Flight Evaluation of a Flight Path Management System for High Density Advanced Air Mobility

Mark G. Ballin,[1] Bryan A. Barrows,[2] Stewart L. Nelson,[3]
Matthew C. Underwood,[4] Tyler D. Fettrow,[5] and David J. Wing[6]
NASA Langley Research Center, Hampton, Virginia, 23681

David A. Karr[7]
Analytical Mechanics Associates, Hampton, VA 23666

Flight Path Management (FPM) is a proposed automation capability for future high-density civil aviation operations. FPM provides strategic and dynamic flight path planning in the presence of other users sharing the airspace, supporting traffic flow management and operator objectives. A flight test of FPM was conducted using a reference prototype automation system developed by NASA. Test goals included a verification of the automation's core functions, exploration of behaviors in a flight environment, and validation of a simulation used to develop the concept and technology. Testing was performed using helicopters as surrogates for future electric vertical take-off and landing aircraft. The two aircraft engaged in conflict encounters in a live-virtual-constructive operating environment populated with up to 330 virtual traffic aircraft. Test results provide positive indications that a vehicle-centric implementation of FPM can be made operational in a future high-density air taxi environment. The FPM-equipped aircraft met its flight objectives while maintaining separation with the real and virtual aircraft. All functional performance success criteria were met or exceeded in an airspace environment containing area restrictions and arrival metering constraints. Test maneuvers involving operations within flow corridors were less successful, indicating further research and development is necessary. Test results and pilot feedback suggest human operators can be in the FPM decision loop at high traffic density.

I. Introduction

Advanced Air Mobility (AAM) is a vision for a future air transportation system that serves new aviation markets, relying on a new generation of technologically advanced aircraft (Ref. 1). One potential new market, referred to as Urban Air Mobility (UAM), involves air transportation within metropolitan areas. The long-term UAM vision anticipates thousands of simultaneous air operations over urban areas, providing economic benefits and relieving urban ground congestion.

UAM introduces significant challenges in the management of flight operations. The Federal Aviation Administration and NASA have proposed concepts for evolutionary development toward high-density urban operations. Initial low-tempo UAM operations will take place in the current regulatory and operational environment, followed by evolutionary changes to increase operational tempo (Ref. 2). As a complement to these activities, this study takes a transformational approach. It investigates one technical challenge associated with management of future high-density operations: in-flight path planning that ensures traffic separation while complying with restrictions and planning constraints necessary in the management of traffic flow.

---

[1] Research Engineer, Crew Systems and Aviation Operations Branch, AIAA Associate Fellow
[2] Research Engineer, Autonomous Integrated Systems Research Branch, AIAA Member
[3] Research Engineer, Crew Systems and Aviation Operations Branch, AIAA Member
[4] Systems Engineer, Engineering Integration Branch, AIAA Senior Member
[5] Human-Machine Scientist, Crew Systems and Aviation Operations Branch
[6] Research Engineer, Crew Systems and Aviation Operations Branch, AIAA Associate Fellow
[7] Lead Principal Scientist, AIAA Senior Member

Safety and efficiency of increasingly complex AAM operations are expected to require extensive use of automation, ranging from controlling revolutionary new aircraft types to managing flights dynamically in high-tempo airspace and aerodrome operations. Airborne and ground automation will play central roles in assisting AAM operators with managing the flight paths of their fleets. Flight Path Management (FPM) is one such function potentially needed for the high-density operations anticipated for UAM at maturity. NASA's UAM maturity scale defines maturity level 4 (UML-4) as having air transportation "accessible and attractive to a significant percentage of the public for travel between high-density, origin-destination pairs (e.g., commercial airport to business district)" (Ref. 3). The UML-4 stage is envisioned to involve hundreds of aircraft aloft over the served urban area in all-weather conditions, a high level of traffic density compared to current operations. Air traffic management functions, which keep aircraft safely separated, maintain high traffic flow, and accommodate operator preferences when possible, cannot support UML-4 in their current state. They currently depend on human situation awareness that is unlikely to be able to keep up with the high tempo and increased traffic density. FPM automation may prove necessary to enable UML-4 operations. A detailed description of the FPM intended function is provided in Ref. 4.

Decision makers need foundational results from operations research to inform the ambitious UAM vision. Concepts similar to FPM have been investigated for airline operations for more than two decades, and their enabling technologies have been prototyped for research. Most of the research has been performed using air traffic simulation. The next logical phase of research and development is to perform evaluations in flight to verify expected performance and to discover unknown challenges. In this study, a reference prototype FPM automation system was developed and evaluated in flight for anticipated future operations at UML-4. The flight test was conducted by a joint undertaking of the Automated Flight and Contingency Management (AFCM) Subproject and the National Campaign Subproject of NASA's AAM Project, in partnership with Sikorsky, a Lockheed Martin Company.

Testing was performed using two helicopters as surrogates for future electric Vertical Take-Off and Landing (eVTOL) aircraft. One helicopter was equipped with the reference prototype FPM automation system, and is referred to as the ownship. Another helicopter was used as a cooperating intruder, sharing its flight state and intent with the FPM automation-equipped aircraft. The two aircraft engaged in conflict (i.e., predicted loss of separation) encounters in a live-virtual-constructive (LVC) UML-4 operating environment. A model of future UML-4 airspace was developed and populated with up to 330 virtual traffic aircraft that shared their flight states and intent.

This paper describes the goals, approach, and conduct of the flight test, and it presents results and insights gained. An overview of FPM is provided, followed by a description of the developed reference FPM automation. The flight test goals and objectives are then described, and the test design, implementation, and execution are covered at a high level. The paper then presents select results and analysis, and it concludes with a summary of insights gained.

II. FPM Concept Overview

UAM is among the most complex operating environments envisioned for commercial aviation. UAM viability will depend on achieving many technical breakthroughs, including possibly introducing a new operating mode that automates traffic separation under exceedingly challenging and diverse constraints. Examples of UAM constraints that must be accounted for include:

- Limited endurance of electric vertical take-off and landing (eVTOL) aircraft, requiring ultra-efficient flight paths
- New small separation standards to enable sufficient traffic density and volume
- Dynamic scheduling of high-demand vertiports, requiring reliable on-time arrivals and departures
- Geographic constraints defined by municipalities, air traffic control, and airport traffic flows
- Altitude constraints defined by buildings, terrain, air traffic control, and airport traffic flows
- Coordination procedures that consume decision-making time
- Wide-ranging aircraft performance of many UAM aircraft configurations
- Interactions with traffic operating under different flight rules.

An approach to dynamically managing flight paths is needed that can address all these constraints simultaneously. Many of these constraints may change during a flight. A sequential approach, such as use of a strategically deconflicted predeparture plan followed by an in-flight tactical "detect and avoid" capability, may be insufficient to provide the required reliability and robustness for sustainable commercial operations. FPM was first proposed in Ref. 5 as one of several functions that are performed in managing the airborne phase of flight, whether performed by human operators, automation, or a combination of humans and automation. FPM considers all constraints simultaneously by strategically replanning the remaining flight each time a constraint is changed or a new constraint is encountered.

Strategic planning achieves mission objectives and retains them when replanning is necessary. Strategic planning can be contrasted with tactical maneuvering, which suspends mission objectives to achieve a time-critical objective, such as resolving an immediate safety issue. FPM is envisioned to work in concert with a tactical maneuvering system, which provides an additional layer of safety.

Mature AAM concepts of operation must define the distribution of system functions and the allocation of responsibility and authority to agents within the system. Design alternatives for the concepts and their supporting system infrastructures are the subjects of ongoing research. Evaluation of these alternatives will depend on the study of system component functions and architectural alternatives involving those functions. Locating the FPM functional capability and its enabling technology with each participating aircraft is one such architectural alternative and is the architecture under study in the flight test. Each aircraft performs the FPM function while sharing information relevant to other system participants and receiving information relevant to itself from other participants.

For such a distributed-agent architecture applied to future UML-4 operations, it is assumed that:

- The vehicle operator is responsible for separation from traffic while conforming to area constraints such as corridors and restricted airspace regions.
- Schedule and airspace constraints are issued from external services. The vehicle operator is responsible for complying with them, even when they change in flight.
- The operator employs collaborative and responsible automation to perform the separation function. Collaborative and responsible automation is defined as automation "assured to perform specified functions such that human monitoring and mitigation of potential failures of those functions is no longer necessary" (Ref. 3).
- Data exchange and any necessary coordination is accomplished through airborne and ground-based data links.

III. FPM Automation Overview

A. Intended Function

An automation system performing dynamic path planning must construct and maintain a flight path with five principal qualities: feasibility, deconfliction, harmonization, flexibility, and optimality (Ref. 6).

1. A *feasible* path is one that conforms to the aircraft performance and range capabilities; complies with the airspace structure, rules, and constraints; avoids the terrain and charted obstacles; and meets the arrival constraints.
2. A *deconflicted* path is one that avoids unsafe proximity to known aircraft, dynamic obstacles, inclement weather, and other emergent airspace hazards.
3. A *harmonized* path is one that follows cooperative rules and procedures to ensure that the use of the airspace is coordinated with other airspace users.
4. A *flexible* path is one that provides adequate maneuvering room to ensure future flight path changes, if needed, are available and feasible.
5. An *optimal* path is one that best achieves the operator's business objectives for the specific flight.

The primary tasks of the FPM automation system are:

- to create the flight path
- to monitor the flight path and the factors that may impact it
- to evaluate ongoing acceptability of the flight path and proposed changes
- to revise the flight path, as needed, to sustain the desired qualities, and
- to coordinate the flight path with other airspace users and service providers.

The computed flight path accounts for mission objectives, pilot inputs, aircraft performance, ownship and traffic flight information, current and predicted atmospheric conditions, obstacles and terrain constraints, community rules, and variations in airspace and aerodrome configurations. The FPM automation system generates a four-dimensional flight path, associated operational intent, and an energy consumption profile, as well as contextual display information, flight conformance status, system health and alerts, and contingency response options.

B. Reference Implementation

A prototype FPM automation system developed by NASA was used for the flight test. This prototype, a modified version of the NASA Autonomous Operations Planner (AOP), has been developed and tested using batch and human-in-the-loop simulations. Development of AOP began in the late 1990s to support the exploration of distributed air/ground traffic management concepts in airline operations (Ref. 7). These concepts had the goal of increasing capacity of the National Airspace System, and were based on the following propositions:

- An inability to predict future airspace system states with high precision requires a management approach capable of timely response to unanticipated changes. Responses must not cause a degradation in stable system operation.
- A distributed system architecture relieves workload bottlenecks associated with centralized architectures that require human operators.
- Placing a management function in the location that contains the most timely and highest-quality information produces a highly responsive system and a reduced need to exchange information between management functions and execution functions.

Reference 8 provides further detail of the distributed air/ground concept of operations for which AOP was originally designed. Specific modifications made to the AOP automation to support AAM UML-4 operations are described in Ref. 9. AOP capabilities deemed ready for flight testing align with the five principal flight path qualities as follows:

1. Feasibility
    - AOP's trajectory generation function uses a model of the ownship's performance capabilities and operating limits.
    - Trajectory generation is aware of the ownship's state, including energy on board.
    - Generated trajectories comply with crossing restrictions, special use airspace restrictions, and instructions to arrive at a required time.
    - Some constraints can be relaxed to satisfy higher-priority constraints.
2. Deconfliction
    - AOP's deconflicted trajectories produced by its conflict detection (CD) and conflict resolution (CR) capabilities avoid traffic, current and forecast convective weather, and current and planned future restricted airspace.
    - Multiple conflict resolution alternatives are identified to meet replanning objectives, as available. Alternatives are offered for lateral path changes, vertical path changes, a combination of lateral and vertical changes, and speed changes.
3. Harmonization
    - AOP's conflict resolution solutions are coordinated with traffic using priority rules (right-of-way rules).
    - Resolution maneuvers employ conflict prevention: computed resolutions will not create new conflicts with traffic.
    - Resolution computations respect instructions to arrive at a required time and respond to changes in required arrival time instructions.
4. Flexibility
    - Not in scope for this effort. In previous research, AOP was used to demonstrate the concept of flexibility preservation, but the capability was not retained in the AOP baseline.
5. Optimality
    - AOP's resolutions provide paths of lowest energy, shortest distance, or shortest time, depending on operator preferences.

These functions are necessary, but probably not sufficient, for full UML-4 operations. Additional capabilities have been envisioned for each of the five automation qualities and will be the subjects of future research and development. As described below, one goal of the flight test was to inform the need for envisioned and unanticipated functional capabilities. References 6 and 9 describe the flight-tested capabilities in detail and provide information about additional capabilities under consideration.

Note that AOP is a research prototype automation system and is therefore only intended as a *reference* capability. The UML-4 definition describes automation as "collaborative and responsible," as defined above. As a research prototype, AOP was not designed to meet the level of assurance required to be considered collaborative and responsible. It was considered suitable for the flight test due to extensive safety procedures employed by the IAS test team.

Major AOP components include a conflict detector, a strategic resolver, and a trajectory generator. Traffic conflicts, future flight into restricted airspace or hazardous weather regions, and noncompliance with trajectory constraints are identified by the conflict detector. It propagates active trajectories of all aircraft forward in time, making use of trajectory intent, if known. Priority (right-of-way) rules are incorporated into logic that determines which aircraft is responsible for resolving the conflict. When the hosting aircraft is responsible for resolving a conflict or one of the other issues identified by the conflict detector, the process of strategic intent-based conflict resolution is triggered. The strategic resolver generates new 4D trajectory candidates to resolve issues identified by the conflict detector. It uses a genetic algorithm (referred to as the Pattern-Based Genetic Algorithm, or PBGA) to perform an iterative search for solutions using a set of parameterized maneuver patterns for the resolution candidates. Several patterns are available, classified into four groups according to the primary degrees of freedom controlled by each pattern: lateral resolutions, vertical resolutions, a combination of lateral and vertical resolutions, and speed-change resolutions. For each degree of freedom, PBGA iterates over 20 generations, each of which has a population of 20 trajectory candidates. A fitness function within the strategic resolver ranks candidates based on specified optimization parameters. The resolver iterates with the conflict detector to remove any trajectory candidates that either do not resolve the identified issue or create new issues. The specific 4D path of each trajectory candidate is created by AOP's trajectory generator, responding to calls from the resolver. The trajectory generator makes use of a model of the hosting aircraft that incorporates performance and procedural constraints to ensure each trajectory created can be flown by the hosting aircraft. Data are provided to these functions in near-real-time from sources internal and external to the hosting aircraft.

Flight test research pilots were provided a user interface to the FPM automation to allow them to evaluate the quality and utility of the technology's solutions. Research of a mature human/automation teaming concept was outside the scope of the flight test activity, so the interface was not designed to support such research. Referred to as the FPM Engineering User Interface (UI) (Fig. 1), the interface was provided to the research pilot on a touchscreen tablet computer that communicated with the AOP software engine residing on a separate computer. It was responsible for providing a graphical display of the current operational situation and a mechanism to preview, select, and execute AOP-generated trajectory change solutions. The interface provides a plan view and profile view of the environment in proximity to the ownship. The plan view displays real and virtual traffic participating in the test. It also shows airspace regions to avoid, which are represented by polygons. Additional information is displayed as it is relevant to the situation, such as a required time of arrival (RTA) constraint. The figure shows a crossing conflict with the intruder, with a predicted loss of separation (LOS) location shown by the yellow circle near the center of the plan view. The intruder aircraft is highlighted because it is in conflict with the ownship. Magenta lines represent the ownship's active trajectory, and cyan lines display one of the conflict resolution advisories provided by AOP.

IV. Flight Testing

The FPM flight test activity was designed to be an initiating element of research focused on at-scale operations envisioned by the UAM/AAM community. For this initial activity that used two helicopters as eVTOL surrogates, only one helicopter was equipped with AOP. The two aircraft and virtual background traffic engaged in conflict encounters in an LVC UML-4 operating environment that provided airspace constraints and traffic management constraints.

A. Goals and Objectives

Within the limitations of the flight test, the research explored the performance and system-level behaviors of the FPM function. The overarching goal was to determine the level of maturity for the concept and its enabling technology. Performance of various AOP functions were examined to understand operational behavior in a relevant environment. The flight test activity sought to:

- assess the FPM automation reference implementation to determine if selected capabilities perform as intended in a flight environment
- establish functional performance benchmarks that will be useful in developing future requirements
- set automation performance benchmarks based on the current state of emerging technologies

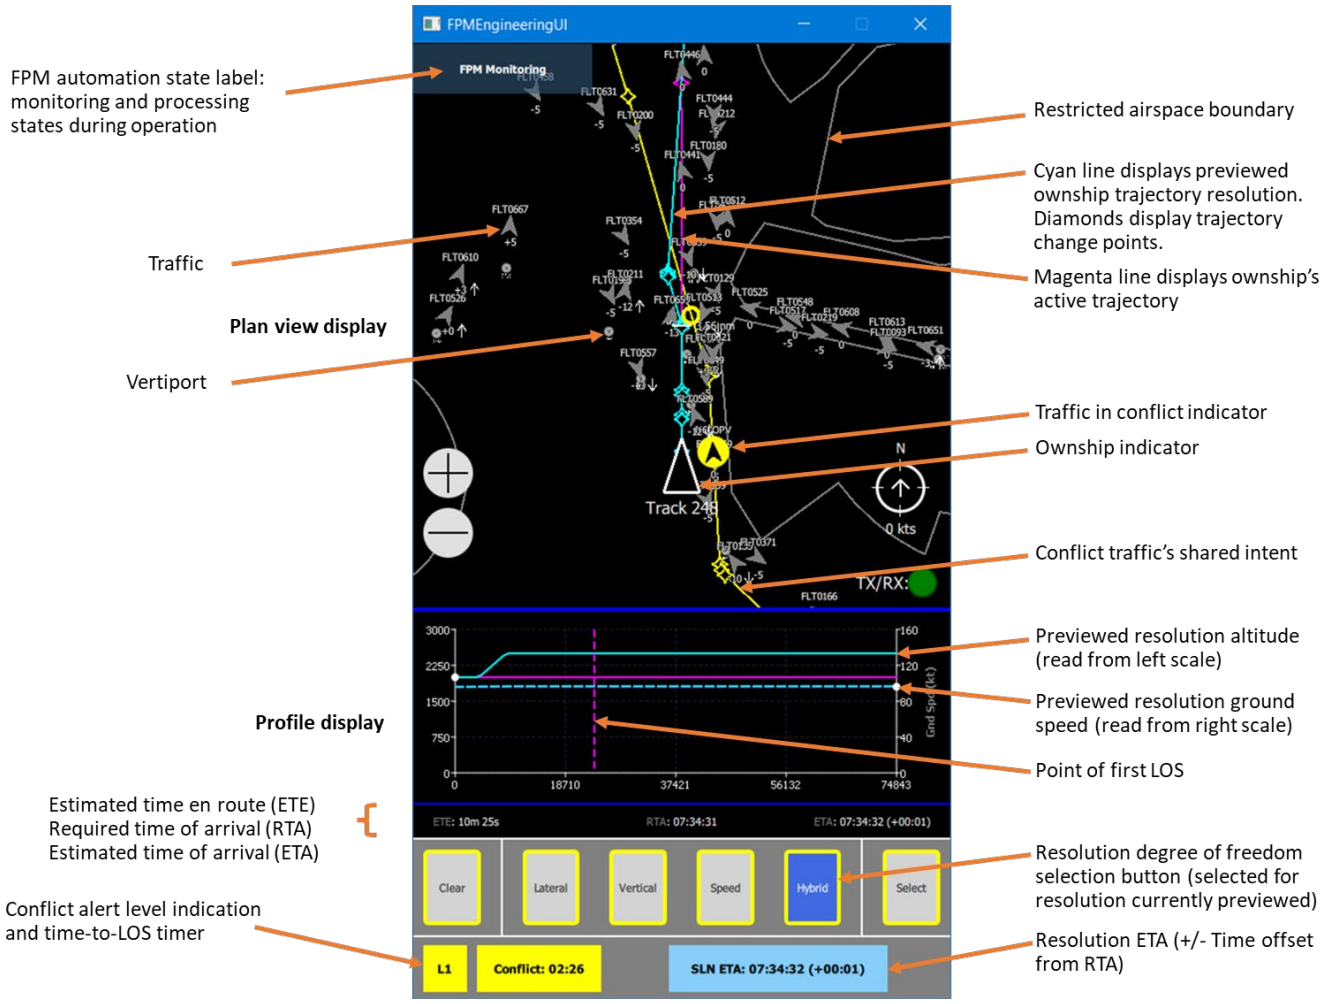


Fig. 1: FPM Engineering User Interface.
Credit: NASA

- provide data leading to further refinement of the automation technology and supporting research tools.

To achieve these goals and objectives, four research elements were defined.

1. Function Verification

   Determine whether expected performance is shown to occur in a relevant environment. Flight test scenarios corresponding to nominal operations were designed to generate data supporting the verification of expected behaviors and the resolution of known research questions.

2. Concept and Technology Advancement

   Provide data to support further research and development of the representative automation technology and the future development of system requirements. The nominal operations test scenarios were used primarily for this element, but stressor scenarios were also used to document behaviors on the boundaries of performance to shed light on the need for additional or more sophisticated automation functionality.

3. Discovery

   Unexpected technology performance and system behaviors may be revealed when tested in a relevant environment. The Discovery research element attempted to uncover some of these unknowns with the use of stressor scenarios. Stressor scenarios were designed to exceed expected performance with the intent that resulting data will reveal unexpected behaviors, trends, or failure modes.

4. Simulation Validation

   Provide data to be used in future validation and refinement of AAM simulations and simulation models. Simulations are needed in research to scale the number of operations to levels that cannot be performed in flight due to high cost. Simulation validations seek to increase confidence in results from these studies. Both nominal and stressor test scenarios were used to uncover model deficiencies and the bounds of simulation validity.

B. Methodology

In preparation for the testing, a model and simulation of a future UML-4 operational environment was developed. The representative FPM technology was modified for application to UAM UML-4 operations and was refined using batch and human-in-the-loop simulations in the two years prior to the test (Ref. 10). Initial flight tests were performed in 2022 and early 2023 to verify and refine ownship and intruder trajectory conformance precision and provide data for calibration of AOP's trajectory prediction uncertainty bounds (TPUBs), described in Ref. 6.

The flight test system was designed and implemented by the National Campaign Subproject's Integration of Automated Systems (IAS) element, with assistance from the FPM development team for specific functions. The test system incorporated agent-to-agent data links, including both vehicle-to-vehicle links and vehicle-ground links, based on an anticipated future operational configuration. Sikorsky's telemetry system hosted the vehicle-ground data links. The airborne component of the test system utilized flight guidance to allow both test aircraft to conform to a four-dimensional trajectory plan within a specified precision. A method for creating reliable conflict encounters in the presence of winds was also developed (Ref. 11).

A hardware-in-the-loop system integration laboratory was developed to provide a high-fidelity simulation of the flight test configuration, duplicating flight test hardware, supporting flight test software, and data exchange links. The laboratory was used to finalize flight test scenarios and all required recordings and configuration files, develop experiment procedures, and verify correct operation of the FPM, ground station, and experiment control functions.

Prior to testing, measures of effectiveness and performance were established for the study based on procedures and guidelines of the NASA System Engineering Handbook. Test scenarios were developed and prioritized to make efficient use of flight resources. The design and implementation of test maneuvers are described in Ref. 11. Tests were conducted and data were analyzed to draw conclusions for each of the objectives and each research element.

Creating the Future Operational Environment

The simulated UML-4 operating environment was developed based on studies conducted by NASA Ames Research Center and the Virginia Tech Air Transportation Systems Laboratory (Ref. 12). The virtual background traffic routes are based on the airspace surrounding Dallas Fort Worth Airport (DFW). The DFW environment was selected as the more challenging of the two cases studied by the reference. Unusable airspace is significant, making up 11.6% of total airspace. Restrictions create major obstructions to the two primary commercial airports.

Future demand was modeled and vertiports were identified and located based on mapping the anticipated demand to the region. Class B, C, and D airspace restrictions were added based on a NASA Ames study. Routes between vertiports were designed to minimize path distance while avoiding restricted airspace. Corridors were added to allow access to DFW by UAM operators and to facilitate less circuitous paths between vertiports. Fig. 2 illustrates the selected DFW environment used to simulate UML-4 operations in the test. DFW terminals are located slightly to the left of the figure center, served by four vertiports that are represented by four tightly grouped red markers. Yellow polygons represent airspace restricted to the eVTOL aircraft. Access to DFW terminals by eVTOL aircraft necessitates operation in corridors due to the close proximity of DFW runways.

Two eVTOL types were modeled, having nominal cruise speeds of 93 kt and 109 kt. Separation minima were assumed to be 1500 ft laterally and 450 ft vertically to allow for high-density operations. All aircraft in cruise fly at 1500 ft, 2000 ft, or 2500 ft geometric altitudes. Initial flight plans for all background traffic were direct routes between waypoints, with altitudes set based on the hemispherical rule. There is no requirement to adhere to the routes or hemispherical altitudes during the flight. All aircraft share their positions, altitudes, and velocities once per second. They also share their flight intent prior to takeoff and throughout the flight upon a change.

Flight Test Platforms

The Sikorsky Autonomy Research Aircraft S-76 helicopter (SARA) served as the ownship and was equipped with the reference FPM system. The Sikorsky S-70 Optionally Piloted Vehicle (OPV) served as the cooperative intruder. Both aircraft were controlled by human pilots on board. They operated with up to 330 virtual traffic aircraft flying routes defined by the UML-4 operating environment. The test aircraft are shown in Fig. 3.

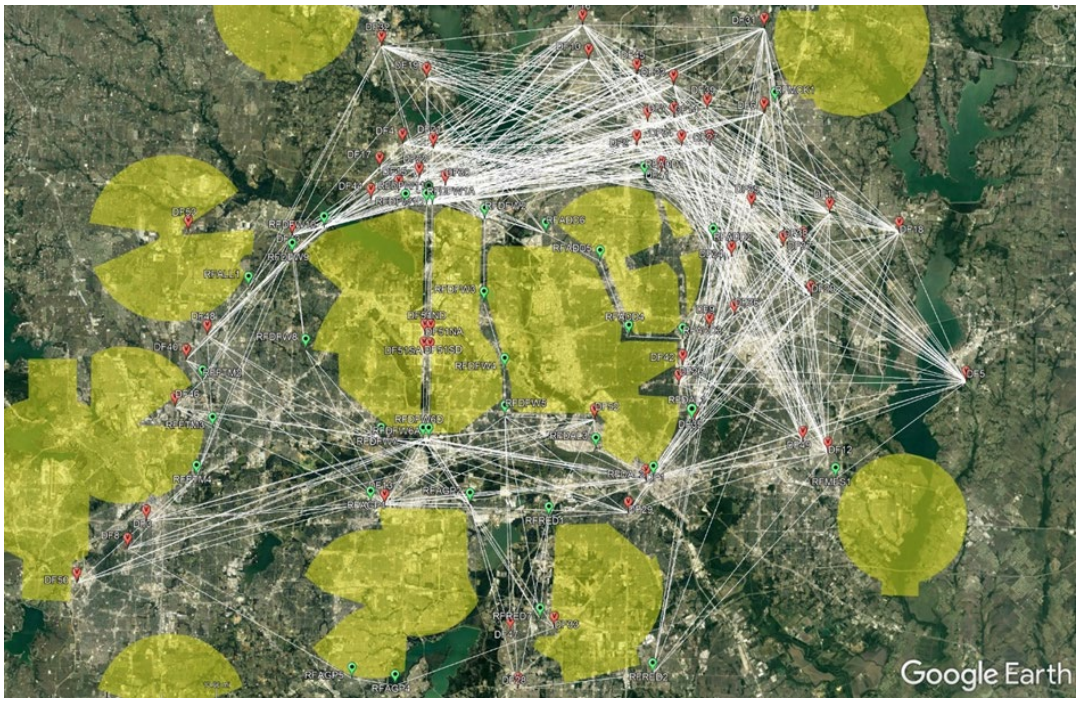


Fig. 2: Virtual UAM UML-4 operating environment.
Credit: NASA, Google



(a) **(b)**

Fig. 3: SARA Ownship (a) and OPV Intruder (b)
Photo credit: NASA.

System Overview

Each of the test aircraft were equipped with Sikorsky's Autonomy Mission Manager system (AMM), which permits the aircraft to automatically execute a 4D trajectory. The AMM system interfaces with NASA Middleware software, which provided 4D guidance to the aircraft based on 4D trajectory plans from AOP. The plans were provided either prior to flight (for OPV and SARA) or during execution (for SARA only, in response to trajectory change needs). Both aircraft were also equipped with two-way data links to a ground station, voice communication links, and dual-band ADS-B surveillance capabilities.

Key FPM Systems include the AOP Engine, the FPM Engineering UI, and an FPM Ground Station. The AOP Engine is the algorithmic engine of the FPM system under test. The AOP Engine's host platform was an Intel

NUC11TNHi7 that contained 64GB of DDR4 memory and a 1TB PCIe M.2 solid state storage device. As previously explained, AOP was responsible for generating 4D trajectory intent information for the host aircraft, detecting issues along the ownship's trajectory, and generating trajectory change solutions in response to the detected issues. Issues detected in the flight test configuration are a) conflicts with traffic, b) future violations of restricted airspace, and c) noncompliance with an assigned RTA. Traffic information needed by AOP consists of state updates at one-second intervals and event-driven intent updates. Although plans called for ADS-B to provide state information from the intruder, reliability issues led to the use of the Sikorsky telemetry system for both state and intent transfer from the OPV. The FPM Engineering UI, described earlier, was hosted by a Getac F110 tablet.

The FPM Ground Station provided situation awareness to researchers and engineers on the ground during the testing events. It was used to monitor system health and data collection processes during test execution. Other FPM systems hosted data important to conducting the research flight activities, including simulated background traffic, AOP configuration parameters, and environmental data such as winds aloft. Background traffic was prerecorded from simulations and was provided to AOP directly from the ownship rather than from an external source. This design decision reduced the burden on the Sikorsky telemetry system.

The IAS team at NASA Armstrong Flight Research Center (AFRC) provided other test components. These systems include the Middleware, the In-Flight Test System User Interface (IFTS UI), and the Ground Control Station (GCS). In addition to providing 4D guidance to the AAM, the Middleware software served as a data hub, passing data between the AOP Engine, the FPM Engineering UI, and the GCS. The Middleware also received ownship state information from the aircraft systems. The IFTS UI provided the flight crew with a human-machine interface that assisted in setting up a test card maneuver. The GCS provided a mechanism to define and transmit test card-specific maneuver and configuration data during test setup, and it provided situation awareness to researchers and engineers in the ground station facility during the testing events. Further details regarding the test system may be found in Ref. 11.

Flight Test Operational Environment

Test flights took place over Long Island Sound. The simulated UML-4 operating environment was overlaid on the test area and rotated as necessary to align virtual background traffic with the two test aircraft, as shown in Fig. 4. The test area, defined by the magenta box, is adjacent to Sikorsky Memorial Airport (KBDR). Research data collection for all maneuvers began and ended while the ownship was inside the test area. The red circle indicates the data link range limit assumed for the test, a 25 NM radius from KBDR.

C. Testing Campaigns

Testing was executed in a series of campaigns over a six-month period from February through October of 2023. Preliminary testing was also conducted in 2022. Each campaign involved test preparation, review and approval, system readiness verification, data collection, and an initial analysis of results. The first campaign was dedicated to verification of testing capabilities, including trajectory conformance verification for two aircraft. Aircraft performance profiles were finalized based on results. The second campaign was conducted in June 2023. It had the goal to verify end-to-end test system operation. The integrated system was tested in flight, with special emphasis on air-ground data link functions, ground station functions, background traffic functions, and test procedures.

Scripted baseline FPM function verification scenarios were then developed. Each test maneuver was made up of several traffic conflict encounter types between the ownship, the real intruder aircraft, and virtual background traffic, in the presence of simulated airspace constraints defined by the DFW database. Encounter types were made up of various level-flight incidence angles, including head-on encounters and in-trail overtakes, and encounters between the ownship and a climbing or descending intruder. All test maneuvers incorporated at least one trajectory change point between the maneuver start and the point of conflict. For instance, there would be a turn or altitude change for each aircraft after the beginning of the maneuver and before the point of first LOS. Other test maneuvers included scenarios involving a change of RTA for the ownship. The assigned arrival times were changed to require earlier or later arrivals by one minute. Some maneuvers included both an RTA change and a traffic conflict.

Another set of maneuvers was designed to support a parametric investigation to explore system behavior and resilience as a function of the automation's time parameters. Four time parameters were varied in matched sets. They were CD look-ahead time (referred to as CD time horizon), CR time horizon (the time ahead for which AOP resolved all conflicts), CR refresh cycle (time allocated to resolving and allowing for pilot execution before starting the next resolution cycle with updated information), and freeze horizon (the earliest future time a trajectory change is allowed to start). Baseline values were:

- CD time horizon: 180 seconds
- CR time horizon: 360 seconds



Fig. 4: DFW FPM operational environment overlaid on test area.
Credit: NASA.

- CR refresh cycle: 20 seconds
- Freeze horizon: 40 seconds

One stressor scenario included a form of blunder in which the intruder changes its intent inside the CD horizon, thereby creating a conflict with the ownship. Another set of stressors performed an initial exploration of deconfliction inside traffic corridors. A final stressor scenario investigated increases in traffic density from a baseline level of approximately 225 traffic aircraft.

The third campaign completed research data collection and was conducted in October 2023. All nominal test scenarios and various stressor scenarios were executed. Thirty-four test points were completed from a planned set of 40 maneuvers and 52 test points. The resulting data set produced dozens of encounters and hundreds of conflict resolutions. Table 1 summarizes the test groups and maneuvers associated with each group. Additional detail is available in Ref. 11.

V. Results and Analysis

The four research elements were analyzed independently. Select results are presented below for Elements 1, 2, and 4, Function Verification, Concept and Technology Advancement, and Simulation Validation, respectively. Surprisingly, Element 3, Discovery, revealed no major unexpected concept behaviors at the airspace system level. A few unexpected technology behaviors were revealed. These led to development of recommendations for functional refinements and the addition of future functions. Select results are presented due to space limitations. Results will be provided more comprehensively in a future publication. Major insights gained for all elements are summarized in the conclusions.

A data validity analysis was performed prior to the analysis of the elements. 4D trajectory conformance was verified for all maneuvers. Both aircraft flew all planned 4D trajectories with high precision. A software integration error that affected several test points was discovered during the post-test examination. AOP did not receive current winds information, so it substituted no winds for actual winds in its trajectory computations. Only test points that involved trajectories at speed limits were affected. The error caused AOP to incorrectly limit the airspeed of its resolution trajectories based on the minimum allowable airspeed specified by the ownship's performance model. Be-

Table 1. Test Groups and Maneuvers

Test Group	Description	Planned Maneuvers	Executed Maneuvers
1. Conflict Detection	Explore AOP's conflict detection functionality via intentional traffic conflicts with the intruder.	M1, M2, M3, M4, M5, M6, M7, M8, M9	M1, M2, M3, M4, M5, M6, M7, M8, M9
2. Conflict Resolution and Prevention	Verify AOP's conflict resolution and prevention functionality via intentional traffic conflicts with the intruder.	M1, M2, M3, M4, M5, M6, M7, M8, M9	M1, M2, M3, M4, M5, M6, M8
3. RTA Change Compliance	Test AOP's ability to detect and comply with RTA changes.	M10a, M10b, M12, M13a, M13b	M10a, M13a, M13b
4. RTA Change Compliance with Traffic Conflict	Evaluate AOP's ability to detect and comply with RTA changes that create a conflict with traffic.	M15, M16, M17	M15, M16, M17
5. Time Parameters Variation	Investigate the effect of altered AOP time parameters on detecting, resolving, and preventing conflicts.	M19, M20, M21, M22, M23, M24, M25, M26, M27, M28, M29, M30	M29, M20, M21, M22, M23
6. Intruder Intent Change Stressor	Examine conflict detection and resolution performance when the intruder's intent changes within the CD horizon.	M31, M32, M33, M34, M35, M36, M37, M38	M31, M32
7. Corridor Operations Stressor	Study AOP's ability to safely navigate airspace corridors using its current open airspace algorithms.	M39, M40, M41	M39, M40, M41
8. High Traffic Density Stressor	Explore functionality and behavior of AOP at higher UML-4 traffic density levels.	M42, M43	M42, M43

cause the starting position of each run was set based on current winds, AOP could not compute solutions in some cases. RTA solutions were most affected. All solutions that AOP did provide were flyable by the test aircraft. Compensation for the error involved removing some test points used in the analysis and performing simulations to understand the impacts of the error for several remaining test points.

A. Function Verification

Assessment was performed using functional performance benchmarks that are expected to be useful in developing future requirements. Automation performance benchmarks were also used to assess the current state of automation technology. Eleven measures of performance (MOPs) were defined, and one or more success criteria were assigned to each MOP. Success criteria values were based largely on engineering judgment and were intended as a starting point for future research. The success values varied from 90% to 100% of perfect performance, depending on the MOP. Researchers recognize that the 90% values are probably not high enough to meet future operational requirements. Researchers considered them appropriate for assessing the current state of the prototype technology.

Based on test results, ten of the eleven measures of performance met their success criteria. The CD function performed well, with no missed detections and a low number of false detections. AOP always detected conflicts within six seconds of the earliest time a detection could have occurred. The technology was able to provide high-quality CR advisories for 97% of all occurrences, and these were provided in the first resolution cycle. Resolution trajectories were always flyable within aircraft performance and procedural limits. No new conflicts were created by the resolution trajectories within the CR time horizon, and there were no repeat conflicts with the original conflicting aircraft. Conflict resolutions met a required time of arrival within three seconds for 92% of all occurrences, and resolving conflicts with traffic was successfully prioritized over RTA compliance when necessary. When RTA constraint relaxation was required, RTA compliance was relaxed to a minimum extent necessary. All CR computations were completed within 15 seconds, even for the high-density traffic scenarios. Pilots had at least 15 seconds to evaluate resolution advisories for 95% of all occurrences. The single unsuccessful measure of performance involved operations within airspace flow corridors. Only three of the four conflicts that occurred within corridors were resolved by AOP.

As expected, test data sets were too small to allow for statistical analysis. The results can only provide an initial indication of success. The MOPs and success criteria were designed to have future application in traffic simulation experiments, which can generate larger and more comprehensive data sets.

B. Concept and Technology Advancement

The advancement element provided data to facilitate future improvements, focusing on the representative automation technology and the FPM UML-4 concept. An overview of findings is presented below for four topic areas.

Conflict Detection

An analysis provided estimates of the number of true positive conflicts, false positive conflicts, true negative conflicts, and false negative conflicts that occurred over the course of the flight test. An AOP detection cycle occurred every three seconds. Results are shown in an error matrix in Table 2. A total of 4529 conflict detection cycles resulted from 41 test points. There were 34 conflicts detected during the test, and 1082 conflict detection cycles were associated with those conflicts. AOP uses TPUBs to drive the number of false negatives (missed alerts) low, potentially at the expense of increased false positives (false alerts). For the TPUBs settings used in the test, less than one percent of false positives were observed. There were no false negatives. The overall prediction accuracy was determined to be 99.78%.

Table 2. Conflict Detection Error Matrix

CD Cycles	Predicted True	Predicted False		
Actual True	True Positives 1072	False Negatives 0 (Type II error)	Precision 100%	False Omission 0%
Actual False	False Positives 10 (Type 1 Error)	True Negatives 3447	False Discovery 0.29%	Negative Prediction 99.71%
	Sensitivity 99.08% (True Pos. Rate)	Misses 0% (False Neg. Rate)	Prediction Accuracy 99.78%	
	Fall-Out 0.92% (False Pos. Rate)	Specificity 100% (True Neg. Rate)		

The test design should be kept in mind when viewing these results. All aircraft in the test were flying a ground-referenced 4D trajectory with high precision, and all aircraft conformed to their shared intent. TPUBs values were set based on the observed conformance precision of the test aircraft. Conflicts occurred by design for almost all test points. Encounters between the ownship and the intruder were also designed to have 0.2 NM lateral offsets for safety.

Conflict Resolution

Flight test results were analyzed to characterize performance of the conflict resolution function. Resolutions, when provided, were found to be flyable by the ownship and were effective in achieving separation. Other performance characteristics of interest include:

- availability of resolution solutions as the aircraft in conflict approach one another
- effectiveness of each of the four resolution degrees of freedom
- performance for encounters that involve an aircraft climbing or descending
- performance consistency across a range of lateral encounter angles.

The complete data set was made up of 111 traffic conflict resolution cycles from all 34 maneuver test points. Up to four resolution advisories were generated from each CR refresh cycle. Some runs involved the pilot selecting and executing a resolution advisory. For the other runs, pilots were asked not to execute a resolution. For the latter runs, resolution cycles were repeated (i.e., “refreshed”) as the range between aircraft decreased, until the run was ended. By not executing a resolution, a test run generated large amounts of data that revealed behaviors as a function of time to LOS. Two examples of the analysis follow to illustrate how results may impact both the technology and the concept designs.

Figure 5 provides an indication of the availability of resolutions as a function of time prior to first LOS, over the portion of the data set that used baseline time horizon settings. Although at least one resolution was available for 97% of test runs, up to four resolution degrees of freedom can be provided in each CR cycle. The figure shows a normalized

value of availability. If all four degrees of freedom are available in a CR cycle, the y-axis value would correspond to 100% for that cycle. The width of each bar corresponds to the duration of the CR cycle. The time of resolution is associated with the start of a cycle, so the first resolutions are shown just prior to the conflict detection horizon of 180 seconds. Baseline alert levels associated with priority rules are shown with extended tick marks on the x-axis. A Level 1 alert (180 seconds) corresponds to a point-out for awareness if the ownship has the right-of-way. A Level 2 alert (130 seconds) corresponds to the time the ownship operator should act if the intruder has not resolved the conflict. Alert levels were not used during the flight test because only the ownship was equipped with AOP, but alert levels are important when considering availability of resolutions.

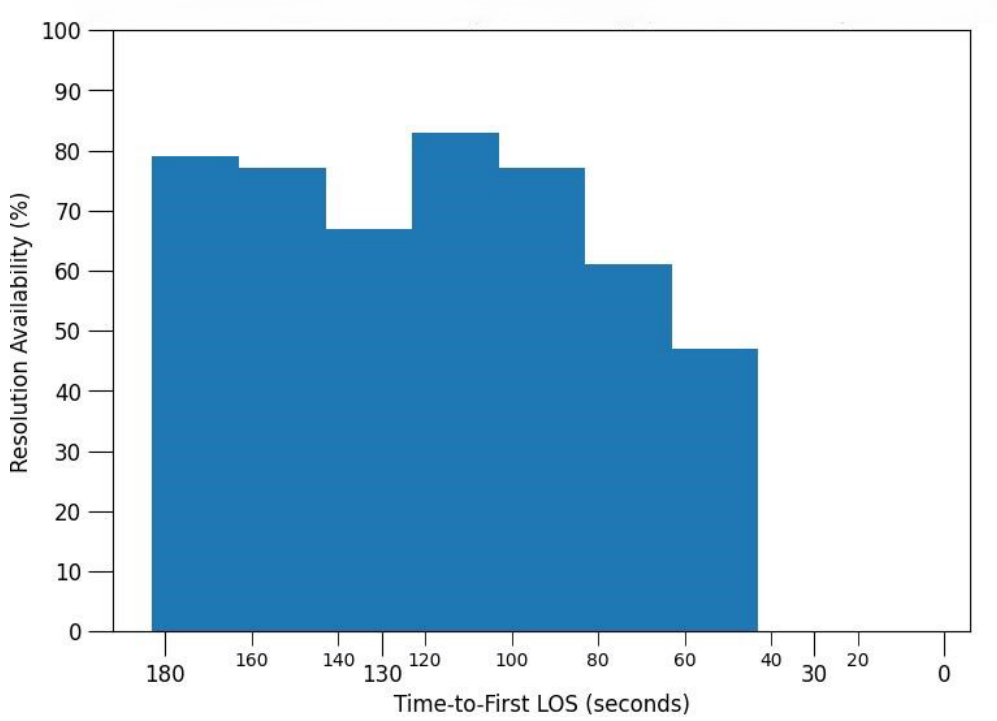


Fig. 5: Average fraction of resolutions available as a function of time to first LOS, for all runs that used baseline time parameters.

As expected, a decline is noticeable as time-to-first-LOS decreases, beginning around 80 seconds. The decline occurs within the Level 2 alert period. Further refinement of the technology may allow for future availability increases in this time range. If that is not possible, it may be beneficial to increase time horizons to allow for Level 2 alerting to begin earlier than 130 seconds prior to LOS. It should be noted that only one resolution is necessary. Multiple resolution advisories are offered to provide optimal solutions and accommodate user preferences.

An analysis was performed to determine times elapsed between a requested resolution computation and its completion for display to the pilot. The analysis data set contained a total of 213 CR cycles across 33 runs. A distribution of elapsed computing times for all CR cycles is shown in Fig. 6. The average processing time for each resolution advisory is 5.41 seconds. The longest resolution solution required more than 11 seconds, and the shortest about one second. Vertical and speed solutions were found to require less time than lateral and hybrid lateral-vertical solutions. This may be due to the time required to compute fitness penalties for conflicts, which are more complex when a lateral degree of freedom is involved. Average times were 2.74 seconds for vertical, 3.87 seconds for speed, 4.55 seconds for lateral, and 5.01 seconds for hybrid solutions.

Arrival Time Compliance

Flight test results were used to evaluate the reference technology's ability to meet an RTA at the destination vertiport. The FPM function is responsible for in-flight replanning to comply with an assigned RTA, even in the presence of traffic and airspace constraints. It is also responsible for complying with a changed RTA assignment. For each test maneuver involving a traffic conflict, AOP was responsible for complying with the ownship's assigned RTA

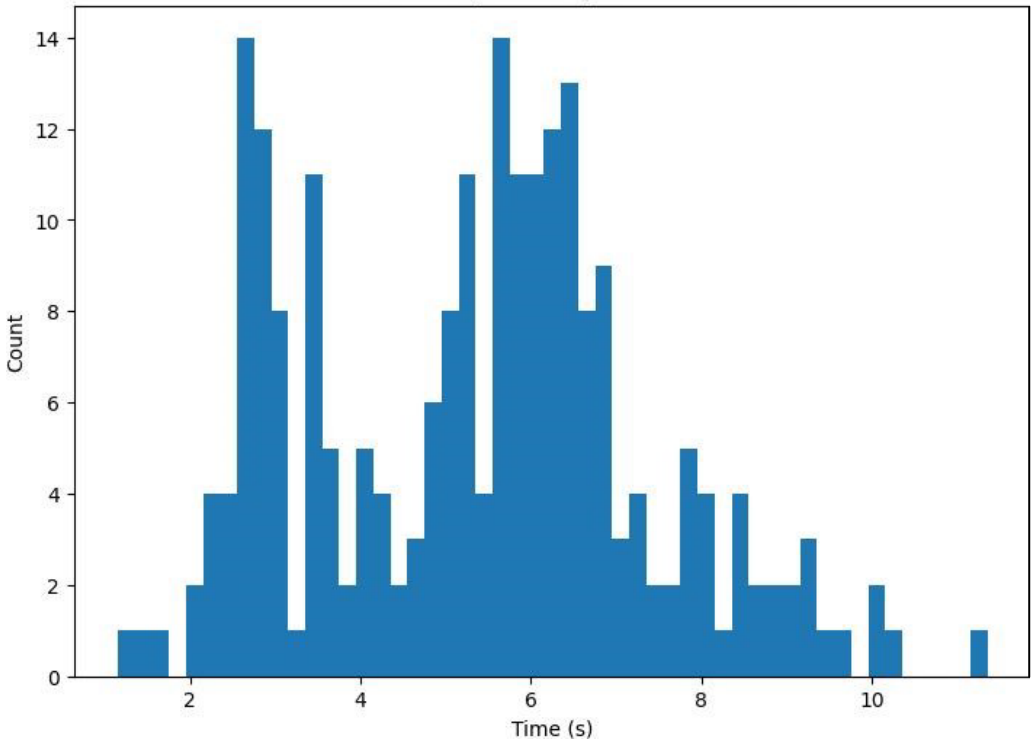


Fig. 6: Resolution compute times for all CR cycles.

while resolving the conflict. Some test maneuvers involved a change of RTA time of up to one minute. RTA compliance was defined for the test as the estimated time of arrival being within +/- 3 seconds of the RTA.

Figure 7 shows RTA conformance for all resolution advisories as a function of time to destination. Again, for most points in the data set, pilots were instructed to not execute an AOP advisory for the duration of the run, thereby creating multiple resolution cycles as the aircraft approached the destination point. Most advisories conformed to AOP's acceptable RTA tolerance, shown by the horizontal lines. The figure also shows some advisories creating late arrivals, which was expected, and others creating early arrivals, which was unexpected. Each of the last three refresh cycles for one run fell progressively later as the aircraft approached the destination, seen in the upper right part of the figure. Close to the destination, AOP was no longer able to generate a resolution advisory capable of absorbing a delay caused by the new RTA. AOP instead produced an advisory that would deliver the ownship to the destination as close to the RTA time as possible. Although such behavior was expected, quantifying the behavior is important in refining both the technology and the integrated FPM/arrival-scheduling concept of operations. The technology can potentially be improved by adding new maneuver patterns to the PBGA algorithm. Using current results, an initial rough estimate of a design parameter for the integrated concept can be derived: small RTA delay updates of up to one minute may be made by a scheduler up to five minutes prior to arrival, but not later. Regarding the unexpected early arrivals, a software defect in AOP was discovered during the analysis that caused AOP to improperly provide a few advisories that were out of tolerance. These advisories can be seen in the lower left area of the figure. They should have been filtered out, especially since other resolutions in the same CR cycle were within tolerance. Several similar small defects were revealed during the flight test. Small opportunities for design improvement were also identified. No major improvement needs were found.

Pilot Assessments

Two evaluation pilots were asked to provide subjective assessments of the reference automation, the vehicle-centric implementation, and the feasibility of UML-4 operations. Human/machine interface assessments specific to the FPM Engineering UI were avoided because the UI was used in lieu of an interface designed to support the operator's FPM roles and responsibilities. Further, a complete set of pilot roles and procedures was not assigned to

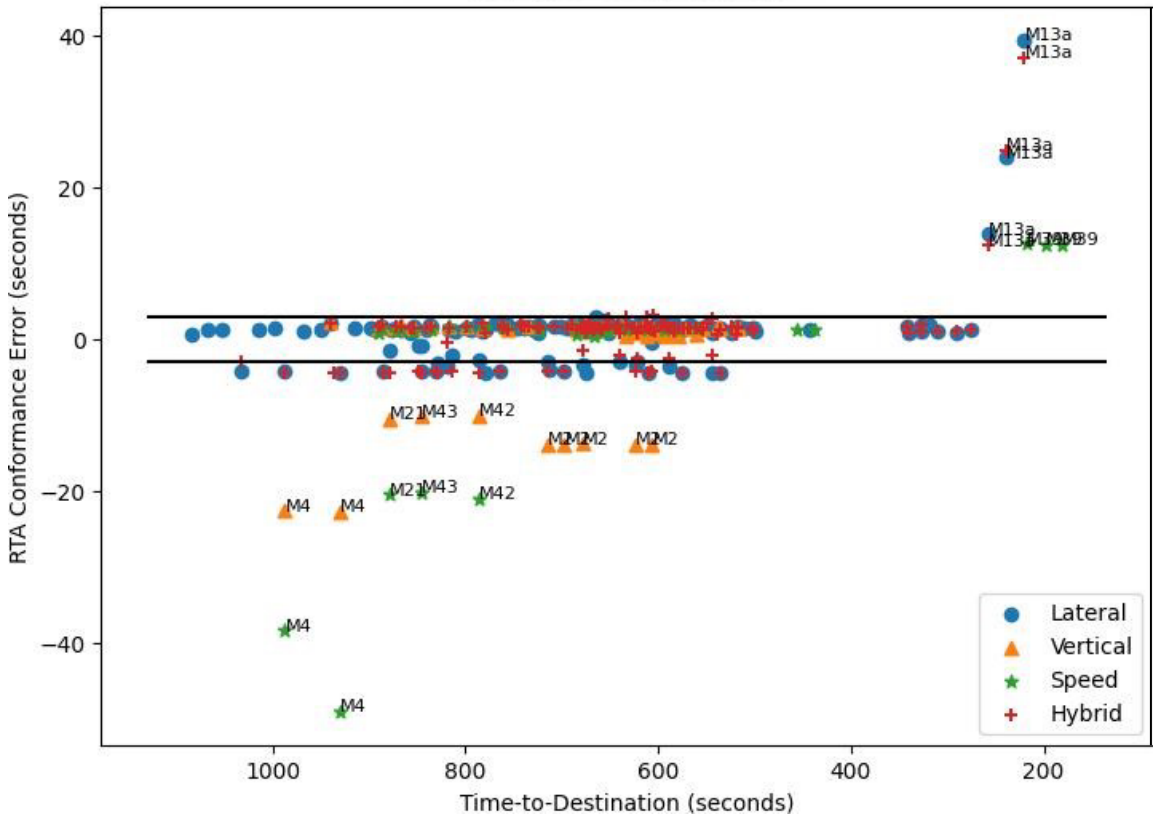


Fig. 7: RTA conformance error as a function of time to destination.

the evaluation pilots. They were asked to dedicate their attention to the FPM task during test runs, with the safety pilot responsible for most other tasks.

The FPM Engineering UI was designed to present a resolution advisory for each of the four degrees of freedom, with no ranking or recommendation to the pilot regarding which advisory to select. For each degree of freedom, the resolution solution with the best fitness score was presented regardless of its score relative to solutions for the other degrees of freedom. This created challenges for both pilots by increasing the time needed to evaluate and select an advisory. One pilot stated that the default time for selection of approximately 20 seconds was not enough to evaluate each advisory. The pilot stated 40 seconds (one value of a parameter varied in the test) was probably enough time. The other pilot made similar comments, but added that if only one resolution was provided, he would probably only need five seconds for evaluation. Both pilots said they would prefer the automation provide a single best solution recommendation. Other possibilities they proposed were to have the automation withhold presentation of advisories that provided little difference from other slightly better advisories, and to have the automation rank advisories and present them one at a time, starting with the top-scoring advisory.

Both pilots indirectly suggested additional information could be used by the algorithm to provide resolution advisories tailored to pilot preferences. As an example, one pilot mentioned that if he was flying at an altitude with good visibility, he would prefer to not select an advisory that would cause him to change altitude and fly into clouds. The other pilot provided the example of a need to avoid turbulence. As a general comment, he said his preferred resolutions depend on the situation he is in. For example, he would prefer not to climb if he is near his destination. He also mentioned he may have personal preferences, such as preferring turns over climbs, which suggests there may be value in adding a capability for the automation to accept individual preference inputs. When asked if passenger comfort would influence his preferences, one pilot mentioned he might prefer resolutions that provide greater separation from traffic and resolutions that are less aggressive, possibly causing him to reject speed-only advisories.

Both pilots acknowledged that the automation system had more information than they did about traffic and hazards in the environment, thereby creating resolutions that were not always intuitive to the operator. Based on their

comments, one pilot may have been willing to rely on trust that a mature system would utilize the information properly, while the other pilot requested more information be provided, allowing him greater situation awareness.

Operating within flow corridors was acknowledged by both pilots as being more challenging than operating in open airspace having a few airspace constraints.

The pilot assessments suggest that FPM automation or a similar function would be necessary in a future UML-4 operating environment. With a collaborative and responsible FPM function, routine operations at UML-4 seem feasible to the pilots. If an operator is in the loop, pilot feedback indicates the automation's advisories must be fully trusted to detect and resolve conflicts, with no expectation that the operator will verify correctness of an advisory. One pilot commented that in a visual environment, the 1500 ft separation standard assumed for the test seemed large enough.

C. Simulation Validation

An initial cursory exploration of simulation fidelity was conducted by replicating flight test conditions and executing test maneuvers using one of NASA's air traffic simulations. A cloud-based multi-aircraft simulation was chosen due to its scalability for high traffic density scenarios and its capability for all simulated aircraft to be equipped with AOP, both of which are necessary in future research. Flight and simulation data were compared, specifically replicating wind conditions observed in the flight test and duplicating the wind input error resulting in AOP not receiving the wind information. Figure 8 is an overview that illustrates the flight test comparison. The y-axis indicates the runs conducted in both simulation and flight. Most runs shown in the figure involved a pilot execution of resolutions, usually in the first refresh cycle, as can be seen in M15 and M32. In other runs, such as M7 and M10a, there was no resolution execution and repeated resolutions were provided by AOP over multiple refresh cycles. The wind condition is shown in parentheses. The flight data are always above the simulation data for a respective run, and the simulation resolutions have black outlines to further distinguish the two. The vertical dashed lines indicate the end of the flight data, and the red lines show the times the ownship was in conflict with another aircraft.

The simulation largely mirrored the flight test data with most runs, exhibiting identical outcomes regarding the number, type, and timing of conflicts and resolutions. However, some deviations between simulation and flight test data were observed. Maneuver M21 shows fewer computed resolutions in simulation compared to flight tests, likely influenced by incorrect wind input. Due to small differences between simulation and flight, the ownship operated on the boundary of a logic condition that triggered PBGA. The missing resolutions were primarily RTA resolutions. Additionally, maneuvers M15, M32, and M40 had slightly different conflict onset times with respect to the same maneuvers in flight. Maneuver M5 had an additional conflict with a virtual traffic aircraft during the flight test, but in simulation the conflict with OPV continued throughout.

To quantitatively assess similarity, a custom metric was developed to describe simulation similarity to flight test data – a figure of merit comprising conflict onset similarity, total cycle similarity, resolution type similarity, and resolution timing similarity. Due to space constraints, detailed definitions for each parameter are omitted in this paper. Based on calculations, the figure of merit yielded a 90% similarity between flight test and simulation data across all validation runs. The cloud-based simulation provides somewhat limited control of simulation timing and processor capability. With tighter control of these factors, the figure of merit is expected to increase significantly.

VI. Conclusions

This activity achieved the first successful flight test of airborne automation that performs strategic flight path deconfliction and in-flight planning to support traffic management. The test produced an initial verification of the automation's functions. It generated insights regarding FPM concept feasibility for future high traffic density operations. It also revealed needs for technology refinement and continued research. Major insights and feasibility conclusions are presented below.

A. Insights Gained

Major insights gained from the activity and considerations for interpretation of results include the following:

Technology Functional and Computational Performance

- The strategic conflict detection, conflict resolution, and conflict prevention functions performed well in flight in the near-open airspace conditions for which they were designed. All success criteria defined for the test were met. Resolution advisories were always flyable.

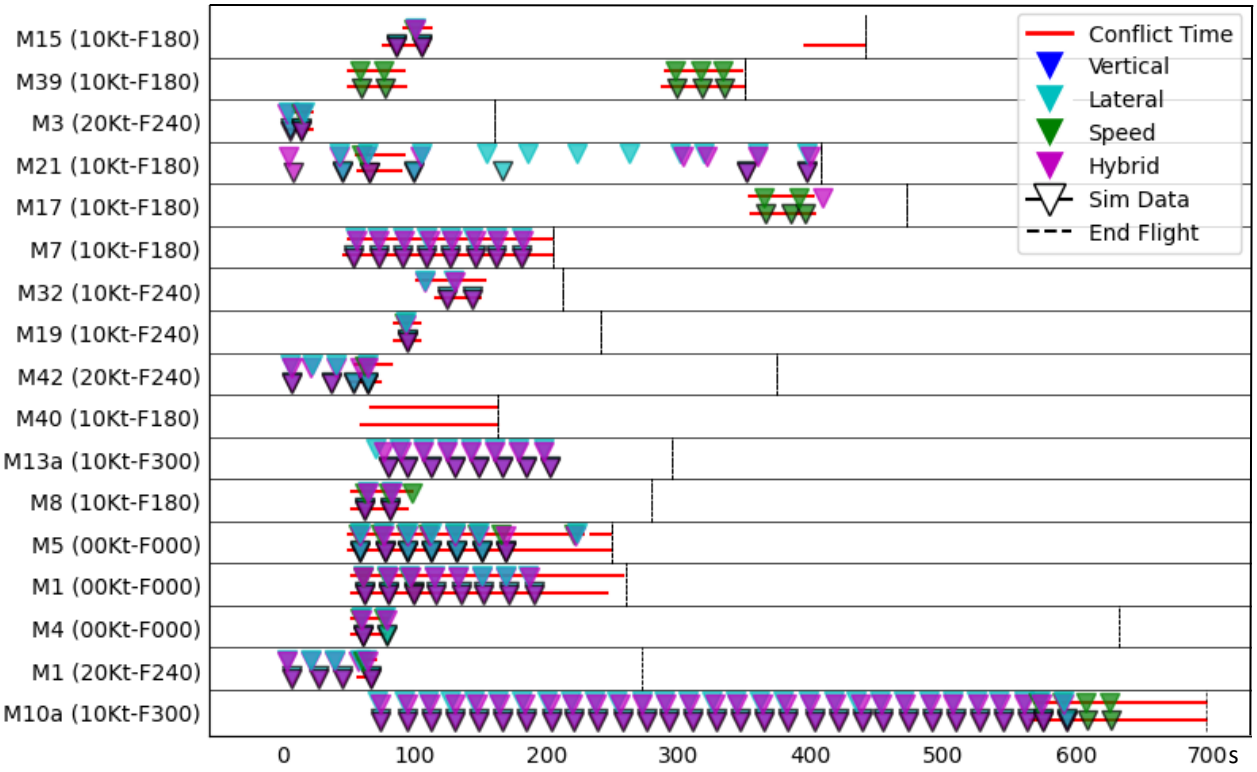


Fig. 8: Comparison of flight and simulation data.

- Strategic resolutions complied with assigned RTAs reliably. RTA change assignments up to one minute, the maximum tested, were handled by the automation up to approximately 300 seconds prior to the destination.
- No traffic density limits to the concept or the enabling FPM technology were found based on the traffic densities tested in simulation and flight.
- No future breakthroughs in computational power will likely be required. The prototype technology easily handled 330 traffic aircraft operating on a consumer-grade four-core processor. There appears to be no need to filter traffic to increase computing speed.

FPM UML-4 Concept

- The flight test produced no new concept behavior discoveries. Simulations conducted prior to the test were highly predictive of flight test results.
- The small separation minima assumed for the future UML-4 environment caused no issue. There were no missed alerts, very few false alerts, and no conflict resolutions that caused separation loss.
- Evaluation pilot feedback suggests human operators can be in or on the FPM loop at UML-4. Pilots had at least 15 seconds to evaluate and execute the automation's advisories for all time horizon parameter settings except the short settings. The short settings tested provided short, potentially unacceptable, resolution display time durations in a few cases.
- AOP never presented an advisory that pilots felt was unsafe or were uncomfortable executing. They did feel that some path changes AOP produced were not intuitive or natural for a pilot to make.
- Pilots expressed interest in having more time to allow for a more complete evaluation of each advisory, or to be provided with automated recommendations for the best option.
- The baseline, long-, and very-long-duration time horizon parameters tested revealed no issues in concept performance and no pilot comments regarding time preference.
- Traffic intent changes revealed no issues for changes as close as one minute prior to first LOS, the minimum tested.

- FPM functions will require substantial refinement for operations within airspace corridors. Test results are inconclusive, but a different approach may be required for high-density corridor operations.
- The ground-referenced four-dimensional aircraft guidance assumed for the concept and implemented for the flight test produced highly successful performance for all tested FPM functions. Satisfactory results may be possible without a requirement for four-dimensional guidance.

Technology Prototype

- The flight test revealed several small software defects and many small opportunities for technology refinement.
- The technology's constraint relaxation capability successfully optimized the level of relaxation for the constraint, but it did not consider other optimization criteria when doing so.
- An allocation of computing resources to balance loads across internal threads may further increase computing speed.

Simulation Validation

- While results were not identical, a multi-aircraft FPM traffic simulation reproduced concept behaviors observed in flight.
- The simulation used in the validation exercise produces outputs that are highly sensitive to computing power of host platforms.

B. Readiness Assessment

All functional performance success criteria for maneuvers in the open airspace environment were met or exceeded. The research prototype automation technology reliably supported the core functions of intent-based conflict detection, strategic conflict resolution, conflict prevention, and arrival time compliance. The technology-equipped test aircraft successfully maintained separation with all traffic. It also kept clear of virtual restricted airspace regions, which made up over 11% of the airspace environment. The FPM technology generated resolution trajectories that were always flyable within procedural and performance limits of the hosting aircraft. Through an engineering user interface that displayed resolution advisories and no other decision assistance, the automation provided pilots enough time to evaluate and execute resolutions. Test maneuvers involving operations within flow corridors were less successful. Results suggest a specific trajectory management approach and its enabling automation technology need to be developed to support high-density operations within corridors. The approach may involve aircraft in-trail spacing additions to the current capabilities. At a minimum, further simulation research and technology refinement is necessary to establish feasibility of the current approach, which was designed to support arrival metering.

C. Feasibility Conclusions

The flight test results provide positive indications that a vehicle-centric implementation of FPM can be made operational in the future. FPM has the potential to address critical challenges in managing the high-density traffic operations envisioned for UML-4 and above. Airspace resource management functions, such as arrival scheduling and aircraft sequencing, are additional challenges posed by high-density operations. FPM does not address these directly, but it may provide an additional benefit by reducing their development complexity. A vehicle-centric FPM function may effectively increase control power at the system level. If aircraft have greater capability and latitude to comply with scheduling constraints, the technical challenge of maintaining those constraints in the presence of disturbances may be less daunting.

Acknowledgments

The authors would like to thank all collaborators that assisted in this work, including members of the Airspace and Traffic Operations Simulation development team: Michael Havlin, Elaine Blount, James Sturdy, Doug Mielke, Joshua Kaurich, Patrick Johnson, and others; contributors to the development and testing of the System Integration Lab: Brent Pickering and Nathan Perreau; the IAS flight test team: Adam Yingling, Nancy Baccheschi, Jan Scofield, Andrew Guion, Ethan Williams, Paul Sampson, Gayle Patterson, Abbigail Waddell, Kevin Monk, Daniel Eng, and Ernest Nwajagu; the flight test pilots: Wayne Ringelberg, Scott Howe, David Zahn, Brent Davis, Bill Fell, John Rucci, and Rachel Simmet; and the flight staff and engineers at Sikorsky.

References

[1] National Academies of Sciences, Engineering, and Medicine, "Advancing Aerial Mobility: A National Blueprint," The National Academies Press, Washington DC, 2020.

[2] Federal Aviation Administration: Urban Air Mobility (UAM) Concept of Operations, v2.0. Federal Aviation Administration, 2023.

[3] Goodrich, K. H., and Theodore, C. R., "Description of the NASA Urban Air Mobility Maturity Level (UML) Scale," AIAA-2021-1627, January 2021.

[4] Sharma, V., and Wing, D. J., "Dynamic Path Planning Automation Concept for Advanced Air Mobility," NASA TM-20220009974, September 2022.

[5] Wing, D. J., Chancey, E. T., Politowicz, M., and Ballin, M. G., "Achieving Resilient In-Flight Performance for Advanced Air Mobility through Simplified Vehicle Operations," AIAA-2020-2915, June 2020.

[6] Karr, D. A., Wing, D. J., Barney, T. L., Sharma, V., Etherington, T. J., and Sturdy, J. L., "Initial Design Guidelines for Onboard Automation of Flight Path Management," AIAA-2021-2326, August 2021.

[7] Ballin, M. G., Sharma, V., Vivona, R. A., Johnson, E. J., and Ramiscal, E., "A Flight Deck Decision Support Tool for Autonomous Airborne Operations," AIAA-2002-4554, August 2002.

[8] Wing, D. J., and Cotton, W. B., "Autonomous Flight Rules: A Concept for Self-Separation in U. S. Domestic Airspace," NASA TP–2011-217174, November 2011.

[9] Karr, D. A., Sturdy, J. L., Barrows, B. A., and Ballin, M. G., "An Experimental System for Strategic Flight Path Management in Advanced Air Mobility," AIAA-2024-0710, January 2024.

[10] Barrows, B. A, Ballin, M. G., Barney, T. L., Nelson, S. L., Underwood, M. C., and Wing, D. J., "Sim to Flight: Evaluating Flight Path Management in High Density Urban Environments," AIAA-2024-1077, January 2024.

[11] Nelson, S. L., Ballin, M. G., Barrows, B. A., Underwood, M. C., Wing, D. J., Williams, E. R., and Sturdy, J. L., "Designing a Flight Test of a Flight Path Management System for Advanced Air Mobility Research," AIAA-2024-1076, January 2024.

[12] Rimjha, M., Hinze, N., Hotle, S., Swingle, H., and Trani, A., "Demand Forecast Model Development and Scenarios Generation for Urban Air Mobility Operations," May 2021. Internal report, NASA NIA Contract NNL13AA08B, TO NNL16AA36T.