A Stackelberg strategy for routing flow over time ☆

Umang Bhaskar [a],[*],[1],[3], Lisa Fleischer [b],[1], Elliot Anshelevich [c],[2]

[a] *California Institute of Technology, Pasadena, CA 91125, United States*
[b] *Dartmouth College, Hanover, NH 03755, United States*
[c] *Rensselaer Polytechnic Institute, Troy, NY 12180, United States*

ARTICLE INFO

ABSTRACT

Routing games are studied to understand the impact of individual users' decisions on network efficiency. Most prior work on efficiency in routing games uses a simplified model where all flows exist simultaneously, and users care about either their maximum delay or their total delay. Both these measures are surrogates for measuring how long it takes to get all of a user's traffic through the network. We attempt a more direct study of network efficiency by examining routing games in a flow over time model. Flows over time are commonly used in transportation research. We show that in this model, by reducing network capacity judiciously, the network owner can ensure that the equilibrium is no worse than a small constant times the optimal in the original network, for two natural measures of optimality. These are the first upper bounds on the price of anarchy in this model for general networks.

1. Introduction

In routing games, players route a fixed amount of flow in a network. A player suffers a cost that depends on its routing and the routing chosen by the other players. A flow in a routing game is an equilibrium flow if no player can choose a different routing and reduce its cost.

Routing games model a variety of problems, including routing on roads (Altman et al., 2002; Wardrop, 1952), computer networks (Cole et al., 2012; Orda et al., 1993; Roughgarden, 2005), and scheduling tasks on machines (Koutsoupias and Papadimitriou, 1999). For many measures of the quality of a routing, the equilibrium in routing games is known to be inefficient compared to a routing that optimizes the measure. This inefficiency is quantified by the *price of anarchy* (Koutsoupias and Papadimitriou, 1999): the worst ratio of the objective evaluated for the equilibrium flow, to the optimal flow. The price of anarchy thus measures the worst-case loss of efficiency due to the lack of coordination among the players.

The term price of anarchy was introduced in Papadimitriou (2001). Since its introduction, a number of papers have focused on obtaining bounds on the price of anarchy as a measure of worst-case behavior. The price of anarchy has also been analyzed for other games besides routing games, such as resource allocation games (Johari and Tsitsiklis, 2004) and network creation games (Fabrikant et al., 2003). In addition, the study of the price of anarchy has also led to various methods to improve the quality of equilibrium by influencing the behavior of players, such as tolls, coordination mechanisms, and Stackelberg strategies (e.g., Christodoulou et al., 2009; Swamy, 2012).


Fig. 1. Flow in excess of capacity on an edge forms a queue at the tail.

Most prior work on the price of anarchy in routing games uses a simplified model of network flow where a player's cost is either its maximum delay or its total delay. For example, players in a routing game have a *bottleneck* objective if a player's cost is the maximum delay on the edges it uses (Banner and Orda, 2007). The bottleneck objective models applications where a player's cost depends largely on the performance of the worst resource used. In routing games, this objective ignores the effect of delay on edges besides the bottleneck edge, which can lead to the counterintuitive situation where players fail to distinguish between two strategies that have the same bottleneck, but have very different delays. This behavior may result in an unbounded price of anarchy, e.g., Akella et al. (2003), Banner and Orda (2007), Cole et al. (2012). The price of anarchy would be 1 in some of these bad instances if player costs depended on edges in addition to the bottleneck edge. Models where a player takes into account the delays on all edges have an improved price of anarchy (Cole et al., 2012).

Even models where a player's cost is an aggregation of the delay on each edge assume that the flow is static: every edge has flow on it instantaneously and simultaneously and, once established, a flow continues indefinitely. However, the flow in networks is often transient. Flow enters a network, uses it for some time, and then exits, and the flow on an edge changes with time.

This time-varying nature of flows is captured by flows over time, introduced in Ford and Fulkerson (1962). In this model, flow traverses the path in finite time, and exits the network at the sink. Thus, the flow on each edge of the network varies with time. Every edge has a capacity which limits the flow rate on the edge.

Flows over time are widely used for modeling traffic in transportation research. The literature on Dynamic Traffic Assignment (DTA) encompasses various models and approaches for studying traffic using flows that vary over time (see, e.g., Peeta and Ziliaskopoulos, 2001). These models are generally large and complex; and problems such as existence and uniqueness of equilibrium are generally intractable. For this reason, heuristic approaches or simulation are commonly used to obtain solutions in these models.

We analyze routing games for flows over time. In our model, every player controls infinitesimal flow, and the cost of a player is the time at which it arrives at the sink. A player's strategy is a path from the source to the destination. On every edge, the flow follows first-in, first-out (FIFO). Although the network is capacitated, the model allows the inflow on an edge to be larger than the capacity of the edge.[4] The excess flow forms a queue at the tail of the edge, and must wait for the preceding flow to enter before it can enter the edge (Fig. 1). Although the capacities and the edge-delays are fixed, the total delay on an edge varies with the size of the queue on the edge. The queue size seen by flow arriving at the edge varies with time; hence the total delay along any path varies with time.

This game, which we call a *temporal routing game*, follows the model of selfish routing of flows over time used in Koch and Skutella (2011). The model possesses a number of interesting characteristics. It follows FIFO, which is a standard assumption in transportation research literature. The model is based on dynamic queueing, first used in Yagar (1971). Further, the equilibrium flow over time can be characterized in terms of static flows that satisfy certain conditions, described in Section 2 of Koch and Skutella (2011).

In flows over time, similar to static flows, various objectives may be used to compare the performance of the equilibrium flow to the optimal flow. A natural objective is the total delay: the sum of the costs of the players. The flow that minimizes the total delay is the *earliest arrival flow*, which maximizes the flow that arrives at the destination by time θ, for every time θ. For networks with a single sink, earliest arrival flows exist (Baumann and Skutella, 2009). The ratio of the total delay of the worst equilibrium flow to the total delay of earliest arrival flow is the *total delay price of anarchy*. A different objective is the time taken to route a fixed amount of flow to the sink. This is the *completion time*, and is minimized by a *quickest flow*. The earliest arrival flow is also a quickest flow. To route a fixed amount of flow, the ratio of the time taken by the worst equilibrium to the time taken by the quickest flow is the *time price of anarchy*. A third objective is the amount of flow that reaches the destination by time θ. The ratio of amount of flow that reaches the destination by time θ in the worst equilibrium flow to the amount of flow that reaches the sink in the earliest arrival flow is the *evacuation price of anarchy*. A fourth objective is the travel time of any player, i.e., the maximum difference for a player between the time it departs from its source and the time it arrives at its destination.

In the *Stackelberg model* introduced in von Stackelberg (1934), different players in a game have different priorities. A *leader* picks a strategy first, and then the followers pick their strategies. Importantly, the leader commits to a strategy before the followers pick theirs. In our setting, the network manager is the leader. Given some physical limit on the capacity of each edge, the network manager acting as leader picks a capacity for each edge that does not exceed this physical limit. The remaining players, acting as followers, then pick a route from source to sink as their strategy in this modified

---

[4] Thus, unlike the static flow model in Correa et al. (2004), the ability to reduce capacities is not sufficient to enforce the optimal flow.