TTIC 31260 Algorithmic Game Theory

Price of Anarchy II: Guiding Dynamics to Higher Quality Equilibria

Your guide:

Avrim Blum

Guiding Dynamics

- So far, we have discussed convergence to some equilibrium.

  - This is usually the best you can hope for with natural dynamics, esp in time poly in the size of the game.

- What about games with a big gap between their best and worst equilibria? (E.g., large PoA, small PoS)

If players are currently in a bad equilibrium, can we "nudge" behavior towards the good ones?

Good equilibria, Bad equilibria

Many games have both good and bad equilibria.

In some places, everyone throws their trash on the street. In some, everyone puts their trash in the trash can.

In some places, everyone drives their own car. In some, everybody uses and pays for good public transit.

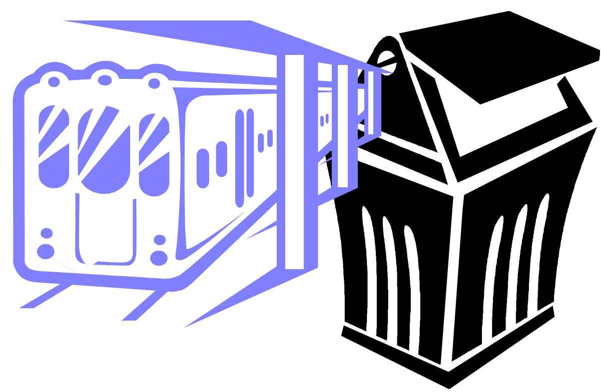

Good equilibria, Bad equilibria

In fair cost-sharing, there exist networks where the worst equilibrium is a factor n more costly than the best equilibrium.

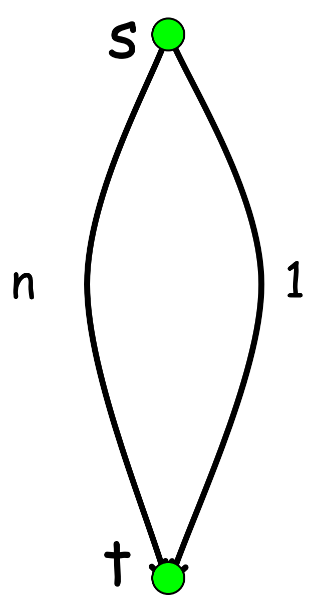

Suppose players enter in one at a time and take the best path given the costs so far. Are we guaranteed a good equilibrium then?

Good equilibria, Bad equilibria

No. Think of the "cars and shared transit" game, where the shared transit has cost 2:

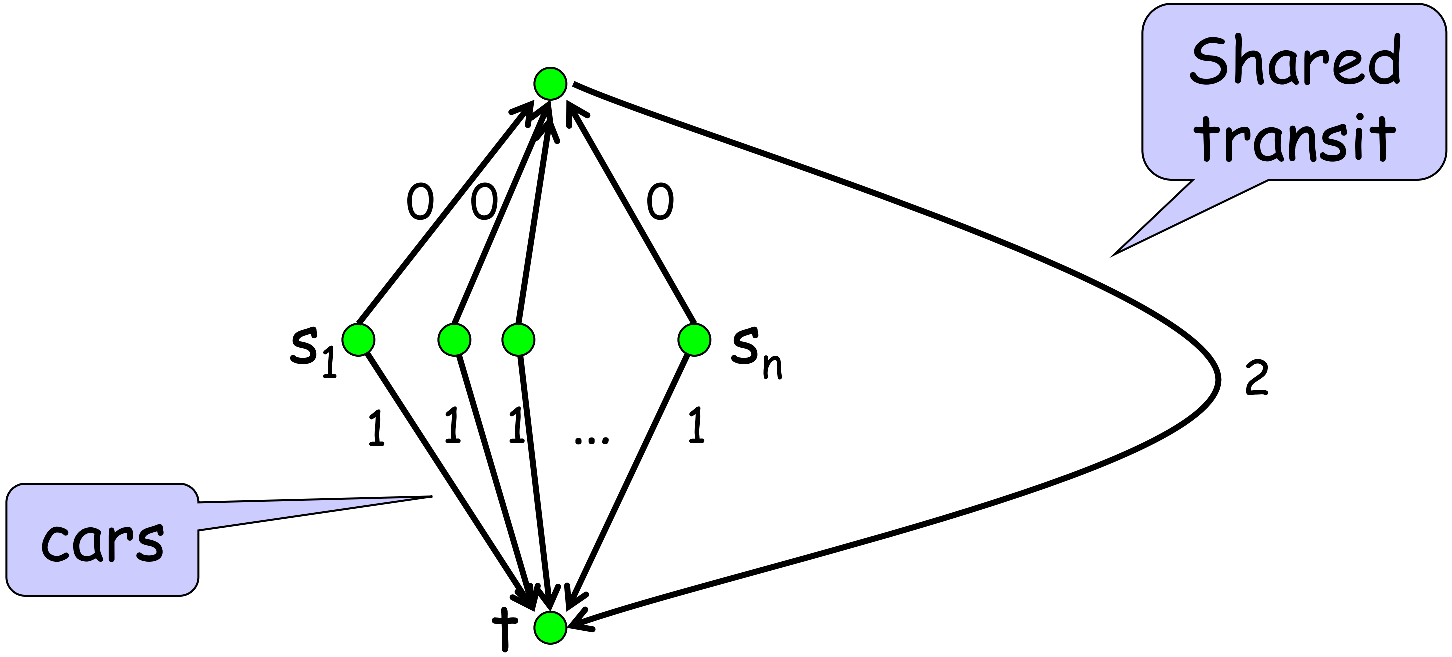

Nudging from bad to good

"Public service advertising model":

0. n players begin in some arbitrary configuration.
1. Authority launches public-service advertising campaign, proposing joint action s^*.

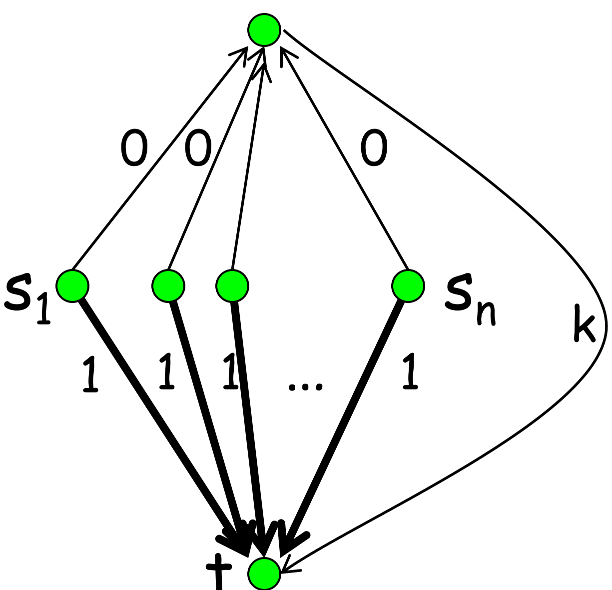

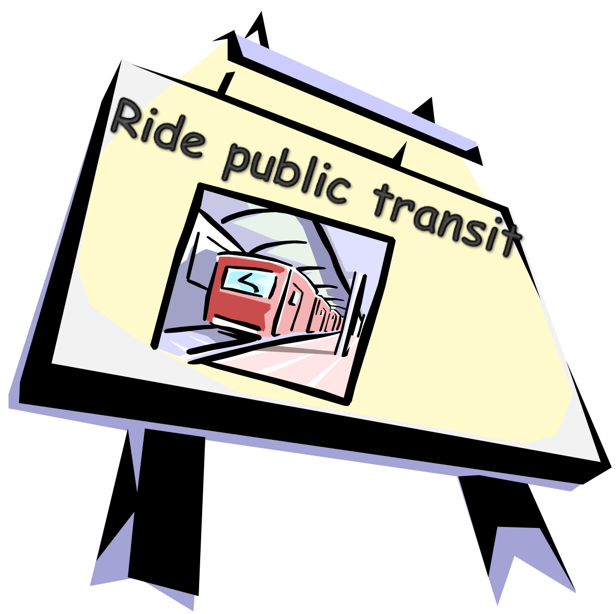

Nudging from bad to good

"Public service advertising model":

0. n players begin in some arbitrary configuration.
1. Authority launches public-service advertising campaign, proposing joint action s^*. Each player i pays attention and follows with probability α. Call these the receptive players
2. Remaining (non-receptive) players fall to some arbitrary equilibrium for themselves, given play of receptive players.

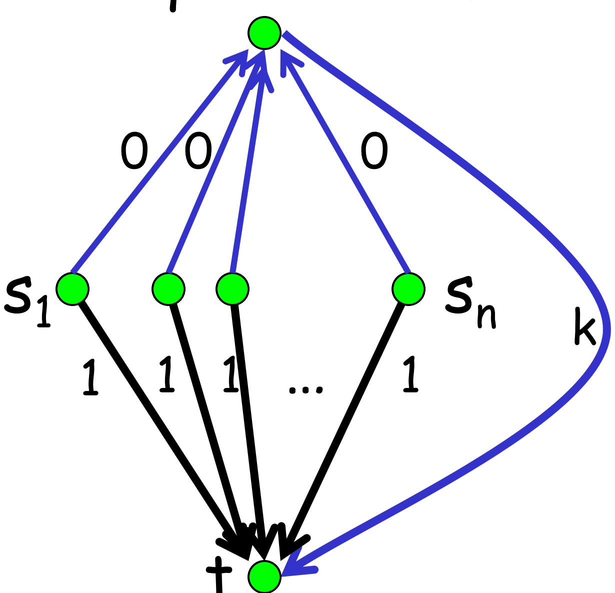

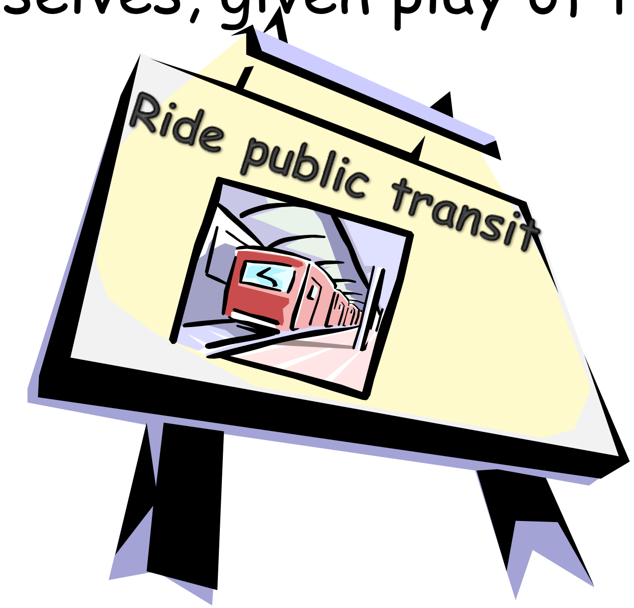

Nudging from bad to good

"Public service advertising model":

0. n players begin in some arbitrary configuration.
1. Authority launches public-service advertising campaign, proposing joint action s^*. Each player i pays attention and follows with probability α. Call these the receptive players
2. Remaining (non-receptive) players fall to some arbitrary equilibrium for themselves, given play of receptive players.
3. Campaign wears off. Entire set of players follows best-response dynamics from then on.

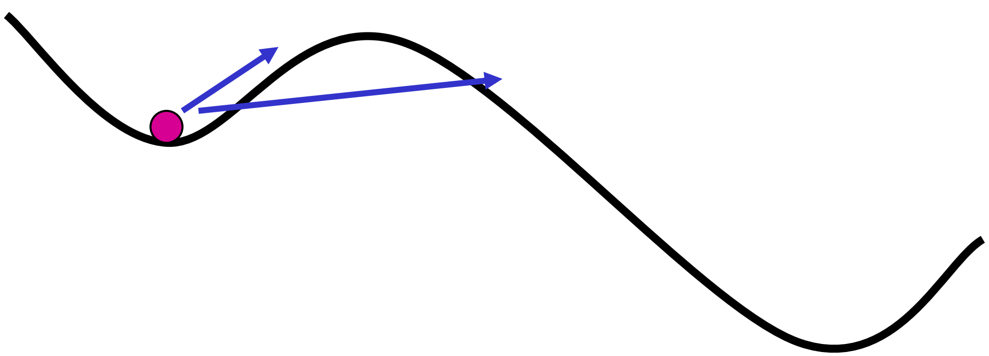

Nudging from bad to good

"Public service advertising model":

0. n players begin in some arbitrary configuration.
1. Authority launches public-service advertising campaign, proposing joint action s^*. Each player i pays attention and follows with probability α. Call these the receptive players
2. Remaining (non-receptive) players fall to some arbitrary equilibrium for themselves, given play of receptive players.
3. Campaign wears off. Entire set of players follows best-response dynamics from then on.

Note: if $\alpha=1$, then it's obvious. Key issue: what if $\alpha < 1$?

Fair Cost Sharing

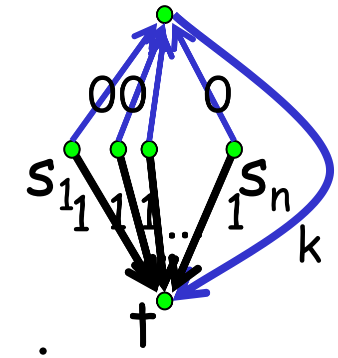

Cost sharing: (PoS = log(n), PoA = n)

If only an α probability of players following the advice, then we get expected cost within $O(\log(n)/\alpha)$ of OPT.

Proof:

- Advertiser proposes OPT (any apx also works)
- In any NE for non-receptive players, any such player i can't improve by switching to his path P_i^{OPT} in OPT.

$$cost_i(s) \leq \sum_{e \in P_i^{OPT}} c_e/(1 + n_{e,R})$$

#receptives on edge e (any extras only helps)

Fair Cost Sharing

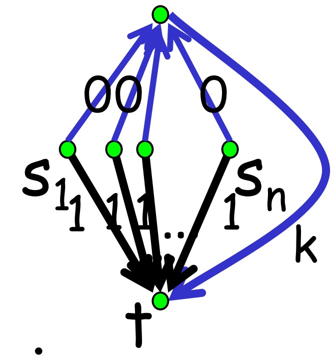

Cost sharing: (PoS = log(n), PoA = n)

If only an α probability of players following the advice, then we get expected cost within $O(\log(n)/\alpha)$ of OPT.

Proof:

- Advertiser proposes OPT (any apx also works)
- In any NE for non-receptive players, any such player i can't improve by switching to his path P_i^{OPT} in OPT.

$$cost_i(s) \leq \sum_{e \in P_i^{OPT}} c_e/(1+n_{e,R})$$

- Calculate total cost of these guaranteed options. $cost_{\bar{R}}(s)$

$$\leq \sum_{i \notin R} \sum_{e \in P_i^{OPT}} c_e/(1+n_{e,R}) \leq \sum_e n_{e,\bar{R}} \cdot c_e/(1+n_{e,R})$$

Rearrange sum...

#non-receptives on edge e in OPT

Fair Cost Sharing

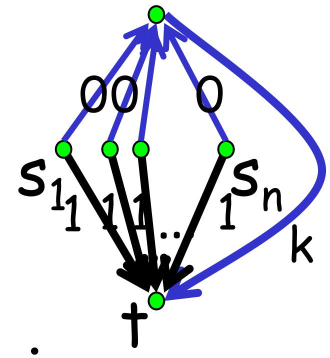

Cost sharing: (PoS = log(n), PoA = n)

If only an α probability of players following the advice, then we get expected cost within $O(\log(n)/\alpha)$ of OPT.

Proof:

- Add in cost of receptives: $cost(s) \leq \sum_e 2n_{e,OPT} \cdot c_e/(1+n_{e,R})$

For receptives, denominator only $n_{e,R}$, giving factor of 2.

- Finally, use: if $X \sim Bin(m,p)$ then $E\left[\frac{1}{X+1}\right] = O\left(\frac{1}{p \cdot m}\right)$.
- Take expectation: get $O(OPT/\alpha)$. (End of phase 2)
- Calculate total cost of these guaranteed options. $cost_{\bar{R}}(s)$

$$\leq \sum_{i \notin R} \sum_{e \in P_i^{OPT}} c_e/(1+n_{e,R}) \leq \sum_e n_{e,\bar{R}} \cdot c_e/(1+n_{e,R})$$

Rearrange sum...

#non-receptives on edge e in OPT

Fair Cost Sharing

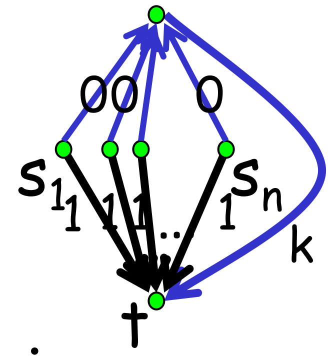

Cost sharing: (PoS = log(n), PoA = n)

If only an α probability of players following the advice, then we get expected cost within $O(\log(n)/\alpha)$ of OPT.

Proof:

For receptives, denominator only $n_{e,R}$, giving factor of 2.

- Add in cost of receptives: $cost(s) \leq \sum_e 2n_{e,OPT} \cdot c_e/(1+n_{e,R})$
- Finally, use: if $X \sim Bin(m,p)$ then $E\left[\frac{c}{X+1}\right] = O\left(\frac{c}{p \cdot m}\right)$.
- Take expectation: get $O(OPT/\alpha)$. (End of phase 2)
- Finally, in last phase, potential argument shows behavior cannot get worse by more than an additional $\log(n)$ factor.

(End of phase 3)

Cost Sharing, Extension

Cost sharing: + linear delays: $f_e(n_e) = c_e/n_e + \ell_e \cdot n_e$

- Problem: can't argue as if remaining NR players didn't exist since they add to delays

Proof Idea:

- Define shadow game: pure affine latency functions (linear plus constant). Offset defined by equilibrium at end of phase 2.

$$\hat{f}_e(n_e) = c_e/(1 + \hat{n}_e) + \ell_e \cdot n_e$$

users on e at end of phase 2

- This has good PoA.

Theorem 18.23 (The price of anarchy in affine atomic instances) *If* (G, r, c) *is an atomic instance with affine cost functions, then the price of anarchy of* (G, r, c) *is at most* $(3 + \sqrt{5})/2 \approx 2.618$.

Cost Sharing, Extension

Cost sharing: + linear delays: $f_e(n_e) = c_e/n_e + \ell_e \cdot n_e$

- Problem: can't argue as if remaining NR players didn't exist since they add to delays

Proof Idea:

- Define shadow game: pure affine latency functions (linear plus constant). Offset defined by equilibrium at end of phase 2.

$$\hat{f}_e(n_e) = c_e/(1 + \hat{n}_e) + \ell_e \cdot n_e$$

users on e at end of phase 2

- This has good PoA.

- State at end of phase 2 is equilibrium for $\bar{R}$ for this game too.

- Why: costs in state are lower, and costs for deviations same.

Cost Sharing, Extension

Cost sharing: + linear delays: $f_e(n_e) = c_e/n_e + \ell_e \cdot n_e$

- Problem: can't argue as if remaining NR players didn't exist since they add to delays

Proof Idea:

- Define shadow game: pure affine latency functions (linear plus constant). Offset defined by equilibrium at end of phase 2.

$$\hat{f}_e(n_e) = c_e/(1 + \hat{n}_e) + \ell_e \cdot n_e$$

users on e at end of phase 2

- This has good PoA.
- State at end of phase 2 is equilibrium for $\bar{R}$ for this game too.
- Also $cost(\hat{f}) \geq cost(f)/2$. So, $cost(f) \leq 2cost(\hat{f}) = O(\widehat{OPT})$.

Cost Sharing, Extension

Cost sharing: + linear delays: $f_e(n_e) = c_e/n_e + \ell_e \cdot n_e$

Finally: $E\left[\widehat{OPT}\right] = O(OPT/\alpha)$. [Because one option for $\widehat{OPT}$ is to use the same paths as OPT, and for each edge e, its expected cost in shadow game is $O(1/\alpha)$ times its cost in real game under OPT]

- Define shadow game: pure affine latency functions (linear plus constant). Offset defined by equilibrium at end of phase 2.

$$\hat{f}_e(n_e) = c_e/(1+\hat{n}_e) + \ell_e \cdot n_e$$

users on e at end of phase 2

- This has good PoA.

- State at end of phase 2 is equilibrium for $\bar{R}$ for this game too.

- Also $cost(\hat{f}) \geq cost(f)/2$. So, $cost(f) \leq 2cost(\hat{f}) = O(\widehat{OPT})$.

Cost Sharing, Extension

Cost sharing: + linear delays: $f_e(n_e) = c_e/n_e + \ell_e \cdot n_e$

Finally: $E[\widehat{OPT}] = O(OPT/\alpha)$. [Because one option for $\widehat{OPT}$ is to use the same paths as OPT, and for each edge e, its expected cost in shadow game is $O(1/\alpha)$ times its cost in real game under OPT]

And then lose $O(\log n)$ in Step 3 as before, since the potential function for cost sharing with linear delays satisfies

$$\frac{1}{2} cost(s) \leq \Phi(s) \leq cost(s) \cdot \log(n).$$

Party affiliation games

- Given graph G, each edge labeled + or -.
- Vertices have two actions: RED or BLUE.

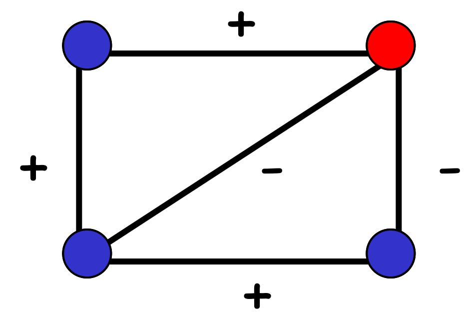

$$\text{cost}_i(s) = \sum_{(i,j)\in +} I_{(s_i \neq s_j)} + \sum_{(i,j)\in -} I_{(s_i = s_j)}$$

Pay 1 for each + edge with endpoint of different color, and each - edge with endpoint of same color.

$$\text{cost}(s) = \sum_i \text{cost}_i(s) + 1$$

+1 to keep ratios finite

- Special cases:
  - All + edges is consensus game. [OPT is all red or all blue]
  - All - edges is cut-game. [OPT is a max cut]

Party affiliation games

Party Affiliation: (PoS = 1, PoA = $\Omega(n^2)$)

- Threshold behavior: for $\alpha > \frac{1}{2}$, can get ratio O(1), but for $\alpha < \frac{1}{2}$, ratio stays $\Omega(n^2)$. (assume degrees $\omega(\log n)$).

Lower bound:

- Consensus game, two cliques, with relatively sparse between them. Players "locked" into place.

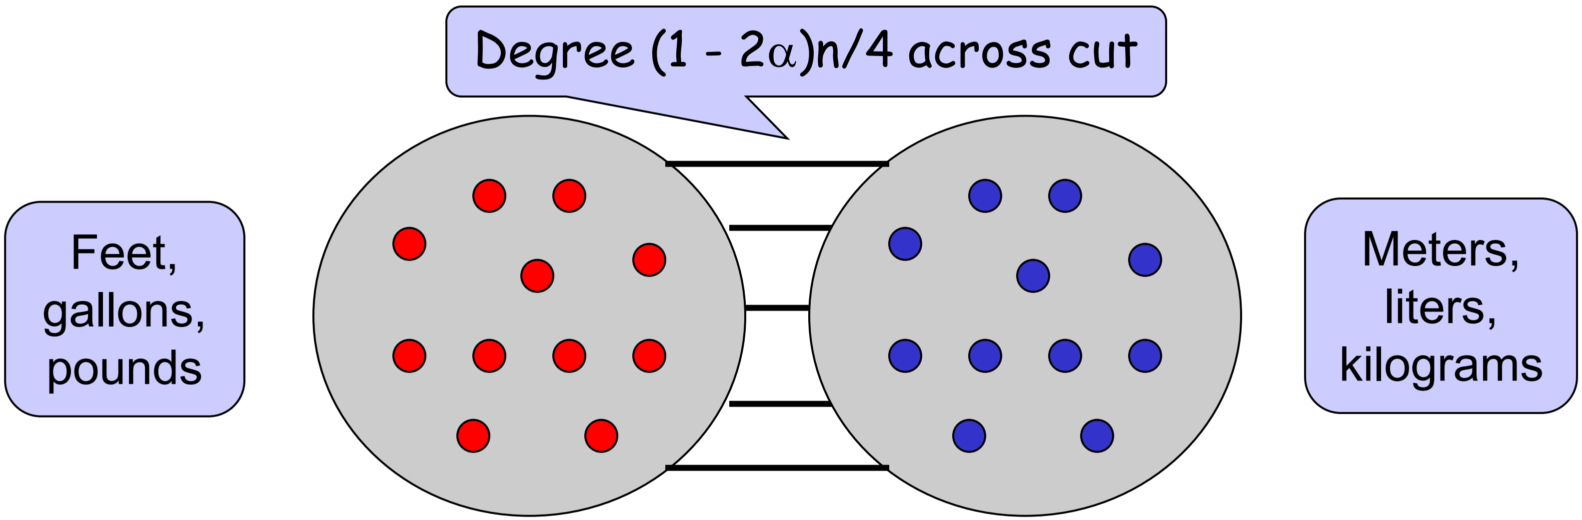

If advertise "blue" then each red still has $\approx (1-\alpha)n/2$ red nbrs compared to $\approx \alpha n/2 + deg$ blue neighbors. So stays red so long as $deg \ll (1-2\alpha)n/2$.

Party affiliation games

Party Affiliation: (PoS = 1, PoA = $\Omega(n^2)$)

- Threshold behavior: for $\alpha > \frac{1}{2}$, can get ratio $O(1)$, but for $\alpha < \frac{1}{2}$, ratio stays $\Omega(n^2)$. (assume degrees $\omega(\log n)$).

Upper bound:

- Split nodes into those incurring low-cost vs those incurring high-cost under OPT.
- Show that low-cost-in-OPT will switch to behavior in OPT. For high-cost-in-OPT, don't care.
- Cost only improves in final best-response process.

References

- M-F Balcan, A Blum, Y Mansour, "Circumventing the Price of Anarchy: Leading Dynamics to Good Behavior", SIAM J Computing 2013.
- Balcan, M. F., Krehbiel, S., Piliouras, G., & Shin, J. (2014). Near-optimality in covering games by exposing global information. *ACM Transactions on Economics and Computation (TEAC)*, *2*(4), 1-22.
- Chapter 18.4.2