Erosion Thickness on Medial Axes of 3D Shapes
(Supplementary Proofs)

1 Properties of ET in 3D

1.1 Remark on infimum

We first make a note on the use of *infimum* instead of minimum in Equation 4 of the paper. This is because BT may only exist in the limit. An example is shown in Figure 1, where M is made up of a horizontal sheet and a vertical, cylindrical sheet close to the horizontal sheet's boundary. The two sheets meet at a non-manifold seam with a circular shape (dotted curve). The shortest exposing tree at a point x located behind the cylinder would be a single path γ to the boundary that lies infinitely close to a portion of the seam (red). Note that γ cannot exactly lie on the seam, or otherwise it would need to spawn paths onto the cylinder (as required by an exposing tree), whose lengths would be much longer. The infimum allows us to capture the length of γ as if it follows the seam.

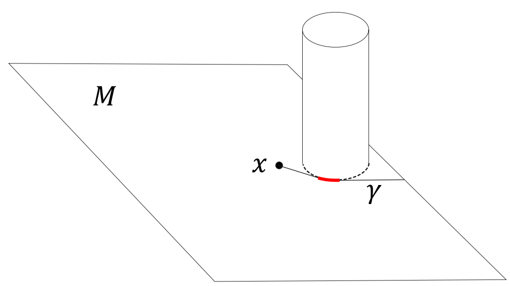

Figure 1: An example where BT is only realized at the limit.

1.2 Finiteness

Proposition 1.1. *$ET(x) = \infty$ if and only if $x \in M_C$ where M_C is the maximal closed subcomplex of M.*

Proof. Consider an exposing tree and traversing a path from its root. If the root is in M_C, then each time the path reaches the singular set at least one of its children must remain on M_C. If the path always remains in M_C it will have infinite combinatorial complexity which is not possible in an exposing tree. Hence points in M_C do not have exposing trees, and by definition they have infinite burn time and erosion thickness.

For points not in M_C we will build exposing trees with finite length showing that both burn time and erosion thickness are finite. Consider the components of M_2 that are not in M_C. We can order them as $C_1, \ldots, C_k$ such that $\partial(M_C \cup \overline{C_k} \cup \overline{C_{k-1}} \cup \cdots \cup \overline{C_i}) \subset \partial M \cup \overline{C_{i-1}} \cup \cdots \cup \overline{C_1}$. Let $X_i = M_C \cup \overline{C_k} \cup \overline{C_{k-1}} \cup \cdots \cup \overline{C_i}$ and define $X_{k+1} = M_C$. Let $Y_i = X_i - X_{i+1}$, which we will refer to as a *skirt*. In principle, this allows us to remove components of the manifold region that meet the boundary one at a time until M_C is obtained.

We will build exposing trees inductively. At any point in Y_1, there is a finite length path to ∂M that does not meet the singular set and such a path is an exposing tree. Assume every point in $Y_1 \cup \cdots \cup Y_{i-1}$ has a finite length exposing tree. Now consider a point in $x \in Y_i$. If $x \in \partial X_i$ then an exposing set of sectors of x lie on one or more skirts in $Y_1 \cup \cdots \cup Y_{i-1}$. We can build an exposing tree for x by combining paths on each sector to some interior points of these skirts with exposing trees at those points. Exposing trees can be built at a point x on the interior of X_i in two steps. First, take a path from x to some point on ∂X_i and combine it with the exposing tree at that point. Next, consider an arbitrary exposing set of sectors of x on M, say E. If $E' = E \setminus X_i$ is not empty, sectors in E' must lie on one or more skirts in $Y_1 \cup \cdots \cup Y_{i-1}$. We combine paths on each sector of E' from x to some interior points of these skirts with exposing trees at those points. This shows that every point in Y_i has an exposing tree and hence finite burn time and erosion thickness. □

1.3 Continuity

Proposition 1.2. *1. ET is continuous and 2-Lipschitz over $M_2 \setminus M_C$.*

2. Consider a singular point $x \in M_s \setminus (\partial M \cup M_C)$. ET over the union $M_2 \cup \{x\}$ is

(a) Upper-semicontinuous at x.

(b) Continuous and 2-Lipschitz at x within some 2-dimensional disk $D \subseteq M$ that contains x in its interior.

To prove these properties, we first present a lemma that bounds the variation of the radius function R. In the following, let $d_A(x, y)$ be the infimum of length over all paths restricted to a set A between two points $x, y \in A$.

Lemma 1.3. *Given any two distinct points $x, y \in M$, $R(x) \leq R(y) + d_M(x, y)$. If $x, y \in \mathcal{M}$, the inequality becomes strict.*

Proof. Suppose on the contrary that $R(x) > R(y) + d_M(x, y)$. Since $d_M(x, y)$ is no smaller than the Euclidean distance between x, y, this implies that the ball centered at y with radius $R(y)$ lies in the interior of the ball centered at x with radius $R(x)$. Since the ball at x is contained in the closure of S, this contradicts the fact that y has at least one nearest neighbor on ∂S whose distance from y is $R(y)$. If $x, y \in \mathcal{M}$, the equality cannot happen, since y needs to have at least two nearest neighbors on ∂S. □

Proof. (Proposition 1.2) Let $A = M_2 \setminus M_C$. We prove each property in order.

1: Consider two points $x, y \in A$. Since A is not closed, we cannot guarantee that a shortest path between x and y exists; however, we can find one that is arbitrarily close. For any $\epsilon > 0$, there is a path γ on A from x to y such that its length satisfies $|\gamma| < d_A(x, y) + \epsilon$. By the definition of BT, there also exists an exposing tree Γ at y such that $L(\Gamma) < BT(y) + \epsilon$. An exposing tree at x can be formed by gluing γ to the root of Γ. This shows that $BT(x) \leq |\gamma| + L(\Gamma) < BT(y) + d_A(x, y) + 2\epsilon$. Limiting ϵ to 0 yields $BT(x) \leq BT(y) + d_A(x, y)$. Combining Lemma 1.3, noting $d_M(x, y) \leq d_A(x, y)$, and switching roles of x and y, we arrive at the following condition:

$$ET(y) - 2d_A(x, y) \leq ET(x) \leq ET(y) + 2d_A(x, y).$$

In other words, the 1-Lipschitz properties of BT and R show that ET is 2-Lipschitz. Note that the Lipschitz condition also implies continuity of ET over A.

2(a): Following the same argument, we have the one-sided Lipschitz condition $ET(y) - 2d_A(x, y) \leq ET(x)$ for any $x \in M_s \setminus (\partial M \cup M_C)$ and $y \in A$. This implies upper-semicontinuity.

2(b): We first choose a suitable disk D. In the burning analogy, this is the last remaining neighborhood of x right before x is burned away. To precisely define D, we consider exposing trees at x that are compact in certain ways.

An exposing tree Γ at x is called *minimal* if there is no subtree of Γ also rooted at x that is an exposing tree for x. In other words, the sectors containing the root edges is a smallest exposing set. We are further interested in those minimal trees that have few, long subtrees. More formally, define the *height* of an root edge, e, which connects x to a subtree Γ_e, as $h(e) = |e| + L(\Gamma_e)$ where $|e|$ is the length of e. The root edge e is said to be *tall* if $h(e) \geq BT(x)$. Note that any tree Γ has at least one tall root edge.

Now, consider an minimal exposing tree at x, denoted as Γ_x, with the fewest tall root edges among all exposing trees at x. Denote its root edges as $\{e_1, \ldots, e_k\}$ ordered in ascending heights. Let C be the set of all sectors at x, and $\{c_1, \ldots, c_k\} \subseteq C$ those sectors containing the root edges. We claim that the remainder after moving the first $k-1$ sectors, $C' = C \setminus \{c_1, \ldots, c_{k-1}\}$, contains a unique set of sectors that forms a disk D. Otherwise, either e_k can be removed from Γ_x (if C' contains no disk), which implies that Γ_x is not minimal, or Γ_x is not an exposing tree (if C' contains multiple disks).

We next show that ET is 2-Lipschitz at x within D. Consider any point $y \in D \cap M_2$. Note that one side of the condition is already provided by our argument above for property 2(a). To show the other side, we follow an argument similar to that for property 1. For any $\epsilon > 0$, there is a path γ on A connecting x to y and an exposing tree Γ_y at y such that $|\gamma| + L(\Gamma_y) < BT(y) + d_A(x, y) + 2\epsilon$. Construct a new tree Γ' from Γ_x by replacing the root edge e_k (and its subtree) with γ and Γ_y. It is easy to see that Γ' is an exposing tree as well. Since Γ_x has the fewest tall root edges and e_k is tall, the replacement of e_k in Γ', γ, has to be tall as well. This yields $BT(x) \leq h(\gamma) < BT(y) + d_A(x, y) + 2\epsilon$. Limiting ϵ to zero, and combining Lemma 1.3 yields $ET(x) \leq ET(y) + 2d_A(x, y)$.

□

1.4 Local minima

Proposition 1.4. *For any $x \in \mathcal{M} \setminus (\partial M \cup M_C)$ and any 2-dimensional disk $D \subseteq \mathcal{M}$ that contains x in its interior, there is some $y \in D$ such that $ET(y) < ET(x)$.*

Proof. The basic idea is to look for some y along the boundary ∂D that is burned before x. Define $\epsilon = \min_{y \in \partial D}(R(y) + d_D(x, y) - R(x))$. By Lemma 1.3, ϵ is strictly positive. There exists an exposing tree Γ at x such that $L(\Gamma) < BT(x) + \epsilon$. Starting from x, we trace along a path in Γ that stays on D until it hits ∂D at a point y. Such tracing is always possible, because at least one child edge at any vertex of Γ on D has to stay on D, and a path only ends on the boundary of M. Let Γ_y be the subtree of Γ rooted at y, we have:

$$\begin{aligned}
BT(y) &\leq L(\Gamma_y) && \text{definition of BT} && (1)\\
&\leq L(\Gamma) - d_\Gamma(x, y) && \text{definition of tree length} && (2)\\
&< BT(x) + \epsilon - d_\Gamma(x, y) && \text{choice of } \Gamma && (3)\\
&\leq BT(x) + \epsilon - d_D(x, y) && \text{path on } \Gamma \text{ from } x \text{ to } y \text{ lies on } D && (4)\\
&\leq BT(x) + \epsilon - (R(x) + \epsilon - R(y)) && \text{definition of } \epsilon && (5)\\
&= BT(x) - R(x) + R(y) && && (6)
\end{aligned}$$

□

2 Error bound of graph-restricted ET

Proposition 2.1. [Proposition 6.2 in paper] *Let $|M|$ count the number of triangles in M, g be the maximal gradient magnitude of R on any triangle edge on ∂M, and ω be the maximal distance between adjacent nodes in G on a triangle edge. For any node v in G,*

$$ET(v) \leq ET_G(v) \leq ET(v) + (2|M| + g)\omega \tag{7}$$

We first introduce a piece-wise representation for an exposing tree Γ on M rooted at some triangle vertex or a point on some triangle edge. Assume that each vertex of Γ is either at a triangle vertex or on a triangle edge (otherwise it will be on M_2 and hence can be treated as an interior point to some edge of Γ). An edge e of Γ can be partitioned into a sequence of curve segments, such that each segment, c, lies interior to a triangle or on a triangle edge and that the end points of c lie on a triangle edge or at a triangle vertex. We call each such segment a *chord*.

To prove the error bound, we first show that Γ can be approximated by some exposing tree restricted to graph G, which is no longer than Γ plus a quantity proportional to the maximal number of chords along any root-to-leaf path in Γ (Lemma 2.2). Next we show that, if Γ is short, it has a bounded number of chords along any of its root-to-leaf path (Lemma 2.3).

Lemma 2.2. *Consider a node v of G. For any exposing tree Γ at v over M, there is some exposing tree Γ' at v restricted to G such that*

$$L(\Gamma') \leq L(\Gamma) + (2K + g)\omega \tag{8}$$

where K is the maximum number of chords along any root-to-leaf path of Γ.

Proof. We first construct a graph G^+ from G by adding, for each triangle, arcs between every pair of non-adjacent nodes along a triangle edge. That is, the subgraph of G^+ on each triangle is a complete graph of its vertices and edge samples. We call the additional arcs *jump-arcs*, as they jump over some edge-arcs connecting adjacent nodes. See Figure 2 for an illustration.

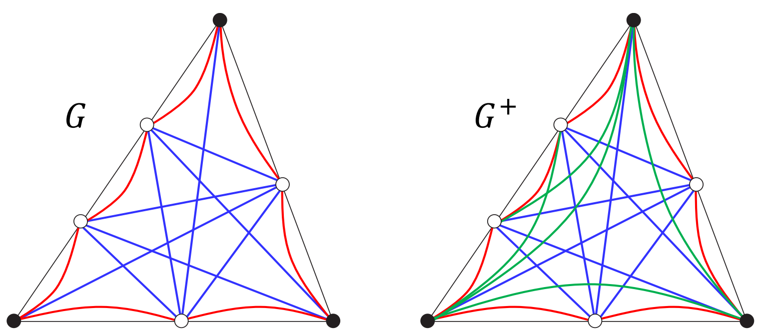

Figure 2: G^+ adds additional jump-arcs (green) to G that connect non-adjacent nodes on a triangle edge.

We can transform any exposing tree Γ^+ at v restricted to G^+ to some exposing tree restricted to G with the same length. Specifically, if Γ^+ does not contain any skip-arcs, no change is needed. Otherwise, consider a skip-arc a in Γ^+. Let the triangle edge containing a be e, and let the subtree of a in Γ^+ be Γ_a^+. If $e \in M_2$ (shared by one or two triangles), we simply replace a by the edge-arcs jumped over by a on *any* triangle sharing e. Note that the replacement arcs remain on M_2 and hence the result is still an exposing tree. Otherwise, if e is non-manifold (shared by more than two triangles), consider the set E of edge-arcs jumped over by a on *all* triangles sharing e. These edge-arcs can form an image of a tree, denoted by Γ_E, where every root-to-leaf path in the tree is represented by a sequence of edge-arcs in E. We replace a in Γ^+ by Γ_E and glue one copy of Γ_a^+ to each leaf of Γ_E. We can verify that the modification also results in an exposing tree with the same length as Γ^+.

Next, given an arbitrary exposing tree Γ over M rooted at node v, we construct an exposing tree Γ^+ restricted to G^+ by replacing each chord of Γ with some arc in G^+. For each chord c with end points p_1, p_2, it is replaced by an arc a connecting nodes n_1, n_2 as follows:

- If $p_i (i = 1, 2)$ is at a triangle vertex, n_i will be that vertex. Otherwise, n_i is the nearest node to p_i on the interior of the triangle edge containing p_i (this is always possible, because our graph construction in Section 6.1 ensures that each triangle edge has at least one interior node).

- If n_1, n_2 lie on different edges of a triangle, a is the unique triangle-arc on that triangle connecting n_1, n_2 (Figure 3 left and middle). Otherwise, n_1, n_2 lie on a same triangle edge, say e, and the graph G^+ has multiple edge-arcs (or jump-arcs) connecting them, one for each triangle sharing e. If c lies interior to one triangle, a is the edge-arc (or jump-arc) on that triangle (Figure 3 right). Otherwise c lies on e, we arbitrarily pick one triangle sharing e and set a to be the edge-arc (or jump-arc) on that triangle connecting n_1, n_2.

Our choice of nodes n_i ensures that arcs constructed for consecutive chords in Γ are connected. It can be further verified that Γ^+ forms an exposing tree at v.

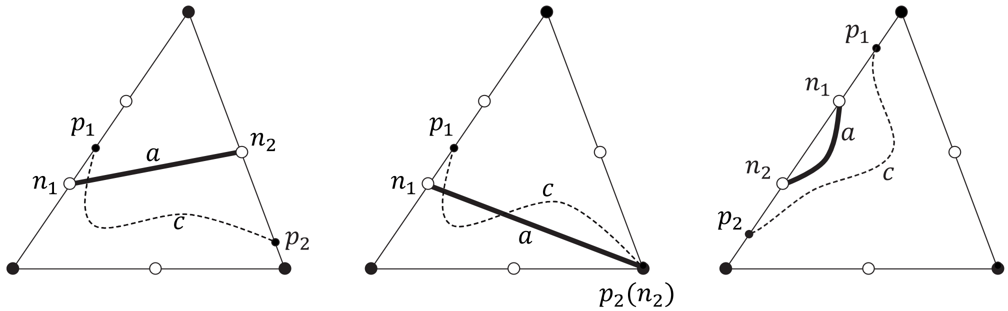

Figure 3: Example of arcs a (solid) chosen for chords c (dashed).

Finally, we will bound the difference between the length of trees Γ and Γ^+. Consider each chord c in Γ with end points p_1, p_2 and its corresponding arc a in Γ^+ with nodes n_1, n_2. Since a is straight, and since the length of both edge segments $\{p_1, n_1\}$ and $\{p_2, n_2\}$ are no greater than ω, the following inequality holds:

$$|a| \leq |\{n_1, p_1\}| + |c| + |\{p_2, n_2\}| \leq |c| + 2\omega \tag{9}$$

Now consider any root-to-leaf path γ of Γ with k chords, and let its corresponding root-to-leaf path in Γ^+ be γ^+. Summing the inequality above over all chords of γ yields

$$|\gamma^+| \leq |\gamma| + 2k\omega \tag{10}$$

Let u be the leaf node at the end of γ, and u^+ be its corresponding leaf on γ^+. Since the gradient magnitude of R over ∂M is bounded by g, we have

$$R(u^+) \leq R(u) + g\omega \tag{11}$$

Combining Equations 10,11 and considering all root-to-leaf paths yields inequality

$$L(\Gamma^+) \leq L(\Gamma) + (2K + g)\omega$$

Since there exists an exposing tree restricted to G that has the same length as $\Gamma+$, this concludes the proof. □

Lemma 2.3. *Consider any exposing tree Γ rooted at v on M where v is either a triangle vertex or a point on a triangle edge. There exists an exposing tree Γ' at v such that*

1. $L(\Gamma') \leq L(\Gamma)$
2. *Each root-to-leaf path of Γ' has at most one chord on any single triangle (including its edges).*

Proof. Consider a root-to-leaf path γ on Γ and a triangle t such that γ has more than one chords on t (if such γ and t do not exist, let $\Gamma' = \Gamma$ and we are done). Let the set of all chords on t be C. Following the root-to-leaf direction on γ, let p be the first end of the first chord of C and q be the second end of the last chord on C. We create a new chord, c, as the straight segment connecting p and q (see Figure 4 left). If p, q both lie on a non-manifold triangle edge, c would lie on the triangle interior but stay infinitely close to that edge (see Figure 4 right). We then replace the entire segment on Γ between p and q, together with any subtrees rooted along the segment, by chord c. The result is an exposing tree that is no longer than Γ.

Note that the operation mentioned above strictly decreases the total number of chords of Γ. Hence we can repeat the operation on Γ until such γ and triangle t cannot be found (in case that Γ has infinite number of chords, we first apply this operation to those triangles that contain infinite number of chords). □

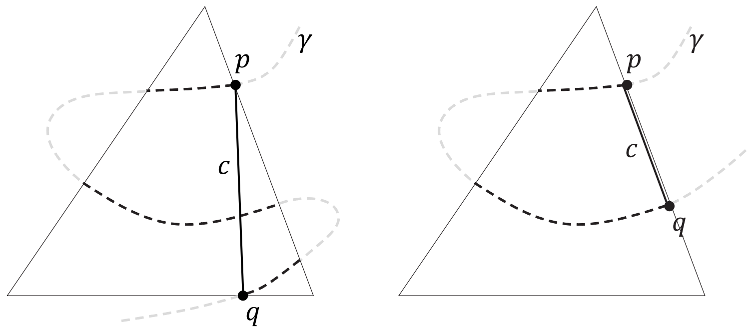

Figure 4: Examples of replacing multiple chords (dark and dashed) on a path γ that lie on the same triangle by a single chord c (solid).

Proof. (Proposition 2.1) A direct corollary of Lemma 2.3 is that the maximum number of chords along any root-to-leaf path of Γ' is bounded by the total number of triangles on M. Combining with Lemma 2.2, there exists some exposing tree Γ'' restricted to G such that

$$L(\Gamma'') \leq L(\Gamma') + (2|M| + g)\omega \leq L(\Gamma) + (2|M| + g)\omega$$

Since the above holds for any exposing tree Γ at v, by definition of burn time, we arrive at the second inequality in 7. The first inequality trivially holds because BT_G considers a subset of exposing trees (restricted to G). □

3 Algorithms

3.1 Accuracy of Burn

Proposition 3.1. [Proposition 6.3 in paper] *At the termination of algorithm **Burn**, $v.time = BT_G(v)$ for every node v.*

We will prove the following Lemma with a stronger assertion that, whenever a node or a sector is marked as burned, its "time" field accurately records the time at which the fire front restricted to G burns away the node or sector. The claim above is a direct corollary of this lemma, since the algorithm terminates with all nodes marked as burned.

Lemma 3.2. *The following holds at the end of each* While *loop in the algorithm **Burn**,*

1. *$\forall s \in v.sectors$, if $s.burned =$True, then $\forall a \in s.arcs$,*

$$s.time \leq a.len + BT_G(v_a) \tag{12}$$

   where v_a denotes the end node of arc a that is not v.

2. *If $v.burned =$True, then $v.time = BT_G(v)$.*

Proof. We prove both properties by induction. Initially, they trivially hold because all nodes and their sectors are marked as unburned. Assuming they hold for all previous loops, we prove the two properties below for the current loop. To simplify the discussion, we drop "restricted to G" when we talk about exposing trees.

1. We only need to consider a sector s that is newly burned in the current loop (otherwise inequality 12 holds by induction hypothesis). We separately consider the case where s is the primary sector and the case where s is exposed by the primary sector.

- $s = v.primeSec$: By induction hypothesis, inequality 12 holds for all arcs a where v_a is marked as burned. We only need to show, for any $a \in s.arcs$ where v_a is not yet burned, and for any exposing tree Γ at v_a, $s.time \leq a.len + L(\Gamma)$.

  To do so, we perform a walk on Γ from its root v_a as follows. At a vertex (node) u of Γ, we follow the child edge (arc) b that is on some un-burned sector of u. We stop if such b does not exist. Figure 5 illustrates the notations.

  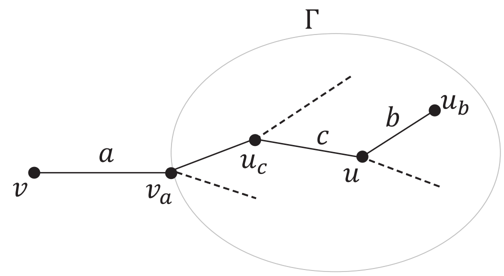

  Figure 5: Illustration for proof of Lemma 3.2.

  We can show that, when we stop, u is either a burned node or it is on ∂G. Otherwise, suppose u is an interior node of Γ (since $u \notin \partial G$) and it is not burned. If u has not been popped from Q before, it would have no burned sectors. Otherwise, the un-burned sectors of u would form closed disks (or they would have been exposed and burned). In either case, some child arc of u in Γ must lie on an unburned sector of u, and we can keep walking.

  Consider the node u where we stop. If $u \notin \partial G$, u must have been marked as burned and hence it cannot be the root, v_a. Let c be the parent arc of u in Γ and t be the sector of the parent node u_c that contains c (see Figure 5). Note that both u_c and t have not been burned (or we would have stopped earlier in the walk). We have:

  $$\begin{aligned} s.time &= v.time && s \text{ is the primary sector} && (13)\\ &\leq u_c.time && v \text{ is at the head of } Q && (14)\\ &\leq t.time && u_c.time \text{ is the smallest } t'.time \text{ for all unburned sectors } t' && (15)\\ &\leq c.len + BT_G(u) && \text{induction hypothesis, since } u \text{ is burned} && (16)\\ &\leq c.len + L(\Gamma_u) && \Gamma_u \text{ is the subtree of } \Gamma \text{ rooted at } u && (17)\\ &\leq L(\Gamma) < a.len + L(\Gamma) && && (18) \end{aligned}$$

  The only remaining case is when $u \in \partial G$ and it is not burned. This implies that u has never been popped from Q before, which in turn implies $u.time \leq u.R$. We therefore have

  $$s.time = v.time \leq u.time \leq u.R \leq L(\Gamma) < a.len + L(\Gamma)$$

- $s \neq v.primeSec$: Let $t = v.primeSec$. Using the same argument above, we have $t.time \leq a.len + L(\Gamma)$ for any arc $a \in s.arcs$ where v_a is not burned and any exposing tree Γ at v_a. Since we set $s.time$ to be $v.time$ inside the loop, and $t.time = v.time$, we arrive at the inequality 12.

2. We need to show, for any exposing tree Γ at v, $v.time \leq L(\Gamma)$. Consider all sectors of v that were not burned before the current loop; denote the set as N. Note that N either contains all sectors of v (if v has not been popped from Q before this loop) or contains closed disks. In either case, Γ must have a root edge a on some sector $s \in N$. Denote the subtree of Γ rooted at v_a as Γ_a. Since v is burned in this loop, all sectors in N are necessarily burned in the same loop. By inequality 12, we have

$$v.time = s.time \leq a.len + BT_G(v_a) \leq a.len + L(\Gamma_a) \leq L(\Gamma)$$

□

3.2 Non-intersecting burn trees

A burn tree at a node v is constructed by following arcs $s.primeArc$ (provided it is not *null*) on all sectors $s \in v.sectors$ back to the boundary ∂G. Here we show that no two arcs, from the same burn tree or two different burn trees, have non-trivial intersection.

Proposition 3.3. *After the algorithm **Burn** terminates, consider any two nodes v, u of G and any pair of sectors $s \in v.sectors$ and $t \in u.sectors$. Let $a = s.primeArc$ and $b = t.primeArc$. If $a \neq$null, $b \neq$null, and $a \neq b$, then $a \cap b$ is either empty or is a subset of end points of a and b.*

Proof. It is easy to check that the only scenario where $a \cap b$ may contain interior points of a or b is when both a, b are edge-arcs of G on a same triangle. See Figure 6 for two examples.

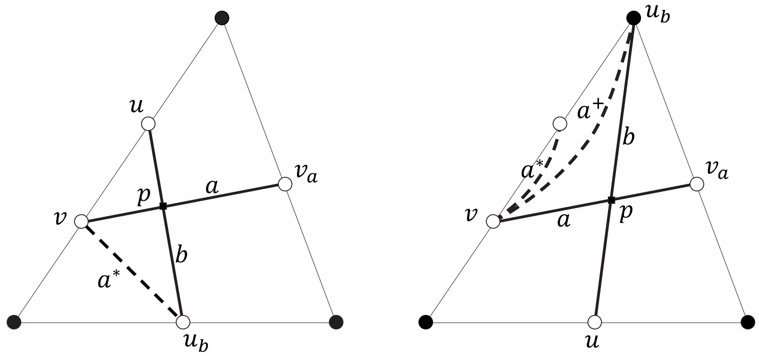

Figure 6: Two examples of intersecting arcs.

Denote the end node of a that is not v as v_a (and similarly u_b). Due to Lemma 3.2 (1) and the definition of primary arc, we have

$$a = \arg \min_{a' \in s.arcs} (a'.len + BT_G(v_{a'})), b = \arg \min_{b' \in t.arcs} (b'.len + BT_G(u_{b'})) \tag{19}$$

Suppose a, b intersects at a non-end point p. Without loss of generality, suppose $|pu_b| + BT_G(u_b) \leq |pv_a| + BT_G(v_a)$, where $|pu_b|$ is the length of the straight segment connecting p and u_b (the $\geq$ case can be handled in a symmetric fashion).

If there is an arc a^* in G that connects v with u_b (Figure 6 left), then by triangle inequality we have

$$\begin{aligned} a^*.len + BT_G(u_b) \quad & < |vp| + |pu_b| + BT_G(u_b) \\ & \leq |vp| + |pv_a| + BT_G(v_a) \\ & = a.len + BT_G(v_a) \end{aligned}$$

which contradicts Equation 19. If such v and u_b are not connected by any arc in G, then they necessarily are non-adjacent nodes on the same triangle edge. Consider the augmented graph G^+ as constructed in the proof of Lemma 2.2 (see also Figure 2). There is an arc a^+ in G^+ that connects v and u_b (Figure 6 right). Using the same argument above and noting that $BT_G \equiv BT_{G'}$ (as shown in the proof of Lemma 2.2), we have $a^+.len + BT_G(u_b) < a.len + BT_G(v_a)$. Using the same transformation technique in the proof of Lemma 2.2, there is an arc a^* in G connecting v to an adjacent node v_{a^*} on the same triangle edge such that $a^*.len + BT_G(v_{a^*}) = a^+.len + BT_G(u_b)$, reaching the same contradiction.

$\square$

3.3 Topology of medial curves

We will show that the medial curve C extracted using our dualization technique in Section 7.1 of the paper preserves the topology of the triangulated medial axis M. We assume that M has a generic structure. By [Giblin and Kimia 2004], this means that each point on M has one of the five local topology as shown for $x_i (i = 1, \ldots, 5)$ in Figure 5 of the paper.

Our approach is to construct a homotopy-preserving deformation retract from M onto C. The retract proceeds in two stages. First, a narrow band around the burn trees is "carved out" from M. Denoting this band as B, the second stage contracts the remainder of the first stage, $M \setminus B$, onto C.

The band B is made up of infinitesimal pieces called *hubs* and *spokes* that lie respectively at the nodes and arcs of G. For each node v of G, the hub is a regular neighborhood of v that is arbitrarily small, and it is made up of *flaps* that correspond to the sectors of v. For each arc a on some burn tree, the spoke is an arbitrarily narrow strip containing a that connects the hubs of the two end nodes of a. Figure 7 (left) illustrates the hubs and spokes for four nodes and three arcs.

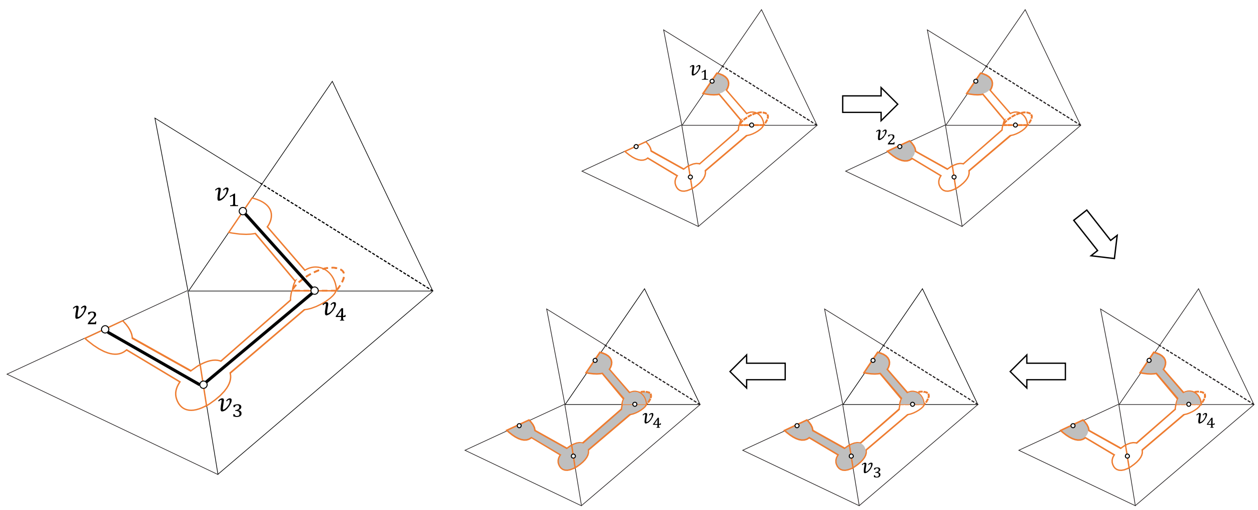

Figure 7: Left: narrow band (orange outline) made up of hubs around four nodes and spokes around three arcs. Right: a sequence of removal of the band.

In the first stage, we retract M to $M \setminus B$ following the progress of the algorithm **Burn**. At each *While* loop where a node v is popped from Q, we remove the spoke of the primary arc of $v.primeSec$ (if it is not *null*) and the flaps of v's hub corresponding to $v.primeSec$ and its exposed sectors (hence the hub of v is completely removed when v is burned). Figure 7 (right) illustrates the process where the nodes are popped from Q in the order v_1, v_2, v_4, v_3, v_4 (v_4 is burned the second time it is popped). It can be shown the removal process preserves homotopy as long as the unburned sectors of any node v forms a single connected component at any time during the algorithm. To show the latter, note that the unburned sectors either comprise of all sectors at v (if v has not been popped) or form closed disks. Since the sectors in each of the five generic local topologies are connected and do not contain multiple, disjoint closed disks, the unburned sectors of v at any time of the algorithm are connected.

In the second stage, we retract $M \setminus B$ onto C. We do so with the aid of a triangulation T_t within each triangle t of M. Specifically, for each face of the subdivision on t defined by the triangle edges and arcs in the burn tree, we create a fan of triangles connecting the dual vertex of the face to each boundary segment of the face. See an illustration in Figure 8 (left, middle). Note that T_t contains the segments of C in its edge graph. There are in fact only two types of triangles in T_t: ones having a segment of C opposite to a node of G (Figure 8 top-right), and ones having an arc of G opposite to a vertex on C (Figure 8 bottom-right). For triangles of the first type, contraction starts from the boundary of the node's hub and proceeds towards the opposite edge. For the second type, contraction starts from the boundary of the arc's spoke and the boundary of its end nodes' hubs and moves towards the opposite vertex. It is easy to verify that such contraction is consistent between neighboring triangles of T_t, and that it continuously deforms the boundary of B onto C.

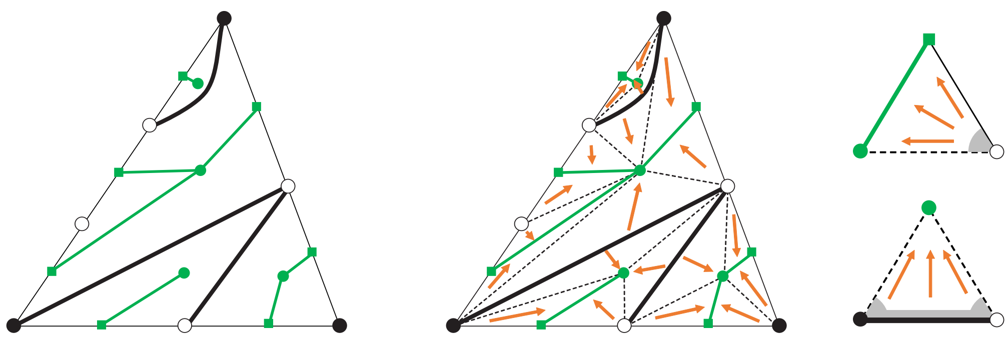

Figure 8: Left: arcs in burn trees (black) and dual medial curve (green) on a triangle t. Middle: triangulation T_t showing retract direction within each triangle. Right: two types of triangles in T_t and the retract direction within each.