A Scheduling Model for Aerial Ride-Sharing Operations

with Limited Infrastructure Capacity

R.V. van Sunten

TU Delft

A Scheduling Model for Aerial Ride-sharing Operations

with Limited Infrastructure Capacity

by

R.V. van Sunten

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday 18 May, 2021 at 09:30 AM.

Student number:	4381807	
Thesis committee:	Prof. dr. ir. J.M. Hoekstra,	TU Delft, chair
	Dr. A. Bombelli,	TU Delft, supervisor
	Ir. P.C. Roling,	TU Delft, examiner

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Preface

The submission of this thesis is a concluding step in obtaining my degree of Master of Science in Aerospace Engineering. A journey which started in September 2014 at the Faculty of Aerospace Engineering at TU Delft. Throughout the past six and a half years, I have learned a lot. Both as part of the curriculum and outside the classroom. In terms of knowledge, as well as skills. Through the study and associated organisations, I have been fortunate to have many unique experiences in the field of aerospace. These experiences have been very valuable: they have broadened my view and have further ignited my passion for aviation.

First of all, I would like to thank my thesis supervisor, Dr. Alessandro Bombelli, for his guidance and feedback throughout the thesis process. Shortly after starting this thesis assignment, COVID-restrictions were imposed, and all education and contact moved online. The lack of in-person discussions and brainstorming was sometimes a detriment to motivation and productivity. I am glad that with Alessandro's coaching, we managed to overcome the challenges of current times.

I would also like to thank my friends with whom I spent countless hours studying and, more importantly, a multitude of trips and activities. Thanks to you, I can look back on an enjoyable period as a student in Delft.

Finally, I would like to express sincere gratitude towards my parents, who have provided me with the independence and opportunity to follow my own study plan and supported me along the way.

R.V. van Sunten
Delft, May 2021

Cover image: Bell concept vehicle above Dallas, *https://www.bellflight.com/*.

Contents

List of Figures

List of Tables

List of Abbreviations

(E)VRP(TW)	(Electrical) Vehicle Routing Problem (with Time Windows)
(e)VTOL	(electric) Vertical Take-Off and Landing
DARP	Dial-a-Ride Problem
DEP	Distributed Electric Propulsion
EV	Electric Vehicle
KPI	Key Performance Indicator
LNG	Liquid/Liquefied Natural Gas
MaaS	Mobility as a Service
MILP	Mixed Integer Linear Programming
OD	Origin-Destination
pax	Passenger
PDP	Pick-up and Delivery Problem
SOC	State of Charge
TSN	Time-Space Network
TW	Time Window
UAM	Urban Air Mobility

Introduction

Rising vehicle usage leads to increasing road congestion in many cities around the world. Some major metropolises experience complete gridlocks, where stationary vehicles gravely contribute to the growing emission of pollutant gasses. Meanwhile, over the past years, drones have increasingly been used for cargo delivery services. The usage of drones has also been proposed for transporting people. This is the concept of Urban Aerial Mobility (UAM). Under some UAM scheme's, drones may partially, or ultimately even completely, replace road-based transportation. Especially in area's with road congestion, the prospect of flying over traffic jams in a drone is attractive as it may offer drastic time savings. Moreover, electrically propelled aircraft can reduce local emission of pollutant gasses. The Vertical Take-Off and Landing (VTOL) vehicles deployed for UAM will most likely not be privately owned by individual users. Instead, they will be exploited by service providers as part of the Mobility as a Service (MaaS) sharing economy. This service will work similarly to current car ride-sharing offerings, like Uber. In order to operate the VTOL vehicles commercially as part of UAM, vehicles must be matched to customer requests as efficiently as possible. An optimised flight schedule is thus required. Various mathematical optimisation models and theories address this type of problem. However, the majority of the existing literature is based on road-based transportation. UAM operations will pose a unique set of constraints and parameters that have not yet explicitly been addressed in a scheduling model.

The aim of this thesis is, therefore, to develop a scheduling model that includes conditions specific to prospective UAM ride-sharing operations. This goal is achieved by identifying the commonalities and differences between road-based and aerial ride-sharing, assessing existing optimisation models for gaps, and ultimately integrating this into a mathematical formulation and a programmed model. The model will schedule flights for the available fleet such that the profit for the provider is maximised. This objective is subject to specific constraints such as a fleet of a fixed amount of specified VTOL vehicles, requirements for battery charging and, in particular, a limited infrastructure capacity. The model developed is used in a case study where different distributions of customer demand are applied. The results are used to assess the model's behaviour for a variety of use cases. The findings are also used to develop recommendations for future UAM operations.

This research project makes the following contributions: first of all, we provide a comprehensive, up-to-date comparison of automotive and aerial ride-sharing concepts, as well as existing scheduling models and their gaps. Secondly, we integrate these theoretical findings and propose, to the best of our knowledge, a novel formulation to address limited infrastructure capacity without using a Time-Space Network (TSN) formulation. The review of different demand distributions applied to the model provides an exploratory insight into the profit behaviour for various possible use cases. Together, these contributions add to the limited body of literature for scheduling future UAM operations. This work may serve as a base for future research on the topic and offers a modelling foundation to investigate the effect of changes in resources and parameters. Even though the formulation developed is tailored to UAM operations, it may also be generalised for other applications in ride-sharing and/or electrical vehicle scheduling with limited infrastructure capacity.

This document is divided into various parts. Part I contains the scientific paper. The paper contains a brief summary of the background and literature review, the methodology with the model formulation, key results, discussion and conclusion. The paper can be read as a stand-alone article. Part II subsequently contains the extensive literature study as well as the integral research plan, which were both previously performed and graded as part of course AE4020. Because the field of UAM and associated scientific areas receive ongoing academic attention, during the thesis process novel and relevant literature emerged. Together with new insights gathered during the model development, this lead to an adjustment in the theoretical focus. As such, the literature presented in the paper may differ from the initial literature review. Lastly, part III encompasses supporting work for this research project which has not made it into the more concise paper.

I

Scientific Paper

A Scheduling Model for Aerial Ride-sharing Operations

R.V. van Sunten,*

Delft University of Technology, Delft, The Netherlands

Abstract

Following the advent of drones for surveillance and cargo delivery purposes, advancements in recent years have also been made towards the development of larger drones for passenger transport. The concept of Urban Air Mobility (UAM) prospects to offer ride-sharing services within and between cities. While trip scheduling and vehicle routing algorithms exist for various forms of road-based transportation services, UAM operations pose specific constraints and requirements that, to the best of our knowledge, have not been addressed comprehensively in academic literature. The purpose of this research is to develop a Mixed-Integer Linear Programming (MILP) model that optimally matches the available VTOL (Vertical Take-off and Landing) vehicle fleet with customer trip requests, subject to (UAM-specific) constraints, and which subsequently provides a vehicle's routing. The model, in particular, addresses the constraint of limited ground infrastructure capacity. A case study is performed where multiple demand distribution scenarios, resembling different use cases, are applied to the model. Results show that wider time windows do not have a clear beneficial effect on the profit and that customer demand distribution has an impact on the efficiency of the operation. Additionally, the results enable the identification of various key input parameters for infrastructure, fleet size and vehicle technology that can improve the overall operation.

Index Terms - Urban Air Mobility, Aerial Ride-sharing, Electric Vertical Take-off and Landing Vehicle, Vehicle Routing Problem, Pick-up and Delivery Problem, Dial-a-Ride Problem, Fleet Scheduling Problem

1 Introduction

Technological advancements in recent years have fast-tracked the usage of drones in various concepts and sizes (Holden et al., 2016). From small personal hobby drones to large vehicles for package transport, the so-called 'vertical mobility services' are set to form a global market worth 74 billion *USD* in 2035, according to Porsche Consulting (2018). Within vertical mobility services, commercial passenger transport is set to be a major component. While the concept of UAM has traditionally been reserved for the ultra-rich, using helicopters to fly in urban settings, multiple players are active to develop vehicles and services that appeal to a much larger consumer base (Boelens and Volocopter, 2019; Booz Allen Hamilton, 2018; EHang, 2020; Holden et al., 2016). Futuristic scenarios picture (electric) VTOL vehicles to transport passengers within and between cities. Complementing and eventually even partly replacing road-based transportation. If the production and operating costs are sufficiently lowered, UAM transportation can have many advantages. Flying over traffic jams, especially in heavily congested metropolises, can yield considerable time savings. Moreover, if the VTOL aircraft are propelled electrically, emission gasses can be reduced locally.

eVTOL aircraft most likely will not be owned by individual consumers. Instead, service providers will offer UAM flights between specific locations, called vertiports, or eventually to any suitable spot in the area of operation. Ad-hoc or scheduled in advance. UAM will be part of the Mobility as a Service (MaaS) framework and will offer an experience similar to ride-sharing operators like Uber, Lyft and Didi Chuxing. For this, matching customer demand and flights supply, and generating optimal routes for the eVTOL vehicles is paramount.

While much research has been performed to optimise road-based transportation, including electric vehicles and on-demand ride providers, there has been little research focused specifically on scheduling commercial UAM operations. The aim of this research is to extend the academic body of literature to the aerial domain, with its specific traits, and to generate a scheduling model. This goal is achieved by finding relevant and representative parameters and constraints specific to UAM operations, integrating these into a MILP optimisation model that finds vehicle routings that best satisfy the customer demand and, consequently, applying different demand distribution scenarios to this model in order to simulate various future use cases.

This paper will have the following contributions. It extends existing research on vehicle scheduling models to the aerial domain, where the current body of academic literature is limited. This addition includes the optimisation model and also an overview of the most important differences and similarities between the road-based and aerial domains. The latter is based on academic literature as well as a comprehensive review of

*MSc Student, Air Transport and Operations, Faculty of Aerospace Engineering, Delft University of Technology

industry players, such as vehicle manufacturers, operators, research institutions and consultancy firms. In the scheduling model, various theoretical and practical concepts from different fields of research are combined. E.g. the vehicle, network and demand submodels. This model can subsequently serve as a base for a broad range of follow-up research projects, diving either deeper into the effects of a single aspect or extending the model to account for other factors.
The structure of this paper is as follows. Section 2 will discuss the state of the art in existing literature relevant to this topic. Following that, in section 3 the methodology for the research is presented. In this section, the mathematical formulation will be declared. The input parameters for the case study will also be explained. Section 4 will consequently provide the model simulations' outcomes according to various demand input scenarios. Section 5 will discuss the research findings, note the strengths and limitations of the model and provide recommendations for further research. The paper is concluded in section 6 by providing a synopsis of the most prominent results and the implications for the further development of UAM.

2 Literature Review

The literature review conducted focused on three area's: mathematical scheduling models in transportation, ride-sharing and UAM operations in conjunction with VTOL vehicles. Considering the latter subject's relatively novel practical state, next to academic literature also industry and media sources were consulted. An overview of the available literature on these subjects will be presented in this section. The section concludes with a synopsis on the limitations of the current literature body, which form the research gap.
Finding a route for a fleet of vehicles to satisfy customer requests is a well-researched problem, also known as the Vehicle Routing Problem (VRP), and overviews are provided by, a.o., Bunte and Kliewer (2009); Toth and Vigo (2000) and Pillac et al. (2013). This problem may also feature time frames imposed on the service of customers (VRP with Time Windows). The associated algorithms can be static or dynamic with respect to the demand data. When transporting goods or people, models can be classified as a Pick-up and Delivery Problem (PDP), defined by, a.o., Dumas et al. (1991). Berbeglia et al. (2007) provide a classification scheme for various PDP variants. The Dial-a-Ride Problem (DARP) is another, more specific, subproblem where users specify a pick-up and drop-off location, which is addressed in Bodin and Sexton (1986); Dumas et al. (1989) and Cordeau and Laporte (2003). Typical objective functions focus on minimising cost, mileage or inconvenience factors such as waiting time and ride time when transporting people. Constraints are usually imposed on the vehicle's capacity, flow continuity and adhering to time windows on departure, arrival, or both.
In recent years, with the increasing use of electric vehicles, another subproblem has been established in academic literature: the Electric Vehicle Routing Problem (EVRP), which is discussed in Afroditi et al. (2014); De Almeida Correia and Santos (2014); Schneider et al. (2014) and Liang et al. (2016). The EVRP imposes additional constraints to account for the vehicle's battery level. Some models force that the battery level does not drop below 0 at every time instance, while other formulations require that there has to be sufficient battery level remaining to drive to the nearest charging station.
The optimisation models, especially the DARP, have been applied to ride-sharing offerings, for which Shaheen and Cohen (2020) give definitions in the context of MaaS. Furuhata et al. (2013) and Chan and Shaheen (2012) discuss the state-of-the-art and future developments of ride-sharing, as well as the (dis)advantages and characteristics. A review on optimising ride-sharing can be found in Agatz et al. (2012). The literature on this subject is mostly limited to ground-based transportation.
Research on (e)VTOL vehicles and UAM operations is extensive but often limited to a specific component. Kohlman and Patterson (2018) compared various fuel and propulsion forms for VTOL vehicles. While the usage of liquefied natural gas has some advantages, the vast majority of literature presumes electrically powered vehicles (Alnaqeb et al., 2018; Boelens and Volocopter, 2019; Holden et al., 2016; Narkus-Kramer, 2013; Pradeep and Wei, 2018). Especially regarding (future) battery technology such as capacity, specific power and specific energy, literature shows a range of figures (Johnson and Silva, 2019; Justin et al., 2017). Silva et al. (2018) give an overview of VTOL vehicle concepts and configurations which, again, shows a wide variety of (theoretical) conceptions. Vehicles in development are described by various manufacturers and operators such as EHang (2020); Uber (2020) and Volocopter (2020). In terms of infrastructure, Vascik and Hansman (2019) investigated vertiport requirements and developed a model to determine the capacity. Comprehensive research to integrated UAM operations has been performed by Dunn (2018); Kohlman and Patterson (2018); Patterson et al. (2018); Vascik and Hansman (2018) and Vascik (2020), which includes determining operational constraints, customer demand, modelling and system analyses. Market studies have been performed by both Hasan and NASA (2018) and consultancy firms, including Porsche Consulting (2018); Roland Berger (2018) and Booz Allen Hamilton (2018). These studies have also put forward different business models for UAM, ranging from offering commuter flights, airport shuttles, air taxis and longer range intercity flights. Operator's strategies vary. Nevertheless, there is consensus that there is a minimum distance from which UAM operations become attractive from a time-saving point of view, considering the additional time required for check-in, boarding and travel to/from

the vertiport. However, there is also an upper bound. For longer distance flights, the vehicle's range and battery capacity become a limiting factor.
From the review of the literature referenced heretofore, similarities and differences between the novel UAM operations and conventional road-based transportation could be determined. These affect the applicability of existing scheduling models to UAM operations. The most profound disparities constitute the network structure and reliance on batteries. Contrary to automotive service providers, which can pick up and drop off customers virtually anywhere, UAM service providers in the near future will focus on providing flights between fixed locations due to the requirements for specific infrastructure such as charging ports and, depending on regulations, also for passenger facilities such as check-in and security screening (Holden et al., 2016). Next to that, factors such as nearby buildings, airspace and noise abatement restrictions, and routing procedures influence where VTOL aircraft can take off and land (National Air Transportation Association, 2018; Thipphavong et al., 2018; Vascik and Hansman, 2018,1). In the VRP, PDP and DARP, a vehicle will typically only visit the same location once during a single route. For UAM operations, a vertiport may be visited multiple times by the same vehicle during a routing to serve multiple customers travelling between the same origin-destination (OD) pair on different instances. This network structure, where multiple vehicles can serve a certain physical vertiport (multiple times), also means that the vertiport's capacity, either in the number of landing pads or charging ports available, imposes an additional constraint that is not present in the traditional VRP, PDP or DARP. In the EVRP literature consulted, a (charging) station's capacity was also not addressed explicitly. Considering the current and prospective characteristics of batteries available for eVTOL aircraft, charging rates, and the commonly stated flight distances, eVTOL aircraft will most likely not be able to perform multiple flights on one full charge, if at all perform more than a single flight before charging is required. eVTOL vehicles will have to spend more time charging relative to their operational flying time compared with automotive transport, occupying the ground infrastructure. This characteristic will subsequently form a limitation that needs to be included in the scheduling model as well.
Shihab and Wei (2019) provide a small-scale scheduling model, one of the very few tailored to UAM operations, based on different transportation service types: on-demand, scheduled and a hybrid combination. The model includes constraints on the vertiport capacity and remaining battery capacity. The model has its limits, however. Due to the Time-Space Network (TSN) formulation, the model is solved at each time step of 30 minutes. Furthermore, the battery level can only take three values: full, half-full or empty. Both of these assumptions limit the accuracy of the model. In subsequent work, Shihab and Wei (2020) have amended the battery capacity constraint such that it depends on the energy required for a particular trip. However, the formulation is still solved in a TSN for time slices of 30 minutes, which is the maximum duration possible to operate a route in the scenario. Therefore, this approach cannot capture the effect of any time windows less than the chosen 30 minutes. Depending on the specific use case, this may be coarse for commercial UAM operations. A higher resolution in the TSN is possible but comes at the expense of an increased problem size. Furthermore, the formulation is quadratic in the constraints and objective function, which complicates solving the model.
A TSN-representation enables monitoring of the spatial position of a vehicle at each time slice. This feature allows counting the number of vehicles at a certain vertiport, which is a practical consideration for UAM operations where ground infrastructure is limited, and the fleet size usually exceeds the number of pads available at a vertiport. This information on the vehicles' position can be used to apply the vertiport capacity constraint.
Another method to impose a limit on a station's capacity, without a TSN-formulation, has been addressed in the truck dock assignment problem where trucks need to be assigned to a specific dock when visiting a warehouse (Bombelli and Tavasszy, 2020; Gelareh et al., 2013; Miao et al., 2009). Trucks can only be assigned to the same dock if their visits do not overlap in time. These constraints, however, are indexed to a specific vehicle and do not allow for the same vehicle to visit the same location more than once.
Summarising, a range of academic and industry literature has been reviewed, which led to an evaluation of the similarities and differences between existing road-based transportation and novel UAM operations. The characteristics were compared with established vehicle scheduling models. It has been found that some aspects of UAM operations are not sufficiently addressed in academic literature, leading to a research gap. The model presented in this paper will explicitly address the vertiport capacity constraint and charging requirements.

3 Methodology

In this section, the methodology of the research will be elaborated upon. First, the problem statement is provided in section 3.1. Next, the resulting graph representation is given in section 3.2. Section 3.3 details the mathematical formulation of the model and explains the constraints, variables, sets and parameters used. The computational solving method is explained in section 3.4. Finally, in section 3.5, the case study where the scheduling model is applied will be discussed. In this section, the various submodels and input parameters used for, a.o., the vehicle, customer demand and infrastructure network are also treated.

3.1 Problem Statement

The UAM scheduling model's aim is to assign vehicles to serve customer requests, thereby generating a routing for each vehicle. A homogeneous fleet is assumed. Customers specify a request for transport that consists of an origin and destination from a finite number of locations and the desired departure time. A fixed time window is imposed on this departure time to provide a broader time range when the customer can be picked up at the origin. Based on the direct travel time between the origin and destination, the earliest and latest allowed delivery time at the destination are determined. All customer demand is known *a-priori*, making the model static. The model's objective function maximises the profit of serving the customer requests, considering the revenue of satisfied requests, (operational) costs and monetary penalties.
All vehicles start and end their operation at a central depot. Vehicles can start leaving the depot one hour before the first possible customer trip request in order to position to a different location and recharge. Vehicles can either be serving a customer request (picking up, flying and delivering passengers), charge at a charging station, idle after a full charge or position to a different location (empty). A customer request can consist of multiple passengers travelling together. If possible within the time windows, multiple customer requests can be consolidated in a single vehicle's flight if the vehicle's passenger capacity is not exceeded. During every phase where the vehicle is on the ground, it needs to be assigned to a pad (similar to a gate). The infrastructure available at each vertiport is assumed to be homogeneous, except for the depot where the pad capacity is sufficient to accommodate the entire fleet.

3.2 Graph Formulation

The physical scenario is first converted to a graph notation, consisting of nodes and edges. The cardinality and directions of the edges will be elaborated upon subsequently. As mentioned in section 2, the UAM scheduling model is characterised by the fact that pick-ups and deliveries take place at specific physical locations, the vertiports, and that there may be multiple customer requests from/to the same physical location. In the graph model, however, every pick-up and delivery point should be distinct. Therefore, each pick-up and delivery point of a customer request is represented by a node. Furthermore, each pick-up and delivery node will have an associated charging node. Multiple nodes may correspond to the same physical location. If a vehicle has finished charging at a charging node and cannot yet proceed to the next node, idling occurs at the charging node. In the current model, a vertiport's capacity is limited by the number of pads available. It is assumed that every pad is equipped with a charging facility. If the number of charging stations differs from the number of pads available, the charging node formulation also allows expanding the formulation to impose a constraint on the available charging stations separately.
There is a single physical depot which is also represented by two distinct nodes in the graph: a starting depot node and a depot node for the end of service. This approach also readily allows the implementation of a network where the origin and destination depots are distinct locations in practice.
A trip to a vertiport may consist of multiple nodes. Each trip to a vertiport starts with an entry node where the vehicle arrives and ends with an exit node where the vehicle departs to the next vertiport.

The representation presented heretofore is depicted in fig. 1, where the routes of two vehicles are expressed, including a recurring trip to the same vertiport 1.

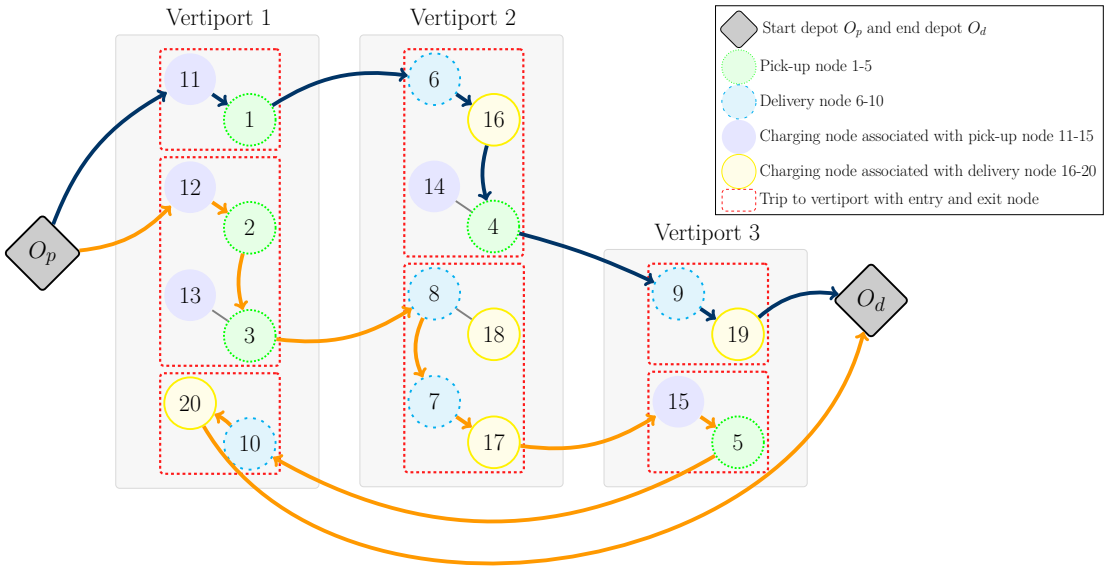

Figure 1: Graphical example of the graph representation with routes of two vehicles shown.

For the notation, the directed graph is represented by $G = (\mathcal{N}, \mathcal{E})$ where $\mathcal{N}$ is the set of nodes and $\mathcal{E}$ is the set of edges. The set of nodes is the union of the sets of pick-up nodes N_p, delivery nodes N_d, charging nodes N_c, pick-up depot node O_p and delivery depot node O_d. I.e. $\mathcal{N} = N_p \cup N_d \cup N_c \cup O_p \cup O_d$. The set of all charging nodes corresponding with the pick-up nodes is noted as N_{c^+} and, analogously, the set of charging nodes corresponding with delivery nodes is noted as N_{c^-}. Consequently, $N_c = N_{c^+} \cup N_{c^-}$.
The physical locations, vertiports, are in the set of customer vertiports V. Both the pick-up and delivery depot nodes O_p and O_d, respectively, correspond to depot vertiport, V_0. Set $V' = V \cup V_0$.
A vehicle k from set K can enter a vertiport through an entry node from set $N_{entry} = N_{c+} \cup N_d$ and exit through an exit node from the set of possible exit nodes $N_{exit} = N_{c^-} \cup N_p$. Both N_{entry} and N_{exit} are subsets of $\mathcal{N}$. From entry through to exit of a vertiport, a vehicle is assigned a single pad p from set P. For clarity, all sets are noted in table 1.

Table 1: Sets for the UAM model

Set	**Description**
N_p	Pick-up nodes
N_d	Delivery nodes
N_c	Charging nodes
N_{entry}	Entry nodes
N_{exit}	Exit nodes
$\mathcal{N}$	All nodes
$\mathcal{E}$	Edges
K	Vehicles
V	Customer vertiports
V'	All vertiports
P	Pads

3.2.1 Indexation and Cardinality

There are n customer requests which may be served. The index of the pick-up depot is 0. The enumeration of the pick-up nodes is from 1 to n. Delivery nodes are indexed from $n+1$ to $2n$. The pick-up charging nodes indices range from $2n+1$ to $3n$, and the delivery charging nodes indices range from $3n+1$ to $4n$. Finally, the delivery depot is indexed to $4n+1$. It follows that the delivery node j corresponding to pick-up node i has the relation $j = i + n$. Furthermore, the charging nodes associated with the pick-up nodes and delivery nodes are $i + 2n$ and $j + 2n$, respectively. In total, there are $|\mathcal{N}| = 4n + 2$ nodes.
It is assumed that a charge takes place before a pick-up and after a delivery. This assumption drastically decreases the number of edges. There are n edges from the pick-up depot to each pick-up charging nodes. From each pick-up charging node, there are n edges to the associated pick-up node. Subsequently, from the pick-up nodes, there are n edges to the associated delivery nodes and $n(n-1)$ edges to other pick-up nodes in the same vertiport or pick-up charging nodes located at other vertiports. From each delivery node, there are n edges to the associated delivery charging node and, at most, $n(n-1)$ edges to other delivery nodes at the same vertiport. Lastly, from the delivery charging nodes, there are n exiting edges to the delivery depot, at most $n(n-1)$ edges to delivery nodes at other vertiports, and $n(n-1)$ edges to pick-up nodes located at the same vertiport or pick-up charging nodes located at other vertiports. This results in a cardinality of $\mathcal{E}$ of $|\mathcal{E}| = 3n^2+2n$.

3.3 Model and Mathematical Formulation

The decision variables for the model, with their respective variable type and description, are listed in table 2. The following subscripts/superscripts are used: i, j for nodes, customer requests or trips, k for vehicles, v for vertiports and p for pads. Furthermore, for the sake of clarity, all decision variables are denoted in capital letters, while parameters and constants are written in lower case.

Table 2: Decision variables for the UAM model

Decision Variable	**Type**	**Description**
$X^k_{i,j}$	Binary	Unitary if vehicle k moves from node i to node j
T_i	Integer	Start of service time at node i
$T^k_{O_p}, T^k_{O_d}$	Integer	Start of service time of vehicle k at starting depot O_p and ending depot O_d, respectively
Q_i	Integer	Vehicle load at node i
E_i	Integer	Energy remaining at node i
Z_i	Binary	Unitary if customer request i is not served
U^k	Binary	Unitary if vehicle k is dispatched to serve customer requests
DT_i	Integer	Time difference between the earliest possible arrival time and actual arrival time at delivery node i
RT_i	Integer	Recharging time to a full battery charge at charging node i
NOT_i, NCT_i	Integer	Time at which node i is occupied/cleared
$VOT^k_{i,v}, VCT^k_{i,v}$	Integer	Time at which vehicle k occupies/clears node i which corresponds to vertiport v
$TAT^k_{i,v}$, $TDT^k_{i,v}$	Integer	Arrival/departure time of vehicle k at vertiport v during a trip with entry node i
$\phi^k_{i,v}$	Binary	Unitary if node i, at vertiport v, is an entry node for vehicle k
$\psi^k_{j,v}$	Binary	Unitary if node j, at vertiport v, is an exit node for vehicle k
$\rho^k_{i,j,v}$	Binary	Unitary if entry node j has an arrival time that is later than the arrival time at entry node i, for vehicle k at vertiport v
$\theta^k_{i,j,v}$	Binary	Unitary if entry node j has an arrival time that is later than the departure time at exit node i, for vehicle k at vertiport v
$\zeta^k_{i,v}$	Integer	Sum of the number of entry nodes at vertiport v for vehicle k that have a later arrival time than the arrival time at entry node i
$\xi^k_{i,v}$	Integer	Sum of the number of entry nodes at vertiport v for vehicle k that have a later arrival time than the departure time at exit node i
$\mu^k_{i,j,v}$	Binary	Unitary if vehicle k, during one trip, enters vertiport v at node i and exits the same vertiport at node j
$W^{k1,k2}_{i,j,v}$	Binary	Unitary if the arrival time of vehicle $k2$, operating trip j, is equal to or later than the departure time of vehicle $k1$, operating trip i, at vertiport v
$Y^k_{i,v,p}$	Binary	Unitary if vehicle k, operating trip i, is assigned to pad p at vertiport v.
$C^{k1,k2}_{i,j,p,v}$	Binary	Unitary if vehicle $k1$, operating trip i, and vehicle $k2$, operating trip j, are both assigned to pad p, at vertiport v

The (input) parameters used for the model are presented in table 3, together with a description.

Table 3: Parameters for the UAM model

Parameter	Description
$[et_i, lt_i]$	Time window for earliest/latest service time at node i
$d_{i,j}$	Distance between nodes i and j
$t_{i,j}$	Travel time between nodes i and j
$e_{i,j}$	Energy required to travel between nodes i and j
δ_i	Passenger demand at pick-up node i
cap	Passenger capacity of eVTOL vehicle
bat	Full battery charge
h	Battery charging rate
f_c	Fare charged to a passenger, per unit distance
c_o	Operational costs, per unit distance
c_e, c_v, c_t, c_r	Cost per unit energy, vehicle dispatched, unit excess time delay and rejected trip, respectively
M	Big-M, for which an upper bound can be determined
ϵ	Small, positive number

The MILP formulation of the UAM scheduling model is as follows:

$$\max \quad \mathcal{P} = \sum_{i\in N_p}\sum_{j\in N_p\cup N_d}\sum_{k\in K}\Big(f_c * \delta_i * d_{i,i+n} - c_o * d_{i,j}\Big)X^k_{i,j} - \sum_{i\in N_d} c_t DT_i - \sum_{i\in N_p} c_r Z_i - \sum_{k\in K} c_v U^k - \sum_{i\in N_c} c_e(bat - E_i) \tag{1.a}$$

s.t.

$$\sum_{j\in\mathcal{N}}\sum_{k\in K} X^k_{i,j} + Z_i = 1, \qquad \forall i \in N_p, i \neq j, \tag{1.b}$$

$$\sum_{j\in\mathcal{N}} X^k_{i,j} - \sum_{j\in\mathcal{N}} X^k_{i+n,j} = 0, \qquad \forall i \in N_p, i \neq j, \forall k \in K \tag{1.c}$$

$$\sum_{j\in\mathcal{N}} X^k_{i,j} - \sum_{j\in\mathcal{N}} X^k_{j,i} = 0, \qquad \forall i \in \mathcal{N}\backslash\{O_p, O_d\}, i \neq j, \forall k \in K \tag{1.d}$$

$$\sum_{j\in N_{c^+}} X^k_{O_p,j} - U^k = 0, \qquad \forall k \in K \tag{1.e}$$

$$\sum_{i\in N_{c^-}} X^k_{i,O_d} - \sum_{j\in N_{c^+}} X^k_{O_p,j} = 0, \qquad \forall k \in K \tag{1.f}$$

$$et_{O_p} \leq T^k_{O_p}, \qquad \forall k \in K \tag{1.g}$$

$$T^k_{O_d} \leq lt_{O_d}, \qquad \forall k \in K \tag{1.h}$$

$$et_i \leq T_i \leq lt_i, \qquad \forall i \in N_p \tag{1.i}$$

$$et_{j-n} \leq T_j \leq lt_{j-n}, \qquad \forall j \in N_d \tag{1.j}$$

$$T_j \geq T^k_{O_p} + t_{O_p,j} - M(1 - X^k_{O_p,j}), \qquad \forall j \in N_{c^+}, \forall k \in K \tag{1.k}$$

$$T^k_{O_d} \geq T_i + RT_i + t_{i,O_d} - M(1 - X^k_{i,O_d}), \qquad \forall i \in N_{c^-}, \forall k \in K \tag{1.l}$$

$$T_j \geq T_i + RT_i + t_{i,j} - M(1 - \sum_{k\in K} X^k_{i,j}), \qquad \forall i \in \mathcal{N}\backslash\{O_p, O_d\}, \forall j \in \mathcal{N}\backslash\{O_p, O_d\}, i \neq j, \tag{1.m}$$

$$RT_i = (bat - E_i)/h, \qquad \forall i \in N_c \tag{1.n}$$

$$DT_j \geq NOT_{j,v} - et_{j-n} - M(1 - \sum_{i\in N_p\vee N_d}\sum_{k\in K} X^k_{i,j}), \qquad \forall j \in N_d, i \neq j, \forall v \in V : j \in v \tag{1.o}$$

$$X^k_{O_p,j} = 1 \implies Q_j = 0, \qquad \forall j \in N_{c^+}, \forall k \in K \tag{1.p}$$

$$X^k_{i,j} = 1 \implies Q_i + \delta_j = Q_j, \qquad \forall i \in \mathcal{N}\backslash\{O_p, O_d\}, \forall j \in N_p, i \neq j, \forall k \in K \tag{1.q}$$

$$X^k_{i,j} = 1 \implies Q_i - \delta_{j-n} = Q_j, \qquad \forall i \in N_p \cup N_d \cup N_{c^-}, \forall j \in N_d, i \neq j, \forall k \in K \tag{1.r}$$

$$X^k_{i,j} = 1 \implies Q_i = Q_j, \qquad \forall i \in \mathcal{N}\backslash\{O_p, O_d\}, \forall j \in N_c, i \neq j, \forall k \in K \tag{1.s}$$

$$\delta_i \leq Q_i \leq cap, \qquad \forall i \in N_p \tag{1.t}$$

$$X^k_{i,j} = 1 \implies E_j = E_i - e_{i,j}, \qquad \forall i \in N_p \cup N_d \vee O_p, \forall j \in \mathcal{N}\backslash\{O_p, O_d\}, i \neq j, \forall k \in K \tag{1.u}$$

$$X^k_{i,j} = 1 \implies E_j = bat - e_{i,j}, \qquad \forall i \in N_c, \forall j \in \mathcal{N}\backslash\{O_p, O_d\}, i \neq j, \forall k \in K \tag{1.v}$$

$$E_{O_p} = bat, \tag{1.w}$$

$$X^k_{O_p,j} = 1 \implies NOT_{j,v} = T^k_{O_p} + t_{O_p,j}, \qquad \forall j \in N_{c^+}, \forall k \in K, \forall v \in V \tag{1.x}$$

$$X^k_{i,j} = 1 \implies NOT_{j,v} = T_j, \qquad \forall i, j \in \mathcal{N}\backslash\{O_p, O_d\}, i \neq j, \forall k \in K, \forall v \in V \tag{1.y}$$

$$X^k_{i,j} = 1 \implies NCT_{i,v} = T_i, \qquad \forall i \in N_p \cup N_d, \forall j \in \mathcal{N}\backslash\{O_p, O_d\}, i \neq j, \forall k \in K, \forall v \in V \tag{1.z}$$

$$X^k_{i,j} = 1 \implies NCT_{i,v} = T_j - t_{i,j}, \qquad \forall i \in N_c, \forall j \in \mathcal{N}\backslash\{O_p, O_d\}, i \neq j, \forall k \in K, \forall v \in V \tag{1.aa}$$

$$X^k_{i,O_d} = 1 \implies NCT_{i,v} = T^k_{O_d} - t_{i,O_d}, \qquad \forall i \in N_{c^-}, \forall k \in K, \forall v \in V \tag{1.ab}$$

$$X^k_{i,j} = 1 \implies \phi^k_{j,v} = 1, \qquad \forall i \in \mathcal{N}\backslash\{O_d\}, \forall j \in N_{entry}, i \neq j, \forall v \in V : i \notin v \vee j \in v, \forall k \in K \tag{1.ac}$$

$$X^k_{i,j} = 1 \implies \psi^k_{i,v} = 1, \qquad \forall i \in N_{exit}, \forall j \in \mathcal{N}\backslash\{O_p\}, i \neq j, \forall v \in V : i \in v \vee j \notin v, \forall k \in K \tag{1.ad}$$

$$\phi^k_{i,v} = 1 \implies VOT^k_{i,v} = NOT_{i,v}, \qquad \forall i \in N_{entry}, \forall v \in V : i \in v, \forall k \in K \tag{1.ae}$$

$$\psi^k_{i,v} = 1 \implies VCT^k_{i,v} = NCT_{i,v}, \qquad \forall i \in N_{exit}, \forall v \in V : i \in v, \forall k \in K \tag{1.af}$$

$$VOT^k_{i,v} - VOT^k_{j,v} + \epsilon + M * \rho^k_{i,j,v} \leq M, \qquad \forall i, j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.ag}$$

$$-VOT^k_{i,v} + VOT^k_{j,v} - M * \rho^k_{i,j,v} \leq 0, \qquad \forall i, j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.ah}$$

$$VCT^k_{i,v} - VOT^k_{j,v} + \epsilon + M * \theta^k_{i,j,v} \leq M, \qquad i \in N_{exit}, \forall j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.ai}$$

$$-VCT^k_{i,v} + VOT^k_{j,v} - M * \theta^k_{i,j,v} \leq 0, \qquad i \in N_{exit}, \forall j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.aj}$$

$$\zeta^k_{i,v} = \sum_{j \in N_{entry}} \rho^k_{i,j,v}, \qquad \forall i \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.ak}$$

$$\xi^k_{i,v} = \sum_{j \in N_{entry}} \theta^k_{i,j,v}, \qquad \forall i \in N_{exit}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.al}$$

$$\zeta^k_{i,v} = \xi^k_{j,v} \vee \phi^k_{i,v} = 1 \implies \mu^k_{i,j,v} = 1, \qquad \forall i \in N_{entry}, \forall j \in N_{exit}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.am}$$

$$\mu^k_{i,j,v} = 1 \implies TAT^k_{i,v} = VOT^k_{i,v}, \qquad \forall i \in N_{entry}, \forall j \in N_{exit}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.an}$$

$$\mu^k_{i,j,v} = 1 \implies TDT^k_{i,v} = VCT^k_{j,v}, \qquad \forall i \in N_{entry}, \forall j \in N_{exit}, i \neq j, \forall v \in V : i, j \in v, \forall k \in K \tag{1.ao}$$

$$\phi^k_{i,v} \implies \sum_{p \in P} Y^k_{i,v,p} = 1, \qquad \forall i \in N_{entry}, \forall v \in V : i \in v, \forall k \in K \tag{1.ap}$$

$$TDT^{k1}_{i,v} - TAT^{k2}_{j,v} + \epsilon + M * W^{k1,k2}_{i,j,v} \leq M, \qquad \forall i, j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k1, k2 \in K, k1 \neq k2 \tag{1.aq}$$

$$-TDT^{k1}_{i,v} + TAT^{k2}_{j,v} - M * W^{k1,k2}_{i,j,v} \leq 0, \qquad \forall i, j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k1, k2 \in K, k1 \neq k2 \tag{1.ar}$$

$$Y^{k1}_{i,v,p} + Y^{k2}_{j,v,p} - C^{k1,k2}_{i,j,p,v} \leq 1, \qquad \forall i, j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k1, k2 \in K, k1 \neq k2, \forall p \in P \tag{1.as}$$

$$C^{k1,k2}_{i,j,p,v} - W^{k1,k2}_{i,j,v} - W^{k2,k1}_{j,i,v} \leq 0, \qquad \forall i, j \in N_{entry}, i \neq j, \forall v \in V : i, j \in v, \forall k1, k2 \in K, k1 \neq k2, \forall p \in P \tag{1.at}$$

$$X^k_{i,j}, Z_i, U^k, \phi^k_{i,v}, \psi^k_{j,v}, \rho^k_{i,j,v}, \theta^k_{i,j,v}, \zeta^k_{i,v}, \xi^k_{i,v}, \mu^k_{i,j,v}, W^{k1,k2}_{i,j,v}, Y^k_{i,v,p}, C^{k1,k2}_{i,j,p,v}, \qquad \in \{0, 1\} \tag{1.au}$$

$$T_i, T^k_{O_p}, Q_i, E_i, DT_i, RT_i, NOT_i, NCT_i, VOT^k_{i,v}, VCT^k_{i,v}, TAT^k_{i,v}, TDT^k_{i,v}, \qquad \in \mathbb{Z}^{0+} \tag{1.av}$$

The objective function, eq. (1.a), maximizes the profit. The first term constitutes the revenue per passenger transported minus the operational costs, per trip operated. Subsequently, the cost of a delayed arrival, the penalty cost for rejected trips, the cost for each vehicle dispatched and the energy cost are subtracted. Constraints (1.b) to (1.f) implement the general PDP/DARP constraints. However, constraint (1.b) has been slightly modified from the general variant: not all customer trips have to be satisfied. In case a request is not served, it is considered rejected, which has the aforementioned penalty in the objective function. Constraint (1.c) ensures that the same vehicle operates both the pick-up and delivery part of a customer request. Constraint (1.d) takes care of the flow continuity at each node, while constraint (1.e) makes sure that a vehicle either leaves the depot for customer trips or is allocated an unused status. Constraint (1.f) ensures that vehicles that leave the depot also return to the depot.

Constraints (1.g) to (1.o) enforce all time-related constraints. The resolution of all time-related decision variables is 1 minute. Constraints (1.g) and (1.h) are the earliest time to leave the depot and the latest time to

arrive back at the depot, respectively. For pick-up and delivery, the time windows are noted in constraints (1.i) and (1.j), respectively. The time between the earliest pick-up time at i and the earliest delivery time at corresponding node j is based on the travel time between the nodes, $t_{i,j}$. The time precedence constraints are provided for the trip following departure from the depot in constraint (1.k), for the return to the depot in constraint (1.l) and for all other trips in between in constraint (1.m). Because of the assumption that charging takes place before and after each flight, and because charging takes place at the pad, no separate service time is added. Boarding and disembarking can take place concurrently with charging (or when the vehicle is idling). As a result, the start of service time is the actual departure/arrival of the vehicle from the pad. Constraint (1.n) is used to compute the recharge time to a full battery charge bat with charging rate h at each charging station. Constraint (1.o) is the time delay between the actual arrival time and the earliest possible arrival time at a delivery node.

Constraints (1.p) to (1.s) are the vehicle's load constraints, determining whether to add, subtract or maintain the load. The limits for the vehicle's load are constrained in constraint (1.t). The load constraints force the model to either accept or rejected the entire customer party δ, i.e. multi-passenger customer requests cannot be split across multiple vehicles. The energy levels are determined using constraints (1.u) and (1.v), while a full battery charge is imposed by constraint (1.w).

Constraint (1.x) through constraint (1.at) together are used to enforce the vertiport capacity. First of all, constraints (1.x) to (1.ab) are used to determine the exact times at which a node is occupied or cleared by a vehicle, for different node sets. This distinction is required because the variable T_i does not always constitute the exact timing at which a node is occupied/cleared. For instance, when there is an idle period following a charge. Also, these constraints add the vertiport index coinciding with the node.

Constraints (1.ac) and (1.ad), if the two nodes in variable $X^k_{i,j}$ correspond to different vertiports, indicate the entry and exit nodes of a vertiport, respectively. If a node is either an entry or exit node, constraints (1.ae) and (1.af) consequently set the vertiport occupying and vertiport clearing time, respectively.

Based on the vertiport occupying and clearing times, a match can be made between an entry and exit node, which then corresponds to a single trip to a vertiport by a vehicle. A trip can be comprised of multiple nodes in between entry and exit. For each entry node i and exit node j, for each vehicle, it is determined how many entries at the same vertiport take place later in time. If the number of entries that takes place after entry node i is the same as the number of entries after exit node j, for the same vehicle at the same vertiport, nodes i and j match to the same vehicle trip at that vertiport. Constraints (1.ag) and (1.ah) indicate if the arrival times at two entry nodes at the same vertiport for the same vehicle take place after one another. The same holds for constraints (1.ai) and (1.aj), which indicate if an entry node arrival takes place after departure from an exit node at the same vertiport. The ϵ ensures the correct value of the indicator in case the two variables under comparison are equal to each other (when both are 0). Constraints (1.ak) and (1.al) count the number of entry node arrivals that occur after arrival at an entry node and following the departure from an exit node, respectively. If the summations match exactly, for the same vertiport and vehicle, the entry and exit node form a pair. This pair match is implemented in constraint (1.am).

For such a match, which constitutes a single trip to a vertiport, constraints (1.an) and (1.ao) indicate the trip arrival and trip departure time, respectively. Note that constraint (1.am) should also count matches where the summation of both the entry and exit nodes is 0. This happens when it is the last visit of a vehicle at a certain vertiport. The summations from constraints (1.ak) and (1.al) may also yield equal values for entry-exit node-combinations where one or more nodes are not actually served (which have a VOT or VCT of 0. All non-served nodes have the same number of nodes that happen later, i.e. the sum of all nodes that are served). Consequently, it is necessary to filter out matches that belong to nodes that are not served. This is done by imposing that $\mu^k_{i,j,v} = 1$ if, and only if, node i is an entry node that is served, i.e. if $\phi^k_{i,v} = 1$.

Constraint (1.ap) assigns a dock to each vertiport trip. Constraints (1.aq) and (1.ar) are used to indicate if two trips overlap in time at the same vertiport while constraint (1.as) provides the indicator when two trips are assigned to the same dock. Lastly, constraint (1.at) guarantees that the same dock at the same vertiport can only be assigned to two trips, for different vehicles, if they do not overlap in time, thereby enforcing the vertiport capacity.

Finally, constraints (1.au) and (1.av) define the binary and integer variables, respectively.

3.3.1 Edge Elimination

In order to reduce the problem size and improve performance, unfeasible edges are discarded before the optimisation process. These include edges where the distance is larger than the vehicle's range or where the energy required exceeds a full battery's capacity. In case the sum of the demand of two pick-up nodes exceeds the maximum vehicle capacity, the corresponding edge is eliminated as well. Furthermore, the time windows for customer requests are used. If the latest arrival time at node i is earlier or equal to the earliest departure time at node j, there can be no flow $j \rightarrow i$, and the corresponding edge is removed. Consequently, only feasible edges are considered.

3.4 Computational Solving

The mathematical optimisation model presented in this section was translated into the *Python* coding language, with *Python* version 3.7. The LP-model files were generated using the *Gurobi 9.1* solver. An academic license was used through the institution. Small scale instances were optimised locally with *Gurobi* on an HP-notebook, featuring an Intel i5-8250 quad-core, eight-thread CPU and 8 GB of RAM. Large scale problem instances were solved on a cluster with an AMD EPYC 7551 processor, which features 64 cores with 128 threads and 256 GB RAM. Optimisation on the cluster was performed by solving the LP-files generated by *Gurobi*, using the *ILOG CPLEX Optimizer 12.10* by *IBM*.

3.5 Case Study

Currently, UAM ride-sharing is not in operation yet. Hence, no real-world scenario is available. Therefore the model will be applied to a case study based on literature. This subsection will present the parameters of the case study.

3.5.1 Vehicle and Battery

The vehicle model is based on the CityAirbus eVTOL vehicle, currently under development by Airbus (2020). This type has also been used as the vehicle in the work of Shihab and Wei (2020). The vehicle's performance matches the majority of proposed UAM operations within the scope of inter-city operations discussed in this research. The CityAirbus, depicted in fig. 2, has four dual-rotor ducted fans and can accommodate up to 4 occupants. Pilot-less operation is presumed such that the entire vehicle's capacity is available to passengers. The cruise speed is 120 *km/h*.

Figure 2: Image of the CityAirbus eVTOL aircraft (Airbus, 2020).

The battery capacity for the vehicle and charging rate are taken from literature and set at 110 kWh and 150 kW, respectively (Shihab and Wei, 2020). The discharging and recharging rates are assumed to be linear. The computation of energy required is based on the cruise power P_{cruise} and the travel time, and independent of the number of passengers on board. From literature, the power for non-cruise flight phases is expressed as a ratio of the cruise power. A typical mission profile is used for the duration of non-cruise flight phases. The cumulative time required for taxiing, taking off, climbing and landing is added to the flying time between vertiports in the calculation of the travel time. No reserve or loiter time is considered. The mission profile is depicted in fig. 3.

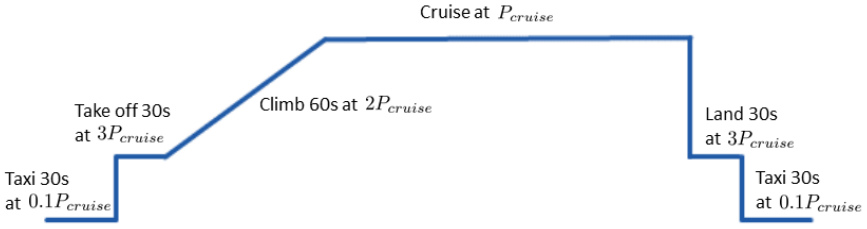

Figure 3: Image showing a typical eVTOL mission profile with the time and cruise power ratio for different flight phases. Adapted from Kohlman and Patterson (2018).

3.5.2 Network and Infrastructure

Various cities around the world have been presented to be a location for UAM operations, such as New York City, San Francisco, Singapore and Dubai. A location that has repeatedly been used in literature is the Dallas-Fort Worth (DFW) area in Texas, US (Bosson and Lauderdale, 2018; Dunn, 2018; Kleinbekman et al., 2020; Kohlman and Patterson, 2018; Shihab and Wei, 2019; Uber, 2020). This location mimics a large metropolitan area with various suburbs and a central business district (CBD), making it an ideal scenario to simulate commuter travel. The network of vertiport locations is often modelled according to the hexagon representation by Patterson et al. (2018) where there is a central vertiport in the CBD and 6 (or more) vertiports in a hexagon shape around the central vertiport. This model can be expanded with fewer or more 'rings' to simulate a different number of vertiports. For this work, the network size has been reduced to half a hexagon with five vertiports, in addition to the depot. The network is shown in fig. 4.
Between each vertiport, the Euclidian distance is computed. The travel time is calculated based on this distance and the cruise speed, and a fixed amount of time is added to account for take-off, approach and landing procedures based on figures by Patterson et al. (2018). The energy required is computed based on the travel time and nominal cruise power inferred from literature, and is also corrected for some energy-intensive flight phases during departure and arrival.

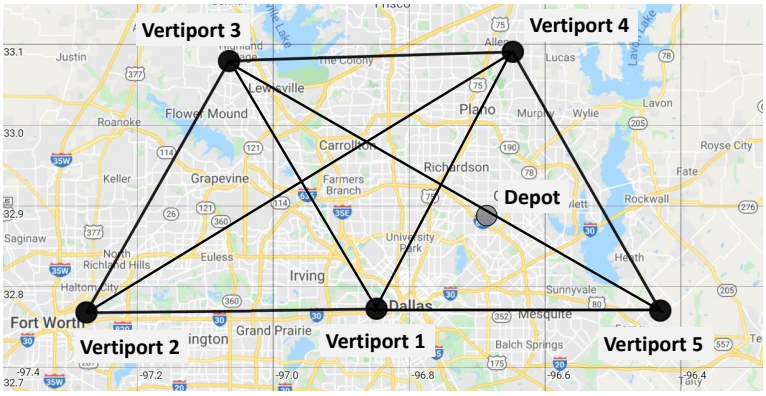

Figure 4: Image showing the vertiport and depot locations, as well as the possible customer routes, overlaid on the map of the Dallas-Fort Worth area. *Image created with maps.co/gis* .

3.5.3 Customer Demand and Price-setting

For simulating flight demand, a model by Kohlman and Patterson (2018) is used, which is based on general demand trends by Uber (2020). The demand is modelled by merging multiple normal distributions centred around the peak hours during the day. In this case, the distributions have peak hours (mean) at 08:00h, 12:00h and 16:00h, with variances of 2, 6 and 2 hours, respectively. The combination of these curves is then normalised to a maximum of 1. A scaling factor can be used to influence the total number of flights requests, as well as the relative demand between vertiports. For generating a request, a random number between 0.0 and 1.0 is compared to the value on the probability curve at a certain time instance (the operational period is discretised using the time resolution of 1 minute). If the random number is less than the value on the curve, a trip request is generated. The probability function of the demand distribution is depicted in fig. 5, for a full 24 hours. The customer requests themselves can take place between 07:00h and 21:00h.

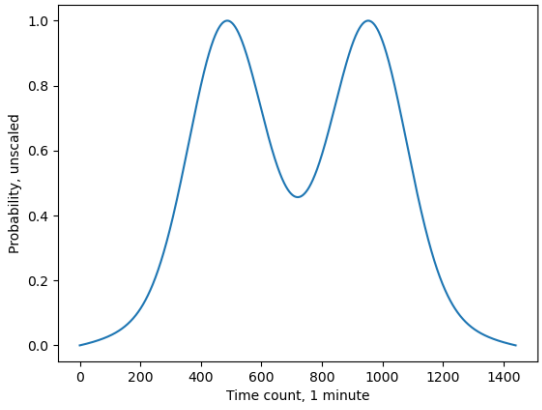

Figure 5: Graph showing the unscaled demand probability for a 24-hour period.

The fare and operating cost figures are based on figures for Uber Elevate (Garrow et al., 2019; Holden et al., 2016). The fare is set at 3.56 *USD/pax/km*, and the operating cost is adjusted for a 4-seater aircraft to 3.06 *USD/km*. The energy cost is separate from the operating cost and amounts to 0.12 *USD/kWh*, which is a commonly used value in literature (Duffy et al., 2017; Holden et al., 2016; Kohlman and Patterson, 2018). The vehicle dispatch cost, c_v, is based on the operating cost for the average distance of positioning return flights to and from the depot, the associated energy cost and a factor to account for maintenance check costs once a vehicle has flown (Brown and Harris, 2020). This totals 268 *USD/vehicle dispatched*. The penalties for excess time delay and rejected trips have been arbitrarily chosen to be 1 *USD/minute delay* and 4 *USD/rejected trip*, respectively. Fare and cost predictions are highly dependent on the size of operation, vehicle utilisation and the number of customers. The monetary values have been chosen conservatively to match the early stage, small scale condition treated in this research.

3.5.4 Scenarios

In order to study how the model responds to different demand distributions, various scenarios are applied to the model. For each scenario, there are a total of 30 customer requests. This number has been found to be an upper limit to solve the algorithm in a reasonable amount of time. The generation of customer requests for each scenario follows the temporal demand distribution shown in fig. 5. However, for some scenarios, the relative demand of certain vertiport(pair)s is adjusted. The seed is modified until the set of customer requests is representative, in time and/or location distribution, for one of the following five scenarios.
First of all, there is the unrestrained scenario where the relative demand is equal between all vertiports. This simulates an "arbitrary" situation, without priorities imposed. The second scenario comprises "clustering" of flights. After determining the OD-pair of a customer request, the subsequent request has double the probability of selecting the same OD-pair. This simulates a case where a certain OD-pair temporarily has a higher demand. In the third, "balanced", scenario, customer requests that take place within (+/-) one hour of the morning peak at 08:00h result in a customer request in the opposite direction within one hour of the afternoon peak at 16:00h. This resembles a case with return journeys the same day. Note, however, that the model currently does not require that the coupled requests both have to be satisfied. The fourth and fifth scenarios mimic the second and third scenario, respectively. However, as an additional adjustment, vertiport 1 (see fig. 4) has twice as much relative demand inbound until noon and twice as much relative demand outbound after midday. These "directed" scenarios resemble a commuter situation, with primarily flow from the suburbs to the city centre in the morning, and the reverse in the afternoon.
The optimisation of the five scenarios presented heretofore will be performed with both a time window of 5 minutes and 10 minutes. A third option of 15 minutes was initially considered as well. However, early results showed that the algorithm had greater difficulty converging to an optimal solution due to the increased solution space. Larger time windows constitute a more relaxed problem, so this model's version should offer the same or better results. However, in the allotted runtime, improvements in the objection function were not always attained. With a nominal flying time of 20-40 minutes, a 15-minute wait is also considerable. Therefore the additional variation of a 15 minutes time window has been rejected.
The above five scenarios and two time window versions lead to ten different optimisation instances.

4 Results

This section will present the results of the optimisation runs for the 10 scenario versions from the case study. First of all, in section 4.1 the profit from the objective function, and the most important performance indicators of vehicles dispatched, customer requests and passengers served will be analysed. Subsequently, in section 4.2, various underlying key performance indicators (KPI's) will be evaluated to compare the different scenarios. A tabulated overview of all KPI values is provided in appendix A. Lastly, in section 4.3, a sensitivity analysis for various penalty cost coefficients, fleet size and vertiport capacity is discussed.

4.1 Profit, Vehicle Use and Request Fulfilment

Before proceeding with the comparison of the actual data, we would like to underline that the scenarios were not solved to optimality due to computational issues. This issue will be elaborated upon further in section 5. To illustrate the optimality gap, in this subsection whiskers indicate the upper bound for the objective function value in graphs depicting the profit.

4.1.1 Profit

The objective function of the model aimed to maximise the total profit of a day's operation. The profit is thus the first parameter to look at. From fig. 6 it can readily be observed that there is a modest to no effect of an

increased time window, and thus theoretically a more relaxed formulation, on the profit. This characteristic may indicate that time delays on arrival were not a substantial constraining factor in these scenarios. Only in the clustering directed scenario does the version with a wider time window have a noticeable larger profit. As will be detailed later, this profit increase is caused by pooled trips that the wider time window allows. Since the clustering scenarios offer a higher probability that consecutive trip requests are between the same OD-pair, it was expected that wider time windows could have a positive effect through pooling trips.
Between the different scenarios, there is more deviation in the profit. The arbitrary scenario has the lowest profit. This scenario also has the lowest number of passengers, cf. table 5, which may be a contributing factor. When the original, arbitrary, demand distribution is adjusted to both clustering and balancing of flights, the profit increases.
It is furthermore apparent that the profit between the directed and undirected balanced scenarios is very similar, while there is a greater disparity between the undirected and directed clustering scenarios. In the clustering scenario, the undirected version has more passengers to transport than in the directed version. Meanwhile, the balanced undirected scenario has a smaller total number of passengers than the balanced directed scenario. Since the difference can not readily be attributed to the different number of customers, the directionality of the customer demand may be an influential factor for the profit, where a directed flow has a depreciating effect on the profit. When comparing the other two distribution conditions, the results do indicate that balancing the flow in- and outbound has a beneficial effect on the profit vis-à-vis clustering.

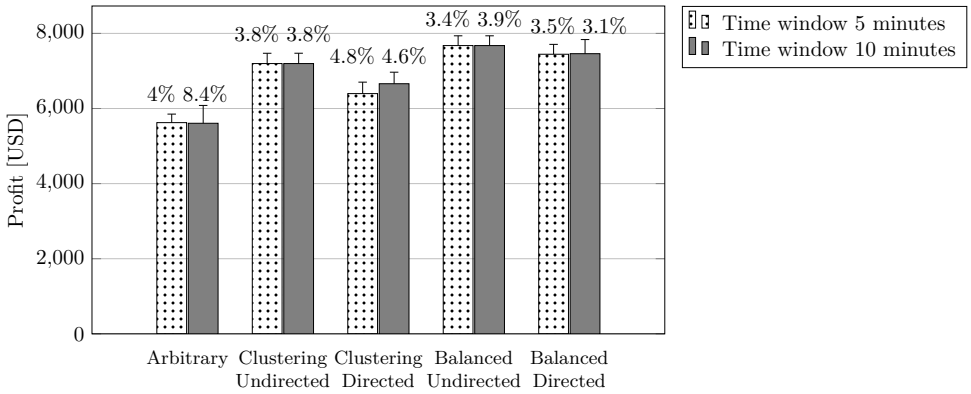

Figure 6: Graph showing the profit for each scenario and version. Whisker and percentage indicate the optimality gap.

4.1.2 Vehicles Dispatched, Customer Requests and Passengers Served

Dispatching vehicles is a prerequisite in order to serve customer requests. However, the model does not impose the use of the entire available fleet. A trade-off is made between the dispatch cost of a vehicle and the revenue it can gain from using that vehicle. For each scenario, the number of vehicles dispatched in the final result is shown in table 4. This table also notes the abbreviations for the different scenarios that will used going forward. There were no differences between the time window versions in each scenario. As can be seen, two out of the five scenarios utilise the entire fleet available. Hence, three scenarios are not saturated fleet-wise and have the potential to transport more passengers. As will become apparent hereafter, this potential depends on the customer request parameters and the potential revenue.

Table 4: Table showing the number of vehicles dispatched, out of a fleet of 6, for each scenario.

Scenario	**# Vehicles dispatched**
Arbitrary (A)	5
Clustering Undirected (CU)	6
Clustering Directed (CD)	5
Balanced Undirected (BU)	5
Balanced Directed (BD)	6

Closely related to the profit is the number of customer requests served and, consequently, the number of passengers transported. These are the main drivers for the revenue. All scenarios feature 30 customer requests. The total number of passengers divided over those 30 requests varies per scenario. The numbers are provided in table 5.

Table 5: Table showing the total number of passengers who may be transported for each scenario.

Scenario	# Passengers
Arbitrary	68
Clustering Undirected	85
Clustering Directed	74
Balanced Undirected	79
Balanced Directed	85

The results for the customer trips served and passengers transported are depicted in fig. 7. Across the board, in most scenarios, the vast majority of trips is served except for the clustering undirected and balanced directed cases. Comparing the graph with the information from table 5, it also transpires that the trips that are not served only have a single customer, except for the clustering undirected cases where 5 rejected trips resulted in 9 spilt passengers. In this scenario, it became apparent that a single trip request of 4 passengers was not compatible with other constraints.
In two scenarios, clustering directed and balanced directed, the version with a longer time window transports one additional customer in total. In the balanced directed case, also one additional customer flight is operated. In the clustering directed case, the total number of customer flights operated remains the same as with a time window of 5 minutes. This implies that a shift in flights has taken place instead of operating an incremental flight, which is consistent with the aforementioned interpretation that operating an additional flight for a single customer is not always worthwhile.
Moreover, taking into account the vehicle dispatch information from table 4, the fact that not all customers are served in scenarios when the entire fleet is not yet utilised points to the notion that (multiple) single-customer trip requests are not worth dispatching an additional vehicle for with the current cost and revenue values. Hence, there may be an opportunity to satisfy more customers by adjusting the pricing or lowering the attached dispatch cost if an operator so desires.

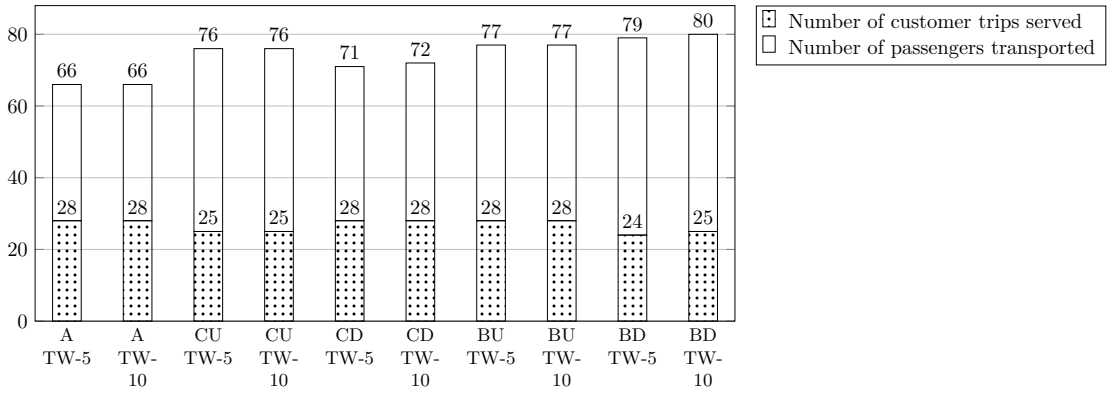

Figure 7: Graph showing both the number of customer flights operated and passengers transported for each scenario and version. 'TW' indicates the time window.

4.2 Underlying Performance Indicators

In addition to the profit from the objective function, the number of vehicles dispatched, and the number of trips and passengers served, many more performance indicators were established from the results. A selection of KPI's is presented in more detail in this subsection. These indicators have been chosen because they either: offer a comparative insight into the different scenarios when applied to the model, or constitute an important feature of the model.

4.2.1 Stage Distribution

The eVTOL vehicles can be engaged in various stages during operation: they can be flying customers, flying empty (positioning flights), charge at a vertiport or idle at a vertiport (waiting before the next pickup or flight). The division between these stages, in terms of total time, is depicted in fig. 8. It becomes clear that the majority of the time, vehicles are either idling or charging. It also signifies that, under the current assumptions, charging takes more time cumulatively than the actual customer flying. Neglecting the distribution of flights and the vertiport capacity for a moment, the considerable amount of idling time also indicates that the available fleet may well be able to operate more than the current 30 customer trips during a day. A faster charging rate may further improve an aircraft's utilisation, but the results *in-casu* do not clearly support this. In relation to the

number of customer trips served as presented in fig. 7, the charging time does not appear to be a limiting factor as the rejected trips most likely stem from the aforementioned cost-revenue trade-off.

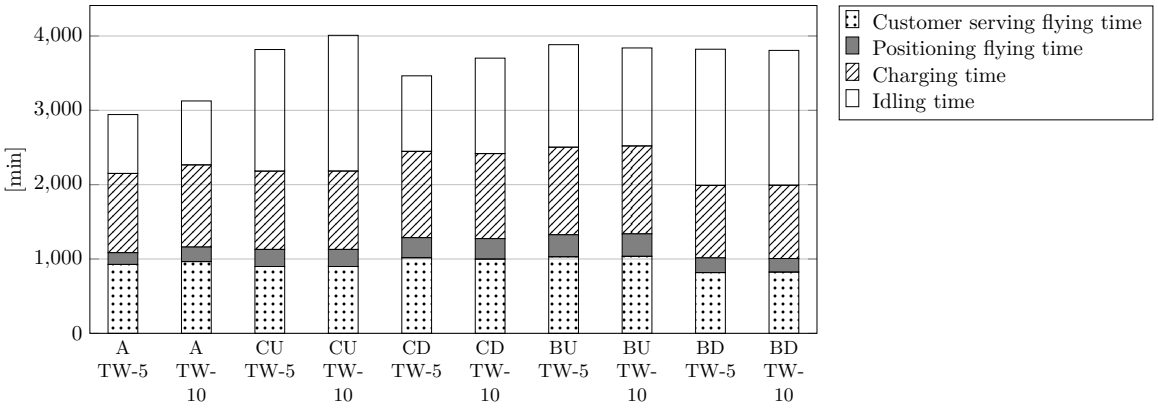

Figure 8: Graph showing the cumulative time spent on customer flying, positioning, charging or idling for each scenario and version. Excludes positioning and idling segments immediately before/after depot-flights.

The model trades off the missed revenue in case of idling at a vertiport for the next flight that departs from the same location or flying empty to another vertiport in order to start another service there. Diving deeper into the positioning flights, fig. 9 details both the number of positioning flights and the cumulative distance flown empty during positioning for all runs.

For all scenarios, the number of positioning trips is consistent for the different time window versions. A comparison of the total positioning distance indicates that the arbitrary and balanced directed scenarios fly noticeably less for positioning. For the former scenario, this is somewhat unexpected. When the flight direction is random, there is a lower probability that an arrival vertiport can serve as a departure vertiport for a subsequent trip, thereby not requiring a positioning trip. Despite this expectation, the arbitrary scenario manages to transport the vast majority of trips and passengers (cf. fig. 7) with the smallest number of positioning trips.

The relatively large differences between the undirected and directed variants of the balanced scenario are remarkable as well. Certainly when remembering that the directed scenario utilises 6 vehicles and the undirected scenario 5. These differences may be explained by the fact that in the directed case, flow is primarily directed to the central vertiport in the first wave in the morning. From the central vertiport, the distance to all other vertiports is equal. The vehicles can subsequently operate other flights in between the morning and afternoon wave before returning to the central vertiport again for the return flights. Once again, regardless of from which vertiport the vehicles fly, the distance to the central vertiport is the same. Without directing flights, vehicles may be at one of the peripheral vertiports in the early afternoon. If they have to position to another peripheral vertiport to operate a return flight, the flying distance may be larger. The distance flown is an important parameter because it directly affects the flying time for positioning flights and, subsequently, also the charging time and charging cost. As a result, it can be argued that a use case with a balanced directed or even arbitrary demand distribution is more efficient: fewer resources are spent on positioning.

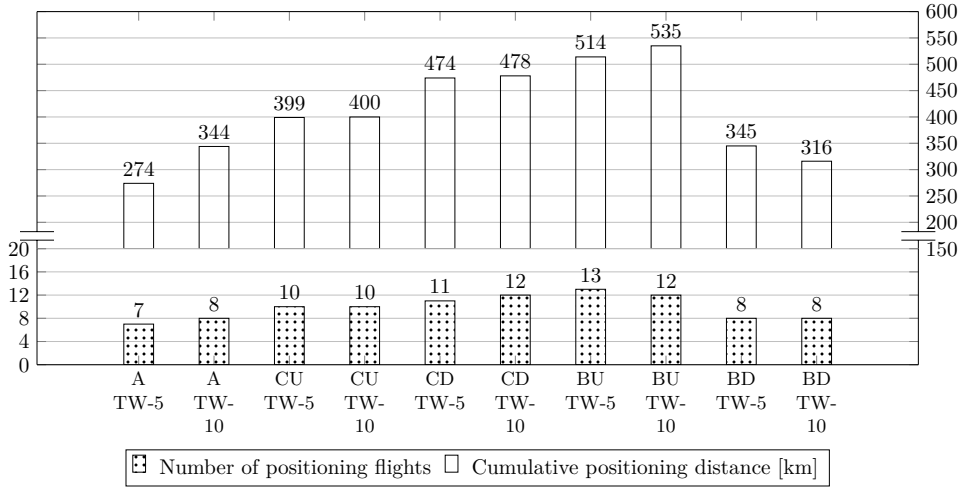

Figure 9: Graph showing the number of positioning flights and the cumulative distance for each scenario and version.

4.2.2 Time Delay

An important attribute from the customer's point of view is the time delay at the destination, with respect to the earliest specified time of arrival. Figure 8 portrays the number of instances where there was a delay on arrival, as well as the cumulative delay in minutes for each case. On the whole, time delays are modest. In relation with fig. 7, a higher cumulative time delay seems correlated with the increased number of customer flights and passengers transported in the clustering directed and balanced directed scenarios. After looking at the route logs, the additional flight in the balanced directed scenario cannot be explained by the longer time window alone since there is no time delay in the end result exceeding the initial 5-minute time window. The in-depth review reveals further that one customer trip has been operated with the maximum allowed delay of 10 minutes in both the arbitrary and clustering directed cases. Only in the clustering directed case this has actually had a noticeable positive effect on the profit.

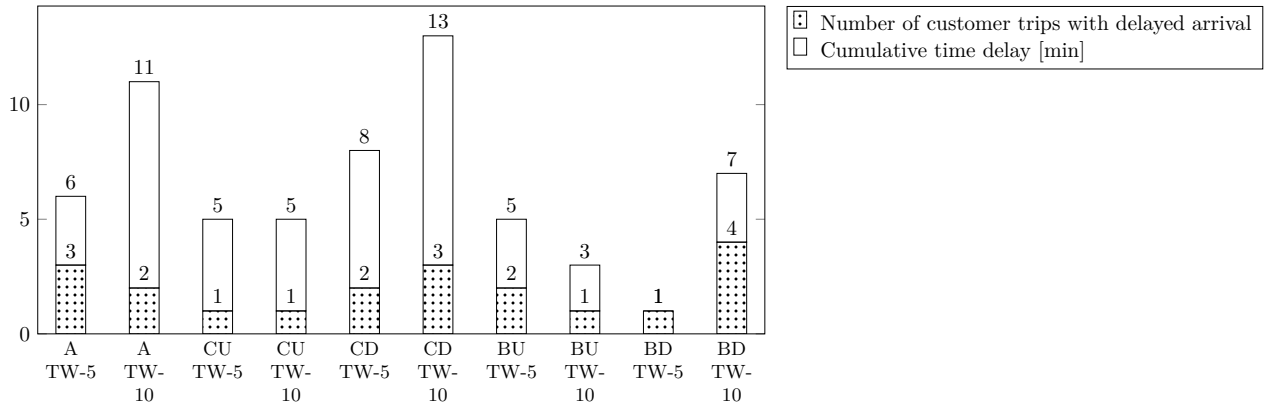

Figure 10: Graph showing both the number of flights with an arrival delay and the cumulative delay in minutes for each scenario and version.

4.2.3 Pooling

The prospective ride-sharing aspect of UAM operations is incorporated in the model by offering the opportunity to pool passengers: if the time window and vehicle capacity constraints are not violated, two or more customer requests can be served by a single flight. The number of instances when this happens as well as the total number of pooled passengers are presented in fig. 11. The results demonstrate the model's ability to combine multiple customer trips in a single flight. Pooling does not occur in every scenario. It was expected that pooling occurs more in the clustering scenarios and with a larger time window. In the clustering directed case, the single pooled trip indeed makes use of the full 10-minute delay allowed in the TW-10 version, demonstrating the added value of the wider time window. However, the other pooling occurrences do not support a conclusion if a change in time window or demand distribution has a clear effect on pooling ability. A reason for this may be the small number of customer trips spread over a multi-hour time horizon. Despite adjusting the demand distribution, the probability of having 2 trips between the same OD-pair within 5-10 minutes of each other is low with only 30 samples. This is a computational limitation that will be further treated in section 5.

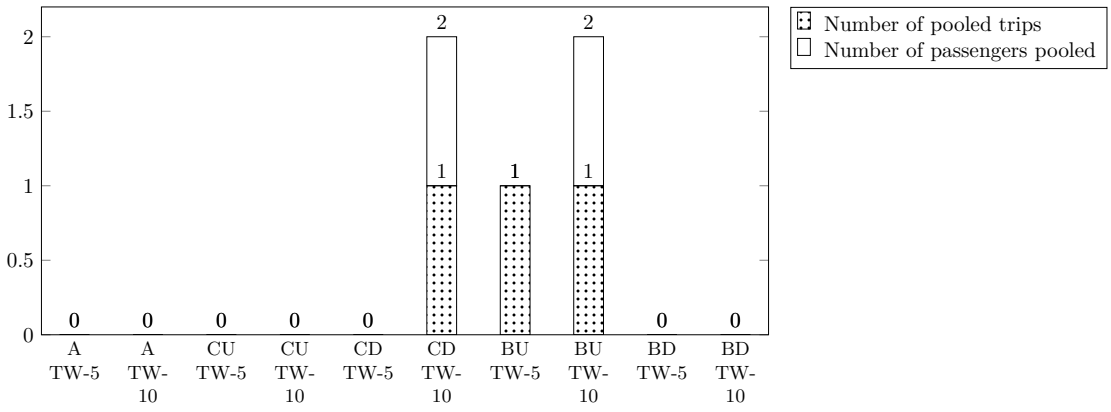

Figure 11: Graph showing how many times multiple customer trips are pooled in one flight and how many passengers have been pooled as a result for each scenario and version.

4.2.4 Pad Saturation

Lastly, it has been checked how many minutes the pad capacity of the 5 vertiports was saturated cumulatively during the day's operation. The results are plotted in fig. 12. The arbitrary runs clearly show the least amount of pad saturation. This behaviour is expected, as flow is not directed towards a certain location and/or within a certain time span. On the other hand, the clustering undirected scenario has around 3 hours of cumulative pad saturation, illustrating a possible bottleneck. It was also expected that the directed scenarios would exhibit noticeable pad saturation, especially at the focus, central vertiport. After reviewing the pad saturation per vertiport, this proved not to be the case. Practically, this implies that in the case of a directed demand distribution, the focus vertiport does not necessarily require additional infrastructure capacity to accommodate more vehicles than the peripheral vertiports.

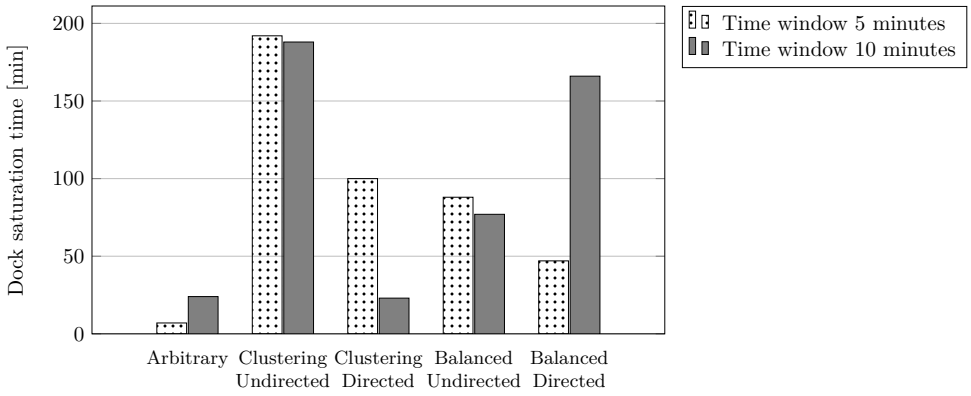

Figure 12: Graph showing the cumulative time when the pad capacity of any vertiport is reached for each scenario and version.

4.3 Sensitivity Analysis

In order to evaluate how the model behaves to changes in some input parameters, a sensitivity analysis has been performed on the balanced undirected scenario with a time window of 10 minutes. This focused, first of all, on various coefficients of the objective function. Additionally, the incremental effect of fleet and infrastructure resources has been investigated.

4.3.1 Penalty Coefficients

As presented in section 3, the objective function includes several penalty factors. The values for the time delay, rejected trip and vehicle dispatch penalties have been arbitrarily chosen for the main scenario optimisation runs. In order to investigate the effect of the penalty coefficients, one of the scenarios has been optimised repeatedly with an altered value for one of the penalty coefficients. The original and revised values are provided in table 6.

Table 6: Table showing the original and revised values for three penalty coefficients of the objective function.

Penalty/version	**Original value**	**Alteration 1**	**Alteration 2**	**Alteration 3**
Time delay c_t	1 *USD/min*	4 *USD/min*	10 *USD/min*	20 *USD/min*
Rejected trip c_r	4 *USD/trip*	40 *USD/trip*	80 *USD/min*	200 *USD/trip*
Vehicle dispatch c_v	268 *USD/vehicle*	536 *USD/vehicle*	1072 *USD/trip*	-*

*A higher dispatch cost than 4 times the original value was not considered realistic and was therefore not implemented.

The resulting profit values are indicated in fig. 13. Once again, the scenarios could not be solved to optimality. The gap between the best incumbent and the upper bound is indicated in the figure for reference. From the values, the profit does not show a drastic decline for increasing rejected trip penalty values. In fact, the solution for all three adapted c_r values is exactly the same as for the original parameter value, and the decrease in the objective function value is solely due to the adjusted coefficients. This attribute would indicate that the rejected trips are constrained by other factors and not by a monetary trade-off through the objective value coefficient.

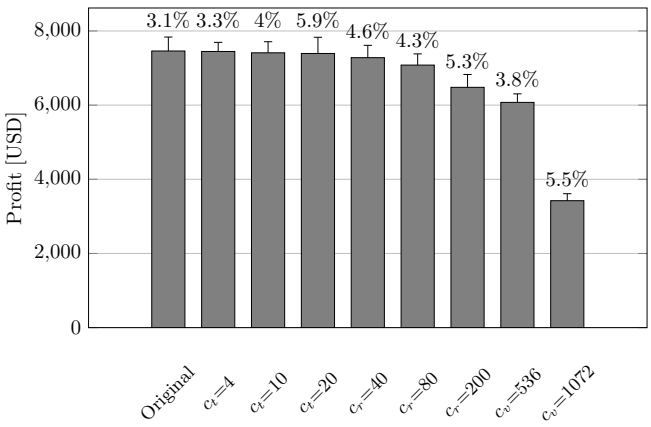

Figure 13: Graph showing the profit for the original and adapted coefficient versions of the balanced directed scenario with time window of 10 minutes. Whisker and percentage indicate the optimality gap.

For the other coefficients, the solution did change. In order to evaluate the consequences, a closer look is taken at some KPI's. In table 7, the values for the revised time delay penalties are shown. As can be seen, from the first increased penalty cost variant, both the number of delayed trips and the cumulative time delay decrease. At the same time, the number of trips and passengers served only decreases, with one, when the time penalty reaches 20 *USD/min*. Hence, the model finds a solution where (almost) the same number of passengers can be served with less time delay. From the number of positioning trips, it can be seen that more positioning flights are required to accomplish this. The energy cost for flying empty to position for a flight without time delay is thus cheaper than operating other flights with a time delay.

Table 7: Table showing a selection of KPI's for the original and revised runs for the time penalty coefficient.

KPI/version	**Original,** $\mathbf{c}_t = 1$	$\mathbf{c}_t = 4$	$\mathbf{c}_t = 10$	$\mathbf{c}_t = 20$
# Trips with time excess	4	1	2	1
Cumulative time delay [min]	7	1	2	1
# Customer trips served	25	25	25	24
# Passengers transported	80	80	80	79
# Positioning trips	8	10	13	12

If the dispatch cost per vehicle increases, by factors 2 and 4, the number of vehicles dispatched decreases, as is shown in table 8. If the original dispatch cost is doubled, one vehicle less is dispatched. The effect on the number of trips and passengers served is small. In case the vehicle use cost is quadruple the original value, only half of the fleet is dispatched. Nevertheless, 20 out of 30 requests are still served, transporting more than 75% of all passengers. Idling time also greatly reduces, showing an increase in utilisation and efficiency. These results indicate that, for 30 customer requests, less than an entire fleet of 6 aircraft can still serve a vast majority of flight requests. An operator can weigh whether or not the cost of an extra vehicle is worth the incremental flights and customers served.

Table 8: Table showing a selection of KPI's for the original and revised runs for the vehicle dispatch penalty coefficient.

KPI/version	**Original,** $\mathbf{c}_v = 268$	$\mathbf{c}_v = 536$	$\mathbf{c}_v = 1072$
# Vehicles dispatched	6	5	3
# Customer trips served	24	23	20
# Passengers transported	80	76	66

4.3.2 Fleet and Infrastructure Resources

Next to the penalty coefficients, different fleet sizes and vertiport pad capacity figures have also been the subject of a sensitivity analysis. Since the problem size greatly scales with the fleet size and computational performance formed a limiting factor in the original optimisation process, only smaller fleet sizes than the standard 6 vehicles have been considered: from 1 to 5 vehicles. For the pad capacity, variations with 2 and 4 pads have been optimised versus the original 3 pads. This approach enables a review of the incremental profit of additional vehicles and pads.

The resulting objective function values are indicated in fig. 14 with the profit values of the fleet variations on the left of the reference value, i.e. the original scenario, and the profit values of the pad capacity variations on the right.

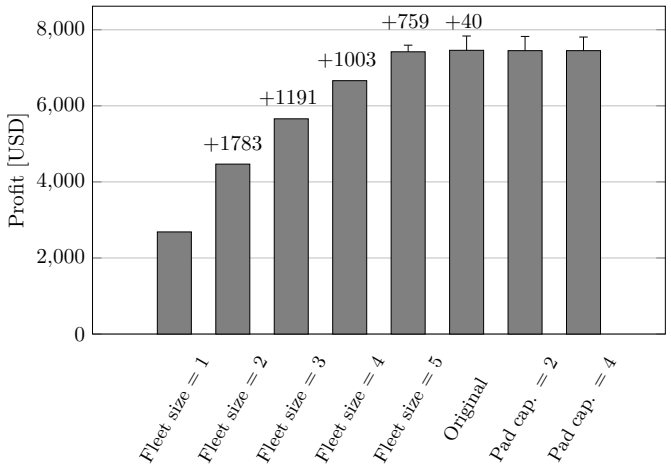

Figure 14: Graph showing the profit for the original and adjusted resources versions of the balanced directed scenario with time window of 10 minutes. For the vehicle variations, the values above each bar indicate the incremental profit increase versus the variation with one vehicle less. Whiskers indicate any optimality gap.

As can be seen in fig. 14, the incremental profit value of an additional vehicle is the greatest from 1 to 2 vehicles. The third, fourth and fifth vehicle still have a noticeable added value to the profit. However, the sixth vehicle has a negligible profit increase. From table 9, it can also be seen that the additional number of customer requests and passengers served between the version with 5 and 6 vehicles is just 4 customers in 2 trips. Between the other, smaller fleet sizes, the number of additional customer requests and passengers served is considerably higher. This is in line with the expected behaviour. With a small number of vehicles, each additional vehicle still has the potential to find a routing consisting of multiple customer trips, spread over the entire day. Once more vehicles are dispatched, there is less potential to combine multiple of the remaining trips in a route. Either because these unserved trips do not fit practically within a larger routing or because other trips are more attractive revenue-wise. These initially rejected trips can subsequently still be served (individually) if their revenue exceeds the dispatch costs of an additional vehicle.
Similarly to the sensitivity analysis on the cost coefficient parameters, the variations with the input fleet size show that a fleet of 5 vehicles is still sufficient to serve the majority of customer trips and passengers, and operating a full fleet of 6 vehicles has a negligible beneficial effect on the total profit.

Table 9: Table showing a selection of KPI's for the original and revised runs for the vehicle fleet size.

KPI/version with fleet size	**Original, 6**	**1**	**2**	**3**	**4**	**5**
# Customer trips served	25	7	13	16	20	23
# Passengers transported	80	24	44	57	66	76
# Positioning trips	8	5	6	6	7	8
Idle time [min]	2596	47	333	714	1023	1619
Flying time [min]	824	309	470	539	687	793

Figure 14 shows further that the profit between runs with a vertiport capacity of 2, 3 and 4 pads is virtually equal. The solution in terms of KPI's for the run with 2 and 4 pads are identical; the original run with 3 pads varies slightly. The most important KPI's are similar across all three runs, though. This can also be seen in the data of the trips and customers served, as shown in table 10. The original scenario has a solution with one positioning trip less, which causes the slightly higher profit. This improvement can most likely be attributed to a longer runtime of the scenario.
A closer comparison of the vehicle routings of the runs with 2 and 3 pads shows that in the run with 2 pads, an idle stage of a vehicle is split between vertiports if the pad capacity would otherwise be violated. In the runs with ample pad capacity, vehicles after charging either: idle at the same vertiport before positioning for the start of the next customer trip, or position to the next vertiport immediately and idle there before the next customer flight starts. If the vertiport pad capacity is reduced, vehicles also exhibit a routing where they idle for a while at the same vertiport as where charging took place, position to the next vertiport once the pad is required for another vehicle, and continue idling the next vertiport. This shows that the model is able to

schedule positioning flights, during an idle stage, in such a way that a vehicle frees up a pad if it is required for another vehicle, thereby making optimal use of the vertiport pad capacity.
All in all, this sensitivity analysis indicates that a pad capacity of 2, at all vertiports, would be sufficient for this scenario. Additional pads do not exhibit an incremental value on the profit, and the model correctly clears a vertiport from an idling vehicle if necessary. This is an important finding for operators, which can save on the cost of additional pad infrastructure. Especially if all pads are also equipped with a charging facility, costs savings can be considerable. The unit cost of a charger can be 450,000 *USD*, according to Kohlman and Patterson (2018).

Table 10: Table showing a selection of KPI's for the original and revised runs for the vertiport pad capacity.

KPI/version with pad capacity	**Original, 3**	**2**	**4**
# Customer trips served	25	25	25
# Passengers transported	80	80	80
# Positioning trips	8	9	9

5 Discussion and Recommendations

While this paper aimed to include operational and technical features specific to UAM ride-sharing operations, it was also necessary to properly scope the research and to make trade-offs. Not all aspects have therefore been (fully) incorporated, and it is important to keep this in mind. First of all, it is highlighted once again that the model is static in nature. All customer demand data is known *a-priori* and the model is optimised a single time based on this information. No stochasticity is implemented in, for instance, the travel time, pricing strategy or customer demand. Neither are possible airspace conflicts, loitering and separation, and energy safety factors and efficiencies considered, which may have an impact on the flying time. These simplifications may limit the accuracy of the model's outcomes. In order to make the model developed in this paper more robust, it is suggested for subsequent research to adapt the formulation for a dynamic solving approach, inclusion of stochasticity for various aspects, and/or a connection to an airspace management function. The latter may be especially relevant for a use case where concurrent operations are expected at a certain vertiport, at a certain time. For instance, in the commuter scenario where flights gravitate to/from the central vertiport at peak hours. Currently, only the capacity on the ground at the vertiport is taken into account, while conflicts may also occur in the airspace surrounding the vertiport. Integration with ground transportation to/from the vertiports may also be incorporated in order to increase the practical usage significance. Including last-mile transportation in the model enables a comprehensive, door-to-door comparison between trips using only road-based operators and the combination of UAM and connecting car rides. Subsequently, this comparison can be used to examine the added value of flying over driving.

Furthermore, the model presented in this paper provides a different mathematical formulation to enforce an infrastructure capacity constraint, compared to the often used TSN-formulation where the capacity needs to be checked for each time slice. This difference in formulation does, however, come at a cost. In the current formulation, a large number of (indicator) constraints is required to impose the capacity constraint. Moreover, the pad capacity is not a right-hand side constant but instead forms an index for the decision variables. This feature means that to account for more pad capacity, the number of variables and constraints increases drastically. The same holds true for the vehicle fleet size. As the entire formulation changes, it is not readily possible to perform a sensitivity analysis using shadow prices. Instead, the entire algorithm must be re-run. Concurrently, the increased problem size can pose greater difficulty for solving the model in a reasonable amount of time and/or to an acceptable optimality gap. Future work may follow up to improve the model's performance, e.g. using a heuristic approach if a suboptimal solution is acceptable. A multi-stage solving method, e.g. using column-and-row generation, may also be of interest to cope with the capacity constraint. Currently, the constraints are imposed for all flights, while a saturated landing pad capacity may not be applicable to many, or even any, flight combinations in a certain scenario. Finding a method that identifies conflicts either *a-priori* or after running a relaxed formulation, and which subsequently imposes the capacity constraints more specifically would form an interesting extension of the model.

Additionally, it is essential to note that the problem size and the resulting runtime formed a limiting factor. First of all, because the scenarios could not be solved to optimality. This issue became more profound with a relaxed problem formulation and, as a result, more constraints and variables. The arbitrary version with a time window of 10 minutes had a worse objective function than the more constrained version with a time window of 5 minutes (cf. fig. 6), as well as quite a higher optimality gap. Even after tweaking the settings to focus on closing the gap, the TW-10 scenario did not get closer to the solution of the TW-5 scenario. However, the solution of the TW-5 version could be applied to the TW-10 version by means of a warm start, leading to a better incumbent value than what *CPLEX* yielded originally. This characteristic is also in line with the

fact that a more relaxed version of the problem should have, at least, the same objective value as the more constrained version, if not better. With a bounded runtime, this does not always materialise. Secondly, because of the relatively small number of customer requests, 30, that could be run in a reasonable amount of time, the demand generation's outcome was very sensitive to the seed used. As a result, it was practically not possible to generate two different random samples that conformed to the same scenario as presented in section 3.5.4. Thus, for each scenario, only one demand sample was used. This fact limits the statistical significance of the analysis of the results and may have affected the opportunities for pooling in particular. With 30 customer requests spread over a time span of 14 hours, there is a low probability that there are two requests on the same route and within 5-10 minutes from each other, even when the demand distribution was adjusted to steer towards this. A larger number of customer requests would allow for more samples per scenario as well as more combinations that allow for pooling. The latter case would subsequently provide a better foundation for analysing the trade-off of pooling multiple customer requests and the resulting time delay. Such insight would be valuable from a business-economical point of view. However, the current methodology and results did not allow for a meaningful sensitivity analysis on this topic.

Finally, additional reservations must be made with regards to the input parameter values used. In particular, the cost and revenue aspects. Since full-scale commercial UAM operations are not yet taking place in practice, it was not possible to take a single reference scenario where all parameters aligned. While a proper effort was made towards establishing a realistic and representative case study, the literature review revealed a wide range of future plans for UAM, with diverging figures for fleet size, battery technology, nominal trip length etc. Monetary amounts are also very much dependent on the location, phase of operation and economies of scale attained. Therefore, the absolute results should be interpreted with caution. Nevertheless, once a comprehensive UAM operation scheme is available, it would be of interest to apply the related input parameters to this model and assess the results. This does not only apply to the technical and operational parameters but also to the demand input. Lacking a real-world dataset, validation was not possible. If a representative customer demand dataset is available and implemented, evaluating the outcome can better serve as an indicator for possible practical usage.

6 Conclusions

In this paper, a scheduling model for UAM ride-sharing operations has been developed. The model's objective is to maximise the profit of a day's operation, subject to a static demand, and to provide routings for the available fleet. This research extends literature and optimisation models from ground transportation to prospective commercial flights using eVTOL vehicles. As a foundation for this, current automotive ride-sharing and future UAM operations have been reviewed and compared to determine the model requirements. The resulting formulation notably includes constraints to account for battery charging, passenger pooling and a limited infrastructure capacity on the ground. The formulation can match multiple-node visits at the same physical location and combine these in a trip. Each trip has a known entry and exit node and time. To ensure a vertiport's capacity is not exceeded at any time, vehicles are assigned to a pad during each trip. The existing model that this approach is based on only addresses a single visit per vehicle, per location. This paper provides, to the best of our knowledge, a novel formulation for the pad assignment to account for recurring trips by the same vehicle to the same location.

The model has been assessed according to a case study which represents a small-scale operation featuring 5 vertiports with 3 pads, 6 CityAirbus vehicles and 30 customer requests per day in the Dallas Fort-Worth area. Different demand scenarios are generated in order to assess the model under different possible use cases: arbitrary (baseline), clustered and balanced (outbound trip in the morning with inbound return trip in the afternoon peaks). The two latter scenarios each had versions with both undirected and directed (double demand probability for requests from the peripheral vertiports to the central vertiport in the morning and the reverse in the afternoon) flow to mimic commuter travel. The total of five scenarios were all run with both time windows of 5 and 10 minutes.

In terms of the profit value, the time windows did not make a clear difference, indicating that it was generally not a limiting factor in the scenarios. Only the 10-minute time window clustering directed version had a noticeable higher profit, caused by a pooled customer request combination allowed by wider time window. Overall, the balanced scenarios had higher profit values, with very similar objective function values between the undirected and directed version. The arbitrary scenario had the lowest profit, which may have resulted from the comparatively lower total number of passengers spread over the 30 requests.

The results indicated further that not in every scenario the entire fleet of 6 vehicles was utilised. In three out of the five scenarios, 5 vehicles were dispatched. Nevertheless, the majority of trips and passengers were served. Trips that were not served in the scenarios generally consisted of a single passenger. Hence, this indicated that the revenue of (multiple) single-customer trips does not outweigh the dispatch cost for an additional vehicle at the current cost and revenue structure. In case operators want to serve single-customer trips, an additional charge to the customer could be considered in order to steer the model towards operating more single-customer

requests.
A comparison of the cumulative time spent in the various stages yielded the insight that actual flying time constitutes a relatively small proportion of the time in operation. The majority of the day, vehicles were idling or charging, signifying that the vehicle's operational use is not close to saturated. Subject to the vertiport capacity and demand distribution, more customer requests may be served by the same fleet size. Additionally, if charging rates are faster, utilisation may also be improved.
A review of the number of positioning flights and the cumulative positioning distance showed that the arbitrary and balanced directed scenarios involved noticeably fewer positioning flights than the other scenarios. Since fewer positioning flights also involve less flying, recharging time and associated costs, this is an important performance indicator for operational efficiency. A use case with a similar demand distribution may, therefore, be favoured by an operator. On the other hand, the undirected version of the balanced scenario featured considerable more positioning than other demand distributions.
The time delay was investigated by determining both the instances and cumulative delay in minutes. In general, the time delays constituted only a few minutes. A longer allowed time window did not have a considerable impact on the effective cumulative time delay. In two scenarios, the larger time window actually enabled more customer flights to be flown since the time delay of some flights exceeded the standard 5 minutes. Thus, the model traded off the additional time delay penalty and the extra revenue in favour of the latter with the standard penalty cost. However, consistent with the aforementioned profit review, a wider time window and thus a longer allowable time delay only resulted in a noticeable profit increase in one scenario.
Passenger pooling occurred only in three of the ten cases. In one case, the entire 10-minute time window was required for pooling, but a general relation between larger time windows and increased pooling could not be established. The limited amount of pooling is attributed to the relatively small number of customer requests compared to the time horizon of the case study. Most importantly, though, these results show the potential of the model for commercial aerial ride-sharing operations by demonstrating the capability to include passenger pooling.
Lastly, for the KPI's, the pad saturation comparison presented that pad saturation was only a few minutes cumulatively during a full day of operation in the arbitrary scenario. For other scenarios, the pad saturation could reach more than 3 hours, possibly indicating that 3 pads are a constraining factor. Between the clustered-balanced and directed-undirected versions, there was quite a large variation in the cumulative saturation. For the directed scenarios, the pad saturation of the central vertiport was reviewed in particular. Saturation was very scarce, contrary to what was expected beforehand. This characteristic implies that, even in a directed demand distribution, the focus vertiport does not necessarily require more infrastructure capacity, which can save costs.
In terms of the sensitivity analysis for various penalty cost coefficients, the comparison showed that the number of trips rejected in the scenario investigated remained constant with an increasing rejected trip penalty coefficient. This behaviour demonstrates that trip rejection may be a result of other constraints and not necessarily by a revenue-cost trade-off. With increasing penalty costs for time delay, the number of trips and customers served stayed constant while the time delay KPI's reduced. It turned out that more positioning trips were performed to maintain the same level of trips and customer served, with lower time delays. The energy cost associated with the extra positioning flights was thus favoured over the increased time delay penalties. In markets where the energy cost is higher, the opportunity to trade time delay for positioning flights may be more limited. Finally, a doubled and quadrupled vehicle dispatch cost demonstrated that with 5 or even 3 vehicles, the majority of trips and passengers could still be served. With fewer vehicles, vehicle utilisation increased. This attribute shows potential for operators to use a smaller fleet to serve the most profitable trips if overall customer spillage is not a paramount concern.
An additional sensitivity analysis varied the input fleet size and vertiport pad capacity. The resulting profit values indicated that, between a fleet size of 1 and 4 vehicles, each additional vehicle had a clear, incremental effect on the profit. The fifth vehicle had a lower positive effect, and the sixth vehicle caused a negligible increase in profit. Running the same scenario with 2 and 4 pads instead of the original 3 yielded no appreciable change in the profit. A comparison of the runs with 2 and 3 pads established that the model is able to make optimal use of a reduced pad capacity by rescheduling positioning flights such that idle stages are split between different vertiports and a pad is freed up for another vehicle. This sensitivity analysis provided insight into the incremental value of additional vehicles and pads, and thus demonstrated that the scenario under consideration could reduce resources to 2 pads and 5 vehicles for a profit value comparable to the original scenario which has 3 pads and 6 vehicles. Operators can consequently use this data for their deliberation whether or not to invest in additional vehicles or pads, both of which have a high purchase cost.

The results analysis is accompanied by the reflection that all scenarios were not solved to optimality. In terms of computational performance, it was found that the vertiport capacity constraints, in particular, caused a large problem size that also scaled with the number of vehicles and pads available, next to the number of customer requests. The restricted number of customer requests that could be run within a reasonable amount of time, together with the fact that not all scenarios could be solved to optimality, limited the ability for a

comprehensive and representative analysis. The customer requests distribution was very sensitive to the input demand distribution and seed, and with the small sample size the probability for e.g. pooling was also limited from the onset. Furthermore, reservations with regards to the fitness of the available input parameters are made as well as the absence of a complete and integral case study for validation.
Summarising, the model demonstrates the capability to find routes for all vehicles in the fleet that satisfy the majority of customer requests. Trade-offs between different components of the objective function are incorporated and can be further tweaked depending on an operator's requirements and preferences. The results analysis of the demand distributions investigated already yielded various focus points for improved efficiency and utilisation. The model may additionally serve as a tool for future operators to determine operational parameters such as fleet size, number of pads and allowed time windows. The formulation's usage is not only limited to eVTOL ride-sharing operators. The model can also be used for comparable vehicle routing problems, especially if they include (the combination of) electrical charging, goods/passenger pooling, recurring trips to the same location and limited infrastructure capacity, all of which have been addressed in this formulation.

Recommendations for future work on this topic propose improving the computational performance such that larger instances can be run. A suggestion is to dynamically add the vertiport capacity constraints only if a relaxed version violates the capacity. More in-depth simulations may then be run in conjunction with sensitivity analyses. Furthermore, the model can serve as a foundation to expand with other operational features that have not been incorporated in the current version.

References

Afroditi, A., Boile, M., Theofanis, S., Sdoukopoulos, E., and Margaritis, D. (2014). Electric Vehicle Routing Problem with industry constraints: Trends and insights for future research. In *Transportation Research Procedia*, volume 3, pages 452–459, Sevilla, Spain. Elsevier B.V.

Agatz, N., Erera, A., Savelsbergh, M., and Wang, X. (2012). Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223(2):295–303.

Airbus (2020). Vahana - Vehicle Demonstrators - Airbus. `https://www.airbus.com/innovation/urban-air-mobility/vehicle-demonstrators/vahana.html`. Date accessed: Jan 21, 2020.

Alnaqeb, A. H., Li, Y., Lui, Y. H., Pradeep, P., Wallin, J., Hu, C., Hu, S., and Wei, P. (2018). Online prediction of battery discharge and flight mission assessment for electrical rotorcraft. *AIAA Aerospace Sciences Meeting, 2018*, (210059):1–12.

Berbeglia, G., Cordeau, J.-F., Gribkovskaia, I., and Laporte, G. (2007). Static Pickup and Delivery Problems: a classification scheme and survey. *Top*, 15(1):1–31.

Bodin, L. and Sexton, T. R. (1986). The multi-vehicle subscriber Dial-a-Ride Problem. *TIMS Studies in the Management Sciences*, 22:73–86.

Boelens, J.-H. and Volocopter (2019). Pioneering the Urban Air Taxi Revolution. Technical report, Bruchsal, Germany. `https://press.volocopter.com/images/pdf/Volocopter-WhitePaper-1-0.pdf`.

Bombelli, A. and Tavasszy, L. (2020). The Multi-Warehouse Dock-Capacitated Pickup and Delivery Problem with Time Windows: Mathematical Formulation and Solution Method. Working Paper.

Booz Allen Hamilton (2018). Urban Air Mobility (UAM) Market Study. Technical report, McLean, VA. `https://ntrs.nasa.gov/search.jsp?R=20190001472`.

Bosson, C. S. and Lauderdale, T. A. (2018). Simulation evaluations of an autonomous Urban Air Mobility network management and separation service. *2018 Aviation Technology, Integration, and Operations Conference*.

Brown, A. and Harris, W. L. (2020). Vehicle design and optimization model for Urban Air Mobility. *Journal of Aircraft*, 57(6):1003–1013.

Bunte, S. and Kliewer, N. (2009). An overview on vehicle scheduling models. *Public Transport*, 1(4):299–317.

Chan, N. D. and Shaheen, S. A. (2012). Ridesharing in North America: Past, Present, and Future. *Transport Reviews*, 32(1):93–112.

Cordeau, J. F. and Laporte, G. (2003). The Dial-a-Ride Problem (DARP): Variants, modeling issues and algorithms. *4or Quarterly Journal of the Belgian, French and Italian Operations Research Societies*, 1(2):89–101.

De Almeida Correia, G. H. and Santos, R. F. G. (2014). Optimizing the use of electric vehicles in a regional car rental fleet. *Transportation Research Record*, 2454:76–83.

Duffy, M. J., Wakayama, S., Hupp, R., Lacy, R., and Stauffer, M. (2017). A study in reducing the cost of vertical flight with electric propulsion. In *Annual Forum Proceedings - AHS International*, pages 187–204, Denver, CO. AIAA.

Dumas, Y., Desrosiers, J., and Soumis, F. (1989). Large scale multi-vehicle Dial-A-Ride Problems. *Les Cahiers du GERAD*.

Dumas, Y., Desrosiers, J., and Soumis, F. (1991). The Pickup and Delivery Problem with Time Windows. *European Journal of Operational Research*.

Dunn, N. N. S. (2018). *Analysis of urban air transportation operational constraints and customer value attributes*. PhD thesis.

EHang (2020). White Paper on Urban Air Mobility Systems. Technical report.

Furuhata, M., Dessouky, M., Ordóñez, F., Brunet, M. E., Wang, X., and Koenig, S. (2013). Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*.

Garrow, L. A., German, B., Mokhtarian, P., and Glodek, J. (2019). A Survey to Model Demand for eVTOL Urban Air Trips and Competition with Autonomous Ground Vehicles. In *AIAA Aviation 2019 Forum*, Dallas, Texas. American Institute of Aeronautics and Astronautics.

Gelareh, S., Goncalves, G., and Monemi, R. N. (2013). On Truck dock assignment problem with operational time constraint within cross docks.

Hasan, S. and NASA (2018). Urban Air Mobility (UAM) Market Study. Technical report, NASA, Crown Consulting Inc., Washington DC. `https://ntrs.nasa.gov/citations/20190026762`.

Holden, J., Goel, N., and Uber Elevate (2016). Fast-Forwarding to a Future of On-Demand Urban Air Transportation. Technical report. `https://www.uber.com/elevate.pdf`.

Johnson, W. and Silva, C. (2019). Observations from exploration of VTOL Urban Air Mobility designs. In *7th Asian/Australian Rotorcraft Forum, ARF 2018*.

Justin, C. Y., Payan, A. P., Briceno, S. I., and Mavris, D. N. (2017). Operational and Economic Feasibility of Electric Thin Haul Transportation. In *17th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, Virginia. American Institute of Aeronautics and Astronautics.

Kleinbekman, I. C., Mitici, M., and Wei, P. (2020). Rolling-horizon electric vertical takeoff and landing arrival scheduling for on-demand Urban Air Mobility. *Journal of Aerospace Information Systems*, 17(3):150–159.

Kohlman, L. W. and Patterson, M. D. (2018). System-Level Urban Air Mobility Transportation Modeling and Determination of Energy-Related Constraints. In *2018 Aviation Technology, Integration, and Operations Conference*, Reston, Virginia. American Institute of Aeronautics and Astronautics.

Liang, X., Correia, G. H. d. A., and van Arem, B. (2016). Optimizing the service area and trip selection of an electric automated taxi system used for the last mile of train trips. *Transportation Research Part E: Logistics and Transportation Review*, 93(2016):115–129.

Miao, Z., Lim, A., and Ma, H. (2009). Truck dock assignment problem with operational time constraint within crossdocks. *European Journal of Operational Research*, 192:105–115.

Narkus-Kramer, M. (2013). On-Demand Mobility (ODM): A Discussion of Concepts and Required Research. In *2013 Aviation Technology, Integration, and Operations Conference*, Reston, Virginia. American Institute of Aeronautics and Astronautics.

National Air Transportation Association (2018). Urban Air Mobility : Considerations for Vertiport Operation. Technical report.

Patterson, M. D., Antcliff, K. R., and Kohlman, L. W. (2018). A proposed approach to studying Urban Air Mobility missions including an initial exploration of mission requirements. In *AHS International 74th Annual Forum & Technology Display*, Phoenix, AZ.

Pillac, V., Gendreau, M., Guéret, C., and Medaglia, A. L. (2013). A review of dynamic Vehicle Routing Problems. *European Journal of Operational Research*.

Porsche Consulting (2018). The Future of Vertical Mobility - Sizing the market for passenger, inspection, and goods services until 2035. Technical report. `https://www.porsche-consulting.com/en/press/insights/detail/study-the-future-of-vertical-mobility/`.

Pradeep, P. and Wei, P. (2018). Energy efficient arrival with RTA constraint for urban eVTOL operations. *AIAA Aerospace Sciences Meeting, 2018*, (210059):1–13.

Roland Berger (2018). Urban Air Mobility; The rise of a new mode of transporation. Technical report. `https://www.rolandberger.com/publications/publication_pdf/Roland_Berger_Urban_Air_Mobility.pdf`.

Schneider, M., Stenger, A., and Goeke, D. (2014). The electric Vehicle-Routing Problem with Time Windows and recharging stations. *Transportation Science*, 48(4):500–520.

Shaheen, S. and Cohen, A. (2020). Mobility on demand (MOD) and mobility as a service (MaaS): early understanding of shared mobility impacts and public transit partnerships. In *Demand for Emerging Transportation Systems*, pages 37–59. Elsevier.

Shihab, S. A. M. and Wei, P. (2019). By Schedule or On-Demand ? - A Hybrid Operational Concept for Urban Air Mobility Services. In *AIAA Aviation 2019 Forum*, pages 1–13, Dallas, Texas. AIAA.

Shihab, S. A. M. and Wei, P. (2020). Optimal eVTOL Fleet Dispatch for Urban Air Mobility and Power Grid Services. In *AIAA AVIATION 2020 FORUM*, pages 1–17.

Silva, C., Johnson, W., Antcliff, K. R., and Patterson, M. D. (2018). VTOL Urban Air Mobility concept vehicles for technology development. In *2018 Aviation Technology, Integration, and Operations Conference*.

Thipphavong, D. P., Apaza, R., Barmore, B., Battiste, V., Burian, B., Dao, Q., Feary, M., Go, S., Goodrich, K. H., Homola, J., Idris, H. R., Kopardekar, P. H., Lachter, J. B., Neogi, N. A., Ng, H. K., Oseguera-Lohr,

R. M., Patterson, M. D., and Verma, S. A. (2018). Urban Air Mobility Airspace Integration Concepts and Considerations. In *2018 Aviation Technology, Integration, and Operations Conference*, Reston, Virginia. American Institute of Aeronautics and Astronautics.

Toth, P. and Vigo, D. (2000). *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics, Philadelphia.

Uber (2020). Uber Elevate. `www.uber.com/elevate`. Date accessed: Jan 20, 2020.

Vascik, P. D. (2020). Systems Analysis of Urban Air Mobility Operational Scaling. Technical report, MIT International Center for Air Transportation, Cambridge, MA.

Vascik, P. D. and Hansman, R. J. (2018). Scaling constraints for Urban Air Mobility operations: Air traffic control, ground infrastructure, and noise. In *Aviation Technology, Integration, and Operations Conference*, Atlanta, GA.

Vascik, P. D. and Hansman, R. J. (2019). Development of Vertiport Capacity Envelopes and Analysis of Their Sensitivity to Topological and Operational Factors. In *AIAA Scitech 2019 Forum*, Reston, Virginia. American Institute of Aeronautics and Astronautics.

Volocopter (2020). Volocopter - Home. `https://volocopter.com/en/`. Date accessed: Jan 21, 2020.

Appendix

A - Overview of KPI Values

Table 11: Table showing the KPI values for all initial scenarios. Abbreviations used for each scenario correspond to table 4. 'TW' indicates the time window.

KPI/Scenario	**A TW-5**	**A TW-10**	**CU TW-5**	**CU TW-10**	**CD TW-5**	**CD TW-10**	**BU TW-5**	**BU TW-10**	**BD TW-5**	**BD TW-10**
Profit [USD]	5624	5607	7197	7197	6397	6660	7676	7676	7445	7460
# Vehicles dispatched	5	5	6	6	5	5	5	5	6	6
# Customer trips served	28	28	25	25	28	28	28	28	24	25
# Passengers transported	66	66	76	76	71	72	77	77	79	80
# Passengers not transported	2	2	9	9	3	2	2	2	6	5
Customer serv. flying time [min]	927	965	899	899	1017	999	1030	1036	817	824
Positioning flying time [min]	158	196	230	230	270	275	296	302	198	182
Charging time [min]	1066	1105	1053	1054	1161	1144	1178	1183	975	986
Idling time [min]	792	860	1636	1826	1016	1284	1379	1318	1832	1814
# Positioning flights	7	8	10	10	11	12	13	12	8	8
Positioning distance [km]	274	344	399	400	474	478	514	535	345	316
# Trips with delayed arrival	3	2	1	1	2	3	2	1	1	4
Cumulative time delay [min]	6	11	5	5	8	13	5	3	1	7
# Pooled trips	0	0	0	0	0	1	1	1	0	0
# Pooled passengers	0	0	0	0	0	2	1	2	0	0
Pad saturation time [min]	7	24	192	188	100	23	88	77	47	166

II

Literature Study
previously graded under AE4020

1 Literature Review

This chapter contains the literature study addressing various aspects of ride-sharing and UAM. In section 1.1, the definition and state-of-the-art of automotive ride-sharing are provided. Section 1.2 dives into the subject of UAM. In the different subsections, sequentially the most important industry players, state-of-the-art and future technology, operational costs, transportation categories, and market studies will be presented. Section 1.3 presents an overview of the existing mathematical models, relevant to this area of research. And finally, this chapter closes with a concluding section on the findings of the literature study, addressing main findings as well as remaining gaps, in section 1.4

1.1. Ride-sharing

Starting with the definition of ride-sharing, Furuhata et al. (2013) refer to ride-sharing as "a mode of transportation in which individual travelers share a vehicle for a trip and split travel costs such as gas, toll, and parking fees with others that have similar itineraries and time schedules." The authors note that ride-sharing 'combines the speed and flexibility (i.e. the possibility to choose a certain pick-up/drop-off location or route without having to adhere to a fixed stop) of private cars with the reduced cost of fixed-line systems (such as public transport) but at the expense of some convenience (primarily a direct routing).' This is mostly in line with the definition by Shaheen and Cohen (2020): "The formal or informal sharing of rides between drivers and passengers with similar origin-destination pairings." Advantages for the driver, passenger, society, and environment entail reduced travel costs and time, congestion, fuel use and consequently air pollution, as well as more travel options and even energy security (Agatz et al., 2012, Amey, 2011, Chan and Shaheen, 2012, Furuhata et al., 2013).

It is important to note the distinction between Transportation Network Companies (also named ride-hailing and ride-sourcing) and ride-sharing or pooling companies (Shaheen and Cohen, 2020). The former includes companies like Uber, Lyft, Careem and Didi Chuxing whereby users can request a private ride. In the case of ride-sharing/pooling, the ride may be shared with other passengers following the definition provided by Furuhata et al. (2013). Uber and Lyft do provide ride-sharing under their uberPOOL and Lyft Shared concept, respectively, but the network and usage are smaller with presence in 14-16 markets in the United Status versus over 300 markets for the ride-hailing options (Shaheen and Cohen, 2020). Other players in the U.S. include Scoop and Waze Carpool.

As mentioned, according to Shaheen and Cohen (2020) ride-sharing may be formal or informal. While the authors in their definitions of ride-sharing and ride-hailing explicitly note that for ride-hailing, driver and passenger matching is performed through digital applications, Agatz et al. (2012) posit that also for ride-sharing the use of technology is growing and is a "key enabler" for matching. Agatz et al. further indicate that ride-sharing is dynamic, which means that rides are often requested ad-hoc and in real-time, "ranging from a couple of hours to mere minutes before the desired departure time." The drivers providing the ride are independent and not employees of a transportation company, contrary to, for example, most taxi companies. This means that the drivers also have a desired route that needs to be taken into account. What further distinguishes ride-sharing from e.g. a taxi according to Furuhata et al. (2013) is the need for "prearrangement" as ride-shares cannot be "sought on the street". However, this is not necessarily valid any more considering services like uberPOOL, which offers on the spot ride-sharing, have been launched since the article by Furuhata

et al.. Another characteristic of ride-sharing is non-recurring trips. Whereas traditional carpooling usually has a pattern with fixed times and routes, for instance colleagues taking turns when commuting to work, this is not a given in ride-sharing (Morency, 2007). If carpooling is coordinated "on the spot", it is called casual carpooling (Kelly, 2007).

There are various objectives in the process of ride-sharing and matching drivers, passengers, and rides (Agatz et al., 2012). The first is the minimisation of the system-wide distance travelled by all vehicles. This is tied to the reduction of congestion and emissions of exhaust gasses. Secondly, there is the objective of minimising the system-wide travel time, i.e. the time spent when travelling between origin and destination. This objective is also related to reducing pollution in the grand scheme but also to the wishes of travellers. The authors emphasise that an objective aimed at a system-wide optimal solution "may not necessarily optimize the cost-savings of all individual ride-share participants." Thirdly, there can be an objective for maximising the number of (satisfied) participants. According to Agatz et al., this objective may be beneficial for the revenue of drivers and also induce future growth. Nourinejad and Roorda (2016) propose two adapted objectives where the vehicle distance travelled savings (the difference between the distance of the original routings and that of the new assignments) are maximised or where the service reliability (ratio of users with a match and total users) is maximised. If a detailed cost breakdown is available, minimising the total costs can also be an objective (Furuhata et al., 2013). However, this is possibly conflicting with objectives such as servicing the most passengers and minimising inconvenience.

In terms of operational constraints, Agatz et al. (2012) note that "time tends to be a more constraining factor than the availability of spare seats." To account for this, time windows for pick-up and/or drop-off are a common tactic. More on the concept of time windows will follow in section 1.3. Furthermore, personal preferences such as smoking and a female driver can be considered. In a review of previous literature, Agatz et al. also came across possibilities where riders can specify "the maximum excess travel time" and whether or not they want to ride with strangers or only with known, other participants. As ride-share participants usually make use of the service to save on costs, a constraint to ensure that only rides that reduce the cost for users are considered is also possible.

1.2. Urban Air Mobility

In this section, the concept of UAM will be discussed in various subsections. To start, major industry players will be presented in section 1.2.1. Subsequently, various aspects of the state-of-the-art technology and prospected future developments are treated in section 1.2.2. A breakdown of various cost elements associated with UAM is provided in section 1.2.3. Following that, various transportation categories tailored to UAM are shown in section 1.2.4. Finally, to provide more context for the future playing field of UAM, projections from various market studies are shared in section 1.2.5. As mentioned in the introduction, the broader vertical mobility services industry comprises several clusters. Most relevant for this thesis is the commercial passenger transport class. Going forward, UAM refers to this class.

1.2.1. Industry Players

Companies active in the UAM vehicle design and production can be roughly divided into four categories: current ride-sharing/ride-hailing providers, current commercial aircraft manufacturers, car manufacturers, and independent companies/start-ups. One of the most prominent companies is Uber, which has a separate business unit responsible for the development of a vehicle and ecosystem to integrate into their existing offerings, called Elevate (Holden et al., 2016). Commercial aircraft manufacturers Boeing and Airbus are both invested as well in the development of VTOL aircraft. Boeing's subsidiary Aurora Flight Sciences is testing an autonomous fixed-wing aircraft while Airbus is developing its Vahana single-seat tilt-wing vehicle and CityAirbus flying taxi (Airbus, 2020, Aurora Flight Sciences, 2020). Various automotive companies also want to diversify and expand their operations to UAM. Some companies work together with aircraft manufacturers, such as Audi and Airbus in the Pop.Up Next collaboration, and Volkswagen and Boeing (Audi Mediencenter, 2018, Volkswagen AG, 2019). Other car companies (planning to be) engaged in UAM development include Rolls-Royce and Aston Martin (NASA Aeronautics Research Mission Directorate (ARMD), 2018). Independent companies include the German Lillium and Volocopter, Chinese Ehang and US-based Kitty Hawk and Joby Aviation (EHang, 2020, Joby Aviation, 2020, Kitty Hawk, 2020, Lillium, 2020, Volocopter, 2020). For illustration, some vehicle designs are shown in fig. 1.1.

(a) CityAirbus by Airbus (2020)

(b) Lillium Jet by Lillium (2020)

(c) VoloCity by Volocopter (2020)

(d) Aurora Personal Air Vehicle by Aurora Flight Sciences (2020)

(e) EHang 184 by EHang (2020)

(f) Elevate concept by Uber (2020)

Figure 1.1: Images showing various VTOL vehicles by different manufacturers/operators.

1.2.2. Technology

The different companies identified in section 1.2.1 above take different approaches to the design of the vehicles, as could also be seen in fig. 1.1. Main design parameters and characteristics of VTOL aircraft include propulsion mechanism, number and size of blades, fuel source and capacity, cabin capacity in terms of passengers and weight, piloted or autonomous operation, range, speed, and charging/refuelling. Various of these parameters will now be discussed more in-depth.

Fuel and Propulsion

Most constructors mentioned in section 1.2.1 have chosen for a VTOL aircraft propelled by electrically driven rotors. This has the obvious advantage, which is highlighted by many manufacturers and operators, that there is no direct emission of exhaust gasses by the vehicle, contrary to most other aircraft and rotorcraft. Nevertheless, it may not be a clear-cut case for electrical energy as fuel, according to Kohlman and Patterson (2018) who compared electrical energy and liquid natural gas (LNG) as fuel for UAM operations. Based on the same mission parameters, LNG in a hybrid system offers "lower operating costs, reduced carbon dioxide emissions, and lower power system weights than pure battery-electric solutions at the same power level while also having fewer infrastructure integration issues." The hybrid systems reviewed include an internal combustion engine fuelled by LNG, which drives a generator and a turbine fuelled by LNG (with or without a solid oxide fuel cell). While the former offers both lower system mass and volume, the downsides of the internal combustion engine setup are the "likely increases in noise, vibrations, and maintenance costs." The LNG and solid oxide fuel cell combination does not offer large gains in system mass compared to a battery-only concept, but due to the higher efficiency and specific energy of the hybrid system, the range for equivalent mass would increase versus an all-battery fuel source. It is important to note that the emission figures have

been computed based on the average CO2 emissions of the electrical grid of the state of Texas. While this does give a specific outcome for the case study of the Dallas Fort Worth area under consideration in the paper, it does lead to a skewed comparison with LNG which cannot be generalised. The authors provided emission figures for the US, Texas, and California and the average emissions of the grid in California are almost 50% less compared to the nationwide and Texas average. Hence if the California emission was used, in terms of CO2 emissions, electricity would be considerably better than LNG instead of considerably worse. Therefore it is important that for costs and emissions of electricity, the location is carefully considered as the local method of energy generation plays a large role. Remarks on the cost comparison between LNG and electrical energy will follow in section 1.2.3.

Johnson and Silva (2019) investigated various VTOL designs and also concluded that all-electric propulsion vehicles are heavier than aircraft with a turboshaft or hybrid propulsion, in line with the findings of Kohlman and Patterson (2018). While most VTOL designs feature rotors only, some designs have a combined fixed-wing and rotor. This is primarily relevant for longer distance sectors, as wings are more efficient for generating lift than rotors. In case this configuration is electrically powered, less energy is required, which saves on battery weight. Nevertheless, this is not sufficient to counter "the increase in structure and propulsion weight", due to the wing and possible tilted-rotor, compared with an all-rotor design.

The propulsion mechanism also influences cost, maintenance, and reliability. Turbine engines, gearboxes, complex rotors, and hydraulic systems require frequent inspection and maintenance, which drives up costs (Boelens and Volocopter, 2019). According to the authors, this is equally the same for hybrid eVTOL aircraft with a turbine engine or turbo-generator, as well as aircraft with complicated tilting mechanisms. Hence direct-drive electric motors and fixed-pitch rotors are preferred from a maintenance cost point of view.

Various, if not most, manufacturers and operators, such as Uber Elevate, Ehang, and Volocopter, rely on multiple, smaller, directly (and independently) electrically-propelled rotors in the so-called Distributed Electric Propulsion (DEP) configuration (Boelens and Volocopter, 2019, EHang, 2020, Holden et al., 2016). These propulsors, as Kim et al. (2018) define, "can be placed, sized, and operated with greater flexibility to leverage the synergistic benefits of aero-propulsive coupling and provide improved performance over more traditional designs." DEP makes use of the "scale-free" nature of the rotors, which means that a larger number of small rotors are not heavier than fewer but larger rotors. But, as mentioned, the larger number of rotors may offer performance improvements as they can be tailored for specific functions. Silva et al. (2018) provide various design possibilities for VTOL aircraft such as a quadrotor with four rotors in each 'corner' next to the vehicle, the side-by-side helicopter with two rotors on either side of the vehicle, and the lift+cruise VTOL aircraft, which uses different rotors, varying from 4 to 12 total, for vertical and horizontal flight. The latter is a good example that demonstrates that DEP offers the possibility to tailor more but smaller rotors to a certain function, which improves efficiency, without system-wide weight gains. Other improvements of DEP include 'redundancy, higher reliability, less mechanical complexity, and more control authority' (Moore, 2012, Nneji et al., 2017)

Batteries

Current battery technology is a limiting factor for the design of eVTOL aircraft (Alnaqeb et al., 2018, Holden et al., 2016, Narkus-Kramer, 2013, Pradeep and Wei, 2018). This is especially due to specific energy (the amount of energy per unit weight) and specific power (the amount of power per unit weight) in comparison with the energy and power required. Boelens and Volocopter (2019) note that "the level of required power in vertical take-off increases more than linearly with the take-off weight," following momentum theory. The author illustrates that a typical design consumes the equivalent of the charge of a current Tesla Model 3 battery of 50 *kWh* in three minutes only for take-off and landing. Next to the take-off, cruise and landing, sufficient reserve must also be taken into account for a possible diversion or, power-intensive, hovering when loitering.

In the example provided, Boelens and Volocopter (2019) assume a specific energy of the cell of 250 *Wh/kg*, which the author claims is "optimistic". Kohlman and Patterson (2018) assume 300 *Wh/kg* in a lithium-sulfur battery cell for a current day comparison with LNG, however stress that this does not account for losses in a complete, assembled power system and notes it is "an advancement over the state-of-the-art today." Justin et al. (2017) project that a battery pack (in which efficiency losses and unusable amounts have been subtracted) featuring a specific energy of 350 *Wh/kg* is "achievable within the 2030 timeframe." Johnson and Silva (2019) even assume a lithium-ion battery pack with a usable 400 *Wh/kg* specific energy for their VTOL design exploration but do not provide a timeline. In another paper, the same authors provide a breakdown of energy loss from the cell to the installed pack (Silva et al., 2018). For the 400 *Wh/kg* installed pack, a single battery cell needs to have a specific energy of 650 *Wh/kg*. This is a very considerable increase compared to the

current 250-300 *Wh/kg* on cell-level as posited by Kohlman and Patterson (2018) and Boelens and Volocopter (2019).

Another limitation identified in batteries is the deterioration of the battery health after charging cycles. Especially fast charging is detrimental to battery life (Boelens and Volocopter, 2019). One way to combat this is by swapping batteries between flights and charging them "at optimal (low) C-rates, while being properly balanced" and with "reduced thermal stress." These difficulties have also been common with electrical vehicles (EV) in the automotive sector. Nie and Ghamami (2013), in their research to EV charging infrastructure, found that swapping batteries yields the highest level of service (quickest journeys with minimal delay), compared to charging and being idle. Yet swapping may also be more expensive than charging, hence a trade-off has to be made.

Next to the battery design, also real-time battery management and evaluation is an important feature. Alnaqeb et al. (2018) describe that, as eVTOL aircraft are relatively light-weighted, "they are subject to wind uncertainties in low-altitude airspace." Hence the authors developed a model which constantly verifies whether the intended mission can be completed with the actual battery levels. The model, however, only takes into account the forward flight (cruise) phase. As noted earlier, the take-off and landing phases are the most power-intensive. Furthermore, the characteristics of the batteries were greatly simplified by assuming the properties of a single cell scale proportionally in combined cells.

Autonomy

Autonomy is an important feature of UAM operations, both for scalability, safety, cost, and capacity. If a pilot can be omitted, there is room for an additional passenger and training and salary expenses are saved. Furthermore, if all VTOL vehicles operate autonomously and are connected through a single guidance system, the number of aircraft operating in a certain zone can increase as all movements are coordinated. The approach to autonomy varies among researchers and industry. Mostly, there is consensus that initially, i.e. in the next few years, VTOL aircraft are still flown by a pilot. If the technology for autonomous flying develops further and becomes available after another few years, the human element will be phased out (Balać et al., 2017, Boelens and Volocopter, 2019, Hemm et al., 2016, Nneji et al., 2017). Next to the technical challenges of autonomous flying, there are also legal and societal factors in play (Booz Allen Hamilton, 2018, NASA, 2018, Yedavalli et al., 2019).

A comparison of a completely autonomous transport system with existing driver-based automotive ride-sharing by Iglesias et al. (2017), showed that "a Model Predictive Control (MPC) algorithm that leverages short-term demand forecasts based on historical data" computes an improved strategy for the repositioning of vehicles. I.e. based on previous demand (location) data, the algorithm computes the optimal, preemptive, location for the autonomous vehicle fleet to operate from. As a result, the mean customer wait time was reduced by up to 89.6%.

Infrastructure

While UAM benefits from the fact that, in principle, it is node-based instead of path-based transportation, the location of these nodes must be carefully chosen. In an urban environment, it is not possible, or allowed, to land at any point. This is in contrast with automotive ride-sharing/ride-hailing where cars can generally pick up and drop off riders at any location. VTOL aircraft, instead, must make use of a so-called vertiport; an airport for VTOL vehicles. These vertiports should be strategically located to offer a comprehensive network in the city at hand and offer easy access to desired locations and/or other, onward, transport options (Yedavalli et al., 2019). Current transport operators can use historical data to their advantage in order to determine attractive locations. There are many factors that influence the location, size, facilities, and design of a vertiport. This includes nearby buildings, airspace limitations, noise abatement restrictions, routing procedures, charging connections and safety and security (National Air Transportation Association, 2018, Thipphavong et al., 2018, Vascik and Hansman, 2018, 2019). UAM operators can build and operate their own vertiports or, for example, make use of existing heliports on high rise buildings. Holden et al. (2016) also suggest that "tops of parking garages and unused land surrounding highway interchanges" can be converted into vertiports. Design criteria include the number of take-off/landing pads, gates for passenger boarding/disembarkation and refuelling/recharging/maintenance stands. Vascik and Hansman (2019) developed a model that can determine the optimal ratio of gates/stands and pads based on various operational parameters. Holden et al. (2016) provide a maximum capacity of 12 aircraft for a typical vertiport in a large city such as New York "to achieve a compact infrastructure size while enabling capacity for multiple simultaneous VTOL takeoff and landings to maximize trip throughput."

Depending on the service level of the operator and regulatory requirements, these vertiports must not only provide facilities for VTOL vehicles but also for passengers. This may include a building with check-in and weighing facilities, a waiting area, security screening, etc.. Holden et al. (2016) posits that, next to vertiports, also vertistops can be introduced, which are "single-vehicle landing locations where no support facilities are provided, but where VTOLs can quickly drop off and pick up passengers without parking for an extended time." Considering the aforementioned possible requirements for check-in, weighing and security, this may not be possible everywhere.

Range & Speed

Boelens and Volocopter (2019) performed an in-depth analysis and found that "most megacities have an urban area spanning less than 30 *km* around the geographic center, while most of the major airports serving these cities are within 30 *km* of the city center." Next to that, "70% of the analyzed megacities have a major airport within 20 *km* of the city while 93% have a major airport within 30 *km* of the city center." Consequently, Boelens and Volocopter estimate the optimal design range of their VTOL vehicle to be between 30 and 35 *km* which will enable both "inner-city taxi and airport shuttle services in more than 90% of the megacities."

Patterson et al. (2018) also computed typical ranges. Based on a review of 28 metropolitan population densities, the usual range for a one-way work commute is 37.5 nautical miles or about 70 *km*. However, the design range for the aircraft should be double that to prevent charging after a single flight. For airport flights, Patterson et al. concluded half the commute distance, approximately 35 *km*, is sufficient for most airports. This is about the same as the upper bound by Boelens and Volocopter (2019).

The cruise speed, in combination with the target range, greatly determines the time savings over conventional transportation and thus the attractiveness. At the same time, a higher speed requires more energy and power, and thus influences the battery characteristics. Holden et al. (2016) project cruise speeds of 150-200 *mph* (241-321 *km/h*), which would apply to the larger distance commutes of Uber Elevate. Within an urban environment, such high speeds are not possible due to noise, collision-avoidance, and risk of bird strike damage (Boelens and Volocopter, 2019). Volocopter, therefore, aims at an average speed of 80 to 100 *km/h* for their flights. Depending on the congestion and regular transport travel speed in a certain city, Boelens and Volocopter projects this speed offers at least a time saving of 50%. Porsche Consulting (2018) prospects different speeds for different types of VTOL vehicles, ranging from 70 to 300 *km/h*. Roland Berger (2018) notes as well that the velocity highly depends on the vehicle design. Aircraft with a tilt-rotor or fixed-wing can fly faster than a vehicle with only vertical lifting rotors.

The cruise speed influence on travel time decreases with shorter distances, as flights need to take-off, land, accelerate and decelerate, which generally takes a fixed amount of time per flight. Kleinbekman et al. (2018) investigated arrival sequencing for UAM operations for an arrival from an altitude of 500 meters, based on the EHang 184 vehicle. The final descent and landing phase takes a fixed 142 seconds. The minimum required time of arrival, from cruise to the final descend phase, was found to be 165 seconds. Depending on the number of concurrent arrivals and battery levels, this time may be increased. But, at the minimum, 307 seconds or about 5 minutes are required for the approach and landing phase alone based on this study. If a comparable amount of time is also required on take-off, about ten minutes per flight are already dedicated to a non-cruise phase. Consequently, the combination of the cruise speed and flight distance needs to be such that there is still an acceptable time saving over ground-based transport.

Noise

Considering the fact that UAM, as the name implies, will operate in (dense) urban environments, an acceptable noise footprint is of paramount importance for social acceptance (Balac et al., 2011, Booz Allen Hamilton, 2018, Porsche Consulting, 2018, Vascik and Hansman, 2017). Uber Elevate has the objective that VTOL noise levels are "barely audible" during cruise, and the noise during take-off and landing should "be comparable to existing background noise" (Holden et al., 2016). This translates in a noise level of 62 $dB\ L_{Amax}$ at 500 feet (approximately 150 *m*) altitude, which is "half of a truck traveling on a residential road."

The noise of rotorcraft is closely related to the blade tip speed, which in turn is related to the vehicle's velocity (Holden et al., 2016). Another influencing factor is the rotor size. Large rotors have low revolutions-per-minute which means that the noise it generates is low frequency. This type of noise travels farther than high-frequency noise, which adds to noise annoyance. A smaller rotor may reduce both sources of noise. However, it may not provide sufficient disc area to generate the required lift. DEP, as introduced in the Fuel and Propulsion part, forms a possible solution for mitigating noise levels whilst still providing enough lift and directional speed (Duffy et al., 2017, Kim et al., 2018). The removal of the turbine or piston engine, which are common in a.o. helicopters, also contributes to a smaller noise footprint.

1.2.3. Costs

In order to make UAM attractive to a large group of consumers, the pricing should be near the price point of premium ground-transport offerings. It can still "command a premium for the time-savings", but should be "much more affordable than today's helicopter services" (Roland Berger, 2018). Roughly, the costs for UAM comprise three categories: construction of the aircraft, infrastructure, and operational costs. Considering the scope of the thesis, the operational costs are most relevant. In the case of an electrical VTOL aircraft, the cost of electricity is highly correlated with the required power and energy. Thus, an 'energy-efficient design enables lower operating costs' (Boelens and Volocopter, 2019). Furthermore, Boelens and Volocopter note there is a clear trade-off between range and the required capacity and weight of batteries, and consequently, the power required and cost. McKinsey&Company et al. (2019) posit that the direct operating costs per seat-mile "must decline by more than 90 percent compared with the estimated 4 to 6 *USD* (...) that helicopter travel costs today" for large scale adoption of UAM, i.e. about 40 to 60 *USD* cents per mile. They envision that reductions in maintenance costs, switch from fuel to electricity, autonomous flying instead of having a pilot, lower infrastructure fees and lower vehicle purchase costs can contribute to lower operating costs. Additionally, McKinsey&Company et al. expect that the average utilisation of eVTOL is 4 times higher than a helicopter: 6.0 vs. 1.5 hours per day.

Kohlman and Patterson (2018) performed more extensive research in the energy-related costs of UAM operations, comparing electricity and liquid natural gas (LNG). The precise costs per unit distance travelled is highly dependent on the energy costs, which can vary by location. The electrical costs based on a case study in the United States range from 5.77 to 16.3 *USD* cents per *km*. For LNG, the range is 2.88 to 13.46 *USD* cents per *km*, depending on the exact powertrain. In the most expensive case, the estimated cost of electricity of Kohlman and Patterson (2018) is 65% of the lowest direct operating cost by McKinsey&Company et al. (2019).

Next to the costs of expendable fuel itself, Kohlman and Patterson (2018) also computed the total cost per vertiport for recharging/refuelling infrastructure. Based on their benchmark of 500 vehicles, the facilities required for recharging electrical vehicles are considerably more expensive than refuelling facilities for LNG-powered VTOL aircraft; 72 million *USD* versus 2.25 million *USD*, respectively. This is largely due to the fact that recharging batteries is estimated to take 5.8 minutes, while refuelling the LNG-tank only takes 4.9 seconds. Consequently, Kohlman and Patterson computed an amount of 160 chargers, valued at 450.000 *USD* each, and 15 pump stations, valued at 150.000 *USD* each. Hence if the number of aircraft per vertiport is different, the infrastructure cost for eVTOL charging can reduce considerably, and the comparison provided may not be representative.

1.2.4. UAM Transportation Categories

As was the case with automotive ride-sharing and ride-hailing explained in section 1.1, also UAM has distinct business models for passenger transportation with different characteristics. The main industry players identified in section 1.2.1 take different approaches to the possible target markets. This affects, for instance, the range and capacity of the vehicles.

Almost all literature agrees that, for UAM to be attractive and be the preferred option over existing transportation modes, there needs to be a clear time-saving advantage. Porsche Consulting (2018) posits that passengers should save "at least 20% in total travel time, notwithstanding transfers." Taking account of the time spent travelling to a vertiport and boarding/disembarking the aircraft, Roland Berger (2018) notes that trips should be "at least 15 to 25 *km* in order to deliver genuine time-savings and become the fastest urban mobility option." Porsche Consulting (2018) concurs as it states a distance of 20 *km* and more to be attractive.

The study by Roland Berger (2018) yielded three possible UAM ride-sharing categories: intercity flights, air taxis, and airport shuttles. Intercity flights are aimed at providing scheduled transport between cities "which are too close even for regional airlines" and where it would form "the only high-speed travel option without much infrastructure need", e.g. when a train service is non-existing or insufficient. The large range required does "pose challenges to technology", especially battery capacity and weight. Porsche Consulting (2018) assumes the range of this category to be between 100 and 400 *km*. For Uber Elevate, this long-distance (commuter) transport category is the focal point. Firstly because it would form an attractive alternative since "the amount of time and money saved increases with the trip length" (Holden et al., 2016). Secondly, there is a relatively small need for (new) infrastructure as "a small number of vertiports could absorb a large share of demand from long-distance commuters since the 'last-mile' ground transportation component will be small relative to the much longer commute distance."

Air taxis would form 'on-demand, point-to-point operations', which is comparable to the automotive ride-hailing/ride-sharing concept. It would (ideally) offer the fastest travel time compared to existing op-

tions. A drawback is the need for a dense network of vertiports in order to generate time savings on shorter distances; the vertiports need to be close to the desired locations of users. An indicative range, according to Porsche Consulting (2018), would be 20-50 *km*. Both Porsche Consulting and Roland Berger note the opportunity for pooling passengers in this form of transportation. Interestingly, Nneji et al. (2017) consider both intercity and intracity flights, with distances up to 300 miles, as part of its air taxi category. The authors omit, however, if there are various vehicle designs within the air taxi category or if they envision a 'one-size-fits-all' approach.

The airport shuttle case is straightforward: providing the fastest travel option from a city to the airport and v.v. (Roland Berger, 2018). These flights can be scheduled instead of on-demand. Drawbacks are the requirement for luggage storage and consequently more payload as well as the need to merge UAM operations in commercially used airspace of the airport.

Patterson et al. (2018) additionally provide other possible missions such as end-to-end city transfers, "allowing passengers to bypass city traffic and quickly move from one side of town to another." Next to that, there may be "metro-like services to connect passengers to other existing forms of mass transit." This is also mentioned as a UAM concept by NASA (2018). Patterson et al. extend the air taxi concept of Roland Berger (2018) to work commutes. Commuter flights may be scheduled or semi-scheduled. Next to that, the authors introduce the Air Pooling concept as a concept distinct from the Air Taxi. The latter would have 1-4 passengers while Air Pooling would have 3-6 passengers per flight.

1.2.5. Market Study

Operators, constructors, research institutions and consultancy firms alike have performed various market studies to size the potential of UAM. The ability to conquer a large and relevant market share will depend greatly on overcoming some of the challenges mentioned in the previous subsections. As has also become clear, the target market varies between companies.

McKinsey & Company, on request of NASA, performed a market sizing study and concluded "there is likely a limited potential market for air taxis in concentrated areas of high net worth individuals and businesses in 2030" (NASA, 2018). This is mainly due to the high investment costs to build a comprehensive infrastructure. Nevertheless, according to NASA (2018), the air metro concept could likely be profitable in 2030 in the 15 largest US cities by population. The number of passengers transported p.a. in 2030 is projected to be 740 million at an average price of 30 *USD* per trip. 23.000 VTOL vehicles are expected to be in service. In a separate McKinsey article, the air taxi market for 2040 is valued at 1.5 billion *USD* with 1000 vehicles (McKinsey&Company et al., 2019).

For 2030, Roland Berger (2018) estimates about 12.000 passenger drones are in service. This consultancy firm does see a potential for air taxis as well as airport shuttles in about a 60-40 ratio. Intercity flights are only expected from 2035 onwards. By 2050 the fleet size is projected to grow to 98.000, and 100 cities will have UAM operations. The number of drones per city can range from 60 to 6000.

In contrast to McKinsey&Company et al. (2019) but in line with Roland Berger (2018), Porsche Consulting (2018) also expects that air taxis will be the first concept to become ubiquitous with 15.000 units by 2035 in a market worth 21 billion *USD*. The number of vehicles projected for 2030 is 2000, so Porsche Consulting expects a very high growth rate in the 5 year period. Porsche Consulting (2018) also expects that by 2035, 8000 UAM vehicles will service the intercity market, valued at 11 billion *USD*. In terms of geographic locations, "the Asia-Pacific region is expected to capture around 45 percent of this market by 2035 (...), followed by the Americas with approximately 30 percent of the market." Europe and other regions will take up the remaining 25%.

Booz Allen Hamilton (2018) also performed a market study including a sensitivity analysis of possible constraints for the US market. The requirement for dedicated and sufficient infrastructure became clear from the quantitative results. For the case study investigated, there would be 13 million passengers daily willing to pay for UAM transportation. If a constraint of only using existing infrastructure was added, this number dropped to 1.5 million passengers per day. If the current infrastructure capacity per hour was also a constraining factor, the number of passengers plummeted further to 120.000.

Most of the references mentioned above assume a near mature situation in a wide variety of cities. However, in the coming years, the early 2020s, UAM operations will be gradually rolled out in a few key cities. Candidate cities mentioned frequently are New York, Dallas, San Francisco, Los Angeles, Dubai, and Singapore. Volocopter has already conducted test flights in both Dubai and Singapore (Ibekwe, 2017, Volocopter GmbH, 2019). Dallas and Los Angeles are the first target markets for Uber (2020). As explained earlier, these candidate cities are in metropolitan areas, characterised by relatively large commuter distances and/or con-

gestion such that UAM operations can offer a competitive advantage.

A model to simulate demand has been developed by Kohlman and Patterson (2018), "based on general demand trends" by Uber (2020). The demand is modelled by 'the sum of three normal distributions, centred around peak hours with a standard deviation based on the time of day.' The combination of these curves is then normalised to a maximum of 1. Per vertiport location, a weight parameter is added to represent the relative demand of each location. To generate a request, a random number between 0.0 and 1.0 is compared to a number on the probability curve. If the random number is less than the value on the curve, there is a flight request. In the model provided, this comparison is run each 10 seconds for an entire day. This can, of course, be tweaked to account for different operating times and demand.

1.3. Mathematical Models

In order to have a theoretical foundation for the development of the mathematical model, existing literature related to the topic has been perused. The findings will be presented in this section. First of all, the taxonomy of relevant mathematical concepts will be treated. Following that, various algorithms per problem category and other characteristics will be presented. Subsequently, various solving techniques from the operations research fields will be explained. It is to be noted that most of these mathematical models and algorithms have been developed based on automotive transport. One model dedicated to UAM vehicle scheduling was found, and will be highlighted separately. Lastly, a separate subsection will treat constraints and models related specifically to electrical transportation, such as batteries, charging and discharging.

1.3.1. Classifications

A large amount of research has been performed on the topic of so-called Vehicle Routing Problems (VRP), mostly by mathematics and operations research academia. According to Pillac et al. (2013), the VRP is "a generalization of the Traveling Salesman Problem (TSP)." In essence, the problem "consists of finding a set of routes for a number of identical vehicles based at the depot, such that each of the requests is visited exactly once while minimizing the overall routing cost." There are multiple variants of the VRP, such as the "VRP with Time Windows (VRPTW), where each customer must be visited during a specific time frame." In the vehicle routing with 'pick-up and delivery problem (PDPs)', "objects or people have to be transported between origins and destinations" (Berbeglia et al., 2007).

An even more specific subproblem is the 'Dial-a-Ride Problem (DARP)', which "consists of designing vehicle routes and schedules for n users who specify pick-up and drop-off requests between origins and destinations" (Agatz et al., 2012, Berbeglia et al., 2007, Cordeau and Laporte, 2003). The objective generally is minimised costs or maximised demand served, considering a set of constraints. Next to the distinction that the DARP concerns passengers as the commodity as opposed to goods, there is an important consequence for the objective function, say Cordeau and Laporte (2007): "when transporting passengers, reducing user inconvenience must be balanced against minimizing operating costs." Indicative criteria for this aspect "include route duration, route length, customer waiting time, customer ride time (i.e., total time spent in vehicles), and difference between actual and desired drop-off times." Additionally, capacity in passenger vehicles is often a constraint while this 'is often redundant in Pick-up and Delivery Vehicle Routing Problem applications.' These considerations would certainly also be applicable to UAM ride-sharing.

In terms of supply, there is also a distinction between the number of vehicles available. There are single-vehicle and multi-vehicle DARP (Cordeau and Laporte, 2003). Next to that, the number of origins and destinations can be used to classify problems (Berbeglia et al., 2007).

1.3.2. Static vs Dynamic

Psaraftis (1980) notes that there are generally two different variants of the VRPs; static and dynamic. In the static case, all information such as supply and demand are known a priori when executing the optimisation model. Meanwhile, in the dynamic variant, parameters are not completely known in advance and can change or emerge during the execution, in real-time (Cordeau and Laporte, 2003). Especially in modern on-demand products, customer requests are a continuous input and the optimisation model requires frequent updates and re-runs to account for this (Agatz et al., 2012). This dynamic nature also increases the complexity (Pillac et al., 2013). There are various methods to account for a dynamic system, according to Pillac et al. (2013). The most basic is "periodic reoptimization", which basically means an algorithm for a static problem is run periodically to account for a new state. This can be "at fixed intervals, or when data changes." The outcome is still a set of routing plans. Continuous reoptimisation "updates a current routing whenever the available

data changes and keeps good solutions in an adaptive memory". Every time a vehicle becomes available, or whenever a new request is made, it is checked in the memory if a good route is available, fitting the new data.

1.3.3. Basic Vehicle Routing Problem

There are different approaches to solving the various vehicle routing problems. Bunte and Kliewer (2009) provide an overview of vehicle scheduling models, based on a public transport bus company that needs to assign vehicles to cover a given timetable. This is thus the most basic VRP where each trip has to be covered once, with minimised operational costs or fleet size. Hence this is different from a ride-sharing model where rides do not necessarily need to be covered if it would result in a sub-optimal solution. Nor is the demand, either static or dynamic, taken into account. The number of depots, where vehicles originate and return, influences the approach. There are models for a single depot and multiple depot case.

As UAM operations can start in one or multiple depots, both categories are treated. Starting with the single depot: in the 'minimal decomposition model', a partial ordered set of all trips in the timetable needs to be defined (Saha, 1970). A relation for consecutive trips is introduced if a trip j starts at the end station of the previous trip i, and the departure time of j is equal to, or later than, the arrival time of trip i. In the first phase, short connections of these trips are generated. In the second phase, the whole problem is solved using fixed connections from phase 1. The resulting solution solves for minimum fleet size. However, operational costs are not taken into account.

This issue is addressed in the 'assignment model'. In this model, each trip has a departure and arrival node and each arc between nodes has a cost factor (Orloff, 1976). If an arc is between two non-compatible trip nodes, the cost of an additional vehicle is awarded. While the operational costs are now integrated, again the number of vehicles available cannot be specified or limited.

The 'transportation model' builds upon the assignment model (Gavish and Shlifer, 1979). However, a compatibility relation that describes if a trip j can be served after a trip i with the same vehicle is incorporated. Consequently, only arcs that satisfy the compatibility relation are considered in contrast with having arcs between every node in the assignment model. Depot nodes are connected to each trip to allow for 'deadhead' trips; rides from a depot to a start station or from an end station back to the depot. The trip nodes have one unit of flow. The sum of the nodes used by vehicles is equal to the number of vehicles available. As not every vehicle is necessary to serve trips, an arc with zero cost is placed between depot nodes. Hence this model does allow for the specification of a maximum number of vehicles available to perform the trips. It is possible to extend the model to account for the case of infeasibility. This requires to add arcs again between all trip nodes, albeit with a penalty cost to "represent the penalty of not serving a trip at all."

The 'network flow model' takes another different approach. Like the transportation model, "each trip is represented by two nodes connected via a trip arc" and two nodes are added to represent depots at the start and finish. The arcs have an associated cost. A solution is reached by "solving a minimum cost flow problem considering all nodes as transshipment nodes", i.e. as an intermediate destination. A feasible flow, therefore, is a flow from the starting depot to the end depot. A constraint ensures that there is as much inflow as outflow on each node. The trip arcs are given a value of one such that all trips are served. Furthermore, "the upper bound of the arc leading back is set equal to the number of vehicles available." The outcome is the arcs that are to be utilised for minimum cost.

For the multi-depot models, there is an additional restriction that vehicles have to return to the same depot from where they started their route. It is to be noted that this is not necessarily a requirement for UAM operations. VTOL aircraft can end their 'shift' at another vertiport from where they started, as long as it offers some facilities such as charging. Instead, this constraint can be replaced by a constraint that makes sure that the vertiport stand capacity is not exceeded.

The number of depots, as well as the vehicles available per depot, are additional variables. Within the multi-depot case, Bunte and Kliewer (2009) mention three further model formulations: single-commodity, multi-commodity and set partitioning.

In the 'single-commodity model', there is "one node per trip and additional nodes for the depots or vehicles." The constraints specify that "each node is covered by exactly one circuit, each circuit contains exactly one depot/vehicle node", and the capacity of a depot is not exceeded. Carpaneto et al. (1989) also add nodes for each available vehicle per depot. Arcs are added from the vehicle nodes to all trip nodes and back, and these are assigned a cost. As not every vehicle needs to be assigned, "additional arcs with zero costs are inserted for each vehicle node", pointing to the same node (Bunte and Kliewer, 2009). The so-called subtour elimination constraints forbid paths, part of feasible routes, which feature more than one depot. This ensures that there are only routes that start and end in the same depot. Because arcs need to be inserted from

every vehicle to all trip nodes, there is such an "extremely high" number of constraints "that MIP-based solution approaches cannot consider these constraints explicitly." Instead of having a node for each vehicle, the 'single-commodity model with assignment variables' groups the vehicle nodes to a single node per depot. This, thus, reduces the number of arcs and consequently the constraints. A new variable is introduced which assigns a trip to a depot.

The 'multi-commodity models' "are an extension of the network flow approach for a single-depot case", which was presented earlier. A network is determined fur each depot, and these are then combined for the multi-commodity formulation (Bunte and Kliewer, 2009).

Starting with the 'connection-based model', "all trip compatibilities are considered explicitly." If a connection between nodes is possible, an arc is inserted where "the number of connection arcs grows quadratically with the number of trips." Just as in the single-depot case flow conservation constraints are added to each node. The main difference with the single-depot case, is that a trip is related to more than one arc. Consequently, the lower bound of an arc cannot be set to one. This is solved by adding 'cover constraints', which still ensure all trips are served by "allowing exactly one arc of the trip arcs to be chosen in a feasible solution."

Finally, in 'set partitioning models', first all feasible routes are to be determined. Subsequently, a subset of these routes is chosen "which fulfills all restrictions" (Bunte and Kliewer, 2009). The set partitioning problem "has only a few constraints (in fast a constraint for each trip to be covered) but a large number of variables since every feasible path through the network for each depot is a variable in the model." Furthermore, this model has the advantage that 'duty-related constraints can be incorporated rather easily by excluding infeasible routes.'

1.3.4. Pick-up and Delivery Problem Models

The above models describe a basic VRP, where a set of trips of a timetable has to be served. As explained earlier, a more advanced version of the VRP is the PDP, where "each transportation request specifies a single origin and a single destination and all vehicles depart from and return to a central depot" (Savelsbergh and Sol, 1995). A framework for solving the general PDP is provided by Savelsbergh and Sol (1995). The model consists of a set with transportation requests where for each request, there is a size of the load, a destination as well as an origin. Compared with the VRP, in the PDP the origin/destination is not (necessarily) the same as the start/end location of a vehicle to signify that a trip does not start/end in a depot. Furthermore, there is a set of vehicles where each vehicle has a capacity, start and end location. Each transportation request also is assigned a distance, travel time and cost. The variables to be determined are which vehicle serves a certain request, the departure time and the load of a vehicle. Constraints are set such that each transportation request is assigned to one vehicle only, the vehicles start and end in the correct location, capacity is not exceeded, the vehicles only leave a location if it corresponds with an assigned trip and departure times are in sequential order. Savelsbergh and Sol (1995) have a fixed load of 1 (passenger) for the DARP, although this need not necessarily be the case.

1.3.5. Time Windows

Practically, and especially for the DARP that is treated hereafter, the pick-up (and/or drop-off) can be variable as this is expected by the rider. The consequences for the solution are two-fold. Depending on the time constraints, it may be difficult to find a feasible solution. On the other hand, the optimisation "may benefit from the presence of time constraints since the solution space may be much smaller" (Savelsbergh and Sol, 1995). Cordeau and Laporte (2003) consider that specifying timings by users on departure and arrival is "unduly constraining." In order to account for variable timings, time windows may be introduced, specifying a time interval during which the pick-up and/or drop-off is required. Bunte and Kliewer (2009) distinguish two types of time windows: discrete and continuous. In the discrete case, additional trip arcs are added for each time window value. Continuous time windows can be applied to the set partitioning model presented earlier, where time windows are considered within the column generation process. Instead of solving the routes using the standard shortest path algorithm, the adjusted shortest path problem with time windows algorithm will be used. Desaulniers and Villeneuve (2000) explain that additional variables and constraints are added to the shortest path formulation. First of all, a continuous time variable is added, which indicates the departure time from a certain node. Together with the travel time on an arc, the waiting time can also be determined. The extra constraints are set such that, for a chosen arc, the departure time from a node cannot be earlier than the previous arrival time at that same node, and the earliest and latest departure time for a certain node according to the time window are set. The objective function combines the (operational) cost of an arc as well as the waiting time as a result of choosing that arc. Because part of the objective function and one of the

constraints are products of decision variables, this problem is nonlinear.

1.3.6. Dial-a-Ride Problem

Cordeau and Laporte (2003) provide an overview of the DARP variants and algorithms. The authors describe that "three decisions are associated with the construction of a DARP solution." These are: '1) determining groups (clusters) of users served by the same vehicle, 2) determining the order of service (sequence) of these users into a vehicle route, and 3) scheduling the pick-up, transport, and drop-off along each route.' Cordeau and Laporte, based on their study of previous work, note that some algorithms perform these steps sequentially, whereas other models "intertwine these decisions." Based on the comparison by Cordeau and Laporte (2003), several references are especially interesting and will be mentioned more in-depth. This will focus on multi-vehicle, static algorithms with varying objective functions, constraints and/other algorithms. Single-vehicle models will not be considered further since they are less relevant to the topic at hand.

Bodin and Sexton (1986) provide a model which is according to the three steps described by Cordeau and Laporte (2003). Users can specify a desired pick-up or drop-off time. The objective function is a weighted function with penalties for both deviations from the desired delivery time and excess ride time (compared to the minimum ride time between two points). The number of vehicles and capacity can be specified. As mentioned, the algorithm starts with forming 'clusters' by assigning customers to vehicles. For each resulting cluster, a route and schedule need to be generated, which is thus based on a single-vehicle problem. For each route, the total customer-inconvenience is determined. Following that, the so-called swapper algorithm attempts to reassign customers to different vehicles to reduce customer inconvenience. This is an iterative process until no more improvements are found. Based on the new clusters, the routes and schedules are determined again. This algorithm was tested on a case with 85 users.

Jaw et al. (1986) have developed a heuristic algorithm for the multi-vehicle, static DARP including time windows that addresses "the most applicable and realistic version of the real-world problem" and at low computation cost. The objective function minimises a "non-linear combination of several types of disutility." The time-windows could either be provided for the pick-up or drop-off with actual pick-up not earlier than the specified time and actual drop-off not later than the specified time. Vehicle capacity was an implemented constraint. Especially interesting was the constraint that "actual ride time cannot exceed a given percentage of the minimum ride time." When considering the multi-modal transport envisioned for UAM, where riders take a car to and/or from a vertiport in addition to the flight itself, the total travel time, including formalities, should still be (much) faster than the alternative drive-only route. This type of constraint could account for such a factor. Next to that, the deviation from the specified time was bounded as well in a constraint, instead of being incorporated in the objective function like in Bodin and Sexton (1986). The algorithm uses heuristics and broadly follows the steps as described by Cordeau and Laporte (2003). After the customer requests are all known, the rides are indexed based upon the "earliest pick-up times." Subsequently, the customers are inserted into the schedule of a certain vehicle until all rides are assigned. This algorithm is also referred to as the 'insertion process'. For each customer, all feasible insertion possibilities are checked for each vehicle schedule. If it is infeasible to assign a certain customer to a certain vehicle, the next vehicle is checked. Next, the insertion of a customer for each vehicle is sought, which results in the minimum additional cost. The different costs of each vehicle are compared, and the customer is finally assigned to the vehicle with the lowest cost. If no feasible insertion is possible, the customer is rejected. The authors also note it is possible to not reject any customers if a vehicle can be added whenever there is an unfeasible condition. The algorithm was tested successfully on a dataset with 2617 users.

Further research on the generation of the clusters has been done by Ioachim et al. (1995). An optimisation model was developed to also explicitly determine the clusters, something the authors found lacking in previous literature because "if this process is performed too simplistically, the overall approach suffers greatly, even if the routing phase is solved optimally." Instead of making large clusters for each vehicle, like in Bodin and Sexton (1986), a larger number of 'local mini-clusters' is generated based on requests which are close together in both distance and time. These mini-clusters already take into account some constraints such as time-windows and vehicle capacity. As a result, 'both the size and the complexity of the problem are reduced.' In terms of relative performance of the clustering approach by e.g. Bodin and Sexton (1986) and the insertion process by e.g. Jaw et al. (1986), Ioachim et al. (1995) demonstrate a 5.4% reduction in total travel time and 4.9% reduction in the number of vehicles in the case of 2500 users due to the mini-clustering viz-a-viz the insertion method, making it a model fit for the scale of real-life cases. The insertion method, however, had faster computation times.

Calvo and Colorni (2007) also provide an algorithm using heuristics, which is able 'to tackle real-life in-

stances.' Most interesting is the different objective function. Instead of only minimising cost or customer inconvenience, like Jaw et al. (1986), Toth and Vigo (2000) and Bodin and Sexton (1986), Calvo and Colorni propose a weighted objective function which maximises both the number of customers served and level of service. The latter is defined as the 'ratio between the travel time proposed by the model and the minimum (theoretical) time to travel on the same origin-destination pair.' For cases where detailed cost information may not be available, this objective function may be very relevant for a company. It maximises the number of customers served and also takes into account the service level. According to the authors, the weighting ratio in the objective function "is high enough to guarantee that maximising the number of serviced customers is the main objective." The algorithm furthermore provides the opportunity to specify time windows for both pick-up and drop-off. Vehicle capacity is a constraint. The model consists of various parts. First of all, an auxiliary graph is generated which has one node instead of two for each customer. Arcs connecting these nodes have a cost function that represents both time and spatial distance. In order to generate this auxiliary graph, "several non-trivial transformations" are performed, "i.e. constraints are relaxed and used to build the cost function." The next step is to determine the minimum amount of vehicles required. Following that, the nodes with vehicles are added, and the auxiliary graph is complete. On this new graph, the assignment problem algorithm is applied. As some constraints have already been implemented in the development of this auxiliary graph, these can now be relaxed. This includes the time window constraints. The solution of the assignment problem algorithm provides a set of cycles: either main paths or 'subtours'. Main paths are "chains of customers starting and ending with the node depot" whereas the subtours "contain only customers." The objective function should yield many main paths and only a few subtours. These subtours need to be inserted in a main path to solve the entire problem. There are two procedures; 'trivial' and 'smart'. Both provide feasible routes. However, the 'smart' procedure takes into account the effect of a routing under consideration on the entire objective function. The use of the 'smart' procedure leads to improved service levels, and more customers served. The entire model was tested for up to 180 users.

The fact that Calvo and Colorni (2007) assign arcs a cost function based on both time and spatial distance caters to Garaix et al. (2010), who note that most models for vehicle assignment have a drawback. When determining paths between nodes, usually the shortest path approach is taken whereby either minimum travel time or shortest distance is considered. Especially in on-demand transportation, Garaix et al. consider this to be possibly disadvantageous in light of other possible quality of service objectives. Depending on the requirements of the service, it is thus possible to tweak the arc cost to account for multiple factors.

As described, some models have incorporated multiple factors in the objective function to account for different aspects of the problem. Chevrier et al. (2012) acknowledge this tactic but present an entirely different approach based on evolutionary multi-objective optimisation. As the name implies, there are multiple objectives, which may be in conflict with each other. In the case of Chen et al., these are minimising the number of vehicles used, delays and journey duration. The multi-objective optimisation falls under vector optimisation. An objective vector and set of feasible solutions are generated, and the resulting outcomes of the objective functions are compared to get the most optimal combination. Several solving algorithms are available, which are 'white-boxes' where "problem-related components need to be defined". A downside of this method is that "large-size problem instances cannot generally be solved by exact methods."

1.3.7. Solving Techniques

In the preceding subsections, various solving techniques have already been presented. Parragh et al. (2008) provide an overview of the different solution methods and their characteristics. The authors distinguish between exact methods, heuristic and metaheuristic methods. Exact methods will always find the global optimum. However, for large scale problems this leads to high computational efforts (Kramer, 2017). Heuristics do not necessarily yield optimality but "generate good solutions in a reasonable computational time." Heuristics and metaheuristics are usually used for so-called (NP-)hard problems. These algorithms are adapted to a certain problem, and 'the effectiveness depends upon its ability to avoid entrapment at local optima' (Festa, 2014). Metaheuristics "combine basic heuristic methods to efficiently explore the set of feasible solutions."

Heuristics, which include the insertion procedures and cluster generation, have already been explained and will not be explained further. For the methods which have not been treated in the previous sections, this subsection will provide more background on the solving techniques used in literature, specifically for the static, multi-vehicle DARP problem, with time windows. The constraints and objective of the related models are generally disregarded, as there is much overlap with the models and concepts already explained. It is to be noted that some models also combine various solving techniques.

The exact methods include branch and bound, which have been used by Cordeau (2006) and Ropke et al.

(2007). Cordeau (2006) notes that brand and bound (or branch and cut) have been successfully applied previously to other routing problems. It consists of two parts. In the branching part, the set of possible solutions is expanded in a tree structure, leading to nodes with subsets (Vlastelica Pogančić, 2019). The subsets have a lower and upper bound. The algorithm goes through the various subsets to compare the subset bounds with a global bound. If the subset bounds are worse than the global bound, the node is disregarded for further exploration. Finally, the algorithm converges to the most optimum solution. While the branch and bound algorithm is automatically used to solve Mixed Integer Linear Programming (MILP) problems in solvers such as *Gurobi* and *CPLEX*, Cordeau (2006) and Ropke et al. (2007) provide an explicit algorithm for branching and bounding (Gurobi, 2020, IBM, 2020). The nodes are generated using inequalities for 'time, load, subtour elimination, precedence, order, and infeasible path constraints.' For the cuts, a heuristic is added, which searches for violated inequalities. This method is subsequently compared with the build-in algorithm of *CPLEX*. The authors found that the explicit method reduced CPU time and the number of nodes explored in the tree. However, this was tested only up to 48 requests and on a system with limited computational power compared to what is available today. Cordeau concludes that "the methodology proposed in this paper obviously cannot be used to solve large-scale instances containing hundreds or thousands of users, as is sometimes the case in large cities." Computation power has been improved since the article, so it is possible the explicit method may be able to solve large instances. However, computational solvers have also improved, so it remains to be seen whether the exact method or build-in algorithm is faster.

The metaheuristic methods encompass the tabu search and genetic algorithm. In the tabu search method, used by Cordeau and Laporte (2003), the algorithm starts at an initial solution and moves at every iteration to the current best solution near the current position, searching for the final optimum. "Solutions possessing some attributes of recently visited solutions" cannot be chosen for "a number of iterations." They are thus *tabu*. Up to 295 requests were solved with this method. A genetic algorithm was utilised by Jorgensen et al. (2007) based on previous "good performance on a number of related routing problems." A set of solutions is called a 'population', and an initial population is required to start the algorithm (Vlastelica Pogančić, 2019). In this case, the initial population was a random cluster of customers. The allocation of customers to vehicles as well as the order of riders on the routes are represented by a matrix with rows for each available vehicle and columns for customers and depots. Each cell can be allocated a binary value, called a 'gene'. A 1 denotes the customer of that column is allocated to the corresponding row of the cell. Each row is called a 'chromosome' and thus a possible routing. Two 'parents' (chromosomes) are then chosen using both a random and stochastic procedure. In each iteration, based on these two parents, one 'offspring' is generated where crossover takes place. In the crossover, genes from two parents are exchanged to form a new chromosome. I.e. customers are swapped to form a new route. This offspring then replaces another chromosome which has a low 'fitness value'; in this case, a high cost. As a result, the "best member of a current population is guaranteed to survive to the next population." Mutations can move customers between chromosomes at random. It is noted that the population size "greatly influences the performance" of the algorithm. A too small size may yield "an under-covered solution space" while a too large population may lead to slow computation. Jorgensen et al. chose for a constant population size (number of solutions), which is 'conventional'. Furthermore, the algorithm stops "after a fixed number of iterations." After this process is completed, there are various clusters of customers. The second phase will construct routes, taking into account costs. The latter in this case is a weighted combination of travel time and time window violations. The algorithm has been applied to the same test cases as Cordeau and Laporte (2003), up to 144 requests. Interestingly, the results of the genetic algorithm were worse for ride time and vehicle waiting time. However, the total route duration was lower compared to the aforementioned tabu search method.

Comparing all methods, Parragh et al. (2008) present that the models by Jaw et al. (1986), Ioachim et al. (1995) and Borndörfer et al. (1999), using the heuristic methods of insertion and clustering algorithms, by far, allow solving for the most requests. Considering these models have been developed quite a while ago, it is difficult to assess what the relative performance using today's mathematical solvers and computing power would be.

1.3.8. UAM Scheduling Model

One scheduling model specific to UAM operations has been found in literature. Shihab and Wei (2019) provide a model "to decide which type of scheduling to offer, how to dispatch the fleet and schedule operations, based on simulated market demand." In terms of services types, on-demand, scheduled and a hybrid combination are investigated using an objective that maximises profit. The network model is represented by 3 vertiports, connected through 6 arcs for the bidirectional routes. The distances are assumed to be 60-80 miles

(96-128 *km*) which are on the high side compared to the ranges presented in section 1.2.2. A space-time network depicts both the temporal and spatial aspects using a graph with nodes, ground and flight arcs, and time intervals of 30 minutes. The demand model is based on previous studies, using "data sources such as helicopter charter services, census data, and consumer wealth data" for Los Angeles and San Francisco. The demand is further divided into commute, family/personal, school/church and recreation/other. This yields an hourly demand forecast for the different routes. The demand model, fleet and network data are combined in the optimisation model, which is Mixed Integer Quadratic Programming. The decision variable is a binary value for the presence of a vehicle on an arc. In the objective function, revenues and costs are incorporated. For the revenue, the product of the fare and passengers transported is summed over all routes. The total cost is computed using the cost per arc and the number of times each arc is traversed. The demand fluctuates according to the fare charges using a demand curve coefficient. The vehicle capacity (4 passengers), flow conservation and start/end at a depot are constraints, which are comparable with models for automotive transportation. UAM-specific constraints include a limit on the vertiport capacity at each instance and several energy-related constraints, which will be elaborated upon in section 1.3.9. The specific constraints imposed differ for the various services. Also, the horizons differ; the scheduling problem - generating timetables in advance for flight departures - is simulated for a week while the dispatch problem - allocating vehicles in the short-term to real-time trip requests - is run for 3 hours. A commercial solver, *Gurobi*, is used to solve the model for 20 eVTOL vehicles. The outcomes are to be compared for various metrics, such as profit and demand served. Also, the input parameters such as "the scheduling horizon, time discretization scheme, eVTOL characteristics, and route connectivity" will be used in a sensitivity analysis to investigate the impact on the objective function outcome and computational cost. Numerical results of the model are not yet known.

1.3.9. Incorporation of Battery-powered Vehicles

Further to the more generic vehicle scheduling models presented earlier in this chapter, the incorporation of electrically-powered transportation imposes additional considerations. In the automotive sector, electrical (autonomous) vehicles are a key area of interest for both researchers and industry. Hence, there is a separate class of the VRP, the Electric Vehicle Routing Problem (with Time Windows) [EVRP(TW)], that already provides a base for constraints that account for battery requirements and charging. A selection of relevant and comparable approaches will be provided. Furthermore, some battery models specific to UAM operations will also be discussed.

First of all, a major difference between the automotive EVRP and the prospective electrical UAM operations, is that EV's usually need to drive to a dedicated charging station, which may not be part of their passenger routing. Meanwhile, with eVTOL vehicles that will normally only fly between vertiports, the charging locations are fixed. Though, depending on the specific infrastructure and other operational considerations of a city, not all vertiports may be equipped with charging facilities. Either way, in the automotive EVRP, additional points need to be added to account for charging facilities in addition to passenger origin/destination locations.

Afroditi et al. (2014) have defined the EVRP with "industry constraints." A notable addition compared to the basic VRP(TW) is a set that represents visits to charging stations, which are separate from the depots. Furthermore, the charge of the battery is defined and the remaining electric charge is a variable. In terms of new constraints, coming from a depot or charging station the battery is always full. If in transport, the battery charge level decreases with distance in another constraint, according to a consumption rate. Finally, a constraint is added that ensures that the vehicle always has sufficient charge left to return to either a depot or charging station. Details on the computation or value of the consumption rate are not provided, however. Afroditi et al. also do not provide numerical computations. Schneider et al. (2014) do have numerical results and use exactly the same constraint with a consumption rate multiplied by the distance travelled to calculate the battery discharge. Nevertheless, the authors state they have set the value of the discharge rate to 1.0 "for sake of simplicity".

For illustration, a possible energy consumption equation for a vehicle is provided by Afroditi et al. (2014), but this does not scale with distance to get energy as in the constraint. The energy consumption depends on the average speed, weight, drag coefficient, road gradient, air density and frontal area of the vehicle. Lin et al. (2016) take an even more elaborate computation approach for the energy use, using engine efficiency, rolling resistance, rolling drag and acceleration in consideration in addition to the parameters also used by Afroditi et al. (2014). This requires extensive knowledge of the vehicle and environmental conditions.

Liang et al. (2016) use a constraint in their model for electric taxis that compares the average driving distance of an operation period (which can be more than a single ride) with the driving range of a full battery.

If the battery range is exceeded, the constraint imposes a longer time at the station with charging facilities for the vehicle to charge for the additional battery charge required. This is a simplification as the taxis are regarded as flows instead of specific vehicles. Thus not the specific battery charge of each vehicle during an operation period is considered.

De Almeida Correia and Santos (2014) take a different approach for EV's in a car rental fleet where "the charge for each time instant of the optimization period" is computed, taking into account both charging and driving (i.e. discharging). Another constraint imposes that the charge is between 0 and a maximum at all time instants. Together, these constraints ensure that vehicles are assigned a charging moment accordingly. In the model, the battery charge level and charging speeds are expressed in time steps such that the discharge is proportional to the driving time.

The model by Shihab and Wei (2019), as stated, has several constraints to account for the electrical propulsion mechanism of the UAM vehicle under consideration. First of all, vehicles have a full charge at the starting time of the scenario. For the other time instances, the charge is either full if the vehicle arrived from a recharging arc in the previous time interval, half full if it arrived from a flight arc with a preceding charging arc, or empty (zero charge) otherwise. This constraint causes the quadratic behaviour of the model. Subsequently, a constraint imposed only enables assignment of a certain vehicle to a flight arc if the energy level matches or exceeds the level required for completion of that flight arc. This algorithm, thus, only considers a limited set of possible battery charge levels; full, half-full or empty. Yet, these fixed values are compared with the energy required for a certain flight arc, which is an input parameter. If the energy required is thus known, the battery levels after performing a flight on a certain arc may be computed more accurately. A possible benefit of computing the energy level after discharge more accurately is the ability to take smaller time intervals for the model. The algorithm of Shihab and Wei is based on intervals of 30 minutes which represents the maximum possible recharging period. More precise values for the charging time of a battery to full consequently also determines the minimum required idle time at a vertiport for that charging, similar to the automotive example of Liang et al. (2016). This can result in charging times less than 30 minutes and, conversely, smaller time intervals which may open up additional solutions and performance improvements.

In terms of battery discharge, Kleinbekman et al. (2018) consider a model for the State of Charge (SOC) during a mission which is based on the total current of all battery cells at a certain time step, the battery capacity, time step and SOC of the previous time step. The total battery cell current is computed by dividing the electric power demand by the nominal battery voltage. For the electrical power demand, the power required is corrected for rotor and mechanical efficiency and a safety factor to account for e.g. "weather conditions and emergency diversions."

A precise, and possibly dynamic, computation of the battery levels may improve accuracy for charging/discharging times and, subsequently, vehicle scheduling. As noted by Alnaqeb et al. (2018), it is also vital that UAM vehicles are able to compute before the flight, but also in real-time and during operation, if battery capacity is sufficient to complete a flight. Nevertheless, accurate computation is dependent on many factors. Next to technical and vehicle/battery specifications also factors such as weather, routing, payload etc. Hence for a vehicle scheduling model, the approach of using fixed, pre-determined energy level requirements for certain routings, as used by Shihab and Wei (2019), or the approach by De Almeida Correia and Santos (2014) where battery discharge scales with time, are fair assumptions.

1.4. Concluding Remarks

In the preceding (sub)sections, existing literature from academia, industry and media has been presented to introduce the topic at hand and to provide an overview of the historic and state-of-the-art body of knowledge, and of future projections in the field of UAM. As the primary goal of the thesis is to provide a mathematical model - which will be elaborated upon in chapter 2 - the mathematical models and algorithms from section 1.3 form the base for further development. Naturally, none of the models already provides all the desired aspects and parameters for the thesis at hand. For instance, only one vehicle scheduling model specifically for UAM has been found. The electrical aspect, where battery capacity is limited and charging may need to happen between flights, is a factor that requires considerate attention. This topic has already been partially addressed in the context of automotive electrical vehicles. Also, in terms of safety, a diversion to another landing spot must always be possible and planned for. Furthermore, UAM flights are most likely part of a multimodal transport model whereby transport to/from the vertiport needs to be considered in the scheduling as well. The transfers introduce both additional travel time but also uncertainty. Considering pooling passengers is an important step to make UAM affordable, the trade-off between waiting for additional passengers to take a flight and losing the time gain due to waiting is an interesting aspect as well.

2

Research Plan

This chapter encompasses the research plan for a Master of Science thesis on the topic of aerial ride-sharing operations. It serves as a base for the planning and subsequent execution of the thesis. This is realised, first of all, by setting out the research objective and research questions in section 2.1. Next, the proposed methodology is explained in section 2.2. Subsequently, remarks on the results, outcome and relevance are provided in section 2.3. This also includes the approach for verification and validation. Finally, the planning of the project is displayed by means of a Gantt chart in section 2.4.

2.1. Research Objective and Questions

In this section, the research objective, as well as the research questions, will be presented sequentially.

2.1.1. Research Objective

Based on the thesis assignment and the preceding literature evaluation, the project goal has been defined. The aim of the research is:

> To provide a framework of recommendations for implementation of aerial ride-sharing by determining parameters relevant for aerial ride-sharing operations, incorporating these into a mathematical model that optimises vehicle allocation considering a number of constraints, testing this model, and evaluating the outcomes.

As can be deduced from the research aim, there are multiple sub-goals to be met. First of all, an overview of relevant parameters for aerial ride-sharing as well as differences vis-à-vis automotive ride-sharing (models) needs to be provided. The second subgoal is including these findings into a supply-demand match/vehicle allocation optimisation model. Thirdly, various scenarios and indicative values of the parameters need to be determined to provide a broad range of results, considering there is no clear or single concept for UAM ride-sharing. The fourth subgoal will be to evaluate the optimisation model using these scenarios for various performance metrics. And finally, based upon the results of the evaluation, recommendations are to be made for a better implementation of aerial ride-sharing.

2.1.2. Research Question

The main research question is formulated as:

> How can aerial ride-sharing be incorporated in a mathematical optimisation model which matches riders and rides and which is subject to operational and technical constraints?

To answer the main research question, various core questions have been generated with sub-questions as well, following Verschuren and Doorewaard (2010). The core questions correspond with the four subgoals presented in section 2.1.1:

1. What are the main differences between automotive ride-sharing and aerial ride-sharing?
   (a) What ride-sharing concepts exist in the automotive sector?

(b) What transportation concepts are envisioned for UAM?

(c) What operational and technical constraints are specific to UAM?

2. How can the mathematical optimisation model be formulated?

(a) What mathematical models do currently exist for ride-sharing?

(b) What characteristics of aerial ride-sharing are missing in existing models?

(c) What objective functions, decision variables, and constraints are relevant for the model for aerial ride-sharing?

3. What are representative scenarios and parameter values for future aerial ride-sharing operations?

(a) What are representative values for:

- VTOL vehicle passenger capacity
- VTOL vehicle speed
- VTOL vehicle range
- eVTOL battery capacity
- Battery charging and discharging

(b) What models exist for UAM passenger transport demand?

(c) What lay-outs are available for vertiport locations and capacity?

4. What is the sensitivity and performance of the model?

(a) What is the effect of the allowable waiting time on the model's outcome?

(b) What is the effect of the charging time on the model's outcome?

(c) What is the effect of different weighting factors on the outcome of the combined objective function?

(d) What is the performance of the model with regards to computational time?

2.2. Methodology

In order to reach the project goal and subgoals presented in section 2.1.1, and ultimately answer the research questions from section 2.1.2, extensive use will be made of the literature study from chapter 1. Especially for the first and third subgoal, data from the articles and journal papers of section 1.1 and section 1.2 will be used. It is acknowledged that parameters relevant to UAM may still be uncertain or ambiguous. The choice for certain values will, therefore, be made after a careful comparison, and motivated where possible. Also, these input parameters may be an adjustable variable to allow for a flexible model. Remarks on the verification and validation will be provided in section 2.3.

Also, the development of the mathematical optimisation model for the second subgoal relies partly on the findings from the literature study. Models and algorithms from section 1.3 will be used in a hybrid approach to incorporate various aspects that are applicable to aerial ride-sharing. As mentioned in section 1.4, the existing literature is incomplete hence the model development also encompasses considerable own input. As such, the approach taken can be considered as theory development (Verschuren and Doorewaard, 2010). The constraints and parameters specific to UAM, as identified for the first subgoal, will be converted into mathematical expressions and added to the model. The model itself will be based on MILP, following industry/academic practice and the aforementioned literature. MILP maximises or minimises a certain objective function, which is subject to constraints, and provides values for variables of interest (decision variables) which yield an optimal solution. The model of this thesis will be static, i.e. all demand requests are known before the algorithm is run. More details on the programming, software and hardware set-up will be provided in the following section. Considering the fact that cost and revenue parameters are unknown at this point in time and are highly dependent on the operator, location and other factors, it is proposed that the objective function is more universal. Following existing research, a weighted objective function incorporating the number of riders served and a convenience factor, such as minimum time deviation/maximum time saved, is considered. The effect of the relative weights can subsequently be investigated.

2.3. Results, Outcome and Relevance

As proposed in section 2.2, a MILP-model is to be created which involves an objective function subject to various constraints. (Input) data used for this model are (e)VTOL vehicle parameters which come from (industry) literature as well as UAM-infrastructure and demand models from existing papers, following core question 3) from section 2.1.2. The main variable will be a binary variable that expresses if a vehicle is assigned a certain flight request.

Since the goal is to provide recommendations for UAM ride-sharing implementation, various qualitative output parameters are of interest. The objective function, as proposed, maximises the users served while also maximising a convenience factor. The effect on both these objectives by altering certain input parameters are results to be investigated. Also, the effect of the weighting factors on the objective function outcome is a main result. Based on these qualitative results, the intention is to provide recommendations to arrive at an optimal and feasible mix of technical and operational aspects for UAM operations.

The findings of the study will be relevant to both industry and academia. The overview of UAM-specific factors in ride-sharing expands the current body of knowledge and can be used for further research on the topic. The model itself can form a base for a broad range of follow-up research projects, diving either deeper into the effects of a single aspect or broadening the model to account for other factors. The quantitative results, as well as the (qualitative) recommendations, can be used by industry in the planning for UAM operations. This includes both the technical domain, such as infrastructure requirements and locations, and also the business model, such as when to offer UAM-based transportation compared with ground-based rides.

In order to check that there are no computation errors in the model, verification will be performed. This will consist of various steps including code verification which aims to find, and fix, programming errors such as incorrect syntax. Subsequently, calculation verification is key to find if the model itself contains errors. Where possible, computations of parameters will be compared with manual calculations. The exact (sub-)model verification method will be dependent on the models used. If an existing model is used, an example scenario from the original literature can be used to check if it has been properly programmed. If multiple models are combined, first of all each distinct model will be verified for correctness separately. After that, the interaction and links between the various parts will be verified to see if parameters properly carry over. Finally, the various model parts will be assembled in order and tested at each step using a simplified reference case for which the solution can be determined separately. A sanity check on all relevant outputs will lastly determine if the values are realistic.

Considering this research relates to future operations, there is currently no real-world data available to compare the outcomes with. As a result, no validation is possible.

2.4. Project Planning and Gantt Chart

The entire thesis work is comprised of two main parts; the literature study, and the research and writing of the thesis itself. The ATO section has provided a timeline for a typical thesis project. There are also various formal milestones such as the kick-off, mid-term and green light meeting. Based on the typical durations, a planning has been made in a Gantt chart. This also accounts for personal absence such as holidays and other obligations.

After approval of this literature review and research plan, the initial research will commence. Based on the literature, an appropriate starting model and various scenarios will be determined. These will be further developed into a programming model. Undoubtedly the literature study will not be sufficient, so concurrently additional literature research will be performed. A design-freeze point is planned where the development of the model will be halted. Based on the model and scenarios available at that time, data will be generated and analysed. This will prevent an incomplete research overview during the mid-term. After the mid-term, comments on the research will be implemented. It is likely that the model has to be adapted, corrected or otherwise it can be further expanded. Again, a design freeze will be held after which the verification process will start. As noted earlier, there is no real-world case as a benchmark, so no validation will be performed. After the verification process, the scenarios will be run another time to generate the final data. Subsequently, the data will be analysed and formatted to form the thesis report. After which all formal parts of the thesis will be written.

Bibliography

[1] Anagnostopoulou Afroditi, Maria Boile, Sotirios Theofanis, Eleftherios Sdoukopoulos, and Dimitrios Margaritis. Electric Vehicle Routing Problem with industry constraints: Trends and insights for future research. In *Transportation Research Procedia*, volume 3, pages 452–459, Sevilla, Spain, 2014. Elsevier B.V. doi: 10.1016/j.trpro.2014.10.026. URL http://dx.doi.org/10.1016/j.trpro.2014.10.026.

[2] Niels Agatz, Alan Erera, Martin Savelsbergh, and Xing Wang. Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223(2):295–303, 2012. ISSN 03772217. doi: 10.1016/j.ejor.2012.05.028. URL http://dx.doi.org/10.1016/j.ejor.2012.05.028.

[3] Airbus. Vahana - Vehicle Demonstrators - Airbus, 2020. URL https://www.airbus.com/innovation/urban-air-mobility/vehicle-demonstrators/vahana.html.

[4] Abdullah H. Alnaqeb, Yifei Li, Yu Hui Lui, Priyank Pradeep, Joshua Wallin, Chao Hu, Shan Hu, and Peng Wei. Online prediction of battery discharge and flight mission assessment for electrical rotorcraft. *AIAA Aerospace Sciences Meeting, 2018*, (210059):1–12, 2018. doi: 10.2514/6.2018-2005.

[5] Andrew Amey. A Proposed Methodology for Estimating Rideshare Viability within an Organization , applied to the MIT Community. In *Transportation Research Board 90th Annual Meeting. No. 11-2585. 2011*, number April 2010, pages 1–16, 2011.

[6] Audi Mediencenter. Audi, Airbus and Italdesign test Flying Taxi Concept, 2018. URL https://www.audi-mediacenter.com/en/press-releases/audi-airbus-and-italdesign-test-flying-taxi-concept-11019.

[7] Aurora Flight Sciences. PAV - eVTOL Passenger Air Vehicle - Aurora Flight Sciences, 2020. URL https://www.aurora.aero/pav-evtol-passenger-air-vehicle/.

[8] Milos Balac, Raoul L. Rothfeld, and Sebastian Hörl. The prospects of on-demand Urban Air Mobility in Zurich, Switzerland. 2011. URL https://doi.org/10.3929/ethz-a-010025751.

[9] Milos Balać, Amedeo R. Vetrella, and Kay W. Axhausen. Towards the integration of aerial transportation in urban settings. Aug 2017. URL https://www.research-collection.ethz.ch/handle/20.500.11850/175251.

[10] Gerardo Berbeglia, Jean-François Cordeau, Irina Gribkovskaia, and Gilbert Laporte. Static Pickup and Delivery Problems: a classification scheme and survey. *Top*, 15(1):1–31, 2007. ISSN 1134-5764. doi: 10.1007/s11750-007-0009-0.

[11] Lawrence Bodin and Thomas R Sexton. The multi-vehicle subscriber Dial-a-Ride Problem. *TIMS Studies in the Management Sciences*, 22:73–86, 1986.

[12] Jan-Hendrik Boelens and Volocopter. Pioneering the Urban Air Taxi Revolution. Technical report, Bruchsal, Germany, 2019. URL https://press.volocopter.com/images/pdf/Volocopter-WhitePaper-1-0.pdf.

[13] Booz Allen Hamilton. Urban Air Mobility (UAM) Market Study. Technical report, McLean, VA, 2018. URL https://ntrs.nasa.gov/search.jsp?R=20190001472.

[14] R. Borndörfer, M. Grötschel, F. Klostermeier, and C. Küttner. Telebus Berlin: Vehicle Scheduling in a Dial-a-Ride System. (July 1995):391–422, 1999. doi: 10.1007/978-3-642-85970-0_19.

[15] Stefan Bunte and Natalia Kliewer. An overview on vehicle scheduling models. *Public Transport*, 1(4): 299–317, 2009. ISSN 1866749X. doi: 10.1007/s12469-010-0018-5.

[16] Roberto Wolfler Calvo and Alberto Colorni. An effective and fast heuristic for the Dial-a-Ride problem. *4or*, 5(1):61–73, 2007. ISSN 16142411. doi: 10.1007/s10288-006-0018-0.

[17] G. Carpaneto, M. Dell'amico, M. Fischetti, and P. Toth. A branch and bound algorithm for the multiple depot vehicle scheduling problem. *Networks*, 19(5):531–548, 1989. ISSN 10970037. doi: 10.1002/net.3230190505.

[18] Nelson D. Chan and Susan A. Shaheen. Ridesharing in North America: Past, Present, and Future. *Transport Reviews*, 32(1):93–112, 2012. ISSN 01441647. doi: 10.1080/01441647.2011.621557.

[19] Chung Min Chen, David Shallcross, Yung Chien Shih, Yen Ching Wu, Sheng Po Kuo, Yuan Ying Hsu, Yuhsiang Holderby, and William Chou. Smart ride share with flexible route matching. *International Conference on Advanced Communication Technology, ICACT*, pages 1506–1510, 2011. ISSN 17389445.

[20] Rémy Chevrier, Arnaud Liefooghe, Laetitia Jourdan, and Clarisse Dhaenens. Solving a Dial-a-Ride Problem with a hybrid evolutionary multi-objective approach: Application to demand responsive transport. *Applied Soft Computing Journal*, 2012. ISSN 15684946. doi: 10.1016/j.asoc.2011.12.014.

[21] Jean François Cordeau. A branch-and-cut algorithm for the Dial-a-Ride Problem. *Operations Research*, 54(3):573–586, 2006. ISSN 0030364X. doi: 10.1287/opre.1060.0283.

[22] Jean François Cordeau and Gilbert Laporte. The Dial-a-Ride Problem (DARP): Variants, modeling issues and algorithms. *4or Quarterly Journal of the Belgian, French and Italian Operations Research Societies*, 1 (2):89–101, 2003. ISSN 16142411. doi: 10.1007/s10288-002-0009-8.

[23] Jean François Cordeau and Gilbert Laporte. A tabu search heuristic for the static multi-vehicle Dial-a-Ride Problem. *Transportation Research Part B: Methodological*, 37:579–594, 2003. ISSN 01912615. doi: 10.1016/S0191-2615(02)00045-0.

[24] Jean François Cordeau and Gilbert Laporte. The Dial-a-Ride Problem: Models and algorithms. *Annals of Operations Research*, 153(1):29–46, 2007. ISSN 02545330. doi: 10.1007/s10479-007-0170-8.

[25] Gonçalo Homem De Almeida Correia and Raquel Filipa Gonçalves Santos. Optimizing the use of electric vehicles in a regional car rental fleet. *Transportation Research Record*, 2454:76–83, 2014. ISSN 03611981. doi: 10.3141/2454-10.

[26] Guy Desaulniers and Daniel Villeneuve. Shortest path problem with Time Windows and linear waiting costs. *Transportation Science*, 34(3):312–319, 2000. ISSN 00411655. doi: 10.1287/trsc.34.3.312.12298.

[27] Michael J. Duffy, Sean Wakayama, Ryan Hupp, Roger Lacy, and Matt Stauffer. A study in reducing the cost of vertical flight with electric propulsion. In *Annual Forum Proceedings - AHS International*, pages 187–204, Denver, CO, 2017. AIAA. doi: 10.2514/6.2017-3442.

[28] EHang. White Paper on Urban Air Mobility Systems. Technical report, 2020. URL https://www.ehang.com/app/en/EHangWhitePaperonUrbanAirMobilitySystems.pdf.

[29] EHang. EHang | UAM - Passenger Autonomous Aerial Vehicle (AAV), 2020. URL https://www.ehang.com/ehang184/.

[30] P Festa. A brief Introduction to Exact, Approximation and Heuristic Algorithms for Solving Hard Combinatorial Optimization Problems. In *2014 16th International Conference on Transparent Optical Networks (ICTON)*, pages 1–20. IEEE, 2014. ISBN 9781479956012. doi: 10.1109/ICTON.2014.6876285.

[31] Masabumi Furuhata, Maged Dessouky, Fernando Ordóñez, Marc Etienne Brunet, Xiaoqing Wang, and Sven Koenig. Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, 2013. ISSN 01912615. doi: 10.1016/j.trb.2013.08.012.

[32] Thierry Garaix, Christian Artigues, Dominique Feillet, and Didier Josselin. Vehicle routing problems with alternative paths: An application to on-demand transportation. *European Journal of Operational Research*, 2010. ISSN 03772217. doi: 10.1016/j.ejor.2009.10.002.

[33] B. Gavish and E. Shlifer. An approach for solving a class of transportation scheduling problems. *European Journal of Operational Research*, 1979. ISSN 03772217. doi: 10.1016/0377-2217(79)90098-5.

[34] Gurobi. Gurobi Optimizer, 2020. URL https://www.gurobi.com/.

[35] Robert V. Hemm, Denise Duncan, and Virginia L. Stouffer. On-Demand Aviation Regulatory Obstacles and Resulting Research Roadmaps. In *16th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, Virginia, Jun 2016. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-440-4. doi: 10.2514/6.2016-3301. URL http://arc.aiaa.org/doi/10.2514/6.2016-3301.

[36] Jeff Holden, Nikhil Goel, and Uber Elevate. Fast-Forwarding to a Future of On-Demand Urban Air Transportation. Technical report, 2016. URL https://www.uber.com/elevate.pdf.

[37] David Ibekwe. Dubai just tested its flying drone taxi for the first time, 2017. URL https://www.businessinsider.com/dubai-flying-taxi-drone-volocopter-test-pilot-unmanned-2017-9?international=true&r=US&IR=T.

[38] IBM. CPLEX, 2020. URL https://www.ibm.com/support/knowledgecenter/SSSA5P_12.10.0/ilog.odms.cplex.help/refcppcplex/html/.

[39] Ramon Iglesias, Federico Rossi, Kevin Wang, David Hallac, Jure Leskovec, and Marco Pavone. Data-Driven Model Predictive Control of Autonomous Mobility-on-Demand Systems. In *International Conference on Robotics and Automation 2018*, sep 2017. URL https://arxiv.org/abs/1709.07032.

[40] Irina Ioachim, Jacques Desrosiers, Yvan Dumas, Marius M. Solomon, and Daniel Villeneuve. A Request Clustering Algorithm for Door-to-Door Handicapped Transportation. *Transportation Science*, 29(1):63–78, 1995. ISSN 0041-1655. doi: 10.1287/trsc.29.1.63.

[41] Jang Jei Jaw, Amedeo R. Odoni, Harilaos N. Psaraftis, and Nigel H.M. Wilson. A heuristic algorithm for the multi-vehicle advance request Dial-a-Ride Problem with Time Windows. *Transportation Research Part B*, 1986. ISSN 01912615. doi: 10.1016/0191-2615(86)90020-2.

[42] Joby Aviation. Joby Aviation, 2020. URL https://www.jobyaviation.com/.

[43] Wayne Johnson and Christopher Silva. Observations from exploration of VTOL Urban Air Mobility designs. In *7th Asian/Australian Rotorcraft Forum, ARF 2018*, 2019. URL https://ntrs.nasa.gov/search.jsp?R=20180007847.

[44] R. M. Jorgensen, J. Larsen, and K. B. Bergvinsdottir. Solving the Dial-a-Ride Problem using genetic algorithms. *Journal of the Operational Research Society*, 58(10):1321–1331, 2007. ISSN 01605682. doi: 10.1057/palgrave.jors.2602287.

[45] Cedric Y. Justin, Alexia P. Payan, Simon I. Briceno, and Dimitri N. Mavris. Operational and Economic Feasibility of Electric Thin Haul Transportation. In *17th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, Virginia, 2017. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-508-1. doi: 10.2514/6.2017-3283. URL https://arc.aiaa.org/doi/10.2514/6.2017-3283.

[46] Kalon Kelly. Casual Carpooling-Enhanced. *Journal of Public Transportation*, 10(4):119–130, dec 2007. ISSN 1077-291X. doi: 10.5038/2375-0901.10.4.6. URL http://scholarcommons.usf.edu/jpt/vol10/iss4/6/.

[47] Hyun D. Kim, Aaron T. Perry, and Phillip J. Ansell. A Review of Distributed Electric Propulsion Concepts for Air Vehicle Technology. *2018 AIAA/IEEE Electric Aircraft Technologies Symposium, EATS 2018*, 2018. doi: 10.2514/6.2018-4998.

[48] Kitty Hawk. Kitty Hawk, 2020. URL https://kittyhawk.aero/.

[49] Imke C. Kleinbekman, Mihaela A. Mitici, and Peng Wei. eVTOL Arrival Sequencing and Scheduling for On-Demand Urban Air Mobility. In *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*, pages 233–240, London, 2018. doi: 10.1109/DASC.2018.8569645.

[50] Lee W. Kohlman and Michael D. Patterson. System-Level Urban Air Mobility Transportation Modeling and Determination of Energy-Related Constraints. In *2018 Aviation Technology, Integration, and Operations Conference*, Reston, Virginia, 2018. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-556-2. doi: 10.2514/6.2018-3677. URL https://arc.aiaa.org/doi/10.2514/6.2018-3677.

[51] Raphael H F Ribeiro Kramer. *Exact and Heuristic Algorithms for Vehicle Routing , Scheduling and Location Problems*. PhD thesis, Università degli studi di Modena e Reggio Emilia, 2017.

[52] Xiao Liang, Gonçalo Homem de Almeida Correia, and Bart van Arem. Optimizing the service area and trip selection of an electric automated taxi system used for the last mile of train trips. *Transportation Research Part E: Logistics and Transportation Review*, 93(2016):115–129, 2016. ISSN 13665545. doi: 10.1016/j.tre.2016.05.006. URL http://dx.doi.org/10.1016/j.tre.2016.05.006.

[53] Lillium. Lillium, 2020. URL https://lilium.com/.

[54] Jane Lin, Wei Zhou, and Ouri Wolfson. Electric Vehicle Routing Problem. In *Transportation Research Procedia*, volume 12, pages 508–521, Tenerife, 2016. Elsevier B.V. ISBN 1312996242. doi: 10.1016/j.trpro.2016.02.007.

[55] McKinsey&Company, Robin Riedel, and Shivika Sahdev. Taxiing for Takeoff : the Flying Cab in Your Future. Technical report, 2019.

[56] Mark D. Moore. Distributed Electric Propulsion (DEP) Aircraft. Technical report, NASA, 2012. URL https://aero.larc.nasa.gov/files/2012/11/Distributed-Electric-Propulsion-Aircraft.pdf{%}0Ahttp://aero.larc.nasa.gov/files/2012/11/Distributed-Electric-Propulsion-Aircraft.pdf.

[57] Catherine Morency. The ambivalence of ridesharing. *Transportation*, 34(2):239–253, 2007. ISSN 00494488. doi: 10.1007/s11116-006-9101-9.

[58] Marc Narkus-Kramer. On-Demand Mobility (ODM): A Discussion of Concepts and Required Research. In *2013 Aviation Technology, Integration, and Operations Conference*, Reston, Virginia, Aug 2013. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-225-7. doi: 10.2514/6.2013-4294. URL http://arc.aiaa.org/doi/10.2514/6.2013-4294.

[59] NASA. Urban Air Mobility (UAM) Market Study. Technical Report November, 2018. URL https://www.nasa.gov/sites/default/files/atoms/files/uam-market-study-executive-summary-pr.pdf.

[60] NASA Aeronautics Research Mission Directorate (ARMD). Urban Air Mobility (UAM) Grand Challenge Industry Day, 2018. URL https://ntrs.nasa.gov/search.jsp?R=20190003422.

[61] National Air Transportation Association. Urban Air Mobility : Considerations for Vertiport Operation. Technical report, 2018.

[62] Y. Nie and Mehrnaz Ghamami. A corridor-centric approach to planning electric vehicle charging infrastructure. *Transportation Research Part B: Methodological*, 57:172–190, 2013. ISSN 01912615. doi: 10.1016/j.trb.2013.08.010. URL http://dx.doi.org/10.1016/j.trb.2013.08.010.

[63] Victoria Chibuogu Nneji, Alexander Stimpson, Mary Cummings, and Kenneth H. Goodrich. Exploring Concepts of Operations for On-Demand. In *17th AIAA Aviation Technology, Integration, and Operations Conference*, pages 1–12, Denver, CO, 2017. NASA.

[64] Mehdi Nourinejad and Matthew J. Roorda. Agent based model for dynamic ridesharing. *Transportation Research Part C: Emerging Technologies*, 64:117–132, Mar 2016. ISSN 0968-090X. doi: 10.1016/J.TRC.2015.07.016. URL https://www.sciencedirect.com/science/article/pii/S0968090X15002661.

[65] Clifford S. Orloff. Route Constrained Fleet Scheduling. *Transportation Science*, 10(2):149–168, 1976. ISSN 00411655. doi: 10.1287/trsc.10.2.149.

[66] Sophie N. Parragh, Karl F. Doerner, and Richard F. Hartl. A survey on pickup and delivery problems: Part II: Transportation between pickup and delivery locations. *Journal fur Betriebswirtschaft*, 58(2):81–117, 2008. ISSN 03449327. doi: 10.1007/s11301-008-0036-4.

[67] Michael D. Patterson, Kevin R. Antcliff, and Lee W. Kohlman. A proposed approach to studying Urban Air Mobility missions including an initial exploration of mission requirements. In *AHS International 74th Annual Forum & Technology Display*, Phoenix, AZ, 2018.

[68] Victor Pillac, Michel Gendreau, Christelle Guéret, and Andrés L. Medaglia. A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 2013. ISSN 03772217. doi: 10.1016/j.ejor.2012.08.015.

[69] Porsche Consulting. The Future of Vertical Mobility - Sizing the market for passenger, inspection, and goods services until 2035. Technical report, 2018. URL https://www.porsche-consulting.com/en/press/insights/detail/study-the-future-of-vertical-mobility/.

[70] Priyank Pradeep and Peng Wei. Energy efficient arrival with RTA constraint for urban eVTOL operations. *AIAA Aerospace Sciences Meeting, 2018*, (210059):1–13, 2018. doi: 10.2514/6.2018-2008.

[71] Harilaos N. Psaraftis. Dynamic Programming Solution To the Single Vehicle Many-To-Many Immediate Request Dial-a-Ride Problem. *Transportation Science*, 14(2):130–154, 1980. ISSN 00411655. doi: 10.1287/trsc.14.2.130.

[72] Roland Berger. Urban Air Mobility; The rise of a new mode of transporation. Technical report, 2018. URL https://www.rolandberger.com/publications/publication_pdf/Roland_Berger_Urban_Air_Mobility.pdf.

[73] Stefan Ropke, Jean-Francois Cordeau, and Gilbert Laporte. Models and Branch-and-Cut Algorithms for Pickup and Delivery Problems with Time Windows Stefan. *Networks*, 49(4):258–272, 2007. ISSN 1097-0037. doi: 10.1002/net.

[74] J.L. Saha. An Algorithm for Bus Scheduling Problems. *Operational Research Quarterly (1970-1977)*, 21 (4):463–474, 1970.

[75] M. W. P. Savelsbergh and M. Sol. The General Pickup and Delivery Problem. *Transportation Science*, 29 (1):17–29, 1995. ISSN 0041-1655. doi: 10.1287/trsc.29.1.17.

[76] Michael Schneider, Andreas Stenger, and Dominik Goeke. The Electric Vehicle-Routing Problem with Time Windows and recharging stations. *Transportation Science*, 48(4):500–520, 2014. ISSN 15265447. doi: 10.1287/trsc.2013.0490.

[77] Susan Shaheen and Adam Cohen. Mobility on demand (MOD) and mobility as a service (MaaS): early understanding of shared mobility impacts and public transit partnerships. In *Demand for Emerging Transportation Systems*, pages 37–59. Elsevier, Jan 2020. ISBN 9780128150184. doi: 10.1016/B978-0-12-815018-4.00003-6. URL https://www.sciencedirect.com/science/article/pii/B9780128150184000036.

[78] Syed A M Shihab and Peng Wei. By Schedule or On-Demand ? - A Hybrid Operational Concept for Urban Air Mobility Services. In *AIAA Aviation 2019 Forum*, pages 1–13, Dallas, Texas, 2019. AIAA. doi: 10.2514/6.2019-3522.

[79] Christopher Silva, Wayne Johnson, Kevin R. Antcliff, and Michael D. Patterson. VTOL Urban Air Mobility concept vehicles for technology development. In *2018 Aviation Technology, Integration, and Operations Conference*, Jun 2018. ISBN 9781624105562. doi: 10.2514/6.2018-3847. URL https://ntrs.nasa.gov/search.jsp?R=20180006683.

[80] David P. Thipphavong, Rafael Apaza, Bryan Barmore, Vernol Battiste, Barbara Burian, Quang Dao, Michael Feary, Susie Go, Kenneth H. Goodrich, Jeffrey Homola, Husni R. Idris, Parimal H. Kopardekar, Joel B. Lachter, Natasha A. Neogi, Hok Kwan Ng, Rosa M. Oseguera-Lohr, Michael D. Patterson, and Savita A. Verma. Urban Air Mobility Airspace Integration Concepts and Considerations. In *2018 Aviation Technology, Integration, and Operations Conference*, Reston, Virginia, Jun 2018. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-556-2. doi: 10.2514/6.2018-3676. URL https://arc.aiaa.org/doi/10.2514/6.2018-3676.

[81] Paolo Toth and Daniele Vigo. *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics, Philadelphia, 2000.

[82] Uber. Uber Elevate, 2020. URL www.uber.com/elevate.

[83] Parker D. Vascik and R John Hansman. Constraint Identification in On-Demand Mobility for Aviation through an Exploratory Case Study of Los Angeles. In *17th AIAA Aviation Technology, Integration, and Operations Conference*, Reston, Virginia, Jun 2017. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-508-1. doi: 10.2514/6.2017-3083. URL https://arc.aiaa.org/doi/10.2514/6.2017-3083.

[84] Parker D. Vascik and R. John Hansman. Scaling constraints for Urban Air Mobility operations: Air traffic control, ground infrastructure, and noise. In *Aviation Technology, Integration, and Operations Conference*, Atlanta, GA, 2018. ISBN 9781624105562. doi: 10.2514/6.2018-3849.

[85] Parker D. Vascik and R. John Hansman. Development of Vertiport Capacity Envelopes and Analysis of Their Sensitivity to Topological and Operational Factors. In *AIAA Scitech 2019 Forum*, Reston, Virginia, Jan 2019. American Institute of Aeronautics and Astronautics. ISBN 978-1-62410-578-4. doi: 10.2514/6.2019-0526. URL https://arc.aiaa.org/doi/10.2514/6.2019-0526.

[86] P Verschuren and H Doorewaard. *Designing a Research Project*. Eleven International Publishing, The Hague, scond edition, 2010. ISBN 978-90-5931-572-3.

[87] Marin Vlastelica Pogančić. Towards Data Science, 2019. URL https://towardsdatascience.com.

[88] Volkswagen AG. Porsche and Boeing to Partner on Premium Urban Air Mobility Market, 2019. URL https://www.volkswagenag.com/en/news/2019/10/Porsche_and_Boeing_to_Partner.html{#}.

[89] Volocopter. Volocopter - Home, 2020. URL https://volocopter.com/en/.

[90] Volocopter GmbH. Volocopter Air Taxi Flies Over Singapore's Marina Bay, 2019. URL https://www.prnewswire.com/news-releases/volocopter-air-taxi-flies-over-singapores-marina-bay-300942843.html.

[91] Pavan Yedavalli, Jessie Mooberry, and Airbus. An Assessment of Public Perception of Urban Air Mobility (UAM). Technical report, 2019. URL https://drive.google.com/file/d/15YTiXudG_eb5IIkqqCMcE2SGvDBJA7_L/view{%}0Ahttps://www.utmblueprint.com/.

III

Supporting work

A

Model Framework & Flowchart

The programmed model was composed of various submodels that merged into the main scheduling model. In addition to the methodology section in the paper in part I, this section will further elaborate on the structure of the model and modules, and provide a flowchart to illustrate the framework.

Figure A.1 illustrates the structure of the programming framework and the relations between the various submodels and the main scheduling model, the input and output, and the intermediate processes. All input parameters are specified in a single *Excel* file per scenario. This makes it straightforward to adjust the formulation if the values change. The first step is to load all parameters from *Excel* to *Python* in the submodels. Starting with the vehicle model, which loads all relevant data of the eVTOL vehicle used in the case study. This submodel also includes parameters for the battery.
The demand model provides the probability density function from the three combined normal distributions that signify the three demand peaks during the day. This method is based on Kohlman and Patterson (2018). This function is for a 24-hour period. In this submodel, based on the user input data, the operational time horizon is also set and the probability function is adjusted accordingly. Finally, in this submodel, a scale factor is applied to the combined probability function. This scale factor ensures that there are 30 customer requests during the operational time horizon. The value is determined experimentally for each scenario.
The network submodel is used for all vertiport related data and computations. The locations of the vertiports, including the depot, are provided using coordinates. In the submodel, the distance between all vertiports is computed. Using the vehicle data, also the travel time and energy required are calculated between all OD-pairs (Origin-Destination). Additionally, various subsets of allowable vertiport OD-pairs are determined based on the scenario. If the scenario includes a directed flow, for example, a subset of OD-pairs contains double the number of combinations that include the focus vertiport.
The request submodel combines parameters from the vehicle, demand and network submodels to generate the trip requests. For each time interval in the operational time horizon, a random value is compared with the demand probability distribution function value for that time instance. If this triggers a trip, an OD-pair is taken from the allowable trips set from the network model. Depending on the scenario, there may be multiple sets of allowable trips during different periods of the day. For instance, in the balanced scenario, which has different flow directions in the morning and in the afternoon. For the clustered scenario, the set of allowable trips is dynamically adjusted. If a trip request is triggered, the same OD-pair will be added to the allowable trips for the next instance such that there is a higher probability to have the same OD-pair again. Afterwards, the double value is removed again. The amount of passengers is determined by taking a random integer between 1 and the maximum passenger capacity value from the vehicle model. The sampling of the passenger demand is based on a standard function without any adjustments. This approach yields 30 trip requests with location, time and passenger demand information.
In the scheduling model, all input parameters, intermediate calculated values and the set of trip requests are used to formulate the MILP model. First of all, the necessary sets and coefficients are computed. Using *Gurobi*, the objective function, decision variables and constraints are formulated and saved in a *.lp* file. The format is adjusted in order to be compatible with *CPLEX*.
The *.lp* file is then solved using *CPLEX* on a server. The server also enabled to run multiple files concurrently. In some cases, the *.lp* file is partly solved before the server-run, or continued after the server-run, on a per-

sonal computer. In these cases, the solution *.sol* file generated is used as a warm start for *CPLEX*. Next to using the previous solution as warm start, other solver settings used are the *MIPemphasis* (set to focus on proving optimality) and the *IncludeThreads* setting to make sure the solver can use the additional processing power of the server. A time limit, and in some cases a gap difference, were specified to influence when the optimisation process was to be terminated.
The optimisation run results are parsed to a text file with the decision variables and their value, the objective function value, the optimality gap and the upper bound, if applicable. Only the total profit is readily available from the raw results. Other, aggregated, results must be computed from the raw data. In the post-processing phase, the values are imported back into *Python* variables. Subsequently, the vehicle routes are constructed and the relevant KPI's are computed. Finally, a verification script is run: the integrity of the routings is checked and, for various decision variable values, it is monitored that they conform to the MILP constraints. For each scenario, the KPI's are combined in an *Excel* file for comparison and further analysis.

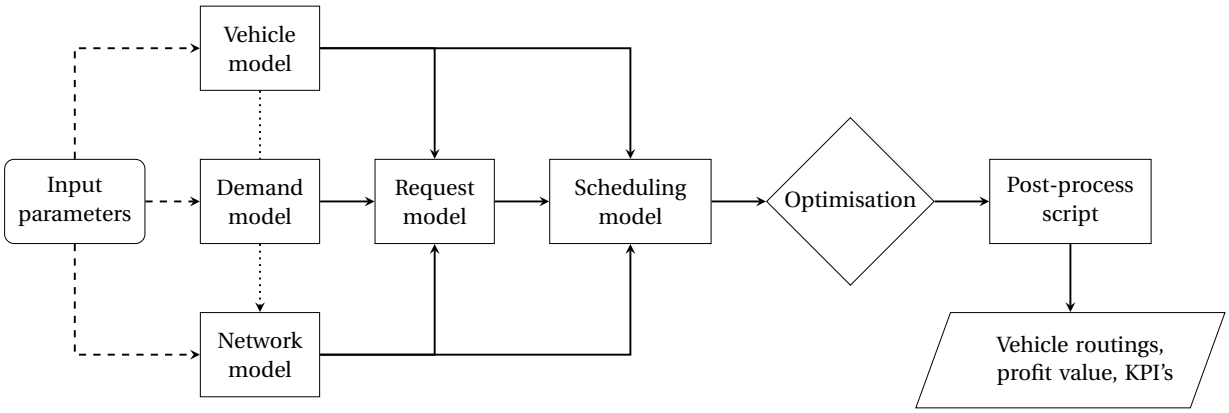

Figure A.1: Flowchart showing the model's structure with the input dataset (rounded rectangle), different *Python* submodels (rectangle), process (diamond) and output (trapezoid).

B

Scenario Trip Request Distributions

For each of the five main scenarios, the distribution of the trip requests is presented in fig. B.1. Each arrow represents a single customer request from the origin to the destination vertiport, based on the earliest possible departure and arrival times.

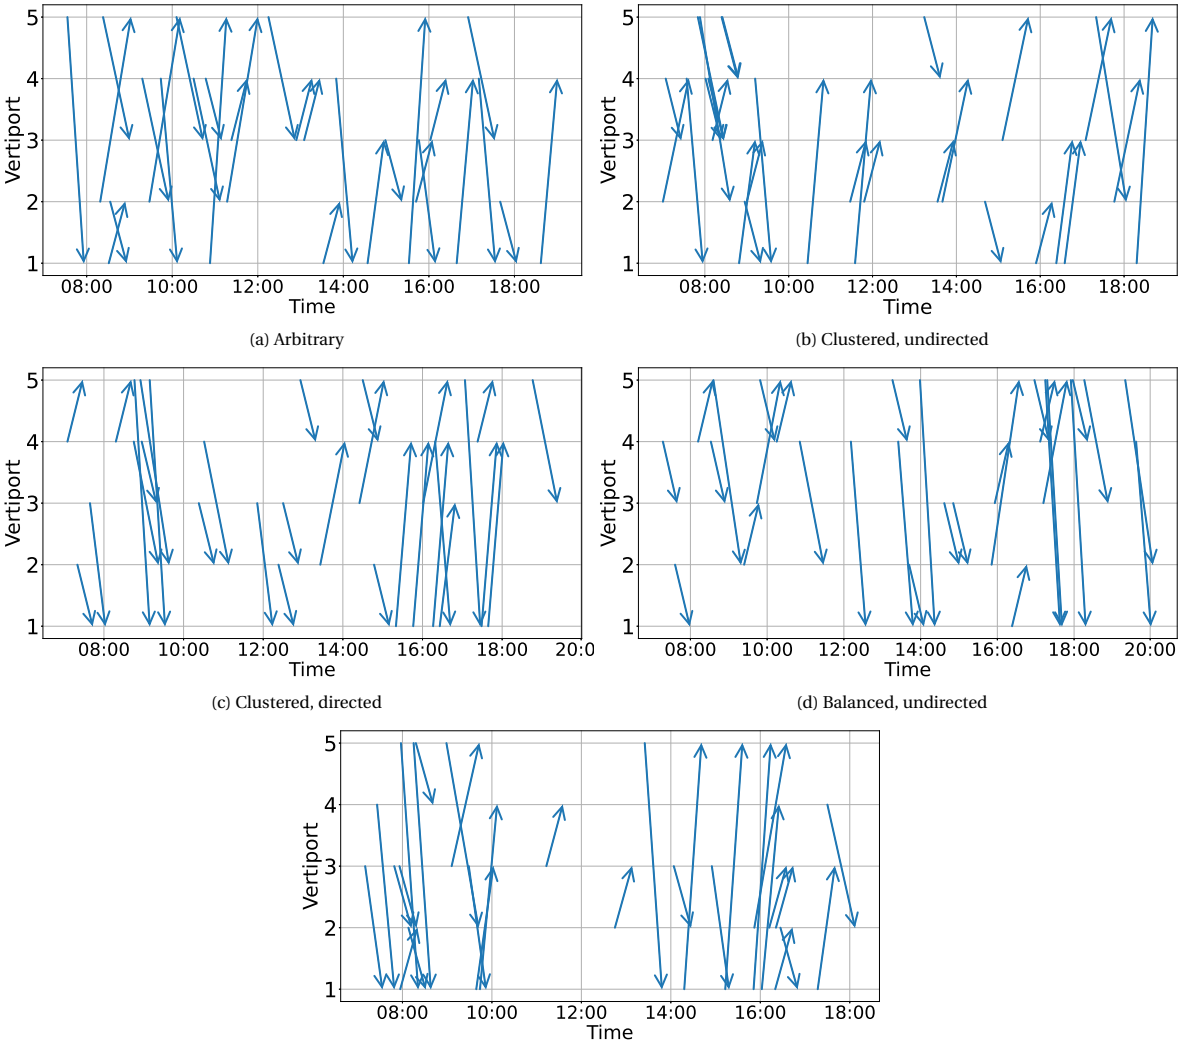

(a) Arbitrary

(b) Clustered, undirected

(c) Clustered, directed

(d) Balanced, undirected

(e) Balanced, directed

Figure B.1: Graphs showing the trip request distribution for each main scenario. Based on the earliest departure time of the customer's request.

C

Model Verification

During the development of the optimisation model, verification was a continuous effort to ensure that the translation from the mathematical to the programming formulation was correct, and that the optimisation model worked as intended. This section will provide a summary of the verification efforts and methods used.

In the verification process, it was checked that the programmed model did not contain any errors. This includes 'technical' errors, such as syntax mistakes or erroneous commands used, as well as more 'fundamental' errors where the programmed formulation does not yield the desired output because it not correctly matches the original mathematical formulation. In order to check the former type, unit testing was primarily used within each of the submodels, as presented in chapter A. First of all, it was checked that each input parameter imported from a central data file had the correct format and unit to ensure that there would be no discrepancy in the value or order of magnitude. Subsequently, each programmed computation step was compared to a manual calculation to ensure that equations were correctly implemented and that the correct input parameters were used. This involved, for example, the computation of the trip distances based on the vertiport coordinates, the travel time between nodes and the energy required. After checking that each submodel was correct internally, special attention was given to the integration of the submodels to ensure that the proper values carry over. This check consisted of feeding a known set of input parameters and comparing the model's final outcome with a manually computed value.

Once the entire formulation was completed, various integral customer demand scenarios were run from start to finish, including the optimisation. These scenarios were composed manually and constituted a reduced problem size (in terms of requests, vehicles and vertiports) such that the correct end result could be determined by hand. The optimisation model results were compared with the manual outcomes to verify that they match and that the intermediate MILP-model is properly programmed. Multiple edge cases were drafted and investigated to ensure that the formulation could correctly handle a wide range of input scenarios. Special attention was drawn to the vertiport capacity constraints: different vertiport-visit overlap situations were applied to confirm that the constraints handle these accurately.

Based on flaws that came to light during the development phase, a verification script was also written in *Python*. This script checks the final outcome, the decision variable values, and the aggregated results of the full-scale optimisation runs for discrepancies to ensure that previously addressed errors do not suddenly occur with an increased problem size. Manual inspection of the final outcomes becomes difficult, if not impossible, with many vehicles and trips. Consequently, the route integrity is checked by the script as well, i.e. is there a proper continuity and precedence of the spatial and temporal parameters between each stage in the route. Graphically, in the route log in chapter D, the departure value in each line should be the arrival value of the previous line. Furthermore, the entry and exit times of vertiport trips and the resulting matches are verified. Also, it is double-checked for multiple variables that these do not violate the constraints. For instance the energy level, vehicle load and time windows. Finally, the different components of the objective function are computed separately based on the decision variables values. The combined value is compared with the profit from the solution file to ensure there is no difference.

D

Example Routing

Figure D.1 shows a time-space plot of vehicle routings from the depot, between customer vertiports and back to the depot for a full day's operation with the balanced directed demand distribution and a time window of 5 minutes. The line styles indicate the different stages a vehicle can be in. For clarity, only the routings of three out of the six vehicles are shown.

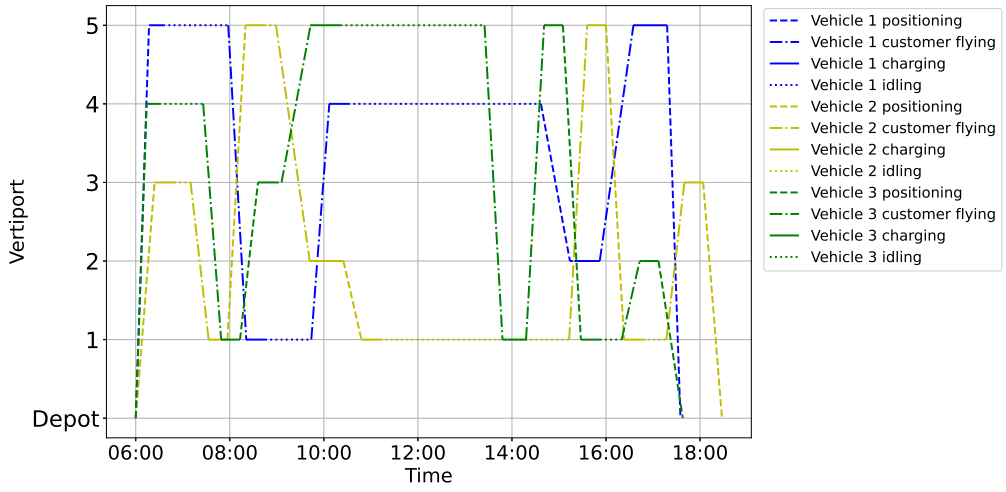

Figure D.1: Time-space plot with the routing of three vehicles for the balanced directed scenario with time window of 5 minutes.

The raw data route log of vehicle 2 from fig. D.1 is provided below. 'Dep.' and 'Arr.' differentiate departure and arrival, respectively. 'Vert.' denotes the vertiport. For the stages, 'PO', 'CH', 'CU' and 'ID' represent positioning (flying or within the same vertiport to/from charging point), charging, customer flying and idling, respectively. The timings provided are based on the interval of 1 minute, from 06:00h onwards. Note that the energy level represents the energy remaining at the (departure) node from the first column at the starting time of the stage. Because the idle stage takes place at the same node as where the charging starts in the preceding stage, the energy level for the idle stage still states the depleted energy level even though, in practice, it has already recharged to full. The same holds for empty positioning flights to another vertiport.

[*Dep. node*	*Arr. node*	*Dep. vert.*	*Arr. vert.*	*Dep. time*	*Arr. time*	*Load*	*Energy [kJ]*	*Stage*]
[[0	61	0	3	0	24	0	396000	$'PO'$]
[61	61	3	3	24	49	0	175804	$'CH'$]
[61	1	3	3	49	70	0	175804	$'ID'$]
[1	31	3	1	70	93	2	396000	$'CU'$]
[31	91	1	1	93	93	0	182721	$'PO'$]
[91	91	1	1	93	117	0	182721	$'CH'$]
[91	70	1	5	117	140	0	182721	$'PO'$]
[70	70	5	5	140	164	0	184883	$'CH'$]
[70	10	5	5	164	179	0	184883	$'ID'$]
[10	40	5	2	179	222	3	396000	$'CU'$]
[40	100	2	2	222	222	0	17094	$'PO'$]
[100	100	2	2	222	265	0	17094	$'CH'$]
[100	81	2	1	265	288	0	17094	$'PO'$]
[81	81	1	1	288	312	0	181583	$'CH'$]
[81	21	1	1	312	553	0	181583	$'ID'$]
[21	51	1	5	553	576	3	396000	$'CU'$]
[51	111	5	5	576	576	0	184883	$'PO'$]
[111	111	5	5	576	600	0	184883	$'CH'$]
[111	89	5	1	600	623	0	184883	$'PO'$]
[89	89	1	1	623	647	0	184883	$'CH'$]
[89	29	1	1	647	677	0	184883	$'ID'$]
[29	59	1	3	677	700	4	396000	$'CU'$]
[59	119	3	3	700	700	0	182721	$'PO'$]
[119	119	3	3	700	724	0	182721	$'CH'$]
[119	121	3	0	724	748	0	182721	$'PO'$]]