
4D Space-Time Delaunay Meshing for Medical Images

Panagiotis Foteinos^{1,2,†} and Nikos Chrisochoides^{2,§}

¹ *Computer Science Department, College of William and Mary, Virginia, USA*

² *Computer Science Department, Old Dominion University, Virginia, USA*

Summary. In this paper, we present a Delaunay refinement algorithm for 4-dimensional (3D+t) segmented images. The output mesh is proved to consist of sliver-free simplices. Assuming that the hyper-surface is a closed smooth manifold, we also guarantee faithful geometric and topological approximation. We implement and demonstrate the effectiveness of our method on publicly available segmented cardiac images.

Key words: Delaunay refinement, space-time, quality, fidelity

1 Introduction

Technological advances in imaging have made the acquisition of 4D medical images feasible [35, 45, 46, 48]. At the same time, pentatope capable FEM solvers [7, 36] operating directly on 4D data have been shown to be effective for advection-diffusion and Navier-Stokes formulations.

In this paper, we describe a 4-dimensional Delaunay mesh algorithm which operates directly on a 4-dimensional image \mathcal{I} . \mathcal{I} represents the domain Ω to be meshed as the temporal evolution of a 3D object. That is, $\Omega = \bigcup_{t_i} \Omega_{t_i}$,

where Ω_{t_i} is the 3D object at time t_i (i.e., the i^{th} slice of Ω).

Volume mesh generation methods can be divided into two categories: *PLC-based* and *Isosurface-based*. The PLC-based methods assume that the *surface* $\partial\Omega$ of the volume Ω (about to be meshed) is given as a *Piecewise Linear Complex* (PLC) which contains linear sub-faces embedded in 3 or 4 dimensions [13, 14, 16, 17, 30, 33, 34, 40, 43]. The limitation of this method is that the success of meshing depends on the quality of the given PLC: if the PLC forms very small angles, then the overall mesh quality deteriorates and termination

[†]pfot@cs.wm.edu

[§]nikos@cs.odu.edu

might be compromised [40, 42]. In Computed Aided Design (CAD) applications, the surface is usually given as a PLC. In biomedical Computer Aided Simulations (CAS), however, there is no reason to use this approach, since it might (depending on the geometry of the image, its segmentation, and its decimation [16]) add the additional small input angle limitation. A workaround for this small input angle limitation is to treat the surface voxels as the PLC of the domain, since those input facets meet at large angles (90 or 180 degrees). This, however, would introduce an unnecessary large number of elements and little control over the density of the domain.

The Isosurface-based methods assume that Ω is known through a function $f : \mathcal{R}^d \rightarrow \mathcal{R}$, such that points in different regions of interest evaluate f differently. This assumption covers a wide range of inputs used in modeling and simulation, such as parametric surfaces/volumes [37], level-sets and segmented multi-labeled images [10, 29, 38]. Of course, these type of functions can also represent PLCs [29], a fact that makes the Isosurface-based method a general approach. Isosurface-based methods ought to recover and mesh both the *isosurface* $\partial\Omega$ and the volume. This method does not suffer from any unnecessary small input angle artifacts introduced by the initial conversion to PLCs, since $\partial\Omega$ is recovered and meshed during refinement.

In this paper, we describe a space-time Delaunay Isosurface-based meshing technique for 4 dimensions. We show that the resulting mesh is sliver free consisting of pentatopes whose boundary is a correct approximation of the underlying isosurface $\partial\Omega = \bigcup_{t_i} \partial\Omega_{t_i}$. Note that space-time meshing is different from dynamic surface simulations (see [28] and the references therein for example). In those simulations, the isosurface is not known; instead, a tetrahedral mesh is adapted on each time step that describes accurately the free surface dynamics.

One way to solve the space-time 4D problem is to mesh separately each 3D object Ω_{t_i} and then connect the elements between two consecutive objects to obtain space-time elements. However, finding such correspondence—which also has to satisfy the quality criteria—is not intuitive, especially when the topology and the geometry of the two objects varies drastically. Alternatively, one could mesh a single object Ω_{t_i} and then deform the mesh to match the shape of the other temporal instances. The limitation of this approach is twofold. First, the quality of the deformed mesh might be much worse than the original; second, there is no control over the mesh density across both the spatial and the temporal direction [7], since the mesh size of the original instance determines the size of the rest of the instances.

Space-time meshing methods have already been proposed in the literature [24, 44]. They assume, however, that the evolving object Ω_{t_i} has the same spatial space across time. Furthermore, the implementation of these techniques is confined to only the 2D+t case (i.e., the space-time elements are tetrahedra). The more general 3D+t meshing has been the focus in [7, 36], but they consider only convex hyper-surfaces such as hyper-cubes or hyper-

cylinders. To our knowledge, the method presented in this paper is the first to address the 3D+t problem where the topology and the geometry of the evolving object may differ substantially through time, and hence, it is allowed to form complex hyper-surfaces.

In the literature [3, 6, 8, 12, 14, 15], it is shown that given a sufficiently dense sample on a surface $\partial\Omega$, the restriction of its Delaunay triangulation to $\partial\Omega$ is a topologically good approximation, or, alternatively, it satisfies the closed-topological-ball property [23]. Their focus, however, was not on volume meshing, but rather, on surface reconstruction. In this paper, we fill the space-time volume Ω with sliver-free pentatopes, such that $\partial\Omega$ is approximated correctly.

Computing the appropriate sample of the surface is a challenging task. In the literature, however, it is assumed that either such a sample is known [3–5] or that an initial sparse sample is given [9, 37, 39]. In this paper, we propose a method that starts directly from labelled images (of one or many more connected components) and computes the appropriate sample on the fly, respecting at the same time the quality and fidelity guarantees.

Our algorithm guarantees that the resulted pentatopes are of bounded aspect ratio. We achieve that by generating elements of low radius-edge ratio and by proving the absence of slivers. We clean the mesh from slivers by integrating into our framework the theory presented in [30]. In [30], the surface is given as an already meshed polyhedral domain (i.e., the method in [30] is a PLC-based method), a different problem than ours, since it is our algorithm’s responsibility to mesh both the underlying zero-surfaces and the bounded volume with topological and geometric guarantees.

The rest of the paper is organized as follows: Section 2 introduces some basic terminology, and in Section 3 we present our algorithm. In Section 4 and Section 5, we prove the guarantees. Section 6 evaluates our method on segmented 4D cardiac data and Section 7 concludes the paper.

2 Preliminaries

The input of our algorithm is a segmented n dimensional image $\mathcal{I} \subset \mathcal{R}^n$. The object $\Omega \subseteq \mathcal{I}$ is assumed to be represented as a cut function $f : \mathcal{R}^n \mapsto \mathcal{R}$, such that its surface $\partial\Omega$ is defined by the set $\{f(p) = 0\}$ [29, 37]. Clearly, from a segmented image, the zero-surface $\{f(p) = 0\}$ can be easily computed by interpolating the voxel values.

We assume that given a point $p \in \mathcal{R}^4$, we can ask for p ’s closest point on $\partial\Omega$. This can be accomplished by an Euclidean Distance Transform (EDT) [20, 32]. Specifically, the EDT returns the voxel $p' \in \partial\Omega$ which is closest to p . Then, we traverse the ray pp' and we compute the intersection between the ray and $\partial\Omega$ by interpolating the positions of different signs [31]. Points on $\partial\Omega$ are referred to as *feature* points.




Fig. 1: A 2D illustration. The simplex $\sigma = \{v, w\}$ and its surface ball $\mathcal{B}_{z,\sigma}$. m is the midpoint of σ . Observe that since the radius $\mathcal{R}_{z,\sigma}$ of $\mathcal{B}_{z,\sigma}$ is larger than the radius $R_\sigma = |m - v|$ of \mathcal{B}_σ , the picking region of σ as defined here is larger than the picking region of [30].

The *local feature size* $\text{lfs}_{\partial\Omega}(x)$ of a point $x \in \partial\Omega$ is defined as the (closest) distance between x and the medial axis of $\partial\Omega$. Since $\partial\Omega$ is smooth, the local feature size is bounded from below by a positive constant $\text{lfs}_{\partial\Omega}$, that is, $\text{lfs}_{\partial\Omega}(x) > \text{lfs}_{\partial\Omega} > 0$. Another useful property is that the local feature size is 1-Lipschitz, that is,

$$\text{lfs}_{\partial\Omega}(p) \leq |p - q| + \text{lfs}_{\partial\Omega}(q). \quad (1)$$

A point set $V \subset \partial\Omega$ is called an ε -sample, if for every point $p \in \partial\Omega$ there is a point $v \in V$ at a distance at most $\varepsilon \cdot \text{lfs}_{\partial\Omega}(p)$ from p [4].

Let V be a finite set of vertices $V = \{v_1, \dots, v_N\} \subset \mathcal{R}^n$. The *Delaunay triangulation* of V is denoted by $\mathcal{D}(V)$. A k -simplex $\sigma_k = \{v_1, \dots, v_{k+1}\} \in \mathcal{D}(V)$ is a simplex defined by $k + 1$ vertices. We denote the length of the shortest edge of a simplex σ with L_σ . The *circumball* \mathcal{B}_σ of a simplex σ is the smallest closed ball circumscribing σ 's vertices. R_σ is the circumradius length of the simplex and C_σ is its circumcenter. The radius-edge ratio of a simplex σ is defined as $\rho_\sigma = \frac{R_\sigma}{L_\sigma}$.

The *voronoi cell* $\mathcal{V}(v)$ of a vertex $v \in V$ is the set $\mathcal{V}(v) = \{p \in \mathcal{R}^n \mid |v - p| \leq |q - p|, \forall q \in V\}$. The *voronoi dual* of a simplex $\sigma \in \mathcal{D}(V)$ is defined as the set $\mathcal{V}(\sigma) = \{\mathcal{V}(v_i) \cap \mathcal{V}(v_j) \mid \forall v_i, v_j \in \sigma\}$.

The restriction of $\mathcal{D}(V)$ to a topological space \mathcal{X} is denoted by $\mathcal{D}_{|\mathcal{X}}(V)$. $\mathcal{D}_{|\mathcal{X}}(V)$ is a *simplicial complex* (as is $\mathcal{D}(V)$) that contains simplices of $\mathcal{D}(V)$ whose voronoi dual intersects \mathcal{X} in a non-empty set. Consider a k simplex σ and let $\mathcal{V}(\sigma)$ intersect \mathcal{X} at a point z . Any ball centered at z circumscribing σ is called a *surface ball* [9]. The corresponding surface ball is denoted by $\mathcal{B}_{z,\sigma}$ and its radius by $\mathcal{R}_{z,\sigma}$, in the sequel. By the definition of Voronoi diagrams, $\mathcal{B}_{z,\sigma}$ does not contain any vertex of V in its interior.

Following the definition of [30], the metric we use to characterize the quality of a simplex σ_k is $\tau_{\sigma_k} = \frac{\text{Vol}_{\sigma_k}}{L_{\sigma_k}^k}$. Low values of τ imply a poor-quality element.

Definition 1 (Sliver [30]). *Simplex σ is a sliver if it contains a k -simplex σ_k ($k \leq 4$) such that $\rho_{\sigma_k} < \bar{\rho}$, $\tau_{\sigma_k} < \bar{\tau}$, and for any m -simplex σ_m of σ ($m < k$), $\rho_{\sigma_m} < \bar{\rho}$, $\tau_{\sigma_m} \geq \bar{\tau}$.*

The *picking region* $\mathcal{PR}(\sigma_4)$ of a 4-simplex σ_4 is defined as the 4-dimensional solid ball centered at C_{σ_4} with radius ζR_{σ_4} , $\zeta < 1$. Consider a restricted k -simplex σ_k and its surface ball $\mathcal{B}_{z,\sigma}$ $k < 4$. Its picking region $\mathcal{PR}(\sigma_k)$ is the intersection between $\partial\Omega$ and the 4-dimensional solid ball centered at z with radius $\zeta R_{z,\sigma}$, $\zeta < 1$. Note that $\mathcal{PR}(\sigma_4)$ and $\mathcal{PR}(\sigma_k)$ are contained in \mathcal{B}_σ and $\mathcal{B}_{z,\sigma}$, respectively. Observe that the picking region of σ_k ($k < 4$) is a topological k -ball and does not belong (necessarily) in the affine k dimensional space defined by σ_k . This is different than the definition in [30], where the picking regions are defined inside the intersection of \mathcal{B}_σ with the affine space of σ . The reason for this change is the fact that the input of our algorithm is not a *Piecewise Linear Complex* (PLC) but a cut function.

A good point $p \in \mathcal{PR}(\sigma)$ is a point that does not introduce smaller slivers. A sliver is small when its radius is less than bR_σ . In [30], it is proved that (a) the number of small slivers $S(\sigma)$ possibly created after the insertion of p is constant, and (b) the volume $|F_\sigma|$ (the *forbidden region*) inside which p forms a small sliver is bounded from above. The same findings hold in our case too, where the picking region of a restricted facet σ_3 is not inside the intersection of \mathcal{B}_{σ_3} and σ_3 's affine space, but inside the intersection of \mathcal{B}_{z,σ_3} and $\partial\Omega$.

Lemma 1. *Given an almost-good mesh, a point p inside the picking region of a σ_k can be found in a constant number of random rounds, such that any new sliver created after the insertion of p has circumradius no smaller than bR_{σ_k} if $k = 4$, or no smaller than bR_{z,σ_k} if $k = 3$.*

Remark 1. The proof is similar to [30], since $|F_\sigma|$ and $S(r)$ do not change and the volume of the intersection of \mathcal{B}_{σ_3} and σ_3 's affine space is smaller than the intersection of \mathcal{B}_{z,σ_3} and $\partial\Omega$. See Figure 1 for an illustration.

3 Algorithm

The user specifies a parameter δ . It will be clear in Section 5 that the lower δ is, the better the mesh boundary will approximate $\partial\Omega$. For brevity, the quantity $\delta \cdot \text{lfs}_{\partial\Omega}(z)$ is denoted by $\Delta_{\partial\Omega}(z)$, where z is a feature point.

Our algorithm initially inserts the 16 corners of a hyper-box that contains the 4 dimensional object Ω , such that the distance between a box corner x and its closest feature point $z = \text{cfp}_{\partial\Omega}(x)$ is at least $2\Delta_{\partial\Omega}(z)$. After the computation of this initial triangulation, the refinement starts dictating which

extra points are inserted. At any time, the Delaunay triangulation $\mathcal{D}(V)$ of the current vertices V is maintained. Note that by construction, $\mathcal{D}(V)$ always covers the entire hyper-volume and that any point on the box is separated from $\partial\Omega$ by a distance at least $2\Delta_{\partial\Omega}(z)$, where z is a feature point.

During the refinement, some vertices are inserted exactly on the box; these vertices are called *box vertices*. The box vertices might lie on 1, 2, or 3-dimensional box faces. We shall refer to the vertices that are neither box vertices nor feature vertices as *free vertices*.

The algorithm inserts new vertices for three reasons: to guarantee that (a) $\partial\Omega$ is correctly recovered, (b) all the elements have small radius-edge ratio, and (c) there are no slivers. Specifically, for a 4-simplex σ_4 in the mesh, the following rules are checked in this order:

- **R1:** Let \mathcal{B}_{σ_4} intersect $\partial\Omega$ and z be equal to $\text{cfp}_{\partial\Omega}(C_{\sigma_4})$. If z is at a distance no closer than $\Delta_{\partial\Omega}(z)$ to any other feature vertex, then z is inserted.
- **R2:** Let \mathcal{B}_{σ_4} intersect $\partial\Omega$ and z be equal to $\text{cfp}_{\partial\Omega}(C_{\sigma_4})$. If $R_{\sigma} \geq 2\Delta_{\partial\Omega}(z)$, C_{σ_4} is inserted.
- **R3:** Let C_{σ_4} lie inside Ω . If $\rho_{\sigma_4} \geq \bar{\rho}$, C_{σ_4} is inserted.
- **R4:** Let C_{σ_4} lie inside Ω . If σ_4 contains a sliver, a good point inside $\mathcal{PR}(\sigma_4)$ is inserted.
- **R5:** Let σ_3 ($\sigma_3 \subset \sigma_4$) be a restricted facet. If the vertices of σ_3 are not feature vertices, then a good point z inside $\mathcal{PR}(\sigma_3)$ is inserted. All the free vertices closer than $\Delta_{\partial\Omega}(z)$ to z are deleted.

For $i < j$, priority is given to R_i over R_j . That is, right before the insertion of a point because of R_j , there is no element that violates a rule R_i . Also, in R_4 , priority is given to the lower dimensional slivers that σ_4 might contain.

Whenever there is no simplex for which R_1 , R_2 , R_3 , or R_4 apply, the refinement process terminates. *The final mesh reported is the set of pentatopes whose circumcenters lie inside Ω .*

In a nutshell, R_1 and R_2 is responsible for generating a sufficiently dense sample on $\partial\Omega$. R_5 makes sure that the vertices of the simplices restricted to $\partial\Omega$ lie on $\partial\Omega$ similarly to [37]. Lastly, R_3 and R_4 deal with the quality guarantees. In Section 4, we will show that there are values for b , ζ , and $\bar{\rho}$ that do not compromise termination.

To prove termination, no vertices should be inserted outside the bounding box. Notice, however, that vertices inserted due to R_2 may lie outside the bounding box. To deal with such cases, C_{σ_4} is rejected for insertion. Instead, its *projection* C'_{σ_4} on the box is inserted in the triangulation. That is, C'_{σ_4} is the closest to C_{σ_4} box point. In Section 4 and Section 5, we prove that the insertion of projected points do not compromise either quality or fidelity. Note that these projections are different than the traditional *encroachment rules* described in [40, 41].

Recall that pentatopes with circumcenters on $\partial\Omega$ or outside Ω are not part of the final mesh, and that is why rule R_3 and R_4 do not check them.




Fig. 2: Proof of Lemma 3, a 3D illustration.

4 Termination and Quality

In this section, we will specify the appropriate values for ζ , $\bar{\rho}$, and b , so that the algorithm terminates. Specifically, we will show during refinement the shortest edge introduced into the mesh cannot be arbitrarily small.

Suppose that σ violates a rule R_i . σ is called an R_i element. R_i dictates the insertion of a point p (and possibly the removal of free points). Point p is called an R_i point. According to [40, 41], the *insertion radius* R_p of p is defined as the length of the shortest edge incident to p created right after the end of R_i and the *parent* $P(p)$ of p as the most recently inserted vertex incident to the shortest edge of σ .

Lemma 2. *Let p and q define the shortest edge of a simplex σ and q being inserted after p . Then $R_q \leq L_\sigma$.*

Proof. Assume that right after the insertion of p , p is the closest point to q . In this case, $R_q = |p - q| = L_\sigma$. Otherwise, there has to be another closest vertex to q , which implies that $R_q < |p - q| = L_\sigma$.

The following Lemma bounds from below the shortest edge introduced into the mesh after the insertion of a box vertex:

Lemma 3. *Let v be a box vertex inserted into the mesh. Then, $R_v \geq 2\Delta_{\partial\Omega}(z)$, where z is a feature point.*

Proof. A box point v is inserted only because of R2. The circumcenter C_σ of a pentatope σ lies on or outside the box and its projection $C'_\sigma = v$ falls on the box. Without loss of generality, assume that the projection lies on the interior of a 3-face (i.e. a box tetrahedron) F . See Figure 2 for a 3D illustration. (If C_σ lies precisely on the box, C'_σ is equal to v .) Consider the (2D) disk (drawn) of B_σ which is coplanar with F . That disk contains v and separates B_σ in two sides: the side that contains C_σ and the side that contains a part of the box.

We claim that the closest vertex —say w — to v lies on the intersection of \mathcal{B}_σ 's boundary and the ray $\overrightarrow{C_\sigma v}$. To see why, note that \mathcal{B}_σ is empty of vertices, and therefore, the closest to v that an arbitrary vertex w' already in the triangulation can be is when it lies on the boundary of \mathcal{B}_σ and on the side of \mathcal{B}_σ that contains a part of the box, as shown. Consider the triangle $\triangle w'vC_\sigma$. From the law of cosines, we have that:

$$\begin{aligned} |v - w'|^2 &= |C_\sigma - w'|^2 + |C_\sigma - v|^2 - 2|C_\sigma - w'| |C_\sigma - v| \cos \omega \\ &\geq |C_\sigma - w'|^2 + |C_\sigma - v|^2 - 2|C_\sigma - w'| |C_\sigma - v| \quad \text{since } \cos \omega \leq 1 \\ &= (|C_\sigma - w'| - |C_\sigma - v|)^2 \\ &= (R_\sigma - |C_\sigma - v|)^2 \quad \text{since } w' \text{ lies on the sphere} \\ &= |v - w|^2, \end{aligned}$$

and our claim is proved.

Therefore, any possible new edge connected to v has length at least $|v - w|$. Since, however, σ triggers R2, \mathcal{B}_σ has to intersect $\partial\Omega$. Therefore, there has to be a feature point $q \in \partial\Omega$ (illustrated) inside \mathcal{B}_σ and on the same side of F as w . Let us denote with q' , the projection of q to the box face F . By construction, $|q - q'|$ is at least $2\Delta_{\partial\Omega}(z)$, where z is a feature point. Observe, however, that $|v - w|$ is always larger than $|q - q'|$, because $vw \parallel qq'$, and the statement holds. Similar reasoning applies in the case where C'_σ lies on a box triangle or a box edge.

The following Lemma proves a lower bound on the lengths created into the mesh because of R1 and R2:

Lemma 4. *Let p be a vertex inserted into the mesh because of R1 or R2. Then, $R_p \geq \Delta_{\partial\Omega}(z)$, where z is the closest feature point to p .*

Proof. If R1 is triggered, then p is equal to z and since there is no other feature point already inserted in the mesh closer than $\Delta_{\partial\Omega}(p)$ to p , the statement holds. Otherwise, R2 applies for a simplex σ_4 and p is equal to C_{σ_4} . Due to the empty ball property, R_p is at least $R_{\sigma_4} \geq 2\Delta_{\partial\Omega}(\text{cfp}_{\partial\Omega}(p))$, and the statement holds.

Lemma 5. *Let p be a vertex inserted into the mesh because R3 applies for an element σ . Then, $R_p \geq \bar{\rho}R_{P(p)}$.*

Proof. Since p is equal to C_σ , $R_p \geq R_\sigma = \rho_\sigma L_\sigma \geq \bar{\rho}L_\sigma$. Lemma 2 suggests that $L_\sigma \geq R_{P(p)}$, and the results follows.

Lemma 6. *Let p be inserted into the mesh because of R4. Then,*

- $R_p \geq \frac{1-\zeta}{2}R_{P(p)}$, if $P(p)$ is neither R4 nor R5,




Fig. 3: Flow diagram depicting the relationship among the rules. No solid cycle should have a product less than 1. The dashed arrows break the cycle.

- $R_p \geq bR_{P(p)}$, otherwise.

Proof. Let σ be the simplex that violates R4. Suppose that $P(p)$ is neither R4 nor R5. Since p belongs to the picking region of σ , $R_p \geq (1 - \zeta)R_\sigma \geq \frac{1-\zeta}{2}L_\sigma$. From Lemma 2, we have that $R_p \geq \frac{1-\zeta}{2}R_{P(p)}$.

Otherwise, consider the case $P(p)$ is an R4 point. From Lemma 1, we know that the circumradius of σ is more than b times the circumradius of the R4 simplex σ' that inserted $P(p)$. Therefore, $R_p \geq (1 - \zeta)R_\sigma \geq (1 - \zeta)bR_{\sigma'}$. However, the quantity $(1 - \zeta)R_{\sigma'}$ is equal to $R_{P(p)}$, and the statement holds.

The exact same logic holds when $P(p)$ is an R5 point, by just substituting $\mathcal{R}_{z,\sigma'}$ for $R_{\sigma'}$, where σ' is an R5 simplex.

Lemma 7. Let p be inserted into the mesh because of R5. Then,

- $\frac{1-\zeta}{2}R_{P(p)}$, if $P(p)$ is not an R5 point,
- $\Delta_{\partial\Omega}(z)$, otherwise.

Proof. Let σ_3 be the simplex that violates R4. Suppose that $P(p)$ is not an R5 point. Because of Lemma 2, the shortest edge of σ_3 is at least $R_{P(p)}$. Therefore, any surface ball of σ_3 has radius at least $\frac{1}{2}R_{P(p)}$. Since the surface ball does not contain any vertex in its interior, $R_p \geq \frac{1-\zeta}{2}R_{P(p)}$.

Suppose that $P(p)$ is an R5 point. Note that when $P(p)$ is inserted, all the free vertices closer than $\Delta_{\partial\Omega}(P(p))$ to $P(p)$ are deleted. Due to R5, σ_3

contains at least one free vertex. Since $P(p)$ is the most recently inserted vertex incident to the closest edge of σ_3 , the edge that contains $P(p)$ and the free vertex has to be at least $\Delta_{\partial\Omega}(P(p))$. Therefore, any surface ball of σ_3 has radius at least $\frac{1}{2}\Delta_{\partial\Omega}(P(p))$. Hence, $R_p \geq \frac{1-\zeta}{2}\Delta_{\partial\Omega}(P(p))$.

Putting all the Lemmas together, the solid arrows of Figure 3 show the insertion radius of the inserted point as a fraction of the insertion radius of its parent. An arrow from R_i to R_j with label x implies that the insertion radius of an R_j point p is at least x times larger than the insertion radius of its R_i parent $P(p)$. The label x of the dashed arrows is the absolute value of R_p . Note that the labels of the dashed arrows depend on the local feature size of $\partial\Omega$, and as such are always positive constants.

Recall that during refinement, free vertices might be deleted (because of R5). Nevertheless, such deletions of vertices are always preceded by insertion of feature points. Considering the fact that feature vertices are never deleted from the mesh, termination is guaranteed if we prove that the insertion radii of the inserted vertices cannot decrease indefinitely. Clearly, [30, 40, 41], if there is no solid cycle of product less than 1, termination is guaranteed.

Theorem 1. *The algorithm terminates producing simplices of bounded aspect ratio, if*

- $\frac{(1-\zeta)^2}{4}\bar{\rho} \geq 1$, and
- $\frac{1-\zeta}{2}b \geq 1$.

Proof. See Figure 3. The smallest product is produced by the solid cycles $R3 \rightarrow R4 \rightarrow R5 \rightarrow R3$ and $R4 \rightarrow R5 \rightarrow R4$. By requiring the label product of these loops to be more than 1, the desired result follows.

5 Accuracy

In this section, we prove that the mesh boundary is equal to the restriction of a $\partial\Omega$ sample Δ to $\partial\Omega$. In the literature, it is proved that these tetrahedra approximate the surface correctly, in geometric and topological sense [3, 8, 14].

First, we show that δ directly controls the density of the feature vertices. Let V be the set of vertices in the final mesh and Δ be equal to $V \cap \partial\Omega$.

Lemma 8. *Let $\delta < \frac{1}{4}$. Then Δ is a $\frac{5\delta}{1-4\delta}$ -sample of $\partial\Omega$.*

Proof. Recall that upon termination, there is no tetrahedron for which R1, R2, R3, R4, or R5 apply.

See Figure 4. Let p be an arbitrary point on $\partial\Omega$. Since $\mathcal{D}(V)$ covers all the domain, point p has to lie on or inside the circumsphere of a pentatope σ (not




Fig. 4: Proof of Lemma 8.

shown). Hence, \mathcal{B}_σ intersects $\partial\Omega$. Let point p' be the feature point closest to C_σ . Note that $|C_\sigma - p| \geq |C_\sigma - p'|$ and therefore p' lies on or inside σ 's circumsphere. We also know that σ 's circumradius has to be less than $2\Delta_{\partial\Omega}(p')$, since otherwise R2 would apply for t . Therefore, we have the following:

$$\begin{aligned} |p - p'| &< 2R_\sigma && \text{(because both } p \text{ and } p' \text{ lie on or inside } \mathcal{B}_\sigma) \\ &< 4\Delta_{\partial\Omega}(p') && \text{(because of R2)} \\ &\leq 4\delta(|p - p'| + \text{lfs}_{\partial\Omega}(p)) && \text{(from Inequality (1)),} \end{aligned}$$

and by reordering the terms, we obtain that:

$$|p - p'| < \frac{4\delta}{1-4\delta} \text{lfs}_{\partial\Omega}(p), \text{ with } \delta < \frac{1}{4}. \quad (2)$$

Moreover, there must exist a feature vertex v in the triangulation closer than $\Delta_{\partial\Omega}(p') = \delta \cdot \text{lfs}_{\partial\Omega}(p')$ to p' , since otherwise R1 would apply for σ . Hence, $|v - p'| < \delta \cdot \text{lfs}_{\partial\Omega}(p')$, and using Inequality (1), we have that:

$$|v - p'| < \delta(|p - p'| + \text{lfs}_{\partial\Omega}(p)) \quad (3)$$

Applying the triangle inequality for $\triangle pvp'$ yields the following:

$$\begin{aligned} |p - v| &\leq |pp'| + |v - p'| \\ &< |p - p'| + \delta(|p - p'| + \text{lfs}_{\partial\Omega}(p)) && \text{(from Inequality (3))} \\ &= |p - p'| (1 + \delta) + \delta \cdot \text{lfs}_{\partial\Omega}(p) \\ &< \frac{4\delta}{1-4\delta} \text{lfs}_{\partial\Omega}(p) (1 + \delta) + \delta \cdot \text{lfs}_{\partial\Omega}(p) && \text{(from Inequality (2))} \\ &= \left(\frac{4\delta(1+\delta)}{1-4\delta} + \delta \right) \text{lfs}_{\partial\Omega}(p) \\ &= \frac{5\delta}{1-4\delta} \text{lfs}_{\partial\Omega}(p), \end{aligned}$$

and the proof is complete.

Let us denote with ω_i one of the n connected components that Ω consists of: $\Omega = \bigcup_{i=1}^n \omega_i$. The next two Lemmas prove a few useful properties for the mesh \mathcal{M}_Ω and its boundary $\partial\mathcal{M}_\Omega$. Our goal is to show that $\partial\mathcal{M}_\Omega$ is always non-empty and does not have boundary (Lemma 10), a fact that will be used for proving the fidelity guarantees (Theorem 2).

Lemma 9. *Let $\delta \leq \frac{1}{4}$. Then, for every ω_i there is a pentatope $\sigma \in \mathcal{D}(V)$, such that C_σ lies inside ω_i .*

Proof. Let us consider a single connected component ω_i . The same reasoning applies for any connected component of Ω .

For the sake of contradiction, assume that there is no pentatope whose circumcenter lies inside ω_i . Since the triangulation $\mathcal{D}(V)$ covers all the domain, the circumballs of the pentatopes in $\mathcal{D}(V)$ also cover the domain ω_i . Therefore, there has to be a circumball \mathcal{B}_σ ($\sigma \in \mathcal{D}(V)$) which intersects a point m on the medial axis of $\partial\omega_i$, such that m lies inside ω_i . By our assumption, the circumcenter C_σ cannot lie inside ω_i . Therefore, \mathcal{B}_σ intersects $\partial\omega_i$. Also, recall that R2 cannot apply to any pentatope. Hence, we have the following:

$$\begin{aligned} 2 \cdot \delta \cdot \text{lfs}_{\partial\Omega}(C_\sigma) &> R_\sigma && \text{(from R2)} \\ &\geq \frac{|\text{cfp}_{\partial\Omega}(C_\sigma) - m|}{2} && \text{(since } m \text{ and } \text{cfp}_{\partial\Omega}(C_\sigma) \text{ do not lie outside } \mathcal{B}_\sigma) \\ &\geq \frac{\text{lfs}_{\partial\Omega}(\text{cfp}_{\partial\Omega}(C_\sigma))}{2} && \text{(since } m \text{ is on the medial axis)} \\ \delta &> \frac{1}{4}, && \Rightarrow \end{aligned}$$

which raises a contradiction.

Lemma 10. *Let $\delta \leq \frac{1}{4}$. Then, $\partial\mathcal{M}_\Omega$ is a non-empty set and does not have boundary.*

Proof. The fact that $\partial\mathcal{M}_\Omega$ is a non-empty set follows directly from Lemma 9: since \mathcal{M}_Ω cannot be empty, its boundary $\partial\mathcal{M}_\Omega$ cannot be empty too. For the other part, since $\partial\mathcal{M}_\Omega$ is the boundary of a set of tetrahedra, it cannot have a boundary.

The following Theorem proves the fidelity guarantees:

Theorem 2. *The mesh boundary $\partial\mathcal{M}_\Omega$ is the restriction to $\partial\Omega$ of $\Delta = V \cap \partial\Omega$.*

Proof. Let f be a tetrahedron σ_3 in $\partial\mathcal{M}_\Omega$. As such, $\mathcal{V}(\sigma_3)$ intersects $\partial\Omega$. Due to R5, the vertices of σ_3 lie on $\partial\Omega$. Recall that the surface ball \mathcal{B}_{z,σ_3} does not contain vertices in its interior. Therefore, \mathcal{B}_{z,σ_3} is empty of vertices in $V \cap \partial\Omega$

also. Without loss of generality, assume that the vertices in V are in general position. Since there is a ball that circumscribes σ_3 and does not contain vertices of $V \cap \partial\Omega$ in its interior, σ_3 has to appear as a simplex in $\mathcal{D}(V \cap \partial\Omega)$. Since the center z of the surface ball lies on $\partial\Omega$, then the voronoi dual of σ_3 intersects $\partial\Omega$ in $\mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$, as well. Hence, $\partial\mathcal{M}_\Omega \subseteq \mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$.

For the other direction, we will prove that $\partial\mathcal{M}_\Omega$ cannot be a proper subset of $\mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$, and therefore, equality between these 2 sets is forced. Toward this direction, we will prove that any proper non-empty subset of $\mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$ has boundary; this is enough, because we have proved in Lemma 10 that $\partial\mathcal{M}_\Omega$ is non-empty and does not have boundary.

$\mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$ is the restriction of a sample of a closed manifold $\partial\Omega$ and therefore it is a 3-manifold without boundary [3]. That means that any 2-simplex in $\mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$ is incident to exactly two 3-simplices of $\mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$. Since any proper non-empty subset \mathcal{A} of $\mathcal{D}_{|\partial\Omega}(\partial\Omega \cap V)$ has fewer 3-simplices, \mathcal{A} contains at least a 2-simplex σ_2 incident to only one 3-simplex. But this implies that σ_2 belongs to the boundary of \mathcal{A} , and the proof is complete.

6 Experimental Evaluation

The algorithm is implemented in C++. We employed the Bowyer-Watson kernel [11, 47] for point insertions. The removal of a point p is implemented by computing the *small* Delaunay triangulation of the vertices incident to p [21], such that the vertices inserted earlier in the triangulation are inserted into the small triangulation first. It can be shown [25] that these new created pentatopes can always be connected back to the original triangulation without introducing invalid elements. For the Euclidean Distance Transform, we made use of the related filter implemented in ITK [2] and described in [32]. Lastly, we borrowed CGAL's [1] exact predicates for the accurate computation of the 4D in-sphere tests.

Table 1: Information about the images of the five patients used in this section. The spacing for all the images is $(1.77, 1.77, 6, 1)mm^4$.

| Case | Pat1 | Pat2 | Pat3 | Pat4 | Pat5 |
|---------|--|--|--|--|--|
| #Voxels | $(100 \times 100 \times 44 \times 15)$ | $(100 \times 100 \times 34 \times 15)$ | $(100 \times 100 \times 26 \times 15)$ | $(100 \times 100 \times 31 \times 15)$ | $(100 \times 100 \times 29 \times 15)$ |

We ran our code on five (segmented) images obtained from the 4D Heart Database [35]. The first three represent the moving left ventricle of the patients, while the last two the ventricle together with the myocardium for 15 cardiac cycles. Information about these data is given in Table 1.

Recall that our algorithm needs the distance of any point on $\partial\Omega$ from the medial axis. The robust computation of the medial axis is a very difficult

Table 2: Statistics of the output meshes generated for each patient.

| | Pat1 | Pat2 | Pat3 | Pat4 | Pat5 |
|--|--------------------|--------------------|--------------------|--------------------|--------------------|
| #Pentatopes | 49,479 | 43,673 | 8,883 | 63,016 | 56,528 |
| #Boundary Tetrahedra | 30,758 | 29,089 | 8,271 | 36,281 | 33,308 |
| #Vertices | 4,709 | 4,314 | 1,362 | 5,567 | 5,132 |
| Shortest edge (mm) | 3.45 | 3.87 | 3.90 | 3.5 | 4.63 |
| Radius-edge ratio (maximum, average, deviation) | (1.93, 1.02, 0.17) | (1.78, 0.98, 0.15) | (1.54, 0.92, 0.10) | (2.20, 1.06, 0.18) | (1.87, 1.05, 0.18) |
| Normalized volume (minimum, average, deviation) | (0.01, 0.34, 0.18) | (0.01, 0.38, 0.18) | (0.02, 0.43, 0.17) | (0.01, 0.32, 0.17) | (0.01, 0.33, 0.17) |

problem (see [22, 27] for computing the exact medial axis, [19] for a review of image-based medial axis methods, and [5] for computing the medial axis given a set of surface points) and out of the scope of this paper. Therefore, in the implementation, we assume that $\text{lfs}_{\partial\Omega}(p)$ is uniform and equal to the unit, which implies that $\Delta_{\partial\Omega}(p)$ becomes equal to δ . That is, in practice, δ determines a uniform and (if small enough) dense sample of the surface. We experimentally verified that a δ value equal to 5 (the length of five consecutive voxels along the fourth dimension) yielded manifold mesh boundaries with vertices lying precisely on the iso-surface in accordance with Theorem 2.

The quantity τ_σ determines the aspect ratio of pentatope σ [30], but it is not normalized, and therefore, it is hard to draw comparative conclusions. For this reason, for a pentatope σ of the final mesh, we report its *normalized volume* $\hat{\tau}_\sigma$ defined as the ratio of its volume over the volume of a regular pentatope with circumradius equal to the circumradius of σ (or alternatively $\hat{\tau}_\sigma = \frac{384\text{Vol}_\sigma}{24R_\sigma^4\sqrt{5}}$). Clearly, $\hat{\tau}_\sigma \in [0, 1]$, where a value of 0 implies a degenerate pentatope, while 1 implies a perfect quality.




Fig. 5: Normalized volume histogram of the output mesh obtained for the input image Pat1.

Table 2 shows quantitative data for the mesh generated on each image. We set the radius of the picking regions equal to $\zeta = \frac{1}{2}$. Theorem 1 suggests that $\bar{\rho}$ be at least 16 and b at least 4. We experimentally observed that by

selecting 4 to 10 random points within the picking regions (both the 4- and the 3-topological balls), no small element σ was created with $\hat{\tau}_\sigma$ less than 0.01. Despite the fact a value of 0.01 is rather small, it is three orders of magnitude larger than the minimum normalized volume reported in the case where no picking regions are employed at all. Also, notice that the average normalized volume is much higher than the minimum value. This fact together with the observed small standard deviation implies that most pentatopes have normalized volume away from the minimum value and very close to the average. Figure 5 shows the histogram of the normalized volumes for the first experiment of Table 2, that is, when the input image Pat1 was used. Similar histograms are observed for all the other inputs as well.

7 Conclusions

In this paper, we present a space-time meshing method for (3D+t) image data. The method is able to provably clean up slivers and recover the hyper-surfaces faithfully. Experiments on five 4D cardiac images show that the resulting meshes consist of elements of bounded aspect ratio.

Efficient Discontinuous Galerkin formulations require that not only the hyper-surface should be recovered but also the evolving 3D object at certain time steps [18]. This is a more challenging task considering the non-manifold nature of the underlying space-time domain and it is left as future work.

Because of the increased memory space needed for high dimensional meshing, our 4D algorithm is rather slow: it is approximately 100 times slower than our three dimensional Delaunay mesher published in the past [26]. In the future, we plan to optimize our sequential 4D code presented in this paper and parallelize it, exploiting in this way the available concurrency across the temporal dimension as well.

Acknowledgements

The authors would like to thank Dr. Marek Behr, RWTH Aachen University, for the constructive discussions. This work is supported in part by NSF grants: CCF-1139864, CCF-1136538, and CSI-1136536 and by the John Simon Guggenheim Foundation and the Richard T. Cheng Endowment.

References

1. CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>, v4.0.
2. ITK, Insight Segmentation and Registration Toolkit. <http://www.itk.org>, v4.1.0.

3. Nina Amenta and Marshall Bern. Surface reconstruction by Voronoi filtering. In *SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry*, pages 39–48, New York, NY, USA, 1998. ACM.
4. Nina Amenta, Sunghee Choi, Tamal K. Dey, and N. Leekha. A Simple Algorithm for Homeomorphic Surface Reconstruction. *International Journal of Computational Geometry and Applications*, 12(1-2):125–141, 2002.
5. Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*, SMA '01, pages 249–266, New York, NY, USA, 2001. ACM.
6. Dominique Attali, Herbert Edelsbrunner, and Yuriy Mileyko. Weak witnesses for delaunay triangulations of submanifolds. In *Proceedings of the 2007 ACM symposium on Solid and physical modeling*, SPM '07, pages 143–150, New York, NY, USA, 2007. ACM.
7. Marek Behr. Simplex space-time meshes in finite element simulations. *International Journal for Numerical Methods in Fluids*, 57:1421–1434, 2008.
8. Jean-Daniel Boissonnat, Leonidas J. Guibas, and Steve Y. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. *Discrete Comput. Geom.*, 42:37–70, May 2009.
9. Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, 2005.
10. Dobrina Boltcheva, Mariette Yvinec, and Jean-Daniel Boissonnat. Mesh Generation from 3D Multi-material Images. In *Medical Image Computing and Computer-Assisted Intervention*, pages 283–290. Springer, September 2009.
11. Adrian Bowyer. Computing Dirichlet tessellations. *Computer Journal*, 24:162–166, 1981.
12. Frédéric Cazals and Joachim Giesen. Delaunay triangulation based surface reconstruction: Ideas and algorithms. In *Effective Computational Geometry for Curves and surfaces*, pages 231–273. Springer, 2006.
13. Siu-Wing Cheng, Tamal K. Dey, Herbert Edelsbrunner, Michael A. Facello, and Shang-Hua Teng. Sliver exudation. *Journal of the ACM*, 47(5):883–904, 2000.
14. Siu-Wing Cheng, Tamal K. Dey, and Edgar A. Ramos. Manifold reconstruction from point samples. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '05, pages 1018–1027, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
15. 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.
16. Andrey Chernikov and Nikos Chrisochoides. Multitissue tetrahedral image-to-mesh conversion with guaranteed quality and fidelity. *SIAM Journal on Scientific Computing*, 33:3491–3508, 2011.
17. Andrey N. Chernikov and Nikos P. Chrisochoides. Generalized insertion region guides for Delaunay mesh refinement. *SIAM Journal on Scientific Computing (SISC)*, 2011. under revision.
18. Bernardo Cockburn, George E. Karniadakis, and Chi-Wang Shu. Discontinuous galerkin methods: theory, computation and applications. *Lecture notes in Computational Science and Engineering*, 11, 2000.
19. David Coeurjolly and Annick Montanvert. Optimal separable algorithms to compute the reverse euclidean distance transformation and discrete medial axis in arbitrary dimension. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29:437–448, March 2007.

20. Per Eric Danielsson. Euclidean Distance Mapping. *Computer Graphics and Image Processing*, 14:227–248, 1980.
21. Olivier Devillers and Monique Teillaud. Perturbations and vertex removal in a 3D Delaunay triangulation. In *Proceedings of the 14th ACM-SIAM Symposium on Discrete algorithms*, SODA '03, pages 313–319, 2003.
22. Tamal K. Dey and Wulue Zhao. Approximate medial axis as a voronoi subcomplex. *Computer-Aided Design*, 36(2):195–202, 2004.
23. Herbert Edelsbrunner and Nimish R. Shah. Triangulating topological spaces. In *SCG '94: Proceedings of the tenth annual symposium on Computational geometry*, pages 285–292, New York, NY, USA, 1994. ACM.
24. Jeff Erickson, Damrong Guoy, John M. Sullivan, and Alper Üngör. Building spacetime meshes over arbitrary spatial domains. *Eng. with Comput.*, 20(4):342–353, August 2005.
25. Panagiotis Foteinos and Nikos Chrisochoides. Dynamic parallel 3D Delaunay triangulation. In *International Meshing Roundtable*, pages 9–26, Paris, France, October 2011.
26. Panagiotis Foteinos and Nikos Chrisochoides. High quality real-time image-to-mesh conversion for finite element simulations. In *27th ACM International Conference on Supercomputing (ICS'13)*, June 2013.
27. Peter Giblin and Benjamin B. Kimia. A formal classification of 3D medial axis points and their local geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:238–251, January 2004.
28. Xiangmin Jiao, Andrew Colombi, Xinlai Ni, and John Hart. Anisotropic mesh adaptation for evolving triangulated surfaces. *Eng. with Comput.*, 26(4):363–376, 2010.
29. François Labelle and Jonathan Richard Shewchuk. Isosurface stuffing: fast tetrahedral meshes with good dihedral angles. *ACM Transactions on Graphics*, 26(3):57, 2007.
30. Xiang-Yang Li. Generating Well-Shaped D-dimensional Delaunay Meshes. In Jie Wang, editor, *Computing and Combinatorics*, volume 2108 of *Lecture Notes in Computer Science*, pages 91–100. Springer Berlin / Heidelberg, 2001.
31. William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *SIGGRAPH Computer Graphics*, 21(4):163–169, 1987.
32. Calvin .R. Maurer, Qi Rensheng, and Vijay Raghavan. A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):265 – 270, feb 2003.
33. Gary L. Miller, Dafna Talmor, Shang-Hua Teng, and Noel Walkington. A Delaunay based numerical method for three dimensions: generation, formulation, and partition. In *Proceedings of the 27th Annu. ACM Sympos. Theory Comput.*, pages 683–692. ACM, 1995.
34. Scott A. Mitchell and Stephen A. Vavasis. Quality mesh generation in higher dimensions. *SIAM J. Comput.*, 29(4):1334–1370, February 2000.
35. L. Najman, J. Cousty, M. Couprie, H. Talbot, S. Clment-Guinaudeau, T. Goissen, and J. Garot. An open, clinically-validated database of 3D+t cine-mr images of the left ventricle with associated manual and automated segmentation. <http://www.laurentnajman.org/heart/index.html>.

36. Martin Neumüller and Olaf Steinbach. Refinement of flexible spacetime finite element meshes and discontinuous Galerkin methods. *Computing and Visualization in Science*, 14:189–205, 2011.
37. Steve Oudot, Laurent Rineau, and Mariette Yvinec. Meshing volumes bounded by smooth surfaces. In *Proceedings of the International Meshing Roundtable*, pages 203–219. Springer-Verlag, September 2005.
38. Jean-Philippe Pons, Florent Ségonne, Jean-Daniel Boissonnat, Laurent Rineau, Mariette Yvinec, and Renaud Keriven. High-Quality Consistent Meshing of Multi-label Datasets. In *Information Processing in Medical Imaging*, pages 198–210, 2007.
39. Laurent Rineau and Mariette Yvinec. Meshing 3D domains bounded by piecewise smooth surfaces. In *Proceedings of the International Meshing Roundtable*, pages 443–460, 2007.
40. Jonathan Richard Shewchuk. Tetrahedral mesh generation by Delaunay refinement. In *Proceedings of the 14th ACM Symposium on Computational Geometry*, pages 86–95, Minneapolis, MN, 1998.
41. Jonathan Richard Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry: Theory and Applications*, 22(1-3):21–74, May 2002.
42. Hang Si. Constrained Delaunay tetrahedral mesh generation and refinement. *Finite Elements in Analysis and Design*, 46:33–46, 2010.
43. Hang Si. TetGen, A Quality Tetrahedral Mesh Generator and a 3D Delaunay Triangulator. <http://tetgen.berlios.de/>, v1.4.3.
44. Shripad Thite. Efficient spacetime meshing with nonlocal cone constraints. In *13th International Meshing Roundtable*, pages 47–58, 2004.
45. Yanghai Tsin, Klaus Kirchberg, Guenter Lauritsch, and Chenyang Xu. A deformation tracking approach to 4d coronary artery tree reconstruction. In Guang-Zhong Yang, David Hawkes, Daniel Rueckert, Alison Noble, and Chris Taylor, editors, *Medical Image Computing and Computer-Assisted Intervention MICCAI 2009*, volume 5762 of *Lecture Notes in Computer Science*, pages 68–75. Springer Berlin / Heidelberg, 2009.
46. M von Siebenthal, G Székely, U Gamper, P Boesiger, A Lomax, and Ph Cattin. 4D MR imaging of respiratory organ motion and its variability. *Physics in Medicine and Biology*, 52(6):1547–1564, 2007.
47. David F. Watson. Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes. *Computer Journal*, 24:167–172, 1981.
48. Ernst Weigang, Fabian A. Kari, Friedhelm Beyersdorf, Maximilian Luehr, Christian D. Etz, Alex Frydrychowicz, Andreas Harloff, and Michael Markl. Flow-sensitive four-dimensional magnetic resonance imaging: flow patterns in ascending aortic aneurysms. *European Journal of Cardio-Thoracic Surgery*, 34(1):11–16, 2008.