CIS 700 Differential Privacy in Game Theory and Mechanism Design February 21, 2014

Lecture 6

Lecturer: Aaron Roth *Scribe: Aaron Roth*

Private Equilibrium Computation Yields Truthful Mediators (Part 2)

1 Introduction

Recall our goal from last class: we wanted to take a game that has nice properties when described as a game of complete information, and then augment it with an extremely weak "mediator", such that in the new game, we can implement the equilibria of the original (complete information) game as simple *ex-post* equilibria, which do not require players to know anything about each other's types. Some brief recapping is in order.

Suppose $\mathcal{G}$ is a game with action set A, type set $\mathcal{T}$, and cost function $c_i(a) \equiv c(t_i, a)$.

Definition 1 *A* mediator *is given by an algorithm $M : (\mathcal{T} \cup \{\perp\})^n \to (A \cup \{\perp\})^n$ mapping reported types (or $\perp$, which denotes declining to report any type) to suggested actions (or $\perp$, which denotes not suggesting any action).*

Given a game $\mathcal{G}$ and a mediator M, we can define the game $\mathcal{G}$ augmented with mediator M, denoted as $\mathcal{G}_M$, as follows. The new game has action space:

$$A' = \{(t', f) : t' \in \mathcal{T} \cup \{\perp\}, f : (A \cup \{\perp\}) \to A\}$$

i.e. each player chooses both a type t_i to report (possibly $\perp$ if not participating), and a function f specifying how to map the advice received by the mediator to an action a to play in the underlying game $\mathcal{G}$ (Note that picking f equal to the identity function corresponds to following the advice of the mediator, but the player can pick anything else, including a constant function which corresponds to completely disregarding the advice. Note also that if a player reports $t' = \perp$ to the mediator, then it will return the advice $\perp$, and a player will then play the action $f(\perp) \in A$, which encodes the strategy he would have used in the original game). The new game has the following cost function:

$$c'(t_i, ((t'_1, f_1), \ldots, (t'_n, f_n))) = \mathrm{E}_{a \sim M(t')}[c(t_i, f(a))]$$

which simply corresponds to their cost in the original game when they pick actions using the mediator and their function f.

We wanted to implement "good behavior" as an ex-post equilibrium of $\mathcal{G}_M$. Recall that we defined the "good behavior" strategy $g_i = (t_i, \mathrm{id})$ where id is the identity function, which simply corresponds to truthfully reporting one's type, and then faithfully following the suggested action of the mediator.

We defined a variant of differential privacy:

Definition 2 (Joint Differential Privacy [KPRU14]) *Let $M : \mathcal{A}^n \to \mathcal{B}^n$. M satisfies ϵ-joint-differential privacy if for every $a \in \mathcal{A}^n$, for every i, for every $a'_i \in \mathcal{A}$, and for every $S \subseteq \mathcal{B}^{n-1}$:*

$$\Pr[M(a)_{-i} \in S] \leq \exp(\epsilon) \Pr[M(a'_i, a_{-i})_{-i} \in S]$$

and then proved our motivating theorem (here stated with the modification that the game has costs bounded in $[0, L]$):

Theorem 3 ([RR13], based on [KPRU14]) *Let $\mathcal{G}$ be any game with costs bounded in $[0, L]$, and let M be a mechanism satisfying ϵ-joint differential privacy such that on any set of reported types t', M outputs an η-approximate pure strategy Nash equilibrium of the complete information game induced by t'. Then good behavior $g = (g_1, \ldots, g_n)$ is an η'-approximate ex-post equilibrium of $\mathcal{G}_M$ for:*

$$\eta' = 2\epsilon \cdot L + \eta.$$

In this class, we will fix an interesting class of games (congestion games), and show that for "large congestion games", we can compute approximate Nash equilibria under the constraint of joint-differential privacy. This will instantiate Theorem 3 and hence give a setting in which we can implement Nash equilibria of the complete information game ex-post using a weak mediator.

Informally, a *large* game is one in which the unilateral deviation of some player $j \neq i$ can have only a small effect on the cost of agent i. Note that if agent i changes his own action, his cost is allowed to change drastically – his cost is merely insensitive to the actions of single *other* players.

Definition 4 *Fix a game $\mathcal{G}$ with action set A and cost functions $c_i : A^n \to \mathbb{R}$. Say that $\mathcal{G}$ is* large *with sensitivity $\Delta_{\mathcal{G}}$ if for every pair of players $i \neq j$, every vector of actions $a \in A^n$ and every deviation a'_j for player j:*

$$|c_i(a) - c_i(a'_j, a_{-j})| \leq \Delta_{\mathcal{G}}$$

We will be interested in large games in which $\Delta_{\mathcal{G}}$ is a diminishing quantity with the number of players n.

2 Congestion Games

Consider the following natural example of a "large" game: every day, each of you wakes up at your home in Philadelphia, and must choose a route to drive to work. In aggregate, everybody's choices result in *traffic*, which cause delays on each road as a function of how much they are used. Suppose that for each road j, there is a delay function $\ell_j(x) = c_j \cdot x_j + b_j$ which maps the fraction of the population x_j using road j to a delay which is some linear function of its usage parameterized by constants c_j and b_j. Your loss is your total delay on your way to work, which is simply the sum of the delays on the roads that you drive on, given traffic.

This is a nice example of the kind of game we are going to be talking about for a couple of reasons. First, it is *large*: although changing your own route can substantially change your delay, if some other player makes a unilateral change in route, it can have only a small effect on your commute time. This is because a single player can only change the *fraction* of players x_j using a road j by $1/n$, and so if $c = \max_j c_j$ represents an upper bound on the coefficients in the cost functions over all roads, and L represents an upper bound on your longest route to work, then a single player can effect your delay by at most $\frac{c \cdot L}{n} = O(1/n)$ for a fixed road network.

Second, it is an example of a *congestion game*, which we will define shortly. In particular, it is a *linear congestion game* (because the cost functions are linear), for which Nash equilibrium in complete information settings are known to be particularly nice. For one, pure strategy Nash equilibria exist in all congestion games. For another, the worst pure strategy Nash equilibrium in linear congestion games has social welfare that is only a factor of at most 2.5 worse than the optimal social welfare [CK05, AAE05] (i.e. the *price of anarchy* is 2.5) – and so we have a good reason to want to implement them.

Definition 5 *A* congestion game *is defined by:*

1. *A set of n players $\mathcal{P}$ and a typespace $\mathcal{T}$.*
2. *A set of m facilities F.*
3. *A set of actions A_i for each type t_i. Each action $s \in A_i$ is a subset of facilities: $s \subseteq F$.*
4. *For each facility $j \in F$, a cost function $\ell_j : \{0, 1 \ldots, n\} \to \mathbb{R}^+$ mapping the number of players "using" facility j to a non-negative cost. For an action profile $s = (s_1, \ldots, s_n)$ write $n_j(s) = |\{i : j \in s_i\}|$ to denote the number of players using facility j. The cost of an agent i is defined to be the sum of the costs of facilities he is playing on:*

$$c_i = \sum_{j \in s_i} \ell_j(n_j(s))$$

000

In the network routing example, F would be the set of edges in a graph representing a road network, types t_i would correspond to source/destination pairs s_i and d_i, and A_i would correspond to the set of all paths between s_i and d_i in the graph.

An appealing property of congestion games is that they posses pure strategy Nash equilibria, which are the inevitable result of a natural game dynamic (best response dynamics). Here, we will describe an efficient algorithm for computing α-approximate pure strategy Nash equilibria in congestion games for any α (which of course proves the existence of pure strategy Nash equilibria in congestion games).

Given an action profile s, write $\text{BR}_i(s) = \arg\min_{s_i' \in A_i} c_i(s_i', s_{-i})$ to denote player i's best response.

BR$(t = (t_1, \ldots, t_n), \alpha)$:
  Let $s \in A_1 \times \ldots, A_n$ be an arbitrary game state.
  while there exists a player i such that $c_i(s) \geq c_i(\text{BR}_i(s), s_{-i}) + \alpha$ **do**
    Let $s_i \leftarrow \text{BR}_i(s)$
  end while
  Output s

This algorithm is extremely intuitive: it picks an arbitrary game state, and then repeatedly, if any player is able to decrease her cost by at least α by making a best response, the algorithm allows her to do so. It is immediate, therefore, that if the algorithm ever halts and outputs a game state s that s must be an α-approximate Nash equilibrium: by definition, no player can decrease her cost by more than α via a unilateral deviation, or else the algorithm would not have halted!

Theorem 6 *If* ***BR****(t, α) halts and outputs s, then s is an α-approximate pure strategy Nash equilibrium of the complete information game defined by t.*

It remains to show that indeed, the algorithm always halts, and does so after a bounded number of steps.

Theorem 7 *Let $c = \max_{j,k} \ell_j(k)$ be an upper bound on the cost on any facility. Then for any t,* ***BR****(t, α) halts after at most T steps for:*

$$T = \frac{n \cdot m \cdot c}{\alpha} = O\left(\frac{n}{\alpha}\right)$$

Proof Define a potential function defined at every game state:

$$\phi(s) = \sum_{j=1}^{m} \sum_{k=1}^{n_j(s)} \ell_j(k)$$

Note that this potential function is *not* social welfare – it doesn't have a clear semantic meaning. But it has the following useful property: every time an agent unilaterally changes her action, she reduces her cost by exactly the amount that she reduces the potential function. Consider agent i, who makes a unilateral deviation from playing s_i to s_i'. We have:

$$\begin{aligned} c_i(s_i, s_{-i}) - c_i(s_i', s_{-i}) &= \sum_{j \in s_i \setminus s_i'} \ell_j(n_j(s)) - \sum_{j \in s_i' \setminus s_i} \ell_j(n_j(s) + 1) \\ &= \phi(s_i, s_{-i}) - \phi(s_i', s_{-i}) \end{aligned}$$

Note that by definition of ϕ, we have that for all s, $\phi(s) \geq 0$, and also:

$$\max_{s \in A_1 \times \ldots \times A_n} \phi(s) \leq \sum_{j=1}^{m} n_j(s) \max_k \ell_j(k) \leq n \cdot m \cdot c$$

Finally, since every round that we do not halt, a single player decreases her cost by at least α (and hence decreases ϕ by at least α), the algorithm can run for at most

$$T = \frac{n \cdot m \cdot c}{\alpha} = O\left(\frac{n}{\alpha}\right)$$

many rounds. ■

Note that the algorithm is a bit underspecified (in terms of the order in which players make moves), so for concreteness let us think about this variant:

BR$(t = (t_1, \ldots, t_n), \alpha)$:
  Let $s \in A_1 \times \ldots, A_n$ be an arbitrary game state.
  while there exists a player i^* such that $c_{i^*}(s) \geq c_{i^*}(\mathrm{BR}_{i^*}(s), s_{-i^*}) + \alpha$ **do**
    for i from 1 to n **do**
      if $c_i(s) \geq c_i(\mathrm{BR}_i(s), s_{-i}) + \alpha$ **then**
        Let $s_i \leftarrow \mathrm{BR}_i(s)$
      end if
    end for
  end while
  Output s

With this variant, there is no ambiguity: players take turns in round robin order considering whether they should make a best response move; they do if they have one that decreases their cost by at least α.

2.1 Making BR(t, α) Private

Ok – we have an algorithm that quickly computes approximate Nash equilibria in congestion games – if we can modify it so that it is jointly differentially private, we are done. The following lemma, which gives a simple way to prove joint differential privacy, will be instructive:

Lemma 8 (Billboard Lemma [HHR$^+$14, RR13]) *Let $M : \mathcal{T}^n \to \mathcal{O}$ be an ϵ-differentially private algorithm, and let $f : \mathcal{T} \times \mathcal{O} \to A$ be an arbitrary function. Then the algorithm $M' : \mathcal{T}^n \to \mathcal{A}^n$ which computes $o = M(t)$ and then outputs $(f(t_1, o), \ldots, f(t_n, o))$ is ϵ-jointly differentially private.*

Proof Fix any input t, player i, deviation t_i' and $S_{-i} \in A^{n-1}$. We have:

$$\Pr_{a_{-i} \sim M'(t)}[a_{-i} \in S_{-i}] = \Pr_{o \sim M(t)}[(f(t_j), o)_{j \neq i} \in S_{-i}]$$

Define $f_j(o) = f(t_j, o)$ and $\phi(j) = (f_j(o))_{j \neq i}$. Note that ϕ is defined independently of i. Hence by the fact (from last class) that post-processing preserves differential privacy, we know:

$$\begin{aligned} \Pr_{a_{-i} \sim M'(t)}[a_{-i} \in S_{-i}] &= \Pr_{a_{-i} \sim \phi(M(t))}[a_{-i} \in S_{-i}] \\ &\leq \exp(\epsilon) \Pr_{a_{-i} \sim \phi(M(t_i', t_{-i}))}[a_{-i} \in S_{-i}] \end{aligned}$$

which is what we wanted. ■

Lets consider what this lemma means. What it says is if we can compute some piece of information with a differentially private algorithm and post it on a "billboard", so that every player, looking at the billboard, and knowing their own type, can figure out what action to play, then we in fact will have a jointly differentially private algorithm mapping types to actions.

So lets consider what information we need to release to run **BR**(t, α) in a way such that at the end, players know what action they should be playing in s. We note that:

1. In rounds, players take turns making best responses. In the output equilibrium s, each player is simply playing the most recent move he made during the run of the algorithm.

2. Each round, for a player to determine what move she will make, she needs to be able to evaluate the costs of each of her actions. Since we are interested in congestion games, it is sufficient for her to be able to see the *count* of the number of players using each facility.

3. Finally, since players move in a well defined order, if we simulate this algorithm and release a transcript of the sequence of player-counts on each facility over the course of the algorithm, every player can (just by looking at the counts) figure out what action they should be playing at the end. This is because the counts each day specify their costs, which in turn specifies whether or not they will move, and which move they will make.

So, to implement our algorithm so that every agent knows her own specified action at the end, it would be sufficient to output a transcript of facility counts over the run of the algorithm. But wait! We know from last lecture how to privately maintain a counter over a stream of numbers! Can we phrase our algorithm so that the counts that we need to release of the number of agents playing on each facility are simply the running counts of a stream of numbers?

Its not hard to see that we can: for each of the m facilities j, we instantiate a separate stream $\sigma_j \in \{-1, 0, 1\}^T$. Whenever it is player i's turn to move, we send a "1" to the stream of any facility that he moves onto, a "-1" to the stream of any facility that he moves off of, and a "0" to the stream of any facility that he neither moves onto nor off of.

Let us now recall the theorem we derived last lecture, generalized slightly to the case of running m simultaneous counters on m distinct streams.

Theorem 9 ([DNPR10, CSS10]) *There is an ϵ-differentially private algorithm that is simultaneously (α, β)-accurate on m streams of length T that jointly have sensitivity k for:*

$$\alpha = O\left(\frac{k \cdot \log\left(\frac{T \cdot m}{\beta}\right)^{5/2}}{\epsilon}\right)$$

There is a little bit of book-keeping we have to do to use this bound. We will be a bit imprecise for convenience – see [RR13] for the precise calculations.

First of all, we will now have two sources of error. If we run $\mathbf{BR}(t, \alpha_1)$ using noisy counters that have error in their counts α_2, we can account for the error as follows. Since the counts may be off by α_2, player cost estimates for any action may be off by as much as $\alpha_2 \cdot \Delta_{\mathcal{G}}$. This means that whey players think they are playing an α_1 best response, they may in fact only be playing an $\alpha_1 + 2\alpha_2 \cdot \Delta_{\mathcal{G}}$ best response. Hence, when the algorithm terminates, it is only guaranteed to output an $\alpha = \alpha_1 + 2\alpha_2 \cdot \Delta_{\mathcal{G}}$-approximate equilibrium.

Second, how many rounds does the algorithm run for? Again, because of the error in the counts, when players think they are decreasing their costs by α_1, they may in fact only be decreasing their costs by $\alpha_1 - 2\alpha_2 \cdot \Delta_{\mathcal{G}}$. In order for us to still make progress at every round, let us set $\alpha_1 = 3\alpha_2 \cdot \Delta_{\mathcal{G}}$. That way, we are converging to an $\alpha = 4\alpha_2 \cdot \Delta_{\mathcal{G}}$-approximate equilibrium, and players indeed decrease their costs by at least $\alpha_2 \cdot \Delta_{\mathcal{G}}$ at every round. Reasoning as before, this means that it will take us at most T rounds to converge to an α-approximate equilibrium, for

$$T = O\left(\frac{n}{\alpha_2 \cdot \Delta_{\mathcal{G}}}\right)$$

It remains to bound the sensitivity of the m facility counters: how many non-zero bits can a single user contribute to all m counters in aggregate? Let $L = \max_{i, s_i \in A_i} |s_i|$. Then a single best-response move can contribute at most $2L$ nonzero entries to the m streams: at most L 1's to the streams of the

facilities s_i' that player i is deviating to, and at most $L-1$'s to the streams of the facilities s_i that player i is deviating from.

But how many best response moves can a single player make? Here we use the largeness condition on the game. Recall that players do not move until they perceive that they can reduce their cost by α_1. Once some player i makes a best response move, how many other players will have to move until i can again decrease her cost by α_1? Each move by the other players can *increase* the difference in cost between player i's current action and her best response action by at most $2\Delta_{\mathcal{G}}$. Hence, there must be at least $\gamma = O\left(\frac{\alpha_1}{\Delta_{\mathcal{G}}}\right)$ best response moves by other players before player i can move again. Since we know there are at most T best response moves in total, this means that player i can best respond no more than:

$$\frac{T}{\gamma} = O\left(\frac{n \cdot \Delta_{\mathcal{G}}}{\alpha_1 \alpha_2 \cdot \Delta_{\mathcal{G}}}\right) = O\left(\frac{n}{\alpha_2^2 \Delta_{\mathcal{G}}}\right)$$

times. Therefore, the total sensitivity of the m counters is at most

$$k = 4L \cdot \frac{T}{\gamma} = O\left(\frac{n}{\alpha_2^2 \Delta_{\mathcal{G}}}\right)$$

Using Theorem 9, we find that we can achieve error:

$$\alpha_2 = \tilde{O}\left(\frac{k}{\epsilon}\right) = \tilde{O}\left(\frac{n}{\alpha_2^2 \Delta_{\mathcal{G}} \epsilon}\right)$$

Solving for α_2, this is consistent when we set:

$$\alpha_2 = \tilde{O}\left(\left(\frac{n}{\Delta_{\mathcal{G}} \epsilon}\right)^{1/3}\right)$$

We recall that this has all allowed us to compute an α-approximate Nash equilibrium for

$$\alpha = 4\Delta_{\mathcal{G}} \cdot \alpha_2 = \tilde{O}\left(\frac{n^{1/3} \Delta_{\mathcal{G}}^{2/3}}{\epsilon^{1/3}}\right)$$

which means we have sketched the proof of the following theorem:

Theorem 10 ([RR13]) *There is an ϵ-jointly differentially private algorithm that in any large congestion game with sensitivity $\Delta_{\mathcal{G}}$ computes an η-approximate pure strategy Nash equilibrium for:*

$$\eta = \tilde{O}\left(\frac{n^{1/3} \Delta_{\mathcal{G}}^{2/3}}{\epsilon^{1/3}}\right)$$

If (for example), $\Delta_{\mathcal{G}} = O(1/n)$, then this gives:

$$\eta = \tilde{O}\left(\frac{1}{n^{1/3} \epsilon}\right)$$

which is diminishing with n.

Ok. But recall, proving this theorem was not our main goal – we designed a private algorithm for computing Nash equilibria so that we could instantiate Theorem 3, which states that if we have an ϵ-jointly differentially private algorithm that computes η-approximate Nash equilibria, then we can use it as a mediator which implements "good behavior" as an η'-approximate ex-post equilibrium for $\eta' = O(\eta + \epsilon)$. So lets instantiate Theorem 3, and recall that we get to pick ϵ.

Picking $\epsilon = n^{1/4} \Delta_{\mathcal{G}}^{1/2}$ to balance the two sources of error, we have that $\eta + \epsilon = \tilde{O}(n^{1/4} \Delta_{\mathcal{G}}^{1/2})$. Thus we have sketched the proof of the following theorem:

Theorem 11 *For any large congestion game $\mathcal{G}$ with sensitivity $\Delta_{\mathcal{G}}$, there is a mediator M such that in the mediated game $\mathcal{G}_M$, good behavior $g = (g_1, \ldots, g_m)$ is an η-approximate ex-post Nash equilibrium for:*

$$\eta = \tilde{O}(n^{1/4}\Delta_{\mathcal{G}}^{1/2})$$

Moreover, when agents play according to this ex-post equilibrium, the resulting play forms an η-approximate pure strategy Nash equilibrium in the original game $\mathcal{G}$ defined by the realized type vector.

If $\Delta_{\mathcal{G}} = O(1/n)$ (for example), then:

$$\eta = \tilde{O}\left(\frac{1}{n^{1/4}}\right)$$

which is diminishing in n.

Bibliographic Information The method of privately maintaining a counter over a stream of numbers was developed independently by [CSS10] and [DNPR10]. Congestion games (equivalently, potential games) were introduced and studied by Rosenthal [Ros73] and Monderer and Shapley [MS96]. The algorithm for privately computing equilibria in large congestion games, and the resulting theorem about mediated congestion games is from [RR13].

References

[AAE05] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 57–66. ACM, 2005.

[CK05] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73. ACM, 2005.

[CSS10] TH Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. In *Automata, Languages and Programming*, pages 405–417. Springer, 2010.

[DNPR10] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 715–724. ACM, 2010.

[HHR+14] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. In *STOC*, 2014.

[KPRU14] Michael Kearns, Mallesh M Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. *Proceedings of the annual 5th Innovations in Theoretical Computer Science (ITCS) conference*, 2014.

[MS96] Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.

[Ros73] Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[RR13] Ryan Rogers and Aaron Roth. Asymptotically truthful equilibrium selection in large congestion games. *arXiv preprint arXiv:1311.2625*, 2013.