Smooth Games:

Price of Anarchy first introduced by Koutsoupias & Papadimitriou in STACS 99

Blum, H., Ligett & Roth in STOC 2008 introduced the concept of price of total anarchy for Regret-minimization Algorithms (sort of learning algorithms)

Roughgarden (STOC 2009) introduced the concept of smooth games when he taught price of total anarchy and try to generalize the results as much as he could.

Generalizations of pure Nash equilibria

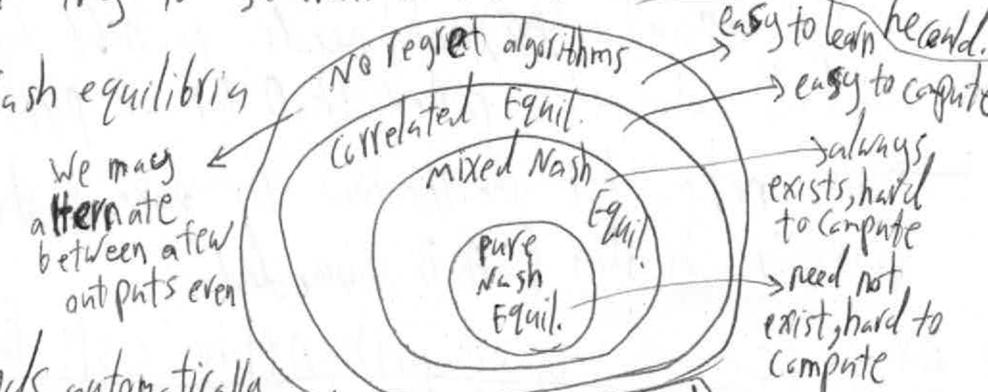

Fig 1

Good news: every bound on the price of anarchy that is derived via a Smoothness argument extends automatically, with no quantitative degration in the bound to all of the more general equilibrium concepts depicted in Figure above.

→ Assume we have n players, each picks a (mixed) strategy s_i and player i incures a cost $C_i(s)$. Assume the objective function is $cost(s) = \sum_i C_i(s)$ (or more generally $C(s) \leq \sum_{i=1}^{k} C_i(s)$)

→ Key definition: A game is (λ,μ)-smooth, if for every pair s, s^* outcomes $(\lambda > 0, \mu < 1)$

$$\sum_i C_i(s_i^*, s_{-i}) \leq \lambda \cdot cost(s^*) + \mu \cdot cost(s) \quad (**)$$

Next we prove:

thm: If a game is (λ,μ)-smooth with $\lambda > 0$ and $\mu < 1$ then each Nash equil. has cost at most $\frac{\lambda}{1-\mu}$ times that of an optimal solution s^*, i.e. $PoA = \frac{\lambda}{1-\mu}$

Proof: Assuming (λ,μ)-smooth, let s be a Nash Equil. and s^* be optimal

$cost(s) \overset{def}{=} \sum_i C_i(s) \overset{\text{since } s \text{ is a Nash}}{\leq} \sum_i C_i(s_i^*, s_{-i}) \overset{\text{def. of smoothness}}{\leq} \lambda \cdot cost(s^*) + \mu \cdot cost(s)$

Thus $cost(s) \leq \frac{\lambda \cdot cost(s^*)}{(1-\mu)}$ and thus $PoA = \frac{\lambda}{1-\mu}$.

Example: for non-atomic congestion games we proved $C(f^{EQ}) \leq C(f^*) + \frac{1}{4} C(f^{EQ})$

Thus $\lambda = 1$ and $\mu = \frac{1}{4}$ and $PoA = \frac{1}{1-\frac{1}{4}} = \frac{4}{3}$ as we have seen.

Why is smoothness stronger? Key point: to derive PoA bound only need $(**)$ to hold in special case where s = a Nash eq. and s^* = optimal but smoothness require $(**)$ for every pair s and s^* outcomes even if s is not a Nash Equil.

Smoothness results so far has been used for non-atomic congestion games, atomic congestion games, weighted congestion games, coordination mechani[sms] and submodular maximization games. ②

→ We can relax definition of smoothness in two ways:

1) by assuming $C(s) \leq \sum_{i=1}^{k} C_i(s)$ (instead of equality)

2) the equation (✪) only needs to hold for some optimal solution s^* and all out comes s, rather than for all pairs (s, s^*) of out comes.

→ Also note that the theorem for pure Nash only says if there IS a pure Nash then PoA is bounded.

Example: Congestion Games with Affine cost functions which are generalizati[on] of non-atomic selfish routing. We have a ground set E of resources, a set of k players with strategy sets $S_1, \dots, S_k \subseteq 2^E$ and a cost function $c_e : Z^+ \to R$ for each $e \in E$.
We always assume that cost functions are non-negative and non-decreasing, in particular when it is affine, i.e., $c_e(x) = a_e x + b_e$ with $a_e, b_e \geq 0$ for every resource $e \in E$.
In selfish routing, E is the edge set and strategies are source-sink pair.
The load induced on edge e is $x_e = |\{i : e \in S_i\}|$, i.e., the number of players using $e \in E$.
The cost of player i, $C_i(s) = \sum_{e \in S_i} c_e(x_e)$, where x is the vector of loads induced by s.
As we showed before (by reversal of sum) $C(s) = \sum_{i=1}^{k} C_i(s) = \sum_{e \in E} c_e(x_e) x_e$.

Thm2: Every congestion game with affine cost functions is $(\frac{5}{3}, \frac{1}{3})$-Smooth.
First a lemma by christodoulou & koutsoupias who noted.

lm 1: $y(z+1) \leq \frac{5}{3} y^2 + \frac{1}{3} z^2$ for all non-negative integers y, z. Thus for all $a, b \geq 0$ and $y, z \geq 0$

$$a y(z+1) + by \leq \frac{5}{3}(a y^2 + by) + \frac{1}{3}(a z^2 + bz) \quad (✪)$$

$\boxed{\text{The right-hand has } \frac{2}{3}by + \frac{1}{3}bz \text{ extra}}$

Now consider a pair s, s^* of outcomes.

To establish smoothness, since the number of player using resource e in the outcome (s_i^*, s_{-i}) is at most one more than in s and this resource contribution is at most x_e^*,

$$\sum_{i=1}^{k} C_i(s_i^*, s_{-i}) \overset{\boxed{X}}{\leq} \sum_{e \in E} (a_e(x_e + 1) + b_e) x_e^* \overset{\text{by } (✪)}{\leq} \sum_{e \in E} \frac{5}{3}(a_e x_e^* + b_e) x_e^* + \sum_{e \in E} \frac{1}{3}(a_e x_e + b_e) x_e = \frac{5}{3} C(s^*) + \frac{1}{3}$$

$\boxed{\begin{matrix} y \leftarrow x_e^* \\ z \leftarrow x_e \end{matrix}}$

Thus by Thm 1, we have upper bound $\frac{5}{2}$ for PoA (and indeed it is tight). Indeed there is an extension Thm for all classes of Fig 1.
for pure Nash

$\boxed{X}$: follows by reversal of sum or double counting

The next example is payoff-maximization games, where player i wants to maximize payoff function $\Pi_i(s)$. We use W to denote objective function of a payoff-maximization game. We call such a game (λ,μ)-smooth if

$$\sum_{i=1}^{k} \Pi_i(s_i^*, s_{-i}) \geq \lambda \cdot W(s^*) - \mu W(s) \quad \text{for every pair } s, s^* \text{ of outcomes.}$$

The same way of minimization, we can show a (λ,μ)-smooth payoff-maximization game, the objective function value of every pure Nash Equil. is at least $\frac{\lambda}{1+\mu}$ fraction of the maximum possible.

Example: <u>Valid Utility Games</u>: We have a ground set E, a non-negative submodular function V defined on subsets of E, and a strategy set $S_i \subseteq 2^E$ and a payoff function Π_i for each player $i=1,2,\dots,k$.

(a function $V: 2^E \to \mathbb{R}$ is submodular if $V(X\cap Y) + V(X \cup Y) \leq V(X) + V(Y)$ for every $X, Y \subseteq E$)

For example, the set E could denote a set of locations where facility can be built, and a strategy $s_i \subseteq E$ could denote the locations at which player i chooses to build facilities.

For an outcome s, let $U(s) \subseteq E$ denote the union $\bigcup_{i=1}^{k} s_i$ of players strategies s.

The objective function value of an outcome s is defined as $W(s) = V(U(s))$

Also we have two extra conditions (i) for each player i, $\Pi_i(s) \geq W(s) - W(\emptyset, s_{-i})$ for every outcome s and (ii) $\sum_{i=1}^{k} \Pi_i(s) \leq W(s)$ for every outcomes [Individual Rationality]

One example of these games is competitive facility location with price-taking markets and profit-maximizing firms by Goemans et al.

Thm 3: Valid utility games with a non decreasing objective function V is $(1,1)$-smooth and hence has (robust) PoA at least $\frac{1}{2}$.

pf: Let s and s^* denote arbitrary outcomes of such a game. Let $Z_i \subseteq E$ denote the union of all of the players' strategies in s, together with the strategies employed by players $1, 2, \dots, i$ in s^*. Then

$$\sum_{i=1}^{k} \Pi_i(s_i^*, s_{-i}) \geq \sum_{i=1}^{k} [V(U(s_i^*, s_{-i})) - V(U(\emptyset, s_{-i}))] \quad \text{(because of (i) condition of Valid. Util. Game)}$$

$$\geq \sum_{i=1}^{k} [V(Z_i) - V(Z_{i-1})] \quad \text{(since } V \text{ is submodular with } X = U(s_i^*, s_{-i}) \text{ and } Y = Z_{i-1}\text{) see App B}$$

$W(s^*) \leq V(Z_k)$

$$\geq W(s^*) - W(s) \quad \text{(since } V \text{ is non-decreasing).}$$

We are done ∎

Extension thm implies the same result for Fig 1 esp. regret minimizing first observed [by] BHLR08

Appendix

A) Note on Regret Minimizing Algorithms, (first introduced in 1950 by Hannan)

In these algorithms at each time the agent selects an action and observes the loss/gain and the goal is to minimize loss or maximize gain, We have often stochastic vs adversarial inputs and we consider the fixed optimum to compare.
There are algorithms with regret approching zero at a rate of $O(\frac{1}{\sqrt{T}})$.
These algorithms are learning algorithms.

B) by submodularity:
$V(X) - V(X \cap Y) \geq V(X \cup Y) - V(Y)$ Thus $V(U_i(s_i^*, S_{-i})) - V(U(\emptyset, S_{-i}))$

$$\geq V(\underbrace{U(s^*_{1,\dots,i}, S_{-i})}_{Z_i}) - V(\underbrace{U(s^*_{1,\dots,i-1}, S_{-i})}_{Z_{i-1}})$$