

# SimUAM: A Comprehensive Microsimulation Toolchain to Evaluate the Impact of Urban Air Mobility in Metropolitan Areas

Pavan Yedavalli\*, Emin Burak Onat<sup>†</sup>, Xin Peng<sup>‡</sup>, Raja Sengupta<sup>§</sup>, Paul Waddell<sup>¶</sup>  
*University of California, Berkeley, Berkeley, CA, 94720*

Vishwanath Bulusu<sup>||</sup>  
*Crown Consulting Inc. at NASA Ames Research Center, Moffett Field, CA, 94035*

Min Xue\*\*  
*NASA Ames Research Center, Moffett Field, CA 94035*

Over the past several years, Urban Air Mobility (UAM) has galvanized enthusiasm from investors and researchers, marrying expertise in aircraft design, transportation, logistics, artificial intelligence, battery chemistry, and broader policymaking. However, two significant questions remain unexplored: (1) What is the value of UAM in a region’s transportation network? and (2) How can UAM be effectively deployed to realize and maximize this value to all stakeholders, including riders and local economies? To adequately understand the value proposition of UAM for metropolitan areas, the authors develop a holistic multi-modal toolchain, SimUAM, to model and simulate UAM and its impacts on travel behavior. This toolchain has several components: (1) Microsimulation Analysis for Network Traffic Assignment (MANTA): A fast, high-fidelity regional-scale traffic microsimulator, (2) VertiSim: A granular, discrete-event vertiport and pedestrian simulator, (3) Flexible Engine for Fast-time Evaluation of Flight Environments (Fe<sup>3</sup>): A high-fidelity, trajectory-based aerial microsimulation. SimUAM, rooted in granular, GPU-based microsimulation, models millions of trips and their movements in the street network and in the air, producing interpretable and actionable performance metrics for UAM designs and deployments. Once the ground-air interface is modeled, the authors find that the market for UAM decreases across all network designs relative to models with static assumptions about transfer times. However, significant improvements can be made to balance the demand and optimize the networks for transfer time, likely increasing the number of benefited trips. The modularity, extensibility, and speed of the platform will allow for rapid scenario planning and sensitivity analysis, effectively acting as a detailed performance assessment tool.

## I. Introduction

When a new mode is introduced into a metropolitan transportation network, its impact on individual travel behavior is often not readily understood [1]. To proponents, Urban Air Mobility (UAM) directly tackles one of the known parameters of mode choice, travel time, a theoretically and empirically known ‘bad,’ or a factor that increases disutility [2, 3]. Existing analyses have attempted to estimate the potential impacts of UAM passenger services on travel behavior, using predetermined vertiport locations and simulating different modes of travel in different metropolitan areas [4? –6]. However, these existing models assume static parameters in either ground traffic, vertiport transfer times, aerial flight, or a combination of the three. For UAM, since its purported value is in travel time savings, simulating the margins at a granular level, incorporating dynamics and scheduling in each leg of the trip, is of utmost importance.

---

\*Ph.D Candidate, Department of City and Regional Planning/Civil and Environmental Engineering, 230 Wurster Hall, Berkeley, CA 94720

<sup>†</sup>Ph.D Student, Department of Civil and Environmental Engineering, 100 McLaughlin Hall, Berkeley, CA 94720


<sup>‡</sup>Ph.D Candidate, Department of Civil and Environmental Engineering, 100 McLaughlin Hall, Berkeley, CA 94720

<sup>§</sup>Professor, Department of Civil and Environmental Engineering, 100 McLaughlin Hall

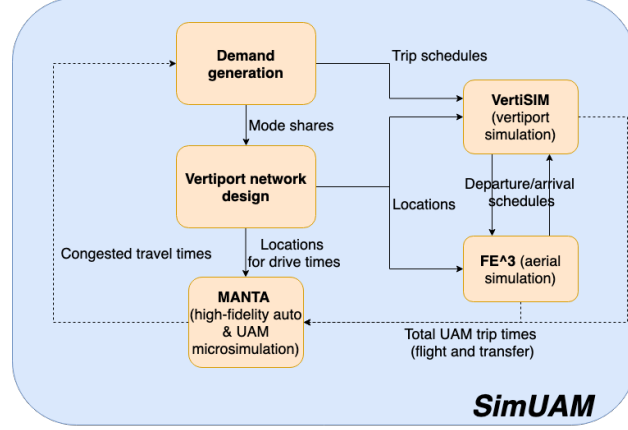
<sup>¶</sup>Professor, Department of City and Regional Planning/Civil and Environmental Engineering, 230 Wurster Hall, Berkeley, CA 94720

<sup>||</sup> Aerospace Research Scientist, Crown Consulting Inc., Aviation Systems Division, NASA Ames Research Center, Moffett Field, CA 94035, USA

\*\* Aerospace Research Engineer, Aviation System Division, NASA Ames Research Center, Moffett Field, CA 94035. AIAA senior member



**Fig. 1 Simulation components for an individual multi-modal UAM trip**




**Fig. 2 The full SimUAM model pipeline**

This paper proposes SimUAM, a holistic multi-modal toolchain that integrates ground traffic microsimulation, ground-air interfacing at the vertiport, and aerial microsimulation, producing accurate travel times for both multi-modal UAM trips and other driving trips across the region. Specifically, the MANTA component of SimUAM models road congestion patterns of all driving trips, including the access and egress legs of the multi-modal UAM trips [7]. VertiSim models the ground-air interface at the vertiport to move the traveler from the ground to a specific aircraft for the multi-modal UAM trips. The  $Fe^3$  component models the flight trajectories of these aircraft, incorporating conflict resolution and UAV mission profiles in the sky [8]. For each multi-modal UAM trip, the *interaction* among the three legs is important in realizing the true value of UAM in a regional transportation network. As a result, SimUAM is the first holistic toolchain, to the best of these authors' knowledge, that models all these interactions together dynamically and granularly, as shown in Fig. 1.

This paper is organized into the following sections. Section II details MANTA, the graphical processing unit (GPU)-based regional-scale road traffic microsimulator. Section III explains VertiSim, the granular, discrete-event vertiport and pedestrian simulator. Section IV discusses  $Fe^3$ , a NASA-developed, GPU-based high-fidelity regional-scale microsimulator for air traffic. These three components comprise the SimUAM architecture, shown in Fig. 2. Section V evaluates various scenarios with the toolchain to show its benefit. Section VI describes the performance metrics measured by SimUAM and experiments showing SimUAM's benefit. Finally, Section VII and Section VIII highlight the limitations, future work, and conclusions of this research.

## II. MANTA: Microsimulation Analysis for Network Traffic Assignment

The SimUAM framework relies on ground traffic microsimulation to produce congested travel times for all the travel demand in the investigated region. To do so, Microsimulation Analysis for Network Traffic Assignment (MANTA), an ultra-fast, highly-parallelized GPU-based microsimulation platform, is developed [7]. Existing simulators have typically revealed a tradeoff among accuracy, computational speed, and geographic scale of simulation [7, 9–12]. However, MANTA exhibits performance benefits in all three of these areas, enabling it to be used for agile scenario planning, particularly with an emergent mode such as UAM, whose deployment is still in its inchoate stages. MANTA's ability to track the microscopic movements of vehicles on the street network, namely lane changes, acceleration, and braking with



**Fig. 3 The San Francisco Bay Area region, defined by the polygonal hull of its nine counties**

respect to other vehicles, enables a granularity that provides important insights in travel time computations [7].

#### **A. Inputs**

MANTA takes two inputs: (1) Geographic network, comprised of edges and nodes in a metropolitan geography’s street network, and (2) Origin-destination (OD) demand showing traveler movements in the selected geography.

The geographic network, as described in [7], is derived from OpenStreetMap. The network library OSMnx is used to determine the nodes that exist within the polygonal hull of a metropolitan region, determined from widely available shapefiles for the region of investigation [13]. The San Francisco (SF) Bay Area is shown in Fig. 3 as an example output of this process, and is also used as the area of investigation in Section VI of this paper.

The OD demand is generated from the local metropolitan planning organization (MPO), which typically calibrates a travel demand model based on estimates from census and household travel data [14]. For the SF Bay Area, the data are derived from the Bay Area Metropolitan Transportation Commission (MTC)’s Travel Demand One model [14].

Once each trip is assigned a mode, these trips are filtered to only driving trips, which include driving alone, shared trips, and driving to transit. In the SF Bay Area, this totals 25 million trips for a full day. This work focuses on the morning commute, from 5am-12pm, and is filtered down to approximately 3M trips. Importantly, these trips may include multiple passengers in the car, but these extra passengers are not considered in the number of trips.

The origin and destination of each trip is at a traffic analysis zone (TAZ) granularity, which is a population density-based geographic unit used by MPOs across the nation, with each TAZ area typically larger than a block but smaller than a zip code. This size is too coarse for MANTA, which operates at the node level. As a result, once the origin and destination TAZs are known, nodes are randomly assigned to the origin and destination within their respective TAZs, sampled from a uniform distribution, producing each trip as one from an origin node to a destination node.

#### **B. Functionality**

In MANTA, there are two major components: (1) routing and (2) microsimulation. The initial routing algorithm developed in MANTA used a parallelized Dijkstra priority queue implementation, which greedily selects the closest vertex that has not yet been processed [15]. This single source shortest path (SSSP) implementation could compute 3M routes, the number of driving trips from 5am to 12pm in the SF Bay Area, in approximately 1 hour.

However, an update was made to integrate an ultra-fast, parallelized Open Source Routing Machine framework using Pandana, based on the contraction hierarchy scheme [7]. The time it takes to compute the same 3M routes in the

SF Bay Area is now under one minute.

MANTA is a time-based microsimulator that accounts for parallel changes that occur due to the dependence of cars' behaviors on one another. The vehicular movement on an edge is dictated by conventional car following, lane changing, and gap acceptance algorithms [7, 16]. The well-known Intelligent Driver Model (IDM), as shown in Eq. 1, is used to control the vehicle dynamics through the network [17].

$$\dot{v} = a \left( 1 - \left( \frac{v}{v_o} \right)^\delta - \left( \frac{s_o + Tv + \frac{v\Delta v}{2\sqrt{ab}}}{s} \right)^2 \right) \quad (1)$$

where  $\dot{v}$  is the current acceleration of the vehicle,  $a$  is the maximum possible acceleration of the vehicle,  $v$  is the current speed of the vehicle,  $v_o$  is the speed limit of the edge,  $\delta$  is the acceleration exponent,  $s$  is the gap between the vehicle and the leading vehicle,  $s_o$  is the minimum spacing allowed between vehicles when they are at a standstill,  $T$  is the desired time headway, and  $b$  is the braking deceleration of the vehicle [7, 18]. The position of each vehicle at the current timestep is computed using this calculated acceleration value  $\dot{v}$ .

In addition to car following, vehicles also follow lane changing rules, based on whether the vehicle is making a mandatory or discretionary lane change. Mandatory lane changes occur when the vehicle must take an exit from the edge, while discretionary lane changes occur during overtaking a slower vehicle [7]. The lane changing model states that the vehicle has an exponential probability of switching from a discretionary lane change to a mandatory lane change as it proceeds through the edge, as shown in Eq. 2.

$$m_i = \begin{cases} e^{-(x_i - x_0)^2} & x_i \geq x_0 \\ 1 & x_i \leq x_0 \end{cases} \quad (2)$$

where  $m_i$  is the probability of a mandatory lane change for vehicle  $i$ ,  $x_i$  is the distance of vehicle  $i$  to an exit or intersection, and  $x_0$  is the distance of a critical location to the exit or intersection [19, 20].

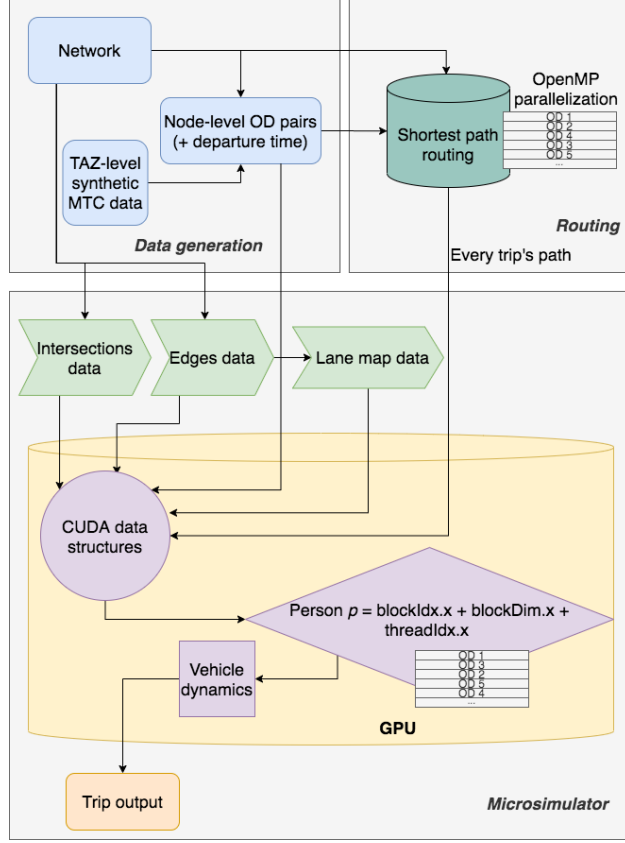
Once a vehicle decides to change lanes, the maneuver is performed if the lead and lag gaps of the cars in the lane to which it is changing are acceptable. The critical lead or lag gap for a successful lane change is defined as the minimum distance to the following or lagging vehicle that allows for a lane change, as shown in Eq. 3.

$$g_{lead} = \max(g_a, g_a + \alpha_{a1}v_i + \alpha_{a2}(v_i - v_a)) + \epsilon_a \quad (3)$$

$$g_{lag} = \max(g_b, g_b + \alpha_{b1}v_i + \alpha_{b2}(v_i - v_b)) + \epsilon_b \quad (4)$$

where  $g_{lead}$  is the critical lead gap for a lane change,  $g_{lag}$  is the critical lag gap for a lane change,  $g_a$  is the desired lead gap for a lane change,  $g_b$  is the desired lag gap for a lane change,  $\alpha$  is a system parameter (typically [0.05, 0.40]) that controls the gap based on speed,  $v_i$  is the speed of the vehicle,  $v_a$  is the speed of the lead vehicle,  $v_b$  is the speed of the lag vehicle, and  $\epsilon_a$  and  $\epsilon_b$  are the random components [7].


MANTA is a time-based simulator, rather than an event-based simulator, and thus accounts for changes that occur for all vehicles in every timestep. A time-based approach makes parallelization easier to implement, as synchronization protocols for event-driven simulators is challenging. In addition, no overhead exists with time-based simulators, which contrasts with event-based simulators, where millions to billions of events must be constantly generated. Without this overhead, this allows for the ability to model more granular movements in edges, not just events such as a vehicle entering or exiting an edge, which often constitutes an event. The MANTA architecture is shown in Fig 4.



**Fig. 4 The MANTA architecture**

### C. Calibration and Validation

MANTA was calibrated against Uber Technologies, Inc. (Uber) Movement data on a per-edge basis for the San Francisco Bay Area. First, 17% of edges in the San Francisco Bay Area network were matched with the Uber Movement data [7]. Then, a mini-batch gradient descent algorithm was used to iteratively converge to the optimal  $a$ ,  $b$ ,  $T$ , and  $s_0$  parameters of the Intelligent Driver Model, shown in Eq. 1. In addition, each driver was outfitted with a different driver profile per edge, based on a Gaussian distribution, with the speed limit of the edge as the mean, and the standard deviation of the Uber speeds on the edge as the standard deviation. This produced closely matched edge speeds for the 17% of edges, as shown in Fig. 5 [7]. Average travel times and speeds were then closely validated with California Household Travel Survey (CHTS) data as well as Uber Movement data [7]. As an example, Fig. 5 shows the mean edge speeds across all edges whose speed limits are 45 mph, for both MANTA and Uber Movement. The average speeds differ by 3 mph and the standard deviations by below 1 mph, suggesting successful calibration.



**Fig. 5 Kernel density estimator plot comparing MANTA microsim edge speeds vs. Uber edge speeds on edges whose speed limits are 45 mph**

From a performance perspective, MANTA can run one full simulation with 3M trips, including routing, in 4.6 minutes on a machine with Intel i9-7940X CPU, 3.10 GHz clock frequency and 28 cores. The GPU that runs the microsimulation component is an NVIDIA GP104 (GeForce GTX 1080) with a 33 MHz clock and 2560 CUDA cores.

MANTA provides the simulation backbone to the entire SimUAM platform, as all trips, both driving and multi-modal UAM, must originate and end through MANTA. More information can be found in [7]. However, as shown in the next section, multi-modal UAM trips must be modeled by multiple interacting components, one of which is the ground-air interface modeled by VertiSim.

### III. VertiSim: Vertiport Ground-Air Interface Modeling

Aircraft need infrastructure to land, taxi-in, park, load and unload passenger/cargo, charge, repair, taxi-out, and take-off. Vertiports have to provide this infrastructure to handle passengers and aircraft within limited land use. Thus, they have limited throughput capacity. This limitation makes vertiports the bottlenecks of the UAM system [21].

Vertiports are the transfer centers for passengers to change from one mode to another. Therefore, transfers at the origin and destination vertiports must be simulated in order to calculate their transfer time. Additionally, aircraft turnaround time and taxiway conflicts must be simulated for throughput calculation.

To understand passenger waiting times and aircraft throughput, VertiSim has been developed. VertiSim is an agent-based discrete-event simulator capable of simulating millions of passenger and aircraft agents.

#### A. Inputs

VertiSim takes vertiport locations, surface layouts (topologies), passenger and aircraft parameters as inputs. These are comprised of touchdown and liftoff (TLOF) pads, parking spaces, number of charging stations, number of security lines at the vertiport entrance, charging time distribution for aircraft, aircraft ground speed, distributions for stochastic walking/boarding/deboarding times for passengers, passenger arrival schedule with passengers' destination vertiport, aircraft arrival schedule. It also takes several inputs for vertiport surface management strategies such as aircraft separation and taxiway blocking during take-off and landing. The parameters of VertiSim used in this study, such as average human walk speed and aircraft ground speed, are determined from the literature and by subject-matter experts [22, 23].

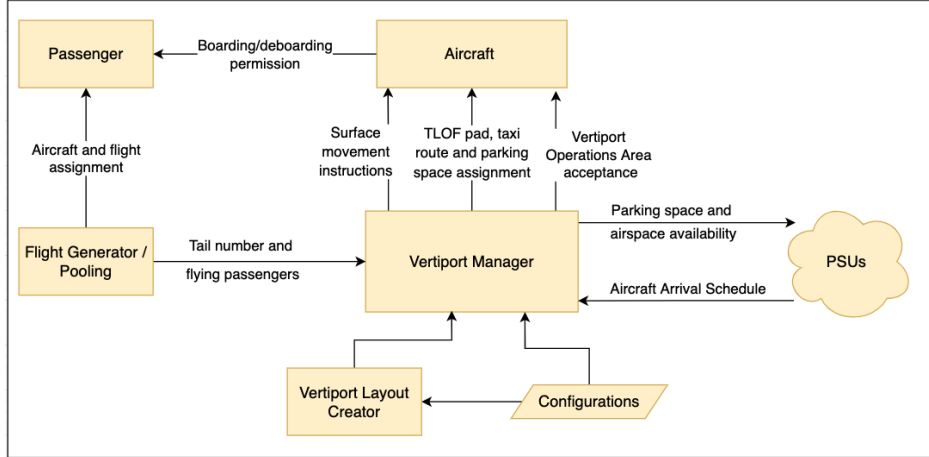
#### B. Functionality

The three major components of VertiSim are the (1) vertiport layout designer, (2) flight generator, and (3) vertiport manager. The responsibility of the vertiport layout designer is to create the node-link model of the topology for the input file. The flight generator then generates flights, as the number of passengers accumulate to fly to the same direction, and assigns passengers to available aircraft. If there is no available aircraft, the generator puts them in the departure queue. It also generates flights for the passengers who waited more than  $T_{max}$ , the maximum waiting time threshold, specified as 10 minutes, without waiting for the aircraft to be full. When there is a flight queue for the departing passengers in the

waiting room, they wait until their turn. The flight generator can assign new passengers to these underbooked flights if passengers who are flying to the same destination arrive to the waiting room before passengers leave the waiting room. Vertisim assumes all flights are direct flights. The vertiport manager oversees the usage of the resources of the vertiport for the aircraft agent. That is, it assigns a parking space and TLOF pad for the aircraft and creates the taxi route for the aircraft. Ultimately, VertiSim is able to simulate the departing passenger processes, arriving passenger processes and the aircraft arrival and departure processes at a vertiport.

VertiSim is capable of simulating a variety of vertiport topologies. Every vertiport in the given vertiport network can have a unique topology. Departing passenger flow in VertiSim is as follows: (1) car to vertiport entrance, (2) security check, (3) security check to waiting room, (4) waiting room to boarding gate, (5) boarding gate to aircraft, (6) taxi-out. Arriving passengers then disembark, go to the exit gate at the vertiport surface, exit the vertiport at the ground level, and enter a car. From the perspective of the aircraft, arriving aircraft request an available parking space. In the rare event that there are no available parking spaces, they loiter until one is reserved. They then land at the TLOF pad assigned by the vertiport manager, taxi-in, unload the passengers, and finally charge. These aircraft then become available for the departing passengers. When a parked aircraft is assigned to a flight, they load passengers, reserve an available TLOF pad, taxi-out, and take-off.

VertiSim outputs aircraft departure schedules, aircraft turnaround time statistics, passenger transfer time statistics, and vertiport efficiency metrics. The VertiSim architecture is shown in Fig. 6. According to FAA's UAM ConOps [24], Provider of Services to UAM (PSUs) are responsible for supporting the UAM operations by ensuring that the UAM operations meet the requirements of safe, efficient and secure use of the airspace. In VertiSim, the user takes the role of the PSUs and provides the aircraft arrival schedule. In return, the user receives parking space and airspace availability as vertiport efficiency metrics. Vertiport Manager processes each aircraft arrival and provides instructions to the aircraft. These instructions include the location of the TLOF pad, taxi routing, parking, charging, debarking and boarding. Flight Generator provides tail number and flying passenger IDs to the Vertiport Manager, which then assigns an aircraft for that flight.



**Fig. 6 The ground/air interface at the vertiport**

### 1. Steady-state data collection

The first passengers on the passenger arrival schedule have priority to fly when the minimum number of passengers required to fly to the same destination has been met, or if they have waited more than  $T_{max}$ . Therefore, the first passengers might wait less than the subsequent passengers because there will not be any other passengers on the departure queue. In order to replicate the reality, the data from the transient-state of the simulation is removed. Here we define the state as the probability of the number of passengers at the waiting room at any given time. Since the state of the queueing system will depend on time at the initial part, the system will be in transient-state. Although there is no general guideline to identify transient-state from steady-state, several methods exist. One option is to choose a large enough simulation time to eliminate the effects of the transient state [25]. A formulation to detect the end of the transient-state region based on our observations is derived. Once VertiSim simulates all the passenger agents, we want to simulate them in MANTA



and  $Fe^3$  as well to calculate their total travel time. Therefore, VertiSim needs to simulate the whole passenger arrival process without any loss of passenger agents. To that end, we add pseudo-passengers, as expressed in Eq. 5 that are randomly sampled from the original demand file and remove them at the end of the simulation.

$$N_A = N_S * N_{TLOF} * 2 \quad (5)$$

where  $N_A$  represents number of pseudo-passengers need to be added on top of the original passenger arrival process of each vertiport,  $N_S$  is aircraft passenger capacity, and  $N_{TLOF}$  represents number of TLOF pads at the vertiport. For instance, on a 4 seat aircraft and 4 TLOF pad scenario, VertiSim creates 32 passengers before the actual passengers arrive, which means that there will be at least 8 take-offs before the actual passengers fly.

Since there are no other vertiport simulators with similar capabilities, and eVTOLs are not practically used, we could not make any calibration or verification of VertiSim. However, the input parameters are determined from the literature and by subject-matter experts.

#### IV. $Fe^3$ : High-fidelity Aerial Traffic Simulator

The aerial component of the trip is simulated using the  $Fe^3$  GPU-based microsimulation platform, developed at NASA. A mission profile is created for every UAM trip, with a departure time, ascent and descent characteristics, cruise altitude, and velocity/directional changes depending on wind patterns and collision avoidance [8]. The  $Fe^3$  platform models these in-flight dynamics and any subsequent flight trajectory shifts that contribute to changes in total flight time, making it ideal for simulating UAM in congested airspace [8].  $Fe^3$  is a dynamic, high-fidelity microsimulator leveraging a parallelized GPU architecture, similar to MANTA. Due to the similarities in fidelity, the loose coupling of  $Fe^3$  and MANTA is both seamless and interpretable.

##### A. Inputs

$Fe^3$  takes several structures as input. As described earlier, MANTA outputs the arrival times of the multi-modal UAM trips to their respective origin vertiports, and VertiSim consolidates these passengers into aircraft based on pedestrian walking times, aircraft arrival times, and vertiport capacities. Once these passengers are assigned aircraft and then board, the aircraft must go through the takeoff and landing process before it departs. As a result,  $Fe^3$  takes as input the aircraft departure times output from VertiSim. The passenger IDs are abstracted away from  $Fe^3$ , which proceeds based on the assumption that all aircraft contain at least 1 passenger. In addition to the flight plans with departure times,  $Fe^3$  also takes vehicle types, rules and parameters for flight scheduling, rules and parameters for both strategic and tactical conflict resolution, and weather information as shown in Fig. 7.

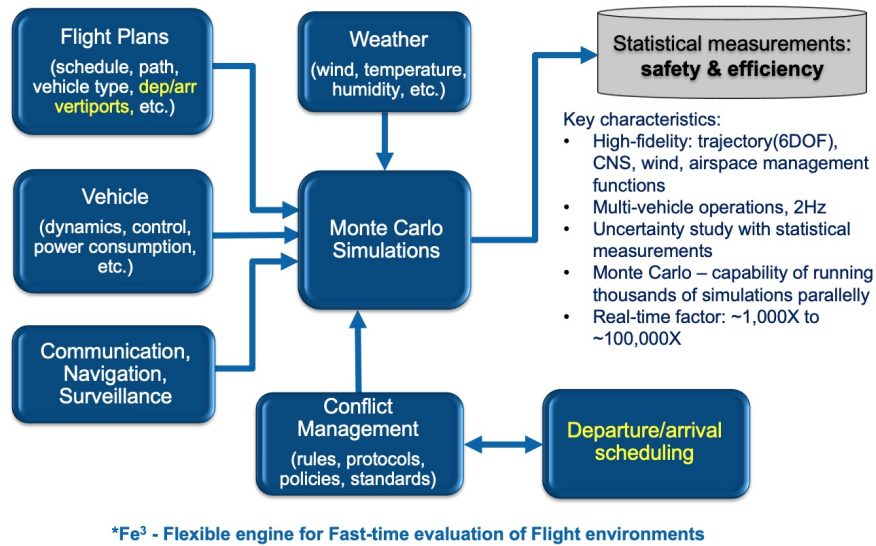


Fig. 7  $Fe^3$  aerial microsimulator architecture



## B. Functionality

Fe<sup>3</sup> uses the CUDA programming language on graphics processing units (GPUs). It is deployed on the Amazon Web Service (AWS) cloud for scalability, such that the number of GPU instances can be dynamically deployed based on simulation needs. Within the simulator, each aircraft's flight in Fe<sup>3</sup> is governed by several models, specifically trajectory, scheduling, conflict resolution, vehicle communication and sensors, and wind models [8, 26, 27].

Fe<sup>3</sup> includes two types of trajectory models, 6 degree-of-freedom (DOF) models where both forces and moments are considered and kinetic models where moments are excluded. For multi-copters used in this work, a kinetic point mass model developed based on NASA's UAM concept vehicle [28] was used, in which proportional and proportional-derivative controllers are used to track desired airspeed, horizontal path, and vertical profile [29].

Both scheduling and conflict resolution in Fe<sup>3</sup> are configurable as described in previous UAM studies [26, 27]. Both take the rules and parameters from the user specified input files and apply through the simulations. The scheduling function and conflict resolution can be turned on and off individually as well. As discussed in previous work [27], because the short duration of UAM operations and limited number of vertiports, the correlation between scheduling and conflict resolution is high. Since the bottleneck happens at the vertiports, scheduling UAM operations at both origin and destination vertiports are critical. Once UAM flights are spaced out at both origin and destination vertiports, the burden on conflict resolution is actually alleviated. When the number of operations increase, the delay at the vertiport is dominant and the airborne delay due to conflict resolution is negligible. Experiments in this work showed that the average airborne delay was usually less than a minute.

Although the communication and sensor models and wind model are implemented in Fe<sup>3</sup>, the communication and sensors were assumed to be a perfect, and calm wind was used to simplify the simulation because the focus of this study is not really on the performance and uncertainties of these models.

## C. Validation and Performance

The individual models in Fe<sup>3</sup> used in this work have been verified in previous works [8, 26, 27] including the trajectory model, scheduling function and conflict resolution function. The running time on an AWS GPU instance for a scenario with 8,000 flights is about 10 minutes, when both scheduling and conflict resolution functions are on and the simulation time step is half a second.

## V. Simulation

The SimUAM toolchain is now applied to the San Francisco (SF) Bay Area as a case study. The SF Bay Area is chosen as the case study due to wide travel data availability and massive geographic scale that reflects useful bidirectional commuting patterns. This toolchain combines network and origin-destination trip generation, ground traffic microsimulation, vertiport simulation, aerial traffic simulation, and UAM trip assignment and modeling. Once these steps are complete, the final step is to simulate both UAM and non-UAM trips [7, 8].

As highlighted in Fig. 1, the first step is a standard MANTA simulation without UAM as a mobility option, and with all trips as driving-only trips. These results are used for later comparison.

The next step involves running a MANTA simulation with UAM as a mobility option. In this scenario, there are several sequential steps. The first requires simulating in MANTA the access leg of each multi-modal UAM trip, with the rest of the trips in the demand assigned to driving.

After the toolchain runs, two experiments are presented, (1)  $S_1$ , Simulated road traffic in MANTA, and static vertiport transfer times and aerial flight times for the multi-modal UAM trips, and (2)  $D_1$ , Simulated road traffic in MANTA, simulated vertiport transfer times in VertiSim, and simulated aerial flight times in Fe<sup>3</sup>.

Once these initial times are determined, we carry out the static experiment,  $S_1$ . A transfer time of 2 minutes is assumed at the origin vertiport. Then, each traveler is assumed to have access to an aircraft after 2 minutes at the origin vertiport, and they embark on their flight. This flight also has static characteristics, specifically 138 mph (120 kts) cruise speed, 2000 feet cruise altitude, ascent/descent angles of 10 degrees, and 100 fpm vertical climb. The flight path assumes a Haversine distance between the origin and destination vertiports. After the flight lands at the destination vertiport, a transfer time of 2 minutes at the destination vertiport is used.

In the dynamic experiment,  $D_1$ , instead of assuming static transfer and aerial times, travelers are simulated within the vertiport with VertiSim and are consolidated into aircraft. These aircraft's trajectories are then simulated in Fe<sup>3</sup>. Once the flight lands at the destination vertiport, VertiSim, simulates the traveler's movement from the aircraft to their egress vehicle, as shown in Fig. 1.

In both experiments, once the traveler reaches the egress vehicle, MANTA takes control and simulates the full demand with all access and egress trips for UAM as well as all the non-UAM driving trips.

## VI. Performance Metrics

SimUAM's value is derived from its ability to produce a comprehensive set of performance metrics for all trips at both the microscopic person scale and the macroscopic regional scale. Once a network design with a specific number of vertiports is input into SimUAM, it is able to simulate all the demand in the region, assigning specific trips to multi-modal UAM and preserving the rest of the demand as driving trips.

In  $S_1$ , as mentioned earlier, the transfer times at the vertiports are statically assigned, and the aerial flights follow standard mission profile assumptions that depart on demand for each traveler. The number of trips that benefit from taking multi-modal UAM, which is defined as a trip whose travel time when taking multi-modal UAM is shorter than the travel time of the same trip if it was driving only, are then determined. Shown in Fig. 8, as the transfer times increase, the number of benefited trips decreases across all  $k$ , where  $k$  is the number of vertiports. The number of benefited trips also increases as the number of vertiports increases, with the largest increase occurring between 10 and 30 vertiports, with more marginal benefit occurring as the number of vertiports increases further, ranging between 10,000 to 250,000 trips. This serves as a useful upper bound of the total addressable market of urban air mobility.





Fig. 8 The number of benefited trips when the ground-air interface is not modeled

### A. Regional impact and addressable market


One of the major limitations of static assumptions is the inaccurate representation of aircraft being on-demand (i.e., being available to fly immediately when the traveler requires it). In experiment  $D_1$ , the ground-air interface at the vertiport, in which the wait times and consolidation times of passengers into aircraft, as well as the aircraft and vertiport capacities, is modeled. Using the number of benefited trips at a transfer time of 2 minutes, the same demand through the full SimUAM toolchain shows that the number of benefited trips at  $k = 10$ ,  $k = 30$ ,  $k = 50$ , and  $k = 70$  decrease significantly. In Fig. 11, at  $k = 10$ , the number of benefited trips using the dynamic model,  $D_1$ , is 5,300 trips, as opposed to 23,000 trips when using the static model,  $S_1$ . At  $k = 30$ , the disparity is more than an order of magnitude, with 11,000 trips using  $D_1$ , as opposed to 125,000 using  $S_1$ . At  $k = 50$ , the  $D_1$  model produces 18,000 trips, while the  $S_1$  model produced 171,000 trips at  $k = 50$ . Finally, at  $k = 70$ , using the  $D_1$  model resulted in 23,000 trips, contrasted with 212,000 trips produced from the  $S_1$  model. This drop in the number of benefited trips immediately shows the impact of real world constraints that will exist for UAM at the vertiports themselves. Specifically, the number of takeoff and landing areas (TOLAs), number of parking spaces for the aircraft, and the number of passengers per aircraft will dictate the efficiency of multi-modal UAM as a beneficial transportation mode for travelers.  $k = 100$  is not calculated due to memory constraints for  $Fe^3$ .

In Fig. 11, the number of benefited trips in the static  $S_1$  experiment shows a logarithmic growth over  $k$ , but once the ground-air interface is modeled in the loosely coupled SimUAM toolchain in  $D_1$ , the number of benefited trips increases relatively linearly across  $k$ , showing that the waiting times and consolidation times at the vertiports significantly decreases the potential of multi-modal UAM for networks of 30 vertiports or fewer.



**Fig. 10** The number of benefited trips once the ground-air interface is simulated in SimUAM significantly decreases across all  $k$

When considering the ground-air interface in  $D_1$ , the trips that benefit are those that have even higher travel times than when the ground-air interface was not considered in  $S_1$ . For instance, at  $k = 10$ , the average driving time of trips that benefit when seeing the ground-air interface statically in the  $S_1$  model is about 95 minutes, but in the  $D_1$  model that integrates the ground-air interface dynamically, this increases to 101 minutes. The multi-modal UAM times for these same trips also increase when considering the ground-air interface, going from 64 minutes to 77 minutes. This is unsurprising given the increase in transfer times due to the ground-air interface bottleneck. This dynamic across  $k$  is shown in Fig. 11.




**Fig. 11** The average travel time of benefited trips once the ground-air interface is simulated in SimUAM significantly increases across all  $k$  (red vs. orange)

## B. Vertiports

Per the existing literature, a throughput of 240 flights per hour per vertiport, including both departures and arrivals, is assumed [8]. At  $k = 10$ , the average transfer time once the ground-air interface is considered dramatically increases to 50 minutes, based on the mean of each vertiport's transfer time in Fig. 12. With a configuration of 4 TLOF pads, 16 parking spaces, and up to 4 seater aircraft, using the initial total demand of 23K at a transfer time of 2 minutes, the real transfer times vary considerably across vertiports. The maximum average waiting time occurs at Vertiport 6, whose traffic of over 4K passengers forces the average transfer time to be over 2 hours. At Vertiport 4, since the demand is below 1K, the average transfer time is only 3.85 minutes.


This shows that demand balancing may be important to decrease average waiting times at each vertiport. However, balancing the demand in a network with few vertiports may also increase the access and egress driving times of these



**Fig. 12** The average transfer time at  $k = 10$  once the ground-air interface is considered (Scenario  $D_1$ )

multi-modal UAM trips.

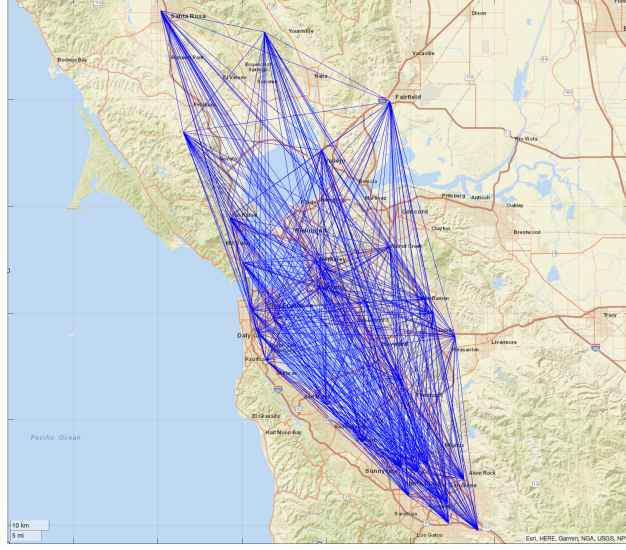
To investigate the effect of demand level on passenger transfer times, seven simulation scenarios, each with the vertiport configuration of 4 TLOF pads and 16 parking spaces per vertiport, are evaluated. Arrivals were modeled with a 30-sec inter-arrival time, meaning that each of the four TLOF pads per vertiport accepted an arrival every 2 min. Thus, the demand per vertiport decreases. Fig. 13 shows that as the demand decreases, the average transfer time decreases down to an inflection point. The decrease in the average transfer time is due to the increased passenger inter-arrival time, which results in shorter queues for the flights. Decreasing the demand further actually results in an increase in the average transfer time because the accumulation time for the required number of passengers for a flight (4 passengers) increases. Also, as the demand decreases, the average aircraft occupancy decreases, due to the ratio of the passengers that wait more than  $T_{max}$  (10 minutes) increases and the flight generator is triggered for those passengers. Ultimately, waiting times decrease as  $k$  increases, but at 30 vertiports, the consolidation time increases as well, due to the demand being distributed across all vertiports rather than at certain vertiports.



**Fig. 13** The average transfer time and average aircraft occupancy with different demand levels

### C. Airspace

After the aircraft are simulated in  $Fe^3$ , there are comparatively many aircraft congesting the airspace, with 7,352 flights for  $k = 30$  vertiports, as shown in Fig. 14. Note that these aircraft trajectories do not consider airspace restrictions, which will further affect the number of benefited trips if the airspace-restricted flight times are higher than the airspace-unrestricted flight times.  $Fe^3$  finds that with federated scheduling and no conflict resolution, the average delay is 21.5 seconds per flight. It is worth mentioning that when VertiSim generated flight schedules at origin vertiports, it already applied departure delay to separate departure flights. That explains the low average delay in the aerial simulation. As discussed before, airborne delay is negligible since the bottleneck exists at the origin and destination due to the limited number of vertiports.



**Fig. 14 The congested airspace of the San Francisco Bay Area at  $k = 30$  vertiports**

## VII. Limitations and Future Research

The SimUAM toolchain has one notable limitation: the inability to do iterative convergence across a loosely coupled pipeline [30]. Due to the complex data input and output of iterative convergence, an ongoing solution is to cross-compile both GPU-based microsimulators, MANTA and  $Fe^3$ , and convert VertiSim into a GPU-based simulator. By synchronizing the time, the convergence issue will be eliminated.

Due to the long transfer times at the vertiports, future research should involve optimal network design to incorporate balanced vertiport networks. In addition, incorporating waiting time and consolidation time into the network design will be able to decrease the average transfer times across vertiports. MANTA is currently being enhanced to carry out dynamic assignment, such that vehicles can adjust their routes based on the current congestion.  $Fe^3$  is currently being refactored to account for larger demand scales and vertiport networks, such as at  $k = 100$ , and also to consider local airspace restrictions. In addition, VertiSim is being restructured as a parallel programming architecture.

## VIII. Conclusion

This paper proposes a high-level multi-component architecture to integrate and model the effects of introducing an emerging mobility technology, urban air mobility, into the transportation network. Scholarship hitherto has tackled different components of this pipeline independently; however, in this research, a full toolchain, SimUAM, to model the surface traffic, vertiport, and aerial traffic with granular, agent-based, high-fidelity simulation, is developed. Initial results show that once the ground-air interface is modeled dynamically, the market for UAM decreases across all network designs relative to models with static assumptions about transfer times, which have often resulted in an optimistic estimate of the number of trips benefiting from use of urban air mobility. However, improvements can be made to balance the demand and optimize the networks for transfer time, likely increasing the number of benefited trips closer to the total addressable market once more. SimUAM is the first of its kind to incorporate fast, high-fidelity, regional-scale microsimulation on the ground and in the air with detailed ground-air interface modeling at the vertiport, and can be used as a powerful planning tool for UAM stakeholders.

## References

- [1] Giuliano, G., and Hanson, S., *The Geography of Urban Transportation*, 2017.
- [2] Stopher, P., and Stanley, J., *Introduction to Transport Policy: A Public Policy View*, 2014.
- [3] Garrow, L. A., Ilbeigi, M., and Chen, Z., "Forecasting Demand for On Demand Mobility," *17th AIAA Aviation Technology*,

- Integration, and Operations Conference*, AIAA AVIATION Forum, American Institute of Aeronautics and Astronautics, 2017. <https://doi.org/10.2514/6.2017-3280>.
- [4] Rothfeld, R., Fu, M., Balac, M., and Antoniou, C., “Potential Urban Air Mobility Travel Time Savings: An Exploratory Analysis of Munich, Paris, and San Francisco,” 2020.
  - [5] Yedavalli, P., “Designing and Simulating Urban Air Mobility Vertiport Networks under Land Use Constraints,” *Transportation Research Board*, 2021.
  - [6] Lim, E., and Hwang, H., “The Selection of Vertiport Location for On-Demand Mobility and Its Application to Seoul Metro Area,” *International Journal of Aeronautical and Space Sciences*, Vol. 20, No. 1, 2019, pp. 260–272. <https://doi.org/10.1007/s42405-018-0117-0>.
  - [7] Yedavalli, P., Kumar, K., and Waddell, P., “Microsimulation Analysis for Network Traffic Assignment (MANTA) at Metropolitan-Scale for Agile Transportation Planning,” *Transportmetrica A: Transport Science*, 2021, pp. 1–22. <https://doi.org/10.1080/23249935.2021.1936281>.
  - [8] Xue, M., Rios, J., Silva, J., Zhu, Z., and Ishihara, A. K., “Fe<sup>3</sup> : An Evaluation Tool for Low-Altitude Air Traffic Operations,” *2018 Aviation Technology, Integration, and Operations Conference*, American Institute of Aeronautics and Astronautics, Atlanta, Georgia, 2018. <https://doi.org/10.2514/6.2018-3848>.
  - [9] Axhausen, K. W., and ETH Zürich, *The Multi-Agent Transport Simulation MATSim*, Ubiquity Press, 2016. <https://doi.org/10.5334/baw>.
  - [10] Chan, C., Wang, B., Bachan, J., and Macfarlane, J., “Mobiliti: Scalable Transportation Simulation Using High-Performance Parallel Computing,” *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, IEEE, Maui, HI, 2018, pp. 634–641. <https://doi.org/10.1109/ITSC.2018.8569397>.
  - [11] Park, B. B., and Schneeberger, J. D., “Microscopic Simulation Model Calibration and Validation: Case Study of VISSIM Simulation Model for a Coordinated Actuated Signal System,” *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 1856, No. 1, 2003, pp. 185–192. <https://doi.org/10.3141/1856-20>.
  - [12] Rasouli, S., and Timmermans, H., “Activity-Based Models of Travel Demand: Promises, Progress and Prospects,” *International Journal of Urban Sciences*, Vol. 18, No. 1, 2014, pp. 31–60. <https://doi.org/10.1080/12265934.2013.835118>.
  - [13] Boeing, G., “OSMnx: New Methods for Acquiring, Constructing, Analyzing, and Visualizing Complex Street Networks,” *Computers, Environment and Urban Systems*, Vol. 65, 2017, pp. 126–139. <https://doi.org/10.1016/j.compenvurbsys.2017.05.004>.
  - [14] Metropolitan Transportation Commission, and Association of Bay Area Governments, “Plan Bay Area 2040,” , 2017.
  - [15] Zhao, B., Kumar, K., Casey, G., and Soga, K., “Agent-Based Model (ABM) for City-Scale Traffic Simulation: A Case Study on San Francisco,” *International Conference on Smart Infrastructure and Construction 2019 (ICSIC)*, ICE Publishing, Cambridge, UK, 2019, pp. 203–212. <https://doi.org/10.1680/icsic.64669.203>.
  - [16] Toledo, T., Koutsopoulos, H., Ben-Akiva, M., and Jha, M., “Microscopic Traffic Simulation: Models and Application,” *Simulation Approaches in Transportation Analysis*, Springer-Verlag, New York, 2005, pp. 99–130. [https://doi.org/10.1007/0-387-24109-4\\_4](https://doi.org/10.1007/0-387-24109-4_4).
  - [17] Treiber, M., and Kesting, A., *Traffic Flow Dynamics: Data, Models and Simulation*, Springer Berlin Heidelberg, 2013.
  - [18] Waddell, P., Boeing, G., Gardner, M., and Porter, E., “An Integrated Pipeline Architecture for Modeling Urban Land Use, Travel Demand, and Traffic Assignment,” *arXiv:1802.09335 [cs]*, 2018.
  - [19] Iqbal, M. S., Choudhury, C. F., Wang, P., and González, M. C., “Development of Origin–Destination Matrices Using Mobile Phone Call Data,” *Transportation Research Part C: Emerging Technologies*, Vol. 40, 2014, pp. 63–74. <https://doi.org/10.1016/j.trc.2014.01.002>.
  - [20] Yang, Q., and Koutsopoulos, H. N., “A Microscopic Traffic Simulator for Evaluation of Dynamic Traffic Management Systems,” *Transportation Research Part C: Emerging Technologies*, Vol. 4, No. 3, 1996, pp. 113–129. [https://doi.org/10.1016/S0968-090X\(96\)00006-X](https://doi.org/10.1016/S0968-090X(96)00006-X).
  - [21] Vascik, P. D., and Hansman, R. J., “Development of Vertiport Capacity Envelopes and Analysis of Their Sensitivity to Topological and Operational Factors,” *AIAA Scitech 2019 Forum*, 2019, p. 0526.
  - [22] Levine, R. V., and Norenzayan, A., “The Pace of Life in 31 Countries,” *Journal of Cross-Cultural Psychology*, Vol. 30, No. 2, 1999, pp. 178–205. <https://doi.org/10.1177/0022022199030002003>.

- [23] Zelinski, S., “Operational Analysis of Vertiport Surface Topology,” *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, 2020, pp. 1–10. <https://doi.org/10.1109/DASC50938.2020.9256794>.
- [24] Vascik, P. D., and John Hansman, R., “Evaluating the Interoperability of Urban Air Mobility Systems and Airports,” *Transportation Research Record: Journal of the Transportation Research Board*, 2021, p. 036119812199150. <https://doi.org/10.1177/0361198121991501>.
- [25] Bose, S., *An Introduction to Queuing Systems*, 2002. <https://doi.org/10.1007/978-1-4615-0001-8>.
- [26] Xue, M., and Kowalski, M., “Parametric Study of Federated Conflict Resolution for UAM Operations,” *AIAA Aviation Forum*, Virtual, 2021.
- [27] Xue, M., “Coordination between Federated Scheduling and Conflict Resolution in UAM Operations,” *AIAA Aviation Forum*, Virtual, 2021.
- [28] Silva, C., Johnson, W., Antcliff, K. R., and Patterson, M. D., “VTOL Urban Air Mobility Concept Vehicles for Technology Development,” *2018 Aviation Technology, Integration, and Operations Conference*, American Institute of Aeronautics and Astronautics, Atlanta, Georgia, 2018. <https://doi.org/10.2514/6.2018-3847>.
- [29] Chatterji, G. B., “Trajectory Simulation for Air Traffic Management Employing a Multirotor Urban Air Mobility Aircraft Model,” *AIAA Aviation*, Virtual event, 2020.
- [30] Yedavalli, P., Peng, X., Burak Onat, E., Sengupta, R., Waddell, P., and Bulusu, V., “Assessing the Value of Urban Air Mobility through Metropolitan-Scale Microsimulation: A Case Study of the San Francisco Bay Area,” *AIAA Aviation 2021*, 2021.