601.436/636 Algorithmic Game Theory **Lecturer:** Michael Dinitz
Topic: Atomic Routing Games **Date:** 3/1/22

11.1 Introduction

Today is the second lecture on routing games. We're going to focus on *atomic* routing games, unlike the nonatomic case we discussed last lecture.

11.2 Setup

This will be almost the same as last lecture, except the players are large rather than infinitesimal, and each player can have a possibly different sink/source pair.

- Directed graph $G = (V, E)$ with edge cost functions $c_e : \mathbb{R} \to \mathbb{R}$ like last time (continuous, nondecreasing, nonnegative)
- k players, each of which has a source s_i and sink t_i in V.
- Let $\mathcal{P}_i$ be the set of all $s_i \to t_i$ paths in G. Then the strategy set of player i is $\mathcal{P}_i$. Let $\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2 \times \cdots \times \mathcal{P}_k$.
- A *flow* is a tuple $f = (P_1, P_2, \ldots, P_k) \in \mathcal{P}$, i.e., a strategy profile. Given a flow f, like last time we can define $f_e = |\{i : e \in P_i\}|$ and $c_P(f) = \sum_{e \in P} c_e(f_e)$.
- The cost to player i of flow f is $c_{P_i}(f)$.
- The global cost function we'll use is the social cost $\sum_{i=1}^k c_{P_i}(f)$.

We proved in Lecture 5 that every atomic routing game has a pure Nash equilibrium. Today we'll analyze the quality of those equilibria.

11.3 AAE Example

Let's first see an example which shows that the price of anarchy can be worse than in the nonatomic case. This is known as the *AAE example* after the people who discovered it: Baruch Awerbuch, Yossi Azar, and Amir Epstein [AAE05]. Baruch Awerbuch was a professor here at JHU for much of his career, so there's actually a local connection to this topic.

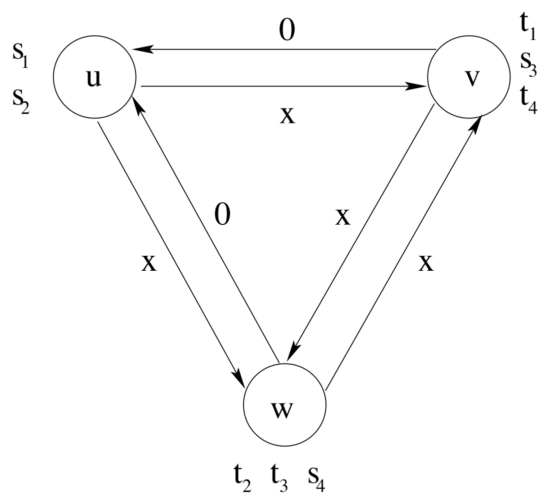

In this example every demand has a one-hop path using an edge with cost function x, but they also have 2-hop paths. First, note that if every player uses their single-hop path then each player pays 1, so the cost of this flow is 4 (this is in fact the optimum flow, although we don't need to prove that).

I claim that all players using their 2-hop paths is a Nash equilibrium. Call this flow f. To see this, let's first figure out the load on each edge:

$$
\begin{aligned}
f_{(u,w)} &= 2 && \text{(players 1 and 3)} \\
f_{(u,v)} &= 2 && \text{(players 2 and 4)} \\
f_{(v,w)} &= 1 && \text{(player 2)} \\
f_{(w,v)} &= 1 && \text{(player 1)}
\end{aligned}
$$

Now we can calculate the cost for each player, as well as the cost if they deviated:

Player	Cost in f	Cost of deviating
1	3	3
2	3	3
3	2	2
4	2	2

So no one has any incentive to deviate, and thus f is a Nash. The cost of f is 10, and thus the price of anarchy is at least $10/4 = 5/2 > 4/3$.

11.4 Upper Bound on Price of Anarchy

Now we're going to prove that the AAE example is the worst case: the pure price of anarchy is always at most 5/2. Unlike the nonatomic case, we won't prove that the AAE example is the worst in a generalized sense – we'll just prove the 5/2 bound.

Theorem 11.4.1. *In every atomic routing game with affine cost functions, the pure price of anarchy is at* 5/2.

Proof. For edge e, let $c_e(x) = a_e x + b_e$ be the affine cost function. Let $f = (P_1, P_2, \ldots, P_k)$ be a pure Nash, and let $f^* = (P_1^*, P_2^*, \ldots, P_k^*)$ be the optimum flow. Since f is a Nash, player i does not want to switch from P_i to P_i^*. This means that

$$\begin{aligned}
\sum_{e \in P_i} c_e(f_e) &\leq \sum_{e \in P_i^* \cap P_i} c_e(f_e) + \sum_{e \in P_i^* \setminus P_i} c_e(f_e + 1) \\
&\leq \sum_{e \in P_i^*} c_e(f_e + 1) && (c_e \text{ is nondecreasing})
\end{aligned}$$

Now we can sum this inequality over all players to get

$$C(f) = \sum_{i=1}^{k} \sum_{e \in P_i} c_e(f_e) \leq \sum_{i=1}^{k} \sum_{e \in P_i^*} c_e(f_e + 1) = \sum_{e \in E} f_e^* c_e(f_e + 1) = \sum_{e \in E} (f_e^* a_e (f_e + 1) + f_e^* b_e)$$

This is starting to look close to what we want, but the right hand side is confusing since f^* and f are combined. We want to "disentangle" them so we can move the f terms to the left side and keep the f^* terms on the right side, giving the kind of inequality that we want. To do this, we're going to use a neat inequality: $y(z+1) \leq \frac{5}{3}y^2 + \frac{1}{3}z^2$ for all $y, z \in \{0, 1, 2, 3 \ldots\}$. It's a good exercise to prove this at home (and is an exercise in Roughgarden's book), but let's just assume it for now. Then we can use this to continue our analysis, applying it to f_e^* and $f_e + 1$ to get:

$$\begin{aligned}
C(f) &\leq \sum_{e \in E} \left(a_e \left(\frac{5}{3}(f_e^*)^2 + \frac{1}{3} f_e^2 \right) + b_e f_e^* \right) \\
&= \sum_{e \in E} f_e^* \left(\frac{5}{3} a_e f_e^* + b_e \right) + \sum_{e \in E} \frac{1}{3} a_e f_e^2 \\
&\leq \sum_{e \in E} f_e^* \left(\frac{5}{3} a_e f_e^* + \frac{5}{3} b_e \right) + \sum_{e \in E} \frac{1}{3} f_e \left(a_e f_e + b_e \right) \\
&= \sum_{e \in E} \frac{5}{3} f_e^* c_e(f_e^*) + \sum_{e \in E} \frac{1}{3} f_e c_e(f_e) \\
&= \frac{5}{3} C(f^*) + \frac{1}{3} C(f)
\end{aligned}$$

Now rearranging this we get that $\frac{2}{3}C(f) \leq \frac{5}{3}C(f^*)$ and thus $C(f) \leq \frac{5}{2}C(f^*)$, proving that the price of anarchy is at most 5/2. $\square$

One thing to note about this proof: it might seem very "lossy", e.g., there were a couple places where we used inequalities that seem rather loose. But the AAE example shows us that in the worst case these inequalities are not lossy, i.e., this is actually a tight analysis.

11.4.1 Extensions

Atomic routing games have been studied pretty extensively, even though much of the original attention focused on nonatomic routing. It's been shown that if different players control different amounts of flow then the price of anarchy is bounded by $\frac{3+\sqrt{5}}{2} \approx 2.618$ rather than 5/2. For non-affine cost functions, the different between atomic and nonatomic games becomes more stark: we saw for nonatomic routing that for degree-p polynomials the price of anarchy is around $\frac{p}{\ln p}$, i.e., it goes sublinearly with p. But for atomic routing, the price of anarchy grows *exponentially* with p.

11.5 Smooth Games

The proof we just did is a "canonical" price of anarchy bound for a certain definition of canonicity that Tim Roughgarden made precise in [Rou15]. Let's think about how this went. Suppose that $s \in S$ is some pure Nash equilibrium, and s^* is the optimal solution, in a cost-minimization game where we care about the social cost.

1. Since our global notion of cost was the social cost, we could sum up over all players of their cost at the Nash. So we have $C(s) = \sum_{i=1}^k C_i(s)$.

2. Since s is a Nash, it wouldn't help any player to deviate, and in particular would not help any player to deviate to what they should be playing in the optimal solution. Thus $C_i(s) \leq C_i(s_{-i}, s_i^*)$. So we have $C(s) \leq \sum_{i=1}^k C_i(s_{-i}, s_i^*)$.

3. $\sum_{i=1}^k C_i(s_{-i}, s_i^*)$ is a weird term that includes both s and s^*. We would like to "disentangle" it. In the atomic routing case, we ended up proving that

$$\sum_{i=1}^{k} C_i(s_{-i}, s_i^*) \leq \frac{5}{3}\sum_{i=1}^{k} C_i(s^*) + \frac{1}{3}\sum_{i=1}^{k} C_i(s).$$

   The important thing to note about this step is that it was just algebra – this disentangling required playing around with some inequalities, but we never used anything game-theoretic. In particular, we didn't use anything about s^* being the optimal solution or s being a Nash.

4. Now we have a disentangled inequality of $C(s) \leq \frac{5}{3}C(s^*) + \frac{1}{3}C(s)$. Basic algebra then gave our price of anarchy bound.

It turns out than an enormous number of price of anarchy proofs follow this pattern (although not all – much of my work has been on games where we don't/can't use this proof outline). What was shown in [Rou15] was that if we formalize this pattern appropriately, we not only get price of (pure) anarchy bounds, but we actually get bounds that are *robust* in a variety of senses. One of those senses is that they extend all the way out to coarse correlated equilibria. This might be surprising and is definitely extremely cool, since it seems crazy that by following a pattern that makes sense for pure Nash we actually get all the way out to CCEs. But that is indeed what happens. One way of viewing this is that in many (though certainly not all) games, the worst CCE is no worse than the worst Nash even though there can be many more CCEs than Nashes.

11.5.1 Basic Definition and Bound

So our first goal is to understand what property of a game makes it possible for us to follow this recipe. Intuitively, what we're going to need is the ability to do the disentangling (which, again, is a purely algebraic property rather than being game-theoretic). Let $\text{cost} : S \to \mathbb{R}$ denote the global cost function that we're using (so we can generalize past social cost). It turns out that what we'll want is the following definition:

Definition 11.5.1. *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 $s, s' \in S$.

Let's prove that if we have a smooth game then we can use the recipe to bound the pure price of anarchy.

Theorem 11.5.2. *The pure price of anarchy of a* (λ, μ)*-smooth game is at most* $\frac{\lambda}{1-\mu}$.

Proof. Let s be a pure Nash equilibrium and let $s^* = \arg\min_{s' \in S}(\text{cost}(s'))$ be the optimal strategy profile. Then we have

$$\begin{aligned}
\text{cost}(s) &\le \sum_{i=1}^{k} C_i(s) && \text{(assumption on cost)} \\
&\le \sum_{i=1}^{k} C_i(s_{-i}, s^*_i) && \text{(s is a pure Nash)} \\
&\le \lambda \cdot \text{cost}(s^*) + \mu \cdot \text{cost}(s) && \text{(definition of smooth)}
\end{aligned}$$

Rearranging this inequality gives $\text{cost}(s) \le \frac{\lambda}{1-\mu}\text{cost}(s^*)$, as claimed. □

And now we can go back and realize that we actually proved that atomic routing games are smooth. Technically we didn't do the full proof, since proving smoothness requires considering all possible pairs of strategy profiles, but since our disentangling step was pure algebra (and not game theory) it's easy to see that we basically did this. But just to be complete, we can do it again.

Theorem 11.5.3. *Atomic routing games with affine cost functions are* $(\frac{5}{3}, \frac{1}{3})$*-smooth.*

Proof. Let $f = (P_1, P_2, \ldots, P_k)$ be one flow (strategy profile) and let $f' = (P'_1, P'_2, \ldots, P'_k)$ be another. Then

$$\begin{aligned}
\sum_{i=1}^{k} C_i(f_{-i}, P'_i) &\le \sum_{i=1}^{k} \sum_{e \in P'_i} c_e(f_e + 1) = \sum_{e \in E} f'_e c_e(f_e + 1) = \sum_{e \in E} (a_e f'_e(f_e + 1) + b_e f'_e) \\
&\le \sum_{e \in E} \left(a_e \left(\frac{5}{3}(f'_e)^2 + \frac{1}{3} f_e^2 \right) + b_e f'_e \right) \le \frac{5}{3} \sum_{e \in E} f'_e c_e(f'_e) + \frac{1}{3} \sum_{e \in E} f_e c_e(f_e) \\
&= \frac{5}{3} C(f') + \frac{1}{3} C(f)
\end{aligned}$$

□

11.5.2 Robustness: coarse correlated equilibria

Now we're going to prove the main shocking fact about smooth games: the bound we proved for the pure price of anarchy also holds for the price of total anarchy! 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 11.5.4. *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 $\mathrm{cost} : S \to R$ be a global cost function with $\mathrm{cost}(s) \le \sum_{i=1}^k C_i(s)$ for all $s \in S$. Let CCE denote the set of all coarse correlated equilibria.

Definition 11.5.5. *The* price of total anarchy *of a cost-minimization game is*

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

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 11.5.6. *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!

11.5.3 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 11.5.7. *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 11.5.8. *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:

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

Rearranging this inequality gives

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

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!

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*, STOC '05, pages 57–66, New York, NY, USA, 2005. Association for Computing Machinery.

[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, page 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.