Diamond-based Models for Scientific Visualization
Summary of Research

Kenneth Weiss*
Department of Computer Science
University of Maryland, College Park

Abstract

Mathematical and scientific computing problems are often approached using a divide and conquer paradigm in which a large problem domain is solved by combining results from smaller subdomains. In this very broad context, we focus on nested spatial decompositions of regularly sampled domains obtained by the *regular simplex bisection* scheme. Specifically, we utilize clusters of such simplices, referred to as *diamonds*, that must be refined concurrently to guarantee crack-free domain decompositions.

This work investigates several diamond-based models for scientific visualization. First, we generalize the notion of diamonds from 2D and 3D to arbitrary dimension, and propose an implicit encoding for all geometric and hierarchical relationships. Second, we propose the *supercube* as a high-level primitive for diamond hierarchies. We use supercubes to compactly encode incomplete scalar fields as well as adaptive simplicial meshes extracted from the hierarchy. Finally, we introduce *isodiamond hierarchies*, a general framework for spatial access structures on a hierarchy of diamonds that decouples the implicit relationships of the hierarchy from the underlying scalar field.

1 Introduction

One of the fundamental problems in computer graphics, scientific visualization, geographic data processing, and shape analysis and understanding is to deal with the huge amount of data that describe the objects of interest. As examples, let us consider terrain modeling and visualization, where we deal with millions to billions of samples describing elevations, or the modeling of a volume data set for visualization and analysis, where we need to deal with huge point data sets or very large meshes describing isosurfaces.

A diverse set of approaches utilizes a hierarchical organization of the field domain to describe subsections of these objects. Among these, a popular class of such approaches is based on the *simplex bisection* operation [8, 6], in which a simplex (i.e., a line segment, triangle, tetrahedron, or their higher dimensional analogues) is split along one of its edges into two simplices covering the same domain (see Figure 1). Although individual bisection operations can introduce cracks into a mesh and, thus, discontinuities in functions

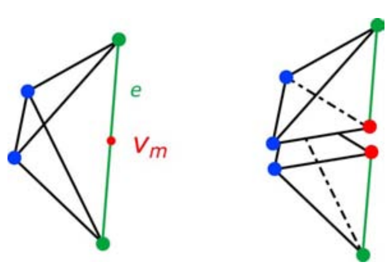

Figure 1: Bisection of a 3-simplex (tetrahedron) along an edge e. Image from [13].

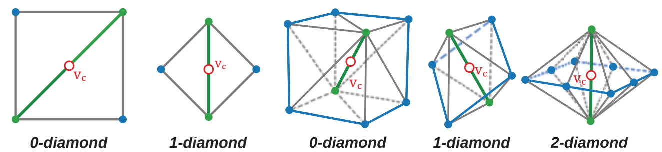

(a) 2D diamonds (b) 3D diamonds

Figure 2: The two classes of diamonds in 2D (a) and the three classes of diamonds in 3D (b). Image from [17].

defined on the mesh, crack-free, or *conforming*, refinements can be applied by concurrently bisecting all simplices sharing the same bisection edge. This group of simplices is often referred to as a *diamond* [4, 5, 13] (see Figure 2). The popularity of these approaches stems from their dimension–independent definition as well as their high degree of adaptability to application-specific constraints.

The primary distinction among all proposed solutions lies in the choice of modeling primitive for such meshes. *Simplex-based* approaches focus on the individual elements of the mesh by modeling the mesh as a collection of simplices. Conforming refinements are typically implemented through an explicit enumeration of all simplices sharing a bisection edge. Although elegant solutions have been proposed for meshes of arbitrary dimension [6, 1], these approaches have not considered the number of such simplices that must be bisected to maintain conforming meshes. In contrast, *diamond-based* approaches focus on the conforming refinements by modeling such meshes as a collection of diamonds. This has led to the development of efficient multiresolution representations for 2D and 3D meshes, but a general understanding of diamonds in higher dimensions has remained an open problem.

2 Contributions

My dissertation research proposes several significant contributions to diamond-based scientific visualization. We first focus on the properties of the domain decomposition induced by simplex bisection, and generalize the notion of diamonds to higher dimensions [13]. In [17, 18], we provide a comprehensive survey that consolidates the literature on simplex-based and diamond-based models as well as their dimension-specific applications to scientific visualization, computer graphics and spatial analysis.

Since many operations pertain to only a portion of the regularly sampled domain, we next consider the efficient encoding of information associated with sparse, coherent subsets of a regular grid [12, 14]. This provides the first general solution to the long-standing open problem of compactly representing *incomplete* scalar fields in a multiresolution context.

Finally, we introduce a general framework, that we call *Isodiamond Hierarchies* [11, 16], for decoupling the implicit domain decomposition induced by diamonds from values defined on the field. This enables the use of the hierarchy as a general spatial access structure.

Our recent work has focused on applications of the implicit hier-

*kweiss@cs.umd.edu

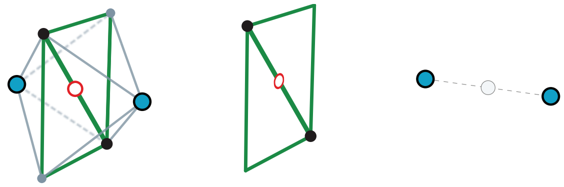

(a) (b) (c)

Figure 3: A 1-diamond in 3D (a) can be decomposed into a triangulated 2-cube (b) and the boundary of a triangulated 1-cube (c). In general, an i-diamond in d dimensions can be decomposed into a triangulated $(d-i)$-cube and the boundary of a triangulated i-cube. Image from [17].

archy, such as the relationship between the resolution of a dataset, and the *distortion* (a generalized notion of discrete curvature) [7] that the scalar field induces on its domain [19, 2]. We have also been interested in the relationship between adaptive simplicial decompositions and those based on hypercubes, such as quadtrees, octrees and their higher dimensional counterparts [15].

2.1 Diamond Hierarchies of Arbitrary Dimension

In [13], we formalize the notion of diamonds to arbitrary dimensions in terms of their decomposition into a pair of triangulated hypercubes (see Figure 3). This enables us to derive the first closed-form equations for the number of vertices and simplices involved in conforming refinements on such meshes and to define a compact pointerless data structure for encoding diamonds of arbitrary dimension.

In particular, we prove that the number of simplices associated with each diamond is *factorial* in the dimension d of the domain, however, due to the regularity of the refinement operation, we can group them into a *linear* number of subclusters that are generated simultaneously. This implies that conforming refinements in simplex-based representations require $O(d!)$ time and space. In contrast, our proposed diamond-based representation requires only $O(d)$ time and space, yielding a significant reduction. The proposed data structure has the potential to make working with higher dimensional domains tractable.

2.2 Representing Incomplete Hierarchies

Although most implementations of nested simplicial meshes attain their efficiency by encoding all samples within a hypercubic domain, many operations on these meshes apply to only a sparse subset of the data, for example, when regions of the domain are oversampled, or unavailable. Alternatively, spatial queries on such datasets, such as isosurface queries or view–dependent visualization, extract large meshes, which can require significant amounts of storage space.

In [12, 14], we introduce a high-level primitive based on the clustering of diamonds into *supercubes*, which tile each level of resolution of the hierarchy (see Figure 4). Each diamond uniquely maps to an edge within a single supercube (see Figure 5).

We propose the use of supercubes to associate information with coherent subsets of multiresolution terrain and volumetric datasets, respectively. When encoding an incomplete mutiresolution scalar field, our supercube-based encoding incurs less than a byte of overhead per retained sample, while still retaining random access. Additionally, our supercube-based encoding of extracted tetrahedral meshes requires only a single byte of overhead per tetrahedron, half as much as a diamond-based encoding.

A particularly important application of supercube-based models is in the field of Geographic Information Systems (GIS), where reg-

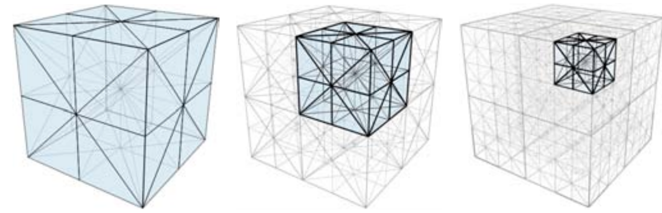

(a) Level ℓ (b) Level $\ell+1$ (c) Level $\ell+2$

Figure 4: Supercubes are structured sets of edges tiling each level of resolution within a diamond hierarchy. Three consecutive levels of resolution covering the same three-dimensional domain are illustrated, containing 1, 8 and 64 supercubes, respectively. Images from [17].

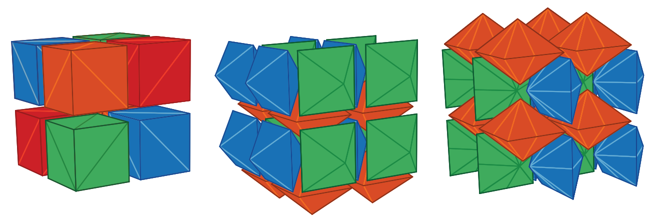

Figure 5: The one to one correspondence between diamonds and supercube edges enables the use of supercubes to associate information with a coherent subset of the diamonds in the hierarchy. Each 3D supercube indexes 56 diamonds – eight *0-diamonds* (left), twenty four *1-diamonds* (middle) and twenty four *2-diamonds* (right).

ular grids are the most common format for terrain datasets. Supercubes provide a solution for a long-standing problem of representing subsets of a regular grid. For example, elevation data for rough terrain and coastlines are required to be sampled at a high resolution, while flatter regions, especially those covering large bodies of water, do not require such high sampling resolution. This is especially relevant for global datasets since approximately 70% of the earth's surface is covered by water (see Figure 6).

2.3 Isodiamond Hierarchies

An important means of analyzing medical and scientific volumetric datasets is through the surfaces of constant field value, known as *isosurfaces*, or through the volume enclosed between two such surfaces, known as *interval volumes*. However, these meshes can require hundreds of millions of triangles and tetrahedra at full resolution. We use our isodiamond hierarchy framework [11, 16] to compactly define multiresolution representations for individual isosurfaces or interval volumes extracted from a multiresolution scalar field.

In contrast to previous approaches, these extracted isosurfaces are guaranteed to be free of self-intersections and have the same representational power as the original multiresolution scalar field in representing the selected isosurface or interval volume while requiring an order of magnitude less space and supporting faster mesh extraction. This is accomplished by only extracting tetrahedra that intersect the isosurface or interval volume (compare the left and right images in Figure 7). Figure 8a illustrates an isosurface extracted using a box-shaped Region of Interest (ROI) query around the head of the model. Figure 8b illustrates an interval volume extracted using an approximation error query, where all cells with an error higher than one percent are refined.

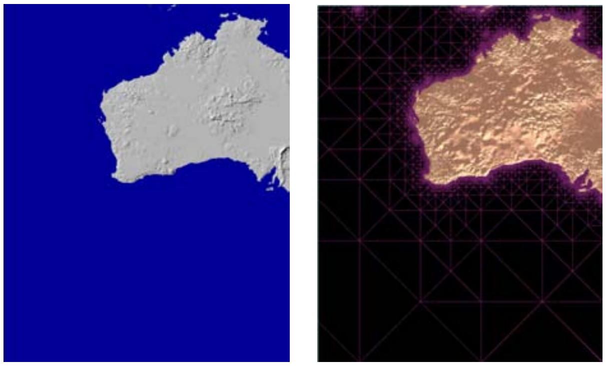

(a) Complete terrain tile (b) Supercube-based terrain

Figure 6: Terrains covering large flat regions such as oceans are oversampled by a regular grid (a). A zero error supercube-based representation of this terrain (b) requires less than $1/6$ of the samples from the original dataset. Image (a) courtesy of USGS [9]. Image (b) from [12].

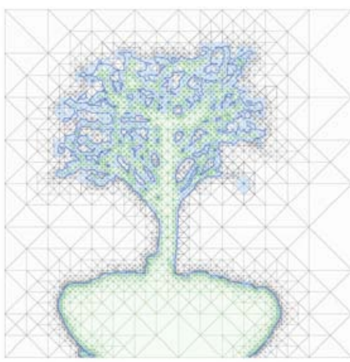

(a) (b)

Figure 7: Traditionally, extracted meshes must cover the entire domain to ensure they are conforming (left). Our isodiamond framework enables the extraction of conforming diamond meshes that cover only the portions of the domain intersecting the desired isosurface or interval volume (right). Image from [16].

3 Ongoing and Future Work

The convergence of these three developments opens the door to the multiresolution analysis and visualization of scalar fields of dimension greater than three. In particular, we intend to explore the feasibility of using a four-dimensional hierarchy of diamonds to analyze and visualize four-dimensional scalar fields representing time-varying volumetric datasets. The latter can be modeled as a four dimensional scalar field by treating the temporal and spatial dimensions homogeneously [10]. Using our supercube-based encoding of the field, we expect to reduce the storage overhead of this representation and of its extracted meshes.

In our ongoing work [19, 2], we have been concerned with the connection between the resolution of an extracted mesh and the *distortion* (a higher-dimensional generalization of curvature) that the retained field values induce on its domain [7]. Preliminary results indicate that a distortion-guided extraction directs the mesh resolution towards the salient features of the field. This enables accurate analysis of complex datasets using significantly fewer resources.

On another front, we are interested in the relationship between adaptive domain decompositions based on hypercubes, such as

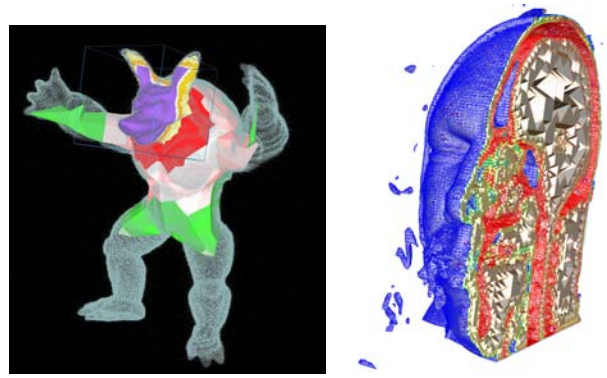

(a) Armadillo Isosurface (b) Head Interval Volume

Figure 8: (a) Isosurface extracted from 512^3 Armadillo dataset using a cubic region of interest focused around the head of the model. Triangle colors indicate the level of resolution and the blue points indicate vertices in the model at full resolution. (b) Interval volume extracted from the 256^3 Visible Human Head dataset with a maximum approximation error of one percent. The mesh is clipped along the median plane to show the internal tetrahedra. Images from [16].

quadtrees, octrees and their higher dimensional counterparts, and nested simplicial meshes generated through a regular simplex bisection scheme. Specifically, we are interested in how simplicial decompositions of such hypercubic meshes using simplex bisection-based rules relate to simplicial meshes generated directly from the hierarchy of diamonds [15]. It has been previously been observed [3] that in the 2D case, the class of triangulations obtainable from a *restricted* quadtree[1] using these rules is a proper subset of those obtainable using a simplex bisection approach. Thus, although hypercube-based decompositions often require fewer cells to satisfy a given query the diamond meshes have a higher expressive power, and are more adaptable.

Due to the highly parallel nature of the domain decomposition, we have also been investigating GPU algorithms for generating, querying, processing and rendering adaptive simplicial meshes extracted from a supercube-based diamond hierarchy [20].

4 Concluding Remarks

Our dimension–independent formulation for diamonds has already yielded significant results in modeling terrain and volumetric datasets as well as in their extracted meshes. Since such spatial decompositions are fundamental in understanding many disparate areas of mathematical and scientific computing and visualization, we anticipate many new applications of these techniques in the coming years.

References

[1] F. Atalay and D. Mount. Pointerless implementation of hierarchical simplicial meshes and efficient neighbor finding in arbitrary dimensions. *International Journal of Computational Geometry and Applications*, 17(6):595–631, 2007.

[2] L. De Floriani, F. Iuricich, P. Magillo, M. Mesmoudi, and K. Weiss. Discrete distortion for 3D data analysis. In L. Linsen, H. Hagen, and B. Hamann, editors, *Visualization in Medicine and Life Sciences (VMLS)*, Mathematics and Visualization. Springer Berlin Heidelberg, 2011.

[3] L. De Floriani and P. Magillo. Multiresolution mesh representation: models and data structures. In M. Floater, A. Iske, and E. Quak,

[1] In a restricted quadtree, also referred to as a balanced quadtree, adjacent squares can differ by at most one level of resolution.

editors, *Principles of Multi-resolution Geometric Modeling*, Lecture Notes in Mathematics, pages 364–418, Berlin, 2002. Springer Verlag.

[4] M. Duchaineau, M. Wolinsky, D. E. Sigeti, M. C. Miller, C. Aldrich, and M. B. Mineev-Weinstein. ROAMing terrain: real-time optimally adapting meshes. In R. Yagel and H. Hagen, editors, *Proceedings IEEE Visualization*, pages 81–88, Phoenix, AZ, October 1997. IEEE Computer Society.

[5] B. Gregorski, M. Duchaineau, P. Lindstrom, V. Pascucci, and K. Joy. Interactive view-dependent rendering of large isosurfaces. In *Proceedings IEEE Visualization*, pages 475–484. IEEE Computer Society Washington, DC, USA, October 2002.

[6] J. M. Maubach. Local bisection refinement for n-simplicial grids generated by reflection. *SIAM Journal on Scientific Computing*, 16(1):210–227, January 1995.

[7] M. M. Mesmoudi, L. De Floriani, and U. Port. Discrete distortion in triangulated 3-manifolds. *Computer Graphics Forum*, 27(5):1333–1340, 2008.

[8] M. Rivara. Local modification of meshes for adaptive and/or multigrid finite-element methods. *Journal of Computational and Applied Mathematics*, 36(1):79–89, 1991.

[9] U. G. Survey. Global 30 arc second elevation data. `http://edc.usgs.gov/products/elevation/gtopo30/gtopo30.html`.

[10] K. Weiss and L. De Floriani. Modeling and visualization approaches for time-varying volumetric data. In G. Bebis, R. Boyle, B. Parvin, D. Koracin, P. Remagnino, F. Porikli, J. Peters, J. Klosowski, L. Arns, Y. Chun, T. Rhyne, and L. Monroe, editors, *Advances in Visual Computing*, volume 5359 of *Lecture Notes in Computer Science*, pages 1000–1010. Springer, 2008.

[11] K. Weiss and L. De Floriani. Multiresolution interval volume meshes. In H.-C. Hege, D. Laidlaw, R. Pajarola, and O. Staadt, editors, *IEEE/EG Symposium on Volume and Point-Based Graphics*, pages 65–72, Los Angeles, California, USA, 2008. Eurographics Association.

[12] K. Weiss and L. De Floriani. Sparse terrain pyramids. In *Proceedings ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 115–124, New York, NY, USA, 2008. ACM.

[13] K. Weiss and L. De Floriani. Diamond hierarchies of arbitrary dimension. *Computer Graphics Forum (Proceedings SGP 2009)*, 28(5):1289–1300, 2009.

[14] K. Weiss and L. De Floriani. Supercubes: A high-level primitive for diamond hierarchies. *IEEE Transactions on Visualization and Computer Graphics (Proceedings IEEE Visualization 2009)*, 15(6):1603–1610, November-December 2009.

[15] K. Weiss and L. De Floriani. Bisection-based triangulations of nested hypercubic meshes. In S. Shontz, editor, *Proceedings 19th International Meshing Roundtable*, pages 315–333, Chattanooga, Tennessee, October 3–6 2010.

[16] K. Weiss and L. De Floriani. Isodiamond hierarchies: An efficient multiresolution representation for isosurfaces and interval volumes. *IEEE Transactions on Visualization and Computer Graphics*, 16(4):583 – 598, July-Aug. 2010.

[17] K. Weiss and L. De Floriani. Simplex and diamond hierarchies: Models and applications. In H. Hauser and E. Reinhard, editors, *EG 2010 - State of the Art Reports*, pages 113–136, Norrköping, Sweden, 2010. Eurographics Association.

[18] K. Weiss and L. De Floriani. Simplex and diamond hierarchies: Models and applications. *Accepted to Computer Graphics Forum*, 2011.

[19] K. Weiss, M. Mesmoudi, and L. De Floriani. Multiresolution analysis of 3D images based on discrete distortion. In *International Conference on Pattern Recognition (ICPR)*, pages 4093–4096, Istanbul, Turkey, August 2010. IEEE Computer Society.

[20] M. Yalçın, K. Weiss, and L. De Floriani. GPU algorithms for diamond-based multiresolution terrain processing. In *Eurographics Symposium on Parallel Graphics and Visualization*, Bangor, Wales, April 10–11 2011.