Research Report:

Planning Delivery-by-Drone Micro-Fulfilment Centres

By

J. S. Lamb, M.Eng PhD Candidate

ENF120, Department of Civil Engineering, University of Calgary, 622 Collegiate Pl NW, Calgary, AB T2N 4V8, Canada Email: <u>Jacob.lamb@ucalgary.ca</u>

Accessible on UCalgary PRISM



Abstract

Delivery drones are a disruptive technology that is spurring logistics system change, such as the adoption of urban micro-fulfilment centres (MFCs). In this paper, we develop and implement a two-stage continuum approximation (CA) model of this disruptive system in a geographic information system (GIS). The model includes common CA techniques at a local level to minimise cost, and then these local solutions are used in a second stage regional location-allocation multiple knapsack problem. We then compare the drone MFC system to a traditional delivery-by-truck system and investigate potential cost or emissions savings by adjusting time-window demand, logistical sprawl, electric truck alternatives, and MFC emissions. Furthermore, we conduct a sensitivity analysis to show that uncertainty in demand and effective storage density both significantly influence the number of MFCs selected and benchmark our model against commercial solvers. This methodology may also be further developed and applied to other new delivery vehicle modes.

Keywords: Drone-based Delivery; Geographic Information System; Micro-Fulfilment; Continuum Approximation; Inventory

Table of Contents

Abstract	2
Table of Contents	3
1. Introduction	4
1.1. State of the Industry	4
1.2. Study Objectives and Contributions	5
2. Literature Review	6
2.1. Drone Delivery Service	6
2.2. Micro-Fulfilment Centres (MFC)	9
2.3. Continuum Approximation Location-Allocation Modelling	10
2.4. Research Gaps	11
3. Methodology	12
3.1. Problem Formulation	12
3.2. Solution Method	13
3.3. Delivery-by-Drone System Objective Function	14
3.3.1. Last-mile Transportation Distance	14
3.3.2. Cycle and Safety Stocks	16
3.3.3. Cost Functions	16
3.4. Traditional Truck System Objective Function	17
3.5. Stage One: Community Optimisation	18
3.6. Stage Two: The Allocation and Location Problem	19
3.6.1. The Multiple Knapsack Allocation Problem	19
3.6.2. The Series of MFC Location Problems	21
4. Case Study	22
4.1. Summary of Parameters	22
4.2. Initial Classical Analysis	23
4.3. Two-stage Method Location-Allocation	24
4.3.1. Benchmarking	27
4.3.2. Operating Emissions	29
4.4. Mixed Drone and Truck System by Time Windows	30
4.5. Sensitivity Analysis	32
5. Conclusion	35
5.1. Future Work	36
References	39
Appendix A. Stage One Notation	44
Appendix B. Graphical Solution for Objective Function Optimisation	45
Appendix C. Allocation Algorithm	46
Appendix D. Parameter Values	47
D.1 Demand	47
D.2 Drone Last Mile Delivery Costs	48
D.3 Fixed Facility and Shelf Cost	48
D.4 Resupply Delivery Costs	49
D.5 Inventory Costs	50
D.6 Truck Last Mile Delivery Costs	50

1. Introduction

Urban logistics, supply chains, and freight transportation worldwide are being disrupted by new technology and changing customer expectations. One such technology is delivery drones. Researchers estimate that current drone technology is advanced enough that drone delivery systems could economically service about 30% of the world's urbanized population (Aurambout, Gkoumas, and Ciuffo 2019), and which investment industries estimate is valued in 2020 between one and two billion USD globally. In North America, at least four large retail firms that have e-commerce presence, Amazon, Walmart, London Drugs, and CVS, are pursuing drone delivery as a delivery option for their customers. In this introduction, we will first cover the state of the industry, then we will outline the objective of this report and our contributions.

1.1. State of the Industry

Over the last several years, many online retail companies have stated their interest in drone delivery. Amazon was an early proponent, which, according to their own promotional material online, has been designing and developing their own drone and delivery system to complement their already growing logistic service capabilities since 2013. Furthermore, Alphabet (the parent company of Google) is conducting trial drone delivery service, called Wing, which has been in operation in the USA, Australia, and Finland, since 2019. In 2020, a pharmacy chain (CVS) paired with a major logistics service provider (UPS) to deliver prescription medicine in Florida via drone delivery. Other major supermarket chains such as Walmart and London Drugs, have declared similar business pairings with new drone delivery service companies, such as Flytrex and Indro Robotics, respectively. This phenomenon is occurring worldwide as, for example, JD.com showcases its drone delivery in China since 2020, and Swiss Post is currently delivering mail from three locations in Switzerland.

With the rise of e-commerce, planners and managers can no longer assume a traditional supply chain from factory to warehouse to store, but rather a fast direct-to-home delivery option from a warehouse is often now expected in developed economies (Perboli et al. 2021). Moreover, today delivery is more frequently performed by the seller themselves, and not a third-party logistics supplier, as the seller seeks greater control over their customer experience and cost savings. Sellers, such as Amazon, Walmart, and JD, enter the logistics market with different motivations and committed to less logistics-only infrastructure than dedicated logistics service providers. These new entrants to the market are thus more interested in adopting emerging technologies, such as drones, than traditional third-party logistics companies. This new perspective changes costs and supply chain structures. Thus, new models are required to understand the changing business environment and to incorporate new transport technologies.

'Seller and service provider' (SSP) is the term given to companies that both sell goods and operate logistical services to make deliveries. Separate companies traditionally handled these two services: sellers utilized the services of a logistic service provider (LSP), such as a national postal service, to deliver their goods to their customers. This arrangement may continue for many years in rural areas where demand density cannot support multiple LSPs. The total cost perspective of SSPs compared to LSPs is different as they not only have to consider the logistic costs of vehicles and processing facilities but also of inventory costs, such as capital interest and opportunity costs. Therefore, the logistical systems of SSPs have different costs to consider and different business values than traditional LSPs.

Drones and other autonomous delivery vehicles are attractive to these Seller and Service Providers (SSPs) due to the public's perceived eco-friendliness of drones (battery-electric powered), and due to drone's ability to deliver in short time windows. According to the Canada Post E-commerce Survey in 2020, consumers are also increasingly influenced and aware of the environmental impacts of their shopping behaviour; for example, over half of Canadians are frustrated with excessive packaging, and over 40% of consumers report shopping with retailers who support an eco-friendly agenda. Drones may have further environmental sustainability benefits when they are paired with advances in micro-fulfilment centre (MFC) warehousing technology promise to address both the transportation emissions issue of urban freight and reduce packaging requirements. In a TED talk, the CEO of Attabotics, a company that designs and manufacturers MFC inventory storage systems, described how MFC technology can reduce packaging; drones and modern MFC conveyor systems carry individual packages, and these items require less protection than if they were packed tightly among many other parcels in a traditional delivery van. However, some items that require packaging due to their nature, such as food, may not benefit from this effect. In addition, due to the COVID-19 pandemic that struck the world in 2019, contactless drop-offs are desirable, which makes robotic drone delivery more attractive than traditional human delivery.

SSPs, such as Amazon, Walmart, and JD, enter the logistics market with different motivations than traditional Logistics Service Providers (LSPs) which make the above benefits of drones even more appealing. SSPs are committed to less logistics-only infrastructure than LSPs and so these new entrants to the market are more interested and able to adopt emerging technologies, such as drones. This new perspective changes costs and supply chain structures. Therefore, new models are required to understand this changing business environment and to incorporate drone delivery.

1.2. Study Objectives and Contributions

Our objective is to develop a model that can solve a multi-facility location-allocation with inventory problem that can be understood by local stakeholders and used as a planning and discussion tool by local area planners and elected decision makers. Furthermore, we wish the model to be easily extended to incorporate other autonomous delivery modes, and other research questions, by future work.

Assuming the proposed benefits of a delivery-by-drone system appeal to an SSP, we consider how they may implement such a system to serve an urban residential community that regularly demands a variety of drone-transportable goods. In our scenario, the SSP wishes to deliver both standard and expedited packages. Expedited packages could be clothing, footwear, grocery items, electronics, books, health products, office supplies, jewellery, or other items that the SSP decides. The SSP controls which products are offered for this expedited delivery; Chen et al. (2021) given heuristics for choosing which products may be profitable for expedited delivery. We do not consider emergency medical items, such as defibrillators or blood. We assume that the last-mile delivery is performed by the item seller who also operates the MFCs. Previously, delivery was the exclusive domain of light trucks and human delivery drivers, but, given advancements in robotics, we consider that it is possible, safe, and legal for these items to be delivered autonomously via drones.

No previous study has considered a cost-optimised location-allocation for an urban delivery-by-drone system or examined the externalities resulting from a profit-seeking SSP perspective. Furthermore, the estimated costs of the two-stage method used to construct the location-allocation solution are for the first

time compared to a GIS estimate of the solution's costs; the comparison validates the two-stage method for this case study and shows the method performs similarly to conventional continuum approximation (CA) techniques. Further development of CA techniques into integrated strategic models allows for future adaptations and extensions of the work, such as in truck-drone hybrids and sidewalk autonomous delivery robots. Adapting existing strategic decision making models for use with emerging technologies in this manner is a research gap as identified by Boysen et al. (2020).

In this report we will present a literature review, a description of the two-stage methodology, and a case study with sensitivity analysis. We will cover literature on drone delivery services, micro-fulfilment centres, and CA location-allocation modelling. We begin the methodology section with a description of the modelled system and its boundaries. We detail the first stage where transportation and inventory costs (among others) are modelled via CA. We then outline a description of the second stage where the allocation heuristic is implemented in a geographic information system (GIS), that is expanded upon in Appendix C. We will show how, by formulating the second stage as a multiple knapsack problem, we can both more easily facilitate design discussions typical of municipal or industry decision making, and facilitate the use of more efficient solver algorithms, than a traditional multi-facility location-allocation formulation. We will then present a case study and conduct a sensitivity analysis that identifies quantifiable barriers to implementing an urban delivery-by-drone system. In the case study, we also scrutinise the belief that drones would be a lower emissions alternative to current truck delivery. Finally, we will discuss the report's findings and present conclusions.

2. Literature Review

In the following thematic literature review we first outline previous studies and literature reviews that investigated drone delivery systems. This section builds the research questions that still exist in the field. We then describe works that developed similar CA location-allocation models to demonstrate the applicability of this methodology to the problem. We then give a definition and description of micro-fulfilment centres alongside recent literature reviews and industry surveys to present an understanding of this facility type. Finally, we describe in further depth works that introduced the specific two-stage model extended in this work and highlight the gaps that our work addresses.

2.1. Drone Delivery Service

In this report, the term 'drone' is used to refer to an airborne vehicle that can carry a payload and has a degree of autonomous self-control. Other terms for this same type of vehicle are 'uncrewed aerial vehicle' (UAV) and 'remotely piloted aircraft system' (RPAS). The term 'UAV traffic management' (UTM) is widely used to describe an air traffic management system that is capable of jointly coordinating multiple drones with existing multiple manned aircraft. It is assumed here that such a UTM system is available.

The study of drone delivery has gained increasing attention in literature as the technology has shown promise in industry. Moshref-Javadi and Winkenbach (2021) provide a broad literature review on drone delivery models. They find that most studies use either heuristic or mathematical programming methods, which aim to minimise system travel distance or financial system cost. They conclude that there is a severe lack of integrated problem studies, where topics such as inventory management or costs are considered. The authors also identify that few multi-facility problems have been adequately addressed for drone

delivery and they recommend that these multiple depot problems should be investigated further as, they argue, they are more likely to be what the logistical system that is implemented in the real world than single facility systems.

At an operational level, there are studies that investigate the optimisation of individual drone flight paths; Coutinho et al. (2018) provides a review of these studies. Moreover, hybrid truck and drone systems (known as "horsefly systems") at a tactical and operational level have been extensively studied, by authors such as Salama and Srinivas (2020). These, and other works, are covered by literature reviews of Rojas Viloria et al. (2020) and Macrina et al. (2020). Such horsefly systems but are out of the scope of our current paper. Our report covers an integrated multi-facility location-allocation problem for urban delivery-by-drone direct from MFCs.

At the tactical level, choosing drones based on battery capacity or weight capacity is also a trade-off, and an optimisation problem that influences drone fleet size requirements has been explored (Dorling et al. 2017). Dorling et al. (2017) conclude that 'hot-swapping' of batteries is more cost and time effective and requires a smaller fleet size than intermittent recharging. Shen et al. (2021) confirm the effectiveness of battery swapping although they highlight the additional capital investment required. This battery swap technique has been assumed in many later studies, such as Chauhan et al. (2019), Arenzana at el. (2020), and Schermer et al. (2020).

At a strategic level, drones require infrastructure (e.g., launch locations) to support their effective deployment, and so many related facility location studies exist in literature. Some studies consider delivery and response for disaster relief. Chowdhury et al. (2017) utilised a CA and GIS model to estimate optimal inventory levels and emergency response logistics locations for three counties in the southern United States of America (U.S.). MacKle et al. (2020) investigated a similar medical response drone system, one specifically designed for long-term facilities that serve cardiac arrest patients across Northern Ireland. This study also used GIS, but it complemented GIS with a genetic algorithm to evaluate the financial costs and estimated lifesaving benefits of the proposed drone system.

Other studies have investigated location and drone network designs for non-emergency medical scenarios. Kim et al. (2017) considered the drone delivery of regular prescription medication to patients' homes in a rural area using a bi-level integer programming model. Arenzana et al. (2020) also used a programming method, but they looked at the inter-hospital case of blood delivery in the city of London, United Kingdom, where congestion makes drone delivery both a faster and cheaper solution than the current road ambulance inter-hospital delivery network. Both studies also included a fleet size estimate for the drones, which added to the modelling complexity and limited the problem size to under ten locations.

Chen et al (2021) investigated cost-optimal drone delivery fleet size for only one MFC location using heuristic solving methods. These authors found that SSPs can profit from having a large fleet of drones to deliver high-value items with short wait times and that consolidating even just two packages per drone per trip can increase profits by over 50%. However, Moon et al. (2020) investigated a similar location-routing problem for multi-compartment last-mile delivery which could be applied to a multi-capacity drone system; as their case study accounted for an added cost to operate the larger and more complex multi-compartment vehicles, they found that this added capacity was often not justified and led to increased system cost.

Hong et al. (2017) showed how the drone battery swap station location problem could be analysed in GIS and how this method allowed for modelling of obstacles and non-Euclidean pathing. Asadi et al. (2021) conducted a detailed analysis of these replacement and recharging strategies at such battery stations to find that a policy of always swapping for fully charged batteries when replenishing the battery stations inventory was the best strategy, for a majority of cases.

Yet other studies have examined package delivery drone systems with various logistical structures. In a case study of Seattle, U.S., Shavarani et al. (2018) proposed a heuristic solution to determine an optimal arrangement of MFC locations and supporting drone recharge stations to completely replace Amazon's current van delivery system. Cokyasar at al. (2021) also developed and applied a heuristic programming method in a Chicago case study; the method located battery-swapping stations and allocated discrete demand points to the stations and allowed for a parallel truck and drone system, where some areas of the city were served by trucks and other areas were served by drones to achieve a low cost system – we also investigate such parallel systems in our case study.

In a statedly unique study, Baloch and Gzara (2020) investigated a drone delivery versus local grocery store delivery problem expressed as a multi-nominal logit market share model that they solve using mixed integer algorithmic solver that interchanges between solving a master problem (MFC location) and sub-problem (customer allocation). These authors find that, in their case study of New York City, that current regulations on the ratio of drone remote operators to drones controlled and the technological capability of drones to deliver to dense urban areas are barriers to drone delivery profitability. Baloch and Gzara also varied customer sensitivity to price, time, and an "inherent attractiveness" of drones to find several instances where drone delivery may be profitable in a range or urban and sub-urban scenarios. Although these authors consider a profit-maximising model, their cost function only includes a piece-wise fixed cost per facility and a fixed cost per delivery whereas our work considers these costs in continuous space plus inventory costs.

Most studies only investigate financial costs as their objective metric, but some investigate carbon emissions. Goodchild and Toy (2018) used GIS to compare the emission intensity of truck and drone delivery systems in Seattle; they adjusted the energy efficiency of the drones in a sensitivity analysis to understand what drone characteristics were required if lower emissions were desired. Stolaroff et al. (2018) included emissions of both vehicles and warehouses in a comparison of drone delivery systems that used MFCs and systems that used traditional van and electric van delivery. Furthermore, Figliozzi (2017) investigated lifecycle emissions of replacing a truck fleet with a drone system using only the CA method and determined that a drone system would result in fewer emissions. Figliozzi (2020) also compared truck and drone delivery system emissions under different operational requirements of a system, such as logistical sprawl distance and time windows. Figliozzi (2020) showed that delivery time windows did affect trucks, with shorter time windows decreasing truck efficiency. These studies found several scenarios in which drones were the most emissions efficient option, but none found that drones were always best, thus showing that the strategic decisions made before the operational ones significantly affected the environmental impact of delivery drones. However, these authors did not consider the total financial costs of a system or inventory

factors, so they could not state the cost changes incurred in the reduced emissions scenarios and thus could not state if profit-seeking companies would pursue these options without external incentive or regulation.

There are further criticisms of using drones for delivery that may prove a barrier for deployment, such as privacy, security, safety, environmental, social, and employment implications. Interested readers are encouraged to see Chung et al. (2020) for further discussion on these topics. For our case study we assume that these barriers are overcome.

2.2. Micro-Fulfilment Centres (MFC)

Urban logistics centres are a topic of growing interest in the urban freight literature and industry practice. Terminology is still not well defined nor agreed upon in the literature (Gunes and Goodchild 2021), but generally, and in this report, the term 'micro-fulfilment centre' (MFC) is used to refer to locations that have deliberate short-to-medium term inventory storage and are owned and operated by one company. 'Micro-hubs' may also be used, but this term may carry connotations of shared use. MFCs are like urban consolidation centres (UCCs), and the two types of centres share many transferrable insights, which have been explored in the literature. However, UCCs generally have no deliberate storage and are intended to reduce delivery vehicle volume into city centres, which can include municipally organised, owned, and even operated facilities, so work that examines inventory considerations is required to understand MFCs specifically (Urban Freight Lab 2020).

Although many UCC and MFC initiatives have been implemented in recent decades, especially in Europe, most have failed; a simulation study by Heeswijk et al. (2019) suggests initiatives fail because such facilities require not just one policy or motivating actor, but rather multiple concurrent policies must be implemented for a UCC facility to be sustainable beyond government subsidy. Another study by Lemardelé et al. (2021), which investigated drone delivery from UCCs, highlighted a conflict between the motivations of private companies to minimise cost and a municipal government's motivation to reduce freight related externalities (such as emissions), a conflict that further strains and weakens UCC initiatives. The impacts from urban logistic facilities affect many stakeholder groups so a model that can include the many risks and costs is needed to determine a sustainable business model.

A literature review by Björklund and Johansson (2018) shows that while many practitioners and municipalities believe there are great societal and environmental benefits that could be gained from implementing UCC initiatives, there is relatively little academic work on the subject. They conclude this lack of knowledge leads to the observed number of failed pilot projects worldwide. This critical knowledge gap is compounded by the dominance of the current state of practice that utilises low-cost industrial land and increases logistic sprawl; a lack of real-world industry practice limits the growth of new knowledge to only modelling studies and the lack of knowledge limits industry acceptance and effective deployment of UCCs.

Interest in UCCs and MFCs has grown recently due to new last mile technologies, such as drones. However, MFCs do not depend on any advanced technology; there are current MFC trials around the world that use cargo-bicycles (Rosenberg et al. 2021). Furthermore, although companies have expressed for many years an interest to take the technology into the city, current trials of drone delivery are exclusive to rural areas.

Many academic studies on urban delivery drones have so far not considered financial costs, including inventory costs, nor have they included the wider strategic decision-making implications of MFCs (deliberate multi-day storage). Rather, many studies focus on UCC models (simple cross-docking and temporary storage) and although some findings are transferrable to MFCs, some findings are not. For instance, Stolaroff et al. (2018) do not evaluate the financial cost of the studied urban drone delivery system, and so the cost of adopting such a system, and whether an SSP may be motivated to implement such a system, are unknown. Although Shavarani et al. (2018) consider the unit transportation cost and facility opening cost, they do not consider the variable sizing costs of each facility, inventory costs, nor upstream logistic costs that an SSP must consider. Lemardelé et al. (2021) evaluate launching drones from UCCs, but they do not consider stored inventory nor variable sizing of facility costs. Figliozzi (2020) evaluates only emissions; furthermore, he does not consider the impacts of multiple facilities in an urban area nor inventory at the facilities. There is a research gap in the literature, and therefore, there is a lack of understanding of the holistic costs of an urban drone delivery system using MFC locations as proposed by industry.

2.3. Continuum Approximation Location-Allocation Modelling

With respect to CA facility location problems, Newell (1973) developed the convex optimisation method and demonstrated how to minimise analytically the combined transportation and warehouse set-up costs. Erlenkotter (1989) accounted for economies of scale in the objective cost functions and introduced three new distance metrics that have various benefits. Building on this work, Rutten et al. (2001) added further cost terms, including inventory stock, and tailored the cost function for a specific case involving trucks in a Manhattan grid. Additional works over the decades are covered in literature reviews by Langevin et al. (1996) and Ansari et al. (2018). Ultimately, these analytical approaches provide the first stage of understanding a location-allocation problem and give estimates of the number of facilities, size of facilities, and the magnitude of the catchment areas.

The analytical outputs can be used in several ways to complete the allocation portion of the locationallocation problem. The information can be used directly to inform decision making and design, as originally recommended and intended by the classical authors. The results can be rounded to an integer amount and used to inform a separate programming solution, such as utilising inbuilt-to-GIS commercial location-allocation solvers. The information can be kept in continuous space, using the method from Ouyang and Daganzo (2006) which allows for the planning of optimally staged openings of several facilities in a region over a time horizon (Wang, Lim, and Ouyang 2017), or the information can be converted to a discrete equivalent as in the two-stage method.

The two-stage method has been applied to different logistical problems. Wirasinghe and Waters (1983) first introduced variable facility costs in their model that determined the optimal locations of waste transfer stations. The model then determined the optimal number of facilities using a CA method, information that was then that integrated in a discrete heuristic, location-allocation model. Waters et al. (1986) used the same method, a CA model combined with a discrete location-allocation model, to show how the optimal locations of bus garages changed depending on whether the authors included local air quality in the CA objective function or not. Wijeratne and Wirasinghe (1986) proposed a different second stage, as the authors used the optimal number of facilities per analytical zone to group communities together heuristically to form catchment areas for fire stations. Wirasinghe and Vandebona (1999) applied this same second stage

heuristic grouping approach to the location-allocation planning of a subway station network. Their analytical model accounted for user walking distance, station opening and variable station size costs. After the location-allocation was complete Wirasinghe and Vandebona then utilised a minimal spanning tree algorithm to design the optimal network link design. However, none of these works validated the estimated costs in the analytical stage nor the final solution costs after the allocation stage. Moreover, the authors did not present a methodology on how to collect these communities together systematically nor did they evaluate different heuristically created solutions.

A more recent use of the two-stage approach, which expanded the method to include inventory costs, was developed by Tsao et al. (2012). The authors examined the whole of the U.S. in a numerical example to demonstrate their approach. They divided the country into many small, equally sized areas that could each be assumed to have a uniform demand. Then, they applied the optimisation methods shown by Newell (1973) to determine optimal numbers of warehouses within these areas. In this instance, the sum of the facilities across these areas then formed the national solution, and they used the case study to investigate iterations of inventory considerations. Tsao et al. (2012) found an almost 12% reduction in total system cost using this two-stage divisive method compared with assigning the whole country a uniform average demand density. The authors, however, did not validate their final solution as they did not compare their estimated values for transportation costs in the optimisation process with the estimated transportation costs arising from their final solution. Furthermore, they did not further allocate demand points to the facilities; rather, they left this step for a future stage. Both points are addressed in the current work.

Chowdhury et al. (2017) used a similar two-stage method in a drone delivery system for a disaster relief study of the south coast of the U.S. The authors described a 'grid-couple-cover' approach that was implemented using a 'trial-and-error' method. This method was used to create a raster-like grid over the study area. There is not enough information on the trial-and-error method as described to replicate the study. Furthermore, the described approach may lose data precision; since every grid square must be of the same geographical size, the data must be gathered for known grid sizes or interpolated for these exact sized areas, which limits the accurate use of past data. By allowing for arbitrarily sized and irregularly shaped areas to form the solution, the method we develop allows for existing data to be more easily used to design systems. Furthermore, similar to Tsao et al. (2012), Chowdhury et al. (2017) did not compare their estimates used to construct their location allocation solution with true costs likely resulting from implementing that solution, and thus, they did not validate their method. Moreover, neither set of authors compare their two-stage CA methods with a discrete method, such as that produced by an in-built solver to a GIS software. In summary, our review of the literature agrees with the results of Ansari et al. (2018), that there are few studies which utilize a CA method to model location-allocation with inventory problems in recent years, which we attribute to the success of discrete and heurticial methods. However, that there has been little attention to the specific area of CA integrated models motivates us to investigate the approach as a research gap, as suggested by Ansari et al. (2018).

2.4. Research Gaps

To address the research gaps identified in the literature review, we present a two-stage CA and algorithm allocation method that has the following features:

1. It can solve a multi-facility location-allocation problem for urban delivery-by-drone;

- 2. It is an integrated model that accounts for inventory costs present in MFCs;
- 3. It can be implemented in common commercial GIS software;
- 4. It is internally validated in the case study to a standard as previously accepted CA works;
- 5. It can be extended to account for parallel or hybrid delivery systems where trucks, drones, sidewalk robots, cargo-bikes, and other modes are deployed from one location.

We then apply this model to a case study to validate the method and explore delivery-by-drone MFCs as an urban logistics solution. We estimate emissions, costs, and inventory factors as previous studies have done and, by considering the strategic problem, are able to make novel insights into the issue.

3. Methodology

In this section, we formulate the problem, including notation and some assumptions, and introduce the CA modelling method used to solve it. We then determine and compare the analytical expressions that are used to evaluate the two modes, drone and truck, including transportation and inventory costs. We present the objective cost function for the drone system and how it is optimised with respect to the number of MFCs is detailed, and finally we show the algorithm for allocating communities in the location-allocation second stage.

3.1. Problem Formulation

In this section, we present the system boundaries and define the system variables. Please see Appendix A for these variables and notation in a convenient table format.

Consider an SSP that desires to implement a multi-commodity delivery-by-drone system in a city. Consider also that the local and national regulations allow these flight operations, and that cost minimisation while serving the entire city is the aim of the SSP; what logistical and engineering factors must they include when they determine their infrastructure network?

The SSP aims to offer their service to every *community* (*subscript C*) within the municipal *region* (*subscript R*) by using *MFCs* (*subscript U*). The SSP wishes to know the best *number of MFCs* (*N*), and the locationallocation arrangement of the MFCs to minimise their *total annual operating costs* (C_{Drone}). Each *community area* (A_C) within the *regional area* (A_R) must be allocated to an MFC and thus be within one *MFC catchment area* (A_U).

The city's adult population demands a typical number of packages per unit area per unit time, which is the *mean regional demand density* (μ_R). The SSP meets this demand with *regional cycle stock* ($W_{\alpha,R}$) held across all MFCs in the region [subscript ordering is where variable subscripts are first, followed by applicable area subscripting second, e.g., $W_{\alpha,R}$ is cycle stock (α) in the region (R)]. Although the SSP has good demand prediction models, future observed demand is uncertain, and so the SSP also holds a *regional safety stock* ($W_{\beta,R}$) related to the *standard deviation of regional demand density* (σ_R). A portion of both stocks are stored locally at each MFC for use within each service area in an arborescent supply chain structure. Each MFC has capacity for the combined *MFC cycle stock* ($W_{\alpha,U}$) and *MFC safety stock* ($W_{\beta,U}$). The former stock is held to meet *catchment expected demand density* (σ_U). The sum of these two stocks prescribes the

total capacity per MFC (W_U), which may be different between MFCs. The sum of total capacities across all MFCs in the region forms the *regional capacity* (W_R).

Considering inventory costs, the average *wholesale value (u)* of the various commodities per delivery to be stored in the MFC must be paid by the SSP before they receive deliveries to the MFCs; they receive these deliveries at a fixed *resupply frequency (E)* from a regional warehouse outside of the city boundaries. The distance from this regional warehouse to a community is the *community logistical sprawl (d_c)*. While in storage, packages can be stored at an *effective storage density (m)*, which includes the need for walkways, sorting areas, item sizes, and the mix of standard and expedited packages; empirically, effective storage density is the observed number of packages sold in the resupply time frame divided by the total area of the storage facility. After being stored for a time, when a package is demanded, it is loaded onto a drone and flown to the customer at an average speed and following a path that can be modelled according to a *configuration factor (\varphi)*. The drone then returns, having not travelled further than its *maximum flight range (MFR)*. Only one customer is serviced per dispatch of a drone meaning that changing time-windows does not affect the total last-mile distance travelled given that we assume an adequate drone fleet size to always serve demand.

These outlined operations come at a cost. The cost of each MFC is separated into a *fixed annual cost per facility* (C_f) and a linear *annual cost per square meter of storage space per facility* (C_s); this cost must be either leased or its purchase financed at a yearly rate. To fill the MFC, the SSP must *pay for each delivery* (C_d) to each MFC, either paid to a third-party supplier or managed themselves. Typically, while packages are in storage, their wholesale cost has been paid by the SSP, but the packages have not yet been bought by a customer. This unmet cost is covered by the SSP, which, incurs an *inventory holding cost rate* (C_h). Other financing methods, such as a seller providing a platform on which others sell goods, and they only facilitate the transaction, (known as "drop shipping") can avoid this cost, but we do not consider these options. Once bought, outbound delivery begins. The drone operation has costs of electricity, equipment, and remote pilot wages that can be expressed as a *last mile cost per kilometre* (C_l).

For the comparison to the traditional truck delivery system, we consider the *truck travel cost* (C_t), which includes the cost of fuel, capital expenditure, driver wages, maintenance, and all other operating expenses expressed as a cost per kilometre travelled. Unlike drones, trucks have a *truck package capacity* (C), and the number of deliveries they can make per time-window is linearly affected by the *number of time-windows per day* (T) set by the SSP.

All above operations constitute the total operating cost of the urban consolidation centre delivery-by-drone system that the SSP wishes to consider. This total cost forms the objective function to be minimised. It is assumed revenue and demand are independent of MFC operations and structure, and so minimising cost is equal to maximising profit. The costs of this operation will be compared to the costs of a traditional delivery-by-truck system, formulated following CA methods outlined by Daganzo (2005).

3.2. Solution Method

A CA cost optimisation model is solved for each community in the region, and the sum of the communities' costs and resulting output parameters form the expected design parameters for a second stage allocation

where the communities are either collected or broken up into service areas for the final design. This method leverages the benefits of CA models while using a framework to introduce an explicit spatial term that can implement the method and provide solutions in a discrete allocation environment, as recommended by Ansari et al. (2018).

This two-stage method is like that used in logistic systems by Chowdhury et al. (2017) and Tsao et al. (2012). While they developed clusters of regularly sized discrete parcel units large enough to contain greater than one facility according to the first stage of the CA optimisation, our method preserves the lowest level of demand data availability, and the second stage is a clustering informed by the values recommended by the first stage of optimisation.

The two-stage method is useful because it can estimate optimal solutions for areas of high spatially varying demand, areas not suitable for conventional CA methods, while also accounting for many costs not easily optimised holistically in discrete methods simultaneously or in commercial GIS software. The two-stage method has not yet been applied to an urban goods delivery problem.

3.3. Delivery-by-Drone System Objective Function

We describe the two most significant terms of the system objective cost function, last mile distance and stock, before we present the resulting objective function and discuss the remaining terms.

3.3.1. Last-mile Transportation Distance

A fundamental aspect of CA modelling is estimating the last mile transportation distances using analytical expressions, which neglect the complexities of underlying routing and touring specifics. Accordingly, assumptions pertaining to delivery operations must be made.

First, we assume that drones operate at a constant rate with little non-operational time besides that normally expected, such as mandated flight restrictions during night-time hours.; we consider a ten-hour operational delivery day. The near constant use of the drones may be achieved by jointly scheduling standard and expedited packages dynamically using an operational heuristic or program. Furthermore, we assume that delivery time-windows for the items can be effectively managed by this scheduling so that time-windows do not impact the delivery cost and that there is a large enough drone fleet to meet the demand. Given these assumptions of constant drone use, the transportation cost term can be determined from drone speed, purchase cost, battery recharging cost, and remote pilot wage cost as the assumption separates these costs from delivery time-windows, and the unknown dwell time that occurs due to business practices and day-to-day demand. Thus, the last mile transportation cost is assumed to be directly and linearly related to the expected daily transportation distance flown by the drones.

The daily distance is calculated using a geometric methodology. In this report, we approximate the service area as hexagonal and approximate routing across a Euclidean plane following the A60 metric proposed by Erlenkotter (1989). This seminal work has since been applied by many authors in the area of logistics, such as Bouchery et al. (2020). Figure 1 shows the approximation; dotted lines show maximal diameters, and the dashed lines show example travel paths according to the metric. Travel is permitted along any of the six axes of the maximal diameters, in a similar manner to a Manhattan grid. This metric results in the edges of the service area being a uniform distance via the metric from the centre (Erlenkotter 1989). This assumption

is a reasonable approximation of a drone flight path given relaxed regulations, yet it still accounts for detours to avoid no-fly zones. Moreover, uniform hexagonal service areas may tesselate together to cover a community completely, and so they more accurately reflect the partition of a region into catchment areas.



Figure 1: A60 Distance Metric (Erlenkotter 1989)

$$L_{Drone,U} = 2. \varphi. \mu_C. \sqrt{A_U^3}; \varphi = 0.414$$
 (1)

The transportation distance, shown in (1), is the expected annual drone delivery distance per MFC (note, subscript "U"). We determined this transportation distance as the expected number of deliveries per unit area (μ_c), multiplied by the service catchment area of an MFC, multiplied by two for a return trip, multiplied by the average distance per delivery (φ) for the given distance metric and assumed catchment shape, multiplied by the square root of the service area (Newell 1973).

Erlenkotter (1989) shows that it is a feature of the A60 metric that any point on the perimeter of a hexagon is an equal distance from the centre when using the metric. Consequently, we estimate the maximum distance a drone would be required to fly in one delivery as twice the circumradius of the hexagon is. We can set this estimate equal to the maximum flight range of the drones to find an upper bound to the size of an MFC catchment area, which gives (2).

$$MFR \ge 2. \sqrt{\frac{2.A_U}{3.\sqrt{3}}} \cong 1.24.\sqrt{A_U}$$
 (2)

However, as we do not yet know the allocation of the MFCs, or even the number of MFCs, we cannot determine (1) or (2) from typically available community level data. So, to determine the transportation distance for a community we multiply the distance per MFC by the number of MFCs we expect in the community and expand the MFC catchment area to be expressed as a function of the known community area. The MFC catchment is the community area (A_C) divided by the number of facilities in that community (N_C). These steps result in the last-mile transportation distance of the drone system per community estimated as follows:

$$L_{Drone,C} = L_{Drone,U} \cdot N_{C}$$

$$= 2. \varphi \cdot \mu_{C} \cdot \sqrt{A_{U}}^{3} \cdot N_{C}; A_{U} = \frac{A_{C}}{N_{C}}$$

$$= 2. \varphi \cdot \mu_{C} \cdot \sqrt{\left(\frac{A_{C}}{N_{C}}\right)^{3}} \cdot N_{C}$$

$$L_{Drone,C} = \frac{2. \varphi \cdot \mu_{C} \cdot \sqrt{A_{C}}^{3}}{\sqrt{N_{C}}}; \varphi = 0.414$$
(3)

3.3.2. Cycle and Safety Stocks

The first and second term of (4) show how cycle and safety stocks are estimated in the model, respectively. Cycle stock is the expected demand density of the community multiplied by the area of the community. Cycle stock includes both standard packages and expedited packages. Standard and expedited packages could be separated into two cycle stocks if data is available in future studies. This expected demand density can be simple and related to population density, as in this report's case study, or it could be more complex and include various socio-economic or customer specific data if available and sufficiently significant in a demand prediction model (Unnikrishnan and Figliozzi 2020; Wang and Zhou 2015). The variation in this demand, captured by the standard deviation in demand density, can also be estimated by such demand models. Inventory planners can then utilise this estimate of variance to hold safety stock according to the relative costs of stock-out events (having too little inventory to service a spike in demand) and the costs of holding inventory (Axsäter 2006). This study's example utilises national sales data to estimate this variation (Aston et al. 2020). The relative cost and the company specific attitudes to risk are accounted for by the inventory stock-out factor (β).

Inventory Stock per Community =
$$A_C \cdot \mu_C + \beta \cdot \sigma_C \cdot A_C \cdot \sqrt{N_C}$$
 (4)

Finally, community safety stock is related to the square root of the number of facilities, as described first by Schwarz (1981), applied recently by Pham et al. (2020), and further investigated by Oeser (2019). Equation (4) assumes an arborescent distribution system, meaning that no trans-shipments between MFCs are considered, i.e., each community is only served by one MFC. The rise in stock requirement is related to the implications of the central limit theorem; as more customers are collected into a catchment, the coefficient of variation becomes smaller; so, inversely, if a community is partitioned into more catchments, and the number of customers per catchment decreases, then the coefficient of variation will increase and each MFC will require relatively more safety stock than cycle stock to maintain the same stock-out factor. This constraint of an arborescent network that does not allow for trans-shipments may be relaxed in future works because, in practice, neighbouring MFCs may deliver across catchment areas in a stock-out event. The current model assumes a conservative approach, which equalises the stock-out probability between truck and drone systems for comparison.

3.3.3. Cost Functions

Equation (5) shows the objective cost function that is to be optimised for each community in the first stage of the model.

$$C_{Drone,C} = C_d. E. N_C + C_l. \frac{2. \varphi. \mu_C. \sqrt{A_C}^3}{\sqrt{N_C}} + C_f. N_C + C_s. \left(\frac{A_C. \mu_C}{E.m} + \frac{\beta. \sigma_C. A_C. \sqrt{N_C}}{E.m}\right) + C_h. \left(\frac{u. A_C. \mu_C}{2. E} + \frac{u. \beta. \sigma_C. A_C. \sqrt{N_C}}{E}\right)$$
(5)

The first term captures the cost of regular deliveries to the MFC; it is the number of resupply trips multiplied by the unit cost of these trips. The regional warehouse where these resupply trips originate from is the furthest "up" the supply chain these costs will go as it is of interest to compare the drone delivery system to an alternative, traditional, delivery-by-truck system, and the distributions from the regional warehouse are the first substantial difference to be noted. The second term captures the expected transportation distance of the drones. The configuration factor φ may be taken from Erlenkotter (1989) to represent a given distance metric and catchment shape.

The third term captures the sum of fixed costs across all MFCs in the community. The fourth and fifth terms capture sizing costs; the former determines the storage space required for the cycle stock, and the latter determines the storage space for the safety stock. The sixth and seventh terms capture the holding costs of inventory, and, like storage costs, the former determines costs resulting from cycle stock, and the latter determines costs resulting from safety stock.

Once the optimisation is conducted, the resulting number of MFCs can be input back into this equation to determine the operating cost of the system in the community.

3.4. Traditional Truck System Objective Function

To compare the proposed drone system to a traditional delivery system, we use an estimate of transportation distance covered by a truck system. Truck delivery distance is estimated as per Daganzo (2005) and adjusted to include the number of time windows per day (T) in which the SSP offers expedited packages to be delivered. We assume that the trucks deliver proportionally fewer packages in a shorter time window; for instance, if a ten-hour day is partitioned into two time windows, then half of the packages delivered in ten hours is delivered in each five-hour window. This proportional effect of the time windows assumption means that the number of time windows per day linearly affects the first term of (6), the number of depotto-service area haul trips required. Secondly, the time window also affects the second term of (6), the local touring distance. Delivery points in each time window are proportionally less dense than in a single, daily, window. There are proportionally more of these local tours per day although each individual tour is shorter.

$$L_{Truck} = \frac{2.d.A.\mu.T}{C} + k.A.\sqrt{\mu.T}$$
(6)

Other estimates of touring distance are available, such as Figliozzi (2008), but we do not consider these because the parameters that require sample data for regression were not available for the study area. Furthermore, local tours could be made more efficient if the split of standard and expedited deliveries, and the split's resulting effect on touring performance, is known, but this data is also not available for the case study.

$$C_{Truck,C} = C_t \cdot \left(\frac{2 \cdot d \cdot A_C \cdot \mu_C \cdot T}{C} + k \cdot A_C \cdot \sqrt{\mu_C \cdot T} \right) + C_f + C_s \cdot \left(\frac{A_C \cdot \mu_C}{E \cdot m} + \frac{\beta \cdot \sigma_C \cdot A_C}{E \cdot m} \right) + C_h \cdot \left(\frac{u \cdot A_C \cdot \mu_C}{2 \cdot E} + \frac{u \cdot \beta \cdot \sigma_C \cdot A_C}{E} \right)$$
(7)

Equation (7) shows the objective cost function used to compare a traditional delivery by truck and a drone delivery system. The first term is the transportation distance, determined as described earlier. The remaining three terms are the costs associated with operating and maintaining a larger regional warehouse than would be needed with the MFC drone system. These costs are mathematically similar to those in (5) but with one fixed warehouse location (conceptually adding to the size of the regional warehouse). Given this formulation, our model implies that the delivery-by-drone system always incurs more warehousing and inventory costs than the traditional truck system does.

If the cost of a delivery-by-drone system is less than the cost of an existing traditional direct-delivery truck system, then drones are a viable and preferable last mile delivery option for the SSP. Even if the cost is greater, the intangible benefits of an eco-friendly brand, providing novelty to customers, and providing expedited delivery, may still make the delivery-by-drone system preferable.

3.5. Stage One: Community Optimisation

To minimise the delivery-by-drone system cost, (5), we determine the value for N_C that sets the first derivative of the function to zero, as in (8). However, this function is not convex, and so numerical, graphical, or heuristic analysis must be used. Later, in the case study, we choose to discretize (8) by evaluating the function over many values of N_C (implemented in Excel) to determine an approximate solution to a reasonable degree of accuracy.

$$\frac{\partial C_{T,C}}{\partial N_c} = C_d \cdot E - C_l \cdot \frac{\varphi \cdot \mu_C \cdot \sqrt{A_C}^3}{\sqrt{N_C^3}} + C_f + \frac{C_s \cdot \beta \cdot \sigma_C \cdot A_C}{2 \cdot E \cdot m \cdot \sqrt{N_C}} + \frac{C_h \cdot u \cdot \beta \cdot \sigma_C \cdot A_C}{2 \cdot E \cdot \sqrt{N_C}} = 0$$
(8)

In Appendix B, we also show a graphical solution to (8), aided by the rearrangement of the terms as shown in (8a). This rearrangement is convenient as we isolate the two different powers of N_C so that the terms can be graphed and compared easily.

$$C_d.E + C_f + \frac{C_s.\beta.\sigma_c.A_c}{2.E.m.\sqrt{N_c}} + \frac{C_h.u.\beta.\sigma_c.A_c}{2.E.\sqrt{N_c}} = C_l.\frac{\varphi.\mu_c.\sqrt{A_c^3}}{\sqrt{N_c^3}}$$
(8a)

For a deterministic demand (σ_c equals zero), (8) gives a familiar equation from Newell (1973) as shown in (8b), in which C_d.E adds to the fixed cost per facility. Assuming deterministic demand gives a closed form upper-bound, which is closer to the solution when the resupply frequency is high, when demand is predictable, and if the cost of goods, cost of shelf space, and cost of holding inventory are low. A graphical

solution method that shows the upper bounding effect is in Appendix A. The upper bound is not used further, and the solution to (8a) is used in the second stage of the method.

$$C_d.E + C_f = C_l.\frac{\varphi.\mu_C.\sqrt{A_c^3}}{\sqrt{N_c^3}}$$
(8b)

$$N_{C,upper \ bound} = A_C \cdot \left(\frac{C_l \cdot \varphi \cdot \mu_C}{C_d \cdot E + C_f}\right)^{\frac{2}{3}}$$
(8c)

With the number of MFCs determined after solving (8a), other outputs of the model can be determined using (4). No optimisation is conducted on the truck system as all the considered parameters monotonically affect the total system cost.

3.6. Stage Two: The Allocation and Location Problem

Although stage one has determined estimates for the number of MFCs required and other parameters, we have not yet utilised any spatial knowledge of the region. In stage two, the optimal number of MFC per community is used as an input weight parameter for a spatially-adapted multiple knapsack optimisation problem (Church and Murray 2008) and series of single facility location problems. This stage is in place of the clustering step in Tsao et al. (2012) and Chowdhury et al. (2017) as our method preserves more spatially-varying demand data. We show how to determine the optimal collection of communities for each MFC catchment area to solve the regional location allocation problem and then how to determine the locations of the MFCs within these catchment areas.

We highlight that this stage could be solved in several different ways, most interestingly is that it could be solved near manually in a stakeholder-led discussion. Although we present an algorithmic process and mathematically exact objective, by separating the problem into two stages then this second stage can be easily presented as a puzzle for discussion; the objectives of compact catchments and one MFC per catchment (as suggested by the sum of N_C values) are simple enough to observe on a map and calculate respectively that many stakeholders could understand and attempt their own solutions. We suggest that this stakeholder method would complement the following algorithmic method in a true decision situation.

3.6.1. The Multiple Knapsack Allocation Problem

Although the multiple knapsack problem is a conventionally NP-hard problem, modern heuristics can solve instances of thousands of items (communities) and over a hundred knapsacks (MFCs) in a reasonable amount of time (Dell'Amico et al. 2019). Joint location-allocation models, however, typically solve smaller instances of less than one thousand demand points and less than ten facilities because they also must compute transportation routing costs for each considered solution combination (Daskin and Tucker 2018). In other word, there exist more efficient optimization algorithms for multiple knapsack problems than there exist optimization algorithms for multiple solution. Therefore, our motivation to

investigate the two-stage method which can approximate a multi-facility location-allocation problem to a multiple knapsack problem, especially when the problem is large.

The typically a-spatial knapsack problem requires a spatial constraint or objective to model a locationallocation problem. This spatial constraint can be formulated in several different ways to enforce varying degrees of strictness to meet the desired catchment characteristics (Church and Murray 2008). In this report, we use two objectives, catchment perimeter and closeness to MFC unity, in a bi-objective method. Furthermore, we enforce that every community must be allocated to a catchment area.

This problem and solution method can be further extended to a method that determines optimal sections of the city to serve given a constraint on MFC number, budget, or to maximise profit. Moreover, it is possible that profit be assigned to each community or to relax the constraint that all communities must be served, to investigate maximum profit scenarios rather than cost minimization. This allocation problem could also be formulated and solved in other ways if desired. For instance, it may be formulated as a modified covering problem and solved using other discrete optimisation techniques, among other approaches (Daskin 2013). We choose to employ an algorithm method, like Hong et al. (2017), that uses a greedy algorithm supplemented by local interchanges that preserve the spatial contiguity of the catchment areas. Table 1 shows the notation that we use in the multiple knapsack allocation problem.

Notation	Value	Description
P _{U,j}	Perimeter of Catchment j	Perimeter of the catchment area in
		kilometres
P _{C,i}	Perimeter of Community i	Perimeter of the community in
		kilometres
P _{C,i,h}	Perimeter of Community i that neighbours	Neighbouring edge. Perimeter of the
	Community h.	community in kilometres
	Equals 0 when communities are not	
	neighbours.	
g _{i,h}	Equals 1 when communities i and h are in	Binary variable denoting whether two
	the same catchment.	communities are allocated to the same
	Equals 0 otherwise.	MFC catchment.
$N_{C,i}^*$	Optimal Number of MFCs in Community	Solution to (8b) in Community i
	i.	
N_R^*	Optimal Number of MFCs in Region	In real number form
$\widetilde{N_R^*}$	Rounded Optimal Number of MFCs in	Rounded to the nearest integer. This is
	Region	"the number of MFCs".
$N_{U,j}^*$	Optimal number of MFCs for Catchment	In real number form
_	j.	
U _{i,j}	Equals 1 when Community i is allocated	Binary variable showing if Catchment j
	to Catchment j.	contains Community i.
	Equals 0 otherwise.	

Table 1: Multiple Knapsack Problem Notation

We determine the perimeter of a catchment area ($P_{U,j}$) by calculating the sum of the perimeters of the communities allocated to the catchment ($P_{C,i}$) and subtracting the internal neighbouring edges. An internal neighbouring edge is the amount of a community's perimeter that neighbours another community's ($P_{C,i,h}$) when both communities are allocated to the same catchment ($g_{i,h} = 1$).

$$P_{U,j} = \sum_{i=1}^{n} P_{C,i} - \sum_{i=1}^{n} g_{i,h} \cdot P_{C,i,h}$$
(9)

Equation (9) allows us to evaluate the total perimeter of each community for a given MFC allocation. We then aim to minimize this total perimeter as this objective makes the catchment areas as compact as possible and, thus, transportation as efficient as possible, as shown in (10).

$$Minimise \sum_{j=1}^{\tilde{N}_R^*} P_{U,j} \tag{10}$$

The second objective of the bi-objective model (11) is to partition the region into MFC catchment zones (subscript label j) such that the sum of the optimal MFCs in all allocated communities (11c) is as close to an even split as possible of the total non-integer number of MFCs required (11b). The number of MFCs for the region is the sum of the MFCs required for each community, first as a real number as in (11a), and second, rounded to the nearest integer as in (11b). This objective, and its related definitions, can be modified if a boundary MFC is desired to be presently underutilised to allow for future increased demand. The optimal number of MFCs in a catchment is the sum of the optimal number of MFCs in the communities that are allocated to the catchment. We do this by using a binary categorical variable ($U_{i,j}$) as in (11c).

$$Minimise \ \sum_{j=1}^{\widetilde{N}_R^*} \left(N_{U,j}^* - \frac{N_R^*}{\widetilde{N}_R^*} \right)^2 \tag{11}$$

$$N_R^* = \sum_{i=1}^n N_{C,i}^* \qquad (11a) \qquad \widetilde{N_R^*} = \left[\frac{1}{2} + \sum_{i=1}^n N_{C,i}^*\right] \qquad (11b) \qquad N_{U,j}^* = \sum \left(U_{i,j} \cdot N_{C,i}^*\right) \qquad (11c)$$

We implemented and solved this stage in a desktop GIS using a greedy algorithm method supplemented by local interchanges, like Hong et al. (2017), so that we preserve the spatial contiguity of the catchment areas. See the supplementary information for further information on the algorithm followed. We also add constraints so that all variables are restricted to non-negative real values. See Appendix C for a description of the algorithm used. The next step in the allocation-location process is to determine the optimal MFC locations within the catchment areas.

3.6.2. The Series of MFC Location Problems

With the MFC catchment areas allocated, the location problem becomes a series of independent location problems. The objective is to minimise the transportation distance to and from all demand points in the

region. In addition, the assumption of uniform demand within each community may even be relaxed to find a locally optimal solution for MFC location within a catchment.

The location may be optimised using any number of solution methods to find the optimal location, but ultimately the decision may be quite restricted by the availability of appropriate space within the catchment area. With the aid of heuristics, the optimal discrete choice may be manually identifiable (especially when given only one MFC catchment area and a finite choice of locations) or the choice may require subjective input of intangible factors of the area. Several discrete optimisation techniques are also available, well known, and available in commercial GIS packages to solve single facility location problems (Daskin 2013; Church and Murray 2008). We use the "Mean Center" tool in ArcMAP in the case study.

4. Case Study

Now that the method has been described, a future drone delivery system, as potentially deployable by an ecommerce retailer (e-retailer) either in close partnership with a drone company or their own developed delivery-by-drone system, is investigated in the city of Calgary. We assume that an Unmanned Aerial System Traffic Management system has been developed and implemented so that the intended operations are legal and safe. The resulting delivery-by-drone system is then compared to a traditional delivery-bytruck system in terms of both cost and expected operational emissions. Finally, a sensitivity analysis is conducted on some estimated parameters, and their significance is discussed. The following results are displayed using Esri ArcMAP version 10.7.1 using 2016 geographical and census data made available by Statistics Canada.

4.1. Summary of Parameters

Table 2 shows a summary of the input variables for the numerical example. See Appendix D for a full explanation.

Population density, community area, and logistical sprawl distance vary per community. Additionally, demand uncertainty at the community level is considered unknown. This level of detail in the data may be available to the e-retailer or other sellers as is assumed in Tsao et al. (2012), but as this data is not available in our study area, we use an approximation, (12). This relationship between regional and community level demands, demand variations, and areas is required to maintain the assumption of arborescent network design and maintenance of the stock-out factor (Schwarz 1981). This formulation also preserves the expected linear scaling of the optimal MFC number in a community area. For example, if a community area is twice as large as another, and all other parameters are the same for both areas, the larger community area will have twice the number of MFCs as the smaller one.

$$\sigma_C = \sigma_R \cdot \frac{\mu_C \cdot \sqrt{A_R}}{\mu_R \cdot \sqrt{A_C}} \tag{12}$$

Parameter	Notation	Value	Units	Source
Mean Demand Density	μ	4.95 *	deliveries per square	(Unnikrishnan and
	$(\mu_R = \mu_C)$	pop.	kilometre per year	Figliozzi 2020)
		density		
Variability in Demand	σ	0.32 *	deliveries per square	(Aston et al. 2020)
		pop.	kilometre per year	
		density		
Last Mile Delivery	C_l	\$0.325	CAD per kilometre	(Figliozzi 2018)
Cost				
Fixed Facility Cost	C_{f}	\$90,000	CAD per MFC	Estimated from local data,
				see Appendix D.
Shelf Cost	C_s	\$180	CAD per square meter	Estimated from local data,
				see Appendix D.
Resupply Delivery	C_d	\$1,100	CAD per resupply per	Estimated from local data,
Cost			MFC	see Appendix D.
Inventory Holding	$C_h.u$	\$7	CAD per safety stock	(Rutten, van Laarhoven,
Cost			unit per year	and Vos 2001)
Truck Travel Cost	Ct	\$1.47	CAD per kilometre	American Transportation
				Research Institute 2019
				Trucking Report
Configuration Factor	φ	0.414	unitless	(Erlenkotter 1989)
Routing Parameter	k	0.82	unitless	(Daganzo 2005)
Stock-out Factor	β	2	unitless	(Axsäter 2006)
Truck Package	С	200	daily deliveries per truck	(Figliozzi 2008)
Capacity			tour	
Resupply frequency	Ε	72	resupplies per year	Estimated – sensitivity
				analysis conducted
Effective Storage	т	20	deliveries per square	Estimated – sensitivity
Density			meter	analysis conducted, see
				Appendix D.
Maximum Flight Rage	MFR	24	kilometres	Amazon self-reported
				technical specification.

Table 2: Case Study Default Parameters

4.2. Initial Classical Analysis

For a baseline understanding of the study area, a classical analysis is first conducted as if the whole region is one community that has slowly varying parameters.

Region Area	Adult	Regional Demand	N*	C _{Drone,R}	C _{Truck,R}
$(A_R; km^2)$	Population	Density (μ_R ;			
		packages/km ² /year)			
481	995,010	10,242	18.3	\$11,351,877	\$3,872,742

This preliminary analysis is also used to set the default resupply frequency. As shown in (5), this rate will have a minimiser for the objective cost. Table 3 shows a cost comparison of the truck and drone systems; as shown, the traditional system is almost three times cheaper. On a per item basis, costs are \$2.30 and \$0.79 for the drone and truck systems, respectively.

Consider a few different scenarios from the baseline. First, logistical sprawl in Calgary is relatively short at an average of only 23 km. The regional warehouse for this study is in Balzac, about 5 km north of the Calgary city limits, because it is the location of similar existing e-retailer facilities. If we consider a regional warehouse in the same Balzac location was to support a delivery-by-drone system in the city of Red Deer (a smaller city about 100 km north of Calgary) then the sprawl distance would increase, and the cost estimate for the truck system overtakes the drone equivalent: \$2.36 for the truck versus \$2.30 for the drone. The expected cost of the drone delivery does not change in the Red Deer scenario as we assume a conservatively high value for the cost to resupply the MFCs which is independent of the logistical sprawl (see Appendix section D.4).

Next, we consider changes to the time windows parameter for the city of Calgary. Imposing time-windows to the deliveries only affects the truck system as we assume the drone deliveries are fast enough, and the standard and expedited deliveries are planned well enough, to always meet the given time-windows. We experimented with the model using different time-windows between one and ten and find that if the truck system offers five two-hour time windows in a day, then the cost per item estimate for the truck increases to \$2.42, which is above the estimate for the drone system (\$2.30). This preliminary analysis agrees with intuition and previous research that these two situations, satellite cities that have large logistical sprawl from the regional hub and in the case of rapid delivery demands, are where the drone plus MFC system will be cost competitive with traditional truck delivery. Furthermore, in section 4.4, we show how the two-stage method can assist in further understanding the time window case.

4.3. Two-stage Method Location-Allocation

There are 202 residential communities in the city of Calgary as of the 2016 census. Table 4 shows a sample of the results from applying the community optimization stage of the two-stage method to the relevant community level data to determine the optimal number of MFCs; the table also includes a comparison of the optimal costs for the drone and truck systems. The first column lists the community number, from 1 to the total number, 201. We excluded one community from the case study due to it being non-contiguous with the other 201.

The second column shows the community area, which we determined as the geodesic area of the community using ArcMap Desktop 10.7.1. We also used ArcMap Desktop 10.7.1 to determine the perimeters of the communities and their neighbours using the "Polygon Neighbours" tool. We obtained population data for

each community from the 2016 census (Statistics Canada, 2016a). Demand density is the population density, calculated from population number and area, multiplied by the mean demand parameter, which is packages per year per person. We calculated the optimal number of MFCs by solving (8) for each community by evaluating the function over many values of N_C (implemented in Excel, version 2110, Build 16.0.14527.20234, 64-bit) to determine an approximate solution to a resolution of +/-0.005 MFC units per community. This numerical method confirmed that the optimizing value of number of MFCs was a cost minimizer. Note that the second stage of the method, where communities are grouped into catchment areas, is most easily conducted when the optimal number of MFCs is lower than one, indeed when it is lower than one-half on average, for each community. The total cost for the drone and truck systems is then calculated using (5) and (7), respectively.

Community	Community	Adult	Demand	Optimal	Total Drone	Total Truck
(i)	Area (A _{C,i} ;	Population	Density	Number of	System Cost	System Cost
	km ²)		(μ _C ;	MFC per	per	per
			packages/	Community	Community	Community
			km ² /year)	$(N_i *_C)$	(C _{Drone,C})	(C _{Truck,C})
1	6.34	19,775	15,446	0.310	\$201,848	\$53,374
2	2.95	19,945	33,514	0.229	\$167,227	\$61,778
3	4.65	16,405	17,475	0.245	\$162,059	\$74,717
•••						
n = 201	0.37	255	3,392	0.007	\$4,014	\$1,265
Region (R)	481	995,010		17.47	\$11,013,855	\$3,689,137

Table 4: Example Community Optimal MFC Number Results

The final row, Region, is the sum of the community values and represents the regional value result of the two-stage method. For instance, the optimal number of MFCs for the region is the sum of the optimal number in the communities (17.47) as per (11a). Rounded to an integer this is 17 as per (11b), one lower than recommended by the classical method, which is 18 (rounded down from 18.3). Furthermore, total system cost for the region is calculated similarly, by summation of the community cost estimates. The cost of the baseline drone system is predicted to be \$11,013,855 CAD and the truck system to be \$3,689,137 CAD, which is within 3% and 5% of the classical method, respectively. The total cost is later investigated in a sensitivity analysis by adjusting the input parameters.

We then solved the second stage of the two-stage method, the multiple knapsack problem by implementing the allocation algorithm (as described in Appendix C) in ArcMap Desktop 10.7.1. using an Intel i7-9750H CPU @ 2.60 GHz, 64-bit, 16GB RAM computer. ArcMap was also used to evaluate the objective functions during the solution process. Figure 2 shows the resulting location-allocation solution, with MFC catchment areas outlined in thick lines and neighbouring catchments textured, in four different patterns, to visually distinguish them.

We determined the optimal locations of each MFC within the catchments (circles with a centre dot) using the "Mean Centre" tool in ArcMap, with the demand of each community as a weight. We then obtained the drone transportation distances in ArcMap by: first estimating an average travel distance from each

community to its allocated MFC using the "Point Distance" tool from the centre of each community to the MFC; multiplying by the annual package demand of the community; multiplying by two for a round trip. We also calculated the maximum distance any drone in the region would be required to fly by using the "Construct Points" and then the "Point Distance" tools in ArcMap to evaluate the maximum distance from the MFCs to their respective boundaries. We found that the maximum distance from an MFC to the perimeter of the allocated catchment found was 11.153km, 22.306km for a return journey, which is less than the maximum flight range (24km) of the drones and so the location-allocation solution is feasible.

The allocation also allowed us to determine the cycle and safety stock required at each MFC, and the related MFC size and cost. For truck travel distance, we selected the location of the regional warehouse (a hollow cross) as the current real-world location of an existing Amazon fulfilment centre. We then determined the logistical sprawl distance from this regional warehouse to each community along road network by using the ArcMap "Make OD Cost Matrix Layer" tool.

The location-allocation solution shown in Figure 2 has objective function values of 1,459 km and 0.015 from the functions (10) and (10a), respectively. The transportation distance estimate of 20,470,629 km compared to 20,880,871 km estimated by ArcMap is approximately 2% lower, which is within a typically accepted range when using a CA method. Similarly, the estimated safety stock of 1,969 packages is about 6% higher than the GIS implemented result of 1,856 packages.

Figure 2 shows the catchment areas of MFCs 5 and 6 that span over Fish Creek Provincial Park. As it is legal for drones to fly over park areas, we labelled communities apposing one another over the park as contiguous in the data preparation stage. Some of these apposing communities, however, are not connected directly by road infrastructure, so models of other modes of transportation may consider these areas as disconnected during the allocation stage. Such modelling decisions are to be made on a case-by-case basis depending on the local conditions and modellers judgement.



Figure 2: Two-stage Model Location-Allocation Result

4.3.1. Benchmarking

It is important to consider the performance of our two-stage method when compared with commercially available methods. Hence, we performed a location-allocation optimization in ArcMAP using the Spatial Analyst extension.

For the 510 potential MFC locations, we used the geometric centres of every industrial and commercial zone in the city, sourced from publicly available data online provided by the City of Calgary. The 201 demand points were represented by the geometric centres of every community, weighted by their respective demand determined as in the two-stage method. We set up the drone travel network by creating a layer using the "XY to Line" tool which directly connected every demand point to every facility point. The ArcMAP commercial solver cannot account for financial cost, so we set the objective to minimize the sum of the weighted distances from each selected MFC to the allocated communities. This gives a location-allocation solution we considered as near-optimal for last-mile transportation distance, and as a benchmark for cost comparison.

Figure 3 shows the resulting location-allocation solution of the commercial solver for the case of seventeen MFCs.



Figure 3: Commercial Solver Location-Allocation Result

Table 5 shows the last mile distance and the estimated annual cost of operating the delivery-by-drone system for the respective location-allocation solutions. As commercial solvers typically require the number of facilities to be an input parameter for the analysis, we also performed the commercial solver analysis for eighteen MFCs, as might have been done if the classical method were used to determine the input number of MFCs.

	Method	Last Mile Distance	Total Annual Cost
		(kilometers)	(\$CAD per year)
Classical Method,	Classical Estimate	21,081,391 (9%)	\$11,352,526 (5%.)
N = 18	Implemented Location-	21,431,849 (11%)	\$11,467,057 (6%)
	Allocation Costs		
	Commercial Solver	19,338,333 (ref.)	\$10,787,306 (ref.)
Two-Stage Method,	Two-stage Estimate	20,470,629 (3%)	\$11,014,678 (2%)
N = 17	Implemented Location-	20,880,871 (5%)	\$11,111,678 (3%)
	Allocation Costs		
	Commercial Solver	19,929,300 (ref.)	\$10,790,430 (ref.)

Table	5:	Benchm	arking	of	Methods
1 4010	\mathcal{I}	Denemin	arking	01	memous

From these results, we can see that the two-stage method better estimates the benchmarks from the commercial solver, the last-mile transportation distance and the total annual cost of the system, than the classical method does. Taking the commercial solver solutions as benchmarks, we can see that the classical analysis overestimates the last mile distance by 9% in the initial estimate, and overestimates by 11% when the solution is implemented into a location-allocation. Whereas the two-stage method does still overestimate the last-mile distance of the benchmark, it does so by only 3% and 5% for the a-spatial estimate, and the implemented estimate respectively. The result is similar for the total system annual costs.

Although the two-stage method does not produce a lower total cost solution than the commercial solver for our dataset, this is due to the specific cost parameters of this scenario as the last-mile transportation distance represents over half of the total system cost in every scenario. We expect that the two-stage method would produce good solutions across a wider variety of cost variables than the commercial solver, given that the solver minimizes only transportation cost whereas the two-stage method accounts for inventory costs. Indeed, we expect there may be scenarios, with high holding costs, where the two-stage method produces a lower-cost solution than a commercial solver. We expect future work to examine the effectiveness of the two-stage method in a wider variety of scenarios.

4.3.2. Operating Emissions

Some companies are concerned about the environmental impact of delivery-by-drone systems, and therefore, it is important to evaluate the systems for carbon emissions, shown in Table 6. The vehicle energy per kilometre parameter in Table 6 is from Figliozzi (2017, 2020), and this parameter is also within the range of parameters as studied by Goodchild and Toy (2018). The carbon dioxide emissions per energy parameter uses emissions per fuel estimates published by the Government of Canada in 2017 adjusted for the portion of fuel currently used in the Alberta energy grid, as published by the Alberta Electric Systems Operator in the 2019 report. We assume that diesel and electric trucks cost the same and conduct routes in the same manner (same routes, capacity, transportation distance) as we assume that electric trucks will only be adopted when they perform similar in cost and routing to the incumbent diesel vehicles; this assumption could be altered and investigated by further studies.

Parameter	Distance	Energy per	Emissions Per	Emissions per km	Emissions
		km	Energy		
Units	km	Wh/km	kg/Wh	kg/km	metric tons
Drone	20,470,62	21.6	0.000201	0.004	88.8
	9				
Diesel Trucks	1,831,658	1016	0.000270	0.274	502.5
Electric Truck	1,831,658	205	0.000201	0.041	75.4

The results support the common assumption that drones are less carbon intensive compared to diesel trucks, even in an urban setting. However, the significantly higher total fleet travel distance of the drones caused by their one-package payload capacity ultimately makes the drone system more emissions intensive than a potential future of electrically powered trucks on the same energy network. Use of multi-capacity drones

may be able to reduce the total distance travelled (Chen, Hu, and Solak 2021). Our findings are consistent with Figliozzi (2020) and Goodchild and Toy (2018).

Absent from Table 6, however, are the energy requirements of the MFCs in the drone system, as discussed by Stolaroff et al. (2018). This value can be estimated for Alberta using information from a commercial building survey from 2013 published by Natural Resources Canada, and the square meterage required for the storage of goods is determined by the two-stage model. The total drone, diesel truck, and electric truck system emissions in tons of carbon dioxide, including warehouse emissions, is 392.1, 702.7, and 275.7, respectively. These results show that warehousing emissions are a significant factor to consider.

4.4. Mixed Drone and Truck System by Time Windows

Another advantage of the developed two-stage model is that a dual system (trucks and MFCs with drones working simultaneously but serving different parts of the same city) can be easily visualised and estimates of optimal systems can be created quickly. The classical analysis suggests that with five, two-hour time windows, the drone-only system is more cost effective than the truck equivalent. However, the classical analysis can only suggest a binary switch for the entire city. The two-stage method is a partitioning process and can help answer the following questions: can part of the city be served by drone and others by truck? and how may the build order of this eventual system be rolled out considering gradually improving time-window offerings to customers? Table 7,Figure 4, and Figure 5 show the two-stage method's answers to these questions. For different delivery window lengths (fractions of a ten-hour delivery day), the optimal cost of drone and truck delivery systems can be estimated and compared.

Number of	1	2	3	4	5	6	7	8	9
Time Windows									
Window	600	300	200	150	120	100	85	75	65
Length									
(minutes)									
Number of	0	0	1	4	9	12	15	15	17
Drone MFCs									
MFCs Added			6	4, 5, 7	8, 9,	3, 16,	2, 13,		1, 14
					10, 11,	17	15		
					12				

Table 7: MFCs by Time Windows

This analysis assumes that the transportation effectiveness of the drone system is unaffected by the time windows due to the effective mixing of standard and expedited packages (see section 3.3.1), but the truck delivery system is affected in terms of delivering fewer parcels per vehicle per tour. An MFC is added and worthwhile when the cost estimate of the catchment is lower for the drone system than the truck system. This analysis also only considers the earlier determined catchments, but new catchment areas, determined using the allocation algorithm, could be created when examining a specific time-window scenario.



Figure 4: Drone, Truck, and Mixed System Costs by Number of Time Windows

Figure 5 shows mixed mode solution for the 150-minute time window problem. The same catchment areas that were regionally optimal for a full delivery-by-drone system (from Figure 2) are shown in bold black lines again. These catchment areas are served by one MFC each, which are numbered and represented by white circles with centre dots. Cross-hatched catchments are served by drones, and the non-hatched catchments are served by trucks. Moreover, the colour-coding of the communities (outlined with thin grey lines) reflects the difference in estimated cost between the two modes per expected package demanded by that community, according to the legend in the figure. Figure 3 also shows that the e-retailer regional hub is located to the north of the city, above and between catchments 1 and 14. The logistical sprawl distance from the hub to the south of the city is the most significant factor in raising the cost of expedited deliveries as the trucks must make this haul multiple times per day, increasing with increasing numbers of time windows.



Figure 5: 100-Minute Time Windows Location-Allocation

4.5. Sensitivity Analysis

We considered three logistical parameters of interest to examine by sensitivity analysis: effective storage density, uncertainty in demand, and resupply frequency. These parameters show the three main relationships that parameters have with the decision variable in (8) and the system cost in (5). Figure 6 shows a linear relationship; Figure 8 is a reciprocal function, and Figure 7 is a convex function that is the union of the linear and reciprocal functions. The figures show the last mile cost (drone distance), the MFC costs (sum of fixed, resupply, and storage costs), the holding cost (related to inventory), the sum of these costs for total system cost and the related optimal number of MFCs, both rounded (\widetilde{N}_R^*) and as the determined CA number (N_R^*). These three figures are representative of how any single variable may affect the number of MFCs in the two-stage method and are expected when examining the system cost and derivative equations as in (5) and (8).



Figure 6: Regional Uncertainty in User Demand



Figure 7: Resupply Frequency



Figure 8: Effective Storage Density

Of these three parameters, demand uncertainty is the parameter that the seller has the least influence over as it is a function of customer behaviour. This factor is shown to near linearly relate to higher optimal drone MFC system cost as less knowledge of customer demand leads to larger safety stock inventory, which leads to increases in both storage space and inventory holding costs. Alternatively, the seller may accept more frequent stock outs and accept the associated indirect costs to branding and customer loyalty. To avoid either increased stock costs or indirect customer costs, our analysis andv Figure 6 suggests that decreasing the number of MFCs is appropriate (Geoffrion 1979). Although decreasing the number of MFCs will increase last-mile transportation costs, this action will also create larger catchment areas, which will aggregate more demand. This lower increase in uncertainty in demand helps contain the increase in total safety stock required in the region, thereby preventing super-linear increases in inventory holding costs. Our results show that the increase holding costs avoided by decreasing the number of MFCs is greater than the trade-off increase in last-mile transportation costs for decreasing the number of MFCs, therefore making decreasing MFCs the correct choice in response to higher demand uncertainty.

Figure 6 can also inform how much effort and resources should be put towards reducing demand uncertainty. For example, the seller could issue a customer survey with the expectation that this survey will reduce demand uncertainty (the standard deviation in predicted user demand) by 0.2 packages per person per year. Figure 6 suggests that this reduced uncertainty could save approximately \$500,000 CAD per year in system cost. This cost is saved because the reduced uncertainty, as provided by the survey, allows for the opening of two more MFCs and a more than commensurate reduction in transportation costs. Thus, such a survey may be a positive investment if the survey and the resulting non-recurring implementation work associated with it, could be conducted for less than the expected amount saved.

Figure 7 shows the resupply frequency, which also affects order quantity. Within realistic operational ranges, the system cost is relatively flat, which suggests that other tactical or operational constraints should be considered to determine the resupply frequency between deliveries once weekly (52) and once every three days (152), the system cost does not increase or decrease by more than 3%. However, over this same range, the number of MFCs to facilitate this system cost changes significantly, 19 to 15. Consequently. this analysis suggests that an effective strategy could be to establish 19 MFCs with have fewer deliveries (once

weekly) at the start of a multi-year plan, and then the seller should expect to increase the resupply frequency over time to accommodate demand growth. This strategy allows for longer-term growth whilst still achieving a near optimal estimated cost in the short-term.

Figure 8 shows that changes in effective storage density above 15 packages per square meter do not result in any changes in the number of MFCs. The system cost savings approach a limit of about 8% beyond this effective storage density. Moreover, effective storage densities above some value will become unreasonable given minimum commercial building lot size, demand, and resupply frequency. This 'cut-off' value, 15 packages per square meter in this case study, however, is affected by the per area shelf space cost, and it is clear by examining the cost function in (5) that a higher shelf space cost leads to a higher storage density cut-off value. Consequently, we define the cut-off value as when decreases in shelving cost become insignificant. This relationship is linear, there is also an effective shelf cost per item (C_s/m) cut-off value. Our case study and analysis suggest that delivery-by-drone systems with an inventory management strategy and rental market combination that can support an effective shelf cost above 0.09 CAD per item are insensitive to further decreases in rental cost or increases in effective storage density.

Effective storage density is achieved by: accurate prediction of demand at a time scale relative to resupply frequency, measured in days; small item sizes; efficient rack placement; improved warehousing technology; and good inventory management. If these densities cannot be achieved, then the effective storage density factor is significant because it is the only parameter that results in MFC numbers below ten and even as low as one or two MFCs, an unfeasible MFC arrangement for current drone technology. Thus, confidence in selling goods stocked in MFCs and achieving a high effective storage density are initial, but not continuing, barriers for companies using a delivery-by-drone system. Furthermore, an effective storage density above a certain range (above 15 in this case) is effectively wasted as the optimal system, as estimated by the first stage in the CA, is likely below the minimal lot size of commercial space to rent or buy and thus cannot be realised. The SSP may instead increase the variety of items offered for expedited delivery until this threshold of effective storage density is reached as this will likely garner more sales, and more profit, from their customers. Understanding the relationship shown by Figure 8 is important to not over or under offer the range of expedited delivery packages.

5. Conclusion

Urban goods delivery has been dominated by fossil fuel powered, human driven trucks and customerattracting physical retail stores of various sizes for the history of modern retail. With the advent of disruptive technologies, such as drones and e-commerce, and the adoption of disruptive business practices, such as seller and service providers and micro-fulfilment centres, the traditional structure of urban retail logistics will change.

Multiple recent literature reviews on urban logistics facilities and urban delivery-by-drone have shown the research gaps that this report addresses. Moshref-Javadi and Winkenbach (2021) support research looking at multi-facility location problems and integrated facility problems that consider inventory management and costs. Boysen et al. (2020) agree with this conclusion, adding that recent studies on emerging delivery technologies have focused largely on routing problems, and that strategic problems of location-allocation and fleet composition remain a valid task for future research. Ansari et al. (2018) conclude that CA models are uniquely positioned to be extended and used to solve these integrated and strategic problems. The

present work solves an integrated location-allocation problem that includes inventory by applying a CA approach that can be further modified to account for different delivery modes in the future.

This report has shown that, for the example seller and service provider studied, a micro-fulfilment centre delivery-by-drone system for typical packages is not yet justifiable on a pure cost effectiveness basis but may be in the future given advances in technology, regulations, and/or implementation of short time-window delivery. We have also shown that the mix of expedited and standard packages the seller will offer is an essential factor for success of a delivery-by-drone and MFC system, and conversely this mix of goods can be a critical barrier if managed incorrectly. This mix of goods affects both drone utilisation efficiency outside of MFCs and effective storage density within each MFC, a factor that determines the size of the MFCs needed. If these challenges are overcome, a cost-optimal drone system can be a lower emissions alternative to a traditional diesel truck system but not lower than a future electric van system.

Furthermore, decision makers in this field need adaptable tools, such as CA models, that will aid in the understanding of these new logistical structures. Within a city, however, the common first assumption of uniform demand density that many CA models make is often invalid due to spatially distributed socioeconomic factors. We have shown that CA methods can be used to understand this coming urban disruption while acknowledging non-uniform demand density space; in addition, CA methods can be used to build allocation maps in a methodological way with commercially available software using this two-stage method. This report has shown that this two-stage method estimates transportation distance and inventory to a similar accuracy as classical single uniform demand approximation methods and thus the insights of CA methods at a local level will approximately hold for cost-optimal regional level solutions. The two-stage methodology could also be expanded at the first stage to include multi-capacity vehicle routing, such as sidewalk robots or local delivery vans, or multi-modal depots.

5.1. Future Work

First, the transportation cost parameter considered in the analysis assumes that multiple drones may be controlled beyond a visual line of sight by a single operator. This technological and regulatory environment is still a few years away at the time of writing. Without this development, the MFC and drone system cannot be cost competitive with trucks in even the best-case scenarios for the drones. The model also assumes fair flying weather at all times. We compared typical drone characteristics and historical Calgary weather data and found this was a good assumption because there were fewer than three days a year on average in the past ten years of data likely to disrupt drone flight operations. Moreover, we assumed drones work nearly constantly over their daily routine, meaning that they were delivering standard parcels, in-between time-sensitive deliveries. This intermixing of expedited and standard deliveries requires effective operational algorithms which have yet to be developed and are of interest for future study.

Second, a barrier to entry, effective storage density, is identified. Offering a larger variety of items for expedited service reduces effective storage density and offering fewer items conversely will increase the density. However, low effective storage density increases the size of MFCs required and thus increases system costs. For a given SSP, the challenge is to maintain the effective storage density above a minimal value (15 in the example) by improving their predictions or changing the mix of expedited and standard packages offered. Investigating how to meet this challenge is a topic for future study. Furthermore, relaxing

the assumption of an arborescent network is likely to affect how effective storage density affects and is affected by the rest of the system, and is also of interest for future study.

Finally, if this system overcomes these challenges, there is potential for carbon emissions savings compared to the traditional diesel delivery truck system. However, these savings are vastly reduced, and potentially completely nullified, when considering the energy demands of the MFCs needed. In the study area of Alberta, Canada, where electricity generation still utilises coal power stations and building heating requirements throughout the year are substantial, the energy impacts of the MFCs is shown to be significant. In these instances, drones supported by an MFC system may be a less effective emissions alternative compared to electric delivery vans. Timely delivery is the main driver behind the adoption of a drone system, without short time-window constraints, and this study suggests that the consolidation efficiencies of electric vans will make them the lower emissions choice. Therefore, it would be interesting if future work could expand on some of these aspects, such as including emissions in the objective cost function for the cost optimization stage.

Further to the parameter and model environment areas of future work stated above, future modelling efforts can incorporate more complex CA optimisation in step one, such as the analysis of a multi-mode MFC. As the two-stage method separates determination of an optimal number of facilities, allocation of communities to catchment areas, and location of final communities into separate problems, each stage can be more complex than if they were considered simultaneously. Future work may also provide improved and automated allocation procedures for the multi-knapsack allocation problem leveraging the Python programming capabilities of ArcMap or comparable languages for other GIS software, such as R or QGIS.

Further modelling expansions of this study may also include a multi-modal autonomous delivery system, accounting for sidewalk delivery robots, multi-capacity drones, electric vans, traditional vehicles, horsefly systems, and joint systems where these modes operate in parallel. Modelling and optimisation of a system over time, as done by Wang et al. (2017), is another avenue for further work for which CA paired with GIS is uniquely suited. Finally, it is possible and of interest to account for emissions in the objective function of the first stage in the two-stage model directly. In this manner, emissions pricing policies may be modelled alongside existing logistical costs. Emissions alone could be considered, and an emissions optimised system and an economically optimised system could be compared to understand complementary features and opportunities for system improvement.

Acknowledgements:

The authors would like to thank the following parties for their support of this work. Firstly, to the Natural Sciences and Engineering Research Council (NSERC) of Canada for their financial support via the Discovery grant program and through the Collaborative Research and Training Experience (CREATE) Program. I also wish to thank my supervisors, Professor Chan Wirasinghe and Professor Nigel Waters, for their help and advice over these years. Thanks also go to Professor Lina Kattan, the director of the Integrated Infrastructure for Sustainable Cities (IISC) CREATE program for her mentorship and funding support. Further financial thanks go to the Government of Alberta for their support through the Alberta Graduate Excellence Scholarship (AEGS). Furthermore, the authors gratefully acknowledge the technical support and advice pertaining to the Geographic Information System components of the project from Renna Truong, of the Spatial And Numeric

Data Services (SANDS) at the University of Calgary. Finally, the authors also gratefully acknowledge and thank the four anonymous journal reviewers for their insightful and helpful comments.

Declaration of Conflicts of Interests:

None

References

- Ansari, Sina, Mehmet Başdere, Xiaopeng Li, Yanfeng Ouyang, and Karen Smilowitz. 2018.
 "Advancements in Continuous Approximation Models for Logistics and Transportation Systems: 1996–2016." *Transportation Research Part B: Methodological*. Pergamon. doi:10.1016/j.trb.2017.09.019.
- Aurambout, Jean Philippe, Konstantinos Gkoumas, and Biagio Ciuffo. 2019. "Last Mile Delivery by Drones: An Estimation of Viable Market Potential and Access to Citizens across European Cities." *European Transport Research Review* 11 (1). European Transport Research Review. doi:10.1186/s12544-019-0368-2.
- Arenzana, Alejandra Otero, Jose Javier, and Escribano Macias. 2020. "Design of Hospital Delivery Networks Using Unmanned Aerial Vehicles." *Transportation Research Record*. doi:10.1177/0361198120915891.
- Aston, Jason, Owen Vipond, Kyle Virgin, and Omar Youssouf. 2020. "Retail E-Commerce and COVID-19: How Online Shopping Opened Doors While Many Were Closing." *Statistics Canada*. https://www150.statcan.gc.ca/n1/pub/45-28-0001/2020001/article/00064-eng.htm.
- Axsäter, Sven. 2006. *Inventory Control Second Edition*. 2nd ed. Lund. https://link-springercom.ezproxy.lib.ucalgary.ca/content/pdf/10.1007%2F0-387-33331-2.pdf.
- Baloch, Gohram, and Fatma Gzara. 2020. "Strategic Network Design for Parcel Delivery with Drones under Competition." *Transportation Science* 54 (1): 204–28. doi:10.1287/trsc.2019.0928.
- Björklund, Maria, and Henrik Johansson. 2018. "Urban Consolidation Centre a Literature Review, Categorisation, and a Future Research Agenda." *International Journal of Physical Distribution and Logistics Management* 48 (8): 745–64. doi:10.1108/IJPDLM-01-2017-0050.
- Bouchery, Yann, Johan Woxenius, and Jan C. Fransoo. 2020. "Identifying the Market Areas of Port-Centric Logistics and Hinterland Intermodal Transportation." *European Journal of Operational Research* 285 (2). Elsevier B.V.: 599–611. doi:10.1016/j.ejor.2020.02.015.
- Boysen, Nils, Stefan Fedtke, and Stefan Schwerdfeger. 2020. "Last-Mile Delivery Concepts: A Survey from an Operational Research Perspective." *OR Spectrum* 43 (1). Springer Berlin Heidelberg: 1–58. doi:10.1007/s00291-020-00607-8.
- Chauhan, Darshan, Avinash Unnikrishnan, and Miguel Figliozzi. 2019. "Maximum Coverage Capacitated Facility Location Problem with Range Constrained Drones." *Transportation Research Part C: Emerging Technologies* 99 (May 2018). Elsevier: 1–18. doi:10.1016/j.trc.2018.12.001.
- Chen, Heng, Zhangchen Hu, and Senay Solak. 2021. "Improved Delivery Policies for Future Drone-Based Delivery Systems." *European Journal of Operational Research* 294 (3). Elsevier B.V.: 1181–1201. doi:10.1016/j.ejor.2021.02.039.
- Chowdhury, Sudipta, Adindu Emelogu, Mohammad Marufuzzaman, Sarah G. Nurre, and Linkan Bian. 2017. "Drones for Disaster Response and Relief Operations: A Continuous Approximation Model." *International Journal of Production Economics* 188 (October 2016). Elsevier B.V.: 167–84. doi:10.1016/j.ijpe.2017.03.024.
- Chung, Sung Hoon, Bhawesh Sah, and Jinkun Lee. 2020. "Optimization for Drone and Drone-Truck Combined Operations: A Review of the State of the Art and Future Directions." *Computers and Operations Research* 123. Elsevier Ltd: 105004. doi:10.1016/j.cor.2020.105004.
- Church, Richard L., and Alan T. Murray. 2008. Business Site Selection, Location Analysis and GIS. Business Site Selection, Location Analysis and GIS. Wiley. doi:10.1002/9780470432761.

- Cokyasar, Taner, Wenquan Dong, Mingzhou Jin, and İsmail Ömer Verbas. 2021. "Designing a Drone Delivery Network with Automated Battery Swapping Machines." *Computers and Operations Research* 129. doi:10.1016/j.cor.2020.105177.
- Coutinho, Walton Pereira, Maria Battarra, and Jörg Fliege. 2018. "The Unmanned Aerial Vehicle Routing and Trajectory Optimisation Problem, a Taxonomic Review." *Computers and Industrial Engineering* 120 (June 2017). Elsevier: 116–28. doi:10.1016/j.cie.2018.04.037.
- Daganzo, Carlos F. 2005. Logistics Systems Analysis. 4th ed. New York: Springer Berlin Heidelberg.
- Daskin, Mark S. 2013. Network and Discrete Location Models, Algorithms, and Applications. Ann Arbor.
- Daskin, Mark S., and Emily L. Tucker. 2018. "The Trade-off between the Median and Range of Assigned Demand in Facility Location Models." *International Journal of Production Research* 56 (1–2). Taylor & Francis: 97–119. doi:10.1080/00207543.2017.1401751.
- Dell'Amico, Mauro, Maxence Delorme, Manuel Iori, and Silvano Martello. 2019. "Mathematical Models and Decomposition Methods for the Multiple Knapsack Problem." *European Journal of Operational Research* 274 (3). Elsevier B.V.: 886–99. doi:10.1016/j.ejor.2018.10.043.
- Dorling, Kevin, Jordan Heinrichs, Geoffrey G. Messier, and Sebastian Magierowski. 2017. "Vehicle Routing Problems for Drone Delivery." *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 47 (1): 70–85. doi:10.1109/TSMC.2016.2582745.
- Erlenkotter, Donald. 1989. "The General Optimal Market Area Model." *Annals of Operations Research* 18 (1). Baltzer Science Publishers, Baarn/Kluwer Academic Publishers: 43–70. doi:10.1007/BF02097795.
- Figliozzi, Miguel A. 2008. "Planning Approximations to the Average Length of Vehicle Routing Problems with Varying Customer Demands and Routing Constraints." *Transportation Research Record* 765 (2089): 1–8. doi:10.3141/2089-01.
- Figliozzi, Miguel A. 2017. "Lifecycle Modeling and Assessment of Unmanned Aerial Vehicles (Drones) CO2e Emissions." *Transportation Research Part D: Transport and Environment* 57: 251–61. doi:10.1016/j.trd.2017.09.011.
- Figliozzi, Miguel A. 2018. "Modeling the Sustainability of Small Unmanned Aerial Vehicles Technologies."
- Figliozzi, Miguel A. 2020. "Carbon Emissions Reductions in Last Mile and Grocery Deliveries Utilizing Autonomous Vehicles." *Transportation Research Part D*. 85.
- Geoffrion, Arthur M. 1979. "Making Better Use of Optimization Capability in Distribution System Planning." *AIIE Transactions* 11 (2): 96–108. doi:10.1080/05695557908974448.
- Goodchild, Anne, and Jordan Toy. 2018. "Delivery by Drone: An Evaluation of Unmanned Aerial Vehicle Technology in Reducing CO2 Emissions in the Delivery Service Industry." *Transportation Research Part D: Transport and Environment* 61 (June). Pergamon: 58–67. doi:10.1016/j.trd.2017.02.017.
- Gunes, Seyma, and Anne Goodchild. 2021. "Characteristics of a Successful Microhub: Taking Public and Private Stakeholders' Perspective into Account." In *Volvo Research and Educational Foundations Conference*.
- Hong, Insu, Michael Kuby, and Alan Murray. 2017. "A Deviation Flow Refueling Location Model for Continuous Space: A Commercial Drone Delivery System for Urban Areas." Advances in Geographic Information Science, 125–32. doi:10.1007/978-3-319-22786-3_12.
- Kim, Seon Jin, Gino J. Lim, Jaeyoung Cho, and Murray J. Côté. 2017. "Drone-Aided Healthcare Services

for Patients with Chronic Diseases in Rural Areas." *Journal of Intelligent and Robotic Systems: Theory and Applications* 88 (1): 163–80. doi:10.1007/s10846-017-0548-z.

- Langevin, André, Pontien Mbaraga, and James F. Campbell. 1996. "Continuous Approximation Models in Freight Distribution: An Overview." *Transportation Research Part B: Methodological* 30 (3 PART B): 163–88. doi:10.1016/0191-2615(95)00035-6.
- Lemardelé, Clément, Miquel Estrada, Laia Pagès, and Mónika Bachofner. 2021. "Potentialities of Drones and Ground Autonomous Delivery Devices for Last-Mile Logistics." *Transportation Research Part E: Logistics and Transportation Review* 149 (March): 102325. doi:10.1016/j.tre.2021.102325.
- MacKle, Conor, Raymond Bond, Hannah Torney, Ronan McBride, James McLaughlin, Dewar Finlay, Pardis Biglarbeigi, Rob Brisk, Adam Harvey, and David McEneaney. 2020. "A Data-Driven Simulator for the Strategic Positioning of Aerial Ambulance Drones Reaching Out-of-Hospital Cardiac Arrests: A Genetic Algorithmic Approach." *IEEE Journal of Translational Engineering in Health and Medicine* 8 (May). doi:10.1109/JTEHM.2020.2987008.
- Macrina, Giusy, Luigi Di Puglia Pugliese, Francesca Guerriero, and Gilbert Laporte. 2020. "Drone-Aided Routing: A Literature Review." *Transportation Research Part C: Emerging Technologies* 120 (February). Elsevier Ltd: 102762. doi:10.1016/j.trc.2020.102762.
- Moon, Ilkyeong, Said Salhi, and Xuehao Feng. 2020. "The Location-Routing Problem with Multi-Compartment and Multi-Trip: Formulation and Heuristic Approaches." *Transport metrica A: Transport Science* 16 (3): 501–28. doi:10.1080/23249935.2020.1720036.
- Moshref-Javadi, Mohammad, and Matthias Winkenbach. 2021. "Applications and Research Avenues for Drone-Based Models in Logistics: A Classification and Review." *Expert Systems with Applications* 177 (March). Elsevier Ltd: 114854. doi:10.1016/j.eswa.2021.114854.
- Newell, Gordon F. 1973. "Scheduling, Location, Transportation, and Continuum Mechanics: Some Simple Approximations to Optimization Problems." SIAM Journal on Applied Mathematics 25 (3): 346–60. doi:10.1137/0125037.
- Oeser, Gerald. 2019. "What's the Penalty for Using the Square Root Law of Inventory Centralisation?" *International Journal of Retail and Distribution Management* 47 (3): 292–310. doi:10.1108/IJRDM-05-2017-0108.
- Ouyang, Yanfeng, and Carlos F. Daganzo. 2006. "Discretization and Validation of the Continuum Approximation Scheme for Terminal System Design." *Transportation Science* 40 (1): 89–98. doi:10.1287/trsc.1040.0110.
- Perboli, Guido, Luce Brotcorne, Maria Elena Bruni, and Mariangela Rosano. 2021. "A New Model for Last-Mile Delivery and Satellite Depots Management: The Impact of the on-Demand Economy." *Transportation Research Part E: Logistics and Transportation Review* 145 (November 2020). Elsevier Ltd: 102184. doi:10.1016/j.tre.2020.102184.
- Pham, Thi Ngan, Albert Tan, and Alvin Ang. 2020. "Determining Safety Stock for an Omni-Channel Environment." *International Journal of Information Systems and Supply Chain Management* 13 (2): 59–76. doi:10.4018/IJISSCM.2020040104.
- Rojas Viloria, Daniela, Elyn L. Solano-Charris, Andrés Muñoz-Villamizar, and Jairo R. Montoya-Torres. 2020. "Unmanned Aerial Vehicles/Drones in Vehicle Routing Problems: A Literature Review." *International Transactions in Operational Research* 0: 1–32. doi:10.1111/itor.12783.
- Rosenberg, Leonardo N., Noemie Balouka, Yale T. Herer, Eglantina Dani, Paco Gasparin, Kerstin Dobers, David Rüdiger, Pete Pättiniemi, Peter Portheine, and Sonja van Uden. 2021. "Introducing the Shared Micro-Depot Network for Last-Mile Logistics." *Sustainability (Switzerland)* 13 (4): 1–21. doi:10.3390/su13042067.

- Rutten, W.G.M.M., P.J.M. van Laarhoven, and B. Vos. 2001. "An Extension of the GOMA Model for Determining the Optimal Number of Depots." *IIE Transactions* 33 (11). Kluwer Academic Publishers: 1031–36. doi:10.1023/A:1010998620314.
- Salama, Mohamed, and Sharan Srinivas. 2020. "Joint Optimization of Customer Location Clustering and Drone-Based Routing for Last-Mile Deliveries." *Transportation Research Part C: Emerging Technologies* 114 (July 2019). Elsevier: 620–42. doi:10.1016/j.trc.2020.01.019.
- Schermer, Daniel, Mahdi Moeini, and Oliver Wendt. 2020. "The Drone-Assisted Traveling Salesman Problem with Robot Stations." *Proceedings of the 53rd Hawaii International Conference on System Sciences* 3: 1308–17. doi:10.24251/hicss.2020.161.
- Schwarz, Leroy B. 1981. "Physical Distribution: The Analysis of Inventory and Location." *AIIE Transactions* 13 (2): 138–50. doi:10.1080/05695558108974546.
- Shavarani, Seyed Mahdi, Mazyar Ghadiri Nejad, Farhood Rismanchian, and Gokhan Izbirak. 2018.
 "Application of Hierarchical Facility Location Problem for Optimization of a Drone Delivery System: A Case Study of Amazon Prime Air in the City of San Francisco." *International Journal of Advanced Manufacturing Technology* 95 (9–12): 3141–53. doi:10.1007/s00170-017-1363-1.
- Shen, Yaohan, Xianhao Xu, Bipan Zou, and Hongwei Wang. 2021. "Operating Policies in Multi-Warehouse Drone Delivery Systems." *International Journal of Production Research* 59 (7). Taylor & Francis: 2140–56. doi:10.1080/00207543.2020.1756509.
- Sheth, Manali, Polina Butrina, Anne Goodchild, and Edward McCormack. 2019. "Measuring Delivery Route Cost Trade-Offs between Electric-Assist Cargo Bicycles and Delivery Trucks in Dense Urban Areas." *European Transport Research Review* 11 (1). European Transport Research Review. doi:10.1186/s12544-019-0349-5.
- Statistics Canada 2016 Census Data Tables. (Spatial And Numerical Data Services, University of Calgary (Distributor), accessed October 01, 2020). https://www12.statcan.gc.ca/census-recensement/2016/dp-pd/dt-td/Rp-eng.cfm?LANG=E&APATH=3&DETAIL=0&DIM=0&FL=A&FREE=0&GC=0&GID=0&GK=0 &GRP=1&PID=109525&PRID=0&PTYPE=109445&S=0&SHOWALL=0&SUB=0&Temporal=2 016&THEME=115&VID=0&VNAMEE=&VNAMEF=.
- Stolaroff, Joshuah K., Constantine Samaras, Emma R. O'Neill, Alia Lubers, Alexandra S. Mitchell, and Daniel Ceperley. 2018. "Energy Use and Life Cycle Greenhouse Gas Emissions of Drones for Commercial Package Delivery." *Nature Communications* 9 (1). Springer US: 1–13. doi:10.1038/s41467-017-02411-5.
- Tsao, Yu Chung, and Jye Chyi Lu. 2012. "A Supply Chain Network Design Considering Transportation Cost Discounts." *Transportation Research Part E: Logistics and Transportation Review* 48 (2): 401–14. doi:10.1016/j.tre.2011.10.004.
- Tsao, Yu Chung, Divya Mangotra, Jye Chyi Lu, and Ming Dong. 2012. "A Continuous Approximation Approach for the Integrated Facility-Inventory Allocation Problem." *European Journal of Operational Research* 222 (2): 216–28. doi:10.1016/j.ejor.2012.04.033.
- Unnikrishnan, Avinash, and Miguel A Figliozzi. 2020. "A Study of the Impact of COVID-19 on Home Delivery Purchases and Expenditures." Working Paper, Portland State University
- Urban Freight Lab. 2020. "Common MicroHub Research Project: Research Scan." Seattle. https://depts.washington.edu/sctlctr/sites/default/files/research_pub_files/SCTL-Microhub-Research-Scan.pdf.
- Van Heeswijk, Wouter, Rune Larsen, and Allan Larsen. 2019. "An Urban Consolidation Center in the City of Copenhagen: A Simulation Study." *International Journal of Sustainable Transportation* 13

(9). Taylor and Francis Ltd.: 675-91. doi:10.1080/15568318.2018.1503380.

- Wang, Xiaokun Cara, and Yiwei Zhou. 2015. "Deliveries to Residential Units: A Rising Form of Freight Transportation in the U.S." *Transportation Research Part C: Emerging Technologies* 58. Elsevier Ltd: 46–55. doi:10.1016/j.trc.2015.07.004.
- Wang, Xin, Michael K. Lim, and Yanfeng Ouyang. 2017. "A Continuum Approximation Approach to the Dynamic Facility Location Problem in a Growing Market." *Transportation Science* 51 (1): 343–57. doi:10.1287/trsc.2015.0649.
- Waters, N. M., S. C. Wirasinghe, A. Babalola, and K. E.D. Marion. 1986. "Location of Bus Garages." *Journal of Advanced Transportation* 20 (2): 133–50. doi:10.1002/atr.5670200204.
- Wijeratne, A., and S. C. Wirasinghe. 1986. "Estimation of the Number of Fire Stations and Their Allocated Areas to Minimize Fire Service and Property Damage Costs." *Civil Engineering Systems* 3 (1): 2–6. doi:10.1080/02630258608970417.
- Wirasinghe, S C, and U Vandebona. 1999. "Planning of Subway Transit Systems." In *The 14th International Symposium on Transportation and Traffic Theory*, edited by Avishai (Avi) Ceder, 759–77. Jerusalem.
- Wirasinghe, S C, and N. M. Waters. 1983. "An Approximate Procedure for Determining the Number, Capacities and Locations of Solid Waste Transfer-Stations in an Urban Region." *European Journal* of Operational Research 12 (1): 105–11. doi:10.1016/0377-2217(83)90185-6.

Appendix A. Stage One Notation

Cost Parameters Notation	Name	Description
C _f	Fixed Facility Cost	Basic cost per MFC
Cs	Storage Space Cost Cost per square meter of storage sp	
		MFC
Cd	Resupply Delivery Cost	Cost for a resupply of the MFC
C _h	Inventory Holding Cost	Cost of holding one unit of capital cost
		per year
Cl	Last mile drone cost	Cost per kilometre of drone movement
Ct	Truck last mile cost	Cost per kilometre of truck movement

Variables Notation	Name	Description
Ν	Number of MFCs	Number of MFCs
Е	Resupply Frequency	Annual number of truck deliveries from
		the regional warehouse to the MFC.

Logistical Parameters	Name	Units
Notation		
μ	Mean Demand Density	Packages per area per year
σ	Standard Deviation of	Packages per area per year
	Demand Density	
А	Area	Square kilometres
φ	Configuration Factor	Unitless
m	Effective Storage Density	Packages per square meter
С	Truck Capacity	Packages per truck
d	Logistical Sprawl	Kilometres
Т	Delivery Time Window	Number of time windows per workday
k	Touring Co-efficient	Unitless
u	Wholesale Value	Unit cost
β	Inventory Stock-out Factor	Unitless
MFR	Maximum Flight Rage	Kilometres

Model Outputs Notation	Name	Units
W_{lpha}	Cycle Stock	Number of packages
\mathbf{W}_{β}	Safety Stock	Number of packages
L _{Drone}	Annual Drone Distance	Kilometres per year
L _{Truck}	Annual Truck Distance	Kilometres per year

Subscript	Zones
С	Communities
U	Service Catchments
R	Region

Appendix B. Graphical Solution for Objective Function Optimisation

Figure 8 shows the graphical solution for the regionally uniform scenario. The left-hand sides (LHS) of equations (8a) and (8b) are similar when the number of facilities is large. 8a LHS is always larger than 8b LHS. Consequently, 8b provides an upper bound for when the right-hand side (RHS) of 8a crosses the LHS of 8a.



Figure 9: Graphical Solution

Appendix C. Allocation Algorithm

The greedy search algorithm employed in the city of Calgary example is a branch and bound algorithm. An overview of the processes is shown in Figure 10 and described in Table 8.





Table 8: Allocation Algorithm

Initial Allocation Process

- 1. For each catchment area, 1 to \widetilde{N}_R^* :
 - a. SEED: Select the community with the fewest neighbours as a starting community for a new catchment area.
 - i. If multiple communities with equal value, then choose one of these at random.
 - b. BRANCH and BOUND: Add into the catchment area the neighbour community of this starting community that has the largest N* value, that does not increase the sum of the catchment's N* values above $\frac{N_R^*}{\tilde{N}_R^*}$.
 - c. Repeat step 1b until there are no more neighbours of the starting community remaining or until none remaining can satisfy the objective function constraint.
 - d. Repeat steps 1b and 1c with the community after the starting community in place of the starting community, until no neighbours of the catchment can be added within the objective function constraint.

Reallocation Process

- 2. ALLOCATE ORPHANS: Randomly allocate all unallocated communities from the initial allocation to a catchment that the community is a neighbour of. For unallocated communities without allocated neighbours, allocate all communities that can be allocated first then repeat until all communities have been allocated.
- 3. SWAP BOUNDARY COMMUNITIES: Select the catchment area contributing most to the objective function. For all communities at the border with another catchment, evaluate the effect of swapping that community with the other catchment on the objective function. Swap the community that most lowers the objective value.
- 4. Repeat step 3 until no swaps of any community between any two catchment areas can improve the objective function.

Appendix D. Parameter Values

D.1 Demand

A seller and service provider (SSP) likely has an estimate of future demand for their products within a region using data from previous years and forecasts several years into the future for this demand. They may even have a model that predicts demand for a community based upon the socio-demographic data available to them. By multiplying the expected demand per person with the known and non-uniform population density across a city, one can calculate the mean and standard deviation for demand density for each community.

To estimate this demand for the e-retailer in the example, this study combines data from Canada Post, an academic prevalence of e-commerce shopping survey, and an industry estimate of an existing similar company's (Amazon) market share. The average number of e-commerce deliveries per e-commerce shopper in Canada is 16.6 according to a Canada Post survey. A study by Unnikrishnan and Figliozzi (2020) found similar results and that 6.8% of adults do not use online retailers at all, and so this number should be discounted from the analysis. We then assume market share equates to share of deliveries and that Amazon accounts for 40% of the e-commerce market. Finally, Amazon stated in their application to the Federal Aviation Administration that their drones could carry 80% of the items they sell. Combined, this gives an average delivery rate of 4.95 deliveries per adult per year. Clearly, no adult orders fractions of packages, and this number is an average of a range of values of individual customers' ordering behaviour. Canada Post categorises online shoppers into six groups: from one-time and occasional shoppers (who order 1 and 2-6 packages annually, respectively) to hyper and hyper-elite shoppers (who order an average of 25-to-40 and 40+ packages annually, respectively).

We analysed e-commerce sales data from January 2016 to February 2020 to determine an estimate for the uncertainty in demand by fitting a normal distribution and accounting for an exponential growth trend over time, we found the standard deviation to be 6.42%, translating to 0.32 deliveries per person per year when applied to the previous average. This demand uncertainty is not equivalent to expedited delivery packages and could be entirely composed of standard packages.



Figure 11: E-commerce Retail Sales Trend

D.2 Drone Last Mile Delivery Costs

A foremost cost to consider in a drone system model is the cost per unit distance to operate the drones. A reliable industry estimate for the cost per kilometre or flight hour for a delivery drone is currently not yet available due to the novel, uncertain, and often proprietary nature of the technology. A report by Figliozzi (2018), however, proposed two values, an optimistic regulatory and pessimistic regulatory environment, resulting in estimates of \$14.98 USD and \$67.64 USD per flight hour, respectively. These estimates include discounted purchase cost, battery replacement cost, electricity cost, and staffing cost. The latter of these costs is the most significant and most uncertain variable, and thus the cause of the wide range of costs between the considered scenarios; an operator overseeing many drones or only one depends on local regulations that can significantly affect cost. In this analysis, we assume the optimistic scenario. Furthermore, given that we assume effective use of drones over time windows, we assume that flight time is directly proportional to flight distance, also implying that take-off and hover times are negligible or accounted for in the average travel speed. We assume drone average travel speed to be 60 km/h, and we assume a USD to CAD conversation rate of 1.3 to give the value of \$0.325 CAD per kilometre.

D.3 Fixed Facility and Shelf Cost

Fixed and shelf costs are discussed together as they are best understood by examining some sample data. Figure 12 shows annual rent data for industrial units in the city of Calgary from a search conducted in March 2020 using an online retail platform, Spacelist. We found 202 industrial units available for rent in the city at the time of the search that were both suitable for conversion into a MFC and under 2000 square meters, a case-specific arbitrary maximum that we selected based on observation of the unsuitability of units above this size for conversion.

The linear relationship fit to the Calgary data, as seen in Figure 7, gives an annual cost per square meter of nearly \$120 CAD (estimate slope) with a constant annual levy (y-intercept) of nearly \$18,000 CAD. This trend line fits the data with an R2 value of 0.7433. Expensive outliers tend to be downtown real estate options and cheaper outliers tend to be in industrial estates; the fixed facility cost varies per community, but we found that this cost does not significantly affect the results, and so it was not included in the main analysis. The levy forms a term of the fixed facility cost, with per MFC personnel wages, per MFC administration costs, and other such charges constituting the rest of the MFC fixed cost. The former per area cost comprises the heating, lighting, security, and other per square area charges that constitute the remainder of the unit shelf cost.

The second part of the fixed facility cost is personnel and administration costs required for each MFC, regardless of size. These costs can be difficult to estimate without knowledge of the specific company being modelled due to the variability between companies' salaries and operations. However, a North American industry report published by Warehousing and Fulfilment online suggests costs to be in the region of \$70,000 CAD per year for employee salaries, which gives a total fixed cost of approximately \$90,000 CAD per year per MFC.



Figure 12: Yearly Rental Cost versus Property Area in Calgary Area, March 2020

The remainder of the shelving cost is comprised of utility costs such as heating, lighting, and water. These costs are similarly difficult to estimate accurately without knowledge of the company to be modelled, but the same industry report suggests that multiplying the unit area cost by 50% is a reasonable estimate. Thus, the shelf cost used in the example is \$180 CAD per square meter per year.

A sensitivity analysis was performed using these parameters, but we found that the parameters affected the solution in a manner expected from previous research, and so we did not include them in the discussion. Increasing the fixed facility cost decreased the number of MFCs and increasing the shelf cost also decreased the number of MFCs, conditional on the effective storage density of the goods as discussed.

D.4 Resupply Delivery Costs

The resupply costs include costs associated with a delivery from the regional warehouse hub to the MFCs. These costs include the wage of the driver, fuel, the discounted ownership of the delivery vehicle, and wages of MFC unloading dock workers.

For this example, the SSP operates the resupply fleet themselves. To determine cost, we assume driver wage is approximately one-quarter of the operating cost of each delivery. This assumption is based on a report by the American Transportation Research Institute (ATRI). Based on an online search, we estimate that the latest average wage for a delivery driver in Calgary is \$25 CAD per hour; thus, the daily wage is \$200 and the daily operating cost of the truck \$600, with one delivery possible per day. The driver is assumed to assist two SSP employees with unloading the items into the MFC; the SSP employees cost \$300 per day, which is a little above the current Alberta minimum wage for eight hours of labour per worker for simplicity. The unloading employees would likely be assigned to a different MFC each day, meeting the delivery driver as they resupply a different MFC every few weeks (the inverse of delivery rate). If the internal handling of goods within the MFC and drone loading can be automated in the future, then this

regular resupply delivery may be the only human intervention in the system besides maintenance. Considering these factors, we use a total resupply delivery cost of \$1100 per delivery per MFC.

This cost is of interest when determining the resupply frequency (E) because the resupply frequency affects the relative significance of the shelf and holding cost terms in (5). However, once the resupply frequency is determined, the cost is directly proportional to the number of MFCs, and it effectively becomes an added fixed facility cost; it also affects the optimal number of MFCs in the same manner.

D.5 Inventory Costs

The cost to hold inventory is typically dominated by the interest rate charged by a bank and other factors, such as refrigeration, or inventory shrinkage depending on the item characteristics (Axsäter 2006). This example uses an inventory cost rate of 20% of item cost per year. Furthermore, the value of the goods sold by this SSP is assumed to be \$35 CAD, as this amount is the average spent on an e-commerce retail purchase according to Canada Post. This amount results in an inventory cost of \$7 CAD per item per year.

The SSP is assumed to be moderately risk averse and wishes to avoid stock outs in 95% of all demand scenarios, and thus they accept a stock-out factor of two, which relates to this risk tolerance. See Schwarz (1981) and Axsäter (2006) for further discussion of this factor.

D.6 Truck Last Mile Delivery Costs

Cost per kilometre for a traditional truck delivery system is taken from a report by the ATRI and adjusted to Canadian dollars and kilometres. This report accounts for vehicle capital costs, maintenance costs, insurance costs, driver wages, and benefits, and has been used by previous researchers for similar estimation purposes (Sheth et al. 2019). The ATRI report estimates the average speed of delivery trucks to be approximately 64 km/h, which is likely higher than an urban delivery truck's true average speed, and so the cost per kilometre is a low conservative estimate, which is consistent with the conservative estimate for drone costs.

Finally, given shorter time windows, different sizes of trucks may be preferrable and thus affect the assumed unit cost. This study assumes that only one size of truck is available to the supplier although different costs related to the necessary capacity can be included if that data were available.