THE GAME OF HEX: A STUDY IN GRAPH THEORY AND ALGEBRAIC TOPOLOGY

MARK SCHACHNER

ABSTRACT. Hex is a game invented in the 1940s with the non-trivial property that no completed game can end in a tie. This paper studies the this fact, dubbed the Hex theorem, at its graph-theoretic roots, as well as its topological consequences, and presents an original generalization to the torus which preserves this property.

CONTENTS

1. INTRODUCTION

Hex is a two-player abstract strategy game played on a rhombic board of hexagons called *hexes*, most commonly of size 11 by 11. The two players, typically denoted by Red and Blue, take turns claiming hexes by placing counters on any unoccupied space on the board. The game continues until either Red has claimed an unbroken chain of hexes connecting the top and bottom boundaries of the board, or Blue has claimed an unbroken chain of hexes connecting the left and right boundaries of the board. Figure 1.1 shows a completed game in which Blue has won.

Hex was invented in 1942 by Piet Hein, and again independently in 1948 by John Nash. Nash studied the game from a game-theoretic standpoint, showing via a simple strategy-stealing argument that the second player has no winning strategy, and, moreover, with perfect play the first player can always win [1]. In this paper, however, we will not concern ourselves with the game-theoretic aspects of Hex, but rather with its topological properties. We will use this paper's foundational result to give elementary proofs of several classical results of algebraic topology, including a proof of the Jordan separation theorem. Next, we will explore generalizations of Hex, both in higher dimensions and on other topological spaces. In order to make such generalizations, we will make use of the notions of homology and cohomology, which we will exposit in Section 6.

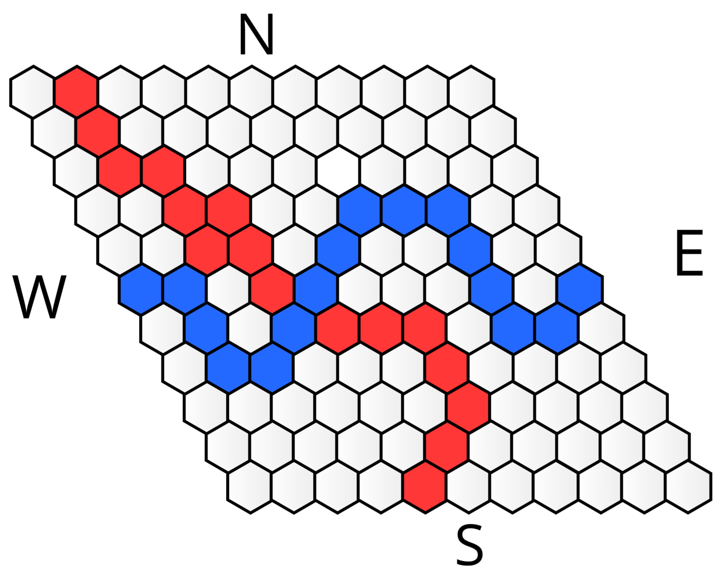

FIGURE 1.1. A completed game of Hex on an 11×11 board. Blue has won by connecting the left side (West) to the right side (East).

We begin by rigorously defining the Hex board as a graph imbedded in the plane, following with some variation the construction introduced by David Gale in [2]:

Definition 1.1. Let $n \in \mathbb{N}$.

(i) We call two points $(a_1, a_2), (b_1, b_2) \in \mathbb{N}^2$ *comparable* if $a_1 < b_1$ and $a_2 < b_2$, or $a_1 > b_1$ and $a_2 > b_2$.

(ii) The $n \times n$ *Hex board* is the graph H_n with vertex set $\{(a_1, a_2) \mid a_1, a_2 \in \mathbb{N}, 1 \leq a_1 \leq n, 1 \leq a_2 \leq n\}$. We call two points $(a_1, a_2), (b_1, b_2) \in H_n$ *adjacent* if they are comparable and $\max\{|a_1 - b_1|, |a_2 - b_2|\} = 1$; between such vertices we draw an edge. We call the sets $\{(i, n)\}, \{(n, i)\}, \{(i, 1)\}$, and $\{(1, i)\}$ *North, East, South,* and *West*, respectively.

(iii) A *game* of $n \times n$ Hex is a map $\Gamma : H_n \to \{0, 1, 2\}$, realized as a coloring of each vertex as uncolored, red, or blue. We say the game is *a win for Red* if there exists a connected subset $Q \subset H_n$ such that $\Gamma(Q) = \{1\}$, and two points $(a, n), (b, 1) \in Q$ which are in North and South, respectively. Similarly, Γ is a *win for Blue* if $Q \ni (1, a), (n, b)$ is connected and $\Gamma(Q) = \{2\}$. Lastly, a *full* game of Hex is a game Γ for which $\Gamma^{-1}(\{0\}) = \emptyset$.

The graph outlined in Definition 1.1 does not resemble exactly the classical Hex board; however, the two boards are graph isomorphic, and in fact it was this square board which Nash originally considered. It is pictured in Figure 1.2.

Some natural questions which emerge in a serious consideration of Hex include:

Question 1.2.

1. *Is a full game of Hex always a win for Red or Blue?*
2. *Can the discrete nature of Hex be used to prove results in indiscrete contexts?*
3. *Can the Hex board be generalized to higher dimensions, other spaces, and higher player counts in such a way that its crucial properties are preserved?*

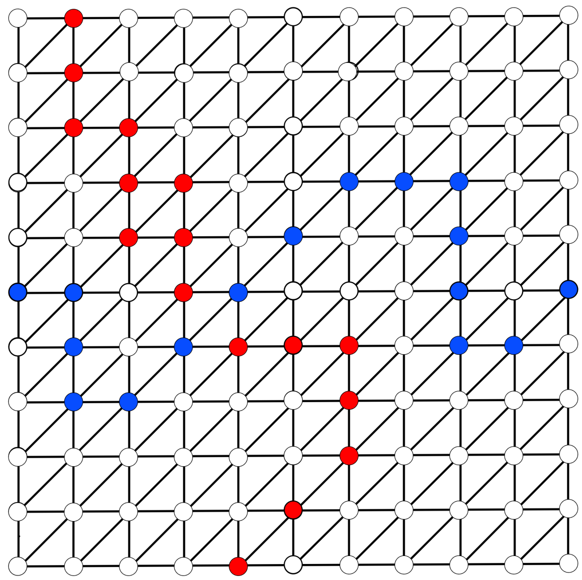

FIGURE 1.2. The same Hex game as in Figure 1.1, represented on the square Hex board.

We reserve a treatment of the last two questions for later in this paper. The first question, however, is a surprisingly deep result, and we waste no more time in investigating it.

2. The Hex Theorem

In playing Hex, one quickly notices that the most "offensive" moves and the most "defensive" moves tend to be the same. Intuitively, this makes sense; we would expect that to connect North to South is to block off our opponent from connecting East to West, and vice versa. Moreover, we might suspect that this local principle is reflected in the overall game structure, so that one player wins if and only if the other loses. This turns out to be true in the following sense:

Theorem 2.1 (Hex theorem). *Let $n \in \mathbb{N}$ and let Γ be a full game of $n \times n$ Hex. Then Γ is a win for exactly one player.*

We have in fact stated the Hex theorem in a slightly stronger form than it is usually stated; in most contexts the Hex theorem only guarantees that G is a win for at least one player. We will make no use of this weaker form.

Proof. We will first show that any full game of Hex has at least one winner, and then that any full game has at most one winner.

To show the former, we begin with some terminology: given a full game Γ we call an edge of the board *uniform* if it connects a red vertex to a red vertex or a blue vertex to a blue vertex, and *mixed* if it connects a red vertex to a blue vertex. We next augment the Hex board, adding $4n$ extra vertices in a "box" around the board, and connect them with edges to make an $(n+2)\times(n+2)$ Hex board without

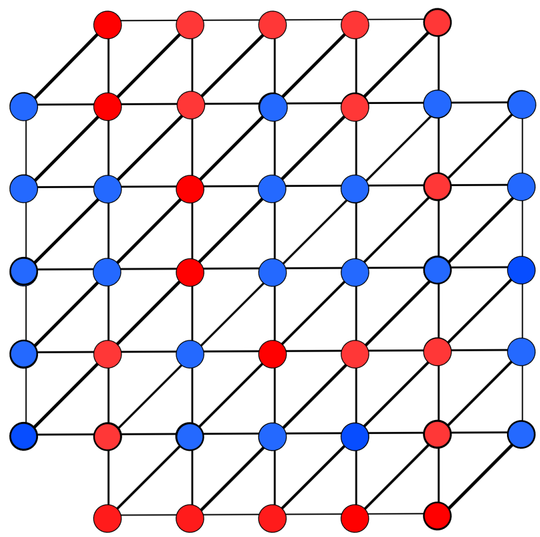

FIGURE 2.1. An augmented version of a full 5×5 game of Hex.

the corner vertices. We color the vertices directly above North and below South red, and those directly left of West and right of East blue. The result of this is shown for an arbitrary full game where $n = 5$ in Figure 2.1.

Now, we present an algorithm for finding a winning path for one of the players, adapted from that of David Gale in [2]. We begin by shading the triangular face in the bottom right. We next shade the triangular face adjacent to that face which shares a mixed edge with that face. Continuing, at each triangular face we encounter, we shade the face adjacent to that face (other than the face we shaded the step before) which shares a mixed edge with that face. The algorithm terminates if no mixed edge exists, or if the path leaves the board or returns to an already shaded face. Figure 2.2 gives the result of this process with our 5×5 example.

Now, we note the following facts about this process:

(i) Since the number of faces is finite, the process will always terminate in a finite number of steps.

(ii) At each step, we enter a triangular face with at least one red vertex and one blue vertex. Regardless of the color of the third vertex, therefore, the face will always have a mixed edge other than the edge we just crossed. Thus the algorithm never terminates in the case that no mixed edge exists.

(iii) In order to enter any face, that face cannot have 3 blue or 3 red vertices; thus it must have 2 of one color and one of the other. This implies that for any face in our path, two of its edges are mixed and the third edge is uniform. In passing through this face, we enter across one of the mixed edges and exit across the other; if we are then to return to this face, we would have to cross the third uniform edge, which by definition the algorithm never does. Thus the algorithm also cannot terminate in the case that we return to a previously shaded face (other than the face at which we began; however, note

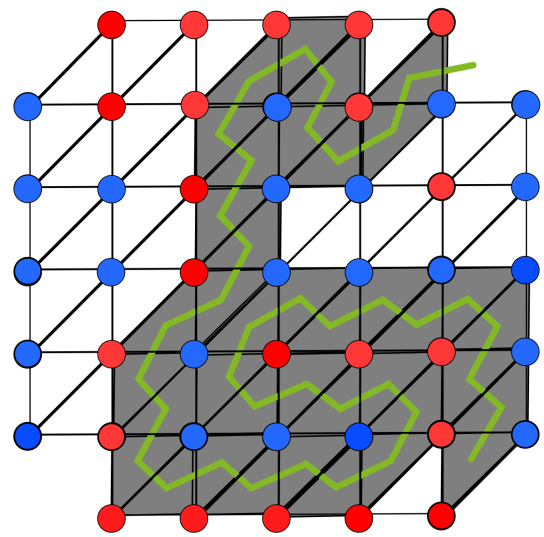

FIGURE 2.2. The same board as Figure 2.1, with the shading algorithm completed and the given path outlined.

that in order to return to where we started we would have to cross from the unbounded face, which the algorithm also never does).

(iv) Lastly, the resultant set of shaded faces has as boundary uniform edges only; since at the beginning of the algorithm the left side of the set is red and the right side of the set is blue, this is preserved throughout the algorithm, so that the final set has all red vertices on one side and all blue vertices on the other.

Together, the facts (i),(ii),(iii) imply that the algorithm terminates when the path leaves the board, which can only happen at the bottom-left, top-left, or top-right mixed edges. By fact (iv), in the first two of these cases, the left side of the set of shaded faces gives a winning set for Red; in the third case, the right side of the set of shaded faces gives a winning set for Blue. Thus a winning set always exists for at least one player.

To complete the proof, we show that any full game of Hex has at most one winner. In order to show this, we suppose for contradiction that an $n \times n$ Hex board exists which is a win for both Red and Blue. We again augment this board, this time adding four single vertices adjacent to North, South, East, and West, and four edges connecting them in a cycle. We then add a fifth vertex outside of the plane, and connect it to each of the four additional vertices. The result of this is depicted in Figure 2.3.

Now, considering the non-intersecting paths by Red and Blue within the board as well as the additional edges we added, we have constructed a planar embedding of the complete graph K_5, which is a contradiction as no such embedding exists. Therefore no Hex board exists which is a win for both players. This result along with the previous one implies the theorem.

$\square$

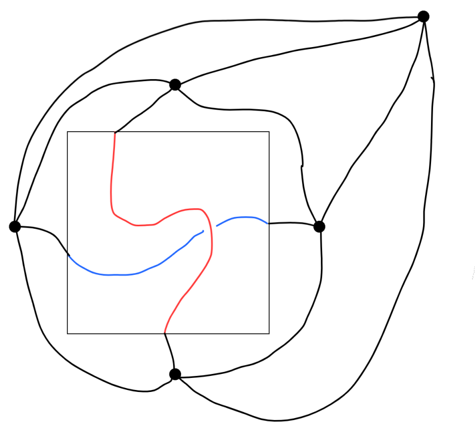

FIGURE 2.3. A different augmentation of a full Hex game, giving an impossible planar embedding of K_5. Note that the Red and Blue paths by hypothesis do not intersect.

The above proof shows that the theorem is fundamentally graph theoretic; however, in the next section we will show it has topological significance as well.

3. Topological Applications of the Hex Theorem

In addition to demonstrating that any game of Hex can be won by at most player, Figure 2.3 suggests a continuous analog, namely that if P_1, P_2 are paths connecting the top and bottom sides and the left and right sides of a square, respectively, then they must intersect in at least one point within the square. As in the proof of Theorem 2.1, if this were untrue then we could construct a planar embedding of the nonplanar graph K_5. In this section, we will follow this idea of applying the discrete Hex theorem to nondiscrete contexts, in order to show two major topological applications. The first of these is the classical result due to L. E. J. Brouwer:

Theorem 3.1 (Brouwer fixed-point theorem). *Let $f : I^2 \to I^2$ be continuous. Then there exists $x \in I^2$ such that $f(x) = x$.*

Again following the work of Gale, we shall in fact show that Theorem 3.1 and Theorem 2.1 are "equivalent" in the sense that either can be proven rather easily using the other. The proofs we present are somewhat streamlined versions of those in [2].

Proof of Theorem 3.1 from Theorem 2.1. For setup, let $\varepsilon > 0$ and write $f = (f_1, f_2)$. Since f is continuous and I^2 is compact, f is uniformly continuous, so that there exists $\delta > 0$ such that $\delta < \varepsilon$ and if $x, y \in I^2$ such that $||x - y|| < \delta$ then $||f(x) - f(y)|| < \varepsilon$. Fix $n \in \mathbb{N}$ such that $1/n < \delta/\sqrt{2}$.

Now, suppose for contradiction that for all $x \in I^2$ we have $||f(x) - x|| \geq \varepsilon\sqrt{2}$. This further implies that for each vertex $v = (v_1, v_2)$ in the $n \times n$ Hex board H_n, we have $\max\{|f_1(v/n) - v_1/n|, |f_2(v/n) - v_2/n|\} > \varepsilon$. We construct a full game of Hex Γ as follows: for each $v \in H_n$, color v red if $|f_1(x) - x_1| \leq |f_2(x) - x_2|$ and color v blue otherwise. We will show that Γ is a win for neither player, contradicting the Hex theorem.

First note that for any blue vertices $v = (v_1, v_2), w = (w_1, w_2) \in H_n$ such that

$$f_1(v/n) - v_1/n \geq \varepsilon \text{ and } w_1/n - f_1(w/n) \geq \varepsilon,$$

we cannot have that v, w are adjacent. This is because if such v, w were adjacent, then we would have $||v/n - w/n|| \leq \frac{\sqrt{2}}{n} < \delta$ but

$$\begin{aligned}
||f(v/n) - f(w/n)|| &\geq \max\{|f_1(v/n) - f_1(w/n)|, |f_2(v/n) - f_2(w/n)|\} \\
&\geq |f_1(v/n) - f_1(w/n)| \\
&\geq |f_1(v/n) - v_1/n + v_1/n - w_1/n + w_1/n - f_1(w/n)| \\
&\geq f_1(v/n) - v_1/n + w_1/n - f_1(w/n) - (w_1/n - v_1/n) \\
&\geq 2\varepsilon - 1/n \\
&\geq \varepsilon,
\end{aligned}$$

violating the uniform continuity of f; thus v, w are not adjacent. Using a similar argument, we may show that if v, w are blue and satisfy $f_2(v/n) - v_2/n, w_2/n - f_2(w/n) \geq \varepsilon$, then v, w cannot be adjacent.

Now, let S be a set of blue vertices in H_n intersecting both East and West. Observe that for any vertex v in East, we cannot have $v_1/n - f_1(v/n) \geq \varepsilon$ as this would imply $f(v/n) \notin I^2$, and similarly for any vertex in West we cannot have $v_1/n - f_1(v/n) < \varepsilon$. Thus S must at least one vertex which satisfies each of the above inequalities. But since such vertices cannot be adjacent by our above argument, S cannot be connected. Hence Γ cannot be a win for Blue. But by a similar argument, Γ also cannot be a win for Red. Thus Γ violates the Hex Theorem, a contradiction; our hypothesis, that for all $x \in I^2$ we have $||f(x) - x|| \geq \varepsilon\sqrt{2}$, must therefore be false. Since our choice of ε was arbitrary and I^2 is compact, therefore, the proof is complete.

□

The proof of Theorem 3.1 from Theorem 2.1 is lengthier, but straightforward given the following two lemmas:

Lemma 3.2. *Let $n \in \mathbb{N}$ and suppose f is a mapping of H_n into $\mathbb{R}^2$. Then there exists a unique piecewise linear map $f^\triangle : [1, n]^2 \to [1, n]^2$ which is continuous.*

We will call this extension $f^\triangle$ the *simplicial extension* of f on $[1, n]^2$.

Proof. Since H_n is a triangulation of $[1, n]^2$, any point $x \in [1, n]^2$ lies in the convex hull of three unique and mutually adjacent vertices v_1, v_2, v_3. By defining f in this convex hull as the unique plane containing the three vertices, and repeating this for each triangular face of H_n, we obtain a piecewise linear map. Moreover, f is trivially continuous on the faces of H_n, and also continuous on its edges since the planes intersect on the edges of of H_n. □

Lemma 3.3. *Let $x_1, x_2, x_3 \in \mathbb{R}^2$, put $f(x_i) = y_i$, and let $f^\triangle$ be the simplicial extension of f on the convex hull of x_1, x_2, x_3. Then $f^\triangle$ has a fixed point if and only if the convex hull of $\{y_i - x_i\}$ contains zero.*

Proof. Since $f^\triangle$ is piecewise linear and continuous, for any point z, we have that z lies in the convex hull of $\{x_i\}$ if and only if $f^\triangle(z) - z$ lies in the convex hull of $\{y_i - x_i\}$. Thus 0 lies in the latter if and only if there exists z in the former such that $f^\triangle(z) = z$. □

We now proceed to the proof of the Hex theorem from the Brouwer fixed-point theorem:

Proof of Theorem 2.1 from Theorem 3.1. Let Γ be a full game of Hex on H_n, and suppose for contradiction that Γ has no winner. This implies no connected set of red vertices intersects both North and South, and similarly no connected set of blue vertices intersects both East and West. We define a function $f : H_n \to H_n$ as follows: for $v \in H_n$,

- if v is red, then f translates v up by one unit if v is connected to South by a set of red vertices, and down one unit otherwise; and
- if v is blue, then f translates v to the left one unit if v is connected to West by a set of blue vertices, and to the right one unit otherwise.

First, note that a vertex in North is never translated up by f, since if it were that would imply that a connected set of red vertices intersects North and South, contradicting our hypothesis that Γ has no winner. Similarly, the vertices in East, South, and West are never translated right, down, or left, respectively. Thus f maps H_n to itself. Secondly, observe that two adjacent vertices of H_n are never translated in opposite directions, since any connected set containing one must contain the other. Thus for any $u, v, w \in H_n$, the convex hull of $\{f(u), f(v), f(w)\}$ cannot contain zero.

By Lemma 3.3, the above facts imply that the simplicial extension $f^\triangle$ of f can have no fixed point, a contradiction of Theorem 3.1. Hence Γ must be a win for at least one player.

□

The second direction of the above proof makes heavy use of the fact that all bounded faces of H_n are triangular; this fact will prove crucial in the following sections.

For a second application of the Hex theorem to general topology, we present a proof of the Jordan separation theorem given the Hex theorem, a modification of an argument due to Maehara [3]. The Jordan separation theorem is a slightly weaker form of the Jordan curve theorem, famous for its simple statement and elusive proof. We state it thus:

Theorem 3.4 (Jordan separation theorem)**.** *Let $\gamma : S^1 \to \mathbb{R}^2$ be continuous and injective. Then the set $\mathbb{R}^2 \setminus \gamma(S^1)$ has exactly two components, exactly one of which is bounded.*

Proof. We begin by considering a Jordan curve in the plane, i.e., the image C of a continuous injection $\gamma : S_1 \to \mathbb{R}^2$. Let $\pi_i\gamma : S^1 \to \mathbb{R}$ denote the projection of γ into the i-th coordinate. Note that both $\pi_1\gamma$ and $\pi_2\gamma$ are continuous functions and their shared domain S^1 is compact; thus we may choose $A_1, A_2 \in$

$\pi_1\gamma(S^1)$ such that $A_1 = \inf\{\pi_1\gamma(x) : x \in S_1\}$ and $A_2 = \sup\{\pi_1\gamma(x) : x \in S^1\}$, and $B \in \mathbb{R}$ such that $C \subset [A_1, A_2] \times [-B, B]$. For later convenience, we fix the following points:

- $a = (a_1, a_2) \in C$ such that $a_1 = A_1$;
- $b = (b_1, b_2) \in C$ such that $b_1 = A_2$;
- $N = (1/2(A_1 + A_2), B)$;
- $S = (-1/2(A_1 + A_2), B)$;
- $l = \gamma(r)$ for $r \in S^1$ such that $\pi_1\gamma(r) = 1/2(A_1 + A_2)$ and $\pi_2\gamma(r)$ is maximal;
- $m = \gamma(s)$ for $s \in S^1$ such that $\pi_1\gamma(s) = 1/2(A_1 + A_2)$, there exists an arc $C' \subset C$ such that $a, b \notin C'$, but $l, m \in C'$, and $\pi_2\gamma(s)$ is minimal;
- $p = \gamma(t)$ for $t \in S^1$ such that $\pi_1\gamma(t) = 1/2(A_1 + A_2)$, there exists an arc $C' \subset C$ such that $a, b, l \notin C'$, but $p \in C'$, and $\pi_2\gamma(t)$ is maximal;
- $q = \gamma(u)$ for $u \in S^1$ such that $\pi_1\gamma(u) = 1/2(A_1 + A_2)$, there exists an arc $C' \subset C$ such that $a, b \notin C'$, but $p \in C'$, and $\pi_2\gamma(u)$ is minimal.

Intuitively, the points a, b represent the leftmost and rightmost points of C, and the points N, S are upper and lower bounds for $\pi_2\gamma$. This divides C into a "top" and "bottom" arc; the points l, m are the y-maximal and y-minimal points of the top arc and p, q are the y-maximal and y-minimal points of the bottom arc. This is depicted in Figure 3.1, adapted from Maehara.

Now, in order for the theorem to be proved, we need to show that (i) $\mathbb{R}^2 \setminus C$ has exactly one unbounded component, and (ii) there exists a unique bounded component of $\mathbb{R}^2 \setminus C$. The proof of fact (i) follows in a straightforward manner from the fact that C is bounded and $\mathbb{R}^2$ is unbounded. We will show that the negation of fact (ii) implies the negation of the Hex theorem, giving a contradiction.

Let U denote the component of $\mathbb{R}^2 \setminus C$ containing the midpoint x of m, p, and suppose for contradiction that U is unbounded. In particular, U is then not bounded by $[A_1, A_2] \times [-B, B]$, so there exists $v \in U \setminus [A_1, A_2] \times [-B, B]$. Since C is closed and $\mathbb{R}^2$ is locally path connected, we have U is path connected, so a path α exists between x and v, which must intersect the boundary D of $[A_1, A_2] \times [-B, B]$. Suppose without loss of generality that α first intersects D on its lower half; there then exists a path α' from this intersection to s containing neither of a, b.

Now, let $\phi_1, \phi_2 : [-1, 1] \to [A_1, A_2] \times [-B, B]$ be two continuous functions with images

$$\phi_1([-1, 1]) = \overline{nl} \cup \widehat{lm} \cup \overline{mx} \cup \alpha',$$
$$\phi_2([-1, 1]) = \widehat{ab},$$

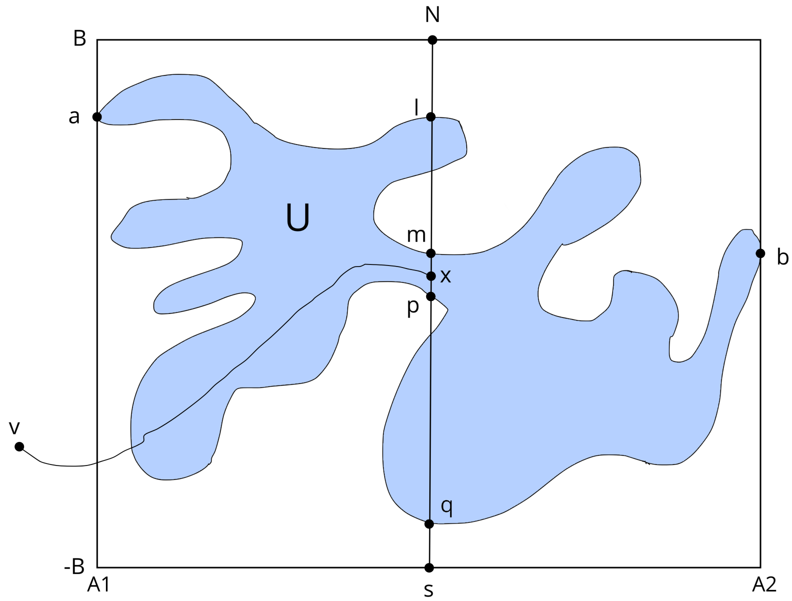

FIGURE 3.1. A Jordan curve C, annotated with bounds and points relative to our proof.

where the notation $\widehat{ab}$ denotes the smallest connected subset of C containing a, b. By construction, ϕ_1 connects the top of $[A_1, A_2] \times [-B, B]$ to the bottom side, and ϕ_2 connects its left side to its right side; moreover, by our hypothesis, these paths do not intersect. Since both ϕ_i are continuous functions on compact sets, their images are compact as well, so that there exists some $\delta > 0$ such that for all $p_1 \in \phi_1([-1,1]), p_2 \in \phi_2([-1,1])$ we have $||p_1 - p_2|| > \delta$. Let $k \in \mathbb{N}$ be such that

$$\frac{||(A_1, -B) - (A_2, B)||}{k-1} < \delta,$$

and let H denote the image of the $k \times k$ Hex board under the linear mapping

$$(x, y) \mapsto \left(\frac{A_2 - A_1}{k} x + A_1, \frac{2B}{k} y + B \right),$$

which intuitively "fits" H_k "onto" $[A_1, A_2] \times [-B, B]$. Now, we form two subsets $P_1, P_2 \subset H$ as follows; for each $z \in H$, we put $z \in P_1$ if there exists

$$z' \in \phi_1([-1,1]) \text{ such that } ||z - z'|| < \delta/2,$$

and $z \in P_2$ if there exists

$$z' \in \phi_2([-1,1]) \text{ such that } ||z - z'|| < \delta/2.$$

Note that, by our choice of k, P_1 and P_2 are disjoint, and by the continuity of ϕ_1, ϕ_2 they are both connected. Lastly, P_1 connects the top part of H to the bottom, and P_2 connects its left and right sides. Since H is graph isomorphic to the Hex board H_k, then, we can use P_1, P_2 to color H_k so that both Red and Blue have winning paths, contradicting the Hex theorem. Thus U must be bounded.

The next part of the proof, that B is the sole bounded component of $\mathbb{R}^2 \setminus C$, follows a similar theme. For if a second bounded component W existed, then we could construct two non-intersecting paths as follows: ϕ_1 is given by the path $\overline{nl} \cup \widehat{lm} \cup \overline{mp} \cup \widehat{pq} \cup \overline{qs}$, and ϕ_2 connects a to b and lies entirely within W. Then, choosing k as above, we can use these paths to construct a full game of Hex with two winners. Thus U is the unique bounded component of $\mathbb{R}^2 \setminus C$, and the theorem is proved. □

The results of Theorem 3.1 and Theorem 3.4 show that the Hex theorem is in fact a rather deep observation which has consequences in contexts outside of its discrete graph-theoretic construction. In the next section, however, in order to get a better idea of what gives the game this property, we will return to the context of graph theory and attempt to find generalizations of Hex, both in the plane and on a bigger class of topological spaces.

4. Generalizing Hex on the Plane

To begin generalizing the ideas of Hex to a broader class of graphs, we will need a few initial definitions:

Definition 4.1. Let G be a planar graph with vertex set V, embedded in the plane without intersection. The *boundary* ∂G of G (with respect to the given embedding) is the set of vertices in V which are adjacent to the unbounded region of the plane.

We will define our generalization of Hex on the set of connected planar graphs. On such a graph G, we can cover ∂G by four sets N, E, S, W, and then "play Hex" on G by coloring vertices red or blue until there exists either a red path connecting N and S or a blue path connecting E and W. Figure 4.1 depicts an example of this sort of game; we will define it more precisely as follows.

Definition 4.2. Let G be a connected planar graph, and let $N, E, S, W \subset \partial G$ be connected sets such that each of $N \cap E, E \cap S, S \cap W$, and $W \cap N$ are nonempty and the four sets together cover ∂G. Then a *game of Hex* on G is a map $\Gamma_G : G \to \{0, 1, 2\}$. We say Γ_G is a *win for Red* if there exists a connected subset $Q \subset G$ such that $\Gamma_G(Q) = \{1\}$, and two vertices $v_1, v_2 \in Q$ which are in N and S, respectively. Similarly, Γ_G is a *win for Blue* if there exists a connected subset $Q \subset G$ such that $\Gamma_G(Q) = \{2\}$, and two vertices $v_1, v_2 \in Q$ which are in E and W, respectively. Again, a *full* game of Hex is a game Γ_G for which $\Gamma_G^{-1}(\{0\}) = \emptyset$.

In order for this to be a "good" generalization of Hex, we would like to preserve the main property of Hex, that no ties can occur. This is given by the following condition on G:

Theorem 4.3. *Let G be a connected planar graph with boundary covered by sets N, E, S, W as given in Definition 4.2, and let Γ_G be a full game of Hex on G. Then Γ_G is a win for exactly one player if G contains an induced subgraph G' such that $\partial G' = \partial G$ and each face of G' is triangular.*

Proof. We will show the conclusion given that each face of G is triangular. This gives the result because if some subgraph of G with the same boundary of G is a win for some player, then G is also a win for that player. Note that by a similar argument as that in Theorem 2.1, we have that Γ_G is a win for at most one player, since if not we could construct a planar embedding of K_5 using the winning paths for Red and Blue. Thus we need only show that Γ_G is a win for at least one

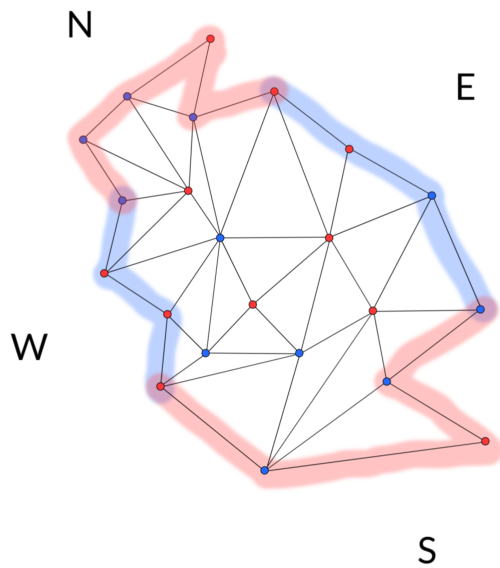

Figure 4.1. (See Definition 4.2) An arbitrary connected planar graph with boundary regions N, E, S, W. Note that the boundary regions must cover ∂G with overlap. By Theorem 4.3, this game is a win for exactly one player (here Blue), since all bounded faces are triangles.

player. To do this, we construct a generalization of the algorithm presented in the proof of Theorem 2.1. This is done by first adding a vertex in the unbounded region of the plane for each vertex in N, E, S, W, with vertices in multiple sets being "double-counted". We connect each of these extra vertices in a cycle, and triangulate the resultant annular region between G and this cycle graph. The augmentation is finished by coloring the vertices corresponding to N and S red and those corresponding to E and W blue. With this augmentation, we can then follow an analagous algorithm on the augmented graph to that on the Hex board, this time beginning on any triangular face whose vertices intersect two opposite-colored boundary regions. The result of such an algorithm is depicted in Figure 4.2.

Now, by similar arguments to those in the proof of Theorem 2.1, the set of shaded faces terminates at one of four edges at the intersections of N, E, S, and W, and aside from these edges must be bounded by uniform edges only. Thus it is bounded by a sequence of red vertices connecting N and S or a sequence of blue vertices connecting E and W, so that Γ_G must be a win for at least one player.

□

Note that Theorem 4.3 implies the strong Hex theorem as a weak corollary. It thus implies a very important fact, namely that the exactly-one-winner property of the original Hex board stems from the fact that it can be thought of as a triangulation of the disk. We will use this idea to generalize Hex in the following section, by considering triangulations of arbitrary compact orientable surfaces.

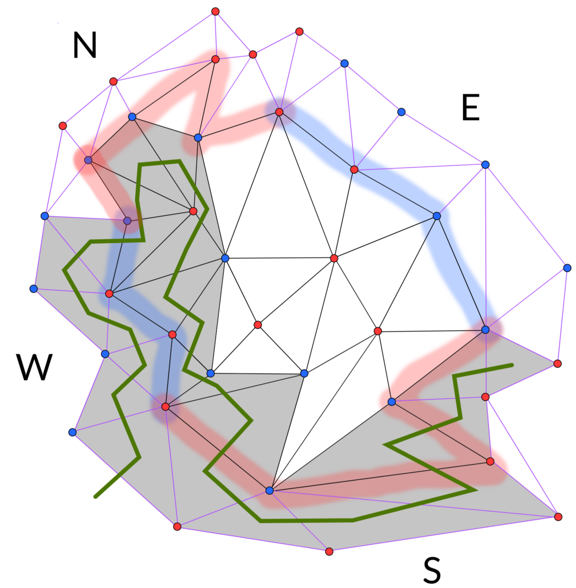

FIGURE 4.2. Here, we have augmented Figure 4.1 with boundary vertices which are colored red or blue depending on which boundary region they correspond to. Then we follow a similar face-shading algorithm as in the proof of the Hex theorem.

5. Homology and Cohomology

Now, we would like to generalize our study of Hex from graphs on the plane to graphs on more diverse spaces. In Section 6, we will define a variation of Hex playable on the class of compact orientable surfaces; in order to do this we will need the following brief exposition on simplicial homology and cohomology.

Definition 5.1. Let $V = \{v_0, v_1, \ldots, v_n\} \subset \mathbb{R}^n$.

(i) We say two orderings of the points in V are *of the same orientation* if they differ by an even permutation. Forming the quotient of V modulo this equivalence relation, we obtain a two-element set; we call each element of this set an *orientation* of V.

(ii) We say V is *affinely independent* if no hyperplane of dimension less than n contains all $v_i \in V$. Equivalently, V is affinely independent if the set $\{v_i - v_0\}_{1\leq i\leq n}$ is linearly independent.

(iii) Let $V \subset \mathbb{R}^n$ be an affinely independent set. The smallest convex set Δ containing all $v \in V$ is called an *n-simplex*. The points in V are called the *vertices* of Δ.

The n-simplex may be broadly thought of as the higher-dimensional generalization of the triangle and tetrahedron. We next define simplicial complexes, which will be the "boards" on which our generalization will be played.

Definition 5.2. An *oriented n-simplex* is an n-simplex together with some orientation of its vertices. If C is a set of oriented simplices such that the face of each simplex in C is in C and the intersection of any two simplices in C is a face of each simplex, then we call C a *simplicial complex.*

The above preliminary definitions, as well as the fact that any closed surface has a decomposition into triangles, begins to suggest that a game analogous to Hex could exist, in which players take turns claiming vertices on a simplicial complex which covers some compact orientable surface. However, since such surfaces have no boundary by definition, the idea of partitioning the boundary as in our generalization in Section 4 fails to generalize. Thus we will need to adjust the win condition so that players attempt to go "all the way around" the surface in perpendicular directions, rather than simply connecting two ends of a boundary. To make this idea more precise, we continue our exposition of H_1, the first homology group:

Definition 5.3. Let C be a simplicial complex.

(i) An *n-chain* in C is a formal linear combination of simplices in C with coefficients in $\mathbb{Z}$.

(ii) Let σ be an oriented n-simplex, given by some ordering of its vertices $(v_0, \dots, v_n)$. The *boundary* of σ is the $(n-1)$-chain defined by

$$\partial_n \sigma = \sum_{0 \le i \le n} (-1)^i (v_0, v_1, \dots, \hat{v_i}, \dots, v_n),$$

where $(v_0, v_1, \dots, \hat{v_i}, \dots, v_n)$ denotes the simplex with all vertices except v_i. If $\Delta = \sum n_i \sigma_i$ is an n-chain, the boundary of Δ is the $(n-1)$-chain

$$\partial_n \Delta = \sum (-1)^i \partial_n \sigma_i.$$

It is now relatively straightforward to see that, for $n = 1$, the set Ker ∂_1 consists exactly of those chains of 1-simplices (i.e., straight lines) which begin where they end; consequently, we call elements of Ker ∂_n *n-cycles.* There is one last stop on our exposition before we can rigorously define H_n :

Lemma 5.4. *Let $n \in \mathbb{N}$. Then $\partial_n \partial_{n+1} = 0$.*

Proof. We proceed through the computation as follows: let σ be an $(n+1)$-simplex. Then we have

$$\begin{aligned} \partial_n \partial_{n+1} \sigma &= \partial_n \left(\sum_{0 \le i \le n+1} (-1)^i (v_0, v_1, \dots, \hat{v_i}, \dots, v_{n+1}) \right) \\ &= \sum_{0 \le j \le n+1} \left(\sum_{\substack{0 \le i \le n+1 \\ i \ne j}} (-1)^i (v_0, v_1, \dots, \hat{v_i}, \dots, \hat{v_j}, \dots, v_{n+1}) \right). \end{aligned}$$

The sum on the inside can then be divided into the cases where $i < j$ and where $j < i$; for given i, j these cancel out perfectly, as the pairs (i, j) and (j, i) both appear in the calculation. Thus the total sum is zero, and the lemma is proved. $\square$

Lemma 5.4 suggests the following definition, which will conclude our exposition of simplicial homology:

Definition 5.5. Let X be a simplicial complex. The n-th *homology group* of X, denoted H_n, is the quotient group Ker $\partial_n/$ Im ∂_{n+1}. Elements of H_n are called *homology classes;* for a cycle c, we write $[c]$ for the class in H_n of which c is a representative.

The abelian groups and boundary homomorphisms between them form a natural structure called a *chain complex*:

$$\cdots \xrightarrow{\partial_3} C_2^* \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \xrightarrow{\partial_0} 0.$$

By replacing the groups with their dual groups, and the homomorphisms with their dual homomorphisms, we may "reverse" this chain and obtain a *cochain complex*

$$\cdots \xleftarrow{\partial^3} C_2^* \xleftarrow{\partial^2} C_1^* \xleftarrow{\partial^1} C_0^* \xleftarrow{\partial^0} 0,$$

from which we can define the n-th *cohomology group* of C by

$$H^n = \operatorname{Ker} \partial^n / \operatorname{Im} \partial^{n-1}.$$

As we generalize Hex to the torus T, we will use without proof the fact that

$$H^n(T) \cong \begin{cases} \mathbb{Z} & \text{if } n = 0, \\ \mathbb{Z} \oplus \mathbb{Z} & \text{if } n = 1, \\ \mathbb{Z} & \text{if } n = 2, \end{cases}$$

and that as a ring $H^*(T)$ is isomorphic to the exterior algebra $\Lambda[a, b]$. We shall also use the fact that H^n, H_n are homotopy invariants, as is the *cup product* $\cup :$ $H^n \times H^m \to H^{n+m}$ given by the bilinear extension of the following operation on basis elements of $H^{n+m}(X)$:

$$(\phi \cup \psi)(\sigma) = \phi(\sigma \mid_{[v_0,\ldots,v_n]})\psi(\sigma \mid_{[v_{n+1},\ldots,v_m]}).$$

This cup product on cohomology groups has a dual multiplicative structure $\cap$ known as the *cap product;* via Poincaré duality this gives, for certain generators v, x, y, σ,

$$H_n(T) \cong \begin{cases} \mathbb{Z}v, & \text{if } n = 0, \\ \mathbb{Z}x \oplus \mathbb{Z}y, & \text{if } n = 1, \\ \mathbb{Z}\sigma & \text{if } n = 2, \end{cases}$$

as well as that $x \cap x = y \cap y = 0$ and $x \cap y = -y \cap x = v$. Applying these facts, a straightforward calculation gives

$$(ax + by) \cap (cx + dy) = ad - bc = \det \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

so that if two elements α, β of $H_1(T)$ are linearly independent then the above determinant is nonzero, and thus their cap product is nontrivial. This means that, for two curves $C_1, C_2 : S^1 \to T$, if we let $[C_1], [C_2]$ be the images of the generator of $H_1(S^1)$ via C_1, C_2 respectively, then $[C_1] \cap [C_2]$ is nontrivial only if C_1, C_2 intersect geometrically.

Given these facts, we are ready to move to our variation of Hex on the torus.

6. Hex on the Torus

In this section, we present one possible variation of Hex which can be played on the torus. Since this space is without boundary, we cannot generalize Hex to it directly; however, by appealing to the tools of homology we can define a playable variation. Recall that a victory in Hex, both in its original form and in our general form on triangulations of the disk, consisted of a monochromatic path connecting two sections of the boundary. For our analog on T^2, we will instead define victory by the completion of a monochromatic cycle representative of a given homology class. In order to make this definition, we need the following lemma:

Lemma 6.1. *Let T be the 2-dimensional torus, triangulated by a simplicial complex X. Then there exist two sets $R, B \subset H_1(X)$ such that $R \sqcup B = H_1(X) \setminus \{0\}$ and, for any classes $[C_1] \in R, [C_2] \in B$, we have that $[C_1] \cap [C_2] \neq \emptyset$.*

Proof. We partition the module $M = H_1(T) \cong \mathbb{Z}x \oplus \mathbb{Z}y$ into two disjoint sets R, B which cover $M \setminus \{0\}$ such that for any $v \in R, w \in B$ we have v, w are linearly independent. This can be accomplished in multiple ways; one way, for instance, is for R to contain all elements (except for the element with all zero coefficients) whose coefficients multiply to a non-negative integer, and for B to contain all elements whose coeffiecients multiply to a negative integer. We then have that any two cycles intersect algebraically if and only if their product in M is nonzero. By the last section, this occurs if and only if the two cycles are independent; we thus have that if $a \in R, b \in B$, then a, b are linearly independent and so must intersect. □

We now summit our exploration of Hex with the following definition:

Definition 6.2. Let T be the 2-dimensional torus, triangulated by a simplicial complex X.

(i) The *n-skeleton* of X, denoted X^n, is the space consisting of all simplices in X of dimension at most n. We consider X^1 as a graph with vertex set X^0 and the 1-simplices of X as its edges.

(ii) A *game of Hex* on X is a map $\Gamma_X : X^1 \to \{0, 1, 2\}$ realized as before as a coloring of each 0-simplex as uncolored, red, or blue, along with a pair (R, B) of subsets of $H_1(X)$ which satisfy the conditions in Lemma 6.1. We say Γ_X is *a win for Red* if there exists a connected subset $Q \subset X^0$ such that $\Gamma(Q) = \{1\}$ and that Q is a cycle representative of a homology class in R. Similarly, we say Γ_X is a win for Blue if a connected cycle $Q \subset X^0$ exists which satisfies $\Gamma(Q) = \{2\}$ and is a cycle representative of a homology class in B.

Theorem 6.3. *Let Γ_X be a full game of Hex on a triangulation X of the torus. Then Γ_X is a win for exactly one player.*

Proof. The fact that Γ_X is a win for at most one player stems from Lemma 6.1; this is because, if we have chosen R, B in the manner specified by the lemma, then, if Red completes a cycle $C \in R$, any cycle in B must intersect C and thus cannot be monochromatically blue. Similarly, if Γ_X is a win for Blue it cannot be a win for Red. Hence Γ_X is a win for at most one player.

To show that Γ_X is a win for at least one player, we follow the algorithm used in our Hex theorem proof one final time. We begin on a triangular face of the 1-skeleton of X with at least one mixed edge; if no such face exists, then the entire graph is monochromatic and the game clearly has at least one winner. We then cross the mixed edge into another triangular face; by similar arguments to those we have made previously regarding this algorithm, there exists one unique way to continue this process until it terminates. However, since X has a finite number of faces and no boundary, the algorithm must terminate at a face previously visited; moreover, this must be the face at which we began, again by similar arguments to those made previously. This means we have completed a cycle of shaded faces; in fact, since the set of faces is again bounded by uniformly red edges on one side and uniformly blue edges on the other, we have completed two homologous cycles, one red and one blue. If the class which these cycles represent is in R or B, we have found a winning path for at least one player, so suppose not. Since R, B cover $H_1(X) \setminus \{0\}$, this implies that the cycle we have completed is homologous to zero. We then repeat this algorithm, each time starting at an unshaded face with at least one mixed edge; if at any time we complete a cycle representative of a nonzero homology class we are done, so we need only consider the case where we cover the entirety of X by cycles homologous to zero, collapsing any triangles with all uniform edges to a point. Now, fix a cycle in such a coloring, and note that it must divide T into two components. Thus any additional cycles we draw must either be surrounded by this cycle or must completely surround it; in either case, the shaded region remains homeomorphic to a disk. However, at the end of the process we have assumed that the shaded region covers the entirety of T; since T cannot be covered by a set homeomorphic to a disk, the process can not terminate in a finite number of steps. This is a contradiction; hence, no such coloring can exist, so that Γ_X is a win for at least one player. □

While this variation of Hex certainly presents a non-trivial application of some of the main tools of algebraic topology, it is worth considering further that the version of the Hex theorem outlined in Theorem 6.3 and its proof retain many of the characteristics of the original Hex theorem. With this in mind, further investigation could involve applying Theorem 6.3 to prove results in algebraic topology similar to or in generalization of those in Section 3 similarly to how Theorem 2.1 can be used to more easily prove difficult theorems.

Acknowledgments

I am highly grateful to my mentor, Weinan Lin, for advising me throughout this project and supplying me with the tools I needed. I also thank Professor May and all the faculty who gave time and energy to organizing and executing the REU program this year.

References

[1] Gardner, M. (1959). The game of Hex. Hexaflexagons and Other Mathematical Diversions: The First Scientific American Book of Puzzles and Games, 73-83.
[2] Gale, D. (1979). The game of Hex and the Brouwer fixed-point theorem. The American Mathematical Monthly, 86(10), 818-827.
[3] Maehara, R. (1984). The Jordan curve theorem via the Brouwer fixed point theorem. The American Mathematical Monthly, 91(10), 641-643.
[4] Hatcher, A. (2002). Algebraic Topology. Cambridge University Press.
[5] Munkres, J. (2014). Topology. Pearson Education.
[6] http://math.ucr.edu/ res/math246A/cuplength.pdf