

MonteCarlo-C4: A Monte Carlo evaluator for Connect-4

Grayson Newell (gln2109)

December 2025

1 Introduction

1.1 Overview

Monte Carlo game playing methods are well suited for parallelization as they rely on many random simulations to create a probability space of results. In this project, we implement a Monte Carlo evaluator for Connect-4 in Haskell and analyze how parallelization affects its performance.

1.2 Hardware

Brand: Apple
Model: M1 Pro
Cores: 8
Threads: 8

2 Sequential Implementation

2.1 Connect-4 Board

Our Connect-4 board representation is simply a nested array of Cells, which have type Maybe Player. Player data can be either Red or Yellow and is implemented as NFData. For working with these data structures, we implemented the following functions:

- **emptyBoard** - Returns an empty board with hardcoded size (6 rows, 7 columns).
- **availableMoves** - Returns an array of available columns given a board.
- **applyMove** - Returns a new board with a move applied given a board, a player, and a column.
- **applyMoveList** - Returns a new board with a list of moves applied given a board, a starting player, and an array of columns.
- **boardString** - Returns a string representation of a board.
- **checkWin** - Returns a winning Player if one exists otherwise Nothing given a board.
- **simulate** - Runs a random simulation and returns the winner if one exists given a board, starting player, and an StdGen.

These are the datatypes and functions on which our Monte Carlo evaluator operates, as described below.




Figure 1: Test Boards

2.2 Monte Carlo Evaluator

The Monte Carlo evaluator returns a predicted best move given a board and a player. Our implementation works as follows:

1. Take all available moves for the given board and create new boards with each move applied.
2. From each new board, run a specified number of simulations where each player makes random moves until the game ends.
3. Count how many simulations resulted in a win for the given player for each available move.
4. Return the move with the most simulated wins.

This is a very simple Monte Carlo implementation, as simulations are played randomly with no heuristics or tree structure. However for a game with a relatively low branching factor like Connect-4, predicting "good" moves is still possible if running enough simulations. We found 4096 simulations to be a sufficient value for strategic play and to emphasize parallel speedup.

2.3 Performance Testing

Figure 1 shows a set of five 7x6 Connect-4 boards that will be used for testing. These were chosen for their balance of progression and win condition.

We tested the performance of our evaluator by measuring the time taken to return an optimal move for all five of these boards. Over five trials, our sequential implementation took an average of 5.787 seconds.

Given our usage of deterministic pseudo-random number generation, the results for each trial were the same: [0, 4, 4, 3, 1]. These numbers represent the most optimal predicted column for each test board and player. Here, its apparent that the evaluator is capable of making educated moves, as it both obtains and blocks a win in the first and third samples respectively.

3 Parallelization

3.1 Initial Approach

Our parallel approach utilized `parMap` and `rdeepseq` from `Control.Parallel.Strategies`. Using these functions, we experimented with evaluating both available moves and game simulations in parallel. The results are shown below for the aforementioned sample boards, running 4096 simulations at 8 threads, averaged over five trials.

Parallelization	Average Runtime (s)
Available Moves	1.182
Game Simulations	0.994
Both	0.984

This data clearly demonstrates that parallelization of game simulations is more effective than doing so for available moves in terms of performance. So, we chose to parallelize game simulations in our final implementation. While the runtimes for using nested parallelization may have been marginally better, this implementation raises scalability concerns and does not offer a justifiable performance advantage.

Figure 2 shows a speedup graph for parallelization of simulations (measured using `getCurrentTime`).

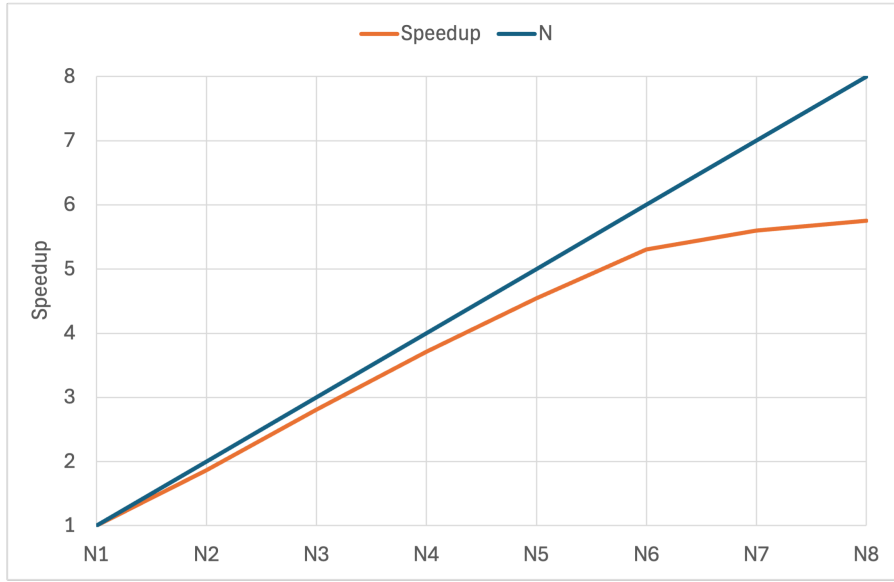


Figure 2: Basic Parallel Speedup

3.2 Simulation Chunking

In efforts to optimise our parallel evaluator, we created an implementation using manual chunking of game simulations. The simulations are split into a specified chunk size and `parMap` is called across chunks instead of all simulations. We performed testing for different chunk sizes as shown below:

Chunk Size	Average Runtime (s)
512	1.199
256	1.185
128	1.133
64	1.114
32	1.143
16	1.092
8	1.135
4	1.088
2	0.945

The apparent downward trend in runtime as chunk size decreases suggests that large chunk size is not an effective method of optimization. However, using a chunk size of 2 does appear to perform slightly better than the pure parMap implementation.

4 Thread Analysis

4.1 Simple parMap

```

1 Parallel GC work balance: 14.70% (serial 0%, perfect 100%)
2
3 TASKS: 18 (1 bound, 17 peak workers (17 total), using -N8)
4
5 SPARKS: 143360 (143360 converted, 0 overflowed, 0 dud, 0 GC'd, 0 fizzled)
6
7 INIT      time      0.005s ( 0.007s elapsed)
8 MUT      time      6.997s ( 0.936s elapsed)
9 GC       time      0.044s ( 0.033s elapsed)
10 EXIT     time      0.002s ( 0.010s elapsed)
11 Total    time      7.049s ( 0.987s elapsed)

```

As shown above, the basic parMap implementation creates a very large amount of sparks. However, all sparks are converted and speedup is roughly 7.14 (based on elapsed versus total CPU time), which is almost optimal for 8 cores. Additionally, the time spent on Garbage Collection appears to be insignificant, so we did not consider it to be a key point for optimization.




Figure 3: Pure parMap implementation

4.2 Chunk size 64

```

1 Parallel GC work balance: 38.61% (serial 0%, perfect 100%)
2
3 TASKS: 18 (1 bound, 17 peak workers (17 total), using -N8)

```

```

4 SPARKS: 2240 (2108 converted, 0 overflowed, 0 dud, 42 GC'd, 90 fizzled)
5
6
7 INIT    time    0.005s ( 0.004s elapsed)
8 MUT     time    6.800s ( 1.131s elapsed)
9 GC      time    0.013s ( 0.012s elapsed)
10 EXIT    time    0.001s ( 0.001s elapsed)
11 Total   time    6.819s ( 1.148s elapsed)

```

Our chunked parallel implementation creates far fewer sparks at a chunk size of 64 and spends less time on garbage collection. Though, its performance at this size is notably worse, which is made apparent by the poor load balancing shown in threadscope.

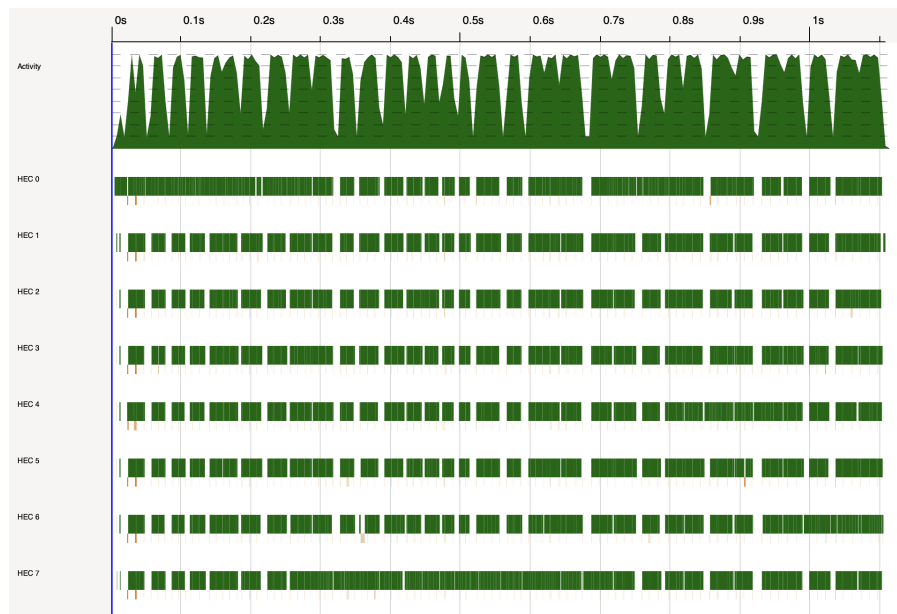


Figure 4: Chunk size 64

4.3 Chunk Size 2

```

1 Parallel GC work balance: 20.78% (serial 0%, perfect 100%)
2
3 TASKS: 18 (1 bound, 17 peak workers (17 total), using -N8)
4
5 SPARKS: 71680 (71667 converted, 0 overflowed, 0 dud, 0 GC'd, 13 fizzled)
6
7 INIT    time    0.006s ( 0.005s elapsed)
8 MUT     time    6.997s ( 0.925s elapsed)
9 GC      time    0.026s ( 0.018s elapsed)
10 EXIT    time    0.002s ( 0.010s elapsed)
11 Total   time    7.031s ( 0.958s elapsed)

```

Chunking with size 2 appears to provide the best load balancing, as shown in Figure 4. This implementation has roughly the same MUT time as the pure parMap one. However it creates half as many sparks and spends roughly half as much time on garbage collection, giving it a slight runtime advantage.




Figure 5: Chunk size 2

5 Final Implementation

5.1 Performance

Given that our chunked parallel implementation with size 2 demonstrated the best initial performance, we chose it as our final implementation.

Further optimization was offered by adding the `-A32m` flag, bringing the garbage collection time from 18ms to 5ms. This improvement was discovered in a project report from 2024: "A* Search For TSP."

Figure 6 shows a speedup graph for this implementation:

N	Runtime (s)	Speedup
1	5.71	1
2	3.04	1.88
3	2.01	2.84
4	1.51	3.77
5	1.22	4.66
6	1.03	5.56
7	0.98	5.83
8	0.92	6.17

5.2 Further Optimization

While this implementation did offer significant speedup at 8 cores, there are a few areas that could be further improved upon. For one, tests were performed on an M1 Pro chip, which has an asymmetric core setup of 6 "Performance" cores and 2 "Efficiency" cores. This could explain the increased gap between speedup and N apparent at N7 and N8. Additionally, our implementation uses a simple array for the board and Int's for the column indices. Using more efficient data structures as well as better methods of applying moves and win checking

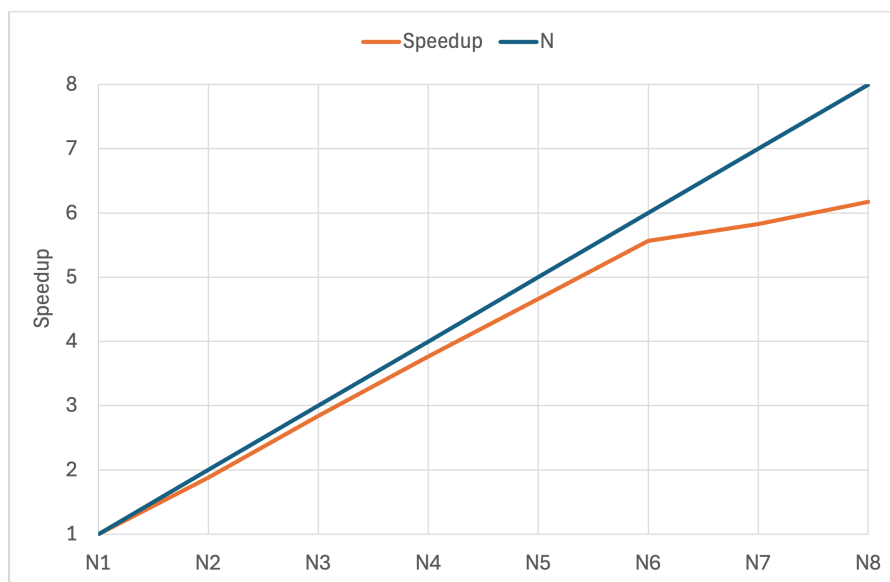


Figure 6: Chunk Size 2 Parallel Speedup

would likely make our implementation much faster.

5.3 Playing Against Evaluator

We created an interactive Connect-4 game in the command line to allow for play against the evaluator. It uses the size 2 chunked implementation with 2048 simulations. We found its moves to be very quick and difficult to beat in most circumstances. Figure 7 shows some screenshots (the evaluator is Yellow).

6 Code

<https://github.com/gln2109/montecarlo-c4>

6.1 Board.hs

```

1 {-# LANGUAGE DeriveGeneric #-}
2 module Board where
3
4 import Data.Maybe
5 import GHC.Generics (Generic)
6 import Control.DeepSeq
7 import System.Random
8
9 data Player = Red | Yellow deriving (Eq, Show, Generic)
10 instance NFData Player
11
12 type Board = [[Maybe Player]]
13
14 -- set board size
15 boardRows, boardCols :: Int
16 boardRows = 6
17 boardCols = 7
18
19 -- helper function for turn swap

```

Yellow's turn: 1 2 3 4 5 6 7 R Y R . . . Y . R Y . . . R R Y R . . . Y R R Y . Y Yellow move: 2 Red's turn: 1 2 3 4 5 6 7 R Y . Y R . . . Y . R Y . . . R R Y R . . . Y R R Y . Y Enter move (1-7): 2 Yellow's turn: 1 2 3 4 5 6 7 R . R Y . Y R . . . Y . R Y . . . R R Y R . . . Y R R Y . Y Yellow move: 3 Red's turn: 1 2 3 4 5 6 7 R . R Y . Y R . . . Y Y R Y . . . R R Y R . . . Y R R Y . Y Yellow wins!	Yellow's turn: 1 2 3 4 5 6 7 . Y . . . R R R Y . . . Yellow move: 4 Red's turn: 1 2 3 4 5 6 7 Y Y . . . R R R Y . . . Enter move (1-7): 5 Yellow's turn: 1 2 3 4 5 6 7 Y Y . . . R R R Y R . . Yellow move: 4 Red's turn: 1 2 3 4 5 6 7 Y Y . . . R R R Y R . . Yellow wins!	Yellow's turn: 1 2 3 4 5 6 7 . R R R Y Y Y R R Y . . . Yellow move: 3 Red's turn: 1 2 3 4 5 6 7 Y R R R Y Y Y R R Y . . . Enter move (1-7): 5 Yellow's turn: 1 2 3 4 5 6 7 Y R R R Y Y Y R R Y R . . Yellow move: 5 Red's turn: 1 2 3 4 5 6 7 Y R R R Y Y Y Y . . . R R Y R . . Yellow wins!
--	---	---

Figure 7: Connect-4 Gameplay

```

20 otherPlayer :: Player -> Player
21 otherPlayer Red = Yellow
22 otherPlayer Yellow = Red
23
24 -- generate empty board
25 emptyBoard :: Board
26 emptyBoard = replicate boardRows (replicate boardCols Nothing)
27
28 -- return array of open cols
29 availableMoves :: Board -> [Int]
30 availableMoves board = [col | col <- [0..boardCols-1], isNothing (board !! 0
    !! col)]
31
32 -- return new Board with move applied
33 applyMove :: Board -> Player -> Int -> Board
34 applyMove board player col =
35     case [row | row <- [0..boardRows-1], isNothing (board !! row !! col)] of
36         [] -> board
37         emptyRows -> let row = maximum emptyRows
38                       newRow = take col (board !! row) ++ [Just player] ++
39                               drop (col+1) (board !! row)
40                       in take row board ++ [newRow] ++ drop (row+1) board
41
42 -- recursively apply an array of moves, alternating players
43 applyMoveList :: Board -> Player -> [Int] -> Board
44 applyMoveList board _ [] = board

```



```

44 applyMoveList board player (col:cols) = applyMoveList (applyMove board player
45   col) (otherPlayer player) cols
46
47 -- return a string representation of the board
48 boardString :: Board -> String
49 boardString board =
50   unlines [
51     concat [playerChar (board !! row !! col) ++ " " | col <- [0..
52       boardCols-1]]
53     | row <- [0..boardRows-1]
54   ]
55   where
56     playerChar Nothing = "."
57     playerChar (Just Red) = "R"
58     playerChar (Just Yellow) = "Y"
59
60 -- return a winning player if one exists
61 checkWin :: Board -> Maybe Player
62 checkWin board = checkCells [(row,col) | row <- [0..boardRows-1], col <- [0..
63   boardCols-1]]
64   where
65     checkFour row col drow dcol =
66       let lastRow = row + 3*drow
67         lastCol = col + 3*dcol
68       in if lastRow < 0 || lastRow > boardRows-1 || lastCol < 0 ||
69         lastCol > boardCols-1 then Nothing
70       else case board !! row !! col of
71         Nothing -> Nothing
72         Just player ->
73           let otherCells = [(board !! (row + d*drow) !! (col + d*
74             dcol)) | d <- [1..3]]
75           in if all (== Just player) otherCells then Just player
76             else Nothing
77
78     checkDirs :: (Int,Int) -> [(Int,Int)] -> Maybe Player
79     checkDirs _ [] = Nothing
80     checkDirs cell (dir:dirs) =
81       case checkFour (fst cell) (snd cell) (fst dir) (snd dir) of
82         Nothing -> checkDirs cell dirs
83         Just player -> Just player
84
85     checkCells :: [(Int, Int)] -> Maybe Player
86     checkCells [] = Nothing
87     checkCells (cell:cells) =
88       case checkDirs cell [(1,0),(0,1),(1,1),(1,-1)] of
89         Nothing -> checkCells cells
90         Just player -> Just player
91
92 -- run a simulation and return the winner
93 simulate :: Board -> Player -> StdGen -> Maybe Player
94 simulate board player gen =
95   case checkWin board of
96     Just winner -> Just winner
97     Nothing ->
98       let moves = availableMoves board
99       in if moves == [] then Nothing
100      else
101        let (moveIndex, newGen) = randomR (0, length moves - 1) gen
102          newBoard = applyMove board player (moves !! moveIndex)
103        in simulate newBoard (otherPlayer player) newGen

```

6.2 EvalSeq.hs

```

1 module EvalSeq (bestMoveSeq) where

```

```

2
3 import Board
4 import System.Random
5 import Data.Function (on)
6 import Data.List
7
8 -- count simulation wins for each move
9 evalMove :: Board -> Player -> Int -> Int -> (Int, Int)
10 evalMove board player simulations move =
11     let gens = [mkStdGen s | s <- [1..simulations]]
12         newBoard = applyMove board player move
13         results = map (simulate newBoard (otherPlayer player)) gens
14     in (move, length (filter (== Just player) results))
15
16 -- return the move with the most wins
17 bestMoveSeq :: Board -> Player -> Int -> Int
18 bestMoveSeq board player simulations =
19     let moves = availableMoves board
20         winCounts = map (evalMove board player simulations) moves
21     in (fst (maximumBy (on compare snd) winCounts))

```

6.3 EvalPar.hs

```

1 module EvalPar (bestMovePar) where
2
3 import Board
4 import System.Random
5 import Data.Function (on)
6 import Data.List
7 import Control.Parallel.Strategies
8
9 -- count simulation wins for each move
10 evalMove :: Board -> Player -> Int -> Int -> (Int, Int)
11 evalMove board player simulations move =
12     let gens = [mkStdGen s | s <- [1..simulations]]
13         newBoard = applyMove board player move
14         results = parMap rdeepseq (simulate newBoard (otherPlayer player))
15             gens
16     in (move, length (filter (== Just player) results))
17
18 -- return the move with the most wins
19 bestMovePar :: Board -> Player -> Int -> Int
20 bestMovePar board player simulations =
21     let moves = availableMoves board
22         winCounts = map (evalMove board player simulations) moves
23     in (fst (maximumBy (on compare snd) winCounts))

```

6.4 EvalChunk.hs

```

1 module EvalChunk (bestMoveChunk) where
2
3 import Board
4 import System.Random
5 import Data.Function (on)
6 import Data.List
7 import Control.Parallel.Strategies
8
9 -- run a chunk of simulations for a given move
10 simulateChunk :: Board -> Player -> Int -> [StdGen] -> Int
11 simulateChunk board player move gens =
12     let newBoard = applyMove board player move
13     in results = map (simulate newBoard (otherPlayer player)) gens

```

```

14     in length (filter (== Just player) results)
15
16 -- split list into length n chunks
17 splitList :: [a] -> Int -> [[a]]
18 splitList [] _ = []
19 splitList li n = take n li : splitList (drop n li) n
20
21 -- count simulation wins for each move
22 evalMove :: Board -> Player -> Int -> Int -> Int -> (Int, Int)
23 evalMove board player simulations chunkSize move =
24     let gens = [mkStdGen s | s <- [1..simulations]]
25         genChunks = splitList gens chunkSize
26         chunkWins = parMap rdeepseq (simulateChunk board player move)
27             genChunks
28         totalWins = sum chunkWins
29     in (move, totalWins)
30
31 -- return the move with the most wins
32 bestMoveChunk :: Board -> Player -> Int -> Int -> Int
33 bestMoveChunk board player simulations chunkSize =
34     let moves = availableMoves board
35         winCounts = map (evalMove board player simulations chunkSize) moves
36     in (fst (maximumBy (on compare snd) winCounts))

```

6.5 Bench.hs

```

1 module Main where
2
3 import Board
4 import EvalSeq
5 import EvalPar
6 import EvalChunk
7 import Control.Monad (forM)
8 import Data.Time.Clock
9
10 testBoards :: [(Board, Player)]
11 testBoards = [(applyMoveList emptyBoard Red [3,3,2,4,1,4], Red),
12               (applyMoveList emptyBoard Red [3,0,2,4,1,4,2], Yellow),
13               (applyMoveList emptyBoard Red [3,0,2,4,1,4,3,4], Red),
14               (applyMoveList emptyBoard Red [3,0,2,4,1,4,3,4,4,0,1], Yellow),
15               (applyMoveList emptyBoard Red [3,0,2,4,1,4,3,4,4,0,1,2], Red)]
16
17 main :: IO ()
18 main = do
19     putStrLn "Test␣Boards:\n"
20     _ <- forM testBoards \(board, player) -> do putStrLn (boardString board
21     ++ "Player:␣" ++ (show player) ++ "\n")
22
23     let sims = 4096
24
25     putStrLn "Sequential"
26     start1 <- getCurrentTime
27     moves1 <- forM testBoards \(board, player) -> return (bestMoveSeq board
28     player sims)
29     putStrLn ("Moves:␣" ++ show moves1)
30     end1 <- getCurrentTime
31     putStrLn ("Time:␣" ++ show (diffUTCTime end1 start1))
32
33     putStrLn "\nBasic␣Parallel"
34     start2 <- getCurrentTime
35     moves2 <- forM testBoards \(board, player) -> return (bestMovePar board
36     player sims)
37     putStrLn ("Moves:␣" ++ show moves2)
38     end2 <- getCurrentTime
39     putStrLn ("Time:␣" ++ show (diffUTCTime end2 start2))

```

```

36
37     putStrLn "\nChunkedParallelSize256"
38     start3 <- getCurrentTime
39     moves3 <- forM testBoards \(board, player) -> return (bestMoveChunk
40         board player sims 256))
41     putStrLn ("Moves:␣" ++ show moves3)
42     end3 <- getCurrentTime
43     putStrLn ("Time:␣" ++ show (diffUTCTime end3 start3))
44
45     putStrLn "\nChunkedParallelSize2"
46     start4 <- getCurrentTime
47     moves4 <- forM testBoards \(board, player) -> return (bestMoveChunk
48         board player sims 2))
49     putStrLn ("Moves:␣" ++ show moves4)
50     end4 <- getCurrentTime
51     putStrLn ("Time:␣" ++ show (diffUTCTime end4 start4))

```

6.6 Play.hs

```

1  module Main where
2
3  import Board
4  import EvalChunk
5  import Text.Read (readMaybe)
6
7  getMove :: Board -> IO Int
8  getMove board = do
9      putStrLn "Enter␣move␣(1-7):"
10     move <- getLine
11     case readMaybe move :: Maybe Int of
12         Just num -> do
13             if notElem (num-1) (availableMoves board) then do
14                 putStrLn "Invalid␣move."
15                 getMove board
16             else return num
17         Nothing -> do
18             putStrLn "Invalid␣move."
19             getMove board
20
21  playGame :: Board -> Player -> IO ()
22  playGame board player = do
23      putStrLn ("\n" ++ show player ++ "'s␣turn:")
24      putStrLn "1␣2␣3␣4␣5␣6␣7"
25      putStrLn (boardString board)
26      case checkWin board of
27          Just winner -> putStrLn (show winner ++ "␣wins!")
28          Nothing ->
29              if (availableMoves board) == [] then putStrLn "Draw."
30              else do
31                  if player == Red then do
32                      playerMove <- getMove board
33                      let playerBoard = applyMove board Red (playerMove-1)
34                      playGame playerBoard Yellow
35                  else do
36                      let evalMove = bestMoveChunk board Yellow 2048 2
37                      evalBoard = applyMove board Yellow evalMove
38                      putStrLn ("Yellow␣move:␣" ++ (show (evalMove+1)))
39                      playGame evalBoard Red
40
41  main :: IO ()
42  main = do
43      playGame emptyBoard Red

```