Applied Algorithms

Algorithmic Game Theory

Part I

CS versus Economics

- What computer science has brought:
  - New problems or new light on old problems: the Internet -- the fundamental, large-scale arena for resource sharing among parties with diverse and selfish interests.
  - New techniques: randomization, reductions.
  - Different performance measures, like worst-case analysis.
  - Computational complexity.
  - Auctions for digital goods: Identical items available in unlimited supply.
  - Routing and multicommodity flow applications.
  - Cost sharing problems.

Example: Network Formation Games

[Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, Roughgarden 2004]

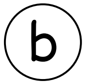

locations.

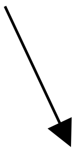

communication channels.

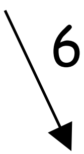

cost of creating the channel.

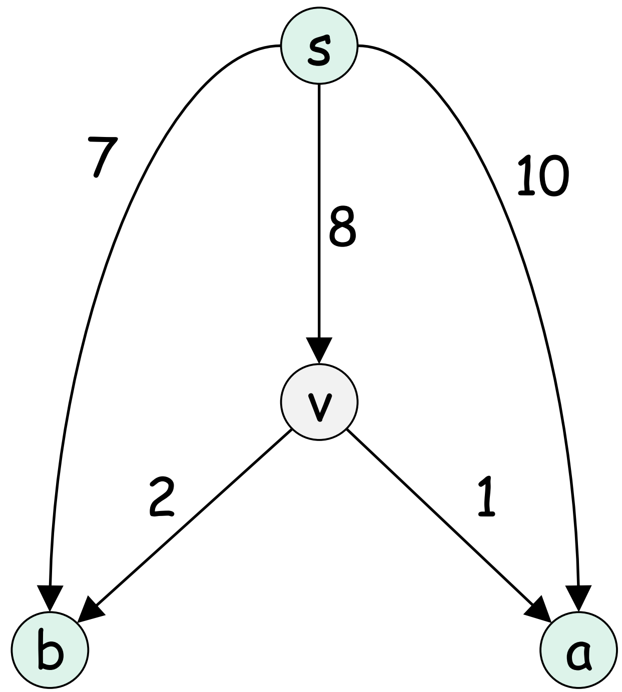

Players that need to transmit messages between locations in the network.

A network formation game

Example: Two players need to transmit messages from s

Player 1  needs to reach a

Player 2  needs to reach b

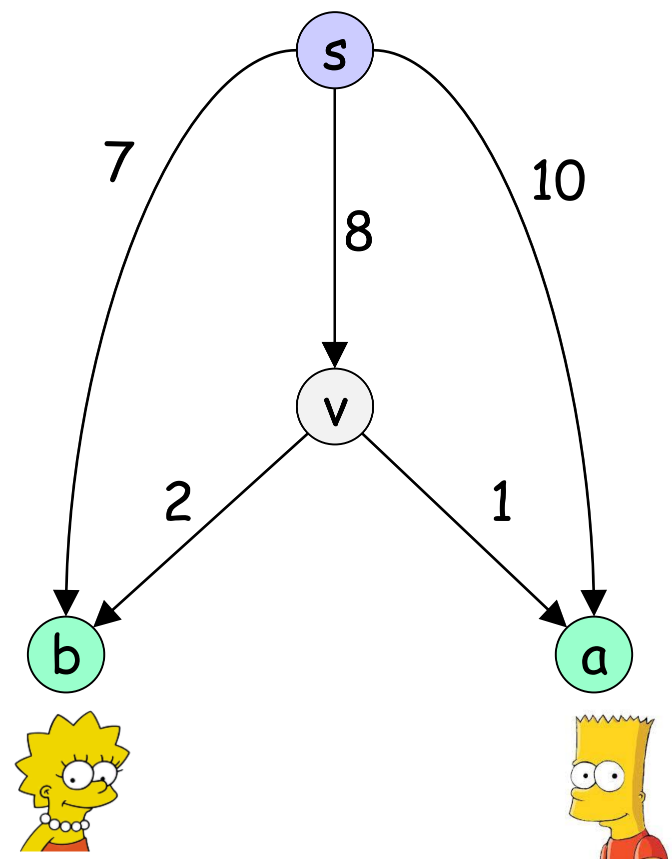

A network formation game

Example: Two players need to transmit messages from s

Player 1  needs to reach a

Player 2  needs to reach b

The strategy space of  :

{ {⟨s,v⟩,⟨v,a⟩} , {⟨s,a⟩} }

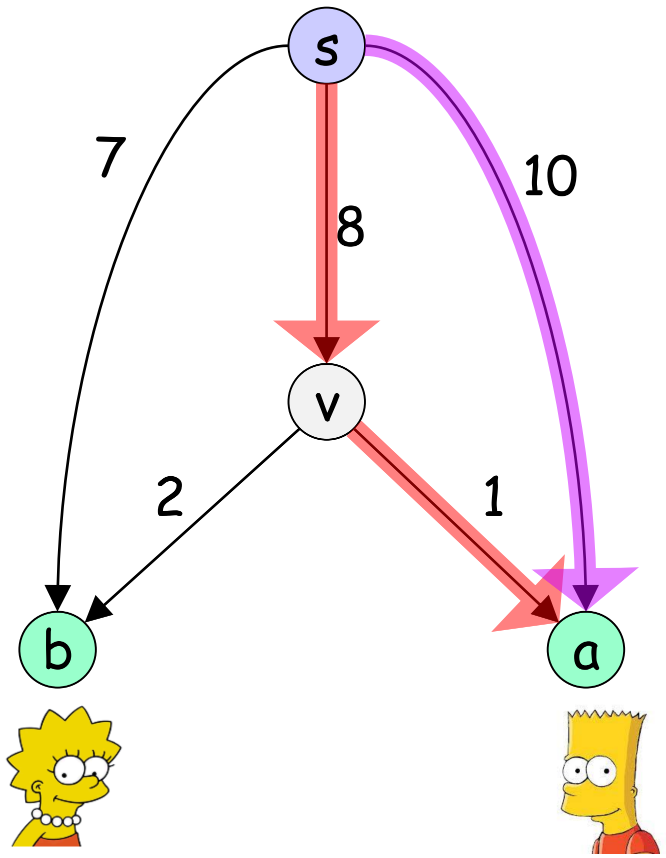

A network formation game

Example: Two players need to transmit messages from s

Player 1  needs to reach a

Player 2  needs to reach b

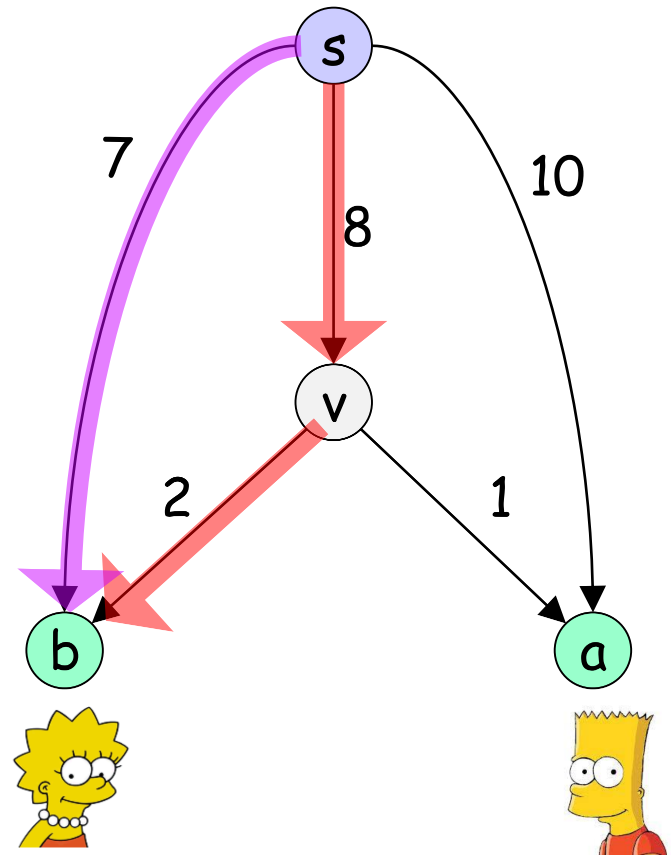

The strategy space of  :
{ {⟨s,v⟩,⟨v,a⟩} , {⟨s,a⟩} }

The strategy space of 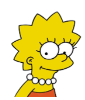 :
{ {⟨s,b⟩} , {⟨s,v⟩, ⟨v,b⟩} }

A profile is a choice of strategy for each player.

Four possible profiles in our example:

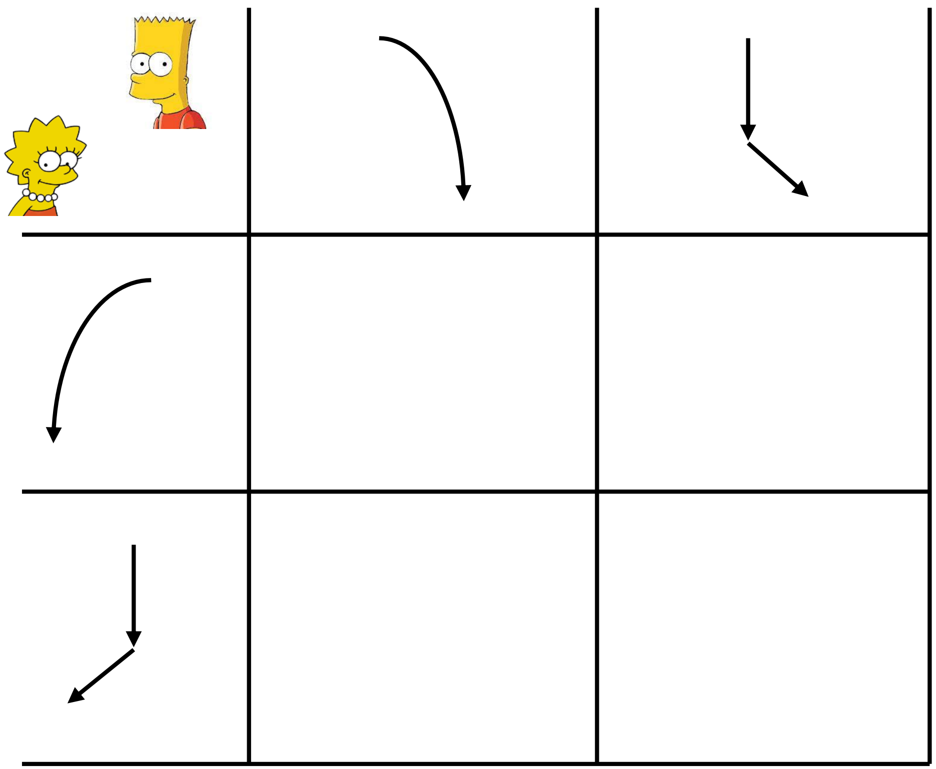

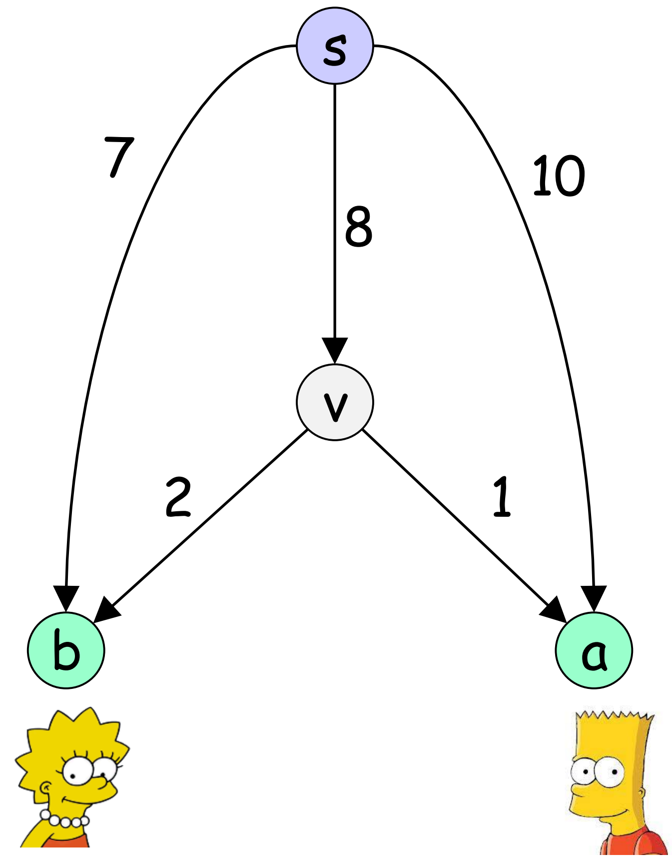

A profile is a choice of strategy for each player.

Four possible profiles in our example:

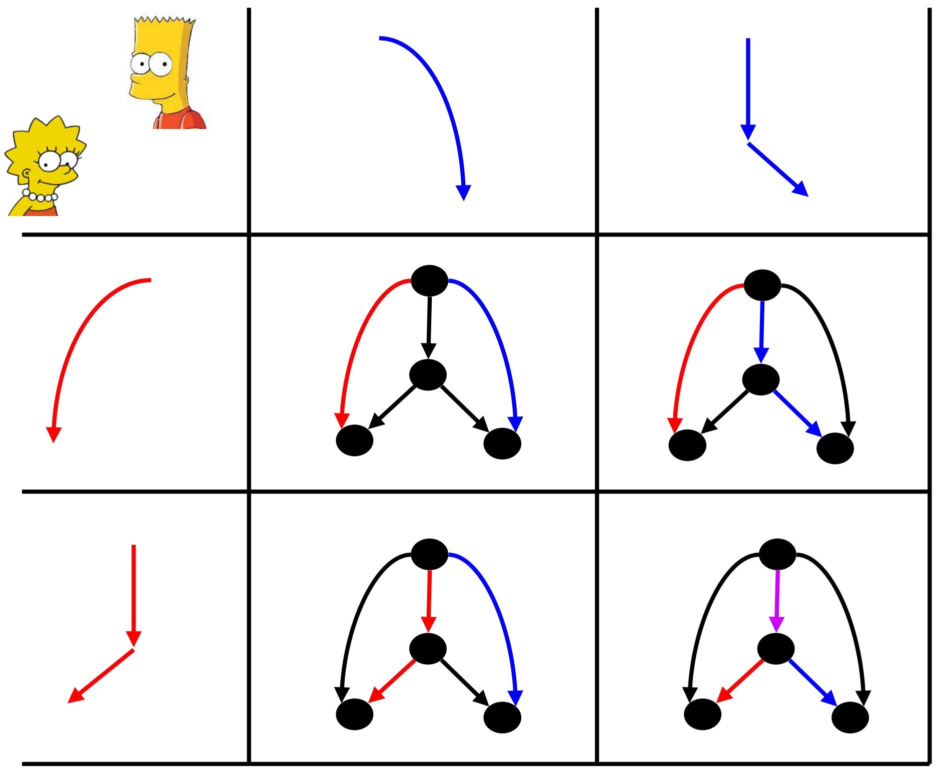

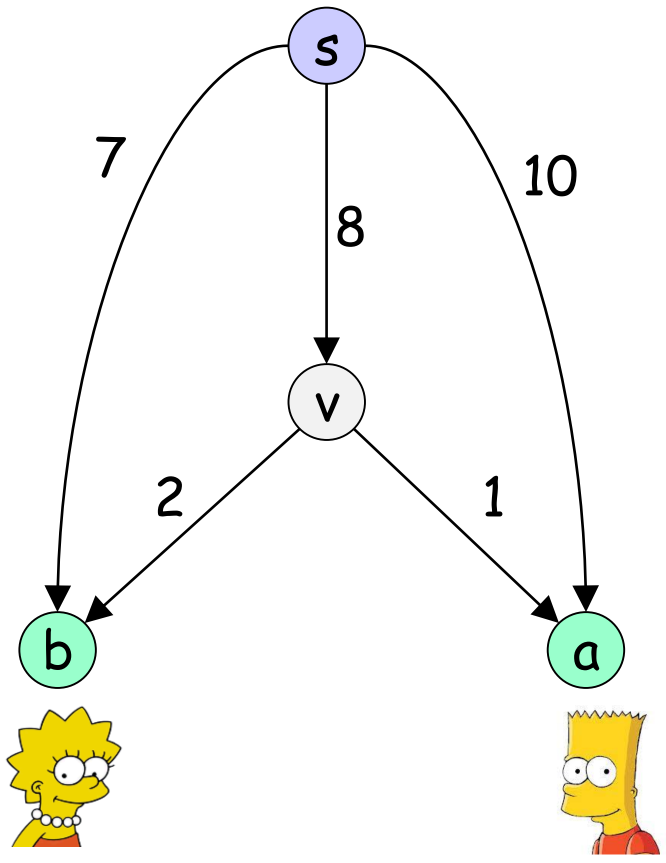

What are the costs?

A profile is a choice of strategy for each player.

Four possible profiles in our example:

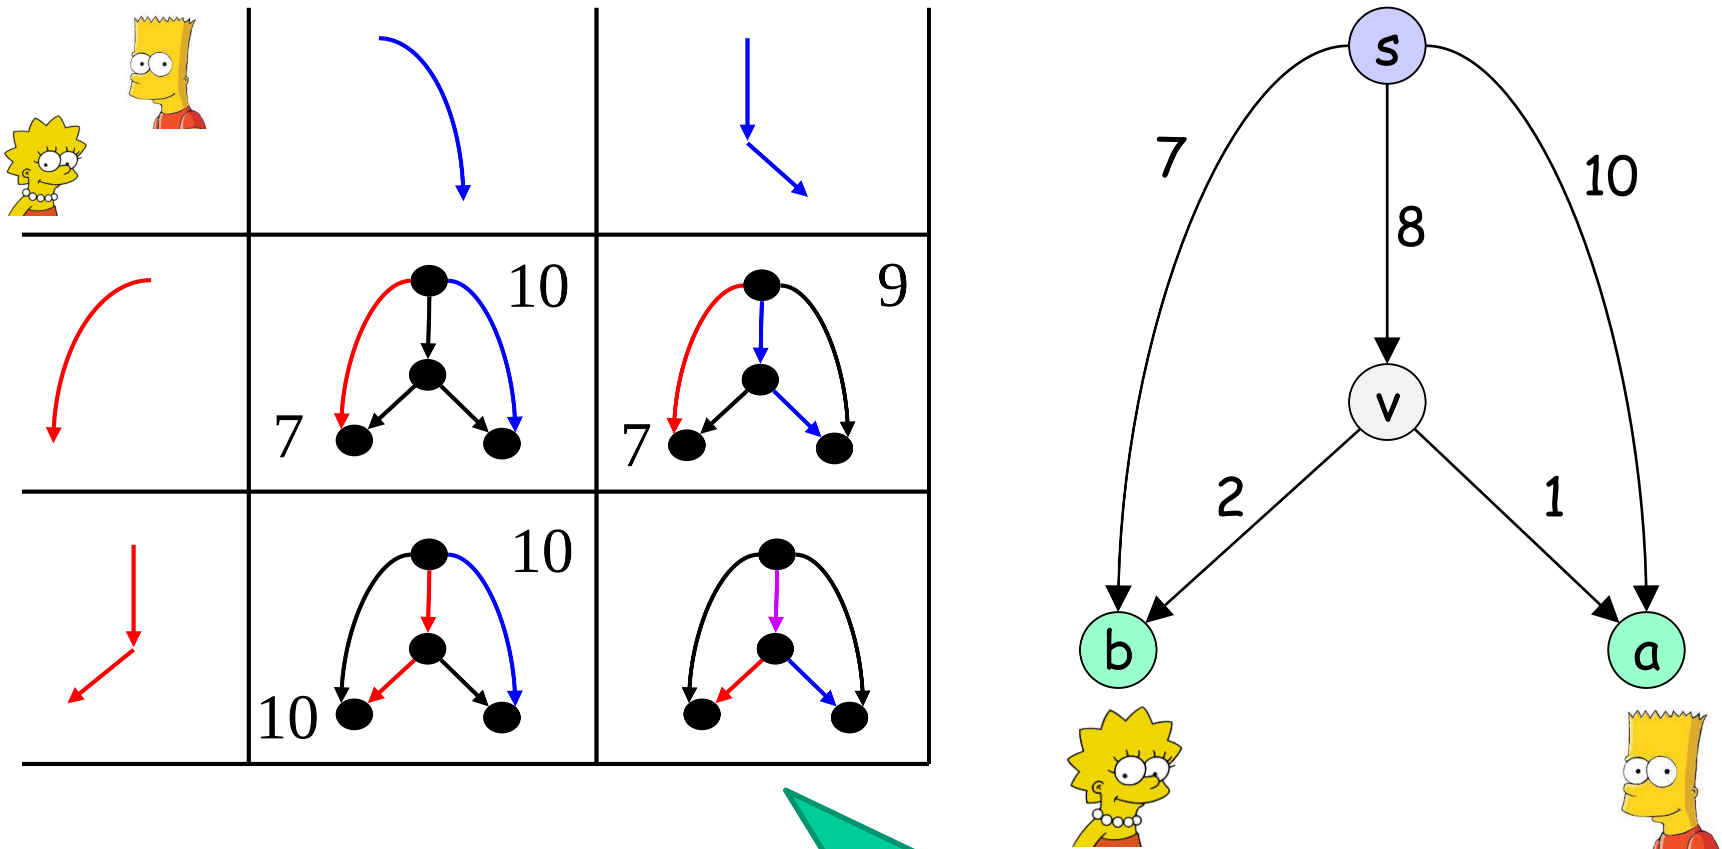

What are the costs?

How is a cost shared?

Players that use the same channel share its cost:

$$\frac{8}{2}+2=6$$

$$\frac{8}{2}+1=5$$

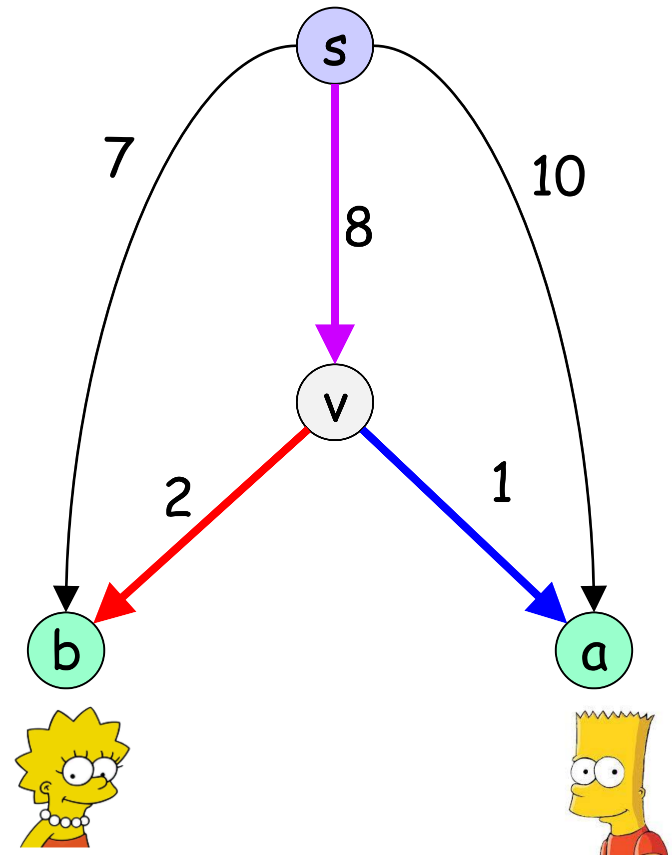

A profile is a choice of strategy for each player.

Four possible profiles in our example:

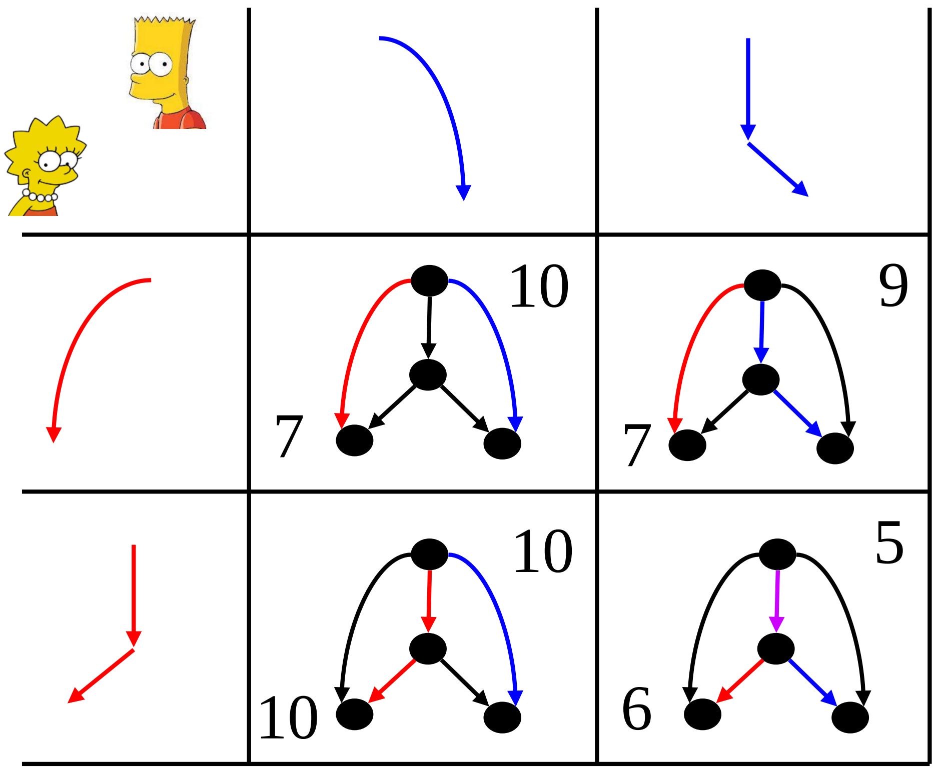

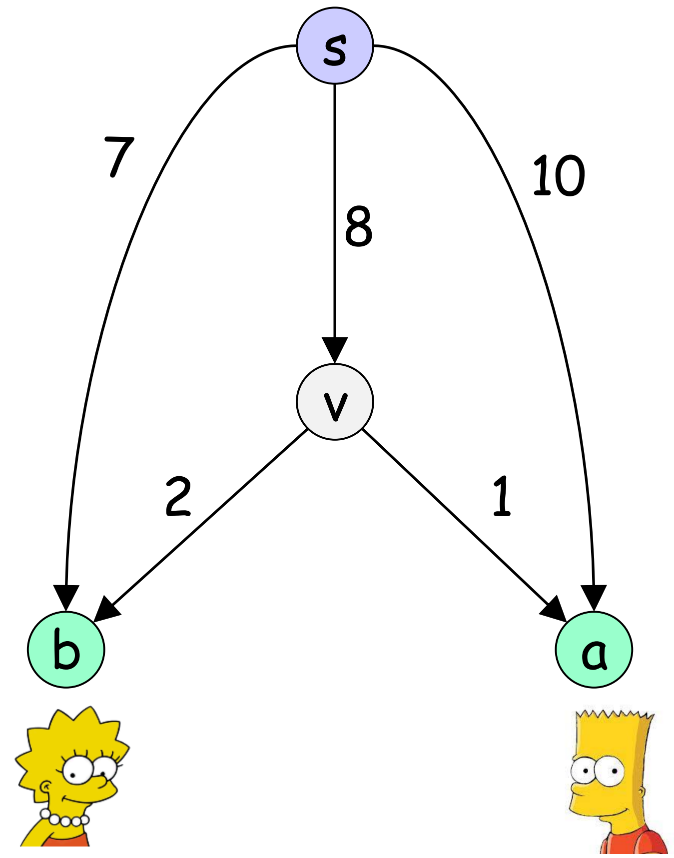

Best Response Dynamics (BRD)

- A local search method.
- Players proceed in turns, each performing a selfish improving step.
- In many scenarios, BRD lead to a pure Nash equilibrium.

A stable profile in which no one has an improving step.

Best response dynamics.

Example: starting from 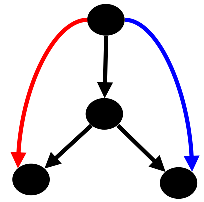

Cost for  :10

Cost for  :7

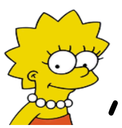, want to change strategy?

No, 7 < 10

, want to change strategy?

Sure ,9 < 10

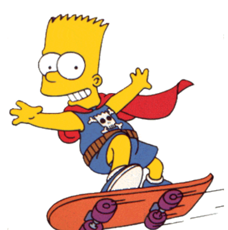

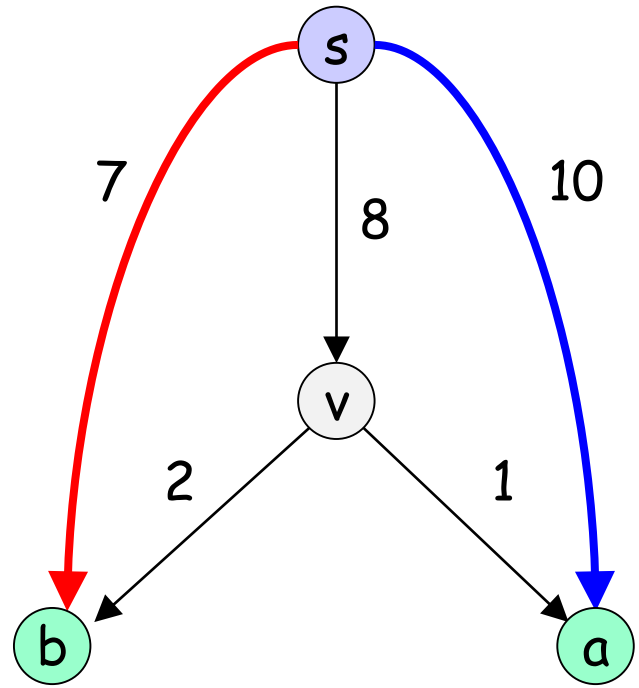

Best response dynamics.

Cost for  :9

Cost for  :7

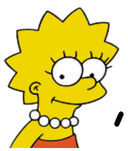, want to change strategy?

Yes, 6 < 7

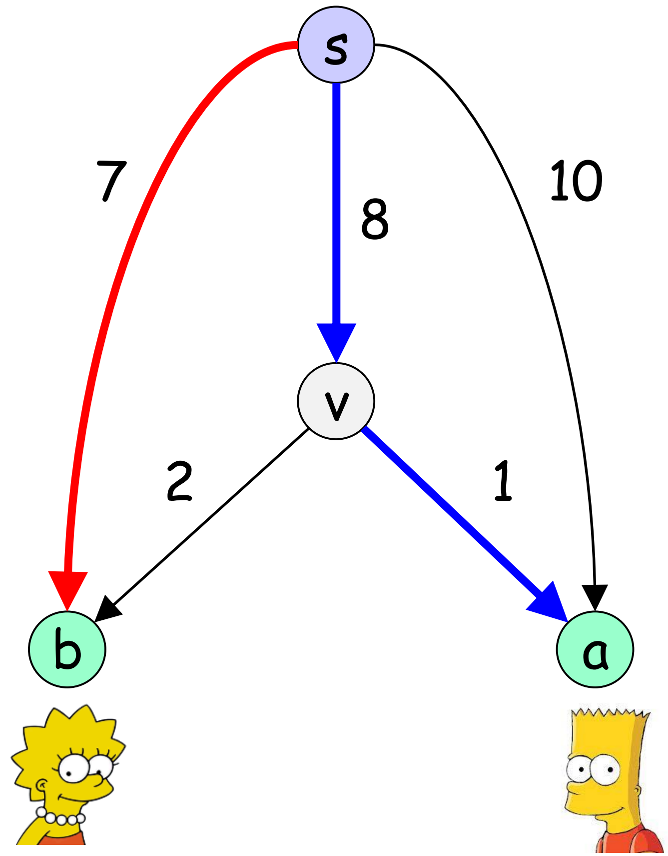

Best response dynamics.

Cost for  :5

Cost for  :6

, want to change strategy?

No, 5 < 10

, want to change strategy?

No, 6 < 7

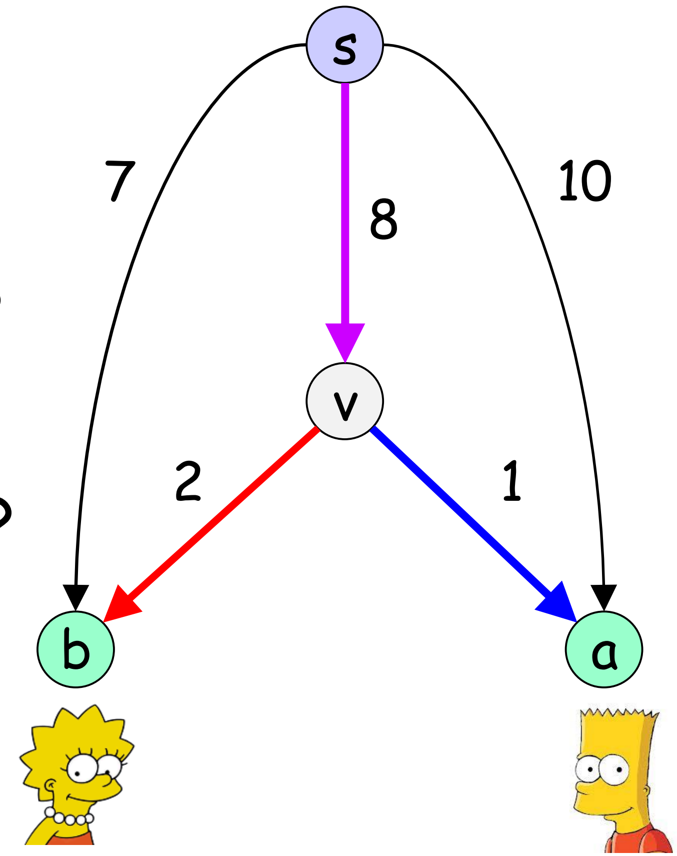

 

BRD halts, we've reached a NE.

Network Formation Games - Formal definition

- Given a directed graph $G = (V, E)$ with edge costs $c_e \geq 0$, a source node s, and k agents located at terminal nodes $t_1, \ldots, t_k$. Agent j must construct a path P_j from node s to its terminal t_j.

- Fair share: If x agents use edge e, they each pay c_e / x.
- The agents are selfish - each agent wants to minimize its cost.
- Agents might modify their selection as a response to actions of other agents.

Slides on multicast routing based on slides by Kevin Wayne.

Nash Equilibrium

- Best response dynamics (BRD): Each agent is continually prepared to improve its solution in response to changes made by other agents.

- Nash equilibrium: Solution where no agent has an incentive to switch.

- Fundamental question: When do Nash equilibria exist? Does BRD terminate?

Social Optimum

- Social optimum: Minimizes total cost to all agents.
- Observation: In general, there can be many Nash equilibria. Even when it's unique, it does not necessarily equal the social optimum.

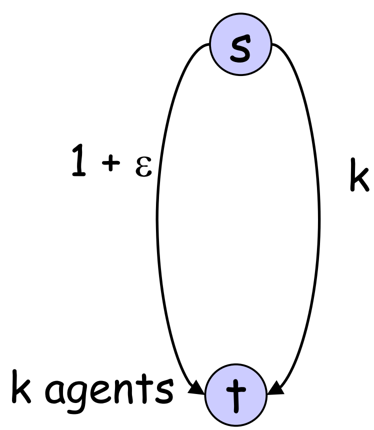

Social optimum = 1 + ε
Nash equilibrium A = 1 + ε
Nash equilibrium B = k

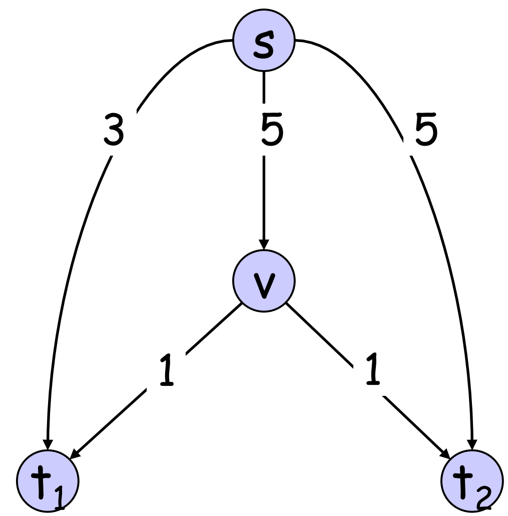

Social optimum = 7
Unique Nash equilibrium = 8

Price of Anarchy

- Price of anarchy: Ratio of worst Nash equilibrium to social optimum.

Theorem: In network formation games with k agents, PoA=k.

Proof:

- 1. PoA $\leq$ k. Assume by contradiction that in some NE, the PoA is more than k. This implies that some agent pays more than the social optimum. This agent can switch to his path in the social optimum and reduce its cost. Contradiction to NE.
- 2. For every $\varepsilon > 0$, there exists an instance with k agents for which PoA > k- ε.

  See left instance in previous slide.

Price of Stability

- Price of stability: Ratio of best Nash equilibrium to social optimum.
- Fundamental question: What is the price of stability?

Example: Price of stability = $\Theta(\log k)$.
Social optimum: Everyone Takes bottom paths (via 'a').
Unique Nash equilibrium: Everyone takes top paths.
Price of stability:

$H(k) / (1 + \varepsilon)$.

$1 + 1/2 + \ldots + 1/k$

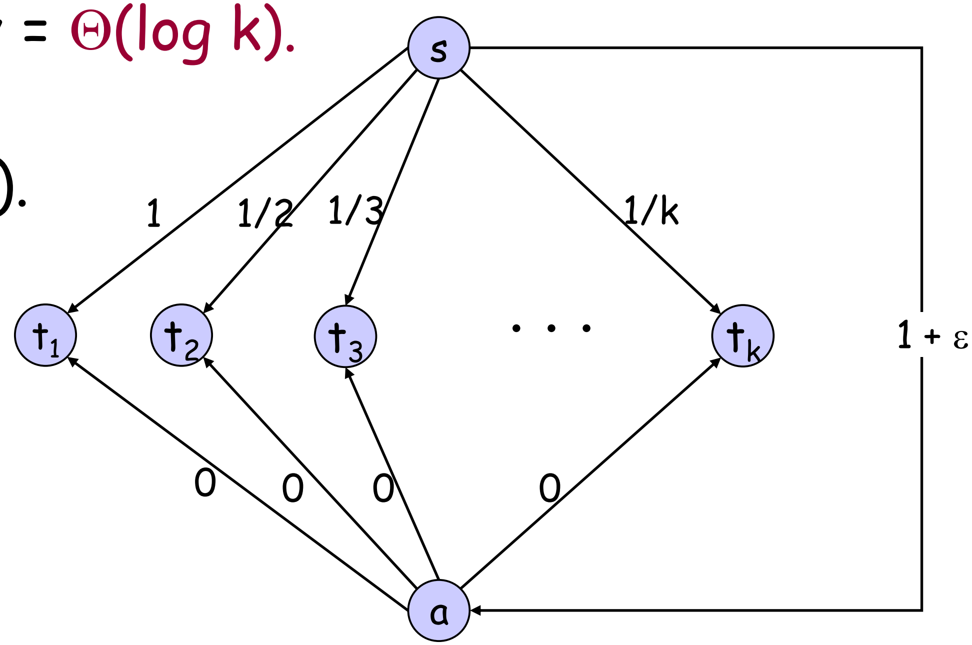

Finding a Nash Equilibrium

Theorem: The following algorithm terminates with a Nash equilibrium.

```
Best-Response-Dynamics(G,c) {
   Pick a path for each agent

   while (not a Nash equilibrium) {
      Pick an agent i who can improve by
       switching paths
      Switch path of agent i
   }
}
```

Finding a Nash Equilibrium

Proof Idea: Define a potential function over the possible solution sets. Show that the potential decrease whenever some agent improve.

Attempt 1:

Let $\Phi(s) = \Sigma_{j=1}^{k} cost(t_i)$ be the potential function.

A problem: The potential might increase when some agent improve.

Example: When all 3 agent use the right path, each pays 4/3 and the potential (total cost) is 4.
After one agent moves to the left path the potential increases to 5.

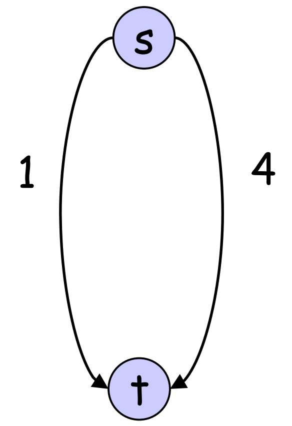

3 agents

Finding a Nash Equilibrium

Attempt 2:

Consider a set of paths $P_1, \ldots, P_k$.

- Let x_e denote the number of paths that use edge e.
- Let $\Phi(P_1, \ldots, P_k) = \sum_{e \in E} c_e \cdot H(x_e)$ be a potential function.

  H(0) = 0, $H(k) = \sum_{i=1}^{k} \frac{1}{i}$
- Consider agent j switching from path P_j to path P_j'.
- Agent j switches because

$$\underbrace{\sum_{f \in P_j' - P_j} \frac{c_f}{x_f + 1}}_{\text{newly incurred cost}} < \underbrace{\sum_{e \in P_j - P_j'} \frac{c_e}{x_e}}_{\text{cost saved}}$$

Finding a Nash Equilibrium

- Φ increases by

$$\sum_{f \in P_j' - P_j} c_f \left[H(x_f + 1) - H(x_f) \right] = \sum_{f \in P_j' - P_j} \frac{c_f}{x_f + 1}$$

- Φ decreases by

$$\sum_{e \in P_j - P_j'} c_e \left[H(x_e) - H(x_e - 1) \right] = \sum_{e \in P_j - P_j'} \frac{c_e}{x_e}$$

- Thus, net change in Φ is negative.
- Since there are only finitely many sets of paths, it implies that the algorithm terminates with a NE.

Bounding the Price of Stability

Claim: Let $C(P_1, \ldots, P_k)$ denote the total cost of selecting paths $P_1, \ldots, P_k$.
For any set of paths $P_1, \ldots, P_k$, we have

$$C(P_1,\ldots,P_k) \le \Phi(P_1,\ldots,P_k) \le H(k) \cdot C(P_1,\ldots,P_k)$$

Proof: Let x_e denote the number of paths containing edge e.

- Let E^+ denote set of edges that belong to at least one of the paths.

$$C(P_1,\ldots,P_k) = \sum_{e \in E^+} c_e \le \underbrace{\sum_{e \in E^+} c_e H(x_e)}_{\Phi(P_1,\ldots,P_k)} \le \sum_{e \in E^+} c_e H(k) = H(k)\ C(P_1,\ldots,P_k)$$

Bounding the Price of Stability

Theorem: There is a Nash equilibrium for which the total cost to all agents exceeds that of the social optimum by at most a factor of $H(k)$.

Proof:

- Let $(P_1^*, \ldots, P_k^*)$ denote set of socially optimal paths.
- Perform BRD starting from P^*.
- Since Φ is monotone decreasing $\Phi(P_1, \ldots, P_k) \leq \Phi(P_1^*, \ldots, P_k^*)$.

$$C(P_1, \ldots, P_k) \underset{\substack{\uparrow \\ \text{previous claim} \\ \text{applied to } P}}{\leq} \Phi(P_1, \ldots, P_k) \leq \Phi(P_1^*, \ldots, P_k^*) \underset{\substack{\uparrow \\ \text{previous claim} \\ \text{applied to } P^*}}{\leq} H(k) \cdot C(P_1^*, \ldots, P_k^*)$$

Summary. k-agent multicast

- Existence: Nash equilibria always exist for k-agent multicast routing with fair sharing.

- Price of stability: Best Nash equilibrium is never more than a factor of H(k) worse than the social optimum. For some networks this is tight.
- Price of anarchy: Any Nash equilibrium is never more than a factor of k worse than the social optimum. For some networks this is tight.

- Fundamental open problem: Find any Nash equilibria in poly-time.

Job Scheduling Games

- m machines
- n jobs.
- Each job has a length (load)
- Each job represents a selfish agent who attempts to optimize its own objective.
- A job that is not happy with its assignment can migrate from one machine to another.

The Scheduling Settings

- The jobs pay for their processing.
- The payment of each job is the total load on the machine it is assigned to.

Example:
m=3, n=7

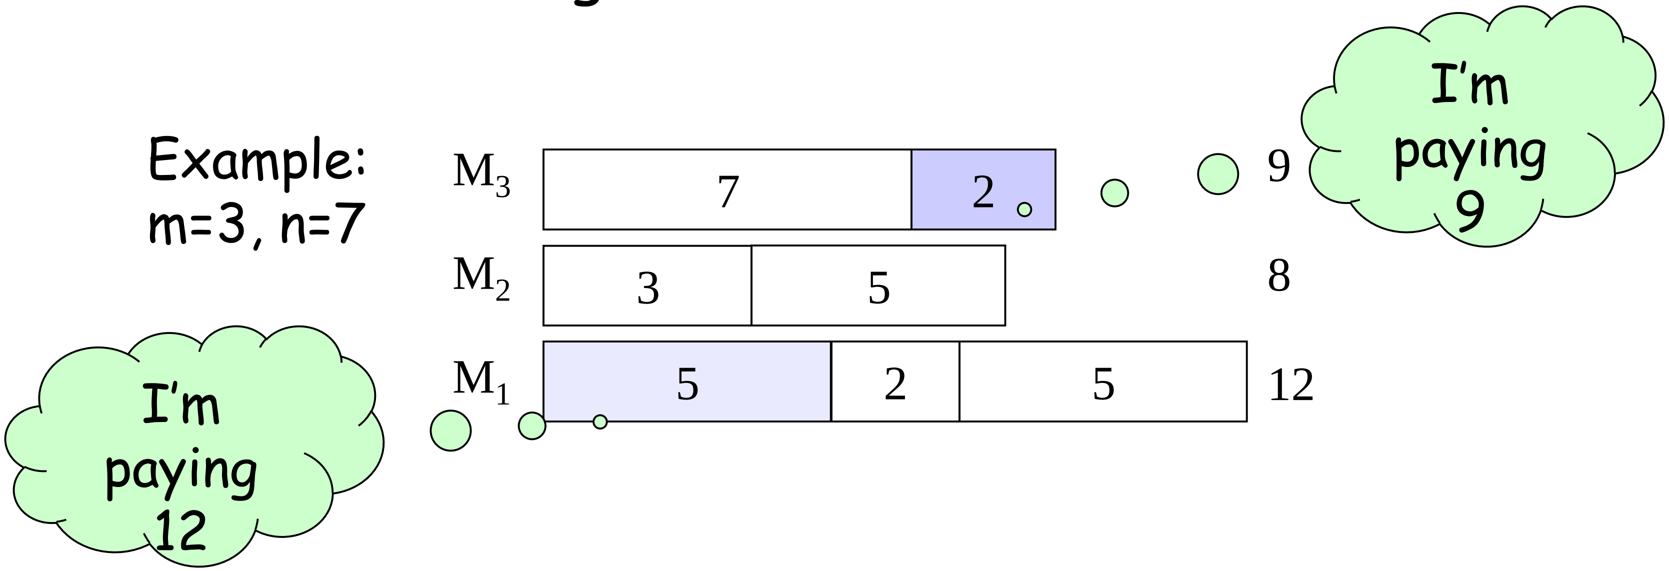

Note: In this payment scheme, the internal job order on each machine has no effect on the individual cost.

Motivation: routing on parallel links, round-robin, more

NE: no single-job migrations

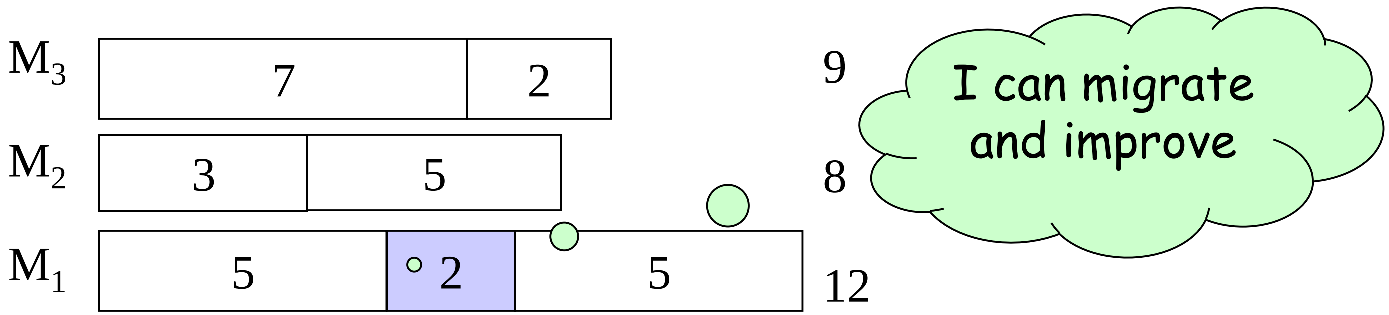

Not a Nash-equilibrium

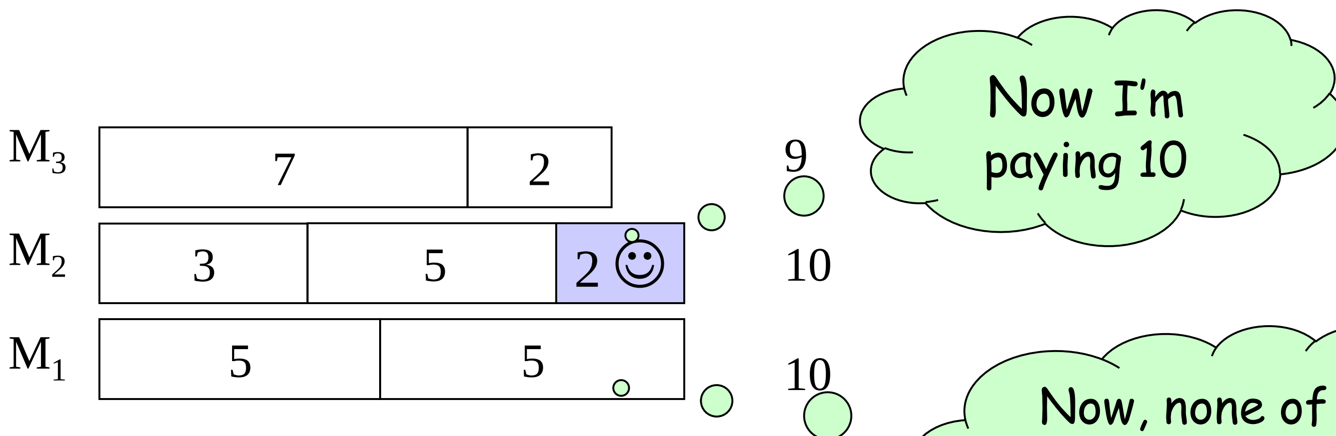

A Nash-equilibrium

NE: no single-job migrations

Notation:

L_i: Load on machine i (this is also the cost of all jobs assigned to i).

cost(j): cost of job j.

Theorem: NE always exist in job scheduling games, and can be found efficiently.

Proof: In class.

Equilibrium inefficiency

Social Optimum: A schedule in which the maximal cost of a job is minimized.
For a schedule s, $cost(s)=max_j\ cost(j)$.
This is equivalent to minimum makespan.

$$\boxed{\text{PoS} = \frac{\text{Minimumum makespan in the best NE}}{\text{Minimumum makespan}}}$$

Theorem: The Price of Stability in the job scheduling games is 1.
Proof: Note that a beneficial move does not increase the makespan. Therefore, by preforming best-response starting from any optimal assignment, we reach a NE whose makespan is equal to the optimum.

Equilibrium inefficiency

Theorem: The Price of Anarchy in the job scheduling game is $2 - \frac{2}{m+1}$.

Example (m=2): Consider the following assignment s on m=2 machines. s is a NE. its cost is 4.

M_2	2	2
M_1	1	1

This instance has an assignment with cost 3

M_2	2	1
M_1	2	1

$$\text{PoA} = \frac{4}{3} = 2 - \frac{2}{3}$$

$$\text{PoA} = \frac{\text{Minimumum makespan in the worst NE}}{\text{Minimumum makespan}}$$

Bounding the PoA

Theorem: The Price of Anarchy in the job scheduling game is $2 - \frac{2}{m+1}$.

Proof: Let s be any NE. Let i be the machine with the highest load (that is cost(s)=L_i) and let j be the shortest job on machine i.

If j is the only job on machine i then PoA=1 (why?).

Otherwise, $p_j \leq \frac{1}{2}$ cost(s).

Observation: For every machine i', $L_{i'} \geq L_i - p_j$

Therefore: $L_{i'} \geq L_i - p_j \geq L_i - \frac{1}{2}\text{cost}(s) = \frac{1}{2}\text{cost}(s)$.

Bounding the PoA

- So for every machine $i' \neq i$, $L_{i'} \geq \frac{1}{2}cost(s)$.

$$cost(OPT) \geq \frac{\sum_k pk}{m} = \frac{\sum_i L_i}{m} \geq \frac{cost(s)+(m-1)\frac{1}{2}cost(s)}{m} =$$

$$\frac{(m+1)cost(s)}{2m}.$$

Therefore, $\frac{cost(s)}{cost(OPT)} \leq 2 - \frac{2}{m+1}$

- The analysis is tight (example of m=2 can be generalized).

SE: no coalition migrations

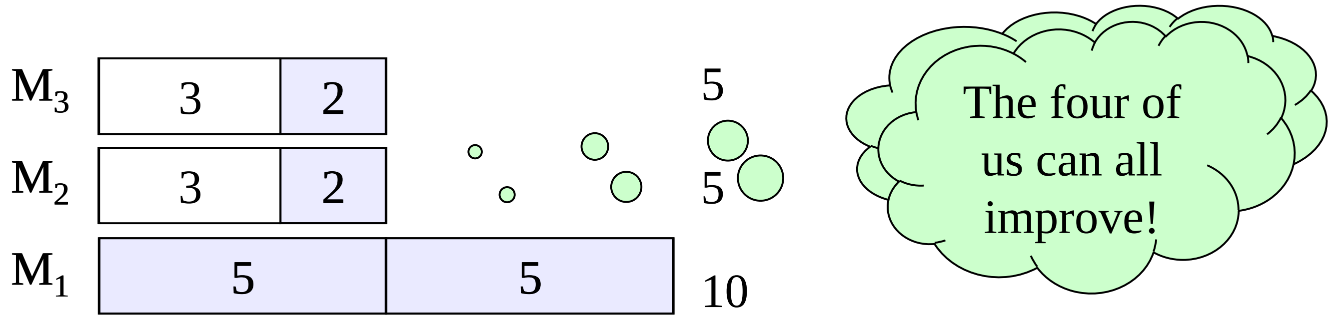

A Nash-equilibrium - But not a Strong NE.

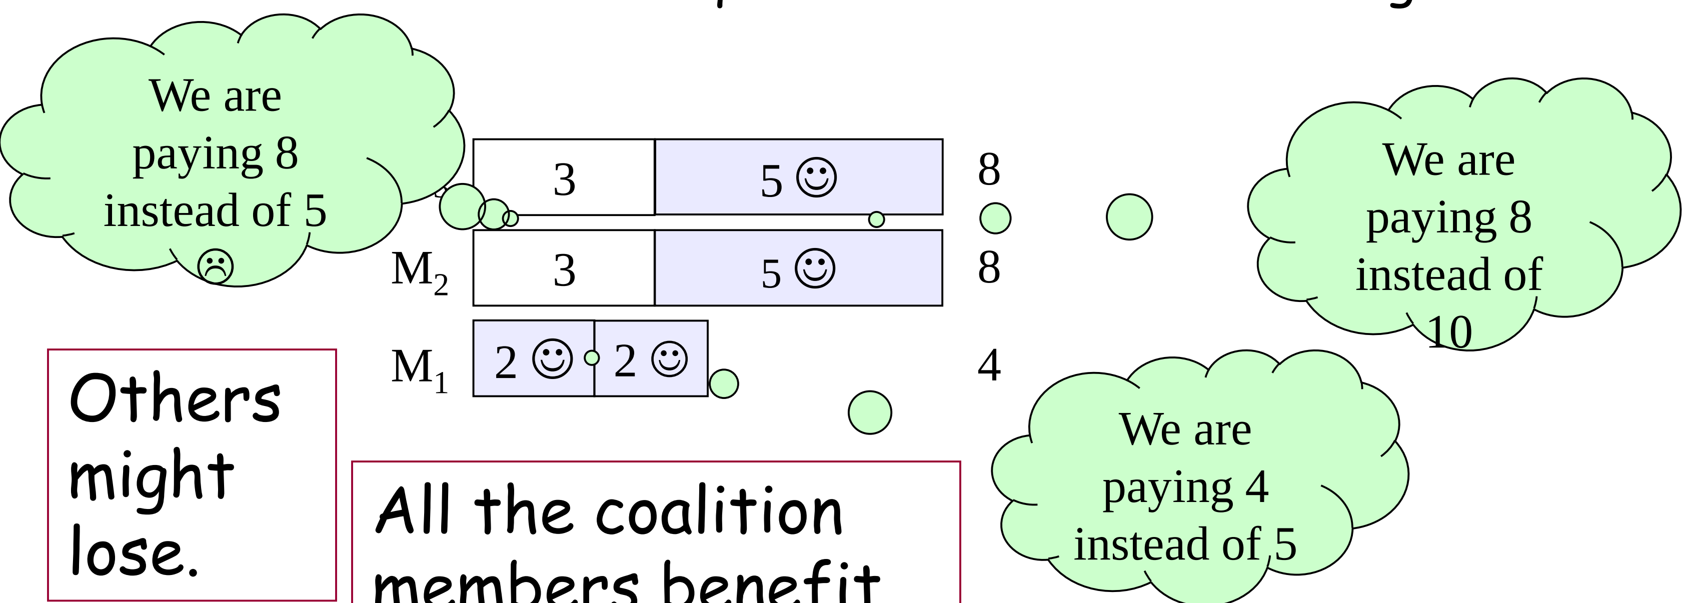

Others might lose.

All the coalition members benefit.

NE vs. SE

Nash Equilibrium: no single player can deviate and improve its utility.

The Global Social Cost might not be achieved due to:

- Players' selfishness
- Lack of coordination

Strong Equilibrium [Aumann'59]: No coalition can deviate and **strictly** improve the utility of **all** of its members

- Separates the effect of selfishness from lack of coordination
- May be a better prediction of rational behavior
- Most games do not admit Strong Eq.

SE: no coalition migrations

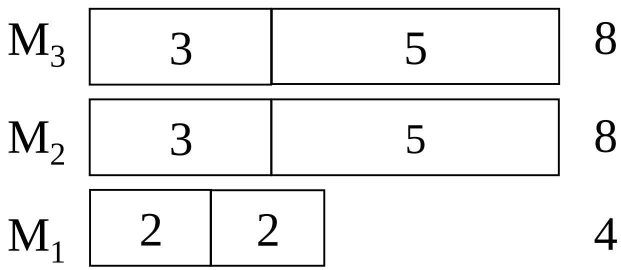

Theorem: SE always exists in job scheduling (even for unrelated machines) [AFM07]

Proof: In class.

Does a NE approximate SE?

In the example, each of the coalition jobs improves by a factor of 5/4 ($10 \rightarrow 8$, $5 \rightarrow 4$)

- Is there a bound on the improvement ratio exhibited by all coalition member?
- Is there a bound on the improvement ratio exhibited by a single coalition member?

In the example, the cost of some jobs is increased by a factor of 8/5 ($5 \rightarrow 8$)

- Is there a bound on the damage ratio exhibited by some non coalition member job?

Evaluating approximate SE

Given a configuration s, and a coalition Γ, we consider 3 measurements:

1. Minimum Improvement Ratio: $IR_{min}(s,\Gamma)$ is the minimal improvement ratio of some job in Γ.

> s is α-SE if there is no coalition Γ for which $IR_{min}(s,\Gamma) > \alpha$.

2. Maximum Improvement Ratio: $IR_{max}(s,\Gamma)$ is the maximal improvement ratio of some job in Γ.

3. Maximum Damage Ratio: $DR_{max}(s,\Gamma)$ is the maximal damage ratio of some job not in Γ.

Evaluating approximate SE

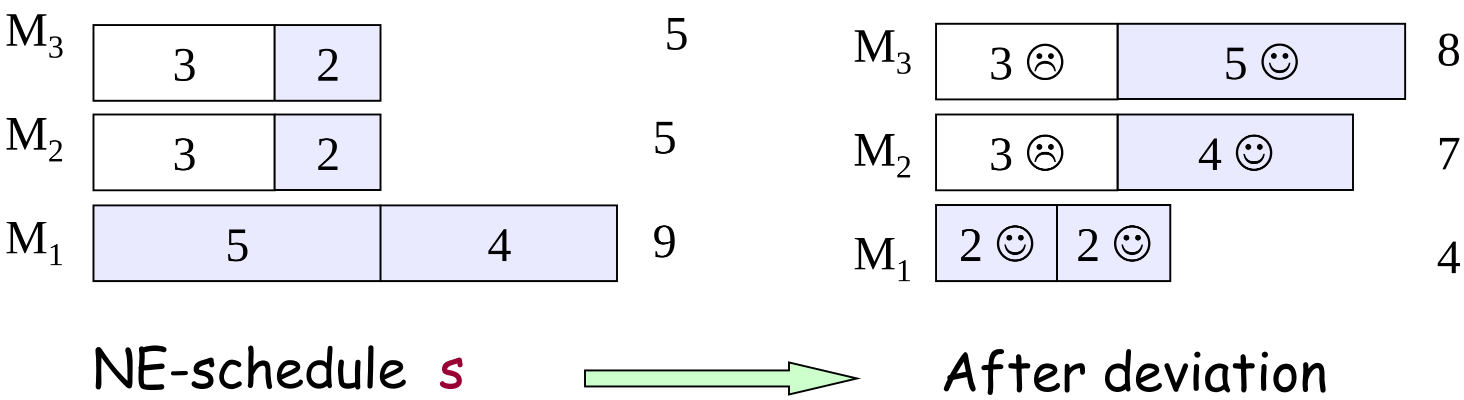

Consider the coalition $\Gamma=\{5,4,2,2,\}$

$IR_{min}(s,\Gamma)$ = 9/8=1.125

$IR_{max}(s,\Gamma)$ = 9/7≈1.28

$DR_{max}(s,\Gamma)$ = 8/5 = 1.6

LPT vs. any-NE

In our example, the initial configuration is a NE.

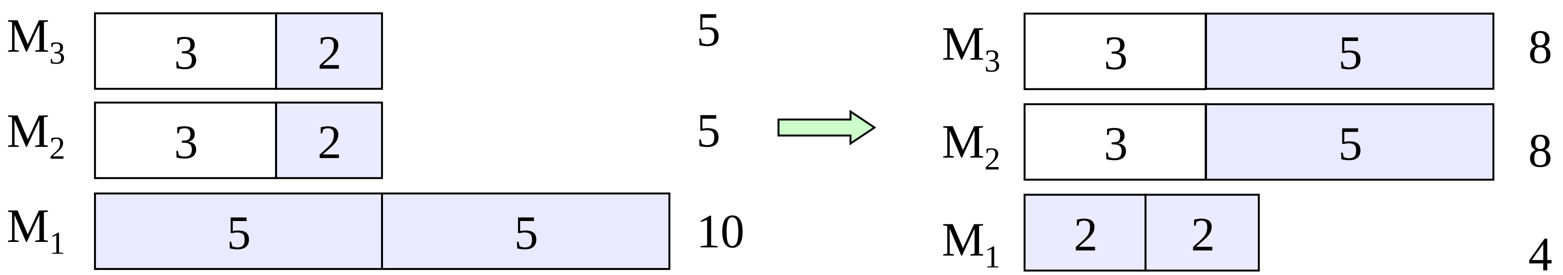

The same set of jobs under LPT:

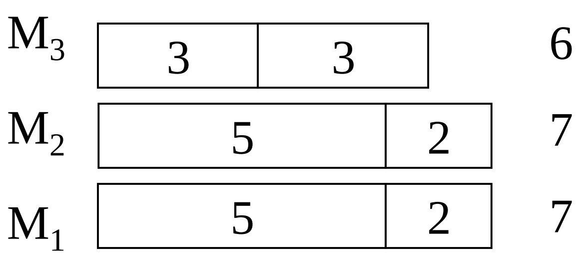

Known Results

	IR_{min}			IR_{max}		DR_{max}	
	Upper bound		Lower bound	Upper bound	Lower bound	Upper bound	Lower bound
	m=3	m≥3					
NE	5/4	$2-\frac{2}{m+1}$	5/4	unbounded		2	2
LPT	$\frac{1}{2}+\frac{\sqrt{6}}{4}$	$\frac{4}{3}-\frac{1}{3m}$	$\frac{1}{2}+\frac{\sqrt{6}}{4}$	5/3 (m=3)	$2-\frac{1}{m}$	3/2	3/2

~1.12

In any schedule produced by LPT, no coalition can improve the cost of all its members by ratio > $\frac{1}{2}+\frac{\sqrt{6}}{4}$ and this is tight.

In any NE schedule, no coalition-move can increase the cost of some job by a factor > 2, and this is tight.