BIDDING GAMES

AARON KLEINMAN

Consider a classic two-person game like Tic-Tac-Toe. Instead of alternating moves, though, you and your friend bid for the right to make the next move. For example, suppose you each start with 5 chips. You bid 2 chips and your friend bids 1, so you pay them 2 and make a move. Now you have 3 chips, your friend has 7 chips, and you repeat the process.

This simple alteration to Tic-Tac-Toe is enough to turn it into a much deeper game (try it!) And many games — chess, hex, go — can be similarly modified. Today, we will explore some of the theory behind such bidding games.

1. Introducton

The theory behind bidding games was first explored by David Richman in the late 1980's. More formally, a *game* is a finite directed graph G with two special vertices labeled a and b, and a token that begins at an initial vertex v of the graph. Two players, Alice and Bob, each start with some nonnegative amount of money. At each step, they both secretly write down some number less than or equal to the total amount of money they have. Then the bids are revealed. The person who bids more pays the other player an amount equal to his or her bid, and moves the token along an adjacent edge to a neighboring vertex. The game ends when the token reaches either a (a win for Alice) or b (a win for Bob). If the bids are ever equal, the tie is broken by a coin toss. The game is ruled a tie if the token never reaches either distinguished vertex. And the players care only about winning the game — the money loses all value once the game is over.

A *winning strategy* is a system of bids and moves that guarantees a win, given the game and the initial state, which includes the board state and the initial amounts of money each player starts with.

Obviously, Alice and Bob cannot both have a winning strategy. What about if *neither* has a wining strategy? The following theorem asserts that this rarely happens:

Theorem 1.1. *For each vertex v of G, there is some number $R(v)$ such that if the token is at v, Alice has a winning strategy if her share of the money is greater than $R(v)$, and Bob has a winning strategy if Alice's share of the money is less than $R(v)$.*

If alice possess exactly $R(v)$, then the outcome may be determined by the tie-breaking coin flips. The surprising part about Theorem 1.1 is that it holds even if G has directed cycles. In such a case, we might expect there to be a range of money Alice can have that will all lead to a tie, but this does not happen.

This also implies that both players have optimal pure strategies, and knowing the other player's strategy does not give a player any advantage when both players play optimally. (Contrast this with games like Rock-Paper-Scissors, where both players move simultaneously).

2. Richman cost functions

Proof. : For each vertex $v \in V(G)$, let $S(V)$ denote the successors of v in G. Given a function $f : V \to [0,1]$, let

$$f^+(v) = \max_{w \in S(V)} f(w), \quad f^-(v) = \min_{w \in S(V)} f(w).$$

A function $R : V \to [0,1]$ is a *Richman cost function* if $f(a) = 0$, $f(b) = 1$, and for all other $v \in V$, $R(v) = \frac{1}{2}(R^+(v) + V^-(v))$.

We will show such a function exists by construction. Define the function $R(v,t)$ as follows. Let $R(b,t) = 0$ and $R(a,t) = 1$ for all $t \in \mathbb{Z}_{\geq 0}$. For $v \notin \{a,b\}$, define $R(v,0) = 1$ and

$$R(v,t) = \frac{1}{2}(R^+(v,t-1) + R^-(v,t-1))$$

for $t > 0$. Then $R(v,t)$ is convergent, and $v \mapsto \lim_{t\to\infty} R(v,t)$ is a Richman cost function. □

Theorem 2.1. *If Bob's share of the total money exceeds $R(v) = \lim_{t\to\infty} R(v,t)$, then Bob has a winning strategy. More specifically, if his share of the money exceeds $R(v,t)$, his victory requires at most t moves.*

As you might expect from the terminology $R(G)$ and $R(v)$, one can also show:

Theorem 2.2. *The Richman cost function of G is unique.*

3. Random turn games

Suppose now that instead of bidding, the right to take the next move is determined by a fair coin toss. Such a game plays very differently from a bidding game. For example, if you are significantly worse than your opponent, a single wrong play in a bidding game will likely lead to a loss. In the corresponding random game, however, you might well still get lucky and pull out a win.

Amazingly, the analysis of a random game is intimately related to the analysis for the corresponding bidding game. Given a finite combinatorial game, recall $R(v)$ is the (unique) Richman cost function on G at vertex v. Let $P(v)$ be the probability that Alice wins in the corresponding random turn game, beginning at vertex v and assuming optimal play. Then:

Theorem 3.1 (Richman's Theorem)**.**

$$R(v) = 1 - P(v).$$

4. Random Hex

This section is adopted from the 2006 paper "Random-Turn Hex and Other Selection Games," by Y. Peres, O. Schramm, S. Sheffield, and D. Wilson.

4.1. **Introduction.**

Hex is a board game, invented in the 1940's by John Nash (among others). Two players alternate turns placing different colored stones on a rhombus-shaped grid, which is tiled by hexagons. The two pairs of opposite sides of the board are colored white and black. Alice wins if she forms a chain linking the whites ides of the board; similarly, Bob wins if he forms a chain linking the black sides of the board.

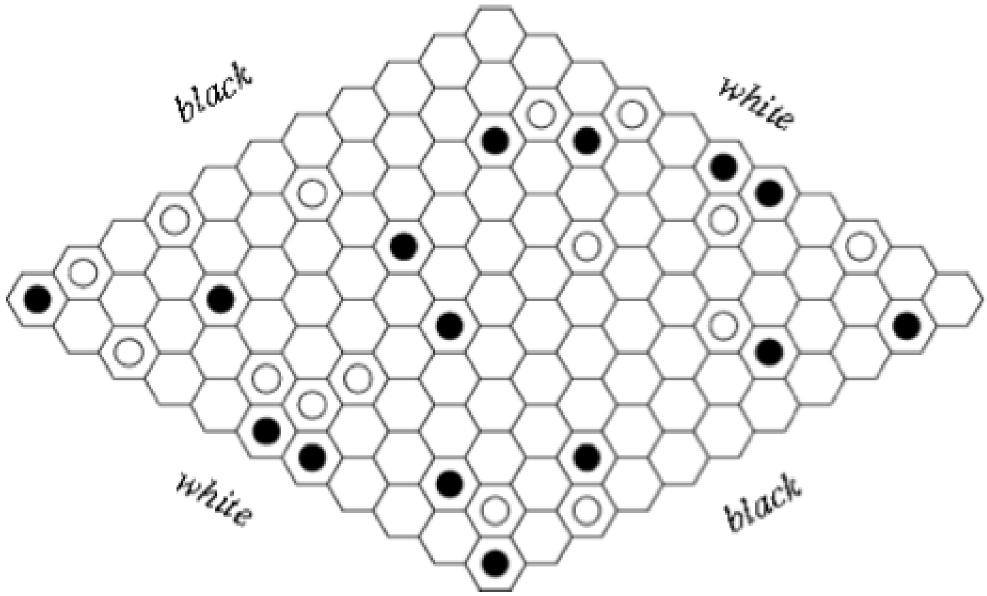

FIGURE 1. A sample game of Hex

Though its rules are simple to state, the game of hex is surprisingly difficult to analyze. Hex is solved for all boards of size 9×9 and below, but the case of the standard 11×11 board remains open.

The random version, however, falls into a class of games known as selection games that are much more amenable to analysis.

4.2. **Selection games.** Let S be an n-element set known as the board, and f a function from the power set of S to $\mathbb{R}$. A *selection game* is played as follows: The first player selects an element of S, the second player selects one of the remaining $n-1$ elements, the first player selects one of the remaining $n-2$ elements, and so on and so forth until S is exhausted. Let S_A, S_B denote the sets chosen by Alice and Bob when they play the selection game, Then Alice receives a payoff of $f(S_A)$ and Bob receives a payoff of $-f(S_A)$.

4.3. **Strategy.** The state of a selection game can be captured by an ordered pair (T_1, T_2) of disjoint subsets of S, corresponding to the current sets of elements selected by A and B. A pure strategy for a player in a random-turn selection game is a map M from such pairs (T_1, T_2) to elements of S.

Let $E(T_1, T_2)$ be the expected payoff for Alice at this stage of the game. E can be computed inductively as follows. If $T_1 \cup T_2 = S$ then $E(T_1, T_2) = f(T_1)$. Now suppose we know $E(T_1, T_2)$ for all T_1 with $|S \backslash (T_1 \cup T_2)| \leq k$. On her turn, Alice's optimal play is to choose an s from $S \backslash (T_i \cup T_2)$ for which $E(T_1 \cup \{s\}, T_2)$ is maximal. Similarly, Bob attempts to minimize $E(T_1, T_2 \cup \{s\})$ whenever he gets to move.

Theorem 4.1. *The expected value of a random-turn selection game is the expectation of $f(T)$ when a set T is selected randomly and uniformly among all subsets of S. Moreover, any optimal strategy for one of the players is also an optimal strategy for the other player.*

Proposition 4.1. *If f is generic, there is a unique optimal strategy and it is the same for both players. Moreover, when both players play optimally, the final S_1 is equally likely to be any one of the 2^n subsets of S.*

4.4. **Win Or Lose.** A game is *win-or-lose* if $f(T)$ takes on precisely two values (which we can assume are -1 and 1). If $S_1 \subset S$ and $s \in S$, then we say s is *critical* for S_1 if

$$f(S_1 \cup \{s\}) \neq f(S_1 \backslash \{s\}).$$

A selection game is *monotone* if $f(S_1) \geq f(S_2)$ for all $S_1 \supset S_2$.

Lemma 4.1. *In a monotone, win-or-lose, random-turn selection game, a first move s is optimal if and only if s is an element of S that is most likely to be critical for a random-uniform subset T of S. When the position is (S_1, S_2), the move s in $S \backslash (S_1, \cup S_2)$ is optimal if and only if s is an element of $S \backslash (S_1 \cup S_2)$ that is most likely to be critical for $S_1 \cup T$, where T is a random-uniform subset of $S \backslash (S_1, S_2)$.*

5. Bidding Hex

This section is adapted from the 2008 preprint "Artificial Intelligence for Bidding Hex", by Sam Payne and Elina Robeva.

The relationship between bidding and random turn games given by Richman's Theorem gives a strategy for how to win at (continuous) bidding games. Here, we give an algorithm used to devastating effectiveness for bidding hex.

Start with a partially filled board. Section 4 shows that it is optimal for both players to play at a hex which is most likely to be critical.

For an empty hex H, let L_H be the probability that H is filled with the losing color when the remainder of the board is filled in at random.

Proposition 5.1. *The probability that H is not critical is $2L_H$.*

The optimal moves are those those that minimize L_H.

Proposition 5.2. *Let H be an open hex such that L_H is minimal from position v. Then an optimal bid for real-valued bidding is the proportion*

$$\delta(v) = \frac{1}{2} - L_H$$

of total bidding resources.

This gives rise to the following AI algorithm:

1. From a given game state, fill in the rest of the board at random a large number of times, recording the outcome each time.
2. Let H be the hex that is least often a losing color. Bid the integer part of $\frac{1}{2} - L_H$ times to the total number of chips in the game.
3. If this wins, move in H.