Minimizing Absolute Gaussian Curvature Locally

Joachim Giesen⋆ and Madhusudan Manjunath†
⋆Friedrich-Schiller-Universität Jena
†Max-Planck Institut für Informatik

Abstract

One of the remaining challenges when reconstructing a surface from a finite sample is recovering non-smooth surface features like sharp edges. There is practical evidence showing that a two step approach could be an aid to this problem, namely, first computing a polyhedral reconstruction isotopic to the sampled surface, and secondly minimizing the absolute Gaussian curvature of this reconstruction globally. The first step ensures topological correctness and the second step improves the geometric accuracy of the reconstruction in the presence of sharp features without changing its topology. Unfortunately it is computationally hard to minimize the absolute Gaussian curvature globally. Hence we study a local variant of absolute Gaussian curvature minimization problem which is still meaningful in the context of surface fairing. Absolute Gaussian curvature like Gaussian curvature is concentrated at the vertices of a polyhedral surface embedded into $\mathbb{R}^3$. Local optimization tries to move a single vertex in space such that the absolute Gaussian curvature at this vertex is minimized. We show that in general it is algebraically hard to find the optimal position of a vertex. By algebraically hard we mean that in general an optimal solution is not constructible, i.e., there exist no finite sequence of expressions starting with rational numbers, where each expression is either the sum, difference, product, quotient or k'th root of preceding expressions and the last expressions give the coordinates of an optimal solution. Hence the only option left is to approximate the optimal position. We provide an approximation scheme for the minimum possible value of the absolute Gaussian curvature at a vertex.

1 Introduction

The problem to reconstruct a surface from a finite sample has been extensively studied for over a decade. From a theoretical point of view the smooth case, i.e., reconstruction of a closed surface embedded smoothly into $\mathbb{R}^3$, is well understood by now. That is, for the smooth case, algorithms are known that are guaranteed to compute a polyhedral surface (surface mesh) from a dense enough sampling that is (1) isotopic to the sampled surface, and (2) closely approximates first order differential properties, i.e., quantities that involve first order derivatives of the embedding like surface normals. For an excellent overview of the state of the art in surface reconstruction with theoretical guarantees see the recent book by Dey [10]. Challenges beyond the smooth case include noise, surfaces with boundaries, the reconstruction of higher order differential properties of the (smooth) embedding, and reconstructing non-smoothly embedded surfaces. This paper contributes to a better understanding of the challenge mentioned last, namely reconstructing non-smoothly embedded surfaces. For the other challenges, noise has been addressed in [12, 16], boundaries have been addressed in [14, 11], and higher order differential properties have been addressed in [6, 8].

Chazal, Cohen-Steiner and Lieutier [7] have presented a sampling theory for compact subsets of Euclidean space (which includes surfaces embedded non-smoothly into Euclidean space). Under this sampling theory a homotopic copy of the sampled compact subset can be obtained from an offset of the sampling. A homotopic reconstruction may be not enough for practical applications, but nevertheless the result by Chazal et al. suggests a two step approach [1] towards reconstructing non-smoothly embedded surfaces: the first step ensures topological correctness of a computed (polyhedral) reconstruction and the second step improves its geometric accuracy without changing its topology. The second step is often referred to as post-processing or mesh faring. Most fairing schemes used in practice aim at optimizing curvature criteria. In the differential geometry of surfaces, embedded in $\mathbb{R}^3$ two curvatures play a major role. First, the mean curvature, which is the average of the two principal curvatures at

Dagstuhl Seminar Proceedings 09111
Computational Geometry
http://drops.dagstuhl.de/opus/volltexte/2009/2031

a surface point, and second, the Gaussian curvature, which is the product of the two principal curvatures. The definitions of Gaussian- and mean curvature can be extended to polyhedral surfaces though this extension is not straightforward: the Gaussian curvature is now concentrated in the vertices of the surface mesh. Integrating over the curvatures gives the total mean and total Gaussian curvature, respectively. One of the most elegant theorems in geometry, the Gauss-Bonnet theorem, states that the total Gaussian curvature is a topological invariant of the surface, i.e., it does not depend on the embedding of the surface. This is different for the total mean curvature and might be the reason that many surface fairing schemes aim at minimizing the total mean curvature of either the whole surface or of a surface patch. As a topological invariant the total Gaussian curvature cannot serve this purpose though it is still possible (and often meaningful) to re-distribute the Gaussian curvature on the surface. In contrast to the total Gaussian curvature the total absolute Gaussian curvature, i.e., the integral over the absolute value of the Gaussian curvature, depends on the embedding. One property that makes minimizing the total absolute Gaussian curvature interesting for mesh fairing is the fact that all surfaces with minimal total absolute Gaussian curvature have the so called *two-piece property*, i.e., any hyperplane cuts them into at most two pieces. Reducing the number of components that a surface can be cut into by a hyperplane can smooth a surface significantly, see Figure 1 for an example. The connection between the absolute Gaussian curvature and the two-piece property suggests that reducing the absolute Gaussian curvature globally or locally can be a means to achieve similar smoothing.

In this paper we focus on minimizing the absolute Gaussian curvature of a polyhedral surface locally. That is, we want to move a single vertex such that the absolute Gaussian curvature at this vertex is minimized. It turns out that this problem is hard in a special sense: in general there exist no finite sequence of expressions starting with rational numbers, where each expression is either the sum, difference, product, quotient or k'th root of preceding expressions and the last expressions give the coordinates of an optimal solution. Bajaj [3] proved an analogous result for the Fermat-Weber problem. Interestingly, our curvature minimization problem can be seen in a certain sense (the analogy does not hold completely) as the Fermat-Weber problem for angles and line segments: instead of finding a point that minimizes the sum of distances to a given set of points, we essentially want to find a point that minimizes the sum of angles at this point in the triangles spanned by the point and the given set of line segments. The hardness proof for the curvature minimization problem turns out to be more involved than the proof for the Fermat-Weber problem due to the following difficulties: (1) the objective function (absolute Gaussian curvature) is continuous everywhere but not everywhere differentiable. So, in contrast to the Fermat-Weber problem a vanishing gradient condition cannot be used directly, and (2) the objective function is not convex. As a consequence, even if the function would be differentiable, the vanishing gradient condition is satisfied at all the local minima and hence does not characterize the global minima. Besides algebraic results we make use of a combination of techniques from integral geometry and real analysis to overcome these difficulties. The only option left after our hardness result is to approximate a point that minimizes the absolute Gaussian curvature. We provide an approximation scheme for the value of the absolute Gaussian curvature at an optimal point. Our approximation scheme also makes use of ideas from integral geometry.

2 Motivation and Definitions

As mentioned in the introduction, this work is motivated by smoothing polyhedral surfaces embedded in $\mathbb{R}^3$.

Polyhedral surfaces. A *polyhedral surface* is the geometric realization of a simplicial complex in $\mathbb{R}^3$, whose underlying topological space is a surface without boundary. The *star* of a vertex of a polyhedral surface is the union of all faces that contain the vertex and the boundary of the star is the *link* of the vertex. Since we are only dealing with surfaces without boundary all the vertex links are polygonal embeddings of the 1-sphere $\mathbb{S}^1$ into $\mathbb{R}^3$.

Surface reconstruction. Polyhedral surfaces are often computed as output of a surface reconstruction algorithm [10]. In the surface reconstruction problem we are given a finite sampling V of a surface $\Sigma \subset \mathbb{R}^3$, i.e., $V \subset \Sigma$. The goal is to compute a polyhedral surface S with vertex set V that is topologi-

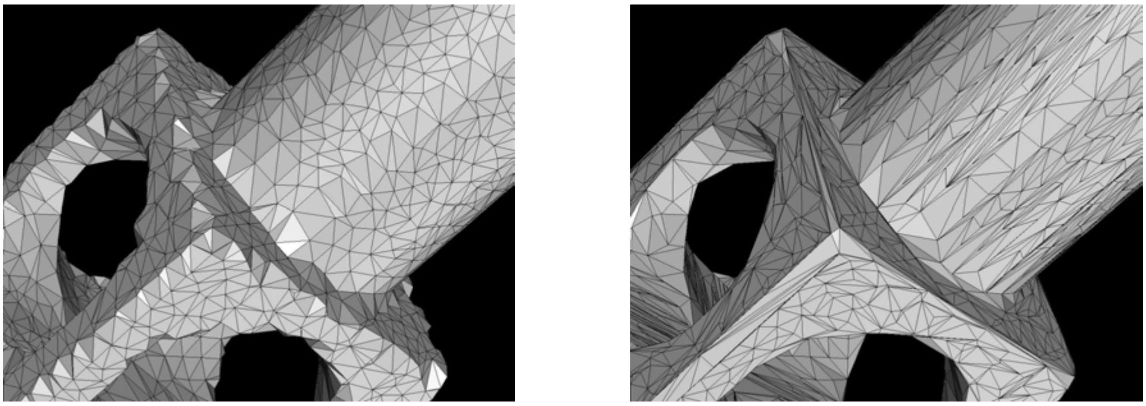

Figure 1: A typical reconstruction of the MECHPART data set (observe that the jaggy pattern at an edge can be cut into many pieces by a hyperplane) and the same surface optimized with respect to absolute Gaussian curvature using a flip heuristic (which gets stuck in a local optimum in general). [Images courtesy of P. Gehr]

cally equivalent and geometrically close to Σ. As of today, there are a number of algorithms [10] that achieve this goal if Σ is closed and smoothly embedded, and if V is a dense enough sample. Interestingly, these algorithms do not fail completely if Σ contains 'sharp edges' (in the sense that they compute a polyhedral surface with the correct topology) though the sharp edges are not preserved. This suggests trying to post-process (smooth) the output of such a reconstruction in order to also reconstruct also the sharp edges. At the sharp edges one often observes a jaggy pattern, see Figure 1. For such a jaggy pattern there is a hyperplane that cuts the surface into many pieces. Hence a promising approach for post-processing seems to be to compute a 'tighter' polyhedral surface that has the same topology and vertex set as the original reconstruction. Here 'tighter' loosely speaking means that the new surface cannot be cut into as many pieces as the original reconstruction.

Tight surface. A closed polyhedral surface is called *tight* if any hyperplane cuts it into at most two pieces (*two-piece-property*). Note that any surface that bounds a convex body is tight. Since a surface of higher genus, e.g., a torus, cannot bound a convex shape, the notion of tightness can be seen as a generalization of the concept of convexity to surfaces of higher genus. Tightness is closely related to the total absolute Gaussian curvature as it turns out that the total absolute Gaussian curvature of a tight closed surface, i.e., the integral of the absolute Gaussian curvature over the surface embedded in $\mathbb{R}^3$, is minimal among all surfaces of the same genus.

Absolute Gaussian curvature [4]. The *Gaussian curvature* of a polyhedral surface S is defined at its vertices. Let v be a vertex of S. For any triangle T_i in S incident to v let α_i be the angle of T_i at v. The Gaussian curvature K_v of S at v is defined as $K_v = 2\pi - \sum_i \alpha_i$. Also the *absolute Gaussian curvature* is defined at the vertices of S as the sum of the *positive curvature* K_v^+ and *negative curvature* K_v^- at a vertex v of S. Let us first define the positive curvature. We distinguish two types of vertices v depending on whether S has a locally supporting hyperplane at v or not. We say that S has a locally supporting hyperplane at v if there exists a neighborhood U of v in S such that U is completely contained in a closed half-space bounded by the hyperplane. If S does not have a locally supporting hyperplane at v we set $K_v^+ = 0$. Otherwise, let T_j' be a triangle in the star of v on the boundary of the convex hull of v and all vertices of S incident to v. Let β_j be the angle of T_j' at v. One defines $K_v^+ = 2\pi - \sum_j \beta_j$, see Figure 2.

The negative curvature at any vertex v of S is now defined as $K_v^- = K_v^+ - K_v$ and the absolute Gaussian curvature of S at v is defined as

$$K_v^* = K_v^+ + K_v^-.$$

There is also an integral geometry formulation of absolute Gaussian curvature at the vertex v of a polyhedral surface (which makes the relationship with tightness and the two-piece property more apparent):

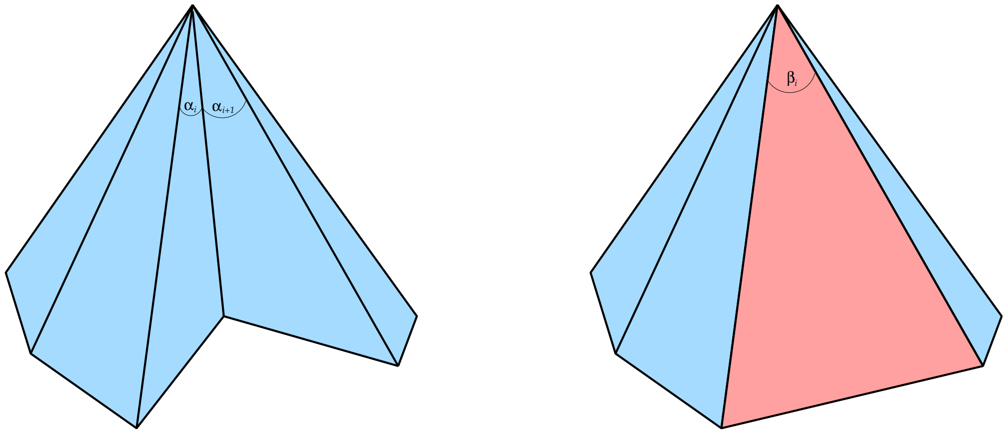

Figure 2: Illustrating the definition of angles α_i (on the left) and β_j (on the right) at the vertex on top.

let C_v be the link of v, for any given direction $d \in \mathbb{S}^2$ let $H_{d,v}$ be the hyperplane passing through v with normal d and let $n_{d,v} = |H_{d,v} \cap C_v|$ be the number of intersections of $H_{d,v}$ with C_v. For almost all directions (up to a set of directions of measure 0 on $\mathbb{S}^2$) the number $n_{d,v}$ is even and $i_{d,v} = n_{d,v}/2$ is called the index of v in direction d. The total curvature at v can be written as

$$K_v^* = \frac{1}{2} \int_{d \in \mathbb{S}^2} |1 - i_{d,v}| \, do.$$

The close connection of tightness and total absolute Gaussian curvature suggests that mesh smoothing for polyhedral surfaces obtained from a surface reconstruction algorithm can be done by transforming the polyhedral surface S into a polyhedral surface S' with smaller total absolute Gaussian curvature. One possible transformation is via constrained optimization, namely, compute S' with minimal total absolute Gaussian curvature such that S' shares the vertex set with S and has the same topology as S. In [5] it is shown that this optimization problem is NP-hard in the case of terrain surfaces. Here we want to take another approach that is quite common in mesh processing, namely local smoothing.

Mesh smoothing. Probably the most commonly used mesh smoothing technique is Laplacian smoothing, where a vertex v is relocated to the average position of the vertices in C_v, i.e., if $v_1, \ldots, v_n$ are the vertices in C_v the new position of v is given as $\frac{1}{n} \sum_{i=1}^{n} v_i$. Since Laplacian smoothing does not always work satisfactorily subsequently other relocation schemes optimizing geometric quantities like the mean curvature were employed [18, 2, 9, 17]. Here we want to transform the polyhedral surface locally by relocating any vertex v of the polyhedral surface such that the absolute Gaussian curvature at v with respect to the link C_v of v is minimized. Note that the absolute Gaussian curvature at a vertex v is completely determined by the location of v and the link C_v. Hence we can state the problem that we want to address in this paper as follows:

Absolute Gaussian curvature minimizing problem. Given a polygonal embedding C of $\mathbb{S}^1$ (or short polygon) into $\mathbb{R}^3$ with n vertices $v_1, \ldots, v_n$, find a position $v \in \mathbb{R}^3 \setminus C$ at which the absolute Gaussian curvature with respect to C is minimized, i.e., when C is considered to be the link of v in a polyhedral surface. In contrast to the related Fermat-Weber problem the solution to the problem need not be unique as demonstrated by the simple example in Figure 3, where the solution space is not even connected. This could make it even hard to approximate a solution to the curvature minimization problem.

In our application, i.e., mesh smoothing, we actually can assume that C is unknotted and that a non-self intersecting surface patch can be glued into C. The hardness result we are going to prove already holds in this special case.

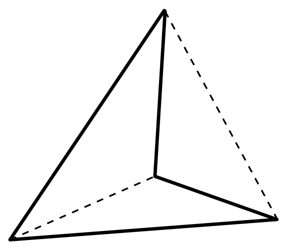

Figure 3: Given the embedding of $\mathbb{S}^1$ by the solid polygon into $\mathbb{R}^3$, then any point on the dashed lines, where the absolute Gaussian curvature is zero, is an optimal solution to the curvature minimizing problem.

3 Algebraic Hardness

Our hardness proof extends the approach taken by Bajaj in his seminal paper [3] where he established the algebraic hardness of the Fermat-Weber problem. We are going to prove the following theorem.

Theorem 1 *The solution to the absolute Gaussian curvature minimization problem is not constructible in general.*

For the proof we will show the existence of a polygon C in $\mathbb{R}^3$ whose vertices have rational coordinates and that has a unique solution to the curvature minimization problem. The solution can be described by the root of a univariate polynomial in $\mathbb{Z}[t]$. We will show that this solution is not constructible, i.e., the problem is *algebraically hard*.

3.1 Non-Constructibility

Definition 1 *An algebraic number $\hat{t}$ (root of a polynomial $p(t) \in \mathbb{Q}[t]$) is not constructible, if there is no finite sequence of expressions $e_1, \ldots, e_n$ with $e_1 \in \mathbb{Q}$ where each e_i is either the sum, difference, product, quotient or k'th root of preceding expressions and $e_n = \hat{t}$.*

Let us briefly summarize Bajaj's approach [3] to prove the algebraic hardness of the Fermat-Weber problem. A first step is the following characterization of non-constructible algebraic numbers.

Lemma 1 *Let be $p(t) \in \mathbb{Q}[t]$ be an irreducible polynomial. No root of $p(t)$ is constructible if and only if the Galois group of $p(t)$ is not solvable.*

The Galois group of a polynomial $p(t) \in \mathbb{Q}[t]$ is a permutation group of its roots, i.e., a subgroup of some symmetric group. For the roots of $p(t)$ to be constructible it is necessary that $p(t)$ is *solvable by radicals* over $\mathbb{Q}$, but this is equivalent for the Galois group of $p(t)$ to be a solvable group [15]. It turns out that the symmetric group S_k, i.e., the group of all permutations of a k-element set, is not solvable for $k > 4$. We show the existence of a polygon C such that any solution to the curvature minimization problem for C can be described by the solution of an irreducible, univariate polynomial in $\mathbb{Z}[t]$ whose associated Galois group is S_k for some $k > 4$.

3.2 Constructing the Polygon

We construct an infinite family of polygons and show that this family contains at least one polygon for which the solution of the curvature minimization problem is not constructible. The family is derived from a base polygon C_0 with six vertices

$$\begin{aligned} v_1 &= (1,-1,0), & v_2 &= (1,0,1), \\ v_3 &= (1,1,0), & v_4 &= (-1,1,0), \\ v_5 &= (-1,0,1), & v_6 &= (-1,-1,0). \end{aligned}$$

Edge e_i is the line segment connecting v_i and $v_{(i+1) \pmod 6}$. For $\eta \in (0,1)$ we perturb the base polygon by introducing a new vertex $v_7^\eta = (\eta, 1-\eta, 0)$ and by replacing e_6 by e_7^η and e_8^η. The edge e_7^η connects the vertices v_6 and v_7^η, and the edge e_8^η connects the vertices v_7^η and v_1. The perturbed polygon is denoted as C_η. See Figure 4.

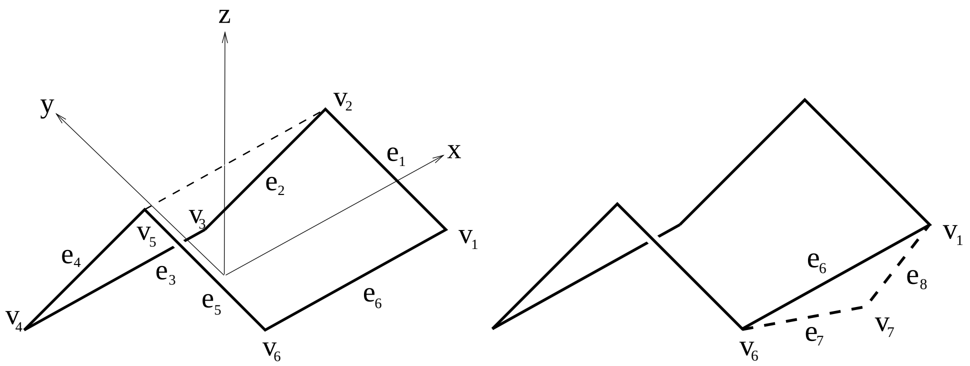

Figure 4: On the left, the base polygon is shown. The line segment ℓ connects v_2 and v_5 and is shown dashed. On the right, the base polygon gets perturbed by replacing the edge e_6 by the edges e_7 and e_8. The vertex v_7 is parameterized by $\eta \in (0,1)$.

At any point v in the relative open line segment ℓ that connects the vertices v_2 and v_5 the absolute Gaussian curvature $K^*(v)$ is 0 (just apply the angle definition for absolute Gaussian curvature), and it can be seen for any other point in $\mathbb{R}^3 \setminus (C_0 \cup \ell)$ the absolute Gaussian curvature is larger than 0 (just apply the integral geometry definition). Hence the solution to the curvature minimization problem for C_0 is ℓ. We will show that if the perturbation η is small enough, then the solution for curvature minimization problem for C_η also has to be in ℓ.

Theorem 2 *For sufficiently small $\eta \in (0,1)$, the solution to curvature minimization problem for C_η is contained in ℓ and there is no additional solution outside ℓ.*

In the following, we will not only establish the existence of a local minimum of K_η^* in a point on ℓ, but that a global minimum of K_η^* (for small enough η) lies on ℓ and there is no other global minimum outside ℓ. For this proof, we need two key lemmas. The statement of these lemmas makes use of the following notation: let e be a line segment in $\mathbb{R}^3$ (this will either be one of the edges e_1 to e_8, the closure of the line segment ℓ, or some line segment derived from a polygon C_η). For a point $v \in \mathbb{R}^3 \setminus e$, let $\alpha_e(v)$ be the angle at v in the triangle spanned by v and e. Note that some of the lines and angles depend on η. The gradients of angle functions $\alpha_e(v)$ are important for discussion, see also Figure 5. We want to restrict our discussion to a closed neighborhood $U_\varepsilon(\ell)$, that contains all points in $\mathbb{R}^3$ with distance at most $\varepsilon > 0$ to the closure of ℓ. Note that by the continuity of the absolute Gaussian curvature with respect to η (easy consequence from the integral geometry formulation) we know that for any $\varepsilon > 0$ there cannot exist a global minimum of the absolute Gaussian curvature K_η^* outside $U_\varepsilon(\ell)$ for η small enough. We denote the open ball of radius ρ centered at $u \in \mathbb{R}^3$ by $B_\rho(u)$ for $\rho > \varepsilon > 0$. Note that we always want to choose ε smaller than ρ.

Lemma 2 *The function $\alpha_e(v)$ is continuously differentiable with respect to the coordinates of v at all $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$.*

Proof. Simple calculation. □

Lemma 3 *(1) At all $v \in U_\varepsilon(\ell) \setminus$ aff(e) the norm of the gradient $\nabla\alpha_e(v)$ is continuously differentiable with respect to the coordinates of v and with respect to $\eta \in (0,1)$.*

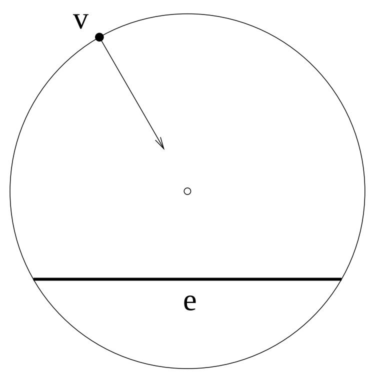

Figure 5: The gradient $\nabla\alpha_e(v)$ of the angle function is in the plane defined by v and e and points towards the midpoint of circum-circle of the triangle spanned by v and e.

(2) *Let u_1 and u_2 be the endpoints of e and let $\|u_1 - u_2\|/2 > \rho > 0$. Then the norm of the derivatives (with respect to the coordinates and with respect to η) of the norm of $\nabla\alpha_e(v)$ can be upper bounded by some constant c for all $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$.*

Proof. Simple calculation. □

Corollary 1 *The norm of $\nabla\alpha_e(v)$ is uniformly continuous in $U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$ with respect to the coordinates of v and with respect to $\eta \in (0, 1)$.*

The outline of the proof of Theorem 2 is now as follows:

(1) We will show that there exist $\varepsilon, \delta > 0$ such that the gradient of the norm of $\nabla K^*(v) > \delta$ for all $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$. Recall that v_2 and v_5 are the endpoints of ℓ (and vertices of every polygon C_η).

(2) We will show that for any given $\zeta > 0$ there exists $\varepsilon, \bar\eta > 0$ such that $\big\|\nabla K^*(v) - \nabla K^*_\eta(v)\big\| < \zeta$ for all $\eta < \bar\eta$ and all $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$.

It follows that there exist $\varepsilon, \delta, \bar\eta > 0$ such that $\nabla K^*_\eta(v) \neq 0$ for all $\eta < \bar\eta$ and all $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$. Hence K^*_η cannot have a local minimum in $U_\varepsilon(\ell)\setminus\big(B_\rho(v_2)\cup B_\rho(v_5)\big)$. Since we also know that for sufficiently small η the global optimum for K^*_η is obtained in $U_\varepsilon(\ell)$ we can conclude that for sufficiently small η the curvature minimization problem with respect to C_η and restricted to $\mathbb{R}^3 \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$ takes its minimum on $\ell \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$. Note that this result holds for any $\varepsilon > 0$, but we still need to establish that the optimal solution for small enough η is actually obtained in the interior of ℓ and not at the vertices v_2 and v_5.

Lemma 4 *There exists $\bar\eta > 0$ such that for $\eta < \bar\eta$, the minimum of $K^*_\eta(v)$ on the closure of ℓ is a point in the interior of ℓ.*

Proof. See appendix. □

We are now going to prove points (1) and (2) from the outline above in the form of two lemmas.

Lemma 5 *Given $\rho > 0$, there exist $\varepsilon, \delta > 0$ such that the gradient of the norm of $\nabla K^*(v) > \delta$ for all $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$.*

Proof. See appendix. □

Lemma 6 *Given $\zeta > 0$, there exists $\varepsilon, \bar\eta > 0$ such that*

$$\big\|\nabla K^*(v) - \nabla K^*_\eta(v)\big\| < \zeta$$

for all $\eta < \bar\eta$ and all $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$.

Proof. See appendix. □

3.3 The Univariate Polynomial

We are now going to describe the solution to the curvature minimization problem for C_η by the root of a univariate polynomial on ℓ. Let $v \in \mathbb{R}^3 \setminus C_\eta$. We denote by $\alpha_i(v)$ the angle in the triangle spanned by the edge e_i and v at the vertex v. For polygon C_η we find for any point $v \in \ell$ that

$$K_\eta^*(v) = K^*(v) - \alpha_6(v) + \alpha_7^\eta(v) + \alpha_8^\eta(v).$$

If we parameterize the line ℓ by $(-1,1) \ni t \mapsto x_t = (t,0,1) \in \ell$, then the existence of a local minimum of K_η^* in $v \in \ell$ implies

$$\frac{\partial}{\partial t} K_\eta^*(x_t) = 0 = -\frac{\partial}{\partial t}\alpha_6(x_t) + \frac{\partial}{\partial t}\alpha_7^\eta(v) + \frac{\partial}{\partial t}\alpha_8^\eta(v).$$

We want to transform each summand in the latter equation into an equivalent term of the form $g_1(t)/g_2(t)$, where $g_1, g_2^2 \in \mathbb{Q}[t]$. Here we explicitly show how to do it for the term $\frac{\partial}{\partial t}\alpha_6(x_t)$. Simple calculation shows that

$$\alpha_6(x_t) = \arccos\left(\frac{t^2+1}{\sqrt{(t+1)^2+2}\,\sqrt{(t-1)^2+2}}\right) = \arccos\left(\frac{f_1(t)}{\sqrt{f_2(t)f_3(t)}}\right),$$

where we used the shorthand notation $f_1(t) = t^2+1$, $f_2(t) = (t+1)^2+2$ and $f_3(t) = (t-1)^2+2$. This gives after another simple calculation

$$\frac{\partial}{\partial t}\alpha_6(x_t) = \frac{2f_1'(t)f_2(t)f_3(t) - f_1(t)\big(f_2'(t)f_3(t) + f_3'(t)f_2(t)\big)}{f_2(t)f_3(t)\sqrt{f_2(t)f_3(t) - f_1(t)^2}} = \frac{g_1(t)}{g_2(t)} =: q_1(t).$$

Obviously, $g_1, g_2^2 \in \mathbb{Z}[t]$. Similarly we get such expressions $q_2(t)$ and $q_3(t)$ for the terms $\frac{\partial}{\partial t}\alpha_7^\eta(v)$ and $\frac{\partial}{\partial t}\alpha_8^\eta(v)$, respectively, and the condition for a local minimum now reads as $q_1(t)+q_2(t)+q_3(t) = 0$. But then we also have

$$\begin{aligned}
0 &= \big(q_1(t)+q_2(t)+q_3(t)\big)\big(-q_1(t)+q_2(t)+q_3(t)\big) \\
&\quad \big(q_1(t)-q_2(t)+q_3(t)\big)\big(q_1(t)+q_2(t)-q_3(t)\big) \\
&= q_1^4(t)+q_2^4(t)+q_3^4(t) - 2q_1^2(t)q_2^2(t) - 2q_1^2(t)q_3^2(t) - 2q_2^2(t)q_3^2(t)
\end{aligned}$$

at a local minimum. The latter expression is a rational function for every $\eta \in (0,1)$. Note that $q_2(t)$ and $q_3(t)$ are parameterized by η though we have not made this explicit in our notation. One can check that the denominator of this rational function is not 0 for $t \in (-1,1)$ and hence can be discarded from our discussion. The numerator can be factored into four factors, three of which do not have real roots for $\eta \in [0,1)$. This leaves us for every $\eta \in (0,1)$ with a polynomial $p_\eta(t) \in \mathbb{Z}[t]$ of degree 16. For these polynomials we can show the following.

Theorem 3 *The Galois group of $p_{1/m}(t)$ is S_{10} for infinitely many $m \in \mathbb{Z}$.*

For the proof of Theorem 3, we will make use of a sufficient condition for the Galois group of a polynomial $p(x) \in \mathbb{Z}[x]$ of degree k to be the symmetric group S_k. The condition makes use of the notion of good primes (with respect to $p(x)$), that do not divide the *discriminant* of $p(x)$.

Lemma 7 (Bajaj [3]) *Let $p(t) \in \mathbb{Z}[t]$ of degree k with $k > 2$ even. If there exist good primes q_1, q_2 and q_3 such that the Galois groups of $p(t) \pmod{q_1}$ is cyclic and generated by a permutation that is a k-cycle, $p(t) \pmod{q_2}$ is cyclic and generated by a permutation that has a $(k-1)$-cycle, and $p(t) \pmod{q_3}$ is cyclic and generated by a permutation with a 2- and a $(k-3)$-cycle, then the Galois group of $p(x)$ is S_k.*

Note that here we choose $\eta = 1/m$ for $m \in \mathbb{Z}$. The polynomial $p_{1/m}(t)$ reads as follows (we used MAPLE to derive this representation in factorized form):

$$p_{1/m}(t) = \phi_1(t)\phi_2(t)\phi_3(t)\phi_4(t),$$

where

$$\begin{aligned}\phi_1(t) &= t^2+2t+3,\\ \phi_2(t) &= t^2-2t+3,\\ \phi_3(t) &= t^2m^2-2tm+2m^2+2m+2,\end{aligned}$$

and

$$\begin{aligned}&\phi_4(t) =\\ &\quad (4m^2)\, t^{10} +\\ &\quad (36m^2)\, t^9 +\\ &\quad (-12m^4-40m^3-16m+143m^2)\, t^8 +\\ &\quad (-110m^4-264m^3+332m^2-128m)\, t^7 +\\ &\quad (16+15m^6-272m^4-730m^3-460m+76m^5+530m^2)\, t^6 +\\ &\quad (682m^5+176m^6+112+1040m^2+28m^4-686m^3-836m)\, t^5 +\\ &\quad (28m^3+1507m^2+236-165m^4-48m^8-196m^7+126m^6-1216m+790m^5)\, t^4 +\\ &\quad (296+434m^3-624m^5-128m^8-324m^6+416m^2-1172m^4-1692m-368m^7)\, t^3 +\\ &\quad (-334m^5+698m^3-716m^2+12m^8-35m^4+156m^7+372-1036m+139m^6)\, t^2 +\\ &\quad (216+180m^3-116m^6-210m^5+154m^4-208m^2+96m+32m^8+48m^7)\, t +\\ &\quad (-9m^4+36m^2-36).\end{aligned}$$

A simple calculation shows that the equations $\phi_1(t) = 0$ and $\phi_2(t) = 0$ do not have real roots and that $\phi_3(t) = 0$ has no real roots for $m > 0$. Hence, the local minimum strictly contained in ℓ satisfies $\phi_4(t) = 0$. This leaves us to prove that $\phi_4(t)$ is not solvable by radicals over $\mathbb{Q}$ for infinitely many $m \in \mathbb{Z}$.

Lemma 8 *For any polynomial $p(t) \in \mathbb{Z}[t]$ we have $p(t+ln) \equiv p(t) \pmod n$, for all integers $l, n \in \mathbb{N}$.*

Proof. Let $p(t) = \sum_{i=1}^{k} a_i t^i$ with $a_i, k \in \mathbb{Z}$. By using the binomial theorem and taking the modulus with respect to n, we get

$$p(t+ln) \pmod n \equiv \sum_{i=0}^{k} a_i (t+kn)^i \pmod n \equiv \sum_{i=0}^{k} a_i t^i \pmod n \equiv p(t) \pmod n.$$

□

Lemma 9 *The Galois group of $\phi_4(t)$ for $m = 10$ is S_{10}.*

Proof. We show that the conditions of Lemma 7 are satisfied for the good primes $23, 29$ and 137 (we checked that these primes are good by MAPLE). For $m = 10$, the polynomial $\phi_4(t)$ reads:

$$\begin{aligned}\phi_4(t) &= 4\big(100t^{10}+900t^9-36465t^8-333020t^7+4799604t^6+60972438t^5-1639120806t^4\\ &\qquad -4219415256t^3+716466603t^2+886175094t-21609\big).\end{aligned}$$

Factorizing $\phi_4(t)$ from $m = 10$ over $\mathbb{Z}_{23}[t], \mathbb{Z}_{29}[t]$ and $\mathbb{Z}_{137}[t]$ we get

$$\begin{aligned}\phi_4(t) \pmod{23} &= 8(t+13)(t^2+3t+1)(t^7+16t^6+19t^5+5t^3+11t^2+15t+22)\\ \phi_4(t) \pmod{29} &= 13(t+4)(t^9+5t^8+17t^7+18t^6+23t^5+17t^4+21t^3+28t^2+15t+20)\\ \phi_4(t) \pmod{137} &= 108+82t+40t^2+105t^3+122t^4+119t^5+20t^6\\ &\quad +101t^7+111t^8+11t^9+100t^{10}\end{aligned}$$

The corresponding Galois groups are cyclic and generated by a permutation with a $2-$ and a 7-cycle, a permutation with a 9-cycle and a permutation with a 10-cycle, respectively. To see this note that the degrees of the factors give the order of the cycles [13]. The claim of the lemma now follows from Lemma 7. □

Lemma 10 *The Galois group of $\phi_4(t)$ is S_{10} for infinitely many $m \in \mathbb{Z}$.*

Proof. See appendix. □

The proof of the Theorem 3 now follows as an immediate corollary to Lemma 10.

This is, $p_{1/m}(t)$ is not solvable by radicals over $\mathbb{Q}$ for infinitely many $m \in \mathbb{Z}$, which in turn implies that a local minimum of $K^*_{1/m}$ in $v \in \ell$ is not constructible for infinitely many $m \in \mathbb{Z}$ by Lemma 1. By Theorem 2 we know that for sufficiently small η (sufficiently large m) a local optimum on ℓ actually is the global optimum for the curvature minimization problem. Hence we have shown that in general the solution to the curvature minimization problem is not constructible, i.e., we have proven Theorem 1.

4 Approximating the Value of the Optimal Solution

Our algebraic hardness result essentially rules out the construction of the exact solution to the curvature minimization problem (in a model of computation where the root of an algebraic equation is obtained only from arithmetic operations and the extraction of k'th roots). Hence the only option left is to approximate an optimal solution. Here we design an approximation scheme for the value of the absolute Gaussian curvature at an optimal solution. The idea behind our approximation scheme is to discretize the integral

$$K^*_v = \frac{1}{2}\int_{d\in\mathbb{S}^2} |1 - i_{d,v}|\, do.$$

for the absolute Gaussian curvature at any $v \in \mathbb{R}^3 \setminus C$, where the index $i_{d,v}$ is defined with respect to the polygon C. Let $v_1, \ldots, v_n$ be the vertices of C. The hyperplanes that contain v and v_i can be parameterized by a great circle (of unit normals) on the sphere of directions $\mathbb{S}^2$. Hence the vertices $v_1, \ldots, v_n$ yield an arrangement $\mathcal{C}_v$ of great circles $C_1, \ldots, C_n$ on $\mathbb{S}^2$.

Lemma 11 *The directions contained in the interior of a cell of the great circle arrangement $\mathcal{C}_v$ on $\mathbb{S}^2$ have all the same index $i_{d,v}$.*

Proof. See appendix. □

Lemma 11 immediately implies that $K^*_v = \frac{1}{2}\sum_{Z \text{ cell of } \mathcal{C}_v} w(Z)|1 - i_{d_Z,v}|$, where $w(Z)$ is the area of the cell $Z \subset \mathbb{S}^2$ and $i_{d_Z,v}$ is the index of v in an arbitrary direction $d_Z \in Z$. Next we are going to choose a finite number of directions $d_1, \ldots, d_m \in \mathbb{S}^2$. Let H_{ij} be the hyperplane in $\mathbb{R}^3$ with normal d_i that contains the vertex v_j, see Figure 6 for an example in the plane with two directions. There are in total nm such hyperplanes. The arrangement $\mathcal{H}$ of these hyperplanes contains $O\big((nm)^3\big)$ cells.

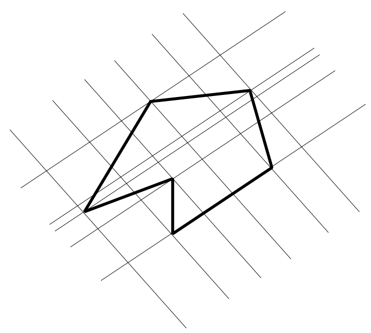

Figure 6: Hyperplane arrangement in the plane induced by two directions and the given polygon.

Lemma 12 *All points v in the same cell of $\mathcal{H}$ have the same index $i_{d_i,v}$ with respect to direction d_i for $i = 1, \ldots, n$.*

Proof. See appendix. □

We choose the directions $d_i \in \mathbb{S}^2$ as follows: start with a regular icosahedron with vertices on $\mathbb{S}^2$. The k'th order triangulation of the icosahedron partitions each face of the icosahedron into k^2 triangles. We project the vertices of these triangulations onto $\mathbb{S}^2$. Since the icosahedron has 12 vertices and 20 triangles, the k'th order triangulation has $20k^2$ triangles, whose projected vertices are approximately evenly spaced on $\mathbb{S}^2$. For each triangle of the k'th order triangulation we take its circum-center and project it onto $\mathbb{S}^2$. These projections are our choice of directions. With each direction we associate the area of the projection of the corresponding triangle onto $\mathbb{S}^2$. Let D_k be the set of directions that we get from the k'th order triangulation, and for any $d \in D_k$ let $w_k(d)$ be the area of the spherical triangle associated with d. For any point $v \in \mathbb{R}^3 \setminus C$ we approximate K_v^* by

$$\frac{1}{2} \sum_{d \in D_k} w_k(d)|1 - i_{d,v}|.$$

Note that by Lemma 12, this expression takes the same value at any point in the same cell of the hyperplane arrangement $\mathcal{H}_k$ induced by the directions in D_k. That is, we can get only $O(n^3k^6)$ different values for K_v^* (remember that $|D_k| = 20k^2$). It remains to bound the maximum error that we make when turning to the approximation above.

Lemma 13 *We have*

$$\left| K_v^* - \sum_{d \in D_k} w_k(d) i_{d,v} \right| \in O(n^2/k)$$

Proof. Given some point $v \in \mathbb{R}^3 \setminus C$. Let Z be a cell of the circle arrangement $\mathcal{C}_v$ induced by v on $\mathbb{S}^2$. By Lemma 11, the point v has the same index in any direction $d \in Z$. Hence, if the spherical triangle associated with $d \in D_k \cap Z$ is completely contained in Z, then we do not make any error on the term $w_k(d)|1 - i_{d,v}|$. We can only make errors for directions $d \in D_k$ whose associated spherical triangles are not completely contained in a cell of the arrangement $\mathcal{C}_v$, i.e., triangles that are intersected by at least one of the great circles in the arrangement $\mathcal{C}_v$. The error we can make for such a direction d is bounded by $(n-1)w_k(d)$ since we can have no more than $2n$ intersections of the polygon C and any hyperplane passing through v (except in a set of directions of measure zero). By construction of the k'th order triangulation we have $w_k(d) \in O(1/k^2)$. Next we want to bound the number of directions on which we can make an error. The spherical triangle associated with such a direction must intersect at least one of the n great circles in the arrangement $\mathcal{C}_v$. Any such circle has length 2π and can intersect at most $O(k)$ of the spherical triangles. Hence we can make an error of at most $(n-1)O(1/k^2)$ in at most $nO(k)$ directions. This implies the bound claimed in the statement of the lemma. □

From Lemma 13, an approximation scheme for the value of the optimal solution of the curvature minimization problem can be derived as follows: given an error bound $\varepsilon > 0$. Choose $k \in \Theta(n^2/\varepsilon)$. Compute the hyperplane arrangement $\mathcal{H}_k$ in $\mathbb{R}^3$ for the directions obtained from the k'th order triangulation of the icosahedron. Pick a point v in any cell of $\mathcal{H}_k$, which has $O\big((nk^2)^3\big) = O\big(n^{15}/\varepsilon^6\big)$ cells, and compute $\frac{1}{2}\sum_{d \in D_k} w_k(d)|1 - i_{d,v}|$. Output the smallest among the computed values. That value can be computed in time $O\big(n^{15}/\varepsilon^6\big)$ and is by Lemma 13 an ε-approximation of the optimal value. That is, we have the following theorem:

Theorem 4 *For any $\varepsilon > 0$, an ε-approximation of the optimal value for the curvature minimization problem can be computed in time $O\big(n^{15}/\varepsilon^6\big)$.*

5 Conclusion

We found the absolute Gaussian curvature to be an intriguing mathematical quantity which is practically relevant in mesh smoothing. So far our analysis does not lead to practical results but we believe that it is worthwhile to further pursue research on the absolute Gaussian curvature—not at least since it is a mathematically beautiful and challenging concept that relates to many areas in mathematics like topology and integral geometry.

References

[1] Lyuba Alboul and Ruud van Damme. Polyhedral metrics in surface reconstruction. In *IMA Conference on the Mathematics of Surfaces*, pages 171–200, 1994.

[2] Nina Amenta, Marshall W. Bern, and David Eppstein. Optimal point placement for mesh smoothing. *J. Algorithms*, 30(2):302–322, 1999.

[3] Chandrajit L. Bajaj. The algebraic degree of geometric optimization problems. *Discrete & Computational Geometry*, 3:177–191, 1988.

[4] Ulrich Brehm and Wolfgang Kühnel. Smooth approximation of polyhedral surfaces regarding curvatures. *Geometriae Dedicata*, 12:435–461, 1982.

[5] Maike Buchin and Joachim Giesen. Minimizing the total absolute gaussian curvature in a terrain is hard. In *CCCG*, pages 195–198, 2005.

[6] Frédéric Cazals and Marc Pouget. Estimating differential quantities using polynomial fitting of osculating jets. In *Computer Aided Geometric Design*, pages 121–146, 2005.

[7] Frédéric Chazal, David Cohen-Steiner, and André Lieutier. A sampling theory for compact sets in euclidean space. In *Symposium on Computational Geometry*, pages 319–326, 2006.

[8] David Cohen-Steiner and Jean-Marie Morvan. Restricted delaunay triangulations and normal cycle. In *Symposium on Computational Geometry*, pages 312–321, 2003.

[9] Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *SIGGRAPH*, pages 317–324, 1999.

[10] Tamal K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, 2006.

[11] Tamal K. Dey and Joachim Giesen. Detecting undersampling in surface reconstruction. In *Symposium on Computational Geometry*, pages 257–263, 2001.

[12] Tamal K. Dey and Samrat Goswami. Provable surface reconstruction from noisy samples. *Comput. Geom.*, 35(1-2):124–141, 2006.

[13] Lisl Gaal. *Classical Galois Theory: With Examples*. Markham, 1971.

[14] Meenakshisundaram Gopi. On sampling and reconstructing surfaces with boundaries. In *CCCG*, pages 49–53, 2002.

[15] I. N. Herstein. *Topics in Algebra, 2nd Edition*. Wiley, 1975.

[16] Boris Mederos, Nina Amenta, Luiz Velho, and Luiz Henrique de Figueiredo. Surface reconstruction for noisy point clouds. In *Symposium on Geometry Processing*, pages 53–62, 2005.

[17] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. Laplacian mesh optimization. In *GRAPHITE*, pages 381–389, 2006.

[18] Gabriel Taubin. A signal processing approach to fair surface design. In *SIGGRAPH*, pages 351–358, 1995.

Appendix

Proof of Lemma 4. We parameterize the line ℓ by $(-1,1) \ni t \mapsto x_t = (t,0,1) \in \ell$ and define

$$\begin{aligned}\left[\frac{\partial}{\partial t} K^*_\eta(x_t)\right]_{t=-1} &= \lim_{t\downarrow -1} \frac{\partial}{\partial t} K^*_{\eta_0}(x_t) \quad \text{and} \\ \left[\frac{\partial}{\partial t} K^*_\eta(x_t)\right]_{t=1} &= \lim_{t\uparrow 1} \frac{\partial}{\partial t} K^*_{\eta_0}(x_t).\end{aligned}$$

It is sufficient to prove that for sufficiently small $\eta > 0$,

$$\left[\frac{\partial}{\partial t} K^*_\eta(x_t)\right]_{t=-1} < 0 \quad \text{and} \quad \left[\frac{\partial}{\partial t} K^*_\eta(x_t)\right]_{t=-1} > 0.$$

since this implies the existence of neighborhoods of $v_2 = x_1$ and $v_5 = x_{-1}$, respectively, on ℓ in which K^*_η decreases if we move away from v_2 and v_5, respectively. To verify the latter inequalities we use the expression

$$K^*_\eta(x_t) = \alpha_7(x_t) + \alpha_8(x_t) - \alpha_6(x_t),$$

where $\alpha_7(x_t)$ and $\alpha_8(x_t)$ depend on η. After some calculations we get

$$\begin{aligned}
\frac{\partial}{\partial t} K^*_\eta(x_t)]_{t=-1} \quad = \quad & \frac{1}{3}\Big(3\sqrt{3\eta^2+2} + 12\eta\sqrt{3\eta^2+2} + 15\eta^2\sqrt{3\eta^2+2} \\
& +6\eta^3\sqrt{3\eta^2+2} + 3\sqrt{4\eta+3\eta^2+2} \\
& -\eta\sqrt{4\eta+3\eta^2+2} + \eta^2\sqrt{4\eta+3\eta^2+2} \\
& -3\sqrt{2}\sqrt{4\eta+3\eta^2+2}\sqrt{3\eta^2+2} \\
& -4\sqrt{2}\eta\sqrt{4\eta+3\eta^2+2}\sqrt{3\eta^2+2} \\
& -2\sqrt{2}\eta^2\sqrt{4\eta+3\eta^2+2}\sqrt{3\eta^2+2}\Big) \\
& \Big((3+4\eta+2\eta^2)\sqrt{4\eta+3\eta^2+2}\sqrt{3\eta^2+2)}\Big)^{-1}
\end{aligned}$$

and

$$\begin{aligned}
\frac{\partial}{\partial t} K^*_\eta(x_t)]_{t=1} \quad = \quad & \frac{1}{3}\Big(-3\sqrt{-4\eta+3\eta^2+2} - 9\eta\sqrt{-4\eta+3\eta^2+2} \\
& -7\eta^2\sqrt{-4\eta+3\eta^2+2} + 4\eta^3\sqrt{-4\eta+3\eta^2+2} \\
& -3\sqrt{8\eta+11\eta^2+2} + 6\eta\sqrt{8\eta+11\eta^2+2} \\
& -9\eta^2\sqrt{8\eta+11\eta^2+2} + 6\eta^3\sqrt{8\eta+11\eta^2+2} \\
& +3\sqrt{2}\sqrt{8\eta+11\eta^2+2}\sqrt{-4\eta+3\eta^2+2} \\
& +2\sqrt{2}\eta^2\sqrt{8\eta+11\eta^2+2}\sqrt{-4\eta+3\eta^2+2}\Big) \\
& \Big((3+2\eta^2)\sqrt{8\eta+11\eta^2+2}\sqrt{-4\eta+3\eta^2+2}\Big)^{-1}
\end{aligned}$$

It is easy to see that both expressions converge to 0 as η goes to 0. From an order analysis (looking at the lowest order terms in η) we get that the first term converges from the negative and the second term from the positive side to 0.

Proof of Lemma 5. We distinguish three cases. The first case deals with points v in the interior of the convex hull $\mathrm{conv}(C_0)$ of the polygon C_0. The second case deals with points v in the interior of the wedge above $\mathrm{conv}(C_0)$, and the third case deals with the case of v in the left- and right wedges above $\mathrm{conv}(C_0)$ (including their boundaries). See also Figure 7.

We start with the first case. Inside the convex hull of C_0, we have $K^*(v) = \sum_{i=1}^6 \alpha_i - 2\pi$. We first show that there exists a neighborhood U_ε of ℓ such that $|\nabla K^*(p)| > 0$ for all $v \in \mathrm{conv}(C_0) \cap U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$. We have $\nabla K^*(v) = \sum_{i=1}^6 \nabla\alpha_i(v)$. Recall that the direction of $\nabla\alpha_i(v)$ is towards the midpoint $c_i(v)$ of the circum-circle of the triangle spanned by v and e_i. The functions $c_i(v)$ are continuous in v. A simple calculation shows that the z coordinate of all $c_i(v) < 1/\sqrt{2}$ for all $v \in \ell \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$. Hence by uniform continuity of the gradients of the angle functions we have for sufficiently small $\varepsilon > 0$ that all the vectors from $c_i(v) - v$ for $v \in \mathrm{conv}(C_0) \cap \big(U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)\big)$ are contained in an open hemisphere. Thus, there also exist $\delta > 0$ such that

$$\|\nabla K^*(p)\| = \left\|\sum_{i=1}^6 \nabla\alpha_i\right\| > \delta$$

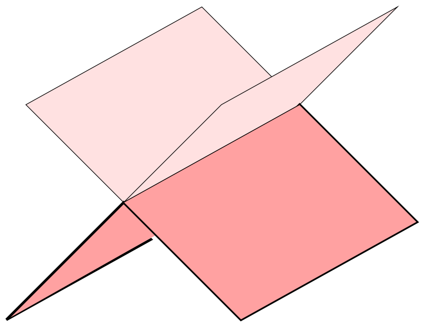

Figure 7: First case: region between the dark shaded facets. Second case: region between the light shaded facets, and third case: regions bounded by dark and light shaded facets.

for all these points v.

Now we handle the second case. Let W_2 be the region we are dealing with here. All points in $v \in W_2 \cap \big(U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)\big)$ see all the edges of $\mathrm{conv}(C_0)$ for all $\varepsilon > 0$. Hence the the absolute Gaussian curvature at these points is $K^*(v) = 2\pi - \sum_{i=1}^{6} \alpha_i(v)$. Now the claim follows by essentially the same argument as in the first case.

Finally, we handle the third case. This case actually has two sub-cases (left- and right wedge). Here we deal only with one of them since the arguments for the other sub-case are essentially the same. We deal with the case that the points v see the face F spanned by the vertices v_1, v_2, v_5 and v_6. At these points the absolute curvature is

$$K^*(v) = 2\pi - 2\alpha_5(v) - 2\alpha_6(v) - 2\alpha_1(v) - 2\alpha_\ell(v) + \sum_{i=1}^{6} \alpha_i(v).$$

Note that for v in the relative interior of the face F we have

$$\nabla\big(\alpha_5(v) + \alpha_6(v) + \alpha_1(v) + \alpha_\ell(v)\big) = 0.$$

Hence, $\nabla K^*(v) = \sum_{i=1}^{6} \nabla\alpha_i(v)$ at these points, and by similar arguments as used them for the first and second case the claim holds for $v \in F \cap \big(U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)\big)$. By the uniform continuity of the gradients of the angle functions we can extend this result to a small neighborhood of F in the wedge we are looking at, such that for small enough $\varepsilon > 0$ this restricted neighborhood contains $U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$ restricted to this wedge.

Proof of Lemma 6. We also do a case distinction for this proof, where we distinguish three cases. The first case is for points $v \in \mathrm{conv}(C_0)$, the second case is for $v \in \mathrm{conv}(C_\eta) \setminus \mathrm{conv}(C_0)$, and the third case is for points outside $\mathrm{conv}(C_\eta)$.

We start with the first case. We get for the difference

$$\begin{aligned} \nabla K^*(v) - \nabla K^*_\eta(v) &= \sum_{i=1}^{6} \nabla\alpha_i(v) - \sum_{i=1}^{5} \nabla\alpha_i(v) - \nabla\alpha_7(v) - \nabla\alpha_8(v) \\ &= \nabla\alpha_6(v) - \nabla\alpha_7(v) - \nabla\alpha_8(v). \end{aligned}$$

Note that $\nabla\alpha_7(v)$ and $\nabla\alpha_8(v)$ actually depend on η, but we find for the limit

$$\lim_{\eta\to 0} \big(\nabla\alpha_6(v) - \nabla\alpha_7(v) - \nabla\alpha_8(v)\big) = 0.$$

That is, we have pointwise convergence of the gradients. Now by Lemma 3(2) we can find for every $v \in U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$ a value $\eta_v > 0$ and small open neighborhood $U(v)$ such that for all points in $u \in U(v)$ we have $\nabla K^*(u) - \nabla K^*_\eta(u) < \zeta$. The neighborhoods $U(v)$ cover the compact set $U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$. Hence there exists a finite sub-cover. If we choose $\bar\eta$ as the minimum of η_v that correspond to the points that correspond to finite cover, then we have for all points $v \in \big(U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)\big)$ and all $\eta < \bar\eta$ that $\big\|\nabla K^*(v) - \nabla K^*_\eta(v)\big\| < \zeta$.

Now we turn to the second case. We get for the difference

$$\begin{aligned}\nabla K^*(v) - \nabla K^*_\eta(v) &= -2\nabla\big(\alpha_1(v) + \alpha_\ell(v) + \alpha_5(v) + \alpha_6(v)\big)\\ &\quad + \sum_{i=1}^{6} \nabla\alpha_i(v) - \sum_{i=1}^{5} \nabla\alpha_i(v) - \nabla\alpha_7(v) - \nabla\alpha_8(v)\\ &= -2\nabla\big(\alpha_1(v) + \alpha_\ell(v) + \alpha_5(v) + \alpha_6(v)\big)\\ &\quad + \nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big).\end{aligned}$$

Observe now that on the face F spanned by the vertices v_1, v_2, v_5 and v_6 we have

$$\nabla\big(\alpha_1(v) + \alpha_\ell(v) + \alpha_5(v) + \alpha_6(v)\big) = 0.$$

Also, we find for the limit

$$\lim_{\eta\to 0} \big(\nabla\alpha_6(v) - \nabla\alpha_7(v) - \nabla\alpha_8(v)\big) = 0.$$

That is, we have pointwise convergence of the gradients on F. Hence, for any point $v \in F \cap \big(U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)\big)$ we find a value $\eta_v > 0$ and small open neighborhood $U(v)$ in $\mathrm{conv}(C_\eta) \setminus \mathrm{conv}(C_0)$ such that for all points in $u \in U(v)$ we have $\nabla K^*(u) - \nabla K^*_\eta(u) < \zeta$. By a similar compactness argument as in the first case we can conclude that there exists $\bar\eta > 0$ and $\varepsilon > 0$ such that we have for all points $v \in \big(\mathrm{conv}(C_\eta) \setminus \mathrm{conv}(C_0)\big) \cap U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)$ and all $\eta < \bar\eta$ that $\big\|\nabla K^*(v) - \nabla K^*_\eta(v)\big\| < \zeta$.

Finally, we turn to the third case. For this case we distinguish three sub-cases (according to the left-, right and top wedges above $\mathrm{conv}(C_\eta)$, see also Figure 7 (note that $\eta = 0$).

We start with the left wedge case; that is the case at the points v do not see the vertex v_7. We get for the difference

$$\nabla K^*(v) - \nabla K^*_\eta(v) = \nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big).$$

The claim now follows by the same arguments as for the first case.

Next we deal with the top wedge case. For this case we need to introduce two new edges on the convex hull of $\mathrm{conv}(C_\eta)$, namely, l_1 that connects v_5 and v_7, and l_2 that connects v_2 and v_7, see also Figure 8. We also need the corresponding angle functions $\alpha_{l_1}(v)$ and $\alpha_{l_2}(v)$, and have to distinguish two cases:

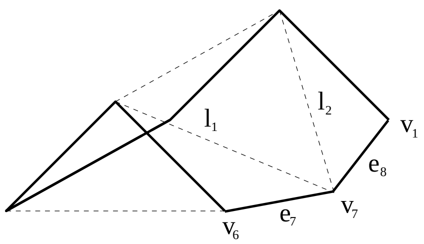

Figure 8: The convex hull of $\mathrm{conv}(C_\eta)$.

either v sees the vertices v_2, v_5 and v_7, or it sees the vertices v_1, v_2, v_5, v_6 and v_7. In the first case we get for the difference

$$\nabla K^*(v) - \nabla K^*_\eta(v) = 2\nabla\big(\alpha_{l_1}(v) + \alpha_{l_2}(v) - \alpha_5(v) - \alpha_6(v) - \alpha_1(v)\big) + \nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big).$$

Let F be the infinite face spanned by ℓ and the rays centered at v_2 and v_5, respectively, shooting in the direction $v_2 - v_1$ and $v_5 - v_6$, respectively. Observe that for $v \in F$,

$$\lim_{\eta\to 0} \nabla\big(\alpha_{l_1}(v) + \alpha_{l_2}(v) - \alpha_5(v) - \alpha_6(v) - \alpha_1(v)\big) = 0$$

and

$$\lim_{\eta\to 0} \nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big) = 0.$$

Now by similar arguments as before the claim also holds for this case. In the second case we get for the difference

$$\nabla K^*(v) - \nabla K^*_\eta(v) \quad = \quad -\nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big).$$

The claim in this case follows from the pointwise convergence $\lim_{\eta\to 0} \nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big)$ and a compactness argument as before.

Finally, we handle the right wedge case. Note that we are only interested in points $v \in \big(\mathbb{R}^3 \setminus \mathrm{conv}(C_0)\big) \cap \big(U_\varepsilon(\ell) \setminus \big(B_\rho(v_2) \cup B_\rho(v_5)\big)\big)$. For all these points v that also lie in the right wedge we need two distinguish two cases: either v sees the vertices v_2, v_5 and v_7, or it sees the vertices v_1, v_2, v_5, v_6 and v_7. In the first case we get for the difference

$$\nabla K^*(v) - \nabla K^*_\eta(v) = 2\nabla\big(\alpha_{l_1}(v) + \alpha_{l_2}(v) - \alpha_5(v) - \alpha_6(v) - \alpha_1(v)\big) + \nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big).$$

Let T be the triangle spanned by v_2, v_5 and v_7 observe that for $v \in T$ we have

$$\lim_{\eta\to 0} \nabla\big(\alpha_{l_1}(v) + \alpha_{l_2}(v) - \alpha_5(v) - \alpha_6(v) - \alpha_1(v)\big) = 0$$

and

$$\lim_{\eta\to 0} \nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big) = 0.$$

The claim follows for this case by arguments as above. In the second case we get for the difference

$$\nabla K^*(v) - \nabla K^*_\eta(v) \quad = \quad -\nabla\big(\alpha_6(v) - \alpha_7(v) - \alpha_8(v)\big).$$

The claim in this case follows immediately by similar arguments as before.

Proof of Lemma 10. For the proof we denote the dependency of $\phi_4(t)$ on m explicitly and write it as $\phi_4(t, m)$. We know that $\phi_4(t, m) = \sum_{i=1}^m a_i^m t^i$, where a_i^m are integer polynomials depending on m. By Lemma 9 the Galois group of $\phi_4(t, m)$ is S_{10} and $p_1 = 23, p_2 = 27$ and $p_3 = 137$ are good primes in this case. By Lemma 8

$$\phi_4(t, 10 + kp_1p_2p_3) \equiv \phi_4(t, 10) \pmod{p_i}$$

for $i = 1, 2, 3$ and $k \in \mathbb{Z}$. In particular, the $\phi_4(t, 10 + kp_1p_2p_3)$ have the same factorizations over $\mathbb{Z}_{p_i}$. Furthermore, the discriminant of these polynomials, that we denote by $\mathrm{disc}(10 + kp_1p_2p_3)$, is an integer polynomial in m and hence,

$$\mathrm{disc}(10 + kp_1p_2p_3) \equiv \mathrm{disc}(10) \pmod{p_i}.$$

Consequently, p_1, p_2 and p_3 are also good primes for $\phi_4(t, 10 + kp_1p_2p_3)$. Hence, $\mathrm{disc}(10 + kp_1p_2p_3)$ satisfies the sufficient conditions of Lemma 7 and thus the Galois group of $\phi_4(10 + kp_1p_2p_3)$ is S_{10} for all $k \in \mathbb{Z}$.

Proof of Lemma 11. Given two directions $d_1, d_2 \in \mathbb{S}^2$ in the same cell and let γ be the shortest path on $\mathbb{S}^2$ connecting d_1 and d_2. By construction, the great circle arc γ does not intersect any of the great circles $C_1, \ldots, C_n$ of the arrangement $\mathcal{C}_v$. The number of intersections of the polygon C with any hyperplane $H_{d,v}$ with normal $d \in \gamma$ and containing v can only change when $H_{d,v}$ passes through one of the vertices $v_1, \ldots, v_n$. The hyperplanes that pass through v and any of the vertices have directions exactly on the great circles. Hence the number of intersections $n_{d,v}$ of $H_{d,v}$ with C does not change for directions along γ and $i_{d_1,v} = i_{d_2,v}$, which proves the lemma.

Proof of Lemma 12. For direction d_i there are n parallel hyperplanes $H_{i1}, \ldots, H_{in}$ that partition $\mathbb{R}^3$ into $n+1$ cells. Each of these cells is a union of cells in $\mathcal{H}$. Let v and v' be two arbitrary points in one of the $n+1$ cells. Sweeping the hyperplane with normal d_i from v to v' does not cross any vertex v_j of the polygon C. Hence $n_{d_i,v} = n_{d_i,v'}$ and $i_{d_i,v} = i_{d_i,v'}$, which proves the lemma. Recall that the number of intersections of any such hyperplane with the polygon C can only change when $H_{d,v}$ passes through one of the vertices $v_1, \ldots, v_n$.