Gabriel meshes and Delaunay edge flips

Ramsay Dyer
GrUVi Lab
School of Computing Science
Simon Fraser University, Canada

Hao Zhang
GrUVi Lab
School of Computing Science
Simon Fraser University, Canada

Torsten Möller
GrUVi Lab
School of Computing Science
Simon Fraser University, Canada

Abstract

We undertake a study of the local properties of 2-Gabriel meshes: manifold triangle meshes each of whose faces has an open Euclidean diametric ball that contains no mesh vertices. We show that, under mild constraints on the dihedral angles, such meshes are Delaunay meshes: the open geodesic circumdisk of each face contains no mesh vertex.

The analysis is done by means of the Delaunay edge flipping algorithm and it reveals the details of the distinction between these two mesh structures. In particular we observe that the obstructions which prohibit the existence of Gabriel meshes as homeomorphic representatives of smooth surfaces do not hinder the construction of Delaunay meshes.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations

Keywords: manifold triangle mesh, Gabriel complex, Delaunay mesh

1 Introduction

We examine and compare two triangle mesh structures for representing smooth surfaces: Gabriel meshes and Delaunay meshes. The definition of each of these structures adapts the Delaunay paradigm to apply to meshes that represent two dimensional non-Euclidean geometry. An appealing aspect of these structures is that they do not depend on a separate reference surface for their definition. This is in contrast to the *restricted Delaunay triangulation* [Edelsbrunner and Shah 1994], a structure which has seen much more attention.

Delaunay meshes have been studied because they present a natural compatibility with recently defined discrete differential structures and operators [Bobenko and Springborn 2007; Desbrun et al. 2005]. Wardetzky et al. [2007] identified four properties that would be desirable to have in a discrete Laplacian operator. They then went on to demonstrate that no discrete Laplacian operator can maintain all four properties on arbitrary meshes. However, on Delaunay meshes the limitations imposed by that theorem do not apply; Laplacian operators based on the cotan formula [Pinkall and Polthier 1993] enjoy all the properties identified by Wardetzky et al.

Gabriel meshes on the other hand have seen some attention in the context of surface reconstruction. Their simple definition ensures that they will be substructures of the three dimensional Delaunay tetrahedralization, which in turn implies that smooth

Gabriel meshes would serve as good representations of well sampled smooth surfaces [Amenta et al. 2000; Petitjean and Boyer 2001]. We loosley define a *smooth mesh* as one in which the angle between the normal vectors of two faces that share a vertex is less than some appropriately small constant.

1.1 Motivation

Our work is inspired by the work of Cheng and Dey [2007], where they describe an edge flipping algorithm that seeks to produce a Gabriel mesh from a given mesh whose vertices are well sampled from a smooth surface. We provide an analysis of this algorithm that reveals a close relationship with the Delaunay edge flipping algorithm we presented in [Dyer et al. 2007]. The local criteria that are required of a Gabriel mesh are slightly stronger than what is required of a Delaunay mesh, and this is important, because a closed Gabriel mesh does not exist in general [Chaine 2003]. In fact, we show here that even when local uniformity constraints are imposed on the sample set, obstructions to the existence of closed Gabriel meshes may remain.

The authors of [Cheng and Dey 2007], which is a technical report, have retracted their result, but the motivation behind that work persists. The appeal of the Gabriel mesh is that it is a *locally defined surface representation*: The connectivity of the vertices is constrained simply in terms of essentially local information and without the need for an independent reference surface. In particular it is not necessary to construct the full Delaunay triangulation of the ambient three dimensional space in order to verify the local connectivity. This characteristic is also shared by Delaunay meshes. We present smooth Delaunay meshes as holding promise in filling the role that has previously been sought in Gabriel meshes. However, the current work does not represent the completion of that program.

1.2 Contribution

Although it was not explicitly done, the tools needed to demonstrate that closed Gabriel meshes will not exist in general, even when local uniformity constraints are imposed on the sample set were already provided in [Chaine 2003]. In Section 7 we explicitly construct an obstruction which demonstrates this fact; but in this respect we wish to draw attention to this implication of Chaine's work, rather than claim originality.

Our principal contribution is the illumination of the close relationship between Delaunay meshes and Gabriel meshes. We show that Gabriel meshes are Delaunay meshes. In Section 6 we demonstrate that the locally Delaunay property of an edge is a slight relaxation of what is locally demanded by a Gabriel mesh: the precise distinction is identified. Delaunay meshes are not hampered by the obstructions that prevent the existence of closed Gabriel meshes, thus they hold promise as an alternative to Gabriel meshes as a locally defined surface representation. However, for Delaunay meshes themselves there is still no existence proof on a fixed sample set.

We also introduce, in Section 4, a definition of the Gabriel complexes that reveals a natural hierarchical decomposition of the Delaunay triangulation, which we express in Theorem 1. To the best

of our knowledge, Theorem 1 is a new result. This description of the Gabriel complexes suggests they may hold promise as a scaffold for a manifold reconstruction algorithm in a high dimensional ambient space, where the full Delaunay triangulation is impractical.

2 Background and related work

In this work, unless stated otherwise, all disks and balls are open: they do not contain their boundary. A *geodesic circumdisk* of a face t on a triangle mesh is a geodesic disk that has the vertices of t on its boundary and whose closure contains t. If t has a geodesic circumdisk empty of mesh vertices, then it is unique and isometric to a planar disk.

2.1 Definitions and planar equivalence

The term *intrinsic* is applied to properties and objects that can be defined in terms of the surface itself, without any reference to its embedding in the ambient space. A *Delaunay mesh*, as defined in [Dyer et al. 2007], is a manifold triangle mesh whose triangle faces are Delaunay with respect to the intrinsic metric of the mesh. This means that the open geodesic circumdisk of each triangle is empty of mesh vertices. *Gabriel meshes* are defined by a more explicit relationship with the ambient space. For a triangle $t \subset \mathbb{R}^3$ the smallest open Euclidean ball that has the vertices of t on its boundary is the *diametric ball* of t. A *Gabriel mesh* is a manifold triangle mesh each of whose faces has a diametric ball empty of mesh vertices. We say the faces have the *Gabriel property* with respect to the vertex set.

The Gabriel property is defined as a global condition, but it is misleading to emphasize this if we are dealing with points well sampled from a smooth surface. In this case the circumradii involved are a small fraction of the distance to the medial axis: geodesically distant points are irrelevant. What is at issue is the local connectivity of the samples.

Consider set of sample points lying in a plane in $\mathbb{R}^3$. A Gabriel mesh on these samples is equivalent to a Delaunay triangulation of the points because the restriction of the diametric ball of triangle t to the plane is just the circumdisk of t. If the circumdisks are empty then so will be the diametric balls, and vise versa. Likewise, the circumdisks of the triangles are geodesic disks on the planar mesh, so for planar meshes, there is no distinction between a Gabriel mesh and a Delaunay mesh. The distinction between Delaunay meshes and Gabriel meshes arises when we consider general manifold surfaces.

2.2 Related work on Gabriel meshes

Petitjean and Boyer [2001] defined the Gabriel complex of a given set of points P as consisting of those triangles in $\mathbb{R}^3$ that have the Gabriel property with respect to P. A reconstruction algorithm was introduced which extracted a manifold triangle mesh from the Gabriel complex. A heuristic argument, based on the planar case mentioned above, was given for why there should be enough triangles in the Gabriel complex to extract a closed manifold mesh if P is well sampled from a smooth surface. The issue of near degenerate configurations was not mentioned, although it was recognized to be a problem with an algorithm described earlier in the paper.

An *umbrella* at a sample point p is a collection of edge-adjacent triangles that all share p as a vertex and that are together homeomorphic to a closed disk. The umbrella is *full* if p is mapped to the interior of the disk. In their surface reconstruction algorithm Adamy et al. [2000] also implicitly assumed that it was possible to extract full umbrellas of Gabriel triangles around each vertex. A

topological clean-up step filled holes in the extracted mesh, some of which were produced by other topological reparation steps. The Gabriel complex was defined as the simplices of codimension one that had the Gabriel property and ambient dimensions two and three were considered.

The surface reconstruction algorithm presented by Attene and Spagnuolo [2000] also exploited the Gabriel property. Interestingly it was not assumed that the Gabriel faces would form a closed surface, but that holes would be limited to isolated missing triangles.

Then Chaine [2003] observed that a certain sliver tetrahedron in the 3D Delaunay tetrahedralization was sufficient to prevent full umbrellas of Gabriel faces at its vertices. We refer to such a tetrahedron as a *tetrahedral obstruction* to Gabriel meshes. The example described has two vertices very close together and one might wonder whether uniformity constraints could be imposed on the sample set so as to permit the construction of closed Gabriel meshes.
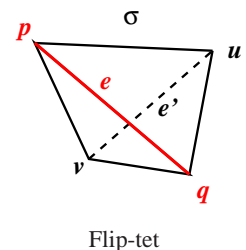
3 Delaunay edge flips

In [Bobenko and Springborn 2007] the intrinsic Delaunay triangulation of a piecewise flat surface was carefully defined and shown to have similar properties to the planar Delaunay triangulation. A Delaunay mesh is a manifold triangle mesh whose edges form an intrinsic Delaunay triangulation of its vertices [Dyer et al. 2007].

Although the Delaunay mesh is defined in terms of its own intrinsic metric and is independent of a reference surface, we emphasize that it is not intrinsically defined. Rather, the definition demands a marriage of intrinsic and extrinsic triangulations. The edges of a mesh can be interpreted as an artifact of the embedding of the piecewise flat surface into $\mathbb{R}^3$. A Delaunay mesh demands that this extrinsic triangulation coincide with the intrinsic Delaunay triangulation.

We review the edge flipping algorithm of [Dyer et al. 2007], which produces a Delaunay mesh from a given mesh M. Note that this is an *extrinsic edge flipping algorithm*: the actual edges of the mesh triangles are flipped, resulting in a modification of the geometry of the mesh. This is in contrast to the intrinsic edge flipping algorithms studied elsewhere [Bobenko and Springborn 2007; Fisher et al. 2006; Glickenstein 2005], where the geodesic triangulation changes, but the geometry remains fixed.

The *hinge* defined by an edge $e = [p, q] \subset M$ is e together with the two triangle faces adjacent to it. Edge e is the *pivot* of the hinge. The four vertices of a hinge define a tetrahedron that we call the *flip-tet*. Let a hinge on e be defined by triangles $t_1 = [u, p, q]$ and $t_2 = [v, q, p]$, and let σ be the flip-tet defined by this hinge. The edge $e' = [u, v]$ is the *opposing edge* to e and it defines, together with its adjacent faces in σ, an *opposing hinge* to that of e.

Flip-tet

Note that an empty geodesic disk on a mesh can be unfolded onto the plane without metric distortion. A *locally Delaunay edge* is defined as in the planar case: e is locally Delaunay if the geodesic circumdisk of one of its adjacent faces does not contain the other adjacent face, otherwise it is *not locally Delaunay* (nlD). An edge e is locally Delaunay if it is locally Delaunay by the usual planar criteria when we unfold its hinge onto a plane. For our purposes a convenient characterization of a locally Delaunay edge is that the sum of the two face angles it subtends must not exceed π. In particular, if e subtends two acute angles, then it is locally Delaunay

and if it subtends two obtuse angles, it is nlD. The opposing edge to an nlD edge is always locally Delaunay, and when an nlD edge with a non-planar hinge is flipped the surface area of the mesh is reduced [Dyer et al. 2007]. It follows that the edge flipping algorithm is sure to terminate. Analogously to the planar case, if all the edges in a mesh are locally Delaunay, then they are all globally Delaunay [Bobenko and Springborn 2007].

4 Gabriel complexes

We begin by defining the Gabriel complexes in arbitrary dimensions. Let $P \subset \mathbb{R}^d$ be a finite set of sample points in general position. Let $k \leq d$ and let σ_k be a k-dimensional simplex with vertices in P. The *diametric d-ball* for σ_k, denoted B_{σ_k}, is the smallest d-dimensional ball containing the vertices of σ_k on its boundary. We say σ_k is k-*Gabriel* if its diametric d-ball does not contain any points of P. The k-Gabriel complex of P is the simplicial complex formed by the k-Gabriel simplices and their faces.

Thus the 1-Gabriel complex is the Gabriel graph of P, and the d-Gabriel complex is the Delaunay triangulation of P. For completeness, define the 0-Gabriel complex to be the sample points themselves. Denote by $\mathcal{G}^k$ the k-Gabriel complex. Then we have

Theorem 1 *The Gabriel complexes form a nested hierarchy of subcomplexes of the Delaunay triangulation of $P \subset \mathbb{R}^d$:*

$$P = \mathcal{G}^0 \subset \mathcal{G}^1 \subset \cdots \subset \mathcal{G}^k \subset \cdots \subset \mathcal{G}^d.$$

That $\mathcal{G}^0 \subset \mathcal{G}^1$ can be established by demonstrating that the Gabriel graph contains the nearest neighbour graph. Our proof of Theorem 1 is an extension of this method. The edge to the nearest neighbour of a point p can be characterized as the 1-simplex on p that has the smallest diametric d-ball. This ball is necessarily empty of sample points. The extension to higher dimensions follows by showing that if σ_k is k-Gabriel, then the $(k + 1)$-simplex with the smallest diametric d-ball amongst those that have σ_k as a face, must be $(k + 1)$-Gabriel. The technical demonstration of this result is provided in a supplemental docmuent, and published in a technical report [Dyer et al. 2008].

According to Theorem 1, the k-Gabriel complex is a subcomplex of the k'-Gabriel complex if $k' \geq k$. However not all the k-simplices of the k'-Gabriel complex need belong to the k-Gabriel complex. It is this latter fact that motivates the introduction of the k prefix.

Indeed, we will be interested in the edges of a substructure of the 2-Gabriel complex, but these edges need not belong to the Gabriel graph. Triangle faces containing an edge e of the 1-Gabriel complex cannot have an obtuse angle subtended by e, but the 2-Gabriel complex may well contain obtuse triangles. This property of 1-Gabriel edges implies that they will never be flipped by a Delaunay edge flip, and so we anticipate their relevance in this context. However, for the purposes of the current work an unspecified mention of the Gabriel property can be understood as a reference to the 2-Gabriel complex. We define a *Gabriel mesh* as a manifold triangle mesh that is a substructure of the 2-Gabriel complex.

In all that follows we will take $d = 3$ for convenience. Note, however, that the ambient dimension has no bearing on the results. Local computations involving an isolated flip-tet may be confined to the affine hull of the tetrahedron. Since Gabriel triangles have empty circumballs, they belong to the Delaunay tetrahedralization by definition. We have the following useful characterization:

Lemma 1 *A triangle $t \subset \mathbb{R}^3$ is Gabriel iff its dual Voronoi edge intersects the affine hull of t.*

Proof If t is Gabriel, the diametric d-ball of t is empty and so its centre must lie on the Voronoi edge dual to t. But the centre of the diametric ball also lies on the plane defined by t. Conversely, if the Voronoi edge dual to t intersects the affine hull of t at c_t, then the empty circumball of t centred at c_t is a diametric ball, and so t is Gabriel. □

4.1 Gabriel faces in a tetrahedron

In the next section we will prepare for a detailed examination of hinges. Every hinge has an associated flip-tet, and we make some preliminary observations here about the Gabriel properties of tetrahedra. We will assume that the tetrahedron is non-degenerate. The case where the affine hull of a flip-tet is a plane corresponds to the usual case of planar Delaunay edge flips. More extreme degeneracies can be dealt with by the same arguments that permit the assumption of general position for planar point sets.

Let σ be a tetrahedron with circumsphere S_σ and circumcentre c_σ. The *interior half-space* with respect to σ of a triangle face $t \subset \sigma$ is the half space bounded by the supporting plane of t and containing the fourth vertex of σ. Likewise, the *exterior half-space* of t is the one that does not contain the fourth vertex. We define these half spaces to be closed: the intersection of the interior and the exterior half spaces of t is the affine hull of t.

The following two lemmas concern the Gabriel properties of the faces of an isolated tetrahedron. In other words we consider the sample set to be comprised solely of the four vertices of the tetrahedron.

Lemma 2 *A face t in a tetrahedron, σ, is Gabriel iff c_σ lies in its interior half-space.*

Proof Consider the Voronoi diagram of the vertices of σ. A Voronoi edge is supported by the line perpendicular to its dual face and through its circumcentre. The Voronoi edges all extend to infinity in one direction and terminate at c_σ in the other. A Voronoi edge dual to t must extend to infinity in the exterior half-space of t: given a sufficiently large radius a ball centred in this half-space and with the vertices of t on its boundary will not contain the fourth vertex of σ. Thus the Voronoi edge dual to t terminates at c_σ and always extends to the exterior half space of t. Therefore it will intersect the plane supporting t iff c_σ lies in interior half-space of t. The result follows from Lemma 1. □

Lemma 2 is one of the two principal observational tools we use to extract our results. This characterization of Gabriel faces in a tetrahedron facilitates an understanding of the four point configurations that obstruct Gabriel meshes.

Lemma 3 *Every tetrahedron has at least two Gabriel faces.*

Proof Suppose σ does not have two Gabriel faces. Lemma 2 implies that c_σ lies in the intersection of the exterior half spaces of three faces of σ. Let v be the vertex common to the three faces. Then c_σ is contained within a space exterior to σ and defined by a solid angle opposite to (thus congruent to) the solid angle of σ at v. Since the tetrahedron is convex, this solid angle is less than 2π and it follows that a vector from c_σ to v would point towards the interior of σ. But the direction of the exterior normal vector of S_σ at v coincides with that of the vector from c_σ to v, and cannot point towards the interior of σ. Thus we have a contradiction. □

5 The anatomy of a hinge

We will examine the distinction between Gabriel meshes and Delaunay meshes by focusing on edges and their adjacent faces:

hinges. A mesh is not a Delaunay mesh if any of its edges is not locally Delaunay. We will compare this local Delaunay criterion with the natural corresponding condition that edges in a Gabriel mesh must possess. We say that e obtains a *Gabriel certificate* from each of its adjacent faces that is Gabriel in the associated flip-tet. Edge e will be *locally Gabriel* only if it has two Gabriel certificates.

In all that follows, the hinge will consist of edge $e = [p, q]$ and adjacent triangles $t_1 = [p, q, u]$ and $t_2 = [q, p, v]$. It is convenient to work with spheres rather than balls. Denote by S_1 and S_2 the boundaries of the diametric balls B_{t_1} and B_{t_2}. The centres of these spheres are denoted c_1 and c_2 respectively. The flip-tet associated with the hinge is σ. Its circumsphere is S_σ with centre c_σ.

Faces t_1 and t_2 have *consistent normals* if they are both oriented towards their respective interior half-spaces, or both towards their exterior half-spaces with respect to σ. A hinge is *non-sharp* if the angle between the consistent normals is less than $\pi/2$. In all that follows we consider only non-sharp hinges. Since our interest is in smooth meshes, this is not a limitation.

5.1 Equivalent hinges

To test the locally Delaunay property of e it is sufficient to check the sum of the angles it subtends. Two different hinges on e will yield the same result on the Delaunay test if their angles are the same. This reflects the fact that what is important is the circumcircles of the triangles, not the triangles themselves.

For the Gabriel property, not only are the angles subtended by e important, but also the dihedral angle between t_1 and t_2. Given two circumcircles on pivot e, the dihedral angle will affect S_σ and the relative position of c_σ. We say two hinges, (t_1, e, t_2) and (t'_1, e, t'_2) on e are *equivalent* if the dihedral angle between t_1 and t_2 is the same as that between t'_1 and t'_2 and the angle subtended by e in t_i is equal to that in t'_i, $i = 1, 2$.
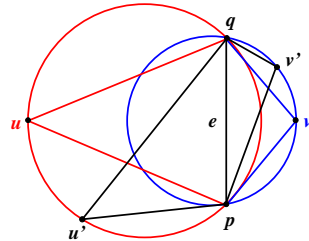
Equivalent hinges

Thus equivalent hinges have the same S_σ, and the affine hulls of their faces coincide. The circumcircles (intersections of the affine hulls with S_σ) are the same, as are the interior and exterior half-spaces. Equivalent hinges yield the same Gabriel certificates and have the same locally Delaunay status. Note that this does not imply equivalent consequences if we were to perform an edge flip. We will not be concerned with the flip-tet itself until Section 7.

5.2 Cross-sectional diagrams

Having established that we only need to consider equivalence classes of hinges, we choose the most convenient representative from each class for the purposes of visualization and analysis. A hinge on e is the *canonical representative* of its equivalence class if u and v lie on the perpendicular bisector plane of e.

This bisector plane defines a cross-section of the hinge and its circumsphere. A *cross-sectional diagram* is our visualization of these cross sections which we now describe. We caution that some objects in a cross-sectional diagram are one dimensional representations of their higher dimensional counterparts in the hinge, yet we

use the same symbol for both. Thus, for example, t_1 may be a segment in the diagram or a triangle in the hinge: we rely on context to resolve the ambiguity.

Refer to Figure 1. The circumsphere of the hinge, S_σ, is depicted as the large solid circle. Triangle t_1 is depicted as a solid segment between u and point e. The latter point is the centre of the pivot edge of the hinge. The circumcircle of t_1 is represented by its diameter, the segment $[u, u^*]$. The portion of this diameter that does not contain t_1 is drawn with a dashed line and denoted t_1^*. If e subtends an acute angle in t_1, then $|t_1| > |t_1^*|$ in the diagram. The other triangle in the hinge, t_2, is represented in the same way.

Figure 1 is used as a visual aid to the proof of Lemma 6 which says that if e has two Gabriel certificates, then it is locally Delaunay. Appealing to Lemma 2, we see that the hinge drawn in the figure has two Gabriel certificates because c_σ lies in the intersection of the interior half-spaces of t_1 and t_2. Indeed, it follows directly from the definition that the interior half-space of t_1 is represented in our diagram by the half-plane bounded by the affine hull of t_1 and containing t_2.

The Delaunay condition of a hinge can be checked by unfolding the faces into a common plane and determining if the circumcircle of one triangle contains the other triangle. In terms of our diagram, Figure 1, we could rotate t_2 clockwise on e until it lies in the affine hull of t_1 and then check if v in this new position lies outside of S_σ. In other words e will be locally Delaunay if u^* lies inside the circle of radius $|t_2|$ centred at e. This is just a check on the relative lengths of t_2 and t_1^*. Thus, an equivalent check is to test whether v itself lies *outside* the circle of radius $|t_1^*|$ centred at e. Since this is the case in Figure 1, the edge e shown there is locally Delaunay.

This leads to another observation we can make from the diagrams. When we draw a hinge, (t_1, e, t_2), in this fashion we at the same time draw another hinge (t_1^*, e, t_2^*) represented with dotted lines. We call this hinge the *complementary hinge*. The dihedral angle of a hinge is the same as that of its complementary hinge. If t_1 is acute, then t_1^* will be obtuse (it has the complementary angle subtended by e). Also, if t_1 yields a Gabriel certificate to e, then t_1^* will not, and vise versa. If e is locally Delaunay with respect to (t_1, e, t_2), then it will be not locally Delaunay with respect to (t_1^*, e, t_2^*), and vise versa. Thus results about the former hinge directly yield "complementary results" about the latter.

So, for example, the first part of Lemma 5 says that it is impossible to draw a non-sharp hinge with t_1 and t_2 acute without having c_σ contained in the intersection of their interior half-spaces. This immediately gives us the complementary result that it is impossible to draw a non-sharp hinge with t_1^* and t_2^* obtuse and such that c_σ is contained in either of their interior half-spaces.

6 Relating Gabriel and Delaunay properties

We now examine the relationship between Gabriel certificates yielded to e and the locally Delaunay property of e. We emphasise that our results here apply only to *non-sharp* edges. We find the results naturally separate into two cases. If t_1 and t_2 both agree on the decision of whether or not to yield a Gabriel certificate, then we call it a *symmetric case*. We find there is no distinction between the locally Delaunay condition and the locally Gabriel condition here. If the faces don't agree, then we have an *asymmetric case*, and it is within this realm that we find the local distinction between Delaunay meshes and Gabriel meshes.

We will make reference to the following observation:

Lemma 4 *Let ℓ be a chord in circle S_σ and e a point on it. If c is the centre of ℓ, then $\angle cec_\sigma$ is acute.*
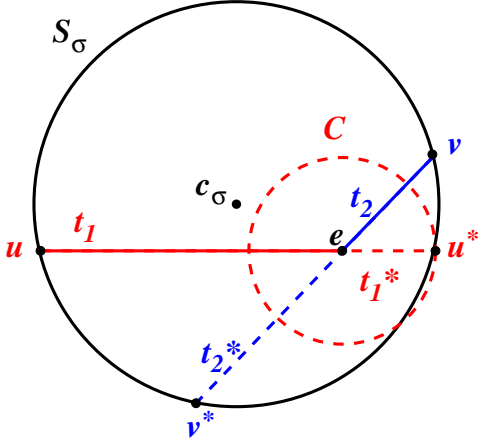
Figure 1: *Lemma 6: If e is locally Gabriel, then it is locally Delaunay.*

Figure 2: *Lemma 7: If an acute triangle yields a Gabriel certificate, so too must the other triangle.*

6.1 Symmetric cases

We observe that much can be said by simply examining the angles subtended by e:

Lemma 5 *If both angles subtended by e are acute, then e has two Gabriel certificates. If these angles are both obtuse, then e has no Gabriel certificates.*

Proof   Assume t_1 and t_2 have acute angles subtended by e. By Lemma 4, $[c_\sigma, e]$ makes an acute angle with both t_1 and t_2. Since the hinge is non-sharp, t_1 and t_2 cannot lie on the same side of the line supporting $[c_\sigma, e]$. Thus c_σ must lie in the intersection of the interior half-spaces of t_1 and t_2, and so by Lemma 2, e is locally Gabriel. The complementary result follows.   □

The intermediate case, where both subtended angles are $\pi/2$, occurs when S_1 and S_2 coincide (c_σ lies on e and S_σ is the diametric ball of both triangles), and so e will be locally Gabriel. Thus if the angles are both obtuse or both non-obtuse, locally Delaunay and locally Gabriel mean the same thing. But we can say more:

Lemma 6 *If e is locally Gabriel, then it is locally Delaunay. If e has no Gabriel certificates, then it is not locally Delaunay.*

Proof   Assume e is locally Gabriel. In light of Lemma 5, we only need to check the case where e subtends an acute angle in t_1 and an obtuse angle in t_2. We refer to Figure 1. We will show that v must lie outside of the circle with centre e and radius $|t_1^*|$. This circle, call it C, intersects S_σ at u^*. If c_σ lies on t_1 then C will be tangential to S_σ at u^* and thus lie entirely within S_σ and e will be locally Delaunay. So assume c_σ lies in the interior of the interior half-space of t_1. In this case C will intersect S_σ transversely at u^*. Furthermore, in a neigbourhood of u^* it will be inside S_σ in the interior half-space of t_1. Since $|t_1^*| < |t_1|$ by hypothesis, C must intersect t_1 and thus must remain inside S_σ within the interior half-space of t_1. Since t_2 lies within this interior half-space, C must intersect t_2 and thus e is locally Delaunay. The complementary result follows.   □

Every edge in a Gabriel mesh must necessarily be locally Gabriel. We reiterate the implication of Lemma 6 for emphasis:
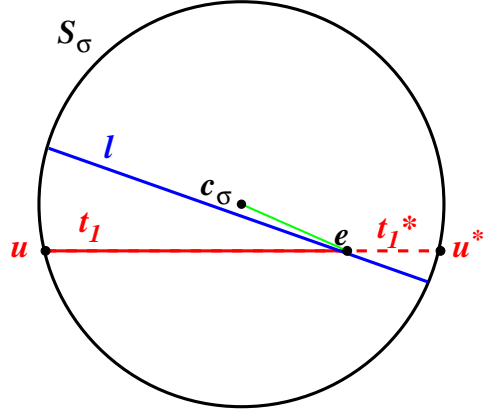
Theorem 2 *A Gabriel mesh is a Delaunay mesh.*

6.2 Asymmetric cases

If e has a solitary Gabriel certificate, it must come from a triangle with an obtuse angle subtended by e:

Lemma 7 *Suppose t_1 has an acute angle subtended by e. If t_1 yields a Gabriel certificate to e, then so too must t_2.*

Proof   Assume t_1 yields a Gabriel certificate to e, so c_σ lies in its interior half-space. We will show that t_2 must also yield a Gabriel certificate to e. Suppose to the contrary that c_σ lies in the exterior half-space of t_2, and refer to the diagram of Figure 2. Let ℓ be the segment composed of t_2 and t_2^*. Then ℓ must separate c_σ and t_1, since by definition t_1 lies within the interior half-space of t_2. By Lemma 4, and our assumption on t_1, $[c_\sigma, e]$ makes an acute angle with t_1. Since the hinge is non-sharp, it must be t_2^* rather than t_2 that lies between $[c_\sigma, e]$ and t_1. But by definition t_2^* lies in the exterior half-space of t_1, contradicting the hypothesis that t_1 yields a Gabriel certificate to e.   □

When e has a single Gabriel certificate, it is not locally Gabriel, but may yet be locally Delaunay. The following observation characterizes the distinction. Let r_t denote the circumradius of t.

Lemma 8 *If edge e is subtended by an obtuse angle in triangle t_1 and an acute angle in t_2, then e is nlD if and only if $r_{t_1} > r_{t_2}$.*

Proof   Unfold the hinge so that it is planar. The result is a direct consequence of the observation that if two circles, C_1 and C_2, intersect such that they share a common chord, e, then C_1 contains C_2 on one side of the line supporting e, and C_2 contains C_1 on the other side. Further, if the centres of both circles lie on the same side of e, then the larger circle contains the other on that side. If C_1 is the circumcircle of t_1, it's centre will be on the same side of e as t_2, and if C_1 contains C_2 on that side, e is nlD.   □

7 Obstructions to Gabriel meshes

In this section we show that given a point set P, which may be nicely sampled from a smooth surface, it is in general not possible to construct a Gabriel mesh whose vertex set is P. We proceed by constructing a flip-tet whose opposing non-sharp hinges each have only a single Gabriel certificate. This is done by placing the tetrahedron, σ, at the equatorial plane, but to the side in such a way
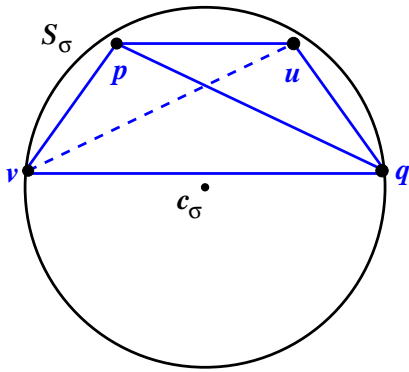
Figure 3: *Looking down on an obstruction to a Gabriel mesh. The figure is projected onto the equatorial plane. The non-sharp hinge on edge $[p, q]$ is on top. Edge $[v, q]$ lies on the equatorial plane while p lies just above it and u lies just below it. Thus c_σ is in the interor half-space of neither $[p, v, q]$ nor $[u, v, q]$. It follows that neither $[p, q]$, nor its opposing edge $[u, v]$ have two Gabriel certificates: neither edge is locally Gabriel.*

that c_σ is exterior to σ, but close to a long sharp edge that is almost a diameter of S_σ. The tetrahedron is described in Figure 3.

This obstruction cannot be avoided by imposing any reasonable uniformity constraints on the sample set. A uniformity constraint imposes a lower limit on the distance between neighbouring sample points and thus imposes a lower bound on the edge lengths of triangles in a mesh. This uniformity constraint is coupled with the sampling radius which puts an upper bound on the triangle circumradius in any Delaunay-based reconstruction. Thus a uniformity constraint serves to put an upper bound on the ratio of the circumradius to shortest edge of the triangles. See [Dey 2007] for details. Note, however, that the shortest edge in the tetrahedron in Figure 3 can be made to be arbitrarily close in length to the radius of S_σ. Also, the hinges can be made arbitrarily flat.

If a Gabriel mesh exists on a sample set P it need not be unique (consider a flip-tet that contains its circumcentre; both hinges are locally Gabriel, so it may be possible to flip that edge and obtain a different Gabriel mesh). Therefore even if we construct a mesh containing this obstruction and in which all faces outside of the obstruction are Gabriel, it doesn't immediately imply that a Gabriel mesh cannot be constructed: there may be a different Gabriel mesh in which no triangle is composed solely of vertices from σ. However, a Gabriel mesh is also a substructure of the Delaunay tetrahedralization, and as such we can constrain the possibilities by creating a sample set, P, in which σ is the only sliver tet in the Delaunay tetrahedralization. In this way we see that a closed Gabriel mesh need not exist. A figure showing a similar obstruction embedded in a mesh of nice triangles can be found in [Chaine 2003].

8 Conclusions

By analyzing the distinction between Gabriel meshes and Delaunay meshes we gain insight into both structures. The Delaunay mesh criteria can be viewed as a relaxation, if not a generalization of the Gabriel mesh criteria. The obstructions which are problematic for Gabriel meshes are not an issue for Delaunay meshes. Delaunay meshes may in fact be the structures to provide the role that has repeatedly been sought in Gabriel meshes; a mesh with natural local connectivity criteria which are independent of any reference surface or ambient Delaunay structure.

However, although the Delaunay edge flipping algorithm presented in [Dyer et al. 2007] was shown to be well defined and guaranteed to terminate, a constructive existence proof for Delaunay meshes on a given fixed vertex set P was not provided. The edge flipping algorithm could encounter an *unflippable edge*: one whose opposing edge already exists in the mesh.

In future work we wish to quantify the sampling conditions that will ensure the existence of a smooth Delaunay mesh, and also the criteria necessary to ensure that the edge flipping algorithm will maintain sufficient mesh smoothness.

References

ADAMY, U., GIESEN, J., AND JOHN, M. 2000. New techniques for topologically correct surface reconstruction. In *IEEE Visualization*, 373–380.

AMENTA, N., CHOI, S., DEY, T. K., AND LEEKHA, N. 2000. A Simple Algorithm for Homeomorphic Surface Reconstruction. In *Symp. Comp. Geom.*, 213–222.

ATTENE, M., AND SPAGNUOLO, M. 2000. Automatic surface reconstruction from point sets in space. *Computer Graphics Forum 19*, 457–465.

BOBENKO, A. I., AND SPRINGBORN, B. A. 2007. A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete and Computational Geometry 38*, 4, 740–756.

CHAINE, R. 2003. A geometric convection approach of 3-D reconstruction. In *Symp. Geometry Processing*, 218–229.

CHENG, S.-W., AND DEY, T. K., 2007. Delaunay edge flips in dense surface triangulations. arXiv.org:0712.1959v1.

DESBRUN, M., HIRANI, A. N., LEOK, M., AND MARSDEN, J. E. 2005. Discrete exterior calculus. arXiv:math.DG/0508341.

DEY, T. 2007. *Curve and Surface Reconstruction; Algorithms with Mathematical Analysis*. Cambridge University Press.

DYER, R., ZHANG, H., AND MÖLLER, T. 2007. Delaunay mesh construction. In *Symp. Geometry Processing*, 271–282.

DYER, R., ZHANG, H., AND MÖLLER, T. 2008. Observations on Gabriel meshes and Delaunay edge flips. Tech. Rep. TR 2008-22, Simon Fraser University.

EDELSBRUNNER, H., AND SHAH, N. R. 1994. Triangulating topological spaces. In *Symp. Comp. Geom.*, 285–292.

FISHER, M., SPRINGBORN, B., BOBENKO, A. I., AND SCHRÖDER, P. 2006. An algorithm for the construction of intrinsic Delaunay triangulations with applications to digital geometry processing. In *SIGGRAPH Courses*, 69–74.

GLICKENSTEIN, D., 2005. Geometric triangulations and discrete Laplacians on manifolds. arxiv:math.MG/0508188v1.

PETITJEAN, S., AND BOYER, E. 2001. Regular and non-regular point sets: Properties and reconstruction. *Computational Geometry: Theory and Applications 19*, 2-3, 101–126.

PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics 2*, 15–36.

WARDETZKY, M., MATHUR, S., KÄLBERER, F., AND GRINSPUN, E. 2007. Discrete Laplace operators: no free lunch. In *Symp. Geometry Processing*, 33–37.