Algorithmic Game Theory	Summer 2016, Week 10

Best-Response Mechanisms II
(TCP, Stable Matching, Single Item Auction)

ETH Zürich — Paolo Penna

In this lecture we give three more applications of best-response mechanisms:

- The TCP protocol for congestion control (why it stabilizes and why users find it convenient to follow the protocol).
- Stable Matching (we re-obtain incentive compatibility of Proposal Mechanism)
- Second-Price Auction (we re-obtain truthfulness by viewing it as repeated First-Price auction).

Because the framework applies to *asynchronous* settings (players do not necessarily move one by one), in each application we get a 'distributed version' of the previous mechanisms.

1 Best-Response Mechanisms (previous lecture)

Let us recall the basic definitions we introduced in the previous lecture on best-response mechanisms.

Base game G	$\Longrightarrow$	Repeated game G^*
$s_i \in S_i$		response strategy $R_i() \in S_i$
$u_i(s)$		total utility $\Gamma_i := \limsup_{t\to\infty} u_i(s^t)$

Definition 1. *Best-response are* ***incentive compatible*** *for G if repeated best-responding is a Nash equilibrium for the repeated game G^*, that is, for every i*

$$\Gamma_i \geq \Gamma_i'$$

where Γ_i i the total utility when all players best respond and Γ_i' is the total utility when all but i best respond (starting from the same initial profile s^0 and applying the same activation sequence).

We have seen the following result in the previous lecture.

Theorem 2. *If the base game is NBR-solvable with clear outcome (according to a prescribed tie breaking rule $\prec$), then best response converge and are incentive compatible.*

This is how these results can be used.

Best-Response Mechanism

Design a game (rules) that satisfy the two conditions of Theorem 2.

2 TCP Games

The Transmission Control Protocol (TCP) is used to reduce congestion in the Internet and it is run on each computer (sender). We describe the protocol by consider the following example. Two players want to send data through a link of a certain capacity C.

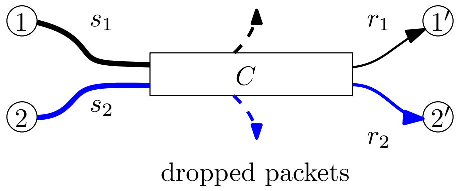

Roughly, the protocol prescribes to adjust the sending rate (number of packets per unit of time) according to these simple rule:

1. No packet loss $\Rightarrow$ increase your rate;
2. Packet loss $\Rightarrow$ decrease your rate.

We view this scenario as the following game. Each sender is a player, who can select a sending rate s_i (the *strategy* of player i) in an interval $[0, M_i]$, and the channel policy (if capacity is exceeded some packets are dropped) determines the actual rate r_i for each player (this amount is the *payoff* of player i). The quantity M_i represents the maximum rate that i is interested in achieving.

The following is an abstract view of what TCP prescribes to do:

Probing Increase Educated Decrease (PIED): Send exactly at the maximum rate that you can get (not more than that).

After gradually increasing the sending rate, at some point some packets are dropped. This is a way for player i to learn the maximum rate he/she can get without packets being dropped. PIED prescribes to send at this maximum rate, that is, to play

$$s_i^* := \max\{s_i \in [0, M_i] | \;\; r_i(s_i, s_{-i}) = s_i\},$$

where the actual rate $r_i()$ depends on the channel policy. There are two natural questions we may ask:

1. If all users run PIED, will the traffic rate stabilize?
2. Are users incetivized to run PIED?

Exercise 1. *Explain why PIED is not incentive compatible if the channel policy is to divide the total capacity* proportionally *to the sending rate of the player (whenever their requests exceed C):*

$$r_i = C\frac{s_i}{\sum_j s_j}.$$

We introduce a channel policy that makes PIED incentive compatible and uses the whole channel capacity:

Strict Priority Queuing: Try to satisfy the players requests one-by-one in a fixed order:

$$\begin{aligned} r_1 &\leftarrow \min(s_1, C) \\ r_2 &\leftarrow \min(s_2, C - r_1) \\ &\vdots \\ r_n &\leftarrow \min(s_n, C - r_1 - r_2 \cdots - r_{n-1}) \end{aligned}$$

Theorem 3. *If the channel uses a Strict Priority Queuing policy then PIED converges and is incentive compatible.*

Proof. We show that the base game is NBR-solvable with clear outcome. The elimination sequence follows the priority order of Strict Priority Queuing: set

$$\begin{aligned} s_1^* &\leftarrow \min(M_1, C) \\ s_2^* &\leftarrow \min(M_2, C - s_1^*) \\ &\vdots \\ s_n^* &\leftarrow \min(M_n, C - s_1^* - s_2^* \cdots - s_{n-1}^*) \end{aligned}$$

and each player i eliminates all strategies different from s_i^* in the order above

$$E_i = \{s_i \neq s_i^*\} \ .$$

To see that this sequence defines an NBR-solvable with clear outcome game we observe the following:

1. Define subgame G_i as the subgame where all players before i have already eliminated their strategies $E_1, \ldots, E_{i-1}$. So G_1 is the original game.

2. In subgame G_i the highest rate available to i is $C - s_1^* - s_2^* - \cdots - s_{i-1}^*$. Since i wants to send at most M_i, strategy s_i^* guarantees i the highest possible payoff in this subgame,
$$\min(M_i, C - s_1^* - s_2^* - \cdots - s_{i-1}^*)$$

3. Sending with rate smaller than s_i^* results in **worse** payoff for i, and sending with higher rate is **not better** than s_i^*. Here we use a simple **tie-breaking rule**, namely

   Prefer smaller sending rate over higher sending rate.

   Then the two cases in Definition 9 correspond to $s_i < s_i^*$ and $s_i > s_i^*$ respectively.

□

Fair Queuing:

1. Allocate the capacity C evenly among all sending "requests"

$$r_i \leftarrow \min(s_i, C/n)$$

2. Recursively allocate the residual capacity among all partially satisfied requests as in previous step:

$$C \leftarrow C - \sum_i r_i \qquad\qquad s_i \leftarrow s_i - r_i.$$

For instance, with $C = 60$ and

$$s_1 = M_1 = 30 \qquad s_2 = M_2 = 40 \qquad s_3 = M_3 = 10$$

we initially allocate $60/3 = 20$ (or less) to each player,

$$r_1 = 20 \qquad r_2 = 20 \qquad r_3 = 10$$

because player 3 does not want more than $M_3 = 10$. The remaining capacity $C' = 10$ is divided equally among the first two:

$$r_1 = 20 + 5 \qquad r_2 = 20 + 5 \qquad r_3 = 10 \ .$$

Exercise 2. *Show that if the channel uses a Fair Queuing policy then PIED converges and is incentive compatible.*

3 Stable Matchings

The general version of the problem considers n players, each of them having preferences over the others. A **stable matching** is a matching such that there are no two players who prefer each other to their matched partners, that is, nothing like this should happen:

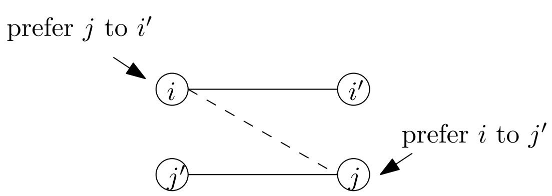

Exercise 3. *Show that in general graphs a stable matching may not exist.*

Example 4 (Interns-Hospitals)**.** *Consider following bipartite restriction of the problem. Players are partitioned into **interns** and **hospitals**:*

- *Hospitals have a **common (same) rank** of interns,*

$$i_1 \succ i_2 \succ \cdots \succ i_n$$

- *Interns rank hospitals differently (e.g., based on salary, location, etc.). So each intern has his/her own (private) rank $\prec_i$ over the hospitals.*

An intuitive mechanism would be to let players propose to the others. A player i **makes a better offer to** j if i proposes to j and j prefers i over all players that currently propose to j. Then a player should try to make a better offer to the player he/she likes the most:

Best-Response Mechanism for Stable Matching:

- Each player checks which players he/she can make a better offer to, and then proposes him/herself to the most preferred one in this set.

One can view the above mechanism as best-responding in the following **stable matching game**:

$$u_i(s) = \begin{cases} rank_i(s_i) & \text{if } i \text{ makes a better offer to } s_i \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where $rank_i(s_i)$ is the natural translation of ranks ($\prec_i$) into utilities: 0 to the least preferred option, 1 to the second-least-preferred, and so on.

Theorem 5. *The Best-Response Mechanism for Stable Matching converges and is incentive compatible for the interns-hospitals matching problem in Example 4.*

Proof. There is a very natural elimination sequence showing that the base game defined by the utilities (1) is NBR-solvable with clear outcome:

- Follow the global order in which interns are ranked by the hospitals,

  $$i_1, i_2, \ldots, i_n$$

  where i_1 is the top-ranked intern for all hospitals.

- At stage t, intern i_t proposes to his/her top-ranked hospital, among those that are still available (not taken by prior interns).

Intuitively, at each step t, intern i_t gets the best hospital, among those currently still available. Formally, we observe the following things:

- Because of the global rank of hospitals, i_t cannot get any of the previously taken hospitals (the utility (1) is 0 because i_t would not make a better offer to such hospitals);

- Because of the global rank of hospitals, i_t can for sure get any of the non-taken hospitals (the utility (1) is nonzero because i_t makes a better offer for any such hospitals);

Let s_t^* be the top-ranked hospital for i_t among those that are not previously taken. Then, the previous two items say that all other strategies are NBR in this subgame (only non-taken hospitals available). The clear outcome comes directly from the last item since s_t^* guarantees i_t the highest payoff in this subgame. □

Incentive compatible here means that it is always convenient to the player to best-respond according to his/her true rank in every iteration. As a corollary we re-obtain the incentive compatibility of the Proposal Algorithm for stable matching in previous lectures.

Corollary 6. *The Proposal Mechanism is incentive compatible.*

Proof. The Proposal Mechanism corresponds to the Best-Response Mechanism above in which players are activated in some order (not important which one).

If the Proposal Mechanism was not incentive compatible, then some player may find it convenient to misreport his/her true preference $\prec_i$ to some $\prec'_i$.

This would mean that in the Best-Response Matching Mechanism this player would improve by repeatedly responding as if her preferences were $\prec'_i$, thus not best responding according to $\prec_i$. □

Example 7 (Correlated Markets)**.** *We have a complete weighted graph. Each player (node) prefers neighbors whose edges weights are higher.*

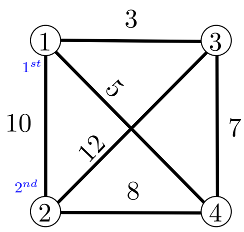

The instances described in the previous exercise satisfy the following definition:

Acyclic Instances: There is no cycle of $\ell \geq 3$ players

$$i_1 \to i_2 \to \cdots \to i_\ell \to i_1$$

such that each player prefers the next one over the previous one.

Exercise 4. *Prove that for acyclic instances the Best-Response Matching Mechanism above converges and is incentive compatible (no player can get matched to a player he/she likes more by misreporting her preferences).*

4 Single Item Auction

We run an auction for selling an item to the players, each player has his/her own valuation v_i for the item. Consider the 2^{nd}**-price auction** in which the highest bid wins the item and the price to pay is the 2^{nd}-highest bid. For example

$$bids: \quad 1, \underbrace{5}_{pays\ 3}, 3$$

The utility of the winner equals to the difference between the valuation and the price to pay (the others have zero utility). This auction "simulates" a **repeated** 1^{st}**-price auction**:

$$bid_1 = 0.1 \to bid_2 = 0.2 \to \cdots \to bid_3 = 3 \to bid_2 = 3.000001$$

where the winner pays his/her final bid.

Exercise 5. *Consider two bidders having* different *valuations and the case bids/valuations are discrete (integers). Describe repeated* 1^{st}*-price auction as a best-response dynamics and prove that it converges and is incentive compatible. Explain how you can deduce from this that* 2^{nd}*-price auction is truthful (reporting a bid different from the true valuation does not improve the utility of the corresponding player).*

Recommended Literature

The three applications presented here are discussed in the same work introducing the best-response mechanisms:

- Noam Nisan, Michael Schapira, Gregory Valiant, and Aviv Zohar. Best-response mechanisms. In *Innovations in Computer Science (ICS)*, pages 155–165, 2011.

More issues on TCP games are described in this work:

- P. Brighten Godfrey, Michael Schapira, Aviv Zo- har, and Scott Shenker. Incentive compatibility and dynamics of congestion control. *SIGMETRICS Perform. Eval. Rev.* , 38(1):95–106, 2010.

  (the setting on general graphs and the need of 'consistent' policies over all edges)

Applications of best-response mechanisms to auctions are in this work:

- Noam Nisan, Michael Schapira, Gregory Valiant, and Aviv Zohar. Best-response auctions. In *ACM conference on Electronic commerce*, pp. 351-360. ACM, 2011.

  (why best-response and the importance of incentive compatibility)

A Definitions from previous lecture

Definition 8 (never best response (NBR)). *A strategy $s_i \in S_i$ is a* never best response (for tie breaking rule $\prec$) *if there is always another strategy that gives a better payoff or that gives the same payoff but is better w.r.t. to this tie breaking rule: for all s_{-i} there exists $s'_i \in S_i$ such that one of these holds*

1. $u_i(s_i, s_{-i}) < u_i(s'_i, s_{-i})$ *or*
2. $u_i(s_i, s_{-i}) = u_i(s'_i, s_{-i})$ *and* $s_i \prec_i s'_i$.

Definition 9 (NBR-solvable). *A game G is* NBR-solvable *if iteratively eliminating NBR strategies results in a game with one strategy per player. That is, there exists a tie breaking rule $\prec$, sequence $p_1, \ldots, p_\ell$ of players, and a corresponding sequence of subsets of strategies $E_1, \ldots, E_\ell$ such that:*

1. *Initially $G_0 = G$ and $G_i + 1$ is the game obtained from G_i by removing the strategies E_i of player p_i;*
2. *Strategies E_i are NBR for $\prec$ in the game G_{i-1}.*
3. *The final game G_ℓ has one strategy for each player (this unique profile is thus a PNE for G).*

A sequence of players and of strategies as above is called an elimination sequence *for the game G.*

Definition 10 (NBR-solvable with clear outcome). *A NBR-solvable game G has a clear outcome if there exists a tie breaking rule $\prec$ such that the following holds. For every player i there exists an elimination sequence consisting of players $p_1, \ldots, p_a, \ldots, p_\ell$ and strategies $E_1, \ldots, E_a, \ldots, E_\ell$ (according to Definition 9) such that,*

1. *p_a denotes the* first appearance *of i in the sequence, that is,*

$$p_a = i \neq p_1, p_2, \ldots, p_{a-1};$$

2. *in the corresponding subgame*

$$G_{a-1} = G \setminus (E_1 \cup E_2 \cup \cdots \cup E_{a-1})$$

the PNE s^* is globally optimal for i, *that is,*

$$u_i(\hat{s}) \leq u_i(s^*) \quad \text{for all } \hat{s} \in G_{a-1}.$$

(Recall that s^* is the unique profile in the final subgame G_ℓ.)