CS 540 Introduction to Artificial Intelligence

Games I

Yudong Chen
University of Wisconsin-Madison

Nov 23, 2021

Announcements

- **Homeworks**: HW9 due next Thursday. Start early!
- Class roadmap:
  - Today: Games I
  - Tuesday: Games II
  - Next: Reinforcement Learning

Outline

- Introduction to game theory
  - Properties of games, mathematical formulation
- Simultaneous Games
  - Normal form, strategies, dominance, Nash equilibrium
- Sequential Games (time permitted)
  - Game trees, minimax, search approaches

Sequential Decision Making

Suppose we have an **agent** **interacting** with the **world**

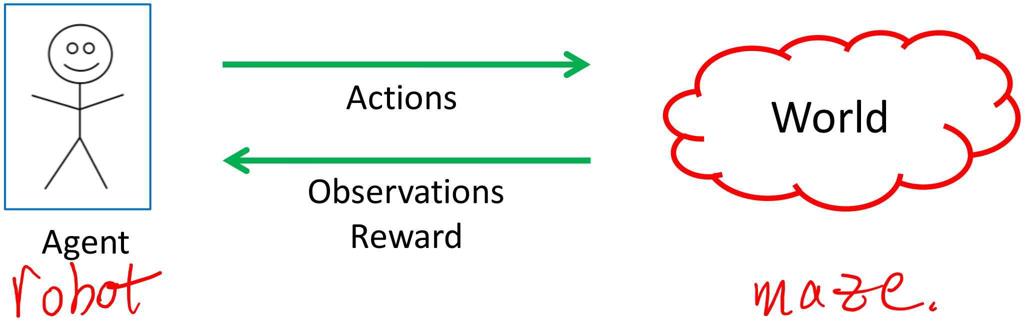

- Agent receives a reward based on state of the world
  - **Goal**: maximize reward/utility (or minimize cost/penalty)
  - Note: now **data** consists of actions & observations

Games: Multiple Agents

robot ⇆ maze
↓
- not maximize reward
- passive.

Games setup: **multiple** agents

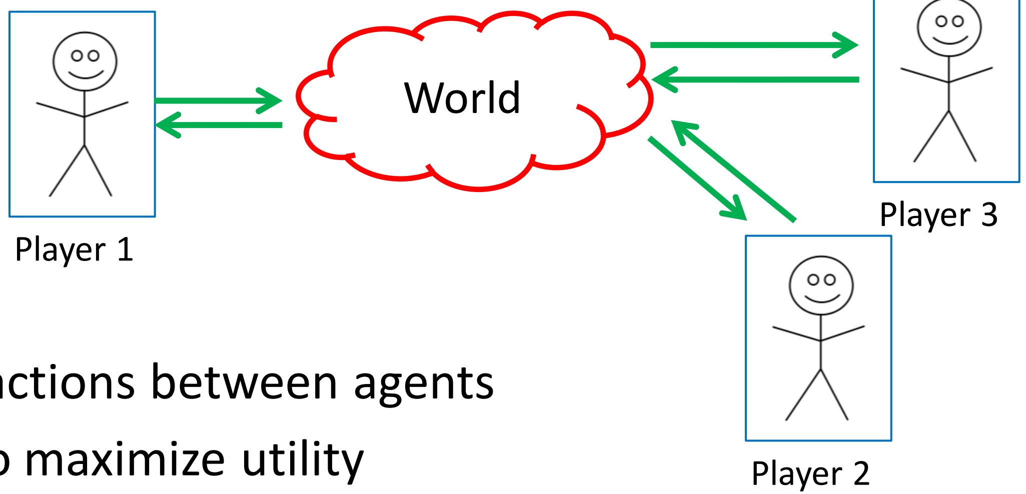

- Now: interactions between agents
- Still want to maximize utility
- **Strategic** decision making

Modeling Games: Properties

Let's work through **properties** of games

- **Number** of agents/players
- State & action spaces: **discrete** or **continuous**
- **Finite** or **infinite** time/ # round
- **Deterministic** or **random**
- **Sum** of rewards: zero or positive or negative
- **Sequential** or **simultaneous**

Wiki

Property 1: **Number** of players

- Games: ≥ 2 players
- Typically a finite number of players

Stock market. → convinient w/ ∞ # agents.

Evolution

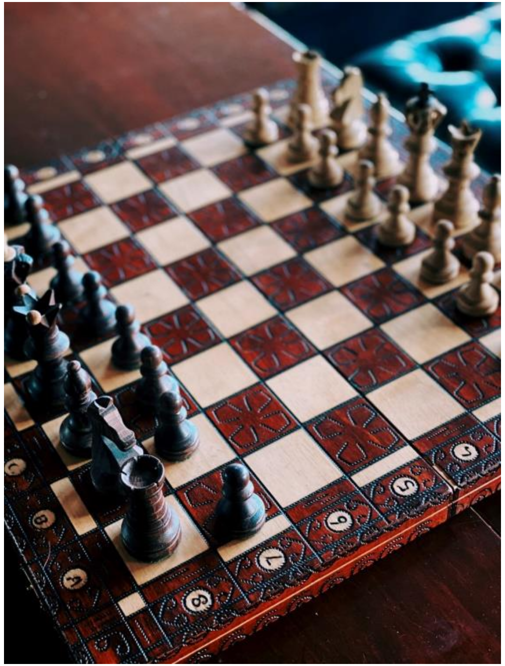

Property 2: Discrete or Continuous

Let's work through **properties** of games

state space

- Recall the **world**. It is in a particular state, from a set of states
- Similarly, the actions the player takes are from an action space
- How big are these spaces? Finite, countable, uncountable?

{win, loss}

Property 3: **Finite** or **Infinite**

Let’s work through **properties** of games

- Most real-world games **finite**
- Lots of single-turn games; end immediately
  - Ex: rock/paper/scissors
- Other games’ rules (state & action spaces) enforce termination
  - Ex: chess under FIDE rules ends in at most 8848 moves
- **Infinite example:** pick integers. First player to play a 5 loses

Property 4: Deterministic or Random

Let's work through **properties** of games

- Is there **chance** in the game?
- Note: randomness enters in different ways
- Not referring to randomness in players' strategies
  - E.g. RPS is deterministic

Property 5: Sums

Let's work through **properties** of games

- **Sum**: zero or positive or negative
- Zero sum: for one player to win, the other has to lose (by same amount)
  - No "value" created

Red \ Blue	*A*	*B*	*C*
1	30, −30	−10, 10	20, −20
2	−10, 10	20, −20	−20, 20

- Can have other types of games: positive sum, negative sum.
  - Example: prisoner's dilemma

Property 6: Sequential or Simultaneous

Let’s work through **properties** of games

- Sequential or simultaneous
- Simultaneous: all players take action at the same time
- Sequential: take turns

- Simultaneous: players do not have information of others’ moves. Ex: **RPS**
- Sequential: may or may not have **perfect information**

Examples

Let's apply this to examples:

1. Chess: **2-player, discrete, finite, deterministic, zero-sum, sequential (perfect information)**
2. RPS: **2-player, discrete, finite, deterministic, zero-sum,** simultaneous
3. Mario Kart: **4-player, continuous, infinite(?), random, zero-sum,** simultaneous

Another Example: Prisoner's Dilemma

Famous example from the '50s.

Two prisoners A & B. Can choose to betray the other or not.

– A and B both betray, each of them serves two years in prison
– One betrays, the other doesn't: betrayer free, other three years
– Both do not betray: one year each

optimal individual action

Properties: **2-player, discrete, finite,**

not necessarily

deterministic, negative-sum,

optimal joint action

simultaneous

Player B / Player A	*Stay silent*	*Betray*
Stay silent	−1, + −1 = −2	−3, + 0 = −3
Betray	0, + −3 = −3	−2, + −2 = −4

Why Do These Properties Matter?

Categorize games in different groups

- Can focus on understanding/analyzing/“solving” particular groups
- **Abstract** away details and see common patterns
- Understand how to produce a “good” overall outcome

- relate abstract concept w/ concrete games.
- Work on ~~toy~~ small examples
- RPS, PD.

Break & Quiz

Q 1.1: Which of these are zero-sum games?

(i) Rock, Paper, Scissors

(ii) Prisoner’s Dilemma

- A. Neither
- B. (i) but not (ii)
- C. (ii) but not (i)
- D. Both

Break & Quiz

Q 1.1: Which of these are zero-sum games?

(i) Rock, Paper, Scissors

(ii) Prisoner’s Dilemma

- A. Neither
- **B. (i) but not (ii)**
- C. (ii) but not (i)
- D. Both

Break & Quiz

Q 1.1: Which of these are zero-sum games?

(i) Rock, Paper, Scissors

(ii) Prisoner's Dilemma

- A. Neither (Rock, Paper, Scissors is, clearly)
- **B. (i) but not (ii)**
- C. (ii) but not (i) (Rock, Paper, Scissors is, clearly)
- D. Both (Prisoner's Dilemma is not, recall the normal form matrix)

Break & Quiz

Q 1.2: Which of these is false?

- A. Monopoly is not deterministic.
- B. A game can be sequential but not have perfect information.
- C. Battleship has perfect information.
- D. Prisoner's dilemma is a simultaneous game.

Break & Quiz

Q 1.2: Which of these is false?

- A. Monopoly is not deterministic.
- B. A game can be sequential but not have perfect information.
- **C. Battleship has perfect information.**
- D. Prisoner’s dilemma is a simultaneous game.

Break & Quiz

Q 1.2: Which of these is false?

- A. Monopoly is not deterministic. (True: you roll dice.)
- B. A game can be sequential but not have perfect information. (True, and in fact Battleship is an example.)
- **C. Battleship has perfect information.**
- D. Prisoner's dilemma is a simultaneous game. (Also true: single round, no turns.)

Simultaneous Games

Simpler setting, easier to analyze

- Can express reward with a simple diagram
- Ex: for prisoner's dilemma

Player 2 / Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

Normal Form

Mathematical description of simult. games. Has:

- n players $\{1,2,\ldots,n\}$
- Player i strategy a_i from A_i.
  - Strategy of **all** players: $a = (a_1, a_2, \ldots, a_n)$
- Player i gets rewards $u_i(a)$ for any outcome
  - **Note**: reward depends on other players!

- Setting: all of these spaces, rewards are **known**

action

→ action space.

Example of Normal Form

Ex: Prisoner's Dilemma

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

- 2 players, 2 actions: yields 2x2 matrix
- Strategies: {Stay silent, betray} (i.e., binary)
- Rewards: {0,-1,-2,-3}

$a_{-i} = \{a_1, a_2, \cdots, a_{i-1}, a_{i+1}, a_{i+2}, \cdots, a_n\}$.

Dominant Strategies

Let's analyze such games. Some strategies are better

- Dominant strategy: if a_i better than a_i' *regardless* of what other players do, a_i is **dominant** (a_i is always the best).
- I.e.,

$$u_i(a_i, a_{-i}) \geq u_i(a_i', a_{-i}) \forall a_i' \neq a_i \text{ and } \forall a_{-i}$$

All of the other entries of *a* excluding *i*

- Doesn't always exist!

Dominant Strategies Example

Back to Prisoner's Dilemma

- Examine all the entries: betray dominates
- Check:

Player 2 Player 1	*Stay silent*	*Betray*
Stay silent	−1, −1	−3, 0
Betray	0, −3	−2, −2

- Note: normal form helps **locate** dominant/dominated strategies.

Dominant Strategies May Not Exist

Rock-Paper-Scissor

- No dominant strategy

Player 2 Player 1	*Rock*	*Paper*	*Scissor*
Rock	0, 0	-1, 1	1,-1
Paper	1,-1	0, 0	-1,1
Scissor	-1,1	1,-1	0,0

Equilibrium

$= (a_1^*, a_2^*, \ldots, a_n^*)$

$a*$ is an equilibrium if all the players do not have an incentive to **unilaterally deviate**

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \forall a_i \in A_i$$

a_i^* is dominant $= u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i}), \forall a_i \in A_i, \forall a_{-i} \in A_{-i}$

- All players dominant strategies -> equilibrium
- Converse doesn't hold (don't need dominant strategies to get an equilibrium)

dominant strategies

equilibrium

Pure and Mixed Strategies

So far, all our strategies are deterministic: "**pure**"

- Take a particular action, no randomness

Can also randomize actions: "**mixed**"

- Assign probabilities x_i to each action

$$x_i(a_i), \text{ where } \sum_{a_i \in A_i} x_i(a_i) = 1, x_i(a_i) \geq 0$$

x_i is a mixed strategy for player i.

- Note: have to now consider **expected rewards**

e.g. 2-player: $u_1(x_1, x_2) = \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} u_1(a_1, a_2)\, x_1(a_1)\, x_2(a_2)$

Nash Equilibrium

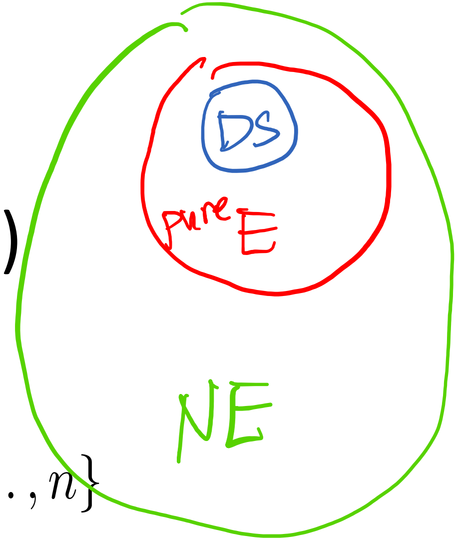

Consider the mixed strategy $x^* = (x_1^*, \ldots, x_n^*)$

- This is a **Nash equilibrium** if

$$u_i(x_i^*, x_{-i}^*) \geq u_i(x_i, x_{-i}^*) \quad \forall x_i \in \Delta_{A_i}, \forall i \in \{1, \ldots, n\}$$

Better than doing anything else, "**best response**"

Space of probability distributions

x_i^* is the best response to x_{-i}^*, $\forall i$.

- Intuition: nobody can **increase expected reward** by changing only their own strategy. A type of solution!

Properties of Nash Equilibrium

Major result: (Nash '51)

- Every finite game has at least one Nash equilibrium
  - But not necessarily **pure** (i.e., deterministic strategy)
- Could be more than one!
- Searching for Nash equilibria: computationally **hard**!

Break & Quiz

Q 2.1: Which of the following is **false**

(i) Rock/paper/scissors has a dominant pure strategy

(ii) There is no pure equilibrium for rock/paper/scissors

- A. Neither
- B. (i) but not (ii)
- C. (ii) but not (i)
- D. Both

Break & Quiz

Q 2.1: Which of the following is **false**?

(i) Rock/paper/scissors has a dominant pure strategy

(ii) There is no pure equilibrium for rock/paper/scissors

- A. Neither
- **B. (i) but not (ii)**
- C. (ii) but not (i)
- D. Both

Rock-Paper-Scissor

- No dominant strategy
- No pure equilibrium
- Has mixed strategy NE

$$a^* = \left(\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right),\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)\right).$$

Player 2 / Player 1	*Rock*	*Paper*	*Scissor*
Rock	0, 0	-1, 1	1,-1
Paper	1,-1	0, 0	-1,1
Scissor	-1,1	1,-1	0,0

$$a = \left((1,0,0),\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)\right)$$

NOT NE

b/c player 2 has incentive to deviate unilaterally.

Break & Quiz

Q 2.2: Which of the following is true

(i) Nash equilibria require each player to know other possible players' strategies

(ii) Nash equilibria require rational play

- A. Neither
- B. (i) but not (ii)
- C. (ii) but not (i)
- D. Both

Break & Quiz

Q 2.2: Which of the following is **true**

(i) Nash equilibria require each player to know other players' possible strategies

(ii) Nash equilibria require rational play

- A. Neither
- B. (i) but not (ii)
- C. (ii) but not (i)
- **D. Both**

Break & Quiz

Q 2.2: Which of the following is true

(i) Nash equilibria require each player to know other players' possible strategies

(ii) Nash equilibria require rational play

- A. Neither (See below)
- B. (i) but not (ii) (Rational play required: i.e., what if prisoners desire longer jail sentences?)
- C. (ii) but not (i) (The basic assumption of Nash equilibria is knowing all of the strategies involved)
- D. **Both**

Sequential Games

More complex games with multiple moves

- Instead of normal form, **extensive form**
- Represent with a **tree**
- Perform search over the tree

- Can still look for Nash equilibrium
  - Or, other criteria like **maximin** / **minimax**

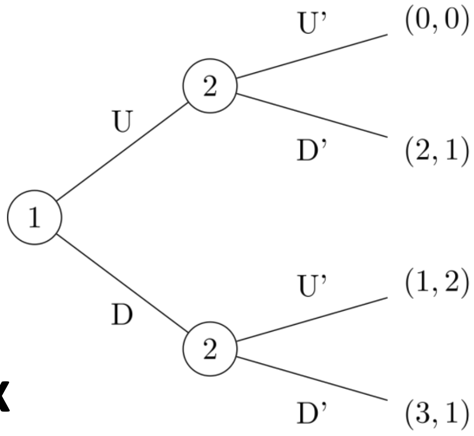

Wiki

II-Nim: Example Sequential Game

2 piles of sticks, each with 2 sticks.

- Each player takes one or more sticks from pile
- Take last stick: lose

(ii, ii)

- Two players: Max and Min
- If Max wins, the score is **+1**; otherwise **-1**
- Min's score is –Max's
- Use Max's as the score of the game

Game Trajectory

(ii, ii)

Game Trajectory

(ii, ii)

Max takes one stick from one pile

(i, ii)

Game Trajectory

(ii, ii)

Max takes one stick from one pile

(i, ii)

Min takes two sticks from the other pile

(i,-)

Game Trajectory

(ii, ii)

Max takes one stick from one pile

(i, ii)

Min takes two sticks from the other pile

(i,-)

Max takes the last stick

(-,-)

Max gets score **-1**

Game tree for II-Nim

Two players:
Max and Min

(ii ii) Max

who is to move at this state

Convention: score is w.r.t. the first player Max. Min's score = – Max

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

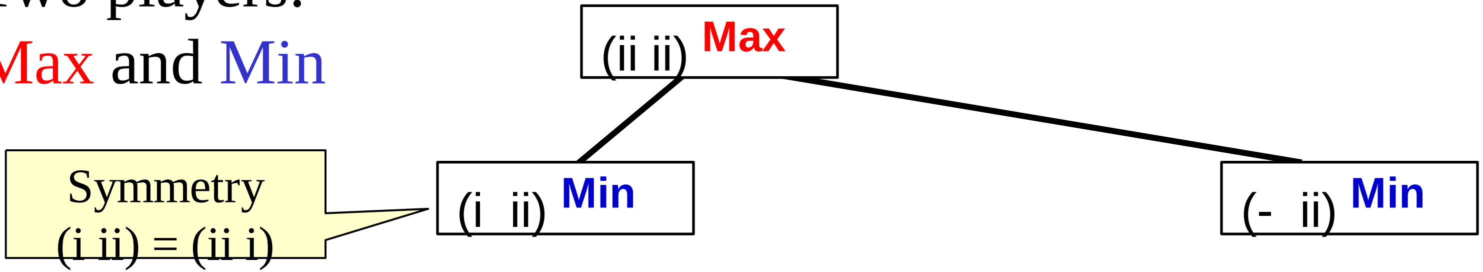

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

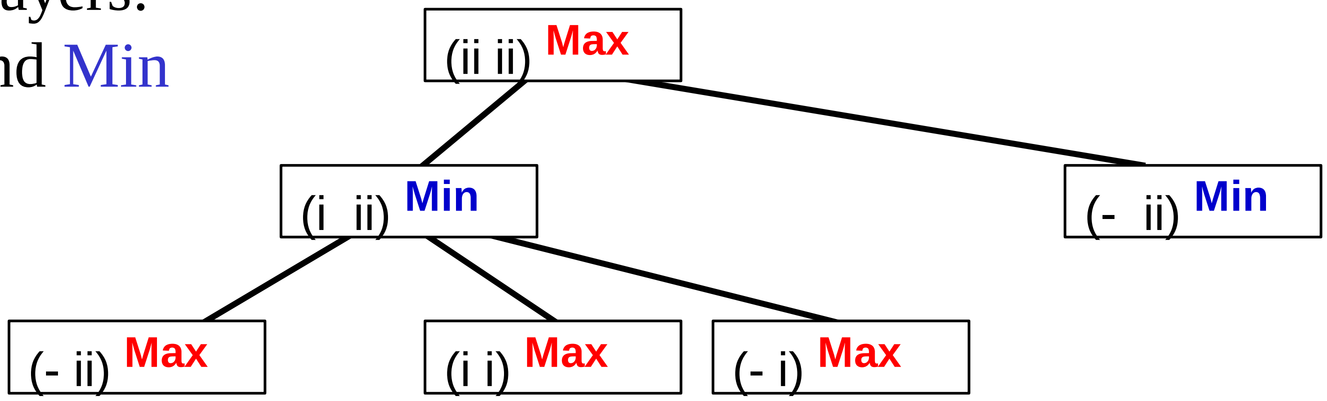

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

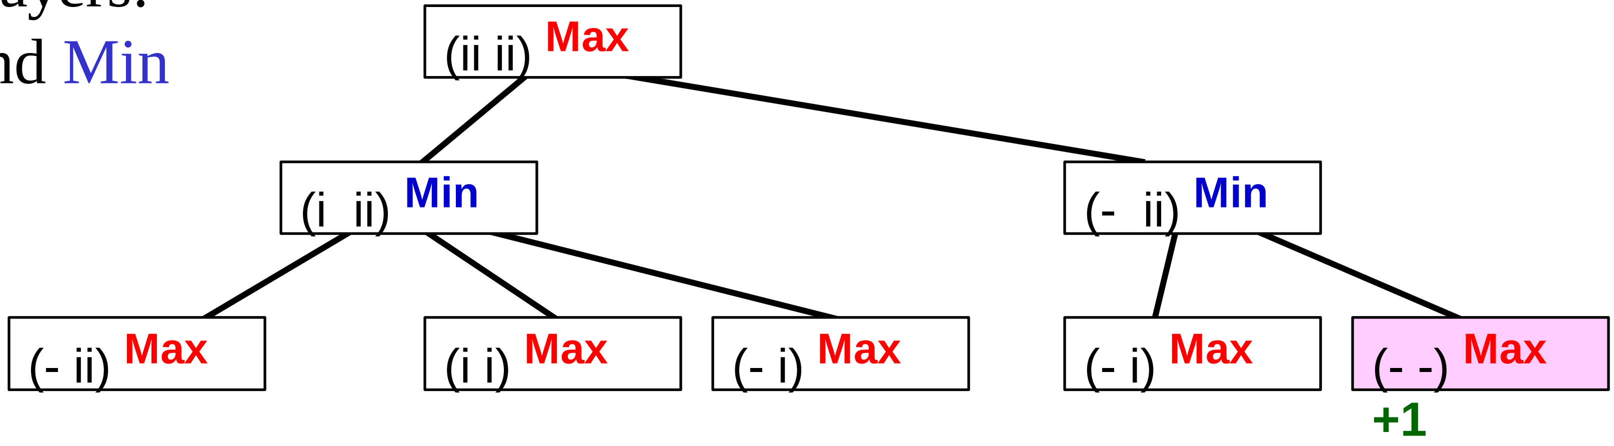

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

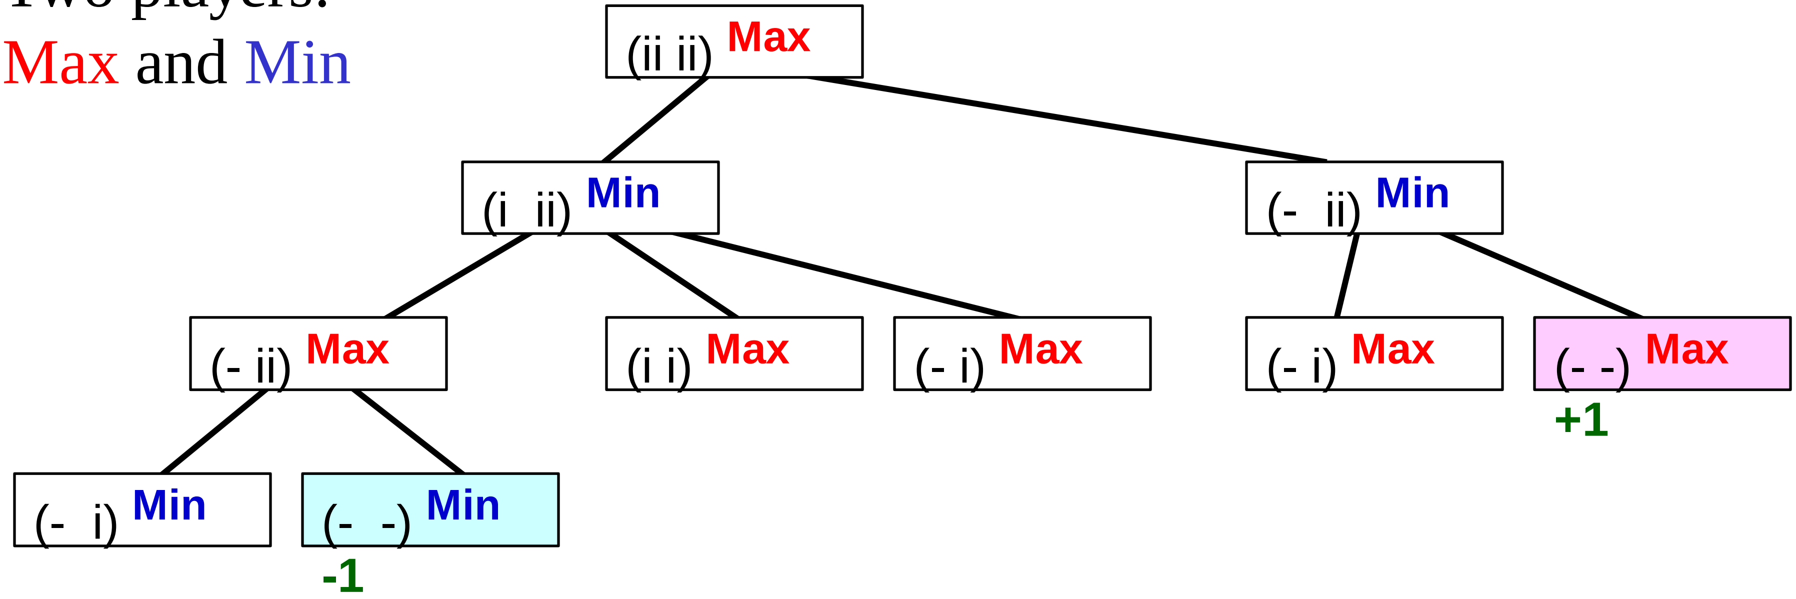

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

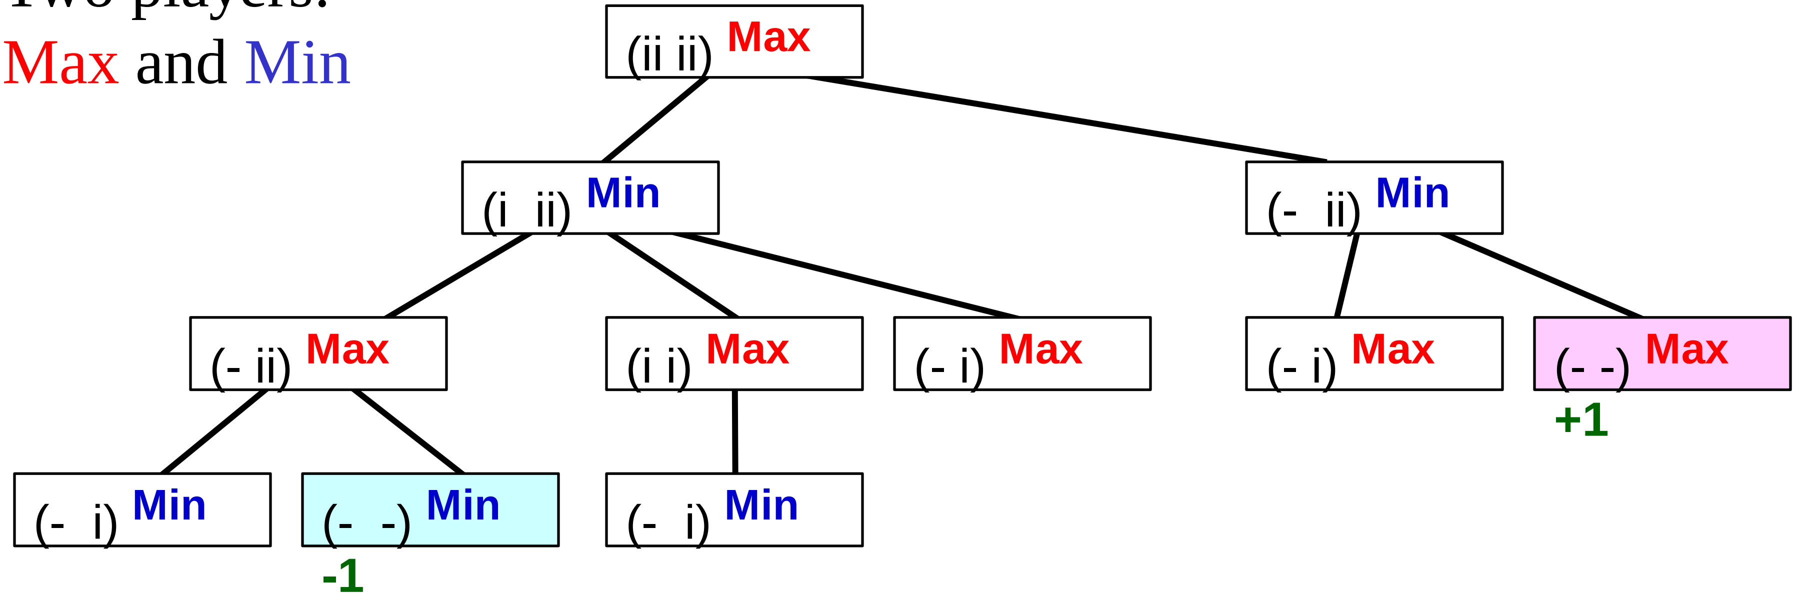

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

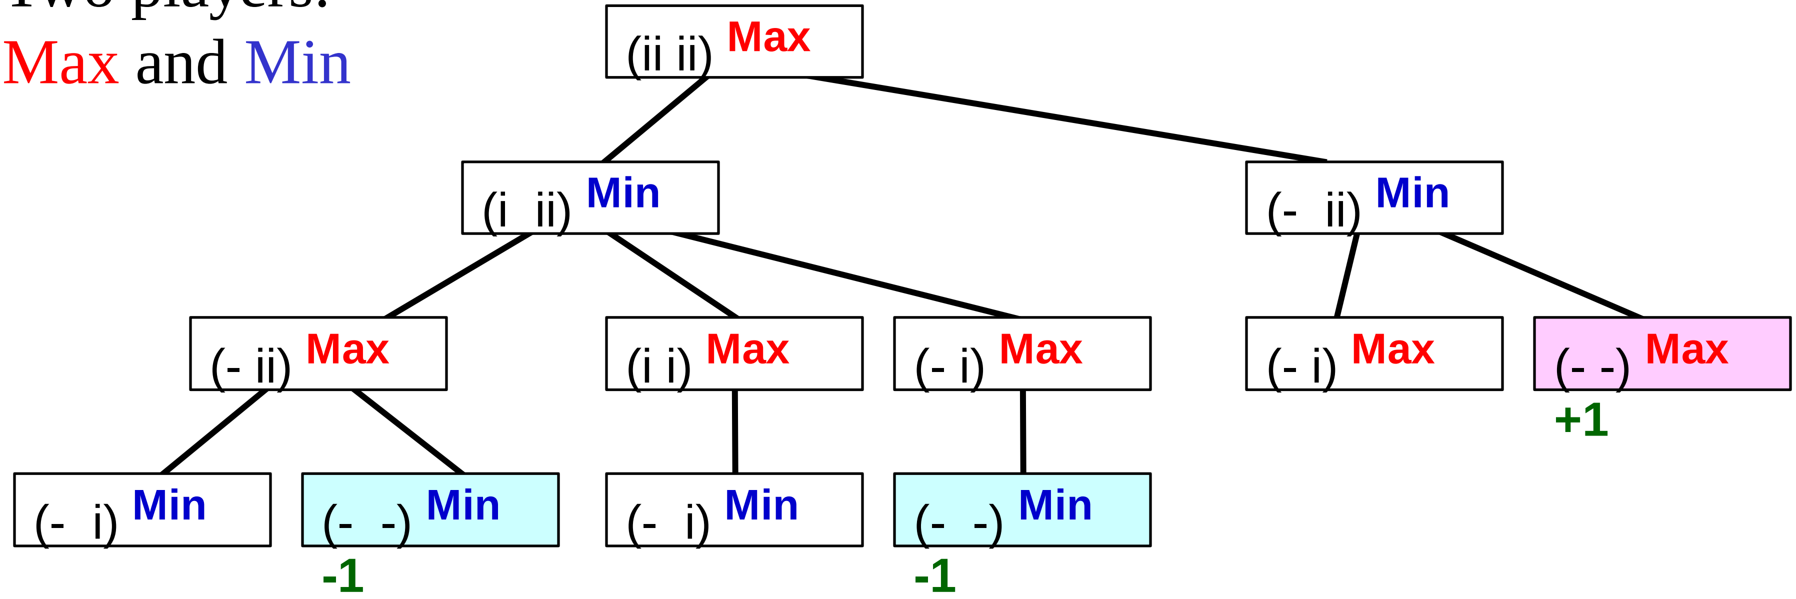

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

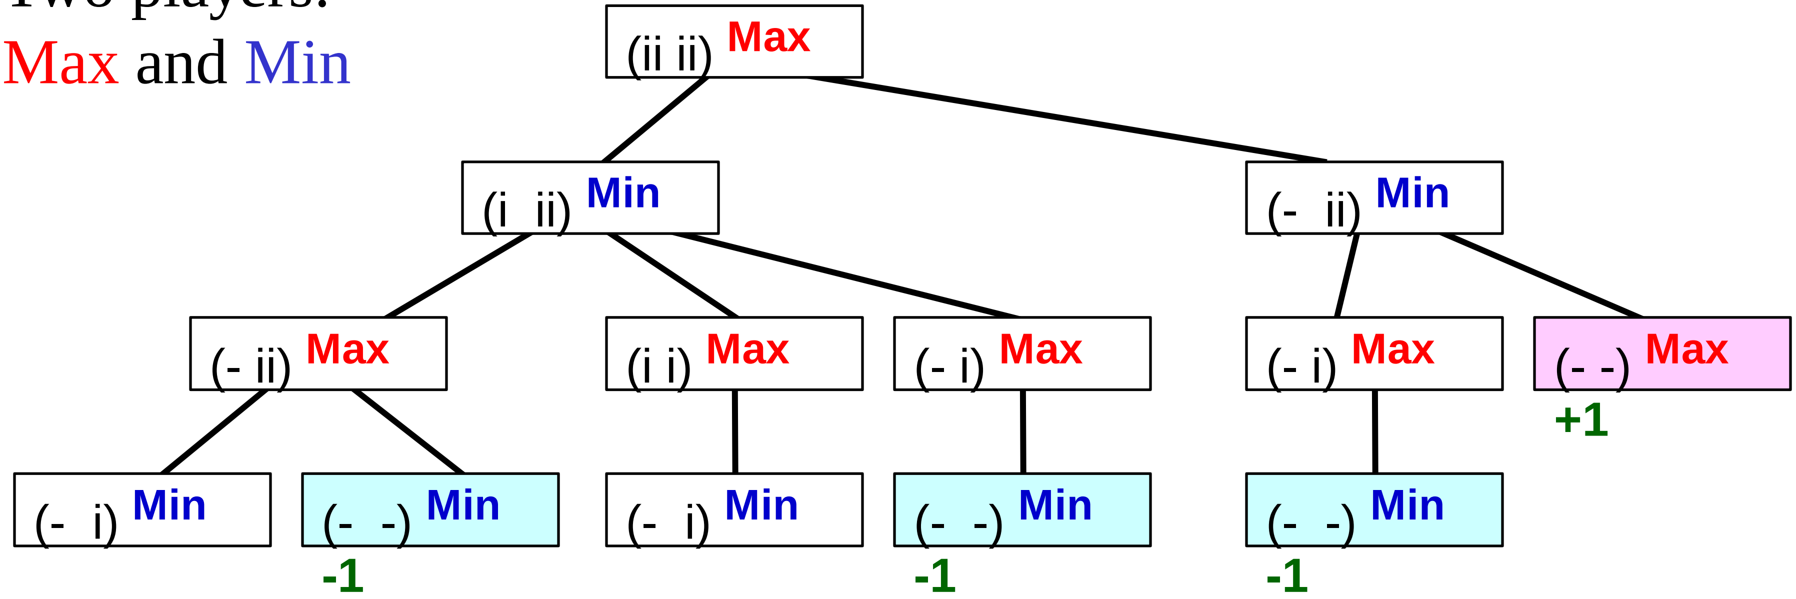

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

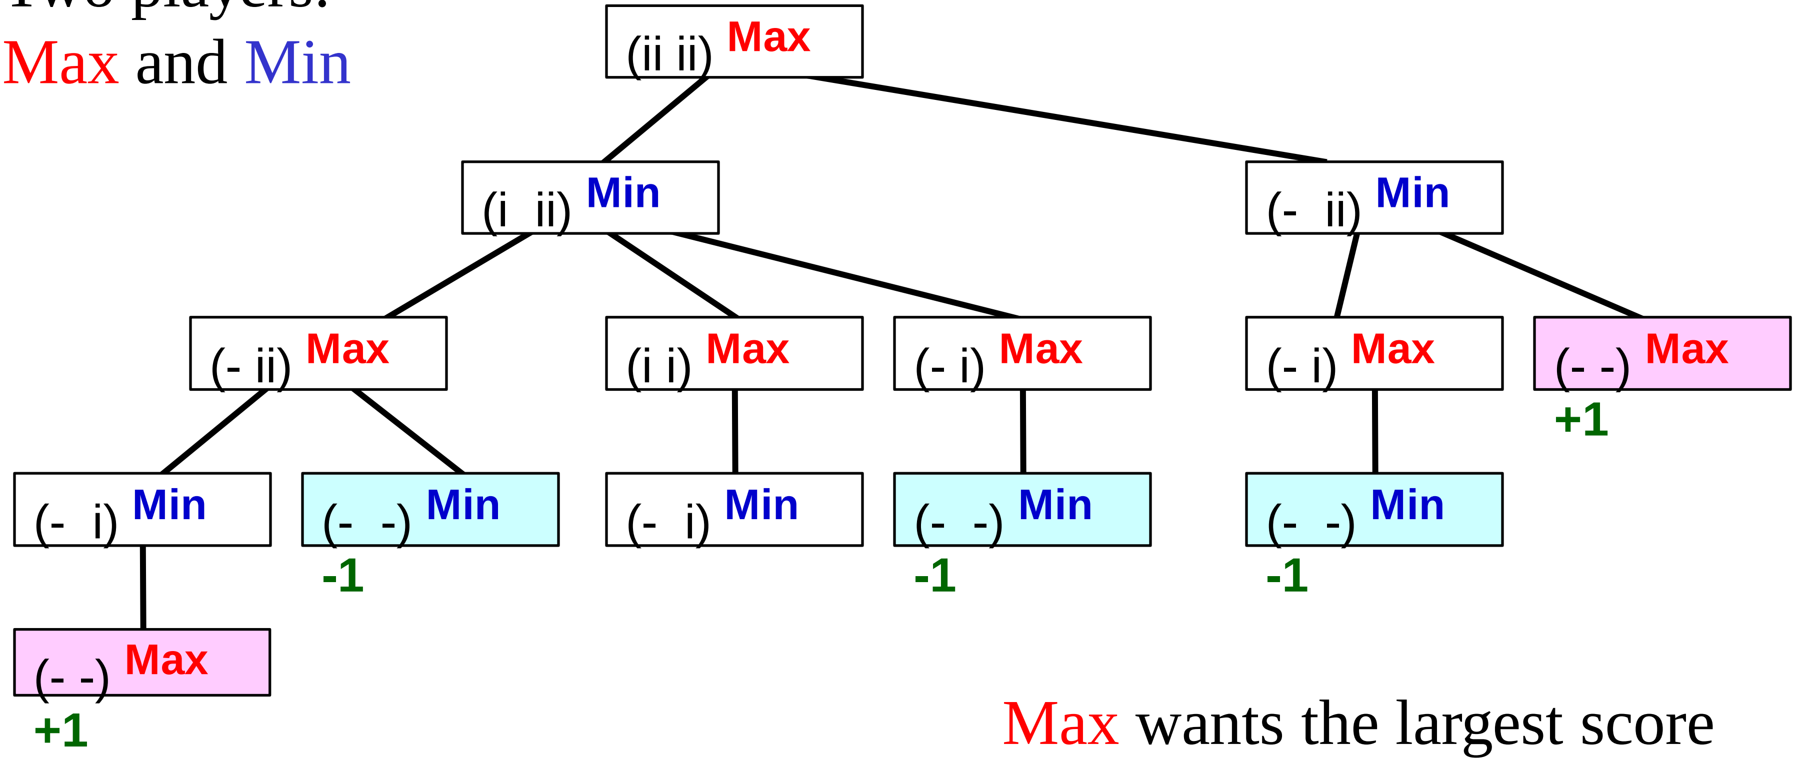

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

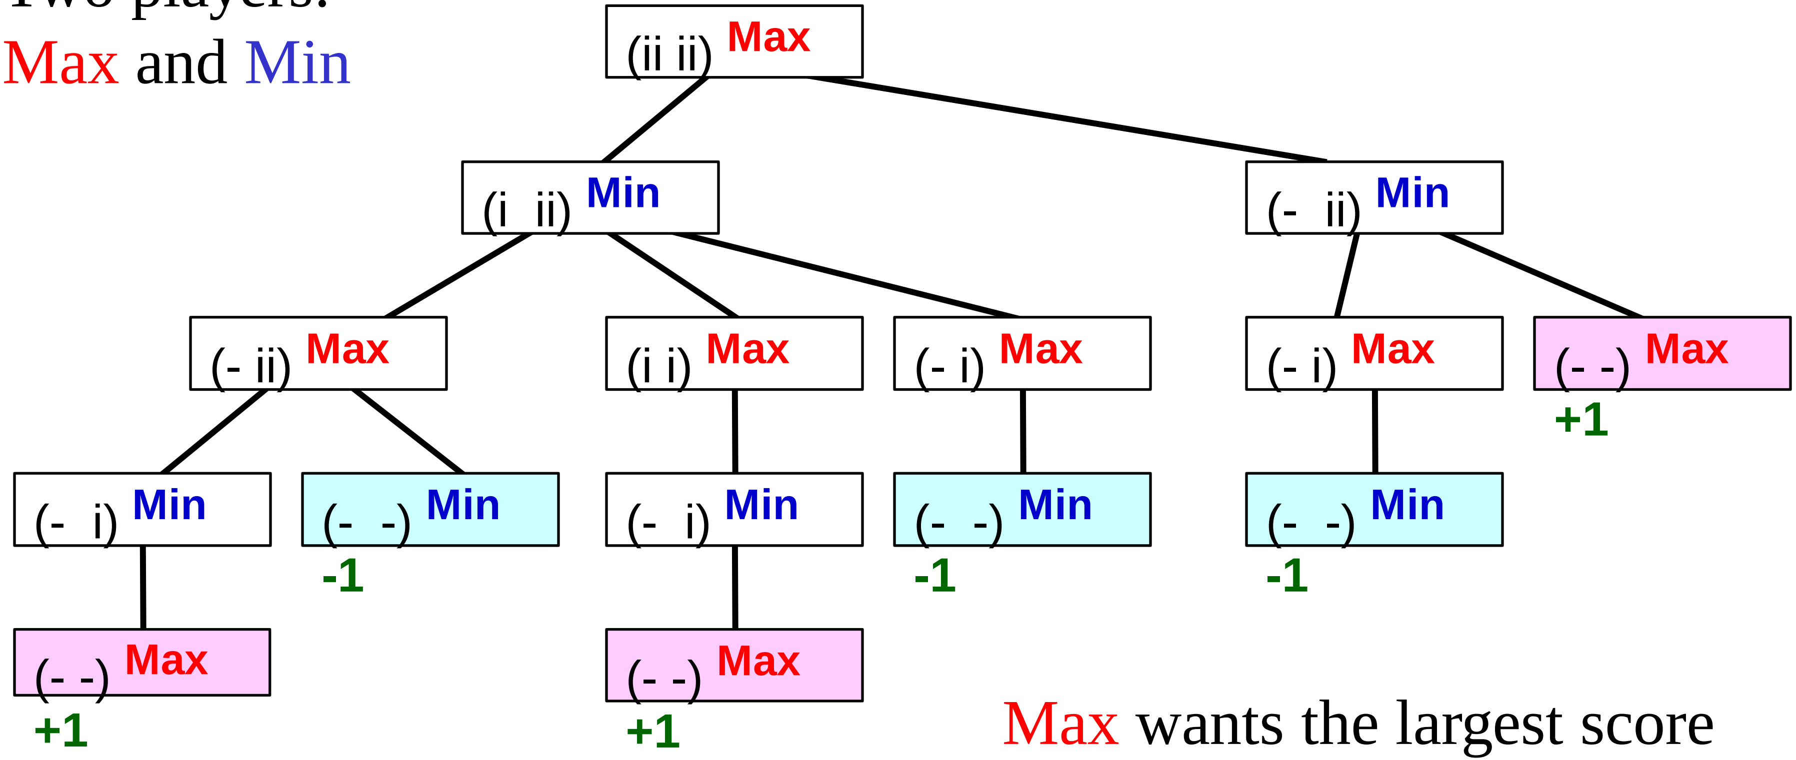

Max wants the largest score
Min wants the smallest score

Minimax Value

Also called **game-theoretic value**.

- Score of terminal node if both players play optimally.
- Computed bottom up; basically search

- Let’s see this for example game

Game tree for II-Nim

Two players:
Max and Min

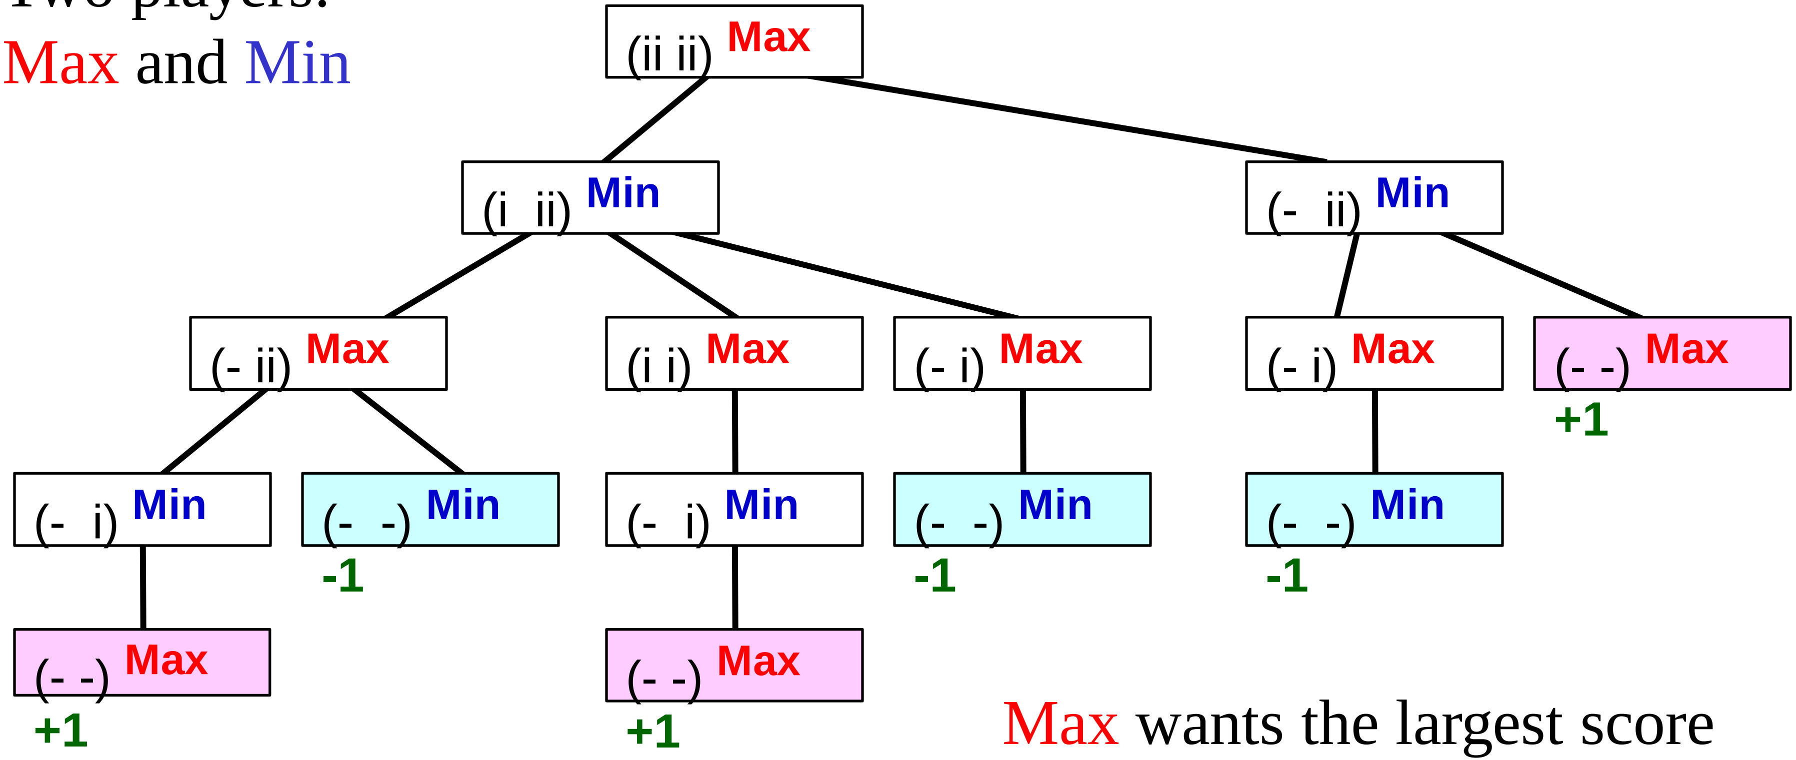

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

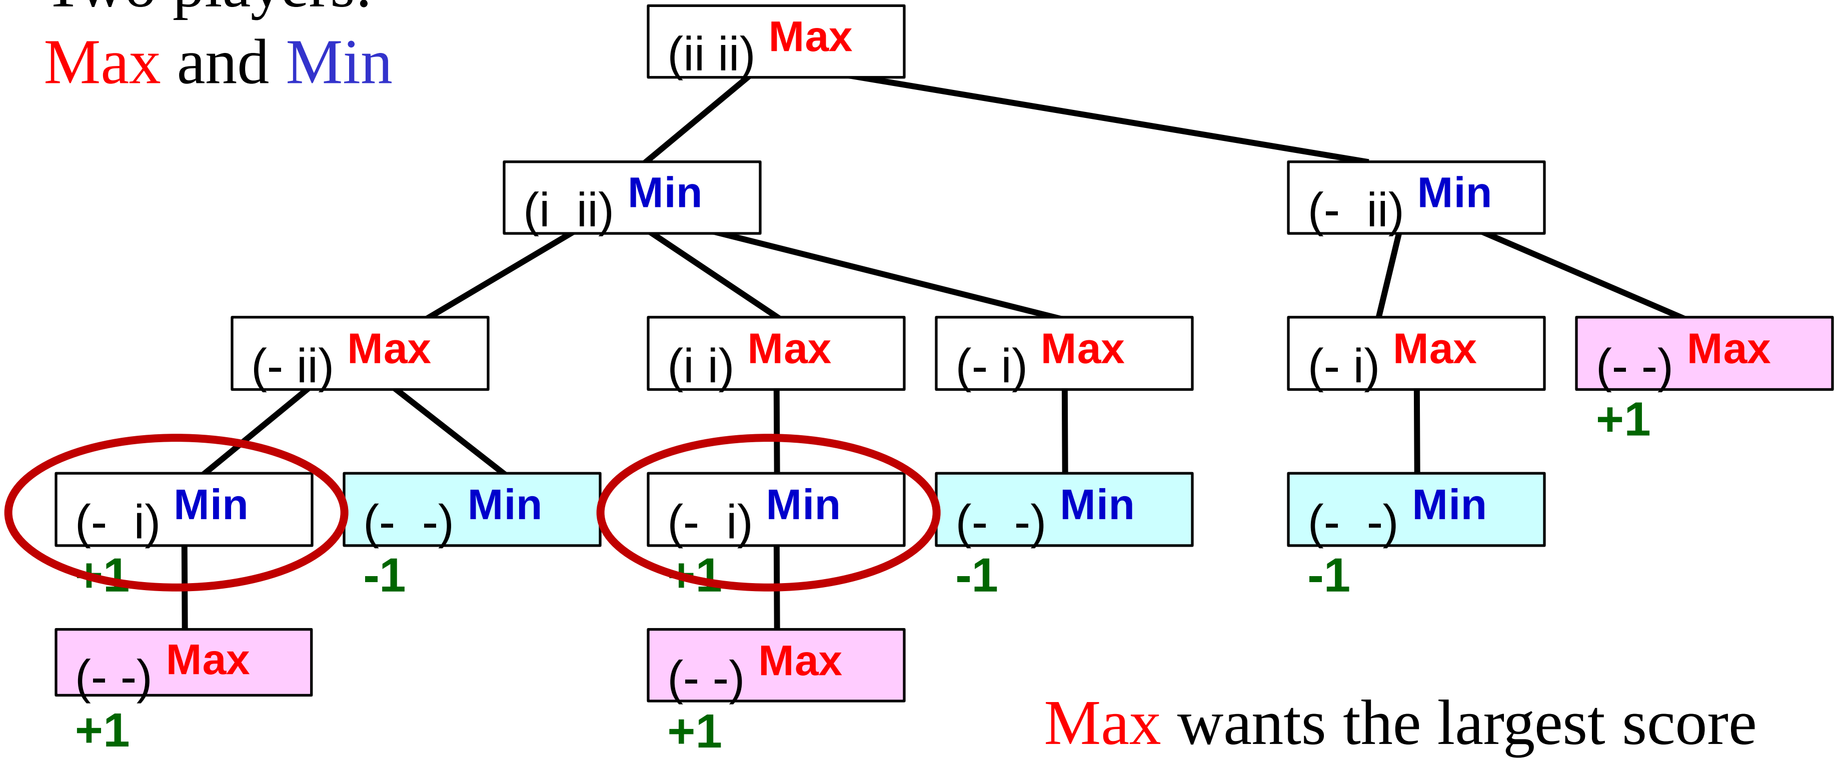

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

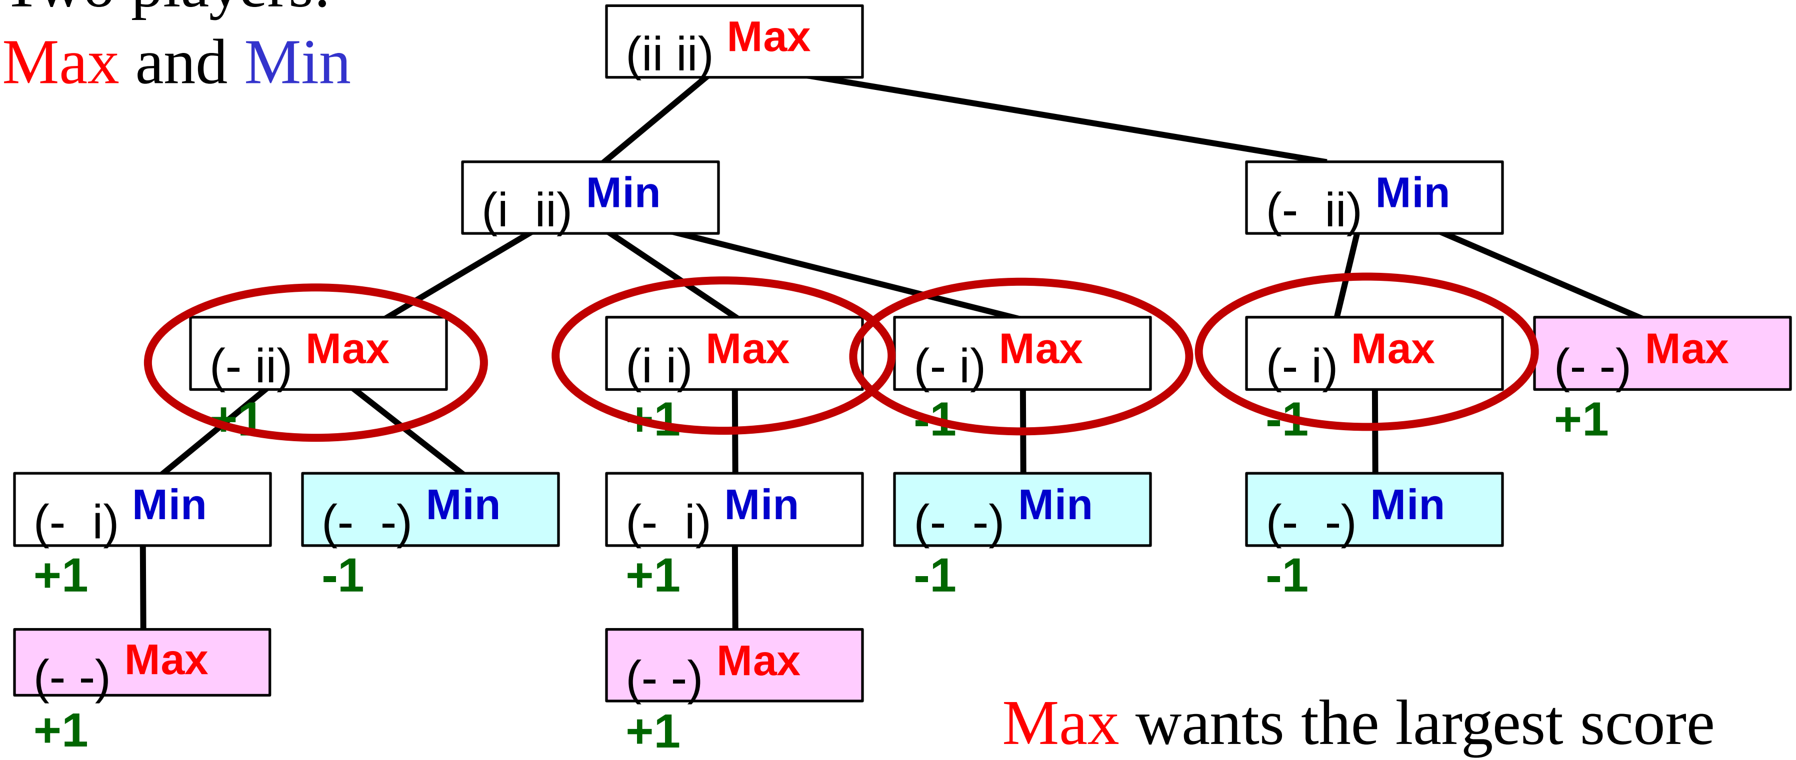

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

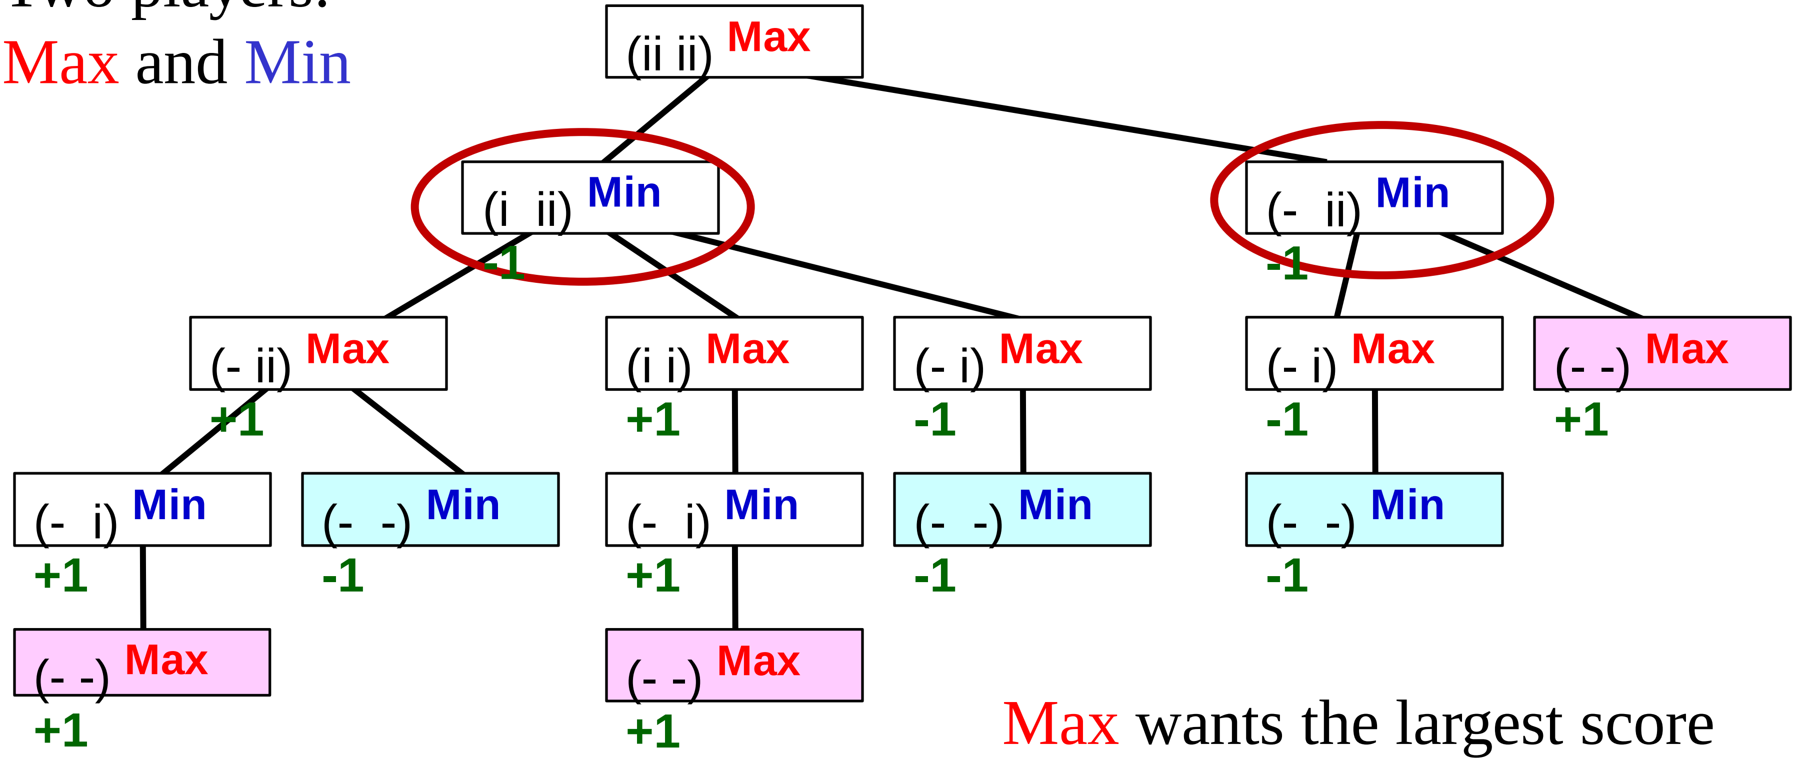

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

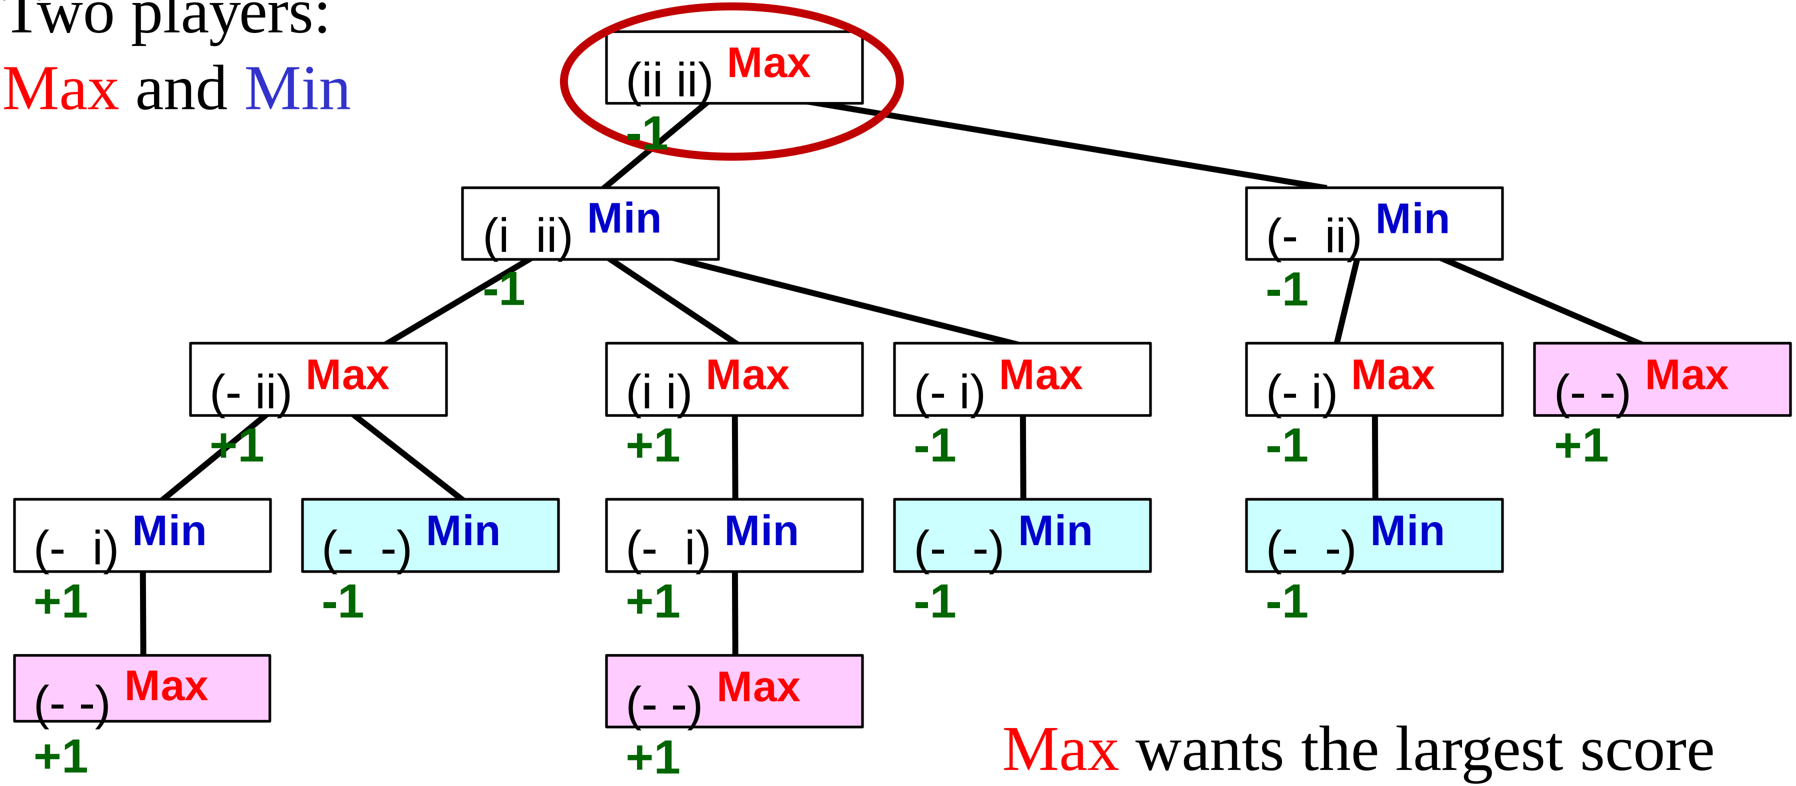

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

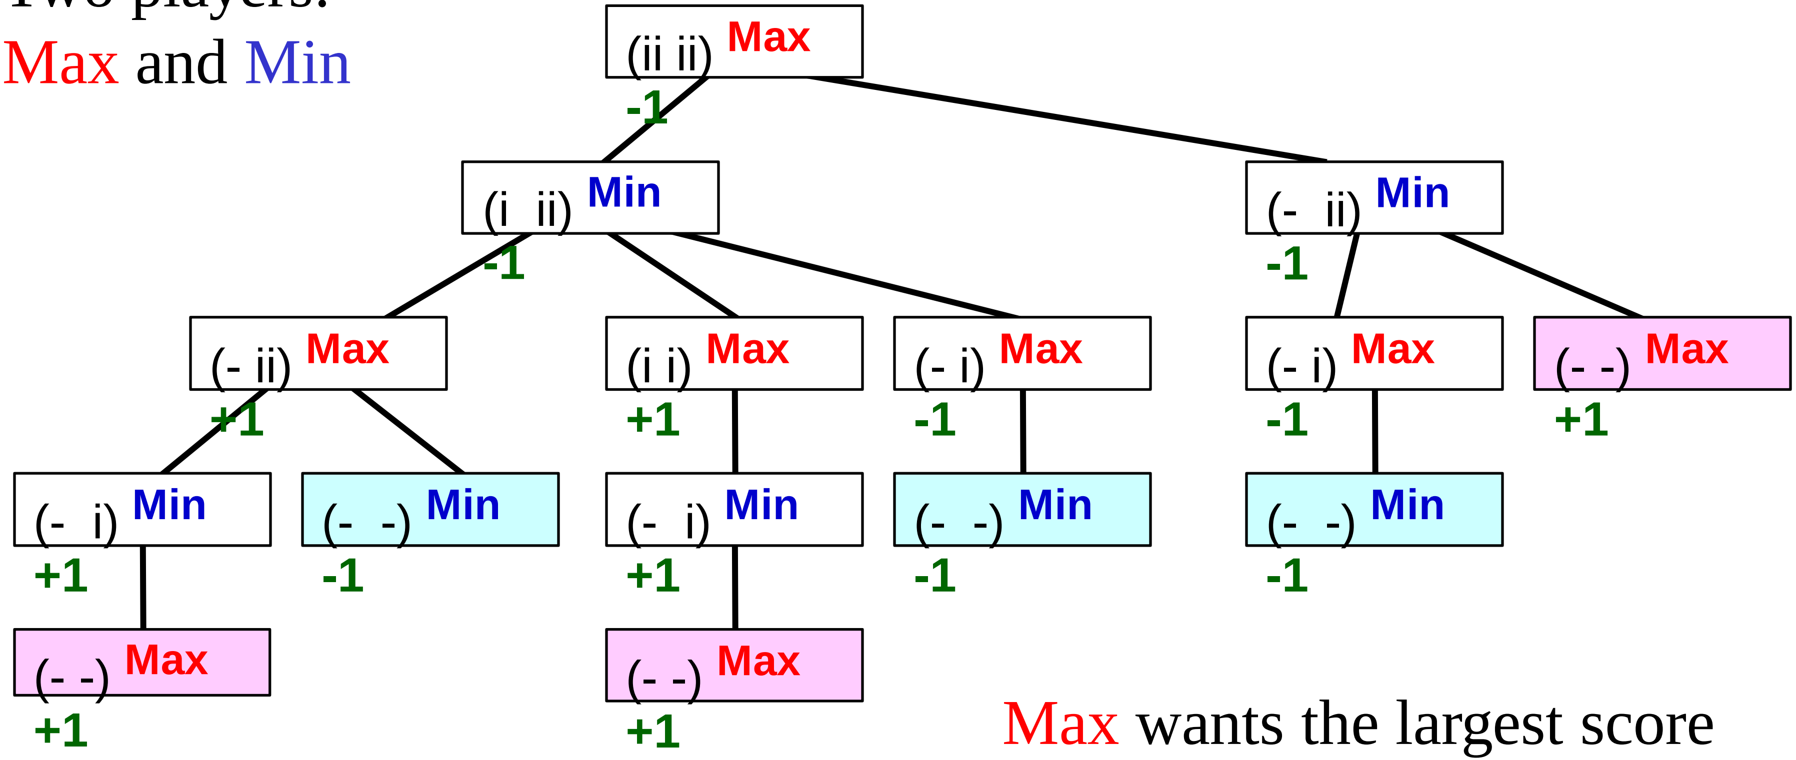

Max wants the largest score
Min wants the smallest score

Game tree for II-Nim

Two players:
Max and Min

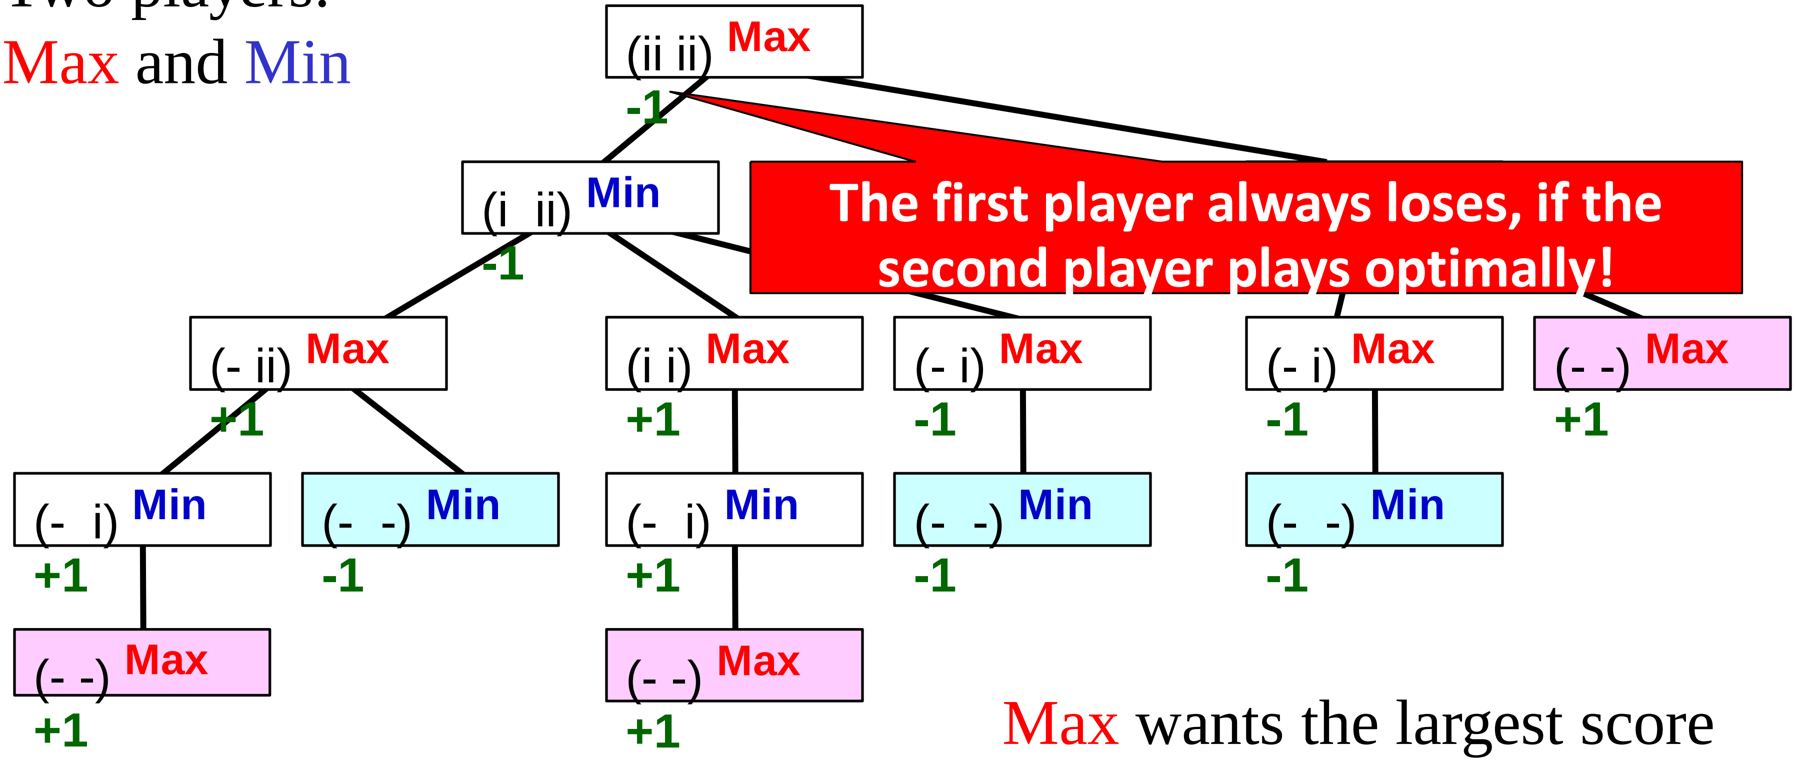

Max wants the largest score
Min wants the smallest score

Summary

- Intro to game theory
  - Characterize games by various properties
- Mathematical formulation for simultaneous games
  - Normal form, dominance, equilibria, mixed vs pure
- Sequential games
  - Game trees, game-theoretic/minimax value

Acknowledgements: Developed from materials by Fred Sala and Yingyu Liang (University of Wisconsin), inspired by Haifeng Xu (UVA).