Geometry & Topology 10 (2006) 2385–2429 **2385**

The global medial structure of regions in $\mathbb{R}^3$

JAMES DAMON

For compact regions Ω in $\mathbb{R}^3$ with generic smooth boundary $\mathcal{B}$, we consider geometric properties of Ω which lie midway between their topology and geometry and can be summarized by the term "geometric complexity". The "geometric complexity" of Ω is captured by its Blum medial axis M, which is a Whitney stratified set whose local structure at each point is given by specific standard local types.

We classify the geometric complexity by giving a structure theorem for the Blum medial axis M. We do so by first giving an algorithm for decomposing M using the local types into "irreducible components" and then representing each medial component as obtained by attaching surfaces with boundaries to 4–valent graphs. The two stages are described by a two level extended graph structure. The top level describes a simplified form of the attaching of the irreducible medial components to each other, and the second level extended graph structure for each irreducible component specifies how to construct the component.

We further use the data associated to the extended graph structures to express topological invariants of Ω such as the homology and fundamental group in terms of the singular invariants of M defined using the local standard types and the extended graph structures. Using the classification, we characterize contractible regions in terms of the extended graph structures and the associated data.

57N80; 68U05, 53A05, 55P55

Introduction

We consider a compact region $\Omega \subset \mathbb{R}^3$ with generic smooth boundary $\mathcal{B}$. We are interested in geometric properties of Ω and $\mathcal{B}$ which lie midway between their topology and geometry and can be summarized by the term "geometric complexity". Our goal in this paper is first to give a structure theorem for the "geometric complexity" of such regions. Second, we directly relate this structure to the topology of the region and deduce how the topology places restrictions on the geometric complexity and how the structure capturing the geometric complexity determines the topology of the region.

Published: 15 December 2006 DOI: 10.2140/gt.2006.10.2385

To explain what we mean by geometric complexity, we first consider $\mathbb{R}^2$. If $\mathcal{B}$ is a simple closed curve as in Figure 1 (a), then by the Jordan Curve and Schoenflies Theorems, Ω in Figure 1 (b) is topologically a 2–disk, and hence contractible. Hence, the region in Figure 1 (b) is topologically simple; however, it has considerable "geometric complexity". It is this geometric complexity which is important for understanding shape features for both 2 and 3 dimensional objects in a number of areas such as computer and medical imaging, biology, etc.

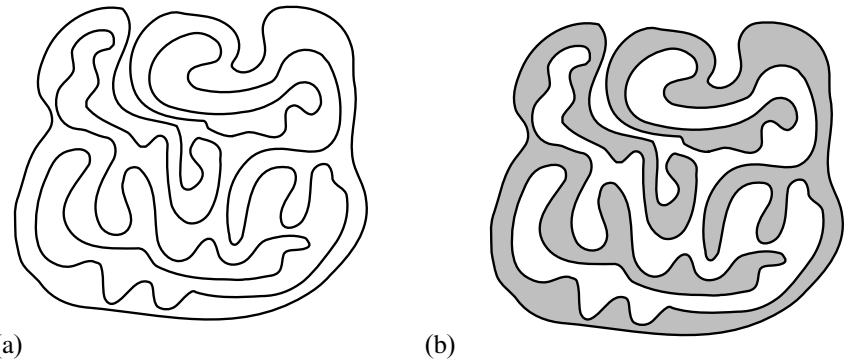

(a) (b)

Figure 1: Simple closed curve in (a) bounding a contractible region in $\mathbb{R}^2$ in (b)

This complexity is not captured by traditional geometrical invariants such as the local curvature of $\mathcal{B}$ nor by global geometric invariants given by integrals over $\mathcal{B}$ or Ω. Nor do traditional results such as the Riemann mapping theorem, which provides the existence of a conformal diffeomorphism of D^2 with Ω, provide a comparison of the geometric complexity of Ω with D^2. For 2–dimensional contractible regions, a theorem of Grayson [13], building on the combined work of Gage and Hamilton [9; 10], shows that under curvature flow the boundary $\mathcal{B}$ evolves so the fingers of the region shrink and the region ultimately simplifies and shrinks to a convex region which contracts to a "round point". Throughout this evolution, the boundary of the region remains smooth. The order of shrinking and disappearing of subregions provides a model of the geometric complexity of the region. Unfortunately, this approach already fails in $\mathbb{R}^3$, where the corresponding mean curvature flow may develop singularities, as in the case of the "dumbell" surface found by Grayson (also see eg Sethian [18]).

The structure theorem for the "geometric complexity" which we give is based on the global structure of the "Blum medial axis" M. The medial axis is a singular space which encodes both the topology and geometry of the region (see Figure 2). It is defined in all dimensions and has multiple descriptions including locus of centers of spheres in Ω which are tangent to $\mathcal{B}$ at two or more points (or have a degenerate tangency) as in Blum and Nagel [1], the shock set for the eikonal/ "grassfire" flow as

in Kimia, Tannenbaum and Zucker [14], which is also a geometric flow on Ω, and the Maxwell set for the family of distance to the boundary functions as in Mather [16]. It has alternately been called the central set by Yomdin [20], and has as an analogue the cut-locus for regions without conjugate points in Riemannian manifolds.

The multiple descriptions allow its local structure to be explicitly determined for regions in $\mathbb{R}^{n+1}$ with generic smooth boundaries: M is an n–dimensional Whitney stratified set (by Mather [16]) which is a strong deformation retract of Ω. The local structure of M is given by a specific list of local models, by Blum and Nagel for $n = 1$, by Yomdin [20] for $n \leq 3$ (where it is called the central set) and Mather [16] for $n \leq 6$, and is given a precise singularity theoretic geometrical description by Giblin [11] for $n = 2$. Results of Buchner [4] show the cut locus has analogous properties for regions without conjugate points (and our structure theorem extends to this more general case). Furthermore, by results in [6; 7; 8; 5], we can derive the local, relative, and global geometry of both Ω and $\mathcal{B}$ in terms of geometric properties defined on M.

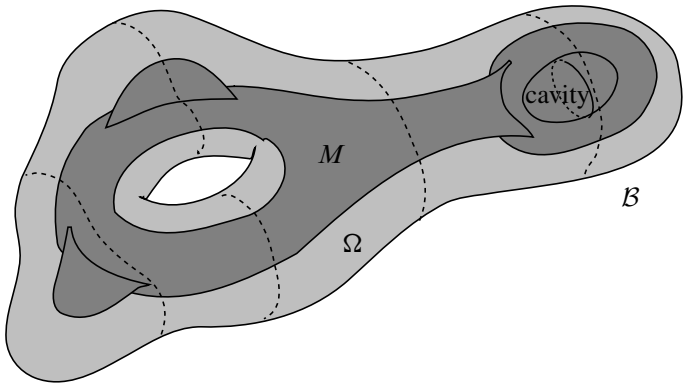

Figure 2: Blum medial axis for a region in $\mathbb{R}^3$

For example, for generic compact regions in $\mathbb{R}^2$, M is a 1–dimensional singular space whose singular points either have Y–shaped branching or are end points. This defines a natural graph structure with vertices for the branch and end points and edges representing the curve segments joining these points (see Figure 3). This graph structure encodes the geometric complexity of the region. Furthermore, Ω is contractible if and only if the graph is a tree. Then the tree structure can be used to contract the region (also for computer imaging a tree structure is a desirable feature, as trees can be searched in polynomial time). A computer scientist Mads Nielsen has asked whether in the case of contractible $\Omega \subset \mathbb{R}^3$, the Blum medial axis still has a "tree structure". We answer this question as part of the structure theorem for the medial axis M for generic regions.

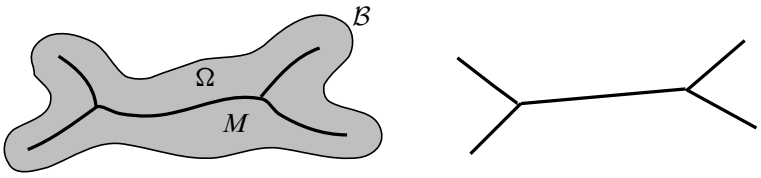

Figure 3: Blum medial axis for a region in $\mathbb{R}^2$ with associated graph structure

While the global structure of the medial axis as a Whitney stratified set with specific local singular structure does capture the geometric complexity of Ω, it is insufficient in this raw form to characterize the geometric complexity or to see directly its relation with the topology. This is what the structure theorem accomplishes.

We give a brief overview of the form of the structure theorem (Theorem 4.1).

The structure theorem for general regions

At the top level we decompose M into "irreducible medial components" M_i which are joined to each other along "fin curves". Geometrically this corresponds to taking connected sums of 3–manifolds with boundaries. A simplified form of the attaching is described by a *top level directed graph* $\Gamma(M)$. We assign vertices representing the M_i, and directed edges from M_i to M_j for each edge component of M_i attached to M_j along a fin curve. This is an "extended graph" in that there may be more than one edge between a pair of vertices or edges from a vertex to itself (see Figure 4). The resulting space obtained by attaching the M_i as indicated by the edges of $\Gamma(M)$ is the simplified form $\widehat{M}$ of the original M. It is homotopy equivalent to M and has the same irreducible components. From the topological perspective, the decomposition is optimal in that the homology and fundamental group of M are decomposed into direct sums, resp. free products, of those for the M_i. Also, M can be recovered from $\widehat{M}$ by certain "sliding operations along fin curves" according to additional data defined from M.

At the second level, we describe the structure of each irreducible medial component M_i by an extended graph $\Lambda(M_i)$, denoted more simply by Λ_i, as in Figure 5. Here, there is an analogy with the resolution graph of an isolated surface singularity, whose vertices correspond to complex curves, ie compact orientable real surfaces, (with data the genera and self-intersection numbers) and edges indicating transverse intersection. Alternately, we can view the vertices as denoting multiply-punctured real surfaces which are attached to the set of intersection points as indicated by the edges. For the irreducible medial components there is an analogous structure but with several added complications.

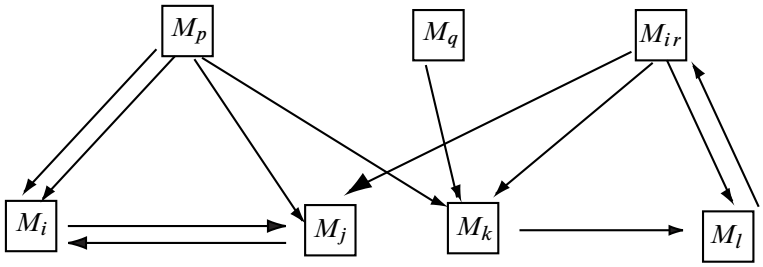

Figure 4: Graph structure given by decomposition into irreducible medial components with edges indicating the attaching along "fin curves"

The extended graph Λ_i corresponding to an irreducible medial component M_i has two types of vertices: "S–vertices" which are attached by edges to "Y–nodes". The S–vertices correspond to the medial sheets S_{ij}, which are (closures of) connected components of the set of smooth points of the component M_i (which are not in general the connected components of the set of smooth points of M). They are compact surfaces with boundaries (possibly nonorientable). The Y–nodes correspond to connected components $\mathcal{Y}_{ij}$ of the "Y–network" $\mathcal{Y}_i$ of M_i. The Y–network of M_i is the collection of "Y–branch curves" together with vertices which are the "6–junction points" where 6 sheets of the Blum medial axis come together in a point along 4 Y–branch curves (and is described by a 4–valent extended graph). Each edge from a vertex for S_{ij} to a node for $\mathcal{Y}_{ik}$ represents the attaching of the sheet S_{ij} along one of its boundary components to the component $\mathcal{Y}_{ik}$.

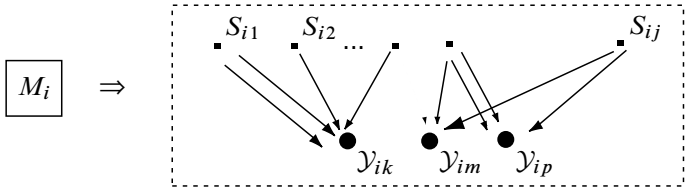

Figure 5: Extended graph structure for irreducible medial component M_i: S–vertices • representing connected medial sheets S_{ij} joined to Y–nodes ■ representing components $\mathcal{Y}_{ik}$ of the "Y–network"

There is associated data attached to the vertices and edges of Λ_i; namely, for S–vertices the genus, orientability, and number of medial edge curves on S_{ij}, for Y–nodes the 4–valent extended graph associated to $\mathcal{Y}_{ik}$, and for each edge the attaching data of the boundary component of S_{ij} to $\mathcal{Y}_{ik}$.

From this structure we obtain two alternate ways to view M_i: either as obtained by attaching compact surfaces with boundaries along (some of the) boundary components

or by an associated CW–decomposition. These allow us to relate the graph structures and their associated data for all of the M_i to the topological structure of Ω.

Topological invariants of M, and hence Ω, are computed by Theorem 3.2 and Theorem 3.4 in terms of M_i and $\Gamma(M)$. In turn, those of each M_i are expressed in terms of the extended graph Λ_i, and the associated data. This data includes the singular invariants of each M_i such as the numbers of 6–junction points, medial edge curves, components of $\mathcal{Y}_i$, the genera of S_{ij}, etc. From this we define a single "algebraic attaching map" for M_i. From this data, we compute in Theorem 3.2 the topological invariants of M_i such as homology groups and Euler characteristic. We then deduce in Theorem 3.4 the corresponding invariants for M and hence, Ω. Also, in Theorem 7.4 we give a presentation of $\pi_1(M_i)$ in terms of the associated data for Λ_i. These provide topological bounds on the geometric complexity of Ω.

The answer to Nielsen's question turns out in general to be no provided we wish to classify the medial axis up to homeomorphism. The tree has to be replaced by an Algorithm we give in Section 1 for decomposing the medial axis into the irreducible components M_i. The data of the algorithm requires more than just a graph showing which medial components are attached to which. However, the simplified version $\widehat{M}$, which sacrifices some of the detail of the attachings, is described by an extended graph $\Gamma(M)$, and for contractible regions $\Gamma(M)$ is a tree. Then the special form which the structure theorem takes for contractible regions in $\mathbb{R}^3$ (Theorem 5.2) gives a complete characterization of contractible regions by the following conditions: the extended graphs $\Gamma(M)$ and the Λ_i are trees, the medial sheets have genus 0 (and so are topological 2–disks with a finite number of holes) with at most one boundary curve representing a medial edge curve, a numerical "Euler relation" is satisfied which involves the basic medial invariants, and a fundamental group condition holds (Condition 5.1) which involves a space formed from $\mathcal{Y}_i$ and Λ_i.

The author is especially grateful to Mads Nielsen for initially raising the question about the tree structure for the contractible case, which led this investigator to these questions and results.

Acknowledgements This research was partially supported by grants from the National Science Foundation DMS-0405947 and CCR-0310546 and a grant from DARPA.

1 Decomposition into irreducible medial components

Generic local structure of Blum medial axis

We consider a region $\Omega \subset \mathbb{R}^3$ with generic smooth boundary $\mathcal{B}$ and Blum medial axis M. Then by Mather [16], M can be viewed as the Maxwell set for the family of

distance functions on $\mathcal{B}$; hence, for generic $\mathcal{B}$, it is a 2–dimensional Whitney stratified set. Also, the generic local structure has one of the following local forms in Figure 6 (see eg [11], where Giblin gives a very explicit geometric description).

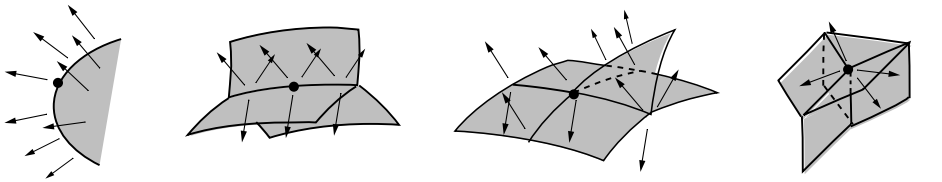

(a) edge (b) Y–branching (c) fin creation point (d) "6–junction"

Figure 6: Local generic structure for Blum medial axes in $\mathbb{R}^3$ and the associated radial vector fields to points of tangency on the boundary

Then M consists of the following: (i) smooth connected (2–dimensional) strata; 1–dimensional strata consisting of (ii) Y–junction curves along which three strata meet in a Y–branching pattern and (iii) edge curves consisting of edge points of M; 0–dimensional strata consisting of (iv) fin points and (v) 6–junction points, where six medial sheets meet along with 4 Y–junction curves. Connected components of Y–junction curves end either at fin points or 6–junction points; while edge curves only end at fin points. We refer to the union of Y–junction curves, fin points, and 6–junction points as the *initial Y–network* $\mathcal{Y}$.

In the simplest form we can view the medial axis as formed by attaching the connected 2–dimensional smooth strata to the 1–complex formed from the Y–branch and medial edge curves and the fin and 6–junction points. In fact, this approach misses a considerable amount of global structure for subspaces of M. To identify this larger structure we will decompose M into "irreducible medial components" M_i by "cutting M along fin curves". Then we further decompose the irreducible medial components M_i by representing them as obtained by attaching smooth medial sheets to the resulting Y–network formed from the union of Y–junction curves and 6–junction points in M_i. From this we will ultimately give a CW–decomposition to compute the topological invariants of M.

Cutting the medial axis along fin curves

In order to proceed, we first explain how we cut along fin curves.

Suppose we have on the Blum medial axis a fin point x_0. Then near x_0 we can distinguish the "fin sheet" which is the sheet that contains a medial axis edge curve ending at x_0. We can begin following the Y–branch curve from x_0, while keeping track of the fin sheet which remains a connected sheet as we move along the curve.

Eventually one of two things must happen: either we reach a 6–junction point or another fin point.

First, if the Y–branch curve meets a 6–junction point, then because the fin sheet is locally connected near the 6–junction point, we can follow the edge of the sheet as it continues through the 6–junction point. After the 6–junction point we have identified both the corresponding continuation of the Y–branch curve, and the sheet. We can do this for each 6–junction point it encounters. As M is compact, eventually the Y–branch curve must meet another fin point. We shall refer to the edge of the sheet from one fin point to the other as a *fin curve*.

At the end of the fin curve, what was identified as the fin sheet (close to the Y–branch curve) from the beginning may or may not be the fin sheet for the end point. If the sheet is a fin sheet at both ends, then we refer to the fin curve as being "essential", while if it is only a fin sheet at one end, then we refer to the fin curve as "inessential", (later discussion will explain the reason for these labels). Examples of these are shown in Figure 7.

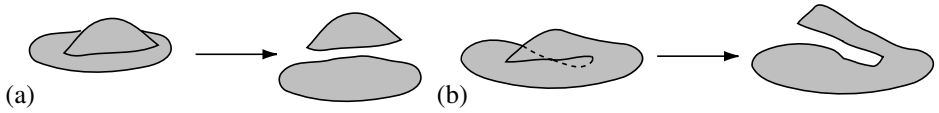

Figure 7: Two possibilities for fin curves on a medial sheet and the results of cutting along the fin sheets: (a) essential fin curve (b) inessential fin curve

An example of a region containing an inessential fin curve is given in Figure 8 and might be called a "Mobius board", a surf board but with a "Mobius band" twist.

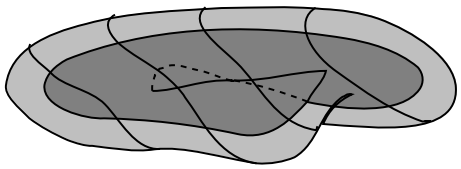

Figure 8: "Mobius board" with a "inessential fin curve"

Then to cut along a fin curve beginning from a fin point x_0, we locally disconnect the fin sheet from the direction of x_0 by locally adding a point of closure to the fin sheet for each point of closure on the fin curve (as followed from x_0). We must take care as a point on the fin curve can locally be a point of closure for more than one part of the sheet. After this step, locally at an added closure point, the fin sheet is now a surface with piecewise smooth boundary; see Figure 7. If the fin curve is essential, then the fin

sheet is now locally disconnected from the remaining two sheets still attached along the Y–junction curve. While if the fin curve is inessential, the fin sheet is still attached at the other fin point to the remaining two sheets; again see Figure 7 (b).

Then we can take the two remaining sheets still attached along that fin curve and smooth them to form a smooth sheet along the curve, with former 6–junction points on the fin curve becoming Y–branch points (for another Y–branch curve).

After having cut along the fin curves as described, the former points on the fin curve have become altered as follows:

(1) Fin points become points of an edge of the fin sheet.

(2) A former Y–branch point becomes a closure point on an edge of a fin sheet.

(3) A former 6–junction point become a (topological) Y–network point.

(4) A Y–branch point on a base sheet becomes a (topological) 2–manifold point of that sheet.

Because of (4) above, the end result depends upon a further distinction for essential fin curves. A *type*-1 *essential fin curve* will be one which only intersects other essential fin curves at 6–junction points; otherwise, it shares a segment of Y–branch curve with another essential fin curve, and it will be *type*-2 *essential fin curve* (see eg Figure 9). If we cut along a type-1 essential fin curve, then the fin sheet becomes disconnected from the other sheets (at least along the curve) and this does not alter any other essential fin curve. If we cut along a type-2 essential fin curve, then it will alter the structure of the other essential fin curves sharing a segment of Y–branch curve with it. Hence, we can prescribe the following algorithm for decomposing the medial axis.

Algorithm for decomposing medial axis into irreducible components

(1) Identify all type-1 essential fin curves and systematically cut along type-1 essential fin curves (it does not matter which order we choose as cutting along one does not alter the fin properties of another).

(2) After cutting along all type-1 essential fin curves, we may change certain inessential fin curves to type-1 essential ones. If so return to step (1).

(3) There only remain type-2 essential fin curves and inessential fin curves. Choose an essential fin curve and cut along it. If a type-1 essential fin curve is created, return to step (1). Otherwise, repeat this step until no essential fin curves remain.

(4) When there are no other essential fin curves, choose an inessential fin curve which crosses a 6–junction point, and cut it from one side until we cut across one 6–junction point.

(5) Check whether we have created an essential fin curve. If so then we cut along it, and repeat the earlier steps (1)–(3).

(6) If no essential fin curve is created, then we repeat step (3) until there are only inessential fin curves which do not cross 6–junction points.

(7) Finally we can cut each such remaining inessential fin curve, producing part of a smooth sheet as in Figure 7 (b).

(8) The remaining connected pieces are the "irreducible medial components" M_i of M.

Remark 1.1 The distinct connected pieces created following steps (1) and (2) are intrinsic to M; while those created using steps (3) and (4) are not because choices are involved. Which choices are made typically depends on the given situation and the importance we subjectively assign to how sheets are attached.

Example 1.2 In Figure 9 (a), we have a contractible medial axis with a pair of type–2 essential fin curves. Depending on which essential fin curve we choose, 1-2 or 3-4, we choose to cut along in step (3), we obtain either (b) or (c), which leads to different attachings (and hence top level graph) for the irreducible medial components. An alternate possibility would be to cut each fin sheet along the fin curves and view them as being attached partially along the edge of a fourth sheet. Again, the exact geometric form of M may suggest one choice being preferred over the others.

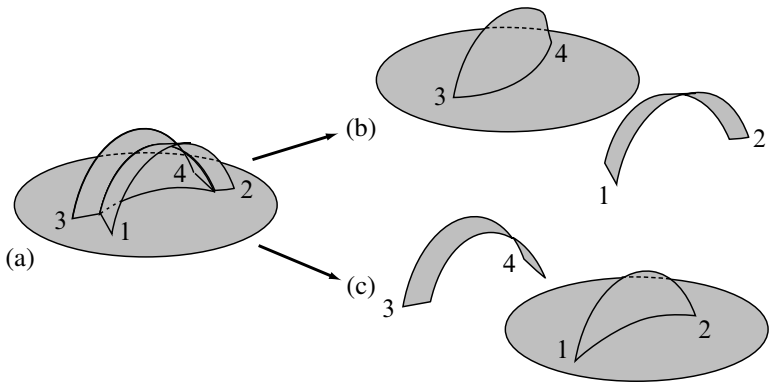

Figure 9: Nonuniqueness of medial decomposition resulting from type–2 essential fin curves: (a) is a contractible medial axis with only type–2 essential fin curves, and (b) and (c) illustrate the results from cutting along the fin curves 1-2 or 3-4.

Example 1.3 In Figure 10 (a), we have a contractible medial axis with 10 fin points 1-10, and all fin curves are inessential. Depending on how we choose cuts in step (4) of the algorithm, we can end up with 1, 2, or 3 irreducible medial components.

If we cut from 4 through the first 6–junction point, then 3-6 becomes an essential fin curve, and we cut away the fin sheet M_1 as in Figure 10 (b). Then further cutting from 8 through the first 6–junction point, we create another essential fin curve 2-9. Cutting along it creates a second fin sheet M_2. The remaining inessential fin curves 1-4, 5-7, and 8-10 can be contracted to points on edges of the third sheet M_3. Each of these 3 medial sheets are then irreducible components.

Alternatively, after the first cut, we could have instead cut from 9, and then from 4 again, and then only inessential fin curves remain without 6–junction points, so they contract to a second sheet, and we only obtain two irreducible components. Thirdly, we could have begun cutting from 7, then 8, and then 4 twice and we would obtain only a single medial sheet with inessential fin curves, leading to a single irreducible component.

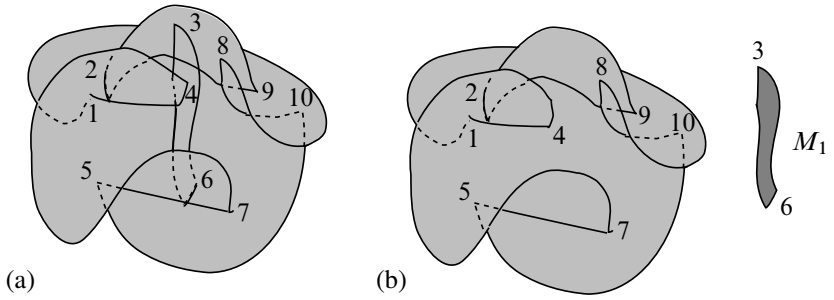

Figure 10: Nonuniqueness of medial decomposition resulting from inessential fin curves: (a) is a contractible medial axis with only inessential fin curves, and (b) illustrates the cutting of irreducible medial component M_1 after first cutting from fin point 4.

Constructing the medial axis by attaching irreducible medial components along fin curves

To reverse the algorithm and reconstruct M from the M_i requires the following additional attaching data:

(1) the list of which segments of medial edge curves which will be attached as either essential or inessential fin curves

(2) the order in which the sheets will be attached

(3) the curve at the j–th stage to which the attaching will be made as a fin curve to obtain the $(j+1)$–st stage.

The curves in step (3) may cross multiple components, and need only be well- defined up to isotopy; however, this isotopy is for the space obtained at the j–th stage, whose singular set must be preserved by the isotopy. Hence for example, in Figure 11, even though the medial axis is contractible, if the attaching curve were on the same side of the two holes then the resulting spaces would not be homeomorphic.

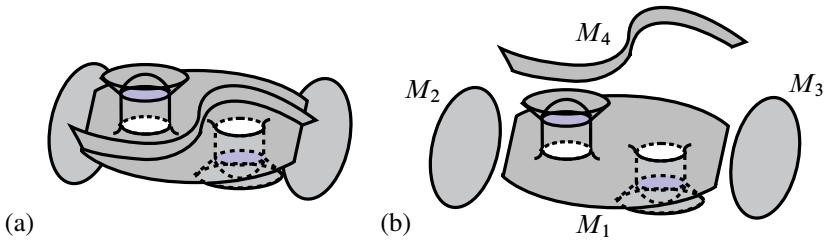

Figure 11: (a) A contractible medial axis, with attaching of component M_4 to multiple components M_1, M_2, M_3 (b) Irreducible medial components for (a)

2 Extended graph structure and medial decomposition

We next turn to a simplified version of M using the same irreducible medial components, but with simplified attaching. The resulting simplified version of M, which we denote by $\widehat{M}$, will still be homotopy equivalent to M. To define $\widehat{M}$, we consider an alternate way to understand the topological effects of attaching along a fin curve γ. Instead, we isotope the attaching map along the support of the original γ, by sliding and shrinking so it now becomes a fin curve γ' which no longer passes through any 6–junction points (which have disappeared after the isotopy as in Figure 12).

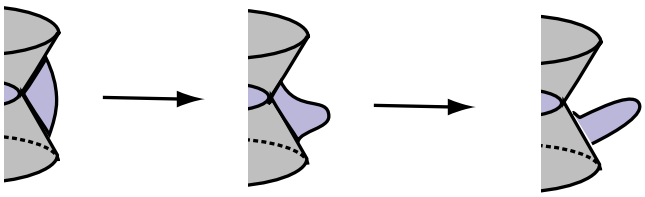

Figure 12: Isotoping a fin curve so it misses 6–junction points

We can achieve this by replacing M by a subcomplex which is a strong deformation retract, and only differs from M in a small neighborhood of the fin curve.

Lemma 2.1 (Isotopy lemma for fin curves) *Let γ be a fin curve in M. Then there is a Whitney stratified set $M' \subset M$, which is a strong deformation retract of M and which only differs from M in a given neighborhood of γ. Furthermore, in that neighborhood, γ has been replaced by a fin curve γ' which does not meet a 6–junction point.*

Proof We obtain M' as the result of a series of deformation retractions $M \supset M^{(1)} \supset M^{(2)} \supset \cdots \supset M^{(k)} = M'$. Here $M^{(j)} \supset M^{(j+1)}$ corresponds to either moving the fin point along a Y–branch curve so it is in a neighborhood of a 6–junction point where it has the normal form as in Figure 13 (a), or when $M^{(j)}$ already has this form then $M^{(j+1)}$ is a deformation retraction across a 6–junction point as in Figure 13 (b).

In the first case we may construct the deformation because M is analytically trivial along a Y–branch curve and so analytically a product in a neighborhood of a compact segment of a Y–branch curve. In the second case, we may use the normal form for 6–junction points to deform along the shown region. □

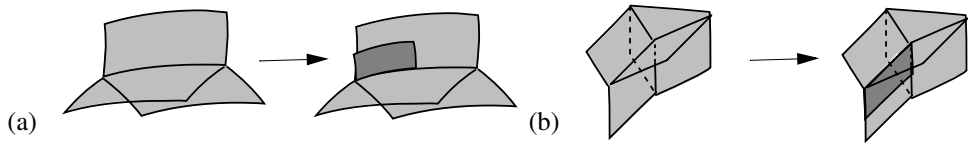

Figure 13: (a) Deformation retraction along a Y–branch curve. (b) Deformation retraction across a 6–junction point. In both cases the deformation is across the dark region.

We view this process as sliding and shrinking the fin curve along its support, across 6–junction points until the fin curve no longer crosses 6–junction points. Also, after isotoping the fin curves, the resulting spaces are still homotopy equivalent.

After first isotoping the fin curves, we can more clearly see the topological effect of cutting along fin curves.

Lemma 2.2 *In the preceding situation, let γ' be a fin curve which does not meet a 6–junction point. There are two possibilities.*

(1) *Suppose the fin curve γ' is essential, with fin sheet S_i at both ends. Then when we cut S_i along γ', the sheet locally becomes disconnected from the remaining sheets of the fin curve as in Figure 7 (a).*

(2) *If instead the fin curve γ' is inessential, then we cut all three sheets along the fin curve; we obtain a single sheet with boundary edge containing the fin curve as in Figure 7 (b). This may be alternately obtained by shrinking the fin curve to a point.*

Proof In the essential case, by the triviality of M along Y–branch sheets, the fin sheet remains a single sheet along the fin curve. When we cut it we disconnect it from the base sheet as in Figure 7 (a). By contrast, in the inessential case when we cut the fin sheet, the edge curve continues along the cut fin curve, joining up with another edge curve again. Hence, the fin points become edge points for a single edge curve.

As an alternate way to view the process, the fin sheet for one fin point becomes part of the base sheet. Then we may follow a path around the second fin point to end up on the base sheet for the first fin point. Then we may continue the curve around the first fin point to be on the sheet which becomes the fin sheet for the second fin point. If we extend these curves so they intersect the edge curve, then together they form a closed curve, so that when the fin curve is contracted to a point we obtain a 2–disk, one edge of which is the edge curve. □

Remark 2.3 This Lemma contrasts cutting along essential versus inessential fin curves. When we cut along the essential fin curve we disconnect the fin sheet from the base sheet, causing a change in homotopy type as in Figure 7 (a). By contrast, for inessential fin curves, when we contract the fin curve to a point p, the sheet then has an edge curve which passes through p. In this case, the "fin curve" can be eliminated, and the pair of fin points cancelled, without any change in the homotopy type of M.

Now beginning with a Blum medial axis, we can repeat the Algorithm given in Section 1, except at each step, instead of cutting along the fin curve, by the Isotopy Lemma for Fin Curves, we slide the fin sheet along the fin curve. We pass all of the 6–junction points (which then become Y–branch points), until the fin curve lies on a single smooth sheet. Then by Lemma 2.2 (b), we may contract the inessential fin curves eliminating their fin points. The resulting space with these simplified attachings is $\widehat{M}$. Finally we cut the essential fin curves using Lemma 2.2 (b). The resulting connected components M_i are again the irreducible medial components of M.

The initial Y–network has been altered by the removal of the fin curves to yield the Y–network $\mathcal{Y}$. Then $\mathcal{Y} = \bigcup_i \mathcal{Y}_i$, where each $\mathcal{Y}_i$ is the resulting Y–network for M_i. Then $\widehat{M}$ is obtained from the irreducible medial components by attaching their edges which came from essential fin curves to the isotoped positions in some M_j. There is some ambiguity in this construction in addition to that coming from the algorithm. We make a choice of the medial sheet that the fin curve passes through. However, the irreducible medial components and the homotopy type of $\widehat{M}$ will remain the same.

This gives rise to a "top level directed extended graph" $\Gamma(M)$ whose vertices correspond to the M_i, with an edge going from M_i to M_j for each segment of an edge curve of M_i attached to M_j along a fin curve.

There is the following relation between the topology of the full medial axis M and that of the irreducible medial components $\{M_i\}_{i=1}^r$ and the top level extended graph $\Gamma(M)$.

Proposition 2.4 *In the preceding situation*

(1) *for any coefficient group G,*

$$H_j(M;G) \simeq H_j(\Gamma(M);G) \oplus (\oplus_{i=1}^r H_j(M_i;G)) \quad \text{for all } j>0;$$

(2) $$\pi_1(M) \simeq \pi_1(\Gamma(M)) * (*_{i=1}^r \pi_1(M_i)).$$

Of course the graph $\Gamma(M)$ only contributes to homology in dimension 1. Although it is an extended graph, we can compute its fundamental group just as for graphs. To prove this we also need the next proposition.

Proposition 2.5 *If Γ is a nonempty connected extended graph, then it contains a maximal tree T. Then $\pi_1(\Gamma)$ is a free group with one generator for each edge (including loops) of Γ which do not belong to T.*

Proof of Proposition 2.5 First, we construct a connected graph Γ' from Γ. For each pair of vertices of Γ joined by an edge we remove all but one of the edges; as well we remove any loops (edges from a vertex to itself). What remains is a connected graph Γ'. Then such a graph has a maximal connected tree T (by eg Spanier [19, Chapter 2]). This is also a maximal connected tree for Γ.

There is only one special case which occurs for Y–networks, and hence which we must allow. It is a single S^1 without any vertex. Then we have to artificially introduce a vertex so S^1 becomes a loop on the vertex. Then the maximal tree is just the introduced vertex. However, we emphasize that we will not count such "artificial vertices" for later numerical relations involving the number of vertices.

Then by repeating the same proof in [19, Chapter 2] as for the fundamental groups of graphs, $\pi_1(\Gamma)$ is a free group with one generator for each edge (including loops) of Γ which do not belong to T. □

We obtain as an immediate corollary.

Corollary 2.6 *If M is contractible, then so is each M_i contractible, and furthermore, $\Gamma(M)$ is a tree.*

Proof of Corollary 2.6 If M is contractible, then M is connected, $\pi_1(M) = 0$, and $H_j(M) = 0$ for all $j > 0$. Thus, by (2) of Proposition 2.4, $\pi_1(\Gamma(M)) = 0$; thus, $\Gamma(M)$ is a tree. In addition, for each i, by the combination of (1) and (2) of Proposition 2.4, $\pi_1(M_i) = 0$, and $H_j(M_i) = 0$ for all $j > 0$. Also, by definition, each M_i is connected. Then by the Hurewicz Theorem, $\pi_j(M_i) = 0$ for all $j \geq 0$. However, M_i is a CW–complex, so by another Theorem of Hurewicz, M_i is contractible. □

Proof of Proposition 2.4 By our earlier discussion, M is homotopy equivalent to the space $\widehat{M}$ obtained by attaching M_i to M_j along an edge segment of M_i to a smooth sheet of M_j. However, this edge segment can be homotoped to a point so M is instead homotopy equivalent to the space obtained by attaching an external curve segment α_{ij} from x_{ij} on the edge of M_i to a point y_{ij} in a smooth sheet of M_j. Then from one fixed point x_{i0} on M_i we can choose disjoint curves β_{ij} from x_{i0} to the other points x_{ij}. Then we choose a maximal tree T in $\Gamma(M)$ and choose the attaching curves α_{ij} corresponding to the edges of T, as well as the curves β_{ij} in each M_i, which give a tree from each x_{i0}. The resulting space X which is the union of these trees and the M_i is homotopy equivalent to the pointed union of the M_i. Hence,

$$\pi_1(X) \simeq *_{i=1}^{r} \pi_1(M_i) \qquad \text{and} \qquad H_j(X) \simeq \oplus_{i=1}^{r} H_j(M_i) \text{ for } j > 0.$$

Then as in the computation of π_1 of a graph as in [19, Chapter 2], we inductively add one of the remaining edges and show using Seifert–Van Kampen that we are taking a free product with a free group on one generator. After k steps, where k is the number of additional edges (including loops) of $\Gamma(M)$ not in T, we obtain that $\pi_1(M)$ is isomorphic to the free product of $\pi_1(X)$ with a free group on k generators, which is exactly the fundamental group of $\Gamma(M)$.

For the case of homology we instead use Mayer–Vietoris in a similar argument. □

Proposition 2.4 is valid for any region Ω. Thus, we have a relation between the topology of M, and hence Ω (as M is a strong deformation retract of Ω), and the individual irreducible medial components M_i and the top level extended graph $\Gamma(M)$. In the special case that Ω is contractible, then so also is M and each M_i, and $\Gamma(M)$ is a directed tree.

Example 2.7 (Knot complement regions) The simplest example of contractible Ω corresponds to boundary $\mathcal{B} \simeq S^2$. A slightly more complicated example would have $\mathcal{B}$ consisting of two components S^2 and a torus T^2. In the case that T^2 belongs to

the compact region bounded by S^2 in $\mathbb{R}^3$, then Ω is the region inside S^2 and outside of the torus T^2. This is a "knot complement region".

If $K \subset \mathbb{R}^3$ is a knot, then the complement $C = \mathbb{R}^3 \backslash K$ is not a region in our sense. However, we can take a small tube T_ε around the knot and a large sphere S^2_R containing the tube. Then the region $\Omega = D^2_R \backslash \text{int}(T_\varepsilon)$ is a compact region with boundary $S^2_R \cup \partial T_\varepsilon$. Also, $\Omega \subset C$ is a strong deformation retract. For a generic knot with $\varepsilon > 0$ and R sufficiently large, the Blum medial axis of Ω will be generic.

To see what form the irreducible medial components $\{M_i\}$ take for such a knot complement region, we note a few results from knot theory. First, by Alexander duality, $H_i(\Omega) \simeq \mathbb{Z}$ for $i < 3$ and 0 otherwise. Also, $\pi_1(C)$ does not split as a free product of nontrivial groups. Hence, by Proposition 2.4 there are just three possibilities: (i) $\pi_1(\Gamma(M)) \neq 0$ so all M_i are simply connected; or $\Gamma(M)$ is a tree and all but one of the M_i, say M_1, are simply connected; and then either (ii) $H_2(M_1) \simeq \mathbb{Z}$, or (iii) $H_2(M_1) = 0$ and for a single other M_i, say M_2, $H_2(M_2) \simeq \mathbb{Z}$.

In the first case, we must have $\pi_1(C) \simeq \pi_1(\Gamma(M)) \simeq \mathbb{Z}$, which implies K is the unknot. In this case, all but one of M_i, say M_1, have trivial reduced homology and hence are contractible, while M_1 is simply connected with 0 reduced homology except for $H_2(M_1) \simeq \mathbb{Z}$. Then $\Gamma(M)$ is homotopy equivalent to S^1, M_1 is homotopy equivalent to S^2, and M (and hence C) is homotopy equivalent to the pointed union $S^1 \vee S^2$ (which is homotopy equivalent to a standard homotopy model for the complement of the unknot, the torus with a disk attached to one of the generators of π_1).

In the second case, $\Gamma(M)$ is a tree and all of the other M_i except M_1 have trivial fundamental group and reduced homology, and hence are contractible. There is then exactly one noncontractible irreducible medial component M_1, and C is homotopy equivalent to it.

In the third case, all but M_1 and M_2 are contractible, M_2 is homotopy equivalent to S^2, and $\Gamma(M)$ is a tree. Then C is homotopy equivalent to $M_1 \vee M_2 \simeq M_1 \vee S^2$.

In all of these cases, all but one or two M_i are contractible and so contribute nothing to the topology of the region but do contribute to its geometric complexity. Next, we shall determine the structure of the M_i and relate this to their topology.

3 Structure of irreducible medial components

Next, we decompose each irreducible medial component M_i, which is connected and locally has only Y–junction points, edge points, and 6–junction points (but no fin

points). We may apply the local closure procedure used in Section 1 to cut the sheets of M_i along the remaining Y–network as in Figure 14. We do this by locally adding closure points to each of the three sheets meeting at a point of the Y–network, or to each of the 6 sheets meeting at 6–junction points. This disconnects the connected components of the set of smooth points of M_i from the Y–network, turning them into compact connected surfaces with piecewise smooth boundaries. We refer to each such compact surface with boundary S as a *medial sheet* of M_i. The data attached to each medial sheet consists of the triple (g, o, e). Here g is the genus of S (after attaching disks to the boundary components), $o = 1$ or 0 corresponding to whether S is orientable or nonorientable (a nonorientable surface with at least one boundary component can be embedded in $\mathbb{R}^3$), and e denotes the number of boundary components of S (ie S^1's) which are medial edge curves of M (keeping in mind that part of the edge might have originally been an essential fin curve).

Then each M_i is formed by attaching certain of the boundary components of medial sheets in M_i to the Y–network $\mathcal{Y}_i$ of M_i, reversing the cutting process above. The Y–network $\mathcal{Y}_i$ has connected components $\{\mathcal{Y}_{ik}\}_{k=1}^{c_i}$, and we denote the collection of medial sheets of M_i by $\{S_{ij}\}_{j=1}^{s_i}$.

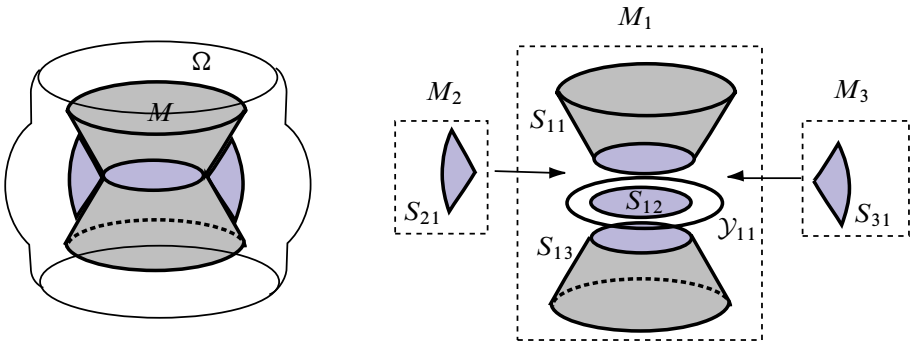

Figure 14: Cutting the medial axis first along fin curves and then along the Y–network $\mathcal{Y}$

We assign an extended graph Λ_i to each such irreducible medial component M_i as follows. The graph will have two type of vertices: to each sheet S_{ij} will be associated an S–vertex, and to each connected component $\mathcal{Y}_{ik}$ we associate a vertex which we refer to as a Y–node. We will associate an edge from the S–vertex associated to S_{ij} to the Y–node corresponding to $\mathcal{Y}_{ik}$, for each boundary component of S_{ij} (ie an S^1) which is attached to $\mathcal{Y}_{ik}$.

Each medial sheet S_{ij} has the associated data (g, o, e). Second, each connected component $\mathcal{Y}_{ik}$ of the Y–network $\mathcal{Y}_i$ can be described by an extended graph Π_{ik}.

Its vertices correspond to the 6–junction points of $\mathcal{Y}_{ik}$, and there are distinct edges between vertices which correspond to the distinct Y–junction curves of $\mathcal{Y}$ joining the corresponding 6–junction points. Because a 6–junction point has exactly four Y–junction curves ending at it, each vertex of the graph will have four edges. Thus, each vertex will have a valence of 4, and we will refer to such an extended graph as a "4–valent extended graph".

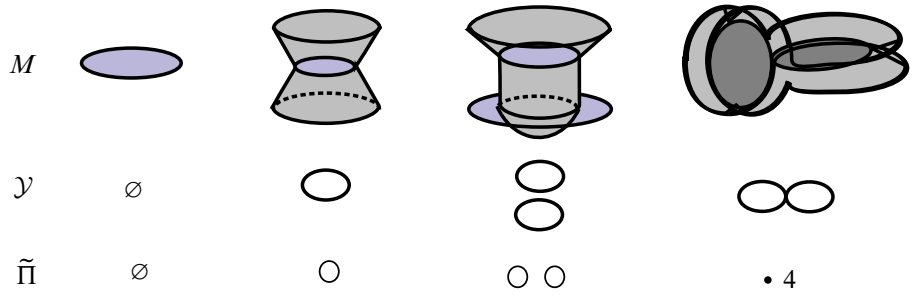

Figure 15: Simplest examples of irreducible medial components, their Y–networks and the reduced network graphs

A basic result we need for 4–valent extended graphs is a consequence of the following result for m–valent extended graphs.

Proposition 3.1 *If Π is a nonempty connected m–valent extended graph with k vertices, then (m or k must be even and) $\pi_1(\Pi)$ is a free group on $\frac{k}{2}\cdot(m-2)+1$ generators.*

Hence, for a nonempty connected 4–valent graph Π with k vertices, $\pi_1(\Pi)$ is a free group on $k+1$ generators.

Proof An extended graph Π is still a 1–complex. Hence, $\pi_1(\Pi)$ is a free group on ℓ generators, where $\ell = \mathrm{rk}H_1(\Pi)$. If V, respectively E, denotes the number of vertices, respectively edges, of Π, then by the m–valence of Π, $2E = mV$. As $V = k$,

$$\begin{aligned}\ell &= 1-\chi(\Pi) = 1-(V-E)\\ &= 1-(k-\frac{k}{2}\cdot m)\end{aligned}$$

which equals the value claimed. □

Associated weighted graphs of extended graphs

In order to work with an extended graph Γ, we simplify the description by defining an associated weighted graph $\widetilde{\Gamma}$. $\widetilde{\Gamma}$ will have the same vertices as Γ; however, for each pair of vertices with at least one edge between them, we remove all but one of the edges. Likewise, we remove all loops. Then we assign to each vertex v_i the number 2ℓ where there are ℓ loops at v_i. Also for the edge between vertices v_i and v_j, we assign the integer m which is the number of edges in Γ between v_i and v_j. Then the valence at v_i is $2\ell + \sum m_i$ where we sum m_i over the edges of $\widetilde{\Gamma}$ which end at v_i. If for a vertex the number of loops is 0 then we suppress it; likewise if there is only 1 edge between two vertices, we usually suppress the 1. For extended graph Π_{ik} we refer to this graph $\widetilde{\Pi}_{ik}$ as the *reduced Y–network graph* of $\mathcal{Y}_{ik}$. For such a reduced graph, at a vertex the sum of the vertex number and edge numbers always equals 4. There are very few possibilities for vertices as shown in Figure 16

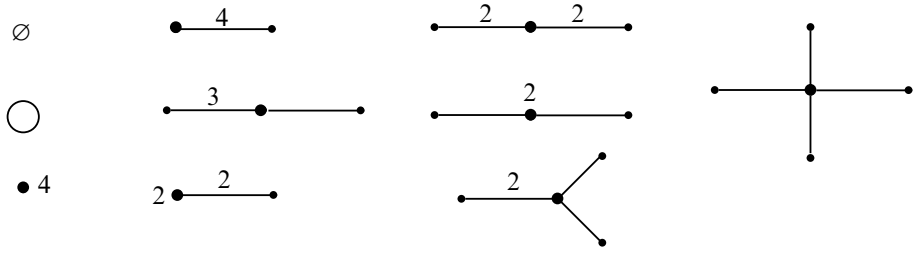

Figure 16: The possible local vertex structures (at the enlarged vertex) for reduced Y–network graphs. The case of S^1 is the exceptional case of nonempty Y–network without any vertices.

Finally, we assign this data associated to M_i to the graph Λ_i as follows. The data of S_{ij} will be attached to its S–vertex, the graph Π_{ik} (or its reduced graph $\widetilde{\Pi}_{ik}$) assigned to the Y–node of $\mathcal{Y}_{ik}$, and the topological attaching data will be assigned to the edge from S_{ij} to $\mathcal{Y}_{ik}$.

Next, just as we did for M, we compute the topological invariants of each irreducible medial component M_i. This time it will be in terms of the extended graph Λ_i and the associated data.

Associated numerical data for irreducible medial components

The numerical invariants can be divided into those that are invariants of the graph Λ_j, and those which are given by the associated data attached to the vertices and edges of Λ_j. In general, for an invariant I_j of the medial component M_j, we let $I = \sum_j I_j$, summed over the irreducible medial components of M, be the associated

global invariant of M. The only exception is our definition of $\nu(M)$ given by (3–2). We summarize these invariants in Table 1.

Graph-theoretic numerical data of Λ_j The key invariants associated to Λ_j are: the number s_j of S–vertices, the number c_j of Y–nodes and the first Betti number λ_j of Λ_j (as a topological space).

We can express other invariants such as the (reduced) Euler characteristic in terms of these. For example, the number of edges of $\Lambda_j = s_j + c_j + \lambda_j - 1$. We introduce another purely graph-theoretic invariant which appears in computing $\chi(M_j)$:

$$\nu(\Lambda_j) \quad = \quad s_j - c_j - \lambda_j. \tag{3–1}$$

Using this we also introduce the global graph-theoretic invariant of M:

$$\nu(M) \quad = \quad \sum_j \nu(\Lambda_j) - \beta_1 \tag{3–2}$$

where the sum is over the irreducible medial components of M and β_1 is first Betti number of $\Gamma(M)$ (as a topological space and $= -\widetilde{\chi}(\Gamma(M))$).

Numerical associated data of Λ_j Second, we list the numerical data defined from the associated data. First of all, an S–vertex is associated to S_{jk}, which is a compact surface with boundary. The boundary will have connected components which are either edge curves or are attached to the Y–network $\mathcal{Y}_j$. We let e_{jk} denote the number which are medial edge curves of M_j. For each connected component $\mathcal{Y}_{jk}$ of the Y–network, we let v_{jk} denote the number of vertices in Π_{jk}. This is the number of 6- -junction points in the Y–network $\mathcal{Y}_j$.

Next, S_{jk} has a genus in the sense that if we attach a 2–disk to each boundary component, then we obtain a compact surface (without boundary). By the classification of surfaces, if it is orientable, it is homeomorphic to a connected sum of tori and the number is its genus; while if the surface is nonorientable, then it is a connected sum of real projective planes, and that number is the genus. In either case we denote the genus by g_{jk}, and refer to it as the genus of S_{jk}. We define the *weighted genus of* S_{jk} to be

$$\widetilde{g}_{jk} \quad = \quad \begin{cases} 2g_{jk} & \text{if } S_{jk} \text{ is orientable} \\ g_{jk} & \text{if } S_{jk} \text{ is nonorientable} \end{cases} \tag{3–3}$$

Then the *total weighted genus of* M_j

$$G_j \stackrel{\text{def}}{=} \sum_k \widetilde{g}_{jk} \quad \text{summed over all medial sheets } S_{jk} \text{ of } M_j. \tag{3–4}$$

Invariants of medial components	Global invariants of medial axis	Description of medial component invariants
Graph invariants		
s_j	s	number S–vertices of Λ_j = number medial sheets of M_j
c_j	c	number Y–nodes of Λ_j = number connected components of $\mathcal{Y}_j$
λ_j	λ	rk $H_1(\Lambda_j)$
$\nu(M_j)$	$\nu(M)$	$\nu(M_j)$ given by (3–1); $\nu(M)$ given by (3–2)
Graph data invariants		
v_j	v	number vertices of Π_j = number 6–junction points of M_j
e_j	e	number medial edge curves of M_j
G_j	G	total weighted genus of M_j (3–4)
q_j	q	sum of q_{jk} given by (3–5)
Q_j	Q	first Betti number of associated 1–complex $\mathcal{SY}'_j \subset M_j$

Table 1: Invariants I_j of Medial Components M_j and their Global Versions I for the medial axis, where $I = \sum_j I_j$.

From the weighted genus, we define q_{jk} as follows.

$$q_{jk} = \begin{cases} \widetilde{g}_{jk} + e_{jk} - 1 & \text{if } e_{jk} > 0 \\ \widetilde{g}_{jk} & \text{if } e_{jk} = 0 \end{cases} \tag{3–5}$$

Note this is the first Betti number of a 1–complex in S_{jk} which will be used in Sections 6 to 8 to compute the homology and fundamental group of M_j. Then we define

$$q_j = \sum q_{jk} \quad \text{summed over all medial sheets } S_{jk} \text{ in } M_j.$$

Although s_j, the number of medial sheets in M_j, is a graph theoretic invariant, there are subsets of S–vertices whose definition depends on associated data. Of these the most important is s_{0j}, *the number of sheets without edge curves*. This is a sum $s_{0o,j} + s_{0n,j}$,

where $s_{0o,j}$, resp. $s_{0n,j}$, denotes the number of orientable, resp. nonorientable, sheets without medial edge curves.

Then by (3–3), (3–4) and (3–5), there is the following relation between these invariants:

$$q_j \;=\; G_j \,+\, e_j \,-\, (s_j - s_{0j}) \tag{3–6}$$

Topological invariants of irreducible medial components

We next describe how to determine the Euler characteristic, homology, and fundamental group of M_j. In Section 8, we introduce a "minimal 1–complex" $\mathcal{SY}'_j$ constructed from the singular data of M_j. It consists of three contributions: the graph structure Λ_j, the Y–network $\mathcal{Y}_j$, and a 1–skeleton from the medial sheets. The $\mathrm{rk}H_1(\mathcal{SY}'_j)$ is given by $Q_j = \lambda_j + (v_j + c_j) + q_j$, which is the sum of the three contributions. Then in Section 8, we define an *algebraic attaching homomorphism*

$$\Psi_j\colon \mathbb{Z}^{s_{0j}} \to \mathbb{Z}^{Q_j}.$$

From this homomorphism we can compute the homology of M_j in terms of the singular data. Second, we shall also define in Section 7 a set of generators and relations to compute the fundamental group of M_j. These provide two ingredients for the topology of M_j.

Theorem 3.2 *Suppose that M_j is an irreducible medial component. Then there are the following properties in homology:*

(1) *M_j has torsion free homology.*

(2) *The reduced Euler characteristic of M_j is given by*

$$\widetilde{\chi}(M_j) \quad = \quad s_{0j} - Q_j.$$

Alternately, it can be written in terms of the extended graph Λ_j, and associated data of Λ_j:

$$\begin{aligned} \widetilde{\chi}(M_j) \quad &= \quad v(\Lambda_j) \,-\, (G_j + e_j + v_j) \\ &= \quad s_j - (e_j + v_j + c_j + G_j + \lambda_j). \end{aligned} \tag{3–7}$$

(3) *The homology groups $H_2(M_j)$, respectively $H_1(M_j)$, are the kernel and cokernel of the algebraic attaching homomorphism Ψ_j; see* (8–4)*. Furthermore, there are bounds*

$$\mathrm{rk}\,(H_2(M_j;\mathbb{Z})) \le s_{0o,j} \quad \text{ and } \quad q_j \le \mathrm{rk}\,(H_1(M_j;\mathbb{Z})) \le Q_j - s_{0n,j}.$$

In addition, there are the properties for the fundamental group:

(4) *There is a continuous map* $\psi_j\colon M_j \to \Lambda_j$ *such that the induced map* $\pi_1(M_j) \to \pi_1(\Lambda_j)$ *is surjective.*

(5) *The fundamental group has a presentation by generators and relations*

$$\pi_1(M_j) \quad \simeq \quad F_{Q_j}/\langle\{r_{ji}\}\rangle$$

where the relations r_{ji} *are given by* (7–3).

Proof For (1), it is sufficient to show $H_*(M)$ is torsion free; for $H_*(M_j)$ is a direct summand by Proposition 2.4. Then since $H_*(M) \simeq H_*(\Omega_0)$ where $\Omega_0 = \Omega\backslash\mathcal{B}$, it is sufficient to show $H_*(\Omega_0)$ is torsion free. We may decompose $\mathbb{R}^3\backslash\mathcal{B} = \Omega_0 \cup \Omega'$ as a disjoint union. Thus, $H_*(\Omega_0)$ is a direct summand of $H_*(\mathbb{R}^3\backslash\mathcal{B})$. However, by Alexander duality $H_j(\mathbb{R}^3\backslash\mathcal{B}) \simeq H^{2-j}(\mathcal{B})$ for $j > 0$. Since $\mathcal{B}$ is a compact orientable surface, $H^*(\mathcal{B})$ is torsion free, giving the conclusion.

The derivation of the formula for the Euler characteristic, and the computation of the homology in terms of the algebraic attaching homomorphism are given in Section 8.

The computation of the fundamental group is given in Section 7, using the CW–decomposition of M_j in terms of the singular structure given in Section 6. □

These results provide the following restrictions on geometric complexity.

Corollary 3.3 *Consider an irreducible medial component* M_j.

First, suppose $H_1(M_j) = 0$. *Then we have the following:*

(1) *The graph* Λ_j *is a tree.*

(2) *Each medial sheet* S_{jk} *has genus* 0, *ie it is a* 2*–disk with a finite number of holes.*

(3) *At most one of the boundary components of* S_{jk} *is an edge curve of* M_j.

(4) *Suppose, in addition,* $H_2(M_j) = 0$. *Then* e_j *is the number of medial sheets with edge curves, and the following "Euler Relation" holds between the number of sheets without edge curves (LHS) and the (RHS), which is* $\operatorname{rk} H_1(\mathcal{Y}_j)$, *a topological invariant of the Y–network*

(3–8) $$s_j - e_j \quad = \quad v_j + c_j.$$

(5) *If* M_j *is simply connected (without requiring* $H_2(M_j) = 0$*), then* M_j *will satisfy the fundamental group relation given in Condition 5.1.*

(6) *Finally, if* M_j *is contractible, then all of the preceding hold.*

The derivation of the Corollary from Theorem 3.2 will be given in Section 8

Finally, we combine the results of Theorem 3.2 with the decomposition in Proposition 2.4 to compute the *homology and reduced Euler characteristic of M from the singular invariants measuring geometric complexity*. We let

$$\Psi \;=\; \oplus_j \Psi_j : \mathbb{Z}^{s_0} \to \mathbb{Z}^{Q},$$

where $s_0 = \sum_j s_{0j}$ is the total number of medial sheets in all of the M_j without edge curves (and $Q = \sum_j Q_j$).

Theorem 3.4 *Let M be the Blum medial axis of a connected region Ω with smooth generic boundary $\mathcal{B}$. Then $H_*(\Omega) \simeq H_*(M)$ have the following properties.*

(1) *M (and Ω) have torsion free homology.*

(2) $H_2(M;\mathbb{Z}) \simeq \ker(\Psi)$ *and* $H_1(M;\mathbb{Z}) \simeq \operatorname{coker}(\Psi) \oplus H_1(\Gamma(M);\mathbb{Z})$.

(3) $rk(H_2(M;\mathbb{Z})) \;\le\; s_{0o}$ *and* $q + \beta_1 \;\le\; rk(H_1(M;\mathbb{Z})) \;\le\; Q - s_{0n} + \beta_1$, *where s_{0o}, resp. s_{0n} denote the number of orientable, resp. nonorientable, medial sheets of M without edge curves and as before $\beta_1 = rk(H_1(\Gamma(M);\mathbb{Z}))$.*

(4) *The reduced Euler characteristic is given by*

$$\widetilde{\chi}(M) \quad = \quad \nu(M) - (G + e + v) \quad = \quad s - (e + v + c + G + \beta)$$

where $\beta = \beta_1 + \sum_j \lambda_j$.

(5) *In the case Ω (and hence M) is contractible, $G = 0$, $\beta = 0$, and we have the relation*

$$s - e \quad = \quad c + v.$$

Proof The proof follows from M being a strong deformation retract of Ω (see eg [6]), using the sum formula given in Proposition 2.4, and substituting the formulas for $H_i(M_j;\mathbb{Z})$ and $\widetilde{\chi}(M_j)$ given in Theorem 3.2. □

Example 3.5 The simplest irreducible medial component M_i has $\mathcal{Y}_i = \varnothing$. Then M_i is a compact connected surface with boundary so $s_i = 1$ and $e_i =$ number of boundary components. The Y–network invariants $c_i = v_i = 0$ and the graph Λ_i consists of a single S–vertex so $\lambda_i = 0$. Then $G_i = Q_i = q_i$ is the rank of $H_1(M_i)$, and $\nu(M_i) = 1$. If $e_i = 0$ then M_i must be orientable, $s_{0i} = 1$, and Ψ_i is the zero map; while if $e_i > 0$ then $s_{0i} = 0$ and Ψ_i is again the zero map.

Example 3.6 The next simplest case has $\mathcal{Y}_i = \circ$. Already there arises a question of which nonintersecting embedded surfaces can be the medial sheets. The easiest case is when the medial sheets are 2–disks and annuli (examples in Figure 2, Figure 14 and the first three examples of Figure 15). However, an example exists of a single medial sheet which is annulus with a single medial edge curve to form a "collapsed version of the umbilic bracelet" as in Figure 17.

4 General structure theorem

We can now state the full structure theorem for the Blum medial axis. and in the next section state the form it takes for contractible Ω.

Theorem 4.1 (General structure theorem) *Suppose $\Omega \subset \mathbb{R}^3$ is a compact connected region with generic smooth boundary $\mathcal{B}$ and Blum medial axis M. Then there is associated to Ω a two level graph structure.*

(1) *At the top level $\Gamma(M)$ is a directed extended graph. It consists of vertices corresponding to the irreducible medial components M_i of M and directed edges from M_i to M_j corresponding to each connection attaching an edge of M_i to M_j along a fin curve.*

(2) *To each M_j is associated a graph Λ_j with two types of vertices: S–vertices ■ corresponding to smooth medial sheets S_{jk} of M_j, and nodes • corresponding to the connected components $\mathcal{Y}_{jk'}$ of the Y–network of M_j. There is an edge from S_{jk} to $\mathcal{Y}_{jk'}$ for each boundary S^1 of S_{jk} which is attached to $\mathcal{Y}_{jk'}$.*

(3) *Each medial sheet S_{jk} is compact surface with boundary (possibly nonorientable). Each Y–network component $\mathcal{Y}_{jk'}$ can be described as a 4–valent extended graph $\Pi_{jk'}$ (or by the corresponding reduced graph $\widetilde{\Pi}_{jk'}$).*

(4) *At the third level, the extended graph Λ_j has data assigned to the S–vertices, Y–nodes, and edges. To an S–vertex for S_{jk} is assigned the data (g, o, e) indicating the genus g, $o = 1$ or 0 denoting orientability or nonorientability and e the number of Blum edge curves. To a Y–node for $\mathcal{Y}_{jk'}$ is assigned the 4–valent extended graph $\Pi_{jk'}$. To an edge from S_{jk} to $\mathcal{Y}_{jk'}$, the topological attaching data of the boundary circle of S_{jk} to $\mathcal{Y}_{jk'}$.*

(5) *Furthermore, the graph theoretic data of each Λ_j, the number of edge curves, 6–junction points and the weighted total genus satisfy the equation* (3–7).

(6) *The fundamental groups of each irreducible medial component are given in terms of generators and relations by* (7–2), *and the homology groups of each irreducible medial component are torsion free and given as the kernel and cokernel of the algebraic attaching map* (8–4).

Proof We have established the decomposition into irreducible medial components, which in turn can be represented by the attaching of medial sheets to the Y–network.

The computation of the fundamental group in terms of generators and relations will be given in Section 7 based on the CW–decomposition given in Section 6, and the computation of both the homology and the reduced Euler characteristic will be given in Section 8. □

Several examples we have already seen illustrate the general structure theorem.

Example 4.2 In Figure 11, M is formed from 4 medial components, of which M_2, M_3, and M_4 are topologically 2–disks ($\mathcal{Y} = \varnothing$ so $c_i = v_i = 0$, and $s_i = e_i = 1$). If we slide M_4 onto M_1, then the graph $\Gamma(M)$ is a tree. M_1 is formed from 5 medial sheets, all of genus 0: two are 2–disks, 2 are annuli, and one is a 2–disk with two holes. Except for the annuli, each of the other 3 medial sheets of M_1 have single medial edge curves. The Y–network is empty except for M_1 where it consists of 2 disjoint $\circ$. Λ_1 is a tree, and the attaching of the 2–disks for M_1 kill the two free generators for $\pi_1(\mathcal{Y}_1) \cup \Lambda_1$. Hence, M_1 is simply- connected. Also, we see $\nu(M_1) = 5-2-0 = 3$ and the other $\nu(M_i) = 1$. Hence, $\widetilde{\chi}(M) = (3 + 3\cdot 1 + 0) - (0 + 3 + 3\cdot 1 + 0) = 0$.

Thus, $H_2(M_1) = 0$, and M_1 and hence all of the M_i and M are contractible.

Example 4.3 The medial axis shown in Figure 18 is irreducible so $M = M_1$. There are five medial sheets ($s = 5$) all of which have genus 0: 3 annuli (two of which have a medial edge curve) and 2 disks (so $G = 0$, $e = 2$, $q = -1$ and the Y–network consists of 2 disjoint $\circ$ ($c = 2$, $v = 0$). Also, Λ_j is a tree so $\lambda = 0$. As above, the attaching of the 2- disks for M_1 kill the two free generators for $\pi_1(\mathcal{Y}_1) \cup \Lambda_1$ so $M = M_1$ is simply connected. We see $\nu(M) = 5-2+0 = 3$. Thus, $\widetilde{\chi}(M) = 3-(0+2+0) = 1$. Hence, $H_2(M_1) = \mathbb{Z}$, and $M = M_1$ is homotopy equivalent to S^2.

5 Structure theorem for contractible regions

For contractible regions Ω, the structure theorem for the Blum medial axis takes a considerably simplified form. We saw in Corollary 3.3 that the medial sheets and extended graphs had to have a special form. As well, we define in Section 7 a set of elements r_{ij} in $\pi_1(\mathcal{S}\mathcal{Y}_i')$ corresponding to medial sheets S_{ij} without edge curves; see (7–3). The computation of the fundamental group in Theorem 7.4 leads to the following condition.

Fundamental Group Condition 5.1 *For an irreducible medial component M_i, the set of elements*

$$\{r_{ij} : \text{all } j \text{ for which } S_{ij} \text{ is without edge curves of } M_i\}$$

form a set of generators for $\pi_1(\mathcal{SY}_i')$.

Taken together the conditions on the extended graphs being trees, the numerical "Euler Relation" and the preceding Condition 5.1 completely characterize contractible regions.

Theorem 5.2 (Structure theorem for contractible regions) *Suppose $\Omega \subset \mathbb{R}^3$ is a contractible bounded region with generic smooth boundary $\mathcal{B}$ and Blum medial axis M. Then there is associated to Ω a multilevel directed tree structure which determines the simplified structure $\widehat{M}$ associated to M.*

(1) *At the top level, $\Gamma(M)$ is a directed tree consisting of vertices corresponding to the irreducible medial components M_i of M, and there are directed edges from M_i to M_j corresponding to the attaching of an edge of M_i to M_j along a fin curve in $\widehat{M}$.*

(2) *At the second level, to each M_j is associated a directed tree Λ_j with two types of vertices: S–vertices ■ corresponding to smooth medial sheets S_{jk} of M_j, and Y–nodes • corresponding to the connected components $\mathcal{Y}_{jk'}$ of the Y–network of M_i. There is an edge from S_{jk} to $\mathcal{Y}_{jk'}$ if a boundary "circle" of S_{jk} is attached to $\mathcal{Y}_{jk'}$. Given S_{jk} and $\mathcal{Y}_{jk'}$, there is at most one such boundary "circle".*

(3) *Each medial sheet S_{jk} is topologically a 2–disk with a finite number of holes. At most one of the boundary circles of a medial sheet is an edge curve of the M_j. Each Y–network component $\mathcal{Y}_{jk'}$ can be described as a 4–valent extended graph $\Pi_{jk'}$ (or by the corresponding reduced graph $\widetilde{\Pi}_{jk'}$.*

(4) *At the third level, the tree Λ_j has data assigned to the S–vertices, Y–nodes, and edges. To an S–vertex for S_{jk} is assigned (h, e) indicating the number of holes and Blum edge curves. To a Y–node for $\mathcal{Y}_{jk'}$ is assigned the 4–valent extended graph $\Pi_{jk'}$. To an edge from S_{jk} to $\mathcal{Y}_{jk'}$, the topological attaching data of the single boundary circle of S_{jk} to $\mathcal{Y}_{jk'}$.*

(5) *Furthermore, for an irreducible medial component M_j, the number of medial sheets, edge curves, 6–junction points and the total number of connected components of the Y–network $\mathcal{Y}_i$ satisfy the Euler relation (3–8), and each M_j satisfies the fundamental group relation Condition 5.1.*

Conversely, suppose we are given a bounded region Ω in $\mathbb{R}^3$ with smooth generic boundary and Blum medial axis M so that: the top level graph $\Gamma(M)$ is a tree; for each irreducible medial component M_j, the graph Λ_j is a tree; the medial sheets are topologically 2–disks with a finite number of holes, having at most one boundary circle an edge curve of M_j; the numerical invariants of M_j satisfy (3–8)*; and each M_j satisfies Condition 5.1. Then M and (hence) Ω are contractible.*

Proof If Ω is contractible, then so is M which is a strong deformation retract of Ω. Then by Corollary 2.6, $\Gamma(M)$ is a tree and each M_j is contractible. Then by Corollary 3.3, each Λ_j is a tree, and each medial sheet S_{jk} is topologically a 2–disk (ie has genus 0), with a finite number of holes. At most one of the boundary circles of a medial sheet is an edge curve of the M.

Furthermore, $\widetilde{\chi}(M_j) = 0$, yielding the Euler relation (3–8). As $\pi_1(M_j) = 0$, by Theorem 7.4, Condition 5.1 must hold.

Conversely, by Theorem 7.4, Condition 5.1 implies that $\pi_1(M_j) = 0$ for each j. Hence, $H_1(M_j;\mathbb{Z}) = 0$. As the graphs Λ_j are trees, $\lambda_j = 0$, and all sheets have $q_j = 0$. Thus, the Euler condition implies by (3–7) that $\widetilde{\chi}(M_j) = 0$. Hence, as $H_2(M_j;\mathbb{Z})$ is torsion free, $H_2(M_j;\mathbb{Z}) = 0$. Thus, as $\Gamma(M)$ is a tree, by Proposition 2.4, M is simply connected and $\widetilde{H}_*(M;\mathbb{Z}) = 0$. As M is a CW–complex, a theorem of Hurewicz implies M is contractible. □

Example 5.3 (Simplest examples) As the simplest examples illustrating the structure theorem, we consider those for which the second level trees are $\Pi_j = \varnothing$ or $\circ$.

(1) "Simple contractible examples" Each $\Pi_j = \varnothing$ so each M_i is topologically a 2–disk. This is the simplest type of region. For example, in Figure 9 and Figure 10, the medial axis is decomposed into irreducible medial components of this type. In computer imaging, the M–rep structure of Pizer and coauthors [17] is based on the region having such a simple medial structure.

(2) The second type of regions would allow both $\Pi_j = \varnothing$ or $\circ$ and the only type of medial sheets are 2–disks and annuli. Besides the example in Figure 2, we also see the example in Figure 14 and the first three examples of Figure 15 are of this type

(3) Figure 17 is a collapsed version of the "umbilic torus" investigated independently by Helaman Ferguson (who has produced sculptures of it) and Christopher Zeeman who called it the "umbilic bracelet". The discriminant for real cubic binary forms is a cone on this space. The "umbilic torus/bracelet" is obtained by rotating the hypocycloid of three cusps about an external axis in its plane, while rotating by $\frac{2\pi}{3}$. If we collapse the hypocycloid onto the three line skeleton as in Figure 17 (a), and rotate then we

obtain Figure 17 (b). It consists of an annulus attached to an inner Y–junction curve and an edge curve. It has $\Pi_j = \circ$ and a single edge curve, and illustrates how only a single medial sheet may meet the Y–junction curve.

(4) The next case would have a reduced network graph of the form "$\bullet$ 4" as for the fourth example of Figure 15, which is contractible.

It appears to be a rather difficult question to determine exactly when examples such as (3) or (4) can be included as part of larger medial structures with special properties such as being contractible.

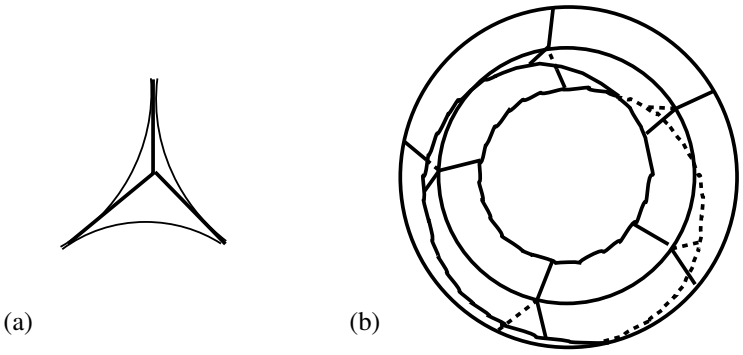

Figure 17: (b) Skeleton of an "umbilic torus/bracelet", obtained by rotating the three intersecting lines in (a) around an axis while rotating by $\frac{2\pi}{3}$

Example 5.4 (Tree structure for a noncontractible region) A region can have tree structures at each of the two levels and consist of medial sheets of genus 0 attached as prescribed by Theorem 5.2 yet not be contractible. An example is given by Figure 18. M is not contractible, and it is only the Euler relation that fails: $s = 5$, $e = 2$, while $m = 2$ and $v = 0$, so $s - e = 3 \neq 2 = v + m$.

6 Representations of the topological structure of the medial sheets

To determine the topology of an irreducible medial component M_j, we give two alternate representations of it as either built up by attaching surfaces with boundaries or by attaching cells. These give two alternate useful ways to decompose M_i to compute homology and the fundamental group. We introduce two subspaces $\mathcal{SY}_j \subset \mathcal{SY}'_j \subset M_j$, such that $\mathcal{SY}_j$ will contain both $\mathcal{Y}_j$ and a subspace homotopy equivalent to the extended graph Λ_j.

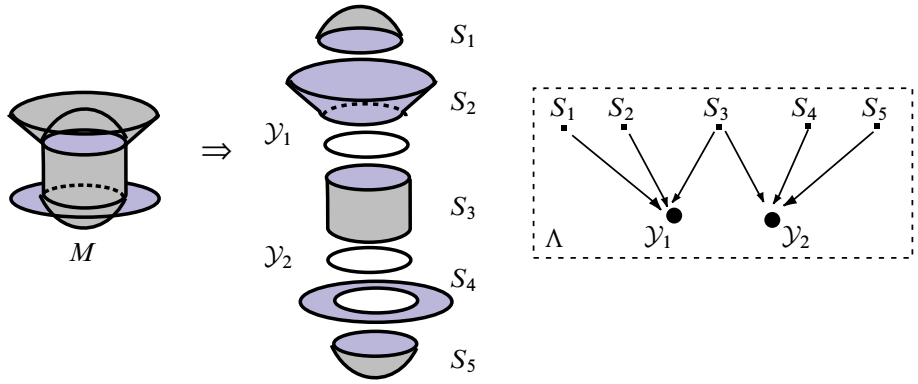

Figure 18: Medial axis M of a noncontractible region with tree structures for $\Gamma(M)$ and Λ

To do this we consider a medial sheet S_{jk} of M_j. We may use the standard representation for compact surfaces as quotients of the 2–disk D^2 after appropriately identifying edges of the boundary as in Massey [15]. In our case, S_{jk} is a quotient of a 2–disk D^2 with a finite number of holes obtained by removing the interiors B_i of embedded 2–disks $\overline{B}_i, i = 1, \dots \ell_{jk} + e_{jk}$. Here ℓ_{jk} is the number of the disk boundaries which will be attached to $\mathcal{Y}_j$, and the other e_{jk} will remain edge curves in M_j. We may assume that for example, the edges of the boundary of D^2 are identified as given by the classification of surfaces; see eg Figure 19 (a).

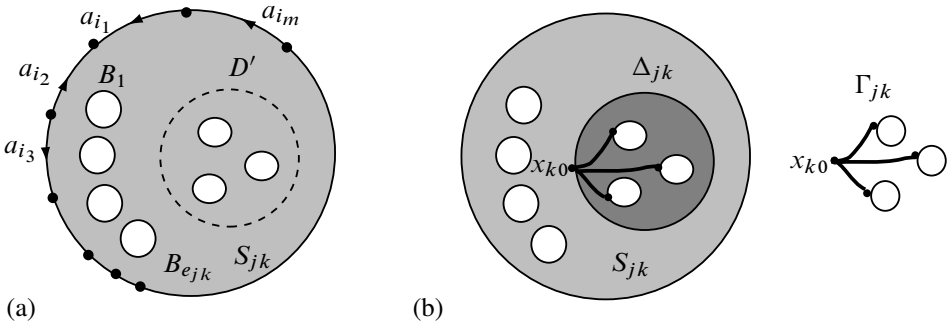

Figure 19: (a) S_{jk} identified as the quotient of a 2–disk with a finite number of holes by identifying the edges of the boundary; (b) Edges of holes to be attached to $\mathcal{Y}$ contained in Δ, and deformation retract Γ_{jk}.

Thus, $S_{jk} = (D^2 \backslash (\bigcup_i B_i))/\sim$. To be specific, we suppose that the boundaries of $\overline{B}_i, i = 1, \dots \ell_{jk}$ are the ones which will be attached to $\mathcal{Y}_j$. Up to isotopy, we may assume that $\overline{B}_i, i = 1, \dots \ell_{jk}$ are contained in a 2–disk D' as shown in Figure 19 (b).

We let $\Delta_{jk} = D' \backslash (\bigcup_{i=1}^{\ell_{jk}} B_i)$. We may construct curves β_i in Δ_{jk} from a point x_{j0} on the boundary of Δ_{jk} to boundary points x_{ji} of each B_i, so the curves are disjoint from each other except at x_{j0}.

We let Γ_{jk} denote the union of the curves β_i and the union of the boundary edges $\partial \overline{B}_i, i = 1, \dots \ell_{jk}$. Then Γ_{jk} is a strong deformation retract of Δ_{jk}. We let

$$\mathcal{S}\mathcal{Y}_j \;=\; (\mathcal{Y}_j \cup (\textstyle\bigcup_k \Delta_{jk}))/\sim \qquad \text{and } \Gamma_j \;=\; (\textstyle\bigcup_k \Gamma_{jk})/\sim$$

where $\sim$ denotes the equivalence for attaching the interior boundary edges in the Γ_{jk} to $\mathcal{Y}_j$. Since every edge of $\mathcal{Y}_j$ is the image of one of the attached edges of some S_{jk}, we see that $\mathcal{Y}_j \subseteq \Gamma_j \subset \mathcal{S}\mathcal{Y}_j$. Also, Γ_j is a strong deformation retract of $\mathcal{S}\mathcal{Y}_j$ by retracting each Δ_{jk} to Γ_{jk} leaving $\mathcal{Y}_j$ fixed.

We next let $$\widetilde{S}_{jk} \;=\; S_{jk} \backslash \text{int}(\Delta_{jk}).$$

Then we can attach each $\widetilde{S}_{jk}$ to $\mathcal{S}\mathcal{Y}_j$ along C_{jk} the outer boundary edge of Δ_{jk}. Attaching all such $\widetilde{S}_{jk}$ recovers M_j. We summarize these statements with the following lemma.

Proposition 6.1 *Let M_j be an irreducible medial component. Then the following hold:*

(1) *There is a subspace $\Gamma_j \subset \mathcal{S}\mathcal{Y}_j$ of M_j which contains the Y–network $\mathcal{Y}_j$.*

(2) *Γ_j is a 1–complex which is a strong deformation retract of $\mathcal{S}\mathcal{Y}_j$.*

(3) *If we collapse the connected components of $\mathcal{Y}_j$ in Γ_j to points, then the resulting quotient space is homeomorphic to the second level extended graph Λ_j.*

(4) *There are compact surfaces with boundaries $\widetilde{S}_{jk}$ obtained from the medial sheets S_{jk} by removing Δ_{jk}, a 2–disk with holes whose edges are identified to $\mathcal{Y}_j$. Attaching the $\widetilde{S}_{jk}$ to $\mathcal{S}\mathcal{Y}_j$ recovers M_j, up to homeomorphism.*

We shall refer to this as the *modified medial sheet representation*.

CW–decomposition of M_j

We further simplify the preceding, by introducing an intermediate space $\mathcal{S}\mathcal{Y}'_j$.

We divide the medial sheets into two types: those with edge curves and those without. If S_{jk} has an edge curve, then in the earlier notation it has $e_{jk} > 0$ edge curves (each surrounding a hole) in the 2–disk. We denote these by E_{jk}. We can choose $e_{jk} - 1$ disjoint paths τ_{ki} from x_{k0} to points on the distinct edge curves E_{ki} for $i < e_{jk}$. We

can then form $\widetilde{E}_{ki} = E_{ki} \cup \tau_{ki}$. Likewise, we choose a curve τ_{k0} from x_{k0} to the boundary of D^2, not intersecting the paths τ_{ki}. We can form loops σ_{ji} centered at x_{k0} by following τ_{ki}, then going around the i–th edge curve counterclockwise, and then following τ_{ki} backwards back to x_{k0} (see Figure 20).

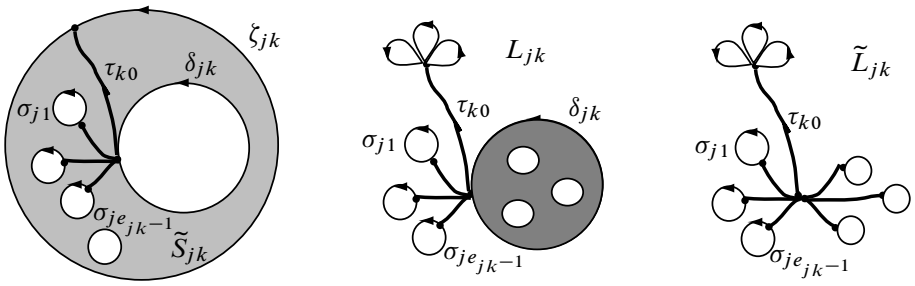

Figure 20: The spaces (a) $\widetilde{S}_{jk}$, (b) L_{jk} and (c) $\widetilde{L}_{jk}$

We let $$L_{jk} \;=\; \Delta_{jk} \underset{k<e_{jk}}{\cup} \widetilde{E}_{jk} \cup \tau_{k0} \cup (\partial D^2/\sim)$$

and $\widetilde{L}_{jk}$ is defined as L_{jk} except we replace Δ_{jk} by Γ_{jk}. Here "$\sim$" denotes the identification on the outer boundary of the 2–disk. In the case that S_{jk} has no edge curves of M_j, then

$$L_{jk} \;=\; \Delta_{jk} \cup \tau_{k0} \cup (\partial D^2/\sim).$$

We define

(6–1) $$\mathcal{SY}'_j \;=\; \mathcal{SY}_j \cup_k L_{jk}$$ where the union is over all medial sheets in M_j.

We describe the relation between M_j and $\mathcal{SY}'_j$ in two steps. First, we relate S_{jk} and L_{jk} (and $\widetilde{L}_{jk}$) and then deduce the consequences for M_j and $\mathcal{SY}'_j$.

In the case S_{jk} has no edge curves, S_{jk} is homeomorphic to the space obtained by attaching a 2–disk to L_{jk} as follows.

Lemma 6.2 *With the preceding notation:*

(1) *If S_{jk} has no edge curves, then S_{jk} is homeomorphic to the space obtained by attaching a 2–disk to L_{jk} along the path given in* (6–2).

(2) *If S_{jk} has edge curves, then L_{jk} is a strong deformation retract of S_{jk}.*

(3) *$\widetilde{L}_{jk}$ is a strong deformation retract of L_{jk}.*

(4) *Furthermore, L_{jk} is homotopy equivalent (rel Δ_{jk}) to $\Delta_{jk} \vee \bigvee_i S^1$ (or equivalently $\widetilde{L}_{jk}$ is homotopy equivalent (rel Γ_{jk}) to $\Gamma_{jk} \vee \bigvee_i S^1$), where $\bigvee_i S^1$ is a bouquet of q_{jk} S^1's (with q_{jk} defined by equation* (3–3)*).*

Proof If S_{jk} has no edge curves then we proceed as follows to recover S_{jk} from L_{jk} by attaching a 2–disk. We cut the annulus $\widetilde{S}_{jk} = S_{jk} \backslash \mathrm{int}(\Delta_{jk})$ along τ_{k0}. We obtain a topological 2–disk. Its boundary is mapped to L_{jk} by first following τ_{k0}, then the boundary of the 2–disk D^2 in a counter clockwise direction, then backwards along τ_{k0} and finally around the boundary C_{jk} of Δ_{jk} but in a clockwise direction. We denote the counterclockwise path around ∂D^2 by ζ'_{jk}, and the counterclockwise path around the boundary Δ_{jk} by δ_{jk}. Then a 2–disk is attached along its boundary to L_{jk} by the path

$$\xi_{jk} = \zeta_{jk} * \overline{\delta}_{jk} \quad \text{where} \quad \zeta_{jk} = \tau_{k0} * \zeta'_{jk} * \overline{\tau}_{k0} \tag{6–2}$$

followed by the identification (as is common, $\overline{\alpha}$ denotes the inverse of a path α).

Next, suppose instead S_{ji} has $e_{jk} > 0$ edge curves. Then we cut $\widetilde{S}_{jk} = S_{jk} \backslash \mathrm{int}(\Delta_{jk})$ along the paths τ_{k0} and τ_{ki} for $i < e_{jk}$. The resulting space is topologically an annulus with one boundary the edge curve $E' \stackrel{\text{def}}{=} E_{je_{jk}}$ of the e_{jk}–th hole. Then we may retract the annulus onto L_{jk} by retracting along curves from the boundary edge of E' as shown, followed by the identification. This gives L_{jk} as a strong deformation retract of S_{jk}.

Since this fixes Δ_{jk}, we may then further retract Δ_{jk} onto Γ_{jk}, while fixing the rest of L_{jk}. This gives $\widetilde{L}_{jk}$ as a strong deformation retract of L_{jk}.

Lastly, the boundary of D^2 after identification is a bouquet of S^1's, where the number is $2g_k$ if S_{jk} is orientable, and g_k if S_{jk} is nonorientable.

Finally, returning to L_{jk}, if we collapse all of the paths τ_{k0} and τ_{ki}, $i < e_{jk}$, to the point x_{j0}, then $\bigcup_{i<e_{jk}} E_{ji} \cup \tau_j \cup (\partial D^2/\sim)$ becomes a bouquet of q_{kj} S^1's. This collapsing induces a projection map from L_{jk} to $\Delta_{jk} \vee \bigvee_{i=1}^{q_{kj}} S^1$ which is the desired homotopy equivalence. An analogous argument works for the space $\widetilde{L}_{jk}$. □

Since the deformation retraction for L_{jk} (resp. $\widetilde{L}_{jk}$) in (3) of Lemma 6.2 fixes Δ_{jk} (resp. Γ_{jk}), we may apply the Lemma to each medial sheet with an edge curve after it is attached to $\mathcal{Y}_j$. If we let $\mathcal{SY}'_j$ denote $\mathcal{SY}_j$ with all of these medial sheets attached, then even though each medial sheet S_{jk} contracts to a bouquet of S^1's attached to Δ_{jk}, as $\mathcal{SY}_j$ is path connected, up to homotopy we may bring these bouquets together and conclude the following "CW–decomposition of M_j".

Proposition 6.3 *For an irreducible medial sheet M_j, we have the following:*

(1) *M_j is homotopy equivalent to a subspace M'_j obtained from $\mathcal{SY}'_j$ by attaching for each sheet S_{jk} without edge curve a 2–disk along the closed path given by* (6–2).

(2) *$\mathcal{SY}'_j$ is homotopy equivalent (rel $\mathcal{SY}_j$) to $\big(\bigvee_{i=1}^{q_j} S^1\big) \vee \mathcal{SY}_j$, and hence to the 1–complex $\big(\bigvee_{i=1}^{q_j} S^1\big) \vee \Gamma_j$ (here q_j is defined in Section 3).*

Proof By Lemma 6.2, we can first strong deformation retract each sheet S_{jk} with an edge curve onto L_{jk}. Then we can obtain each sheet S_{jk} without an edge curve by attaching a 2–disk to L_{jk}. Because this is done while fixing $\mathcal{SY}'_j$, we may do this all at once obtaining the strong deformation retract M'_j.

Also, by Lemma 6.2, each L_{jk} has as a strong deformation retract $\widetilde{L}_{jk}$ (rel Γ_{jk}), which is homotopy equivalent (rel Γ_{jk}) to $\big(\bigvee_{i=1}^{q_{jk}} S^1\big) \vee \Gamma_{jk}$. Again doing this for each L_{jk} gives $\big(\bigvee_{i=1}^{q_j} S^1\big) \vee \Gamma_j$ as a strong deformation retract of $\mathcal{SY}'_j$. Thus, M_j is homotopy equivalent to the space obtained from $\big(\bigvee_{i=1}^{q_j} S^1\big) \vee \Gamma_j$ by attaching one 2–disk for each medial sheet S_{jk} without edge curves. □

7 Fundamental group of irreducible medial components

To compute the fundamental group of an irreducible medial component M_j, we use the CW–complex decomposition we have just given in Section 6. We successively use the form of the fundamental group for Y–network components given in Proposition 3.1 and that of the graph Λ_j to compute the fundamental group of $\mathcal{SY}_j$. We then compute the changes obtained by attaching the L_{jk} to give that for $\mathcal{SY}'_j$, and lastly determine the effect of attaching the 2–disks corresponding to S_{jk} without edge curves.

We first compute the fundamental group of $\mathcal{SY}_j$ as follows.

Proposition 7.1 *With the preceding notation,*

$$\pi_1(\mathcal{SY}_j) \;\simeq\; \pi_1(\Gamma_j) \;\simeq\; \pi_1(\Lambda_j) * (*_{i=1}^{r} \pi_1(\mathcal{Y}_{ji})) \tag{7–1}$$

where there are r components for $\mathcal{Y}_j$.

Proof By the preceding remarks, it is sufficient to compute $\pi_1(\Gamma_j)$, viewing Γ_j as an extended graph. We choose a maximal tree T_{jk} for each component $\mathcal{Y}_{jk}$ and a maximal tree R_j for Λ_j. Furthermore, as the image of an edge $\partial \bar{B}_{ji}$ will cover completely each edge of $\mathcal{Y}_{jk}$ which it meets, we may choose the points x_{jk} so their images lie in T_{jk}, eg they are vertices y_m.

We also have to treat the special case where one of the $\mathcal{Y}_{jk}$ is an $\circ$ without any vertices. As in the proof of Proposition 2.5, we choose a point in $\circ$ as an artificial vertex, so the $\circ$ can be viewed as a loop on the vertex and then the vertex itself becomes the maximal

tree. As earlier, we do not count such "vertices" in determining v_j for numerical relations.

Then we let $\widetilde{\Gamma}_{jk}$ denote the union of the curves β_i in Γ_{jk} and form

$$\widetilde{\Gamma}_j \;=\; (\textstyle\bigcup_k \widetilde{\Gamma}_{jk} \cup_i T_{ji})/\sim .$$

Then each tree T_{jk} is contractible and if we collapse each T_{jk} to a point then the quotient is isomorphic to the extended graph Λ_j.

We may lift the edges of the tree R_j to a collection of edges in $\widetilde{\Gamma}_j$, by which we mean we can choose curves β_i in the $\widetilde{\Gamma}_{jk}$ which correspond to the edges of R_j. We denote the union of those curves by R'_j. Then we form

$$\mathcal{T}_j \;=\; (R'_j \cup (\textstyle\bigcup_i T_{ji}))/\sim$$

where now "$\sim$" only involves attaching the points x_{jk} of R'_j to the corresponding points y_m of the T_{ji}.

Lemma 7.2 *$\mathcal{T}_j$ is a maximal tree in the extended graph $\widetilde{\Gamma}_j$.*

Proof First, $\mathcal{T}_j$ is a tree because each T_{ji} is contractible and if we collapse each to a point we obtain R_j which is contractible; hence, $\mathcal{T}_j$ is contractible and is a graph.

To see $\mathcal{T}_j$ is maximal in $\widetilde{\Gamma}_j$, we note any other edge of $\widetilde{\Gamma}_j$ either corresponds to an edge of one of the $\mathcal{Y}_{jk}$, or to another edge of Λ_j under the collapsing map. In the first case, if the edge η is in $\mathcal{Y}_{jk}$, then as T_{ji} is maximal $T_{ji} \cup \eta$ is no longer simply connected; thus nor will be $\mathcal{T}_j \cup \eta$. If instead the added edge η maps to another edge η' of Λ_j under the collapsing map, then $R_j \cup \eta'$ is no longer simply connected. Then we may lift a noncontractible loop γ of $R_j \cup \eta'$ to a loop in $\mathcal{T}_j \cup \eta$ (joining edges of γ lift to edges which share a common Y–node, say $\mathcal{Y}_{jk}$, and hence can be connected by a path in T_{ji}), destroying the simple connectivity of $\mathcal{T}_j$. Thus, $\mathcal{T}_j$ is maximal. □

Then by Proposition 2.5, $\pi_1(\mathcal{S}\mathcal{Y}_j)$ is a free group with a free generator for each edge of $\widetilde{\Gamma}_j$ not in $\mathcal{T}_j$. There are λ_j edges of Λ_j, and $v_{jk}+1$ edges for each $\mathcal{Y}_{jk}$ by Proposition 3.1. Thus, there are a total of $\lambda_j + v_j + c_j$ free generators (recall $v_j = \sum_k v_{jk}$ and c_j is the number of connected components $\mathcal{Y}_{jk}$ of $\mathcal{Y}_j$). Such a free group is the free product of the free groups on $v_{jk}+1$ generators which are isomorphic to $\pi_1(\mathcal{Y}_{jk})$, and a free group on λ_j generators which is isomorphic to $\pi_1(\Lambda_j)$. □

The $\pi_1(\mathcal{Y}_{jk})$ are naturally identified as subgroups of $\pi_1(\mathcal{S}\mathcal{Y}_j)$. Also, the collapsing map $\varphi\colon \widetilde{\Gamma}_j \to \Lambda_j$ is a homotopy equivalence; hence, $\pi_1(\widetilde{\Gamma}_j) \simeq \pi_1(\Lambda_j)$. Then the inclusion $\widetilde{\Gamma}_j \subset \mathcal{S}\mathcal{Y}_j$, identifies a subgroup isomorphic to $\pi_1(\Lambda_j)$.

The second step is to compute the fundamental group of $\mathcal{SY}'_j$. It is homotopy equivalent to $\left(\bigvee_{i=1}^{q_j} S^1\right) \vee \mathcal{SY}_j$. Hence, again by the Seifert–Van Kampen theorem,

$$\pi_1(\mathcal{SY}'_j) \;\simeq\; \pi_1(\mathcal{SY}_j) * F_{q_j}$$

where F_{q_j} is a free group on q_j generators corresponding to the generators of $\bigvee_{i=1}^{q_j} S^1$ obtained from the E_{jk} and the identifications on the boundaries $\partial D^2/\sim$. Hence,

Corollary 7.3 *$\pi_1(\mathcal{SY}'_j)$ is a free group on $Q_j(=\lambda_j+v_j+c_j+q_j)$ generators.*

Computing the fundamental group of an irreducible medial component

We now use the proceeding to compute the fundamental group of M_j by determining the effect of attaching each medial sheet S_{jk}.

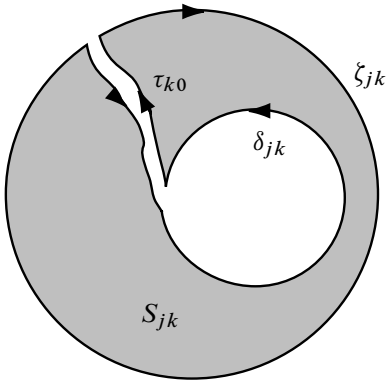

Figure 21: Representing S_{jk} by attaching a 2–disk to $\mathcal{SY}'_j$

Then we can compute the fundamental group of M_j by attaching the remaining 2–disks to $\mathcal{SY}'_j$; see Figure 21. First, we recall the loops at x_{k0}: ζ_{jk} for each medial sheet without edge curves, and δ_{jk} (obtained by following C_{jk} in the counterclockwise direction). We again let $\mathcal{T}_j$ denote a maximal tree in $\widetilde{\Gamma}_j$ and we choose a base point $x_0 \in \widetilde{\Gamma}_j$, and for each x_{k0}, a path α_k in the tree $\mathcal{T}_j$ from x_0 to x_{k0}. As $\mathcal{T}_j$ is a tree, the path homotopy class of α_k only depends on x_{k0}. We form

$$\widetilde{\delta}_{jk} = \overline{\alpha}_k * \delta_{jk} * \alpha_k \quad \text{and} \quad \widetilde{\zeta}_{jk} = \overline{\alpha}_k * \zeta_{jk} * \alpha_k.$$

These define elements $r_{jk} = \widetilde{\delta}_{jk} * \widetilde{\zeta}_{jk}^{-1}$ in $\pi_1(\mathcal{SY}'_j, x_0)$, which after composing with the inclusion map yields elements of $\pi_1(M_j, x_0)$.

Theorem 7.4 *$\pi_1(M_j)$ has the representation*

$$\pi_1(M_j) \quad \simeq \quad F_{Q_j}/N \tag{7–2}$$

where F_{Q_j} is a free group on Q_j generators, and N is the normal subgroup generated by the elements $r_{jk} = \widetilde{\zeta}_{jk} * (\widetilde{\delta}_{jk})^{-1}$, one for each medial sheet without edge curves. Equivalently, $\pi_1(M_j)$ has exactly the relations

$$r_{jk} = \widetilde{\delta}_{jk} * \widetilde{\zeta}_{jk}^{-1} = 1 \quad \text{or equivalently} \quad \widetilde{\delta}_{jk} \;=\; \widetilde{\zeta}_{jk}. \tag{7–3}$$

Also, for F_{Q_j}, q_j of the generators correspond to the generators from the identified boundaries of the 2–disks, and $v_j + c_j$ generators come from $\mathcal{Y}_j$.

Also, λ_j generators come from the generators for $\pi_1(\Lambda_j)$. In fact, the collapsing map $\varphi_j\colon \mathcal{S}\mathcal{Y}_j \to \Lambda_j$ extends to a continuous map $\widetilde{\varphi}_j\colon M_j \to \Lambda_j$ which is a surjective map on π_1.

Proof M_j is obtained up to homotopy equivalence by attaching 2–disks to $\mathcal{S}\mathcal{Y}'_j$ along the closed paths given by (6–1); see Figure 21. By the Seifert–Van Kampen theorem this introduces relations killing the corresponding homotopy classes given by (6–1), but with the base point moved to y_j. These are exactly the relations given in (7–3). The identifications of the various generators from the $\mathcal{Y}_{jk}$, the edge curves of the S_{jk}, and the identified boundaries of the 2–disks follow from the Seifert–Van Kampen theorem.

To show that the collapsing map $\varphi_j\colon \mathcal{S}\mathcal{Y}_j \to \Lambda_j$ extends to M_j, we first construct a quotient $\widetilde{M}_j$ of M_j by collapsing the components $\mathcal{Y}_{jk}$ to points. Within $\widetilde{M}_j$ is the subspace $\widetilde{\mathcal{S}}\mathcal{Y}_j$, formed from $\mathcal{S}\mathcal{Y}_j$ by collapsing the components $\mathcal{Y}_{jk}$ to points. The collapsing map $\varphi_j\colon \mathcal{S}\mathcal{Y}_j \to \Lambda_j$ factors through $\widetilde{\mathcal{S}}\mathcal{Y}_j$. Hence, it is sufficient to extend this collapsing map to $\widetilde{M}_j$ and then compose with the quotient map $M_j \to \widetilde{M}_j$.

Then in $\widetilde{\mathcal{S}}\mathcal{Y}_j$, Δ_{jk} has its inner boundary edges collapsed to points to form $\widetilde{\Delta}_{jk}$, which is topologically a 2–disk. As a 2–disk is an absolute retract, we can extend the identity map on $\widetilde{\Delta}_{jk}$ to a continuous map $\widetilde{S}_{jk} \to \widetilde{\Delta}_{jk}$, where $\widetilde{S}_{jk}$ is S_{jk} with each boundary edge in Δ_{jk} collapsed to a point. Hence, these continuous maps together give a retract $\widetilde{M}_j \to \widetilde{\mathcal{S}}\mathcal{Y}_j$. Composed with the collapsing map gives the continuous extension $\widetilde{\varphi}_j$.

Finally the collapsing map $\widetilde{\Gamma}_j \to \Lambda_j$ is a homotopy equivalence, and factors through $\widetilde{\varphi}_j\colon M_j \to \Lambda_j$. Thus, the induced map on $\widetilde{\varphi}_{j*}\colon \pi_1(M_j) \to \pi_1(\Lambda_j)$ is surjective. □

Hence, in the special case when M_j is simply connected, we obtain Condition 5.1.

8 Exact sequence for attaching sheets

Next, we compute the homology and Euler characteristic of an irreducible medial component M_j. We use the two ways of obtaining M: by the medial sheet representation, attaching surfaces with boundaries to $\mathcal{SY}_j$, or the CW-representation obtained by attaching cells to $\mathcal{SY}'_j$. Using these we may compute the exact sequences of pairs $(M_j, \mathcal{SY}_j)$ and $(M_j, \mathcal{SY}'_j)$. Furthermore, we have the inclusion of pairs $(M_j, \mathcal{SY}_j) \subset (M_j, \mathcal{SY}'_j)$ This leads to the homomorphism of exact sequences of pairs using reduced homology. Since all spaces have dimension at most 2 the sequence ends at H_2.

(8–1)
$$\begin{array}{ccccccccc}
0 & \longrightarrow & H_2(\mathcal{SY}_j) & \longrightarrow & H_2(M_j) & \longrightarrow & H_2(M_j, \mathcal{SY}_j) & \xrightarrow{\delta} & \\
\| & & \downarrow & & \| & & \downarrow & & \\
0 & \longrightarrow & H_2(\mathcal{SY}'_j) & \longrightarrow & H_2(M_j) & \longrightarrow & H_2(M_j, \mathcal{SY}'_j) & \xrightarrow{\delta} & \\
& & H_1(\mathcal{SY}_j) & \longrightarrow & H_1(M_j) & \longrightarrow & H_1(M_j, \mathcal{SY}_j) & \longrightarrow & 0 \\
& & \downarrow & & \| & & \downarrow & & \| \\
& & H_1(\mathcal{SY}'_j) & \longrightarrow & H_1(M_j) & \longrightarrow & H_1(M_j, \mathcal{SY}'_j) & \longrightarrow & 0
\end{array}$$

First, using Proposition 6.1 and Proposition 6.3 we can compute the groups $H_i(\mathcal{SY}_j)$ and $H_i(\mathcal{SY}'_j)$, which are zero for $i > 1$ and for $i = 1$

$$H_1(\mathcal{SY}_j) \simeq \mathbb{Z}^{\lambda_j + v_j + c_j}$$

and
$$H_1(\mathcal{SY}'_j) \simeq \mathbb{Z}^{\lambda_j + v_j + c_j + q_j} \simeq \mathbb{Z}^{Q_j}.$$

Second, we can compute the relative groups by excision using Proposition 6.1 and Proposition 6.3:

$$H_i(M_j, \mathcal{SY}_j) \simeq \oplus_k H_i(\widetilde{S}_{jk}, C_{jk}) \tag{8–2}$$

where the sum is over all medial sheets, and

$$H_i(M_j, \mathcal{SY}'_j) \simeq \oplus_k H_i(D^2, \partial D^2)$$

where the second sum is over medial sheets without edge curves. From this, we conclude $H_i(M_j, \mathcal{SY}'_j) = 0$ if $i \neq 2$, and

$$H_2(M_j, \mathcal{SY}'_j) \simeq \mathbb{Z}^{s_{0j}}$$

(recall s_{0j} denotes the number of medial sheets of M_j which are without edge curves). Likewise, from (8–2), we compute $H_i(M_j, \mathcal{SY}_j)$ using the next lemma.

Lemma 8.1 *Suppose N is a compact surface with boundary, with a distinguished boundary component C. Let N' denote the surface obtained from N by attaching a 2–disk to N along the boundary component C. Then*

$$H_i(N, C) \simeq \widetilde{H}_i(N').$$

Proof By excision

$$H_i(N', D^2) \simeq H_i(N, C).$$

Then by the exact sequence of the pair (N', D^2) using reduced homology, we conclude

$$\widetilde{H}_i(N') \simeq H_i(\widetilde{N}, D^2). \qquad \square$$

Hence, by Lemma 8.1:

(1) If N has more than one boundary component, or N is nonorientable, then $H_2(N, C) = 0$; while if N is orientable with only the single boundary component C, then $H_2(N, C) = \mathbb{Z}$.

(2) For $(N, C) = (\widetilde{S}_{jk}, C_{jk})$, $\mathrm{rk}(H_1(\widetilde{S}_{jk}, C_{jk})) = q_{jk} - \varepsilon_{jk}$, where $\varepsilon_{jk} = 1$ if $e_{jk} = 0$ and $\widetilde{S}_{jk}$ is nonorientable, otherwise $\varepsilon_{jk} = 0$.

Thus, by (8–2), $H_i(M_j, \mathcal{SY}_j) = 0$ for $i > 2$, and

$$H_i(M_j, \mathcal{SY}_j) \simeq \begin{cases} \mathbb{Z}^{s_{0o,j}} & i = 2 \\ \mathbb{Z}^{(q_j - s_{0n,j})} \oplus (\mathbb{Z}_2)^{s_{0n,j}} & i = 1 \end{cases}$$

where we recall $s_{0o,j}$, resp. $s_{0n,j}$, denotes the number of orientable, resp. nonorientable, medial sheets with no edge curves.

Hence, the homomorphism of exact sequences (8–1) takes the form

(8–3)
$$\begin{array}{ccccccccc}
0 & \longrightarrow & 0 & \longrightarrow & H_2(M_j) & \longrightarrow & \mathbb{Z}^{s_{0o,j}} & \xrightarrow{\delta} & \\
\| & & \| & & \| & & \downarrow{\scriptstyle i} & & \\
0 & \longrightarrow & 0 & \longrightarrow & H_2(M_j) & \longrightarrow & \mathbb{Z}^{s_{0j}} & \xrightarrow{\delta} & \\
\mathbb{Z}^{(\lambda_j + v_j + c_j)} & & \longrightarrow & & H_1(M_j) & \longrightarrow & \mathbb{Z}^{(q_j - s_{0n,j})} \oplus (\mathbb{Z}_2)^{s_{0n,j}} & \longrightarrow & 0 \\
\downarrow{\scriptstyle i} & & & & \| & & \downarrow{\scriptstyle 0} & & \| \\
\mathbb{Z}^{(\lambda_j + v_j + c_j + q_j)} & & \longrightarrow & & H_1(M_j) & \longrightarrow & 0 & \longrightarrow & 0
\end{array}$$

Recall $\mathbb{Z}^{s_{0j}} = \mathbb{Z}^{s_{0n,j}} \oplus \mathbb{Z}^{s_{0o,j}}$ is the number of medial sheets without medial edge curves, and i denotes inclusion of $\mathbb{Z}^{s_{0o,j}}$.

We can draw a number of conclusions from the diagram (8–3). We let b_i denote the i–th Betti number of M_j, and we can take the Euler characteristics of either row and obtain the same formula. For example, from the first row of (8–3),

$$b_2 - s_{0o,j} + (\lambda_j + v_j + c_j) - b_1 + (q_j - s_{0n,j}) \quad = \quad 0.$$

This yields

$$\widetilde{\chi}(M_j) \quad = \quad s_{0j} - (\lambda_j + v_j + c_j + q_j) \quad = \quad s_{0j} - Q_j.$$

Substituting (3–6) into this equation yields (3–7) as asserted.

Computing homology by the algebraic attaching homomorphism

Next, we use the second row. The middle homomorphism is exactly the algebraic attaching homomorphism

$$\Psi_j \colon \mathbb{Z}^{s_{0j}} \xrightarrow{\delta} \mathbb{Z}^{Q_j}. \tag{8–4}$$

It sends the generator of $H_2(\widetilde{S}_{jk}, C_{jk})$ for orientable S_{jk} without edge curves to the image of its boundary (ie C_{jk}) under the attaching map to $\mathcal{Y}_j$. Here $\mathbb{Z}^{Q_j}$ represents the first homology of $\mathcal{SY}'_j$. Then by the exactness of the bottom row of (8–3), the kernel and cokernel of the Ψ_j are $H_2(M_j)$, respectively $H_1(M_j)$, yielding (2) of Theorem 3.2. Furthermore, by the top row of (8–3), we obtain $\operatorname{rk}(H_2(M_j;\mathbb{Z})) \le s_{0o,j}$. As $H_*(M_j;\mathbb{Z})$ is torsion free, we can tensor the top row with $\mathbb{Z}/2\mathbb{Z}$ and obtain $\operatorname{rk}(H_1(M_j;\mathbb{Z})) \ge (q_j - s_{0n,j}) + s_{0n,j} = q_j$. Also, by the bottom row, or the formula for $\widetilde{\chi}(M_j)$,

$$\operatorname{rk}(H_1(M_j;\mathbb{Z})) \;=\; Q_j - s_{0j} + b_2 \;\le\; Q_j - s_{0j} + s_{0o,j} \;=\; Q_j - s_{0n,j}.$$

This completes the proof of Theorem 3.2

The third consequence of diagram (8–3) is for the cases when various homologies of M_j vanish.

Proof of Corollary 3.3 First, suppose $H_1(M_j) = 0$. From the first row, we conclude that both $q_j - s_{0n,j} = 0$ and $s_{0n,j} = 0$. Hence, $q_j = s_{0n,j} = 0$. Then $q_j = \sum_k q_{jk}$ is a sum of nonnegative integers, so $q_{jk} = 0$. By (3–3), each $g_{jk} = 0$ and if $e_{jk} > 0$, then $e_{jk} = 1$. In particular, S_{jk} must be a 2–disk with a finite number of holes, and

the edge of at most one of the holes is an edge curve of M_j. In particular, all of the medial sheets are orientable so

$$s_{0o,j} = s_{0j} = s_j - e_j. \tag{8–5}$$

This establishes (2) and (3) of Corollary 3.3. Furthermore, by Theorem 7.4, the collapsing map $p\colon M_j \to \Lambda_j$ induces a surjection on π_1 and hence on the abelianized π_1. Hence, $H_1(\Lambda_j;\mathbb{Z}) = 0$, so Λ_j is a tree and $\lambda_j = 0$, establishing (1) of Corollary 3.3.

Second, suppose in addition that $H_2(M_j;\mathbb{Z}) = 0$. Then

$$\delta\colon \mathbb{Z}^{s_{0o,j}} \simeq \mathbb{Z}^{(\lambda_j+v_j+c_j)}.$$

Thus,

$$s_{0o,j} = \lambda_j + v_j + c_j.$$

Hence, as $\lambda_j = 0$, this equation and (8–5) together imply

$$s_j - e_j = v_j + c_j.$$

The remaining conditions on the fundamental group follow from Theorem 7.4. □

9 Algorithm for contracting contractible medial axes

In the contractible case, the structure suggests a method for contracting a region. We would like a specific algorithm for contracting the medial axis when it is contractible. First, we can deform it by sliding along fin curves so we have the simplified structure $\widehat{M}$. Then we want to simplify each irreducible medial component. If the medial component has a nonempty Y–network , then we choose a medial sheet with a medial edge curve, and deform the medial sheet so the medial edge curve touches the Y–network. At that point a transition occurs with two fin points being created as in Figure 22. We can then slide the fin curve until it doesn't meet the remaining Y–network. This does not change the homotopy type of the medial component.

We repeat this process as long as the Y–network of the component is nonempty. Applying this to all components, we reach a point where all Y–networks are empty. In the process we have created more irreducible components attached to each other along fin curves. By contractibility, all of the medial sheets are 2–disks. The attaching is described by a tree. We then proceed to simplify the tree.

We take one of the outermost branches. It corresponds to a 2–disk attached to another sheet along a segment of its boundary as a fin curve. We contract down the sheet and

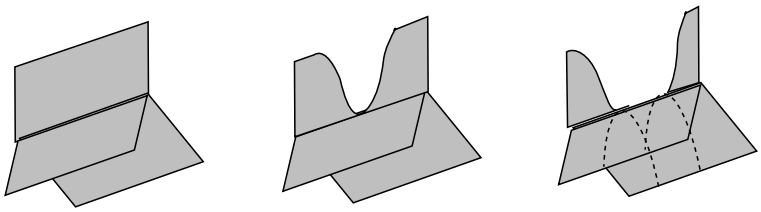

Figure 22: Creation of a pair of fin points from a medial edge curve

the fin curve to a point. Then we repeat this process until there is only a single medial component left. It is a 2–disk and hence contractible.

The only point about the preceding argument is the assertion that if a medial component is contractible then it has an edge curve. This step has not been completed and we are left with a conjecture.

Conjecture If an "irreducible medial component" is contractible then it has an edge curve.

The "irreducible medial component" is in quotes because after the first step we no longer know that it really comes from the medial axis of a region. What we really are asking is whether a singular space which is embedded in $\mathbb{R}^3$ and has the local singular structure of a medial axis satisfies the conjecture.

10 Discussion and open questions

There are a number of basic questions which are unanswered.

(1) The conditions we gave are for a medial axis which already exists. Given an abstract model for a medial axis, satisfying all of the properties, this model may not be realized in $\mathbb{R}^3$. In particular, we must understand the conditions allowing embeddings of the medial sheets so they do not intersect.

(2) Likewise, there are related restrictions on the attaching maps to Y–network so that three sheets are attached at each point. What constraints does this place on possible structures? This and the preceding point appear to involve subtle questions in topology.

(3) In order to complete the algorithm for contracting regions, we need to prove the existence of edge curves on contractible irreducible "medial components".

(4) The decomposition into irreducible medial components requires operations on fin curves. These local operations (with the exception of the cutting of essential fin curves) are examples of local deformations of medial axes resulting from deformations of regions by results of Giblin and Kimia [12] building on the results of Bogaevsky [2; 3]. Can we obtain the operations on fin curves as global deformations of the region?

(5) The structure of the medial axis for specific types of regions such as, for example, knot complement regions contains topological information about the region. For example, if we look for a topological representative of the region for which the number of medial sheets is minimal, then as the reduced Euler characteristic is fixed, we obtain from (3–7) an expression for $s_j - \widetilde{\chi}(M_j)$ as a sum of five nonnegative integers $e_j + v_j + c_j + G_j + \lambda_j$. What possible partitions of $s_j - \widetilde{\chi}(M_j)$ into five nonnegative integers are actually possible, and how does this relate to the structure of the original knot?

References

[1] **H Blum**, **R Nagel**, *Shape description using weighted symmetric axis features*, Pattern Recognition 10 (1978) 167–180

[2] **I A Bogaevskiĭ**, *Reconstructions of singularities of minimum functions, and bifurcations of shock waves of the Burgers equation with vanishing viscosity*, Algebra i Analiz 1 (1989) 1–16 MR1027457

[3] **I A Bogaevsky**, *Perestroikas of shocks and singularities of minimum functions*, Phys. D 173 (2002) 1–28 MR1945478

[4] **M A Buchner**, *The structure of the cut locus in dimension less than or equal to six*, Compositio Math. 37 (1978) 103–119 MR0501100

[5] **J Damon**, *Global Geometry of Regions and Boundaries via Skeletal and Medial Integrals*, to appear in Comm. Anal. and Geom.

[6] **J Damon**, *Smoothness and geometry of boundaries associated to skeletal structures. I. Sufficient conditions for smoothness*, Ann. Inst. Fourier (Grenoble) 53 (2003) 1941–1985 MR2038785

[7] **J Damon**, *Smoothness and geometry of boundaries associated to skeletal structures. II. Geometry in the Blum case*, Compos. Math. 140 (2004) 1657–1674 MR2098407

[8] **J Damon**, *Determining the Geometry of Boundaries of Objects from Medial Data*, Int. Jour. Comp. Vision 63 (2005) 45–64

[9] **M E Gage**, *Curve shortening makes convex curves circular*, Invent. Math. 76 (1984) 357–364 MR742856

[10] **M E Gage**, **R S Hamilton**, *The heat equation shrinking convex plane curves*, J. Differential Geom. 23 (1986) 69–96 MR840401

[11] **P Giblin**, *Symmetry sets and medial axes in two and three dimensions*, from: "The mathematics of surfaces, IX (Cambridge, 2000)", Springer, London (2000) 306–321 MR1846302

[12] **P Giblin**, **B Kimia**, *Transitions of the 3D Medial Axis Under a One-Parameter Family of Deformations*, from: "Proc. ECCV 2002", Lecture Notes in Computer Science 2351, Springer, London (2002) 718–734

[13] **M A Grayson**, *The heat equation shrinks embedded plane curves to round points*, J. Differential Geom. 26 (1987) 285–314 MR906392

[14] **B B Kimia**, **A Tannenbaum**, **S Zucker**, *Toward a computational theory of shape: An overview*, from: "Three Dimensional Computer Vision", (O Faugeras, editor), MIT Press, Boston (1990) 402–407

[15] **W S Massey**, *A basic course in algebraic topology*, Graduate Texts in Mathematics 127, Springer, New York (1991) MR1095046

[16] **J N Mather**, *Distance from a submanifold in Euclidean space*, from: "Singularities, Part 2 (Arcata, CA, 1981)", Proc. Sympos. Pure Math. 40, Amer. Math. Soc., Providence, RI (1983) 199–216 MR713249

[17] **S Pizer et al**, *Deformable M-reps for 3D Medical Image Segmentation*, Int. Jour. Comp. Vision 55 (2003) 85–106

[18] **J A Sethian**, *Level set methods*, Cambridge Monographs on Applied and Computational Mathematics 3, Cambridge University Press, Cambridge (1996) MR1409367Evolving interfaces in geometry, fluid mechanics, computer vision, and materials science

[19] **E H Spanier**, *Algebraic topology*, McGraw-Hill Book Co., New York (1966) MR0210112

[20] **Y Yomdin**, *On the local structure of a generic central set*, Compositio Math. 43 (1981) 225–238 MR622449

Department of Mathematics, University of North Carolina
Chapel Hill, NC 27599-3250, USA

jndamon@math.unc.edu

Proposed: Colin Rourke Received: 2 February 2006
Seconded: Robion Kirby, Walter Neumann Accepted: 30 August 2006