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

Daniel Dumitriu[1], Stefan Funke[2], Martin Kutz, and Nikola Milosavljević[3]

[1] Max-Planck-Institut für Informatik
Campus E 1.4, 66123 Saarbrücken, Germany
dumitriu@mpi-inf.mpg.de
[2] Ernst-Moritz-Arndt-Universität
Jahnstr. 15a, 17487 Greifswald, Germany
stefan.funke@uni-greifswald.de
[3] Stanford University, Stanford CA 94305, USA
nikolam@stanford.edu

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 prove 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, described in a companion paper [3], preforms well on typical sample sets. Due to its local mode of computation, it might be particularly suited for parallel computing or external memory scenarios.

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 and 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 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* $\mathrm{lfs}(p)$ as the distance of p to the *medial axis*[4]

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

of Γ. A set of points $V \subset \Gamma$ is called a ε-sample of Γ if $\forall p \in \Gamma\ \exists s \in V : |sp| \leq \varepsilon \cdot \mathrm{lfs}(p)$. For sufficiently small ε, Amenta and Bern define a *canonical correct reconstruction* of V with respect to Γ as the set of Delaunay triangles whose dual Voronoi edges in the Voronoi diagram of V intersected 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 wrt to Γ without knowing Γ, no matter how dense the sampling V is. However, there are algorithms that can 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 and in terms of the surface normals, as the sampling density goes to infinity ($\varepsilon \to 0$). The CoCone algorithm [2] is one example; it proceeds in 4 stages: 1) The Voronoi diagram of V is computed. 2) For every point $p \in V$ the surface normal $\overrightarrow{n_p}$ at p is estimated as a vector pointing from p to the furthest point in p's Voronoi cell. 3) A set of *candidate triangles* $\mathcal{T}$ is determined by selecting all Delaunay triangles whose dual Voronoi edge intersects the *CoCones*[5] of all three respective sample points. 4) From the set of candidate triangles that form a "thickened" layer near the real surface, the final piecewise-linear surface approximating Γ is extracted.

The last step of the CoCone algorithm 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 contradictions arise. Not only in theory, but even in practice, the manifold extraction step is still quite challenging and requires nontrivial engineering to actually work as desired.

One potential way of obtaining 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 the respective dual Voronoi edge/face essentially completely pierces the CoCone region and hence are certainly part of the canonical reconstruction as well as any good approximation to it. 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

[5] The CoCone region of a point p with estimated surface normal $\overrightarrow{n_p}$ is the part of p's Voronoi cell that makes an angle close to $\pi/2$ with $\overrightarrow{n_p}$ at p.

the original surface. That is, it is connected, locally planar, and contains no large holes.

We want to point out that the CoCone algorithm (like many other algorithms in that area) has an inherent (theoretical) quadratic worst-case running time since it computes a Voronoi diagram/Delaunay triangulation of a point set in $\mathbb{R}^3$ – the algorithm by Funke and Ramos [4] is an exception since it runs in near-linear time by enforcing the relevant Voronoi computations to take place locally (nevertheless, this algorithm also requires a non-local manifold extraction step). We borrow two ingredients of [4]. In a first step, the algorithm by Funke and Ramos computes a function $\phi(p)\ \forall p \in V$ such that ϕ is Lipschitz[6] and $\phi(p) \leq \varepsilon\,\mathrm{lfs}(p)$. They use this function to "prune" the original point set V to obtain a set S which has certain nice properties. For our theoretical analysis we assume that this function ϕ has been computed in the same manner and use it to construct a local neighborhood graph on which our algorithm operates. For S the algorithm in [4] then locally computes candidate triangles (which due to the nice properties of S can be done locally in near-linear time), and uses the standard (non-local!) manifold extraction step to obtain a reconstruction of S with respect to Γ. Finally, all samples in $V - S$ are reinserted to produce the final reconstruction; this can be done very elegantly using the reconstruction of S as a "reference surface" with respect to which the restricted Voronoi Diagram/Delaunay triangulation of $V - S$ is considered. The restricted VD/DT can easily be computed locally and efficiently via a 2-d weighted Delaunay triangulation on planes supporting the faces of the reconstruction of S. We borrow this last step for our algorithm as we also compute an intermediate reconstruction for a subset of sample points.

In [5] Funke and Milosavljević 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 called *sliver pumping* has been proposed by Cheng et al. in [6]. 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 derived

[6] More precisely their algorithm computes a δ-approximate ω-Lipschitz function ϕ, that is for x, y we have $\phi(x) \leq (1+\delta)(\phi(y) + \omega|xy|)$.

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 that conservative edge creation only leads to small, constant-size faces in the corresponding 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. at any point relying only on adjacency information of points that are geometrically close to the part of the manifold being extracted.

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 [3]) suggests that the approach is viable even for practical use. The results are quite promising, and there is potential for considerable speedup e.g. in parallel computing or external memory scenarios due to the local nature of computation in our new method.

From a technical point of view, two insights are novel in this paper (and not a result of the mere combination of the two 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 [5]. Second, we provide a more elegant and much stronger result about the density of the extracted planar graph based on a connection between the β-skeleton and a power-spanner property; this insight also improves the overall result in [5].

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:

1. Determine a Lipschitz function $\phi(v)$ for every $v \in V$ which lower-bounds $\varepsilon \operatorname{lfs}(v)$ (as in [4])
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 [5])
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 [4])

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. Note that the first and last item from above are original and novel to this paper and do not follow from our previous results in [4] and [5] (the last item makes the theoretical result in [5] much stronger).

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. The main rationale of our approach is the "conservative" creation of adjacencies; that is, we only create an edge between two samples if in any good reconstruction the two points are adjacent, which can be interpreted as creating edges only in the absence of slivers in the vicinity.

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 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, in G all nodes at distance at most α have to be adjacent, while all nodes at distance more than 1 cannot be adjacent. For nodes with distances in between, either is possible.

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 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.

The following algorithm determines adjacencies between nodes in S based on a *Graph Voronoi diagram* such that the induced graph on S remains planar.

Graph-Based Conservative Adjacencies The idea for construction and the planarity property of our construction are largely derived from the geometric intuition. To be specific, the planarity follows from the fact that our constructed graph – we call it *combinatorial Delaunay map* of S, short 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 on the graph connectivity as described in [5]. The reasoning relies on the fact that our graph instance is not an arbitrary graph but reflects the geometry of the underlying domain by being a quasi-unit-disk graph.

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. *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 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 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 [5].

Theorem 1. *If G is an α-qUDG with $\alpha \geq \frac{1}{\sqrt{2}}$ and S is 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. But we will show in the following that this is not the case; in particular we show that the respective graph is connected and all its faces are bounded by a constant number of edges. The following lemmas and proofs are *not* taken from [5], but also apply there, improving the (weaker) statements about the density of CDM(S) in [5].

CDM(S) is Dense(!) Let us consider the β-skeleton [7] of the points corresponding to the node set S. The β-skeleton of a point set has an edge between two points p, q iff any ball of radius $\beta|pq|/2$ touching p and q is empty of other points. For $\beta = 1$ we obtain the well-known Gabriel graph (there is exactly one ball touching p and q with radius $|pq|/2$). For $\beta > 1$ we get a subgraph thereof (there are always two balls of radius $> |pq|/2$ touching p and q).

First we will show that the graph obtained via the β-skeleton is connected[7] and all internal faces of this graph (when using the obvious straight-line embed-

[7] In general the β-skeleton need not be connected; in our case it is due to our choice of β and the local uniformity of S.

ding in the plane) have constant complexity. Then we argue that for suitable parameters k and ε, every β-skeleton edge is also present in $CDM(S)$.

β-skeleton is dense. We establish both connectivity as well as bounded face complexity by showing that the β-skeleton for the point set S is a *σ-power-spanner*, more precisely we show that all power distance optimal paths only use edges of the β-skeleton. For the point set S, the *σ-power distance* $d^\sigma(p,q)$ between two points $p, q \in S$ is determined by a sequence of points $p = p_0 p_1 \dots p_l = q$ such that $d^\sigma(p,q) = \sum_{i=0}^{l-1} |p_i p_{i+1}|^\sigma$ is minimal, where $p_i \in S$. Intuitively, the power distance between two points is the minimum amount of energy required to transmit a message between the two points (potentially using intermediate points as relays) assuming that direct communication between two points at distance d has cost d^σ.

Observation 2. *For $\sigma \geq 2$ every $p_i p_{i+1}$ is a Gabriel edge.*

What we want to show is that for a suitable choice of σ dependent on β, all links $p_i p_{i+1}$ in a power minimal path are even edges of the β-skeleton. We start with a simple observation about the distribution of S.

Lemma 1. *For any pair $s_1, s_2 \in S$ we have $|s_1 s_2| \geq (\alpha - \varepsilon)k$.*

Proof. Follows from the fact that V is an ε-sampling and S being a tight k-sample in a α-qUDG.

The fact that V is an ε-sampling of a disk also implies that all Gabriel (and hence β-skeleton) edges cannot be too long.

Lemma 2. *For any $s_1, s_2 \in S$ such that $|s_1 s_2| > 2(k + \varepsilon)$, $s_1 s_2$ cannot be a Gabriel edge.*

Proof. Assume otherwise. Consider the Gabriel ball of $s_1 s_2$, with center c and radius $r > k + \varepsilon$; due to convexity of a disk and V being a ε-sampling thereof, there must be a point $p \in V$ at distance at most ε from c. But for this p there must exist a $s \in S$ at distance at most k, violating the Gabriel ball property.

Let us now consider one potential edge between nodes p and q which is not part of the β-skeleton and show that this edge cannot be part of any power minimal path.

Lemma 3. *Let $p, q \in S$ but pq not an edge in the β-skeleton. Then if $\sigma \geq \frac{4 \ln 2(k+\varepsilon)}{(\alpha-\varepsilon)k}$ and $\beta \leq \frac{\left(1+\frac{(\alpha-\varepsilon)k}{4(k+\varepsilon)}\right)^2}{2\sqrt{\left(1+\frac{(\alpha-\varepsilon)k}{4(k+\varepsilon)}\right)^2-1}}$, pq cannot be part of any power minimal path.*

Proof. Due to space restrictions we can only provide a very rough sketch. See the authors' homepages for a long version. Due to Observation 2 and Lemma 2, pq cannot be too long. On the other hand, because pq is not part of the β-skeleton, there must be other landmarks inside the two balls of radius $\beta|pq|/2$ touching p and q. Then σ can be chosen such that any power minimal path would rather go via those other landmarks than taking the direct hop pq.

It follows that for the respective choice of σ and β, all power-efficient paths use only edges of the β-skeleton and that the β-skeleton-graph is connected. It remains to show that the graph is somewhat dense, more precisely we want to show that the induced faces (when embedding using the original geometric coordinates) are of constant size.

Lemma 4. *Any bounded face of the β-skeleton is of size $O\left(\left(\frac{2(k+\varepsilon)}{(\alpha-\varepsilon)k}\right)^{2(\sigma+1)}\right)$.*

Proof. Again, we only give a sketch, due to limited space. Fix any bounded face f. Observe that some Delaunay edge $e = (u, v)$ cuts f in a balanced way. Lemmas 1, 2 (a version for Delaunay edges), and 3 upper-bound the hop-distance between u and v in the β-skeleton. A shortest uv-path p need not go along the face boundary, but p and e "enclose" a path p' that does. A packing argument (involving Lemma 1 again) shows that p' cannot be too long either.

At this point we have shown that the β-skeleton of S induces a connected planar graph which has constant-size faces only. It remains to show that all edges of the β-skeleton of S are also identified as adjacencies in $CDM(S)$. To achieve that, observe that any value $\beta > 1$ implies that points $s_1, s_2 \in S$ adjacent in the β-skeleton of S share a "relatively long" Voronoi edge in the Voronoi diagram of S; more precisely s_1 and s_2 share a Voronoi edge of length $|s_1 s_2|\sqrt{\beta^2 - 1}$. Hence it suffices to show that this long Voronoi edge is also reflected in the α-qUDG by respective witness paths. This can be easily achieved by choosing a large enough k (dependent on β). The following corollary follows immediately:

Corollary 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 so far has been concentrating on a flat, planar setting. Let us now consider the actual setting in $\mathbb{R}^3$. Let V be an ε-sampling of a smooth closed 2-manifold Γ in $\mathbb{R}^3$.

Here the steps of our algorithm are as follows: (1) we have to compute a Lipschitz[8] function ϕ with $\phi(p) \leq \varepsilon\,\mathrm{lfs}(p)$ for all $p \in V$. This can be done using the procedure given in [4] in near-linear time. Then in step (2) the graph $G(V, E)$ is constructed by creating edges between samples p_1, p_2 iff $|p_1 p_2| \leq 6 \cdot \phi(p_1)$ or $|p_1 p_2| \leq 6 \cdot \phi(p_2)$ (Note that the constant 6 is somewhat arbitrary and only chosen to make every sample connected to its neighbors in the canonical correct reconstruction). The following steps (3) to (5) are exactly the same as in the 2-dimensional case. We now want to argue that locally around a sample point p the constructed graph looks like an α-quasi-unit-disk graph. To keep the representation simple, we assume that $\phi(p) = \varepsilon\,\mathrm{lfs}(p)$. Although the procedure in [4] yields a ϕ which is a poitwise lower bound for the latter, the argumentation remains the same. The following basic observations are easy to derive (think of $\varepsilon \ll \gamma \ll 1$):

[8] In fact one computes a δ-approximate ω-Lipschitz function.

Lemma 5. *Let $p \in \Gamma$ be a point on the surface, T_p the tangent plane at p, $q \in \Gamma$ some point on the surface with $|pq| \leq \gamma \operatorname{lfs}(p)$, $\gamma < 1$, q' its orthogonal projection on T_p. Then we have $|qq'| < \gamma^2 \operatorname{lfs}(p)$.*

Proof. Consider the two balls of radius $\operatorname{lfs}(p)$ tangent at p. By definition of the local feature size, no point of Γ is contained in the *interior* of either of the two balls. Let q, q' and T_p be defined as above, see Fig. 1, left. In the worst case, $q \in \Gamma$ lies on one of the two balls tangent at p as in Fig. 1, right. The angle at p in triangle $\Delta pqq'$ is equal to the angle at c in the triangle Δcmp – call this angle θ. Since both these triangles are right-angled, we have that $\sin\theta = |pm|/|pc|| = |qq'|/|pq|$, or in other words $|qq'| = |pm||pq|/|pc|$. But since $|pm| = |pq|/2 \leq \gamma \operatorname{lfs}(p)/2$, $|pc| = \operatorname{lfs}(p)$, and $|pq| \leq \gamma \operatorname{lfs}(p)$, we obtain $|qq'| \leq \gamma^2 \operatorname{lfs}(p)/2$.

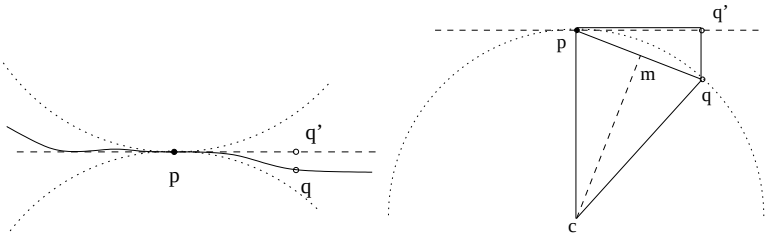

Fig. 1. Sandwiching balls.

Observation 3. *Let $p \in \Gamma$ be a point on the surface, T_p the tangent plane at p, $q_1, q_2 \in \Gamma$ two points on the surface with $|pq_i| \leq \gamma \operatorname{lfs}(p)$, $\gamma < 1$, q_i' their orthogonal projections on T_p. Then we have $|q_1q_2| - |q_1'q_2'| \leq 2\gamma^2 \operatorname{lfs}(p)$.*

The latter observation means that when projecting into the tangent plane at p, distances are shortened by at most an additive amount of $2\gamma^2 \operatorname{lfs}(p)$. We are interested in the largest distance between any two samples s_1, s_2 with $|ps_i| \leq \gamma \operatorname{lfs}(p)$ and the smallest (projected!) distance of a *non-adjacent* pair. If the ratio of these two distances is at most $1/\alpha$ we know that locally in a $\gamma \operatorname{lfs}(p)$-neighborhood of p the constructed graph is an α-quasi-unit-disk graph.

Lemma 6. *For $\gamma \leq 1/16$ the $\gamma \operatorname{lfs}(p)$-neighborhood of p in the constructed graph is an α-quasi-unit-disk graph with $\alpha > 1/\sqrt{2}$.*

Proof. Within a distance of $\gamma \operatorname{lfs}(p)$ from p on the surface the local feature size can *increase* by at most $2\gamma^2 \operatorname{lfs}(p)$ (again a sandwiching argument as in the previous Lemma). Hence the distance of samples identified as adjacent can increase by a factor of $(1 + 2\gamma^2)$. On the other hand, the local feature size can actually *decrease* by $\gamma \operatorname{lfs}(p)$, hence the smallest distance of a non-adjacent pair (taking

into account the projection) can be as little as $6\varepsilon \operatorname{lfs}(p)(1-\gamma-2\gamma^2)$. Therefore the ratio between these two distances is $\frac{1+2\gamma^2}{1-\gamma-2\gamma^2}$. Choosing e.g. $\gamma = 1/16$ makes this ratio less than $\sqrt{2}$.

Now we can invoke Corollary 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.
Note: The faces or more precisely the local embedding can be simply obtained by reusing the geometry again and locally projecting the adjacent points s_i of a point $s \in S$ into an (almost) tangent plane and reading off the cyclic order around s. Global connectivity is ensured by choosing γ large enough (compared to ε) such that large faces are completely contained in the $\gamma \operatorname{lfs}(p)$-neighborhood of any node bounding the face.

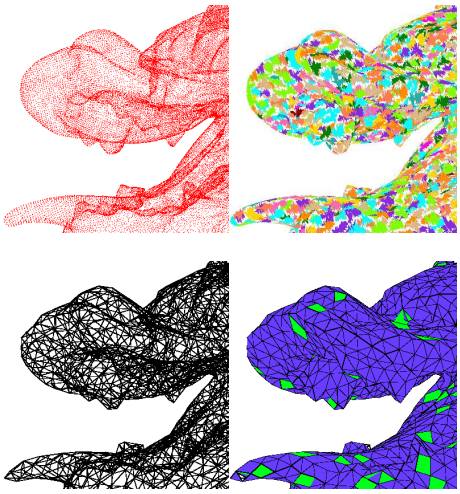

Fig. 2. Main steps of our algorithm: Point cloud, Graph Voronoi Diagram, $CDM(S)$, identified faces.

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.

Algorithm Epilog We did not talk about steps (6) and (7) of our approach since they follow exactly the description in [4] and are not novel to this work; we nevertheless give a brief summary here. Essentially in step (6) we triangulate non-triangular faces by projecting them into a nearby (almost) tangent plane and computing the Delaunay triangulation. The resulting triangulated faces behave nicely since all faces have small size (and hence their vertices are almost coplanar) and because S is a *locally uniform* sampling of the surface. In step (7) the points pruned in step (3) are reinserted by computing a weighted Delaunay triangulation on the supporting planes of the respective faces. The resulting triangulations are guaranteed to patch up. The proofs for convergence both point-wise as well as with respect to triangle normals can be carried over from [4] 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 4. *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 $\widetilde{\Gamma}$ output by our algorithm satisfies the following conditions:*

1. Bijection*: $\mu : \widetilde{\Gamma} \to \Gamma$, determined by closest point, is a bijection*
2. Pointwise Approximation*: For all $x \in \widetilde{\Gamma}$, $d(x, \mu(x)) = O(\varepsilon^2 \operatorname{lfs}(\mu(x)))$*
3. Normal Approximation*: For all $x \in \widetilde{\Gamma}$, $\angle n_{\widetilde{\Gamma}}(x) n_\Gamma(\mu(x)) = O(\varepsilon)$ where $n_F(y)$ denotes the (outside) normal of F at y.*[9]
4. Topological Correctness*: Γ and $\widetilde{\Gamma}$ have the same topological type.*

3 Implementation and Experimental Evaluation

We have prototyped the novel steps of our algorithm in `C++`. This implementation is the topic of a companion paper to this paper ([3]). This companion paper also highlights another novelty of our approach in that most of our computation does *not* use any geometry information: after establishing neighborhood relations between nearby points, the main steps of our algorithm operate combinatorially on a graph structure. As such, they are by far less susceptible to robustness problems due to round-off errors in floating-point arithmetic.

In Fig. 2 we have visualized the main steps of our algorithms: starting with a point cloud we compute a Graph Voronoi diagram, based on that the $CDM(S)$ and finally we inspect the $CDM(S)$ to identify faces, some of which might be non-triangular (here in light color) due to the conservative edge creation. In Fig. 3 you can find a picture of a reconstruction of a complete object (the standard "Dragon" dataset).

4 Outlook

Theoretically our approach has the potential to work for reconstructing 2-manifolds even in higher dimensions. It does not extend to non-2-manifolds, though, as the "local planarity property" of a graph that our algorithm crucially depends upon, has no equivalent for non-2-manifolds.

A more in-depth discussion on the advantages/disadvantages of dropping geometry information early-on can be found in the companion paper to this submission [3], further studies are required, though, whether this can be extended to a general paradigm when designing geometric algorithms.

In parallel computing or external memory scenarios, it is much easier to obtain efficient algorithms if the performed operations require only *local* access to data. In the former, making non-local data available typically incurs a runtime penalty for the data transfer or more complicated access control mechanisms, in the latter, local data can be cached in internal memory, while non-local data has to be read from external memory again incurring considerable latency. The

[9] For $\widetilde{\Gamma}$ 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.

manifold extraction step as for example employed by the CoCone algorithm is a global, highly non-local operation. It remains to be seen whether the localization property exhibited in this paper leads to practically more efficient algorithms in the parallel computing or external memory scenario.

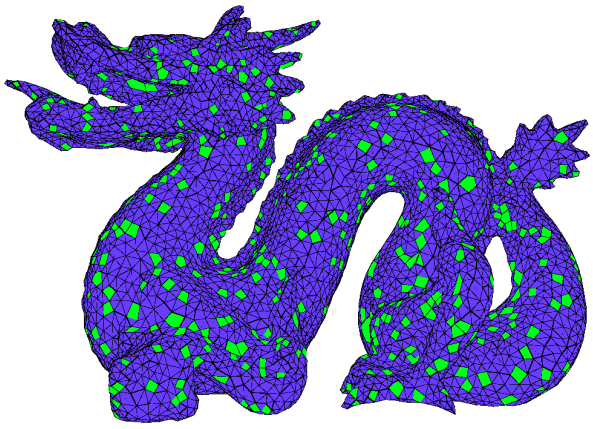

Fig. 3. Output of our implementation for the standard "Dragon" model. Non-triangular faces are denoted in light color.

Acknowledgment

In memory of our dear friend and colleague Martin Kutz.

References

1. Amenta, N., Bern, M.: Surface reconstruction by Voronoi filtering. In: Proc. 14th ACM SoCG. (1998)
2. Amenta, N., Choi, S., Dey, T.K., Leekha, N.: A simple algorithm for homeomorphic surface reconstruction. In: Proc. 16th SoCG. (2000)
3. Dumitriu, D., Funke, S., Kutz, M., Milosavljević, N.: How much Geometry it takes to Reconstruct a 2-Manifold in $\mathbb{R}^3$. In: Proc. 10th ACM-SIAM ALENEX. (2008) 65–74
4. Funke, S., Ramos, E.: Smooth-surface reconstruction in near-linear time. In: Proc. ACM-SIAM SODA. (2002)
5. Funke, S., Milosavljević, N.: Network Sketching or: "How Much Geometry Hides in Connectivity? – Part II". In: Proc. ACM-SIAM SODA. (2007) 958–967
6. Cheng, S.W., Dey, T., Ramos, E.: Manifold reconstruction from point samples. In: ACM-SIAM SODA. (2005) 1018–1027
7. Kirkpatrick, D.G., Radke, J.D.: A framework for computational morphology. In Toussaint, G.T., ed.: Computational Geometry. North-Holland, Amsterdam, Netherlands (1985) 217–248