Graphs and Networks	Lecture 15
The Price of Anarchy	
Daniel A. Spielman	October 17, 2013

15.1 Disclaimer

These notes are not necessarily an accurate representation of what happened in class. They are a combination of what I intended to say with what I think I said. They have not been carefully edited.

15.2 Overview

We state more precisely exactly what we mean by social optima and Nash equilibria. We then show that Pigou's and Braess's examples are essentially worst cases: with linear cost functions the cost of counterintuitive effects is at most 4/3. This material comes from the work of Roughgarden and Tardos [RT02].

15.3 Road Networks

We will consider traffic flow in an idealization of road networks. We view a road network as a directed graph. In our model, the time that it takes to traverse a road will be a function of the number of cars using that road. The more cars, the higher the traffic. In this lecture, we will assume that the time it takes to traverse a road is a linear function of the form

$$af + b,$$

where f is the traffic on the road, and $a, b \geq 0$.

In fact, it suffices to just consider roads with cost functions like ax or b, as we can make a more general road by just putting these two end-to-end.

We view a road network as a directed graph $G = (V, E)$. For each edge e we define the cost function of road e to be $c_e(f) = a_e f + b_e$. Our road network will have two distinguished vertices, s and t, and we assume that all traffic wants to go from node s to node t. We assume that the total traffic is 1. Traffic must be conserved (no crashes), so the amount of traffic entering each node must equal the amount leaving. A traffic flow is any function $f : E \to \mathbb{R}^{\geq 0}$ that obeys this conservation condition and sends one unit from s to t. Any given car will take one particular path P from s to t. The time taken by such a car will be the sum of the cost of all the edges on its path:

$$c(P) \stackrel{\text{def}}{=} \sum_{e \in P} c_e(f_e),$$

where f_e is the flow on edge e.

15.4 Selfish (Nash) Routing

Assuming that everyone is greedy, everyone will try to take the path of least cost. Of course, this might make that path slower. A natural question is what routing will result if everyone tries to optimize for themselves. Such a routing is called a Nash equilibrium. In this situation the Nash equilibrium is unique. But, we don't know that yet.

Let the set of s-t paths be $P_1, \ldots, P_k$, and let α_i be the fraction of the traffic that flows on path P_i. In the Nash equilibrium, no car will go along a sub-optimal path. Assuming that each car has a negligible impact on the traffic flow, this means that every path P_i that has non-zero flow must have minimal cost. That is, for all i such that $\alpha_i > 0$ and all j

$$c(P_i) \leq c(P_j).$$

An equivalent formulation is

$$c(P_j) > c(P_i) \quad \Longrightarrow \quad \alpha_j = 0.$$

15.5 Flows and Paths

There are two natural ways to describe s-t flows: either by a function f that gives the amount of flow on every edge in the network, or by giving the fractions $\alpha_1, \ldots, \alpha_k$ of flow that traverse each s-t path in the network. We will switch between these two perspectives often in the lecture. Given $\alpha_1, \ldots, \alpha_k$, the flow is determined. However, there may be many ways of decomposing one flow into fractions along each path. But, it will make no difference for our purposes.

15.6 Social optimum

Society in general cares more about the average time its takes to get from s to t. If we have a flow that makes this average time low, everyone could rotate through all the routes and decrease the total time that they spend in traffic. So, the social cost of the flow f is

$$\begin{aligned} c(\alpha_1, \ldots, \alpha_k) &\stackrel{\text{def}}{=} \\ \sum_i \alpha_i c(P_i) &= \sum_i \alpha_i \sum_{e \in P_i} c_e(f_e) \\ &= \sum_e c_e(f_e) \sum_{i: e \in P_i} \alpha_i \\ &= \sum_e c_e(f_e) f_e. \end{aligned}$$

I recall that a function $c(x)$ is *convex* if for all x_1 and x_2, and all $0 < t < 1$

$$c(tx_1 + (1-t)x_2) \leq tc(x_1) + (1-t)c(x_2).$$

For example, a linear function of the form

$$c(x) = bx$$

is convex. A function $c(x)$ is *strongly convex* if the inequality above is strict whenever $x_1 \neq x_2$. A quadratic function with positive leading term, such as

$$ax^2 + bx$$

is strongly convex.

Theorem 15.6.1. *The social cost of a s-t flow is a convex function.*

Proof. Recall that the social cost of a flow f is

$$\sum_e c_e(f_e) f_e = \sum_e a_e f_e^2 + b_e f_e.$$

This is a sum of convex function, and so it is convex. □

Theorem 15.6.2. *Let $c(x)$ be a convex function (of one variable or many). Then, every local minimum of c is a global minimum.*

Proof. Let x^1 be a global minimum of C. If x^0 is a local minimum of c, the gradient of c at x^0 is zero in every direction. However, if $c(x^1) < c(x^0)$, then the gradient in the direction $x^1 - x^0$ will be negative: we have

$$\frac{c(x^0 + \epsilon(x^1 - x^0)) - c(x^0)}{\epsilon} \leq \frac{(1-\epsilon)c(x^0) + \epsilon c(x^1) - c(x^0)}{\epsilon} = \frac{\epsilon(c(x^1) - c(x^0))}{\epsilon} = c(x^1) - c(x^0) < 0.$$

□

Corollary 15.6.3. *Let f^0 and f^1 be socially optimal flows. Then for all $0 < t < 1$, the flow $tf^0 + (1-t)f^1$ is socially optimal too. Moreover, for every edge e for which $a_e > 0$, $f_e^0 = f_e^1$.*

Proof. The first part follows immediately from the previous two theorems. The second part follows from the strict convexity of quadratics. □

15.7 Social optimum as Nash equilibrium

Theorem 15.7.1. *Let G be a graph with edge costs given by $c_e(f) = a_e f + b_e$. Let $\tilde{G}$ be the same graph but with with edge costs given by*

$$\tilde{c}_e(f) \stackrel{\text{def}}{=} 2a_e f + b_e.$$

Then, a flow is a social optimum for G if and only if it is a Nash equilibrium for $\tilde{G}$.

Proof. Let f be a socially optimal flow for G. It can be represented as having flow α_i over path P_i for some set of coefficients $\alpha_1, \ldots, \alpha_k$. Well, there might be many such representations. We only need to know that there is at least one. We will prove that for all i such that $\alpha_i > 0$ and for all j,

$$\tilde{c}(P_i) \leq \tilde{c}(P_j),$$

which establishes the claimed result.

To see this, we will prove that if we shift a little bit of the flow from path i to path j. If we decrease α_i by ϵ and increase α_j by ϵ, for small ϵ, then the change in the cost of the social optimum approaches

$$\epsilon \left(\frac{\partial c(\alpha_1, \ldots, \alpha_k)}{\partial \alpha_j} - \frac{\partial c(\alpha_1, \ldots, \alpha_k)}{\partial \alpha_i} \right) \geq 0,$$

as $\alpha_1, \ldots, \alpha_k$ is a social optimum.

So, for $\alpha_i > 0$, we know that for every j

$$\frac{\partial c(\alpha_1, \ldots, \alpha_k)}{\partial \alpha_j} \geq \frac{\partial c(\alpha_1, \ldots, \alpha_k)}{\partial \alpha_i}.$$

Now, let's compute those derivatives.

We have

$$\begin{aligned}
\frac{\partial c(\alpha_1, \ldots, \alpha_k)}{\partial \alpha_j} &= \sum_e \frac{\partial f_e c_e(f_e)}{\partial \alpha_j} \\
&= \sum_{e \in P_j} \frac{\partial f_e c_e(f_e)}{\partial \alpha_j} \\
&= \sum_{e \in P_j} \frac{\partial a_e f_e^2 + b_e f_e}{\partial \alpha_j} \\
&= \sum_{e \in P_j} \frac{\partial a_e f_e^2 + b_e f_e}{\partial f_e} \frac{\partial f_e}{\partial \alpha_j} \\
&= \sum_{e \in P_j} \frac{\partial a_e f_e^2 + b_e f_e}{\partial f_e} && (\text{for } e \in P_j) \\
&= \sum_{e \in P_j} 2a_e f_e + b_e \\
&= \tilde{c}(P_j).
\end{aligned}$$

Conversely, we see that if α is a Nash equilibrium for $\tilde{G}$ then it is a local social optimum. As local social optima are also global social optimum, it must be one of those as well. □

While I claimed the following corollary in class, I neglected to prove it.

Corollary 15.7.2. *Every Nash equilibrium has the same social cost.*

Proof. Follows from Theorem 15.7.1 and Corollary 15.6.3. □

So, the cost of a Nash equilibrium is now well defined. Let $\mathrm{Nash}(G)$ denote the value of a Nash equilibrium in G and let $\mathrm{Opt}(G)$ denote the value of the social optimum.

Theorem 15.7.3. *Let G be a directed road network in which every edge e has a cost function of the form $c_e(f_e) = a_e f_e + b_e$, with $a_e, b_e \geq 0$. Let f be any Nash flow in G. Then*

$$c(f) \leq \frac{4}{3}\mathrm{Opt}(G).$$

Proof. The key to the proof is the construction of another road network $\widehat{G}$ that has the same vertices as G, all the edges of G, and an additional set of edges. Let f be a Nash flow for G. Our construction of $\widehat{G}$ will depend upon f. For each edge $e \in G$, $\widehat{G}$ contains both e and another edge $\hat{e}$ that has the same endpoints as e, but the cost function

$$c_{\hat{e}}(x) \stackrel{\text{def}}{=} c_e(f_e).$$

That is, the cost of edge $\hat{e}$ is a constant function, and equals the cost of edge e under the Nash flow f.

We now observe that f, when treated as a flow in $\widehat{G}$, is a Nash flow in $\widehat{G}$. The reason is that every s-t path in $\widehat{G}$ will consist of a combination of both original and hatted edges. For each of the hatted edges $\hat{e}$, there is an original edge e that has the same cost under f. So, there is no s-t path in $\widehat{G}$ that has a lower cost than the paths being used in the flow f.

We will now construct an optimal flow in $\widehat{G}$, which we will call g. This flow g will make equal use of each of the original and the hatted edges. That is, for every e

$$g_e = \frac{1}{2} f_e \quad \text{and} \quad g_{\hat{e}} = \frac{1}{2} f_e.$$

To show that g is an optimal flow in $\widehat{G}$, we will construct a graph $\tilde{G}$ that has the same edges as $\widehat{G}$, but costs

$$\tilde{c}_e(x) = 2a_e x + b_e \quad \text{and} \quad \tilde{c}_{\hat{e}}(x) = c_e(f_e).$$

By Theorem 15.7.1, g is an optimal flow in $\widehat{G}$ if and only if it is a Nash flow in $\tilde{G}$. To show that g is a Nash flow in $\tilde{G}$, observe that

$$\begin{aligned} \tilde{c}_e(g_e) &= 2a_e(f_e/2) + b_e = c_e(f_e), && \text{and} \\ \tilde{c}_{\hat{e}}(g_{\hat{e}}) &== c_e(f_e). \end{aligned}$$

So, g is a Nash flow in $\tilde{G}$ precisely because f is a Nash flow in G.

Now that we know both a Nash flow and an optimal flow in $\widehat{G}$, we can show

$$\frac{3}{4}\mathrm{Nash}(\widehat{G}) \leq \mathrm{Opt}(\widehat{G}).$$

This follows from

$$\begin{aligned}
\mathrm{Nash}(\widehat{G}) &= \sum_e f_e c_e(f_e) + \sum_e f_{\hat{e}} c_{\hat{e}}(f_{\hat{e}}) \\
&= \sum_e f_e c_e(f_e) \\
&= \sum_e a_e f_e^2 + b_e f_e,
\end{aligned}$$

and

$$\begin{aligned}
\mathrm{Opt}(\widehat{G}) &= \sum_e g_e c_e(g_e) + \sum_e g_{\hat{e}} c_{\hat{e}}(g_{\hat{e}}) \\
&= \sum_e \frac{f_e}{2} c_e(\frac{f_e}{2}) + \frac{f_e}{2} c_{\hat{e}}(\frac{f_e}{2}) \\
&= \sum_e \frac{f_e}{2}(a_e \frac{f_e}{2} + b_e) + \frac{f_e}{2}(a_e f_e + b_e) \\
&= \sum_e \frac{3}{4} a_e f_e^2 + f_e b_e \\
&\geq \frac{3}{4} \sum_e a_e f_e^2 + f_e b_e \\
&= \frac{3}{4}\mathrm{Nash}(\widehat{G}).
\end{aligned}$$

To finish the proof, observe that $\mathrm{Nash}(G) = \mathrm{Nash}(\widehat{G})$ and, as $\widehat{G}$ has every edge that G does,

$$\mathrm{Opt}(\widehat{G}) \leq \mathrm{Opt}(G).$$

□

15.8 Bounding Braess's Paradox

In the previous lecture, I showed how to use Theorem 15.7.3 to bound the effect of Braess's paradox. I'll now do that formally.

Theorem 15.8.1. *Let G be a road network and let $\widehat{G}$ be a road network obtained by adding roads to G. Then*

$$\mathrm{Nash}(\widehat{G}) \leq \frac{4}{3}\mathrm{Nash}(G).$$

Proof.

$$\mathrm{Nash}(G) \geq \mathrm{Opt}(G) \geq \mathrm{Opt}(\widehat{G}) \geq \frac{3}{4}\mathrm{Nash}(\widehat{G}).$$

□

References

[RT02] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.