Algorithmic Game Theory and Applications

Lecture 16: a brief glimpse beyond

Kousha Etessami

School of informatics

warning, again

- The subjects we will briefly glimpse at today are part of a fast emerging field at the intersection of game theory, CS, economics and e-commerce, and "internetology".

  Numerous advanced courses are being taught on these subjects around the world, and there are conferences and workshops dedicated to them. See, e.g., the proceedings of the ACM conference on Electronic Commerce.

  See the recent textbook *Algorithm Game Theory*, edited by N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Cambridge Press, 2007. (A collection of chapters written by different experts).

- We won't even scratch the surface (of the surface) of these subjects, so please do explore them further if they interest you.

games and the internet

- Basic idea: "The internet is a huge experiment in interaction between agents (both human and automated)".
- Such interactions can profitably be viewed from a game theoretic viewpoint: agents trying to maximize their own payoffs, etc.
- How do we set up the rules of these games to harness "socially optimal" results?
- These vague notions can best be illustrated by examples.

a flow network game

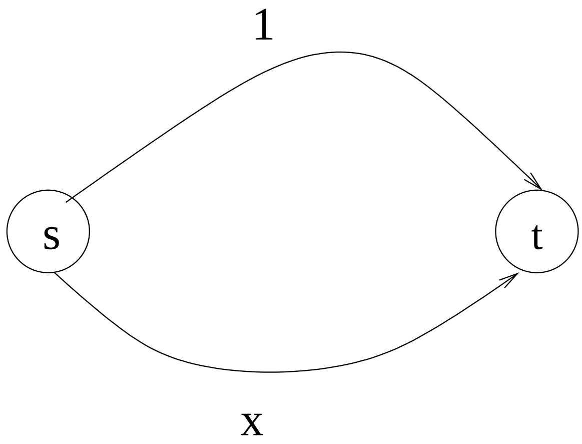

(from [Roughgarden-Tardos'00])

- A large number n of customers in the network want to go from s to t.
- Each can either take the edge labeled 1, with "latency" 1 (delay of crossing edge), or edge labeled with latency x. Here x represents the "congestion", given by the ratio of the number of customers that are using that edge divided by the total n.
- Assume n is very large, (basically, $n \to \infty$).
- What is the Nash Equilibrium?
  (NEs in such a setting can be shown to be essentially unique [Beckmann, et. al. '56].)
- What is a "globally optimal" customer routing strategy profile that minimizes average delay?
  What is the globally optimal average delay?

a modified game

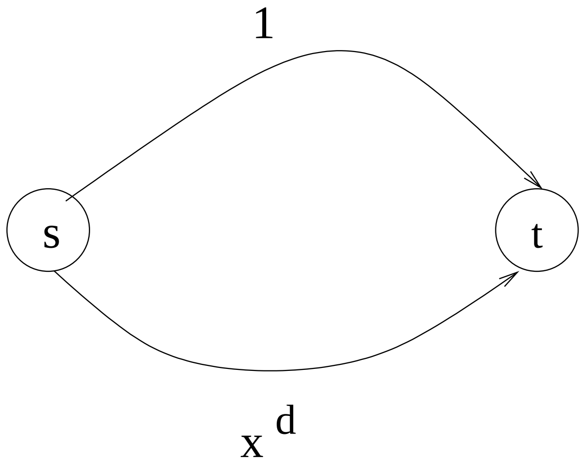

- What is the NE, and what is the average delay it induces?

- What is the globally optimal average delay?

a different network

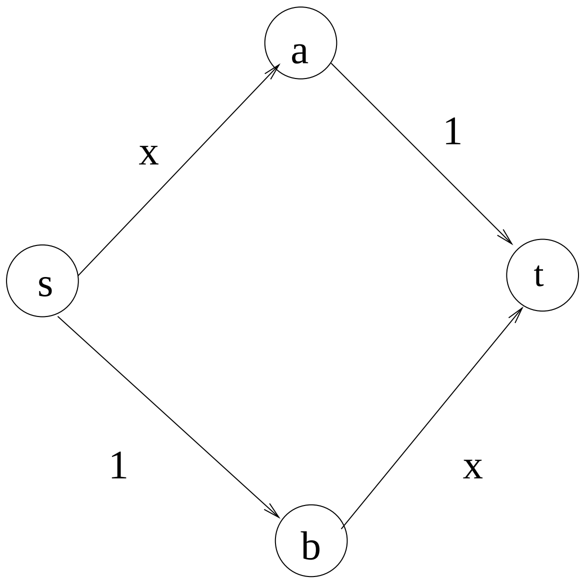

- What is the NE, and what is its average delay?
- What is a globally optimal strategy profile and optimal average delay?
- What if an ambitious "network service provider" wanted to build additional "superfast" lines?

Braess's paradox

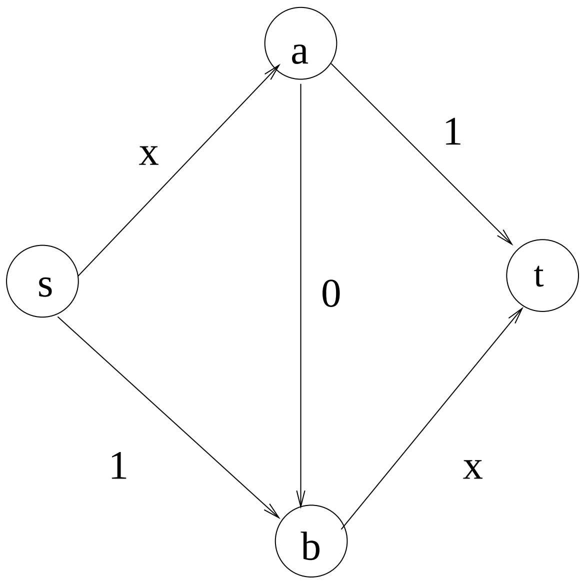

- What is the NE and its average delay?
- What is the globally optimal average delay?

social welfare and the price of anarchy

Recall that in a strategic game Γ, we may have different measures of the "social welfare" *welfare*(x) under a particular profile of (mixed) strategies $x \in X$. For example, "utilitarian" social welfare is *welfare*$(x) := \sum_{i=1}^{n} U_i(x)$. For a game Γ, let $NE(\Gamma)$ be the set of NE's of Γ.
For our next definition suppose *welfare*$(x) > 0$ for all $x \in X$. (In many games we could enforce this by, e.g., "shifting" all payoffs by an additive factor.)

A version of "the price of anarchy" can be defined as: ([Koutsoupias-Papadimitriou'98,Papadimitriou'01])

$$PA(\Gamma) := \frac{\max_{x \in X} \mathit{welfare}(x)}{\min_{x \in \mathrm{NE}(\Gamma)} \mathit{welfare}(x)}$$

Thus, the "price of anarchy" is the ratio of best "global" outcome to the the worst NE outcome. Note: this ratio is ≥ 1 and larger means "worse". (Perhaps "price" is not the best word for this variant.)

It would be comforting to establish that in various situations the "price of anarchy" isn't too high.

price of anarchy in the flow network game

- For flow f let *welfare*$(f) := 1/$(average s-t-delay).

- In Braess's paradox, the price of anarchy is $4/3$: by playing the NE the average delay is 2, but playing half-and-half on the upper and lower route, the average delay is $3/2$ (and that's optimal).

- We have seen that the price of anarchy in network games can be arbitrarily high, when x^d is an edge label.

- Remarkably, [Roughgarden-Tardos'00] showed that in a more general flow network setting (where there can be multiple source-destination pairs (s_j, t_j)), as long as "congestions" labeling edges are linear functions of x, the worst-case price of anarchy is $4/3$.

- In other words, for linear latencies, Braess's paradox is the worst-case scenario.

auctions as games

Auctions have been studied game theoretically for a long time. Consider one such formulation:

- Each of n bidders is a player.
- Each player i has a "valuation" $v_i \in \mathbb{R}$ for the item being bid on.
- The payoff to player i, if he/she wins the auction at price pr_i is
$$u_i(outcome) := v_i - pr_i$$
- All other players get payoff 0.
- Under the following constraints, the auctioneer is free to set up the rules of the auction: given bids $(b_1, \ldots, b_n)$ one of the highest bidders must win, and at a price $\leq \max_i b_i$.
- Question: What rules should the auctioneer employ, so that for each player i, revealing the "true worth" v_i is a dominant strategy (i.e., $b_i = v_i$ is a dominant strategy)?

vickery auctions and mechanism design

- In a Vickery auction, a.k.a., *second-price, sealed bid* auction, the highest bidder wins but the price paid by the highest bidder is the second highest bid price.

- Convince yourself: bidding your true “valuation” v_i is a dominant strategy in this game.

- This is the starting point of a vast sub-discipline of game theory called “Mechanism Design”, where the goal is to design a game where selfish player will behave in a desired way (e.g., a “socially optimal” way).

- Many algorithmic issues impinge on mechanism design. See, e.g., [Nisan-Ronen’99,Nisan’99-00, . . .]. See Nisan’s chapters on Mechanism Design in the AGT book referenced on page 1 of these notes.

a formal setting for mechanism design

Formally, the input to a mechanism design problem can be specified by giving three things:

- An **environment**, $\mathcal{E}$, in which the game designer operates.
- A **choice function**, f, which describes the designer's preferred outcomes in the given environment.
- A **solution concept**, *Sol*, which describes the kinds of strategy profiles of games (such as their Nash Equilibria) which will be considered "solutions" in which the desired choice function should be implemented.

We describe each of these three inputs separately, using auctions as a running example to illustrate each concept.

Once this is done, we will be able to state the mechanism design problem formally.

environments

An **environment**, $\mathcal{E} = (N, C, \Theta, \mathcal{M})$, consists of:

- A set $N = \{1, \ldots, n\}$ of players.
- A set C of "outcomes".
  auction example: C could be a set of pairs (i, pr), meaning that player i wins and pays price *pr*.
- A cartesian product $\Theta = \Theta_1 \times \Theta_2 \times \ldots \times \Theta_n$, where each set Θ_i is a set of possible "utility functions" for player i. Each $u_i \in \Theta_i$, is a function $u_i : C \mapsto \mathbb{R}$, that maps outcomes to payoffs for player i.
  auction example: u_i could be given by
  $$u_i((i', \mathit{pr})) = \begin{cases} 0 & \text{if } i \neq i' \\ v_i - \mathit{pr} & \text{if } i = i' \end{cases}$$
- A set $\mathcal{M}$ of "mechanisms". Each $M \in \mathcal{M}$ has form $M = (S_1, \ldots, S_n, g)$, where S_i is a set of strategies for player i, and $g : S_1 \times \ldots \times S_n \mapsto C$, is a function from strategy profiles to "outcomes".
  auction example: the strategy sets S_i give the possible bids for each player, and g gives a function that maps a profile of bids to an "outcome" of who wins and at what price. Such outcome functions can be constrained (using $\mathcal{M}$), e.g., to only allow top bidders to win, and to ask for at most the maximum bid price to be paid.

choice functions

Given an environment $\mathcal{E}$, a **choice function**

$$f : \Theta \mapsto 2^C$$

for the designer, specifies what set of outcomes the designer would prefer if it knew exactly what utility (payoff) function each player has.
auction example: the choice function of the designer could be, e.g., optimize social welfare, i.e., $f((u_1, \ldots, u_n)) = \{(i, pr) \in C \mid i = arg\,\max_i v_i \;\&\; pr = \max_{j \neq i} v_j\}$.

Note that a mechanism $M = (S_1, \ldots, S_n, g) \in \mathcal{M}$ together with the tuple $u \in \Theta$, gives a complete description of a strategic game $\Gamma_{M,u}$, where the payoff function for player i can now be described as a function of strategy profiles: by abuse of notation, for the game $\Gamma_{M,u}$ we can write $u_i(g(s_1, \ldots, s_n))$ as simply $u_i(s_1, \ldots, s_n)$.

solution concepts

We lastly need the notion of a "solution concept" which specifies what kind of solutions of a game we are trying to implement the choice function in. Such solutions might be Nash Equilibria, or more strong conditions such as profiles of dominant strategies.

Given an environment $\mathcal{E}$, a **solution concept** is given by a mapping *Sol*, whose domain is $\mathcal{M} \times \Theta$, and such that for each mechanism $M = (S_1, \ldots, S_n, g)$ and each tuple $u = (u_1, \ldots, u_n)$ of utility functions, *Sol*$((S_1, \ldots, S_n, g), u) \in 2^S$, where $S = S_1 \times \ldots \times S_n$.
auction example: in the auction example we may wish to implement the choice in dominant strategies. In this setting, we ask that *Sol*(M, u) be the set of strategy combinations $(s_1, \ldots, s_n)$ such that each s_i is a dominant strategy for player i in $\Gamma_{M,u}$.

finally: the mechanism design problem statement

The mechanism design problem can now be stated:

> Given environment $\mathcal{E} = (N, C, \Theta, \mathcal{M})$, a choice function f, and a solution concept *Sol*, find $M = (S_1, \ldots, S_n, g) \in \mathcal{M}$ such that for all tuples $u \in \Theta$ of utility functions, *Sol*(M, u) is a non-empty set and for all $s \in$ *Sol*(M, u), $g(s) \in f(u)$.

If such a mechanism M exists, we say M *Sol*-**implements** choice function f in environment $\mathcal{E}$. We then also say f is *Sol*-**implementable** in $\mathcal{E}$.

This was a rather long-winded definition. There is room to argue whether some of this formalization adequately captures "real life" situations of mechanism design. Nevertheless, these formalizations have proved useful to auction design and the mechanism design community. See books [Mas-Colell-Whinston-Green'95,Osborne-Rubinstein'94].

A fact which may seem surprising, particularly if one knows *Arrow's* and *Gibbard-Satterthwaite's* Theorems, is the "revelation principle".

a little social choice theory

Before discussing the revelation principle, we discuss some classic social choice theory to motivate it.

Let C be a set of *outcomes*, and let L be the set of *linear orderings* on C.

A *social choice function* is a function $f : L^n \mapsto C$. (Thus f aggregates the ordering on C of n individuals into a choice of some "preferred" outcome from C.)

A social choice function f can be *strategically manipulated* by player (voter) i, if for some $\prec_1, \ldots, \prec_n \in L$, we have $a \prec_i a'$, and $a = f(\prec_1, \ldots, \prec_n)$, and there exists some other ordering $\prec_i'$ such that $a' = f(\prec_1, \ldots, \prec_i', \ldots, \prec_n)$.

f is called *incentive compatible* if it can not be strategically manipulated by anybody.

f is a *dictatorship* if there is some player i such that for all $\prec_1, \ldots, \prec_n$, $f(\prec_1, \ldots, \prec_n) = a$ where a is the maximum outcome for player i, i.e., $\forall b \neq q$, $b \prec_i a$.

Theorem (Gibbard-Satterthwaite) If $|C| \geq 3$, and $f : L^n \mapsto C$ is an incentive compatible social choice function for which all outcomes in C are possible (i.e., the function f is onto C), then f is a dictatorship.

We remark that a closely related Theorem is the famous:

Arrow's Impossibility Theorem: There is no "really good" way to aggregate people's preference orders (over more than 2 alternatives) into a societal preference order. "Really good" here means it should satisfy the following three desired criteria:

Unanimity: if everybody has the same preference order, then that should become the societal preference order.

Non-dictatorship: no individual's preferences should dictate the societies preferences.

Independence of irrelevant alternatives: The societal preference between any pair of outcomes a and b should depend only on the preference between a and b of all individuals, and not on their preferences for other alternatives $c \neq a, b$.

the revelation principle

Define *Dom* to be the solution concept where $s = (s_1, \ldots, s_n) \in S$ is in *Dom*(M, u) if and only if for every player i, s_i is a dominant strategy in $\Gamma_{M,u}$. Such a profile s is called a dominant strategy equilibrium (DSE) of $\Gamma_{M,u}$.

In environment $\mathcal{E} = (N, C, \Theta, \mathcal{M})$, a mechanism M is called a **direct revelation mechanism** if $M = (\Theta_1, \ldots, \Theta_n, g)$. In other words, players' strategies amount to "announcing" their utility. Players can in general "lie", but.....

Proposition (*Revelation Principle*) Suppose f is *Dom*-implemented by $M = (S_1, ..., S_n, g)$ in $\mathcal{E} = (N, C, \Theta, \mathcal{M})$. Then in environment $(N, C, \Theta, \mathcal{M}')$, where $\mathcal{M}'$ consists of all direct revelation mechanisms, there is some $M' = (\Theta_1, \ldots, \Theta_n, g') \in \mathcal{M}'$ such that for all $u \in \Theta$, $u \in$ *Dom*(M', u) and there is some $s \in Dom(M, u)$, such that $g'(u) = g(s) \in f(u)$.

Thus, each player revealing their true utility function is a DSE in $\Gamma_{M',u}$. We then say M' **truthfully** *Dom*-**implements** f. Intuitively, RP says if f can be *Dom*-implemented by M then there is a "equivalent" revelation mechanism M' where players "might as well" reveal their true utility.

Note the contrast to the Gibbard-Satterthwaite Theorem, which says that there is no non-dictatorial social choice function for player's preferences among 3 or more alternatives for which "truth revealing" is always optimal.

The difference is that in the setting of the revelation principle, the players do not have a preference order over all alternatives, but a utility function which describes their payoff under each given alternative.

combinatorial auctions

As mentioned, many algorithmic issues impinge on mechanism design. We unfortunately haven't time to get into these. See the AGT book and, e.g., [Nisan-Ronen'99,Nisan'99-00]. The next thing we will look at is:

- Consider an auction where instead of one item, a set of k items, K, is up for auction at the same time and each bidder i may only want to bid on some subset $D_i \subseteq K$, for a price b_i.

  Question: Given bids $(D_1, b_1), (D_2, b_2), \ldots, (D_n, b_n)$ what is the most profitable allocation for the auctioneer?

  Obviously NP-hard. (Max Clique even when all bids b_i are the same.)
- If bidders have more complicated criteria for which subsets they want to bid on, how should they express this? I.e., with what "bidding language"?
- Again, there is a large literature on combinatorial auctions and its connections to mechanism design. See, e.g., [Nisan'00,Parkes'00, . . .].

THE END

(hope you enjoyed the course)