CSE 6319 Notes 2: Price of Anarchy

(Last updated 2/11/24 12:47 PM)

Karlin & Peres 8

Nisan 17/18/19/20

Roughgarden 1/11/12/13/14/15

2.A. RELATING EQUILIBRIA AND OPTIMALITY

Selfish Routing (infinitesimal, non-atomic)

Braess's Problem (R 1, 11.1; KP p. 148)

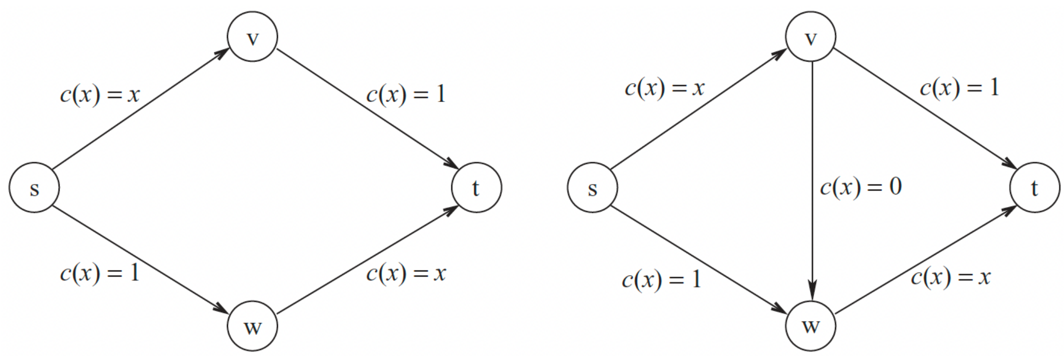

(a) Initial network (b) Augmented network

Initial network: Suppose load is 1.0, so random traffic will split between two routes (1.5 hours each way)

Augmented network: Upon "signal" that v $\rightarrow$ w exists, each driver assumes that s $\rightarrow$ v $\rightarrow$ w $\rightarrow$ t is a 2ϵ hour path. But, s $\rightarrow$ v and w $\rightarrow$ t saturate, so a 2-hour drive results and POA = cost of worst NE/optimal cost = 2/1.5 = 4/3

Pigou's Problem (R 11.2)

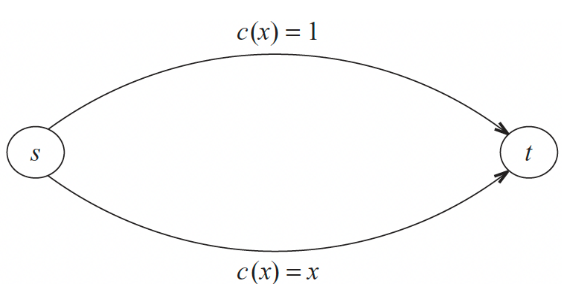

Selfish choice is bottom path with cost of 1.0.

Optimal routing is 1/2 of the traffic on each path, average cost is .75. POA is 4/3

Nonlinear Pigou (R 11.1.3; N p. 464)

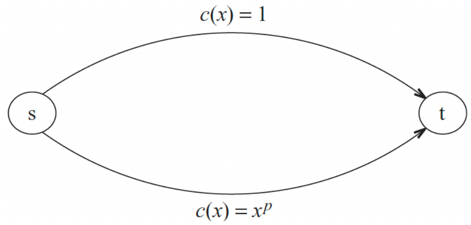

Selfish choice is bottom path with load of 1.0, average cost is 1.

Putting 1/2 the traffic on each path, average cost approaches 0.5.

Optimal: As p increases, the fraction on lower edge should increase with almost no traffic on the upper edge. POA grows unbounded as p increases.

Table 11.1 demonstrates POA for polynomials of various degrees (R p. 149, $p/\ln p$)

Theorem 11.1: Worst-case POA for a set of cost functions is achieved in a Pigou-like network. (R p. 148). *It is unnecessary to construct more complicated networks!*

Definition 18.18 (Pigou bound) Let $\mathcal{C}$ be a nonempty set of cost functions. The *Pigou bound* $\alpha(\mathcal{C})$ for $\mathcal{C}$ is

$$\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \sup_{x,r \geq 0} \frac{r \cdot c(r)}{x \cdot c(x) + (r - x)c(r)}, \tag{18.7}$$

(N p. 473)

where r is the traffic rate and x is the amount of traffic on the lower edge.

Theorem 11.2 generalizes Theorem 11.1 using Pigou bound.

Equilibrium Flows (R p. 153; KP p. 152; N 463)

Usual *conservation of flow* at non-source, non-sink vertices holds

(Shortest path is called "minimum cost path" in some presentations)

In an equilibrium flow, every source-to-sink path (using edges with non-zero flow) is a ***shortest path*** (R Definition 11.3) of the same length as all other SPs.

Travel time along path P: $c_P(f) = \sum_{e \in P} c_e(f_e)$

Total time (cost of a flow) for all paths: $$C(f) = \sum_{P \in \mathcal{P}} c_P(f) f_P$$

Alternative definition in terms of edges: $$C(f) = \sum_{e \in E} c_e(f_e) f_e$$

For a given (non-atomic) network, the equilibrium flow (value) is *unique* (R p. 154, inequality 11.9; N p. 468)

Proof of Theorem 11.2: Especially R, top of p. 155 (skim)

Resource Augmentation Bound (R p. 161; KP 8.1.5)

Traffic-Anarchy Trade-Off (doubling traffic) (KP p. 156; R p. 162; N p. 479)

Theorem 8.1.16. *Let G be a road network with a specified source s and sink t where r units of traffic are routed from s to t, and let $\mathbf{f}$ be a corresponding equilibrium flow with total latency $L(\mathbf{f})$. Let $\mathbf{f}^*$ be an optimal flow when $2r$ units of traffic are routed in the same network, resulting in total latency $L(\mathbf{f}^*)$. Then*

$$L(\mathbf{f}) \leq L(\mathbf{f}^*).$$

Remark 8.1.17. An alternative interpretation of this result is that doubling the capacity of every link can compensate for the lack of central control. See Exercise 8.6.

Atomic Selfish Routing (R p. 163; N p. 465; KP p. 160)

Agents are no longer "infinitesimal". Agents contribute to the count (n) on each edge (or may be generalized to agent i having r_i units of traffic on their path).

Each agent has an origin and destination.

Each edge has a cost function. Often these are affine: $ax + b$

Equilibrium Flow: No path can be replaced to give a decrease in cost (R p. 164)

Unlike non-atomic networks, equilibrium flows ***<u>are not unique</u>*** (R bottom of p. 164)

Theorem 12.3: In every atomic selfish network with affine cost functions, the POA is at most 5/2. (R p. 165; KP p. 161) (skim)

Lemma 12.4 is the critical part of the proof.

Since atomic selfish routing games are potential games, an optimal flow is an equilibrium flow.

AAE Example (N p. 467; R p.166; KP p. 160 Figure 8.11)

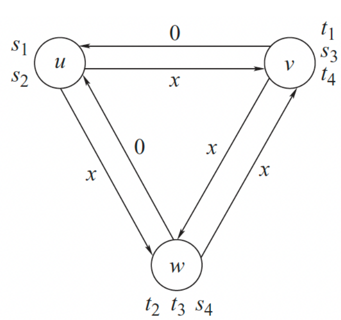

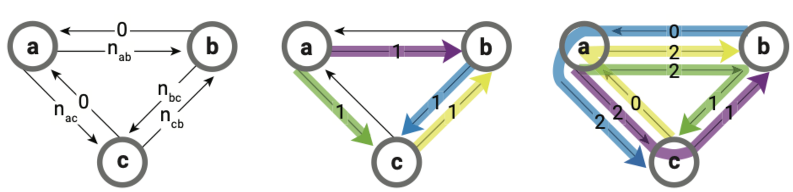

Instance with no equilibrium flow when the players have weights specifying different amounts of traffic (N p. 467; R p. 185)

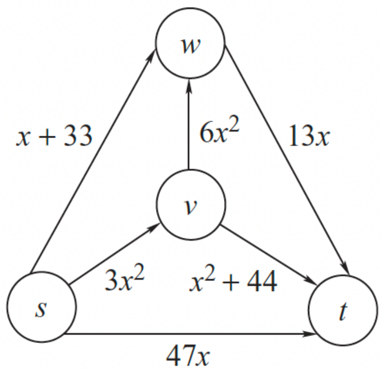

2.B. Hierarchy of Equilibria (R 13.1)

Problem: Various optimization games lack "convenient" equilibria

Pure Nash Equilibria - no unilateral deviation may be an improvement

Mixed Nash Equilibria - no unilateral deviation in mixed strategy gives an expected improvement. (revised POA definition, R p. 175)

Correlated Equilibria - replace product distribution of MNE with distribution over (limited) set of outcomes. Agents know the distribution and (to initiate play) receive an indication of their strategy, but not the strategies of other agents.

Coarse Correlated Equilibria - Agents know only the distribution of outcomes and make the decision to deviate without knowing what (i.e. a strategy) they are deviating from.

Example (R 13.1.6, p. 178) - atomic selfish network example

Single origin and destination for four agents. Six edges labeled 0 ... 5, each with cost function $c(x) = x$.

Agents choose distinct edges ⇒ pure Nash equilibrium (one unit of cost each)

Agents choose edge at random ⇒ mixed Nash equilibrium (3/2 unit of expected cost each, easy simulation, `https://ranger.uta.edu/~weems/NOTES6319/tr13.1.6.c`)

Coordinator informs agents of their edges: two solos and one duet ⇒ correlated equilibrium (still 3/2 as expected cost, avoids worse possibilities)

Still coordinated with two solos and one duet, but edges used are either { 0, 2, 4 } or { 1, 3, 5 } ⇒ not a correlated equilibrium, since a duet agent may move to the next edge. Coarse correlated equilibrium since it is not useful to change edge without knowing the coordinator's decisions.

MAKESPAN SCHEDULING (ON IDENTICAL MACHINES)

`https://en.wikipedia.org/wiki/Balls_into_bins_problem` (Aside: `https://awards.acm.org/kanellakis` 2020)

Selfish Load Balancing (N p.451; N 20; R Problems 12.3 and 13.1)

n tasks (jobs, agents), each with positive weight w_i (distribution?)

m identical machines ("speed" $s_i = 1$ for all machines)

Load on a machine is the sum of the assigned weights

Makespan is the maximum load over all machines

Goal: Minimize the makespan

Pure Nash equilibrium:

Assignment of tasks so no task wishes to change machines to experience a decreased load (N p. 519, Proposition 20.2)

Construction: (N p. 519, Proposition 20.3)

Sorted load vector (decreasing order)   Improvement step . . .

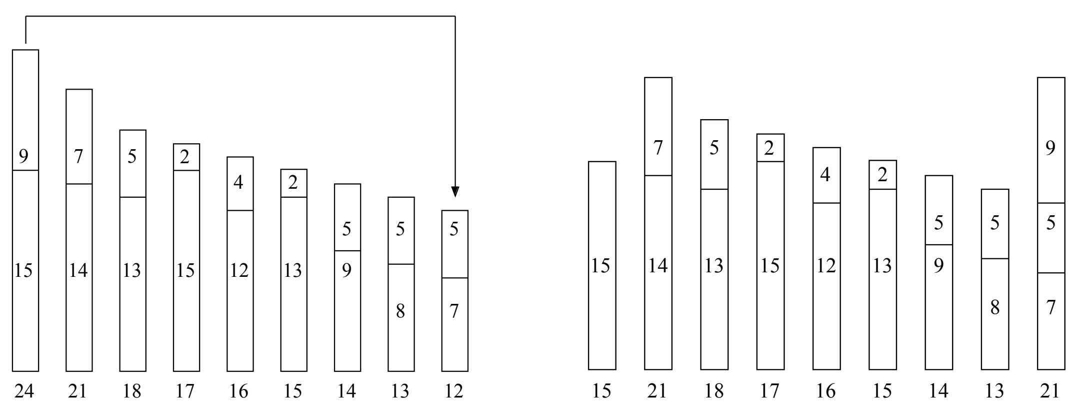

. . . yields a sorted load vector lexicographically smaller than the original

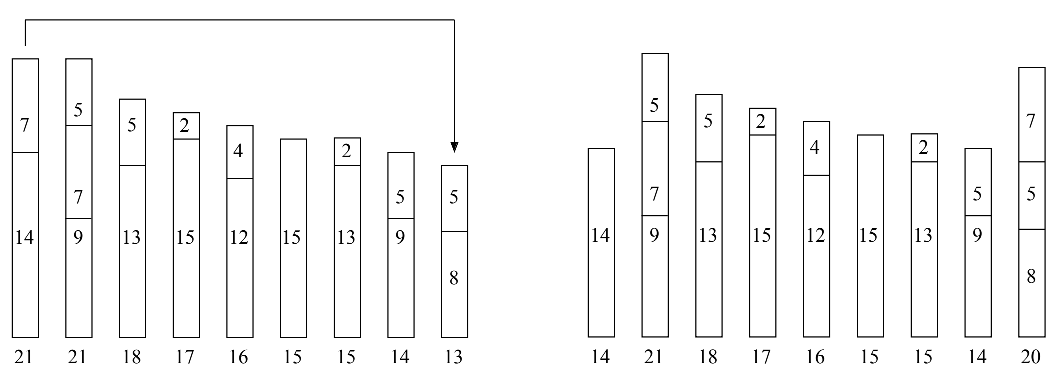

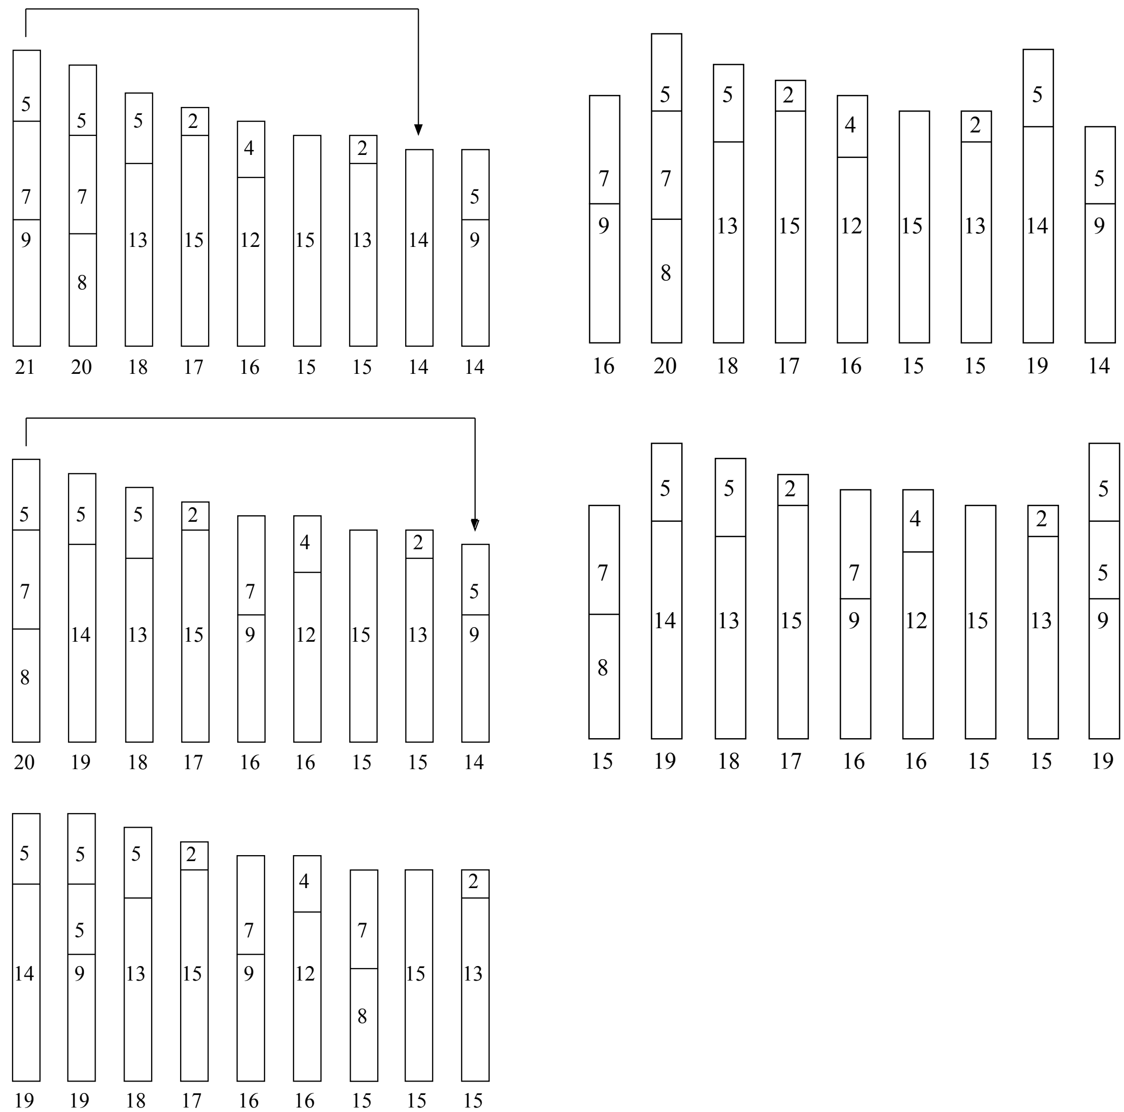

Convergence Time: (N p. 524)

Consider PNE construction:

> Satisfied task = task that cannot unilaterally reduce its cost
>
> Max-weight best response policy = move the unsatisfied job with maximum weight to the machine with minimum load

Theorem 20.6 *Let $A : [n] \rightarrow [m]$ denote any assignment of n tasks to m identical machines. Starting from A, the max-weight best response policy reaches a pure Nash equilibrium after each agent was activated at most once.*

Assigning tasks to machines in non-increasing weight order (LPT) is very effective. (N p. 529)

Previous example using LPT:

```
                        2
 2      4   5   5   5   7   5   7
15  15  14  13  13  12   9   9   8
```

Definition 20.4 (Price of anarchy) For $m \in \mathbb{N}$, let $\mathcal{G}(m)$ denote the set of all instances of load balancing games with m machines. For $G \in \mathcal{G}(m)$, let $\mathrm{Nash}(G)$ denote the set of all strategy profiles being a Nash equilibrium for G, and let $\mathrm{opt}(G)$ denote the minimum social cost over all assignments. Then the price of anarchy is defined by

$$\mathrm{PoA}(m) = \max_{G \in \mathcal{G}(m)} \max_{P \in \mathrm{Nash(G)}} \frac{\mathrm{cost}(P)}{\mathrm{opt}(G)}.$$

Theorem 20.5 *Consider an instance G of the load balancing game with n tasks of weight $w_1, \ldots, w_n$ and m identical machines. Let $A : [n] \to [m]$ denote any Nash equilibrium assignment. Then, it holds that*

$$cost(A) \leq \left(2 - \frac{2}{m+1}\right) \cdot opt(G).$$

Mixed Nash equilibrium:

Strategy profile: Probability p_i^j of each task i being assigned to each machine j

For job i, the expected cost on machine j is: $c_i^j = w_i + \sum_{k \neq i} w_k \, p_k^j$ (N p. 519, (20.1))

Proposition 20.1 *A strategy profile P is a Nash equilibrium if and only if* $\forall i \in [n] : \forall j \in [m] : p_i^j > 0 \Rightarrow \forall k \in [m] : c_i^j \leq c_i^k$.

Fully mixed Nash equilibrium: Every task is assigned with equal probability to each of the machines $p_i^j = \frac{1}{m}$. (N p. 529)

Suppose all n tasks have weight 1. Optimal makespan is $\lceil \frac{n}{m} \rceil$.

Proposition 20.11 *Suppose that $n \geq 1$ balls are placed independently, uniformly at random into $m \geq 1$ bins. Then the expected maximum occupancy is*

$$\Theta\left(\frac{\ln m}{\ln\left(1 + \frac{m}{n} \ln m\right)}\right).$$

Theorem 20.13 *Consider an instance G of the load balancing game with n tasks of weight $w_1, \ldots, w_n$ and m identical machines. Let $P = (p_i^j)_{i \in [n], j \in [m]}$ denote any Nash equilibrium strategy profile. Then, it holds that*

$$cost(P) = \mathcal{O}\left(\frac{\log m}{\log \log m}\right) \cdot opt(G).$$

2.C. (More) Potential Games

Market Sharing Game (KP p. 158)

There are k NBA teams, and each of them must decide in which city to locate.[1] Let v_j be the profit potential, i.e., the number of basketball fans, of city j. If ℓ teams select city j, they each obtain a utility of v_j/ℓ. See Figure 8.10.

Proposition 8.3.1. *The market sharing game is a potential game and hence has a pure Nash equilibrium. (See Exercise 4.18.)*

For any set S of cities, define the total value

$$V(S) := \sum_{j \in S} v_j.$$

Assume that $v_j \geq v_{j+1}$ for all j. Clearly $S^* = \{1, \ldots, k\}$ maximizes $V(S)$ over all sets of size k.

We use $\mathbf{c} = (c_1, \ldots, c_k)$ to denote a strategy profile in the market sharing game, where c_i represents the city chosen by team i. Let $S := \{c_1, \ldots, c_k\}$ be the set of cities selected.

Lemma 8.3.2. *Let $\mathbf{c}$ and $\tilde{\mathbf{c}}$ be any two strategy profiles in the market sharing game, where the corresponding sets of cities selected are S and $\tilde{S}$. Denote by $u_i(c_i, \mathbf{c}_{-i})$ the utility obtained by team i if it chooses city c_i and the other teams choose cities $\mathbf{c}_{-i}$. Then*

$$\sum_i u_i(\tilde{c}_i, \mathbf{c}_{-i}) \geq V(\tilde{S}) - V(S).$$

Proof. Let $\tilde{c}_i \in \tilde{S} \setminus S$. Then $u_i(\tilde{c}_i, \mathbf{c}_{-i}) = v_{\tilde{c}_i}$. Thus

$$\sum_i u_i(\tilde{c}_i, \mathbf{c}_{-i}) \geq V(\tilde{S} \setminus S) \geq V(\tilde{S}) - V(S).$$

Sum over all cities.
Different vector for each team.
For i, the term is $\dfrac{v_{\tilde{c}_i}}{1 + \text{multiplicity of } \tilde{c}_i \text{ in } \mathbf{c}_{-i}}$

If $\tilde{c}_i \in \tilde{S} \setminus S$ then $u_i(\tilde{c}_i, \mathbf{c}_{-i}) = v_{\tilde{c}_i}$ (multiplicity is 0)

Theorem 8.3.4. *Suppose that $\mathbf{c} = (c_1, \ldots, c_k)$ is a Nash equilibrium in the market sharing game and $S := S(\mathbf{c})$ is the corresponding set of cities selected. Then the price of anarchy is at most 2; i.e., $V(S^*) \leq 2V(S)$.*

Proof is similar to R p. 165 in the notion of "disentangling" (KP Lemma 8.3.2)

Same vectors

Proof. We claim that

$$V(S) = \sum_i u_i(c_i, \mathbf{c}_{-i}) \geq \sum_i u_i(c_i^*, \mathbf{c}_{-i}) \geq V(S^*) - V(S)$$

By definition NE to left Lemma 8.3.2

For i, the term is $\dfrac{v_{c_i}}{1 + \text{multiplicity of } c_i \text{ in } \mathbf{c}_{-i}}$ For i, the term is $\dfrac{v_i}{1 + \text{multiplicity of } i \text{ in } \mathbf{c}_{-i}}$

Suppose 5 teams with cities (10, 8, 6, 4, 2). S^* is all cities. $\mathbf{c}^*$ is (1, 2, 3, 4, 5).

$\mathbf{c}$ is (1, 1, 2, 2, 3). (Convenient to order indices.)

Left (NE) Σ is 5 + 5 + 4 + 4 + 6 = 24. Right Σ is 10/2 + 8/3 + 6/2 + 4 + 2 = 16.667.

(Facility) Location Games (R 14.2 p. 188; N 19.4)

Model:

L, set of possible locations

k agents (players), each agent i choosing exactly one location from $L_i \subseteq L$. (Pointless, payoff-wise, for two agents to choose the same location . . .)

M, set of markets, each market j having a public value v_j (i.e. maximum price)

c_{lj}, the cost (possibly ∞) of serving market j from location l

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 T is the set of chosen locations and i chooses $\ell \in T$

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

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 T and j.

Maximize social welfare: $$W(\mathbf{s}) = \sum_{j \in M} (v_j - d_j(\mathbf{s}))$$

$$(v_j - d_j(\mathbf{s})) = v_j - d_j^{(2)}(\mathbf{s}) + d_j^{(2)}(\mathbf{s}) - d_j(\mathbf{s})$$

where $d_j(\mathbf{s})$ is the minimum of v_j and the smallest cost between a chosen location and j

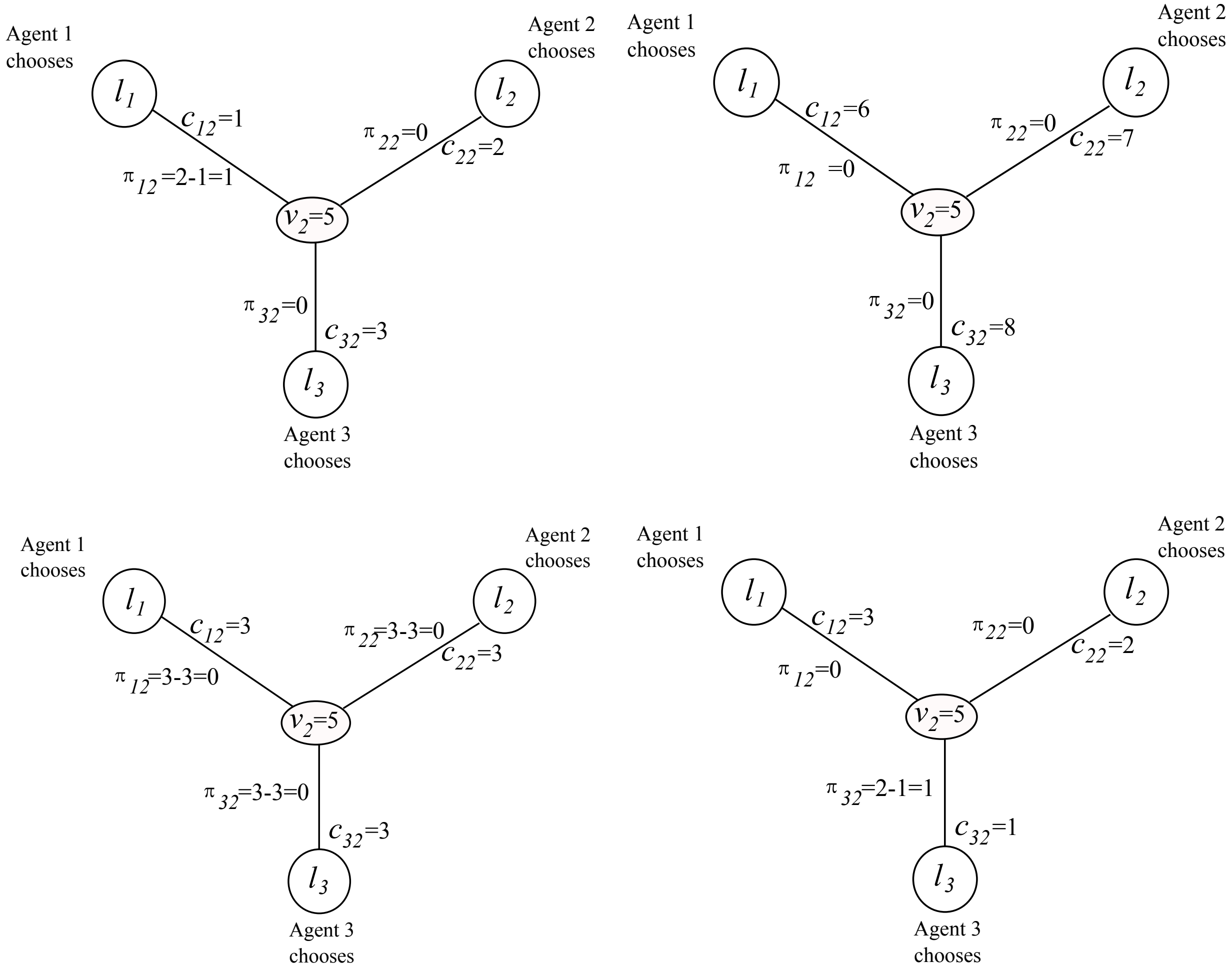

Example:

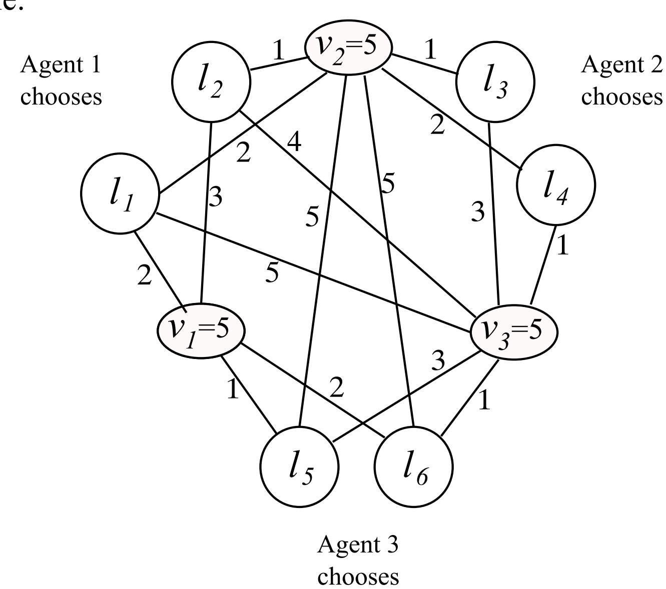

Agent Choices with Payoffs $\pi_i(\mathbf{s})$						$W(\mathbf{s})$	$\sum_{i=1}^{k} \pi_i(\mathbf{s})$	
l_1 (0)		l_3 (1)		l_5 (1)		10	2	
l_1 (0)		l_3 (1)			l_6 (2)	11	3	PNE
l_1 (0)			l_4 (2)	l_5 (1)		11	3	
l_1 (0)			l_4 (0)		l_6 (0)	10	0	
	l_2 (0)	l_3 (0)		l_5 (2)		10	2	
	l_2 (0)	l_3 (0)			l_6 (3)	11	3	PNE
	l_2 (1)		l_4 (2)	l_5 (2)		12	5	PNE, OPT
	l_2 (1)		l_4 (0)		l_6 (1)	11	2	

Properties:

At least one pure Nash equilibrium.

The POA of every location game is at least 1/2. (R p. 191, Theorem 14.1)

(P1) For every strategy profile, sum of the agents' payoffs is at most the social welfare.

(P2) An agent's payoff equals the extra welfare from their taking on their location.

(P3) Social welfare, W, is:

Monotone: $T_1 \subseteq T_2 \subseteq L$ implies $W(T_1) \leq W(T_2)$

Submodular: For any location $l \in L$ and $T_1 \subseteq T_2 \subseteq L$ implies

$$W(T_2 \cup \{ l \}) - W(T_2) \leq W(T_1 \cup \{ l \}) - W(T_1)$$

(Including a location earlier has more impact than including it later)

(Aside: "FAP", `https://dl-acm-org.ezproxy.uta.edu/doi/10.1145/1478873.1478955`)

Proof of Theorem 14.1 (much like the proof of Theorem 12.3)

2.D. SMOOTH GAMES (R 14.3 p. 194; KP p. 162 - leaves out "for all strategy profiles" detail)

(This is useful for R p. 223-226. Restatement of (12.10) R p. 168.)

Comparison between an arbitrary strategy profile (**s**) and an optimal one (**s***):

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

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

   for all strategy profiles $\mathbf{s}, \mathbf{s}^*$. 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 $\mathbf{s}$.

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

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

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

Robust Bounds for POA

Pure Nash Equilibria: In a (λ, μ)-smooth cost-minimization game with $\mu < 1$, every PNE **s** has cost at most $\frac{\lambda}{1-\mu}$ times that of an optimal outcome **s***. (R p. 195; KP p. 163)

Coarse Correlated Equilibria in Smooth Games: In every (λ, μ)-smooth cost-minimization game with $\mu < 1$, the POA of CCE is at most $\frac{\lambda}{1-\mu}$. (R p. 196, Theorem 14.4; KP p. 163)

Approximate Pure Nash Equilibria: In every (λ, μ)-smooth cost-minimization game with $\mu < 1$, for every , the POA of ∈-PNE is at most $\frac{(1+\epsilon)\lambda}{1-\mu(1+\epsilon)}$. (R p. 198 Theorem 14.6)

NETWORK COST-SHARING (FORMATION) GAMES

Model:

Directed (or undirected) graph with each edge e having nonnegative cost γ_e.

k agents, each with origin o_i and destination d_i along with strategy set of $o_i \ldots d_i$ paths

Agents using an edge e split the cost. f_e is the number of agents using e

Cost of path chosen by agent i is $C_i(\mathbf{P}) = \sum_{e \in P_i} \frac{\gamma_e}{f_e}$

Goal is to minimize total cost $\sum_{e \in E : f_e \geq 1} \gamma_e$

Example:

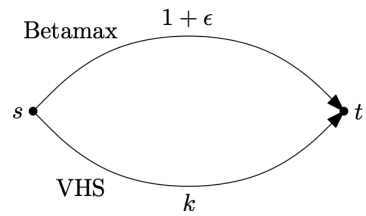

Figure 1: VHS or Betamax. The price of anarchy in a network cost-sharing game can be as large as the number k of players.

Example:

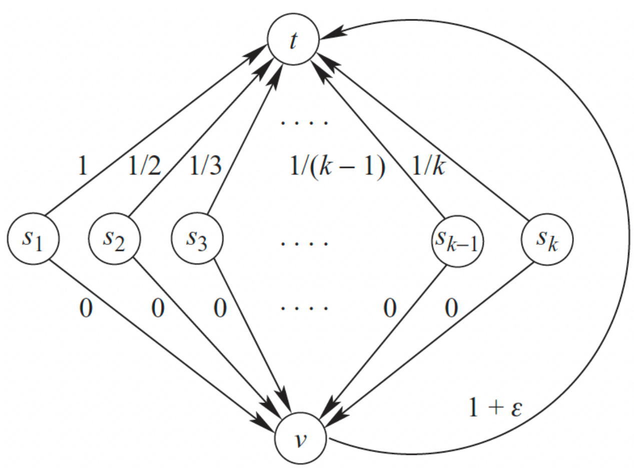

k agents could split cost $1 + \varepsilon$ by using vertex v, but agents will use their alternate edges 1, 1/2, 1/3, 1/4, . . . as the unique pure Nash equilibrium.

POA of $H_k = \ln k$ for network cost-sharing games.

PRICE OF STABILITY (N p. 445)

Price of Stability = cost of best equilibrium / cost of optimal outcome

POA in the previous example is the worst-case for all network cost-sharing games (R p. 206; KP p. 157)

POA OF STRONG NASH EQUILIBRIA

A strong Nash equilibrium is a pure Nash equilibrium having no improving coalition, i.e. a set of agents with a beneficial deviation, i.e. no deviating agent is worse-off and at least one deviating agent is strictly better-off. (Some call this Pareto Optimal. Some say a strong NE requires all defectors to benefit.)

If a strong NE exists for a network cost-sharing game with k agents, it has a cost at most H_k times that of an optimal outcome. (R p. 210, Theorem 15.3)