



**HAL**  
open science

# CGALmesh: a Generic Framework for Delaunay Mesh Generation

Clément Jamin, Pierre Alliez, Mariette Yvinec, Jean-Daniel Boissonnat

► **To cite this version:**

Clément Jamin, Pierre Alliez, Mariette Yvinec, Jean-Daniel Boissonnat. CGALmesh: a Generic Framework for Delaunay Mesh Generation. [Research Report] RR-8256, 2013. hal-00796052v1

**HAL Id: hal-00796052**

**<https://inria.hal.science/hal-00796052v1>**

Submitted on 1 Mar 2013 (v1), last revised 27 Jan 2014 (v2)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



# CGALmesh: a Generic Framework for Delaunay Mesh Generation

Clément Jamin, Pierre Alliez, Mariette Yvinec, Jean-Daniel  
Boissonnat

**RESEARCH  
REPORT**

**N° 8256**

February 2013

Project-Teams Geometrica





## CGALmesh: a Generic Framework for Delaunay Mesh Generation

Clément Jamin<sup>\*†</sup>, Pierre Alliez<sup>\*</sup>, Mariette Yvinec<sup>\*</sup>, Jean-Daniel  
Boissonnat<sup>\*</sup>

Project-Teams Geometrica

Research Report n° 8256 — February 2013 — 28 pages

**Abstract:** 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 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. 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.

**Key-words:** Mesh generation, Delaunay refinement, mesh optimization, isosurface extraction

---

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

<sup>\*</sup> Inria

<sup>†</sup> Université Lyon 1, LIRIS, UMR5205

**RESEARCH CENTRE  
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93  
06902 Sophia Antipolis Cedex

## CGALmesh : un cadre générique pour la génération de maillages de Delaunay

**Résumé :** CGALmesh est le composant logiciel de génération de maillages de la bibliothèque CGAL (Computational Geometry Algorithm Library). Ce composant génère des maillages simpliciaux isotropes – maillages de surface triangulaires ou maillages volumiques tétraédriques – à partir de domaines à mailler fournis en entrée. Ces domaines peuvent être des surfaces ou des domaines 3D, avec ou sans arêtes vives. L’algorithme s’appuie sur la notion de triangulation de Delaunay restreinte pour approcher les domaines et les surfaces, et sur le raffinement de Delaunay pour fournir une approximation précise et s’assurer de la qualité des maillages produits. CGALmesh offre des garanties sur la qualité d’approximation, ainsi que sur la taille et la forme des éléments du maillage. Quatre algorithmes d’optimisation de maillage sont proposés pour améliorer la qualité du maillage. Une propriété distinctive de CGALmesh est sa grande flexibilité quant aux domaines fournis en entrée. Cette flexibilité est obtenue grâce à une conception logicielle générique qui rassemble en un seul concept abstrait, appelé oracle, toute l’interface nécessaire entre le maillage et le domaine fourni en entrée. Des oracles pour les domaines définis par des polyèdres et par des surfaces implicites sont fournis avec la bibliothèque.

**Mots-clés :** Génération de maillages, raffinement de Delaunay, optimisation de maillages, extraction d’isosurfaces

# 1 Introduction

Surface triangle 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 [36] and spatio-temporal coherent segmentation of time-varying data sets such as beating hearts. As for images, meshes are now commonly 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 triangular 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 [41], lattice-based methods [31, 29, 17], and Delaunay refinement. The latter can be further classified into two classes of algorithms. The first class of Delaunay refinement algorithms first 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 [22, 23, 21, 45]. The second class of Delaunay refinement algorithms, the one of CGALmesh, meshes both the input surface and volume in a unified manner, with the ability to mesh multi-domains and curved surfaces while approximating them [35, 38]. For a detailed review of mesh generation methods we refer the reader to comprehensive books and surveys [27, 44, 14].

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 [28]. The open source offer is much more limited, with, e.g., NETGEN from the University of Linz (Advancing Front) [42], GRUMMP from the University of British Columbia (boundary-preserving, Delaunay-based) [34], and TETGEN from the Weierstrass Institute of Berlin [45]. These solutions all share the common requirement that the domain boundary must be provided as a defect-free polyhedral surface mesh or parametric surface, and hence are not very generic in that respect. Furthermore, most of them 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.

## 1.2 Positioning

The main paradigm behind our algorithms is the refinement of a Delaunay triangulation restricted to the input domain and to its subdividing and boundary surfaces. The paradigm of *restricted Delaunay refinement* has received special attention due to its versatility and the existence of 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 [19], then in 3D for polyhedral domains [33], for smooth surfaces [20], for 3D domains bounded by smooth surfaces [35, 8] as well as for 3D domains bounded by piecewise smooth surfaces [38, 15, 16]. One of the main strength of restricted Delaunay refinement is to provide through a single process a mesh approximating all the components of a subdivided domain and of its subdividing and bounding surfaces.

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 [2], sharp features are handled through so-called protecting balls placed before refinement [16] 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 upper bound on the radius-edge ratio of the mesh elements (surface triangles and volume tetrahedra). The radius-edge ratio relates the circumradius of the mesh element to 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 on a 3D domain, the Delaunay refinement process yield a final mesh including 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 [1, 25, 12, 46] to reduce the number of slivers.

The CGALmesh meshing algorithm is not limited to a specific type of input domain. The required interface between the meshing engine and the input domain is encapsulated into an abstract concept denoted by *oracle*. The oracle provides answers to a small number of predicate and construction queries – such as “*does a query point belongs to the input domain?*” or “*construct the intersection point between a query segment and the input surface*”. The meshing algorithm thus queries the input domain – and simultaneously probes it – in an abstract and generic manner during meshing. Dealing with a new type of input only requires implementing a new model of oracle. CGALmesh already provides oracles for domains described through implicit functions, 3D images and polyhedral meshes [2].




Figure 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.



## 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 insert 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  $\Omega$  may be subdivided into several subdomains, each denoted by  $\Omega_i$ .

The *output* to the meshing algorithm is an isotropic simplicial mesh  $\mathcal{M}$  formed by triangle facets (surface mesh) and tetrahedral cells (volume mesh).

Before stepping into the details, we precise 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 Delaunay triangulation and the restrictions of this triangulation to the surface and volume of the input domain. Figure 2, 3 and 4 depict the 2D counterparts of the notion of restricted Delaunay triangulations. More specifically, we denote by:

$\Omega$	The input 3D domain (Figure 3a and 4a);
$\Omega_i$	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 Delaunay triangulation of $P$ (Figure 2c);
$Vor(P)$	Its dual Voronoi diagram (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 $\Omega$ (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 centered on the surface and empty of points of  $P$ . 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  $\Omega$ . In the case of a mono-domain, 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 gallblader).

### 2.2 The Oracle

A model of the oracle concept provides a mean for the meshing algorithm to probe the input domain  $\Omega$  and surface  $S$ .




Figure 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  $Del(P)$ .




Figure 3: (a) Input domain  $\Omega$  (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  $\Omega$ , is formed by triangles (in blue) whose circumcenters (blue dots) lie inside  $\Omega$ .

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  $\Omega$  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  $\Omega$ , and in the positive, in which subdomain  $\Omega_i$ ;
- 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.




Figure 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)  $Del_{|\Omega}(P)$ , the Delaunay triangulation restricted to the domain  $\Omega$ , is formed by triangles (in pink for  $\Omega_1$ , in blue for  $\Omega_2$ ) whose circumcenters (blue dots) lie inside  $\Omega$ .

### 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, ( $\theta$  depicted by Figure 5b).
3. **Approximation:** an upper bound for the approximation error of surface facets ( $\epsilon$  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).

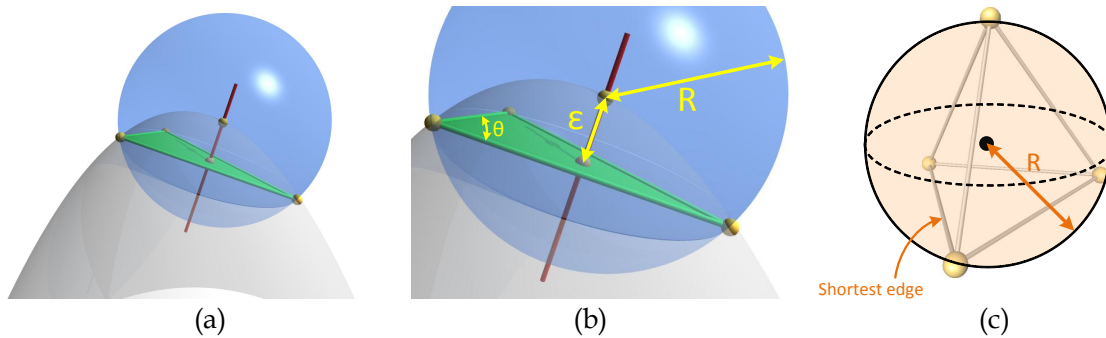


Figure 5: (a) Surface Delaunay ball. (b) Surface Delaunay ball radius ( $R$ ), approximation error of the surface ( $\epsilon$ ) and triangle angles (e.g.  $\theta$ ) are used by the facet criteria. (c) Circumsphere radius ( $R$ ) and shortest tetrahedron edge are used by the cell criteria.

Enforcing such 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 any of the criteria.

**Definition 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 is there are several.*

**Definition 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 any of 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 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 take care to handle bad cells and bad facets through inserting Steiner points either inside  $\Omega$  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  $\Omega$ :

1. Add  $c$  to  $Del_{|\Omega}(P)$ ;
2. Test  $c$  against the cell criteria;
3. If  $c$  does not meet any of 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 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 any of 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. 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  $\Omega$  or on  $S$ . These Steiner points are either generated by probing  $\Omega$  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. [13]. 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.

Figures 6 and 11 depict examples of CAD models meshed without and with protecting the sharp features.




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

## 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  $\geq 2$ , then the above algorithm is shown to terminate after inserting a finite number of Steiner vertices [20, 39, 43].

**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 [26, 5, 10] 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  $\Omega$ ;
- $\text{Hausdorff-distance}(Del_{|S}(P), S) \approx 0$ ;
- $Del_{|S}(P)$  provides a good estimation of normals, area as well as curvature of  $S$ .

Further results [6] 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 [6]. 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 seen 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.




Figure 7: Slivers after Delaunay refinement. The final tetrahedral mesh contains slivers, which are shown for a dihedral angle bound of respectively 5 (middle) and 10 degrees (right).

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 [24, 25] and Optimal Delaunay Triangulation (ODT for short) relaxation [11, 1]. The second category includes vertex perturbation [30, 46] and sliver exudation [12].

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 locations of mesh vertices, and aim at minimizing this energy by relocating all mesh vertices at once.

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 hyperplans 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 [25, 1]. 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.




Figure 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 shaded.

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 treat them individually. The focus is on improving first the worst elements.

The vertex perturber [30, 46] improves the mesh by relocating the vertices in turn, starting with vertices incident to the worst sliver tetrahedra. For each vertex, the algorithm searches for a new location that yields an improvement of the vertex star when the Delaunay mesh connectivity is updated. The new location is first searched so as to increase the incident sliver circumradii. If this fails, the new location is search so as to invert the worst incident tetrahedron. At last if this fails again, random perturbations of the vertex location are tried.

The sliver exuder [12] turns the Delaunay triangulation into a weighted Delaunay triangulation. It tries 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 has been carefully designed and implemented in order to obtain a flexible and reliable algorithm. 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. Any 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. CGALmesh follows this common 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.

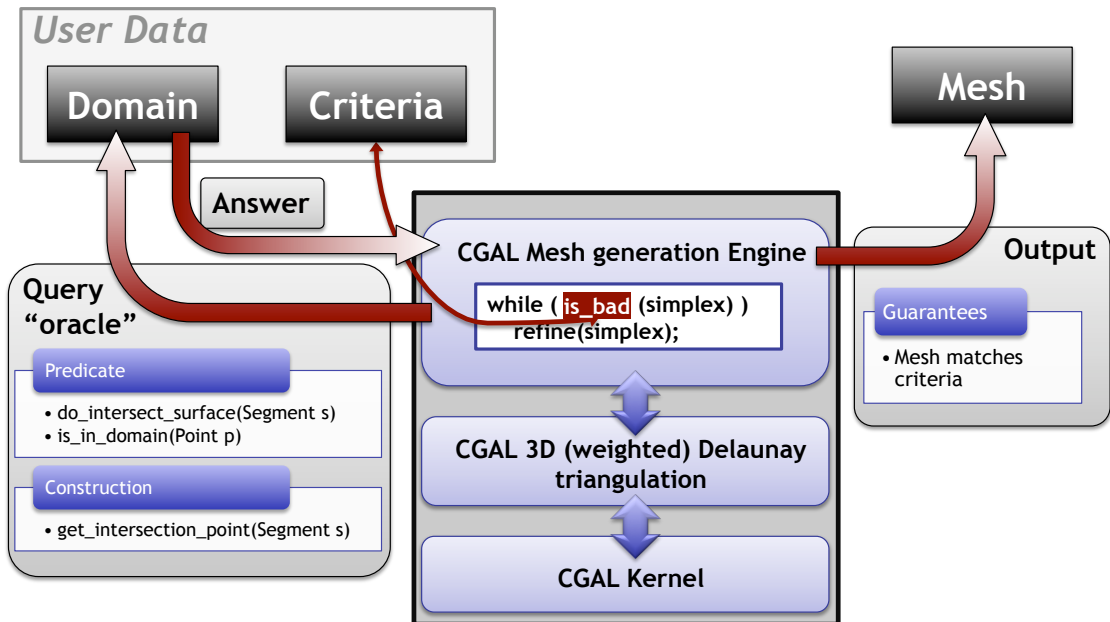


Figure 9: Overall design of the mesh generation process.

Figure 9 depicts the overall modular design of the mesh generation algorithm and the interaction between the distinct modular components. The core mesh generation engine is described in a dedicated publication [37].

#### 3.2 Oracles

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

The polyhedral mesh oracle takes a surface triangle mesh as input. To provide efficient surface intersections during refinement a hierarchical data structure of axis-aligned bounding boxes (AABB tree) is first created [4]. The oracle for implicit functions uses the concept of

bipolar Voronoi edges to detect intersections. A Voronoi edge is said bipolar if the implicit function evaluates to values with opposite signs at its endpoints [6]. Only bipolar Voronoi edges are considered to intersect the surface and the intersection points are obtained through dichotomy. The 3D image oracle computes implicit functions from input 3D images, and relies on the so-constructed implicit function oracle for predicates and constructions.

The complexity of the predicates and constructor required in the oracle greatly varies among the different models of 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. However, 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 known through measurements [40, 32, 3] 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 [9]. 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.

### 3.4 Performance Optimization

One of the main claims of CGAL is to offer efficient algorithms with no compromise on the reliability. CGALmesh has been designed so as to reach satisfactory performances. The algorithms have been run successfully on a series of input domains, with different parameter values.

## 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 2.8 GHz CPU and 32GB RAM.

### 4.1 Mesh Refinement

Table 4.1 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 includes CAD polyhedral surfaces with sharp features. The third set includes domains defined by implicit functions. The fourth set includes 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.

### 4.2 Mesh Optimization

Figure 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 1: Refinement results of the sequential algorithm – domains are depicted by Figure 10

Domain	#V	#F	#C	S.R. <sup>1</sup> (s)	C.S. <sup>2</sup> (s)	V.R. <sup>3</sup> (s)	Total (s)	#V/s
Elephant	536K	127K	3,217K	20.0	5.6	40.2	65.9	8,136
	4,151K	507K	25,703K	81.2	23.6	319.8	424.5	9,778
Bimba	965K	139K	5,942K	18.8	6.4	68.4	93.6	10,311
	4,400K	387K	27,484K	54.3	18.5	334.4	407.2	10,804
Lucy	566K	179K	3,328K	32.7	16.9	52.4	102.0	5,549
	4,350K	737K	26,610K	137.9	71.8	372.1	581.8	7,477
David	749K	232K	4,414K	38.5	16.5	55.1	110.1	6,799
	3,301K	576K	20,166K	92.2	39.8	251.5	383.5	8,607
Ecrou-larousse	853K	381K	4,805K	173.2	78.5	90.9	342.6	2,490
	3,595K	826K	21,620K	339.0	154.9	368.2	862.2	4,169
Pump carter	604K	205K	3,530K	32.3	11.8	45.2	89.4	6,759
	4,582K	789K	28,008K	116.4	45.2	354.6	516.2	8,876
Turbine	451K	206K	2,553K	47.2	12.5	42.6	102.3	4,413
	3,370K	793K	20,259K	169.7	47.8	281.9	499.4	6,748
Cheese	1,044K	502K	5,845K	211.3	107.0	146.3	464.6	2,247
	5,202K	1,442K	30,893K	589.6	324.0	686.1	1,599.7	3,252
Thin cylinder function	316K	406K	1,149K	23.8	0.7	3.7	28.2	11,218
	1,585K	1,665K	6,607K	98.8	3.1	32.6	134.5	11,788
Klein function	754K	639K	3,734K	35.1	1.6	40.4	77.1	9,784
	3,062K	2,014K	16,226K	120.2	4.1	170.1	294.4	10,403
Pancake function	485K	881K	1,645K	52.2	1.9	7.4	61.5	7,894
	4,402K	6,676K	15,896K	420.7	14.6	89.8	525.1	8,383
Tanglecube function	389K	482K	1,664K	19.5	1.1	18.7	39.3	9,896
	1,478K	1,522K	6,905K	63.1	3.6	73.5	140.2	10,539
Liver kidney gallbladder	1,289K	215K	7,915K	15.2	0.9	76.4	92.6	13,928
	4,630K	525K	28,781K	36.4	2.0	306.7	345.2	13,413
Visible human	346k	714K	2,056K	67.3	2.2	32.7	102.1	3,383
	3201k	4,328K	19,799K	394.5	16.7	367.0	778.2	4,113

<sup>1</sup> S.R. = Surface Mesh Refinement <sup>2</sup> C.S. = Cell Scan <sup>3</sup> V.R. = Volume Mesh Refinement

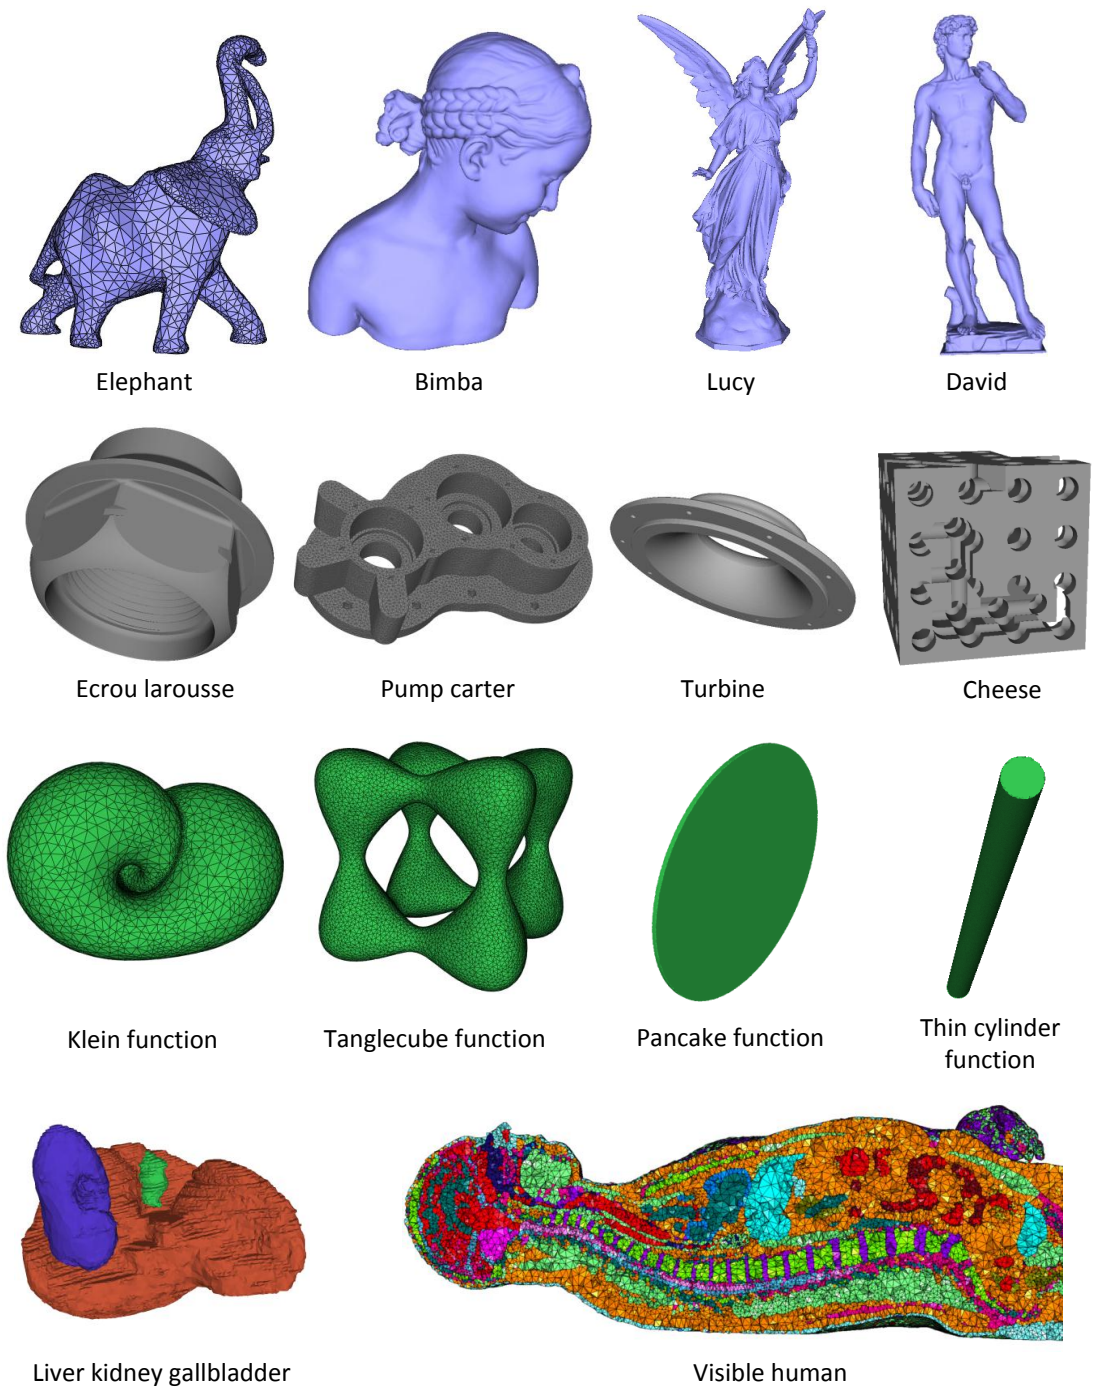


Figure 10: Input domains: domains bounded by smooth surfaces (top row, in blue), CAD models with sharp features (second row, in grey), implicit functions (third row, in green), 3D images (bottom row, multicolored).




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




Figure 12: Refinement times, with respect to the complexity of the output mesh.



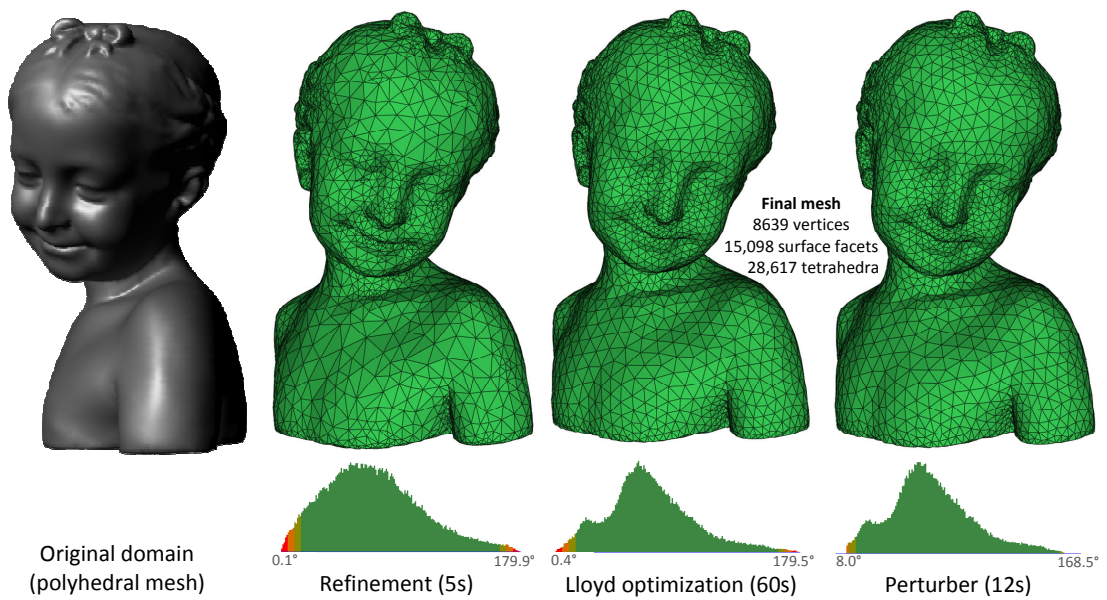


Figure 13: Mesh refinement and optimization on the Bimba model. The histogram of dihedral angles, with both min and max values of these angles, is shown under each mesh.




Figure 14: Mesh refinement and optimization of the Pump\_carter model. The histogram of dihedral angles, with both min and max values of these angles, is shown under each mesh.

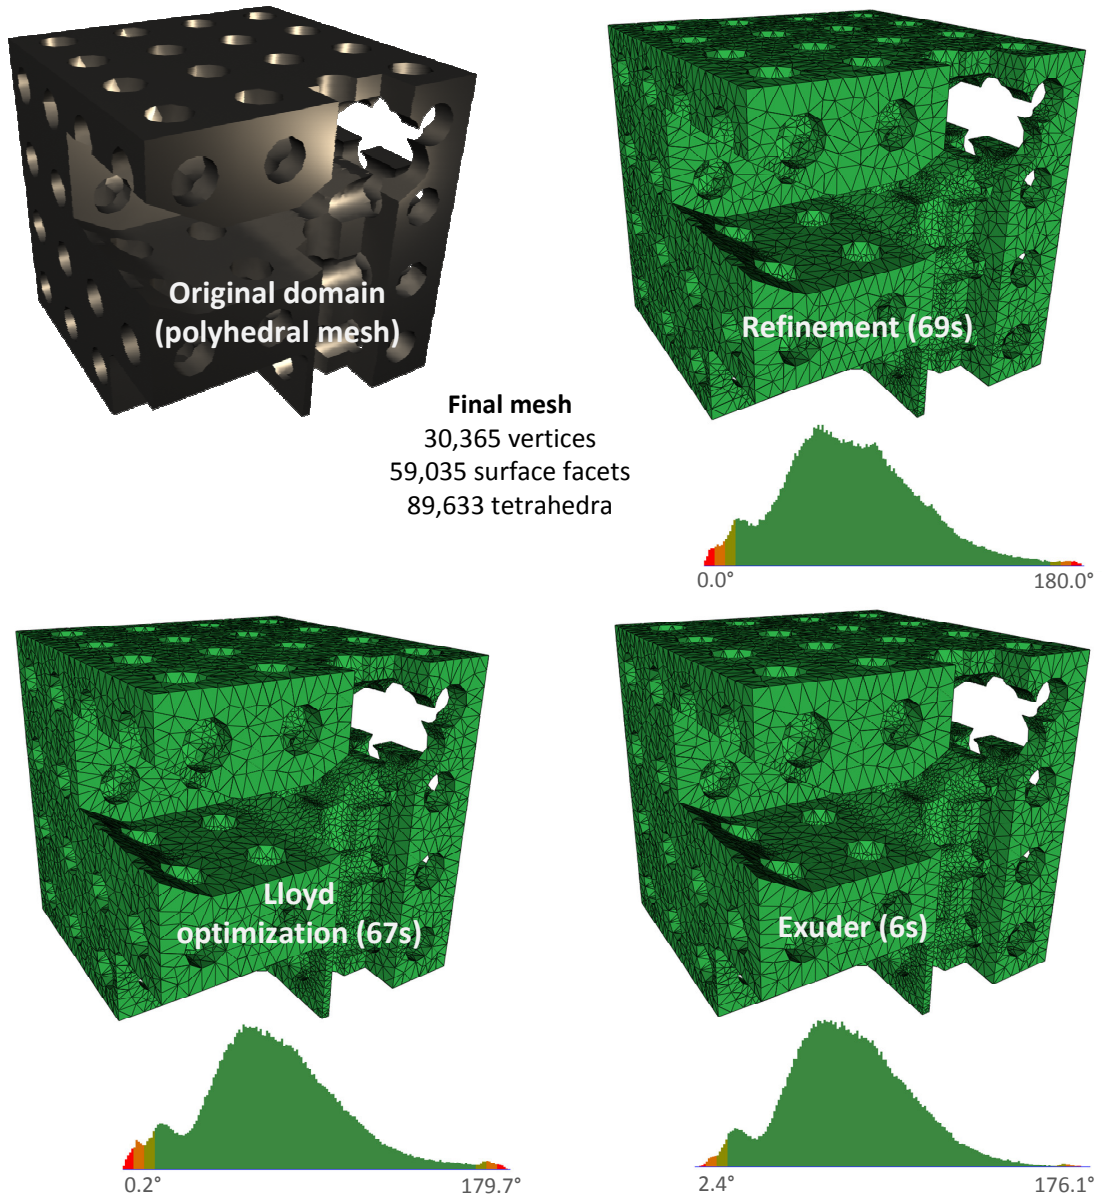


Figure 15: Mesh refinement and optimization of the Cheese model. The histogram of dihedral angles – not taking into account dihedral angles at input creases –, with both min and max values of these angles, is shown under each mesh.

## 5 Conclusion

CGALmesh offers algorithms to generate isotropic surface triangle 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 images<sup>1</sup>.

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 [18] mesh generation is far from being an embarrassingly parallel problem.

Another stimulating future work is the generation of anisotropic simplicial surface and volume meshes [7].

## 6 Acknowledgments

This work was partially funded by the European Research Council (ERC Starting Grant “Robust Geometry Processing”, Grant agreement 257474).

## References

- [1] P. Alliez, D. Cohen-Steiner, Mariette Yvinec, and Mathieu Desbrun. Variational Tetrahedral Meshing. *ACM Transactions on Graphics (SIGGRAPH)*, 24:617–625, 2005.
- [2] Pierre Alliez, Laurent Rineau, Stéphane Tayeb, Jane Tournois, and Mariette Yvinec. 3d mesh generation. in CGAL user and reference manual. CGAL editorial board, 4.1 edition, 2012. [http://www.cgal.org/Manual/latest/doc\\_html/cgal\\_manual/packages.html#Pkg:Mesh\\_3](http://www.cgal.org/Manual/latest/doc_html/cgal_manual/packages.html#Pkg:Mesh_3).
- [3] Pierre Alliez, Laurent Saboret, , and Gael Guennebaud. Surface reconstruction from point sets. in cgal user and reference manual. cgal editorial board, 4.1 edition, 2012. [http://www.cgal.org/Manual/latest/doc\\_html/cgal\\_manual/packages.html#Pkg:SurfaceReconstructionFromPointSets](http://www.cgal.org/Manual/latest/doc_html/cgal_manual/packages.html#Pkg:SurfaceReconstructionFromPointSets).
- [4] Pierre Alliez, Stéphane Tayeb, and Camille Wormser. 3d fast intersection and distance computation (aabb tree). in cgal user and reference manual. cgal editorial board, 4.1 edition, 2012. [http://www.cgal.org/Manual/latest/doc\\_html/cgal\\_manual/packages.html#Pkg:AABB\\_tree](http://www.cgal.org/Manual/latest/doc_html/cgal_manual/packages.html#Pkg:AABB_tree).
- [5] Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. In *Proc. 14th Annu. Sympos. Comput. Geom.*, pages 39–48, 1998.
- [6] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67:405–451, 2005.

---

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

- [7] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Locally uniform anisotropic meshing. In *Proc. 24th Annual Symposium on Computational Geometry*, pages 270–277, 2008.
- [8] C. Boivin and C. Ollivier-Gooch. Guaranteed-quality triangular mesh generation for domains with curved boundaries. *Int. J. Numer. Methods Eng.*, 55:1185–1213, 2002.
- [9] Ricard Campos, Rafael Garcia, Pierre Alliez, and Mariette Yvinec. Splats-based surface reconstruction from defect-laden point sets. To appear.
- [10] F. Chazal, D. Cohen-Steiner, A. Lieutier, and B. Thibert. Stability of Curvature Measures. *Computer Graphics Forum (proc. SGP 2009)*, pages 1485–1496, 2008.
- [11] Long Chen and Jinchao Xu. Optimal Delaunay triangulations. *J. Comput. Math.*, 22(2):299–308, 2004.
- [12] Siu-Wing Cheng, Tamal K. Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. Sliver exudation. *J. ACM*, 47:883–904, 2000.
- [13] Siu-Wing Cheng, Tamal K. Dey, and Edgar A. Ramos. Delaunay refinement for piecewise smooth complexes. In *Proc. 18th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 1096–1105. ACM Press, 2007.
- [14] Siu-Wing Cheng, Tamal K. Dey, and Jonathan R. Shewchuk. *Delaunay Mesh Generation*. CRC Press, 2012.
- [15] S.W. Cheng, T.K. Dey, and J. Levine. A practical delaunay meshing algorithm for a large class of domains. In *Proc. of the 16th Int. Meshing Roundtable*, pages 477–494, 2007.
- [16] S.W. Cheng, T.K. Dey, and E.A. Ramos. Delaunay refinement for piecewise smooth complexes. In *Proceedings of the 18th annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1096–1105, 2007.
- [17] Nuttapon Chentanez, Bryan E. Feldman, François Labelle, James F. O’Brien, and Jonathan R. Shewchuk. Liquid simulation on lattice-based tetrahedral meshes. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation 2007*, pages 219–228, 2007.
- [18] Andrey N. Chernikov and Nikos P. Chrisochoides. A template for developing next generation parallel delaunay refinement methods. *Finite Elements in Analysis and Design*, 46(1-2):96–113, 2010. Mesh Generation - Applications and Adaptation.
- [19] L. P. Chew. Guaranteed-quality triangular meshes. Technical Report TR-89-983, Dept. Comput. Sci., Cornell Univ., Ithaca, NY, April 1989.
- [20] L.P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *SCG '93: Proceedings of the ninth annual symposium on Computational geometry*, pages 274–280. ACM New York, NY, USA, 1993.
- [21] Tamal K. Dey, Firdaus Janoos, and Joshua A. Levine. Meshing interfaces of multi-label data with Delaunay refinement. *Engineering with Computers*, 28(1):71–82, January 2012.
- [22] Tamal K. Dey and Joshua A. Levine. Delaunay meshing of piecewise smooth complexes without expensive predicates. *Algorithms*, 2(4):1327–1349, November 2009.
- [23] Tamal K. Dey, Joshua A. Levine, and A. Slatton. Localized Delaunay refinement for sampling and meshing. *Computer Graphics Forum (Special Issue of Eurographics SGP)*, 29(5):1723–1732, July 2010.

- [24] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal voronoi tessellations : Applications and algorithms. *SIAM Review*, 41:637–676, 1999.
- [25] Qiang Du and Desheng Wang. Tetrahedral mesh generation and optimization based on centroidal Voronoi tessellations. *International Journal for Numerical Methods in Engineering*, 56(9):1355–1373, 2003.
- [26] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.*, 7:365–378, 1997.
- [27] Pascal Jean Frey and Paul-Louis George. *Mesh Generation: Application to Finite Elements*. ISTE, 2007.
- [28] P.-L. George. *Tetmesh-GHS3D, Tetrahedral Mesh Generator*. INRIA - Simulog-Technologies, 2005.
- [29] François Labelle and Jonathan Richard Shewchuk. Isosurface stuffing: fast tetrahedral meshes with good dihedral angles. *ACM Trans. Graph.*, 26(3), July 2007.
- [30] X. Y. Li. Spacing control and sliver-free Delaunay mesh. In *Proc. 9th Int. Meshing Round Table*, 2000.
- [31] William E. Lorensen and Harvey E. Cline. Seminal graphics. chapter Marching cubes: a high resolution 3D surface construction algorithm, pages 347–353. ACM, New York, NY, USA, 1998.
- [32] Patrick Mullen, Fernando de Goes, Mathieu Desbrun, David Cohen-Steiner, and Pierre Alliez. Signing the Unsigned: Robust Surface Reconstruction from Raw Pointsets. *Computer Graphics Forum*, 29:1733–1741, 2010. Special issue 7th Annu. Sympos. Geometry Processing.
- [33] D mian Nave, Nikos Chrisochoides, and L. Paul Chew. Guaranteed-quality parallel delaunay refinement for restricted polyhedral domains. *Comput. Geom. Theory Appl.*, 28(2-3):191–215, June 2004.
- [34] C. Ollivier-Gooch. *GRUMMP User Guide*. Department of Mechanical Engineering, University of British Columbia, 2010.
- [35] S. Oudot, L. Rineau, and M. Yvinec. Meshing volumes bounded by smooth surfaces. In *Proc. of the 14th Int. Meshing Roundtable*, pages 203–219, 2005.
- [36] Jean-Philippe Pons and Jean-Daniel Boissonnat. Delaunay deformable models: Topology-adaptive meshes based on the restricted Delaunay triangulation. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2007.
- [37] L. Rineau and M. Yvinec. A generic software design for Delaunay refinement meshing. *Comput. Geom. Theory Appl.*, 38:100–110, 2007.
- [38] L. Rineau and M. Yvinec. Meshing 3d domains bounded by piecewise smooth surfaces. In *Proc. of the 16th Int. Meshing Roundtable*, pages 443–460, 2007.
- [39] J. Ruppert. A delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms*, 18(3):548–585, 1995.

- [40] Nader Salman, Mariette Yvinec, and Quentin Mérigot. Feature Preserving Mesh Generation from 3D Point Clouds. *Computer Graphics Forum*, 29:1623–1632, 2010. special issue for EUROGRAPHICS Symposium on Geometry Processing.
- [41] Joachim Schöberl. Netgen an advancing front 2d/3d-mesh generator based on abstract rules. *Computing and Visualization in Science*, 1:41–52, 1997.
- [42] Joachim Schöberl. *NETGEN - 4.3*. Johannes Kepler University Linz, 2003.
- [43] J.R. Shewchuk. Tetrahedral mesh generation by delaunay refinement. In *Proc. of the Fourteenth Annual Symp. on Comp. Geometry*, pages 86–95, 1998.
- [44] J.R. Shewchuk. *Combinatorial Scientific Computing*, chapter 10: Unstructured Mesh Generation, pages 257–297. CRC Press, 2012.
- [45] H. Si. *TetGen A Quality Tetrahedral Mesh Generator and Three-Dimensional Delaunay Triangulator*. Weierstrass Institute for Applied Analysis and Stochastics (WIAS), Mohrenstr. 39, 10117 Berlin, Germany, 2006.
- [46] Jane Tournois, Camille Wormser, Pierre Alliez, and Mathieu Desbrun. Interleaving Delaunay refinement and optimization for practical isotropic tetrahedron mesh generation. *ACM Transactions on Graphics*, 28(3):75:1–75:9, 2009. SIGGRAPH '2009 Conference Proceedings.



**RESEARCH CENTRE  
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93  
06902 Sophia Antipolis Cedex

Publisher  
Inria  
Domaine de Voluceau - Rocquencourt  
BP 105 - 78153 Le Chesnay Cedex  
[inria.fr](http://inria.fr)

ISSN 0249-6399