In science there are no 'depths'; there is surface everywhere: all experience forms a complex network, which cannot always be surveyed and can often be grasped only in parts.

— Rudolf Carnap, Hans Hahn, and Otto Neurath,
The Scientific Conception of the World: The Vienna Circle (1929)

The perfidious lemma of Dehn
Was every topologist's bane
'Til Christos Papa-
kyriakopou-
los proved it without any strain.

— John Milnor (c. 1956)

14 Normal Surfaces and Knots

14.1 Definitions

Following Kneser [8] and Haken [4], we say that an embedded surface S in a 3-manifold M is ***normal*** with respect to a fixed triangulation T of M if (1) all intersections between S and T are transverse, and (2) the intersection of S with each tetrahedron in T is a set of disjoint elementary disks. An ***elementary disk*** in a tetrahedron intersects either three or four edges of the tetrahedron, each exactly once; in other words, it looks just like the intersection of a Euclidean plane with a Euclidean tetrahedron. Each tetrahedron can support exactly seven different types of elementary disks: four triangles (each cutting off one vertex of the tetrahedron) and three quadrilaterals (each separating two vertices from two others). Because S is embedded, at most one type of quadrilateral can appear in any tetrahedron.

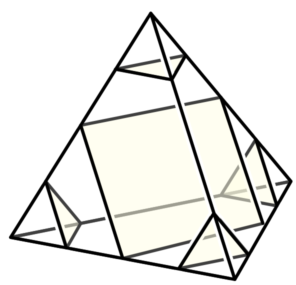

Five of the seven types of elementary disks inside a tetrahedron

For any normal surface S and any tetrahedron $pqrs$, let $S_{p|qrs}$ denote the number of elementary triangles that separate vertex w from the other vertices, and let $S_{pq|rs}$ denote the number of elementary quadrilaterals separating p and q from r and s. The normal surface S can be described by a vector $\langle S\rangle$ of $7t$ ***normal coordinates*** where t is the number of tetrahedra in T; these include $4t$ triangle coordinates $S_{w|xyz}$ and $3t$ quadrilateral coordinates $S_{wx|yz}$.

Haken proved that a vector $\langle S\rangle \in \mathbb{N}^{7t}$ is a normal coordiante vector if and only if it satisfies two constraints. First, different tetrahedra must agree on the number of intersections with each edge, as well as the number of each type of elementary segment on each triangle. It suffices for $\langle S\rangle$ to satisfy the following ***consistency constraint***; if $pqrz$ and $apqr$ are two tetrahedra sharing a common face pqr, then

$$S_{p|aqr} + S_{ap|qr} = S_{p|qrz} + S_{pz|qr}.$$

Second, within each tetrahedron, at most one of of the quadrilateral coordinates can be non-zero.

The normal-curve algorithms of Schaefer, Sedgwick, and Štefankovič [11] are easily adapted to similar problems on normal surfaces.

Theorem 14.1. *Given the normal coordinate vector $\langle S \rangle$ of a normal surface S, we can compute the number of components of S and the Euler characteristic of each component in polynomial time.*

14.2 Knots and Spanning Disks

A ***knot*** is a simple cycle in $\mathbb{R}^3$. We consider only ***polygonal*** knots, which are composed of a finite number of line segments. In fact, we will represent any knot by a planar ***knot diagram***, which is a connected 4-regular planar graph, possibly with loops and parallel edges, where each vertex represents a crossing and indicates which of the two strands goes 'over' the other. It is straightforward (if a bit tedious) to transform any polygonal knot into an equivalent knot diagram, or vice versa, in polynomial time. (One direction is described in the proof of Lemma 14.3 below.)

Two knots K and L are considered equivalent if there is an ***ambient isotopy*** from one to the other. An ambient isotopy is a function $h\colon [0,1] \times \mathbb{R}^3$, such that $h(0,\cdot)$ is the identity map, $h(t,\cdot)$ is a homeomorphism for all t, and $h(1,K) = L$. A knot K is ***trivial*** if it is ambient-isotopic to the standard unit circle S^1 in the xy-plane.

It is obvious, but surprisingly difficult to prove, that nontrivial knots exist at all. The first proof that *any* knot is nontrivial was given by Dehn [3]. In the same paper, Dehn claimed that a knot K is trivial if and only if $\pi_1(S^3 \setminus K) \cong \mathbb{Z}$; that is, K is trivial if and only if its complement $S^3 \setminus K$ is homotopy-equivalent to a circle. Kneser [8] observed that Dehn's proof was flawed; the first correct proof of 'Dehn's lemma' was given by Papakyriakopoulos almost 50 years after Dehn's paper [9]. It is relatively straightforward to compute a presentation of the group $\pi_1(S^3 \setminus K)$. Unfortunately, determining whether an *arbitrary* group presentation describes the infinite cyclic group $\mathbb{Z}$ is an undecidable problem [10], so any algorithm for deciding whether a knot is trivial must exploit some topological properties specific to knots.

Haken [4] developed the first algorithm to determine whether a given knot K is trivial, as one of the earliest applications of normal surface theory. Haken's algorithm is based on the following topological characterization of trivial knots:

Lemma 14.2. *A knot K is trivial if and only if it is the boundary of an embedded disk.*

One direction of the proof is immediate: If K is a trivial knot, reversing the ambient isotopy from K to S^1 gives an ambient isotopy from the unit disk B^2 to an embedded disk bounded by K. The other direction is also relatively straightforward, but much more tedious; any embedded disk with boundary K can be used to construct an ambient isotopy from K to the circle.

Intuitively, Haken's algorithm builds a triangulation of $S^3 \setminus K$ and then looks for a spanning disk that is ***normal*** with respect to this triangulation.

Lemma 14.3. *Given an n-vertex abstract knot diagram J, we can construct a polygonal knot K with $O(n)$ edges whose knot diagram is isomorphic to J, as well as a triangulation T with integer vertex coordinates between 0 and $8n$ that contains K in its 1-skeleton, all in $O(n)$ time.*

Proof: First subdivide each edge in the graph J into a path of length 3, to eliminate loops and parallel edges. Then embed the subdivided graph in the square $[n,7n] \times [n,7n]$ in the xy-plane, so that the coordinates of every vertex are integer multiples of 3, and every edge is a path of horizontal and vertical line segments [12, 13]. The total number of segments required for such an embedding is $O(n)$.

Modify the planar drawing in the neighborhood of each crossing point $(x,y,0)$ as follows. Add the vertices of a regular octahedron around the crossing point:

$$(x+1,y,0),\ (x-1,y,0),\ (x,y+1,0),\ (x,y-1,0),\ (x,y,1),\ (x,y,-1)$$

Now redirect the edges incident to $(x, y, 0)$ through this octahedron, so that the graph becomes a knot whose knot diagram is isomorphic to J. The new paths are

$$\ldots, (x-1, y, 0), (x, y, \pm 1), (x+1, y, 0), \ldots \quad \text{and} \quad \ldots, (x, y-1, 0), (x, y, \mp 1), (x, y+1, 0), \ldots;$$

the $\pm/\mp$ choice depends on whether the horizontal or vertical strand goes 'over' the crossing. Fill each octahedron with eight right-angled tetrahedra. To complete the planar drawing, we triangulate the (truncated) faces of the planar embedding, including the intersection of the outer face with the bounding square $[0, 8n] \times [0, 8n]$. (This is the only part of the construction that is difficult to complete in only $O(n)$ time [2].)

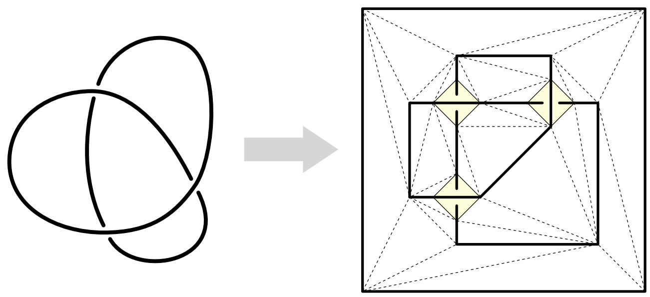

Transforming a knot diagram into a triangulated planar drawing.

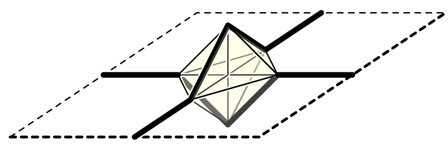

Local 3D triangulation at a crossing in the planar drawing.

To create a full 3d triangulation, add two vertices $a = (4n, 4n, 8n)$ and $z = (4n, 4n, 8n)$ and join them to every triangular face of the planar drawing. In particular, we join the apex a to the upper faces of each crossing octahedron, and join the zenith z to the lower faces of each crossing octahedron. The resulting 3D triangulation fills a large regular octahedron; to complete the required triangulation T of S^3, join every facet of this octahedron to the point at infinity. □

Let T be the triangulation provided by the previous lemma. Let $T'' = Sd(Sd(T))$, the result of barycentrically subdividing T twice. Finally, let $\boldsymbol{T_K}$ be the subcomplex of T'' consisting of simplices that do not intersect K. The repeated subdivision guarantees that $S^3 \setminus T_K$ is an open tubular neighborhood of K. Let $\boldsymbol{\partial T_K}$ denote the induced triangulation on the boundary of this tubular neighborhood, which is a torus. Haken's algorithm is formally based on the following refinement of Lemma 14.2.

Lemma 14.4 (Hass, Lagarias, and Pippenger [6]). *A knot K is trivial if and only if there is a disk D that is normal with respect to T_K and whose boundary ∂D is noncontractible on the torus ∂T_K.*

Call any normal surface satisfying the conditions of Lemma 14.4 an *spanning disk* of T_K.

14.3 The Haken Normal Cone

Given a vector $\langle D \rangle \in \mathbb{N}^{7t}$, we can determine in polynomial time (in $\|D\|$) whether $\langle D \rangle$ is the normal coordinate vector of a spanning disk, by checking the following conditions:

- **$\langle D \rangle$ is a normal coordinate vector.** We can easily check the consistency and quadrilateral conditions in $O(t)$ time.

- ***D* is connected.** This is a straightforward adaptation of the connectivity algorithm for normal curves, using either word equations [11] or orbit-counting [1].

- ***D* has Euler characteristic 1.** The Euler characteristic of D can be computed directly from the normal coordinates, just as we did to determine whether a normal curve is contractible. For example, the number of faces in D is just the sum of the normal coordinates in $\langle D\rangle$.

- **∂D is noncontractible on the torus ∂T_K.** The boundary curve ∂D is itself a normal curve in the torus triangulation ∂T_K, so we can check contractibility via word equations, as described in the previous lecture. However, a much simpler approach is to compute the ***homology*** class of ∂D (with $\mathbb{Z}_2$ coefficients), almost exactly as we did in the min-cut algorithm for surface graphs. We compute a tree-cotree decomposition (T', L', C') of ∂T_K. The set L' contains only two edges; call them e and e'. Let γ and γ' be the unique cycles in $T' \cup e$ and $T' \cup e'$. A normal cycle in ∂T_K is null-homologous if and only if it crosses both γ and γ' an even number of times.

Haken's unknottedness algorithm simply checks all vectors $\langle D\rangle \in \mathbb{N}^{7t}$ whose coordinates are smaller than a certain function $f(n)$; thus, to complete the description of the algorithm, we need to prove that a suitable function $f(n)$ exists. Hass, Lagarias, and Pippenger [6] prove that $f(n) = 2^{O(t)}$; it follows immediately that the unknottedness problem is in the complexity class NP.

To prove that a *small* spanning disk exists, Hass, Lagarias, and Pippenger closely examine the structure of the set of normal coordinate vectors. The ***Haken normal cone*** $\mathcal{H}(T)$ of a triangulation T is the set of all *real* vectors in $\mathbb{R}^{7t}$ that satisfy the nonnegativity and consistency constraints of a normal coordinate vector:

$$\begin{aligned} S_{p|qrs} &\geq 0 && \text{for all } pqrs \in T \\ S_{pq|rs} &\geq 0 && \text{for all } pqrs \in T \\ S_{p|aqr} + S_{ap|qr} &= S_{p|qrz} + S_{pz|qr} && \text{for all } apqr, pqrz \in T \end{aligned}$$

$\mathcal{H}(T)$ is an unbounded convex polyhedron with one vertex: the origin.

The Haken normal cone (like any polyhedral cone) is closed under addition: the sum of any two vectors in $\mathcal{H}(T)$ is another vector in $\mathcal{H}(T)$. An integral vector in $\mathcal{H}(T)$ is ***fundamental*** if it is not the sum of two other integral vectors in $\mathcal{H}(T)$. Similarly, a normal surface S is ***fundamental*** if there are no other normal surfaces S_1 and S_2 such that $\langle S\rangle = \langle S_1\rangle + \langle S_2\rangle$. Haken proved that any trivial knot has a *fundamental* spanning disk, and that any triangulation T has a finite number of fundamental normal surfaces. (See also below.) Thus, Haken's algorithm terminates in finite time. In fact, the following simple observation considerably simplifies Haken's algorithm:

Lemma 14.5. *Every fundamental normal surface is connected.*

Proof: If S is the union of two normal surface $S_1 \cup S_2$, then $\langle S\rangle = \langle S_1\rangle + \langle S_2\rangle$. □

The converse of this lemma is not true; not every connected surface—in particular, not every spanning disk of a trivial knot—is a fundamental surface.

14.4 Vertex Surfaces

The notion of fundamental surfaces was refined further by Jaco and Tollefson [7]. A ***vertex surface*** is a fundamental normal surface S whose normal coordinate vector $\langle S\rangle$ lies on an extreme ray of $\mathcal{H}(T)$.

Lemma 14.6 (Jaco and Tollefson [7]). *If a knot K has a spanning disk, then it has a spanning disk that is a vertex surface in $\mathcal{H}(T_K)$.*

In particular, Jaco and Tollefson proved that any spanning disk D that lexicographically minimizes the vector $(\#\text{vertices}(D), \#\text{faces}(D))$ is a vertex surface. Again, this is a nontrivial restriction; not every fundamental spanning disk is a vertex surface.

Lemma 14.7. *Let T be a triangulation of a 3-manifold with boundary. If S is a vertex surface in $\mathcal{H}(T)$, then every normal coordinate of S is at most $2^{O(t)}$; thus, $\|S\| = O(t^2)$.*

Proof: Consider the polytope $\overline{\mathcal{H}}(T)$ defined by intersecting the Haken normal cone $\mathcal{H}(T)$ with the hyperplane $\sum_i S_i = 1$. (Jaco and Tollefson [7] call $\overline{\mathcal{H}}(T)$ the ***projective solution space of T***.) Every vertex surface in T has normal coordinates λv for some vertex v of $\overline{\mathcal{H}}(T)$ and some scalar λ. Specifically, because each coordinate of v is a rational number less than or equal to 1, the scalar λ is the lowest common denominator of the coefficients of v.

Let v be any vertex of $\overline{\mathcal{H}}(T)$. The vector v is the solution to a system of $7t$ linear equations, which include all of (at most $2t$) the consistency constraints $v_{p|aqr} + v_{ap|qr} = v_{p|qrz} + v_{pz|qr}$, the normalizing constraint $\sum_i v_i = 1$, and at least $5t-1$ hyperplanes of the form $v_i = 0$. (In particular, at least $2t$ of the $3t$ quadrilateral coordinates of v are equal to zero.) Thus, we can write

$$Mv = (0,0,\dots,0,1)^\top$$

for some matrix $M \in \{-1,0,1\}^{7t\times 7t}$ that has one row of 1s, at most $2t$ rows with four nonzero entries, and at least $5t-1$ rows with exactly one nonzero entry. Cramer's rule implies that

$$v_i = \frac{\det M_i}{\det M},$$

where M_i is the matrix obtained by replacing the ith column of M with $(0,0,\dots,0,1)^\top$. Again, each M_i is a matrix in $\{-1,0,1\}^{7t\times 7t}$ with one row of 1s, at most $2t$ rows with four nonzero entries, and at least $5t-1$ rows with exactly one nonzero entry. It is now straightforward to prove that $|\det M| \le 7t\,4^{2t}$ and $|\det M_i| \le 7t\,4^{2t}$ for all i.

Finally, the point $(\det M)v$ is the normal coordinate vector of a vertex surface. Thus, every normal coordinate of every vertex surface is at most $(7t\,4^{2t})^2 = 2^{O(t)}$. □

Corollary 14.8 (Hass, Lagarias, and Pippenger [6]). *The problem of deciding whether a given polygonal knot is trivial is in the complexity class NP. There is an algorithm to determine whether a given polygonal knot with n vertices is trivial in $2^{O(n)}$ time.*

It is still open whether knot triviality is NP-complete, or if there is a polynomial time algorithm, or even if knot *non*-triviality is in NP.[1]

References

[1] I. Agol, J. Hass, and W. P. Thurston. The computational complexity of knot genus and spanning area. *Trans. Amer. Math. Soc.* 358(9):3821–3850, 2006.

[1]Hara, Tani, and Yamamoto [5] announced a 2-round interactive proof system that verifies whether a knot is non-trivial, by adapting a similar interactive proof system for graph non-isomorphism. Unfortunately, there is a subtle flaw in their proof [A. Sidiropoulos, personal communication, 2008].

[2] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.* 6(5):485–524, 1991.

[3] M. Dehn. Über die Topologie des dreidimensionalen Raumes. *Math. Ann.* 69(1):137–168, 1910.

[4] W. Haken. Theorie der Normalflächen: Ein Isotopiekriterium für den Kreisknoten. *Acta Mathematica* 105:245–375, 1961.

[5] M. Hara, S. Tani, and M. Yamamoto. UNKNOTTING is in AM ∩ co-AM. *Proc. 16th Ann. ACM-SIAM Symp. Discrete Algorithms*, 359–364, 2005.

[6] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *J. ACM* 46(2):185–211, 1999.

[7] W. Jaco and J. L. Tollefson. Algorithms for the complete decomposition of a closed 3-manifold. *Illinois J. Math.* 39(3):358–406, 1995.

[8] H. Kneser. Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten. *Jahresbericht Deutscher Math.-Verein.* 38:248–260, 1929.

[9] C. D. Papakyriakopoulos. On Dehn's lemma and the asphericity of knots. *Ann. Math.* 66:1–26, 1957.

[10] M. O. Rabin. Recursive unsolvability of group-theoretic problems. *Ann. Math.* 67:172–194, 1958.

[11] M. Schaefer, E. Sedgwick, and D. Štefankovič. Algorithms for normal curves and surfaces. *Proc. 8th Int. Conf. Comput. Combin.*, 370–380, 2002. Lecture Notes Comput. Sci. 2387, Springer-Verlag.

[12] J. A. Storer. On minimal node-cost planar embeddings. *Networks* 14(2):181–212, 1984.

[13] R. Tamassia and I. G. Tollis. Planar grid embedding in linear time. *IEEE Trans. Circ. Syst.* CAS-36(9):1230–1234, 1989.