NASA/TM–20230018665

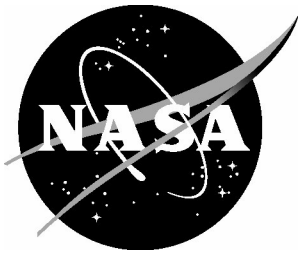

Testing of Advanced Capabilities to Enable In-time Safety Management and Assurance for Future Flight Operations

Andrew J. Moore, Steven Young, George Altamirano, Ersin Ancel, Kaveh Darafsheh, Evan Dill, John Foster, Cuong Quach, Anne Mackenzie, Justin Matt, Truong Nguyen, Patricia Revolinsky, Kyle Smalling, and Sixto Vazquez
NASA Langley Research Center, Hampton, Virginia

Lynne Martin, Jolene Feldman, Vimmy Gujral, Lilly Spirkovska, Becky Hooey, Kevin Bradner, David Iverson, Shawn Wolfe, and Nikunj Oza
NASA Ames Research Center, Mountain View, California

Ryan Condotta, Christopher Morris, and J. Sloan Glover
Analytical Mechanics Associates, Inc., Hampton, Virginia

Nicholas Rymer and Andrew Turner
ViGYAN, Inc., Hampton, Virginia

Charles Walter
ASRC Federal Data Solutions, Moffett Field, California

Portia Banerjee, Matteo Corbetta, and Chetan Kulkarni
KBR, Inc., Moffett Field, California

Jane Cleland-Huang
University of Notre Dame, South Bend, Indiana

James Jones and Tim Bonin
MIT Lincoln Laboratory, Lexington, Massachusetts

April 2024

NASA STI Program Report Series

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

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

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

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

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

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

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

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

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

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

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

- Help desk contact information:

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

NASA/TM–20230018665

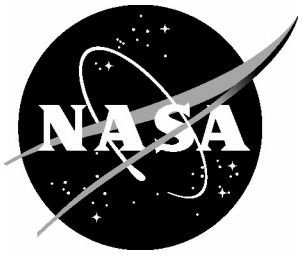

Testing of Advanced Capabilities to Enable In-time Safety Management and Assurance for Future Flight Operations

Andrew J. Moore, Steven Young, George Altamirano, Ersin Ancel, Kaveh Darafsheh, Evan Dill, John Foster, Cuong Quach, Anne Mackenzie, Justin Matt, Truong Nguyen, Patricia Revolinsky, Kyle Smalling, and Sixto Vazquez
NASA Langley Research Center, Hampton, Virginia

Lynne Martin, Jolene Feldman, Vimmy Gujral, Lilly Spirkovska, Becky Hooey, Kevin Bradner, David Iverson, Shawn Wolfe, and Nikunj Oza
NASA Ames Research Center, Mountain View, California

Ryan Condotta, Christopher Morris, and J. Sloan Glover
Analytical Mechanics Associates, Inc., Hampton, Virginia

Nicholas Rymer and Andrew Turner
ViGYAN, Inc., Hampton, Virginia

Charles Walter
ASRC Federal Data Solutions, Moffett Field, California

Portia Banerjee, Matteo Corbetta, and Chetan Kulkarni
KBR, Inc., Moffett Field, California

Jane Cleland-Huang
University of Notre Dame, South Bend, Indiana

James Jones and Tim Bonin
MIT Lincoln Laboratory, Lexington, Massachusetts

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

April 2024

Acknowledgments

The Research and Development described herein was conducted over the course of the System Wide Safety project, part of NASA's Aviation Operations and Safety Program, and would not have been possible without the leadership, support, and sponsorship of NASA's Aeronautics Research Mission Directorate. The authors are grateful to Associate Administrator Robert (Bob) Pearce, Program Director Akbar Sultan, and Project Managers John Koelling, Misty Davies, and Kyle Ellis for their vision and constancy. They, along with their teams, established a strategic vision that recognizes the difficulty and importance of safety assurance in aviation as it enters an era that requires new safety approaches that can adapt to the evolution of aircraft and the expansion of flight into new operational environments.

In addition to the authors, key contributors to the R&D and testing were Scott Dorsey, Bailey Ethridge, Ken Eure, Alwyn Goodloe, Julian Gutierrez, Rafia Haq, Ed Hogge, Tom Johnson, Natasha Neogi, Mike Scherner, Laura Smith, Chris Thames, John Ossenfort, Jason Watkins, and Tyler Willhite. Thank you all for your dedication, creativity, and can-do spirit.

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA STI Program / Mail Stop 050
NASA Langley Research Center
Hampton, VA 23681-2199

Abstract

In order to refine an initial Concept of Operations, explore Concepts of Use, and expose/validate requirements for future In-Time Aviation Safety Management Systems (IASMS), testing architectures were created, along with a set of capabilities and underlying information exchange protocols. These systems were conceived and developed based on hazards associated with two envisioned urban area flight domains: (1) highly autonomous small uncrewed aerial systems (sUAS) operating at low altitudes, and (2) highly autonomous air taxis. The initial scope of this development is described in [1]; this report provides an update, focusing on the subsequent developments and test activities.

As stated in [1], it is important to note that there are many capabilities already in use by the industry (or soon to be in use) that will play critical roles in future IASMS designs. Those reported here were developed to address a gap in the current state-of-the-art regarding specific hazards/risks, and/or to allow for investigation of the interplay between and across hazard types — particularly regarding how overall safety risk can be reduced or managed effectively. Results of testing and development activities are organized by the operational phase wherein a particular capability would be employed (i.e., preflight, in-flight, and post-flight/off-line).

Pre-flight: A set of capabilities were developed to help mitigate safety risk prior to flight (e.g., during flight and mission planning). Results of testing summarize (1) validation activities to raise the Technology Readiness Level (TRL) and (2) evaluation activities where the capabilities were applied to flight/mission planning procedures and used by operators/pilots. For the latter, flight plans were automatically assessed, and operators/pilots were notified of hazardous flight segments so as to enable adjustment of the flight plan and re-evaluation, and/or to better inform go/no-go decisions. Capabilities addressed hazards associated with power consumption, third-party risk, wind, navigation system performance, radiofrequency interference, and proximity to geo-spatial threats (e.g., buildings, trees, and no-fly zones).

In-flight: Flight experiments tested capabilities that detect and respond to hazards encountered during flight. In the first series, safety hazards were monitored and assessed onboard, and system-generated mitigation maneuvers were recorded (but not acted upon by the vehicle). In the second series, mitigation maneuver commands directed the aircraft in response to safety hazards (i.e., auto-mitigation). The sUAS used for testing is described in full, as is the test architecture, which included commercial avionics, research avionics, and onboard software designed to detect, assess, and respond to hazards. The onboard system was designed as a run-time assurance framework, consistent with [2] and supportive of both supervisory and automated modes. The primary functions included: real-time risk assessment (RTRA), auto-pilot monitoring, constraint monitoring, and contingency select/triggering. RTRA performs integrated risk assessment considering data from several hazard-related monitors (e.g., battery, motors, navigation, communications, population density, and loss-of-control).

Post-flight/off-line: Data monitored and recorded during flights can enable IASMS capabilities that execute after flights have completed (or “off-line”). These include: (1) the ability to identify anomalies and trends that may only be observable when comparing data spanning a number of similar flights; (2) the ability to update and validate pre-flight and in-flight capabilities and any underlying models to improve their performance; (3) the ability to report anomalies/off-nominals that may indicate design changes or maintenance actions are needed; and (4) the ability for humans involved in operations to report safety-relevant observations to help in understanding the flight data and/or the operational context of a flight. Progress on three such capabilities is summarized; the first investigates anomaly detection given a limited set of flight logs and applies an approach previously used for space operations. The second explores what could be identified using a larger set of flight logs, including from web-based forums where flight logs are posted by sUAS autopilot users. The third creates a new means of collecting information on UAS incidents and accidents via the Aviation Safety Reporting System (ASRS).

Contents

1 Introduction

1.1 Research Goals and Scope

The second century of flight will include a broadening of vehicle types, mission types, operational environments, and automation applications. In some cases, traffic densities will be very high. More flights will occur at lower altitudes, particularly in and around densely populated areas. To assure safety in this era of expansion, new approaches to safety assurance will need to be added to those already in place for traditional aviation (e.g., today's commercial airline operations).

With this challenge in mind, NASA's Aeronautics Research Mission Directorate (ARMD) and the National Academies have suggested a set of research and development goals and a concept to proactively mitigate safety risks for future aviation operations [3][4]. For simplicity, in this document we use the term defined in [4], In-Time Aviation Safety Management Systems (IASMS). As shown in Fig. 1, an IASMS framework is "data driven" and includes (a) continuous monitoring of a set of safety metrics, (b) integrated predictive assessment to identify elevated risk states, and (c) safety assurance actions which may include supervised as well as automated risk mitigation strategies. This is described in [4] as:

> "The concept of real-time system-wide safety assurance should be approached in terms of an in-time aviation safety management system (IASMS) that continuously monitors the national airspace system, assesses the data that it has collected, and then either recommends or initiates safety assurance actions as necessary. Some elements of such a system would function in real time or close to real time, while other elements would search for risks by examining trends over a time frame of hours, days, or even longer."

In order to refine an initial Concept of Operations (ConOps) ([1][5][6]), explore various Concepts of Use (ConUses), and expose/validate requirements for future IASMS(es), testing architectures were created. Then, a set of capabilities and an underlying information system was conceived and developed based on hazards associated with two domains: (1) highly autonomous small uncrewed aerial systems (sUAS) operating at low altitudes within and over urban areas and (2) highly autonomous air taxis operating within and over urban areas. The initial scope of this development is described in [1]. This report provides an update, focusing on the subsequent developments and test activities.

As stated in [1], it is important to note that there are many capabilities already in use by the industry (or soon to be in use) that will play critical roles in future IASMS designs. The ones reported here were developed to address a gap in the current state-of-the-art regarding specific hazards/risks and/or to allow for investigation of the interplay between and across hazard types — particularly regarding how overall safety risk can be reduced or managed effectively. Results of testing and development activities are organized by the operational phase wherein a particular capability would be employed (i.e., preflight, in-flight, and post-flight/off-line).

1.2 Terminology, Conceptual Framework, and Research Overview

Initial research on IASMS has led to a terminology and conceptual framework that helps to describe various traits, abstractions, and implementations. As described in [3] and shown in Fig. 1, three high-level "functions" (Monitor, Assess, Mitigate) are envisioned that would be tailored

to a specific IASMS application domain (i.e., a four-tuple consisting of the vehicle type/equipage, mission type, intended operational environment, and safety risk tolerance). The new terminology discussed below provides a way to represent lower levels of detail for specific domains. We consider IASMS designs as consisting of a tailored set of Services, Functions, and Capabilities (SFCs). These terms provide the flexibility when articulating a design and are defined in [7] and [8] as follows:

Capability – The ability to perform or achieve certain actions or outcomes.

- May require service(s), function(s), an underlying system, and/or human involvement
- Includes procedures, training, and an interface for any required human involvement[1]
- Provides safety-relevant benefits

Service – Information generated at a remote site[2] and available for use during an operation.

- Presumes the information is needed or useful to the operation
- Requires connection to a remote server and service provider
- Provides information via 'request-reply', 'publish-subscribe', and/or 'broadcast'
- Produces safety-relevant information

Function – One or more actions or means of translating a set of inputs to a desired set of outputs.

- Resides in all elements of the system (e.g., servers, ground stations, and onboard)
- Produces safety-relevant outputs (metrics) or data needed to compute safety-relevant metrics

The next key IASMS concept is operational. Different capabilities may be employed to help reduce, or better manage, safety risk during the three phases of operations: pre-flight, in-flight, and post-flight/off-line (see Fig. 1). Examples include: the ability to check forecast winds during pre-flight planning; the ability to automatically divert to an alternate landing location during flight; and the ability to detect unsafe trends when looking across several similar flights during post-flight (or off-line). In general:

- Pre-flight capabilities support risk assessment and mitigation during flight planning (e.g., when selecting route to fly and time window);
- In-flight capabilities support monitoring, assessment, and mitigation during flight based on the current context and predicted trajectory; and
- Post-flight/off-line capabilities support identifying precursors, anomalies, and trends when comparing across a set of similar flights.

In terms of human role(s), procedures and training must be designed as well (e.g., how to complete a pre-flight checklist that requires use of a particular service or function). This procedure design and training, while critical to effective implementation of IASMS designs, are out of scope of the research described in this document.

---

[1] Services and functions may also require human involvement; but for simplicity, this involvement is captured here within the overarching capability that employs the service(s) and/or function(s).

[2] Remote in this context means not embedded in onboard systems, on a local ground station, or in a pilot remote control (RC) device. As technology advances, computing for some services may be portable to local computers/avionics, in which case they would become functions (in this context).

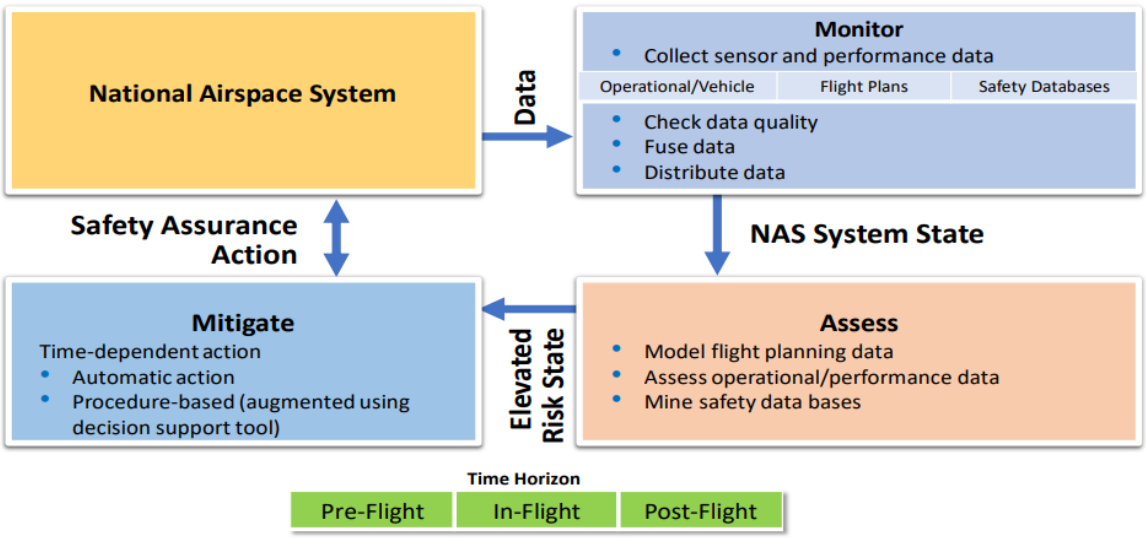

Fig. 1. Data gathered and fused (Monitor) are evaluated in the context of targeted safety metrics for the vehicle and flight domain (Assess) to enable actions intended to lessen safety risk (Mitigate). Action outcomes flow back, with the loop checking that they have the intended affect and do not introduce new safety risks. Mitigations may be introduced as design updates as well as during operations. From [6].

To illustrate how these concepts and terminology are used, the most abstract IASMS capabilities are explained below.

- IASMS Monitoring Capability – The ability to acquire, integrate, and assure the quality of data that may come from one or more sources regarding a safety-relevant metric or metrics. Information is provided in accordance with data quality requirements; recognizing that some data may require protections to de-identify the source, assure integrity, and defend against corruption by unauthorized or unauthenticated sources. Consumers and producers of the information may be human or automated agents. When there is a human role, appropriate procedures and training should be defined and validated as part of the capability. A monitoring capability requires embedded function(s) and an underlying architecture for managing the data/information produced by the capability.

- IASMS Assessment Capability – The ability to provide timely detection, diagnosis, and predictions regarding changes in safety risk and/or hazard states. Where necessary, assessment should span hazard types to judge how the overall safety margin is changing based on the current operational context, recent cascading event sequences, and longer-term trends that can become evident with access to historical data maintained by monitoring capabilities. An assessment capability requires embedded function(s), connection to one or more monitoring capabilities, and an underlying architecture for managing the data/information produced by the capability.

- IASMS Mitigation Capability – The ability to plan and execute timely responses to hazardous events or event sequences when/if safety margins are observed or are predicted to deteriorate below acceptable levels. Mitigation capabilities can span many types of responses (e.g., manual control inputs, multi-agent coordinated actions, supervisory interventions, and fully automated actions). A mitigation capability may require embedded function(s), connection to one or more monitoring or assessment capabilities, and an underlying architecture for enabling execution of planned response(s).

To proceed further in exploring the IASMS concept, specific application domain(s) must be selected. For flight operations, an IASMS domain is defined by the four-tuple of (a) vehicle type/equipage, (b) mission type, (c) intended operational environment, and (d) safety risk tolerance. Domain selection is necessary in order to determine what to monitor, what to assess, and what to mitigate; and determine how to enable these capabilities and how well they should perform. Two domains were selected for the research described in this document:

(Domain 1) Highly autonomous sUAS operating at low altitudes within and over urban areas

a. Vehicle type/equipage: sUAS (< 25 kg, or 55 lb) [9], auto-flight capable
b. Mission type: package delivery and reconnaissance (< 3 mi); lift-off to landing
c. Operational environment: low altitude in and above urban areas (< 120 m, or 400 ft)[3]; light-to-moderate winds; no rain or storms; an air traffic management (ATM) system such as the UAS traffic management (UTM) system [6]; auto-flight enabled lift-off through landing; remote operator providing supervisory oversight with ability to intervene if necessary; operator may not always be within line-of-sight
d. Safety risk tolerance: Moderate. Equivalent or better than what is expected by FAA for Part 135 certificates such as currently held by sUAS package delivery companies (e.g., UPS, Google, Zipline). FAA Part 108 is in draft but also is expected to imply an acceptable risk tolerance for sUAS for Beyond Visual Line-of-Sight operations.

(Domain 2) Highly autonomous occupied air taxi operating within and over urban areas

a. Vehicle type/equipage: 2-4 passenger electric vertical takeoff and landing (eVTOL) aircraft, auto-flight capable
b. Mission type: short range air taxi (< 20 mi); pad-to-pad (lift-off to landing)
c. Operational environment: Urban Air Mobility Maturity Level 3/4 (UML-3/4) [10]; low altitude in and above urban areas, i.e., up to 1220 m (4000 ft) AGL; light-to-moderate winds; no rain or storms; extensible traffic management (xTM)-like air traffic management (ATM) system; auto-flight enabled lift-off through landing; onboard pilot or remote operator providing supervisory oversight with ability to intervene if necessary; remote operator may not always be within line-of-sight.
d. Safety risk tolerance: Low. Equivalent or better than what is expected by FAA for Part 135 for limited operations (e.g., low density; low number of vertiports); however, at scale, likely direct correspondence to Part 121 (wherein Safety Management Systems (SMS) are required). More stringent Part 121/Part 25 type operational/certification requirements may result from safety/hazard assessments (e.g., the CFR 21.17b tailoring process and application of Special Operations Risk Assessment (SORA) / Safety Assurance and Integrity Levels (SAIL) methods).

[3] All altitudes are stated as above ground level (AGL) unless otherwise noted explicitly as above mean sea level (MSL).

These domains and their associated hazards determine a diverse set of SFCs that can help to expose or validate IASMS requirements and considerations. To cover this span, four research, development, test, and evaluation (RDT&E) activities were initiated by four teams:

1. NASA internal team
2. University of Notre Dame and Iowa State University team
3. Virginia Commonwealth University (VCU) and National Institute of Standards and Technology (NIST) team
4. George Washington University (GWU), Massachusetts Institute of Technology (Lincoln Laboratory) (MIT/LL), Vanderbilt University, and University of Texas (Austin) team

Each team defined an architecture and a different set of SFCs that addressed complementary aspects of the overarching IASMS concept. The remainder of this report primarily summarizes recent testing conducted by the NASA internal team. However, a short summary of work by the other teams is provided below along with references to their reports.

Notre Dame-led team (Dronolytics). The Notre Dame-led team focused on the sUAS domain; developing novel onboard and off-board data analytic capabilities to detect, diagnose, and potentially mitigate emergent anomalies observed in critical data streams or complex operational contexts. For example, two deep-learning techniques were developed as onboard functions to detect problems related to vibration, excessive attitude fluctuations, and compass interference. The team also explored methods to search/analyze online forum discussions and flight logs to identify common causal and contributing factors to incidents, and to combine run-time monitoring with off-line postmortem monitoring. Additional details of some of the team's work on the latter is given in Section 4.2. See also [10][11][12][13][14][15][16][17][18].

VCU-led team (Pervasive Monitoring). The VCU-led team focused on what they dubbed a "pervasive monitoring capability." This included a strong emphasis on cybersecurity-related issues, such as developing security solutions to protect from data errors caused by faults or cyberattacks that could lead to system malfunction or unexpected unsafe behavior. One example is monitoring of embedded system functions within a flight control system of a UAS. Developers of these systems are dependent on a diverse supply chain for sensors and processors and, in addition, they cannot always ensure the trusted delivery of verified firmware updates to the end user. These factors increase vulnerability to various types of cyberattacks. The team developed a cybersecurity monitor for such control systems called the HECAD (Hierarchical Embedded Cyber-Attack Detection). HECAD is designed to monitor the system for the effects of cyber security attacks and system faults at all levels of the implementation, from the physical component interconnect level to the information level. The team used two testbeds for evaluations; wherein cyberattacks can be simulated and inserted into the control system. As the name implies, several other types of monitors were also part of the pervasive monitoring capability, including monitoring for hazards in the operational environment. For more information on this project see [19].

GWU-led team (Learning-based IASMS). The GWU-led team focused on application of IASMS concepts to both sUAS and eVTOL-like air taxis (such as are envisioned for future Advanced Air Mobility (AAM) and Urban Air Mobility (UAM) operations). Three high-level functions were investigated: (1) trajectory risk prediction and re-routing; (2) health management and vehicle-adaptive control; and (3) scalable automated conflict detection and resolution. Learning-based techniques are applied across these functions to detect, assess, and mitigate three types of hazards while also investigating a unique run-time assurance technique based on a shielding concept. For example, a path planning and risk assessment algorithm based on machine learning was developed and simulated using adverse winds, aircraft battery/motor faults and performance degradation, and airspace corridor incursions by non-cooperative aircraft. A high-resolution low altitude wind modeling and forecast service was developed and

integrated along with a means of translating forecasts into probabilistic representations of wind hazard areas. The Dallas-Fort Worth metropolitan area was used to test (via simulation) AAM/UAM scenarios and SFCs. Flight testing of sUAS SFCs is planned for later in 2024 to occur in the Boston area. For more information on this project see [20][21][22].

NASA internal team. As described in [1], the NASA internal team began by establishing a scalable UTM-like architecture to support the R&D. Then, a diverse set of SFCs were selected and implemented. These were chosen for three reasons: (1) the SFC addressed one or more hazards identified for at least one of the selected domains; (2) the SFC addressed a specific gap or vulnerability in state-of-the-practice for safety assurance and/or (3) the selected set covered a diversity of SFC building blocks such that integration issues across a broad IASMS ecosystem could be investigated (e.g., compute platforms, communications requirements, data package size/protocol, a mix of onboard/offboard function and function allocations).

The implemented SFCs are given in Tables 1-4 along with a short descriptor, and in some cases, acronyms that are used later in the document (e.g., Fig. 2). The nomenclature '2x' or '3x' indicates cases where multiple approaches were implemented/tested. The 'Usage phase' column in Tables 1-3 indicate the flight phase in which an SFC is used: 'Pre' indicates preflight usage, 'In' indicates in-flight usage, and 'Post/Offline' indicates post-flight usage. Further details and reference publications are provided in [1] and Appendix A. Selected results from evaluation testing are given in Sections 2, 3, and 4.

Table 1. IASMS Service Implementations Tested by the NASA Internal Team

Service Name	Short name, note	Usage phase
3rd party casualty risk assessment [23][24]	NPCRA; GRASP	Pre, In
Navigation quality assessment	NavQ, 2x	
Corridor Assessment of Positioning System [25][26][27]	CAPS	Pre
Geometric Assessment of Positioning System [25][28]	GAPS	Pre
Proximity to threat assessment [24][29][30]	PtT	Pre, In
Obstacle collision risk assessment	ROC	Pre, In
RF environment and interference monitor [8][31][32][33][34]	RF/RFE/RFI	Pre
Weather/wind monitoring and forecasting	3x	
Wind modeling [35][36]	WM	Pre
LaRC networked weather stations [36]	WxSt	Pre, In
Wind information prediction service[4] [37][38]	WIPS	Pre
Airspace dynamic density assessment [39]	DD	Pre, In
Battery prognostics [40] and health management	BP/BHM	Pre, In

[4] Developed and tested by partners at MIT/LL.

Environment observation station network (continuous measures of radio frequency (RF) spectrum use; weather conditions; and GPS parameters)	EOSN	Pre, In
Flight performance assessment service [41]	FPAS	Pre, In
Aircraft and flight system state telemetry (service provider is onboard) [42]	TM	In

Table 2. IASMS Function Implementations Tested by the NASA Internal Team

Function Name	Short name, note	Usage phase	On/off board
Ground control station (GCS) operator display and interaction support	3x		Off
Supplemental Data Service Provider-Consolidation Dashboard [43][44]	SDSP-CD	Pre	Off
Human Autonomy Teaming Interface System [44][45]	HATIS	Pre, In	Off
Operation Planning Tool [46]	OPT	Pre, In	Off
Auto-pilot monitor	APMon	In	On
Battery health monitor [47][48]		In	On
Constraint monitor (e.g., geofences, flight path deviation, airspeed, altitude) [49][50]	Safeguard	In	On
Contingency select and trigger [7][8]	CST	In	On
In-flight anomaly detector [51][52]	MMS, IMS	In	On/off
Link monitor		In	On
Proximity to threat monitor (e.g., buildings, trees, no-fly zones, high risk areas)	PtT	In	On
Real-time risk assessment [23][24]	RTRA	In	On, Off
Off-line anomaly detector [51][52]	MMS, IMS	In	Off
Traffic detect-and-avoid; merging-and-spacing; re-routing[5] [53][54]	ICAROUS; DAIDALUS	In	On

Table 3. Assessment Capabilities Enabled by the Tested Services and Functions

Name	Usage phase
Assess airspace density safety margin or risk	Pre, In
Assess weather/wind safety margin or risk	Pre, In
Assess battery and propulsion system health	Pre, In

[5] Tested prior to 2020; these functions not used during 2021-2023 testing.

Assess in-flight integrated safety risk	Pre, In
Assess navigation performance risk (i.e., assess loss of navigation system risk)	Pre, In
Assess 3rd party casualty risk	Pre, In
Assess obstacle collision risk Assess proximity to threats (e.g., obstacles, high-risk areas, no-fly zones)	Pre, In
Assess RF environment and interference risk (i.e., assess loss of link risk)	Pre
Assess traffic collision risk	Pre, In
Identify precursors, anomalies, and trends	Pre, In, Post/Off-line
Assess flight plan risks across all monitored hazards	Pre, In

Table 4. Mitigation Capabilities Enabled by Developmental Services and Functions

Adjust airspace usage and/or airspace constraints
Select and execute contingency (maneuver) (onboard, automated)
Select and execute contingency (flight plan) (onboard, automated)
Select and execute contingency (maneuver) (air/ground, supervised)
Select and execute contingency (flight plan) (air/ground, supervised)
Switch to backup system/support/procedure (e.g., switch to alternative navigation source or alternative telemetry frequency)
Change design (system, procedure, training) (feed back data/findings to designers)

The IASMS architecture shown in Fig. 2 was used to support testing of the services and functions listed in Tables 1-2; early versions of the architecture are described in [1][7]. Two elements of the system shown in Fig. 2 were not part of the testing described in this report:

- USS/PSU – UAS Traffic Management Service Supplier (USS) [55] and Provider of Services for UAM (PSU) [56]. These emerging capabilities provide air traffic management services.
- FIMS – Flight Information Management System (FIMS) [55] that provides access to FAA-provided information services within UTM ecosystems.

The service providers for the test architecture were at four locations:

- Service provider 1: hosted at NASA Ames Research Center (ARC)
- Service provider 2: hosted at NASA Langley Research Center (LaRC)
- Service provider 3: hosted at NASA Johnson Space Center (JSC) using APPDAT[6] [57]

[6] Application Platform, Packaged Deployment and Analytics Technologies (APPDAT) is a platform hosted at NASA Johnson Space Center (JSC) for developing cloud-native applications and making services discoverable and accessible by both NASA and non-NASA users.

- Service provider 4: hosted at Mitre (external partner)

Four services listed in Table 1 and Appendix A were not evaluated as part of the flight tests detailed in Section 3 and were instead evaluated independently as standalone services: ROC, WIPS, BHM, and FPAS.

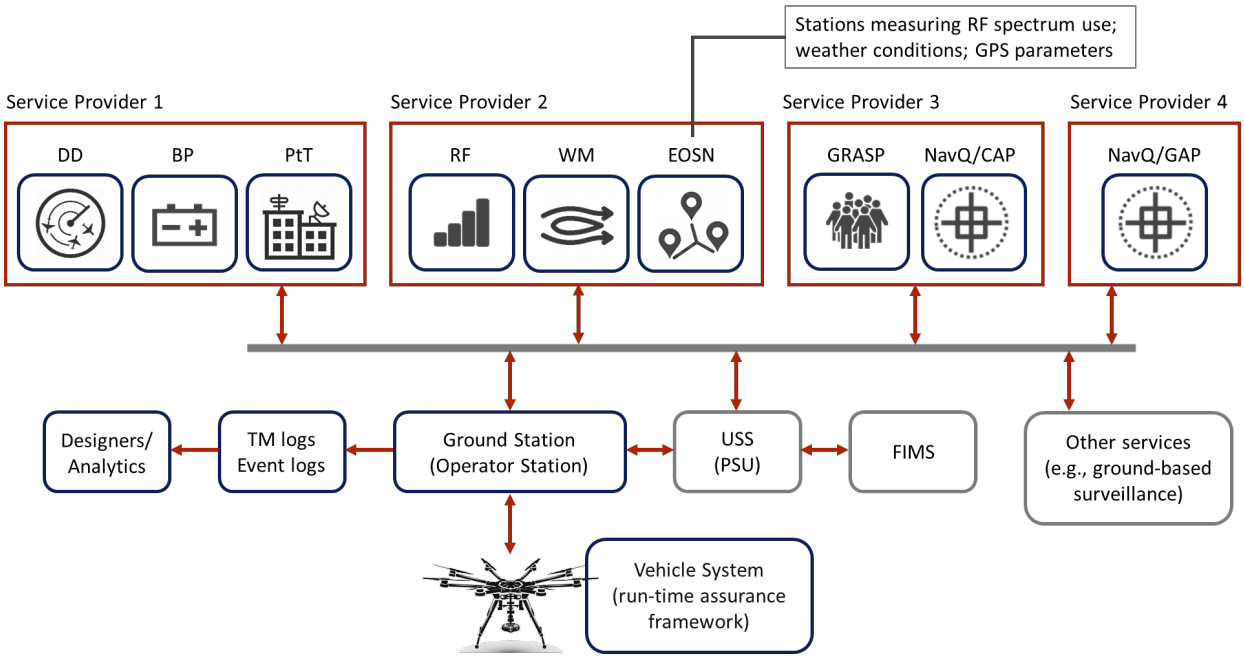

Fig. 2. Testing architecture used by the NASA internal team. IASMS service and service providers (top icons), utilizing data streams from observation stations (top), operator ground station (middle), UTM service providers (middle right), and the vehicle (bottom) continually monitor flight conditions, assess safety risks, and provide for supervised and automated mitigations. Flight data and outputs from services and functions are archived (middle left) for post-flight analysis and design improvement.

1.3 Report Organization

The remainder of this report summarizes results of NASA testing by operational phase: pre-flight (Section 2), in-flight (Section 3), and post-flight/off-line (Section 4). Testing by partner teams is described in separate reports. Concluding remarks are provided in Section 5.

Brief descriptions of nine services are provided in Appendix A. The experimental aircraft, ground equipment, and flight checklists used in test flights are detailed in Appendix B. Additional flight test results are summarized in Appendix C. A list of acronyms is provided in Appendix D.

2 Testing Pre-flight SFC Implementations

This section is organized into two subsections. Section 2.1 summarizes standalone testing of some of the pre-flight services and functions listed in Tables 1 and 2 with the remainder of results covered by several previous publications (as listed in the Tables). Section 2.2 summarizes testing of some of the pre-flight capabilities listed in Tables 3 and 4 where sets of services and functions were applied and utilized as part of flight and mission planning.

2.1 Standalone Testing

As mentioned, results of most standalone testing are available in the referenced publications provided in Tables 1 and 2, and in Appendix A. Below are selected additional test results for some services and functions.

2.1.1 RF Monitoring

RF interference has a known role in compromising safety of flight operations. Hazardous effects include coupling within and across on-board equipment and degradation of air-ground communication links. These hazards can have safety-critical manifestations in highly automated UAS and UAM flights where they can lead to (1) unintended autopilot mode switching on the vehicle, (2) loss of operator situational awareness when intervention may be most needed, and/or (3) the inability to re-command a vehicle entering restricted air space (e.g., due to an out-of-control condition). Adding to performance uncertainty, RF energy levels vary in time and are influenced by weather, local terrain/structures, and human activity. There is a higher likelihood that these factors will co-exist during low altitude urban flight operations, which suggests a critical need for effective RF environment monitoring, modeling, and/or forecasting.

To help mitigate this risk, a monitoring and assessment service was explored. To evaluate the feasibility of this concept, a set of commercial off-the-shelf (COTS) observations stations were installed at LaRC and an RF propagation model [34][58] was developed for the LaRC flight test range. The model was based on locations and types of known emitters and local-area digital surface models (DSMs). Data collected from the spectrum monitors served multiple purposes: (1) to compare with model-based estimates for validation or quality improvement, (2) to allow for human observation/assessment of the RF environment prior to and during operations, and (3) to characterize nominal and off-nominal states of the RF environment such that metrics can be developed for automated alerting functions. Significant testing was completed supporting the latter two purposes in collaboration with NASA's AAM) and High Density Vertiport (HDV) projects. Test data were also used by HDV researchers to support the flight operations approval process.

A prototype RF monitoring service was developed and is functionally described in Section A.5 in Appendix A. During the course of research and development over 2021-present, tests were conducted to raise the Technology Readiness Level (TRL) of this service. These tests are summarized below.

2.1.1.1 Transmitter Mapping

Electromagnetic (EM) wave propagation modeling (Winprop [58]) was applied to map RF levels emitted by a collection of known sources at NASA LaRC and validated by field measurements [34]. Terrain elements (buildings, towers, and trees) were assigned EM properties at the frequencies of interest, including transmission, reflection, diffraction, and scattering parameters (Fig 2.1.1-1, top left). The locations, power levels, and antenna patterns of nine NextNav towers distributed across an area of 6 square kilometers were specified; these towers radiated at levels up to 30 W in two bands centered on 926.227 MHz and 924.442 MHz. The predicted power for the 926.227 MHz band at two elevations was computed at altitudes of 100 meters AGL (Fig 2.1.1-1, top right) and 2.82 meters AGL (Fig 2.1.1-1, bottom left).

To validate the service product, a mobile measurement van was equipped with a receiving antenna and spectrum analyzer and driven across the site (blue trace in Fig 2.1.1-1, bottom left). The measured signal power tracked the predicted power within about 5 dB, including at locations predicted to be shielded by the terrain.

Mapping the RF environment in this way is also possible using RF source information, which is publicly available from government and private databases [59].

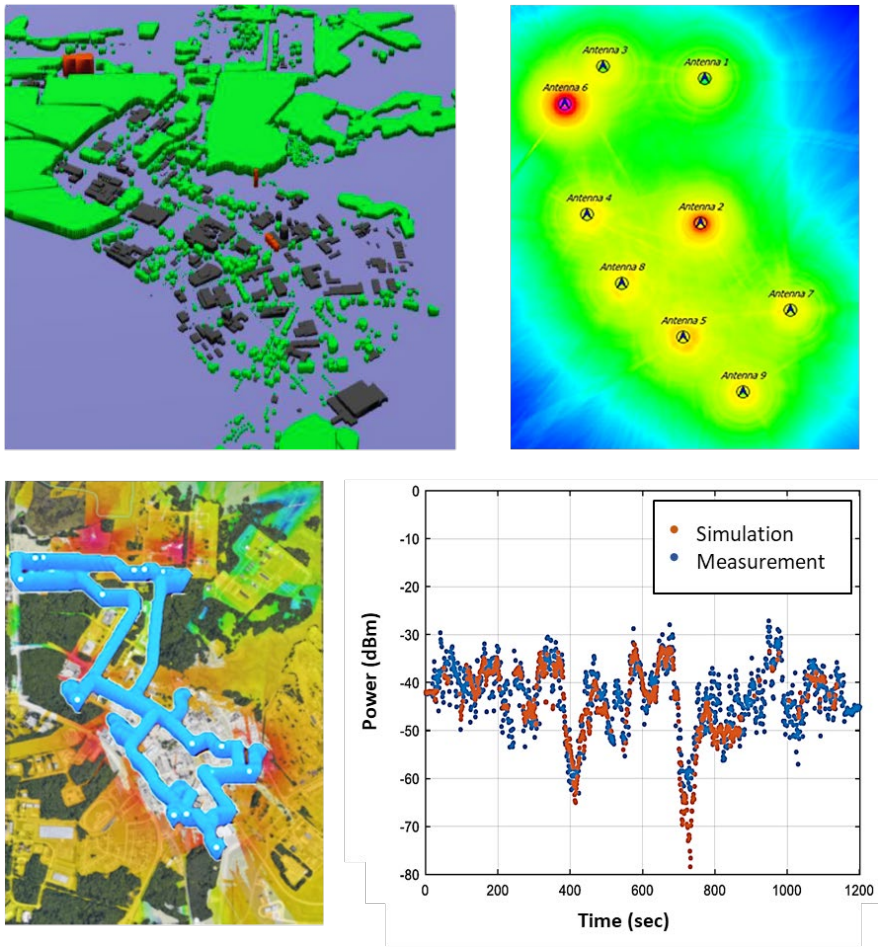

Fig. 2.1.1-1 RF Interference Mapping. Top left: Terrain property assignment. Buildings are shown in grey, metallic structures in orange, and trees in green. Top right: simulated power at 100 meters AGL from nine terrestrial navigation system transmitters emitting at 926.227 MHz. Color scale from low (blue) -50 to high (red/pink) -30 dBm. Lower left: simulated power at 2.82 meters AGL. Color scale from low (blue) -100 to high (red/pink) -30 dBm. The measurement drive path is superimposed (light blue trace). Lower right: comparison of predicted (brown dots) and measured (blue dots) power along the drive path showing prediction fidelity. Image credits: Google Earth, Open Street Map. From [34].

2.1.1.2 Spectrum Occupancy Monitor

A metric to help detect lost-link hazards based on spectrum occupancy in a common uncrewed aerial vehicle (UAV) telemetry band [31] was developed and tested. Spectrum occupancy is a measure of how much of a frequency band is in use in a given time interval. When occupancy is high, there is a higher likelihood that communications using that band will experience dropouts (i.e., a lost-link condition). Testing demonstrated the feasibility of this metric to discern the number of UAVs that were simultaneously transmitting in a commonly used telemetry band.

During testing, a commercial spectrum analyzer recorded signals in the 900 MHz range while the number of transmitting vehicles was systematically varied (Fig. 2.1.1-2). A standard signal processing method, frequency-based signal occupancy, was applied to these recordings

to evaluate its ability to discriminate the number of transmitters. In particular, the feasibility of using a simple threshold of the processed signal was demonstrated.

A set of five identically configured vehicles (designated A, B, C, D, and E) was operated in the LaRC flight range near a prototype monitoring station (Fig. 2.1.1-2, top right). The vehicles were configured with a Pixhawk® flight controller running the ArduCopter™ variant of the ArduCopter autopilot software [42], connected to a Mission Planner ground station [60] using an RFD900 [61] radio. The RFD900 radios operate in the 900-930 MHZ band, with a maximum power of 30 mW, using a frequency hopping link protocol. Each vehicle/ground station pair communicated via MavLink [62] messages over the RFD900 link. The UAVs were manually powered up and down individually and in increasing numbers in the sequence shown (Fig. 2.1.1-2, top left).

# Vehicles Transmitting	Local Time
0	11:11:32 AM
1(A)	11:21:01 AM
1 (B)	11:30:16 AM
1 (C)	11:36:35 AM
1 (D)	11:43:01 AM
1 (E)	11:49:52 AM
2 (A,B)	12:02:43 AM
3 (A,B,C)	12:14:42 AM
4 (A,B,C,D)	12:20:17 AM
5 (A,B,C,D,E)	12:26:44 AM

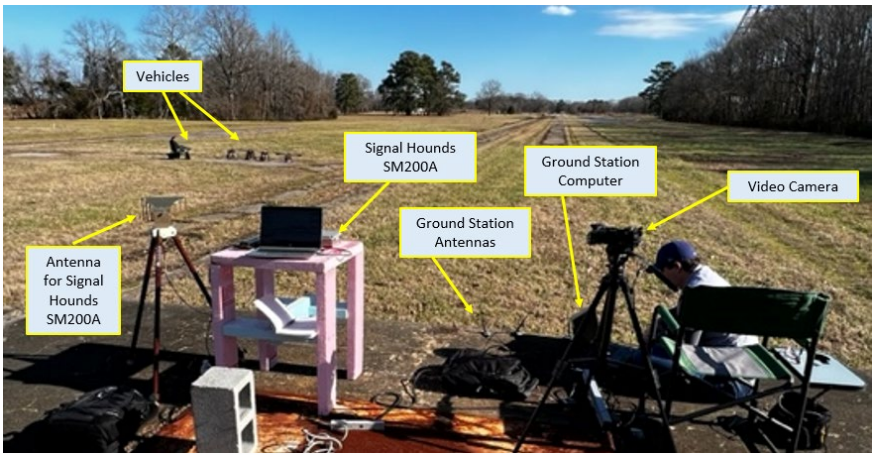

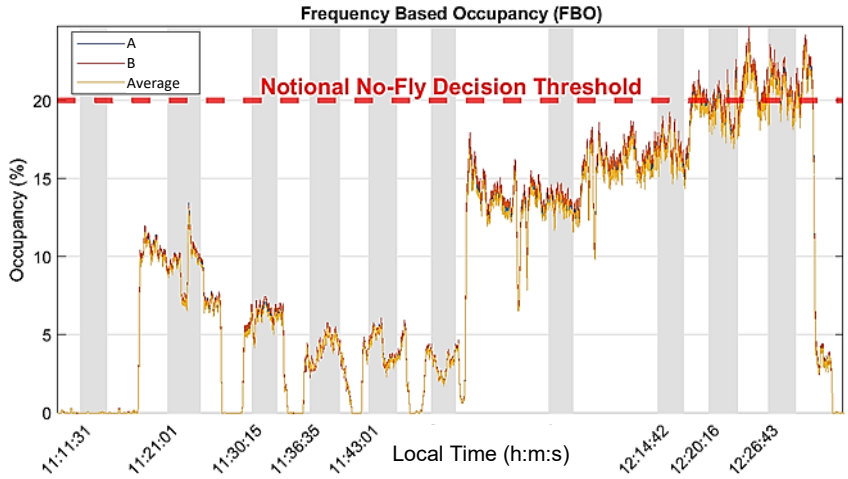

Fig. 2.1.1-2. Spectrum occupancy alert test. Five vehicles (A-E) were powered up in sequence (table, top left) and the ambient RF (900 MHz to 930 MHz) was recorded for setups such as shown (top right). A frequency based occupancy metric (bottom) was developed to alert if RF levels approach a nominal level; a 20% level, corresponding to four transmitting vehicles, served as a proxy for a level that could result in a loss of control link condition. In the bottom graph, time windows with grey shading cover vehicle powerup times associated with entries in the table at top left. Measurement traces are shown for a generic vehicle (Average) and for vehicles A and B.

2.1.2 NavQ Prediction

For low altitude flights, such as those for small UAS or larger UAM vehicles, positioning, navigation and timing (PNT) solutions typically employ inertial systems, optical systems, and global navigation satellite systems (GNSS) such as the global positioning system (GPS). Often

there is a critical reliance on satellite-based systems to provide continuous, independent measurements. Given the RF nature of satellite-based systems, this can result in degraded navigation system performance in complex environments (e.g., urban canyons) due to effects such as shadowing, multipath, and signal attenuation. As a means of reducing this risk, predictive performance capabilities were developed to support both pre-flight planning and in-flight contingency management. The service-based capability builds upon a tool that leverages available real-time satellite state information (e.g., ephemeris) and a digital surface model of the operational area. These are used to generate a discretized time-varying 4D volume (or set of waypoints) of estimates and forecasts of quality metrics associated with satellite-based measurements. Estimates may be generated across a user-selected flight environment, flight plan, and time period. These capabilities are used to help (a) during pre-flight planning to assist in the creation of safe flight paths and (b) in-flight to provide contingency management agents with navigation system-related risk of proximal flight corridors.

Two types of testing activities were completed to help evaluate/validate two navigation fidelity prediction services (NavQ CAPS [27] and GAPS [28], Figure 2.1.2-1). The first evaluated both services through comparison with over 6000 readings from GNSS sensors collected along an 8 km (5 mile) path through urban areas of Corpus Christi, Texas, on three dates in 2022. Predictions are based on satellite line of sight through 3D terrain data collected in 2018. Each service predicts a set of navigation fidelity metrics over a user-specified time period. One metric estimated by both is the number of visible satellites. A direct comparison of the number of predicted visible satellites with the number sensed by the receiver is used to validate the prediction services. Results showed an exact match in the number of predicted satellites for 60% of the measurements, and a match within ± 4 satellites for 95% of the measurements. As expected, agreement improves when farther away from (or above) vertical blocking terrain/structures. Most cases of mismatch are due to a lower predicted count than measured (false negatives) and can be accounted for by receiver pickup of stray signals caused by multipath propagation. About 10% of mismatches are false positives and are mostly accounted for by foliage effects. The two services predicted visibility of the same set of satellites 80% of the time, differed by two or less satellites 95% of the time, and could compute predictions for one hour of observations in one minute or less. Validation is analyzed statistically and in detailed case studies of selected observation times in [25].

The second set of tests were operational evaluations by end users. This was done for NavQ GAPS via a partnership with another government agency. Evaluations were completed at several large-scale events including the Macy's Thanksgiving Day Parade, New Year's Eve in Times Square, New York, New York, and the NFL Draft in Las Vegas, Nevada. Feedback from the partner was overwhelmingly positive, including one instance in which the service provided insight into the temporal aspect of GPS coverage, allowing real-time strategic decisions to be made to regain GPS availability in an acceptable timeframe.

For more information on NavQ services see Appendix A.

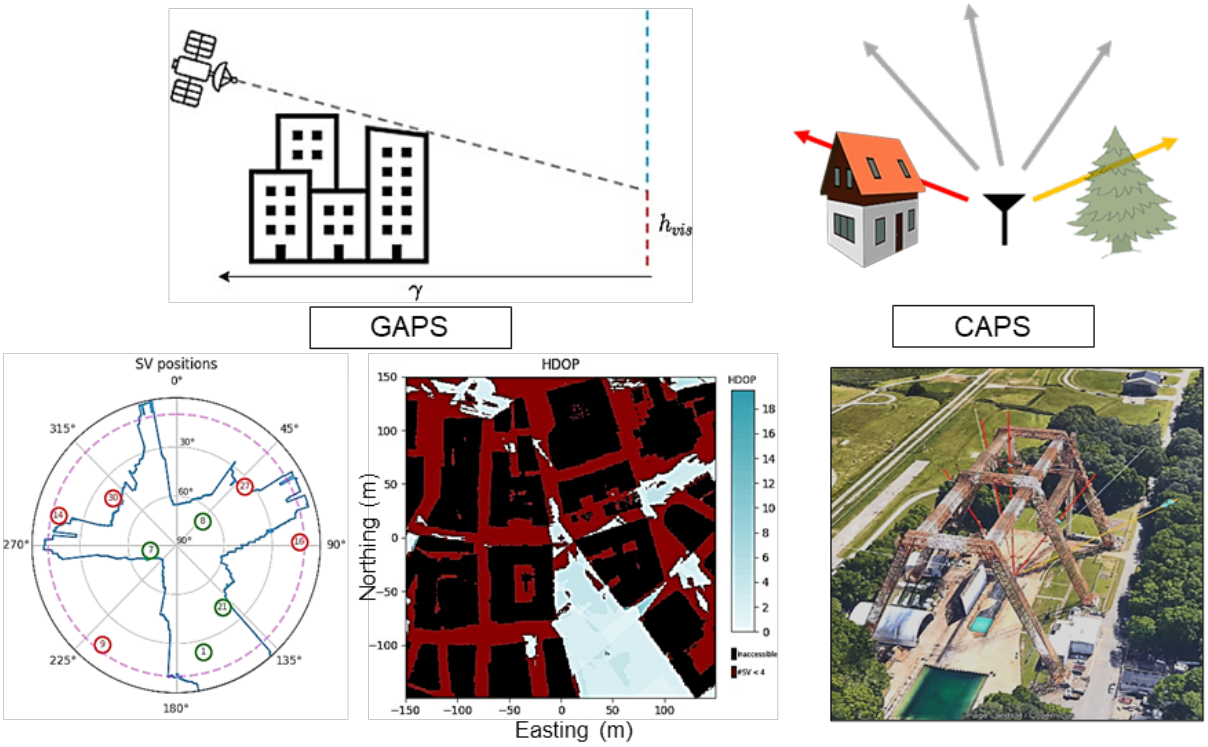

Fig. 2.1.2-1. Principles and test results for two navigation quality prediction services. Top left: The Geometric Assessment of Positioning Systems (GAPS) service casts shadows of terrain from overhead satellites to a 3D location. Top right: The Corridor Assessment of Positioning Systems (CAPS) services casts rays from a receiver location to the orbital satellites and determines whether each ray is blocked by terrain. Bottom left: For a location within an areas of interest, GAPS produce a skyplot (left) with the terrain skyline (blue trace) and the computed satellite availability, blocked (red circles) or in clear line of sight (green circles). Horizontal position accuracy (scale bar from white to blue) is computed for the entire area of interest, in this case downtown Boston. Areas with no satellites are colored red and buildings are colored black in this result. Bottom right: CAPS result at the NASA Langley Lunar Lander Research Facility. Rays to satellites are colored red if completely blocked by terrain, yellow if partially blocked by foliage, and grey if in clear line of sight. ©Graphics: NASA, NOAA, USDA; @Map data: Google, Landsat/Copernicus

2.1.3 Flight Performance Assessment

Multiple IASMS capabilities require flight dynamics and performance models in order to predict and assess feasible trajectories, power consumption, and operational envelope limits. Wind tunnel tests were conducted to improve the accuracy of multirotor vehicle aerodynamic models and to investigate stability and power usage characteristics relevant to flight safety. Models have been created for both sUAS and eVTOL vehicle types. A specific activity used NASA LaRC's 12-Foot Low-Speed Tunnel to help understand flight performance characteristics of an octocopter vehicle such as was used during the flight testing described in Section 3. Both isolated-airframe (no propellers) and powered-airframe tests were completed. A one-factor-at-a-time approach was used to explore the effect of propeller advance ratio and vehicle angle of attack on vehicle forces, moments, and power consumption. Results showed nonlinear aerodynamic behavior and indicated regions where aerodynamic interaction effects may be significant [63][64].

Model improvements based on test results have been (or are being) implemented for future simulation and flight tests. Results and models are used as part of a flight performance service,

which allows users/operators to request flight performance-related data based on a set of vehicle-specific parameters and other information if available (e.g., winds). Outputs include: predicted aircraft states (e.g. position, velocity, attitude); battery state of charge (SOC) and remaining useful life (RUL) estimates; and warnings for locations or times when envelope boundaries are predicted to be exceeded (see A.9). Similarly, the trajectory prediction functions are applied within developed risk assessment services (see A.1) [65].

2.2 Integrated Capability Testing

Several IASMS pre-flight capabilities are listed in Tables 3 and 4 that can be implemented using an integrated set of services and functions, a human-machine interface, and a defined human role/procedure. Preflight capabilities were integrated and tested using three graphical user interface (GUI) concepts that provide preflight planning predictive support allowing the operator to assess and mitigate flight hazards and risks:

- Supplemental Data Service Provider-Consolidation Dashboard (SDSP-CD) [43][44]
- Human Autonomy Teaming Interface System (HATIS) Operator Interface [44][45][66]
- Operation Planning Tool (OPT) [46]

Researchers conducted a usability study using two of the GUI concepts, (SDSP-CD and HATIS) to explore how individuals interpret and interact with different information provided by a collection of services during mission and flight planning [43][44]. As part of the testing, a series of preflight risk-assessment tasks was developed. Participants were trained to use each interface and their performance was evaluated. Sixteen participants worked through tasks using the GUIs and connected services. Objective data were collected on performance tasks across different simulated scenarios, as well as self-reports of interactions and subjective experiences using the GUIs. Scores on the system usability scale (SUS) and on simple/complex tasks were analyzed [43][44], as well as user feedback on open-ended questions, to inform development and identify potential improvements to the interfaces. The third interface (OPT) [46] was used during flight test operations but was not part of the usability study.

2.2.1 Integration Goals and Scope

In addition to the formal usability assessment for GUI designs, integration testing was performed to address the following technological and human factor considerations:

- Communications: Establishing and maintaining links to services
- Capacity: Managing the real-time data feed(s) of multiple services, each of which may have different update rates and product payload size/format
- Interpretation: Each information service issues data/alerts in dimensions pertinent (and potentially unique) to the type of hazard it addresses (e.g., numerical, geospatial, or temporal data/alerts)
- Saliency: Visual rendition of multiple alerts must be graded so that the most crucial ones are most visually conspicuous while others are available but not as demanding of operator attention

For the pre-flight services tested, data streams adhered to a common message protocol informed by the SDSP framework developed for UTM services [23][44]. Information exchange protocols were established as stream-based (request, reply) with all exchanges of packets in a custom format initiated by the GCS. As more services were added, a hypertext transfer protocol (i.e., HTTP-based POST, GET protocol) with standardized JavaScript Object Notation (i.e., JSON-based) packets was adopted to allow the GCS to poll several services at once and ingest the results of the request as needed.

2.2.2 Services Evaluated

Six preflight services were tested. As expected with independent development efforts, the specific subset of services tested with a particular ground station varied. To allow flexibility to ground station developers and researchers, all services employed a common communication protocol (stream- or HTTP-based; see below). For integrated testing of pre-flight capabilities, the following services were available. Descriptions are available in Appendix A. See also Fig. 2 and Table 1.

1. Battery prognostics and health management (BP/BHM): Given the flight plan, the service provided an estimated remaining flight time to an end-of-discharge event, so that the GCS could show the state of battery as the projected flight progressed and the probability of reaching end of discharge before the end of the mission.
2. Proximity to threat (PtT): Given the flight plan, the service provided points along the (possibly uncertain) flight trajectory where proximity falls below a client-defined threshold, nearest approach points (i.e., points along the trajectory where proximity is smallest), distance to nearest approach point, and a measure of the violation severity.
3. Third party casualty risk assessment (GRASP): Given the flight plan and a spatial distribution of unshielded ground population, the service provided estimated impact points if the UASs were to experience power loss or motor failure at points along the flight path and the associated probability of human casualty for each impact point. The service also reported the ground population distribution as a heat map of population density showing areas of high, medium, low, and sparse population density in the local area.
4. Radio frequency environment and interference monitoring (RFE/RFI)
5. Navigation quality prediction (CAPS)
6. Wind (WM): Given the ambient wind speed and direction at the time of mission planning, the service accessed and provided 3D windfields throughout the flight range. These fields were precomputed for a range of prevailing wind speeds and directions via a set of CFD (computational fluid dynamics) simulations on the flight range terrain populated with a static 3D database of building and tree geodata. For ease of display, wind fields were reported as polygons within 10m slices of altitudes from ground level to 100m above ground. Polygon boundaries were found by thresholding windspeed as low (< 5 m/s), medium (between 5 m/s and 10 m/s) and high (10 m/s or higher).

2.2.3 Tests with SDSP-CD

The SDSP-CD displays route planning guidance from all of the developmental services listed in 2.2.2, except for the Wind service. Two examples of preflight planning tests using these services with SDSP-CD are given below.

In the first test, waypoints for four flight plans were checked for proximity (<10m distance) from ground obstacles. Collision hazard results for a flight plan in a semi-urban area at NASA LaRC are shown in Figure 2.2.3-1. The PtT service received the flight plan waypoints from the ground station, computed the flight path and proximity risks, and transmitted risk-labelled flight segments back to the ground station. Safe segments were displayed as green and unsafe segments as red on the operator display. Hazard results for the other three flight plans and services (battery reserve, population overflight risk, and GPS fidelity) were similarly retrieved and displayed. An interface usability study gathered feedback from twelve experts in UAS hazard and risk analysis, and suggested improvements to display layout and style were incorporated. An updated ground station display was then tested using the same services. Example population overflight (i.e., via GRASP service) hazard results are shown in the updated display (Figure 2.2.3-2).

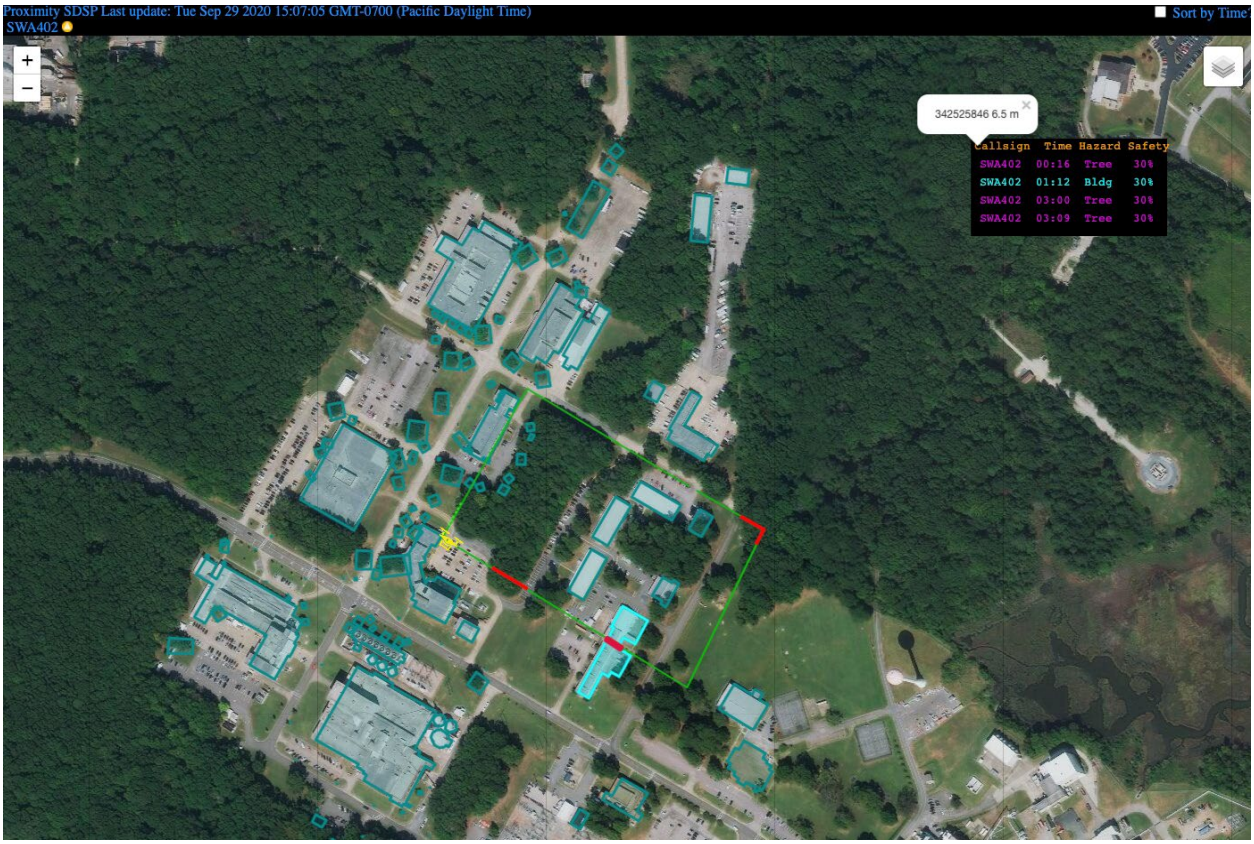

Fig. 2.2.3-1. Image of SDSP-CD display interface with proximity to obstacle hazard alerts displayed for a rectangular flight path over buildings and trees (from [43]). While most of the flight path is clear of obstacles (green segments), three collision hazards with trees and one collision hazard with a building (red segments) alert the flight planner to adjust the intended path. A hazard list overlay (top right) includes information about the distance from the flight path to the obstacle (NASA image)

Fig. 2.2.3-2. The updated SDSP-CD GUI shows alerts for four flights in the table (inset) and in a map display. Population overflight hazards are shown on the map for one of the four flight paths. From [43].

At the time of the first test, services were implemented with a stream-based (request, reply) protocol. Subsequently, an HTTP-based (POST, GET) protocol was implemented for preflight services, and the test was repeated. The RF hazard alert service was also added for the second test. Ancillary information, such as population density map layers, were added as well. Fig. 2.2.3-3 shows the same flight path with the underlying population cluster that gives rise to the hazard warning as an aid to flight replanning.

Fig. 2.2.3-3. Image of SDSP-CD display interface showing underlying population density that gives rise to the hazard alert highlighted in Fig. 2.2.3-2. Population density (yellow hexagons) and a probability of casualty heat map (red hexagons) is overlaid on a map of buildings (blue polygons). From [44]. (NASA image)

2.2.4 Tests with HATIS

The HATIS GCS Operator Interface [66] leveraged the same service information exchange protocol to create an interactive preflight methodology (Fig. 2.2.4-1). Upon flight waypoint entry (i.e., the initial flight plan), each service was queried to assess risks/hazards along the path. If the 'Auto Evolution' capability was enabled, a repositioning (or change) of any waypoint automatically queried all services again to re-check for risk/hazard potential, thus immediately providing feedback to the crew about the safety efficacy of the change.

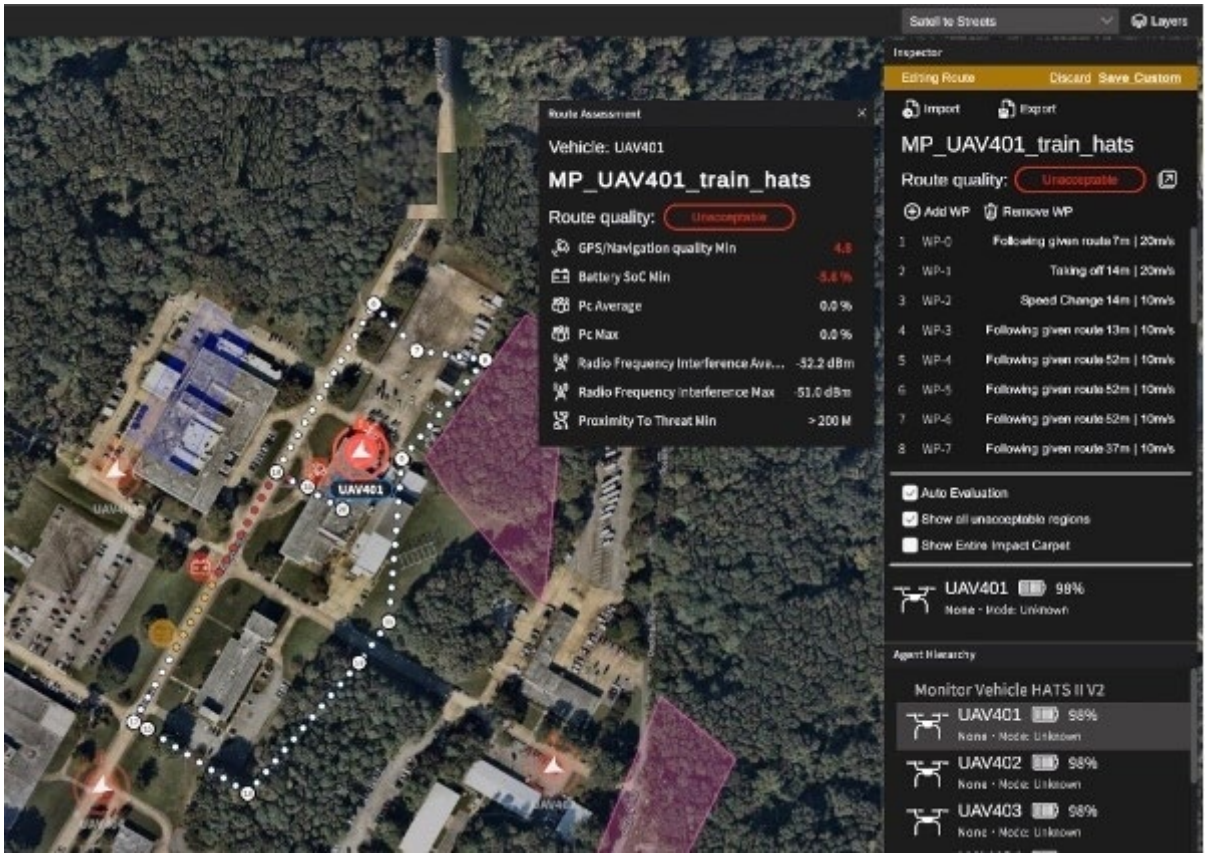

Fig. 2.2.4-1. HATIS display interface with preflight route assessment selected for a UAS with the callsign "UAV401". While population casualty, RF environment, and collision risks are assessed as acceptable, navigational and battery life preflight safety checks result in an 'Unacceptable' rating for the route. The waypoints of the proposed flight path (white, yellow, and red dots) can be edited (top right panel), and the safety checks are run again if a waypoint is changed. From [44]. (NASA image)

2.2.5 Tests with OPT

The OPT user interface was tested for preflight safety checking with all services listed in section 2.2.2. The services were accessed in sequence and produced graphical hazard overlays similar to the other ground stations (Fig. 2.2.5-1). OPT was tightly integrated with the GroundWatch research ground station used in flight [46], which allowed service safety parameters (e.g., navigational satellite count or safe proximity distance) checked in the preflight stage to be passed directly to in-flight safety monitors.

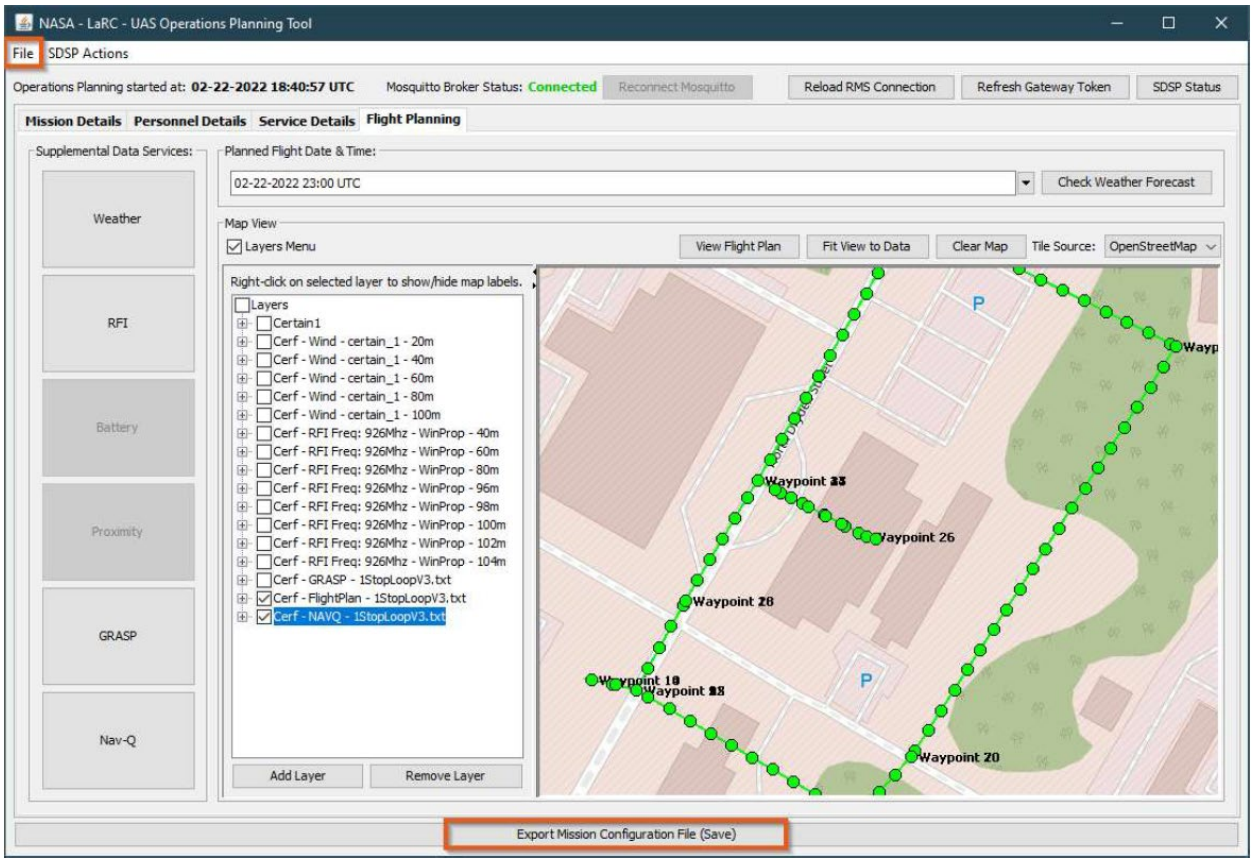

Fig. 2.2.5-1. OPT preflight route assessment for six SDSP-style services. The waypoints of the proposed flight path and intermediate points are shown as green dots. As services are invoked in sequence (buttons in left panel), with resulting hazardous portions along the flight corridor overlaid. No hazards are found in this example of a navigation system performance check (by NavQ). After all checks, a mission configuration that matches the preflight checks are passed to the Research Ground Station for use in flight. From [46].

3 Testing In-flight SFC Implementations

Both simulation and flight testing were conducted to evaluate a set of in-flight SFCs. Here we focus on flight testing conducted on NASA LaRC's test range to evaluate a set of onboard functions intended to provide a highly assured capability to automatically mitigate risk during flight (see [7] and [67]). Several functions worked in concert to provide this capability, including, for example, an integrated real-time risk assessment function and an automated contingency select and trigger function. The experimental aircraft, ground equipment, and flight operation checklists may be found in Appendix B.

3.1 Test Architecture: Monitor/Assess/Mitigate Information Flow

Fig. 3.1-1 represents the airborne functions tested and the information flow organized, from left to right, through the chain of monitor, assess, and mitigate. Note that some flight software modules contained elements of more than one function, and the following description is function-based rather than module-based. Output from a set of hazard monitors was routed to a hazard monitor function, which formed a risk assessment that could initiate a risk mitigation maneuver. These stages are described in turn below.

3.1.1 Hazard monitor functions

The scope of tested hazard monitor functions is as follows:

- Proximity to Threat. The service described in Appendix A was implemented as an onboard function to aid in hazard monitoring. Static obstacle boundaries of buildings and trees in Geographic JSON (GeoJSON) format are loaded at mission start by the Research GCS, and the UAV position from the navigation GPS is compared to these boundaries twice per second [1][29]. Obstacles are comprised of actual buildings and trees in the flight range supplemented with artificial buildings, all represented as right perpendicular polyhedra extruded from ground footprints to the maximum obstacle elevation (28 buildings and 130 trees). In addition, boundaries that enclose the flight path within an approximately 60-meter corridor are loaded to monitor flight containment. The Proximity to Threat function issues an alert whenever the three-dimensional distance between the UAV and an obstacle is less than a specified length. In these experiments, the alert distance was set to 15.2 meters (50 feet). The Hazard Likelihood function (described below) considers such an alert as a 100% likelihood of a hazardous event and recommends a hover flight maneuver to the autopilot (ArduCopter 'Position Hold').
- Battery Health. A battery health monitor function (based on the BP service described in Appendix A) was implemented using an electrochemical model of the propulsion battery [47][48]. The function continuously estimated the stage of charge (SOC) and remaining flight time (RFT) based on battery temperature, voltage, and current draw. The resulting RFT was continuously compared to a minimum value of 200 seconds set by the operator for the testing, and a battery capacity falling below this value prompted a 100% hazard likelihood assessment with a recommendation of an immediate land maneuver to the autopilot (ArduCopter LAND). The 200 second minimum was determined empirically to suffice for standard propulsion batteries but can be varied if needed.
- Pilot Radio Link Monitor. Pilot control integrity was monitored by comparing one channel of the 2.4 GHz remote control (RC) link to a minimum acceptable pulse duration (950 microseconds, determined from inspection of the pulse width modulation signal characteristics). Loss of this link prompts a 100% hazard likelihood assessment with a

recommendation of an assigned location landing maneuver to the autopilot (ArduCopter 'Return to Launch')

- Navigation GPS Monitor. The number of received satellites (GPS and GLONASS[7] constellations) and horizontal position uncertainty (HDOP) were compared to preset warning levels (8 satellites and 5 meters) and failure levels (6 satellites and 10 meters). The HDOP

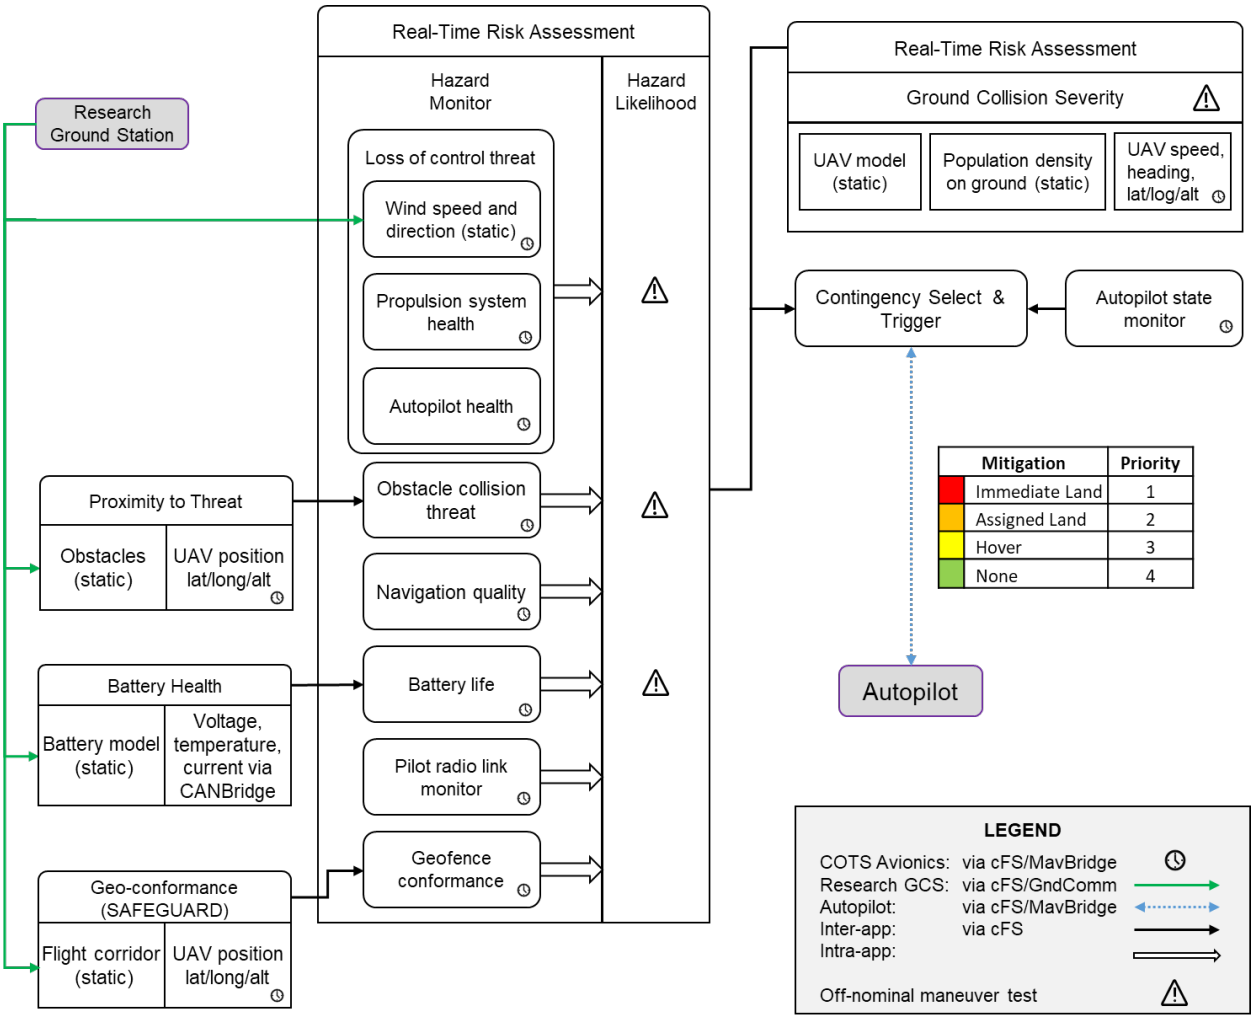

Fig. 3.1-1. Software architecture for active autopilot control test flights. A Core Flight System (cFS) backplane passes messages between airborne software modules and manages bus connections to avionics. Three onboard hazard monitor functions (Proximity to Threat, Battery Health, Real-Time Risk Assessment - Hazard Monitor) receive mission-critical readings from vehicle avionics and pass alerts to a hazard likelihood assessment function (Real-Time Risk Assessment – Hazard Likelihood). Each alert and a recommended hazard-specific safety maneuver is reported to a mitigation function (Contingency Select and Trigger). In the case of a loss of control threat, the maneuver varies according to the potential severity to people on the ground at the projected collision location. The mitigation function selects the most urgent recommended action and commands the autopilot to abandon its current course and maneuver to safety. The monitor/assess/mitigate framework of all safety hazards shown were verified in simulation and/or logged-maneuver flight tests; as indicated by the off-nominal icon, three were further verified in active-maneuver flight tests.

[7] GLObalnaya NAvigatsionnaya Sputnikovaya Sistema

warning and failure minimums of 5 meters and 10 meters were determined by considering safe obstacle distances (See Section A.3, Proximity to Threat Assessment) and a typical HDOP performance of 1-2 meters. Determination of satellite count minimums is necessarily less rigorous, as position fidelity depends on the constellation geometry as well as satellite count (see Section A.2 Navigation Quality Assessment). Since a minimum of four satellites is required for a position fix, satellite counts 50% and 100% higher were chosen. Failure-level navigation impairment prompted a 100% hazard likelihood assessment with a recommendation of an immediate land maneuver to the autopilot (ArduCopter LAND).

- Geospatial Conformance Monitoring. Trajectory conformance to the intended flight volume was monitored by an independent highly assured conformance monitoring function (aka, the Safeguard system [49][50]). The function uses predefined geo-fence polygons that represent the boundaries of no-fly zones and stay-in areas. When these boundaries are approached by the vehicle, warning signals are generated. As an added layer of safety, Safeguard continuously computes the predicted vehicle impact trajectory in the event of total power loss. If the predicted trajectory crosses defined boundaries (i.e., geo-fence polygon edges), two types of flags, Warning and Critical Warning, can be issued. A Critical Warning flag prompts a 100% hazard likelihood assessment with a recommendation of an immediate land maneuver to the autopilot.
- Loss of Control Threat. The risk of aircraft failure is computed based on three hazard monitors:
    - Autopilot Health. Either a loss of the autopilot heartbeat signal (over 10 seconds, to allow for radio intermittency) or excessive aircraft vibration (as measured by accelerometer saturation and clipping) are regarded as an autopilot failure.
    - Propulsion System Health. The experimental aircraft is propelled by eight electric motors driven via individual electronic speed controllers (ESCs). Propulsion system health is gauged at a whole-aircraft level, at the component level, and at an intermediate (collection of components) level. Both warning and failure levels of health are reported. For brevity, only the failure levels are listed here. See [7] for a full description of the computation. Parameters values were determined by examining archived flight logs and by interviewing subject matter experts.

      Excessive aircraft vibration (> 55 m/s) serves as the whole-aircraft measure of propulsion system health.
      At the component level, a propulsion failure is reported if either excessive motor speed (10,000 revolutions per minute for 5 seconds) or excessive ESC temperature (85 degrees Celsius, 8 amperes) is detected.

      Propulsion failure at the collection of components level is reported if one or more of the following are detected:
        - high ESC temperature variation (two or more ESCs 25°C different from the mean ESC temperature);
        - high ESC current variation (two or more ESCs drawing current 3.8 amperes different from the mean ESC current); or
        - high ESC voltage variation (two ESC voltages 10% different from main battery voltage).
    - Wind speed and direction. Three levels of wind (none, low, high) wind speed are reported. High winds are those greater than 5 m/s.

Recommended maneuver per hazard monitor

Hazard Detected	Mitigation
Navigation Loss	Immediate Land
Command/control link loss	Assigned Land
Proximity to obstacle	Hover
Loss of control	see table

Maneuver selection priority

Mitigation	Priority
Immediate Land	1
Assigned Land	2
Hover	3
None	4

Recommended maneuvers for loss of control hazard

Likelihood of Loss of Control	Loss of Control Severity: Minimal	Minor	Major	Catastrophic
Frequent	Immediate Land	Immediate Land	Immediate Land	Immediate Land
Probable	None	Assigned Land	Assigned Land	Immediate Land
Remote	None	Assigned Land	Assigned Land	Assigned Land
Improbable	None	None	None	None

Fig. 3.1-2. Risk assessment and mitigation logic. Top left: Maneuvers which minimize vehicle damage are specific to each assessed hazard. Bottom: A vehicle loss of control condition may present a hazard to people on the ground. Maneuvers which minimize population risk are recommended based on ground population at a crash point determined by the vehicle's ballistic trajectory combined with the likelihood of vehicle failure. Top right: Maneuver selection is prioritized by urgency.

3.1.2 Risk Assessment Functions

Risk assessment is immediate for all hazard monitors except the Loss of Control Threat monitor. As shown in the top left table of Fig. 3.1-2, navigation loss (Navigation Quality monitor) is judged as most acute and an immediate land maneuver is recommended; a lost command and control link (Link Monitor) is less acute and aircraft flight to an assigned landing point is recommended; and an obstacle collision threat (Proximity to Threat Monitor) implies that further aircraft traversal along its flight plan is risky and that the pilot should take control, and so a hover maneuver is recommended. These maneuvers correspond to ArduCopter flight modes as follows: Immediate Land is LAND mode; Assigned Land is RTL mode; and Hover is POSHOLD mode. An entry of 'None' in Fig. 3.1-2 means that no command should be issued to the autopilot, i.e., the current flight mode should continue.

If a loss of control condition is likely (Loss of Control Threat Monitor), the Real-Time Risk Assessment function projects the point of ground collision based on the wind speed and the aircraft location, heading, and flight speed. Population density, sheltering effects (i.e., whether the population is indoors), casualty impact area, and the kinetic energy at impact are evaluated [7][23][24] to estimate the probability of a casualty (P_C) caused by the falling vehicle. A recommendation which minimizes harm to people on the ground is determined by considering two assessments (bottom left table of Fig. 3.1-2): a severity rank (minimal, minor, major, and catastrophic, based on quartiles of P_C) and a ground collision likelihood rank[8] (improbable,

[8] Mapping of loss of control likelihood (P_{LOC}) values to likelihood rank: Improbable $0 \le P_{LOC} < 0.01$; Remote $0.01 \le P_{LOC} < 0.1$; Probable $0.1 \le P_{LOC} < 0.5$; Frequent $0.5 \le P_{LOC} \le 1$

remote, probable, frequent). Low likelihood and low severity assessments result in no recommended command change to the autopilot. A combination of moderate to high assessments drive a recommendation to either maneuver to an assigned landing point or to land immediately.

A static wind from the east with speed of 8 m/s at 10m altitude was set at the start of these test flights. In this high wind condition, the loss of control likelihood is elevated to a 'Probable' or 'Frequent' level, depending on autopilot and propulsion health metrics. To avoid equipment loss, actual autopilot/propulsion health was not impaired in flight experiments. Based solely on the high constant wind speed, the loss of control likelihood was effectively set to at least the 'Probable' level (second row of bottom table in Fig. 3.1-2). This was convenient for testing because vehicle movement from areas with no ground population to areas with ground population would change the recommended mitigation from None to Assigned Land or Immediate Land, depending only on population density.

3.1.3 Risk Mitigation Functions

The Contingency Select and Trigger function continuously collects maneuver recommendations from the hazard/risk monitor and assessment functions, as well as from an autopilot state monitor (i.e., APMon). As multiple hazards may occur at any particular time, this function prioritizes the most urgent safety action consistent with the current flight state (top right table of Fig. 3.1-2) and within the capability of the autopilot given its current state. A recommendation to land immediately is considered the most urgent/acute maneuver and is prioritized highest; a recommendation to maneuver to an assigned landing point is prioritized as the next most important mitigation; a recommendation to hover in place is prioritized last. If no hazard alert reaches the Contingency Select and Trigger (CST) function, a "no-op" status is logged, and no maneuver command is issued to the autopilot.

3.1.4 Maneuver Execution Verification

An Autopilot Monitor function, developed using the Co-Pilot formal methods tool [68], keeps track of the current autopilot flight mode and assures that a given mitigation action command from the Contingency Select and Trigger function is executable and valid given the current context of the flight. A switch to Immediate Land (ArduCopter LAND) is allowed from any autopilot state; Hover (ArduCopter POS_HOLD) is allowed as long as the autopilot reports healthy navigation and velocity; Assigned Land (ArduCopter RTL) is allowed as long as the autopilot reports healthy navigation and velocity and valid landing coordinates; and a return to the waypoint flight plan (ArduCopter AUTO) is allowed as long as the autopilot reports (a) healthy navigation and velocity, (b) a valid set of flight waypoints, and (c) positive pilot permission (as indicated by a throttle setting over the RC command link).

In the test flights, mitigation commands generated by CST were correct by construction, and so the Autopilot Monitor did not deny any actual commands, though the logic of the safety checks was verified preflight in laboratory tests. Current development includes a fault injection regime at the startup of the Autopilot Monitor function to allow effectiveness checks of the command verification logic on each flight day.

3.2 Flight Tests with Mitigation Maneuvers Logged

Twelve flight tests were conducted in the spring of 2022 with airborne research avionics disconnected from the autopilot. Performance of the hazard monitoring, risk assessment, APMon, and CST functions were logged and evaluated for validity post-flight. Experiments varied the number and combination of hazards in flight to check that the monitor/assess/mitigate architecture resulted in the expected mitigations. As described in [7] and below, an isolated occurrence of each hazard type was tested and occurrences of some (but not all) combinations of two or more of simultaneous hazard types were tested. Flights were conducted in a test range free from structures and ground population. A set of 'virtual buildings' were added to pose obstacle collision hazards and a set of 'virtual crowds' were added to pose population overflight risk (Fig. 3.2-1). Full details of these experiments are available in [7].

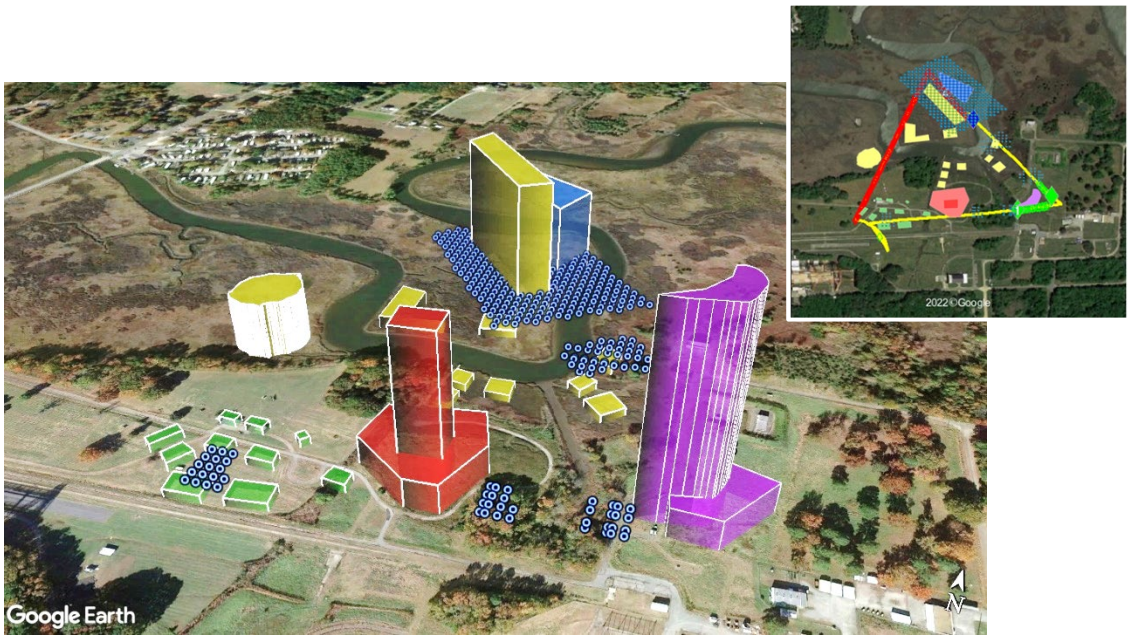

Fig. 3.2-1. Collision and population overflight virtual hazards. Population is shown as blue circles and structures as colored geometries in 3D (foreground) and 2D (inset).

A central objective of this flight campaign was to verify that automatically generated mitigation actions in the face of multiple hazards were prioritized correctly and executed immediately. Fig. 3.2-2 shows the safety situation at one moment during flight 68 (from [7]). This image was rendered to produce a 3D visualization using actual flight data and log data collected using the test architecture of Fig. 3.1-1. At this flight position, three hazards were assessed as acute. One of them – Proximity to Obstacle – requires a Hover (POSHOLD) maneuver, and two of them – Battery Health and Loss of Control – require an Immediate Land (LAND) maneuver. As designed for the system under test, the Contingency Select and Trigger mitigation function issued the most urgent maneuver, Immediate Land.

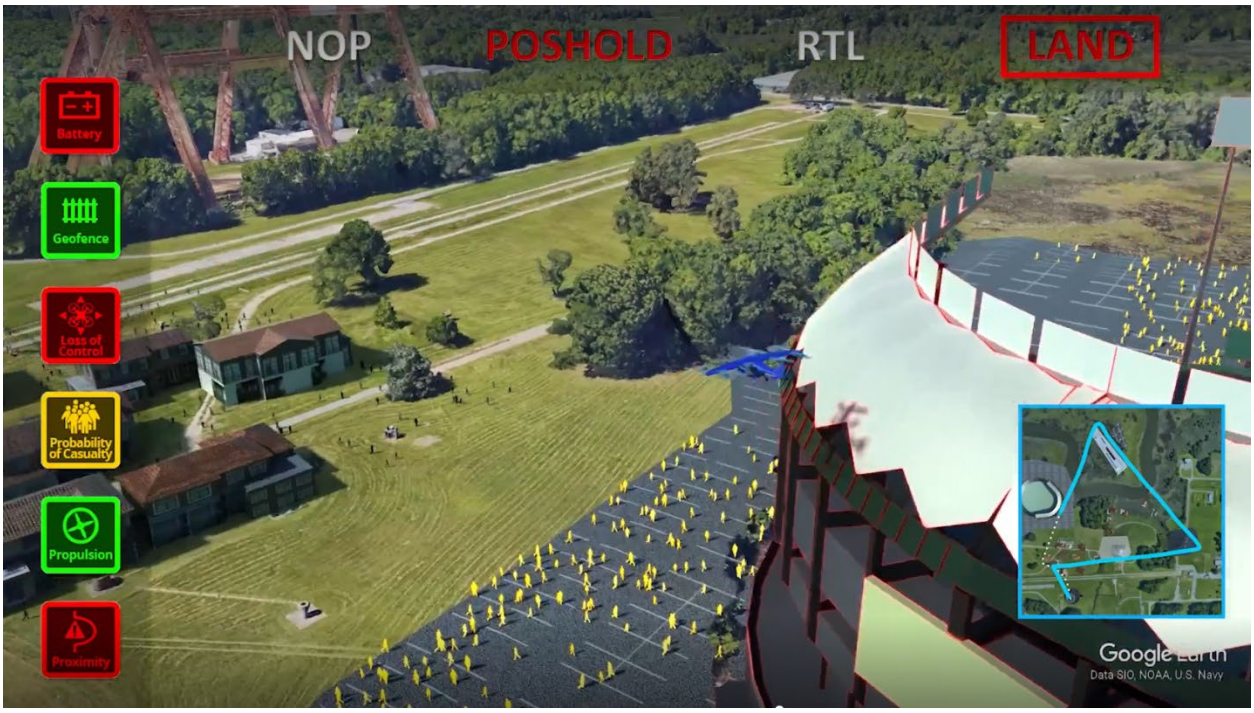

Fig. 3.2-2. Example of hazard states and decision making at one moment in Flight 68. The aircraft (middle right) is near a building/structure and there is ground population in the vicinity. Estimated hazard states are indicated at left (green: no threat; yellow: possible threat; red: immediate threat); these suggest three hazards require action. The mitigations available to maneuver safely are shown at top (grey text: not a suggested action; red text: maneuver options suggested; red text (boxed): highest priority maneuver option given all hazard states). In this test, the Contingency Select and Trigger mitigation function issued the most urgent maneuver, Immediate Land. From [7].

3.3 Flight Tests with Mitigation Maneuvers Active

A series of nine test flights on four flight days in the spring of 2023 further exercised the monitor/ assess/mitigate architecture. In addition to logging the outputs of functions (for post-flight analysis), the recommended command maneuvers actively controlled the vehicle autopilot (i.e., closed-loop operation).

With confidence gained in the logging flights (previous section) that mitigation actions are timely and effective, autopilot control was tested. Fig. 3.3-1 shows an example from flight 102, rendered in a 3D visualization from actual flight data (from [44]). At this flight position, two hazards were assessed as acute. One of them – Proximity to Obstacle — requires a Hover (POSHOLD) maneuver, and one of them — Probability of Casualty — requires diverting to an Assigned Landing point (RTL[9]). As designed for the system under test, the Contingency Select and Trigger mitigation function issued the most urgent maneuver, divert to an Assigned Landing point.

[9] The RTL sequence is: a) climb to preset safe altitude; b) fly direct to the nearest pre-planned rally point; c) descend and land.

For most flights, several conditions were tested in succession. Waypoint-based flight paths were constructed so that aircraft would encounter a series of hazard conditions that triggered a mitigation action (e.g., Hover). After each hazard encounter that triggered a mitigation, the pilot would recover aircraft control, direct it along the flight path past the hazard zone, and then reinitiate autonomous vehicle traversal along the flight path. A dedicated channel of the pilot R/C controller switched an onboard relay ('Comms Kill Switch' in Fig. B1.1 in Appendix B, referred to hereafter as the 'research relay') that established or interrupted commands sent from the research computer to the autopilot. At vehicle launch and during pilot recovery, the research

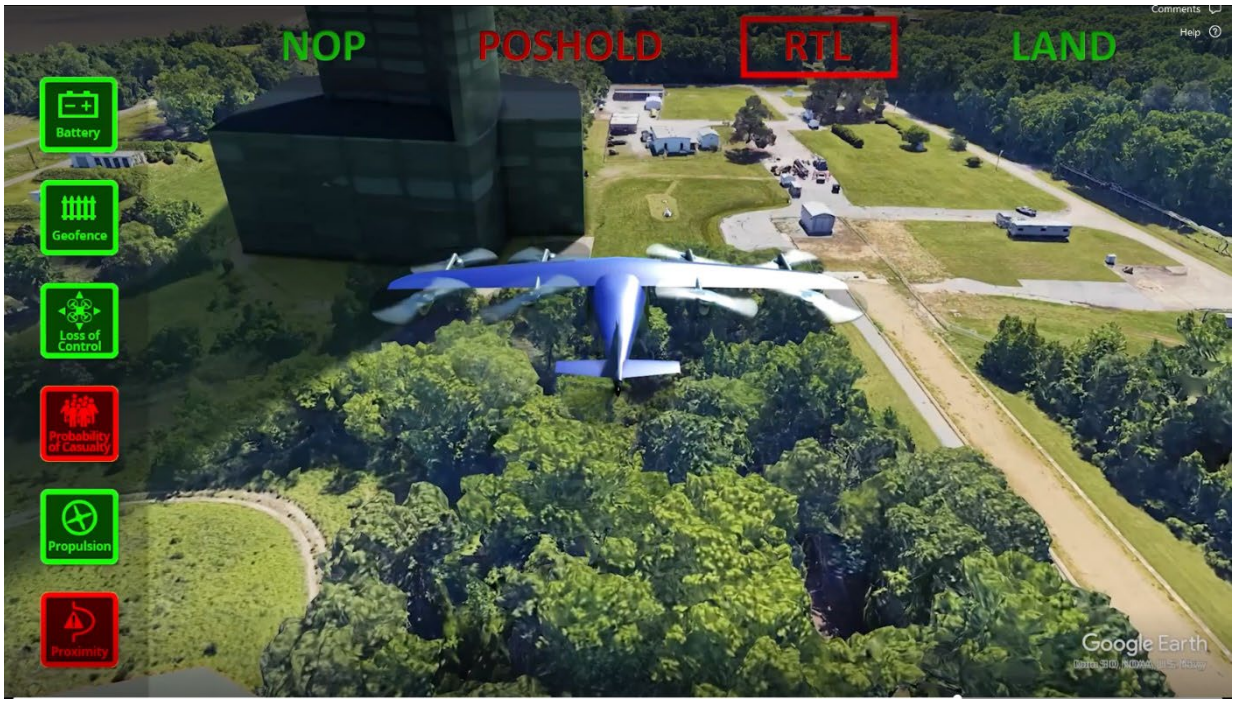

Fig. 3.3-1. Example of prioritization of the most urgent action with two safety alerts to direct the most urgent flight command at one moment in flight 102. Collision with a virtual building is threatened and vertical descent to land would threaten the underlying ground population. The monitor/assess/mitigate architecture selected RTL, diverting the aircraft to land at an alternate safe landing site. From [44].

relay was opened to prevent research commands from reaching the autopilot. Waypoint-based flight could therefore operate in one of two ways: via default automation (i.e., autopilot modes subject to interruption only by autopilot failsafe settings) or via safety-enhanced automation (i.e., autopilot modes additionally subject to interruption by the mitigation commands issued by the onboard research computer).

For two research flights, a detailed analysis of the hazards and the monitor/assess/mitigate safety response is included below. Flight test results for an additional seven research flights are presented in brief form in Appendix C.

3.3.1 Flight 97 Hazards and Monitor-Assess-Mitigate Operation

A representative flight (F097) is shown in Fig. 3.3-2. In this flight, four hazardous conditions were encountered:

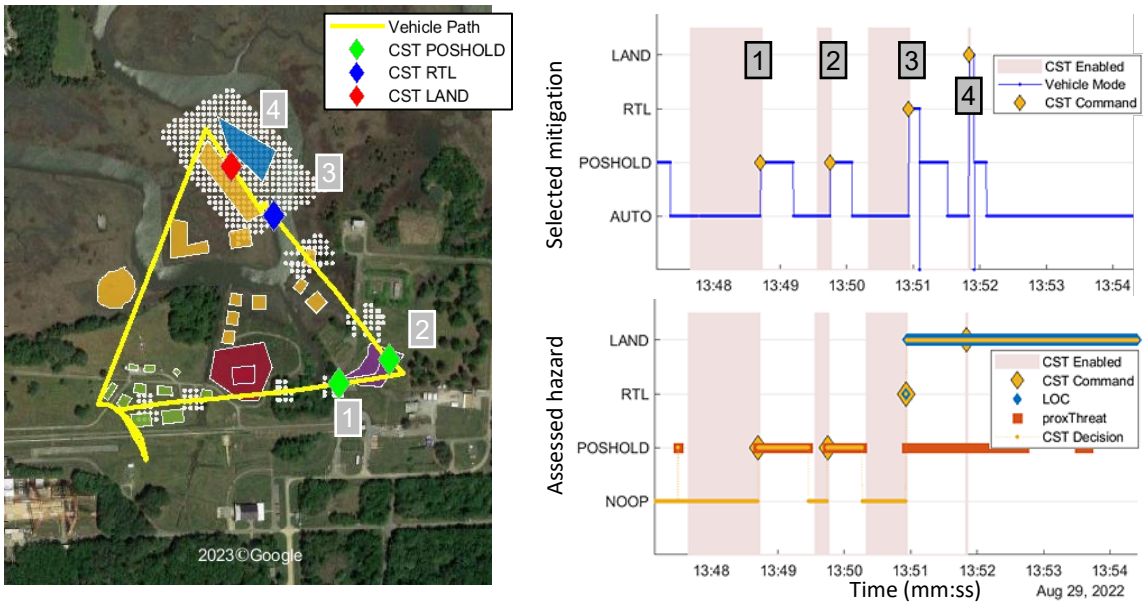

Fig. 3.3-2. Left: Flight path (yellow) for flight F097, annotated with mitigation maneuvers (green, blue, and red diamonds). Virtual ground population is shown as white squares and virtual buildings are shown as colored polygons. Onboard safety monitors detected four hazard conditions during safety-enhanced autonomous flight. Onboard assessment led to two Hover mitigation maneuvers (green diamonds labeled 1 and 2), one Assigned Land (or RTL) maneuver (blue diamond labeled 3) and one Immediate Land (or LAND) maneuver (red diamond labeled 4). Right: Monitor-assess-mitigate events during the flight. At four intervals (solid pink shading), experimental commands were enabled until a maneuver command (yellow diamonds) halted safety-enhanced autonomous traversal along the triangular flight path in response to a hazard condition. After the first three maneuvers were triggered, the pilot took control, disabled experimental maneuver commands, directed the vehicle past the hazard area, and reinitiated autonomous flight with the experimental connection to the autopilot. After the last maneuver, the experimental connection was kept disabled for the remaining flight and landing. The top diagram shows the prioritized mitigation and flight mode. The bottom diagram shows all assessed hazards (and the corresponding safety mitigation), which were prioritized to compute the safety maneuver. For events 1 and 2, only obstacle proximity hazards were encountered, and so Hover (POSHOLD) maneuvers were the sole priority. For event 3, a population overflight hazard and an obstacle proximity hazard were encountered simultaneously, and the higher priority Assigned Land (RTL) mitigation was issued to the autopilot. For event 4, a more severe population overflight hazard and an obstacle proximity hazard were encountered simultaneously, and the higher priority Immediate Land (LAND) mitigation was issued to the autopilot.

1) Upon launch, the vehicle climbed and flew to the first waypoint of a triangular flight path. As the vehicle began flight along the bottom leg, the research relay was closed, initiating safety-enhanced autonomous flight. Following a hazard-free traversal of the bottom of the path, an imminent collision with a virtual building (solid green diamond labeled *1* in Fig. 3.3-2) was detected by the Proximity to Threat monitor; a Hover maneuver was recommended by the Hazard Likelihood Assessment function, and the Contingency Select and Trigger mitigation function issued a Hover command to the vehicle. The pilot opened the relay to stop further research-generated commands to the autopilot, and using default autonomy, flew the aircraft through the corridor (that is very close to a virtual building) to the next waypoint at the bottom right of the triangle. The pilot then closed the relay to re-initiate autonomous waypoint traversal with the research system engaged.
2) Within a few seconds the Proximity to Threat monitor flagged another close approach hazard, triggering a second Hover event (solid green diamond labeled *2* in Fig. 3.3-2) and the pilot again took control and flew the aircraft out of the virtual danger zone. The

pilot closed the research relay to enable safety-enhanced autonomy within a few seconds.

3) Enroute to the waypoint at the top of the triangle, more than a minute of safe (hazard-free) flight followed until a ground population was approached. With loss of control likelihood elevated by a constant high wind condition, the projected ground collision severity computed by the Real-Time Risk Assessment – Ground Collision Severity function exceeded a safe level, and an Assigned Land maneuver was recommended. The Contingency Select and Trigger mitigation function issued an Assigned Land (RTL) command to the vehicle (solid blue diamond labeled *3* in Fig. 3.3-2). Though a Proximity to Threat alert was occurring, its mitigation (Hover) is of lower priority and was disregarded. The pilot then opened the relay connection between the research systems and the autopilot, commanded a vehicle Hover briefly, and then set the vehicle to resume waypoint traversal using default autonomy.
4) After about a further half minute of flight, the pilot re-established the research-to-autopilot link. In the interim, two mitigation recommendations were being issued once per second by the hazard assessment function in response to hazard alerts. First, the vehicle had moved to an even higher loss of control severity area, so that an Immediate Land recommendation was pending from the Real-Time Risk Assessment – Ground Collision Severity function. Second, at this location in a narrow corridor between two buildings, a Hover recommendation was pending from the Proximity to Threat function. Once the control relay was closed, the Contingency Select and Trigger function acted on the higher priority Immediate Land recommendation and took the craft into a descent maneuver (solid blue diamond labeled *4* in Fig. 3.3-2). The pilot retook control within seconds, disabled research systems communication with the autopilot, and resumed waypoint traversal via the autopilot's native autonomous flight capability. The flight proceeded through the remaining waypoints and landed at the planned landing coordinates.

3.3.2 Flight 115 Hazards and Monitor-Assess-Mitigate Operation

Figure 3.3-3 shows a flight experiment in which the monitor-assess-mitigate flight safety architecture responded six times to hazardous conditions. In flight F115, mitigation actions again resulted from obstacle collision threats and loss of control (population overflight) threats. Additionally, a battery health alert prompted a mitigation maneuver, and a low telemetry level signal triggered an autopilot failsafe maneuver.

The flight path is again triangular in a corridor with virtual buildings and ground populations (Fig. 3.3-3, left). The vehicle flew under the control of the autopilot's default waypoint-based autonomy except when the research relay was closed (shaded pink intervals in time plot on the right in Fig. 3.3-3). For brevity, pilot recovery actions to place the vehicle back on the flight path after a maneuver are not included in some of the following event descriptions. Seven off-nominal events occurred in flight F115:

1) After the vehicle took off and climbed, it proceeded to the first waypoint of the triangular flight path and began eastward flight. An obstacle collision threat was issued by the Proximity to Threat monitor when distance from a virtual building was less than the 15.2 m (50 ft) threshold. The Real-Time Risk Assessment – Hazard Likelihood Assessment function recommended a Hover maneuver and the Contingency Select and Trigger mitigation function issued a Hover command to the vehicle.

2) Further east on the flight path, with a wind blowing to the east, a population casualty threat was detected by the Real-Time Risk Assessment – Ground Collision Severity function as it projected the likely landing location of the vehicle as a point within a virtual ground population and the perceived impact was expected to cause casualties. The combination of a 'Probable' loss of control likelihood and a 'Minor' loss of control severity (bottom table in Fig. 3.1-2) dictates an Assigned Land mitigation, and as the Contingency Select and Trigger mitigation function was receiving no competing recommendations, it initiated an Assigned Land maneuver.
3) At the start of northwesterly flight, the threat of collision with a building resulted in a Hover maneuver.
4) Further northwest, a population overflight threat resulted in an Assigned Land maneuver. The pilot opened the research relay so that the vehicle flew using default autonomy until the vehicle passed the ground population.
5) As the vehicle neared a virtual building, a collision hazard was imminent, but a transient low signal level for the 900 MHz telemetry link to the ground station was first detected by the autopilot. The failsafe maneuver (Assigned Land) set in the autopilot preempted commands from the research systems[10] and the vehicle began flying to the nearest safe landing location. The pilot opened the research relay after a few seconds and directed the vehicle to hover. The ground station operator alerted the pilot of the low signal alert and the pilot waited for about one minute for the radio signal to recover. The pilot then commanded a change from hover flight to default autonomous flight, and the vehicle moved back to the centerline of the flight path. Once satisfied that the vehicle was on the correct course, the pilot closed the research relay and safety-enhanced autonomous flight commenced.
6) The vehicle again neared the virtual building and a cluster of virtual people on the ground. Two recommended mitigations were sent to the mitigation function: an Assigned Land maneuver from the Real-Time Risk Assessment – Ground Collision Severity function and a Hover maneuver triggered by an alert from Proximity to Threat function. The Contingency Select and Trigger mitigation function prioritized the more urgent Assigned Land mitigation and issued a maneuver command. The pilot opened the research relay and safety-enhanced autonomous flight ceased, so that the vehicle flew using the default autopilot autonomy. Flight progressed to the next waypoint (top of triangle in Fig. 3.3-3) and then to the southwest.

[10] Off-nominal event number 5 in F115 was classified as an autopilot failsafe event because at the time of event 5:
a) The Contingency Select and Trigger function was recommending NOOP (no change to current flight mode, labelled as 'None' in Fig. 3.1-1 and Fig. 3.1-2).
b) A system status variable available in the autopilot heartbeat message transitioned from ACTIVE (which is the normal state) to CRITICAL (which is an off-nominal state).
c) A status text message from the autopilot reported "GCS Failsafe."
The first observation (a) eliminates the possibility that research avionics initiated the maneuver and the other two observations confirm that a native autopilot failsafe occurred.

7) A third of the way along the last leg of the flight path, the Battery Health monitor detected a low charge condition and the Real-Time Risk Assessment – Hazard Assessment function accordingly issued an Immediate Land recommendation. The pilot opened the research relay and put the vehicle into a hover. The ground station operator confirmed that the battery was low and called for an end to the flight. The pilot manually flew the vehicle toward the launch point and landed it using the default autopilot autonomy.

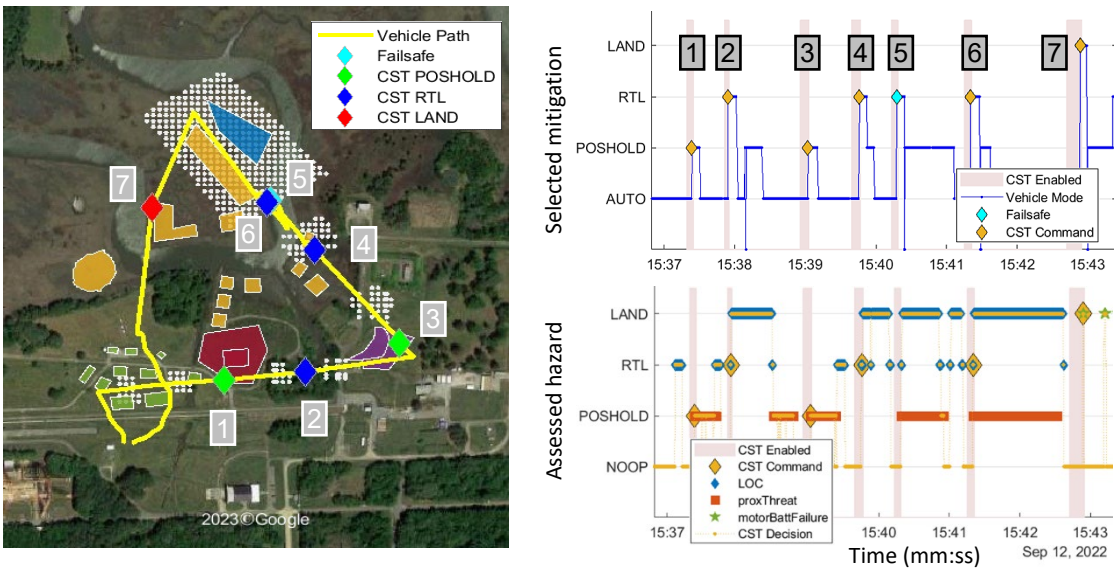

Fig. 3.3-3. Left: Flight path (yellow) for flight F115, annotated with mitigation maneuvers (green, blue, and red diamonds) and an autopilot failsafe maneuver (cyan diamond). As in flight F097 above, the pilot took command after each maneuver and advanced the vehicle along the flight path before reinitiating safety-enhanced autonomous flight. Hazard detection led to six mitigation events: two Hover mitigation maneuvers (green diamonds labeled 1 and 3), three Assigned Land maneuvers (blue diamonds labeled 2, 4, and 6) and one Immediate Land maneuver (red diamond labeled 7). Right: Monitor-assess-mitigate and failsafe events during the flight. At seven intervals (solid pink shading) experimental commands were enabled. During one of these intervals, a lost telemetry link to the ground station triggered an autopilot failsafe (event 5); in the other six intervals, a maneuver command halted safety-enhanced autonomous traversal along the triangular path (yellow diamonds) in response to a hazard condition. The top diagram shows the prioritized mitigation and flight mode. The bottom diagram shows assessed hazards and the corresponding safety mitigations considered to compute each maneuver. For events 1 and 3, only obstacle proximity hazards were encountered, and so Hover maneuvers were the sole priority. For events 2 and 4, only population overflight hazards of moderate probability/severity were encountered, and so Assigned Land maneuvers were the sole priority. For event 6, a population overflight hazard and an obstacle proximity hazard were encountered simultaneously, and the higher priority Assigned Land mitigation was issued to the autopilot. The battery health monitor detected a low charge condition near the end of the flight (event 7), triggering an Immediate Land maneuver.

4 Testing Post-flight/Off-line SFC Implementations

As described in [1], data monitored and recorded during flights can enable safety assurance capabilities that execute after flights have completed (or “off-line”). These include: (1) the ability to identify precursors, anomalies, and trends that may only be observable when comparing data collected across a number of similar flights; (2) the ability to update and validate pre-flight and in-flight capabilities, and any underlying models to improve their performance; (3) the ability to report anomalies or off-nominal conditions that may indicate that design changes or maintenance actions are needed; and (4) the ability for humans involved in operations to report

safety-relevant observations or metadata to help in understanding the flight data and/or the operational context of a flight. In an IASMS, these “off-line” capabilities may be highly automated and supported by appropriate services and functions such that, when needed, post-flight activities can be completed in a timely manner (i.e., supporting “in-time” mitigation of safety risk).

Development and integration of specific services and functions that would enable such capabilities was limited during recent testing by the NASA team. This is primarily due to the limited availability of data (i.e., records/logs spanning many similar flights for the chosen application domains), and to external activities aimed at determining how such data collection and sharing could be implemented on a broad scale (i.e., something akin to the Aviation Safety Information Analysis System program [69]).

However, three relevant activities were initiated. The first investigates anomaly detection given a limited set of data/flight logs and applies an approach previously used for space operations. The second explores what could be identified using a larger set of flight logs, including from web-based forums where flight logs are posted by COTS sUAS autopilot users. The third creates a new means of collecting information on UAS incidents and accidents via the Aviation Safety Reporting System (ASRS). A brief summary of each of these activities follows.

4.1 Anomaly Detection from Limited Datasets Using IMS, MMS, and Active Learning

The pipeline of functions shown in Fig. 4.1-1 were implemented to investigate the applicability of a machine learning anomaly detection (AD) capability originally developed for aerospace systems [51][70], and to determine its utility given the constraints of a limited parameter set and a limited number of data collection flights (e.g., such as were conducted during NASA’s internal testing using sUAS).

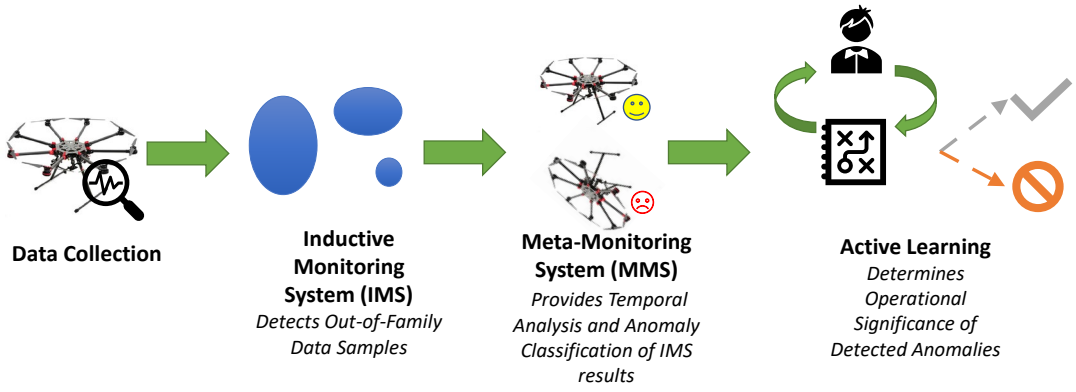

Fig. 4.1-1. From [51] Anomaly Detection Pipeline using IMS, MMS, and Active Learning.

In this application, the data collection portion of the pipeline obtains values from multiple sUAS parameters captured onboard by the ArduPilot software and either fed directly to the anomaly detection software for real time monitoring (e.g., executing onboard), or archived for off-line (and off-board) evaluation and model-building. After filtering out non-flight segments, these data are grouped into vectors representing snapshots of values at particular times and passed to the Inductive Monitoring System (IMS).

As described in [52], IMS is a data-driven anomaly detection software that uses nominal system data to build models of parameter interactions, dependencies, and behaviors seen during typical operations. Subsequently, the IMS monitoring module compares new data to these models and produces “deviation from nominal” scores to numerically indicate how far

each snapshot data vector falls from nominal operations (as defined by the training data used to build the models). The model building/training portion and analysis/monitoring portion are two independent software modules. The model building/training module would run offline, while the analysis/ monitoring portion may be configured to run offline as well (on a ground-based system) to evaluate archived data or configured to run onboard for real-time monitoring.

Features of IMS:

- Analyzes multiple parameter interactions
  - Automatically extracts system parameter relationships and interactions
  - Determines highly focused, context sensitive parameter limits
  - Detects variations not readily apparent with individual parameter monitoring practices
- Automatically derives models off-line from archived nominal operations data
  - Does not require off-nominal data for model building
  - Does not require knowledge engineers or modelers to capture details of system operations
- Inputs are observations about the physical system (parameter values)
- Outputs a “distance from nominal” anomaly score

The IMS results are then passed to the Meta Monitoring System (MMS), which was developed to help interpret IMS deviation scores. While IMS provides a relative measure of how anomalous a particular set of parameter values is, MMS considers the time history and trends of the IMS scores to categorize the recent system behavior as nominal or off-nominal. Like IMS, MMS has two phases: a model-building/training phase and a monitoring phase. MMS takes IMS's deviation scores as its only input. MMS primarily uses deviation scores from nominal data but can also make statistical use of some results from anomalous data if available. In the second (training) phase, it builds two models: one of nominal deviation scores and one of anomalous deviation scores. Each model consists of a probability distribution of deviation scores. These models are built in tandem in order to fit the nominal data well, and to give good classification performance on the training data. In the second (monitoring) phase, incoming IMS deviation scores are passed to the learned MMS models, and probabilities of producing the observed sequence of deviation scores are calculated for both models. These are combined in a ratio that gives the probability that an anomaly is occurring at the given observation. A threshold on this ratio is used to provide a binary nominal/off-nominal classification. Further details on MMS are available in [51].

Features of MMS:

- Post processes IMS output
- Analyzes one dimensional signal
  - Automatically extracts expected behavior and transitions
  - Detects anomalous trends over time not readily apparent with IMS alone
- Automatically derives models from IMS processed archived operations data
  - Does not require off-nominal data, but can use it to refine models if available
  - Does not require knowledge engineers or modelers to capture details of system operations

The final step in the AD pipeline is to determine if any detected anomalies are operationally significant (OS) or not operationally significant (NOS) – i.e., is this an event the operators should be concerned about as a potential safety issue or just something that happens occasionally but is acceptable behavior? This step falls to the Active Learning module. Active learning is a method to model subject-matter expert (SME) knowledge in order to classify anomaly detection results (in this case IMS and MMS scores) as either OS or NOS and, as a result, require

minimal time from often-busy system experts. If IMS/MMS detects some system behavior(s) that are statistically anomalous, Active Learning can use its model to match the current situation with previous system expert input and report whether (and why) the event is OS/NOS.

Features of Active Learning:

- Obtains domain expert feedback on whether the anomalies found are operationally significant (OS) or “false alarms” (i.e., not operationally significant (NOS)), and why
- Judiciously chooses samples for the domain expert to review and more accurately re-classifies all event examples as OS/NOS based on these reviews
- Respects domain expert time by finding examples for expert labeling that are most critical; there are no minimum required number of examples to label
- Resulting active learning model is used to minimize “false alarms” during operations

Unfortunately, resource limitations did not allow implementation of the entire pipeline through Active Learning for this project, but the data collection/IMS/MMS portions were evaluated (offline) using archived data from 178 flights of two different sUAS. The approach was evaluated using parameters collected during flights performed on the LaRC test range. The set of 56 parameters included MavLink [62] aircraft state variables (e.g., position, velocity, acceleration, attitude, and attitude rates) and system state indicators (e.g., battery measures, communication link status, and vibration levels). The archived data records were filtered to exclude non-flight data segments. Due to limited availability of data parameters from UAV hardware components, the anomaly detection efforts concentrated on unusual flight characteristics rather than onboard mechanical issues.

The first efforts considered data from 55 flights of two DJI S1000 drones and built models incorporating data from both aircraft. These data included brief flight log entries recording events from each flight, including two flights with noted issues. The anomaly detection system was able to identify the two noted anomalies, as well as four other out-of-family events, as shown in Fig. 4.1-2. Captions provide some details on the circumstances.

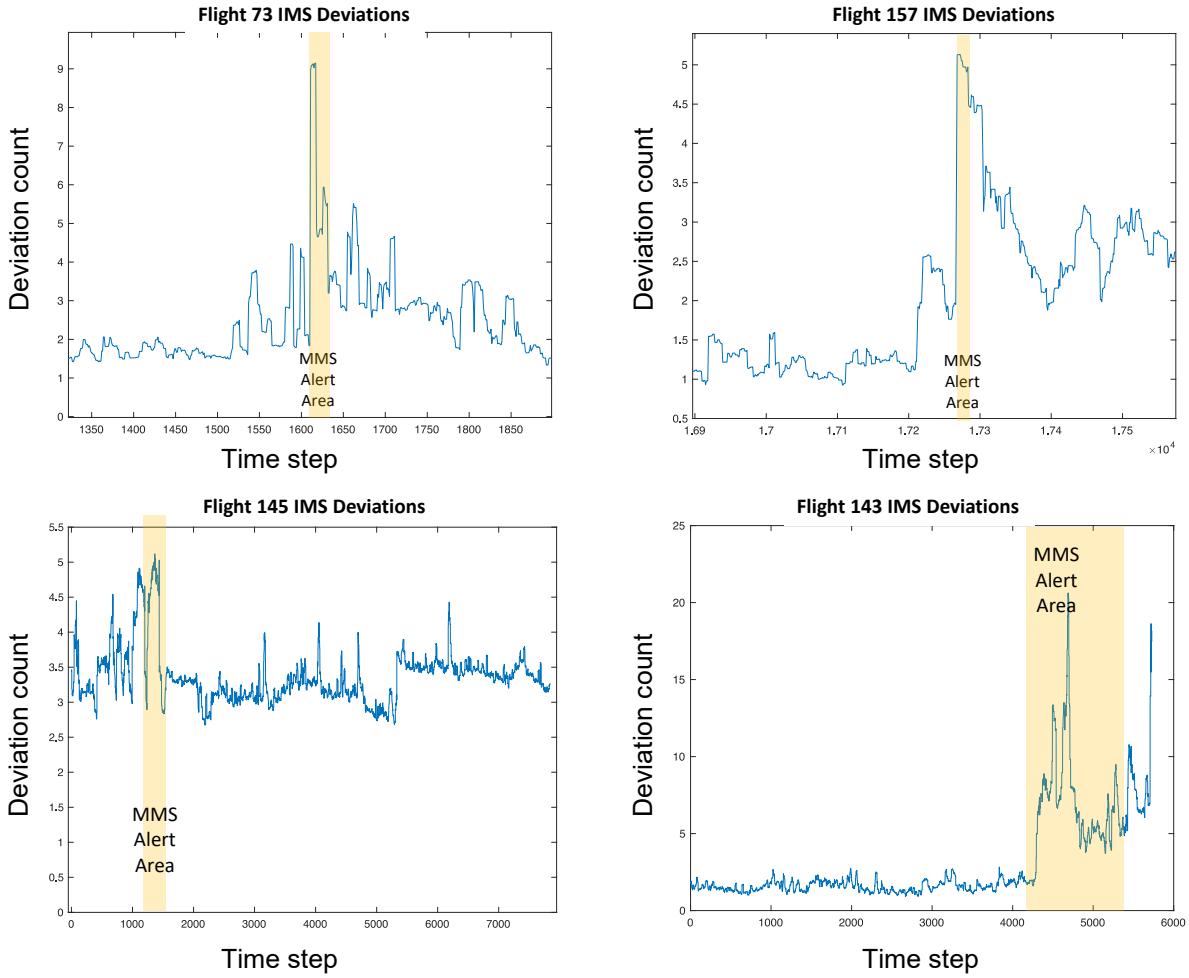

Fig. 4.1-2. IMS and MMS outputs for four anomalous flights: (top left) sudden heading error, possibly due to compass issue; (top right) spike in RC channel for yaw command; (bottom left) erroneous heading information; (bottom right) GPS failure.

Following the DJI drone flights, there were flights using two Tarot T18 sUAS referred to as Charlotte and Aragog. Charlotte had the most active flight schedule early on and provided data from 44 flights for anomaly detection development. Unfortunately, Charlotte had a relatively short operational life span. The final flight concluded with a fault that put the vehicle permanently out of service. Although there was limited sensor data covering Charlotte's mechanical systems, the nascent failure may have revealed itself via steadily increasing IMS deviation scores in the flights leading up to the final flight. The last flight showed an even more dramatic upward deviation trend when compared to its predecessors.

The Aragog drone took Charlotte's place and eventually provided 79 flights worth of data. These data files were not accompanied by any flight logs or other notes, so the machine learning team took on the role of pseudo-SMEs to classify extreme flight activity, such as unusual attitudes or abrupt maneuvers, as "off nominal." The test objectives of the sUAS flight program required a wide variety of maneuvers and flight conditions, so there was a wide range of complex flight activity from which to select unusual "anomalous" events. Penultimately, ten out of 61 of the initial Aragog flights were labelled as "anomalous" and used to construct IMS & MMS machine learning models intended for deployment onboard the drone for in-flight monitoring. A late influx of data from 18 new Aragog flights provided an excellent test of the deployment models. These data were processed with these models as if they had been running

onboard during the 18 unseen flights, with very good analysis results. Fifteen flights were categorized as nominal and three as anomalous. Each of the off-nominal segments in those three flights exhibited extreme maneuvers that easily qualified as out of family with significant statistical extreme values in one or more parameter values. Data from the new nominal flights were subsequently incorporated into the deployment models to improve their performance.

4.2 Off-line Model Creation and Management Using Flight Logs

Substantial products of the Aragog data analysis were the IMS and MMS models slated for on-board deployment for near real-time inflight systems monitoring. A similar model building process could be used during everyday operations with frequent UAS/UAM/etc. flight activity. Data logs would be collected from vehicles during each flight and formatted for processing by (and eventually training) IMS and MMS. These data would be analyzed using existing anomaly detection models to find off-nominal occurrences that would be brought to the attention of flight managers, perhaps after processing through the Active Learning tool to classify each event as operationally significant or not. The flight managers could take appropriate action, such as taking a vehicle offline and sending it for maintenance if warranted. As in the case of the Aragog study, the new flight data could then be incorporated into the IMS and MMS models to improve their performance for use during the next cycle. Models could be built for each specific vehicle or generalized for use with multiple vehicles depending on the homogeneity of the fleet and flight operations.

4.2.1 On-Board Anomaly Detection

Each of the experimental sUAS vehicles described in section 4.0 includes an onboard computer running core Flight Software (cFS) [71][72], a software framework designed to promote development of reusable embedded flight software. The cFS environment used in these vehicles includes, among other functionality, applications which communicate sensor values, current system status, etc. The data on the cFS bus will include all the parameters which have so far been collected and distributed offline for the purpose of training the machine learning models used in this study. A prototype cFS application is under development to allow live, onboard monitoring of vehicle data using IMS and MMS.

The application logic monitors each cFS data channel carrying pieces of information used for IMS and MMS. Whenever a new measurement arrives on one of these channels its value is stored by the application. Once per iteration of the application's main loop, the latest value of each monitored field is captured and stored. Some of these values are used to calculate derived parameters, such as lateral acceleration or roll and pitch errors. All the values collected this way are then concatenated with a timestamp and that vector of values is forwarded to the IMS monitoring process for analysis. The cFS application manages an instance of the IMS monitor which runs for the duration of each flight. Every time a valid IMS data row is constructed during portions of data collection deemed to be part of a flight segment it is passed to the monitor process to be assigned a deviation score along with optional parameter contribution data. Currently, this IMS output is streamed to a file on the onboard computer's disk. Future development will add MMS processing of the IMS results to this application, which should be a comparatively simple addition in terms of input/output complexity.

The current version of this application only stores results onboard and does not send its outputs to other onboard functions (e.g. RTRA or CST) or to the ground station during flight. However, this capability is planned for flight tests to occur later in 2024. Monitoring of the output by ground personnel (or automation) could influence decisions to land early or otherwise mitigate risks posed by detected anomalies. Sending IMS and MMS output to a ground station

would also facilitate the addition of live monitoring by a ground based Active Learning module in the future.

This concept could also scale up for full scale emerging operations. Each vehicle in the airspace would run anomaly detection processes similar to the cFS application developed for this project. Their flight data would also be collected onboard for later analysis and AD model updates, as was done with the Aragog data. If an anomaly was detected in flight the vehicle could inform pilots, dispatchers, and flight controllers of the situation for proper management within the established airspace protocols.

4.2.2 Anomaly Detection and the Drone Response™ Framework

For the sUAS domain, one of the project's partner teams has completed significant work under the umbrella of in-time safety assurance and IASMS concepts. Several new capabilities have been investigated and advanced [11][12][13], with the research continuing into early 2024. Follow-on research has also been awarded through NASA's University Leadership Initiative (ULI) which extends the scope and continues until 2026. The current IASMS project is entitled "Safe Deployment of Small Unmanned Aerial Systems through On-Board Monitoring and Assessment" [14]; while the ULI project is entitled "A Safety-Aware Ecosystem of Interconnected and Reputable sUAS" [15]. Both activities are led by the University of Notre Dame and are based in part on the Drone Response™ framework [16][17]. A brief summary of aspects related to anomaly detection and the use of flight log files is provided below.

As discussed in [11] and [12], the team needed data (e.g., flight logs) to investigate various approaches to the general anomaly detection problem for sUAS flight behavior. A two-fold approach was taken: (1) collect flight data using COTS sUAS for a (large) set of flights conducted by the team; and (2) collect flight data via online user forums.

For (1), flight data was collected using five baseline flight plans executed repeatedly by multiple sUAS, as well as some additional flight plans which deviated from their baselines. Flights were completed across multiple timeframes and weather/wind conditions. Logs were annotated with observed anomalies and off-nominal conditions/behaviors. At the time of publication of [11], 20 flights had been completed with significantly more planned. Each individual flight log was parsed using a suite of anomaly detectors, with detected anomalies checked to confirm whether they actually occurred during the flight but not observed/noted at the time. The flight logs were then augmented to add explanations of any such behaviors. Data parameters recorded and monitored aligned with the list in [1] which includes 16 classes of information. However, not all classes (and parameters) were available to record. The team has made the dataset publicly available in order to get feedback from the community on its utility and completeness with respect to the in-time safety assurance concept [18].

For (2), more than 53,000 flight logs were collected from the PX4 user forum website [73] and included logs from both simulations and physical flights. This set is extremely rich in terms of the number and types of flights flown by users across the globe. It is also useful in that data from this source are fairly consistent in form and format, often including discussions about the anomalous behavior which was compared against the actual flight logs to determine the most likely cause of the anomaly. However, these log files provide very limited metadata about conditions of the flights and whether discussion/observations are conjecture or evidentiary.

The team's general approach is to provide in-flight onboard, in-flight offboard, and post-flight offboard capabilities for anomaly detection and analysis (Fig. 4.2-1). Each option exhibits different tradeoffs in terms of resource constraints, and latency and response time requirements. To help determine key parameters to monitor, search-based techniques identified unsafe flight controller configurations and reviewed documentation and online forum discussions to identify common sUAS failures modes, hazards, and contributing factors. For example, analysis of ArduPilot documentation generated a list of 19 groups of attributes, with up to 20

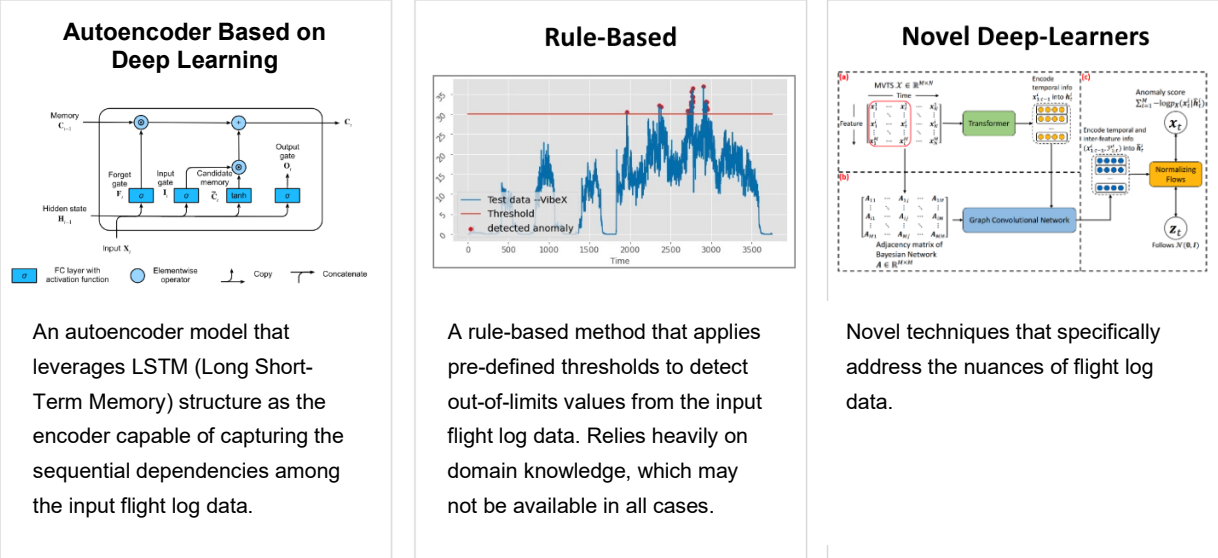

Fig. 4.2-1. Anomaly detection types evaluated by the Notre-Dame-led team. (left – Deep-Learning Autoencoder [74]; center – Rule-Based [75]; right – Novel Deep Learners)

attributes in each group. Reference [12] summarized the most-common failure issues mentioned across the ArduPilot and PX4 forums.

The initial anomaly detection capability focused on three types of failures related to mechanical problems, vibration, and compass interference. Tests used autopilot logs collected from the two sources mentioned above. Several hundred (670) logs were collected for training and validation purposes from 23 sUAS operated by the team. An additional 46 logs from the ArduPilot user forum were used. From these logs, 135 were randomly selected to look for vibration anomalies, 112 for attitude anomalies, and 38 for compass anomalies. 80% of the data were used for training and 20% for validation. Seven anomaly detectors were evaluated within the classes shown in Fig. 4.2-2. Results are provided in [11].

Ultimately the goal is to not only detect anomalies, but to also generate explanations that are useful to non-experts. This involves developing explanation models (currently based on heuristics, augmented by a diagnostic model learned from the forum data, and informed by operational or technical SME input), for each type of anomaly. The approach taken by the NASA team will be investigated as one means to achieve this end goal (e.g., see Section 4.1 on IMS, MMS, and Active Learning).

The above research is part of a larger ecosystem concept aligned with NASA's in-time safety assurance and IASMS concepts. The team refers to this as the Drone Response™ platform [16][17] which was primarily designed to support complex emergency response operations involving multiple autonomous sUAS. Each sUAS is capable of detecting anomalies as they occur; but anomalies that may only be exposed when looking across flights can also be monitored. To date, testing of this concept has used both MAVLink (Micro Air Vehicle Link [62]) and PX4's uORB (Micro Object Request Broker [73]) asynchronous publish-subscribe messaging application programming interface (API). Development is toward a dedicated onboard micro-service that collects, aggregates, and analyzes data provided by all sUAS in the ecosystem. Information about relevant detected anomalies is then made available via the service to support safety risk mitigation in advance of future flights. The team recently field-tested the approach, demonstrating the viability of running a set of seven anomaly detectors onboard, detecting anomalies in close-to-real-time, and either raising an alert or performing a simple mitigation such as reverting from an incorrectly configured flight parameter.

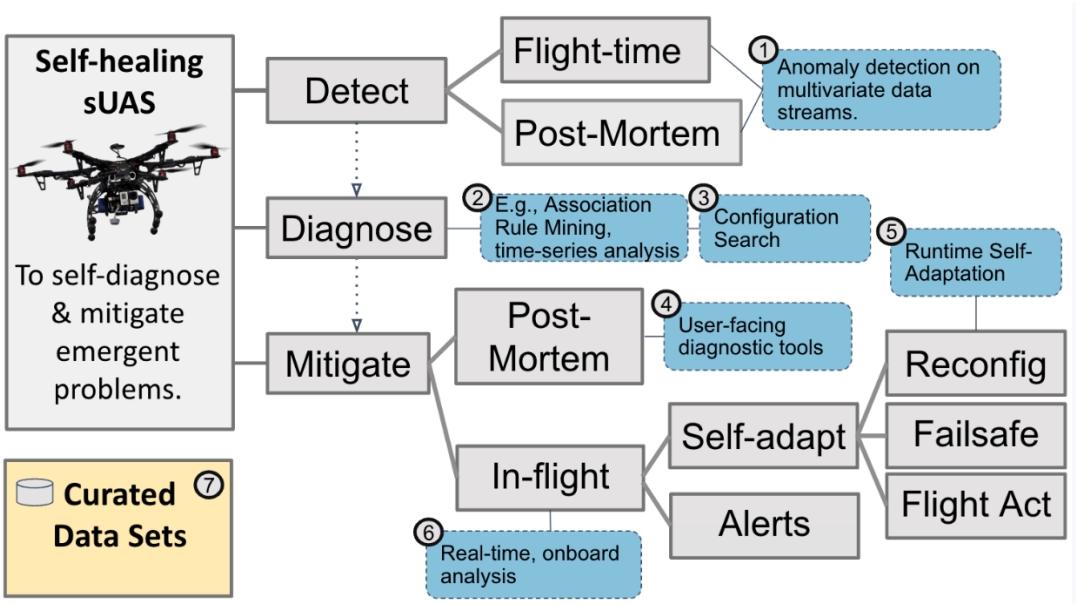

Fig. 4.2-2. Anomaly Detection, Diagnosis, and Mitigation Dataflow

4.3 New Data Collection via the Aviation Safety Reporting System (ASRS)

In conjunction with the FAA and industry, the NASA ASRS was updated to provide a unique post-flight reporting format for UAS operators. This resulted from a Safety Enhancement (SE) activity within the government-/industry-based Drone Safety Team (DST) wherein there was consensus across the team that this update could have significant impact on safety. Although the ASRS has been in place for decades and has been highly valued by the commercial airline operations sector, it has received relatively few reports from UAS operators. The new form, coupled with outreach and awareness activities, enables a larger set of reports to be provided by this community. And as the number of reports grows, the data should be more representative of the types of incidents and accidents occurring in the UAS domain. Likewise, analytics tools and capabilities can search this expanded data set at regular intervals to identify precursors, anomalies, and trends that should be addressed to improve safety.

The new form is available for submissions at https://asrs.arc.nasa.gov/uassafety.html [76]. The form was designed with inputs from a cross-section of FAA and industry organizations with the goal to provide a simple easy-to-understand easy-to-use tool for reporting events by the broad range of UAS operators and pilots (i.e., from novice to expert). It requests information on what happened, weather conditions, airspace and air traffic conditions, the UAS platform, location of the event, and contributing factors. Related, the FAA worked with NASA to provide the same protections for reporters of UAS events as for reporters of events involving traditional passenger-carrying flights (e.g., airline pilots). Most notably, confidentiality and protection against civil penalty and certificate suspension are provided in exchange for valuable safety information (see [77]).

Since the UAS reporting form went “live” in 2021, as of the date of this Technical Memorandum approximately 488 reports have been provided by and about UAS operations in

the U.S. National Airspace System. These are used in combination with the traditional ASRS reports to support analyses looking for precursors, anomalies, and trends across these types of flights. Automated SFCs that make use of these reports remain in development and are being investigated to support the IASMS goal of more timely ("in-time") safety risk mitigation to hazardous situations such as those reported via the ASRS.

5 Conclusion

This report summarizes three years of testing to explore a broad concept for future safety management systems (i.e., IASMS), as applied to two emerging flight domains where higher levels of automation are anticipated. Though the primary emphasis of the reporting herein is on the work of NASA's internal team, three additional sponsored teams significantly advanced understanding of different aspects of the concept (see Section 1.2). To fully appreciate the many and varied In-Time Aviation Safety Management advances achieved by this collaboration, please refer to the cited published reports.

Several IASMS services, functions, and capabilities were tested in select pre-flight, in-flight, and post-flight/offline scenarios. The scenarios were chosen to test the ability and effectiveness of the developed technologies to monitor, assess, and mitigate key safety risks anticipated for low-altitude flight operations using emerging sUAS and air taxi vehicles.

In terms of historical context, this report should be viewed as (a) an update to preliminary work published in [1], (b) a guidepost to a significant body of work by an extended team, and (c) a stepping-stone to research-informed guidelines and recommended practices for future operational systems aimed at more timely proactive safety risk mitigation.

Three areas warrant further effort. Most immediately, adoption of the services, functions, and/or capabilities described will require some work to customize and operationalize for a particular aviation activity. Second, the aviation community would benefit from standardization in some areas (e.g., performance requirements, service information exchange protocols, function allocation, and contingency management policies). Lastly, these systems and concepts are likely applicable to additional flight domains and hazards; exploring the extent of this applicability would benefit from the described architectural approaches, testing methods/capabilities, and lessons learned.

References

[1] Young, Steven, Ersin Ancel, Andrew Moore, Evan Dill, Cuong Quach, John Foster, Kaveh Darafsheh et al. "Architecture and information requirements to assess and predict flight safety risks during highly autonomous urban flight operations." NASA Technical Report, NASA-TM/2020-220440, January 2020.

[2] "ASTM F3269-21, Standard practice for methods to safely bound behavior of aircraft systems containing complex functions using run-time assurance," ASTM International, https://www.astm.org/f3269-21.html. Accessed November 2023.

[3] "Aeronautics Research Strategic Implementation Plan (2019 Update)," National Aeronautics and Space Administration, www.nasa.gov/aeroresearch/strategy/. Accessed April 2022.

[4] National Academies of Sciences, Engineering, and Medicine, and Aviation Safety Assurance Committee. In-time Aviation Safety Management: Challenges and Research for an Evolving Aviation System. National Academies of Sciences, Engineering, and Medicine. 2018.

[5] Young, Steven D., Cuong Quach, Kai Goebel, and Jessica Nowinski. "In-time safety assurance systems for emerging autonomous flight operations." In 2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC), pp. 1-10. IEEE, 2018.
[6] Ellis, Kyle K., Paul Krois, John Koelling, Lawrence J. Prinzel, Misty Davies, and Robert Mah. "A Concept of Operations (ConOps) of an in-time aviation safety management system (IASMS) for Advanced Air Mobility (AAM)." In AIAA SciTech 2021 Forum, p. 1978. 2021.
[7] Ancel, Ersin, Steven D. Young, Cuong C. Quach, Rafia F. Haq, Kaveh Darafsheh, Kyle M. Smalling, Sixto L. Vazquez et al. "Design and Testing of an Approach to Automated In-Flight Safety Risk Management for sUAS Operations." In AIAA AVIATION 2022 Forum, p. 3459. 2022.
[8] Young, Steven D., Ersin Ancel, Evan T. Dill, Andrew Moore, Cuong C. Quach, Kyle M. Smalling, and Kyle K. Ellis. "Flight Testing In-Time Safety Assurance Technologies for UAS Operations." In AIAA AVIATION 2022 Forum, p. 3458. 2022.
[9] “PART 107—SMALL UNMANNED AIRCRAFT SYSTEMS,” Electronic Code of Federal Regulations, https://www.ecfr.gov/current/title-14/chapter-I/subchapter-F/part-107. Accessed November 2023.
[10] Goodrich, Kenneth H., and Colin R. Theodore. "Description of the NASA Urban Air Mobility Maturity Levels (UML)." In AIAA AVIATION Forum 2020. 2020.
[11] Al Islam, Md Nafee, Muhammed Tawfiq Chowdhury, Pedro Alarcon, Jane Cleland-Huang, and Lilly Spirkovska. "Towards an Annotated All-Weather Dataset of Flight Logs for Small Uncrewed Aerial System." In AIAA AVIATION 2023 Forum, p. 3856. 2023.
[12] Cleland-Huang, Jane, Nitesh Chawla, Myra Cohen, Md Nafee Al Islam, Urjoshi Sinha, Lilly Spirkovska, Yihong Ma, Sulil Purandare, and Muhammed Tawfiq Chowdhury. "Towards Real-Time Safety Analysis of Small Unmanned Aerial Systems in the National Airspace." In AIAA AVIATION 2022 Forum, p. 3540. 2022.
[13] Al Islam, Md Nafee, Muhammed Tawfiq Chowdhury, Ankit Agrawal, Michael Murphy, Raj Mehta, Daria Kudriavtseva, Jane Cleland-Huang, Michael Vierhauser, and Marsha Chechik. "Configuring mission-specific behavior in a product line of collaborating Small Unmanned Aerial Systems." Journal of Systems and Software 197 (2023): 111543.
[14] “Drone Response, NASA.” Software and Requirements Engineering Center (SAREC) at Notre Dame. http://sarec.nd.edu/pages/DR_NASA.html. Accessed Sept 2023.
[15] “Making the skies safer with smarter drones.” University of Notre Dame College of Engineering. March 2023. https://engineering.nd.edu/news/making-the-skies-safer-with-smarter-drones/
[16] “Drone Response.” Drone Response, Inc. https://droneresponse.ai/. Accessed September 2023.
[17] Agrawal, Ankit, Sophia J. Abraham, Benjamin Burger, Chichi Christine, Luke Fraser, John M. Hoeksema, Sarah Hwang et al. "The next generation of human-drone partnerships: Co-designing an emergency response system." In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. 2020.
[18] “Datasets for Small Uncrewed Aerial Systems.” Software and Requirements Engineering Center (SAREC) at Notre Dame. https://github.com/SAREC-Lab/sUAS-UseCases/blob/master/README.md. Accessed September 2023.
[19] Elks, Carl, Patrick Martin, Robert H. Klenke, Smitha Gautham, Brandon Simon, Alexander Will, Peter Truslow, and Evan T. Dill. "Pervasive Runtime Monitoring for

Detection and Assessment of Emerging Hazards for Advanced UAM Systems." In AIAA AVIATION 2022 Forum, p. 3541. 2022.
[20] Baheri, Ali, Hao Ren, Benjamin Johnson, Pouria Razzaghi, and Peng Wei. "A verification framework for certifying learning-based safety-critical aviation systems." In AIAA AVIATION 2022 Forum, p. 3965. 2022.
[21] Guo, Wei, Marc Brittain, and Peng Wei. "Safety Validation for Deep Reinforcement Learning Based Aircraft Separation Assurance with Adaptive Stress Testing." In 2023 IEEE/AIAA 42nd Digital Avionics Systems Conference (DASC), pp. 1-10. IEEE, 2023.
[22] Taye, Abenezer, Ellis L. Thompson, Peng Wei, Timothy Bonin, and James C. Jones. "Probabilistic Evaluation for Flight Mission Feasibility of a Small Octocopter in the Presence of Wind." In AIAA AVIATION 2023 Forum, p. 3964. 2023.
[23] Ancel, Ersin, Tyler Helsel, and Christina M. Heinich. "Ground Risk Assessment Service Provider (GRASP) Development Effort as a Supplemental Data Service Provider (SDSP) for Urban Unmanned Aircraft System (UAS) Operations." In 2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC), pp. 1-8. IEEE, 2019.
[24] Ancel, Ersin, Francisco M. Capristan, John V. Foster, and Ryan C. Condotta. "In-time non-participant casualty risk assessment to support onboard decision making for autonomous unmanned aircraft." In AIAA AVIATION 2019 Forum, p. 3053. 2019.
[25] Moore, Andrew, Julian Gutierrez, Evan Dill, Michael Logan, J. Sloan Glover, Steven Young, and Nathan Hoege. "Accuracy Assessment of Two GPS Fidelity Prediction Services in Urban Terrain." In 2023 IEEE/ION Position, Location and Navigation Symposium (PLANS), pp. 25-33. IEEE, 2023.
[26] Moore, Andrew, Nicholas Rymer, J. Sloan Glover, and Derin Ozturk. "Predicting GPS Fidelity in Heavily Forested Areas." In 2023 IEEE/ION Position, Location and Navigation Symposium (PLANS). IEEE, 2023.
[27] Moore, Andrew, Matthew Schubert, Nick H. Rymer, Daniel Villalobos, J. Sloan Glover, Derin Ozturk, and Evan T. Dill. "Volume raycasting of GNSS signals through ground structure lidar for UAV navigational guidance and safety estimation." In AIAA SciTech 2022 Forum. 2022.
[28] Dill, Evan, Julian Gutierrez, Steven Young, Andrew Moore, Arthur Scholz, Emily Bates, Ken Schmitt, and Jonathan Doughty. "A predictive GNSS Performance Monitor for Autonomous Air Vehicles in Urban Environments." In Proceedings of the 34th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2021), pp. 125-137. 2021.
[29] Banerjee, Portia, Matteo Corbetta, and Katelyn Jarvis. "Probability of Obstacle Collision for UAVs in presence of Wind." In AIAA AVIATION 2022 Forum. 2022.
[30] Corbetta, Matteo, Portia Banerjee, Wendy Okolo, George Gorospe, and Dmitry G. Luchinsky. "Real-time uav trajectory prediction for safety monitoring in low-altitude airspace." In AIAA AVIATION 2019 Forum, p. 3514. 2019.
[31] Vazquez, Sixto, Cuong Quach, Laura J. Smith, Truong X. Nguyen, Kyle Smalling, Derek Dietz. Monitoring Spectrum Occupancy to Assess Communications Link Health for UAS Flight Operations, in preparation, 2024.
[32] Vazquez, Sixto, Kaveh Darafsheh, Anne Mackenzie, Christopher M. Morris, Truong X. Nguyen, Kyle Smalling, Haylee Winters, Laura J. Smith, Andrew Turner, Cuong Quach, “Ground Based RFI Mitigation Approaches for UAS Flight Operations in Urban Environment,” in preparation, 2024.

[33] Mackenzie, Anne I. “Simulation of Radio Frequency Power Received by a UAV Along Its Flight Path,” NASA Technical Report, NASA-TM/2020-220440, January 2020.
[34] Mackenzie, Anne I., Truong X. Nguyen, Kevin N. Barnes, and Michael J. Scherner. “Simulated Versus Measured UHF Radiated Power at LaRC Street Level.” NASA/TM-20210013542. 2021.
[35] Turner, Andrew. “User’s Guide for Creating a Wind Model for UAS Missions using QUIC-URB,” in preparation, 2024.
[36] Young, Steven D. "SWS Project Relevant Research." In NASA/FAA/Agility Prime Weather Technical Interchange Meeting. 2022. ntrs.nasa.gov/citations/20220018489. Accessed September 2023.
[37] Jones, James C., Timothy Bonin, and Erin Mitchell. "Evaluating Wind Hazards for Advanced Air Mobility Operations." In AIAA AVIATION 2023 Forum, p. 4104. 2023.
[38] Taye, Abenezer, Ellis L. Thompson, Peng Wei, Timothy Bonin, and James C. Jones. "Probabilistic Evaluation for Flight Mission Feasibility of a Small Octocopter in the Presence of Wind." In AIAA AVIATION 2023 Forum, p. 3964. 2023.
[39] Spirkovska, Lilly, Chetan S. Kulkarni, Jason Watkins, and Lynne Martin. "Urban Air Mobility Airspace Dynamic Density." In AIAA AVIATION 2022 Forum, p. 3403. 2022.
[40] Paez, Gina Katherine Sierra, Marcos Orchard, Chetan Kulkarni, and Kai Goebel. "A hybrid battery model for prognostics in small-size electric UAVs." In Annual Conference of the PHM Society, vol. 10, no. 1. 2018.
[41] Hartman, David, Christopher L. Hartman, and John V. Foster. "Performance Modeling of Urban Air Mobility Vehicles to Support Air Traffic Management Research." In AIAA SCITECH 2023 Forum, p. 0548. 2023.
[42] “ArduPilot Copter,” ArduPilot Dev Team. https://ardupilot.org/copter/index.html.
[43] Feldman, Jolene, Lynne Martin, Julia Bradley, Charles Walter, and Vimmy Gujral. "Developing a dashboard interface to display assessment of hazards and risks to sUAS flights." In AIAA AVIATION 2022 Forum, p. 3462. 2022.
[44] Feldman, Jolene, Lynne Martin, Vimmy Gujral, Charles Walter, Dorrit Billman, Patricia Revolinsky, and Gregory Costedoat. "Developing and testing two interfaces for Supplemental Data Service Provider (SDSP) tools to support UAS Traffic Management (UTM)." In AIAA AVIATION 2023 Forum, p. 3966. 2023.
[45] Ho, Nhut, Walter Johnson, Karanvir Panesar, Kenny Wakeland, Garrett Sadler, Nathan Wilson, Bao Nguyen, Joel Lachter, and Summer Brandt. "Application of human-autonomy teaming to an advanced ground station for reduced crew operations." In 2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC), IEEE, 2017.
[46] Turner, Andrew, Cuong Quach, Yamira Santiago-Espada, Kyle Smalling, Sixto Vazquez, Laura J. Smith, “Groundwatch User Guide,” in preparation, 2024.
[47] Daigle, Matthew, and Chetan S. Kulkarni. "Electrochemistry-based battery modeling for prognostics." In Annual Conference of the PHM Society, vol. 5, no. 1. 2013.
[48] Eure, Kenneth, and Edward Hogge. “Mathematical Characterization of Battery Models.” NASA/TM-20205008059. 2021.
[49] Dill, Evan T., Russell V. Gilabert, and Seth S. Young. "Safeguard – Flight Test Results of an On-board System Designed to Assure Conformance to Geospatial Limitations." In 2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC), pp. 1-8. IEEE, 2018.

[50] Dill, Evan T., Kelly J. Hayhurst, Steven D. Young, and Anthony J. Narkawicz. "UAS hazard mitigation through assured compliance with conformance criteria." In 2018 AIAA Information Systems-AIAA Infotech@ Aerospace. 2018.
[51] Oza, Nikunj, Kevin Bradner, David L. Iverson, Adwait Sahasrabhojanee, and Shawn R. Wolfe. "Anomaly Detection, Active Learning, Precursor Identification, and Human Knowledge for Autonomous System Safety." In AIAA SciTech 2021 Forum. 2021.
[52] Iverson, David L. "Inductive system health monitoring." In International Conference on Artificial Intelligence. 2004.
[53] Consiglio, María, César Muñoz, George Hagen, Anthony Narkawicz, and Swee Balachandran. "ICAROUS: Integrated configurable algorithms for reliable operations of unmanned systems." In 2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC), pp. 1-5. IEEE, 2016.
[54] Peters, Andrew, Swee Balachandran, Brendan Duffy, Kyle Smalling, María Consiglio, and César Muñoz. "Flight test results of a distributed merging algorithm for autonomous UAS operations." In 2020 AIAA/IEEE 39th Digital Avionics Systems Conference (DASC), pp. 1-7. IEEE, 2020.
[55] Federal Aviation Administration, “UTM Concepts of Operations,” Version 2.0, FAA, March 2020.
[56] Federal Aviation Administration, “Urban Air Mobility Concept of Operations,” Version 1.0, FAA, June 2020.
[57] Shoemaker, Paul, and Collin Estes. "Accessible Telemetry Streams using a Zero Trust Architecture for the Flight Operations Directorate." 2021. https://ntrs.nasa.gov/citations/20210010837. Accessed April 2022.
[58] “WinProp Overview.” Altair Engineering. https://help.altair.com/winprop/topics/winprop/user_guide/introduction/overview_winprop_c.htm. Accessed November 2023.
[59] Nguyen, Truong X. "High-Intensity Radiated Field (HIRF) Map-An Avoidance Approach for UAM, AAM, and UAS Vehicles." In 2023 IEEE/AIAA 42nd Digital Avionics Systems Conference (DASC). IEEE, 2023.
[60] “Mission Planner Home.” ArduPilot Dev Team. https://ardupilot.org/planner/index.html.
[61] “RFD900+ Modem.” RFDesign Pty Ltd. https://rfdesign.com.au/products/rfd900-modem. Accessed November 2023.
[62] Koubâa, Anis, Azza Allouch, Maram Alajlan, Yasir Javed, Abdelfettah Belghith, and Mohamed Khalgui. "Micro Air Vehicle Link (MAVlink) in a nutshell: A survey." IEEE Access 7 (2019): 87658-87680.
[63] Foster, John V., and David Hartman. "High-fidelity multi-rotor unmanned aircraft system (UAS) simulation development for trajectory prediction under off-nominal flight dynamics." In 17th AIAA Aviation Technology, Integration, and Operations Conference, p. 3271. 2017. doi: 10.2514/6.2017-3271.
[64] Altamirano, George, Justin J. Matt, Ronald C. Busan, and John V. Foster. "Wind Tunnel Testing of Static Aerodynamic and Power Consumption Characteristics of an Octocopter." In AIAA AVIATION 2023 Forum, p. 3967. 2023. doi: 10.2514/6.2023-3967
[65] Ancel, Ersin, Francisco M. Capristan, John V. Foster, and Ryan C. Condotta. "Real-time risk assessment framework for unmanned aircraft system (UAS) traffic management (UTM)." In 17th AIAA Aviation Technology, Integration, and Operations Conference, p. 3273. 2017.

[66] “HATIS (Human Autonomy Teaming Interface System),” Human Automation Teaming Solutions, Inc., https://www.hats.solutions/hatis. Accessed December 2023.
[67] Neogi, Natasha A., Steven D. Young, and Evan T. Dill. "Establishing the Assurance Efficacy of Automated Risk Mitigation Strategies." In AIAA AVIATION 2022 Forum, p. 3538. 2022.
[68] Pike, Lee, Alwyn Goodloe, Robin Morisset, and Sebastian Niller. "Copilot: A hard real-time runtime monitor." In International Conference on Runtime Verification, pp. 345-359. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010.
[69] “Aviation Safety Information Analysis and Sharing Program,” Federal Aviation Administration, https://www.faa.gov/newsroom/aviation-safety-information-analysis-and-sharing-program-1. Accessed December 2023.
[70] Frank, Jeremy D., David Iverson, Christopher Knight, Sriram Narasimhan, Keith Swanson, Michael S. Scott, May Windrem, Kara M. Pohlkamp, Jeffery M. Mauldin, Kerry McGuire and Haifa Moses.Frank, J., Iverson, D., Knight, C., Narasimhan, S., Swanson, K., Scott, M., Windrem, M., Pohlkamp, K., Mauldin, J., McGuire, K. and Moses, H., “Demonstrating autonomous mission operations onboard the international space station,” In AIAA Space Forum 2015, American Institute of Aeronautics and Astronautics, 2015.
[71] McComas, David. “NASA/GSF’s Flight Software Core Flight System.” In Flight Software Workshop Flight Workshop, November 2012, (No. GSFC. CPR. 7525.2013).
[72] “core Flight System (cFS),” NASA. https://cfs.gsfc.nasa.gov. Accessed December 2023.
[73] “PX4 Autopilot Flight Review.” Dronecode Project, Inc. https://review.px4.io/browse. Accessed September 2023
[74] Ma, Yihong, Md Nafee Al Islam, Jane Cleland-Huang, and Nitesh V. Chawla. "Detecting anomalies in small unmanned aerial systems via graphical normalizing flows." IEEE Intelligent Systems (2023). p. 46-54. 2023.
[75] Al Islam, Md Nafee, Yihong Ma, Pedro Alarcon, Nitesh Chawla, and Jane Cleland-Huang. "RESAM: Requirements elicitation and specification for deep-learning anomaly models with applications to UAV flight controllers." In 2022 IEEE 30th International Requirements Engineering Conference (RE), pp. 153-165. IEEE, 2022.
[76] “UAS Safety Reporting,” NASA Aviation Safety Reporting System, https://asrs.arc.nasa.gov/uassafety.html. Accessed December 2023.
[77] “Advisory Circular 00-46F,” Federal Aviation Administration. https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_00-46F.pdf. Accessed December 2023.
[78] “DO-364, Minimum Aviation System Performance Standards for Aeronautical Information/Meteorological Data Link Services,” Radio Technical Commission for Aeronautics, 2016. https://global.ihs.com/doc_detail.cfm?document_name=RTCA%20DO-364&item_s_key=00701504. Accessed November 2023.
[79] “Truweather Solutions,” TruWeather Solutions, Inc., https://truweathersolutions.com. Accessed November 2023.
[80] “Activity Density,” AirSage Inc., https://airsage.com/solutions/population-density. Accessed November 2023.
[81] Hansen, Andrew, Stephen Mackey, Hadi Wassaf, Vaibhav Shah, Eric Wallischeck, Christopher Scarpone, Michael Barzach, and Elliott Baskerville. Complementary PNT and GPS backup technologies demonstration report: Sections 1 through 10. No. DOT-

VNTSC-20-07. John A. Volpe National Transportation Systems Center (US), 2021. https://www.transportation.gov/administrations/assistant-secretary-research-and-technology/complementary-pnt-and-gps-backup. Accessed November 2023.
[82] “Annex 15 to the Convention on International Civil Aviation – Aeronautical Information Services, 16th Edition,” International Civil Aviation Organization, 2018. https://store.icao.int/en/annex-15-aeronautical-information-services. Accessed November 2023.
[83] “Aeronautical Information Services Manual (Doc 8126), 7th Edition,” International Civil Aviation Organization, 2022. https://store.icao.int/en/aeronautical-information-services-manual-doc-8126. Accessed November 2023.
[84] Spirkovska, Lilly., and Indranil Roychoudhury. “Computing Proximity to Threat Along Uncertain Trajectory to Support Urban Air Mobility.” NASA Technical Report, NASA-TM/2024-0000892, February 2024.
[85] Banerjee, Portia, George Gorospe, and Ersin Ancel. "3D representation of UAV-obstacle collision risk under off-nominal conditions." In 2021 IEEE Aerospace Conference (50100), pp. 1-7. IEEE, 2021.
[86] Banerjee, Portia, Matteo Corbetta, Katelyn Jarvis, Kyle Smalling, and Andrew Turner. "Probability of Trajectory Deviation of Unmanned Aerial Vehicle in Presence of Wind." Journal of Air Transportation (2023): 1-12.
[87] “WindNinja,” Missoula Fire Sciences Laboratory, United States Forest Service, https://www.firelab.org/project/windninja. Accessed November 2023.
[88] Wagenbrenner, Natalie S., Jason M. Forthofer, Wesley G. Page, and Bret W. Butler. "Development and evaluation of a Reynolds-averaged Navier–Stokes solver in WindNinja for operational wildland fire applications." Atmosphere 10, no. 11 (2019): 672.
[89] Wagenbrenner, Natalie S., Jason M. Forthofer, Brian K. Lamb, Kyle S. Shannon, and Bret W. Butler. "Downscaling surface wind predictions from numerical weather prediction models in complex terrain with WindNinja." Atmospheric Chemistry and Physics 16, no. 8 (2016): 5229-5241. doi:10.5194/acp-16-5229-2016.
[90] “What is QUIC?” Los Alamos National Laboratory. https://www.lanl.gov/projects/quic/. Accessed November 2023.
[91] Skamarock, William C., Joseph B. Klemp, Jimy Dudhia, David O. Gill, Zhiquan Liu, Judith Berner, Wei Wang et al. "A description of the advanced research WRF version 4." NCAR tech. note ncar/tn-556+ str 145 (2019).
[92] Dowell, David C., Curtis R. Alexander, Eric P. James, Stephen S. Weygandt, Stanley G. Benjamin, Geoffrey S. Manikin, Benjamin T. Blake et al. "The High-Resolution Rapid Refresh (HRRR): An hourly updating convection-allowing forecast model. Part I: Motivation and system description." Weather and Forecasting 37, no. 8 (2022): 1371-1395.
[93] “Observation Networks,” National Weather Service, https://www.weather.gov/lmk/observation_networks. Accessed November 2023.
[94] “The NOAA CORS Network (NCN),” NOAA/National Geodetic Survey, https://geodesy.noaa.gov/CORS. Accessed November 2023.
[95] Spirkovska, Lilly, “Vertiport Dynamic Density,” NASA Technical Report, NASA-2023-0012622, September 2023.
[96] Hogge, Edward F., Brian M. Bole, Sixto L. Vazquez, Chetan S. Kulkarni, Thomas H. Strom, Boyd L. Hill, Kyle M. Smalling, and Cuong C. Quach. "Verification of prognostic

algorithms to predict remaining flying time for electric unmanned vehicles." International Journal of Prognostics and Health Management 9, no. 1 (2018).
[97] Eure, Kenneth W., and Edward F. Hogge. Exploration of an Adaptive Routine for Battery Modeling. NASA/TM-20220000668. 2022.
[98] Corbetta, Matteo., “Li-ion Battery Prognosis Based on Hybrid Bayesian PINN,” [Source code]. 2023. https://github.com/nasa/Li-ion-Battery-Prognosis-Based-on-Hybrid-Bayesian-PINN.
[99] Nascimento, Renato G., Felipe AC Viana, Matteo Corbetta, and Chetan S. Kulkarni. "A framework for Li-ion battery prognosis based on hybrid Bayesian physics-informed neural networks." Scientific Reports 13, no. 1 (2023): 13856.

Appendix A. Service Descriptions

As discussed in Section 1, the IASMS ConOps envision systems that provide a set of capabilities that proactively monitor, assess, and mitigate safety risks, and these capabilities are tailored to address hazards associated with specific application domains. Capabilities are enabled by connectivity across a set of relevant information services and/or by functions embedded in the aircraft and operations/ground stations. This appendix provides descriptions for the information services used for the R&D described in Section 1 and listed in Table 1.

In general, an information service may be configured to operate in a request-reply, publish-subscribe, or broadcast mode. Except where noted, services described in this section were tested in a request-reply mode. For this mode, information flow is consistent with [78] and is illustrated below in Figure A-1 as a sequence diagram. The telemetry service was enabled as part of the COTS autopilot system [42] and is not described here.

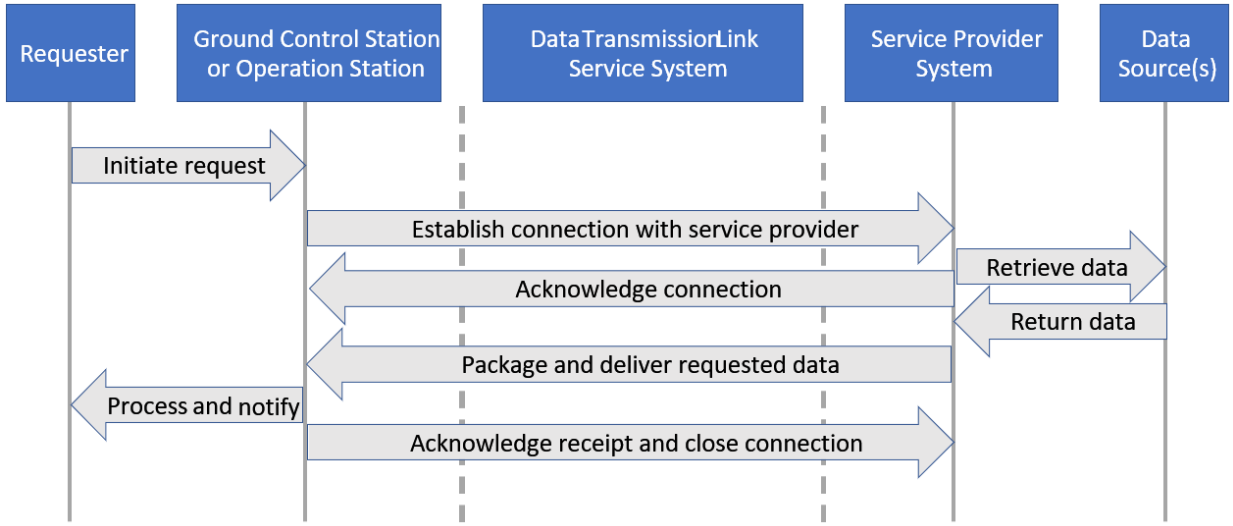

Fig. A-1. Sequence diagram representing service configuration information flow for developmental services (i.e., request-reply).

A brief description with estimated TRL status[11] is provided for the following services along with reference documents where available:

1. 3rd party casualty risk assessment (also in [1])
2. Navigation quality assessment
3. Proximity to threat assessment (also in [1])
4. Obstacle collision risk assessment
5. RF environment and interference monitoring
6. Weather/wind monitoring and forecasting (3 types)
7. Airspace dynamic density assessment
8. Battery prognostics and health management (also in [1])
9. Flight performance assessment

[11] Technology Readiness Level (TRL) estimates follow definitions in https://www.nasa.gov/wp-content/uploads/2017/12/458490main_trl_definitions.pdf

A.1 Third Party Casualty Risk Assessment (*GRASP*)

Purpose: As described in [1][23][65], this service helps to mitigate the risk of third-party casualties should the vehicle experience loss-of-control or a critical system failure (e.g., power loss). Service functionality is based on the UTM Risk Assessment Framework (URAF) [65] which includes various models (e.g., population density model, off-nominal trajectory and impact point prediction model, probability of casualty estimation model, and mishap likelihood model). The impact point prediction model takes wind into account as this can significantly affect sUAS trajectories (e.g., following a power loss). Wind speed may be inputted by the operator/client or may be read from an external service (i.e., a wind service). At least one such service is currently available ([79]) and others are anticipated. Population density data may be provided by the operator/client or may be read from an external service. The current implementation reads population data from an externally provided service ([80]) with coverage across the continental US, including 20 metropolitan areas.

During pre-flight planning, the service assists operators in understanding and minimizing the associated casualty risk of alternative flight plans [23]. It then selects a flight plan that meets the operation's third-party risk requirement. In addition to assessing 3rd party risk along the route, this service may also be useful to regulatory authorities when evaluating operational risk assessments and proposed mitigations as part of operational approval processes. Further, by using downlinked flight parameters from the vehicle [65], the service may also be utilized during flights to provide the operator with contemporaneous risk estimates; and/or to support automation that may be responsible for contingency management and execution (e.g., avoiding densely populated areas when diverting). [Note: Ongoing and planned improvements to the service include estimating separation risk, adding aircraft/equipage reliability models, and predictive population density models.]

Inputs: The service simulates a flight using a client-provided flight plan (i.e., where the flight will go and when it will occur). To reduce uncertainty in the output product(s), additional parameters may be provided by the client in the request. Each product requires some combination of:

1. Flight plan (i.e., list of waypoint coordinates and speeds or estimated time of arrivals (ETAs) at waypoints)
2. Date/time period of planned flight
3. Simulation resolution (e.g., estimate risk every 10m along the flight path)
4. Aircraft information (e.g., make/model, size, weight, type, and aero model information if available)
5. Expected winds (static vector, static database, or streamed data) (optional)
6. Terrain and population density models [Note: Some urban area models are accessible via the service archive, or available via Aeronautical Information Service (AIS) providers or others (e.g., [80]).]

Outputs: Based on client requests, data products may include:

1. Casualty probability estimates (Fig. A.1-1, Fig. A.1-2) along the planned route of flight (pre-flight applications)
2. GUI with color coded impact location/area visualizations (pre-flight applications)
3. 24-hour population density distribution, controlled via a slider bar (pre-flight applications)
4. Casualty risk (mishap likelihood and casualty probability estimates) (during flight)
5. Mishap likelihood and recommended action (e.g., abort, land, RTL, continue) (during flight)

TRL assessment: TRL = 6-8, based on performance during several high-fidelity simulation tests, flight tests, and usability/user evaluations spanning a five-year period.

For more information: [65][23][24]

Fig. A.1-1. Example visualization of population density colormap and casualty risk assessment values for flight plan over the city of Hampton, VA.

Fig. A.1-2. Closeup view of flight segment over city of Hampton, VA. In this example, colored dots indicate locations along the flight path where the casualty risk is estimated/predicted to be the higher due to elevated population presence below and near the flight, e.g., on the highway and streets.

A.2 Navigation Quality Assessment (*NavQ*)

Purpose: This service helps to mitigate the risk of navigation system performance degradation by predicting the availability of GNSS signals and dilution of precision (DOP). The service is intended primarily to support pre-flight planning of routes or launch times in urban areas to (1) avoid expected high-risk areas and/or (2) identify where GNSS augmentation or alternate types of positioning systems may be needed (e.g., [81]). Two methods were developed and deployed as services: a waypoint-based service (Corridor Assessment of Positioning Systems, or CAPS); and an area-based service (Geometric Assessment of Positioning Systems, or GAPS). [Note: Although not tested to date, the NavQ service is envisioned to also provide access to GNSS augmentation and alternate types of positioning systems that may be available in the region of interest (see [1])].

Inputs:

1) CAPS: List of flight plan waypoint coordinates
   GAPS: Bounding volume coordinates (e.g., NE/SW corners; min/max altitudes)
2) Date/time period of planned flight
3) Estimated satellite orbital positions (e.g., via GPS almanacs) [Note: User/client selects the constellation(s); service automatically retrieves the relevant almanac(s).]
4) Terrain model [Note: Some urban area models are accessible via the service archive, or available via AIS providers or others.]
5) GAPS: Desired output resolution within the coverage volume (x, y, z, t)

Outputs:

1) The expected number of unobstructed satellites (i.e., satellites within direct line-of-sight) at each waypoint in the client-provided flight plan (CAPS) or
2) The expected number of unobstructed satellites and the horizontal dilution of precision (HDOP) within the client-specified coverage volume and time window (GAPS)

TRL assessment: TRL = 6-8; based on independent in-the-field evaluations by partners in various locations (e.g., New York, Boston, Las Vegas) and on validation tests completed using data collected in Corpus Christ, TX (a representative result is shown in Fig. A-2).

For more information: [25][26][27][28]

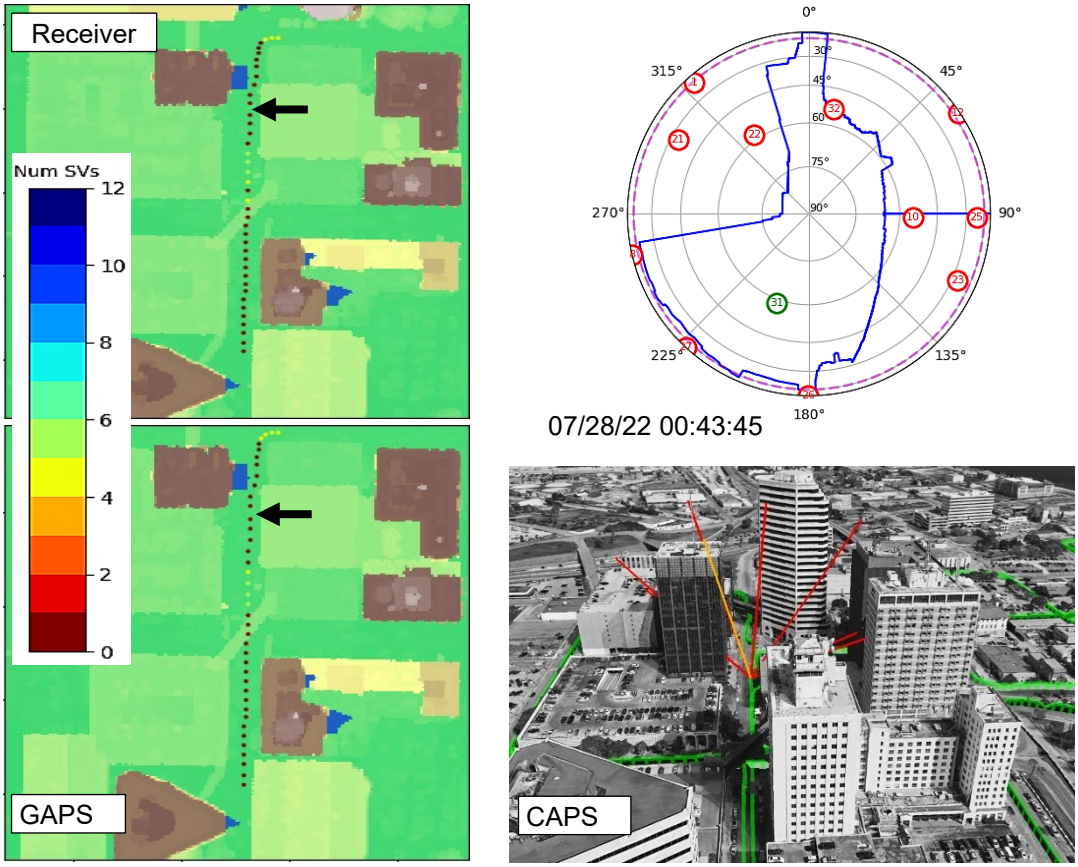

Fig. A.2. Example visualizations of NavQ predicted metrics in urban canyon from a validation test in Corpus Christi, TX on July 28, 2022. As expected from the sky plot (top right), few satellites are available to a GNSS receiver at this location due to urban blockage (blue trace), and post-processed GNSS receiver computations show complete loss of navigation (top left, arrow). Both GAPS (bottom left) and CAPS (bottom right) accurately predict the set of visible and blocked satellites. The scale at left assigns color according to the number of satellites (space vehicles or SVs) in the Receiver and GAPS images. @Map data: Google

A.3 Proximity to Threat Assessment (*PtT*)

Purpose: This service helps to reduce the risk of collision with (or entry into) geospatial threats along the flight path (e.g., buildings and no-fly zones). As described in [1], functions within the service estimate, track, and predict proximity metrics associated with hazardous areas that have been modeled and stored in geo-referenced mapping databases. The example used during development has been the perimeter of stationary vertical features (e.g., buildings/trees) but these may also represent no-fly zones or other hazardous areas defined by operators, ATM/UTM service providers, or Aeronautical Information Service (AIS) providers [82][83]. Outputs indicate portions of the vehicle trajectory that are predicted to violate client-specified proximity thresholds or safety margins to the hazardous areas (see Fig. A.3-1). The service may be used pre-flight for route planning or in flight for continuously updated route assessment. The route may be defined either as a flight plan, as a precise 4-D trajectory (position and altitude at closely spaced time steps), or as a probability distribution of uncertain 4-D trajectories together with the client-specified confidence level for representing acceptable risk.

Inputs:

1) Configuration settings (e.g., thresholds for proximity warnings, output resolution, trajectory confidence level)
2) Aircraft information (e.g., size, weight, make/model, number of rotors/engines, drag coefficient)
3) Flight plan (i.e., list of 3D waypoint coordinates and speeds at waypoints) [Note: If used during flight, the flight plan may be a predicted trajectory from current position to a client-selected look-ahead time or landing location.]
4) Geo-referenced database of extruded polygons that encapsulate hazardous vertical and horizontal features of threats [Note: Some urban area databases are accessible via the service archive, or available via AIS or other service providers.]

Outputs:

1) Points along the (possibly uncertain) flight trajectory where proximity falls below the client-defined threshold (e.g., Fig. A.3-2)
2) Nearest approach points (i.e., points along the trajectory where proximity is smallest)
3) Distance to nearest approach point
4) Severity of violation (a safety margin ranging from 0-100%, defined as in Fig. A.3-1)

TRL assessment: TRL = 5-6; Simulation and flight testing completed using moderate to complex models of hazardous areas representative of urban environments. No testing of automated interoperability with service providers for terrain, obstacle, and restricted airspaces (e.g., via ATM/USS or AIS providers).

For more information: [1][84]

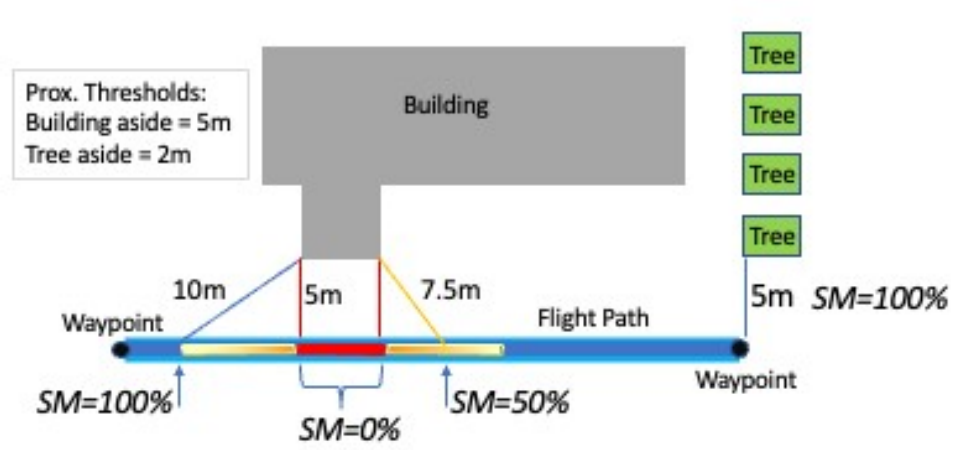

Fig. A.3-1. From [1]. Example illustrating method for computing Proximity to Threat Safety Margin (SM) metric for a flight path near trees and buildings. SM is computed using the percent error formula, SMproximity = (distance - threshold)/threshold*100. A margin of 0% means the flight path is closer to the object than the given threshold, i.e., there is no margin. A margin of 100% means the flight path is at least twice as far from the object as the given threshold. For the example shown, the threshold is set as 5m, which means that a point on the flight path that is 7.5m from a building will have a margin of 50%, or 2.5m more than the closest allowed approach distance.

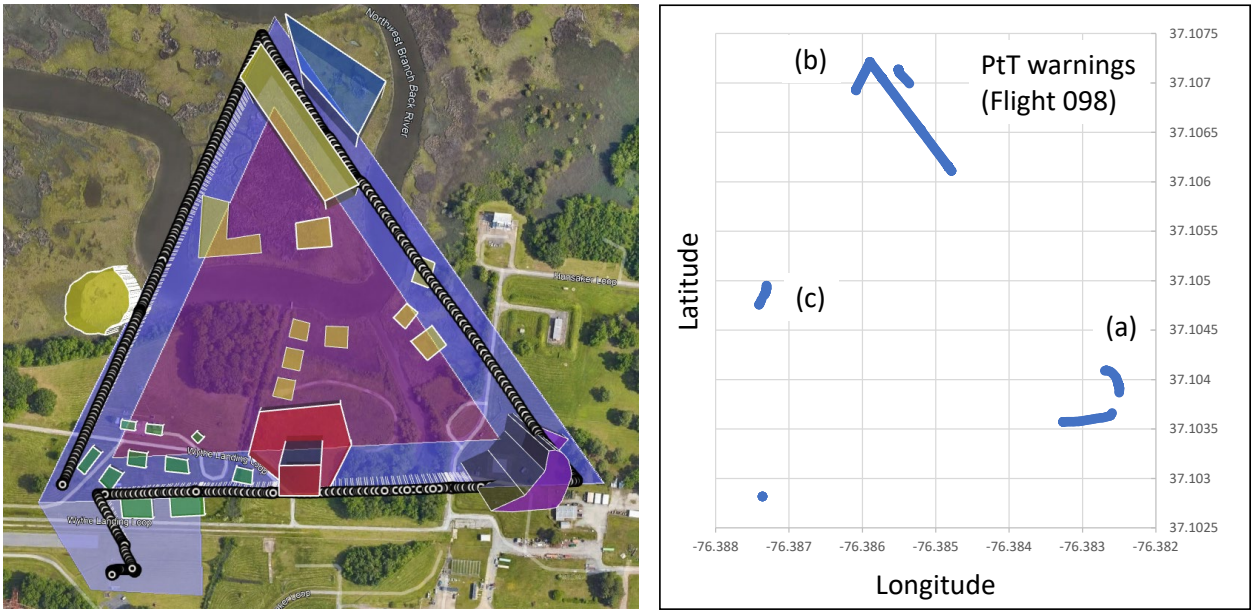

Fig. A.3-2. Example proximity to threat performance Flight 098. (Left) Nadir view of flight track as flown with threats as 3D building volumes (shown as colored shapes here). (Right) Points during the flight where PtT warnings occurred: (a) when flying east and north around tower building; (b) when flying northwest between two tall buildings; and (c) when flying southwest near the circular stadium structure.

A.4 Obstacle Collision Risk Assessment (*ROC*)

Purpose: This service helps to mitigate the risk of collision with vertical structures along or near the flight path. It applies a function (based on PtT) to monitor and track proximity metrics to stationary obstacles (e.g., buildings, trees, and utility poles) and then considers a variety of factors to quantify two risk-related metrics. The service applies a detailed formulation of obstacle collision risk likelihood that incorporates the effects of component failures, reduced vehicle controllability, and environmental disturbances such as wind gusts. The trajectory deviation due to wind is estimated and predicted using a UAS trajectory simulation function. Obstacle collision risk for the flight plan is estimated based on generating the probability of trajectory deviation at each point along the predicted trajectory, then checking whether these predictions coincide with obstacle locations. When used during flight planning, the service helps users/clients to see (1) how wind speed magnitudes and directions will affect the probability of deviation; (2) how risk of collision changes with the planned speed profile; and (3) how deviations (and collision risk) will be higher at turning waypoints versus straight flight path segments. The service may also be employed in-flight to re-assess collision risk and trajectory deviations when/if a contingency plan is considered that re-directs the aircraft to an alternate destination.

Inputs:

1) Configuration settings (e.g., threshold for trajectory deviation)
2) Aircraft information (e.g., size, weight, make/model, number of rotors/engines, drag coefficient)
3) Flight plan (i.e., list of waypoint coordinates and speeds or ETAs at waypoints) [Note: If used during flight, the flight plan may be a predicted trajectory from current position to a client-selected look-ahead time or landing location.]
4) Wind magnitude and direction along the flight path (as mean and standard deviation if available)
5) Geo-referenced database of obstacle locations defined as 2D polygons with height values and inclusive of features where collision risk is of concern). [Note: Some obstacle databases are available via AIS providers or others.]
6) Obstacle location data accuracy or measurement noise (if available)
7) Cross-sectional area of aircraft exposed to obstacle (if available)
8) Date/time period of planned flight (if using a wind forecast service)
9) Desired simulation resolution (e.g., estimate risk every 10m along the flight path)

Outputs:

1) Probability of trajectory deviation (can be computed without obstacle location data)
2) Probability of obstacle collision

TRL assessment: TRL = 2-3. Initial validation using log files from experimental flights under varying wind conditions compared with simulation results. See Figs. A.4-1 and A.4-2.

For more information: [29][85]][86]

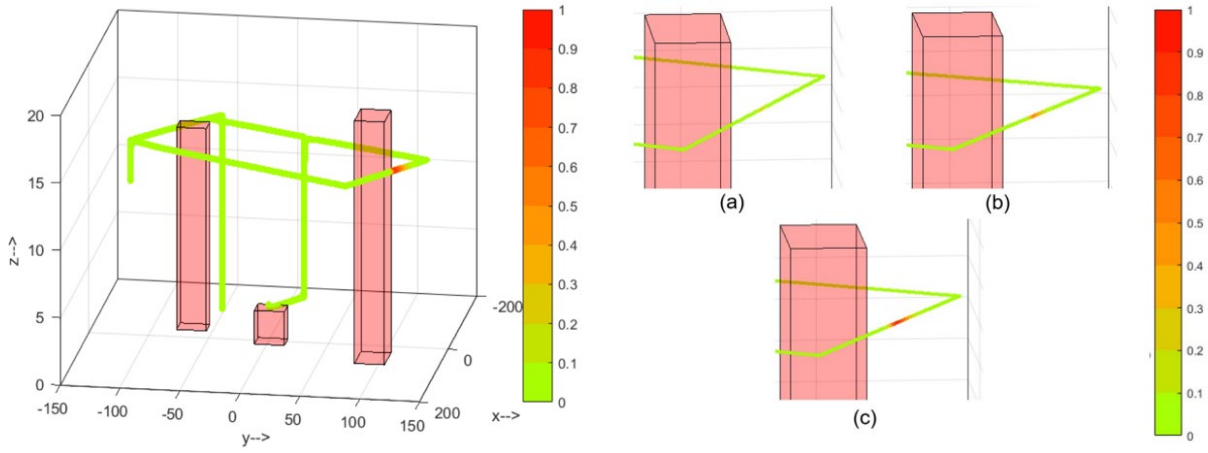

Fig. A.4-1. From [85] Example simulation results for obstacle collision probability: (green) 0%; red 100% (Left) |V_wind | = N (8,2)m/s, χ_wind = -120° and γ_wind = 2° (Right) (a) Wind: |V_wind | = N (2, 2) m/s, (b) Wind: |V_wind | = N (5, 2) m/s, (c) Wind: |V_wind | = N (8, 2) m/s

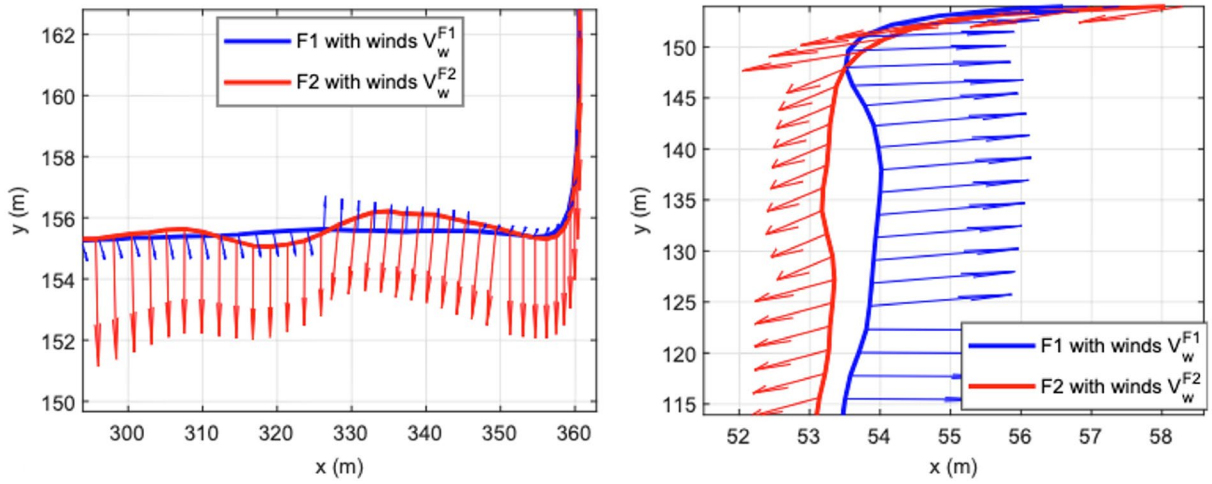

Fig. A.4-2. From [86] Example experimental results - Effect of wind on probability of trajectory deviation. Flight path with arrows depicting wind showing (a) higher deviations in Flight 2 (F2) due to higher winds and (b) trajectory deviations in Flight 1 (F1) and F2 due to winds in opposing directions.

A.5 RF Environment and Interference Monitoring (*RF/RFE/RFI*)

Purpose: This service helps mitigate hazards associated with the RF environment that may be encountered during operations, particularly when flying at low altitudes in urban areas and using RF bands not dedicated to aviation use or when using aviation bands which employ multi-access protocols. Service products are tailored to user-provided domain-specific information (e.g., planned route, frequencies of interest, and equipage). Several RF-related metrics and imagery can inform mitigation decisions during ops planning, pre-flight, and in-flight phases.

Inputs: Various products can be delivered based upon information provided in a client request. Each product requires some combination of:

1) Flight plan or coverage area
2) Frequencies of interest
3) Link equipage sensitivities
4) Date/time period of interest (for historical trend)
5) Frequency hop range

Outputs: Based on the client/user request, data products may include:

1) Pre-computed RF Environment Maps (REM) for requested frequencies and elevations
2) Along-path expected received power for selected communication link bands and interference power ratios
3) Average measured power for requested frequency over date range by hour or by day
4) Recommended stay out regions around FCC-registered transmitters for given frequencies
5) Spectrum occupancy per sensor node (within coverage of spectrum monitoring stations)

Example outputs for (1)-(4) are shown in Fig. A.5. An example for (5) is shown in Section 2.1.

TRL assessment: TRL = 3 for output product (2) with developmental testing ongoing; TRL = 5 for (1) with limited testing of two frequencies to date; TRL = 7 for (3,4) by leveraging COTS spectrum monitors but testing within limited coverage area; TRL = 5 for (5) and with ongoing work to address unregistered transmitters.

For more information: [8][31][32][33][34][59]

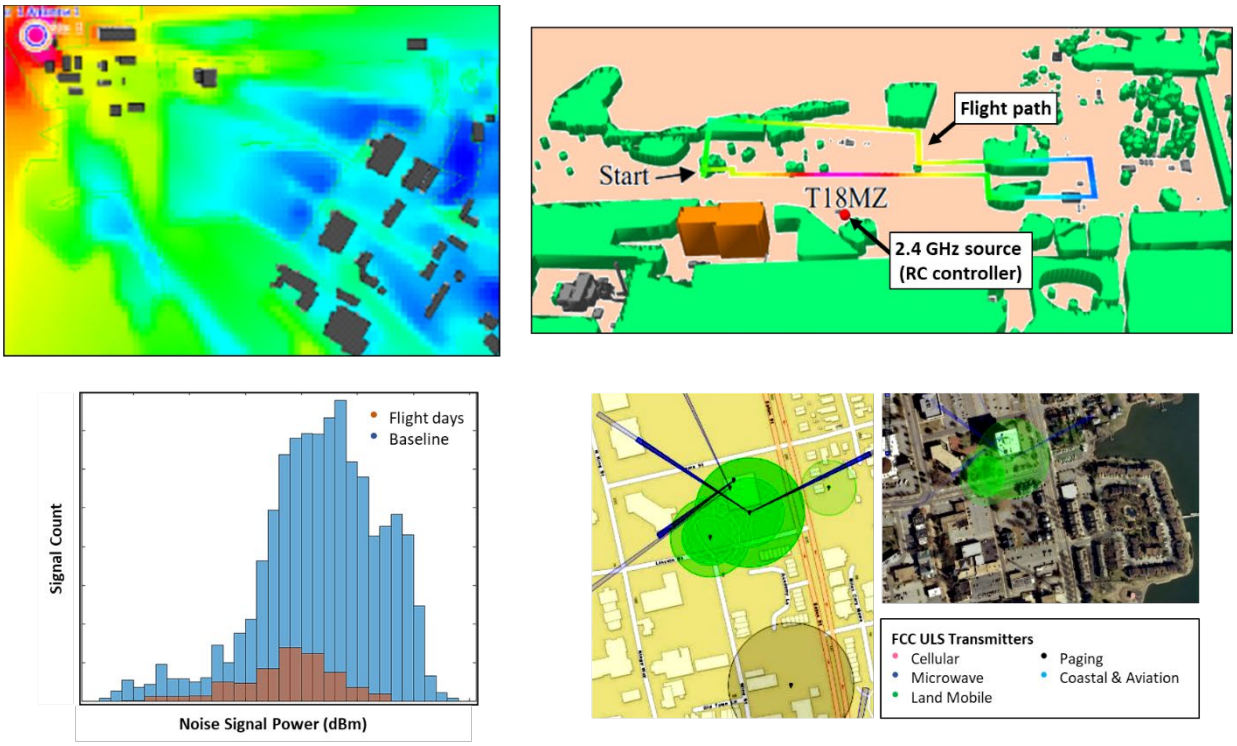

Fig. A.5. RF Environment and Interference Monitoring Service Data Products. Top left: Example RF environment map in the 900MHz telemetry band at 1m altitude. The radiating source is a navigational signal tower at top left. Color scale from low (blue) -50 to high (red/pink) -30 dBm. Top right: Along-path expected received power from an interference source at the 2.4 GHz control frequency. Color scale from low (blue) -90 to high (red/pink) -72 dBm. The RF power is highest near the interfering RC transmitter ('T18MZ'). Bottom left: Average measured power per time-of-day window. Bottom right: Recommended stay out regions around FCC-registered transmitters.

A.6 Weather/Wind Monitoring and Forecasting (*WM/WIPS/WxSt*)

There were three services developed and tested within this topic by the project. A fourth was developed with support by NASA's SBIR program. Because this is a commercial product, it is not described here. However, additional information can be found at [79]. The three services are denoted here as the Wind modeling service (WM), the Wind information prediction service (WIPS), and the Weather station observations service (WxSt).

A.6.1 Wind Modeling Service (WM)

Purpose: This service helps mitigate hazards associated with wind conditions that may impact controllability, stability, and deviations from a safe flight envelope or intended route of flight. The service is designed specifically to model wind affects at low altitudes near vertical features (e.g., structures and treelines). Air movements around such features can be estimated to 1 meter resolution.

Inputs: Client requests must include some combination of:

1) Prevailing wind speed (at altitude just above tallest feature)
2) Prevailing wind direction (at altitude just above tallest feature)
3) Wind tolerance (wind speed threshold of interest (e.g., max wind speed for safe flight))
4) Altitude range – Requested upper/lower altitudes for wind contours (e.g., 0-100m AGL)
5) Altitude resolution – Requested vertical separation of wind contours (e.g., every 5m)
6) Horizontal resolution – Requested contour resolution in the x-y plane (default is 1m)
7) Location or coverage area – Area of interest (e.g., NW and SE corner coordinates)
8) Terrain or feature model for area of interest (if not available within the service)

Outputs: Data products can include:

1) Estimate wind values within requested volume (wind speed and direction, relative or absolute, Fig. A.6.1-1)
2) Estimated wind value contours at the specified resolution and altitude values
3) Regions where windspeed values exceed specified wind tolerance value (as polygons, see Fig. A.6.1-2)

TRL assessment: TRL = 5-6; This service is an application of WindNinja [87][88][89] and QUIC-URB [90], which have been used by the U.S. National Forest Service and Los Alamos National Labs, respectively, for several years to predict winds at small scales in complex terrain. Validation of this application is limited to model-to-model comparisons, SME subjective evaluations, and measurement-based validation at the NASA LaRC test range. Additional validation is planned for a more complex urban environment.

For more information: [35][87][88][89][90]

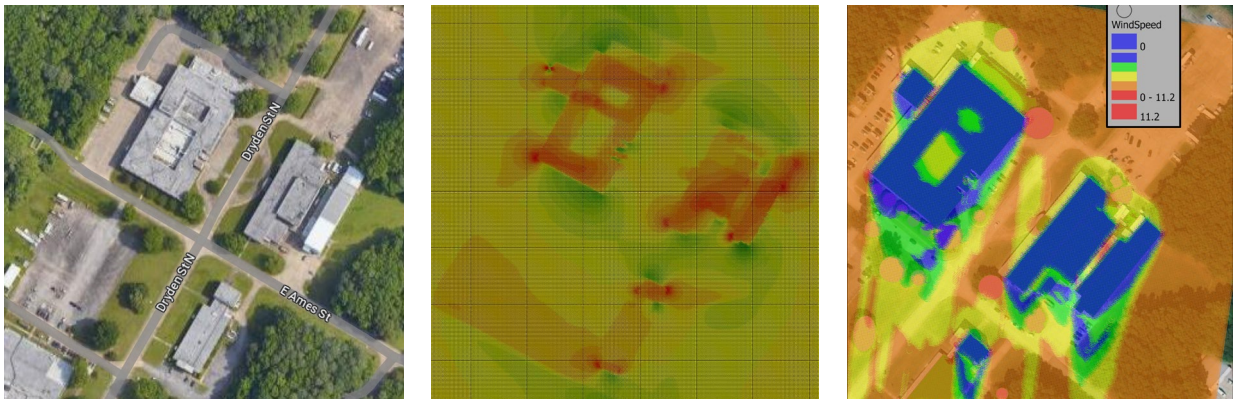

Fig. A.6.1-1. Example comparison of wind model visualizations (center, right) to geographic features (left). Input is a prevailing wind value of 15 m/s at 0 deg (north to south) measured at 10m AGL. Output shown as colormap of relative wind speeds (center) and absolute wind speeds (right). Altitude slice shown is at 10m AGL.

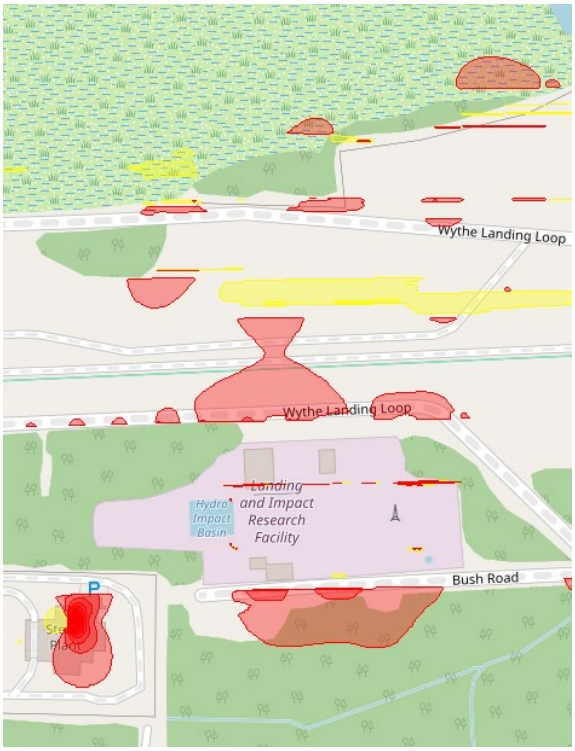

Fig. A.6.1-2. Example image of wind modeling service data product
(Inputs: prevailing wind value: 7m/s at 0 deg; model elevation: 30m AGL; wind tolerance: 6 m/s)
(Red: exceeds threshold; Yellow: at threshold)

A.6.2 Wind Information Prediction Service (WIPS)

Purpose: This service helps mitigate hazards associated with wind conditions that may impact controllability, stability, and/or deviations from a safe flight envelope or intended route of flight. The service generates high-resolution low altitude wind forecasts and translates these forecasts into probabilistic representations of wind hazard regions. These products may be used during flight simulations to evaluate various ATM and vehicle systems that can be affected by wind. Hazard regions are identified based on 3D gridded wind fields generated using the open-sourced Weather Research and Forecasting (WRF) model [91]. WRF simulates weather conditions based on initial and boundary conditions and produces weather data, including temperature, wind, humidity, precipitation, and other diagnostic fields, for a volumetric grid out to a client-selected forecast horizon. For WIPS, WRF was customized to run at large eddy simulation scales with a horizontal grid spacing of 100 meters using high-resolution terrain and land use characteristics to capture local variability in the flow not resolved by the operational 3-kilometer High-Resolution Rapid Refresh (HRRR) model [92] (Fig. A.6-1). The WIPS service runs a version of WRF on a local server, along with additional software, to generate derived high-resolution low altitude wind forecasts and wind hazard areas in a timely manner in accordance with information provided in client requests. Wind hazard areas (Fig. A.6-2) are generated as polygons representing regions of the forecast where the probability of client-defined hazardous wind conditions exceeds a client-defined threshold. These hazard models can be produced with either a deterministic forecast or an 11-member ensemble forecast. This capability is intended to help clients/users visualize where wind conditions may exceed safety risk tolerances for a particular vehicle (e.g., a maximum wind speed). These hazard areas may also be used by automation to avoid these areas (e.g., geofencing systems).

Inputs:

1) Initial and boundary conditions for WIPS WRF are automatically pulled using HRRR by default, and based on user/client requests that should include:
    a. Spatial resolution (horizontal and vertical) and interval/time resolution
    b. Forecast time window
    c. Coverage volume
    d. Physical parameterizations
2) For hazard polygons: Altitude(s) of interest, wind limits, and probability threshold
3) Reference model version (date/time)

Outputs:

1) 3D wind fields at deterministic forecast time epochs
2) 11-member ensemble 3D wind field forecast epochs
3) Wind hazard polygons

TRL assessment: TRL = 4-5 for simulated forecasts based on historical data. Retrospective forecasts were generated for four case days in the Dallas-Fort Worth (DFW) metropolitan area and then used to/evaluate the service implementation. The highest resolution used during testing was 100 meters over a coverage area of 100 kilometers by 100 kilometers.

For more information: [37][91][92]

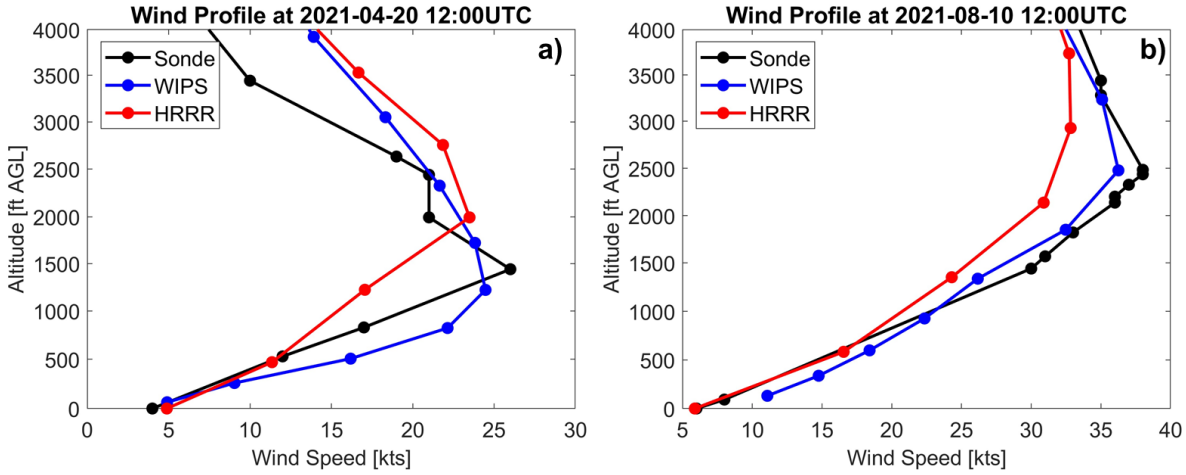

Fig. A.6-1. Example comparisons of simulated winds by WIPS WRF (blue) and HRRR (red) to radiosonde (Sonde) measurements (black) in the DFW Area; (a) April 20, 2021 (12:00 UTC), (b) August 10, 2021 (12:00 UTC). From [37].

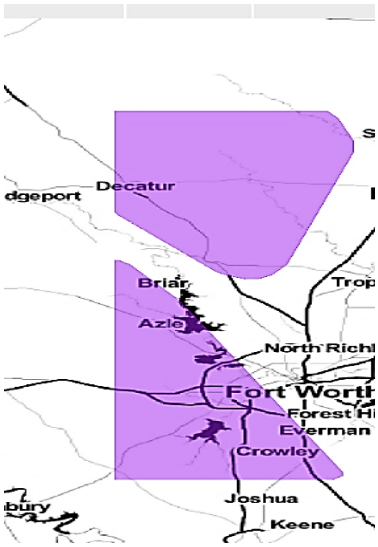

Fig. A.6-2. Example wind hazard region polygons at 305 m (1000 ft) AGL generated for the DFW Area. From [37].

A.6.3 Weather Station Observations Service (WxSt)

Purpose: This service helps mitigate hazards associated with weather conditions that may impact controllability, stability, and/or deviations from a safe flight envelope or intended route of flight. In its current configuration, the service provides contemporaneous and archived weather conditions measured at several locations across the NASA Langley campus (Fig. A.6.3-1). Weather stations installed at these locations are designed to be portable such that re-location to other areas is possible; however, network connectivity would be needed. Each station is designed as a low-cost version of Automated Weather Observing System (AWOS) and Automated Surface Observing System (ASOS) stations, which are installed and operational at many sites across the U.S. [93]. Measurements are sent to a centralized server at pre-defined rates and stored as log files. Clients/users may access these files via internet connection (internal NASA URL). This protocol is similar to data provided by the network of Continuously Operating Reference Stations (CORS) [94].

Inputs: None for the current installation; however, there are configuration settings that can be adjusted to client/user preferences (e.g., update rate, units, filename/structure for archive). The service is not set up as a request/reply protocol.

Outputs:

1) Minimum, average, and maximum wind speed (m/s) (at 5 sec intervals)
2) Minimum, average, and maximum wind direction (deg) (at 5 sec intervals)
3) Air temperature (C); relative humidity (%); air pressure (hPA) (at 1 min intervals)
4) Rain accumulation (mm); rain duration (secs); rain intensity (mm/hr)
5) Hail accumulation (m); hail duration (sec); hail intensity (hits/cm^2)
6) Cloud ceiling (m) and horizontal visibility (m)

[Note: Above archived on a centralized server and broadcast within the NASA network.]

TRL assessment: TRL = 8-9; The network of stations has been operational for multiple years, sensing technology is COTS, and data has been used by multiple teams/projects.

For more information: [36]

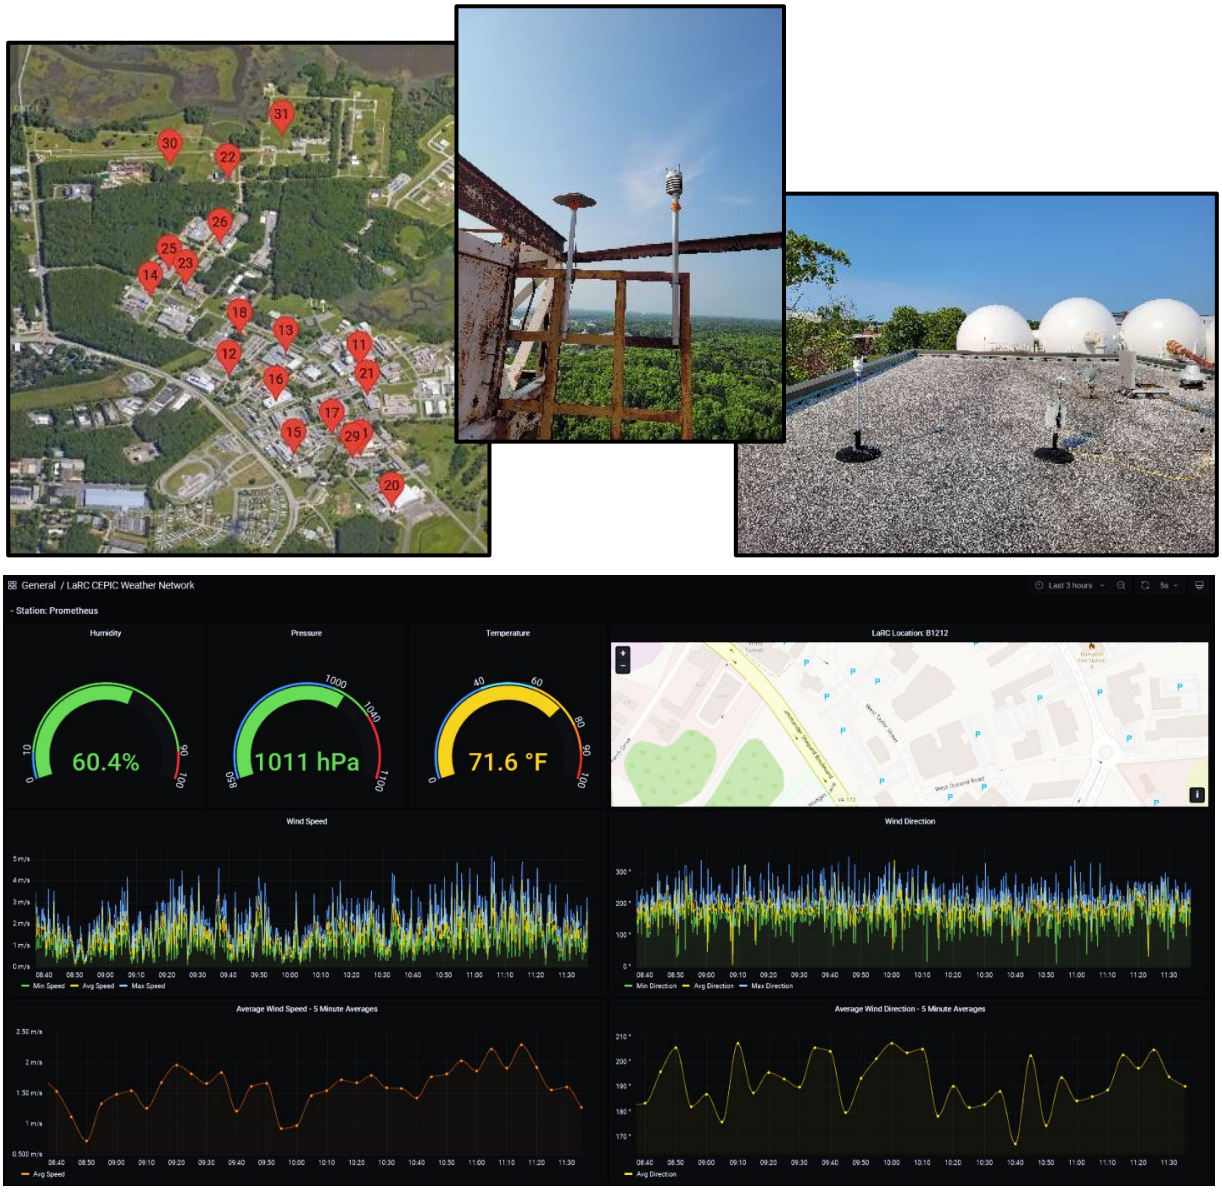

Fig. A-6.3-1 Top left: locations of networked observation stations at NASA LaRC. At each location, weather conditions are monitored and recorded (wind speed and direction at 5 sec intervals; temperature, humidity, and air pressure at 1 min intervals). Each location is also instrumented to serve as a base station for differential GPS navigation. Top middle: Station atop the NASA Lunar Lander Facility, at ~ 70 m altitude. Top right: Typical rooftop station. Bottom: broadcast of station weather readings for the station shown at top right.

A.7 Airspace Dynamic Density Assessment (*DD*)

Purpose: This service helps with management of congested airspaces that may be more likely to occur in future airspace designs for highly autonomous AAM operations. Metrics produced by the service provide insight into the traffic situation to support more informed decision-making (e.g., re-route or divert decisions). Computed metrics are based on the sociology concept of dynamic density (DD) that considers not only the number of flights, but also the interdependence of flights sharing limited resources (e.g., the number of landing pads and parking spots at a vertiport, and the volume and tracks of enroute corridors). The DD outputs could support automation systems as well as human roles (e.g., pilots) when either are making decisions (e.g., whether to proceed, expedite, delay, or divert). Likewise, DD outputs could inform air traffic control (ATC) and vertiport operators when deciding how to manage the airspace and/or vertiport assets. The current version includes two functions that assess congestion – one that predicts congestion-related metrics for vertiport airspace, and one that predicts similar metrics for the enroute airspace.

Depending on which function is selected by the user/client, different information will be needed as input (see Fig. A.7-1 and list below). For vertiport congestion, two metrics are produced: (1) the expected delay for flights that are within the estimated time of arrival (ETA) of a client/user-specified airspace 'freeze horizon'; and (2) a categorical value that represents the anticipated congestion for flights that are within a client/user-specified airspace 'intent sharing horizon' more distant than the freeze horizon (Figure A.7-1). For enroute congestion, DD metrics are produced for each corridor or airspace segment. For each function, DD produces a numeric value that is then categorized into four bins: negligible, low, moderate, and high. The thresholds delineating the bins are set empirically to reflect operational severity of expected delays. AAM/UAM corridors were used for proof-of-concept for the enroute function but other airspace designs may be supported. For either function, metrics are forecasts (i.e., predictions) based on near real-time updates to airspace usage data and trends. As with any prediction, uncertainty grows with the user-specified look-ahead time (i.e., horizons).

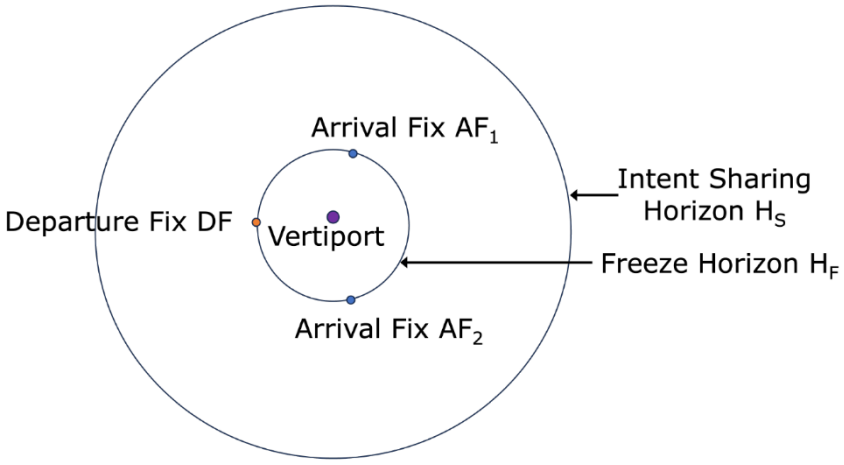

Fig. A.7-1. Relevant airspace features and vertiport congestion metrics. From [95].

Inputs:

1) Vertiport function configuration settings (Figure A.7-1)
    a. Intent sharing horizon (H_S) – the lookahead time at which a flight's proposed ETA is shared with the service
    b. Freeze horizon (H_F) – the time at which a flight commits to a conflict-free ETA issued by the service
    c. Departure/landing pad occupancy – duration of an aircraft's exclusive occupancy of the pad (a.k.a. minimum separation between flights departing/landing from the same pad)
    d. DD-to-category thresholds, to specify alternate mapping of delays to operational severity (optional)
2) Vertiport function live data, for each landing pad of interest (Fig. A.7-2)
    a. ETA for all arrivals (within H_S and updated as necessary until H_F)
    b. ETD for all departures (updated if flights are delayed)
    c. Number of empty available parking spots
3) Enroute function configuration settings
    a. Prior to takeoff and for each aircraft: 4-D trajectory (time-series of position and altitude) or flight plan (aircraft type, route of flight, expected time of departure, speed, and altitude at each waypoint)
    b. Airspace geo-spatial design information (e.g., corridor volumes, tracks, and waypoints)
    c. Maximum proximity between two aircraft considered part of same cluster
    d. DD-to-category thresholds, to specify alternate mapping of delays to operational severity (optional)
4) Enroute function live data
    a. Aircraft position estimates (every 5 sec or better)

Outputs:

1) Vertiport function (Figure A.7-3)
    a. For schedule window defined by H_F (aka, the frozen part of the schedule, or frozen "chunk"):
        i. Categorical DD value (negligible, low, moderate, or high)
        ii. Revised conflict-free ETA for each flight inside H_F
        iii. Delay from proposed ETA for each flight inside H_F
        iv. Residual delay for the frozen chunk that will propagate to the next chunk
    b. For schedule windows/chunks beyond H_F but within H_S, with each chunk of duration H_F:
        i. Categorical DD value (negligible, low, moderate, or high) for flights with proposed ETA in each chunk
        ii. Residual delay predicted to propagate forward to the next chunk
2) Enroute function, for every time step until a look-ahead horizon, and for each client/user-specified corridor or airspace segment:
    a. Categorical DD value (negligible, low, moderate, or high)

b. Aircraft density, the minimum and mean separation distance between aircraft in the corridor/airspace, cluster density, and mean cluster size. [Note: A cluster is defined as three or more aircraft in close proximity to each other.]

3) Enroute function, for a client/user specified flight:

a. Categorical DD value (negligible, low, moderate, or high) predicted for each time step of the flight's trajectory

TRL assessment: TRL = 4-5, testing/validation to date limited to synthetic data representative of variable traffic densities. Usability of service connection was validated for enroute DD on notional DFW UAM corridor-based airspace with low traffic densities during multi-partner simulation testing. Usability of products for vertiport DD is planned for an upcoming human-in-the-loop study.

For more information: [39][95]

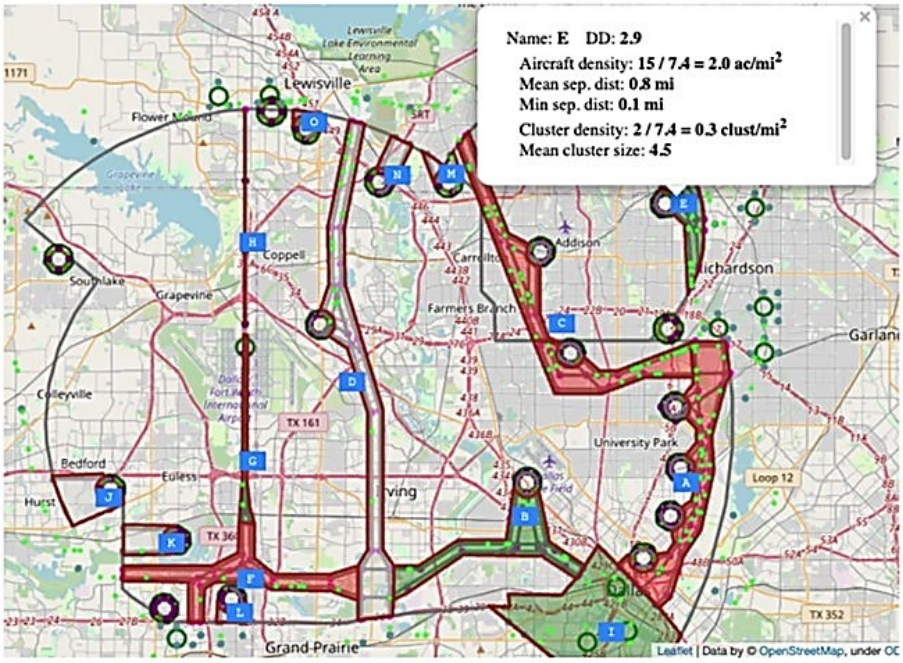

Fig. A.7-2. Enroute example visualization of DD for the AAM/UAM corridor nowcast. Bright green dots show most recent known positions of all aircraft. Corridors are color-coded for DD category: negligible (gray), low (green), moderate (yellow), and high (red). More detailed computed information is available, such as shown for corridor E (upper right text box). From [39].

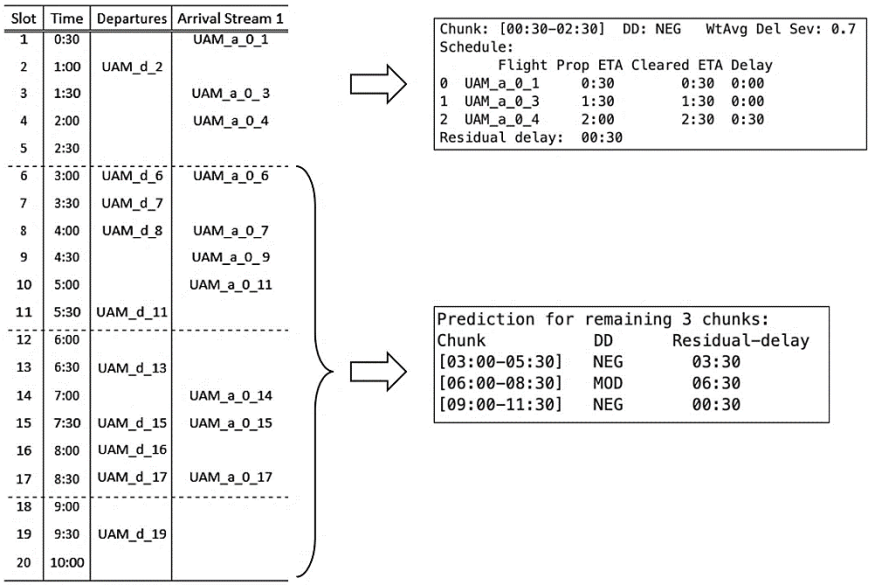

Slot	Time	Departures	Arrival Stream 1
1	0:30		UAM_a_0_1
2	1:00	UAM_d_2	
3	1:30		UAM_a_0_3
4	2:00		UAM_a_0_4
5	2:30		
6	3:00	UAM_d_6	UAM_a_0_6
7	3:30	UAM_d_7	
8	4:00	UAM_d_8	UAM_a_0_7
9	4:30		UAM_a_0_9
10	5:00		UAM_a_0_11
11	5:30	UAM_d_11	
12	6:00		
13	6:30	UAM_d_13	
14	7:00		UAM_a_0_14
15	7:30	UAM_d_15	UAM_a_0_15
16	8:00	UAM_d_16	
17	8:30	UAM_d_17	UAM_a_0_17
18	9:00		
19	9:30	UAM_d_19	
20	10:00		

Fig. A.7-3. Example of DD service predictions for vertiport with one departure stream and one arrival stream. Freeze horizon (i.e., chunk duration) is 3 min and intent sharing horizon is 10 min. Left: Synthetic schedule used as input. Right: Resulting service predictions at time 0.00 for first chunk (top) and remaining chunks (bottom). In lower right, abbreviations are: NEG (negligible), MOD (moderate). From [95].

A.8 Battery Prognostics and Health Management (*BP/BHM*)

Purpose: These two services help to mitigate hazards arising from control or propulsion system shutdown during flight due to insufficient battery power (which may be caused by use of unqualified flight batteries or incorrect battery power consumption estimation). Each service may be used individually, or they may be combined. (1) Battery prognostics (BP) provides clients/users with a battery's estimated state-of-charge (SOC) or remaining flight time (RFT) at each waypoint in a client/user-provided flight plan or along a predicted trajectory ahead of the aircraft's current position. When used during pre-flight planning, this allows operators to plan flights with appropriate safety margin regarding battery power and to set power margins for alerting functions. In-flight use provides SOC or RFT estimates for the current aircraft location and the remaining points along the planned or predicted trajectory. This supports re-planning in-flight as part of contingency management. The BP request types and service outputs for each flight phase are summarized in Table 5. (2) Battery health management (BHM) helps operators manage usage and life thresholds by (a) alerting operators if an unqualified battery is selected for a flight, (b) helping to automate and enforce battery use safety policy so operators can know which batteries are due for recertification or removal from flight stock, and (c) helping to automate, track, and maintain battery parameter update cycles over a battery's lifetime. The BHM request types and service outputs are summarized in Table 6. When used together with BP, this helps to ensure predictions of SOC or RFT are tuned to the latest parameter updates to reduce uncertainty in the predictions (i.e., better accuracy). Two types of battery models are used in the current implementation: an electro-chemistry model (EChM) and an equivalent circuit model (ECM) (see Fig. A.8-1 and Fig. A.8-2).

Table 5 Battery Prognostics Service Request Types

Service	Request Type	User/Client Inputs	Data Returned
BP Preflight	Request SOC estimates at points along flight path	Battery ID EChM parameters Points along flight path	Projected battery SOC at each point along flight path
BP Inflight	Request live SOC estimate at current position	Battery ID EChM parameters Points along flight path Live stream of battery voltage and current	Live updates to SOC/RFT estimate at current position
BP Inflight	Request live SOC estimate at remaining points along flight path	Battery ID EChM parameters Points along flight path Live stream of battery voltage and current	Live updates to predicted battery SOC/RFT at remaining points along flight path
BP Postflight/Offline	Project state of health to end of flight (using neural network approach)	Battery ID EChM parameters Points along flight path Live stream of battery voltage and current (may be log files)	Computed battery degradation parameters for resistance and capacity. Used for updating EChM or ECM parameters to use in subsequent flights. (See Fig. A.8-1)

Table 6. Battery Health Management Service Request Types

Service	Request Type	User/Client Inputs	Data Returned
BHM Preflight	Request battery use status	Battery ID	Flight qualified status Flight system used Cycle capacity Number cycles used Last cycle date
BHM Preflight	Request EChM parameters	Battery ID	EChM parameters (see [47][96])
BHM Preflight	Request ECM parameters	Battery ID	ECM parameters (see [96][97])

[Note: If the on-board BP function is used, SOC/RFT estimates can be provided to autopilot contingency management functions and/or downlinked to ground operator(s). Alerts may be generated based on operator policy.]

TRL assessment: TRL=4-6; For BP, development and testing has been ongoing for many years resulting in good validation of the methods; however, the methods require tailoring to some battery types, vehicle types, and expected flight profiles. For BHM, development and testing are limited to NASA's flight test environment, facilities, equipment, policies, and procedures. Application by non-NASA operators has not been validated.

For more information: [47][96][97][98][99]

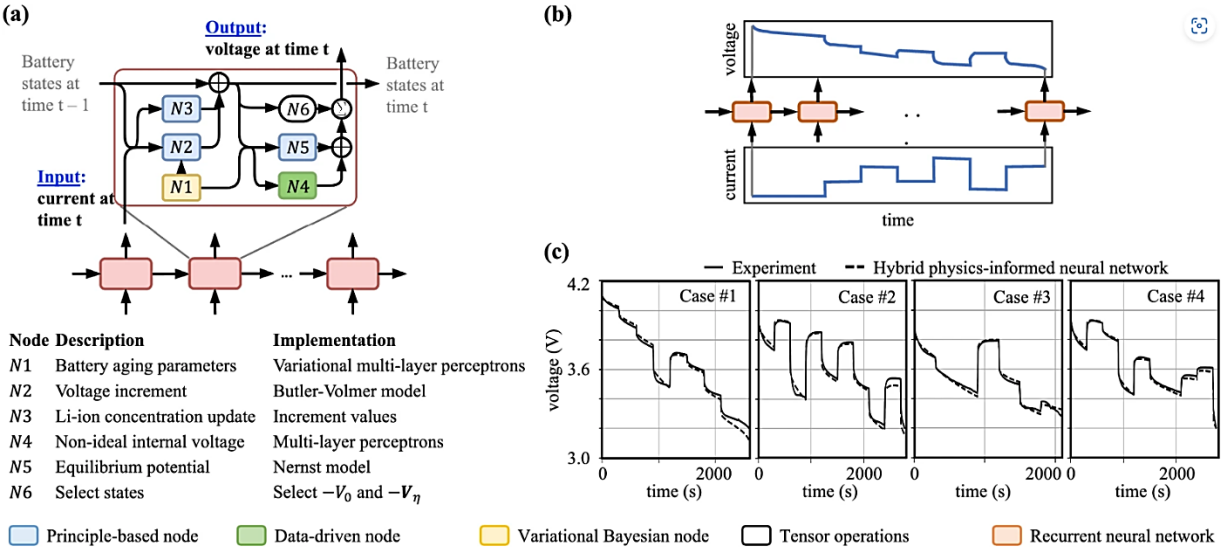

Fig. A.8-1. Implementing a hybrid physics-informed neural network for Li-ion battery prognosis. (a) Recurrent neural network implements numerical integration of governing equations in the state-space representation. The recurrent unit is composed by surrogate models describing the main phenomena driving the battery electrochemistry, a data-driven node which captures the non-ideal internal voltage, and a variational Bayesian node which models aging through degradation of battery parameters. (b) The hybrid model takes a current measurement time series as input and returns battery voltage. This allows use of the model for tracking of discharge cycles as well as forecasting of future missions. (c) Comparison between experimental data and predictions from the hybrid physics-informed neural network. Predictions are obtained for batteries not used in the training set. From [99].

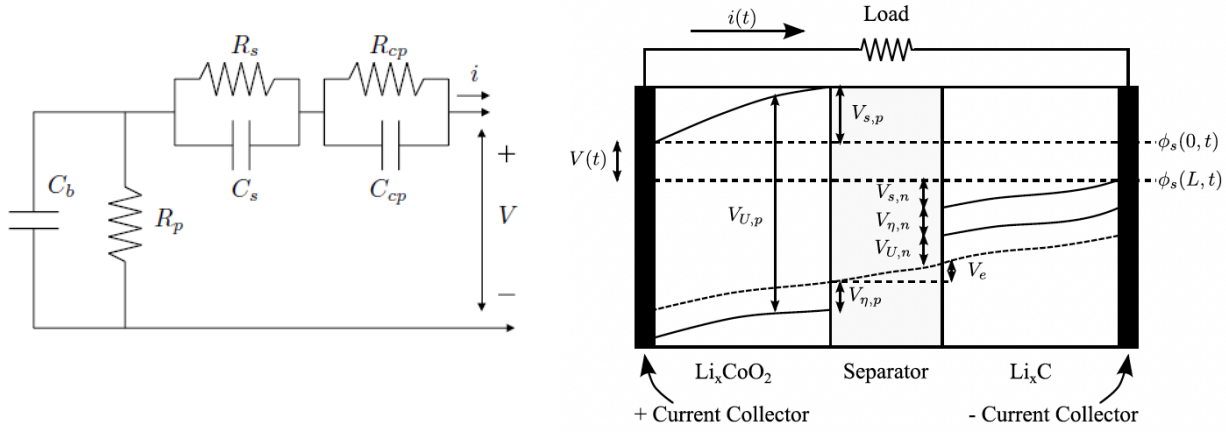

Fig. A.8-2. BP/BHM battery models and parameters. Left: Equivalent Circuit Model (ECM) (from [96]). Right: Electro-Chemistry Model (EChM) (from [47]).

A.9 Flight Performance Assessment (*FPAS*)

Purpose: This service employs nonlinear aerodynamic and power consumption models integrated into a 6-degree of freedom (DOF) flight dynamics simulation. This simulation uses a control structure typical for sUAS multirotor vehicles to enable waypoint tracking and other highly augmented control modes (e.g., position and velocity control). There is an inner loop that regulates attitude while an outer loop provide attitude commands to control position. The service can provide other applications/users with a predicted trajectory, including power consumption, for a user-selected aircraft and flight plan (Figure A.9-1). This service is designed to support multiple sUAS aircraft platforms (quadcopter, octocopter), as well as passenger-carrying configurations (lift+cruise, multirotor, tilt wing). Aerodynamic performance models are currently based (in part) on wind tunnel testing data [63] and are updated frequently [64], and future updates will include models based on computational fluid dynamics (CFD) lookup tables and analytical methods to support eVTOL configurations. The service can provide predicted performance for numerous flight conditions such as nominal, degraded, and off-nominal due to failures.

Inputs:

1) Desired spatial and time resolution for outputs (e.g., 5m at 2 sec intervals)
2) Aircraft information (e.g., type and weight) (See Note 1)
3) Flight plan information (e.g., 3D waypoints, airspeed along legs, climb/descent rates for segments with altitude change (e.g., takeoff/landing))
4) Environment information (e.g., wind and temperature) (See Note 2)
5) Expected battery performance information (See Note 2)
6) Failure conditions (e.g., motor failure(s) along the flight plan) (optional) (See Note 3)

[Note 1: If type and weight do not match class of available models, then user/client may need to provide model parameters for the aircraft or be advised that predictions will have unknown uncertainties due to the mismatch.]

[Note 2: These inputs may be manually entered or provided by connection to other information services (e.g., Wind or BP services).]

[Note 3: If this information is provided, predictions will be based on simulation of the failure(s) at the requested time(s) or point(s) in the planned flight.]

Outputs: At desired spatial and time resolution, the client/user is provided

1) Predicted aircraft states (e.g., position, velocity, attitude)
2) Battery state of charge (SOC) and remaining flight time (RFT)
3) Warnings for locations or times when envelope boundaries are predicted to be exceeded

TRL assessment: TRL = 3-4. Validations to date include a 6-DOF model for an octocopter and a low-order model for a six-passenger quadrotor UAM vehicle. Uncertainty quantification and inclusion is not fully complete. Testing of connections to commercially available or in-house wind service APIs anticipated in 2024.

For more information: [41][63][64]

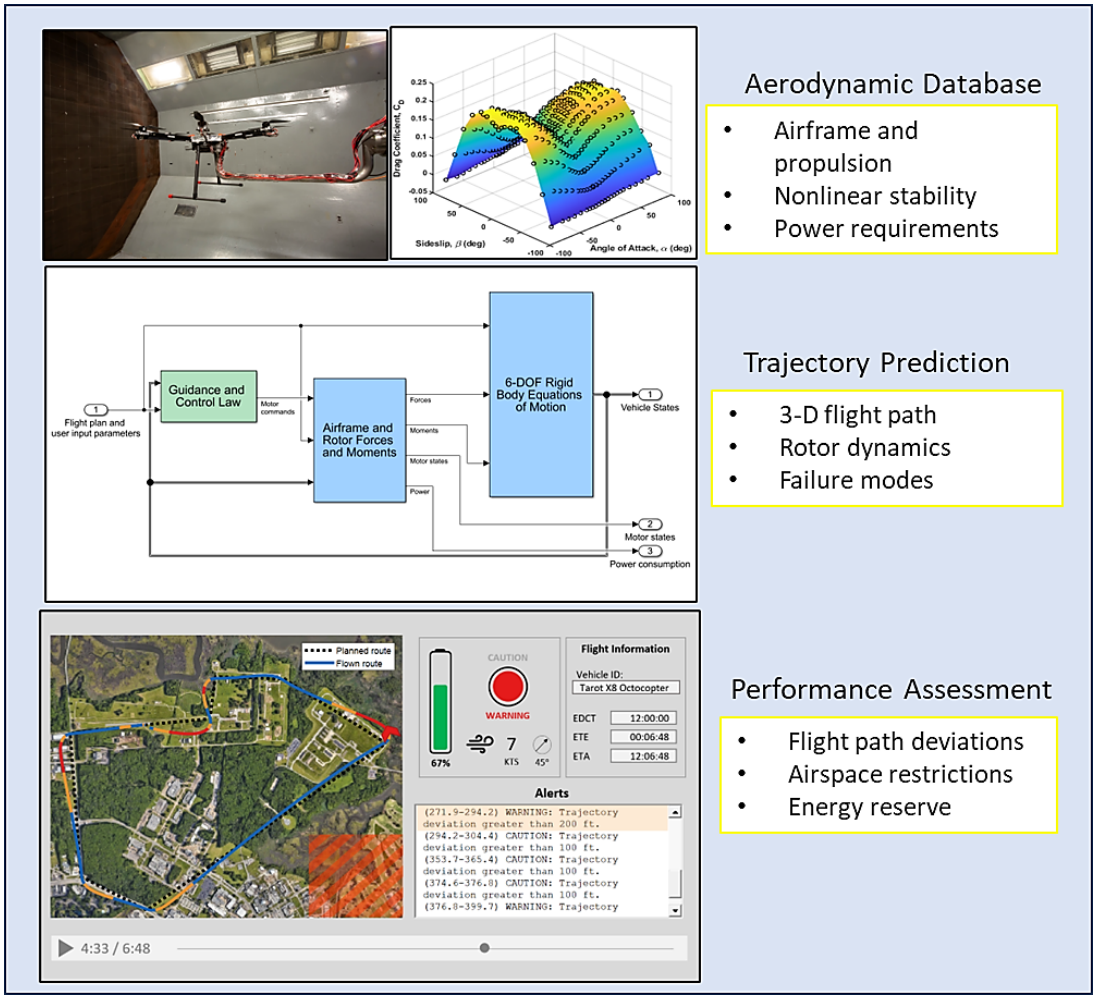

Fig. A.9-1. Flight performance assessment: Aerodynamics of representative vehicles measured in the LaRC 12-Foot Low-Speed Wind Tunnel (top) and vehicle-specific propulsion and kinematic characteristics (middle) are used to estimate power consumption and trajectory deviations from a user-specified flight plan (bottom).

Appendix B. Experimental Aircraft, Ground Equipment, and Flight Checklists

B.1 Aircraft and Ground Equipment

A Tarot T18 Octocopter frame was outfitted for flight with COTS hardware required for flight:

- an autopilot (Pixhawk Cube running ArduCopter version 4.0.7)
- a navigation GNSS receiver and compass (u-Blox M8N)
- motors and speed controllers (KDE Direct 4213XF-360 motors and UAS40UVC controllers)
- a lithium-ion battery for propulsion (22 Ah 6S 22.2v)
- radio transceivers for command and control (Futaba R7008SB receiver) and flight telemetry (RFD 900x)

Research sensors, computers and associated electronics were attached to the frame or on payload trays, including:

- radio-controlled cut-off relays to power down research and sever communications to autopilot from research hardware (Pololu #2804)
- two compute nodes (Intel NUC7i7BNH)
- a radio transceiver for research telemetry (Teltonika RUTX11 Cellular Router)
- battery monitoring electronics (NASA custom 'MiDDAS Battery Node') to measure and log propulsion battery temperature, pack voltage, cell voltage, and pack current
- a research Battery (Tattu 4s 8000 mAh lithium polymer)
- a geolocation and containment monitor with independent GPS receiver (NASA 'Safeguard')
- software defined radio electronics (Signal Hound BB60C)
- a full-signal GPS receiver and logger (u-Blox Zed-F9P, Sparkfun SerialDataLogger)

With this payload, the vehicle takeoff weight was 13.2 kg (29.93 lb). A photo of the aircraft and a block diagram of avionics interconnection are shown in Fig. B.1-1.

Ground equipment included COTS hardware required for flight:

- an RC transceiver (Futaba T18MZ)
- a laptop running Windows 10 with groundstation software (Mission Planner version 1.3.7) and telemetry receiver (RFD 900x)
- a Linux computer (Intel NUC7i7BNH) running research ground station software (NASA custom 'GroundWatch')

The preflight checklists for the vehicle and ground equipment are included below.

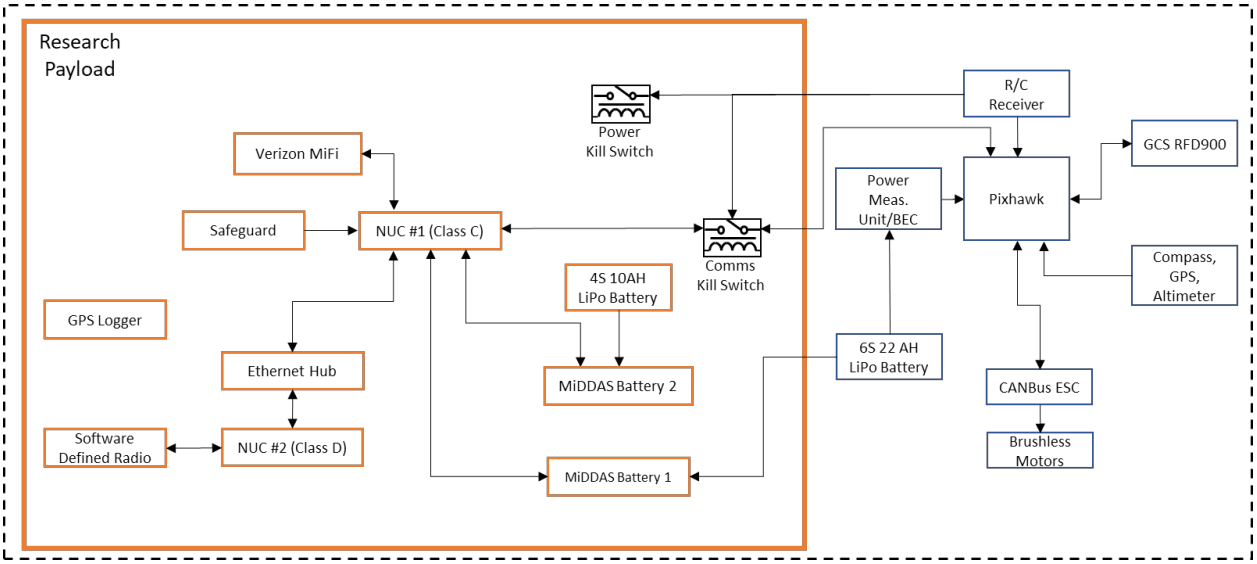

Fig. B.1-1. Experimental aircraft (top) and avionics configuration (bottom).

B.2 Flight Checklists

Date: _____

First Flight of the Day

Lab	Field	
☐	☐	Batteries charged and logged
☐	☐	Transmitter inspected for damage
☐	☐	Inspect each motor arm for damage
☐	☐	Propellers inspected for damage and propeller screws secured
☐	☐	Inspect landing gear for damage
☐	☐	Inspect airframe for damage
☐	☐	Check that all screws are secured
☐	☐	Inspect batteries and battery tray for damage
☐	☐	Ensure there is no interference between hardware components when in flight orientation
☐	☐	Check all device wires for damage, are securely attached, and connectors seated
☐	☐	Check all whisker antennae are straight and aligned in the correct axis
☐	☐	Connect batteries required for vehicle operation
☐	☐	Check RC failsafe program. (Power Tx off and wait for low beep)
☐	☐	Verify GCS connection over telemetry radio
☐	☐	Check operation of research cutoff relay (if applicable)
☐	☐	Turn on Applicable Research Equipment
☐	☐	Perform RC transmitter range check

Aragog_T18 Checklist V1.0 November 2021

GROUND CONTROL STATION

Software Version

Firmware	Arducopter v4.0.7 or	_____
Mission Planner	V1.3.68 or	_____
Research Software		_____

Configuration

	Parameter	**Normal Value**		**Other Value**
Frame	FRAME_TYPE:	1 (X)	☐	
	FRAME_CLASS:	3 (Octa)	☐	
Flight Modes	Channel 1	POS Hold	☐	
	Channel 2	Stabilize	☐	
Optional Channels	Channel 7	RTL	☐	
	Channel 8	Land	☐	
Power	BATT_CAPACITY:	20000 mAh	☐	
	BATT_MONITOR:	4 (V and C)	☐	
Radio	Model:	Aragog V1.1	☐	
Sensors	AHRS_ORIENTATION:	0 (None)	☐	
	COMPASS_ORIENTATION:	0 (None)	☐	
Parameters	MOT_SPIN_ARMED:	0.1	☐	
	WP_YAW_BEHAVIOR:	0	☐	
	WPNAV_SPEED:	500cm/s	☐	
	LOG_FILE_DSRMROT:	1 (1 DFL/FLT)	☐	
Failsafe Triggers	BATT_FS_LOW_ACT	2 (RTL)	☐	
	BATT_LOW_VOLTAGE	19.0 Volts	☐	
	FS_EKF_ACTION	1 (LAND)	☐	
	FS_GCS_ENABLE	1 (Enabled Always RTL)	☐	
	FS_THR_ENABLE	1 (Enabled Always RTL)	☐	
GeoFence	FENCE_ACTION	RTL or LAND	☐	
	FENCE_ALT_MAX	100 m	☐	
	FENCE_ENABLE	1 (Enabled)	☐	
	FENCE_TYPE	Alt and Poly	☐	
Return to Launch	RTL_ALT	3000	☐	
	RTL_SPEED (0 uses WPNAV_SPEED)	700	☐	
	LAND_SPEED:	50 cm/s	☐	
	Arming Check	1 (All enabled)	☐	

FLIGHT CHECK LIST

Record Vehicle Flight Number:				
Record motor battery number:				
Record motor battery voltage:				
Record research battery number:				
Record research battery voltage:				
Inspect vehicle for damage				
Check Tightness of motor arm pivot bolts				
Install batteries into vehicle				
Ensure PPE is being worn				
Plug in motor & research battery				
Turn on research system				
Connect GCS and verify connection				
Verify all flight modes				
Load flight plan if required				
Verify GPS status				
Verify vehicle position and heading on map				
Verify altitude				
Verify Roll and Pitch				
Verify battery voltage and current				
Turn RC Relay ON				
Turn Master Power Switch OFF				
Move vehicle to flight line				
Record flight time				
Was any damage incurred during the flight?				
Record motor battery post flight voltage				
Record research battery post flight voltage				

Appendix C. Summaries of Selected Test Flights of Risk-Informed Auto-Mitigation Capability

The monitor/assess/mitigate architecture with active autopilot control (Fig. 3.1-1) produced mitigation maneuvers in response to flight hazards in nine test flights in course of a campaign in the spring of 2023. Two of those flights are described in detail in Section 3.3. Results from an additional seven flight tests (numbered 95, 96, 98, 102, 112, 113 and 114 in the campaign) are presented in brief form below, in the same format as Fig. 3.3-1 and Fig. 3.3-2.

Flight 095

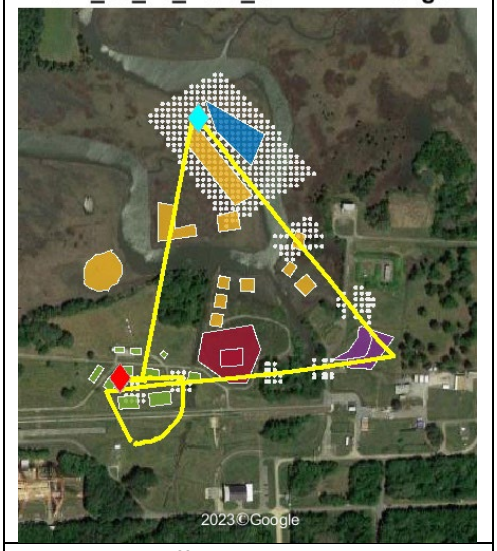

Manual take off in POS_HOLD. Started AUTO mission, lost GCS connection halfway through the flight. Aborted RTL and moved to first waypoint. Engaged AUTO, research Triggered land mode when enabled. Moved further east, battery was low at that point so brought in and manually landed.

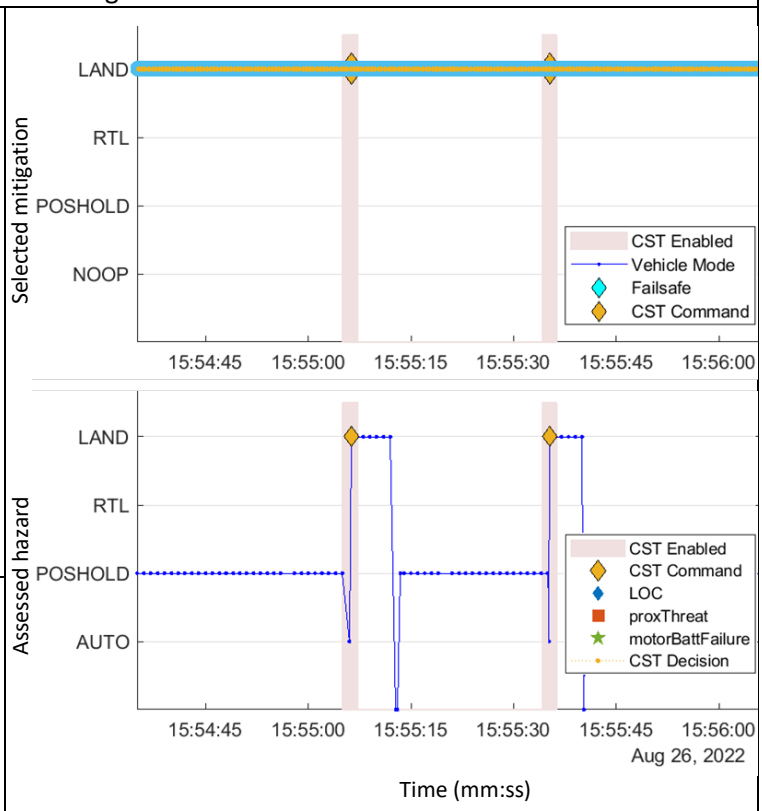

Fig. C-1. Flight 95 map (top left), brief description (bottom left), assessed hazards (bottom right) and selected hazard mitigation (top right). See Fig. 3.3-2 and Fig. 3.3-3 for detailed descriptions of all annotations.

Flight 096

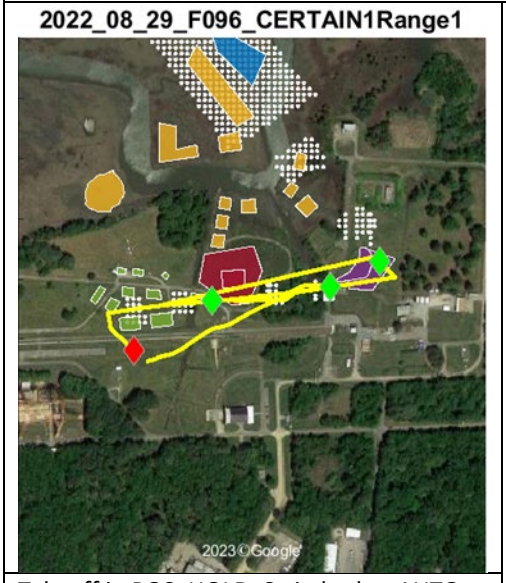

Takeoff in POS_HOLD. Switched to AUTO and vehicle immediately switched to LAND mode. Pilot took over and positioned over waypoint #3. Went back to AUTO but with comms relay open. Closed comms relay shortly after turn to the North. Research system commanded a change to POS_HOLD. GCSO set active waypoint and the vehicle flew back to near the start of the flight path. Back to AUTO mode and began flying to the East. Research switched mode to POS_HOLD. Pilot positioned vehicle to the East and went back to AUTO mode. Research computer changed mode back to POS_HOLD. Pilot manually flew vehicle back to takeoff site and did a manual landing.

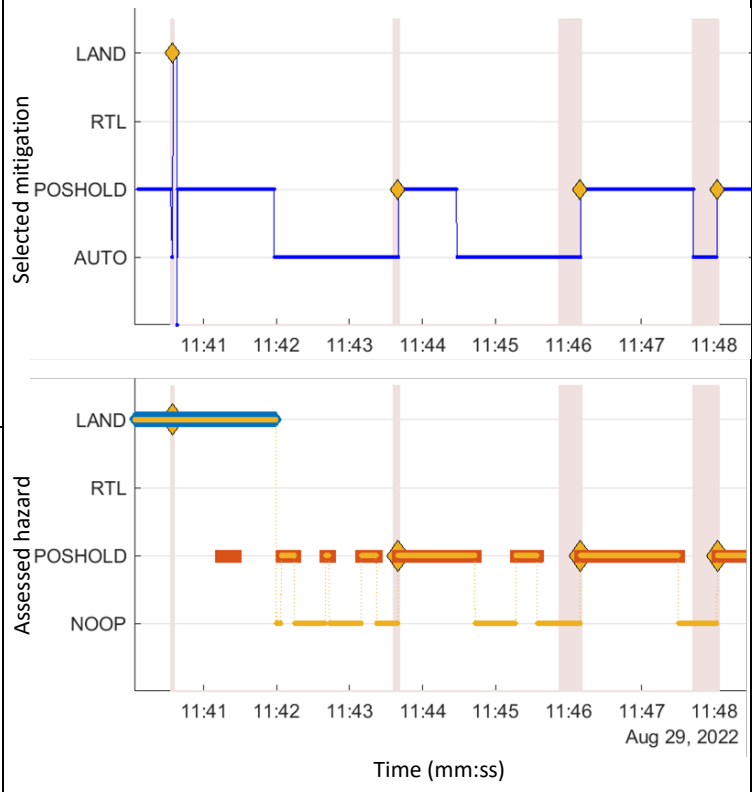

Fig. C-2. Flight 96 map (top left), brief description (bottom left), assessed hazards (bottom right) and selected hazard mitigation (top right). See Fig. 3.3-2 and Fig. 3.3-3 for detailed descriptions of all annotations.

Flight 098

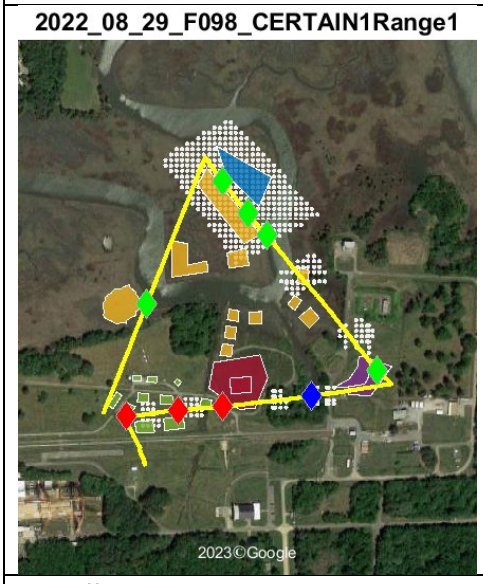

Takeoff in POS_HOLD. Pilot switched to AUTO with comms relay open. Closed relay after reached. Had many switches to POS_HOLD RTL, and LAND during the flight. Some triggered right after closing the relay, some were delayed. Hovered over last waypoint expecting LAND mode but it never triggered. Pilot initiated landing after reaching about 20.5V. Landed using LAND mode due to breezy winds conditions.

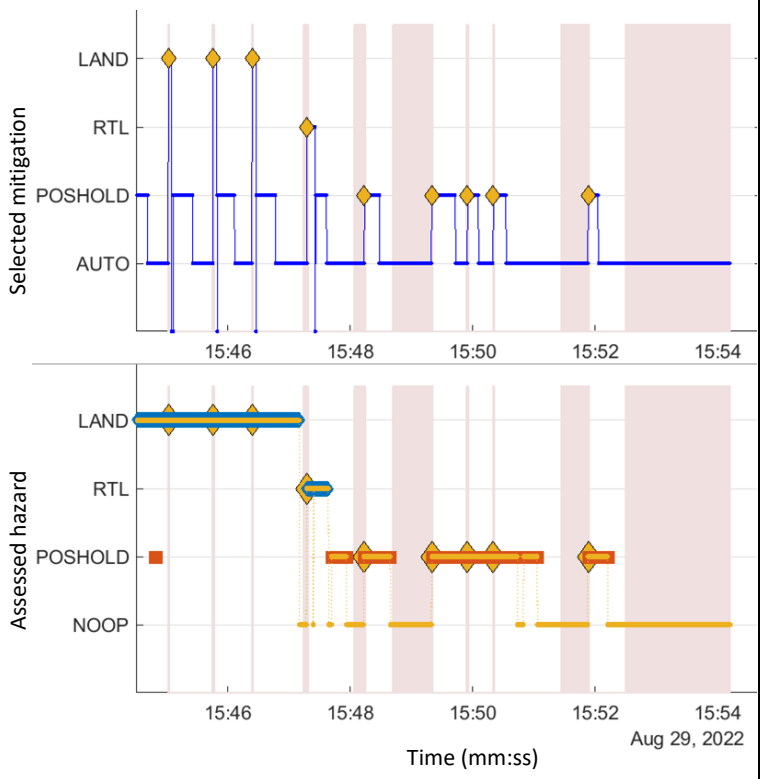

Fig. C-3. Flight 98 map (top left), brief description (bottom left), assessed hazards (bottom right) and selected hazard mitigation (top right). See Fig. 3.3-2 and Fig. 3.3-3 for detailed descriptions of all annotations.

Flight 102

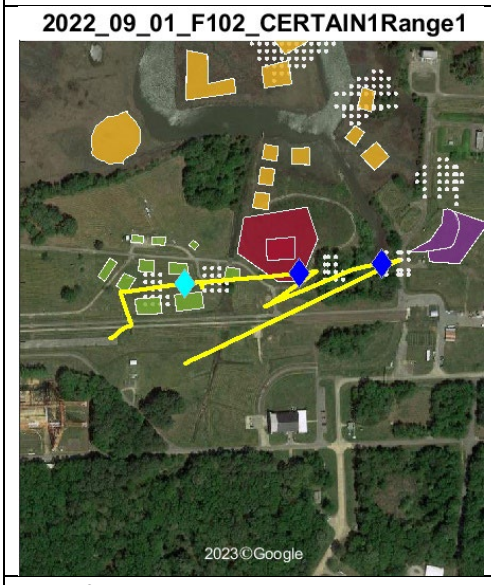

Test of research system with research telemetry unplugged. Takeoff in POS_HOLD. Switch to AUTO. Had GCS Failsafe. Pilot changed flight mode back to AUTO. Had 2 instances of research commanded RTL. Second RTL ended with vehicle doing a full stop landing.

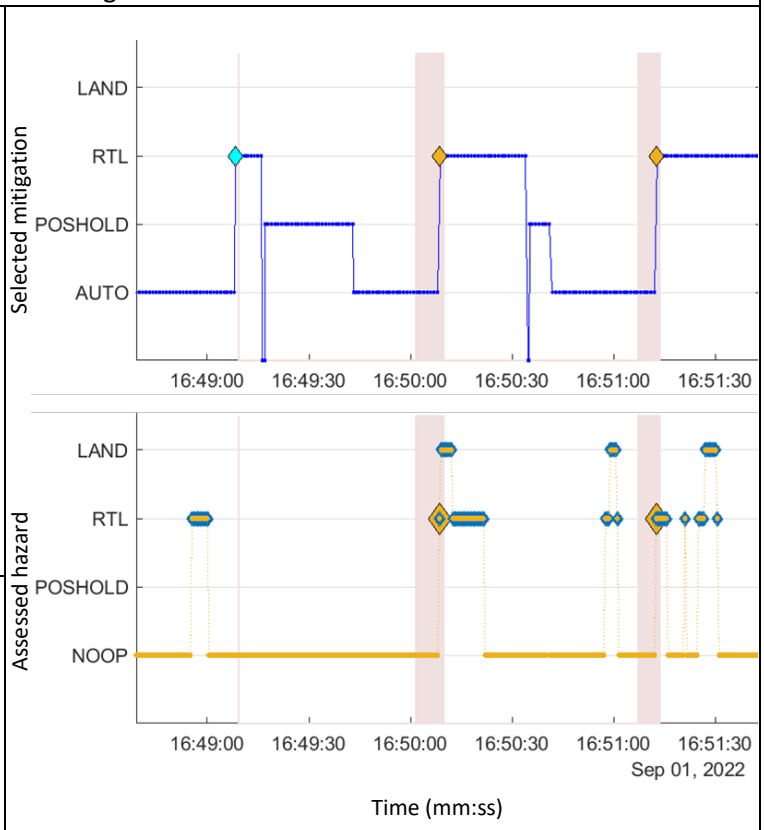

Fig. C-4. Flight 102 map (top left), brief description (bottom left), assessed hazards (bottom right) and selected hazard mitigation (top right). See Fig. 3.3-2 and Fig. 3.3-3 for detailed descriptions of all annotations.

Flight 112

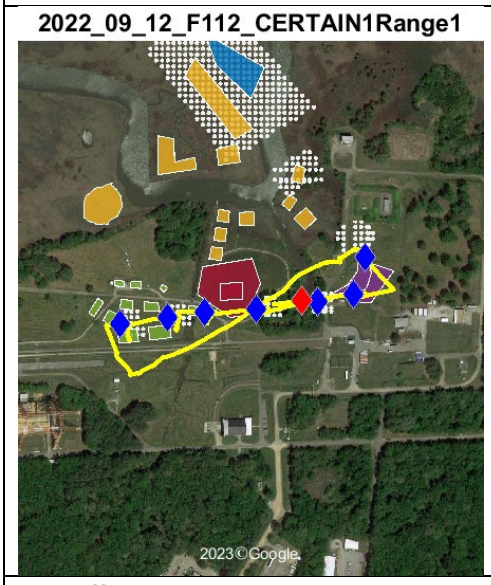

Take off in POS_HOLD. Switch to AUTO. Had multiple research commanded RTL and LAND commands but no POS_HOLD mode. Manually flew vehicle back to takeoff site in POS_HOLD and did manual landing.

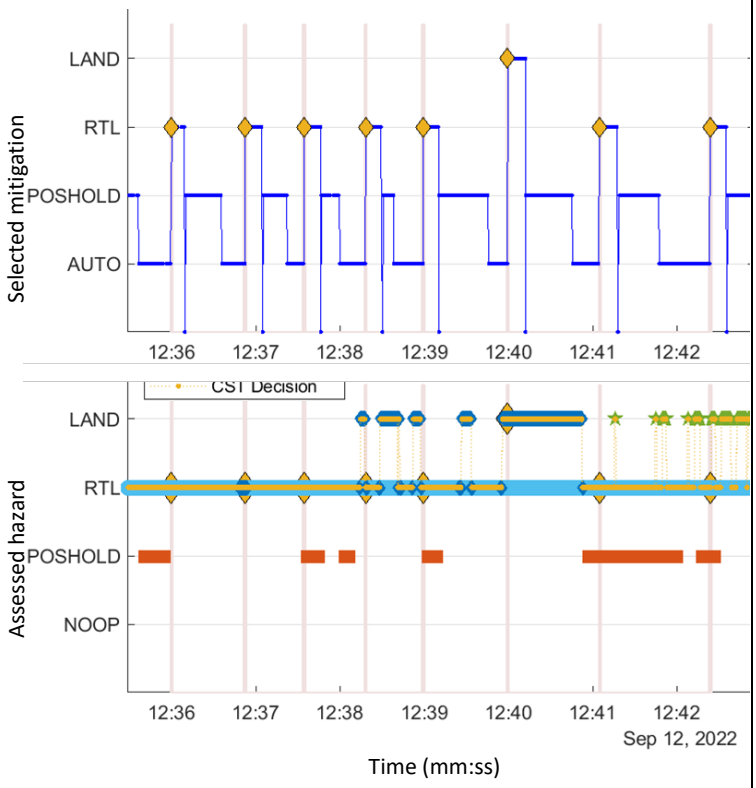

Fig. C-5. Flight 112 map (top left), brief description (bottom left), assessed hazards (bottom right) and selected hazard mitigation (top right). See Fig. 3.3-2 and Fig. 3.3-3 for detailed descriptions of all annotations.

Flight 113

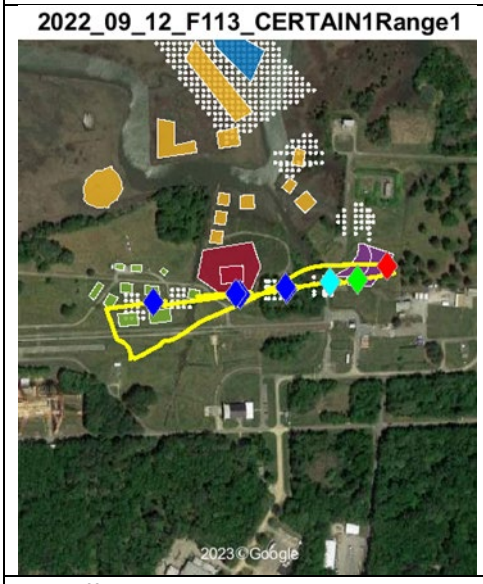

Takeoff in POS_HOLD. Went to AUTO mode and vehicle flew path. Had multiple RTL's caused by research. 2 RTL's caused by lost GCS link. LAND mode issued by research. Vehicle was flown back manually. While hovering the vehicle seemed to fall a bit after hovering. Manual landing by pilot.

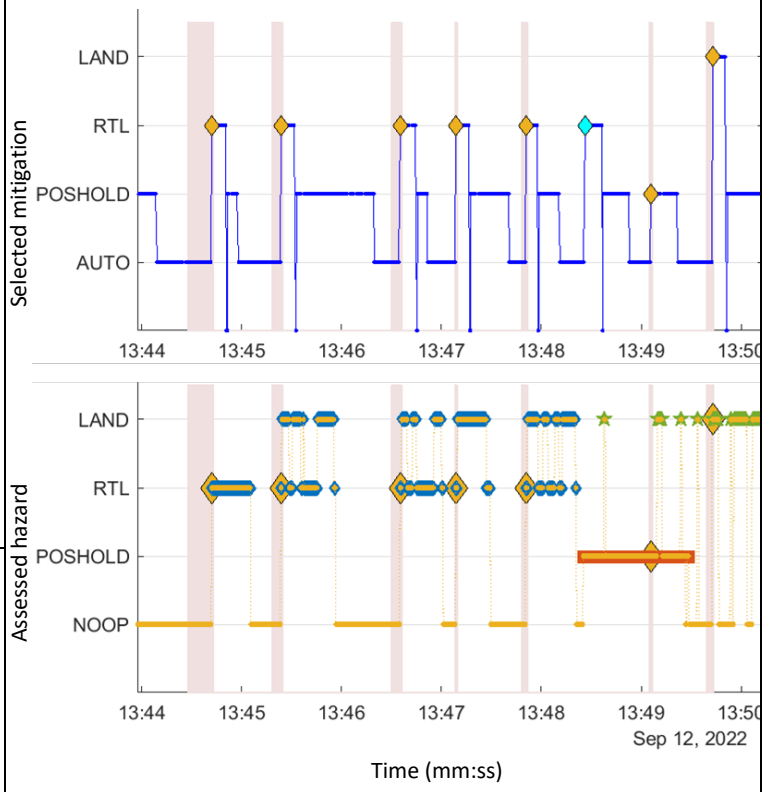

Fig. C-6. Flight 113 map (top left), brief description (bottom left), assessed hazards (bottom right) and selected hazard mitigation (top right). See Fig. 3.3-2 and Fig. 3.3-3 for detailed descriptions of all annotations.

Flight 114

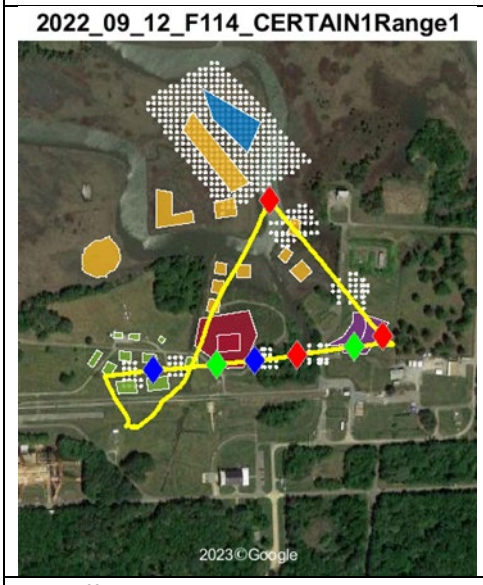

Takeoff in Pos Hold. Engaged AUTO shortly after reaching first waypoint. Had multiple changes to RTL, POS_HOLD, and LAND during flight. Made it about half way through flight plan before GCSO gave call to land due to low battery. Manual Landing.

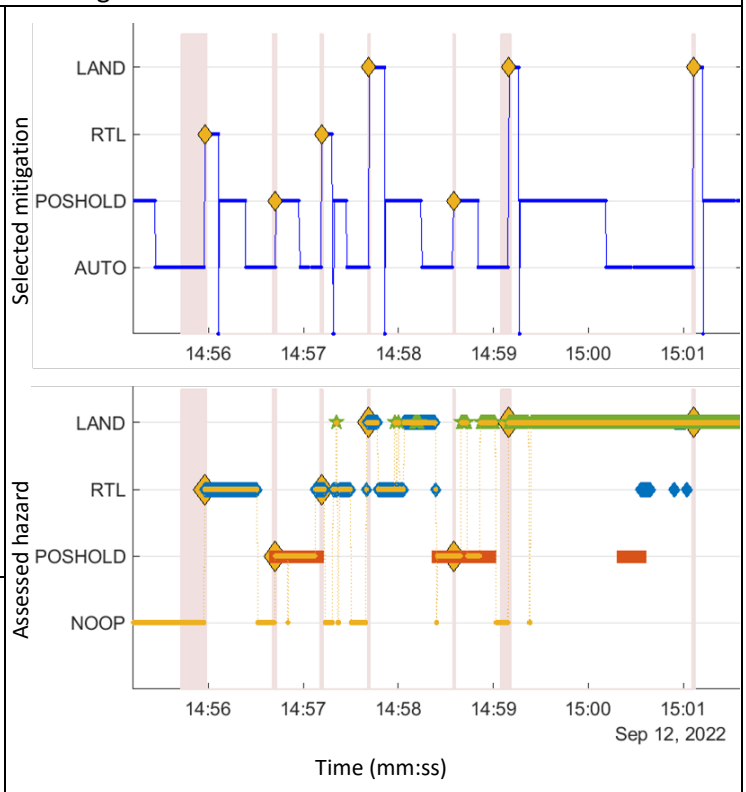

Fig. C-7. Flight 114 map (top left), brief description (bottom left), assessed hazards (bottom right) and selected hazard mitigation (top right). See Fig. 3.3-2 and Fig. 3.3-3 for detailed descriptions of all annotations.

Appendix D. Acronyms

AAM	Advanced Air Mobility
AD	Anomaly Detection
AGL	Above Ground Level
AIS	Aeronautical Information Service
API	Application Programming Interface
APmon	Autopilot Monitor
APPDAT	Application Platform, Packaged Deployment and Analytics Technologies
ARC	Ames Research Center
ARMD	Aeronautics Research Mission Directorate
ASOS	Automated Surface Observing System
ASRS	Aviation Safety Reporting System
ATM	Air Traffic Management
AWOS	Automated Weather Observing System
BHM	Battery Health Management
BP	Battery Prognostics
CAPS	Corridor Assessment of Positioning Systems
CFD	Computational Fluid Dynamics
CFR	Code of Federal Regulations
cFS	coreFlight System
ConOps	Concept of Operations
ConUse	Concept of Use
CORS	Continuously Operating Reference Stations
COTS	Commercial Off-the-shelf
CST	Contingency Select and Trigger
DAIDALUS	Detect and Avoid Alerting Logic for Unmanned Systems
DB	Decibel
DD	Dynamic Density
DFW	Dallas-Fort Worth International Airport
DOF	Degrees of Freedom
DOP	Dilution of Precision
DSM	Digital Surface Model
DST	Drone Safety Team
EChM	Electro-chemical Model
ECM	Electric Circuit Model
EM	Electromagnetic
EOSN	Environment Observation Station Network
ESC	Electronic Speed Controller
ETA	Estimated Time of Arrival
eVTOL	electric Vertical Takeoff and Landing
FAA	Federal Aviation Administration
FCC	Federal Communications Commission
FIMS	Flight Information Management System
FPAS	Flight Performance Assessment Service
GAPS	Geometric Assessment of Positioning Systems
GeoJSON	Geographic JavaScript Object Notation
GLONASS	GLObalnaya NAvigatsionnaya Sputnikovaya Sistema
GNSS	Global Navigation Satellite Systems

GPS	Global Positioning System
GRASP	Ground Risk Assessment Service Provider
GUI	Graphical User Interface
GWU	George Washington University
HATIS	Human Autonomy Teaming Interface System
HDOP	Horizontal Dilution of Precision
HDV	High Density Vertiport
HECAD	Hierarchical Embedded Cyber-Attack Detection
HRRR	High Resolution Rapid Refresh model
HTTP	Hypertext Transfer Protocol
IASMS	In-time Aviation Safety Management System
ICAROUS	Integrated Configurable Algorithms for Reliable Operation of Unmanned Systems
IMS	Inductive Monitoring System
JSC	Johnson Space Center
JSON	JavaScript Object Notation
LaRC	Langley Research Center
LSTM	Long Short Term Memory
MavLINK	Micro Air Vehicle Link
MIT/LL	Massachusetts Institute of Technology / Lincoln Laboratory
MMS	Meta Monitoring System
MSL	Mean Sea Level
NASA	National Aeronautics and Space Administration
NavQ	Navigation Quality Service
NOS	Not Operationally Significant
NPCRA	Non-Participant Casualty Risk Assessment
OOL	Out-of-limits
OPT	Operations Planning Tool
OS	Operationally Significant
Pc	Probability of casualty
PNT	Positioning, Navigation, and Timing
PSU	Provider of Services for Urban Air Mobility
PtT	Proximity to Threat
R&D	Research and Development
RC	Remote control
RDT&E	Research, Development, Test, and Evaluation
RF	Radio Frequency
RFE	Radio Frequency Environment
RFI	Radio Frequency Interference
RFT	Remaining Flight Time
ROC	Risk of Obstacle Collision
RTL	Return-to-Launchpoint or Return-to-Land
RTRA	Real-time Risk Assessment
RUL	Remaining Useful Life
SAIL	Safety Assurance and Integrity Levels
SDSP-CD	Supplemental Data Service Provider – Consolidation Dashboard
SE	Safety Enhancement
SFC	Service, Function, and Capability
SME	Subject Matter Expert
SMS	Safety Management System

SOC	State of Charge
SORA	Special Operations Risk Assessment
sUAS	small Uncrewed Aerial System
SUS	System Usability Scale
TM	Telemetry service
TRL	Technology Readiness Level
UAM	Urban Air Mobility
UAV	Uncrewed Aerial Vehicle
UML	Urban Air Mobility Maturity Level
uORB	Micro Object Request Broker
URAF	UTM Risk Assessment Framework
USS	UTM Service Supplier
UTM	Uncrewed Aerial System Traffic Management
VCU	Virginia Commonwealth University
W	Watts
WIPS	Wind Information Prediction Service
WM	Wind Modeling service
WRF	Weather Research and Forecasting service
WxSt	Networked Weather Station service
xTM	extensible Traffic Management