Approximation in Algorithm Game Theory: The Price of Anarchy

Tim Roughgarden
Stanford University

Pigou's Example

Example: one unit of traffic wants to go from s to t

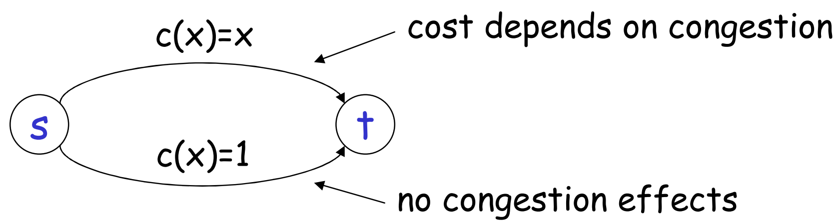

Question: what will selfish network users do?

- assume everyone wants smallest-possible cost
- [Pigou 1920]

Motivating Example

Claim: all traffic will take the top link.

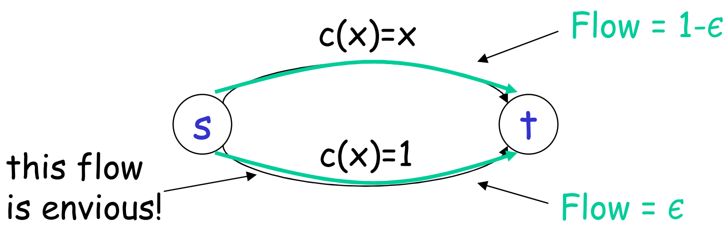

Reason:

- $\epsilon > 0 \Rightarrow$ traffic on bottom is envious
- $\epsilon = 0 \Rightarrow$ equilibrium
  - all traffic incurs one unit of cost

Can We Do Better?

Consider instead: traffic split equally

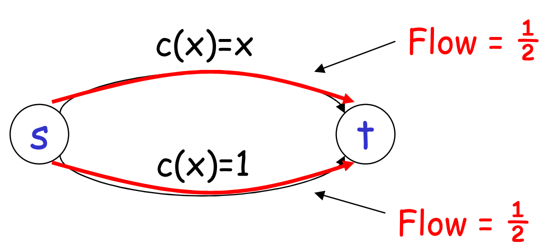

Improvement:

- half of traffic has cost 1 (same as before)
- half of traffic has cost $\frac{1}{2}$ (much improved!)

Braess's Paradox

Initial Network:

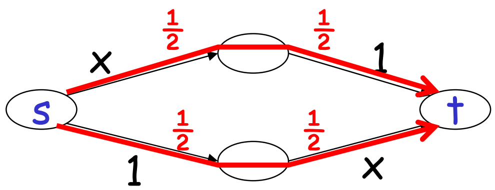

Cost = 1.5

Braess's Paradox

Initial Network:

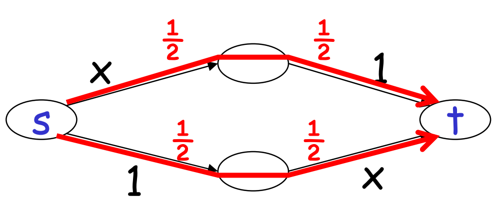

Cost = 1.5

Augmented Network:

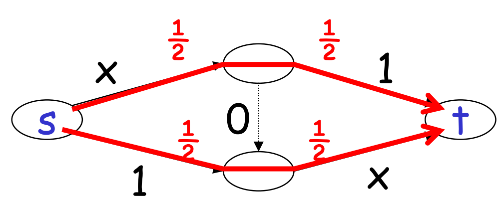

Now what?

Braess's Paradox

Initial Network:

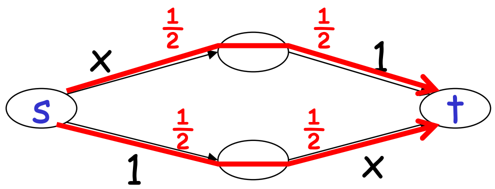

Cost = 1.5

Augmented Network:

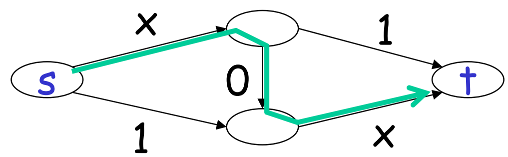

Cost = 2

Braess's Paradox

Initial Network:

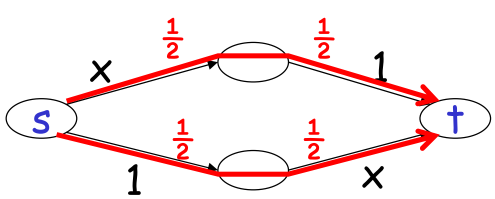

Cost = 1.5

Augmented Network:

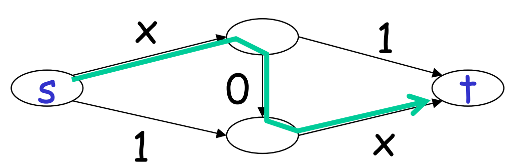

Cost = 2

All traffic incurs more cost! [Braess 68]

- also has physical analogs [Cohen/Horowitz 91]

High-Level Overview

Motivation: equilibria of noncooperative network games typically inefficient

- e.g., Pigou's example + Braess's Paradox
- don't optimize natural objective functions

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

Our goal: when is the price of anarchy small?

- when does competition approximate cooperation?
- benefit of centralized control is small

Nonatomic Selfish Routing

- directed graph $G = (V,E)$
- source-destination pairs $(s_1,t_1), \ldots, (s_k,t_k)$
- r_i = amount of traffic going from s_i to t_i
- for each edge e, a cost function $c_e(\cdot)$
  - assumed continuous and nondecreasing

Defn: a multicommodity flow is an *equilibrium* if all traffic routed on shortest paths.

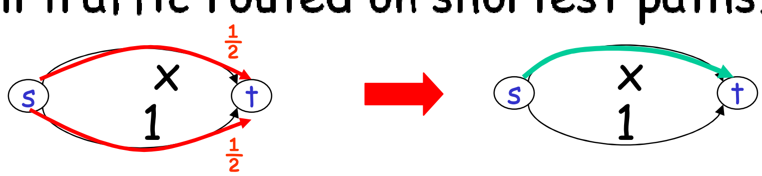

Our Objective Function

Definition of social cost: total cost $C(f)$ incurred by the traffic in a flow f.

Formally: if $c_P(f)$ = sum of costs of edges of P (w.r.t. flow f), then:

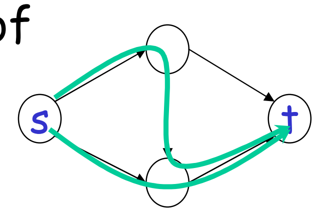

$$C(f) = \Sigma_P f_P \cdot c_P(f)$$

Our Objective Function

Definition of social cost: total cost $C(f)$ incurred by the traffic in a flow f.

Formally: if $c_P(f)$ = sum of costs of edges of P (w.r.t. flow f), then:

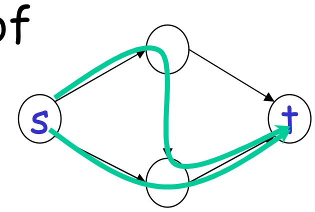

$$C(f) = \Sigma_P f_P \cdot c_P(f)$$

Example:

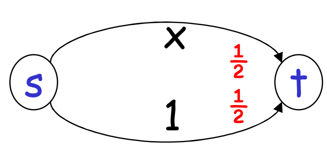

Cost = $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = \frac{3}{4}$

The Price of Anarchy

Defn: price of anarchy of a game = $\frac{\text{obj fn value of worst equilibrium}}{\text{optimal obj fn value}}$

- definition from [Koutsoupias/Papadimitriou 99]

Example: POA = 4/3 in Pigou's example

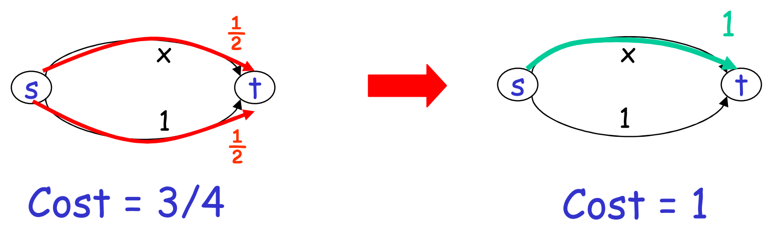

A Nonlinear Pigou Network

Bad Example: (d large)

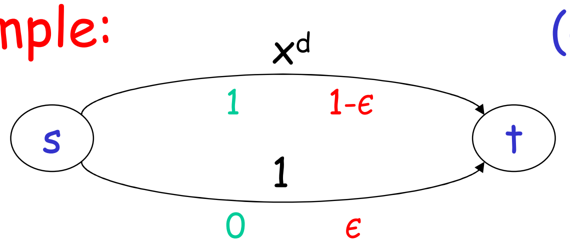

equilibrium has cost 1, min cost ≈ 0

A Nonlinear Pigou Network

Bad Example: (d large)

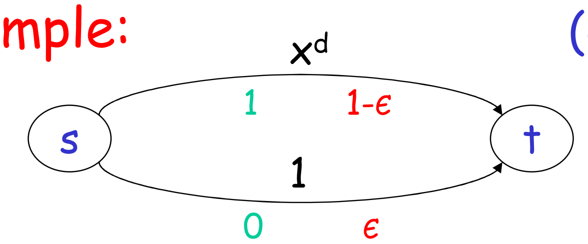

equilibrium has cost 1, min cost ≈ 0

⇒ price of anarchy unbounded as d -> infinity

Goal: weakest-possible conditions under which P.O.A. is small.

When Is the Price of Anarchy Bounded?

Examples so far:

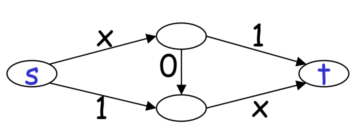

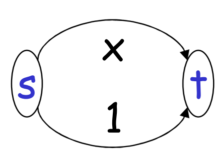

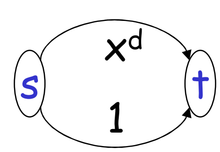

Hope: imposing additional structure on the cost functions helps

- worry: bad things happen in larger networks

Polynomial Cost Functions

Def: linear cost fn is of form $c_e(x)=a_e x+b_e$

Theorem: [Roughgarden/Tardos 00] for every network with linear cost functions:

cost of Nash flow $\leq$ 4/3 × cost of opt flow

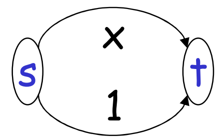

Polynomial Cost Functions

Def: linear cost fn is of form $c_e(x)=a_e x+b_e$

Theorem: [Roughgarden/Tardos 00] for every network with linear cost functions:

cost of Nash flow $\leq$ 4/3 × cost of opt flow

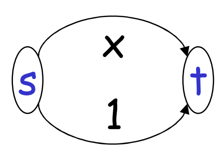

Bounded-deg polys: (w/nonneg coeffs) replace 4/3 by $\approx d/\ln d$

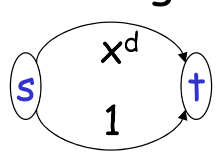

tight example

A General Theorem

Thm: [Roughgarden 02], [Correa/Schulz/Stier Moses 03] fix any set of cost fns. Then, a Pigou-like example 2 nodes, 2 links, 1 link w/constant cost fn) achieves worst POA

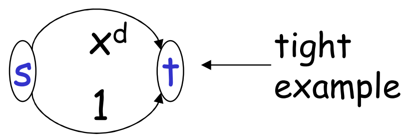

Pigou Bound

Recall goal: want to show Pigou-like examples are always worst cases.

Pigou bound: given set of cost functions (e.g., degree-d polys), largest POA in a network:

- two nodes, two links
- one function in given set
- one constant function
  - constant = cost of fully congested top edge

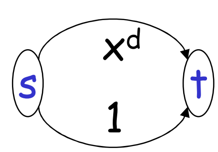

Pigou Bound

Defn: the Pigou bound $\alpha(S)$ for S is:

$$\max \frac{r \cdot c(r)}{y \cdot c(y) + (r-y) \cdot c(r)}$$

- max is over all choices of cost fns c in S, traffic rate $r \geq 0$, flow $y \geq 0$

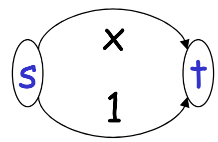

- choose $c(x) = x$; $r = 1$; $y = 1/2 \Rightarrow$ get 4/3
- calculus: $\alpha(S) = 4/3$ when S = affine functions [d/ln d for deg-d polynomials]

Main Theorem (Formally)

Theorem: [Roughgarden 02, Correa/Schulz/Stier Moses 03]: For every set S, for every selfish routing network G with cost functions in C, the POA in G is at most $\alpha(S)$.

- POA always maximized by Pigou-like examples

That is, if f and f^* are Nash + optimal flows in G, then $C(f)/C(f^*) \leq \alpha(S)$.

- example: POA $\leq$ 4/3 if G has affine cost fns

Interpretation

Bad news: inefficiency of selfish routing grows as cost functions become "more nonlinear".

- think of "nonlinear" as "heavily congested"
- recall nonlinear Pigou's example

Good news: inefficiency does not grow with network size or # of source-destination pairs.

- in lightly loaded networks, no matter how large, selfish routing is nearly optimal

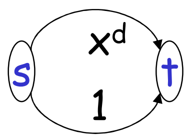

tight example

Benefit of Overprovisioning

Suppose: network is overprovisioned by $\beta > 0$ (β fraction of each edge unused).

Then: Price of anarchy is at most $\frac{1}{2}(1+1/\sqrt{\beta})$.

- arbitrary network size/topology, traffic matrix

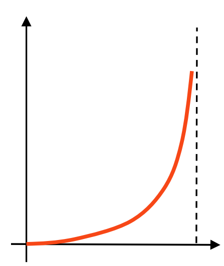

Moral: Even modest (10%) over-provisioning sufficient for near-optimal routing.

Variational Inequality

Claim:

- if f is a Nash flow and f^* is feasible, then

$$\Sigma_e f_e \cdot c_e(f_e) \leq \Sigma_e f_e^* \cdot c_e(f_e)$$

- proof: use that Nash flow routes flow on shortest paths (w.r.t. costs $c_e(f_e)$)

Variational Inequality

Claim:

- if f is a Nash flow and f^* is feasible, then

$$\Sigma_e f_e \cdot c_e(f_e) \leq \Sigma_e f_e^* \cdot c_e(f_e)$$

- proof: use that Nash flow routes flow on shortest paths (w.r.t. costs $c_e(f_e)$)

Thus: $\Sigma_e f_e \cdot c_e(f_e) \leq \Sigma_e f_e^* \cdot c_e(f_e^*) +$

$\Sigma_e f_e^* \cdot [c_e(f_e) - c_e(f_e^*)]$

relation to $C(f)$? →

Geometry of Affine Case

Assume: $c_e(x) = a_e x + b_e$

Goal: compare

$$f_e^* \cdot [c_e(f_e) - c_e(f_e^*)] \text{ vs. } f_e \cdot c_e(f_e)$$

Interesting case: when $c_e(f_e) > c_e(f_e^*)$:

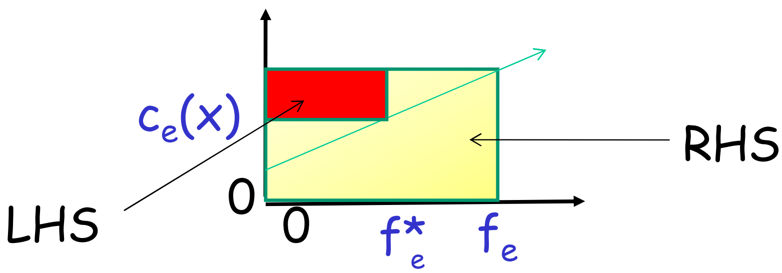

POA = 4/3 for Affine Costs

Assume: $c_e(x) = a_e x + b_e$

Thus: $f_e^* \cdot [c_e(f_e) - c_e(f_e^*)] \leq [f_e \cdot c_e(f_e)]/4$

Thus: $\Sigma_e f_e \cdot c_e(f_e) \leq \Sigma_e f^* \cdot c_e(f^*) +$

$$\Sigma_e f_e^* \cdot [c_e(f_e) - c_e(f_e^*)]$$

$$\leq \Sigma_e [f_e \cdot c_e(f_e)]/4$$

Thus: $C(f) \leq 4/3 \cdot C(f^*)$

- proof from [Correa/Schulz/Stier Moses 08]

Atomic Selfish Routing

Atomic networks: each of (finitely many) players picks a path on which to route one unit of traffic. (otherwise identical model)

AAE example: [refer to whiteboard for details] shows that the POA can be as high as 2.5 in this model, with affine cost functions.

Potential Functions

So: potential fn tracks deviations by players

Thus: equilibria of game = local optima of Φ

- so finite potential games have pure-strategy Nash equilibria (proof: just do "best-response dynamics") [Monderer/Shapley 96]
  - precursors: [Rosenthal 73], [Beckmann et al 56]

Claim: every atomic selfish routing game has a potential function.

Proof of Potential Function

Define $\Phi_e(k_e) = c_e(1) + c_e(2) + c_e(3) + \ldots + c_e(k_e)$

where k_e is # players using e.

Let $\Phi(S) = \sum_{e \in S} \Phi_e(S)$

Consider some solution S (a path for each player).
Suppose player i is unhappy and decides to deviate.
What happens to $\Phi(S)$?

Proof of Potential Function

$\Phi_e(k_e) = c_e(1) + c_e(2) + \ldots + c_e(k_e)$

Suppose player i's new path includes e.

i pays $c_e(k_e+1)$ to use e.

$\Phi_e(k_e)$ increases by the same amount.

If player i leaves an edge e', $\Phi_{e'}(k_e)$ exactly reflects the change in i's cost.

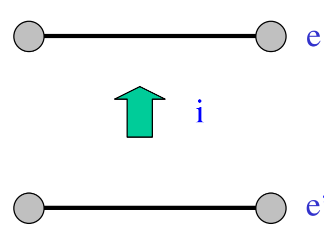

Consequences for the Price of Anarchy?

Example: linear cost functions.

Compare cost + potential function:

$C(f) = \Sigma_e f_e \cdot c_e(f_e) = \Sigma_e [a_e f_e^2 + b_e f_e]$
$\Phi(f) = \Sigma_e \Sigma_i c_e(i)dx \approx \Sigma_e [(a_e f_e^2)/2 + b_e f_e]$

- cost, potential fn differ by factor of ≤ 2
- gives upper bound of 2 on price on anarchy?
  - $C(f) \leq 2\times\Phi(f) \leq 2\times\Phi(f^*) \leq 2\times C(f^*)$

POA in Atomic Model

Catch: only bounds the cost of the *global* potential fn minimizer, not all Nash equilibria (≈ *local* minimizers).

Instead: use variational inequality, modified for the atomic case:

$$\Sigma_e f_e \cdot c_e(f_e) \leq \Sigma_e f_e^* \cdot c_e(f_e+1)$$

A Technical Lemma

Claim:

- [Christodoulou/Koutsoupias 05]: for all integers y,z:
  $y(z+1) \leq (5/3)y^2 + (1/3)z^2$

A Technical Lemma

Claim:

- [Christodoulou/Koutsoupias 05]: for all integers y,z:
  $y(z+1) \leq (5/3)y^2 + (1/3)z^2$

- so: $ay(z+1) + by \leq (5/3)[ay^2 + by] + (1/3)[az^2 + bz]$
  - for all integers y,z and $a,b \geq 0$

A Technical Lemma

Claim:

- [Christodoulou/Koutsoupias 05]: for all integers y,z:
  $y(z+1) \le (5/3)y^2 + (1/3)z^2$
- so: $ay(z+1) + by \le (5/3)[ay^2 + by] + (1/3)[az^2 + bz]$
  - for all integers y,z and $a,b \ge 0$
- so: $\Sigma_e [a_e(f_e+1) + b_e)f_e^*] \le (5/3)\, \Sigma_e [(a_e f_e^* + b_e)f_e^*]$
  $+ (1/3)\, \Sigma_e [(a_e f_e + b_e)f_e]$

A Technical Lemma

Claim:

- [Christodoulou/Koutsoupias 05]: for all integers y,z:
  $y(z+1) \le (5/3)y^2 + (1/3)z^2$
- so: $ay(z+1) + by \le (5/3)[ay^2 + by] + (1/3)[az^2 + bz]$
  - for all integers y,z and $a,b \ge 0$
- so: $\Sigma_e [a_e(f_e+1) + b_e)f_e^*] \le (5/3)\, \Sigma_e [(a_e f_e^* + b_e)f_e^*]$
  $+ (1/3)\, \Sigma_e [(a_e f_e + b_e)f_e]$
- so: $C(f) \le \Sigma_e f^* \cdot c_e(f_e+1) \le (5/3)C(f^*) + (1/3)C(f)$
- so: $POA \le 5/2$

Key Points of the Day

Key Examples: Pigou's example + nonlinear variant; Braess's Paradox; AAE example

- POA depends on cost fns, nonatomic vs. atomic
  - but proofs "structurally" the same (more on Friday)

Key proof techniques: (1) equilibrium satisfies a "variational inequality", which can be related to Nash & OPT costs; (2) parameterize POA bounds via "universal worst-case examples"

Also: equilibria (locally) minimize potential

Approximation in Algorithmic Game Theory: Revenue-Maximizing Auctions

Tim Roughgarden (Stanford)

Example: Single-Good Auctions

Assume: 1 good, n bidders, bidder i has "valuation" v_i for good [like eBay]

- v_i = maximum willingness to pay

Design space: given sealed bids, pick:

- (1) a winner; and (2) a selling price.

Example: first-price auction.

- winner = highest bidder; price = winner's bid
- how would you bid in this auction?

Example: Posted Price

Assume: 1 good, 1 bidder with valuation v

- utility = v - price paid; or 0 if no sale

Posted price: seller compares the bid b to a "take-it-or-leave it" offer at some fixed price p.

Note: truthful bidding (b = v) is "foolproof"

- i.e., a false bid never outperforms a true bid
- case 1: (v ≤ p) max utility = 0, achieved when b = v
- case 2: (v ≥ p) max utility = v-p, achieved when b = v
- called a *truthful* auction or "mechanism"

Example: The Vickrey Auction

Second-Price (or Vickrey) Auction [Vickrey 61]:

- winner = highest bidder; price = *2nd-highest* bid
- note: truthful bidding ($b_i = v_i$) is "foolproof"

Proof: each bidder i effectively faces posted price p_i = highest bid by someone else.

k-Vickrey: With *k* copies of a good: winners = top k bidders; all pay (k+1)-highest price.

Variation: can add a *reserve price* (≈ a dummy bidder).

Auction Benchmarks

Goal: prove results of the form (for revenue):

"Theorem: auction A is (approximately) optimal."

Auction Benchmarks

Goal: prove results of the form (for revenue):

"Theorem: auction A is (approximately) optimal."

Auction model: focus on multi-item auctions

- n bidders, k identical goods (mostly k = n)
- *allocation rule:* b_i's $\longrightarrow$ x_i's (probability of winning)
- *payment rule:* b_i's $\longrightarrow$ p_i's [require $0 \le p_i \le b_i x_i$]
- truthful (i.e., truthful bidding [b_i=v_i] dominant)
  - equivalent: each i faces bid-independent posted price

Auction Benchmarks (con'd)

Goal: prove results of the form:

"Theorem: for every valuation profile v:
auction A's revenue on v is at least OPT(v)/α."
(for a hopefully small constant α)

Auction Benchmarks (con'd)

Goal: prove results of the form:

"Theorem: for every valuation profile v:
auction A's revenue on v is at least OPT(v)/α."
(for a hopefully small constant α)

Idea for OPT(v): sum of k largest v_i's.

Problem: too strong, not useful.

- makes all auctions A look equally bad.
- every A has a bad v [no constant α possible]

The Obvious Idea Fails

Claim: no auction always has revenue at least a constant fraction of the sum of k largest v_i's.

Proof sketch: by probabilistic method. Take k = n.

- pick each v_i i.i.d. from distribution with CDF $F(z) = 1\text{-}1/z$ on $[1,\infty)$ [density $f(z) = 1/z^2$]
- expected revenue of every posted price $p_i \geq 1$ for bidder i = $p_i\,[1\text{-}(1\text{-}1/p_i)] = 1$.
- expected revenue of every auction: $\leq n$
- expected sum of v_i's: unbounded

The Fixed Price Benchmark

Solution: [Goldberg/Hartline/Karlin/Saks/Wright GEB 06]

- define OPT(v) := best *fixed-price* revenue:

$$F^{(2)}(v) := \max_{2 \leq i \leq k} i v_i \qquad \text{(assume sorted } v_i\text{'s)}$$

Justification?: for now, "seems to work".

- α-competitive auctions exist for small α
  - will prove today with $\alpha = 4$
- no auction has α smaller than 1 (or even 2.42)

Friday: a fundamental explanation of why it works.

Auction Benchmarks (con'd)

Goal: prove results of the form:

"Theorem: for every valuation profile v:
auction A's revenue on v is at least $F^{(2)}(v)/\alpha$."
(for a hopefully small constant α)

Note: the Vickrey auction achieves α=2 when n=2.

- but no constant factor for larger n
- to de better: need a more "operational" understanding of truthful auctions (next)

Two Definitions

Implementable Allocation Rule: is a function x (from bids to winners/losers) that admits a payment rule p such that (x,p) is truthful.

- i.e., truthful bidding $[b_i:=v_i]$ always maximizes a bidder's (expected) utility

Monotone Allocation Rule: for every fixed bidder i, fixed other bids b_{-i}, probability of winning only increases in the bid b_i.

- example: highest bidder wins
- non-example: 2nd-highest bidder wins

Myerson's Lemma

Myerson's Lemma: [1981; also Archer-Tardos 01] an allocation rule x is implementable if and only if it is monotone.

Myerson's Lemma

Myerson's Lemma: [1981; also Archer-Tardos 01] an allocation rule x is implementable if and only if it is monotone.

Moreover: for every monotone allocation rule x, there is a unique payment rule p such that (x,p) is truthful and losers always pay 0.

Myerson's Lemma

Myerson's Lemma: [1981; also Archer-Tardos 01] an allocation rule x is implementable if and only if it is monotone.

Moreover: for every monotone allocation rule x, there is a unique payment rule p such that (x,p) is truthful and losers always pay 0.

Explicit formula for $p_i(b)$:

- keep b_{-i} fixed, increase z from 0 to b_i
- consider breakpoints $y_1,\ldots,y_q$ at which x_i jumps
- set $p_i(b) := \Sigma_j\, y_j \bullet$ [jump in x_i at y_j]

Myerson's Lemma (Proof)

Proof of "if" direction: let x be monotone, fix i and b_{-i}. Write x(z), p(z) for $x_i(z, b_{-i})$, $p_i(z, b_{-i})$.

Swapping trick: if (x,p) is truthful, p satisfies:

- [take true value = z, false bid = z + ε]:
  $z \circ x(z) - p(z) \geq z \circ x(z + \varepsilon) - p(z + \varepsilon)$
- [take true value = z + ε, false bid = z]:
  $(z + \varepsilon) \circ x(z + \varepsilon) - p(z + \varepsilon) \geq (z + \varepsilon) \circ x(z) - p(z)$

Thus: $p(z + \varepsilon) - p(z)$ lies between $z \circ [x(z + \varepsilon) - x(z)]$ and $(z + \varepsilon) \circ [x(z + \varepsilon) - x(z)]$

Myerson's Lemma (Proof con'd)

The story so far: $p(z + \varepsilon) - p(z)$ lies between

$z \circ [x(z + \varepsilon) - x(z)]$ and $(z + \varepsilon) \circ [x(z + \varepsilon) - x(z)]$

So: taking ε to zero get,

- $p'(z) = z \circ x'(z)$ [if x differentiable at z] or
- jump in p at $z = z \circ$ [jump in x at z]

Integrating from 0 to b_i, get:

$$p_i(b) := \Sigma_j\, y_j \bullet \text{[jump in } x_i \text{ at } y_j]$$

To finish proof (exercise): verify that auction (x,p) is truthful if and only if x monotone.

A Competitive Auction

Theorem: [Fiat/Goldberg/Hartline/Karlin 02]

There is a randomized auction for n-bidder n-item auctions that, for every input v, has expected revenue at least $F^{(2)}/4$.

- works also for $k \geq 2$ items, see exercises

Recall:

$$F^{(2)}(v) := \max_{2 \leq i \leq k} i v_i \qquad \text{(assume sorted } v_i\text{'s)}$$

Subroutine: Profit Extractor

Given: (truthful) bids v + revenue target R.:

- initialize S = all bidders
- while there is an i in S such that $v_i < R/|S|$:
  - remove such a bidder from S
- return final set S and charge all winners (if any) a price of $p = R/|S|$

Note: allocation rule is monotone; prices are correct (p = min bid s.t. a winner still wins)

- => truthful by Myerson's Lemma

Profit Extractor (con'd)

Claim: ProfitExtract has revenue R if there is a common posted price that extracts R; and has revenue 0 otherwise.

Proof Sketch: 2nd statement is clear (ProfitExtract only uses common posted prices).

For 1st statement: suppose common posted price p works, define $T = \{ i \mid v_i \geq p \}$. All such bidders can pay $R/|T|$. Inductively, no bidder of T ever gets deleted.

The RSPE Auction

- collect (truthful) bid v
- randomly partition v = (x,y)
  [each bidder placed 50%/50%, independently]
- let R_1 = max revenue from x via common posted price;
  R_2 = max revenue from y via common posted price
- ProfitExtract(x) with revenue target R_2
- ProfitExtract(y) with revenue target R_1

Example: n = 2, v = (1,1/2)

- $F^{(2)}(v) = 1$
- expected revenue of RSPE = ½(½+0) = 1/4

The RSPE Auction

Claim #1: RSPE is a truthful auction.

Proof sketch: Say bidder i part of x. i can't change R_2. For i, RSPE is same as ProfitExtract(x, R_2), where truthful bidding is optimal.

Claim #2: RSPE's revenue is at least $\min\{R_1, R_2\}$.

Proof : E.g., if $R_1 \le R_2$, ProfitExtract(y, R_1) will successfully extract revenue R_1.

- recall key property of ProfitExtract subroutine

The Final Lemma

Claim #3: For all inputs v, $E[\min\{R_1,R_2\}] \geq F^{(2)}/4$.

- expectation is over the random split v=(x,y)
- recall $F^{(2)}(v) := \max_{2 \leq i \leq k} iv_i$ (assume sorted v_i's)

Proof : Fix v. Let $i \geq 2$ satisfy $F^{(2)}(v) = iv_i$.

Let a,b = number of top i bidders in x,y.

Note: $R_1 \geq av_i$ and $R_2 \geq bv_i$

- So: just need $E[\min\{a,b\}] \geq i/4$.

The Final Lemma (con'd)

Need to show : For every $i \geq 2$, a random split into a,b (with a+b=i) satisfies $E[\min\{a,b\}] \geq i/4$.

Case i = 2: min{a,b} is either 0 (50% probability) or 1 (50% probability) => $E[\min\{a,b\}] = 1/2 = i/4$.

Case i = 3: min{a,b} is either 0 (25% probability) or 1 (75% probability) => $E[\min\{a,b\}] = 3/4 = i/4$.

For general case: imagine throwing balls into two urns. Each new pair of balls increases the smaller population by 1 with 50% probability.

Key Points of the Day

From Bayesian to Worst-Case Optimal Auctions

Tim Roughgarden (Stanford)
[mostly joint work with
Jason Hartline (Northwestern)]

Example: Multi-Item Auctions

Setup: n bidders, k identical goods.

- bidder i has private "valuation" v_i for a good
- v_i = maximum willingness to pay

Design space: decide on:

- (1) at most k winners; and (2) selling prices.

Example: Vickrey auction.

- top k bidders win; all pay (k+1)th highest bid

Variant: Vickrey with a *reserve*. [≈extra bid by seller]

Auction Benchmarks

Goal: prove results of the form (e.g., for revenue):

"Theorem: auction A is (approximately) optimal."

Auction Benchmarks

Goal: prove results of the form (e.g., for revenue):

"Theorem: auction A is (approximately) optimal."

Auction model: single-round, sealed bid auctions.

- truthful (i.e., truthful bidding [$b_i=v_i$] is foolproof)
  - equivalent: selling price to i independent of b_i

Fact: [Myerson 81] these assumptions are "WLOG".

- "Revelation Principle" + "Revenue Equivalence Thms"

The Fixed Price Benchmark

Solution: [Goldberg/Hartline//Wright SODA 01]

- define OPT(v) := best *fixed-price* revenue:

$$RB(v) := \max_{i \leq k} i v_i \qquad \text{(assume sorted } v_i\text{'s)}$$

Usual justification: "seems to work".

- α-competitive auctions exist for small α
  - assuming no "dominant bidder"
- no auction has α smaller than 1 (or even 2.42)

Question: is there a fundamental explanation?

Bayesian Profit Maximization

Example: 1 bidder, 1 item, v ~ known distribution F

- truthful auctions = posted prices p
- expected revenue of p: $p(1-F(p))$
  - given F, can solve for optimal p^*
  - e.g., $p^* = ½$ for v ~ uniform[0,1]
- but: what about k,n >1 (with i.i.d. v_i's)?

Bayesian Profit Maximization

Example: 1 bidder, 1 item, v ~ known distribution F

- truthful auctions = posted prices p
- expected revenue of p: p(1-F(p))
  - given F, can solve for optimal p^*
  - e.g., $p^* = ½$ for v ~ uniform[0,1]
- but: what about k,n >1 (with i.i.d. v_i's)?

Theorem: [Myerson 81] auction with max expected revenue is Vickrey with above reserve p^*.

- note p^* is *independent of k and n*

Proof of Myerson's Theorem

Given: a truthful auction (x,p), denoting the:

- *allocation rule*: v_i's $\longrightarrow$ x_i's [who wins]
- *payment rule*: v_i's $\longrightarrow$ p_i's [who pays what]

Key Lemma: characterize expected revenue in terms of "virtual valuations“ of the winner:

$$= E_{\mathbf{v}}[\Sigma_i \, \varphi(v_i) \cdot x_i(\mathbf{v})]$$

where

"virtual valuation" $\longrightarrow$ $\left\{ \varphi(v_i) = v_i - \dfrac{1-F(v_i)}{f(v_i)} \right.$

E.g., $\varphi(v_i) = 2v_i-1$ when F = Unif[0,1]

Proof of Myerson's Theorem

So far: expected revenue of any (x,p):

$$= E_{\mathbf{v}}[\Sigma_i \varphi(v_i) \cdot x_i(\mathbf{v})]$$

E.g., $\varphi(v_i) = 2v_i - 1$ when $F = \mathrm{Unif}[0,1]$

- to maximize: *for each vector* $\mathbf{v}$, *set* x_i*'s to maximize sum of virtual valuations.*
  - fine print: need F to be "regular" to be truthful
- multi-item auctions: award k items to the top k $\varphi(v_i)$'s that are also positive
  - i.e., Vickrey with reserve price $\varphi^{-1}(0)$ [for all k,n]

Opt Fixed-Price via Myerson

Recall question: meaning of the optimal fixed-price revenue for (non-Bayesian) auctions?

$$RB(v) := \max_{i \le k} i v_i \qquad \text{(assume sorted } v_i\text{'s)}$$

Recall: "seems to work" (even with apples vs. oranges).

Myerson: *for all F*, Vickrey + a reserve is optimal.

Corollary 1: *for all F and all* $\mathbf{v}$, ex post behavior of optimal auction for F is to charge a fixed price.

- namely: max{reserve price, (k+1)th highest bid of **v**}

Opt Fixed-Price via Myerson

Corollary 2: If auction A is α-competitive w.r.t benchmark RB, then it is *simultaneously competitive with all Bayesian optimal auctions!*

I.e.: For every F, corresponding opt auction A_F:

A's expected revenue $\geq$ (A_F's expected revenue)/α

Proof: inequality holds for every v:

A's revenue on v $\geq$ RB(v)/α $\geq$ (A_F's revenue on v)/α

Interpretation: ignorance of F costs only α factor.

From Old to New Results

So far: re-interpretation of old results for worst-case profit-maximization in multi-item auctions.

Next: new applications [Hartline/Roughgarden 08, 09]

- beyond multi-item auctions
- beyond identical bidders
- novel objectives (will skip, see [HR STOC 08])

Analysis Template

Moral: Bayesian auction design yields strong worst-case performance benchmarks.

- characterize ex post Bayesian optimal behavior
  - for all distributions of interest, all valuation profiles
- want to simultaneously compete with all such behaviors (on each valuation profile)
- automatic corollary: competitive with expected performance of every Bayesian-optimal auction

Beyond Multi-Item Auctions

Example: n bidders (valuations v_i), feasible subsets of winners = independent sets of some matroid.

- e.g., spanning trees; multi-item = uniform matroid

Beyond Multi-Item Auctions

Example: n bidders (valuations v_i), feasible subsets of winners = independent sets of some matroid.

- e.g., spanning trees; multi-item = uniform matroid

Myerson's Revenue Formula: given a prior distribution F (with virtual value φ), expected revenue of an auction: $E_{\mathbf{v}}[\Sigma_i \varphi(v_i) \cdot x_i(\mathbf{v})]$

To maximize: given v, choose independent set maximizing $\Sigma_i \varphi(v_i)$ [e.g., via greedy algorithm]

A Prior-Free Benchmark

So: ex post behavior of a Bayesian optimal auction:

- choose independent set with max $\Sigma_i \varphi(v_i)$

Equivalent: the "VCG mechanism" with a common reserve price.

- VCG allocation rule: pick feasible set with max $\Sigma_i v_i$

Upshot: prior-free benchmark RB(v) := max revenue achievable via VCG with a common reserve.

- [HR]: can be 8-competitive with this benchmark
- randomize between VCG and ProfitExtract

More Prior-Free Benchmarks

Beyond matroids?: e.g., each bidder wants a bundle of goods, can only allocate each good once.

- the optimal mechanism is complicated

But: [Hartline/Roughgarden EC 09] VCG mechanism with a common reserve is a 2-approximation.

  - needs somewhat stronger distributional assumption
  - offers simple and provably good alternative to the (complex) optimal auction
  - justifies VCG + optimal reserve benchmark in general

Beyond Symmetric Bidders

Asymmetric bidders (Bayesian): different prior F_i (and corresponding φ_i) for each bidder i.

- Myerson's formula: $E_{\mathbf{v}}[\Sigma_i \varphi_i(v_i) \cdot x_i(\mathbf{v})]$
- arbitrary F_i's => all prices can arise ex post
- *ordered* F_i's => optimal prices *monotone*
  - "ordered" = φ_i^{-1}'s can be consistently ordered

Prior-free version: only know bidder ordering .

- RB(v) := max revenue via monotone prices
- Open: can you O(1)-compete with this?

Conclusions

Take-home point: new template for generating meaningful worst-case auction benchmarks.

- automatic: simultaneous competitive guarantee with all Bayesian-optimal auctions
- enables new theorems for money-burning mechanisms, asymmetric allocations and/or bidders

Open Questions:

- thoroughly understand "single-parameter" problems
- multi-parameter? (e.g., combinatorial auctions)
- general definitions of "more informed opponent"

Bulow-Klemperer ('96)

Observation: for every F, $E[\varphi(v_i)] = 0$.

- proof #1: consider Vickrey with k = n = 1
- proof #2: integrate $\varphi(v_i) = v_i - (1-F(v_i)/f(v_i))$

Bulow-Klemperer ('96)

Observation: for every F, $E[\varphi(v_i)] = 0$.

- proof #1: consider Vickrey with $k = n = 1$
- proof #2: integrate $\varphi(v_i) = v_i - (1-F(v_i)/f(v_i))$

Corollary [BK96]: for $k = 1$, every $n \geq 1$, every F:

Vickrey's revenue OPT's revenue

Bulow-Klemperer ('96)

Observation: for every F, $E[\varphi(v_i)] = 0$.

- proof #1: consider Vickrey with $k = n = 1$
- proof #2: integrate $\varphi(v_i) = v_i - (1-F(v_i)/f(v_i))$

Corollary [BK96]: for $k = 1$, every $n \geq 1$, every F:

Vickrey's revenue $\geq$ OPT's revenue

Bulow-Klemperer ('96)

Observation: for every F, $E[\varphi(v_i)] = 0$.

- proof #1: consider Vickrey with k = n = 1
- proof #2: integrate $\varphi(v_i) = v_i - (1-F(v_i)/f(v_i))$

Corollary [BK96]: for k = 1, every n ≥ 1, every F:

Vickrey's revenue ≥ OPT's revenue

[with (n+1) i.i.d. bidders] [with n i.i.d. bidders]

Interpretation: small increase in market size more important than running optimal auction.

Bulow-Klemperer (Proof)

Proof idea:

- OPT's expected revenue [n bidders]:

  $$E_{\mathbf{v}}[\max \{\max_{i \le n} \varphi(v_i), 0\}]$$

- Vickrey's expected revenue [(n+1) bidders]:

  $$E_{\mathbf{v}}[\max \{\max_{i \le n} \varphi(v_i), \varphi(v_{n+1})\}]$$

- condition on $\varphi(v_1),\ldots,\varphi(v_n)$, use observation that $E[\varphi(v_i)] = 0$

Intrinsic Robustness of the Price of Anarchy

Tim Roughgarden
Stanford University

The Price of Anarchy

Recall: price of anarchy of a game = $\dfrac{\text{obj fn value of worst equilibrium}}{\text{optimal obj fn value}}$

Example: POA = 4/3 in Pigou's example

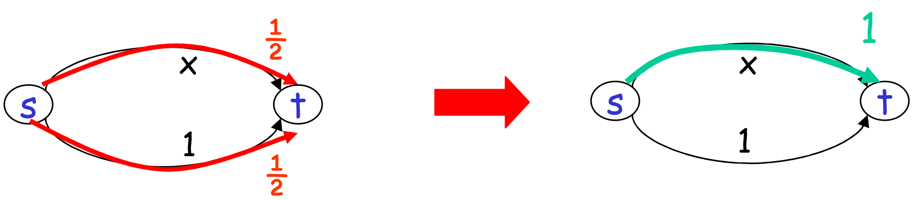

Cost = 3/4   Cost = 1

Key Points for Lecture

- 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

Variational Inequalities

Recall nonatomic variation inequality:

- if f is a Nash flow and f^* is feasible, then

$$\Sigma_e f_e \cdot c_e(f_e) \leq \Sigma_e f_e^* \cdot c_e(f_e)$$

Atomic variational inequality:

$$\Sigma_e f_e \cdot c_e(f_e) \leq \Sigma_e f_e^* \cdot c_e(f_e+1)$$

Proof: for each player i, current cost $c_P(f)$ only increases from switching to path in f^*

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 **s**,**s*** outcomes (λ > 0; μ < 1):

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

Smooth => POA Bound

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

- notation: **s** = a Nash eq; **s*** = optimal

Assuming (λ,μ)-smooth:

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

Then: POA (of pure Nash eq) ≤ λ/(1-μ).

Why Is Smoothness Stronger?

Key point: to derive POA bound, only needed

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

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 "price of total anarchy" implies bound of inefficiency of mixed + correlated equilibria

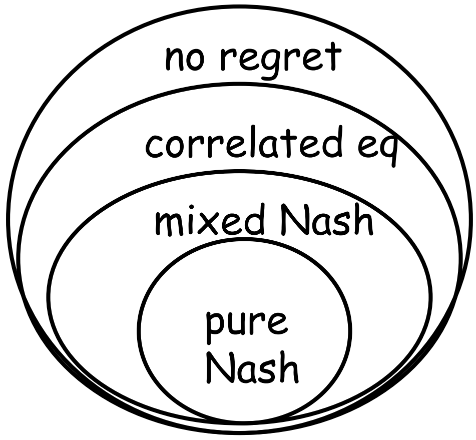

- bound applies even to sequences that don't converge in any sense
  - no regret much weaker than reaching equilibrium
  - [Blum/Even-Dar/Ligett PODC 06], [Blum/Hajiaghayi/Ligett/Roth STOC 08]

Smooth => POTA Bound

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

Assuming (λ,μ)-smooth:

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

Smooth => POTA Bound

- notation: $s^1, s^2, \dots, s^T$ = 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]

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

Smooth => POTA Bound

- notation: $s^1, s^2, \dots, s^T$ = 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

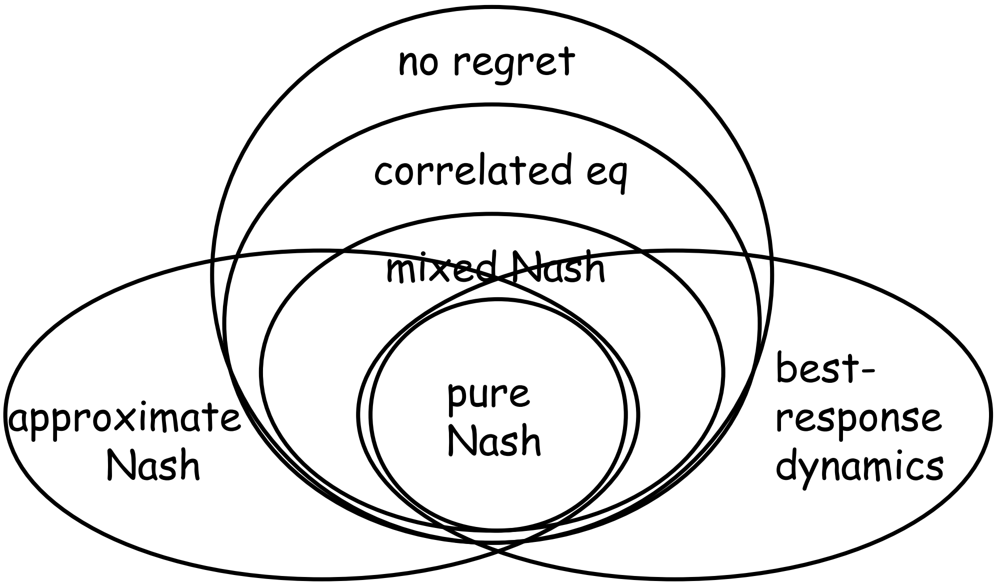

Theorem: in a (λ,μ)-smooth game, everything in these sets costs (essentially) $\lambda/(1-\mu) \times$ OPT.

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)

Geometry of Affine Case

Assume: $c_e(x) = a_e x + b_e$

Goal: compare

$$f_e^* \cdot [c_e(f_e) - c_e(f_e^*)] \text{ vs. } f_e \cdot c_e(f_e)$$

Interesting case: when $c_e(f_e) > c_e(f_e^*)$:

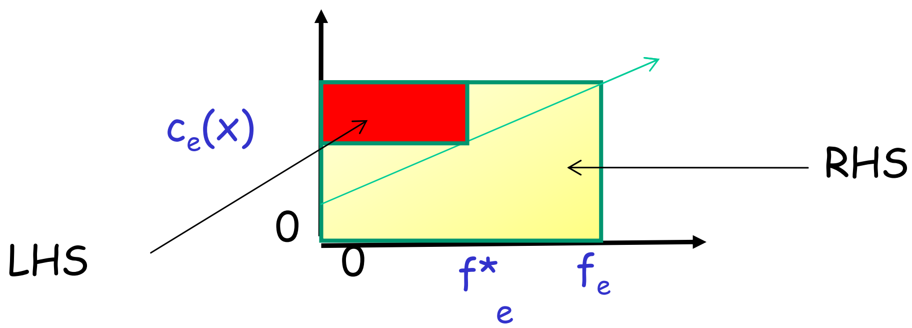

A Technical Lemma

Recall Claim for Atomic Case:

- [Christodoulou/Koutsoupias 05]: for all integers y,z:

$$y(z+1) \le (5/3)y^2 + (1/3)z^2$$

- so: $ay(z+1) + by \le (5/3)[ay^2 + by] + (1/3)[az^2 + bz]$
  - for all integers y,z and $a,b \ge 0$
- so: $\Sigma_e [a_e(f_e+1) + b_e)f_e^*] \le (5/3)\, \Sigma_e [(a_e f_e^* + b_e)f_e^*] + (1/3)\, \Sigma_e [(a_e f_e + b_e)f_e]$
- so: $C(f) \le \Sigma_e f^* \cdot c_e(f_e+1) \le (5/3)C(f^*) + (1/3)C(f)$

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

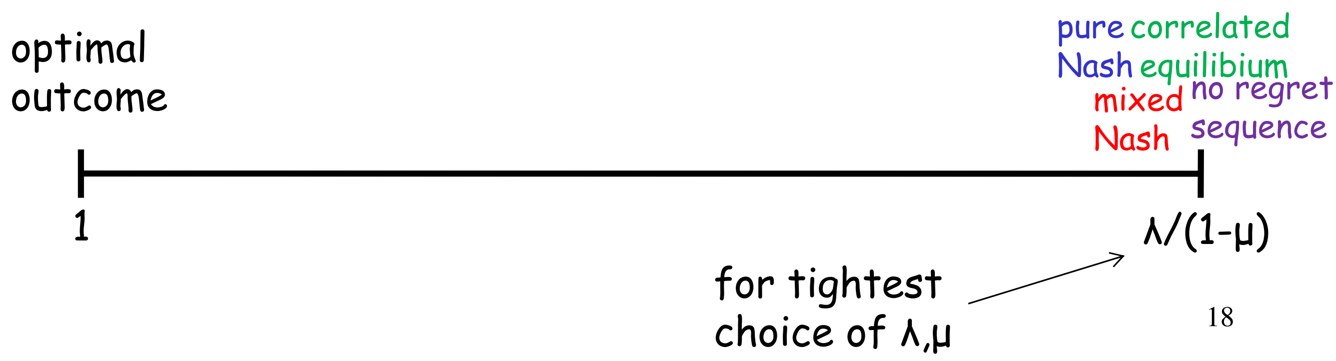

First Main Proof Step

Step 1: characterize optimal smoothness parameters λ,μ as vertex of a 2-D polyhedron

- intersection of halfplanes of the form

$$yc(z+1) \le \lambda \cdot c(y)y + \mu \cdot c(z)z$$

for all integers y,z and cost fns $c \in C$

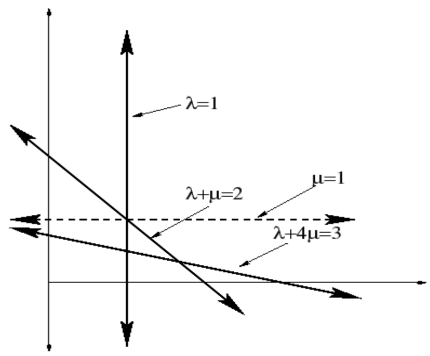

Second Main Proof Step

Step 2: exhibit example with POA = $\lambda/(1-\mu)$

- use two "parallel cycles" (one per tight halfplane)
- each player has "short" and "long" strategy
  - each strategy uses resources of *both* cycles

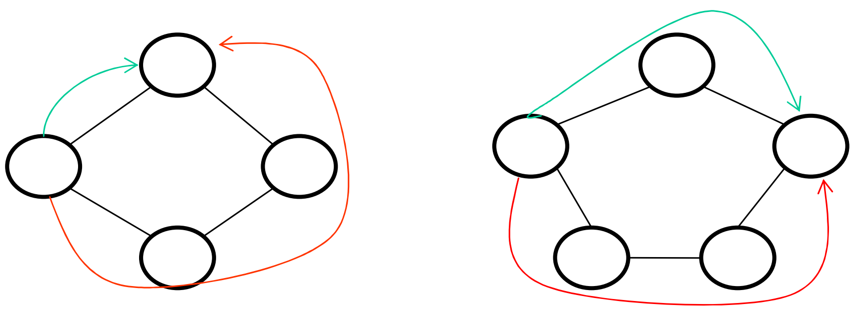

- OPT = all use short strategies;
- worst Nash = all use long strategies

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]

Take-Home Points

- the most common way of proving POA bounds automatically yields a much more robust guarantee
- and this technique often gives tight bounds
- future work: characterize tight game classes (where smoothness gives optimal POA bounds, even for pure NE)