

On the Locality of Extracting a 2-Manifold in \mathbb{R}^3

Daniel Dumitriu* Stefan Funke† Martin Kutz* Nikola Milosavljevic‡

Abstract

Algorithms for reconstructing a 2-manifold from a point sample in \mathbb{R}^3 based on Voronoi-filtering like *CRUST* [1] or *CoCone* [2] still require – after identifying a set of candidate triangles – a so-called *manifold extraction* step which identifies a subset of the candidate triangles to form the final reconstruction surface. Non-locality of the latter step is caused by so-called *slivers* – configurations of 4 almost cocircular points having an empty circumsphere with center close to the manifold surface.

We show that under a certain mild condition – local uniformity – which typically holds in practice but can also be enforced theoretically, one can compute a reconstruction using an algorithm whose decisions about the adjacencies of a point only depend on nearby points. While the theoretical proof requires an extremely high sampling density, our prototype implementation – which is described in a companion paper [4] – exhibits pretty good results on typical sample sets and might have some potential in particular in parallel computing or external memory scenarios due to its local mode of computation. The full version of this paper is available under [5].

1 Introduction

Reconstructing a surface Γ in \mathbb{R}^3 from a finite point sample V has attracted a lot of attention both in the computer graphics community as well as in the computational geometry community. While in the former the emphasis is mostly on algorithms that work ‘well in practice’, the latter has focused on algorithms that come with a theoretical guarantee: if the point sample V




Figure 1: Output of the first 5 stages of our algorithm for the Stanford Bunny. Due to the conservative adjacency creation, some faces (light) are non-triangular.

satisfies a certain *sampling condition*, the output of the respective algorithm is guaranteed to be ‘close’ to the original surface.

In [1], Amenta and Bern proposed a framework for rigorously analyzing algorithms reconstructing smooth closed surfaces. They define for every point $p \in \Gamma$ on the surface the *local feature size* $\text{lfs}(p)$ as the distance of p to the *medial axis*¹ of Γ . A set of points $V \subset \Gamma$ is called a ε -sample of Γ if $\forall p \in \Gamma \exists s \in V : |sp| \leq \varepsilon \cdot \text{lfs}(p)$. For sufficiently small ε , Amenta and Bern define a *canonical correct reconstruction* of V with respect to Γ as the set of Delaunay triangles that are dual to Voronoi edges in the Voronoi diagram of V that are intersected by the surface Γ . Unfortunately, due to certain point configurations called *slivers* – 4 (almost) cocircular points that are nearby on the surface and have an empty, (almost) diametral circumsphere – it is not possible to algorithmically determine the canonical correct reconstruction of V with-

¹The medial axis of Γ is defined as the set of points which have at least 2 closest points on Γ .

*Max-Planck-Institut für Informatik, Saarbrücken, Germany, dimitriu@mpi-inf.mpg.de

†Ernst-Moritz-Arndt-Universität, Greifswald, Germany, stefan.funke@uni-greifswald.de

‡Stanford University, Stanford, CA, U.S.A., nikolam@stanford.edu

out knowing Γ . Algorithms have been proposed, though, that determine a collection of Delaunay triangles which form a piecewise linear surface that is topologically equivalent to the canonical correct reconstruction and converges to the latter both point-wise as well as in terms of the surface normals as the sampling density goes to infinity ($\varepsilon \rightarrow 0$).

The CoCone algorithm [2] is one example; in its last step, it first removes triangles with free edges and then determines the final reconstruction as the outside surface of the largest connected component of the remaining triangles. Observe that this is a highly non-local operation. There have been attempts to locally decide for each sample p which of the candidate triangles to keep for the final reconstruction; such local decisions might disagree, though, and hence the selected triangles do not patch up to a closed manifold. Again, the reason why local decisions might disagree is the presence of *slivers* which induce a Voronoi vertex inside the CoCone region of the involved sample points. Each involved sample point has to decide whether in ‘its opinion’ the true surface Γ intersects above or below the Voronoi vertex and create the respective dual Delaunay triangles. If these decisions are not coordinated contradicting decisions are made. Not only in theory but even in practice the manifold extraction step is still quite challenging and requires deliberate engineering to actually work as desired.

One potential way to obtain a local manifold extraction step is to decide on triangles/adjacencies in a conservative manner by only creating those triangles/adjacencies which are ‘safe’, i.e., where essentially the respective dual Voronoi edge/face completely pierces the CoCone region. It is unclear, though, how much connectivity is lost — whether the resulting graph is connected at all and how big potential holes/faces are. The main contribution of this paper is to show that it is actually possible to make local decisions but still guarantee that the resulting graph exhibits topological equivalence to the original surface. That is, it is connected, locally planar, and contains no large holes.

In [6] Funke and Milosavljevic present an algorithm for computing *virtual coordinates* for the nodes of a wireless sensor network which are themselves unaware of their location. Their approach crucially depends on a subroutine to identify a provably planar subgraph of a communication graph that is a quasi-Unit-Disk graph.

The same subroutine will also be used in our surface reconstruction algorithm presented in this paper.

While we deal with the problem of slivers in some sense by avoiding or ignoring them, another approach coined *sliver pumping* has been proposed by Cheng et al. in [3]. Their approach works for smooth k -manifolds in arbitrary dimension, though its practicality seems uncertain. There are, of course, other non-Voronoi-filtering-based algorithms for manifold reconstruction which do not have a manifold extraction step; they are not in the focus of this paper, though.

Our Contribution

We propose a novel method for extracting a 2-manifold from a point sample in \mathbb{R}^3 . Our approach fundamentally differs from previous approaches in two respects: first it mainly operates combinatorially on a graph structure, which is derived from the original geometry; secondly, the created adjacencies/edges are “conservative” in a sense that two samples are only connected if there is a safe, sliver-free region around the two samples. Interestingly we can show, though, that conservative edge creation only leads to small, constant-size faces in the respective reconstruction, hence completion to a triangulated piecewise linear surface can easily be accomplished using known techniques. The most notable advantage compared to previous Voronoi-filtering based approaches is that the manifold extraction step can be performed locally, i.e., it only relies on adjacency information of geometrically nearby points.

While the theoretical analysis requires an absurdly high sampling density – like most of the above mentioned algorithms do – our prototype implementation of the novel local manifold extraction step (see companion paper [4]) suggests that the approach is viable even for practical use.

From a technical point of view, two insights are novel in this paper (and not a result of the mere combination of previous results): first, we show that the neighborhood graph that our algorithms constructs is locally a *quasi-unit-disk graph*; it is this property that allows us to actually make use of the machinery developed in [6]. Second, we provide a more elegant and much stronger result about the density of the extracted planar graph based on the β -skeleton

and power-spanner properties; this insight also improves the overall result in [6].

2 Graph-based, conservative Adjacencies

In this Section we present an algorithm that given a ε -sample V from a closed smooth 2-manifold Γ in \mathbb{R}^3 computes a faithful reconstruction of V with respect to Γ as a subcomplex of the Delaunay tetrahedralization of V . The outline of our method is as follows (with the novel steps being 2.–5.):

1. Determine a Lipschitz function $\phi(v)$ for every $v \in V$ which lower-bounds $\varepsilon \text{fs}(v)$ (as in [7])
2. Construct a local neighborhood graph $G(V)$ by creating an edge from every point v to all other points v' with $|vv'| \leq O(\phi(v))$.
3. Compute a subsample S of (V)
4. Identify adjacencies between elements in S based on the connectivity of $G(V)$ (as in [6])
5. Use geometric positions of the points in S to identify faces of the graph induced by certified adjacencies when embedded on the manifold
6. Triangulate all non-triangular faces
7. reinsert points in $V - S$ by computing the weighted Delaunay triangulations on the respective faces (as in [7])

The core components of the correctness proof of this approach are:

- We show that the local neighborhood graph corresponds locally to a quasi-unit-disk graph for a set of points in the plane.
- The identified adjacencies locally form a planar graph.
- This locally planar graph has faces of bounded size.

Essentially this means that we cover Γ by a mesh with vertex set S consisting of small enough cells that the topology of Γ is faithfully captured.

We first discuss the 2-dimensional case, where we are given a uniform ε -sampling (i.e. the local feature size is 1 everywhere) of a disk and

show that steps 2. to 5. yield a planar graph with ‘small’ faces. Then we show how the same reasoning can be applied to the 3-dimensional case.

2.1 Conservative adjacencies in \mathbb{R}^2

Let V be a set of n points that form a ε -sampling of the disk of radius R around the origin o , that is, $\forall p \in \mathbb{R}^2$ with $|po| \leq R$, $\exists v \in V : |vp| \leq \varepsilon$.

Definition 2.1 *A graph $G(V, E)$ on V is called a α -quasi-unit-disk-graph (α -qUDG) for $\alpha \in [0, 1]$ if for $p, q \in V$*

- if $|pq| \leq \alpha$ then $(p, q) \in E$
- if $|pq| > 1$ then $(p, q) \notin E$

That is, for distances between α and 1 the presence of an edge is left open.

Within G we consider the distance function d_G defined by the (unweighted) graph distances in $G(V, E)$. Let $k \geq 1$, we call a set $S \subseteq V$ a *tight k -subsample* of V if

- $\forall s_1, s_2 \in S: d_G(s_1, s_2) > k$
- $\forall v \in V: \exists s \in S$ with $d_G(v, s) \leq k$.

A tight k -subsample of V can easily be obtained by a greedy algorithm which iteratively selects a so far unremoved node v into S and removes all nodes at distance at most k from consideration.

2.1.1 Graph-based conservative Adjacencies

The idea for construction and the planarity property of our construction are largely derived from the geometric intuition. The planarity follows from the fact that our constructed graph – we call it *combinatorial Delaunay map* of S , short $\text{CDM}(S)$ – is the *dual graph* of a suitably defined partition of the plane into simply connected disjoint regions. In the following we use the method for identifying adjacencies between nodes in S purely based

First we introduce a *labeling* of $G(V, E)$ for a given set $S \subseteq V$ assuming that all elements in V (and hence in S) have unique IDs that are totally ordered.

Definition 2.2 *Consider a vertex $a \in S$ and a vertex $v \in V - S$. We say that v is an a -vertex (or: labeled with a) if a is one of the elements in S which is closest to v (in graph distance), and a has the smallest ID among such.*

Clearly, this rule assigns unique labels to each vertex and edge, due to the uniqueness of nodes' IDs. Also note that any $a \in S$ is an a -vertex. Next we present a criterion for creating adjacencies between vertices in S .

Definition 2.3 $a, b \in S$ are adjacent in $CDM(S)$ iff there exists a path from a to b whose 1-hop neighborhood (including the path itself) consists only of a and b vertices, and such that in the ordering of the nodes on the path (starting with a and ending with b) all a -nodes precede all b -nodes.

We have the following result of [6].

Theorem 2.1 If G is an α - q UDG with $\alpha \geq \frac{1}{\sqrt{2}}$ and S a tight k -subsample of G , then $CDM(S)$ is a planar graph.

Of course, just planarity as such is not too hard to guarantee – one could simply return a graph with no edges.

2.1.2 CDM(S) is dense(!)

Interestingly we can show that in spite of the conservative adjacency creation, $CDM(S)$ is relatively dense, in particular it exhibits (internal) faces of size $O(1)$. Due to space restrictions we cannot provide a proof in this paper but instead refer to the complete version under [5], therefore we only state the following corollary:

Corollary 2.1 The graph induced by S and the adjacencies identified by our algorithm is planar, connected and has (internal) faces of size $O(1)$.

2.2 Conservative adjacencies in \mathbb{R}^3

All the reasoning which has been concentrating on a flat, planar setting can be translated with little effort to our actual setting in \mathbb{R}^3 as we can show that locally the neighborhood graph we construct looks like an α -UDG.

Lemma 2.1 For $\gamma \leq 1/16$ the γ lfs(p)-neighborhood of p in the constructed graph is an α -quasi-Unit-disk graph with $\alpha > 1/\sqrt{2}$.

Now we can invoke Corollary 2.1 which implies that locally for any $p \in S$ the graph constructed by our algorithm is planar, connected and has internal faces of constant size.

What does this mean? The graph that we constructed on the subsample of points S is a *mesh* that is locally planar and covers the whole 2-manifold. The mesh has the nice property that all its *cells* (aka faces) have constant size (number of bounding vertices). The edge lengths of the created adjacencies between S are proportional to the respective local feature sizes. Therefore its connectivity structure faithfully reflects the topology of the underlying 2-manifold.

2.2.1 Algorithm epilog

We did not talk about steps (6) and (7) of our approach since they follow exactly the description in [7] and are not novel to this work. The proofs for convergence both point-wise as well as with respect to triangle normals can be carried over from [7] since S can be made an arbitrarily good, locally uniform ε' -sampling (the original ε -sampling V has to be accordingly denser, i.e. $\varepsilon \ll \varepsilon'$). Therefore, the same theorem holds for the result of our algorithm:

Theorem 2.2 There exists ε^* such that for all $\varepsilon < \varepsilon^*$, smooth surfaces Γ in \mathbb{R}^3 and ε -samplings $V \subset \Gamma$, the triangulated surface $\tilde{\Gamma}$ output by our algorithm satisfies the following conditions:

1. BIJECTION: $\mu : \tilde{\Gamma} \rightarrow \Gamma$, determined by closest point, is a bijection
2. POINTWISE APPROXIMATION: For all $x \in \tilde{\Gamma}$, $d(x, \mu(x)) = O(\varepsilon^2 \text{lfs}(\mu(x)))$
3. NORMAL APPROXIMATION: For all $x \in \tilde{\Gamma}$, $\angle n_{\tilde{\Gamma}}(x) n_{\Gamma}(\mu(x)) = O(\varepsilon)$ where $n_F(y)$ denotes the (outside) normal of F at y^2 .
4. TOPOLOGICAL CORRECTNESS: Γ and $\tilde{\Gamma}$ have the same topological type.

References

- [1] N. Amenta and M. Bern. Surface reconstruction by voronoi filtering. In *Proc. 14th ACM SoCG*, 1998.
- [2] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th SoCG*, 2000.

²For $\tilde{\Gamma}$ the normal is well-defined in the interior of triangles; at edges and vertices it can be defined as an interpolation from that at the incident triangles.

- [3] S.-W. Cheng, T.K. Dey, and E.A. Ramos. Manifold reconstruction from point samples. In *ACM-SIAM SODA*, pages 1018–1027, 2005.
- [4] D. Dumitriu, S. Funke, M. Kutz, and N. Milosavljevic. How much geometry it takes to reconstruct a 2-manifold in \mathbb{R}^3 . To appear at ALENEX 2008. <http://www.mpi-inf.mpg.de/~funke/Papers/GraphRecon.pdf>.
- [5] D. Dumitriu, S. Funke, M. Kutz, and N. Milosavljevic. On the locality of reconstructing a 2-manifold in \mathbb{R}^3 . manuscript. <http://www.mpi-inf.mpg.de/~funke/Papers/LocalRecon.pdf>.
- [6] S. Funke and N. Milosavljevic. Network Sketching or: "How Much Geometry Hides in Connectivity? – Part II". In *Proc. ACM-SIAM SODA*, pages 958–967, 2007.
- [7] S. Funke and E. Ramos. Smooth-surface reconstruction in near-linear time. In *Proc. ACM-SIAM SODA*, 2002.