

CGALmesh: a Generic Framework for Delaunay Mesh Generation

CLEMENT JAMIN, Inria and Université Lyon 1, LIRIS, UMR5205

PIERRE ALLIEZ, Inria

MARIETTE YVINEC, Inria

JEAN-DANIEL BOISSONNAT, Inria

CGALmesh is the mesh generation software package of the Computational Geometry Algorithm Library (CGAL). It generates isotropic simplicial meshes – surface triangular meshes or volume tetrahedral meshes – from input surfaces, 3D domains as well as 3D multi-domains, with or without sharp features. The underlying meshing algorithm relies on restricted Delaunay triangulations to approximate domains and surfaces, and on Delaunay refinement to ensure both approximation accuracy and mesh quality. CGALmesh provides guarantees on approximation quality as well as on the size and shape of the mesh elements. It provides four optional mesh optimization algorithms to further improve the mesh quality. A distinctive property of CGALmesh is its high flexibility with respect to the input domain representation. Such a flexibility is achieved through a careful software design, gathering into a single abstract concept, denoted by the oracle, all required interface features between the meshing engine and the input domain. We already provide oracles for domains defined by polyhedral and implicit surfaces.

Categories and Subject Descriptors: I.3.5 [Computational Geometry and Object Modeling]: Geometric algorithms, languages, and systems

General Terms: Algorithms

Additional Key Words and Phrases: Mesh generation, Delaunay refinement, mesh optimization, isosurface extraction.

ACM Reference Format:

Clément Jamin, Pierre Alliez, Mariette Yvenc and Jean-Daniel Boissonnat. 2014. CGALmesh: a Generic Framework for Delaunay Mesh Generation. *ACM Trans. Math. Softw.* 1, 1, Article 1 (June 2014), 26 pages. DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Surface triangular meshes and 3D tetrahedral meshes play a central role in computational engineering for simulating and visualizing physical phenomena. They are also commonplace for modeling and animating complex scenes in special effects or multimedia applications. For rendering, surface meshes must provide a good approximation of object boundaries. For numerical simulation, 3D volume meshes must fulfill additional constraints over the shape, orientation and size of the elements.

Meshes are used not only for rendering and simulation, but also as means to solve other problems. Example of such problems include image segmentation for medical applications [Pons and Boissonnat 2007] and spatio-temporal coherent segmentation of time-varying data sets such as beating hearts. As for images, meshes are now com-

Author's addresses: Inria Sophia Antipolis - Méditerranée 2004 route des Lucioles BP69 06902 Sophia Antipolis, France; e-mail: firstname.lastname@inria.fr

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 0098-3500/2014/06-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>




Fig. 1. Mesh generation from 3D images. The usual pipeline applies marching cubes to extract isosurfaces, simplifies, remeshes, generates 3D meshes separately for each domain and finally merges the resulting 3D sub-meshes. Our mesh generation technique generates a multi-domain 3D mesh at once from 3D images. The output mesh is by construction conforming at all domain interfaces, and the whole process does not require trial-and-error processes.

monly used as central data structures for a variety of algorithms. This poses new challenges in terms of complexity and genericity of the mesh generation algorithms.

The description of a domain can take many forms, ranging from implicit surfaces to triangular meshes through 3D images and point sets. This variety of inputs is a notoriously enduring challenge for mesh generation algorithms. Most techniques are designed to handle a single type of domain description and are not suited to other descriptions. Choosing the proper algorithm often reduces to selecting the one which can deal with the input description. Such choice may be too limited, especially when restricting to open-source software solutions. For this reason, each type of input corresponds to a specific methodology: meshing from implicit surfaces is referred to as *isosurface extraction*, meshing from point sets as *reconstruction*, meshing from trian-

gular meshes as *remeshing*, to name a few. These methodologies correspond to different research areas, each with its own challenges.

1.1. Related Work

Mesh generation is an active research field. The mesh generation algorithms divide into the following meshing strategies: advancing front [Schöberl 1997; Marcum 1998], lattice-based methods [Lorensen and Cline 1998; Labelle and Shewchuk 2007; Chentanez et al. 2007], and Delaunay refinement. The latter can be further classified into two classes of algorithms. The first class of Delaunay refinement algorithms requires to mesh or remesh the input domain boundary (often a trial and error process) before meshing the input domain while preserving the boundary mesh [Dey and Levine 2009; Dey et al. 2010; Dey et al. 2012; Si 2006]. The second class of Delaunay refinement algorithms, our choice for CGALmesh, meshes the input surface and volume altogether in a unified manner, with the ability to mesh multi-domains and curved surfaces while approximating them [Oudot et al. 2005; Rineau and Yvinec 2007b; Pons et al. 2007]. For a detailed review of mesh generation methods we refer the reader to comprehensive books and surveys [Frey and George 2007; Shewchuk 2012; Cheng et al. 2012; Dey 2007b].

There is a plethora of commercial solutions for the automatic generation of 3D tetrahedral meshes. Some of them are used in CAD or simulation software packages for computational fluid dynamics or mechanical analysis [George 2005]. The open source offer is more limited, with, e.g., NETGEN from the University of Linz (Advancing Front) [Schöberl 2003], GRUMMP from the University of British Columbia (boundary-preserving, Delaunay-based) [Ollivier-Gooch 2010], QMG from the Cornell University [Vavasis 1999], and TETGEN from the Weierstrass Institute of Berlin [Si 2006]. These solutions all share the common requirement that the domain boundary must be provided in a specific form, e.g. as a defect-free polyhedral surface mesh or parametric surface, and hence are not so generic. A more flexible design appears, e.g., in the form of a common geometry model (CGM) [Gosselin and Ollivier-Gooch 2011]. CGM offers a wide range of functions, but is restricted to the class of models that are the output of solid modelers. Furthermore, most software solutions require a quality surface mesh as input (with well-shaped triangles) as they preserve the input surface mesh exactly, instead of approximating it. The preparation of quality input surface meshes from defect-laden inputs requires applying a sequence of algorithms, which is notoriously labor-intensive and prone to robustness issues.

The main paradigm behind our algorithms is the refinement of a Delaunay triangulation restricted to the input domain and surfaces (in the sequel the input surfaces refer to both subdividing and boundary surfaces). The restricted Delaunay triangulation approach, detailed in Section 2.1, consists in extracting from the 3D Delaunay triangulation a subset of its simplices: vertices, facets approximating the input surfaces, edges approximating the sharp creases (if any), and tetrahedra within the input domain. The refinement process is devised to sample the sharp features, surfaces as well as 3D volumes so that the restricted Delaunay triangulation provides a faithful approximation of surfaces and sharp features, together with a subdivision of the input 3D volume with well-shaped tetrahedra. The paradigm of *restricted Delaunay refinement* has received a special attention due to its versatility and theoretical foundations for mesh generation and shape approximation. It is amenable to algorithms that are proven to terminate and that provide control on the size and quality of the mesh elements. Delaunay refinement has been first used in 2D [Chew 1989], then in 3D for polyhedral domains [Nave et al. 2004], for smooth surfaces [Chew 1993], for 3D domains bounded by smooth surfaces [Oudot et al. 2005; Boivin and Ollivier-Gooch 2002] as well as for 3D domains bounded by piecewise

smooth surfaces [Rineau and Yvinec 2007b; Cheng et al. 2007a; Cheng et al. 2007b; Boltcheva et al. 2009].

The restricted Delaunay refinement approach is more than just a methodological feature of our approach. It is amenable to a novel functionality of the mesh generation software: in case of a multi-component domain, the restricted Delaunay refinement provides in a single refinement process a mesh approximating all components of a subdivided domain and of its subdividing and bounding surfaces. This approach departs from the common pipeline where a surface mesh is first generated for each input surface, then 3D meshes respecting these surface meshes are generated separately for each domain component before merging them to form the final mesh (Figure 1).

Despite its versatility, Delaunay refinement alone is facing two issues. The first issue occurs in the presence of sharp features (creases, corners, darts, cusps, tips) on the input domain boundary. Sharp features subtending either small or large angles may jeopardize the termination of the refinement process. In CGALmesh [Alliez et al. 2013a], sharp features are handled through so-called protecting balls placed before refinement [Cheng et al. 2007b] so as to ensure both the termination of the refinement process and a faithful approximation of the sharp features in the final mesh. The second issue is related to the shape quality of the mesh elements. More specifically, the Delaunay refinement algorithm is shown to terminate when the shape quality governing the refinement is an upper bound on the radius-edge ratio of the mesh elements (surface triangles and volume tetrahedra). The radius-edge ratio is the ratio between the circumradius of the mesh element and the length of its shortest edge. Unfortunately, a bound on radius-edge ratios does not provide a bound on the dihedral angle of tetrahedra. Applied to a 3D domain, the Delaunay refinement process yields a final mesh that may include quasi-degenerate tetrahedra of a special kind, referred to as slivers. Slivers are ill-shaped tetrahedra formed by four vertices close to a circle and roughly equally spaced on this circle; slivers have small radius-edge ratio. To overcome this issue CGALmesh offers optimization processes [Alliez et al. 2005; Du and Wang 2003; Cheng et al. 2000; Tournois et al. 2009] to reduce the number of slivers.

1.2. Contributions

Our main contribution is a generic software design for mesh refinement and optimization. CGALmesh is able to handle a wide range of domains (surfaces, multi-domain volumes, sharp features), and has no equivalent in terms of flexibility. Our means to achieve such flexibility is the notion of so-called “oracle”: all required interactions between the algorithms and the input domain are embedded in a single oracle. Combined with a general restricted Delaunay refinement process, this oracle enables generating in a single refinement process the mesh of a multi-component domain, departing from the common approach which requires generating a surface mesh for each boundary/subdividing surface before generating a 3D mesh preserving those boundary meshes for each volume component.

Our notion of domain oracle departs in the sense that it is completely independent from the input representation: its design results from a careful analysis of the atomic geometric queries performed by the meshing algorithm. Our goal was to embed into the C++ concept of the oracle the minimal number of geometric queries in order to ease future implementations of this concept for a wide range of inputs. Some users have already implemented oracles for application domains ranging from geological modeling to simulation of fluid dynamics through medical imaging. We provide example implementations for common inputs such as polyhedral domains, implicit domains or domains described by grey level or segmented 3D images.

2. ALGORITHMS

The meshing algorithm includes three main phases: initialization, Delaunay refinement and optimization. The initialization phase provides an initial set of mesh vertices and deals with sharp 0 and 1-dimensional input features, if any. The Delaunay refinement phase gradually inserts new vertices into the triangulation, guided by a set of *criteria*. Finally, the optimization phase aims at improving the distribution of tetrahedron dihedral angles.

The *input* to the meshing algorithm is a domain description provided as a model of the *oracle* concept, and a set of parameters tuning to the user's needs the Delaunay refinement *criteria* and the optimization phase. The input domain Ω may be subdivided into several subdomains, each denoted by Ω_i .

The *output* of the meshing algorithm is an isotropic 3D simplicial mesh \mathcal{M} formed by tetrahedral cells (volume mesh) in which triangle facets that form an approximation of the domain boundary surfaces (surface mesh) are marked. CGALmesh can thus be used to generate either 3D meshes or surface meshes.

Before stepping into the details, we describe below the notion of *restricted Delaunay triangulation*, central in CGALmesh algorithms.

2.1. Restricted Delaunay Triangulations

The meshing algorithm maintains the current set P of vertices, its 3D Delaunay triangulation and the restrictions of this triangulation to the surface and volume of the input domain. Figures 2, 3 and 4 depict the 2D counterparts of the notion of Delaunay triangulations and restricted Delaunay triangulations. More specifically, we denote by:

Ω	The input 3D domain (Figure 3a and 4a);
Ω_i	The i^{th} subdomain (Figure 4a);
$S = \bigcup_i bd(\Omega_i)$	The surface, defined as the union of all subdomain boundaries (Figure 3a and 4a);
P	The current set of vertices (Figure 3a and 4a);
$Del(P)$	The 3D Delaunay triangulation of P , <i>i.e.</i> , the triangulation T such that no point in P lies inside the circumsphere of any tetrahedra of T (Figure 2c);
$Vor(P)$	The Voronoi diagram of P , dual of $Del(P)$ (Figure 2b);
$Del _{\Omega}(P)$	The restriction of the Delaunay triangulation to the domain, <i>i.e.</i> , the subcomplex of $Del(P)$ formed by all tetrahedra whose circumsphere centers (circumcenters) lie inside Ω (Figure 3c and 4c);
$Del _S(P)$	The restriction of the Delaunay triangulation to the surface, <i>i.e.</i> , the subcomplex of $Del(P)$ formed by triangles whose dual Voronoi edges intersect S (Figure 3b and 4b). Those triangles are referred to as <i>surface facets</i> .

By definition, each surface facet f has at least one circumscribed ball empty of points of P and whose center lies on the surface. Such a ball, referred to as a *surface Delaunay ball* and denoted by $SDB(f)$, is centered at a point where the dual Voronoi edge intersects the surface (Figure 5a).

At any step of the algorithm, $Del|_S(P)$ is the current approximation of the surface S and $Del|_{\Omega}(P)$ is the current approximation of the domain Ω . In the case of a monodomain, at the end of the Delaunay refinement, $Del|_S(P)$ is the boundary of $Del|_{\Omega}(P)$. In the case of a multi-domain the triangles whose two adjacent tetrahedra are from different subdomains belong to $Del|_S(P)$ (Figure 4). Examples of such multi-domains are depicted by Figure 1 and 10 (liver kidney gallbladder).




Fig. 2. (a) Points P sampling the surface. (b) Voronoi diagram of P . (c) Delaunay diagram of P . We depict in black the circumcircle of the gray triangle, highlighting the empty sphere property. (d) $Vor(P)$ is the dual diagram of $Del(P)$.




Fig. 3. (a) Input domain Ω (gray area), boundary S (dark gray curve), and points P (black dots) sampling the surface. (b) $Del_{|S}(P)$, the Delaunay triangulation restricted to surface S , is formed by edges (thick blue segments) whose dual Voronoi edges (thick red segments) intersect S . (c) $Del_{|\Omega}(P)$, the Delaunay triangulation restricted to the domain Ω , is formed by triangles (in blue) whose circumcenters (blue dots) lie inside Ω .




Fig. 4. (a) Input domain $\Omega = \Omega_1 \cup \Omega_2$ (pink and light blue areas), boundary S (dark gray curve), and points P (black dots) sampling the surface. (b) $Del_{|S}(P)$, the Delaunay triangulation restricted to surface S , is formed by edges (thick blue segments) whose dual Voronoi edges (thick red segments) intersect S . (c) $Del_{|\Omega_i}(P)$, the Delaunay triangulation restricted to the domain Ω , is formed by triangles (in pink for Ω_1 , in blue for Ω_2) whose circumcenters (blue dots) lie inside Ω .




Fig. 5. (a) Surface Delaunay ball. (b) Surface Delaunay ball radius (R), approximation error of the surface (ϵ) and triangle angles (e.g. θ) are used by the facet criteria. (c) Circumsphere radius (R) and shortest tetrahedron edge are used by the cell criteria.

2.2. The Oracle

A model of the oracle concept provides a means for the meshing algorithm to probe the input domain Ω and surface S .

First, the oracle is required to construct a set of initial points on the surface S . When S contains sharp features, (creases and corners), the oracle is further required to enumerate corners and to sample creases.

The second requirement on an oracle is to provide the predicates required to extract from a Delaunay triangulation its restrictions to the domain Ω and the surface S , as well as constructors to further sample the surface when requested during refinement:

- Predicates to know whether a query point is inside Ω , and in the positive, in which subdomain Ω_i ;
- Predicates to know if a given line, ray or line segment query intersects S ;
- Constructors to compute the intersection points between S and a line, ray or line segment query.

2.3. Criteria

The Delaunay refinement phase is governed by a set of user-tuned criteria to be met by the mesh elements (domain tetrahedron cells and surface triangle facets). The elements in $Del_{|\Omega}(P)$ and $Del_{|S}(P)$ are said to be *bad* when they do not meet some of the criteria. The refinement process computes and inserts new vertices – referred to as Steiner points – to remove these bad elements. Facet and cell criteria are designed to control the quality of the output mesh, albeit only to a certain extent in order to ensure the termination of the Delaunay refinement process. As there is additional room for tuning however, the user can provide a set of parameters to tune these criteria to the targeted application.

Facet criteria. Three criteria are used to control the properties of the surface facets:

- (1) **Size:** an upper bound for the radii of surface Delaunay balls (R depicted by Figure 5b), specified either as a uniform or as a non-uniform scalar sizing field (i.e. a function from \mathbb{R}^3 to \mathbb{R}).
- (2) **Shape:** an upper bound on the radius-edge ratio of surface facets equivalent to a lower bound on the angles of surface facets, (θ depicted by Figure 5b).
- (3) **Approximation:** an upper bound for the approximation error of surface facets (ϵ depicted by Figure 5b). The local approximation error for each surface facet f is estimated as the distance between the circumcenter of f and the center of its

surface Delaunay ball. The bound can be specified either as a uniform or as a non-uniform scalar field.

Cell criteria. For tetrahedral volume mesh generation, two additional criteria are used to control the properties of the mesh tetrahedron cells:

- (1) **Size:** an upper bound for the circumradii of the cells in $Del_{|\Omega}(P)$ (R depicted by Figure 5c), specified either as a uniform or as a non-uniform scalar sizing field.
- (2) **Shape:** an upper bound for the radius-edge ratio of the cells in $Del_{|\Omega}(P)$. The radius-edge ratio relates the circumradius of a cell to the length of its shortest edge (Figure 5c). Enforcing such an upper bound on the radius-edge ratio leads to a vast majority of well-shaped (isotropic) cells in the final mesh, and possibly a few slivers.

2.4. Delaunay Refinement

Starting from a set of points computed in the initialization phase (Section 2.5), the Delaunay refinement process gradually inserts new Steiner points into P until the mesh elements (in $Del_{|S}(P)$ and $Del_{|\Omega}(P)$) meet the criteria. The location of each inserted Steiner point is computed so as to suppress from $Del_{|S}(P)$ and $Del_{|\Omega}(P)$ the surface facets or domain cells that do not meet the criteria.

Definition 2.1. The Steiner point computed to suppress a bad surface facet f is the center of its surface Delaunay ball $SDB(f)$ (Section 2.1), or the center of one of those balls if there are several.

Definition 2.2. The Steiner point computed to suppress a bad cell c is the circumcenter of c .

The Delaunay refinement phase involves two steps: the *surface mesh step* and the *volume mesh step*. The surface mesh step handles bad surface facets, adding only Steiner points onto the surface, and running until all surface facets meet the criteria. The surface mesh step first performs a *facet scan* process to initialize the restricted Delaunay triangulation $Del_{|S}(P)$ and a priority queue Q_F including the bad surface facets of $Del_{|S}(P)$.

Facet scan. All triangle facets of $Del(P)$ are scanned as follows: for each triangle t whose dual edge intersects the input surface:

- (1) Add surface facet t to $Del_{|S}(P)$;
- (2) Test t against facet criteria;
- (3) If t does not meet all the facet criteria, add t to queue Q_F .

Queue Q_F now contains all bad facets of $Del_{|S}(P)$. The *surface mesh refinement* step is then started and runs until Q_F is empty. Note that Q_F is sorted such that facets with large surface Delaunay balls are popped and refined first.

Surface mesh refinement. For each bad facet f of Q_F the following operations are performed:

- (1) Compute the Steiner point p of f (Definition 2.1);
- (2) Insert p into $Del(P)$;
- (3) Update $Del_{|S}(P)$;
- (4) Update Q_F by removing destroyed facets and adding newly created bad surface facets if any.

The surface mesh step terminates when Q_F is empty for the first time. The surface mesh $Del_{|S}(P)$ now meets all facet criteria. If only a surface mesh is sought after by

the user, then $Del_{|S}(P)$ is the final interpolating output mesh. If a 3D mesh is sought after, the algorithm proceeds to the *volume mesh step*. The latter consists of a *cell scan* process to initialize a queue Q_C of bad cells, followed by the *volume mesh refinement* that will handle bad cells *and* bad facets through inserting Steiner points either inside Ω or on S .

Cell scan. P now contains the initial points as well as all Steiner points inserted during surface mesh refinement. All bad cells are now collected as follows. For each cell c of $Del(P)$ whose dual circumcenter is in Ω :

- (1) Add c to $Del_{|\Omega}(P)$;
- (2) Test c against the cell criteria;
- (3) If c does not meet all the cell criteria, add it to the cell refinement queue Q_C .

Q_C now contains all bad cells of $Del_{|\Omega}(P)$. Q_C is a priority queue sorted by decreasing size of circumradius so that cells with large circumsphere radii are popped and refined first,.

Volume mesh refinement. The volume mesh refinement requires maintaining the two priority queues Q_C and Q_F for bad cells and facets. The two following rules are applied as long as one of Q_C and Q_F is not empty. Rule 1 has priority and is repeatedly applied as long as Q_F is not empty.

Rule 1 If Q_F is not empty, pop a facet f then:

- (1) Compute Steiner point p of f (Definition 2.1);
- (2) Insert p into $Del(P)$;
- (3) Update $Del_{|S}(P)$ and $Del_{|\Omega}(P)$;
- (4) Update Q_F by removing destroyed facets and adding all newly created surface facets which do not meet all the criteria.
- (5) Update Q_C by removing destroyed cells, and adding newly created bad cells.

Rule 2 If Q_C is not empty, pop a cell c and perform the following operations:

- (1) Compute the refinement point p of c (Definition 2.2);
- (2) If p is inside the surface Delaunay ball of a surface facet f (p is said to *encroach* f), add f to Q_F and leave c in Q_C .
- (3) Otherwise:
 - (a) Insert p into $Del(P)$;
 - (b) Update $Del_{|\Omega}(P)$;
 - (c) Update Q_C by removing destroyed cells, and adding newly created bad cells.

The algorithm terminates when both Rule 1 and Rule 2 do not apply anymore, *i.e.*, when both Q_F and Q_C are empty. The final 3D mesh $Del_{|\Omega}(P)$ now meets both facet and cell criteria.

2.5. Initialization and Sharp Features

The initialization phase generates through the oracle a small set of points lying on the surface, which are inserted into P . Depending on the input, these initial points may be obtained by probing the surface through, *e.g.*, random ray shooting, by using a user-provided point set or by any other sampling method best suited to the targeted application.

During refinement Steiner points are inserted either inside Ω or on S . These Steiner points are either generated by probing Ω with Voronoi vertices (cell circumcenters) or by probing S with Voronoi edges (dual of Delaunay facets). Sharp features (creases and corners) are not probed during refinement. While probing sharp creases with

Voronoi faces (dual of Delaunay edges) is conceptually possible, it is known to jeopardize the termination of the refinement when sharp features subtend small or large angles. Sharp features are thus not faithfully approximated by the refined mesh unless they are dealt with during the initialization phase (Figure 6a). To ensure an accurate representation of sharp features in the final mesh, CGALmesh relies on the so-called protecting-balls approach introduced by Cheng et al. [Cheng et al. 2007b]. The main idea behind the protecting-balls is to discretize sharp features during initialization phase and ensure – through the protecting balls – that this discretization appears unchanged in the final mesh. This is achieved as follows. During initialization phase we add one vertex per corners, sample the creases and compute all protecting balls – a set of balls centered either at corners or at crease sample points, and such that:

- The union of the balls entirely covers the sharp features;
- Three balls do not mutually intersect;
- Each ball does not include any other ball center;
- Two balls centered on different creases do not intersect.

The refinement process is then modified to use a weighted Delaunay triangulation. It is started with an initial set of weighted points including weighted points corresponding to the protecting balls: those weighted points are located at the ball centers and their weights are set to the squared radii of the balls. Any other Steiner point later added by the refinement process is given a zero weight. Such weights ensure that each segment joining the centers of two consecutive protecting balls along a sharp crease exists as an edge in the final mesh. Protecting balls guarantee the termination of the Delaunay refinement process whatever may be the angles formed by input surface patches or input creases. Furthermore, the input surface S is never probed by the refinement process inside the union of protecting balls, which is an advantage in cases where input data are noisy around sharp creases.




Fig. 6. Sharp features on a CAD model: (a) Meshing without features. (b) Protecting balls. (c) Meshing with features.

Figures 6 and 11 depict examples of CAD models meshed without and with protecting the sharp features. Note that in case of domains with sharp features, the domain oracle must be enriched with the description of these features. More specifically, the oracle must provide the list of sharp corners and capabilities to sample sharp creases by computing a sample point on a sharp crease at a given distance from another sample point. The oracle must also provide topological information relating to the incidence

graph of corners, creases and surface patches. Depending on the oracle implementation, these features may be provided by the user, or automatically detected – *e.g.* based on an angle criterion.

2.6. Termination Proof and Guarantees

Termination proof. If the minimum facet angle criterion is set to be $\leq 30^\circ$ and the maximum cell radius-edge-ratio criterion is set to be ≥ 2 , then the above algorithm is shown to terminate after inserting a finite number of Steiner vertices [Chew 1993; Ruppert 1995; Shewchuk 1998].

Guarantees. The resulting mesh includes surface facets and cells meeting all specified criteria about the element size, element shape and boundary approximation error.

Note however that the approximation criterion met by all surface facets does not imply the accuracy of the approximation of S by the surface mesh. Indeed, together with the sizing field, this criterion provides a bound on the one-sided Hausdorff distance from the surface mesh to the input surface but not the other way round.

Further guarantees about the accuracy of the surface approximation may be obtained, relying upon the notion of *medial axis* and *local feature size*. Given an input surface S , the medial axis of S is defined as the locus of points that have more than one closest point on S . The local feature size $lfs(x, S)$ is a sizing field equating for each point x its distance to the medial axis of S . The local feature size is thus related to both the surface curvature – a local property – and the thickness or separation of the domain – a global property.

It has been shown [Amenta and Bern 1998; Cohen-Steiner and Morvan 2003] (see also [Dey 2007a]) that, if a point set P is a sampling of the surface S , locally dense enough with respect to the feature size of S ,

- $Del_{|S}(P)$ is homeomorphic to S and $Del_{|\Omega}(P)$ is homeomorphic to Ω ;
- Hausdorff-distance($Del_{|S}(P), S$) ≈ 0 ;
- $Del_{|S}(P)$ provides a good estimation of normals, area as well as curvature of S .

Further results [Boissonnat and Oudot 2005] show that the Delaunay refinement process described above provides a sampling of the surface that is dense enough with respect to the local feature size, provided that the sizing field input to the refinement process is itself small enough with respect to this local feature size. Unfortunately, in general the local feature size of the surface is unknown. The approximation criterion is a way to adapt locally the sizing of the mesh to the surface curvature but this criterion is insensitive to thickness and separation. From the user point of view a conservative solution is to get an estimation for the minimum thickness/separation of the domain and to specify a uniform sizing function that equates to this estimation.

In addition, the initial point set P must be dense enough to avoid that the mesh refinement process misses completely some small connected components of the surface. The proposed oracles use random probing and/or sampling of sharp features for this purpose, which yields very good results in most cases. Some specific initialization is however required if one needs absolute guarantees that no surface components are missed. One way to provide such a guarantee is the so-called *persistent facet* approach [Boissonnat and Oudot 2005]. The latter consists in adding to the initial set of vertices, three points per surface connected component, sufficiently close to construct a surface facet that persists throughout the refinement process. Note however that such a solution may yield small overly refined areas in the final mesh.

2.7. Mesh Optimization

As discussed above, the Delaunay refinement process provides control upon the topological and geometric accuracy of the domain and surfaces approximation, on the size of mesh elements, and on the shape of triangles in the surface approximation. It also provides control upon the shape of mesh tetrahedra through the radius-edge criterion, albeit this is not enough to provide a non-trivial bound on the dihedral angles of tetrahedra. The radius-edge criterion ensures that the output mesh contains no ill-shaped tetrahedra except for one class of them denoted by slivers. Slivers are formed by four vertices close to a circle and roughly equally spaced along this circle. Slivers are the only class of quasi degenerate tetrahedra that do not have a large radius-edge ratio, which explains why Delaunay refinement is blind to slivers. Figure 7 illustrates the slivers remaining after meshing a sphere. Removing slivers is the main purpose of the mesh optimization phase that we describe next.




Fig. 7. Slivers after Delaunay refinement. The final tetrahedral mesh contains about 38K tetrahedra, some of them being slivers, which are shown for a dihedral angle bound of respectively 5 degrees (middle – 35 slivers) and 10 degrees (right – 136 slivers).

Many mesh quality improvements techniques have been proposed in the literature [Freitag and Plassmann 2000; Klingner and Shewchuk 2008; Park and Shontz 2010; Sastry et al. 2012]. CGALmesh offers four different mesh optimization methods which can be combined sequentially. These methods fall into two main categories: global optimization and local optimization. The first category includes Lloyd relaxation [Du et al. 1999; Du and Wang 2003] and Optimal Delaunay Triangulation (ODT for short) relaxation [Chen and Xu 2004; Alliez et al. 2005]. The second category includes vertex perturbation [Li 2000; Tournois et al. 2009] and sliver exudation [Cheng et al. 2000].

While global optimizers substantially reduce the total number of slivers, local optimizers focus on the worst tetrahedra. Our experiments show that a Lloyd global optimization followed by vertex perturbation often provides us with the best results when seeking for the largest minimum dihedral angle for the tetrahedra.

Global optimization. Global optimizers consider an energy that is a function of the mesh vertices and connectivity, and aim at minimizing this energy by relocating all mesh vertices at once, before updating the mesh connectivity.

Both Lloyd and ODT smoothers use an energy defined as the L_1 -norm of the approximation error between the isotropic paraboloid function $f(x) = x^2$ (whose graph is in \mathbb{R}^4) and a piecewise linear interpolation of this function on the domain.

In the Lloyd smoother, the linear interpolation f_{PWL}^{dual} is defined on each Voronoi cell as the linear expansion of $f(x)$ from its value at the Voronoi site. Consider the unit isotropic paraboloid $y = x^2$ of \mathbb{R}^4 , the interpolating points (x_i, x_i^2) obtained by lifting the mesh vertices on the paraboloid, and the polytope P_{Lloyd} defined as the intersection of the half-spaces above the hyperplanes tangent to the paraboloid at the lifted mesh vertices. The L_1 -norm of the approximation error, $\|f - f_{PWL}^{dual}\|_{L_1}$, measures the volume enclosed between the paraboloid and the boundary of the polytope P_{Lloyd} (Figure 8).

In the ODT smoother, the linear interpolation f_{PWL}^{primal} is defined as the linear interpolation of $f(x)$ on each mesh tetrahedron. The L_1 -norm of the approximation error, $\|f_{PWL}^{primal} - f\|_{L_1}$, measures the volume enclosed between the unit paraboloid of \mathbb{R}^4 and the lower boundary of the convex hull of the lifted vertices (Figure 8).

Note that after Delaunay refinement the final mesh may have a non-uniform density of vertices, reflecting a non-uniform sizing field that is the pointwise minimum between a (possibly non-uniform) user-defined sizing field and the local feature size of the meshed domain. To preserve this non-uniform density throughout the optimization process, the Lloyd and ODT energy integrals are computed using a weighted version of the error, where the weights are locally estimated from the average length of edges incident to each vertex of the mesh after refinement. For both optimizers, at each optimization step, closed form formulas provide the new location of the mesh vertices as a function of the current mesh vertices and connectivity [Du and Wang 2003; Alliez et al. 2005]. Each optimization step computes the new position of all mesh vertices, relocates them and updates the Delaunay triangulation as well as both restricted Delaunay triangulations.




Fig. 8. 1D depiction of the mesh energy minimized by Lloyd (left) and ODT (right) smoothers. Black dots on the x axis represent the mesh vertices. White circles represent Voronoi vertices. The energy to be minimized is the volume measure of the shaded region.

The user is provided control upon the termination of the optimization process through the following parameters:

- A time limit;
- A maximum number of optimization steps;
- A “convergence” ratio r . The optimization process stops when, during the last iteration, the displacement of any vertex v is lower than a given fraction r of the length of the shortest edge incident to v .

Local optimization. The vertex perturber and the sliver exuder act locally in the sense that they first parse the mesh in order to find the worst elements and process them individually. The focus is on improving the worst elements first.

The vertex perturber [Li 2000; Tournois et al. 2009] improves the mesh by relocating the vertices in turn, starting with vertices incident to the worst sliver tetrahedra. Each vertex relocation is followed by updating the mesh connectivity to restore the Delaunay and restricted Delaunay properties of the mesh. For each vertex, the algorithm searches for a new location that yields an improvement of the vertex star. More specifically, the new vertex location is not computed so as to improve the shape of the incident sliver, but instead so as to make this sliver disappear when the Delaunay mesh connectivity is restored. The new vertex location is thus first searched so as to increase the incident sliver circumradius. If this fails to yield a change in the mesh connectivity which improves the star, another vertex location is searched so as to invert the worst incident tetrahedron. If this again fails, random perturbations of the vertex location are attempted.

The sliver exuder [Cheng et al. 2000] turns the Delaunay triangulation into a weighted Delaunay triangulation. It aims at removing slivers through re-weighting the mesh vertices with optimal weights that trigger combinatorial changes inducing sliver exudation. As for the global optimizers, the user is provided control upon the termination of local optimization processes, through the following parameters:

- A time limit;
- A lower bound on dihedral angles of the mesh tetrahedra.

3. IMPLEMENTATION DETAILS AND SOFTWARE DESIGN

CGALmesh [Alliez et al. 2013a] is the 3D Mesh Generation package of the CGAL library [CGAL Editorial Board b], which is available on the CGAL website [CGAL Editorial Board a]. For Linux systems, CGAL is part of the package repositories of Arch Linux, Debian-based distributions (Ubuntu, Mint, KNOPPIX, etc.) and distributions using the RPM package manager.

CGALmesh has been carefully designed and implemented in order to achieve both reliability and flexibility. We now go through the main design choices made during the implementation process.

3.1. Generic Programming

CGAL is a C++ library, designed to be highly generic through the use of templates. Each CGAL package provides the users with the ability to use their own data types for elementary geometric objects as long as they provide the required attributes types and functions. Being part of CGAL, CGALmesh follows such generic design and provides an additional level of genericness through the use of an abstract oracle to interface with the input domain representation. This oracle encapsulates the required domain knowledge. Supporting a novel type of domain representation boils down to implementing a new oracle. The current software provides oracles for polyhedral domains, implicit surfaces and 3D labeled images.

Figure 9 depicts the overall modular design of the mesh generation algorithm and the interaction between the different modular components. The core mesh generation engine is described in a dedicated publication [Rineau and Yvinec 2007a].

3.2. Oracles

CGALmesh already provides oracles for domains described through polyhedral meshes, implicit functions and 3D images.

The polyhedral mesh oracle takes as input a triangle surface mesh. To provide efficient surface intersections during refinement a hierarchical data structure of axis-aligned bounding boxes (AABB tree) is first created [Alliez et al. 2013b]. The oracle for implicit functions uses the concept of bipolar Voronoi edges to detect intersections. A




Fig. 9. Overall design of the mesh generation process.

Voronoi edge is said to be bipolar if the implicit function values have opposite signs at its endpoints [Boissonnat and Oudot 2005]. Only bipolar Voronoi edges are considered to intersect the surface and the intersection points are obtained through dichotomy, *i.e.*, by recursively subdividing the edge into two equal edges, keeping the one whose endpoints have opposite signs, until reaching a specified precision. The 3D image oracle computes implicit functions from input 3D images, and relies on the so-constructed implicit function oracle for predicates and constructions [Pons et al. 2007].

The complexity of the predicates and constructor required in the oracle greatly varies among the different models for the oracle. For the case of implicit function and 3D images, testing for intersection boils down to evaluating the function at two locations, whereas the polyhedral oracle requires a compute-intensive tree search. Note however that computing the intersection by dichotomy is more complex than computing the intersection point between a Voronoi edge and a triangle in the leaf of the AABB tree.

The flexibility offered by the oracle design makes CGALmesh particularly well suited to applications where the input domain is either known through measurements [Salman et al. 2010; Mullen et al. 2010; Alliez et al.] or described with complex application-specific data structures such as composed implicit functions in geophysics. As the input domain is queried only at meshing time in a coarse-to-fine manner the CGALmesh algorithms are also favored in constrained-memory applications, where complex data processing such as robust denoising can be delegated to the oracle [Campos et al. 2013]. We also performed prototype experiments with oracles designed to mesh smooth subdivision surfaces and smooth parametric NURBS surfaces, where the oracles embed complex numerical procedures to compute intersection points.

3.3. Data Structures

The CGALmesh package primarily relies on the CGAL 3D Delaunay triangulation package. The latter provides fast and reliable algorithms to create, modify and access the Delaunay triangulation of a set of 3D points. The meshing algorithm maintains the

restricted Delaunay triangulations to surfaces and volumes which are respectively a 2D and a 3D complex embedded in the 3D Delaunay triangulation of the vertices. Two additional data structures are thus used to represent 2D and 3D complexes embedded in a 3D triangulation.

4. MESHING EXPERIMENTS

We now provide a series of experimental results on a variety of input domains (Figure 10). Timings are measured on a PC running Windows 7 with an Intel Xeon 2.4 GHz CPU and 128GB RAM, using CGAL 4.4, 64-bit version. Note that CGAL – which includes CGALmesh – is also available for Linux, MacOS X and Solaris; it can be downloaded on the CGAL website [CGAL Editorial Board a].

4.1. Mesh Refinement

Table I and Figure 12 focus on the refinement step of the mesh generation process, using the three provided oracles: polyhedral surface meshes, implicit functions and 3D images. Timings for the facet scan process, negligible, are not shown. The first set of input domains contains polyhedral surfaces that are piecewise linear approximations of smooth surfaces. The second set contains CAD polyhedral surfaces with sharp features. The third set contains domains defined by implicit functions. The fourth set contains 3D images. Figure 12 shows that most domains exhibit nearly similar refinement speed, except for the *cheese* and the *ecrou-larousse* models which are slower due to a large number of sharp creases. Note that the *cheese* model, containing many sharp creases incident to small angles, is correctly meshed, albeit slowly, by the meshing process (Figure 11). These results detail the volume and surface meshing performances.

Table II compares CGALmesh with LOCPSC [Dey et al. 2010], a Delaunay-based surface mesh refinement algorithm based on CGAL triangulations. By not specifying any tetrahedron-related criteria, we use CGALmesh to produce a triangular mesh of similar size and quality than the ones produced by LOCPSC. Those results show that the main strength of CGAL is its processing times, while LocPSC consumes less memory through a localized Delaunay refinement process. We observe that CGALmesh is slightly more parsimonious.

Regarding tetrahedral mesh generation, most available tools, such as TETGEN [Si 2006], require a triangle surface mesh as input, and the quality of the generated tetrahedral mesh directly relates to the quality of the input. Although comparing two different kinds of methods is difficult, we provide through table III a comparison between the volume refinement step of CGALmesh to the mesh generation process of TETGEN. CGALmesh is not as fast as TETGEN, but yields honorable speed when considering the fact that it maintains a Delaunay triangulation of the surface while refining the volume.

4.2. Mesh Optimization

Figures 13, 14 and 15 illustrate how combining global and local optimization helps improving the distribution of tetrahedron dihedral angles. The global optimization step is effective at concentrating the distribution, but leaves a few extreme dihedral angles. The local optimization process removes most extreme angles while leaving other angles unchanged. On the smooth Bimba model the Lloyd optimization step improves the whole distribution of angles. On the CAD models the Lloyd optimization step improves mostly near curved areas, as the elements on flat areas are already well shaped straight after refinement. On CAD models the protecting balls prevent the vertices on the sharp creases to be relocated during optimization. The local optimization is thus not as effective as for smooth models, and a few slivers remain.

Table I. Timings for mesh refinement – domains are depicted by Figure 10

Domain	#V ¹	#F ¹	#T ¹	S.R. ² (s)	C.S. ³ (s)	V.R. ⁴ (s)	Total (s)	#V/s ⁵	Mem(MB)
Elephant (5k triangles)	536k	126k	3,218k	8.8	0.9	36.4	46.1	11,622	1,187
	4,152k	507k	25,709k	35.5	3.5	352.0	391.1	10,615	8,638
	9,764k	903k	60,927k	64.0	6.9	782.4	853.3	11,443	19,997
Bimba (384k triangles)	965k	139k	5,941k	7.8	1.3	53.5	62.6	15,426	2,235
	4,400k	387k	27,486k	22.1	3.7	277.3	303.0	14,520	9,505
	14,736k	874k	92,734k	49.9	8.6	1067.3	1125.8	13,089	30,410
lucy (100k triangles)	566k	179k	3,327k	14.0	1.9	31.6	47.5	11,923	1,275
	4,343k	723k	26,585k	56.2	8.5	289.1	353.7	12,277	9,159
	28,042k	2,550k	175,063k	196.2	30.7	2192.9	2419.8	11,588	57,625
david (200k triangles)	731k	198k	4,350k	12.7	1.9	41.3	55.9	13,063	1,668
	3,288k	554k	20,119k	35.6	5.6	211.1	252.3	13,031	6,916
	10,936k	1,250k	67,866k	80.3	13.0	799.1	892.4	12,255	22,701
ecrou-larousse (27k triangles)	836k	346k	4,760k	59.9	8.1	45.5	113.5	7,363	1,531
	3,581k	799k	21,584k	128.3	16.5	226.5	371.3	9,645	6,418
	11,784k	1,703k	72,553k	271.3	33.4	912.6	1217.3	9,680	20,659
pump.carter (30k triangles)	600k	197k	3,519k	12.5	1.3	30.2	44.0	13,645	1,132
	4,578k	780k	28,002k	48.0	5.5	278.4	331.9	13,792	8,093
turbine (18k triangles)	449k	201k	2,544k	19.3	2.0	24.1	45.4	9,885	884
	3,367k	788k	20,243k	71.6	8.0	202.8	282.5	11,919	6,001
cheese (18k triangles)	1,025k	463k	5,797k	92.8	9.7	65.8	168.3	6,092	1,952
	5,179k	1,396k	30,836k	268.6	28.5	363.7	660.7	7,838	9,097
	39,851k	5,505k	245,751k	1053.6	117.2	3657.9	4828.7	8,253	69,635
Thin.cylinder_function	316k	406k	1,149k	16.1	0.7	3.4	20.3	15,569	655
	1,585k	1,665k	6,607k	68.4	3.3	29.2	100.8	15,722	3,187
	6,115k	6,607k	24,888k	343.5	13.7	143.8	500.9	12,207	12,267
Klein_function	754k	639k	3,734k	24.5	1.4	33.9	59.8	12,608	1,554
	3,062k	2,014k	16,226k	81.1	4.4	147.5	233.0	13,142	6,320
	21,721k	9,817k	122,315k	394.8	26.6	1871.6	2292.9	9,473	44,164
Pancake_function	485k	881k	1,645k	34.1	1.9	6.6	42.7	11,371	1,002
	4,402k	6,676k	15,896k	290.4	14.6	76.9	382.0	11,523	8,781
Tanglecube_function	389k	482k	1,664k	18.4	1.0	14.8	34.2	11,379	818
	1,478k	1,522k	6,905k	60.1	3.2	59.0	122.3	12,085	3,091
liver_kidney_gallbladder	1,289k	215k	7,915k	10.6	0.5	65.0	76.2	16,934	2,722
	4,630k	525k	28,781k	26.3	1.5	262.4	290.1	15,957	9,623
	19,666k	1,459k	123,437k	74.0	4.7	1319.0	1397.7	14,071	40,459
VisibleHuman 1mm	345k	715k	2,055k	25.3	1.3	17.6	44.2	7,811	1,116
	3,200k	4,328k	19,795k	155.5	10.0	204.9	370.4	8,640	7,036
	10,448k	10,812k	65,426k	424.6	31.1	725.9	1181.6	8,842	21,859

¹#V = Number of vertices / #F = Number of triangles / #T = Number of tetrahedra²S.R. = Surface mesh refinement³C.S. = Cell scan⁴V.R. = Volume mesh refinement⁵#V/s = Number of vertices created per second




Fig. 10. Input domains: domains bounded by smooth surfaces (top row, in blue), CAD models with sharp features (second row, in grey), implicit functions (third row, in green), 3D images (bottom row, multicolored). All non implicit models are available from the VISIONAIR shape repository, except for the Visible human model which is available from the Visible Human Project (http://www.nlm.nih.gov/research/visible/getting_data.html).




Fig. 11. Cheese model. (left) Meshing without sharp features. (right) Meshing with sharp features. Input domain courtesy Distene.




Fig. 12. Refinement times with respect to the complexity of the output mesh.

Table II. Refinement results for surface mesh generation, compared to the LocPSC software.

Model	Algorithm	Criteria	#V ¹	#F ¹	Time (s)	Mem (MB)
Fandisk	LocPSC	$\lambda = 0.0035$	468,652	937,300	309	129
	CGAL 4.4	$FS^2 = 0.001788$	455,649	911,262	63	893
Rocker	LocPSC	$\lambda = 0.0045$	468,543	937,086	757	137
	CGAL 4.4	$FS^2 = 0.001365$	459,674	919,318	79	1,000
3 Holes	LocPSC	$\lambda = 0.0011$ & $AT^3 = 0.01$	8,165,426	16,330,860	9,470	2,138
	CGAL 4.4	$FS^2 = 0.03764$	8,042,469	16,084,914	1,268	15,976
Lucy	LocPSC	$\lambda = 0.004$ & $AT^3 = 0.01$	371,656	743,308	309	143
	CGAL 4.4	$FS^2 = 0.001341$	365,579	731,110	87	864
Fertility	LocPSC	$\lambda = 0.0035$ & $AT^3 = 0.01$	611,537	1,223,086	606	338
	CGAL 4.4	$FS^2 = 0.2568$	603,847	1,207,658	111	1,553

¹#V = Number of vertices / #F = Number of surface triangles

²FS = Facet size

³AT = Angle threshold

Table III. Refinement results for volume mesh generation, compared to the TETGEN software.

Model	Algorithm	#V ¹	#F ¹	#T ¹	Time (s)
Fandisk	Tetgen	42,654	16,194	247,376	1.79
	CGAL 4.4	43,171	19,178	244,542	$C.S.^2 = 0.02$ & $V.R.^3 = 3.4$
Lucy	Tetgen	72,157	144,310	248,081	1.88
	CGAL 4.4	78,415	145,022	245,536	$C.S.^2 = 2.5$ & $V.R.^3 = 1.5$
David	Tetgen	100,979	201,974	335,144	2.71
	CGAL 4.4	111,103	203,730	335,834	$C.S.^2 = 3.1$ & $V.R.^3 = 4.9$

¹#V = Number of vertices / #F = Number of triangles / #T = Number of tetrahedra

²C.S. = Cell scan

³V.R. = Volume mesh refinement




Fig. 13. Mesh refinement and optimization on the Bimba model. The distribution of dihedral angles, with both min and max values, is depicted under each mesh.




Fig. 14. Mesh refinement and optimization of the Pump_carter model. The distribution of dihedral angles, with both min and max values, is depicted under each mesh.




Fig. 15. Mesh refinement and optimization of the Cheese model. The distribution of dihedral angles – not taking into account dihedral angles at input creases –, with both min and max values, is depicted under each mesh.

5. CONCLUSION

CGALmesh is the 3D mesh generation package provided by the CGAL library. It offers algorithms to generate isotropic surface triangular meshes and 3D tetrahedral meshes. The user is provided with control over the output meshes through several criteria, and four optional optimization algorithms are used to trade mesh quality for time. The Delaunay refinement process embedded in the main mesh generation algorithm comes with several theoretical and practical guarantees which translate from the end user point of view into reliability and conservative matching of the user-specified criteria.

The input domain is queried in an abstract and generic manner through *oracles*. This results in a versatile computational software tool which can be used, among other applications, for meshing CAD models – preserving sharp edges and creases –, extracting isosurfaces, remeshing polyhedral surfaces or meshing from 3D segmented images¹.

In the near future we plan to release parallel versions of the proposed algorithms so as to take advantage of modern multi-core processor architectures. As already shown [Chernikov and Chrisochoides 2010] mesh generation is far from being an embarrassingly parallel problem.

Another stimulating work for CGALmesh is the generation of anisotropic simplicial surface and volume meshes [Boissonnat et al. 2008; Boissonnat et al. 2012], which will be integrated into CGAL in the near future.

Regarding optimization, we plan to implement new algorithms, such as methods which combine non-Delaunay swapping and optimization-based smoothing. See [Klingner and Shewchuk 2008] for an example of such techniques.

6. ACKNOWLEDGMENTS

This work was funded by Inria Technology Transfer and Innovation, and by the European Research Council (ERC Starting Grant “Robust Geometry Processing”, Grant agreement 257474).

REFERENCES

- ALLIEZ, P., COHEN-STEINER, D., YVINEC, M., AND DESBRUN, M. 2005. Variational Tetrahedral Meshing. *ACM Transactions on Graphics (SIGGRAPH)* 24, 617–625.
- ALLIEZ, P., RINEAU, L., TAYEB, S., TOURNOIS, J., AND YVINEC, M. 2013a. 3D mesh generation. In CGAL user and reference manual. CGAL Editorial Board, 4.2 edition. http://www.cgal.org/Manual/latest/doc.html/cgal_manual/packages.html#Pkg:Mesh.3.
- ALLIEZ, P., SABORET, L., AND GUENNEBAUD, G. Surface reconstruction from point sets. In CGAL user and reference manual. CGAL Editorial Board, 4.2 edition.
- ALLIEZ, P., TAYEB, S., AND WORMSER, C. 2013b. 3D fast intersection and distance computation (AABB tree). In CGAL user and reference manual. CGAL Editorial Board, 4.2 edition. http://www.cgal.org/Manual/latest/doc.html/cgal_manual/packages.html#Pkg:AABB.tree.
- AMENTA, N. AND BERN, M. 1998. Surface reconstruction by Voronoi filtering. In *Proc. 14th Annu. Sympos. Comput. Geom.* 39–48.
- BOISSONNAT, J.-D. AND OUDOT, S. 2005. Provably good sampling and meshing of surfaces. *Graphical Models* 67, 405–451.
- BOISSONNAT, J.-D., SHI, K.-L., TOURNOIS, J., AND YVINEC, M. 2012. Anisotropic Delaunay meshes of surfaces. *ACM Transactions on Graphics*.
- BOISSONNAT, J.-D., WORMSER, C., AND YVINEC, M. 2008. Locally uniform anisotropic meshing. In *Proc. 24th Annual Symp. on Computational Geometry*. 270–277.
- BOIVIN, C. AND OLLIVIER-GOOCH, C. 2002. Guaranteed-quality triangular mesh generation for domains with curved boundaries. *Int. J. Numer. Methods Eng* 55, 1185–1213.
- BOLTCHIEVA, D., YVINEC, M., AND BOISSONNAT, J.-D. 2009. Mesh generation from 3D multi-material images. In *Medical Image Computing and Computer-Assisted Intervention*. Lecture Notes in Computer Science Series, vol. 5762. 283–290.

¹Projects using CGAL: <http://www.cgal.org/projects.html>

- CAMPOS, R., GARCIA, R., ALLIEZ, P., AND YVINEC, M. 2013. Splats-based surface reconstruction from defect-laden point sets. *Graphical Models*. to appear.
- CGAL EDITORIAL BOARD. Downloading cgal. <http://www.cgal.org/download.html>.
- CGAL EDITORIAL BOARD. CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>.
- CHEN, L. AND XU, J. 2004. Optimal Delaunay triangulations. *J. Comput. Math.* 22, 2, 299–308.
- CHENG, S., DEY, T., AND LEVINE, J. 2007a. A practical Delaunay meshing algorithm for a large class of domains. In *Proc. of the 16th Int. Meshing Roundtable*. 477–494.
- CHENG, S.-W., DEY, T. K., EDELSBRUNNER, H., FACELLO, M. A., AND TENG, S.-H. 2000. Sliver exudation. *J. ACM* 47, 883–904.
- CHENG, S.-W., DEY, T. K., AND RAMOS, E. A. 2007b. Delaunay refinement for piecewise smooth complexes. In *Proc. of the 18th annual ACM-SIAM Symp. on Discrete Algorithms*. 1096–1105.
- CHENG, S.-W., DEY, T. K., AND SHEWCHUK, J. R. 2012. *Delaunay Mesh Generation*. CRC Press.
- CHENTANEZ, N., FELDMAN, B. E., LABELLE, F., O'BRIEN, J. F., AND SHEWCHUK, J. R. 2007. Liquid simulation on lattice-based tetrahedral meshes. In *ACM SIGGRAPH/Eurographics Symp. on Computer Animation 2007*. 219–228.
- CHERNIKOV, A. N. AND CHRISOCHOIDES, N. P. 2010. A template for developing next generation parallel Delaunay refinement methods. *Finite Elements in Analysis and Design* 46, 1-2, 96 – 113. Mesh Generation - Applications and Adaptation.
- CHEW, L. 1993. Guaranteed-quality mesh generation for curved surfaces. In *SCG '93: Proc. of the 9th Annual Symp. on Computational Geometry*. ACM New York, NY, USA, 274–280.
- CHEW, L. P. 1989. Guaranteed-quality triangular meshes. Technical Report TR-89-983, Dept. Comput. Sci., Cornell Univ., Ithaca, NY. Apr.
- COHEN-STEINER, D. AND MORVAN, J.-M. 2003. Restricted Delaunay triangulations and normal cycle. In *Proc. 19th Annual Symposium on Computational Geometry*. 237–246.
- DEY, T. K. 2007a. *Curve and surface reconstruction: algorithms with mathematical analysis*. Vol. 23. Cambridge University Press.
- DEY, T. K. 2007b. Delaunay mesh generation of three dimensional domains. *Tessellations in the sciences: Virtues, Techniques and Applications of Geometric Tilings*.
- DEY, T. K., JANOOS, F., AND LEVINE, J. A. 2012. Meshing interfaces of multi-label data with Delaunay refinement. *Engineering with Computers* 28, 1, 71–82.
- DEY, T. K. AND LEVINE, J. A. 2009. Delaunay meshing of piecewise smooth complexes without expensive predicates. *Algorithms* 2, 4, 1327–1349.
- DEY, T. K., LEVINE, J. A., AND SLATTON, A. 2010. Localized Delaunay refinement for sampling and meshing. *Computer Graphics Forum (Special Issue of Eurographics SGP)* 29, 5, 1723–1732.
- DU, Q., FABER, V., AND GUNZBURGER, M. 1999. Centroidal Voronoi tessellations : Applications and algorithms. *SIAM Review* 41, 637–676.
- DU, Q. AND WANG, D. 2003. Tetrahedral mesh generation and optimization based on centroidal Voronoi tessellations. *International Journal for Numerical Methods in Engineering* 56, 9, 1355–1373.
- FREITAG, L. A. AND PLASSMANN, P. 2000. Local optimization-based simplicial mesh untangling and improvement. *International Journal for Numerical Methods in Engineering* 49, 1-2, 109–125.
- FREY, P. J. AND GEORGE, P.-L. 2007. *Mesh Generation: Application to Finite Elements*. ISTE.
- GEORGE, P.-L. 2005. *Tetmesh-GHS3D, Tetrahedral Mesh Generator*. INRIA - Simulog-Technologies.
- GOSSELIN, S. AND OLLIVIER-GOOCH, C. 2011. Constructing constrained Delaunay tetrahedralizations of volumes bounded by piecewise smooth surfaces. *International Journal of Computational Geometry & Applications* 21, 05, 571–594.
- KLINGNER, B. M. AND SHEWCHUK, J. R. 2008. Aggressive tetrahedral mesh improvement. In *Proceedings of the 16th International Meshing Roundtable*. 3–23.
- LABELLE, F. AND SHEWCHUK, J. R. 2007. Isosurface stuffing: Fast tetrahedral meshes with good dihedral angles. *ACM Trans. Graph.* 26, 3, 57.1–57.10.
- LI, X. Y. 2000. Spacing control and sliver-free Delaunay mesh. In *Proc. 9th Int. Meshing Round Table*.
- LORENSEN, W. E. AND CLINE, H. E. 1998. Seminal graphics. ACM, New York, NY, USA, Chapter Marching cubes: a high resolution 3D surface construction algorithm, 347–353.
- MARCUM, D. L. 1998. Advancing-front/local-reconnection (aflr) unstructured grid generation. *Computational Fluid Dynamics Review* 1, 140–157.
- MULLEN, P., DE GOES, F., DESBRUN, M., COHEN-STEINER, D., AND ALLIEZ, P. 2010. Signing the unsigned: Robust surface reconstruction from raw pointsets. *Computer Graphics Forum* 29, 1733–1741. Special issue 7th Annu. Sympos. Geometry Processing.

- NAVE, D., CHRISOCHOIDES, N., AND CHEW, L. P. 2004. Guaranteed-quality parallel Delaunay refinement for restricted polyhedral domains. *Comput. Geom. Theory Appl.* 28, 2-3, 191–215.
- OLLIVIER-GOOCH, C. 2010. *GRUMMP User Guide*. Department of Mechanical Engineering, University of British Columbia.
- OUDOT, S., RINEAU, L., AND YVINEC, M. 2005. Meshing volumes bounded by smooth surfaces. In *Proc. of the 14th Int. Meshing Roundtable*. 203–219.
- PARK, J. AND SHONTZ, S. M. 2010. Two derivative-free optimization algorithms for mesh quality improvement. *Procedia Computer Science* 1, 1, 387 – 396. ICCS 2010.
- PONS, J.-P. AND BOISSONNAT, J.-D. 2007. Delaunay deformable models: Topology-adaptive meshes based on the restricted Delaunay triangulation. In *IEEE Conf. on Computer Vision and Pattern Recognition*.
- PONS, J.-P., SÉGONNE, F., BOISSONNAT, J.-D., RINEAU, L., YVINEC, M., AND KERIVEN, R. 2007. High-quality consistent meshing of multi-label datasets. In *Information Processing in Medical Imaging*. 198–210.
- RINEAU, L. AND YVINEC, M. 2007a. A generic software design for Delaunay refinement meshing. *Comput. Geom. Theory Appl.* 38, 100–110.
- RINEAU, L. AND YVINEC, M. 2007b. Meshing 3D domains bounded by piecewise smooth surfaces. In *Proc. of the 16th Int. Meshing Roundtable*. 443–460.
- RUPPERT, J. 1995. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms* 18, 3, 548–585.
- SALMAN, N., YVINEC, M., AND MÉRIGOT, Q. 2010. Feature preserving mesh generation from 3D point clouds. *Computer Graphics Forum* 29, 1623–1632. Special issue for EUROGRAPHICS Symp. on Geometry Processing.
- SASTRY, S., SHONTZ, S., AND VAVASIS, S. 2012. A log-barrier method for mesh quality improvement and untangling. *Engineering with Computers*, 1–15.
- SCHÖBERL, J. 1997. NETGEN an advancing front 2D/3D-mesh generator based on abstract rules. *Computing and Visualization in Science* 1, 41–52.
- SCHÖBERL, J. 2003. *NETGEN - 4.3*. Johannes Kepler University Linz.
- SHEWCHUK, J. 1998. Tetrahedral mesh generation by Delaunay refinement. In *Proc. of the Fourteenth Annual Symp. on Comp. Geometry*. 86–95.
- SHEWCHUK, J. 2012. *Combinatorial Scientific Computing*. CRC Press, Chapter 10: Unstructured Mesh Generation, 257–297.
- SI, H. 2006. *TetGen A Quality Tetrahedral Mesh Generator and Three-Dimensional Delaunay Triangulator*. Weierstrass Institute for Applied Analysis and Stochastics (WIAS), Mohrenstr. 39, 10117 Berlin, Germany.
- TOURNOIS, J., WORMSER, C., ALLIEZ, P., AND DESBRUN, M. 2009. Interleaving Delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. *ACM Transactions on Graphics* 28(3), 75:1–75:9. SIGGRAPH '2009 Conference Proc.
- VAVASIS, S. A. 1999. QMG 2.0 reference manual.

Received June 2014; revised ; accepted