Inefficiency in Networks with Multiple Sources and Sinks

Hyejin Youn[1], Michael T. Gastner[2,3,4], and Hawoong Jeong[1]

[1] Department of Physics, Korea Advanced Institute of Science and Technology, Daejeon 305-701, Korea
[2] Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA
[3] Department of Computer Science, University of New Mexico, Albuquerque, NM 87131, USA
[4] Institute for Chemistry and Biology of the Marine Environment, Carl von Ossietzky Universität, 26111 Oldenburg, Germany
visang@kaist.ac.kr

Abstract. We study the problem of optimizing traffic in decentralized transportation networks, where the cost of a link depends on its congestion. If users of a transportation network are permitted to choose their own routes, they generally try to minimize their personal travel time. In the absence of centralized coordination, such a behavior can be inefficient for society and even for each individual user. This inefficiency can be quantified by the "price of anarchy", the ratio of the suboptimal total cost to the socially optimal cost. Here we study the price of anarchy in multi-commodity networks, (i.e., networks where traffic simultaneously flows between different origins and destinations).

Keywords: flow optimization, transportation network, Nash equilibrium, multi-commodity flow.

The past few years have witnessed dramatic advances in finding, understanding, and characterizing complex networks [1,2]. One frontier for scientists is now to understand dynamic processes on networks such as the flow of matter or information in technological or social networks, vehicles traveling in transportation networks, electricity exchanged through the power grid, and the spread of diseases in biological networks [3,4,5,6]. If users of a network can decide freely which paths they take, then it is important to understand the users' behaviors and their interactions in order to optimally design the network and control the flows [7,8]. If the users' paths do not interfere with each other, users with even a small amount of local information are able to navigate on the network almost as efficiently as if they possessed global knowledge [9]. However, if the users' decisions are mutually dependent, the flow can in reality still be far from optimal even if all individuals have complete information about the network and other users' behaviors [10].

Consider, for instance, the simple network depicted in Fig. 1 [11]. Suppose that there is a constant flow of travellers F between the nodes s and t which

J. Zhou (Ed.): Complex 2009, Part I, LNICST 4, pp. 334–338, 2009.

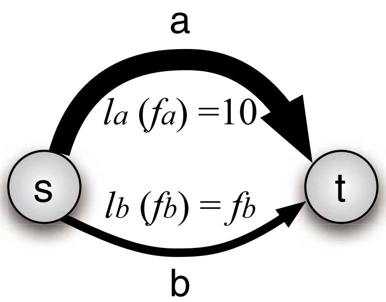

Fig. 1. Schematic depiction of a simple network with an origin s and a destination t. If 10 users travel from s to t, the flows are $f_a^{SO} = f_b^{SO} = 5$ in the social optimum and $f_a^{NE} = 0$ and $f_b^{NE} = 10$ in the Nash equilibrium.

are connected by two different types of links: a short but narrow bridge a where the effective speed becomes slower as more cars travel on it, and a long but broad multi-lane freeway b where delays due to congestion are negligible. Suppose further that the delay on link a is proportional to the flow, $l_a(f_a) = f_a$, while the delay on b is flow-independent, $l_b(f_b) = 10$, where $f_{a(b)}$ is the flow on link $a(b)$. The total time spent by all users is given by the "cost function" $\mathcal{C}(f_a) = l_a(f_a) \cdot f_a + l_b(f_b) \cdot f_b$ where the flow on b is equal to $f_b = F - f_a$. The socially optimal flow $f_{a(b)}^{SO}$ is defined as the minimum of $\mathcal{C}$. If, for instance, the total flow $F = 10$, it can be easily calculated that $f_a^{SO} = f_b^{SO} = 5$ and $\mathcal{C} = 75$.

On the other hand, every user on link b could reduce his delay from 10 to 6 by switching paths, which poses a social dilemma: as individuals, users would like to reduce their own delays, but this reduction comes at an additional cost to the entire group. In our example, as long as l_a is not equal to l_b, there will be an incentive for the users experiencing longer delays to shift to another link and finally flows reach a Nash equilibrium $f_{a(b)}^{NE}$, where no single user can make any individual gain by changing his own strategy unilaterally. All users take the link a at the total cost of $\mathcal{C} = 100$. The *price of anarchy* (POA) is defined as the ratio of the total cost of the Nash equilibrium to the total cost of the social optimum [12],

$$\text{POA} = \frac{\sum_z l_z(f_z^{NE}) \cdot f_z^{NE}}{\sum_z l_z(f_z^{SO}) \cdot f_z^{SO}}, \tag{1}$$

where, for a general network, the sums are over all links. This ratio indicates the relative inefficiency of the decentralized system; in our simple example POA=100/75 = 1.33.

There have been a number of theoretical studies about the POA (see [11] and references therein). Recently, we published calculations of the POA in real transportation networks [13]. Like most previous studies, we based our analysis on flows with a single origin and a single destination. This approach was motivated by analogies to electric circuits where even multiple current sources or sinks can be mapped onto a network with a single origin-destination pair by connecting

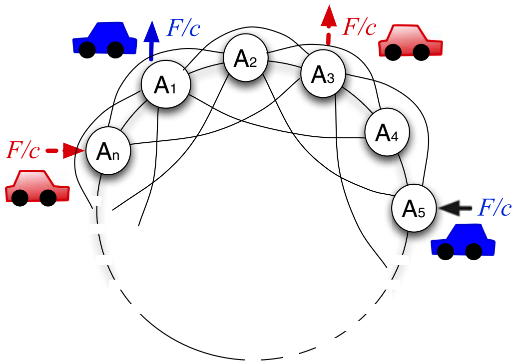

Fig. 2. Schematic depiction of a multi-commodity flow in a one-dimensional lattice with n nodes and a periodic boundary. Each node is connected to its six nearest neighbors and each commodity is assigned a flow F/c where c is the number of commodities. In the figure, red and blue cars symbolize different commodities. Origins and destinations of different commodities are not interchangeable. For example, red cars starting at node A_n are bound for A_3. A_1 is a closer destination, but it is not a destination for red cars.

all sources (sinks) to a fictitious auxiliary source (sink) and solving for the optimal and equilibrium flows. In other words, an electric current can flow to any of the sinks present in the network. However, if the flows are constrained to move from specific sources to specific destinations, the situation becomes more complicated [14]. Unfortunately, flows in human transportation networks or in the Internet fall in this class of so-called "multi-commodity flows". An important question for future research is whether the POA is affected by the number of commodities c.

In order to develop some intuition, we have studied the effect of multiple commodities in regular one-dimensional lattices (Fig. 2). All the networks contain 100 nodes; each node is connected to its nearest, next-nearest, and second-next-nearest neighbors, and periodic boundaries are applied (i.e., the lattice has the topology of a circle). Every link between nodes i and j has a delay of the form $l_{ij} = a_{ij} f_{ij} + b_{ij}$, where $a_{ij} = a_{ji}$ is a random integer equal to 1, 2, or 3, and $b_{ij} = b_{ji}$ between 1 and 100. The coefficient a_{ij} denotes how steeply the delay on the link increases with the flow present on $i \to j$. The constant b_{ij} implies a given delay regardless of congestion. The delays depend only on the total flows $f_{ij} = \sum_{k=1}^{c} f_{ij,k}$, where $f_{ij,k}$ is the flow of commodity k on the link $i \to j$ and c is the number of commodities. We then randomly choose c origins s_k and destinations t_k $(k = 1, \ldots, c)$, among the nodes. For each commodity, origin and destination must be distinct, i.e., $s_k \neq t_k$, but the origin or destination of one commodity may be the origin or destination of another commodity. To each commodity we assign a traffic flow F/c, i.e., all commodities transport an equal share of the total traffic volume F. Every commodity is permitted to split its

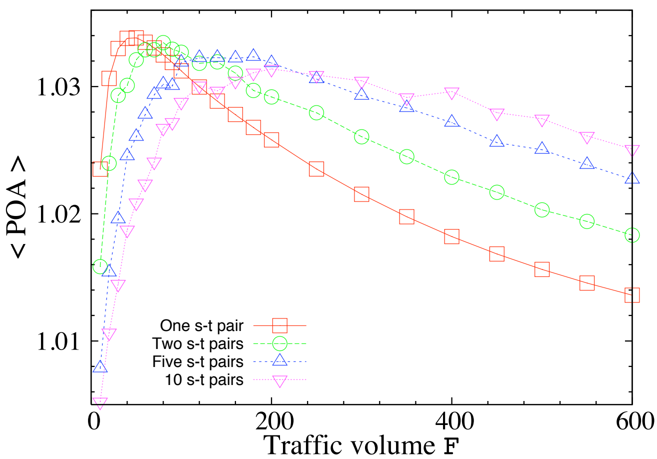

Fig. 3. The price of anarchy in regular lattices with multiple random sources and sinks ("multi-commodity flows") averaged over 100 to 400 networks. Each pair contributes equally to F.

flow on different paths from s_k to t_k. We then calculate the multi-commodity POA

$$\mathrm{POA} = \frac{\sum_{\mathrm{link}\ i \to j} \left[l_{ij} \left(\sum_k f_{ij,k}^{NE} \right) \cdot \left(\sum_k f_{ij,k}^{NE} \right) \right]}{\sum_{\mathrm{link}\ i \to j} \left[l_{ij} \left(\sum_k f_{ij,k}^{SO} \right) \cdot \left(\sum_k f_{ij,k}^{SO} \right) \right]}, \tag{2}$$

with the straightforward Frank-Wolfe algorithm [15].

The results, averaged over 100 to 400 random instances of such networks, are shown in Fig. 3. Our preliminary results indicate that the POA as a function of F is a unimodal curve for all c, but the values depend on c: the peak becomes broader and shifts to higher values as c increases. We currently hypothesize that it is possible to collapse the POA for all c to a universal scaling function $f(F)$. Unfortunately, the quality of our data is not yet sufficient to test this hypothesis. The convergence of the Frank-Wolfe algorithm is too slow to extend our calculations to $c \gg 1$. We are currently redesigning our code to implement the PARTAN algorithm of Ref. [16] which should significantly reduce the computation time.

References

1. Albert, R., Barabási, A.-L.: Statistical mechanics of complex networks. Reviews of Modern Physics 74, 47–97 (2002)
2. Newman, M.E.J.: The structure and function of complex networks. SIAM Review 45, 167–256 (2003)
3. Argollo de Menezes, M., Barabási, A.-L.: Fluctuations in Network Dynamics. Physical Review Letters 92, 028701 (2004)

4. Hufnagel, L., Brockmann, D., Geisel, T.: Forecast and control of epidemics in a globalized world. Proceedings of the National Academy of Sciences 101, 15124–15129 (2004)
5. Wu, F., Huberman, B.A., Adamic, L.A., Tyler, R.J.: Information flow in social groups. Physica A 337, 327–335 (2004)
6. Kossinets, G., Watts, D.J.: Empirical Analysis of an Evolving Social Network. Science 311, 88–90 (2006)
7. Jahn, O., Möhring, R.H., Schulz, A.S., Stier-Moses, N.E.: System-Optimal Routing of Traffic Flows with User Constraints in Networks with Congestion. Operations Research 53, 600–616 (2005)
8. Kim, D.H., Motter, A.E.: Resource allocation pattern in infrastructure networks. Journal of Physics A 41, 224019 (2008)
9. Kleinberg, J.M.: Navigation in a small world. Nature 406, 845 (2000)
10. Pigou, A.C.: The Economics of Welfare. Macmillan, London (1920)
11. Roughgarden, T.: Selfish Routing and the Price of Anarchy. MIT Press, Cambridge (2005)
12. Papadimitriou, C.: Algorithms, games, and the Internet. In: Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, pp. 749–753 (2001)
13. Youn, H., Gastner, M.T., Jeong, H.: The price of anarchy in transportation networks: Efficiency and optimality control. Physical Review Letters 101, 128701 (2008)
14. Carmi, S., Wu, Z., López, E., Havlin, S., Eugene Stanley, H.: Resource allocation pattern in infrastructure networks. European Physical Journal B 57, 165–174 (2007)
15. Frank, M., Wolfe, P.: An algorithm for quadratic programming. Naval Research Logistics Quarterly 3, 95–110 (1956)
16. Arezki, Y., van Vliet, D.: A Full Analytical Implementation of the PARTAN/Frank Wolfe Algorithm for Equilibrium Assignment. Transportation Science 24, 58–62 (1990)