Mission Planner Algorithm for Urban Air Mobility – Initial Performance Characterization

Nelson M. Guerreiro[1], Ricky W. Butler[2], Jeffrey M. Maddalon[2], and George E. Hagen[2]
NASA Langley Research Center, Hampton, VA 23681, USA OK

In this paper, an initial characterization was performed of the Mission Planner algorithm developed by NASA for Urban Air Mobility (UAM) operations research. The algorithm plans conflict-free trajectories for flights to support a given set of UAM passenger trips. The UAM trips are planned in an on-demand, first-come, first-served manner, such that any given trip is subject to the constraints imposed by previously planned trips. For this analysis, the mission planning algorithm considered only the trajectory constraints from previously-planned trips in one test condition and added vertiport constraints for the second test condition. The conflict and constraint resolution strategies used by the Mission Planner were characterized by their percentage contribution to planning iterations, their percentage effectiveness in those iterations, and their contributions to the departure delay applied to each UAM trip's flight. With the exception of the climb and descent vertical speed strategies, most strategies showed reasonable or good performance in all test scenarios. In the test condition with vertipad constraints enabled, both the total number of iterations executed, and the number of flights that required planning iterations, was reduced for all scenarios. This was the result of the natural conditioning of the traffic achieved with scheduling and the additional information available to the Mission Planner from the vertiport scheduler. The next steps for this work will include improvements to the mission planning strategies and analyses with additional constraints and under other demand scenarios.

I. Introduction

NASA's Air Traffic Management-Exploration (ATM-X) project[1] is studying the impact that new entrants may have on the National Airspace System (NAS). These new entrants include Urban Air Mobility operations, or the carrying of goods and people in and around cities. Research studies currently underway are focused on understanding the impact to existing operations, refining existing or new concepts of operations, and identifying the air traffic and other services required to enable safe and efficient use of the NAS for these new operations.[2]

UAM represents a new type of air traffic operations. These operations will need to safely operate in a near-term (within 5-10 years) environment but also require a well-defined concept of operations for how the air traffic system can support the predicted level of demand[3] of a far-term (20-30 years) environment. One way that NASA is exploring the potential to support both of these environments is through the development of services in support of a service-oriented architecture (SOA). A SOA provides many advantages including reliability, scalability, and simplified maintenance. One example of a SOA is NASA's Unmanned Aerial System (UAS) Traffic Management system, or UTM.[4]

II. Background

Mission planning is one service that could be available in a SOA UAM architecture. A mission planning function plans the details of how a UAM flight will get from origin to destination. Mission planning ensures that any known system constraints are considered in the planning of flight trajectories, including constraints of capacity-limited resources, and constraints from previously planned flight trajectories. Mission planning can also take into account

---

[1] Research Aerospace Engineer, Crew Systems and Aviation Operations Branch, MS 152, AIAA Member.
[2] Research Engineer, Safety-Critical Avionics Systems Branch, Mail Stop 130.

operator-specific constraints, such as vehicle availability. A mission planning service would be one element in a system that includes, among other services, conformance monitoring and a re-planning service as one method to handle poor performance, uncertainty, or off-nominal situations in the air traffic system. One example of such an architecture is shown in Figure 1. Another example would be the UTM framework, where mission planning could be part of a UAM operator's functions or offered as a service to those operators.

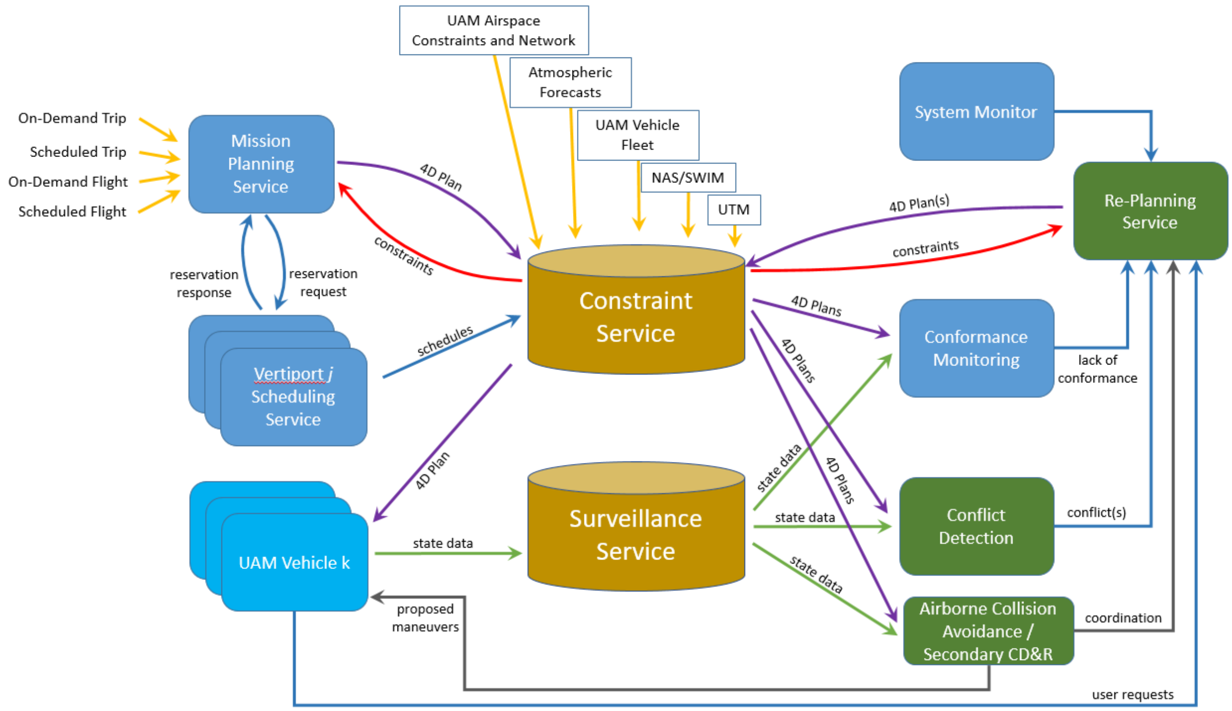

Figure 1. One possible architecture supporting the need for a mission planning and mission re-planning service.

In this work, a prototype mission planning algorithm has been developed for UAM operations. The algorithm plans 4-dimensional (4-D) trajectories for a set of UAM trips in the presence of any known constraints. The possible set of constraints includes vehicle availability, vertiport availability, other planned trajectories, airspace avoidance regions, and terrain or other obstructions. In this paper, an initial characterization of this algorithm is described.

A. ATS-TIGAR

The Mission Planner algorithm is hosted within the Advanced Trajectory Services – Toolkit for Integrated Ground and Air Research (ATS-TIGAR) tool. This tool, shown in Figure 2, is a derivative of the TBO-TIGAR[5] tool previously developed for Trajectory-Based Operations research. ATS-TIGAR allows for the rapid prototyping of algorithms and for the analysis of those algorithms under different scenarios.

A UAM scenario input to ATS-TIGAR defines a set of desired trips, the vertiport system, the vehicle fleet, and other scenario parameters. Each trip includes an origin, a destination, the number of passengers, and the desired trip start time. The vertiport system includes the location of each vertiport and their basic configuration, including the number of vertipads and the number of parking spots. The vehicle fleet is defined by the type and quantity of each vehicle, their performance characteristics, their passenger capacity, and the initial allocation of the vehicles to the available vertiports.

ATS-TIGAR implements a sub-set of the functions or services described in Figure 1. First, and foremost, ATS-TIGAR implements the Mission Planner algorithm that is the subject of this analysis, and which is described in the next section. Second, ATS-TIGAR implements the equivalent of a Constraints Service. The Constraints Service stores information about the constraints that should be considered in mission planning, which are the 4-D trajectories of already planned trips for this analysis. ATS-TIGAR also implements a version of the Vertiport Scheduling Service, where the Mission Planner Service is able to query and make reservations on the vertiport resources (e.g., the vertipads and parking spaces).

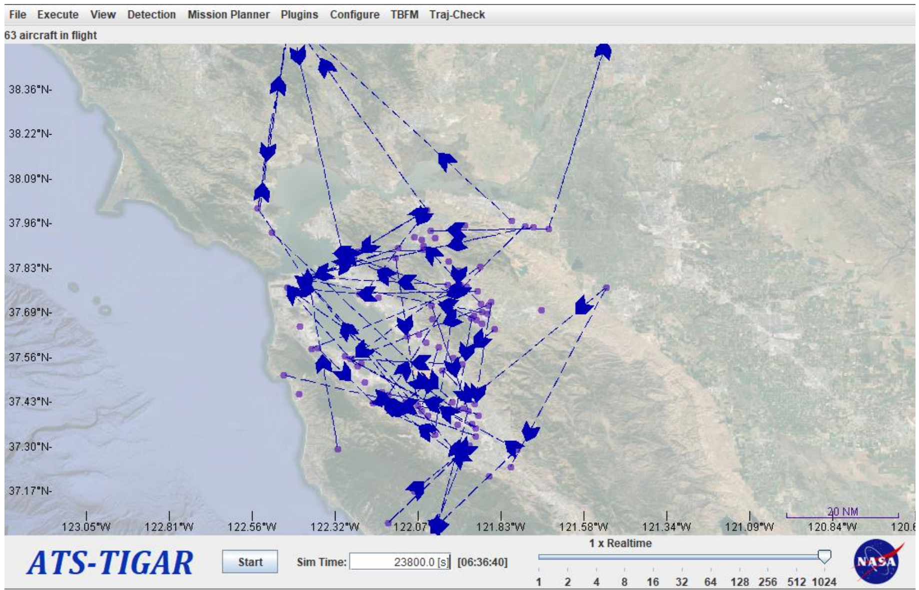

Figure 2. ATS-TIGAR tool visualization.

B. Mission Planner

The Mission Planner algorithm plans the flights required to support a given UAM trip, and the 4-D trajectories for those flights, in the presence of known constraints. The steps of the algorithm are to identify a suitable vehicle for the given trip, to initiate the movement of that vehicle to the origin vertiport (if not already there), to identify suitable origin and destination takeoff and landing times, and to plan a trajectory that avoids other trajectories, restricted airspace, and other obstructions.

The first step of the Mission Planner algorithm is to identify a suitable vehicle for a given trip. This fleet management feature is one that can be enabled or disabled in the algorithm. When disabled, the algorithm assumes that a vehicle is available immediately at the origin vertiport. When enabled, the algorithm finds a vehicle, from a given fleet of vehicles, with a passenger capacity equal to or greater than the number of passengers on the requested trip, and with earliest availability at the origin vertiport. In some situations, this vehicle is already at the origin vertiport. In other instances, a repositioning flight is required to bring that vehicle to the origin vertiport. In some cases, a clearing flight may be required at the origin vertiport to make room for the repositioning vehicle, assuming the vertiport has limited parking capacity. Similarly, a clearing flight may be required at the destination vertiport to make room for the requested trip's main flight, again assuming limited parking capacity. Consequently, for any requested UAM trip, there may be a need to plan trajectories for up to four flights that support the execution of the primary trip when fleet and vertiport constraints are enabled.

The second step of the Mission Planner algorithm is to find suitable takeoff and landing times at the origin and destination vertiports, based on the availability of vertipads and parking spaces at those vertiports. ATS-TIGAR implements a simple, first-come, first-served, reservation-based scheduler for each vertiport in a scenario. These vertiport schedulers are queried for available takeoff time and landing time (while considering parking space constraints) based on the availability of the UAM vehicle and, initially, based on the unimpeded trajectory time from origin to destination. At the end of the mission planning cycle, the vertiport schedulers are queried once again to confirm the availability of the takeoff and landing times and reservations are made with those schedulers.

The third step of the Mission Planner is to find conflict-free, 4-D trajectories from the origin to the destination, for each flight, that adhere to airspace constraints. Because vehicle availability and vertiport availability constraints are handled in the other steps of mission planning, this step is primarily concerned with enforcing the constraints introduced by other planned trajectories, airspace avoidance regions, and terrain or other obstructions. This step in the Mission Planner must also take into consideration the constraints of operating rules or procedures that may exist, such as pre-defined routings or other special procedures. The Mission Planner algorithm finds all conflicts along a proposed trajectory with those constraints and uses a set of strategies to resolve those conflicts. A conflict is defined as the projected loss of separation with another aircraft, the entry into a restricted volume of airspace, or the violation of any other constraint. The resolution strategies are targeted at resolving the conflicts sequentially starting at the origin and ending at the destination. The strategies currently used by the Mission Planner are discussed in Section B.

The implementation of the Mission Planner algorithm used for this analysis assumes a pre-departure planning service. That is, one of the mechanisms available for conflict resolution (or constraint satisfaction) is a departure delay. The algorithm implements an iterative approach to trial planning until a solution is found that produces a conflict-free and constraint-satisfying trajectory. An extension of this algorithm useful for re-planning after departure is also being developed but is not discussed in this paper.

III. Analysis Design

In this analysis, the objective was to characterize the effectiveness of each of the Mission Planner's strategies in resolving trajectory conflicts or constraints for a given set of scenarios. This initial performance characterization will drive future development of additional strategies, or improvement of current strategies, to more effectively resolve the conflicts within a given phase of flight along of a trajectory. The efficiency of a mission planning algorithm will play a critical role in the ability to support high-density UAM operations in the presence of constrained resources, such as vertiports and limited airspace.

The approach used to perform this initial performance characterization was to exercise the Mission Planner algorithm with a set of demand scenarios. The scenarios were analyzed for each resolutions strategy's usage, effectiveness, and contribution to departure delays.

A. Scenarios

The Mission Planner algorithm was exercised with a set of demand scenarios for the Northern California region. These scenarios originated from prior work completed by Virginia Tech for NASA in estimating the potential demand for UAM passenger trips using a mode choice model, given a set of socio-economic factors and historical commuting patterns[6]. The high-demand data set obtained from this work described nearly forty thousand UAM flights to nearly one thousand vertiports during a 3-hour period of the morning, with an assumed load factor of three passengers per flight.

The high-demand data set was modified for this analysis by:

- Re-allocating flights from the lowest-demand vertiports to the 100 highest-demand vertiports (Figure 3), based on the shortest distance to a high-demand vertiport;
- Randomly down-sampling for 1000, 2000, 3000, 4000, 5000, 7500, and 10000 flights scenarios; and
- Converting each flight to a single trip with three passengers.

The final scenario set consisted of seven scenarios with 1000, 2000, 3000, 4000, 5000, 7500, and 10000 UAM 3-passenger trips distributed over a 3-hour window. Each of these UAM trips was defined by five parameters: origin vertiport, destination vertiport, number of passengers, demanded time (time at which Mission Planner would know about the trip), and the desired trip start time (at or after the demanded time).

The seven demand scenarios were run in ATS-TIGAR with the Mission Planner in two conditions. In the first test condition, the only constraints were the trajectories of previously planned flights, with no vertiport constraints (equivalent to unlimited parking spaces and unlimited vertipads). In the

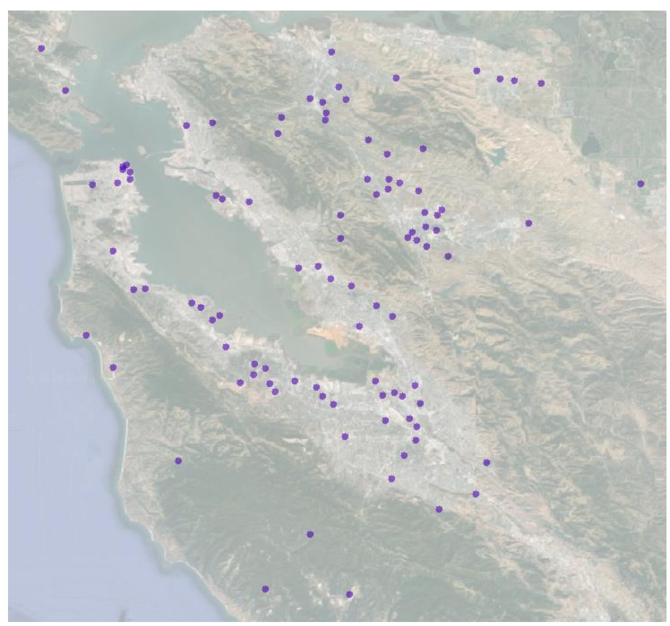

Figure 3. Vertiport placement for the Northern California scenarios (note: two vertiports not shown here, one near Sacramento, CA and one near Angwin, CA).

second test condition, the additional constraint of two independent vertipads per vertiport was introduced, still with unlimited parking spaces. The Mission Planner planned a flight for each of the scenario trips at the demanded time, leveraging its various resolutions strategies to handle the constraints in the system. Metrics were gathered regarding the usage, effectiveness, and delay contribution of each strategy.

B. Mission Planner Resolution Strategies and Planning Algorithm

The set of Mission Planner strategies used to resolve conflicts with traffic or other constraints is described in Table 1. With the exception of the departure delay strategy, each resolution strategy has an associated “short name” that is used in the discussion and in many of the results figures in the Results section.

Table 1. Mission Planner Conflict Resolution Strategies.

Strategy	Short Name	Description
Departure Delay		When all suitable combinations of other strategies fail to resolve a conflict, the Mission Planner adds a departure delay
Change Departure Direction	*depDir*	Change the departure direction away from the nominal departure direction
Change Climb Vertical Speed	*depVS*	Increase the climb vertical speed
Offset Climb Path	*climbOffset*	Offset the climb path by an angle to create an offset on the top-of-climb point relative to its prior location
Change Cruise Altitude	*cruiseAltChg*	Change the cruise altitude to one stratification lower than the nominal cruise altitude
Use En Route Conflict Resolution	*enrouteCR*	Use Stratway[7] strategies to resolve any en route conflicts
Change Arrival Direction	*arrDir*	Change the arrival direction away from the nominal arrival direction
Change Descent Vertical Speed	*arrVS*	Increase the descent vertical speed magnitude
Offset Descent Path	*descentOffset*	Offset the descent path by an angle to create an offset on the top-of-descent point relative to its prior location

The departure delay strategy is used to delay a trip’s start time by a certain amount and is used in one of two ways. First, departure delay is used when the vertiport schedules indicated the lack of availability of a takeoff vertipad at the origin vertiport or a landing vertipad at the destination vertiport, at the takeoff or landing times of each trajectory trial plan. In this use case, the amount of delay required to achieve the next available landing or takeoff time is added to the trip start time. Second, departure delay serves as the last available measure to resolve a conflict when the other resolution strategies fail to solve a given set of conflicts. In this use case, a pre-defined amount of departure delay is added to the trip start time before beginning the next trial plan iteration. Whenever any departure delay is added to a UAM trip, a new mission planning iteration is started. Note that vehicle availability constraints are typically handled by the departure delay strategy, but that vehicle availability was not a constraint in this analysis. Once vehicle availability constraints are addressed by the Mission Planner, the next step is to identify an available takeoff time at the origin vertiport. A departure delay is added to the trip’s start time if the next available takeoff time is greater than the current planning iteration’s takeoff time.

After resolving vehicle availability and origin vertiport constraints, the Mission Planner computes a nominal trajectory from the origin vertiport to the destination vertiport. This trajectory is evaluated for conflicts with traffic and other constraints. The next steps of the algorithm attempt to resolve the conflicts that may exist in different portions of the trajectory.

The Mission Planner uses a strategy that changes the departure direction as the method for resolving departure conflicts (the *depDir* strategy). The objective of the Mission Planner is to obtain a conflict-free trajectory at some radial distance from the origin vertiport. The approach for handling UAM departure and arrival operations near the vertiports, where vehicles may be using reduced separation criteria, has yet to be defined. Consequently, this analysis ignores trajectory conflicts within this radial distance of the vertiports. A radial distance equal to the minimum horizontal separation criterion was used in this analysis. The algorithm iterates on up to six departure directions (including the nominal Great Circle departure direction), equally distributed amongst all possible directions, to try to resolve departure conflicts. When not all conflicts can be resolved with a departure direction change, the default strategy (departure delay) is used to delay the trip start time and a new mission planning iteration is started.

When all departure conflicts are resolved, the Mission Planner's objective is to resolve any conflicts in the climb phase of flight using one or more of three available strategies. The first strategy that is attempted is to change the vertical rate in climb (*depVS*). The *depVS* strategy iterates over a set of scale factors applied to the nominal climb vertical rate. The second strategy that is attempted is to introduce an offset to the lateral path in the climb (*climbOffset*). The *climbOffset* strategy introduces a change in the lateral path in the climb phase of the trajectory that implements an offset of the top-of-climb point equal to the horizontal separation criterion, relative to the nominal trajectory's top-of-climb location. Finally, a strategy attempts to change the cruise altitude (*cruiseAltChg*). The *cruiseAltChg* strategy changes the planned cruising altitude to one altitude stratification lower than the nominal cruise altitude. When one of these strategies, or a combination of them, is unable to resolve all climb conflicts, the default strategy (departure delay) is again used to delay the trip start time and a new mission planning iteration is started.

After resolving all conflicts in the departure and climb phases of the trajectory, the Mission Planner works on resolving any conflicts present in the level, cruise phase of the trajectory. The Mission Planner uses the *enrouteCR* strategy, which leverages the NASA-developed Stratway[7] conflict resolution algorithms. Stratway uses a combination of lateral path and vertical path changes to resolve, sequentially, the conflicts in this phase of the trajectory. As before, if the *enrouteCR* strategy is unable to resolve all en route conflicts, the default, departure delay, strategy delays the trip start time and a new mission planning iteration is started.

In the descent and arrival phases of the trajectory, the Mission Planner uses strategies similar to those used for the departure and climb phases of the trajectory. That is, the Mission Planner uses an arrival direction change strategy (*arrDir*) to resolve conflicts at a radial distance from the arrival vertiport, and uses descent vertical speed (*arrVS*), descent path offset (*descentOffset*), and cruise altitude change (*cruiseAltChg*) strategies to resolve trajectory conflicts in the descent phase of the trajectory. Departure delay is the default resolution strategy when the Mission Planner is unable to resolve conflicts in the descent or arrival phases of the trajectory.

Once all trajectory conflicts are resolved, the Mission Planner computes a landing time at the destination vertiport and attempts to make the appropriate vertiport reservations. Using the projected landing time, the Mission Planner evaluates the availability of the destination vertiport. If the next available landing time is greater than the projected landing time, a departure delay is added to the trip start time and a new mission planning iteration is started. If the vertiport is available at the projected landing time, the mission planning iterations terminate and the takeoff and landing reservations are made at the origin and destination vertiports.

The ATS-TIGAR tool provides a timer, which cycles through the time window of a given scenario. ATS-TIGAR calls the Mission Planner algorithm at each scenario time to plan the UAM trips on a first-come, first-served (or on-demand) approach based on each trip's demanded time. Each trip planning iteration considers the constraints introduced by all previously planned trips (their associated flight trajectories and vertipad reservations).

C. Assumptions

The assumptions used by the Mission Planner algorithm, as well as the assumptions specific to this analysis, are presented in this section. The assumptions are focused around the two test conditions for the seven scenarios. In the first condition, only the trajectories of previously planned trips were constraints to the planning of any given UAM trip – no vertiport or other airspace constraints were active. This allowed for the analysis of the Mission Planner strategies in resolving only traffic conflicts. In the second condition, the trajectories of the previously planned trips were constraints but the vertiports were assumed to have a limited number of vertipads – this enabled vertipad scheduling while still assuming infinite parking capacity at the vertiports.

Mission Planner algorithm assumptions:

- Each trip is planned on a first-come, first-served basis at the demanded time
- Each trip's start time represents the start of passenger boarding – the flight's takeoff time is some time after the trip start time
- When a trip is delayed due to a failed planning strategy, the next trial planning iteration may exercise any previously attempted strategies on the new set of conflicts

Analysis-specific assumptions:

- Only one Mission Planner instance was used to plan all trips in a scenario
- Each trip's demanded time was the same as the desired trip start time (on-demand trip planning)
- Each trip assumed 3 passengers
- The passenger loading and un-loading time used was 60 seconds per passenger

- The vehicle fleet management functionality was disabled (vehicles were available immediately at the origin vertiport)
- Only one flight was planned for each trip (no vehicle repositioning or clearing trips were added)
- No avoidance airspace regions were used
- No terrain, building, or other obstructions were considered
- Nominal trajectories from origin to destination vertiports used Great Circle routing
- Nominal routing was changed by the various resolution strategies as needed to resolve conflicts
- Trajectory conflicts inside a 1000 foot radius from each vertiport were ignored
- The radial distance from the vertiports used for the departure and arrival direction strategies (*depDir* and *arrDir*) was 1000 feet
- The vertiport constraints were:
    - Unlimited vertipads and parking spaces for the first test condition,
    - Two vertipads with unlimited parking spaces per vertiport for the second test condition, and
    - Vertipad usage time slots of 60 seconds were used in the second condition
- The vertical speed strategies (*depVS* and *arrVS*) used vertical speed multipliers of [1.1, 1.2, 1.3, 1.4, 1.5]
- The default departure delay value used was 15 seconds
- The horizontal separation criterion used was 1000 feet
- The vertical separation criterion used was 450 feet
- The cruise altitude stratification used was 500 feet
- The minimum cruise altitude was 500 feet and maximum cruise altitude was 2000 feet, determined as a function of origin-to-destination distance, with trips beyond ~15 NM achieving the maximum altitude

D. Metrics

The objective of this analysis was to perform an initial characterization of the performance of the Mission Planner algorithm and its associated strategies. The performance was characterized in each of the two test conditions, with and without vertipad constraints, as a function of the demand scenario, defined by the number of flights. The analysis characterizes the total number of planning iterations completed in each demand scenario, the percentage of flights requiring planning iterations, the percentage of planning iterations used by each strategy, the success rate, or percentage effectiveness of each strategy, and the delays attributed to each strategy.

The contribution of each Mission Planner strategy to the trip departure delays was captured by a set of delay metrics, as shown in Figure 4. The *origPortDelay* and *destPortDelay* metrics captured the delay implemented by the Mission Planner due to origin or destination vertiport availabilities, respectively. The departure and arrival direction

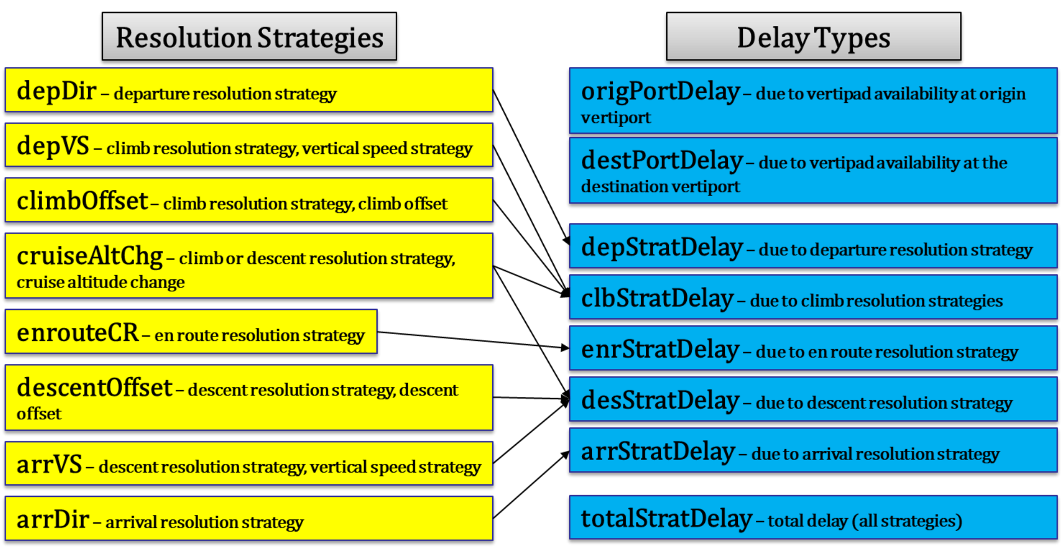

Figure 4. Mission Planner strategy to delay metric mapping.

strategies (*depDir* and *arrDir*) were captured in the *depStratDelay* and *arrStratDelay* metrics. In the climb phase of the trajectory, the *clbStratDelay* captured the contributions to delay from the strategies of *depVS*, *climbOffset*, and *cruiseAltChg*. Similarly, in the descent phase of the trajectory, the *desStratDelay* captured the contributions to delay from the strategies of *arrVS*, *descentOffset*, and *cruiseAltChg*. Note that *cruiseAltChg* could only have been used once, either by the climb resolution strategy, or by the descent resolution strategy. The *enrStratDelay* metric captured the contribution to delay of the *enrouteCR* strategy. Finally, the *totalStratDelay* is the sum of the delay contribution from all strategies.

IV. Results

In this section, the results obtained from the two run conditions, and from each demand scenario, are presented. The figures from each of the two conditions are plotted together in order to aid in discussion and comparison.

Figure 5 shows the percentage of flights in each scenario that required planning iterations for the condition with no vertiport constraints (left), and the condition with two vertipads per vertiport (right). An iteration of mission planning implies that the unimpeded or nominal trajectory for a flight was not free from traffic or other constraint conflicts. In both conditions, the number of flights requiring planning iterations is small in the 1000 flights scenarios (about 12%). This represents a low demand scenario, where most of the requested trips were able to start at the requested time and followed the nominal Great Circle route. As the demand increased from 1000 flights to 10000 flights, an increasingly greater percentage of flights required planning iterations. This increasing pattern had nearly the same magnitude in both constraint conditions, but only up to 5000 flights. In the 7500 and 10000 flights scenarios, the condition with vertipad constraints had slightly smaller percentages of flights requiring planning iterations. This is because the Mission Planner is able to take larger delay steps (more than the default 15 seconds) based on the vertiport availability reported from each vertiport's scheduler, which leads to fewer iterations. As will be discussed below, these higher demand scenarios also led to a greater number of scheduling conflicts at the origin and destination vertiports when compared to the lower demand scenarios.

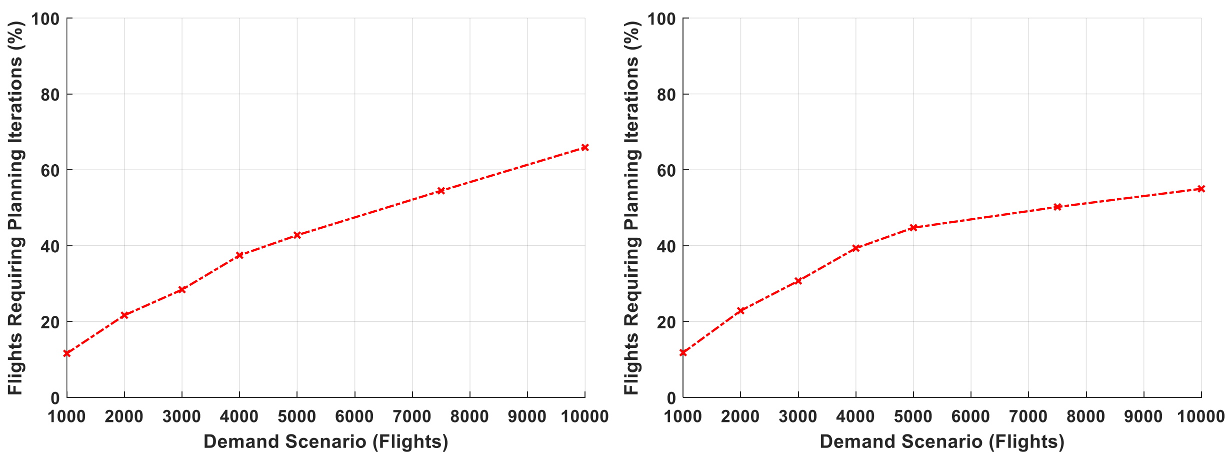

Figure 5. Percentage of flights requiring planning iterations for each demand scenario, with no vertiport constraints (left), and with two vertipads per vertiport (right).

Figure 6 shows the number of planning iterations executed in each of the demand scenarios for the condition with no vertiport constraints (left), and the condition with two vertipads per vertiport (right). Similar to the percentage of flights requiring planning iterations, the number of overall planning iterations in each scenario was comparable in both conditions but only up to the demand scenario with 5000 flights. In the scenarios with 7500 and 10000 flights, the number of total planning iterations was smaller in the condition with vertipad constraints enabled, again, due to the ability of the Mission Planner to delay the amount required to meet a given vertiport's availability versus delaying a default value.

Figure 7 shows the average number of planning iteration per flight, for those flights requiring planning iterations, for both test conditions. It is clear from this figure that the enabling of the vertipad scheduling constraints impacted mission planning by reducing the average number of iterations required for each flight, especially with the demand scenarios of 2000 flights and higher. One reason for this reduction is the Mission Planner's ability to delay the required amount, based on vertipad availability, as mentioned above. Another reason is the fact that, with scheduling enabled, the likelihood of traffic conflicts near a vertiport is reduced because the traffic rate into and out of the vertiport is limited and conditioned.

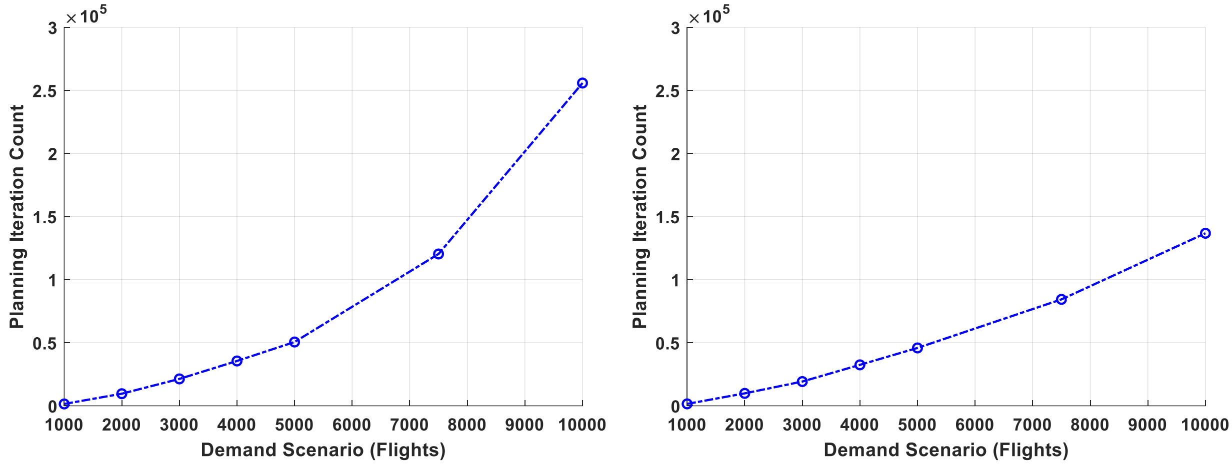

Figure 6. Total number of planning iterations executed for each demand scenario, with no vertiport constraints (left), and with two vertipads per vertiport (right).

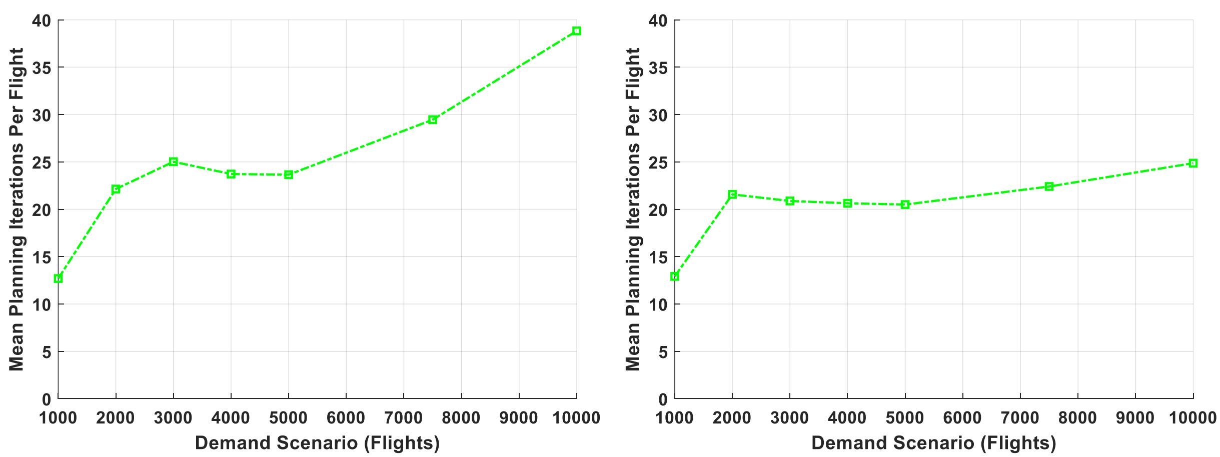

Figure 7. Mean planning iterations per flight (only flights requiring planning iterations) for each demand scenario, with no vertiport constraints (left), and with two vertipads per vertiport (right).

In Figure 8, the percentage of all planning iterations attributed to each of the Mission Planner's resolution strategies is shown for the condition with no vertiport constraints (left) and the condition with two vertipads per vertiport (right). There were no significant differences in these percentages between the two conditions. Most strategies showed slightly lower usage in the condition with vertipad scheduling constraints enabled, especially in the higher demand scenarios. Conversely, the descent vertical speed strategy (*arrVS*) did show higher percentage usage in the conditions with scheduling enabled. In both conditions, the strategy with the highest percentage of planning iterations was the climb vertical speed iteration (*depVS*). This is somewhat unbalanced but is to be expected because this strategy was the first used in resolving trajectory conflicts in a climb, and iterated over five vertical speed multipliers before the alternative strategies (*climbOffset* and *cruiseAltChg*) were attempted.

Figure 9 shows the percentage effectiveness for each of the Mission Planner strategies in both test conditions. From this figure, it is clear that the climb and descent vertical speed strategies (*depVS* and *arrVS*) were ineffective in resolving the trajectory conflicts in the climb and descent phases of the trajectory. The en route resolution strategy (*enrouteCR*) was highly effective in resolving en route trajectory conflicts, with only a few percent effectiveness reduction as the demand increased. All other strategies had effectiveness ranging from 11 percent to 57 percent.

The en route resolution strategy did have two advantages over the other Mission Planner strategies, with respect to strategy effectiveness. First, *enrouteCR* had its own iterative loop for solving the traffic conflicts that was not counted in the Mission Planner's iteration count. Second, *enrouteCR* leveraged conflict information – namely the time into and out of a conflict – in order to identify a suitable conflict resolution.

The strategy effectiveness data does indicate that the Mission Planner strategies as tested could benefit from more information about the specific traffic conflict or conflicts in order to make more informed trajectory changes.

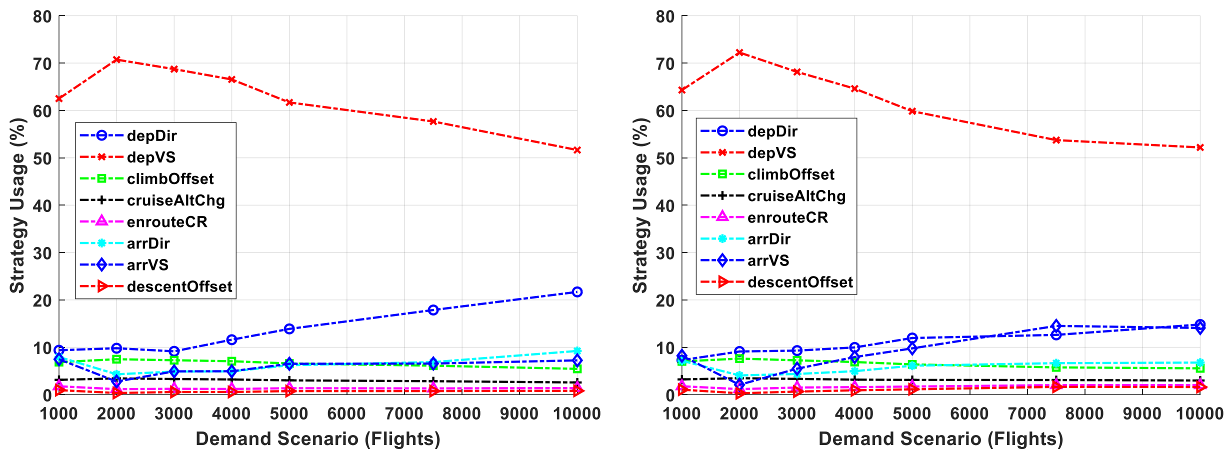

Figure 8. Percentage of all planning iterations used by each strategy for each demand scenario, with no vertiport constraints (left), and with two vertipads per vertiport (right).

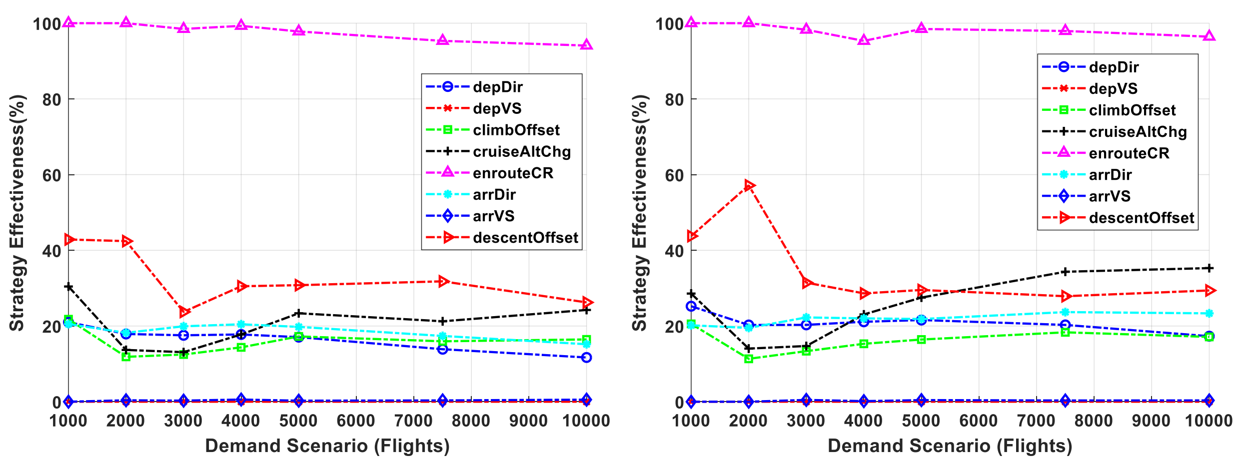

Figure 9. Percentage strategy effectiveness for each demand scenario, with no vertiport constraints (left), and with two vertipads per vertiport (right).

Figure 10 shows the percentage of flights impacted by departure delays from each of the strategies, or by any strategy (*totalStratDelay*), for both test conditions. Note that these figures include the impacts of scheduling constraints in the origin and destination vertiport delay metrics (*origPortDelay* and *destPortDelay*). These metrics are zero in the condition with no scheduling constraints but very quickly dominate the percentage of flights with delay as the demand increases when scheduling constraints are introduced. This indicates that, at the higher demand scenarios, the scheduling constraints become more dominant than the traffic constraints when it comes to mission planning, and the vertiports can become the bottlenecks in the system.

There was a greater percentage of flights with destination vertiport scheduling delay as compared to origin vertiport delay in the 7500 and 10000 flights demand scenarios. This is a limitation of the current Mission Planner algorithm that can be overcome with the addition of a speed change strategy to satisfy, simultaneously, an available departure slot and an available arrival slot. In the current implementation, the arrival slot was identified first, the trajectory was computed next, and then the query was made for an available arrival slot at the resulting landing time. If an arrival slot were not available at that landing time, departure delay would be attributed to the destination vertiport's delay metric, which could lead to more delays due to destination vertiport availability than due to origin vertiport availability, especially in high demand conditions.

The impact of scheduling constraints on mission planning can be further validated by considering the mean flight delays, for delayed flights, from each strategy, shown in Figure 11. In most conditions with no scheduling constraints, each mission planning strategy introduced less than one minute of delay on average to each delayed flight, even in the higher demand scenarios. The average total delay only exceeded one minute (64 seconds) in the 10000 flights scenario. In comparison, when vertipad scheduling was introduced, the origin, destination, and total mean delays per flight increased rapidly in the 5000 flights and higher demand scenarios.

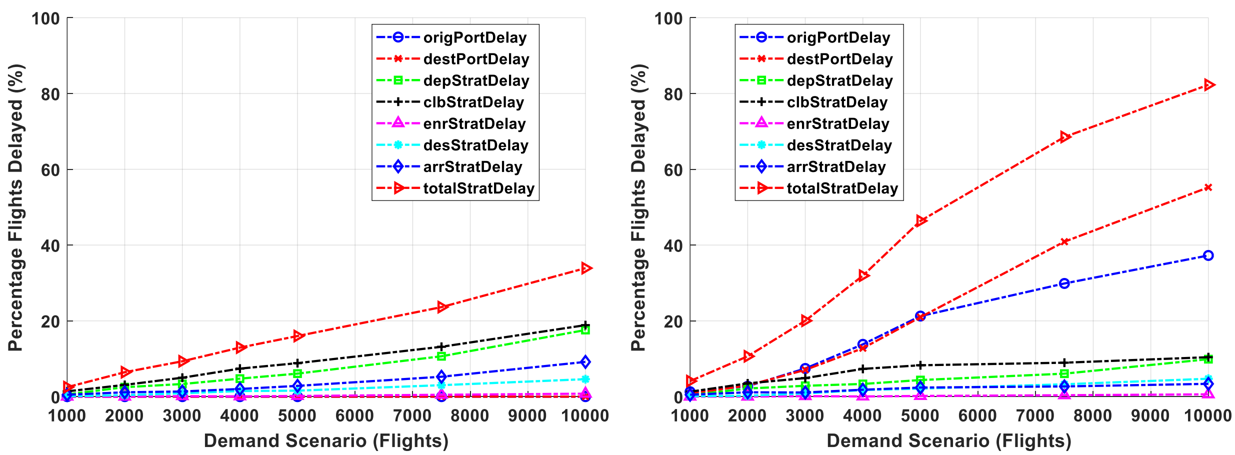

Figure 10. Percentage of flights delayed with each type of delay for each demand scenario, with no vertiport constraints (left), and with two vertipads per vertiport (right).

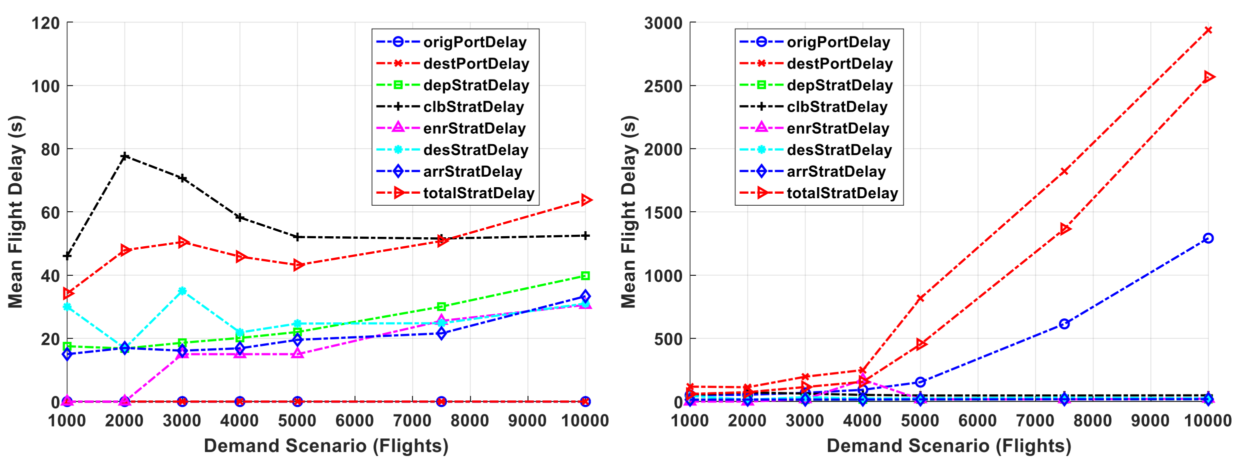

Figure 11. Mean delay per flight (only delayed flights) by delay type for each demand scenario, with no vertiport constraints (left), and with two vertipads per vertiport (right).

The results presented in this paper did not consider the additional constraints that could be present at the vertiports, especially in terms of available parking spaces independent of the vertipads, or vehicle availability, which could further negatively affect the delays that would be required for each flight.

V. Conclusion

In this paper, an initial characterization of a mission planning algorithm for UAM operations research was performed. The mission planning algorithm planned conflict-free trajectories for flights to support a given set of UAM passenger trips. The UAM trips were planned in an on-demand, first-come, first-served manner, such that any given trip was subject to the constraints imposed by previously planned trips. For this analysis, the mission planning algorithm considered only the trajectory constraints from previously-planned trips in one test condition and added the constraints of a limited number of vertipads at the vertiports for the second condition. The Mission Planner algorithm, in combination with the ATS-TIGAR tool, is a capability that supports real-time and fast-time analysis of UAM demand scenarios in order to understand the impacts of these operations on the NAS.

The introduction of the vertipad constraints in the second test condition reduced the number of flights that required planning iterations and the number of iterations required by the Mission Planner, both as a per flight average, and the total iterations in a single demand scenario. This was the result of a reduced number of conflicts near the vertiports due to the natural conditioning of the traffic achieved with scheduling. Additionally, with scheduling constraints enabled, the Mission Planner was able to implement the required amount of departure delay to achieve the vertiport's availability versus incrementally iterating with the default, departure delay value.

Overall, the Mission Planner strategies were reasonably effective in resolving the traffic constraints with the exception of the climb and descent vertical speed multiplier strategies (*depVS* and *arrVS*). Most strategies showed an effectiveness between 11 and 57 percent over both test conditions and all demand scenarios. The en route conflict resolution strategy (*enrouteCR*) showed greater than 94 percent effectiveness in all scenarios, while the climb and descent vertical speed strategies showed nearly 100 percent ineffectiveness in all scenarios.

The departure delays implemented by the Mission Planner were largely driven by the strategies to resolve conflicts near a vertiport, or in the climb or descent phases of flight. In the condition with no vertiport constraints (unlimited capacity) the average delay per flight introduced by each strategy, and for overall delay from all strategies, was typically less than 60 seconds. This was not the case in the condition with the added vertipad constraints where, for the demand scenarios with more than 5000 flights, the average delay per flight increased to as high as 48 minutes.

Future work with the Mission Planner will focus on improving the effectiveness of the various constraint resolution strategies. Additional analyses also need to be conducted to understand the impact of other airspace constraints, such as limited parking spaces, vehicle constraints, restricted airspace regions, and pre-defined route networks, to UAM operations under various demand scenarios.

VI. References

[1]Chan, W. N., et al., "Overview of NASA's Air Traffic Management - eXploration (ATM-X) Project," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, AIAA 2018-3363, Atlanta, Georgia, 2018.

[2]Thipphavong, D. P., et al., "Urban Air mobility Airspace Integration Concepts and Considerations," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, AIAA 2018-3676, Atlanta, Georgia, 2018.

[3]Goyal, R., "Urban Air Mobility (UAM) Market Study," Booz Allen Hamilton, NASA Technical Report HQ-E-DAA-TN64105, 2018.

[4]Kopardekar, P., et al., "Unmanned Aircraft System Traffic Management (UTM) Concept of Operations," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, AIAA 2016-3292, Washington, DC, 2016.

[5]Guerreiro, N.M., et al., "Prototype Tool and Focus Group Evaluation for an Advanced Trajectory-Based Operations Concept," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, AIAA 2017-4096, Denver, Colorado, 2017.

[6]Syed, N., et al., "Preliminary Considerations for ODM Air Traffic Management based on Analysis of Commuter Passenger Demand and Travel Patterns for the Silicon Valley Region of California," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, AIAA 2017-3082, 2017.

[7]Hagen, G.E., Butler, R.W., Maddalon, J.M., "The Stratway Program for Strategic Conflict Resolution: User's Guide," NASA TM-2016-219196, 2016.