Lecture 10

Monday, 26 September 2022 11:32 AM

Recall:

- CCE : $\sigma \in \Delta_{|S|}$ s.t. $\forall i, s_i' \in S_i$,
$$\mathbb{E}_{s\sim\sigma}[u_i(s)] \geq \mathbb{E}_{s\sim\sigma}[u_i(s_i', s_{-i})]$$

- External regret: (single agent)
$$\frac{1}{T} \max_{a\in[n]} \mathbb{E}\left[\sum_{t\leq T} c^t(a^t) - \sum_{t\leq T} c^t(a)\right]$$

- if all players play no external regret algorithms, the average distribution over S obtained is an approximate CCE.

- CE : $\sigma \in \Delta_{|S|}$ s.t. $\forall i, \forall \delta: [n] \to [n]$
$$\mathbb{E}_{s\sim\sigma}[u_i(s)] \geq \mathbb{E}[u_i(\delta(s_i), s_{-i})]$$

Internal Regret / Swap Regret:

Single agent:

At each time step $t = 1 \dots T$

- agent picks $p^t \in \Delta_n$
- adversary sees p^t, picks cost $c^t : [n] \to [0,1]$
- nature picks $a^t \sim p^t$
- agent gets cost $c^t(a^t)$, learns c^t

Internal Regret / Swap Regret:

Defn. :
$$\frac{1}{T} \max_{\delta:[n]\to[n]} \mathbb{E}\left[\sum_{t\leq T} c^t(a^t) - \sum_{t\leq T} c^t(\delta(a^t))\right]$$

hence, IR $\geq$ ER

Note : δ same for each time step.
Randomization affects adversary cost too:

$$IR = \frac{1}{T}\mathbb{E}\left[\sum_{t\leq T}\sum_{a\in[n]} p^t(a)\, c^t(a) - \sum_{t\leq T}\sum_{a\in[n]} p^t(a)\, c^t(\delta(a))\right]$$

Dynamics: Each player i uses a no internal regret algorithm to choose mixed strategy σ_i^t at time step t.
Let $\sigma^t = (\sigma_1^t, \sigma_2^t, \dots, \sigma_m^t)$ be the mixed strategy profile thus obtained.

Theorem: Suppose at time T, the internal regret for each player is at most ε. Then $\hat{\sigma} \in \Delta_{|S|} = \sum_{t=1}^{T} \sigma^t$ is an ε-Correlated Equilibrium
(prove yourself)

Theorem: There exist algorithms with no internal regret

Will use ER minimization algorithms as a subroutine.
In fact, will use n such ER minimization algorithms, where $n = |S_i|$ $\forall i$

$$P^t = \begin{bmatrix} p_1^t \\ \vdots \\ p_n^t \end{bmatrix} \in \mathbb{R}^{n\times n}$$

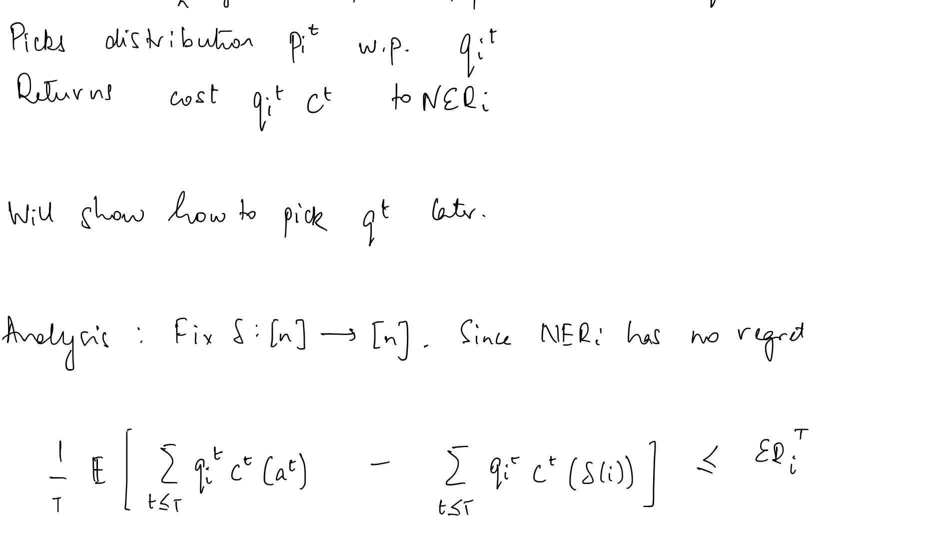

At time t, NIR algo uses $p_1^t, \dots, p_n^t$ to obtain q^t
Picks distribution p_i^t w.p. q_i^t
Returns cost $q_i^t c^t$ to NER$_i$

Will show how to pick q^t later.

Analysis : Fix $\delta : [n] \to [n]$. Since NER$_i$ has no regret

$$\frac{1}{T}\mathbb{E}\left[\sum_{t\leq T} q_i^t c^t(a^t) - \sum_{t\leq T} q_i^t c^t(\delta(i))\right] \leq ER_i^T$$

$$\Rightarrow \frac{1}{T}\left[\sum_{t\leq T}\sum_{a\in[n]} q_i^t p_i^t(a)\, c^t(a) - \sum_{t\leq T} q_i^t c^t(\delta(i))\right] \leq ER_i^T$$

Summing over all n NER algos,

$$\frac{1}{T}\left[\sum_{t\leq T}\sum_{a\in[n]}\sum_{i\in[n]} q_i^t p_i^t(a)\, c^t(a) - \sum_{t\leq T}\sum_{i} q_i^t c^t(\delta(i))\right]$$
$$\leq \sum_i ER_i^T$$

We want $\hat{q}$ s.t.

$$\frac{1}{T}\left[\sum_{t\leq T}\sum_{a\in[n]} \hat{q}^t(a)\, c^t(a) - \sum_{t\leq T}\sum_{a\in[n]} \hat{q}^t(a)\, c^t(\delta(a))\right] \leq \dots$$

so, can we choose q^t : $\forall a$ $\sum_{i\in[n]} q_i^t\, p_i^t(a) = q_i^t$?

$$\begin{bmatrix} & q^t & \end{bmatrix} \begin{bmatrix} p_1^t \\ p_2^t \\ \vdots \\ p_n^t \end{bmatrix} = \begin{bmatrix} & q^t & \end{bmatrix}$$
(columns $a_1\ a_2\ \dots\ a_n$)

Claim: Every stochastic matrix has a left eigenvector that is a stationary distribution.
(proof by BFPT)
Thus, $\exists\, q \in \Delta_n$: $q^T P = q^T$

Given P^t, the NIR algo then returns this left eigenvector q^t.

Hence, algo :
At time $t = 1 \dots T$
- obtain p_i^t from NER$_i$, $i \in [n]$
- compute q^t : $(q^t)^T P^t = q^t$, pick p_i^t w.p. q_i^t
- return cost $q_i^t c^t$ to NER$_i$

Best - Response Dynamics

Dynamics:
At time $t = 1 \dots T$
- if $t = 1$, each player plays an arbitrary pure strategy
- else, pick a player i
- update $s_i^t \leftarrow \arg\min_{s_i' \in S_i} c_i(s_i', s_{-i}^t)$

(if s_i^{t-1} minimizes $c_i(s_i', s_{-i}^t)$, choose $s_i^t = s_i^{t-1}$)

- $s_j^t \leftarrow s_j^{t-1}$ $\forall j \neq i$

- does it converge?
- what does it converge to?
- how fast does it converge?

(Atomic) Network Congestion Games

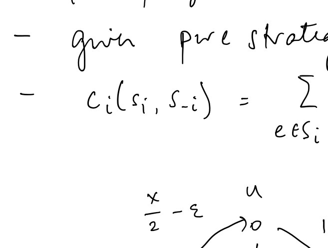

Given: (i) directed graph G
(ii) source s_i, sink t_i for each player $i \in N$ $(n := |N|)$
(iii) cost function $c_e : [n] \to \mathbb{R}_+$

- for player i, $S_i = \{$set of all $s_i - t_i$ paths$\}$
- given pure strategy profile s, $n_e(s) = |\{i : e \in s_i\}|$
- $c_i(s_i, s_{-i}) = \sum_{e\in s_i} c_e(n_e(s))$

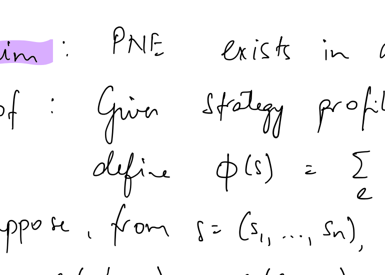

This is a (small) example of Braess' Paradox: without the u-v edge, every player is better off at equilibrium.

Theorem: In an atomic NCG, BRD converges to a PNE if $\exists T$ s.t. in every T time steps, each player gets to update its strategy at least once.

Claim: PNE exists in a NCG.
Proof : Given strategy profile $s \in S$,
define $\phi(s) = \sum_e \sum_{i=1}^{n_e(s)} c_e(i)$
Suppose, from $s = (s_1, \dots, s_n)$, player i deviates to s_i'. Let $s' = (s_i', s_{-i})$
Then $c_i(s_i', s_{-i}) - c_i(s_i, s_{-i}) = \sum_{e\in s_i'\setminus s_i} c_e(n_e(s)+1) - \sum_{e\in s_i\setminus s_i'} c_e(n_e(s))$

Further : $\phi(s_i', s_{-i}) - \phi(s) = \sum_e c_e(n_e(s')) - \sum_e c_e(n_e(s))$

$$= \sum_{e\in s_i'\setminus s_i} c_e(n_e(s)+1) - \sum_{e\in s_i} c_e(n_e(s))$$

Hence, $c_i(s_i', s_{-i}) - c_i(s_i, s_{-i}) = \phi(s_i', s_{-i}) - \phi(s)$
That is, the change in cost for player i if it deviates is exactly the change in ϕ.
Now consider a strategy profile $s^* \in \arg\min_{s\in S} \phi(s)$

Then s^* must be a pure Nash equilibrium, since if deviating from s_i^* can lower the cost for player i, then this deviation must also further lower the potential, which is impossible.

Claim: BRD cannot cycle in a potential game.
Proof: Note that every time any player deviates to play a best response, the potential $\phi(s)$ must decrease. Hence BRD cannot cycle.

Claim: Let $\hat{s}$ be the strategy profile at time t. If $\hat{s}$ is the strategy profile after T further steps of BRD.
Proof: Since in every T steps, every player has the option to deviate but chose not to.

Claim: Let $\hat{s}$ be the strategy profile after $T|S|$ steps of BRD, then $\hat{s}$ is a PNE.
Proof: Since BRD cannot cycle, $\exists s$ that must have persisted for T time steps. By the previous claim, this must be a PNE.

A game is a **potential game** if $\exists$ a potential fn.
$\phi : S \to \mathbb{R}_+$ s.t. $\forall s \in S$, $i \in N$, & $s_i' \in S_i$,
$$c_i(s_i', s_{-i}) - c_i(s_i, s_{-i}) = \phi(s_i', s_{-i}) - \phi(s_i, s_{-i})$$

As is clear from above, every potential game has a PNE, which can be obtained by best-response dynamics.