The fence around a cemetery is foolish,
for those inside can't get out
and those outside don't want to get in.
— Arthur Brisbane, *The Book of Today* (1923)

Outside of a dog, a book is man's best friend.
Inside of a dog, it's too dark to read.
— attributed to Groucho Marx

1 The Jordan Polygon Theorem

The Jordan Curve Theorem and its generalizations are the formal foundations of many results, if not *every* result, in surface topology. The theorem states that any simple closed curve partitions the plane into two connected subsets, exactly one of which is bounded. Although this statement is intuitively clear, perhaps even obvious, the generality of the phrase 'simple closed curve' makes the theorem incredibly challenging to prove formally. According to most classical sources, even Jordan's original proof of the Jordan Curve Theorem [3] was flawed; most sources attribute the first correct proof to Veblen almost 20 years after Jordan [6]. (But see also the recent defense and updated presentation of Jordan's proof by Hales [2].)

However, we can at least sketch the proof of an important special case: simple polygons. Polygons are by far the most common type of closed curve in practice, so this special case has immediate practical consequences. Moreover, most proofs of the Jordan Curve Theorem rely on this special case as a key lemma. (In fact, Jordan dismissed this special case as obvious.)

1.1 First, A Few Definitions

A ***homeomorphism*** is a continuous function $h\colon X \to Y$ with a continuous inverse $h^{-1}\colon X \to Y$. Two topological spaces are ***homeomorphic*** (or ***topologically equivalent***) if there is a homeomorphism from one to the other.

A ***simple path*** is a subset of the plane that is homeomorphic to the unit interval $[0,1] \subset \mathbb{R}$, or equivalently, the image of a continuous injective function from $[0,1]$ into the plane.[1] A subset X of the plane is ***(path-)connected*** if for any two points in X, there is a simple path in X from one point to the other. A ***connected component*** of X is a maximal path-connected subset of X.

Similarly, a ***simple closed curve*** is a subset of the plane that is homeomorphic to the unit circle $\boldsymbol{S}^1 := \{(x,y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$, or equivalently, the image of a continuous injective function from S^1 into the plane. The full-fledged Jordan curve theorem states that for any simple closed curve C in the plane, the complement $\mathbb{R}^2 \setminus C$ has exactly two connected components.

Finally, a simple path or closed curve is ***polygonal*** if it is the union of a finite number of line segments (called ***edges***). An endpoint of an edge is called a ***vertex***. A polygonal path

A simple polygonal closed curve is also called a ***simple polygon***.

1.2 The Theorem

The Jordan Polygon Theorem. *The complement $\mathbb{R}^2 \setminus P$ of any simple polygon P in the plane has exactly two components.*

[1] In later lectures, to allow for self-intersections, we will formally define a path as a continuous function from $[0,1]$ into the ambient space, rather than the image of such a function.

Our proof essentially follows a recent proof of Thomassen (see also Hales), but with some additional details in the interest of formality. Throughout the proof, we fix an arbitrary simple polygon P. Let $d(p,q)$ denote the Euclidean distance between points p and q, and let $d(P,q) := \min_{p\in P} d(p,q)$. The compactness of P implies that $d(P,q) = d(p,q)$ for at least one point $p \in P$.

Lemma ≤ 2. *$\mathbb{R}^2 \setminus P$ has at most two connected components.*

Proof: For any $\varepsilon > 0$, let P_ε be the set of points at distance ε from P, that is, $P_\varepsilon := \{q \in \mathbb{R}^2 \mid \min_{p\in P} d(p,q) = \varepsilon\}$. We call this set an *offset curve* of P. Let δ be the minimum distance from any vertex v of P to any edge of P that is not incident to v. In particular, δ is at most the length of the shortest edge of P. (This distance is sometimes called the *minimum local feature size* of P.)

Fix a distance $\varepsilon < \delta/2$. The offset curve P_ε consists of straight line segments, which occur in pairs at distance ε on either side of each edge of P, and circular arcs of radius ε, which are centered at each vertex of P. These segments and arcs meet at common endpoints, clustered in triples around each vertex of P, as shown in the right half of the figure below. Because $\varepsilon < \delta/2$, the triangle inequality implies that these offset segments and arcs intersect *only* at their shared endpoints. Moreover, if we give the polygon an orientation, the offset segment just to the left of any edge of P is connected (either directly or through a single offset arc) to the offset segment just to the left of the next edge of P. Thus, by induction, all the segments and arcs to the left of the polygon are connected into a single cycle; symmetrically, all the right curves are connected into a single cycle.

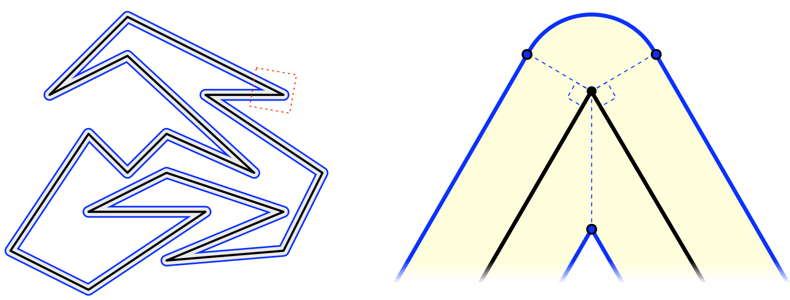

An offset curve P_ε and its structure near of a vertex of P.

Now let D be a closed circular disk centered at some point of P, such that $D \cap P$ is a single line segment, and let ρ be the radius of D. The set $D \setminus P$ clearly has exactly two connected components (open half-disks). Without loss of generality, we can assume that $\rho < \delta/2$.

Finally, let q be any point in $\mathbb{R}^2 \setminus P$, and let $\varepsilon = \min\{\rho, d(p,Q)/2\}$. The minimum-length line segment from q to P intersects the offset curve P_ε. Because $\varepsilon < \delta/2$, the offset curve P_ε has two components, each intersecting a component of D. Thus, there is a path in $\mathbb{R}^2 \setminus P$ from q to P_ε, and then along P_ε to a point in D. □

Lemma ≥ 2. *$\mathbb{R}^2 \setminus P$ has at least two connected components.*

Proof: Let p be an arbitrary point in $\mathbb{R}^2 \setminus P$, and let ρ be an arbitrary ray (infinite half-line) based at p. The intersection $\rho \cap P$ consists of a finite number of line segments, each of which is either a single point or an entire edge of P. We call one of these segments s a ***crossing*** if the edges of P preceding and following s lie on opposite sides of (the line through) ρ.

If we rotate ρ continuously around p, the number of crossings changes only when ρ encounters a vertex of P. Whenever the moving ray hits a single vertex, the number of crossings either stays the same, increases by 2, or decreases by 2. If the moving ray collides with several vertices simultaneously, the number of crossing may change by a larger number, but that number is always even. Thus, all rays based at p have the same number of crossings modulo 2; we call this bit the *parity* of the point p. A similar argument implies that moving p continuously along any path in $\mathbb{R}^2 \setminus P$ does not change its parity. Thus, the parity of any component of $\mathbb{R}^2 \setminus P$ is well-defined.

Finally, let σ be any line segment that crosses P exactly once. The endpoints of σ have different parities, and therefore lie in different components of $\mathbb{R}^2 \setminus P$. □

Notice that these two lemmas require complementary assumptions. The proof of Lemma ≥ 2 applies without modification if we replace P with any locally finite[2] collection of closed (not necessarily simple or disjoint) polygonal curves in the plane. On the other hand, the proof of Lemma ≤ 2 applies almost without modification to if we replace $\mathbb{R}^2$ with any 2-manifold (after imposing a suitable metric).

1.3 Point in Polygon Test

In light of Lemma ≤ 2, the proof of Lemma ≥ 2 can be easily converted into the standard algorithm to test whether a point is inside a simple polygon in the plane in linear time. Shoot an arbitrary ray from the query point, count the number of times this ray crosses the polygon, and return TRUE if and only if this number is odd. The calculations are simplified if we always shoot the ray directly to the right. This algorithm has been rediscovered several times, but the earliest publication seems to be a 1962 paper of Shimrat [5] (later corrected by Hacker [1]).

To make this algorithm concrete, we need one numerical primitive from computational geometry. A triple (q, r, s) of points in the plane are *oriented counterclockwise* if walking from q to r and then to s requires a left turn, or oriented clockwise if the walk requires a right turn. We can check this condition by computing the determinant

$$\Delta(q,r,s) := \det \begin{bmatrix} 1 & q.x & q.y \\ 1 & r.x & r.y \\ 1 & s.x & s.y \end{bmatrix} = (r.x - q.x)(s.y - q.y) - (r.y - q.y)(s.x - q.x).$$

The triple (q, r, s) is oriented counterclockwise if $\Delta(q, r, s) > 0$ and clockwise if $\Delta(q, r, s) < 0$. If $\Delta(q, r, s) = 0$, then the three points are collinear.

Finally, here is the algorithm. The input polygon P is represented by an array of consecutive vertices, which are assumed to be distinct. The algorithm returns $+1$, -1, or 0 to indicate that the query point q lies inside, outside, or directly on P, respectively. The RIGHTCROSS subroutine treats any polygon vertex that lies on the ray as though it were slightly above; this trick automatically takes care of degenerate cases. The algorithm clearly runs in $O(n)$ time.

```
PointInPolygon(P[1..n], q):
    sign ← −1
    P[0] ← P[n]
    for i ← 0 to n − 1
        sign ← sign · RightCross(q, P[i], P[i + 1])
    return sign
```

```
RightCross(q, r, s):
    if r.y > s.y
        swap r ⟷ s
    if (q.y ≤ r.y) or (q.y > s.y)
        return +1
    return sgn(Δ(q, r, s))
```

1.4 The Jordan-Schönflies Theorem

The following stronger version of the Jordan Curve Theorem, originally due to Arthur Schönflies [4], is also incredibly useful:

[2] Every point in the plane has an open neighborhood U that intersects only a finite number of the curves, and those curves have at most one point of common intersection in U. This condition is necessary to avoid pathological counterexamples.

The Jordan-Schönflies Theorem. *For any simple closed curve C in the plane, there is a homeomorphism from the plane to itself that maps C to the unit circle S^1.*

The Jordan-Schönflies Theorem implies not only that $\mathbb{R}^2 \setminus C$ has two components, but also that C is the boundary of both components, and that the closure of the bounded component is homeomorphic to the disk $B^2 := \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$. A polygonal version of this stronger theorem can be proved by triangulating the interior of the polygon, and then mapping the polygon first to a combinatorially equivalent triangulation of a regular convex n-gon, and then to the disk. We leave the remaining details as an exercise for the reader.

References

[1] R. Hacker. Certification of Algorithm 112: Position of point relative to polygon. *Commun. ACM* 5(12):606, 1962.

[2] T. Hales. Jordan's proof of the Jordan curve theorem. *From Insight to Proof: Festschrift in Honour of Andrzej Trybulec*, 2007. Studies in Logic, Grammar and Rhetoric 10(23), University of Białystok.

[3] C. Jordan. Courbes continues. *Cours d'Analyse de l'École Polytechnique*, 587–594, 1887. vol. 3.

[4] A. Schönflies. Beiträge zur Theorie der Punktmengen. III. *Math. Ann.* 62:286–328, 1906.

[5] M. Shimrat. Algorithm 112: Position of point relative to polygon. *Commun. ACM* 5(8):434, 1962.

[6] O. Veblen. Theory on plane curves in non-metrical analysis situs. *Trans. Amer. Math. Soc.* 6:83–98, 1905.