Selfish routing - infinitely divisible flow model

December 19, 2006

Networks

Model: Tim Roughgarden and Eva Tardos

$G = (V, E)$ a finite directed graph.

$s_i, t_i, (i = 1, \ldots, k)$ couples of source - sink (destination) vertices (commodities).

$\mathbb{P}_i$ the set of all (simple) paths from s_i to t_i.

$\mathbb{P} = \cup_{i=1}^{k} \mathbb{P}_i$ - the set of all paths between a source and a corresponding destination.

Flows

$f : \mathbb{P} \longrightarrow \mathbb{R}^+$ the flow mapping.

f_P denotes the flow through a path $P \in \mathbb{P}$.

$f_e = \sum_{\{P|e \in P\}} f_P$ - the flow through an edge e

r_i - a finite and positive *traffic rate* from s_i to t_i, $i = 1, ..., k$.

A flow f is feasible if

$$\sum_{P \in \mathbb{P}_i} f_P = r_i$$

for all i.

In the sequel we are mostly (only?) interested in feasible flows.

Latency (cost) mappings

A latency or cost mapping for an edge e:

$$\ell_e : \mathbb{R}^+ \longrightarrow \mathbb{R}^+$$

We assume that such a mapping ℓ_e is **continuous** and **non-decreasing**.

Intuitively, $\ell_e(f_e)$ gives the delay over the edge e if the flow going through e is equal to f_e. Thus latency depends on congestion.

$$\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$$

- the latency/cost of a path P with respect to a flow f.

An instance:

$$(G, r, \ell).$$

The **cost** of the flow f:

$$C(f) = \sum_{P \in \mathbb{P}} \ell_P(f) \cdot f_P = \sum_{e \in E} \ell_e(f_e) \cdot f_e$$

$$\sum_{P \in \mathbb{P}} \ell_P(f) f_P = \sum_{P \in \mathbb{P}} (\sum_{e \in P} \ell_e(f_e)) f_P =$$

$$\sum_{e \in E} (\sum_{\{P \in \mathbb{P} | e \in P\}} f_P) \ell_e(f_e) = \sum_{e \in E} \ell_e(f_e) \cdot f_e$$

Definition. *Given an instance* (G, r, ℓ) *the flow minimizing* $C(f)$ *is called optimal.*

The optimal flow always exists: the space of feasible flows is compact and the cost function is continuous.

Flows at Nash equilibrium

Definition. *A feasible flow* f *in* (G, r, ℓ) *is at Nash equilibrium (is a Nash flow) if for all commodities* $i \in \{1, \ldots, k\}$ *and for all paths* $P_1, P_2 \in \mathbb{P}_i$ *with* $f_{P_1} > 0$*, for all amounts* $\delta \in (0; f_{P_1}]$ *of traffic on* P_1

$$\ell_{P_1}(f) \leq \ell_{P_2}(\overline{f})$$

where

$$\overline{f}_P = \begin{cases} f_P - \delta & \text{if } P = P_1 \\ f_P + \delta & \text{if } P = P_2 \\ f_P & \text{otherwise} \end{cases}$$

Intuitively, if a player controls the amount δ of the flow going through the path P_1 then his flow suffers the latency $\ell_{P_1}(f)$. If he redirects this flow to the path P_2 then $\overline{f}$ will be the new flow and he will suffer the latency of $\ell_{P_2}(\overline{f})$ on P_2.

Thus at a Nash equilibrium there is no incentive to redirect any part of the flow.

Lemma. *A flow f feasible for the instance (G, r, ℓ) is at Nash equilibrium if for every commodity i and all $P_1, P_2 \in \mathbb{P}_i$ with $f_{P_1} > 0$*

$$\ell_{P_1}(f) \leq \ell_{P_2}(f)$$

Proof. By continuity and monotonicity of ℓ_e. $\square$

Pigou's example

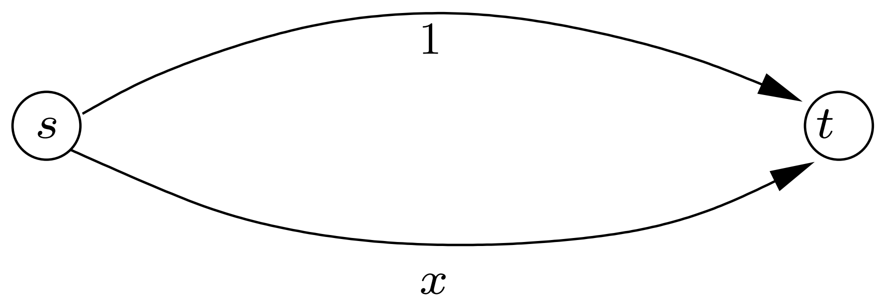

Figure 1: The latency of upper edge is $\ell(x) = 1$, the latency of lower edge is $\ell(x) = x$. The traffic rate between s and t is 1.

In this example the Nash equilibrium f is attained only if all flow goes through the upper edge, the cost is $C(f) = 1$.

If the flow a goes through the the lower edge and $1 - a$ through the upper one then the cost is $C(f_a) = a^2 + 1 - a$ and the minimum is attained for

$a = \frac{1}{2}$, thus the cost of the optimal flow f^* is $C(f^*) = \frac{3}{4}$.

Braess's paradox

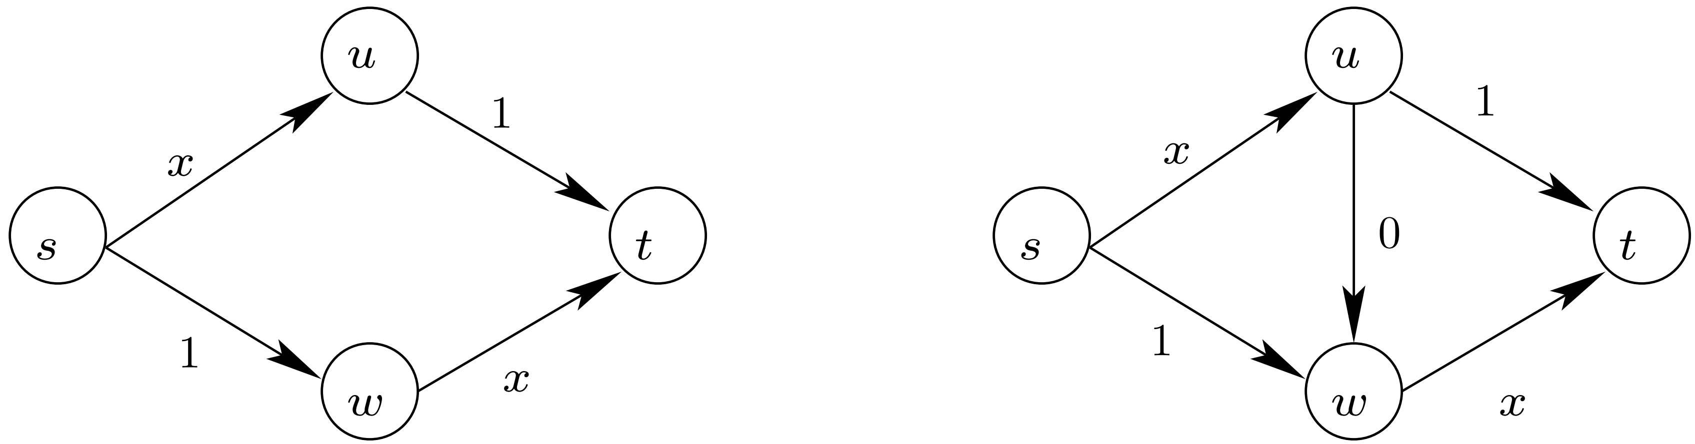

(a) (b)

Figure 2: The traffic rate is 1. Latency functions are the same as in the previous example. In (b) the latency on (u, w) edge is 0.

On figure (a) both Nash and optimal flow pass $\frac{1}{2}$ of the traffic by each path. The cost is then $\frac{3}{2}$.

On (b) the new edge (u, w) has latency 0. Now the Nash flow routes all the traffic by $s \to u \to w \to t$ with the cost 2. Adding a new low latency edge results in higher cost of Nash equilibrium! What is the cost of the optimal flow in (b)?

The price of anarchy

Fix an instance (G, r, ℓ). Suppose that f is the **worst** Nash equilibrium flow (worst case equilibria – Papadimitriou), i.e. the Nash equilibrium with the highest cost $C(f)$. Suppose that f^* is an optimal flow.

The price of anarchy of (G, r, ℓ) is defined as

$$\rho(G, r, \ell) = \frac{C(f)}{C(f^*)}$$

where f^* is an optimal flow and f a flow at Nash equilibrium.

If $C(f^*) = 0$ then f^* is also a Nash equilibrium and we set $\rho(G, r, \ell) = 1$.

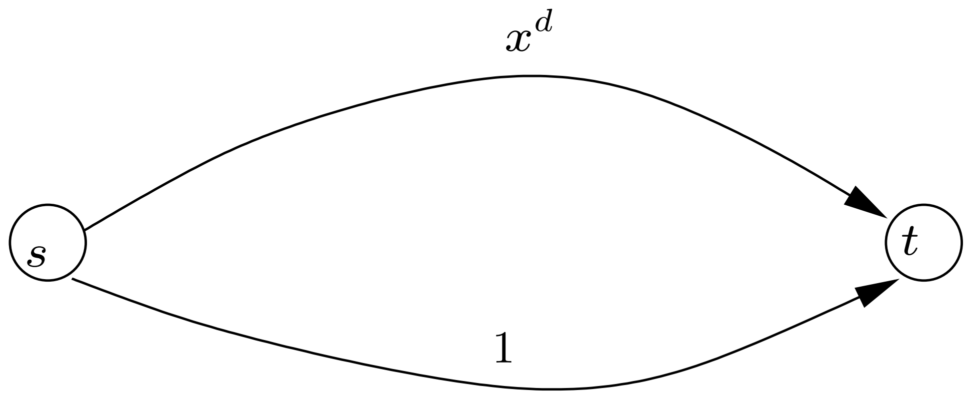

Figure 3: The price of anarchy can be arbitrarily high. Traffic rate is 1. The figure shows latency mappings for both edges. In Nash equilibrium all flow goes through the upper edge. Calculate the optimal flow.

Convex optimization problem

It turns out that the problem of finding the optimal flow and the problem of finding a Nash flow are related to the same non-linear programming problem:

$$\text{minimize} \sum_{e \in E} h_e(f)$$

subject to:

$$\sum_{P \in \mathbb{P}_i} f_P = r_i, \quad \forall i \in \{1, \ldots, k\}$$

$$f_e = \sum_{\{P \in \mathbb{P}_i | e \in P\}} f_P, \quad \forall e \in E \tag{1}$$

$$f_P \geq 0, \quad \forall P \in \mathbb{P},$$

We can solve this problem if h_e are convex functions.

r_i are constants fixed by the instance (G, r, ℓ), f_e and f_P are variables variables.

A *feasible solution* of (1) is any flow f satisfying all the constraints of (1).

A solution of (1) is any feasible solution f minimizing the objective function.

Convex functions

A subset A of $\mathbb{R}^n$ is convex if for all $x, y \in A$ and $0 \leq \alpha \leq 1$,

$$\alpha x + (1 - \alpha) y \in A.$$

A function $g : \mathbb{R}^n \longrightarrow \mathbb{R}$ is convex if for all $x_1, x_2 \in \mathbb{R}^n$

$$g(\alpha x_1 + \beta x_2) \leq \alpha g(x_1) + \beta g(x_2)$$

whenever $\alpha, \beta \geq 0$ and $\alpha + \beta = 1$.

If g_e are convex functions then $\sum_{e \in E} g_e(f)$ is convex.

If g is a convex function and $\alpha > 0$ then αg is a convex function.

Lemma 1. *Let f be convex differentiable function. Then* $(x-a) \cdot f'(a) \leq f(x) - f(a)$.

Proof.

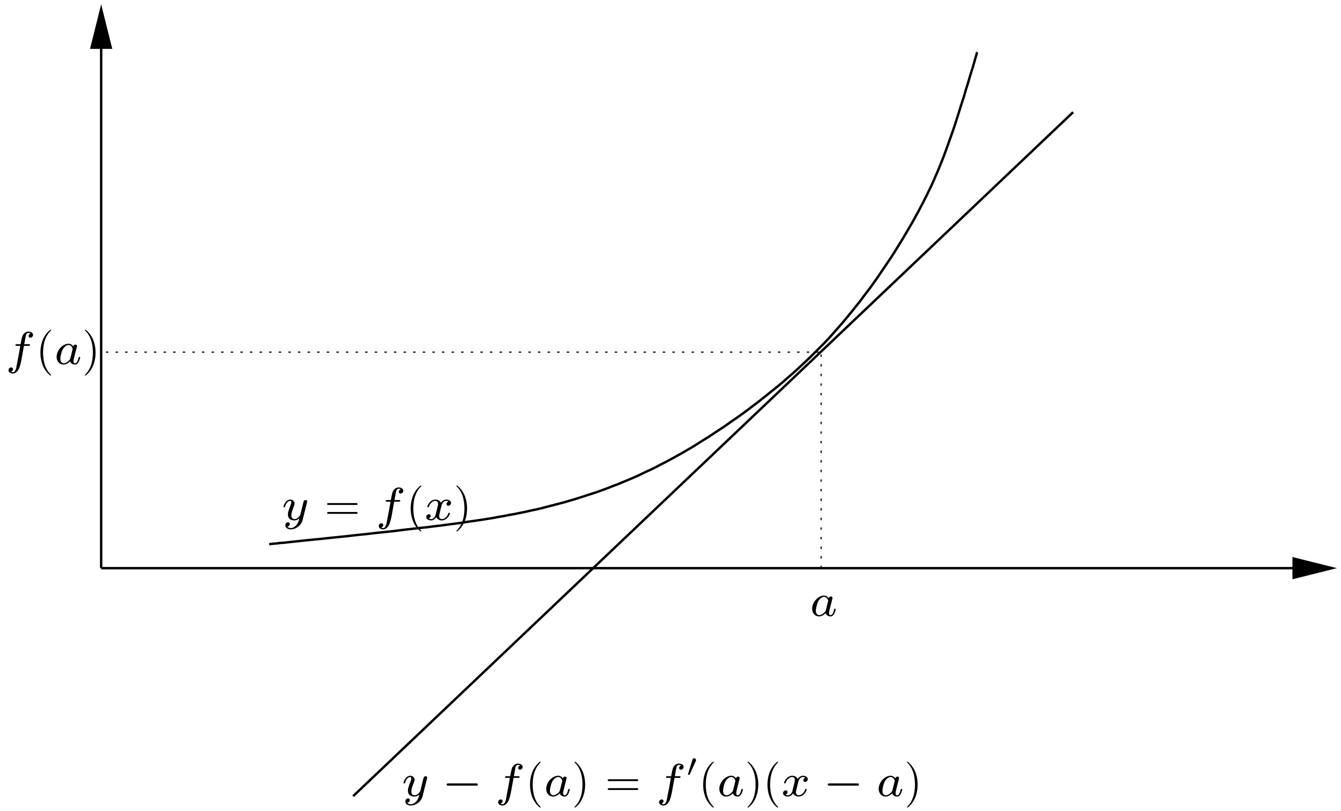

The tangent line to f at the point $(a, f(a))$: $y - f(a) = f'(a) \cdot (x - a)$ lies below the graph of f. This yields the thesis. $\square$

Lemma. *If f is a convex function over a convex domain $A \subset \mathbb{R}^n$ and $a \in A$ is a local minimum of f then a is also a global minimum.*

If f is strictly convex then there is only one global minimum.

Notation: $h'_e(x) = \frac{d}{dx}h_e(x)$ and $h'_P(x) = \sum_{e \in P} h'_e(x)$.

Theorem. *Let f^* be a feasible solution to (1), where all functions h_e are continuously differentiable and convex. Then the following are equivalent:*

(a) f^* is an optimal solution of (1),

(b) for every i, $1 \le i \le k$, and $P_1, P_2 \in \mathbb{P}_i$ with $f^*_{P_1} > 0$,

$$h'_{P_1}(f^*) \le h'_{P_2}(f^*)$$

(c) for every feasible flow f

$$\sum_{P \in \mathbb{P}} h'_P(f^*) f^*_P \le \sum_{P \in \mathbb{P}} h'_P(f^*) f_P$$

(d) for every feasible flow f,

$$\sum_{e \in E} h'_e(f^*_e) f^*_e \leq \sum_{e \in E} h'_e(f^*_e) f_e$$

Remark In the preceding theorem

1. setting for all $e \in E$
$$h_e(f_e) = \ell_e(f_e) f_e$$
in (a) we obtain the optimal flow problem.

2. taking $h_e(f_e)$ such that
$$h'_e(f_e) = \ell_e(f_e)$$
then in (b) w obtain the condition of a Nash equilibrium.

Corollary. *Finding an optimal flow for an instance* (G, r, ℓ) *is equivalent to finding a Nash flow for an instance* (G, r, ℓ^*) *where*

$$\ell^*_e(f_e) = \frac{d}{dx}(\ell_e(x)x) = \ell_e(x) + x\ell'_e(x)$$

Optimal equilibria

Latency functions ℓ_e are *semiconvex* if ℓ_e are differentiable and $\ell_e^*(x) = x\ell_e(x)$ are convex on $[0; \infty)$.

Existence and uniqueness of a Nash equilibrium

If $\ell_e(x)$ are continuous and nondecreasing then $h_e(x) = \int_0^x \ell_e(t)dt$ is convex and there exists a Nash equilibrium (convex programming problem on bounded convex domain has a solution).

Unicity
If f and $\widetilde{f}$ are Nash flows for (G, r, ℓ) then $\ell_e(f_e) = \ell_e(\widetilde{f}_e)$ for each edge e.

(Since f and $\widetilde{f}$ are solutions for a convex programming problem all linear combinations $\alpha f + (1-\alpha)\widetilde{f}$ are also optimal and the objective function is therefore constant for all these combinations. This is only possible if all $h_e(x)$ are linear between f_e and $\widetilde{f}_e$. Thus $\ell_e(x) = \frac{d}{dx}h_e(x)$ are constant.)

If f and $\widetilde{f}$ are Nash flows for (G, r, ℓ) and $\ell_e(x)$ is strictly increasing then $f_e = \widetilde{f}_e$.