

Routing of Heterogeneous Fleets for Flash Deliveries via Vehicle Group Assignment

Maximilian Kronmueller*, Andres Fielbaum* and Javier Alonso-Mora*

Abstract—This paper presents a novel approach to route heterogeneous fleets for flash delivery operations. Flash deliveries offer to serve customers’ wishes in minutes. We investigate a scenario that allows to pick up orders at multiple depots with a heterogeneous vehicle fleet leveraging different modes of transportation. We propose the Heterogeneous Vehicle Group Assignment (HVGA) method, which, given a problem state, identifies potential pick-up locations, calculates potential trips for all modes of transportation and last chooses from the set of potential trips. Experiments to analyze the proposed method are executed using a fleet featuring two modes of transportation, trucks and drones. We compare to a state-of-the-art method. Results show that HVGA is able to serve more orders while requiring less total traveled distance. Further, the effects of the fleet size and fleet composition between drones and trucks are examined by simulating three hours of a flash delivery operation in the city center of Amsterdam.

I. INTRODUCTION

Using heterogeneous fleets for last-mile delivery operations allows leveraging the strengths of different modes of transportation. For example, a drone can maneuver independently of roads and traffic and thus deliver quickly in hard-to-reach areas and a truck can load many parcels simultaneously and deliver multiple orders in one neighborhood. In addition, grocery deliveries within minutes have established themselves within many cities. In the Netherlands alone, consumers spent around 40 million euros per month on flash deliveries at the end of 2021, a trend that is continuously rising [1]. Combining these two aspects poses a highly relevant and interesting question: How do we route heterogeneous fleets for on-demand last-mile deliveries?

This paper proposes a novel optimization-based approach to route heterogeneous fleets for on-demand last-mile deliveries, considering multiple depots and short delivery times. The proposed method is able to combine many types and forms of transportation while staying scalable. We specifically investigate the use case of supporting ground-based vehicles with drones.

The main virtues are threefold. First, we generalize a well-known method for ride-sharing, called Vehicle Group Assignment, to be able to handle heterogeneous fleets, including various modes of transportation. Second, this work combines dynamic multi-depot vehicle routing with heterogeneous fleets. Third, to the best of our knowledge, this work is first in investigating the use of drones for flash delivery operations.

II. RELATED WORK

The problem we face can be classified as a dynamic and heterogeneous vehicle routing problem (pick-up and delivery problem). Thus it is related to the two broad fields of Dynamic Vehicle Routing Problems (DVRP) and Vehicle Routing Problems with heterogeneous fleets. For an overview of dynamic routing problems, see [2]–[5], and for an overview of routing for heterogeneous fleets see [6]. Heterogeneous routing problems are further divided by the type of vehicles they are considering and by the way these can interact with each other. For a more detailed analysis, we exclusively focus on dynamic problems, as approaches deployed between static and dynamic vehicle routing problems differ strongly.

Proposed methods to tackle the dynamic vehicle routing problem with a heterogeneous fleet are mainly based on heuristics. Large Neighbourhood Search based strategies were applied by [7] considering a fleet of trucks with different capabilities and by [8] tackling the technician routing problem (technicians differ in skills and repair parts carried). [9] used a Tabu Search to improve routes while dynamically integrating new incoming orders. They consider different vehicle speeds and capacities. A problem of dynamically refuelling aeroplanes by trucks with different speeds, capacities and fit to planes was studied by [10] using a genetic ant colony algorithm.

More specifically, the problem studied in this paper is a specific case of a Same-Day Delivery (SDD) problem [11]–[16]. Works that combine an SDD problem and a heterogeneous fleet are sparse. We identified three works, analyzed in the following.

[13] proposes a method routing a heterogeneous fleet of trucks and drones for a Same-Day Delivery (SDD) operation, splitting it into two. For each new order, it is decided if it is served by a drone or a truck or is ignored. Subsequently, each mode of transportation is routed independently. To determine if a drone or a truck should be used, they apply policy function approximation based on geographical districting. Orders which have a travel time longer than a given threshold are preferably served by drones. [15] builds upon [13] by adapting the way orders are assigned to a mode of transportation. They leverage Q-learning to do so and improve on previous results. Improvements are most substantial for small fleets but vanish fully if more resources are provided. [17] investigates a SDD problem, in which drones are used to resupply trucks, which exclusively deliver to customers. They examine two approaches for resupplying: first, only empty trucks can receive a resupply and second, resupplies are

*Delft University of Technology, Mekelweg 5, 2628 CD Delft, The Netherlands m.kronmueller@tudelft.nl

possible anytime. They optimize on finding the best resupply locations, serving as many orders as possible.

In contrast, this work investigates a specific variation of SDD, a flash delivery operation featuring multiple depots. HVGA further differs from the works proposed for heterogeneous SDD by not splitting into assigning orders to vehicles and routing but does so jointly.

The proposed approach in this paper is building on a ride-sharing method called Vehicle Group Assignment (VGA) [18], which shows great scalability. A retail context variation considering multiple depots was investigated by [19]. This work extends VGA to heterogeneous fleets.

As we test our method considering scenarios in which drones and trucks are utilized, we want to point interested readers to an overview on routing problems featuring the usage of drones (mostly static), see [20].

III. PROBLEM DESCRIPTION

This section presents the problem statement covering the whole operation and introduces the used notation (Section III-A). As our problem is dynamic, we model it as a Markov Decision Process (MDP), explicitly capturing the problem's dynamics (Section III-B).

A. Notation and Problem Statement

We consider a heterogeneous fleet of vehicles V consisting of multiple sub-fleets, each having a different mean of transport, noted by the subscript m . Each vehicle, independent of its mode, is denoted by $v \in V$. For the sake of simplicity, we describe the method assuming a truck fleet L of H identical trucks and a drone fleet D of G identical drones. The generalization towards more than two sub-fleets is straightforward. Vehicles move on a weighted and directed graph $G_m = (N; A_m)$, thereby the arcs and arcs' weights are mode-dependent. Apart from the network they travel on, the fleets differ in their maximum capacity of each vehicle C_m and their traveling speed S_m . A set of depots De , locations where orders can be picked-up, is placed within the graph at specific nodes $\{de_1; \dots; de_j; De_j\}^1$. Each customer is defined as an order o . Customers can request the delivery of goods to a location of their desire. We denote the goal location by g_o , which is assumed to be a node. The time a customer orders is called the request time t_o . All orders are summarized within the demand O . Each order is unique and has a size of one. We additionally introduce time-dependent demand subsets: First, the set of loaded orders O_t^{load} consisting of all orders currently loaded to a truck or drone. Further, O_t^{deli} denotes the set of delivered orders and O_t^{igno} the set of ignored orders (i.e., orders that were not delivered within the given constraints). Last, the set of known open orders O_t^{open} consisting of all orders currently known and unloaded nor delivered nor ignored.

The operation starts at $t = 0$ and ends at $t = T$, with t being the current time. During a short time span of duration

τ before T no new orders are accepted. Further, τ_{pick} and τ_{drop} are fixed time spans needed to load or to deliver an order.

General Problem Statement: Consider a heterogeneous fleet V consisting of two sub-fleets, a truck fleet L of H identical trucks and a drone fleet D of G identical drones. Each mode of transportation operates on a weighted and directed graph $G_m = (N; A_m)$. Multiple depots $de \in De$ are placed within the graph. Each truck and drone is in a known initial state (starting locations and no prior load). Customers continuously place orders, summarized as the demand O , where each order is specified by a request time $t_o \in [0; T - \tau]$ and a goal location g_o . The operation starts at $t = 0$ and ends at $t = T$. Find a set of routes for each truck and drone to pick up, thus including the choice of a depot, and to deliver the orders such that a given cost function J is minimized subject to a set of given constraints.

We consider two constraints: First, each vehicle's maximum capacity C_m may not be exceeded. Second, each order has a maximum delivery time $t_{o,max;m}$. We calculate the maximum delivery time as the sum of the request time t_o , the optimal delivery time $t_{o,opt;m}$ and a fixed time interval, called maximum delay τ_{max} , i.e., $t_{o,max;m} = t_o + t_{o,opt;m} + \tau_{max}$. Thereby, the optimal delivery time $t_{o,opt;m}$ is calculated as the time a vehicle needs to travel from the closest depot to the goal location of the order plus the time to pick-up τ_{pick} and to deliver τ_{drop} the order. Note that the optimal delivery time is dependent on the used mode of transportation. An order might not be delivered within the given constraints and is thus considered ignored.

B. Markov Decision Process

An MDP captures the dynamics of a problem by modelling subsequent states S_t connected by a decision taken and a transition between them.

Decision Points: Decisions are taken at specific points in time, summarized in the set \mathcal{T} . Individual decisions and according states are enumerated by k . We make decisions in fixed time steps of Δt , i.e., $t_{k+1} = t_k + \Delta t$.

Problem State: The problem state S_k at time t_k is fully characterized by the time itself t_k , the state of both sub-fleets and the set of open orders O_t^{open} . To describe the fleet's states each individual vehicle v is described by its location $l_{v;t}$, its set of loaded orders $O_{v;t}^{load}$ and the plan it follows currently $\psi_{v;t}$. For example, the truck fleet's state is described as $L_t = ((l_{1;t}; O_{1;t}^{load}; \psi_{1;t}) \mid \dots \mid (l_{H;t}; O_{H;t}^{load}; \psi_{H;t})) \mid L$. This results in the overall state definition as

$$S_t = (t; L_t; D_t; O_t^{open}).$$

Note that a plan $\psi_{v;t}$ consists of an ordered sequence of actions (pick-up and drop-off of orders or rebalancing) with associated locations. Between locations, the vehicle follows the shortest path. As soon as an action is executed, it is

¹Also shared between trucks and drones.

² ψ can be determined during operation or beforehand.

removed from the plan. Further, a plan can be updated at a later point in time.

Decision: The decision at decision point k is to update the plan $v;t$ of each vehicle v , which it will follow until the next decision point t_{k+1} . Note that a vehicle's plan can change if a subsequent decision updates it, including newly obtained information.

Transition: The transition can be split into a deterministic part and an unknown part. In the deterministic part, we update the truck and drone fleet's status. Each vehicle follows its plan $v;t$ determined within the taken decision. Doing so, a vehicle's location and its loaded orders change (orders can get delivered and new orders can be loaded). As time propagates, between subsequent states, customers may place new orders (which is the unknown aspect). As a result, new orders are added to the set of open orders O_{k+1}^{open} . If an order can not be delivered within its constraints anymore, we consider it ignored. The order is removed from the set of open orders³ O_{k+1}^{open} .

Objective: The goal of the posed problem is to minimize a combination of costs, considering the total driven distance, service quality measured as the delay d_o , and a penalty term for orders that are not delivered (Equation 1).

$$J_T = 4 \left(\sum_{o \in O_{k+1}^{deli}} \dots \right) \times \sum_{o \in O_{k+1}^{open}} d_o + \sum_{v \in V} v;T + \sum_{o \in O_{k+1}^{igno}} \dots \quad (1)$$

Thereby, the delay d_o is defined per order as the difference between the actual delivery time and the optimal delivery time. The total driven distance of a vehicle $v \in V$ at time t is denoted by $v;t$. α is a constant predefined penalty for ignoring an order. Note that if α is set to a large constant, the objective function puts the highest priority onto serving as many orders as possible. This penalty can be interpreted as the cost to hire a third party to deliver the respective order, such that no customer is neglected. β is a tuneable weighting parameter between operator cost and service quality. Note that we can not directly map a decision taken at a specific state S_t to Equation 1. Each vehicle will follow its plan $v;t$ until the next decision state, then each plan may be altered due to new information. Thus, a current plan can not be directly evaluated by Equation 1.

IV. METHOD

In this section, we describe the proposed method, called Heterogeneous Vehicle Group Assignment (HVGA). Following the MDP (Section III-B), each time a decision point is hit, a decision given the current state S_k , determining a set of routes $v;t$ for every vehicle $v \in V$, needs to be taken. To do so, we follow a sequence of steps explained in the following. First, we determine multiple potential pick-up locations for each order (Section IV-A). Second, we calculate sets of potential trips each truck and drone could take (Section IV-B). Last, we pick from those sets of potential

³Note that when vehicles load orders following their plan, they also get removed from the set of open orders.

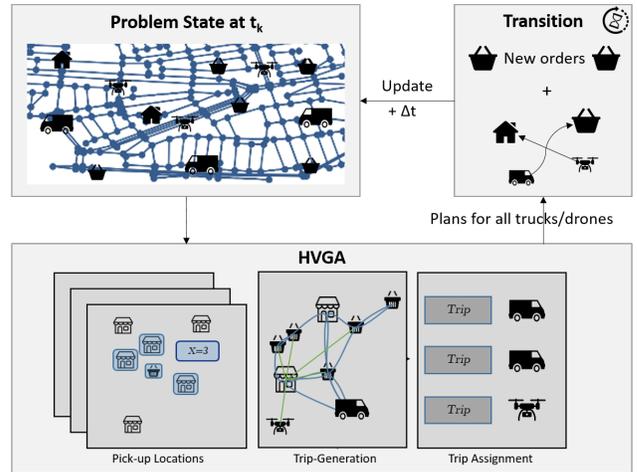


Fig. 1: Given a problem state, a decision is taken on how to update the plans of all trucks and drones. Heterogeneous Vehicle Group Assignment (HVGA) creates plans for all vehicles by first identifying potential pick-up locations for each order. Subsequently, a large set of potential trips is calculated for each vehicle. Which of these trips are executed is determined in the last step called trip assignment. During a transition phase, all trucks and drones follow their determined plan and new orders are received until a new decision state is met.

trips, which each vehicle should carry out by solving an Integer Linear Problem (ILP) (Section IV-C).⁴ An illustrative overview is depicted in Figure 1. The proposed method builds on top of previous work for flash deliveries for homogeneous fleets, introduced in [19].⁵

A. Selecting Potential Pick-up Locations

As orders only specify their delivery location g_o , the decision of where to pick up the corresponding goods is raised. To acknowledge this, we introduce a concept called candidate. A candidate combines an order with a potential pick-up location. We define it as follows: A candidate c is a tuple containing an order $o_c \in O$ and an associated pick-up location $p_c \in De$. Thus a candidate is described as $c = (o_c; p_c)$. Each order can have multiple candidates associated with it. Further, we introduce a heuristic that reduces the number of candidates by only considering a subset of depots. We consider the x depots closest in delivery time by truck to the orders goal location. x is a predetermined tuneable parameter.

B. Finding Potential Trips

A trip, denoted as t , is defined as a set of candidates $\{c_1; \dots; c_j\}$, a vehicle and a plan, that serves all candidates

⁴Note that the proposed method does not take any future information into account and thus operates myopic. Doing so could improve obtained results [21], [22], but is regarded as future work.

⁵[19] also provides a more detailed analysis of their proposed method. Interested readers are referred to their work for further details.

of the trip, $\tau = (f; c_1; \dots; c_j; g; v; v; t)$. The goal of the trip generation step is to find all potential trips for each truck and drone, summarized as the set Z^{all} . This set is formed by combining the sets of all potential trips Z_v for each vehicle, i.e. $Z^{all} = \bigcup_{v \in V} Z_v$. We calculate the sets Z_v separately for each truck and drone and thus also separately for each mode of transportation. Each trip is further tagged with the cost to execute it.

The general workflow for a truck and drone are identical but differ in the used network G_m , speed s_m and vehicle's capacity C_m . If a trip τ is feasible for a specific vehicle v , a plan $\tau_{v;t}$ can be found, picking up and delivering all candidates of this trip without violating any constraint. This includes orders that are already on board of the considered vehicle. To generate the complete set of all feasible trips Z_v , the method builds onto the idea that a trip can only be feasible if all its sub-trips are feasible as well. As a result, we start searching for trips of size one. Subsequently, we combine obtained trips to form larger trips, successively increasing in size. The cost of a trip τ is given by $c(\tau)$ and is derived from Equation 1:

$$c(\tau) = (1 - \sum_{o \in \tau} x_o) \times \sum_{o \in \tau} c_o(\tau) + travel(\tau) \times s_m \quad (2)$$

$travel(\tau)$ determines the required distance traveled to serve the according trip τ . Note that the delay δ_o of an order, as well as the time needed to complete a trip, depends on the trip and the used mode of transportation. The way to navigate between different trip stops is determined by the used graph G_m , and the time needed to traverse a given arc is shaped by the speed of the vehicle s_m . To determine the cost and plan $\tau_{v;t}$ of a trip τ , we perform an exhaustive search on all possible sequences and continue solely with the cheapest option.

C. Assigning Trips

Given the set Z^{all} of all potential trips each truck and drone can take, this step decides which of them are executed. Thereby we want to coordinate the individual decisions to maximize performance and minimize cost. This problem is formalized and solved as an Integer Linear Problem (ILP), see Equations 3-6.

$$\text{argmin} \sum_{\tau \in Z^{all}} (c(\tau) \times x_{\tau}) + \sum_{o \in O_t^{open}} c_o \times o \quad (3)$$

$$\sum_{\tau \in Z_v} x_{\tau} \leq 1 \quad \forall v \in V \quad (4)$$

$$\sum_{\tau \in Z^{all}} x_{\tau} \times c_o = 1 \quad \forall o \in O \quad (5)$$

$$x_{\tau} \in \{0, 1\} \quad \text{and} \quad o \in O \quad (6)$$

Equation 3 presents the cost function. Note it differs slightly from Equation 2 as we only account for changes in vehicles plans by subtracting the costs needed to serve the already loaded orders of the according vehicle $loaded;v$.

We introduce the constraints that each truck and drone is maximally used once (Equation 4). Also, each order is maximally served once or ignored for now (Equation 5). Equation 6 introduces the binary variable x_{τ} , which takes the value of one if a trip τ is executed; and the binary variable o taking the value one if an order o is not served by the chosen trips.⁶

As a result, each truck and drone are either assigned a new trip, which they execute until the next decision is taken, or they follow their plan as previously determined. If a vehicle becomes idle (i.e., it has no orders to serve after the trips have been assigned), we assign it a special plan, we call rebalancing. The vehicle is routed to its closest depot, such that it is in a promising position for future decisions.

Remark: Note that HVGA has a large potential to unify different types of vehicles and modes of transportation due to assigning a cost to each trip and using it to make decisions. Trips and their associated costs can thereby be calculated using entirely different approaches. If approaches are similar, this ensures more straightforward comparability. For example, one can only adapt the graph vehicles operate on, from a road network to a water canal system or the speed can be adjusted to the capabilities of an individual vehicle.

V. RESULTS

First, we compare HVGA to a state-of-the-art method to evaluate its performance in Section V-A. Section V-B investigates the effect of the size and composition of a heterogeneous fleet of trucks and drones.

All experiments presented in this section use a time window of 3 hours and 10 minutes, whereby no new orders are placed within the last 10 minutes. During this time, 1828 orders are placed randomly within the operation area.⁷ As an operation area, we simulate the city centre of Amsterdam, with trucks driving along the street network with a speed of 10 meters per second, and assuming that drones can fly directly to their goal location with a speed of 15 meters per second. Trucks can load up to six orders simultaneously, whereas drones have a maximum capacity of one. Each order is allowed a maximum delay of 8 minutes. To load an order 15 seconds are needed and 30 seconds to deliver it. HVGA was performed every 100 seconds, i.e., $t = 100$. The penalty for ignoring an order was chosen as a high constant $\beta = 10000$ [sec] and the cost weighting parameter as $\alpha = 0.3333$.

As a first impression, a snapshot of an area of a particular state is depicted in Figure 2. The plan for one truck and one drone are highlighted.

A. Comparison to Other Approaches

Next to HVGA, we implemented a second approach resembling the work by Ulmer et. al. [13]. This method first decides which mode of transportation is going to be used for which order, followed by a separate routing step. To

⁶To solve this problem, standard software, like Mosek or Gurobi, can be used. We used Mosek [23].

⁷We keep the used demand distribution and number constant to allow for better comparison between all analyzed scenarios.

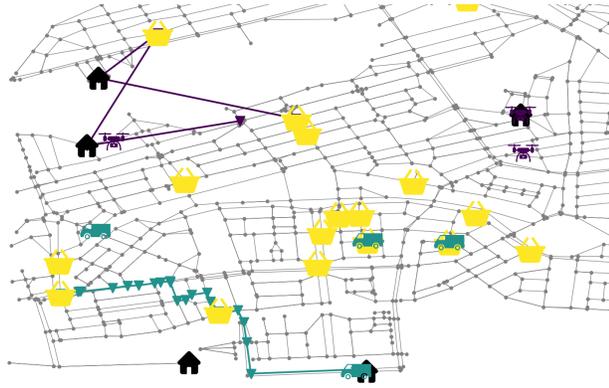


Fig. 2: A cutout of one particular state of the simulation is illustrated. The road network graph, on which trucks move is shown in grey. The plan of one drone and one truck are highlighted in purple and green. Goal locations of known and non-delivered orders are shown in yellow.

assign the mode of transportation, they use a policy function approximation based on geographical districting, preferably serving orders with long travel times by drones. To allow for a fair comparison, we adapted the approach [13] to use the routing method as introduced in this paper. Details on this adapted approach can be found in the Appendix (Section VII). In the following, we call this approach Split&Route. Here we investigate a scenario in which a single depot placed in the centre of the graph is used, as Split&Route is not specifically designed to work with multiple depots. Obtained results are depicted in Figure 3. HVGA improves on Split&Route by increasing the service rate by about 17%. Despite serving more orders, the total driven distance reduces by 108 km. On the other hand, the average delay increases by 24 seconds. Figure 4 illustrates the difference in

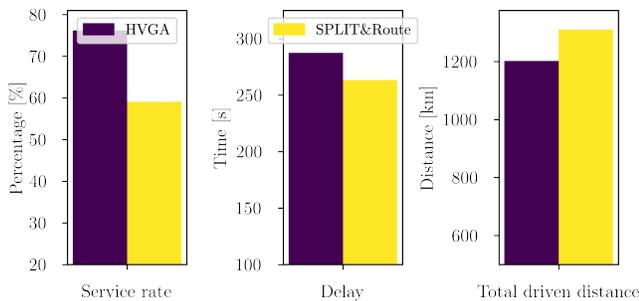


Fig. 3: Comparison of HVGA and Split&Route

the orders that get served in regard to their distance between the depot and drop-off location, based on the street network. A histogram of orders against travel distance is shown, where orders served by trucks are shown purple and orders served by drones in yellow (HVGA: Figure 4a, Split&Route: Figure 4b). HVGA primarily utilizes drone to serve short-distance orders, standing in direct contrast to Split&Route. This results in a greater amount of orders being served by drones (HVGA: 341, Split&Route: 138). The results suggest that the assumption to serve long-distance orders using drones does

not hold for flash delivery operations. Not doing so allows drones to serve a larger share of the requests.

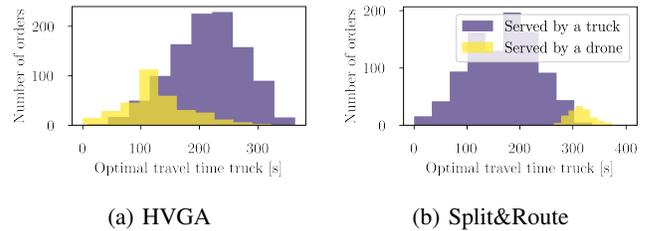


Fig. 4: Illustration highlighting the difference in served orders regarding their distance between depot and drop-off location between HVGA, left, and Split&Route, right.

B. Fleet Composition

As this work is the first (to the best of our knowledge) in investigating the effect of deploying drones within a flash delivery operation, we analyzed the fleet composition in more detail. We increase the number of pick-up locations in the service area to 20, to leverage that our method can handle multiple depots; three are considered per order ($x = 3$). Figure 5 shows the change in service rate, delay and total travelled distance for three fleets of total size 15, 20 and 25 and their composition in various ratios of trucks to drones. Independent of the total fleet size, we see an increase in service rate the more drones are used. This comes with an increase in total travelled distance. We see one exception for a fleet size of 25, when the service rate comes close to 100%. Then resources become available, which the approach can use to stronger optimize on traveled distance and delay. Generally, delay decreases the higher the percentage of drones used.

For flash delivery operations drones hold great potential. Short times between ordering and delivery and short distances between drop-off and pick-up location (strengthened by considering multiple depots) fit well with drones' benefits. High capacities are less important if comparing flash deliveries to traditional next-day operations, where times between leaving and returning to a depot can cover many hours.

VI. CONCLUSIONS

This paper presented an optimization-based approach to route a heterogeneous fleet of vehicles for an on-demand last-mile delivery operation. Orders are served within minutes after ordering, posing a special variant of a Same-Day Delivery problem. To analyze the results, a fleet of trucks and drones was studied and compared to a method inspired by [13]. HVGA was able to serve more orders while driving fewer kilometres in total. Further, the size and composition of various fleets have been studied. A larger amount of drones increases obtained service rates at the cost of increased travelled distance.

Future work involves a more detailed representation of used vehicle types, enabling a more accurate study about their use for flash deliveries. Further, the scope of experimental

