Final Report for

Aviation Global Demand Forecast Model Development and ISAAC Studies

Task 3.7: Concept of Operations for ODM VTOL Aircraft Package Delivery

Brian J. German (PI) and Matthew J. Daskilewicz
Georgia Institute of Technology

July 2018

Contents

1. Introduction

This report documents the research conducted for "Task 3.7: Concept of Operations for ODM VTOL Aircraft Package Delivery" of NIA Task Order #6563, "Aviation Global Demand Forecast Model Development and ISAAC Studies" for NASA contract number NNL13AA08B. The research project was conducted at Georgia Tech with Brian German as the principal investigator. Ty Marien is the NASA technical representative (TR), and Jeremy Smith is the associate TR.

The overall goal of the research was to study the potential of using electric vertical takeoff and landing (VTOL) passenger aircraft in the context of urban air mobility (UAM) to carry cargo within cities during times of low passenger demand. The aircraft being considered are envisioned to be capable of autonomous operations and of carrying four passengers or approximately 800-900 lbs of cargo. Compared to small UAS, this size class of aircraft could potentially enable new modalities of cargo delivery for commercial and industrial supply chain applications or new possibilities for package delivery upstream of the last-mile distribution to the consumer. The research focused on a case study of cargo delivery in a 17-county region of northern California including the San Francisco Bay Area.

The project was conducted in collaboration with Anthony Trani and his research team at Virginia Tech, Sam Dollyhigh at Analytical Mechanics Associates, and our technical monitors at NASA, Jerry Smith and Ty Marien. Georgia Tech focused primarily on investigating concepts of operations (CONOPS) in terms of (1) cargo business/market concepts, (2) vertiport placement throughout the city (in a coarse-grained sense based on promising census tracts for eVTOL service), (3) modeling of aircraft performance and energy usage, (4) time–speed–distance tradeoffs and their impact on cargo delivery operations, and (5) a Monte Carlo simulation of a UAM cargo network. Virginia Tech focused on (1) assessing eVTOL cargo demand, (2) vertiport placement (in terms of more detailed site studies at particular types of locations, e.g. department store parking lots), (3) economics modeling, (4) carbon dioxide emissions modeling, and (5) comparison to other transportation modes such as delivery trucks. Mr. Dollyhigh conducted a review of public domain information about cargo delivery business models of online retailers and carried out analyses supporting the team's work in demand estimation.

The report is organized as follows. First, we describe several guiding principles that we followed in our investigation of potentially viable eVTOL cargo delivery business models and several archetypes for operational concepts. Next, we describe the first CONOPS we investigated in detail: eVTOL package delivery from a remote Amazon distribution center to dedicated cargo vertiports throughout the San Francisco area. Finally, we describe a CONOPS based on delivery of packages from Amazon distribution centers to Amazon Lockers co-located with vertiports for a passenger UAM network.

2. Operational Concepts for UAM Package Delivery

2.1 The Value Proposition of eVTOL Cargo Delivery

As we considered potential concepts of operations that would result in viable business models, we approached the problem from the viewpoint that the most suitable eVTOL cargo applications are those for which the end user ascribes a high value to the timeliness or immediacy of delivery. Examples of products that are potentially applicable in this context include:

- Prepared fresh/hot food, e.g. restaurant delivery
- Fresh produce
- Time-critical items needed in the workplace, e.g. replacing failed production line equipment parts

- Items needed to support leisure activities in progress, e.g. replacing failed consumer electronics such as cameras on a vacation

Additionally, we believe that eVTOL cargo delivery, like sUAS drone delivery, has significant potential for motivating latent demand of retail products, i.e. impulse buying, in part because of the novelty of the delivery modality (in the near term) and in part because the timeliness of delivery reduces the waiting threshold and increases convenience.

However, cargo delivery with passenger-class eVTOL aircraft is not small UAS "drone delivery"; any viable business model must aggregate sufficient cargo volume or value to justify a flight in a rather large aircraft in terms of operating costs. Additionally, being larger than sUAS aircraft, eVTOL aircraft cannot land directly at a consumer's doorstep. The time required for aggregation and "the first/last mile" from a vertiport must not increase the overall door-to-door delivery time to a level that negates the advantage of eVTOL aircraft flight speed relative to shipment with other transportation modes such as trucks. The value proposition of eVTOL cargo delivery improves as ground traffic congestion worsens.

The cost of eVTOL cargo delivery could be allowed to be somewhat higher than that of trucks or other delivery modes, to the extent that the customer is willing to pay extra for the timeliness of delivery and to the extent to which costs can be offset by additional latent demand/revenue generated by this value of immediacy. Nonetheless, to create a vibrant market, eVTOL operating costs must be sufficiently low, and we believe that because similar cargo delivery services have not been previously adopted with current-generation helicopters or fixed wing aircraft, there is a need for significant operating cost reductions via technologies such as electric aircraft propulsion and autonomous operation. Additionally, rapid delivery requires high flight dispatch frequencies and reasonably high flight speeds.

In the context of these guiding principles, we conceptualized four promising categories of eVTOL UAM cargo services. These categories are described briefly below.

2.2 Intra-City Point-to-Point Courier

Courier services of small packages in metropolitan areas are promising applications for eVTOL delivery. We first consider a point-to-point model with direct eVTOL cargo flights operating between pairs of vertiports in the network. Considering a network with n vertiports, there are $n(n-1)/2$ unique direct flight routings. This $\mathcal{O}(n^2)$ scaling implies that in most cases it would be impractical to offer a cargo service between all origin–destination (O–D) vertiport pairs. Service would likely be offered only for O–D pairs with sufficient demand and sufficient time savings relative to ground transport. Additionally, because eVTOL aircraft have rather large payload capability compared to typical package sizes, the service would likely not be on-demand but rather, scheduled, with flight frequencies of 1-2 hours between select O–D pairs. First- and last-mile pickup/delivery would either not be provided (in the case of customer dropoff/pickup models) or by courier van, sUAS, rideshare delivery, or other relevant mode.

The primary advantage of a point-to-point model is that it could potentially provide the minimum door-to-door delivery time possible with an eVTOL aircraft because the flights are direct. However, a disadvantage is that is difficult to serve all routes economically, i.e. many of the point-to-point routes will have low demand.

We envision that vertiport locations near existing package delivery retail facilities, e.g. FedEx and UPS stores, are promising for this business model, presuming that these facilities are optimally geographically distributed to capture consumer package delivery demand. This concept of operations is illustrated in Figure 2.1 by considering all O–D pairs defined by the UPS store locations in the San Francisco area.

2.3 Intra-City Hub-and-Spoke Courier

A hub-and-spoke courier model connecting n vertiports requires only $n-1$ flight routes to connect all O–D pairs. All flights from each "spoke" vertiport connect through a "hub" vertiport. The linear scaling has the advantage of much more economically serving the demand in terms of number of aircraft and flights required because *all* of the cargo from a given origin vertiport destined to all other vertiports in a particular time window can be carried in a *single flight* (subject to payload capacity limits, of course).

Figure 2.1: Intra-city point-to-point courier. San Francisco UPS Store locations, showing all routes between stores.

A disadvantage of a hub-and-spoke model, however, is that the hub vertiport must now be used as a package sorting location to transfer packages from arriving aircraft to the an aircraft deparating to the destination vertiport. This sorting operation incurs not only labor costs but also a time delay to process the packages that reduces the timeliness of delivery.

In this model, packages would be aggregated at each vertiport and a flight would be dispatched to the hub when an aircraft can be fully filled and/or at the end of a specified time window. We envision this model to be a scheduled service, with flights every 30 minutes to 1 hour (likely more frequent than the point-to-point model because aircraft can be filled more quickly at the spoke vertiports). First- and last-mile pickup/delivery would either not be provided (in the case of customer dropoff/pickup models) or by courier van, sUAS, rideshare delivery, or other relevant mode.

This concept of operations is illustrated in Figure 2.2 by considering all O–D pairs defined by the UPS store locations in the San Francisco area.

2.4 Regional Hub-and-Spoke Courier

We envision a regional hub-and-spoke courier model that is in most respects similar to the intra-city hub-and-spoke model, except that one or several origins or destinations are located at longer distances from most spokes in the network. We consider this model to be distinct from the intra-city model because the long distance to some vertiports will likely have a significant bearing on fleet allocation challenges and on aircraft energy requirements. This concept of operations is illustrated in Figure 2.3.

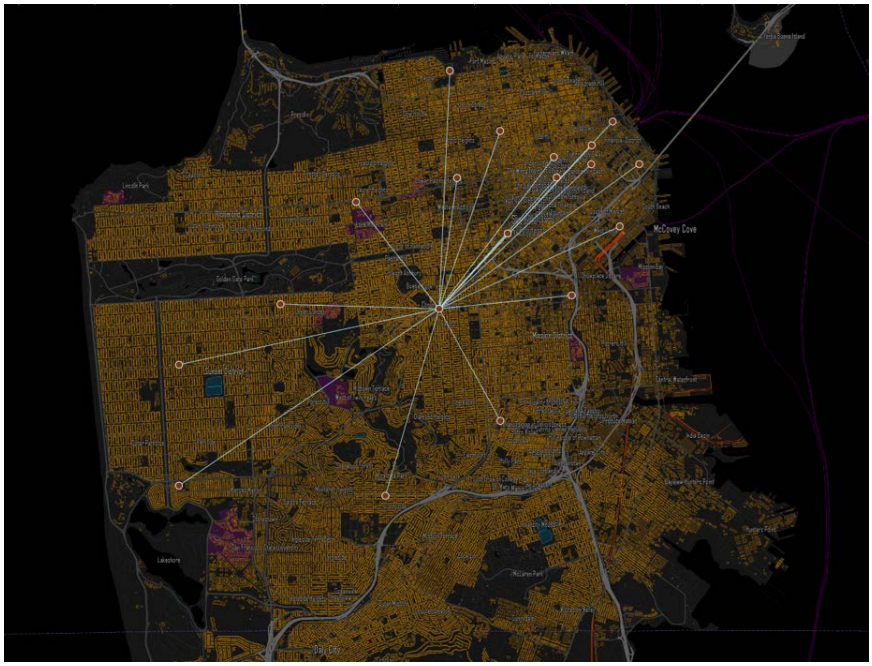

Figure 2.2: Intra-city hub-and-spoke courier. San Francisco UPS Store locations showing routes between stores and a presumed central hub.

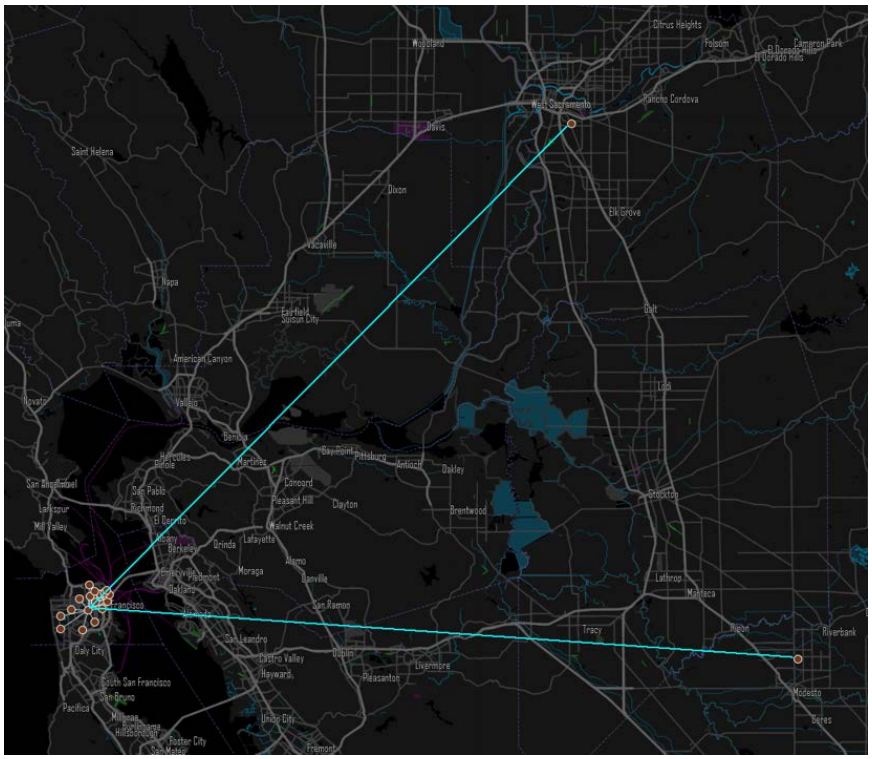

Figure 2.3: Regional hub-and-spoke courier. Routes between San Francisco hub and regional delivery nodes in Sacramento and Modesto.

2.5 City-to/from-Airport Courier

Another potentially promising eVTOL cargo concept of operations is to deliver overnight freight packages from vertiports at FedEx/UPS stores or similar locations to a package shipment facility at a nearby airport (such as the UPS facility at Oakland International Airport). The advantage of this concept is that the higher speed of eVTOL during the evening rush hour could allow a later a last drop-off time for domestic or global overnight packages by avoiding traffic and reducing travel time compared to vans/trucks. The last drop-off time is dependent on both time zone and the flight time required for express packages to reach a national sorting hub, e.g. Memphis, TN for FedEx or Louisville, KY for UPS. For example, most locations in the San Francisco area have a last drop-off time of 5:30 PM local time, whereas east coast locations typically have 7:00 PM last drop-offs. This concept of operations would likely offer considerable value to customers by allowing later drop-offs and would also likely prove profitable for express shippers. This concept of operations is illustrated in Figure 2.4 for delivery from vertiports sited near UPS Store locations to Oakland International Airport.

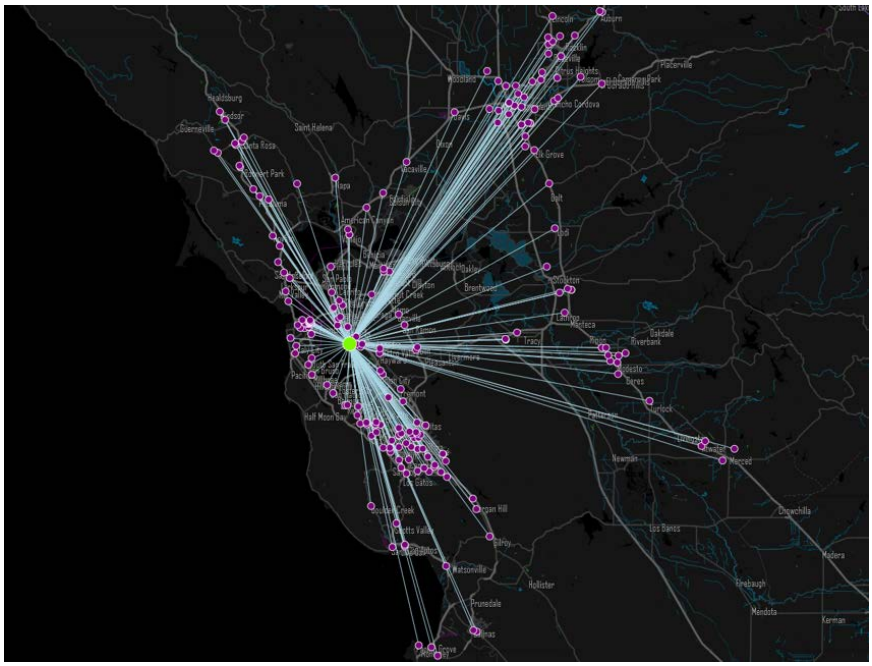

Figure 2.4: City-to/from-airport courier. Routes between UPS Stores and OAK airport UPS hub.

2.6 Selection of CONOPS for Further Study

Based on our initial study, we felt that package delivery as a business itself, such as the courier services discussed above, has considerable potential for a viable eVTOL cargo market. However, we felt that courier models and similar business concepts may not offer as much value as the application of eVTOL cargo delivery as a part of a vertically-integrated retail sales operation. Motivated by the drone delivery model, we wondered whether eVTOL cargo delivery could offer more significant value by delivering packages for online retailers such as Amazon or brick-and-mortar retailers such as Target or Walmart.

We therefore considered two CONOPS in more detail, both based on leveraging eVTOL for delivering packages from Amazon fulfillment centers to consumers. First, we examined shipping packages from a fulfillment center to dedicated cargo vertiports located in the San Francisco area. Next, we leveraged and extended the modeling and simulation tools developed in this first (and simpler) analysis to explore the more complicated CONOPS of delivering packages from fulfillment centers to Amazon lockers situated at

passenger vertiports by using idle aircraft capacity from the passenger eVTOL network. Our analyses of these two CONOPS are presented in the following sections.

3. Package Delivery from an Amazon Distribution Center to Dedicated Cargo Vertiports

The first concept of operations we modeled in detail is to deliver packages rapidly from a distribution center situated in outlying suburban or rural area to several vertiport package distribution hubs located within a metro area. Last mile delivery from a vertiport to individuals would be carried out by other means such as van, small UAS, or customer pickup. Our motivation for this operational concept is that, by providing high speed, high payload, and low cost flights from remote distribution centers to urban areas, eVTOL UAM aircraft could enable faster delivery of a wider range of products via same-day or 2-hour delivery services. The discussion of this CONOPS in this section is drawn largely from an AIAA SciTech 2018 paper prepared based on this research project [1].

Examples of the types of services that could potentially benefit from rapid eVTOL cargo delivery include Amazon Prime Same Day and Amazon Prime Now. Amazon Prime Same Day offers guaranteed delivery of certain items in the Amazon inventory by 9:00 PM for orders placed before noon. Pricing for Amazon Prime members is set at $5.99 for orders totaling less than $35 and free for orders of $35 or above. For customers without a Prime membership, pricing is $8.99 per order plus $0.99 per item [2]. Amazon Prime Now is a more rapid delivery service offering Prime members free 2-hour delivery for selected items and 1-hour delivery for $7.99. Both services are currently available in 30 or more U.S. cities.

We have two motivations for a UAM electric VTOL aircraft application based on these Amazon services and similar rapid delivery services. First, relatively few products available from Amazon are available via Prime Now. Recent reports indicate that Prime Now offers approximately 1 million items, whereas approximately 400 million unique items are available for purchase from Amazon [3, 4]. Prime Now offers time-critical items by sourcing from local retail stores such as supermarkets that are attracted to the appeal of being featured in the Amazon Prime Now app and, likely, from small Amazon-owned urban warehouse facilities stocked with high-demand items. We hypothesize that, in many cases, the narrow 1- to 2-hour delivery window for Prime Now and similar services currently precludes the ability to ship the greater variety of items from larger outlying fulfillment centers to serve demand in urban areas. Fulfillment centers are often located in suburban and rural areas where real estate and labor costs are less than in the urban core.

However, UAM passenger eVTOL delivery could be sufficiently rapid to allow orders to be filled from remote distribution centers and still meet a 2-hour order-to-door delivery threshold. The time savings would be achieved by higher aircraft speeds compared to truck delivery and by overflying traffic congestion. Additionally, passenger class eVTOL aircraft are expected to have a greater range than small UAS, in addition to a considerably larger payload capability. By enabling access to remote large fulfillment centers, eVTOL aircraft could potentially allow services similar to Prime Now to offer a much greater variety of products.

Our second motivation is that UAM electric VTOL aircraft could enable a reduction in the delivery time window for services such as Amazon Same Day delivery. For example, the latest order time for guaranteed delivery by 9:00 PM could perhaps be extended from noon to mid-afternoon and/or orders placed by noon could be guaranteed to be delivered by earlier in the evening.

We examine this concept of operations for UAM eVTOL aircraft by considering the San Francisco Bay Area as an example. The Bay Area is interesting because (1) it has a tech-savvy population who are early adopters of new tech-enabled business models, (2) the high-income population is predominantly concentrated near the city center, (3) it has significant geographic constraints and traffic that complicate ground transport, and (4) the Bay Area Amazon fulfillment centers are located approximately 50 mi east of the city in Tracy, CA. Travel time by road from a fulfillment center in Tracy can be more than an hour to downtown San Francisco during high traffic times of day.

3.1 Optimizing Vertiport Placement

We consider a case study of delivering small packages from a warehouse located in Tracy, CA to customers in the San Francisco Bay Area. The package transit is considered to be via two legs: an eVTOL trip from the warehouse to a vertiport located near the consumer, followed by last-mile delivery by car, small UAS, or other mode from the vertiport to the consumer. As illustrated in Figure 3.1, this concept of operations offers the potential for reducing delivery time to a consumer by as much as 40 minutes to 1.5 hours, depending on traffic congestion.

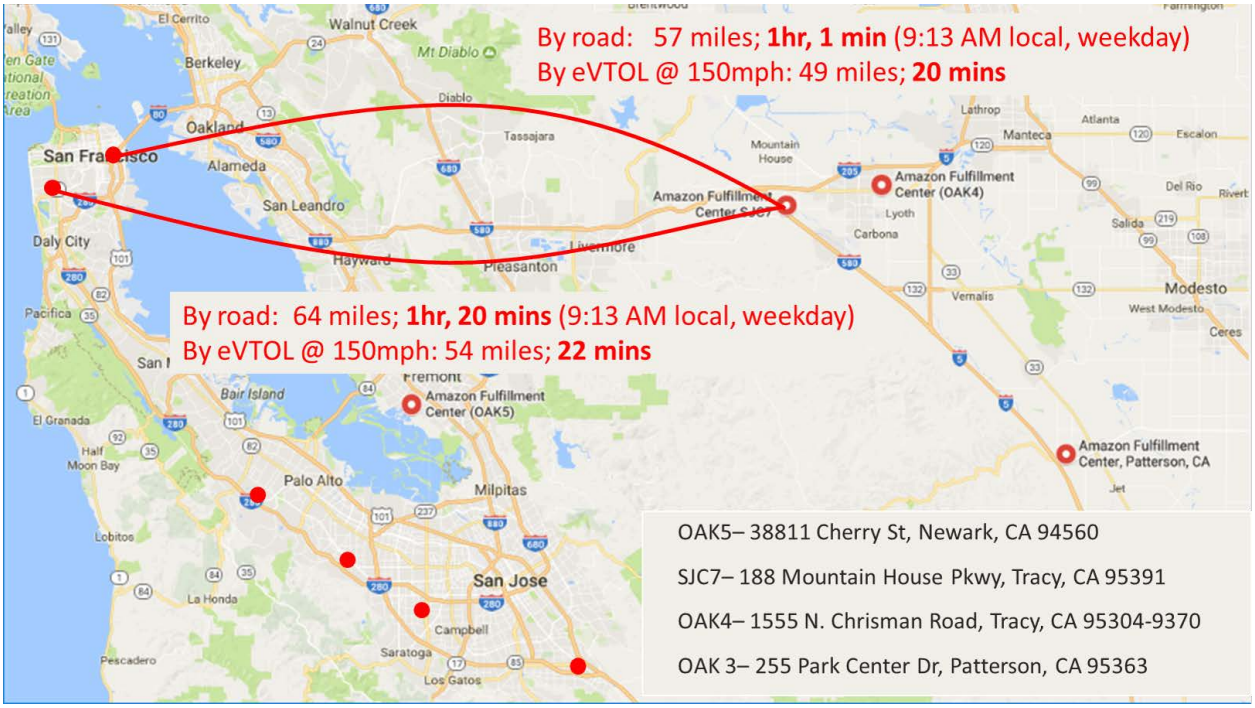

Figure 3.1: Concept of eVTOL delivery of packages from Tracy, CA Amazon fulfillment center to dedicated cargo vertiport locations.

The cost of building a vertiport is expected to be significant, and the number of suitable building locations is expected to be limited. Both of these factors encourage the building of relatively few vertiports. However, because the primary incentive for delivering cargo via VTOL aircraft is decreased delivery time, it is also desirable to have vertiports located close to the consumer in order to reduce the time required for the relatively slow last-mile delivery leg. Considering this tradeoff, we have therefore formulated and solve an optimization problem to select cargo vertiport locations. Our optimization approach is formulated to maximize the package demand served, subject to limits on the number of vertiports.

Accurately estimating demand for same-day package delivery by eVTOL aircraft requires detailed knowledge of customer purchasing habits, and how those habits may change with the introduction of this service. This type of detailed package demand model is beyond the scope of the present study. Instead, we have chosen to maximize a demand surrogate based on population and income level. Rather than try to predict the precise number of packages customers will order, we assume that demand for this service will be highest in places with larger populations (more potential customers) and higher incomes (more customers willing to pay a premium for faster delivery).

To develop this demand surrogate, we first discretized the San Francisco Bay Area into its component census tracts, which are geographic regions defined by the US Census, and which have an average of approximately 4000 residents. In city centers, census tracts may be as small as a few city blocks, whereas in remote areas, they may cover several square miles. For each census tract, we multiply the total population by the

average per capita income to obtain a "total income" measure for the tract. The population and income data is obtained from the American Community Survey 5-year estimate (2011–2015)[5]. The total income measure is then scaled to a range between approximately 0 and 20 for each census tract. This scaled total income measure is our demand surrogate, intended to be representative of the relative demand for eVTOL package delivery in each census tract. The census tract boundaries, colored according to this relative demand metric, are shown in Figure 3.2. To give a sense of how tracts are sized near a city center, the inset in the figure shows a closeup of the San Jose region.

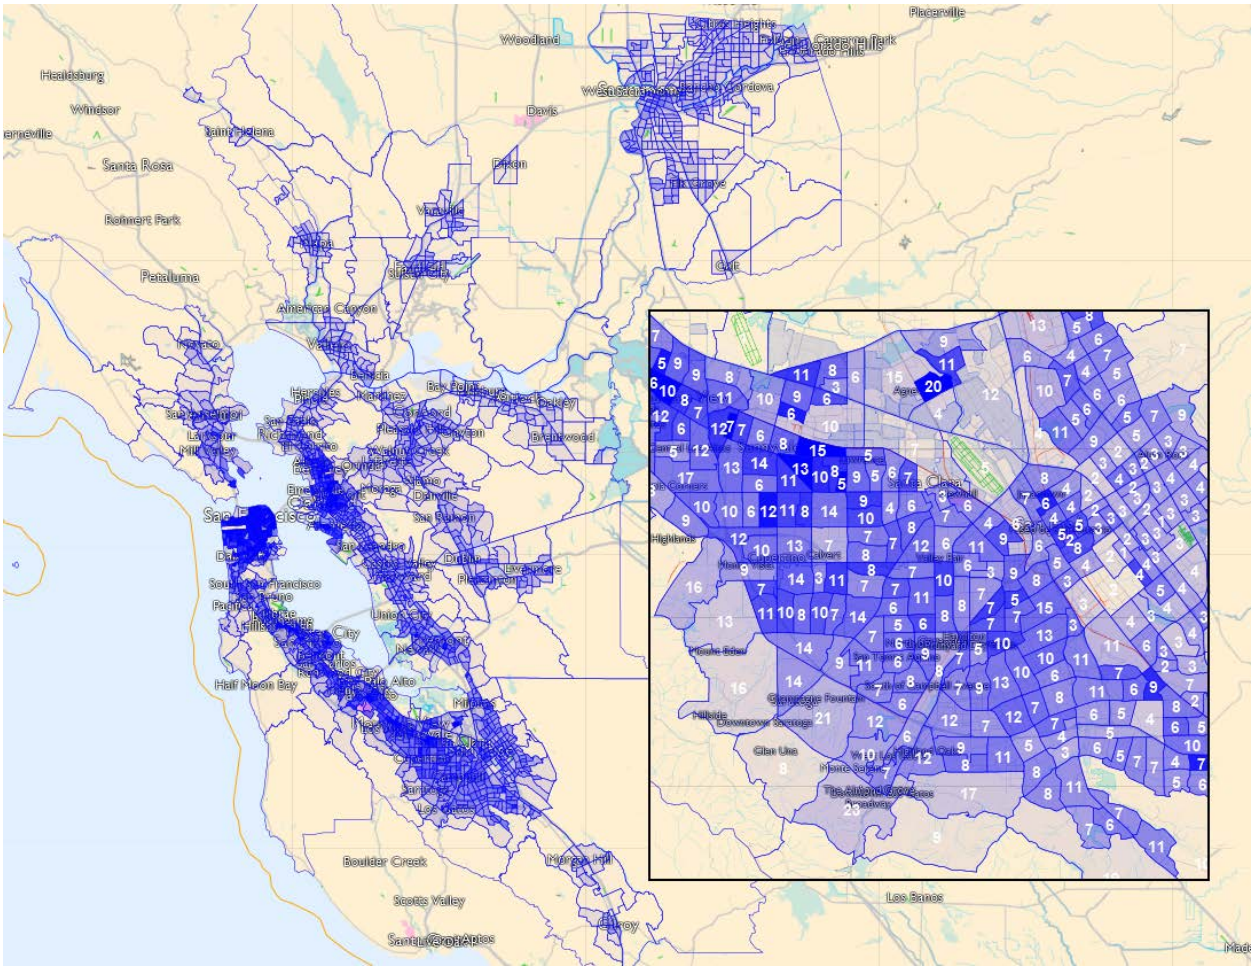

Figure 3.2: Census tract boundaries, colored by demand per unit area. Inset shows San Jose region, with demand values labeled.

Given the small size of the census tracts in the densely populated regions where demand is highest, we can assume without much loss of fidelity that each census tract may hold at most one vertiport. This simplifies the vertiport placement problem to a problem of determining which census tracts should be selected to have a vertiport. Notionally, we envision the vertiports to be placed at the centroids of census tracks. Later studies would be required to refine the specific locations.

Next, we need a model of what demand can be served by each potential vertiport. To develop this model, we assume that the vertiports are able to serve all census tracts within a local neighborhood defined by a driving time threshold from the vertiport to the census tract centroid. Our reasoning is that regardless of whether the last-mile delivery is via customer pickup directly at the vertiport or delivery via a local Uber-like delivery service, the ground-travel time from the vertiport to the end customer is an important limiting factor in determining which customers can be effectively served by each vertiport.

In order to determine each vertiport's neighborhood, the shortest path driving time between each pair of census tracts is needed. Accurate driving times can be obtained from the Google Maps API, but it is desirable to minimize the number of search queries needed. The nine-county metro area contains 1800 census tracts, resulting in over 3.2 million potential pairwise distances that could be computed. However, this number far exceeds the number of roads in the metro area. A more efficient approach is to query driving times along major roadways and use these to build a graph-based model of driving times along particular roadways. A shortest-path algorithm can then be used to traverse this graph between any two points in the metro area.

To construct the transit time graph, we executed a few thousand queries using the Google Maps Directions API. The endpoints of the queries were randomly selected pairs of census tract centroids, distributed in a way that ensures each census tract appeared in multiple queries. Travel times reflect Google's estimated traffic at 6pm on a weeknight – a high-traffic time of day. The travel time along each segment of the trip is added to the roadway graph. The full coverage of the resulting graph is illustrated in Figure 3.3.

An all-pairs shortest path algorithm is then used to quickly compute travel times between all pairs of census tracts. The servable neighborhood around any potential vertiport is then easily computed by filtering the pairs based on a defined driving time threshold.

The San Francisco Bay Area is home to three major airports, and consequently there are significant airspace restrictions in the area. A map of these restrictions is shown in Figure 3.4. For the purpose of this study, we enforce airspace restrictions by forbidding the placement of vertiports in any census tract whose centroid lies inside the portion of the Class B or C airspace that extends to the surface (SFC) around SFO, OAK, and SJC airports. These areas are shaded red in Figure 3.6. We do not consider the impact of airspace restrictions in flight path planning at higher altitudes.

We consider optimization problems of the form "maximize demand served, subject to a limit on how many vertiports can be constructed." In this way we avoid the need for a detailed economics model to balance the cost of constructing vertiports with the revenue they generate. Instead, we optimize multiple scenarios that differ by how many vertiports are constructed and what the driving time threshold is.

Let C denote the set of census tracts, R denote the set of tracts in which vertiports cannot be constructed due to airspace restrictions, N_i denote the set of tracts neighboring tract i, d_i denote the demand of the census tract i, x_i denote the binary variable indicating whether tract i contains a vertiport, y_i the binary variable indicating whether the demand d_i is served, t_{ij} the driving time from tract i to tract j, and v the maximum number of vertiports that may be constructed.

Then we have:

$$j \in N_i \iff t_{ij} < time\_threshold$$

We assume that driving times are symmetric, so that:

$$j \in N_i \iff i \in N_j$$

And we obtain the following optimization problem:

$$\max \quad \sum_{i \in C} d_i y_i \tag{3.1a}$$

$$\text{s.t.} \quad y_i \leq \sum_{j \in N_i} x_j \quad \forall i \in C \tag{3.1b}$$

$$y_i = 0 \quad \forall i \in R \tag{3.1c}$$

$$\sum_{i \in C} x_i \leq v \tag{3.1d}$$

$$x_i, y_i \in \{0, 1\} \quad \forall i \in C \tag{3.1e}$$

The form of this problem ensures that the demand in any census tract can only be satisfied once, regardless of how many vertiports neighbor it.

We consider eight cases, corresponding to building 1 to 8 vertiports. The optimization algorithm solves to completion in each case: the results shown are the true optima. In all cases, a time threshold of 10 minutes is

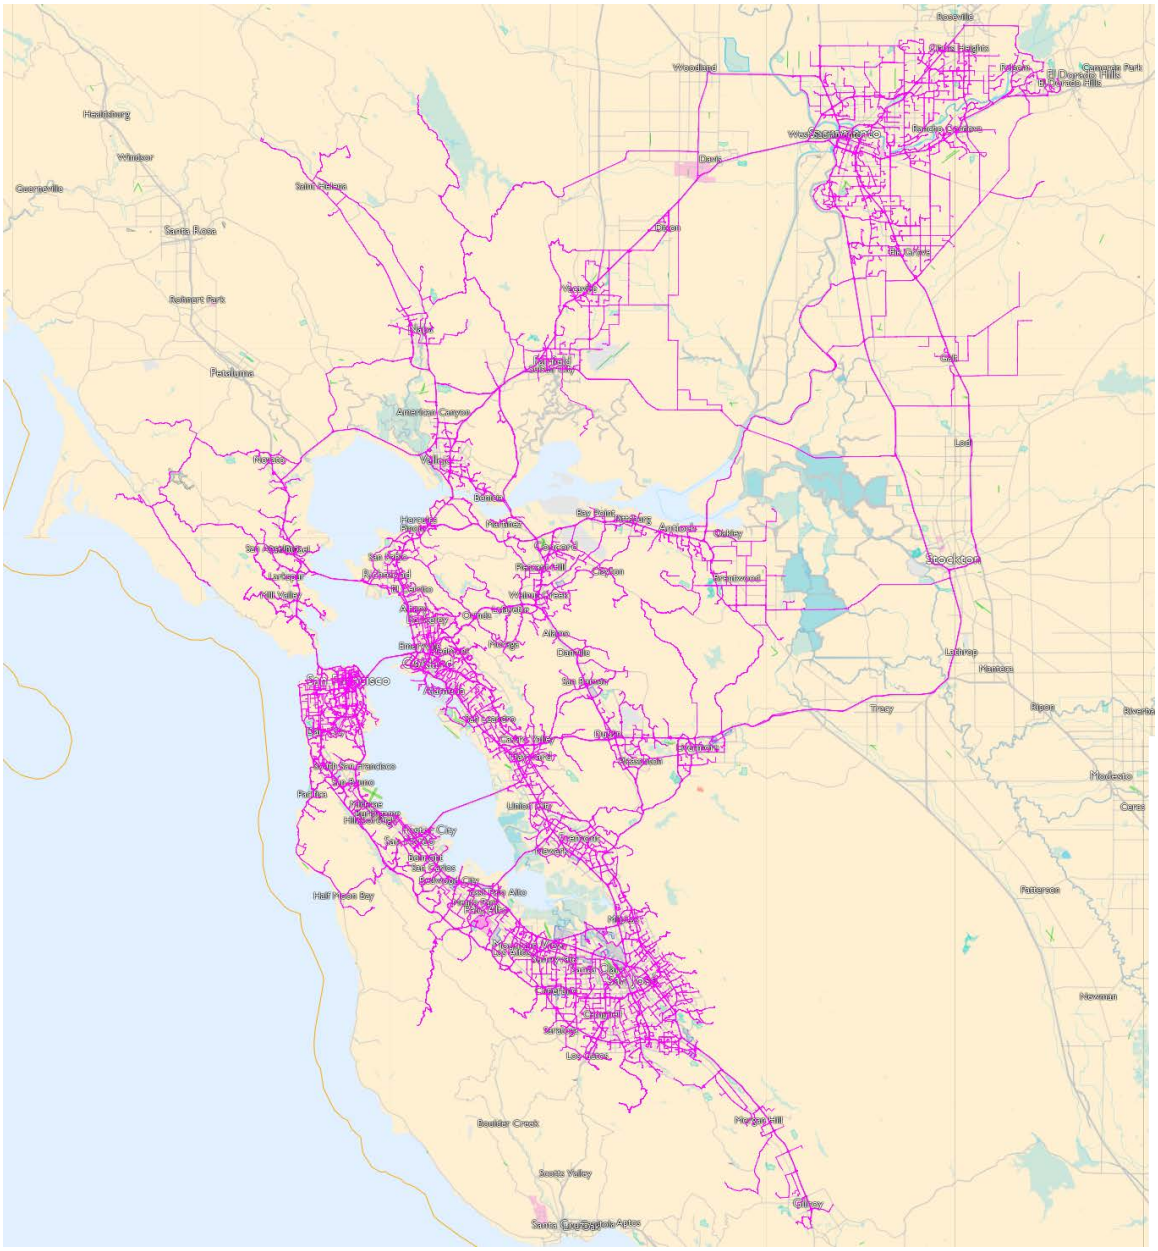

Figure 3.3: Roadway coverage of driving time model.

used for the allowable driving time from the vertiport to the end customer. The optimal vertiport locations for each case are mapped in Figure 3.6, and the incremental demand served by each additional vertiport is plotted in Figure 3.5.

Perhaps unsurprisingly, the first three vertiports are placed near each of the major cities: San Francisco, San Jose, and Oakland. Figure 3.5 shows a precipitous drop in the demand served for the 3rd+ vertiports compared to the first two, and the 5th+ vertiports each serve less than half the demand of the first two. A surprising result is that even with 8 vertiports, large portions of the map remain unserved while three vertiports are clustered near downtown San Francisco. As shown in Figure 1, the area under consideration extends as far north as Sacramento, but the optimizer never places vertiports outside of the area immediately surrounding the bay. The implication is that the population (and income) is too diffuse in the outlying areas to make a vertiport worthwhile with a 10 minute time threshold. Increasing the time threshold will result

Figure 3.4: Airspace restrictions in the San Francisco Bay Area.

in a wider geographical distribution of vertiports, as fewer vertiports are needed to serve the areas directly around the bay.

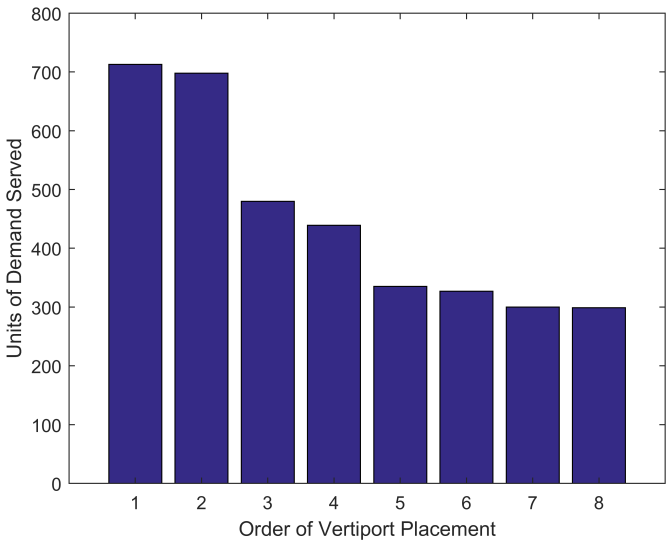

Figure 3.5: Incremental new demand served by placing each subsequent vertiport

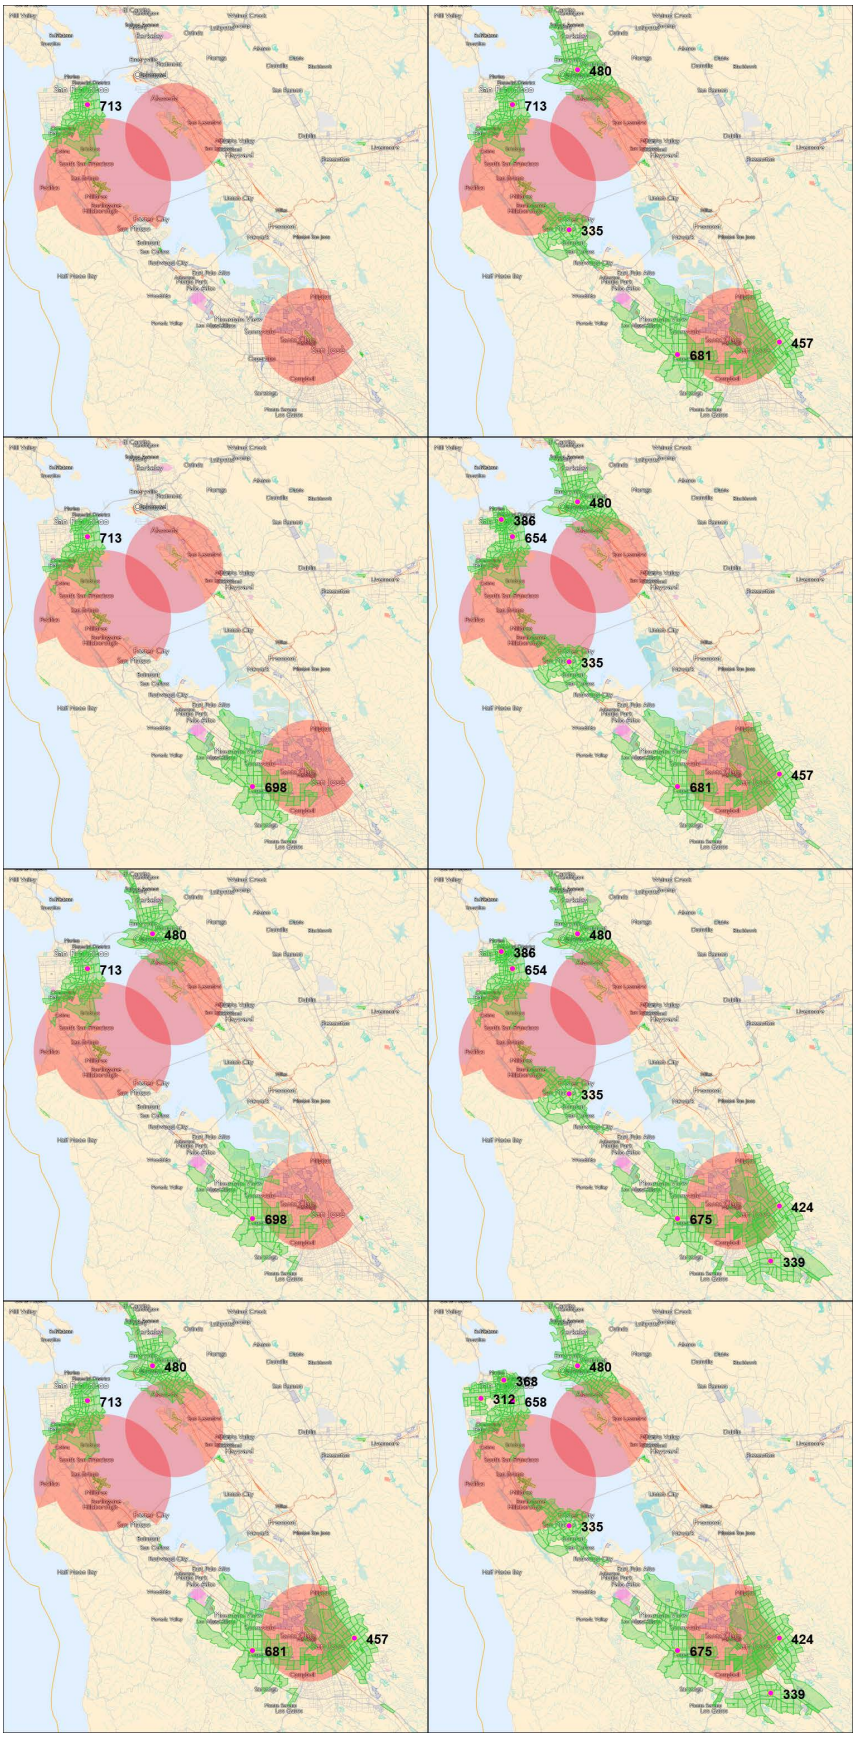

Figure 3.6: Maximal-demand vertiport locations for 1-8 vertiports for a 10-minute driving time neighborhood. Vertiports shown by pink circles, served demand shown by green shading. Red shaded areas show restricted air space. Numbers indicate number of demand units served by each vertiport.

To construct representative missions for sizing eVTOL aircraft, we choose the case with six vertiports because this case shows a good geographic distribution of vertiports and contains a reasonable number to represent a network that might exist within the first few years of introducing this service. The ranges from Tracy, CA, the location of the presumed distribution center, to these vertiport locations is shown in Figure 3.7.

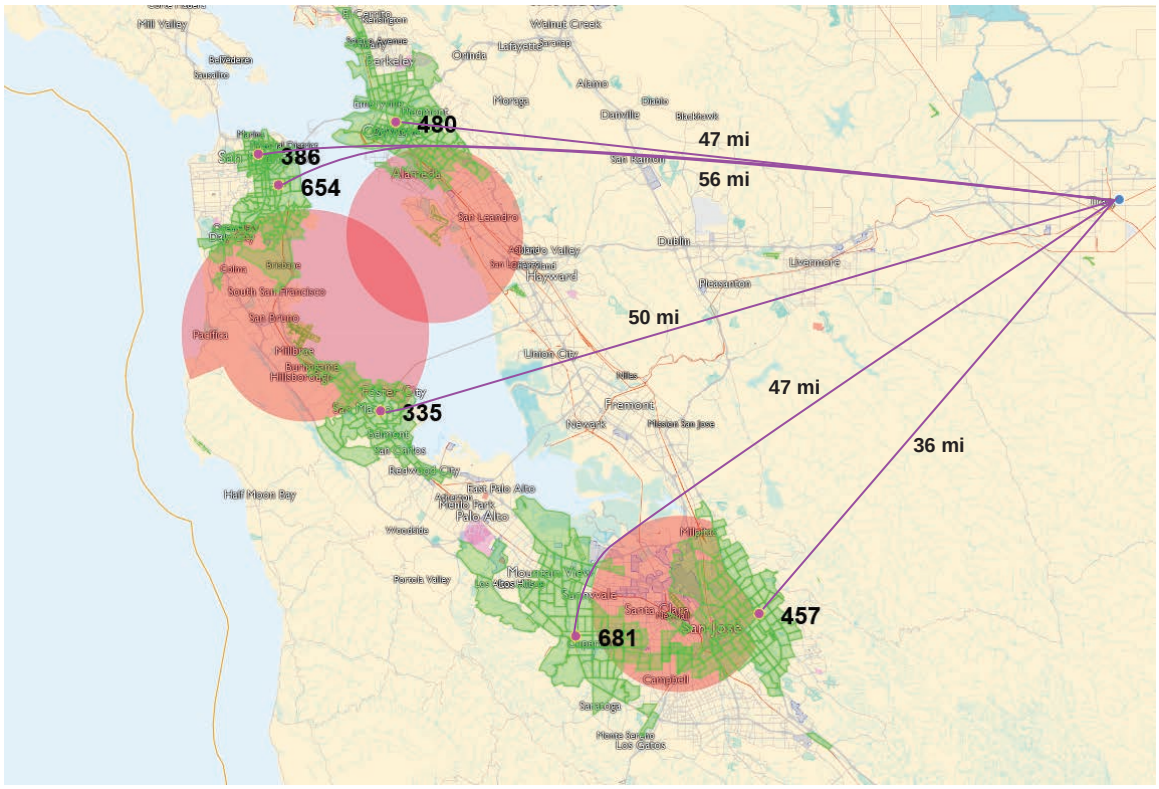

Figure 3.7: Flight ranges from Tracy, CA to six potential vertiports

A comparison of the driving time and flying time from Tracy, CA to these vertiport locations is shown in Table 3.1. We model the flight time as $4\,\text{min} + 60\,\text{min} \times (Range/150\,\text{mph})$, where the fixed time of 4 min is intended to account for climb/descent and terminal area operations during departure and arrival. The table shows a time savings of 32–44 minutes depending on the particular vertiport distance.

Table 3.1: Comparison of driving and flying times

Vertiport	Driving Distance	Driving Time	Flight Distance	Flight Time	Time Diff.
Oakland	55 mi	56 min	47 mi	23 min	32 min
San Francisco, North	64 mi	78 min	56 mi	26 min	38 min
San Francisco, South	65 mi	76 min	56 mi	26 min	39 min
San Mateo	59 mi	61 min	50 mi	24 min	35 min
Cupertino	67 mi	75 min	47 mi	23 min	44 min
San Jose	59 mi	63 min	36 mi	18 min	41 min

3.2 Aircraft Configurations and Characteristics

To examine the concept of operations from the standpoint of aircraft sizing and energetics, we consider two representative UAM aircraft configurations explored in a recent paper by Duffy et al [6]. The first aircraft

concept, often termed a lift+cruise configuration, has eight separate vertically oriented lift propulsors, similar to a multicopter configuration, a wing to offload the lift motors during forward flight, and one horizontally oriented cruise propulsor in a pusher configuration at the tail for forward flight propulsion. Examples of similar lift+cruise configurations that have been discussed include the concepts by Zee Aero [6] and Aurora Flight Sciences [7]

The second aircraft concept is a distributed tiltrotor configuration with six tractor configuration proprotors mounted on the wing and four on the tailplane. In forward flight, these rotors serve in the role of propellers. In vertical flight, the rotors tilt upward to provide lift. Unlike the V-22 and conventional helicopters, the proprotors on this distributed tiltrotor, like those of the lift+cruise concept discussed above, are envisioned to lack cyclic and collective control, and three-axis control of the vehicle in vertical flight is provided by differential motor speed. An example of a similar distributed tiltrotor configuration is the Joby S2 concept [8]. The lift+cruise and tiltrotor configurations examined by Duffy, et al. are shown in Figure 3.8.

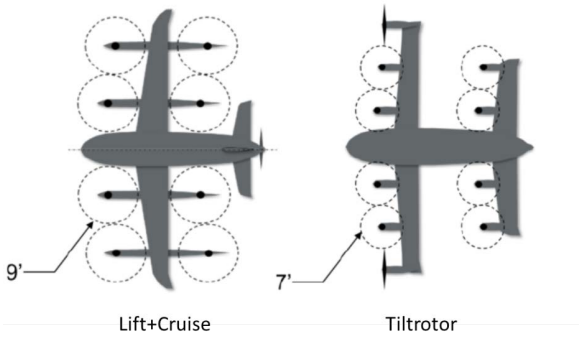

Figure 3.8: Aircraft configurations.

Table 3.2 provides technical characteristics of these aircraft configurations presented or derived from Duffy, et al. Both aircraft configurations are envisioned to have battery electric power systems without internal combustion engines.

Table 3.2: Aircraft characteristics

Characteristic	Lift+Cruise	TiltRotor
Nominal cruise speed (mph)	150	150
Equivalent L/D at nominal cruise	9.1	11
Forward flight efficiency (propulsor and electrical losses)	0.82	0.76
Disk loading (lbs/ft2)	7.3	12.8
Figure of merit in vertical flight	0.68	0.68
Empty weight fraction (w/o battery)	0.53	0.55

3.3 Sizing and Mission Performance

Although the results presented by Duffy et al. provide aircraft sizing and certain mission performance results for the lift+cruise and tiltrotor concepts discussed above, we implemented our own aircraft sizing model to examine the effects of different assumptions regarding battery specific energy, battery discharge behavior, and reserve requirements. We also implemented an off-design performance model to estimate the mission energy expenditure required to fly flights of shorter distances than the design range and also to assess the impact of different strategies for battery recharging between flights. The off-design performance model provides a

Table 3.3: Aircraft weight breakdown

Parameter	Lift+Cruise	Tiltrotor
Gross weight	3980 lb	4860 lb
Empty weight	2190 lb	2673 lb
Payload weight	800 lb	800 lb
Battery weight	1071 lb	1387 lb

mechanism for evaluating the specific flights envisioned in our San Francisco package delivery concept of operations.

The sizing and off-design performance models are formulated at a high level of abstraction, require few high-level inputs, and are suitable for requirements studies and for the early phases of aircraft conceptual design. Vertical flight performance, including hover and vertical climb, is estimated with actuator disk theory based on representative disk loadings and figures of merit. Forward flight performance is modeled based on an equivalent lift-to-drag ratio and/or flight power required at an airspeed consistent with the nominal cruise conditions for a particular aircraft type. The weight buildup approach is to solve for the available battery pack weight based on a prescribed (and fixed) empty weight fraction, payload weight, and gross weight; weight sizing is then completed by iteratively adjusting gross weight. Battery performance is modeled with a simple discharge model based on battery cell manufacturer data sheets, similar in many respects to the Shepherd model [9] and battery cell mass can then be scaled to reflect improvements possible with increased specific energy. The overall approach is broadly consistent with the models developed by McDonald and German [10]. As inputs to our models, we use the technical characteristics for the lift+cruise and tiltrotor configurations examined by Duffy et al. and presented in Table 3.2.

As discussed in Section 3.1, the maximum flight distance between the fulfillment center and any of the vertiports, with consideration of surface airspace restrictions, is 56 miles. Considering the potential for additional vertiport locations to be added in the Bay Area and to account for the range penalty associated with winds and unmodeled airspace restrictions, we therefore selected 65 miles as the sizing range for lift+cruise and tiltrotor aircraft.

In computing the sizing range and off-design performance, we presumed that the battery has a pack-level rated specific energy of 300 Wh/kg, an aggressive goal, but consistent with recent assumptions for UAM aircraft batteries [10]. The battery was presumed to have discharge behavior similar to the Samsung INR18650-25R, a state-of-the-art high power (low internal resistance) cell [11]. The battery is presumed to operate over a range of state of charge between a maximum of 90% and a minimum of 10%. The maximum is intended to reflect the fact that many lithium ion chemistries degrade more quickly if fully charged, and the minimum is intended as a buffer for flight safety by avoiding the steeply nonlinear discharge behavior at low states of charge.

We also modeled a reserve mission that must be flown in its entirety without violating the minimum state of charge constraint of 10%. Currently, there is considerable uncertainty in the specific regulatory reserve requirements that will be established for electric VTOL aircraft. Some studies have examined the implications of following the current Part 91 rotorcraft VFR reserve requirements that stipulate that a flight may not be initiated unless there is "enough fuel to fly to the first point of intended landing and, assuming normal cruising speed, to fly after that for at least 20 minutes" [12]. The General Aviation Manufacturer's Association (GAMA) is currently developing a recommendation for performance measurement for electric VTOL aircraft by requiring a reserve segment consisting of a balked landing followed by a 2 nautical mile flight to an alternate landing site [13]. This short distance to an alternate is reflective of the capability of VTOL aircraft to land in any suitable open area in emergency situations and of the severe challenge of extensive reserve requirements on electric aircraft feasibility. We adopt the GAMA reserve requirements in our sizing approach.

Applying these state of charge and reserve requirements and the technical assumptions provided in Table 2, the aircraft weight results for this sizing range are as shown in Table 3.3.

After sizing the aircraft, we simulated mission performance for various distances appropriate for flights from the Tracy fulfillment center to the vertiport locations defined in Section 3.1. The distances spanned

Table 3.4: Charging time with 300 kW charger

Distance (mi)	Lift+Cruise Time (min)	Tiltrotor Time (min)
35	12.5	16.0
45	14.7	18.4
55	16.9	20.7
65	19.1	23.1

Table 3.5: Charging time with 400 kW charger

Distance (mi)	Lift+Cruise Time (min)	Tiltrotor Time (min)
35	9.5	12.1
45	11.1	13.9
55	12.8	15.7
65	14.4	17.4

from 35 miles, corresponding approximately to the minimum vertiport distance, to 65 miles, the aircraft sizing range. Our goal in these simulations was to determine the battery recharge times that are required after each flight in order to fully recharge the battery to the 90% maximum state of charge threshold. From an energetic perspective, by charging to the maximum capacity after each flight, the aircraft can fly back-and-forth between the fulfillment center and the vertiport indefinitely. If the battery were charged less than this amount of time, the net battery charge between flights would decrease, resulting in a maximum time window during which repeated flights could be carried out before the battery charge depleted to the minimum state of charge threshold. Results for the recharge time necessary for full recharge are shown in Tables 3.4 and 3.5 for charger power capabilities of 300 kW and 400 kW, respectively. These charger power levels are quite aggressive when viewed from the lens of the installed base of current electric car chargers; however, a new generation of automotive chargers are now approaching or exceeding these power levels [14].

3.4 Package Delivery Throughput

Based on the results presented in Section IV, we can estimate an upper bound on daily package delivery throughput per aircraft. First, we presume that one aircraft is assigned exclusively to flights between the distribution center and one specific vertiport. Next, we consider "a day in the life" of each aircraft in conducting back-and-forth repeating trips by building a timeline of its arrivals and departures. We presume that the aircraft recharges to the 90% maximum state of charge at each landing with a 400 kW charger, requiring the times listed in Table 5. We additionally presume that these recharge times are sufficient for package loading or unloading to be fully completed during recharge. Considering the 150 mph nominal cruise speed and 4 min for climb or descent and terminal area operations, the total time required for a round trip flight (two flight legs and two recharges) can be computed. Additionally, from initial cabin sizing studies, we presume a cabin volume of 125 cubic feet and a volumetric efficiency of 65% for packaging typical small rectangular delivery boxes. We presume that the maximum payload capability of 800 lb may be fully utilized. We also presume a 12 hr daily period of continuous operations. The corresponding throughput results are shown in Table 6 and Table 7 for the lift+cruise and tiltrotor configurations, respectively.

This throughput analysis must be recognized to suffer high degrees of uncertainty in the assumptions outlined above. Additionally, for reasonable and typical implementations of this concept of operations, aircraft will likely be scheduled with high frequency departures to achieve adequate order-to-door time savings, and payload load factors are therefore likely to be considerably lower than the maximum values indicated.

In the next section, we extend the methods presented in this section to optimize placement of vertiports for a UAM passenger service for commuters. The resulting UAM network is then simulated with both

Table 3.6: Lift+Cruise daily package throughput.

Distance (mi)	Flight (min)	Recharge (min)	Round Trips per Day	Maximum ft^3/day	Maximum lb/day
35	18	9.5	13	1056	10400
45	22	11.1	10	813	8000
55	26	12.8	9	731	7200
65	30	14.4	8	650	6400

Table 3.7: Tiltrotor daily package throughput.

Distance (mi)	Flight (min)	Recharge (min)	Round Trips per Day	Maximum ft^3/day	Maximum lb/day
35	18	12.1	11	894	8800
45	22	13.9	10	813	8000
55	26	15.7	8	650	6400
65	30	17.4	7	569	5600

passenger and cargo demand.

4. Package Delivery to Amazon Lockers at Passenger Vertiports

A goal of this research project is to assess how cargo service can be incorporated into an existing passenger UAM network. For this reason, we first needed to design a passenger UAM network in the San Francisco Bay Area, and then develop a UAM operational simulation capability to "fly" aircraft in the UAM network and aggregate statistical results including wait times, throughput, aircraft repositioning, and demand that can be served. The resulting network and simulation was then used to interleave cargo flights to determine whether and how cargo service could improve aircraft utilization. The goal of improving aircraft utilization with cargo flights is primarily to reduce the ticket price for UAM passenger air service.

Intra-city passenger air service in the context of UAM is envisioned to serve trip purposes including airport transfer, leisure trips, and daily commuting to work [15]. We considered how vertiports should be placed within cities to maximize the time savings potential of UAM compared to driving for daily commuting. The discussion about modeling commuter demand and vertiport placement this section is drawn largely from an AIAA Aviation 2018 paper prepared based on this research project [16]. The results regarding the network simulation of passenger and cargo flights are new and have not been previously published. We present results for both the San Francisco Bay Area (the primary focus of this research project) and also the Los Angeles area. The Los Angeles study provided an additional data set to assess our modeling and optimization approach.

4.1 Modeling Potential eVTOL Commutes with Census Data

Data from three different sources are considered in this study. The first data source is the American Community Survey (ACS), prepared by the U.S. Census Bureau [17, 18]. The ACS is a rolling survey of approximately 3.5 million households per year, with data published in single-year or past-five-year estimates. The variables surveyed by the ACS that are relevant to our problem of modeling commutes include population, household income, mode of transportation for commuting, time leaving home to go to work, and travel time to work. ACS data is published only as aggregated statistics – individual data points are not available. Because of

this, certain correlations between variables cannot be determined. For example, the ACS publishes a data table that contains both earnings and mode of transportation, so this correlation can be explored, but the data set does not provide a table indicating time leaving home to go to work and travel time to work, so this correlation cannot be explored. All ACS data used in this study are from the 2012–2016 5-year estimates.

The second data source is the Longitudinal Employer-Household Dynamics Origin-Destination Employment Statistics (LODES) survey [19]. This data set provides detailed information that can be used to infer "origin-destination" home-work pairs: for each ordered pair of census blocks, the LODES data tells how many people live in the "origin" block and work in the "destination" block. However, caution is needed in interpreting LODES data for understanding commuting because unlike ACS, LODES data is based on administrative records, not surveys of workers themselves. One consequence is that the LODES data does not directly measure commutes, since there is no guarantee that the destination address is co-located with the worker's office (vs., e.g., a company's headquarters) or that the worker does not work remotely. Nevertheless, the LODES data is the closest surrogate for point-to-point commute counts that is freely available, and we consider it to be sufficient for this study, when used with caution and after applying relevant filtering.

The third data source is the TIGER/Line data sets from the U.S. Census Bureau which provide geographical boundaries for the census blocks and tracts [20]. TIGER files were used to prepare the maps in this report and to identify interior points within each census block for the purposes of calculating straight-line block-to-block distances that are used to model flight ranges.

4.1.1 Estimating current commute time and number of commuters

For each pair of census blocks, (H, W), the LODES data provides a count of the number who reside in block H and whose work address is in block W. Because of the small size of a census block, these pairwise counts are usually only one or two people. As discussed above, this data is not exactly equivalent to the number of people who commute between blocks H and W, since it does not guarantee that the person's physical place of work is at W, only that their employer's recorded address is at W. We describe how we account for this discrepancy below.

First, however, we compute commute distance and time measures between H and W. We treat each census block as a single latitude/longitude point, corresponding to the "interior point" datum listed in the 2017 TIGER/Line shapefiles. This is not necessarily the centroid of the block; it is only guaranteed to be a point inside the block. Denoting these points h and w, we compute the "flight distance" between the census blocks, $d_{h,w,air}$ using the haversine formula. This represents the potential shortest path from H to W using eVTOL service, but is not a good estimate of the current commute distance by car. The drive distance by car, $d_{h,w,car}$ is estimated by using the open source GraphHopper routing engine and OpenStreetMap data to calculate turn-by-turn driving directions from h to w – similar to those provided by any GPS navigation app. This provides $d_{h,w,car}$ as well as an estimate of the driving time, $t_{h,w,car}$.

4.1.2 Modeling traffic

The driving times calculated by GraphHopper are primarily based on speed limits and do not take into account traffic. In order to simulate the effect of traffic, we scale the drive times obtained from GraphHopper to more closely match the driver-reported commute times in the ACS data. Recall that the ACS data provides a distribution of commute times as reported by respondents, but does not provide information on the destination of their commute. For this analysis we use the ACS data for commute time by mode of transit (ACS table B08134), considering only the mode "car, truck, or van." We do not further reduce our consideration to only those who reported driving alone at this step. This approach to modeling traffic differs from our approach described in Section 3.1 and was adopted based on considerations of scalability and use of fully public domain data instead of the Google Maps API.

The ACS data gives the number of respondents who report commute times within differently-sized bins between 0 and 60 minutes. All commutes over 60 minutes are grouped into a single bin. Our first step in scaling the commutes is to reproduce a population of commutes from the binned ACS data such that the binning of this population matches the distribution found in the ACS data. First, we scale the counts in the ACS data to match the count of trips we have in the LODES data. (All we are concerned with in the traffic scaling analysis is the *distribution* of times and this allows for a simple 1-to-1 comparison between the data

sets; we correct for the count of commuters in the final analysis as described below.) For example, if the total ACS count of commuters who reside at H and commute by car, truck, or van is 1000, and the count of LODES trips originating at H is 1500, we multiply all ACS bin counts for H by 1.5 so that we sample an artificial ACS population of 1500 members.

Next, we sample the population by simply assuming a uniform distribution of commute times within each bin. For example, if for a particular census tract, our corrected bin-count for the "15-19 minutes" bin is 253 commuters, we add 253 uniformly distributed values between 15 and 20 to our artificial ACS population. (We use 20 rather than 19 because 20 is the lower limit of the next bin; ACS data is rounded to the nearest minute but our population is not.) For the 60+ minute bin, we use the 98th percentile of the LODES/GraphHopper trip times as the bin's upper limit for the purpose of sampling.

We now have two sampled populations of commute times originating at H: The reconstructed ACS population described in this section, and the drive times obtained with GraphHopper and LODES data described above. Sorting each of these populations and plotting them against each other for each origin census tract yields the image in Figure 4.1 for the San Francisco metro area. In this plot, each (monotonically increasing) curve represents the commutes originating from one of the 1800 census tracts in the study area. Two effects are prominently seen. First, the bulk of the curves fall below the one-to-one parity line of a perfect fit. This indicates that, in general, the ACS-reported commute times are longer than the LODES/GraphHopper calculated driving times. This result is expected because the GraphHopper routing is based on OpenStreetMap data that does not contain traffic information (only speed limits); therefore, our LODES/GraphHopper calculation underestimates actual commute times. This trend confirms our motivation for correcting the drive times to account for traffic by calibrating to ACS-reported commute times.

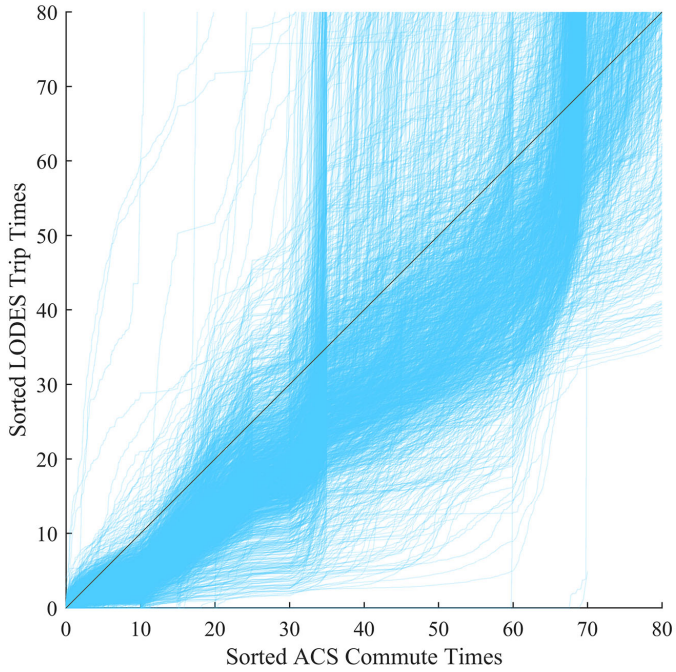

Figure 4.1: Sorted reconstructed ACS commutes vs. corresponding LODES trip time computed with GraphHopper

Second, we see that many of the curves become nearly vertical as time increases. This behavior indicates that the GraphHopper results are vastly *over-estimating* the commute time for these cases. We believe this is because of the aforementioned shortcoming in the LODES data – it lists the workplace according to administrative records, not a worker-reported job site, implying that it includes some reported workplace

locations that are far from actual work locations, perhaps even several states away. For example, Figure 4.1 has many nearly vertical lines near $x = 35$. This trend indicates that no commuters in these census tracts reported commutes exceeding 35 minutes in the ACS, but according to the LODES records, some of their workplaces are far enough away that it would take ¿80 minutes to drive. The workplace records in LODES for these commuters apparently do not represent the commuters' actual workplace.

To overcome this shortcoming, we discard the top 30% of data points with respect to commute time within each census tract when computing the effect of traffic. In other words, our traffic calculation is based only on shorter-distance trips, which are more likely to represent actual commutes. Trimming the data in this way also avoids the difficulty of modeling the ACS bin for commutes of 60 minutes or greater, for which the upper limit is unknown. (As stated above, we model this upper limit as the 98th percentile of the trip times computed from LODES.) A similar plot to Figure 4.1 for this reduced data set is shown in Figure 4.2, which does not exhibit the problematic sharp increases seen previously.

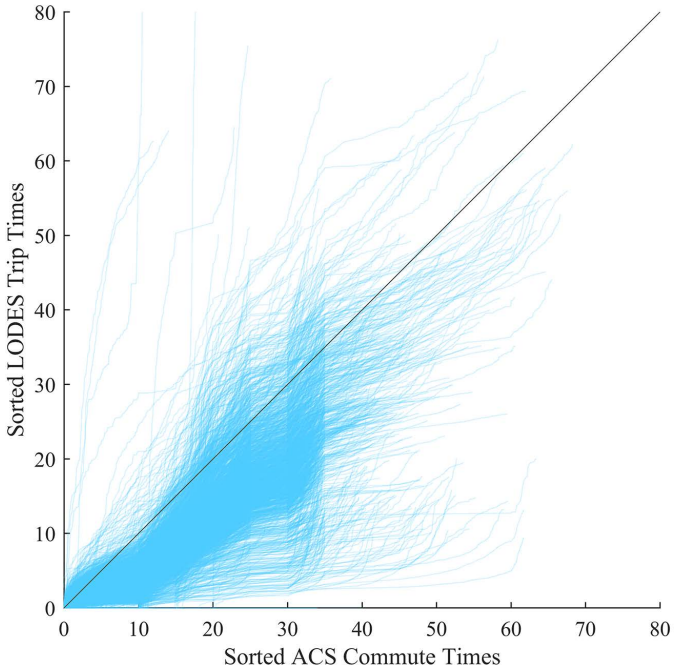

Figure 4.2: Sorted reconstructed ACS commutes vs. corresponding LODES trip time computed with GraphHopper

Finally, from these trimmed populations, we calculate a traffic scaling factor for each census tract by using a least-squares fitting procedure to calculate the scalar value s_H that minimizes the squared errors between the ith member of the populations of the reconstructed ACS data and the corresponding member of the LODES trip data, summed over all i. This computation is completed independently for each census tract, yielding a different traffic scalar for each tract. The LODES trip times, computed using GraphHopper, are then multiplied by this scalar to obtain the final effective driving time between each pair of census tracts. The resulting driving times between tracts are those that will be used in our vertiport placement optimization.

The effect of the traffic scalar is easily illustrated with an example. Figure 4.3 shows the cumulative distributions of commute times for a representative census tract in the San Francisco area. The LODES/GraphHopper data is binned according to the same commute times used in the ACS data; the ACS curve shows the actual ACS data, not the reconstructed population. The curves show that the scaled LODES/GraphHopper trip times data much more closely matches the distribution of commute times reported in the ACS data than the unscaled LODES/GraphHopper trip times, which severely underestimate

commute times.

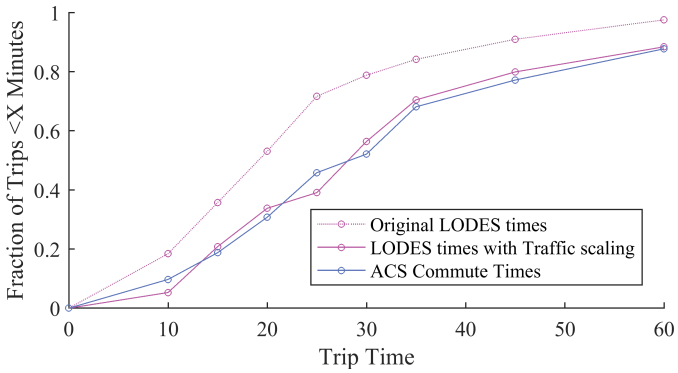

Figure 4.3: Effect of traffic scaling on modeled commute times for a single census tract

One shortcoming of this approach is that it applies the same scaler to all trips from the home census tract. Consequently, it attributes traffic to the location of a person's home, rather than with the route they travel or the location of their work.

4.1.3 Modeling the number of commuters

Next, we scale the LODES counts to attempt to better represent the actual number of people commuting, and even further, to represent the number of people who would be potential eVTOL riders. First, we scale the counts so that the total number of counts originating at H matches the corresponding ACS total estimate of commuters whose home is in H. Next, we reduce this number to attempt to better represent the potential eVTOL passenger pool. For the purpose of this study, we assume the potential eVTOL passenger pool to be comprised of high income individuals who currently drive to work alone. We assume that there is no relation between income and driving alone, and distance or time of commute, so we simply multiply every LODES count originating at H by the same scaling factor to uniformly reduce the counts. ACS data table B08119 provides counts of commuters split by mode and income level; from this data, we use the count of commuters who drove to work alone and who have an individual annual earnings greater than $75,000. Our model for the final effective count of potential eVTOL commuters is

$$count_{H,W} = (\text{LODES count from H to W}) \frac{(\text{Total ACS commuters from H})}{(\text{Sum of LODES counts from H})} \frac{(\text{ACS drive alone + ¿\$75k})}{(\text{Total ACS commuters from H})}$$

Generally, this estimate will be a decimal number, which may be thought of as the "typical" number of daily trips by high income individuals driving alone on any given day. The geographical distribution of potential commuters is shown in Figure 4.4.

4.1.4 Data modeling summary

For the purposes of this study, we assume that the typical near-future eVTOL commuters will be people who have a high income, currently drive alone, and have a long-duration commute relative to the straight-line distance between their home and work. In order to model these commuters, we need to know several things: How many high income commuters that drive alone live in each census tract? Where do they commute to? And how long does their commute currently take? The number of high income commuters is directly estimated by the ACS data. The question of where they commute to is not directly answered by any census data set, but is most closely approximated by the LODES data, which lists workplace addresses according to administrative records. Finally, the question of how long the commute takes is the most difficult to answer. ACS provides distributions of commute times for each home census tract, but does not show where each commuter is going. LODES provides information about where people are going, but not how long it

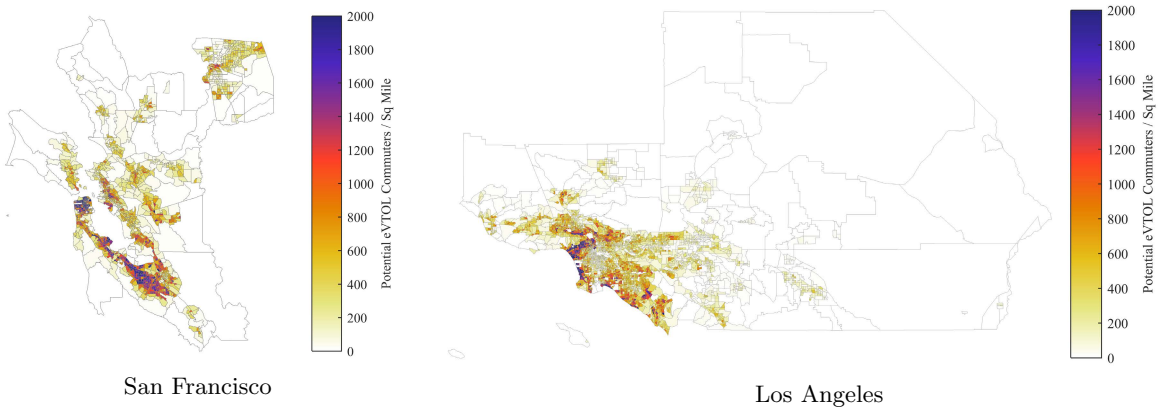

Figure 4.4: Density of potential eVTOL commuters, defined as individuals with annual earnings greater than $75,000 who drive to work alone

takes to get there. We have used the LODES location data as an input to a driving-directions algorithm, GraphHopper, based on OpenStreetMap data, to approximate a driving time for each trip. We then scale these driving times so that their distribution more closely matches the self-reported commute times in the ACS data.

Finally, we note that the LODES data used was tabulated at a census-block level, while the ACS data is available only at the coarser census-tract level. The block-level LODES data was used while calculating the commute distances and times with GraphHopper, since it provides more accurate estimates of locations. After these calculations were completed, we aggregated the block-level commute data to the tract level by taking the count-weighted average. This tract-level data was then used to compute the traffic scalers and throughout the rest of the study.

4.2 Optimizing Vertiport Placement

From the data described above, we formulate a (binary) integer programming optimization model to place vertiports to optimally serve the potential commuters.

To formulate the problem, we discretize vertiport placements by allowing a vertiport to be placed in each census tract; i.e. there is one variable per census tract representing whether or not it contains a vertiport. Each vertiport is presumed to have unlimited capacity; we later examine the resulting number of flight operations at each vertiport. Vertiports are allowed to serve not only their own census tract but all census tracts within a small neighborhood around them. This formulation ensures that the optimizer does not place too many vertiports in a nearby area when really one (larger) vertiport would suffice. This allowance is especially necessary because we limit the total number of vertiports that the optimizer may place. The "catchment" neighborhood around a vertiport is defined as the set of census tracts that are within 3 miles *and* within a 5 minute drive.

A second, much larger, set of variables represents whether the commute connecting each *pair* of census tracts can be served by eVTOL trips. These variables can only be set to "true" if there is a vertiport placed within the neighborhood of both the origin (home) and destination (work) census tracts. We simplify the problem by only creating these variables for pairs of census tracts separated by more than 30 minutes of driving (with traffic), reasoning that trips shorter than this threshold are less appealing as eVTOL flights. We also exclude any pair of census tracts with less than 0.1 potential eVTOL commuters per day.

The objective function of the optimizer is to maximize the cumulative time saved in commuting throughout the entire network. For each of the commute variables described above, we associate a value equal to the number of potential eVTOL commuters who live in the origin census tract multiplied by the time savings that would be achieved by completing this trip by eVTOL. In this study, we model the time required for an eVTOL flight as

$$4min + \frac{\text{Straight-line distance between census tracts}}{150mph} \frac{60min}{1hr}$$

presuming a 4 minute time total time for climb/descent and hover (inclusive of both takeoff and landing), a 150 mph eVTOL nominal cruise speed, no wind effects, and no routing penalty relative to direct haversine distance. This flight time is compared to the drive time with traffic to determine the time savings achievable with eVTOL.

We consider case studies for two cities, San Francisco and Los Angeles, and we complete three optimization studies for each city: a small network of 10 vertiports, a medium network of 20 vertiports, and a large network of 40 vertiports. The increase in number of vertiports may be thought of as approximating the growth of a network over its first several years of operation.

The optimization problems were modeled and solved using the Gurobi optimization solver. We ran the cases for approximately 4 hours on a modern 24-core workstation. These are very large, difficult to solve problems; they contain over a million variables since the number of variables scales as (number of census tracts)^2. They are also poorly bounded by their linear relaxation. The poor bounds make it difficult to tell how well the optimization solver has converged to the (global) optimal solution. However, since we see similar geographical patterns in the different-sized cases, we are confident that the results are representative of "good" networks, even if they may not be strictly "optimal."

Results of the case studies are shown in Figures 5 through 7. In Figures 5 and 6, the blue dots show the locations of vertiports placed by the optimization algorithm. In the second and third columns, these dots are sized according to the number of potential trips (i.e. the number of trips if all of the potential commuters were served) for which that vertiport serves as the commuter's "home" or "work" location. In the first column, the purple lines show connections between census tracts of eVTOL trips served by these vertiports. Thicker, darker lines indicate more trips being flown between the corresponding tracts. Note that these lines are drawn to connect the commuter's home and work tracts, which may be any tract within a 5 minute drive of the vertiport location.

Figure 4.5 shows the results of the San Francisco case study. Each row of the figure shows one of the three networks considered: small, medium, and large. The results show that the most appealing trips based on the existing census commute data are relatively short range trips around San Francisco Bay. For "work" destinations, downtown San Francisco, near the Financial District, is the most important vertiport location. The "home" locations are more evenly spread around the metro area. Notably, few home-work trips originate from downtown San Francisco. This is to be expected – people who live downtown probably also work downtown, and therefore have little need for an eVTOL service in the context of commuting.

As the network size is increased, we see a few outlying vertiports, with one placed near Sacramento and another situated south of San Jose. Although the Sacramento vertiport is only lightly used, it serves the longest possible commutes in the area, so it offers great benefits to its passengers.

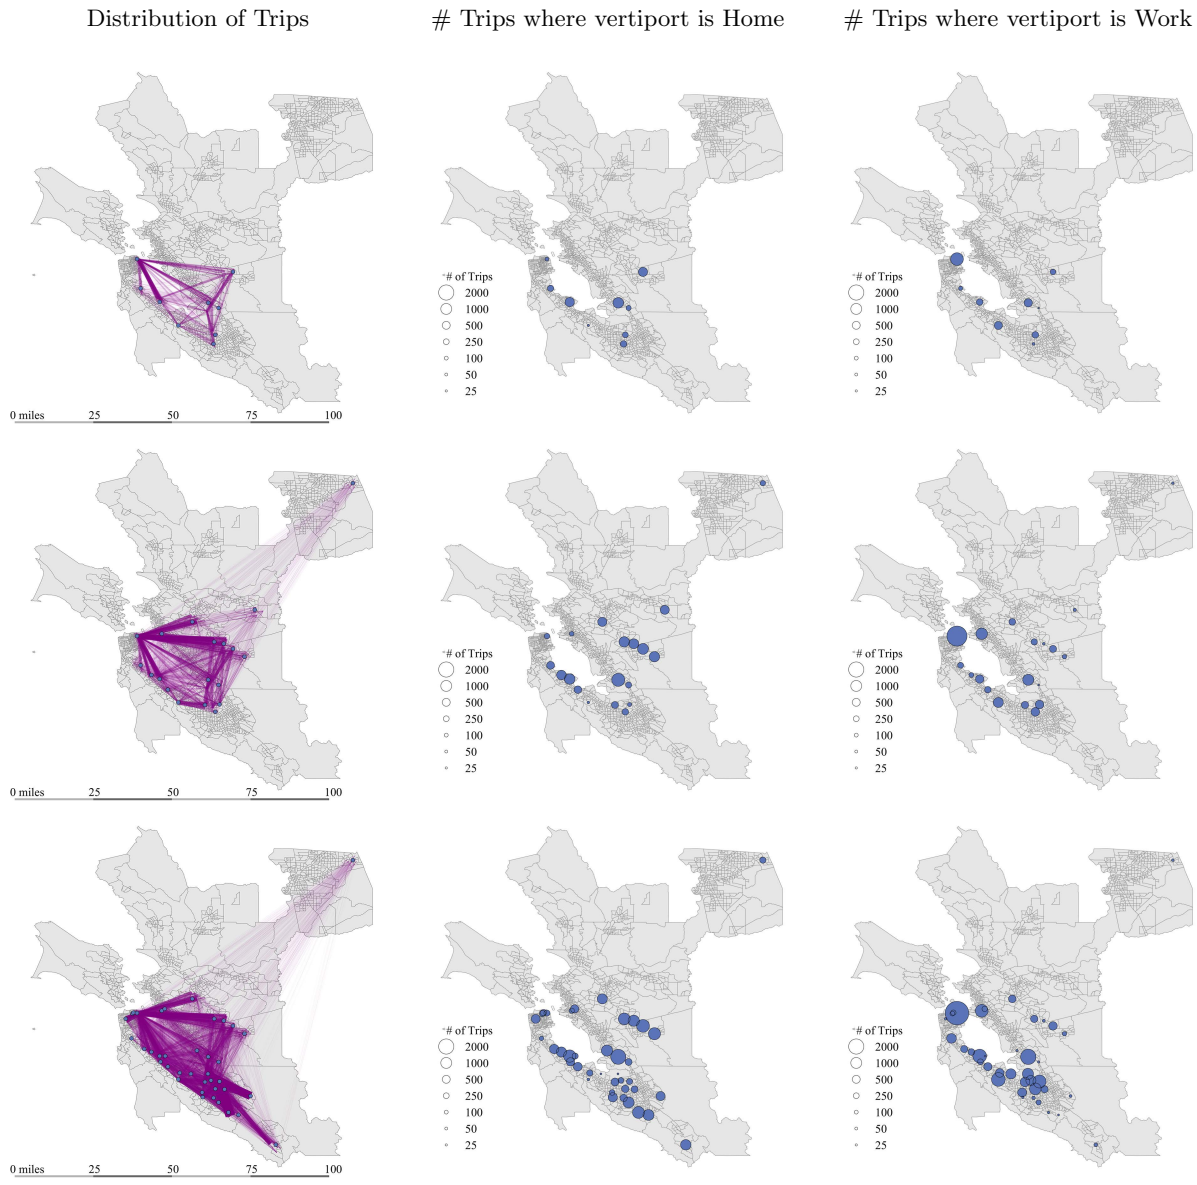

Figure 4.5: Optimization results for San Francisco area. Row 1: Small network; Row 2: Medium Network; Row 3: Large Network

Figure 4.6 shows the results for the Los Angeles case studies. Somewhat surprisingly, the network is very strongly aggregated in a small area around Santa Monica, Burbank, and downtown Los Angeles. Comparing with the distribution of likely commuters shown in Figure 4, this seems to be explained by the density of the high-income commuters in these areas. The ranges of flights in this central cluster are generally very short – as little as 5 miles. However, the severe traffic in the area can make even such short trips difficult by car. As we increase the network size, we begin to see more flights serving the outlying areas, but the strongest growth in eVTOL use remains in the central area.

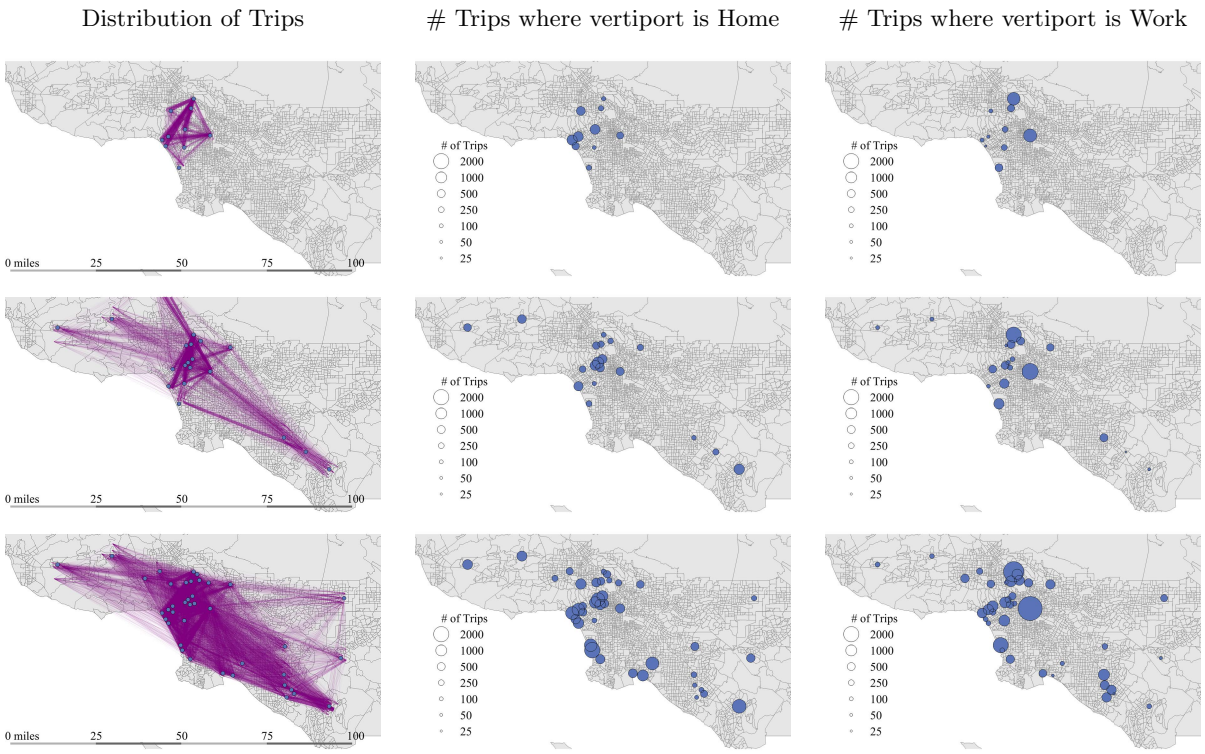

Figure 4.6: Optimization results for Los Angeles area. Row 1: Small network; Row 2: Medium Network; Row 3: Large Network

Figure 4.7 and Figure 4.8 show the distribution of haversine ranges in trips servable by eVTOL for each of the six cases. Detour penalties based on airspace restrictions or other considerations are not modeled. These results show that regardless of network size, short trips under 30 miles dominate the eVTOL commuter market in these cities. Only a few trips extend beyond 60 miles. However, since this study is based on existing census data records, the ranges are skewed toward where people live and work now. Likely, once an eVTOL service is in reliable at-scale operation, people may choose to move farther away from their jobs, resulting in longer-distance commutes.

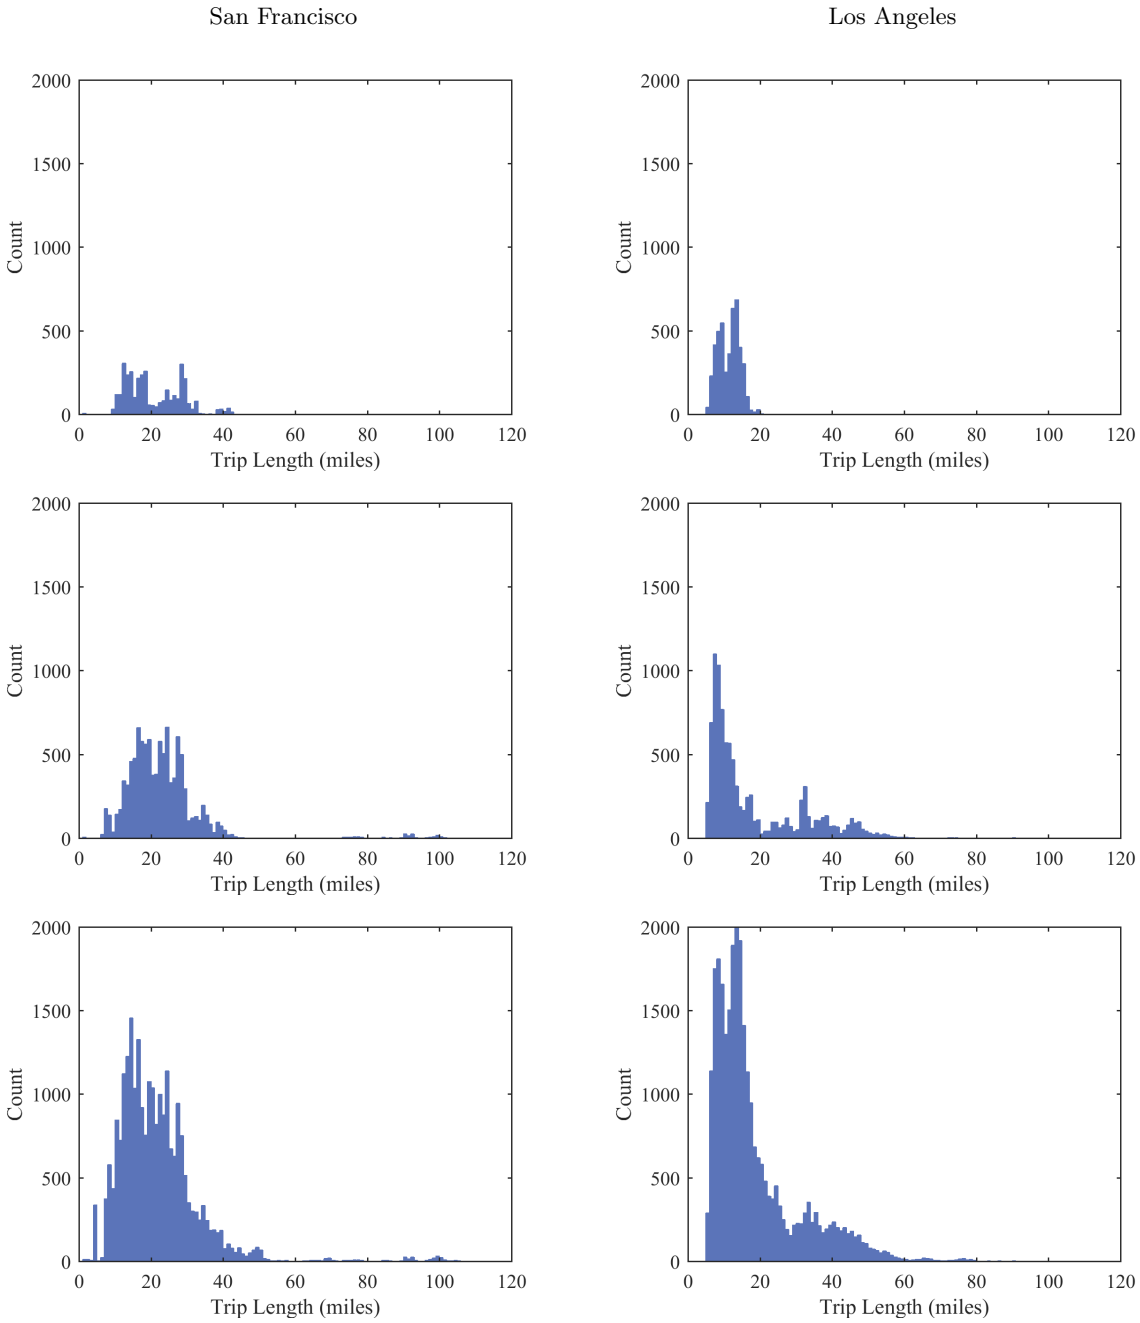

Figure 4.7: Distribution of ranges of trips served by eVTOL. Row 1: Small Network; Row 2: Medium Network; Row 3: Large Network

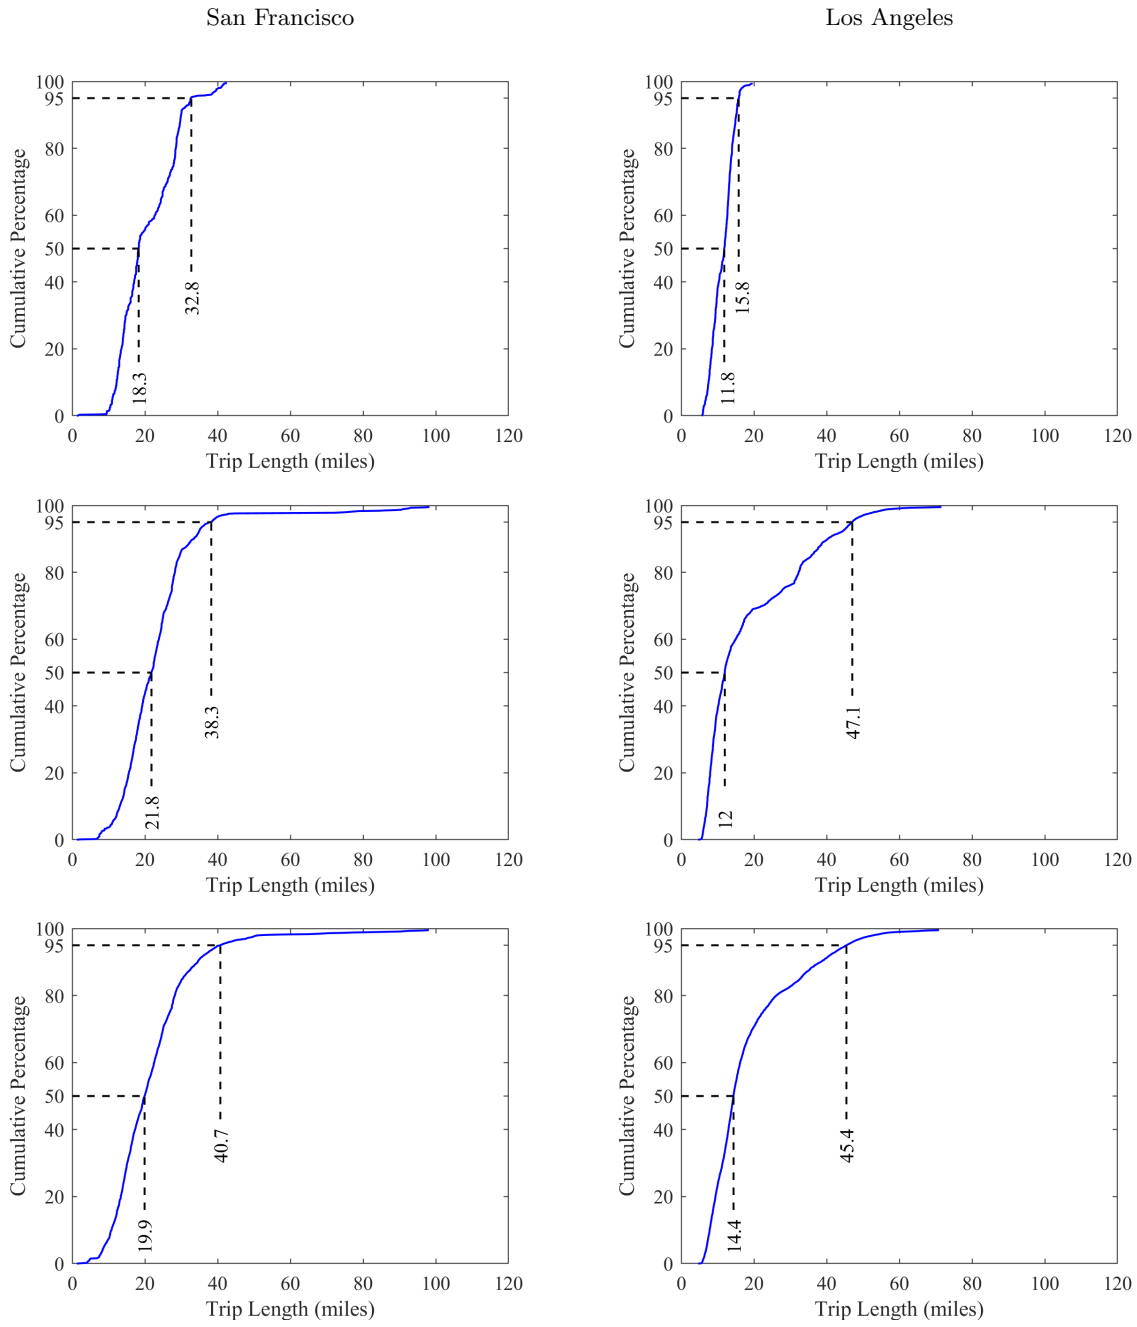

Figure 4.8: Cumulative distribution of ranges of trips served by eVTOL with indication of 50th percentile (mean) and 95th percentile ranges. Row 1: Small Network; Row 2: Medium Network; Row 3: Large Network

4.3 Simulating Passenger and Cargo Flights in the Network

4.3.1 Simulation Capability

Having obtained a good arrangement of vertiports from the optimization model, we next seek to simulate the performance of a fleet of aircraft serving these vertiports. To facilitate this, we have created a continuous-time network simulator in Java.

The network simulator manages a fleet of aircraft, a schedule of demanded trips, and a set of vertiports. Demanded trips are defined by a number of passengers and/or volume of cargo to be moved between a pair of vertiports at a particular time. The fleet is controlled by a centralized dispatching algorithm that determines when an aircraft should be dispatched to serve a demanded trip or to reposition (empty) from one vertiport to another. Aircraft can only be dispatched if they are idle and have a high enough state of charge to serve the desired trip. Aircraft performance and state of charge are modeled with the techniques described in Section 3.3.

The simulator continuously loops over small time steps, each representing about 15 simulated seconds, and updates the state of all aircraft, vertiports, and demanded trips at each time step. Aircraft on the ground recharge at a constant power, and aircraft in flight move toward their destination and discharge their battery according to a low-order physics-based flight model. The simulator runs long enough to simulate a 24-hour day of operations, which takes about 7 minutes of wall time.

At each time step, newly demanded trips appear according to a fixed schedule defined in an input file. These trips are defined by an origin, destination, desired departure time, number of passengers, and pounds of cargo. The dispatching algorithm attempts to assign each demanded trip to an aircraft as soon as possible by first attempting to use an idle aircraft already on the ground at the origin, and second trying to find a nearby idle aircraft that can reposition to the trip's origin and then serve the trip. Trips that remain unassigned after 20 minutes beyond their desired start time are considered "failed" and deleted from the queue. (Note that the passenger's actual wait time may exceed 20 minutes if an aircraft has been assigned but needs to reposition and/or recharge before serving the passenger.)

The simulator records a log of the state of each aircraft at each time step, including its location, state of charge, payload, and (if in flight) speed and altitude. A second log records information about which demanded trips were served and which failed. These logs allow many simulation performance metrics to be computed, from detailed information about a single aircraft to network-wide summary statistics.

4.3.2 Description of Scenarios

For this study, we examined the performance of a mixed passenger/cargo network that involves carrying passengers in their daily commutes and carrying Amazon packages moving between distribution centers and Amazon Lockers located at the commuter vertiports. The reasoning for this CONOPS is that by locating the lockers at vertiports, passengers can pick up their Amazon packages as they pass through the vertiports, eliminating the need for a "last-mile" delivery.

Simulations were run for the 20-vertiport case examined in the optimization studies described above. The first step of setting up the simulation is to construct a fleet and demand schedule to be simulated. Four fleets were considered, consisting of 20, 40, 60, and 80 aircraft. All aircraft are identical and are based on the Uber eCRM1 aircraft model.

Three demand scenarios were considered. In each of scenario, the passenger demand is set to be identical, consisting of two normally-distributed peaks around morning and evening rush hour. (Peaks are centered at 8am and 5pm, with a 90 minute standard deviation.) The total number of trips demanded was set at 5% of the number of potential commuters that live and work within 5 minutes of a vertiport as calculated in the optimization studies described above. For this study, we do not attempt to model exact number of passengers on each flight, rather, we simply model the total number of trips demanded; each trip could involve 1-4 passengers. The total number of passenger trips demanded for the day is 994.

The demand scenarios differ by how much cargo is demanded. In the first scenario, we presume no cargo demand, only passenger demand. This scenario serves as a passenger-only baseline for comparison purposes. In the second scenario, one cargo flight is demanded from Amazon's OAK5 fulfillment center in Newark, CA to each commuter vertiport every hour, for a total of 480 cargo flights per day. In the third scenario, one cargo flight per hour is demanded from Amazon's OAK4 fulfillment center in Tracy, CA to each commuter vertiport. The Tracy fulfillment center is farther from most of the vertiports than the Newark distribution center, but it is a larger Amazon fulfillment center and can presumably serve a wider variety of product orders. The locations of the commuter vertiports and the two Amazon fulfillment centers are shown in Figure 4.9.

We believe that one flight per hour from the fulfillment centers to each vertiport is a reasonable assumption to provide the level of timely service that might be expected from eVTOL package delivery, particularly for

delivery to Amazon Lockers situated at the vertiports. Longer delivery times would likely motivate the use of truck delivery for lower costs. An hourly delivery rate would likely allow Amazon to guarantee delivery within the 2 hour delivery window for Amazon Prime Now (providing sufficient time for order processing at the fulfillment center). With flights from the fulfillment centers, the Prime Now service could offer access to a much greater percentage of the Amazon product inventory than would likely be possible by sourcing from smaller urban warehouses.

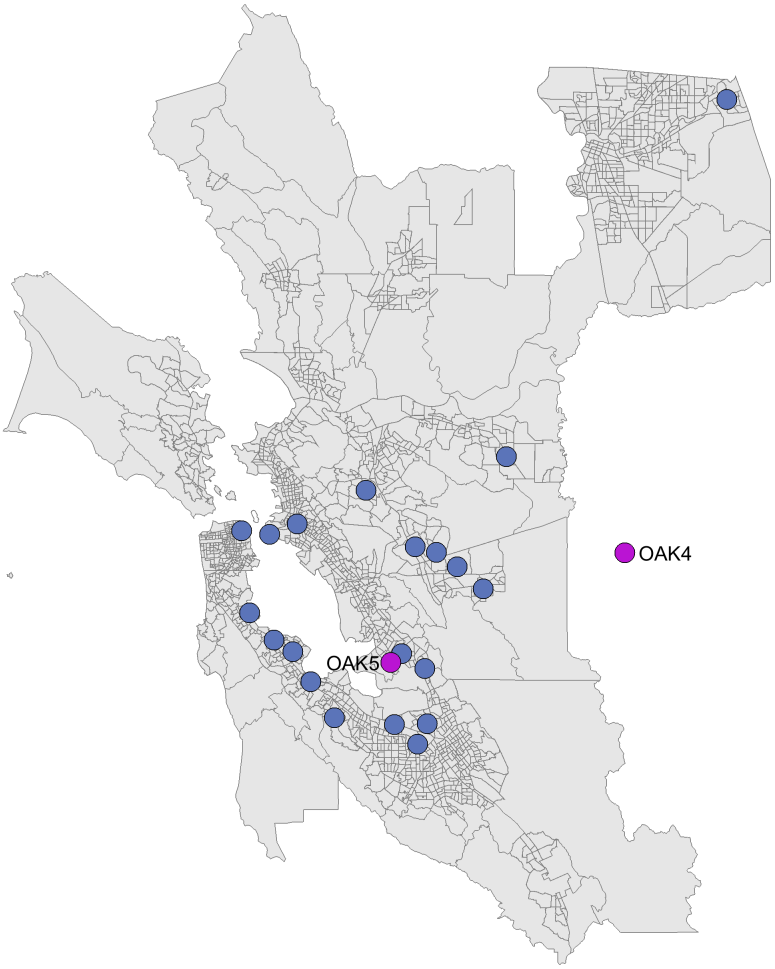

Figure 4.9: Map of commuter vertiports and Amazon fulfillment centers. Background shows extent of census tracts considered for commuter travel.

Finally, we note that in this simulation we assume that each vertiport has unlimited parking pad and FATO areas available, i.e. the vertiports are unconstrained in terms of aircraft capacity. This assumption avoids the need for complex additional logic that would redirect flights in the event that a vertiport is full.

4.3.3 Simulation Results

The results of the simulations are shown in Figures 4.10, 4.11, and 4.12, and Table 4.1 indicates the number of trips for each simulation that the network failed to serve, i.e. trips that involved a wait time greater than 20 mins.

Figure 4.10 shows the results of the passenger-only simulation. Each subplot shows the time-history of the fleet for one of the four fleet scenarios (20, 40, 60, or 80 ships), indicating the number of aircraft in different operational states at each time step.

At small fleet sizes, large amounts of demand go unserved (see Table 4.1) and the passenger flights are spaced over a wider time period as wait times for flights increase. Nearly half of the overall flying time is spent repositioning as the fleet is in a constant struggle to meet the demand, and during the morning and afternoon rush hours the entire fleet is either flying or charging at all times. Moving to larger fleet sizes improves overall service, with the 80-aircraft fleet being able to serve all 994 passenger flights with relatively little repositioning. However, this improved service comes at a large cost – at any given time, at least 30 aircraft are on the ground, and overall, the fleet spends the vast majority of the day sitting idle.

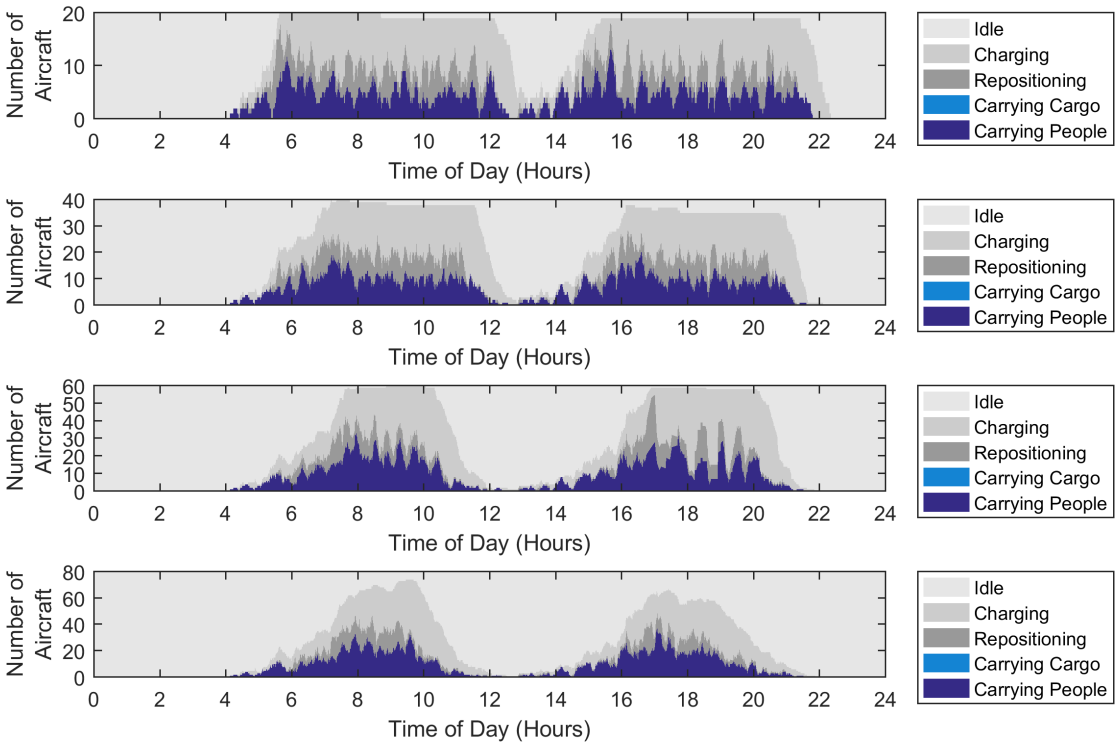

Figure 4.10: Simulated time history for passenger-only network with fleet size = 20; 40; 60; 80 ships. (Note different y-axis scaling.)

Figure 4.11 shows the results of the simulation when cargo demand is added from a dedicated vertiport situated at Amazon's OAK5 distribution center in Newark, CA, which is fairly centrally located relative to the 20 vertiports that serve commuters. At each hour of the day, a flight from this distribution center to each of the commuter vertiports is demanded. These flights are served at a "lower priority" than the passenger flights, meaning that during the peak rush hours the cargo flights will tend to go unserved. (Note that every cargo flight is preceded by a repositioning flight to OAK5, since by definition OAK5 is never a revenue destination.)

In the 20 aircraft case, almost all of the cargo demanded between 6am-11am and 4pm-9pm goes unserved. Overall, nearly half of the total cargo for the day goes unserved, and passenger performance also suffers compared to the no-cargo scenario. Still, the fleet is more productive overall because cargo demand increases fleet utilization in the late-night hours. Although it is unclear whether delivery to Amazon lockers would nominally be carried out in the night hours, we view this as a particularly compelling opportunity to leverage the aircraft fleet for deliveries of consumer orders placed the night before for pickup at the beginning of a morning commute.

Increasing the fleet to 40 aircraft roughly halves the number of failed trips, but increasing to 60 aircraft has less incremental impact. Finally, increasing to 80 aircraft yields good results overall, resulting in 74 total failed trips for the day and overall low waiting times for both passengers and cargo. Even with 80 aircraft, there is a 4-hour period each morning where every aircraft in the fleet is either flying or recharging.

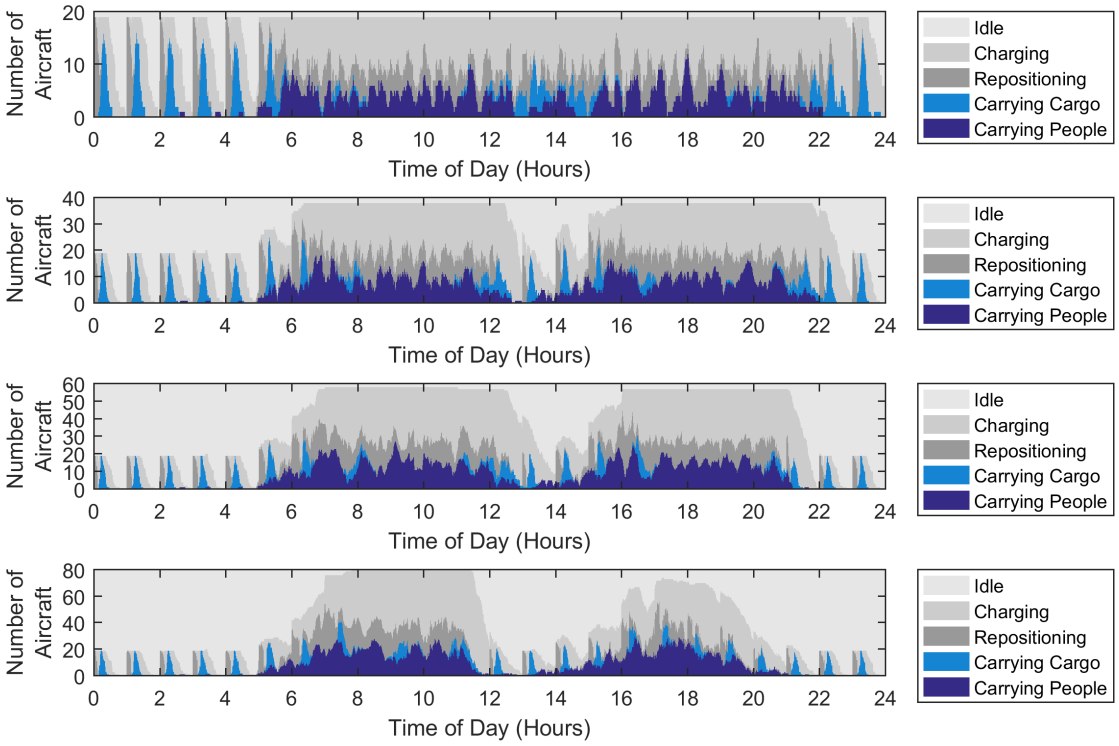

Figure 4.11: Simulated time history for mixed passenger/cargo network with fleet size = 20; 40; 60; 80 ships. Cargo demand originating from Amazon's "OAK5" distribution center in Newark, CA (Note different y-axis scaling.)

Figure 4.12 shows the results when the cargo demand is moved farther from the population centers to the OAK4 distribution center in Tracy, CA. The primary difference between this case and the above case is that flights to and from the Amazon distribution center are now generally longer. Surprisingly, at larger fleet sizes, this case actually outperforms the OAK5 case in terms of fewer numbers of failed trips – at a fleet size of 80 this case has zero failed trips for the day. This behavior seems to be a result of the rudimentary algorithm we use to reposition flights to serve the cargo demand, i.e. the flights are dispatched from as nearby the distribution center as possible. In this case, those nearby vertiports are less useful for commuting, and so taking idle aircraft from these locations to serve cargo has less impact on the performance of the passenger network. With improved dispatching logic, similar results should be achievable in the OAK5 case.

4.3.4 Demand Served

Table 4.1 shows the counts of trips that failed to be served for each simulation. Assessing these results suggests that for a passenger-only network, at least 60 aircraft are necessary to make the service acceptably reliable (failing to serve approximately 10% of trips). When cargo demand is added to the network, at least 80 aircraft are needed for good performance in the simulation.

The results show, however, that at fleet sizes of 60 or 80 aircraft, most of the fleet is idle at any given time. This finding suggests that with improved dispatch logic or cargo demand scheduling, much better fleet utilization could be achieved.

4.3.5 Fleet Utilization

The average utilization of each ship in the fleet over the course of the day is shown in Figure 4.13. Each bar segment represents the fraction of the day that the average ship in the fleet spends in the corresponding state.

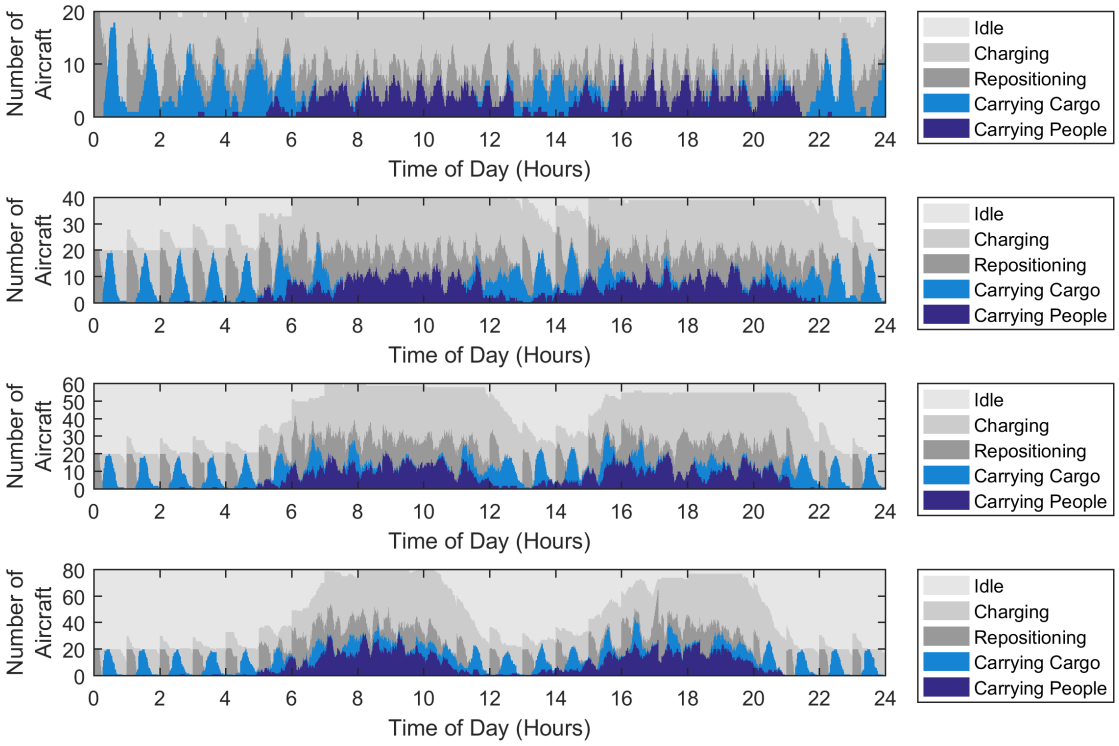

Figure 4.12: Simulated time history for mixed passenger/cargo network with fleet size = 20; 40; 60; 80 ships. Cargo demand originating from Amazon's "OAK4" distribution center in Tracy, CA (Note different y-axis scaling.)

Table 4.1: Count of demanded trips not served for each simulation

Fleet Size/Cargo	Passenger Not Served (980 available)	Cargo Not Served (480 available)
20/none	597	-
40/none	239	-
60/none	100	-
80/none	0	-
20/OAK5	667	239
40/OAK5	309	125
60/OAK5	223	112
80/OAK5	56	18
20/OAK4	685	277
40/OAK4	352	148
60/OAK4	157	76
80/OAK4	0	0

Depending on how utilization is defined, different measures of effectiveness can be derived from Figure 4.13. If utilization is defined as the fraction of time spent carrying passengers or cargo, it varies between about 10% and 25% depending on the scenario. Scenarios with more aircraft have lower per-aircraft utilizations, but, as shown in Table 4.1, are also able to serve more customers.

Adding cargo flights increases utilization significantly. In the cases for which cargo originates at OAK4,

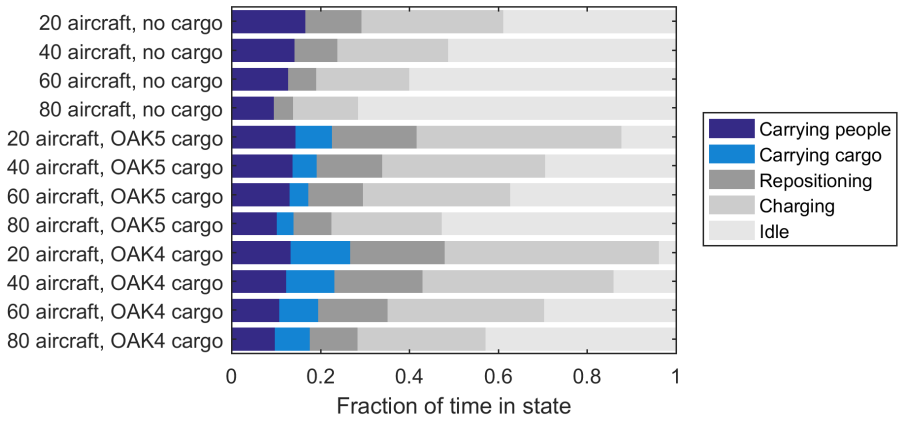

Figure 4.13: Average utilization for each ship in the fleet, for each scenario

the amount of time spent carrying cargo is almost equal to the time spent carrying passengers without significantly reducing the number of passengers served compared to the no-cargo cases.

It could be argued that repositioning flights, or some fraction of them, should also be included in utilization, since these are necessary to serve customers and could potentially be turned into revenue flights.

The need to recharge between flights creates a strict upper limit on utilization. Figure 4.13 shows that this limit is around 50%, as aircraft spend about 1 minute charging for each minute of flight.

Another perspective on utilization is to examine the cumulative number of flights in a day for the entire fleet up to a given time of day, as shown in Figure 4.14.

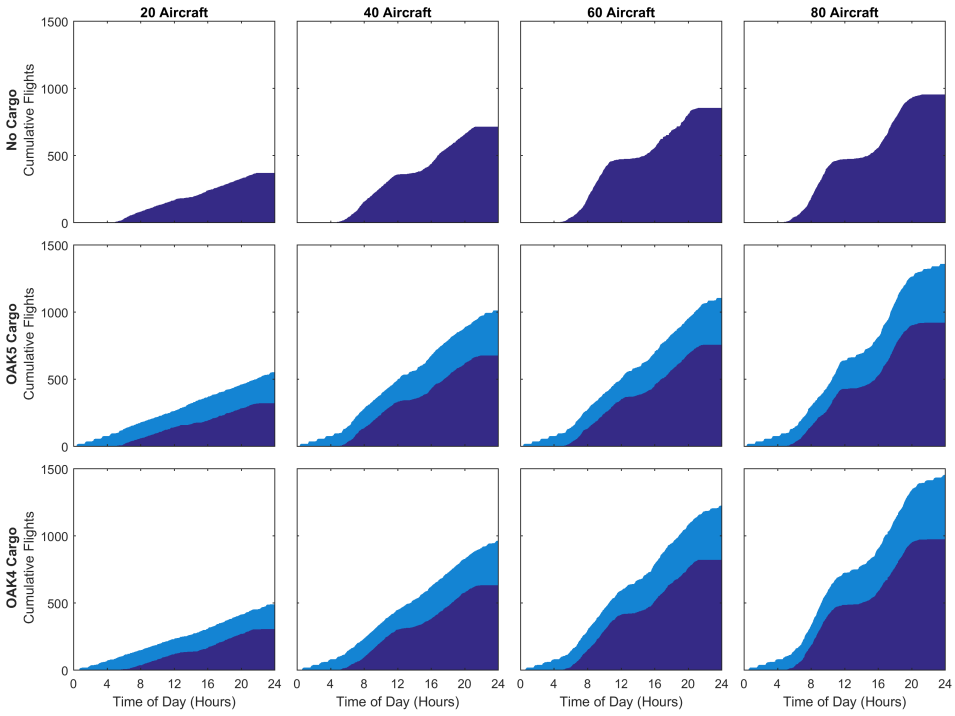

Figure 4.14: Cumulative number of passenger and cargo flights for the fleet at a particular time of day. Dark blue indicates passenger flights, and light blue indicates cargo flights.

5. Conclusions

In this research project, we have investigated potential concepts of operation for using passenger-class eVTOL UAM aircraft to carry cargo. The primary contributions of our work have been:

- Envisioning different canonical types of concepts of operation that take advantage of the potential of eVTOL at reducing cargo delivery time
- Investigating two Amazon package delivery CONOPS in detail in terms of package delivery throughput and aircraft utilization
- Developing an optimization methodology for deteremining promising vertiport locations for UAM (in the coarse-grained sense of selecting which census tracts should have vertiports)
- Developing an approach to calibrate travel times implied from LODES home/work locations based on self-reported ACS commuting times
- Developing eVTOL aircraft performance models suitable for studying energy requirements for eVTOL operations
- Developing of a UAM network operations simulation to track aircraft location and energy state, demand served, and vertiport occupancy throughout a "day in the life" of a UAM fleet
- Writing two AIAA conference papers (at SciTech 2018 and Aviation 2018) documenting our results

Future research is needed in the area of eVTOL cargo delivery to better understand its potential. One area of importance is determining better ways to predict cargo demand. Demand prediction was a primary challenge that we faced in this research because of the lack of public domain data for cargo shipments of time-critical products in metropolitan areas. Lacking adequate credible demand data, we were primarily restricted to estimating the potential for cargo throughput in UAM networks. Although we could estimate aircraft operating costs (via our collaborator Virginia Tech's eVTOL cost estimation tool), the lack of demand data made predictions of the revenue potential of eVTOL cargo, and particularly of the induced demand aspect, largely unobtainable. If new demand data can be obtained, then additional concepts of operations may be worth exploring based on insights from the data.

Another research area that would be fruitful would be to improve the rules for our Monte Carlo UAM network simulation model. In particular, there is a need for improved optimization-based aircraft dispatch logic to reduce repositioning flights and to better serve more demand more quickly. There is also a need for improvements in how demand is generated probabilistically, e.g. whether passengers or packages appear in groups or individually, and how this demand in served temporally, i.e. how long do aircraft wait in anticipation of more passengers or cargo arriving.

Lastly, there is a need for additional work that explores the impact of vertiport capacity in UAM operations for both cargo and passengers. In our studies in the present project, we typically presumed unlimited vertiport capacity because finite capacity substantially complicates the aircraft dispatch logic. However, it is likely that vertiport capacity will prove to be a significant constraint on the tempo and scale of operations, as a vertiport in a popular area with inadequate capacity can effectively throttle the entire network operation.

References

[1] German, B. J., Daskilewicz, M. J., Hamilton, T. K., and Warren, M. M., "Cargo Delivery by Passenger eVTOL Aircraft: A Case Study in the San Francisco Bay Area," *AIAA SciTech*, 2018.

[2] "Amazon Prime Same Day," Tech. rep., https://www.amazon.com/gp/help/customer/display.html?nodeId=201910880], accessed December 3, 2017.

[3] Tech. rep., https://www.scrapehero.com/how-many-products-are-sold-on-amazon-com-january-2017-report/, accessed December 3, 2017.

[4] Tech. rep., https://www.scrapehero.com/how-many-products-does-amazon-prime-now-sell-september-2017-update/, accessed December 3, 2017.

[5] Ruggles, S., Genadek, K., Goeken, R., Grover, J., and Sobek, M., "Integrated Public Use Microdata Series: Version 7.0 [dataset]." Tech. rep., Minneapolis: University of Minnesota, 2017, https://doi.org/10.18128/D010.V7.0.

[6] Duffy, M. J., Wakayama, S., Hupp, R., Lacy, R., and Stauffer, M., "A Study in Reducing the Cost of Vertical Flight with Electric Propulsion," *AIAA Aviation Forum, 5-9 June 2017, Denver, Colorado.*, 2017.

[7] "Aurora Unveils New eVTOL Aircraft at Uber Elevate Summit," Tech. rep., http://www.aurora.aero/wp-content/uploads/2017/04/APR-342_Aurora-Uber-Partnership-Announcement.pdf, accessed December 3, 2017.

[8] "Joby S2," Tech. rep., http://www.jobyaviation.com/S2/, accessed December 3, 2017.

[9] Shepherd, C. M., "Design of Primary and Secondary CellsPart 2: An Equation Describing Battery Discharge," *Journal of the Electrochemical Society*, Vol. 112, 1965, pp. 657–664.

[10] McDonald, R. and German, B., "eVTOL Energy Needs for Uber Elevate," *Uber Elevate Summit April 25-26, 2017, Dallas,TX*, 2017.

[11] Tech. rep., http://www.samsungsdi.com/lithium-ion-battery/power-devices/power-tool.html, accessed December 3, 2017.

[12] "Fuel requirements for flight in VFR conditions," Tech. rep., https://www.law.cornell.edu/cfr/text/14/91.151, accessed December 3, 2017.

[13] GAMA, "Gama Publication No. 16: Hybrid & Electric Propulsion Performance Measurement, draft document," Tech. rep., 2017.

[14] "Express Plus Specifications for Power Cube, Power Modules and Station," Tech. rep., https://www.chargepoint.com/files/datasheets/ds-expressplus.pdf, accessed December 3, 2017.

[15] Thomsen, M., "Perspectives on Prospective Markets Panel," *5th Annual AHS Transformative VTOL Workshop*, Airbus, 2018.

[16] German, B. J., Daskilewicz, M. J., Warren, M. M., Garrow, L. A., Boddupalli, S.-S., and Douthat, T. H., "Progress in Vertiport Placement and Estimating Aircraft RangeRequirements for eVTOL Daily Commuting," *AIAA Aviation*, 2018.

[17] U.S. Census Bureau, "American Community Survey B08119 MEANS OF TRANSPORTATION TO WORK BY WORKERS' EARNINGS IN THE PAST 12 MONTHS (IN 2016 INFLATION-ADJUSTED DOLLARS)," 2016 5-year estimates, Retrieved from: https://factfinder.census.gov/.

[18] U.S. Census Bureau, "American Community Survey B08134 MEANS OF TRANSPORTATION TO WORK BY TRAVEL TIME TO WORK," 2016 5-year estimates, Retrieved from: https://factfinder.census.gov/.

[19] U.S. Census Bureau, "LEHD Origin-Destination Employment Statistics (LODES)," 2015, Retrieved from: https://lehd.ces.census.gov/data/lodes/.

[20] U.S. Census Bureau, "TIGER/Line Shapefiles," 2017, Retrieved from ftp://ftp2.census.gov/geo/tiger/TIGER2017/.