

# Artificial intelligence for Bidding Hex

SAM PAYNE AND ELINA ROBEVA

We present a Monte Carlo algorithm for efficiently finding near optimal moves and bids in the game of Bidding Hex. The algorithm is based on the recent solution of Random-Turn Hex by Peres, Schramm, Sheffield, and Wilson together with Richman's work connecting random-turn games to bidding games.

## 1. Introduction

Hex is well-known for the simplicity of its rules and the complexity of its play. Nash's strategy-stealing argument shows that a winning strategy for the first player exists, but finding such a strategy is intractable by current methods on large boards. It is not known, for instance, whether the center hex is a winning first move on an odd size board. The development of artificial intelligence for Hex is a notoriously rich and challenging problem, and has been an active area of research for over thirty years [Davis 1975/76; Nishizawa 1976; Anshelevich 2000; 2002a; 2002b; Cazenave 2001; Rasmussen and Maire 2004], yet the best programs play only at the level of an intermediate human [Melis and Hayward 2003]. Complete analysis of Hex is essentially intractable; the problem of determining which player has a winning strategy from a given board position is PSPACE-complete [Reisch 1981], and the problem of determining whether a given empty hex is *dead*, or irrelevant to the outcome of the game, is NP-complete [Björnsson et al. 2007]. Some recent research has focused on explicit solutions for small boards [Hayward et al. 2004; Hayward et al. 2005], but it is unclear whether such techniques will eventually extend to the standard  $11 \times 11$  board.

Bidding Hex is a variation on Hex in which, instead of alternating moves, the players bid for the right to move. Suppose, for example, that Alice plays against Bob, and both start with 100 chips. If Alice bids seventeen for the first move and Bob bids nineteen, then Bob gives nineteen chips to Alice and makes a move. Now Alice has 119 chips and Bob has 81, and they bid for the second move. The total number of chips in the game remains fixed, and the chips have no value

---

Supported by the Clay Mathematics Institute, NSF DMS-1068689, NSF CAREER DMS-1149054, and a Stanford University VPUE grant.

at the end of the game — the goal is simply to win the game of Hex. The main goal of this paper is to present an efficient algorithm for near-optimal play of Bidding Hex. We have implemented this algorithm and it is overwhelmingly effective — undefeated against human opponents. Code is available from the authors on request.

## 2. Bidding games

Bidding games are variations on traditional two-player games in which, instead of alternating moves, the players bid for the right to move (or for the right to determine who moves next, in situations where moving may not be desirable). Richman developed an elegant theory of such games with real-valued bidding in the late 1980s; his theory is presented in [Lazarus et al. 1996; Lazarus et al. 1999]. For recreational play, integer valued bidding is preferred. Bidding Chess has also achieved some popularity among fans of Chess variations [Beasley 2008a; 2008b].

For mathematical purposes, the coarseness of integer bidding with small numbers of chips creates additional subtleties. For instance, the set of optimal first moves in Bidding Tic-Tac-Toe depends on the total number of chips — the center is the unique optimal first move when the total number of chips is greater than 26, but the center is not an optimal first move when the total number of chips is five [Develin and Payne 2010, Theorem 6.17].

One of Richman’s insights was the fundamental connection between bidding games with real-valued bidding and random-turn games in which players flip a coin to determine who moves next. Say  $G$  is a finite, loop-free game played by Alice and Bob. Let  $R(G)$  be the critical threshold, between zero and one, such that Alice has a winning strategy for  $G$  with real-valued bidding if her proportion of the total bidding resources is greater than  $R(G)$ , and she does not have a winning strategy if her proportion is less than  $R(G)$ . Let  $P(G)$  be the probability that Alice wins  $G$  as a random-turn game, assuming optimal play. Intuitively, one expects that if the game favors Alice then  $R(G)$  will be closer to zero and  $P(G)$  will be closer to one, but the precise relation discovered by Richman remains surprising.

**Richman’s Theorem.** *Let  $G$  be a finite, loop-free game. Then*

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

Furthermore, the set of optimal moves for  $G$  with real-valued bidding is the same as the set of optimal moves for random-turn play, and optimal bids may be determined as follows. For any position  $v$  in  $G$ , let  $G_v$  denote the game played

starting from  $v$ . Define

$$R^+(v) = \max\{R(G_w) \mid \text{Bob can move from } v \text{ to } w\},$$

$$R^-(v) = \min\{R(G_{w'}) \mid \text{Alice can move from } v \text{ to } w'\}.$$

The function  $R$  is then determined by the relation

$$R(G_v) = (R^+(v) + R^-(v))/2,$$

together with the conditions that  $R(G_v)$  is equal to zero if  $v$  is a winning position for Alice, and is equal to one if  $v$  is a winning position for Bob. From position  $v$ , an optimal bid for both players is  $\delta(v) = |R^+(v) - R^-(v)|/2$ , which may be thought of, roughly speaking, as the value of moving from position  $v$ .

Since  $R^+(v)$  and  $R^-(v)$  can be interpreted, through Richman's Theorem, as probabilities of winning certain random-turn games, results and insights from random-turn games can be applied directly to bidding games with real-valued bidding. For discrete bidding, such probabilistic methods are also useful, though less directly. For any fixed positive  $\epsilon$ , if Alice's proportion of the total number of chips is at least  $R(G) + \epsilon$  and the total number of chips is sufficiently large with respect to  $G$  and  $\epsilon$ , then Alice has a winning strategy in which all of her moves are optimal for real-valued bidding and hence also for random-turn play [Develin and Payne 2010, Theorem 3.10]. The bounds on the number of chips required are astronomical for most interesting games and grow exponentially with  $1/\epsilon$  and the number of possible turns in the game but, in practice, following real-valued bidding strategy is highly effective even when the number of chips is small.

### 3. Hex background

The game of Hex never ends in a draw. In other words, the only way a player can block the other is by completing a winning chain. We sketch a proof of this fact following [Gale 1979] in the form of an algorithm for determining the winner of a completed game. This algorithm is a key ingredient in the implementation of step (1) in our artificial intelligence program for Bidding Hex, presented in Section 5.

Suppose Alice and Bob place amber and blue hexes, respectively, until the board is full. Color the regions northwest and southeast of the board amber, and the regions northeast and southwest of the board blue, so Alice wins if she makes a chain of amber hexes connecting the two amber regions, and Bob wins if he creates a chain of blue hexes connecting the two blue regions (Figure 1).

Now the boundary between the amber and blue regions of the board is a finite union of loops, plus two paths whose four endpoints are at the north, south, east,



**Figure 1.** Alice, playing amber, has successfully blocked Bob, playing blue.

and west corners of the board. The directed path beginning at the west corner of the board always has amber on its left and blue on its right, so its endpoint must be at either the north or the south corner of the board. If its endpoint is at the north corner, the blue hexes adjacent to this path on the right form a winning chain for Bob; if its endpoint is south, as in Figure 2, the amber hexes adjacent to this path on the left form a winning chain for Alice.

In summary, to determine the winner of a completed game, start at the west corner and follow the boundary between amber and blue regions. If the path leads to the north corner of the board, then Bob wins. Otherwise, the path leads to the south corner of the board, and Alice wins.

If the board is only partially completed, we can determine the state of the game as follows. First, color all of the empty hexes amber and determine the winner. If Bob is the winner, then he already has a winning chain, so the game is over. Otherwise, color all empty hexes blue to check whether Alice already has a winning chain. If not, then neither player has won yet, and the game continues.



**Figure 2.** It follows that Alice has a winning chain joining her two external regions, given by following the component of the boundary between amber and blue that starts at the left corner.

#### 4. Random-Turn Hex

Random-Turn Hex has recently been solved by the probabilists Peres, Schramm, Sheffield, and Wilson; their study was apparently motivated by suggestions of Werner pointing toward a conformal invariance property for Hex played on unusually shaped boards with vanishingly small hexes, based on relations with percolation theory. One of their main results is the Random-Turn Hex Theorem below.

In a completely filled hex board, a hex is *critical* if changing the color of that hex would change the outcome of the game. Equivalently, a hex is critical if every winning chain contains that hex. If the board is only partially filled in, then the *probability* of an open hex being critical is the number of fillings of the remainder of the board for which the given hex is critical, divided by the total number of possible fillings, which is  $2^N$ , where  $N$  is the number of empty hexes.

**Random-Turn Hex Theorem** [Peres et al. 2007]. *The set of optimal moves from any position in Random-Turn Hex is the same for both players, and is exactly the set of open hexes whose probability of being critical is maximal.*

The Random-Turn Hex Theorem gives an exact solution to Random-Turn Hex. From any position, fill in the empty hexes in every possible way and count the number of times that each hex is critical. The hexes that are critical most often are the optimal moves. When the number of empty hexes is large, the number of possible fillings is astronomical, so finding the optimal moves exactly is practically impossible. Nevertheless, the Random-Turn Theorem suggests a Monte Carlo algorithm — if one fills in the board at random a large number of times, then the hexes that are critical most often in this sampling are likely to be optimal or near-optimal. This algorithm has been implemented by David Wilson in the program Hexamania [Wilson 2007]. Wilson’s program beats a skilled human player more than half the time, but a person can still beat the computer fairly often, by randomly winning most of the coin flips. We have implemented a similar algorithm for Bidding Hex, finding not only near-optimal moves but also near-optimal bids, and the resulting program plays exceedingly well.

Cohen [2004a; 2004b] independently experimented with an analogous Monte Carlo method for playing traditional alternating move Hex, in which random hexes are filled in with alternating colors, so the end filling has equal numbers of hexes of both colors if the board size is even, and first player has exactly one extra hex in the final filling if the board size is odd.

#### 5. Bidding Hex

We now describe our artificial intelligence program for playing Bidding Hex. Given a partially filled board, we must efficiently find a near optimal move and

a near optimal bid. Start with a partially filled board. 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  $2L_H$ .*

*Proof.* In any filling, if  $H$  is of the losing color then it is not critical. Furthermore, there is a natural bijection between fillings in which  $H$  is of the losing color and those in which  $H$  is of the winning color but not critical, given by switching the color of  $H$ . Therefore, the number of fillings in which  $H$  is not critical is exactly twice the number of fillings in which  $H$  is of the losing color.  $\square$

The proposition shows that the optimal moves are those hexes  $H$  such that  $L_H$  is as small as possible. An optimal bid can also be determined in terms of  $L_H$ , as follows.

**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 the total bidding resources.*

*Proof.* The probability that  $H$  is critical is  $1 - 2L_H$ , and is equal to  $R^+(v) - R^-(v)$ , the difference between the conditional probability that Alice wins if she gets  $H$  and the probability that she wins if Bob gets  $H$ . Therefore,  $\delta(v) = (1 - 2L_H)/2$ , as required.  $\square$

Here is our algorithm for finding near optimal bids and moves from a partially filled Hex board.

- (1) Fill in the remainder of the board at random a large number of times, and determine the outcome of the game each time.
- (2) Keep track of the number of times that each open hex is of the losing color.
- (3) If  $H$  is the open hex that is least often of the losing color, bid the integer part of  $\frac{1}{2} - L_H$  times the total number of chips in the game.
- (4) If this bid wins, move in hex  $H$ . Otherwise, the opponent has probably overbid. Take his chips, and let him move wherever he wants.

A moderately fast desktop computer, by 2008 standards, can check roughly fifty thousand random fillings of an  $11 \times 11$  board per second. A few hundred thousand fillings per move suffice to consistently defeat the best human players today.

## References

- [Anshelevich 2000] V. V. Anshelevich, “The game of Hex: An automatic theorem proving approach to game programming”, pp. 189–194 in *Proceedings of the Seventeenth National Conference on Artificial Intelligence* (Austin, TX, 2000), AAAI Press, Menlo Park, CA, 2000.
- [Anshelevich 2002a] V. V. Anshelevich, “The game of Hex: The hierarchical approach”, pp. 151–165 in *More games of no chance* (Berkeley, CA, 2000), edited by R. Nowakowski, Math. Sci. Res. Inst. Publ. **42**, Cambridge Univ. Press, Cambridge, 2002.
- [Anshelevich 2002b] V. V. Anshelevich, “A hierarchical approach to computer Hex”, *Artificial Intelligence* **134**:1–2 (2002), 101–120.
- [Beasley 2008a] J. Beasley, “Bidding Chess”, *Variant Chess* **7**:57 (2008), 42–44.
- [Beasley 2008b] J. Beasley, “Isolated pawns”, *Variant Chess* **7**:58 (2008), 71.
- [Björnsson et al. 2007] Y. Björnsson, R. Hayward, M. Johanson, and J. van Rijswijk, “Dead cell analysis in Hex and the Shannon game”, pp. 45–59 in *Graph theory in Paris*, edited by A. Bondy et al., Birkhäuser, Basel, 2007.
- [Cazenave 2001] T. Cazenave, “Retrograde analysis of patterns versus metaprogramming”, pp. 57–75 in *Computational intelligence in games*, edited by N. Baba and L. C. Jain, Studies in Fuzziness and Soft Computing **62**, Physica-Verlag, Heidelberg, 2001.
- [Cohen 2004a] B. Cohen, “Board evaluation in hex”, 2004, <http://www.advogato.org/person/Bram/diary/112.html>.
- [Cohen 2004b] B. Cohen, “Hex”, 2004, <http://www.advogato.org/person/Bram/diary/116.html>.
- [Davis 1975/76] M. Davis, “On artificial machine learning: Some ideas in search of a theory”, *Internat. J. Comput. Math.* **5**:4 (1975/76), 315–329.
- [Develin and Payne 2010] M. Develin and S. Payne, “Discrete bidding games”, *Electron. J. Combin.* **17**:1 (2010), R85, 40.
- [Gale 1979] D. Gale, “The game of Hex and the Brouwer fixed-point theorem”, *Amer. Math. Monthly* **86**:10 (1979), 818–827.
- [Hayward et al. 2004] R. Hayward, Y. Björnsson, M. Johanson, M. Kan, N. Po, and J. van Rijswijk, “Solving  $7 \times 7$  Hex: Virtual connections and game-state reduction”, pp. 261–278 in *Advances in computer games* (Graz, 2003), edited by H. J. van den Herik et al., IFIP Int. Fed. Inf. Process. **135**, Kluwer, Boston, 2004.
- [Hayward et al. 2005] R. Hayward, Y. Björnsson, M. Johanson, M. Kan, N. Po, and J. van Rijswijk, “Solving  $7 \times 7$  Hex with domination, fill-in, and virtual connections”, *Theoret. Comput. Sci.* **349**:2 (2005), 123–139.
- [Lazarus et al. 1996] A. J. Lazarus, D. E. Loeb, J. G. Propp, and D. Ullman, “Richman games”, pp. 439–449 in *Games of no chance* (Berkeley, CA, 1994), edited by R. J. Nowakowski, Math. Sci. Res. Inst. Publ. **29**, Cambridge Univ. Press, Cambridge, 1996.
- [Lazarus et al. 1999] A. J. Lazarus, D. E. Loeb, J. G. Propp, W. R. Stromquist, and D. H. Ullman, “Combinatorial games under auction play”, *Games Econom. Behav.* **27**:2 (1999), 229–264.
- [Melis and Hayward 2003] G. Melis and R. Hayward, “SIX wins Hex tournament”, *ICGA* **26**:4 (2003), 277–280.
- [Nishizawa 1976] T. Nishizawa, “An improved version of the program for Hex”, pp. 40–61 in *Studies on puzzles by games and computers*, Research Institute for Mathematical Sciences, Kyoto, 1976.
- [Peres et al. 2007] Y. Peres, O. Schramm, S. Sheffield, and D. B. Wilson, “Random-turn hex and other selection games”, *Amer. Math. Monthly* **114**:5 (2007), 373–387.

- [Rasmussen and Maire 2004] R. Rasmussen and F. Maire, “An extension of the  $H$ -search algorithm for artificial Hex players”, pp. 646–657 in *AI 2004: Advances in artificial intelligence*, edited by G. I. Webb and X. Yu, Lecture Notes in Comput. Sci. **3339**, Springer, Berlin, 2004.
- [Reisch 1981] S. Reisch, “Hex ist PSPACE-vollständig”, *Acta Inform.* **15**:2 (1981), 167–191.
- [Wilson 2007] D. Wilson, “Hexamania”, 2007, <http://research.microsoft.com/~dbwilson/hex/>.

sam.payne@yale.edu

*Mathematics Department, Yale University,  
10 Hillhouse Ave., New Haven, CT 06511, United States*

robeva@stanford.edu

*Department of Mathematics, Stanford University,  
Stanford, CA 94305, United States*