The Price of Fragmentation in Mobility-on-Demand Services*

Thibault Séjourné
Ecole Polytechnique
France

thibault.sejourné@polytechnique.
edu

Samitha Samaranayake
Cornell University
USA

samitha@cornell.edu

Siddhartha Banerjee
Cornell University
USA

sbanerjee@cornell.edu

ABSTRACT

Mobility-on-Demand platforms are a fast growing component of the urban transit ecosystem. Though a growing literature addresses the question of how to make individual MoD platforms more efficient, less is known about the cost of market fragmentation, i.e., the impact on overall welfare due to splitting demand between multiple independent platforms. Our work aims to quantify how much platform fragmentation degrades the efficiency of the system. In particular, we focus on a setting where demand is exogenously split between multiple platforms, and study the increase in supply rebalancing costs incurred by each platform to meet this demand, vis-a-vis the cost incurred by a centralized platform serving the aggregate demand. We show under a large-market scaling, this Price-of-Fragmentation undergoes a phase transition, wherein, depending on the nature of the exogenous demand, the additional cost due to fragmentation either vanishes or grows unbounded. We provide conditions that characterize which regime applies to any given system, and discuss implications of these findings on how such platforms should be regulated.

CCS CONCEPTS

• **Networks** → **Network performance analysis; Network performance modeling;**

KEYWORDS

Mobility on Demand systems; platform competition; large-market scaling; phase transition

ACM Reference Format:

Thibault Séjourné, Samitha Samaranayake, and Siddhartha Banerjee. 2018. The Price of Fragmentation in Mobility-on-Demand Services. In *SIGMETRICS '18 Abstracts: ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems Abstracts, June 18–22, 2018, Irvine, CA, USA*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3219617.3219623>

1 INTRODUCTION

Traditionally, urban mobility has been satisfied via either personal vehicle ownership (which is convenient, but has high per-passenger congestion/pollution costs), or mass transit systems (which are

efficient, but prone to first-and-last-mile problems). Mobility-on-Demand (MoD) systems (e.g. Uber, Lyft, Didi, Ola) provide a new alternative that preserves the convenience of personal vehicles, while reducing its burdens. Their growing popularity is a testament to their success, and they have the potential to help the overall efficiency of the transportation system as a whole.

The popularity of MoD systems has given rise to a growing literature on fleet management, dispatch and pricing problems in these systems. However, there is limited work in understanding the effects of having multiple service providers (i.e., platforms). The economic benefits of competition for passengers and drivers (in terms of prices, wages, service quality, etc.) are well-established; sustaining competition may however result in higher overall operational costs. In particular, when multiple platforms fragment demand, and myopically optimize to serve their own customers, they may result in a higher cost to society (in terms of pollution/fuel consumption/congestion, etc) as compared to a monopoly.

Our work aims to understand and quantify the increased operational costs due to competition in MoD ecosystems. There are two challenges to doing so: First, since MoD systems involve complex network dynamics involving many vehicles, any analysis requires careful modeling of demand/supply/rebalancing and stochastic dynamics. Secondly, given the growth of MoD systems, the pertinent question is not so much what the costs are currently, but rather, how they will behave as demand scales in future; such counterfactual analysis is not directly accessible from data. Our techniques address both these questions.

To this end, we define a novel metric we call the *Price of Fragmentation* (PoF), that captures how the efficiency of a MoD ecosystem degrades as demand is randomly split between platforms. We focus on a setting where passengers are exogenously split between multiple platforms (i.e. we ignore endogenous competition effects), and study the increase in system costs incurred by each platform to meet this demand, vis-a-vis the cost incurred by a monopolist serving the aggregate demand. This captures the multiplexing loss (i.e., the net increase in vehicle miles traveled when Lyft customers are matched to the nearest Lyft driver although there is a closer Uber driver, and vice versa) due to having competing MoD providers. By studying the scaling behavior of the PoF with increasing demand, we uncover a surprising phase-transition phenomena (see Figure. 1), wherein the PoF either vanishes or grows unbounded, depending on a structural property of the underlying aggregate demand. We corroborate these results via synthetic experiments, as well as data-driven studies using the NYC taxi-cab data [1].

Our work is the first study on the loss in efficiency due to demand fragmentation across multiple MoD platforms. While more work is needed to understand platform competition, we argue that the

*This document is an extended abstract for [3].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '18 Abstracts, June 18–22, 2018, Irvine, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5846-0/18/06.

<https://doi.org/10.1145/3219617.3219623>

systemic inefficiencies uncovered in our work form an important first step for guiding public policy for the MoD ecosystem.

1.1 Modeling the MoD Ecosystem

We consider an MoD ecosystem with either a monopolist platform, or two (or more) competing firms. All firms operate in a city represented as a complete directed graph $G(V, E)$, with nodes corresponding to *stations* or neighborhoods, and edges corresponding to *fastest routes* between any pair of stations. Each edge has an associated *cost* τ_{ij} (in terms of driving time or miles). Firms have independent and exogenous demand, and strive to serve all their demand with equal priority.

Formally, in the monopolist setting, we let $\{\Lambda_{ij}\}_{(i,j) \in E}$ denote the *aggregate demand vector*, where Λ_{ij} denotes the (random) number of customers requesting rides from i to j . The aim of the platform is to minimize the *cost of rebalancing* $RC(\Lambda)$, i.e., the operational cost of rerouting empty supply to meet the demand. This given by the following *min-cost circulation LP*:

$$\begin{aligned} \min \quad & \sum_{(i,j)} \tau_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_j (x_{ji} - x_{ij}) = \Lambda_i \quad \forall i \in V \quad ; \quad x_{ij} \geq 0 \quad \forall (i,j) \in E, \end{aligned} \quad (1)$$

where $\Lambda_i = \sum_j (\Lambda_{ij} - \Lambda_{ji})$ is the net inflow of demand at station i , and x_{ij} is the number of rebalancing trips from station i to j . The formulation of the rebalancing cost in terms of the min-cost circulation (or *optimal transport*) problem arises in many studies of MoD systems, and is in a sense the smallest ex-post cost that a platform can achieve via rebalancing and dispatch to satisfy all the demand. Moreover, recent work has shown that simple dynamic rebalancing policies can match this bound very closely, in particular as the number of vehicles increases [2].

To understand the effects of platform competition, we consider a *duopolist* setting, where for any edge $(i, j) \in E$, the demand is fragmented as $(\Lambda_{ij}^a, \Lambda_{ij}^b)$ between two firms a and b , such that $\Lambda_{ij}^a + \Lambda_{ij}^b = \Lambda_{ij}$. We assume the split is random and exogenously determined (i.e., unaffected by firms' strategies): as a canonical example, let Λ_{ij} be integral and given *market-share* $\rho \in [0, 1/2]$ of firm a , we assume that firm a receives $\Lambda_{ij}^a \sim \text{Binomial}(\Lambda_{ij}, \rho)$ requests, and firm b receives the remaining $\Lambda_{ij}^b = \Lambda_{ij} - \Lambda_{ij}^a$ requests. Each firm then tries to minimize its own costs, leading to a total rebalancing cost that is greater than the monopolist cost. We quantify this by defining the *Price of Fragmentation* (PoF) to be the expected excess rebalancing cost incurred by the duopoly. Finally, to understand the change in the PoF as MoD platforms grow, we analyze the system under the so-called *large-market scaling*, wherein we scale all the incoming demands by $\theta \in \mathbb{N}$.

Formally, let (Λ^a, Λ^b) denote the random vector of demands for the two firms. Now, under the large-market scaling with demand splits $\Lambda^{a,\theta} \sim \text{Binomial}(\theta\Lambda, \rho)$, and $\Lambda^{b,\theta} = \theta\Lambda - \Lambda^{a,\theta}$ (for $\theta \in \mathbb{N}$), we define the *Price of Fragmentation* as follows:

Definition 1.1. (Price of Fragmentation) Given demand Λ and market-share ρ , the Price of Fragmentation is given by:

$$\gamma^\theta \triangleq \mathbb{E} \left[RC(\Lambda^{a,\theta}) + RC(\Lambda^{b,\theta}) \right] - RC(\theta\Lambda)$$

Figure 1: Simulation demonstrating the scaling behavior of the Price of Fragmentation (the difference between the rebalancing cost of the duopoly and the monopoly) in the fragmentation-resilient (green) and fragmentation-affected (blue) regimes. The demand distributions and costs are based the NYC taxi data; in particular, the two regimes shown here correspond to two successive hours (noon to 1 p.m and 1 to 2p.m.) on May 10, which have have similar PoF at current demand, but for which our analysis suggests different scaling behavior.

To get some insight into the PoF, observe that $RC(\cdot)$ is convex, and hence Jensen's inequality implies that $\gamma \geq 0$; in other words, competition necessarily reduces system efficiency under stochastic splitting. More precisely, since $\mathbb{E}[\Lambda^{a,\theta}] = \rho\theta\Lambda$, we get that:

$$\mathbb{E}[RC(\Lambda^{a,\theta}) + RC(\Lambda^{b,\theta})] \geq RC(\rho\theta\Lambda) + RC((1-\rho)\theta\Lambda) = RC(\theta\Lambda)$$

Since $RC(\cdot)$ is continuous and $\Lambda^{a,\theta}$ is binomial (and hence, has sub-Gaussian tails), the law of large numbers guarantees that $\gamma^\theta/\theta \rightarrow 0$. Our main results characterize the rate of this convergence; in particular, we show that depending on the underlying demand vector, the stochastic $\text{PoF}(\theta) = \gamma^\theta$ either decays to 0 or grows unbounded. We refer to the former as the *fragmentation-resilient* regime, and the latter as the *fragmentation-affected* regime.

1.2 Summary of our Results

For ease of presentation, we focus here on the case of homogeneous market shares, wherein the average expected fraction of demand going to each firm is the same across the network; in our full paper, we extend our results to heterogeneous market shares and to any number of companies. It is important to reiterate that our setup assumes an independent and sufficient vehicle supply, and exogenous demand splitting. In other words, we assume that consumers have idiosyncratic demand preferences and we do not model competition between the platforms. While we acknowledge that this is an important consideration for real world systems, the insights we gain under our simpler model are by themselves significant and non-obvious, and will likely continue to hold under more complex models as well.

Our results are summarized as follows: for general networks under the large-market scaling, we show that the *PoF* undergoes a *phase transition*, depending on the structure of the underlying demand vector. In the first regime the PoF decays exponentially under scaling

Figure 2: Probability of observing fragmentation-affected demand in NYC based on the NYC taxi dataset. Fig. 2(a) shows the influence of the number of stations, while Fig. 2(b) demonstrates the influence of the length of the time windows. Confidence intervals are plotted based on aggregation of the data into time windows of length 2 hours.

– we refer to this as the *fragmentation-resilient* regime. In contrast, in the second regime, the PoF grows under scaling, at a rate equal to the square root of the scaling parameter – we refer to this as the *fragmentation-affected* regime. Figure. 1 shows an example of these two cases, based on demand distributions estimated using the NYC taxi data over two different one-hour time slots; note that though both the settings appear to have identical PoF at the original demand, the two curves are markedly different as we scale the average demands.

More specifically: our results show that the scaling behavior of the PoF is determined by structural features of the minimum-cost circulation problem in Eq. 1 for the rebalancing cost $RC(\Lambda)$ under the aggregate demand. In particular, we show that under homogeneous demand splitting, the system is in the *fragmentation-affected regime* if and only if the aggregate demand vector is dual-degenerate with respect to the min-cost circulation program in Eq. 1. Our results also extend beyond Binomial splits to more general demand-fragmentation processes, and also to non-homogeneous splits, where the behavior may be more complex.

Though our results provide an explicit characterization of the phase transition, they are less interpretable when stated in terms of structural properties of the optima of an LP. To this end, we also provide an elegant combinatorial characterization for the fragmentation-affected regime. Our characterization essentially shows that the transition from fragmentation-resilient to fragmentation-affected regimes occurs when the demand induces *local balanced clusters*. Intuitively, the inefficiency arises when the demand is such that the network can be partitioned in two or more clusters with (roughly) balanced underlying demand going to and fro between clusters, while the cost of traveling between clusters is high.

Though one can create artificial scenarios that lead to such balanced partitions, it is unclear if they exist in practice. To this end, we provide empirical evidence using the NYC taxi data, wherein we show that local balanced clusters do often arise in practice. Our experiments use taxi data recorded for the months of May and June 2016. The dataset has about 335,000 trips per day; on an hourly basis, the number of trips ranges from 3,600 to 20,000 with an

average number of around 14,000 thousands trips per hour. We model the city as consisting of a finite number of pickup and dropoff stations (20-80 in our experiments), where we assume each trip begins and ends. The stations are determined by taking all of the user demand for the full period of study, clustering the data (pickup and dropoff locations) into the desired number of stations (clusters), and picking the centroid for each cluster. For each time window of interest, the demands are then aggregated to the total demand per origin-destination station pair Λ_{ij} . For convenience, we use the distance d_{ij} as a proxy for the travel-times τ_{ij} to parametrize edge-costs; we compute the Manhattan distance between the stations to get the distances d_{ij} . Figure 2 provides an example of our simulation results, wherein we use our characterization to determine the fraction of time-slots where the underlying demand profile was fragmentation-affected. For more, refer to our full version [3].

Finally, we discuss implications of our results on how MoD ecosystems should be regulated, as well as extensions of our work. In particular, we argue that our work highlights the importance of having a significant population of *multi-homing drivers*, i.e., those working simultaneously on multiple platforms. Our results also suggest that demand aggregation at fixed *pickup points* can help reduce the PoF. Finally, our work highlights the importance of modeling, as phase-transition phenomena can lead to erroneous findings in data-driven observational studies.

ACKNOWLEDGMENTS

SB acknowledges support from Army Research Office grant W911NF-17-1-0094.

REFERENCES

- [1] 2017. NYC Taxi Data. (2017). http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml
- [2] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. 2017. Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 517–517.
- [3] Thibault Séjourné, Samitha Samaranyake, and Siddhartha Banerjee. 2018. The Price of Fragmentation in Mobility-on-Demand Services. *arXiv preprint arXiv:1711.10963* (2018).