
The Jordan Curve Theorem, Formally and Informally

Thomas C. Hales

1. INTRODUCTION. The Jordan curve theorem states that every simple closed planar curve separates the plane into a bounded interior region and an unbounded exterior. One hundred years ago, Oswald Veblen declared that this theorem is “justly regarded as a most important step in the direction of a perfectly rigorous mathematics” [13, p. 83]. Its position as a benchmark of mathematical rigor has continued to our day.

Many vastly underestimate the logical gulf that separates a typical published proof from a fully formal mathematical proof, in which every single logical inference has been generated and checked by computer. A striking example of this logical gulf can be found in Bourbaki’s *Elements of Sets*. How many primitive symbols of logic does it require to represent the number “1” in fully expanded form? Bourbaki estimates that the fully expanded representation may require “several tens of thousands of symbols.” The story has it that when A. R. Mathias learned of Bourbaki’s estimate, he thought “that must be false, surely only a couple of hundred” are required. This led him to make a careful calculation and to uncover the astounding fact that Bourbaki’s own estimate is wrong by several orders of magnitude. Over four trillion symbols are needed to express the number “1” [7]. Such is the logical gulf. It is fortunate that not all systems suffer from the same inefficiencies.

Who hasn’t heard the one about the mathematician who wakes up at night from the smoke of a fire in his hotel? Seeing a fire extinguisher, he notes that a solution to the problem exists and falls peacefully back into a deep sleep. Researchers from Frege to Gödel, who solved the problem of rigor in mathematics, found theoretical solutions but did not extinguish the fire, because they omitted the practical implementation. Some, such as Bourbaki, have even gone so far as to claim that “formalized mathematics cannot in practice be written down in full” and call such a project “absolutely unrealizable” [1, pp. 10, 11].

While it is true that formal proofs may be too long to print, computers—which do not have the same limitations as paper—have become the natural host of formal mathematics. In recent decades, insomniacs at the same hotel have reworked the foundations of mathematics, putting them in an efficient form designed for real use on real computers. We can look back at Bourbaki’s trillion-symbolized definitions and claims of absolute unrealizability as historical curiosities.

Today, there is an ambitious long-term endeavor to translate a large body of mathematics into formal proofs that can be read and checked by computer. Such computer systems check every logical inference of every lemma, theorem, and corollary all the way back to the fundamental axioms of mathematics. Even a simple arithmetic identity such as $1 + 1 = 2$ is checked thoroughly from the basic principles of arithmetic, rather than relying on the arithmetic capabilities of the computer chip. While it is impossible to avoid every potential source of defect in a real-world computer, the designers of these systems continue to make every effort to provide the greatest mathematical rigor attainable by current technology.

One computer system that checks the correctness of mathematical proofs is called MIZAR. The MIZAR group has made the computer verification of the Jordan curve theorem a pet project.

To give an idea of the magnitude of a fully computer verified proof of the Jordan curve theorem within MIZAR, Trybulec (the designer of MIZAR) has observed that “the time needed for formalization is extremely difficult to estimate. Proving the Hahn-Banach theorem took three days, if I recall. On the other hand, the proof of the Jordan curve theorem we started in 1991 and as yet it is proved only for polygonal curves (actually a special case for polygons with edges parallel to axes)” [12, p. 1]. The MIZAR webpage reports that as of 2003 “74 articles devoted to this theorem have been collected so far, which makes about 10% of the Mizar Mathematical Library” [9].

Because of the strong attraction of the problem and the tendency of those who undertake this project to disappear beyond the event horizon and never be seen again, G. Gonthier has described the Jordan curve theorem as the black hole into which formalizing mathematicians fall [3]. It was in this same lecture that Gonthier announced the completion of the formal proof of the four color theorem.

In January 2005, I completed a formal proof of the Jordan curve theorem in a different computer system, called HOL Light [4], [5]. (HOL Light is so named for being a lightweight implementation of *Higher-Order-Logic*.) The implementation took me months of rather intensive work at a computer. In September 2005, the Mizar team completed their formal proof of the same theorem. The formal proof of the Jordan curve theorem in HOL-light consists of 138 definitions, 1381 lemmas, and over 44,000 proof steps spread over 59,000 lines of computer code. There are approximately 20 million primitive logical inferences in this proof of the Jordan curve theorem.

While these numbers may initially make the proof seem unwieldy, it is unwise to make judgements of complexity (notwithstanding Kolmogorov complexity) based on simple word counts. Does the 138-word vocabulary of the proof of the Jordan curve theorem make it far less complex than *Hamlet*, which contains a vocabulary of 4605 words? Do 59,000 lines indicate that it is that many times more difficult than the one-line unix command that counted the distinct words in *hamlet.txt* or half that many times more difficult than Adam Back’s two lines of Perl code that implement the RSA algorithm? Or could we argue in the opposite direction that the proof has trivial complexity compared with the forty million source lines of code in Windows XP?

A formal proof typically begins with a careful conventional proof. Each lemma is then expanded in meticulous detail and then transcribed to computer for exhaustive checking.

My starting point is an elegant proof by C. Thomassen [11]. We will wait to give a sketch of Thomassen’s proof until section 3. The proof first proves the Jordan curve theorem for polygons and then uses an approximation argument to derive the Jordan curve theorem in general.

2. THE CASE OF POLYGONS. To ease our way into the formal proof, let us recall Thomassen’s argument showing that if C is a simple closed polygonal curve in the plane, then its complement has two connected components at most. The argument that he gives is a standard proof that can be found in many places [2]. The argument is brief and convincing. Pick a disk D such that $D \cap C$ is a line segment (Figure 1). Starting from any point x in the plane not lying on C , we can “walk along a simple polygonal arc (close to C but not intersecting C)” from the point x into one of the two components of $D \setminus C$. Hence, there are two components at most.

Nobody doubts the correctness of this argument. Every mathematician knows how to walk without running into walls. Detailed figures indicating how to “walk along a simple polygonal arc” would be superfluous, if not downright insulting.

Yet, it is quite another matter altogether to train a computer to navigate its way around a maze-like polygon without collisions. To translate this argument into an ar-




Figure 1. A simple closed polygon curve separates the plane into two components at most, because it is always possible to walk to a given disk that is separated into two halves by the curve.

argument a computer can follow, far more instruction would be required. In fact, in our relentless pursuit of a simple logical argument, we do not even attempt the Jordan curve theorem for general polygons. Rather, we restrict to polygons whose edges have integer length and whose edges run parallel to the coordinate axes. This restriction was introduced in [10].

Graph paper geometry. Let us now drill down to a level close to the formal level. The following definitions are not fully formal because they are not capable of being parsed by computer, yet they are definitions that a trained human could without hesitation rewrite in computer readable form.

- The *plane* is the set of functions f from $\mathbb{N} = \{0, 1, \dots\}$ to \mathbb{R} that are identically 0.0 for all inputs except possibly 0 and 1.

This definition of plane will strike most mathematicians as peculiar. Why not represent the plane in the usual way as ordered pairs? As we point out in greater detail at the end of the paper, each mathematical term in the logical system we use has a *type*. Phrasing the definition this way will make it easy to define n -dimensional Euclidean space in such a way that n has the type of a natural number.

- Elements of the plane are *points*. For any two real numbers x and y write $p(x, y)$ for the point taking value x at 0 and y at 1.
- An *integer point* is a point $p(x, y)$ in the plane for which x and y are integers (or more precisely, the image of integers under the canonical embedding of the set of integers into the set of real numbers). If i and j are integers, write $I(i, j)$ for the corresponding integer point.
- A *horizontal edge* is a planar segment, excluding its two endpoints, from an integer point $I(i, j)$ to $I(i + 1, j)$.
- A *vertical edge* is a planar segment, excluding its two endpoints, from an integer point $I(i, j)$ to $I(i, j + 1)$.
- An *edge* is a horizontal edge or a vertical edge.

- A *square* is a subset of the plane consisting of all points whose first coordinate lies strictly between the integers i and $i + 1$ and whose second coordinate lies strictly between the integers j and $j + 1$ for some integers i and j .
- A *cell* is either a singleton set consisting of an integer point, an edge, or a square.

(The narrow definition of *edge* as a horizontal or vertical segment of unit-length in the plane is useful in the context of graph-paper geometry. This paper consistently uses the term in this narrow sense except in the specific contexts of the edges of a graph or polygon.)

Earlier, we stated that the proof of the Jordan curve theorem relies on 1381 lemmas. We now state a number of lemmas at the same degree of granularity as in a formal proof. The only difference is that they have been transcribed from a strict computer syntax to idiomatic English. At the end of this paper, we give some examples of statements expressed in the syntactically precise notation used by the computer.

- Two integer points are equal if and only if their integer coordinates are equal.
- A horizontal edge is not equal to a vertical edge.
- The closure of an edge contains two integer points.
- Two horizontal edges are equal if and only if their left endpoints are equal.
- The closure of an edge does not contain a square.
- The only square in the closure of a square s is s itself.
- Each square is flanked by four edges.
- Each square has four corners.
- Two of the edges that flank a square are horizontal.
- The two horizontal edges that flank a square have disjoint closures.
- Two cells are disjoint or equal.
- The closure of a cell is a union of cells.

From these examples, we see that each lemma of the formal proof is smaller in scope than what is traditionally called a lemma.

The complaint is sometimes made (naturally only by people who have never studied an actual formal proof) that formal proofs would be too complex to understand. And yet, there is nothing inordinately complex about these statements, even though they are lemmas coming directly out of a formal proof. Perhaps we need to be regularly reminded that formal proofs—although machine readable and checkable—are designed by humans and are implemented in high-level languages that have been designed for human convenience and understanding.

In the formal proof, each of these statements must be accompanied by a sequence of instructions to the computer telling it how to generate a proof of the statement. Here are some individual statements, again at the same level of granularity as a formal proof, but rewritten into English. There are over 44,000 such instructions to the computer that generate the proof of the Jordan curve theorem.

- Take the statement “every edge is a horizontal edge or a vertical edge” and rewrite it by replacing the term “horizontal edge” with its definition.
- Take an inequality of the form $x < |y|$ and rewrite it into the form $(0.0 \leq y \wedge x < y) \vee (y < 0.0 \wedge x < -y)$, justified by a built-in decision procedure for linear real arithmetic.

- Show that a closed ball with center a is contained in an open ball of larger radius with the same center, using the definitions of the terms “closed_ball,” “open_ball,” “subset,” and the built-in decision procedure for linear real arithmetic.
- Break a statement to be proved of the form $A \wedge B$ into two separate subcases: first considering A and then B .

We should confess that this is the size of *our* steps in the formal proof of the Jordan curve theorem, but those that are at the grand-master level of formal theorem proving achieve much more in a single step, sometimes with considerable ingenuity.

A simple case of the Jordan curve theorem. Finite sets of edges can be represented on a sheet of graph paper by tracing along its guidelines from intersection to intersection, excluding points of intersection. The *support* of the finite set of edges is the subset of the plane consisting of the closures of all the edges.

We call a finite set of edges a *segment* if its support is a *simple arc* in the plane, that is, a continuous injective image of an interval in the real line. The endpoints of the simple arc are called the *endpoints* of the segment. A simple closed curve in the plane is the continuous image of a unit interval under a function that is injective on the complement of an endpoint and that takes the same value at the two endpoints. We call a finite set of edges a *rectagon* if its support is a simple closed curve. The term is intended to suggest rectilinear polygons. It turns out that a Jordan curve theorem for rectagons will deliver the full Jordan curve theorem.

The walk that we make around a general polygon (Figure 1) disappears from the rectagon version. Instead, a unit square is moved—like a token on an infinite chessboard—from square to square, up or down, to the left or right, or diagonally, from an arbitrary starting square until it reaches one of the two unit squares along a given edge of a given rectagon (Figure 2).

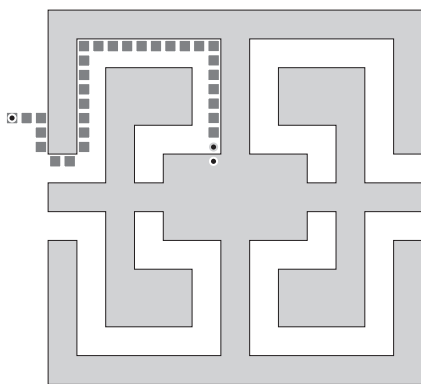


Figure 2. A rectagon separates the plane into two components at most, because it is always possible to move a token in unit steps from any unit square to one of the two unit squares along a fixed edge of the rectagon, represented by two dots in the center of the figure.

A little lemma. To give further insight into the degree of detail required in formal proofs, we focus for a few moments on a single lemma. The result in question reads: *a unit square can be moved from a position along one edge of a rectagon to a position along an adjacent edge of the rectagon, while remaining in the same connected component of its complement.* This lemma is what justifies the movement of the token

in Figure 2. Its proof threatens to degenerate into an unpleasantly long procession of cases depending on the various relative positions of the two edges and the unit square. The stock phrases of mathematicians such as “it is easy to see” or “the other cases are similar” have no formal effect. We discuss a couple of possibilities, beginning with the approach that we actually use in the formal proof.

In the formal proof, we argued by symmetry. Up to symmetries of the plane, there are only three cases: the edges turn toward the initial square to flank it, the edges are collinear, and the edges turn away from from the initial square so that only one flanks it (see Figure 3).



Figure 3. To move a token around a rectagon, up to symmetry, there are only three local possibilities: no movement required, a step to an adjacent square, and a step kitty-corner.

We came to regret this approach, because formal proofs by symmetry are much harder than anticipated. It was necessary to give a total of nearly a hundred lemmas, showing that the symmetries preserve all of the relevant structures, all the way back to the foundations. The necessary facts include the following: that the symmetries are homeomorphisms of the plane, that the image under a symmetry of the topological closure of a set is the closure of the image, that the number of edges meeting at an integer point is preserved by a symmetry, explicit facts about how the symmetries act on the unit squares, edges, and integer points. The bloat from this little symmetry argument occupies about 6 percent of the total proof of the Jordan curve theorem. All this to justify the three words “up to symmetry” in Figure 3. All this for a fact that is intuitively obvious.

A preferable way to prove this little lemma might be to examine the integer point at which the two edges of the rectagon meet and to study the four squares and four edges (two of which belong to the rectagon) around it (Figure 4). If the given unit square is along both edges, then there is no need to move the square, and we are done. If not, then one edge of the rectagon is along the given initial square and the other is along the diagonally opposite unit square. If this diagonally opposite square is in the same connected component of the complement, then we move diagonally to that square, and we are done. Otherwise, each of the four squares is adjacent to another square in the same connected component. A simple count then shows that the four squares occupy at most two connected components. The square s that is adjacent to the diagonally opposite one and that shares an edge of the rectagon must be in a different component from the diagonally opposite one. The square s thus lies in the same connected component as the initial square. This completes the alternate proof.



Figure 4. The argument justifying a movement around a rectagon focuses on one integer point, the four edges (two belonging to the rectagon), and the four squares around it.

Even though this argument is better from a formal point of view than the original approach based on symmetry, it, too, requires a series of definitions and lemmas. It would probably take at least a couple of days of hard work to transform it from its present state to a formal proof.

Two components at most in the complement of a rectagon. It is not our intention to describe the thousands of lemmas that enter into the proof, but for the sake of completeness we briefly describe the remainder of the proof of the Jordan curve theorem for rectagons. We follow that with an even sketchier outline of the rest of the proof.

To show that the complement of a rectagon has at least two components, we define a parity function on the set of unit squares in the plane. A square is said to be *odd* if there are an odd number of horizontal edges of the rectagon directly below the given square. A square is *even* if the number is even (Figure 5). It follows immediately from this definition that a unit square and the unit square directly below it have opposite parity precisely when the rectagon passes between them.

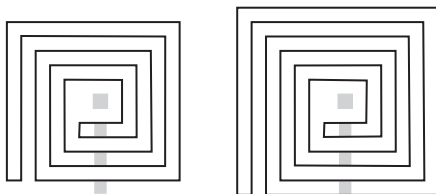


Figure 5. The central unit square is odd in the first frame, because there are an odd number of edges below it; but the central unit square is even in the second frame.

It is also true that a unit square and the unit square directly to its left have opposite parity, precisely when the rectagon passes between them. This statement becomes clear on inspection of Figure 6.

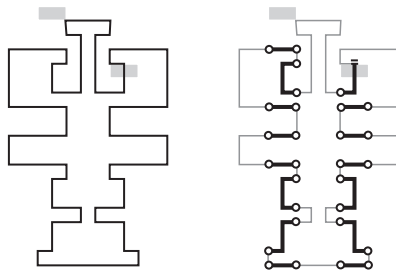


Figure 6. If a rectagon does not pass between a unit square and the square to its left (shaded 1×2 rectangle on the left of both figures), the total number of horizontal edges below the two squares is even, being in bijection with the endpoints (shown as circles) of the segments (shown as heavy lines) forming the part of the rectagon below the shaded rectangle. If the rectagon passes between the squares (shown on the right of both figures), the one stray endpoint of the segment passing between them reverses the parity.

The parity function can be extended to any edge that does not lie on the rectagon by the common value of the two unit squares it flanks. It can similarly be extended to any integer point (except those on the rectagon) as the common value of the unit squares that surround it. The parity function is then locally constant, hence constant on each connected component of the complement of the rectagon. Both parities exist, since the parity is reversed as a unit square crosses the rectagon. Accordingly, there are at least two connected components.

3. A SKETCH OF THOMASSEN'S PROOF. Although Thomassen avoided diagrams to “emphasize that the proofs are rigorous,” geometrical intuition guides his proofs. The same intuition guides the expanded version of the proof that was used as a script for the formal proof. Geometrical intuition guides the proof so strongly that the entire detailed proof of the Jordan curve theorem could be presented as a sequence of “proofs without words” in the style of a column in *Mathematics Magazine*. In fact, even the finished product—the formal proof itself—is nonverbal to an extraordinary degree. It completely bypasses the false opposition (that many profess) between picture and logical rigor, by specifying one through the other with absolutely minimal intervention of human language. Although we do not go so far as to give proofs without words, a few of the key lemmas of the proof are illustrated in this section.




Figure 7. The utility plant puzzle asks one to join each utility plant to each house with noncrossing paths. The impossibility is essentially equivalent to the Jordan curve theorem.

Thomassen’s proof is based on the fact that the Jordan curve theorem is nearly equivalent to the nonplanarity of the complete bipartite graph $K_{3,3}$ (the graph obtained by joining each of one set of three vertices to each of a second set of three vertices). This graph is also known as the “utility graph,” because it is the graph that occurs in the utility graph puzzle—a children’s puzzle of drawing paths from three houses to three utility plants in such a way that none of the paths cross (Figure 7). From a slightly different point of view, the utility graph can be represented as a hexagon (or a general Jordan curve) with three diagonals (Figure 8).




Figure 8. The utility graph can be described as a hexagon with three diagonals.

The first step of the proof is to prove the Jordan curve theorem for rectangles. This is treated in some detail in section 2.

The proof then gives a construction that shows how to produce a planar embedding of the utility graph from any counterexample to the Jordan curve theorem. There are two figures illustrating the idea. The first (Figure 9) treats the case of a Jordan curve with a connected complement, and the second (Figure 10) treats the case of a Jordan curve whose complement has at least three connected regions.

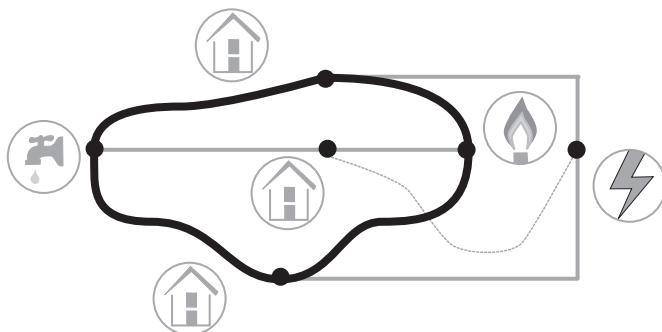


Figure 9. A Jordan curve (heavy line) whose complement is connected could be used to solve the utility plant puzzle as shown. Under the assumption of a connected complement, the house along the central line can be joined to the utility plant outside the curve by a path (faint line) that does not cross the Jordan curve.

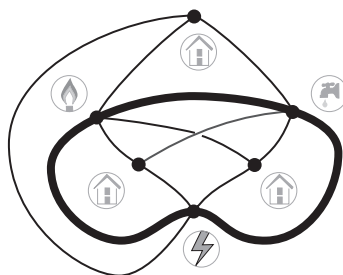


Figure 10. A Jordan curve (heavy line) whose complement has at least three connected components could be used to solve the utility plant puzzle as shown. Under this assumption, each house can be situated in a different connected component, and the apparent crossing in the center of the figure can be avoided, because the paths lie in different connected components of the complement.

These figures gloss over what is perhaps the most important technical point in the proof of the Jordan curve theorem: a simple arc A (A is a set that is homeomorphic to the real interval $[0, 1]$) in the plane has a connected complement. This technical result is required to show that the three points (utility plants) on the Jordan curve in Figure 10 are accessible from each of three connected components in the complement. The idea of the proof is as follows. Given two points p and q that do not lie on A , we construct a *scaled* rectangle¹ with A in its interior and so tightly around the simple arc that both p and q are in the exterior of the rectangle (Figure 11). Other than this subtle argument, the figures adequately convey the idea of the proof.

A graph approximation theorem is then proved that takes any planar embedding of the utility graph and constructs from it a planar embedding with piecewise linear

¹A scaled rectangle is defined as rectangle, except that its basic unit of length for edges is a real number r , rather than fixed at 1. At this stage of the argument, the Jordan curve theorem is known for rectangles. This allows us to refer to the (scaled) rectangle's interior and exterior.

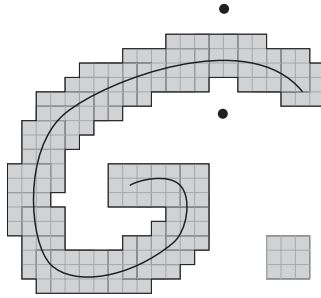


Figure 11. Given a simple arc and two points (dots) not on the arc, a scaled rectangle can be constructed as a cage around the simple arc, fitting so closely that both points lie in the exterior region. The scaled rectangle can be constructed as the outer perimeter of the figure formed by sliding a sufficiently small 3×3 square along the arc.

edges. In our modification of Thomassen's proof, we show that the embedding can be made with a finite number of horizontal and vertical segments, each with integer length (Figure 12).

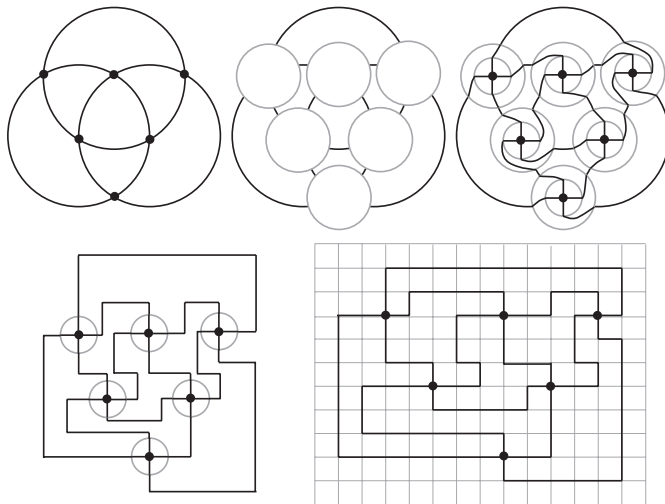


Figure 12. (The stages of a planar embedding.) A finite planar graph with degrees at most four admits an embedding with a finite number of integer-length horizontal and vertical edges. First cut away a disk centered around each vertex, and remove any orphaned edges. Redraw the graph inside each disk in such a way that the edges are horizontal and vertical near each vertex. Next, replace the part of each edge outside the inner disk with horizontal and vertical segments. Finally, stretch the graph between vertices in horizontal and vertical directions to give the edges integer lengths. The final graph is isomorphic to the first.

To finish the proof, we argue that if the Jordan curve theorem is false, then there is a solution to the utility graph puzzle. By the graph approximation theorem, there is a solution to the utility graph puzzle in which edges are horizontal and vertical segments, each of integral length. This graph can be described as a rectangle with three diagonals. By the Jordan curve theorem for rectangles, every diagonal lies in the interior or exterior region formed by the rectangle. By the pigeon-hole principle, at least two of the three diagonals traverse the same region. Regardless of whether that region is the interior or exterior, Figure 13 shows that the two diagonals in the same region intersect, contrary to the requirements of the utility graph puzzle. This contradiction

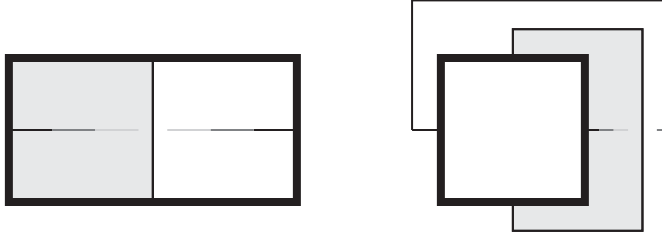


Figure 13. In the figure on the left, if two diagonals with alternating endpoints on a rectangle (heavy line) both lie in the interior of the rectangle, then the first diagonal starts in the interior of a second rectangle (interior shaded), formed in part by the second diagonal. However, it finishes in the exterior of the second rectangle. By the Jordan curve theorem for rectangles, the two diagonal cross. A similar argument (on the right) shows that the diagonals cross when both lie in the exterior of a given rectangle.

establishes the Jordan curve theorem. Thomassen’s original proof can be consulted for further details [11].

4. A FORMAL STATEMENT. The formal statement of the Jordan curve theorem takes the following form. The expression within the single quotes is written in a formal mathematical syntax that can be parsed by computer.

$$\begin{aligned}
 & \text{'}\forall C. \text{ simple_closed_curve } top2\ C \Rightarrow \\
 & (\exists A\ B. \text{ top2 } A \wedge \text{ top2 } B \wedge \\
 & \quad \text{connected } top2\ A \wedge \text{ connected } top2\ B \wedge \\
 & \quad (A \neq \emptyset) \wedge (B \neq \emptyset) \wedge \\
 & \quad (A \cap B = \emptyset) \wedge (A \cap C = \emptyset) \wedge (B \cap C = \emptyset) \wedge \\
 & \quad (A \cup B \cup C = \text{euclid } 2)\text{' }
 \end{aligned}$$

A rather literal translation of this HOL Light code into English is as follows: [Here, *top2* is the standard metric space topology on \mathbb{R}^2 .] Let *C* be a simple closed curve, with respect to *top2*. Then there exist sets *A* and *B* with the following properties: *A* and *B* are open in the topology *top2*; *A* and *B* are connected with respect to the topology *top2*; *A* and *B* are nonempty; the sets *A*, *B*, and *C* are pairwise disjoint; and the union of *A*, *B*, and *C* is \mathbb{R}^2 . Or more idiomatically, a Jordan curve *C* partitions the plane into the three sets *A*, *B*, and *C* itself, where *A* and *B* are nonempty, connected, and open.

The formal statement of the Jordan curve theorem depends on a number of other definitions. Logical definitions (such as \wedge , \Rightarrow , \forall , and \exists) and basic set-theoretic definitions (such as \emptyset , \cap , \cup) are part of the standard HOL Light package. The definitions that have been added to the standard package to express the Jordan curve theorem are precise to the point of being machine readable.

$$\begin{aligned}
 & \text{'}(\forall x. \text{euclid } 2\ x = (\forall n. 2 \leq n \Rightarrow (x\ n = 0.0))) \wedge \\
 & (\text{top2} = \text{top_of_metric } (\text{euclid } 2, d_euclid)) \wedge \\
 & (\forall (X : A \rightarrow \text{bool})\ d. \text{top_of_metric } (X, d) = \\
 & \quad \{A \mid \exists F. F \subset \text{open_balls } (X, d) \wedge (A = \bigcup F)\}) \wedge \\
 & (\forall (X : A \rightarrow \text{bool})\ d. \text{open_balls } (X, d) = \\
 & \quad \{B \mid \exists x\ r. (B = \text{open_ball } (X, d)\ x\ r)\}) \wedge \\
 & (\forall X\ d\ (x : A)\ r. \text{open_ball } (X, d)\ x\ r = \\
 & \quad \{y \mid X\ x \wedge X\ y \wedge d\ x\ y < r\}) \wedge
 \end{aligned}$$

$$\begin{aligned}
& (\forall U (Z : A \rightarrow \text{bool}). \text{connected } U Z = \\
& \quad Z \subset \bigcup U \wedge \\
& \quad (\forall A B. U A \wedge U B \wedge (A \cap B = \emptyset) \wedge Z \subset A \cup B \Rightarrow \\
& \quad \quad Z \subset A \vee Z \subset B)) \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall (C : A \rightarrow \text{bool}) U. \text{simple_closed_curve } U C = \\
& \quad (\exists f. (C = \text{IMAGE } f \{x \mid 0.0 \leq x \wedge x \leq 1.0\}) \wedge \\
& \quad \text{continuous } f (\text{top_of_metric } (UNIV, d_real)) U \wedge \\
& \quad \text{INJ } f \{x \mid 0.0 \leq x \wedge x < 1.0\} (\bigcup U) \wedge \\
& \quad (f(0.0) = f(1.0)))) \wedge
\end{aligned}$$

$$\begin{aligned}
& (\forall (f : A \rightarrow B) U V. \text{continuous } f U V = \\
& \quad (\forall v. V v \Rightarrow U \{x \mid (\bigcup U) x \wedge v (f x)\})) \wedge
\end{aligned}$$

$$(\forall x y. \text{d_real } x y = \text{abs}(x - y)) \wedge$$

$$\begin{aligned}
& (\forall x y. \text{euclid } 2 x \wedge \text{euclid } 2 y \\
& \quad \Rightarrow (\text{d_euclid } x y = \\
& \quad \quad \text{sqr}t(\text{sum}(0, 2)(\lambda i. (x i - y i) * (x i - y i))))'
\end{aligned}$$

In these definitions, every term and every subterm has a *type*. In most cases, the type of a term can be inferred from its context. For example, in the definition of *euclid 2*, the variable n appears in the expression $2 \leq n$. From this, we can infer that n is a natural number. The quantifier $(\forall n)$ then runs over all natural numbers n . Similarly, from the expression $x n = 0.0$, we see that the variable x runs over functions from natural numbers to real numbers.

The term *UNIV* (short for *universe*) represents the set of all elements of a given type. The given type depends on the context. We can infer from its context in the definition of *simple_closed_curve* that it represents the set of real numbers. Thus, *top_of_metric (UNIV, d_real)* is a long-winded expression for the standard metric space topology on the real line.

In a few cases, it is not possible for the computer to infer the type, and the type is explicitly provided. For example, the variable x in the definition of *open_ball* has a type that cannot be inferred. The expression $(x : A)$ in that definition indicates that the type is to be left indeterminate and is represented by a type variable A .

ACKNOWLEDGMENTS. The author's research has been supported by NSF grant 0503447.

REFERENCES

1. N. Bourbaki, *Elements of Sets*, Addison-Wesley, Reading, MA, 1968.
2. R. Courant and H. Robbins, *What is Mathematics?*, Oxford University Press, Oxford, 1996.
3. G. Gonthier, lecture at the TYPES conference, Campus Thalès (17 December, 2004).
4. T. Hales, The formal proof of the Jordan curve theorem (HOL Light code), available at <http://www/math.pitt.edu/~thales/>.
5. J. Harrison, HOL Light, available at <http://www.cl.cam.ac.uk/users/jrh/hol-light/index.html>.
6. C. Jordan, *Cours d'analyse de l'École Polytechnique*, Paris, 1893.
7. A. R. D. Mathias, A term of length 4,523,659,424,929, *Foundations of the formal sciences 1* (Berlin, 1999), *Synthese* **133** (2002) 75–86.
8. R. Maehara, The Jordan curve theorem via the Brouwer fixed point theorem, this *MONTHLY* **91** (1984) 641–643.

9. MIZAR formalization of the Jordan curve theorem, available at <http://mizar.uwb.edu.pl/jordan/>
10. Y. Nakamura and Y. Takeuchi, On the Jordan curve theorem (1980, preprint).
11. C. Thomassen, The Jordan-Schönflies theorem and the classification of surfaces, this MONTHLY **99** (1992) 116–130.
12. A. Trybulec, MIZAR proofs, April 1999, available at <http://www.cs.nyu.edu/pipermail/fom/1999-April/002996.html>.
13. O. Veblen, Theory on plane curves in non-metrical analysis situs, *Trans. Amer. Math. Soc.* **6** (1905) 83–98.

THOMAS C. HALES received his M.S. from Stanford in the School of Engineering and his Ph.D. (1986) from Princeton in mathematics. He has taught at Harvard, the University of Chicago, and the University of Michigan. He is currently the Andrew Mellon Professor of Mathematics at the University of Pittsburgh. His honors include the Chauvenet Prize (2003) of the MAA and the Moore Prize (2004) for applications of interval analysis. His research interests include representation theory, motivic integration, discrete geometry, and formal proofs.

Math Department, University of Pittsburgh, Pittsburgh PA 15260
hales@pitt.edu

Emily Dickinson and Euclid

“I finished my examination in Euclid last eve & without a failure at any time. You can easily imagine how glad I am to get through with 4 books, for you have finished the whole *forever*.”

—Emily Dickinson, letter to her brother Austin, Dec. 10, 1847 (her 17th birthday).

“Emily was never floored. When the Euclid examination came and she had never studied it, she went to the blackboard and gave such a glib exposition of imaginary figures that the dazed teacher passed her with the highest marks.”

—Mabel Todd, notes on conversations with Emily’s sister Lavinia Dickinson.

Both quotations can be found in: Jay Leyda, *The Years and Hours of Emily Dickinson*, Vol. 1, Yale University Press, New Haven, 1960, p. 131.

—Submitted by Robert Haas, Cleveland Heights, OH