Immersion of Self-Intersecting Solids and Surfaces: Supplementary Material 1: Topology

Yijing Li and Jernej Barbič, University of Southern California

May 15, 2018

1 Topological definitions

A mapping $f : X \to Y$ is **injective** (also called "one-to-one") if for every $x, y \in X$ such that $x \neq y$, we have $f(x) \neq f(y)$. It is **surjective** (also called "onto") if for each $y \in Y$, there exists $x \in X$ such that $f(x) = y$, i.e., $f(X) = Y$.

A **topological space** is a set X equipped with a topology. A topology is a collection τ of subsets of X satisfying the following properties: (1) the empty set and X belong to τ, (2) any union of members of τ still belongs to τ, and (3) the intersection of any finite number of members of τ belongs to τ. A neighborhood of $x \in X$ is any set that contains an open set that contains x.

A **homeomorphism** between two topological spaces is a bijection that is continuous in both directions. A mapping between two topological spaces is continuous if the pre-image of each open set is continuous. Homeomorphic topological spaces are considered equivalent for topological purposes.

An **immersion** between two topological spaces X and Y is a continuous mapping f that is locally injective, but not necessarily globally injective. I.e., for each $x \in X$, there exists a neighborhood $\mathcal{N}$ of x such that the restriction of f onto $\mathcal{N}$ is injective.

Topological closure of a subset A of a topological set X is the set $\overline{A}$ of elements of $x \in X$ with the property that every neighborhood of x contains a point of A.

A topological space X is **path-connected** if, for each $x, y \in X$, there is a continuous mapping $\gamma : [0, 1] \to X$ (a curve), such that $\gamma(0) = x$ and $\gamma(1) = y$.

A **manifold** of dimension d is a topological space with the property that every point x has a neighborhood that is homeomorphic to $\mathbb{R}^d$ (or, equivalently, to the open unit ball in $\mathbb{R}^d$), or a halfspace of $\mathbb{R}^d$ (or, equivalently, to the open unit ball in $\mathbb{R}^d$ intersected with the upper halfspace of $\mathbb{R}^d$). Intuitively, this means that the manifold is a space that is "locally Euclidean". The points that fall into the second case form the **manifold boundary**. In addition, one requires that manifolds are Haussdorff spaces, i.e., distinct points have disjoint neighborhoods, and second-countable, i.e., there exists a countable collection $\mathcal{U} = \{U_i\}_{i=1}^{\infty}$ of open sets such that each open set is a union of a subfamily of $\mathcal{U}$. These additional conditions just ensure that the manifold is a sufficiently "nice" space; all subsets of Euclidean spaces $\mathbb{R}^k$ satisfy these additional conditions. For $d = 2$, important examples of manifolds in this paper are the sphere and torus with $k \geq 1$ handles, both of which have no boundary. For $d = 3$, in this paper, manifolds have boundaries. Typical examples are solid balls, or solid tori with $k \geq 1$ handles.

A manifold M embedded into an Euclidean space is **compact** if, intuitively, there exists a bounding box that completely contains it (i.e., the manifold is finite in size), and M is closed, i.e., contains all of its limits. Formally, a topological space X is compact if for any infinite family of open sets $\{U_\alpha\}_\alpha$, that covers X, there exists a finite subfamily that also covers X. By the Whitney embedding theorem [Whitney, 1944] and its extensions [Munkres, 2000], every compact d–manifold (even if non-smooth) can be embedded into $\mathbb{R}^N$, for some properly high N, e.g., $N = 2d + 1$ is sufficient. Bijective continuous mapping from a compact topological space X onto a Haussdorff topological space (such as an $\mathbb{R}^k$) are homeomorphisms.

Disjoint union $X \sqcup Y$ of sets X and Y is a union of X and Y where we have ensured that the elements of $X \cap Y$ are present twice. Formally, it can be formed as $X \sqcup Y = (X \times \{0\}) \cup (Y \times \{1\})$, where $\times$ denotes the Cartesian set product, and $\cup$ is the usual set union.

Equivalence relation $\sim$ on a set X is a relation with the properties that (1) $x \sim x$ for each $x \in X$, (2) if $x \sim y$ for some $x, y \in X$, then also $y \sim x$, and (3) if $x \sim y$ and $y \sim z$ for some $x, y, z \in X$, then $x \sim z$. An equivalence class is a maximal set such that all elements are equivalent to one another.

A **quotient space** of a set X with respect to an equivalent relation $\sim$ on X is the set of equivalence classes of $\sim$ on X.

A **quotient topological space** of a topological space X with respect to an equivalence relation on X, is the quotient space of X with respect to $\sim$, equipped with the quotient topology. The quotient topology consists of all sets with an open pre-image under the projection map that maps each element of X to its equivalence class. Intuitively speaking, a quotient topological space is the result of identifying or "gluing together" certain points of a given topological space. This is commonly done in order to construct new spaces from the existing ones.

Self-touching input triangle meshes are defined as follows. Let $x \in \hat{M}$ where $\gamma(x) \geq 2$, i.e., $x \in \Gamma$. Then, we call x to be a *crossing intersection* if there exists a (sufficiently small) neighborhood $\mathcal{N}$ of $\hat{\rho}(x)$ such that $\hat{\rho}^{-1}(\mathcal{N})$ consists of $\gamma(x)$ components in $\hat{M}$, and the $\hat{\rho}$-projections of any two of these compoments cross each other in $\mathbb{R}^d$, as opposed to merely touching. Then, we call the input mesh M self-touching if there exists at least one point $x \in \Gamma$ that is *not* a crossing intersection. The intent is to capture degenerate inputs where a surface touches itself at a single point, or along a loop, but does not penetrate. For $d = 2$, an example of such a degeneracy is a circle touching a square. For $d = 3$, an example is a solid bowl placed upside down on a solid box. We note that self-touching input triangle meshes should not be confused with self-touching cells. The former is a degenerate input that our algorithm does not permit, whereas the latter is a common occurence in our algorithm, even with non-degenerate inputs.

A covering projection between path-connected topological spaces X and Y is a continuous surjective mapping $p : X \to Y$, such that for each $y \in Y$, there exists a neighborhood U so that $p^{-1}(U)$ is a union of disjoint open sets, each of which is mapped homeomorphically onto U.

The **total turning number** of a planar curve is the angle, divided by 2π, by which a passenger riding down this curve will rotate when they make one loop around the curve. The total turning number is a non-zero integer. It should not be confused with the winding number. The winding number is a property of a point with respect to the curve. The total turning number is an inherent property of the curve independent of any particular point.

A **group** is a set G equipped with a binary operation $\circ$, such that (1) $x \circ y \in G$ for each $x, y \in G$, and (2) $\circ$ is associative ($(x \circ y) \circ z = x \circ (y \circ z)$ for all $x, y, z \in G$), (3) $x \circ e = e \circ x = x$ for a special element $e \in G$ (the "unity"), and (4) for each $x \in G$, there exists the inverse y so that $x \circ y = y \circ x = e$.

A **homomorphism** is a mapping between groups G and H that respects the group operation, i.e., $f(a \circ b) = f(a) \circ f(b)$ for any $a, b \in G$.

An **isomorphism** is a bijective homomorphism. Groups are isomorphic if there exists an isomorphism between them. Isomorphic groups are considered equivalent for algebraic purposes.

A **subgroup** H of a group G is a subset H of G which is closed under the group operation. The **index** of H in G is the number of distinct left (or, equivalently, right) cosets of H in G. Left and right cosets of H in G are sets $aH = \{ah\,;\, h \in H\}$ and $Ha = \{ha\,;\, h \in H\}$, respectively, where $a \in G$. The number of left and right cosets is always the same.

The **free product** $G_1 * G_2 * \ldots * G_n$ of groups $G_1, \ldots, G_n$ is the group of finite words (of arbitrary length) $a_{i_1}^{j_1} a_{i_2}^{j_2} \ldots a_{i_\ell}^{j_\ell}$, where $a_i \in G_i$, $\ell \geq 0$, $j_k \geq 1$, and i_k are arbitrary integer indices between 1 and n.

The **fundamental group** $\pi_1(X)$ of a path-connected topological space X is the set of loops, starting and ending at some basepoint $b \in X$, under the equivalence relation that equates loops that can be continuously morphed into one another. Group operation is path concatenation.

2 Topological proofs

(Section 3) $\hat{M}$ **is a manifold:** If the equivalence class $\hat{x} \in \hat{M}$ originates from $x \in M$ that is in the interior of the triangle face (i.e., $\hat{x} = \{x\}$), then x trivially has a neighborhood homeomorphic to the open ball in $\mathbb{R}^{d-1}$. If $\hat{x}$ originates from the interior of an edge between two triangles, then, because M has manifold connectivity, there are exactly two triangle joining at that edge. Therefore, it is possible to find a neighborhood homeomorphic to the open ball in $\mathbb{R}^{d-1}$. The remaining case is that

$\hat{x}$ originates from a vertex x of M. Then, take any triangle that has x as its vertex. Because M has no boundary, there must be two adjacent triangles that share an edge and vertex x with M. We can continue adding such adjacent triangles. Because the number of triangles is finite, we must eventually form a complete 360 degree loop around x. Because, due to the manifold connectivity assumption, the triangles touching x must form a continuous fan, the patch of triangles looping around x is homeomorphic to a closed disk in $\mathbb{R}^{d-1}$, and therefore there exists a neighborhood of $\hat{x}$ homeomorphic to the open ball in $\mathbb{R}^{d-1}$. The Haussdorff and countable separability axioms are trivially satisfied because $\hat{M}$ is constructed from triangles in the Euclidean space. ∎

(Section 3) **Augmentation of ρ to an immersion $\hat{\rho}$ of $\hat{M}$ onto M** : As stated in paper, we assume that M is non-degenerate. In order for ρ to be augmentable, it needs to respect the gluing equivalence relation of $\hat{M}$, i.e., it maps all members of each equivalence class of $\hat{M}$ into the same point in $\mathbb{R}^3$. This is true because all the "glued" equivalent points are co-located in $\mathbb{R}^d$. We can therefore un-ambiguously define $\hat{\rho}(\hat{x}) = \rho(x)$. Because ρ is continuous, $\hat{\rho}$ is also continuous. It remains to be proven that $\hat{\rho}$ is locally injective. Given $\hat{x} \in \hat{M}$, we can form the manifold neighborhood N of x in M, as in the proof of manifoldness of $\hat{M}$. Because the restriction of the mapping ρ onto this neighborhood is injective, it follows that $\hat{\rho}$ is injective on the augmentation $\hat{N} = \{[x]; x \in N\}$ of N, where $[x]$ is the equivalence class of x. ∎

(Section 4.1) **Topology of cells:** By construction, cells are open and connected. For connected inputs, because M is not degenerate, the boundary B of the closure of each cell is a $d-1$ manifold without boundary. B is also orientable, and a closed and bounded subset of $\mathbb{R}^d$. For $d = 3$, by the characterization theorem of 2-manifolds [Brahana, 1921], B must topologically be either a sphere, or a torus with $k \geq 1$ handles. For $d = 2$, it must be a topological circle. ∎

(Section 5.1) **Finiteness of cells:** If there exists a volume immersion $\hat{\sigma}$ from a compact manifold $\hat{S}$ onto $\mathbb{R}^d$ whose boundary is $\hat{M}$, then the pre-image of each cell $\hat{\sigma}^{-1}(\mathcal{C}) \subset \hat{S}$ consists of a finite number of components in $\hat{S}$.
Proof: Suppose the pre-image of a cell $\mathcal{C}$ consists of an infinite number of components. Then, construct a sequence $y_i \in \hat{S}$ such that each y_i is in a different connected component of $\hat{\sigma}^{-1}(\mathcal{C})$. Because $\hat{S}$ is compact, there exists a subsequence $\{x_i\}_i$ of the sequence $\{y_i\}_i$ such that x_i has a limit $x \in \hat{S}$. Suppose x is not a boundary point of the manifold $\hat{S}$. Then, there exists a neighborhood $\mathcal{N}$ of x that is homeomorphic to the unit ball in $\mathbb{R}^d$, such that the restriction of $\hat{\sigma}$ onto $\mathcal{N}$ is a homeomorphism onto its image. This last property holds because $\hat{\sigma}$ is an immersion. Because $\hat{\sigma}$ is continuous, $\hat{\sigma}(x)$ must be in $\overline{\mathcal{C}}$. Now, we need to consider several cases depending on whether $\hat{\sigma}(x)$ lies in the interior of $\overline{\mathcal{C}}$ (i.e., inside $\mathcal{C}$; note that by our definition, cells are open sets), or on a patch, arc or corner point on the boundary of $\mathcal{C}$. Suppose $\hat{\sigma}(x) \in \mathcal{C}$. Then, it is possible to select a neighborhood $\mathcal{N}' \subset \mathcal{N}$ of x such that $\hat{\sigma}(\mathcal{N}') \subset \mathcal{C}$, and the restriction of $\hat{\sigma}$ onto $\mathcal{N}'$ is a homeomorphism onto its image. However, this implies that $\mathcal{N}'$ does not contain any points that map to $\rho(M)$. Because $x_i \to x$, the sequence will eventually enter $\mathcal{N}'$, and so we will find x_{i_0} and x_{i_0+1} that are both inside of $\mathcal{N}'$. We can then join these two points with an arc contained completely inside $\mathcal{N}'$. None of the points on this arc map to $\rho(M)$, and therefore x_{i_0} and x_{i_0+1} are in the same connected component of $\hat{\sigma}^{-1}(\mathcal{C})$. This is a contradiction, proving the original statement. The other cases where $\hat{\sigma}(x)$ lies on a patch, arc or corner point, or where x is a boundary point of $\hat{S}$, can be handled in the same way; we just have to consider "half-disk", "quarter-disk", or "eight-disk"

neighborhoods $\mathcal{N}'$ of x. ■

(Section 5.1) **Restrictions of immersions to cells are covering projections.** In addition, we will simultaneously prove the following. Let $p = \hat{\sigma}_{\mathcal{S}} : \mathcal{S} \to \mathcal{C}$ be the restriction of the immersion $\hat{\sigma}$ onto a pre-cell $\mathcal{S} \subset \hat{S}$. Then, for any $y \in \mathcal{C}$ the cardinality of $p^{-1}(y)$ is finite and constant on $\mathcal{C}$ (the "number of sheets").

Proof: By the definition of $\mathcal{S}$, p is surjective and continuous. Suppose the cardinality of $p^{-1}(y)$ would be infinite for some $y \in \mathcal{C}$. Then by the same compactness argument as in the proof of the finiteness of cells (above), we can construct a sequence $x_i \in \mathcal{S}$ such that $p(x_i) = y$ for all i, and x_i has a limit $x \in \overline{\mathcal{S}} \subset \hat{S}$. Because $\hat{\sigma}$ is an immersion, there exists a neighborhood of x such that the restriction of $\hat{\sigma}$ is injective. However, because the sequence x_i eventually has to enter this neighborhood, injectivity is impossible, as $p(x_i) = y$ for all i. Therefore, the cardinality of $p^{-1}(y)$ is finite for all $y \in \mathcal{C}$. Let us now prove the existence of the neighborhood U from the definition of the covering space, and that the cardinality of $p^{-1}(y)$ is constant on $\mathcal{C}$. Let $c(y) = |p^{-1}(y)|$ for $y \in \mathcal{C}$. For each of the $c(y)$ distinct elements $z_i \in \mathcal{S}$, such that $p(z_i) = y$, there exists a neighborhood Z_i of z_i that maps homeomorphically onto a neighborhood of y. We can always shrink Z_i so that they are pairwise disjoint. Because the number of z_i is finite, we can set

$$U = \bigcap_i p(Z_i).$$

By construction, $p^{-1}(U)$ is a union of disjoint open sets, each of which is mapped homeomorphically onto U. It also follows that for any $y' \in U$, $c(y') = |p^{-1}(y')| = c(y)$, i.e., $c(y)$ is locally constant on U. Any locally constant scalar function on a path-connected space is constant [Munkres, 2000]. Because the cell $\mathcal{C}$ is path-connected, $c(y)$ is constant on the entire cell $\mathcal{C}$. ■

(Section 5.1) **Theorem 1:** If a valid input M is simply volume-immersible, then for each cell $\mathcal{C}$ of the cell complex of M, the number of connected components of $\hat{\sigma}^{-1}(\mathcal{C})$ equals the winding number $\mathrm{wind}(x, M)$ of any point $x \in \mathcal{C}$ with respect to M. The winding number $\mathrm{wind}(x, M)$ is the same for all points $x \in \mathcal{C}$.

Proof: Although we numbered this theorem as Theorem 1 for pedagogical reasons, it actually depends on Theorems 4 and 5, whose proof must logically happen first. Note that Theorems 4 and 5, and Corollary 6 do not discuss or need the winding number, or this Theorem 1, so there is no circular dependency. This is because Theorems 4 and 5 never specifically postulated any number of copies of each cell in the immersion graph. So, we defer this proof until later on in this document, when both Theorem 4 and 5 are proven.

(Section 5.1) **Lemma (total turning number):** Let i be an immersion of the unit disk in $\mathbb{R}^2$ into $\mathbb{R}^2$, such that its restriction on the unit circle is a surjective immersion onto a closed 2D loop $\gamma : [0,1] \to \mathbb{R}^2$; $\gamma(0) = \gamma(1)$. Then, the total turning number of γ must be 1 or -1.

Proof: Because i is an immersion, there exists $\epsilon > 0$ such that the the restriction of i on a disk of radius ϵ centered at the origin is injective. Now, consider the following homotopy of immersions:

$$H \; : \; [0,1] \times [0,1] \to \mathbb{R}^2, \tag{1}$$

$$H(s,t) = \left((1-s) + \frac{s\epsilon}{2}\right) \begin{bmatrix} \cos(2\pi t) \\ \sin(2\pi t) \end{bmatrix}. \tag{2}$$

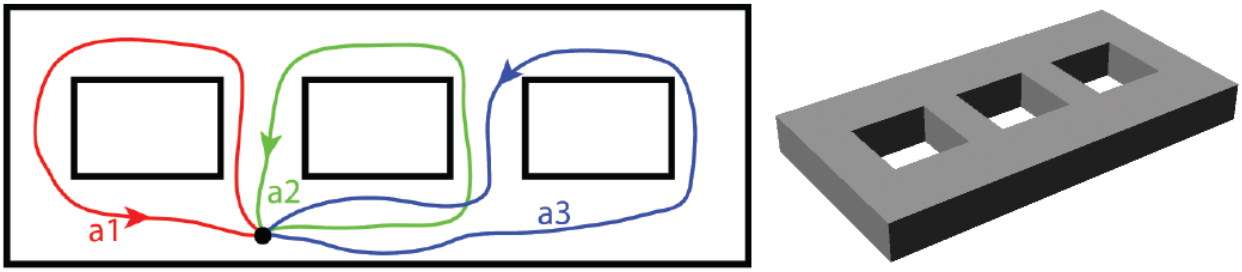

Figure 1: The generators of $\pi_1(\mathcal{C})$. In this case, $\mathcal{C}$ is a solid torus with 3 handles ($k = 3$), viewed (on the left) under an orthographic projection from the top.

Because on the ϵ-disk, i is a homeomorphic onto its image, this is a homotopy of immersions between the immersion onto a closed 2D loop γ ($s = 0$) and an immersion of circle onto S^1 (i.e., a topological circle; $s = 1$). By the Whitney-Graustein theorem [Whitney, 1937], two immersions are homotopic through a family of immersions if and only if they have the same turning number. This means that the total turning number of γ must be 1 or -1. ■

(Section 5.1) **Uniqueness and existence of lifted paths:** Let $p : X \to Y$ be a covering space projection. Pick an arbitrary starting point $x \in X$, and let γ be a path in Y that begins at $p(x)$. Then, there exists a unique path $\tilde{\gamma}$ in X that starts at x and that "lifts" γ, i.e., $p \circ \tilde{\gamma} = \gamma$. This is a standard result in algebraic topology [Glickenstein, 2018].

(Section 5.1) **Theorem 2:** If a valid input M is the boundary of an immersion from a disk in $\mathbb{R}^d$, then the immersion is simple. This holds both for $d = 2$ and $d = 3$.
Proof: Let $\mathcal{C}$ be a cell of the cell complex of M, and $\mathcal{S}$ be its "pre-cell", i.e., one of the connected components of $\hat{\sigma}^{-1}(\mathcal{S}$. Let $p = \hat{\sigma}_{\mathcal{S}} : \mathcal{S} \to \mathcal{C}$ be the restriction of the immersion $\hat{\sigma}$ onto $\mathcal{S} \subset \hat{S}$. By the "Restrictions of immersions" lemma (above), $\mathcal{S}$ is a covering space over $\mathcal{C}$. In order to prove simple immersibility, we must prove that the number of sheets of $\mathcal{S}$ over $\mathcal{C}$ is 1. Topologically, cell $\mathcal{C}$ is the interior of a torus with $k \geq 0$ handles. This is because the boundary of the cell is a connected orientable compact 2-manifold, and therefore, by the classification theorem of 2-manifolds [Brahana, 1921], has to be homeomorphic to a torus with $k \geq 0$ handles (the case $k = 0$ denotes the sphere). We can homeomorphically deform cell $\mathcal{C}$ into a canonical shape, whereby the walls are planar at 90-degree angles (see Figure 1), and the holes are perpendicular to the y-axis. The space $\mathcal{S}$ can then be seen as "hovering" above $\mathcal{C}$ along the y-axis.

For $k = 0$, the fundamental group $\pi_1(\mathcal{C})$ is $\{e\}$. For $k > 0$, the fundamental group of $\mathcal{C}$ can be computed using the Seifert-van Kampen theorem [Seifert, 1931, Van Kampen, 1933] using the fundamental group of the solid torus, which is $\mathbb{Z}$. We therefore have $\pi_1(\mathcal{C}) = \mathbb{Z} * \ldots * \mathbb{Z}$ (free product; k times), i.e., $\pi_1(\mathcal{C})$ is the set of all words formed by k letters $a_1, \ldots, a_k$, i.e., all finite words of the form $a_{i_1}^{j_1} a_{i_2}^{j_2} \ldots a_{i_\ell}^{j_\ell}$. Each of the words corresponds to a loop which starts at some fixed base point b, and loops around one of the holes exactly once; the finite words corresponds to all possible loops in $\mathcal{C}$ (see Figure 1). The projection p induces a homomorphism $p_\#$ of $\pi_1(\mathcal{S})$ into $\pi_1(\mathcal{C})$, whereby every loop γ in $\mathcal{S}$ with a basepoint in $p^{-1}(b)$ is mapped to $p \circ \gamma$. A well-known result in algebraic topology is that the number of sheets equals the index of the subgroup $p_\#(\pi_1(\mathcal{S}))$ in $\pi_1(\mathcal{C})$ [Glickenstein, 2018].

We therefore need to prove that this index is 1. For $k = 0$, because $\pi_1(\mathcal{C})$ is trivial, so is $p_\#(\pi_1(\mathcal{S}))$, and therefore, this index is 1. For $k > 0$, we use another standard result in algebraic topology, which states that $p_\#(\pi_1(\mathcal{S}))$ is isomorphic to the subgroup of loops in $\pi_1(\mathcal{C})$ which have the property that they lift to loops in $\mathcal{S}$ [Glickenstein, 2018]. We will now prove that every loop in $\pi_1(\mathcal{C})$ lifts to a loop in $\mathcal{S}$. Then, it follows that $p_\#(\pi_1(\mathcal{S}))$ is isomorphic to $\pi_1(\mathcal{C})$, and therefore, the number of sheets equals 1. Because the generator loops $a_1, \ldots, a_k$ generate $\pi_1(\mathcal{C})$, it is sufficient to prove that they each lift to a loop in $\mathcal{S}$. Suppose one of them does not (say, a_1 without loss of generality). Let us place the coordinate system so that the origin is in the middle of the handle of a_1, and the y-axis is pointing through the handle. This means that the lifting of a_1 into $\mathcal{S}$ is not a loop, but an arc joining two different points B_1 and B_2 from $p^{-1}(b) \subset \mathcal{S}$. We can then perform the lifting again starting from B_2, obtaining another arc, from B_2 to B_3. This process will eventually terminate with $B_M = B_1$, for $M \geq 3$, producing a loop in $\mathcal{S}$ that winds itself around the y-axis once, and whose projection via immersion p into $\mathcal{C}$ winds itself $M-1$ times around the y-axis. Because p is an immersion, this is not possible by the total turning lemma (above). Hence, all the generator loops lift to a loop in $\mathcal{S}$. ∎

(Section 5.2) **Theorem 5:** Let M be a valid input. Assume that there is a cell graph and a B-patch ownership assignment that satisfy Rules 1-7 given in Section 5.1. Then, M is volume-immersible, and the graph can be used to construct the immersion.

Proof: By gluing the cells according to the cell graph, we will form a d-manifold $\hat{S}$, and an immersion $\hat{\sigma}$ from $\hat{S}$ into $\mathbb{R}^d$ such that $\hat{M}$ is the boundary of $\hat{S}$, and the restriction of $\hat{\sigma}$ onto $\hat{M}$ is $\hat{\rho}$. We first define $\hat{S}$, as follows. Note that $\hat{S}$ will be an "abstract" manifold, and if M has self-intersections, not embedded into $\mathbb{R}^d$. Number the cell graph nodes as $1, 2, ...$, in arbitrary order. Then, define S as the disjoint union

$$S = \coprod_{i \in \text{nodes}(G)} \overline{C(i)} \times \{i\}, \tag{3}$$

where i runs over all the nodes of the cell graph, $\times$ is the set Cartesian product, $C(i)$ is the cell corresponding to node i, and $\overline{}$ is the topological closure operation. We then form the set $\hat{S}$ as a quotient set of S with respect to a gluing relation whereby we identify points on a closure of a patch shared by two connected nodes as equivalent. We define the mapping $\sigma : S \to \mathbb{R}^d$ as $\sigma(x, i) = x$, where $x \in \overline{C(i)} \in \mathbb{R}^d$. We then define our volume immersion $\hat{\sigma} : \hat{S} \to \mathbb{R}^d$, by augmenting σ with respect to the gluing relation. Because each patch is owned by exactly one node (Rule 7), we can embed $\hat{M}$ into $\hat{S}$. The immersion $\hat{\sigma}$ trivially matches $\hat{\rho}$ on $\hat{M}$.

We now prove that $\hat{S}$ is a d-manifold. Let $x \in \hat{S}$. We consider the different cases. (1) If x is in the interior of a node c, then since cells are d-manifolds, x trivially has a neighborhood homeomorphic to $\mathbb{R}^d$. (2) If x is on the boundary of c, we have several subcases. (2.1) If x is on a B-patch p that the node owns, we consider whether it is on the boundary of the patch or not. (2.1.1) If it is in the interior of p, then x has a neighborhood homeomorphic to the halfspace of $\mathbb{R}^d$ due to manifoldness of cells. (2.1.2) Else, it is on the boundary of p, and therefore in the closure of one or more arcs. Hence, x is shared with other topologically neighboring patches. By Rule 4, the graph nodes owning the patches are connected. By Rule 5, no other B-patches contribute to the neighborhood of x. So, there is a neighborhood of x homeomorphic to the halfspace of $\mathbb{R}^d$, bounded by the patches and emanating into the interior of the cells. (2.2) Else, node c declined p. By Rule 2, there must be another node d connecting c across p. By Rule 1, there are no other nodes connected across p. (2.2.1) If x is in the interior of p, then there is a neighborhood of x homeomorphic to $\mathbb{R}^d$ straddling p and emanating into the interiors of c and d. (2.2.2) Else, x is on the boundary of p. Since the boundary of a patch consists of arcs, and each arc is shared by two pairs of topologically neighboring patches, there are four patch connections surrounding x. If the four patch connections are from four different patches, we are either in the situation in Figure 15 (in main paper), or (2.1.2) applies. By Rule 6, the nodes surrounding the arcs whose closure contains x, are all connected. So, there is a neighborhood of x homeomorphic to $\mathbb{R}^d$ straddling patches and emanating into the interior of cells.

Finally, we prove that $\hat{\sigma}$ is an immersion. The mapping $\hat{\sigma}$ is continuous because the glued cells are adjacent in space, and is trivially locally injective if x is in the interior of a node. If x is on the boundary of a node, then local injectivity follows from the fact that the glued cells occupy distinct regions of space in $\mathbb{R}^d$. ■

(Section 5.1) **Theorem 1:** If a valid input M is simply volume-immersible, then for each cell $\mathcal{C}$ of the cell complex of M, the number of connected components of $\hat{\sigma}^{-1}(\mathcal{C})$ equals the winding number $\text{wind}(x, M)$ of any point $x \in \mathcal{C}$ with respect to M. The winding number $\text{wind}(x, M)$ is the same for all points $x \in \mathcal{C}$.

Proof: We now give our "deferred" proof of Theorem 1. By Corollary 6, there exists an immersion graph G and a B-patch ownership assignment that satisfy Rules 1-7 given in Section 5.2. Let $\mathcal{C}_1, \ldots, \mathcal{C}_n$ be the cells of the cell complex of M. By the "Finiteness of Cells" lemma, the number of connected components (the "pre-cells") of $\mathcal{C}_i$ is finite; denote it by c_i. Let $x \in \mathcal{C}_{i_0}$, for some i_0, $1 \le i_0 \le n$. Then, we have

$$\text{wind}(x, M) = \int_{y \in M} \Omega(x, y) dS = \sum_{i=1}^{n} \sum_{j=1}^{c_i} \int_{y \in \partial \mathcal{C}_i} \Omega(x, y) dS = \sum_{i=1}^{n} \sum_{j=1}^{c_i} \Big(\textbf{if } (x \in \mathcal{C}_i) \textbf{ then } 1 \textbf{ else } 0\Big) = \quad (4)$$

$$= \sum_{i=1}^{n} c_i \Big(\textbf{if } (x \in \mathcal{C}_i) \textbf{ then } 1 \textbf{ else } 0\Big) = c_{i_0}. \quad (5)$$

where $\partial \mathcal{C}_j$ is the boundary of cell $\mathcal{C}_j$, and where $\Omega(x, y)$ is the "winding number kernel", namely $\Omega(x, y) dS$ is the solid angle seen from x towards y, due to a small surface patch dS on M, as shown in [Jacobson et al., 2013]. The second equality in (5) holds because of Rules 1, 2 and 7. Namely, every boundary patch of a node is either shared by two nodes that are connected in G (in which case the contributions by these two nodes on this patch cancel each other in the sum), or it is owned by the node and therefore appear exactly once and forms a part of the boundary of M. ■

(Section 6.1) **Our 3D immersion problem is NP-complete:** For $d = 2$, it has been proven in [Eppstein and Mumford, 2009] that the problem of determining if the given planar curve is a boundary of an immersed surface, is NP-complete. The size of the problem is measured as a number of 2D intersections (2D arcs), or equivalently, 2D patches or 2D cells. For $d = 3$, it is easiest to measure the problem size as a number of patches. For $d = 3$, we can easily prove that the problem is also NP-complete. Given a 2D problem, M_{2D}, we can pick a 3D line that does not intersect the bounding box of M_{2D} and revolve M_{2D} around it 360°, producing a 3D problem for a 3D input M. Due to the geometry of our 3D surface of revolution, M_{2D} is a boundary of an immersed surface for $d = 2$ if and only if M is a boundary of an immersed surface for $d = 3$. Therefore, the immersibility problem for $d = 2$ is polynomial-time reducible to the immersibility problem for $d = 3$. It now follows from Eppstein's result that the problem for $d = 3$ is NP-complete. ■

References

[Brahana, 1921] Brahana, H. R. (1921). Systems of circuits on two-dimensional manifolds. *Annals of Mathematics*, pages 144–168.

[Eppstein and Mumford, 2009] Eppstein, D. and Mumford, E. (2009). Self-overlapping curves revisited. In *Proc. of ACM-SIAM Symp. on Discrete Algorithms*, pages 160–169.

[Glickenstein, 2018] Glickenstein, D. (2018). Covering space. Technical report, University of Arizona.

[Jacobson et al., 2013] Jacobson, A., Kavan, L., and Sorkine-Hornung, O. (2013). Robust inside-outside segmentation using generalized winding numbers. *ACM Transactions on Graphics (TOG)*, 32(4):33.

[Munkres, 2000] Munkres, J. (2000). *Topology*. Featured Titles for Topology Series. Prentice Hall, Incorporated.

[Seifert, 1931] Seifert, H. (1931). *Konstruktion drei dimensionaler geschlossener Räume*. Hirzel.

[Van Kampen, 1933] Van Kampen, E. R. (1933). On the connection between the fundamental groups of some related spaces. *American Journal of Mathematics*, 55(1):261–267.

[Whitney, 1937] Whitney, H. (1937). On regular closed curves in the plane. *Compositio Mathematica*, 4:276–284.

[Whitney, 1944] Whitney, H. (1944). The self-intersections of a smooth n-manifold in 2n-space. *Ann. of Math*, 45(2):220–246.