HEX

SP.268
SPRING 2011

1. Introduction

The game of Hex was first invented in 1942 by Piet Hein, a Danish scientist, mathematician, writer, and poet. In 1948, John Nash at Princeton re-discovered the game, which became popular among the math graduate students at Princeton. They called Hex either "Nash" or "John", though the latter referred to the hexagonal bathroom tiles that they played the game on. In 1952, Parker Brothers, Inc. popularized the game as "Hex" [1]. We will show that given perfect play, the first player wins. We will also show that Hex always has a winner – as in, Hex cannot end in a tie.

2. Game Rules

Hex is played with two people, on a diamond-shaped board made up of hexagonal cells. The board dimensions can vary, but the typical size is 11×11. Two opposite sides of the board are labeled "black", and the remaining two sides are labeled "white". One of the players has a supply of black tiles while the other player has a supply of white tiles. The players alternate turns to place their tile on any unoccupied space on the game board, with the goal of forming an unbroken chain of tiles (of his own color of course) linking his two regions. Figure 1 shows a Hex game board with the black and white regions. Some game pieces have already been played. Although the rules of Hex are simple, the game can provide insights for many mathematical concepts.

The Hex Theorem states that a game of Hex cannot end in a draw. The only way to keep the opponent from building a winning chain is to build a winning chain first. Although the Hex Theorem is intuitive, proving it can require invoking complicated topological results. John Nash is said to have proven the Hex Theorem, but he may not have bothered to publish the proof. David Gale gave a simple proof of the Hex Theorem based on graph theory and also showed the equivalence between the Hex Theorem and the Brouwer Fixed Point Theorem [2].

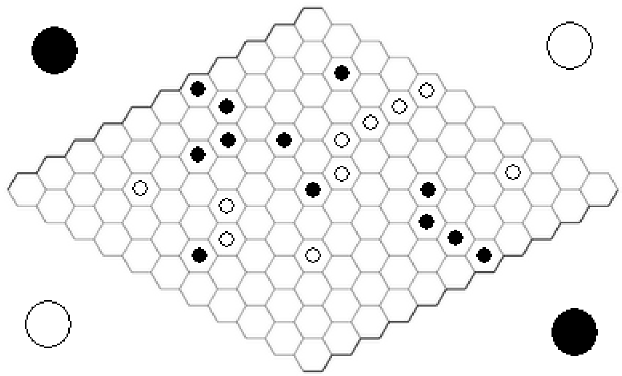

FIGURE 1

3. First-Player Win

For the typical 11×11 board, Hex is ultra-weakly solved. That means given perfect-play, we know at the beginning of the game who will win – the first player. The proof for this is a simple strategy-stealing argument.

First, we state two lemmas:

(1) Having extra/random pieces of your own color lying on the board cannot hurt you.
(2) Hex cannot end in a draw.

To prove the first lemma, suppose that there is an extra piece at position x on the board. If x is part of your winning strategy, then on the turn when you should be playing at position x, you could instead lay down another piece somewhere else. If x is not part of your winning strategy, then you would not care that it is occupied.

We defer the proof of the second lemma to the next section.

To prove that the first player must win Hex, let there be two players A and B: player A goes first and B goes second. Suppose for the sake of contradiction that B, the second player, has a winning strategy. Then A can just play a random move somewhere on the board. When B plays, he is effectively the first player. Now the original first player, A, is effectively the "second" player and can play the winning strategy for the rest of the game. By Lemma 1, having an extra piece from that first random move will never hurt A. Since A has "stolen" the winning strategy from B, B must not win. By Lemma 2, Hex cannot end in a draw, and since B cannot win, then A must be the winner. Then the first player must win Hex.

Comments:

- Someone in class asked, what is keeping B from stealing the strategy back? As in, what if B plays another bogus move, and we're back at A being the effective first player? We are assuming perfect play, so if A's winning strategy beats B when B is playing the best he can, that strategy will also beat B when he is playing less than perfectly.
- Given this, we would then wonder, how can A steal the strategy? Player B had the winning strategy, so if A started out playing a random move, then B can just keep being smart and beat A, right? This is the heart of the contradiction and our proof. Given that A played a random move and has effectively made B the first player, B is no longer in the position to use the winning strategy even though our assumption would imply that B can just keep playing perfectly and beat A. Hence, our assumption that B had the winning strategy must have been flawed.
- This strategy-stealing argument can be applied to any other symmetric game where having an extra move or game piece on the board can never hurt you. An example would be tic-tac-toe, though that game is strongly solved anyway, so we would not need strategy-stealing to know who will win.
- We also needed Lemma 2 to rule out the case of a tie, because otherwise, knowing that the second player will not win does not imply that the first player will win.
- Knowing that the first player wins is great, but unfortunately, this proof doesn't give us an actual winning strategy that we can use.

In this proof, we skipped over the proof of Lemma 2, but for good reason, because it is a theorem all in its own, and will take some work to prove.

4. The Hex Theorem

To prove the Hex Theorem, we begin with a graph lemma.

Lemma 4.1. *A finite graph whose vertices have degree at most two is the union of disjoint subgraphs, each of which is either* `(i)` *an isolated node,* `(ii)` *a simple cycle,* `(iii)` *a simple path.*

Proof. We induct on the number of edges in a graph. Consider a graph g with N nodes. Each node can have degree at most two, so g can have at most N edges. For simplicity, we denote a graph with k edges as g_k.

In the base case, g_0, all the nodes are isolated. When a graph has $n+1$ edges, we randomly choose an edge to remove, call it (u, v). The the nodes u and v now have degree at most 1 since they had degree at most two before we removed edge (u, v). Therefore u and v cannot be on any cycles. By assumption, g_n is the union of disjoint isolated nodes, simple cycles, and simple paths. We now add (u, v) back into the graph. The subgraphs that were disjoint from u and v in g_n are unchanged by the addition of (u, v), and the nodes u and v are now either on the same simple path or cycle. Therefore, g_{n+1} is also the union of disjoint isolated nodes, simple cycles, and simple paths. Hence the lemma is true for all g_k with $0 \leq k \leq N$. □

For simplicity, we substitute the black and white colors in the game with `x`'s and `o`'s, respectively. We represent the game board as a graph $G = (V, E)$, with a set of nodes V and a set of edges E. Each vertex of a hexagonal board space is a node in V, and each side of a hexagonal board space is an edge in E. We create four additional nodes, one connected to each of the four corners of the core graph; call these new nodes u_1, u_2, u_3, and u_4 and the edges that connect them to the core graph e_1, e_2, e_3, and e_4. An X-face is either a tile marked with an `x` or one of the regions marked X or X'. Similarly, O-face is either a tile marked with an `o` or one of the regions marked O or O'. Hence, the edges e_1, e_2, e_3, e_4 lie between an X-face and an O-face since the regions X, X', O', and O' are considered 'faces' as well. Figure 2 shows a Hex board with the X and O notation.

To prove the Hex Theorem, it suffices to show that two vertices out of u_1, u_2, u_3, and u_4 are connected by a simple path. The hexagonal tiles traced out by this simple path contain a winning chain.

Theorem 4.2 (Hex Theorem)**.** *If every tile of the Hex board is marked either* `x` *or* `o`*, then there is either an* `x`*-path connecting regions X and X' or an* `o`*-path connecting regions O and O'.*

Proof. First we construct a subgraph $G' = (V, E')$ of G, with the same nodes but a subset of the edges. We define an edge to belong in E' only if it lies between a X-face and an O-face. Therefore, e_1, e_2, e_3, and e_4 belong in E'. The nodes u_1, u_2, u_3, and u_4 have degree one.

If all three hexagons around a node are marked the same, then the node is isolated in G' and has degree zero. If a node is surrounded

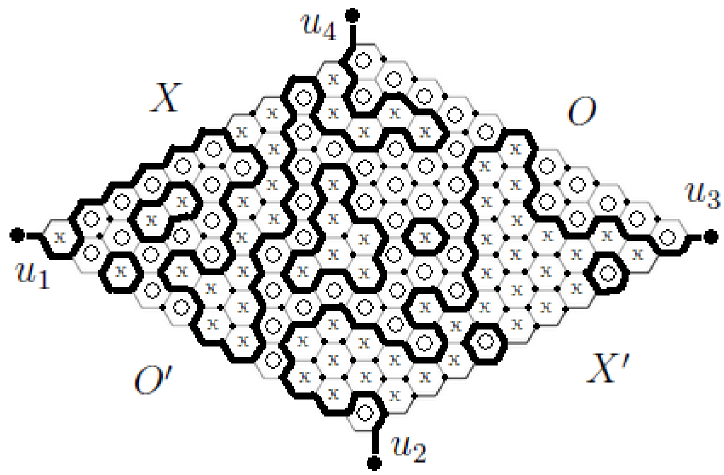

FIGURE 2

by two hexagons of one pattern and one hexagon of the other pattern, then that node has two incident edges. Hence, each node in the core graph has degree either zero or two.

Since G' has nodes with degree at most two, by the lemma, G' is a union of disjoint subgraphs, each of which are isolated nodes, simple cycles, or simple paths. Each of the nodes u_1, u_2, u_3, and u_4 are ends of some path because they have degree one. The disjointness of subgraphs in G' ensures that these paths do not cycle. Therefore, there exist two simple paths in G', each connecting two of u_1, u_2, u_3, and u_4. Although the winner depends on the orientation of the paths, the paths do trace out a winning chain of hexagons. Therefore, for any arbitrary configuration of the Hex board, a winning path for one of the players exists.

For example, in Figure 2, the darkened edges and vertices belong to G'. There are two simple paths, one connecting u_1 and u_4, and another connecting u_2 and u_3. The paths mark out a winning chain for the O-player. □

Using the Hex Theorem, John Nash proved that the first player always has a winning strategy. This is by a simple strategy-stealing argument: By the Hex Theorem, a winning strategy always exists. Suppose for the purposes of contradiction that a winning strategy exists for the second player. Then the first player can always place his first piece randomly on the board, thus effectively becoming the second player. He can proceed by playing the winning second-player strategy. The sum of a random first move and the second player's winning strategy makes a winning strategy for the first player.

The piece that was randomly placed on the board can never hurt the first player, for if it is in the winning strategy, then the he cannot be hurt by having already played it. If that position is not in the winning strategy, then it does not matter. Therefore, the first player in Hex always has a winning strategy.

Note that this proof does not actually give us the winning strategy. We can solve smaller boards by brute force, but the commonly-played 11x11 board size remains unsolved.

5. The Equivalence of Hex and Brouwer's Fix Point Theorems

For simplicity, we once again change the representation of the Hex board for this section.

Let Z^n denote the lattice points of R^n. For $x \neq y \in R^n$, let $|x-y| = \max_i(x_i - y_i)$; $x < y$ if $x_i \leq y_i$ for all i. The points x and y are *comparable* if $x < y$ or $y < x$.

The two-dimensional Hex board of size k, call it B_k, is a graph whose vertices is the set of all $z \in Z^2$ with $(1,1) \leq z \leq (k,k)$. Two vertices z and z' are *adjacent* (i.e., an edge in B_k connects z and z') if $|z - z'| = 1$ and z and z' are comparable.

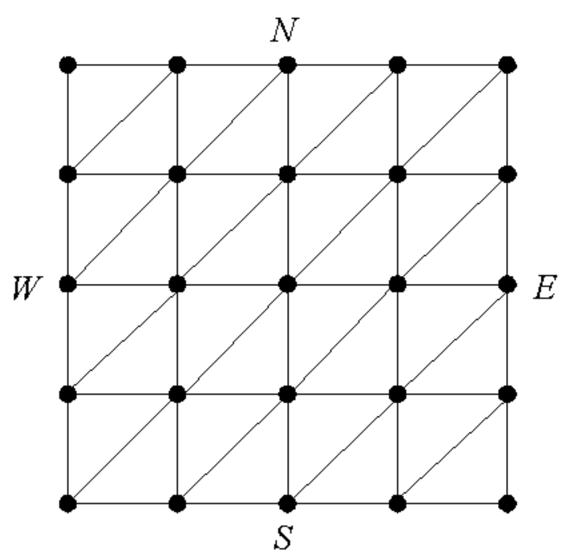

Figure 3

Figure 3 shows a Hex board of size 5. The boundary edges are labeled by the cardinal directions, N, S, E, W. For a board of size k, the vertices on the boundary are all $z = (z_1, z_2)$ that satisfy $z_2 = k$, $z_2 = 0$, $z_1 = k$, $z_1 = 0$, respectively. Rather than the x and o player, now we can think of them as the *horizontal* and *vertical* players.

With the new board representation, we can restate the Hex Theorem:

Theorem 5.1 (Hex Theorem)**.** *Let B_k be covered by two sets H and V. Then either H contains a connected set meetin E and W or V contains a connected set meeting N and S.*

Our goal is to show that the Hex Theorem is equivalent to

Theorem 5.2 (Brouwer Fixed-Point Theorem)**.** *Let f be a continuous mapping from the unit square I^2 into itself. Then there exists $x \in I^2$ such that $f(x) = x$.*

Proof. First we show that the Hex Theorem implies the Brouwer Theorem. Let $f : I^2 \to I^2$ be given by $f(x) = (f_1(x), f_2(x))$. The set I^2 is compact, so it suffices to show that for any $\varepsilon > 0$, there exists $x \in I^2$ for which $|f(x) - x| < \varepsilon$. The compactness of I^2 also implies that f is uniformly continuous, so we know that given $\varepsilon > 0$, there exists $\delta > 0$ such that if $|x - x'| < \delta$, then $|f(x) - f(x')| < \varepsilon$. Without loss of generality, we can pick $\delta < \varepsilon$.

Consider a Hex board B_k, large enough so that $1/k < \delta$. We will also define four subsets H^+, H^-, V^+, V^- of B_k as follows:

$$\begin{aligned} H^+ &= \{z \mid f_1(z/k) - z_1/k > \varepsilon\} \\ H^- &= \{z \mid z_1/k - f_1(z/k) > \varepsilon\} \\ V^+ &= \{z \mid f_2(z/k) - z_2/k > \varepsilon\} \\ V^- &= \{z \mid z_2/k - f_2(z/k) > \varepsilon\} \end{aligned}$$

The goal is to show that these four sets do not cover B_k. The points not covered by these four sets, are the desired fixed points (since if z lies on none of them, then $|f(z/k) - z/k| < \varepsilon$).

The sets H^+ and H^- (V^+ and V^-) are disjoint. The key observation, however, is that they are not *contiguous* (two subsets A and B of a graph are contiguous if there exist $a \in A$ and $b \in B$ where a and b are adjacent; recall the definition for *adjacent* given earlier.)

Suppose that $z \in H^+$ and $z' \in H^-$ are adjacent. Then by definition,

$$f_1(z/k) - z_1/k > \varepsilon$$

and

$$z_1'/k - f_1(z'/k) > \varepsilon.$$

Adding the two gives

$$f_1(z/k) - f_1(z'/k) + z_1'/k - z_1/k > 2\varepsilon. \tag{5.1}$$

By our assumption that z and z' are adjacent, and the choice of k such that $1/k < \delta$, we have $z_1'/k - z_1/k \leq |z_1'/k - z_1/k| = 1/k < \delta < \varepsilon$. Hence

$$z_1/k - z_1'/k > -\varepsilon. \tag{5.2}$$

Adding (5.1) and (5.2) gives

$$f_1(z/k) - f_1(z'/k) > \varepsilon. \tag{5.3}$$

If z and z' were adjacent, then $|z/k - z'/k| = 1/k < \delta$ would imply that $|f(z/k) - f(z'/k)| < \varepsilon$. Hence (5.3) is a contradiction to our choice of δ, and H^+ and H^- cannot be contiguous.

Similarly, V^+ and V^- are not contiguous. Now let $H = H^+ \cup H^-$, and $V = V^+ \cup V^-$. Suppose Q is a connected set lying on H. Then set Q must lie entirely on either H^+ or H^- because they are not contiguous and Q is connected. Now note that H^+ does not meet E, because f maps I^2 onto itself. For all possible $z \in I^2$, $f_1(z/k) \leq 1$, so $f_1(z/k) - 1 \leq 0 < \varepsilon$. Similarly, H^- cannot meet W, so Q cannot meet both E and W. By the same arguments, V cannot contain a connected set that meets both N and S. By the Hex Theorem, H and V do not cover B_k, and the points not covered by H and V are the fixed points.

□

Now we need to prove that the Brouwer Theorem implies the Hex Theorem. We use the fact that any point x in the $k \times k$ square I_k^2 in R^2 can be uniquely expressed as a convex combination of some set of (at most three) vertices in B_k.

Note also that for any mapping f from B_k to R^2, we can extend it to a continuous *simplicial* map $\hat{f}$ on I_k^2. That is, if $x = \lambda_1 z^1 + \lambda_2 z^2 + \lambda_3 z^3$, where $\lambda_i > 0$ and $\sum_i \lambda_i = 1$, then by definition $\hat{f}(x) = \lambda_1 f(z^1) + \lambda_2 f(z^2) + \lambda_3 f(z^3)$.

Before we proceed with the rest of the proof, we prove an algebraic lemma.

Lemma 5.3. *Let z^1, z^2, z^3 be vertices of any triangle in R^2. Let ρ be a mapping given by $\rho(z^i) = z^i + v^i$, where v^1, v^2, v^3 are given vectors* *and $\hat{\rho}$ is its simplicial extension. Then $\hat{\rho}$ has a fixed point if and only if* 0 *lies in the convex hull of* v^1, v^2, v^3.

Proof. Let $x = \lambda_1 z^1 + \lambda_2 z^2 + \lambda_3 z^3$. Then

$$\begin{aligned} \hat{\rho}(x) &= \lambda_1(z^1 + v^1) + \lambda_2(z^2 + v^2) + \lambda_3(z^3 + v^3) \\ &= (\lambda_1 z^1 + \lambda_2 z^2 + \lambda_3 z^3) + (\lambda_1 v^1 + \lambda_2 v^2 + \lambda_3 v^3) \\ &= x + (\lambda_1 v^1 + \lambda_2 v^2 + \lambda_3 v^3). \end{aligned}$$

Thus $\hat{\rho}(x) = x$ if and only if $\lambda_1 v^1 + \lambda_2 v^2 + \lambda_3 v^3 = 0$. □

Assume that B_k is partitioned by two sets H and V. Let an H-path (V-path) be a connected set in H (V). We define four subsets of B_k as follows: let $\hat{W}$ be all vertices connected to W by an H-path and let $\hat{E} = H - \hat{W}$. Let $\hat{S}$ be all vertices connected to S by a V-path, and $\hat{N} = V - \hat{S}$. We have defined $\hat{W}$ and $\hat{E}$ such that they are not contiguous. We assume there is no H-path from E to W and no V-path from N to S and show a contradiction.

Let e^1 and e^2 be the unit vectors of R^2 and let $f : B_k \to B_k$ be given by:

$$f(z) = \begin{cases} z + e^1, & \text{for } z \in \hat{W} \\ z - e^1, & \text{for } z \in \hat{E} \\ z + e^2, & \text{for } z \in \hat{S} \\ z - e^2, & \text{for } z \in \hat{N} \end{cases}$$

For each case, we can check that $f(z)$ is in B_k. The only way $z + e^1$ can lie outside B_k is if some $z \in \hat{W}$ belongs to E. However, we assumed that there is no H-path from W to E, so $\hat{W}$ cannot meet E. Since $\hat{E}$ and $\hat{W}$ are not contiguous, z cannot meet W and $z - e^1$ belongs in B_k for all $z \in \hat{E}$.

Now we turn our attention to $\hat{f}$, which is the simplicial extention of f on I_k^2. Note that $\hat{f}$ is continuous. The non-contiguousness of $\hat{W}$ and $\hat{E}$ implies that the mapping f would translate any triangle with mutually adjacent vertices exclusively by e^1 or $-e^1$, never both. The non-contiguousness of $\hat{S}$ and $\hat{N}$ also implies translation in a single direction for the second coordinate. Then f translates the three vertices by two vectors that lie in a single quadrant of R^2, without 0 in their convex hull. By the lemma, $\hat{f}$ is a continuous function on I_k^2 without a fixed point, which is a contradiction to the Brouwer Theorem. Then the Brouwer Theorem must imply the Hex Theorem. □

John Nash later won the 1994 Nobel Prize for his work in non-cooperative game theory, and the Brouwer Theorem is key in proving one of his most celebrated theorems, the existence of Nash Equilibria in games. How fitting then that Nash's earlier invention, Hex, can be used to prove the Brouwer-Fixed Point Theorem.

References

[1] Martin Gardner, *The Scientific American Book of Mathematical Puzzles and Diversions*, Simon & Schuster, New York, 1959, pp.73-83.

[2] David Gale, The Game of Hex and The Brouwer Fixed-Point Theorem, *American Mathematical Monthly*, 1979, **86**(10), pp. 818-827.

E-mail address: `lijing@mit.edu`