Initial Design Guidelines for Onboard Automation of Flight Path Management

David A. Karr, PhD*
National Institute of Aerospace, Hampton, VA 23666, USA

David J. Wing,† Terique L. Barney,‡ and Vivek Sharma, PhD§
NASA Langley Research Center, Hampton, VA 23681, USA

Timothy J. Etherington¶
Analytical Mechanics Associates, Hampton, VA 23666, USA

James L. Sturdy‖
KBR Inc., Hampton, VA 23681, USA

Achieving the National Academy of Science's vision of advanced aerial mobility will depend on significant developments in automation to achieve safe and efficient operations. Flight path management (FPM) is a major category of automation functionality needed to achieve this vision. FPM will provide dynamic management of an aircraft's flight path, ensuring that it continues to meet five objectives: it is feasible to fly to mission completion, deconflicted from hazards, coordinated with other traffic, flexible to accommodate future disturbances, and optimized to meet business objectives. While efforts are underway to advance FPM technology for the Urban Air Mobility application, the authors present initial design guidelines for FPM automation capabilities to achieve each of its five objectives based on 15 years of prior FPM automation research and development. We describe methods to efficiently account for uncertainty in the prediction of trajectories and discuss additional considerations for prioritizing safety in the design of FPM automation capabilities and interactions between aircraft. We make recommendations supported by extensive experience gained via previous work with the Autonomous Operations Planner, an FPM reference automation system developed by NASA. By employing capable FPM automation supported by cooperative operational flight rules and information sharing, future aircraft operators will benefit from an increased ability to plan and execute safe and efficient flights to achieve mission success in a dynamic airspace.

I. Introduction

Fully achieving the National Academy of Sciences' vision of advancing aerial mobility [1] will likely require onboard flight path management (FPM) automation operating in partnership with ground systems that supply complementary services. Onboard FPM automation will assist in, and under future operational flight rules may be responsible for, managing various aspects of an aircraft's flight path to ensure safety in a traffic environment and conformance to operational constraints [2]. In the application of Urban Air Mobility (UAM), for example, a traffic management system is envisioned where fleets of electric vertical takeoff and landing (eVTOL) vehicles rely on UAM operators to manage flight paths and traffic separation through a combination of onboard vehicle automation capabilities and ground-based services from so-called Providers of Services to UAM (PSU). Early UAM operations may be restricted to specified air corridors [3], whereas later operations would also be conducted in widespread operational areas [4]. Even though UAM operations are expected to operate independently from the traditional air traffic control (ATC) system, which provides separation and other services to aircraft operating under Instrument Flight Rules (IFR), UAM operational

*Principal Engineer, AIAA senior member.
†ATM Research Engineer, Crew Systems & Aviation Operations Branch, AIAA associate fellow.
‡ATM Research Engineer, Crew Systems & Aviation Operations Branch.
§Aerospace Research Engineer, Crew Systems & Aviation Operations Branch.
¶Technical Fellow, Simulation and Visualization, AIAA Senior Member.
‖Software Systems Engineer.

concept documents are not yet fully specific in defining the entity ultimately responsible for separation. Instead, these operational concepts refer to the various entities (e.g., PSU, UAM vehicle, UAM operator) that are expected to have contributing roles in assuring a safe flight path that separates the vehicle from traffic aircraft and other airspace hazards. Each of these entities may therefore be required to employ various FPM capabilities (discussed later in this paper) in their automation systems.

As aerial mobility concepts evolve through detailed examination and community engagement, and new flight rules are developed to support UAM and other emerging operations, greater clarity will emerge regarding where the responsibilities and authorities reside for flight path safety. Regardless of the outcome of these concept-level decisions, onboard automation functionality is expected to play a critical operational role. For instance, if the UAM concept were to dictate that safety responsibility resides with a PSU or UAM operator dispatch service (requiring them to incorporate FPM functionality in their automation systems), and the PSU or dispatch transmits safety-critical flight path instructions to the vehicle, the Pilot in Command (PIC) must have the ability to assess and then either accept or reject the instruction as the flight's final authority [5], and will likely require the support of onboard FPM automation using data local to the vehicle for that assessment (regardless of whether the PIC is onboard the vehicle). If the communications link to the ground station is disrupted, the vehicle's assessment-support automation now becomes primary and must perform all required functions to assure flight path safety independent of ground support. Alternatively, the concept may allow for the UAM operator to manage the vehicle's flight path safety directly through onboard automation [6]. While this distributed architecture reduces the dependency on communication links and enables greater scalability, it may also increase the certification requirements for the onboard automation.

Such onboard FPM automation systems require careful design and development, given that they will have a non-traditional, safety-critical role in ensuring that the vehicle's flight path remains sufficiently clear of hazards (including proximate traffic) and be coordinated in the airspace. While current flight management systems are capable of ensuring that a specified flight path is contained within the aircraft's performance envelope and conforms to certain operational constraints (e.g., crossing altitudes, speeds, and times), they are not designed to assure safety of the flight path with respect to external hazards such as traffic aircraft, weather, and special use airspace (SUA). Furthermore, they are not designed to replan and coordinate the flight path safely in an active airspace environment with dynamically changing hazards. These functions will need to be performed by advanced FPM automation systems with access to broad information regarding the vehicle and the airspace environment and having a purposeful design and certification that enables the automation to be fully responsible for the functions for which these systems are designed (i.e., the certification cannot rely on human intervention). In this regard, the UAM Maturity Level (UML) scale defines UML-4 as relying on "collaborative and responsible" automation to enable operations in non-visual conditions with medium traffic density and medium complexity [7]. In its mature state, FPM automation is intended to partially fulfill this role (along with other automation supporting functions such as flight control, systems health monitoring, and detect-and-avoid (DAA)).

The NASA Advanced Air Mobility (AAM) Project aims to accelerate air mobility concepts and automation by delivering validated system architectures and requirements. Included in this goal are developing and demonstrating key automation functions, with a specific emphasis on Automated Flight and Contingency Management (AFCM). AFCM focuses on vehicle automation functions that safely enable simplified piloting and high-density flight operations in low-visibility conditions and obstacle-challenged urban environments. As an output to stakeholders, AAM AFCM research intends to produce design guidelines, standards, and methods of compliance for select vehicle automation systems, pilot interfaces, and operational requirements. FPM will be one of the key focus areas of AFCM research.

In part, FPM design guidelines to industry will be initially informed by knowledge gained through previous research and development (R&D) on advanced vehicle automation systems designed to enable vehicle autonomy in the National Airspace System. The Autonomous Operations Planner (AOP) is a NASA-developed research prototype of an integrated, onboard FPM automation system designed to manage the safety and efficiency of the host vehicle's flight path in high-density, dynamic, complex airspace [8]. Developed initially for commercial transport operations, AOP is adaptable to the UAM mission and eVTOL vehicle type (an activity currently underway in the AAM Project). As a research tool, AOP features a broad suite of integrated FPM functionality to help illuminate design requirements for onboard automation systems responsible for flight path safety and conformance to operational constraints. Though AOP has so far only been exercised in a simulation environment, it was developed with enough sophistication and fidelity to be operable in flight. In fact, a derivative of the AOP software, the Traffic Aware Planner (TAP), a pilot's advisory tool for flight path optimization, was operationally tested in airline revenue service and is undergoing commercialization by industry [9]. Given its integrated FPM functionality and high fidelity, AOP is ideally suited to begin the process of informing the AAM community on the functional design of onboard FPM automation systems for the UAM mission. As such, it serves as an initial "reference" automation system for AFCM and for this paper.

In this paper, the authors document initial functional design guidelines, at a summary level, for onboard automation systems performing or supporting FPM in the AAM mission. Though the focus is on UAM, the guidelines presented here will be applicable to many future operational contexts that exercise various degrees of vehicle operational autonomy. In the following sections, we first put FPM in context with other functions of automation in AAM, then describe five high-level functional objectives for onboard FPM automation. We introduce AOP as a reference system for FPM automation, then discuss functional capabilities associated with the five FPM objectives, making key design recommendations based on the experience gained through extensive R&D with AOP. We then make design recommendations to prioritize safety across the five capability sets. The paper concludes after a discussion of future work.

II. The Context of Flight Path Management

Wing et al. [10] abstracted the flight management functionality relevant to the roles of onboard automation and human pilots for AAM operations into four representative categories of functions: mission management, flight path management, tactical operations, and vehicle control. The four categories are illustrated in Fig. 1.

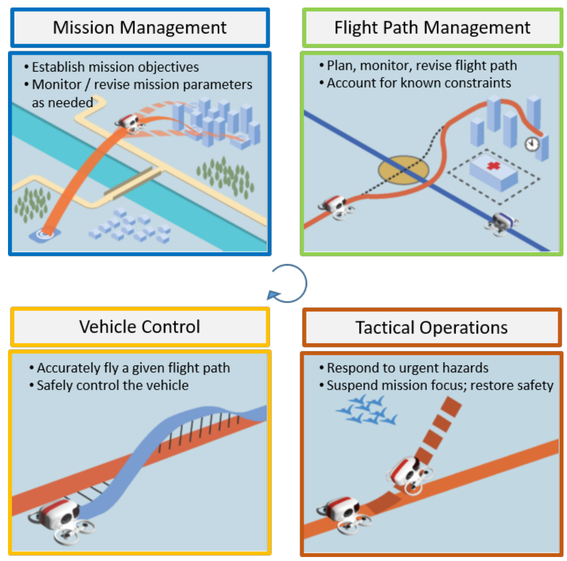

Fig. 1 The context of FPM within a framework of four categories of flight management functionality.

- **Mission Management.** Mission management is concerned with decision-making regarding the achievement of mission goals and the management of mission parameters. A mission is defined here as the safe transport of an air vehicle to an intended destination within a set of mission parameters. Such parameters can range from minimizing turbulence to arriving reasonably close to the advertised time. Mission management generally operates one level removed from the specific flight path and includes monitoring vehicle systems, destination availability, and energy reserves; adjusting mission parameters under evolving conditions such as a change in flight optimization priorities; and maximizing mission success by developing contingency and risk mitigation plans. While these functions do not manage the actual flight path, they leverage FPM in predicting arrival states and generating viable candidate contingency plans.
- **Flight Path Management.** FPM functionality directly supports decision-making for the current flight path and any prospective flight path changes, within the context of established (and potentially evolving) mission parameters. These functions maintain awareness of trajectory constraints in the airspace environment and they monitor the current flight path for conflicts that would compel a change or for optimization opportunities that could improve the flight within known constraints. Safely replanning in response to emergencies is a critical function of FPM. FPM is concerned with completing the mission safely, coordinating with other airspace users, and optimizing the flight path within the available degrees of freedom.
- **Tactical Operations.** As a layer of safety independent of FPM, a separate category of functions detects and responds immediately to near-term and generally unshared hazards. Categorized here as tactical operations, and identified in the AAM Project as hazard perception and avoidance, these functions ideally use data sources independent of FPM to perceive hazards and detect conflicts along the aircraft's immediate trajectory and to produce guidance for a rapid response to avoid the hazard. In these situations, mission objectives are suspended until safety is restored. Once the aircraft is on a safe and stable path, FPM functionality may be invoked to establish the aircraft on a flight path that recovers mission objectives.
- **Vehicle Control.** Vehicle control addresses the execution of the planned flight path but not its generation or modification. These functions monitor the aircraft's guidance and control systems, conformance to stable flight, and navigation performance, while responding to flight deviations with controls appropriate for the functioning

condition of the vehicle. Any changes to vehicle performance (e.g., resulting from a failed motor) would be communicated to FPM to assess the feasibility of the current flight path and to ensure newly generated flight paths account for the altered performance envelope.

Given this context, we focus this paper on the FPM category of flight management functionality. Section ref-sec:capabilities defines a representative set of capabilities and functions of an onboard FPM automation system. Section IV introduces our reference implementation of an FPM automation system, which is the basis of design guidelines presented in Section V through Section X.

III. The Capabilities of Flight Path Management

The goal of flight path management is to ensure that a safe and operationally acceptable flight path for a vehicle is available to the vehicle operator and flight control system through the life of the flight. An FPM automation system achieves this goal by meeting one or more of the following objectives: constructing a feasible path (discussed in Section V), deconflicting the path (Section VI), coordinating the path with other airspace users (Section VII), preserving the flexibility of the flight path (Section VIII), and optimizing the flight path (Section IX). A fully capable FPM automation system therefore simultaneously creates a *feasible*, *deconflicted*, *coordinated*, *flexible*, and *optimal* path, monitors it continuously, and revises it automatically in response to evolving mission objectives, constraints, and other stimuli. This process ensures that the flight path remains safe and operationally acceptable, continues to meet mission objectives, and is sufficiently adaptable to manage potential contingencies.

Figures 2–6 depict examples of flight paths meeting each of the five key objectives of FPM. In summary:

- A *feasible* path (Fig. 2) is one that is consistent with the aircraft's performance capabilities in the predicted atmospheric conditions and complies with restrictions due to airspace structure, requirements of operational rules, and aerodrome constraints.
- A *deconflicted* path (Fig. 3) is one that avoids unsafe proximity to known traffic, terrain, obstacles, weather, and airspace hazards.
- A *coordinated* path (Fig. 4) is one that harmonizes with other airspace users for local and system-wide safety, capacity, efficiency, and equity.
- A *flexible* path (Fig. 5) is one that proactively maximizes robustness and adaptability to evolving mission objectives, constraints, and external stimuli.
- An *optimal* path (Fig. 6) is one that achieves business objectives and preferences of the pilot and fleet operator.

An onboard FPM automation system (hosted by and focused exclusively on an ownship) generates a four-dimensional flight path (for flying), the associated operational intent (for sharing), and an energy consumption profile (for monitoring). This system also provides users with contextual information, flight conformance status, system health and alerts, and contingency response options. The generated flight path accounts for mission objectives, pilot preferences, aircraft performance capabilities, ownship and traffic flight information, current and predicted atmospheric conditions, obstacles and terrain constraints, relevant operational rules, and airspace configuration. Many functions contribute to meeting these objectives of the automation system, including (but not limited to):

- four-dimensional (4D) flight path construction;
- energy consumption prediction;
- conflict detection, prevention and resolution;
- coordination mechanisms;
- contingency response generation;
- flight path change evaluation;
- constraint conformance and energy reserves monitoring;
- automated flight path revision;
- flight path optimization; and
- interactive user interfacing.

An operator may have a choice where to host the FPM automation system: onboard the vehicle, in a ground-based control station, or some combination of the two. It may also be an option for a third party to manage a *feasible*, *deconflicted*, *coordinated*, *flexible*, and *optimal* flight path on behalf of a fleet operator. However, an FPM automation system hosted primarily onboard the vehicle (perhaps with some ground-station redundancy) benefits not only from direct access to onboard real-time data and the shortest loop-closure time but also from access to the fleet operator's proprietary business goals (via the pilot) and the vehicle's proprietary performance model. Such a system would be capable of computing flight paths with greater accuracy and tighter conformance than a system used by a third party

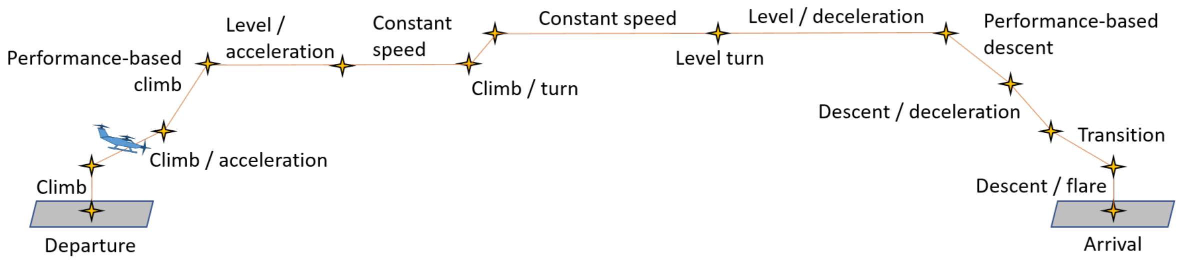

Fig. 2 Creating a feasible flight path is accomplished through trajectory generation capabilities accounting for aircraft performance, energy use, atmospheric conditions, airspace design, and destination constraints.

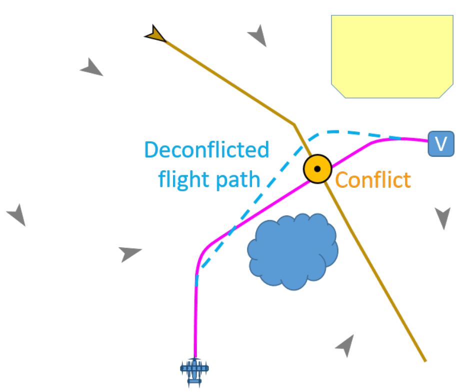

Fig. 3 Deconflicting a flight path is accomplished through trajectory prediction, conflict detection, and resolution, taking prediction uncertainty into account.

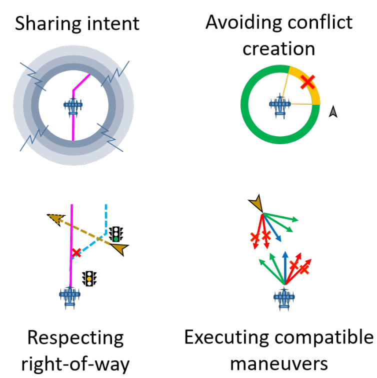

Fig. 4 Coordinating a flight path is achieved in a variety of ways such as sharing intent, avoiding conflict creation, respecting right-of-way rules, and maneuvering in complementary directions.

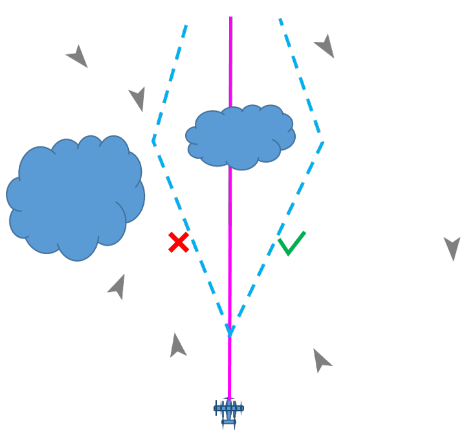

Fig. 5 Preserving flexibility in a flight path involves placement of the flight path to maximize ability to accommodate future disturbances, thereby increasing mission completion probability.

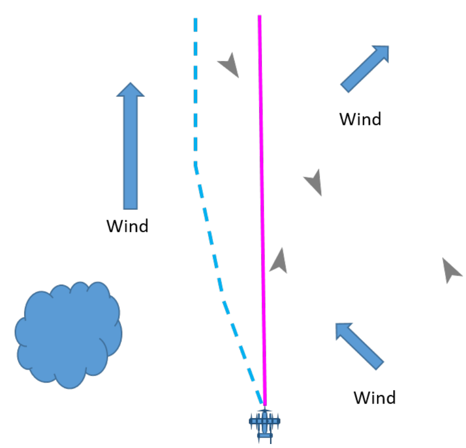

Fig. 6 Optimizing a flight path involves seeking the best available cost-benefit using the degrees of freedom remaining after other objectives have been met.

without access to proprietary data. These attributes of computed flight paths result in less conservative planning, reduced airspace buffer requirements, tighter separation standards, and improved responsiveness, all of which contribute to increased system capacity and efficiency. The reference automation system described next is an example of an onboard FPM system.

IV. Reference Implementation of Flight Path Management Automation

The design guidelines presented in this paper at a summary level are derived primarily from the knowledge and experience NASA has achieved in approximately 15 cumulative years of development, testing, and refinement of a sophisticated research-prototype FPM automation system, the Autonomous Operations Planner (AOP) [8], which was designed to enable vehicle operational autonomy within the research concept of Autonomous Flight Rules (AFR) [11]. In this concept, AFR aircraft, while in the en-route phase of flight, flexibly manage their own flight path to achieve their mission objectives, conform to operational constraints, and assure separation from traffic and airspace hazards. They operate without segregation from non-AFR aircraft, applying appropriate rules of engagement with non-AFR aircraft and with other AFR aircraft. As in UAM, AFR aircraft operate largely independently from ATC except for departure and arrival operations.

Through iterative R&D, AOP implemented prototype functionality for each of the five capability sets described in the previous section. The design philosophy for AOP was to err on the side of higher fidelity and excess functionality where possible, rather than build too little and risk an incorrect research finding of infeasibility of AFR operational autonomy. The design approach was to develop an integrated suite of FPM functionality compatible with current-day commercial transport avionics architectures and current or anticipated industry standards. The integrated functionality enabled evaluation by current airline pilots in a familiar and realistic (but simulated) operating environment (namely, existing commercial transport aircraft conducting typical airline missions, but under AFR rules).

Experiments were conducted both in human-in-the-loop [12, 13] and in batch simulation modes [14–18]. In the former case, airline pilots were trained with procedures that placed FPM automation (i.e., AOP) in a responsible role; the pilot's role was to make mission-level decisions and select between alternative flight path solutions offered by the FPM automation. Therefore, high levels of functional reliability were needed, which was achieved in part through extensive AOP batch testing and experimentation, using a pilot model in lieu of a pilot. These batch simulations were designed to stress-test AOP in large, randomly generated, high-density, multi-vehicle simulations with various real-world constraints (such as limited information sharing, saturated surveillance characteristics, excessive execution delays, and uncertain winds or weather). Though many research questions were yet to be explored and associated AOP functionality development remained incomplete, the levels of functionality and maturity achieved in AOP make it well suited to serve as a reference system to provide initial design guidance to industry for commercial development of FPM automation.

V. Constructing a Feasible Flight Path

Constructing a feasible flight path is accomplished through trajectory generation capabilities that account for aircraft performance, energy use, atmospheric conditions, airspace design, and destination constraints. It considers the vehicle and its capabilities in achieving a specific mission in the absence of any traffic or airspace hazards. Several requirements must be met in creating and maintaining a feasible flight path:

- The vehicle must not operate outside its performance abilities or safety limits.
- The vehicle must conform to operating procedures of the airspace and to the operating procedures of the fleet operator.
- The vehicle must satisfy all explicit constraints placed on its flight, such as a required time of arrival (RTA) at its destination.
- The vehicle must satisfy all constraints required by regulations or operating procedures, for example a speed limit below a certain altitude.
- The vehicle must complete its flight with adequate energy reserves and land safely.

The ability to construct and maintain a feasible flight path depends on the quality of trajectory generation (TG) functions implemented in the FPM automation system. The TG functions provide predictions of the ownship's flight path that conform to constraints and accurately reflect the flight performance characteristics of the vehicle. The TG is the central functionality within an FPM automation system, as it supports all five capability sets. Several key applications and associated design considerations of the TG are discussed in the following subsections.

A. Four-Dimensional Trajectory Prediction

The primary TG function is the conversion of a vehicle's *strategic intent* with constraints into a complete and accurate *four-dimensional trajectory* (4DT). For a vehicle that has not yet taken off, the operator produces a strategic intent that may include, among other data, an origin aerodrome, a takeoff time, desired cruising altitude and cruising speed, a destination, additional constraints such as an RTA, and any other operationally appropriate constraints defined by arrival and departure procedures. The essence of a 4DT is that it predicts the latitude, longitude, and altitude of the aircraft as functions of time, considering vehicle performance and operator preferences.

The 4DT of an airborne aircraft is considered complete if it predicts the path from the aircraft's present position to the destination or to a time or distance horizon defined by the particular client that requested a 4DT. A 4DT computed prior to takeoff starts from the instant of takeoff based on an assumed departure time.

The required accuracy and therefore fidelity of the 4DT depends upon its purpose. In the implementation of AOP, multiple TG algorithms have been used. In early versions of AOP, the primary purpose of the TG was to support research on deconfliction and coordination in a single simulation environment. To ensure FPM solution feasibility, AOP employed the same TG used by the research prototype flight management system (FMS) in the simulation environment. Thus, any solution produced by the conflict resolution algorithm could be loaded directly into the FMS and executed, and the (simulated) aircraft would be capable of flying it.

In the AOP-derivative software application, TAP, the research focus was flight path optimization, where accurate predictions of flight time and energy use for the en-route portion of flight were the most important factors. In addition, during flight testing and operational trials, TAP was required to be quickly adapted to multiple aircraft types, for which performance data came in a variety of formats and levels of fidelity. Therefore, a new TG algorithm design was implemented to meet these requirements: Behavior-Based Trajectory Generation (BBTG) [19]. Its flexible design allowed it to be adapted to three aircraft types (a turboprop, a business jet, and a commercial transport jet) [9]. Flight testing confirmed that BBTG prediction accuracy was comparable to predictions of airline flight planning systems and commercial FMS.

An advantage of BBTG's design approach is the ability to swap new aircraft behaviors into the TG framework, such as the vertical takeoff and landing procedures of an eVTOL aircraft. Over the past year, AOP was updated to use BBTG to model UAM vehicles, specifically two eVTOL variants: a quadrotor and a lift-plus cruise configuration. This flexible design that adapts readily to significantly different aircraft types provides the advantage of potentially certifying one generic TG implementation and then easily adapting it to multiple aircraft types with little to no recertification.

Experience with AOP indicates that point-mass modeling of the ownship provides a 4DT with adequate fidelity. A full six-degree-of-freedom model is generally not necessary to meet FPM objectives, though a TG will provide better fidelity if it accounts for the attitude of the aircraft in maneuvers such as banked turns. Hosting FPM automation onboard the vehicle facilitates the ability to use suitable and available (i.e., proprietary and/or local) data for point-mass modeling at a reasonable level of fidelity, and thereby to produce good-quality 4DTs.

One penalty for lower fidelity would be the need to incorporate greater restrictions on range, speed, altitude, and other performance metrics in order to provide a sufficient margin of error for the worst-case performance model. Another penalty occurs during deconfliction, as will be seen in Section VI.A. In our experience, however, attempting to provide an extremely high level of fidelity is neither necessary nor desirable due to several inherent sources of error in the input of the TG, such as wind forecasts, sensed position data, and the precision of performance data. Instead, computing the 4DT by numerical integration with a large step in time (or some other convenient variable of integration) has not been shown in our research to compromise the intended function of the TG (namely, creating feasible trajectories in support of other FPM objectives such as deconfliction). A TG that makes only the necessary computations can compute a 4DT and the energy consumption over the 4DT while using a relatively small amount of processing power, which may be important in low Size Weight and Power (SWAP) applications.

B. Energy Usage Prediction

A vital function of TG is to predict energy consumption over the predicted 4DT. When operating an eVTOL vehicle, for example, in order to ensure the vehicle has adequate energy reserves, the operator must estimate how much charge will remain in the on-board batteries if the vehicle completes its flight as planned. In addition, flights are often planned with a cost function that uses criteria such as a cost index to specify the desired tradeoff between time and energy usage to meet operational needs. Energy usage prediction is required in order to compare the total trip costs of alternative flight paths. Predicted energy use for a given 4DT is therefore an important factor in safety (ensuring adequate reserves) and optimization (maximizing reserves).

C. Identification and Prediction of Over-Constrained Trajectories

Because a strategic intent, by definition, is a specification of a desired flight, it may not be possible to satisfy all constraints of the strategic intent within the physical capabilities of the vehicle or within other limits on operations. For example, the vehicle may not be capable of meeting an assigned RTA, the flight path angle implied by two altitude constraints may be too steep for its segment of the flight, or the aircraft may not have sufficient energy reserves to carry out the complete strategic intent.

In cases such as these, the strategic intent is over-constrained. A function of TG, therefore, is to identify and report over-constraints and their consequences so they can be resolved. Moreover, even in the presence of over-constraints the TG must still predict a 4DT the aircraft will follow. In order to do this, the TG can "relax" constraints (in effect, predict which constraints vehicle guidance will violate and the degree to which it will violate them) until a feasible flight path is found.

D. Prediction of Tactical Trajectories

If there is no strategic intent but the aircraft is following shorter-term guidance (a tactical intent), it is useful (and possibly necessary for deconfliction purposes) for the TG function to use this tactical intent to produce a partial 4DT up to a configurable time horizon. Such a 4DT will stop short of the destination and will likely not refer to a destination at all. Tactical intent can include target states, such as turning to an intended heading or climbing to an intended altitude. In the case where no intent is available, a relatively simple function can compute a partial 4DT known as a state projection by assuming that the vehicle travels at a constant groundspeed and constant vertical rate without turning. Though it is impossible to decide the feasibility of a complete flight path using a 4DT predicted from tactical intent or state projections, such partial 4DT predictions still serve practical purposes such as detecting and resolving near-term conflicts. The resolution output is in the same tactical intent form (e.g., target state guidance to turn and capture a different heading or capture a different altitude). Because AOP was developed for commercial transport research, it included such TG modes, as well as hybrid modes such as strategic lateral intent (e.g., lateral path navigation) paired with tactical vertical intent (e.g., flight level change). Not surprisingly, accounting for the many possible hybrid modes complicated AOP's design.

Table 1 summarizes our initial design guidelines for FPM automation in creating a feasible flight path. The first column identifies the locations in the text the guidelines are found. The second column assigns an identification (ID) number to each guideline. The third column provides a short description identifying each guideline, though we recommend reviewing the extended discussion in the text of the paper. The fourth column provides reminders of some (but not necessarily all) of the benefits of following the guideline. The last two columns indicate the nature of the guideline: a mark occurs in the column labeled BP (best practices) on each row that describes a recommended best-practices guideline for implementing FPM automation, whereas a mark in the column labeled SD (standards development) means the recommendation on that row requires the development of standards.

VI. Deconflicting the Flight Path

Deconflicting a flight path is accomplished through trajectory prediction, conflict detection, and conflict resolution, taking prediction uncertainty into account.

In the interest of safety, airborne vehicles must maintain adequate separation from all *hazards*, which include other vehicles, special activity airspace, operationally restrictive weather, terrain, obstacles, and any other airspace region the vehicle's flight path should avoid. Attached to each such hazard is a *separation zone* that the vehicle must not enter (or should not enter, depending on the severity of the hazard). A *loss of separation* (LOS) occurs when the vehicle enters the separation zone of a hazard. A *conflict* is a predicted (future) LOS. *Deconfliction* is the process of *detecting* conflicts prior to LOS, *resolving* conflicts with a flight path change, and *preventing* new conflicts when changing flight paths for any reason. Over-constrained situations may require temporary trades, for instance, resolving the most urgent conflict first or shedding certain non-safety-related objectives until reaching a conflict-free state.

The following four subsections discuss important design considerations for functions supporting deconfliction.

A. Accounting for Trajectory Prediction Uncertainty in Strategic Conflict Detection

Virtually all FPM deconfliction functions rely on the ability to perform conflict detection (CD). Strategic CD operates on the ownship's complete 4DT, testing it for separation from hazards of all kinds.

Table 1 Summary of functional design recommendations for creating a feasible flight path.

Section number	ID	Design recommendation	Benefit	BP	SD
V.A	FE-1	Predict a complete 4DT to the destination for a strategic intent.	Supports strategic planning and decision-making for the other four FPM objectives.	•	
	FE-2	Employ a TG framework that is readily adaptable to different aircraft types.	Portability with reduced recertification requirements.	•	
	FE-3	Employ a TG framework that enables varying the modeling fidelity for different flight segments.	Computational efficiency and the ability to add future fidelity if needed.	•	
	FE-4	Use a point-mass aircraft model with suitable level of fidelity for TG.	Fewer restrictions on range, speed, or altitude.	•	
	FE-5	In TG, use step sizes small enough to provide an adequate reference path for the guidance system, but not smaller.	Accurate guidance; low computational cost.	•	
V.B	FE-6	Predict fuel and/or energy usage of every trajectory generated.	Managing energy reserves and optimizing flight path changes to preserve energy	•	
V.C	FE-7	Detect over-constraints in strategic intent and predict a feasible flight path with “relaxed” constraints.	Recovery from contingencies and more accurate prediction.	•	
V.D	FE-8	In the absence of strategic intent, use tactical or state-based prediction.	Detection of near-term conflicts.	•	

For any 4DT, however, there will be some deviations between the time sequence of vehicle states described by the predicted 4DT and the time sequence of actual vehicle states. We refer to these deviations as trajectory prediction uncertainty. Sources of trajectory prediction uncertainty include trajectory generation errors (such as wind prediction errors, modeling approximations, simplified vehicle performance data, and any lack of fidelity in the TG algorithm) and flight errors (such as sensor errors and guidance tracking errors). While a nominal predicted state (representing an approximation of the expected future position of the ownship) is adequate for some purposes, deviations of the actual future states from the predicted states can cause “surprise” LOS events or cause once-resolved conflicts to reappear if the CD algorithms do not account for the uncertainty. In the case of aircraft-to-aircraft separation, this effect is compounded by trajectory prediction uncertainties in the 4DTs of other aircraft in the airspace. As long as the ownship vehicle’s guidance system closes the loop in all four dimensions, some of the trajectory generation errors, especially those resulting from wind errors, can be mitigated (with an efficiency penalty) as long as the mitigation does not require exceeding the vehicle’s operating envelope.

The uncertainty associated with trajectory predictions needs to be considered in three dimensions: lateral, vertical, and along-path. The lateral or cross-track uncertainty is the most straightforward, represented by actual navigation performance (ANP) and any guidance-following error, typically called flight technical error (FTE). For operating environments such as UAM that may require extremely small values of required navigation performance (RNP), automated flight guidance systems must control the vehicle to achieve this performance. In these situations the uncertainty bound for cross-track prediction errors (left or right of the ground track) may be set based on RNP, for example, twice the RNP value (assuming the trajectory generation errors are kept relatively small in comparison).

The vertical path uncertainty is more critical than lateral path uncertainty because vertical separation standards are typically much smaller than lateral separation standards. Because the vertical and lateral paths are continuously updated in flight and ANP is monitored, two-times-vertical-RNP is a candidate for the trajectory prediction uncertainty bound, which would then be checked against the separation standard for deconfliction, at least during cruise. (The uncertainty in other flight modes depends on additional factors such as the type of guidance.)

Along-path prediction uncertainty will be negligible in the near field and grow along the trajectory unless mitigated by the guidance system. Significant along-path prediction errors are likely from two sources. Firstly, trajectory predictions must be made against estimated or forecast winds and temperatures, but actual atmospheric conditions during the flight may be quite different. The effect of wind and wind error on a vehicle's ground track is a function of heading and airspeed and thus does not affect all vehicles equally. Relative prediction uncertainty due to wind prediction errors for vehicles traveling along the same path may be small (assuming the predictions use a common wind forecast), whereas crossing and merging vehicles may show dissimilar time prediction errors at the crossing point. This effect is likely to be larger for guidance systems that control to airspeed (where the relevant measure is true airspeed) than for guidance systems that control to a 4DT. Secondly, auto-throttle control of speed is not precise and there will always be some speed error. Along-path uncertainty can be mitigated procedurally with conformance management in one of three ways: 1) require the vehicle to adjust airspeed to maintain time conformance plus or minus some maximum value (forcing a tradeoff between fuel/battery consumption and the precision of along-path guidance); 2) reserve a large airspace block along track (a very conservative approach, as much of this block will likely be unused); or 3) provide procedural timing constraints or pairwise speed control measures to keep vehicles separated.

Lateral, vertical, and along-path errors combine during flight to cause the actual vehicle path to deviate from the predicted path through 4D space with three degrees of freedom. Because actual vehicle path performance may result in unpredicted LOS between aircraft, penetration of special activity airspace, or other adverse events, FPM automation must be able to account for maximum expected deviation from the 4DT in each degree of freedom due to the combined effects of all sources of prediction uncertainty. In AAM applications where the traffic density may be elevated (e.g. UML-4), it is desirable to mitigate these errors in a more general manner within the CD algorithm. There are several ways this may be done. We omit the discussion of methods that are likely to be computationally intensive, such as Monte Carlo analysis, in favor of known methods with relatively small computational cost.

1. Separation buffers

One way to account for prediction error is to add a buffer to the separation requirement, increasing the effective horizontal separation standard equally in all lateral directions and increasing the vertical separation requirement equally above and below. This method, however, will likely place excessive restrictions on the flight paths of passing aircraft and lead to an overall reduction in allowable airspace capacity. For example, since along-path errors (in the direction of travel) will typically be larger than cross-track errors (perpendicular to the trajectory), if the horizontal separation buffer is sufficient to account for along-path errors, the buffer would result in excessive separation in some crossing or passing maneuvers. On the other hand, if the trajectory prediction uncertainty can be characterized in the dimensions of the trajectory (i.e., laterally, vertically, and along path relative to a predicted future position of the aircraft), then a CD function can be designed to account for these uncertainties, providing both greater protection and more efficient use of airspace.

In summary, separation buffers that are adequate in some directions at some points along the four-dimensional path are likely to be grossly excessive in other directions or at other points along the path. This is likely to cause good flight paths to be rejected, to interfere with equitable sharing of airspace, to force the assurance of CD to be compromised, or some combination of all three effects. Time-varying separation buffers can mitigate some but not all of these problems. For these reasons, we do not recommend separation buffers as the primary mitigation for trajectory prediction uncertainty.

2. Operational volumes

Different levels of uncertainty in along-track, cross-track, and vertical positions may be described by conformance concepts described as Operational Volumes and Conformance Volumes. In these concepts, each vehicle is required to operate inside at least one Operational Volume and only when the uncertainty can be limited by using Operational Volumes. Conformance Volumes are considered optional and if specified they are smaller than Operational Volumes and can be used to further bound predicted uncertainty.

In collaboration with Uber Technologies, Inc., NASA has investigated a way to control for trajectory prediction errors in UAM operations by using Transit Based Operational Volumes (TBOVs), originally developed for Unmanned

Aircraft System (UAS) Traffic Management [20]. A TBOV encloses a portion of a predicted trajectory within a fixed block of airspace that provides a buffer on each side of the predicted path as well as buffers above and below the predicted altitude in order to account for cross-track errors and altitude errors. TBOVs account for along-path errors by their extended length, essentially reserving its entire elongated volume for the one vehicle. By stringing a series of TBOVs along the trajectory (each TBOV assigned to a straight segment), FPM automation can use this sequence of TBOVs to cover an arbitrarily long trajectory, including turns. TBOVs are fixed with reference to the airspace, and they are essentially a trajectory-aligned airspace reservation method. A conflict occurs when FPM automation predicts two vehicles will occupy overlapping TBOVs at the same time. Each TBOV can overlap the next TBOV along the path in order to account for along-path uncertainty and a hard time constraint at each boundary. Using Operational Volumes such as TBOVs is potentially viable but can lead to inefficiently excessive blocking of along-path airspace. We recommend consideration of the following approach, which is not subject to this limitation.

3. Trajectory Prediction Uncertainty Bounds

Trajectory Prediction Uncertainty Bounds (TPUBs) [21] offer a highly customizable and efficient way to balance the need to cover prediction errors in all dimensions with the desire to minimize the amount of airspace "blocked off" by a vehicle's prediction at any time, as well as the need to represent the error bounds concisely. Like TBOVs, TPUBs surround nominal straight paths between selected points along a trajectory, but unlike TBOVs, TPUBs bound the maximum along-path error at each point along the trajectory as well as the maximum cross-track and altitude errors, all of which are defined relative to predicted future positions of the aircraft on the 4DT rather than fixed structures in the airspace. Moreover, the error bounds in any given dimension may be specified to grow or shrink between any two selected points in order to account for growth or reduction of estimated errors. The TPUBs surround a trajectory in a trajectory-oriented four-dimensional "tube" whose bounds in time and space interact as equal partners, rather than TBOV's fixed airspace orientation. It therefore more efficiently uses airspace, protecting only the true uncertainty in four dimensions and allowing tailoring as uncertainty varies by trajectory segment type. An illustration of TPUBs is shown in Fig. 7a.

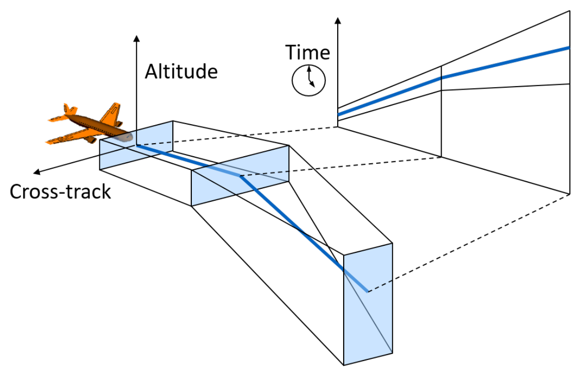

(a) Tailored, 4D specifications of prediction uncertainty oriented along each segment of an aircraft's predicted trajectory, enabling efficient use of airspace.

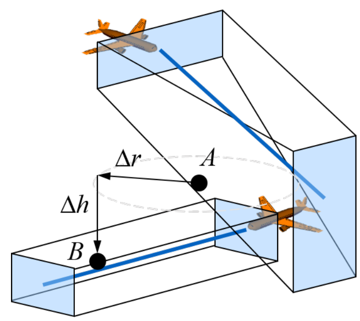

(b) Deconfliction of two aircraft, conservatively ensuring the outer edges of the 4D bounds (points A and B) meet the separation criteria ($\Delta h > h_{\min}$ or $\Delta r > r_{min}$) if the two aircraft might pass those points simultaneously ($\Delta t = 0$).

Fig. 7 Definition and usage of Trajectory Prediction Uncertainty Bounds (TPUBs).

Empirical work is required in order to size the TPUBs dimensions appropriately. If the maximum prediction error can be estimated closely, then TPUBs provide an extremely conservative result while using nearly the minimum airspace for that level of assurance. If a less conservative result is acceptable (as determined in the development of industry standards), the size of the TPUBs may be reduced in order to enable higher density operations. In any case, TPUBs will reserve less airspace at any given instant than an equally conservative algorithm that requires stationary blocks of airspace to be reserved for significant intervals of time.

B. Efficient Implementation of Strategic CD

AOP implemented an efficient CD algorithm with TPUBs [21] that determines whether an ownship "4D tube" is adequately separated from the 4D tube around any other vehicle's trajectory or from the separation zone of any other hazard (Fig. 7b). This TPUBs implementation replaced a previous algorithm that simply plotted each aircraft's position at synchronized 10-second intervals ("breadcrumbs") and measured the separation between the positions at each time. The TPUBs implementation used less processing time than the "breadcrumb" method while probing for conflicts at every instant during a trajectory (not just once every 10 seconds), accounting for different prediction errors in the along-path and cross-track directions, and accounting for growth in prediction errors over time or distance. This performance gain was attributed mainly to the ability of the TPUBs algorithm to cover a relatively long segment of a trajectory, either in level flight, climb, or descent, with a single constant-velocity segment (a linear path in both space and time). The CD algorithm probes such a segment end-to-end for conflict with a given hazard in a small bounded sequence of mathematical operations rather than requiring an iteration for each "breadcrumb."

A 4DT with TPUBs can be simplified by removing small non-linear effects, provided that the predicted deviation from a linear path in space and time, added to the estimated prediction error to be protected, would not result in actual aircraft states outside the TPUBs surrounding the simplified trajectory. Navigation over the (curved) surface of the Earth introduces non-linear relationships between the predicted path and an actual path, but these effects are negligible over small distances. In AOP the effect of the Earth's curvature is controlled by limiting the lengths of linear segments of the 4DT to a maximum of 40 nautical miles.

In cases where the velocity vector of the aircraft changes more rapidly, such as in a turn or vertical acceleration, the trajectory can be approximated by a few linear segments around which the dimensions of the TPUBs are temporarily increased in order to account for the predicted deviations from the linear path.

In summary, in order to perform CD rapidly and conservatively (for safety) while supporting stable operations in a densely populated airspace, we recommend using a mathematical algorithm on a 4DT described by linear segments of various lengths surrounded by TPUBs. This works best if community standards specify the TPUBs to apply to the 4DTs of traffic aircraft, especially if the operational intents of those aircraft include clues about the sizes of TPUBs (for example, an RNP level). If TBOVs are provided instead, they will provide some mitigation of prediction errors. Use of TBOVs rather than TPUBs may incur the cost of lower traffic density, lower confidence in CD, or both. Nevertheless, TBOVs are preferable to operational intents with no protection of operational volumes.

Further optimizations to CD can be gained by preprocessing and high-level filtering. Any aircraft whose reported position is too far from the ownship to close with the ownship within the time period of interest can be completely excluded from consideration in CD. Hazards with fixed locations (such as SUAs or fixed obstacles) can be indexed by a regular grid superimposed on the airspace; prior to CD, the grid cells that the ownship's 4DT passes through or near can be identified and any fixed hazards that do not overlap those cells can also be excluded from CD processing.

C. Strategic Conflict Resolution and Conflict-Free Flight Planning

Before activating the flight path before departure, or before executing a flight path change in flight, FPM automation should predict a 4DT and perform CD on this 4DT against all known hazards in the airspace. This can proactively accommodate even the most recent hazard updates., providing an opportunity to substitute an alternative flight path that may be conflict-free. AOP applied this technique by performing CD on the FMS "mod route" immediately before executing the flight path change and was shown to prevent new conflicts in a dynamic airspace environment.

Even during stable portions of the flight, FPM automation should repeat CD as knowledge of hazards, atmospheric conditions, or the ownship's own state data are updated. This suggests a cyclic approach to CD with an adaptable or small cycle period. For example, en-route commercial transport applications of AOP used a ten-second cycle.

When CD finds a conflict with the already-activated flight path due to updated information (or discovery of the conflict as it emerges within a look-ahead time horizon), FPM automation invokes conflict resolution (CR) in order to obtain a conflict-free flight path modification. If the ownship is following a complete 4DT and if the time until LOS is sufficient to allow a new complete 4DT to be generated, coordinated with other agents (such as a PSU), and activated, FPM automation should perform strategic CR, which provides a new, deconflicted flight path to the destination while preserving other mission objectives. Regular use of strategic CR helps retain predictability and promote stability in the airspace, resulting from the ownship sharing operational intent representing the complete 4DT.

Ideally, strategic CR will have information on the complete flight paths of all other aircraft that might interact with any new flight path of the ownship. That is, the operational intents of other aircraft will ideally represent strategic 4DT predictions of their flight paths rather than some less-detailed specification. Achieving this requires both for all aircraft to remain on their strategic 4DT as much as possible and for the sharing of strategic operational intent to be either called for in standards or incentivized through mechanisms such as elevated priority over those not sharing strategic operational intent.

In AOP, strategic CR is performed by the Pattern-Based Genetic Algorithm (PBGA) [22]. The PBGA algorithm in AOP was tested in multiple simulation experiments [12–18] and performed well from a research perspective. PBGA includes several deterministic components, but its core "genetic algorithm" prevents a typical deterministic performance analysis for certification. While PBGA is a possible model for other implementations of strategic CR, for certification purposes the algorithm will likely need to be partitioned from software modules that need to be certified to a higher standard. Alternatively, other algorithms for motion planning may be considered. However, because a primary source of the complexity of PBGA is the requirement that strategic CR provide "global" resolutions that avoid conflicts with all hazards, which can occur in many forms and may be arbitrarily distributed throughout the airspace, other strategic CR algorithms are likely to raise certification concerns as well. Focusing certification on CD performance (to confirm deconfliction of CR's ultimate solution, regardless of how it is produced) may be an appropriate approach. A deterministic backup CR algorithm may be included and invoked if the non-deterministic algorithm fails to produce an acceptable solution, but we are not aware of any such algorithms that themselves guarantee an acceptable solution. AOP's approach is to use non-strategic CR as a backup (see Section VI.D below on tactical CR and Section X on prioritizing safety.)

The performance of strategic CR is subject to the level of detail provided by operational intents of other aircraft. Operational intents that identify only takeoff times, landing times, and the corridors used may be difficult to deconflict when the corridors cross, but knowledge of the altitudes used can enable safe vertical separation. Additional knowledge of times and speeds along the path could better enable separation of vehicles even in cases where their paths cross at the same altitude. A 4DT can be derived from sparse route and time data if reasonable assumptions can be made about traffic aircraft performance. It is far preferable, however, if community standards call for sufficient operational intent to be shared, ideally the full 4DT but at least the route plus altitudes and speeds or time along the route, supporting accurate estimation of a 4DT without incurring too much uncertainty. Preferably, shared operational intent also includes some specification of prediction uncertainty containment (such as RNP levels) to increase stability of deconfliction.

If the operational intent of the ownship is known to FPM automation but useful operational intents of other vehicles are not, a form of strategic CR can be performed in which the resolution provides a strategic maneuver for the ownship that avoids conflict with the target state or state projections of other vehicles. In this case, occasional short-notice conflicts should be expected. AOP's design includes the use of the highest level of intent available from traffic aircraft, with automatic switching to target state or state projections if trajectory-level intent is not provided.

D. Tactical Conflict Resolution

If there is not sufficient time before LOS to generate and implement a new strategic flight path, some form of tactical CR should be performed. (This case also will occur if the ownship has left its strategically planned flight path due to previous use of tactical CR to avoid either traffic or airspace hazards.) One method of tactical CR is to search for a conflict-free partial trajectory based on limited intent data such as a target state; in AOP this method was called tactical intent-based CR. When that method is unavailable or unsuccessful, CR must then rely on state projections and state-based algorithms alone. Routine reliance on tactical CR rather than strategic CR, however, may tend to result in an increased number and severity of avoidance maneuvers, while also leaving aircraft more frequently in states of poor 4DT predictability.

State-based CD still has a useful role even when all operational intents are well known: it provides a degree of blunder protection by pointing out where conflicts could occur if a vehicle does not execute the turn, altitude change, or level-off that it is approaching. In AOP, state-based CD and CR were used to "override" intent-based CD and CR for close-in encounters.

If the ownship must maneuver tactically, FPM automation should attempt to resume strategic flying as soon as practicable. In AOP, the strategic reroute capability performed this function by planning a new deconflicted flight path from the ownship's present state to the destination. AOP used the PBGA algorithm with patterns designed specifically for strategic reroute. This function performed best when the velocity vector of the ownship achieved a steady state after tactical maneuvering.

Table 2 summarizes our initial design guidelines for FPM automation for deconflicting the flight path.

Table 2 Summary of functional design recommendations for deconflicting the flight path.

Section number	ID	Design recommendation	Benefit	BP	SD
VI.A.3	DE-1	Use aircraft-position-oriented, trajectory-segment-tailored uncertainty bounds for CD.	Protects actual uncertainty and supports efficient and flexible airspace use in dense operations.	•	
VI.B	DE-2	Simplify control points of flight path within conformance bounds in operational intent.	Faster processing, less data to exchange.	•	
	DE-3	Preprocess and filter hazards.	Faster processing.	•	
VI.C	DE-4	Perform CD immediately before departure and prior to any flight path change, even if generated by CR.	Greater conflict prevention in dynamic airspace conditions.	•	
	DE-5	Perform CD cyclically during flight, at least as often as new data are received.	Continued deconfliction.	•	
	DE-6	Deconflict using strategic operational intent (complete 4DT) of ownship and traffic aircraft when possible, or the highest available level of intent.	Greater predictability and stability in airspace operations	•	
	DE-7	Require or incentivize through standards an increased priority for aircraft that share increased specifications of intent	Greater predictability and stability in airspace operations		•
	DE-8	Perform strategic CR to produce a new flight path when time to conflict permits.	Preserves deconflicted strategic flight.	•	
	DE-9	Provide global resolutions in strategic CR, resolving not just the triggering conflict but creating no other conflicts in the process.	Reduced invocations of CR and greater stability in airspace operations.	•	
	DE-10	Partition non-deterministic strategic CR algorithms from safety-critical systems and include deterministic backup CR algorithms.	Reduced certification burden while gaining the advantages of advanced non-deterministic algorithms.	•	
	DE-11	Include complete 4DT in shared operational intent, or if not possible, altitudes and speeds or times along the entire path.	Ability to deconflict a complete flight path.		•
	DE-12	Include predicted conformance levels along entire path in operational intent.	More stable deconfliction of the flight path.		•
	DE-13	In CR, use available traffic data in the absence of strategic flight paths.	Deconfliction regardless of the amount of intent data available on traffic aircraft.	•	
VI.D	DE-14	Perform tactical CR when strategic CR is not possible.	Persistent and adaptable deconfliction.	•	
	DE-15	Use state-based CD to point out (and even proactively resolve) consequences of a missed turn.	Blunder anticipation and protection.	•	
	DE-16	Provide a strategic reroute capability during tactical flight.	Safely return to strategic flight after tactical maneuver	•	

VII. Coordinating the Flight Path

When several vehicles are using the same airspace, some mechanism is required to coordinate the flight paths of these vehicles with each other in order to avoid excessive replanning and maneuvering, ensure equity among operators, and arbitrate capacity-constrained resources. FPM automation can be designed to provide this coordination through a variety of mechanisms, such as sharing intent, avoiding conflict creation while changing flight paths, respecting right-of-way rules, and resolving tactical conflicts with maneuvers in complementary directions. Industry standards for flight path coordination will likely be needed to ensure design compatibility across different FPM implementations.

Sharing intent is the most fundamental form of coordination, akin to a turn signal in automobiles. It communicates the intended future changes in the trajectory and contributes directly to creating a dynamic common operating picture amongst users sharing the airspace. Without knowing intent of nearby aircraft, FPM automation must compare its own 4D intent (which an ownship has access to by definition) with state vectors broadcast by traffic, a dissimilar comparison to the one that would be performed by the traffic aircraft's FPM automation.

Previous AOP research identified the benefits of intent data sharing (e.g., reduced conflict rate, increased warning time, and reduced need for tactical maneuvering), while also showing safe operations are achievable without it [18]. Intent may be specified in varying manners and degrees, from target state (i.e., a target heading in a turn, or a target altitude to be captured in a climb or descent) to the next one or several trajectory change points (TCPs) to full 4D intent to the destination. AOP research in operations at the scale of transport category aircraft found diminishing returns beyond two TCPs. Similar studies should be performed on the costs and benefits of intent sharing in UAM-scale operations.

Conflict prevention functionality in FPM automation provides coordination by enabling the operator to maneuver while maintaining conflict-free operational intent, reducing the need for subsequent deconfliction by other operators or external services. FPM automation should therefore incorporate conflict prevention measures in all flight path changes, regardless of cause, in order to ensure that such maneuvers are coordinated by only using "unclaimed" airspace as indicated by published operational intents. A corollary is that the more flight path is included in shared operational intent, the greater claim an operator would have to that airspace.

A suitable set of right-of-way rules implicitly coordinates the flight paths of vehicles by reducing the number of actors required to resolve conflicts. As in maritime operations, a vehicle having priority may "stand on" by maintaining a steady course, whereas the burdened aircraft performs the avoidance maneuver. In AOP, this concept was applied by staggering the first alerting to the conflict based on right-of-way determined by the conflict geometry. In cases of conflicts to be resolved by strategic CR, the relative positions of the vehicles at the time the conflict is detected may be quite different from the relative positions predicted at the time of LOS; in this case, right-of-way rules can be applied according to the predicted future states of the aircraft just prior to LOS. If the conflict is detected early enough, the vehicle with right of way under those rules can delay strategic CR, allowing the burdened vehicle to resolve the conflict unilaterally. In simulations with AOP, this practice has been found to reduce the average number of changes to flight paths in flight and to reduce the chance of interference between changes to different vehicles' flight paths.

For conflicts that remain unresolved or "pop up" within an established time horizon, simultaneous maneuvering by both aircraft to ensure safety is warranted. In AOP's implementation of right-of-way, conflict alerting was staggered but not eliminated for the priority aircraft. After all, the priority aircraft may not have assurance that the burdened aircraft has detected the conflict and will maneuver. Though this could be solved with explicit communications between the two vehicles, AOP's design avoided the need for this explicit coordination by implementing implicit coordination through the design of the tactical resolution algorithms [23]. This approach uses a shared set of algorithmic rules to direct tactical maneuvering by both aircraft in complementary directions so that the maneuvers mutually add to separation rather than subtract or cancel their effects. We recommend further investigation and development of an industry standard for some form of implicit coordination to increase the integrity of separation assurance while reducing the expense and complexity of peer-to-peer communications.

Table 3 summarizes our initial design guidelines for FPM automation in coordinating the flight path.

VIII. Preserving Flight Path Flexibility

Creating a flexible flight path involves placement of the flight path to maximize ability to accommodate future disturbances, thereby increasing mission completion probability. A flight path that preserves flexibility thereby preserves options for modification, serving as a proactive mitigation against entering excessively complex situations that risk mission success (i.e., not arriving at the intended destination and roughly on time). This requirement for flexibility derives from the need to ensure that operations based on distributed control, in which each aircraft operator (rather than a central authority) holds responsibility for separation and can meet that responsibility reliably without the risk of

Table 3 Summary of functional design recommendations for coordinating the flight path.

Section number	ID	Design recommendation	Benefit	BP	SD
VII	CO-1	Produce 4DT operational intent (at least two TCPs) for sharing with other operators directly or through a service provider.	Increased predictability and stability of flight paths.	•	
	CO-2	Incorporate conflict prevention measures in all flight path changes.	Increased stability of airspace operations.	•	
	CO-3	Stagger CR action times based on right-of-way (ROW) rules.	Reduced requirement to change operational intent.		•
	CO-4	Use rules that implicitly coordinate the directions of resolution maneuvers when two aircraft are in conflict.	Maneuvers that increase the separation of vehicles rather than canceling each other.		•

compromising mission success. Meeting the separation responsibility requires maintaining enough degrees of freedom to alter the flight path, should a new disturbance arise that requires a flight path change (such as a conflict resolution or a path stretch to absorb a schedule delay). The more degrees of freedom available along the future flight path, the less likely the mission objectives would need to be shed in the process of addressing the disturbance. In other words, under highly constrained (inflexible) conditions, resolving a conflict, for example, may require abandoning a scheduled arrival slot, since safety takes precedence over schedule adherence. The outcome would be a safety success but a mission failure. While these situations may occasionally be tolerated, they should not be the norm.

Prior NASA-sponsored research postulated "trajectory flexibility preservation" as a potential FPM mitigation to such scenarios, serving in effect as a distributed complexity management tool for operator-responsible trajectory management concepts [24, 25]. According to this research, trajectory flexibility preservation would be implemented as an FPM capability that results in the production of flight paths that maximize two characteristics, robustness and adaptability, to enhance the likelihood that the flight path can tolerate future disturbances. Robustness is defined as the ability of the aircraft to keep its planned flight path unchanged in response to the occurrence of a disturbance. A flight path that can withstand a disturbance without having to change is more robust than other flight paths that become infeasible when the disturbance occurs. Adaptability is defined as the ability of the aircraft to change its planned flight path in response to the occurrence of a disturbance that renders the current planned flight path infeasible. A flight path that positions the aircraft so that other feasible flight paths remain accessible to it if a disturbance occurs and renders the current flight path infeasible is more adaptable than another flight path for which the disturbance leaves fewer or no feasible flight paths.

A flexibility preservation function was implemented in AOP on an exploratory basis, demonstrating the feasibility of integrating trajectory flexibility preservation into an FPM automation system. An initial implementation applied flexibility planning only beyond a planning horizon specified for strategic CR [24]; later a flexibility formula was integrated into the fitness function of AOP's strategic CR algorithm, PBGA.

While the research proposed quantitative metrics for robustness and flexibility, and modeling demonstrated the technical feasibility of this approach, the capability is not sufficiently mature for the development of specific FPM design guidelines that could be offered here. Further exploration of the trajectory flexibility preservation capability will be needed to determine its viability. In the absence of such a capability, it will be necessary to provide other mitigations against excessively constrained operations in which mission objectives must be frequently abandoned to ensure flight safety, for instance, limiting the quantity of aircraft permitted in a given airspace or providing built-in degrees of freedom (e.g., an extra "lane" in corridor operations).

Table 4 summarizes our initial design guidelines for FPM automation in preserving flight path flexibility.

IX. Optimizing the Flight Path

The final capability set for FPM is the functionality to create optimal flight paths. An optimal flight path is one that achieves a desired optimization objective within the available degrees of freedom defined by the other applicable FPM objectives for flight path generation (feasible, deconflicted, etc.). For instance, the requirements to generate flight paths

Table 4 Summary of functional design recommendations for preserving flight path flexibility.

Section number	ID	Design recommendation	Benefit	BP	SD
VIII	FL-1	Include flexibility as a criterion for acceptability of a flight path.	Increase likelihood of accomplishing mission goals in complex airspace environments.	•	

that are both feasible and deconflicted may still yield a substantial range of viable solutions. While unlikely to be a required functionality (at least from a regulatory perspective), FPM algorithms that can find an optimal solution among a range of viable solutions will yield added benefit at little cost, creating a potential market discriminator among FPM automation systems.

The optimization objective would typically be set by the operator, either by general policy or specifically determined for each flight. It could be a vehicle specific objective, such as minimizing the energy used in completing the mission, or minimizing exposure to light/moderate turbulence for passenger comfort. Alternatively, it could be a fleet-related business objective that applies different objectives across multiple vehicles, for example, to ensure greater overall throughput or to ensure that higher priority flights have the best possible schedule conformance. Optimization can also consider community impact, for example, to minimize the transmission of noise over sensitive areas.

The search for the most optimal flight path may explore any of the degrees of freedom available to the flight. In corridor operations, where lateral maneuvering may be procedurally restricted, the search may compare the efficiency of two or more corridors (if the choice exists), and it may analyze different altitudes and speeds within a single corridor. In more flexible airspace, a variation in lateral path may also be explored by itself or in combination with altitude and speed. The more degrees of freedom available, the more sophisticated the algorithm may need to be to achieve a reasonably comprehensive search for an optimal solution.

As a design guideline, optimization should generally be integrated with one or more of the other FPM capability sets, the minimum being the feasible flight path generation capability. Such integration can take advantage of solution space within the aircraft's performance envelope, such as selecting an optimal descent profile or an optimal balance of speed and energy conservation (e.g., cost index). Integration with the deconflicted capability enables two functional modes: optimizing as a primary goal, and optimizing as a secondary goal. The difference is the trigger event for flight path modification. The first mode's trigger may be a cyclic check for more optimal routes; by integrating deconfliction capabilities into this optimization search, one ensures that only conflict-free airspace is considered in the search for a more optimal flight path. In effect, it also integrates with the coordinated capability by ensuring that flight-optimizing path changes occur without creating new conflicts. The second mode's trigger is the detection of a needed flight path change, such as a conflict requiring resolution or a path stretch to absorb delay. By integrating optimization into the deconfliction capability set, one ensures that the "best" of all viable flight path solutions is found, whether the criterion is minimum energy use, optimized arrival time, total flight cost (considering both energy use and flight time), or turbulence exposure.

The reference FPM automation system, AOP, includes prototype functionality for both modes. The "strategic intent-based conflict resolution" (SICR) function employs a pattern-based genetic algorithm (PBGA) that explores a defined envelope of permissible lateral and vertical flight path changes to resolve the conflict [22]. As is typical for genetic algorithms, a population of candidate solutions—each of which AOP ensures is feasible—compete using a "fitness function," and the competition ensues over multiple generations of ranking, mating, and mutation, until an ultimate winning solution emerges. The fitness function is mathematically specified to apply penalties of varying magnitudes to first weed out conflicted solutions (highest penalty) but also to favor those that best achieve the user-defined optimization objective, e.g., minimum energy use (smaller penalties). By implementing multiple fitness functions, AOP allows the operator to switch optimization objectives in midflight, a useful feature for adapting to changing mission objectives as flight conditions change.

Optimization was more directly implemented in a derivative application of AOP called "Traffic Aware Planner" (TAP) [26] in support of NASA's "Traffic Aware Strategic Aircrew Requests" (TASAR) concept [9]. The capability was later added back into AOP. Using the same PBGA algorithm as used in SICR, TAP triggers a cyclic search for more optimal flight paths, again using the approach of multiple fitness functions allowing the user to switch objectives midflight. Because TASAR is intended for use under IFR wherein ATC ensures flight paths are deconflicted, it leverages traffic state data (received through an onboard receiver of Automatic Dependent Surveillance Broadcast data) and other data sources to increase the likelihood that a requested flight path change (for optimization purposes) is free of conflict. This integration of optimization and deconfliction increases the likelihood of ATC approval of the requested change. TASAR was operationally validated in airline revenue service [27] and is undergoing commercialization.

Table 5 summarizes our initial design guidelines for FPM automation for optimizing the flight path.

Table 5 Summary of functional design recommendations for optimizing the flight path.

Section number	ID	Design recommendation	Benefit	BP	SD
IX	OP-1	Include trajectory optimization criteria in TG.	Meet business objectives better.	•	
	OP-2	Include trajectory optimization in strategic CR.	Meet business objectives better.	•	
	OP-3	Periodically attempt trajectory optimization with deconfliction.	Better meet objectives when environment changes favorably.	•	

X. Prioritizing Safety in Flight Path Management

FPM automation nominally seeks to achieve multiple objectives simultaneously, by establishing and maintaining a flight path that is feasible, deconflicted, coordinated, flexible, and optimal. In unexceptional situations, meeting all five objectives should normally be possible, though some tradeoffs can be expected. For instance, the most flexible flight path may not also be the most optimal flight path, if the former is measured in degrees of freedom and the latter measured in economic efficiency. The appropriate balance could be determined by the operator's risk tolerance (i.e., an insufficiently flexible flight path that has greater cost savings may risk the fulfillment of all mission objectives if later disturbances impact the flight path). However, what must be assured, and what cannot be traded off, is that safety be given ultimate priority in FPM.

The types of processes that promote a safe outcome are illustrated in Fig. 8. The diagram shows the five objectives in a radial depiction, where the outer regions depict the portion of the flight farthest in time from the nearest hazard, even before detecting a conflict. In this regime, conflict detection is periodically performed and intent data is shared, while flight path changes focus on feasibility, preserving flexibility, and performing optimization. Once a conflict has been detected within a specified time horizon (which may vary depending on the type of hazard), the focus of FPM increases in deconfliction and coordination, while reducing priority on flexibility and optimization. Under nominal conditions, the conflict is resolved soon after detection, using 4D intent as a detection basis and producing a feasible full 4D trajectory as the resolution output. The resolution is coordinated by creating no other conflicts and respecting right of way in the timing of the action.

However, we must allow that some confluence of events in complex operations may result in situations in which a conflict with a hazard (e.g., a traffic aircraft) remains unresolved and eventually becomes urgent. This urgency increases as the remaining time until LOS decreases, corresponding to a shrinking distance to the center of Fig. 8. The principle of prioritizing safety has two important implications: some objectives may need to yield to others based on safety assurance, and individual objectives may need to adjust their approach as a hazard grows nearer to promote a safe outcome. In other words, the functional capabilities designed to achieve each FPM objective must be prepared to adapt to the increasing hazard in order to drive to an ultimately safe outcome.

First, because some objectives become less important, specifically the objectives to produce flexible and optimal flight paths, they may need to have decreasing influence. The principal goal of the flexibility objective is to provide adequate degrees of freedom to resolve future conflicts, an objective that can be maintained in initial deconfliction but loses relevance as the time-to-hazard decreases. A similar argument of decreasing relevance exists for the optimization objective.

Second, individual objectives may need to adapt their approach as an unresolved hazard grows nearer. Generally, this involves simplifying the problem in various ways. For instance, the objective to create a feasible flight path change is simplified by relaxing non-safety-critical constraints, such as procedural airspace boundaries (e.g., corridors), or

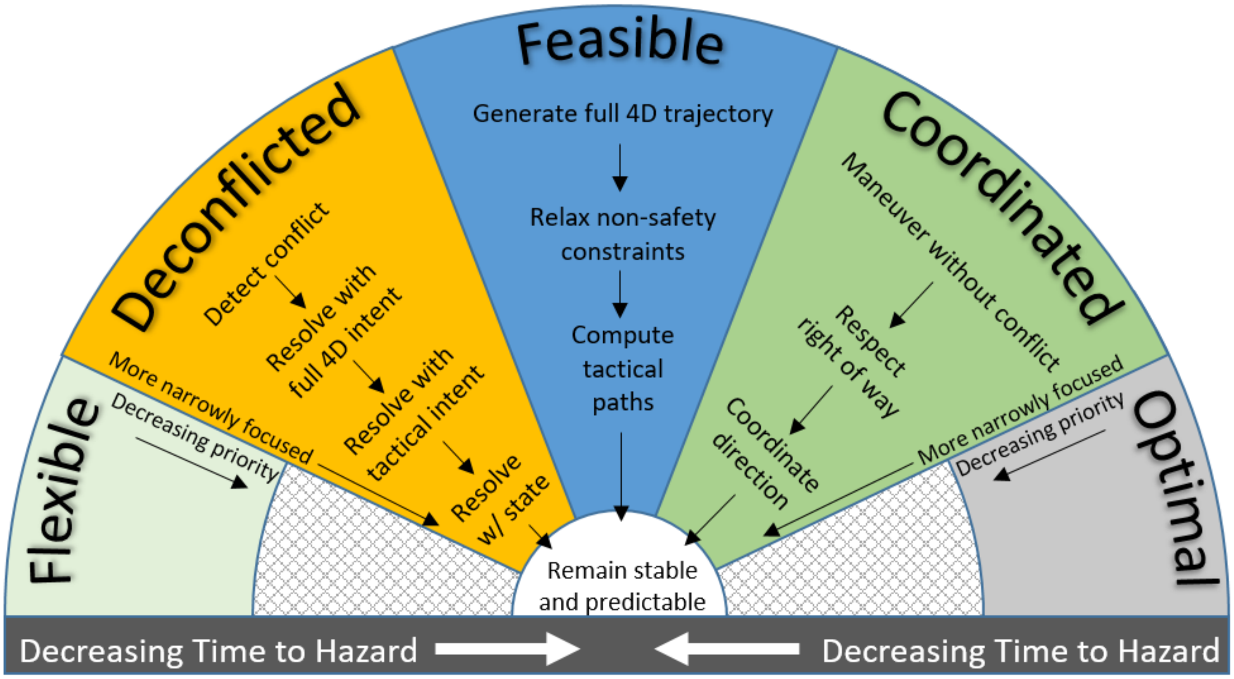

Fig. 8 FPM automation supports multiple objectives but should be designed to prioritize safety as unresolved hazards grow close. Some objectives are shed, whereas others take on greater focus, as the time to hazard decreases (moving with decreasing radius toward the center of the graphic).

non-urgent constraints, such as an RTA, thus opening up additional degrees of freedom for deconfliction. It can be further simplified by switching to the generation of tactical (open loop) maneuvers. Similarly, the deconfliction approach can adapt from strategic intent-based resolutions (producing a conflict-free 4DT solution) to tactical intent-based resolutions (producing conflict-free but open-loop maneuvers, i.e., tactical CR) to state-based (focusing tactical CR to atomic maneuvers, e.g. turning or climbing, that address only the most urgent conflict). Coordination's increasing focus begins with maneuvering without creating new conflicts and respecting right-of-way rules, but with an unresolved conflict would switch to an "all hands on deck" mode where both vehicles are expected to resolve the conflict with implicitly coordinated maneuver directions (achieved through common use of state-based CR algorithms designed for this purpose). In the very rare case where no maneuver by the ownship will resolve the urgent conflict, the best option may be to assume a stable and predictable path so that the other aircraft may take the required action or other safety-dedicated algorithms like DAA take over to ensure the aircraft pass "well clear."

NASA's research with AOP in very high-density scenarios of transport category aircraft with random traffic geometries (i.e., no pre-emptive deconfliction) has shown this layered approach to prioritizing safety to be highly effective and scalable, even under conditions of wind uncertainty and limited surveillance performance [14–18]. It was rare that even intermediate layers of capability were invoked to prevent losses of separation. It is expected that similar performance will be seen in UAM-scale operations at UML-4.

XI. Future Work

A. Planned Research Activities

Several efforts are planned to continue the work of FPM automation R&D for UAM applications. A formal FPM automation concept of operation will expand the current definition of flight path management to include (at a minimum) corridor operations, PSU interactivity, and contingency management. Our reference FPM automation system, AOP, will be expanded to study these factors and will be evaluated in a series of planned research activities, including simulations and flight tests, in UAM-relevant scenarios that support transition to UML-4. Pertinent information gained from these activities will enable further design guidelines to be developed and provided to industry and will inform FPM standards development.

Multiple factors in FPM automation functionality that could impact airspace system performance are being considered for analysis in the planned FPM research efforts. Although existing functionality of the reference system, AOP, was designed to be hosted onboard the aircraft, application of FPM functionality is not limited to the flight deck. Future efforts will examine multiple system architecture configurations in which FPM automation functions are distributed across airborne and ground-based systems, enabling evaluation of performance impacts on FPM objectives and UAM airspace operations. The studies will also quantify constraining factors influencing communication requirements. Communication needs in traditional commercial-transport operations have been met by relatively higher bandwidth between ground-based systems, lower latencies between airborne systems, and minimal reliance on air-ground data links. At higher UML with significant increases in traffic density, the same communication capabilities may not be assumed and could have significant implications for FPM functionality not hosted onboard. Performance of in-flight deconfliction services impacted by potentially saturated, range-limited, or unreliable data exchanges between air-ground systems and between vehicles will be evaluated.

B. Recommendations for Future Research

While the current work of FPM automation R&D focuses on UAM, the FPM functionality is applicable to operational contexts beyond UAM operations, including commercial transportation, General Aviation, and UAS. It is recommended that FPM be broadly studied across these domains.

Separation standards developed for commercial transport operations are likely inappropriate for universal application and should be reexamined for the UAM and other operational contexts. Emerging automation capabilities may enable reduced or tailored separation standards that lead to greater operational flexibility. Such capabilities should be evaluated in studies that examine the parameters used to set deconfliction thresholds for separation assurance, such as strategic deconfliction, DAA, and conflict prevention. Airspace performance using these reduced separation standards is quantifiable and has multiple performance indicators and potential benefits, including increased airspace capacity, increased flight efficiency, increased vehicle energy reserves, decreased reliance on DAA systems, and decreased reliance on excessive maneuvering for safe mission completion.

The feasibility of operations that leverage the full capabilities of emerging airborne FPM automation to minimize reliance on ground-based service providers (yet still benefit from their services) should be investigated. Such systems could support airborne FPM while also providing pilots and fleet operators with a maintainable common operating picture, which is of particular importance for AAM operations in regions not fully or reliably supported by PSUs or similar service providers. Furthermore, contingencies for inaccessible ground-based services (e.g., PSU, SDSPs) under off-nominal communications should be examined in studies that include models for automation functions that compute alternative flight paths for diversion contingencies. Such studies should examine the impacts of critical communication loss on vehicle energy reserves and mission objectives. Benefits of such studies would include increased operational flexibility, reduced dependency on external services for flight path changes, decreased ground-operator workload for nominal and off-nominal operations, and increased trust in automation.

XII. Conclusions

The role of FPM in AAM is to achieve mission objectives by creating, monitoring, and revising the intended flight path of an aircraft in the presence of evolving mission objectives, constraints, and external stimuli. FPM is distinct from mission management (which sets the objectives and monitors mission feasibility), tactical operations (which temporarily intervenes in urgent hazard scenarios), and vehicle control (which executes the intended flight path). Vehicle-hosted FPM is consistent with emerging concepts of operation (for example, UAM) which hold the operator responsible for flight path safety including in-flight deconfliction. FPM automation provides five sets of capabilities, enabling it to generate flight paths that are feasible, deconflicted, coordinated, flexible, and optimal. Each capability set plays a critical or important operational role and therefore needs either to be incorporated into FPM automation or provided by an external service.

This paper presents design guidelines for FPM capability sets. The FPM automation functional design guidelines are based on knowledge and experience gained through 15 cumulative years of development, simulation and testing of a reference FPM automation system, AOP, originally developed for autonomy research in commercial transport operations and currently being adapted to the UAM application. The design guidelines for each of the five FPM capability sets are intended to support industry stakeholders in developing and pursuing best practices and may assist in development of

FPM standards. Based on our experience with AOP, in order to create feasible flight paths, the trajectory generation function should be easily adaptable to adding future fidelity where needed and supporting adaptation to different aircraft types. Flight path deconfliction algorithms should account for trajectory prediction uncertainty without sacrificing airspace capacity, using an approach similar to AOP's TPUBs. Multiple mechanisms of implicit coordination that collectively reduce or eliminate the need for explicit real-time coordination and negotiation should be implemented within FPM automation. Emerging concepts that enable an operator to ensure adequate degrees of freedom always remain for producing feasible, deconflicted flight paths, such as trajectory flexibility preservation, will become important at higher UML and warrant further exploration. Finally, while the other FPM capabilities are being developed, matured and certified, optimization functions that offer the opportunity to create market-discriminating FPM functionality can be fielded today, given that this capability is unlikely to be regulated. These represent the most salient guidelines for the FPM capability sets discussed in this paper. Future development, simulation/flight testing, and analysis of UAM functionality supporting corridor operations and contingency management will result in future design guidelines to be shared with the FPM community.

References

[1] National Academies of Sciences, Engineering, and Medicine, *Advancing aerial mobility: A national blueprint*, National Academies Press, Washington DC, 2020.

[2] Wing, D. J., and Levitt, I. M., "New Flight Rules to Enable the Era of Aerial Mobility in the National Airspace System," NASA TM–20205008308, November 2020.

[3] Federal Aviation Administration, "Urban Air Mobility (UAM) Concept of Operations v1.0," , June 2020. URL `https://assets.evtol.com/wp-content/uploads/2020/07/UAM_ConOps_v1.0.pdf`.

[4] Hill, B. P., DeCarme, D., Metcalfe, M., Griffin, C., Wiggins, S., Metts, C., Bastedo, B., Patterson, M. D., and Mendonca, N. L., "UAM Vision Concept of Operations (ConOps) UAM Maturity Level (UML) 4," , December 2020. URL `https://ntrs.nasa.gov/citations/20205011091`.

[5] "General Operating and Flight Rules," *Title 14 Code of Federal Regulations, Federal Aviation Regulations, Part 91*, U.S. Government, 2021. URL `https://ecfr.federalregister.gov/current/title-14/chapter-I/subchapter-F/part-91/subpart-A/section-91.3`.

[6] Cotton, W. B., and Wing, D. J., "Airborne trajectory management for urban air mobility," *AIAA Paper 2018–3674*, June 2018. https://doi.org/10.2514/6.2018-3674.

[7] Goodrich, K. H., and Theodore, C. R., "Description of the NASA Urban Air Mobility Maturity Level (UML) Scale," *AIAA Paper 2021–1627*, January 2021. https://doi.org/10.2514/6.2021-1627.

[8] Karr, D. A., Vivona, R. A., Roscoe, D. A., Depascale, S. M., and Wing, D. J., "Autonomous operations planner: A flexible platform for research in flight-deck support for airborne self-separation," *AIAA Paper 2012–5417*, September 2012. https://doi.org/10.2514/6.2012-5417.

[9] Wing, D. J., "The TASAR Project: Launching Aviation on an Optimized Route Toward Aircraft Autonomy," NASA TP–2019-220432, December 2019.

[10] Wing, D. J., Chancey, E. T., Politowicz, M. S., and Ballin, M. G., "Achieving resilient in-flight performance for advanced air mobility through simplified vehicle operations," *AIAA Paper 2020–2915*, June 2020. https://doi.org/10.2514/6.2020-2915.

[11] 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.

[12] Wing, D. J., Prevot, T., Murdoch, J. L., Cabrall, C. D., Homola, J. R., Martin, L. H., Mercer, J. S., Hoadley, S. T., Wilson, S. R., Hubbs, C. E., Chamberlain, J. P., Chartrand, R. C., Consiglio, M. C., and Palmer, M. T., "Comparison of ground-based and airborne function allocation concepts for NextGen using human-in-the-loop simulations," *AIAA-2010-9293*, September 2010. https://doi.org/10.2514/6.2010-9293.

[13] Wing, D., Prevot, T., Lewis, T., Martin, L., Johnson, S., Cabrall, C., Commo, S., Homola, J., Sheth-Chandra, M., Mercer, J., and Morey, S., "Pilot and controller evaluations of separation function allocation in air traffic management," *Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)*, Chicago, IL, 2013.

[14] Consiglio, M., Hoadley, S., Wing, D., and Baxley, B., “Safety Performance of Airborne Separation: Preliminary Baseline Testing,” *AIAA Paper 2007–7739*, September 2007. https://doi.org/10.2514/6.2007-7739.

[15] Consiglio, M., Hoadley, S., Wing, D., Baxley, B., and Allen, D., “Impact of Pilot Delay and Non-Responsiveness on the Safety Performance of Airborne Separation,” *AIAA Paper 2008–8882*, September 2008. https://doi.org/10.2514/6.2008-8882.

[16] Consiglio, M. C., Hoadley, S. T., and Allen, B. D., “Estimation of separation buffers for wind-prediction error in an airborne separation assistance system,” *Eighth USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*, Napa, CA, 2009.

[17] Khan, Z., Idris, H., Consiglio, M., and Wing, D., “Surveillance range and interference impacts on self-separation performance,” *2011 IEEE/AIAA 30th Digital Avionics Systems Conference*, IEEE, Seattle, WA, 2011, pp. 4D5–1–4D5–16.

[18] Lewis, T. A., Phojanamongkolkij, N., and Wing, D. J., “The effects of limited intent information availability on self-separation in mixed operations,” *12th Integrated Communications, Navigation and Surveillance Conference (ICNS)*, IEEE, Herndon, VA, 2012, pp. B6–1–B6–14.

[19] Karr, D. A., Vivona, R. A., Woods, S., and Wing, D. J., “Point-Mass Aircraft Trajectory Prediction Using a Hierarchical, Highly-Adaptable Software Design,” *AIAA Paper 2017–3336*, June 2017. https://doi.org/10.2514/6.2017-3336.

[20] Verma, S. A., Monheim, S. C., Moolchandani, K. A., Pradeep, P., Cheng, A. W., Thipphavong, D. P., Dulchinos, V. L., Arneson, H., Lauderdale, T. A., Bosson, C. S., Mueller, E. R., and Wei, B., “Lessons learned: using UTM paradigm for urban air mobility operations,” *2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC)*, San Antonio, TX, 2020, pp. 1–10.

[21] Karr, D. A., and Vivona, R. A., “Conflict Detection Using Variable Four-Dimensional Uncertainty Bounds to Control Missed Alerts,” *AIAA Paper 2006–6057*, August 2006. https://doi.org/10.2514/6.2006-6057.

[22] Vivona, R. A., Karr, D. A., and Roscoe, D. A., “Pattern-Based Genetic Algorithm for Airborne Conflict Resolution,” *AIAA Paper 2006–6060*, August 2006. https://doi.org/10.2514/6.2006-6060.

[23] Narkawicz, A. J., and Muñoz, C. A., “State-Based Implicit Coordination and Applications,” NASA TP–2011-2117067, March 2011.

[24] Idris, H., Vivona, R. A., and Al-Wakil, T., “Trajectory-Oriented Approach to Managing Traffic Complexity: Trajectory Flexibility Metrics and Algorithms and Preliminary Complexity Impact Assessment,” NASA CR-2009-215933, September 2009.

[25] Idris, H. R., Shen, N., and Wing, D. J., “Complexity management using metrics for trajectory flexibility preservation and constraint minimization,” September 2011. https://doi.org/10.2514/6.2011-6809.

[26] Woods, S., Vivona, R. A., Henderson, J., Wing, D. J., and Burke, K. A., “Traffic Aware Planner for Cockpit-based Trajectory Optimization,” *AIAA Paper 2016–4067*, June 2016. https://doi.org/10.2514/6.2016-4067.

[27] Henderson, J., Wing, D. J., and Ballard, K. M., “Alaska Airlines TASAR Operational Evaluation: Achieved Benefits,” NASA TM-2019-220400, September 2019.