NASA/TM–20220017551

Analysis of Strategic Conflict Management Approaches as Applied to Simulated UAM Operations

Nelson M. Guerreiro, George E. Hagen, Jeffrey M. Maddalon
Langley Research Center, Hampton, Virginia

Ricky W. Butler
Analytical Mechanics Associates, Hampton, Virginia

April 2023

NASA STI Program Report Series

The NASA STI Program collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- Help desk contact information:

https://www.sti.nasa.gov/sti-contact-form/
and select the "General" help request type.

NASA/TM–20220017551

Analysis of Strategic Conflict Management Approaches as Applied to Simulated UAM Operations

Nelson M. Guerreiro, George E. Hagen, Jeffrey M. Maddalon
Langley Research Center, Hampton, Virginia

Ricky W. Butler
Analytical Mechanics Associates, Hampton, Virginia

National Aeronautics and
Space Administration

Langley Research Center
Hampton, VA, 23681

April 2023

Available from:

NASA STI Program / Mail Stop 148
NASA Langley Research Center
Hampton, VA 23681-2199
Fax: 757-864-6500

Contents

Executive Summary

This report presents the results of a comparison analysis of candidate Strategic Conflict Management (SCM) strategies as applied to Urban Air Mobility (UAM) operations. The SCM strategies investigated included: no SCM, resource scheduling (RS), resource flow rates (FR), area-based flow rates (AR), and conflict detection and resolution (CR). The study evaluated each strategy against the same 3 levels of flight demand in a representative airspace construct for the Dallas/Fort Worth region. The study also accounted for different levels of uncertainty in operational planning and equivalent levels of trajectory following error. In the analysis, we compared the SCM strategies' effectiveness in reducing the need for the tactical conflict management layer to act and looked at the metrics of unmitigated losses-of-separation (LOS), flight delays imposed by strategic planning, and throughput of the overall airspace.

The study results indicated that the CR strategy was the most effective at reducing LOS and percentage of flights with LOS, even in the presence of trajectory error when uncertainty is accounted for in planning. Thus, if the objective is to reduce the number of actions that the tactical conflict management layer will have to take, in terms of conflicts that may need to be resolved, CR is the strategy to use. The FR strategy was found to be the worst at reducing LOS and percentage of flights with LOS. Thus, it is not very effective at reducing the actions required by a tactical conflict management layer.

In the airspace tested, the demand was high enough to produce unacceptable levels of flight delay and reductions of throughput when the SCM strategies were implemented. This was especially evident at the higher levels of trajectory uncertainty and error. The need to account for expected levels of uncertainty becomes more and more important as the demand level increases and as the level of trajectory error increases. This is because the likelihood of flights with LOS increases as the density of operations increases and as the level of trajectory error increases.

Accounting for uncertainty in the SCM strategies improves the strategy effectiveness with respect to LOS but increases delays and reduces throughput. Thus, there is a tradeoff between scalability and allowable levels of uncertainty. That is, we can implement an air traffic management construct that allows high levels of uncertainty, but we can expect that same system to have limits on scalability that may be evident even at small demand levels, such as those used in this study. Therefore, the results in this study indicate that an air traffic system should attempt to implement mechanisms appropriate for reducing uncertainty where possible in order to increase the chances for scalability. And this increased level of predictability of operations needs to be balanced with mechanisms for ensuring flexibility when operational conditions and plans need to change, even though those types of changes should be the exception rather than the rule under normal conditions.

The study also introduced a trade space that could serve as a mechanism for selecting the appropriate SCM strategy in a trade-off between uncertainty and error, mean flight delay, and the LOS metrics (which this study equates to the potential for tactical conflict management actions). The optimal solution for a given airspace, demand level, and other factors, could likely be a combination of SCM strategies, although this study only compared the use of a single SCM strategy at a time. The CR strategy appeared to have the best opportunity for scalability by limiting the number of potential tactical actions to nearly zero.

1 Introduction

The nation and the world are currently seeing a drive to enable emerging aviation markets. These markets, categorized broadly as Advanced Air Mobility, intend to serve areas of the aviation industry not currently served, or underserved, by today's aviation operations. Urban Air Mobility (UAM) is one example of these markets, where small, electric or hybrid propulsion vehicles with the ability to takeoff and land in short distances (and typically with vertical takeoff and landing capabilities) can transport small numbers of passengers, or cargo, around an area of interest, typically a city, in shorter periods of time than current transportation methods. The 'urban' nature of this mode is related to the ability to ease traffic congestion (or provide quicker travel times within that congestion) around a major city, but the mode could have applications in other areas of the country not centered on an urban area. In addition, urban areas often feature highly complex airspace that any new air traffic management mechanisms for these emerging markets must be able to support.

UAM is a mode of transportation that will increase the complexity of the airspace, in terms of numbers of aircraft over a given area, to a scale that is much greater than can be seen in the airspace today. This level of complexity and scale introduces the need to expand or augment today's air traffic system to support these new operations.

The proposed UAM air traffic management architecture follows, in part, the same federated and service-oriented architecture developed by NASA and industry for small, unmanned aircraft. Known as the Unmanned Aerial System (UAS) Traffic Management system [1], the architecture creates an environment where vehicle operators can take on more responsibility for the coordination and management of their operations through shared airspace resources, while potentially alleviating the need to add to the workload of today's air traffic managers (e.g., air traffic controller and traffic flow managers). This architectural influence can be seen in NASA's vision document for UAM [2] as well as the Federal Aviation Administration's NextGen UAM Concept of Operations v1.0 [3].

Strategic Conflict Management (SCM) must be a part of any air traffic system. SCM is focused on enabling orderly flow of operations through the shared airspace resources. In today's airspace system, SCM is achieved through many mechanisms, such as traffic initiatives, airspace structure, scheduling, etc. In this study, we look at a few candidate mechanisms that can be leveraged for SCM with UAM operations. We compare those mechanisms in their ability to reduce the need for conflict management actions that may be required beyond the SCM layer.

2 Background

The International Civil Aviation Organization (ICAO), in their Global Air Traffic Management Operational Concept [4], defines conflict management as: "[t]he function of conflict management will be to limit, to an acceptable level, the risk of collision between aircraft and hazards." They further reduce the conflict management layers into strategic conflict management, separation provision, and collision avoidance and indicate that the conflict management process can be applied anywhere along the conflict horizon, from flight planning all the way to flight execution time. In addition, their definition of strategic is used to mean in advance of tactical, where the last two layers of conflict management together comprise the tactical part of conflict management. SCM can be applied both pre-departure, as well as during the flight operation. In relation to UAM, the ICAO conflict management definitions serve as a useful starting point and research will dictate what applies, what does not apply, or what needs to change to support these new types of operations or relevant and emergent technologies, procedures, or rules.

In ICAO's description of conflict management, a conflict is "any situation involving aircraft and hazards in which the applicable separation minima may be compromised." Given this definition, a conflict is a predicted loss-of-separation (LOS) using the applicable separation minima, which may be different depending on whether the conflict is with another aircraft, with weather, with restricted airspace, with a structure, or any other real or virtual object (i.e., an airspace region could be considered a virtual object). It is a 'predicted' LOS, because the objective is to use conflict management strategies to try to prevent the actual LOS. We also note the use of the qualifier 'applicable' on separation minima, which indicates that this separation minima may be different for a variety of conditions.

In today's air traffic system, there exist established separation minima that air traffic controllers must apply under various conditions [5]. Most of these apply to radar separation under Instrument Flight Rules (IFR) as well as to procedural separation, such as on arrival and departure. Visual Flight Rules (VFR) are an area where the separation minima are not defined because the rule is for the pilot to maintain visual separation, although rules do exist for controllers to use when separating IFR traffic from VFR traffic. The appropriate separation minima for UAM have not yet been determined and are a research topic in NASA's UAM Research Roadmap [6]. If the operations fall purely under IFR or VFR, then those rules apply, but there may be an opportunity to define new separation criteria as appropriate if the designated separator changes from the air traffic controller and the pilot to another entity, or if automation is certified to aid or assist the designated separator in maintaining separation. An example of this would be in the area of small unmanned aerial vehicles, where well-clear definitions have been identified to support the requirement that these operations remain well clear of each other [7]. In this study we are not defining alternate separation minima, but we are using alternate separation criterion (a separation volume) in defining what constitutes a LOS. As described here, this loss could very well be equated to a loss of well-clear that an assistive detect-and-avoid system is using to ensure operations remain well clear of each other. In addition, as will be described later, the separation volume that defines the LOS region for this study is largely driven by the constraints of the airspace construct used in the study.

We should note that, often times, terminology may only be applicable in a certain context, as is the case with the commonly used term in air traffic management: a scheduling conflict. In a scheduling conflict, the applicable constraint is the scheduling interval at the scheduling point and the conflict exists when one aircraft's estimated time of arrival at that scheduling point is predicted to be less than the scheduling interval from the previous aircraft's estimated or scheduled time of arrival. We note that scheduling intervals can sometimes include a design buffer to support uncertainty of operations. We also note that this type of conflict is not the same as a predicted LOS, but rather a predicted violation of a scheduling interval, thus the term 'conflict' here must be taken in context.

Within the Global Air Traffic Management Operational Concept document [4], ICAO also defined the aim of SCM and general mechanisms used in SCM. SCM's objective is to reduce the need to apply the second layer of conflict management (separation provision) to an appropriate level. Note that this does not mean that a separation provision layer is not required. A tactical layer is always required and serves to supplement the SCM layer. Per ICAO, SCM is achieved through airspace organization and management, demand and capacity balancing, and traffic synchronization.

Airspace organization and management establishes airspace structures, procedures, and processes that help ensure orderly use of the airspace. With airspace design, for example, pre-defined arrival and departure tracks into and out of aerodromes, or other pre-defined route structures, help to reduce aircraft interactions. Similarly, airspace access may implement specific hours of access, or minimum capability requirements, for given airspace regions.

In demand and capacity balancing, various scheduling or management approaches can be used to establish predictable allocation of resources. Resource scheduling, for example, can establish required or nominal sequences of traffic into resources with appropriate allocated intervals based on resource availability and requirements. Flow rate restrictions, on the other hand, can be used to efficiently manage air traffic flows by limiting traffic to ‘manageable’ levels with the expectation that tactical management actions resolve any remaining resource constraint violations. Dynamic adjustments to capacities and scheduling constraints can be implemented to handle predicted changes in airspace or air traffic demand conditions.

Yet another mechanism for SCM is traffic synchronization, which is also designed to establish or maintain the orderly flow of traffic. Examples of this mechanism includes optimized sequencing into choke points or aerodromes, 4-dimensional trajectory control, such as interval spacing, and negotiated conflict-free trajectories.

In this paper, we describe a SCM study with the objective of understanding the impact of various SCM approaches applied to the UAM air traffic management problem. The study provides a comparative analysis of different SCM approaches. We note that the optimal SCM solution for UAM may be a combination of these and other SCM strategies, but this study only examined SCM approaches independently. The design and approach for this analysis are described in the next section, where we provide details of the modeling of the various SCM approaches, the simulation tools used, and the scenarios evaluated.

3 Analysis Description and Modeling

The objective of this study was to compare and contrast a set of SCM approaches. A simulation-based approach was used to achieve the comparison via a set of pre-defined metrics. The simulation was executed in two parts: (1) using a deterministic operational planning stage, and (2) using a Monte Carlo simulation stage. The deterministic stage implemented the SCM approaches and accounted for expected levels of uncertainty in flight execution while the Monte Carlo portion introduced actual variability in flight execution. The strategies were compared using a set of scenarios with different levels of UAM demand.

3.1 Research Question

The primary research question for this study was as follows: what are the advantages and/or disadvantages of various strategic conflict management approaches? The approaches analyzed were strategies for demand capacity balancing and traffic synchronization and are described in more detail later in this document. The UAM Airspace Research Roadmap’s [6] Strategic Conflict Management Services component describes the need to implement demand capacity balancing services and traffic synchronization services but does not go into detail on what those exact service functions should be. This study attempts to provide concrete examples of those mechanisms and provides some initial comparison metrics that help validate or inform the next level of those roadmap requirements.

Considering the ICAO’s description of the objective of SCM (to reduce to an acceptable level the need to apply separation provision), we compared and contrasted each strategy in its effectiveness in reducing the need for separation provision by measuring the unmitigated LOS. In addition, we studied the impacts to delays or throughput for the selected demand scenarios in the presence of each SCM strategy. We also looked at the effectiveness of each strategy as a function of the different traffic demand levels. Finally, we analyzed the effectiveness of each strategy in the presence of trajectory uncertainty and trajectory following accuracy.

3.2 Analysis Design

The analysis involved simulating the operational planning of a set of UAM flights and the execution of those flight operations in the presence of trajectory error relative to each flight’s established operational plan. The SCM approaches implemented included the airspace design and constraints (common to all scenarios) and four different resource coordination strategies: resource scheduling, resource flow rate restrictions, area-based flow rate restrictions, and conflict detection and resolution. Each of these strategies are described in detail in later sections. Scenario variations were run without any of these four SCM strategies to create baseline conditions for comparison. The UAM traffic demand was simulated at three different levels. Trajectory uncertainty was accounted for in planning at 4 different levels and trajectory following error was also modeled at 4 different levels.

The independent variables for this study and their variations are show in Table 1. Note that each of the variables and its levels is described in the following sections of the paper and are shown here for the sake of introducing the various scenario variations for the analysis. The full factorial of this analysis would result in 240 different scenario variations. After removing variations that are less useful (such as assuming a non-zero level of trajectory uncertainty in planning but then assuming perfect trajectory execution), the study featured 84 scenario variations sufficient for studying the impact of each SCM strategy, the different demand levels, trajectory planning uncertainty, and trajectory execution error.

Table 1. Independent variables and levels.

Variable	Levels
SCM Approach	None (base), resource scheduling (RS), resource flow rates (FR), area-based flow rates (AR), conflict detection and resolution (CR)
Demand Level	Level 1 (D1), level 2 (D2), level 3 (D3)
Trajectory Planning Uncertainty	None (U0), level 1 (U1), level 2 (U2), level 3 (U3)
Trajectory Execution Error	None (E0), level 1 (E1), level 2 (E2), level 3 (E3)

3.3 Simulation Tools

The primary analysis tool used for this study was the Advanced Trajectory Services – Toolkit for Integrated Air and Ground Research (ATS-TIGAR). ATS-TIGAR, described in [8], is a research and development environment suitable for prototyping various air traffic management functions and algorithms. ATS-TIGAR implements the UAM Mission Planner algorithm [8], which serves to model UAM operational planning, as well as a suite of libraries and functions suitable for other fast-time scenario analysis. At its core, the UAM Mission Planner implements demand management, fleet management, flight planning, constraint management, vertiport coordination, and fleet operator coordination, among other functions. Together these functions allow the Mission Planner to plan the flights required to meet a given UAM mission (group of passengers going from one vertiport to another) in the presence of fleet operator, fleet vehicle, vertiport, airspace, and other constraints. Multiple Mission Planner algorithms can be used simultaneously to represent a set of UAM fleet operators that must coordinate their operations at shared airspace resources. The UAM Mission Planner was expanded to implement the SCM strategies for this study.

A Monte Carlo algorithm in ATS-TIGAR was used to simulate the planned operations in the presence of trajectory errors. Each of the scenario variations with trajectory execution error was replicated 1000 times using this Monte Carlo approach.

3.4 Scenarios

3.4.1 Airspace

The scenarios for this analysis leveraged a research representation of a possible UAM airspace construct for the Dallas/Fort Worth area. The airspace construct, seen in Figure 1, was developed as part of NASA's airspace research activities [9] for the purposes of the X4 simulation. In this construct, portions of the Class B airspace with very low levels of air traffic were identified as potential useable airspace regions for UAM operations where interactions with existing traffic could be minimized. A set of corridors with directional tracks were inserted into these useable airspace regions to provide the first layer of SCM for the Dallas airspace region. A set of 34 vertiports were selected, using engineering judgement, from a set of demand scenarios created for NASA by Virginia Tech in prior work [10].

The majority of the UAM tracks within the corridors of this airspace design were spaced 1500 feet apart for tracks of opposite direction traffic. This spacing was chosen under the assumption of aircraft navigating a track with a width of 1200 feet (+/- 600 feet from the track centerlines) under nominal operations, plus providing an additional 300 feet of buffer between the tracks. We recognize that this spacing is somewhat arbitrary and preliminary, but it does introduce a constraint that guides the selection of SCM parameters for this study as will be described later. Namely, it leads to a separation volume that must be consistent with this track spacing, with a lateral separation requirement between operations of 1200 feet.

3.4.2 Demand Levels

The demand levels selected for this study originated from a single representative UAM demand scenario. The source scenario does not represent an actual demand prediction but, rather, a randomized generation of demand targeting 10 operations per vertiport per hour. This source scenario is assigned to the middle demand level for this analysis (demand level D2) and represents a 2-hour period of traffic consisting of 44 simultaneous operations at peak time and a total of 339 flights over the scenario duration.

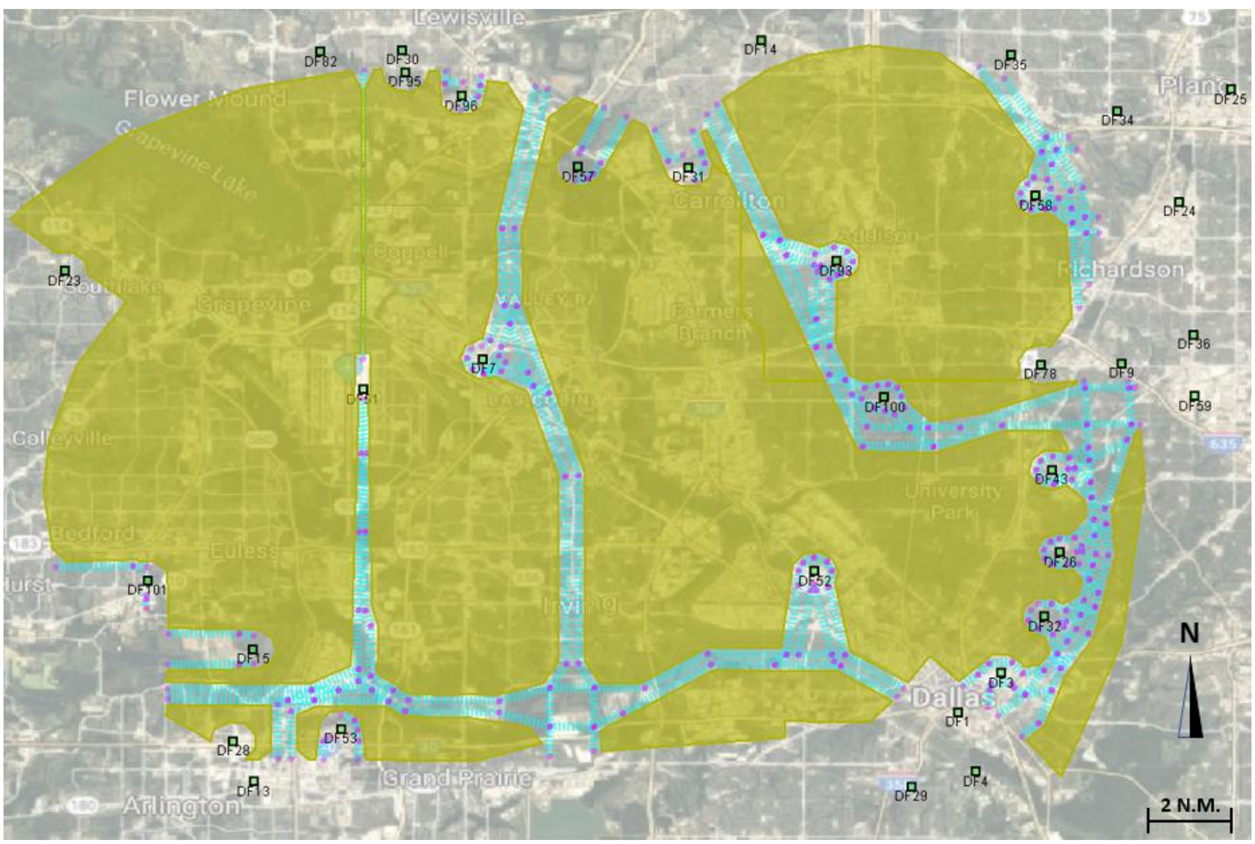

Figure 1. Dallas/Ft. Worth airspace showing UAM corridor tracks (cyan), vertiport locations (green squares with black boundaries), and restricted airspace regions (yellow polygons) [background map credit: Google].

The two additional demand levels (D1 and D3) were derived from the demand level D2, as shown in Table 2. Demand level D1 represents a lower demand and was generated by randomly down-sampling one-fourth of the demand level D2. The resulting demand level D1 had 15 simultaneous operations at peak time and a total of 84 flights over the two-hour duration. Similarly, demand level D3 represents a higher demand level than demand level D2 and was achieved by doubling of the scenario demand level D2 with new random times assigned to the duplicate demand flights within the 2-hour scenario time window. Note that the times for the duplicated flights were not deconflicted with original demanded flights because these scenarios represent random demand. This scenario resulted in up to 89 simultaneous operations and 678 total flights. We note that these numbers of simultaneous operations represent the unimpeded demand, which changed as SCM and other constraints were applied in planning and coordination of the flights.

Table 2. UAM demand levels D1-D3.

	Peak Simultaneous Flights	Total Flights (over 2-hour period)
Demand Level 1 (D1)	15	84
Demand Level 2 (D2)	44	339
Demand Level 3 (D3)	89	678

3.5 Strategic Conflict Management Modeling

The UAM Mission Planner was expanded for this study by enabling it to adhere to additional SCM constraints. The four SCM strategies implemented are described next.

3.5.1 Resource Scheduling

The resource scheduling approach (tagged with 'RS' in scenario names and results figures) is a mechanism to ensure a defined spacing between operations at shared network resources in the airspace (i.e., network waypoints or intersections). At these key points of the airspace, such as vertiports or other common points, a spacing requirement is established as the resource constraint. An algorithm must then be used to ensure that the operations using that resource meet the required spacing. For this study, a first-scheduled, first-served approach was used. Using the scheduled time-of-arrival (STA) from previously scheduled operations, an operator with a distributed scheduler (or, equivalently, a centralized scheduler) identified the existing resource usage and evaluated the delay required to add a new or proposed operation, given that proposed operation's estimated time-of-arrival (ETA) at the same resource. That delay was then absorbed prior to arriving at that resource. In this study, the delay was implemented as a pure departure delay. We note that there are other possible strategies for delay absorption that were not explored with this study, such as pass-back delay which may be handled by speed adjustments for the flight at the various resources along the flight trajectory. The ETAs used to build the schedule became STAs once the flight operation was approved within the requirements of the UAM airspace system.

In this study, the resource scheduling SCM approach was implemented at every vertiport (34 in total), every entry and exit point to the UAM track system, and every crossing or merge point within the UAM track system (80 resource points), as seen in Figure 2. The spacing used for all vertiports was 120 seconds between operations, which was an engineering assumption that provided a capacity at each vertiport of 30 operations per hour on a single vertipad and was sufficient to support the demand levels chosen. The spacing of the entry, exit, crossing, and merge points was determined by the time required to traverse the lateral separation distance (1200 feet) at some assumed speed. In addition, the spacing was dependent on the crossing angle of the resource, thus the spacing was different at most of the resources. For example, a crossing angle of 90 degrees used approximately 16 seconds spacing whereas the in-trail time spacing for zero intersection angle would have been only around 11 seconds. More details of the spacing assumptions can be found in Appendix A.

Figure 3 shows an example of a resource timeline with three scheduled operations. In this example, the required spacing is approximately 36 seconds between operations. The spacing is the space between the operation's STAs (the back dots on the timeline) and the figure depicts a reservation window for each operation with a pre- and post-operation time that is half the magnitude of the spacing.

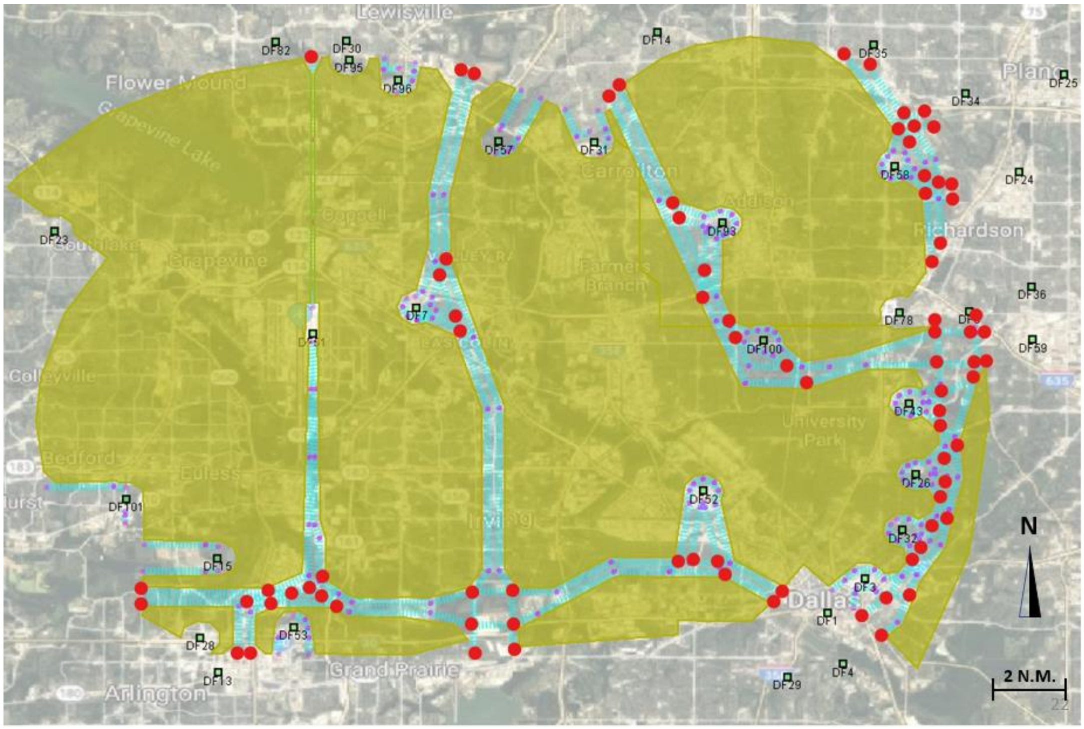

Figure 2. Map showing 34 vertiport (green squares with black boundaries) and 80 resource points (red circles) used for RS and FR SCM approaches [background map credit: Google].

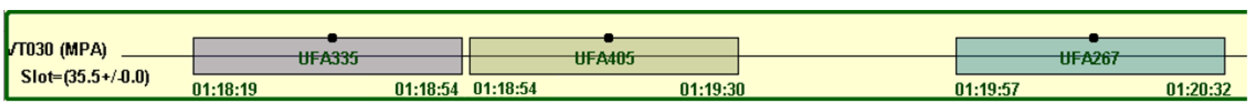

Figure 3. Example resource scheduling timeline.

3.5.2 Flow Rate Constraints

The flow rate constraints approach (tagged with 'FR' in scenario names and results figures) imposed rate restrictions at key points within the airspace. The rates were imposed as a maximum number of operations per given time interval (e.g., 10 operations per hour) where the objective was not to exceed the defined rate during strategic planning. This flow rate approach is similar to demand-capacity balancing approaches used in the existing airspace system to establish capacity limits on resources.

The implementation of the flow rate approach amounted to a binning of operations, as depicted in Figure 4. The resources were modeled using a set of time bins. Strategic planners used the STAs of already scheduled operations to assess the impact to the resource of adding a new operation. The strategic planners impose delay as required for the proposed operation in order to resolve the potential demand-capacity imbalances. In the implementation for this study, the flow rate constraints were resolved solely with departure delay.

The flow rate based SCM approach was implemented at every vertiport, every entry and exit point to the UAM track system, and every crossing or merge point within the UAM track system, as seen in Figure 2. The time bin size for this analysis was set at 12 minutes, which is somewhat arbitrary but consistent with NASA's X4 simulation work. The vertiport capacity (or rate) was set to be consistent with the RS approach, whose theoretical maximum capacity was 30 operations per hour with 120 seconds between operations, which resulted in a capacity of 6 operations per 12 minutes using the bin size chosen. For the key points in the UAM track system, the capacity (or rate) was computed based on the required time spacing between operations used in the RS approach but with a smaller, one-minute time window. As an example, crossing points at a 90-degree intersection used a capacity of 3 operations per minute while merge points with shallower intersection angles had higher rates and crossing points with higher intersection angles had lower rates. More details on this rationale can be found in Appendix A.

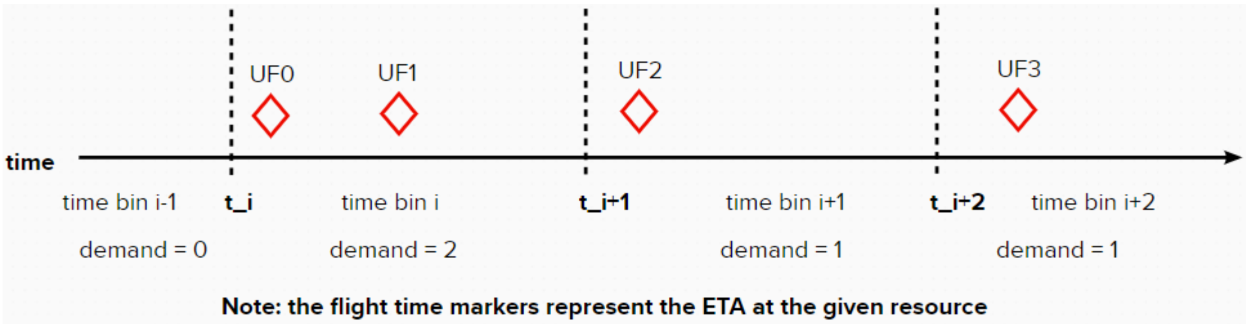

Figure 4. Example flow rate constraints time bins.

3.5.3 Area-Based Flow Rate Constraints

In the area-based flow rate constraints SCM approach (tagged with 'AR' in scenario names and results figures), an area (or volume) of airspace was defined as the resource rather than a single point in space, as in the standard flow rate approach above. The usage of this area by an operation was described as a time interval required to transit that area or volume, rather than a single ETA. The intervals can then be assessed in one of two ways: number of operations in a given time bin (as shown in Figure 5), or as the number of simultaneous operations (can be thought of abstractly as an infinitesimally small bin). The capacity or rate limit, therefore, was specified as the number of operations per unit time, or as maximum number of simultaneous operations.

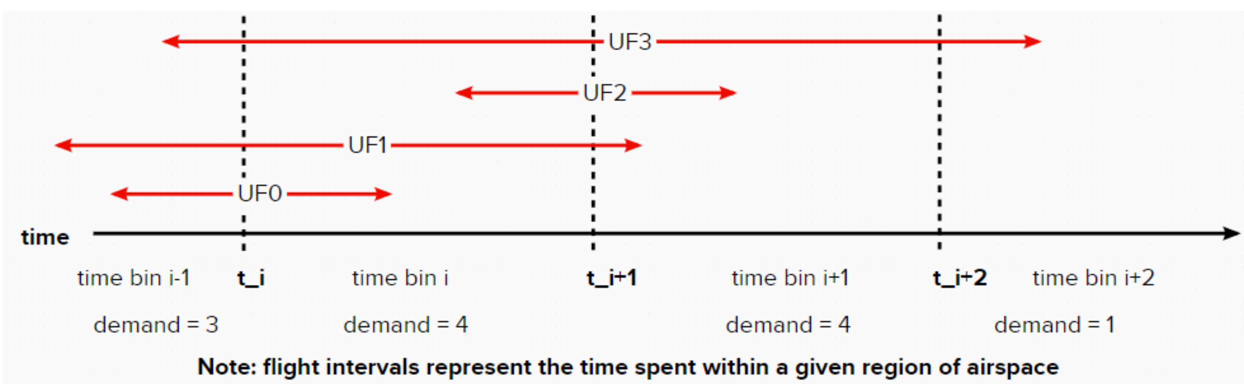

Figure 5. Example area-based rate constraints time bins with operation intervals.

In the implementation for this study, both number of operations per time, and number of simultaneous operations, were used depending on the resource. The binned approach was used for vertiports, with no more than one operation per 120-second time window, which is similar to, but slightly more conservative than, the RS SCM approach with 30 operations per hour of capacity. This approach was used because the interval portion of a trajectory between takeoff and the exit of an area around the vertiport was observed during scenario development to be just over 30 seconds. This short time interval would lead to a vertiport capacity approaching 120 operations per hour, which would not be a fair comparison with the other strategies. Realistically, one would want to account for additional surface time prior to takeoff as well as immediately after landing.

The area polygons for this SCM strategy we selected to overlap the vertiports as well as the UAM tracks. The 34 vertiports were enveloped in a hexagonal region with vertices that were 1386 feet from the centroid and sides at 1200 feet from the centroid. The choice of 1200 feet is derived from the effective lateral separation criterion, which was set to this value. The UAM tracks, however, were enveloped in polygon regions that extended 600 feet on either of the track and were discretized into 2 nautical mile long segments or smaller. A total of 137 polygon regions covered the UAM track network, as can be seen in Figure 6. The capacities for each of the track polygons were set based on the worst-case track crossing angle of points within a given region and the number of operations that could be supported within that region. For example, for a track polygon with a 90-degree crossing intersection inside of it, the capacity was set to a maximum of 3 simultaneous operations. More details on this rationale can be found in Appendix A.

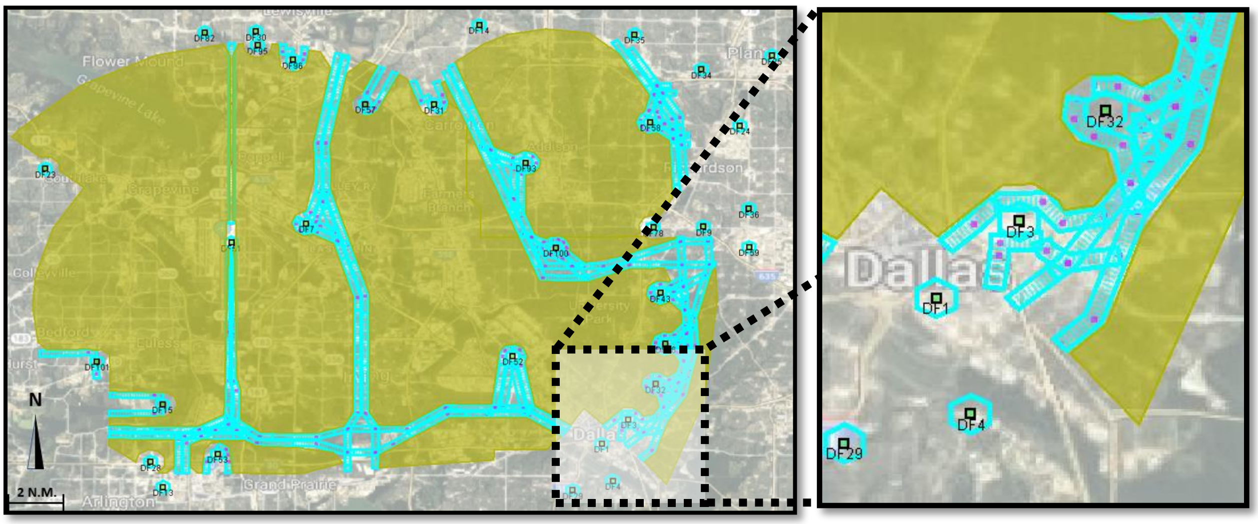

Figure 6. Area-based SCM approach polygons on 34 vertiports and UAM tracks (137 segments); all of Dallas area (left) and detailed downtown area (right) [background map credit: Google].

3.5.4 4-D Trajectory Conflict Detection and Resolution

The conflict detection and resolution approach (tagged with 'CR' in scenario names and results figures) used 4-dimensional trajectory deconfliction to resolve anticipated conflicts between flight trajectories through trajectory modifications. The proposed operational intent trajectory for a given flight was evaluated against other planned trajectories to identify trajectory conflicts. A set of strategies was used to effectively deconflict the proposed trajectory from the other trajectories using speed changes, vertical maneuvers, lateral maneuvers, and departure delay. In the region of UAM tracks, where the space for maneuvers was limited, trajectory modifications were limited to speed resolutions and departure delay. In regions outside of the UAM track system, all resolutions were enabled. Departure delay was the strategy implemented when other strategies failed to provide a solution due to the difficulty of the encounter, for example. The effective separation criteria used in this SCM approach were defined as 1200 feet lateral and 450 feet vertical, primarily based on the spacing between UAM tracks. A violation of these separation criteria constitutes a LOS in the metrics for this study.

The CR SCM approach was disabled when flights were within the lateral separation distance from the vertiports. That is, trajectory conflicts within this distance from the vertiport were ignored because this would be an area where procedural separation would be required. In order to be consistent with the vertiport constraints of the other SCM strategies, resource schedulers with 120 second slots were introduced at all vertiports alongside the CR SCM approach.

The CR strategy was implemented as a pre-departure deconfliction mechanism. It is worth nothing that this type of SCM approach is not widely used in today's air traffic system but provides a potential opportunity in the UAM space. Most operations in today's air traffic system are on a scale that is much different than UAM, where operations are traveling hundreds or thousands of miles on flights with durations on the scale of hours, and often experience delays on the order of many minutes. Pre-departure deconfliction may not make sense in these instances. Conversely, UAM operations are typically considered to be under 100 miles and with flight times that range on the order of minutes. In addition, UAM operations will likely benefit from higher levels of data sharing and data availability, given the proposed federated and service-based architecture, in a way that does not happen today. Furthermore, possible implementation of functionality such as conformance requirements and conformance monitoring may provide an environment more conducive to pre-departure deconfliction. Also, all pre-departure planning with any SCM strategy must be complemented by an element of continuous strategic re-evaluation due to the need for re-coordination should the conditions change before the flight departs (e.g., for vertiport closures, flight cancellations, etc.). It is the recommendation of the authors of this report that this should be a requirement for any SCM strategy.

3.6 Trajectory Planning Uncertainty and Trajectory Execution Error

Trajectory planning uncertainty refers to the amount of uncertainty that was used when planning the trajectories for each flight. The intent was for the uncertainty to account for the expected level of error between the planned trajectory and the as-flown trajectory. Therefore, planning uncertainty and execution error are connected and, ideally, one would account for a level of uncertainty consistent with the observed errors in flight operations.

Trajectory execution errors were assumed to follow zero-mean normal distributions. The error components modeled were lateral, vertical, and temporal error. Table 3 shows the characteristics of the uncertainty levels and error levels modeled in this study. Within Table 3, the mean of the distribution is represented by μ and the standard deviation of that same distribution is represented by σ. The selection of these error magnitudes was driven partly by the available spacing buffer between the UAM tracks (300 feet). The largest lateral magnitude $\sigma = 75$ feet assumes that the aircraft would be able to follow a given UAM track to within +/- 150 feet (half of the spacing allotted to each of the opposite direction tracks) approximately 95% of the time. The largest vertical magnitude assumes the aircraft can maintain altitude with +/- 50 feet approximately 95% of the time. Finally, the largest temporal error magnitude assumes that the aircraft could maintain +/- 40 seconds relative to the nominal planned trajectory approximately 95% of the time. The medium and low error levels were selected using engineering judgement in order to create other simulated and tighter levels of performance. Whether these performance levels could be achieved in practice is unclear given the lack of data for new vehicle types and is beyond the scope of this study; nonetheless, they serve as useful points of reference.

Table 3. Trajectory uncertainty (U1-U3) and execution error (E1-E3) means and standard deviations.

Level	Temporal Error (sec)	Lateral Error (ft)	Vertical Error (ft)
U1, E1 (low)	$\mu = 0$; $\sigma = 1$	$\mu = 0$; $\sigma = 10$	$\mu = 0$; $\sigma = 2$
U2, E2 (medium)	$\mu = 0$; $\sigma = 5$	$\mu = 0$; $\sigma = 25$	$\mu = 0$; $\sigma = 10$
U3, E3 (high)	$\mu = 0$; $\sigma = 20$	$\mu = 0$; $\sigma = 75$	$\mu = 0$; $\sigma = 25$

Trajectory uncertainty was implemented by selecting a 2σ magnitude from the expected error. For the RS SCM approach, for example, only the temporal uncertainty was a factor. Given the magnitudes of Table 3 for the high uncertainty level, the ETAs of each flight were assumed to have +/-40 seconds of uncertainty, which translated to an additional 80 second spacing between operations on top of the already established spacing requirement. For example, a 90-degree crossing intersection changed from a 16 second spacing to a 96 second spacing in order to support this level of uncertainty. The FR SCM approach accounted for ETA uncertainty in a similar way, which meant that an operation could potentially be scheduled into multiple bins due to the time uncertainty, thereby reducing the capacity of the resource. For the AR SCM approach, the time intervals for an operation inside of a given area were also buffered to assume a longer interval consistent with the 2σ level of uncertainty. Finally, the CR SCM approach accounted for uncertainty by assuming the effective separation volume was increased by twice the lateral and vertical error magnitudes (to account for ownship and traffic with similar levels of uncertainty) and the trajectories were evaluated for conflicts at the nominal times as well as various forward and backward time-shifts for both traffic and ownship to account for the temporal error. Resolutions were required to resolve against these uncertain trajectories, therefore covering more airspace than nominal resolutions with no uncertainty.

Trajectory execution error was implemented in the second part of the analysis, where the nominal planned trajectories were given a simulated error according to the error levels of Table 3. Each of the component errors was implemented using the standard sinusoidal pattern given by the following equation:

$$e(t) = e_{max} * \sin\left(2 * \pi * \frac{t - t_0}{T} + \varphi\right) \quad \textbf{\textit{(1)}}$$

Here, $e(t)$ is the amount of error (lateral, vertical, or temporal) at any given trajectory time, t, in seconds, e_{max} is the maximum magnitude of error sampled from the normal distribution, t_0 is the first trajectory time, in seconds, T is the period for the error, in seconds, and φ is the phase of the error, in radians. The phase was sampled from the uniform distributions while the period was computed using the equations of Table 4. These were derived under assumptions described in Appendix B and were implemented with the values of: $GS_{nom} = 130\ kts$, $\Delta GS = 10$ knots, $VS_{nom} = 500\ ft/min$, $\Delta VS = 250$ feet/min, and $LateralRateRatio = 2$.

Table 4. Trajectory error period equations and phase distributions.

Category	Period (sec)	Phase (radians)
Temporal Error	$T = \frac{GS_{nom} * 4 * e_{max}}{\Delta GS}$	φ = [-π/2 : π/2]
Lateral Error	$T = LateralRateRatio * 4 * e_{max}$	φ = [-π/2 : π/2]
Vertical Error	$T = \frac{VS_{nom} * 4 * e_{max}}{\Delta VS}$	φ = [-π/2 : π/2]

Figure 7 shows an example of one of the trajectories with the implemented trajectory errors. The left image shows the lateral profile of the trajectory with sinusoidal error superimposed on the nominal trajectory. Similarly, the right image shows the vertical profile of the trajectory with altitude oscillations. We note that the temporal error can also be seen in the right image, where the descent from 1600 feet to 1100 feet altitude happens slightly later in time for this particular trajectory, indicating the flight is late in comparison to the coordinated operational plan at this point in the flight.

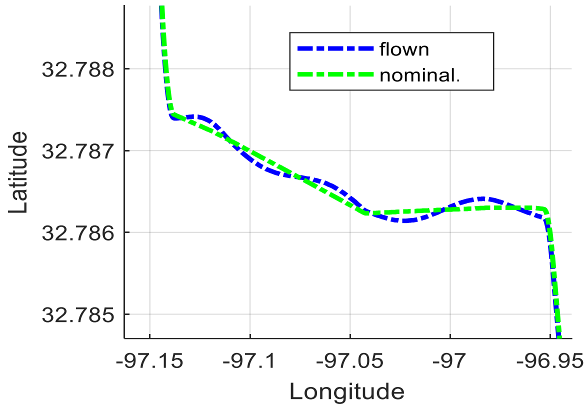

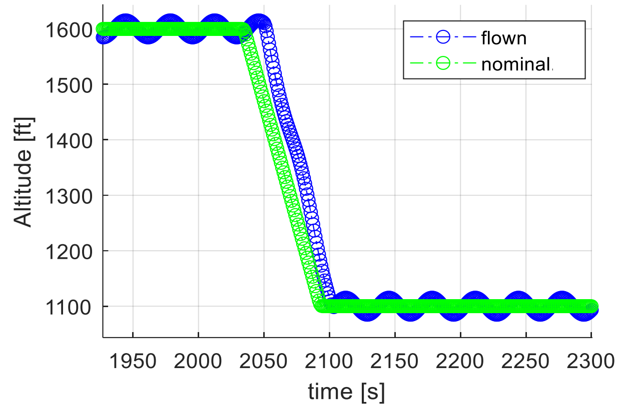

Figure 7. Sample nominal and flown trajectory; lateral profile (left) and vertical profile (right).

3.7 Metrics

Table 5 lists the metrics collected in this SCM study. The number of flights for each scenario remained constant and equal to the input number of trips because every flight was allowed to receive whatever amount of delay was required without cancellation. In addition, the model constraints that would typically results in additional flights for fleet movements were disabled. The number of flights as a function of time provided the throughput of the given airspace under a given SCM strategy and other scenario assumptions.

Table 5. Metrics collected in this SCM study.

Metric	Description
NumFlights	The total number of flights planned in a scenario.
Throughput	The number of observed operations as a function of time (e.g., number of airborne flights, number of departures). May be measured over the whole airspace as well as for various specific resources.
NumLOS	The total number of unmitigated losses of separation for a given scenario.
NumFlightsWithLOS, PercentFlightsWithLOS	The number and percentage of flights with one or more losses of separation.
NumFlightsWithDelay, PercentFlightsWithDelay	The number and percentage of flights with non-zero delay with respect to their unimpeded plan.
MeanDelay, MeanDelayOnlyDelayedFlights	The average delay over all flights and only for flights that had non-zero delay.

The primary measure for assessing the need for action from the tactical conflict management layer was unmitigated LOS. The effective separation criteria used were defined as 1200 feet lateral and 450 feet vertical and a violation of these separation criteria constituted a LOS. The overall number of LOS provides one means for comparison, but the number of flights with LOS, number of flights with delay, and mean delay per flight were other key metrics for comparison of the SCM approaches.

As mentioned previously, the metric of LOS could equivalently be thought of as a loss of well clear if a loss of well clear is what constitutes the unmitigated conflict, especially given the assumption of a separation volume in this study that does not match existing separation minima, or in a visual separation environment.

4 Results

The results of this analysis are presented in two main sections below. First, we describe the characteristics of the scenarios and the characteristics of the SCM strategies in the absence of trajectory uncertainty and trajectory execution error. These runs provide an effective baseline comparison for the second part of the analysis, where we introduce this uncertainty and error.

In the results to follow, the scenario labels follow the convention shown in Figure 8. The scenario name begins with a 'scn_' label, followed by the SCM strategy shorthand (base, AR, CR, FR, or RS), then the demand level (D1, D2, or D3), next the uncertainty level (U0, U1, U2, or U3), and finally the trajectory error level (E0, E1, E2, or E3). Note that, for brevity, sometimes the U0 and E0 labels are omitted, and no uncertainty or error is accounted for in those scenarios and associated data.

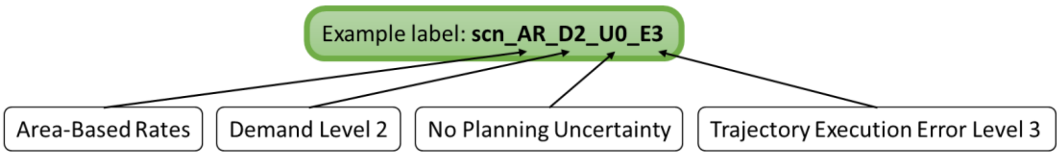

Figure 8. Scenario labeling convention example.

4.1 Baseline Cases – No Uncertainty or Error

Figure 9 shows the unimpeded throughput in terms of simultaneous airborne flights for the three demand levels prior to implementing any of the SCM strategies. The figure shows the 2-hour period of demand, which does not begin until about 15 minutes into the scenario due to planning time, vehicle surface time, and other model assumptions. There is also a ramp-down time for flights that are still airborne at the end of the 2-hour period, making the overall nominal period of demand approximately 2.25 hours long.

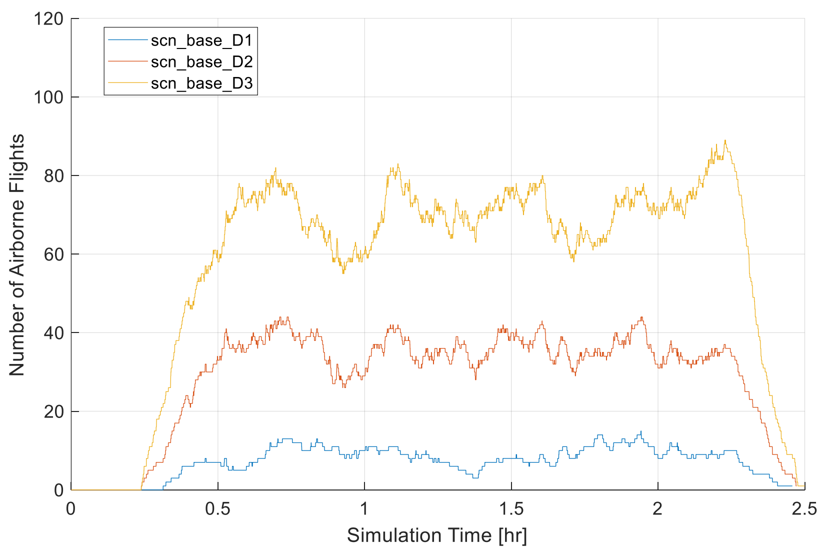

Figure 9. Unimpeded flights for the 3 demand levels.

Figure 10 through Figure 12 show the results of implementing the SCM strategies without accounting for any uncertainty. With demand levels 1 and 2, the number of airborne flights does not appear to be significantly impacted by the use of the SCM strategies - the throughput is approximately the same. This would indicate that the delay required to implement the strategies is relatively small, which can be confirmed in the mean delay figures. Demand level 1 resulted in delay for less than 18% of flights with mean delays under 2 minutes, and the FR strategy produced the least amount of delay (<1 minute) and the smallest percentage of flights with delay when compared with other strategies (~1%). Demand level 2 resulted in delay for nearly 55% of flights with the mean delay staying under 3 minutes, and the FR strategy again had the least amount of delay and percentage of delayed flights.

Demand level 3 produced considerably more delay when implementing the same SCM strategies, and the mean delays were much higher than the two lower demand levels. Some impact can be observed in the throughput, where the number of airborne flights was slightly lower than the unimpeded levels at times. In some cases, the demand peaks can be seen to shift to the right, as can be seen with the CR SCM strategy. Nonetheless, the CR and FR SCM strategies appeared to be able to support the demand within the same overall time period as the unimpeded demand (~2.25 hours), whereas, the RS strategy required ~2.6 hours and the AR strategy required ~2.9 hours to meet the same demand. The mean delays exceeded 10 minutes with the AR SCM strategy, and most strategies required nearly 89% of flights to implement delay. The FR strategy produced the lowest mean delays and lowest percentage of flights with delay due to its less constraining nature as compared to the other strategies.

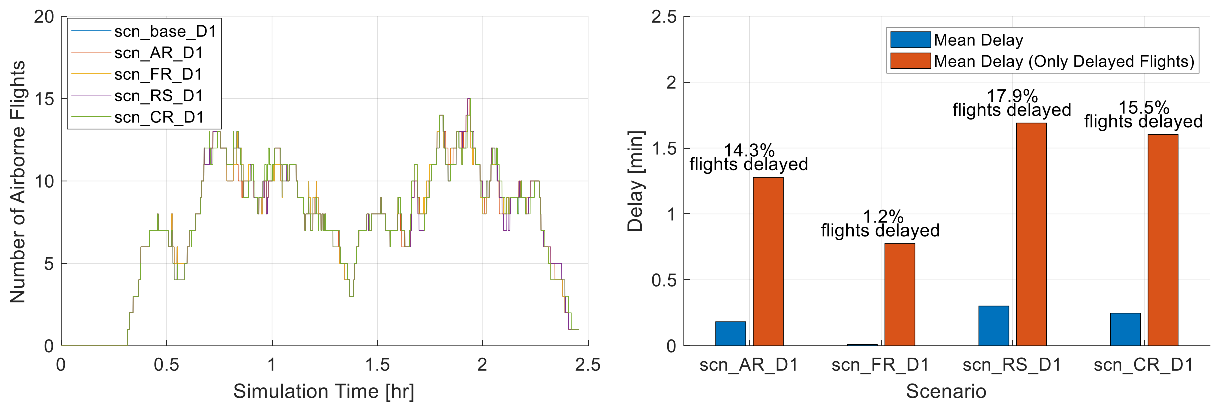

Figure 10. Number of airborne flights as a function of time (left) and mean flight delays (right) for demand level 1 under each SCM strategy, without trajectory uncertainty.

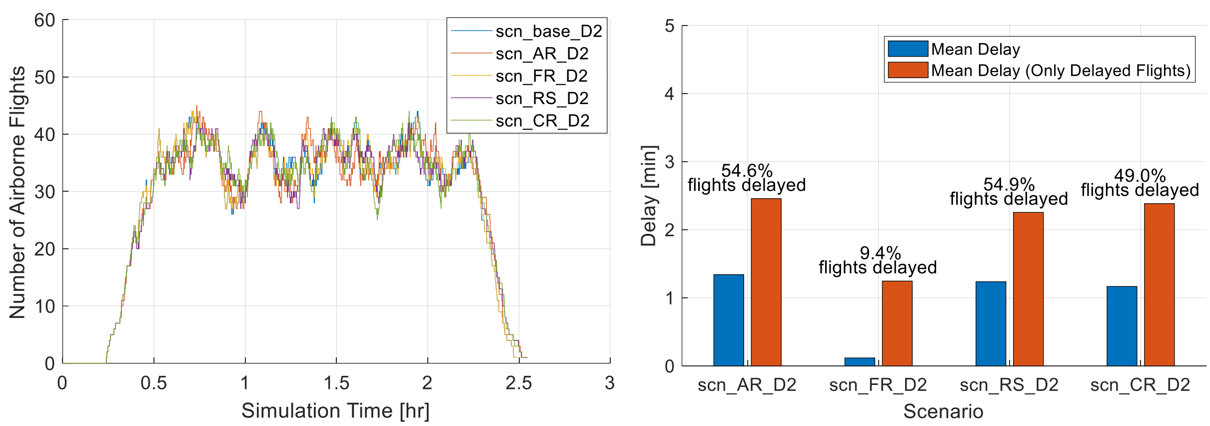

Figure 11. Number of airborne flights as a function of time (left) and mean flight delays (right) for demand level 2 under each SCM strategy, without trajectory uncertainty.

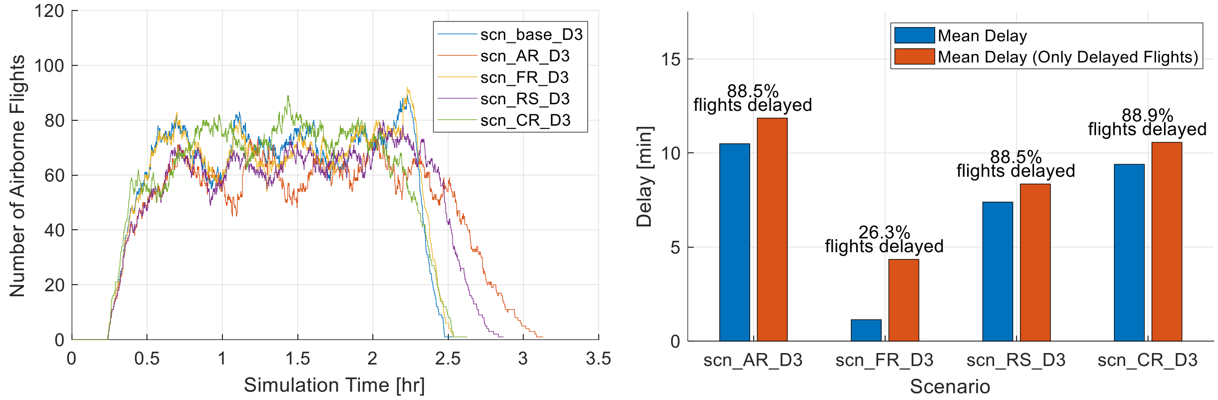

Figure 12. Number of airborne flights as a function of time (left) and mean flight delays (right) for demand level 3 under each SCM strategy, without trajectory uncertainty.

4.2 SCM With Trajectory Uncertainty and Error

4.2.1 Throughput

Figure 13 through Figure 18 show the throughput impacts of the various SCM strategies under the different demand levels and the various levels of uncertainty. At demand level 1 (Figure 13 and Figure 14), there was little impact to the throughput due to uncertainty. This indicates that, at this demand level, any of the strategies are easily capable of coping with these various levels of uncertainty because the airspace has not reached a saturation point where the uncertainty starts translating into significant levels of delay.

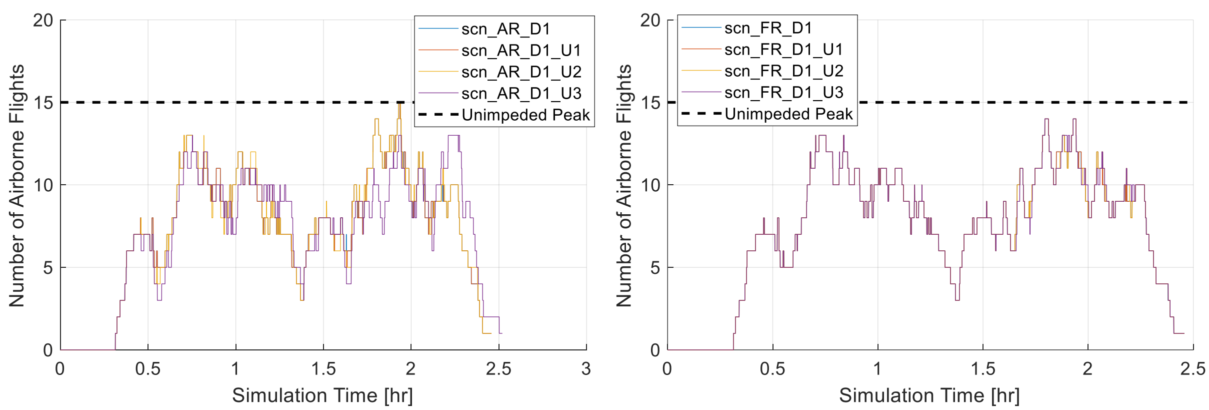

Figure 13. Number of airborne flights as a function of time for the AR strategy (left) and the FR strategy (right) for demand level 1, with trajectory uncertainty.

We observed that the uncertainty at demand level 2 (Figure 15 and Figure 16) does have some impact on throughput, but mostly at the highest uncertainty level. With all SCM strategies, we saw flights that go beyond the nominal 2.25 hours of the scenario due to delays at the highest uncertainty level. This indicates that significant amounts of delay were being imposed on flights due to this high level of uncertainty. We also saw that the AR, RS, and CR strategies experienced reduced throughput at the highest level of uncertainty within the nominal 2-hour period of demand, which is a clear indication of an over-demand situation (the airspace is not able to absorb this much demand under the given constraints). This over-demand, and excess delay, condition was also seen in the FR strategy, but the effects appear to be less pronounced.

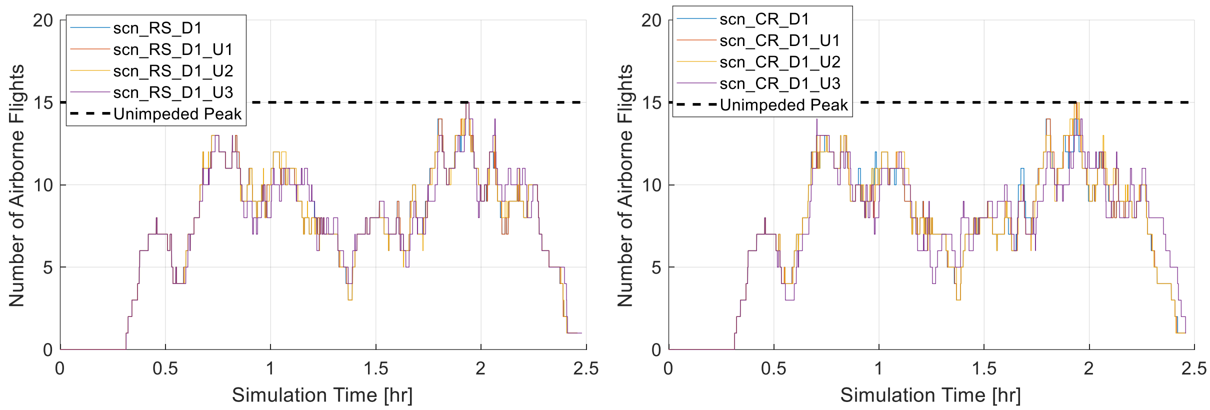

Figure 14. Number of airborne flights as a function of time for the RS strategy (left) and the CR strategy (right) for demand level 1, with trajectory uncertainty.

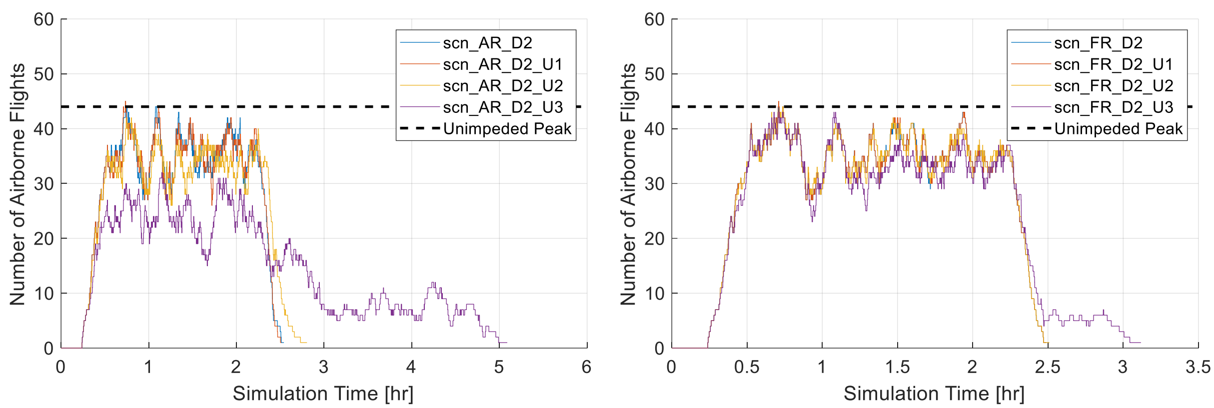

Figure 15. Number of airborne flights as a function of time for the AR strategy (left) and the FR strategy (right) for demand level 2, with trajectory uncertainty.

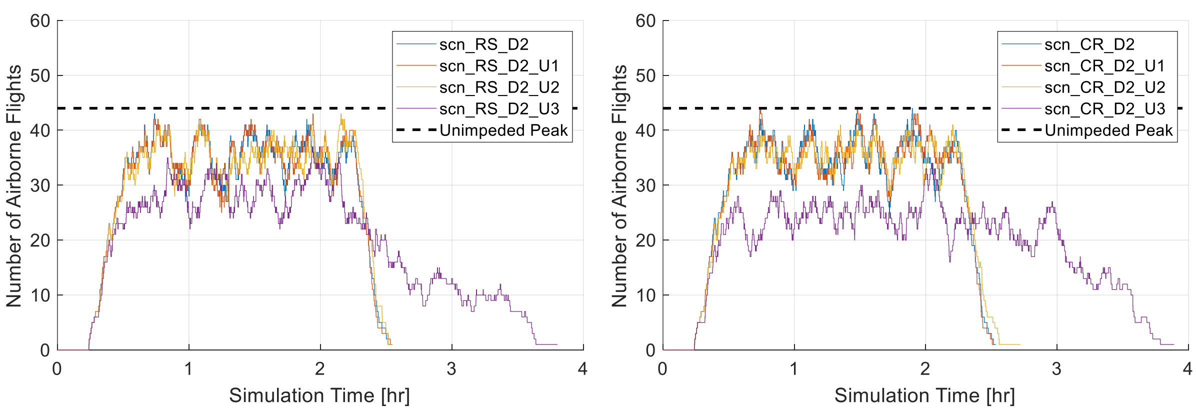

Figure 16. Number of airborne flights as a function of time for the RS strategy (left) and the CR strategy (right) for demand level 2, with trajectory uncertainty.

At demand level 3 (Figure 17 and Figure 18), we observed that the impacts of uncertainty on throughput are even more pronounced. The FR strategy showed the least impact to throughput, with little to no impact to throughput in the absence of uncertainty. As we increased the uncertainty, the peak number of flights was nearly unchanged at uncertainty levels 1 and 2 but was reduced to a peak of 69 for the uncertainty level 3. In addition, at uncertainty level 3, the scenario required an additional 3 hours of simulation time in order to support the demand. We note that the part of the throughput curve with smaller magnitudes at hours 2.5 to 5.5 as compared to the initial 2 hours is an indication that there were specific parts of the airspace constraints that had more demand than others.

In comparison to the FR strategy, the AR strategy had a throughput reduction from the peak 89 simultaneous operations to between 60 and 80 simultaneous operations for the strategy under no uncertainty and uncertainty levels 1 and 2. Under uncertainty level 3, the peak simultaneous operations dropped to 47 and the period of demand increased to just over 9 hours. Even uncertainty level 2 indicated a period of demand of nearly 4 hours as compared to the nominal 2 hours.

The CR and RS SCM strategies showed similar throughput profiles under demand level 3. Uncertainty levels 1 and 2 reduced the throughput only slightly and extended the demand period less than one hour. Uncertainty level 3 reduced the peak throughput from 89 simultaneous operations to 46 in the RS strategy and 36 for the CR strategy. Both strategies had a demand period of nearly 6.5 hours, indicating significant delays were required to meet the given demand. Nonetheless, when compared to the FR and AR strategies, the overall demand period had a more uniform distribution of simultaneous operations with the RS and CR strategies (perhaps even more obvious in the CR case as compared to the RS case).

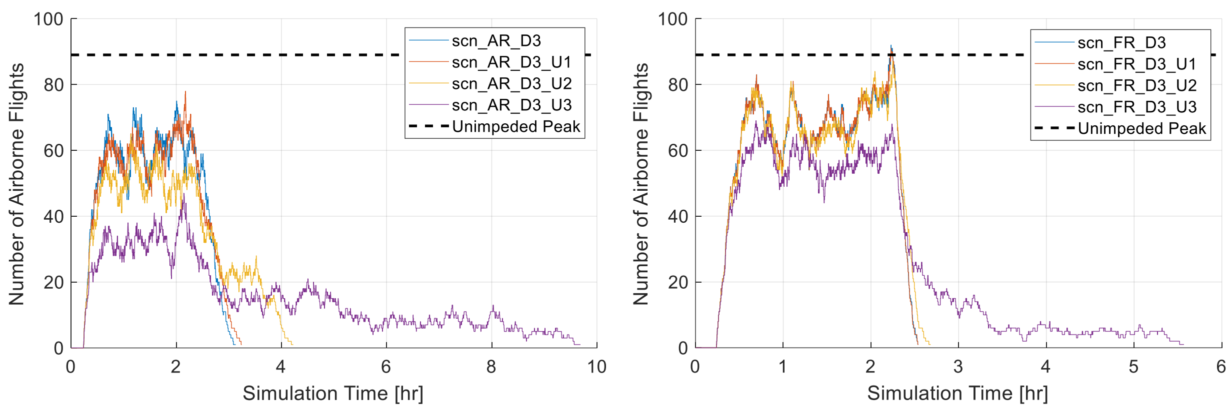

Figure 17. Number of airborne flights as a function of time for the AR strategy (left) and the FR strategy (right) for demand level 3, with trajectory uncertainty.

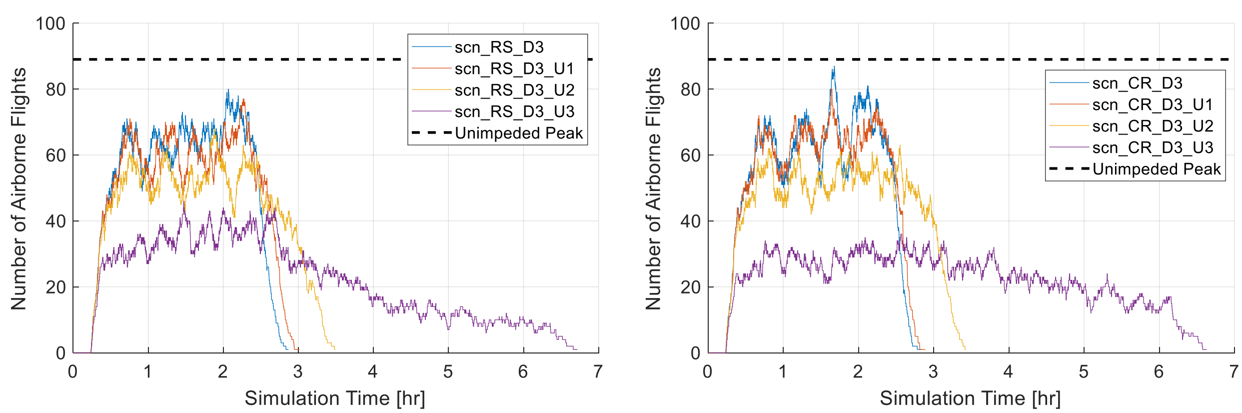

Figure 18. Number of airborne flights as a function of time for the RS strategy (left) and the CR strategy (right) for demand level 3, with trajectory uncertainty.

The inferred delays observed in the previous throughput figures can be confirmed in Figure 19 and Figure 20. These figures show the mean flight delay (only delayed flights) for the different strategies as function of demand level and by uncertainty level. The increase in delay as a function of demand, and as a function of uncertainty, can be easily compared for the 4 different strategies. As observed already, we saw that the FR strategy had the least amount of delay, both with and without uncertainty, as compared to the other three strategies. We also saw that the AR strategy had the highest delays of any strategy and that the CR strategy had only slightly higher delays than the RS strategy. In some conditions, the mean flight delays far exceeded the average flight time for flights in these scenarios (12.6 minutes), making the scenario likely unacceptable and pointing to the airspace constraints (or SCM mechanism) as unable to support the given demand under some levels of uncertainty.

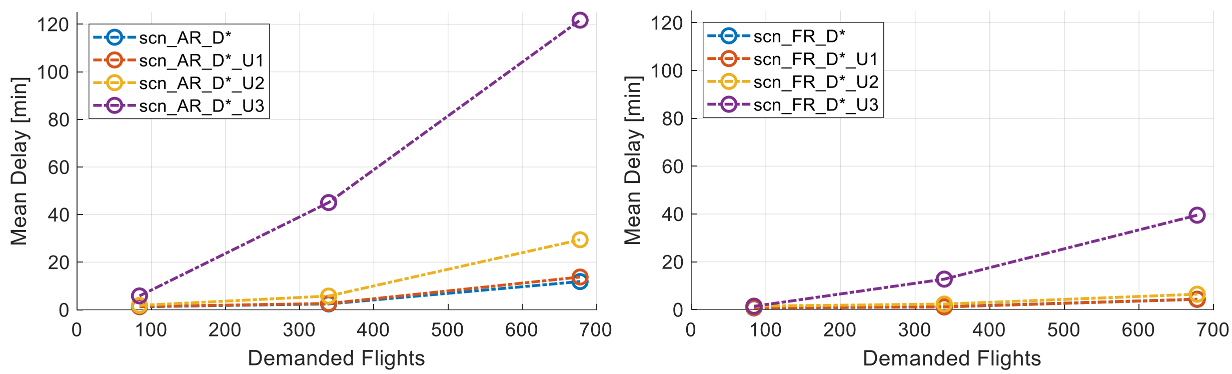

Figure 19. Mean flight delays as a function of demand level and trajectory uncertainty for the AR strategy (left) and the FR strategy (right).

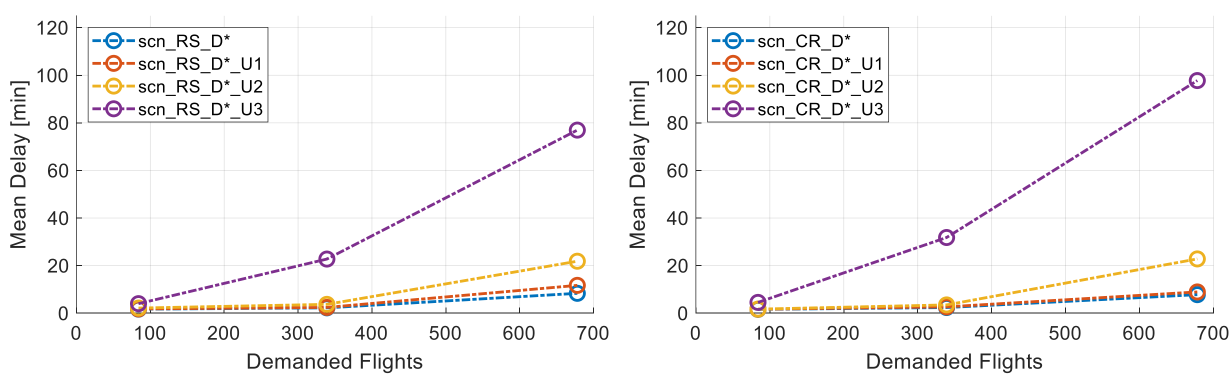

Figure 20. Mean flight delays as a function of demand level and trajectory uncertainty for the RS strategy (left) and the CR strategy (right).

4.2.2 LOS

The LOS information for each scenario run, along with the percentage of flights with LOS, can be seen in Figure 21 through Figure 29. Figure 21, Figure 22, and Figure 23 show the data for demand level 1 and trajectory error levels 1, 2, and 3, respectively. Similarly, Figure 24, Figure 25, and Figure 26 show the data for demand level 2 while Figure 27, Figure 28, and Figure 29 show the data for demand level 3, and for the three trajectory error levels. In each of these figures, the number of pair-wise LOS and percentage of flights with LOS for the unimpeded scenario without any error or SCM is shown as the 'baseline run.' Along the x-axis of the plot are scenario variations with or without uncertainty in planning and for the various SCM strategies. The left-most scenario in each plot indicates the variability in that metric when the unimpeded demand was executed with trajectory execution error but without any SCM strategy. The next eight scenarios in each figure show the variability on the parameter when each of the four SCM strategies was used both without and with uncertainty in planning. The data for each scenario shows the distribution of the metric in the 1000 Monte Carlo samples for trajectory execution error and provides one mechanism for assessing whether results are significantly different from each other.

At demand level 1 (Figure 21, Figure 22, and Figure 23), the baseline results without SCM or any trajectory uncertainty or error indicated that the number of pair-wise LOS was 17. Theoretically, this number of pair-wise LOS could lead to approximately 40% of flights with LOS, but some flights had multiple LOS, and the observed percentage of flights with LOS was 32.1%. When trajectory error was introduced into the baseline scenario without SCM, the number of LOS increased slightly while the number of flights with LOS remained nearly unchanged (as judged by the median values from the Monte Carlo runs). We saw that the variability in these metrics increased as the level of trajectory error increased, as one might expect. At trajectory error level 1, we observed that the CR strategy had a significantly better chance of reducing the number of LOS and the percentage of flights with LOS, with or without taking trajectory uncertainty into account in strategic planning. At trajectory error levels 2 and 3, the CR strategy without uncertainty did not appear to have a significantly better chance of reducing these metrics (based on inter-quartile range overlaps) but the CR strategy with uncertainty did still show better performance in reducing the number of LOS and the percentage of flights with LOS. For the AR, FR, and RS strategies, there did not appear to be a significant improvement in these metrics when using the strategy without taking into account uncertainty versus not using any strategy at all. At the trajectory error level 3, however, we did see a benefit in these metrics when using the AR and RS strategies and accounting for trajectory uncertainty. The RS strategy with uncertainty also showed some improvement at trajectory error level 2.

In the demand level 2 scenarios (Figure 24, Figure 25, and Figure 26), we saw that the baseline run had 155 LOS and 55.8% flights with LOS. When adding uncertainty to this baseline scenario with no SCM, the median number of LOS increased to nearly 200 while the median percentage flights with LOS increased to 57.5-64.3%. At trajectory error level 1, we saw that the AR, RS and CR strategies showed improvements in these metrics relative to the baseline case with error, with the CR strategy showing the biggest improvement compared to the others, but the FR strategy did not show any improvement. With trajectory error level 2, FR showed only marginal improvement when uncertainty is taken into account. The same was true for FR with error level 3. Also, at error level 3 we now saw less significant improvements in the metrics from the CR strategy without uncertainty whereas CR accounting for uncertainty still showed significantly reduced number of LOS (~191 -> 0) and percentage of flights with LOS (~64% -> 0%).

In the demand level 3 scenarios (Figure 27, Figure 28, and Figure 29), we saw that the baseline scenario with no SCM had 674 LOS and 79.8% flights with LOS. When trajectory error was added to the baseline scenario, the number of LOS increased to 826, 803, and 805 median LOS for the three error levels. Similarly, the median percentage of flights with LOS increased to 80.5, 82.3, and 84.6% for the three error levels. At trajectory error level 1, all strategies showed improvements in these metrics, with FR showing the worst improvement and CR showing the best improvement. At trajectory error level 2, we saw similar trends in improvement. At error level 3, the same trend was observed but only when comparing equivalent runs with or without accounting for uncertainty. At error level 3 and demand level 3, the benefits of accounting for uncertainty can be clearly seen in the improvements in number of LOS and percent flights with LOS. These metrics could be improved by a reduction in both LOS and percent flights with LOS to nearly half the magnitude of those metrics, even when the SCM strategies don’t account for uncertainty.

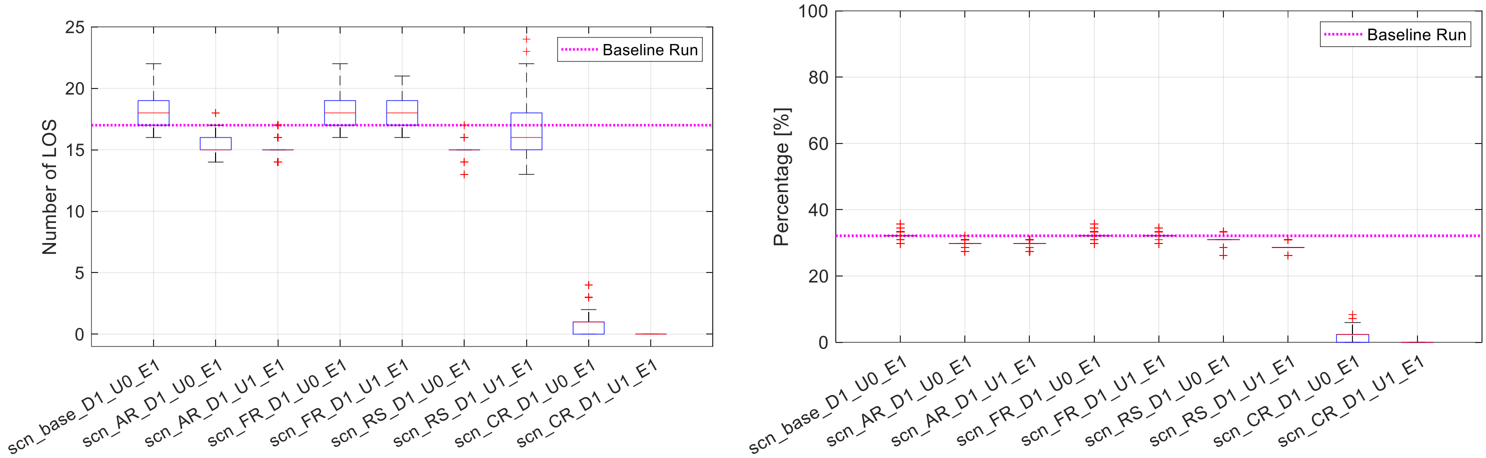

Figure 21. LOS counts and percentage flights with LOS for demand level 1 and trajectory error level 1.

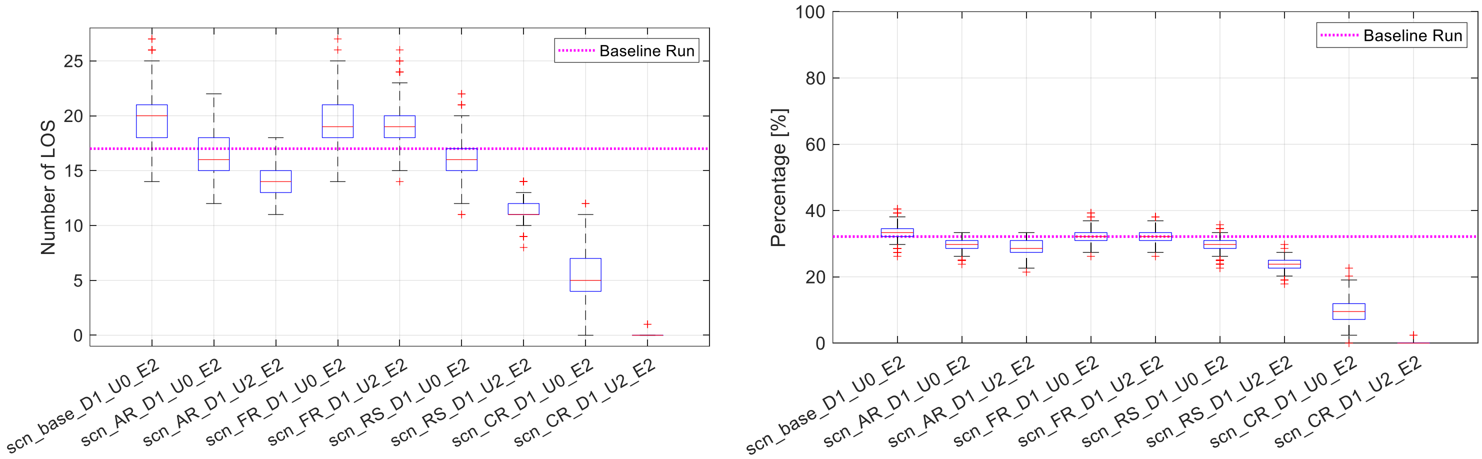

Figure 22. LOS counts and percentage flights with LOS for demand level 1 and trajectory error level 2.

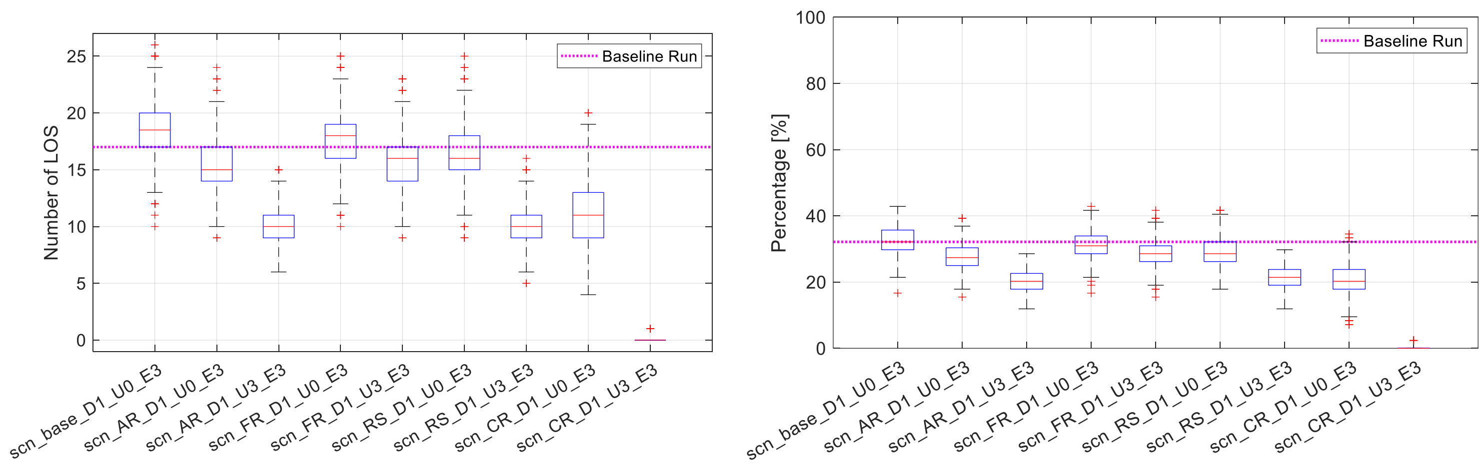

Figure 23. LOS counts and percentage flights with LOS for demand level 1 and trajectory error level 3.

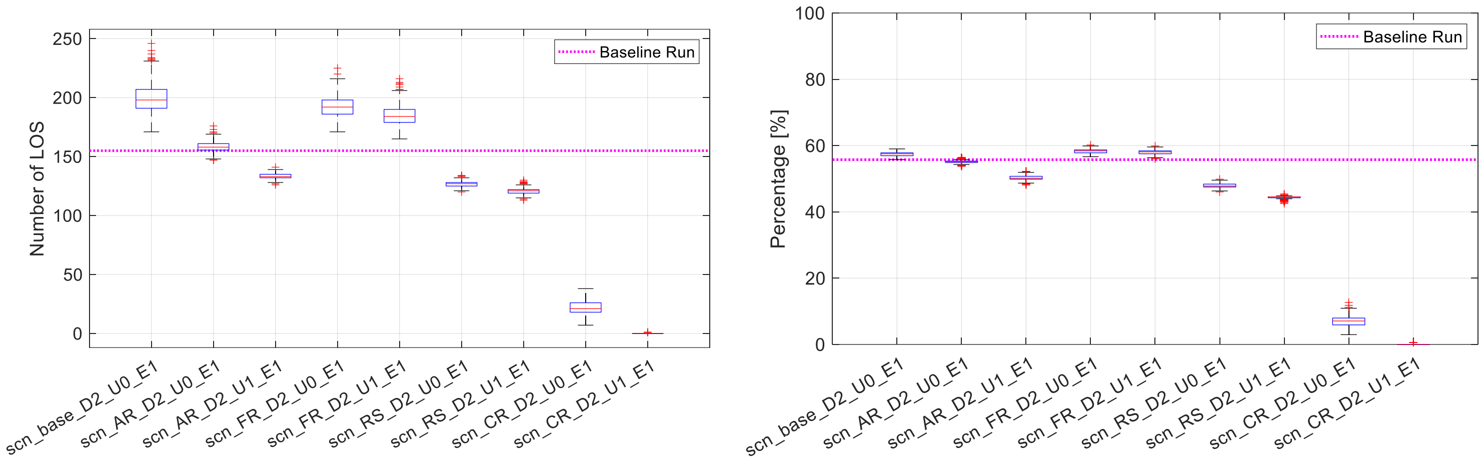

Figure 24. LOS counts and percentage flights with LOS for demand level 2 and trajectory error level 1.

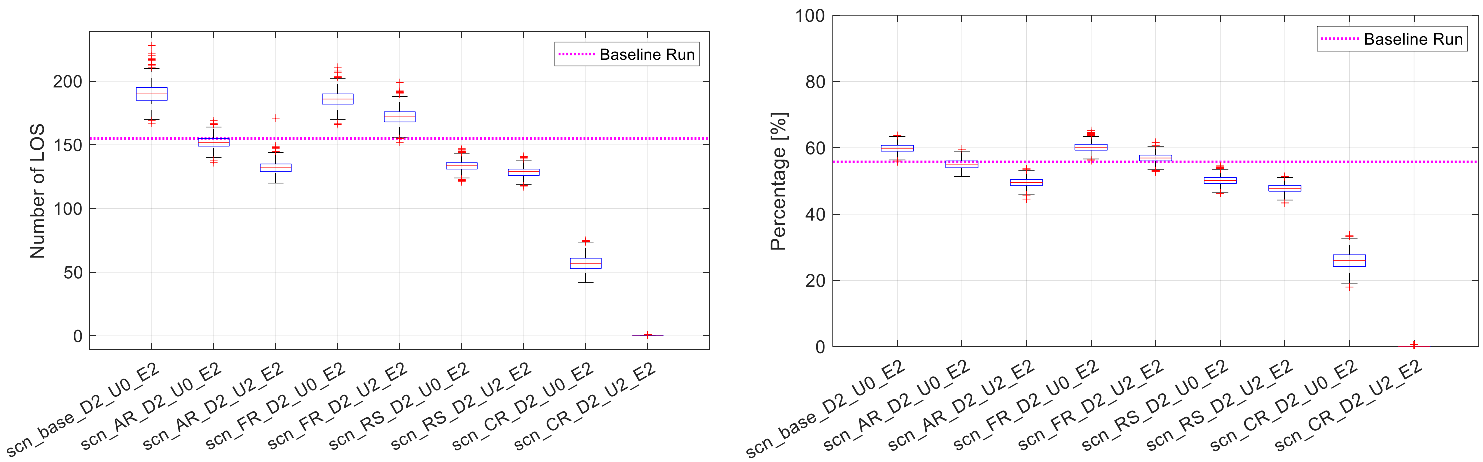

Figure 25. LOS counts and percentage flights with LOS for demand level 2 and trajectory error level 2.

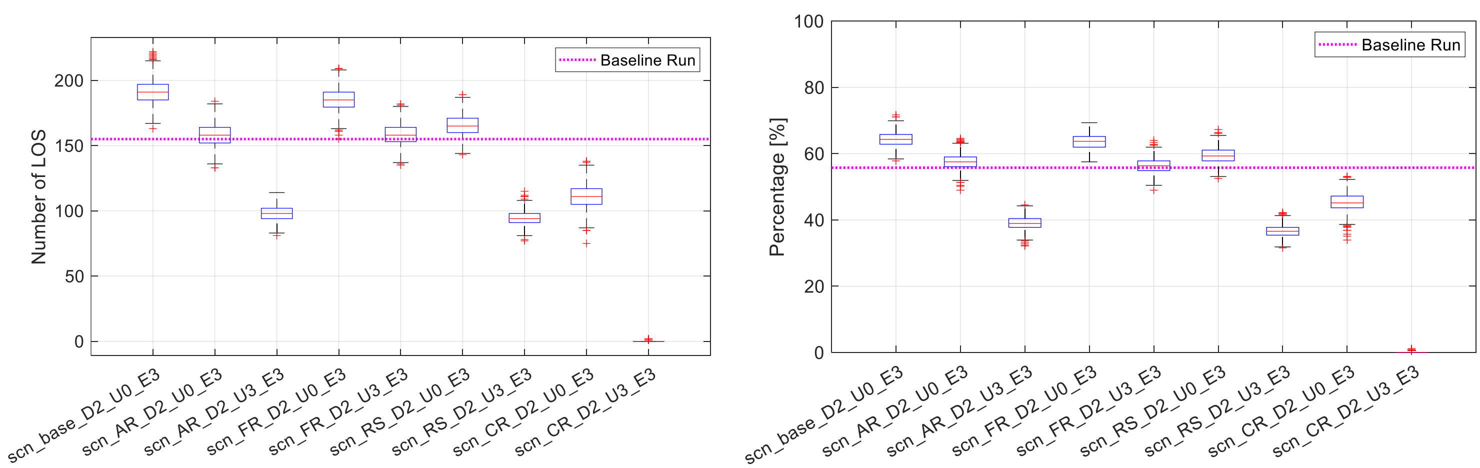

Figure 26. LOS counts and percentage flights with LOS for demand level 2 and trajectory error level 3.

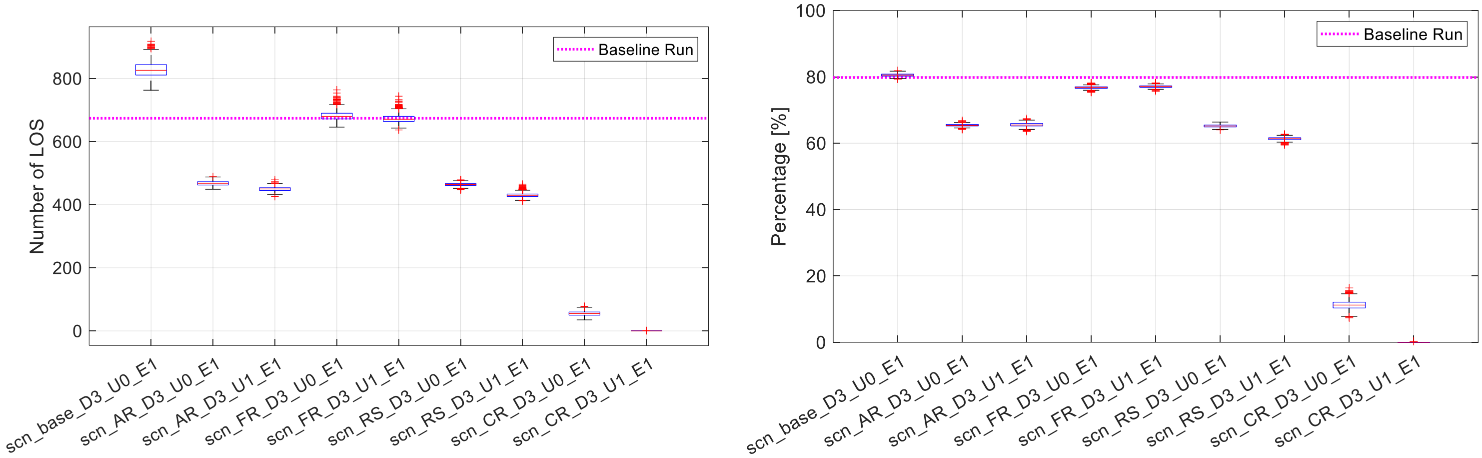

Figure 27. LOS counts and percentage flights with LOS for demand level 3 and trajectory error level 1.

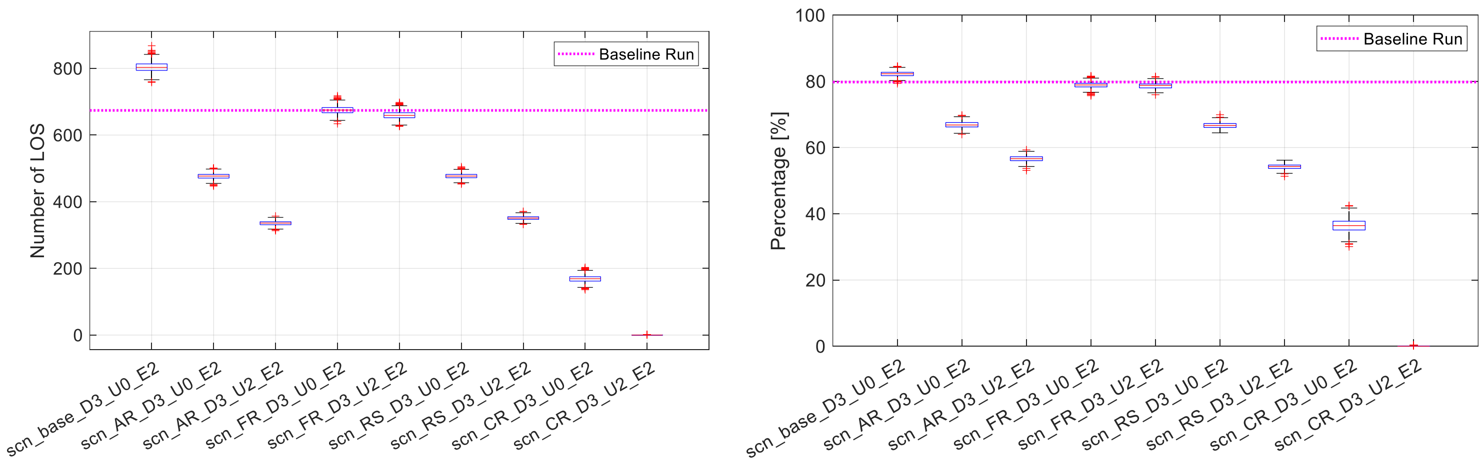

Figure 28. LOS counts and percentage flights with LOS for demand level 3 and trajectory error level 2.

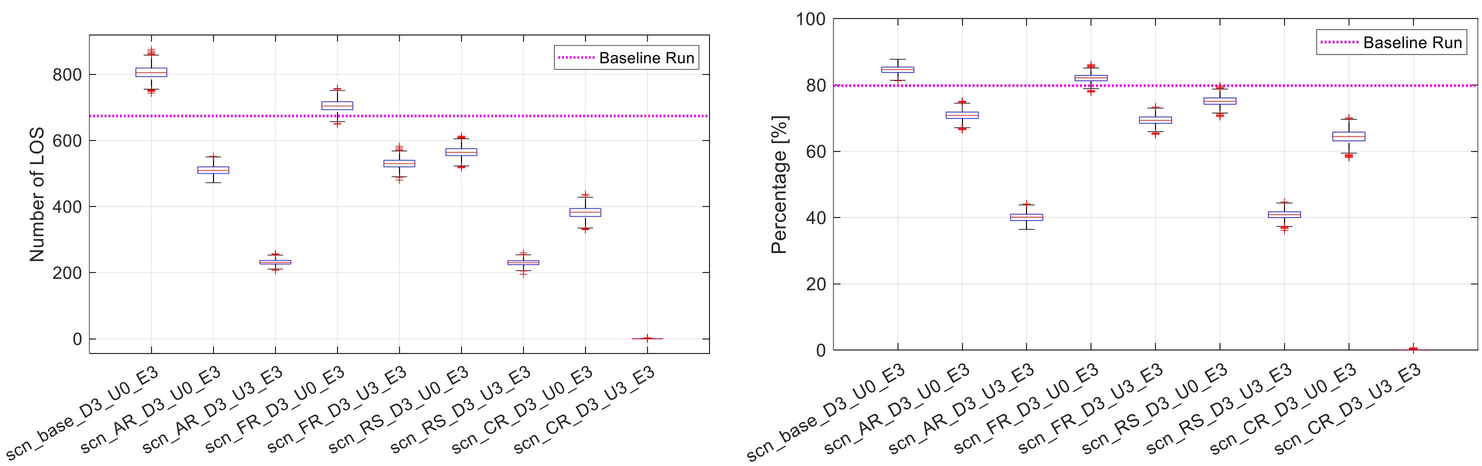

Figure 29. LOS counts and percentage flights with LOS for demand level 3 and trajectory error level 3.

4.2.3 Delays

Figure 30, Figure 31, and Figure 32 show the percentage of flights with delay and the mean delay for those delayed flights for each of the strategies as a function of the total flights in each demand level and by trajectory error level. Here, all strategies accounted for trajectory uncertainty. We note that the demand levels here can be thought of as demand rates as these were the number of total flights in a 2-hour period. Equivalently, because the airspace area covered by all of the demand sets was the same, one could also equate the demanded flights variable in these figures to a density measure. The FR strategy showed a consistently lower percentage of flights requiring delay and lower mean delay per delayed flight across all demand levels and trajectory error levels. All other strategies showed similar trends and magnitudes for the percentage of flights requiring delays. At the higher demand levels, the AR strategy showed the highest levels of mean delay for all three error levels, with RS and CR strategies showing only slightly lower delays. The percentage of flights requiring delay increased rapidly with increases in demand level and trajectory error level as did the mean delay. In fact, these scenarios were executed without delay limits but, realistically, some of this demand would not be executed if the delay levels were of this magnitude.

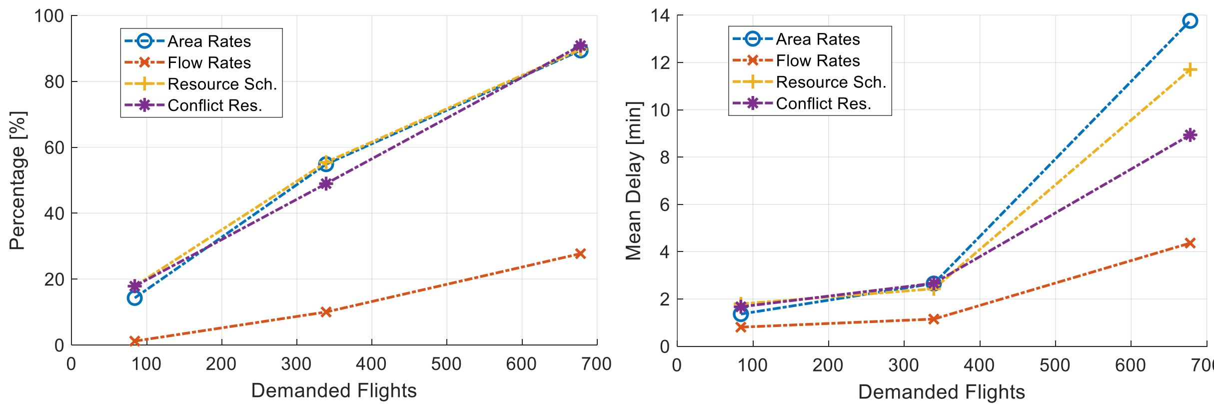

Figure 30. Percentage flights with delay (left) and mean delay (right) as a function of demand level, for trajectory uncertainty and error level 1.

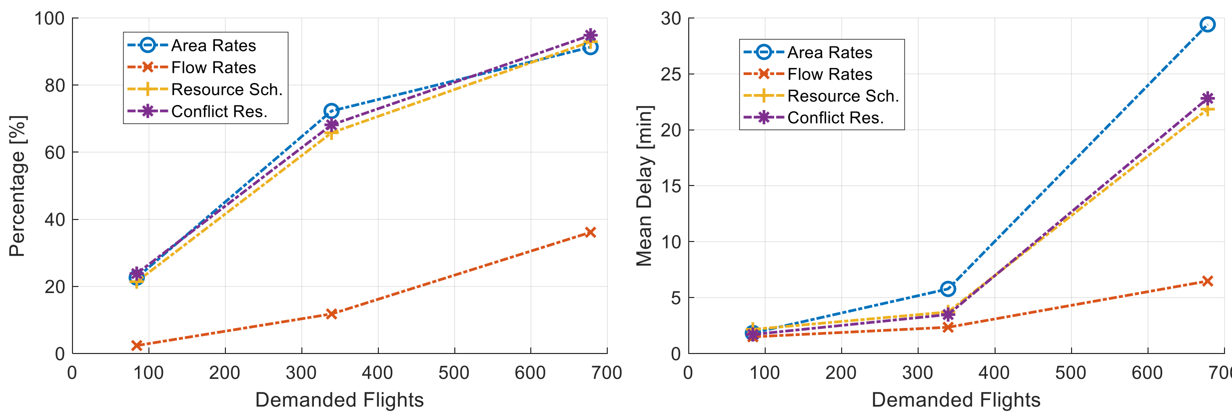

Figure 31. Percentage flights with delay (left) and mean delay (right) as a function of demand level, for trajectory uncertainty and error level 2.

4.2.4 Trade Space

The results indicated that a trade-off may be required in order to select the best SCM approach. If the selection criterion were based solely on LOS or percentage of flights with LOS, the preferred choice would be the CR strategy. If we were not concerned about LOS, the ideal choice when comparing flight delays would be the FR strategy. Figure 33, Figure 34, and Figure 35 show a possible trade space using the SCM strategies accounting for uncertainty and at the three different trajectory error levels. Suppose we could set thresholds for maximum allowable average delay (based on what makes sense for the business case) and either number of LOS or percentage of flights with LOS (based on the sophistication of the tactical conflict management mechanisms), we could then use this trade space to determine the achievable demand and suitable SCM strategies.

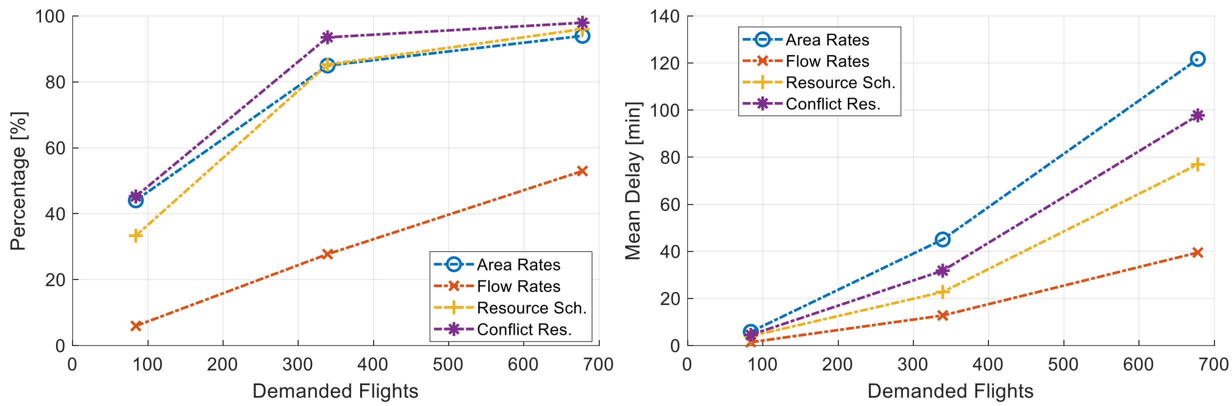

Figure 32. Percentage flights with delay (left) and mean delay (right) as a function of demand level, for trajectory uncertainty and error level 3.

As an example, suppose we wanted to ensure that any flight delayed by SCM was not delayed more than 5 minutes, and that we also wanted to ensure that no more than 40% of flights showed strategic conflicts that would likely lead to tactical conflict management actions after departure. We see that, under low trajectory uncertainty and error, any strategy would be suitable at demand level 1, only CR could support demand level 2, and no strategy could support demand level 3. The same is still true under uncertainty and error level 2. At uncertainty and error level 3, however, demand level 1 could be achievable with the CR, FR, and RS strategies, but not with the AR strategy as that would lead to a higher than acceptable mean delay. In addition, the CR strategy is also no longer feasible for demand level 2 with this high uncertainty and error.

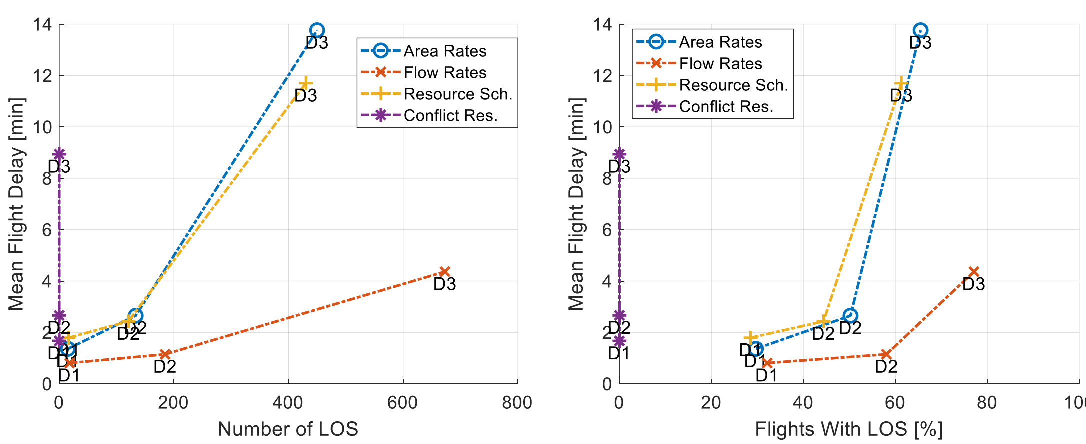

Figure 33. Mean delay versus mean number of LOS (left) and mean percentage flights with LOS for the various SCM strategies and the various demand levels, with trajectory uncertainty and error level 1.

We note that the source demand was somewhat contrived to achieve uniform vertiport demand levels and that a demand set based on a realistic business model and passenger demand would produce results somewhat different from what was observed in this study. Nonetheless, there does appear to be a saturation of the airspace design used for this analysis. We observed this with the high percentage of flights with LOS (~80%) at the highest demand level before implementing SCM. This would certainly be an extremely challenging problem for the tactical conflict management system to have to solve. Thus, under these conditions it would be desirable to implement a SCM approach that reduces these occurrences without introducing significant delays or limiting throughput.

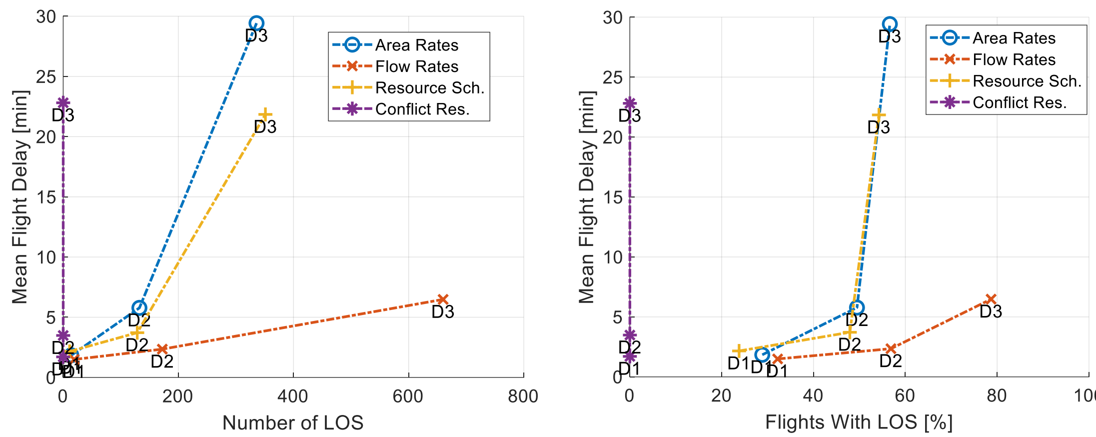

Figure 34. Mean delay versus mean number of LOS (left) and mean percentage flights with LOS for the various SCM strategies and the various demand levels, with trajectory uncertainty and error level 2.

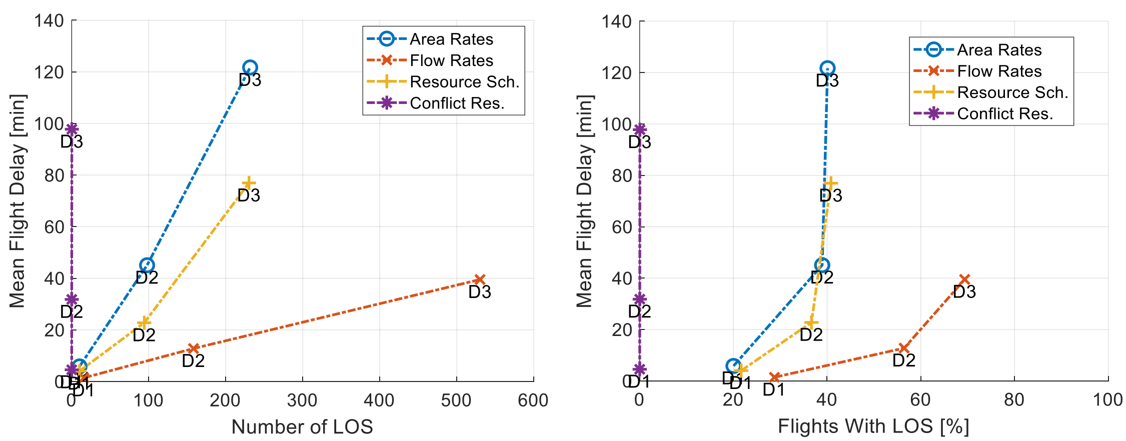

Figure 35. Mean delay versus mean number of LOS (left) and mean percentage flights with LOS for the various SCM strategies and the various demand levels, with trajectory uncertainty and error level 3.

5 Conclusions and Recommendations

In this analysis we compared various SCM strategies and their effectiveness in reducing the need for the tactical conflict management layer to act. We analyzed the strategy effectiveness using the following metrics: (1) unmitigated LOS, (2) flight delays imposed by strategic planning, and (3) throughput of the overall airspace. We introduced trajectory uncertainty in the SCM implementations that impacted the throughput of the airspace, but that improved the effectiveness of each SCM strategy when flight operations were executed with equivalent levels of error. Finally, we introduced a trade space using the results of the study that could serve as a mechanism for selecting the appropriate SCM strategy in a trade-off between uncertainty and error, mean flight delay and the LOS metrics (which this study equates to potential tactical conflict management actions).

The CR strategy was found to be the most effective at reducing LOS and percentage of flights with LOS, even in the presence of trajectory error, when uncertainty is accounted for in planning. Thus, if the objective is to reduce the number of actions that the tactical conflict management layer will have to take to resolve conflicts, CR is the strategy to use. Some of the other strategies did show the ability to at least reduce the number of LOS and percentage of flights with LOS, but perhaps not to levels that might be acceptable to a tactical conflict management layer. Thus, the decision to use a strategy other than CR comes down to the assessment of performance for the tactical conflict management layer in resolving traffic conflicts, which should be the focus of future studies. In that assessment, consideration needs to be made not only of the strategy's effectiveness, but the overall impact to strategic operational plans (e.g., the impact at downstream coordination points, such as vertiports).

The FR strategy was shown to be the worst at reducing LOS and percentage of flights with LOS. Thus, it is not very effective at reducing the actions required by a tactical conflict management layer. Nonetheless, it may still have effectiveness in other ways not measured by this study. For example, the rate-based conditioning of traffic may increase the likelihood of tactical conflict management layer success the same way it ensures manageable traffic levels for air traffic managers in today's air traffic system. A study to validate this assumption, or identify other important benefits, could be a future research activity if FR is the chosen SCM strategy.

The effectiveness of each strategy was measured in terms of other metrics, such as flight delays and throughput. In the airspace tested, the demand was high enough to produce unacceptable levels of flight delay and reductions of throughput when the SCM strategies were implemented. This was especially evident at the higher levels of trajectory uncertainty and error. Nominally, flights with this unacceptable level of delay would be canceled even though this study retained those flights in order to support a comparison of equivalent satisfied demand by the various SCM strategies. Nonetheless, a more accurate evaluation of throughput would need to be assessed in a simulation with a threshold set for the maximum allowable delay, which would lead to some flights being canceled. The simulation tools used in this analysis are capable of these types of analyses in future studies.

The need to account for expected levels of uncertainty becomes more and more important with increasing demand level, and as the level of trajectory error increases. That is because the likelihood of flights with LOS is increased with higher densities of operations, and as the level of trajectory error rises. This is especially true in the route structure used here where more trajectory error means a higher likelihood of causing a conflict with an operation on an adjacent track.

Accounting of uncertainty improves the strategy effectiveness with respect to LOS but increases delays and reduces throughput. Thus, there is a tradeoff between scalability and allowable levels of uncertainty. That is, one can implement an air traffic management construct that allows high levels of uncertainty, but we can expect that same system to have limits on scalability that may be evident even at small demand levels, such as those emulated in this study. Therefore, the results in this study indicate that an air traffic system should attempt to implement mechanisms appropriate for reducing uncertainty where possible in order to increase the chances for scalability. And this increased level of predictability of operations needs to be balanced with mechanisms for ensuring flexibility when operational conditions and plans need to change even though those types of changes should be the exception rather than the rule under normal operations.

The optimal solution for a given airspace, demand level, and other factors, could likely be a combination of SCM strategies. This study only compared the use of a single strategy at a time, but there could be benefits to using one strategy for one reason and another strategy for a different reason. For example, we observed that the RS and CR strategies had similar performance in terms of the other metrics, such as delay and throughput. Although not highlighted in the figures presented here, the RS strategy was very effective in reducing the LOS in the region of the track network even though it was unable to reduce LOS outside the track network. This is primarily due to the large number of scheduling constraints used and because the scheduling intervals were based on the lateral separation criterion. Under those conditions, most of the LOS with the RS strategy were outside the track network. Thus, it may be feasible to use this strategy within a given track network and to use the CR strategy where there is no track network.

Nevertheless, based on the CR strategy's ability to reduce the unmitigated LOS to nearly zero, even where there is no new structure in the airspace, it is obvious that only that strategic management approach has any real potential to scale up beyond very low levels of demand. This is especially true if the potential for large number of unmitigated LOS is consistent with what was observed in this study. By generating conflict-free trajectories before flight initiation, the tactical conflict layer must only deal with non-conforming aircraft, which, under a conformance-based air traffic management system, should occur infrequently. Other yet unpublished and related work and preliminary investigations have shown that aircraft (in simulation) can follow the generated 4-D trajectories with great accuracy, even in the presence of some level of disturbances. In addition, because the flight times are short for these types of operations, it is anticipated that weather predictions will also be very accurate and thus the aircraft will not have to fly at sub-optimal airspeeds very often. We believe that even high demand levels can be achieved with a combination of the CR approach and a tactical conflict management solution system, such as DAIDALUS [11] or DANTi [12]. We recognize that is a drastic change from how the current system operates today and that much research is still required to demonstrate how this mode of operating would be impacted by existing operational constraints, non-participating traffic, off nominal conditions, and many other factors.

References

[1] P. Kopardekar, J. Rios, T. Prevot, M. Johnson, J. Jung and J. Robinson, "Unmanned Aircraft System Traffic Management (UTM) Concept Of Operations," in *AIAA Aviation Forum*, Washington, DC, 2016.

[2] Deloitte Consulting LLP, "UAM Vision Concept Of Operations (Conops) UAM Maturity Level (UML) 4," NASA, Doc. ID:20205011091, 2020.

[3] Federal Aviation Administration (FAA), "Urban Air Mobility Concept of Operations v1.0," US Department of Transportation, Federal Aviation Administration, Washington, DC, 2020.

[4] International Civil Aviation Organization (ICAO), "Global Air Traffic Management Operational Concept," International Civil Aviation Organization (ICAO), 2005.

[5] Federal Aviation Administration (FAA), "ORDER JO 7110.65Z," US Department of Transportation, Federal Aviation Administration, Washington, DC, 2021.

[6] I. Levitt, N. Phojanamongkolkij, A. Horn and K. Witzberger, "UAM Airspace Research Roadmap Rev 1.2," NASA, Doc. ID:20220008917, 2022.

[7] RTCA, "DO-365 - Minimum Operational Performance Standards (MOPS) for Detect and Avoid (DAA) Systems," RTCA, 2017.

[8] N. M. Guerreiro, R. W. Butler, J. M. Maddalon and G. E. Hagen, "Mission Planner Algorithm for Urban Air Mobility - Initial Performance Characterization," in *AIAA Aviation Forum*, Dallas, TX, 2019.

[9] S. Verma, V. Dulchinos, R. Wood, A. Farrahi, R. Mogford, M. Shyr and R. Ghatas, "Design And Analysis Of Corridors For UAM Operations," in *DASC 41st Digital Avionics Systems Conference*, Portsmouth, VA, 2022.

[10] M. Rimjha, M. Li, N. Hinze, S. Tarafdar, S. Hotle, H. Swingle, A. Trani and J. Smith, "Demand Forecast Model Development And Scenarios Generation For Urban Air Mobility Concepts," NASA, 2020.

[11] C. Munoz, A. Narkawicz, G. Hagen, J. Upchurch, A. Dutle, M. Consiglio and J. Chamberlain, "DAIDALUS: DETECT AND AVOID ALERTING LOGIC FOR UNMANNEDSYSTEMS," in *2015 AIAA/IEEE Digital Avionics Systems Conference*, Prague, 2015.

[12] V. Carreno, "Evaluation, Analysis and Results of the DANTi Flight Test Data, the DAIDALUS Detect and Avoid Algorithm, and the DANTi Concept for Detect and Avoid in the Cockpit," NASA, 2020.

Appendix A - Capacity Estimation

The following describes the approach used to set the capacity for each of the SCM approaches. The objective was to derive and set capacities that were relatively equivalent for all SCM strategies such that the comparison of strategies was fair.

A.1 Resource Scheduling Capacities

Given the lack of capacity information for vertiports for these new type of operations, and assuming that vertiports are virtual sources and sinks of traffic (i.e., no limits on parking exist), we can assume a value for the vertipad usage time. A reasonable, yet conservative, value to use was selected based on 120-second slot sizes for the vertipads. This slot size translates to a capacity of 30 operations per hour, assuming a vertiport with a single vertipad. This value is reasonable for the scenarios used in this study, where the maximum expected level of operations was roughly 20 operations per vertiport per hour with the highest demand scenario. We note that, even at this capacity value of 30 operations per hour, delays will occur in a scenario of 20 operations per vertiport per hour due to the fact that the departure times of the given demand were random and could conflict in time. In addition, delays on departure could result from conflicting arrivals at the destination vertiport.

For the resource schedulers at the merging and crossing points of the UAM track system, one can assume that the time required to traverse a given distance determines the slot size. For example, to traverse 1200 feet (the route separation distance without buffer) at a speed of 130 knots, the time required would be ~5.5 seconds. However, if the vehicle speed is 70 knots, for example, the time to traverse that same distance is 10.2 seconds. For the purposes of this analysis, we used the more conservative time, and we added a 5% buffer for a spacing time of 10.7 seconds assuming the operations would be in trail of each other.

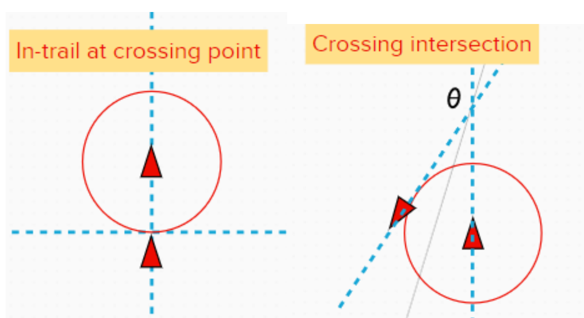

Figure A-1. In-trail at intersection (left) and crossing intersection (right).

The interval spacing between operations was dependent on the crossing angle. In the above example, we computed the spacing for an in-trail pair of vehicles. For waypoints that are crossing or merging waypoints, more time needed to be taken into account. That is, given the example in Figure A-1, more distance is required when the intersection is at a large angle. In fact, the distance between the crossing aircraft and the intersection to take this effect into account can be given by:

$$D_{eff} = \frac{D_{min}}{sin\left(\frac{180-\theta}{2}\right)} \qquad \textbf{\textit{(A-1)}}$$

Where, D_{min} is the distance required for in-trail sequencing, $180 - \theta$ is the route crossing angle (θ is the complementary angle; in-trail route crossing angle is zero; θ < 180), and D_{eff} is the distance that should be used for the crossing route in order to keep the aircraft at least D_{min} apart at their closest point. Note that both aircraft are assumed to be traveling at the same speed in this derivation. Thus, for a route crossing point with 90-degree angle between routes, the effective distance for a minimum distance of 1200 feet is 1697 feet. In order to provide that much distance in the same 130 knots example above, would require ~7.7 seconds; and at 70 knots that time would be ~14.4 seconds. Accounting for a small buffer, we arrive at a required spacing of 15.1 seconds.

Table A-1 lists the effective distances and spacing intervals for various crossing angles. For this analysis, the conservative spacing interval assumed 70 knots and included the 5% buffer, as shown in the last column of the table.

Table A-1. Effective distance and time for various route crossing angles.

Route Crossing Angle [deg]	Effective Distance [ft]	Time @ 130 knots	Time @ 70 knots	Time with 5% Buffer
0	1200	5.5	10.2	10.7
10	1205	5.5	10.2	10.7
20	1219	5.6	10.3	10.8
30	1242	5.7	10.5	11
40	1277	5.8	10.8	11.3
50	1324	6	11.2	11.8
60	1386	6.3	11.7	12.3
70	1465	6.7	12.4	13
80	1566	7.1	13.3	13.9
90	1697	7.7	14.4	15.1
95	1776	8.1	15	15.8
100	1867	8.5	15.8	16.6
105	1971	9	16.7	17.5
110	2092	9.5	17.7	18.6
115	2233	10.2	18.9	19.8
120	2400	10.9	20.3	21.3
125	2599	11.8	22	23.1
130	2839	12.9	24	25.2
135	3136	14.3	26.5	27.9
140	3509	16	29.7	31.2
145	3991	18.2	33.8	35.5
150	4636	21.1	39.2	41.2
155	5544	25.3	46.9	49.3
160	6911	31.5	58.5	61.4
165	9194	41.9	77.8	81.7
170	13768	62.8	116.5	122.4

A.2 Flow Rate Capacities

For the flow rate SCM approach, we defined rates based on the spacing intervals selected for the resource schedulers SCM approach. For vertiports, we selected a time window of 12 minutes to be consistent with other NASA research activities, resulting in a capacity of 6 operations per 12 minutes. This was consistent with the 30 operations per hour achieved with 120 second vertiport slot sizes in the RS strategy.

The capacity of UAM track network points was set using a smaller time window of 60 seconds and with the time intervals defined in the last column of Table A-1. The number of operations per time window was computed as follows:

$$Capacity = floor\left(\frac{TimeWindow}{TimeInterval}\right) \quad \textbf{\textit{(A-2)}}$$

For example, for a time window of 60 seconds and a time interval of 15.1 seconds at a 90 degree route crossing, the capacity was set to a rounded value from 3.97, or 3 operations per minute (rounding down was used in order to be conservative). The capacities as a function of route crossing angle can be seen in Table A-2.

Table A-2. Flow rate capacities as a function of route crossing angle.

Route Crossing Angle [deg]	Flow Rate/Capacity [Ops/min]
0	5
10	5
20	5
30	5
40	5
50	5
60	4
70	4
80	4
90	3
95	3
100	3
105	3

Route Crossing Angle [deg]	Flow Rate/Capacity [Ops/min]
110	3
115	3
120	2
125	2
130	2
135	2
140	1
145	1
150	1
155	1
160	1
165	1
170	1

A.3 Area-Based Flow Rate Capacities

Area-based capacity limits were implemented either in a binned approach or in a simultaneous operations limit approach. In the binned approach we limited the number of operations that could overlap with a given time bin window (e.g., capacity of 3 operations in a 5-minute window through a given area). In the simultaneous operations limit approach, we limited the maximum number of operations within the specified area at any given time (e.g., capacity of 3 simultaneous operations inside an area at any given time).

Vertiports were given a hexagonal region around them (vertices at 1386 feet, slides at 1200 feet distance from the vertiport). These vertiport areas were assumed to have a capacity of no more than one operation per 120 seconds, consistent with the RS SCM approach. The vertiport volumes can be seen in Figure A-2.

The UAM tracks were segmented into 2 nautical mile long segments and given a 600 feet buffer on either side to define an area/polygon 2 nautical miles long and 1200 feet wide where possible. The choice of 2 nautical miles was arbitrary and the 1200 feet width was based on the track spacing and the lateral separation criterion used in this analysis.

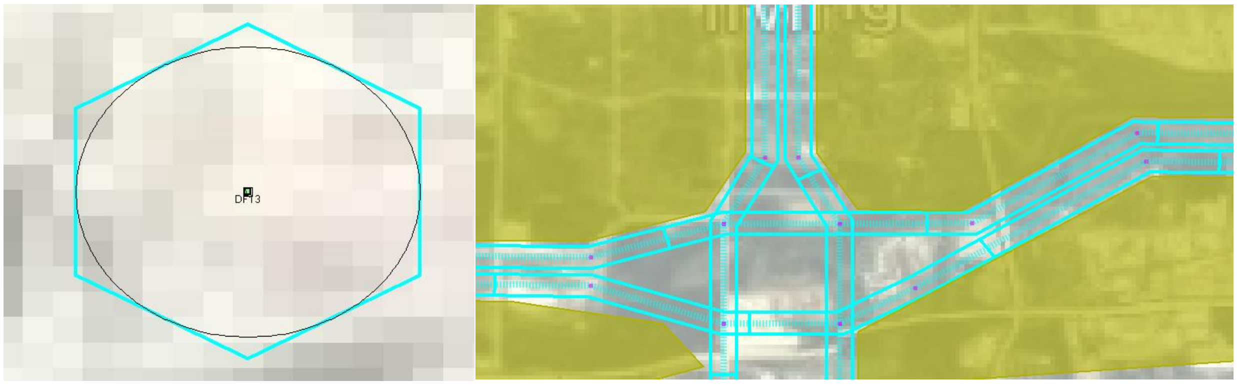

Figure A-2. Hexagonal volume around vertiports (left) and 2 nautical mile long polygon segments along the UAM track system (right) [background map credit: Google].

The capacity for UAM track polygons was determined using time intervals from the RS SCM approach. Assuming a cruise speed of 130 knots (the nominal for the flights in the scenarios for this analysis), we computed the distance traveled over the time intervals. The 2 nautical miles was then divided by this distance to determine the capacity for the segment. In addition, the time interval was selected based on the wort case crossing angle inside the polygon region. For example, in the case where there are no crossing intersections in the polygon region, we selected the time interval of 10.7 seconds and an aircraft traveling at 130 knots would traverse ~2348 feet. The capacity, then, is estimated to be 12,152/2,348, which equates to 5 when rounded down. Thus, the capacity would be 5 simultaneous operations within the polygon area at any given time. Table A-3 shows the simultaneous operations capacity assigned to the 2 nautical mile segments as a function of the worst-case route crossing angle inside the polygon area.

Table A-3. Simultaneous operations capacity for a 2 nautical mile long polygon segment.

Route Crossing Angle [deg]	Equivalent Distance @ 130 knots [ft]	Simultaneous Operations Capacity
0	2347.7	5
10	2347.7	5
20	2369.6	5
30	2413.5	5
40	2479.3	4
50	2589.1	4
60	2698.8	4
70	2852.3	4
80	3049.8	3
90	3313.1	3
95	3466.7	3
100	3642.2	3
105	3839.7	3

Route Crossing Angle [deg]	Equivalent Distance @ 130 kts [ft]	Simultaneous Operations Capacity
110	4081.0	2
115	4344.3	2
120	4673.5	2
125	5068.4	2
130	5529.2	2
135	6121.6	1
140	6845.6	1
145	7789.1	1
150	9039.7	1
155	10817.0	1
160	13471.8	1
165	17925.9	1
170	26855.9	1

A.4 Conflict Detection and Resolution Capacities

The CR SCM approach did not require any explicit capacities to be set. The throughput of the airspace was limited by the choice of effective separation criteria, which were set to 1200 feet lateral and 450 feet vertical. The exception to this was at vertiports, where resource schedulers were used and the scheduling interval was set at 120 second between operations.

Appendix B- Trajectory Uncertainty and Error

In this study, we consider the impacts of trajectory following error on the performance of the various strategic conflict management approaches. Trajectory following error represents the difference between the actual flown trajectory and the established operational plan trajectory. These trajectory following errors can be due to many factors, including pilot or automation execution error, or atmospheric disturbances. The purpose of this study was not to model all of those factors; rather, we abstracted those component errors into a cumulative error. Thus, the trajectory following errors were modeled as the temporal error, lateral error, and vertical error with respect to the reference trajectory.

The strategic conflict management strategies accounted for this trajectory following error during flight planning; we refer to this as trajectory uncertainty. The trajectory uncertainty in planning was modeled as expected levels of trajectory following error. That is, a single value for each of the temporal, lateral, and vertical error was based on the sigma values of the assumed normal distributions of the expected error levels.

The trajectory following error was modeled using Eq. (1), as described in section 3.6 of the report.

Reasonable values for temporal period were derived from the temporal error magnitude and the speed of the vehicle. For a given nominal cruise speed, GS_{nom}, the distance travelled, ds, during one complete cycle of an error magnitude, e_{max}, can be given by:

$$ds = GS_{nom} * 4 * e_{max} \qquad \textbf{\textit{(B-1)}}$$

Assuming a small average delta speed is available for temporal error correction, ΔGS, the speed control logic must traverse that same distance in a single period, T, is:

$$ds = \Delta GS * T \qquad \textbf{\textit{(B-2)}}$$

Substituting, we arrive at the equation for the period as:

$$T = \frac{GS_{nom} * 4 * e_{max}}{\Delta GS} \qquad \textbf{\textit{(B-3)}}$$

An example, assuming $GS_{nom} = 130\ kts$, $\Delta GS = 10\ kts$, and $e_{max} = 2\ sec$, the period is 104 seconds. For $e_{max} = 10\ sec$, the period is 520 seconds, and for $e_{max} = 40\ sec$, the period is 2080 seconds.

Note that e_{max} is the magnitude of error sampled from the normal distribution, thus, the period becomes a function of the error magnitude. This makes sense because, correcting for a 40 second error would take much longer that correcting for a 2 second error, given the same amount of control authority.

A similar argument can be made for the vertical period using the nominal vertical speed, VS_{nom}, and the delta available vertical speed, ΔVS. Thus, the vertical error period takes a similar form to that of the temporal period:

$$T = \frac{VS_{nom} * 4 * e_{max}}{\Delta VS} \qquad \textbf{\textit{(B-4)}}$$

Although a similar derivation could be made for the period of the lateral errors, the lateral drift rate is something that is more difficult to estimate. Thus, we use a simple ratio, similar to $GS_{nom}/\Delta GS$ and

$VS_{nom}/\Delta VS$, simply as "$LateralRateRatio$" and select a reasonable value for that parameter that is greater than 1 (indicating less lateral movement for correction relative to maximum or nominal lateral movement).

In the interest of clarity, the list below provides a set of assumptions for the error model. Although the validity of assumptions may be questioned (in particular, regarding error independence), it is unlikely that such differences will have notable impact on this paper's conclusions. In addition, as more and more performance data from UAM operations with new types of vehicles is made available, more sophisticated error models may be possible that could increase the level of fidelity of these types of simulation analyses and provide opportunities for re-analyzing the SCM problem.

- Each component of error is independent of the others
- Each component of error is modeled from a normal distribution
- Each trajectory's error is independent from the errors of other trajectories
- Each component of error is parameterized as a magnitude (sampled from a normal distribution), a sinusoidal period, and a sinusoidal phase
- A "bread-crumb" trajectory is created at 1Hz in order to implement the error model; this trajectory starts as a sampling of the expected trajectory/plan
- The vertical error is added first, the lateral error is added second, and the temporal error is added last: a magnitude, period, and phase are sampled from a distribution for each of these components
- The error normal distribution parameters are given in Table 3 and Table 4
- For strategic planning, +/-2σ is used as the uncertainty value
- For period computations of the above equations we assume: $GS_{nom} = 130\ kts$, $\Delta GS = 10$ knots, $VS_{nom} = 500\ ft/min$, $\Delta VS = 250$ feet/min, $LateralRateRatio = 2$