Algorithms, Games, and Networks February 5, 2013

Lecture 7

Lecturer: Avrim Blum *Scribe: Fan Xiang*

1 Overview

Today's lecture is on **Non-atomic selfish routing**. Readings are 18.1, 18.2.1, 18.3.1, 18.4.1 in AGT.

2 Model

The non-atomic selfish routing problem can be viewed as a game-theoretic type of *multi-commodity flow* problem. The setup is as follows. We are given a directed graph G, where each edge e has a cost (or congestion) function $c_e(x)$ based on how much flow x is on the edge. The cost functions $c_e(x)$ are assumed to be non-decreasing, continuous and with the property that $xc_e(x)$ is convex. For example, we might have a constant cost function, a linear cost function, a quadratic cost function, etc. We are also given a set of flow requirements: for each pair of vertices (s_i, t_i) we have r_i units of flow that want to travel from s_i to t_i. These are assumed to be infinitely divisible, which is why this is called *non-atomic*. E.g., think of an individual player or packet as being infinitesimally small compared to the overall flow wanting to get from s_i to t_i. The players each want to travel on a shortest (i.e., least costly) path for themselves, regardless of how this impacts others, which is why this is called *selfish* routing.

So, a non-atomic selfish routing problem is defined by a triple $\{G, c, r\}$: G is the directed graph with edges e, c is the set of cost functions c_e for each edge e, and r is the set of flow requirements: r_i units of flow that need to travel from s_i to t_i.

Definitions: For a flow f, we will use f_e to denote the amount of flow on edge e. Given a flow f, the cost of a path p is $C_p(f) = \sum_{e \in p} c_e(f)$. The overall social cost of the flow is $Cost(f) = \sum_e c_e(f_e) \cdot f_e$. We will use P_i to denote the set of paths from s_i to t_i, and f_p to denote the amount of flow using some path p.

In this model, we are interested in

- what do equilibria look like?
- what is the social optimal solution?

- what is the price of anarchy?

3 Equilibrium and optimal flow

Definition 1 (Non-atomic equilibrium flow) *A flow f is at equilibrium if everyone is traveling on a cheapest path given the flow. That is, for all i, all $p \in P_i$ s.t. $f_p > 0$, all $\widetilde{p} \in P_i$, $C_p(f) \leq C_{\widetilde{p}}(f)$.*

Definition 2 (optimal flow) *A flow f is optimal flow if it minimizes the overal cost $Cost(f)$.*

4 Pigou's example

Let's take a look at Pigou's example for non-atomic selfish routing.
We want to route 1 unit of flow from s to t in figure 1.
The equilibrium flow routes all of the traffic on the second link and has cost 1. The optimal

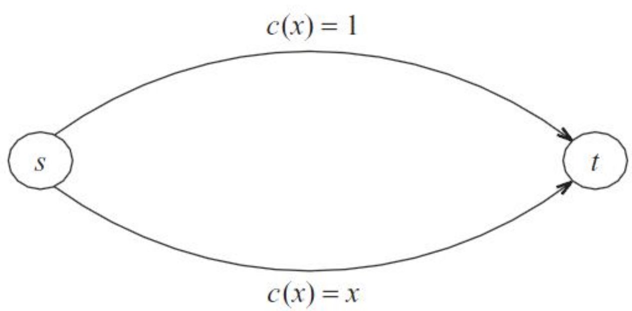

Figure 1: Pigou's example

flow has 0.5 units on each link. The total cost is 3/4. The price of anarchy is 4/3.

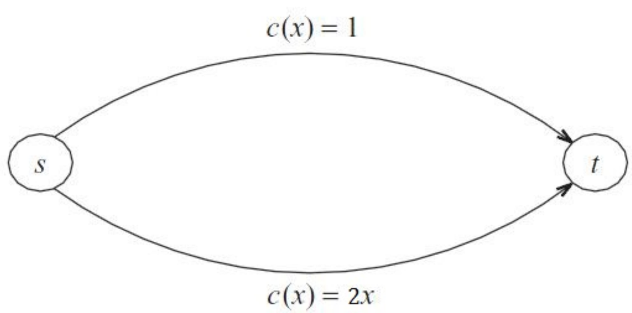

Figure 2: Pigou's example variant

In figure 2, the equilibrium is to send 0.5 unit of flow on each path.

5 Price of anarchy in non-atomic selfish routing

5.1 Pigou's example: general case

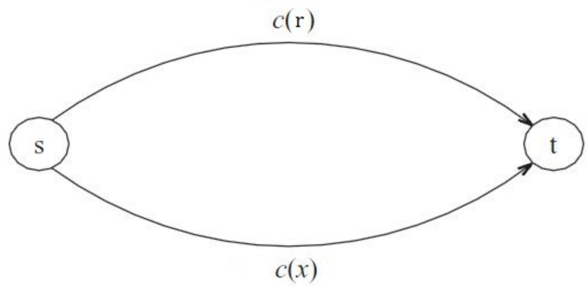

Figure 3: general example

In the general case of Pigou's example, we want to flow r units of flow from s to t, and c is some strictly increasing function. The equilibrium is to flow all r units on the bottom link, since otherwise anyone going on the top link would want to travel on the bottom instead (since the cost function is increasing). Thus the cost of the equilibrium flow is $rc(r)$. The optimal cost is

$$OPT = \min_{x \leq r} (r-x)c(r) + xc(x).$$

Actually, since the function $c(x)$ is strictly increasing, even if somehow OPT were allowed to use $x > r$ in the above expression, it would not be advantageous to do so, so we can drop the "$x \leq r$" condition (this will be useful later). Given a class of Pigou networks specified by a class C of cost functions (e.g., linear, quadratic, etc), the *price of anarchy* is the maximum (over cost functions in the class) of the ratio between the cost of the equilibrium flow and that of the optimal flow. So, the price of anarchy is:

$$PoA = \sup_{c \in C} \sup_{x,r \geq 0} \frac{rc(r)}{(r-x)c(r) + xc(x)}.$$

Note that the Price of Anarchy is larger for cost functions that are higher degree polynomials, see Figure 4.

We now take a slight detour and then get back to the issue of Price of Anarchy in general networks.

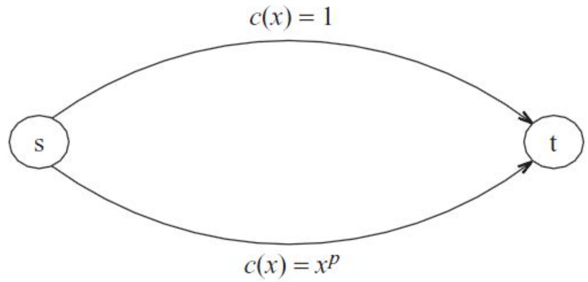

Figure 4: Degree p polynomial. In this example we have $r = 1$. The unique equilibrium flow routes all traffic on the bottom link for a cost of 1. The optimal flow routes a small fraction ε on the top link and all the rest on the bottom link. The cost of flow is $\varepsilon + (1-\varepsilon)^p$. When p tends to infinity and ε tends to 0, the cost tends to 0. The price of anarchy tends to infinity. This example shows the inefficiency of equilibrium flow.

5.2 Equivalence of equilibrium and optimal flows

We now show that the optimal flow in (G, c, r) is an equilibrium flow for a different set of cost functions (G, c^*, r) where

$$c_e{}^*(x) = \frac{d}{dx}(xc_e(x)) \text{ is the } \textit{marginal cost function.}$$

Note that $xc_e(x)$ can be viewed as the social cost of edge e (the contribution of edge e to the overall social cost, when x amount of flow is flowing on edge e).

Claim 3 (Equivalence of equilibrium and optimal flows) *Let (G, c, r) be a non-atomic instance such that, for every edge e, the function $x \cdot c_e(x)$ is convex and continuously differentiable. Let c_e^* denote the marginal cost function of the edge e. Then f^* is an optimal flow for (G, c, r) if and only if, for every commodity $i \in \{1, 2, \cdots, k\}$ and every pair p, $\widetilde{p} \in P_i$ such that $f_p^* > 0$, $C_p^*(f^*) \leq C_{\widetilde{p}}^*(f^*)$*

Proof: If $\widetilde{p}$ has lower $C_{\widetilde{p}}^*$, we move dx from p to $\widetilde{p}$. The cost in (G, c, r) is going to drop by $[C_p^*(f^*) - C_{\widetilde{p}}^*(f^*)]dx$. So, this says that if f^* is an optimal flow for (G, c, r) then such $\widetilde{p}$ can't exist. In the other direction, because $xc(x)$ is convex function, we know that a local optimum (f^* such that no such $\widetilde{p}$ exists) is also a global optimum. ■

5.3 Potential function

We want to show that in non-atomic selfish routing games, equilibrium flows always exist and are essentially unique. Specifically, if the cost function is strictly increasing, we only have one equilibrium flow. If the cost function is not strictly increasing, all the equilibrium flows have the same cost on every edge.

In the previous section we showed that the optimal flow in the original game is an equilibrium flow in a different game. Here we will prove this result by showing that the equilibrium flow in the original game is the optimal flow in a different game. In particular, this new game will be one defined by the potential function, which is just a continuous version of what we saw in the last two lectures.

Definition 4 *The potential function of non-atomic selfish routing instance* (G, c, r) *is*

$$\Phi(f) = \sum_e \int_0^{f_e} c_e(y)dy$$

The potential function for edge e is $\Phi(e) = \int_0^{f_e} c_e(y)dy$

5.4 Existence and uniqueness of equilibrium flow

Next, let's use the potential function to prove equilibrium flow exists and is unique. First, we prove the equilibria exist by proving the optimal value exists for the potential function. The potential function is a continuous function because $c_e(y)$ is continuous. The set of feasible flows of (G, c, r) is a closed and bounded subset of P-dimensional Euclidean space. So the potential function has a minimum value on this set. That is equivalent of saying there exists a equilibrium flow for (G, c, r) according to claim 3. Second, we prove the equilibrium flow is unique. If f and $\widetilde{f}$ are equilibrium flows for (G, c, r), by claim 3, they both minimize the potential function. We consider linear combination of f and $\widetilde{f}$, $\lambda f + (1-\lambda)\widetilde{f}$. The fact that these are convex functions guarantees that $\Phi(\lambda f + (1-\lambda)\widetilde{f}) \leq \lambda\Phi(f) + (1-\lambda)\Phi(\widetilde{f})$. So the equality must occur. This can occur only if every summand $\int_0^x c_e(y)dy$ is linear between the values f and $\widetilde{f}$. This implies that the cost function c_e is constant between f and $\widetilde{f}$.

5.5 Pigou bound

Claim 5 (Variational inequality characterization) *f is equilibrium iff for any other flow $\widetilde{f}$,*

$$\sum_e c_e(f_e) \cdot f_e \leq \sum_e c_e(f_e) \cdot \widetilde{f}_e$$

Proof: Let's for convenience assume the total flow is 1. Then the LHS in the equation above corresponds to the expected cost of a random person if they choose to use f, given that everyone else is using f. The RHS in the equation above corresponds to the expected cost of a random person if they choose to use $\widetilde{f}$, given that everyone else is using f. In particular, you can view $\widetilde{f}_e$ as the probability you would travel on edge e if you woke up as a random flow-path according $\widetilde{f}_e$, so then $c_e(f_e)\widetilde{f}_e$ is the contribution of edge e to your total expected cost if everyone else is using f. Flow f is an equilibrium exactly when nobody would prefer any other $\widetilde{f}$ to f given that everyone else is using f. To address the case that the total flow is $r \neq 1$ just divide the LHS and RHS above by r. ■

Claim 6 (Tightness of the Pigou bound) *Let C be a set of cost functions and $\alpha(C)$ the Pigou bound for C. If (G, c, r) is a non-atomic instance with cost functions in C, then the price of anarchy of (G, c, r) is at most $\alpha(C)$.*

Proof: For given arbitrary graph G,

$$cost(f^*) = \sum_e c_e(f_e{}^*) \cdot f_e{}^* \geq \sum_e \frac{c_e(f_e) \cdot f_e}{\alpha(C)} - (f_e - f_e^*)c_e(f_e)$$

$$= \frac{1}{\alpha(C)} \cdot cost(f) - \sum_e (f_e - f_e^*)c_e(f_e)$$

The inequality follows from the definition of Pigou's bound.

$$(r-x)C(r) + xC(x) \geq \frac{rC(r)}{\alpha(C)}$$

replacing with $x = f_e^*, r = f_e$.
$\sum_e (f_e - f_e^*)c_e(f_e)$ is negative. We can see that in claim 5 replacing $f_e{}^*$ for $\widetilde{f}_e$.
So,

$$cost(f^*) \geq \frac{1}{\alpha(C)} \cdot cost(f)$$

$$\frac{cost(f)}{cost(f^*)} \leq \alpha(C)$$

■