

# *Game Theory*

## *Part 2, Chapter 4*



*Roger Wattenhofer*

# Overview

- Selfish Caching
- Nash Equilibrium
- Price of Anarchy
- Rock Paper Scissor
- Mechanism Design

# Selfish Peers

- Peers may not try to destroy the system, instead they may try to benefit from the system without contributing anything
- Such **selfish behavior** is called **free riding** or **freeloading**
- Free riding is a common problem in file sharing applications:
- Studies show that most users in the P2P file sharing networks do not want to provide anything
- Protocols that are supposed to be “incentive-compatible”, such as **BitTorrent**, can also be exploited
  - The **BitThief** client downloads without uploading!



# Game Theory


- Game theory attempts to mathematically capture behavior in strategic situations (games), in which an individual's success in making choices depends on the choices of others.
- “Game theory is a sort of umbrella or 'unified field' theory for the rational side of social science, where 'social' is interpreted broadly, to include human as well as non-human players (computers, animals, plants)”  
[Aumann 1987]



game theory

# Selfish Caching

- P2P system where node  $i$  experiences a demand  $w_i$  for a certain file.
  - Setting can be extended to multiple files
- A node can either
  - cache the file for cost  $\alpha$ , or
  - get the file from the nearest node  $l(i)$  that caches it for cost  $w_i \cdot d_{i,l(i)}$
- Example:  $\alpha = 4$ ,  $w_i = 1$



What is the global „best“ configuration?  
Who will cache the object?  
Which configurations are „stable“?

# Social Optimum & Nash Equilibrium

- In game theory, the „best“ configurations are called social optima
  - A social optimum maximizes the social welfare

## Definition


A strategy profile is called **social optimum** iff it minimizes the sum of all cost.

- A strategy profile is the set of strategies chosen by the players
- „Stable“ configurations are called (Nash) Equilibria

## Definition


A **Nash Equilibrium (NE)** is a strategy profile for which nobody can improve by unilaterally changing its strategy

- Systems are assumed to magically converge towards a NE



## Selfish Caching: Example 2

- Which are the social optima, and the Nash Equilibria in the following example?
  - $\alpha = 4$



- Nash Equilibrium  $\not\leftrightarrow$  Social optimum
- Does every game have
  - a social optimum?
  - a Nash equilibrium?

# Selfish Caching: Equilibria

Theorem

Any instance of the selfish caching game has a Nash equilibrium

- **Proof by construction:**

- The following procedure always finds a Nash equilibrium


1. Put a node  $y$  with highest demand into caching set
2. Remove all nodes  $z$  for which  $d_{zy}w_z < \alpha$
3. Repeat steps 1 and 2 until no nodes left

- The strategy profile where all nodes in the caching set cache the file, and all others chose to access the file remotely, is a Nash equilibrium.




# Selfish Caching: Proof example

1. Put a node  $y$  with highest demand into caching set
2. Remove all nodes  $z$  for which  $d_{zy}w_z < \alpha$
3. Repeat steps 1 and 2 until no nodes left




# Selfish Caching: Proof example

1. Put a node  $y$  with highest demand into caching set
2. Remove all nodes  $z$  for which  $d_{zy}w_z < \alpha$
3. Repeat steps 1 and 2 until no nodes left




# Selfish Caching: Proof example

1. Put a node  $y$  with highest demand into caching set
2. Remove all nodes  $z$  for which  $d_{zy}w_z < \alpha$
3. Repeat steps 1 and 2 until no nodes left



# Selfish Caching: Proof example

1. Put a node  $y$  with highest demand into caching set
2. Remove all nodes  $z$  for which  $d_{zy}w_z < \alpha$
3. Repeat steps 1 and 2 until no nodes left



- Does NE condition hold for every node?


# Proof

- If node  $x$  not in the caching set
  - Exists  $y$  for which  $w_x d_{xy} < \alpha$
  - No incentive to cache because remote access cost  $w_x d_{xy}$  are smaller than placement cost  $\alpha$
- If node  $x$  is in the caching set
  - For any other node  $y$  in the caching set:
    - Case 1:  $y$  was added to the caching set before  $x$ 
      - It holds that  $w_x d_{xy} \geq \alpha$  due to the construction
    - Case 2:  $y$  was added to the caching set after  $x$ 
      - It holds that  $w_x \geq w_y$ , and  $w_y d_{yx} \geq \alpha$  due to the construction
      - Therefore  $w_x d_{xy} \geq w_y d_{yx} \geq \alpha$
  - $x$  has no incentive to stop caching because all other caching nodes are too far away, i.e., the remote access cost are larger than  $\alpha$

# Price of Anarchy (PoA)

- With selfish nodes any caching system converges to a stable equilibrium state

- Unfortunately, NEs are often not optimal!



- Idea:

- Quantify loss due to selfishness by comparing the performance of a system at Nash equilibrium to its optimal performance

- Since a game can have more than one NE it makes sense to define a **worst-case Price of Anarchy (PoA)**, and an **optimistic Price of Anarchy (OPoA)**

Definition

$$PoA = \frac{\text{cost}( \textit{worst NE} )}{\text{cost}( \textit{social Opt} )}$$

Definition


$$OPoA = \frac{\text{cost}( \textit{best NE} )}{\text{cost}( \textit{social Opt} )}$$

- $PoA \geq OPoA \geq 1$
    - A  $PoA$  close to 1 indicates that a system is insusceptible to selfish behavior


# PoA for Selfish Caching

- How large is the (optimistic) price of anarchy in the following examples?


1)  $\alpha = 4$ ,  $w_i = 1$



2)  $\alpha = 4$




3)  $\alpha = 101$




# PoA for Selfish Caching with constant demand and distances

- PoA depends on demands, distances, and the topology
- If all demands and distances are equal (e.g.  $w_i = 1$ ,  $d_{ij} = 1$ ) ...
  - How large can the PoA grow in cliques?



- How large can the PoA grow on a star?




- How large can PoA grow in an arbitrary topology?



# PoA for Selfish Caching with constant demand

- PoA depends on demands, distances, and the topology
- Price of anarchy for selfish caching can be linear in the number of nodes even when all nodes have the same demand ( $w_i = 1$ )




$$\text{cost}( NE ) = \alpha + \frac{n}{2}(\alpha - \epsilon)$$

$$\text{cost}( OPT ) = 2 \cdot \alpha$$

$$PoA = OPoA \stackrel{\epsilon \rightarrow 0}{=} \frac{1}{2} + \frac{n}{4} \in \Theta(n)$$

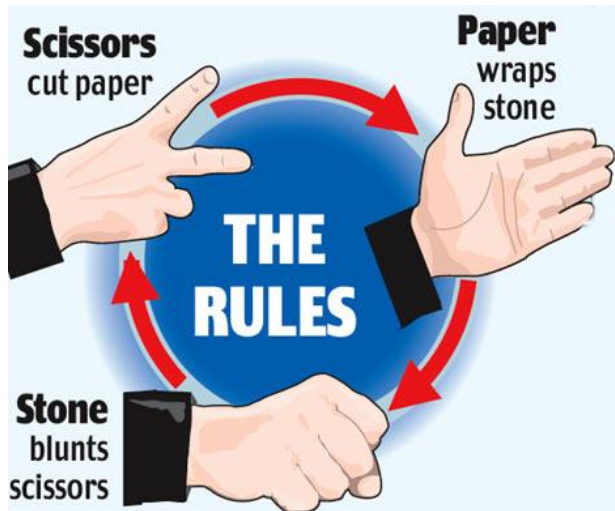
## Another Example: Braess' Paradox







- Flow of 1000 cars per hour from A to D
- Drivers decide on route based on current traffic
- Social Optimum? Nash Equilibrium? PoA?






- Is there always a Nash equilibrium?

# Rock Paper Scissors



			
	0	-1	1
	1	0	-1
	-1	1	0

- Which is the best action:  ,  , or  ?
- What is the social optimum? What is the Nash Equilibrium?
- Any good strategies?

# Mixed Nash Equilibria

- Answer: Randomize !
  - Mix between pure strategies. A **mixed strategy** is a probability distribution over pure strategies.
  - Can you beat the following strategy in expectation?  
(  $p[\text{Rock}] = 1/2$ ,  $p[\text{Paper}] = 1/4$ ,  $p[\text{Scissors}] = 1/4$  )
  - The only (mixed) Nash Equilibrium is  $(1/3, 1/3, 1/3)$
  - Rock Paper Scissors is a so-called Zero-sum game



Theorem [Nash 1950]

Every game has a mixed Nash equilibrium

# Solution Concepts

- A solution concept predicts how a game turns out

## Definition

A **solution concept** is a rule that maps games to a set of possible outcomes, or to a probability distribution over the outcomes

- The **Nash equilibrium** as a solution concept predicts that any game ends up in a strategy profile where nobody can improve unilaterally.  
If a game has multiple NEs, then the game ends up in any of them.
- Other solution concepts:
  - **Dominant strategies**
    - A game ends up in any strategy profile where all players play a dominant strategy, given that the game has such a strategy profile
    - A strategy is dominant if, regardless of what any other players do, the strategy earns a player a larger payoff than any other strategy.
  - There are more, e.g. correlated equilibrium

# Prisoner's Dilemma

- One of the most famous games in game theory is the so called Prisoner's Dilemma
  - Two criminals A and B are charged with a crime, but only circumstantial evidence exists
  - Both can cooperate (**C**), i.e., stay silent or they can defect (**D**), i.e., talk to the police and admit their crime
  - If both cooperate, each of them has to go to prison for one year
  - If both defect, each of them has to go to prison for three years
  - If only A defects but B chooses to cooperate, A is a crown witness and does not have to serve jail time but B gets three years (and vice versa)

- Dominant strategy is to defect

	<b>C</b>	<b>D</b>
<b>C</b>	-1 -1	-3 0
<b>D</b>	0 -3	-2 -2

# How can Game Theory help?

- Economy
  - Understand markets?
  - Predict economy crashes?
  - Sveriges Riksbank Prize in Economics (“Nobel Prize”) has been awarded many times to game theorists
- Problems
  - GT models the real world inaccurately
  - Many real world problems are too complex to capture by a game
  - Human beings are not really rational
- GT in computer science
  - Players are not exactly human
  - Explain unexpected deficiencies (emule, bittorrent etc.)
  - Additional measurement tool to evaluate distributed systems

# Mechanism Design

- **Game Theory** describes existing systems
  - Explains, or predicts behavior through solution concepts (e.g. Nash Equilibrium)
- **Mechanism Design** creates games in which it is best for an agent to behave as desired by the designer
  - incentive compatible systems
  - Most popular solution concept: dominant strategies
  - Sometimes Nash equilibrium
  - Natural design goals
    - Maximize social welfare
    - Maximize system performance

Mechanism design  $\approx$  „inverse“ game theory




# Incentives

- How can a mechanism designer change the incentive structure?
  - Offer rewards, or punishments for certain actions
    - Money, better QoS
    - Imprisonment, fines, worse QoS
  - Change the options available to the players
    - Example: fair cake sharing (MD for parents)




# Selfish Caching with Payments

- Designer enables nodes to reward each other with **payments**
- Nodes offer **bids** to other nodes for caching
  - Nodes decide whether to cache or not after all bids are made




- $OPoA = 1$
- However,  $PoA$  at least as bad as in the basic game



# Selfish Caching: Volunteer Dilemma

- Clique
  - Constant distances  $d_{ij} = 1$
  - Variable demands  $1 < w_i < \alpha = 20$
- Who goes first?
  - Node with highest demand?
  - How does the situation change if the demands are not public knowledge, and nodes can lie when announcing their demand?




# First-Price Auction

- Mechanism Designer
  - Wants to minimize social cost
  - Is willing to pay money for a good solution
  - Does not know demands  $w_i$



## Idea: Hold an auction

- Auction should generate competition among nodes. Thus get a good deal.
  - Nodes place private bids  $b_i$ . A bid  $b_i$  represents the minimal payment for which node  $i$  is willing to cache.
  - Auctioneer accepts lowest offer.  
Pays  $b_{min} = \min_i b_i$  to the bidder of  $b_{min}$ .
- What should node  $i$  bid?
    - $\alpha - w_i \leq b_i$
    - $i$  does not know other nodes' bids



# Second-Price Auction

- The auctioneer chooses the node with the lowest offer, but pays the price of the second lowest bid!
- What should  $i$  bid?
  - Truthful ( $b_i = \alpha - w_i$ ), overbid, or underbid?



Theorem


Truthful bidding is the dominant strategy in a second-price auction

# Proof

- Let  $v_i = \alpha - w_i$ . Let  $b_{min} = \min_{j \neq i} b_j$ .
- The payoff for  $i$  is  $b_{min} - v_i$  if  $b_i < b_{min}$ , and 0 otherwise.
- „truthful dominates underbidding“
  - If  $b_{min} > v_i$  then both strategies win, and yield the same payoff.
  - If  $b_{min} < b_i$  then both strategies lose.
  - If  $b_i < b_{min} < v_i$  then underbidding wins the auction, but the payoff is negative. Truthful bidding loses, and yields a payoff of 0.
  - Truthful bidding is never worse, but in some cases better than underbidding.
- „truthful dominates overbidding“
  - If  $b_{min} > b_i$  then both strategies win and yield the same payoff
  - If  $b_{min} < v_i$  then both strategies lose.
  - If  $v_i < b_{min} < b_i$  then truthful bidding wins, and yields a positive payoff. Overbidding loses, and yields a payoff of 0.
  - Truthful bidding is never worse, but in some cases better than overbidding.
- Hence truthful bidding is the dominant strategy for all nodes  $i$ .

## Another Approach: 0-implementation

- A third party can implement a strategy profile by offering high enough „insurances“
  - A mechanism implements a strategy profile  $S$  if it makes all strategies in  $S$  dominant.
- Mechanism Designer publicly offers the following deal to all nodes except to the one with highest demand,  $p_{max}$ :
  - „If nobody choses to cache I will pay you a millinillion.“
- Assuming that a millinillion compensates for not being able to access the file, how does the game turn out?



Theorem

Any Nash equilibrium can be implemented for free

# MD for P2P file sharing


- Gnutella, Napster etc. allow easy free-riding
- BitTorrent suggests that peers offer better QoS (upload speed) to collaborative peers
  - However, it can also be exploited
  - The BitThief client downloads without uploading!
    - Always claims to have nothing to trade yet
    - Connects to much more peers than usual clients



- Many techniques have been proposed to limit free riding behavior
  - Tit-for-tat (T4T) trading
    - Allowed fast set (seed capital),
    - Source coding,
    - indirect trading,
    - virtual currency...
  - Reputation systems
    - shared history




increase trading opportunities





# MD in Distributed Systems: Problems

- Virtual currency
  - no trusted mediator
  - Distributed mediator hard to implement
- Reputation systems
  - collusion
  - Sibyl attack



- Malicious players
  - Nodes are not only selfish but sometimes Byzantine

# Credits

- The concept for a Nash Equilibrium is from John Nash, 1950
- The definition of a Price of Anarchy is from Koutsoupias and Papadimitriou, 1999
- The Selfish Caching Game is from Chun, Chaudhuri, Wee, Barreno, Papadimitriou, and Kubiawicz, 2004
- The Prisoner's Dilemma was first introduced by Flood and Dresher, 1950
- A generalized version of the second-price auction is a VCG auction, introduced by Vickrey, Clarke, and Groves, 1973

# *That's all, folks!*

*Questions & Comments?*



*Roger Wattenhofer*