Network Formation Games

Network Formation Games

- NFGs model distinct ways in which *selfish* agents might create and evaluate networks
- We'll see two models:
  - Global Connection Game
  - Local Connection Game
- Both models aim to capture two competing issues: players want
  - to minimize the cost they incur in building the network
  - to ensure that the network provides them with a high quality of service

Motivations

- NFGs can be used to model:
  - social network formation (edge represent social relations)
  - how subnetworks connect in computer networks
  - formation of networks connecting users to each other for downloading files (P2P networks)

Setting

- What is a stable network?
  - we use a NE as the solution concept
  - we refer to networks corresponding to Nash Equilibria as being *stable*
- How to evaluate the overall quality of a network?
  - we consider the *social cost*: the sum of players' costs
- **Our goal**: to bound the efficiency loss resulting from stability

Global Connection Game

E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, T. Roughgarden,
The Price of Stability for Network Design with Fair Cost Allocation, FOCS'04

The model

- $G=(V,E)$: directed graph
- c_e: non-negative cost of the edge $e \in E$
- k players
- player i has a source node s_i and a sink node t_i
- player i's goal: to build a network in which t_i is reacheable from s_i while paying as little as possible
- Strategy for player i: a path P_i from s_i to t_i

The model

- Given a strategy vector S, the constructed network will be $N(S) = \cup_i P_i$
- The cost of the constructed network will be shared among all players as follows:

$$\mathrm{cost}_i(S) = \sum_{e \in P_i} c_e / k_e(S)$$

$k_e(S)$: number of players whose path contains e

sometimes we write k_e instead of $k_e(S)$ when S is clear from the context

this cost-sharing scheme is called *fair* or *Shapley cost-sharing mechanism*

Remind

- We use Nash equilibrium (NE) as the solution concept
- A strategy vector S is a NE if no player has convenience to change its strategy
- Given a strategy vector S, $N(S)$ is *stable* if S is a NE
- To evaluate the overall quality of a network, we consider the *social cost*, i.e. the sum of all players' costs

$$cost(S)=\Sigma_i \, cost_i(S)$$

- a network is *optimal* or *socially optimal* if it minimizes the social cost

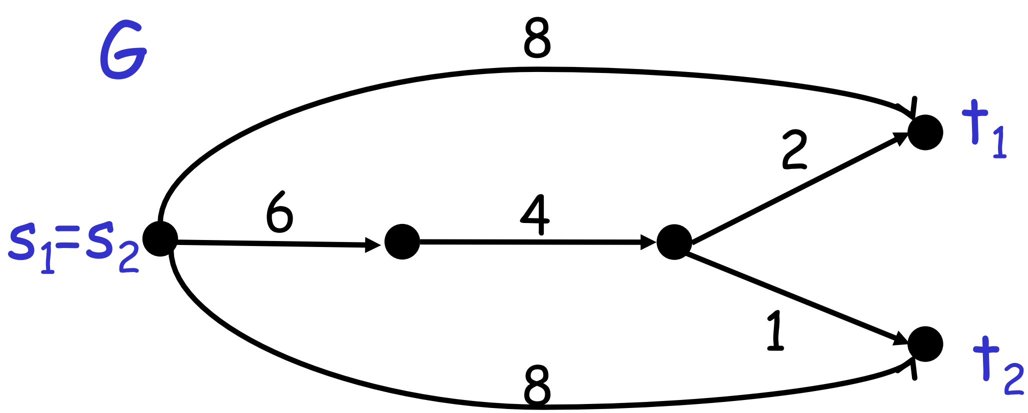

Remind

- We use Nash equilibrium (NE) as the solution concept
- A strategy vector S is a NE if no player has convenience to change its strategy
- Given a strategy vector S, N(S) is *stable* if S is a NE
- To evaluate the overall quality of a network, we consider the *social cost*, i.e. the sum of all players' costs

$$\text{cost}(S)=\Sigma_i\ \text{cost}_i(S)$$

- a network is *optimal* or *socially optimal* if it minimizes the social cost

Notice: $\text{cost}(S)=\Sigma_{e\in N(S)} c_e$

the optimal network is a cheapest subgraph of *G* containg a path from s_i to t_i, for each i

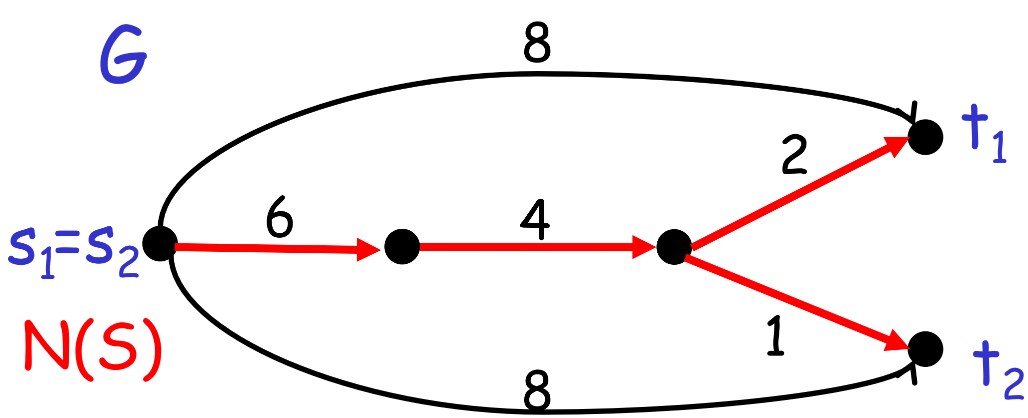

$\text{cost}_1=7$
$\text{cost}_2=6$

an example

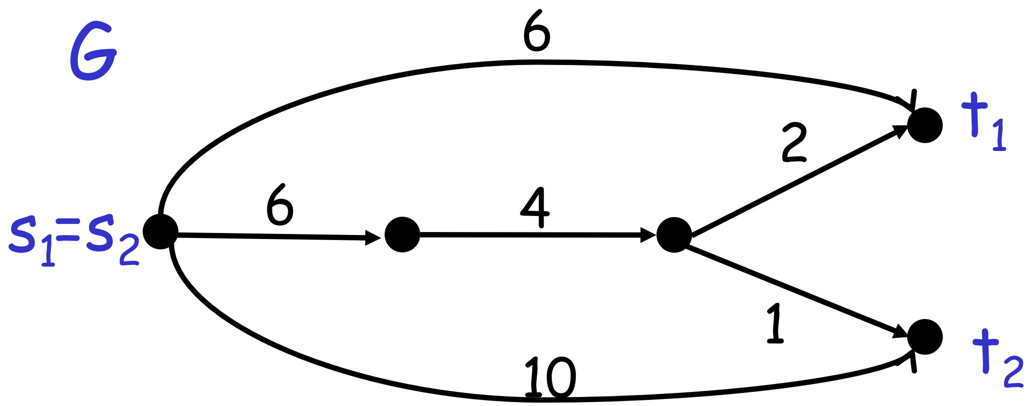

what is the socially optimal network?

an example

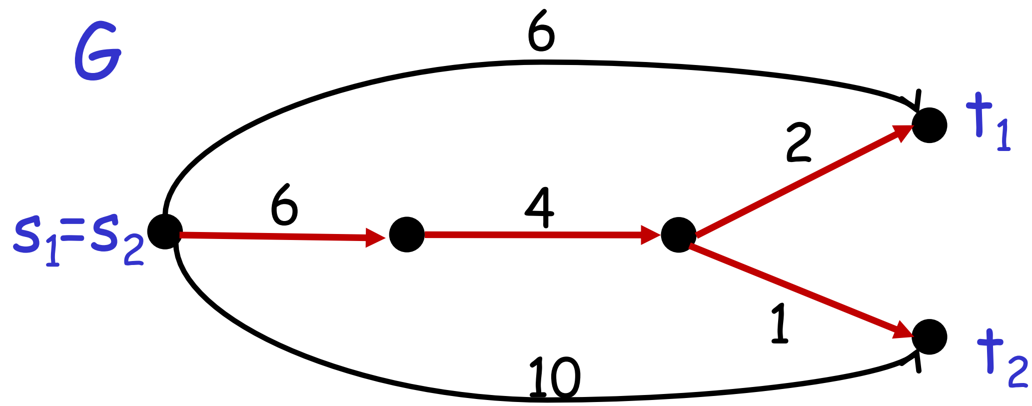

$cost_1$=7
$cost_2$=6

social cost
of the network
13

what is the socially optimal network?

cost of the social optimum: 13

is it stable?

...no!

an example

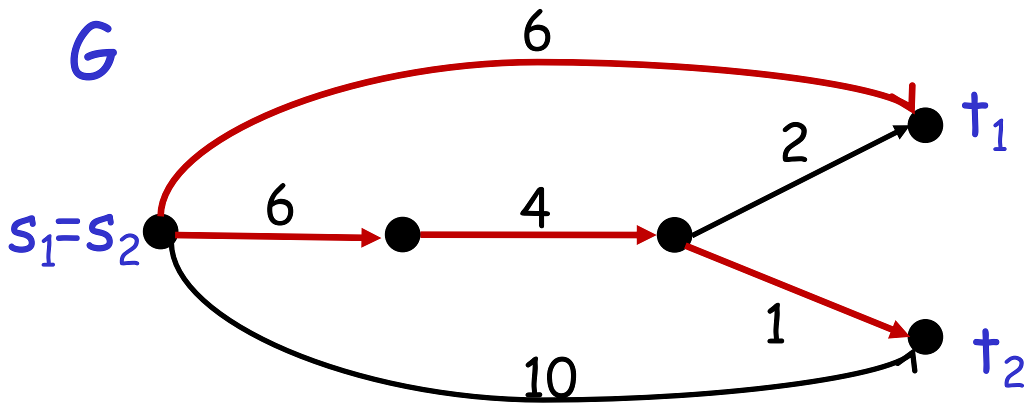

$cost_1$= 6
$cost_2$= 11

social cost
of the network
17

what is the socially optimal network?

cost of the social optimum: 13

is it stable?

...no!

an example

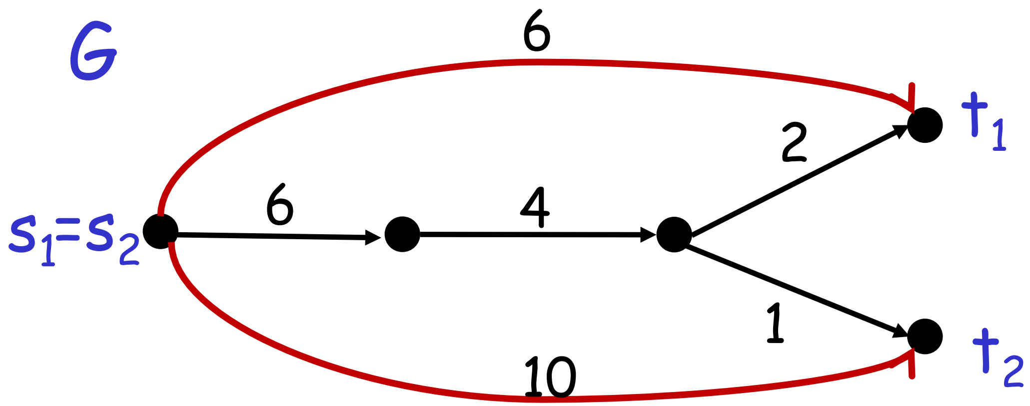

$cost_1= 6$
$cost_2= 10$

social cost
of the network
16

what is the socially
optimal network?

cost of the social
optimum: 13

is it stable?

...yes!

one more example

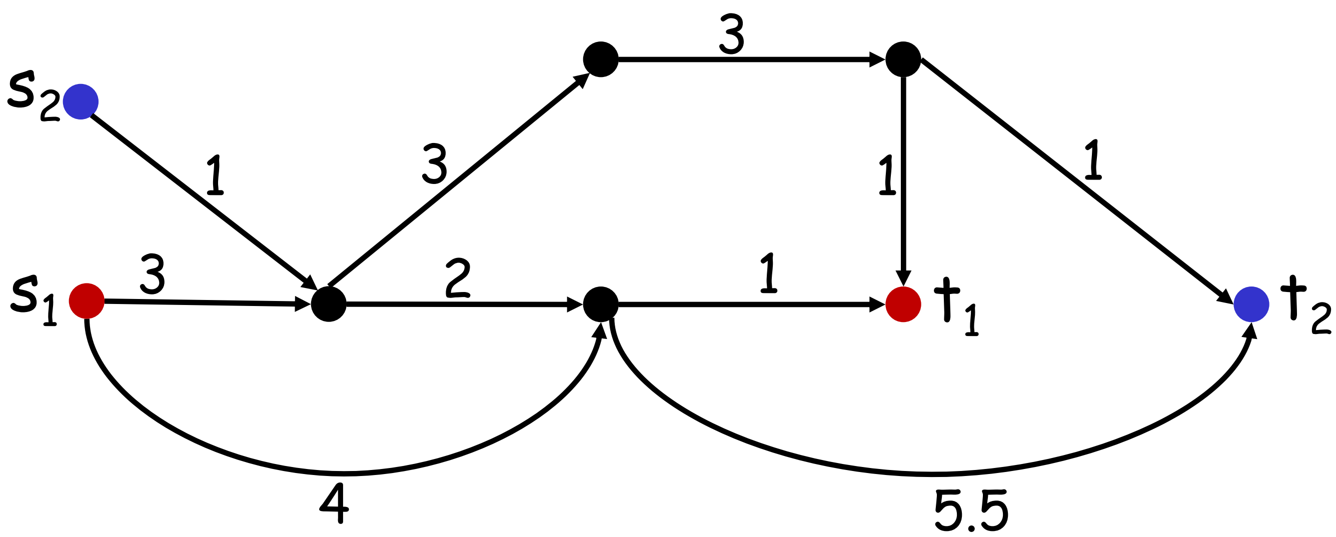

one more example

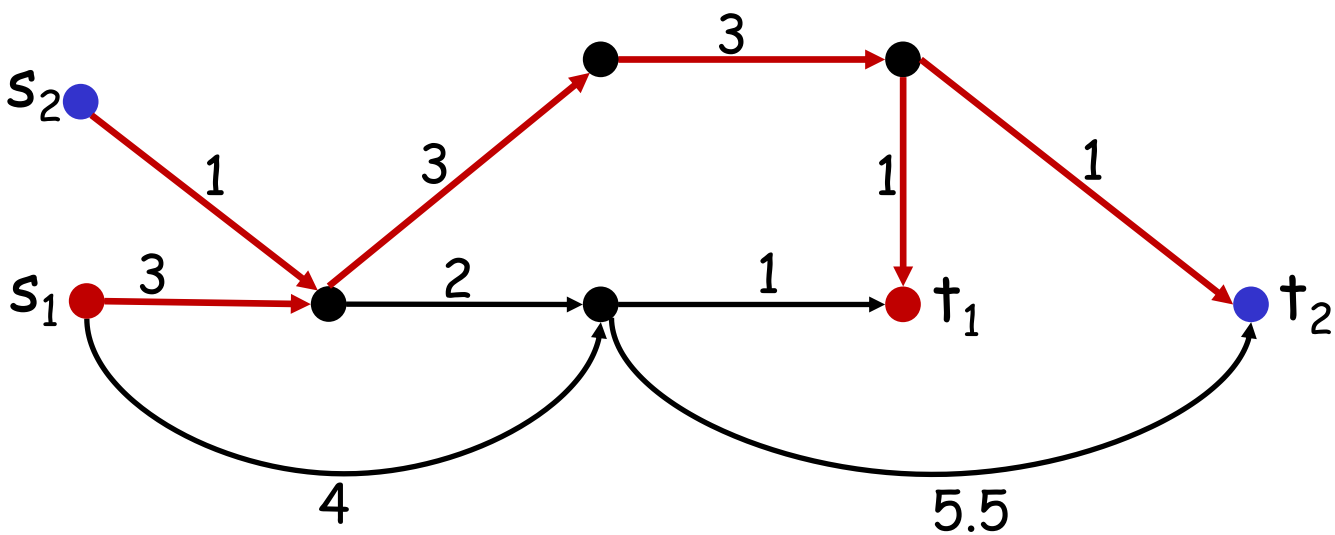

optimal network has cost 12

$cost_1$=7
$cost_2$=5

is it stable?

one more example

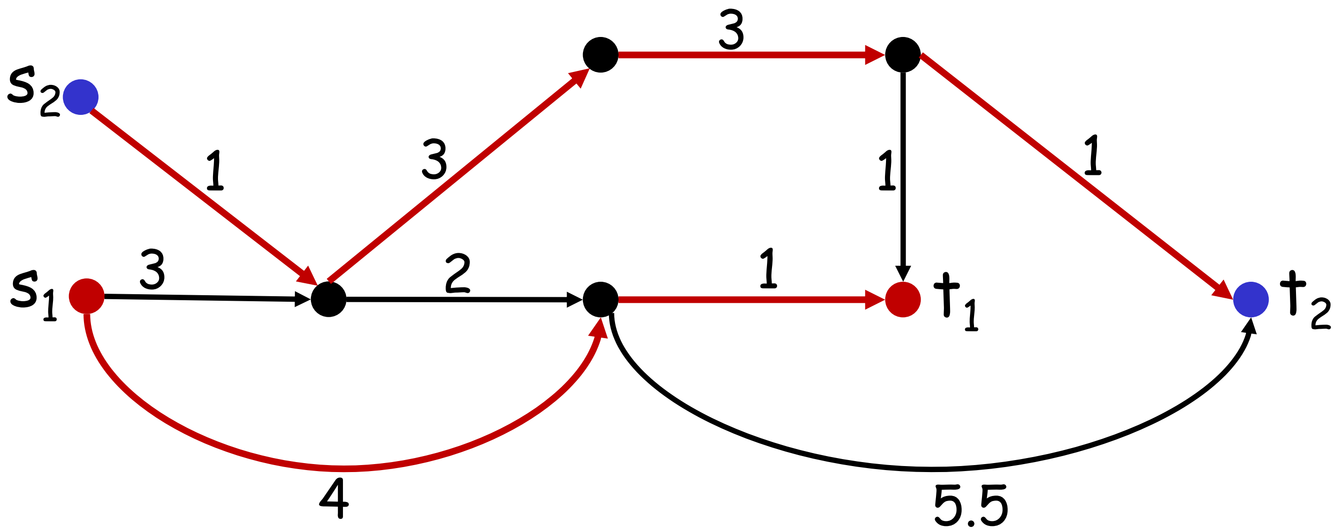

...no!, player 1 can decrease its cost

$cost_1=5$
$cost_2=8$

is it stable? ...yes!

the social cost is 13

one more example

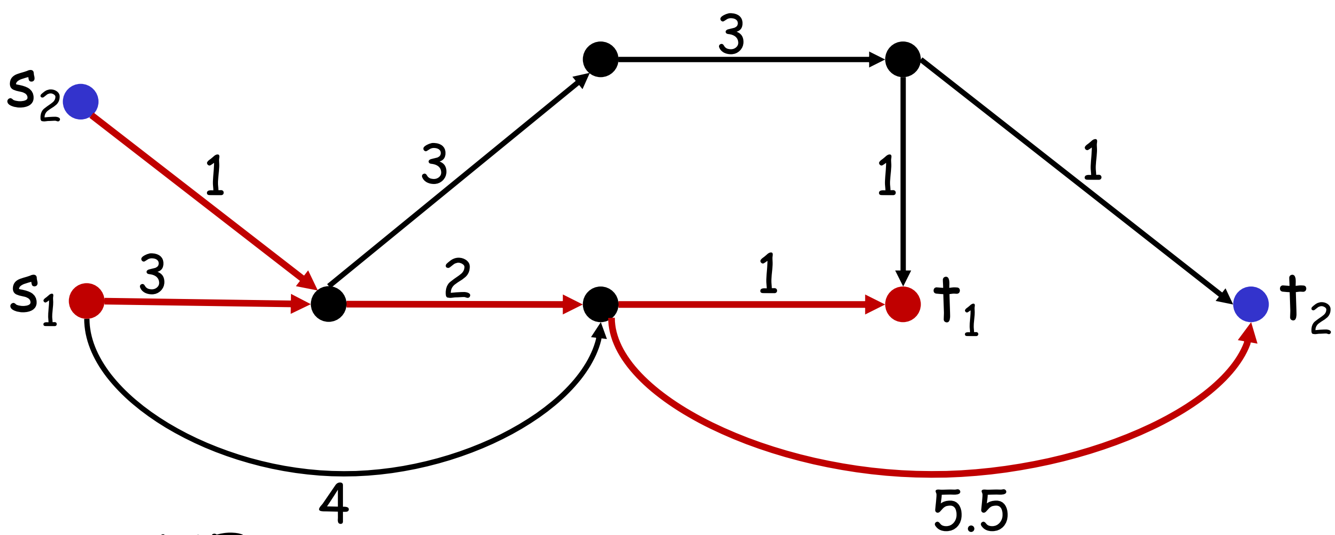

...a better NE...

$cost_1=5$
$cost_2=7.5$

the social cost is 12.5

Addressed issues

- Does a stable network always exist?
- Can we bound the price of anarchy (PoA)?
- Can we bound the price of stability (PoS)?
- Does the repeated version of the game always converge to a stable network?

PoA and PoS

S^*_G : socially optimum for G

for a given network G, we define:

PoA of the game in G $= \max\limits_{\substack{S \text{ s.t.} \\ S \text{ is a NE}}} \dfrac{\text{cost}(S)}{\text{cost}(S^*_G)}$

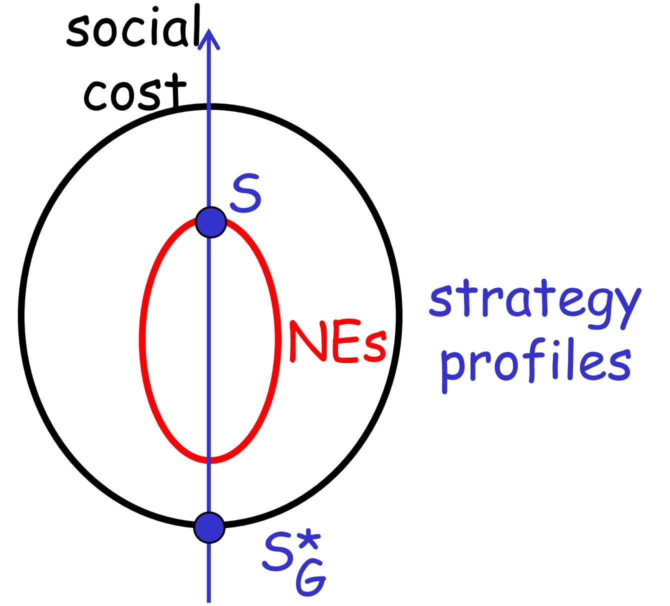

PoS of the game in G $= \min\limits_{\substack{S \text{ s.t.} \\ S \text{ is a NE}}} \dfrac{\text{cost}(S)}{\text{cost}(S^*_G)}$

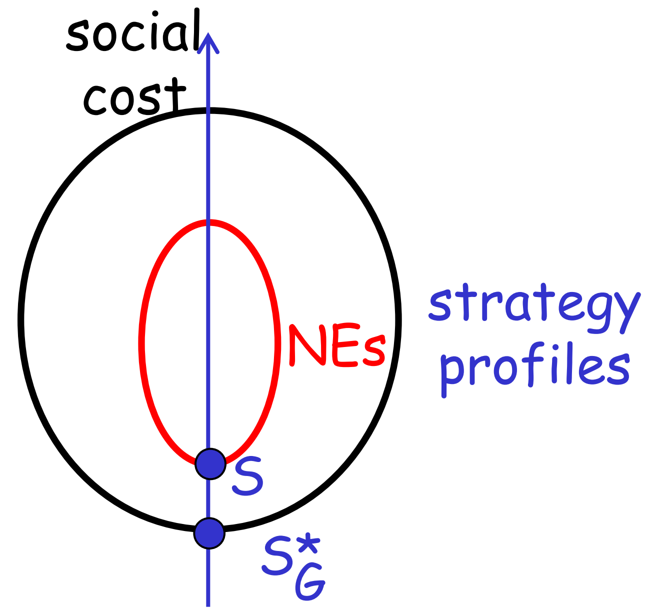

PoA and PoS

we want to bound PoA and PoS in the worst case:

$$\text{PoA of the game} = \max_{G} \text{ PoA in } G$$

$$\text{PoS of the game} = \max_{G} \text{ PoS in } G$$

some notations

we use:

$x=(x_1,x_2,\dots,x_k)$; $x_{-i}=(x_1,\dots,x_{i-1},x_{i+1},\dots,x_k)$; $x=(x_{-i},x_i)$

G: a weighted directed network

cost or length of a path π in G from a node u to a node v : $\sum_{e\in\pi} c_e$

$d_G(u,v)$: distance in G from a node u to a node v : length of any shortest path in G from u to v

Price of Anarchy

Price of Anarchy: a lower bound

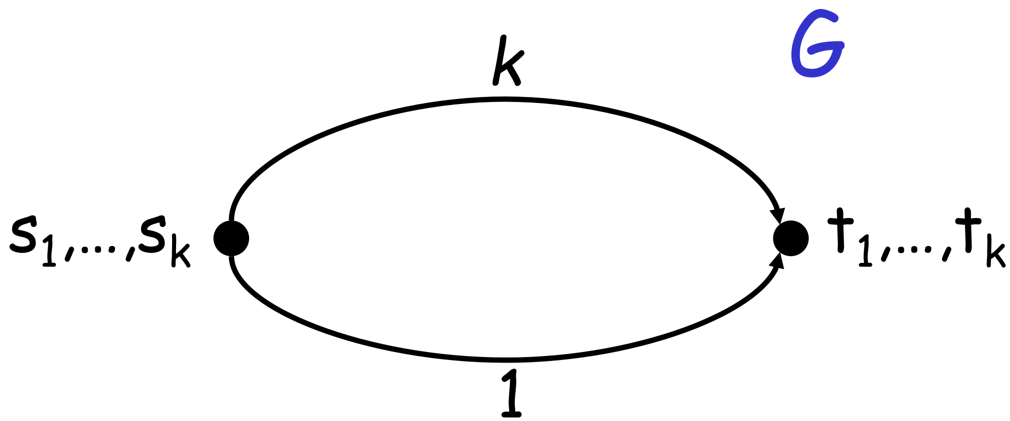

optimal network has cost 1

best NE: all players use the lower edge ⇒ PoS in G is 1 ☺

worst NE: all players use the upper edge ⇒ PoA in G is k ☹

⇒ $\boxed{\text{PoA of the game is } \geq k}$

Theorem

> The price of anarchy in the global connection game with *k* players is at most *k*

proof

S: a NE S*: a strategy profile minimizing the social cost

for each player i,

let π_i be a shortest path in G from s_i to t_i

we have

$$cost_i(S) \leq cost_i(S_{-i}, \pi_i) \leq d_G(s_i,t_i) \leq cost(S^*)$$

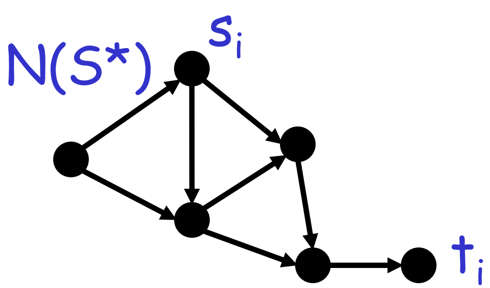

Theorem

The price of anarchy in the global connection game with k players is at most k

proof

S: a NE S^*: a strategy profile minimizing the social cost

for each player i,

let π_i be a shortest path in G from s_i to t_i

we have

$$cost_i(S) \leq cost_i(S_{-i}, \pi_i) \leq d_G(s_i,t_i) \leq cost(S^*)$$

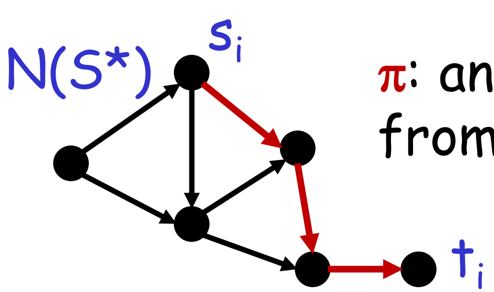

π: any path in $N(S^*)$ from s_i to t_i

$$d_G(s_i,t_i) \leq \text{cost of } \pi \leq cost(S^*)$$

$$cost(S)=\Sigma_i\, cost_i(S) \leq k\, cost(S^*)$$

Price of Stability & potential function method

Price of Stability: a lower bound

$\varepsilon > 0$: small value

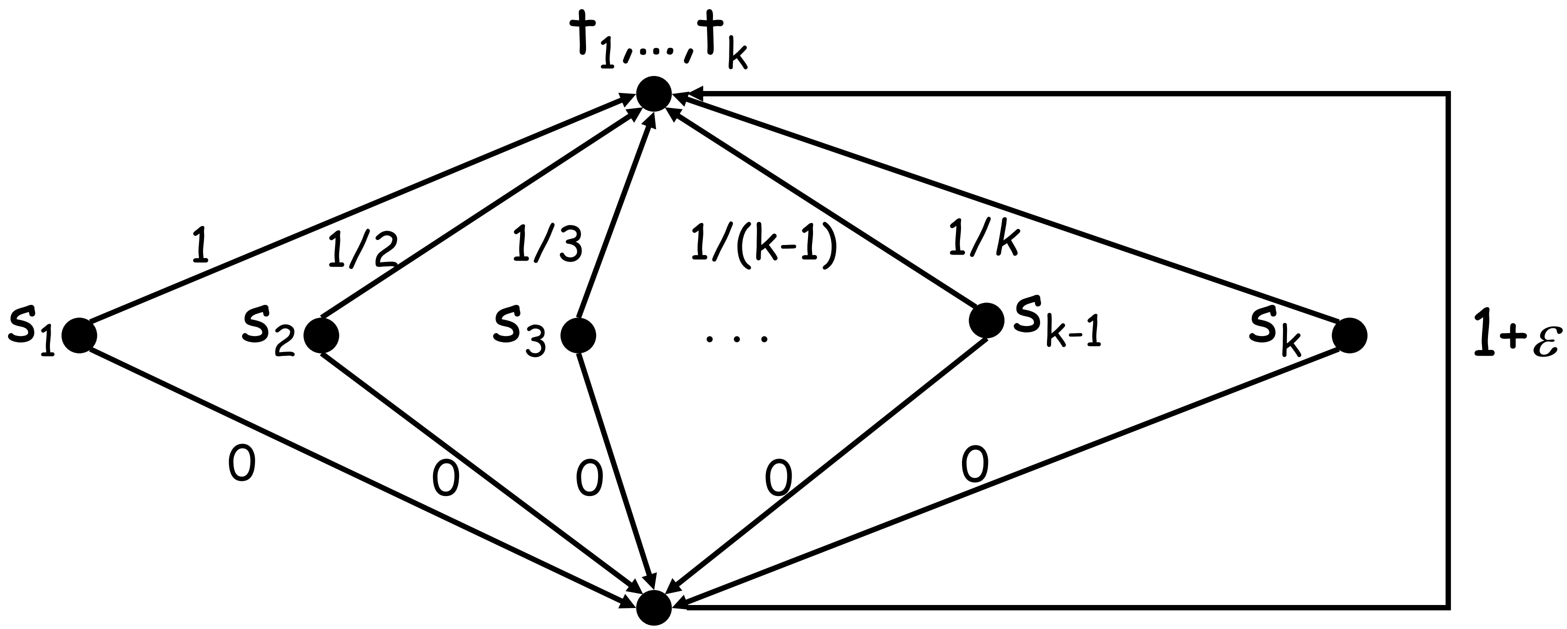

Price of Stability: a lower bound

ε>0: small value

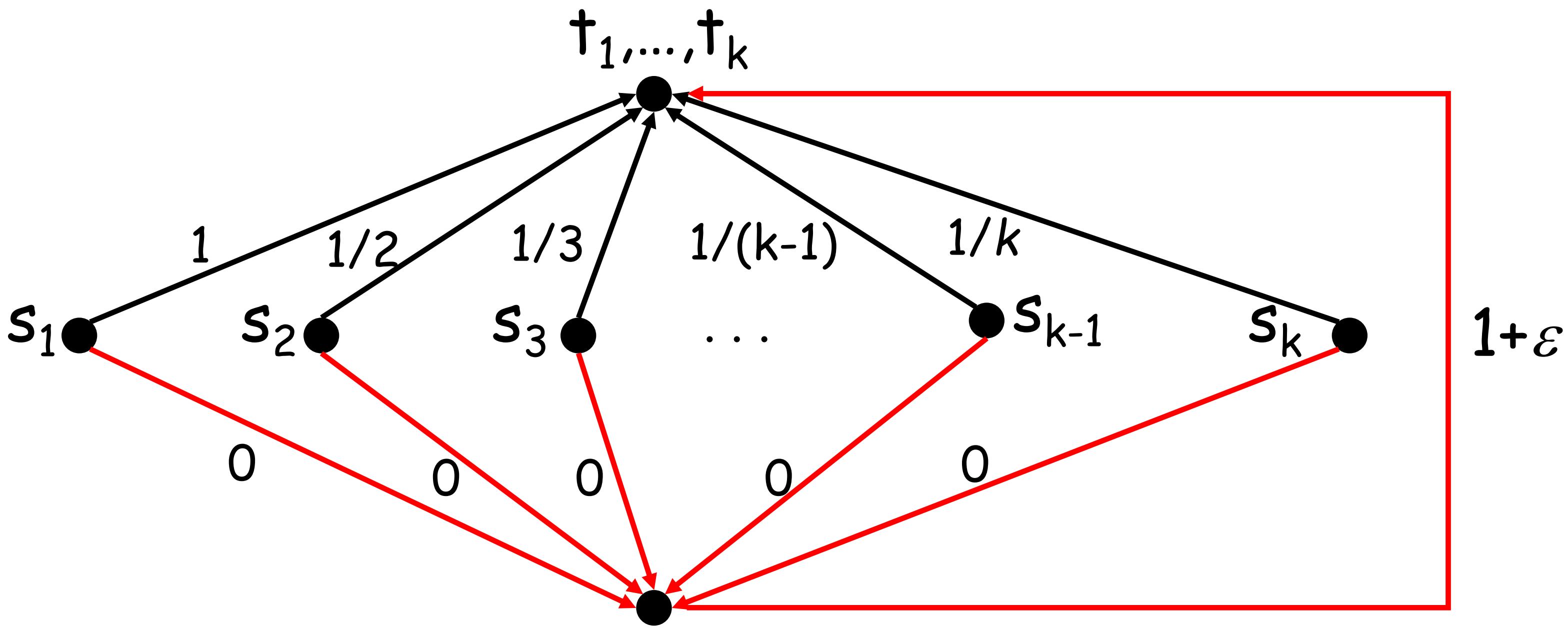

The optimal solution has a cost of $1+\varepsilon$

is it stable?

Price of Stability: a lower bound

ε>0: small value

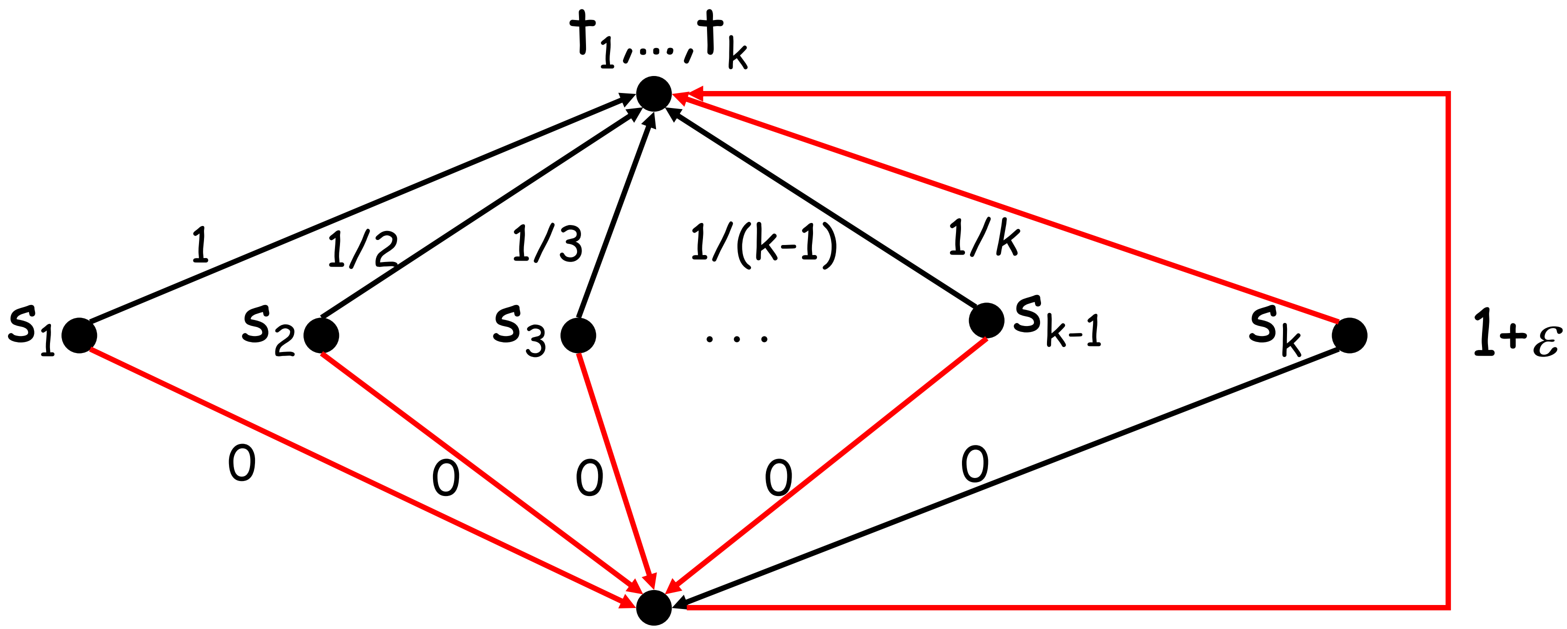

…no! player k can decrease its cost…

is it stable?

Price of Stability: a lower bound

ε>0: small value

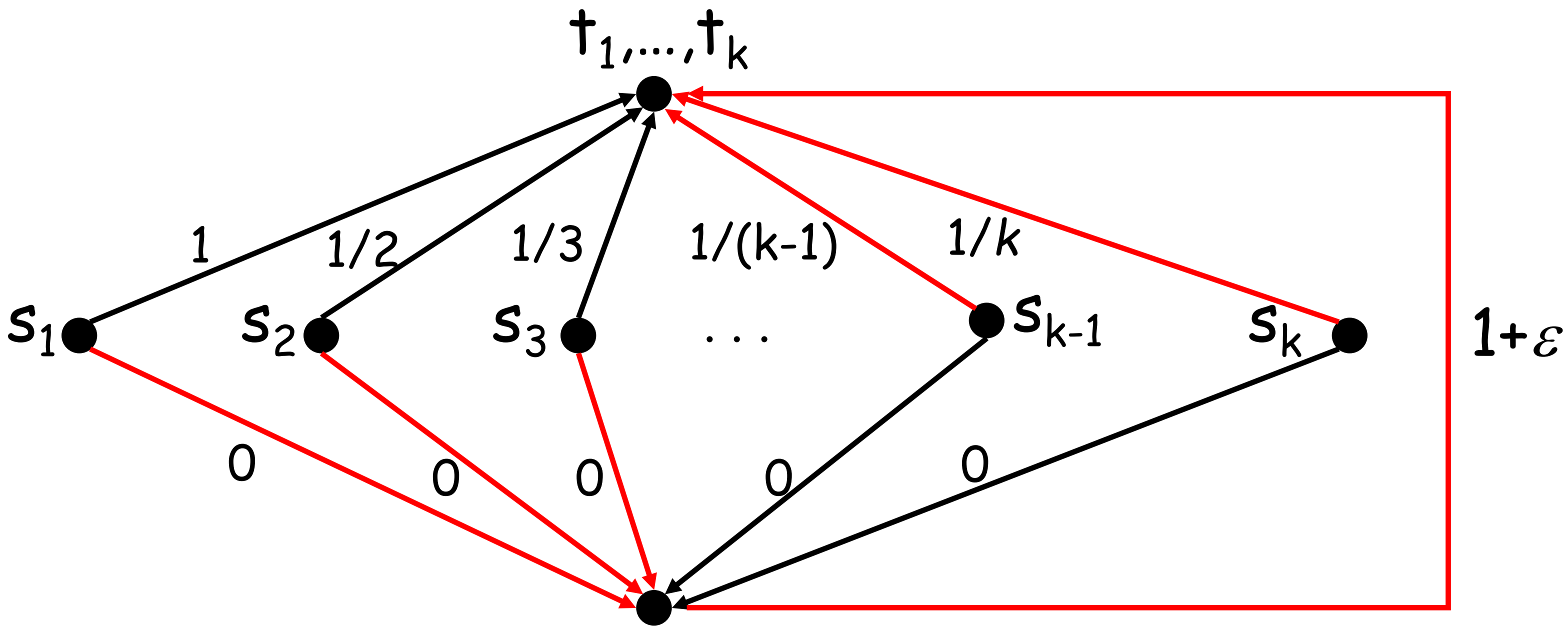

…no! player k-1 can decrease its cost…

is it stable?

Price of Stability: a lower bound

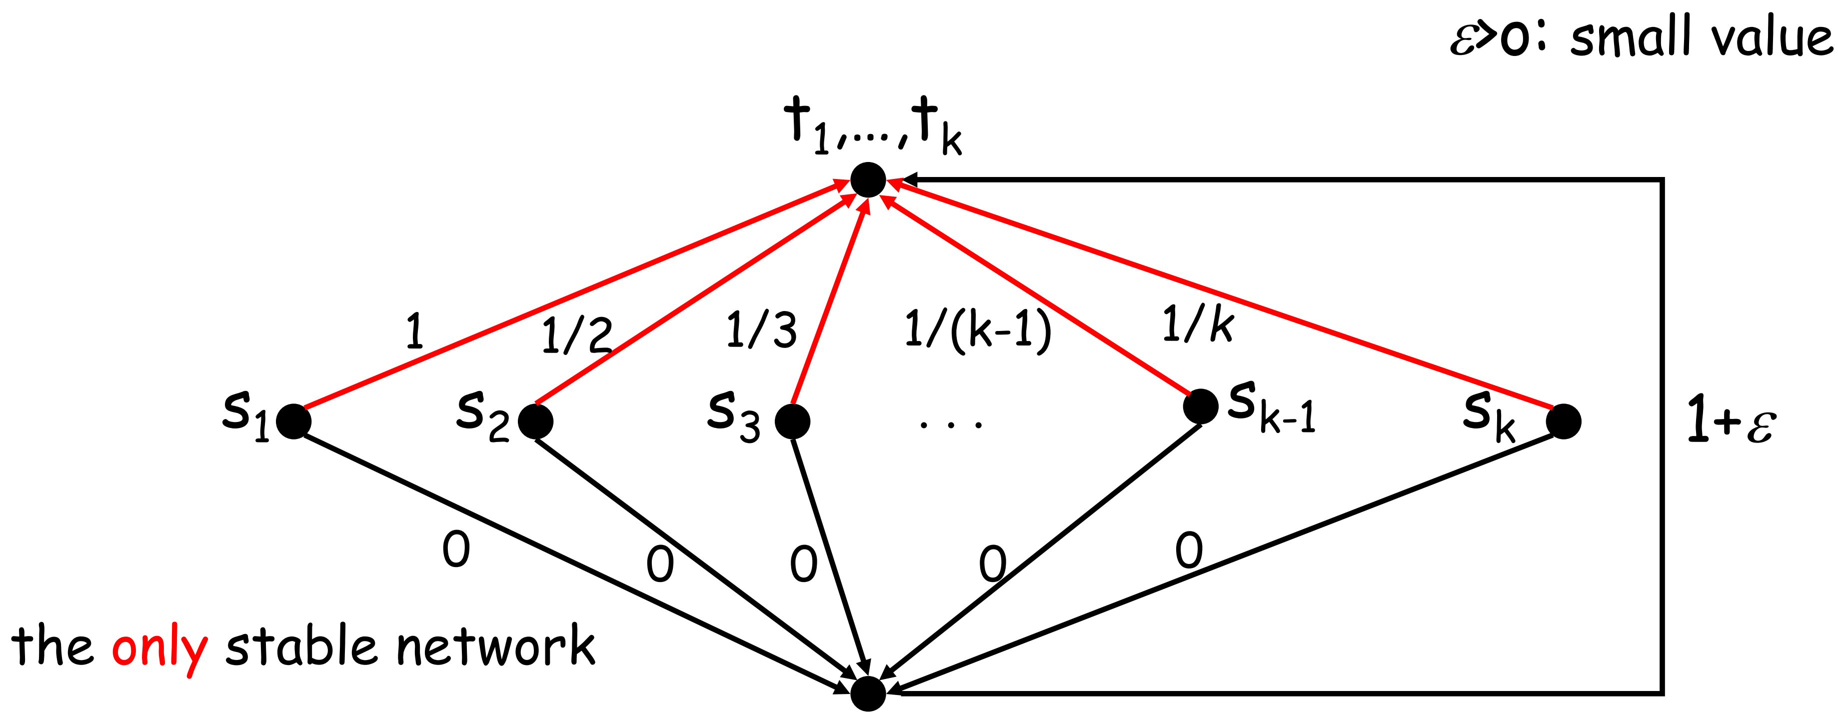

social cost: $\sum_{j=1}^{k} 1/j = H_k \le \ln k + 1$ k-th *harmonic number*

the optimal solution has a cost of $1+\varepsilon$ ➡ PoS of the game is $\geq H_k$

Theorem

Any instance of the global connection game has a pure Nash equilibrium, and better response dynamic always converges

Theorem

The price of stability in the global connection game with k players is at most H_k, the k-th harmonic number

To prove them we use the *Potential function method*

Notation:

$x=(x_1,x_2,\ldots,x_k)$; $x_{-i}=(x_1,\ldots,x_{i-1},x_{i+1},\ldots,x_k)$; $x=(x_{-i},x_i)$

Definition

For any finite game, an *exact potential function* Φ is a function that maps every strategy vector S to some real value and satisfies the following condition:

$\forall S=(S_1,\ldots,S_k)$, $S'_i \neq S_i$, let $S'=(S_{-i},S'_i)$, then

$$\Phi(S)-\Phi(S') = cost_i(S)-cost_i(S')$$

A game that posses an exact potential function is called *potential game*

Theorem

Every potential game has at least one pure Nash equilibrium, namely the strategy vector S that minimizes $\Phi(S)$

proof

consider any move by a player i that results in a new strategy vector S'

we have:

$$\underbrace{\Phi(S)-\Phi(S')}_{\leq 0} = cost_i(S)-cost_i(S')$$

$\Rightarrow$ $cost_i(S) \leq cost_i(S')$ $\Rightarrow$ player i cannot decrease its cost, thus S is a NE

Theorem

In any finite potential game, better response dynamics always converge to a Nash equilibrium

proof

better response dynamics simulate local search on Φ:

1. each move strictly decreases Φ
2. finite number of solutions $\blacksquare$

Note: in our game, a best response can be computed in polynomial time

Theorem

Suppose that we have a potential game with potential function Φ, and assume that for any outcome S we have

$$\mathrm{cost}(S)/A \le \Phi(S) \le B\ \mathrm{cost}(S)$$

for some A,B>0. Then the price of stability is at most AB

proof

Let S' be the strategy vector minimizing Φ

Let S* be the strategy vector minimizing the social cost

we have:

$$\mathrm{cost}(S')/A \le \Phi(S') \le \Phi(S^*) \le B\ \mathrm{cost}(S^*)$$

…turning our attention to the global connection game…

Let Φ be the following function mapping any strategy vector S to a real value:

$$\Phi(S) = \Sigma_{e\in E}\, \Phi_e(S)$$

where

$$\Phi_e(S) = c_e\, H_{k_e(S)}$$

$H_k = \sum_{j=1}^{k} 1/j$ k-th *harmonic number*

[we define $H_0 = 0$]

Lemma 1

Let $S=(P_1,\ldots,P_k)$, let P'_i be an alternative path for some player i, and define a new strategy vector $S'=(S_{-i},P'_i)$. Then:

$$\Phi(S) - \Phi(S') = cost_i(S) - cost_i(S')$$

Lemma 2

For any strategy vector S, we have:

$$cost(S) \leq \Phi(S) \leq H_k\, cost(S)$$

…from which we have:

PoS of the game is $\leq H_k$

Lemma 2

For any strategy vector S, we have:

$$\mathrm{cost}(S) \le \Phi(S) \le H_k\,\mathrm{cost}(S)$$

proof

$$\mathrm{cost}(S) \le \Phi(S) = \sum_{e \in E} c_e H_{k_e(S)}$$

$$= \sum_{e \in N(S)} c_e H_{k_e(S)} \le \sum_{e \in N(S)} c_e H_k = H_k \mathrm{cost}(S)$$

$$1 \le k_e(S) \le k \quad \text{for } e \in N(S)$$

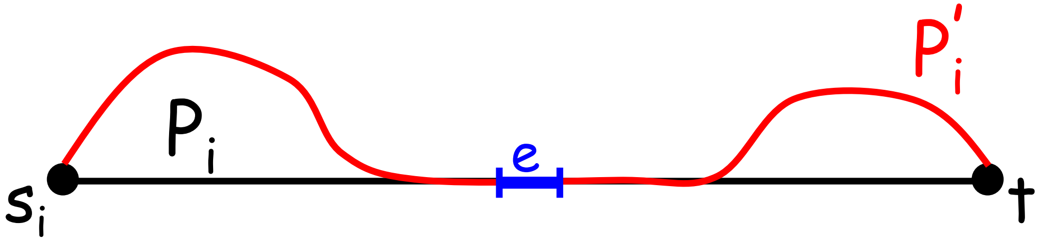

for each $e \in P_i \cap P_i'$

term e of $cost_i()$ & potential Φ_e remain the same

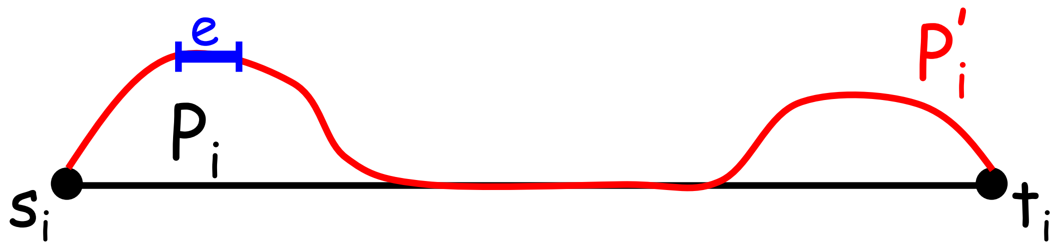

for each $e \in P_i' \setminus P_i$

term e of $\text{cost}_i()$ increases by $c_e/(k_e(S)+1)$

potential Φ_e increases from $c_e\left(1 + \frac{1}{2} + \ \ldots + \frac{1}{k_e(S)}\right)$

to $c_e\left(1 + \frac{1}{2} + \ \ldots + \frac{1}{k_e(S)} + \frac{1}{k_e(S)+1}\right)$

$\Rightarrow \quad \Delta\Phi_e = c_e/(k_e(S)+1)$

(proof of Lemma 1)

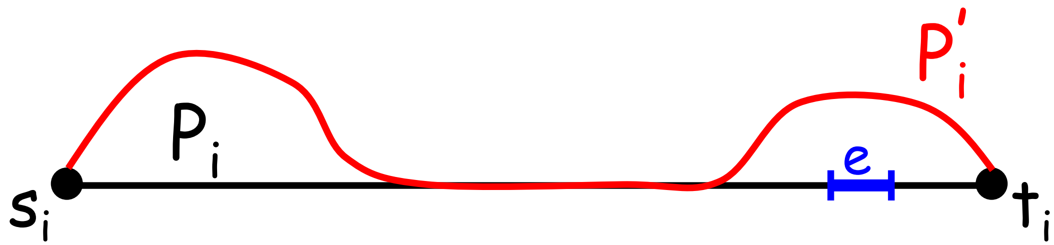

for each $e \in P_i \setminus P_i'$

term e of $\text{cost}_i()$ decreases by $c_e / k_e(S)$

potential Φ_e decreases from $c_e \left(1 + \frac{1}{2} + \ . \ . \ . + \frac{1}{k_e(S)-1} + \frac{1}{k_e(S)}\right)$

to $c_e \left(1 + \frac{1}{2} + \ . \ . \ . + \frac{1}{k_e(S)-1}\right)$

$\Rightarrow$ $\Delta\Phi_e = - c_e / k_e(S)$

Theorem

Given an instance of a GC Game and a value C, it is NP-complete to determine if a game has a Nash equilibrium of cost at most C.

proof

Reduction from 3-dimensional matching problem

3-dimensional matching problem

- Input:
  - disjoint sets X, Y, Z, each of size n
  - a set $T \subseteq X \times Y \times Z$ of ordered triples
- Question:
  - does there exist a set of n triples in T so that each element of $X \cup Y \cup Z$ is contained in exactly one of these triples?

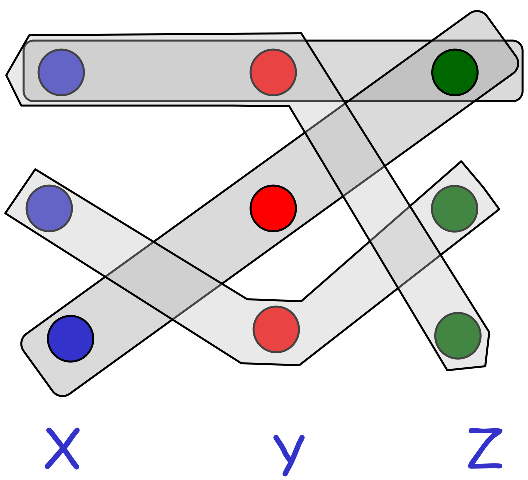

3-dimensional matching problem

- Input:
  - disjoint sets X, Y, Z, each of size n
  - a set $T \subseteq X \times Y \times Z$ of ordered triples
- Question:
  - does there exist a set of n triples in T so that each element of $X \cup Y \cup Z$ is contained in exactly one of these triples?

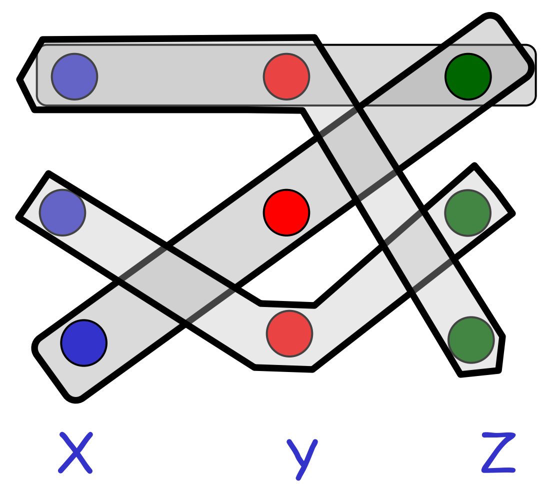

the reduction

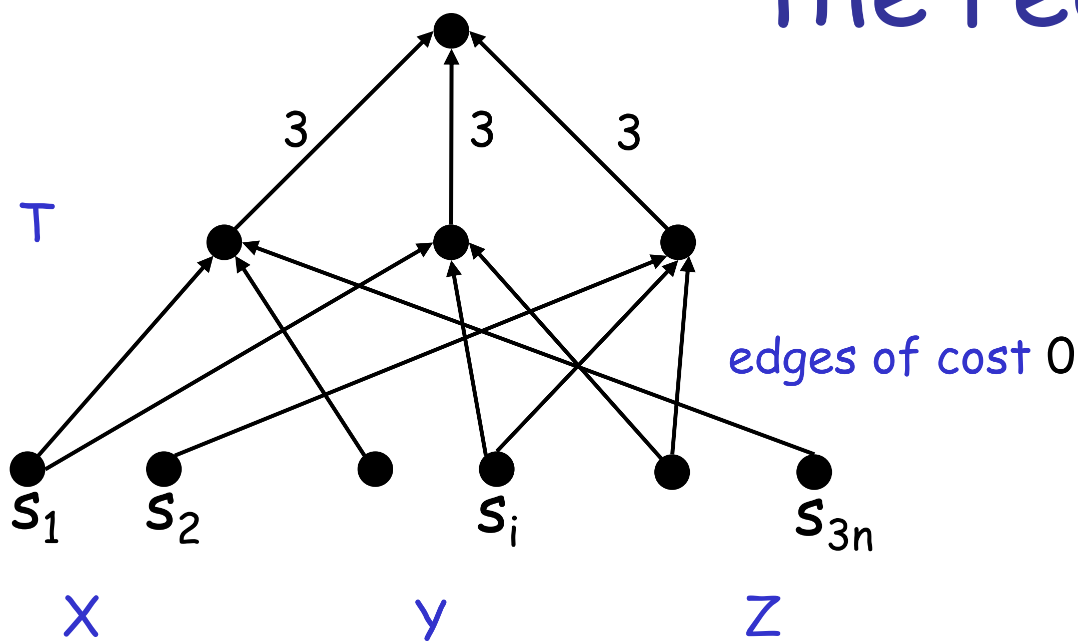

There is a 3D matching if and only if there is a NE of cost at most C=3n

the reduction

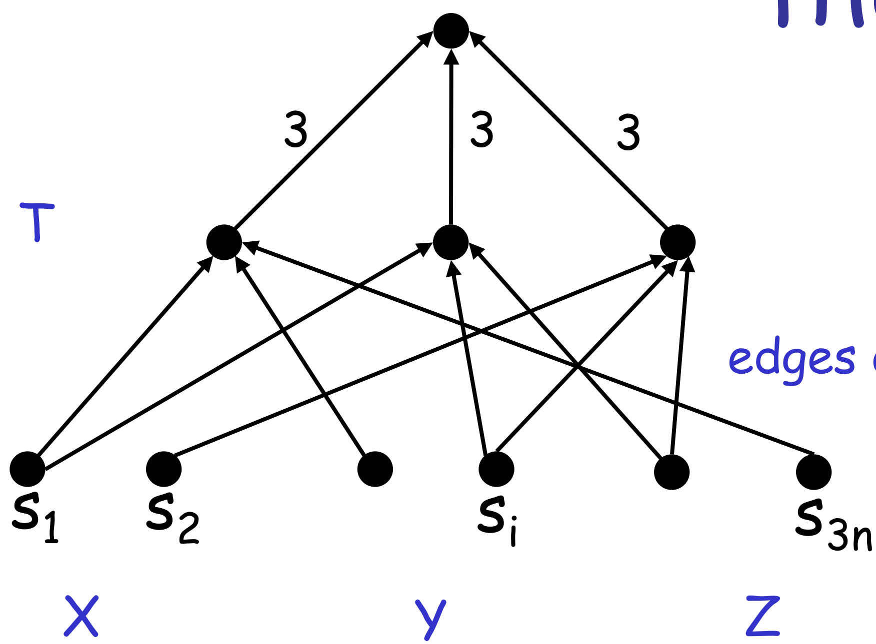

Assume there is a 3D matching.

S: strategy profile in which each player choose a path passing through the triple of the matching it belongs to

the reduction

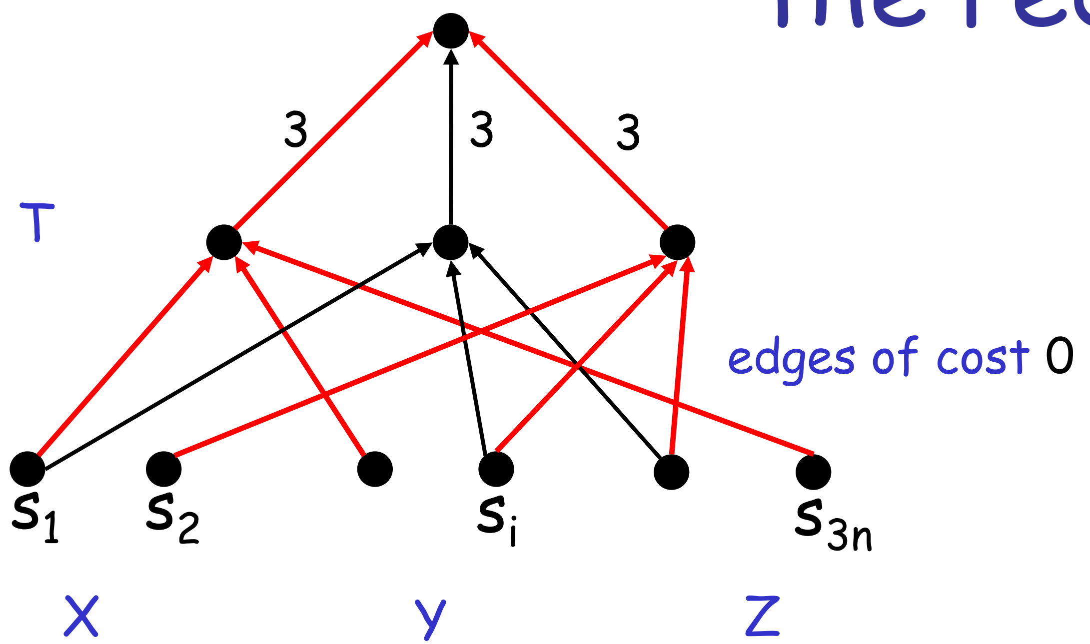

Assume there is a 3D matching.

S: strategy profile in which each player choose a path passing through the triple of the matching it belongs to

cost(S)= 3n

S is a NE

the reduction

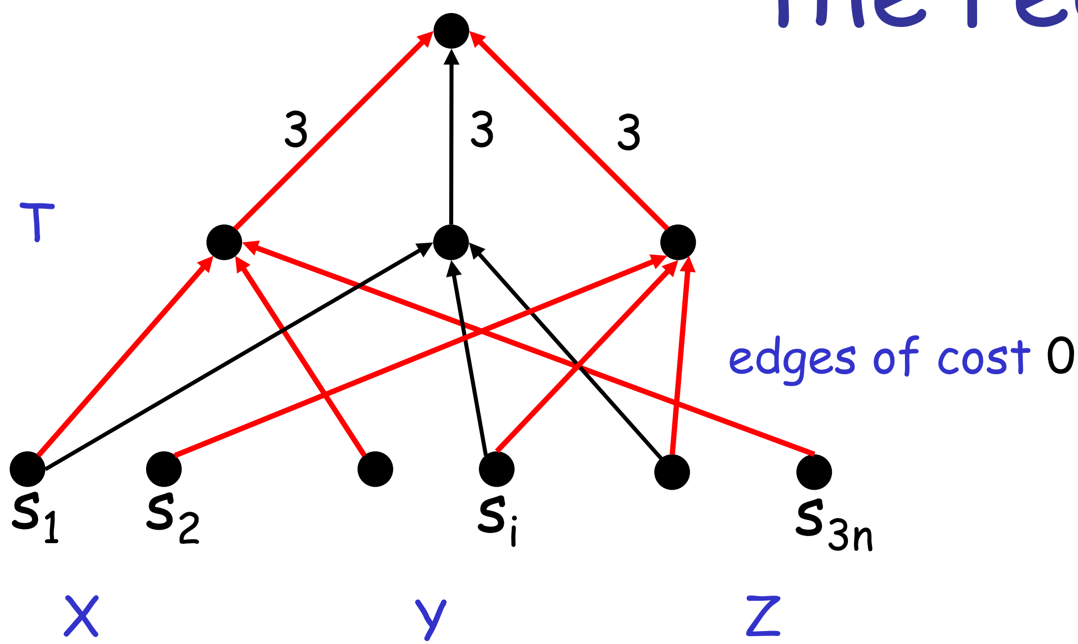

Assume there is a NE of cost $\leq 3n$

N(S) uses at most n edges of cost 3

each edge of cost 3 can "serve" at most 3 players

then, the edge of cost 3 are exactly n

…and they define a set of triples that must be a 3D-matching

What is the PoS of the game for *undirected* networks?

PoS for undirected graphs: State of the art

UB LB

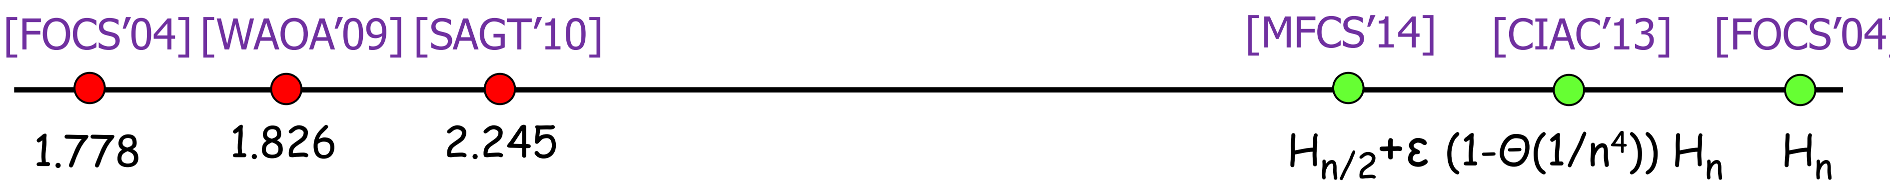

one single terminal (multicast)

+ all sources (broadcast)

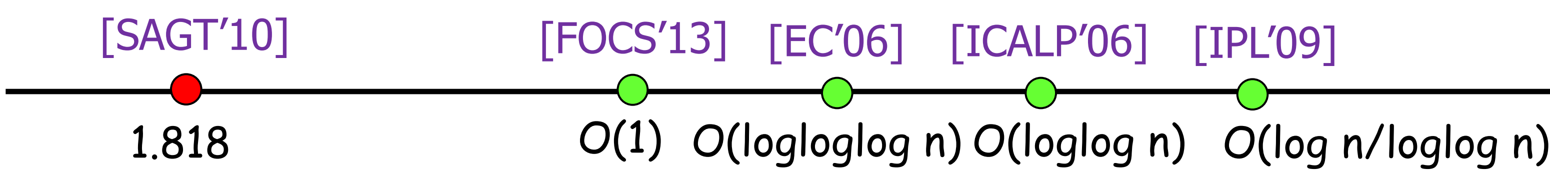

Max-cut game

- $G=(V,E)$: undirected graph
- Nodes are (selfish) players
- Strategy S_u of u is a color {red, green}
- player u's payoff in S (to maximize):
  - $p_u(S)=|\{(u,v)\in E : S_u \neq S_v\}|$

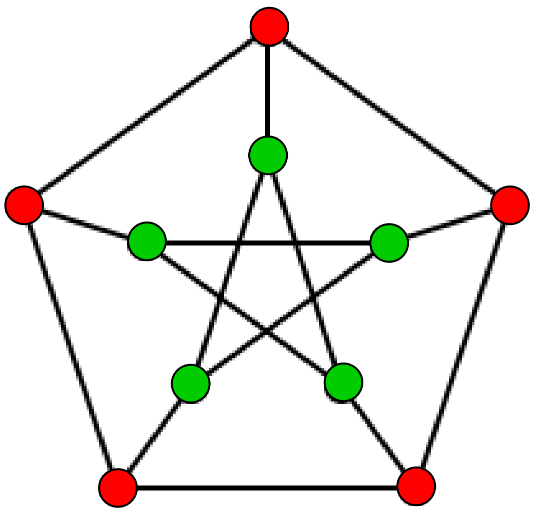

social welfare of strategy vector S
$\Sigma_u\, p_u(S)=$
2 #edges crossing the red-green cut

Max-cut game

does a Nash Equilibrium always exist?

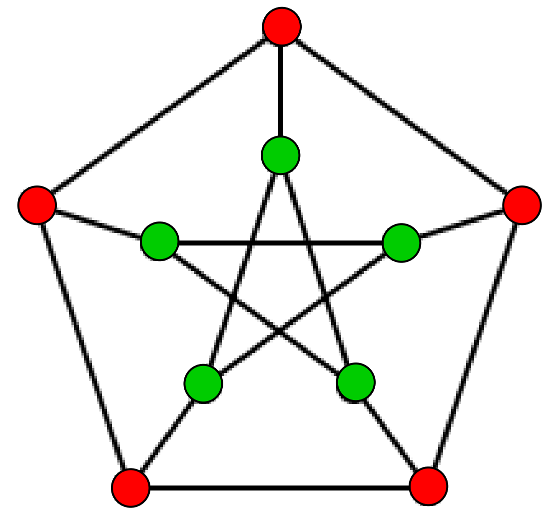

how bad a Nash Equilibrium Can be?

does the repeated game always converge to a Nash Equilibrium?

...let's play Max-cut game on Petersen Graph

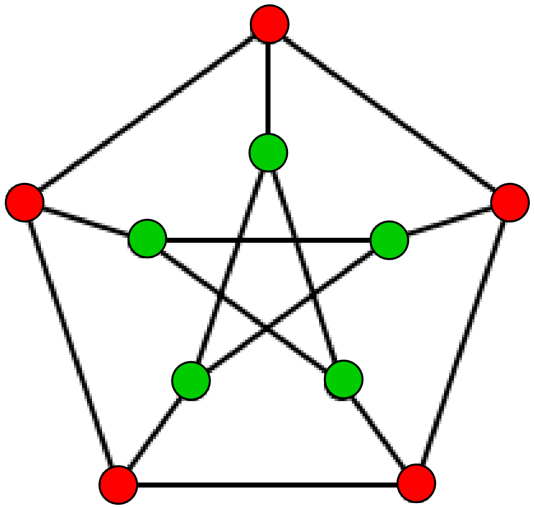

...is it a NE?

...let's play Max-cut game on Petersen Graph

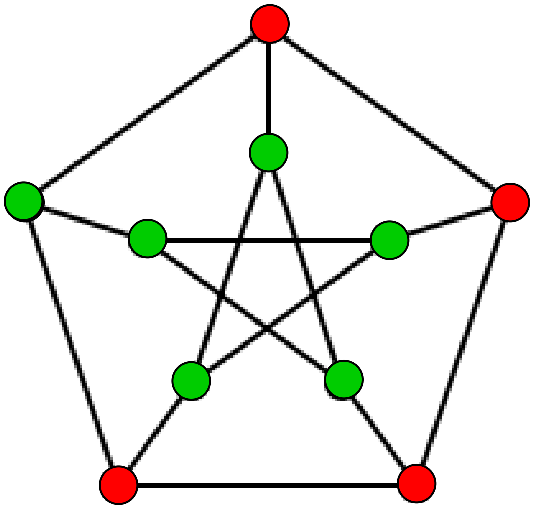

...is it a NE?

…let's play Max-cut game
on Petersen Graph

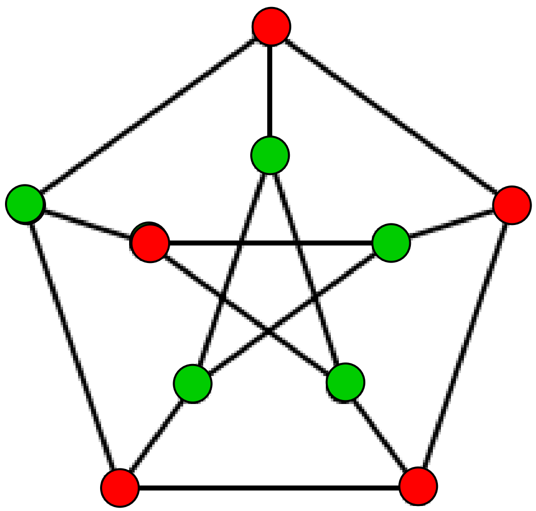

…is it a NE?

…let's play Max-cut game on Petersen Graph

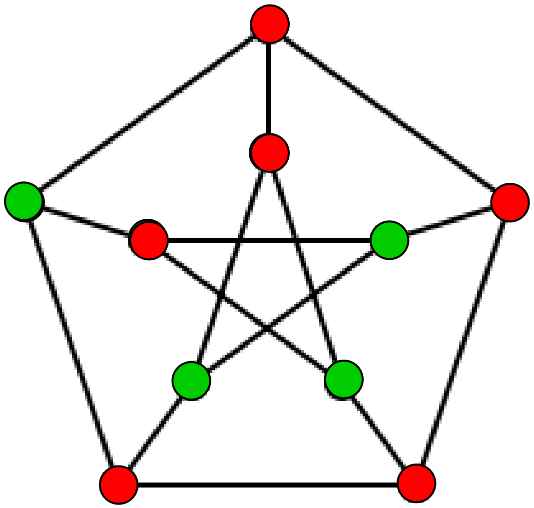

…is it a NE?

...let's play Max-cut game
on Petersen Graph

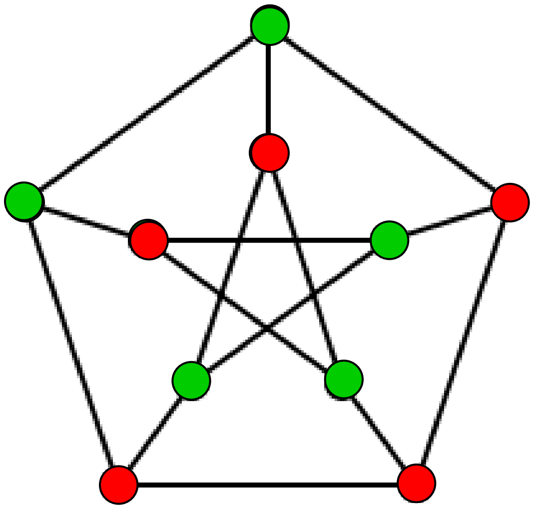

...is it a NE?

...let's play Max-cut game on Petersen Graph

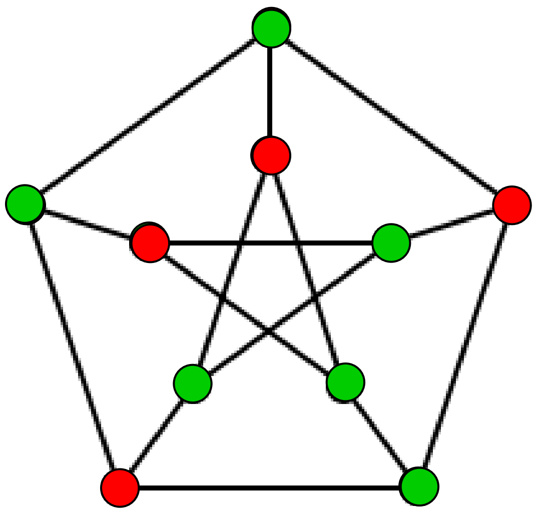

...is it a NE?

...yes!

# of edges crossing the cut is 12

Exercise

Show that:

(i) Max-cut game is a potential game

(ii) PoS is 1

(iii) PoA $\geq \frac{1}{2}$

(iv) there is an instance of the game having a NE with social welfare of $\frac{1}{2}$ the social optimum