Cycle Bases of Graphs and Sampled Manifolds

Craig Gotsman [a] Kanela Kaligosi [b] Kurt Mehlhorn [b]
Dimitrios Michail [b] Evangelia Pyrga [b]

[a] *Technion-Israel Institute of Technology, Haifa, Israel*
[b] *Max-Planck-Institut für Informatik, Saarbrücken, Germany*

Abstract

Point samples of a surface in $\mathbb{R}^3$ are the dominant output of a multitude of 3D scanning devices. The usefulness of these devices rests on being able to extract properties of the surface from the sample. We show that, under certain sampling conditions, the minimum cycle basis of a nearest neighbor graph of the sample encodes topological information about the surface and yields bases for the trivial and non-trivial loops of the surface. We validate our results by experiments.

Key words: cycle basis, homology basis, manifold, nearest neighbor graph, non-trivial and trivial loops, sample points

1 Introduction

Point samples of a surface in $\mathbb{R}^3$ are the dominant output of a multitude of 3D scanning devices. The usefulness of these devices rests on being able to extract properties of the surface from the sample. Undoubtedly, the ultimate application is to form a geometric approximation of the surface based on the sample. The objective is typically to form a piecewise linear surface (i.e. triangle mesh) whose geometry and topology are as similar to the original as possible. For smooth surfaces and sufficiently dense samples, good and efficient reconstruction algorithms are available, see e.g., the papers [1–8] and

Email addresses: gotsman@cs.technion.ac.il (Craig Gotsman), kaligosi@mpi-inf.mpg.de (Kanela Kaligosi), mehlhorn@mpi-inf.mpg.de (Kurt Mehlhorn), michail@mpi-inf.mpg.de (Dimitrios Michail), pyrga@mpi-inf.mpg.de (Evangelia Pyrga).

the survey by Dey [9]. For non-smooth surfaces or sparse samples, the problem is open.

In some situations, e.g., when the above algorithms do not work satisfactorily, it is desirable to compute some topological or geometric information about the surface directly from the sample (e.g. the genus of the surface) *without* first performing an explicit reconstruction. Moreover, sometimes this information can facilitate the reconstruction process itself.

A recent paper by Tewari et al. [10] has shown how to construct a piecewise linear approximation (i.e. a 3D triangle mesh) of a scanned surface of genus one by first parameterizing the sample points. This parameterization relies on the computation of so-called *discrete harmonic one-forms*, which may be generated once the basic topological structure of the surface is determined. More specifically, this is equivalent to computing a basis of the linear subspace of all trivial loops on the surface. Beyond this application, it is sometimes necessary to compute a basis of the so-called *first homology group* of the surface (containing $2g$ loops for a surface of genus g) directly from the point cloud. Informally, the trivial loops are those which bound topological disks on the surface, otherwise they are non-trivial. This paper addresses these problems, and shows under which sampling conditions can correct results be obtained.

Our main result is as follows: Let S be a compact manifold in $\mathbb{R}^3$ and let P be a finite set of points in S. We call P a *point sample* of S. We show that if S is a smooth surface and P a sufficiently dense sample of S, the minimum cycle basis of the nearest neighbor graph of P gives information about the genus of S and yields bases for the set of trivial and non-trivial loops. Our experiments suggest that the algorithm also works for some non-smooth surfaces.

For an integer k, the k-nearest neighbor graph $G_k = G_k(P)$ on P is an undirected graph with vertex set P. It contains an edge between sample points a and b if b is one of the k points closest to a and a is one of the k points closest to b. The k-nearest neighbor graph is relatively easy to compute [11] and is a popular starting point for many algorithms extracting information about S from P. Some researchers define G_k in an unsymmetric way by requiring that either b is k-closest to a or a is k-closest to b. We do not know whether our results apply to the alternative definition.

A generalized cycle or simply cycle in an undirected graph $G = G(V, E)$ is a set C of edges with respect to which all nodes have even degrees. The sum of two cycles is their symmetric difference. With the addition operation defined in this manner, the set of cycles forms a vector space over the field of two elements, the *cycle space* of G. A *cycle basis* is a basis of this vector space. Every spanning tree gives rise to a basis. Each non-tree edge (i.e. chord) plus the tree path connecting its endpoints forms a cycle. For a connected graph

containing n vertices and m edges, every cycle basis contains exactly $m-(n-1)$ cycles. See [12] for more details. The *weight* of a cycle is the number of edges in the cycle and the weight of a cycle basis is the total weight of the cycles in it. A *minimum cycle basis (MCB)* is a cycle basis of minimum weight; in general, it is not induced by a spanning tree. MCBs can be computed efficiently: the currently best algorithm [13] runs in time $O(m^2n)$ using $O(m^2 + mn)$ space; the practical performance seems to be much better [14].

In this paper we show that for suitably nice samples of smooth manifolds of genus g and sufficiently large k, the k-nearest neighbor graph G_k has a cycle basis consisting only of *short* (= length at most $2(k+3)$) and *long* (= length at least $4(k+3)$) cycles. Moreover, the MCB is such a basis and contains exactly $m-(n-1)-2g$ short cycles and $2g$ long cycles. The short cycles span the subspace of trivial loops and the long cycles form a homology basis; see the next section for a definition. Thus, the MCB of G_k reveals the genus of S and also provides a basis for the set of trivial cycles and a set of generators for the non-trivial cycles of S. These cycles may then be used to parameterize P and ultimately generate a piecewise linear manifold surface approximating S, see Section 4.2.

We prove the statement for smooth manifolds, sufficiently dense samples, and sufficiently large k. A dense sample is defined in a manner similar to that of Amenta and Bern [3], as follows: The *medial axis* of a manifold S embedded in $\mathbb{R}^3$ is the closure of the set of points in $\mathbb{R}^3$ with more than one nearest neighbor on S. The *local feature size* $f : S \mapsto \mathbb{R}$ assigns to every point in S its least distance to the medial axis of S. The point sample P is called an ϵ-sample of S, if every point $x \in S$ has a point in P at distance at most $\epsilon f(x)$. If in addition $||p-q|| \geq \delta f(p)$ for all distinct $p, q \in P$ the sample P is called an (ϵ, δ)-sample. This definition requires the sample to be dense with respect to the local feature size, but at the same time samples cannot be arbitrarily close. Our main result is now as follows: If P is an (ϵ, δ)-sample of a smooth manifold for sufficiently small ϵ and k is sufficiently large (the meaning of sufficiently large depends on ϵ and δ), the claim of the preceding paragraph is true. The claim is also true if we use a weaker (more realistic in practice) sampling assumption, namely *locally uniform* ϵ-samples [5].

2 Related Work

Constructing a geometric approximation of a surface from a set of n samples of the surface has been the topic of many papers over the past decade (e.g. [1–3,5,4,6–8] and survey by Dey [9]). The objective is typically to form a piecewise linear surface (i.e. triangle mesh) whose geometry and topology are as similar to the original as possible. The algorithms above solve the problem for smooth

surfaces and sufficiently dense samples. The best asymptotic running time is $O(n \log n)$ [5] and efficient implementations are available [6,15]. Certainly, once this is done, it is relatively easy to extract less detailed information, such as the genus of the surface or a homology basis. For example, a homology basis of a manifold 3D mesh may be computed in $O(n)$ time by an algorithm of Vegter and Yap [16], and a *shortest* homology basis may be computed in $O(n^2 \log n + n^2 g + n g^3)$ by an algorithm of Erickson and Whittlesey [17].

Our method is an alternative way for extracting topological information. It provably works under the same conditions as the methods above and may work more often in practice. Also, it shows that the nearest neighbor graph suffices to deduce the topology and sometimes also the geometry, see Section 4.2. Interestingly enough, the latter requires knowledge of the basis of *trivial* loops, which are usually considered less interesting than the basis of non-trivial loops, and have not been addressed in prior work.

Another approach to reveal the topological structure of a sampled surface is to form an abstract simplicial complex on the point set (e.g. the Cech complex [18], see also [19]), and then compute the homology of this combinatorial object using *simplicial homology* [20]. This is defined in an manner analogous to singular homology and can be formulated in a linear algebraic setup. A theorem of Niyogi et al. [21] shows that for dense enough uniform samples, the homology of a Cech complex of a sampled manifold is isomorphic to that of the manifold, thus it suffices to compute these for the complex. More precisely, if P is $\epsilon/2$-dense in S, i.e., every point in S has a sample point within distance at most $\epsilon/2$, and $\epsilon \leq \sqrt{5/3}\tau$ where $\tau = \min_{p \in S} f(p)$ is the minimum distance of S from its medial axis, and the complex C contains a $j-1$-simplex, $1 \leq j \leq 4$, for every set $\{x_1, \ldots, x_j\}$ of sample points with non-empty common intersection of the ϵ-balls centered at the points, C has the same homology as the surface. In general, it is possible to compute the genus (and other Betti numbers) of a simplicial complex embedded in $\mathbb{R}^d$ as the co-rank of a Laplacian matrix derived from the complex, as defined by Friedman [22]. This will require at least $O(s^2)$ time, where s is the size of the complex.

For the special case of a geometric simplicial complex in $\mathbb{R}^3$, the genus can be computed in time $O(s)$ time [23,24]. However, the Cech complex defined above is not a geometric complex and hence the results do not apply.

The size s of a Cech complex is in $[\Omega(n), \Omega(n^4)]$. In practical situations, one can hope for linear size, but probably with a fairly large constant. We conjecture that our approach via the nearest neighbor graph is faster and also simpler. Furthermore, it facilitates the computation of a basis for the set of trivial cycles, which leads to a natural parameterization of the point cloud.

3 Some Theory

We prove that our approach works for smooth surfaces, sufficiently dense samples, and sufficiently large k.

3.1 The Basic Idea

Let S be a compact 2-manifold of genus g with no boundary embedded in $\mathbb{R}^3$. For the sequel we will need to define a number of topological concepts. Rather than give a complete formal exposition, which may be found in most algebraic topology textbooks (e.g. [20]), we provide a more intuitive informal set of definitions.

A closed simple curve on S is called a *loop*. The elements of the *first homology group* of S are equivalence classes of loops. The identity element of this homology group is the equivalence class of *separating loops*, namely, loops whose removal disconnects the surface. Two homology loops are in the same *homology class* if one can be continuously deformed into the other via a deformation that may include splitting loops at self-intersection points, merging intersecting pairs of loops, or adding and deleting separating loops. A loop is *trivial* if it is a separating loop and is *non-trivial* otherwise. A homology basis of S is any set of $2g$ loops whose homology classes span all non-trivial loops.

Let P be a sample of S and G_k its k-nearest neighbor graph. For any edge (a, b) of G_k we define a curve in S connecting a and b. Let us parameterize the segment ab by length. Let $p(t) = a + t(b - a)$ be the point with parameter value t, $0 \leq t \leq 1$ and let $q(t)$ be a point on S nearest to $p(t)$. *We assume that $q(t)$ is unique.* Then, $q(t)$ traces a continuous curve on S connecting a and b, which we denote by γ_{ab}. In this way, we obtain a drawing of G_k on S and cycles in G_k *induce* loops in S. In particular, cycles can induce trivial or non-trivial loops, in which case we call them trivial or non-trivial cycles, respectively. The assumption that $q(t)$ is unique is not very stringent. Observe that if $q(t)$ is not unique, $p(t)$ is a point on the medial axis and hence $\max(f(a), f(b)) \leq ||a - b||/2$.

We next give general conditions under which every minimum cycle basis of G_k contains exactly $2g$ long cycles, the long cycles induce a homology basis of S over the reals, and all short (= non-long) cycles are trivial.

(1) $L(\gamma_{ab}) \leq c_1||a-b||$ where c_1 is a constant, i.e., the length of the curve γ_{ab} is not much larger than the length of the segment ab.
(2) There is a subgraph M of G_k (all vertices and a subset of the edges) such that

- the drawing restricted to M is an embedding,
- M is a mesh of genus g for S,
- the faces of M have bounded size, say at most c_2 edges, and
- for every edge $e \in G_k \setminus M$, the shortest path in M connecting a and b has bounded length, say bounded by c_3.

(3) The minimal length of a non-trivial loop of S is L_{min} and for every edge $(a, b) \in G_k$ the distance between its endpoints a and b is at most $L_{min}/(2c_1 \max(c_2, c_3 + 1))$.

Theorem 1 *If the conditions above hold, every MCB of G_k contains exactly $m - (n-1) - 2g$ short (length less than $\max(c_2, c_3 + 1)$) and exactly $2g$ long (length at least $2\max(c_2, c_3 + 1)$) cycles. Moreover, the long cycles induce a basis of the first homology group of S over the reals.*

PROOF. The embedding of M has m_M edges and f faces where $f - m_M + n = 2 - 2g$. Consider the following set $\mathcal{B}$ of cycles: all face cycles of M but one and for each edge $e = (a, b)$ in $G_k \setminus M$ the cycle consisting of e plus the shortest path in M connecting a and b. Any cycle in $\mathcal{B}$ has length at most $\max(c_2, c_3+1)$ and the cycles in $\mathcal{B}$ are independent. There are $f-1+m-m_M = m_M - n + 2 - 2g - 1 + m - m_M = m - (n-1) - 2g$ cycles in $\mathcal{B}$.

Any cycle basis of G_k must contain at least $2g$ non-trivial cycles and these cycles induce loops which span the homology group of S. Consider any non-trivial cycle. It has length at least L_{min}. For any edge (a, b) of G_k, the length of γ_{ab} is at most $L_{min}/(2\max(c_2, c_3+1))$ and hence any non-trivial cycle must contain at least $2\max(c_2, c_3 + 1)$ edges.

We have now shown that there are $m - (n-1) - 2g$ independent cycles of length at most $\max(c_2, c_3 + 1)$ and that every non-trivial cycle consists of at least $2\max(c_2, c_3+1)$ edges. Consider now any MCB $\mathcal{B}^*$ of G_k. It must contain at least $2g$ long cycles. Assume that it contains less than $m-(n-1)-2g$ short cycles. Then, at least one cycle in $\mathcal{B}$, call it C, is not spanned by the short cycles in $\mathcal{B}^*$, i.e., the representation of C as a sum of cycles in $\mathcal{B}^*$ contains a cycle D which is not short. Thus we can improve the total length of B^* by replacing D by C.

In Sections 3.3 and 3.4 we substantiate the theorem for smooth curves. We will actually prove a stronger result. Note that in the previous theorem we did not use the first condition. This condition will be useful later on when we will also replace the global condition of maximal edge length (third condition above) by a local condition depending on the local feature size.

The local feature size is 1-Lipschitz continuous, i.e.,

Lemma 2 (Amenta and Bern [3]) *For any two points p and q on S, $|f(p)-f(q)| \leq ||p-q||$.*

Let D_P and V_P denote the Delaunay and the Voronoi diagram of P. The Voronoi cell $V_p \subset V_P$ for a point $p \in P$ is defined as the set of points $x \in \mathbb{R}^3$ such that $||x-p|| \leq ||q-x||$ for any $q \in P$ and $q \neq p$. The Delaunay triangulation of P is the dual of V_P. It has an edge pq if and only if V_p, V_q share a face, has a triangle pqr if and only if V_p, V_q, V_r share an edge, and a tetrahedron $pqrs$ if and only if V_p, V_q, V_r and V_s share a Voronoi vertex. We assume that the input sample $P \in \mathbb{R}^3$ is in general position and that no vertex of V_P lies on S.

Consider the restriction of the Voronoi diagram V_P to the surface S. This defines the *restricted Voronoi diagram* $V_{P|S}$, with *restricted Voronoi cells* $V_{p|S} = V_p \cap S$. It is said to satisfy the *ball* property if each $V_{p|S}$ is a topological 2-ball, each nonempty pairwise intersection $V_{p|S} \cap V_{q|S}$ is a topological 1-ball, and each nonempty triple intersection $V_{p|S} \cap V_{q|S} \cap V_{r|S}$ is a single point.

The dual of the restricted Voronoi diagram defines the *restricted Delaunay triangulation* $D_{P|S}$. In more detail, an edge pq is in $D_{P|S}$ if and only if $V_{p|S} \cap V_{q|S}$ is nonempty and a triangle pqr is in $D_{P|S}$ if and only if $V_{p|S} \cap V_{q|S} \cap V_{r|S}$ is nonempty. Our general position assumptions guarantee that there is no tetrahedron in $D_{P|S}$. It is known that the restricted Delaunay edges for an ϵ-sample are short.

Lemma 3 (Amenta and Bern [3], see also Giesen and Wagner [25]) *Let P be an ϵ-sample of S with $\epsilon < 1$. Then (1) The distance between $p \in P$ and its nearest neighbor in $P \setminus \{p\}$ is at most $\frac{2\epsilon}{1-\epsilon} f(p)$ and (2) if pq is an edge of the restricted Delaunay triangulation, $||p-q|| \leq \frac{2\epsilon}{1-\epsilon} min\{f(p), f(q)\}$.*

For $\epsilon \leq 0.08$, the restricted Voronoi diagram is known to have the ball property [3]. In this case the restricted Delaunay triangulation is a simplicial surface homeomorphic to S [26]. The k-neighborhood graph contains the restricted Delaunay triangulation if P is a sufficiently nice sample of S.

Lemma 4 (Andersson et al. [27]) *Let P be an (ϵ, δ)-sample of S, let $w = \frac{2\epsilon}{1-\epsilon}$ and let $k \geq \frac{(\delta(1+w)+w)^2}{\delta^2(1-w)^2-w^4}$. Then, the restricted Delaunay triangulation $D_{P|S}$ is a subgraph of G_k.*

Funke and Ramos [5] have shown how to turn any ϵ-sample into a *locally uniform* ϵ-sample by using decimation. Locally uniform ϵ-samples are related

to (ϵ, δ)-samples, but technically more involved. Locally uniform ϵ-samples are considered a more realistic assumption in practice.

We have chosen to state our results in terms of (ϵ, δ)-samples, for ease of exposition. We remark that the need for an (ϵ, δ)-sample instead of simply an ϵ-sample, is the property that each sample point has only a constant number of restricted Delaunay neighbors. Since locally uniform ϵ-samples also have this property, our results are valid in this case as well. We remark that locally uniform ϵ-samples are much more practical than (ϵ, δ)-samples.

We also state one more useful fact. At each point $p \in S$, there are two tangent *medial balls* centered at points of the medial axis. The vectors from p to the centers of its medial balls are normal to S, and S does not intersect the interiors of the medial balls. Since $f(p)$ is at most the radius of the smaller medial ball, S is also confined between the two tangent balls of radius $f(p)$.

3.3 Short Cycles

In this section we show that for an appropriate sampling density there exists a linearly independent set of $m - (n-1) - 2g$ short cycles.

3.3.1 Warm-Up: The Planar Case

Consider a finite set P of points in the plane, its nearest neighbor graph G_k, its Delaunay triangulation $D(P)$, and assume that $D(P)$ is contained in G_k. We describe a cycle basis consisting only of short cycles. In the next section, we will generalize the approach of this section to manifolds in $\mathbb{R}^3$.

Consider the following set $\mathcal{B}$ of cycles. It contains (a) the face cycles of all bounded faces of the Delaunay triangulation and (b) for each edge $e = (a, b) \in G_k \setminus D_P$ the cycle formed by e and the shortest path from a to b in D_P.

Lemma 5 *$\mathcal{B}$ is a cycle basis and any cycle in $\mathcal{B}$ has length at most $k+1$.*

PROOF. First we show that we have the right number of cycles. Let m_p be the number of edges of D_P and let m_k be the number of remaining edges. The dimension of the cycle space is $N = m - n + 1 = m_p + m_k - n + 1$. D_P is a planar graph and therefore $m_p - n + 1 = f - 1$ where f is the number of faces. Thus, $N = f - 1 + m_k = |\mathcal{B}|$.

The bounded faces of D_P are clearly linearly independent. The remaining cycles are also linearly independent since each of them contains an edge which

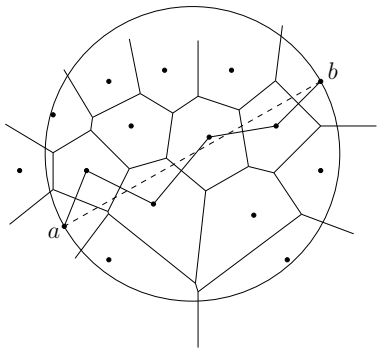

Fig. 1. An induced path from a to b formed with edges of the triangulation.

is not contained in any other cycle in $\mathcal{B}$. This proves the first part of our lemma.

We come to the second part. Bounded faces of D_P have length 3. Consider next an edge $e = (a, b) \in G_k \setminus D_P$. The straight line segment ab from a to b crosses some cells of the Voronoi diagram of P and induces a path $a = b_0, b_1, \ldots, b_{l-1}, b_l = b$ in D_P, namely the path through the sites owning these cells, see Figure 1. The path is entirely contained in the circle with ab as its diameter [28].

It remains to show that this cycle is short. Since $e = (a, b)$ is an edge of G_k, there can be at most $k - 1$ other points in this circle and hence any cycle in $\mathcal{B}$ has length at most $k + 1$.

This cycle basis is not necessarily the minimum cycle basis of the graph.

3.3.2 The 3-dimensional Case

We consider essentially the same set of cycles as in the preceding section: (a) all but one faces of the restricted Delaunay triangulation, and (b) for each remaining edge $e = (a, b) \in G_k \setminus D_{P|S}$ the cycle consisting of e plus the shortest path from a to b in the restricted Delaunay triangulation. As in the planar case, these cycles are linearly independent. It remains to prove that they are short. In this section we will prove the following theorem.

Theorem 6 *Let P be an ϵ-sample of S, let $k < \log_{\frac{1+\epsilon}{1-\epsilon}} \frac{3}{2}$, and assume that $D_{P|S} \subseteq G_k$. Let $(a, b) \in G_k \setminus D_{P|S}$. Then, there is a path from a to b in $D_{P|S}$ of length at most $2k + 5$.*

Consider an edge $e = (a, b) \in G_k \setminus D_{P|S}$. The curve γ_{ab} crosses some cells of the restricted Voronoi diagram and induces a path $p_{ab} = (a = b_0, b_1, \ldots, b_{l-1}, b_l = b)$ in the restricted Delaunay diagram, namely the path through the sites

owning the cells. Since G_k contains the restricted Delaunay triangulation, p_{ab} exists in G_k.

Lemma 7 *Let $e = (a, b) \in G_k \setminus D_{P|S}$ and let B be the ball with ab as its diameter. If the induced path p_{ab} is contained in B, then the cycle consisting of e plus the direct path from a to b on the triangulation has length at most $k + 1$.*

PROOF. Since $e \in G_k$, the ball B contains at most $k - 1$ other points apart from a and b.

The above lemma generalizes the planar case. Unfortunately, we are unable to prove that p_{ab} runs within B. We therefore argue somewhat differently. We first show that either there is a very short path in $D_{P|S}$ from a to b or both a and b are far away from the medial axis. In the latter case we show that p_{ab} is contained in a slightly bigger ball but still sufficiently small for our purposes.

Lemma 8 *Let a and b be two points in P and assume that there is a path of length l in $D_{P|S}$ from a to b. Let $\alpha = (1+\epsilon)/(1-\epsilon)$. Then, $\alpha - 1 = 2\epsilon/(1-\epsilon)$ and*

$$||a - b|| \leq (\alpha^l - 1) \min(f(a), f(b)).$$

PROOF. We use induction on the length of the path. Let the path from a to b in $D_{P|S}$ be $a = q_0, q_1, q_2, \ldots, q_l = b$. For the base case $l = 1$ we have $||a - q_1|| \leq (\alpha - 1)f(a)$ by Lemma 3. Assume that the statement holds for paths of length $l - 1$, then

$$\begin{aligned}
||a - b|| &= ||a - q_l|| \\
&\leq ||a - q_{l-1}|| + ||q_{l-1} - q_l|| \\
&\leq ||a - q_{l-1}|| + (\alpha - 1)f(q_{l-1}) && \text{by Lemma 3} \\
&\leq ||a - q_{l-1}|| + (\alpha - 1)(f(a) + ||a - q_{l-1}||) && \text{by Lemma 2} \\
&= (\alpha - 1)f(a) + \alpha||a - q_{l-1}|| \\
&\leq (\alpha - 1)f(a) + \alpha(\alpha^{l-1} - 1)f(a) && \text{by induction} \\
&= (\alpha^l - 1)f(a).
\end{aligned}$$

The same argument applies to b instead of a.

Lemma 9 *Let $(a, b) \in G_k$, let $\lambda \geq 1$ and assume*

$$k < \log_{\frac{1+\epsilon}{1-\epsilon}} \frac{1 + \lambda}{\lambda}. \tag{1}$$

Then, either there is a path of length at most k from a to b in the restricted Delaunay triangulation or

$$f(a), f(b) > \lambda ||a - b||. \tag{2}$$

PROOF. Recall that $\alpha = \frac{1+\epsilon}{1-\epsilon}$. Thus, inequality (1) implies $\alpha^k - 1 < 1/\lambda$. Let $a = q_0, q_1, q_2, \ldots, q_{l-1}, q_l = b$ be a shortest path from a to b in the restricted Delaunay triangulation. Furthermore, assume $l > k$ and $f(a) \leq \lambda ||a - b||$. Lemma 8 implies that

$$||a - q_i|| \leq (\alpha^i - 1) f(a) \leq (\alpha^k - 1) f(a) < f(a)/\lambda \leq ||a - b|| \quad \text{for } 1 \leq i \leq k,$$

and hence there are k points closer to a than b is, a contradiction. The argument works symmetrically for b.

The above lemma states that it is enough to prove Theorem 6 when $f(a), f(b) > \lambda ||a-b||$. From now on, we proceed under this assumption for some $\lambda \geq 1$. Let us parameterize the segment ab by length. Let $p(t) = a + t(b-a)$ be the point with parameter value t, $0 \leq t \leq 1$ and let $q(t)$ be the point on M nearest to $p(t)$. Note that $q(t)$ is unique, because otherwise $p(t)$ would be a point of the medial axis contradicting the fact that $f(a) > ||a - b||$. Finally, let c denote the mid-point of the segment ab and $s(t)$ denote the site of the Voronoi cell containing $q(t)$.

Our goal is to prove that $s(t)$ belongs to a ball of radius $\frac{\sqrt{3}}{2}||a - b||$ centered at c (Lemma 13) and that this ball contains at most $2(k + 3)$ sample points. We begin with the latter.

Lemma 10 *For $(a, b) \in G_k$ and c as defined above, the ball B' of radius $\frac{\sqrt{3}}{2}||a - b||$ centered at c contains at most $2(k + 3)$ sample points.*

PROOF. Consider the ball B_a with center a and radius $||a-b||$. Every sample point in the interior of this ball is closer to a than b is. Thus, B_a has at most k points in its interior. Also B_a has at most four points in its boundary by our non-degeneracy assumption. Similarly, the ball B_b with center b and radius $||a - b||$ also contains at most $k + 4$ points.

The ball B' is completely contained in the union of B_a and B_b and thus it contains at most $2(k + 4) - 2$ points. The -2 accounts for the fact that a and b are contained in both balls.

Next we estimate the distance from c to $s(t)$.

Lemma 11 $||c - s(t)|| \leq ||a - b||/2 + 2||p(t) - q(t)||$

PROOF. Assume w.l.o.g that $||a - p(t)|| \leq ||b - p(t)||$, otherwise we do the computation with b. By the triangle inequality, $||c-s(t)|| \leq ||c-p(t)||+||p(t)-q(t)|| + ||q(t) - s(t)||$. Since $q(t)$ is closer to $s(t)$ than any other sample point, $||q(t) - s(t)|| \leq ||q(t) - a||$. Moreover, $||q(t) - a|| \leq ||q(t) - p(t)|| + ||p(t) - a||$. Finally $||a - p(t)|| + ||p(t) - c|| = ||a - b||/2$.

It remains to bound $||p(t) - q(t)||$ as a function of $||a - b||$. We first estimate the distance of $q(t)$ from the medial axis.

Lemma 12 $f(q(t)) > (\lambda - 1)||a - b||$

PROOF. Assume w.l.o.g that $||a - p(t)|| \leq ||b - p(t)||$, otherwise we do the computation with b. Since $q(t)$ is the point in S closest to $p(t)$, we have $||a-q(t)|| \leq ||a-p(t)|| + ||p(t)-q(t)|| \leq 2||a-p(t)|| \leq ||a-b||$. By Lemma 2 $f(q(t)) \geq f(a)-||a-q(t)||$ and hence $f(q(t)) \geq f(a)-||a-b|| > (\lambda-1)||a-b||$.

Lemma 13 *For* $\lambda \geq 2$, $||p(t)-q(t)|| \leq \frac{\sqrt{3}-1}{4}||a-b||$ *and* $||c-s(t)|| \leq \frac{\sqrt{3}}{2}||a-b||$.

PROOF. Consider the point $q(t)$. By Lemma 12 there are two medial balls with radius at least $(\lambda-1)||a-b||$ tangent to $q(t)$. The surface passes between these balls and does not intersect their interior, in particular, a and b do not lie in the interior of these balls. Thus, the worst case (when $||p(t) - q(t)||$ compared to $||a - b||$ is maximized) occurs, when both lie on the boundary of one of these balls (see Figure 2). Let m be the center of this ball and use B_m to denote the ball. Consider the perpendicular bisector of segment ab passing through m. It intersects segment ab at c and ball B_m at v. Also, $p(t)$ is on the segment $q(t)m$.

Distance $||p(t) - q(t)||$ is upper bounded by $||c - v||$ and thus we are left with bounding $||c - v||$. Referring to Figure 2 we see that the triangle acm is right. Thus, $||c - m||^2 = ||a - m||^2 - ||a - c||^2$. Moreover,

$$||a - m|| = ||v - m|| = \zeta||a - b|| \text{ for some } \zeta \geq \lambda - 1,$$

and $||a - c|| = ||a - b||/2$. Combining all these,

$$||c - m|| = \sqrt{\zeta^2 - \frac{1}{4}} \cdot ||a - b||.$$

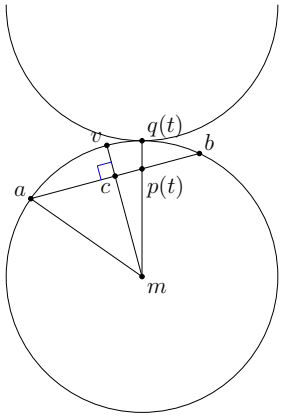

Fig. 2. Bounding $||p(t) - q(t)||$ in terms of $||a - b||$.

Finally, $||c - v|| = \zeta||a - b|| - ||c - m||$ and hence

$$||p(t) - q(t)|| \leq ||c - v|| = \left(\zeta - \sqrt{\zeta^2 - \frac{1}{4}}\right) ||a - b|| \\ \leq (1 - \sqrt{3/4})||a - b|| < \frac{\sqrt{3} - 1}{4}||a - b||.$$

This proves the first part of the lemma. The second part follows now from Lemma 11 since, $||c - s(t)|| \leq ||a - b||/2 + 2||p(t) - q(t)|| \leq \sqrt{3}||a - b||/2$.

It is now easy to complete the proof of Theorem 6. Set $\lambda = 2$. Then, p_{ab} is contained in the ball B' and B' contains at most $2(k+3)$ sample points. Thus, p_{ab} has length at most $2k+5$. Together with the edge e we get a cycle of length at most $2(k + 3)$.

Recall that our goal is to satisfy the assumptions of Section 3.1. If we combine Theorem 6 and Lemma 8 we get that the edges of G_k are small in length.

Corollary 14 *Let $\alpha = (1 + \epsilon)/(1 - \epsilon)$. For any edge $e = (a, b) \in G_k$,*

$$||a - b|| \leq (\alpha^{2k+5} - 1) \min\{f(a), f(b)\}. \tag{3}$$

Moreover, we will need the following lemma which can be easily derived from [25, Lemma 10 and Theorem 4].

Lemma 15 (Giesen and Wagner [25]) *Let a and b be two points of S such that $||a-b|| \leq \eta \cdot \min(f(a), f(b))$ with $\eta \leq 1/4$. Then, $L(\gamma_{ab}) \leq 4||a-b||$ where $L(\gamma_{ab})$ denotes the length of γ_{ab}.*

Using Lemma 15 we can set c_1 to 4 in Section 3.1.

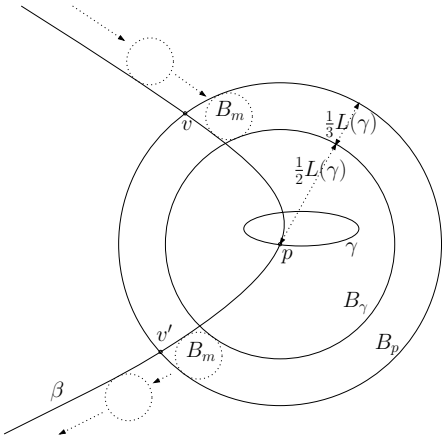

Fig. 3. Proving Theorem 16.

3.4 Long cycles

In this section we make precise how non-trivial cycles are long. The idea is simple; non-trivial cycles have a certain minimum length and edges of G_k are short. We will actually prove a stronger result. The length of a non-trivial cycle is bounded from below by the maximum feature size of any point on the loop. Combined with Lemma 15, we will obtain the desired result.

Assume that $\eta = \alpha^{2k+5} - 1 \leq \frac{1}{4}$. Let C be any non-trivial cycle of G_k. Substituting each edge $(a, b) \in C$ by the curve γ_{ab} gives us a non-trivial loop γ on S. By Lemma 15 and Corollary 14 we get

$$L(\gamma) \leq 4 \sum_{(a,b)\in C} ||a - b|| \leq 4 \sum_{(a,b)\in C} \eta \min(f(a), f(b)),$$

and if $|C|$ denotes the number of edges of C,

$$L(\gamma) \leq 4\eta|C| \max_{a\in C} f(a). \tag{4}$$

In order to get a lower bound on $|C|$ we need to relate γ to its distance to the medial axis. More precisely, we are going to show the following theorem which might be of independent interest.

Theorem 16 *Let γ be any non-trivial loop on S, then $L(\gamma) \geq \max_{p\in\gamma} f(p)$.*

PROOF. Let p be a point on γ with maximum distance from the medial axis and assume $L(\gamma) < f(p)$ for the sake of contradiction. Let β be a non-trivial loop of different homology class going through p (see Figure 3). At each point $x \in \beta$ there are two tangent balls with radius $f(x)$ which do not contain any point of S in their interior. One of these tangent balls when moving it along

β (and adjusting its size accordingly) produces an object $\mathcal{T}$, topologically equivalent to a torus, around which γ loops non-trivially.

Let B_p be the ball with center p and radius $5L(\gamma)/6$. For all $x \in B_p \cap \beta$, the local feature size is large, namely, $f(x) \geq f(p) - ||p-x|| > L(\gamma) - 5L(\gamma)/6 = L(\gamma)/6$ and hence for $x \in B_p \cap \beta$ the ball defining $\mathcal{T}$ has radius at least $L(\gamma)/6$.

The loop γ stays inside a ball of radius $L(\gamma)/2$ centered at p and hence well inside B_p. Since γ loops around $\mathcal{T}$ its length it least $2\pi L(\gamma)/6 > L(\gamma)$, a contradiction.

We can now establish that non-trivial cycles in G_k are long.

Theorem 17 *For appropriate values of ϵ, δ and k any non-trivial cycle $C \in G_k$ has length $|C| \geq \frac{1}{4\eta}$ where $\eta = \alpha^{2k+5} - 1 \leq \frac{1}{4}$.*

PROOF. Using inequality (4) and Theorem 16 we obtain

$$L(\gamma) \leq 4\eta |C| \max_{a \in C} \{f(a)\} \leq 4\eta |C| L(\gamma),$$

and the theorem follows.

Corollary 18 *If $\eta = \alpha^{2k+5} - 1 < \frac{1}{16(k+3)}$ then all non-trivial cycles in G_k have length larger than $4(k+3)$.*

PROOF. We fix k and ϵ to some constants according to (a) our assumptions in Section 3.3, and (b) such that $\eta < \frac{1}{16(k+3)}$. Then, by Theorem 17 we have $|C| \geq \frac{1}{4\eta} > 4(k+3)$.

Putting everything together establishes Theorem 1.

Corollary 19 *If the conditions in Section 3.1 hold, the sampling density is high enough and k is large enough: every MCB of G_k contains exactly $m - (n-1) - 2g$ short (length less than $2(k+3)$) and exactly $2g$ long (length at least $4(k+3)$) cycles.*

PROOF. Use $D_{P|S}$ as M in the assumptions of Section 3.1. By Lemma 15 we can set $c_1 = 4$. Theorem 6, Corollary 14 and Theorem 17 prove the remaining assumptions. The corollary follows by the proof of Theorem 1 for $c_2 = 3$ and $c_3 = 2k + 5$.

ϵ	valid k	chosen k	η	short-cycles upper bound	long-cycles lower bound
10^{-2}	$[15, 20]$	15	1.014	36	-
5×10^{-3}	$[14, 40]$	14	0.391	34	-
10^{-3}	$[14, 202]$	14	0.068	34	4
5×10^{-4}	$[14, 405]$	14	0.033	34	8
10^{-4}	$[14, 2027]$	14	0.0066	34	38
5×10^{-5}	$[14, 4054]$	14	0.0033	34	76

Table 1
Evaluation of the various conditions for the separation of the minimum cycle basis for different values of the sampling density ϵ.

3.5 Putting It All Together

Our assumptions so far suggest that given a (ϵ, δ)-sample and $w = \alpha - 1 = \frac{2\epsilon}{1-\epsilon}$, we should choose k such that:

$$\frac{(\delta(1+w)+w)^2}{\delta^2(1-w)^2 - w^4} \leq k < \log_{\frac{1+\epsilon}{1-\epsilon}} \frac{3}{2}. \tag{5}$$

There are values of ϵ such that inequality (5) cannot be satisfied. However, as ϵ decreases the right hand side increases while the left hand side decreases. Thus, both conditions can always be simultaneously satisfied for some dense enough sample. The above conditions are what is required for the trivial cycles of the MCB to have length at most $2(k+3)$.

For the lower bound on the length of the non-trivial cycles of the MCB, we also require that $\eta = \alpha^{2k+5} - 1 \leq \frac{1}{4}$, and in order for the length of the non-trivial cycles to reach the desired number $|C| \geq \frac{1}{4\eta} > 4(k+3)$ we require that $\eta = \alpha^{2k+5} - 1 < \frac{1}{16(k+3)}$.

We fix δ to $3w/8 \approx 3\epsilon/4$ and evaluate, in Table 1, the bounds for different values of ϵ.

Remark The lower and upper bounds presented in Table 1 are not tight. Somewhat large constants appear due to the proof technique used. Perhaps, by using some other proximity graph instead of the k-neighborhood or by performing a different analysis, the bounds can be improved. This is especially true for the value of ϵ required by the theory in order for the non-trivial cycles to be longer than the trivial ones. In practice this is true for smaller values of

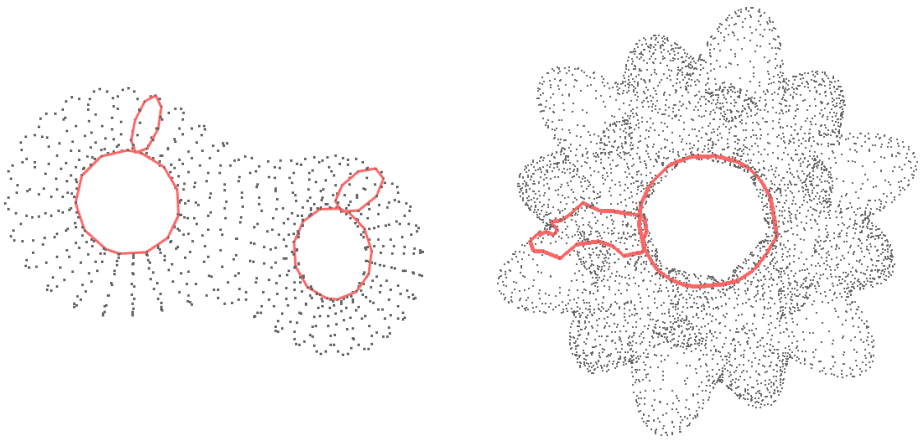

Fig. 4. Point clouds of (a) double torus model with 767 points and (b) "bumpy torus" model with 5044 points. The red cycles are the long non-trivial cycles extracted from the MCB of G_k for $k = 10$.

k and sampling densities. The next section presents experimental data which confirms this.

4 Experimental Validation

4.1 Genus Determination

This section presents experimental data on the size of trivial and non-trivial cycles in the MCBs of point clouds sampled from compact manifolds. The main observation is that the MCB cycles are separated into the two categories, short trivial and long non-trivial, for rather small values of k and sampling density. Moreover, the upper bound on the length of the trivial cycles is much less than $2(k + 3)$ and the method also works for some non-smooth samples.

We study three different examples: (a) a genus 2 double torus with a sparse point cloud (Figure 4), (b) a genus 1 surface with a dense point cloud (Figure 4), and (c) a genus 1 non-smooth surface (Figure 6).

Double torus: Since the model has genus 2 we expect the MCB to reveal exactly 4 non-trivial cycles. The minimum value of k for this event to happen is 6. All cycles but four of the MCB have length at most 8, two cycles have length 13 and two 24. As k increases, this gap grows, for $k = 8$ two cycles have length 11, two others have length 24 and all the rest at most 5. This continues to be true as we increase k as long as the edges of the k-neighborhood graph do not shortcut a handle.

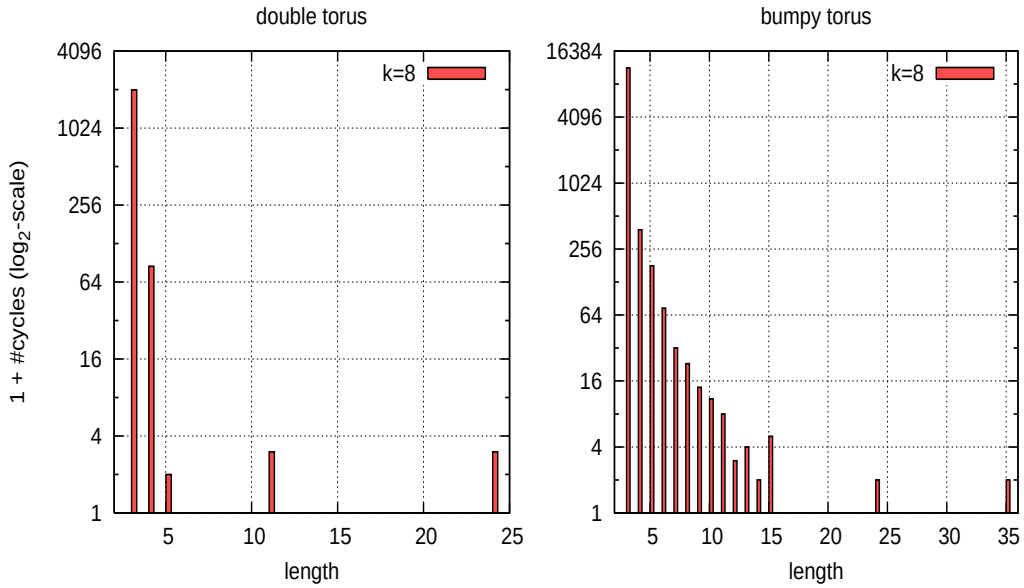

Fig. 5. Distribution of the lengths of the MCB cycles of G_k of the double and bumpy torus models for $k = 8$.

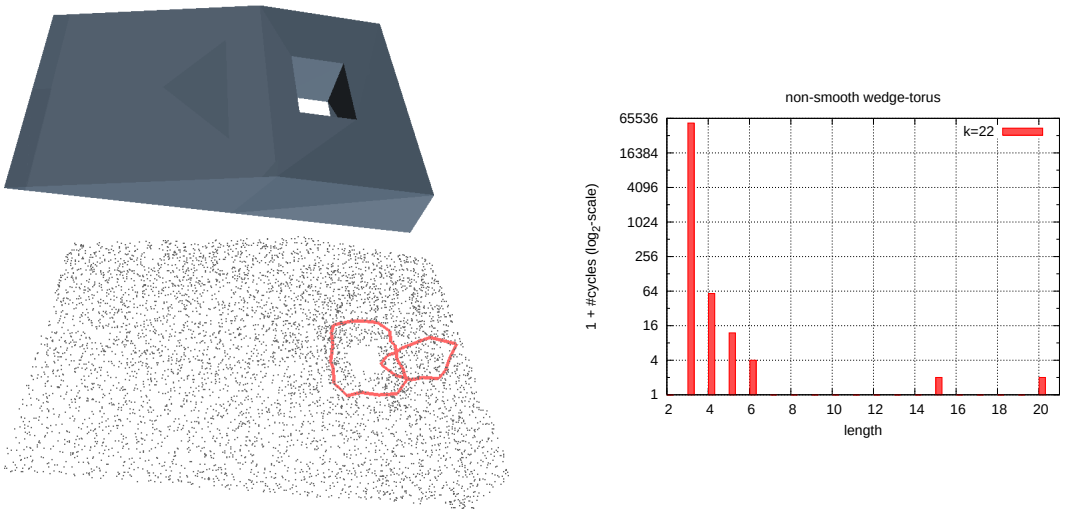

Fig. 6. Point cloud of the non-smooth model with 6078 points. On the left is the original surface and the point cloud with the 2 non-trivial cycles revealed by the MCB. On the right is the distribution of the lengths of the MCB cycles. All for $k = 22$.

Although the proof of Lemma 10 does not apply in the case of the *unsymmetric* k-neighborhood graph, where an edge is added even if only one endpoint is a k-nearest neighbor of the other, in practice we observe the same behavior. The values of k are even smaller for this case. For minimum $k = 5$ there are two cycles of length 11, two of length 24 and the rest at most 6. For $k = 9$ all MCB cycles are triangles except two with length 9 and two with length 14.

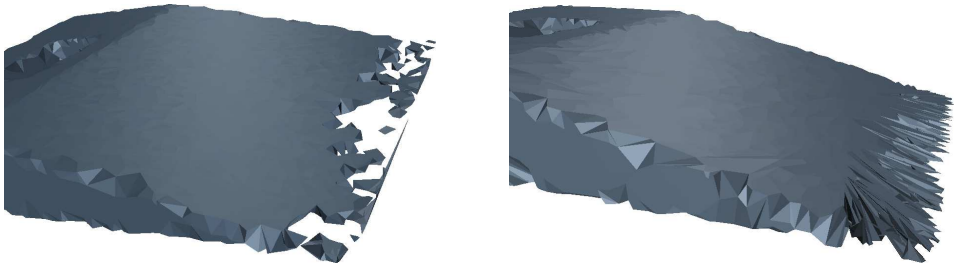

Fig. 7. Attempts to reconstruct the non-smooth model by Tight-Cocone [6] on the left and by the algorithm in [10] on the right.

Bumpy torus: The situation improves if the sampling is dense. The "bumpy torus" model has genus one and thus we expect the MCB to reveal two non-trivial cycles. For $k = 6$ the two non-trivial cycles have length 35 and 42 and the rest at most 27. Due to the density of the sample as we increase k this difference becomes more noticeable. For $k = 10$ the non-trivial cycles have length 22 and 30 and the rest at most 12. For $k = 12$ the two non-trivial have length 20 and 26 and the rest at most 9. Note also that in all these examples almost all trivial cycles have length 3 or 4. For example, when $k = 12$ about 99% of the trivial cycles are triangles. See Figure 5 for a histogram of the cycles length when $k = 8$.

Wedge torus: The non-smooth surface has genus one. We expect that as long as k is not too large, our method should reveal two non-trivial cycles. Figure 6 shows the two non-trivial cycles of the MCB for $k = 22$. Note that this is a difficult instance for surface reconstruction. Even Cocone [29] based implementations fail on this example, and the same is probably true for most Delaunay methods. See Figure 7.

Discussion: In the examples above, the MCB was able to reveal the genus. There were exactly $2g$ long cycles and long and short cycles are clearly discernible by length. However, in practice, if g is unknown, and the sampling density not as high as it should be, how can g be determined from the MCB? We suggest the following heuristic. Let l be the minimal integer such that the MCB contains no cycle of length between l and $2l$ inclusive. Then, $2g$ is the number of cycles of length larger than l. If this number is odd, this is an indication of insufficient sampling density or a wrong value of k.

What can go wrong when the sample is not sufficiently dense or the value of k is not chosen properly? When k is too small, the MCB might contain long trivial cycles. When k is too large, G_k may contain edges between points distant from each other in S and hence spurious long cycles may enter the basis, see Figure 8. The figure also shows that non-smoothness by itself is not an obstacle.

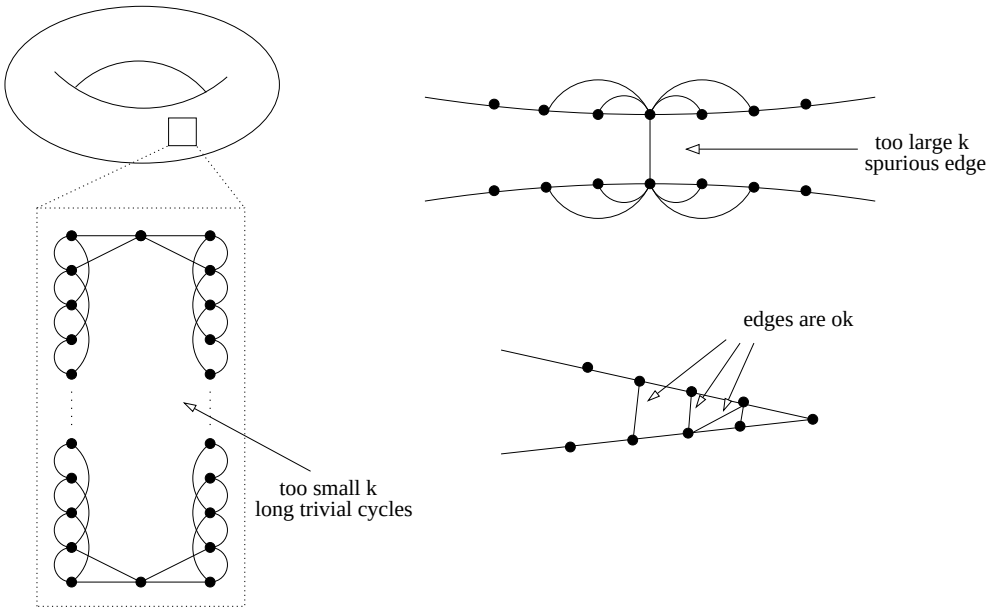

Fig. 8. Several difficult situations. In the left figure due to symmetry a small choice for k ($k = 4$) leads to long trivial cycles. In the upper right figure a large value of k results in an edge which connects two parts of the surface which are distant from each other. The lower right figure shows that edges near non-smooth features are not a problem as long as k is not too large.

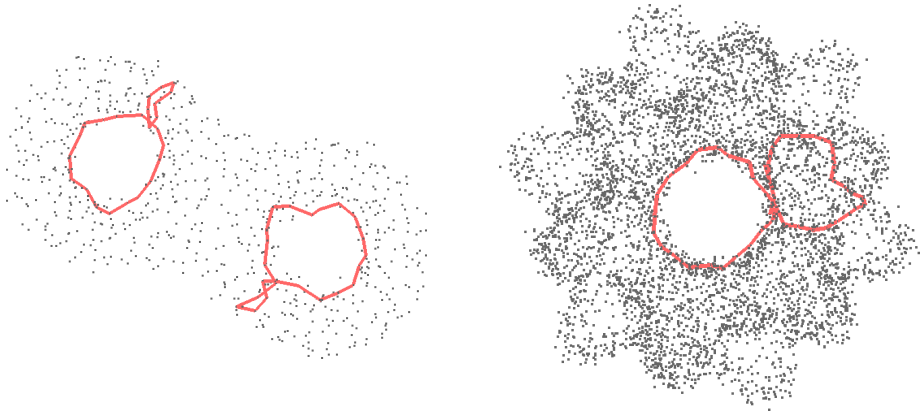

Fig. 9. Double and bumpy torus under Gaussian noise. In red color are the non-trivial cycles extracted from the MCB of G_k.

What about noisy samples? Figure 9 shows the double and bumpy torus under a small amount of Gaussian noise. The MCB was still able to distinguish trivial and non-trivial cycles. The main observation here is that a small amount of noise has little effect in instances where the non-trivial cycles are relative long w.r.t the sampling density. On the other hand we expect problems if the length of the non-trivial cycles is relatively close to the length of the trivial cycles.

Running time: The current fastest MCB algorithm requires $O(m^2n)$ time in unweighted graphs. Since the k-nearest neighbor graph has $O(kn)$ edges we

need $O(k^2n^3)$ time to compute an MCB. Due to this running time our method is restricted to small-medium instances.

This running time can be reduced if we use an approximate MCB. Note that a constant factor approximate MCB is bound to have similar properties as the MCB itself, at least for higher sampling densities. A $2q-1$ approximate MCB can be computed in time $O(qmn^{1+2/q} + mn^{(1+1/q)(\omega-1)})$ [30] where $\omega < 2.376$ [31] is the exponent of matrix multiplication. For denser graphs a $2q-1$ approximate MCB can be computed in time $O(n^{3+2/q})$.

We performed our experiments on a Pentium M 1.7GHz with 1GB of memory. For $k = 8$ the MCB computation on the double torus took 0.36 seconds while on the bumpy torus 12.01 seconds. The wedge torus required 1242 seconds since k was 22 and thus the graph was a lot denser. We also experimented with a 3-approximate MCB which in all cases managed to extract the non-trivial loops. For $k = 8$ a 3-approximate MCB on the double torus took 0.15 seconds and on the bumpy torus 2.66 seconds. For $k = 22$ on the wedge torus we spend 15.57 seconds to compute a 3-approximate MCB. We remark that even a 5-approximate MCB was able to separate properly the cycles on the wedge torus. The 5-approximate MCB computation required 5.16 seconds.

4.2 Application to Surface Reconstruction

In this section we outline the surface reconstruction algorithm of Tewari et al. [10] for surface reconstruction. The interested reader is referred to their paper for more details.

Tewari et al. [10] show that if a basis for the trivial loops of the manifold may be computed from the sample of a 2-manifold of genus 1, it is possible to parameterize the sample set, and then construct a piecewise-linear approximation to the surface. They use the MCB of the k-nearest neighbor graph to extract this basis, assuming that the non-trivial cycles are the two longest ones. They observed that this is correct if the sample is dense enough, but did not prove anything in this respect. Theorem 1 above shows under which conditions this reconstruction algorithm provably constructs a triangulation homeomorphic to the surface.

The parameterization based approach has its origins in Tutte's "spring embedder" for planar graphs [32]. Tutte introduced a simple, yet powerful method for producing straight-line drawings of planar graphs. The vertices of the outer face are mapped to the vertices of a convex polygon and all other vertices are placed at the centroid of their neighbors. Algorithmically, this amounts to solving a linear system of equations. If we use p_v to denote the location of

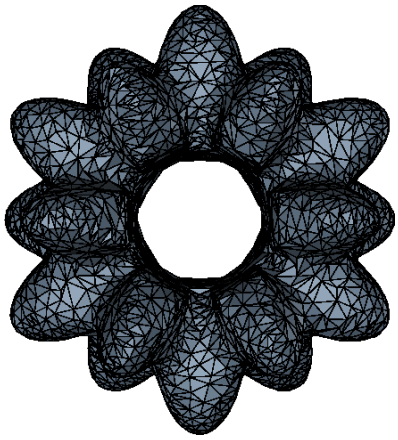

Fig. 10. Reconstruction of the bumpy torus model by the algorithm in [10].

vertex v and N_v to denote the set of neighbors of v, this means

$$p_v = \sum_{w \in N_v} \lambda_{vw} p_w \quad \text{and} \quad \lambda_{vw} = 1/|N_v| \tag{6}$$

for every interior vertex v. Tutte proved that the coordinates computed in this way define a non-degenerate embedding for any 3-connected planar graph. Floater [33] showed that the result stays true if vertices are placed at arbitrary convex combinations of their neighbors, i.e.,

$$\sum_{w \in N_v} \lambda_{vw} = 1 \quad \text{and} \quad \lambda_{vw} \geq 0.$$

For the sequel, it is convenient to rewrite Equation (6) as

$$\sum_{w \in N_v} \lambda_{vw}(p_w - p_v) = 0$$

and to introduce x_{vw} for the vector from v to w in the embedding. Gortler et al. [34] extended the method to embeddings onto the torus. Given a 3-connected map (= graph + cyclic ordering on the edges incident to any vertex) of genus one, they viewed undirected edges $\{v, w\}$ as pairs of directed edges and associated a variable x_{vw} with every directed edge. They used the equations:

$$\begin{aligned}
x_{vw} + x_{wv} &= 0 && \text{for all edges } (v, w) && \text{(symmetry)} \\
\sum_{w \in N_v} \lambda_{vw} x_{vw} &= 0 && \text{for all vertices } v && \text{(center of gravity)} \\
\sum_{(w,v) \in \delta f} x_{wv} &= 0 && \text{for all faces } f && \text{(face sums)}
\end{aligned}$$

The first class of equations ensures that the vector from w to v is the same as the vector from v to w, the second class ensures that v is a convex combination

of its neighbors, and the third class ensures that the vectors along any face boundary sum to zero. There are $2m$ unknowns and $m+n+f$ equations. Two equations are redundant (one center of gravity constraint and one face sum constraint) and hence the rank of the system is $m+n+f-2 = 2m-2$ (this uses the Euler theorem $f - m + n = 2 - 2g = 0$). Gortler et al. , extending results of Gu and Yau [35], proved that two independent solutions can be used as the x and y-coordinates of an embedding onto a torus.

Floater and Reimers [36] observed that Tutte's method can also be used to reconstruct surfaces with boundary of genus zero and Tewari et al. [10] extended the observation to closed surfaces of genus one, as follows. Construct the k-nearest neighbor graph G_k of P and then set up the equations introduced above. Face sum constraints are needed for a basis of the trivial cycles and this is exactly what an MCB yields. The solution of the system defines an embedding of P onto the torus. A triangulation, say the Delaunay triangulation, of the embedded point set is computed and then lifted back to the original point set. In this way a genus-1 surface interpolating P is obtained. The surface may have self-intersections. Postprocessing can be used to improve the quality of the mesh.

See Figure 7 and 10 for an attempt to reconstruct the non-smooth model and the reconstruction of the bumpy torus model by this algorithm. Reconstruction of non-smooth surfaces is an open problem.

5 Conclusions

In this work we have shown that given a suitably nice sample of a smooth manifold of genus g and sufficiently large k, the k-nearest neighbor graph of the sample has a cycle basis consisting only of *short* (= length at most $2(k+3)$) and *long* (= length at least $4(k+3)$) cycles. Moreover, the MCB is such a basis and contains exactly $m-(n-1)-2g$ short cycles and $2g$ long cycles. The short cycles span the subspace of trivial loops and the long cycles form a homology basis. Thus, the MCB reveals the genus of S and also provides a basis for the set of trivial cycles and a set of generators for the non-trivial cycles of S. These cycles may then be used to parameterize P and ultimately generate a piecewise linear manifold surface approximating S.

In our experiments we observe that the length separation of trivial and non-trivial cycles happens already for relatively sparse samples. In addition, this threshold is less than $2(k+3)$. Furthermore, our experiments suggest that the method also works for some non-smooth surfaces.

References

[1] J.-D. Boissonnat, F. Cazals, Smooth surface reconstruction via natural neighbour interpolation of distance functions, in: Proc. 16th Annual ACM Sympos. Comput. Geom., 2000, pp. 223–232.

[2] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, W. Stuetzle, Surface reconstruction from unorganized points, in: Proceedings of SIGGRAPH, 1992, pp. 71–78.

[3] N. Amenta, M. Bern, Surface reconstruction by voronoi filtering, in: SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry, ACM Press, New York, NY, USA, 1998, pp. 39–48.

[4] N. Amenta, S. Choi, R. Kolluri, The power crust, in: Proceedings of 6th ACM Symposium on Solid Modeling, 2001, pp. 249–260.

[5] S. Funke, E. A. Ramos, Smooth-surface reconstruction in near-linear time., in: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms, 2002, pp. 781–790.

[6] T. K. Dey, S. Goswami, Tight cocone: A water-tight surface reconstructor, Journal of Computing and Information Science in Engineering 3 (2003) 302–307.

[7] Y. Ohtake, A. Belyaev, M. Alexa, G. Turk, H.-P. Seidel, Multi-level partition of unity implicits, ACM Transactions on Graphics 22 (3) (2003) 463–470.

[8] R. Kolluri, J. Shewchuk, J. O'Brien, Spectral surface reconstruction from noisy point clouds, in: Proc. Symposium on Geometry Processing, 2004, pp. 11–21.

[9] T. K. Dey, Curve and surface reconstruction, in: J. E. Goodman, J. O'Rourke (Eds.), Handbook of Discrete and Computational Geometry, CRC Press, 2004.

[10] G. Tewari, C. Gotsman, S. J. Gortler, Meshing genus-1 point clouds using discrete one-forms, Computer and Graphics 30 (6).

[11] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, A. Y. Wu, An optimal algorithm for approximate nearest neighbor searching fixed dimensions, J. ACM 45 (6) (1998) 891–923.

[12] B. Bollobas, Modern Graph Theory, Graduate Texts on Mathematics, Springer, 1998.

[13] T. Kavitha, K. Mehlhorn, D. Michail, K. E. Paluch, A faster algorithm for minimum cycle basis of graphs., in: 31st International Colloquium on Automata, Languages and Programming, Finland, 2004, pp. 846–857.

[14] K. Mehlhorn, D. Michail, Implementing minimum cycle basis algorithms, in: 4th International Workshop on Efficient and Experimental Algorithms (WEA'05), Vol. 3503 of Lecture Notes in Computer Science, Santorini, Greece, 2005, pp. 32–43.

[15] T. K. Dey, J. Giesen, J. Hudson, Delaunay based shape reconstruction from large data, in: Proc. IEEE Symposium in Parallel and Large Data Visualization and Graphics, 2001, pp. 19–27.

[16] G. Vegter, C. Yap, Computational complexity of combinatorial surfaces, in: Sixth Annual Symposium on Computational Geometry, 1990, pp. 102–111.

[17] J. Erickson, K. Whittlesey, Greedy optimal homotopy and homology generators., in: Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, 2005, pp. 1038–1046.

[18] T. K. Dey, H. Edelsbrunner, S. Guha, Computational topology, in: B. Chazelle, J. E. Goodman, R. Pollack (Eds.), Advances in Discrete and Computational Geometry, Vol. 223 of Contemporary Mathematics, AMS, 1999, pp. 109–143.

[19] G. Carlsson, V. de Silva, Topological approximation by small simplicial complexes, manuscript (2003).

[20] J. Munkres, Elements of Algebraic Topology, Perseus, 1984.

[21] P. Niyogi, S. Smale, S. Weinberger, Finding the homology of submanifolds with high confidence from random samples, Tech. Rep. TR-2004-08, Computer Science, Univ. Chicago (2004).

[22] J. Friedman, Computing Betti numbers via combinatorial Laplacians, Algorithmica 21 (1998) 331–346.

[23] C. Delfinado, H. Edelsbrunner, An incremental algorithm for Betti numbers of simplicial complexes, in: Proceedings of the Ninth Annual Symposium on Computational Geometry, 1993, pp. 232–239.

[24] T. K. Dey, S. Guha, Computing homology groups of simplicial complexes in R^3, Journal of the ACM 45 (2) (1998) 266–287.

[25] J. Giesen, U. Wagner, Shape dimension and intrinsic metric from samples of manifolds with high co-dimension, in: SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry, ACM Press, New York, NY, USA, 2003, pp. 329–337.

[26] H. Edelsbrunner, N. Shah, Triangulating topological spaces, in: Proc. 10th ACM Symposium on Computational Geometry, 1994, pp. 285–292.

[27] M. Andersson, J. Giesen, M. Pauly, B. Speckmann, Bounds on the k-neighborhood for locally uniformly sampled surfaces, in: Symposium on Point-Based Graphics, 2004.

[28] D. P. Dobkin, S. J. Friedman, K. J. Supowit, Delaunay graphs are almost as good as complete graphs, Discrete Comput. Geom. 5 (4) (1990) 399–407.

[29] N. Amenta, S. Choi, T. K. Dey, N. Leekha, A simple algorithm for homeomorphic surface reconstruction, Int. J. Comput. Geometry Appl. 12 (1-2) (2002) 125–141.

[30] T. Kavitha, K. Mehlhorn, D. Michail, New approximation algorithms for minimum cycle bases of graphs, manuscript (June 2006).

[31] D. Coppersmith, S. Winograd, Matrix multiplications via arithmetic progressions., Journal of Symb. Comput. 9 (1990) 251–280.

[32] W. T. Tutte, How to draw a graph, Proceedings of the London Mathematical Society 13 (3) (1963) 743–768.

[33] M. S. Floater, Parametrization and smooth approximation of surface triangulations, Computer Aided Geometric Design 14 (3) (1997) 231–250.
URL http://dx.doi.org/10.1016/S0167-8396(96)00031-3

[34] S. Gortler, C. Gotsman, D. Thurston, One-forms on meshes and applications to 3D mesh parameterization, Computer Aided Geometric Design.

[35] X. Gu, S.-T. Yau, Computing conformal structures of surfaces, Communications in Information and Systems 2 (2) (2002) 121–146.

[36] M. S. Floater, M. Reimers, Meshless parameterization and surface reconstruction, Computer Aided Geometric Design 18 (2) (2001) 77–92.
URL http://dx.doi.org/10.1016/S0167-8396(01)00013-9