How bad is Selfish Routing?
- Tim Roughgarden and Eva Tardos

Yogesh Balaji

University of Maryland

April 17, 2018

Overview

Motivation

- Traffic routing problem - Given the rate of traffic between each pair of nodes in a network, find an assignment of traffic to paths so that the sum of all travel times (the total latency) is minimized
- Latency function for each edge is load dependent
- Hard to impose optimal routing strategies on users. Results in selfish behavior
- *How much does network performance suffer from this lack of regulation?*

Braees's Paradox

Assume we need to transport 1 unit from node s to node t. The figure in the right has one additional edge of 0 latency running from v to w.

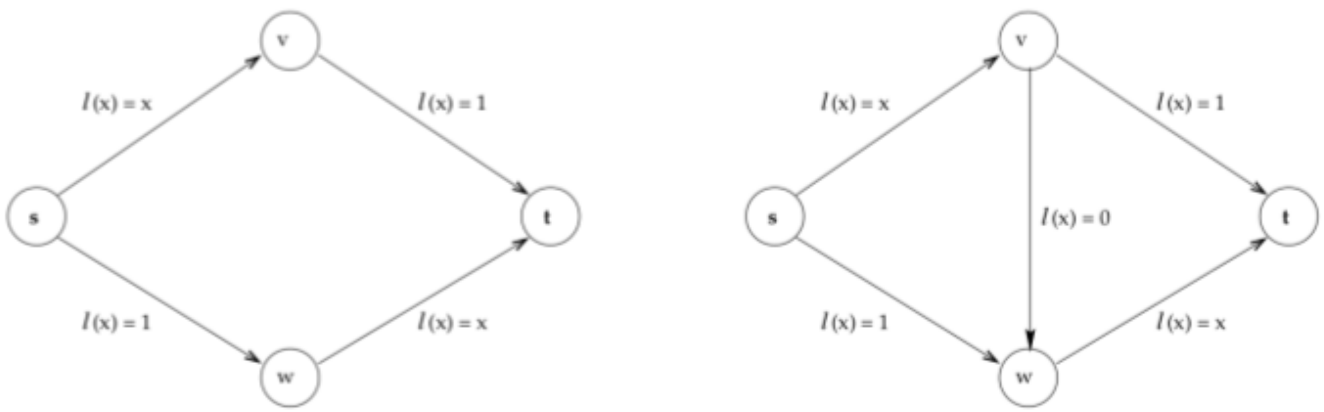

Braees's Paradox

Additional of an edge negatively impacts the performance of the system

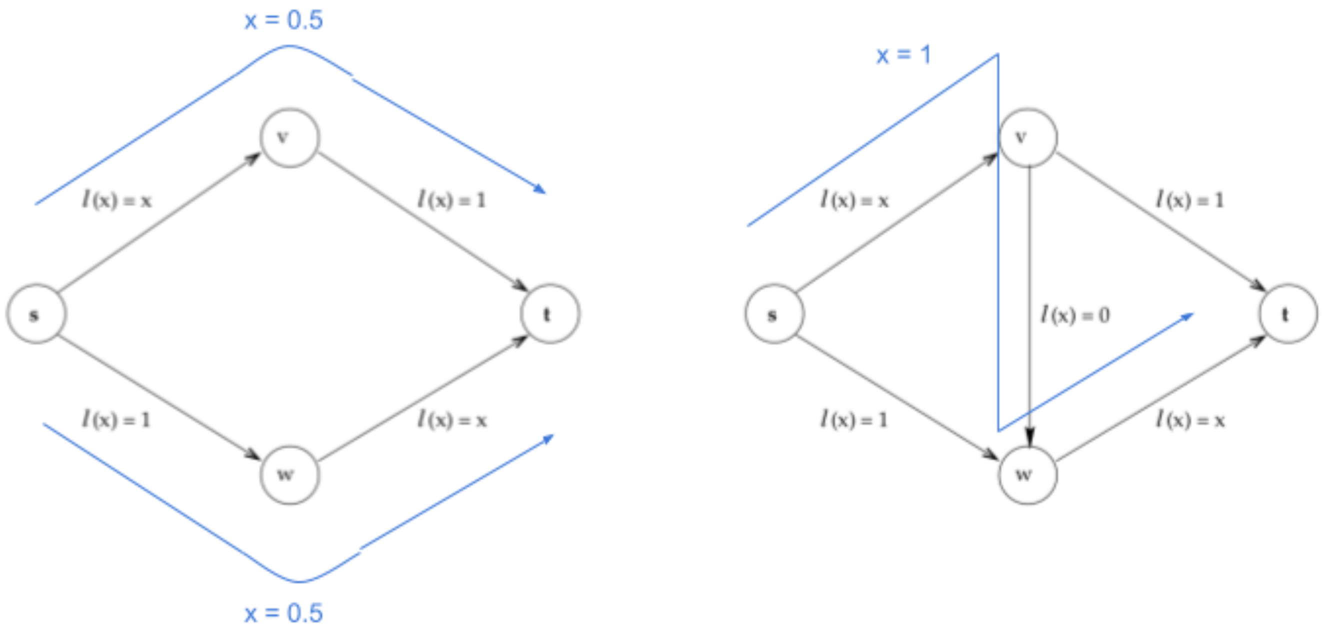

Optimal cost: 3/2
Equilibrium cost: 3/2

Optimal cost: 3/2
Equilibrium cost: 2

Notation

- Directed Graph $G = (V, E)$, V - vertex set, E - edge set
- $\{(s_1, t_1), (s_2, t_2), \dots (s_k, t_k)\}$ - set of source destination pairs
- $\mathcal{P}_i$ - Set of simple s_i-t_i paths; $\mathcal{P} = \cup_i \mathcal{P}_i$
- $f : \mathcal{P} \rightarrow \mathcal{R}^+$ - flow for a given path
- $f_e = \sum_{P:e \in P} f_P$ - net flow flowing through a given edge e
- r_i - Amount of flow to be sent from s_i to t_i
- $l_e(f)$ - Latency for the edge e, function of the flow f. l_e is assumed to be nonnegative, differentiable and nondecreasing
- $l_P(f) = \sum_{e \in P} l_e(f_e)$ - latency for the path f
- $C(f) = \sum_{P \in \mathcal{P}} l_P(f) f_P$ - Total latency of flow f: Path formulation
- $C(f) = \sum_{e \in E} l_e(f_e) f_e$ - Total latency of flow f - edge formulation

Flows at Nash Equilibrium

Definition 1

A flow f feasible for instance (G, r, l) is at Nash Equilibrium if $\forall i \in \{1, 2, \ldots k\}$, $P_1, P_2 \in \mathcal{P}_i$ and $\delta \in [0, f_{P_1}]$, we have $l_{P_1}(f) \leq l_{P_2}(\tilde{f})$, where

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

This is because at Nash Equilibrium, users do not have any incentive to deviate.

Flows at Nash Equilibrium

Letting $\delta \to 0$ in the above definition, we find that at Nash Equilibrium, all $s_i - t_i$ flow paths (paths to which f assigns a positive flow) will have equal latency. Denote this by $L_i(f)$.

Lemma 1

If f is a flow at Nash Equilibrium for instance (G, r, l), then

$$C(f) = \sum_{i=1}^{k} L_i(f) r_i$$

Optimal Flow

The optimal flow can be found by minimizing the following non-linear program

Optimal Flow Formulation

$$\min \sum_{e \in E} c_e(f_e)$$

subject to

$$\sum_{P \in \mathcal{P}_\rangle} f_P = r_i \quad \forall i \in \{1, 2, \ldots k\}$$

$$f_e = \sum_{P \in \mathcal{P} : e \in P} f_P \quad \forall e \in E$$

$$f_P \geq 0 \quad \forall P \in \mathcal{P}$$

Bicriteria result

- We will analyze the ratio between the cost of a flow at Nash Equilibrium to that of minimum latency flow.
- Price of Anarchy $= \rho(G, r, l) = \frac{\text{Cost of Flow at Nash Equilibrium}}{\text{Cost of Minimum Latency Flow}}$
- In the Braees's Paradox example, $\rho(G, r, l) = 4/3$
- We will upper bound the cost of the Nash flow by the cost of optimal flow at *increased rates*

Bicriteria result

Theorem (Thm 3.1)

If f is a flow at Nash equilibrium at (G, r, l) and f^* is feasible for $(G, 2r, l)$, then $C(f) \leq C(f^*)$

Proof

We know from Lemma 1 that

$$C(f) = \sum_i L_i(f) r_i$$

Construct new latency function $\bar{l}_e(x)$ such that

$$\bar{l}_e(x) = \begin{cases} l_e(f_e) & \text{if } x \leq f_e \\ l_e(x) & \text{if } x > f_e \end{cases}$$

Proof cont.

$$x(\bar{l}_e(x) - l_e(x)) \leq l_e(f_e)f_e \quad \forall x \geq 0 \tag{1}$$

$$\sum_e \bar{l}_e(f_e^*)f_e^* - C(f^*) = \sum_{e \in E} f_e^*(\bar{l}_e(f_e^*) - l_e(f_e^*)) \tag{2}$$

$$\leq \sum_{e \in E} l_e(f_e)f_e \tag{3}$$

$$= C(f) \tag{4}$$

In other words, evaluating f^* with latency function $\bar{l}$ (rather than l) increases its cost by at most an additive factor of $C(f)$

Proof cont.

Let f_0 denote zero flow in G. By construction, $\bar{l}_P(f_0) \geq L_i(f)$ for any path $P \in \mathcal{P}_i$. Since l_e is nondecreasing for each edge e, $\bar{l}_P(f^*) \geq L_i(f)$. So,

$$\sum_e \bar{l}_e(f_e^*)f_e^* = \sum_P \bar{l}_P(f^*)f_P^* \geq \sum_i \sum_{P \in \mathcal{P}_i} L_i(f)f_P^* \tag{5}$$

$$= \sum_i 2L_i(f)r_i \tag{6}$$

$$= 2C(f) \tag{7}$$

Proof cont.

Combining the previous two slides,

$$
\begin{aligned}
C(f^*) &\geq \sum_P \bar{l}_P(f^*) f_P^* - C(f) \\
&\geq 2C(f) - C(f) = C(f) \\
& QED.
\end{aligned}
$$

More general form of previous result:

Theorem (Thm 3.2)

If f is a flow at Nash equilibrium at (G, r, l) and f^* is feasible for $(G, (1+\gamma)r, l)$, then $C(f) \leq \frac{1}{\gamma} C(f^*)$

References

Roughgarden, Tim and Tardos, Éva
How Bad is Selfish Routing?
J. ACM March 2002, 236–259.

Thank You