Network Improvement for Equilibrium Routing

UMANG BHASKAR
University of Waterloo
and
KATRINA LIGETT
California Institute of Technology

Routing games are frequently used to model the behavior of traffic in large networks, such as road networks. In transportation research, the problem of adding capacity to a road network in a cost-effective manner to minimize the total delay at equilibrium is known as the *Network Design Problem*, and has received considerable attention. However, prior to our work, little was known about guarantees for polynomial-time algorithms for this problem. We obtain tight approximation guarantees for general and series-parallel networks, and present a number of open questions for future work.

Categories and Subject Descriptors: F.2.0 [**General**]: Analysis of Algorithms and Problem Complexity—*Approximation Algorithms*

General Terms: Algorithms; Theory; Economics

Additional Key Words and Phrases: Routing games, network design, Wardrop equilibrium

1. INTRODUCTION

Routing games model network traffic in applications where users choose paths in the network to minimize their delay. The cost, or total delay, at equilibria in routing games is very well-studied: tight bounds are known on the price of anarchy, as well as in many cases on techniques to influence player strategies (e.g., tolls, Stackelberg strategies) to improve the total delay at equilibria. In common practical applications, the primary means of improving the total delay is to add capacity to the network. However, techniques to compute the optimal allocation of additional capacity are primarily heuristic or computationally inefficient. Our work addresses the problem of adding capacity to a network under a budget to improve the total delay at the resulting equilibrium, and obtains guarantees for computationally efficient techniques [Bhaskar et al. 2014].

The problem of figuring out where to add capacity in a network is clearly a non-trivial computational problem: Braess' paradox is a prominent example where adding capacity can actually worsen the total delay at equilibrium. Given the importance of this problem for the optimization of road networks, it is unsurprising that it has received considerable attention in the literature on transportation research, where a general version of the problem discussed in this article is called the *Continuous Network Design Problem* [Abdulaal and LeBlanc 1979]. While a number of algorithms for the problem are given in the literature, these algorithms either offer no performance guarantees on the solution obtained, or are highly inefficient computationally (see, e.g., [Yang and H. Bell 1998; Luathep et al. 2011;

Authors' addresses: `umang.bhaskar@uwaterloo.ca`, `katrina@caltech.edu`

Li et al. 2012]). Prior to our work, little was known about the performance of computationally efficient algorithms.

2. PROBLEM DEFINITION

Formally, a routing game is defined by a directed graph $G = (V, E)$, a vector of non-decreasing delay functions on the edges $(l_e)_{e \in E}$, and a set of triples $\{(s_i, t_i, d_i)\}_{i \in [k]}$ where each triple denotes the source, sink, and demand for a commodity. We assume players are *nonatomic*: each player controls infinitesimal traffic. Players of commodity i choose a route from s_i to t_i, and the (infinite) set of such players constitute the demand for commodity i. A strategy profile for the players corresponds to a multicommodity flow f, and given a flow f, each player of commodity i picks an s_i-t_i path P to minimize her delay $\sum_{e \in P} l_e(f_e)$.

A multi-commodity flow $f := (f^i)_{i \in [k]}$ where each f^i is an s_i-t_i flow of value d_i is called a valid flow. A valid flow f is called a Wardrop equilibrium if for all $i \in [k]$ and any s_i-t_i paths P, Q with $f^i_P > 0$, $\sum_{e \in P} l_e(f_e) \leq \sum_{e \in Q} l_e(f_e)$. The Wardrop equilibrium can also be obtained as the solution to a convex program. If all delay functions are strictly increasing, the program is strictly convex, and hence the equilibrium is unique.

To model network improvement, every edge has a capacity $c_e(\beta_e) := c^0_e + \mu_e \beta_e$, where β_e is the money allocated to improving the edge (initially zero), c^0_e is the initial capacity of the edge, and μ_e is the marginal rate of improvement. The latency function on every edge then has the form $l_e(x, \beta_e) = (x/c_e(\beta_e))^{n_e} + b_e$. For a given budget $B \in \mathbb{R}_+$, a valid allocation is a vector $\beta = (\beta_e)_{e \in E} \geq 0$ with $\sum_e \beta_e \leq B$.

For a valid allocation β, let $f(\beta)$ be the resulting equilibrium with delay functions $l_e(x, \beta_e)$. The network improvement problem can then be written as the following mathematical program:

$$\min_{f,\, \beta} \sum_e f_e\, l_e(f_e, \beta_e) \quad \text{s.t.} \quad \beta \geq 0, \quad \sum_e \beta_e \leq B\,, \quad \text{and} \quad f = f(\beta). \qquad (1)$$

In the program above, both the flow f and the allocation β are treated as variables, although the constraints restrict f to be the unique equilibrium flow. Despite being continuous, the objective function in (1) is non-convex. The feasible region may also be disconnected; both of these make the problem difficult to solve. Despite this, our results give tight approximation guarantees for this problem in a number of cases.

3. OUR CONTRIBUTION

We first consider the problem in general graphs. Here, we show that a simple algorithm that operates by relaxing the constraint that f be an equilibrium flow in (1) and solves the resulting convex optimization problem, in fact obtains the best possible approximation ratio for instances with affine delay functions ($n_e = 1$ on all edges).[1]

THEOREM 3.1. *We can obtain in polynomial time a 4/3-approximate allocation for instances with affine delay functions, and an $O(p/\log p)$-approximate allocation*

[1] The results of Theorem 3.1 were also independently obtained by Gairing, Harks and Klimm [Gairing et al. 2014].

for instances with delay functions of degree p. It is NP*-hard to obtain an approximation ratio better than 4/3, even in single-commodity instances with affine delays.*

The somewhat surprising optimality of the simple algorithm described is proved by showing connections to previous work on routing games in algorithmic game theory. The proof of the approximation ratio uses well-known bounds on the price of anarchy [Roughgarden 2003]. The proof of hardness uses techniques developed, motivated by Braess' paradox, for the problem of *removing* capacity in a network to reduce the total delay at equilibrium [Roughgarden 2006].

Given the practical relevance of the problem, an immediate question is if in special cases, the 4/3-approximation can be improved upon. In order to circumvent the inapproximability, it is obvious that we must consider restricted topologies. We hence consider the problem in *series-parallel* graphs.[2] Here, we show that in the single-commodity setting, we can in fact get arbitrarily close to the optimal solution in polynomial time (i.e., get an FPTAS).

THEOREM 3.2. *For single-commodity routing games on series-parallel graphs with polynomial delay functions, for any* $\epsilon > 0$*, we can obtain a* $(1+\epsilon)$*-approximate allocation in time polynomial in the size of the input and* $1/\epsilon$*.*

Our algorithm for Theorem 3.2 uses dynamic programming as well as the structure of series-parallel graphs to obtain the FPTAS. The algorithm consists primarily of three steps. The first step is to modify the mathematical program (1) in two ways. Firstly, we relax the equilibrium constraint in (1) as we did in the proof of Theorem 3.1. Secondly, instead of minimizing the total delay, we modify our objective to minimize the maximum delay over all paths with strictly positive flow. This new objective is a discontinuous function of the flows. However, we show that in series-parallel graphs, the two problems are equivalent: an optimal solution to the second problem is optimal for the original problem as well. The advantage of modifying the problem in this way is to simplify the feasible set. Instead of all valid allocations and flows at *equilibrium*, our feasible set is now the set of all valid allocations and valid flows.

We then show that we can appropriately discretize this simplified feasible set and search in this discretized space using dynamic programming to obtain a near-optimal solution to the modified problem. Finally, we show that a near-optimal solution to the modified problem is a near-optimal solution to the original problem as well. A number of technical issues arise in the discretization, such as edges where the marginal rate of improvement is large (and thus a small allocation significantly improves the capacity). We show that these can be dealt with, and obtain a polynomial-time algorithm.

A consequence of Theorem 3.2 is that restricting network topologies can circumvent the 4/3-inapproximability. However, we show that even in series-parallel graphs, the FPTAS is the best possible: the problem is NP-hard. In fact, we show

[2]Series-parallel graphs are defined recursively. A single edge $e = (s, t)$ is a series-parallel graph, with terminals s and t. Two series-parallel graphs can be combined to obtain a new series-parallel graph, either by merging the s-terminals together and the t-terminals together (a parallel-join) or by merging the t-terminal of one graph with the s-terminal of the other (a series-join).

NP-hardness for a special case of series-parallel graphs, obtained by joining in series subgraphs consisting of two parallel edges. The proof of hardness is via reduction from PARTITION.

THEOREM 3.3. *The network improvement problem is NP-hard, even for single-commodity routing games on series-parallel graphs with affine delay functions.*

4. OPEN PROBLEMS

Our work [Bhaskar et al. 2014] presents tight bounds for the settings we consider. There are many open problems and generalizations that follow, and are of interest both theoretically and practically. We present two questions immediately motivated by applications.

Variable demand. If the capacity of a network were expanded, it is natural to suppose that more traffic would use the network. This is called "induced demand," and is frequently a practical concern when expanding capacity in road networks. The price of anarchy with induced demand has been studied previously [Cole et al. 2012]. An open problem is to develop algorithms for network improvement that address induced demand as well.

Other network topologies. Our reduction that establishes a lower bound of 4/3 on the approximation guarantee is not applicable if we restrict the graph topology, e.g., if the graph is planar. However, road networks are often planar or nearly so. Can a better approximation guarantee be obtained for planar graphs?

More generally, there are numerous problems in economics and engineering that are naturally expressed as mathematical programs with equilibrium constraints (see, e.g.,[Luo et al. 1996]); network improvement is just one such problem. Such problems are computational in nature, and have received considerable attention in fields such as operations research and engineering. We believe that, as for the network design problem, the tools of algorithmic game theory can contribute to a better understanding of these natural and well-motivated problems.

Acknowledgments. K. Ligett gratefully acknowledges support of an NSF CAREER Award (1254169), the Charles Lee Powell Foundation, and a Microsoft Research Faculty Fellowship.

REFERENCES

ABDULAAL, M. AND LEBLANC, L. J. 1979. Continuous equilibrium network design models. *Transportation Research Part B: Methodological 13,* 1, 19–32.

BHASKAR, U., LIGETT, K., AND SCHULMAN, L. J. 2014. Network improvement for equilibrium routing. In *IPCO*. 138–149.

COLE, R., DODIS, Y., AND ROUGHGARDEN, T. 2012. Bottleneck links, variable demand, and the tragedy of the commons. *Networks 60,* 3, 194–203.

GAIRING, M., HARKS, T., AND KLIMM, M. 2014. Complexity and approximation of the continuous network design problem. In *APPROX/RANDOM*. 226–241.

LI, C., YANG, H., ZHU, D., AND MENG, Q. 2012. A global optimization method for continuous network design problems. *Transportation Research Part B: Methodological 46,* 9, 1144–1158.

LUATHEP, P., SUMALEE, A., LAM, W. H., LI, Z.-C., AND LO, H. K. 2011. Global optimization method for mixed transportation network design problem: a mixed-integer linear programming approach. *Transportation Research Part B: Methodological 45,* 5, 808–827.

Luo, Z.-Q., Pang, J.-S., and Ralph, D. 1996. *Mathematical programs with equilibrium constraints.* Cambridge University Press.

Roughgarden, T. 2003. The price of anarchy is independent of the network topology. *J. Comput. Syst. Sci. 67,* 2, 341–364.

Roughgarden, T. 2006. On the severity of braess's paradox: designing networks for selfish users is hard. *Journal of Computer and System Sciences 72,* 5, 922–953.

Yang, H. and H. Bell, M. G. 1998. Models and algorithms for road network design: a review and some new developments. *Transport Reviews 18,* 3, 257–278.