

Lecture 2: Selfish Routing

Lecturer: Sergei Vassilvitskii

Scribe: Etienne Vouga

Recall from the first lecture that we are working graphs $G : (V, E)$ with vertices V and directed edges E . We also specify two distinct special vertices, $s \in V$ and $t \in V$, termed the *source* and *sink* respectively, and associate monotonic *edge latency functions* $\ell_e : [0, 1] \rightarrow \mathbb{R}^+$ to each edge $e \in E$. We let $P = \{p_1, \dots, p_n\}$ be the set of paths from s to t , and define a *flow* f to be an assignment to each path p_i a nonnegative real number f_{p_i} , with $\sum_{i=1}^n f_{p_i} = 1$. One intuitive interpretation of flow is to imagine 1 unit of infinitely splittable supply starting at s , and trying to travel to t . Then an f_{p_i} fraction of supply travels on each path p_i .

We consider only simple paths, that is each path p_i has no cycles, and passes through each edge at most once. We write $e \in p_i$ if p_i passes through edge e . Then for a flow f of G , we can compute the flow f_e through an edge e by summing the contribution of each path through e :

$$f_e = \sum_{i|e \in p_i} f_{p_i}.$$

The latency contributed by edge e is then $\ell_e(f_e)$, and the cost $c_{p_i}(f_{p_i})$ induced by f on a path p_i is

$$c_{p_i}(f) = \sum_{e \in p_i} f_{p_i} \ell_e(f_e).$$

The total cost $C(f)$ is then the sum of the costs of all of the paths:

$$\begin{aligned} C(f) &= \sum_i c_{p_i}(f) \\ &= \sum_i \sum_{e \in p_i} f_{p_i} \ell_e(f_e) \\ &= \sum_{e \in E} \sum_{p_i \in P} f_{p_i} \ell_e(f_e) \\ &= \sum_{e \in E} f_e \ell_e(f_e) \end{aligned}$$

The last equality follows after switching the order of summations, and will prove useful to us later.

Given the above structure, there are two canonical problems we can pose: first, we can try to find an *optimal routing*, a flow f^* which minimizes the total cost $C(f^*)$. Second, we can look for a *Nash equilibrium* f , in which, intuitively, every infinitesimal “piece” of supply acts selfishly and takes the cheapest available path to the sink, regardless of the consequences this choice has on the rest of the supply. We give a more precise formulation of Nash equilibria below.

Given the two solutions flows f^* and f , we can consider the ratio $\frac{C(f)}{C(f^*)}$, the price paid by the supply for acting greedily instead of in concert. This is also known as the *price of anarchy*. The remainder of the lecture proves the following theorem:

Theorem 1 *In the case of linear latency functions, ℓ_e , the price of anarchy, $\frac{C(f)}{C(f^*)} \leq \frac{4}{3}$.*

1 Nash Equilibrium

Suppose we have a flow f on G , and any two paths p_1, p_2 with $f_{p_1} > 0$. If you're flowing along p_1 and acting greedily, and flowing along p_2 instead would be cheaper, you would switch; hence, a necessary condition on f being a Nash equilibrium is that

$$\sum_{e \in p_1} \ell_e(f_e) \leq \sum_{e \in p_2} \ell_e(f_e) \quad \forall p_2 \in P, p_1 \in P | f_{p_1} > 0. \quad (1)$$

It can be shown the above condition is also sufficient.

2 Optimal Solution

Again, consider a flow f on G , p_1, p_2 paths, with $f_{p_1} > 0$. Suppose we were to move some amount of flow δ from p_1 to p_2 . On the one hand, the cost along p_1 will decrease since less supply is moving through it, but on the other hand, the cost along p_2 will increase. If f is optimal, such a switch cannot improve the total cost, so the benefit cannot outweigh the cost.

Writing it down formally, the total cost of the original flow f is:

$$C(f) = \sum_e c_e(f_e) = \sum_{\substack{e \in p_1 \\ e \in p_2}} c_e(f_e) + \sum_{\substack{e \in p_1 \\ e \notin p_2}} c_e(f_e) + \sum_{\substack{e \notin p_1 \\ e \in p_2}} c_e(f_e) + \sum_{\substack{e \notin p_1 \\ e \notin p_2}} c_e(f_e).$$

Consider a flow f' where we take a δ fraction of flow from p_1 and route it on p_2 instead. Notice that for edges e contained both in p_1 and p_2 or neither p_1 nor p_2 the flow doesn't change. Formally, :

$$f'_e = \begin{cases} f_e & e \in p_1, e \in p_2 \\ f_e - \delta & e \in p_1, e \notin p_2 \\ f_e + \delta & e \notin p_1, e \in p_2 \\ f_e & e \notin p_1, e \notin p_2 \end{cases}$$

The total cost of f' is:

$$C(f') = \sum_e c_e(f'_e) = \sum_{\substack{e \in p_1 \\ e \in p_2}} c_e(f_e) + \sum_{\substack{e \in p_1 \\ e \notin p_2}} c_e(f_e - \delta) + \sum_{\substack{e \notin p_1 \\ e \in p_2}} c_e(f_e + \delta) + \sum_{\substack{e \notin p_1 \\ e \notin p_2}} c_e(f_e).$$

If f is the optimal solution, then:

$$\begin{aligned} C(f) &\leq C(f') \\ \sum_{\substack{e \in p_1 \\ e \notin p_2}} c_e(f_e) + \sum_{\substack{e \notin p_1 \\ e \in p_2}} c_e(f_e) &\leq \sum_{\substack{e \in p_1 \\ e \notin p_2}} c_e(f_e - \delta) + \sum_{\substack{e \notin p_1 \\ e \in p_2}} c_e(f_e + \delta) \\ \sum_{\substack{e \in p_1 \\ e \notin p_2}} (c_e(f_e) - c_e(f_e - \delta)) &\leq \sum_{\substack{e \notin p_1 \\ e \in p_2}} (c_e(f_e + \delta) - c_e(f_e)) \end{aligned}$$

Dividing through by δ , and taking the limit as $\delta \rightarrow 0$, we have:

$$\sum_{\substack{e \in p_1 \\ e \notin p_2}} c'_e(f) \leq \sum_{\substack{e \notin p_1 \\ e \in p_2}} c'_e(f)$$

where $c'_e(f) = \frac{dc_e}{df_e}$ is the derivative of the cost function. By adding $\sum_{\substack{e \in p_1 \\ e \in p_2}} c'_e(f)$ to both sides we conclude that a solution f is optimal if, for a path p_1 for $f_{p_1} > 0$,

$$c'_{p_1}(f) < c'_{p_2}(f).$$

It turns out that this condition, too, is both necessary and sufficient. Observe that this condition looks superficially similar to the condition for a flow f to be an equilibrium flow, the only difference is the functions – latency for the equilibrium, marginal costs for the optimum.

3 Linear Latency Functions

Suppose now our latency functions are linear, that is, of the form $\ell_e(x) = a_e x + b_e$. Then the above Nash equilibrium condition (1) becomes, assuming $f_{p_1} > 0$,

$$\sum_{e \in p_1} a_e f_e + b_e \leq \sum_{e \in p_2} a_e f_e + b_e. \quad (2)$$

Similarly, for the optimal solution, using the chain rule we get

$$\begin{aligned} \sum_{e \in p_1} \left(\frac{d}{df_e} f_e \ell_e(f_e) \right) (f_e) &\leq \sum_{e \in p_2} \left(\frac{d}{df_e} f_e \ell_e(f_e) \right) (f_e) \\ \sum_{e \in p_1} (\ell_e(f_e) + f_e \ell'_e(f_e)) (f_e) &\leq \sum_{e \in p_2} (\ell_e(f_e) + f_e \ell'_e(f_e)) (f_e) \\ \sum_{e \in p_1} a_e f_e + b_e + f_e a_e &\leq \sum_{e \in p_2} a_e f_e + b_e + f_e a_e \\ \sum_{e \in p_1} 2a_e f_e + b_e &\leq \sum_{e \in p_2} 2a_e f_e + b_e \end{aligned} \quad (3)$$

Notice that if $b_{e_1} = b_{e_2} = 0$, equations (2) and (3) are multiples of each other, and so any flow that is a Nash equilibrium is optimal, and vice versa.

The key insight for analyzing the general case $b_{e_i} \neq 0$ is to notice that *half of the Nash flow is optimal for routing half a unit of supply*, since if f satisfies equation (2), then $\frac{f}{2}$ satisfies (3). In particular, for each edge e , the marginal cost over e after routing half of the flow optimally is the latency of the Nash flow $\ell_e(f_e)$. Our strategy now will be to bound $C(f^*)$, the cost of the optimal solution, relative to the Nash solution, by taking the Nash route for half of the supply, and bounding the cost of routing the other half.

Before we proceed, we prove a technical lemma. Let $L^*(f)$ be the marginal cost of increasing the flow starting with the flow f .

Lemma 1 *Let f^* be an optimal partial flow of some amount $k < 1$ of supply, and f any flow of $k + \delta$ supply. Then*

$$C(f) \leq C(f^*) + \delta L^*(f^*),$$

where $L^*(f^*)$ is the marginal cost of increasing the flow starting with f^* , $L^*(f^*) = \sum_i c'_{p_i}(f^*)$.

Proof. Notice that since f^* is optimal, we know that for any two paths p_1 and p_2 , the marginal costs are the same, $\sum_{e \in p_1} c'_e(f^*) = \sum_{e \in p_2} c'_e(f^*) = L^*(f^*)$. Consider a flow f , with cost

$$\begin{aligned} C(f) &= \sum_{e \in E} f_e \ell_e(f_e) \\ &\geq \sum_{e \in E} f_e^* \ell_e(f_e) + \sum_{e \in E} (f_e - f_e^*) c'_e(f_e^*) \end{aligned}$$

The inequality follows due to the convexity of the cost function. It is each to check that $f_e > f_e^*$ then $c_e(f_e) > c_e(f_e^*) + (f_e - f_e^*) c'_e(f_e^*)$. On the other hand, if $f_e < f_e^*$ then $c_e(f_e) < c_e(f_e^*) + (f_e - f_e^*) c'_e(f_e^*)$. Continuing, we have:

$$\begin{aligned} C(f) &\geq \sum_{e \in E} f_e^* \ell_e(f_e) + \sum_{e \in E} (f_e - f_e^*) c'_e(f_e^*) \\ &= C(f^*) + \sum_{p_i \in P} (f_{p_i} - f_{p_i}^*) c'_{p_i}(f_{p_i}^*) \\ &= C(f^*) + \sum_{p_i \in P} (f_{p_i} - f_{p_i}^*) L^*(f^*) \\ &= C(f^*) + \delta L^*(f^*) \end{aligned}$$

□

We are now ready to prove the theorem.

Proof of Theorem 1: Letting f^* be the optimal flow, f the Nash flow, and $\frac{f}{2}$ half of the Nash flow. We know from the discussion above that for the path p_i , the marginal cost $c'_{p_i}\left(\frac{f}{2}\right)$ is equal to the Nash latency $\ell_{e_i}(f_{e_i})$. Thus $L^*\left(\frac{f}{2}\right) = C(f)$, and so, applying the lemma with $\delta = \frac{1}{2}$,

$$C(f^*) \geq C\left(\frac{f}{2}\right) + \frac{1}{2} L^*\left(\frac{f}{2}\right) \geq C\left(\frac{f}{2}\right) + \frac{1}{2} C(f).$$

It remains to calculate $C\left(\frac{f}{2}\right)$:

$$\begin{aligned}
C\left(\frac{f}{2}\right) &= \sum_e \frac{f_e}{2} l_e\left(\frac{f_e}{2}\right) \\
&= \sum_e \frac{f_e}{2} \left(a_e \frac{f_e}{2} + b_e\right) \\
&= \sum_e \left(\frac{a_e f_e^2}{4} + \frac{b_e f_e}{2}\right) \\
&\geq \sum_e \frac{1}{4} (a_e f_e^2 + b_e f_e) \\
&= \frac{1}{4} C(f),
\end{aligned}$$

so $C(f^*) \geq \frac{3}{4}C(f)$, and

$$\frac{C(f)}{C(f^*)} \leq \frac{4}{3}.$$