than a given constant C , or, equivalently, its smallest angle is larger than a given constant angle θ .

The el function is a measure of the mesh element size at a given point in the domain; the aspect-ratio is a measure of the element shape. These two qualities are often at odds: a better aspect-ratio mesh conforming to a given boundary usually has more elements.

2.2 Mesh coarsenings and coarsening sequences

Definition 2.3 (coarsening) *A **coarsening** $M' = (\mathcal{P}', E', \mathcal{B})$ of a mesh $M = (\mathcal{P}, E, \mathcal{B})$ is a mesh whose edge-length function $\text{el}_{M'}$ is point-wise larger than el_M but still conforms to the same domain. The mesh M is then a **refining** of M' .*

Mesh coarsening is the inverse problem of mesh refining. Adaptive refinement is often used to generate an unstructured mesh: a new mesh is obtained by refining a previously generated mesh at regions where more accurate information is desired, or where large errors are suspected. If all intermediate changes have been kept, a coarsening sequence can be obtained by undoing the changes made by reversing the refinement process. However, it is unreasonable to assume that such a mesh refinement history is always available, or if available, that it generates an optimal coarsening sequence. In addition, not all meshes are generated by a refinement method. In a black-box approach to mesh coarsening, the only information available is the input mesh M itself.

A coarsening can be **element-nested**, **node-nested**, or **non-nested**. In an element-nested coarsening, each element of M' can be represented as a union of elements in its refinement M , while elements in the other two types may be unrelated. A coarsening is **node-nested** if each node of mesh M' is also a node of mesh M ; otherwise it is **non-nested**.

In general, a triangular mesh does not have any element-nested coarsening, unless it was carefully crafted as such. Even mesh refinement methods that mostly use element decomposition in the refinement process are not purely decomposition based; a certain amount of edge flipping and deleting must be used to maintain good aspect-ratio.

Coarsening, even in the relaxed sense of node-nested meshes, might cause a degradation in the aspect-ratio of the coarser mesh compared to that of the finer one. This degradation should be

avoided. Our goal is to automatically generate a sequence of coarsened meshes from an initial high-quality mesh, with tight controls on the aspect-ratio, the number of elements in each coarsened mesh, and the number of meshes in the sequence.

Definition 2.4 (coarsening sequence) *A coarsening sequence of a mesh M_0 is a sequence of meshes $\mathcal{M} = M_0, M_1 \dots M_k$, such that M_{i+1} is a coarsening of mesh M_i .*

Definition 2.5 (length and width of a coarsening sequence) *The length of a coarsening sequence is the number of meshes (levels) in the sequence; the width of the sequence is a function of its level: at level i , the width is $|M_i|$, the number of nodes and elements of mesh M_i .*

Meshes in a coarsening sequence should approximate each other. The approximation condition can be expressed by their edge-length functions:

Definition 2.6 (local similarity) *Two meshes M_1 and M_2 are \mathcal{I} -locally similar if*

$$\frac{1}{\mathcal{I}} \text{el}_{M_2} \leq \text{el}_{M_1} \leq \mathcal{I} \text{el}_{M_2}$$

Definition 2.7 (well-shaped coarsening sequence) *A coarsening sequence $\mathcal{M} = \{M_0 \dots M_k\}$ is (θ, \mathcal{I}, b) -well-shaped if*

- \mathcal{M} is a bounded aspect-ratio coarsening sequence with smallest angle bound θ .
- Mesh M_i and M_{i+1} are \mathcal{I} -locally similar.
- $|M_k| = O(b)$

The first condition, bounded aspect-ratio, is motivated by iterative methods: at each level of the multigrid method an iterative method is used for a few rounds to “smooth” the error. The convergence properties of iterative methods are related to the aspect-ratio of the underlying mesh. The second condition, local similarity, is motivated by the restriction and interpolation phases of the multigrid methods, used to transform partial solutions between meshes in adjacent levels of the hierarchy. To reduce the interpolation and restriction errors, adjacent meshes should approximate each other, i.e., they should be locally similar.

3 Previous work

The problem of mesh coarsening has two aspects: numerical analysis and computational geometry. Most of the previous research is concentrated on the numerical analysis aspect of mesh coarsening. The numerical analysis investigation must naturally take precedence, to uncover the necessary and sufficient conditions that a desirable coarsening sequence should satisfy. However, as the multigrid method has been gaining some acceptance and maturity, software development requires better understanding of the underlying computational geometry issues. Therefore, one objective of this work is to initiate the study of mesh coarsening from its computational geometry aspects.

To the best of our knowledge, all current algorithms for mesh coarsening provide no guarantees on the size and shape of the elements that they produce. In fact, for most algorithms, we can provide counterexamples that show that these algorithms can fail to produce a good coarsening sequence.

Arguably the most popular approach is the topologically based coarsening paradigm. Other approaches can be found in Bank and Xu [2], and Barth [3] for example. In topological coarsening, the topological structure of the mesh, i.e., the neighborhood graph structure of either its elements or nodes, is used. MIS-based coarsening is the prominent topological based method:

MIS-based methods: A maximal independent subset of the node neighborhood graph is selected to be the coarser mesh nodes. This node set is then triangulated, and the process is repeated iteratively to generate the coarsening sequence. The coarse nodes may be triangulated by any triangulation method. For example, Guillard [12] uses the Delaunay triangulation, whereas Ollivier-Gooch [18] iteratively removes the nodes not selected, and locally repairs the mesh structure.

Definition 3.1 (maximal independent set (MIS)) *The MIS of a graph $G = (V, E)$, where E is the edge set and V is the node set, is a set $U \subseteq V$ such that the following are true:*

1. **independence:** if $u_1, u_2 \in U$ then $(u_1, u_2) \notin E$, and
2. **maximality:** no node $v \in V$ can be added to U without violating independence, i.e., all nodes $v \in V \setminus U$ are incident to some edge (v, u) such that $u \in U$.

Agglomeration methods: This topological method is applied to control-volume discretizations, a description of which can be found in Mavriplis' survey [14]. In a control-volume discretization, each mesh node is assigned a volume; these may be the Voronoi diagrams when the mesh is a Delaunay triangulation, or the volumes formed by connecting all the midpoints of edges and centers of triangles incident to a node.

Agglomeration methods coarsen this partition by joining neighboring volumes together. The complexity of the cell or volume description increases with the levels of the coarsening process. The coarsening proceeds by iteratively considering all volumes. If a volume has been coarsened already, then it is skipped. Otherwise, the volume grabs some or all of its un-coarsened neighboring volumes, to form a new, larger cell. As in the MIS-based schemes, the order in which the volumes are traversed can be prioritized to proceed from the boundary cells inward, or in an advanced front fashion. See [13, 14, 15, 22].

3.1 Problems of topological coarsening: repeated degradation

An MIS-based method does not guarantee the geometric quality of the resulting coarsening mesh sequence. This 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 highly uniform meshes, as in the grid-like mesh of Figure 1. Agglomeration methods are in a sense dual to MIS-based methods, and can be shown to be subject to the same aspect-ratio degradation problem. For general unstructured meshes, the problem of repeated degradation is even more severe as stated in the following Lemma.

Lemma 3.2 *There exists an unstructured mesh such that any selection of a maximal independent set results in a coarsening sequence with increasingly worse aspect-ratio.*

The mesh in Figure 6 is an example of one such mesh. It is characterized by rapidly increasing element sizes. Taking this example to the extreme, one can construct a bounded aspect-ratio mesh with $O(n)$ elements, and edge-lengths varying from 2 to 2^n . Such a mesh does not have a well-shaped coarsening sequence with only $O(\log n)$ levels. The results of this paper show that unless one compromises either on the local-similarity property or on the bounded aspect-ratio requirement, any well-shaped coarsening sequence of this extreme example must contain $O(n)$ levels. Since topological coarsening methods always reduce the size of the mesh by a constant factor at each level, they must fail to produce a well-shaped coarsening sequence of this mesh.




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


The two examples, the grid of Figure 1 and the mesh of Figure 6 demonstrate two different points: for the grid, a simple $O(\log n)$ length coarsening sequence exists. Some choices of an MIS may yield an optimal well-shaped coarsening sequence of the grid, and in practice they indeed perform favorably. However, the example showed that not all choices of an MIS lead to good grid coarsenings. Furthermore, one can show that even when selecting nodes for the MIS randomly, the aspect-ratio of the grid coarsening sequence degrades with probability approaching one as the size of the initial grid is increased [20]. The second example involved a mesh for which an $O(\log n)$ well-shaped coarsening sequence simply does not exist. Topological approaches, including the MIS, consider a mesh solely based on its neighborhood structure. However, highly graded meshes can not be as easily coarsened as their quasi-uniform counterparts even when they have the same topological structure. To obtain elements with high geometric quality, one may have to coarsen less aggressively in areas of high geometric gradation. Notice that there is a necessary tradeoff between the size of the coarsening sequence and its geometric quality. Our method uniformly handles both examples, providing size and shape guarantees while generating a coarsening sequence of optimal size.

4 New approach: function based coarsening


In this section, a new approach for mesh coarsening is introduced: a representation we term “spacing functions” is used to capture the geometric structure of the mesh, and to generate an optimal coarsening sequence while avoiding the pitfall of repeated degradation.

Let M_0 be the initial mesh. Its spacing function representation is a function f_0 which describes the typical size and node spacing of M_0 : f_0 is defined by its values at the mesh nodes. The idea is to compute a spacing function f_i for each level, and use it to generate M_i from the original (or previous) node set. Given the spacing functions, the task is then to create a node set that is “spaced” according to that function, and triangulate it. We call this coarsening technique **function-based coarsening**. It contains four steps:

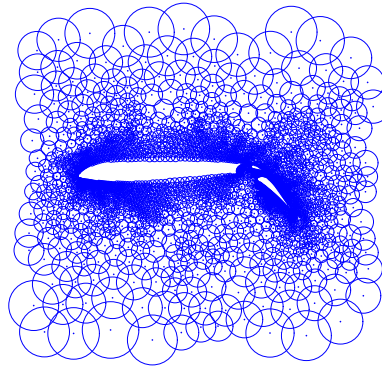
1. recover the spacing function of the initial mesh;
2. increase the spacing value of the mesh nodes smoothly to obtain the new spacing functions;
3. delete some mesh nodes so that the remaining nodes are spaced according to the new spacing function; and
4. generate the new mesh by triangulating the nodes obtained in Step 3.




(a)




(b)



(c)



(d)



(e)

Figure 2: Function-based coarsening: the spacing function value at a node is denoted by a sphere whose radius is proportional to the function value. The initial mesh, generated by Barth and Jespersen, is plotted in (a). The first step is recovering the spacing function of the mesh using the NN function (b). The function is then coarsened (c), and a subset of the nodes forming a maximal set of disjoint spheres is picked (d) and triangulated (e).

In the function-based coarsening approach presented in this paper, we use a simplifying assumption that all the meshes of the coarsening sequence conform to the same boundary description. Thus, the problem of simplifying the domain description itself is not handled theoretically in this paper. We allow boundary nodes lying on the segments to be removed, but the vertices of the boundary segments are always retained. For example, if the domain is a square, the square vertices are always retained but boundary nodes placed on the sides of the square may be pruned away.

It is sometimes desirable to combine mesh coarsening with boundary description simplification. In boundary simplification, some of the vertices of the boundary segments may be removed and a new boundary description is formed, under the restriction that the new boundary approximates the old boundary. This might enable continuing the coarsening beyond the original boundary restrictions. However, it is easy to find domain counter-examples where no reasonable boundary simplification that approximates the original boundary allows for good aspect-ratio meshing. Examples and a more detailed discussion of this problem can be found in [20].

Nonetheless, for many interesting and practical domains the ideas of function-based coarsening can be easily adapted to handle simplification as well. Figure 2 shows function-based coarsening employed on a domain with a complex boundary, where the boundary is allowed to be coarsened.

The following sections describe the coarsening algorithm in more detail. Section 4.1 shows how to recover the spacing function from the original mesh. Section 4.2 describes how to coarsen the spacing function. Finally, Section 4.3 shows how to generate the coarse mesh using the coarse spacing function. This basic coarsening scheme is described in algorithm ONE_LEVEL_COARSEN of Figure 3.

4.1 Recovering the spacing function

We first formalize the notion of spacing functions. Any slowly changing function may serve as a spacing function.

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

To space nodes according to a spacing function, we draw spheres around the nodes and restrict the spheres to be disjoint. The radius of a sphere at a node is proportional to the spacing function. We found it more convenient to deal with functions normalized to be 1-Lipschitz, at the cost of introducing a parameter that controls how to space the nodes according to the function.

Definition 4.2 (f^β -spaced points) *Let $\beta > 1$ be a real number, f a 1-Lipschitz function. A point set \mathcal{P} is f^β -spaced if for any two points $p, q \in \mathcal{P}$, $\frac{f(p)}{\beta} + \frac{f(q)}{\beta} < \|p - q\|$, i.e. the spheres of radius $f(p)/\beta$ and $f(q)/\beta$ centered at p and q respectively are required to be disjoint. In that case, we say f β -spaces the point set, and we refer to the spheres as the f^β -spheres. When β is clear from the context, we omit it from the notation.*

The initial spacing function is recovered from the natural spacing of the original mesh as follows:

Definition 4.3 *The nearest neighbor (NN) function of a node set $\mathcal{P} \subset \Omega$ assigns to each node $p \in \mathcal{P}$ the distance to the node $q \in \mathcal{P}$ that is nearest to it. It can be extended to the domain Ω by assigning to a point $x \in \Omega$ the radius of the smallest sphere centered at x that contains at least two nodes of \mathcal{P} .*

4.2 Coarsening the spacing function

Each node has a desirable coarsening goal which, because of the Lipschitz condition, can also be used to derive restrictions on the coarsened function values at other nodes. The final coarsening function is the largest Lipschitz function that obeys all these restrictions. The following definition describes this coarsening function.

Definition 4.4 (coarsening function) *Let \mathcal{P} be a node set in a domain $\Omega \subseteq \mathbb{R}^2$. Let g be a spacing function over Ω . Let $\mathcal{C} > 1$ be an arbitrary real number which is the desired coarsening factor. The **local \mathcal{C} -coarsening** of g at a node $q \in \mathcal{P}$ is a cone-shaped function centered at q :*

$$\hat{f}_{g,\mathcal{C},q}(x) = \mathcal{C} \cdot g(q) + \|q - x\|$$

*The **\mathcal{C} -coarsening** of f with respect to \mathcal{P} is defined as a minimum of all the local coarsenings:*

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

When clear from the context, g , \mathcal{P} and \mathcal{C} are omitted from the notation. For complex domains, the distance metric used is the geodesic distance metric.

Sequence of coarsening functions: to generate the sequence of coarsening meshes from the initial mesh M_0 , a corresponding sequence of spacing functions is defined:

$$f_i = f_{\text{NN},2^i,\mathcal{P}}$$

4.3 Coarsening the meshes

Let M_0 be the initial mesh and \mathcal{C} be the coarsening factor. Refer to Figure 3 for the basic algorithm. To coarsen the mesh, we first generate the coarsened spacing function values for each node of the mesh using \mathcal{C} , and then select a subset of the mesh nodes which is f^β -spaced according to this spacing function. We select the coarsened mesh nodes using a conflict graph.

Definition 4.5 *The **conflict graph** $CG(\mathcal{P})$ of a point set \mathcal{P} with respect to a β -spacing function f and a boundary description \mathcal{B} is a graph $CG(\mathcal{P}) = (\mathcal{P}, E)$ with the following edge set:*

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

such that (p_i, p_j) is contained in the domain (i.e., node p_i visible to node p_j)

The coarsened point-set \mathcal{Q} is essentially an MIS of the conflict graph with respect to the coarsened spacing function. This basic scheme is enhanced in order to handle the boundary. The coarse point set, \mathcal{Q} , is built in increasing order of dimensions: initially the boundary segment vertices and points are included in \mathcal{Q} , then an MIS of the conflict-graph of the boundary nodes is added to \mathcal{Q} , and new boundary edges are formed by edge-contraction of removed boundary nodes. Finally, \mathcal{Q} is completed to include the interior nodes that are not too close to the boundary edges, and that form an MIS of the corresponding conflict-graph. To prevent interior nodes that are too close to a boundary edge from joining \mathcal{Q} , a protective zone around each boundary edge is used:

Definition 4.6 (γ protective zone) *Let e be an edge. The γ protective zone around e is a rectangle of height $2\gamma|e|$, extending to a height of $\gamma|e|$ on each side of e and of width $|e|$. If this rectangle intersects any boundary edge, a truncated protective zone is used. (However, γ can be set to guarantee intersection occurs only with incident boundary edges).*

Procedure: ONE_LEVEL_COARSEN($M_0, f_0, \mathcal{C}, \beta, \gamma$)

Input: $M_0 = (\mathcal{P}, E, \mathcal{B})$, a mesh over a domain described by \mathcal{B} .

f_0 , the spacing function of M_0 .

\mathcal{C} , the coarsening factor.

β , the spacing constant.

γ , the boundary edges protective zone constant. (See Section 4.3 and definition 4.6.)

Output: $M_1 = (\mathcal{Q}, E_1, \mathcal{B})$, a coarser triangular mesh conforming to \mathcal{B} .

Method:

1. Compute $f_1(p_i) = f_{f_0, \mathcal{C}, \mathcal{P}}(p_i)$ for each $p_i \in \mathcal{P}$.
2. Let \mathcal{P}^0 be the boundary vertices; \mathcal{P}^1 M_0 's nodes located on the segments of \mathcal{B} ; \mathcal{P}^2 the mesh nodes. Note that $\mathcal{P}^0 \subseteq \mathcal{P}^1 \subseteq \mathcal{P}^2 = \mathcal{P}$.
3. Set $\mathcal{Q}^0 = \mathcal{P}^0$.
4. Let $G_1 = \text{CG}(\mathcal{P}^1)$ be a conflict-graph with respect to f_1 over \mathcal{P}^1 . (see Definition 4.5)
Set \mathcal{Q}^1 to be a completion of \mathcal{Q}^0 to an MIS of G_1 .
5. Form \hat{E} , the set of boundary edges of M_1 . \hat{E} is formed by edge contraction of the nodes $\mathcal{P}^1 \setminus \mathcal{Q}^1$.
Let $\hat{\mathcal{Q}}^2$ be the set of nodes from \mathcal{P}^2 that are outside the γ -protective zones of \hat{E} .
6. Let $G_2 = \text{CG}(\mathcal{Q}^1 \cup \hat{\mathcal{Q}}^2)$ be a conflict-graph with respect to f_1 .
Set \mathcal{Q}^2 to be a completion of \mathcal{Q}^1 to an MIS of G_2 .
7. Set $E_1 = \text{CDT}(\mathcal{Q}^2, \hat{E})$, the edges of a constrained Delaunay triangulation of \mathcal{Q}^2 with constraints \hat{E} .
Return $M_1 = (\mathcal{Q}^2, E_1, \mathcal{B})$.

Figure 3: One level function based coarsening.

Procedure: MULTILEVEL_COARSEN(M_0, f_0, β, γ)

1. Let k be the length of the required coarsening sequence:
2. Let ONE_LEVEL_MODIFIED($M_0, f_0, \mathcal{C}, \beta, \gamma, \mathcal{P}$) be a modification of ONE_LEVEL_COARSEN that forces \mathcal{P} into the output coarsened mesh. (see Lemma 4.7).
3. Let $\mathcal{P}_{k+1} = \emptyset$.
4. For $i = k$ to 1
 - $M_i = \text{ONE_LEVEL_MODIFIED}(M_0, f_0, 2^i, \beta, \gamma, \mathcal{P}_{i+1})$
 - Set \mathcal{P}_i to the node set of mesh M_i .
5. Return (M_1, \dots, M_k) .

Figure 4: Multi-level function-based coarsening: coarse-to-fine approach.

Multi-level schemes: The basic one-level scheme is outlined in Figure 3. We have two approaches to apply this one-level scheme to build a multi-level scheme:

1. Constructing the coarsening meshes in a fine-to-coarse order: mesh M_1 is generated first, followed by M_2 , etc. To generate M_{i+1} , algorithm `ONE_LEVEL_COARSEN` is called with M_i as the input mesh and with coarsening factor $\mathcal{C} = 2$.
2. Using a geometrically increasing coarsening constant $\mathcal{C}_i = 2^i$ to generate the coarsening mesh M_i directly from the initial mesh M_0 with coarsening factor \mathcal{C}_i .

The first scheme is simple and generates a node-nested coarsening sequence. Our experiments (see Section 9 and also in [20]) show that in practice this scheme successfully avoids repeated degradation on both quasi-uniform and highly graded meshes. However, only for the quasi-uniform case were we able to show that the scheme produces optimal coarsening sequences [20]. It is interesting to compare this approach to MIS approaches. Our experiments show that MIS approaches fail in practice on highly graded meshes. Moreover, as we have shown previously, MIS approaches fail even quasi-uniform meshes.

In the rest of this paper, we show that our second scheme always generates an optimal well-shaped coarsening sequence. A variant of this scheme can be used to generate a coarsening sequence which is node-nested as well. The variant generates first the coarsest mesh and proceeds to generate finer meshes level by level. At each level it ascertains that coarser mesh nodes are included in the current mesh. Lemma 4.7 shows that the node set of M_k can then be forced to be included in M_{k-1} . Figure 4 contains a pseudo-code description of the scheme.

Lemma 4.7 *Let \mathcal{P}_i be the node set of the coarse mesh M_i generated by the algorithm of Figure 3 with input $M_0, f_0, \mathcal{C}_i, \beta, \gamma$. Let $\mathcal{C}_j < \mathcal{C}_i$ be a smaller coarsening factor. The algorithm can be modified to complete \mathcal{P}_i into a coarse mesh M_j , using input parameters $M_0, f_0, \mathcal{C}_j, \beta, \gamma$.*

Proof: $\mathcal{C}_j < \mathcal{C}_i$ implies that $f_j < f_i$. Therefore, any conflict-graph generated for f_i is a super-graph of the corresponding graph for f_j . By simply adjusting the priorities of the node set \mathcal{P}_0 such that the nodes of \mathcal{P}_i are always considered first, it is straightforward to complete \mathcal{P}_i into an MIS of the less constrained conflict-graph.

Similarly, since the boundary edges of M_j are smaller too, the protective zones are smaller and fewer nodes are deleted by them. □

5 The equivalence between spacing functions and meshes

Spacing functions are at the heart of our coarsening approach. This section discusses spacing functions and their connection to bounded aspect-ratio meshes. The main result of this section, Theorem 5.3, shows that the concept of points spaced according to a spacing function with appropriate restrictions on the formation of gaps among points, is equivalent to the concept of bounded aspect-ratio meshes.

The following definitions enable the handling and controlling of gaps:

Definition 5.1 (gap) *Let \mathcal{L} be a positive constant. A point set \mathcal{P} (or its mesh M) is said to have an \mathcal{L} -gap at x with respect to a spacing function f if there exists a sphere S with radius $\mathcal{L}f(x)$ that contains x on its circumference but does not contain any node of \mathcal{P} visible to x in its interior. Let c be the center of S . \mathcal{L} is called the **gap parameter**.*

- S is an **interior \mathcal{L} -gap** at x if the radius from x to c is inside the domain.
- a one-dimensional sphere S is a **boundary \mathcal{L} -gap** at x if there exists a boundary segment B of the mesh such that $x, c \in B$.

Definition 5.2 (bounded gaps property) Let \mathcal{L} be a positive constant. A mesh $M = (\mathcal{P}, E, \mathcal{B})$ is said to have the bounded gaps property with gap parameter \mathcal{L} if all its interior and boundary gaps have a gap parameter smaller than \mathcal{L} .

Points spaced according to a spacing function, together with the concept of bounded gaps, are equivalent to bounded aspect-ratio meshes in the following sense:

Theorem 5.3 (equivalence) The following two statements show that bounded aspect-ratio meshes and points spaced by a spacing function with bounded gaps are equivalent:

1. **From a mesh to a spacing function:** For every bounded aspect-ratio mesh $M = (\mathcal{P}, E, \mathcal{B})$ there exists a spacing function f 2-spacing the mesh nodes, such that for some positive constants \mathcal{L} and b :
 - (a) M has the bounded \mathcal{L} -gaps property with respect to f .
 - (b) Let p be a node of \mathcal{P} , let B be a boundary segment such that p and B are disjoint. The geodesic distance between p and B is greater than $bf(p)$.
2. **From a spacing function to a mesh:** If a point set \mathcal{P} of a domain Ω is β -spaced with respect to a spacing function f , and has properties (a) and (b) above for some positive constants b and \mathcal{L} , then a constrained Delaunay triangulation of \mathcal{P} is a bounded aspect-ratio mesh.

The rest of this section is organized as follows: Section 5.1 discusses previous usages of mesh functions. Sections 5.2 and Section 5.3 prove the two parts of Theorem 5.3. Finally, Section 5.4 proves some properties of spacing functions that are used in later sections.

5.1 Previous usages of mesh functions

There are many previous instances of associating local size, or local scale functions with meshes. The two examples most directly related to this paper are the work of Mitchell and Vavasis [17] and the work of Ruppert [19]. In both works edge length functions describing the typical local size are used to bound the number of elements in the meshes that their algorithms generate. Ruppert also introduced a function called the **local feature size**. He showed that all bounded aspect-ratio meshes of the same domain have edge length functions that are at most a constant factor larger than the local feature size function, and thus provided an absolute way to measure (up to a constant factor) the size of the smallest bounded aspect-ratio mesh conforming to the domain. The following definitions and results by Ruppert are used here:

Definition 5.4 (Ruppert: Local feature size (lfs)) Given a PSLG \mathcal{B} , the local feature size at point x , $\text{lfs}_{\mathcal{B}}(x)$, is the radius of the smallest disk centered at x that intersects two unrelated features of \mathcal{B} . Two features (segments or vertices) are unrelated if they are disjoint (non-incident).

Lemma 5.5 (Ruppert) The lfs function is 1-Lipschitz.

Under certain restrictions on the boundary \mathcal{B} (e.g., the smallest domain angle between segments of \mathcal{B} has to be larger than 90°), Ruppert showed the following theorem:

Theorem 5.6 (Ruppert) *Given a PSLG \mathcal{B} , and $\theta \leq 20^\circ$, suppose T is a triangulation of \mathcal{B} with minimum angle bound θ . There exists a constant C such that: $\text{el}(x) \leq C \text{lfs}(x)$, $\forall x \in \Omega$.*

Note that lfs depends only on the boundary \mathcal{B} and not on the actual discretization of the domain, whereas the nearest neighbor function $\text{NN}_{\mathcal{P}}$ depends on the discretization itself. Therefore, Theorem 5.6 can be strengthened for $\text{NN}_{\mathcal{P}}$ (under the same assumptions):

Theorem 5.7 *Let $M = (\mathcal{P}, E, \mathcal{B})$ be a mesh with smallest angle θ . There exist two positive constants C_1, C_2 depending on θ only such that $C_1 \text{el}_M(x) \leq \text{NN}_{\mathcal{P}}(x) \leq C_2 \text{el}_M(x)$.*

Proof: Create a new boundary description $\mathcal{B}_1 \equiv \mathcal{P}$. The mesh M conforms to the boundary \mathcal{B}_1 as well. Note that $\text{lfs}_{\mathcal{B}_1} \equiv \text{NN}_{\mathcal{P}}$. Theorem 5.6 implies that $C_1 \text{el}_M(x) \leq \text{lfs}_{\mathcal{B}_1}(x) = \text{NN}_{\mathcal{P}}(x)$. For the other direction, first consider the inequality for some $p \in \mathcal{P}$: $\text{NN}_{\mathcal{P}}(p) \leq \text{el}(p)$ since the longest edge at p must be larger than the distance to its nearest neighbor. Now consider $x \in \Omega \setminus \mathcal{P}$, and suppose x is in a triangle T . Let p be one of the vertices of T . Since the edges of neighboring triangles in a bounded aspect-ratio mesh can differ by at most a constant factor, there exists a constant C , depending on θ only, such that $\text{el}(p) \leq C \text{el}(x)$. The distance between x and p is at most $\text{el}(x)$, therefore: $\text{NN}_{\mathcal{P}}(x) \leq \text{NN}_{\mathcal{P}}(p) + \text{el}(x) \leq \text{el}(p) + \text{el}(x) \leq (C + 1)\text{el}(x)$. \square

5.2 From a mesh to a spacing function

Let $M = (\mathcal{P}, E, \mathcal{B})$ be a bounded aspect-ratio mesh. This section shows that the function $\text{NN}_{\mathcal{P}}$, of Definition 4.3, is a spacing function exhibiting all the desirable properties stated in Theorem 5.3.

Lemma 5.8 *\mathcal{P} is 2-spaced according to $\text{NN}_{\mathcal{P}}$.*

Proof: Let $p, q \in \mathcal{P}$. $\|p - q\| \geq 0.5(\text{NN}(p) + \text{NN}(q))$. \square

We now proceed to show that interior and boundary gaps of $\text{NN}_{\mathcal{P}}$ are bounded. By definition, if S is an interior gap, the radius at some point x on the boundary of S is inside the domain. The triangles intersecting that radius form a connected component. The following lemma bounds the number of triangles in such a connected component.

Lemma 5.9 *Let $M = (\mathcal{P}, E, \mathcal{B})$ be a mesh with smallest angle θ . Let S be a sphere containing no nodes from \mathcal{P} . Let \mathcal{T} be a connected component of triangles of M intersecting S . Then, \mathcal{T} contains at most $2 + \frac{12}{\sin^4 \theta}$ triangles of M .*

Proof sketch: We give a high-level proof. The derivation of the constants involved can be found in [20]. A triangle can intersect S with one, two, or three edges. In a connected component, a triangle with one edge intersecting S is either incident to a triangle with two or three edges intersecting S , or only two triangles intersect S . Triangles that intersect S with two or three edges, and have smallest angle larger than θ , contain at least a constant fraction of the area of S . Hence by a volume argument, at most a constant number of them exist. \square

The following theorem bounds the size of gaps:

Theorem 5.10 *Let $M = (\mathcal{P}, E, \mathcal{B})$ be a mesh with smallest angle θ . There exists a constant \mathcal{L} , depending on θ only, such that all interior gaps have a gap parameter smaller than \mathcal{L} with respect to the spacing function $\text{NN}_{\mathcal{P}}$.*

Proof: Let $x \in \Omega$ be an arbitrary point and let S be a sphere through x with radius r and center c , such that the radius from x to c is contained in Ω . Consider \mathcal{T} , the connected component of triangles that intersect the radius. If S contains no node of \mathcal{T} , then, by Lemma 5.9, \mathcal{T} contains at most $C(\theta)$ triangles of M , where C is a constant depending only on θ . The center c is in some triangle (p, q, t) . Let $|(p, q)| = \ell$ be the longest edge of the triangle. Clearly, $r \leq |c - p| \leq \ell$, and therefore $r \leq \text{el}(c)$. The number of triangles separating x from c is at most $C(\theta)$. The rate at which el can increase from one triangle to its neighbor is bounded by another constant $D(\theta)$. Therefore, $\text{el}(x) \geq D(\theta)^{C(\theta)} \text{el}(c) \geq D(\theta)^{C(\theta)} r$. By Theorem 5.7, $\text{NN}_{\mathcal{P}}(x) \geq C_1 \text{el}(x) \geq C_1 D(\theta)^{C(\theta)} r$, and the radius of a gap at x is thus bounded in terms of $\text{NN}_{\mathcal{P}}$. \square

The following theorem bounds boundary gaps:

Theorem 5.11 *Let $M = (\mathcal{P}, E, \mathcal{B})$ be a mesh with smallest angle θ . There exists a constant \mathcal{L} depending on θ only such that all boundary gaps, with respect to $\text{NN}_{\mathcal{P}}$, have a gap parameter smaller than \mathcal{L} .*

Proof: Let x be a point on a boundary edge B . Let $p, q \in B \cap \mathcal{P}$ be the nodes of B . By the definition of the el function, the size of the gap at x that contains either p or q must be smaller than $\text{el}(x)$. By Theorem 5.7, the gap size is at most $\frac{1}{C_1} \text{NN}_{\mathcal{P}}(x)$. To obtain the result, set $\mathcal{L} > \frac{1}{C_1}$. \square

Theorem 5.12 *The geodesic distance between a node p of mesh $M = (\mathcal{P}, E, \mathcal{B})$ and a boundary segment B , such that p is not incident to B , is at least $b \text{NN}_{\mathcal{P}}(p)$, where $b = \sin \theta$.*

Proof: The geodesic distance between p and B must be at least as large as the distance between p and e , where e is the closest opposite edge to p in a triangle T incident to p . Let H be the height to e , let L be the length of e_1 , one of the other edges of T . Since the angle between e_1 and e is larger than θ , $H > L \sin \theta > \text{NN}_{\mathcal{P}}(p) \sin \theta$. \square

5.3 From a spacing function to a mesh

We now show that points spaced according to a spacing function, with bounded gaps, can be used to build a bounded aspect-ratio mesh. Furthermore, the el function of the resulting mesh is similar up to a constant factor to the original spacing function.

Theorem 5.13 *Let \mathcal{P} be a set of points β -spaced according to a 1-Lipschitz function f . Let $b < \mathcal{L}$ be two positive constants with the following properties:*

- I. *The geodesic distance between any point $p \in \mathcal{P}$ and a boundary segment B not incident to p is greater than $bf(p)$.*
- II. *\mathcal{P} has the bounded \mathcal{L} -gaps property with respect to f .*

Then the constrained Delaunay triangulation of \mathcal{P} , $CDT(\mathcal{P})$, has a bounded aspect-ratio. In particular, let T be a triangle of $CDT(\mathcal{P})$, let S be its circumsphere, and let θ be its smallest angle:

1. *if S is an interior gap then $\sin \theta \geq \frac{1}{\mathcal{L}(2+\beta)}$; and*
2. *if S is not an interior gap then $\sin \theta \geq \frac{1}{4\mathcal{L}^2(\mathcal{L}+1/b)(2+\beta)}$.*

Proof: Let $T \in \text{CDT}(\mathcal{P})$, and o, p, q be its vertices. Let (p, q) be the shortest edge of T and assume $\|p - q\| = d$. Because \mathcal{P} is β -spaced, $f(p) + f(q) \leq \beta d$. Without loss of generality, assume $f(p) \leq \beta d/2$, then since f is 1-Lipschitz $f(q) \leq f(p) + d < d(1 + \beta/2)$. Let R be the radius of the Delaunay ball S of T . Note that S is empty in the constrained sense: no other node of \mathcal{P} is visible to p, q or o inside S . Let c be the center of S . If c is visible to either p or q , then S is an interior gap, otherwise S contains a boundary gap. We now argue according to these two cases:

1. S is an interior gap: without loss of generality, since our bound on $f(q)$ is weaker, assume c is visible to q . By property II, $R \leq \mathcal{L}f(q) \leq \mathcal{L}d(1 + \beta/2)$. The corresponding angle bound is:

$$\sin \theta = \frac{d}{2R} \geq \frac{2}{\mathcal{L}(2 + \beta)}$$

2. S is not an interior gap: There exist boundary edges blocking c from p and q . Let B be one of the edges that intersect the radius from q to c . Note that B can not be the edge (p, q) itself, for otherwise, it would imply that the shortest edge of T separates the circum-center away from T . Consequently, the smallest triangle angle is larger than 90° , a contradiction. Also, the vertices of B are on or outside S : if B is one of the triangle edges (p, o) or (q, o) , then this is true; Otherwise, assume the triangle has no boundary edges. Then if B has a node inside S , there must exist a node of some boundary edge visible to either p or q within S .

Without loss of generality, since p and q can not both be incident to B , assume q is not incident to B . By property I, the distance l between q and B is at least $bf(q)$. The distance l must be smaller than $\mathcal{L}f(q)$, for otherwise S contains an interior \mathcal{L} -gap through q . As a result it contains a node u visible to q , which is a contradiction. Hence $bf(q) < l < \mathcal{L}f(q)$.

If B is at distance l , by the bounded boundary gaps property, the length of B is at most $2(f(q) + l)\mathcal{L}$. Let $2K$ be the length of the chord at distance l from q . Note that $K^2 = 2Rl - l^2$. If $K > (f(q) + l)\mathcal{L}$ then the gap at B is shorter than the chord and one of the boundary nodes of edge B are inside S , a contradiction. The proof proceeds by showing that if $R = 4(\mathcal{L}^3 + \mathcal{L}^2/b)f(q)$ (assuming \mathcal{L} is a large enough constant) then $K > (f(q) + l)\mathcal{L}$. It suffices to observe that $K > (f(q) + l)\mathcal{L}$ at distance $l = bf(q)$ and $l = \mathcal{L}f(q)$. At the distances between them, the size of B is a linear interpolation and will be smaller than the chord length as well.

Applying the bound on $f(q)$ to R , we obtain $R \leq 2\mathcal{L}^2(\mathcal{L} + 1/b)(\beta + 2)d$, and

$$\sin \theta \geq \frac{1}{4\mathcal{L}^2(\mathcal{L} + 1/b)(2 + \beta)}$$

□

The following theorem connects the spacing function f to the NN function of \mathcal{P} . By Theorem 5.7, a similar connection between $\text{el}_{\text{CDT}(p)}$ and f hold as a consequence.

Theorem 5.14 *Let \mathcal{P} be a set of nodes β -spaced according to a 1-Lipschitz function f , with the bounded gap property parameterized by \mathcal{L} . Then:*

$$\frac{\text{NN}(x)}{3\mathcal{L}^2(2 + \mathcal{L})^2} \leq f(x) \leq (2\beta + 1)\text{NN}(x)$$

Proof:

- $f(x) \leq (2\beta + 1)\text{NN}(x)$: First assume $p \in \mathcal{P}$, and let $q \in \mathcal{P}$ be the node nearest to p . Then $\text{NN}(p) = \|p - q\|$. On the other hand, since \mathcal{P} is β -spaced according to f , $f(p) \leq \beta\|p - q\|$. Therefore, $f(p) \leq \beta\text{NN}(p)$.

Now let $x \in \Omega$. By definition, $\text{NN}(x)$ is the radius r of the smallest ball that contains two (visible from x) nodes, say p and q , from \mathcal{P} . Let $\|x - p\| = r$ and $\|x - q\| = r_1 \leq r$. Because f is 1-Lipschitz, $f(x) \leq f(p) + r \leq \beta\text{NN}(p) + r$. But $\text{NN}(p) \leq 2r$, for q is no more than $2r$ (geodesic) distance away from p . Thus, $f(x) \leq 2\beta r + r \leq (2\beta + 1)\text{NN}(x)$.

- $\text{NN}(x)/(3\mathcal{L}^2(2 + \mathcal{L})^2) \leq f(x)$: Let $x \in \Omega$. Consider a ball S of radius $\mathcal{L}f(x)$ through x . If the center of S is visible from x , then by the bounded gaps property it contains a point of \mathcal{P} visible to x . Otherwise, a boundary edge B blocks the center of S . One of the nodes of B , p , is then at distance at most $f(x)\mathcal{L}(\mathcal{L} + 2)$ from x , using the bound on boundary gaps. Repeating the argument with respect to p , a point q exists at distance $f(p)\mathcal{L}(\mathcal{L} + 2)$ from p . The distance of q from x is at most $f(x)\mathcal{L}(\mathcal{L} + 2) + f(p)\mathcal{L}(\mathcal{L} + 2)$. Using the Lipschitz derived bound on $f(p)$, and assuming that $\mathcal{L} \geq 2$, the distance from x to q is at most $3f(x)\mathcal{L}^2(\mathcal{L} + 2)^2$. So within a geodesic distance of at most $3f(x)\mathcal{L}^2(\mathcal{L} + 2)^2$ two non-incident features are found, providing a bound on $\text{NN}(x)$.

□

5.4 Properties of coarsening functions

This section establishes some of the properties of these coarsening functions. First, the coarsening functions are shown to be 1-Lipschitz as well.

Lemma 5.15 (1-Lipschitz) *If g is 1-Lipschitz, then for any $\mathcal{C} > 1$, $f_{g,\mathcal{C},\mathcal{P}}$ is 1-Lipschitz.*

Proof: Let $f = f_{g,\mathcal{C},\mathcal{P}}$. Let x, y be two points in Ω . Without loss of generality, $f(x) < f(y)$. Assume further $f(x) = \hat{f}_{\mathcal{C},p_i}(x)$, and $f_{\mathcal{C},\mathcal{P}}(y) = \hat{f}_{\mathcal{C},p_j}(y)$. Then,

$$f(y) - f(x) = \hat{f}_{\mathcal{C},p_j}(y) - \hat{f}_{\mathcal{C},p_i}(x) \leq \hat{f}_{\mathcal{C},p_i}(y) - \hat{f}_{\mathcal{C},p_i}(x) \leq \|x - y\|$$

The last transition uses the obvious fact that $\hat{f}_{\mathcal{C},p_i}$ is 1-Lipschitz. □

Lemma 5.16 (linearity and monotonicity) *For any $\mathcal{C}_1, \mathcal{C}_2 \geq 1$*

$$f_{g,\mathcal{C}_1,\mathcal{P}}(x) \leq f_{g,\mathcal{C}_1\mathcal{C}_2,\mathcal{P}}(x) \leq \mathcal{C}_2 f_{g,\mathcal{C}_1,\mathcal{P}}(x).$$

Proof: Since $f_{g,\mathcal{C}_1,\mathcal{P}}$ is a minimum over a collection of functions $\{\hat{f}_{p_i} \mid p_i \in \mathcal{P}\}$ it suffices to show the inequality for each such function. The first inequality is obvious. The second one holds because:

$$\hat{f}_{\mathcal{C}_1\mathcal{C}_2,p_i}(x) = \mathcal{C}_1\mathcal{C}_2 g(p_i) + \|p_i - x\| \leq \mathcal{C}_2(\mathcal{C}_1 g(p_i) + \|p_i - x\|) \leq \mathcal{C}_2 \hat{f}_{\mathcal{C}_1,p_i}(x).$$

□

The following Lemma strongly motivates the choice of the coarsening functions: they are the largest possible 1-Lipschitz functions that coarsen the original function by a factor of \mathcal{C} at the points \mathcal{P} .

Lemma 5.17 (maximality of $f_{\mathcal{C},\mathcal{P}}$) *Let $\mathcal{C} > 1$. Let h be a 1-Lipschitz function over the domain Ω such that for all $p_i \in \mathcal{P}$ $h(p_i) \leq \mathcal{C}g(p_i)$. Then for all $x \in \Omega$, $h(x) \leq f_{g,\mathcal{C},\mathcal{P}}(x)$.*

Proof: $f_{\mathcal{C},\mathcal{P}}$ is a minimum over a collection of functions $\{\hat{f}_{p_i} \mid p_i \in \mathcal{P}\}$. It suffices to show $h(x) \leq \hat{f}_{\mathcal{C},p_i}(x)$. Recall that $\hat{f}_{\mathcal{C},p_i}(x) = \mathcal{C}g(p_i) + \|p_i - x\|$. Because h is 1-Lipschitz,

$$h(x) \leq h(p_i) + \|p_i - x\| \leq \mathcal{C}g(p_i) + \|p_i - x\| = \hat{f}_{\mathcal{C},p_i}(x).$$

□

6 Bounds for one-level mesh coarsening

Given an initial bounded aspect-ratio mesh M_0 , algorithm `ONE_LEVEL_COARSEN` of Figure 3 is used to produce the meshes of the coarsening sequence. To produce the i -th mesh in the sequence M_i , the algorithm is called with coarsening factor $\mathcal{C} = 2^i$. Therefore, obtaining bounds on the one-level scheme translates with little extra effort into bounds for the entire coarsening scheme.

Let \mathcal{C} stand for an arbitrary coarsening factor. In this section we prove the following two bounds on the quality of the coarsening meshes:

1. **aspect-ratio bound:** if mesh M_0 is a bounded aspect-ratio mesh, and the coarsening parameters are β and γ (chosen carefully as a function of the aspect-ratio of M_0) then the mesh M_1 resulting from the coarsening algorithm is a bounded aspect-ratio mesh. The bound depends on the quality of M_0 only.
2. **size optimality:** let M be some other bounded aspect-ratio mesh coarsening M_0 by coarsening factor \mathcal{C} , then $|M_1| = O(|M|)$.

6.1 Aspect-ratio bound

The main tool used in proving the aspect-ratio bound is the one developed in the previous section, the equivalence between bounded aspect-ratio meshes and points spaced according to a spacing function with bounded gaps. The input to algorithm `ONE_LEVEL_COARSEN` is a bounded aspect-ratio mesh $M_0 = (\mathcal{P}_0, E_0, \mathcal{B})$, therefore the equivalence theorem implies that the mesh nodes, \mathcal{P}_0 , are spaced according to a spacing function f_0 and its gaps are bounded using a gap parameter \mathcal{L}_0 . The resulting coarsened mesh $M_1 = (\mathcal{P}_1, E_1, \mathcal{B})$, is, by construction, spaced according to a spacing function f_1 . To employ the equivalence theorem, and show that M_1 has bounded aspect-ratio, all that remains is to prove that \mathcal{P}_1 has bounded gaps in terms of f_1 . This subsection is aimed at showing that the parameters of algorithm `ONE_LEVEL_COARSEN` can be set to guarantee the bounded gaps property for mesh M_1 , with a gap parameter \mathcal{L}_1 depending on \mathcal{L}_0 only.

M_1 is constructed by first coarsening the spacing function f_0 . The resulting function f_1 is a \mathcal{C} coarsening of the function f_0 over the point set \mathcal{P}_0 . In this section, the only assumption made is that M_1 is constructed over M_0 using a spacing function $f_1 \geq f_0$. The coarsening parameter \mathcal{C} is arbitrary and plays no role in the proofs. Furthermore, it is assumed that $f_1 \leq \text{lfs}_{\mathcal{B}}$, i.e., that coarsening is subjected to the limitation posed by the local feature size of the boundary description, and that the angle in the domain formed between two incident segments of \mathcal{B} is greater than 60° . This is an often used restriction, see for example [4, 19].

The proof proceeds in increasing order of dimensions. First, the boundary gaps of M_1 are shown to be bounded, and then the interior gaps are bounded.

The following lemma establishes the fact that nodes on different boundary elements do not share conflict graph edges, so different segments can be coarsened independently.

Lemma 6.1 (Boundary independence) *Let $M_0 = (\mathcal{P}_0, E_0, \mathcal{B})$. Let f_1 be a spacing function such that $f_1 \leq \text{lfs}_{\mathcal{B}}$. Let $\alpha > 60^\circ$ be the minimal angle between two segments of \mathcal{B} . Let the spacing parameter $\beta \geq \frac{1}{2 \sin \frac{\alpha}{2} - 1}$. The following are true during the coarsening algorithm:*

1. f_1 -spheres placed on two non-incident boundary segments, or a segment and a boundary point not on the segment, or two boundary points, do not intersect.
2. The f_1 -sphere of a node retained on segment $B_1 \in \mathcal{B}$ can not intersect the f_1 -sphere of a node retained on segment $B_2 \in \mathcal{B}$ such that B_1 and B_2 are incident.

Proof:

1. Let $p \in \mathcal{P}_0$ be a node on a boundary point or segment. Let q be another such node on a disjoint boundary point or segment. Since they are on disjoint boundary elements, $f_1(p) \leq \text{lfs}_{\mathcal{B}}(p) \leq \|p - q\|$ and similarly for q . Clearly, if $\beta \geq 2$, the spheres with radius $f_1(p)/\beta$ and $f_1(q)/\beta$ do not intersect.

2. Let $s \in \mathcal{P}_0 \cap \mathcal{B}$ be the segment vertex common to B_1 and B_2 . By the coarsening algorithm, \mathcal{P}_1 is built in increasing order of dimensions, and s is added to \mathcal{P}_1 before any point q on $B_1 \cup B_2$.

When $q \in B_1$ is added to \mathcal{P}_1 , its f_1 -sphere does not intersect that of s , i.e., let $d = \|s - q\|$ then $f_1(s) + f_1(q) \leq \beta d$.

To show q 's sphere does not intersect any sphere of a similar node on B_2 it suffices to show that q 's f_1 -sphere lies below the bisector between B_1 and B_2 . This happens when the radius of the sphere $r = f_1(q)/\beta \leq d \sin(\alpha/2)$, or, equivalently, when $f_1(q) \leq \beta d \sin(\alpha/2)$. Because f_1 is 1-Lipschitz, for some $-1 \leq \gamma \leq 1$, $f_1(q) = f_1(s) + \gamma d$. The previous facts imply $2f_1(q) - \gamma d < \beta d$, or $f_1(q) \leq d(\beta + \gamma)/2 \leq d(\beta + 1)/2$. It suffices to require then that $d \frac{\beta+1}{2} \leq \beta d \sin \frac{\alpha}{2}$ or that $\beta \geq 1/(2 \sin \frac{\alpha}{2} - 1)$. This also clarifies the need to restrict $\alpha > 60^\circ$.

□

Theorem 6.2 (Bounded boundary gaps) *Let mesh $M_0 = (\mathcal{P}_0, E_0, \mathcal{B})$ be a bounded aspect-ratio mesh with the bounded \mathcal{L}_0 -gaps property. Let $M_1 = (\mathcal{P}_1, E_1, \mathcal{B})$ be a coarsening of M_0 by algorithm ONE_LEVEL_COARSEN using coarsening function f_1 , and a spacing constant $\beta \geq 1$. The boundary gaps of M_1 can be bounded using the boundary gap parameter $\hat{\mathcal{L}}_1$:*

$$\hat{\mathcal{L}}_1 = \mathcal{L}_0 + \frac{4(1 + \mathcal{L}_0)}{\beta - 1} \left[1 + \frac{1}{\beta - 1} \right]$$

Proof: Let $B \in \mathcal{B}$ be some boundary segment. Using Lemma 6.1 only the nodes $p \in \mathcal{P}_0 \cap S$ need to be considered, and the vertices of B belong to \mathcal{P}_1 . Let $x \in B$ be an arbitrary domain point. Without loss of generality, let $f_1(x) = 1$. Let d be the length of \mathcal{I} , a sub-segment of B starting at x and extending to x 's right. Segments extending to the left are treated similarly. The proof proceeds to show that if $d \geq \hat{\mathcal{L}}_1$ then $\mathcal{I} \cap \mathcal{P}_1$ is not empty.

Split \mathcal{I} into three smaller intervals $\mathcal{I}_1, \mathcal{I}_2$ and \mathcal{I}_3 . The scheme of the proof is to argue that the sizes d_1, d_2 and d_3 of the three intervals can be set such that: (1) \mathcal{I}_2 must contain a node of \mathcal{P}_0 , (2) a node of $\mathcal{P}_0 \cap \mathcal{I}_2$ can not be pruned out by any node of \mathcal{P}_0 which is outside \mathcal{I} since their f_1 -spheres can not intersect.

If (1) and (2) are valid, this implies $\mathcal{I} \cap \mathcal{P}_1$ is not empty, since either \mathcal{I}_2 contains a node of \mathcal{P}_1 , or $\mathcal{I} \setminus \mathcal{I}_2$ contains a node of \mathcal{P}_1 which is responsible for \mathcal{I}_2 being empty.

By the Lipschitz condition, $f_1(w) \leq 1 + \|x - w\|$ for any point $w \in \Omega$, and without loss of generality we assume that $f_1(w) = 1 + \|x - w\|$. Since f_1 in reality must be smaller, this assumption provides us with an upper bound for $\hat{\mathcal{L}}_1$.

Setting d_1 : if $d_1 = \frac{2}{\beta-1}$, then for any $y \in \mathcal{I}_2$ the sphere $\mathcal{S}(y, f_1(y)/\beta)$ is disjoint from any sphere placed left of x . This can be proven by a straightforward application of the value of f_1 .

Setting d_2 : let $d_2 = \mathcal{L}_0(1 + \frac{2}{\beta-1})$, and let y be the point at the intersection of \mathcal{I}_1 and \mathcal{I}_2 . Mesh M_0 has boundary gaps bounded by \mathcal{L}_0 , hence an interval of length $\mathcal{L}_0 f_0(y)$ must contain a point of \mathcal{P}_0 . Since d_2 was set to be $\mathcal{L}_0 f_0(y)$, $\mathcal{P}_0 \cap \mathcal{I}_2$ is not empty.

Setting d_3 : let $d_3 = \frac{2+2\mathcal{L}_0}{\beta-1} \left(1 + \frac{2}{\beta-1}\right)$ then for any $z \in \mathcal{I}_2$ the sphere $\mathcal{S}(z, f_1(z)/\beta)$ is disjoint from any sphere placed right of \mathcal{I} . This can be proven by a straightforward application of the value of f_1 . Summing the interval lengths and letting $\hat{\mathcal{L}}_1 = d_1 + d_2 + d_3$, yields the result. \square

The bound on interior gaps is more complex. However, the proof of the bound on interior gaps also follows the ideas of the proof above. The difficulty comes from the fact that the interior gaps take the form of spheres which are harder to work with than segments. As a consequence, our constant bounds for interior gaps are weaker.

Theorem 6.3 (Interior gaps of 2D coarsening) *Let mesh $M_0 = (\mathcal{P}_0, E_0, \mathcal{B})$ be a bounded aspect-ratio mesh with bounded gaps parameter \mathcal{L}_0 . Let $f_1 \geq f_0$ be the coarsening function. Let $M_1 = (\mathcal{P}_1, E_1, \mathcal{B})$ be the coarser mesh, obtained with with spacing parameter $\beta = \max(3, 12\sqrt{\mathcal{L}_0})$, and boundary protection parameter $\gamma = 1/(\hat{\mathcal{L}}_1(8 + \mathcal{L}_0))$. Then, mesh M_1 has the bounded gaps property with parameter \mathcal{L}_1 , where $\mathcal{L}_1 = \max\{4, 12\mathcal{L}_0, (4\hat{\mathcal{L}}_1^2 + 1)(3\mathcal{L}_0 + 2)\}$.*

Proof: Let $u \in \Omega$ be an arbitrary point, and let $S = \mathcal{S}(c, R)$ be a sphere through u . Consider Figure 5. Without loss of generality, let $f_1(u) = 1$, $u = (0, 0)$, and let the center c lie on the y -axis. S is an interior gap of M_1 at u , so the radius (u, c) is in the domain.

This proof shows that R can be set to a constant large enough such that S must contain a point of \mathcal{P}_1 visible to u , and thus all gaps are bounded.

Let S_0 be a smaller ball nested in S of radius $3\mathcal{L}_0$ centered at $(0, 2 + 3\mathcal{L}_0)$. Let $v = (0, 2)$ be a point on the circumference of S_0 . Then $f_1(v) \leq f_1(u) + 2 = 3$, so S_0 , being of radius $3\mathcal{L}_0$, is an interior gap of M_0 containing a node $q \in \mathcal{P}_0$ visible to v . By Lemma 6.4, either q is visible to u as well, or S contains a boundary segment's vertices, which are always chosen by the algorithm to remain in \mathcal{P}_1 . In the latter case, S is not empty with respect to \mathcal{P}_1 and the proof is complete. Therefore, assume q is visible to u . The proof proceeds as follows:

1. If $R \geq \max(12\mathcal{L}_0, 4)$ then by Lemma 6.5 node q can not be eliminated from \mathcal{P}_1 by any \mathcal{P}_1 node outside S .
2. If $R \geq (4\hat{\mathcal{L}}_1^2 + 1)(3\mathcal{L}_0 + 2)$ then no boundary segment protection can eliminate q from \mathcal{P}_1 , by Lemma 6.8.

Hence, $q \in \mathcal{P}_1$ or it has been eliminated because of a conflict with some node $p \in S \cap \mathcal{P}_1$ visible to q . Then again by Lemma 6.4 S must contain a node of \mathcal{P}_1 visible to u , and M_1 possesses the bounded gaps property. \square

Lemma 6.4 (Visibility) *Let S be a sphere through some domain point $u \in \Omega$. Let v be some node visible to u , and w is a node visible to v . Assume further that $v, w \in S$. If w is not visible to u then some vertex of a boundary segment is in S and visible to u .*




Figure 5: After one level coarsening, a circle of radius $12\mathcal{L}_0 f_1(u)$ must contain a node of \mathcal{P}_1

Proof: Consider the triangle with end-points (u, v, w) . This triangle is contained in S by convexity. Let B be the segment intersecting triangle edge (u, w) closest to u . Since w is not visible to u , such a segment exists. No segments can intersect the edges (u, v) and (v, w) . Therefore, one of the vertices of B , call it p , must be in the triangle. Moreover, no other segment intersecting (u, w) can obscure p from u .

Note however that segments totally contained in the triangle can obscure p from u . In that case, let q be the vertex of the closest such segment and it must be visible to u . \square

Lemma 6.5 *If $R \geq \max(4, 12\mathcal{L}_0)$, and $\beta \geq \max(3, 12\sqrt{\mathcal{L}_0})$, then no node p placed outside S can conflict with a node q placed in S_0 .*

Proof: Refer back to Theorem 6.3 for the setting. Without loss of generality, let $q = (x, y)$ be on the boundary of S_0 (all bounds derived hold if q is internal to B_0 as well).

The following properties of q are used:

- I) By Lemma 6.6, the distance from q to the boundary of S is larger than $y/2$.
- II) By Lemma 6.7, if $\mathcal{L}_0 \geq 4$ and $\beta = 12\sqrt{\mathcal{L}_0}$, then $f_1(q)/\beta \leq y/6$. (It is also assumed that $\beta \geq 3$. Note that \mathcal{L}_0 can be assumed to be as large as necessary — M_0 still has the bounded gaps property with a larger \mathcal{L}_0).

Let S_1 be the sphere centered at q and of radius $f_1(q)/\beta$ (see the dashed ball in Figure 5). Let $p \in \mathcal{P}_1$ be some node placed outside S . We prove, by way of contradiction, that S_1 can not intersect the f_1 sphere of p . Assume the two spheres do intersect. Property I implies $d = \|p - q\| \geq y/2$. Thus,

$$f_1(p) \leq d + f_1(q) \quad (\text{because } f \text{ is Lipschitz})$$

$$\begin{aligned}
f_1(q) + f_1(p) &\geq \beta d && \text{(because they share a conflict graph edge)} \\
2f_1(q) + d &\geq \beta d && \text{(follows from the two previous lines)} \\
2f_1(q) &\geq (\beta - 1)d \geq (\beta - 1)y/2 && \text{(by property I)} \\
\beta y/3 &\geq 2f_1(q) \geq (\beta - 1)y/2 && \text{(by property II)} \\
\beta/3 &\geq (\beta - 1)/2
\end{aligned}$$

This is a contradiction when $\beta > 3$ and hence q and p have disjoint f_1 -spheres. \square

Lemma 6.6 *The distance from any point (x, y) on the boundary of S_0 to the boundary of S is greater than $y/2$.*

Proof: Note that $y > 2$. The distance from (x, y) to the boundary of S is equal to $R - \sqrt{x^2 + (y - R)^2}$, where recall $R = 12\mathcal{L}_0$. Because (x, y) is on the boundary of S_0 , we have $x^2 + (y - r - 2)^2 = r^2$, where $r = 3\mathcal{L}_0 = R/4$. By substituting $x^2 = r^2 - (y - r - 2)^2$, it is easily shown that $x^2 + (y - R)^2 < (R - y/2)^2$. \square

Lemma 6.7 *Let $q = (x, y)$ be on the boundary of S_0 . If $\mathcal{L}_0 > 4$, and let $\beta \geq \max(12\sqrt{\mathcal{L}_0}, 3)$, then $f_1(q)/\beta \leq y/6$.*

Proof: $f_1(u) = 1$ and f_1 is 1-Lipschitz, therefore $f_1(q) \leq 1 + \|u - q\|$. It suffices to show that $1 + \|u - q\| \leq \beta y/6$. If $\beta > 3$ then $\beta y/6 - 1 > 0$, since $y \geq 2$. Since $u = (0, 0)$, $(\|u - q\|)^2 = x^2 + y^2$, and the goal is to show $x^2 + y^2 \leq (\beta y/6 - 1)^2$. q is on the sphere $S_0 = \mathbf{S}((0, 2 + r), r)$ where $r = 3\mathcal{L}_0$, therefore $x^2 + y^2 = -4 - 4r + 2yr + 4y$. Using again that fact that $y \geq 2$, this leads to the bound on β . \square

Lemma 6.8 *If $R \geq (4\hat{\mathcal{L}}_1^2 + 1)(3\mathcal{L}_0 + 2)$ and $\gamma \leq \frac{1}{\hat{\mathcal{L}}_1(8 + 2\mathcal{L}_0)}$ then $q \in S_0$ can not be eliminated from \mathcal{P}_1 by the protection zone around any boundary edge e of M_1 .*

Proof: The proof proceeds in two parts: first, R is set so it can be shown that no boundary edge can be too close to S_0 , and second, γ is set so that the protection zone around e can not eliminate q when e is not too close to S_0 .

It can be assumed that the nodes of e are not inside S . The reasoning is as follows: assume that e is some edge such that q is within its protective zone. Note that since e must be visible to q , and the nodes of no other boundary edge can intrude on the protective zone of e (for example, by Lemma 6.1), the two nodes of e are visible to q . If any of these vertices are inside S , then by Lemma 6.4, they are visible to u as well contradicting the assumption that S is a gap at u .

e can not be close to S_0 : Consider a ring G of width 1 around S_0 . Using a constant $0 \leq \epsilon \leq 1$, a point $z = (x, y) \in G$ has the following property: $x^2 + (y - (3\mathcal{L}_0 + 2))^2 = (3\mathcal{L}_0 + \epsilon)^2$, or, $x^2 + y^2 = (3\mathcal{L}_0 + \epsilon)^2 - (3\mathcal{L}_0 + 2)^2 + 2y(2\mathcal{L}_0 + 2)$.

e 's nodes are outside S , and furthermore, e is a boundary gap through z . The chord is the shortest segment through z with vertices outside S . The proof proceeds by showing that a boundary gap through z must be shorter than the chord through z , thus a node of e must be inside S , contradicting the assumptions on e . Therefore, no boundary edge can pass through z .

Let l be half the length of a chord through z , then $l^2 = 2yR - (x^2 + y^2)$. The length of an empty 1D gap at z , by Theorem 6.2, must be smaller than $\hat{\mathcal{L}}_1 f_1(z) \leq \hat{\mathcal{L}}_1(1 + \sqrt{x^2 + y^2})$. However, if $R > (4\hat{\mathcal{L}}_1^2 + 1)(3\mathcal{L}_0 + 2)$ it is easy to show, using the formulation for $x^2 + y^2$ above, that $\sqrt{2yR - (x^2 + y^2)} > \hat{\mathcal{L}}_1(1 + \sqrt{x^2 + y^2})$.

if e is far from S_0 , its protective zone can not eliminate q : Let d be the distance between q and e . Since $q \in S_0$, $f_1(q) \leq 3 + 6\mathcal{L}_0$. Let $z \in e$ be the point of e closet to q . $f_1(z) \leq 3 + 6\mathcal{L}_0 + d$. Therefore, the length of e is bounded (see Theorem 6.2): $|e| \leq 2\hat{\mathcal{L}}_1(3 + 6\mathcal{L}_0 + d)$. The protective zone around e is a rectangle projecting a distance $\gamma|e|$ around it. γ should be set so that $\gamma|e| < d$, or $\gamma 2\hat{\mathcal{L}}_1(3 + 6\mathcal{L}_0 + d) \leq d$. Rearranging, this reduces to: $(\gamma\hat{\mathcal{L}}_1(6 + 2\mathcal{L}_0))/(1 - 2\hat{\mathcal{L}}_1\gamma) \leq d$, or, since $d > 1$, $\gamma \leq 1/(\hat{\mathcal{L}}_1(8 + 2\mathcal{L}_0))$. \square

6.2 Size optimality of one-level coarsening

Since the number of elements in a bounded aspect-ratio mesh can be directly related to its spacing function, the size optimality of mesh M_1 is a result of the optimality of the coarsening functions used.

The following Lemma and its corollary relate the size of a bounded aspect-ratio mesh to its spacing function NN.

Lemma 6.9 *Let M be a mesh with smallest angle θ . Let T be a triangle of M , then:*

$$\frac{\sin^2 \theta}{2C_2^2} \leq \int_T \frac{dA}{\mathbf{NN}_M^2} \leq \frac{1}{2 \sin \theta C_1^2}$$

Proof: Let $\mathbf{el}(T)$ be the longest side of T , α , β , and γ its angles. The area of T is $A(T) = 0.5\mathbf{el}^2(T) \sin \alpha \sin \beta / \sin \gamma$. Using the bound on the smallest angle, $0.5\mathbf{el}^2(T) \sin^2 \theta \leq A(T) \leq 0.5\mathbf{el}^2(T) / \sin \theta$. By Theorem 5.7

$$\frac{A(T)}{C_2^2 \mathbf{el}^2(T)} \leq \int_T \frac{dA}{\mathbf{NN}_M^2} \leq \frac{A(T)}{C_1^2 \mathbf{el}^2(T)}$$

or

$$\frac{\sin^2 \theta}{2C_2^2} \leq \int_T \frac{dA}{\mathbf{NN}_M^2} \leq \frac{1}{2C_1^2 \sin \theta}$$

\square

Corollary 6.10 *The number of elements of mesh M with smallest angle θ is:*

$$a_1(\theta) \int_{\Omega} \frac{dA}{\mathbf{NN}^2} \leq |M| \leq a_2(\theta) \int_{\Omega} \frac{dA}{\mathbf{NN}^2}$$

where $a_1(\theta)$, $a_2(\theta)$ are constants depending on θ only.

Theorem 6.11 (Size optimality of one-level coarsening) *Consider the following three meshes:*

- M_0 , a mesh with smallest angle θ .
- M_1 , the coarsening of M_0 via algorithm ONE_LEVEL_COARSEN with coarsening factor \mathcal{C} .
- M'_1 , a mesh with smallest angle θ' , whose elements are at most a factor \mathcal{C} greater than M_0 , i.e., $\mathbf{el}_{M'_1} \leq \mathcal{C}\mathbf{el}_{M_0}$.

There exists a constant b , depending on θ and θ' only, such that $|M_1| \leq b|M'_1|$.

Proof: Since $\text{el}_{M'_1} \leq \mathcal{C}\text{el}_{M_0}$, Theorem 5.7 implies:

$$\text{NN}_{M'_1} \leq C_2(\theta')\text{el}_{M'_1} \leq \mathcal{C}C_2(\theta')\text{el}_{M_0} \leq \frac{\mathcal{C}C_2(\theta')}{C_1(\theta)}\text{NN}_{M_0}$$

Therefore, we can assume that $\text{NN}_{M'_1} \leq k\mathcal{C}\text{NN}_{M_0}$ where $k = \max\left(1, \frac{C_2(\theta')}{C_1(\theta)}\right)$. Since $\text{NN}_{M'_1}$ is 1-Lipschitz, by Lemma 5.17:

$$\text{NN}_{M'_1} \leq f_{\text{NN}_{M_0}, k\mathcal{C}, \mathcal{P}_0}$$

By Lemma 5.16: $\text{NN}_{M'_1} \leq k f_{\text{NN}_{M_0}, \mathcal{C}, \mathcal{P}_0}$. By Theorem 5.14, $\text{NN}_{M'_1} \leq k(2\beta+1)\text{NN}_{M_1}$ and Corollary 6.10 implies the result. \square

7 Optimality of the coarsening sequence

The optimality of the coarsening sequence we produce is mostly due to the optimality properties of the one-level coarsening algorithm. It only remains to show that neighboring meshes in the sequence approximate each other well. This section states the optimality theorem for the coarsening sequence, and shows the meshes are locally similar as well.

Theorem 7.1 (Bounded aspect-ratio coarsening sequence) *Let M_0 be a mesh with smallest angle θ . The coarsening sequence (M_1, \dots, M_k) constructed by Algorithm MULTI_LEVEL_COARSEN 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 below by θ_1 .*
2. **local similarity:** *There is a constant \mathcal{I} depending on θ only such that for each $1 \leq i \leq k$, $\text{el}_{M_{i+1}} \leq \mathcal{I}\text{el}_{M_i}$.*
3. **size guarantee:** *Let (M'_1, \dots, M'_k) be any other bounded aspect-ratio coarsening sequence of M_0 with smallest angle bound θ' , then there exists a constant b depending on θ , and θ' only such that $|M_i| \leq b|M'_i|$ $1 \leq i \leq k$*

Proof:

1. Direct corollary of Theorem 6.3.
2. Mesh M_i is generated using $f_{\text{NN}_{\mathcal{P}_0}, a, \mathcal{P}_0}$, and mesh M_{i+1} is generated using spacing function $f_{\text{NN}_{\mathcal{P}_0}, 2a, \mathcal{P}_0}$. By Lemma 5.16: $f_{\text{NN}_{\mathcal{P}_0}, 2a, \mathcal{P}_0} \leq 2f_{\text{NN}_{\mathcal{P}_0}, a, \mathcal{P}_0}$. By Theorem 5.14: $\text{NN}_{M_{i+1}} \leq c_1 f_{\text{NN}_{\mathcal{P}_0}, 2a, \mathcal{P}_0} \leq 2c_1 f_{\text{NN}_{\mathcal{P}_0}, a, \mathcal{P}_0} \leq 2c_1 c_2 \text{NN}_{M_i}$. By Theorem 5.7, for some constant \mathcal{I} : $\text{el}_{M_{i+1}} \leq \mathcal{I}\text{el}_{M_i}$.
3. It follows from Theorem 6.11. \square

8 Run-time complexity and efficient variations

The function-based coarsening algorithm generates one mesh of the coarsening sequence using the following two main computational steps: (1) computing the new coarsening function at the set of

candidate nodes, and (2) computing an MIS of the conflict graph over the candidate nodes with respect to the coarsening function.

The first step can be done in $O(n \log n)$ using the additive Voronoi diagram (see for example the algorithm by Fortune [10]), where n is the size of the candidate node set over which the coarsening is performed. Note that the resulting mesh size, n_i , may be much smaller than n , depending on the choice of the candidate node set and the coarsening function. The second step can be expensive, since the conflict graph can be as large as $O(n^2)$.

The coarsening functions f_1, f_2, \dots, f_k suggested in previous sections are optimal for coarsening, (see for example Lemma 5.17). However, this optimality carries over to any sequence of coarsening functions g_1, g_2, \dots, g_k such that $a_1 g_i \leq f_i \leq a_2 g_i$.

This section suggests such a sequence of coarsening functions that approximates the spacing functions, and results in a more efficient algorithm. These functions are based on the balanced quad-tree data structure. We describe, at a high level, these functions and how to use them so that the complexity of generating the mesh at level i becomes $O(n_i)$, plus some pre-processing cost. A more detailed description can be found in Talmor [20].

Definition 8.1 (Balanced Quad-Tree [4]) *A quad-tree is a balanced quad-tree if no quad-tree box b neighbors a box whose side length is smaller than half the side length of b .*

A spacing function q_T can be associated with a balanced quad-tree T , such that values q_T achieves over a quad-tree box are similar to the side length of the box. For example, letting P stand for the quad-tree vertices, the function NN_P is one such function.

Furthermore, given a general spacing function g over a square domain, a balanced quad-tree T can be constructed such that its spacing function q_T approximates g , i.e., there exist two constants a_1 and a_2 such that $a_1 q_T \leq g \leq a_2 q_T$. The algorithm would therefore proceed by first converting the input mesh M to a quad-tree (rather than a spacing function) at a pre-processing cost of $O(n_0 \log n_0)$.

Coarsening a quad-tree function: To coarsen a quad-tree function q_T , its associated balanced quad-tree is coarsened. The coarsening method tries to replace four siblings that are all leaf boxes by their parent box. This is possible only if the four siblings do not have a small neighbor, hence coarsening proceeds from smaller leaf boxes to larger. If the box b_i has a smaller neighbor that could not be coarsened, do not coarsen b_i . Otherwise, replace box b_i and its other three siblings by their parent box. The new balanced quad-tree is composed of all the new coarsened boxes and the boxes that could not be coarsened. Note that if the boxes of T were sorted by size, it takes linear time to sort the boxes of the new tree: the boxes can be kept in a linked list of buckets, and if a box is replaced by a larger box, it is simply moved to the next bucket. The only non-linear cost is that of the creation and sorting of the initial quad-tree. This cost is $O(n_0 \log n_0)$.

Generating the coarsening sequence meshes: To generate the coarsening sequence meshes, the corresponding sequence of coarsening quad-trees is used. It is intuitively clear that even if a coarsened quad-tree box contains many nodes of the candidate node set, there is no need to consider them all while generating the mesh coarsening at the current level. It suffices to consider a subset of the candidate points that includes a constant number of points at each quad-tree box. The extra processing necessary can be easily amortized into the generation of the initial quad-tree, q_{T_0} , by associating a constant number of points with each quad-tree box. Using the pruned candidate point-set, the conflict graph is now linear in size.




Figure 6: crack plate mesh

9 Experiments

This paper introduced a coarsening strategy guaranteed to generate a bounded aspect-ratio coarsening sequence. The bounds however were theoretical in nature, and were merely shown to be some constant that depends on the quality of the original mesh without explicitly stating the size of this constant. The goal of this section is to demonstrate, using a simple experiment, that our approach produces coarsening sequences with good aspect-ratio in practice. The quality of our coarsening sequence is also compared with the quality of the topological-base MIS coarsening sequences introduced in 5.1. More experimental results, covering a larger variety of mesh types and sizes, can be found in Talmor [20].

The mesh being coarsened is 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 [11]. The mesh itself is plotted in Figure 6. This is a simple, medium-sized (about 10000 nodes) high quality mesh that is extremely graded, i.e., it contains elements with vastly different edge lengths.

We implemented both a variant of the function-based approach, and an MIS-based coarsening algorithm. Figure 7 displays the resulting coarsening sequences. The top row shows the function-based coarsening, the bottom row the MIS coarsening. The MIS reduces the size quickly, but as a result the smallest angle quality degrades by the second level of coarsening. The small angles typically form near the center, in the transition from the denser to the less dense areas. The function-based approach also offers a (smaller) geometric size reduction, but with a much better angle control. This example was chosen to demonstrate the inherent trade-off between mesh quality and mesh size reduction.

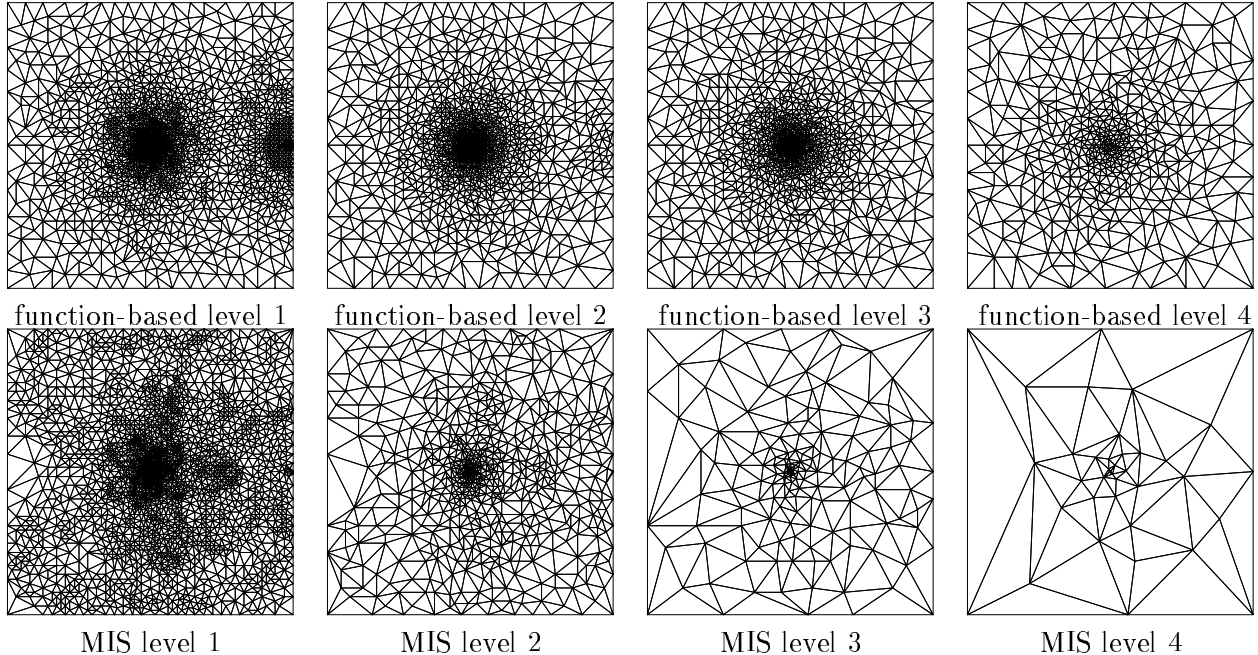


Figure 7: Coarsening sequences produced by the function-based approach and the MIS-based approach.

	function-based		MIS-based	
level	number of nodes	smallest angle	number of nodes	smallest angle
0	10183	37.01	10183	37.01
1	3635	15.07	3028	8.13
2	2379	17.1	726	3.97
3	1300	14.03	175	3.22
4	466	16.7	42	3.46
5	80	17.35	12	2.29

10 Conclusions and future work

In this paper, we have presented a new approach to mesh coarsening that generates a bounded aspect-ratio coarsening sequence that is optimal in both depth and width. To the best of our knowledge, this is the first provably good algorithm for unstructured mesh coarsening.

The complexity of our algorithm is asymptotically linear in the output size, plus a pre-processing cost that is amortized over the length of the coarsening sequence. However, we note that some steps in our algorithm are still too complicated and therefore they may lead to high constant factors in running time and make the algorithm hard to implement.

Recently, we have been working on a much simpler $O(n \log n)$ variant of our function-based coarsening approach. We have not yet obtained a full proof that this new variant provides the same theoretical guarantees, though we have obtained quality bound on the simpler case of threshold coarsening, where only the smallest features are repeatedly coarsened. Nonetheless, the algorithm has been implemented and preliminary experiments have shown that it is efficient in practice while generating high quality coarsening sequences. Details of these experiments and the algorithm can

be found in Talmor [20]. Future work includes a companion paper on practical mesh coarsening algorithms.

Acknowledgments We would like to thank Tony Chan, Jim Ruppert, and Noel Walkington for helpful discussions at different stages of this research. We would like to thank the referees for their very helpful comments and suggestions.

References

- [1] I. Babuška and A.K. Aziz. On the angle condition in the finite element method. *SIAM Journal on Numerical Analysis*, 13(2):214–226, 1976.
- [2] Randolph E. Bank and Jinchao Xu. An algorithm for coarsening unstructured meshes. *Numerische Mathematik*, 73:1–36, 1996.
- [3] Timothy J. Barth. Randomized multigrid. *AIAA paper*, 95-0207, 1995.
- [4] Marshall Bern, David Eppstein, and John Gilbert. Provably good mesh generation. In *Proceedings of the 31st Annual IEEE Symposium on the Foundations of Computer Science*, pages 231–241, 1990.
- [5] J. H. Bramble, J. E. Pasciak, and J. Xu. Convergence estimates for product iterative methods with applications to domain decomposition. *Mathematics of Computation*, 55:1–22, 1991.
- [6] A. Brandt. Multi-level adaptive solutions to boundary value problems. *Mathematics of Computation*, 31:333–390, 1977.
- [7] T. F. Chan and B. Smith. Domain decomposition and multigrid algorithms for elliptic problems on unstructured meshes. *Contemporary Mathematics*, pages 1–14, 1993.
- [8] 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.
- [9] L. P. Chew. Guaranteed-quality triangular meshes. Technical Report TR 89-983, Department of Computer Science, Cornell, Ithaca, NY, 1989.
- [10] S. Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica*, 2:153–174, 1987.
- [11] Omar Ghattas and Xiaogang Li. Private communications, 1996.
- [12] H. Guillard. Node nested multigrid with Delaunay coarsening. Tech. report, INRIA, Valbonne, France, 1993.
- [13] Marie-Helene Lallemand, H. Steve, and Alain Dervieux. Unstructured multigriding by volume agglomeration: current status. Technical Report 1224, Inria-Sophia Antipolis, 1990.
- [14] D. J. Mavriplis. Multigrid techniques for unstructured meshes. Technical Report 95-27, ICASE, 1995.
- [15] D. J. Mavriplis and D. Venkatakrisnan. Agglomeration multigrid for viscous turbulent flows. *AIAA paper*, 1994.
- [16] Gary L. Miller, Dafna Talmor, Shang-Hua Teng, and Noel Walkington. A Delaunay based numerical method for three dimensions: generation, formulation and partition. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, 1995.
- [17] S. A. Mitchell and S. A. Vavasis. Quality mesh generation in three dimensions. In *8th ACM Symposium on Computational Geometry*, pages 212–221, 1992.
- [18] Carl F. Ollivier-Gooch. Upwind acceleration of an upwind Euler solver on unstructured meshes. *AIAA Journal*, 33(10):1822–1827, 1995.

- [19] J. Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Proceedings of the 4th ACM-SIAM Symposium on Discrete Algorithms*, pages 83–92, 1993.
- [20] Dafna Talmor. *Well-spaced points for numerical methods*. PhD thesis, Carnegie Mellon University, 1997.
- [21] 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.
- [22] D. Venkatakrishnan and D. J. Mavriplis. Agglomeration multigrid for the 3D Euler equation. *AIAA paper*, 1994.