11/9/2017

Mean-Payoff games

Lecture Notes - II

AlgoLabs

http://www.algolabs.org.in

Organization:

1. Recall previous results
2. Understanding finite tree games; ~~per~~ effects of perturbing the values
3. Intermediate game; proof of equivalence of values with std. mean-payoff
4. Proof of positional determinacy
5. Computing values of all vertices uniformly
6. Uniform positional strategies.

Part 1: Recall previous results

Mean-Payoff game (infinite) :→ G

First-cycle mean-payoff :→ F

Both start from specified vertex a_0.

(Infinite) **Result:** There exist numbers $\tilde{v}(a_0)$, $\tilde{\sigma}$, $\tilde{\tau}$ s.t. when ~~Max~~ in G:

$\forall \tau \quad \text{Payoff}(\tilde{\sigma}, \tau) \geq \tilde{v}(a_0)$

$\forall \sigma \quad \text{Payoff}(\sigma, \tilde{\tau}) \leq \tilde{v}(a_0)$

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

(Finite) **Result:** Similar $v(a_0)$, σ, τ for first-cycle version.

Transfer result: $v(a_0) = \tilde{v}(a_0)$. $\tilde{\sigma}$ can be obtained from σ by the stack-based method.

$\tilde{\tau}$ can be obtained from τ ———

To show: ~~σ, τ~~ There $\exists$ exist positional strategies giving the optimal value.

H1, SIPCOT IT Park, Siruseri, Kelambakkam 603 103, INDIA
Phone: +91 44 6748 0900 Email: office@algolabs.org.in

Part 2: Some properties of the finite tree games

1). F

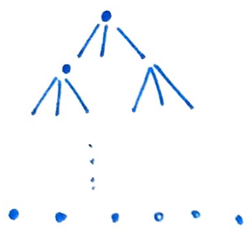

Each leaf has a payoff.
Min-Max computation gives
v, σ, τ

2) Restrict the tree to σ
- Root has value v.
  if it is min, vertex, every child has value $\geq v$.
- if max vertex, then the single edge out of it is given by σ, and value equals max

3) Inductively we can show that all labels in F_σ have $\geq v$.

Similarly for F_τ we can show all values $\leq v$.

Exercise: Suppose some leaves of F have ~~values~~ payoffs changed to v. Will the value of the game change?

Proof: No. Consider same strategies σ, τ. In F_σ, all leaves have $\geq v$ so playing σ ensures $\geq v$. Similarly playing $\tau \leq v$. Value of the game is unique. No other no. can have this property. So, modifying with v does not affect the value.

Part 3: An intermediate game.

(G, a_0) (F, a_0). $(F, a_0, a) \to$ new game.

(F, a_0, a) : ~~pro~~ starts ~~at~~ a_0 and proceeds like F. If a does not come in the play, then it stops when there is a repetition (identical to F). When there is an 'a'. the game F starts from 'a'. Essentially the prefix till now is forgotten and a new first cycle game started from a.

F :

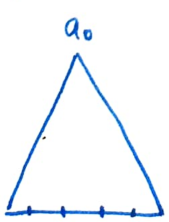

$F(a_0, a)$

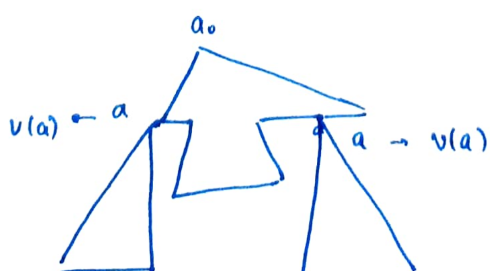

What is value of (F, a_0, a) ?

AlgoLabs
http://www.algolabs.org.in

<u>Claim 1:</u> Let σ be optimal strategy for Max in F.
If 'a' occurs in F_σ, then $v(a) \geq v(a_0)$.

<u>Proof:</u> Pitfall: ~~The~~ The min-max label on 'a' in F_σ need not be $v(a)$. So this is not immediate.

We will use the infinite version for this and show that $\tilde{v}(a) \geq \tilde{v}(a_0)$.

Consider the stack strategy $\tilde{\sigma}(a_0)$, and consider the infinite tree $G_{\tilde{\sigma}}$

We have seen that every play following $\tilde{\sigma}$ has value $\geq \tilde{v}(a_0)$.

- a is a node in this tree.
- a occurs at multiple positions.
- Pick one occurrence of 'a' and look at the infinite tree below it.

- The tree below 'a' gives a strategy for ~~'a'~~ Max for the game starting at 'a'. ~~By pla~~ Call this $\tilde{\sigma}_a$. ~~By playing +~~ since payoffs are not dependent on the prefix, each play starting from 'a' has the same payoff $\geq \tilde{v}(a_0)$. ~~So~~

~~- Similarly, we can look at $G_{\tilde{\tau}}$~~

Therefore, by playing this strategy, Max can ensure at least $\tilde{v}(a_0)$.

$\Rightarrow$ optimal strategy for Max for game starting at 'a' ~~is~~ at least $\tilde{v}(a_0)$.

$$\Rightarrow \quad \tilde{v}(a) \geq \tilde{v}(a_0)$$

$$\Rightarrow \quad v(a) \geq v(a_0)$$

H1, SIPCOT IT Park, Siruseri, Kelambakkam 603 103, INDIA
Phone: +91 44 6748 0900 **Email:** office@algolabs.org.in

Claim 2: Let τ be optimal strategy for Min in F. If 'a' occurs in F_τ, then $v(a) \leq$ ~~$v(a_0)$.~~ $v(a_0)$

Similar proof as before.

Claim 3: Consider game (F, a_0, a). Value of this game equals $v(a_0)$.

Let $\mu(a_0)$ be the value.

Consider optimal strategy σ_1 in F. Suppose a never occurs in the tree F_σ, then play σ will result in a game similar to F. So $\mu(a_0) = v(a_0)$

However, suppose 'a' occurs in F_σ. Then from Claim 1, $v(a) \geq v(a_0)$

~~Hence (from discussion~~ — The tree starting at 'a' will have value $v(a)$.

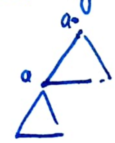

- ~~This~~ From the discussion in Part 2, we get that $\mu(a_0) \geq v(a_0)$.
- Similarly, ~~we~~ when Min plays τ, by Claim 2, we can show that $\mu(a_0) \leq v(a_0)$.

Result: Value of the intermediate game $\mu(a_0) = v(a_0)$.

AlgoLabs

http://www.algolabs.org.in

Part 4: Proof of positional determinacy

Game starts at a_0. We will show positional strategies σ, τ for G and F.

~~Proof~~ First we show positional strategy σ for Max giving optimal value.

Proof is by induction on $|E| - |V_{max}|$. We assume $|E| - |V_{max}| \geq 0$
From every vertex, at least one outgoing edge.

Base case: $|E| - |V_{max}| = 0$

Every Max strategy is positional. So done.

Induction case: Assume $|E| - |V_{max}| = k+1$, and the thm. true for $|E| - |V_{max}| = k$.

Pick a vertex 'a' of Max s.t. there are at least 2 outgoing edges.

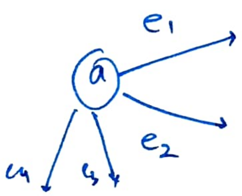

~~G_1 : restrict G to e_1~~

Consider game ~~G_a star~~ (F, a_0, a) starting from a_0.

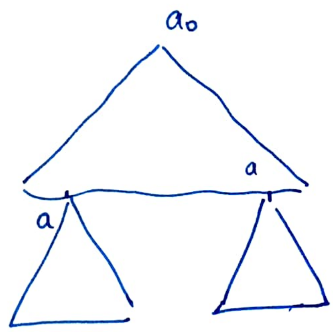

consider optimal strategy σ_a for Max

We can assume that Max plays the same way in all trees starting with 'a'.

Moreover 'a' occurs only once in a non-leaf node. Let 'e_1' be the edge played ~~in~~ from 'a'. Restrict the graph to 'e_1' alone out of 'a'. Value $v(a_0)$ will not change in this graph because $\mu(a_0)$ does not change! By induction hyp. there is a positional strategy for Max for this game. The same strategy works in the bigger game ($\geq v(a_0)$)

H1, SIPCOT IT Park, Siruseri, Kelambakkam 603 103, INDIA
Phone: +91 44 6748 0900 **Email:** office@algolabs.org.in

Similar proof shows positional strategy for Min.

Part 5: Computing values of vertices uniformly for all vertices.

In the previous discussion, we have considered the game starting at a_0. We have shown that there exist positional strategies $\tilde{\sigma}, \tilde{\tau}$ that achieve $\hat{v}(a_0)$.

This gives a map:

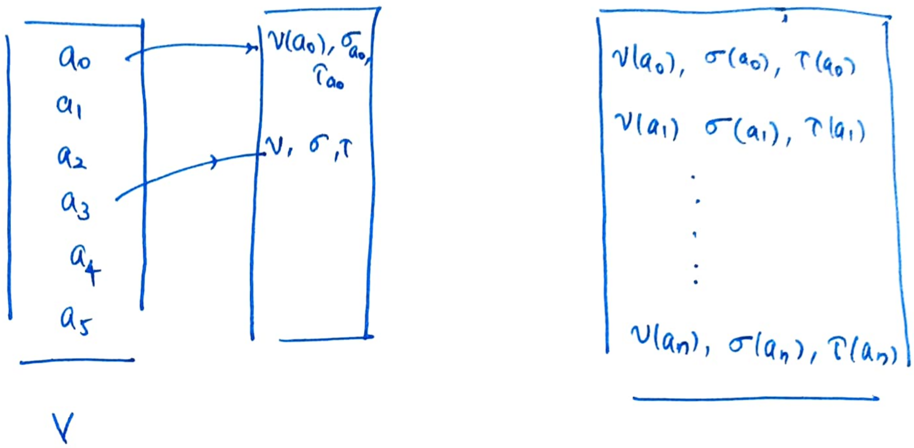

Questions: 1. Can we compute $\langle v(a_0), \dots v(a_n) \rangle$ all together?

2. Do there exist uniform positional strategies σ, τ that work for every vertex, and give the required value?

Uri-Zwick's k-step approximation procedure:

So far we have looked at values for games starting from a specific vertex. To get a global picture, we need to be able to relate value of a vertex to values of its successors.

→ Consider a game starting from 'a' ~~and~~ which goes for k-rounds. The payoff after k-rounds is the sum of weights on the k-edges seen. Min wants to minimize this and Max wants to maximize this.

→ Since this is a finite duration game, values exist by the min-max computation. Let $v_k(a)$ denote value of k-step game starting from 'a'.