News

- Good news
    - I probably won't use 1:30 hours.
    - The talk is supposed to be easy and has many examples.
    - After the talk you will at least remember how to prove one nice theorem.
- Bad news
    - Concerning algorithmic concepts, this talk (and the part in the book it is about) won't be extraordinary fascinating.
- Ugly news
    - - (none).

The inefficiency of equilibria

Chapters 17,18,19 of AGT

---

12 April 2010

University of Bergen

Outline

The subject

- Part II of the book is about the question

How inefficient are equilibria?

- Or, still very vague but in more down to earth terminology:

If we let selfish people do what they want without much control, will they be much less happy than if we would impose more rules?

- But, of course the first question only make sense if equilibria do exist.

UNIVERSITETET I BERGEN
Institutt for informatikk

A real world example: Street crossing

vs

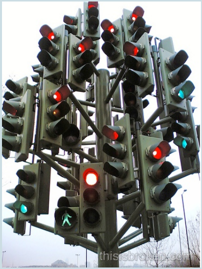

A real world example: Street crossing

P1 \ P2	Cross	Wait
Cross	-100 / -100	1 / 0
Wait	0 / 1	0 / 0

- If we allow mixed strategies, there are at least 3 equilibria:
  - Player 1 lets player 2 cross, the other way around, and
  - both players cross with probability $\frac{1}{101}$.
- Note that while the total payoff in the first two equilibria is 1, in the third it is very small and there is a small probability of a car crash.
- So typically in this situation some control is needed.

Formalization

- Like usual, we consider a game with
  - n players $\{1, \ldots, k\}$,
  - sets of strategies S_i for each player,
  - a utility function $u_i : S \rightarrow \mathbb{R}$ for each player, where $S = S_1 \times \ldots \times S_k$ is the set of all strategy vectors.
- unlike before, we also introduce a social function $\sigma : S \rightarrow \mathbb{R}$.
- Denote $E \subseteq S$ for the set of all equilibria $s^* \in S$ as the social optimum (the strategy vector maximizing σ).

Definition

The price of anarchy is $\max_{s \in E} \frac{\sigma(s)}{\sigma(s^*)}$.

Formalization

- Like usual, we consider a game with
  - n players $\{1, \dots, k\}$,
  - sets of strategies S_i for each player,
  - a utility function $u_i : S \to \mathbb{R}$ for each player, where $S = S_1 \times \dots \times S_k$ is the set of all strategy vectors.
- unlike before, we also introduce a social function $\sigma : S \to \mathbb{R}$.
- Denote $E \subseteq S$ for the set of all equilibria, and $s^* \in S$ as the social optimum (the strategy vector maximizing σ).

Definition

The price of stability is $\min_{s \in E} \frac{\sigma(s)}{\sigma(s^*)}$.

A real world example: Street crossing

P1 \ P2	Cross	Wait
Cross	-100, -100	1, 0
Wait	0, 1	0, 0

- Social cost function is expectancy of sum of payoffs of both players.

A real world example: Street crossing

P1 \ P2	Cross	Wait
Cross	-100, -100	1, 0
Wait	0, 1	0, 0

- Social cost function is expectancy of sum of payoffs of both players.
- Price of stability:

A real world example: Street crossing

P1 \ P2	Cross	Wait
Cross	−100, −100	1, 0
Wait	0, 1	0, 0

- Social cost function is expectancy of sum of payoffs of both players.
- Price of stability: 1

A real world example: Street crossing

P1 \ P2	Cross	Wait
Cross	-100 / -100	1 / 0
Wait	0 / 1	0 / 0

- Social cost function is expectancy of sum of payoffs of both players.
- Price of stability: 1
- Price of anarchy:

A real world example: Street crossing

P1 \ P2	Cross	Wait
Cross	-100 / -100	0 / 1
Wait	1 / 0	0 / 0

- Social cost function is expectancy of sum of payoffs of both players.
- Price of stability: 1
- Price of anarchy:

$$\frac{2\frac{1}{101}\frac{100}{101} - 100(\frac{1}{101})^2 \approx 0.01}{1}$$

Our tool: Potential function method

In general, the potential function method is the following:

- Suppose we want prove some property of some implicitly given subset E of a set S (for example, it is nonempty).
- Define a potential function $\phi : S \to \mathbb{R}$ such that E are exactly the (global) minima of ϕ.
- Since ϕ has a global minimum, E is non-empty.
- Algorithmically, this is also useful since an element of E can be found by minimizing ϕ (but in chapters 17,18 and 19 of the book people don't care).

Nonatomic Selfish routing

- A directed graph $G = (V, E)$, and a source-sink pair (s_i, t_i) for every player (commodity) i.
- A requirement vector $r \in \mathbb{R}^k$, where r_i represents the traffic of commodity i, and a nondecreasing, continuous cost function $c_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$.
- Let $\mathcal{P}_i$ be the set of all paths from s_i to t_i, and $\mathcal{P} = \cup_{i=1}^k \mathcal{P}_i$.
- A flow f is a non-negative vector indexed by $\mathcal{P}$. f is feasible for i if $\sum_{P \in P_i} f_P \leq r_i$.
- Let $f = (f^1, \ldots, f^k)$ be a strategy vector, and f_e be the total amount of flow of f on e; the cost function c_i of player i and social function $\sigma : S \rightarrow \mathbb{R}^+$ are defined as

$$c_i(f) = \sum_{P \in \mathcal{P}_i} \sum_{e \in P} c_e(f_e) f_e^i \quad \text{and} \quad \sigma(f) = \sum_{i=1}^k c_i(f^i) = \sum_{e \in \mathcal{P}} c_e(f_e) f_e$$

For a path P, we shorthand $c_P(f) = \sum_{e \in P} c_e(f_P)$.

Observation

A strategy vector $f = (f_1, \ldots, f_k)$ is an equilibrium if and only if for every commodity i and every pair $P, \tilde{P} \in \mathcal{P}_i$ with $f_P > 0$

$$c_P(f) \leq c_{\tilde{P}}(f)$$

Proof.

The right to left direction follows from definition of equilibrium. For the other direction, note that since c_P is "nice", rerouting any amount of flow from P to $\tilde{P}$ decreases the costs if the inequality does not hold. □

UNIVERSITETET I BERGEN
Institutt for informatikk

Theorem (Beckmann et al. (1956))

A nonatomic selfish routing game admits at least one equilibrium flow, and if f and $\tilde{f}$ are equilibrium flows, $c_e(f) = c_e(\tilde{f})$ for every edge e.

Taking the observation into consideration, it is natural to ask whether equilibria and social optima are the same, but:

Example (Pigou (1920))

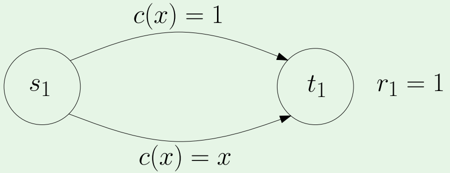

Theorem (Beckmann et al. (1956))

A nonatomic selfish routing game admits at least one equilibrium flow, and if f and $\tilde{f}$ are equilibrium flows, $c_e(f) = c_e(\tilde{f})$ for every edge e.

Taking the observation into consideration, it is natural to ask whether equilibria and social optima are the same, but:

Example (Pigou (1920))

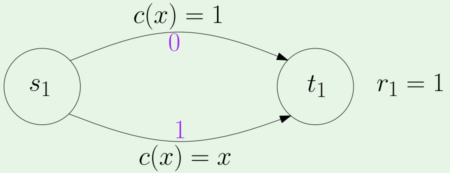

Theorem (Beckmann et al. (1956))

A nonatomic selfish routing game admits at least one equilibrium flow, and if f and $\tilde{f}$ are equilibrium flows, $c_e(f) = c_e(\tilde{f})$ for every edge e.

Taking the observation into consideration, it is natural to ask whether equilibria and social optima are the same, but:

Example (Pigou (1920))

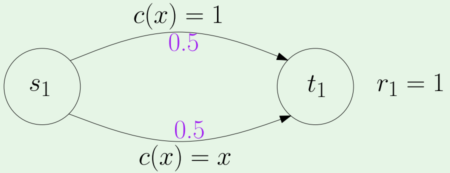

But, it appears that if we modify the cost function to $c^*(x) = (x \cdot c(x))'$, the social optima of the old came are exactly the equilibria of the new game

But, it appears that if we modify the cost function to $c^*(x) = (x \cdot c(x))'$, the social optima of the old came are exactly the equilibria of the new game:

Example

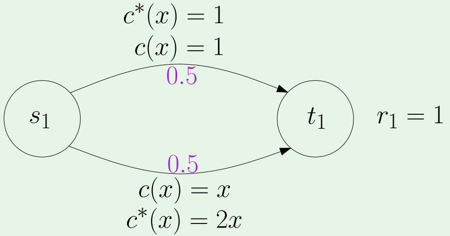

Lemma

Let (G, r, c) be instance of nonatomic selfish routing, where c is a nondecreasing, continues (differentiable) function. Than f^* is an optimal flow if and only if it is an equilibrium in the instance (G, r, c^*).

Proof.

- f^* is optimal iff for every i and $P, \tilde{P} \in \mathcal{P}_i$

$$\sum_{e \in P} (f_e^* \cdot c_e(f_e^*))' \leq \sum_{e \in \tilde{P}} (f_e^* \cdot c_e(f_e^*))'$$

- now the equivalence follows from the previous observation.

□

Proof of Theorem

Now, we define a potential function ϕ such that $\phi* = c$ and hence all minimizers of ϕ are exactly the equilibria of the instance (G, r, c).

$$\phi(f_e) = \frac{1}{f_e} \int_0^{f_e} c_e(x)dx$$

Then $\phi * (f_e) = (f_e\phi(f_e))' = c_e(f_e)$.

Proof of Theorem

Now, we define a potential function ϕ such that $\phi* = c$ and hence all minimizers of ϕ are exactly the equilibria of the instance (G, r, c).

$$\phi(f_e) = \frac{1}{f_e} \int_0^{f_e} c_e(x) dx$$

Then $\phi * (f_e) = (f_e \phi(f_e))' = c_e(f_e)$.
The second part of the theorem, follows from the convexity of the cost function and the set of feasible flows.

Proof of Theorem

Now, we define a potential function ϕ such that $\phi* = c$ and hence all minimizers of ϕ are exactly the equilibria of the instance (G, r, c).

$$\phi(f_e) = \frac{1}{f_e} \int_0^{f_e} c_e(x) dx$$

Then $\phi * (f_e) = (f_e \phi(f_e))' = c_e(f_e)$.
The second part of the theorem, follows from the convexity of the cost function and the set of feasible flows.
The proof can probably be used to obtain an equilibrium in polynomial time using convex programming.

Braess's paradox

Example

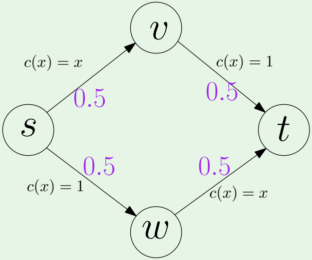

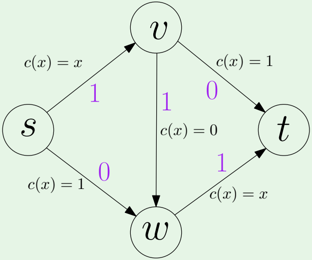

$t_1 = 1$

Congestion game

Definition

A congestion game is a game k players, a ground set of resources R, a cost function $c_r : \{1, \dots, k\} \to \mathbb{R}$ for each $r \in R$ and each player has a strategy set $\mathcal{S}_i \subseteq R$. In a strategy profile $S = (s_1, \dots, s_k)$, the cost of a player is defined as $c^i(S) = \sum_{r \in s_i} = c_r(n_r)$, where n_r is the number of strategies containing r.

Theorem (Rosenthal (1973), not in the book.)

Every congestion game has at least one pure equilibrium.

Proof.

For strategy profile $S = (s_1, \ldots, s_k)$, denote (S_{-i}, s') for the strategy vector $(s_1, \ldots, s_{i-1}, s', s_{i+1}, \ldots, s_k)$. Define $\phi : \mathcal{S} \rightarrow \mathbb{R}^+$:

$$\phi(s_1, \ldots, s_k) = \sum_{r \in R} \sum_{i=1}^{n_r} c_r i$$

then

$$c^i((S_{-i}, s')) - c(S) = \sum_{r \in s' \setminus s_i} c_r(n_r + 1) - \sum_{r \in s_i \setminus s'} c_r(n_r)$$
$$= \phi((S_{-i}, s')) - \phi(S)$$

□

Atomic selfish routing

- atomic selfish routing game is atomic selfish routing restricted to integral flows.
- note that, if the players are restricted to $0-1$ flows, the existence directly follows from Rosenthals' Theorem. However:

Example

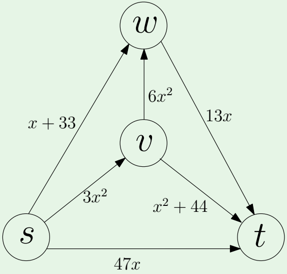

$r_1 = 1$
$r_2 = 2$

Example

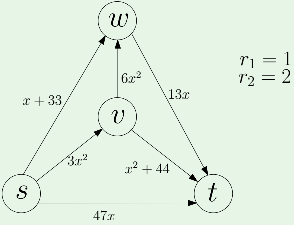

Let P_1, P_2, P_3, P_4 be respectively $st, sv, swt, svwt$ then

Example

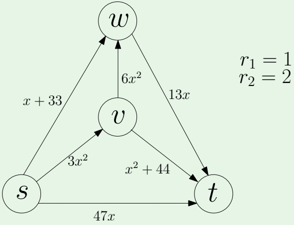

Let P_1, P_2, P_3, P_4 be respectively $st, sv, swt, svwt$ then

- player 2 P_1 or P_2 $\longrightarrow$ player 1 P_4

Example

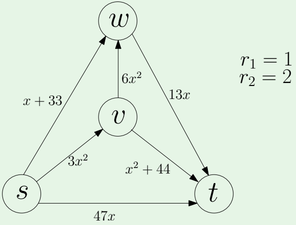

Let P_1, P_2, P_3, P_4 be respectively $st, sv, swt, svwt$ then

- player 2 P_1 or P_2 $\longrightarrow$ player 1 P_4
- player 2 P_3 or P_4 $\longrightarrow$ player 1 P_1

Example

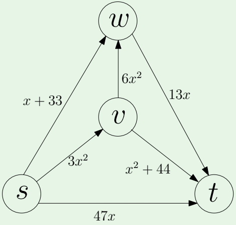

$r_1 = 1$
$r_2 = 2$

Let P_1, P_2, P_3, P_4 be respectively $st, sv, swt, svwt$ then

- player 2 P_1 or P_2 $\longrightarrow$ player 1 P_4
- player 2 P_3 or P_4 $\longrightarrow$ player 1 P_1
- player 1 P_4 $\longrightarrow$ player 2 P_3

Example

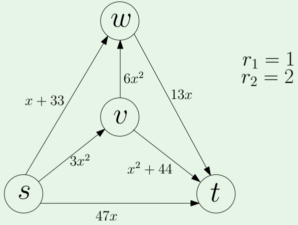

Let P_1, P_2, P_3, P_4 be respectively $st, sv, swt, svwt$ then

- player 2 P_1 or P_2 $\longrightarrow$ player 1 P_4
- player 2 P_3 or P_4 $\longrightarrow$ player 1 P_1
- player 1 P_4 $\longrightarrow$ player 2 P_3
- player 1 P_1 $\longrightarrow$ player 2 P_2

Cliffhanger

Next time:

- Bounds on the price of anergy.
- Network formation: The same as selfish routing, but with decreasing cost functions.

Thanks for attending!!!

Any questions?