CS540 Introduction to Artificial Intelligence
Lecture 21

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

August 2, 2021

Pirate Game Example

Quiz

- 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)

need 3 votes

need 3 votes

P5 : (0, 0, 0, 0, 100)

P4 : (0, 0, 0, 0, 100)

P3 : (0, 0, 99, 1, 0)

P2 : (0, 97, 0, 2, 1)

P1 : (97, 0, 1, 0, 2)

Pirate Game Example Diagram

Quiz

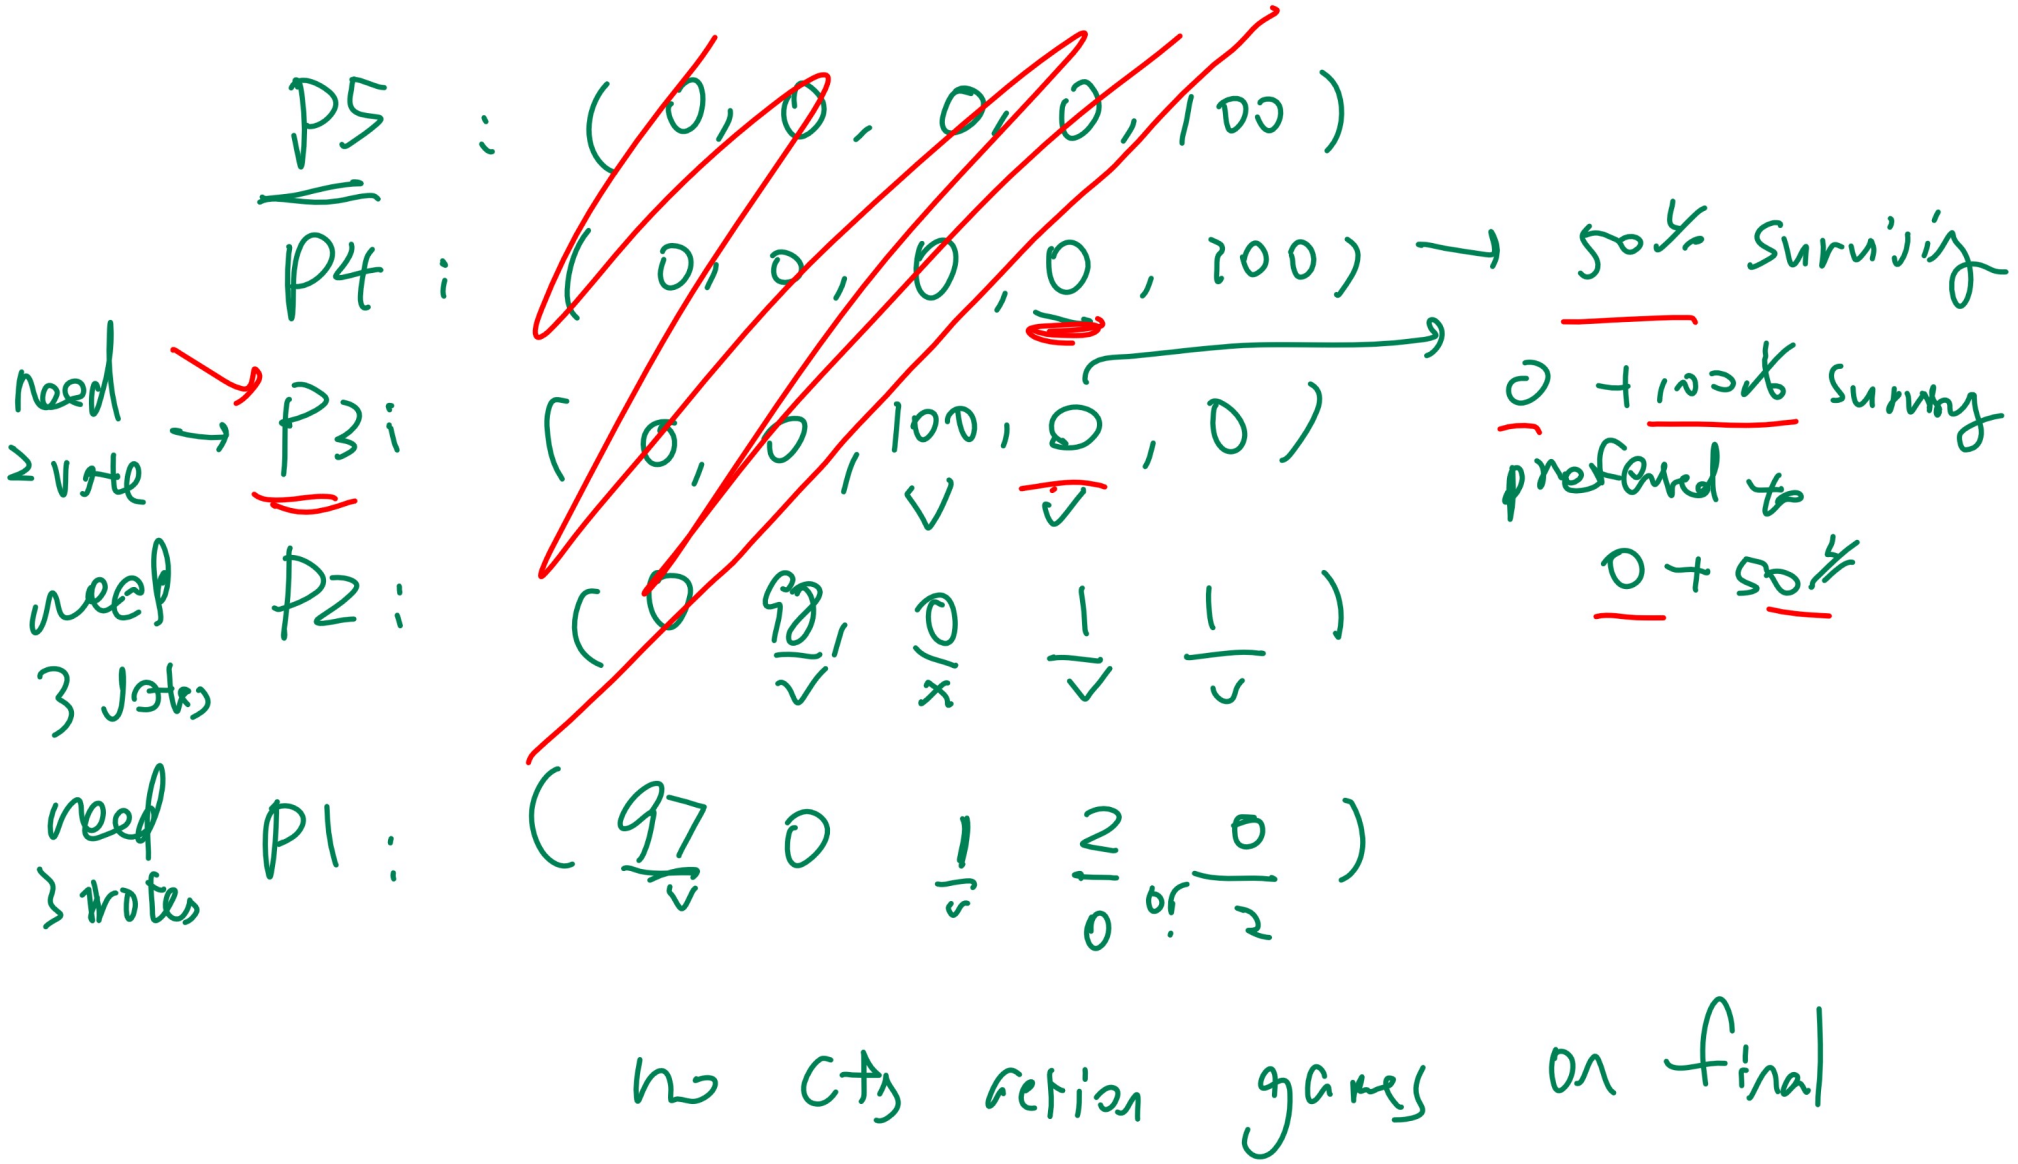

Midterm Adjustments

Admin

M16 Q11-14 ✓

M8 Q8

- Midterm stats posted on W4 page.
- Adjustments will be uploaded to Canvas tonight.
- M15 and M16 complete (not sure if auto-grading is working).
- Review Sessions and Friday Discussion to go through M15 and M16.
- Next week: play and solve more games, adversarial machine learning, brief intro to mechanism design, no new materials for the exams.

Remind Me to Start Recording

Admin

- The messages you send in chat will be recorded: you can change your Zoom name now before I start recording.

Lion Game Example

Quiz

- There are N lions, ordered by size, $i = 1, 2, 3, ..., N$, and a bunny. N takes an integer between 1 and 10 with equal probability (known to all the lions). Each lion i can choose to jump out and eat the slightly smaller lion $i - 1$, or stay hidden, and only lion 1 can eat the bunny. Each lion prefers eating to staying hungry to being eaten.
- What is the probability that the bunny is eaten?
- A: 0, B: $\frac{1}{3}$, C: $\frac{1}{2}$, D: $\frac{2}{3}$, E: 1

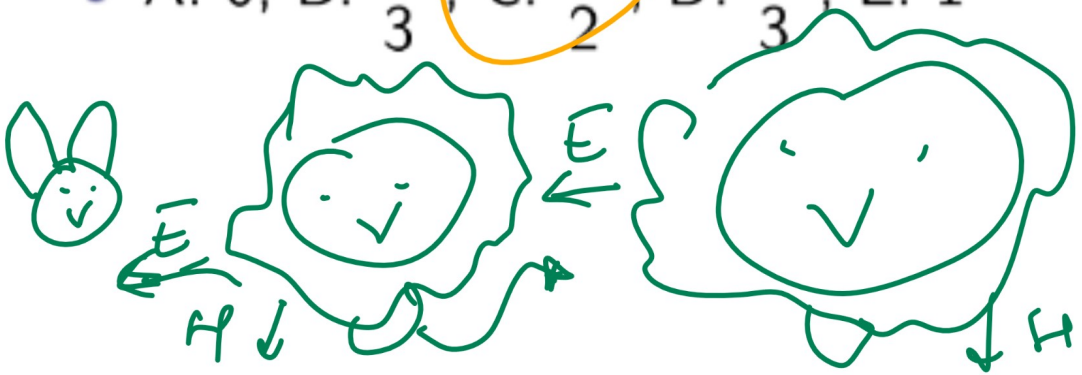

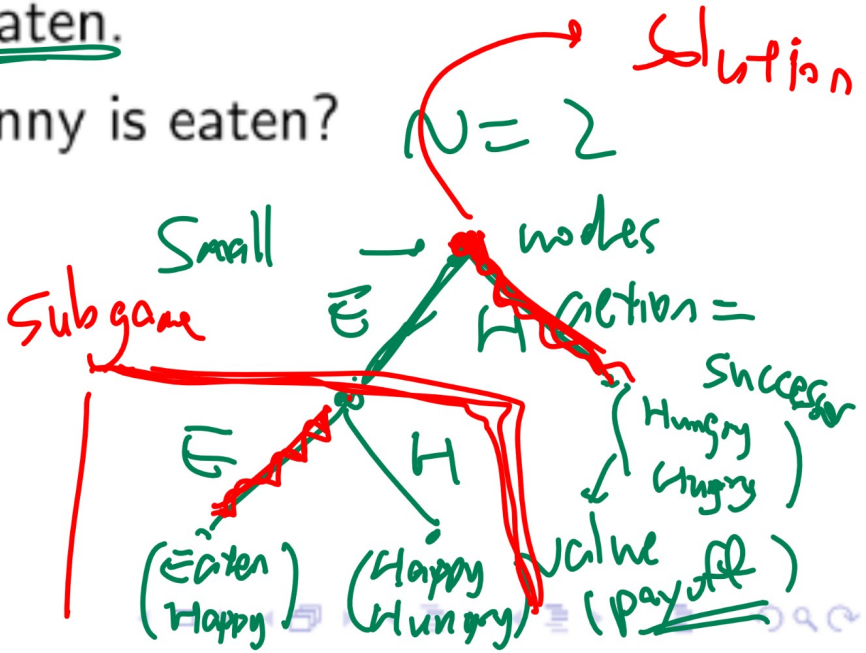

Lion Game Example Diagram

Quiz

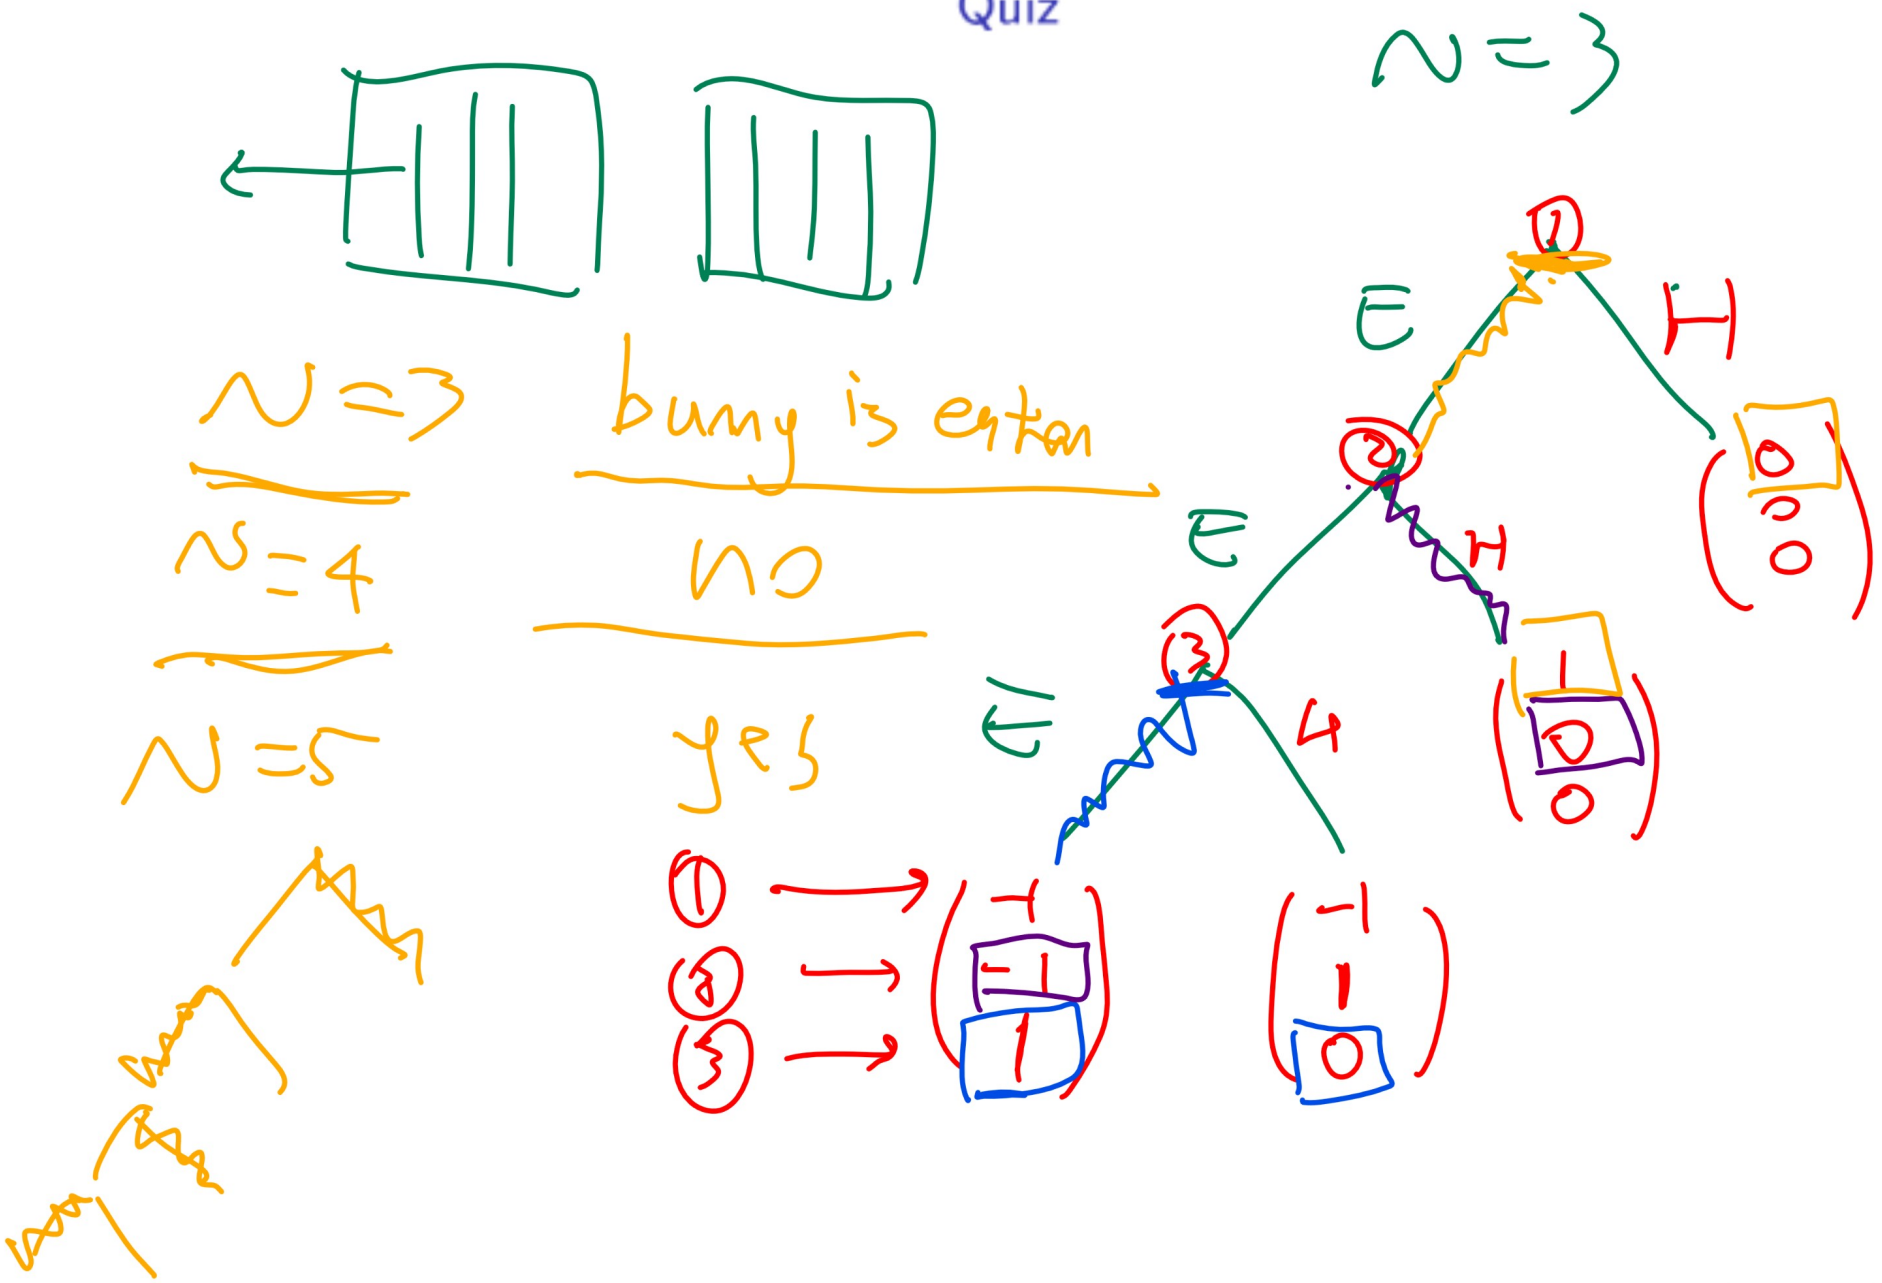

Tic Tac Toe Example

Motivation

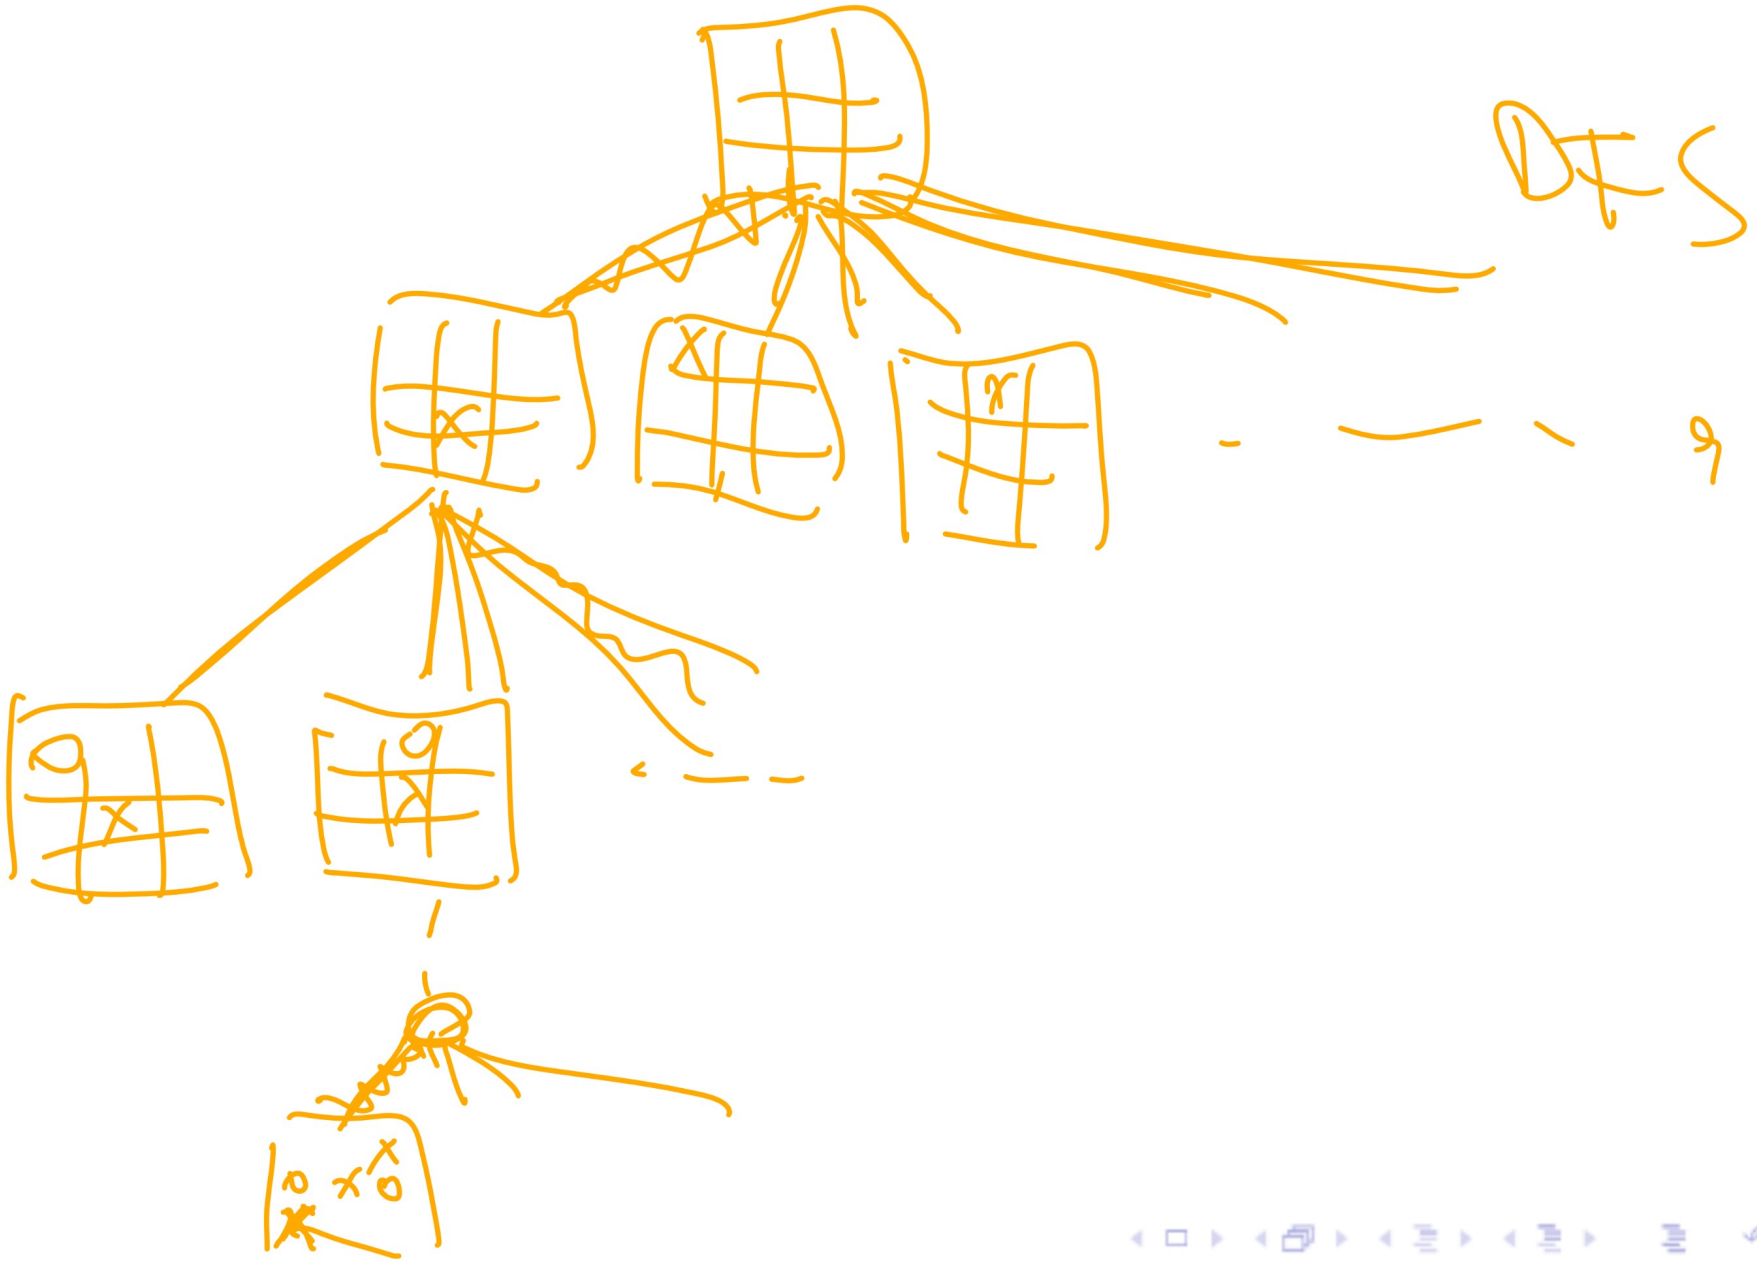

Minimax Algorithm

Description

- Use DFS on the game tree.

Minimax Example

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

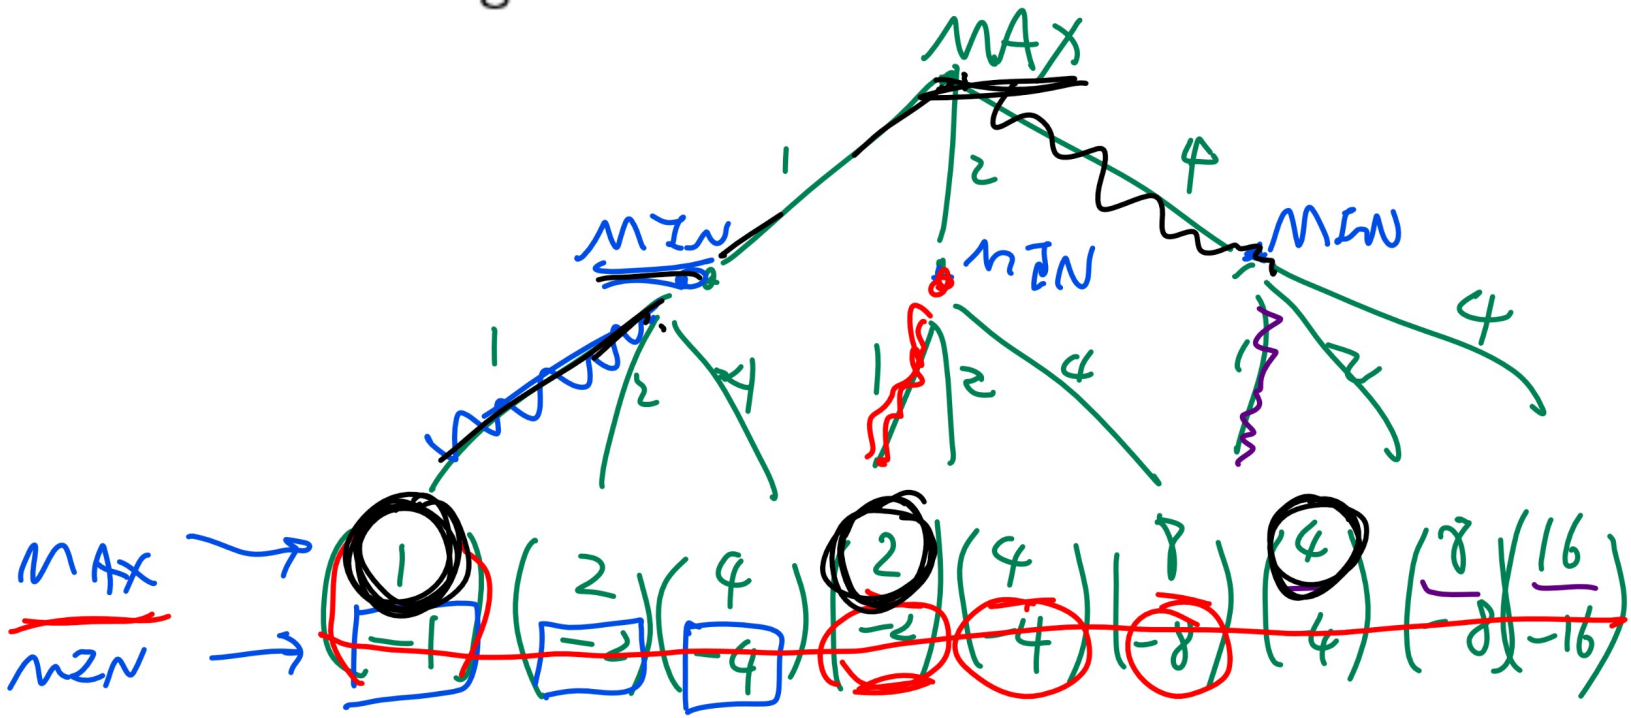

Minimax Performance

Discussion

- The time and space complexity is the same as DFS. Note that $D = d$ is the maximum depth of the terminal states.

$$T = 1 + b + b^2 + ... + b^d$$

$$S = (b - 1) \cdot d$$

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?

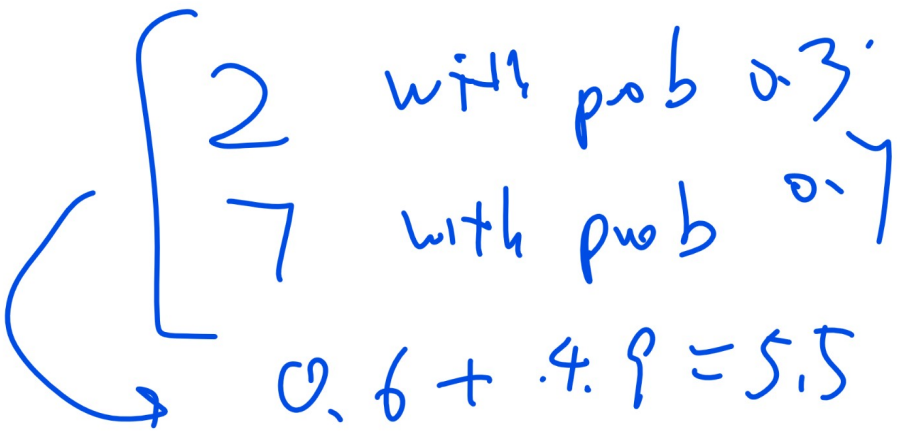

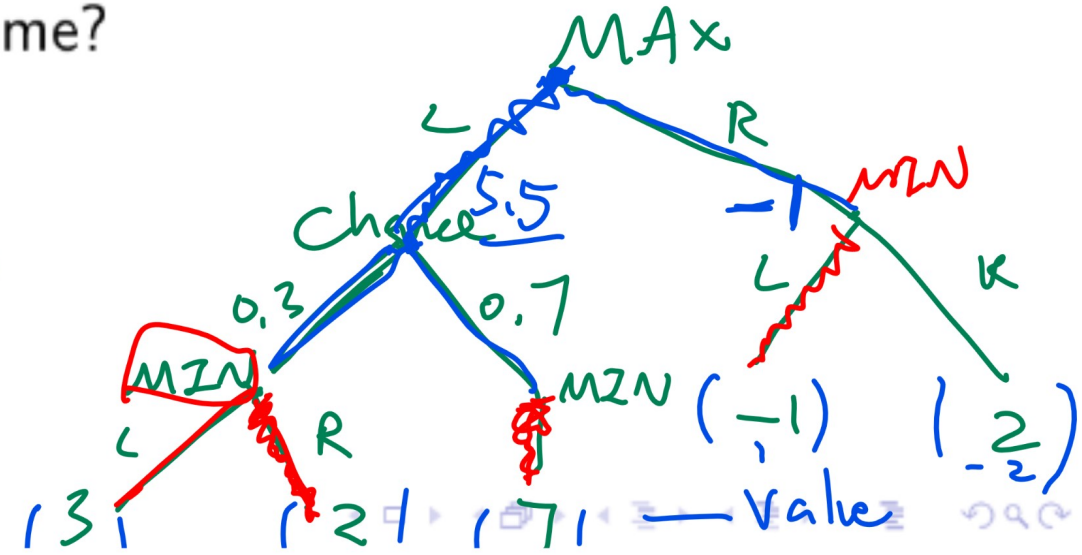

(=3) (-2 -7

Game Tree with Chance Example 1 Diagram

Quiz

Pruning

Motivation

- Time complexity is a problem because the computer usually has a limited amount of time to "think" and make a move.
- It is possible to reduce the time complexity by removing the branches that will not lead the current player to win. It is called the Alpha-Beta pruning.

Alpha Beta Pruning

Description

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

value for MAX

value for MIN

Alpha Beta Example 1

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

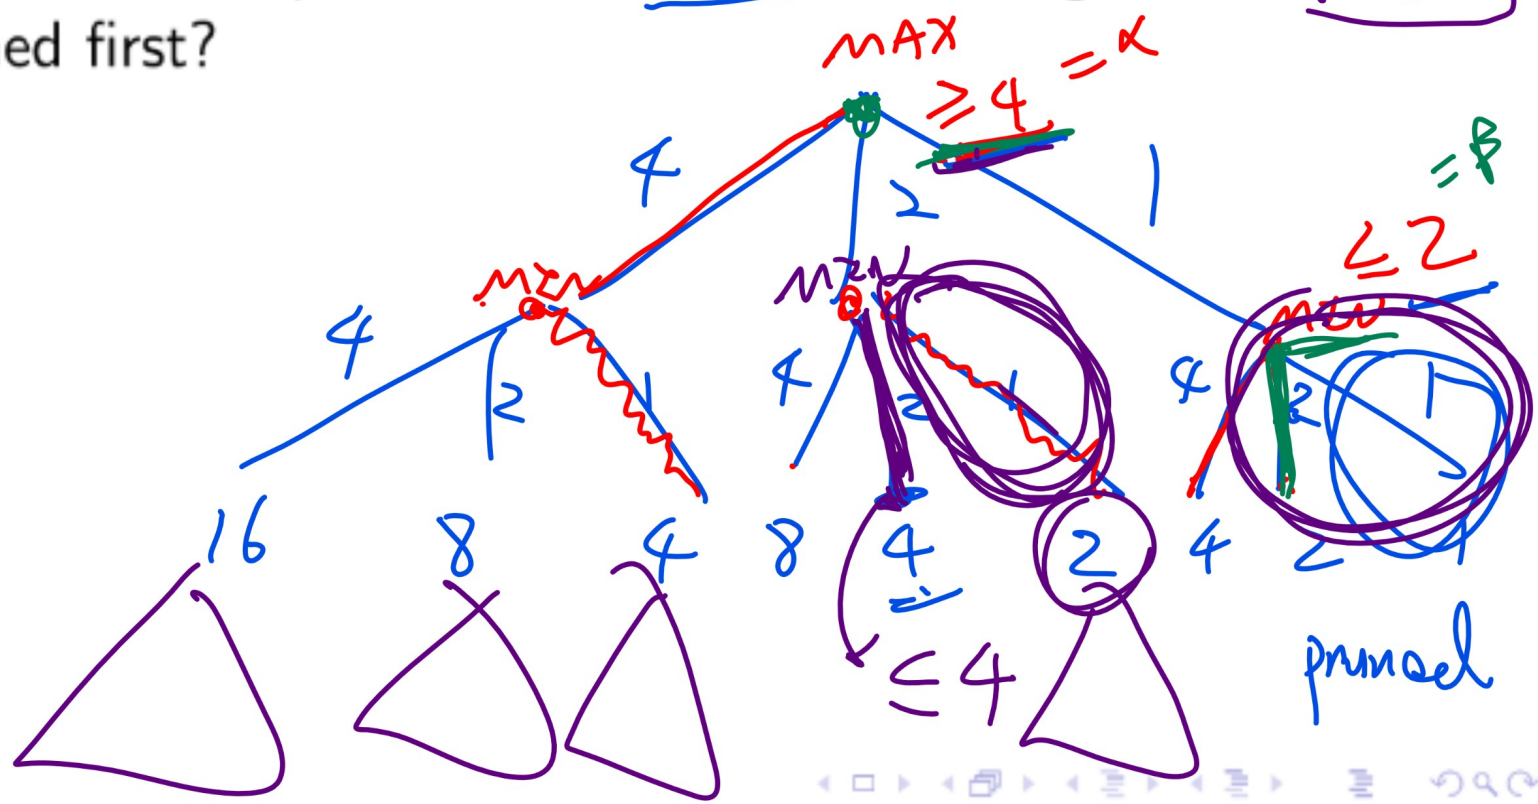

Alpha Beta Example 2

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

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 maximum number of branches (states) that can be pruned if the actions can be searched in any order?
  - A: 2
  - B: 3
  - C: 4
  - D: 5
  - E: 6

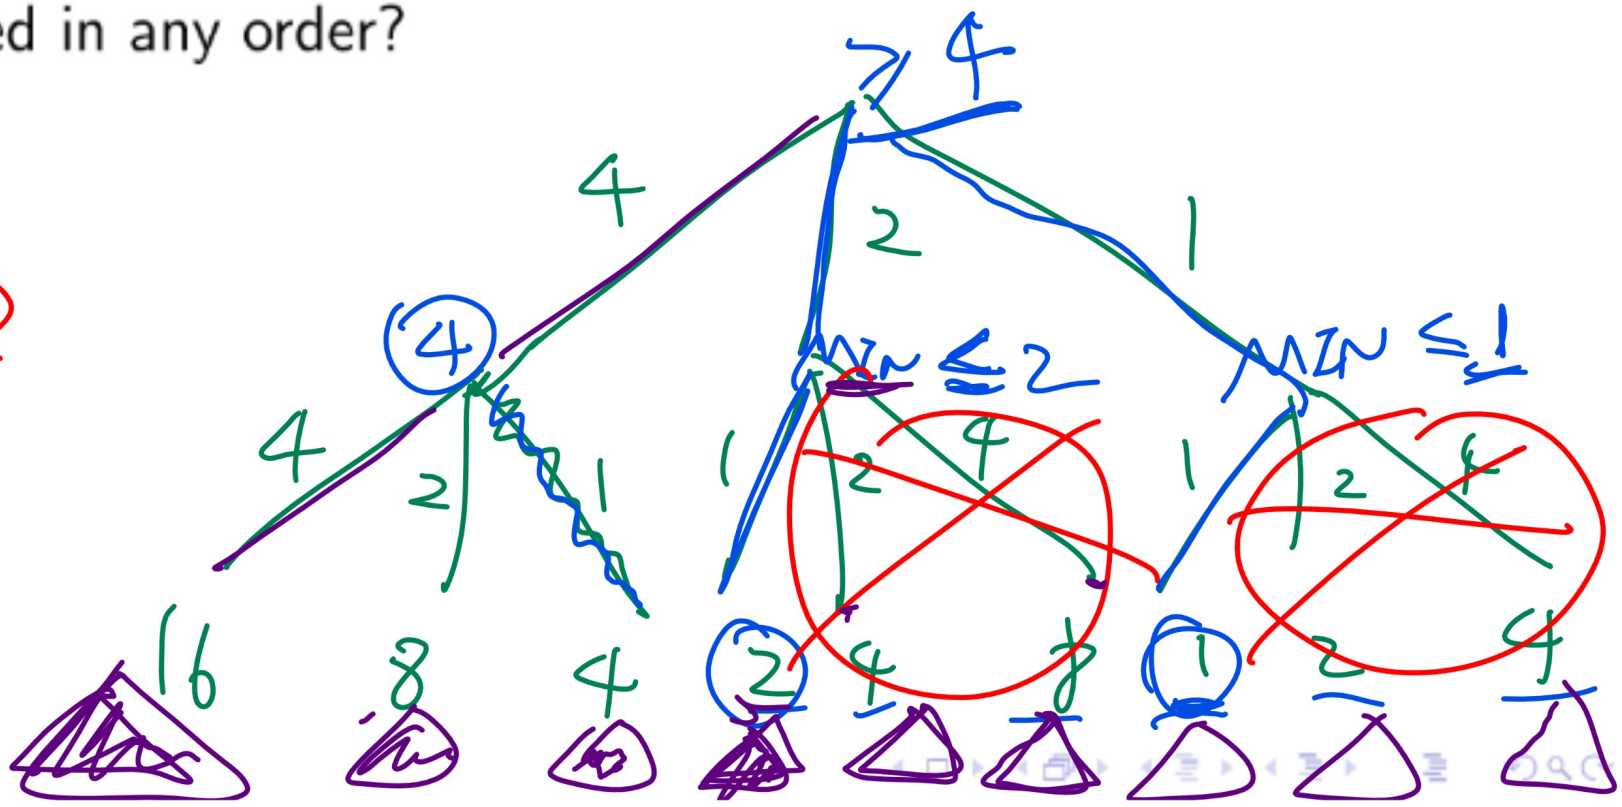

Alpha Beta Example 4

Quiz

Alpha Beta Example 4 Continued

Quiz

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.

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.