

Optimizing Halma: Parallel Minimax and Alpha-Beta Pruning in Haskell

Catherine Lyu (hl3553), Luana Liao (ll3637)

December 2024


1 Introduction

Halma is a strategic board game played on a checkerboard. Unlike the more well-known Chinese Checkers (Figure 1a), which can be played on a hexagonal board with up to six players, Halma is a two-player precursor invented in the 1800s, which uses square tiles. The goal of each player is to move all their pieces from their starting corner to the opponent's corner (Figure 1b).

Players take turns moving one piece per turn across the grid, either to an adjacent empty square or by jumping over an adjacent piece. Multiple consecutive jumps over other pieces are allowed in a single turn if the piece can land in an open space.



(a) A Chinese Checkers board.



(b) Example starting setup for 2-player Halma.

Figure 1: Game layout examples.

1.1 Movement Rules


To elaborate, there are two ways to move in Halma:

- **Single Move:** A piece moves to an adjacent unoccupied square. This ends the player's turn.


- **Jump:** A piece jumps over an adjacent piece (either the player’s own or the opponent’s) to land in a blank square on the other side. Consecutive jumps are allowed in a single turn, but they are optional. Players can stop jumping even if more jumps are available.

Pieces are never ”captured” or removed from the board, even when jumped. Additionally, jumping is always optional and not enforced by the rules.

The rules and examples of movement mechanics were adapted from Pocket Monkey’s Halma Guide [1]. This website also provided visual examples used in Figure 2.



(a) A piece making a single move.



(b) A piece making two jumps and stopping.

Figure 2: Examples of movement mechanics in Halma.

2 Project Overview

The Halma game presents an intriguing computational challenge due to its complex decision space and unique gameplay mechanics. The high branching factor of the game, especially with the possibility of multiple jumps in a single turn, makes it an ideal problem to study algorithms such as Minimax and optimizations such as parallelization. Beyond its gameplay, Halma offers an excellent opportunity to explore computational techniques and evaluate their performance in a real-world scenario.

In this project, our objective was to tackle the problem of efficiently solving Halma using Minimax, a classic algorithm for two-player games that evaluates potential moves by simulating decision making for both players. By combining Minimax with parallelization, we aimed to accelerate the search process, leveraging Haskell’s functional programming paradigm and robust concurrency features. Our approach also incorporated alpha-beta pruning, a technique that reduces the search space by eliminating unproductive branches, further improving the efficiency of the algorithm.

The way an AI for Halma works is by representing the board as a game state and potential moves as transitions between states. The objective of such algorithms is to determine the optimal sequence of moves that would lead to victory, transferring all pieces from a starting camp to the opponent’s camp.

The challenge lies in the exponential growth of the search space. The game’s branching factor is high, especially due to its allowance for multiple jumps in a single turn.

If the algorithm is too slow or resource-intensive, it becomes impractical as a real-time adversary. By integrating parallelization, we aim to distribute computations across multiple cores, enhancing performance.

To address these challenges, we explored optimizations for the Minimax algorithm, a staple in two-player game AI. Specifically, our project goals included:

1. Implement a sequential version of Minimax to serve as a baseline.
2. Parallelization of the algorithm to distribute computational workload across multiple cores, reducing runtime.
3. Enhance the algorithm with alpha-beta pruning to cut off unproductive branches of the search tree.
4. Evaluate the effectiveness of these approaches by measuring runtime improvements and decision quality.

In this work, we designed a Halma-specific game state graph where nodes represent board configurations and edges represent possible moves. Using Haskell, we implemented these algorithms to test and optimize performance. Our solution leverages Haskell’s powerful concurrency features, including the Par monad and Strategies library, to parallelize computations effectively.

We found that each optimization layer—from alpha-beta pruning to parallelized computation—yielded significant improvements in runtime while maintaining the quality of decisions. However, the project also revealed areas for further exploration, such as dynamic workload balancing and hybrid strategies.

In the following sections, we discuss the design and implementation of our algorithms, analyze their performance, and reflect on the challenges and future directions for Halma AI development.

2.1 Minimax Algorithm for Game Agents

The minimax algorithm is a fundamental approach for creating game-playing agents. It simulates all possible moves and counter-moves in a game, constructing a tree of game states. Each node in the tree represents a possible configuration of the game, and the edges represent the moves taken to reach those states. The goal is to identify the optimal move by evaluating the terminal states using a heuristic function.

2.1.1 Game State Representation

In our Halma AI implementation, we use a `GameState` class to store properties like the board configuration, current player, and move history. Here is a snippet from our code:

```
1 data Color = Black | White deriving (Eq, Show)
2
3 type Position = (Int, Int)
4
5 data Piece = Piece {
6     position :: Position, -- (x,y) coord of piece
7     color    :: Color -- which player's piece
```

```

8 } deriving (Eq, Show)
9
10 type Board = [[Maybe Piece]]
11
12 data GameState = GameState {
13     board      :: Board
14     currentPlayer :: Color
15 } deriving (Show)

```

To determine the best move, the Minimax algorithm explores all possible future game states, constructing a decision tree. At each level of the tree, it alternates between maximizing and minimizing the potential outcomes. The "minimax" value of a position represents the best achievable score if both players make the best possible moves.

The algorithm evaluates which moves lead to the best outcomes for each player using some sort of heuristic function.




Figure 3: Minimax Example for TicTacToe

2.2 Heuristic Function

Since Halma has a large state space and may not be solved exhaustively within reasonable time limits, we rely on heuristic evaluations. One such heuristic is the Manhattan distance of pieces to their target zone. Specifically, we used the following evaluating function:

```

1 evaluateBoard :: Board -> Int
  evaluateBoard b = let whiteScore = sum [
                        distance (position p) (0,0) | ... ]
                    blackScore = sum [ distance (
                        position p) (7,7) | ... ]
                    in blackScore - whiteScore

```

This calculate the difference between the sum of black pieces' Manhattan distances from the bottom right corner and the sum of white pieces' Manhattan distances from the top left corner. The higher score means a bigger total distance for the black player or a smaller total distance for a white player, so a higher score favors white.

3 Sequential Solution

Starting board:

	0	1	2	3	4	5	6	7
0	B	B	B	B				
1	B	B	B					
2	B	B						
3	B							
4								W
5							W	W
6						W	W	W
7					W	W	W	W

Figure 4: Starting Board

	0	1	2	3	4	5	6	7
0	B			B				
1	B			B		B		
2	W	B	W					
3				B				W
4			B			B		
5			W	W		B	W	
6							W	
7				W		W	W	

Figure 5: Mid Game Board

Initially, we implemented a basic minimax algorithm with alpha-beta pruning. The algorithm evaluates the board state and prunes branches that do not influence the final decision. This serves as the baseline for performance comparisons. The following describes in high level how the algorithm works:

- Enumerates all possible moves for the current player.
- Recursively applies Minimax to the resulting states.
- Uses alpha and beta bounds to prune unpromising branches.

4 Parallel Solutions

While alpha-beta pruning improved performance significantly, we wanted to further reduce the runtime by exploiting parallelism. Parallelization is natural here because evaluating different moves (child states) at a given level is often independent. Haskell’s parallel constructs make it straightforward to spawn computations in parallel and combine results.

We explored two main parallelization strategies:

4.1 Top-Level Parallelism

In the *top-level parallelism* approach, we evaluate all immediate child states of the root node in parallel. The steps are:

1. At the root state, generate all possible moves for the current player.
2. Use `parMap` and `rpar` to run minimax on each child state concurrently.
3. Gather their evaluations and pick the best move.

This strategy can deliver substantial speedups if the branching factor at the root is large. However, top-level parallelism alone does not help once we move deeper into the tree, as later levels are evaluated sequentially (or less parallelized) once we commit to a branch. For our test configuration, we used a depth of 3 and ran on a machine with 12 or 8 cores.

Top-level parallelism is simple to implement and reason about. It reduces latency for one move decision if the root has many possible moves.

On the flip side, its benefits diminish if the number of top-level moves is small or if alpha-beta pruning quickly eliminates most moves, since this parallelism does not exploit deeper parallelism.

4.2 Chunked Parallelism

The *chunked parallelism* ideas are summarized as follows:

1. Split the top-level moves into chunks (e.g., groups of total number of possible moves / maximum core number each).
2. Evaluate one chunk in parallel, update global alpha-beta bounds, and potentially prune future chunks if the pruning condition is met.
3. Proceed to the next chunk only if necessary, using updated alpha and beta values.

This approach tries to combine parallelism with more effective pruning. By chunking, we do not waste resources evaluating all moves if a strong pruning opportunity arises early. We can stop processing remaining chunks once a prune condition ($\beta \leq \alpha$) is reached, saving time.

A chunk size of number of all possible moves divided by 12 or 8 was used in our experiments depending on the maximum of cores of the machine. On a mid-state board with close to 500 possible moves at the top level, chunked parallelism allowed early pruning after evaluating the first few chunks, potentially speeding up the decision.

Chunked parallelism retains the benefit of parallelism while not committing to evaluating all moves upfront. It potentially reduces wasted computations by applying alpha-beta updates incrementally.

However, it is more complex to implement and requires careful tuning of chunk size. Too large a chunk behaves like the naive top-level parallelism and too small leads to overhead and less parallel efficiency.

4.3 Implementation Details

Both parallel approaches rely on similar code structures to the sequential version with parallel combinators used in the respective Minimax functions to make the computations parallel. The key difference is how the results are combined and how alpha-beta bounds are updated.

In top-level parallelism, we compute all moves at once and pick the best:

```
1 alphaBeta :: [GameState] -> (Int, GameState) -> Int -> Int -> Bool ->
  Int -> (Int, GameState)
2 alphaBeta [] bestEval _ _ _ _ = bestEval
3 alphaBeta gameStates (bestVal, bestState) alpha beta maximizingPlayer
  depth =
4   let results = parMap rpar (\childState -> minimax childState (
      depth - 1) alpha beta (not maximizingPlayer)) gameStates
5       evals = map fst results
6       bestIndex = if maximizingPlayer
7                   then snd $ maximumBy (\(v1,_) (v2,_) -> compare v1
      v2) (zip evals [0..])
8                   else snd $ minimumBy (\(v1,_) (v2,_) -> compare v1
      v2) (zip evals [0..])
9       bestEval = evals !! bestIndex
10      chosenChild = gameStates !! bestIndex
11      finalVal = if maximizingPlayer then max bestVal bestEval else
  min bestVal bestEval
12      finalSt = if (maximizingPlayer && bestEval > bestVal) || (not
      maximizingPlayer && bestEval < bestVal)
13                then chosenChild else bestState
14  in (finalVal, finalSt)
```

In chunked parallelism, we compute in batches, updating alpha and beta after each batch before deciding to continue:

```
1 alphaBetaChunked :: Int -> [GameState] -> (Int, GameState) -> Int ->
  Int -> Bool -> Int -> (Int, GameState)
2 alphaBetaChunked _ [] bestEval _ _ _ _ = bestEval
3 alphaBetaChunked cSize gs (bestVal, bestState) alpha beta
  maximizingPlayer depth =
4   let chunks = chunkList cSize gs
5       in processChunks chunks (bestVal, bestState) alpha beta
  where
6   processChunks [] (curVal, curSt) _ _ = (curVal, curSt)
7   processChunks (chunk:rest) (curVal, curSt) curA curB
8     | curB <= curA = (curVal, curSt)
9     | otherwise =
10      let results = parMap rpar (\childState -> minimax
  childState (depth - 1) curA curB (not maximizingPlayer)) chunk
11          evals = map fst results
12          -- Pick best immediate childState from 'chunk', not
  from deeper states:
13          bestIndex = if maximizingPlayer
14                      then snd $ maximumBy (\(v1,_) (v2,_) ->
  compare v1 v2) (zip evals [0..])
15                      else snd $ minimumBy (\(v1,_) (v2,_) ->
  compare v1 v2) (zip evals [0..])
16          bestEval = evals !! bestIndex
17          chosenChild = chunk !! bestIndex
18  in (bestVal, bestState)
```

```

19         finalVal = if maximizingPlayer then max curVal
bestEval else min curVal bestEval
20         finalSt  = if (maximizingPlayer && bestEval > curVal)
|| (not maximizingPlayer && bestEval < curVal)
21                 then chosenChild else curSt
22         newA = if maximizingPlayer then max curA finalVal else
curA
23         newB = if not maximizingPlayer then min curB finalVal
else curB
24     in processChunks rest (finalVal, finalSt) newA newB

```

5 Performance Evaluation

We evaluated the two parallel solutions: top-level parallel, and chunked parallel across a variety of test states, and compared them to the sequential solution. We specifically focused on a generated mid-state test case. Our metrics include execution time for one move at a fixed depth (e.g., depth = 3) and speedup as we increase the number of cores.

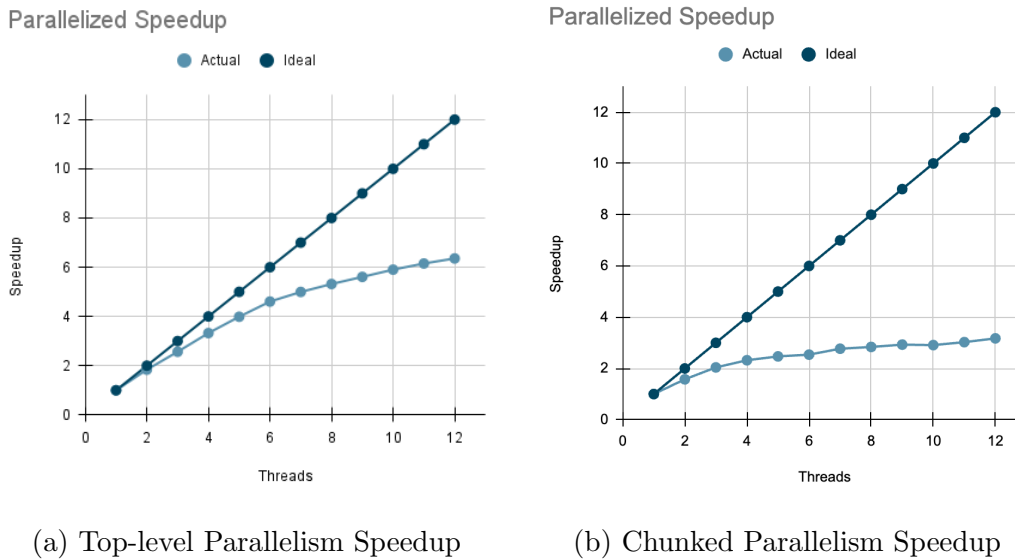


Figure 6: Speedup of (a) Top-level Parallelism and (b) Chunked Parallelism vs. Number of Cores

Top-Level Parallel Minimax												
Threads	1	2	3	4	5	6	7	8	9	10	11	12
Total Run Time(s)	2.473	1.345	0.964	0.744	0.620	0.537	0.495	0.465	0.441	0.419	0.402	0.389

Chunked Parallel Minimax												
Threads	1	2	3	4	5	6	7	8	9	10	11	12
Total Run Time(s)	0.457	0.290	0.224	0.197	0.185	0.180	0.165	0.161	0.156	0.157	0.151	0.144

Figure 7: Tables of Total Run Time of (a) Top-level Parallelism and (b) Chunked Parallelism vs. Number of Cores

As shown in Figure 6, while our parallel solutions do not achieve perfect linear scaling, the speedup continues to improve as we increase the number of cores, even beyond 12. This happens for both top-level and chunked parallelism, where top-level

parallelism achieves a closer to ideal speedup curve. This suggests that, for a complex mid-state Halma board with a chaotic set of moves, the parallelization strategies can continue to leverage additional computational resources. This trend indicates that with further tuning or more substantial computational resources, the performance gains could become even better.




Figure 8: Threadscope of top-level parallelism on 12 cores

Figure 8 depicts a Threadscope visualization of top-level parallelism running on a 12-core machine. We see that all twelve cores remain busy for a significant portion of the runtime. Some cores start a bit earlier, likely due to initial task allocation, and one core eventually becomes responsible for aggregating and returning the results at the end. The relatively solid and continuous workload across all cores suggests that top-level parallelism effectively distributes initial moves at the root, allowing the system to exploit available parallel capacity. This distribution helps shorten the decision time, though some slight idle periods and synchronization points are visible as the computation nears completion.




Figure 9: Threadscope of chunked parallelism on 12 cores

Figure 9 shows the Threadscope output for chunked parallelism using 12 cores. In this scenario, we observe more frequent transitions between active and idle phases on each core. This pattern aligns with the chunked parallel approach, where computations are done in batches. After each batch, alpha-beta global bounds are updated before proceeding to the next chunk. While the workload is still reasonably well spread out and cores remain active most of the time, these incremental synchronization steps lead to more intermittent workload patterns. Despite the pauses, the cores still do a balanced amount of work, and pruning between chunks can significantly cut down on unnecessary computations. Thus, chunked parallelism has shown to substantial speed improvements, especially if pruning occurs effectively.

	Avg. Run Time (s)
Sequential Minimax	14.81
Sequential Minimax + Alpha Beta Pruning	0.79
Top Level Parallelism (12 threads) + Latter Layers Alpha Beta Pruning	0.39
Chunk Parallelism with Global Bounds Updating Alpha Beta Pruning	0.15

Figure 10: Time comparison between the different methods

In Figure 10, we see a direct runtime comparison of the different methods. The sequential Minimax approach without alpha-beta pruning is very slow, and adding alpha-beta pruning drastically reduces execution time. Introducing parallelism at the top level further cuts the runtime, and chunked parallelism with global alpha-beta updating yields even better performance. The table highlights how each layer of optimization, including alpha-beta pruning, parallelism, and chunking, contributes to speed improvements. The final chunked approach runs substantially faster than the original sequential baseline, demonstrating the effectiveness of combining pruning with smartly managed parallel workloads.

6 Discussion and Further Considerations

6.1 Alpha-Beta Pruning

Alpha-beta pruning remains critical in reducing the search space. With parallelizing the top level of the search space, pruning the following layers remain important because it saves time by skipping obviously inferior moves. Therefore, with pruning, parallelism is more efficient since we evaluate fewer states overall.

6.2 Parallelism Trade-offs

While parallelization clearly improves performance, it brings the challenge to chunked parallelism because the algorithm requires iterative synchronization to update alpha-beta bounds.

7 Conclusion

We explored Minimax and alpha-beta pruning for the game Halma, and evaluated both sequential and parallel solutions. Parallel strategies, particularly top-level parallelism and chunked parallelism, showed promising runtime improvements over a pure sequential approach. While chunked parallelism adds complexity, it can yield better pruning effectiveness, reducing wasted work.

Overall, our results suggest that parallelization is a valuable tool for complex board games like Halma, but fitting alpha-beta pruning to the parallelized structures is also very important in reducing the size of the overall search space.

Appendix

Code Listing

Because our codes are lengthy with the game settings, we only included the code from `halma_gameplay_ai_vs_input.hs`, and the rest of the code can be found in our submission `.tar.gz` file.

```
1 import System.IO
2
3 -- Data Types and Constants
4 data Color = Black | White deriving (Eq, Show)
5
6 type Position = (Int, Int)
7
8 data Piece = Piece {
9     position :: Position,
10    color     :: Color
11 } deriving (Eq, Show)
12
13 type Board = [[Maybe Piece]]
14
15 data GameState = GameState {
16     board           :: Board,
17     currentPlayer  :: Color
18 } deriving (Show)
19
20 rows, cols :: Int
21 rows = 8
22 cols = 8
23
24 blackStart, whiteStart :: [Position]
25 blackStart = [(0, 0), (0, 1), (0, 2), (0, 3),
26              (1, 0), (1, 1), (1, 2),
27              (2, 0), (2, 1),
28              (3, 0)]
29
30 whiteStart = [(4, 7),
31              (5, 6), (5, 7),
32              (6, 5), (6, 6), (6, 7),
33              (7, 4), (7, 5), (7, 6), (7, 7)]
34
35 -- Initialize Board
36 initializeBoard :: Board
37 initializeBoard = [ [ initialPieceAt (r, c) | c <- [0..cols-1] ] | r
38                   <- [0..rows-1] ]
39
40 initialPieceAt :: Position -> Maybe Piece
41 initialPieceAt pos
42     | pos `elem` blackStart = Just (Piece pos Black)
```

```

42 | pos 'elem' whiteStart = Just (Piece pos White)
43 | otherwise           = Nothing
44
45 -- Access Board Elements
46 getPiece :: Board -> Position -> Maybe Piece
47 getPiece b (r, c) =
48     if inBounds (r, c) then (b !! r) !! c else Nothing
49
50 inBounds :: Position -> Bool
51 inBounds (r, c) = r >= 0 && r < rows && c >= 0 && c < cols
52
53 -- Move Pieces
54 movePiece :: Board -> Position -> Position -> Board
55 movePiece b from to =
56     let piece = getPiece b from
57         updatedPiece = fmap (\p -> p { position = to }) piece
58         b1 = updateBoard b from Nothing
59         b2 = updateBoard b1 to updatedPiece
60     in b2
61
62 updateBoard :: Board -> Position -> Maybe Piece -> Board
63 updateBoard b (r, c) val =
64     take r b ++
65     [take c (b !! r) ++ [val] ++ drop (c + 1) (b !! r)] ++
66     drop (r + 1) b
67
68 -- Generate Valid Moves
69 directions :: [Position]
70 directions = [(dr, dc) | dr <- [-1,0,1], dc <- [-1,0,1], (dr, dc) /=
71     (0, 0)]
72
73 getValidMoves :: Board -> Piece -> [Position]
74 getValidMoves b p =
75     let singleMoves = getSingleMoves b p
76         jumpMoves = getJumpMoves b (position p) []
77     in singleMoves ++ jumpMoves
78
79 getSingleMoves :: Board -> Piece -> [Position]
80 getSingleMoves b p =
81     [ (r, c)
82     | (dr, dc) <- directions
83     , let (r, c) = addPos (position p) (dr, dc)
84       , inBounds (r, c)
85       , isEmpty b (r, c)
86     ]
87
88 addPos :: Position -> Position -> Position
89 addPos (r1, c1) (r2, c2) = (r1 + r2, c1 + c2)
90
91 isEmpty :: Board -> Position -> Bool
92 isEmpty b pos = getPiece b pos == Nothing
93
94 getJumpMoves :: Board -> Position -> [Position] -> [Position]
95 getJumpMoves b pos visited =
96     concatMap (\dir -> jumpInDirection b pos dir (pos : visited))
97     directions

```

```

96
97 jumpInDirection :: Board -> Position -> Position -> [Position] -> [
    Position]
98 jumpInDirection b (r, c) (dr, dc) visited =
99     let midPos = (r + dr, c + dc)
100         landingPos = (r + 2*dr, c + 2*dc)
101     in if inBounds landingPos &&
102         not (landingPos `elem` visited) &&
103         not (isEmpty b midPos) &&
104         isEmpty b landingPos
105     then
106         let newVisited = landingPos : visited
107             furtherJumps = getJumpMoves b landingPos newVisited
108         in landingPos : furtherJumps
109     else []
110
111 -- Check for Game Over
112 isGameOver :: Board -> Maybe Color
113 isGameOver b
114     | allPiecesInZone b White blackStart = Just White
115     | allPiecesInZone b Black whiteStart = Just Black
116     | otherwise                          = Nothing
117
118 allPiecesInZone :: Board -> Color -> [Position] -> Bool
119 allPiecesInZone b colorPiece zone =
120     let pieces = [ p | row <- b, Just p <- row, color p == colorPiece
121 ]
122     in not (null pieces) && all (\p -> position p `elem` zone) pieces
123
124 -- Evaluate Board
125 evaluateBoard :: Board -> Int
126 evaluateBoard b =
127     let whiteScore = sum [ manhattanDistance (position p) (0, 0) | row
128 <- b, Just p <- row, color p == White ]
129         blackScore = sum [ manhattanDistance (position p) (7, 7) | row
130 <- b, Just p <- row, color p == Black ]
131     in blackScore - whiteScore -- Lower score favors White
132
133 manhattanDistance :: Position -> Position -> Int
134 manhattanDistance (r1, c1) (r2, c2) = abs (r1 - r2) + abs (c1 - c2)
135
136 -- Minimax Algorithm with Alpha-Beta Pruning
137 minimax :: GameState -> Int -> Int -> Int -> Bool -> (Int, GameState)
138 minimax gameState depth alpha beta maximizingPlayer =
139     let b = board gameState
140     in case isGameOver b of
141         Just winner -> if winner == White then (10000, gameState) else
142         (-10000, gameState)
143         Nothing ->
144             if depth == 0
145             then (evaluateBoard b, gameState)
146             else
147                 let moves = getAllMoves gameState
148                     initialEval = if maximizingPlayer then (minBound,
149 gameState) else (maxBound, gameState)

```

```

145         in alphaBeta moves initialEval alpha beta
           maximizingPlayer depth
146
147 alphaBeta :: [GameState] -> (Int, GameState) -> Int -> Int -> Bool ->
           Int -> (Int, GameState)
148 alphaBeta [] bestEval _ _ _ _ = bestEval
149 alphaBeta (gameState:rest) (bestVal, bestState) alpha beta
           maximizingPlayer depth =
150     let (eval, _) = minimax gameState (depth - 1) alpha beta (not
           maximizingPlayer)
151         (newBestVal, newBestState) =
152             if maximizingPlayer
153             then if eval > bestVal then (eval, gameState) else (
           bestVal, bestState)
154             else if eval < bestVal then (eval, gameState) else (
           bestVal, bestState)
155         newAlpha = if maximizingPlayer then max alpha eval else alpha
156         newBeta  = if not maximizingPlayer then min beta eval else
           beta
157     in if newBeta <= newAlpha
158        then (newBestVal, newBestState) -- Prune remaining moves
159        else alphaBeta rest (newBestVal, newBestState) newAlpha newBeta
           maximizingPlayer depth
160
161 getAllMoves :: GameState -> [GameState]
162 getAllMoves gameState =
163     let b = board gameState
164         colorPiece = currentPlayer gameState
165         pieces = [ p | row <- b, Just p <- row, color p == colorPiece
           ]
166         moves = [ (p, dest) | p <- pieces, dest <- getValidMoves b p ]
167         nextPlayer = switchPlayer colorPiece
168         gameStates = [ GameState (movePiece b (position p) dest)
           nextPlayer | (p, dest) <- moves ]
169     in gameStates
170
171 switchPlayer :: Color -> Color
172 switchPlayer Black = White
173 switchPlayer White = Black
174
175 -- Command-Line Interface
176 displayBoard :: Board -> IO ()
177 displayBoard b = do
178     putStrLn " 0 1 2 3 4 5 6 7"
179     mapM_ displayRow (zip [0..] b)
180
181 displayRow :: (Int, [Maybe Piece]) -> IO ()
182 displayRow (i, row) = do
183     putStrLn (show i ++ " ")
184     putStrLn $ concatMap displayCell row
185
186 displayCell :: Maybe Piece -> String
187 displayCell Nothing           = ". "
188 displayCell (Just (Piece _ Black)) = "B "
189 displayCell (Just (Piece _ White)) = "W "
190

```

```

191 -- Main Game Loop
192 main :: IO ()
193 main = do
194     hSetBuffering stdout NoBuffering
195     let initialState = GameState initializeBoard Black
196         gameLoop initialState
197
198 gameLoop :: GameState -> IO ()
199 gameLoop gameState = do
200     displayBoard (board gameState)
201     case isGameOver (board gameState) of
202         Just winner -> putStrLn $ show winner ++ " wins!"
203         Nothing      -> do
204             if currentPlayer gameState == Black
205             then playerTurn gameState
206             else aiTurn gameState
207
208 playerTurn :: GameState -> IO ()
209 playerTurn gameState = do
210     putStrLn "Your turn. Enter move as 'row col newRow newCol':"
211     input <- getLine
212     let inputs = words input
213         if length inputs == 4
214         then case map read inputs :: [Int] of
215             [row, col, newRow, newCol] ->
216                 let piece = getPiece (board gameState) (row, col)
217                     in case piece of
218                         Just p -> do
219                             let validMoves = getValidMoves (board gameState) p
220                                 if (newRow, newCol) `elem` validMoves
221                                 then do
222                                     let newBoard = movePiece (board gameState) (
223                                         row, col) (newRow, newCol)
224                                         newGameState = GameState newBoard White
225                                         gameLoop newGameState
226                                     else do
227                                         putStrLn "Invalid move. Try again."
228                                         playerTurn gameState
229                                 Nothing -> do
230                                     putStrLn "No piece at that position. Try again."
231                                     playerTurn gameState
232                             _ -> do
233                                 putStrLn "Invalid input format. Try again."
234                                 playerTurn gameState
235         else do
236             putStrLn "Invalid input format. Try again."
237             playerTurn gameState
238
239 aiTurn :: GameState -> IO ()
240 aiTurn gameState = do
241     putStrLn "AI is thinking..."
242     let (eval, newGameState) = minimax gameState 2 (minBound :: Int) (
243         maxBound :: Int) True
244         putStrLn $ "AI evaluated the board with a score of: " ++ show eval
245         gameLoop newGameState

```


References

- [1] Pocket Monkey. *Help: Halma*. Available at: <http://www.pocket-monkey.com/help-halma.jsp>
- [2] Nuevo Foundation. *TicTacToe Minimax*. <https://workshops.nuevofoundation.org/java-tictactoe/activity-5/>
- [3] Cameron Chafin. *HalmaGame*. GitHub Repository. Available at: <https://github.com/cameronchafin/HalmaGame>