Who Has Access to E-Commerce and When?
Time-Varying Service Regions in Same-Day Delivery

D. Banerjee A.M. Stroh A.L. Erera A. Toriello*

Industrial and Systems Engineering, Georgia Tech, Atlanta, Georgia, USA
* Corresponding author (atoriello@isye.gatech.edu)

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

December 9, 2021

Keywords: e-commerce, same-day delivery, last mile

1 Motivation

In the past decade, e-commerce has driven many innovations in logistics. One such advance, same-day delivery (SDD), allows a consumer to purchase and receive goods within the same service day. SDD allows e-commerce firms to more directly compete with brick-and-mortar retail by providing the customer with near-instant gratification; however, this pressure to deliver under a same-day deadline implies high last-mile logistical costs and complex operations. The question of where to offer SDD and how late to accept orders is of paramount importance in this context: A service region that is too small or an early deadline may lose SDD customers and market share, while a large region or late acceptances may result in very costly operations or even worse, failed deliveries and loss of customer goodwill.

With this motivation in mind, our goal in this study is the design of an SDD service region and the choice of order cutoff deadline, from the perspective of a single warehouse or fulfillment center that wishes to maximize the order volume it serves. In particular, we study the question of whether the system gains by "segmenting" the service region; i.e. by offering different deadlines to different parts of the region. Our results indicate that the system may indeed gain substantially by such a segmentation, and distinguish our study from the growing body of work in e-commerce service region and vehicle zone design (Banerjee *et al.*, 2021, Carlsson *et al.*, 2021, Stroh *et al.*, 2021), which assumes static regions that do not vary over the course of the service day. The intuition is clear: Customers that are farther away, e.g. in suburban areas, may need to place orders earlier in the day to obtain SDD, while we can increase our SDD order volume by allowing nearby in-town customers to place SDD orders until later in the day.

2 Problem Statement

We consider an SDD system in which a depot and its vehicle fleet offer SDD in a region whose size is to be determined. Our goal is to choose the region's size and the order acceptance deadline during the service day in order to maximize the expected number of served SDD orders. We consider both the region size and order deadline to be tactical design variables the system would hold fixed for a medium-term horizon of at least a few months, so we make approximations that capture the system's average behavior over a service day and simplify our analysis. In particular, we assume a constant order arrival rate (per time and per area unit) over the potential service region and use continuous approximation techniques (Beardwood *et al.*, 1959) to model vehicle dispatching times. These approximations have been extensively applied and validated in the tactical design of logistics systems, and once the design variables are chosen, the system can be optimized at the operational level using more detailed models, e.g. Klapp *et al.* (2018), Ulmer *et al.* (2019), Voccia *et al.* (2017).

The length of the service day is T; SDD orders begin accumulating and delivery vehicles may depart any time after $t = 0$, and all vehicles must return to the depot by $t = T$. SDD orders accumulate in the region at a continuous rate of λ per time and area unit. We consider a potential service region centered around the depot with a fixed geometry (e.g. a circular service region, or a "pie wedge"), but allow the region's radius to vary. In realistic settings, the region's radius defines a maximum driving time or distance from the depot, as we discuss below in our computational case study. We model the time a vehicle requires to serve n orders in a region of area A as $f(A,n) = c_0\sqrt{An}$ for a constant c_0 (Beardwood *et al.*, 1959). Equivalently, it is more convenient for our analysis to define routing time as $f(A,\tau) = cA\sqrt{\tau}$, where τ is the accumulation time since the previous order dispatch and $c = c_0\sqrt{\lambda}$.

The model's objective is to choose a set of feasible dispatches for each vehicle in the fleet to serve a maximum total number of orders. We allow each dispatch to serve a region of different size, and thus dispatch d is specified by a tuple (t_d, A_d). For example, if we have a single vehicle and allow it to make only one dispatch, the model becomes

$$\max_{t,A\geq 0} \lambda At \qquad \text{s.t. } t + cA\sqrt{t} \leq T. \tag{1}$$

The constraint stipulates that if we accept SDD orders in an area A for a time t and then dispatch the vehicle with these orders, the vehicle must return by T. The optimal t^* specifies the SDD order cutoff deadline (as we only make one dispatch) and the optimal A^* specifies the area to serve. It is not difficult to show that $t^* = T/3$; i.e. under the single-vehicle and single-dispatch assumption, we should allow orders to accumulate for the first third of the service day and choose the area so the vehicle spends all of the remaining two thirds serving these orders. For example, if we assume the SDD service day lasts from 9AM to 6PM, our model suggests accepting SDD orders until noon and delivering them between noon and 6PM.

Naturally, the model becomes more complex as we increase the number of vehicles, the number of dispatches per vehicle, or both. We have results for a single delivery vehicle making multiple dispatches, and for a larger fleet where each vehicle is dispatched once. For brevity's sake, we focus on the former case in this abstract.

3 Results

The general case of the single-vehicle model performing D dispatches is captured by

$$\max_{t_d,A_d\geq 0} \sum_{d=1}^{D} \lambda A_d t_d \qquad \text{s.t. } t_d + cA_d\sqrt{t_d - t_{d-1}} \leq t_{d+1},\ d < D; \qquad t_D + cA_D\sqrt{t_D - t_{D-1}} \leq T, \tag{2}$$

where $t_0 = 0$. The constraints stipulate that the d-th dispatch, leaving at t_d, carrying orders for area A_d accumulated between t_{d-1} and t_d, must return by the next dispatch t_{d+1}, or the end of the service day T in the case of the last dispatch. We briefly summarize our results:

Single Dispatch The optimal solution of (1) is a $1/2$-approximation of (2) for any $D > 0$. In other words, a single dispatch is able to serve at least 50% of the orders that a single vehicle can serve with any number of dispatches. This suggests that even a relatively simple setup with a single vehicle and single dispatch can command a significant fraction of the potential SDD market.

Two Dispatches When we allow the vehicle to make $D = 2$ dispatches, we increase the number of orders served by 15.5% from the single-dispatch case. The optimal solution of this model has $t_1^* = T/9$, $t_2^* = 5T/9$ and $A_1^* = 2A_2^*$. The first dispatch serves orders for a relatively short window (one ninth of the service day) over a comparatively large area, while the second serves orders in a region of half the size, but orders in this region accumulate until after the halfway point (5/9) of the service day. For example, returning to a service day that ranges from 9AM to 6PM, our model suggests allowing a relatively large service region (e.g. including outlying suburbs) to place SDD orders until 10AM, while customers in a smaller, in-town service region of half the area can continue placing orders until 2PM. Figure 1 diagrams the solution.

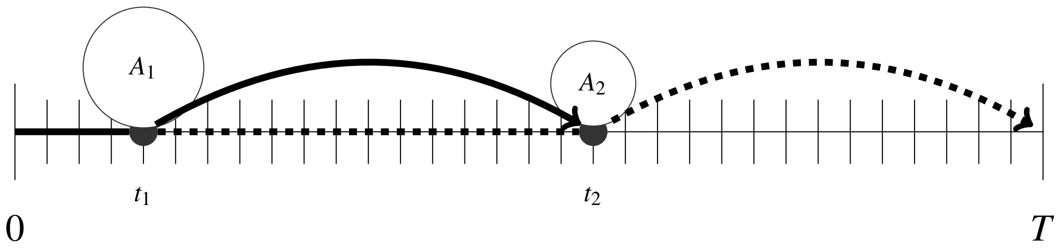

Figure 1 – *Optimal single-vehicle, two-dispatch dispatch policy with scaled service areas.*

Additional Dispatches Empirically, we observe a significant decrease in the marginal benefit of adding dispatches beyond two. For example, adding a third dispatch increases the number of orders served by less than 1.6% over the two-dispatch optimal solution, and a fourth dispatch provides essentially no benefit, less than 0.1%. This suggests that two dispatches should suffice for an SDD system that is served using a single vehicle.

Our results for multiple dispatches also confirm the intuition that larger areas should be served earlier, while smaller areas can enjoy later SDD service: all optimal solutions satisfy $A_d^* > A_{d+1}^*$. Similarly, there are decreasing marginal returns for every additional vehicle we add to the service fleet.

4 Case Study

We briefly discuss the practical application of our results in a case study based in the Phoenix, Arizona metro area. We assume the SDD depot is located in the Glendale suburb and has a four-vehicle fleet, with each vehicle serving one of the four quadrants surrounding the depot; we focus on the southeastern quadrant. As in our examples, the service day starts at 9AM and ends at 6PM; we use an order accumulation rate of 0.2 per square mile per hour, and calibrate the constant c_0 in our driving time approximation using true driving times in the Phoenix road network queried from openrouteservice.org.

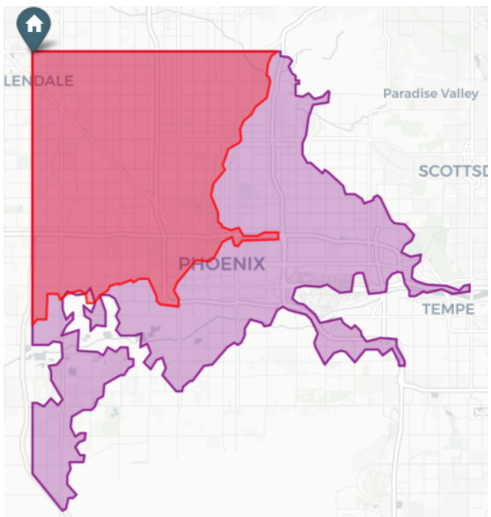

Figure 2 – *Service areas for case study in Phoenix, AZ metro area.*

Figure 2 details the resulting service areas; we generated the figure with VeRoViz (Peng & Murray, 2020, veroviz.org). As the discussion in the previous section indicates, customers in the larger area (red and purple, roughly 113 square miles) can place orders from 9AM to 10AM, after which a first vehicle dispatch delivers these orders, with an expected duration of four hours. From 10AM to 2PM, customers in the smaller region (red only, roughly 57 square miles or about half of the larger area) continue placing orders. Upon its return at 2PM, the vehicle departs again immediately to deliver these orders, with an expected return at the end of the service day at 6PM. The expected number of orders served by each dispatch are 22.6 and 45.3, respectively.

To validate the model's predictions, we perform a simulation of the depot's operations in which order arrivals are governed by a Poisson process with the same rate (0.2 per hour per square mile). We simulate 120 service days and implement an operationalized version of the dispatch policy, in which actual cutoff times, dispatch times and durations may vary based on that day's arrival process: As orders accumulate in the larger region A_1, we continually solve for θ_1, the duration of the TSP route that delivers these orders, until the first time t in which $t+\theta_1+f(A_2,\theta_1)$ equals or exceeds the end of the service day. At that point the first dispatch is executed and we accept orders only from the smaller region A_2, continually solving for θ_2, the duration of the TSP that delivers the second set of orders. We accept orders until $t+\theta_1+\theta_2$ equals or exceeds the end of the service day (dropping the last order if it causes the total time to exceed the service day), and execute this dispatch upon the vehicle's return at time $t+\theta_1$.

The simulation's sample means for orders served and total dispatch time are both within roughly 5% of our predictions. Interestingly, in both cases the model's predictions are more conservative than the simulated results (i.e. the model predicts slightly fewer orders served and slightly longer dispatch time), a function of some necessary conservatism in our driving time approximation and its calibration to the Phoenix road network.

References

Banerjee, Dipayan, Erera, Alan, & Toriello, Alejandro. 2021. *Fleet sizing and service region partitioning for same-day delivery systems*. In review, http://www.optimization-online.org/DB_HTML/2020/12/8183.html.

Beardwood, J., Halton, J., & Hammersley, J. 1959. The Shortest Path Through Many Points. *Mathematical Proceedings of the Cambridge Philosophical Society*, **55**(4), 299–327.

Carlsson, John Gunnar, Liu, Sheng, Salari, Nooshin, & Yu, Han. 2021. *Provably good region partitioning for on-time last-mile delivery*. In review, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3915544.

Klapp, M., Erera, A., & Toriello, A. 2018. The Dynamic Dispatch Waves Problem for Same-Day Delivery. *European Journal of Operations Research*, **271**(2), 519–534.

Peng, Lan, & Murray, Chase. 2020. *VeRoViz: A Vehicle Routing Visualization Toolkit*. https://ssrn.com/abstract=3746037.

Stroh, A., Erera, A., & Toriello, A. 2021. Tactical Design of Same-Day Delivery Systems. *Management Science*. In press.

Ulmer, M., Thomas, B., & D., Mattfeld. 2019. Preemptive Depot Returns for dynamic same-say delivery. *EURO Journal on Transportation and Logistics*, **8**(4), 327–361.

Voccia, S., Campbell, A., & Thomas, B. 2017. The Same-Day Delivery Problem for Online Purchases. *Transportation Science*, **53**(1), 1–39.