

Untangling Triangulations through Local Explorations*

Pankaj K. Agarwal

Bardia Sadri

Hai Yu

January 2, 2010

Abstract

In many applications it is often desirable to maintain a valid mesh within a certain domain that deforms over time. During a period for which the underlying mesh topology remains unchanged, the deformation moves the vertices of the mesh and thus potentially turns a mesh invalid, or as we call it, *tangled*. We introduce the notion of *locally removable region*, which is a certain tangled area in the mesh that allows for local removal and re-meshing. We present an algorithm that is able to quickly compute, through local explorations, a minimal locally removable region containing a seed tangled region in the 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. Our algorithm exploits several novel insights into the structure of the tangled mesh, which may be of independent interest in contexts beyond mesh untangling.

1 Introduction

In many applications such as physical simulation, scientific computing, and computer graphics, it is often required to maintain a valid mesh (triangulation) within a certain domain that deforms over time. This problem has attracted much attention from computational geometry in recent years, and a number of techniques have been proposed. For example, the spacetime meshing method [5, 11] builds a mesh over the entire spacetime domain whose resolution is adaptive to the movement of the underlying space; numerical simulation can then be carried out directly on the spacetime mesh. The drawback of this method is that one needs to mesh within a domain that is one-dimensional 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 domain of influence and of dependence of a point in spacetime.

Another popular 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, 3, 2] uses an event-driven framework (i.e., the kinetic data structure framework by Basch *et al.* [4]) to proactively detect when the mesh will become invalid, and repair it immediately when it occurs, thus maintaining a valid mesh all the time. This method usually requires accurate knowledge on the deformation of the domain to be able to predict critical events, and a significant amount of extra storage to keep track of these events.

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

*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.

two consecutive time steps, the mesh can be either valid or invalid. Since many numerical algorithms also discretizes the time domain, and computation is done at these discrete time instances, this method is especially suitable for these algorithms. Between two time steps, the deformation moves the vertices of the mesh while the underlying mesh topology remains unchanged, thus potentially turning a mesh invalid, or as we call it, *tangled*, because the elements of mesh may intersect inadmissibly. When this happens, an “untangling” process needs to be invoked at the next time step so as 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 are moved. The moved vertices and the inverted triangles have been marked in both figures.

Problem statement. Specifically in this paper, spurred by an abundance of work on maintaining triangulations in a deforming planar domain, we study local untangling of planar triangulations. 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 piece-wise linear map that extends a “motion map” defined on the vertex set of a plane triangulation. This is formalized as follows.

Let \mathcal{T} and \mathcal{E} respectively be the set of triangles and edges of a triangulation of the convex hull of a set of points (*vertices*) $V \subset \mathbb{R}^2$ in general position. We will use the shorthand S for $\text{conv } V = \bigcup_{T \in \mathcal{T}} T$. Let $f_V : V \rightarrow \mathbb{R}^2$ be a *motion* map that represents the new position of each vertex after a motion period. The map f_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} . It is folklore that f defined this way is continuous. It is also linear restricted to the closure of every edge or triangle. For the rest of this paper we fix the triangulation $(V, \mathcal{E}, \mathcal{T})$ and the map f and define subsequent notions relative to them.

For simplicity, throughout the paper we make the following boundary assumption: (i) f is identity on ∂S , and (ii) $f(\text{int } S) \subseteq \text{int } S$. In other words, the map $f : S \rightarrow S$ is surjective, identity on ∂S and moves the interior of S without taking any point outside its boundary. The assumption can be enforced on any mesh by introducing three dummy vertices at infinity into the mesh.

The orientation of a triangle $T \in \mathcal{T}$ with vertices v_1, v_2 , and v_3 under the map f is defined as the sign of the cyclic permutation $[f(v_1), f(v_2), f(v_3)]$ relative to $[v_1, v_2, v_3]$, taking the latter permutation as positive. For any triangle $T \in \mathcal{T}$, we label T as *inverted* if T is oriented negatively under f , and *upright* otherwise. In other words, if v_1, v_2 , and v_3 are the vertices of T in clockwise order, T is labeled upright or inverted depending on whether $f(v_1), f(v_2)$, and $f(v_3)$ appear around

$f(T)$ in clockwise order as well, or not.

Given a tangled triangulation, the goal of untangling is to remove inverted triangles (unavoidably together with a few other upright triangles) and replace them with upright triangles so as to restore the validity of the triangulation. To provide more flexibility and compatibility for different applications, we mainly consider the following form of the problem: given a seed inverted triangle $T \in \mathcal{T}$, compute a region in the tangled mesh that contains T and can be locally removed and re-meshed with upright triangles.

Related work. Edge flip is probably the most elegant atomic operation for transforming triangulations, whose power is well illustrated in Lawson’s celebrated algorithm [9] 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 untangling triangulation can also be accomplished 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 at all whether this sequence can be found by applying a number of simple guiding rules based on local geometric information, as in Lawson’s algorithm.

Shewchuk and Wallace [10] study the problem of untangling triangulations by performing local surgeries to the troubled region of the mesh. The primary local operations used by their algorithm is edge flips, but occasionally their algorithm may perform other more complicated local operations. As a result, 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 always decreases as the algorithm proceeds.

A popular method for untangling triangulation in scientific computing community is based on optimization techniques [6, 7, 8]. 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 [6] untangles the mesh by maximizes 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 steepest descent, which is guaranteed to converge because of the convexity of the objective function. Using a more complicated form of objective function that also takes into account the quality of the mesh (e.g., minimum angle, aspect ratio of the triangles, etc.), one can accomplish mesh untangling and quality improvement simultaneously.

Our results. Unlike the two methods mentioned above, our approach for untangling triangulation strictly respects the vertex set of the triangulation: no vertex is inserted, deleted or moved after untangling. To this end, we introduce a class of regions called *locally removable regions* in the tangled mesh, that can be locally removed from \mathcal{T} and re-meshed with a set of upright triangles defined over the original set of vertices. Given a seed inverted triangle $T \in \mathcal{T}$, we show how to find such a region $R \subseteq \mathcal{T}$ with $T \in R$, which is minimum in the sense that, for any $R' \subseteq \mathcal{T}$ containing T and with similar properties, $f(R) \subseteq f(R')$. Our algorithm is output-sensitive in the sense that it never explores beyond the output region R .

As a result of the intricate interactions among inverted components of the tangled mesh, a locally removable region is often induced by a collection of inverted components rather than just the inverted component containing the seed triangle T . How to identify these inverted components through local exploration is not straightforward at all. Furthermore, during the local exploration process, the region that has been explored so far may have complicated structures (e.g., its boundary

may have large winding 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 novel insights into the structure of tangled meshes, which we deem as another main contribution of the paper. To the best of our knowledge, these structural properties have not been studied or exploited before, and may be of independent interest in contexts beyond mesh untangling.

The rest of the paper is organized as follows. In Section 2 we introduce the notion of *conflict sets* and *primary conflict sets* of inverted components and prove a number of their useful properties. In Section 3 we define *locally removable* regions 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. We conclude in Section 6.

2 Signs, Conflict Sets, and Shadows

In a tangled triangulation, an edge common to two inverted or two upright triangles is called a *regular* edge. A regular edge between two upright triangles is called an *upright edge*. An *inverted edge* is defined similarly. In contrast, we call an edge common to one upright and one inverted triangle, a *crease* edge. A vertex for which all the incident triangles are upright is called an *upright vertex*. An *inverted vertex* is defined analogously. A vertex that is neither upright nor inverted is called a *crease vertex*. Observe that a crease vertex must be incident to at least two crease edges.

All proofs in this section have been deferred to the appendix.

Signed incidence function. We like to talk about the interior points of an inverted triangle as negative points and those on the upright ones as positive ones. More precisely we define a sign function $s : S \rightarrow \{+1, 0, -1\}$ by letting

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

Notice that unless x belongs to a crease edge, the value of the signed incidence function remains the same in a sufficiently small neighborhood of x . Given a set $U \subseteq S$ we define the *signed incidence function relative to U* , as $\gamma_U : \mathbb{R}^2 \rightarrow \mathbb{Z}$ as

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

where $f^{-1}(x) = \{y \in S : f(y) = x\}$.

Lemma 2.1 *For any $x \in S$, $\gamma_S(x) = +1$.*

Conflict sets of inverted components. It is beneficial to deal with the untangling problem at a coarser granularity than individual triangles. Define the binary relation “ \sim ” between triangles by letting $T_1 \sim T_2$ if (1) T_1 and T_2 share an edge, and (2) T_1 and T_2 have the same orientation under f . Let “ \sim^* ” be the transitive closure of \sim . It is easy to see that \sim^* is an equivalence relation. Each equivalence class of \sim^* consists of triangles that all have the same orientation under f and are also connected through edges — or more precisely, their induced dual subgraph is connected. In

the sequel, by an *inverted (upright) component* we refer to an equivalence class of the \sim^* relation consisting of inverted (upright) triangles. We shall use \mathcal{I} to denote the set of all inverted components under the map f . As a convention, for $J \subseteq \mathcal{I}$, $f(J)$ should be interpreted as $f(\bigcup_{I \in J} I)$.

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.,

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

In other words, $\phi(U)$ consists of the set of all points in S that share the same space as the points in U under the map f . Notice that $U \subset \phi(U)$. Even when U is connected $\phi(U)$ may be not so. We are particularly interested in those connected components of $\phi(U)$ that intersect U . We define $\hat{\phi}(U)$ as the union of these connected components and call it the *primary conflict set* of U . We are particularly interested in sets of the form $\phi(J)$ where $J \subseteq \mathcal{I}$ is a collection of inverted components. The following lemma states some important properties concerning these sets.

Lemma 2.2 *Let $J \subseteq \mathcal{I}$ be a collection of inverted components. Then*

- (1) $f(\phi(J)) = f(J)$;
- (2) $f(\partial\phi(J)) \subseteq \partial f(J)$;
- (3) $\partial\phi(J)$ may only intersect J in $V \cap \partial J$;
- (4) If $x \in f^{-1}(\partial f(J)) \setminus \partial\phi(J)$, then x lies on a crease edge belonging to the boundary of some inverted component $I \in J$.

A simple yet important observation is that relative to each connected component of $\phi(J)$ the signed incidence function γ is constant.

Lemma 2.3 *Let $J \subseteq \mathcal{I}$ be a set of inverted components. Then for each connected component U of $\phi(J)$, γ_U is a constant function, i.e., there is a constant $c_U \in \mathbb{Z}$ such that $\gamma_U(x) = c_U$ for all $x \in f(J)$.*

At first sight, it seems hard to imagine γ_U of a connected component $U \in \phi(J)$ as described in Lemma 2.3 to be anything other than -1 , 0 , or $+1$. However, it happens that γ_U can indeed be an arbitrarily large integer. In fact this is true even for primary conflict set $\hat{\phi}(I)$ of a single inverted component I .

Lemma 2.4 For any integer $k \in \mathbb{Z}$, there exists a tangled triangulation and an inverted component I of it, for which $\gamma_{\hat{\phi}(I)} = k$.

Shadows and neighborhoods. For a subset $R \subseteq S$, we sometimes refer to $f(R)$ as the *shadow* of R . For a point $x \in S$ and a set $X \subset \mathbb{R}^2$, the *neighborhood of x under shadow X* is defined as

$$\mathcal{N}(x, X) = \{y \in S \mid \exists \text{ a path } p \subset S \text{ connecting } x \text{ and } y \text{ s.t. } f(p) \subset X\}.$$

Intuitively, $\mathcal{N}(x, X)$ is the region that can be reached from x without stepping out of the set X . Note that for any $y \in \mathcal{N}(x, X)$, $\mathcal{N}(y, X) = \mathcal{N}(x, X)$, because $\mathcal{N}(x, X)$ is path-connected. So, for a path-connected set $R \subseteq S$, we can write $\mathcal{N}(R, X) = \mathcal{N}(x, X)$ for any $x \in R$. The following simple lemma provides a convenient way to characterize primary conflict sets in terms of their shadows and neighborhoods.

Lemma 2.5 Let $J \subseteq \mathcal{I}$ be a collection of inverted components so that $\hat{\phi}(J)$ is path-connected. Let x be an arbitrary point in $\hat{\phi}(J)$. Then $\hat{\phi}(J) = \mathcal{N}(x, f(J))$.

3 Locally Removable Regions

We call a set $R \subseteq S$ *locally removable* if (i) R is path-connected, (ii) ∂R does not intersect the interior of any inverted components, and (iii) $\gamma_R(x)$ is identically $+1$ for all $x \in f(R)$. We are particularly interested in locally removable sets of the form $\hat{\phi}(J)$ for some collection $J \subseteq \mathcal{I}$ of inverted components.

Lemma 3.1 Let $J \subseteq \mathcal{I}$ be a set of inverted components. If $\hat{\phi}(J)$ is locally removable, then $f(\partial\hat{\phi}(J))$ consists of simple disjoint plane polygonal curves that identically make the boundary of $f(\hat{\phi}(J))$. Thus $f(\partial\hat{\phi}(J)) = \partial f(\hat{\phi}(J))$.

Proof. Lemma 2.2 (2) already implies that $f(\partial\hat{\phi}(J)) \subseteq \partial f(J)$. Observe that the restriction of the map f to $\partial\hat{\phi}(J)$ is surjective. This is because by Lemma 2.2 (4) for any point $y \in \partial f(\hat{\phi}(J))$, every point of the set $X = f^{-1}(y) \cap \hat{\phi}(J)$ is either contained in $\partial\hat{\phi}(J)$ or belongs to a crease edge. Since $\gamma_{\hat{\phi}(J)}(y) = 1$, there has to be one such point x of the former type in X . On the other hand, since X contains no negative points, there can only be one positive point in X and every other point must belong to crease edges. Thus the considered restriction of f indeed constitutes a bijection between $f(\partial\hat{\phi}(J))$ and $\partial f(\hat{\phi}(J)) = \partial f(J)$. The latter being a simple plane polygon, implies the former is so as well. \blacksquare

Let $J \subseteq \mathcal{I}$ be a collection of inverted components so that $R = \hat{\phi}(J)$ is locally removable. We define the *locally re-meshable region induced by R* as $R^* = \{T \in \mathcal{T} \mid R \cap \text{int } T \neq \emptyset\}$. The following lemma states that one can locally remove and re-triangulate the region R^* with upright triangles, and as a result, all inverted components in J are removed from the triangulation without introducing any new inverted triangle.

Lemma 3.2 There exists a triangulation \mathcal{T}' of the vertex set V , such that (i) $\mathcal{T} \setminus R^* \subseteq \mathcal{T}'$, and (ii) triangles in $\mathcal{T}' \setminus (\mathcal{T} \setminus R^*)$ are all upright.

Proof. We first construct a somewhat refined triangulation \mathcal{T}'' and then later coarsen it to the desired \mathcal{T}' . Specifically, we introduce a set of additional vertices

$$A = \{p \in S \mid p \text{ is a vertex of } \partial R\} \cap \{p \in S \mid p = \partial R \cap \text{int } e \text{ for some edge } e \in \mathcal{E}\}.$$




Figure 3: A locally re-meshable region induced by a locally removable region R . The shaded region is $f(R)$, and only triangles intersecting ∂R are drawn. The vertex set A is denoted by empty circles.

See Figure 3 for an illustration. Let $V'' = V \cup A$. For each triangle $T \in \mathcal{T}$ that intersects the boundary of R , T must be upright, and the region $T_R = T \setminus \text{int } R$ consists of a set of simple polygons with vertices in $V'' \cap T_R$. We can therefore triangulate each such T_R with upright triangles only. We further triangulate within $f(R)$ on the vertex set $V'' \cap R$, using upright triangles. By Lemma 3.1, the triangulations within T_R for each $T \in \mathcal{T}$ intersecting ∂R , the triangulation of $f(R)$, and $\mathcal{T} \setminus R^*$, together form a triangulation \mathcal{T}'' on the vertex set V'' .

Pick a vertex $p \in V'' \setminus V$, we remove p from the triangulation \mathcal{T}'' , together with all triangles incident upon it. Observe that $p \in \partial R$. Therefore a sufficiently small neighborhood of p is completely contained in upright components of the triangulation, and hence so does the star of p (the union of the triangles incident upon p). In particular, the star of p is a simple polygon, and we can re-triangulate it by upright triangles. We repeat the above process for each vertex in $V'' \setminus V$. In the end, we are left with a triangulation \mathcal{T}' of V with the desired properties. ■

4 Local Untangling

In this section we present an algorithm for computing a locally re-meshable region R^* induced by a locally removable region $R \subseteq S$ that contains a given seed inverted component $I \in \mathcal{I}$. By Lemma 3.2, we can then remove I from the triangulation by re-meshing R^* . The region R will be minimum in the sense that, for any $R' \subseteq S$ that is locally removable and contains I , $f(R) \subseteq f(R')$. We first present the algorithm at a high-level, followed by a more elaborate description of the algorithm.

Ultimately, the algorithm computes a collection $J \subseteq \mathcal{I}$ of inverted components, with $I \in J$, such that $\hat{\phi}(J)$ is locally removable. Let X be the shadow of J . At intermediate stages of the algorithm, an incomplete set J and an incomplete shadow X (not necessarily the shadow of the current J) are maintained. Specifically, J is set to $\{I\}$ and X is set to $f(I)$ initially, and at each subsequent step, the algorithm either expands the set J , or expands the shadow X , and then continues to the next stage with the new J or X . Otherwise, the algorithm declares that the desired J is found and terminates.

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

$$E(J) = J \cup \left\{ I \in \mathcal{I} : I \cap \partial \hat{\phi}(J) \neq \emptyset \right\}.$$

For any positive integer j , recursively define $E^j(J) = E(E^{j-1}(J))$ with the base case $E^0(J) = J$. Since there are only a finite number of inverted components and the operator E is *expansive*, i.e. $J \subseteq E(J)$, there is a finite $j \geq 0$ such that $E^j(J) = E^{j+1}(J)$. Let $E^*(J)$ denote $E^j(J)$ for this value of j , and call $E^*(J)$ the *extended inverted set* of J . We define $\psi(J)$ as $\hat{\phi}(E^*(J))$.

Lemma 4.1 For any $J \subseteq \mathcal{I}$,

- (1) $\partial\psi(J)$ intersects the interior of no inverted components.
- (2) for any $J' \subseteq \mathcal{I}$ with $J' \subseteq \hat{\phi}(J)$ we have $\hat{\phi}(J') \subseteq \hat{\phi}(J)$. In particular, if $J' \subseteq \psi(J)$, then $\psi(J') \subseteq \psi(J)$.

Proof. (1) Let $j \geq 0$ be an integer for which $E^*(J) = E^j(J)$ and suppose for contradiction that $\partial\psi(J)$ intersects the interior of some inverted component I . By Lemma 2.2 (3), $\partial\psi(J)$ does not intersect the interior of any of the inverted components in J . Therefore, I has to be an inverted component not included in J . By definition, $I \in E^{j+1}(J)$. Thus $I \in E^{j+1}(J) \setminus E^j(J)$ which contradicts the definition of $E^*(J)$.

(2) Observe that for any two subsets $J, J' \subseteq \mathcal{I}$ it holds that if $f(J') \subseteq f(J)$, then $\phi(J') \subset \phi(J)$ which means that every connected component of $\phi(J')$ has to be contained in some connected component of $\phi(J)$. In particular $\hat{\phi}(J')$ is contained in some connected component of $\phi(J)$. Thus if $J' \subset \hat{\phi}(J)$, which implies $f(J') \subseteq f(J)$, it has to be that $\hat{\phi}(J') \subseteq \hat{\phi}(J)$. The second half of the claim is a simple corollary of what we have just proved. ■

Lemma 4.2 Let $J \subseteq \mathcal{I}$ be such that $\partial\hat{\phi}(J)$ intersects the interior of no inverted component. Then for every $I \in J$, $E^*(I) \subseteq J$ and $\psi(I) \subseteq \hat{\phi}(J)$.

Proof. Since $I \in J$, $\hat{\phi}(I) \subset \hat{\phi}(J)$. If every point of $\partial\hat{\phi}(I)$ has non-negative sign, then $E^*(I) = \{I\}$ and $\hat{\phi}(I) = \psi(I)$ and we have nothing to prove. Otherwise, the boundary of $\hat{\phi}(I)$ intersects some inverted components other than I . Let I' be one such inverted component. Since $I' \cap \hat{\phi}(I) \neq \emptyset$ and $\hat{\phi}(I) \subseteq \hat{\phi}(J)$, $I' \cap \hat{\phi}(J) \neq \emptyset$. On the other hand, the interior of I' cannot intersect the boundary of $\hat{\phi}(J)$ by our assumption. Thus the interior of I' is entirely contained in $\hat{\phi}(J)$. By Lemma 4.1, $\hat{\phi}(\{I, I'\}) \subseteq \hat{\phi}(J)$. Applying this argument for all inverted components I' that intersect $\partial\hat{\phi}(I)$, we have shown that $E(I) \subset \hat{\phi}(J)$. We can now repeat this argument for $E(I)$, $E^2(I)$, and so forth. Thus by induction $E^*(I) \subset J$ and therefore by Lemma 4.1 (2), $\psi(I) \subseteq \hat{\phi}(J)$. ■




Figure 4: (i) Extending inverted components; (ii) Growing shadows.

Growing shadows. Consider a collection $J \subseteq \mathcal{I}$ of inverted components such that $\hat{\phi}(J)$ is path-connected and $\partial\hat{\phi}(J)$ intersects the interior of no inverted component. Clearly $\gamma_{\hat{\phi}(J)} > 0$. If $\gamma_{\hat{\phi}(J)} = +1$, then by definition $\hat{\phi}(J)$ is locally removable. However, if $\gamma_{\hat{\phi}(J)} \geq +2$ (which may indeed occur by Lemma 2.4), the extended inverted set of J , which remains to be J , does not help to extend $\hat{\phi}(J)$ further. In this case, instead of extending the set J directly, we shall expand the shadow of J , as follows. For a path $p \subset f(S)$, its preimage $f^{-1}(p)$ generally consists of a number of connected components. For a set $R \in S$, let $L(p, R)$ denote the subset of connected components of $f^{-1}(p)$ whose intersection with R is nonempty. We define the *extended shadow* of J , denoted by $G(J)$, as the union of $f(J)$ and the collection of points $x \in f(S) \setminus f(J)$ for which there exist $y \in \partial f(J)$ and a path p connecting x and y such that at least two connected components in $L(p, \hat{\phi}(J))$ intersect the interior of no inverted components. The following lemma justifies this definition.

Lemma 4.3 *Let $J \subseteq \mathcal{I}$ be a collection of inverted components such that $\hat{\phi}(J)$ is path-connected and $\gamma_{\hat{\phi}(J)} \geq +2$. For any set $J' \supseteq J$ for which $\hat{\phi}(J')$ is locally removable, $G(J) \subseteq f(J')$.*

Proof. Clearly $f(J) \subseteq f(J')$. Fix an arbitrary point $y \in G(J) \setminus f(J)$; we need to show that $y \in f(J')$. Let x and p be the point and the path that together serve as the witness for $y \in G(J)$. We claim that $p \subseteq f(J')$; since $y = p(1) \in p$, this claim would immediately imply $y \in f(J')$ and prove the lemma.

Now suppose for the sake of contradiction that p does not lie completely in $f(J')$. Observe that $p(0) = x \in f(J')$, and as such, $p \cap \partial f(J') \neq \emptyset$. Let $t \in [0, 1]$ be the smallest number for which $p(t) \in \partial f(J')$. Therefore, $p[0, t] \subseteq f(J')$. It then follows from Lemma 2.5 that $f^{-1}(p(t)) \cap L(p, \hat{\phi}(J)) \subset \hat{\phi}(J')$. Consider the connected components in $L(p, \hat{\phi}(J))$ that intersect the interior of no inverted components. Each such component contains a preimage of $p(t)$ whose sign function evaluates to $+1$. On the other hand, by Lemma 2.2 (4) and the fact that $\partial\hat{\phi}(J')$ intersects the interior of no inverted component, $f^{-1}(p(t)) \cap \hat{\phi}(J')$ contains no point whose sign function evaluates to -1 . Therefore $\gamma_{\hat{\phi}(J')}(p(t))$ is at least $+2$, contradicting with the assumption that $\hat{\phi}(J')$ is locally removable. \blacksquare

Algorithm. Recall that $I \in \mathcal{I}$ is the given seed inverted component. The following algorithm computes a locally re-meshable region R^* containing I .

```

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

```

By Lemma 2.5, $R = \hat{\phi}(J)$ and $f(R) = X$ at the end of the algorithm. Furthermore, $\gamma_R = +1$ and ∂R intersects the interior of no inverted component. Thus R is locally removable and R^* is locally re-meshable as desired. In addition, for any locally removable region R' containing I , by

Lemmas 4.2 and 4.3, it is easy to show that $f(R) \subseteq f(R')$. Therefore, the algorithm produces a locally removable region with the smallest shadow. There are different ways to implement the above algorithm in an output-sensitive manner. We sketch an implementation that is conceptually simple though by no means the most efficient. It is presented here for the sake of completeness and to illustrate that the algorithm never explores beyond the output region R^* .

Instead of maintaining R explicitly, we maintain the subset of triangles in \mathcal{T} whose interior intersects R . Slightly abusing the notation, we use R^* to denote this set; in the end, R^* coincides with the locally re-meshable region induced by R . The subset of triangles in R^* whose interior intersects ∂R is called the *frontline* F . By examining which triangles in the frontline are inverted, it is straightforward to compute the set $E(J)$ in step 3.

In each iteration when the shadow X is updated, we need to update R^* as in step 2. Starting from each triangle in the current frontline, we perform a depth-first search on the dual graph of the triangles in $\mathcal{T} \setminus (R^* \setminus F)$ to search for triangles whose image intersects X . To this end, we preprocess X for point location and ray shooting queries. Then for each triangle T encountered during the search, we can decide whether one of the vertices of $f(T)$ lies inside X , or one of the edges of $f(T)$ intersects ∂X . If either case happens, then $f(T)$ intersects X and we add T to the set R^* . In particular, when the latter case happens, T also belongs to the new frontline.

We still need to describe how to compute the extended shadow in step 5. To this end, we define a graph in which, each vertex represents a pair $\{T_1, T_2\}$ of upright triangles with $T_1, T_2 \in \mathcal{T} \setminus (R^* \setminus F)$ and $f(T_1) \cap f(T_2) \neq \emptyset$, and each edge connects two pairs $\{T_1, T_2\}$ and $\{T_3, T_4\}$ if $T_1 = T_3$ and T_2, T_4 share a common edge. We start from vertices in the graph corresponding to pairs of triangles in the current frontline (note that all triangles in the current frontline are necessarily upright), and perform a depth-first search on the graph to identify the connected components containing the starting vertices. Then one can show that the extended shadow $G(J)$ is the union of the current shadow X and the intersection $f(T_1) \cap f(T_2)$ for each visited pair $\{T_1, T_2\}$. Note that every triangle T involved in any one of these visited pairs will need to be added to the set R^* because $f(T) \cap G(J) \neq \emptyset$, and therefore the algorithm does not explore beyond the output region for computing extended shadows.

The running time of the above implementation can be easily bounded by $O(k_R \cdot C_R \cdot \log n_R)$, where n_R is the number of triangles in R^* , k_R is the number of inverted components in R^* , and C_R is the complexity of the overlay of $\phi(T) \cap R^*$ for all $T \in R^*$. By being more careful, it is possible to improve the running time to $O(C_R \cdot \log n_R)$.

Theorem 4.4 *Given any seed inverted component $I \in \mathcal{I}$, the algorithm outputs a locally re-meshable region R^* containing I , induced by a locally removable region R . Furthermore, for any $R' \subseteq S$ that is locally removable and contains I , $f(R) \subseteq f(R')$. The algorithm never explores beyond the output region R^* .*

5 Global Untangling

Using Lemma 2.1, it is not hard to prove that each connected component of $\phi(\mathcal{I})$ is a locally removable region. Therefore if we wish to untangle the entire mesh so as to remove all inverted components in \mathcal{I} , we can simply compute $R = \phi(\mathcal{I})$, 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 suggests that the algorithm presented in the previous section can be used to compute $\phi(\mathcal{I})$, and furthermore, step 5 can be skipped altogether, as its only purpose there is to ensure the condition $\gamma_R = +1$ which will not be needed here.

Theorem 5.1 Let $\mathcal{J} = \{J_1, \dots, J_m\}$, where, for each $J_i \subseteq \mathcal{I}$, $\partial\hat{\phi}(J_i)$ intersects the interior of no inverted component, and $\bigcup_{i=1}^m J_i = \mathcal{I}$. Then $\phi(\mathcal{I}) = \bigcup_{i=1}^m \hat{\phi}(J_i)$.

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

The proof works by induction on the number of inverted components under the map f . If $\mathcal{I} = \{I\}$ where I is the only inverted component of f , then the statement of the theorem is true by the following argument. Take a point $y \in \partial f(I)$ and consider the set $X = f^{-1}(y)$. By the choice of y , no point in X is negative. Therefore, since $\gamma_S(y) = +1$ and therefore the sum of the signs of the points in X is 1, there is precisely one point $x_1 \in X$ with $s(x_1) = 1$ and for every point $x \in X \setminus \{x_1\}$, $s(x) = 0$. Since I is the only inverted component, $E^*(I) = I$ and therefore $\gamma_{\hat{\phi}(I)} > 0$. Thus the sum of the signs of the points in $f^{-1}(y) \cap \hat{\phi}(I) \subset X$ and this implies that $x_1 \in \hat{\phi}(I)$. Therefore, $\gamma_{\hat{\phi}(I)} = +1$. Now, take any arbitrary point $y' \in f(I)$ and let $x' \in f^{-1}(y')$ be positive, i.e. $s(x') = 1$. We claim that $x' \in \hat{\phi}(I)$. This is because otherwise, since $\gamma_{\hat{\phi}(I)}(y') = 1$, the sum of the points in $f^{-1}(y') \cap \hat{\phi}(I)$ would be 1 and given that there are no other negative points in $f^{-1}(y')$ than those already in I , the sum of the signs of the points in $f^{-1}(y')$ would be at least 2 when $s(x')$ is accounted for. But $\gamma_S(y') = 1$, a contradiction.

Now suppose the statement of the theorem is true for any map f with $m-1$ inverted components and let f be a map with m inverted components $\mathcal{I} = \{I_1, \dots, I_m\}$. Note that $f(\mathcal{I}) = \bigcup_i f(I_i) = \bigcup_i f(\hat{\phi}(J_i))$. Take an arbitrary point $y \in \partial f(\mathcal{I})$ and let $X = f^{-1}(y)$. Since $f^{-1}(y)$ intersects the interior of no inverted components, no point in X is negative. On the other hand $\gamma_S(y) = +1$ by Lemma 2.1. Thus for a single point $x_1 \in X$, $s(x_1) = 1$ and for every other $x \in X \setminus \{x_1\}$, $s(x) = 0$. Now, take an arbitrary inverted component $I \in \mathcal{I}$ for which $y \in \partial f(I)$, and a set $J_i \in \mathcal{J}$ for which $I \in J_i$. Let $X_i = X \cap \hat{\phi}(J_i)$. By definition $\gamma_{\hat{\phi}(J_i)}(y) > 0$. Since $X_i \subseteq X$ and X contains exactly one positive point, i.e. x_1 , $x_1 \in X_i$ and $\gamma_{\hat{\phi}(J_i)} = +1$. Now let $y' \in f(J_i)$ be a positive point that is not in $f(J_j)$ for any $j \neq i$. We argue similar to the base case argument given above that every positive point $x' \in f^{-1}(y')$ is in $\hat{\phi}(J_i)$. Suppose to the contrary that $x' \notin \hat{\phi}(J_i)$. Since $\gamma_{\hat{\phi}(J_i)}(y') = 1$, the sum of the points in $f^{-1}(y') \cap \hat{\phi}(J_i)$ is 1. Since there are no other negative points in $f^{-1}(y')$ other than those in J_i , the sum of the signs of the points in $f^{-1}(y')$ with $s(x')$ considered must be at least 2. But $\gamma_S(y') = 1$, and the contradiction proves that $f^{-1}(y') \in \hat{\phi}(J_i)$.

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 $\partial f(J_i)$. Let P denote the set of points inside this plane polygon. Define a map $f_0 : S \rightarrow S$ by letting

$$f_0(x) = \begin{cases} f(x) & x \notin P, \\ x & \text{otherwise.} \end{cases}$$

Let \mathcal{I}_0 be the set of inverted components of f_0 . For any inverted component $I \in \mathcal{I}_0$, let $\phi_0(I)$, $\hat{\phi}_0(I)$, $E_0(I)$, and $E_0^*(I)$ respectively be the counterparts of $\phi(I)$, $\hat{\phi}(I)$, $E(I)$, and $E^*(I)$ only with respect to f_0 instead of f .

Observe that since $\partial\hat{\phi}(J_i)$ intersects no inverted components of \mathcal{I} , the inverted components in \mathcal{I}_0 are precisely the inverted components in \mathcal{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 $\phi(\mathcal{I}_0)$ is contained in $\hat{\phi}_0(J)$ for some $J \in \mathcal{I}$. A crucial observation is that for any $J \in \mathcal{I}_0$ and its counterpart $J' \in \mathcal{I}$, $\hat{\phi}_0(J) \subseteq \hat{\phi}(J')$. Thus every positive point in $\phi(\mathcal{I}_0)$ are contained in $\hat{\phi}(J_j)$ for some $j \neq i$.

On the other hand, we showed above that positive points in $\phi(I_i)$ are also contained in $\hat{\phi}(I_i)$. This prove the theorem. ■

6 Conclusions

Since our untangling algorithm first indentifies tangled regions of the mesh and then simply re-mesh these regions, it does not involve any edge flip operation. The algorithm of Shewchuk and Wallace [10] relies on a set of other local operations, besides edge flips, for untangling mesh. It is legitimate to ask whether untangling triangulation can also be accomplished by a pure edge-flip algorithm. Note that such a sequence of edge flips always exists, but it is not clear at all whether it can be found under the guide of a few simple rules based on local geometric information, as in Lawson’s algorithm. We leave it as an intriguing open question for future research.

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*, 21:179–197, 2002.
- [2] P. K. Agarwal, Y. Wang, and H. Yu. A two-dimensional kinetic triangulation with near-quadratic topological changes. *Discrete Comput. Geom.*, 36:573–592, 2006.
- [3] G. Albers, L. J. Guibas, J. Mitchell, and T. Roos. Voronoi diagrams of moving points. *Internat. J. Comput. Geom. Appl.*, 8:365–380, 1998.
- [4] J. Basch, L. J. Guibas, and J. Hershberger. Data structures for mobile data. *J. Algorithms*, 31(1):1–28, 1999.
- [5] J. Erickson, D. Guoy, J. Sullivan, and A. Üngör. Building space-time meshes over arbitrary spatial domains. In *Proc. 11th Intl. Meshing Roundtable*, pages 391–402, 2002.
- [6] L. Freitag and P. Plassmann. Local optimization-based simplicial mesh untangling and improvement. *Intl. J. Numer. Meth. Engin.*, 49:109–125, 2000.
- [7] P. Knupp. Hexahedral and tetrahedral mesh untangling. *Engineering with Computers*, 17:261–268, 2001.
- [8] K. Kovalev, M. Delanaye, and C. Hirsch. Untangling and optimization of unstructured hexahedral meshes. In *Proc. Workshop on Grid Generation: Theory and Applications*, 2002.
- [9] C. Lawson. Software for C^1 surface interpolation. In John R. Rice, editor, *Mathematical Software*, volume III, pages 161–194. Academic Press, NY, 1977.
- [10] J. Shewchuk and D. T. Wallace. Untangling triangulations, 2007. unpublished manuscript.
- [11] A. Üngör and A. Sheffer. Pitching tents in space-time: Mesh generating for discontinuous Galerkin method. In *Proc. 9th Intl. Meshing Roundtable*, pages 111–122, 2000.

Appendix

Proof of Lemma 2.1. By our boundary assumption for any point $y \in \partial S$, $\gamma_S(y) = 1$. Let z be an arbitrary point in S . To simplify the argument we assume for the moment that $z \notin f(\mathcal{E})$. Choose an arbitrary point $y \in \partial S$ and take a path p from y to z that does not self-intersect and is entirely contained in S . Clearly, p can be chosen in such a way to intersect $f(\mathcal{E})$ transitively, and avoid $f(V)$ altogether. To complete the proof we verify that the function $\gamma_S(p(t))$ remains constant as t varies from 0 to 1 and therefore $\gamma_S(z) = \gamma_S(y) = 1$.

To this end, 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 a section $p' = p([t_1, t_2])$, $0 \leq t_1 < t_2 \leq 1$ of p that does not intersect $f(\mathcal{E})$. Every connected component of $f^{-1}(p')$ is contained in a single triangle in \mathcal{T} and therefore every such component is a path bijectively mapped to p' 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 p' . Now consider the case that p' intersects the image $f(e)$ of a regular edge e . Without loss of generality assume that e is incident to two upright triangles T_1 and T_2 . The above argument can be extended to this case by observing that when we look at $f^{-1}(p')$, we can take $T_1 \cup T_2$ as one component for which $f^{-1}(p') \cap (T_1 \cup T_2)$ is mapped to p' bijectively. Thus the value of γ_S cannot change when p' crosses $f(e)$.

Thus γ_S can potentially change only when p crosses $f(e)$ of some crease edge e . Consider a section p' of p that intersects the image of no edge other than e . Once again consider the connected components of $f^{-1}(p')$. With the exception of one component, every other one of these components 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 to e , one is upright and the other is inverted. Every point of p' has precisely one inverse image in T_1 and another in T_2 . Therefore, the combined contribution of T_1 and T_2 to γ_S of a point in p' is zero. Therefore, the latter component has no effect on the value of γ_S in any point of p' . Thus γ_S is constant on p' with the same argument as used in the previous case.

Now to complete the argument, it suffices to break the path p between y and z into sections p' that each intersects the image of at most one edge and apply the above arguments to show that γ_S does not change along each section.

Proof of Lemma 2.2. (1) Since $J \subset \phi(J)$, $f(J) \subseteq f(\phi(J))$. On the other hand $f(\phi(J)) = f \circ f^{-1} \circ f(J) \subseteq f(J)$.

(2) The function f is continuous and therefore the inverse image of every open set is open. This implies that every point y in the preimage $f^{-1}(x)$ of an interior point x of $f(J)$ is an interior point of $f^{-1} \circ f(J) = \phi(J)$. Thus the image of the boundary of $\phi(J)$ under f is contained in the boundary of $f(J)$.

(3) Notice that $J \subset \phi(J)$. The boundary of J is made of crease edges. Let x be a point in the relative interior of an edge $e \subset \partial J$. The edge e is a crease edge and is therefore incident to an inverted triangle $T_1 \in J$ and an upright triangle $T_2 \notin J$. It is easy to see that there is a small enough neighborhood N of x such that $f(N) \subset f(T_1) \cap f(T_2)$. This implies that $N \subset \phi(T_1) \subseteq \phi(J)$. Thus x cannot be a boundary point of $\phi(J)$. It is therefore only possible for vertices of ∂J to appear in $\partial\phi(J)$.

(4) This follows from the fact that every interior point x of an inverted (or upright) component is mapped to $f(x)$ in a locally homeomorphic manner, i.e. there is a neighborhood N of x such that f maps N to $f(N)$ bijectively. This is clearly true for points in the relative interior of every triangle. For a point x in the relative interior of an upright (inverted) edge e , we observe that this holds as long as the neighborhood N is chosen small enough to fit inside the union of the two upright (inverted) triangles incident to e .

Proof of Lemma 2.3. Let U be any connected component of $\phi(J)$. Since f is continuous, $f(U)$ is connected. Let x and y be two points in $f(U)$ and consider a path $p : [0, 1] \rightarrow f(U)$ with $p(0) = x$ and $p(1) = y$. We can choose p such that its interior points, i.e. $\{p(t) : 0 < t < 1\}$ are contained in the interior of $f(J)$. The value of $\gamma_U(p(t))$ when t goes from 0 to 1 can only change when $p(t)$ reaches $f(\partial U)$. The argument is essentially identical to the one given for Lemma 2.1:

$\gamma_U(p(t))$ is unchanged as long as the number of positive and negative points in $f^{-1}(p(t)) \cap U$ remains unchanged, which is the case when $p(t)$ does not intersect the image of any edges that intersect U , or when it intersects the image of a regular edge that intersects U . Moreover, it is also the case when $p(t)$ crosses $f(e)$ for a crease edge e that intersects U since this adds or removes a pair of points from $f^{-1}(p(t))$ of which one is positive and one is negative and this does not affect the sum. Thus unless $f^{-1}(p(t)) \cap \partial U \neq \emptyset$, $\gamma_U(p(t))$ does not change. But by Lemma 2.2, $p(t)$ is not a boundary point of $f(J)$ and therefore p is contained in the interior of $f(U)$ as assumed. Thus if points $f(x)$ and $f(y)$ can be connected by a path in $f(J)$ which is the case when x and y belong to the same connected component U of $\phi(J)$, $\gamma_U(x) = \gamma_U(y)$. This proves that the function γ_U is constant in $f(U)$.

Proof of Lemma 2.4. 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 5. First in the figure on the left the gray area is an inverted component I . All other triangles are upright. In the middle figure, we add one extra layer of triangles to the star of one of the two spiraling vertices, in such a way that the entire inverted component is covered one more time. Clearly all the added triangles are upright. Now if you look at the boundary of this piece (the right figure), it 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 $\gamma_{\hat{\phi}(I)}$ having constant value $+2$.




Figure 5: 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).

Of course we need to show that this 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 red edges in the right figure will turn out to be crease edges. We have now a topological sphere. Now, poke a hole in the original inverted component I (which is part of this sphere) and a similar whole in a large plane triangulation and glue the punctured construction to the punctured plane triangulation along the boundaries of these holes.

Proof of Lemma 2.5. By Lemma 2.2 (1), the shadow of $\hat{\phi}(J)$ is exactly $f(J)$. Hence for any point $y \in \hat{\phi}(J)$, the path $p \subset \hat{\phi}(J)$ connecting x and y satisfies $f(p) \subset f(J)$, thus serves as a witness for $y \in \mathcal{N}(x, f(J))$. As such, $\hat{\phi}(J) \subseteq \mathcal{N}(x, f(J))$. On the other hand, fix a point $y \in \mathcal{N}(x, f(J))$ and a path p connecting x and y with $f(p) \subset f(J)$. It follows that $p \subset \phi(J)$. Therefore y lies in the same connected component of $\phi(J)$ as x , implying $y \in \hat{\phi}(J)$. Thus $\mathcal{N}(x, f(J)) \subseteq \hat{\phi}(J)$. The lemma then follows.