Who wins Misère Hex?

Jeffrey Lagarias Danny Sleator

Hex is an elegant and fun game that was first popularized by Martin Gardner[1]. Two players alternate placing white and black stones into the hexagons of an N by N rhombus shaped board. A hexagon may contain at most one stone.

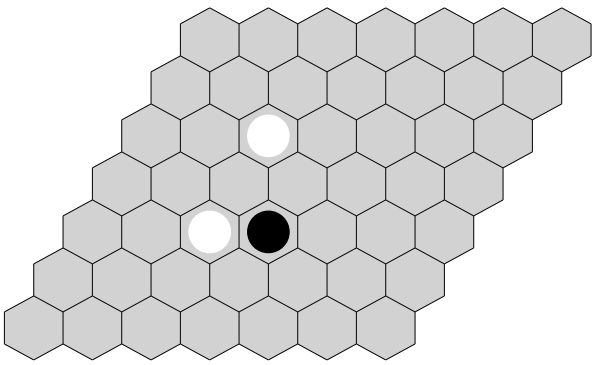

A game of 7 by 7 Hex after three moves.

White's goal is to put white stones in a set of hexagons that connect the top and bottom of the rhombus, and Black's goal is to put black stones in a set of hexagons that connect the left and right sides of the rhombus. Martin Gardner credits John Nash with the observation that there exists a winning strategy for the first player in a game of hex.

The proof goes as follows. First we observe that the game cannot end in a draw, for in any Hex board filled with white and black stones there must be either a winning path for white, or a winning path for black [1, 2]. Since the game is finite, there must be a winning strategy for either the first or the second player. Assume, for the sake of contradiction, that the second player has a winning strategy. The first player can make an arbitrary first move, then follow the winning strategy (reflected) for a second player (imagining that the hexagon containing the first move is empty). If the strategy requires the first player to move in the non-empty cell, the player simply chooses another empty cell in which to play, and now imagines that this one is empty. Since the extra stone can *only help* the first player, the winning strategy will work, and the first player wins. This contradicts our assumption that the second player has a winning strategy.

The purpose of this note is to analyze a variant of Hex that we call *Misère Hex*. The only difference between normal Hex and Misère Hex is that the outcome of the game is reversed: White wins if there is a black chain from left to right, and Black wins if there is a white chain from top to bottom.

Contrary to one's intuition, it is *not* the case that the second player can always win at Misère Hex.

[1]See his Mathematical Games column in Scientific American for July 1957, October 1957, July 1975 and December 1975, or his first book of mathematical puzzles and diversions [3].

Theorem: *The first player has a winning strategy for Misère Hex on an $N \times N$ board when N is even and the second player has a winning strategy when N is odd. Furthermore, the losing player has a strategy which guarantees that every cell of the board must be played before the game ends.*

Proof. It suffices to prove the second assertion, because it shows that the parity of the number of cells on the board determines which player has the winning strategy.

Because the game cannot end in a draw, either the second player or the first player has a winning strategy. Let P be the player who has a winning strategy, and let Q be the other player. For any winning strategy $\mathcal{L}$ for P define $m(\mathcal{L})$ to be the minimum (over all possible games of Misère Hex in which P plays strategy $\mathcal{L}$) of the number of cells left uncovered at the end of the game. We must show that $m(\mathcal{L}) = 0$.

We shall make use of the following *monotonicity property* of the game. Consider a terminal position of a game that is a win for Q. By definition, such a position contains a P-path. Suppose the position is modified by filling in any subset of the empty cells with Q's stones, and further modified by changing any subset of Q's stones into P's stones. The position is still a win for Q, because none of these changes can interfere with the P-path.

We are now ready to prove the theorem. We shall argue by contradiction, supposing that $m(\mathcal{L}) \geq 1$. The contradiction will be to show that under this assumption Q has a winning strategy. The basic idea resembles Nash's proof that the first player has a winning strategy for Hex, in that we will describe a new strategy for Q in which (in effect) Q makes an extra move and then plays the reflected version $\mathcal{L}^R$ of P's hypothetical winning strategy $\mathcal{L}$. (Note that $m(\mathcal{L}) = m(\mathcal{L}^R)$.) However the proof is complicated by the fact that it is not clear *a priori* that having an extra stone on the board is either an advantage or a disadvantage. The proof splits into two cases depending on whether Q is the first player or the second.

Suppose that Q is the first player. Player Q applies the following strategy. She makes an arbitrary first move, and draws a circle around the cell containing this move. From now on she applies strategy $\mathcal{L}^R$ in what we shall call the *imaginary game*. The state of this game is exactly like that of the real game, except that in the imaginary game the encircled cell is empty, while in the real game, that cell contains a Q-stone. This relationship will be maintained throughout the game. When the strategy $\mathcal{L}^R$ requires Q to play in the encircled cell, she plays instead into another empty cell (chosen arbitrarily), erases the circle, and draws a new circle around the move just played. Because $m(\mathcal{L}^R) \geq 1$, when it is P's turn to move there must be at least two empty cells in the imaginary game, and there must be at least one empty cell in the real game. Therefore it is possible for P to move. (We'll see below that P will not have won the real game.) Similarly when it is Q's turn to move there must be at least three empty cells in the imaginary game, so there are at least two empty cells in the real game. Thus the real game can continue.

Eventually Q will win the imaginary game because $\mathcal{L}^R$ is a winning strategy. When this happens she has also won the real game, because of the monotonicity property. This contradicts our assumption that P has a winning strategy.

Now, suppose that Q is the second player. Let p_0 be P's first move. Player Q begins

by encircling p_0, and playing out $\mathcal{L}^R$ in an imaginary game. The imaginary game and the real game differ in up to two places, as follows. The imaginary game is obtained from the real game by first changing p_0 from P's stone to Q's stone, then by erasing the stone in the encircled cell. If the strategy $\mathcal{L}^R$ requires a move into the encircled cell, then Q arbitrarily chooses a different empty cell in which to move, and transfers the circle from its current location to the new cell. The fact that $m(\mathcal{L}^R) \geq 1$ ensures that both players can continue to move. It is easy to see that the relationship between the real game and the imaginary game is maintained.

Player Q eventually wins the imaginary game. The position in the real game is obtained from the position in the imaginary game by putting Q's stone in the encircled cell, and changing the contents of p_0 from a Q-stone to a P-stone. The position in the imaginary game is a winning position for Q, and the monotonicity property ensures that the corresponding position in the real game is also a win for Q. This contradicts our assumption that P has a winning strategy. □

References

[1] A. Beck, M. Bleicher, and J. Crow, *Excursions into Mathematics* Worth, New York, 1969, pp. 327–339.

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

[3] M. Gardner, *The Scientific American Book of Mathematical Puzzles and Diversions*, Simon and Schuster, New York, 1959, pp. 73–83.