In-Time Safety Assurance Systems for Emerging Autonomous Flight Operations

Steven D. Young and Cuong (Patrick) Quach
NASA Langley Research Center
Hampton, VA, USA

Kai Goebel and Jessica Nowinski
NASA Ames Research Center
Mountain View, CA, USA

***Abstract*— As aviation adopts new operational paradigms, new vehicle types, and new technologies to broaden airspace capability and efficiency, maintaining a safe system will require recognition and timely mitigation of new safety issues as they emerge and before they become hazards. A shift toward a more predictive risk mitigation capability becomes critical to meet this challenge. In-time safety assurance comprises monitoring, assessment, and mitigation functions that proactively reduce risk in complex operational environments wherein the interplay of hazards may not be known, and cannot be accounted for at design time. They also can help to understand and predict emergent effects caused by the increased use of automation or autonomous functions that may exhibit unexpected non-deterministic behaviors. The envisioned monitoring functions can observe these behaviors and apply model-based and data-driven methods to drive downstream assessment and mitigation functions, thereby providing a level of run-time assurance. This paper presents a preliminary conceptual design of such an in-time safety assurance system for highly-autonomous aircraft operating at low altitudes near and over populated areas. Research, development, and evaluation tests are initially aimed at public-use surveillance missions such as those needed for infrastructure inspection, facility management, emergency response, law enforcement, and/or security. A longer term goal is to support transportation missions such as medical specimen delivery and urban air mobility. Safety-critical risks initially addressed within the system concept were identified in previous work by NASA and others in industry. These include: flight outside of approved airspace; unsafe proximity to people or property; critical system failures including loss of link, loss or degraded positioning system performance, loss of power, and engine failure; loss-of-control due to envelope excursion or flight control system failure; and cyber-security related risks.**

Keywords—risk mitigation; hazard monitoring; real-time assessment; unmanned aircraft; information services; model-based predictive capability

I. BACKGROUND

Commercial airline operations have maintained an exemplary safety record in recent years. This is achieved by a plethora of safety assurance processes that have been applied and evolved over many decades. These processes can be viewed at a high level as two cyclical loops that execute continuously and simultaneously to mitigate risk (Fig. 1). Both loops execute the same functions with many humans and systems involved. However, they operate on very different time-scales. The 'design' loop (on the left) can take anywhere from months to decades to reduce a safety risk that has been identified. Examples of this are the Traffic Collision Avoidance System (TCAS) [1] and the Terrain Awareness Warning System (TAWS) [2]. Both of these systems have been shown to be very powerful risk-reduction technologies. But development and widespread implementation took many years. Whenever technology-based 'design' mitigations such as these are employed, extensive system development time and effort is required (including verification, validation testing, and certification), as is time and effort to develop procedures for using the system and training to execute these procedures in relevant conditions. Any attempt to speed up this process runs the risk of overlooking an unintended outcome or consequence that ultimately reduces safety.

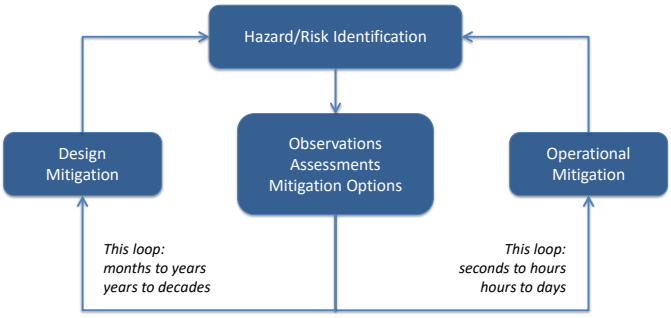

Fig. 1. Safety assurance loops and timescales.

On the other hand, the 'operational' loop (on the right in Fig. 1) can take anywhere from seconds to a few days to mitigate safety risk. This loop involves human observers who are participating in the operation (e.g., pilots and air traffic controllers) and are supported by technology that is designed to issue warnings or alerts of impending unsafe conditions. Many systems have been designed and implemented to help provide relevant observations (shown in the middle of Fig. 1), but humans are responsible for making most of the 'operational' assessments based on warnings or indicators issued by these systems. Likewise, mitigation options are most-often trained procedures for the humans to execute under certain conditions. Referring back to the examples of TCAS and TAWS; once these systems were implemented, they became part of the 'operational' risk mitigation loop. They generate warnings or alerts; pilots act on these alerts in

accordance with their training. The entire cycle, from warning to action may take only a few seconds to avoid an accident or close-call, thereby mitigating risk.

Fig. 2 illustrates the large number and span of tools and technologies that support the safety assurance process in the domain of airline-like operations. Here, the axes of the figure consider three dimensions of the multi-dimensional design space: (1) level of system/human authority and autonomy; (2) the types of hazards or risks; and (3) the aerospace application domain. We define these domains as the three-tuple of vehicle type, mission, and operational environment. As shown, for the airline-like domain, many real-time (RT) warning systems support pilots, controllers, and dispatchers (Level 1). Each of these systems represents a complex hardware/software system with associated training and procedures (e.g., TCAS and TAWS). However, there are fewer systems that function at the higher levels of authority/autonomy, where systems are making their own assessments, and in rare cases, taking their own actions to mitigate safety risk.

In Fig. 2, Level 2 represents an automated/autonomous control function that must be engaged by a human and can be disengaged by a human. An example is the auto-brakes function on aircraft. The pilot can enable this function to help mitigate the risk of runway over-run during landing roll-out. In contrast, Level 3 is a function that 'automatically' takes action, but the human can intervene to over-ride or takeover. An example is envelope protection systems. These are always active, but the pilot can push through (over-ride) if he/she feels envelope excursion is warranted. Level 4 is the rarest type in today's aircraft, representing functions that once armed or activated, have some degree of authority that cannot be over-ridden. The automatic ground collision avoidance system is an example from military aviation. Once this system is armed, the aircraft will maneuver to avoid an accident (i.e., roll wings-level and climb) if it determines a collision with the ground is eminent [3]. While the examples given above are onboard aircraft, there are also many tools supporting air traffic controllers and individuals at the airline operations centers. Here, there are even fewer examples when moving to higher levels of system authority and autonomy.

Consider also the floor of Fig. 2. This region corresponds to the 'design' loop of Fig. 1. During current airline operations, a great deal of data and information are recorded and used for subsequent analysis (i.e., not in real-time). This analysis can, and does, identify unknown risks and developing trends as well as come up with solutions to mitigate known risks. Many organizations are involved in this process, including manufacturers, airline operators, airport facility management, the FAA, the National Transportation Safety Board (NTSB), the Flight Safety Foundation (FSF) and the Commercial Aviation Safety Team (CAST). The best example is perhaps the CAST, which over the course of two decades of consensus-based data-driven work across government and industry sectors, has published and implemented more than 200 Safety Enhancements (SEs) [4]. In the U.S., much of the aviation data are collected and maintained as part of the Aviation Safety Information Analysis and Sharing (ASIAS) system [5].

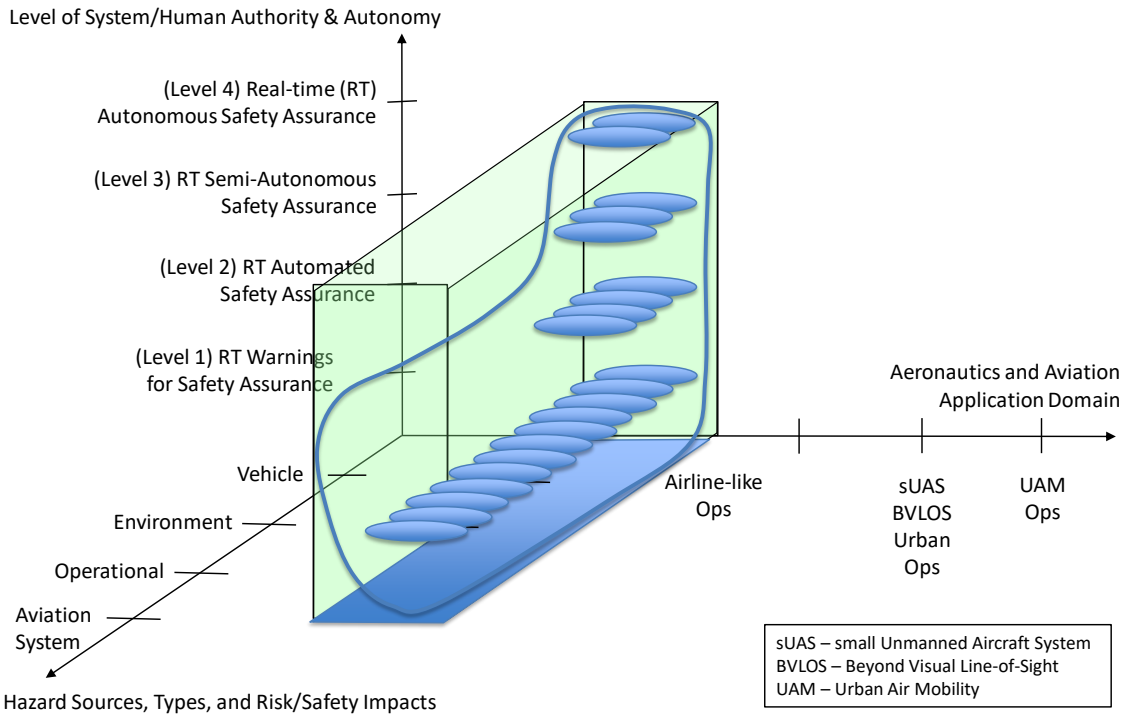

Fig. 2. Safety assurance during airline-like operations.

ASIAS manages a vast repository of data and information and when used in conjunction with data mining tools has been shown to be effective in finding vulnerabilities regarding safety margins such as anomalies and trends that may otherwise be imperceptible to human analysts. ASIAS is continually working to expand its ability to detect safety issues in the U. S. National Air Space (NAS) by taking advantage of an increasing number of data sources and fusing data from multiple sources to provide a clearer picture of context-specific confluence of events or factors that indicate a vulnerability. Findings have led to design changes to mitigate risk, including technological, procedural, and training changes. Findings have also led to improvements in business practices (e.g., maintenance).

While ASIAS continues to evolve as a powerful tool toward safety improvement, one drawback remains (as indicated on the left side of Fig. 1). This process can take anywhere from months to years to decades to get from hazard and risk identification to risk mitigation.

Further, as aviation adopts new technologies, new vehicles, and new missions to enhance the capacity, efficiency, and uses of the airspace, maintaining a safe system will require timely recognition and mitigation of safety issues as they emerge. A shift toward more prognostic hazard identification and proactive risk mitigation will become critical to maintaining safety. 'Big-data' techniques, such as those already employed by ASIAS, combined with on-line machine learning and model-based techniques can serve not only to identify unknown unknowns, but also to predict loss of safety margin.

This notion of 'in-time' safety assurance posits to overcome some of the limitations of current processes while enabling such a proactive capability. More specifically, the R&D seeks to answer two questions:

(1) Can an ASIAS-like 'big-data'-driven process be automated/sped-up to provide timely information to agents operating in the system? (i.e., can it identify, observe, assess, and mitigate risks in seconds to a few days, moving this process to the operational loop?)

(2) Can such a process be designed and implemented for emerging domains (e.g., see right side of Fig. 2) and capture risks by monitoring for safety as these operations come online, rather than rely on a forensic retroactive approach to hazard identification?

II. IN-TIME SAFETY ASSURANCE CONCEPT

NASA's Aeronautics Research Mission Directorate describes a strategic plan for investigating and advancing in-time safety assurance capabilities [6]. Within this plan, it is envisioned that advanced safety assurance tools can be introduced to take advantage of increasing availability of aviation system data. The identification of safety issues will focus on scaling currently available data mining methods to process a broad range of data; thereby enabling a disparate set of high-fidelity model-based information services that can inform and track the changing nature of risk during flights. By using these models and information services, combined with increased speed and accuracy of analysis tools, safety assurance can progress toward more timely identification of precursors to emerging safety issues.

Further, the vision implies that system-of-systems modeling, prognostic tools, and run-time assurance techniques will enable in-time safety assurance capabilities at all levels, spanning from the vehicle-level to the airspace system level. To achieve this vision, R&D is suggested along four themes:

(1) Monitoring – Development of information services and derived architectural requirements to support comprehensive safety monitoring through acquisition, integration and quality assurance of heterogeneous data coming from a diverse set of sources (including the vehicles); recognizing that some data may require protections that de-identify the source and defend against corruption by unauthorized or unauthenticated sources.

(2) Assessment – Development of tools and techniques to improve the accuracy and integrity of timely detection, diagnosis and predictive capability regarding risk and hazard states. Assessment technology should be capable of spanning hazard types to judge how overall safety margin is changing based on context and cascading event sequences, as well as based on longer-term trends that can become evident with access to historical data maintained by the monitoring functions.

(3) Mitigation – Development of methods for multi-agent or automated planning and execution of timely responses to hazardous events or event sequences when/if safety margins deteriorate below acceptable levels.

(4) Implementation, experimentation, demonstration, cost-benefit analyses, and feasibility assessments of these new capabilities.

A recent report published by the National Academies outlines challenges that must be addressed in development of a similar but broader capability – an In-time Safety Assurance Management System (IASMS) [7]. This capability is defined within the larger context of SMS [8], which covers all aspects of managing safety including organizational structures, accountabilities, policies, procedures, and business practices. As reported in [7], the committee's vision for an IASMS is summarized in the following recommendation:

> "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."

The report has both validated NASA's general approach toward developing tools and techniques underlying an in-time safety assurance capability and also served to prioritize early R&D efforts. Indeed, many of the technical challenges outlined

in the report [7] are addressed by NASA's initial plans within the System-Wide Safety project, including identifying and prioritizing risks, developing in-time algorithms and mitigation techniques, identifying/addressing emergent risks, verification and validation, and data quality and fusion.

It is understood that there will be differences in solutions across the various possible application domains. Initial R&D will focus on two domains of interest: airline-like operations in the terminal area (e.g., such as at the major U.S. airports) and emerging highly autonomous operations at low altitudes near and over populated urban areas (e.g., such as envisioned for Urban Air Mobility (UAM) and several small UAS use-cases). The remainder of this paper considers the emerging operations domain.

III. APPLICATION DOMAIN – HIGHLY AUTONOMOUS LOW ALTITUDE URBAN OPERATIONS

New airspace operations are emerging with a variety of proposed civil and commercial applications that have garnered significant interest due to their anticipated benefits. While safety will play a key role in either constraining or enabling these benefits, an acceptable level of safety remains to be determined. However, early data suggests new safety assurance technologies are needed [9-12].

A. State-of-the-Practice

While not addressing the urban domain specifically, industry and the FAA are taking several steps to seamlessly integrate unmanned aircraft into the NAS while maintaining an acceptable level of safety. Examples include:

- Developing advanced separation and collision avoidance capability, including sensors, and Sense-/Detect-And-Avoid (SAA/DAA) systems [13-15]
- Advancing the performance and security protocols of command-and-control (C2) communication links for remotely-piloted aircraft systems (RPAS)
- Developing human-system interfaces to support situation awareness, command-and-control, and automation/autonomy management
- Developing standards and procedures for the safe integration of UAS operations into the airspace used by manned commercial aviation [15-18]
- Developing systems and geo-fencing technology for sUAS; but not to high design assurance levels (DALs) and Data Processing Assurance Levels (DPALs) such as required for manned aircraft [19-21]
- Developing industry consensus on high-priority safety enhancements based on a data-driven process [22]; the Unmanned Aircraft Safety Team (UAST) is a government/industry group that is working towards this consensus; they also envision a data environment analogous to ASIAS, but dedicated to UAS. Such a data environment could serve as a critical element of the in-time safety assurance system described herein.

Despite all these developments as well as significant market interest, several unique risks to the low altitude urban flight environment remain largely unaddressed. As a result, operational constraints have been imposed (e.g., no flight beyond visual line-of-sight and no flight over people) [23]. Further, for both UAM vehicles and sUAS, some of these developments may not be applicable due to size, weight, and power (SWAP) constraints (e.g., C2 or DAA equipment may be too large or heavy for some sUAS vehicles). Finally, there is no solution that can provide a timely predictive capability that considers context and the changing nature of risk, particularly in off-nominal or complex situations.

B. Concept of Operations

For the emerging urban operations domain, the in-time safety assurance concept of operations can be thought of as an extension to the way current operations are conducted, with the provision that an information sharing infrastructure will be in place. This infrastructure, with both on-board and off-board elements, will enable the sharing of safety-relevant information, the assessment of risk (including predictions), and mitigation option generation and/or automated control actions (as required). To illustrate the concept, consider three operational phases: pre-flight, in-flight, and post-flight.

Pre-flight – As part of a pre-flight checklist, a trained, qualified, and licensed ground control system (GCS) operator 'connects' his/her vehicle(s) to one or more authorized safety-relevant information service providers. The operator may opt for generally-available 'broadcast' information, or mission-specific 'request-reply' information. For the latter, the operator can send flight plan or intended region of flight along with requested information types. In either case, the latest information (which may include a model and a forecast) is transmitted by the service provider. Upon receipt, the information is checked for validity by the GCS operator and he/she has the opportunity to check for any high-risk areas that may cause reconsideration of flight plans or launch window. If found acceptable, the information is loaded onto the vehicle system for use as the basis for on-board monitoring, assessment, and mitigation of risk.

Such a pre-flight procedure is already being used for some UAS and sUAS operations for limited information types and services [24-26], and is consistent with the UAS Traffic Management (UTM) concept [27]. However, in this case several new safety-related information elements would be available and associated new service provider functions in operation.

In-flight – Once vehicle(s) are launched, the GCS operator monitors flight status including risk(s) or other observables reported by the vehicle(s) or other service providers. Monitoring functions may be executing in parallel at three locations: on-board the vehicle(s), at the GCS, and/or at the service provider(s). Most-likely each will be operating at different rates and with forecasts of different resolutions and time horizons. The choice of where monitoring should occur can be based on operational safety requirements for a particular domain (i.e., the mission/vehicle/operational environment three-tuple), the risks therein, and the bandwidth limits of links

between locations. Risk assessment and mitigation functions would be implemented in a similar manner. They may execute on-board, at the GCS, and/or at the service provider during flight.

The above describes a concept with only one human in-the-loop during operations, the GCS operator. It is possible that other humans will be (at least initially) required to ensure safety. For example, a remote 'safety' pilot who is part of the mitigation function under some circumstances. As such, a 'fully-autonomous' operation is not envisioned at this time (i.e., no human in-the-loop during flight).

Post-flight – After each flight, the primary activity to perform, relative to in-time safety assurance, is to off-load all flight data recorded during the flight. These data, along with any data recorded at the GCS, will be uploaded to the relevant service providers to support updating and validating their services and relevant models. Performance anomalies can be reported which may lead to design changes or maintenance actions. The GCS operator may also report safety-relevant observations or metadata to help providers understand the data and/or the operational context of the flight. This process may be highly automated and supported by appropriate tools such that post-flight procedures can be completed in a timely manner. Other post-flight activities may include active probing of flight-critical equipment with ground-based inspection tools to determine any deterioration not sensed during flight.

C. Use-Case

To better illustrate the concept of operations, it can be helpful to walk through scenarios that utilize these constructs. UAM and urban sUAS-based use-cases will vary with complexity and boundary conditions and examples span the transport of goods/supplies, infrastructure inspection, fire department and law enforcement support, and air taxi. As a low-complexity example, the transport of medical specimens from a suburban medical office to a large downtown laboratory for testing at a hospital will be used for illustrative purposes here.

Consider an urban setting, the weather is benign, with low winds, clear and unlimited visibility during daylight. A direct route between a medical office's drone operations loading area (roof top) and a testing lab at a local hospital covers about three miles. The route crosses over large businesses that are within walking distance of popular lunch destinations. An elementary school is near the departure zone but not on the direct route between the office and hospital. An event is in progress at a stadium near the hospital. The vehicle is owned by a company that provides specialized transport services of small cargo (<10 lbs) using multi-copter sUAS. Flight profiles remain below 400 ft above the ground, and sufficient airspace has been allocated by UTM such that there should be no other air traffic operating during the desired flight window.

During pre-flight, the GCS operator connects to UTM to coordinate airspace access and also subscribes the vehicle to two safety-relevant information services (i.e., SDSs) that provide mission-specific 'request-reply' information to customers. In this case, the services are a flight route risk evaluator and a vehicle health state assessor. The GCS operator sends a flight plan to the risk evaluator service. A risk estimate and forecast is sent by the service provider, and the GCS operator checks for any predicted high-risk areas that may suggest changing the flight plan or launch time. Because of the lunch crowd on the direct route and the proximity to the school and a baseball game that will be underway at the stadium during the flight window, the risk estimate is high so the operator defines a new route that is not the shortest distance but one that minimizes risk due to flight over people while maximizing the operator's objectives of meeting desired launch time (when the specimen is available) and duration of flight (which should be minimized because the specimen's quality deteriorates with time and the UAS battery life is limited).

Because the operator has subscribed to the vehicle health state assessor, the risk evaluator also considers a system health estimation and prediction (e.g., battery state). Use of this service requires transmitting the current state of systems along with metadata regarding the vehicle type and the type of onboard systems (e.g., battery type). This data would be automatically annotated to the flight plan when the service request is initiated, so that only one request-for-information is sent. Once the GCS operator confirms 'latest info received', the information is auto-loaded onto the vehicle system which will use this data (and models) as the basis for on-board monitoring, assessment, and mitigation of risk.

After the vehicle is launched, onboard systems maintain situational awareness and predict proximity to hazards along the route, taking into consideration changes to the environment (e.g., winds, RF interference, population density dynamics), and performs predictive monitoring of the health of the vehicle's critical assets that could change within the duration of a flight (e.g., navigation system). The GCS operator also monitors flight status including updates reported by the vehicle and risk tracking functions operating at the GCS. Mid-way through the flight because of unexpected head-winds, the available remaining charge in the battery causes the risk assessment for the remainder of the flight to change dynamically. This information is used by another safety service that determines alternative routes can achieve acceptable risk exposure based on the projected vehicle state. A shortcut is suggested over the stadium. This remains within the reserved airspace but provides lower risk as the crowd from the stadium has dispersed due to the game ending earlier than expected. The vehicle lands at the hospital on time and the specimen is retrieved by authorized personnel.

After the flight, vehicle flight data recorded on-board (i.e., data not transmitted during flight due to bandwidth limitations) is uploaded automatically to the relevant service providers so that models can be refined/updated (e.g., vehicle performance, battery dissipation, winds, RFI, and population density). Per an agreement with the service provider for 'no-cost' service, the vehicle is connected to a docking station where the battery is probed to better assess and model capacity deterioration and reported to the manufacturer.

IV. FUNCTIONAL DECOMPOSITION AND ARCHITECTURE

For the selected domain, a reference architecture has already begun to evolve. The UAS Traffic Management (UTM) ecosystem is shown in Fig. 3 [27]. For the purposes of research, development, test, and evaluation activities, it is assumed that in-time safety assurance functions must reside within the UTM construct for the selected domain.

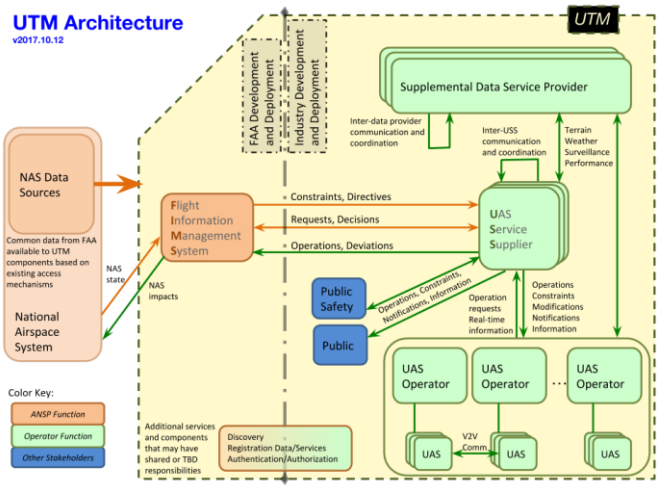

Fig. 3. UAS Traffic Management (UTM) architecture [27].

As previously-described, in-time safety assurance systems provide integrated dependable monitoring, assessment and mitigation capabilities for safety-critical risks. But, these three high-level functions must be decomposed to address domain-specific safety concerns (i.e., what to monitor, assess, and mitigate). Initial research targets five safety risks identified by the community:

(1) flight outside of approved airspace

(2) unsafe proximity to people or property

(3) critical system failures (including loss of link, loss or degraded GPS, loss of power, engine failure)

(4) loss-of-control (including envelope excursions and flight control system failures)

(5) cyber-security related risks

[Note: Air traffic-related risks are not listed as they are implicitly and explicitly addressed by UTM.]

Risks (1) and (2) are recognized as the most undesirable unsafe outcomes of flight in this domain. They are also largely responsible for the current constraints on operations (e.g., the vehicle must remain within line-of-sight and must not fly over people) [23]. Risks (3) through (5) are the most common and likely causal or contributing factors to these undesirable outcomes. Additionally, these risks can potentially be mitigated using the data-driven in-time safety assurance concept, and do so with a positive cost-benefit over traditional methods applied to large manned aircraft.

Comprehensive statistics on the frequency and severity of these risks do not yet exist; and they will likely be comprehensively measured and tracked over the coming years; however, some limited analyses have been done using limited data sets. As mentioned earlier, the UAST's data working group is looking to create an ASIAS-like capability for the UAS domain. This capability will allow the industry to both confirm whether these are the right set of risks, and to evolve a more data-driven model-based predictive capability as is envisioned here for in-time safety assurance systems.

A. *Monitor Function and Supplemental Data Services*

Initially, model-based predictive capabilities are defined as sub-functions within the monitor function for each safety-critical risk. These sub-functions may operate at different rates, and look-ahead horizons based on user/operator requirements. Fig. 4 lists example monitor sub-functions implemented as UTM Supplemental Data Services (SDSs) and hosted by UTM SDS providers (SDSPs). As described in the concept of operations, there may also be instances of one or more such sub-functions executing at the Ground Control Station (GCS) and/or on the vehicle.

The SDSs currently under investigation as part of the monitor function for the selected domain include:

- Aircraft state information and aerodynamic model
- Positioning system state information and performance model
- Communications system state information and radio frequency interference (RFI) model
- Population density information and dynamics model
- Vehicle system health state information and model (i.e., engine and battery health)
- Aeronautical Information Services (AIS) [28-33] (e.g., special use airspace, temporary flight restrictions, weather, and geographic data representing terrain, obstacles, and airport mapping features); this type of service already exists and is transitioning to a more timely update rate such as would be needed here; however, it is not yet tailored to low altitude sUAS urban operations

As described in the concept of operations, similar monitoring sub-functions may execute at the GCS and/or on the vehicle, albeit at different rates, resolutions, and look-ahead horizons. Further, the coordination, synchronization, and interaction of these instances may need to be addressed differently based on operational state: pre-flight, in-flight, or post-flight.

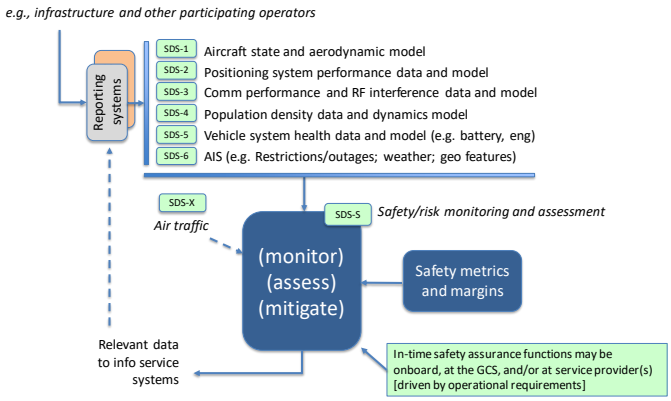

Fig. 4. In-time safety assurance functions and SDSs.

B. *Assess Function*

Two potentially powerful assessment capabilities become possible within the described system construct. First, monitoring functions at the SDSP can take in reports from many (all) operators and reporting systems that 'connect' to the system. Given this capability, models can evolve. This data-driven operational validation can reduce uncertainty over time and allow for unique models for various equipment types (e.g., vehicle, engine, battery); operating environments (e.g., weather, 3D structures); and mission profiles. Further, the quality of the information provided by the service(s) should improve over time as more and more flight data is added to the archive. This in turn allows for assessments that consider not only prior experiences of a particular user; but similar experiences of other users.

Second, as alluded to previously, assessment can now look across hazard/risk types. Historically in manned aviation, accidents are the result of complex and often cascading events. Unusual circumstances that combine in ways that are either unanticipated by designers or viewed as extremely improbable. Assessment can now consider over-arching risk and safety margin [34-38].

The 'assess' function is defined as the processing of information obtained from the 'monitor' function with the goal of detecting, diagnosing and predicting risk and hazard states. As such, it will leverage information from available sources as well as the various SDSs outlined in the previous section. Assess functions may be operating concurrently onboard, at the GCS, and/or at the SDSP. And, as such, outputs of the assessment function may be an SDS. In this way, operators may choose to receive an overall risk assessment, rather than connecting to all the services that may be contributing factors.

Initially, the assessment SDSs under investigation include:

- SDS-X provides information relative to air traffic and airspace constraints; this may include position reports, warnings, and/or advisories. This capability is largely (if not completely) covered by UTM developments and the USS connection [27]
- SDS-S is envisioned as a service that would provide an over-arching report and assessment of the evolution of safety risk vis-à-vis a desired safety margin [34-37]. This service would make use of output from other safety-relevant services to estimate, track, and predict safety risk. In cases where communication bandwidth is limited, this service may be used in lieu of the others. However, the operator may have less in-sight as to which elements are most affecting the reported risk if not connected to the other services.
- SDS-R performs a real-time risk assessment based primarily on aircraft state, vehicle system states, weather factors, and population density in the region of flight [38].

Generally, any function that reasons over data and provides advisories and information regarding safety state or risk considerations is covered by this function.

C. *Mitigate Function*

The 'mitigate' function resolves either current or impending situations that exceed a defined safety threshold. While much of the R&D for this function is planned for future years, it is important to keep in mind that the monitoring and assessment functions ultimately determine how well mitigation can occur for any safety-adverse situation that develops. Decision-making is the task of choosing a course of action among multiple alternatives, and therefore the tools that will be employed will likely utilize a suite of optimization techniques. For in-time decision-making, speed of execution is key and needs to be considered in the presence of possibly limited on-board computational resources. Another key challenge will be defining roles and responsibilities between human(s) and machine. Likewise, the distribution of authority and autonomy between human(s) and machines. There is a significant amount of prior work in this area that can be leveraged and applied. However, the degree to which this can be done, versus discovering completely new approaches, will depend on the specific use-case, associated hazards, and target level of safety.

V. BASELINE CAPABILITY DEVELOPMENT AND TESTING

To test and evaluate information exchange protocols as well as distribution of function across SDSPs, GCSs, and vehicles, a developmental system for the vehicle has been designed. Fortunately, ground-based elements can be based on previous work to advance UTM Technology Capability Level (TCL) [27]. This architecture allows sufficient flexibility to conduct the planned R&D. What is missing is an analogous on-board system architecture supportive of the R&D. For this aspect, we build upon the core Flight System (cFS) [38] and a previous activity to develop an operating system for unmanned aircraft, sponsored by NASA's Convergent Aeronautics Solutions (CAS) project [40].

cFS is a platform, software framework, and environment that allows for development and re-use of flight software applications [39]. Essentially, it's a form of middle-ware that resides between the Operating System (OS) and application layers. This middleware provides a dynamic run-time environment that allows for an independent component-based design. cFS has been used by a number of NASA flight projects that require complex embedded software systems. One of its features is that software applications can be developed

and functionally tested independently of other applications. This is particularly useful for our R&D as some the safety assurance functions and sub-functions are at various levels of maturity and will mature at different rates as the research progresses.

Among other things, cFS allows independently-executing functions to access a common 'shared' information bus. In a carefully controlled and coordinated way, functions may read or write information to this bus, much like 'apps' interface to cloud-based data storage. cFS has matured over many years and is in wide-spread use across many domains where real-time OS performance can be mission- and safety-critical. As such, it provides the stable and robust platform we require for conducting R&D on each of the envisioned functions, as well as the aggregate system-of-systems aimed at safety assurance. For example, we can create an asynchronous 'app' associated with each of the monitoring sub-functions mentioned previously, as well as for any automated or autonomous assessment and mitigation functions we may want to evaluate. This also allows each 'app' to be designed to unique data quality requirements (DQRs) as well as a DAL, DPAL, or Specific Assurance and Integrity Level (SAIL) [18-21].

Fig. 5 shows the CFS-based architecture to be used during initial testing of a baseline capability. Details on this flight system design will be published separately, along with test results.

Referring to Fig. 5, thirteen vehicle system functions are defined that correspond to the monitoring and assessment functions previously described (mitigation functions will be introduced in the future). Some of these functions may be operating concurrently at the GCS and/or the SDSP, likely at different rates and look-ahead horizons. Information exchange between the aircraft and GCS safety assurance functions should be kept to a minimum during flight to (1) reduce bandwidth demands and associated costs (downlink) and (2) reduce vulnerability to data corruption or loss-of-integrity effects (uplink). Each use-case and associated risks would determine which (if any) of these functions would be required, as well as whether they should be onboard, at the GCS, and/or at the SDSP. Also, use-case risks and desired safety level would determine DQR, DAL, DPAL, and/or SAIL for the functions [18-21].

Of the developmental flight system functions ('apps') shown in Fig. 5, the following will be tested initially:

- Battery monitor (based on [41])
- Safety/risk monitor (based on [34-38])
- Constraint monitor (based on [42-44])
- Traffic monitor (based on [14,45])

Data will also be recorded to support off-line development and testing of the navigation, engine, and link monitoring functions. These function descriptions and findings will be reported elsewhere.

Testing will be conducted on NASA's City Environment for Range Testing of Autonomous Integrated Navigation (CERTAIN) test range [46,47], where scenarios and mission profiles will emulate various urban and suburban use-cases.

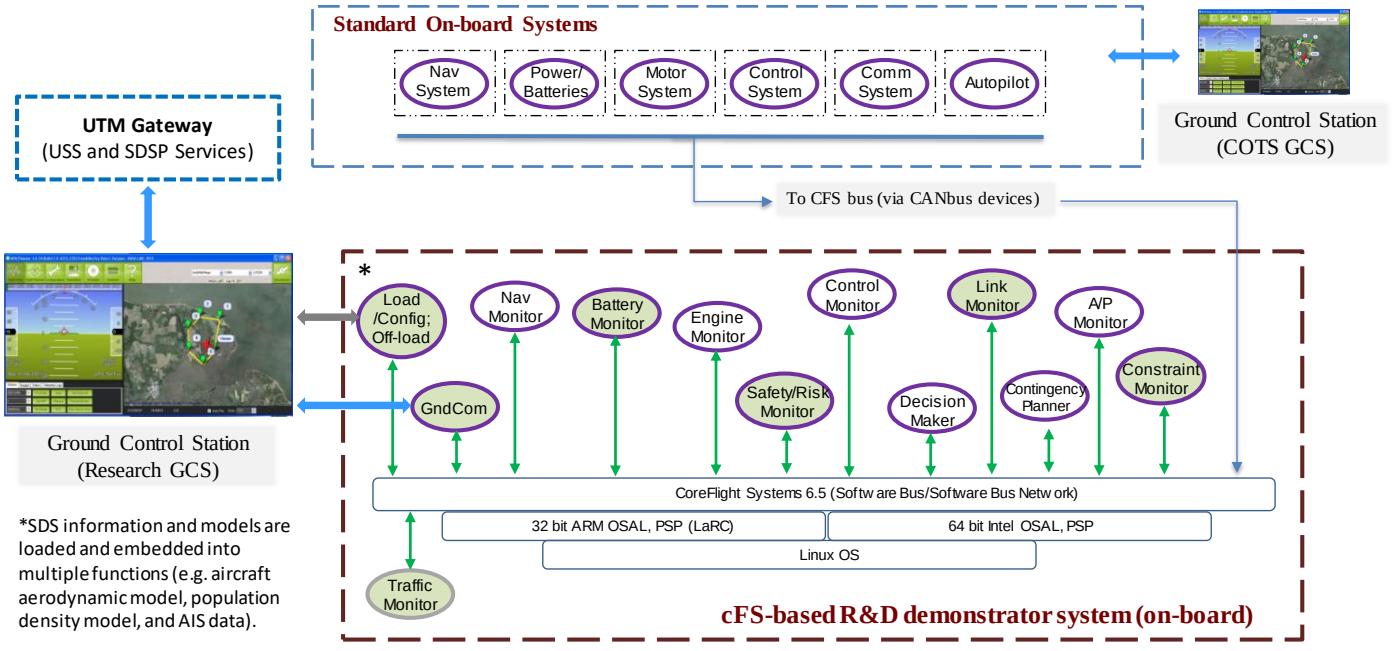

Fig. 5. cFS-based flight system elements supporting R&D.

VI. NEXT STEPS

Research will continue toward advancing the monitoring functions described herein, while additional efforts will address automated assessment challenges. In particular, determining DQRs [20,21] for any inputs to safety-critical decision-making functions, and in parallel, characterizing the uncertainties associated with information sources and services. Incremental testing is planned on an annual basis within environments that can host Smart-City-like infrastructure elements of the system architecture [47,48].

VII. SUMMARY

A safety assurance system concept is presented that comprises monitoring across a broad set of information types, data-driven automated assessment, and in-time mitigation functions. These functions are intended to proactively reduce risk in complex operational environments wherein the interplay of hazards may not be known, and cannot be accounted for at design time. Such a system can also help to understand and predict effects caused by the increased use of automation or autonomous functions that may exhibit unexpected non-deterministic behaviors. The envisioned system can observe these behaviors and apply model-based and data-driven methods to provide a level of run-time assurance. This paper presents a preliminary conceptual design of such a system being tested for the emerging domain of highly-autonomous aircraft operating at low altitudes near and over populated urban areas.

REFERENCES

[1] RTCA, Minimum Operational Performance Standards (MOPS) for Traffic Alert and Collision Avoidance System II (TCAS-II) Hybrid Surveillance, RTCA document DO-300A, December 15, 2015.

[2] RTCA, Minimum Operational Performance Standards (MOPS) for Terrain Awareness Warning Systems Airborne Equipment, RTCA document DO-367, March 21, 2017.

[3] AvWeek, "Automatic Ground Collision Avoidance System Explained," Aviation Week and Space Technology, September 19, 2016.

[4] Commercial Aviation Safety Team (CAST), Safety Enhancement Plan, [Online] https://skybrary.aero/index.php/Portal:CAST_SE_Plan , April 2018.

[5] Federal Aviation Administration, Aviation Safety Information Analysis and Sharing System, [Online] , https://www.asias.faa.gov, August 2018.

[6] National Aeronautics and Space Administration, Aeronautics Research Mission Directorate Strategic Implementation Plan, [Online] https://www.nasa.gov/aeroresearch/strategy, March 23, 2017.

[7] National Academies of Sciences, Engineering, and Medicine. 2018. In-time Aviation Safety Management: Challenges and Research for an Evolving Aviation System. Washington, DC: The National Academies Press. https://doi.org/10.17226/24962.

[8] International Civil Aviation Organization, Safety Management, Standards and Recommended Practices, Annex 19 to the Convention on International Civil Aviation, 2nd edition, July, 2016.

[9] Federal Aviation Administration, UAS Sightings Report, [Online] https://www.faa.gov/uas/resources/uas_sightings_report/, August, 2018.

[10] National Transportation Safety Board, Aviation Incident Final Report, Incident Number DCA-17IA202A, Dec 14, 207 (Sikorsky UH-60M Black Hawk helicopter collides with a privately owned Dà-Jiang Innovations (DJI) Phantom 4, Hoffman Island near New York City, New York, Sep 21, 2017).

[11] Arkin, D., "Rogue Drone Slams Into Ferris Wheel Near Downtown Seattle's Waterfront," [Online] Available at: www.nbcnews.com, November 12, 2015.

[12] Belcastro, C., et. al., "Hazards Identification and Analysis for Unmanned Aircraft System Operations," 17th AIAA Aviation Technology, Integration, and Operations Conference, AIAA AVIATION Forum, AIAA 2017-3269, Denver, CO, June 5-9, 2017.

[13] Manfredi, G., and Jestin, Y., "An introduction to ACAS Xu and the Challenges Ahead," AIAA/IEEE 35th Digital Avionics Systems Conference, Sacramento, CA, Sep 2016.

[14] Munoz, C., et. al., "DAIDALUS: Detect and Avoid Alerting Logic for Unmanned Systems," 34th 2015 AIAA/IEEE Digital Avionics Systems Conference, Prague, Czechoslovakia, September 13-17, 2015.

[15] RTCA, "Minimum Operational Performance Standards (MOPS) for Detect and Avoid (DAA) Systems," RTCA Document DO-365, May 31, 2017.

[16] RTCA, "Guidance Material and Considerations for Unmanned Aircraft Systems," RTCA Document DO-304, March, 2007.

[17] RTCA, "Operational Services and Environment Definition for Unmanned Aircraft Systems," RTCA Document DO-320, June, 2010.

[18] Joint Authorities for Rulemaking of Unmanned Systems, "Guidelines on Specific Operations Risk Assessment," JARUS Document JAR-DEL-WG6-D.04, June 26, 2017.

[19] RTCA, "Software Considerations in Airborne Systems and Equipment Certification," RTCA Document DO-178C, December 2011.

[20] Federal Aviation Administration, "Acceptance of Aeronautical Data Processes and Associated Databases," FAA Advisory Circular AC-20-1538, April 2016.

[21] RTCA, "Standards for Processing Aeronautical Data," RTCA Document DO-200B, June 2015.

[22] Unmanned Aircraft Safety Team, Mission and Safety Enhancements, [Online] https://www.unmannedaircraftsafetyteam.org/, Aug 2018.

[23] United States Government, Title 14 Code of Federal Regulations, Part 107, Small Unmanned Aircraft Systems, [Online], available at http://www.ecfr.gov, August 2018.

[24] Federal Aviation Administration, FAA UAS Data Exchange and the Low Altitude Authorization and Notification Capability, [Online] https://www.faa.gov/uas/programs_partnerships/uas_data_exchange/, August 2018.

[25] Airmap. Complete Situational Awareness for Drone Pilots, [Online] https://www.airmap.com/airmap-for-drones/, August 2018.

[26] Skyward. [Online] https://skyward.io/how-it-works/, August 2018.

[27] Aweiss, A., et. al., "Unmanned Aircraft Systems (UAS) Traffic Management (UTM) National Campaign II," Proceedings of AIAA SciTech Forum, AIAA Infotech @ Aerospace, Kissimmee, FL, 8-12 January 2018.

[28] International Civil Aviation Organization, Aeronautical Information Services, Standards and Recommended Practices, Annex 15 to the Convention on International Civil Aviation, 15th edition, July, 2016.

[29] RTCA, "Minimum Aviation System Performance Standards (MASPS) for Aeronautical nformation/Meteorological Data Link Services," RTCA Document DO-364, December 15, 2016.

[30] RTCA, "User Requirements for Aerodrome Mapping Information," RTCA Document DO-272D, November 2015.

[31] RTCA, "User Requirements for Terrain and Obstacle Data," RTCA Document DO-276C, November 2015.

[32] RTCA, "Interchange Standards for Terrain, Obstacle and Aerodrome Mapping Data," RTCA Document DO-291C, November 2015.

[33] EUROCONTROL, Aeronautical Information Exchange Model (AIXM), [Online] www.aixm.aero, August 2018.

[34] Roychoudhury, I., Spirkovska, L., Daigle, M., Sankararaman, S., Balaban, E., Kulkarni, S., Poll, S., and Goebel, K., "Predicting Real-Time Safety Margins in the National Airspace System," Proceedings of AIAA SciTech 2016, San Diego, CA, January 4-8, 2016.

[35] Roychoudhury, I., Daigle, M., Goebel, K., Spirkovska, L., Sankararaman, S., Ossenfort, J., Kulkarni, S., McDermott, W., and Poll, S., "Initial Demonstration of the Real-Time Safety Monitoring

Framework for the National Airspace System Using Flight Data," Proceedings of AIAA AVIATION 2016, AIAA-2016-4216, Washington, DC, June 13-17, 2016.

[36] Roychoudhury, I., Spirkovska, L., Ossenfort, J., Sankararaman, S., Kulkarni, S., Goebel, K., Poll, S., McDermott, W., and Daigle, M., "Real-Time Prediction of Safety Margins in the National Airspace," Proceedings of AIAA AVIATION 2017, AIAA-2017-4388, Denver, CO, June 5-9, 2017.

[37] Spirkovska, L., Roychoudhury, I., Daigle, M., and Goebel, K., "Real Time Safety Monitoring: Concept for Supporting Safe Flight Operations," Proceedings of AIAA AVIATION 2017, AIAA-2017-4494, Denver, CO, June 5-9, 2017.

[38] Ancel, E., et. al., "Real-time Risk Assessment Framework for Unmanned Aircraft System (UAS) Traffic Management (UTM)," AIAA Aviation 2017 Conference, Denver, CO, June 5-9, 2017.

[39] National Aeronautics and Space Administration, Goddard Flight Research Center, Core Flight System Background and Overview, [Online] https://cfs.gsfc.nasa.gov/ and https://cfs.gsfc.nasa.gov/cFS-OviewBGSlideDeck-ExportControl-Final.pdf , 2018.

[40] Lowry, M., et. al., "Autonomy Operatng System for UAVs: Pilot-in-a-Box", 17th AIAA Aviation Technology, Integration, and Operations Conference, AIAA AVIATION Forum, (AIAA 2017-4272), Denver, CO, Jun 5-9, 2017.

[41] Kulkami, C., et. al., "Modeling for Battery Prognostics," NASA Aerospace Battery Workshop, Huntsville, AL, [Online] Available at: https://ntrs.nasa.gov/, November 14-16, 2017.

[42] Gilabert, R., et. al., "Safeguard – Progress and Test Results for A Reliable Independent On-board Safety Net for UAS," Proceedings of the 36th AIAA/IEEE Digital Avionics Systems Conference, St. Petersburg, FL, Sep 2017.

[43] ASTM International, ASTM F3269-17, "Standard Practice for Methods to Safely Bound Flight Behavior of Unmanned Aircraft Systems Containing Complex Functions," 2017.

[44] European Aviation Safety Agency, "Study and Recommendations regarding Unmanned Aircraft Systems Geo-Limitations", EASA/NAA Task Force Report, 02 September 2016.

[45] Consiglio, M., et. al., "ICAROUS – Integrated Configurable Algorithms for Reliable Operations of Unmanned Systems," 35th AIAA/IEEE Digital Avionics Systems Conference, Sacramento, CA, September 25-30, 2016.

[46] National Aeronautics and Space Administration, "NASA Langley Drone Flying Site Open for Testing," [Online] Available at: https://www.nasa.gov/langley/nasa-langley-drone-flying-site-open-for-testing , August 2017.

[47] Moore, A., et. al., "Testing Enabling Technologies for Safe UAS Urban Operations," Proceedings of AIAA AVIATION 2018, Atlanta, GA, June 25-29, 2018.

[48] National Institute of Science and Technology, "IoT-Enabled Smart City Framework," [Online] https://pages.nist.gov/smartcitiesarchitecture/, 2018.