

A Conforming Nonpropagating Adaptive Mesh Refinement Strategy for Quadrilateral and Hexahedral Meshes

Paul F. Fischer and Chaman S. Verma

October 17, 2005

1 Introduction

In this note, we present an approach to all-quad mesh h -type adaptation that provides the exponential scale reduction of quad-trees, avoids the need for nonconforming interfaces (i.e., hanging nodes), and yields guaranteed uniform mesh quality.

This approach arose out a need to perform external field computations where the solution is smoothly varying in a homogeneous far field and there are small scale geometric features that need to be resolved in the near-field region of interest. Such situations arise, for example, in the computation of flow over multicomponent airfoils or in electromagnetic (EM) scattering off of objects. In the case of external fluid flows, the resolution requirements (points per unit length) go to zero as $|\mathbf{x}| \rightarrow \infty$, while in the EM case the resolution requirements approach a fixed value of > 5 points per wavelength. The background meshes are thus qualitatively different, and this has some impact on propagation of mesh refinement, as we will discuss later.

In our particular case, we are considering the solution of EM scattering using quadrilateral spectral elements. In anticipation of wide spatial variation in resolution requirements, we are developing a p -type refinement strategy in which the polynomial order within each element may vary over typical ranges of $N=4$ to 16. One advantage of such an approach is that it provides more control over the resolution than one based on quad-refinement alone. For example, with p -refinement one may have an element of order 12 adjacent to another of order 16 whereas with standard quad-refinement, the minimal jump in resolution is 2-to-1. We also note that p -type refinement alone is not sufficient for the far-field problems that we are interested in. One needs the exponential variation in resolution that is provided by quad-refinements to rapidly scale the element size $|\mathbf{x}| \rightarrow \infty$. The governing constraints on our meshing problem are as follows. In all parts of the domain, we would like to have elements as large as possible, so that requisite resolution is obtained through increased approximation order. (This is the most efficient approach to wave propagation—in the extreme limit, one recovers a global spectral method.) Thus, in the far field, we should have large elements and, in the near field, we should have elements small enough to capture the geometry of interest, but no smaller. In each element, error estimates are used to determine the polynomial degree necessary to attain the prescribed level of accuracy.

We note that it is primarily a matter of computational convenience that drives our development of a conforming mesh refinement strategy. Conformity is not requisite for full $h - p$ refinement strategies, as has been shown in numerous prior works (e.g., [?, ?, ?]). However, by eliminating geometric nonconformity (hanging nodes), coding complexity is considerably reduced as one only

*Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, USA.

needs to address two cases: adjacent elements A and B are of the same order, or A and B are of different orders.

Another effective approach to mixed p and h refinement is to use triangular high-order elements. This is possible within the standard p -type finite element framework. However, the costs generally scale as $O(N^{2d})$ per element for each operator evaluation in \mathbb{R}^d , $d = 2$ or 3 , and the polynomial degree for p -type finite elements rarely exceeds $N = 5$ in practice. Other high-order triangles have been developed in recent years. There are *modal* spectral elements of Sherwin and Karniadakis, which are based on local tensor-product function expansions and have two-dimensional work and storage (W, S) complexities of $O(N^3, N^2)$ per element, per operator evaluation. Efficient high-order *nodal* bases for triangles have been developed based on Fekete points, which minimize the determinate of the associated Van der Monde matrix [?], and on generalizations of Gauss point distributions derived from equilibrium distributions of electrostatic charges [?]. Efficient realizations using nodal bases have been extensively developed by Hesthaven and Warburton [?] and, despite having $O(N^{2d})$ complexity, the nodal bases often outperform the modal triangles for polynomial degrees typically employed ($N=4-8$). (We note that the choice of modal or nodal bases should not be overemphasized, as it is a simple matter to switch between the equivalent representations during the course of a computation at the cost of a single operator evaluation.) Most recently, high-order simplicial bases yielding $O(N^d)$ complexity have been developed by the group of Schoeberl that are quite promising for extending generalized finite element methods to high degree [?]. Numerous examples with N as high as 15 have been considered by this group.

Without question, triangles and tetrahedral meshes offer significant flexibility in mesh generation. For configurations that can be appropriately meshed, however, there remain significant advantages for the use of quadrilateral and hexahedral spectral elements, which have $O(N^{d+1}, N^d)$ leading order work and storage complexities. We consider the arguments in \mathbb{R}^3 , as these are the essential ones for large-scale applications. First, the hex elements require significantly less indirect addressing than their simplicial counterparts. This results from the fact that nodes internal to a hex element can be accessed lexicographically and because there are only 6 faces that need to be exchanged when coupling elements together. By contrast, a five-tetrahedra decomposition of the unit cube results in 20 faces that need to be exchanged, each of half the size (i.e., $(N + 1)^2/2$). Second, the leading-order $O(N^4)$ work complexity associated with the hex-based spectral element method can be cast as cache and vector-efficient matrix-matrix product, which effectively masks much of the additional work overhead—the actual evaluation time is dominated by memory access costs, which are only $O(N^3)$, with a very small constant. (See [?] for details. We remark that the same efficiencies apply to the approach of Warburton and Hesthaven [?].)

Here, we present a meshing technique that significantly extends the applicability of quadrilateral-based spectral elements in \mathbb{R}^2 . We are currently investigating extensions of this approach to \mathbb{R}^3 that will be presented in a future article. We emphasize that our present focus is a “top down” approach. That is, we are considering only the ability to provide an efficient, high-quality, far- to near-field mesh transition. We are not considering the much broader (and more difficult) problem of adapting quad meshes to arbitrary geometries. Such approaches have been extensively developed by other authors (e.g. [?, ?, ?] CHAMAN.) Our goal here is to present a simple technique that puts the conforming all-quad mesh on an equal footing with classic two-dimensional quad-tree based AMR without having to introduce hanging nodes.

2 Methodology

Two of the central components of quad-refinement are scale invariance and exponential transition. The first component guarantees mesh quality—elements at one scale are the same shape as those at another and element quality is thus preserved as a range of scales is traversed. The second component allows for rapid refinement the vicinity of singularities and is achieved by using a 2-to-1 transition step. A first level of refinement doubles the local linear resolution. A second refinement in one of the new cells doubles this resolution, and so on, such that one has an exponential rise in resolution (in this case, 2^k) within a fixed region. This process is illustrated for the case of quad-refinement near a boundary in Fig. ?? a. Our conforming approach to the same problem derives from the observation that a hexagon, partitioned into two isosceles trapezoids, can be refined into eight isosceles trapezoids similar to the first and at half the scale. Moreover, adjacent refined hexagonal regions yield smaller hexagons comprising a similar trapezoid pair. We thus have a 2-to-1 self-similar scale reduction that yields the same essential properties of quad-based AMR.

It is traditional (though not necessary) to keep quad-tree resolution steps at a 2-to-1 ratio. Thus, if a once-refined domain, Ω^e , is to be refined a second time and is also adjacent to an unrefined domain, $\Omega^{e'}$, one insists that $\Omega^{e'}$ be refined first. This restriction applies to all adjacent element pairs and can result in refinement propagation. The propagation is generally limited, however, except in isolated pathological cases and generally poses no significant difficulty. A similar (necessary) refinement-propagation requirement arises in our conforming-quadrilateral strategy. The refinement procedure is illustrated in Fig. ??, which shows a sequence of successive refinements about the point X . We see that the region of refinement is bounded, provided that the far field elements are of bounded size.

The refinement rule for the hexagonal-based mesh is

- (i) If an element is marked for refinement, then the element adjacent to its longest edge must also be marked.
- (ii) An element may be refined only if ... (SEE JOSE THESIS) All new elements introduced by the refinement step must satisfy If an element is marked for refinement, then the element adjacent or if its longest edge is coincident with the domain boundary.

Applying these criteria recursively yields the refinement procedure. The refinement criterion is completely analogous to the Rivara bisection criterion for triangles [?] and all of the properties of that procedure are thus inherited by the present method. If the far field elements are of uniform size H , then propagation from a fixed region of interest will be bounded within the first layer of elements of size H surrounding the region.

To maintain conformity

... grab scale similarity text that s already written..

3 Introduction II

We are concerned with the development of computational meshes that employ as their atomic units quadrilateral elements in \mathbb{R}^2 and hexahedral elements in \mathbb{R}^3 . Such elements are commonly used in finite element methods and are the basis for the tensor-product efficiencies of spectral element methods.

A common difficulty with computational meshes in general, and quadrilateral meshes in particular, is that mesh refinement originating from isolated features of interest can propagate all the way to the computational domain boundary if one desires to retain a conforming mesh. Isolated

features of interest are common in “external” field problems where one is interested in the response of a device exposed to ambient conditions. Such problems arise, for example, in aeronautics, aeroacoustics, and electromagnetics. For quadrilateral meshes, the propagation problem is well-known and has led to widespread adoption of nonconforming quad-tree discretizations. For triangles, the Rivara edge-bisection refinement strategy provides quality refinement (i.e., bounded Jacobians) and typically terminates, but can, under certain pathological conditions, lead to refinement of every element in the mesh.

Presently, we are interested in developing adaptive mesh refinement (AMR) that preserves element conformity and employs only quadrilateral elements (or hexahedral in \mathbb{R}^3). As an example, we consider the spectral element meshes around a few isolated cylinder illustrated in Fig. ???. The mesh on the left shows the cylinders inserted into a hex lattice, while on the right the underlying lattice is based on squares. We make several remarks about the two mesh configurations. First, the number of sites where one might place a cylinder is limited by the lattice structure. Our intent is to initially place the cylinders (or other objects of interest) on the lattice and to then move them using mesh deformation to the final desired location. Second, both meshes give relatively uniform resolution throughout the computational domain. Computationally, the square lattice is advantageous because the Jacobian that transforms from the element reference domain, $\hat{\Omega} := [-1, 1]^2$, to the element Ω^e in physical space is constant, whereas it is bilinear for the trapezoids that form the cells in the hex case. In the region of interest, we note that the hex lattice essentially preserves the quality of the Jacobian because no new sharp angles are introduced, while the square lattice does incur some penalty resulting from the 45 degree angles in the mesh. As configured, the hex lattice can be transformed into a square (or rectangular) lattice by a simple transformation that straightens out the vertical dividing lines. One could thus recover the attractive feature of a constant Jacobian in the far-field. It is interesting to note that one can change the hex-mesh by simply “turning” a cell by 60 degrees. If we then apply the straightening, we recover elements that have vanishing Jacobians. So, while there is an isomorphism between our trapezoidal-based hex lattice and the square lattice, it only holds under restricted circumstances.

An interesting feature of the hex lattice is that it possesses a self-similarity property that is key to AMR. Figure ?? shows a sequence of refinements that focuses on a particular region of interest (ROI). Starting at the coarsest level, one identifies the hexagon containing the ROI. This hexagon is refined, producing eight trapezoids that are similar to the original two and thus yielding the fourfold increase in areal resolution that would be expected from a single localized bisection refinement step. To go to the next level, one needs to also refine the neighboring hexagon, as seen in Fig. ??c. It is clear, however, that the refinement does not propagate beyond the nearest neighbor. The self-similarity feature, which ensures bounded Jacobians, and nonpropagation, imply that the hexagonal mesh provides a sound basis for AMR. Figure ??d illustrates several levels of refinement around an array of cylinders.

With an understanding of how we can generate conforming quad-based AMR in \mathbb{R}^2 , it is natural to ask if there is an equivalent extension in \mathbb{R}^3 . Such analogies do not always exist. For example, a quad-refinement of an equilateral triangle yields four equilateral triangles with half the linear dimension of the original. By recursion, this implies that equilateral triangles can fill space in \mathbb{R}^2 . In \mathbb{R}^3 , however, an oct-refine of a regular tetrahedron does not yield eight self-similar tetrahedra, which explains why regular tetrahedra cannot fill space in three dimensions. What is needed in the present case is the three-dimensional equivalent to the hexagon. This is given by the rhombic dodecahedron (RD), which is the Voronoi cell associated with a face-centered cubic (fcc) lattice packing of spheres in \mathbb{R}^3 [?]. Like the hexagon, the RD fills space. Fortunately, the faces of the RD are all quadrilaterals, which makes it a good candidate for 3D AMR based on hexahedral elements. The RD can be decomposed in to four hexahedral elements, each is a regular rhomboid

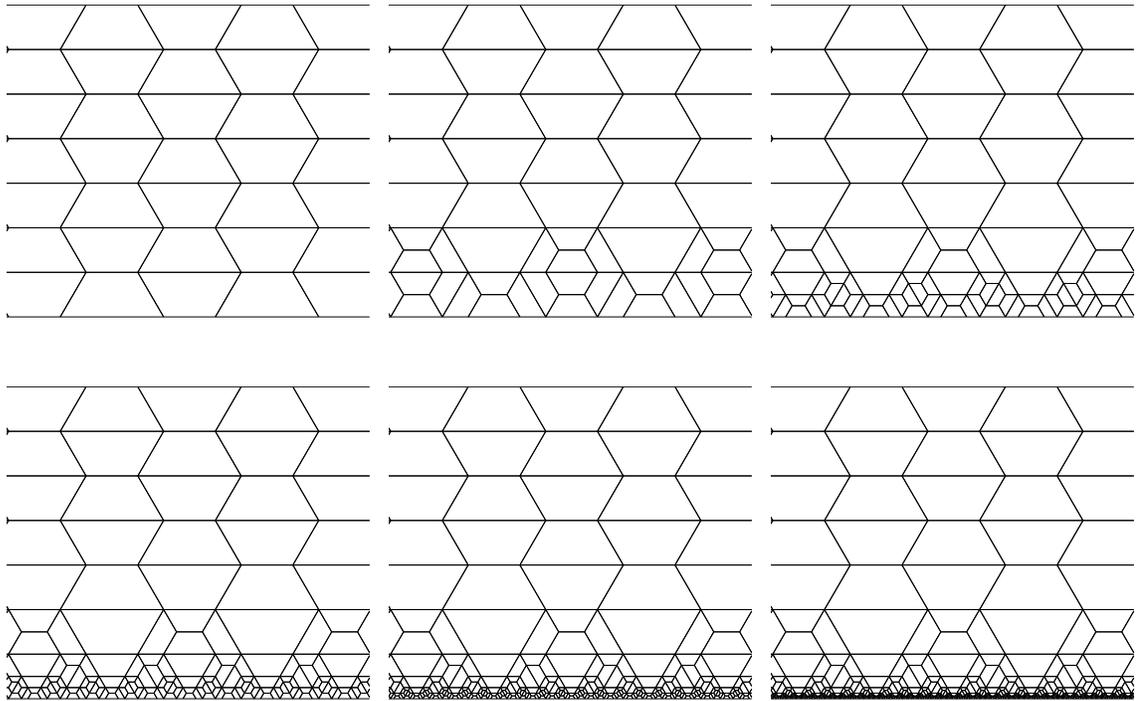


Figure 1: Hexagonal-based mesh showing six generations of refinement along a boundary. The mesh is conforming, the elements are uniform in shape and quality and the refinement is confined within the first layer of hexagons at the coarsest level.

with equal edge lengths. Under local refinement, each of the 12 faces is extended inwards and a new rhomboid is placed at the center,

So far, we have only addressed AMR from the top-down (or, more precisely, coarse-to-fine) perspective. We are not attempting to resolve the general problem of all-quad meshing but simply to address the question of AMR requirements for external field problems. We are assuming that the applications specialist will have techniques that are specific to the geometry at hand. For the case of cylinders (and) spheres, however, this approach is immediately applicable.

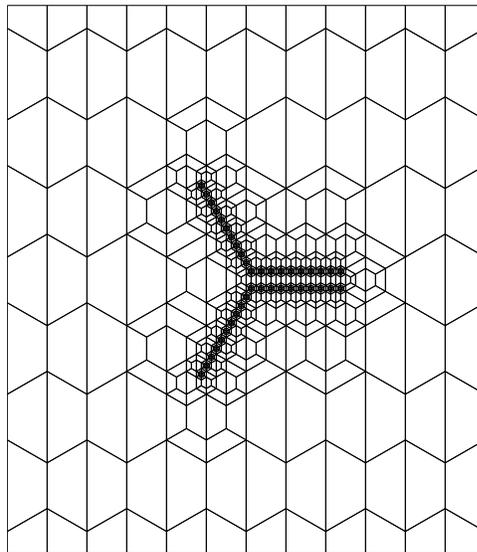


Figure 2: Hexagonal-based mesh refinement in the vicinity of a nanodevice. Here, nanoparticles have been sited at 4th-generation lattice points. Mesh motion will be employed to adjust the particles to more general positions.