



AlgoLabs

<http://www.algolabs.org.in>

Mean-Payoff games:

7/9/2017

Lecture Notes: I

Organization:

Part 1: Definition of the game, objectives of the Players

Part 2: Illustrative example

Part 3: Overview of results

Part 4: First-cycle version of the game

Part 5: Equivalence of values between finite and infinite versions -

Part 1: Mean-Payoff game definition

Arena: $G = (V_0, V_1, E)$

$w: E \rightarrow \{-W, \dots, +W\}$ (weight function)
integers

Two Players Maximizer and Minimizer

Initial vertex $a_0 \in V$

Play: Sequence $a_0 a_1 a_2 \dots$
of vertices, edges

$\pi: a_0 \xrightarrow{w_0} a_1 \xrightarrow{w_1} a_2 \xrightarrow{w_2} \dots$

Play continues ~~for~~ indefinitely

- Suppose Play stops at 'n' steps:

$$a_0 \xrightarrow{w_0} a_1 \xrightarrow{w_1} a_2 \rightarrow \dots \rightarrow a_{n-1} \xrightarrow{w_{n-1}} a_n$$

(Average) Mean-~~Pay~~ weight = $\frac{\sum_{i=0}^{n-1} w_i}{n}$

Minimizer pays $\frac{\sum_{i=0}^{n-1} w_i}{n}$ to Maximizer

- In the infinite ~~pay~~ game, we need a definition of the average of the play (play).

Mean-Payoff (π) $\stackrel{?}{:=} \lim_{n \rightarrow \infty} \frac{\sum_{i=0}^{n-1} w_i}{n}$ [Can we define it this way?]

↓
play

This seems like a natural definition ~~at~~ for the infinity case, but unfortunately the above limit might not exist always.

Example: (credits to Soumyajit).

w _i sequence:	2	-2	2	2	-2	-2	2	2	2	2	2	...
		↓		↓		↓				↓		...
$\frac{\sum w_i}{n}$		0		1		0				1		...

$\lim_{n \rightarrow \infty} \frac{\sum w_i}{n}$ oscillates between 0 and 1.

- So we consider $\lim_{n \rightarrow \infty} \inf \frac{\sum w_i}{n}$

- How is this defined?

$$\lim_{n \rightarrow \infty} \inf x_n = \lim_{n \rightarrow \infty} \left(\inf_{m \geq n} \{x_m\} \right)$$

In the above example $\lim_{n \rightarrow \infty} \inf \left(\frac{\sum w_i}{n} \right) = 0$

Why should $\lim \inf$ always exist in our case?

- Note that each $\frac{\sum w_i}{n} \geq -W$ and $\sum w_i$

- Sequence $\left\{ \inf_{m \geq n} \{x_m, x_{m+1}, \dots\} \right\}$ is increasing.

- sequence is ~~non~~ increasing and bounded. So its limit should exist.



AlgoLabs

http://www.algolabs.org.in

Objective in Mean Payoff:

Maximizer wants to maximize $\liminf_{n \rightarrow \infty} \frac{\sum w_i}{n}$ of a play.

Minimizer wants to minimize it.

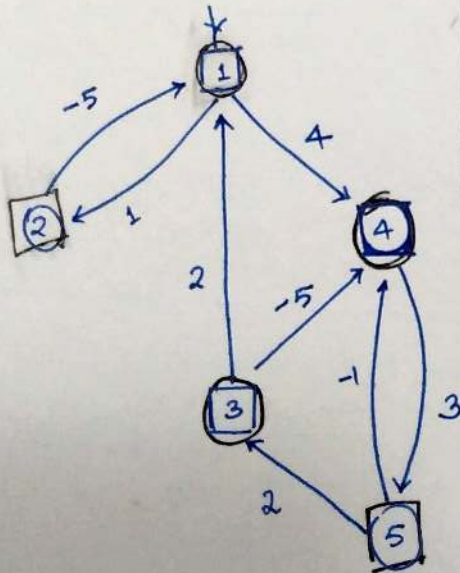
We can think of it as: Min paying the "Mean Payoff" to Max.

Strategies: σ for Max } convention.
 τ for Min }

$\sigma: V^* V_1 \mapsto V$ } ~~not necessary~~ positional if $\sigma: V \rightarrow V$
 $\tau: V^* V_0 \mapsto V$ } $\tau: V_0 \mapsto V$

Objective for Min is to play a strategy that minimizes payoff.
 Objective for Max is to play a strategy that maximizes payoff.

Part 2: Example (from Raskin's tutorial, with notations changed).



Play: $1 \xrightarrow{1} 2 \xrightarrow{-5} 1 \xrightarrow{1} 2 \xrightarrow{-5} \dots$

$\frac{\sum w_i}{n}$: $\frac{1}{1} \xrightarrow{-2} \frac{-1}{2} \xrightarrow{-3} \frac{-4}{3} \xrightarrow{-4} \frac{-8}{4} \xrightarrow{-5} \frac{-12}{5} \dots$

$\frac{1}{1}$	$\frac{-4}{2}$	$\frac{-3}{3}$	$\frac{-8}{4}$	$\frac{-7}{5}$	$\frac{-12}{6}$...
1	-2	-1	-2	$-\frac{7}{5}$	-2	...

$\frac{-4n+1}{2n+1} = \frac{1-4n}{2n+1} \quad n: 0 \text{ to } \infty$
 $\liminf_{n \rightarrow \infty} = -2$
 $\Rightarrow = \frac{(-4n-2) + 3}{2n+1}$
 $= -2 + \frac{3}{2n+1}$

Ⓢ: Maximizer



Ⓜ: Minimizer.

Consistent with

Raskin's notes + Jurdzinski's paper on w-n-up




Another play:

$$\pi_2: 1 (45)^\omega$$



$$\frac{4 + 2n}{2n + 1}$$

$n=2$
 $2n-2=2$
 $2 \cdot 2 = 4$




$$4 + 2(2n-2) + 3 = \frac{4 + 2n + 1}{2n}$$

$$\frac{4}{2n+1} + \frac{2n}{2n+1}$$

$$\frac{4}{2n} + \frac{(2n+1)}{2n}$$

First look at:




lim int $\frac{4}{2n+1} = 1$
 $n \rightarrow \infty$

$$\lim_{n \rightarrow \infty} \frac{(4+2n)}{2n+1} = 1$$

$$\lim_{n \rightarrow \infty} \frac{(4+2n+1)}{2n} = 1$$

What happens by adding 4 in front?



$\frac{4 + 2n}{2n + 1} \rightarrow > 1$ at these points the numbers are > 1



AlgoLabs

<http://www.algolabs.org.in>

So the payoff of $1(45)^\omega = 1$

→ so would be the payoff of $(45)^\omega$.

→ ~~even~~ intuitively, the finite prefix does not play a role in the limit average payoff.

Exercise: Suppose π is a path play. ^{show that:} Removing a finite prefix does not change the payoff. Similarly show that adding a finite prefix does not change the payoff. [in the sequences, the differences will be due to a constant factor]. → Use fact that: 1) if $\{a_n\}_{n \geq 0}$, $\{b_n\}_{n \geq 0}$ are 2 bdd. real sequences

and $\{b_n\}_{n \geq 0}$ converges to b

then: $\liminf (a_n b_n) = b \liminf (a_n)$.

2) $\liminf (a_n + b_n) = \liminf (a_n) + \liminf (b_n)$

Another play:

$$\pi_3: (1 \ 4 \ 5 \ 3)^\omega$$

$$= \frac{11}{4} \rightarrow \text{mean payoff.}$$

How should Maximizer play? From 1 if she ^{always} plays 2, the payoff

that she can get is -2. ~~So~~ If she plays 4 always, the

payoffs can be 1 or $\frac{11}{4}$ provided Min plays positionally.

→ It turns

How should Minimizer play? At 5, she should choose to go to 4.

Currently, we have only described positional strategies and we have seen what payoffs you can get from such positional strategies.

Will history-dependent strategies help? Can any of them play better?
 with such strategies

Part 3: Overview of results.

Given

Theorem: (G, w, a_0) is a mean-payoff game. There exists a

number $\tilde{v}(a_0)$ and positional strategies $\tilde{\sigma}, \tilde{\tau}$ s.t.

Payoff $(\tilde{\sigma}, \tau) \geq \tilde{v}(a_0)$ for every strategy τ of Min.

Payoff $(\tilde{\sigma}, \tilde{\tau}) \leq \tilde{v}(a_0)$ for every strategy τ of Max.

The value $\tilde{v}(a_0)$ is called the value of the game.

→ also called Minimax equilibrium.

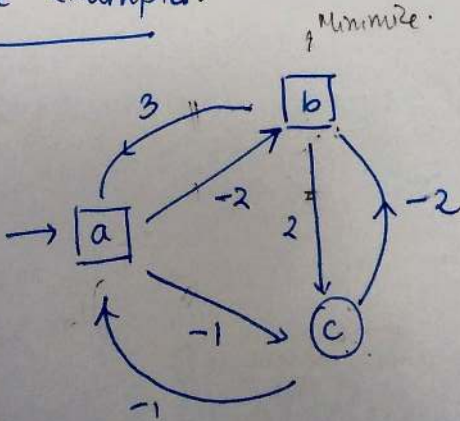
Illustration on example:

Max: $1 \rightarrow 4$ she can ensure ≥ 1
 $3 \rightarrow 1$ no matter however Min plays.

Min: $5 \rightarrow -1$ she can ensure ≤ 1
 no matter how Max plays.

Value 1 is the value of the game.

More examples:



What is the value of this game?

- $(ac)^w = 0$
- $a(cb)^w = 0$
- ~~$a(c)$~~ $(ac)^w = -1/2$
- $(acb)^w = 0$
- $(abc)^w = (-1/3)$
- $(cb)^w = 1/2$
- $a(bc)^w = 0$

- definition of game
- computing values manually
- using result abt positional
- limit of sequences

$\frac{1}{2}$



AlgoLabs


<http://www.algolabs.org.in>

Part 4: First cycle version of the game

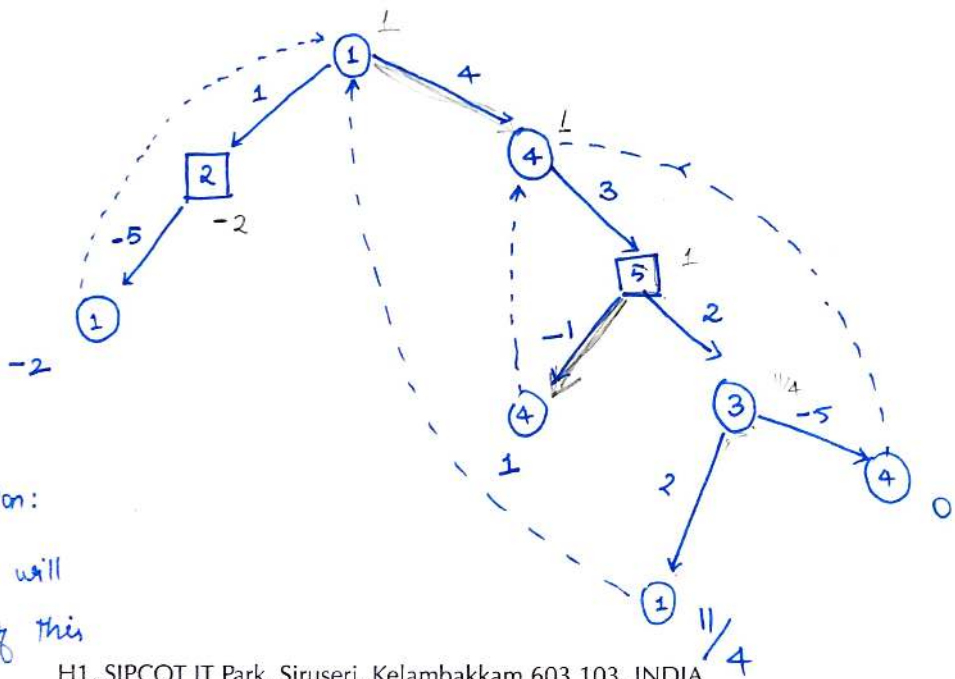
In the previous two examples, the value that we computed turns out to be the mean weight of some cycle. Is this a coincidence or does it always happen? (if the theorem about positional strategies is true, this should always happen).

Motivated by this observation, we investigate a finite version of mean payoff games called the first-cycle version.

- Given (G, w, a_0) . Play starts at a_0 and continues as before. However the play stops as soon as a vertex which was seen previously is encountered again.



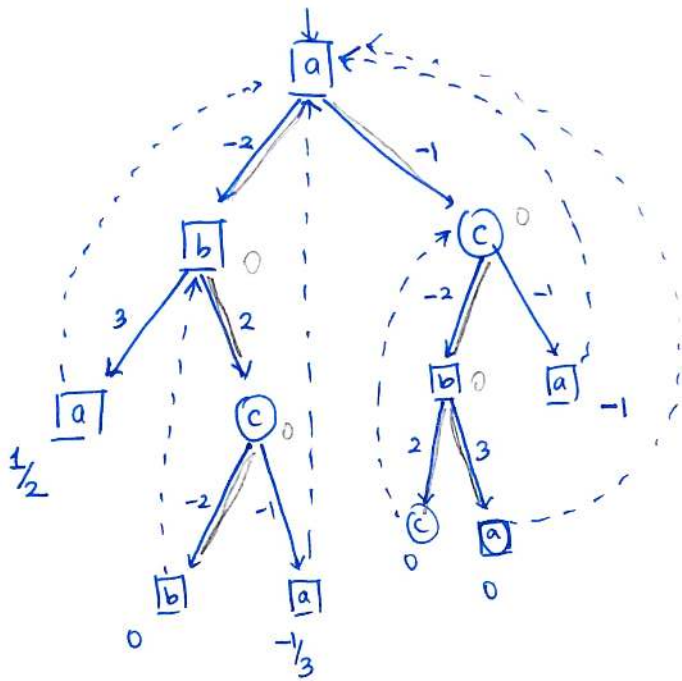
Payoff = Mean-weight of the cycle formed.



Idea of backward min max computation:

- Marked strategies will give the value of this finite game.

⑧



Observations: Different occurrences of a node could mean different possible values from each occurrence. The mean weight of the cycle formed depends on the history. For instance, for ⑧, on the left there is a possibility to go to $-\frac{1}{3}$; and on the right, there is a possibility for -1 .

→ The min-max computation starting from leaves all the way up to root will give a unique value. * The arg min / arg max gives strategies τ and σ for Minimizer and Maximizer that achieve this value.

→ It is not immediately clear that ~~there~~ there exist positional strategies σ, τ for this game.

→ It is also not still clear why the value computed this way should be same as that for the infinite version game.



AlgoLabs

<http://www.algolabs.org.in>

Summary so far:

- 1. Mean-payoff games (infinite duration)
Payoffs objectives, Result about existence of a value and positional strategies achieving this value.
- 2. First-cycle ~~Mean-payoff~~ mean payoff games
Algorithm to compute value and (not necessarily positional) strategies.

Coming next:

- 1. Value computed in FC game = value in Mean Payoff game
- 2. There exist positional strategies σ_{a_0}, τ_{a_0} which can achieve this value (in both the games).

Part 5: Equivalence between first cycle game and mean payoff game.

$F \rightarrow$ First cycle game starting at a_0

$G \rightarrow$ Infinite duration mean payoff game starting at a_0 .

$\tilde{v}(a_0) \rightarrow$ value in infinite game (existence not proved yet)

$v(a_0) \rightarrow$ value in finite game (algorithm to compute $v(a_0)$
strategies σ, τ securing $v(a_0)$)

Pick the strategy σ for Max in F (not necessarily positional).

- We will define strategy $\tilde{\sigma}$ as follows. It uses some finite memory (in the form of a bounded stack).

(10)

σ is a strategy defined on "simple paths".

- Game starts at a_0 . Start playing using σ

Initial memory $m_0 = a_0$.


At the $(k+1)^{th}$ step, play $a_{k+1} = \sigma(m_k)$ if it is Max's turn.

Update m_{k+1} as follows.

- Suppose $m_k = u_0 u_1 u_2 \dots u_r$. $r \geq 0, r \leq n$

$$m_{k+1} = \begin{cases} u_0 u_1 \dots u_s & \text{if } v_{k+1} = u_s \text{ for some } s \leq r \\ u_0 u_1 \dots u_r v_{k+1} & \text{otherwise.} \end{cases}$$

Idea:



Think of this as a stack, each time there is a repetition pop ^{the} contents cycle, except the root.

Example: ~~σ~~ (first running example). $\sigma [1 \rightarrow 4, 14 \rightarrow 5, 1453 \rightarrow 1]$

$\tilde{\sigma} [1 \rightarrow 4]$

~~1 4 5 4 5 3 1 4~~

$\tilde{\sigma} [14 \rightarrow 5]$

$\tilde{\sigma}$ uses a ^{bounded} stack as a memory

$\tilde{\sigma} [1454 \rightarrow 5]$

(finite memory strategy).

$\tilde{\sigma} [145453 \rightarrow 1]$

⋮

Claim: Playing $\tilde{\sigma}_n$ in G ensures that the value obtained for any play according to this strategy $\geq v(a_0)$.



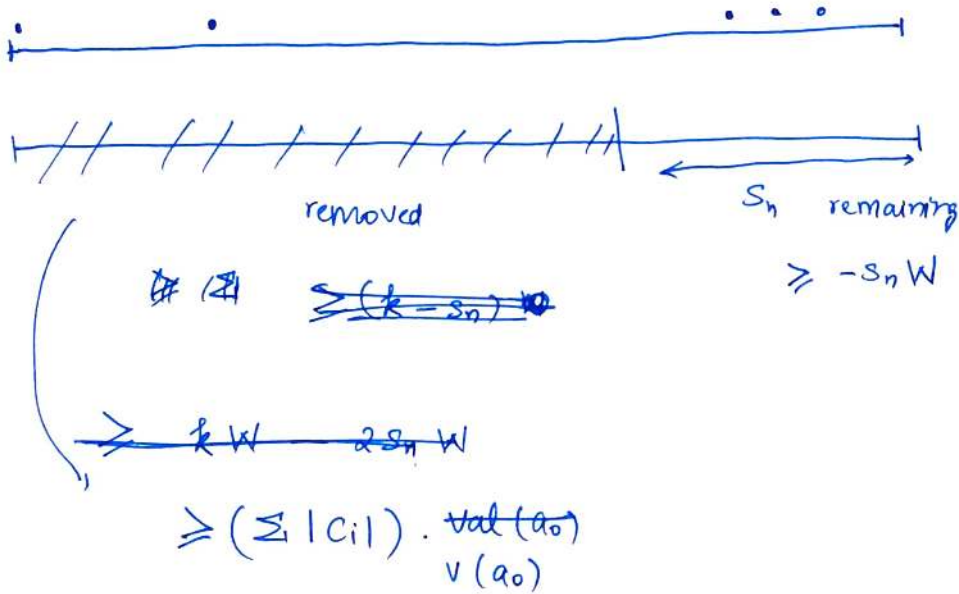
AlgoLabs

http://www.algolabs.org.in

Proof: Consider the mean payoff game played upto k steps, according to $\tilde{\sigma}$.

What is $\sum_{i=0}^{k-1} \frac{w_i}{k}$?

After k steps In this process of playing k steps, say C_1, C_2, \dots, C_m were the cycles removed. Let there be s_k elements left in the memory.



Sum of wts. $\geq (k - s_n) v(a_0) - s_n W$

$\geq (k - n) v(a_0) - \frac{n}{k} W$

$s_n \leq n$
 $-s_n \geq -n$

Mean $\frac{\sum w_i}{k} \geq v(a_0) - \frac{n}{k} \cdot v(a_0) - \frac{n W}{k}$

$\geq v(a_0) - \frac{2n W}{k}$

$v(a_0) \geq -W$

(12)

$$\liminf_{k \rightarrow \infty} \frac{\sum w_i}{k} \geq \liminf_{k \rightarrow \infty} \left(v(a_0) - \frac{2nW}{k} \right) \geq v_0(a_0)$$

This proves that using $\tilde{\sigma}$ gives value at least $v(a_0)$.

- Using similar arguments, we can construct $\tilde{\tau}$ which ensures that the value of any play has payoff $\leq v(a_0)$.
- This gives us two strategies $\tilde{\sigma}, \tilde{\tau}$ and a value $\tilde{v}(a_0) = v(a_0)$ s.t. $\tilde{v}(a_0)$ is a minimax equilibrium.
- We also know how to get $\tilde{\sigma}, \tilde{\tau}$ using algorithm to compute σ and τ .
- However we still have not obtained positional strategies for any of these games. To do this, we need to go back and forth between the finite and infinite versions. This is a ~~new~~ ^{novel} technique. ~~to prove such str~~