601.436/636 Algorithmic Game Theory **Lecturer:** Michael Dinitz
Topic: Smooth Games and Facility Location **Date:** 3/5/20
Scribe: Michael Dinitz

12.1 Introduction

Today we're going to finish up our discussion of smooth games, and start discussing a more complex game (the facility location game) which we will analyze via smoothness.

12.2 Smooth Games

Let's remember our definition of smooth games:

Definition 12.2.1. *A cost-minimization game with objective function* $cost : S \to \mathbb{R}$ *is* (λ, μ)*-smooth if* $cost(s) \le \sum_{i=1}^{k} C_i(s)$ *for all* $s \in S$ *and*

$$\sum_{i=1}^{k} C_i(s_{-i}, s'_i) \le \lambda \cdot cost(s') + \mu \cdot cost(s)$$

for all $s, s' \in S$.

Last class we proved that every (λ, μ)-smooth game has pure price of anarchy at most $\lambda/(1-\mu)$, and we proved that atomic routing games with affine cost functions are $(5/3, 1/3)$-smooth.

12.2.1 Robustness: coarse correlated equilibria

Last class I promised that we would get to mixed Nash, and we're going to do that by actually going much further: all the way out to coarse correlated equilibria! We're going to prove that smooth games, and in fact a very small extensions of the same analysis we did for pure Nash, will let us bound the gap between the worst CCE and OPT.

Let's start by remembering the definition of a CCE.

Definition 12.2.2. *A distribution* σ *on* S *is a coarse correlated equilibrium if*

$$\mathop{\mathbf{E}}_{s \sim \sigma}[C_i(s)] \le \mathop{\mathbf{E}}_{s \sim \sigma}[C_i(s_{-i}, s'_i)]$$

for all $i \in [k]$ *and* $s'_i \in S_i$.

Let $\text{cost} : S \to R$ be a global cost function with $\text{cost}(s) \le \sum_{i=1}^{k} C_i(s)$ for all $s \in S$. Let CCE denote the set of all coarse correlated equilibria.

Definition 12.2.3. *The* price of total anarchy *of a cost-minimization game is*

$$\frac{\max_{\sigma \in CCE} \mathbf{E}_{s \sim \sigma}[cost(s)]}{\min_{s \in S} cost(s)}$$

In other words, the price of total anarchy is the extension of the price of anarchy to CCEs. This was introduced by [BHLR08], who noticed that for many games it was possible to get the same bounds on the price of anarchy as on the price of anarchy. The following theorem of [Rou15] is basically an explanation for this.

Theorem 12.2.4. *The price of total anarchy of a (λ,μ)-smooth game is at most $\frac{\lambda}{1-\mu}$.*

Proof. Let σ be a coarse correlated equilibrium, and let s^* be the optimal strategy profile. Then

$$\begin{aligned}
\mathop{\mathbf{E}}_{s\sim\sigma}[\mathrm{cost}(s)] &\le \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^{k} C_i(s)\right] = \sum_{i=1}^{k}\mathop{\mathbf{E}}_{s\sim\sigma}[C_i(s)] \le \sum_{i=1}^{k}\mathop{\mathbf{E}}_{s\sim\sigma}[C_i(s_{-i},s_i^*)] && \text{(CCE)}\\
&= \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^{k} C_i(s_{-i},s_i^*)\right] \le \mathop{\mathbf{E}}_{s\sim\sigma}[\lambda\cdot\mathrm{cost}(s^*)+\mu\cdot\mathrm{cost}(s)] && \text{(smooth)}\\
&= \lambda\cdot\mathrm{cost}(s^*)+\mu\cdot\mathop{\mathbf{E}}_{s\sim\sigma}[\mathrm{cost}(s)].
\end{aligned}$$

Now rearranging this inequality gives $\mathbf{E}_{s\sim\sigma}[\mathrm{cost}(s)] \le \frac{\lambda}{1-\mu}\cdot\mathrm{cost}(s^*)$, as claimed. □

So if we proved a pure price of anarchy bound using smoothness, we automatically get a price of total anarchy bound. In particular, even though we only thought about pure Nash equilibria for atomic routing, our bound of 5/2 actually holds even for coarse correlated equilibria. Moreover, because we proved that no-regret dynamics converge to CCEs, this implies that simple and rational behavior (no-regret) actually leads to solutions that aren't too far from optimal!

12.2.2 Robustness: approximate equilibria

Another way that smooth games exhibit "robust" bounds is that the same recipe works even when we're only at an "approximate" equilibrium. This is important in practice (since we're often not *exactly* at equilibrium), and also important for CCEs in particular since at any fixed point in time, the time-averaged history is only an *approximate* CCE. So let's show that smooth games also work well for approximate equilibria.

Definition 12.2.5. *A distribution σ over S is an* ϵ-approximate coarse correlated equilibrium *(ϵ-CCE) if*

$$\mathop{\mathbf{E}}_{s\sim\sigma}[C_i(s)] \le (1+\epsilon)\mathop{\mathbf{E}}_{s\sim\sigma}[C_i(s_{-i},s_i')]$$

for all $i\in[k]$ and $s_i'\in S_i$.

Note that this is slightly different from the definition of an ϵ-approximate CCE we used in Lecture 6: here we have a multiplicative notion of approximate, while there we used an additive notion. It's not too hard to convert between them, but for today for simplicity we're only going to deal with the multiplicative version.

Theorem 12.2.6. *For any (λ,μ)-smooth game and $\epsilon < \frac{1}{\mu}-1$, for every ϵ-CCE σ and strategy profile s^*,*

$$\mathop{\mathbf{E}}_{s\sim\sigma}[cost(s)] \le \frac{(1+\epsilon)\lambda}{1-(1+\epsilon)\mu}cost(s^*)$$

Proof. We can basically just use the same proof that we did before:

$$\begin{aligned}\mathop{\mathbf{E}}_{s\sim\sigma}[\mathrm{cost}(s)] &\le \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^k C_i(s)\right] = \sum_{i=1}^k \mathop{\mathbf{E}}_{s\sim\sigma}[C_i(s)] \le \sum_{i=1}^k (1+\epsilon) \mathop{\mathbf{E}}_{s\sim\sigma}[C_i(s_{-i}, s_i^*)] && (\epsilon\text{-CCE}) \\ &= (1+\epsilon) \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^k C_i(s_{-i}, s_i^*)\right] \le (1+\epsilon) \mathop{\mathbf{E}}_{s\sim\sigma}[\lambda\cdot\mathrm{cost}(s^*) + \mu\cdot\mathrm{cost}(s)] && (\text{smoothness}) \\ &= (1+\epsilon)\lambda\cdot\mathrm{cost}(s^*) + (1+\epsilon)\mu\cdot \mathop{\mathbf{E}}_{s\sim\sigma}[\mathrm{cost}(s)].\end{aligned}$$

Rearranging this inequality gives

$$\mathop{\mathbf{E}}_{s\sim\sigma}[\mathrm{cost}(s)] \le \frac{(1+\epsilon)\lambda}{1-(1+\epsilon)\mu}\mathrm{cost}(s)$$

as claimed. □

For a concrete example of this, consider atomic routing games again, which we know are $(\frac{5}{3}, \frac{1}{3})$-smooth. Then for $\epsilon < \frac{1}{1/3} - 1 = 3 - 1 = 2$, any ϵ-CCE is at most

$$\frac{(1+\epsilon)\frac{5}{3}}{1-(1+\epsilon)\frac{1}{3}} = \frac{5(1+\epsilon)}{3-(1+\epsilon)} = \frac{5+5\epsilon}{2-\epsilon}$$

away from the optimal routing. So even for pretty weak equilibria, like when $\epsilon = 1$ (so no player can do better than *halve* their cost by deviating), we're still only a factor of 10 away from optimum!

12.2.3 Utility Maximization

We're going to want to talk about smooth utility maximization games, rather than cost minimization games. Everything works exactly as you would expect, so I'll just state the main definitions and results that we'll use.

Definition 12.2.7. *A utility maximization game with global objective function $V : S \to \mathbb{R}$ (usually called the value) is (λ, μ)-smooth if*

- $V(s) \ge \sum_{i=1}^k u_i(s)$ *for all* $s \in S$, *and*
- $\sum_{i=1}^k u_i(s_{-i}, s_i') \ge \lambda \cdot V(s') - \mu \cdot V(s)$ *for all* $s, s' \in S$.

Since we want to keep thinking about the *price* of anarchy/total anarchy, for a utility maximization game let's define the price of total anarchy as

$$\frac{\max_{s\in S} V(s)}{\min_{\sigma\in CCE} \mathbf{E}_{s\sim\sigma}[V(s)]}$$

Theorem 12.2.8. *The price of total anarchy in any (λ, μ)-smooth utility-maximization game is at most $\frac{1+\mu}{\lambda}$.*

Proof. Let σ be a CCE and let s^* be the optimal profile. Then

$$\mathop{\mathbf{E}}_{s\sim\sigma}[V(s)] \geq \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^{k} u_i(s)\right] = \sum_{i=1}^{k} \mathop{\mathbf{E}}_{s\sim\sigma}[u_i(s)] \geq \sum_{i=1}^{k} \mathop{\mathbf{E}}_{s\sim\sigma}[u_i(s_{-i}, s_i^*)] = \mathop{\mathbf{E}}_{s\sim\sigma}\left[\sum_{i=1}^{k} u_i(s_{-i}, s_i^*)\right]$$
$$\geq \mathop{\mathbf{E}}_{s\sim\sigma}[\lambda \cdot V(s^*) - \mu \cdot V(s)] = \lambda \cdot V(s^*) - \mu \cdot \mathop{\mathbf{E}}_{s\sim\sigma}[V(s)].$$

Rearranging yields

$$V(s^*) \leq \frac{1+\mu}{\lambda} \mathop{\mathbf{E}}_{s\sim\sigma}[V(s)]$$

as claimed. □

12.3 Facility Location Game

For the rest of today we're going to talk about another game (not routing) which we will also prove is smooth, but which is a bit more complex and so is good practice to see how to use these ideas. This called the "location game", or the "facility location game", or (if we want to be complete accurate and descriptive) "competitive facility location with price-taking markets and profit-maximizing firms".

12.3.1 Definition

- There is a set F of possible "locations".
- There are k players. Player i has as its strategy set some $F_i \subseteq F$. Think of these as "places where player i might build a facility". We'll slightly abuse notation and let $\varnothing \in F_i$, so all players also have the ability to build "nowhere".
- There is a set M of *markets*. Each $j \in M$ has some value $v_j \geq 0$. Think of this as how much the customer at market j is willing to pay for the service that the facilities provide.
- For each $x \in F$ and $j \in M$, there is some cost $c_{xj} \geq 0$ of serving market j from location x.

12.3.1.1 Utilities: example

The utilities are going to end up looking a bit complicated and weird formally, but are pretty intuitive if we work through an example. They're basically the profits that a player gets by putting their facility at some location and trying to serve the markets using their facility. Let's motivate this by an example. We'll let $F = [3]$ and $M = [2]$.

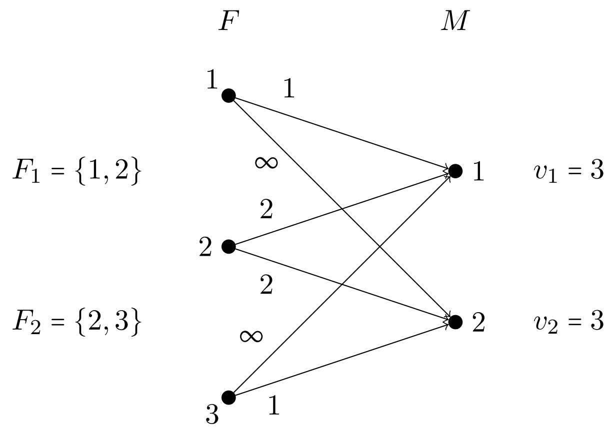

Each player gets to choose a location for their facility, and can also choose a price to charge each market. Each market will choose the player offering it the cheapest price (if it's below their value), and then gives that amount of money to the player, who has to pay the cost for servicing the market from their facility and gets to keep the leftover as profit.

So, for example, suppose that $s_1 = 1$ and $s_2 = 3$ (player 1 chooses location 1 and player 2 chooses location 3). Then player 2 would have to pay an infinite amount fo service market 1, so the price it offers market 1 is ∞ (which player 1 will never choose since it's larger than $v_1 = 3$). Player 1, on the other hand, can charge 3 to market 1, and market 1 will take it. This gives profit (utility) of $3 - c_{11} = 3 - 1 = 2$ to player 1. And note that the same thing happens symmetrically with player 2 and market 2, so under this strategy profile each player gets utility 2.

But now suppose that player 1 changes to instead play location 2. Then player 1 can still charge 3 to market 1, but if player 2 tries to charge 3 to market 2 then player 1 can charge less than 3 to undercut them. The lowest price that player 1 can afford to charge market 2 is 2, in order to cover the cost of servicing, and the lowest price that player 2 can afford to charge is 1 (to cover the cost of servicing). So player 2 can charge 2 (or really $2 - \epsilon$) to market 2, and player 1 won't bother. So in this new strategy profile, the utility of player 1 is $3 - 2 = 1$ and the utility of player 2 is $2 - 1 = 1$.

12.3.1.2 Utilities: Definition

Let's formally define utilities now that we've built up some intuition. To simplify notation, let's assume that $c_{xj} \le v_j$ for all $x \in F$ and $j \in M$. This is without loss of generality, since if we change any service cost larger than v_j to v_j then nothing will change (e.g., in our example changing the ∞ costs to 3 doesn't actually change anything).

Let $s \in S$ be a strategy profile. Let's start with a simple observation: player i at location s_i can get profit from market j only if it's the "closest" (minimum service cost): $c_{s_i j} \le c_{s_x j}$ for all $x \in [k]$. If it *is* the closest, then the highest price it can charge is the second-smallest cost: $p_{ij}(s) = \min_{x \neq i} c_{s_x j}$. So then player i will get profit from market j equal to

$$\pi_{ij}(s) = \begin{cases} p_{ij}(s) - c_{s_i j} & \text{if } c_{s_i j} \le c_{s_x j} \text{ for all } x \in [k] \\ 0 & \text{otherwise} \end{cases}$$

Putting all this together, we get that the utility of player i under strategy profile s is

$$u_i(s) = \sum_{j \in M} \pi_{ij}(s).$$

References

[BHLR08] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 373–382, New York, NY, USA, 2008. Association for Computing Machinery.

[Rou15] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32:1–32:42, 2015.