CS540 Introduction to Artificial Intelligence
Lecture 19

Young Wu
Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

July 14, 2020

Pirate Game Example

Quiz

Q1

- 5 pirates got 100 gold coins. Each pirate takes a turn to propose how to divide the coins, and all pirates who are still alive will vote whether to accept the proposal or reject the proposal, kill the pirate, and continue to the next round. Use strict majority rule for the vote, and use the assumption that if a pirate is indifferent, he or she will vote reject with probability 50 percent.
- How will the first pirate propose?
- A: $(0, 0, 0, 0, 100)$
- B: $(20, 20, 20, 20, 20)$
- C: $(94, 0, 1, 2, 3)$
- D: $(97, 0, 1, 0, 2)$
- E: $(98, 0, 1, 0, 1)$

Pirate Game Example Diagram

Quiz

round \ player		1	2	3	4	5
1						100
2					0	100
3				99	1	0
4			97	0	2	1
5		97	0	1	0	2
~~6~~	~~96~~	~~0~~	~~1~~	~~2~~	~~1~~	~~0~~

need 4

Nim Game Example

Motivation

- Ten objects. Pick 1 or 2 each time. Pick the last one to win.
- A: Pick 1.
- B: Pick 2.
- C, D, E: Don't choose.

Nim Game Example Diagram

Motivation

suppose 5 objs.

win 1

lose -1

zero-sum

~~cost~~ payoff for me

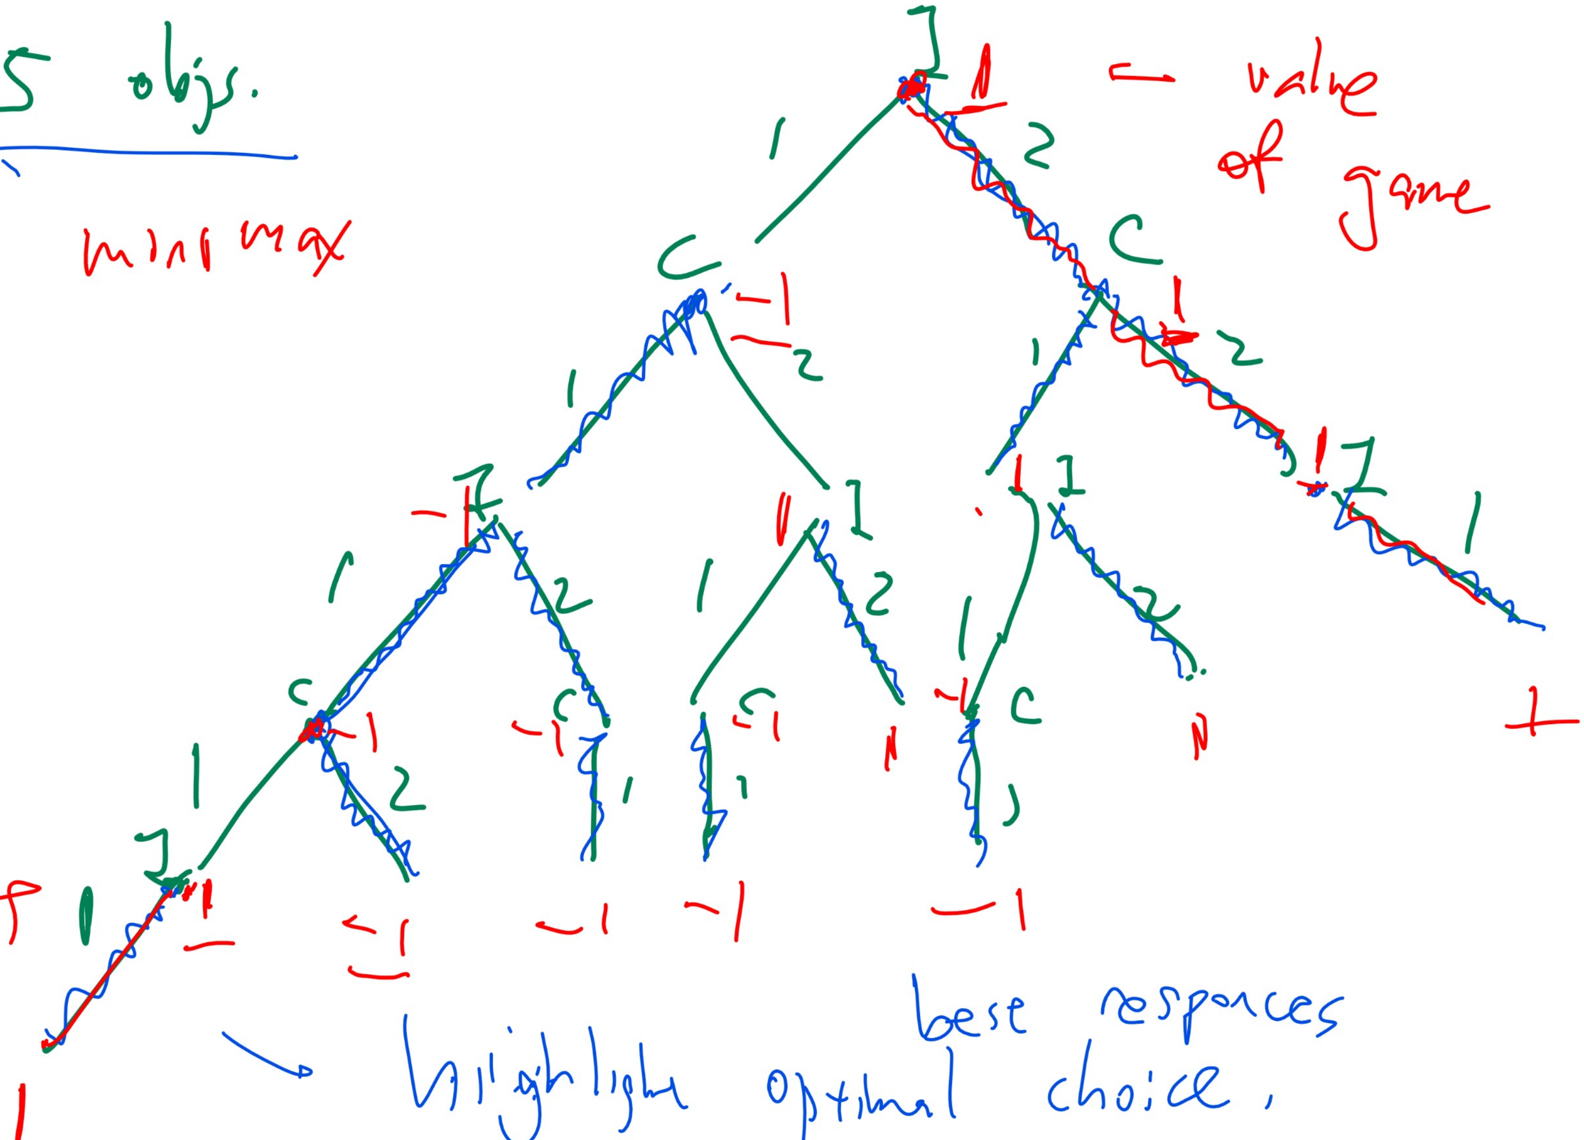

Minimax Algorithm

Description

- Use DFS on the game tree.

Minimax Example 1

Quiz

two player

zero - sum

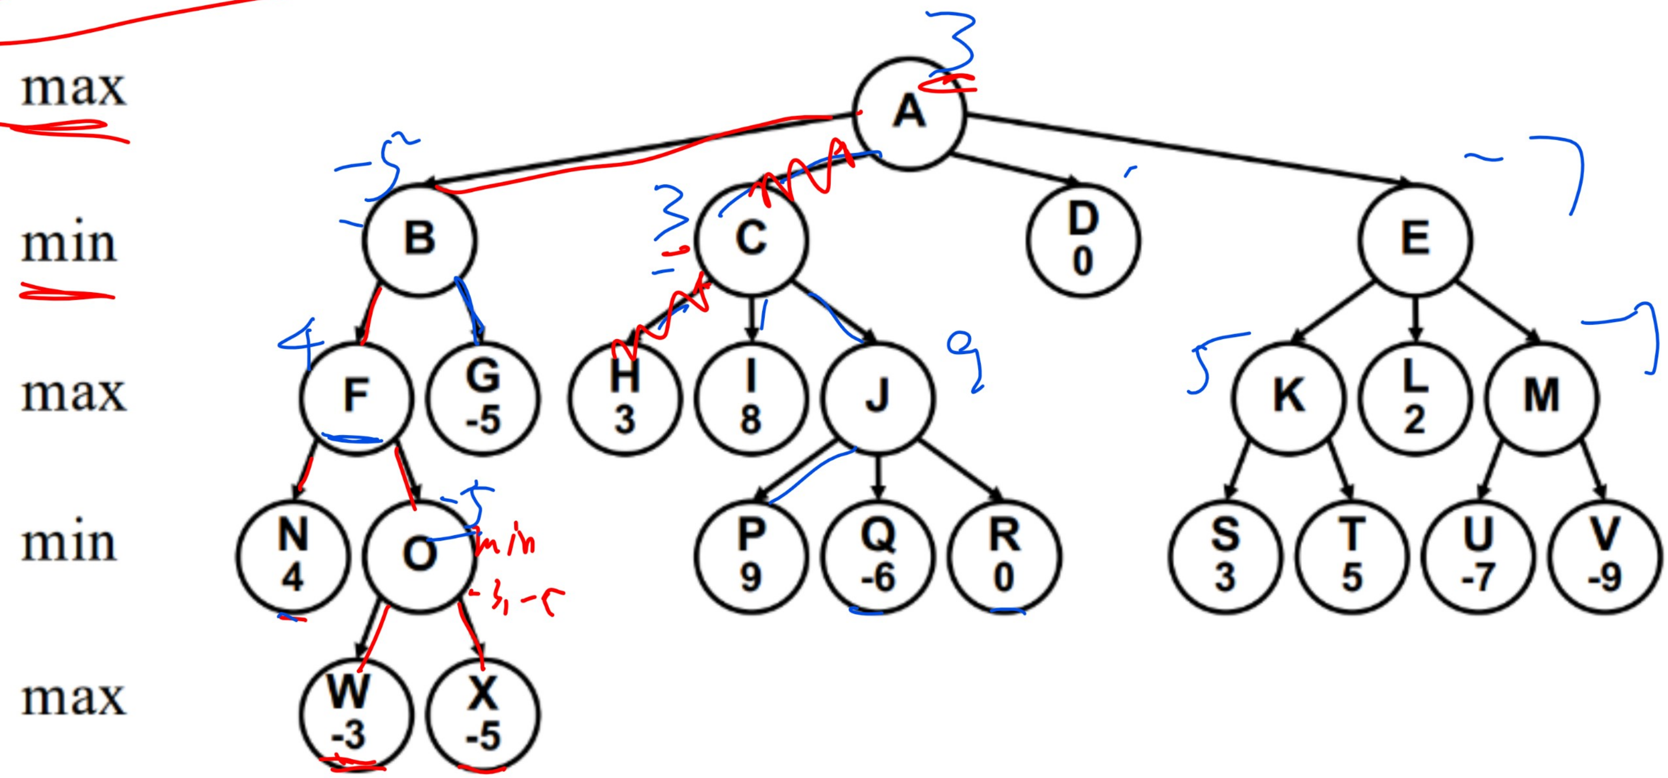

Minimax Example 2

Quiz

Q2

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. What is the value of the game?
- A: 1
- B: 2
- C: 4
- D: 8
- E: 16

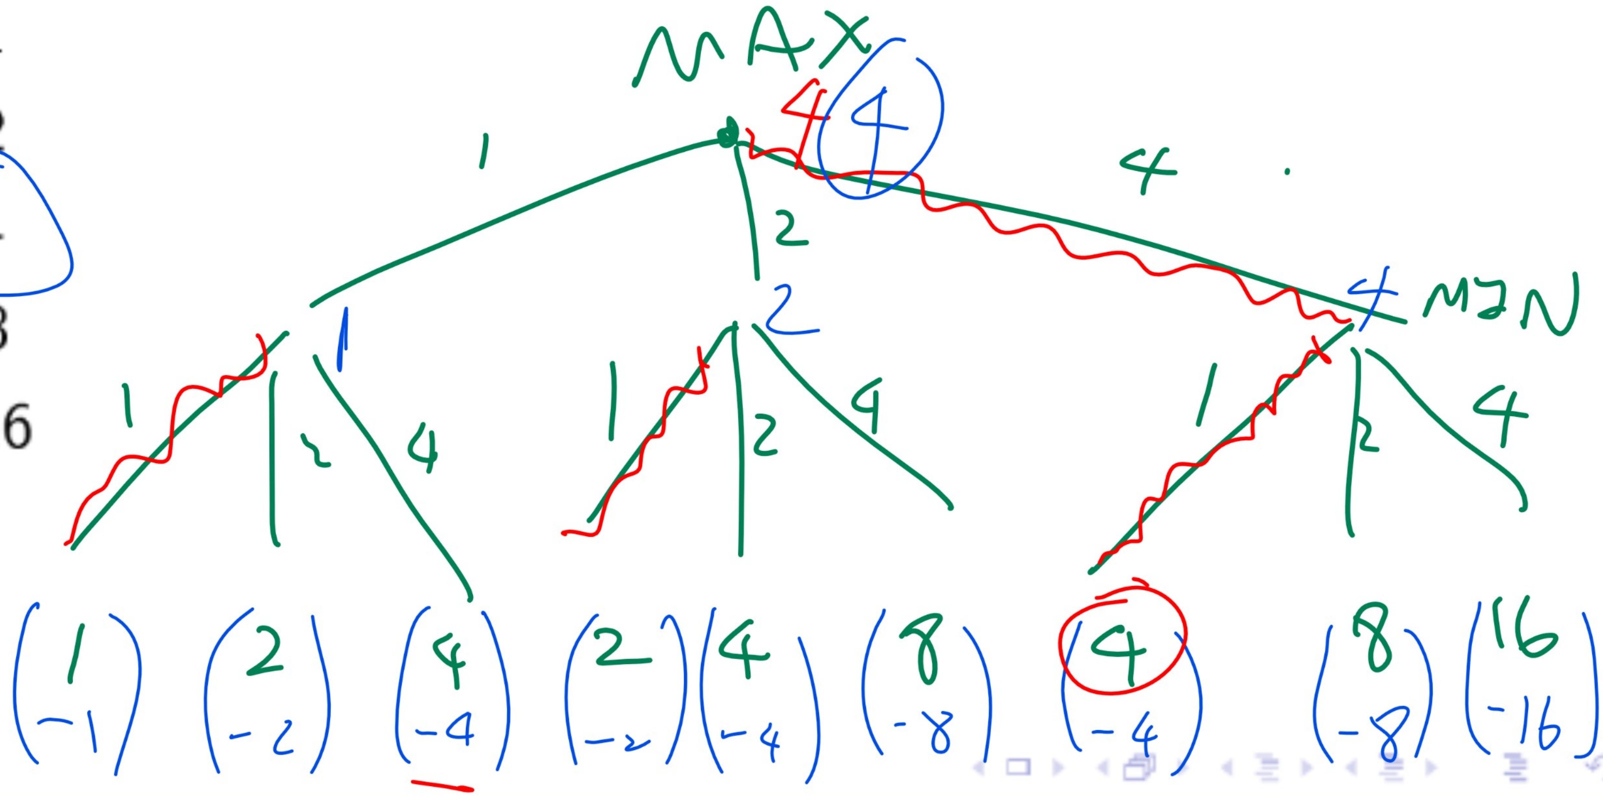

Minimax Algorithm

Algorithm

- Input: a game tree (V, E, c), and the current state s.
- Output: the value of the game at s.

- If s is a terminal state, return $c(s)$.
- If the player is MAX, return the maximum value over all successors.

$$\alpha(s) = \max_{s' \in s'(s)} \beta(s')$$

- If the player is MIN, return the minimum value over all successors.

$$\beta(s) = \min_{s' \in s'(s)} \alpha(s')$$

Backtracking

Discussion

- The optimal actions (solution paths) can be found by backtracking from all terminal states as in DFS.

$$s^{\star}(s) = \arg \max_{s' \in s'(s)} \beta(s') \text{ for MAX}$$

$$s^{\star}(s) = \arg \min_{s' \in s'(s)} \alpha(s') \text{ for MIN}$$

Non-deterministic Game

Discussion

- For non-deterministic games in which chance can make a move (dice roll or coin flip), use expected reward or cost instead.
- The algorithm is also called expectiminimax.

Game Tree with Chance Example 1

Quiz

- Fall 2005 Midterm Q7
- Max can pick L or R. If Max picks L, Chance picks L with probability 0.3 and R with probability 0.7. If Chance picks L, Min picks L to get 3, R to get 2, and if Chance picks R, Min gets 7. If Max picks R, Min picks L to get -1 and R to get 2. What is the value of the game?

Game Tree with Chance Example 1 Diagram

Quiz

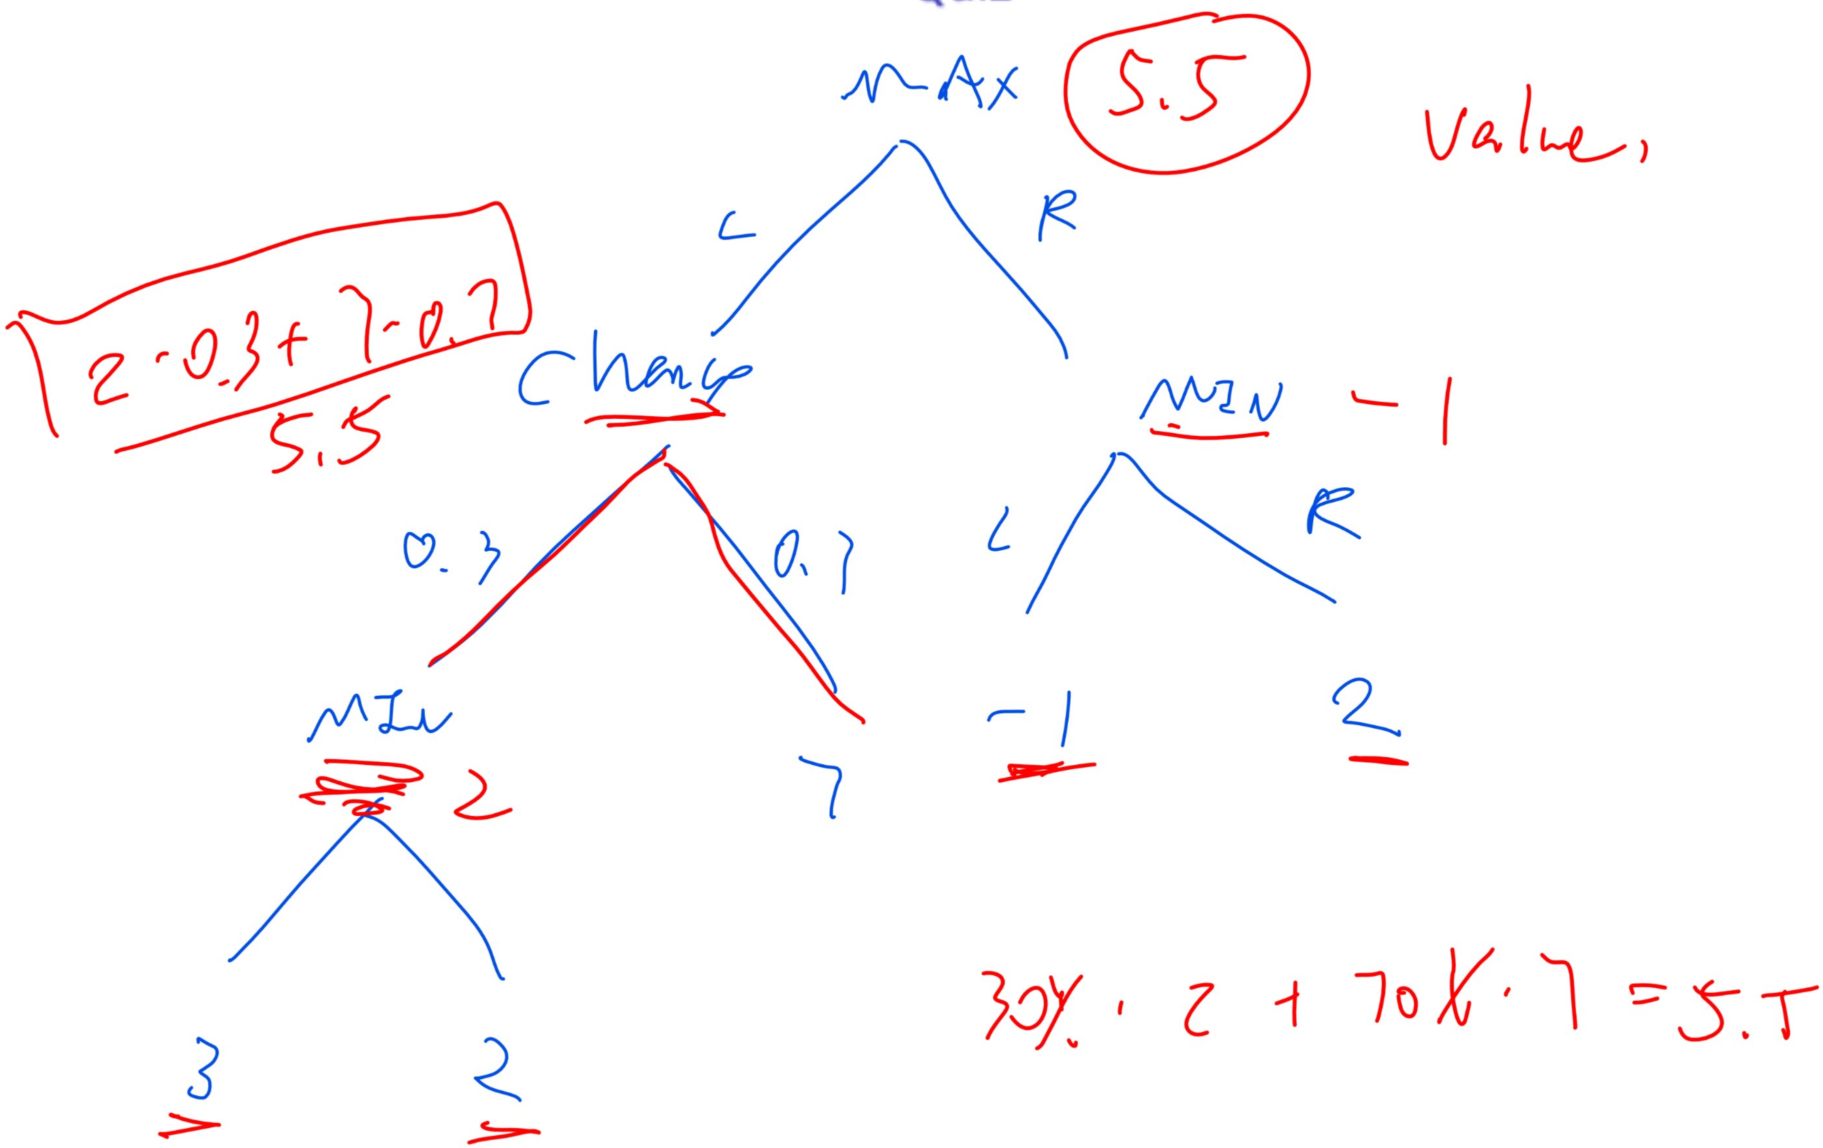

Game Tree with Chance Example 2

Quiz

Q3

- MAX, Chance (half-half), MIN sequentially choose an action H or T. The value of the terminal state is $\max\{x, 3 - x\}$, where x is the number of actions that are H. What is the expected value of the whole game?
- A: 1
- B: 1.5
- C: 2
- D: 2.5
- E: 3

if MIN moves first.

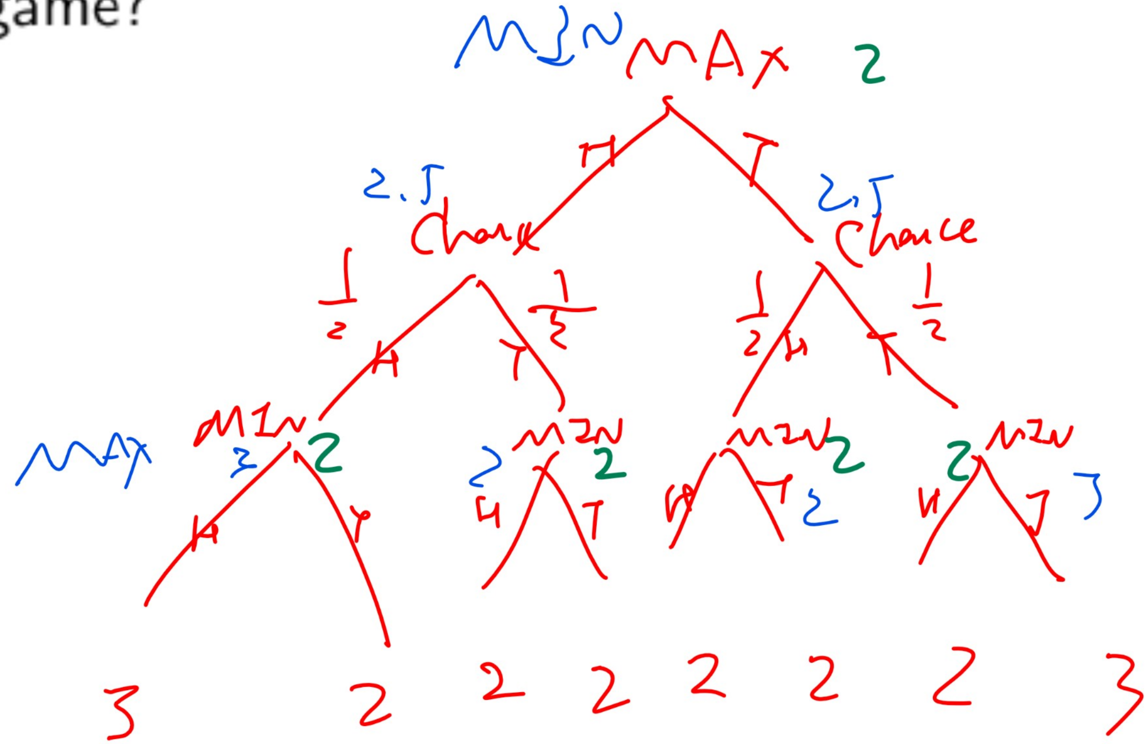

Game Tree with Chance Example 3

Quiz

Q4

- MIN, Chance (half-half), MAX sequentially choose an action H or T. The value of the terminal state is $\max\{x, 3 - x\}$, where x is the number of actions that are H. What is the value of the whole game?
- A: 1
- B: 1.5
- C: 2
- D: 2.5
- E: 3

Alpha Beta Pruning

Description

- During DFS, keep track of both α and β for each vertex.
- Prune the subtree with $\alpha \geqslant \beta$.

Alpha Beta Simple Example 1

Quiz

- Fall 2014 Final Q13
- After MAX picks L, MIN can pick L, M, R to get $3, 12, 8$.
  After MAX picks R, MIN can pick L, M, R to get $2, 15, 6$.
  Which vertices can be pruned.

α β pruning

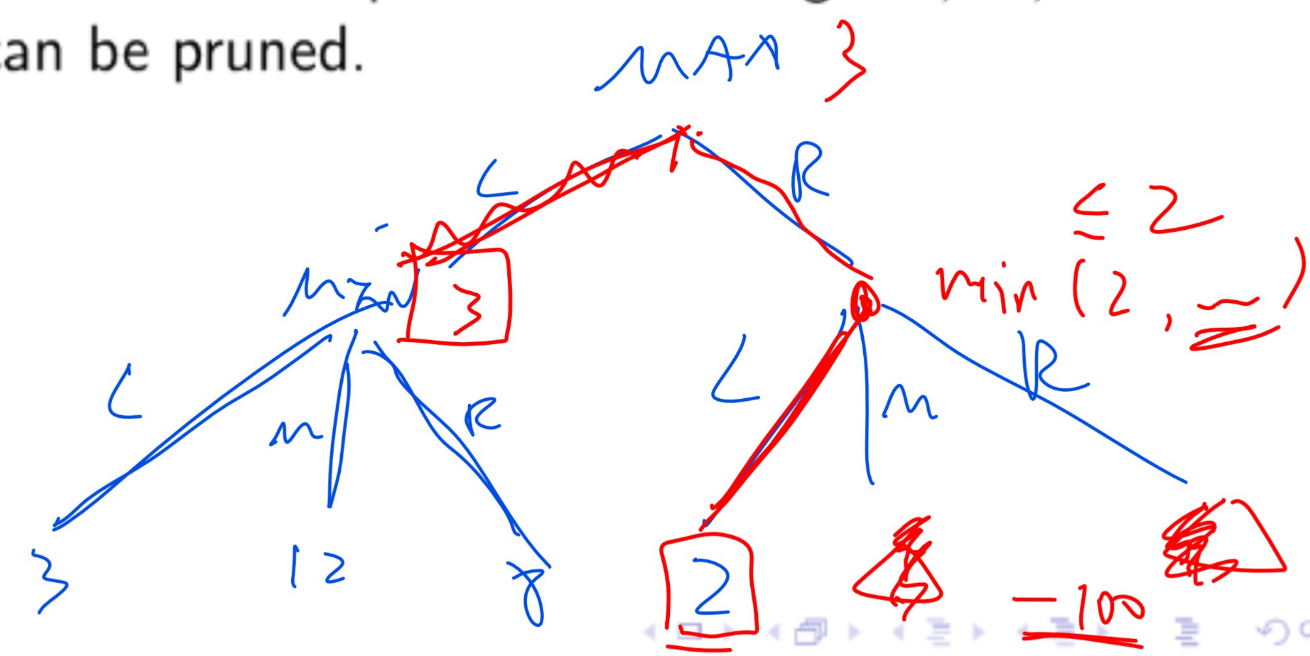

Alpha Beta Example 1

Quiz

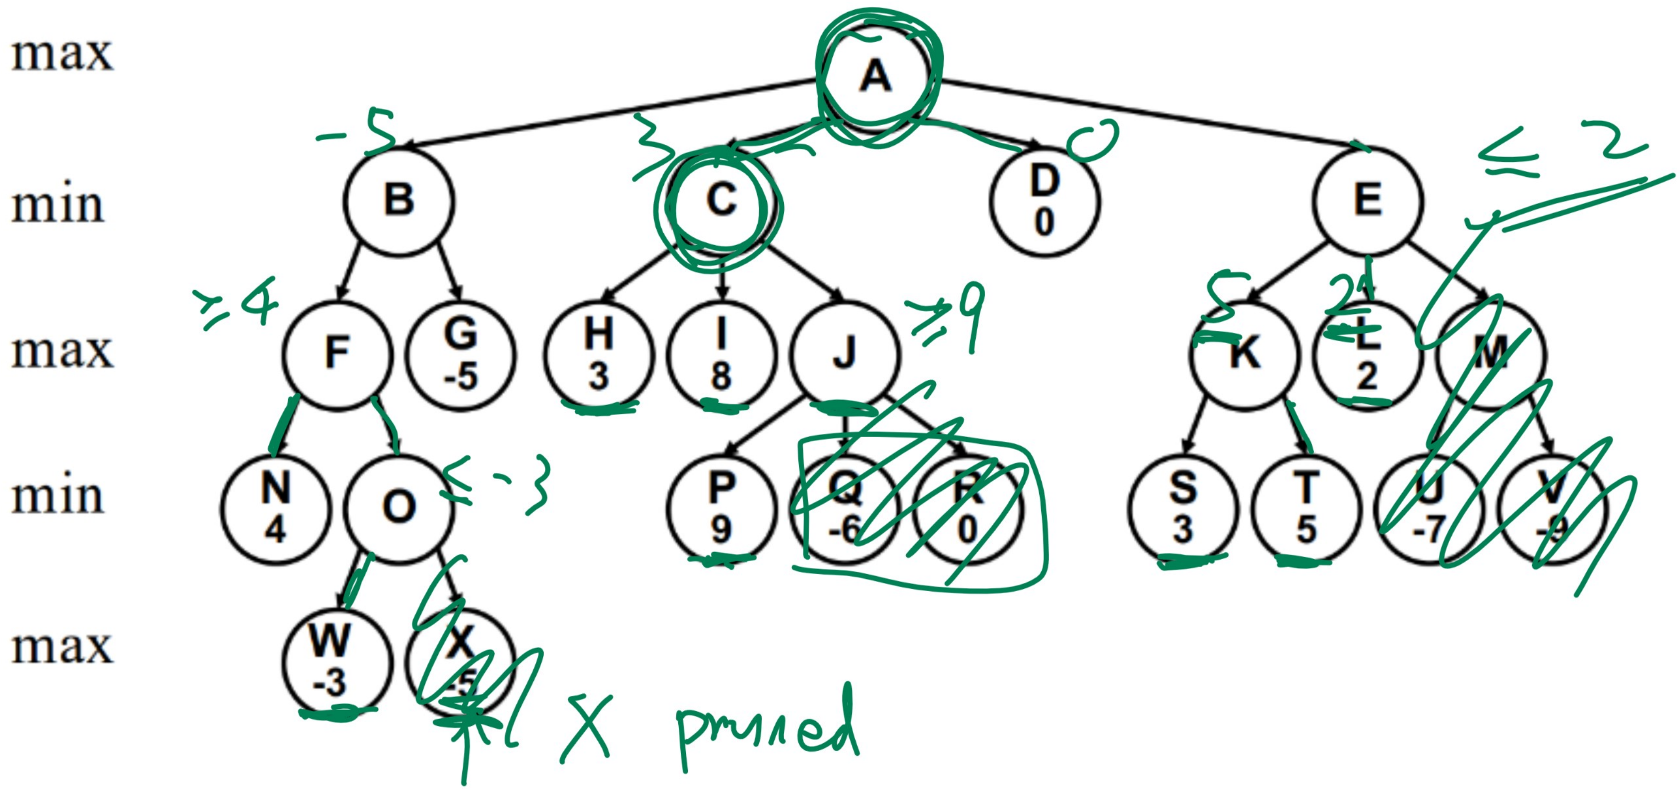

Alpha Beta Example 1 Continued

Quiz

Alpha Beta Example 2

Quiz

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. Alpha-Beta pruning is used. What is the number of branches (states) that can be pruned if the actions with smaller labels are searched first?
- A: 0
- B: 1
- C: 2
- D: 3
- E: 4

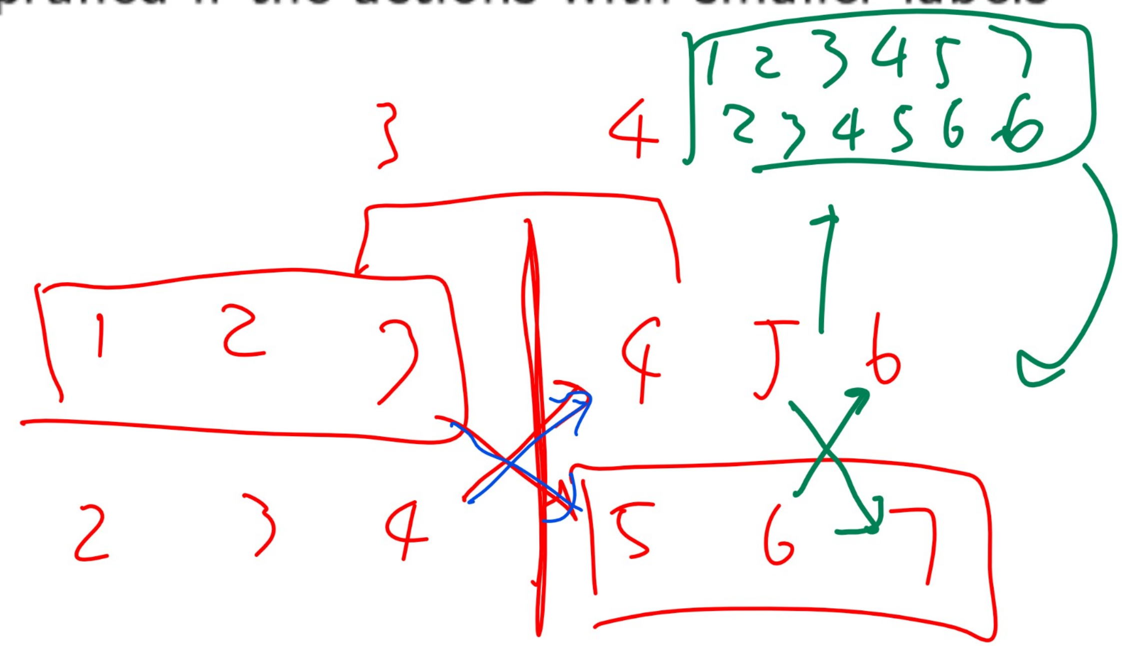

child 1 : 1 2 3 5 6 7

2 : 2 3 4 4 5 6

Alpha Beta Example 3

Quiz

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. Alpha-Beta pruning is used. What is the number of branches (states) that can be pruned if the actions with larger labels are searched first?
- A: 0
- B: 1
- C: 2
- D: 3
- E: 4

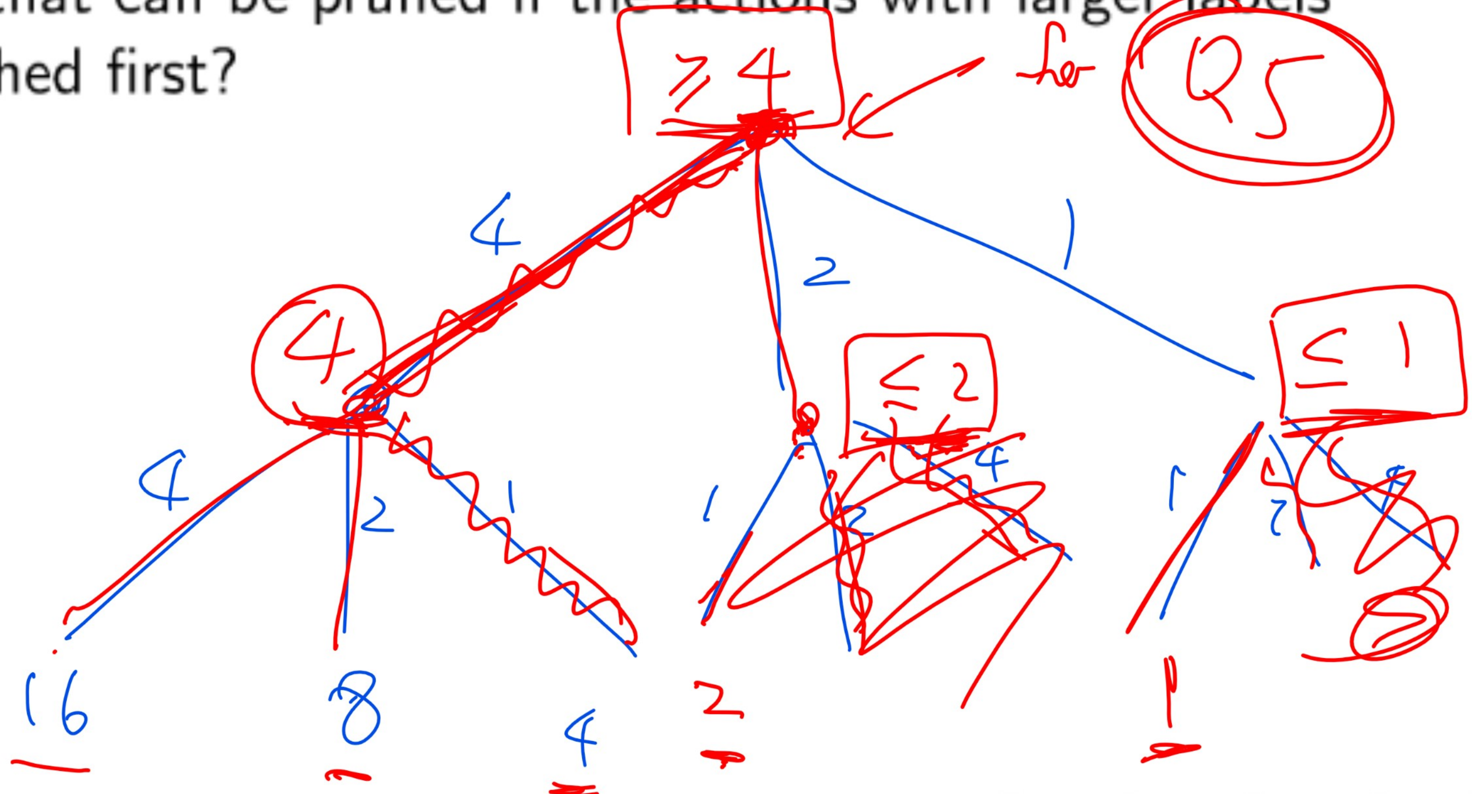

Alpha Beta Example 3

Quiz

Q5

- For a zero-sum game, the value to the MAX player if MAX plays $x_1 \in \{1, 2, 4\}$ and MIN plays $x_2 \in \{1, 2, 4\}$ is $x_1 \cdot x_2$. Alpha-Beta pruning is used. What is the maximum number of branches (states) that can be pruned if the actions can be searched in any order?

pruned = 0

- A: 2
- B: 3
- C: 4
- D: 5
- E: 6

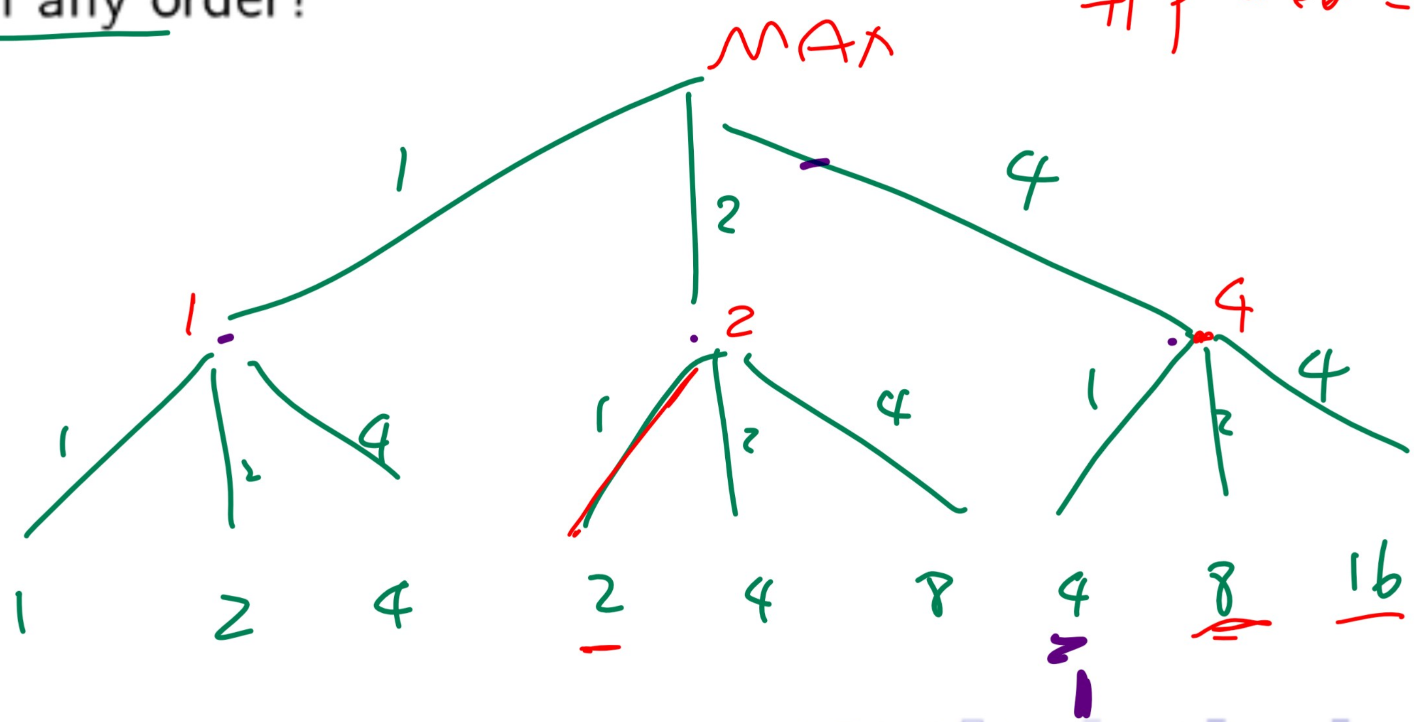

Alpha Beta Pruning Algorithm, Part I

Algorithm

- Input: a game tree (V, E, c), and the current state s.
- Output: the value of the game at s.
- If s is a terminal state, return $c(s)$.

Alpha Beta Pruning Algorithm, Part II

Algorithm

- If the player is MAX, return the maximum value over all successors.

$$\alpha(s) = \max_{s' \in s'(s)} \beta(s')$$

$$\beta(s) = \beta(\text{ parent } (s))$$

- Stop and return β if $\alpha \geqslant \beta$.
- If the player is MIN, return the minimum value over all successors.

$$\beta(s) = \min_{s' \in s'(s)} \alpha(s')$$

$$\alpha(s) = \alpha(\text{ parent } (s))$$

- Stop and return α if $\alpha \geqslant \beta$.

Alpha Beta Performance

Discussion

- In the best case, the best action of each player is the leftmost child.
- In the worst case, Alpha Beta is the same as minimax.

Static Evaluation Function

Definition

- A static board evaluation function is a heuristics to estimate the value of non-terminal states.
- It should reflect the player's chances of winning from that vertex.
- It should be easy to compute from the board configuration.

Linear Evaluation Function Example

Definition

- For Chess, an example of an evaluation function can be a linear combination of the following variables.

1. Material.
2. Mobility.
3. King safety.
4. Center control.

- These are called the features of the board.

Iterative Deepening Search

Discussion

- IDS could be used with SBE.
- In iteration d, the depth is limited to d, and the SBE of the non-terminal vertices are used as their cost or reward.

IDS with SBE Diagram

Discussion

Non Linear Evaluation Function

Discussion

- The SBE can be estimated given the features using a neural network.
- The features are constructed using domain knowledge, or a possibly a convolutional neural network.
- The training data are obtained from games between professional players.