Towards a Synthetic Proof of the Polygonal Jordan Curve Theorem

Towards a synthetic proof of the Polygonal Jordan Curve Theorem (Extended Abstract)

Phil Scott and Jacques Fleuriot
phil.scott@ed.ac.uk, jdf@inf.ed.ac.uk

School of Informatics, University of Edinburgh, UK.
phil.scott@ed.ac.uk, jdf@inf.ed.ac.uk

1 Introduction

The Jordan Curve Theorem effectively says that if a closed curve does not intersect itself, then it must divide its plane into an inside and an outside. The first rigorous formulation of this statement was given by Jordan in 1894 [2], but his proof was criticised eleven years later by Veblen [14]. In Veblen's own proof, generally accepted as the first rigorous proof of the theorem, he points out that Jordan assumed without proof that the theorem held in the special case of the polygon.

In his 1899 edition of the celebrated *Foundations of Geometry* [8], David Hilbert gave his own formulation of the special case in terms of just three primitive relations on three primitive domains: an *incidence* relation on *points* and *lines*; an incidence relation on points and *planes*; and finally, a linear ordering relation on triples of points. These primitives are sufficient to formulate the theorem, and Hilbert claimed that his axioms allowed one to prove the theorem "without much difficulty."

In 1904, Veblen gave a standalone formulation and synthetic proof of the special case based on his own axiomatic system, using similar primitives. He gave a detailed two page proof broken down into several lemmas. Even so, according to Reeken and Kanovei [9], the proof was deemed "inconclusive". Solomon Feferman, writing after a long history of proofs of the theorem, gives caution by pointing out that the theorem "turned out to be devilishly difficult to prove even for reasonably well-behaved simple closed curves, namely those with polygonal boundary" [5]. It is little wonder that in the tenth edition of the *Foundations of Geometry*, Bernays had edited out the phrase "without much difficulty."

The Jordan Curve Theorem has a long history within the formal verification community. The MIZAR [3] community first began its verification in 1991, and completed the special case in 1996. The full proof was completed in 2005. In the same year, Hales completed the proof in HOL Light [6, 7]. Both proofs use the special case for polygons, though in a restricted form: in the case of the MIZAR proof, only polygons with edges parallel to axes are considered. In Hales' proof, the polygon is restricted to lie on a grid. The formulations are algebraic rather than synthetic, and so are outside the scope of Hilbert's and Veblen's formulations.

We are currently verifying Hilbert's *Foundations of Geometry*, which was the first rigorous axiomatic treatment of Euclidean geometry, and has been called the most influential work in geometry on the 20th century [1]. There have been partial attempts to formalise it, one utilising an intuitionistic approach by Dehlinger et al [4] in the Coq theorem prover [15], while our own project started from a partial formalisation by Meikle and Fleuriot [10] in the Isabelle theorem prover [11]. We have since migrated to HOL Light [7] where we have found it easier to rapidly prototype automated tools.

We have reached the point where we must verify Hilbert's formulation of the Polygonal Jordan Curve Theorem, a daunting prospect given its history. In this paper, we discuss aspects of our ongoing proof, and describe some of the representations and automated tools that have assisted us so far.

2 Formulation of the Theorem

Hilbert states the theorem under very weak geometric assumptions, before any notions of angle, distance, parallels or continuity are introduced. The only notions Hilbert had recourse to were those of incidence and the linear ordering of points. This setting rules out traditional proofs of the theorem, including the one by Hales, and therefore makes the proof particularly challenging. Indeed, we can intuitively point out that a proof of the polygonal Jordan curve theorem is effectively a means to navigate a maze whose walls are defined by a single edge, and then note that in a world in which there is no notion of distance, orientation or even a notion of what it means for a path to run parallel to an edge, our ability to navigate is likely to be significantly compromised.

> THEOREM 9. Every single polygon lying in a plane α separates the points of the plane α that are not on the polygonal segment of the polygon into two regions, the *interior* and the *exterior*, with the following property: If A is a point of the interior **(an inner point)** and B is a point of the exterior **(an exterior point)** then every polygonal segment that lies in α and joins A and B has at least one point in common with the polygon. On the other hand if A, A' are two points of the interior and B, B' are two points of the exterior then there exist polygonal segments in α which join A with A' and others which join B with B', none of which have any point in common with the polygon. By suitable labelling of the two regions there exist lines in α that always lie entirely in the exterior of the polygon. However, there are no lines that lie entirely in the interior of the polygon.

In developing a formalised proof of this theorem which appeals only to properties of incidence and linear ordering, we needed automation and representations specific to geometry theorem proving, but general enough to cover the very weak axiom system in which we are working. The only axioms we can appeal to govern two incidence relations among points, lines and planes, and a betweenness relation ordering three points on a line.

3 Linear Reasoning

Much of the proof relies on reasoning about ordering, and many of the difficult case-splits occur in particular with *linear orders*. To help solve these linear problems, we use as our main workhorse a discovery system tailored to reasoning about incidence and which we describe elsewhere [12, 13]. In this section, we describe how we can automatically reduce the problems to a decision procedure for linear arithmetic. The reduction is justified by our formalisation of Hilbert's sixth theorem:

> THEOREM 6 (generalisation of Theorem 5). Given any finite number of points on a line it is always possible to label them A, B, C, D, E,. . ., K in such a way that the point labelled B lies between A and C, D, E, . . ., K, the point labelled C lies between A, B, and D, E, K, D lies between A, B, C and E, . . ., K, etc. Besides this order of labelling there is only the reverse one that has the same property.

The term "labelling" here is not a primitive in Hilbert's theory. It is metatheoretical. So to formally verify the theorem over all possible labellings, we need to bring this notion down to the object level, and so we define an "ordering" as follows (Hilbert's primitives and some other derived terms are given in Table 1).

$$\begin{aligned}\text{ordering}\, f\, X \longleftrightarrow\, & f(X) = \{n |\, \text{finite}\, X \longrightarrow n < |X|\} \\ & \wedge \forall n\; n'\; n''. \,(\text{finite}\, X \longrightarrow n < |X| \wedge n' < |X| \wedge n'' < |X|) \\ & \qquad \wedge\, n < n' \wedge n' < n'' \\ & \qquad \longrightarrow \text{between}\, (f\, n)\, (f\, n')\, (f\, n'')\end{aligned}$$

Primitive	Meaning
A online a	The point A lies on the line a
between $A\,B\,C$	The point B lies between A and C
ordering $f\,X$	The function f orders the points of X
P intriangle (A, B, C)	The point P lies in the interior of $\triangle ABC$.
P ontriangle (A, B, C)	The point P lies on the sides of $\triangle ABC$.
connected $P\,Q$	The points P and Q can be connected by a polygonal segment.

Table 1. Primitive relations

We have formalised Theorem 6 as follows:

$$\forall X.\ \text{finite}\ X \wedge (\exists a. \forall P.\, P \in X \longrightarrow P\ \text{online}\ a) \longrightarrow \exists f.\ \text{ordering}\ f\ X$$

and from this, we obtain the corollary concerning the inverse of the "labelling":

$$\begin{aligned}
&\forall X.\ \text{finite}\ X \wedge (\exists a. \forall P. P \in X \longrightarrow P\ \text{online}\ a) \qquad (1)\\
&\quad \longrightarrow \exists g. \forall A\ B\ C.\ A \in X \wedge B \in X \wedge C \in X\\
&\qquad\qquad \longrightarrow (\text{between}\ A\,B\,C \longleftrightarrow (g\,A < g\,B \wedge g\,B < g\,C)\\
&\qquad\qquad\qquad\qquad \vee (g\,C < g\,B \wedge g\,B < g\,A))\\
&\quad \wedge \forall A\ B.\ A \in X \wedge B \in X \longrightarrow (A = B \longleftrightarrow g\,A = g\,B)
\end{aligned}$$

We have implemented a tactic which converts goals involving betweenness into goals involving linear arithmetic. The tactic takes a concrete enumerated set of points as a parameter and tries to apply it to corollary 1. To do so, it first proves that the enumerated set is collinear, using our incidence discoverer. After this, all formulas involving betweenness of points can be rewritten to inequalities, and any equations and inequations of points can be lifted into equations and inequations of the image under g. Provided that the original goal is solvable entirely by linear reasoning on the chosen set, the proof can be solved by decision procedures for linear arithmetic.

We have applied this tactic routinely during our ongoing formalisation of the proof of the Polygonal Jordan Curve Theorem. Part of this proof follows Veblen's idea of decomposing a polygon into triangles, where the argument that a polygon separates the plane into *at least* two regions reduces to the claim that a triangle separates the plane into at least two regions. This can be proven by showing that if a segment crosses a triangle at a single point between two of its vertices, then one of those points must lie in the interior of the triangle (see Figure 1):[1]:

$$\begin{aligned}
&\neg(\exists a. A\ \text{online}\ a \wedge B\ \text{online}\ a \wedge C\ \text{online}\ a)\\
&\wedge \neg P\ \text{ontriangle}\ (A, B, C) \wedge \neg Q\ \text{ontriangle}\ (A, B, C)\\
&\text{between}\ A\,R\,B \wedge \text{between}\ P\,R\,Q\\
&\wedge (\forall X. X\ \text{ontriangle}\ (A, B, C) \wedge \text{between}\ P\,X\,Q \longrightarrow R = X)\\
&\longrightarrow P\ \text{intriangle}\ (A, B, C) \vee Q\ \text{intriangle}\ (A, B, C)
\end{aligned}$$

The proof runs to 18 steps (excluding steps for reasoning about planes) and 3 of these use the linear ordering tactic. We start by considering the case that C lies on the line PQ. To use our linear reasoning tactic, we must first prove the following:

$$\neg\text{between}\ P\,C\,Q \qquad (2)$$

$$P \neq C \wedge Q \neq C \wedge R \neq C \qquad (3)$$

The first condition follows because we assume that PQ intersects the triangle only once and at R. The second condition follows because we assume that P and Q are not on the triangle.

[1] For space, we have omitted assumptions about all points being planar.

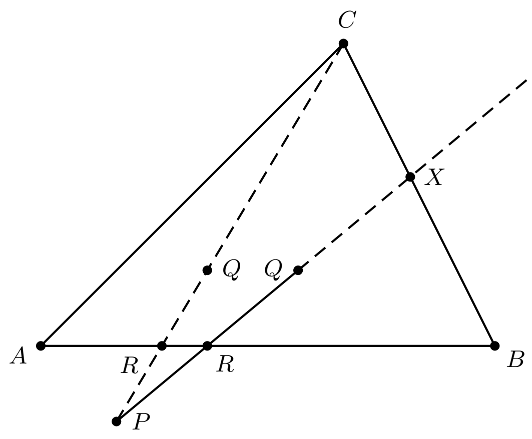

Fig. 1. Crossing a Triangle

We now set as our goal the conclusion $\text{between}\, C\, P\, R \lor \text{between}\, C\, Q\, R$, and apply our linear reasoning tactic. The tactic solves the goal, and allows us to conclude from either case that one of P or Q is inside the triangle.

Next, we consider the possibility that C is not on PQ. Our plan here is to apply Pasch's Axiom to the triangle and the line of PQ, and so obtain a point at which the line PQ exits the triangle. But to do this, we must show that the vertices A and B do not lie on the line PQ. We do this by contradiction.

Supposing that one of the vertices lies on PQ, it follows that PQ *is* the line AB. But we know that P and Q do not lie on the triangle, so we must have:

$$P \neq A \land P \neq B \land Q \neq A \land Q \neq B \land \neg \text{between}\, P\, A\, Q \land \neg \text{between}\, P\, B\, Q$$

At this point, a contradiction must follow by linear reasoning alone, and is deduced using a tactic.

We can now apply Pasch's axiom to find a point X where the line PQ emerges from the triangle. In other words, we obtain a point X that is either between B and C or between A and C.

Now we prove the following

$$\neg \text{between}\, P\, X\, Q \tag{4}$$

$$P \neq X \land Q \neq X \tag{5}$$

Again, the first condition follows because we assume that PQ intersects the triangle only once. The second follows because P and Q are not on the triangle. We set as our goal the conclusion $\text{between}\, R\, P\, X \lor \text{between}\, R\, Q\, X$, and then apply our linear reasoning tactic. The first disjunct tells us that P is inside the triangle. The second tells us that Q is inside. This concludes all the cases of the theorem.

4 Triangle Symmetry

In our proof of the Polygonal Jordan Curve Theorem, we have found a need to exploit symmetries in the argument. Our proof of the theorem relies on triangulating a polygon, and the properties of triangles in which we are interested (indeed, the properties that we can *define* in the weak setting of Hilbert's early axioms) are invariant up to any permutation of the vertices of a triangle. One way to proceed is therefore to repeat every formal proof for each symmetry. This however, is inefficient when running the proofs, and is not robust to later refactoring.

This was an issue when trying to route between any two points in the exterior of a triangle. The interior of a triangle $\triangle ABC$ is defined as the intersection of three half-planes on the lines AB, AC and BC. Any exterior point is then defined with respect to one or more half-planes, and proving that we can navigate between these points requires reasoning carefully about the betweenness relation when applied to the relevant half-planes.

There is a great deal of symmetry in the proofs, and to abstract over this, we introduced a notation to describe the position of a point relative to the three half-planes defining a triangle.

For example, the line AB divides the plane into two half-planes. One half-plane contains the interior of the triangle and the other contains the exterior. We will use the notation $\mathcal{I}_{AB}(P)$ to say that the point P lies on the same side of AB as the interior of the triangle. We use the notation $\mathcal{X}_{AB}(P)$ to say that P lies on the same side as the *exterior* of the triangle. Finally, we use the notation $\mathcal{S}_{AB}(P)$ to say that the point P lies on on the *line* AB but not on the triangle's edge.

Since a triangle is defined by three lines, every point on the plane which is not on the edge of the triangle is defined by a triple. For instance, a point can be stated to lie in the interior with $\mathcal{I}_{AB}(P) \wedge \mathcal{I}_{AC}(P) \wedge \mathcal{I}_{BC}(P)$, which we abbreviate to $\mathcal{I}_{AB}\mathcal{I}_{AC}\mathcal{I}_{BC}(P)$.

We now have the following theorems which completely characterise these triples (we use x, y, z and w to denote variables ranging over $\mathcal{I}, \mathcal{X}, \mathcal{S}$):

$$\vdash \forall A\,B\,C.\ \neg(\mathcal{I}_{AB}x_{AC}y_{BC}(P) \wedge \mathcal{S}_{AB}x_{AC}y_{BC}(P)) \tag{6}$$
$$\vdash \forall A\,B\,C.\ \neg(\mathcal{I}_{AB}x_{AC}y_{BC}(P) \wedge \mathcal{X}_{AB}x_{AC}y_{BC}(P)) \tag{7}$$
$$\vdash \forall A\,B\,C.\ \neg(\mathcal{S}_{AB}x_{AC}y_{BC}(P) \wedge \mathcal{X}_{AB}x_{AC}y_{BC}(P)) \tag{8}$$
$$\vdash \forall A\,B\,C.\ P \text{ intriangle } (A, B, C) \longleftrightarrow \mathcal{I}_{AB}\mathcal{I}_{AC}\mathcal{I}_{BC}(P) \tag{9}$$
$$\vdash \forall A\,B\,C.\ \mathcal{X}_{AB}\mathcal{X}_{AC}x_{BC}(P) \longrightarrow \mathcal{X}_{AB}\mathcal{X}_{AC}\mathcal{I}_{BC}(P) \tag{10}$$
$$\vdash \forall A\,B\,C.\ \mathcal{S}_{AB}\mathcal{I}_{AC}x_{BC}(P) \longrightarrow \mathcal{S}_{AB}\mathcal{I}_{AC}\mathcal{X}_{BC}(P) \tag{11}$$
$$\vdash \forall A\,B\,C.\ \neg\mathcal{S}_{AB}\mathcal{S}_{AC}x_{BC}(P) \tag{12}$$
$$\vdash \forall A\,B\,C.\ \mathcal{X}_{AB}x_{AC}y_{BC}(P) \wedge \mathcal{X}_{AB}z_{AC}w_{BC}(Q) \longrightarrow \text{connected } P\,Q \tag{13}$$
$$\vdash \forall A\,B\,C.\ \mathcal{X}_{AB}x_{AC}y_{BC}(P) \longrightarrow \exists Q.\, \text{connected } P\,Q \wedge \mathcal{X}_{AB}\mathcal{X}_{AC}\mathcal{I}_{BC}(Q) \tag{14}$$

Theorems (6)–(8) are injectivity lemmas for the notation. Together with (9)–(12), they allow us to narrow down the possible triples from 27 to 13, while theorems (13) and (14) allow us to navigate from any of the 12 exterior regions to any other exterior region.

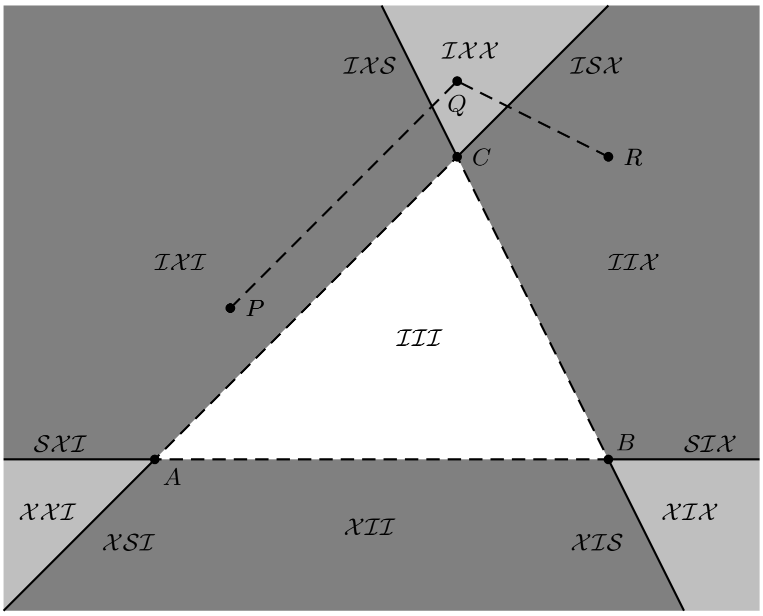

Fig. 2. Regions of a Triangle

Consider the points P and R in Figure 2. The point P is notated by $\mathcal{I}_{AB}\mathcal{X}_{AC}\mathcal{I}_{BC}(P)$. We want to apply 14, but to do so we must permute the first two symbols. To do this, we can rotate the triangle clockwise and notate the point by $\mathcal{X}_{CA}\mathcal{I}_{CB}\mathcal{I}_{AB}(P)$. This done, we can find a connecting point Q notated by $\mathcal{X}_{CA}\mathcal{X}_{CB}\mathcal{I}_{AB}(P)$.

We apply a similar argument to R which is notated by $\mathcal{I}_{AB}\mathcal{I}_{AC}\mathcal{X}_{BC}(P)$. We first apply a reflection and notate the point by $\mathcal{X}_{CB}\mathcal{I}_{CA}\mathcal{I}_{BA}(R)$. We now use 14 to find a connection point Q' notated by $\mathcal{X}_{CB}\mathcal{I}_{CA}\mathcal{I}_{BA}$. Applying a second reflection gives us the representation $\mathcal{X}_{CA}\mathcal{I}_{CB}\mathcal{I}_{AB}$.

Finally, we apply Theorem (13) to show that Q and Q' are connected, and thus that P and R are connected by transitivity.

5 Conclusion

The Jordan Curve Theorem for polygons is challenging. When we restrict ourselves to a weak subset of Hilbert's synthetic axioms from the *Foundations of Geometry*, its proof is particularly difficult, and Hilbert did not even provide an informal one. To aid our formalised proof, we need a repertoire of automated tools and convenient representations to handle the symmetries involved. We have harnessed decision procedures for arithmetic to handle linear reasoning, based on our formalisation of one of Hilbert's theorems and a tactic we have implemented to rewrite goals. We have also formalised a succinct notation to completely abstract over the complex details of navigating around the exterior of a triangle, allowing us to push the symmetries of triangles into the symmetries of a simple notation.

References

1. Garrett Birkhoff and Mary Bennett. Hilbert's Grundlagen der Geometrie. *Rendiconti del Circolo Matematico di Palermo*, 36:343–389, 1987.
2. Camille Jordan. *Science and Method*. Paris, Gauthier-Villars et fils, 1894.
3. N.G. de Bruijn. The Mathematical Vernacular, a language for Mathematics with typed sets. In P. Dybjer et al, editor, *Proceedings from the Workshop on Programming Logic*, volume 37, 1987.
4. Christophe Dehlinger, Jean-François Dufourd, and Pascal Schreck. Higher-Order Intuitionistic Formalization and Proofs in Hilbert's Elementary Geometry. In *Automated Deduction in Geometry: Revised Papers from the Third International Workshop on Automated Deduction in Geometry*, volume 2061, pages 306–324, London, UK, 2001. Springer-Verlag.
5. Solomon Feferman. Mathematical intuition vs. mathematical monsters. *Synthese*, 125(3):317–332, 2000.
6. Thomas C. Hales. Jordan's Proof of the Jordan Curve Theorem. *Studies in Logic, Grammar and Rhetoric*, 10(23), 2007.
7. John Harrison. HOL Light: a Tutorial Introduction. In *Proceedings of the First International Conference on Formal Methods in Computer-Aided Design*, volume 1166, pages 265–269. Springer-Verlag, 1996.
8. David Hilbert. *Foundations of Geometry*. Open Court Classics, 10th edition, 1971.
9. Vladimir Kanovei and Michael Reeken. A nonstandard proof of the Jordan curve theorem. *Pacific Journal of Mathematics*, 36(1):219–229, 1971.

10. Laura I. Meikle and Jacques D. Fleuriot. Formalizing Hilbert's Grundlagen in Isabelle/Isar. In *Theorem Proving in Higher Order Logics*, volume 2758, pages 319–334. Springer, 2003.
11. Lawrence C. Paulson. *Isabelle: a Generic Theorem Prover.* Number 828 in Lecture Notes in Computer Science. Springer – Berlin, 1994.
12. Phil Scott and Jacques D. Fleuriot. An Investigation of Hilbert's Implicit Reasoning through Proof Discovery in Idle-Time. In *Automated Deduction in Geometry*, Lecture Notes in Computer Science, pages 182–200. Springer, 2010.
13. Phil Scott and Jacques D. Fleuriot. Composable discovery engines for interactive theorem proving. In *Interactive Theorem Proving*, volume 6898 of *Lecture Notes in Computer Science*, pages 370–375. Springer, 2011.
14. Oswald Veblen. Theory on Plane Curves in Non-Metrical Analysis Situs. *Transactions of the American Mathematical Society*, 6(1):83–98, 1905.
15. Pierre Castéran Yves Bertot. *Interactive Theorem Proving and Program Development.* Springer, 2004.