Same-Day Delivery with Drone Resupply

Iman Dayarian[1], Martin Savelsbergh[1], and John-Paul Clarke[2]

[1]H. Milton Stewart School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, USA

[2]Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, USA

Abstract

Unmanned Aerial Vehicles (UAVs), commonly referred to as drones, have recently seen an increased level of interest as their potential use in same-day home delivery has been promoted and advocated by large retailers and courier delivery companies. We introduce a novel way to exploit drones in same-day home delivery settings: drone resupply. We consider a home delivery system in which delivery vehicles are regularly resupplied by drones. Resupply can take place whenever a delivery vehicle is stationary and a drone can land on the vehicle's roof. We introduce the vehicle routing problem with drone resupply to capture and investigate this setting. We develop different algorithms and compare their performance. Finally, we quantify the potential benefits of drone resupply and generate valuable insights for advancing this concept.

Keywords— same-day delivery; dynamic decision-making; drone delivery; vehicle routing problem

1 Introduction

The markets for both durable and perishable goods have shifted in focus over the past decade to direct-to-consumer (D2C) deliveries. This shift has been driven by the rapid and significant growth in e-commerce, specifically online shopping followed by home delivery, which has growth by between 7 and 10 percent in mature markets, e.g., the US and Germany, and close to 300 percent in developing markets, e.g., India (see Joerss et al., 2016). Market share in the e-commerce space is primarily a function of the following pre-purchase considerations of online shoppers, i.e., speed, flexibility, security and the cost of delivery. That said, even though a growing number of online shoppers are placing greater importance on the speed of delivery, they are also highly price sensitive and often unwilling to pay more for faster delivery. Thus, retailers will have to absorb most of the cost of more speedy delivery services, and in a highly competitive market where many retailers are struggling to stay ahead of the competition, all retailers are looking for innovative and cost-effective ways to provide faster D2C deliveries.

To date, much of the activity with respect to rapid delivery has been in relation to urban and suburban environments. This focus on the urban environment is no surprise given demographic trends which indicate that, by 2050, 70% of the world's population, approximately 6.3 billion people, will live in major cities (Bretzke, 2013), with "megacities" being at the forefront of the trend (megacities refers to cities with a population exceeding 10 million inhabitants). The high population densities in urban environments, in particular megacities, results in higher and geographically concentrated demand, which in turn results in greater economies of scale, especially within the context of the traditional terrestrial delivery network.

However, there are significant challenges to a purely terrestrial rapid delivery network in urban environments. Specifically, traffic congestion is frequently reported to be the most pressing infrastructure problem in megacities, with freight traffic being simultaneously one of the causes of the infrastructure overload and one of its victims. Moreover, a common reaction from the city administrations has been imposing restrictive access regulations to city centers for larger vehicles, either permanently or during certain time periods. These challenges drive retailers and carriers to seek new technologies and new concepts to provide cost-effective, flexible, and fast D2C delivery services. Examples of new concepts being pursued are crowdshipping (van Cooten, 2016; Dayarian and Savelsbergh, 2017), trunk delivery (Reyes et al., 2016; Gozbayin et al., 2017), side-walk robots (Wong, 2017), and drone delivery (Hegranes, 2017). The use of drones, especially, has obvious advantages in home delivery systems because they do not require a driver and, thus, may result in cost savings, and they do not get stuck in traffic and, thus, may result in time savings.

Large retailers such as Amazon and Alibaba as well as courier delivery companies such as UPS and DHL are already considering home delivery services partially or fully performed by drones. In fact, on 1 December 2013, in an interview on 60-Minutes, the CEO of Amazon, Jeff Bezos, revealed that his company had plans for an ultra-rapid package delivery service entitled Amazon Prime Air. The vision for Amazon Prime Air is that a customer within a ten-mile radius of a participating Amazon order fulfillment center will be able to receive their purchase within 30 minutes of order via an autonomously-operated multi-rotor drone if the total order weighs less than five pounds (2.26 kg) and can fit in the cargo compartment of the drone (Pierce, 2013). Since that date, Google has announced plans for a drone-based package delivery system, and several other companies around the world are either developing such systems or researching related concepts.

Research in the transportation science community, specifically the vehicle routing and scheduling research community, has in large part mirrored the focus of the large retailers and courier companies. For example, in the last few years, several papers on models and algorithms to support this operational use of drones have appeared (Murray and Chu, 2015; Agatz et al., 2016; Campbell et al., 2017). Two settings have been investigated: (1) drones deliver a single package on an out-and-back trip from a fulfillment center and supplement the delivery capacity provided by a regular fleet of delivery vehicles, and (2) drones deliver a single package on an out-and-back trip launched from a delivery vehicle, where the delivery vehicle carries one or more drones, also supplementing the delivery capacity provided by a regular fleet of delivery vehicles, but, of course, also reducing that capacity since drones have to be carried. The prevailing assumption in existing concepts of operation for drone-based package delivery is that each package will be delivered via a single drone, and that package recipients will be adjacent to a location where the drone can land or in a location where the package may be lowered to them via a rope or similar mechanism. However, the vast majority of the concepts of operation that have been proposed will be difficult to realize in extreme urban environments (such as Manhattan).

Drone operations will be particularly challenging in extreme urban environments where the vast majority of the windows in tall buildings are inoperable (Steemers and Yannas, 2000); the wind conditions at potential landing zones, on the roof-top of building or at ground-level near and between tall buildings (Baskaran and Kashef, 1996), are outside the operating envelope of the small- to medium-sized drones that have been proposed to-date for package delivery (Zarovy et al., 2013; Da C. Siqueira, 2017); the so-called "urban canyons" between tall buildings create conditions for multiple conflicting surveillance signals (e.g. GPS multipath) and the loss of communication due to line-of-sight blockages (Adjrad and Groves, 2017); and the high population density increases the risk of fatalities if a drone should fail (Weibel and Hansman, 2004). Thus the overarching technical challenge is to develop a concept of operation where the drones are sufficiently large to operate in all the airspace where they are required to operate to fulfill their mission, but, at the same time, they are not so large that the amount of cargo required to make their operation viable is such that the ground segments in the delivery itinerary will be too long. Efforts to overcome this technical challenge are starting to pay off and drones carrying more than a ton have already been deployed (McDonald, 2017).

Therefore, we are considering a different use of drones, namely *drones that resupply delivery vehicles*. This not only eliminates some of the safety concerns, but also solves the issue of how to handle delivery addresses without an obvious place for a drone to drop off a package. Furthermore, it may be especially useful in large metropolitan area, where it may be difficult or costly to have many fulfillment centers, which would be required to overcome the limited range of small drones. We envision a delivery system in which delivery vehicles make deliveries in certain areas and are regularly resupplied by a drone. Resupply can take place anywhere as long as the deliver vehicle is stationary (the drone will land on the roof of the delivery vehicle). The envisioned delivery system has the advantage that (1) it is less costly, and (2) the area in which same-day delivery service can be offered can be expanded, because drones can travel faster – because of their speed, because they are not impeded by traffic congestion, or both.

We introduce the *vehicle routing problem with drone resupply* (VRPDR) to represent this form of drone-assisted delivery, but focus our algorithmic efforts on the special case with a single drone and a single delivery vehicle. This special case already captures many, but certainly not all, of the algorithmic challenges encountered when developing operational decision support technology for such a home delivery system. Synchronization, another characteristic that has attracted a lot of attention in the routing and scheduling community recently, see e.g., Drexl (2012), is the most challenging feature, as the delivery vehicle and the drone need to meet at a particular location for resupply to take place. For example, determining the latest feasible dispatch time of a package at fulfillment center to reach its destination at a commitment service time is far from trivial.

Our contributions can be summarized as follows:

- We introduce an innovative home delivery system concept, especially designed to support a highly dynamic environment with tight service guarantees, based on the use of a fleet of drones to resupply delivery vehicles. The system offers advantages over other drone-based delivery systems in that it better supports fast deliveries in a large coverage area at low cots.
- The viability of the proposed home delivery system has been demonstrated for the special case in which the system consists of a single drone and a single delivery vehicle.
- The insights obtained from an extensive computational study indicate that the proposed home delivery system is effective under a wide range of system parameters.

The remainder of the paper is organized as follows. In Section 2, we provide a brief overview of the relevant literature. In Section 3, we provide a detailed description of the general problem setting as well as the variant considered in this paper. In Section 4, we introduce our solution approach for the single vehicle, single drone version of the problem. The results of an extensive computational study are presented in Section 5. Finally, We provide some concluding remarks and future directions in Section 6.

2 Relevant literature

UAVs are a compelling alternative to other modes of transport, especially in time-sensitive environments, because they are less sensitive to ground conditions and congestion (Kim et al., 2017). This realization has accelerated the maturation of the technology and has expanded the set of domains in which the use of UAVs is contemplated, e.g., medicine and vaccine delivery, organ transport, defense, search and rescue, aerial imaging, environmental surveillance, and agriculture (Markoff, 2016; Thiels et al., 2015; Francisco, 2016; Khosiawan and Nielsen, 2016; Barrientos et al., 2011; Avellar et al., 2015; Yakıcı, 2016).

Our research focuses on the use of UAVs in home delivery systems, and, consequently, our literature review focuses on papers related to last-mile delivery. In the last-mail delivery literature, drones are being employed in different ways, either independently or in conjunction with delivery vehicles, and are launched from and recovered from different locations, either a central facility, a mobile facility, or even a (specially designed and equipped) delivery vehicle.

Murray and Chu (2015) introduced the flying sidekick traveling salesman problem (FSTSP), in which a set of customers must be served exactly once by either a delivery vehicle or a drone operating in conjunction with the delivery vehicle. The delivery vehicle and the drone must depart from and return to a single depot exactly once. Over the course of a delivery cycle, the drone may make multiple out-and-back trips, each time making a single delivery, starting and ending at either the depot or the delivery vehicle when it is parked at a customer location. While the trip of the delivery vehicle has multiple stops at customer locations (where deliveries are made), a trip of the drone has a single single stop at a customer location (making a single delivery). The objective of the FSTSP is to minimize the time required to serve all customers, i.e., the difference between the time of departure and time of return of the delivery vehicle and the drone from the depot, and accounts for vehicle travel time, launch and recover time of the drone, and any waiting time of the delivery vehicle for the return of the drone. The concept of a flying sidekick is similar to the concept of transportable delivery resources introduced in Lin (2011), where on-foot couriers traveling on a delivery vehicle perform deliveries in conjunction with the delivery vehicle. Murray and Chu (2015) also consider a simpler version of the problem, in which the deliveries are divided into two groups, one set delivered by the delivery vehicle and the other set delivered by the drone traveling out-and-back from the distribution center. They develop simple construction and improvement heuristics for both settings.

Ponza (2016) proposes an improved mathematical formulation for the FSTSP and develops a simulated annealing solution approach. Agatz et al. (2016) also consider a last-mile delivery setting involving a combination of a delivery vehicle and a drone, which they refer to as the traveling salesman problem with drone (TSP-D). They present solutions obtained by a route first — cluster second heuristic and show that the vehicle-drone delivery system can be more effective than the vehicle-only delivery system. Ha et al. (2015) consider a TSP-D in which once the drone is launched from either the depot or the vehicle, it either returns to the vehicle at the vehicle's next stop, or it returns to the depot. The main difference with the version considered by Agatz et al. (2016) is that the drone is not launched and recovered at the same stop. Later, Ha et al. (2016) consider a TSP-D in which the objective is to minimize the total transportation cost, comprising of vehicle and drone costs. They present heuristic solution methods based on GRASP and modifying an optimal TSP tour. Ferrandez et al. (2016) introduced a vehicle-drone in tandem delivery system and analyze the system in terms of time and energy. In their model, all demands are known in advance and the drone is launched from the vehicle, which follows a TSP tour, i.e., the vehicle acts as a moving hub. Each drone can carry at most one package and therefore must traverse a star distance of ingress and egress from vehicle to delivery location and back to vehicle. The solution approach uses K-means clustering to find the vehicle stops, and a genetic algorithm for finding the vehicle route. Their analysis indicates that the energy efficiencies achieved when employing a drone are significant, but that time improvements only occur when the speed of the drone is at least twice the speed of the vehicle. Wang et al. (2017) address a setting similar to the FSTSP, which they refer to as the vehicle routing problem with drones (VRPD). Specifically, they incorporate multiple vehicles, assume that each vehicle can launch two drones, that drones can only fly along the same routes as the vehicles, and at the return of a drone to a vehicle, both the drone and the vehicle can wait for each other. Campbell et al. (2017) examine drone deliveries, using small payload capacity drones (e.g., 5 lbs.), from fixed depots in conjunction with hybrid vehicle-drone routes and vehicle-only delivery. In contrast to the discrete TSP-based optimization models, they develop strategic models for the design of vehicle-drone delivery systems using continuous approximation modeling techniques. This allows them to treat the demand for deliveries as a continuous spatial density over a service region.

Dorling et al. (2017) focus on developing vehicle routing problems (VRPs) specifically for drone delivery scenarios, considering the effect of battery and payload weight on energy consumption. They propose two multi-trip VRPs for drone delivery, one minimizing costs subject to a delivery time limit, and one minimizing the delivery time subject to a budget constraint. An energy consumption model for multi-rotor drones is derived and experimentally validated demonstrating that energy consumption varies approximately linearly with payload and battery weight.

To the best of our knowledge, no previous work has considered a drone-assisted delivery setting in which drones resupply delivery vehicles. This setting can be viewed as a two-echelon VRP, in which the location of the intermediate facilities (satellites) is not fixed. Our intent is to demonstrate that such a setting offers many practical advantages over the settings studied in the literature so far.

3 Problem statement

In this section, we formally introduce the VRPDR, in which a fleet of drones and a fleet of vehicles collaboratively perform home deliveries of online orders from a fulfillment center.

3.1 The vehicle routing problem with drone resupply

We consider a home delivery problem where online shoppers dynamically place orders throughout a time horizon T and these orders are subsequently dispatched from a fulfillment center and delivered to the online shoppers at their homes. The problem is defined on a graph $\mathcal{G} = (\mathcal{V}_0, \mathcal{A})$, where $\mathcal{V}_0$ represents the set of relevant locations and $\mathcal{A}$ the set of links between these locations. We have $\mathcal{V}_0 = \{0\} \cup \mathcal{V} \cup \mathcal{M}$ with $\mathcal{V}$ the set of home locations of online buyers, $\mathcal{M}$ the set of locations where a drone can resupply a delivery vehicle, and 0 the fulfillment center. We make no assumptions on $\mathcal{M}$, but in our computational experiments, we have that $\mathcal{M} \cap \mathcal{V} \neq \emptyset$ and $\mathcal{M} \setminus \mathcal{V} \neq \emptyset$. The distance between two locations $i, j \in \mathcal{V}_0$ is given by d_{ij}. The order placement rate at location $j \in \mathcal{V}$ is given by λ_j. Let $\mathcal{N}$ be the set of online orders placed over the planning horizon T (i.e., a realization of order placements), then we denote the placement time of order o at location j by τ_j^o. We consider a common service time guarantee S, i.e., an order o from location j placed at time τ_j^o has to be delivered at or before $\rho_j^o = \tau_j^o + S$, and we refer to ρ_j^o as the due time of the order.

The deliveries are performed collaboratively by a fleet of drones $\mathcal{D} = \{1, \ldots, D\}$ and a fleet of delivery vehicles $\mathcal{K} = \{1, \ldots, K\}$. The drones and the delivery vehicles differ in terms of their speed and their capacity. The speed of a drone and a delivery vehicle are v^D and v^K, respectively, where $v^D \geq v^K$. The fact that the speed of a drone is assumed to be higher than the speed of a delivery vehicle reflects, among other things, that (1) a drone does not get stuck in traffic, and (2) the path a drone takes between two locations may be shorter than the path a delivery vehicle takes between the same two locations. Thus, a distance d_{ij} is covered in $t_{ij}^D = \frac{d_{ij}}{v^D}$ and $t_{ij}^K = \frac{d_{ij}}{v^K}$ units of time by a drone and by a delivery vehicle, respectively. (In practice, a vehicle's speed is likely to vary more than a drone's speed, and, thus, the ratio will not be constant; for simplicity, we will assume it is.) The capacity of a delivery vehicle is assumed to be infinite, since we are dealing with small packages, while the capacity of a drone is limited. We assume that all packages have the same weight, and, thus, that the capacity of a drone, P^D, represents the maximum number of packages that can be carried. Finally, we assume that drones can land on and take off from a vehicle only when the vehicle is stationary and that the transfer of packages from a drone to a delivery vehicle takes a fixed amount of time τ_0.

On a typical day, the following events and actions take place.

1. Online shoppers place orders throughout the day.
2. Delivery vehicles are dispatched from the fulfillment center to serve a set of orders in a pre-determined route.
3. Drones are dispatched from the fulfillment center to transfer a set of orders to a delivery vehicle at a pre-determined meeting location and time (after which the drone returns to the fulfillment center). The route of the delivery vehicle receiving the orders is re-optimized such that the service time guarantee of the orders on-board the delivery vehicle (the as-yet undelivered orders already on the delivery vehicle as well as the orders just transfered to the delivery vehicle) are met.

Since drone landing and take-off are only possible when the vehicle is stationary, two package transfer options are conceivable:

Stationary transfer: The vehicle remains at the meeting location until the transfer of packages has been completed, after which the drone returns to the fulfillment center and the vehicles departs for the next location.

In-motion transfer: After the drone lands on the vehicle, the vehicle immediately departs for its next destination, and package transfer takes place while the vehicle is in motion. The drone takes off for the fulfillment center at the first stop after the transfer of packages has been completed.

The goal is to deliver as many orders as possible respecting the promised service time guarantee and doing so at minimum cost, where the cost is a function of the distance traveled by the drones and the distance traveled by the delivery vehicles. For simplicity, we assume that the per-mile travel cost of a drone and a delivery vehicle is the same. As a consequence, the main benefit of using drones is their speed, which may allow serving more orders and/or expanding the service area. (In practice, the operating costs of a drone may be significantly less than that of a delivery vehicle, because a drone does not require a driver.)

Dispatching a delivery vehicle involves decisions regarding the dispatch time of the vehicle, the subset of orders to load on the vehicle, and the sequence in which to deliver these orders. Dispatching a drone involves decisions regarding the dispatch time of the drone, the delivery vehicle to resupply, the subset of orders to load on the drone, the meeting location of the drone and the delivery vehicle, and the sequence in which to deliver the orders on the vehicle after resupply has taken place (i.e., remaining and transfered orders). The decisions regarding the dispatch time of a drone, the delivery vehicle to resupply, and the subset of orders to load on the drone, strongly depends on the delivery locations of the orders at the fulfillment center and the as-yet undelivered orders on-board the delivery vehicles. The dispatch time of a drone must be such that it arrives at the targeted meeting location at or after the delivery vehicle arrives there. That is, a delivery vehicle can wait at a meeting location, but a drone cannot.

For the remainder of the paper, we focus on the single-drone, single-vehicle variant. We elaborate on this setting and discuss some further assumptions below.

3.2 The single-drone, single-vehicle variant

We envision a fulfillment center located at the periphery of the city and orders that have to be delivered within the city boundaries. The fundamental idea of drone resupply is that the delivery vehicle performs the actual delivery of orders at online buyers' home locations, while the drone resupplies the delivery vehicle from the fulfillment center with recently placed orders. Thus, hopefully, more orders respecting the service time can be delivered and the coverage area can be expanded. Resupplying the delivery vehicle can take place before all on-board orders are delivered. After a delivery vehicle is resupplied, its delivery route will be reoptimized (considering all on-board orders – as-yet undelivered orders and orders transfered via the drone). The process is guided by a series of practical assumptions/rules:

- No late delivery is accepted. That is, if based on the system state, delivering an order on time is not possible, it is not dispatched from the fulfillment center. Orders that cannot be delivered on time are assumed to be transferred to a subcontractor at a fixed cost per order. This can be viewed as a penalty for failing to deliver an order on time. Consequently, the main objective is to maximize the number of orders delivered using the drone-delivery vehicle combination.
- If the delivery vehicle finishes its route before it is resupplied by the drone, three different actions are possible: 1) return to the fulfillment center, 2) wait at the location of the last delivery, and 3) move to the nearest meeting location (or an advantageous meeting location based on predicted future delivery locations). Since the fulfillment center is likely to be far away from the delivery locations, returning to the fulfillment center is not likely to be a good option (furthermore, this option is likely dominated by a drone resupply action given the higher speed of the drone). Moving to the closest or an advantageous meeting location may be a good option, however, it may impact the set of orders that can be transferred to the vehicle with the next drone resupply. Therefore, in our implementation, we assume that the vehicle remains at the location of its last delivery waiting for new instructions from central dispatch.
- No orders are placed after $T - S$. The vehicle is allowed to return to the fulfillment center after T (after delivering last order at or before T).

4 Solution approach

In this section, we discuss our optimization-based solution approaches to order delivery planning. We consider two strategies, which differ in the options considered for the location and time of drone resupply. In what follows, we describe each of the two strategies: *restricted-resupply*, where a drone resupply can only take place after all orders on-board a delivery vehicle have been delivered, and *flexible-resupply*, where a drone resupply can take place at any time, even before all orders on-board a delivery vehicle have been delivered.

4.1 Restricted-resupply

In this strategy, a drone dispatch from the fulfillment center is planned in such a way that a meeting with the vehicle takes place at the end of the active delivery route, i.e., once all on-board orders are delivered. In other words, the active vehicle route is never modified, but a new delivery route for the set of orders transferred to the vehicle, starting at the meeting location and ending at one of the home locations associated with the transferred orders, "accompanies" the orders.

In this strategy, new decisions are made at specific times throughout the time horizon T. The first decision corresponds to the initial vehicle dispatch from the fulfillment center. It is assumed that the vehicle starts and ends its daily route at the fulfillment center. The vehicle is first dispatched some time shortly after the start of the planning period, carrying a subset of orders placed before the vehicle's dispatch time. It will return to the fulfillment center only at the end of the planning period. In between, the vehicle may be resupplied multiple times by the drone. As mentioned before, if a vehicle finishes its active route before it is resupplied, it waits at the location of the last delivery until new instructions are issued by central dispatch (indicating at which meeting location the next resupply will take place). The decision making process throughout the planning period follows the steps described below. Each of these steps incorporates multiple algorithmic modules which are detailed in the subsequent sections.

Vehicle dispatch

The first set of decisions are associated with the first vehicle dispatch from the fulfillment center. Decisions regarding the subset of active orders to be loaded on the vehicle as well as the routing decisions (i.e., the sequence of the visits to the home locations of the orders loaded on the vehicle) are made concurrently. Since no late delivery is allowed, only orders which are guaranteed to be delivered on time are actually loaded on the vehicle. More specifically, for a given vehicle dispatch time, a route visiting all active orders is generated. This may result in a route with due date violation for one or more of the orders. If that is the case, then a minimum set of orders is deleted from the route, ensuring on-time delivery of the order that remaining. Even though the vehicle does not return to the fulfillment center until the end of the planning period, the vehicle route created still assumes a return to the fulfillment center. The first vehicle dispatch is planned to be at time h_1, δ minutes into the planning period (unless no order is placed in the interval $[0, h_1]$, in which case h_1 is set to 2δ). Consequently, only an order j with $\tau_j^o \leq h_1$ can be loaded on the vehicle.

Drone dispatch

In the restricted-resupply strategy, a vehicle is resupplied only at the end of its active route. Since no waiting time is allowed for the drone at a meeting location, the drone dispatch time at the fulfillment center is scheduled so that its arrival at the meeting location is guaranteed to be greater than or equal to the vehicle's arrival time. Note that the drone can be dispatched only when it is at the fulfillment center. Therefore, drone dispatch decisions are always made at time h^D, the next time at which the drone returns to the fulfillment center.

After delivering all on-board orders and receiving instructions from central dispatch, the vehicle moves to a meeting location $m \in \mathcal{M}$. To decide on the meeting location, we determine for every meeting location m, the earliest possible meeting time at m, θ_m, which is the time the vehicle arrives at m, if it drives to m immediately after delivering the last order of its active route. The earliest possible drone dispatch time to resupply the vehicle at location m is $\mu_m = \max\{\theta_m - t^D_{0,m}, h^D\}$. The chosen meeting location will depend on the location of the last order on the vehicle's active route as well as on the set of active orders at the fulfillment center. The goal is to choose the meeting location which allows the largest number of orders to be delivered on time. Our proposed strategy to choose a meeting location and time to resupply the vehicle follows the steps of Algorithm 1. It is worth mentioning that the package transfer option, i.e., stationary transfer or in-motion transfer, has an impact on the selected meeting location as well as on the vehicle's route following a transfer. With stationary transfer, the transfer time delays the vehicle's departure from the meeting location, which may result in a smaller number of orders that can be delivered (on time) on the vehicle's route after the drone resupply. With in-motion transfer, however, since the transfer of the packages takes place while the vehicle is in motion, the vehicle's departure from the meeting location is not delayed. The impact of in-motion package transfer is on the drone's return time to the fulfillment center, since the drone's departure from the vehicle may not happen immediately upon completion of the transfer of packages, and the drone's departure may take place at a location further away from the fulfillment center (as the vehicle has moved).

Algorithm 1: Determining a meeting location and time in the restricted-resupply strategy

Step 1: Let $h_1, \ldots, h_M$ correspond to the earliest drone dispatch times μ_m for $m \in \mathcal{M}$ in ascending order.
Step 2: Set $i := 1$.
Step 3: Set tentative dispatch time $\mu^* := h_i$.
Step 4: Identify the set of active orders, Ω_{h_i}, placed by time h_i.
Step 5: For each meeting location $m \in \mathcal{M}$, a route is generated serving the largest subset of orders in Ω_{h_i}. Select the meeting location m^* that allows delivery of the largest number of orders from Ω_{h_i}.
Step 6: If $\mu_{m^*} > \mu^*$, then set $i := i + 1$, and go to **Step 3**, otherwise stop.
Result: m^* and μ^*

In **Step 5** of Algorithm 1, for each meeting location $m \in \mathcal{M}$, a route delivering a subset of active orders in Ω_{h_i} is generated, ensuring that each order is delivered on-time (taking the chosen package transfer option into account). This is done by initially generating a route delivering all active orders in Ω_{h_i}, and then greedily deleting a minimal subset of orders so that the remaining orders are delivered on time and the number of orders transferred by the drone is no more than the drone capacity. We will discuss how the routes are generated and how the subset of orders to be deleted is selected in detail later. The location m allowing the largest number of on-time order deliveries is selected.

The set of active orders depends on the time a drone dispatch decision is made. More precisely, $\Omega_{h_1} \subseteq \Omega_{h_2} \subseteq \ldots \subseteq \Omega_{h_M}$. Deciding later may allow more active orders to be dispatched with the drone, but it will likely also result in a later return of the drone to the fulfillment center. Since an earlier return of the drone increases the chance of handling future order placements, we use the drone return time as a tie breaker.

Note that the vehicle starts moving towards the chosen meeting location m^* only when we commit to that meeting location. Before then, the vehicle remains at the location of the last delivery in its active route. Also, note that we may decide to dispatch the drone to meet at m^* at a time μ^* with $\mu^* \gneqq \mu_{m^*}$. As mentioned before, the return time of the drone to the fulfillment center, which defines the next drone availability time, depends on the package transfer option. With stationary transfer, the next drone availability time, h^D, is $\mu^* + 2t^D_{0m^*} + \tau_0$. For in-motion transfer, let ω be the location from where, and θ_D be the time at which, the drone departs from the vehicle. Then, with in-motion transfer, the next drone availability time, h^D, is $\theta_D + t^D_{0\omega}$.

4.2 Flexible-resupply

Contrary to the restricted-resupply strategy, in the flexible-resupply strategy the meeting between a drone and a vehicle does not necessarily take place after delivering the last order on the vehicle. Instead, the vehicle may be instructed to perform a detour at some point along its current route towards a meeting location where it can meet with a drone and receive new orders to deliver. A meeting before the end of vehicle's current route may result in rerouting the vehicle to deliver the set of remaining on-board orders as well as the resupplied active orders by the drone, altogether constituting the new set of on-board orders. The vehicle can move towards a meeting location only if the driver has received instructions to do so. Further, if the vehicle is en-route to deliver an order when instructions are received, moving towards a meeting location can occur only after completing the delivery of this order. To decide on the meeting location and time, we determine for every possible meeting location, m, and every possible last order delivered before moving towards the meeting location, υ, the earliest possible meeting time, θ^e_m, and the latest possible meeting time θ^l_m. The latest possible meeting time depends on the set of on-board orders that are scheduled to be delivered after υ in the active route, and, for stationary package transfer, the transfer time τ_0. When determining the latest possible meeting time, we assume that on-board orders are delivered in the sequence specified by the active route. (Once a decision on the meeting location and time has been made and the active route is updated to include the orders delivered by the drone, this sequence may change.) If $\theta^e_m \leq \theta^l_m$, then a drone resupply at meeting location m after delivering order υ is feasible. We denote the delivery time of order υ by θ_υ and we set $\theta^l_m = \infty$ when υ is the last order delivered on the active route. Note that a drone resupply at meeting location m may be feasible for more than one order υ in the active route. Feasible drone resupply options are characterized by tuples $(\upsilon, \theta_\upsilon, m, \theta^e_m, \theta^l_m)$.

Similar to the restricted-resupply strategy, the decision times depend on the drone availability at the fulfillment center. That is, decisions are made upon the drone's return to the fulfillment center, at time h^D. At a decision time h, the set of admissible drone resupply options, denoted by Π_h, are identified. An admissible drone resupply option π at decision point h is characterized by two conditions: (a) $\theta^\upsilon < h$, unless υ is the last delivery on the active route, and (b) $h + t^D_{0,m} \leq \theta^l_m$. Condition (a) ensures that the driver can be instructed to divert to the meeting location after completing the delivery at υ. In fact, θ^υ, for all but the last order on the active route represents the latest time to decide to dispatch a drone to meet right after υ. On the other hand, $\mu_\pi = \max\{\theta^e_{m_\pi} - t^D_{0,m_\pi}, h^D\}$ denotes the earliest drone dispatch time to meet at m_π, given that the drone may not wait for the vehicle at the meeting location. Let $\sigma_\pi = \min\{\theta_{\upsilon_\pi}, \mu_\pi\}$. Our proposed strategy to choose a meeting location and time to resupply the vehicle follows the steps of Algorithm 2.

Algorithm 2: Determining a meeting location and time in the flexible-resupply strategy

Step 1: Let $h_1, \ldots, h_M$ correspond to the values σ_π for $\pi \in \Pi_{h^D}$ in ascending order.
Step 2: Set $i := 1$.
Step 3: Set tentative dispatch time $\mu^* := h_i$.
Step 4: Identify the set of active orders Ω_{h_i}, placed by h_i. Update Π_{h_i}.
Step 5: For each $\pi \in \Pi_{h_i}$, a route, r_π, is generated serving the largest subset of orders in Ω_{h_i} and the on-board orders. Select the π^* that allows delivery of the largest subset of active orders from Ω_{h_i}.
Step 6: If $\sigma_{\pi^*} > \mu^*$, then set $i := i + 1$, and go to **Step 3**, otherwise stop.
Result: m^* and μ_*

Similar to the restricted-resupply strategy, in **Step 5** of Algorithm 2, for each feasible meeting location, a route is generated that delivers all active orders at the fulfillment center and all (remaining) orders on-board of the delivery vehicle (taking the chosen package transfer option into account). If it is not possible to deliver all orders on time, a minimal subset of active orders is deleted to ensure on-time delivery of the remaining orders and to respect the drone's capacity. The meeting location and time resulting in the maximum number of deliveries is selected. In the case of a tie, the drone return time to the fulfillment center is used as tiebreaker.

The active route is updated as follows: the portion up to and including order υ^* is kept, and the remainder is replaced by r_{π^*}. The next decision time, h^D, occurs when the drone returns to the fulfillment center. We allow drone dispatches after $T - S$ subject to drone availability at the fulfillment center, to make sure that all orders are considered at least once for delivery.

4.3 Algorithmic components

The algorithmic components of the two strategies are further detailed in what follows.

4.3.1 Route generation

All routes are generated using the same mechanism. The initial route departs from the fulfillment center and carries a subset of the active orders, i.e., those orders placed prior to the vehicle's departure. The updated portion of the active route departing from a meeting point and carrying the on-board orders, if any, and the transfered orders, i.e., those orders brought to the delivery vehicle by the drone. The main difference is the starting point of the route. Note all generated routes have the the fulfillment center as the ending point, even though the vehicle only returns to the fulfillment center at the end of the day.

In the single-drone, single-vehicle variant, generating routes involves solving a traveling salesman problem with time windows (TSPTW), where an order's time window starts at the order's placement time and ends at the order's due time. Starting from an empty route, an initial route is generated using the *best insertion* paradigm. Then, inspired by the large neighborhood search scheme (Shaw, 1998; Schrimpf, 2000), in an iterative process, at each iteration, part of the current solution is destroyed and then reconstructed with the goal of minimizing due time violation and route duration. Due time violation is taken to be the primary goal, while minimizing route duration is taken as the secondary goal. Destruction consists of disconnecting a number q of orders from their current locations. We use a *worst removal* operator to destroy the current solution by identifying and removing the q worst placed orders (incurring the highest detour cost). The partial route is then reconstructed using *best insertion* operator, which selects the node with the lowest insertion cost, and inserts it in the cheapest position. See Dayarian et al. (2016) for more details on the routing algorithm used.

4.3.2 Order release

Recall that orders will not be assigned to (included in) the vehicle's route unless they can be delivered on time. Furthermore, the number of packages loaded on a drone is limited by the drone's capacity. As mentioned before, our approach to developing the routes to be executed is to first generates a route that includes all on-board, if any, and all active orders at the fulfillment center, and then, after doing so, remove a minimal subset of active orders to ensure that the remaining orders can be delivered on time and their number does not exceed the drone's capacity, P^D. A greedy approach is used to select active orders to be "released" (deleted) from the route. Active orders are released (one by one) based on the following hierarchical criteria. First, we look for the active order whose release results in the largest number of on-time orders in the resulting route. Second, in case of ties, we look for the order whose release results in the smallest total lateness of late orders in the resulting routes.

When there are no more late deliveries, but the number of packages transferred exceeds the drone capacity, a second heuristic is employed to reduce the number of transferred orders. This heuristic has two steps:

Step 1 Release orders in non-increasing order of their slack, where slack is the time between the arrival of the vehicle at the customer's location and the order's due time. If none of the orders has slack and the number of orders still exceeds the drone's capacity, then go to **Step 2**.

Step 2 Release orders in non-increasing order of their flexibility, where flexibility for order j is $\rho_j^o - (t_{0m}^D + t_{mj}^K)$, i.e., the difference between the due time of the order and an estimate of the time to reach its delivery location from the fulfillment center via the current meeting location m.

5 Computational Results

In this section, we conduct a series of computational studies to investigate the potential benefits of a same-day home delivery system with drone resupply, in terms of the number of orders served, for different service areas, different service guarantees, and different order arrival patterns. More specifically, we will evaluate the benefits of incorporating drone resupply in a home delivery system, and compare the performance of the restricted-resupply and flexible-resupply strategies.

5.1 Instance generation

We consider a set of randomly generated instances. Each instance has a planning horizon $T = 480$ minutes and is characterized by the following elements:

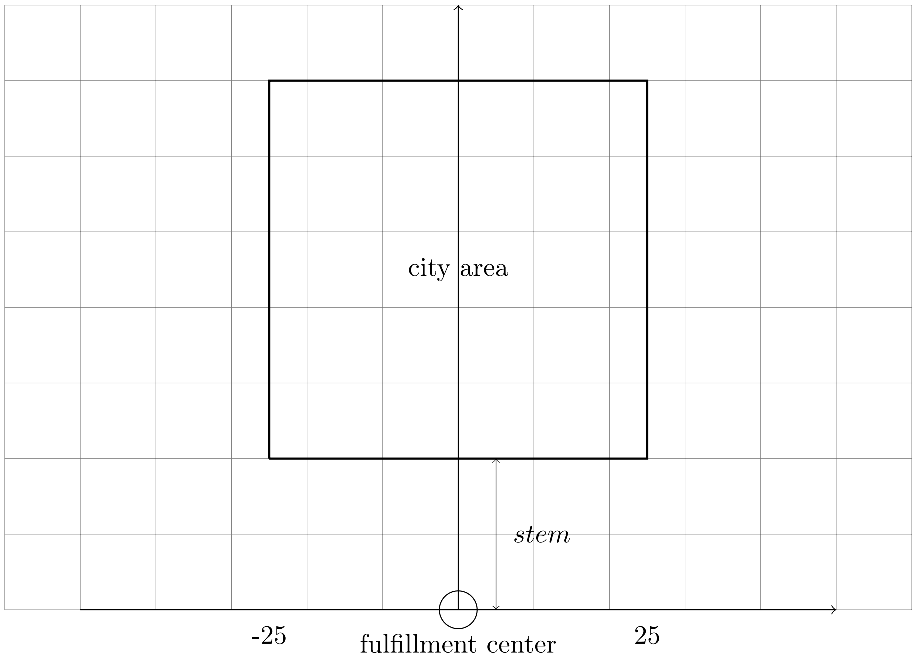

Figure 1: Geographical representation of an instance

- Fulfillment center: located at (0,0);
- Number of customers: N;
- Customer home locations: located uniform randomly in a 50×50 square region with its center at $(0, \mathit{stem} + 25)$, where *stem* represents the distance from the fulfillment center to the service area (see Figure 1);
- Number of meeting locations: M;
- Meeting locations: 70% located uniform randomly in the service area and 30% selected uniform randomly from the customer home locations;
- Order placement times: generated using a Poisson process with placement rate λ for each customer – the customer order placement lists are concatenated and then sorted in ascending order;

We generate eight instance classes as follows:

- Two values for *stem*: 10 and 40;
- Two sets of locations;
- For each set of locations, two order placement rates: $\lambda = 0.03$ and $\lambda/2$. The second set of order placements is a subset of the first set, with each order placement having 0.5 probability of occurring.

For each class, we generate 10 instances. Each instance is solved for service guarantee either $S = 60$ or $S = 120$ and drone capacity either $P^D = 50$ or $P^D = 30$. We set fixed epoch length $\delta = 10$, package transfer time $\tau_0 = 2$, and speed ratio $v^D/v^K = 1.5$, with v^K equal to 2 distance units per time unit.

5.2 Comparing the performance of different approaches

Tables 1 and 2 provide results for the different instance classes when solved with three strategies: (1) vehicle-only, (2) restricted-resupply, and (3) flexible-resupply, and with a drone capacity of 50 packages and the stationary package transfer option. In the vehicle-only strategy, the vehicle needs to return to the fulfillment center to pick up new orders before delivering them. We assume that the vehicle does not wait at the fulfillment center and is dispatched with the largest subset of active orders, allowing on-time deliveries. Each vehicle dispatch is designed in a similar manner that the first vehicle dispatch from the fulfillment center in the case of drone assisted strategies. Each row in Tables 1 and 2 represents the average over the 10 instances in each class. The strategies are compared based on the following metrics:

Cost: This represents the duty time of the vehicle and the travel time of the drone (if applicable). The duty time of the vehicle represents the time between the vehicle's first dispatch from the fulfillment center to the vehicle's last return to the fulfillment center. The drone's travel time only captures the time the drone was traveling between the fulfillment center and a meeting location.

Served %: This represents the percentage of the orders served, i.e., delivered on time (given the service guarantee).

time/dec (s): This represents the average time spend per dispatch decision. In the vehicle-only setting, this corresponds to vehicle dispatch decisions, while in the drone resupply settings, this is the average time spent deciding when to dispatch the drone and which active orders to transfer.

disp.: This represents the number of drone dispatches in the drone resupply settings.

Table 1: Results for instances in the 8 instance classes using different solution approaches, stationary transfer, drone capacity = 50, $S = 60$

			vehicle-only			**Restricted-resupply**						**Flexible-resupply**					
Stem	Instance class		Cost	Served %	time/dec (S)	Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)	Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)
Stem 10	Rate1	Loc1	525.3	19.8	1.0	507.7	137.6	645.3	31.8	10.8	10.5	507.3	134.2	641.5	32.4	11.2	25.3
		Loc2	533.0	22.3	0.9	509.7	114.3	624.0	32.2	10.9	8.6	511.4	117.5	628.9	31.6	11.2	26.0
	Rate2	Loc1	520.0	27.7	0.1	509.9	131.1	641.0	43.3	11.3	0.9	500.0	158.6	658.6	44.3	13.0	2.1
		Loc2	531.1	29.2	0.1	501.8	107.7	609.5	44.3	10.8	1.1	498.6	130.8	629.4	43.4	12.5	2.2
Stem 40	Rate1	Loc1	537.8	10.3	0.9	478.1	221.2	699.3	24.6	10.9	28.3	477.0	222.9	699.9	24.8	10.9	28.3
		Loc2	526.9	11.5	0.9	499.0	210.6	709.6	25.9	11.1	23.2	492.3	215.6	707.9	26.5	11.4	22.3
	Rate2	Loc1	525.0	13.3	0.1	481.5	222.5	704.0	33.2	11.1	2.8	477.4	226.5	703.9	33.5	11.4	2.1
		Loc2	523.6	14.7	0.1	479.9	214.2	694.1	34.9	11.4	2.6	480.4	218.3	698.7	35.8	11.6	2.6

Table 2: Results for instances in the 8 instance classes using different solution approaches, stationary transfer, drone capacity = 50, $S = 120$

			vehicle-only			**Restricted-resupply**						**Flexible-resupply**					
Stem	Instance class		Cost	Served %	time/dec (S)	Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)	Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)
Stem 10	Rate1	Loc1	598.6	39.3	5.0	581.0	68.5	649.5	54.7	5.4	20.6	580.1	86.0	666.1	55.6	6.9	143.9
		Loc2	595.1	43.9	4.6	586.9	58.6	645.5	53.7	5.5	29.2	576.5	71.3	647.8	53.5	7.4	243.2
	Rate2	Loc1	583.8	52.3	0.6	582.6	73.5	656.1	71.1	6.1	3.1	572.4	100.8	673.2	74.1	8.7	11.7
		Loc2	593.1	54.8	0.6	579.3	61.9	641.2	69.8	6.1	4.0	569.7	94.3	664.0	71.6	9.2	18.6
Stem 40	Rate1	Loc1	600.9	29.5	4.8	595.3	135.9	731.2	49.8	6.0	36.9	591.1	153.9	745.0	52.3	7.1	275.4
		Loc2	595.2	28.3	4.6	601.5	124.2	725.7	48.9	5.7	52.2	595.7	143.1	738.8	49.8	7.1	215.0
	Rate2	Loc1	602.3	35.1	0.6	591.3	140.6	731.9	63.9	6.2	3.2	588.0	168.1	756.1	68.7	7.7	19.6
		Loc2	607.1	37.6	0.7	604.6	118.0	722.6	64.1	6.0	10.1	580.8	144.0	724.8	67.3	7.5	20.6

The results show that there are significant benefits to employing drone resupply strategies. The reasons for this are:

1. The use of drone resupply allows more efficient use of the delivery vehicle, which results in a significant increase in the percentage of orders served.

2. The use of drone resupply allows for expansion of the service area. Without drone resupply, there is a significant drop in the percentage of orders served when the stem is increased from 10 to 40, but the drop in the percentage of orders served is minimal when drone resupply is employed.

The results also show that the flexible-resupply strategy performs better than the restricted-resupply strategy, but that the added flexibility does not result in a large increase in the percentage of orders served, especially not when the service guarantee is tight. When the service guarantee is tight ($S = 60$), then little and often no "buffer" time exists between consecutive visits to customer home locations along the route. This decreases the flexibility of the decision maker to insert a drone resupply between two consecutive deliveries. As a consequence, also in the flexible-resupply strategy, a drone resupply is typically pushed to the end of the route, after delivering all on-board orders. We see too that the average number of drone dispatches is higher when the flexible-resupply strategy is used. Resupplying a vehicle along its route rather than waiting for all on-board orders to be delivered before resupplying the vehicle may allow the drone to return to the fulfillment center earlier. An earlier return to the fulfillment center allows for more frequent drone dispatches, but will also provide more flexibility to handle unexpected variations in order placements (which was not part of our experiments). Of course, more frequent drone dispatches increase the cost associated with employing drones.

As expected, a larger order placement rate decreases the percentage of order served. The impact of customer locations, on the other hand, appears to be minimal.

We next study the impact of the chosen package transfer option and the drone capacity. Because the results presented above suggest that flexible-resupply outperforms restricted-resupply, in this set of computational experiments we only use flexible-resupply. The results can be found in Tables 3 and 4.

The results indicate that in-motion package transfer allows the on-time delivery of a larger number of orders. This is due to the fact that when the package transfer takes place while the vehicle is at the meeting location (stationary transfer), a possibly smaller number of orders can be accommodated in the vehicle route following the resupply compared to when package transfer takes place while the vehicle is in motion (in-motion transfer). Even though in-motion package transfer may cause the return of the drone to the fulfillment center to be delayed, the impact appears to be less significant (likely because the drone has a higher speed than the delivery vehicle).

The results also show that when the service guarantee is aggressive ($S = 60$), the drone capacity has a relatively small impact on the number of packages delivered; the percentage of orders served when the drone capacity is 30 is identical to the percentage of orders served when the drone capacity is 50. However, when the service guarantee is less aggressive ($S = 120$), drone capacity does matter, especially when the order placement rate is high; the number of orders served when the drone capacity is 30 can be more than 20 percent less than the number of orders served when the drone capacity is 50. As expected, when the order placement rate is low, the impact of a smaller drone capacity is much less detrimental.

Table 3: Results for instances in the 8 instance classes using flexible-resupply, in-motion transfer, $S = 60$

			Drone capacity = 50						**Drone capacity = 30**					
Stem	Instance class		Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)	Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)
Stem 10	Rate1	Loc1	530.6	148.3	678.9	34.5	11.4	14.6	530.6	148.3	678.9	34.5	11.4	678.9
		Loc2	523.6	131.4	655.0	35.4	12.4	17.9	523.6	131.4	655.0	35.4	12.4	18.4
	Rate2	Loc1	527.0	153.9	680.9	48.5	12.5	2.2	527.0	153.9	680.9	48.5	12.5	2.4
		Loc2	524.9	146.8	671.7	49.6	13.8	1.8	524.9	146.8	671.7	49.6	13.8	1.9
Stem 40	Rate1	Loc1	486.9	201.6	688.5	25.6	10.1	32.8	486.9	201.6	688.5	25.6	10.1	32.2
		Loc2	501.5	205.6	707.1	29.2	11.0	28.5	501.5	205.6	707.1	29.2	11.0	29.8
	Rate2	Loc1	464.7	199.5	664.2	31.4	10.2	3.5	464.7	199.5	664.2	31.4	10.2	3.9
		Loc2	494.8	206.8	701.6	38.0	11.2	3.8	494.8	206.8	701.6	38.0	11.2	4.3

Table 4: Results for instances in the 8 instance classes using flexible-resupply, in-motion transfer, $S = 120$

			Drone capacity = 50						**Drone capacity = 30**					
Stem	Instance class		Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)	Vehicle cost	Drone Cost	Total cost	Served %	# disp.	time/dec (s)
Stem 10	Rate1	Loc1	602.8	96.9	699.7	58.9	8.0	216.6	593.9	91.5	685.4	48.4	8.1	685.4
		Loc2	593.0	81.3	674.3	55.2	7.5	204.7	594.8	83.6	678.4	43.7	7.7	443.6
	Rate2	Loc1	587.2	109.6	696.8	77.3	9.1	10.6	585.1	109.7	694.8	76.5	9.1	9.8
		Loc2	586.1	105.5	691.6	75.2	9.7	15.4	588.1	101.4	689.5	74.1	9.5	18.4
Stem 40	Rate1	Loc1	603.2	167.5	770.7	55.6	7.6	256.9	605.9	163.3	769.2	46.0	7.5	242.6
		Loc2	599.2	154.4	753.6	52.4	7.3	317.6	596.8	151.9	748.7	44.7	7.5	429.3
	Rate2	Loc1	598.2	171.0	769.2	72.8	7.9	17.3	598.9	168.7	767.6	72.0	7.8	17.5
		Loc2	597.8	159.2	757.0	72.0	8.1	15.0	598.3	153.6	751.9	71.2	7.9	18.0

6 Discussion

We introduced and (empirically) analyzed a novel drone-assisted last-mile delivery system. The proposed system differs from other drone-assisted last-mile delivery systems that have been investigated in terms of the drone's responsibility. Rather than being responsible for the final delivery to the customer, in our proposed system, the drone is responsible for resupplying the delivery vehicle. The proposed system has two important benefits. First, it overcomes any accessibility restrictions associated with residential/commercial locations in dense urban areas. Secondly, since larger drones will have to be employed, safety will be enhanced as the drone will be more stable and less vulnerable to adverse weather conditions.

We focused on a simplified scenario in which a single drone and a single vehicle collaboratively perform deliveries in an urban area. The insights obtained, however, are meaningful, as delivery systems in large cities often divide the service area into smaller zones, and where each zone is assigned to a single driver (see, for example, Huang et al. (2017)). In such systems, a zone may be served more effectively by a single drone and a single vehicle.

We presented an initial study of the potential benefits of drone resupply in last-mile delivery systems focused on the decision making aspects, but ignoring certain practical aspects. Further analysis is needed to investigate, for example, the impact of vehicle capacity, drone loading time at the fulfillment center, drone battery recharge/exchange at the fulfillment center, and drone travel range. Most of these considerations, however, can easily be incorporated in the current approach.

Using drones to resupply delivery vehicles (rather than, for example, a dedicated fleet of resupply vehicles) has two main advantages. First, drones are independent autonomous aerial vehicles which offer cost savings. Second, drones can likely travel faster than vehicles, as they are not (less) constrained by the road network and traffic congestion.

References

M. Adjrad and P. Groves. 3d-mapping-aided gnss exploiting galileo for better accuracy in dense urban environments. In *2017 European Navigation Conference (ENC)*, pages 108–118, 2017.

N. Agatz, P. Bouman, and M. Schmidt. Optimization approaches for the traveling salesman problem with drone. *ERIM Report Series Reference No. ERS-2015-011-LIS*, 2016. URL https://ssrn.com/abstract=2639672.

G. S. C. Avellar, G. A. S. Pereira, L. C. A. Pimenta, and P. Iscold. Multi-uav routing for area coverage and remote sensing with minimum time. *Sensors*, 15(11):27783–27803, 2015. ISSN 1424-8220.

A. Barrientos, J. Colorado, J. d. Cerro, A. Martinez, C. Rossi, D. Sanz, and J. Valente. Aerial remote sensing in agriculture: A practical approach to area coverage and path planning for fleets of mini aerial robots. *Journal of Field Robotics*, 28(5):667–689, 2011. ISSN 1556-4967.

A. Baskaran and A. Kashef. Investigation of air flow around buildings using computational fluid dynamics techniques. *Engineering Structures*, 18(11):861 – 875, 1996. ISSN 0141-0296.

W.-R. Bretzke. Global urbanization: a major challenge for logistics. *Logistics Research*, 6(2):57–62, 2013. ISSN 1865-0368.

J. Campbell, D. Sweeney, and J. Zhang. Strategic design for delivery with trucks and drones. 2017. URL https://www.researchgate.net/publication/316169890_Strategic_Design_for_Delivery_with_Trucks_and_Drones.

J. C. Da C. Siqueira. *Modeling of Wind Phenomena and Analysis of Their Effects on UAV Trajectory Tracking Performance*. PhD thesis, West Virginia University, 2017.

I. Dayarian and M. Savelsbergh. Crowdshipping and same-day delivery: Employing in-store customers to deliver online orders. Technical report, 2017. optimization-online 2017/07/6142.

I. Dayarian, T. G. Crainic, M. Gendreau, and W. Rei. An adaptive large-neighborhood search heuristic for a multi-period vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 95:95 – 123, 2016. ISSN 1366-5545.

K. Dorling, J. Heinrichs, G. G. Messier, and S. Magierowski. Vehicle routing problems for drone delivery. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(1):70–85, 2017.

Michael Drexl. Synchronization in vehicle routing—a survey of vrps with multiple synchronization constraints. *Transportation Science*, 46(3):297–316, 2012.

S. M. Ferrandez, T. Harbison, T. Weber, R. Sturges, and R. Rich. Optimization of a truck-drone in tandem delivery network using k-means and genetic algorithm. *Journal of Industrial Engineering and Management*, 9(2), 2016.

M. Francisco. Organ delivery by 1,000 drones. *Nature Biotechnology*, 34(684), 2016. ISSN 1087-0156.

G. Gozbayin, O. Karasan, M. Savelsbergh, and H. Yaman. A branch-and-price algorithm for the vehicle routing problem with roaming deliveries. *Transportation Research Part B*, 100:115–137, 2017.

Q. M. Ha, Y. Deville, Q-. D. Pham, and M. H. Hà. Heuristic methods for the traveling salesman problem with drone. September 2015. URL https://www.researchgate.net/publication/282403410_Heuristic_methods_for_the_Traveling_Salesman_Problem_with_Drone.

Q. M. Ha, Y. Deville, Q-. D. Pham, and M. H. Hà. On the min-cost traveling salesman problem with drone. *CoRR*, abs/1509.08764, 2016. URL http://arxiv.org/abs/1509.08764.

J. Hegranes. Separating fiction from feasibility in the future of drone delivery. *Tech Crunch*, August 2017. URL https://techcrunch.com/2017/08/15/separating-fiction-from-feasibility-in-the-future-of-drone-delivery/.

Y. Huang, M. Savelsbergh, and L. Zhao. Business-to-consumer e-commerce: Home delivery in megacities. 2017. Optimization Online 2017/08/6170.

M. Joerss, J. Schröder, F. Neuhaus, C. Klink, and F. Mann. Parcel delivery, the future of last mile. *Travel, Transport & Logistics | McKinsey & Company*, September 2016.

Y. Khosiawan and I. Nielsen. A system of uav application in indoor environment. *Production & Manufacturing Research*, 4(1):2–22, 2016.

S. J. Kim, G. J. Lim, J. Cho, and M. J. Côté. Drone-aided healthcare services for patients with chronic diseases in rural areas. *Journal of Intelligent & Robotic Systems*, 2017. ISSN 1573-0409.

C.K.Y. Lin. A vehicle routing problem with pickup and delivery time windows, and coordination of transportable resources. *Computers & Operations Research*, 38(11):1596 – 1609, 2011. ISSN 0305-0548.

J. Markoff. Drones marshaled to drop lifesaving supplies over rwandan terrain. *New York Times*, April 2016. URL http://www.nytimes.com/2016/04/05/technology/drones-marshaled-to-droplifesaving-supplies-over-rwandan-terrain.html?_r=0.

J. McDonald. One of china's biggest online retailers is developing drones that can carry over a ton. *Associated Press*, 2017.

C. C. Murray and A. G. Chu. The flying sidekick traveling salesman problem: Optimization of drone-assisted parcel delivery. *Transportation Research Part C: Emerging Technologies*, 54:86 – 109, 2015. ISSN 0968-090X.

D. Pierce. Delivery drones are coming: Jeff bezos promises half-hour shipping with amazon prime air. *The Verge*, December 2013. URL https://www.theverge.com/2013/12/1/5164340/delivery-drones-are-coming-jeff-bezos-previews-half-hour-shipping.

A. Ponza. *Optimization of Drone-Assisted Parcel Delivery*. PhD thesis, Universita Degli Studi Di Padova, 2016.

Damián Reyes, MWP Savelsbergh, and Alejandro Toriello. Vehicle routing with roaming delivery locations. *Transportation Research Part C*, 80:71–91, 2016.

G. Schrimpf. Record breaking optimization results using the ruin and recreate principle. *Journal of Computational Physics*, 159(2):139–171, 2000.

P. Shaw. Using constraint programming and local search methods to solve vehicle routing problems. In M. Maher and J.-F. Puget, editors, *Principles and Practice of Constraint Programming – CP98*, volume 1520 of *Lecture Notes in Computer Science*, pages 417–431. Springer Berlin Heidelberg, 1998.

K. Steemers and S. Yannas. Architecture, city, environment. In *Proceedings of PLEA 2000*, Cambridge, United Kingdom, 2000. Routledge.

C. A. Thiels, J. M. Aho, S. P. Zietlow, and D. H. Jenkins. Use of unmanned aerial vehicles for medical product transport. *Air Medical Journal*, 34(2):104 – 108, 2015. ISSN 1067-991X.

C. van Cooten. Crowdsourced delivery : the traditional delivery method reinvented. Master's thesis, Technische Universiteit Eindhoven, 2016.

X. Wang, S. Poikonen, and B. Golden. The vehicle routing problem with drones: several worst-case results. *Optimization Letters*, 11(4):679–697, 2017. ISSN 1862-4480.

R. Weibel and R. J. Hansman. chapter Safety Considerations for Operation of Different Classes of UAVs in the NAS. Infotech@Aerospace Conferences. American Institute of Aeronautics and Astronautics, Sep 2004.

J. C. Wong. Delivery robots: a revolutionary step or sidewalk-clogging nightmare? *The Guardian*, 2017.

E. Yakıcı. Solving location and routing problem for uavs. *Computers & Industrial Engineering*, 102: 294 – 301, 2016. ISSN 0360-8352.

S. Zarovy, M. Costello, and A. Mehta. Experimental method for studying gust effects on micro rotorcraft. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 227(4):703–713, 2013.