"How Bad is Selfish Routing?" by Tim Roughgarden and Éva Tardos

Enoch Lau

14 May 2007

Outline

Note about notation: The notation used is from Roughgarden's book *Selfish Routing and the Price of Anarchy* while the content of the presentation is based on the paper.

The Motivation

What route will take *me* to work in the shortest time?

- Time taken depends on congestion on roads
- Congestion depends on choices of other users on the road
- Other users are selfish too!

The problem: With a lack of central coordination, how much worse is everyone off in total?

Results to be Discussed

- If latency of each edge is a linear function of its congestion, the total latency caused by selfish users is at most $\frac{4}{3}$ times the minimum possible total latency.
- For more general edge functions, the total latency caused by selfish users is no more than the total latency incurred by optimally routing twice as much traffic.

Pigou's Example

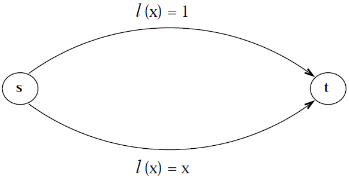

- Route traffic from s to t
- x is the fraction of traffic that uses the road
- Cost functions $c(\cdot)$ give travel time as function of fraction of traffic
- What will selfish users do?
- All traffic will follow the lower road
- Optimal routing: half on lower route

Selfish behaviour need not produce a socially optimal outcome.

Braess' Paradox

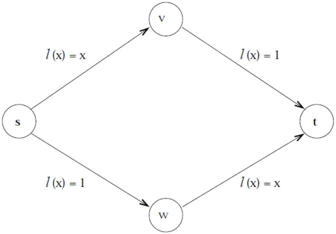

(a) Before

- Both routes are identical to start with
- Traffic should be split evenly between the two routes
- Travel time: 1.5 hours
- Effect of adding a road (v, w) with $c(x) = 0$?

Braess' Paradox

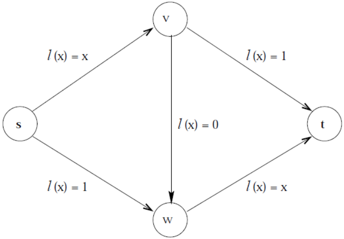

(b) After

- Everyone will use $s \longrightarrow v \longrightarrow w \longrightarrow t$
- If any traffic fails to use it, cost is strictly less than cost of any other path

With selfish routing, network improvements can degrade network performance.

Graph

- Directed graph: $G = (V, E)$
- Parallel edges allowed
- Sources: $s_1, \ldots, s_k$
- Destinations: $t_1, \ldots, t_k$
- All traffic from s_i travel to t_i
- Commodity: $\{s_i, t_i\}$

Flow

- $\mathcal{P}_i$: set of s_i-t_i paths
- $\mathcal{P} = \cup_i \mathcal{P}_i$: set of all paths
- Flow f: aggregated routes chosen by a large number of agents
- f_P: amount of agents using P
- f_e: amount of flow on edge e
- Each agent controls a negligible amount of overall traffic

Traffic

- Traffic rate r_i: amount of traffic that needs to flow from s_i to t_i
- $0 < r_i < \infty$
- f is feasible if, for all i, all traffic gets routed along some path:

$$\sum_{P \in \mathcal{P}_i} f_P = r_i$$

Costs

- Cost function of edge e: $c_e(\cdot)$
- $c_e(\cdot)$ is non-negative, continuous and non-decreasing
- Instance: (G, r, c)
- $c_P(f)$: the sum of the cost of edges in path P
- $C(f)$: the sum of all costs incurred by the traffic:

$$C(f) = \sum_{P \in \mathcal{P}} c_P(f) f_P = \sum_{e \in E} c_e(f_e) f_e$$

- Flow is optimal if it minimises $C(f)$ over all feasible flows

Game Theory 101

- Game theory is the discipline where players make decisions that maximise their returns
- A game consists of:
  - Players
  - Strategies
  - Payoffs for each combination of strategies
- A Nash equilibrium is a set of strategies from which no player has an incentive to deviate

Prisoner's Dilemma

	Co-operate	Defect
Cooperate	3, 3	0, 5
Defect	5, 0	1, 1

- Two choices: co-operate and defect
- Defect is the only stable strategy
- Regardless of what the other player chooses, defecting gives a higher payoff

Rational decisions may not be sensible.

Nash Flows

- A flow f feasible for (G, r, c) is at Nash equilibrium if for all commodities i, s_i-t_i paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$,

$$c_{P_1}(f) \leq c_{P_2}(\tilde{f})$$

  where $\tilde{f}$ is obtained from f by moving some flow from P_1 to P_2
- Some nice properties:
  - All s_i-t_i flow paths have a common cost
  - Every instance (G, r, c) admits a Nash flow
  - $C(f) = C(\tilde{f})$ for two Nash flows f and $\tilde{f}$

The Price of Anarchy

- The price of anarchy is the worst-possible ratio between the cost of a Nash flow and that of an optimal flow
- If f^* is an optimal flow and f is a Nash flow, then

$$\rho(G, r, c) = \frac{C(f)}{C(f^*)}$$

- All Nash flows of an instance have the same cost, so it is well defined

Example of the Price of Anarchy

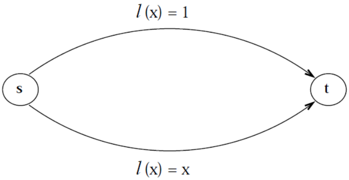

- $C(f) = 1$
- $C(f^*) = \frac{3}{4}$
- $\rho = \frac{4}{3}$

Optimal Flows via Non-Linear Programs

Finding optimal feasible flow for (G, r, c) via (NLP)

- Min $\sum_{e \in E} h_e(f_e)$
- Subject to:

$$\begin{aligned} \sum_{P \in \mathcal{P}_i} f_P &= r_i & \forall i \in \{1, \ldots, k\} \\ f_e &= \sum_{P \in \mathcal{P} : e \in P} & \forall e \in E \\ f_P &\geq 0 & \forall P \in \mathcal{P} \end{aligned}$$

with $h_e(f_e) = c_e(f_e) f_e$

$h_e(x) = x \cdot c_e(x)$ is a convex function and (NLP) is a convex program.

Lemma

A flow f is optimal for (NLP) iff for every commodity i $P_1, P_2 \in \mathcal{P}_i$ *with* $f_{P_1} > 0$

$$h'_{P_1}(f) \leq h'_{P_2}(f)$$

where $h'_P(f) = \sum_{e \in P} h'_e(f_e)$ *and* $h'_e = \frac{d}{dx} h_e(x)$

Proof.

A flow is locally optimal iff moving flow from one path to another (belonging to the same commodity) increases the cost. This happens when the marginal benefit of decreasing flow along one path is at most the marginal cost of increasing flow along another path. The local and global minima of a convex function on a convex set coincide. The lemma is true if the objective function is convex. □

From the definition of a flow at Nash equilibrium, we obtain the following lemma.

Lemma

A flow f feasible for the instance (G, r, c) is at Nash equilibrium iff for every commodity i and $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$

$$c_{P_1}(f) \leq c_{P_2}(f)$$

Definition

If c is a differentiable cost function, then the corresponding marginal cost function is

$$c^* = \frac{d}{dx}(x \cdot c(x))$$

Relationship Between Optimal and Nash Flows

Putting the two lemmas together, we obtain the following corollary:

Corollary

Let (G, r, c) be an instance with continously differentiable, semi-convex cost functions. A flow f feasible for (G, r, c) is optimal iff it is at Nash equilibrium for the instance (G, r, c^*).

Properties of Nash Flows

- Existence: Follows from the previous corollary
- Uniqueness: If f and $\tilde{f}$ are Nash flows for (G, r, c), then $c_e(f_e) = c_e(\tilde{f}_e)$ for every edge e
- All Nash flows have equal cost

Bicriteria Results

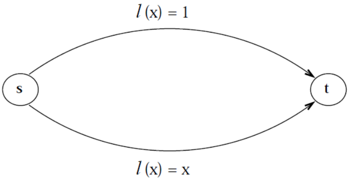

If we double the flow, the additional unit of flow goes on the top edge if we have an optimum flow, and the cost is still 1.

The Theorem

Theorem (Bicriteria Result for General Latency Functions)

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

Proof: see paper

Generalisation

Theorem (Generalisation of Bicriteria Result)

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

The Theorem

Theorem

Let (G, r, c) be an instance with edge latency functions, then $\rho(G, r, c) \leq \frac{4}{3}$.

This bound is a tight bound.

Preliminary Observations

- Edge latency functions: $c_e(x) = a_e x + b_e$ for some $a_e, b_e \geq 0$
- Total latency of a flow f: $C(f) = \sum_e a_e f_e^2 + b_e f_e$ (which implies (NLP) is a convex program and the previous characterisation holds)
- Marginal cost function: $c_e^*(x) = 2a_e x + b_e$

Lemma (4.1)

(a) *A flow f is at Nash equilibrium iff for each commodity i and $P, P' \in \mathcal{P}_i$, with $f_P > 0$*

$$\sum_{e \in P} a_e f_e + b_e \leq \sum_{e \in P'} a_e f_e + b_e$$

(b) *A flow f^* is optimal iff for each commodity i and $P, P' \in \mathcal{P}_i$, with $f^*_P > 0$*

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

This follows as a specialisation of the previous lemmas.

Lemma (4.3)

If f is a Nash flow, then

(a) *The flow $f/2$ is optimal for $(G, r/2, c)$*

(b) *The marginal cost of increasing the flow on a path P with respect to $f/2$ equals the latency of P with respect to f*

(a) follows from Lemma 4.1(a).
For (b), because $c_e(x) = a_e x + b_e$ and $c_e^*(x) = 2a_e x + b_e$, $c_e^*(f_e/2) = c_e(f_e)$ and hence $c_P^*(f/2) = c_P(f)$ for each path P.

Proof Outline

We will create an optimal flow for (G, r, c) in two stages:

1. Send flow optimal for $(G, r/2, c)$ through G
2. Augment the flow to one optimal for (G, r, c)

The first flow has cost at least $\frac{1}{4}C(f)$ and the augmentation has cost at least $\frac{1}{2}C(f)$, where f is some Nash flow. Overall, this means $C(f) \leq \frac{4}{3}C(f^*)$.

Bounding the cost of $f/2$

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

Augmentation Cost

Lemma

Let $c_i^*(f^*)$ be the minimum marginal cost of increasing flow on an s_i-t_i path with respect to f^*. Then, for any $\delta > 0$, a feasible flow for the problem instance $(G, (1+\delta)r, c)$ has cost at least

$$C(f^*) + \delta \sum_{i=1}^{k} c_i^*(f^*) r_i$$

Intuitively: the per-unit cost of increasing the amount of flow through a network is at least the marginal cost of increasing flow on any path with respect to the current optimal flow.

Proof: see paper

Proof of Theorem

- Let f be a flow at Nash equilibrium. Let $c_i(f)$ be the latency of an s_i-t_i flow path, and

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

  because all flow paths for a commodity in a Nash flow have the same latency.
- By Lemma 4.3(a), $f/2$ is an optimal solution to $(G, r/2, c)$.
- By Lemma 4.3(b), $c_i^*(f/2) = c_i(f)$ for each i.

Proof of Theorem

- Taking $\delta = 1$ in Lemma 4.4, the cost of any flow f^* feasible for (G, r, c) satisfies

$$\begin{aligned} C(f^*) &\geq C(f/2) + \sum_{i=1}^{k} c_i^*(f/2)\frac{r_i}{2} \\ &= C(f/2) + \frac{1}{2}\sum_{i=1}^{k} c_i(f)r_i \\ &= C(f/2) + \frac{1}{2}C(f) \end{aligned}$$

- Substitute the bound for cost of $f/2$:

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

Generalisation

In any instance (G, r, c), where for some p, $c_e(x) = a_e x^p + b_e$, $a_e, b_e \geq 0$ for each e

$$\rho(G, r, c) = \Theta\left(\frac{p}{\ln p}\right)$$

Flows at Approximate Nash Equilibrium

Theorem

If f is at ϵ-approximate Nash equilibrium with $\epsilon < 1$ for (G, r, c) and f^* is feasible for $(G, 2r, c)$, then $C(f) \leq \frac{1+\epsilon}{1-\epsilon} C(f^*)$.

Finitely Many Agents: Splittable Flow

Theorem

If f is at Nash equilibrium for the finite splittable instance (G, r, c) with $x \cdot c_e(x)$ convex for each e, and f^* is feasible for the finite splittable instance $(G, 2r, c)$, then $C(f) \leq C(f^*)$.

Finitely Many Agents: Unsplittable Flow

- No such analogue of the bicriteria bound
- Optimal flows can send arbitrarily more flow at arbitrarily less cost than a flow at Nash equilibrium

Summary

- A Nash equilibrium is a stable situation where no agent (or flow) has any incentive to deviate from its current strategy.
- However, for selfish routing, Nash flows are, in general, not optimal.
- If latency of each edge is a linear function of its congestion, the total latency caused by selfish users is at most $\frac{4}{3}$ times the minimum possible total latency.
- For more general edge functions, the total latency caused by selfish users is no more than the total latency incurred by optimally routing twice as much traffic.