

COMP/MATH 553: Algorithmic Game Theory


Lecture 1

Fall 2016

Yang Cai

www.cs.mcgill.ca/~cai

*what we **won't** study in this class...*



I only mean this as a metaphor of what we usually study in Eng.:


- *central design*
- *cooperative components*
- *rich theory*

*what we **will** study in this class...*

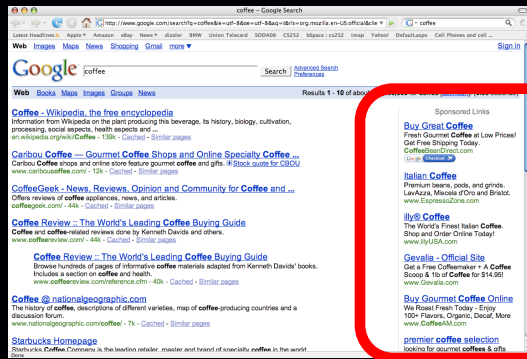
Markets



Routing in Networks



Online Advertisement



Evolution






Elections



Social networks

facebook



facebook

Game Theory

- *central design ?*
- *cooperative components ?*
- *rich theory ?*

we will study (and sometimes question) the algorithmic foundations of this theory

Game Theory



$0,0$	$-1,1$	$1,-1$
$1,-1$	$0,0$	$-1,1$
$-1,1$	$1,-1$	$0,0$

Game Theory

Games are thought experiments to help us learn how to *predict rational behavior* in *situations of conflict*.

Rational Behavior: a player will always prefer to improve her situation regardless of the consequences.

- Encoded via a utility function.
- Players want to maximize their own expected utility.
- Much of the “*irrationality*” can also be incorporated ...
- Assume other people are as smart or smarter than you ...

Game Theory

Games are thought experiments to help us learn how to *predict rational behavior* in *situations of conflict*.

Situation of conflict: Everybody's action affects others.


- In traditional optimization, the goal is optimizing a single objective.
- In game theory, every player has a utility function to optimize.
- The outcome depends on the player's own choice and the choices of the others.

Game Theory

Games are thought experiments to help us learn how to *predict rational behavior* in *situations of conflict*.

Predict: We want to know what happens in a game. Such predictions are called solution concepts (e.g., Nash equilibrium).

Algorithmic Game Theory



0,0	-1,1	1,-1
1,-1	0,0	-1,1
-1,1	1,-1	0,0

Mechanism Design: How can we design a system that will be launched and used by competitive users to optimize our objectives ?

?

Can we predict what will happen in a large system?

Game theory says yes, through its prediction tools (solution concepts).

Can we *efficiently* predict what will happen in a large system?

Are the predictions of Game Theory *plausible*, i.e. likely to arise?



An overview of the class

- *Administration*
- *Solution Concepts*
- *Equilibrium Computation*
- *Price of Anarchy*
- *Mechanism Design*

An overview of the class

- *Administration*
- *Solution Concepts*
- *Equilibrium Computation*
- *Price of Anarchy*
- *Mechanism Design*

Administrivia

TA: Mingfei Zhao (Office Hours: consult webpage)

Prereq: A course in algorithms (COMP 360 or equivalent) and probability (MATH 323 or equivalent). No prior knowledge of Economics or Game Theory is required.

Attendance/Evaluation:

Everybody is welcome

Requirements, if registered for credit:

- (10%) Class attendance and participation (including playing a few games in class)
- (45%) Solve problem sets: 3 problems sets, two weeks to solve (15% each).
- (45%) Final exam.

If just auditing: Become a listener. Why not?

An overview of the class

- *Administration*
- *Solution Concepts*
- *Equilibrium Computation*
- *Price of Anarchy*
- *Mechanism Design*

Battle of the Sexes

	Theater!	Football fine
Theater fine	1, 5	0, 0
Football!	0, 0	5, 1

Think of this game as a **metaphor** of real-life examples, not necessarily in the context of a couple's decision making, not necessarily about football vs theater, and not necessarily with exactly these numerical values associated to different outcomes.

Nash Equilibrium: A pair of strategies (deterministic or randomized) such that the strategy of the row player is a *Best Response* to the strategy of the column player and vice versa.

Aims to capture the behavioral norms to which a society where this game occurs may converge.

Battle of the Sexes

		5/6	1/6
		Theater!	Football fine
1/6	Theater fine	1, 5	0, 0
5/6	Football!	0, 0	5, 1

Think of this game as a **metaphor** of real-life examples, not necessarily in the context of a couple's decision making, not necessarily about football vs theater, and not necessarily with exactly these numerical values associated to different outcomes.

Nash Equilibrium: A pair of strategies (deterministic or randomized) such that the strategy of the row player is a *Best Response* to the strategy of the column player and vice versa.

Aims to capture the behavioral norms to which a society where this game occurs may converge.

(Theater fine, Theater!)

Matriarchic society

(Football!, Football fine)

Patriarchic society

In medias res: There is always an odd number of Nash Equilibria

Remark:

Matrix-form games intend to model repeated occurrences of the same conflict, provided that there are no strategic correlations between different occurrences of the game. If such correlations exist, we exit the realm of matrix-form games, entering the realm of **repeated games**.

How can repeated occurrences happen without inter-occurrence correlations?

Imagine a population of **blue players** (these are the ones preferring football) and **orange players** (these are those preferring theater). Members of the **blue** population encounter members of the **orange** population at random and need to decide whether to watch football or theater.

What do the Nash equilibria represent?

The Nash equilibria predict what types of behaviors and (in the case of randomized strategies) at what proportions will arise in the two populations at the steady state of the game.

Battle of the Sexes

Suppose now that the blue player removes a strategy from his set of strategies and introduces another one:

	Theater!	Football fine
Theater fine	1, 5	0, 0
Football!	0, 0	5, 1
Theater great, I'll invite my mom	2, -1	0, 0

unique Equilibrium

(Football!, Football fine)

Moral of the story: The player who knows game theory managed to eliminate her unwanted Nash equilibrium from the game.

Back to Rock-Paper-Scissors



0,0	-1,1	1,-1
1,-1	0,0	-1, 1
-1,1	1, -1	0,0

The unique Nash Equilibrium is the pair of uniform strategies.

Contrary to the battle of the sexes, in RPS randomization is necessary to construct a Nash equilibrium.

Rock-Paper-Scissors Championship



The behavior observed in the RPS championship is very different from the pair of uniform strategies; indeed, the matrix-form version of RPS did not intend to capture the repeated interaction between the same pair of players---recall earlier remark; rather the intention is to model the behavior of a population of, say, students in a courtyard participating in random occurrences of RPS games

Guess Two-Thirds of the Average

- k players $p_1, p_2, p_3, \dots, p_k$
- each player submits a number in $[0, 100]$

$$x_1, x_2, \dots, x_k$$

- compute

$$\bar{x} := \frac{1}{k} \sum_{i=1}^k x_i$$

Let's Play!

- find x_j , closest to $\frac{2}{3}\bar{x}$
- player p_j wins \$100, all other players win nothing

Guess Two-Thirds of the Average

Is it rational to play above $\frac{2}{3} \cdot 100$?

A: no (why?)

Given that no rational player will play above $\frac{2}{3} \cdot 100$ is it rational to play above $(\frac{2}{3})^2 \cdot 100$?

A: no (same reasons)

⋮

All rational players should play 0.

The all-zero strategy is the only Nash equilibrium of this game.

Rationality versus common knowledge of rationality

historical facts: 21.6 was the winning value in a large internet-based competition organized by the Danish newspaper [Politiken](#). This included 19,196 people and with a prize of 5000 Danish kroner.

OK, Nash equilibrium makes sense and is stable, but does it always exist?

2-player Zero-Sum Games

	 1/3	 1/3	 1/3
 1/3	0,0	-1,1	1,-1
 1/3	1,-1	0,0	-1, 1
 1/3	-1,1	1, -1	0,0

von Neumann '28:

In two-player zero-sum games, it always exists.

[original proof used analysis]

Danzig '47



LP duality

Poker



von Neuman's predictions are in fact accurate in predicting players' strategies in two-player poker!

Poker




von Neuman's predictions are in fact accurate in predicting players' strategies in two-player poker!

But what about larger systems (more than 2 players) or systems where players do not have directly opposite interests?

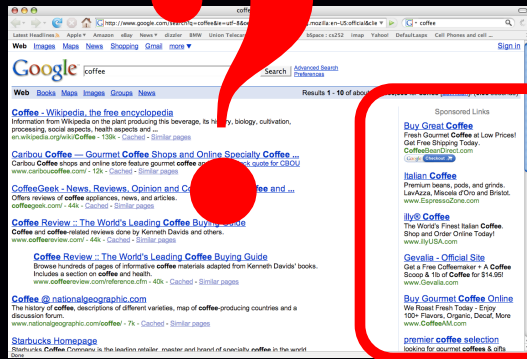
Markets



Routing in Networks



Online Advertisement




Evolution



Elections









Social networks



Modified Rock Paper Scissors

Not zero-sum any more

	25% 	50% 	25% 
33% 	0,0	-1, 1	2,-1
33% 	1,-1	0,0	- 1 , 1
33% 	- 2, 1	1 , -1	0,0

Is there still an equilibrium?
[that is a pair of randomized strategies so that no player has incentive to deviate given the other player's strategy ?]

John Nash '51:


There always exists a Nash equilibrium, regardless of the game's properties.

Nobel 1994, due to its large influence in understanding systems of competitors...

Markets



Routing in Networks




and every other game!

Evolutionary Biology



Elections



Social Networks



Applications...

game =

market  price equilibrium







Internet  packet routing

roads  traffic pattern

facebook,
twitter, linkedin, ...  structure of the social network

Modified Rock Paper Scissors

Not zero-sum any more

	25% 	50% 	25% 
33% 	0,0	-1, 1	2,-1
33% 	1,-1	0,0	-1, 1
33% 	-2, 1	1, -1	0,0

Highly Non-Constructive

Is there still an equilibrium?


[that is a pair of randomized strategies so that no player has incentive to deviate given the other player's strategy ?]

Brouwer's Fixed Point Theorem

John Nash '51:

There always exists a Nash equilibrium, regardless of the game's properties.

Nobel 1994



How can we compute a Nash equilibrium?

- if we had an *algorithm* for equilibria we could predict what behavior will arise in a system, before the systems is launched

- in this case, we can easily compute the equilibrium, thanks to gravity!



An overview of the class

- *Administration*
- *Solution Concepts*
- *Equilibrium Computation*
- *Price of Anarchy*
- *Mechanism Design*

2-player zero-sum vs General Games

1928 Neumann:

- existence of min-max equilibrium in *2-player, zero-sum* games;
- proof uses analysis;
- + Danzig '47: equivalent to LP duality;
- + Khachiyan '79: poly-time solvable;
- + a multitude of distributed algorithms converge to equilibria.

1950 Nash:

- existence of an equilibrium in *multiplayer, general-sum* games;
- Proof uses Brouwer's fixed point theorem;
- intense effort for equilibrium computation algorithms:
 - Kuhn '61, Mangasarian '64, Lemke-Howson '64, Rosenmüller '71, Wilson '71, Scarf '67, Eaves '72, Laan-Talman '79, etc.
- Lemke-Howson: simplex-like, works with LCP formulation;
- no efficient algorithm is known after 50+ years of research.
- hence, also no efficient dynamics ...

the Pavlovian reaction

“Is it NP-complete to find a Nash equilibrium?”

Why should we care about the complexity of equilibria?

- First, if we believe our equilibrium theory, efficient algorithms would enable us to make predictions:

Herbert Scarf writes...

“[Due to the non-existence of efficient algorithms for computing equilibria], general equilibrium analysis has remained at a level of abstraction and mathematical theoretizing far removed from its ultimate purpose as a method for the evaluation of economic policy.”

The Computation of Economic Equilibria, 1973

- More importantly: If equilibria are supposed to model behavior, computational tractability is an important modeling *prerequisite*.

“If your laptop can't find the equilibrium, then how can the market?”

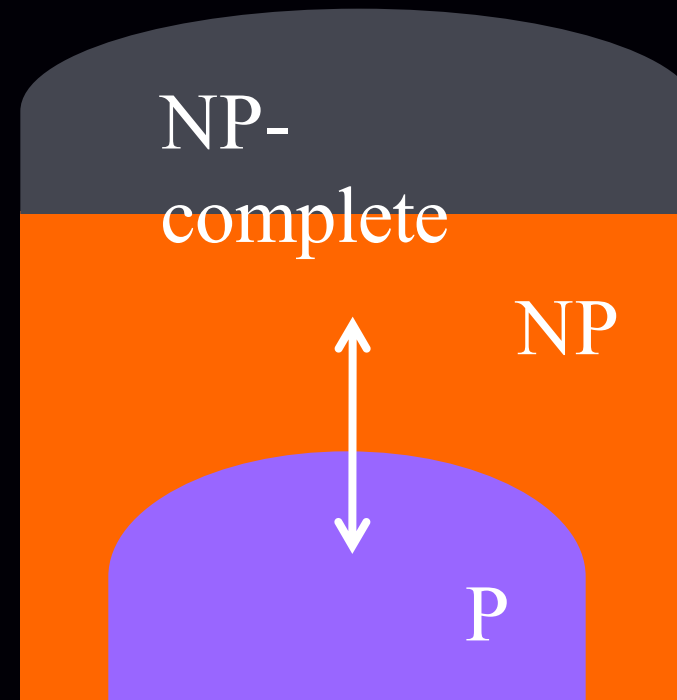
Kamal Jain, EBay

N.B. computational intractability implies the non-existence of efficient dynamics converging to equilibria; how can equilibria be universal, if such dynamics don't exist?

the Pavlovian reaction

“Is it NP-complete to find a Nash equilibrium?”


- the theory of NP-completeness does not seem appropriate;
- in fact, NASH seems to lie below NP-complete;
- **Stay tuned!** we are going to answer this question later this semester.



An overview of the class

- *Administration*
- *Solution Concepts*
- *Equilibrium Computation*
- *Price of Anarchy*
- *Mechanism Design*

Traffic Routing



Delay is 1.5 hours for everybody at the unique Nash equilibrium

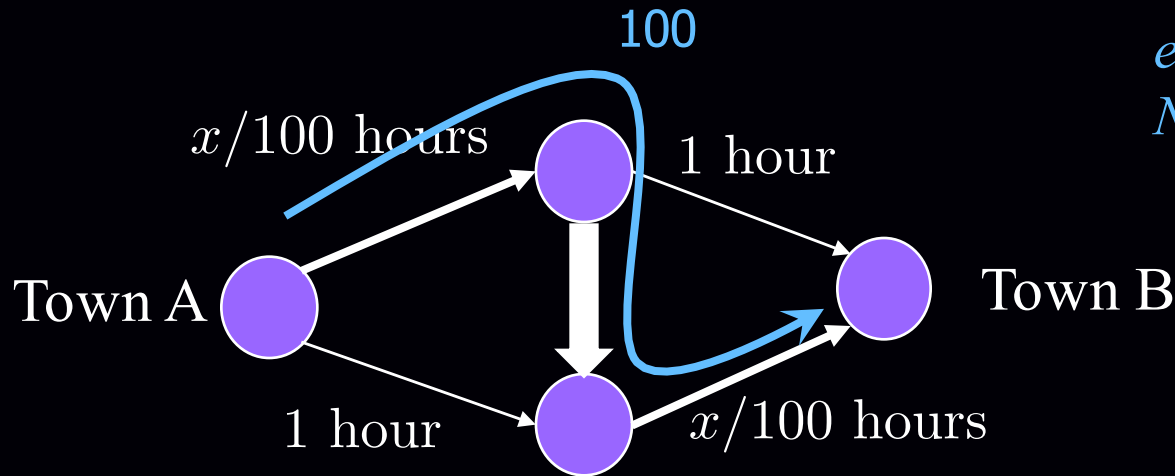
Suppose 100 drivers leave from town A towards town B.

Every driver wants to minimize her own travel time.

What is the traffic on the network?

In any unbalanced traffic pattern, all drivers on the most loaded path have incentive to switch their path.

Traffic Routing




Delay is 2 hours for everybody at the unique Nash equilibrium

A benevolent mayor builds a superhighway connecting the fast highways of the network.

What is now the traffic on the network?

No matter what the other drivers are doing it is always better for me to follow the zig-zag path.

Traffic Routing



Adding a fast road on a road-network is not always a good idea!

Braess's paradox

In the RHS network there exists a traffic pattern where all players have delay 1.5 hours.

$$PoA = \frac{\text{performance of system in worst Nash equilibrium}}{\text{optimal performance if drivers did not decide on their own}}$$

Price of Anarchy [Koutsoupias-Papadimitriou '99]: measures the loss in system performance due to free-will.

4/3

Traffic Routing

Obvious Questions:

What is the PoA in a system?

How do we design a system whose PoA is small?

In other words, what incentives can we provide to induce performance that is close to optimal?

E.g. tolls?

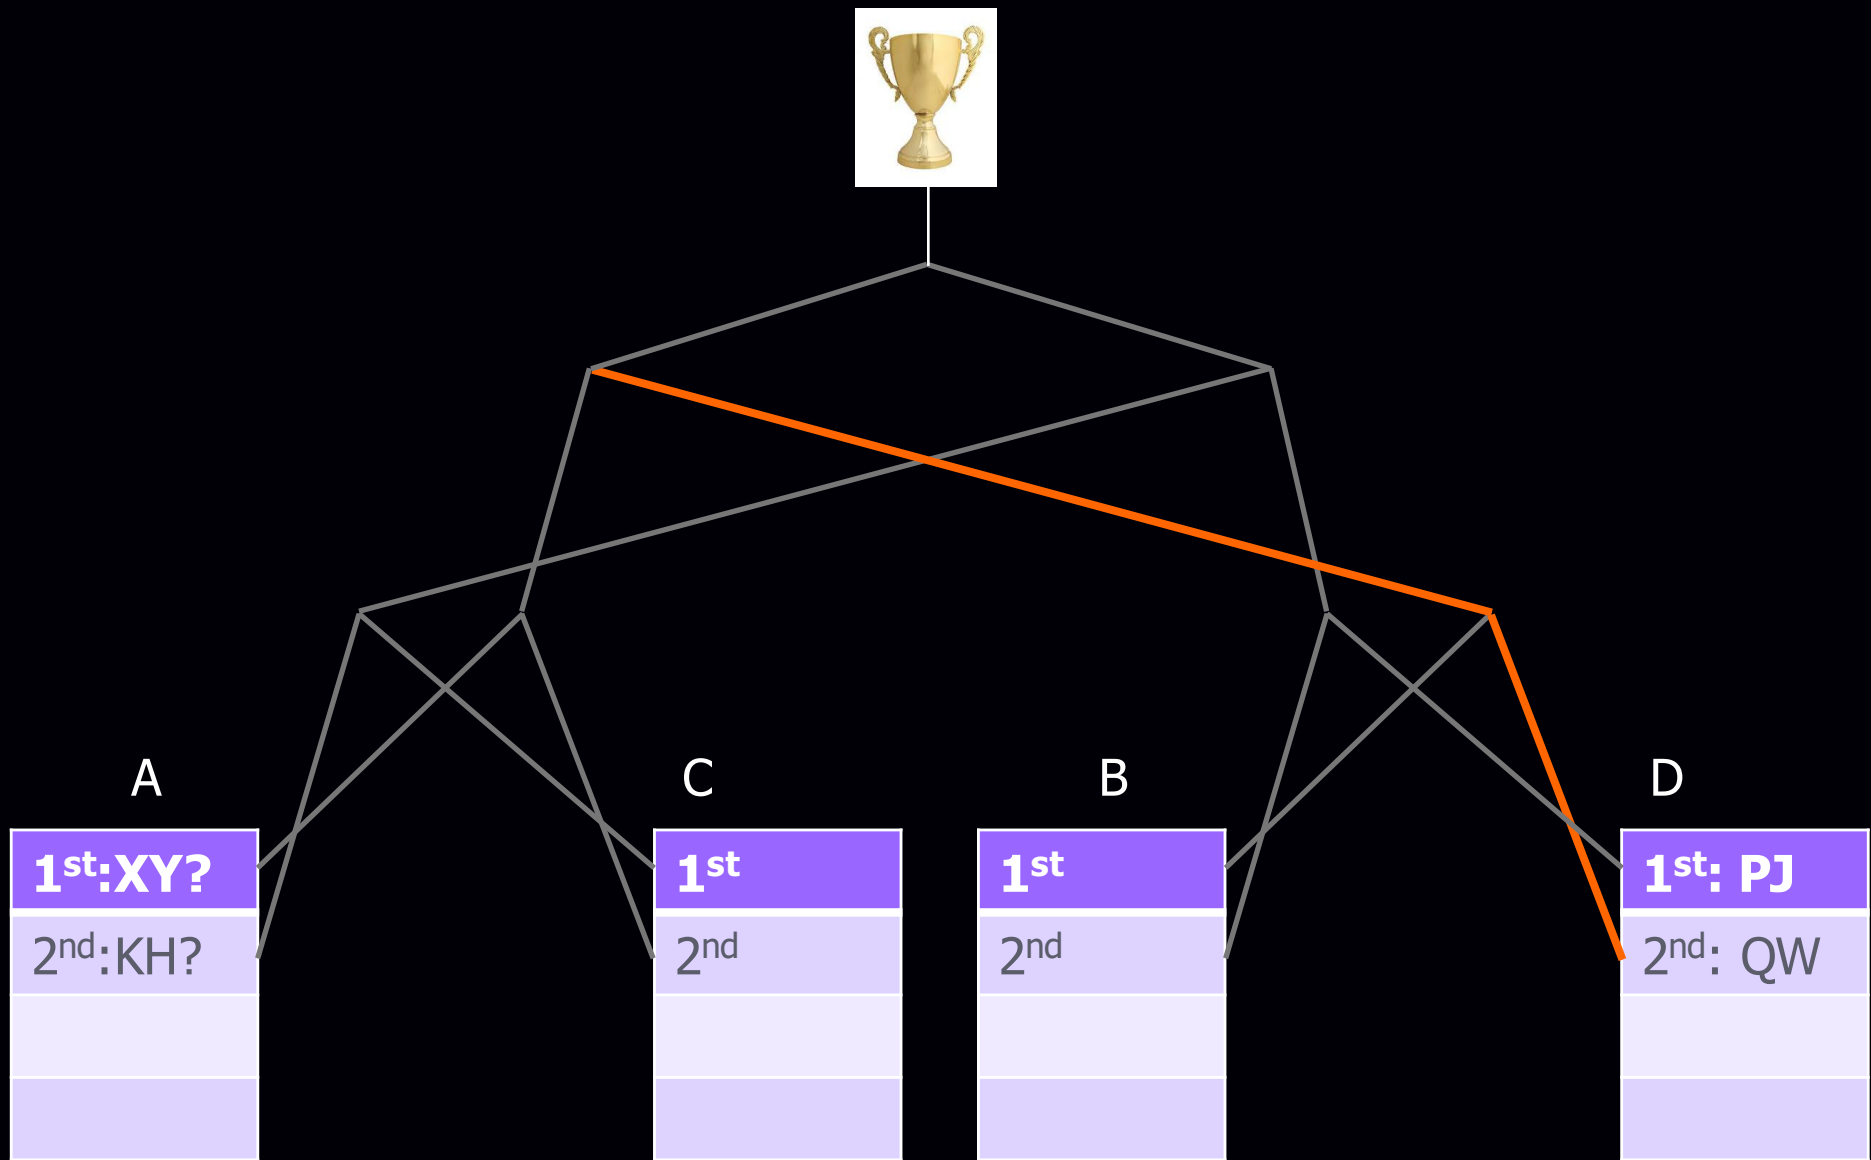
An overview of the class

- *Administration*
- *Solution Concepts*
- *Equilibrium Computation*
- *Price of Anarchy*
- *Mechanism Design*

Mechanism Design

- Understand how to design systems with strategic participants that have good performance guarantees
- Let's start with a poorly designed system...
- Biggest scandal in 2012 London Olympics
 - Women's badminton double
 - No doping
 - Failed tournament design

2012 London Olympics Scandal



2012 London Olympics Scandal

- Four teams are disqualified, due to “not using one’s best effort to win a game”
- Ironically, they are using their best effort to win a Olympic medal
- System rule matters: poorly designed systems suffer from unexpected and undesirable results
- To quote Hartline and Kleinberg:

“The next time we bemoan people exploiting loopholes to subvert the intent of the rule makers, instead of asking ‘What’s wrong with these people?’ let’s instead ask, ‘What’s wrong with the rules?’ and then adopt a scientifically principled approach to fixing them.”

In conclusion

- *We are going to study, question and develop the algorithmic foundations of Game Theory*
- *Modeling systems of strategic individuals*
- *Predicting strategic behavior*
 - solution concepts, dynamics, ...*
- *Complexity of finding equilibria*
 - NP-completeness theory not relevant, new theory below NP ...*
- *System Design*
 - auctions, voting, policy making*
 - robustness against strategic manipulation*
- *Price of Anarchy*
 - Network routing.*