Lecture 16

Lecturer: Bo Waggoner *Scribe: Aaron Roth, Bo Waggoner*

The Price of Anarchy and Stability

Now we'll transition to questions in mechanism design, where we design or change games so that, when agents play an equilibrium, good things happen. We saw some examples of good things in matching and voting; one was Pareto optimality. But now we will make things more quantitative. What can we say about the quality of the outcome that has been reached?

This is where the *price of anarchy* and *price of stability* come in. They measure how bad things *can* and *must* get respectively, when players play according to an equilibrium. We will focus on *Nash* equilibrium, but more generally, it makes sense to study the price of anarchy and stability over any class of equilibria.

In order to talk about the quality of a game state, we must define what our objective function is. Recall that a game has n agents, each with a utility function $u_i(a)$ giving the utility when the action profile is $a = (a_1, \ldots, a_n)$. **We will give one definition for games with utilities and one for games with costs.** It is important to know which context you are in.

In this case, games with utility functions that players want to maximize, the objective is

$$\mathrm{SW}(a) = \sum_{i=1}^{n} u_i(a)$$

and this is often called the *social welfare* for action profile a. One goal in designing a game is to maximize this objective. When agents play according to the mixed strategy profile $p = (p_1, \ldots, p_n)$, then we can extend the definition to

$$\mathrm{SW}(p) = \sum_{i=1}^{n} u_i(p)$$

where, recall, $u_i(p)$ is player i's expected utility when every plays according to p. We also refer to this as (expected) social welfare.

Given this objective, define OPT to be the optimal value, i.e.

$$\mathrm{OPT} = \max_a \mathrm{SW}(a).$$

This is the quality of the solution we could obtain if we had dictatorial control, and could mandate the action that everyone took.

On the other hand, in a game, players can make decisions independently, and we are interested in how much worse things can be in rational solutions. In other words: **How much worse are all stable (equilibrium) strategy profiles?** This is the price of stability. Formally:

Definition 1 *Let G be a game with nonnegative utility functions and let $NE(G)$ be the set of its Nash equilibria. The* ***price of stability*** *of G is*

$$PoS = \frac{\max_{p \in NE(G)} \mathrm{SW}(p)}{OPT}.$$

Notice we required the utility functions to be nonnegative, because otherwise, one could have a negative ratio, which doesn't make much sense. The PoS for games with utilities is always between 0 and 1, with 1 meaning that equilibria can achieve the optimal welfare. So it measures the fraction of the optimal welfare that is obtained at equilibrium (in the best case).

Prisoner's Dilemma:

	Confess	Silent
Confess	$(1,1)$	$(5,0)$
Silent	$(0,5)$	$(3,3)$

For example, consider the Prisoner's Dilemma. OPT $= 6$ because the social welfare is $3+3=6$ if both players stay silent. However, the only Nash equilibrium is for both players to Confess. In this case, social welfare is 2. So the Price of Stability of this game is $\frac{1}{3}$. Notice that PoS is a ratio, so if we multiplied all utilities in the game by 100, it would not change. However, if we added 100 to all utilities, it would change. This makes sense: If we add 100 to all utilities, then relatively speaking, players don't care very much about the differences between Confess and Silent, and their utility with be a large percentage of optimal no matter what happens.

Battle of the Sexes:

	Football	Opera
Football	$(5,1)$	$(0,0)$
Opera	$(0,0)$	$(1,5)$

Now consider Battle of the Sexes. It has three Nash equilibria: (F,F), (O,O), and a fully mixed equilibrium where the row player plays $(\frac{5}{6}, \frac{1}{6})$ while the column player plays $(\frac{1}{6}, \frac{5}{6})$. (Recall this means that each player independently randomizes, with the row player picking Football with probability $\frac{5}{6}$ and the column player picking Football with probability $\frac{1}{6}$.)

Here we have OPT $= 6$, which is achieved by the profiles (F,F) and (O,O). Since these are both Nash equilibria, we have $PoS = 1$ for this game.

Price of Anarchy. But, continuing the Battle of the Sexes example, what if the players play the third equilibrium? In this case, they do not get the optimal social welfare. In fact, we can calculate the expected social welfare when they play the fully mixed equilbrium by taking the probability of each action profile times the social welfare at that profile:

$$\frac{5}{36}(6) + \frac{25}{36}(0) + \frac{1}{36}(0) + \frac{5}{36}(6) = \frac{5}{3}.$$

The ratio of this value to OPT is only $\frac{5/3}{6} = \frac{5}{18}$, which isn't so great. This is the *price of anarchy* of the game.

Definition 2 *The* price of anarchy *of a game G with nonnegative utility functions is*

$$PoA = \frac{\min_{p \in NE(G)} SW(p)}{OPT}$$

The price of anarchy is the ratio of social welfare of the *worst* equilibrium to OPT. Notice we have

$$0 \leq PoA \leq PoS \leq 1$$

The price of anarcy measures how much worse things *can be* compared to OPT, while the price of stability measures how much worst things *must be.*

The names are appropriate. The price of anarchy measures how bad things can get if we let everyone act for themselves – if we have anarchy – and assume only that they are rational enough to reach equilibrium. In contrast, if we have the power to suggest to players how they should play, we could suggest that they play the *best* Nash equilibrium. But if we want our suggestions to be stable, we must suggest some stable state (i.e. an equilibrium). The price of stability tells us how bad things must be, even if we get to pick our favorite stable state.

Games with cost functions instead of utilities. We will also define PoS and PoA for games with cost functions $c_i : A \to \mathbb{R}$ where each player's objective is to minimize cost. In fact, this is how they are more frequently defined. It is important, when discussing the terms, to always know which scenario you are in – a game with utilities or costs!

In a game with costs, the objective is social cost:

$$SC(a) = \sum_{i=1}^{n} c_i(a).$$

We define $SC(p)$ for mixed strategy profiles p the same way. The goal is to minimize social cost, i.e.

$$\text{OPT} = \min_a SC(a)$$

Definition 3 *For a game G with nonnegative cost functions c_i and Nash equilibria $NE(G)$, we define the* ***Price of Stability*** *as*

$$PoS = \frac{\min_{p \in NE(G)} SC(p)}{OPT}$$

and the ***Price of Anarchy*** *as*

$$PoA = \frac{\max_{p \in NE(G)} SC(p)}{OPT}$$

This time, note that

$$1 \leq PoS \leq PoA \leq \infty$$

So now the ratio measures how many times worse (or higher) costs are in equilibria than at OPT.

Congestion game example. We will give one quick example, the well-known Braess paradox. Consider the following graph. n agents are trying to travel from the leftmost vertex to the rightmost. Each has cost equal to their total travel time. On some of the edges the travel time is a constant 1, while on the others, the travel time is $\frac{k}{n}$ where k is the total number of agents traveling on that edge.

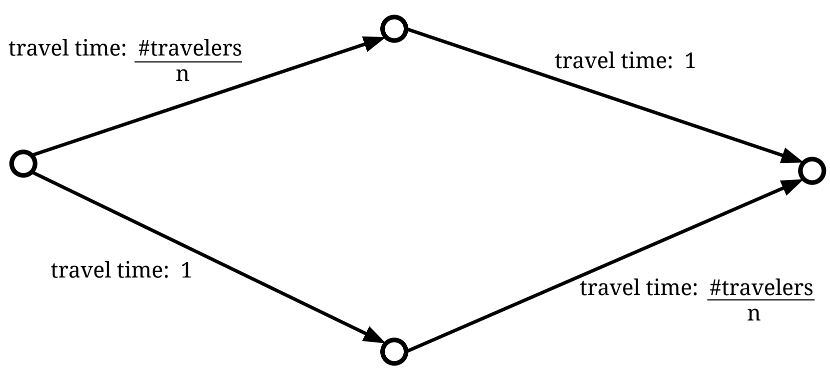

One can verify that the only Nash equilibrium (if n is even) is for half the agents to take the top path while half take the bottom path. In this case, social cost is $\frac{3}{2}n$ because everyone's cost is $\frac{3}{2}$. This is also the optimal solution, i.e. OPT $= \frac{3}{2}n$, so the price of anarchy and price of stability are both 1.

But now, suppose we add a new edge with zero travel time. One might think that this additional edge can only improve congestion and shorten travel times.

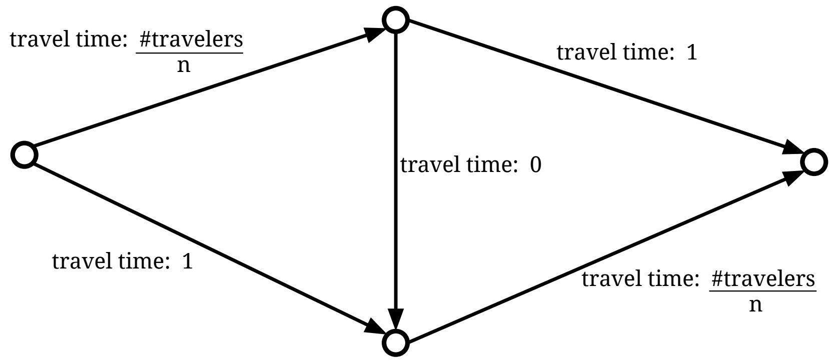

However, this is not the case. One can check that, while the socially optimal cost remains $\frac{3}{2}n$, it as an equilibrium when all players follow the new cut-through edge[1], leading to a social cost of $2n$. The price of anarchy and, roughly, price of stability now become

$$\frac{2n}{(3/2)n} = \frac{4}{3}.$$

Ironically, by adding a new road, we have made the price of stability and total social cost worse!

In general, one can ask how to best design or modify routing graphs and other congestion games. PoA and PoS are useful concepts for reasoning about such modifications, e.g. a change that greatly improves OPT may not help much if the PoS also increases.

[1]To see this, note that if any player is not following the new edge, then she can improve by switching to it, unless she is the only one.