Game Theory and Algorithms*
Lecture 18: Routing Games and Potentials

May 12, 2011

Summary: We show that potential functions can also be used to analyze the traffic *routing* game. This brings us to a discussion of convexity, and sufficient conditions for the Nash equilibrium to be unique. Then, we show *marginal costs*, a sort of inverse of the potential function, which can be used as taxes to induce greedy users to route themselves optimally.

1 Price of Stability

Recall the traffic routing game: we are given a directed graph with a source, a sink, and $\text{delay}_e : [0,1] \to \mathbb{R}_+$ for all edges e. A flow f is Nash if every path used is at least as fast as every alternate path, whereas the social cost of the flow equals the "sum" $\sum_e f_e \text{delay}_e(f_e)$ of all delays.

In the previous class, we saw a potential function for the discrete network formation game, where decreases in the potential function correspond to decreases in cost due to a unilateral change. This idea generalizes in a straightforward way to this continuous routing game: assume delays are integrable (e.g., increasing) and define

$$\Phi(f) := \sum_e \int_{u=0}^{f_e} \text{delay}_e(z) dz.$$

The interpretation is that if we move an ϵ amount of flow from path P to path P', the decrease in delay from the old path to the new path equals the decrease in the potential function. So analogously to last class[1],

$$f \text{ is Nash} \Leftrightarrow f \text{ is a local minimum of } \Phi.$$

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

[1] Technically we showed this for "component-wise local minima" but it is not hard to show every component-wise local minimum is a local minimum.

Moreover, the social cost for a flow is

$$\text{social-cost}(f) := \sum_e f_e \text{delay}_e(f_e).$$

Analogously to last class, if $C \cdot \text{social-cost}(f) \le \Phi(f) \le C' \cdot \text{social-cost}(f)$ for all flows f, then the Price of Stability (ratio of "socially best NE" cost to optimal social cost) is at most C'/C.

Exercise. Show that if all edge delay functions are nondecreasing nonnegative affine functions, then the Price of Stability is at most 2. What bound can you get if each edge delay function is a quadratic of the form $z \mapsto a_e z^2 + b_e z + c_e$ for $a_e, b_e, c_e \ge 0$?

1.1 Anarchy is Stability! (If you're unique)

Assuming nondecreasing delays, the routing game turns out to have (essentially) only one Nash equilibrium. In this case, the worst NE is also the best NE, which means that the prices of anarchy and stability are the same.

We need to cover a little bit of background material on convex functions. A function $g(z)$ of one variable is *strictly convex* if

$$g(\lambda z + (1-\lambda)z') < \lambda g(z) + (1-\lambda)g'(z)$$

whenever $z \ne z'$ and $0 < \lambda < 1$. Notice that if $z = z'$ or $\lambda \in \{0, 1\}$ then the inequality holds with equality. The definition of convex functions has a nice geometric interpretation: a function is convex if and only if a chord of the graph always lies strictly above the graph.

Next, we need the following fact from analysis:

Fact 1. *A differentiable function is (strictly) convex if and only if its derivative is (strictly) increasing.*

In other words, the second derivative (assuming it exists) is positive.

This gets us to the main result of this section:

Theorem 2 (Uniqueness of NE)**.** *Assume each edge delay function is strictly increasing. Then any two Nash equilibria f, f' have $f_e = f'_e$ for all edges e.*

To simplify notation, let $\text{Delay}_e(z)$ denote $\int_{u=0}^{z} \text{delay}_e(u)du$, so $\Phi(f) = \sum_e \text{Delay}_e(f_e)$.

Proof. The crux is that the potential $\Phi : [0,1]^E \to \mathbb{R}_+$, expressed as a function of the edge loads, is a strictly convex (multivariable) function. This means that for any two distinct flows f, f' and any $0 < \lambda < 1$,

$$\Phi(\lambda f + (1-\lambda)f') < \lambda\Phi(f) + (1-\lambda)\Phi(f')$$

where $\lambda f + (1-\lambda) f'$ is a vector sum, i.e. a weighted-average flow "between" flows f and f'. To prove this in detail, we expand the definition of Φ: we want to prove

$$\sum_e \text{Delay}_e(\lambda f_e + (1-\lambda) f'_e) < \sum_e \lambda \text{Delay}_e(f_e) + (1-\lambda)\text{Delay}_e(f'_e)$$

but this holds since each Delay_e is a strictly convex univariate function and at least one $f_e \neq f'_e$.

Now that we have shown Φ is convex, we claim it has a unique local minimum. Indeed, if there were two distinct local minima f and f', without loss of generality $\Phi(f) \leq \Phi(f')$, the flows "between" them would contradict the local-minimality of f'.

Since there is only one local minimum, and Nash equilibria are the same as local minima of Φ, we are done. □

Aside: since there is only one local minimum, it is the global minimum.

Remark: Under the weaker assumption that the edge delay functions are nondecreasing, there might be more than one Nash equilibrium, but we get the weaker theorem (Theorem 18.8 in Nisan et al.) that $\text{delay}_e(f_e) = \text{delay}_e(f'_e)$ which (with a little work) implies all Nash equilibria have the same cost. So again, the prices of anarchy and stability are the same.

Another useful fact about convex functions is that there are efficient algorithms to compute their minima. Therefore, we can efficiently compute a Nash equilibrium for this game.

2 Potentials Backwards = Marginal Costs

Nash equilibria are, by definition, local characterizations. Potential functions give a characterization of Nash equilibria in terms of global minima, provided that the potential function is convex. For *social cost minimizers* in place of *Nash equilibria*, we can do the reverse operation: instead of the natural global-minimizer definition of social cost, we can write it as a local-minimizer.

For this to work, we need the social cost function to be convex. So for the rest of this section we assume:

Assumption 1. *For each edge e, the function $z \cdot \text{delay}_e(z)$ is convex.*

This assumption has the following effects, similar to what we saw previously:

1. The function social-cost(f) is a convex function (of the f_e variables),
2. a socially-optimal flow can be found in polynomial time, and
3. the socially-optimal flows are a convex set (it is unique if the convexity is strict).

To proceed, now let us compare social-cost to Φ. Recall

$$\text{social-cost}(f) = \sum_e f_e \cdot \text{delay}_e(f_e) \qquad \text{and} \qquad \Phi(f) = \sum_e \int_{z=0}^{f_e} \text{delay}_e(z) dz.$$

Also recall that minimizers of Φ (Nash equilibria) are characterized by

$$\sum_{e \in P} \text{delay}_e(f_e) \leq \sum_{e \in P'} \text{delay}_e(f_e) \tag{1}$$

for all paths P used by f and all other paths P'. Just as delay is the derivative of $\int$ delay, let the *marginal social cost* $\texttt{msc}_e(f_e)$ be the derivative of $f_e \cdot \text{delay}_e(f_e)$ with respect to f_e:

$$\texttt{msc}_e(z) := \frac{d}{dz}(z \cdot \text{delay}_e(z)) = z \cdot \text{delay}'_e(z) + \text{delay}_e(z).$$

Then we would expect by analogy that social optima are characterized by:

Theorem 3 (Local characterization of social optima)**.** *Assume each function $z \mapsto z \cdot \text{delay}_e(z)$ is convex and differentiable. Then f is a socially optimal flow if and only if*

$$\sum_{e \in P} \texttt{msc}_e(f_e) \leq \sum_{e \in P'} \texttt{msc}_e(f_e)$$

for all paths P used by f and all other paths P'.

This is really a theorem! I.e., the analogy's final destination is correct.

	Nash Equilibrium	Social Optimum
Local-min of	delay, by definition	msc (assuming msc increasing)
Global-min of	Φ (assuming Delay convex)	social-cost, by definition

Remark: You can run the analogy in reverse, starting from first-order convex optimization conditions and then deriving what the potential function should look like. See §18.3.1 of Nisan et al.

2.1 Application: Taxes to Induce Optimality

Suppose we are allowed to increase some of the edge costs by adding *taxes* to them. If we are allowed to have the taxes depend on load of an edge, then we basically can replace the original delay functions with whatever we like. So we focus on the more interesting case that we are allowed to add only a *constant* tax to each delay function: we replace the delays with

$$\widetilde{\text{delay}}_e(z) = \text{delay}_e(z) + t_e$$

where t_e is some constant for each edge e.

The surprising application of marginal social costs is that we can coerce greedy users into forming a socially-optimal flow.

Theorem 4. *In a network flow game (where* msc *exists and is increasing), let f^* be a socially optimal flow. Define the tax e for each edge by $t_e = f^*_e \cdot \text{delay}'_e(f^*_e)$. Then in the new instance with taxes, f^* is a Nash flow.*

Proof. By Theorem 3,

$$\sum_{e\in P} \texttt{msc}_e(f_e^*) \le \sum_{e\in P'} \texttt{msc}_e(f_e^*)$$

for all paths P used by f^* and all other paths P'. Let us expand the definition of $\texttt{msc}$:

$$\texttt{msc}_e(z) = \frac{d}{dz} z \cdot \mathrm{delay}_e(z) = \mathrm{delay}_e(z) + z \cdot \mathrm{delay}_e'(z).$$

So,

$$\sum_{e\in P} \underbrace{\mathrm{delay}_e(f_e^*) + f_e^* \cdot \mathrm{delay}_e'(f_e^*)}_{\widetilde{\mathrm{delay}}_e(f_e^*)} \le \sum_{e\in P'} \underbrace{\mathrm{delay}_e(f_e^*) + f_e^* \cdot \mathrm{delay}_e'(f_e^*)}_{\widetilde{\mathrm{delay}}_e(f_e^*)}$$

for all paths P used by f^* and all other paths P'. This (under the horizontal brackets) shows f^* satisfies the characterization (1) of Nash equilibria for the taxed delays. □

Q: Is f^* also socially optimal for the taxed delays?

Exercise. The *unfairness* of a flow f is the ratio

$$\max\{\sum_{e\in P} \mathrm{delay}_e(f_e) \mid f_P > 0\} / \min\{\sum_{e\in P} \mathrm{delay}_e(f_e) \mid f_P > 0\}.$$

What is the unfairness for a Nash flow? Show that for affine nonnegative nondecreasing delay functions, the unfairness of any socially optimal flow is at most 2.

3 Wrap-up

Exercise. In this exercise, you will prove the stronger bound of 4/3 on the price of anarchy for nonnegative nondecreasing affine delay functions.

1. Let f^* be a flow. Pin the edge delays at constants $\mathrm{delay}_e(f_e^*)$. Show that f^* was an equilibrium flow for the original delays iff it is an optimal flow for the new constant delays. In other words, show f^* is Nash iff for all other flows f,

$$\sum_e f_e^* \mathrm{delay}_e(f_e^*) \le \sum_e f_e \mathrm{delay}_e(f_e^*).$$

2. Show that for any nonnegative nondecreasing affine delay function $d(z)$,

$$rd(r) \le \frac{4}{3}(xd(x) + (r-x)d(r))$$

for $0 \le r, x \le 1$.

3. Show that for these delay functions, the price of anarchy is at most 4/3.

Exercise. Consider any class of nonnegative nondecreasing delay functions which include all of the nonnegative constant cost functions. Let α be minimal such that for all functions d from this class and all $0 \le r, x \le 1$,

$$rd(r) \le \alpha(xd(x) + (r - x)d(r)).$$

By the previous exercise, the price of anarchy for instances with this class of delay functions is at most α. Show that in fact the price of anarchy is *exactly* α.

One other very interesting fact about this network flow game has to do with *resource augmentation*. Namely, no matter how bad the price of anarchy is for a collection of delay functions (it can be arbitrarily bad, due to Pigou's example), the bad cost due to greedy users can be offset by scaling up all edge capacities by a function of 2. Formally,

Theorem 5. *Assume delay functions are nondecreasing; define* $\widetilde{\mathrm{delay}}_e(z) = \mathrm{delay}_e(z/2)$. *Then the worst social cost of any Nash equilibrium for the new delays does not exceed the optimum social cost for the original delays.*

This is Theorem 18.29 in Nisan et al.