Non-Increasing Selfish Routing

Enoch Lau, elau8465@it.usyd.edu.au
Supervisor: Dr Tasos Viglas
Advanced Networks Research Group, School of Information Technologies

Introduction

We spend a lot of our time worrying about congestion – here is a screenshot of the traffic in Seattle from Google Maps (green is fast, yellow is slower):

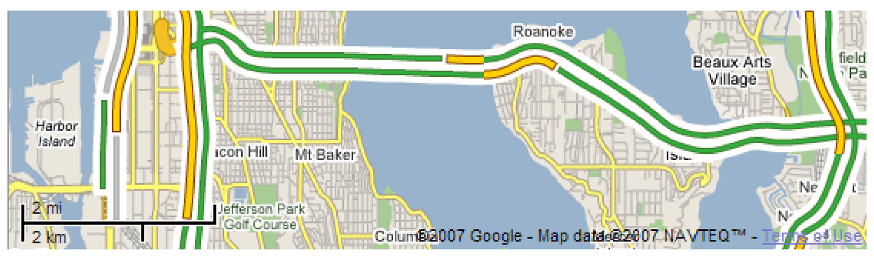

In an ideal world, we might have a central figure who directs traffic to minimise travel time globally. However, in reality, we don't have this, and everyone just minimises their *own* travel time – how much worse can that be? That is the question that we look at with *selfish routing*, a topic at the intersection of game theory and computer science. This analysis applies to computer networks as well as more traditional networks like roads.

Motivating Examples

Take a look at the first graph: 1 unit of selfish traffic wants to go from s to t on either the top or the bottom route. Each edge has a cost function, which expresses the latency along that edge in terms of the amount of traffic along that edge. Here, the average delay is 1.5.

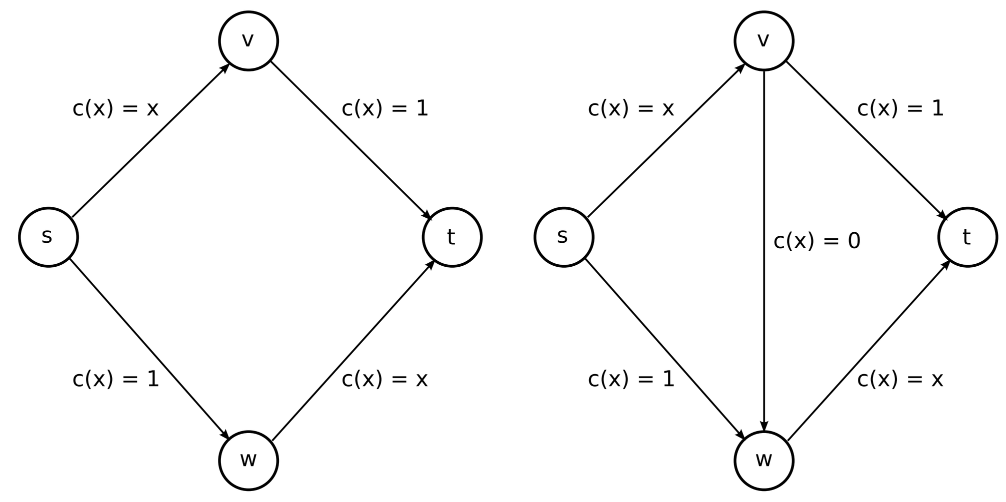

What if we wanted to make things better by adding vw? In fact, this causes all the selfish users to go via $svwt$, which means that the average delay is 2. Selfish routing need not produce socially optimal outcomes, and network improvements can cause congestion.

Selfish routing also has a mechanical interpretation. If we cut the middle string, the weight at the bottom will rise instead of fall:

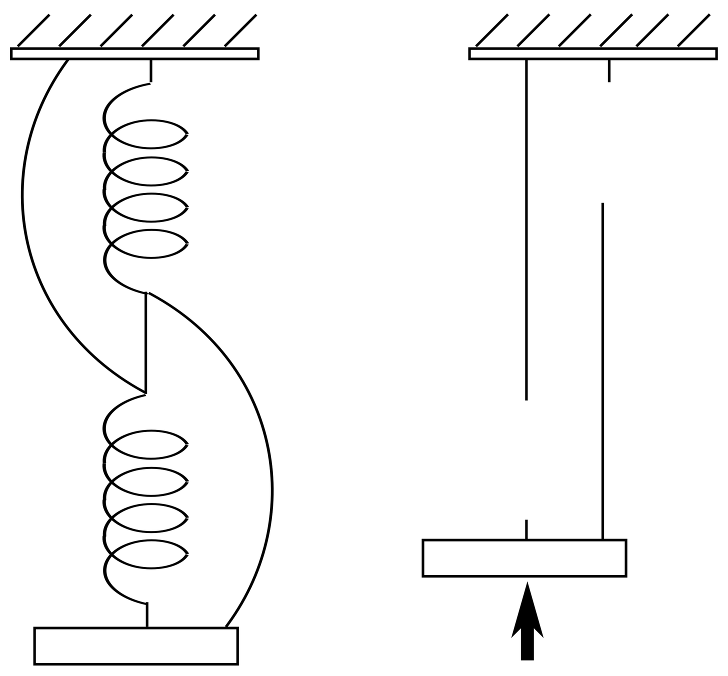

The length of a string or spring is analogous with the edge latencies, and the weight carried by each string or spring is like the amount of flow it must endure.

Non-Increasing Edge Costs

Recent developments in selfish routing have made a number of assumptions, which, while restrictive, have a natural justification. One typical restriction has been the non-decreasing nature of edge cost functions – this models a network such as a road network well, because congestion increases with the number of cars on the road. In this project, we seek to characterise the behaviour in networks where the cost functions are non-*increasing*. For example, this can be used to model economies of scale, user recommendations in e-commerce systems, and multicast routing, where data duplication along edges is minimised.

The Model

There are a number of *commodities*, which route a specified amount of traffic between two vertices in a graph, using one of the paths that join the source and the destination. Flow consists of an infinite number of agents that carry an infinitesimal amount of flow.

We are interested in characterising the equilibrium (stable) routing of the network, and comparisons between the equilibrium and optimum routing. A Nash equilibrium is a situation where no user can receive a higher payoff by unilaterally deviating. Here, it is equivalent to a Wardrop equilibrium, where all paths for a commodity have the same cost, and no unused path has a lower cost than any used path. We use a non-linear complementarity formulation of the traffic equilibrium problem [1] to prove the existence of equilibria in this model.

Contributions: We formulate the traffic equilibrium problem with non-increasing cost functions, and explore how it is different from previous formulations, and where their assumptions fall false. We provide analogies of the Braess' Paradox and the price of anarchy bounds, and we characterise the selfish and optimal flows analytically and algorithmically.

Equivalent of Braess' Paradox

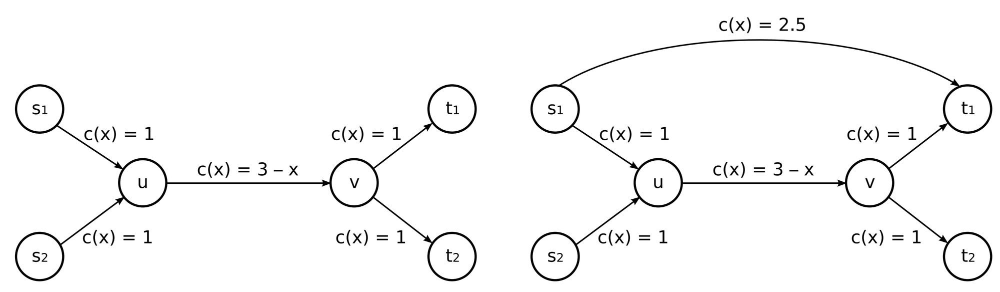

When adding the new edge above, the aim is to increase the social cost – because sharing edges lowers the cost, the intuition here is that we need to entice paths away from sharing edges.

Price of Anarchy

The *price of anarchy* ρ is a measure of how bad selfish routing can get. It is defined as:

$$\rho = \frac{\text{Cost of worst Nash flow}}{\text{Cost of optimal flow}}$$

This ratio is bounded if we apply additional restrictions on the network. For example, if the cost functions are all linear and non-decreasing, ρ is tightly upper bounded by 4/3 (Pigou's example: see left figure below). This ratio is quite resilient – for example, the presence of malicious users do not make the ratio worse [2].

If we try to find an analogous bound in this model, we find that ρ is unbounded. In the right figure:

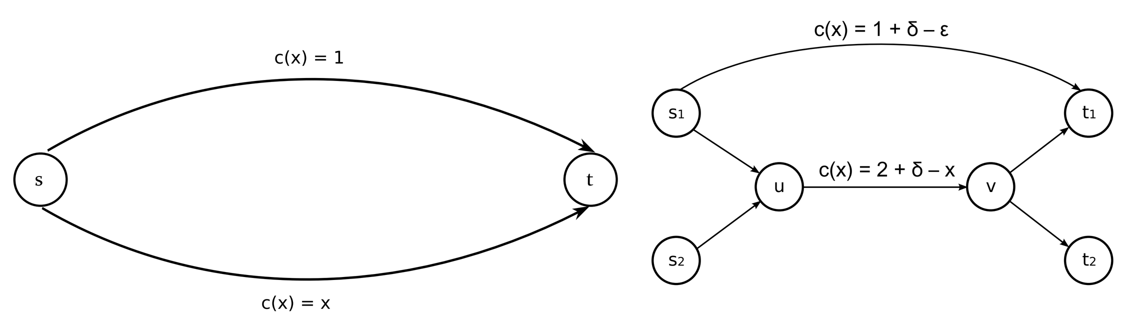

take $\delta \to 0$, and observe that the cost of the optimal routing goes to 0 while the cost of an equilibrium stays positive.

Characterisation of Nash & Optimal Flows

We have been able to provide some insight into the nature of selfish and optimal flows in this model, if we take *linear* non-increasing cost functions:

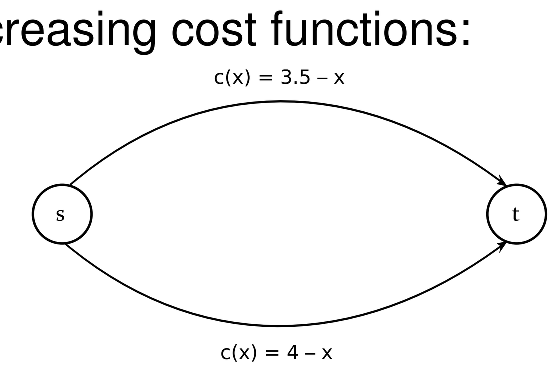

whose social cost function is plotted below on the left. The "slice" represents the flow feasibility constraint.

We have found that:

- Both types of flow only exist as pure strategies – that is, each commodity routes all flow through one path. This is a consequence of the fact that minima of concave functions over convex sets must appear at the extrema (see right figure below).
- The bicriteria bound (which for linear non-decreasing cost functions states that the Nash flow is no worse than optimally routing twice as much flow) is not as meaningful, because by routing more, we do not necessarily increase the total social cost (see left figure below).

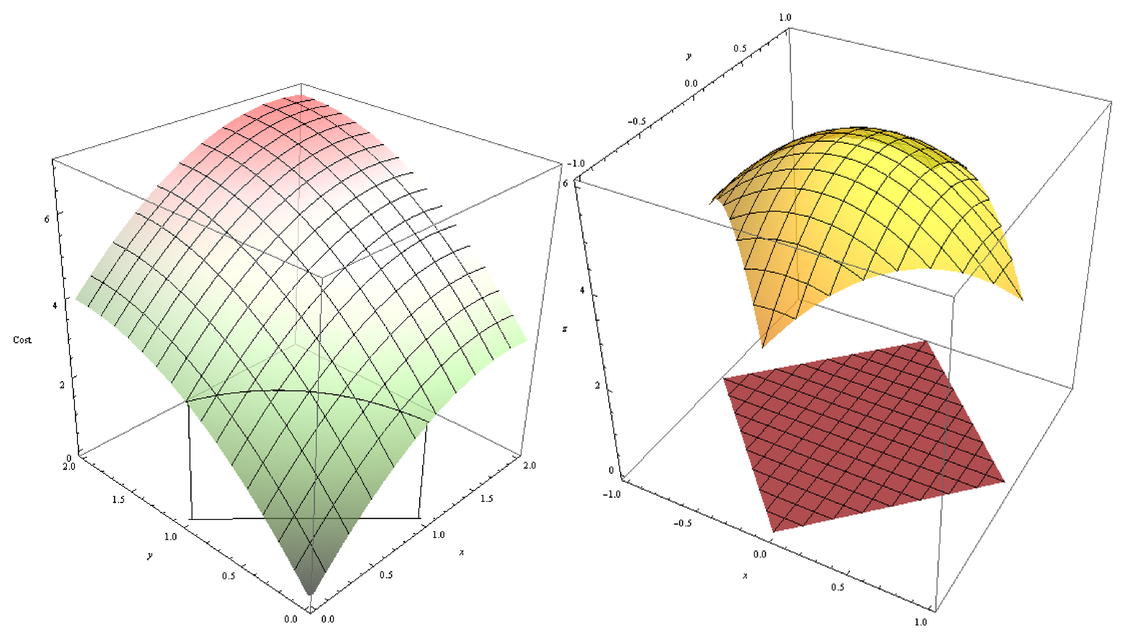

Algorithm for Finding Equilibria

Existing software for examining traffic models were found to be inappropriate for our work, primarily because of the assumption of non-decreasing cost functions, or convexity of certain functions, neither of which we guarantee. To assist our work, we formulated an exact, but inefficient algorithm to find all equilibria in a system. This relies on the fact that in the networks we consider, each commodity will only send its flow down one path.

Generate a list of paths for each commodity
for all permutations of selecting one path for each commodity **do**
    Set the flow rate on the selected paths to the commodity's rate
    Run the simulator on the flow allocations
    if the flow rates have not changed **then**
        This selection of paths is at equilibrium
    end if
end for

The simulation step uses the definition of an equilibria directly, by simulating the shift of users from higher cost paths to lower cost paths. Although we cannot simulate infinitely many users, swapping sufficiently small amounts of flow suffices for verifying flow as equilibria.

Acknowledgements

I would like to thank my supervisor Dr Tasos Viglas for his support and guidance during the year, and Dr Joachim Gudmundsson for some insightful discussions.

References

[1] H. Aashtiani and T. Magnanti. Equilibrium on a congested transport network. *SIAM Journal of Algebraic and Discrete Methods*, 2:213–216, 1983.

[2] G. Karakostas and A. Viglas. Equilibria for networks with malicious users. *Mathematical Programming*, 110(3):591–613, 2007.

[3] R. Rockefeller. *Convex Analysis*. Princeton University Press, 1970.

[4] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

The University of Sydney

University of Sydney, Camperdown, NSW 2006, Australia