Intrinsic Robustness of the Price of Anarchy

Tim Roughgarden
Stanford University

Some Context

Obvious fact: many modern applications in CS involve autonomous, self-interested agents

- motivates noncooperative games as modeling tool

Unsurprising fact: equilibria of noncooperative games typically inefficient

- i.e., don't optimize natural objective functions
- e.g., *Nash equilibrium*: an outcome such that no player better off by switching strategies

Price of anarchy: quantify inefficiency w.r.t some objective function.

Price of Anarchy

Definition: *price of anarchy (POA)* of a game (w.r.t. some objective function):

$$\frac{\text{equilibrium objective fn value}}{\text{optimal obj fn value}}$$

← the closer to 1 the better

Well-studied goal: when is the POA small?

- benefit of centralized control is small
- can suggest engineering rules of thumb: [Roughgarden STOC 02]: *10% extra network capacity guarantees POA for network routing ≤ 2*

The Price of Anarchy

Network w/2 players:

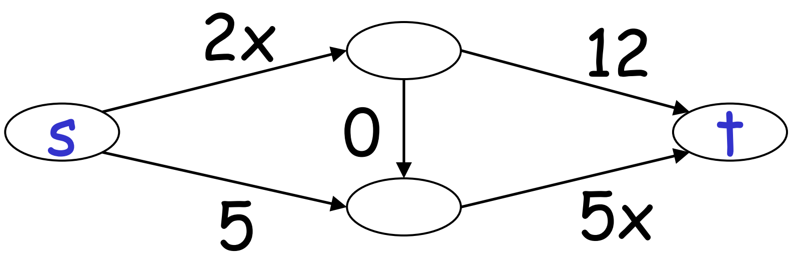

The Price of Anarchy

Nash Equilibrium:

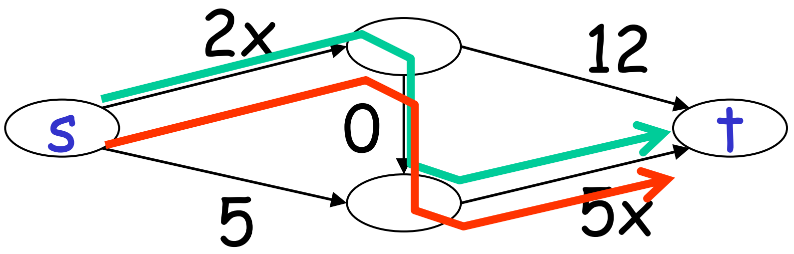

cost = 14+14 = 28

The Price of Anarchy

Nash Equilibrium:

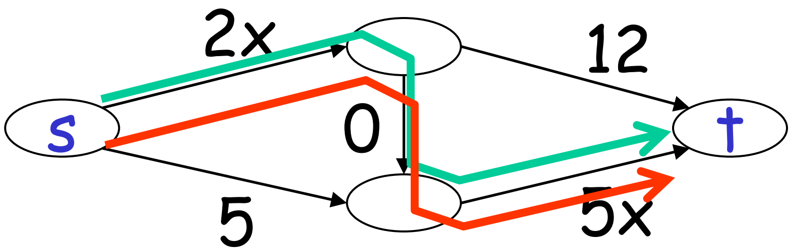

cost = 14+14 = 28

To Minimize Cost:

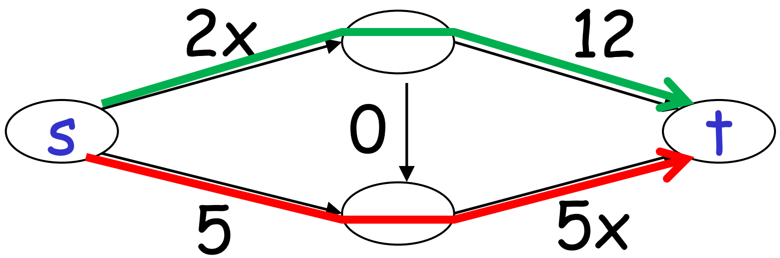

cost = 14+10 = 24

Price of anarchy = 28/24 = 7/6.

- if multiple equilibria exist, look at the *worst* one

Key Points

- main definition: a "canonical way" to bound the price of anarchy (for pure equilibria)
- theorem 1: every POA bound proved "canonically" is *automatically far stronger*
  - e.g., even applies "out-of-equilibrium", assuming no-regret play
- theorem 2: canonical method provably yields optimal bounds in fundamental cases

Abstract Setup

- n players, each picks a strategy s_i
- player i incurs a cost $C_i(\mathbf{s})$

Important Assumption: objective function is $cost(\mathbf{s}) := \Sigma_i C_i(\mathbf{s})$

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

$$\Sigma_i C_i(s^*_i, s_{-i}) \leq \lambda \bullet cost(\mathbf{s}^*) + \mu \bullet cost(\mathbf{s}) \quad [(*)]$$

Smooth => POA Bound

Next: "canonical" way to upper bound POA (via a smoothness argument).

- notation: $\mathbf{s}$ = a Nash eq; $\mathbf{s}^*$ = optimal

Assuming (λ,μ)-smooth:

$$\begin{aligned} \text{cost}(\mathbf{s}) &= \Sigma_i C_i(\mathbf{s}) && \text{[defn of cost]} \\ &\leq \Sigma_i C_i(s^*_i, s_{-i}) && \text{[}\mathbf{s}\text{ a Nash eq]} \\ &\leq \lambda \bullet \text{cost}(\mathbf{s}^*) + \mu \bullet \text{cost}(\mathbf{s}) && \text{[(*)]} \end{aligned}$$

Then: POA (of pure Nash eq) $\leq \lambda/(1-\mu)$.

Why Is Smoothness Stronger?

Key point: to derive POA bound, only needed

$\Sigma_i C_i(s^*_i, s_{-i}) \leq \lambda \bullet cost(\mathbf{s}^*) + \mu \bullet cost(\mathbf{s})$ [(*)]

to hold in special case where **s** = a Nash eq and **s*** = optimal.

Smoothness: requires (*) for *every* pair **s**,**s*** outcomes.

- even if s is *not* a pure Nash equilibrium

Example Application

Definition: a sequence $s^1, s^2, \ldots, s^T$ of outcomes is *no-regret* if:

- for each player i, each fixed action q_i:
  - average cost player i incurs over sequence no worse than playing action q_i every time
  - simple hedging strategies can be used by players to enforce this (for suff large T)

Theorem: in a (λ,μ)-smooth game, average cost of every no-regret sequence at most $[\lambda/(1-\mu)]$ x cost of optimal outcome.

Why Important?

- bound on no-regret sequences implies bound on inefficiency of mixed and correlated equilibria
- bound applies even to sequences that don't converge in any sense
  - no regret much weaker than reaching equilibrium
  - [Blum/Even-Dar/Ligett PODC 06], [Blum/Hajiaghayi/Ligett/Roth STOC 08]

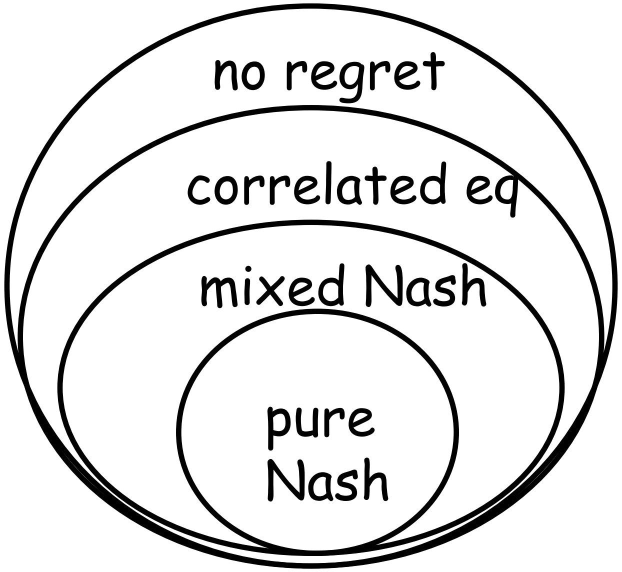

Smooth => POTA Bound

- notation: $s^1, s^2, \ldots, s^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ,μ)-smooth:

$\Sigma_t \text{cost}(\mathbf{s}^t) = \Sigma_t \Sigma_i C_i(\mathbf{s}^t)$ [defn of cost]

Smooth => POTA Bound

- notation: $s^1, s^2, \ldots, s^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ,μ)-smooth:

$\Sigma_t \text{cost}(\mathbf{s}^t) = \Sigma_t \Sigma_i C_i(\mathbf{s}^t)$ [defn of cost]

$= \Sigma_t \Sigma_i [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}]$ $[\Delta_{i,t} := C_i(\mathbf{s}^t) - C_i(s^*_i, s^t_{-i})]$

Smooth => POTA Bound

- notation: $s^1, s^2, \ldots, s^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ,μ)-smooth:

$\Sigma_t \, cost(\mathbf{s}^t) = \Sigma_t \Sigma_i \, C_i(\mathbf{s}^t)$ [defn of cost]

$= \Sigma_t \Sigma_i \, [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}]$ [$\Delta_{i,t} := C_i(\mathbf{s}^t) - C_i(s^*_i, s^t_{-i})$]

$\leq \Sigma_t \, [\lambda \bullet cost(\mathbf{s}^*) + \mu \bullet cost(\mathbf{s}^t)] + \Sigma_i \Sigma_t \, \Delta_{i,t}$ [(*)]

Smooth => POTA Bound

- notation: $s^1, s^2, \ldots, s^T$ = no regret; $\mathbf{s}^*$ = optimal

Assuming (λ,μ)-smooth:

$\Sigma_t \, cost(\mathbf{s}^t) = \Sigma_t \Sigma_i \, C_i(\mathbf{s}^t)$ [defn of cost]

$= \Sigma_t \Sigma_i \, [C_i(s^*_i, s^t_{-i}) + \Delta_{i,t}]$ $[\Delta_{i,t} := C_i(\mathbf{s}^t) - C_i(s^*_i, s^t_{-i})]$

$\leq \Sigma_t \, [\lambda \bullet cost(\mathbf{s}^*) + \mu \bullet cost(\mathbf{s}^t)] + \Sigma_i \Sigma_t \, \Delta_{i,t}$ [(*)]

No regret: $\Sigma_t \, \Delta_{i,t} \leq 0$ for each i.

To finish proof: divide through by T.

Further Applications

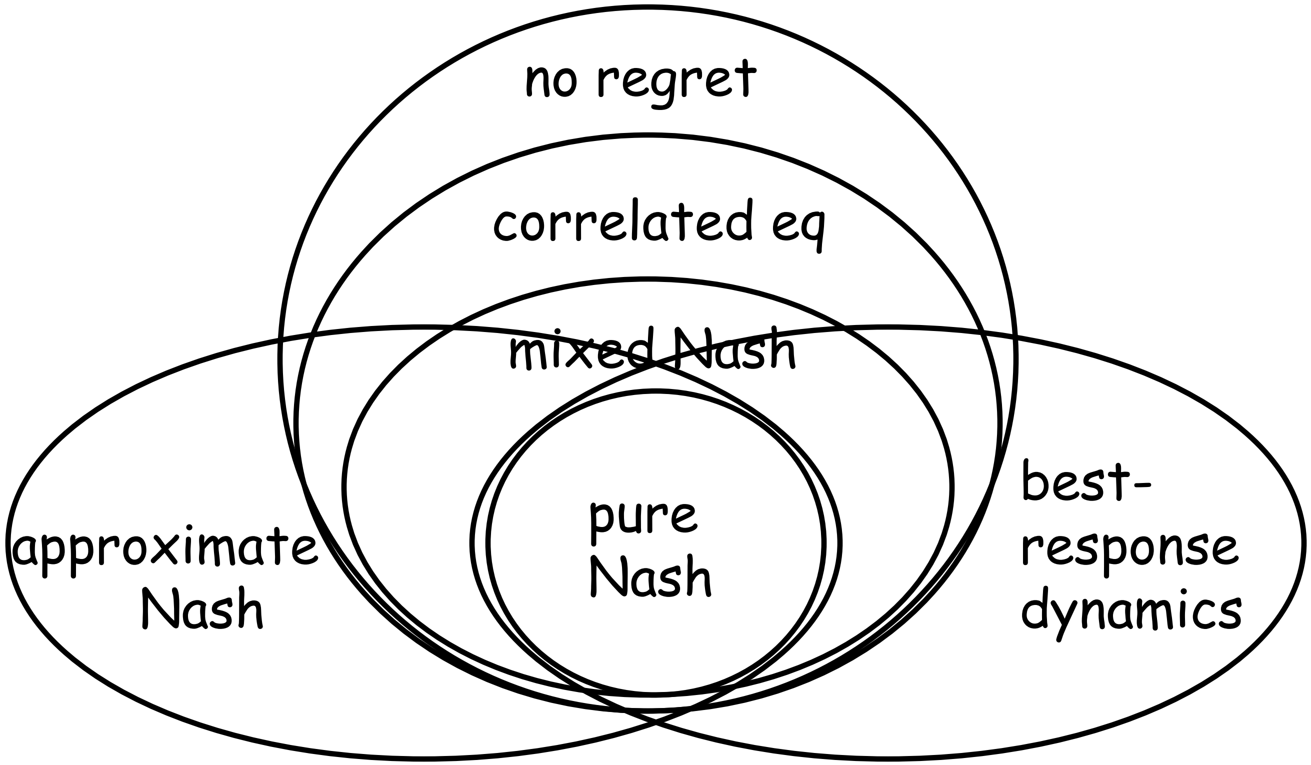

Theorem: in a (λ,μ)-smooth game, everything in these sets costs (essentially) λ/(1-μ) x OPT.

Some Smoothness Bounds

Examples: selfish routing, linear cost fns.

- every nonatomic game is (1,1/4)-smooth
  - implicit in [Roughgarden/Tardos 00]
  - less implicit in [Correa/Schulz/Stier Moses 05]
  - implies bound of 4/3 (tight even for pure eq)
- every atomic game is (5/3,1/3)-smooth
  - follows directly from analysis in [Awerbuch/Azar/Epstein 05], [Christodoulou/Koutsoupias 05]
  - implies bound of 5/2 (tight even for pure eq)

Tight Game Classes

Theorem: for every set C, congestion games with cost functions restricted to C are *tight*:

maximum [pure POA] = minimum [λ/(1-μ)]

congestion games w/cost functions in C

(λ ,μ): all such games are (λ ,μ)-smooth

Corollaries

Corollary 1: first characterization of "universal worst-case congestion games" in the atomic case.

- analog of "Pigou-like (2-node, 2-link) networks are the worst" in nonatomic case [Roughgarden 03]
- here: "2 parallel cycles always suffice"
  - and are generally necessary for minimal worst-case examples

Corollary 2: first (tight) POA bounds for (atomic) congestion games with general cost functions.

- previous exact bounds for polynomials +w/nonnegative coefficients: [Aland et al 06], [Olver 06]

Wrap-Up

Summary: the most common way of proving POA bounds automatically yields a much more robust guarantee

- and this technique often gives tight bounds

Ongoing work: weighted congestion games [with Bhawalkar & Gairing]

- splittable congestion games [with Schoppman]
- "inexpressive" auctions [with Bhawalkar]
- limitations of smoothness [with Nadav]