A1
Anisotropic Voronoi Diagrams and Guaranteed-Quality Anisotropic Mesh Generation

Francois Labelle Jonathan Richard Shewchuk

Abstract
We introduce *anisotropic Voronoi diagrams*, a generalization of multiplicatively weighted Voronoi diagrams suitable for generating guaranteed-quality meshes of domains in which long, skinny triangles are required, and where the desired anisotropy varies over the domain. We discuss properties of anisotropic Voronoi diagrams of arbitrary dimensionality—most notably circumstances in which a site can see its entire Voronoi cell. In two dimensions, the anisotropic Voronoi diagram dualizes to a triangulation under these same circumstances. We use these properties to develop an algorithm for anisotropic triangular mesh generation in which no triangle has an angle smaller than 20 degree, as measured from the skewed perspective of any point in the triangle.

A2
Laplacian Surface Editing

Olga Sorkine Daniel Cohen-Or Yaron Lipman Marc Alexa

Abstract
Surface editing operations commonly require geometric details of the surface to be preserved as much as possible. We argue that geometric detail is an intrinsic property of a surface and that, consequently, surface editing is best performed by operating over an intrinsic surface representation. We provide such a representation of a surface, based on the Laplacian of the mesh, by encoding each vertex relative to its neighborhood. The Laplacian of the mesh is enhanced to be invariant to locally linearized rigid transformations and scaling. Based on this Laplacian representation, we develop useful editing operations: interactive free-form deformation in a region of interest based on the transformation of a handle, transfer and mixing of geometric details between two surfaces, and transplanting of a partial surface mesh onto another surface. The main computation involved in all operations is the solution of a sparse linear system, which can be done at interactive rates. We demonstrate the effectiveness of our approach in several examples, showing that the editing operations change the shape while respecting the structural geometric detail.

A3
Fast Collision Detection between Massive Models using Dynamic Simplification

Sung-Eui Yoon Brian Salomon Ming Lin Dinesh Manocha

Abstract
We present a novel approach for collision detection between large models composed of tens of millions of polygons. Each model is represented as a clustered hierarchy of progressive meshes (CHPM). The CHPM is a dual hierarchy of the original model; it serves both as a multiresolution representation of the original model, as well as a bounding volume hierarchy. We use the cluster hierarchy of a CHPM to perform coarse-grained selective refinement and the progressive meshes for fine-grained local refinement. We present a novel conservative error metric to perform collision queries based on the multiresolution representation. We use this error metric to perform dynamic simplification for collision detection. Our approach is conservative in that it may overestimate the set of colliding regions, but never misses any collisions. Furthermore, we are able to generate these hierarchies and perform collision queries using out-of-core techniques on all triangulated models. We have applied our algorithm to perform conservative collision detection between massive CAD and scanned models, consisting of millions of triangles at interactive rates on a commodity PC.

A4
Filling Holes in Meshes

Peter Liepa

Abstract
We describe a method for filling holes in unstructured triangular meshes. The resulting patching meshes interpolate the shape and density of the surrounding mesh. Our methods work with arbitrary holes in oriented connected manifold meshes. The steps in filling a hole include boundary identification, hole triangulation, refinement, and fairing.

A5
Domain Decomposition for Multiresolution Analysis

Ioana M. Boier-Martin

Abstract
This paper describes a method for converting an arbitrary mesh with irregular connectivity to a semi-regula rmultiresolution representation. A shape image encoding geometric and differential properties of the input model is computed. Standard image processing operations lead to an initial decomposition of the model that conforms to its salient features. A triangulation step performed on the resulting partition in image space, followed by resampling and multiresolution analysis in object space, complete the procedure. The conversion technique is automatic, takes into account surface properties for deriving a base domain, and is computationally ef_cient as the bulk of the processing is carried out in image space. Besides domain decomposition, our image-based approach to handling geometry may be used in the context of related applications, including model simpli_cation, remeshing, and wireframe generation.

A6
Approximated Centroidal Voronoi Diagrams for Uniform Polygonal Mesh Coarsening

Sébastien Valette and Jean-Marc Chassery

Abstract
We present a novel clustering algorithm for polygonal meshes which approximates a Centroidal Voronoi Diagram construction. The clustering provides an ef_cient way to construct uniform tessellations, and therefore leads to uniform coarsening of polygonal meshes, when the output triangulation has many fewer elements than the input mesh. The mesh topology is also simpli_ed by the clustering algorithm. Based on a mathematical framework, our algorithm is easy to implement, and has low memory requirements. We demonstrate the ef_ciency of the proposed scheme by processing several reference meshes having up to 1 million triangles and very high genus within a few minutes on a low-end computer.

A7
Single-Strip Triangulation of Manifolds with Arbitrary Topology

M. Gopi and David Eppstein

Abstract
Triangle strips have been widely used for efficient rendering. It is NP-complete to test whether a given triangulated model can be represented as a single triangle strip, so many heuristics have been proposed to partition models into few long strips. In this paper, we present a new algorithm for creating a single triangle loop or strip from a triangulated model. Our method applies a dual graph matching algorithm to partition the mesh into cycles, and then merges pairs of cycles by splitting adjacent triangles when necessary. New vertices are introduced at midpoints of

edges and the new triangles thus formed are coplanar with their parent triangles, hence the visual fidelity of the geometry is not changed. We prove that the increase in the number of triangles due to this splitting is 50% in the worst case, however for all models we tested the increase was less than 2%. We also prove tight bounds on the number of triangles needed for a single-strip representation of a model with holes on its boundary. Our strips can be used not only for efficient rendering, but also for other applications including the generation of space filling curves on a manifold of any arbitrary topology.

A8
Star Splaying: An Algorithm for Repairing Delaunay Triangulations and Convex Hulls

Jonathan Richard Shewchuk

Abstract
Star splaying is a general-dimensional algorithm that takes as input a triangulation or an approximation of a convex hull, and produces the Delaunay triangulation, weighted Delaunay triangulation, or convex hull of the vertices in the input. If the input is "nearly Delaunay" or "nearly convex" in a certain sense quantified herein, and it is sparse (i.e. each input vertex adjoins only a constant number of edges), star splaying runs in time linear in the number of vertices. Thus, star splaying can be a fast first step in repairing a high-quality finite element mesh that has lost the Delaunay property after its vertices have moved in response to simulated physical forces. Star splaying is akin to Lawson's edge flip algorithm for converting a triangulation to a Delaunay triangulation, but it works in any dimensionality.

A9
The Power Crust

Nina Amenta Sunghee Choi Ravi Krishna Kolluri

Abstract
The *power crust* is a construction which takes a sample of points from the surface of a three-dimensional object and produces a surface mesh and an approximate medial axis. The approach is to first approximate the medial axis transform (MAT) of the object. We then use an inverse transform to produce the surface representation from the MAT. This idea leads to a simple algorithm with theoretical guarantees comparable to those of other surface reconstruction and medial axis approximation algorithms. It also comes with a guarantee that does not depend in any way on the quality of the input point sample. *Any* input gives an output surface which is the 'watertight' boundary of a three-dimensional polyhedral solid: the solid described by the approximate MAT. This unconditional guarantee makes the algorithm quite robust and eliminates the polygonalization, hole-filling or manifold extraction post-processing steps required in previous surface reconstruction algorithms. In this paper, we use the theory to develop a power crust implementation which is indeed robust for realistic and even difficult samples. We describe the careful design of a key subroutine which labels parts of the MAT as inside or outside of the object, easy in theory but non-trivial in practice. We find that we can handle areas in which the input sampling is scanty or noisy by simply discarding the unreliable parts of the MAT approximation. We demonstrate good empirical results on inputs including models with sharp corners, sparse and unevenly distributed point samples, holes, and noise, both natural and synthetic.

A10
Image Reconstruction Using Data-Dependent Triangulation
Xiaohua Yu, Bryan S. Morse, Thomas W. Sederberg

Abstract
This article presents an algorithm for digital image reconstruction based on data dependent triangulation. The method enables resampling operations, such as magnification, to create images that are generally of higher quality than those obtained through traditional bilinear or bicubic spline reconstruction. The article presents a new cost function that is invariant to the intensity range. Also, an improved optimization algorithm is presented for data-dependent triangulation -based on a strategy called "look-ahead"-that performs better than Lawson 's edge-swapping algorithm.

A11
Simplification and Improvement of Tetrahedral Models for Simulation
B. Cutler, J. Dorsey, and L. McMillan

Abstract
Most 3D mesh generation techniques require simpli_cation and mesh improvement stages to prepare a tetrahedral model for ef_cient simulation. We have developed an algorithm that both reduces the number of tetrahedra in the model to permit interactive manipulation and removes the most poorly shaped tetrahedra to allow for stable physical simulations such as the _nite element method. The initial tetrahedral model may be composed of several different materials representing internal structures. Our approach targets the elimination of poorly-shaped elements while simplifying the model using edge collapses and other mesh operations, such as vertex smoothing, tetrahedral swaps, and vertex addition. We present the results of our algorithm on a variety of inputs, including models with more than a million tetrahedra. In practice, our algorithm reliably reduces meshes to contain only tetrahedra that meet speci_ed shape requirements, such as the minimum solid angle.

A12
Tetrahedron Based, Least Squares, Progressive Volume Models With Application to Freehand Ultrasound Data
Tom Roxborough, Gregory M. Nielson

Abstract
In this paper we present a new method for the modeling of freehand collected three-dimensional ultrasound data. The model is piece-wise linear and based upon progressive tetrahedral domains created by a subdivision scheme which splits a tetrahedron on its longest edge and guarantees a valid tetrahedrization. Least squares error is used to characterize the model and an effective iterative technique is used to compute the values of the model at the vertices of the tetrahedral grid. Since the subdivision strategy is adaptive, the complexity of the model conforms to the complexity of the data leading to an extremely efficient and highly compressed volume model. The model is evaluated in real time using piece-wise linear interpolation, and gives a medical professional the chance to see images which would not be possible using conventional ultrasound techniques.