Food-Delivery Platforms: A Near-Optimal Policy for Capacity Sizing, Order Batching, and Spatial Routing

Yang Bo[a,*], Milind Dawande[b], and Ganesh Janakiraman[b]

[a]CUHK Business School, The Chinese University of Hong Kong
yangbo@cuhk.edu.hk
[b]Naveen Jindal School of Management, The University of Texas at Dallas
milind@utdallas.edu, ganesh@utdallas.edu

Extended abstract submitted for presentation at the 12th Triennial Symposium on Transportation Analysis conference (TRISTAN XII)
June 22-27, 2025, Okinawa, Japan

February 14, 2025

Keywords: on-demand delivery, spatial operations, capacity sizing, asymptotic analysis

1 INTRODUCTION

The food-delivery business is low-margin and logistically challenging. For example, a recent Wall Steet Journal article reports that, after accounting for all expenses, DoorDash is left with just 2.5% of the customers' total bill (Wall Street Journal, 2021). With payment to drivers being the biggest expense for food-delivery platforms, improving drivers' efficiency in every possible way – by choosing the correct number of drivers, by cleverly batching spatially-distributed orders, and by optimizing delivery routes to shave-off delivery time – is critical for survival in this scale-driven business. Our goal in this paper is to characterize the fundamental trade-offs in this stochastic, dynamic optimization problem faced by a platform and offer an algorithm that is not only near-optimal but also asymptotically tight; i.e., whose long-run average cost per unit time matches a lower bound.

2 Problem Statement

We consider a continuous-time, infinite-horizon, stochastic, dynamic food-delivery system on the Euclidean plane. Located at the origin $z^0 := (0, 0)$ is a hub of restaurants (or a cloud-kitchen) that serves customers living within a connected, bounded region $\mathcal{R}$, who place orders from these restaurants through a delivery-app of a platform. Assume that z^0 also lies within $\mathcal{R}$ without loss of generality. The cloud-kitchen operator (i.e., the firm that owns and operates the restaurant hub) foresees an infinite planning horizon and employs $n_s \in N$ drivers/servers upfront for the pick-up and delivery of all potential orders; n_s is a one-time decision at the beginning of the horizon. Drivers are salaried at a constant wage of $c_s > 0$ per driver per unit time.

During the planning horizon, orders from customers arrive at the system according to a Poisson process with rate $\lambda > 0$, where the arrival time and location of every customer become known at the time of her arrival. Upon arrival, every order, initially unprepared, begins preparation instantaneously and spends a deterministic, constant preparation time of $t_p > 0$ at the cloud-kitchen. This order then becomes "prepared" and waits at the cloud-kitchen before it enters delivery service, i.e., is loaded/picked up by some driver who sets out for delivery. Customers introduce waiting costs to the platform at $c_o > 0$ per customer per unit time of waiting, where the waiting time of a customer counts from her arrival until her departure from the system.

Enroute to delivering orders, each driver drives at a velocity of either 0, i.e., not moving, or $v > 0$, on the Euclidean plane with the shortest traveling distance between any two locations being the Euclidean distance. The delivery service of an order completes once this order is delivered at its arrival location. Assume that the pick-up operation (resp., the drop-off operation) of every order at z^0 (resp., at its arrival location) is instantaneous. In a single delivery trip, a driver can carry/deliver multiple orders to exploit spatial economies of scale.Assume that a driver cannot preemptively return to the cloud-kitchen to pick up newly available orders before completing the delivery of all orders currently loaded in his vehicle. Each driver has a constant capacity of $q \in N$, which is the maximum number of orders the driver can carry in a trip.

The objective of the platform is to minimize the expected long-run average cost incurred per unit time by finding the one-time capacity investment decision on n_s and real-time order-batching and driver-routing decisions for the spatial delivery of orders, where the cost includes the wages to the drivers and the waiting costs of the orders.

2.1 Information Structure: Anticipation of Order Arrivals

In our setting, the presence of a positive food preparation time t_p causes a special information structure for decision making: at any time $t > 0$, the platform, as the decision maker, effectively *anticipates* the arrival times and locations of prepared orders[1] with full certainty for a future time duration of t_p, as these prepared orders correspond precisely to the unprepared orders that arrive between $\max(t - t_p, 0)$ and t, whose arrival information is available by t. We see that this information could be abundant when the demand arrival rate is high and/or the food preparation time t_p is long. Besides, this information is collected “for free", i.e., without incurring extra cost beyond the cost necessarily incurred while these orders are in the process of preparation. The special information structure in our model causes significant technical challenges, and requires the development of novel techniques.

3 Methodology and Results

We analyze exact spatial routing in the stochastic, dynamic food-delivery system with multiple vehicles. Our direct analysis of the spatial system differs significantly from that in the prevailing literature on stochastic, dynamic spatial systems, which oftentimes involves a stylized queuing approximation of the original spatial system that allows for a convenient use of many existing ideas/results from the queuing literature; see, e.g, Chen *et al.* (2022) and Besbes *et al.* (2022). Our setting, by contrast, poses significantly new technical challenges. To overcome these technical challenges, we develop novel combinatorial arguments that relate our stochastic, dynamic problem to several static Euclidean vehicle-routing problems and exploit the non-trivial geometric and probabilistic properties therein. Our analytical framework, while developed for the study of stochastic, dynamic spatial delivery systems, can also be helpful for studying other complex on-demand spatial systems. Our analysis establishes the following three main results, which we explain in Sections 3.1-3.3 below.

3.1 The Spatial-Temporal Trade-off Inequality

At the core of our entire analysis is an inequality that characterizes the fundamental spatial-temporal trade-off between balancing the spatial economies of scale and orders' waiting before they enter the delivery service: for any given policy, if the long-run average travel distance per order is $\overline{L}$ and the long-run average waiting time at the hub, i.e., z^0, per order (resp., number of

[1]A prepared order arrives at the system upon the completion of its preparation.

orders waiting in the system) is $\overline{T}_w$ (resp., $\overline{N}_w$) under this policy, then the spatial economies of scale, as measured[2] by $\overline{L}$, cannot be better than $\frac{2\bar{r}}{q} + \Omega(\frac{1}{\sqrt{\lambda \overline{T}_w + c}})$, where $c > 0$ is some constant, and $\bar{r}$ is the expected distance of the location of an order from the cloud-kitchen. This inequality is an interesting result by itself, while it is also essential for the understanding and analysis of subsequent results/manegerial insights.

To the best of our knowledge, Bertsimas & Van Ryzin (1993a,b) (BV (1993a,b) hereafter) are the only earlier papers that establish a similar result, where the authors study the dynamic traveling repairman problem (DTRP) with exact spatial routing and multiple vehicles. Though related, we note a crucial difference in the information structure between their setting and our setting: in BV (1993a,b), every order instantaneously becomes "ready" to receive service upon arrival in the system, i.e., can be assigned to some server (if available) at any moment after its arrival; by analogy, their setting corresponds to the food preparation time $t_p = 0$ in our setting. As a result, There is absolutely no anticipation of the arrival time and location of a future unassigned order at any time prior to the arrival of this future order in BV (1993a,b), as compared to the anticipation of the arrival information of unassigned orders for a future time duration of $t_p > 0$ in our setting; see Table 1. The spatial-temporal inequality in BV (1993a,b)

Table 1 – (a) t_p denotes the order preparation time, and (b) "anticipation" refers to the anticipation of the arrival times and locations of future unprepared orders

	system	t_p	information structure	Spatial-Temporal trade-off
BV (1993a,b)	DTRP	$= 0$	no anticipation	$\overline{L} \geq \frac{2\bar{r}}{q} + \Omega\left(\frac{1}{\sqrt{\lambda \overline{T}_w + c}}\right)$
Our paper	food-delivery	> 0	anticipation	$\overline{L} \geq \frac{2\bar{r}}{q} + \Omega\left(\frac{1}{\sqrt{\lambda \overline{T}_w + c}}\right)$

does not apply to our setting because substantially more information is available at every moment for decision making, which could potentially bring down the spatial inefficiency as captured by $\overline{L}$. For example, while an unassigned order arrives at the system, the existing unassigned orders at z^0 could be precisely those with close proximity to that order's arrival location, as carefully chosen by the decision maker who anticipates the arrival of that order t_p time units earlier. Surprisingly, we show that, $\overline{L} \geq \frac{2\bar{r}}{q} + \Omega\left(\frac{1}{\sqrt{\lambda \overline{T}_w + c}}\right)$ remains true irrespective of the degree of anticipation, even for clairvoyant policies that can perfectly anticipate the arrival times and locations of all the future orders at time 0.

3.2 A Near-Optimal Policy

We develop a novel anticipating, group scheduling policy (AGS) for the platform's problem, and evaluate the performance of AGS based on two measures: (a) The performance gap of AGS, which measures the extra cost incurred under AGS beyond the platform's total cost under a clairvoyant optimal policy that has full information of the arrival times and locations of all orders during the entire planning horizon; and (b) the excess delay under AGS, which measures the extra time an order spends in the system beyond the least time that this order has to spend in the system. We show that AGS simultaneously ensures (a) a vanishing performance gap and (b) a vanishing excess delay in a meaningful asymptotic regime.

3.3 Safety Staffing Rule

For service systems without the spatial feature, the square-root safety staffing rule (Halfin & Whitt, 1981), i.e., staffing $\Theta(\lambda^{\frac{1}{2}})$ drivers above the nominal offered load, is known to balance servers' salary costs and customers' waiting costs. For our spatial food-delivery system, however,

[2] The smaller the value of this measure, the better the spatial economies of scale.

our analysis shows that a higher safety level, namely $\Theta(\lambda^{\frac{2}{3}})$, is needed to balance these two costs. This finding echoes that of an earlier paper by Besbes *et al.* (2022), who study capacity sizing for a spatial ride-hailing platform and establish that a safety level of $\Theta(\lambda^{\frac{2}{3}})$ balances capacity utilization and customers' waiting under the nearest-neighbor dispatching protocol for drivers. Thus, echoing Besbes *et al.* (2022), our paper points to a fundamental departure in the capacity-sizing of spatial service systems from conventional service systems without the spatial feature. Though related in this respect, we note that our paper differs significantly from Besbes *et al.* (2022) in both the setting and the analysis. In particular, Besbes *et al.* (2022) establish mathematical results for a state-dependent queuing system that is a stylized approximation of the original spatial system, and regard these results as qualitative characterizations of the spatial system. By contrast, we directly study the food-delivery system with exact spatial routing, and establish all our results for this spatial system. To the best of our knowledge, the present work is the first to show the optimality of the $\lambda^{\frac{2}{3}}$-Staffing Rule in an actual spatial service setting.

3.4 Figures

We adopt the modified tour partitioning policy (MTP) designed in BV (1993a,b) to our setting as the benchmark, and refer to it as the BV benchmark. We numerically test our policy AGS against the BV benchmark in practically relevant parameter regimes. The results demonstrate the good performance of AGS in all instances. For example, in Figure 1, when[3] $q = 2$, $c_s = 1$, and $c_o \in \{0.3, 1, 3\}$, the platform's cost under AGS is lower than that under the BV benchmark for almost all values of the order arrival rate λ.

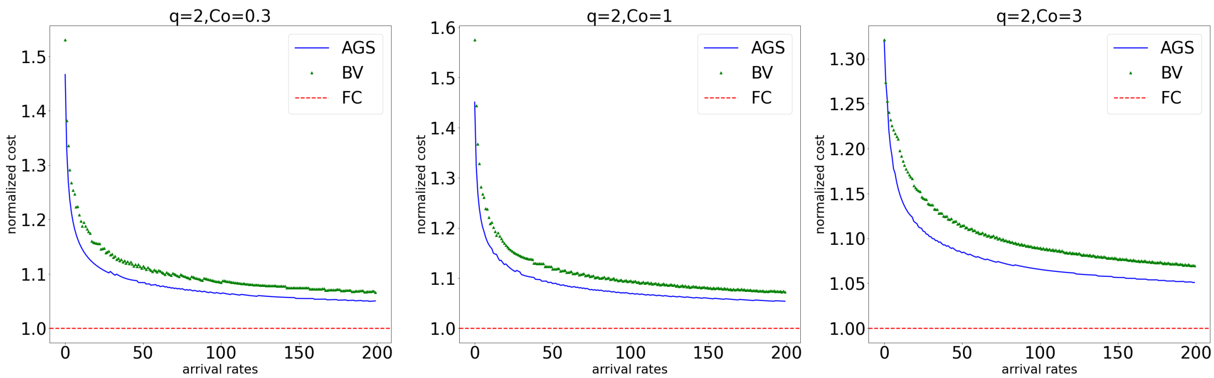

Figure 1 – The Performances of Policies in Instances $(q, c_o) = (2, 0.3), (2, 1)$, and $(2, 3)$. "FC" denotes a theoretical lower bound on the platform's cost normalized to one. The platform's cost under AGS (resp., the BV benchmark) is normalized by FC as well in the plots.

References

Bertsimas, D. J., & Van Ryzin, G. 1993a. Stochastic and Dynamic Vehicle Routing in the Euclidean Plane with Multiple Capacitated Vehicles. *Operations Research*, **41**(1), 60–76.

Bertsimas, D. J., & Van Ryzin, G. 1993b. Stochastic and Dynamic Vehicle Routing with General Demand and Interarrival Time Distributions. *Advances in Applied Probability*, **25**(4), 947–978.

Besbes, O., Castro, F., & Lobel, I. 2022. Spatial Capacity Planning. *Operations Research*, **70**(2), 1271–1291.

Chen, M., Hu, M., & Wang, J. 2022. Food Delivery Service and Restaurant: Friend or Foe? *Management Science*, **68**(9), 6539–6551.

Halfin, S., & Whitt, W. 1981. Heavy Traffic Limits for Queues with Many Exponential Servers. *Operations Research*, **29**(3), 567–588.

Wall Street Journal. 2021. DoorDash and Uber Eats Are Hot. They're Still Not Making Money. *Wall Street Journal (May 28, 2021)*, **28**.

[3] The values of other parameters are hidden in the interest of space.