Optimizing Pricing, Repositioning, En-Route Time, and Idle Time in Ride-Hailing Systems

Anton J. Kleywegt[a,*], Hongzhang Shao[a]

[a] School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332-0205, USA
anton@isye.gatech.edu, steveshao@gatech.edu
* Corresponding author

Extended abstract submitted for presentation at the 11th Triennial Symposium on Transportation Analysis conference (TRISTAN XI)
June 19-25, 2022, Mauritius Island

April 4, 2022

Keywords: ride-hailing, transportation network companies, optimal pricing, optimal repositioning

1 RIDE-HAILING OPERATIONS

In ride-hailing systems, available vehicle time is one of the most important resources. Vehicle time can be partitioned into four types of activities:

1. The time used to transport riders from their origins (pickup locations) to their destinations (dropoff locations). This is the major utility (and revenue) generating activity of ride-hailing systems. This time is called on-trip time.
2. The time used to reposition empty from one location to another location. In most transportation systems, travel demand is not balanced over space over time scales of the order of the duration of a trip. Imbalance in travel demand can be mitigated somewhat by pricing incentives, but most imbalance in travel demand is accommodated either by repositioning of vehicles without riders or by parking of vehicles. Under typical costs, repositioning of vehicles is preferred over parking of vehicles, and therefore, although repositioning does not generate revenue, it is an essential activity in the operation of most ride-hailing systems.
3. The time that elapses from the moment a vehicle is dispatched to pick up a rider until the rider is picked up. This time is called en-route time. Typically, the farther a vehicle has to travel from its location when it is dispatched to the rider's pickup location, the longer the en-route time.
4. The time that a vehicle waits to be dispatched. This time is called idle time.

En-route time and idle time are both in some sense unproductive uses of vehicles' time, and therefore one would like to minimize en-route time and idle time. A fundamental phenomenon in ride-hailing systems is that there is a trade-off between en-route time and idle time — if one of these times is reduced, the other time increases. In short, if vehicles spend little time waiting idle for a dispatch, then few vehicles are available when a rider makes a request, and thus the mean distance between a rider and the closest available vehicle is long, which means that en-route time is long. This trade-off was pointed out by Arnott (1996), who showed for a stylized setting with pickup locations and vehicles uniformly distributed over a space without boundaries that the mean en-route distance is proportional to the inverse square root of the density of idle vehicles. This phenomenon is of great importance in ride-hailing, because en-route time increases rapidly as the number of idle vehicles decreases, and every minute that a vehicle spends en-route is one minute less that the vehicle can transport riders. In spite of this, much of the existing literature on price optimization for ride-hailing Banerjee *et al.* (2016a,b), Bimpikis *et al.* (2019), and on repositioning optimization for ride-hailing Braverman *et al.* (2019), ignores en-route time — it is implicitly assumed that en-route time is zero whatever number of idle vehicles are available. Exceptions include Castillo *et al.* (2017) and Xu *et al.* (2020), that consider the trade-off in the stylized setting with pickup locations and vehicles uniformly distributed over a space without boundaries mentioned above.

2 A NEW MODEL OF RIDE-HAILING OPERATIONS

We consider a continuous time, infinite horizon Markov decision process (MDP) model of a ride-hailing system with average profit per unit time objective. Space is partitioned into zones. Each ride request is associated with a pickup zone and a dropoff zone. Ride requests for each origin-destination pair arrive according to a Poisson process. Instead of just considering the mean en-route time for each zone, we consider the distribution of en-route time for each zone which depends on the current number of idle vehicles in the zone. More specifically, for each zone we choose a finite number of en-route time classes. The probability that the en-route time of a ride request belongs to class k depends on the current number of idle vehicles in the zone of the ride request. Given that the en-route time of a ride request belongs to class k, the en-route time of the ride request is exponentially distributed (because it is a continuous time MDP) with mean $1/\nu_k$. Thus the distribution of en-route time for each zone is approximated by a mixture of exponential distributions that depends on the current number of idle vehicles in the zone. The state of the MDP includes information about the number of vehicles on-trip between each origin-destination pair, the number of vehicles repositioning between each origin-destination pair, the number of vehicles en-route in each pickup zone for each en-route time class, and the number of vehicles idle in each zone. This gives a more accurate model than models that use only the mean en-route time (and especially models that assume that en-route time is zero), and it also results in a more tractable fluid optimization model than for models that use only the mean en-route time.

The decisions in our model include both origin-destination pricing decisions, as well as repositioning decisions. This is in contrast with previous work that considered pricing without repositioning Banerjee *et al.* (2016a,b), Castillo *et al.* (2017), Bimpikis *et al.* (2019), or repositioning without pricing Braverman *et al.* (2019).

The MDP is intractable, partly due to the large state space. A widely used approach to develop approximately optimal (and under appropriate conditions, asymptotically optimal) policies, is to formulate and solve a deterministic fluid optimization problem associated with the MDP, and then to use an optimal solution of the deterministic fluid optimization problem to compute a policy for the MDP. Such a deterministic fluid optimization problem is obtained by replacing all random variables by their means, and by allowing discrete variables to take fractional values (hence the name fluid optimization problem). An issue is that even the deterministic fluid optimization problem obtained in the usual way, is intractable for the MDP described above. As mentioned above, instead of replacing the en-route time by its mean, we approximated the distribution of the en-route time. We showed that the resulting stochastic fluid optimization problem can be solved in polynomial time, by solving an associated conic optimization problem.

The solution of the stochastic fluid optimization problem can be used in various ways to compute a policy for the MDP. A simple alternative is the following static (open loop) policy: The optimal prices for the stochastic fluid optimization problem are used as static prices in the MDP. The optimal repositioning flows for the stochastic fluid optimization problem are used to compute repositioning probabilities for the MDP. Another alternative is the following state-dependent (closed loop) policy: The optimal prices for the stochastic fluid optimization problem are used as static prices in the MDP. Periodically, the state of the MDP is used as input to a linear program that determines the optimal repositioning decisions to move the state to the optimal state for the stochastic fluid optimization problem while satisfying flow balance constraints.

We also consider an extension of the models described above that take the following into account. In practice, vehicles can be (and many are) dispatched while repositioning. For example, while a vehicle is repositioning from zone A to zone B, it may pass through zone C and be dispatched to pickup a rider in zone C. In that case, it does not complete its planned repositioning move. As far as we are aware, this is the first work that incorporates this important feature.

3 NUMERICAL RESULTS

A number of policies are compared in terms of long-run average profit per unit time and in terms of their success in controlling the number of available vehicles. The static and state-dependent policies described above are based on the stochastic fluid optimization problem FP_1 that takes the dependence of the en-route time distribution on the number of idle vehicles into account. To demonstrate the importance of repositioning, we also present results for a policy based on a similar stochastic fluid optimization problem that takes the dependence of the en-route time distribution on the number of idle vehicles into account, but without repositioning. To demonstrate the importance of taking the dependence of the en-route time distribution on the number of idle vehicles into account, we present results for a static policy based on a similar fluid optimization problem FP_2, that ignores en-route time, as is done in much of the existing literature on pricing and repositioning for ride-hailing systems.

Figure 1 and Figure 2 present numerical results for three instances (different instances have different origin-destination demand rates and different origin-destination travel times) of a five-zone city with 100 vehicles (20 in each zone at initialization). The instances are the same as those used in Braverman *et al.* (2019). Each column of plots in the figures corresponds to a specific policy, while each row of plots corresponds to a specific instance. We generated 10 sample paths for each combination of policy and instance. Each sample path consisted of 20,000 transitions of the continuous-time Markov chain under the considered policy, with the first 10,000 transitions discarded as transient (warm-up), and the second 10,000 transitions used for result collection. Each black dot in the plots corresponds to one of these sample paths. The dashed lines in Figure 1 show the average profit per unit time of the corresponding stochastic fluid optimization problem, which is an upper bound on the long-run average profit per unit time of any policy for the corresponding MDP. The dashed lines in Figure 2 show the optimal number of available vehicles in each zone according to the corresponding stochastic fluid optimization problem.

Figure 1 shows that: (1) Without repositioning, the objective values are low; even the upper bounds on the objective values are low. (2) The static repositioning policy that takes the dependence of the en-route time distribution on the number of idle vehicles into account performs well, and the simulated revenues are quite close to the upper bounds given by FP_1. (3) The state-dependent repositioning policy performs slightly better than the static repositioning policy, and the corresponding simulated revenues are even closer to the upper bounds. (4) The upper bounds given by FP_2 are larger and much looser than those of FP_1, i.e. ignoring en-route time leads to more optimistic upper bounds. Also, with high and/or imbalanced demand as is typical in practice (Instance 1 and Instance 3), the policy that takes the dependence of the en-route time distribution on the number of idle vehicles into account performs much better than the policy that ignores this dependence. With relatively low and balanced demand (Instance 2), the performance of the policy resulting from FP_2 is comparable with that from FP_1.

Figure 2 shows that under both the static and the state-dependent policy based on FP_1, the numbers of available vehicles are much better controlled than under the other policies, and are close to the optimal numbers given by FP_1. The state-dependent policy based on FP_1 gives slightly better results than the static policy based on FP_1.

Figure 1 – *Simulated average profits per unit time under different policies (columns) for three instances (rows).*

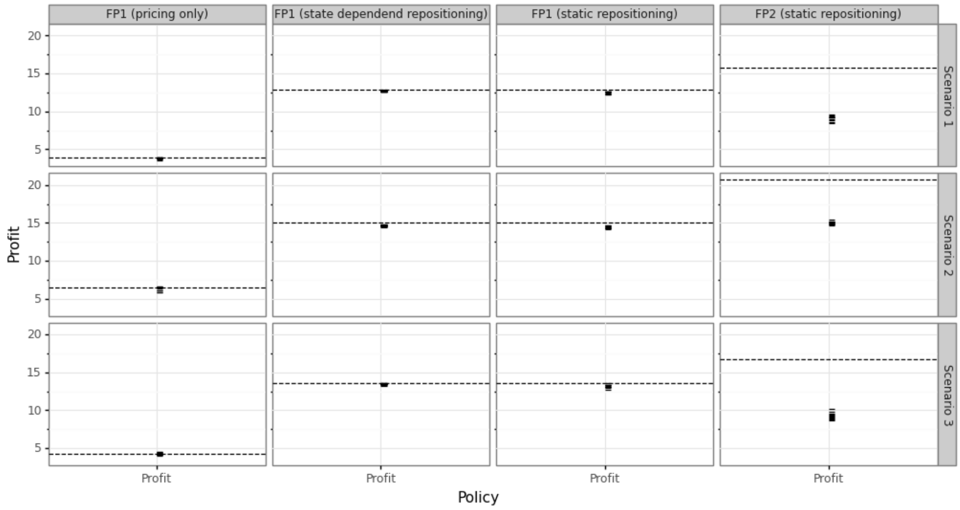

Figure 2 – *Average number of idle vehicles at different zones (x-axis), under different policies (columns) for three instances (rows).*

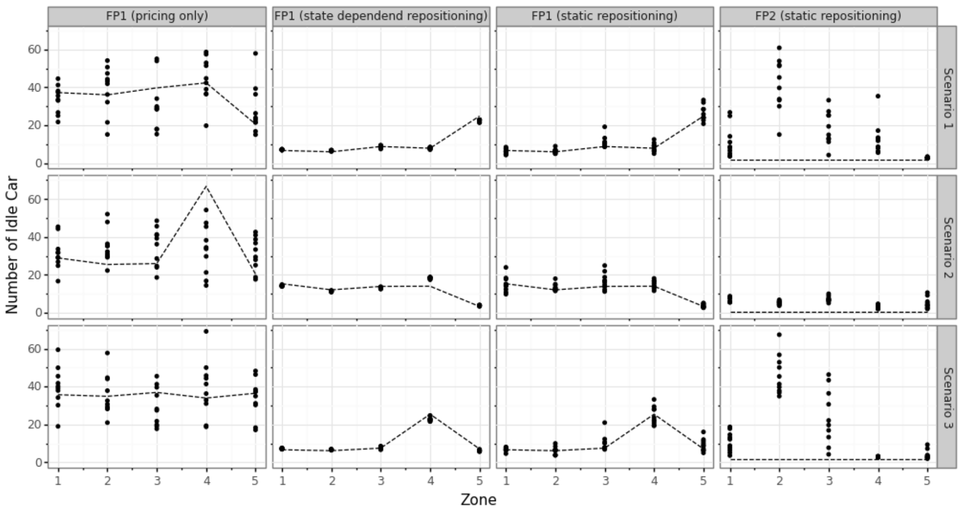

References

Arnott, R. 1996. Taxi Travel Should be Subsidized. *Journal of Urban Economics*, **40**(3), 316–333.

Banerjee, Siddhartha, Freund, Daniel, & Lykouris, Thodoris. 2016a. Multi-Objective Pricing for Shared Vehicle Systems. *arXiv preprint arXiv:1608.06819*.

Banerjee, Siddhartha, Freund, Daniel, & Lykouris, Thodoris. 2016b. Pricing and Optimization in Shared Vehicle systems: An Approximation Framework. *arXiv preprint arXiv:1608.06819*.

Bimpikis, Kostas, Candogan, Ozan, & Saban, Daniela. 2019. Spatial Pricing in Ride-Sharing Networks. *Operations Research*, **67**(3), 744–769.

Braverman, Anton, Dai, Jim G, Liu, Xin, & Ying, Lei. 2019. Empty-car Routing in Ridesharing Systems. *Operations Research*, **67**(5), 1437–1452.

Castillo, Juan Camilo, Knoepfle, Dan, & Weyl, Glen. 2017. Surge Pricing Solves the Wild Goose Chase. *Pages 241–242 of: Proceedings of the 2017 ACM Conference on Economics and Computation.*

Xu, Zhengtian, Yin, Yafeng, & Ye, Jieping. 2020. On the Supply Curve of Ride-Hailing Systems. *Transportation Research Part B: Methodological*, **132**, 29–43.