

CS364A: Algorithmic Game Theory

Lecture #14: Robust Price-of-Anarchy Bounds in Smooth Games*

Tim Roughgarden[†]

November 6, 2013

1 Canonical POA Proofs

In Lecture 12 we proved that the price of anarchy (POA) in every atomic selfish routing game with affine cost functions is at most $\frac{5}{2}$. To review, the proof had the following high-level steps.

1. Given an arbitrary pure Nash equilibrium (PNE) \mathbf{s} , the PNE hypothesis is invoked once per player i with the hypothetical deviation s_i^* , where \mathbf{s}^* is an optimal outcome, to derive the inequality $C_i(\mathbf{s}) \leq C_i(s_i^*, \mathbf{s}_{-i})$ for each i . Importantly, this is the only time that the PNE hypothesis is invoked in the entire proof.
2. The k inequalities that bound individuals' equilibrium costs are summed over the players. The left-hand side of the resulting inequality is the cost of the PNE \mathbf{s} ; the right-hand side is a strange entangled function of \mathbf{s} and \mathbf{s}^* (involving terms of the form $f_e f_e^*$).
3. The hardest step is to relate the entangled term $\sum_{i=1}^k C_i(s_i^*, \mathbf{s}_{-i})$ generated by the previous step to the only two quantities that we care about, the costs of \mathbf{s} and \mathbf{s}^* . Specifically, we proved an upper bound of $\frac{5}{3}\text{cost}(\mathbf{s}^*) + \frac{1}{3}\text{cost}(\mathbf{s})$. Importantly, this step is just algebra, and is agnostic to our choices of \mathbf{s} and \mathbf{s}^* as a PNE and an optimal outcome, respectively.
4. Solve for the POA. Subtracting $\frac{1}{3} \cdot \text{cost}(\mathbf{s})$ from both sides and multiplying through by $\frac{3}{2}$ proves that the POA is at most $\frac{5}{2}$.

*©2013, Tim Roughgarden.

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




Figure 1: Location game with 3 locations (F), 2 markets (M) and 2 players.

This proof is canonical, in a sense that we formalize in this lecture. Many other POA proofs are canonical in the same sense. The main point of this lecture is that POA proofs that follow this general recipe automatically generate “robust POA bounds.” That is, the proved guarantee applies not only to all PNE of the game, but also to, among other things, all of its coarse correlated equilibria — the biggest set of equilibria defined in the last lecture.

2 A Location Game

Before proceedings to the general theory, it will be helpful to have another concrete example under our belt. Consider a *location game* with the following ingredients:

- A set F of possible locations. These could represent possible locations to build a Web cache in a network, an artisanal chocolate shop in a gentrifying neighborhood, and so on.
- A set of k players. Each player i chooses one location from a set $F_i \subseteq F$ from which to provide a service. All players provide the same service; they differ only in where they are located. There is no limit on the number of markets that a player can provide service to.
- A set M of markets. Each market $j \in M$ has a value v_j that is known to all players. This is the market’s maximum willingness-to-pay for receiving a service.
- For each location $\ell \in F$ and market $j \in M$, there is a known cost $c_{\ell j}$ of serving j from ℓ . This could represent physical distance, the degree of incompatibility between two technologies, and so on.

Given a strategy profile — a location choice by each player — each player tries to capture as many markets as possible, at the highest prices possible. To define the payoffs precisely, we start with an example. Figure 1 shows a location game with $F = \{1, 2, 3\}$ and $M = \{1, 2\}$. There are two players, with $F_1 = \{1, 2\}$ and $F_2 = \{2, 3\}$. Both markets have value 3. The cost between location 2 and either market is 2; locations 1 and 3 have cost 1 to the nearer market (1 and 2, respectively) and infinite cost to the other market.

Continuing the example, suppose the first player chooses location 1 and the second player chooses location 3. Then, each player has a monopoly in the market that they entered. The

only thing restricting the price charged is the maximum willingness-to-pay of the market. Thus, each player can charge 3 for its service to its market. Since the cost of service is 1 in both cases, both players have a payoff of $3 - 1 = 2$.

Alternatively, suppose the first player switches to location 2, while the second player remains at location 3. Player 1 still has a monopoly in market 1, and thus can still charge 3. Its service cost has jumped to 2, however, so its payoff from that market has dropped to 1. In market 2, player 2 can no longer charge a price of 3 without consequence — at any price strictly bigger than 2, the player 1 can profitably undercut the price and take the market. Thus, player 2 will charge the highest price it can get away with, which is 2. Since its cost of serving the market is 1, player 2's payoff is $2 - 1 = 1$.

In general, in a strategy profile \mathbf{s} of a location game, the payoff of player i is defined as

$$\pi_i(\mathbf{s}) = \sum_{j \in M} \pi_{ij}(\mathbf{s}),$$

where, assuming that C is the set of chosen locations and i chooses $\ell \in C$,

$$\pi_{ij}(\mathbf{s}) = \begin{cases} 0 & \text{if } c_{\ell j} \geq v_j \text{ or } \ell \text{ is not the closest location of } C \text{ to } j \\ d_j^{(2)}(\mathbf{s}) - c_{\ell j} & \text{otherwise,} \end{cases} \quad (1)$$

where $d_j^{(2)}(\mathbf{s})$ is the highest price that player i can get away with, namely the minimum of v_j and the second-smallest cost between a location of C and j .¹

The payoff $\pi_{ij}(\mathbf{s})$ is thus the “competitive advantage” that i has over the other players for market j , up to a cap of v_j minus the service cost. The definition in (1) assumes that each market is served by the potential provider with the lowest service cost, at the highest competitive price. This assumption can also be justified from first principles by setting up a three-stage game and proving that its “subgame perfect equilibria” have these properties; see [2] for more details.

The objective function in a location game is to maximize the social surplus. The surplus $V(\mathbf{s})$ of a strategy profile \mathbf{s} — a location choice by each player — is defined as

$$V(\mathbf{s}) = \sum_{j \in M} v_j - d_j(\mathbf{s}), \quad (2)$$

where $d_j(\mathbf{s})$ is the minimum of v_j and the smallest cost between a chosen location and j . The definition (2) assumes that each market j is served by the chosen location with the smallest cost of serving j , or not at all if this cost is at least v_j .

Note that the surplus $V(\mathbf{s})$ depends on the strategy profile \mathbf{s} only through the set of locations chosen by some player in \mathbf{s} . Indeed, the definition (2) makes sense more generally for any subset of chosen locations T , and we sometimes write $V(T)$ for this quantity.

The rest of this section proves that every PNE of every location game has social surplus at least 50% times that of an optimal outcome.

¹In contrast to selfish routing games, these location games are most naturally described as payoff-maximization games. POA bounds are equally interesting in both formalisms. The POA of a payoff-maximization game is at most 1, the closer to 1 the better.

Theorem 2.1 ([3]) *The POA of every location game is at least $\frac{1}{2}$.*

We next identify three properties possessed by every location game; these are the only properties that our proof of Theorem 2.1 relies on.²

- (P1) For every strategy profile \mathbf{s} , the sum $\sum_{i=1}^k \pi_i(\mathbf{s})$ of players' payoffs (i.e., the net revenue) is at most the social surplus $V(\mathbf{s})$.

This follows from the fact that each $j \in M$ contributes $v_j - d_j(\mathbf{s})$ to the surplus and $d_j^{(2)}(\mathbf{s}) - d_j(\mathbf{s})$ to the payoff of the closest location, and that $d_j^{(2)}(\mathbf{s}) \leq v_j$ by definition.

- (P2) For every strategy profile \mathbf{s} , $\pi_i(\mathbf{s}) = V(\mathbf{s}) - V(\mathbf{s}_{-i})$. That is, a player's payoff is exactly the extra surplus created by the presence of its location.³

To see this property, observe that the contribution of a particular market j to the right-hand side $V(\mathbf{s}) - V(\mathbf{s}_{-i})$ is the extent to which the closest chosen location to j is closer in \mathbf{s} than in \mathbf{s}_{-i} (with the usual upper bound v_j), namely $\min\{v_j, d_j(\mathbf{s}_{-i})\} - \min\{v_j, d_j(\mathbf{s})\}$. This is zero unless player i 's location is the closest to j in \mathbf{s} , in which case it is

$$\min\{v_j, d_j^{(2)}(\mathbf{s})\} - \min\{v_j, d_j(\mathbf{s})\}. \quad (3)$$

Either way, this is precisely market j 's contribution $\pi_{ij}(\mathbf{s})$ to player i 's payoff in \mathbf{s} . Summing over all $j \in M$ proves the property.

- (P3) The function $V(\cdot)$ is *monotone* and *submodular*, as a function of the set of chosen locations. Monotonicity just means that $V(T_1) \leq V(T_2)$ whenever $T_1 \subseteq T_2$; this property is evident from (2). Submodularity is a set-theoretic version of diminishing returns, defined formally as

$$V(T_2 \cup \{\ell\}) - V(T_2) \leq V(T_1 \cup \{\ell\}) - V(T_1) \quad (4)$$

for every location ℓ and subsets $T_1 \subseteq T_2$ of locations (Figure 2).

Submodularity follows immediately from our expression (3) for the surplus increase caused by one new location ℓ — the only interesting case is when ℓ is closer to a market j than any location of T_2 , in which case the bigger the set of locations to which ℓ is being added, the smaller the value of $d_j^{(2)}(\mathbf{s})$ and hence of (3).

Proof of Theorem 2.1: We follow the same four-step outline we used for atomic selfish routing games in Lecture 12 (see Section 1). Let \mathbf{s} denote an arbitrary PNE and \mathbf{s}^* a surplus-maximizing outcome. In the first step, we invoke the PNE hypothesis once per player, with the outcome \mathbf{s}^* providing hypothetical deviations. That is, since \mathbf{s} is a PNE,

$$\pi_i(\mathbf{s}) \geq \pi_i(\mathbf{s}_i^*, \mathbf{s}_{-i}) \quad (5)$$

²Games that possess these three properties are sometimes called *basic utility games* [3].

³Note the similarity to VCG payments (Lecture 7). This equality implies that every location game is a potential game (see Exercises). In our proof of Theorem 2.1, we only need the inequality $\pi_i(\mathbf{s}) \geq V(\mathbf{s}) - V(\mathbf{s}_{-i})$. Games satisfying this inequality and properties (P1) and (P3) are called *valid utility games* [3].




Figure 2: Submodularity: adding ℓ to T_2 yields a lower increase in value than adding ℓ to T_1 because $T_1 \subseteq T_2$.

for every player i . This is the only step of the proof that uses the assumption that \mathbf{s} is a PNE.

The second step is to sum (5) over all the players, yielding

$$V(\mathbf{s}) \geq \sum_{i=1}^k \pi_i(\mathbf{s}) \geq \sum_{i=1}^k \pi_i(s_i^*, \mathbf{s}_{-i}), \quad (6)$$

where the first inequality is property (P1) of location games.

The third step is to disentangle the final term of (6), and relate it to the only two quantities we really care about, $V(\mathbf{s})$ and $V(\mathbf{s}^*)$. By property (P2) of location games, we have

$$\sum_{i=1}^k \pi_i(s_i^*, \mathbf{s}_{-i}) = \sum_{i=1}^k [V(s_i^*, \mathbf{s}_{-i}) - V(\mathbf{s}_{-i})]. \quad (7)$$

To massage the right-hand side into a telescoping sum, we add extra locations to the terms. By submodularity of $V(\cdot)$ (property (P3)), we have

$$V(s_i^*, \mathbf{s}_{-i}) - V(\mathbf{s}_{-i}) \geq V(s_1^*, \dots, s_i^*, \mathbf{s}) - V(s_1^*, \dots, s_{i-1}^*, \mathbf{s}).$$

Thus, the right-hand side of (7) can be bounded below by

$$\sum_{i=1}^k [V(s_1^*, \dots, s_i^*, \mathbf{s}) - V(s_1^*, \dots, s_{i-1}^*, \mathbf{s})] = V(s_1^*, \dots, s_k^*, s_1, \dots, s_k) - V(\mathbf{s}) \geq V(\mathbf{s}^*) - V(\mathbf{s}),$$

where the inequality follows from the monotonicity of $V(\cdot)$ (property (P3)). This completes the third step of the proof.

The fourth and final step of the proof is to solve for the POA. We've proved that

$$V(\mathbf{s}) \geq V(\mathbf{s}^*) - V(\mathbf{s}),$$

so

$$\frac{V(\mathbf{s})}{V(\mathbf{s}^*)} \geq \frac{1}{2}$$

and the POA is at least $\frac{1}{2}$. This completes the proof of Theorem 2.1. ■

3 Smooth Games

There is a general recipe for deriving POA bounds, which includes our analyses of atomic selfish routing games and location games as special cases. There are also many other examples of this recipe that we won't have time to talk about. The goal of formalizing this recipe is not generalization for its own sake; as we'll see, POA bounds established via this recipe are automatically “robust” in several senses.

The following definition is meant to abstract the third “disentanglement” step in the POA upper bound proofs for atomic selfish routing games and location games.

Definition 3.1 (Smooth Games)

1. A cost-minimization game is (λ, μ) -smooth if

$$\sum_{i=1}^k C_i(s_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot \text{cost}(\mathbf{s}^*) + \mu \cdot \text{cost}(\mathbf{s}) \quad (8)$$

for all strategy profiles \mathbf{s}, \mathbf{s}^* .⁴ Here $\text{cost}(\cdot)$ is an objective function that satisfies $\text{cost}(\mathbf{s}) \leq \sum_{i=1}^k C_i(\mathbf{s})$ for every strategy profile \mathbf{s} .⁵

2. A payoff-maximization game is (λ, μ) -smooth if

$$\sum_{i=1}^k \pi_i(s_i^*, \mathbf{s}_{-i}) \geq \lambda \cdot V(\mathbf{s}^*) - \mu \cdot V(\mathbf{s}) \quad (9)$$

for all strategy profiles \mathbf{s}, \mathbf{s}^* . Here $V(\cdot)$ is an objective function that satisfies $V(\mathbf{s}) \geq \sum_{i=1}^k \pi_i(\mathbf{s})$ for every strategy profile \mathbf{s} .⁶

Justifying a new definition requires examples and consequences. We have already seen two examples of classes of smooth games; two consequences are described in the next section. In Lecture 12, we proved that every atomic selfish routing game with affine cost functions is a $(\frac{5}{3}, \frac{1}{3})$ -smooth cost-minimization game. In Section 2, we proved that every location game is a $(1, 1)$ -smooth payoff-maximization game. The conditions (8) and (9) were established in the third, “disentanglement” steps of these proofs. At the time, we had in mind the case where \mathbf{s} and \mathbf{s}^* are a PNE and an optimal outcome of the game, respectively, but the corresponding algebraic manipulations never used those facts and hence apply more generally to all pairs of strategy profiles.

⁴As long as (8) holds for some optimal solution \mathbf{s}^* and all strategy profiles \mathbf{s} , all of the consequences in Section 4 continue to hold.

⁵In atomic selfish routing games, this inequality holds with equality.

⁶This is property (P1) of location games.

4 Robust POA Bounds in Smooth Games

In a (λ, μ) -smooth cost-minimization game with $\mu < 1$, every PNE \mathbf{s} has cost at most $\frac{\lambda}{1-\mu}$ times that of an optimal outcome \mathbf{s}^* . To see this, use the assumption that the objective function satisfies $\text{cost}(\mathbf{s}) \leq \sum_{i=1}^k C_i(\mathbf{s})$, the PNE condition (once per player), and the smoothness assumption to derive

$$\text{cost}(\mathbf{s}) \leq \sum_{i=1}^k C_i(\mathbf{s}) \leq \sum_{i=1}^k C_i(s_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot \text{cost}(\mathbf{s}^*) + \mu \cdot \text{cost}(\mathbf{s}), \quad (10)$$

and rearrange terms. Similarly, every PNE of a (λ, μ) -smooth payoff-maximization game has objective function value at least $\frac{\lambda}{1+\mu}$ times that of an optimal outcome. These are generalizations of our POA bounds of $\frac{5}{2}$ and $\frac{1}{2}$ for atomic selfish routing games with affine cost functions and location games, respectively.

We next describe two senses in which the POA bound of $\frac{\lambda}{1-\mu}$ or $\frac{\lambda}{1+\mu}$ for a (λ, μ) -smooth game is “robust.” For the first, we recall last lecture’s definition of coarse correlated equilibria (CCE).

Definition 4.1 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 $i \in \{1, 2, \dots, k\}$ and every unilateral deviation $s'_i \in S_i$,

$$\mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(s'_i, \mathbf{s}_{-i})]. \quad (11)$$

The equilibrium condition in (11) compares the expected cost of playing according to the distribution σ to that of an unconditional deviation. In Lecture 17, we’ll show that simple and natural learning algorithms drive the time-averaged history of joint play to the set of CCE. In this sense, CCE are a quite tractable set of equilibria, and hence a relatively plausible prediction of realized play. See also Figure 3.

A drawback of enlarging the set of equilibria is that the POA, which is defined via a worst-case equilibrium, can only get worse. In smooth games, however, CCE are a “sweet spot” — permissive enough to be highly tractable, and stringent enough to enable good worst-case bounds.

Theorem 4.2 ([1]) *In every (λ, μ) -smooth cost-minimization game, the POA of CCE is at most $\frac{\lambda}{1-\mu}$.*

That is, the exact same POA bound that we derived in (10) for PNE holds more generally for all CCE. Similarly, in (λ, μ) -smooth payoff-maximization games, the POA bound of $\frac{\lambda}{1+\mu}$ applies more generally to all CCE (the details are left as an exercise). Our POA bounds of $\frac{5}{2}$ and $\frac{1}{2}$ for atomic selfish routing games and location games, respectively, may have seemed specific to PNE at the time, but since the proofs established the stronger smoothness condition (Definition 3.1), these POA bounds hold for all CCE.

Given the definitions, the proof of Theorem 4.2 is not difficult; let’s just follow our nose.




Figure 3: The Venn-diagram of the hierarchy of equilibrium concepts.

Proof of Theorem 4.2: Consider a (λ, μ) -smooth cost-minimization game, a coarse correlated equilibrium σ , and an optimal outcome \mathbf{s}^* . We can write

$$\mathbf{E}_{\mathbf{s} \sim \sigma}[\text{cost}(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s} \sim \sigma} \left[\sum_{i=1}^k C_i(\mathbf{s}) \right] \quad (12)$$

$$= \sum_{i=1}^k \mathbf{E}_{\mathbf{s} \sim \sigma} [C_i(\mathbf{s})] \quad (13)$$

$$\leq \sum_{i=1}^k \mathbf{E}_{\mathbf{s} \sim \sigma} [C_i(s_i^*, \mathbf{s}_{-i})] \quad (14)$$

$$= \mathbf{E}_{\mathbf{s} \sim \sigma} \left[\sum_{i=1}^k C_i(s_i^*, \mathbf{s}_{-i}) \right] \quad (15)$$

$$\leq \mathbf{E}_{\mathbf{s} \sim \sigma} [\lambda \cdot \text{cost}(\mathbf{s}^*) + \mu \cdot \text{cost}(\mathbf{s})] \quad (16)$$

$$= \lambda \cdot \text{cost}(\mathbf{s}^*) + \mu \cdot \mathbf{E}_{\mathbf{s} \sim \sigma} [\text{cost}(\mathbf{s})], \quad (17)$$

where inequality (12) follows from the assumption on the objective function, equalities (13), (15), and (17) follow from linearity of expectation, inequality (14) follows from the definition (11) of a coarse correlated equilibrium (applied once per player i , with the hypothetical deviation s_i^*), and inequality (16) follows from the assumption that the game is (λ, μ) -smooth. Rearranging terms completes the proof. ■

Theorem 4.2 is already quite satisfying — we now have good POA bounds for an equilibrium concept that is guaranteed to exist and is easy to compute. It turns out that smooth games have a number of other nice properties, as well. We conclude this lecture by noting that the POA bound of $\frac{\lambda}{1-\mu}$ or $\frac{\lambda}{1+\mu}$ for a smooth game applies automatically to approximate equilibria, with the bound degrading gracefully as a function of the approximation parameter. For instance, define an ϵ -pure Nash equilibrium (ϵ -PNE) of a cost-minimization game

as a strategy profile \mathbf{s} in which no player can decrease its cost by more than a $(1 + \epsilon)$ factor via a unilateral deviation:

$$C_i(\mathbf{s}) \leq (1 + \epsilon) \cdot C_i(s'_i, \mathbf{s}_{-i}) \quad (18)$$

for every i and $s'_i \in S_i$. Then, the following guarantee holds.

Theorem 4.3 *For every (λ, μ) -smooth cost-minimization game G , every $\epsilon < \frac{1}{\mu} - 1$, every ϵ -PNE \mathbf{s} of G , and every outcome \mathbf{s}^* of G ,*

$$C(\mathbf{s}) \leq \frac{(1 + \epsilon)\lambda}{1 - \mu(1 + \epsilon)} \cdot C(\mathbf{s}^*).$$

The proof of Theorem 4.3 is not difficult and we leave it as an exercise. Similar results hold for smooth payoff-maximization games, and for approximate versions of other equilibrium concepts.

To illustrate Theorem 4.3, consider atomic selfish routing games with affine cost functions, which are $(\frac{5}{3}, \frac{1}{3})$ -smooth. Theorem 4.3 implies that every ϵ -PNE of such a game with $\epsilon < 2$ has expected cost at most $\frac{5+5\epsilon}{2-\epsilon}$ times that of an optimal outcome.

References

- [1] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *41st ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.
- [2] É. Tardos and T. Wexler. Network formation games and the potential function method. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 19, pages 487–516. Cambridge University Press, 2007.
- [3] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.