

Untangling Triangulations through Local Explorations ^{*}

Pankaj K. Agarwal

Bardia Sadri

Hai Yu

March 18, 2008

Abstract

The problem of maintaining a valid mesh (triangulation) within a certain domain that deforms over time arises in many applications. During a period for which the underlying mesh topology remains unchanged, the deformation moves vertices of the mesh and thus potentially turns a mesh invalid, or as we call it, *tangled*. We introduce the notion of *locally removable regions*, which are certain tangled regions in the mesh that allow for local removal and re-meshing. We present an algorithm that is able to quickly compute, through local explorations, a minimum locally removable region containing a “seed” tangled region in an invalid mesh. By re-meshing within this area, the “seed” tangled region can then be removed from the mesh without introducing any new tangled region. The algorithm is output-sensitive in the sense that it never explores outside the output region.

1 Introduction

In physical simulation, scientific computing, computer graphics and many other applications, it is often required to maintain a valid mesh (triangulation) within a certain domain that deforms over time. This problem has attracted much attention in recent years, and several techniques have been proposed. For example, the space-time meshing method [9, 17] builds a mesh over the entire space-time domain whose resolution is adaptive to the movement of the underlying space; numerical simulation can then be carried out directly on the space-time mesh. The drawbacks of this method are that one needs to mesh a domain that is one dimension higher than the original space, as well as *a priori* knowledge on the physical property of the underlying space (e.g., wave speed) in order to determine the domains of influence and of dependence of a point in space-time.

Another approach is to maintain the mesh incrementally, by discretizing the time axis and updating the mesh only at these discrete time instances. For example, the kinetic triangulation method [1, 2, 3] uses an event-driven framework (i.e., the kinetic data structure framework by Basch *et al.* [5]) to proactively detect when the mesh will become invalid, and repair it immediately after this occurs, thus maintaining a valid mesh all the time. See also [6, 8] for similar work on dynamic skin triangulations. Such methods usually require accurate knowledge of the deformation of the domain to predict critical events, and a significant amount of extra storage to keep track of these events. A more recent work of Cheng and Dey [7] on maintaining a provably good mesh over a set of sample points from a deforming surface also loosely follows the above approach. Their algorithm assumes weaker knowledge of the motion, specifically, an upper bound on the velocities

^{*}Work on this paper is supported by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376, by an NIH grant 1P50-GM-08183-01, by a DOE grant OEG-P200A070505, and by a grant from the U.S.–Israel Binational Science Foundation. Authors’ affiliation: Department of Computer Science, Duke University, Durham, NC 27708-0129, USA. Email: {`pankaj,sadri,fishhai`}@cs.duke.edu.

of the points, to discretize the time appropriately, so that within each time step, the points move at most a fraction of the smallest feature size of the surface. Their algorithm requires the sample points to be sufficiently uniformly dense over the entire deformation process.

A lazy version of the kinetic triangulation method, which is perhaps more popular among practitioners, is to ensure correctness of the mesh only at fixed or adaptive time steps; in between two consecutive time steps, the mesh can be either valid or invalid. Since many numerical algorithms also discretize the time domain, and computation is done at these discrete time instances, this approach is especially suitable for these algorithms. Between two time steps, the deformation moves vertices of the mesh while the underlying mesh topology remains unchanged, thus potentially turning a mesh invalid, or as we call it, *tangled*, because elements of the mesh may intersect inadmissibly. When this happens, an “untangling” process needs to be invoked at the next time step to restore the validity of the mesh. Of course one may simply re-mesh the entire domain. However, because of the continuity of the deformation and the fine granularity of the time discretization in most situations, the extent of tangling is rather small compared to the size of the entire mesh. In this case, one may substantially benefit from performing local explorations to quickly detect and remove tangled regions in the invalid mesh, followed by local re-meshing.




Figure 1: (i) A valid plane triangulation. (ii) The triangulation becomes tangled after three vertices moved. In particular, the colored triangles become inverted. (iii) The untangled triangulation.

Problem statement. In this paper we study local untangling of planar triangulations. Intuitively, a tangled triangulation is the image of a plane triangulation under a continuous map that maps each triangle linearly but is not injective. It is easy to see that these maps precisely arise from moving the vertices of a triangulation freely while retaining the same abstract triangulation. We therefore treat a tangled triangulation as the image of a piecewise linear map that extends a “motion map” defined on the vertex set of a plane triangulation. This is formalized as follows.

Let $\mathcal{V} \subset \mathbb{R}^2$ be a finite set of points (vertices) in general position. Let S be the convex hull of \mathcal{V} . Let $(\mathcal{V}, \mathcal{E}, \mathcal{T})$ be a triangulation of S spanned over \mathcal{V} , where \mathcal{E} and \mathcal{T} denote the set of edges and triangles, respectively, in the triangulation. Let $f_{\mathcal{V}} : \mathcal{V} \rightarrow \mathbb{R}^2$ be a *motion* map that represents the new position of each vertex after a motion period. The map $f_{\mathcal{V}}$ extends to a map $f : S \rightarrow \mathbb{R}^2$ through linear interpolation in the relative interior of edges and triangles of \mathcal{T} . Clearly, the map f is continuous. For a subset $R \subseteq S$, we sometimes refer to $f(R)$ as the *shadow* of R . Let $\tilde{\mathcal{V}} = \{f(v) \mid v \in \mathcal{V}\}$, $\tilde{\mathcal{E}} = \{f(e) \mid e \in \mathcal{E}\}$ and $\tilde{\mathcal{T}} = \{f(T) \mid T \in \mathcal{T}\}$. We call $(\tilde{\mathcal{V}}, \tilde{\mathcal{E}}, \tilde{\mathcal{T}})$ a *tangled triangulation* of $\tilde{\mathcal{V}}$. When no confusion arises, we simply call \mathcal{T} a triangulation of \mathcal{V} and $\tilde{\mathcal{T}}$ a tangled triangulation of $\tilde{\mathcal{V}}$; the vertex set and the edge set of the triangulation are automatically understood as those of the collection of triangles.

We assume that vertices in $\tilde{\mathcal{V}}$ are in general position: no three vertices are collinear, and no three line segments spanned by different vertices intersect at the same point. We also make the following boundary assumption: (i) f is identity on ∂S , and (ii) $f(\text{int } S) \subseteq \text{int } S$. The assumption can be

enforced on any mesh by introducing three dummy vertices at infinity. Note that the assumption implies $f(S) = S$; therefore in the following we use S to denote both the domain and the range of the map f . For the sake of clarity, we will use Roman letters (v, x, y , etc.) to denote points in the domain S , and Greek letters (ξ, η , etc.) to denote points in the range S . For the rest of this paper, we fix \mathcal{T} and f , and define subsequent notions relative to them. We emphasize that, for an input tangled triangulation $\tilde{\mathcal{T}}$, we do not assume its corresponding \mathcal{T} and f are given; the use of \mathcal{T} and f is only for the clarify of presentation.

We orient the boundary ∂S in counterclockwise order, and extend this orientation to all triangles in \mathcal{T} in the standard way. For a triangle $T \in \mathcal{T}$ with vertices oriented in the order $[v_1, v_2, v_3]$, we call T as well as its image $f(T) \in \tilde{\mathcal{T}}$ *upright* if $[f(v_1), f(v_2), f(v_3)]$ is in counterclockwise order, and *inverted* otherwise. By the above boundary assumption, every triangle $T \in \mathcal{T}$ adjacent to ∂S and its image $f(T)$ are upright.

Given a tangled triangulation $\tilde{\mathcal{T}}$, the goal of untangling is to remove a subset $\mathcal{U} \subseteq \tilde{\mathcal{T}}$ of triangles from $\tilde{\mathcal{T}}$, which includes all inverted triangles (and unavoidably a few upright triangles), and replace \mathcal{U} with a new set \mathcal{U}' of triangles so that $(\tilde{\mathcal{T}} \setminus \mathcal{U}) \cup \mathcal{U}'$ is a valid triangulation of $\tilde{\mathcal{V}}$. In this paper we mainly consider the following form of the problem: given a “seed” inverted triangle $T \in \tilde{\mathcal{T}}$, compute a subset $\mathcal{U} \subseteq \tilde{\mathcal{T}}$ of triangles with $T \in \mathcal{U}$ and another set \mathcal{U}' of triangles, so that $(\tilde{\mathcal{T}} \setminus \mathcal{U}) \cup \mathcal{U}'$ is a tangled triangulation of $\tilde{\mathcal{V}}$ in which all triangles in \mathcal{U}' are upright. An untangling algorithm for this form of problem provides flexibility and compatibility for different applications. For example, many algorithms for local mesh smoothing that work only on valid meshes (e.g., [4, 10, 13] and the Delaunay edge-flip algorithm) can now be extended to tangled meshes as follows: whenever an inverted triangle is encountered during mesh smoothing, the untangling algorithm is invoked to remove that inverted triangle so that smoothing can be continued in that local region.

Related work. Edge flip is probably the most elegant atomic operation for transforming triangulations, whose power is well illustrated in Lawson’s celebrated algorithm [15] for converting an arbitrary planar triangulation to the Delaunay triangulation of its vertices, guided only by local geometry. However, it is an open question whether a tangled triangulation can be untangled by a simple edge-flip-based algorithm. On the optimistic side, for every tangled triangulation there exists a sequence of edge flips that converts it into a valid one. However, it is not clear whether this sequence can be found by applying a number of simple guiding rules based on local geometric information of the tangled mesh, in a way similar to Lawson’s algorithm. (Note that, however, if the motion map $f_{\mathcal{V}}$ is known, then one can find this edge-flip sequence easily.)

Shewchuk and Wallace [16] study the problem of untangling triangulations by performing local surgeries to the tangled mesh. The primary local operation used by their algorithm is edge flip. Occasionally their algorithm may perform other more complicated local operations in which vertices defining the triangulation may be moved or deleted, and new vertices may be inserted. They showed that, by repeatedly applying one of these operations, the mesh can be untangled monotonically, in the sense that the total area of the inverted triangles never increases as the algorithm proceeds and eventually becomes zero.

Some popular methods for untangling triangulations in scientific computing community are based on optimization techniques [11, 12, 14, 18]. In these approaches, the vertices of the triangulation are allowed to move, and the mesh is untangled by moving vertices to a new configuration that locally optimizes a certain objective function. For example, the algorithm of Freitag and Plassmann [11] untangles the mesh by maximizing the minimum (signed) area of the triangles in the tangled regions of the mesh. The optimization problem can be solved using a technique analogous to gradient descent, which is guaranteed to converge because of the convexity of the level sets of

the objective function. Using more complicated objective functions that take into account various parameters of the quality of the mesh (e.g., minimum angle, aspect ratio of the triangles, etc.), their algorithm accomplishes mesh untangling and mesh smoothing simultaneously.

Our results. Our approach for untangling triangulations strictly preserves the vertex set of the triangulation: no vertex in $\tilde{\mathcal{V}}$ is inserted, deleted or moved after untangling. To this end, we introduce a class of regions in \mathcal{T} called *locally removable regions* which, intuitively speaking, can be locally removed from \mathcal{T} and re-meshed with a set of upright triangles. Given a “seed” inverted triangle $T \in \mathcal{T}$, we show how to find a minimum locally removable region $R \subseteq \mathcal{T}$ containing T , that is, for any locally removable region $R' \subseteq \mathcal{T}$ containing T , $R \subseteq R'$. Our algorithm is output-sensitive in the sense that it never explores beyond the output region R .

As a result of the intricate structure of the tangled mesh, a locally removable region is often induced by a wide collection of inverted triangles rather than just those in the immediate neighborhood of the “seed” triangle. Identifying these inverted triangles through local exploration is not straightforward. Furthermore, during the local exploration process, the region that has been explored up to any stage may have a complicated structure (e.g., the shadow of its boundary may have a large turning number), and it is not always clear how to continue the search until a desired re-meshable region is reached. Our algorithm overcomes these difficulties, by exploiting several interesting observations about the structure of tangled meshes, which we deem as one of the main contributions of the paper.

The rest of the paper is organized as follows. In Section 2 we introduce the notion of *conflict sets* and *primary conflict sets* and prove a number of their useful properties. In Section 3 we define *locally removable sets* that allow for local removal and re-meshing. In Section 4 we make use of primary conflict sets to characterize the structure of locally removable regions, and provide an algorithm for computing these regions. In Section 5 we prove a structural theorem relating untangling the entire mesh to local untanglings studied in Section 4. Some concluding remarks are given in Section 6.

2 Signs, Conflict Sets, and Shadows

In the triangulation \mathcal{T} , an edge common to two inverted or two upright triangles is called a *regular edge*. A regular edge between two upright (resp. inverted) triangles is called an *upright edge* (resp. *inverted edge*). In contrast, we call an edge common to one upright and one inverted triangle a *crease edge*. If all the triangles adjacent to a vertex v are upright (resp. inverted), v is called an *upright vertex* (resp. *inverted vertex*). If v is upright nor inverted, v is also called a *regular vertex*; otherwise v is called a *crease vertex*. A crease vertex must be incident upon at least two crease edges.

Lemma 2.1 *For any set $R \subseteq S$, if $x \in f^{-1}(\partial f(R)) \cap R$, then either $x \in \partial R$ or x lies on a crease edge.*

Proof. Let $x \in \text{int } R$ that does not lie on a crease edge. Let $T_1, \dots, T_k \in \mathcal{T}$ be the triangles adjacent to x . Either all of them are upright or all of them are inverted. Therefore the union of their shadows form a star-shaped polygon $\bigcup_{i=1}^k f(T_i)$ with $f(x)$ contained in its interior. It is clear that f is bijective restricted to a sufficiently small neighborhood $N \subseteq \bigcup_{i=1}^k T_i$ of x . Hence, $x \in \text{int } f(U)$, which completes the proof. ■

Let $Q = \{\xi \in S \mid \xi = f(v) \text{ for some crease vertex } v \in \mathcal{V}\}$. Note that both Q and $f^{-1}(Q)$ are finite. For technical reasons, we need to exclude Q from the range S and $f^{-1}(Q)$ from the domain

S . For the rest of the paper, we slightly abuse the notation and continue to use S to denote the domain $S \setminus f^{-1}(Q)$ and the range $S \setminus Q$.

Signed incidence function. We define a sign function $\chi : S \rightarrow \{+1, 0, -1\}$ by letting

$$\chi(x) = \begin{cases} +1 & \text{if } x \text{ is an upright vertex or in the interior of an upright edge or triangle,} \\ -1 & \text{if } x \text{ is an inverted vertex or in the interior of an inverted edge or triangle,} \\ 0 & \text{if } x \text{ is on a crease edge.} \end{cases}$$

In particular, we refer to $x \in S$ as a negative point if $\chi(x) = -1$ and a positive point if $\chi(x) = +1$. If $\chi(x) \neq 0$, then $\chi(x) = \chi(y)$ for all points y in a sufficiently small neighborhood of x . Given a set $U \subseteq S$, we define $\gamma_U : f(U) \rightarrow \mathbb{Z}$, the *signed incidence function relative to U* , as

$$\gamma_U(\xi) = \sum_{x \in f^{-1}(\xi) \cap U} \chi(x),$$

where $f^{-1}(\xi) = \{y \in S \mid f(y) = \xi\}$. If for a set $U \subseteq S$, there is a constant $c_U \in \mathbb{Z}$ such that $\gamma_U(\xi) = c_U$ for every regular point $\xi \in f(U)$, we write $\gamma_U = c_U$.

Lemma 2.2 $\gamma_S = +1$.

Proof. By our boundary assumption, $\gamma_S(\eta) = +1$ for any point $\eta \in \partial S$. Let ξ be a point in S . Choose an arbitrary point $\eta \in \partial S$ and take a path $\Pi : [0, 1] \rightarrow S$ from η to ξ that does not self-intersect and is entirely contained in S . Clearly, Π can be chosen so that it intersects each edge of $\tilde{\mathcal{E}}$ transversally and avoids $\tilde{\mathcal{V}}$ altogether. We next show that the function $\gamma_S(\Pi(t))$ does not change as t varies from 0 to 1. Therefore $\gamma_S(\xi) = \gamma_S(\eta) = +1$.

First observe that for any individual triangle $T \in \mathcal{T}$, the restricted map $f : T \rightarrow f(T)$ is a linear homeomorphism and in particular a bijection. Consider any section Π' of Π that does not intersect $\tilde{\mathcal{E}}$. Every connected component of $f^{-1}(\Pi')$ is contained in a single triangle in \mathcal{T} and therefore every such component is a path bijectively mapped to Π' by f . Moreover, all the points in such a component have the same sign, thereby implying that the value of the function γ_S is constant on Π' .

Now consider any section Π' of Π that intersects the image of no edge other than $f(e)$ of some $e \in \mathcal{E}$. First consider the case in which e is a regular edge; without loss of generality, assume that e is incident upon two upright triangles $T_1, T_2 \in \mathcal{T}$. The previous argument can be extended to this case by observing that when we look at $f^{-1}(\Pi')$, we can take $T_1 \cup T_2$ as one component for which $f^{-1}(\Pi') \cap (T_1 \cup T_2)$ is mapped to Π' bijectively. Thus the value of γ_S cannot change when Π' crosses $f(e)$.

Next consider the case in which e is a crease edge. Once again consider the connected components of $f^{-1}(\Pi')$. With the exception of one component, every other component is fully contained in a single triangle and therefore has a fixed sign. The exception is for the component that intersects e . Of the two triangles T_1 and T_2 incident upon e , one is upright and the other is inverted. Every point of Π' has precisely one inverse image in T_1 and another in T_2 . As such, the combined contribution of T_1 and T_2 to γ_S of a point in Π' is zero. Therefore, this component has no effect on the value of γ_S in any point of Π' . Thus γ_S is constant on Π' .

In summary, $\gamma_S(\Pi(t))$ does not change as t varies from 0 to 1, thereby implying $\gamma_S(\xi) = \gamma_S(\eta) = +1$. ■

Inverted components and conflict sets. It is beneficial to deal with the untangling problem at a coarser granularity than individual triangles. Define the binary relation \sim between triangles in \mathcal{T} by letting $T_1 \sim T_2$ if T_1 and T_2 share a regular edge, i.e., T_1 and T_2 are both upright or both inverted. Let $\overset{*}{\sim}$ be the transitive closure of \sim . It can be proved that $\overset{*}{\sim}$ is an equivalence relation. Each equivalence class of $\overset{*}{\sim}$ consists of triangles that are all upright or all inverted, and are connected through edges — or more precisely, their induced dual subgraph is connected. In the sequel, by an *inverted component* (resp. *upright component*) we refer to an equivalence class of the $\overset{*}{\sim}$ relation consisting of inverted (resp. upright) triangles. We use \mathbb{I} to denote the set of all inverted components.

In the following, it is advantageous to think of f as a continuous map independent of the triangulation based on which it is defined. Indeed, the same map can be defined using several triangulations; in particular, the triangulation in the domain and range of the map can be refined consistently using subdivisions without affecting the map itself.




Figure 2: A one-dimensional example of conflict set of an inverted component. Upright and inverted components are respectively represented by horizontal and slanted segments. The grayed region shows the conflict set of the inverted component shown in heavier black. The darker gray represents the primary conflict set of the mentioned inverted component.

We define the *conflict set* $\phi(U)$ of a subset $U \subseteq S$ as the preimage of $f(U)$, i.e.,

$$\phi(U) = f^{-1} \circ f(U)$$

By definition, $U \subset \phi(U)$. Note that even if U is connected $\phi(U)$ may be not so. We are particularly interested in the connected components of $\phi(U)$ that intersect U .¹ We define $\hat{\phi}(U)$ as the union of these connected components and call it the *primary conflict set* of U . The following lemma states a few properties of conflict sets.

Lemma 2.3 *For any set $U \subseteq S$,*

- (1) $f(\phi(U)) = f(\hat{\phi}(U)) = f(U)$;
- (2) $f(\partial\phi(U)) \subseteq \partial f(U)$;
- (3) if $x \in f^{-1}(\partial f(U)) \setminus \partial\phi(U)$, then x lies on a crease edge;
- (4) for any $U' \subseteq \hat{\phi}(U)$, $\hat{\phi}(U') \subseteq \hat{\phi}(U)$.

Proof. (1) Since $U \subseteq \phi(U)$, $f(U) \subseteq f(\phi(U))$. On the other hand, $f(\phi(U)) = f \circ f^{-1} \circ f(U) \subseteq f(U)$. Hence $f(U) = f(\phi(U))$. Since $U \subseteq \hat{\phi}(U) \subseteq \phi(U)$, the claim follows.

¹Throughout the paper, by connectedness we always mean path-connectedness.

(2) Since f is continuous, the inverse image of every open set is open. Therefore, for every $\xi \in \text{int } f(U)$, every point in $f^{-1}(\xi)$ is contained in the interior of $f^{-1} \circ f(U) = \phi(U)$. Thus the image of $\partial\phi(U)$ under f is contained in $\partial f(U)$.

(3) This follows directly from (1) and Lemma 2.1.

(4) By (1), $U' \subseteq \hat{\phi}(U)$ implies $f(U') \subseteq f(U)$. Therefore, $\phi(U') \subseteq \phi(U)$, i.e., every connected component of $\phi(U')$ is contained in some connected component of $\phi(U)$. In particular, every connected component W of $\hat{\phi}(U')$ is contained in one connected component of $\phi(U)$. Since W intersects a connected component of U' by definition, W also intersects $\hat{\phi}(U)$, which further implies W is contained in $\hat{\phi}(U)$. Hence $\hat{\phi}(U') \subseteq \hat{\phi}(U)$. ■

In the rest of the paper, as a convention, for $\mathcal{J} \subseteq \mathbb{I}$, $f(\mathcal{J})$ (resp., $\phi(\mathcal{J})$, $\hat{\phi}(\mathcal{J})$) should be interpreted as $f(\bigcup_{I \in \mathcal{J}} I)$ (resp., $\phi(\bigcup_{I \in \mathcal{J}} I)$, $\hat{\phi}(\bigcup_{I \in \mathcal{J}} I)$).

Lemma 2.4 *Let $\mathcal{J} \subseteq \mathbb{I}$ be a collection of inverted components. Then*

- (1) $\partial\phi(\mathcal{J})$ may only intersect \mathcal{J} in $\mathcal{V} \cap B$, where $B = \bigcup_{I \in \mathcal{J}} \partial I$;
- (2) there is no consecutive sequence of points along $\partial\phi(\mathcal{J})$ whose signs are zero.

Proof. Since $\mathcal{J} \subseteq \phi(\mathcal{J})$, $\partial\phi(\mathcal{J}) \cap \mathcal{J} \subseteq B$. B consists of crease edges as \mathcal{J} is a collection of inverted components. Let x be a point in the relative interior of an edge $e \subset B$ that is incident upon an inverted triangle $T_1 \in \mathcal{J}$ and an upright triangle $T_2 \notin \mathcal{J}$. There is a sufficiently small neighborhood N of x such that $f(N) \subset f(T_1) \cap f(T_2)$. This implies that $N \subset \phi(T_1) \subseteq \phi(\mathcal{J})$ and that $x \notin \partial\phi(\mathcal{J})$. Hence, only vertices on B can appear on $\partial\phi(\mathcal{J})$, thus proving (1). (2) can be proved by a similar argument. ■

Lemma 2.5 *Let $\mathcal{J} \subseteq \mathbb{I}$ be a set of inverted components, and let U be a connected component of $\phi(\mathcal{J})$. There is a constant $c_U \in \mathbb{Z}$ such that $\gamma_U = c_U$. Furthermore, if ∂U does not intersect the interior of any inverted component, then $c_U > 0$.*

Proof. Let η and ξ be two points in $f(U)$. Since $f(U)$ is connected, we choose a path $\Pi : [0, 1] \rightarrow f(U)$ with $\Pi(0) = \eta$ and $\Pi(1) = \xi$. We can choose Π such that its interior points $\text{int } \Pi = \{\Pi(t) \mid 0 < t < 1\}$ are contained in the interior of $f(\mathcal{J})$. By Lemma 2.3 (2), $f(\partial U) \subseteq \partial f(\mathcal{J})$ and thus $\text{int } \Pi$ avoids $f(\partial U)$. Now following a similar argument given for Lemma 2.2, it can be shown that the value of $\gamma_U(\Pi(t))$ does not change when t goes from 0 to 1, implying $\gamma_U(\eta) = \gamma_U(\xi)$. This proves the first half of the lemma.

We next prove the second half of the lemma. Suppose ∂U does not intersect the interior of any inverted component, that is, ∂U contains no negative points. Lemma 2.4 (2) further shows that ∂U cannot entirely consist of points whose signs are 0. Therefore, we can always pick a positive point $x \in \partial U$. Lemma 2.3 (3) and the fact that ∂U does not intersect the interior of any inverted component imply that, for $\xi = f(x)$, no points in $f^{-1}(\xi) \cap U$ are negative. Since x is positive and $x \in f^{-1}(\xi) \cap U$, it follows that $\gamma_U(\xi) \geq \chi(x) > 0$, implying $c_U > 0$. ■

At first sight, it seems hard to imagine γ_U of a connected component $U \in \phi(\mathcal{J})$, as described in Lemma 2.5, to be anything other than -1 , 0 , or $+1$. However, it happens that γ_U can indeed be an arbitrarily positive or negative integer. In fact this is true even for primary conflict set $\hat{\phi}(I)$ of a single inverted component I .

Lemma 2.6 *For any integer $k \in \mathbb{Z}$, there exists a tangled triangulation and an inverted component I of it, for which $\gamma_{\hat{\phi}(I)} = k$.*

Proof. We only describe a construction for $k = +2$, which easily extends to other values of k . The stages of the construction are shown in Figure 3. First the figure on the left illustrates a valid triangulation; the thick edges represent the links of p and p' . We now move the vertices a, b, c and a', b', c' to their new positions as depicted in the middle figure. The gray area I becomes inverted, and all other triangles remain upright. In the right figure, we add one extra layer of triangles to the star of p' , in such a way that the entire inverted component is covered one more time; all the added triangles are upright. Now the boundary of the entire piece in the right figure is a curve winding around the inverted component twice. Indeed the boundary of the primary conflict set $\hat{\phi}(I)$ of this inverted component will have the same property and runs entirely inside upright triangles. This translates to $\gamma_{\hat{\phi}(I)} = +2$.




Figure 3: A construction showing the possibility of having large γ_U even when $U = \hat{\phi}(I)$ of a single inverted component I (the gray region).

It remains to show that the right figure can indeed be extended to a tangled triangulation conforming to our boundary assumption. Take two copies of the constructed piece, turn one of them upside down (reverse the role of inverted and upright triangles) and glue them together along their boundaries. So, the boundary edges in the right figure will turn out to be crease edges. We have now a topological sphere. Finally, poke a hole in the original inverted component I (which is part of this sphere) and a similar hole in a large plane triangulation and glue the punctured construction to the punctured plane triangulation along the boundaries of these holes. ■

Shadows and neighborhoods. For a point $x \in S$ and a set $X \subseteq S$, the *neighborhood of x under shadow X* is defined as

$$\mathcal{N}(x, X) = \{y \in S \mid \exists \text{ a path } \Pi \subset S \text{ from } x \text{ to } y \text{ such that } f(\Pi) \subset X\}.$$

Intuitively, $\mathcal{N}(x, X)$ is the region that can be reached from x without stepping out of X under the map f . Note that for any $y \in \mathcal{N}(x, X)$, $\mathcal{N}(y, X) = \mathcal{N}(x, X)$. So, for a connected set $R \subseteq S$, we can write $\mathcal{N}(R, X) = \mathcal{N}(x, X)$ for any $x \in R$. If R consists of multiple connected components $\{R_1, R_2, \dots, R_k\}$, we let $\mathcal{N}(R, X) = \bigcup_{i=1}^k \mathcal{N}(R_i, X)$.

The following simple lemma provides a convenient way to characterize conflict sets in terms of their shadows and neighborhoods.

Lemma 2.7 *Let $\mathcal{J} \subseteq \mathbb{I}$ be a collection of inverted components, U a connected component of $\phi(\mathcal{J})$, and x an arbitrary point in U . Then $U = \mathcal{N}(x, f(\mathcal{J})) = \mathcal{N}(x, f(U))$ and $\hat{\phi}(\mathcal{J}) = \mathcal{N}(\mathcal{J}, f(\mathcal{J}))$.*

Proof. By Lemma 2.3 (1), $f(U) \subseteq f(\mathcal{J})$. Hence for any point $y \in U$, the path $\Pi \subset U$ from x to y satisfies $f(\Pi) \subset f(\mathcal{J})$, thus serves as a witness for $y \in \mathcal{N}(x, f(\mathcal{J}))$. As such, $U \subseteq \mathcal{N}(x, f(\mathcal{J}))$. On the other hand, fix a point $y \in \mathcal{N}(x, f(\mathcal{J}))$ and a path Π from x to y with $f(\Pi) \subset f(\mathcal{J})$. It follows that $\Pi \subset \phi(\mathcal{J})$. Therefore y lies in the same connected component of $\phi(\mathcal{J})$ as x , implying $y \in U$. Thus $\mathcal{N}(x, f(\mathcal{J})) \subseteq U$ and thereby $U = \mathcal{N}(x, f(\mathcal{J}))$. An identical argument shows $U = \mathcal{N}(x, f(U))$. Since each connected component of $\hat{\phi}(\mathcal{J})$ contains at least an inverted component in \mathcal{J} , $\hat{\phi}(\mathcal{J}) = \mathcal{N}(\mathcal{J}, f(\mathcal{J}))$. ■

3 Locally Removable Regions

We call a region $R \subseteq S$ *locally removable* if

- (i) R is a connected, compact set bounded by simple curves;
- (ii) ∂R does not intersect the interior of any inverted component or the interior of any crease edge;
- (iii) $\gamma_R = +1$.

Note that (ii) is equivalent to that ∂R consists of positive points (recall that crease vertices have been excluded from the domain). By (ii), if R contains a negative point $x \in S$, then R contains the entire inverted component $I \in \mathbb{I}$ containing x . A locally removable region R is called *canonical* if $R = \hat{\phi}(\mathcal{J})$ for some $\mathcal{J} \subseteq \mathbb{I}$.

Lemma 3.1 *If $R \subseteq S$ is a locally removable region, then the restriction of f on ∂R is a bijection between ∂R and $\partial f(R)$.*

Proof. We first prove that $f(\partial R) \subseteq \partial f(R)$. Let ξ be a point in $\text{int } f(R)$, and N be a sufficiently small (open) neighborhood of ξ . The preimage $f^{-1}(N)$ consists of disjoint open sets U_1, \dots, U_k , where each U_i contains one single preimage x_i of ξ . If $x_i \notin R$ or $x_i \in \text{int } R$, then since N is sufficiently small, it can be shown that $\gamma_{U_i \cap R}$ is constant over N . For the sake of contradiction, suppose there exists at least one $x_i \in \partial R$. Then x_i is positive and in fact every point in U_i is positive, and f is a bijection from U_i to N , since N is sufficiently small. Let $H = f(U_i \setminus R) = N \setminus f(U_i \cap R)$, which is a nonempty open set. Then $\gamma_{U_i \cap R}$ is equal to 0 over H , and equal to +1 over $N \setminus H$. Let $\phi = \gamma_R - \gamma_{U_i \cap R}$ be a function from N to \mathbb{Z} . Because γ_R is constant over N , ϕ is equal to k over H and equal to $k - 1$ over $N \setminus H$, for some $k \in \mathbb{Z}$. On the other hand, ϕ is also the sum of a finite number of functions $\gamma_{U_j \cap R}$ (for $j \neq i$), each of which is either constant over N , or equal to 0 over H' and equal to +1 over $N \setminus H'$ for some nonempty open set H' . It can be shown that $\arg \min_{\xi \in N} \phi(\xi)$, which is exactly $N \setminus H$, must be a nonempty open set. This contradicts the fact that N is connected. Therefore, $f^{-1}(\xi) \cap R \subset \text{int } R$ for every $\xi \in \text{int } f(R)$, implying $f(\partial R) \subseteq \partial f(R)$.

By Lemma 2.1, for any point $\xi \in \partial f(R)$, every point x of the set $X = f^{-1}(\xi) \cap R$ is either contained in ∂R or belongs to a crease edge. Since $\gamma_R(\xi) = +1$, there is at least one such point x of the former type in X . Hence the restriction of the map f to ∂R is surjective. Furthermore, since X does not contain any negative point, there is only one positive point in X and every other point in X belongs to crease edges. Thus the considered restriction of f indeed constitutes a bijection between $f(\partial R)$ and $\partial f(R)$. ■

The following lemma suggests that it suffices to focus on canonical locally removable regions.

Lemma 3.2 *Let $R \subseteq S$ be a locally removable region, and let $\mathcal{J}_R \subseteq \mathbb{I}$ be the collection of inverted components contained in R . Then for any $I \in \mathcal{J}_R$, there exists a canonical locally removable region R_I so that $I \subseteq R_I \subseteq R$.*

Proof. By Lemma 2.7, $\hat{\phi}(R) = \mathcal{N}(R, f(R))$. However, Lemma 3.1 implies $\mathcal{N}(R, f(R)) = R$, because whenever a path $\Pi \subset S$ crosses ∂R (at a positive point), $f(\Pi)$ also crosses $\partial f(R)$. Hence $\hat{\phi}(R) = R$. Now since $\mathcal{J}_R \subseteq \hat{\phi}(R)$, we have by Lemma 2.3 (4) that $\hat{\phi}(\mathcal{J}_R) \subseteq \hat{\phi}(R) = R$.

We claim that, for any connected component U of $\hat{\phi}(\mathcal{J}_R)$:

- (a) ∂U does not intersect the interior of any inverted component or the interior of any crease edge;
- (b) $\gamma_U = +1$.

By Lemma 2.4 (1), ∂U does not intersect the interior of any inverted component in \mathcal{J}_R . Since $U \subseteq R$ and there is no other inverted component in R except those in \mathcal{J}_R , the first half of (a) then follows. Using this and Lemma 2.4 (2), it follows that ∂U does not intersect the interior of any crease edge either. Thus (a) is proved. For (b), by (a) and Lemma 2.5, we know that $\gamma_U = c_U$ for some constant $c_U > 0$. Then for any $\xi \in f(U)$,

$$1 = \gamma_R(\xi) = \gamma_{\hat{\phi}(\mathcal{J}_R)}(\xi) + \gamma_{R \setminus \hat{\phi}(\mathcal{J}_R)}(\xi) \geq \gamma_{\hat{\phi}(\mathcal{J}_R)}(\xi) \geq \gamma_U(\xi) > 0, \quad (1)$$

implying $c_U = 1$.

In fact, Eqn. (1) implies that shadows of the connected components of $\hat{\phi}(\mathcal{J}_R)$ are interiorly disjoint, because otherwise $\gamma_{\hat{\phi}(\mathcal{J}_R)}(\xi)$ would be strictly greater than $\gamma_U(\xi)$ for any ξ in the interior of the intersection of any two such shadows, leading to a contradiction in Eqn. (1). As such, by $f(\hat{\phi}(\mathcal{J}_R)) = f(\mathcal{J}_R)$, we then have $f(U) = f(\mathcal{J}_U)$ for any connected component U of $\hat{\phi}(\mathcal{J}_R)$, where $\mathcal{J}_U \subseteq \mathbb{I}$ is the subset of inverted components in U . (Note that $\mathcal{J}_U \subseteq \mathcal{J}_R$ because $U \subseteq R$.) Hence by Lemma 2.7, $U = \mathcal{N}(x, f(U)) = \mathcal{N}(x, f(\mathcal{J}_U))$ for any $x \in U$, and in particular, for any $x \in I \in \mathcal{J}_U$. Since $\hat{\phi}(\mathcal{J}_U) = \mathcal{N}(\mathcal{J}_U, f(\mathcal{J}_U))$ by Lemma 2.7, we have $U = \hat{\phi}(\mathcal{J}_U)$.

Let U be the connected component of $\hat{\phi}(\mathcal{J}_R)$ containing I . Setting $R_I = \hat{\phi}(\mathcal{J}_U)$ proves the lemma. ■

For any $R \subseteq S$, we write $R^* = \{T \in \mathcal{T} \mid R \cap \text{int} T \neq \emptyset\}$. If R is locally removable, we call R^* the *locally re-meshable set induced by R* .




Figure 4: A locally removable region R (delimited by dashed lines in the left figure) containing an inverted component (colored triangles) and its R^* (delimited by shaded boundaries).

Lemma 3.3 *Let $\tilde{\mathcal{T}}$ be a tangled triangulation of $\tilde{\mathcal{V}}$, let R be a locally removable region, and let $K = \tilde{\mathcal{T}} \setminus \{f(T) \mid T \in R^*\}$. Then there exists another tangled triangulation $\tilde{\mathcal{T}}'$ of $\tilde{\mathcal{V}}$, such that (i) $K \subseteq \tilde{\mathcal{T}}'$ and (ii) all triangles in $\tilde{\mathcal{T}}' \setminus K$ are upright.*

Intuitively, the lemma states that one can replace all triangles in a locally re-meshable set R^* with upright triangles while keeping triangles outside the region intact, thus removing all inverted components in R^* from the triangulation. If all triangles in K are already upright, then $\tilde{\mathcal{T}}'$ becomes a plane triangulation of $\tilde{\mathcal{V}}$.

Proof. Assume without loss of generality that ∂R consists of polygonal curves. We first construct a refined triangulation $\tilde{\mathcal{T}}_0$ and then coarsen it to the desired $\tilde{\mathcal{T}}'$. For each triangle $T \in \mathcal{T}$ that intersects ∂R , let $U_T = f(T \setminus R)$, and

$$A_T = \{\xi \in S \mid \xi = f(v) \text{ for some vertex } v \text{ of the region } T \setminus R\}.$$

The region U_T consists of a set of simple polygons with vertices from A_T . Hence we can triangulate U_T on the vertex set A_T ; let U_T^∇ be the set of triangles in the resulting triangulation. Set $U^\nabla = \bigcup U_T^\nabla$, where the union is taken on all triangles $T \in \mathcal{T}$ that intersect ∂R . We further triangulate within $f(R)$ on the vertex set

$$A_R = f(\mathcal{V} \cap R) \cup \{\xi \in S \mid \xi = f(\partial R \cap \text{int } e) \text{ for some edge } e \in \mathcal{E}\};$$

let R^∇ denote the set of triangles in the resulting triangulation. We define $\tilde{\mathcal{T}}_0 = U^\nabla \cup R^\nabla \cup K$. It can be shown by Lemma 3.1 that $\tilde{\mathcal{T}}_0$ forms a tangled triangulation of $\tilde{\mathcal{V}} \cup \mathcal{V}'$, where \mathcal{V}' represents the collection of new vertices introduced in the above process other than those in $\tilde{\mathcal{V}}$. Furthermore, using the fact that ∂R consists of positive points, it can be shown that all triangles in $\tilde{\mathcal{T}}_0 \setminus K$ are upright. See Figure 5 for an illustration.

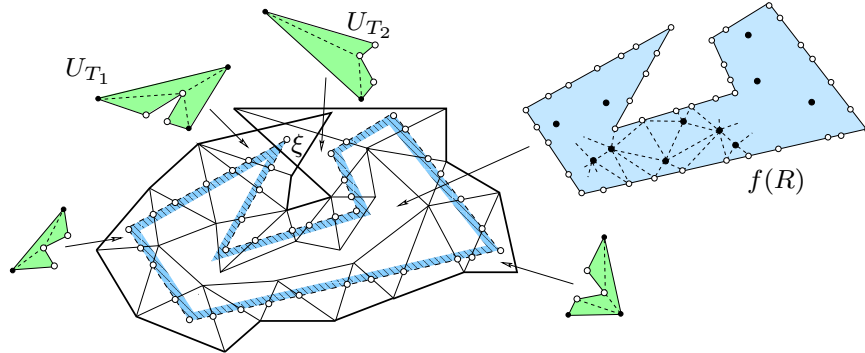


Figure 5: Retriangulating a locally re-meshable set induced by a locally removable region R . The shaded region in the center is $f(R)$. The empty circles represent the set \mathcal{V}' of newly introduced vertices other than those in $\tilde{\mathcal{V}}$. Note that for the vertex ξ , $\xi \in A_{T_1}$, but $\xi \notin A_{T_2}$ although $\xi \in f(T_2)$.

Pick a vertex $\xi \in \mathcal{V}'$ of $\tilde{\mathcal{T}}_0$. Observe that a sufficiently small neighborhood of ξ is completely contained in the upright components of the tangled triangulation $\tilde{\mathcal{T}}_0$, and the star of ξ (the union of the triangles incident upon ξ) is a star-shaped polygon P with ξ in the interior of P . We can re-triangulate P by upright triangles, using only vertices on the boundary of P . We now obtain a tangled triangulation $\tilde{\mathcal{T}}_1$ of $\tilde{\mathcal{V}} \cup (\mathcal{V}' \setminus \{\xi\})$ such that all triangles in $\tilde{\mathcal{T}}_1 \setminus K$ are upright. By repeating the above process for each vertex in \mathcal{V}' , we obtain a tangled triangulation $\tilde{\mathcal{T}}'$ of $\tilde{\mathcal{V}}$ with desired properties. ■

4 Local Untangling

In this section we present an algorithm for computing a locally removable region R that contains a given “seed” inverted component $I \in \mathbb{I}$. By Lemma 3.3, we can remove I from the triangulation

by re-meshing R^* . The region R is minimum in the sense that, for any $R' \subseteq S$ that is locally removable and contains I , $R \subseteq R'$. We first present the algorithm at a high level, followed by a more elaborate description of the algorithm.

Ultimately, the algorithm computes a collection $\mathcal{J} \subseteq \mathbb{I}$ of inverted components, with $I \in \mathcal{J}$, such that $\hat{\phi}(\mathcal{J})$ is locally removable. Let X be the shadow of \mathcal{J} . At intermediate stages of the algorithm, an incomplete set \mathcal{J} and an incomplete shadow X (not necessarily the shadow of the current \mathcal{J}) are maintained. Initially, \mathcal{J} is set to $\{I\}$ and X is set to $f(I)$. At each step, the algorithm either expands the set \mathcal{J} or the shadow X . If neither \mathcal{J} nor X is expanded, the algorithm declares that the desired \mathcal{J} is found and terminates.

Extending inverted components. For a set $\mathcal{J} \subseteq \mathbb{I}$ of inverted components, let $E(\mathcal{J})$ be the set of all inverted components whose interiors intersect $\hat{\phi}(\mathcal{J})$, i.e.,

$$E(\mathcal{J}) = \left\{ I \in \mathbb{I} \mid \hat{\phi}(\mathcal{J}) \cap \text{int } I \neq \emptyset \right\}.$$

We call $E(\mathcal{J})$ the *extension* of \mathcal{J} . See Figure 6 (i). Clearly, $\mathcal{J} \subseteq E(\mathcal{J})$.

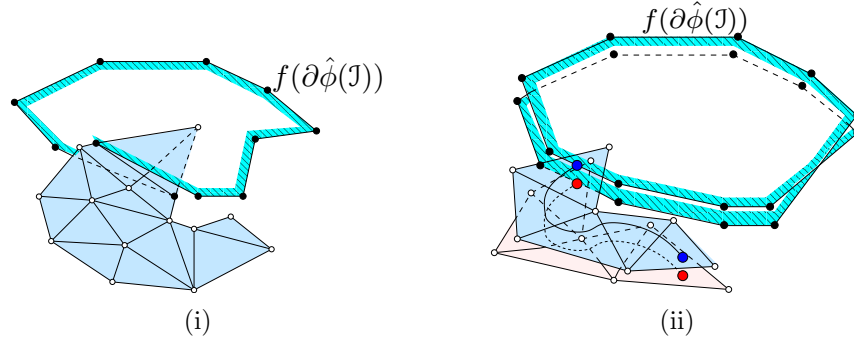


Figure 6: (i) Extending inverted components. (ii) Growing shadows.

Growing shadows. Consider a collection $\mathcal{J} \subseteq \mathbb{I}$ of inverted components such that $\hat{\phi}(\mathcal{J})$ is connected and $\partial\hat{\phi}(\mathcal{J})$ does not intersect the interior of any inverted component. Clearly $\gamma_{\hat{\phi}(\mathcal{J})} > 0$. If $\gamma_{\hat{\phi}(\mathcal{J})} = +1$, then by definition $\hat{\phi}(\mathcal{J})$ is locally removable. However, if $\gamma_{\hat{\phi}(\mathcal{J})} \geq +2$ (which may indeed occur by Lemma 2.6), the extension of \mathcal{J} will not help to extend $\hat{\phi}(\mathcal{J})$ further. In this case, instead of extending the set \mathcal{J} directly, we expand the shadow of \mathcal{J} , as follows. For a path $\Pi \subset S$, its preimage $f^{-1}(\Pi)$ generally consists of a number of connected components. For a set $R \subseteq S$, let $L(\Pi, R)$ denote the subset of connected components of $f^{-1}(\Pi)$ that intersect R . We define the *extended shadow* of \mathcal{J} , denoted by $G(\mathcal{J})$, as $f(\mathcal{J})$ plus the collection of points $\xi \in S \setminus f(\mathcal{J})$ for which there exist a point $\eta \in \partial f(\mathcal{J})$ and a path Π from ξ to η such that at least two connected components in $L(\Pi, \hat{\phi}(\mathcal{J}))$ do not intersect the interior of any inverted component. See Figure 6 (ii).

Algorithm. Recall that $I \in \mathbb{I}$ is the given “seed” inverted component. The following algorithm computes a locally removable region R containing I . As mentioned earlier, we make use of \mathcal{J} and f and other relevant notations only for the clarify of the description of the algorithm; \mathcal{T} and f are not assumed to be given in association with the input tangled triangulation $\tilde{\mathcal{T}}$.

1:	$\mathcal{J} \leftarrow \{I\}, X \leftarrow f(I)$
2:	$R \leftarrow \mathcal{N}(I, X)$
3:	$E(\mathcal{J}) \leftarrow \{I \in \mathbb{I} \mid R \cap \text{int } I \neq \emptyset\}$
4:	if $E(\mathcal{J}) \neq \mathcal{J}$ $\mathcal{J} \leftarrow E(\mathcal{J}); X \leftarrow f(\mathcal{J})$ goto step 2
5:	else if $\gamma_R \neq +1$ $X \leftarrow G(\mathcal{J})$ goto step 2
6:	else return R

At the end of the algorithm, by Lemma 2.7, we have $R = \hat{\phi}(\mathcal{J})$ and $f(R) = X$. Next we show that R is indeed a minimum locally removable region containing the “seed” inverted component.

Lemma 4.1 *Let $\mathcal{J} \subseteq \mathbb{I}$ be a collection of inverted components. For any set $\mathcal{J}' \subseteq \mathbb{I}$ for which $\hat{\phi}(\mathcal{J}')$ is locally removable and $\mathcal{J} \subseteq \hat{\phi}(\mathcal{J}')$, $\hat{\phi}(E(\mathcal{J})) \subseteq \hat{\phi}(\mathcal{J}')$.*

Proof. Since $\mathcal{J} \subseteq \hat{\phi}(\mathcal{J}')$, it follows from Lemma 2.3 (4) that $\hat{\phi}(\mathcal{J}) \subseteq \hat{\phi}(\mathcal{J}')$. For any inverted component $I \in E(\mathcal{J}) \setminus \mathcal{J}$, I intersects $\hat{\phi}(\mathcal{J})$ and therefore also intersects $\hat{\phi}(\mathcal{J}')$. But by assumption, $\hat{\phi}(\mathcal{J}')$ is locally removable and therefore $\partial\hat{\phi}(\mathcal{J}')$ does not intersect the interior of any inverted component. As such, $I \subseteq \hat{\phi}(\mathcal{J}')$. Hence, $E(\mathcal{J}) \subseteq \hat{\phi}(\mathcal{J}')$. By Lemma 2.3 (4) again, we then have $\hat{\phi}(E(\mathcal{J})) \subseteq \hat{\phi}(\mathcal{J}')$. ■

Lemma 4.2 *Let $\mathcal{J} \subseteq \mathbb{I}$ be a collection of inverted components such that $\hat{\phi}(\mathcal{J})$ is connected and $\gamma_{\hat{\phi}(\mathcal{J})} \geq +2$. For any set $\mathcal{J}' \subseteq \mathbb{I}$ for which $\hat{\phi}(\mathcal{J}')$ is locally removable and $\mathcal{J} \subseteq \hat{\phi}(\mathcal{J}')$, $\mathcal{N}(\mathcal{J}, G(\mathcal{J})) \subseteq \hat{\phi}(\mathcal{J}')$.*

Proof. Clearly $f(\mathcal{J}) \subseteq f(\mathcal{J}')$. Fix an arbitrary point $\xi \in G(\mathcal{J}) \setminus f(\mathcal{J})$; we next show that $\xi \in f(\mathcal{J}')$. Let η and $\Pi[0, 1]$ be the point and the path that together serve as the witness for $\xi \in G(\mathcal{J})$ in the definition of $G(\mathcal{J})$. We claim that $\Pi \subseteq f(\mathcal{J}')$; since $\xi = \Pi(1)$, this claim would immediately imply $\xi \in f(\mathcal{J}')$ and prove the lemma.

Now suppose for the sake of contradiction that Π does not lie completely in $f(\mathcal{J}')$. Observe that $\Pi(0) = \eta \in f(\mathcal{J}')$, and as such, $\Pi \cap \partial f(\mathcal{J}') \neq \emptyset$. Let $t \in [0, 1]$ be the smallest value for which $\Pi(t) \in \partial f(\mathcal{J}')$. Therefore, $\Pi[0, t] \subseteq f(\mathcal{J}')$. It then follows from Lemma 2.7 that $f^{-1}(\Pi(t)) \cap L(\Pi, \hat{\phi}(\mathcal{J})) \subset \hat{\phi}(\mathcal{J}')$. Consider the connected components in $L(\Pi, \hat{\phi}(\mathcal{J}))$ that do intersect the interior of any inverted component. Each such component contains a preimage of $\Pi(t)$ whose sign function evaluates to $+1$. On the other hand, by Lemma 2.3 (3) and the fact that $\partial\hat{\phi}(\mathcal{J}')$ does not intersect the interior of any inverted component, $f^{-1}(\Pi(t)) \cap \hat{\phi}(\mathcal{J}')$ does not contain any negative point. Therefore $\gamma_{\hat{\phi}(\mathcal{J}')}(\Pi(t))$ is at least $+2$, contradicting the assumption that $\hat{\phi}(\mathcal{J}')$ is locally removable.

Therefore $G(\mathcal{J}) \subseteq f(\mathcal{J}')$. It follows that $\mathcal{N}(\mathcal{J}, G(\mathcal{J})) \subseteq \hat{\phi}(\mathcal{J}')$. ■

Lemma 4.3 *At the end of the algorithm, R is locally removable. Furthermore, for any locally removable region R' containing I , $R \subseteq R'$.*

Proof. When the algorithm stops, step 4 ensures $E(\mathcal{J}) = \mathcal{J}$. Recall that $R = \hat{\phi}(\mathcal{J})$. Suppose for the sake of contradiction ∂R intersects the interior of some inverted component $I \in \mathbb{I}$. By Lemma 2.4 (1), ∂R does not intersect the interior of any of the inverted components in \mathcal{J} . As such, I has to be an inverted component not included in \mathcal{J} . By definition of $E(\mathcal{J})$, $I \in E(\mathcal{J}) \setminus \mathcal{J}$, contradicting $E(\mathcal{J}) = \mathcal{J}$. Therefore, ∂R does not intersect the interior of any inverted component. Since step 5 further ensures $\gamma_R = +1$, R is locally removable. The second half of the lemma follows directly from Lemmas 3.2, 4.1 and 4.2. ■

Implementation. There are many ways to implement the above algorithm in an output-sensitive manner. We sketch an implementation that, though by no means the most efficient, is conceptually simple and illustrates that the algorithm never explores beyond the output R^* .

Data structure. As mentioned above, the algorithm maintains a shadow X . Instead of maintaining $R = \mathcal{N}(I, X)$, we maintain the set R^* (recall that $R^* = \{T \in \mathcal{T} \mid R \cap \text{int} T \neq \emptyset\}$); when the algorithm stops, R^* coincides with the locally re-meshable set induced by R . We identify the “boundary” triangles of R^* , i.e., the other triangle adjacent to one of its edges is not in R^* . We also mark a subset of inverted triangles in R^* as “active”; they correspond to those that were added to R^* in the last iteration of the algorithm.

We represent X as the union of a family \mathcal{X} of planar polygons possibly with holes such that (i) $X = \bigcup_{P \in \mathcal{X}} P$ and (ii) each edge of every $P \in \mathcal{X}$ is a portion of the shadow of a crease edge in \mathcal{E} whose endpoints are the vertices of the arrangement of the shadows of crease edges in \mathcal{E} . We preprocess each $P \in \mathcal{X}$ into a data structure $\mathcal{D}(P)$ so that we can quickly determine whether a query segment ℓ intersects P . Theoretically, such a query can be answered in logarithmic time using a linear-size data structure, but in real applications we may want to use a more practical data structure. Let $\tau(|P|)$ denote the query time of $\mathcal{D}(P)$, and assume $\mathcal{D}(P)$ can be constructed in $O(|P| \log |P|)$ time.

Updating R^ .* Whenever X changes, we need to update R^* . Starting from “boundary” triangles in R^* , we perform a depth-first search on the dual graph of the triangles in \mathcal{T} to search for triangles whose shadow intersects X . At each step we move from one triangle T_1 to another triangle T_2 in the dual graph if and only if $T_1 \in R^*$, $T_2 \notin R^*$, and the shadow of the edge e shared by T_1 and T_2 intersects X , in which case we then add T_2 into R^* . To determine whether $f(e)$ intersects X , for each $P \in \mathcal{X}$, we test whether $f(e)$ intersects P using $\mathcal{D}(P)$. If the answer is yes for any of them, we conclude that $f(e)$ intersects X . In fact, we only query $f(e)$ against those $\mathcal{D}(P)$ ’s that have not been queried by $f(e)$ before. During the expansion of R^* , we mark newly added inverted triangles as “active” and also update the set of “boundary” triangles.

Extending inverted components. This step is executed if R^* has active inverted triangles. For each active triangle T , we compute the inverted component $I \in \mathbb{I}$ containing T , add I to \mathcal{J} and all triangles in I to R^* . X now becomes $X \cup f(I)$. We set $P = f(I)$, construct $\mathcal{D}(P)$ and add P to \mathcal{X} . We construct $f(I)$ as follows. ∂I consists of a set $C_I \subseteq \mathcal{E}$ of crease edges. Let $\Gamma = \{f(e) \mid e \in C_I\}$ be the set of shadows of these edges. We construct the arrangement $\mathcal{A}(\Gamma)$. It can be checked that $\gamma_I(\xi)$ is the same for all points ξ in a face z of $\mathcal{A}(\Gamma)$, which we denote by c_z . Moreover $z \subseteq f(I)$ if and only if $c_z < 0$, and $c_z = 0$ otherwise. We can compute c_z for each face z of $\mathcal{A}(\Gamma)$ by doing a traversal of $\mathcal{A}(\Gamma)$. By identifying the faces z for $c_z < 0$, we can compute $f(I)$. After processing I , we mark all triangles of I inactive in R^* . After processing all active triangles in R^* , we update R^* as described above.

Growing shadows. This step is executed if R^* contains no active triangles. We compute the extended shadow $G(\mathcal{J})$ as follows. Define a graph \mathcal{G} in which, each vertex represents a pair $\{T_1, T_2\}$ of upright triangles in \mathcal{T} with $f(T_1) \cap f(T_2) \neq \emptyset$, and each edge connects two pairs $\{T_1, T_2\}$ and $\{T_1, T_3\}$ if T_2, T_3 share a common edge e with $f(e) \cap f(T_1) \neq \emptyset$. We first merge all polygons in \mathcal{X} to compute $X = \bigcup_{P \in \mathcal{X}} P$ using the same procedure that we used to compute $f(I)$ and set $\mathcal{X} = \{X\}$. For each connected component of ∂X , We choose an arbitrary point from it, find the subset of triangles in R^* whose shadows contain the point (note that these triangles are necessarily upright), and form all pairs of these triangles. Starting from vertices corresponding to these pairs, we perform a depth-first search in \mathcal{G} to identify the connected components containing the starting vertices. Then the extended shadow $G(\mathcal{J})$ is the union of the current shadow X and the intersection polygon

$\Delta = f(T_1) \cap f(T_2)$ for each visited pair $\{T_1, T_2\}$. Let $\Delta_1, \dots, \Delta_u$ be the intersection polygons computed by the depth-first search. We compute $P = X \cup \bigcup_{i=1}^u \Delta_i$, which is $G(\mathcal{J})$, similarly to the way we compute $f(I)$, and set $\mathcal{X} = \{P\}$. It can be shown that \mathcal{X} satisfies the two conditions mentioned above. Finally, we update R^* as described above.

Running time. Let n_R denote the number of triangles in the final output R^* . The total time spent in constructing the polygons that ever appear in \mathcal{X} is bounded by $O(k_R^c \log n_R)$, where k_R^c is the complexity of the arrangement of the shadows of all the crease edges in the output R^* . This follows from property (ii) of \mathcal{X} and the observation that once a polygon has been merged in the computation of the extended shadow, it no longer shows up on ∂X . During the expansion of R^* , an upright triangle may be queried against $O(\delta_R)$ data structures in total, where δ_R denotes the number of inverted components in the output R^* . We do not perform these queries for inverted triangles encountered during the search as they are directly added to R^* . Hence the total time for updating R^* is $O(n_R^- + \delta_R n_R^+ \tau(n_R))$, where n_R^+ (resp. n_R^-) is the number of upright (resp. inverted) triangles in R^* . Finally, when computing extended shadows, for each visited pair of triangles, if the shadows of their boundaries intersect, we charge the cost of visiting the pair to the intersection; otherwise we charge it to the triangle whose shadow is entirely contained in the other. For the latter case, the total number of such charges to a triangle is at most $(d_R + 1)/2$, where d_R is the maximum number of preimages a point in X can have in R . As such, the total time for computing extended shadows can be bounded by $O((k_R^+ + d_R n_R^+) \log n_R)$, where k_R^+ is the complexity of the arrangement of the shadows of all the upright triangles in R^* . Overall, the total running time is $O(n_R^- + \delta_R n_R^+ \tau(n_R) + (k_R^+ + d_R n_R^+) \log n_R)$; the term k_R^c is absorbed by k_R^+ as clearly $k_R^c \leq k_R^+$. If $\tau(n_R) = O(\log n_R)$, it becomes $O(n_R^- + (n_R^+ \cdot \max\{d_R, \delta_R\} + k_R^+) \log n_R)$. The total space requirement is $O(n_R^- + n_R^+ \cdot \max\{d_R, \delta_R\} + k_R^+)$.

Using Lemma 4.3 and the above algorithm, we obtain the main result of this section.

Theorem 4.4 *Given any inverted triangle $T \in \mathcal{T}$, the algorithm outputs a locally re-meshable set R^* containing T , induced by a locally removable region R . Furthermore, for any $R' \subseteq S$ that is locally removable and contains T , $R \subseteq R'$. The algorithm never explores beyond the output R^* .*

5 Global Untangling

Using Lemma 2.2, it is not hard to prove that each connected component of $\phi(\mathbb{I})$ is locally removable. Therefore if we wish to untangle the entire mesh, i.e., to remove all inverted components in \mathbb{I} , we can simply compute $R = \phi(\mathbb{I})$, remove all triangles in $R^* = \{T \in \mathcal{T} \mid R \cap \text{int } T \neq \emptyset\}$, and re-mesh R^* with upright triangles. The following theorem provides a characterization of $\phi(\mathbb{I})$ in terms of the primary conflict sets of a collection of subsets of \mathbb{I} .

Theorem 5.1 *Let $\mathcal{C} = \{\mathcal{J}_1, \dots, \mathcal{J}_m\}$ be a collection of subsets of \mathbb{I} so that for each \mathcal{J}_i , $\partial \hat{\phi}(\mathcal{J}_i)$ does not intersect the interior of any inverted component, and $\bigcup_{i=1}^m \mathcal{J}_i = \mathbb{I}$. Then $\phi(\mathbb{I}) = \bigcup_{i=1}^m \hat{\phi}(\mathcal{J}_i)$.*

In particular, the theorem suggests that the algorithm presented in the previous section can be used to compute $\phi(\mathbb{I})$, and furthermore, step 5 can be skipped altogether, as its only purpose there is to ensure the condition $\gamma_R = +1$ which is no longer needed for computing $\phi(\mathbb{I})$ by the theorem.

Proof. For every $i = 1, \dots, m$, $\hat{\phi}(\mathcal{J}_i) \subseteq \phi(\mathbb{I})$. We thus only need to show that $\phi(\mathbb{I}) \subseteq \bigcup_{i=1}^m \hat{\phi}(\mathcal{J}_i)$. Since every negative point in $\phi(\mathbb{I})$ is contained in some inverted component I and therefore in $\hat{\phi}(\mathcal{J})$ for some $\mathcal{J} \in \mathcal{C}$, we only need to show for every positive point $x \in f^{-1}(\mathbb{I})$ that $x \in \hat{\phi}(\mathcal{J})$ for some $\mathcal{J} \in \mathcal{C}$.

The proof works by induction on the number of inverted components under the map f . If I is the only inverted component, then the theorem is true by the following argument. Take a point $\xi \in \partial f(I)$ and consider the set $f^{-1}(\xi)$. By the choice of ξ , no point in $f^{-1}(\xi)$ is negative. Since $\gamma_S(\xi) = +1$, the sum of the signs of the points in $f^{-1}(\xi)$ is $+1$. Therefore there is precisely one point $x \in f^{-1}(\xi)$ with $\chi(x) = +1$ and for every other point $y \in f^{-1}(\xi)$, $\chi(y) = 0$. Since I is the only inverted component, $\partial \hat{\phi}(I)$ does not intersect the interior of any inverted component and therefore $\gamma_{\hat{\phi}(I)} > 0$ by Lemma 2.5. Thus the sum of the signs of the points in $f^{-1}(\xi) \cap \hat{\phi}(I)$ is at least $+1$, which implies that $x \in \hat{\phi}(I)$. Therefore, $\gamma_{\hat{\phi}(I)}(\xi) = +1$, implying $\gamma_{\hat{\phi}(I)} = +1$ by Lemma 2.5. Now, take any arbitrary point $\xi' \in f(I)$ and let $x' \in f^{-1}(\xi')$ be positive, i.e., $\chi(x') = +1$. We claim that $x' \in \hat{\phi}(I)$. This is because otherwise, since $\gamma_{\hat{\phi}(I)}(\xi') = +1$, the sum of the signs of the points in $f^{-1}(\xi') \cap \hat{\phi}(I)$ would be $+1$, and given that there are no other negative points in $f^{-1}(\xi')$ than those already in I , the sum of the signs of the points in $f^{-1}(\xi')$, i.e., $\gamma_S(\xi')$, would be at least $+2$ when $\chi(x')$ is accounted for — a contradiction.

Now suppose the statement of the theorem is true for any map with $m - 1$ inverted components and let f be a map with m inverted components $\mathbb{I} = \{I_1, \dots, I_m\}$. Note that $f(\mathbb{I}) = \bigcup_i f(I_i) = \bigcup_i f(\hat{\phi}(J_i))$. Take an arbitrary point $\xi \in \partial f(\mathbb{I})$ and consider the set $f^{-1}(\xi)$. An argument similar to the above shows that for a unique point $x \in f^{-1}(\xi)$, $\chi(x) = +1$, and for every other $y \in f^{-1}(\xi)$, $\chi(y) = 0$. Now, take an arbitrary inverted component $I \in \mathbb{I}$ for which $\xi \in \partial f(I)$, and a set $J_i \in \mathcal{C}$ for which $I \in J_i$. Again, a similar argument as in the previous paragraph shows that $\gamma_{\hat{\phi}(J_i)} = +1$. Let $\xi' \in f(J_i)$ be a point that is not in $f(J_j)$ for any $j \neq i$. We can argue similarly to the previous paragraph that every positive point $x' \in f^{-1}(\xi')$ is in $\hat{\phi}(J_i)$. Therefore, $f^{-1}(\xi') \subset \hat{\phi}(J_i)$.

On the other hand, since $\gamma_{\hat{\phi}(J_i)} = +1$, J_i is locally removable. By Lemma 3.1, $\partial \hat{\phi}(J_i)$ is a simple plane polygon possibly with holes whose boundary is mapped identically by f into $\partial f(J_i)$. We can extend the restriction $f : \partial \hat{\phi}(J_i) \rightarrow \partial f(J_i)$ to a bijection $g : \hat{\phi}(J_i) \rightarrow f(J_i)$. Define a new map $f_0 : S \rightarrow S$ by letting

$$f_0(x) = \begin{cases} f(x) & x \notin \hat{\phi}(J_i), \\ g(x) & x \in \hat{\phi}(J_i). \end{cases}$$

Let \mathbb{I}_0 be the set of inverted components of f_0 . For any inverted component $I \in \mathbb{I}_0$, let $\phi_0(I)$ and $\hat{\phi}_0(I)$ respectively be the counterparts of $\phi(I)$ and $\hat{\phi}(I)$ only with respect to f_0 instead of f .

Observe that since $\partial \hat{\phi}(J_i)$ does not intersect the interior of any inverted component of \mathbb{I} , the inverted components in \mathbb{I}_0 are precisely the inverted components in \mathbb{I} minus those (fully) contained in $\hat{\phi}(J_i)$. Thus f_0 has fewer inverted components than f and therefore by induction, every positive point in $\phi_0(\mathbb{I}_0)$ is contained in $\hat{\phi}_0(J)$ for some $J \in \mathbb{I}$. Observe that for any $J' \in \mathbb{I}_0$ and its counterpart $J \in \mathbb{I}$, $\hat{\phi}_0(J') \subseteq \hat{\phi}(J)$. Thus every positive point in $\phi_0(\mathbb{I}_0)$ is contained in $\hat{\phi}(J_i)$ for some $j \neq i$. On the other hand, we showed above that positive points in $\phi(\mathbb{I}) \setminus \phi_0(\mathbb{I}_0)$ are contained in $\hat{\phi}(J_i)$. This proves the theorem. \blacksquare

6 Conclusions

Since our untangling algorithm first identifies tangled regions of the mesh and then simply remeshes these regions, it does not involve any edge-flip operation. The algorithm of Shewchuk and Wallace [16] relies on a set of other local operations, in addition to edge flips, for untangling a mesh. It is legitimate to ask whether mesh untangling can also be accomplished by a pure edge-flip algorithm. As mentioned before, such a sequence of edge flips always exists, but it is not clear whether it can be found under the guide of a few simple rules based on local geometric information in the tangled mesh, in a way similar to Lawson's algorithm.

The locally re-meshable set computed by our algorithm is by no means the smallest subset of triangles that can be removed and replaced with upright triangles to restore the validity of the mesh. A relatively efficient algorithm for computing the smallest such set would be very interesting, although practically it may not be worthwhile to do so. In practice, some simple heuristics may be combined with our algorithm to reduce the number of triangles to be removed for re-meshing; for example, large upright components in the explored region should be retained as much as possible to avoid retriangulation.

The running time of our algorithm depends on the complexity of the overlay of the explored region. This seems to be inherent in steps for maintaining shadows and for computing extended shadows. If one can find a way to get around these steps, then an algorithm with running time (nearly)-linear in the number of output triangles might be possible.

Finally, we do not foresee any difficulty, at least conceptually, in extending the untangling algorithm to three and higher dimensions. However, it is somewhat unclear how to formulate the question in the context of surface triangulations and to extend the untangling algorithm there. We leave all these open questions for future research.

Acknowledgments. We thank Jonathan Shewchuk for sending us a copy of his manuscript. We also thank Tamal Dey, Leo Guibas, Alper Üngör, and the anonymous referees for many useful comments.

References

- [1] P. K. Agarwal, J. Basch, L. J. Guibas, J. Hershberger, and L. Zhang, Deformable free space tiling for kinetic collision detection, *Intl. J. Robotics Research*, 21 (2002), 179–197.
- [2] P. K. Agarwal, Y. Wang, and H. Yu, A two-dimensional kinetic triangulation with near-quadratic topological changes, *Discrete Comput. Geom.*, 36 (2006), 573–592.
- [3] G. Albers, L. J. Guibas, J. Mitchell, and T. Roos, Voronoi diagrams of moving points, *Internat. J. Comput. Geom. Appl.*, 8 (1998), 365–380.
- [4] N. Amenta, M. Bern, and D. Eppstein, Optimal point placement for mesh smoothing, *J. Algorithms*, 30(2) (1999), 302–322.
- [5] J. Basch, L. J. Guibas, and J. Hershberger, Data structures for mobile data, *J. Algorithms*, 31(1) (1999), 1–28.
- [6] H.-L. Cheng, T. K. Dey, H. Edelsbrunner, and J. Sullivan, Dynamic skin triangulation, *Discrete Comput. Geom.*, 25(4) (2001), 525–568.
- [7] S.-W. Cheng and T. K. Dey, Maintaining deforming surface meshes, *Proc. 19th ACM-SIAM Sympos. Discrete Algorithms*, 2008, pp. 112–121.
- [8] H. Edelsbrunner and A. Üngör, Relaxed scheduling in dynamic skin triangulation, *Proc. Japan Conf. Discrete Comput. Geom.*, 2002, pp. 135–151.
- [9] J. Erickson, D. Guoy, J. Sullivan, and A. Üngör, Building space-time meshes over arbitrary spatial domains, *Proc. 11th Intl. Meshing Roundtable*, 2002, pp. 391–402.
- [10] L. Freitag, On combining laplacian and optimization-based mesh smoothing techniques, *Proc. Sympos. Trends in Unstructured Mesh Generation*, Vol. AMD-Vol 220, Amer. Soc. Mechanical Engineers, 1997, pp. 37–44.
- [11] L. Freitag and P. Plassmann, Local optimization-based simplicial mesh untangling and improvement, *Intl. J. Numer. Meth. Engin.*, 49 (2000), 109–125.

- [12] P. Knupp, Hexahedral and tetrahedral mesh untangling, *Engineering with Computers*, 17 (2001), 261–268.
- [13] P. Knupp, L. Margolin, and M. Shashkov, Reference Jacobian optimization-based rezone strategies for Arbitrary Lagrangian Eulerian methods, *J. Comput. Phys.*, 176 (2002), 93–128.
- [14] K. Kovalev, M. Delanaye, and C. Hirsch, Untangling and optimization of unstructured hexahedral meshes, *Journal of Computational Mathematics and Mathematical Physics*, 43(6) (2003), 847–855.
- [15] C. Lawson, Software for C^1 surface interpolation, in: *Mathematical Software* (J. R. Rice, ed.), Vol. III, Academic Press, NY, 1977, pp. 161–194.
- [16] J. Shewchuk and D. T. Wallace. Untangling triangulations, 2007. unpublished manuscript.
- [17] A. Üngör and A. Sheffer, Pitching tents in space-time: Mesh generating for discontinuous Galerkin method, *Proc. 9th Intl. Meshing Roundtable*, 2000, pp. 111–122.
- [18] P. Vachal, R. Garimella, and M. Shashkov, Untangling of 2D meshes in ALE simulations, *J. Comput. Phys.*, 196(2) (2004), 627–644.