Lecture 5: How Bad is Selfish Routing?

September 13, 2016

Given:

- A directed network $G(V, E)$
- A set of source-destination pairs s_i-t_i with a fixed traffic rate r_i
- A load dependent latency function $l_e : R^+ \rightarrow R^+$ for each edge.
  - non-decreasing
  - continuous

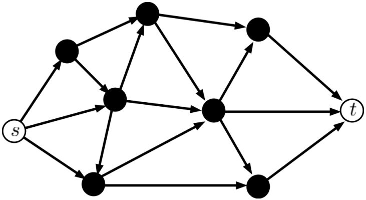

- SWM: find a way to route all traffic such that total cost $\sum_e f_e l_e(f_e)$ is minimized

A routing game

- Infinitely many players
- Each arriving at some source s_i at rate r_i
- The routes chosen by every player can be described using a vector f (the flow)

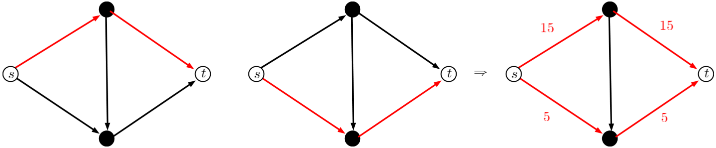

A routing game

- Let $\mathcal{P}$ be the set of all s_i-t_i paths. For a flow f and a s_i-t_i path P:
  - f_P = flow through P
  - Cost of each player $l_P(f) = \sum_{e \in P} l_e(f_e)$, where $f_e = \sum_{P \in \mathcal{P}: e \in P} f_P$
  - Cost of flow $C(f) = \sum_{P \in \mathcal{P}} f_P l_P(f) = \sum_{e \in E} f_e l_e(f_e)$
- In the absence of external traffic control, assume every player plays a selfish game

Social Welfare Maximization

Given instance (G, r, l).

$$
\begin{aligned}
\text{minimize} \quad & \sum_{e} f_e l_e(f_e) \\
\sum_{P \in \mathcal{P}_i} f_P &= r_i && \forall i \\
\sum_{P \in \mathcal{P} : e \in P} f_P &= f_e && \forall e \in E \\
f_P &\geq 0 && \forall P \in \mathcal{P}
\end{aligned}
$$

Example

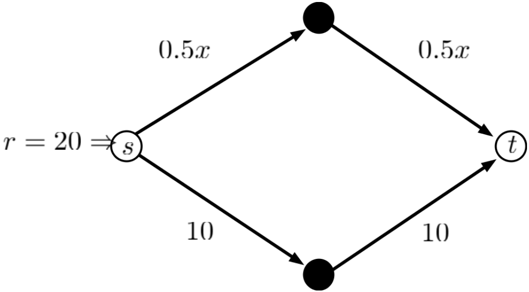

Example

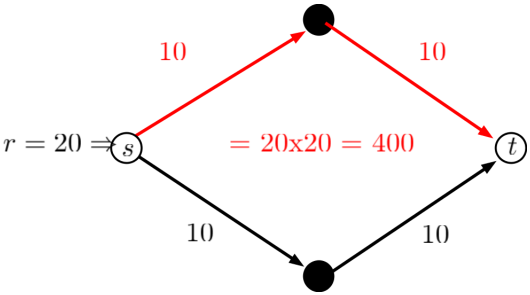

Example

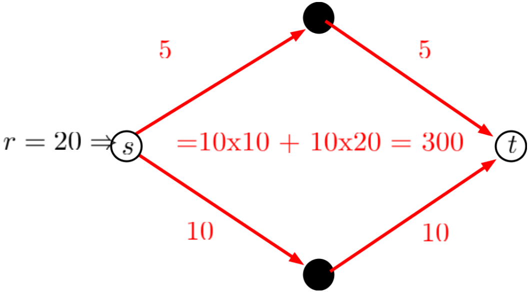

Price of Anarchy

For an instance (G, r, l), let f be the flow at equilibrium and f^* be the flow at social welfare max. How far is $C(f)$ from $C(f^*)$?

Result

Theorem

If f and f^* are the Nash and optimal flows for (G, r, l) with linear latency functions, then $\frac{C(f)}{C(f^*)} \leq \frac{4}{3}$

Convex Program

Given instance (G, r, l). Assume $l_e(x) = a_e x + b_e$. Then cost function $C(f) = \sum_e a_e f_e^2 + b_e f_e$ is convex.

$$
\begin{aligned}
\text{minimize} \ \sum_e a_e f_e^2 + b_e f_e & \\
\sum_{P \in \mathcal{P}_i} f_P = r_i & \qquad \forall i \\
\sum_{P \in \mathcal{P} : e \in P} f_P = f_e & \qquad \forall e \in E \\
f_P \geq 0 & \qquad \forall P \in \mathcal{P}
\end{aligned}
$$

Flows at Equilibrium

- Lemma 1.1: A feasible flow f for (G, r, l) is at Nash equilibrium (Nash flow) if for each s_i-t_i pair, $P_1, P_2 \in \mathcal{P}_i$ and $f_{P_1} > 0$, we have:

$$\sum_{e \in P_1} a_e f_e + b_e \leq \sum_{e \in P_2} a_e f_e + b_e$$

- $\Rightarrow$ All s_i-t_i paths must have the same latency $= L_i(f)$.

$$\Rightarrow C(f) = \sum_{P \in \mathcal{P}} f_P l_P(f) = \sum_i L_i(f) r_i$$

Flows at Equilibrium

- $\frac{d}{dx}C = \sum_e 2a_e x + b_e \Rightarrow$ marginal cost.
- Lemma 1.2: A feasible flow f^* for (G, r, l) is (globally) optimal if for each pair s_i-t_i, $P_1, P_2 \in \mathcal{P}_i$, and $f^*_{P_1} > 0$, we have:

$$\sum_{e \in P_1} 2a_e f^*_e + b_e \leq \sum_{e \in P_2} 2a_e f^*_e + b_e$$

- in other words, f^* is a Nash Flow for instance (G, r, l^*), where $l^*_e(x) = 2a_e x + b_e$

Flows at equilibrium

- Lemma 1.3: If f is a Nash flow for (G, r, l)
  a $f/2$ is optimal for $(G, r/2, l)$
  b marginal cost of increasing flow on a path P with respect to $f/2 =$ latency of P with respect to f

Proof.
Observe $l_e(x) = a_e x + b_e = 2a_e \frac{x}{2} + b_e = l_e^*(x/2)$.
For part a: Since f is a Nash flow it satisfies conditions of Lemma 1.1. Also, $f/2$ is feasible for instance $(G, r/2, l)$ and from Lemma 1.2, it is also optimal.
For part b: Summing over all edges in P, we have
$l_P^*(f/2) = l_P(f)$ □

Cost of Augmentation

- Lemma 1.4: If f^* is the optimal flow for (G, r, l) and $L_i^*(f^*)$ is the minimum marginal cost of increasing flow on a s_i-t_i path w.r.t f^*, then for any $\delta > 0$ a feasible flow for $(G, (1+\delta)r, l)$ has cost $\geq C(f^*) + \delta \sum_i L_i^*(f^*) r_i$

Proof.

Let f be a feasible flow for $(G, (1+\delta)r, l)$. Since the per edge cost $xl_e(x)$ is convex,

$$f_e l_e(f_e) \geq f_e^* l_e(f_e^*) + (f_e - f_e^*) l_e^*(f_e^*)$$

$$\therefore C(f) \geq C(f^*) + \sum_i \sum_{P \in \mathcal{P}_i} l_P^*(f^*)(f_P - f_P^*)$$

□

Proof Outline

- From Lemma 1.3 $f/2$ is an optimal flow for $(G, r/2, l)$, where f is a Nash flow for (G, r, l).

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

- We then augment the optimal flow $f/2$ for $(G, r/2, l)$ to the optimal flow f^* for (G, r, l). Then from Lemma 1.4 ($\delta = 1$):

$$\begin{aligned} C(f^*) &\geq C(f/2) + \sum_i L_i^*(f/2) \frac{r_i}{2} \\ &\geq \frac{1}{4} C(f) + \frac{1}{2} C(f) = \frac{3}{4} C(f) \end{aligned}$$

Extensions to general latency functions

- A bad example (for large k):

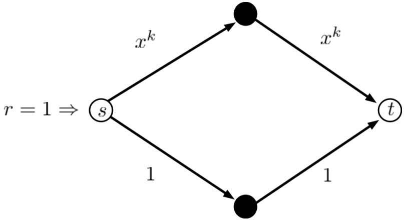

- Cost of Nash flow $= 2$. Cost of optimal flow ≈ 0

Extensions to general latency functions

Theorem

If f is a Nash flow for (G, r, l) with continuous, non-decreasing latency functions. Then the $C(f)$ is at most the cost of an optimal flow for $(G, 2r, l)$

Selfish routing and network design - Braess's Paradox

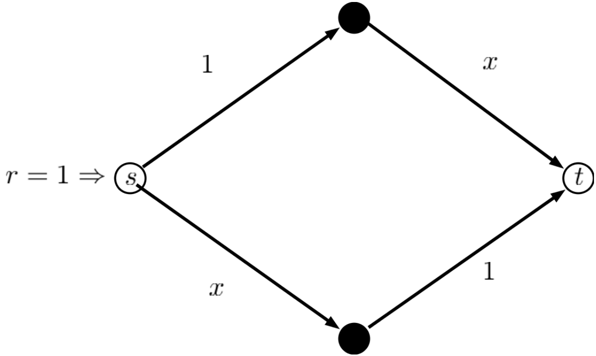

$C(f) = C(f^*) = \frac{3}{2}$

Selfish routing and network design - Braess's Paradox

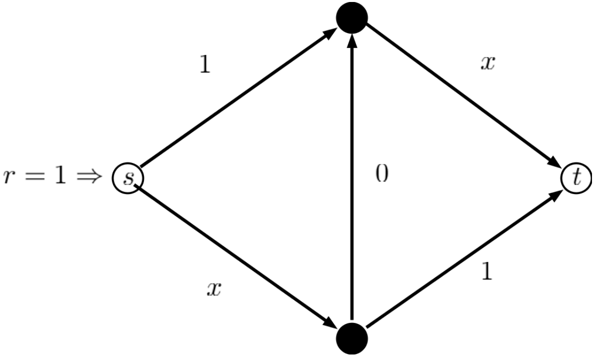

$C(f) = 2, C(f^*) = \frac{3}{2}$