Testing a Run-Time Assurance Framework Coupled with Integrated Risk Mitigation Capabilities for Autonomous Urban UAS Flights

Ersin Ancel
Aeronautics Systems Analysis Branch
NASA Langley Research Center
Hampton, VA, USA
ersin.ancel@nasa.gov

Andrew J. Moore
Dynamic Systems & Controls Branch
NASA Langley Research Center
Hampton, VA, USA
andrew.j.moore@nasa.gov

Steven D. Young
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, VA, USA
steven.d.young@nasa.gov

Evan T. Dill
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, VA, USA
evan.t.dill@nasa.gov

Cuong (Patrick) Quach
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, VA, USA
cuong.c.quach@nasa.gov

Kyle M. Smalling
Safety-Critical Avionics Systems Branch
NASA Langley Research Center
Hampton, VA, USA
kyle.m.smalling@nasa.gov

***Abstract*—The In-Time Aviation Safety Management System (IASMS) Concept of Operations (ConOps) envisions new capabilities to monitor, assess, and mitigate flight safety risks. Systems will be tailored to mission type, vehicle/equipage type, operational environment, and safety risk tolerance. Within an IASMS framework, several capabilities may be implemented spanning three operational phases (pre-flight, in-flight, and post-flight/off-line); consisting of lower level functions and information services which may reside on board the aircraft, on third-party server(s), and/or on ground/operator station(s). Each capability will be designed to produce and disseminate safety-relevant information; perform detection, diagnosis, and prediction of unsafe situations; and/or execute mitigation actions when hazardous events warrant such changes. This paper focuses on recent testing of airborne capabilities that demonstrate *inflight* aspects of the overarching concept for autonomous unmanned aircraft systems (UAS) operations in urban environments. A flight test architecture is described that applies run-time assurance principles (e.g., executes independent of the unassured autopilot), real-time risk assessment, and a technique to execute contingencies if necessary either automatically or via pilot intervention. Several tests using small UAS were conducted to verify the assured in-flight risk mitigation capability. The paper draws significantly from a larger NASA technical report and recent prior conference papers, providing additional details. Data are analyzed for two representative flights to illustrate the performance for various sequential and simultaneous hazards used during testing. During each automated flight, several hazards are encountered at various points along the flight path. At each point, the hazard is mitigated by the system, with the vehicle then continuing to subsequent points. The paper concludes with lessons-learned regarding relevant aspects of the overarching IASMS concept and how it may be updated and further advanced in the future.**

***Index Terms*—risk, population activity, unmanned aircraft systems**

I. Introduction

Based upon NASA's Aeronautics Research Mission Directorate Strategic Implementation Plan and recommendations by the National Academies, the System-Wide Safety Project has been investigating In-Time Aviation Safety Management System (IASMS) Concepts of Operations (ConOps) and enabling technologies. The IASMS ConOps envisions expanded access to safety-relevant data from a broad set of information sources. The data, in turn, enable more timely integrated analysis and predictive capabilities, improved real-time detection and alerting of domain-specific hazards, decision support, and new forms of automated and supervisory risk mitigation strategies. Tailored IASMS architectures, with both ground-based and onboard elements, are envisioned to be employed in the future based on the four-tuple of mission type, vehicle/equipage type, operational environment, and safety risk tolerance [1], [2].

Within an IASMS framework, three high-level capabilities (Monitor, Assess, and Mitigate) are envisioned spanning three operational phases (pre-flight, in-flight, and post-flight/off-line). The IASMS framework is further decomposed into lower-level Services, Functions, and Capabilities (SFCs), which may reside on board the aircraft, on third-party server(s), and/or on ground/operator station(s). These SFCs work together to produce and disseminate safety-relevant information, perform detection, diagnosis, and prediction of unsafe situations, and execute mitigation actions when hazardous events warrant such changes.

This paper reviews and expands upon previous IASMS research published at the DASC ([1]–[3]), at other AIAA forums ([4]–[6]), and in NASA Technical Memorandums ([7], [8]), with a focus on recent testing of airborne capabilities that demonstrate in-flight aspects of the overarching concept.

Over 100 flight experiments were conducted to test and further develop capabilities that detect and respond to hazards encountered during flight. First, safety hazards were monitored and assessed on board, and system-generated mitigation maneuvers were recorded (but not acted upon by the vehicle) with the objective of verifying the in-flight risk mitigation capability in the face of multiple hazards. Preliminary results for these tests were highlighted by Young et al. [5] and Ancel et al. [6]. Next, mitigation maneuver commands directed the aircraft in response to safety hazards (i.e., auto-mitigation). The second set of tests were described in great detail in Moore et al. [8] including the test architecture; which included commercial avionics, research avionics, and onboard software designed to detect, assess, and respond to hazards.

The test architecture for the flight campaign is comprised of several onboard SFCs constructed using run-time assurance (RTA) approaches, highlighted by Neogi et al. [9]: a novel method for assessing risks during flight (e.g., risk of loss of control), autopilot monitoring, proximity to threat monitoring (e.g., man-made structures), and conformance monitoring. The risk assessment function tracks a set of risk-related metrics and considers contingencies per each hazard type (i.e., navigation loss, command and control link loss, unsafe proximity to obstacle, and loss-of-control of the aircraft). An independent risk mitigation function tracks available autopilot flight mode changes, prioritizes among contingency options, and triggers execution of maneuvers in order to reduce risk. Execution may be performed automatically or via pilot intervention (i.e., supervisory).

The paper draws significantly from a larger NASA technical report [8] and the results of testing the RTA framework combined with in-flight risk mitigation (i.e., independent monitoring of an unassured COTS system via a highly assured system and intervening when/if established risk thresholds are exceeded during flight).

The remainder of the paper is organized as follows. The flight system software architecture is given in Section II. Representative flight data is given in Section III illustrating performance of the tested capabilities for various sequential and simultaneous hazards. Section IV provides summary conclusions with lessons-learned regarding relevant aspects of the overarching IASMS concept and how it may be updated and further advanced as part of future work.

II. Flight System Software Architecture

The onboard system was designed (as an RTA framework) to operate independent of the COTS autopilot and, supportive of both supervisory and automated mechanisms, to monitor, assess, and mitigate risk. As described in [5], [6], [8], and [9], the primary functions included Real-Time Risk Assessment (RTRA), auto-pilot monitoring, constraint monitoring, and contingency select/triggering (CST). RTRA performs integrated risk assessment considering data from several hazard-related monitors (e.g., battery, motors, Global Navigation Satellite System (GNSS), communications, population density, and loss-of-control).

Figure 1 provides the overall software architecture used during testing. The functions executing within this architecture are described in the following subsections.

A. Hazard Monitor Functions

The scope of tested hazard monitor functions is as follows:

1) Proximity to Threat: The previously developed Proximity to Threat (PtT) preflight planning service [10] is implemented as an onboard function to aid in hazard monitoring. Static obstacle boundaries of both real and artificial buildings and trees are loaded at mission start by the Research Ground Control Station (GCS), and the UAV position from the navigation is compared to these boundaries twice per second. 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 (set to 50 feet, or 15.24 meters, for these experiments). The Hazard Likelihood function considers such an alert as a 100% likelihood of a hazardous event and recommends a hover flight maneuver to the autopilot.[1]

2) Battery Health: A battery health monitor function is implemented using an electrochemical model of the propulsion battery [11]. The function continuously estimates the state of charge (SOC) and remaining flight time (RFT) based on battery temperature, voltage, and current draw. The resulting RFT is compared to a minimum value of 200 seconds set by the operator for the testing; an RFT falling below this value prompts a 100% hazard likelihood assessment with a recommendation of an immediate land maneuver sent to the autopilot.[1]

3) Pilot Radio Link Monitor: Pilot control integrity is monitored by comparing one channel of the 2.4 GHz R/C link to a minimum acceptable pulse duration (950 microseconds). Loss of this link prompts a 100% risk likelihood assessment with a recommendation of an assigned location landing maneuver to the autopilot.[1]

4) Global Navigation Satellite System (GNSS) Monitor: The number of received satellites (GPS and GLObalnaya NAvigatsionnaya Sputnikovaya Sistema (GLONASS) constellations) and horizontal position uncertainty (in terms of horizontal dilution of precision (HDOP)) are compared to preset warning levels (8 satellites and 5 meters) and failure levels (6 satellites and 10 meters). Failure-level navigation impairment prompts a 100% hazard likelihood risk assessment with a recommendation of an immediate land maneuver to the autopilot.[1]

5) Geospatial Conformance Monitor: Trajectory conformance to the intended flight volume is monitored by an independent highly assured conformance monitoring function (a.k.a, the SAFEGUARD system [12]). The function uses predefined geo-fence polygons that represent the boundaries

[1]These maneuvers correspond to ArduCopter flight modes as follows: Immediate Land is LAND mode, Assigned Land is RTL or Return to Launch mode, and Hover is POSHOLD or Position Hold mode. An entry of "None" means that no command should be issued to the autopilot, i.e., the current flight mode should continue.

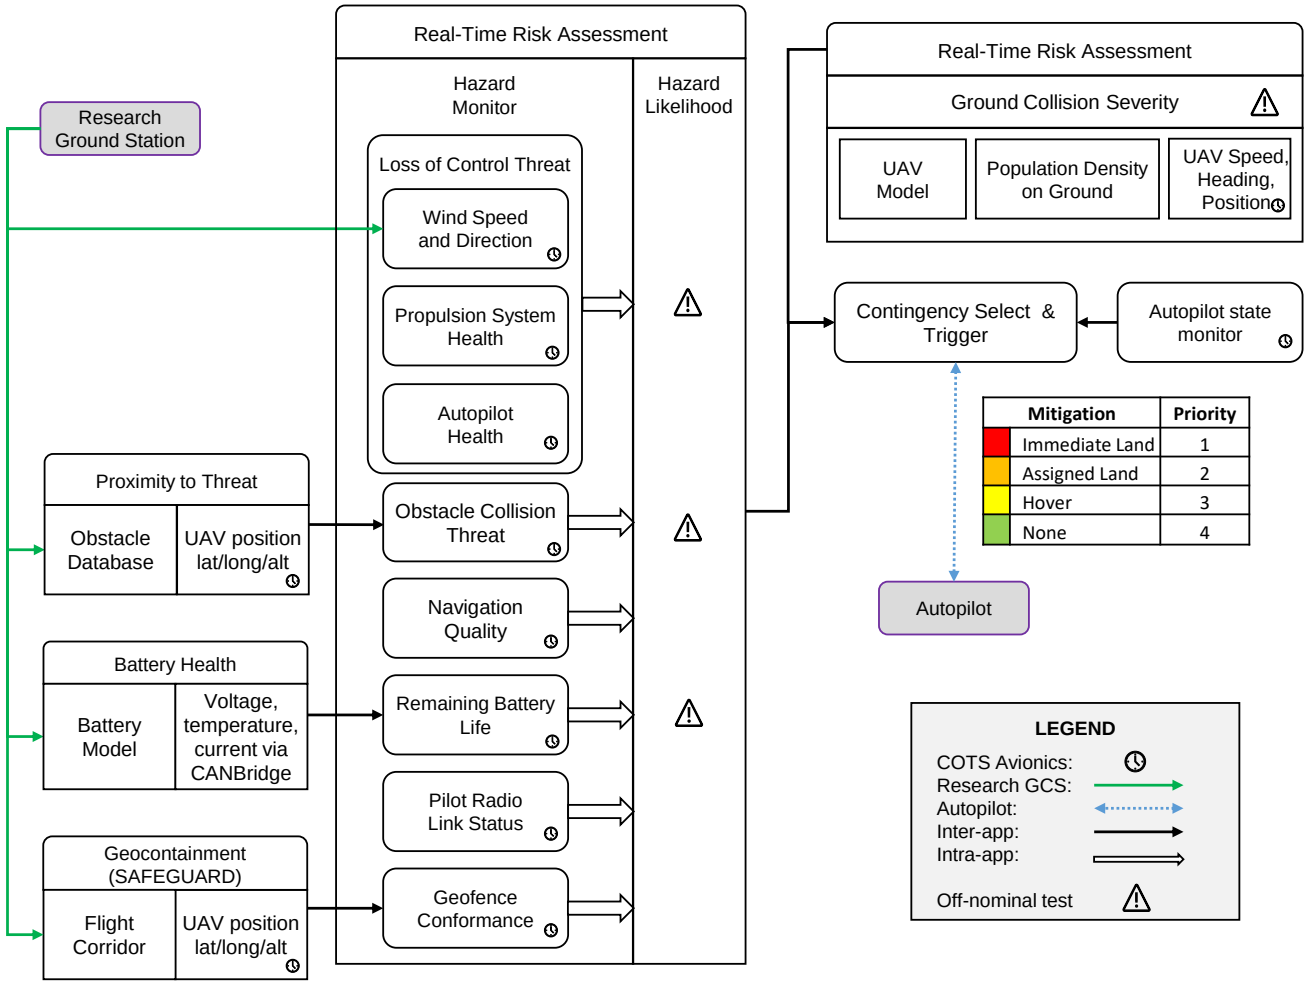

Fig. 1. Onboard software architecture for independent hazard monitoring, risk assessment, and contingency selection and triggering (i.e., auto-migitation) [8].

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), an immediate land command is issued to the autopilot.

6) 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) 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). These ESCs provide individual motor speed, temperature, voltage, and current information. 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 where warning state indicates one inoperative motor and failure state is triggered due to more than one inoperative motor. For brevity, propulsion system failure levels and their determinants are not listed here; see [6] for a detailed description of this SFC.
- Wind speed and direction: Three levels of wind speed (none, low, high) are reported. High winds are those greater than 5 m/s. The presence of high winds is assumed to increase the likelihood of loss of control threat.

B. Risk Assessment Functions

Risk assessment is immediate for all hazard monitors except the Loss of Control Threat monitor. As shown in Table I, navigation loss (Navigation Quality monitor) and available motor battery status (Battery Health) are judged as most acute and an immediate land maneuver is recommended; a lost command and control link (Pilot Radio Link Status) 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 to the autopilot.

TABLE I
MITIGATION ACTIONS FOR TRACKED HAZARDS [8]

Hazard Detected	Mitigation	Priority
Navigation Loss	Immediate Land	1
Battery Loss	Immediate Land	1
Command/control link loss	Assigned Land	2
Proximity to obstacle	Hover	3
Loss of control	see Table II	

The loss of control condition is continuously monitored via the Loss of Control Threat Monitor at a 1 Hz rate, and as part of the assessment, the RTRA function projects the point of ground collision based on the wind speed and the aircraft 3D location, heading, and flight speed. Population density, sheltering effects (i.e., whether the population is protected indoors), casualty impact area, and the kinetic energy at impact are evaluated to estimate the probability of a casualty (P_C) caused by the falling vehicle [13], [14]. A recommendation which minimizes harm to people on the ground is determined by considering two assessments, given in Table II. The severity ranking steps (minimal, minor, major, and catastrophic) are based on quartiles of P_c values and a probability of loss of control ranking, classified as improbable, remote, probable, and frequent[2]. 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.

TABLE II
LOSS OF CONTROL RISK TABLE WITH ASSOCIATED MITIGATION ACTIONS [8].

	Loss of Control Severity			
Loss of Control Likelihood	Minimal	Minor	Major	Catastrophic
Frequent				
Probable				
Remote				
Improbable				

Key: Immediate Land | Assigned Land | None

A static wind from the east with speed of 8 m/s at 10 m altitude was set at the start of the test flights. In this high wind condition, it is assumed that 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 in Table II). 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.

[2]Mapping of loss of control likelihood P_{LOC} values to likelihood ranking: Improbable $0 \leq P_{LOC} < 0.01$, Remote $0.01 \leq P_{LOC} < 0.1$, Probable $0.1 \leq P_{LOC} < 0.5$, and Frequent $0.5 \leq P_{LOC} \leq 1$.

C. Risk Mitigation Functions

The Contingency Select and Trigger (CST) function continuously collects maneuver recommendations from the hazard/risk monitor and assessment functions, as well as from an autopilot state monitor. As multiple hazards may occur at any particular time, this function prioritizes the most urgent safety action consistent with the current flight state (Table I) 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; and a recommendation to hover in place is prioritized last. If no hazard alert reaches the Contingency Select and Trigger function, a No-Operation (NOOP)/None status is logged and no maneuver command is issued to the autopilot.

D. Maneuver Execution Verification

An Autopilot Monitor (APMon) function, developed using formal methods software development tools [15], keeps track of the current autopilot flight mode and assures that a proposed mitigation action command from the CST function is executable and valid given the current context of the flight. A switch to Immediate Land mode is allowed from any autopilot state; Hover mode is allowed as long as the autopilot reports healthy navigation and velocity; Assigned Land mode is allowed as long as the autopilot reports healthy navigation and velocity and valid landing coordinates; and a return to the waypoint flight plan (AUTO mode in ArduCopter) 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 R/C 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.

III. FLIGHT TESTING AND SAMPLE RESULTS

This section focuses on the flight testing conducted to evaluate numerous in-flight SFCs intended to provide a highly assured capability to automatically mitigate risk during flight (see [6] and [5]). As previously discussed, the flight testing was conducted in two phases. The first phase (61 flights conducted between August 2021 and July 2022) focused on verification of research software without mitigation actions relayed to the autopilot. The second phase of testing (54 flights performed between July 2022 and October 2023) involved

execution of mitigation actions that were proposed by the onboard research hardware.

A. Flight Test Range

Flights were conducted at NASA Langley Research Center's City Environment Range Testing for Autonomous Integrated Navigation (CERTAIN) test range, free from structures and ground population. A set of virtual buildings was added to pose obstacle collision hazards and a set of virtual crowds was added to pose population overflight risk. Fig. 2 depicts the CERTAIN flight range where the blue shaded triangle (approximately 0.6 km per leg) is the geo-fence limit used by SAFEGUARD, the population is shown as white circles, and structures are shown as colored 3D geometries.

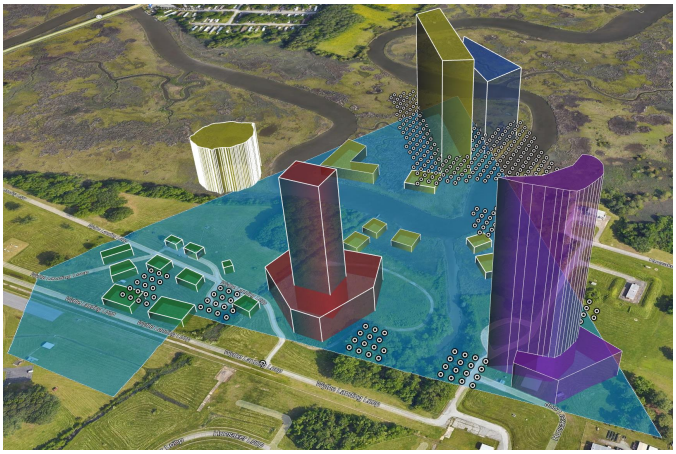

Fig. 2. NASA Langley Research CERTAIN flight range and virtual urban structures used during testing.

B. Research Aircraft

The experimental aircraft for these sets of tests was a modified Tarot T18 Octocopter frame outfitted with COTS hardware as well as research sensors, computers, and associated electronics attached to the frame and on payload trays. The research equipment included two separate Intel® Next Unit Computing (NUC) nodes, radio-controlled cut-off relays that power down the research system and sever communications to the autopilot as a safety measure, the SAFEGUARD module, and other hardware (for a full list of the hardware and software suite, see [8]). Ground support equipment included a dedicated GCS computer running Mission Planner software and a separate Linux computer running NASA-developed research GCS that tracks and displays onboard research payload activity. Fig. 3 depicts the modified Tarot T18 Octocopter used for majority of the flight testing reported in this paper.

C. Phase I Testing: Flight Tests with Logged Mitigation Actions

A total of 61 flight tests were conducted with airborne research avionics disconnected from the autopilot. Performance of the hazard monitoring, risk assessment, APMon, and CST functions was 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 mitigation actions. Ancel et al. [6] provide results from a sample Phase I test where the central objective of the test was to verify that automatically generated mitigation actions in the face of multiple hazards were prioritized correctly.

Fig. 3. Modified Tarot T18 Octocopter Research Aircraft [8].

D. Phase II Testing: Flight Tests with Executed Mitigation Actions

With confidence gained from the logging flights that mitigation actions are timely and effective, autopilot control via the independent system was tested. For this set of flights, in addition to logging the outputs of functions, the recommended command maneuvers actively controlled the vehicle autopilot (i.e., closed-loop operation). Waypoint-based flight paths were constructed so that aircraft would encounter a series of hazard conditions to trigger mitigation actions resulting in autopilot mode change (e.g., Hover). After each hazard-triggered mode change, the pilot would recover aircraft control, direct it along the flight path past the hazard zone, and then re-initiate autonomous vehicle traversal along the flight path. This enabled testing multiple hazard encounters and mitigation actions during the same flight. Tests were not intended to evaluate performance of the autopilot after the desired mode change occurred.

A detailed analysis of the hazards encountered and the monitor/assess/mitigate safety response is included for two research flights next. See [8] for similar analyses of an additional seven research flights.

1) Sample Flight 097 (F097): The sample flight F097 trajectory and its autopilot modes are shown in Figs. 4, 5, and 6. During this flight, the aircraft was flown over the flight range with the relay switch in the closed position, which allowed the CST application to transmit mitigation actions to the autopilot[3].

[3]A dedicated channel of the pilot R/C controller switches an onboard research relay that allows or interrupts commands sent from the research computer to the autopilot. This relay switch was implemented as an additional safety assurance measure and to comply with NASA Class C software requirements.

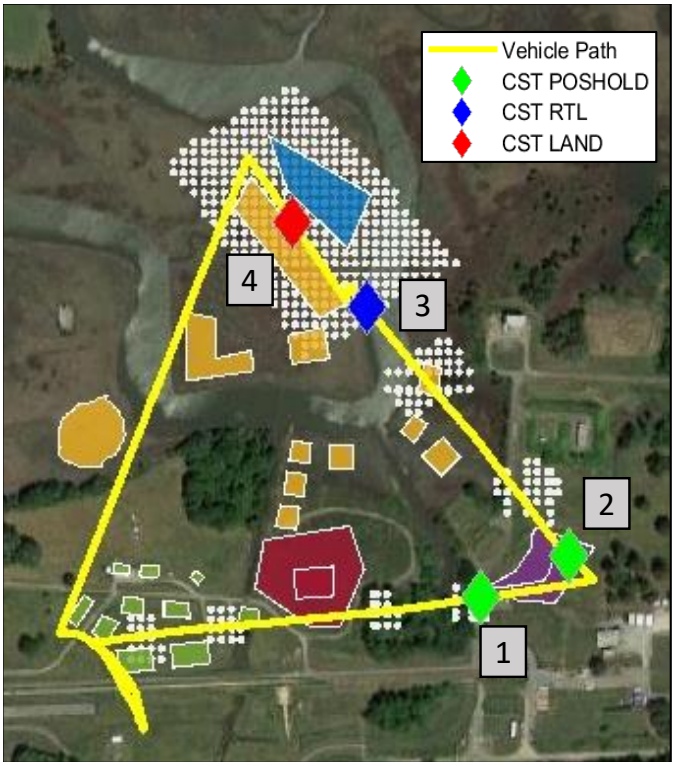

Fig. 4. Flight trajectory and encountered hazardous conditions for F097 [8].

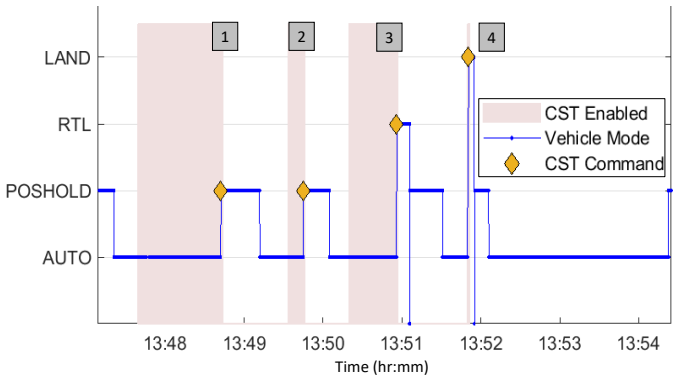

Fig. 5. Vehicle autopilot mode and executed mitigation actions for F097 [8].

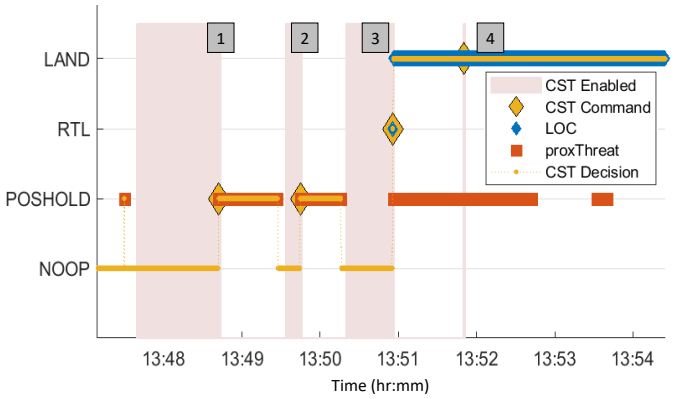

Fig. 6. CST-commanded autopilot maneuvers and observed hazards for F097 [8].

Fig. 4 provides the vehicle flight path and its position relative to several simulated hazards, i.e., virtual buildings shown as colored polygons and ground population in varying densities shown in white circle clusters within the flight range. Also shown are the vehicle's automated mitigation responses: two Hover mitigation maneuvers (green diamonds for POSHOLD, labeled "1" and "2"), one Assigned Land maneuver (blue diamonds for RTL, labeled "3") and one Immediate Land maneuver (red diamond for LAND, labeled "4")

Figs. 5 and 6 provide autopilot mode and CST decision making sequences for flight F097. Specifically, Fig. 5 depicts the automation status (i.e., autopilot mode) with a blue line, mitigation maneuvers issued by the CST with orange diamonds, and status of the relay switch that allows CST commands to be executed, in solid pink bands[4].

Fig. 6 contains the observed hazard condition (i.e., loss of control in blue and proximity to threat in red bands) and the arbitrated/prioritized CST decision (in orange line). Note that during the concurrent occurrence of loss of control and obstacle collision hazards, the CST Decision prioritizes the respective mitigation actions as intended/designed.

Test Point (1): Upon launch, the vehicle was manually navigated near the first waypoint in POSHOLD mode and, subsequently, was switched to AUTO mode upon reaching the flight path. As the vehicle began flight along the bottom leg, the research relay was closed, initiating safety-enhanced autonomous flight (i.e., CST Enabled given in solid pink shading in Fig. 5). First, the vehicle flew over a few buildings and populated areas located on the bottom of the path; however, the altitude and observed P_C values were above the threshold values to trigger a mitigation action. Next, an unsafe proximity to a virtual building condition (solid green diamond labeled "1" in Fig. 4 and the matching solid orange diamond in Figs. 5 and 6) was detected by the Proximity to Threat monitor. Subsequently, a Hover maneuver was recommended by the RTRA function and the CST mitigation function issued a Hover/POSHOLD command to the vehicle, indicating a successful autonomous mitigation execution. The pilot then opened the relay to inhibit further research-generated commands reaching the autopilot and used default autonomy to fly the aircraft through the corridor (a very close proximity to a virtual building) to the next waypoint at the bottom right of the triangular flight path. The pilot then closed the relay to re-initiate autonomous waypoint traversal with the research system engaged.

Test Point (2): Within a few seconds into this stage of autonomous flight (AUTO mode) with CST enabled, the Proximity to Threat monitor flagged another close approach hazard, triggering a second Hover event (solid green diamond

[4]CST commands given with orange diamonds can only transmitted to the autopilot when the research relay is in closed status, where the CST is enabled, given in pink bands.

labeled "2" in Fig. 4). The pilot then took control and flew the aircraft out of the virtual danger zone (depicted by the red proximity threat bar given in Fig. 6). The pilot closed the research relay to enable safety-enhanced autonomy within a few seconds.

Test Point (3): Enroute to the waypoint at the top of the triangle, the aircraft traversed over lower population and short buildings without triggering mitigation actions (hazard-free) until approaching overflight of a densely populated area. With loss of control likelihood elevated by a constant high wind condition, the projected ground collision severity computed by the RTRA – Ground Collision Severity function exceeded a safe level and an Assigned Land maneuver was recommended. The CST mitigation function issued an Assigned Land (RTL) command to the vehicle (solid blue diamond labeled "3" in Fig. 4). Though a Proximity to Threat alert was also occurring simultaneously, 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.

Test Point (4): After about a further half minute of flight, the pilot re-engaged the research-to-autopilot link via the relay switch. In the interim, two mitigation recommendations were being issued once per second by the hazard assessment function, RTRA, in response to hazard alerts. First, the vehicle had moved to an even higher loss of control severity condition due to flying over a densely populated area (depicted as a collection of white circles in Fig. 4), so that an Immediate Land recommendation was pending from the RTRA – Ground Collision Severity function. Second, at this location in a narrow corridor between two buildings (orange and blue buildings in Fig. 4), a Hover recommendation was pending from the Proximity to Threat function. Once the control relay was closed, the CST function acted on the higher priority Immediate Land recommendation and directed the vehicle into a descent maneuver (solid red diamond labeled "4" in Fig. 4). Upon confirming the successful mitigation execution, the pilot reclaimed 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.

2) Sample Flight 115 (F115): Figs. 7, 8, and 9 show the details of F115 in which the monitor/assess/mitigate flight safety architecture responded six times to hazardous conditions. In this flight, mitigation actions again resulted from obstacle collision 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, given in orange, is triangular and resides in a corridor with virtual buildings and ground populations. The mitigation maneuvers are enumerated "1" through "7" with cyan, green, blue, and red diamonds in Fig. 7.

Similar to the previous discussion of flight F097, Fig. 8

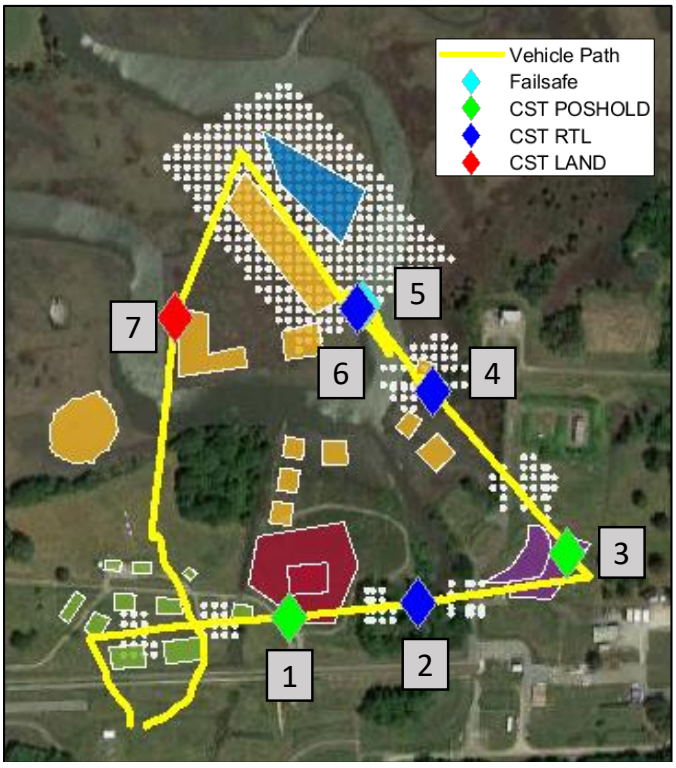

Fig. 7. Vehicle autopilot mode and executed mitigation actions for F115 [8].

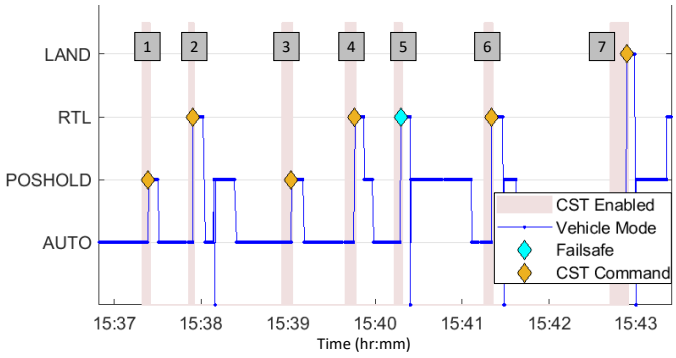

Fig. 8. CST-commanded autopilot maneuvers and observed hazards for F115 [8].

shows the vehicle autopilot mode with a blue line, CST-issued mitigation maneuvers with orange diamonds, and the relay status indicating the CST's ability to change the autopilot modes in solid pink bands. Fig. 9, in turn, provides observed hazard conditions (i.e., loss of control in blue bands, obstacle collision in red bands, and instances of motor battery failures with green stars) and the arbitrated/prioritized CST decision (orange line). As before, orange diamonds in Fig. 9 indicate mitigation maneuvers issued by the CST function throughout the flight.

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:

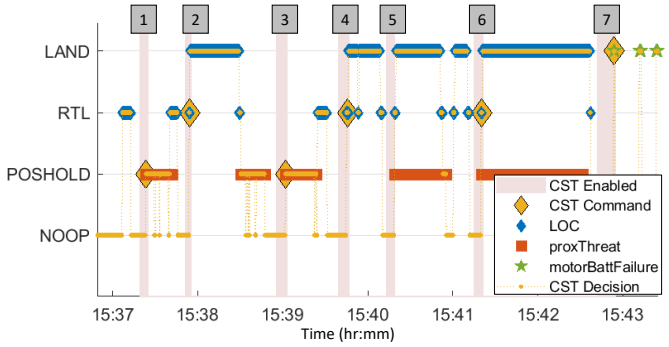

Fig. 9. Vehicle autopilot mode, proposed and executed mitigation actions for F115 [8].

Test Point (1): After the vehicle took off and climbed, it proceeded to the first waypoint of the triangular flight path and began eastward flight. Similar to the F097 scenario, the vehicle flew over short buildings and sparsely populated areas without triggering a mitigation action. Next, an obstacle collision threat was issued by the Proximity to Threat monitor when distance from a virtual building was less than the 50 foot threshold, marked with a green diamond and "1" in Fig. 7. The RTRA recommended a Hover maneuver and the CST mitigation function issued a Hover command to the vehicle.

Test Point (2): Further east on the flight path, with a wind blowing to the east, a population casualty threat was detected by the RTRA's 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 (refer to Table II) dictates issuing an Assigned Land mitigation, and since the CST mitigation function was receiving no competing recommendations, it initiated an Assigned Land maneuver.

Test Point (3): At the start of northwesterly flight, the threat of collision with a building resulted in a Hover maneuver.

Test Point (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.

Test Point (5): As the vehicle neared a virtual building, an unsafe proximity to an obstacle was expected, but a transient low signal level for the 900 MHz telemetry link to the ground control station was first detected by the autopilot. The native autopilot failsafe maneuver for GCS link loss (Assigned Land) was automatically activated 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 resumed.

Test Point (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 RTRA – Ground Collision Severity function and a Hover maneuver triggered by an alert from the Proximity to Threat function. The CST 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. 7) and then to the southwest.

Test Point (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 RTRA – 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.

IV. Summary, Observations, and Future Work

Over the past six years, NASA and its partners have been investigating and advancing the IASMS concept [1], [4], [16], which posits a set of capabilities intended to enable more timely mitigation of safety risk during operations. A subset of these capabilities would be used in-flight and be installed on board aircraft. Two such capabilities are described here along with detailed analysis of how they were tested using small UAS and how they performed. These capabilities are: (1) a run-time assurance capability that operates similar to what is described in Neogi et al. [9] and is suggested for highly autonomous aircraft that utilize unproven/unverified autopilots (i.e., the assurance level of the autopilot software is less than the assurance level of the RTA monitor) and (2) a real-time risk management capability that monitors a set of safety-related metrics (e.g., hazard states), assesses risk, and if/when necessary selects and triggers contingency maneuvers. A few general observations follow. These are based on testing described here and in prior publications which provide additional details.

Observation 1: Although implemented and tested as an independent (RTA-like) system, each function (Fig. 1) may also be useful as embedded within the autopilot. The choice is driven by the design assurance level requirement of the autopilot vice the flight system as a whole, as well as the types of hazards expected (e.g., if operating in urban areas near vertical structures, a function like PtT may be essential within the autopilot to reduce obstacle collision risk).

Observation 2: Test scenarios and some findings are constrained by the equipment used and the available test environment. For example, the COTS autopilot had a limited number

of alternate flight modes to choose from for contingency selection and triggering (e.g., LAND, HOVER, RTL), and for a full implementation of an RTA framework, an independent highly assured controller would be available for situations where there was a controllability failure within the autopilot. Likewise, a combination of onboard COTS and NASA-developed hardware/software was used in concert to capture the aircraft health status, which partially hampers dissemination and repeatability of the research conducted here. Finally, the test range was limited in terms of vertical structures, population density, and wind limits. This was overcome by using virtual constructs, which yielded good verification results yet did not achieve full validation.

Observation 3: Although several complex software functions were implemented, real-time performance was easily achieved using the onboard Intel® NUCs. Outputs were generated at or better than a 1 Hz rate and no significant latencies or errors were encountered during testing. Likewise, autopilot responsiveness to the CST-generated commands was nearly immediate and as expected in almost all cases. It is unknown whether this performance would hold for alternate research systems (e.g., RTA framework) on other, non-NUC platforms or autopilots. However, further testing of the system presented here is ongoing on a different UAS platform via a partner activity.

In terms of ongoing and future work, there are multiple directions being pursued. One of particular interest and the subject of 2024 simulation and flight testing is coordination of independent decision-making functions that may be operating simultaneously onboard a highly automated vehicle. Two examples are being initially investigated: (1) the coordination of behavior of a system such as described here with behavior of a detect-and-avoid (DAA) system and (2) the application of techniques based on artificial intelligence/machine learning. The latter is considering the tradeoffs of using these techniques to optimize across safety and efficiency goals as compared to heuristic methods, which may be easier to assure, but, at times, can be limiting in terms of operational flexibility and efficiency. A second area of ongoing work investigates how to apply this construct to new operational domains and vehicle types. These activities are intended to inform guidance and standards with respect to the tailoring of IASMS designs (i.e., not all of the capabilities or functions described in this paper would be required for some domains). A third direction is related to automatically applying flight data and observations to post-flight capabilities, looking for precursors, anomalies, and trends that may only be distinguishable when considering multiple similar flights.

In conclusion, this paper summarizes recently completed flight testing ([5], [6], [8]) by providing an overview of tests specifically aimed at demonstrating the ability to detect and respond to a set of hazards encountered during flight (both in automated and supervisory modes). A step-by-step walkthrough of created scenarios also provides a means that others may use to increase the efficiency of conducting such complex tests in the future.

ACKNOWLEDGMENT

The work presented here was supported by the System-Wide Safety project in NASA's Aeronautics Research Mission Directorate's Airspace Operations and Safety Program (AOSP). The authors would like to thank the extended team of researchers, engineers, developers, and operators who contributed to this work, at NASA's Langley and Ames Research Centers.

REFERENCES

[1] S. D. Young, C. Quach, K. Goebel, and J. Nowinski, "In-time safety assurance systems for emerging autonomous flight operations," in *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*, (London, UK), pp. 1–10, Sept. 2018.

[2] P. Krois, "An approach for identifying IASMS services, functions, and capabilities from data sources," in *2021 IEEE/AIAA 40th Digital Avionics Systems Conference (DASC)*, (San Antonio, TX), pp. 1–10, Oct. 2021.

[3] E. Ancel, T. Helsel, and C. 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)*, (San Diego, CA), pp. 1–8, Sept. 2019.

[4] K. K. Ellis, P. Krois, J. H. Koelling, L. J. Prinzel, M. D. Davies, and R. W. 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*, no. AIAA-2021-1978, American Institute of Aeronautics and Astronautics, Jan. 2021.

[5] S. D. Young et al., "Flight testing in-time safety assurance technologies for UAS operations," in *Aviation Forum 2022*, no. AIAA 2022-3458, (Chicago, IL), American Institute of Aeronautics and Astronautics, June 2022.

[6] E. Ancel et al., "Design and testing of an approach to automated in-flight safety risk management for suas operations," in *Aviation Forum 2022*, no. AIAA 2022-3459, (Chicago, IL), American Institute of Aeronautics and Astronautics, June 2022.

[7] S. D. Young et al., "Architecture and information requirements to assess and predict flight safety risks during highly autonomous urban flight operations," NASA TM-2020-220440, Jan. 2020.

[8] A. J. Moore et al., "Testing of advanced capabilities to enable in-time safety management and assurance for future flight operations," NASA TM-20230018665, Apr. 2024.

[9] N. A. Neogi, S. D. Young, and E. T. Dill, "Establishing the assurance efficacy of automated risk mitigation strategies," in *Aviation Forum 2022*, no. AIAA 2022-3538, (Chicago, IL), American Institute of Aeronautics and Astronautics, June 2022.

[10] L. Spirkovska, I. Roychoudhury, and C. C. Quach, "Computing proximity to threat along uncertain trajectory to support urban air mobility," NASA TM-20240000892, Feb. 2024.

[11] K. W. Eure and E. F. Hogge, "Mathematical characterization of battery models," NASA TM-20205008059, Jan. 2021.

[12] E. T. Dill, R. V. Gilabert, and S. 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)*, (London, UK), pp. 1–8, Dec. 2018.

[13] E. Ancel, F. Capristan, J. V. Foster, and R. Condotta, "Real-time risk assessment framework for unmanned aircraft system (UAS) traffic management (UTM)," in *Aviation Technology, Integration, and Operations (ATIO) Conference*, no. AIAA-2017-3273, (Denver, CO), American Institute of Aeronautics and Astronautics, June 2017.

[14] E. Ancel, F. M. Capristan, J. V. Foster, and R. C. Condotta, "In-time non-participant casualty risk assessment to support onboard decision making for autonomous unmanned aircraft," in *AIAA Aviation 2019 Forum*, no. AIAA-2019-3053, (Dallas, TX), American Institute of Aeronautics and Astronautics, June 2019.

[15] L. Pike, A. Goodloe, R. Morisset, and S. Niller, "Copilot: A hard real-time runtime monitor," in *Runtime Verification* (H. Barringer, Y. Falcone, B. Finkbeiner, K. Havelund, I. Lee, G. Pace, G. Roşu, O. Sokolsky, and N. Tillmann, eds.), (Berlin, Heidelberg), pp. 345–359, Springer Berlin Heidelberg, 2010.

[16] National Academies of Sciences, Engineering, and Medicine, *In-Time Aviation Safety Management: Challenges and Research for an Evolving Aviation System*. National Academies Press, Mar. 2018.