On the design of crowdsourced delivery fleets: Strategic decisions and operational implications

J. Luy[a,*], G. Hiermann[a], M. Schiffer[a,b]

[a] Technical University of Munich, School of Management, Munich, Germany
[b] Munich Data Science Institute, Technical University of Munich, Munich, Germany
{julius.luy,gerhard.hiermann,schiffer}@tum.de
* 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 3, 2022

Keywords: Crowdsourced delivery, dynamic programming, stochastic workforce planning

1 Introduction

Logistics service providers (LSP) increasingly use crowdsourced workforce on the last-mile to fulfill customers' expectations regarding same-day or on-demand delivery. Crowdsourced couriers, typically independent contractors, flexibly decide when and how much to work (see, e.g., Amazon Flex or DoorDash), thereby introducing uncertainty to an LSP's strategical and operational planning. Accordingly, current research on crowdsourced deliveries focuses on developing algorithms for operational decision-making under supply uncertainty. Both, static (see, e.g., Archetti *et al.* (2016)) and dynamic approaches (see, e.g., Arslan & Zuidwijk (2016)) exist and consider mostly two different crowdsourced driver types: *gigworkers* and *occasional drivers*. The former are primarily financially incentivized and the work for the LSP constitutes a significant amount of their income. The latter are financially independent of the LSP and typically only accept deliveries that overlap with routes they intend to take anyway. To mitigate supply uncertainty, LSPs rely on full-time employees, in the following called *fixed drivers*, to compensate for periods with low availability of crowdsourced drivers. This yields a planning problem with a partially uncertain workforce, which has been analyzed in the context of shift scheduling (Ulmer & Savelsbergh, 2020).

So far, the strategic workforce planning question of how many fixed drivers to hire for a certain time horizon, while anticipating operational implications, has not been studied in a crowdsourced delivery context. Answering this question entails a trade-off between the costs of hiring a new employee and relying on crowdsourced workforce. Against this background, we study a novel approach that integrates the strategic workforce planning problem with its operational implications. We consider an LSP that operates a mixed fleet consisting of fixed drivers, gigworkers and occasional drivers in an area $\mathcal{M}$, knowing the expected demand at every location $i \in \mathcal{M}$. The LSP has to make decisions on two planning levels: On the *strategic* level, she has to decide on the number of fixed drivers to hire in every time step t over a time horizon $\mathcal{T}$. On the *operational* level, defined on time horizon $\widetilde{\mathcal{T}}$, she has to find an optimal routing policy for her fixed drivers while outsourcing requests to crowdsourced drivers. We find that integrating these two planning levels and taking into account future crowdsourced driver capacity in the workforce planning problem yields a significant total cost reduction compared to an approach that only considers immediate operational implications.

2 Methodology

This section details the introduced problem setting. In a first step we describe the strategic and in a second step the operational level.

2.1 Strategic level

We formulate the decision making problem under uncertainty faced by the LSP as a MARKOV Decision Process. We describe the state of the fleet S_t in every time step $t \in \mathcal{T}$ by $S_t = \left(n_t^{\mathrm{FD}}, n_t^{\mathrm{GW}}, n_t^{\mathrm{OD}}\right) \in \mathbb{N}_0^3$. State variable n_t^{FD} represents the number of fixed drivers (FD), n_t^{GW} the number of gigworkers (GW), and n_t^{OD} the number of occasional drivers (OD). We denote the total costs of the LSP in time step t with C_t^{tot}. These consist of fixed costs C^{fix} per fixed driver and time step and of operational costs C_t^{ops}, that are evaluated on the operational level:

$$C_t^{\mathrm{tot}}(n_t^{\mathrm{FD}}, n_t^{\mathrm{GW}}, n_t^{\mathrm{OD}}) = C_t^{\mathrm{ops}}(n_t^{\mathrm{FD}}, n_t^{\mathrm{GW}}, n_t^{\mathrm{OD}}) + C^{\mathrm{fix}}\, n_t^{\mathrm{FD}}. \tag{1}$$

At the beginning of each time step t the LSP decides how many fixed drivers to add to her fleet. We define the action space by $\mathcal{A} = \left\{0, ..., a_{\max}^{\mathrm{FD}}\right\} \in \mathbb{N}_0$ with $a_{\max}^{\mathrm{FD}}$ being the maximum fixed driver fleet size increment. Upon the hiring decision, we evaluate the operational problem and subsequently the joining and resignation process of crowdsourced drivers. The joining process models new private individuals intending to work as crowdsourced drivers for the LSP, starting in time step t. We describe the joining process by a fleet size increase rate q^α ($\alpha \in \{\mathrm{GW}, \mathrm{OD}\}$), expressed as percentage of n_t^α, and model it as a normally distributed random variable. The resignation process describes crowdsourced drivers that decide not to work for the LSP anymore, e.g., because they found a better outside option. We denote the probability of drivers of type α to resign at the end of time step t with p^α and the total number of drivers of type α resigning with X^α. The resulting number of drivers in the next time step then reads

$$n_{t+1}^\alpha = n_t^\alpha - X^\alpha + q^\alpha \cdot n_t^\alpha. \tag{2}$$

Starting in S_0, the LSP's objective is to minimize expected future total costs over $\mathcal{T}$

$$v(S_0) = \min_{\pi \in \Pi} \mathbb{E}\left[\sum_{t=0}^{t_{\max}} \gamma^t \cdot C_t^{\mathrm{tot}}\left(n_t^{\mathrm{FD}}, n_t^{\mathrm{GW}}, n_t^{\mathrm{OD}}\right) | S_0\right]. \tag{3}$$

Here Π denotes the set of all policies, π refers to a single policy and γ is the discount factor. We solve this problem with backward dynamic programming.

2.2 Operational level

The problem on the operational level consists of finding optimal routing policies for the fixed drivers, while possibly outsourcing some requests to crowdsourced drivers. We define the operational problem on a different time horizon $\tilde{\mathcal{T}}$ and we denote its time variable with $\tilde{t}$. The strategic level's time variable t can be segmented into N representative time intervals. The operational problem is then solved for each interval, yielding $C_{t,n}^{\mathrm{ops}}$ for $n \in \{1, ..., N\}$. The resulting $C_{t,n}^{\mathrm{ops}}$ is obtained via $C_{t,n}^{\mathrm{ops}} = \sum_n d_n \cdot C_{t,n}^{\mathrm{ops}}$, where d_n is the number of occurrences of time interval n in t.

We face a pick-up and delivery problem, where requests arise at some location $i \in \mathcal{M}$ and need to be delivered to a destination $j \in \mathcal{M}$. We denote with $E_{ij}(\tilde{t})$ the number of FDs driving from location $i \in \mathcal{M}$ to location $j \in \mathcal{M}$ without serving a request. The variable $E_{ii}(\tilde{t})$ describes waiting FDs at location i. Further, $F_{ij}(\tilde{t})$ describes the number of FDs delivering a request from i to j. Requests follow a POISSON arrival process, with arrival rate λ_i^R. Upon arrival at location i, a request is matched to an FD, GW, or OD, or not matched at all. We denote routing costs by c_{ij}^β ($\beta \in \{\mathrm{FD}, \mathrm{GW}, \mathrm{OD}, \emptyset\}$), where $c_{ij}^\emptyset$ is a penalty that arises when a request cannot be served and leaves the system. While routing costs for FDs and GWs depend on the trip length, ODs receive a fixed compensation per request. Similarly, crowdsourced driver arrivals follow a Poisson process with rate λ_i^{α}, where λ_i^{GW} depends on n_t^{GW} and the distribution of requests. The arrival rate λ_i^{OD} depends on n_t^{OD} and area specific mobility patterns.

We denote with Q the empty FD-routing policy, i.e., a decision rule determining whether an FD should remain at station i or drive without freight to a location j. The objective is to find a policy Q that minimizes total routing costs

$$\min_{Q} \sum_{i,j\in\mathcal{M}} \left[E_{ij}(\tilde{t}) \cdot c_{ij}^{\mathrm{FD}} + \sum_{\beta} c_{ij}^{\beta} \cdot A_{i\beta}(\tilde{t}) \right] \qquad \beta \in \{\mathrm{FD}, \mathrm{GW}, \mathrm{OD}, \emptyset\}\,. \tag{4}$$

The variable $A_{i\beta}(\tilde{t})$ denotes the number of requests matched to option β at station i. This problem becomes intractable for large fleets sizes. Therefore, we follow a fluid approximation approach based on the results from Braverman *et al.* (2019). We interpret the system as a closed queueing network with $|\mathcal{M}|$ single server queues representing FDs not carrying load and waiting at station i (E_{ii}) and service rate λ_i^R, $|\mathcal{M}|^2 - |\mathcal{M}|$ infinite server queues representing fixed drivers driving from i to j (E_{ij}) without serving a request, and $|\mathcal{M}|^2$ infinite server queues for fixed drivers serving a request from i to j (F_{ij}) with service rate μ_{ij}. We represent with e_{ij} and f_{ij} the fluid scaled queue lengths E_{ij} and F_{ij}, i.e., $e_{ij} = \frac{E_{ij}(\tilde{t}\to\infty)}{n_t^{\mathrm{FD}}}$ and $f_{ij} = \frac{F_{ij}(\tilde{t}\to\infty)}{n_t^{\mathrm{FD}}}$. We denote with a_i^{β} the fraction of requests served by fixed drivers, by crowdsourced drivers, or unserved requests. The terms P_{ij}^{R}, P_{ij}^{G} and P_{ij}^{O} describe request, GW, and OD route patterns respectively. The operational problem is then formulated as the following linear program:

$$\min_{e,f,a^{\beta}} \sum_{i} \sum_{j} \left[\sum_{\beta} \left(c_{ij}^{\beta} \cdot \lambda_i^{\mathrm{R}} \cdot a_i^{\beta} \cdot P_{ij}^{\mathrm{R}} \right) + c_{ij}^{\mathrm{FD}}\, n_t^{\mathrm{FD}}\, e_{ij} \right] \tag{5}$$

$$(\lambda_i^{\mathrm{R}}/n_t^{\mathrm{FD}}) \cdot a_i^{\mathrm{FD}} \cdot P_{ij}^{\mathrm{R}} = \mu_{ij} \cdot f_{ij} \qquad \forall i,j \in \mathcal{M} \tag{6a}$$

$$\lambda_i^{\mathrm{R}} \cdot a_i^{\mathrm{GW}} \cdot P_{ij}^{\mathrm{R}} \le \lambda_i^{\mathrm{GW}} \cdot P_{ij}^{\mathrm{GW}} \qquad \forall i,j \in \mathcal{M} \tag{6b}$$

$$\lambda_i^{\mathrm{R}} \cdot a_i^{\mathrm{OD}} \cdot P_{ij}^{\mathrm{R}} \le \lambda_i^{\mathrm{OD}} \cdot P_{ij}^{\mathrm{OD}} \qquad \forall i,j \in \mathcal{M} \tag{6c}$$

$$\mu_{ij}\, e_{ij} \le \sum_{k} \mu_{ki} f_{ki}, \quad i \ne j \qquad \forall i,j \in \mathcal{M} \tag{6d}$$

$$\sum_{k,\,k\ne i} \mu_{ki}\, e_{ki} \le (\lambda_i^{\mathrm{R}}/n_t^{\mathrm{FD}})\, a_i^{\mathrm{FD}} \le \sum_{k,\,k\ne i} \mu_{ki}\, e_{ki} + \sum_{k} \mu_{ki}\, f_{ki} \qquad \forall i \in \mathcal{M} \tag{6e}$$

$$(\lambda_i^{\mathrm{R}}/n_t^{\mathrm{FD}})\, a_i^{\mathrm{FD}} + \sum_{j,\,j\ne i} \mu_{ij}\, e_{ij} = \sum_{k,\,k\ne i} \mu_{ki}\, e_{ki} + \sum_{k} \mu_{ki}\, f_{ki} \qquad \forall i \in \mathcal{M} \tag{6f}$$

$$a_i^{\mathrm{FD}} + a_i^{\mathrm{GW}} + a_i^{\mathrm{OD}} + a_i^{\emptyset} = 1 \qquad \forall i \in \mathcal{M} \tag{6g}$$

The objective describes the minimization of operational costs arising from delivering requests and empty-FD-routing. Constraints (6a) to (6c) describe Little's law of flow conservation. Constraint (6d) is the relaxed Little's law, stating that the mass of out-going empty FDs in one direction j at one location i cannot be higher than the mass of incoming full FDs. Constraints (6e) and (6f) state equality between the total rate of outflow from location i to the total rate of inflow. Finally, constraint (6g) ensures that at every location i a request is either matched or not matched. The objective of this linear program yields the approximated operational cost of a representative time interval, which can then be used to compute C_t^{ops} as input for the strategic level's problem.

3 Preliminary results

In our preliminary experiments, we study an LSP with an initial fleet of ten FD, one GW, and one OD. We consider costs per request of 4$ and 2$ respectively for GWs and ODs; FDs receive a fixed wage of 20\$/h. Furthermore, we consider a joining rate with distribution $\mathcal{N}(0.3, 0.05)$ to model an average increase of the crowdsourced driver fleet of 30% per time step. We choose a resignation probability of $p^{\alpha} = 0.1$, which seems to be a reasonable choice when comparing to established crowdsourced platforms, e.g., Uber (Hall & Krueger, 2018). The area consists of 6 locations, and we consider a demand of on average 6 requests per location per hour for the operational level. We choose a time horizon of 30 weeks on the strategic level to account for the long-term costs of hiring fixed drivers and assume constant demand levels over the strategic time horizon. To benchmark our approach, which integrates future crowdsourced driver capacity via backward dynamic programming (BDP), we compare it with a myopic approach, that bases hiring decisions only on immediate operational costs.

Figure 1 shows the accumulated total cost over time for both the BDP and the myopic approach. One can see that the hiring policy drawn from the BDP approach leads to lower total costs at the end of the time horizon. In the beginning, however, total costs for BDP are higher since it does not hire as many FDs as the myopic approach and has to pay penalties on not served requests. Figure 2 shows the accumulated total costs in the last time step. The two right bars result from a different study, in which the degree of synchronization of OD-specific mobility patterns and request patterns is lowered. The second case's higher costs emphasize the influence of mobility patterns on the success of crowdsourced delivery services.

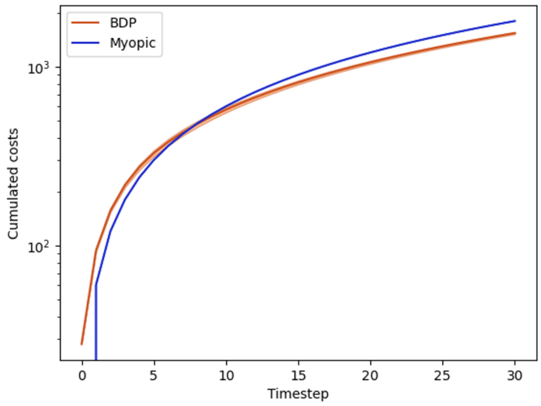

Figure 1 – *Accumulated total costs C_t^{tot} over time (based on 1000 runs)*

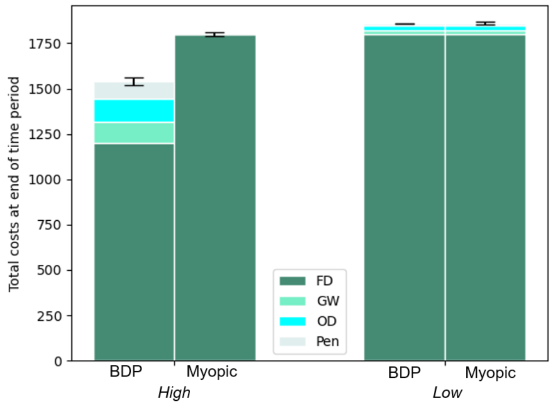

Figure 2 – *High vs. low degree of synchronization (based on 1000 runs)*

References

Archetti, C., Savelsbergh, M., & Speranza, G. 2016. The Vehicle Routing Problem with Occasional Drivers. *European Journal of Operational Research*, **254**(2), 472–480.

Arslan, A.M., & Zuidwijk, R. 2016. Crowdsourced Delivery, A Dynamic Pickup and Delivery Problem with Ad Hoc Drivers. *Transportation Science*, **53**(1), 472–480.

Braverman, A., Dai, J. G., Liu, X., & Ying, L. 2019. Empty-Car Routing in Ridesharing Systems. *Operations Research*, **67**(5), 1209–1502.

Hall, J. V., & Krueger, A. B. 2018. An Analysis of the Labor Market for Uber Driver-Partners in the United States. *ILR Review*, **71**(3), 705–732.

Ulmer, M. W., & Savelsbergh, M. 2020. Workforce Scheduling in the Era of Crowdsourced Delivery. *Transportation Science*, **54**(4), 1113–1133.