

Airspace Network Design for Urban UAV Traffic Management with Congestion

Leanne Stuive, Fatma Gzara*

Department of Management Science and Engineering, University of Waterloo

Abstract To support the safe and widespread use of unmanned aerial vehicles (UAVs) in urban environments, industry stakeholders and regulatory authorities are partnering to develop urban airspace traffic management systems (UTMs). UTM system providers face strategic decisions in how to design and manage airspace available to UAV flights. We consider a provider that plans to open an urban airspace in which UAV flights are routed above existing roads in 3D corridors corresponding to segmented altitude levels. The provider aims to select a subset of the road network to form an air-network with the goal of providing safe and cost effective service for UAV traffic. The air-network selected must provide routes that respect UAV technology restrictions, and must have adequate capacity to support the expected flight volume. We develop a 3D airspace network design model that selects a subset of roads whose 3D projection into the sky will be used for routing flights. The constrained system optimum (CSO) traffic assignment model is used to evaluate the quality of the network; the CSO user constraints represent battery restrictions while minimizing the total travel time ensures realistic routing in the face of congestion. To incorporate the 3D nature of flights, we use simulation to calibrate a Bureau of Public Roads capacity parameter that reflects the multiple vertical layers of airspace made available when a road is selected for the network. We introduce a methodology to derive candidate maps for urban areas and use it on open-source data to build a case study for Chicago city centre. We assess the impact of budget, congestion, minimum-path deviation, and demand patterns on network designs.

Keywords: Airspace network design, UAV traffic management, airspace congestion

1 Introduction

As unmanned aerial vehicle (UAV), or drone, technology evolves, UAVs have become an increasingly viable mode for transporting goods. UAVs are not limited by road traffic and provide a greener alternative to trucks for last mile delivery of consumer packages in some settings [31]. In 2022, Amazon Prime Air [49] announced trials in Lockeford California, Alphabet’s Project Wing trialed UAV delivery in Ireland [39], and Walmart made UAV delivery available to four million customers [51]. Smaller companies are also active in the UAV

*Corresponding Author: fgzara@uwaterloo.ca

delivery space, with Flytrex receiving Federal Aviation Administration (FAA) permission to increase the flight radius of their delivery zones [20] and SkyDrop partnering with Dominos Pizza in New Zealand [48]. Continued advances in UAV battery life and control capability - especially in inclement weather - will further the case for wide-spread UAV usage.

However, a major hindrance to UAVs at scale is the lagging development of government regulations to support the safe and fair operation of many simultaneous flights. Government regulations surrounding UAV usage differ by country and continue to evolve; commercial users must be aware of and adapt to local requirements [50]. Various frameworks for regulating UAV flights have been proposed [30]; many of these frameworks depend on a UAV traffic management system (UTM) to track ongoing UAV flights and facilitate the approval of new flight requests. The UTM concept of operations provides necessary roles (ex. UAV pilot, traffic manager), processes (ex. flight request, monitor traffic), and technologies (ex. supporting software and datasets) to manage UAV traffic in an airspace [32]. The concept is similar to existing air traffic management systems and the civil aviation authority (ex. FAA in USA, NavCanada in Canada) is leading the integration of UAVs in many countries. We focus here on a pressing technical design issue for UTMs; for an economics and policy perspective see [18], and for a management perspective see [36].

Within the UTM concept of operations, we use the term *UTM provider* for an entity that provides the tools required to develop and administer a UAV airspace. In other words, a UTM provider builds the tools, software, and system required to implement an urban airspace within the UTM framework. One key process in a UTM provider solution is the submission and approval of flight requests via a Web API, or similarly automated system. In some cases the corresponding aviation authority ultimately grants the flight request approval. For example, in the United States, UAV flights under 400ft in controlled airspace must receive approval from the FAA. This approval can be obtained through a UTM provider that has integrated with the automated Low Altitude Authorization and Notification Capability system. In other cases the approval of flights may be required only at the UTM level; flight notification and approval allow the traffic manager to ensure that there are no conflicting flights. Although detect-and-avoid capability for UAVs is improving, reservation of airspace coupled with flight request approval is the next step towards unlocking UAV flights beyond visual line of sight, and at scale.

Current iterations of UTMs are focused on flight request approvals in low risk areas. As flight requests in densely populated cities increase, the design of airspace and management of simultaneous UAV flights in urban environments is a critical area of UTM design. The “best”, or *greedy*, path for a single UAV to travel between two locations is point-to-point (i.e. via the Euclidean shortest path), avoiding obstacles, and at its ideal altitude (which depends on flight length, weather, etc.). If many flights are requested in a short period of time, their greedy paths may conflict and UAV pilots would have to either wait long times to start with

a cleared path, or drain battery while hovering en route. A well-designed urban airspace could segment and separate traffic to prevent the likelihood of such path conflicts. Bauranov and Rakas [5] summarize recent academic and industry literature on urban airspace design and various UTM design proposals. They define the *physical structure of urban airspace* to refer to the position and size of airspace elements such as flying trajectories, tubes, corridors, and layers, as well as the associated rules of operation.

The design of airspace for traditional aircraft provides guidance for UTM airspace, but key differences apply. As UAVs fly at low altitudes, tall buildings and ground features are physical obstacles that must be avoided and can cause communication interference when flying nearby. The technical capabilities of UAVs are more disparate than for traditional aircraft because a UAV may be sized, built, and equipped for a very specific task. One defining technical difference may be the propulsion mechanics of the vehicle; some UAVs generate lift using wings (“fixed-wing”), but the more maneuverable UAVs are rotorcraft that rely on rotating blades to generate lift. Many UAVs rely on electric power, meaning flight length - in terms of both time and distance - is limited by battery capacity, rather than a conventional fuel supply. Most consumer UAVs need to recharge after 30-45 minutes of flight time. The distance between vehicles required for safe operation, or *separation standards*, are much smaller for UAVs than for traditional airplanes. Exact separation standards are still in development and expected to evolve over time to reflect the maturity of the underlying systems and technology.

Although UAVs are capable of direct point-to-point flight, an airspace design that restricts UAVs to fly along existing road networks provides many advantages in urban settings. A similar design is discussed by Jang et al. [29] who propose using “sky lanes” in urban centres to mimic street organization. Doole et al. [19] focus on flying above roads, observing it provides separation, reduces privacy concerns, and enables ground vehicle coordination. To test the airspace design, they develop one-way and two-way airspace configurations with altitude levels segmented by turning/through traffic and heading direction. Then, they simulate air traffic over all the streets of Manhattan, assuming conflict avoidance procedures instead of flight approval processes. As UAVs in the city scale, a version of this design which integrates flight approvals is a more practicable next step. A UAV airspace above roads is also a practical choice because it enables the rapid collection of the telecommunication, location, and weather data required to provide precise safe paths for UAV flight. Some of this data must be obtained with physical measurements made by traversing the underlying terrain with a vehicle. This procedure is easiest to perform when the ground below is an already navigable street. The term *mapping* will refer to the procedure of collecting data on the streets above which UAVs will fly.

Mapping and maintaining the data required for safe UAV flight over every city street would be cost-prohibitive. UTM providers may prefer to map a portion of streets so that they can safely scale capacity, be responsive to a change in regulation, and build infrastructure that is supported by demand. In this work,

we take the point of view of our industry partner, whose needs align with a UTM provider opening up UAV airspace by implementing a UTM system. The tactical decision they face is determining which roads of a city to map for UTM service first. We develop an optimization model to select a set of roads in a city to support UAV flights. Once selected and mapped, the roads, sky above, and supporting data provide the initial 3D pathways of a scalable UAV traffic network, or *UTM network*. UAV pilots will request to fly between points on the network and, with approval, execute flight plans restricted to a layer of the airspace above the mapped network.

The quality of a candidate UTM network is measured by its ability to service expected flight requests. UAV pilots provide flight requests by specifying their origin, destination, and battery limit. A network has adequate connectivity to service a flight request if there is an origin to destination 3D route whose trajectory is limited to airspace above mapped roads and whose trajectory can be traversed by the UAV within the time limit imposed by its battery life. If certain roads are oversubscribed, however, a flight request may not be accommodated in real time because the network has inadequate capacity and therefore there is congestion along the route. Either the UAV will need to wait too long for traffic to clear, or experience untenable delays en route due to hovering and excessive altitude switching. To incorporate both route length limits and congestion effects in our model, we use constrained system optimum (CSO) traffic assignment to evaluate routings allowed by a candidate network. CSO traffic assignment minimizes total travel time, where travel time along a road section is a function of the amount of traffic using the road. For example, The Bureau of Public Roads (BPR) function [40] is often used to model travel time as a function of traffic flow on ground highways.

The main contribution of this paper is the first UAV network design model for UTM providers. The model differs from existing UAV research in that it considers the 3D nature of flights through a congestion function, assumes UAVs fly over existing streets rather than via point-to-point flight, and explicitly evaluates routings which are designed in respect of regulatory constraints and airspace design. We use simulation to develop BPR congestion function parameters consistent with routing UAVs in 3D space in respect of separation standards and along multiple corridors according to altitude. The single level model we develop is the first to extend the constrained system optimum road traffic model to include a network design component. As demand for UAV flights increases, UTM providers need to be agile to expand their network rapidly. To this end, the models and methodology developed can also be used in a network expansion context. We focus on initial network planning in this paper because UTM's are in early stages of development.

The remainder of the paper is structured hence. Section 2 summarizes related literature in UAV Operations Research, traffic modeling, and network design. Section 3 provides an exact mathematical formulation of the model we develop, explains how it captures various real-world constraints, and introduces the approximation we solve which uses piecewise linear approximation. Section 4 discusses designing congestion functions that

account for the projection of 2D ground roadways into 3D pathways in the sky. Section 5 contains the results of numerical experiments. Its primary focus is a case study of UTM network design for Chicago city centre, covering end-to-end modeling methodology from instance generation to sensitivity analysis to runtimes. Section 6 concludes with a summary of the results and future research directions.

2 Literature Review

To the best of our knowledge, no academic literature addresses selecting ground roadways in the interest of managing UAV traffic by projecting the chosen roads to 3D pathways in the sky. Of the abundant UAV-specific Operations Research work, the most closely related models focus on designing a UAV distribution network for a company that is integrating a UAV fleet into their logistics operations. These models assume point-to-point flight with recharging, and disregard regulatory constraints and path conflicts. In fact, the UTM network design problem we formulate shares many similarities to transportation planning for ground vehicles. In our literature review, we focus first on existing UAV-specific Operations Research models and then discuss the most relevant concepts in ground traffic routing and network design.

2.1 Operations Research Models for UAV Applications

As the use of UAVs in civilian and commercial enterprises has become more economically viable, many Operations Research practitioners and researchers have studied the related strategic, tactical, and operational challenges. The authors of [37], [41], and [45] survey the breadth of logistics models for civilian UAV applications. Each survey points to integrating regulatory constraints into models within its discussion of future research directions. In particular, Moshref-Javadi and Winkenbach [37] observe that risk, chaos, and safety issues are of “growing concern” as UAV usage increases. They identify dedicated aerial highways, standard routing protocols, and stochastic risk models as possible ways to mitigate these concerns. Agatz et al. [41] observe that it is “important to work out technological and regulatory solutions that will address public concerns of privacy and safety without impeding value-creating drone deployment”. Lastly, Poikonen and Campbell [45] observe that “thus far, there has been little research on drone routing related to UTM and constraints UTM may impose”.

The UAV logistical challenges receiving the most attention have been operational problems and, more specifically, problems related to routing (i.e. choosing flight paths and schedules). Models and solutions have been proposed for: routing fleets of delivery UAVs, coordinating UAV and truck deliveries, routing UAVs for inspecting physical infrastructure, and routing UAVs for surveillance purposes. For surveys on some of these problems, see [13] and [33]. Most routing-based work assumes UAVs fly directly point-to-point and few consider flight approval requirements, flight delays caused by waiting for traffic, and regulatory restrictions on where a UAV may fly. One exception is [11] which considers optimization of traffic managed within a UTM

as an extension of classical air traffic flow management. The authors consider approaches to fairly integrate flights known in advance with on-demand requests, showing fairness can be improved while maintaining efficiency. They perform a case study for package delivery in Toulouse France.

The managerial models most closely related to UTM network design pertain to designing distribution networks that integrate a UAV fleet. The strategic challenge faced is determining the commercial viability of using a UAV fleet for certain deliveries. To perform the analysis, optimized decisions regarding infrastructure placement are considered. For example, Shavarani et al. [46, 47] use genetic algorithms to solve models minimizing cost, travel distance, and lost demand for UAV delivery systems. These models capture: the strategic costs of building charging stations, building warehouses, and procuring UAVs; and the operational costs incurred when operating the UAVs to serve demand. Chauhan et al. [10] address a similar problem using heuristics, but do not incorporate as many operational features such as congestion at facilities. Other work addresses only the problem of charging station placement; for example [14], [25], and [44]. Baloch and Gzara [4] also model strategic network design for UAV parcel delivery and study sensitivity to technological limitations and government regulation. They deviate from other strategic models in that the objective is not to minimize delivery cost but rather to maximize return from a set of delivery services offered by an e-retailer where customer demand is utility driven. All of these models incorporate battery restrictions by limiting the capabilities of a UAV within the model constraints or instance considered.

Hou et. al. [26] provide the first UAV distribution network design model that incorporates lane capacity in addition to charging needs. The model is also differentiated in its use of distributionally robust modeling to deal with demand uncertainty. The authors observe that “total infrastructure costs can be saved by pooling drone flows into a small number of high-capacity channels/transfer airports”. However, the lanes they consider are Euclidean shortest paths between candidate facility locations, rather than 3D corridors above existing ground roadways. To model airspace congestion, the capacity of these lanes is taken to be a hard limit determined by an investment decision.

All the distribution network design models cited so far assume, often implicitly, that: (1) the fleet operator does not need to compete with other airspace traffic, and (2) UAVs may fly freely in the local airspace. Our work provides an alternative point of view - that of the UTM provider which will eventually operate the high demand airspace in urban centres and impose system constraints on UAV flight. This point of view investigates a research gap identified by the UAV logistic model surveys, and may inform future operational models which will need to integrate UTM constraints.

2.2 Related Concepts in Traffic Modeling and Network Design

Determining how UAV traffic would be routed on a fixed network is, in and of itself, a nontrivial problem. The class of mathematical models that determine routings are referred to as *traffic assignment* problems. What is common to all these models is that they predict traffic flow patterns while accounting for congestion effects. Traffic assignment problems date back to the seminal work of Wardrop, who introduced the user equilibrium (selfish routing on shortest path) and system optimal (smallest possible total travel time) routing principles [52]. From a practitioner’s point of view, traffic assignment is one aspect of a transportation forecasting model, alongside trip generation, trip distribution, and mode choice; for more information see [42], [12] and [43]. Here we restrict our discussion to CSO traffic assignment because it closely models the routing requirements we choose for our UTM network design problem. The CSO traffic assignment problem is a static model which performs system optimal routing subject to user constraints. As a static model, the CSO problem assesses steady-state behaviour via a congestion function. The goal is to route vehicles in a system optimal way; that is, the objective is to minimize the total travel time on the network. However, the paths chosen are subject to a user constraint which requires the individual travel time from origin to destination for each vehicle to respect some corresponding limit.

The CSO problem first appeared in Jahn et. al [27, 28], where the authors studied traffic assignment models that minimize total travel time subject to route lengths being within a given multiplicative factor (usually between 105% and 150%) of the length of the shortest path. The length of a path is determined by the sum of the lengths of the arcs en route. The model is restricted in that the arc lengths must be *normal*, meaning that the length of an arc cannot depend on the amount of traffic assigned, and further implying that the set of allowable paths is fixed prior to the optimization. The authors use a Frank Wolfe type algorithm alongside column generation to solve the model and they analyze the results for a subset of the Transportation Networks for Research [22] instances. Angelelli et. al [2, 1] also study practical instances of the CSO problem. They solve the problem using piecewise linear approximation and heuristic path generation techniques. These techniques allow for the arc lengths to depend on the amount of traffic assigned, and computational results show that the solutions found are of good quality.

Study of the CSO model emerged alongside the widespread availability of navigation devices that direct drivers to their destination. The devices could, in theory, act as a control mechanism by guiding drivers on routes that alleviate congestion to boost system efficiency. Route lengths limits are considered to ensure the system is fair, and therefore no user would be inclined to disregard the navigation device’s advice. Our interest in the CSO model differs from that of real-time traffic directors. We take the perspective of transportation planners; the traffic assignment results are a means to evaluate the routes available on a network when safety requirements are considered. UTM network users (i.e. UAV pilots) will be required to follow the system optimal routing for regulatory reasons and route length limits primarily reflect real physical concerns

pertaining to battery life. Fairness and competitiveness with alternative transport modes are secondary considerations.

While traffic assignment pertains to routing vehicles/flow on a fixed network, *network design problems* are primarily concerned with selecting arcs to form a network. The corresponding models typically include arc-based binary variables that are used to model the inclusion or exclusion of arcs in the network. Network design models are prevalent in transportation planning; for more information see [17]. Here, we restrict our discussion to network design models which incorporate congestion and path length limits. Bayram et al. [6] consider a disaster response application where individuals are instructed to evacuate their homes and drive along roadways to a specific shelter location. The objective is to minimize total congestion subject to CSO routing of evacuees and a budget limit on the number of shelters that may be opened. The network design decision is which nodes to open; a node represents the candidate location of a shelter. The authors solve the model with second order cone programming (SOCP) techniques. SOCP techniques are also used by Gürel [24] to minimize congestion plus arc capacity installation cost subject to routing all flow on a network. However, the network routing is not subject to path length limits. Most transportation modeling work that considers congestion but not path length limits is done respecting user equilibrium routing [34]. For example, Mathew and Sharma [35] consider capacity expansion decisions for a road network with the goal of minimizing total congestion.

In capacity expansion applications, an arc’s capacity may be increased incrementally at a per unit cost. Our UTM network design problem differs in that UTM network capacity decisions are binary. Either the arc is open, and the corresponding layers of airspace are available, or, the arc is closed and no airspace is available. The capability of open airspace is not a model decision but a parameter determined by separation requirements coupled with flight ceilings, and is reflected in the congestion function. Contrary to work on expanding existing road traffic networks, our UTM network design model applies to transportation planning for an emerging mode of transport with no existing infrastructure. We compare CSO traffic assignment for any networks within budget, as opposed to weighed against capacity installation cost.

3 UTM Network Design Model

The core decision to model is simple: which roads should be selected for use by UAV flights? There is a cost of selecting a section of road which reflects both the initial mapping cost and any ongoing maintenance costs. Throughout, we will refer to the roads selected as forming a *UTM network* because this represents the primary network design decision. However, it’s important to recognize that mapping a road will make multiple vertical layers of airspace available to UAV traffic because UAVs operate in three dimensional space.

To assess a candidate UTM network, we require as input an indication of where UAV flights are expected

to originate and terminate. This information will be represented as an OD-list where each element is a tuple that indicates: the starting location (*origin*), ending location (*destination*), and a demand number specifying the expected number of flights per hour from origin to destination. This is similar to static road traffic assignment problems where an OD-matrix is used to specify road traffic flows between every pair of intersections in a road network. A list of *OD-pairs* is a more appropriate data structure than a matrix in our application because the demand is sparse; it will usually include locations of known UTM clients (often early adopters) and other pairs that target a coverage-based objective. This is in contrast to road traffic modeling which is based on existing, observed traffic patterns of much larger volumes (i.e. where the number of vehicles in the system per hour is magnitudes larger). We return to the discussion of demand patterns for UAV networks in Section 5.1. Note that a sparser demand pattern is also required for the models developed to be solvable.

A UTM provider will have many competing criteria for evaluating the quality of a UTM network. In this paper, we focus on *connectivity* and *capacity*; in Section 6 we outline other criteria which provide future research directions. In evaluating connectivity we aim to answer: for how many OD-pairs is there an acceptable route on the UTM network? In evaluating capacity, we aim to answer: how many UAV flights can the UTM network support? The answers to these questions, in large part, depend on how the UTM provider would route UAV flight requests on the network as they arrive in real-time. However, we do not embed an exact routing protocol into our optimization model because traffic rules and routing strategies for UAVs are in early stages of development. Instead, we use industry-informed assumptions and approximations to capture an indicative assessment of the network’s capabilities. This approach is consistent with road traffic methodology where static macroscopic traffic assignment models are used for high-level assessment and coupled with further investigation using simulation and queuing models. The time-dimension, or *dynamic*, nature of traffic is not considered in such models; a steady-state behaviour is assumed. Similarly, our model does not capture the propagation of traffic over time - for example, a bottleneck slowing traffic on incoming roads. The model does consider the high-level flight routing strategy in that it concurrently assigns to each OD-pair an acceptable path through the UTM network it selects.

A fundamental choice in routing strategy is whether a UAV flight request will: (1) give both the end points and a specific flight path between them, or, (2) only the endpoints and allow the UTM provider to dictate the flight path. The second option may be preferred by both the UTM provider and UAV pilot, assuming the path offered by the UTM provider is chosen judiciously. For example, the flight path offered must respect the UAV’s battery restrictions which limit total distance travelled. The UTM provider maps and maintains the data required to safely route UAVs on exact paths and is, therefore, well-equipped to find such a good path. From the UAV pilot’s point of view, relying on expertise to plan flight paths is preferable because it increases the likelihood that the flight will be approved by the aviation authority. However, the flight

path cannot differ too greatly from the greedy path or shortest route on the road, because the company funding the UAV flight may otherwise find substitute ground transport. As UTM providers are incentivized to service as many flights as safe capacity allows, they must offer relatively short paths, and may also try to offer the same path for flights with the same endpoints to build customer trust. In summary, the UTM provider is best equipped to provide routes in real-time and therefore our network design model will assume system-optimal routing subject to user constraints restricting flight lengths.

A second fundamental choice in routing strategy is how *congestion* (i.e. traffic buildup at a network component) is ameliorated. When a UAV flight request is made, the flight path chosen by the UTM provider must not conflict with airspace reservations made by past flights. The UAV will need to either wait to begin its journey, or, hover en route and deplete its battery while it waits for preceding traffic to clear. This dynamic problem manifests in a steady-state model whenever there is a road segment that appears on the assigned paths of many OD-pairs. One technique for limiting congestion is to use a capacitated model which enforces a strict limit on the number of OD-pairs assigned paths through a given arc. Such hard limits are not ideal for capturing congestion trade-offs in traffic models, however, because capacity is not so rigid. It is more realistic to consider the observed effect of congestion on the model, rather than strictly forbid it.

In road traffic networks, congestion effects are captured using *congestion* / latency / volume-delay functions which measure the effect of the volume of traffic flow on travel time. The amount of time to traverse a road increases gradually as the number of vehicles using it increases, up until the point at which the number of vehicles on the road exceeds its *saturation capacity* (hypothetical maximum capacity). In a dynamic setting, the number of vehicles that can traverse the road would then decrease. However, in traditional steady-state traffic modeling, the objective is to minimize the total travel time on the network, where travel time is obtained by evaluating the congestion function for the amount of demand assigned to the road (which may not represent a steady-state flow). Above the saturation capacity, the congestion function is designed to increase sharply, to indicate the disutility incurred by traffic flows outside of steady state behaviour. However, such flows are not strictly forbidden. Congestion functions are also appropriate in our UAV context and we choose a model objective of minimizing total travel time, where travel time accounts for congestion effects. We choose a BPR congestion function; see the discussion in Section 4.

Following the model design decisions just discussed, Section 3.1 present a mathematical formulation of the model as a the mixed integer convex optimization problem. Section 3.2 describes how to use piecewise linear approximation to obtain a mixed integer linear problem which approximates the nonlinear model.

3.1 Formulation

Underlying Network Data: The underlying city street network is represented by a graph denoted by $G = (N, A)$ where N denotes the set of nodes (corresponding to intersections) and A denotes the set of undirected edges (corresponding to sections of road). Let $c_{ij} \geq 0$ denote the cost of including edge $(i, j) \in A$ in the UTM network, indicating the price of mapping the road segment to obtain the precise navigational data required to support UAV flights. Currently we expect the length of an edge - measured by driving distance along the centre of the road segment, and not Euclidean distance between its endpoints - to be a reasonable proxy for mapping cost. However, the model extends to generic costs. Let $B \geq 0$ denote the mapping budget which represents a hard limit on the total cost of the edges mapped. Once mapped, an edge may be traversed in either direction. For the purposes of representing routing decisions, we introduce notation for the graph $G' = (N, A')$, which is the symmetric directed graph obtained from G by replacing each edge in A with an arc in either direction.

UAV Flight Demand Data: Let K denote the set of OD-pairs, representing expected UAV flight requests. For each OD-pair $k \in K$, there is an affiliated origin $O_k \in N$, destination $D_k \in N$, and demand $d^k \in \mathbb{Z}_+$ which we interpret as the expected number of flight requests per hour from origin to destination. In this paper we assume that all UAV flights originate and terminate at an intersection, but the model can be extended to allow UAV flights to originate anywhere along a road segment.

Distance Data: Let $\ell_{ij} \geq 0$ denote the length of arc $(i, j) \in A'$. Let path length limit L_k for OD-pair $k \in K$ denote the maximum distance a UAV may fly to travel from O_k to D_k . The *path length limit* bound represents the restriction imposed by the battery life of an “average UAV” flying between the two locations. Note that this model can be extended to consider vehicle-specific arc distances and path length limits in a straightforward way by introducing arc lengths ℓ_{ij}^k for $(i, j) \in A$ and $k \in K$.

Travel Time Data: Let $t_{ij} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be the congestion function which provides the time a UAV takes to traverse arc $(i, j) \in A'$ as a function of the number of flights using the arc. It is used to evaluate the total travel time on the network. We assume congestion is one-way; that is to say, counter-flow traffic does not affect the travel time on an arc. Note that as arcs become congested, UAVs may become “stuck” behind slower traffic and thus modifying the travel times to be vehicle-specific is more complicated than vehicle-specific path length limits.

Decision Variables: Let design variable $z_{ij} \in \{0, 1\}$ indicate if edge $(i, j) \in A$ is selected for the network. This model restricts all flight requests for an OD-pair to be assigned to the same path in the network. Let routing variable $x_{ij}^k \in \{0, 1\}$ indicate if arc $(i, j) \in A'$ is on the path from origin O_k to destination D_k for OD-pair $k \in K$. Let helper flow variable f_{ij} indicate the amount of traffic assigned to arc $(i, j) \in A'$.

Objective Function and Constraints: For node $i \in N$, let $N_i^+ := \{j \in N : (i, j) \in A'\}$ denote the out-nodes from i and let $N_i^- = \{j \in N : (j, i) \in A'\}$ denote the in-nodes to i . Note that $(i, j) \in A' \iff (j, i) \in A'$, and therefore $N_i^+ = N_i^-$. We use the distinct notation for consistency with standard directed models.

The mathematical formulation of our model, which will be denoted (UTM-TT), is as follows.

$$\underset{z, x, f}{\text{minimize}} \quad \sum_{(i,j) \in A'} [t_{ij}(f_{ij})] \cdot f_{ij} \quad (1a)$$

$$\text{subject to} \quad \sum_{j \in N_i^+} x_{ij}^k - \sum_{j \in N_i^-} x_{ji}^k = \begin{cases} 1 & \text{if } i = O(k) \\ -1 & \text{if } i = D(k) \\ 0 & \text{otherwise} \end{cases} \quad \forall k \in K, \forall i \in N \quad (1b)$$

$$f_{ij} = \sum_{k \in K} d^k x_{ij}^k \quad \forall (i, j) \in A' \quad (1c)$$

$$x_{ij}^k, x_{ji}^k \leq z_{ij} \quad \forall (i, j) \in A, \forall k \in K \quad (1d)$$

$$\sum_{(i,j) \in A'} \ell_{ij} \cdot x_{ij}^k \leq L_k \quad \forall k \in K \quad (1e)$$

$$\sum_{(i,j) \in A} c_{ij} \cdot z_{ij} \leq B \quad (1f)$$

$$x_{ij}^k \in \{0, 1\} \quad \forall (i, j) \in A', \forall k \in K \quad (1g)$$

$$f_{ij} \geq 0 \quad \forall (i, j) \in A' \quad (1h)$$

$$z_{ij} \in \{0, 1\} \quad \forall (i, j) \in A \quad (1i)$$

Equation (1a) is the objective function; it measures the total travel time experienced by vehicles when routing on the given paths in the selected network considerate of congestion effects. Constraints (1b) are the flow conservation constraints which choose a path in the network from each origin to its corresponding destination. As long as travel times are positive, these constraints will always result in a simple path being chosen because the objective is to minimize total travel time. Constraints (1c) calculate the flow on each arc. Constraints (1d) are the design forcing constraints; they ensure that every arc used in an OD-pair's path is open. Constraints (1e) ensure that the time it takes each vehicle to traverse its path satisfies the corresponding static travel time limit. Constraint (1f) ensures the selected network's cost does not exceed the budget. Constraints (1g), (1h), and (1i) are the variable bound constraints.

3.2 Piecewise Linear Approximation

The BPR function is convex and therefore our objective function is amenable to piecewise linear (PWL) approximation techniques. For a general overview of using PWL functions for solving mixed integer non-linear

programs, see [23]. Some of the related works discussed in Section 2 employ PWL approximation techniques; we note two particularly relevant examples here. Firstly, Avella et al. [3] use PWL congestion functions within a network design problem, citing the benefit of providing a formulation that is suited to commercial mixed integer linear programming solvers. Secondly, Angelelli et. al [1, 2] present a piecewise linearization of the CSO model to capture flow-dependent travel times within the constraints. They analyze the quality of the approximation, and investigate various options for discretization technique.

To approximate the objective function with a PWL function, we introduce additional variables, constants, and parameters to the formulation. We follow the notation used in [2]. For each arc $(i, j) \in A'$:

- Total travel time variable σ_{ij} denotes the approximation of the total travel time experienced by traffic on the arc.
- Tunable parameter n_{ij} denotes the number of pieces used in the PWL approximation.
- Tunable parameter U_{ij} denotes the maximum amount of flow that may be assigned to the arc.
- Constant breakpoint b_{ij}^s for $s \in \{0, \dots, n_{ij}\}$ denotes the s -th breakpoint used in the PWL approximation. Necessarily $b_{ij}^0 = 0$ and $b_{ij}^{n_{ij}} = U_{ij}$. Usually we will use evenly spaced breakpoints.
- Constant $\Delta_{ij}^s := b_{ij}^s - b_{ij}^{s-1}$ for $s \in \{1, \dots, n_{ij}\}$ gives the length of the s -th interval.
- Arc-flow variable λ_{ij}^s for $s \in \{1, \dots, n_{ij}\}$ denotes the amount of flow assigned to the s -th interval.
- Constant C_{ij}^s for $s \in \{0, \dots, n_{ij}\}$ gives the value of $x \rightarrow x \cdot t_{ij}(x)$ evaluated at b_{ij}^s .

The resulting model, which will be denoted (UTM-TT-PWL), is as follows.

$$\begin{aligned} & \underset{z, x, f, \lambda, \sigma}{\text{minimize}} && \sum_{(i,j) \in A'} \sigma_{ij} && (2a) \end{aligned}$$

$$\text{subject to} \quad (1b), (1c), (1d), (1e), (1f), (1g), (1h), (1i)$$

$$\sum_{s \in \{1, \dots, n_{ij}\}} \lambda_{ij}^s = f_{ij} \quad \forall (i, j) \in A' \quad (2b)$$

$$\sigma_{ij} = \sum_{s \in \{1, \dots, n_{ij}\}} \left(\frac{C_{ij}^s - C_{ij}^{s-1}}{\Delta_{ij}^s} \right) \lambda_{ij}^s \quad \forall (i, j) \in A' \quad (2c)$$

$$0 \leq \lambda_{ij}^s \leq \Delta_{ij}^s \quad \forall (i, j) \in A', \forall s \in \{1, \dots, n_{ij}\} \quad (2d)$$

$$\sigma_{ij} \geq 0 \quad \forall (i, j) \in A' \quad (2e)$$

Equation (2a) is the objective function which sums the approximated total travel time experienced for each arc. Constraints (2b) decompose the arc-wise flow for the purposes of PWL approximation. Constraints (2c) compute the arc-wise contribution to the objective function using the PWL approximation. Constraints (2d) and (2e) are the additional variable bound constraints. As the congestion function is convex, the interval arc-flow variables will be assigned flow in the correct order; i.e. in any optimal solution $(z^*, x^*, f^*, \lambda^*, \sigma^*)$, if $(\lambda^*)_{ij}^s < \Delta_{ij}^s$ for $(i, j) \in A'$ and $s \in \{1, \dots, n_{ij}\}$, then $(\lambda^*)_{ij}^{s'} = 0$ for all $s' \in \{s+1, \dots, n_{ij}\}$.

4 Congestion Functions for UAVs

Congestion functions (a.k.a. volume-day, travel time, or link capacity functions) have been used since the 1950s for road traffic modeling. Various mathematical forms for congestion functions were proposed early on. See [8] for a survey of early congestion function proposals and [24, 2.2] for a summary of common mathematical requirements for congestion functions. The BPR function developed by the United States Bureau of Public Roads in the early 1960s [40] continues to be used by practitioners and in Operations Research studies today. The parameterized form of the BPR function, which maps flow to travel time, is

$$t(f) = t_{FF} \left(1 + \alpha \left(\frac{f}{\kappa} \right)^\beta \right) \quad (3)$$

where the parameters are t_{FF} - the road’s *free flow travel time*, κ - road “capacity”, and $\alpha \geq 0$, $\beta \geq 1$ which influence the functional form. The free flow travel time is the amount of time it takes a vehicle to traverse the road when there is no congestion. As discussed in Section 2, the “capacity” gives an estimate of the critical point where the volume of traffic begins to significantly impede travel time. The values for α and β recommended by the BPR engineers [40] were $\alpha := 0.15$ and $\beta := 4$, which were computed by a linear regression based on highway data. These parameter values remain a common choice in modeling today. Various researchers have worked to improve the parameter choices, often by calibrating the function to match observed traffic patterns in a specific region; see, for example, [9, 38, 54, 53].

Using observed traffic patterns to calibrate a congestion function for UAV traffic above a road will not be possible until data is gathered from early UTM implementations. Accordingly, we take a more theoretical approach and use the common choice for parameters (α, β) to take on values $(0.15, 4)$. By increasingly penalizing flows over capacity κ , this form of congestion function captures the trade-off between relieving a highly congested arc and opening a shorter route for a UAV assigned a long path. Since our modeling objective is to provide a realistic high level assessment of network capability, we need to choose an appropriate capacity value κ to reflect the critical point for the trade-off. Unlike automobile traffic, safety regulations require UAVs to maintain strict separation standards, indicating the vertical and lateral distance required between UAVs in the air. In essence, each UAV is modeled as occupying an ellipsoid of 3D volume, and intersections of these ellipsoids are forbidden. Assuming traffic respects these standards, we will use simulation to impute capacity numbers which penalize the flow values that would result in UAVs waiting long times for clear airspace.

Before describing the simulation, we discuss the remaining BPR function parameter - free-flow travel time t_{FF} . For automobile traffic, this value is usually determined by the speed limit and road characteristics because most vehicles maintain a speed near the posted (or socially accepted) limit. Road vehicles may differ in acceleration/deceleration speeds and ability to maneuver to change lanes. Models considering these differences often modify the total flow on the road using the concept of an “equivalent passenger car unit” (ECU), which specifies a vehicle type’s contribution to total flow relative to a passenger car. Less agile

vehicles (ex. trucks) may be assigned an ECU above one and more agile vehicles (ex. motorcycles) may be assigned an ECU below one. Whether specific speed limits will apply within UTMs is, as of yet, undetermined. It is expected that UAVs operating within the managed airspace will be required to meet minimum operating standards. Above that standard, however, UAV speeds may vary greatly depending on the vehicle type, as well as the payload onboard. In this paper, we make the simplifying assumption that all UAVs travel at the same constant speed and therefore use free flow travel times directly proportional to road length.

As we are restricting to flight above roads in segmented layers of vertical airspace, there are only two UAV separation standards to consider: horizontal and vertical. Along a single segment of road, these separation standards have straightforward interpretations. The *horizontal distance* (H) indicates that a UAV entering a section of road at a fixed altitude must wait until the UAV “ahead of it” is H units of distance away before beginning its traversal. The UAV must continue to maintain a distance of H from the preceding vehicle while it traverses the road. The next UAV to enter must do the same. The *vertical distance* (V) indicates the distance between the vertically segmented layers of airspace. The number of layers of airspace available is given by the flight ceiling less the flight floor divided by V . Each of these layers must be assigned a fixed direction of traversal. A typical UAV flight ceiling is 400ft; specific values are dictated by the corresponding regulatory authority. For example, flight ceilings for UAVs in the United States can be obtained from the FAA Unmanned Aerial System Facility Map (UASFM) grid cells [21]. A typical UAV flight floor is 100ft; specific values depend on the underlying ground terrain.

The horizontal distance, vertical distance, flight ceiling, flight floor, average UAV speed, and a *lane change penalty*, can be used to determine an approximate capacity for the BPR function. First, the flight ceiling, flight floor and vertical distance are used to determine the *number of lanes* available for traversal (in a fixed direction). Then, the horizontal distance and average UAV speed are used to determine the *clearance time*: how long it takes a UAV to clear the horizontal distance. We simulate f flights per hour arriving according to a Poisson process and then determine the average waiting plus traversal time of a flight. When there is only one lane, this is a simple queuing process. However, when there are multiple lanes, we assume that the vehicles arrive at a given lane (i.e. vertical layer) uniformly at random and will traverse at that layer unless its queue is too long. A *lane change penalty* is added to the travel time of a UAV when it changes lane to model the time it takes to switch altitude. These lane changes are unique to UAV flight; car traffic on a highway does not experience the same disutility in terms of speed and battery drain when switching lanes. For $f \in \{0, 1, \dots\}$ flights per hour, we can obtain a simulated travel time $t(f)$. We find the value of f for which the travel time $t(f)$ is 1.15 times the free flow time t_{FF} and use it as the BPR capacity for the given parameters. Pseudocode for one simulation is provided in Algorithm 1; the simulation can be repeated sufficiently many times to account for the random Poisson arrival times and random arrival lanes.

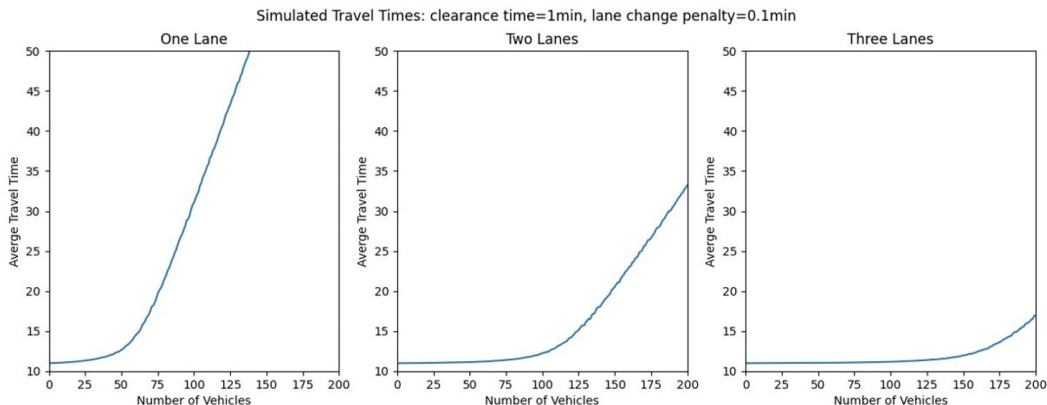


Figure 1: Results of simulated travel times for parameters as indicated and a free flow travel time of 11min.

In the next section, numerical experiments are performed in which we fix arc capacities “consistent with the network instances being solved”; the simulation here provides a simple link back to the physical constraints imposed by separation standards, available lanes, and lane change costs. This link is important as UTM providers determine the “rules of the road” both in early UTM development and at maturity. The congestion function translates these rules into a simple travel time penalty that can be used to evaluate network quality.

5 Numerical Experiments

This section first discusses a process for arriving at practical test instances. By following the process, we generate test instances for Chicago City Center. We first experiment with a specific instance class, using the network design model to analyze design tradeoffs. Then, we investigate whether these managerial insights hold or differ for other instance classes. Lastly, we discuss model runtimes and accuracy. All models described in this section were implemented using Gurobi (v10.0.1) via the Python API and instances were run on an Apple M1 Pro Version macOS Ventura 13.2 with 16GB of memory. Each optimization problem solved is given a runtime limit of three hours.

5.1 Generating and Sizing Instances

The first step in generating an instance is obtaining an appropriate base road network. Although we tested our models with the standard transportation research networks from [22], we found these networks weren’t ideal for extending to UAV flight. The geographic regions covered were often too large and the (aggregate) roads too sparse because they represented major arteries. In the early stages of UTM implementation, a UTM provider may make a strategic decision regarding which area of a city to service first because opening up airspace in an entire metropolitan area could be both costly and risky. We expect the service area covered

Algorithm 1 Simulating Travel Time as a Function of Number of Flights

Inputs:

- f : the number of flights per hour
- n_{lanes} : flight ceiling less flight floor, divided by $2V$
- t_{clear} : time it takes a vehicle to traverse the horizontal separation standard
- t_{switch} : lane change penalty
- t_{FF} : free flow travel time

Outputs: $t(f)$: the simulated average travel time when f flights arrive in an hour

$t_{\text{entry}} \leftarrow [0.0 \text{ for } i \in \{1, \dots, f\}]$

$t_{\text{current}} \leftarrow 0.0$

$t_{\ell} \leftarrow [0.0 \text{ for } i \in \{1, \dots, n_{\text{lanes}}\}]$

for $i \in \{1, \dots, f\}$ **do**

t_{wait} is a randomly drawn wait time for a Poisson process with f arrivals per hour

$t_{\text{current}} \leftarrow t_{\text{current}} + t_{\text{wait}}$

a is drawn uniformly at random from $\{1, \dots, n_{\text{lanes}}\}$

if $(t_{\ell}[a] \leq t_{\text{current}})$ **then**

$t_{\ell}[a] = t_{\text{current}} + t_{\text{clear}}$

$t_{\text{entry}}[i] \leftarrow t_{\ell}[a] - t_{\text{current}}$

else if $(t_{\ell}[a] \leq t_{\text{current}} + t_{\text{switch}})$ **then**

$t_{\ell}[a] = t_{\ell}[a] + t_{\text{clear}}$

$t_{\text{entry}}[i] = t_{\ell}[a] - t_{\text{current}}$

else

$a^* = \text{argmin}\{t_{\ell}(j) \text{ for } j \in \{1, \dots, n_{\text{lanes}}\}\}$

if $(a^* == a)$ or $(t_{\ell}[a^*] + t_{\text{switch}} > t_{\ell}[a])$ **then**

$t_{\ell}[a] = t_{\ell}[a] + t_{\text{clear}}$

$t_{\text{entry}}[i] = t_{\ell}[a] - t_{\text{current}}$

else

if $(t_{\ell}[a^*] < t_{\text{current}})$ **then**

$t_{\ell}[a^*] \leftarrow t_{\text{current}} + t_{\text{clear}} + t_{\text{switch}}$

else

$t_{\ell}[a^*] \leftarrow t_{\ell}[a^*] + t_{\text{clear}} + t_{\text{switch}}$

end if

$t_{\text{entry}}[i] \leftarrow t_{\ell}[a^*] - t_{\text{current}}$

end if

end if

end for

$t(f) \leftarrow t_{\text{FF}} - t_{\text{clear}} + \text{avg}(t_{\text{entry}})$

to align with delivery UAV flight capabilities and accordingly suggest a test instance whose geographic size is around 10-15km by 10-15km. In an area of this size, a “typical delivery UAV” can fly point-to-point from most origins to most destinations and back without recharging. Operations Research work relating to UAV flight assumes similar ranges. For example, the authors of [25] study locating UAV recharging stations and consider two possible ranges of five miles and ten miles. However, they remark that determining a realistic flight range for delivery UAVs was a challenge.

Having narrowed in on a geographic size for our test instances, we turn to the problem of identifying candidate roads. OSMNX [7] is a Python library that can download open source map data in graph format. We use it to obtain the underlying road network and filter based on the “highway” tag to restrict the set of roads considered. For example, the base network for the Chicago instance was obtained by restricting candidate roads to those with a “highway” tag of “motorway”, “motorway_link”, “primary”, or “secondary”. According to the OSM guidelines for tagging highways in Illinois [15] (and the US more generally [16]), these candidate roads include interstate highways and expressways with similar characteristics, main surface streets in urban areas, and the roads in urban areas that extend long distances through multiple neighbourhoods and have high volumes. These roads are the best candidates for UAV traffic because flying down major streets provides access to the busiest areas of the city and is preferential to flying on residential roads from a privacy and societal acceptance perspective. Note that the choice of “highway” tag may depend on the features of the area being mapped; in the suburbs, it may be appropriate to use additional tags. Besides providing the topology - i.e. the relationships between nodes (intersections) and arcs (roads) - the OSM data also provides street geometries by representing roads with a sequence of line segments that capture any turns or bends. Therefore, we obtain arcs lengths that reflect the true driving distance along the road, and thus the horizontal distance a UAV will travel if it flies above the road.

The raw driving networks provided by OSMNX are not suited for direct input into the UTM network design model; some modifications are prudent. The driving network represents connections according to feasible (possibly one-way) driving paths which may represent specific on-the-ground or car-minded infrastructure. Some of these features are not used by UAVs in the same way. If a road passes over another (ex. a highway overpass with no on-ramp), then the road network will not contain a node representing that intersection. However, a UAV flying above the roads is not restricted in the same way. Major road intersections are often managed by segmenting traffic using turning lanes or curving ramps. However, the intersection may be managed for UAV traffic by creating vertical separation, rather than horizontal because the lane created to segment car traffic may not be appropriate for UAVs, which have greater separation requirements. The capacity provided by major expressways may correspond to separated roads - for example, when there are express lanes and local/collector lanes for a highway. These roads are represented by separate, but similar, arcs in the raw driving network. However, a UTM network may not use both these roads because simul-

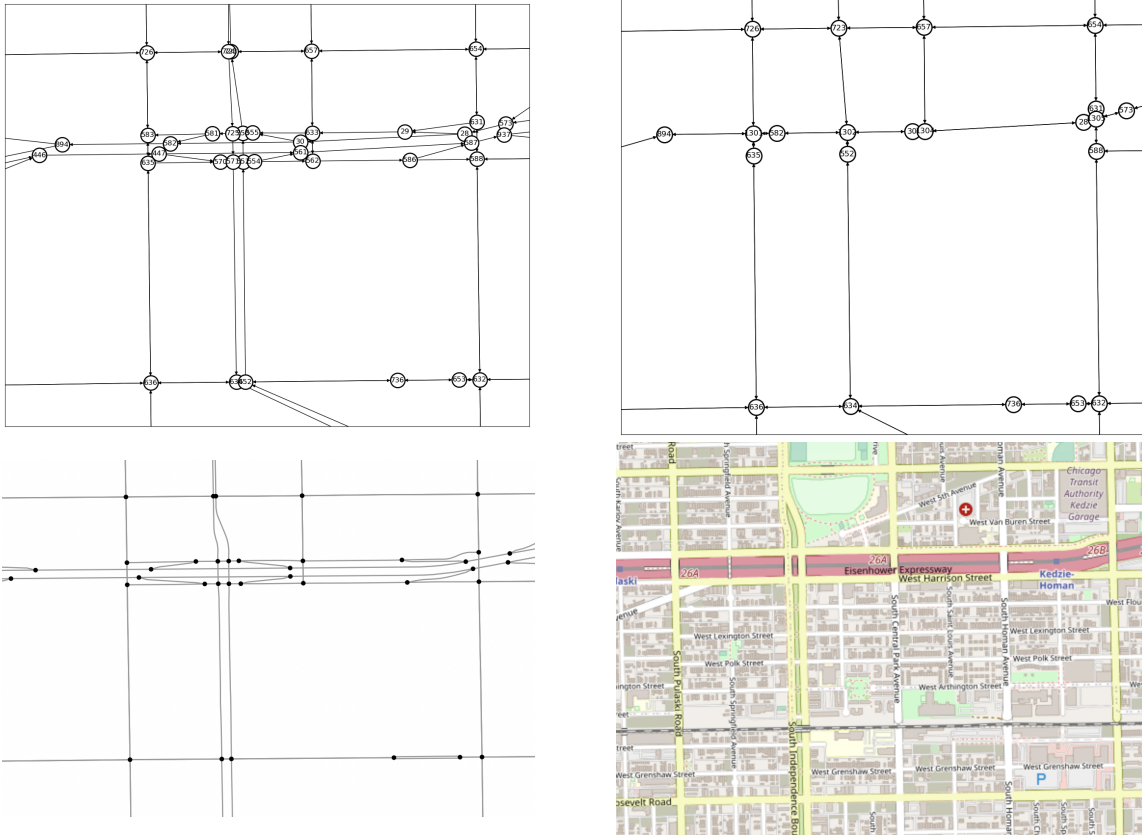


Figure 2: A section of the road network before (top left) and after (top right) simplification, as well as the geometry of the underlying roads (bottom left), and map of corresponding city streets (bottom right).

taneous flights would violate horizontal separation standards. In light of these differences, we perform a sequence of simplifying operations to adjust the raw road network. An example of these operations is given in Figure 2 where three main modifications were made. The highway overpasses for all roads running north to south are modified into intersections. The two-lane road represented by two parallel paths running north to south is replaced with one path. The highway running from east to west is replaced with one path. An example of simplifying an intersection is provided in Figure 3. The simplifications made preserve distances as much as possible. A complete simplification is performed to arrive at the base network in Section 5.2.

The second step in generating an instance is determining the origins and destination of flights, together with the affiliated flight volumes and path-length limits. Publicly available information about current UAV flight patterns is scarce and we therefore propose to randomly generate instances according to probable demand patterns, where a *demand pattern* specifies how flights requests are spatially distributed. We consider three different demand patterns. In the *space-uniform* (SU) demand pattern, pairs of coordinates are drawn at random within the bounding box of the area. The origin is the nearest node (geographically) to the first




Figure 3: An intersection before (left) and after (center) simplification, as well as the geometry of the underlying roads (right).




Figure 4: The SU (left) and RNP (center) demand pattern with both origins and destinations indicated by triangles. The WC (right) demand pattern with the warehouse indicated by a green triangle and the delivery points indicated by red triangles.

coordinate pair while the destination is the nearest node to the second coordinate pair. This demand pattern provides the largest geographic coverage, and may arise when the area considered is homogeneous (e.g. in terms of population features). In the *random-node-pair* (RNP) demand pattern, the origin and destination are sampled at random from all nodes. This method has the benefit of concentrating demand in areas with many major streets (e.g. the downtown core), which is natural. However, the method is sensitive to how OSMNX chooses to place nodes, as well as the post-processing performed. In the *warehouse-customer* (WC) demand pattern, one node is presumed to be the location of a warehouse within the city. The origin of each OD-pair is the warehouse and the destination is a randomly selected node representing a location where a customer may request a delivery. This demand pattern is natural if a specific company intends to offer UAV delivery service, or if there is a fixed UAV launch point within the city. The demand patterns are illustrated in Figure 4.

Although OD-pairs will be drawn at random to generate an instance, the model itself is solved for a fixed set of OD-pairs. The underlying assumption is that the random pairs are adequate to achieve good coverage

(which may be analyzed), or represent a pattern that could arise when a specific city knows its needs. In Section 6, we discuss research directions related to stochastic demand scenarios or optimizing for the amount of demand served. Now, if we are generating an instance with k OD-pairs, then the origins and destinations of the OD-list are drawn without replacement from all node pairs that fit the demand pattern. The demand for each OD-pair is drawn uniformly at random from $\{1, 2, \dots, f_{max}\}$, where f_{max} is some appropriately determined upper bound. The expected total number of flight requests is $\frac{k \cdot f_{max} \cdot (f_{max} + 1)}{2}$.

For each OD-pair, we require a corresponding path-length limit. As in [28], we consider path-length deviations which impute sets of path-length limits. An instance is said to be solved for *path-length deviation* $\delta\%$ if the path length limit for each OD-pair is equal to $(100 + \delta)\%$ times the length of the shortest path (on the road network) between origin and destination. Since we chose an area size in which UAVs can fly between any two points while respecting battery life, path length deviations up to 25% yield limits within a reasonable range. For close together OD-pairs, the limit may seem overly restrictive. However, unduly long routes may not be acceptable to UAV pilots because long flight times introduce unnecessary air risk. Note that we measure path-length deviation with respect to the shortest path on the existing road network, which is almost always longer than the Euclidean shortest path, and thus introduces further deviation from the greedy choice.

Arc congestion functions and a budget are the last two pieces of data required to complete an instance. As per Section 4, we will use the BPR function with standard α and β . The road length provided by OSMNX can be used as both the mapping cost for a road segment, as well as the free flow travel time (following the constant speed assumption). Hence, it remains to determine an appropriate budget and arc capacities. As discussed in Section 4, known separation standards, drone speeds and flight ceilings may be used to determine arc capacities. These standards are still in development, and are not yet publicly available. However, it should be readily apparent that the budget, path-length limits, demand volume, and arc capacities must be consistent and aligned to obtain practically interesting instances. For example, if every network that provides each OD-pair a path within its path length limit (1e) is above the budget (1f), then the model is infeasible. Similarly, the total demand, budget, and arc capacities must be aligned so that there exists a network within budget for which most arcs don't have flow "too far" above capacity.

Pre-analysis can ensure that the inputs to the (UTM-TT) model are consistent and aligned. The idea is to set parameters in succession to produce reasonable instances. We choose to fix demand first, and impute appropriate arc capacities last. Similar techniques may back out consistent inputs by fixing budget or arc capacity parameters first. Now, given a base road network and OD-list together with a path length deviation, we can compute a minimum cost network satisfying the corresponding path length limits by solving

$$\underset{z, x}{\text{minimize}} \quad \sum_{(i,j) \in A} c_{ij} z_{ij} \tag{4a}$$

subject to (1b), (1d), (1e), (1g), (1i).

Denote this model by (UTM-Cost). Its optimal value provides the cost of the cheapest network which connects every OD-pair by a short enough path. The corresponding (UTM-TT) model is feasible only if the budget is at least the optimal value of (UTM-Cost). After determining a reasonable budget, we solve

$$\begin{aligned} & \underset{z, x}{\text{minimize}} && \sum_{(i,j) \in A} \ell_{ij} f_{ij} && (5a) \\ & \text{subject to} && (1b), (1c), (1d), (1e), (1f), (1g), (1h), (1i). \end{aligned}$$

Denote this model by (UTM-Dist). The model minimizes the total distance traveled when flights are routed with no regard for congestion effects. By analyzing the arc usage in the model results, we can determine the capacity level at which arcs start to be oversubscribed and use it to determine interesting arc capacities. Note that (UTM-Cost) and (UTM-Dist) provide insights into candidate networks and their capability. For example, when the total flight volume is low and therefore congestion effects are negligible, the (UTM-Dist) model results mimic the (UTM-TT) model. Varying the flight volume, the (UTM-Dist) model gives an indication of the throughput level at which bottlenecks occur and congestion effects should be modeled.

To summarize, instance generation requires the following inputs: **city**, **demand pattern**, number of OD pairs **k**, max number of flights between an OD-pair **f_{max}**, and path length deviation δ . The instance generation procedure is as follows.

1. Obtain base road network for **city** from OSMNX.
2. Generate **k** OD-pairs according to the **demand pattern**. Assign each OD-pair a demand uniformly at random from $\{1, \dots, \mathbf{f}_{\max}\}$.
3. Calculate path length limits corresponding to path length deviation δ .
4. Use the (UTM-Cost) model to determine an appropriate budget range.
5. Use the (UTM-Dist) model to determine an appropriate arc capacity range.

Instances with the same city, demand pattern and number of OD pairs will be said to belong to the same *instance class*.

5.2 Case Study of Chicago City Center RNP Fifty OD-Pair Instance Class

We analyze UTM networks for an area of Chicago consisting of the downtown core and adjacent neighbourhoods and focus first on instances with fifty OD-pairs drawn according to the RNP demand pattern. We measure the impact of budget, congestion, and path length limits on the network design across five instances drawn from the instance class. Table 1 describes the abbreviations and conventions we adopt to report model results. For runtime analysis, see Section 5.4. Here we report only timeouts and PWL approximation parameters as applicable because they inform the validity of the analysis.

Heading	Description	Applies To
avg	average taken over five instances in the instance class	all
std	standard deviation taken over five instances in the instance class	all
dmd	demand pattern considered (SU, RNP, or WC)	all
pld	path length deviation considered	all
num ods	number of OD-pairs in the instance class	all
bd	indicates the budget deviation from the minimum budget from (UTM-Cost)	(UTM-Dist), (UTM-TT-PWL)
cap	arc capacity considered	(UTM-TT-PWL)
cost	total length of roads in the chosen network in kilometers	all
num arcs	number of arcs in the chosen network	all
avg pld	weighted average of the realized path length deviation over all OD-pairs	(UTM-Dist), (UTM-TT-PWL)
num on sp	number of OD-pairs provided a path equal to the length of their shortest path	(UTM-Dist), (UTM-TT-PWL)
max flow	amount of demand on the arc which carries the most demand	(UTM-Dist), (UTM-TT-PWL)
travel time	indicates the objective function value (1a)	(UTM-TT-PWL)

Table 1: Summary statistics reported for UTM network design models.

The area of Chicago analyzed is bounded by: the I50 on the west, the Lake Michigan harbour on the east, the I64 on the north, and 47th Street on the south. The precise bounding box has corners with (longitude, latitude) given by $(-87.581972, 41.806857)$ and $(-87.750543, 41.914249)$. The resulting area is 14km by 12km. The region considered and initial set of roads downloaded are illustrated in Figure 5. The simplifications described in Section 5.1 are performed on the base network. Other straightforward simplifications are made to reduce the problem size: degree two vertices are suppressed and some nearby parallel roads are removed from the downtown core to eliminate similar solutions. After all simplification, the raw driving network with 1283 nodes and 2735 (directed) arcs is reduced to a graph with 500 nodes and 863 edges. The final network is illustrated in Figure 6.

To assess appropriate budget ranges, (UTM-Cost) is solved for all five instances and aggregated results are reported in Table 2. There are a total of 424.8km of roads that may be mapped. For path length deviations of 10% and 20%, the minimum budget network for all instances is around 100km. The cheapest network that routes every OD-pair on a shortest path between its origin and destination is calculated using path length deviation 0%. This optimal value gives the maximum sensible budget for the (UTM-Dist) model, which can do no better than shortest path routing. The (UTM-TT) model budget could reasonably be set higher as mapping more roads may alleviate congestion. For the fifty OD-pair RNP instance class, however, we found the rate of objective value improvement diminished as budget increased and therefore present results to a maximum budget of 150% the minimum cost network. In particular, we first vary the budget in increments of about 10km by considering budgets equal to 110%, 120%, . . . , 150% of the optimal value of the corresponding (UTM-Cost) model. These instances will be said to have *budget deviation* 10%, 20%, . . . , 50% respectively.




Figure 5: A map of the area of Chicago considered on the left, and the corresponding major roads obtained from OSMNX on the right. The data was obtained on 7 November 2022.




Figure 6: Final Chicago Road Network after UAV-specific adjustments and simplification.

		cost		num arcs	
dmd	pld	avg	std	avg	std
rnp	0	170.85	11.73	394.60	12.58
rnp	10	106.30	3.25	252.80	12.40
rnp	20	92.26	3.79	227.20	6.53

Table 2: Minimum cost networks satisfying path length constraints for fifty OD-pair RNP instances.

		cost		num arcs		travel dist		avg pld		num on sp		max flow	
pld	bd	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
10	10	116.89	3.55	273.80	10.38	1774.52	232.86	1.65	0.25	14.40	1.82	36.80	1.79
10	20	127.50	3.91	296.20	15.02	1755.44	229.09	0.71	0.13	21.80	3.42	33.80	2.17
10	30	138.09	4.25	318.60	12.84	1748.12	228.79	0.33	0.14	25.40	4.93	35.00	4.00
10	40	148.69	4.60	342.60	17.70	1744.13	228.04	0.14	0.09	30.80	6.46	33.20	5.72
10	50	158.76	5.17	364.40	18.01	1742.32	227.63	0.04	0.04	39.00	9.92	31.80	4.32
20	10	101.47	4.17	246.00	10.00	1814.97	233.46	3.66	0.54	9.00	2.12	44.40	9.40
20	20	110.67	4.53	265.00	12.59	1782.24	230.45	2.05	0.37	12.60	2.61	34.80	4.02
20	30	119.90	4.90	280.60	13.22	1762.80	227.62	1.15	0.23	19.00	2.83	33.60	3.71
20	40	129.08	5.27	301.40	10.11	1753.11	226.833	1.64	0.22	21.80	3.56	33.80	2.17
20	50	138.30	5.64	319.60	11.97	1747.85	227.15	0.33	0.14	25.40	5.03	35.20	4.09

Table 3: Minimum travel distance with path length and budget constraints for fifty OD-pair RNP instances.

To assess appropriate arc capacities, (UTM-Dist) is solved for all five instances and aggregated results are reported in Table 3. As the budget increases, the number of arcs included in the network and the number of OD-pairs routed on their shortest path increases while the average path length deviation decreases. These relationships don't just hold on aggregate; each can be observed within a given instance. Such a simple increasing (or decreasing) relationship does not hold for the maximum arc flow; one possible explanation is that as more arcs are mapped, more OD-pairs are able to connect to a shared short route through an area. Based on the (UTM-Dist) results, we choose to analyze the (UTM-TT) model when the BPR arc capacities are equal to 30. At this level, we expect incorporating congestion should force the model away from overloading any arc without being too restrictive. Note that using Section 4, we can back out possible sets of corresponding separation standards and lane numbers.

First, we evaluate the impact of budget on networks designed by the (UTM-TT) model. To do so, (UTM-TT-PWL) is solved with a maximum arc capacity of 50 and fifty pieces in the PWL congestion functions. Since the flow on each arc is restricted to be integer (and provided the true optimal solution does not route over 50 units of demand on any one arc), (UTM-TT-PWL) models (UTM-TT) exactly. Except for two instances with budget deviation 10% and path length deviation 20% (with mipgaps of 1.3% and 0.8% at three hours), the underlying optimization problems were solved to optimality. The results are summarized in Table 4. Observe that the improvement in total travel time as a function of budget is more pronounced

		cost		num arcs		travel dist		avg pld		num on sp		max flow		travel time	
pld	bd	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
10	10	116.91	3.57	276.80	8.93	1785.30	233.29	2.03	0.32	13.20	3.11	29.00	3.39	1809.74	246.34
10	20	127.52	3.91	298.20	9.71	1763.44	230.34	1.14	0.21	17.00	4.64	27.00	2.74	1781.54	239.10
10	30	138.15	4.19	321.60	15.50	1754.59	230.16	0.64	0.22	23.00	5.39	26.00	1.22	1768.61	236.61
10	40	148.78	4.54	345.40	17.66	1749.16	228.92	0.43	0.08	26.80	5.89	26.00	2.00	1761.74	235.46
10	50	159.37	4.94	371.00	21.05	1746.23	228.63	0.26	0.08	29.80	5.76	26.00	1.41	1757.83	234.44
20	10	101.47	4.17	247.00	11.73	1830.34	235.40	4.39	1.03	8.80	2.68	31.80	3.27	1871.47	250.64
20	20	110.68	4.55	268.40	11.59	1790.33	230.98	2.37	0.42	11.20	2.28	30.00	3.32	1818.54	242.89
20	30	119.92	4.93	285.80	9.81	1772.20	228.90	1.56	0.33	14.80	4.15	27.40	4.51	1792.24	237.77
20	40	129.13	5.32	303.40	10.78	1760.90	228.11	1.00	0.19	19.20	3.49	26.40	1.67	1776.38	235.07
20	50	138.38	5.69	323.00	10.98	1754.12	228.23	0.62	0.13	22.00	5.61	25.60	1.34	1767.41	234.35

Table 4: Minimum travel time with path length and budget constraints for 50 OD-pair RNP instances

for low budget deviations. That is to say, the value obtained from slightly increasing the mapping budget is not linear and the largest impact occurs when raising the budget slightly above the cost of the cheapest feasible network. Another interesting take-away from these instances is that the travel time minimizing model is able to find a network that has small total travel distance and simultaneously avoids arcs that carry large amounts of flow. Comparing the average path length deviation and number on shortest path between (UTM-Dist) in Table 3 and (UTM-TT) in Table 4 for the same budget and path length deviation, we see that the total travel time is improved at the expense of routing fewer OD-pairs on their shortest paths. This should be expected as shortest paths for different OD-pairs may share a short route through a specific area, and congestion may be relieved by choosing an alternative path for one such OD-pair.

Next, we evaluate the impact of path length deviation on the network designed by the (UTM-TT) model. We choose our budget based on the minimum cost network for path length deviation of 10%; in particular, we use a budget of 110% of the corresponding minimum cost network. We continue to use an arc capacity of 30, max arc capacity of 50 and fifty pieces in the PWL congestion functions. For this budget with path length deviation 5%, the resulting model is infeasible for all five instances. The remainder of the problems are solved to optimality and the results are summarized in Table 5. Observe that at this budget level, the networks designed share very similar summary statistics; this observation hold true even analyzing within an instance. It's interesting to note that the average path length deviation decreases as the allowable path length deviation increases. The result may seem counter-intuitive when phrased this way. However it is natural when considering the increasingly less restrictive optimization problems being solved; the ability to deviate further from the shortest paths to the disadvantage of certain OD-pairs can only lead to system improvement which comes from routing more distant OD pairs on their shortest path.

Lastly we evaluate the impact of arc capacity on the networks designed by the (UTM-TT) model. Setting the budget to 110% of the minimum and the path length deviation to 10%, (UTM-TT-PWL) is solved for

pld	cost		num arcs		travel dist		avg pld		num on sp		max flow		travel time	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
10	116.91	3.57	276.80	8.93	1785.30	233.29	2.03	0.32	13.20	3.11	29.00	3.39	1809.74	246.34
15	116.92	3.59	277.00	11.20	1780.25	235.09	1.86	0.32	13.20	3.56	28.20	3.35	1801.97	246.70
20	116.91	3.57	280.60	12.42	1779.58	236.84	1.80	0.45	13.20	3.96	28.20	3.35	1801.12	246.78
25	116.91	3.57	279.80	15.16	1778.95	236.76	1.83	0.45	13.40	3.21	29.20	3.96	1800.78	247.03

Table 5: Minimum travel time fixing budget and varying path length limit for 50 OD-pair RNP instances.

cap	cost		num arcs		travel dist		avg pld		num on sp		max flow		travel time	
	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
25	116.92	3.57	275.00	11.05	1790.96	237.67	2.36	0.40	12.00	2.12	28.20	3.35	1833.47	258.17
30	116.91	3.57	276.80	8.93	1785.30	233.29	2.03	0.32	13.20	3.11	29.00	3.39	1809.74	246.34
35	116.91	3.57	274.80	7.76	1780.36	232.27	1.86	0.22	13.80	2.95	30.60	3.78	1797.32	240.01
40	116.91	3.59	275.20	11.19	1776.70	232.63	1.75	0.16	14.40	2.51	32.00	3.00	1789.14	236.79
45	116.92	3.59	275.20	11.88	1775.73	232.60	1.71	0.18	14.00	1.87	32.80	3.63	1784.39	235.22

Table 6: Minimum travel time fixing budget and varying arc capacity for 50 OD-pair RNP instances.

arc capacities of 25, 30, . . . , 45. We use a maximum arc capacity of 50 and fifty pieces in the PWL congestion functions. All except three instances solved to optimality within three hours; of the instances that timed out, the maximum MIP gap was 0.2%. The results are summarized in Table 6. The results show that for lower arc capacities, the solution deviates from the minimum travel distance solution to route more flow on congestion-avoiding paths. One natural question to ask here is: are the different solutions a result of choosing different networks at different arc capacity levels, or simply representative of a different choice of routing? Comparing the arc capacity 25 solutions to the arc capacity 35 solutions, there is an average of 29.2 different arcs in the lower capacity solution. Comparing the arc capacity 25 solutions to the arc capacity 45 solutions, there is an average of 37.0 different arcs in the lower capacity solution. Three of the instances select the same network for arc capacity 40 and 45, showing that the arc capacity impact on network designs levels off when congestion is low.

Lastly, we narrow in on one instance and compare the networks designed. Figure 7 includes an image of the nodes appearing in some OD-pair for the instance (left), and depicts the minimum cost network (middle) and minimum distance network for budget deviation 10% (right). The networks chosen are somewhat similar - with more budget small adjustments can be made to route more OD-pairs on shorter paths. Figure 8 depicts the solutions of the (UTM-TT) model with arc capacity 30, max capacity 50 and budget deviations 10% (left), 20% (center), and 30% (right). The minimum distance and minimum travel time network choices for budget deviation 10% differ as expected; when congestion is considered, more arcs in the downtown core area are opened as opposed to outer roads. As the budget increases within the (UTM-TT) model, the larger budget allows the model to both alleviate oversubscribed arcs and open additional outer roads used only by a few OD-pairs. Figure 9 depicts the solutions of the (UTM-TT) model with budget deviation 10%, max




Figure 7: One instance demand pattern and minimum cost and minimum distance networks.




Figure 8: One instance minimum travel time networks, varying budget.

capacity 50 and arc capacities 25 (left), 35 (center), and 45 (right). Again we see that the importance placed on congestion (via arc capacity choice) impacts the trade-off between mapping underused outer roads and relieving congestion on city center roads.

Many of the networks chosen by the models have a topology consisting of a backbone grid in the central region complemented by access paths to outlying nodes. This topology arises in other network design problems and there are two driving forces prompting it to arise here. Firstly, absent congestion effects, the path length limits are sufficient to induce multiple ways to cross the backbone region. Secondly, considering congestion effects leads to opening alternative routes that divert flow from shared short paths. A cursory glance at the networks designed show they have good survivability. Roughly, this means that there are multiple paths from origin to destination and therefore failure of an arc or node does not result in complete network failure. Determining survivability in the UTM context also requires evaluating the length of the paths that remain after a node or arc is removed; see Section 6 for further discussion.




Figure 9: One instance minimum travel time networks, varying arc capacity.

5.3 Insights from Other Chicago City Center Demand Patterns

The fifty OD pair SU instance class results shares many similarities with the observations made for the RNP fifty OD pair instance. The specific results of running this model are provided in Section 6.1. Here we highlight the major differences observed in the results.

This fifty OD pair WC instance class diverges from the RNP and SU instances. To analyze the WC demand pattern, we select a warehouse location south-west of the city centre, near to the Amazon warehouse’s location. The main difference is that the oversubscribed roads are evident before running any model. Congestion effects near the warehouse will dominate.

The one hundred OD pair RNP instance class is similar to the fifty OD pair instance. Here we focus

instance params	cost (km)		num arcs	
	avg	std	avg	std
pld (%)				
0	233.32	17.83	538.20	31.22
10	145.53	8.74	351.40	12.66
20	122.16	6.10	308.60	13.94

Table 7: Minimum cost networks satisfying path length limits for 100 OD-pair RNP instances.

5.4 Analysis of Model Runtime and PWL Accuracy

We analyze the runtime of the (UTM-TT-PWL) model, as well as the quality of the piecewise linearization in approximating (UTM-TT). Table 8 describes the additional abbreviations and conventions we adopt to report model results.

First we briefly discuss runtimes for the (UTM-Cost) and (UTM-Dist) models, which are summarized in Table 9 and ???. Both of these models are mixed integer linear programs. We found that the (UTM-Cost) model may be difficult to solve, especially when the path length deviation is large. Our implementation

Heading	Description
nto	the number of instances that timed due to the 3 hour runtime limit
runtime	runtime (wall-clock time) in seconds, obtained from Gurobi
work	Gurobi-specific parameter providing a deterministic measure of the work required to solve the problem
mipgap	absolute difference between the incumbent and dual bound, divided by the incumbent objective value, obtained from Gurobi

Table 8: Runtime statistics reported for UTM network design models.

			runtime		work		mipgap		nto
num ods	dmd	pld	avg	std	avg	std	avg	std	
50	rnp	0	0.01	0.00	0.01	0.00	0.00	0.00	0
50	rnp	10	196.54	213.04	454.76	499.24	0.00	0.00	0
50	rnp	20	3952.38	2005.58	11382.01	6046.72	0.00	0.00	0
50	su	0	0.01	0.00	0.01	0.00	0.00	0.00	0
50	su	10	713.42	528.22	1416.49	1571.98	0.00	0.00	0
50	su	20	9242.41	3959.38	15021.44	11081.29	1.96	1.28	4
50	wc	0	0.01	0.00	0.01	0.00	0.00	0.00	0
50	wc	10	10.26	10.46	18.28	21.19	0.00	0.00	0
50	wc	20	1005.14	770.41	2632.12	2146.32	0.00	0.00	0
100	rnp	0	0.01	0.00	0.01	0.00	0.00	0.00	0
100	rnp	10	8044.23	3863.89	13137.38	9728.25	0.51	0.85	2
100	rnp	20	11287.31	404.67	9420.42	7657.36	4.16	1.51	5

Table 9: Runtimes for (UTM-Cost) model.

applied the standard technique of removing the x_{ij}^k variables for all arcs $(i, j) \in A'$ and OD-pairs $k \in K$ for which every path from O_k to D_k that includes (i, j) has length greater than L_k . This pre-processing step is most effective when the path length deviation is small allowing a large number of variables to be eliminated.

6 Conclusion

In this work we introduced the first UAV network design model for UTM providers. The model differs from existing UAV research in that it considers the 3D nature of flights through a congestion function, assumes UAVs fly over existing streets rather than via point-to-point flight, and explicitly evaluates routings which are designed in respect of regulatory constraints and airspace design. We used simulation to develop BPR congestion function parameters consistent with routing UAVs in 3D space in respect of separation standards and along multiple corridors according to altitude. While we used the models to design new networks, the models and methodology developed in this work apply in a network expansion context. There are many further operational research problems that arise from this point of view and in this domain. Firstly, the

model studied herein may be extended to capture survivability needs and demand uncertainty. The networks designed in the case study often contained more than one short enough route from an origin to a destination; however, actively modeling multiple paths as a requirement may lead to alternative topologies. As in conventional survivable network design models, relaxing the path length limit for a “back-up path” may be appropriate; definitions of survivability should be considered with input from UTM implementers for whom safe operation is a main concern. The case study examined multiple instances drawn from each instance class; however, our objective was to show the validity of the observed results across inputs rather than design models that were robust to the actual realization of demand. In early stage development, UTM providers may want to design a network with good coverage regardless of the realized demand. In this case, developing stochastic models would provide valuable insights. The UTM network design problem and case study led to a high-level model with static path assignment, appropriate for the time frame of the decisions being made; however, the dynamic variant of the model gives rise to many operational questions. For example, once a UTM network consisting of 3D corridors above roads is established, which routing protocols will be best for handling flight requests in a busy urban airspace?

6.1 Appendix A: SU Fifty OD-Pair Results

pld%	cost		num arcs	
	avg	std	avg	std
0	196.99	12.361	393.80	21.16
10	125.45	6.64	244.00	6.04
20	107.81	4.50	211.80	3.56

Table 10: Minimum cost networks satisfying path length constraints for fifty OD-pair SU instances.

pld%	bd%	cost		num arcs		travel dist		avg pld		num on sp		max flow	
		avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
10	10	137.93	7.32	267.80	8.01	2169.39	296.44	1.27	0.49	18.60	4.34	38.20	9.76
10	20	150.48	8.00	294.20	7.29	2152.17	296.52	1.58	0.28	22.80	3.96	39.60	10.21
10	30	162.82	8.66	319.60	4.98	2144.52	295.15	0.24	0.13	27.80	5.81	38.60	9.53
10	40	175.40	9.06	344.20	9.09	2141.26	294.44	0.09	0.07	35.20	6.76	36.00	10.77
10	50	186.50	7.9	374.00	5.15	2139.77	293.81	0.02	0.02	42.40	6.27	34.00	12.65
20	10	118.57	4.97	234.60	8.76	2229.69	306.94	3.69	0.97	13.20	4.09	46.80	13.03
20	20	129.34	5.42	252.40	7.77	2185.88	299.31	2.03	0.87	17.20	3.83	36.80	2.28
20	30	140.09	5.87	271.80	9.12	2163.11	295.18	1.09	0.63	19.20	4.27	39.80	10.18
20	40	150.87	6.27	293.60	13.05	2151.52	295.40	0.57	0.36	23.20	5.07	39.60	10.21
20	50	161.53	6.57	316.20	9.78	2145.11	294.83	0.29	0.23	27.80	5.36	38.00	9.90

Table 11: Minimum travel distance with path length and budget constraints for fifty OD-pair SU instances.

		cost		num arcs		travel dist		avg pld		num on sp		max flow		travel time	
pld	bd	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std	avg	std
10	10	137.97	7.30	268.60	7.77	2177.75	300.03	1.57	0.46	16.20	2.86	29.40	5.98	2210.95	318.34
10	20	150.49	7.95	294.40	4.04	2160.23	302.85	0.89	0.37	20.80	2.86	28.80	6.61	2184.62	312.94
10	30	163.05	8.63	321.40	5.86	2152.18	299.12	0.53	0.19	24.20	3.42	27.00	7.00	2171.40	309.31
10	40	175.39	9.02	349.60	4.98	2148.22	298.00	0.35	0.19	28.00	6.28	25.60	4.72	2164.92	305.59
10	50	188.09	9.90	377.20	7.82	2146.38	296.82	0.28	0.15	31.20	7.50	25.20	5.07	2161.57	303.55
20	10	118.72	5.74	233.50	5.92	2115.70	129.40	4.20	0.97	12.00	2.94	31.00	1.41	2156.19	137.84
20	20	129.34	5.39	253.00	8.66	2199.70	311.57	2.52	1.03	14.20	4.15	29.00	6.00	2237.93	335.07
20	30	140.12	5.85	272.80	9.91	2171.77	298.64	1.47	0.65	17.00	2.45	30.20	5.50	2200.65	315.31
20	40	150.92	6.30	293.80	10.80	2159.33	300.42	0.90	0.35	19.80	2.05	29.20	6.76	2183.25	312.45
20	50	161.67	6.66	317.00	9.27	2153.28	298.99	0.61	0.22	23.20	4.02	26.80	6.98	2172.12	308.51

Table 12: Minimum travel time with path length and budget constraints for 50 OD-pair SU instances

7 References

- [1] Enrico Angelelli, Valentina Morandi, Martin Savelsbergh, and Maria Grazia Speranza. System optimal routing of traffic flows with user constraints using linear programming. *European journal of operational research*, 293(3):863–879, 2021.
- [2] Enrico Angelelli, Valentina Morandi, and Maria Grazia Speranza. Minimizing the total travel time with limited unfairness in traffic networks. *Computers & Operations Research*, 123:105016, November 2020.
- [3] Pasquale Avella, Giacomo Bernardi, Maurizio Boccia, and Sara Mattia. An optimization approach for congestion control in network routing with quality of service requirements. *Networks*, 74(2):124–133, 2019.
- [4] Gohram Baloch and Fatma Gzara. Strategic Network Design for Parcel Delivery with Drones Under Competition. *Transportation Science*, 54:204–228, January 2020.
- [5] Aleksandar Bauranov and Jasenka Rakas. Designing airspace for urban air mobility: A review of concepts and approaches. *Progress in Aerospace Sciences*, 125:100726, August 2021.
- [6] Vedat Bayram, Barbaros Ç. Tansel, and Hande Yaman. Compromising system and user interests in shelter location and evacuation planning. *Transportation Research Part B: Methodological*, 72:146–163, February 2015.
- [7] Geoff Boeing. Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems*, 65:126–139, 2017.
- [8] David Branston. Link capacity functions: A review. *Transportation Research*, 10(4):223–236, August 1976.

- [9] Gerard Casey, Bingyu Zhao, Krishna Kumar, and Kenichi Soga. Context-specific volume-delay curves by combining crowd-sourced traffic data with Automated Traffic Counters (ATC): a case study for London. Technical Report arXiv:2011.01964, arXiv, November 2020. arXiv:2011.01964 [stat] type: article.
- [10] Darshan Chauhan, Avinash Unnikrishnan, and Miguel Figliozzi. Maximum coverage capacitated facility location problem with range constrained drones. *Transportation Research Part C: Emerging Technologies*, 99:1–18, February 2019.
- [11] Christopher Chin, Karthik Gopalakrishnan, Hamsa Balakrishnan, Maxim Egorov, and Antony Evans. Efficient and fair traffic flow management for on-demand air mobility. *CEAS Aeronautical Journal*, October 2021.
- [12] Yi-Chang Chiu, Jon Bottom, Michael Mahut, Alexander Paz, Ramachandran Balakrishna, Steven Waller, and Jim Hicks. *Dynamic traffic assignment: A primer (transportation research circular e-c153)*. Transportation Research Board, 2011.
- [13] Sung Hoon Chung, Bhawesh Sah, and Jinkun Lee. Optimization for drone and drone-truck combined operations: A review of the state of the art and future directions. *Computers & Operations Research*, 123:105004, November 2020.
- [14] Taner Cokyasar, Wenquan Dong, Mingzhou Jin, and İsmail Ömer Verbas. Designing a drone delivery network with automated battery swapping machines. *Computers & Operations Research*, 129:105177, May 2021.
- [15] Open Street Map Contributors. Osm wiki: Illinois/road classification. https://wiki.openstreetmap.org/wiki/United_States/Road_classification. Accessed: 2022-11-10.
- [16] Open Street Map Contributors. Osm wiki: United states/road classification. https://wiki.openstreetmap.org/wiki/United_States/Road_classification. Accessed: 2022-11-10.
- [17] Teodor Gabriel Crainic, Michel Gendreau, and Bernard Gendron, editors. *Network design with applications to transportation and logistics*. Springer, 2021.
- [18] Christopher Decker and Paul Chiambaretto. Economic policy choices and trade-offs for Unmanned aircraft systems Traffic Management (UTM): Insights from Europe and the United States. *Transportation Research Part A: Policy and Practice*, 157:40–58, March 2022.
- [19] Malik Doole, Joost Ellerbroek, Victor L. Knoop, Victor L. Knoop, Victor L. Knoop, and Jacco M. Hoekstra. Constrained Urban Airspace Design for Large-Scale Drone-Based Delivery Traffic. *Aerospace*, 8(2):38, 2021.
- [20] David Edwards. Flytrex receives FAA approval to double drone delivery radius to two nautical miles. <https://roboticsandautomationnews.com/2022/07/28/flytrex-receives-faa-approval->

- [to-double-drone-delivery-radius-to-two-nautical-miles/53615/](#), July 28 2022. Accessed: 2022-12-23.
- [21] FAA. Uas facility maps. https://www.faa.gov/uas/commercial_operators/uas_facility_maps. Accessed: 2023-03-31.
- [22] Transportation Networks for Research Core Team. Transportation networks for research. <https://github.com/bstabler/TransportationNetworks>. Accessed: 2022-11-29.
- [23] Björn Geißler, Alexander Martin, Antonio Morsi, and Lars Schewe. Using piecewise linear functions for solving minlp s. In *Mixed integer nonlinear programming*, pages 287–314. Springer, 2012.
- [24] Sinan Gürel. A conic quadratic formulation for a class of convex congestion functions in network flow problems. *European Journal of Operational Research*, 211(2):252–262, June 2011.
- [25] Insu Hong, Michael Kuby, and Alan T. Murray. A range-restricted recharging station coverage model for drone delivery service planning. *Transportation Research Part C: Emerging Technologies*, 90:198–212, May 2018.
- [26] Wenjuan Hou, Tao Fang, Zhi Pei, and Qiao-Chu He. Integrated design of unmanned aerial mobility network: A data-driven risk-averse approach. *International Journal of Production Economics*, 236:108131, June 2021.
- [27] Olaf Jahn, Rolf H Möhring, and Andreas S Schulz. Optimal routing of traffic flows with length restrictions in networks with congestion. In *Operations research proceedings 1999*, pages 437–442. Springer, 2000.
- [28] Olaf Jahn, Rolf H. Möhring, Andreas S. Schulz, and Nicolás E. Stier-Moses. System-Optimal Routing of Traffic Flows with User Constraints in Networks with Congestion. *Operations Research*, 53(4):600–616, 2005.
- [29] Dae-Sung Jang, Corey A Ippolito, Shankar Sankararaman, and Vahram Stepanyan. Concepts of airspace structures and system analysis for uas traffic flows for urban areas. In *AIAA Information Systems-AIAA Infotech@ Aerospace*, page 0449, 2017.
- [30] Tao Jiang, Jared Geller, Daiheng Ni, John Collura, and John Collura. Unmanned Aircraft System traffic management: Concept of operation and system architecture. *International journal of transportation science and technology*, 5(3):123–135, October 2016.
- [31] Thomas Kirschstein. Comparison of energy demands of drone-based and ground-based parcel delivery services. *Transportation Research Part D: Transport and Environment*, 78:102209, January 2020.
- [32] Parimal H. Kopardekar. Unmanned Autonomous Systems (UAS) Traffic Management. In *Explore Flight: Caltrans Planning Horizons Forum*, number ARC-E-DAA-TN68454, June 2019.

- [33] Giusy Macrina, Luigi Di Puglia Pugliese, Francesca Guerriero, and Gilbert Laporte. Drone-aided routing: A literature review. *Transportation Research Part C: Emerging Technologies*, 120:102762, November 2020.
- [34] P. Marcotte. Network design problem with congestion effects: A case of bilevel programming. *Mathematical Programming*, 34(2):142–162, March 1986.
- [35] Tom V. Mathew and Sushant Sharma. Capacity Expansion Problem for Large Urban Transportation Networks. *Journal of Transportation Engineering*, 135(7):406–415, July 2009.
- [36] Rico Merkert and James Bushell. Managing the drone revolution: A systematic literature review into the current use of airborne drones and future strategic directions for their effective control. *Journal of Air Transport Management*, 89:101929, 2020.
- [37] Mohammad Moshref-Javadi and Matthias Winckenbach. Applications and Research avenues for drone-based models in logistics: A classification and review. *Expert Systems with Applications*, 177:114854, September 2021.
- [38] Enock Mtoi and Ren Moses. Calibration and Evaluation of Link Congestion Functions: Applying Intrinsic Sensitivity of Link Speed as a Practical Consideration to Heterogeneous Facility Types Within Urban Network. *Journal of Transportation Technologies*, 4:141–149, April 2014.
- [39] Ciara O’Brien. Alphabet’s Wing to trial drone deliveries in Ireland. <https://www.irishtimes.com/business/2022/10/20/alphabets-wing-to-trial-drone-deliveries-in-ireland/>, October 20 2022. Accessed: 2022-12-23.
- [40] Bureau of Public Roads. *Traffic assignment manual for application with a large, high speed computer*, volume 2. US Department of Commerce, Bureau of Public Roads, Office of Planning, Urban, 1964.
- [41] Alena Otto, Niels Agatz, James Campbell, Bruce Golden, and Erwin Pesch. Optimization approaches for civil applications of unmanned aerial vehicles (UAVs) or aerial drones: A survey. *Networks*, 72(4):411–458, 2018.
- [42] Michael Patriksson. *The traffic assignment problem: models and methods*. Courier Dover Publications, 2015.
- [43] Olga Perederieieva, Matthias Ehrgott, Andrea Raith, and Judith Y.T. Wang. A framework for and empirical study of algorithms for traffic assignment. *Computers & Operations Research*, 54:90–107, February 2015.
- [44] Roberto Pinto and Alexandra Lagorio. Point-to-point drone-based delivery network design with intermediate charging stations. *Transportation Research Part C: Emerging Technologies*, 135:103506, February 2022.

- [45] Stefan Poikonen and James F. Campbell. Future directions in drone routing research. *Networks*, 77(1):116–126, 2021.
- [46] Seyed Mahdi Shavarani, Mahmoud Golabi, and Gokhan Izbirak. A capacitated biobjective location problem with uniformly distributed demands in the UAV-supported delivery operation. *International Transactions in Operational Research*, 28(6):3220–3243, October 2019.
- [47] Seyed Mahdi Shavarani, Sam Mosallaeipour, Mahmoud Golabi, and Gokhan Izbirak. A congested capacitated multi-level fuzzy facility location problem: An efficient drone delivery system. *Computers & Operations Research*, 108:57–68, August 2019.
- [48] Skydrop. SkyDrop, Domino’s gear up to launch commercial drone delivery trial in New Zealand. <https://www.prnewswire.com/news-releases/skydrop-dominos-gear-up-to-launch-commercial-drone-delivery-trial-in-new-zealand-301573821.html>, June 23 2022. Accessed: 2022-12-23.
- [49] Amazon Staff. Amazon Prime Air prepares for drone deliveries. <https://www.aboutamazon.com/news/transportation/amazon-prime-air-prepares-for-drone-deliveries>, June 13 2022. Accessed: 2022-12-23.
- [50] Claudia Stöcker, Rohan Bennett, Francesco Nex, Markus Gerke, and J. Zevenbergen. Review of the Current State of UAV Regulations. *Remote Sensing*, 9:459, May 2017.
- [51] Walmart. We’re Bringing the Convenience of Drone Delivery to 4 Million U.S. Households in Partnership with DroneUp. <https://corporate.walmart.com/newsroom/2022/05/24/were-bringing-the-convenience-of-drone-delivery-to-4-million-u-s-households-in-partnership-with-droneup>, May 24 2022. Accessed: 2022-12-23.
- [52] John Glen Wardrop. Road paper. some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(3):325–362, 1952.
- [53] Wai Wong and S. C. Wong. Network topological effects on the macroscopic Bureau of Public Roads function. *Transportmetrica A: Transport Science*, 12(3):272–296, 2016. Publisher: Taylor & Francis.
- [54] Ammar Šarić and Ivan Lovrić. Improved Volume–Delay Function for Two-Lane Rural Highways with Impact of Road Geometry and Traffic-Flow Heterogeneity. *Journal of Transportation Engineering, Part A: Systems*, 147(10):04021066, October 2021.