CS364A: Algorithmic Game Theory
Lecture #13: Potential Games; A Hierarchy of Equilibria*

Tim Roughgarden†

November 4, 2013

Last lecture we proved that every pure Nash equilibrium of an atomic selfish routing game with affine cost functions (of the form $c_e(x) = a_e x + b_e$ with $a_e, b_e \geq 0$) has cost at most $\frac{5}{2}$ times that of an optimal outcome, and that this bound is tight in the worst case. There can be multiple pure Nash equilibria in such a game, and this bound of $\frac{5}{2}$ applies to all of them. But how do we know that there is at least one? After all, there are plenty of games like Rock-Paper-Scissors that possess no pure Nash equilibrium. How do we know that our price of anarchy (POA) guarantee is not vacuous? This lecture introduces basic definitions of and questions about several equilibrium concepts. When do they exist? When is computing one computationally tractable? Why should we prefer one equilibrium concept over another?

1 Potential Games and the Existence of Pure Nash Equilibria

Atomic selfish routing games are a remarkable class, in that pure Nash equilibria are guaranteed to exist.

Theorem 1.1 (Rosenthal's Theorem [4]) *Every atomic selfish routing game, with arbitrary real-valued cost functions, has at least one equilibrium flow.*

Proof: We show that every atomic selfish routing game is a *potential game*. Intuitively, we show that players are inadvertently and collectively striving to optimize a "potential function." This is one of the only general tools for guaranteeing the existence of pure Nash equilibria in a class of games.

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: `tim@cs.stanford.edu`.

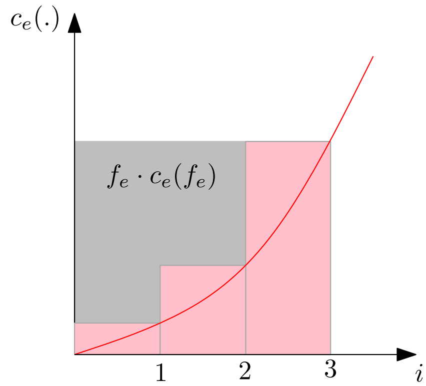

Figure 1: The function c_e and its corresponding (underestimating) potential function.

Formally, define a *potential function* on the set of flows of an atomic selfish routing game by

$$\Phi(f) = \sum_{e \in E} \sum_{i=1}^{f_e} c_e(i), \tag{1}$$

where f_e is the number of player that choose a path in f that includes the edge e. The inner sum in (1) is the "area under the curve" of the cost function c_e; see Figure 1. Contrast this with the corresponding term $f_e \cdot c_e(f_e)$ in the cost objective function we studied last week, which corresponds to the shaded bounding box in Figure 1. The similarity between the potential function and the cost objective function can be useful, as we'll see later.

The defining condition of a potential function is the following. Consider an arbitrary flow f, and an arbitrary player i, using the s_i-t_i path P_i in f, and an arbitrary deviation to some other s_i-t_i path $\hat{P}_i$. Let $\hat{f}$ denote the flow after i's deviation from P_i to $\hat{P}_i$. Then,

$$\Phi(\hat{f}) - \Phi(f) = \sum_{e \in \hat{P}_i} c_e(\hat{f}_e) - \sum_{e \in P_i} c_e(f_e); \tag{2}$$

that is, the change in the potential function under a unilateral deviation is exactly the same as the change in the deviator's individual cost. In this sense, the single function Φ simultaneously tracks the effect of deviations by each of the players.

Once the potential function Φ is correctly guessed, the property (2) is easy to verify. Looking at (1), we see that the inner sum of the potential function corresponding to edge e picks up an extra term $c_e(f_e + 1)$ whenever e is newly used by player i (i.e., in $\hat{P}_i$ but not P_i), and sheds its final term $c_e(f_e)$ when e is newly unused by player i. Thus, the left-hand side of (2) is

$$\sum_{e \in \hat{P}_i \setminus P_i} c_e(f_e + 1) - \sum_{P_i \setminus \hat{P}_i} c_e(f_e),$$

which is exactly the same as the right-hand side of (2).

Given the potential function Φ, the proof of Theorem 1.1 is easy. Let f denote the flow that minimizes Φ — since there are only finitely many flows, such a flow exists. Then, no unilateral deviation by any player can decrease Φ. By (2), no player can decrease its cost by a unilateral deviation and so f is an equilibrium flow. ■

2 Extensions

The proof idea in Theorem 1.1 can be used to prove a number of other results. First, the proof remains valid for arbitrary cost functions, nondecreasing or otherwise. We'll use this fact in Lecture 15, when we discuss a class of games with "positive externalities."

Second, the proof of Theorem 1.1 never uses the network structure of a selfish routing game. That is, the argument remains valid for *congestion games*, the generalization of atomic selfish routing games in which there is an abstract set E of resources (previously, edges), each with a cost function, and each player i has an arbitrary collection $S_i \subseteq 2^E$ of strategies (previously, s_i-t_i paths), each a subset of resources. We'll discuss congestion games at length in Lecture 19.

Finally, analogous arguments apply to the *nonatomic* selfish routing networks introduced in Lecture 11. We sketch the arguments here; details are in [5]. Since players have negligible size in such games, we replace the inner sum in (1) by an integral:

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} c_e(x)dx, \tag{3}$$

where f_e is the amount of traffic routed on edge e by the flow f. Because cost functions are assumed continuous and nondecreasing, the function Φ is continuously differentiable and convex. The first-order optimality conditions of Φ are precisely the equilibrium conditions of a flow in a nonatomic selfish routing network (see Exercises). This gives a sense in which the local minima of Φ correspond to equilibrium flows. Since Φ is continuous and the space of all flows is compact, Φ has a global minimum, and this flow must be an equilibrium. This proves existence of equilibrium flows in nonatomic selfish routing networks. Uniqueness of such flows follows from the convexity of Φ — its only local minima are its global minima. When Φ has multiple global minima, all with the same potential function value, these correspond to multiple equilibrium flows that all have the same total cost.

3 A Hierarchy of Equilibrium Concepts

How should we discuss the POA of a game with no pure Nash equilibria (PNE)? In addition to games like Rock-Paper-Scissors, atomic selfish routing games with varying player sizes need not have PNE, even with only two players and quadratic cost functions [5, Example 18.4]. For a meaningful POA analysis of such games, we need to enlarge the set of equilibria to recover guaranteed existence. The rest of this lecture introduces three relaxations of PNE, each more permissive and more computationally tractable than the previous one (Figure 2). All three of these more permissive equilibrium concepts are guaranteed to exist in every finite game.

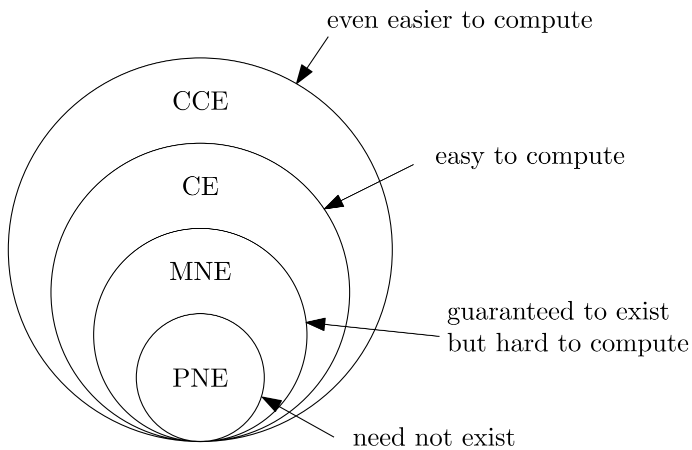

Figure 2: The Venn-diagram of the hierarchy of equilibrium concepts.

3.1 Cost-Minimization Games

A *cost-minimization* game has the following ingredients:

- a finite number k of players;
- a finite strategy set S_i for each player i;
- a cost function $C_i(\mathbf{s})$ for each player i, where $\mathbf{s} \in S_1 \times \cdots \times S_k$ denotes a *strategy profile* or *outcome*.

For example, atomic routing games are cost-minimization games, with $C_i(\mathbf{s})$ denoting i's travel time on its chosen path, given the paths $\mathbf{s}_{-i}$ chosen by the other players.

Conventionally, the following equilibrium concepts are defined for payoff-maximization games, with all of the inequalities reversed. The two definitions are completely equivalent.

3.2 Pure Nash Equilibria (PNE)

Recall the definition of a PNE: unilateral deviations can only increase a player's cost.

Definition 3.1 A strategy profile $\mathbf{s}$ of a cost-minimization game is a *pure Nash equilibrium (PNE)* if for every player $i \in \{1, 2, \ldots, k\}$ and every unilateral deviation $s_i' \in S_i$,

$$C_i(\mathbf{s}) \leq C_i({s_i}', \mathbf{s}_{-i}). \tag{4}$$

PNE are easy to interpret but, as discussed above, do not exist in all games of interest. We leave the POA of pure Nash equilibria undefined in games without at least one PNE.

3.3 Mixed Nash Equilibria (MNE)

When we discussed the Rock-Paper-Scissors game in Lecture 1, we introduced the idea of a player randomizing over its strategies via a *mixed strategy.* In a mixed Nash equilibrium, players randomize independently and unilateral deviations can only increase a player's expected cost.

Definition 3.2 Distributions $\sigma_1, \ldots, \sigma_k$ over strategy sets $S_1, \ldots, S_k$ of a cost-minimization game constitute a *mixed Nash equilibrium (MNE)* if for every player $i \in \{1, 2, \ldots, k\}$ and every unilateral deviation $s_i' \in S_i$,

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i({s_i}', \mathbf{s}_{-i})]\,, \tag{5}$$

where σ denotes the product distribution $\sigma_1 \times \cdots \times \sigma_k$.

Definition 3.2 only considers pure-strategy unilateral deviations; also allowing mixed-strategy unilateral deviations does not change the definition (Exercise).

By the definitions, every PNE is the special case of MNE in which each player plays deterministically. The Rock-Paper-Scissors game shows that, in general, a game can have MNE that are not PNE.

Here are two highly non-obvious facts that we'll discuss at length in Lecture 20. First, every cost-minimization game has at least one MNE; this is Nash's theorem [3]. Second, computing a MNE appears to be a computationally intractable problem, even when there are only two players. For now, by "seems intractable" you can think of as being roughly NP-complete; the real story is more complicated, as we'll discuss in the last week of the course.

The guaranteed existence of MNE implies that the POA of MNE is well defined in every finite game — this is an improvement over PNE. The computational intractability of MNE raises the concern that POA bounds for them need not be meaningful. If we don't expect the players of a game to quickly reach an equilibrium, why should we care about performance guarantees for equilibria? This objection motivates the search for still more permissive and computationally tractable equilibrium concepts.

3.4 Correlated Equilibria (CE)

Our next equilibrium notion takes some getting used to. We define it, then explain the standard semantics, and then offer an example.

Definition 3.3 A distribution σ on the set $S_1 \times \cdots \times S_k$ of outcomes of a cost-minimization game is a *correlated equilibrium (CE)* if for every player $i \in \{1, 2, \ldots, k\}$, strategy $s_i \in S_i$, and every deviation $s_i' \in S_i$,

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s}) \,|\, s_i] \leq \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i({s_i}', \mathbf{s}_{-i}) \,|\, s_i]\,. \tag{6}$$

Importantly, the distribution σ in Definition 3.3 need not be a product distribution; in this sense, the strategies chosen by the players are correlated. Indeed, the MNE of a game correspond to the CE that are product distributions (see Exercises). Since MNE are guaranteed to exist, so are CE. Correlated equilibria also have a useful equivalent definition in terms of "switching functions;" see the Exercises.

The usual interpretation of a correlated equilibrium [1] involves a trusted third party. The distribution σ over outcomes is publicly known. The trusted third party samples an outcome $\mathbf{s}$ according to σ. For each player $i = 1, 2, \ldots, k$, the trusted third party privately suggests the strategy s_i to i. The player i can follow the suggestion s_i, or not. At the time of decision-making, a player i knows the distribution σ, one component s_i of the realization σ, and accordingly has a posterior distribution on others' suggested strategies $\mathbf{s}_{-i}$. With these semantics, the correlated equilibrium condition (6) requires that every player minimizes its expected cost by playing the suggested strategy s_i. The expectation is conditioned on i's information — σ and s_i — and assumes that other players play their recommended strategies $\mathbf{s}_{-i}$.

Believe it or not, a traffic light is a perfect example of a CE that is not a MNE. Consider the following two-player game:

	stop	go
stop	0,0	0,1
go	1,0	-5,-5

If the other player is stopping at an intersection, then you would rather go and get on with it. The worst-case scenario, of course, is that both players go at the same time and get into an accident. There are two PNE, (stop,go) and (go,stop). Define σ by randomizing 50/50 between these two PNE. This is not a product distribution, so it cannot correspond to a MNE of the game. It is, however, a CE. For example, consider the row player. If the trusted third party (i.e., the stoplight) recommends the strategy "go" (i.e., is green), then the row player knows that the column player was recommended "stop" (i.e., has a red light). Assuming the column player plays its recommended strategy (i.e., stops at the red light), the best response of the row player is to follow its recommendation (i.e., to go). Similarly, when the row player is told to stop, it assumes that the column player will go, and under this assumption stopping is a best response.

In Lecture 18 we'll prove that, unlike MNE, CE are computationally tractable. One proof goes through linear programming. More interesting, and the focus of our lectures, is the fact that there are distributed learning algorithms that quickly guide the history of joint play to the set of CE.

3.5 Coarse Correlated Equilibria (CCE)

We should already be quite pleased with positive results, like good POA bounds, that apply to the computationally tractable set of CE. But if we can get away with it, we'd be happy to enlarge the set of equilibria even further, to an "even more tractable" concept.

Definition 3.4 ([2]) A distribution σ on the set $S_1 \times \cdots \times S_k$ of outcomes of a cost-minimization game is a *coarse correlated equilibrium (CCE)* if for every player $i \in \{1, 2, \ldots, k\}$ and every unilateral deviation $s_i' \in S_i$,

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i({s_i}', \mathbf{s}_{-i})]\,. \tag{7}$$

In the equilibrium condition (7), when a player i contemplates a deviation ${s_i}'$, it knows only the distribution σ and *not* the component s_i of the realization. Put differently, a CCE only protects against unconditional unilateral deviations, as opposed to the conditional unilateral deviations addressed in Definition 3.3. Every CE is a CCE — see the Exercises — so CCE are guaranteed to exist in every finite game and are computationally tractable. As we'll see in a couple of weeks, the distributed learning algorithms that quickly guide the history of joint play to the set of CCE are even simpler and more natural than those for the set of CE.

3.6 An Example

We next consider a concrete example, to increase intuition for the four equilibrium concepts in Figure 2 and to show that all of the inclusions can be strict.

Consider an atomic selfish routing game (Lecture 12) with four players. The network is simply a common source vertex s, a common sink vertex t, and 6 parallel s-t edges $E = \{0, 1, 2, 3, 4, 5\}$. Each edge has the cost function $c(x) = x$.

The pure Nash equilibria of this game are the $\binom{6}{4}$ outcomes in which each player chooses a distinct edge. Every player suffers only unit cost in such an equilibrium. One mixed Nash equilibrium that is obviously not pure has each player independently choosing an edge uniformly at random. Every player suffers expected cost $3/2$ in this equilibrium. The uniform distribution over all outcomes in which there is one edge with two players and two edges with one player each is a (non-product) correlated equilibrium, since both sides of (6) read $\frac{3}{2}$ for every i, s_i, and s_i' (see Exercises). The uniform distribution over the subset of these outcomes in which the set of chosen edges is either $\{0, 2, 4\}$ or $\{1, 3, 5\}$ is a coarse correlated equilibrium, since both sides of (7) read $\frac{3}{2}$ for every i and s_i'. It is not a correlated equilibrium, since a player i that is recommended the edge s_i can reduce its conditional expected cost to 1 by choosing the deviation s_i' to the successive edge (modulo 6).

3.7 Looking Ahead: POA Bounds for Tractable Equilibrium Concepts

The benefit of enlarging the set of equilibria is increased tractability and plausibility. The downside is that, in general, fewer desirable properties will hold.

For example, consider POA bounds, which by definition compare the objective function value of the *worst* equilibrium of a game to that of an optimal solution. The larger the set of equilibria, the worse (i.e., further from 1) the POA. Is there a "sweet spot" equilibrium concept that is simultaneously big enough to enjoy tractability and small enough to permit strong worst-case guarantees? In the next lecture, we give an affirmative answer for many interesting classes of games.

References

[1] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[2] H. Moulin and J. P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3/4):201–221, 1978.

[3] J. F. Nash. Equilibrium points in N-person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.

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

[5] T. Roughgarden. Routing games. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 18, pages 461–486. Cambridge University Press, 2007.