JOURNAL OF AEROSPACE INFORMATION SYSTEMS
Vol. 18, No. 1, January 2021

Generating Certification Evidence for Autonomous Aerial Vehicles Decision-Making

Donald H. Costello, III* and Huan Xu†
University of Maryland, College Park, Maryland 20742

https://doi.org/10.2514/1.I010848

The last 15 years have seen a large uptick in the use of unmanned aircraft. However, the current safety of flight clearances for unmanned aircraft requires a qualified operator who can make decisions and ultimately bear the responsibly for the safe operations of the vehicle. The future of aviation is unmanned, and ultimately autonomous. Yet, a clear path for certifying an autonomous vehicle to make decisions currently reserved for qualified pilots does not exist. This paper presents a preliminary approach for certifying an autonomous controller to select an appropriate landing site for a large rotorcraft in an unprepared landing zone. In particular, this paper will decompose the steps currently used by qualified pilots to the basic requirements to define an envelope where the vehicle will be allowed to operate autonomously while landing. These requirements are the basis for a specification that we examine to ensure it met the requirements. A protocol is developed based on the analyzed specification that will ensure what the vehicle "will not do" while operating autonomously. Finally, we describe how this protocol can be used as the safety of flight evidence, and eventually for clearing an autonomous controller to complete a task reserved for qualified pilots.

I. Introduction

UNMANNED aviation is expected to continue to increase over the next decade [1]. Unmanned aircraft can operate far beyond the limitations of human endurance. However, unmanned aircraft are currently required to have a qualified operator in the loop. This operator, who controls the vehicle and makes decisions, is ultimately responsible for the safe operations of the vehicle [2]. Future systems are expected to allow these vehicles to operate autonomously. Yet, an approved process for certifying an autonomous vehicle to accomplish tasks that are currently reserved for qualified pilots does not exist [3,4].

Regulations for certifying a manned aircraft have evolved since the beginning of powered flight. Depending on the ultimate mission of the aircraft, a well-defined process can be identified and used for certification. The qualification process for pilots is also well defined and has gone through numerous iterations over the years [5]. Many modern aircraft can, and are, operated through a set of pilot relief modes (autopilots) that allow the aircraft to complete nearly the entire flight without a pilot touching the controls. However, the pilot in command still has the responsibility for the aircraft. The pilot in command (or controller in the case of unmanned aircraft) is required to operate the vehicle under current certification standards. Federal Aviation Administration (FAA) certification for unmanned vehicles only deals with small vehicles (referred to as quadcopters or similar small drones) and requires the operator to be within line of sight of the vehicle [6]. Autonomous aircraft will not have an operator in the loop and will ultimately require a new process [2,7].

Before safety of flight certification, officials require data to justify such a flight clearance [8]. These data are referred to as certification evidence. This paper describes the development of certification evidence for safety of flight certification of a well-defined task: autonomous landing of a helicopter in an unprepared landing zone. An unprepared landing zone is a location that is not certified for rotorcraft operations (not an aerodrome or helipad). We use the unprepared confined area landing (CAL)/landing zone (LZ) mission currently carried out by H-60 variant helicopters by the United States Navy (USN) and United States Marine Corps (USMC) as a running example [9]. This mission can be as simple as landing in an open field adjacent to a highway, or as difficult as landing between buildings in an urban setting. The process for choosing a landing spot is complicated and, before being certified as a helicopter aircraft commander (HAC), a candidate is expected to be able to accurately complete this task [5].

Since the dawn of aviation, many of the innovations we currently take for granted came from the military (some examples include radar [10], medevac air ambulance [11], jet engines [12], glow sticks [13], and advanced night vision technology [14]). Many military applications can transition easily to the civilian sector because their functionality is similar. For this reason, we chose a military application that can be easily translated into a civilian sector for this research. The evidence generated can be use for certification of future autonomous vehicles.

For naval aviation, airworthiness certification authority is delegated to Naval Air Systems Command (NAVAIR) 4.0 Engineering (4.0P is the branch assigned) [8]. When a new capability/software/weapon/airframe is acquired, and before naval personnel operate it, 4.0P must grant a flight clearance (also referred to as a safety of flight certification). The certification process for naval aircraft is a risk mitigation process. Aircraft subsystems, software, components, and ultimately the aircraft itself are certified through an established process. Technical Area Experts (TAEs) are tasked with reviewing certification evidence (referred to as artifacts) in their individual technical areas. These reviews are rolled up into a larger flight clearance that certification officials use to certify the vehicle as a whole. When a vehicle is certified safe for flight, NAVAIR 4.0P is certifying that when given to a qualified pilot, they can safely complete the desired mission of the aircraft [8].

All modern aircraft have some level of automation, and this automation is thoroughly tested during the certification process. In this paper, a distinction has been made between automation (such as a pilot relief mode in an autopilot) and autonomy. For automation, a system functions with no/little human operator involvement; however, the system performance is limited to the specific actions it has been designed to do. Typically these are well-defined tasks that have predetermined responses (such as "maintain altitude" or "fly the published approach for the duty runway"). For autonomy, a system has a set of intelligence-based capabilities that allows it to respond to situations that were not preprogrammed or anticipated (i.e., decision-based responses) before system deployment. Autonomous systems have a degree of self-government and self-directed behavior [15]. This difference can be further deconstructed into deterministic behavior (based on known input conditions, the vehicle will exhibit

Received 7 April 2020; revision received 15 September 2020; accepted for publication 20 September 2020; published online 20 October 2020.

*Commander United States Navy, Naval Air Warfare Center Aircraft Division/UMD Ph.D. Candidate, Mechanical Engineering; donald.h.costello@navy.mil.

†Assistant Professor Aerospace Engineering/Institute for Systems Research; mumu@umd.edu.

a known behavior) and nondeterministic behavior (the exact behavior of the system cannot be determined based upon the input conditions).

As NAVAIR 4.0P certifies aircraft to be operated by qualified pilots, it is important to understand how the process to qualify a pilot differs from the aircraft safety of flight certification process. The qualification process for naval aviators (pilots) is considered to be a trust process. Unlike the civilian sector, military pilots are trusted by their commanding officers (COs) to complete missions critical to national interests. While each pilot is required to log a minimum amount of flight time and show competency in aircraft procedures before qualification, a commanding officer will not designate them as fully qualified until the individual has earned the trust of the CO in their decision-making abilities in off-nominal conditions [5].

In an attempt to provide a path forward for certifying autonomy in aviation, this paper provides a limited approach for providing evidence that can be used for certifying an autonomous controller to exhibit nondeterministic behavior when selecting a LZ autonomously during the unprepared CAL/LZ mission. This mission (the task of selecting and continuously evaluating a landing spot during the approach and landing phase of flight) is currently carried out by H-60 variant helicopters by the USN and USMC [9]. Before certification, TAEs need to be provided certification evidence that the system can complete tasks currently reserved for pilots [8]. This paper will decompose the tasks currently completed by a pilot during the CAL/LZ mission to their basic requirements. To develop these requirements, we consulted (over several interview sessions) multiple senior naval officers (those that currently certify a pilot as a HAC), and we followed several junior aviators during the qualification process. Through our conversations and observations, we gained insight as to what was expected of a fully qualified HAC during the mission. Ultimately, we propose a clearance envelope where the system can exhibit nondeterministic behavior. This means the actions of the system cannot be exactly predicted by evaluating the systems parameters, and the system is clear to make decisions currently reserved for qualified pilots, providing it does not reach one of the limits of the clearance envelope. For the CAL/LZ mission, this implies that the autonomous controller can pick its landing spot, providing it does not violate restrictions put in place. If the system were to reach one of these limits, it would revert to predetermined behavior. We examine the correctness of the specification in an effort to show that a path forward exists in which formal verification could be used to certify autonomous systems to complete tasks currently reserved for qualified pilots [16]. We used Prototype Verification System (PVS; a theorem proving tool) to examine a high-level specification for correctness. Then, the analyzed specification was used to develop a protocol for the actions the autonomous controller would take when selecting and controlling the aircraft during the CAL/LZ mission. The protocol was then evaluated against a sample set of possible LZ conditions to ensure that only an LZ that met all of the requirements of the specification would be allowed to be selected by the autonomous controller (eliminate corner cases). Software developers can use the protocol as a guideline for developing the specific code that will control the aircraft. We also presented the protocol to the same senior naval officers that helped develop the requirements for the specification to ensure that it met their criteria for qualification of HACs. All four naval officers agreed that, provided the assumptions were valid, the protocol was adequate for modeling the behavior of a fully qualified HAC in the CAL/LZ mission. The evaluation can be used by certification officials as evidence for the ultimate certification of the system [8].

This paper is structured as follows. In Sec. II, in addition to a review of related research in the area, the current NAVAIR certification process and the HAC qualification process are summarized. In Sec. III, the actions taken by a qualified pilot are deconstructed to their requirements. These requirements are then used as the basis for a specification. In Sec. IV, we analyze the specification to ensure it meets the requirements. Finally, we will show that specified behavior will satisfy the requirements (given the assumptions). In Sec. V, a protocol is presented that can be used by certification officials for the possible certification evidence of autonomous behavior in a naval aircraft. In Sec. VI, we describe how the process of generating the protocol is just one of the first steps toward the certification of autonomous behavior to complete tasks currently reserved for qualified pilots. Directions for future research are also provided.

II. Background

Currently, a formalized/approved process does not exist for naval aircraft/systems that exhibit autonomous behavior (the system is able to respond to situations that were not explicitly preprogrammed) because there has never been a requirement for one to be developed. Several possible approaches have been proposed, but none have been vetted through the naval flight clearance authorities [17–19]. Several issues have been identified for certificating autonomy (i.e., the complexity of autonomous systems results in an inability to test under all known conditions, difficulties in objectively measuring risk, and an ever-increasing cost of rework/redesign due to errors found late in the verification and validation (V&V) process [15]). The decision space for certifying a vehicle to complete all tasks assigned is extremely complex. This work focused on developing evidence for certification officials to certify autonomous functionality during a specific mission: executing a safe landing of a large rotorcraft (capable of transporting passengers) within an unprepared (not an aerodrome or helipad) CAL/LZ. This will enable an exercise of a methodology for just one mission normally reserved for fully qualified rotorcraft pilots (other missions would include tasks such as power line avoidance, see and avoid, formation flying, and visual navigation), thus limiting the complexity of the autonomous functionality requirements.

Academia has proposed several approaches for certification of unmanned/autonomous systems. A majority of the work deals with small unmanned aerial vehicles (UAVs) or abstract methods for certifying large vehicles. One common theme is to identity errors in the software early in the design cycle, since the later a defect is found, the more resources (both in time and money) are required to correct the issue [15,20–23]. Many of the approaches involved modeling and simulation (M&S) to determine if the software was adequate for the system requirements [21,24–31]. Another common approach involves employing formal methods for safety-critical software V&V (run time verification [16,32–42], model checking [19,43–54], and theorem proving [43,54–60]). Some proposals have detailed methodologies for V&V for the unmanned see and avoid requirement but only for a two-dimensional problem [61,62]. One drawback of these approaches is the limited focus of their work. As an approved methodology does not exist, their work was limited to one or two pieces of the V&V process, and most did not consult aviation certification officials. One notable exception is the work done by the Formal Methods Group at National Aeronautics and Space Administration (NASA) Langley Research Center. Currently, NASA is working on (and has published) several papers on obtaining flight clearances for unmanned aerial systems (UASs) to operate within the national airspace. Their work focuses on formally defining the specification from the requirements of operation within the national airspace, and then V&V via theorem provers [63–65]. This is designed to give certification officials confirmation that the software will perform per the requirements.

Most of the current work to certify autonomy is based off easily definable, black and white regulations for operating in the public airspace. One example is collision avoidance, where aircraft are required to maintain a safety bubble around them to avoid collision. This involves an easily definable and well documented set of requirements (such as lateral and vertical separation). These requirements do not involve pilot judgment, and they can be accomplished by using data currently available via onboard systems (such as the traffic collision avoidance system and Mode C transponders). The CAL/LZ mission does not fall into that category; it requires a HAC to evaluate several variables continuously and make a judgment-based decision on an ever-changing situation. This work focuses on developing a protocol that can be used in certifying that "judgment" task. The CAL/LZ mission, as we will reduce the requirements down to, may appear to be a simplified "decision tree" mission.

However, embedded within the assumptions are judgment calls that a qualified HAC will need to complete during the mission.

The protocols/control laws used by an autonomous vehicle are essentially the decision engine for determining where the system will go and how it will react to input conditions. A simplified example would be "if the landing zone is no longer clear, abort landing and hold at predetermined point Alpha." Formal methods offer the ability to build the protocols/control laws based on a formally verified specification. The specification is based on the requirements of the software that will control the system. Using a formal methods approach for the development of the protocols offers flight clearance authorities' insight as to what the system will not do. This insight can be used in the safety of flight certification of an autonomous controller [64,65].

A. Current Certification Process for Naval Aircraft/Systems

Currently, when an aircraft is certified safe for flight (when operated safely, they will not break down or cause a danger to the general public), it is implied that they will be operated by a qualified pilot (or operator in the case of large UAVs such as the Global Hawk or Predator). As an example of a currently fielded system, the USN currently operates the MQ-8 Fire Scout UAV. NAVAIR has certified the large rotorcraft to fly without a qualified HAC. However, an Air Vehicle Operator (AVO) is ultimately responsible for the safe operation of the vehicle. During preflight mission planning, the AVO programs the vehicle to complete parts of the mission without operator input (similar to an autopilot). In the event of loss link, the system will fly to a preplanned point, and land. The system does not perform any evaluation of the landing point; it simply executes a preplanned route to a LZ and autolands [66].

NAVAIR 4.0P has established processes where TAEs, who have been given the authority in their subject fields, review relevant artifacts before approving their portion of a flight clearance. Artifacts can range from subject matter expert opinion to detailed engineering analysis. Often, an artifact is a dataset characterizing the performance of a system. In the end, artifacts exist to quantify the system and allow the certification official to determine the risk they will be accepting.

For the respective TAEs to certify autonomy in their subject area, several challenges will need to be overcome. In the words of the former chief engineer of the United States Air Force, "It is possible to develop systems having high levels of autonomy, but it is the lack of suitable V&V methods that prevents all but relatively low levels of autonomy from being certified for use [67]." The U.S. Air Force Research Laboratory funded a study asking a question regarding the state of possible processes for certification of unmanned aerial systems that employ machine learning or autonomous functionality through some sort of evidence-based licensure process. The report summarized several categories that may lead to the certification of UASs. These categories were formal methods; requirements and metrics; normative oracle generation; coactive design; implications of learning autonomous systems; and modeling and simulation considerations for licensure of autonomous systems [68]. All or some of these categories will be required for the individual TAEs to accept the risk associated with certifying the autonomous functionality.

B. Current Certification Process for Helicopter Aircraft Commander

The overarching purpose of this research is to determine a path forward for certifying a decision engine to act as a HAC in the USN/USMC. To accomplish this, the current HAC qualification process must be understood. This process is formally established, but full qualification depends on a subjective decision of a CO (typically an O-5 or O-6) [5]. Following graduation from the helicopter replacement air group, a pilot will be assigned to a fleet squadron for approximately 36 months. During this time, they will be expected to qualify as a second pilot, complete a HAC syllabus, complete the prerequisite flight experience in model (such as a H-60), pass a HAC oral board, and ultimately earn their CO's trust in their decision-making process before they are considered a fully qualified HAC [5].

To qualify a candidate as a HAC, the CO is placing trust in the pilot's judgment. Any pilot can follow directions or complete a simple mission when everything goes as planned. The question is how they will respond when things do not go as planned. By designating a pilot as a HAC, the CO is putting their stamp of approval on the pilot's ability to cope with the unexpected.

III. Requirements Definition and the Specification

A. Development of the Basic Requirements

The first step in a path for a flight clearance of an autonomous system to complete tasks currently reserved for a qualified pilot is to define the requirements the decision engine must complete. Landing in an unprepared LZ is a difficult mission for qualified HACs. The last 15 years have seen several fatal mishaps where naval aviators have made decisions that led to unsuccessful landing attempts. The Chief of Naval Air Forces (CNAF) established a procedure for pilots to complete when attempting a landing in such a location. The procedure is abbreviated as SWEEP (which stands for size/slope, wind, elevation, escape route, and power) [9]. Several syllabus flights are dedicated to mastering this task, and these flights must be passed before a pilot can be designated a HAC. These flights consist of 17 events totaling 36 flight hours. The experience gained by completing the syllabus events, in addition to the experience the HAC candidate obtains during other events, is used to train the judgment of the aviator before their CO designates them as a HAC [5].

If a decision engine were to be allowed to make the decision on where to land, it would need to demonstrate the ability to complete the SWEEP procedure. This work attempts to program the judgment via the SWEEP into the decision engine, and then allows the decision engine to select a landing point (provided SWEEP is valid). Any protocol used to control its action must prove that it can accurately complete the procedure, every time, before it is certified. It is important to understand each part of SWEEP:

1) The first part is size. The S in SWEEP has two meanings: the first is size of the LZ. The HAC must be able to define the size of the LZ from altitude (nominally 200 ft above ground level). This includes obstacles and the actual area and orientation available for the vehicle to touch down in. An obstacle within the LZ may not negate the suitability of the LZ. Rotor wash may blow some items out of the way during landing (such as tumbleweeds). A HAC uses their experience and judgment to identify which objects may pose a threat. West coast helicopter pilots normally train in the desert of eastern San Diego. The biggest threats to defining a LZ are tall bushes that can cause the vehicle to tip over if they are under the aircraft on landing. A confined area, such as an urban setting, offers still other issues dealing with the actual dimensions of the LZ. Buildings and fences confine the available space to land in. HACs are expected to be able to visually identify the LZ and determine the suitability for landing. All helicopters differ in size.

2) The S in SWEEP also stands for slope. Most prepared LZs are flat and clear of any obstacles. When a helicopter touches down on a flat surface, both skids (or landing gear) touch down at nearly the same time. The greater the slope, the greater the risk that the vehicle may tip over on landing/touchdown due to dynamic rollover. The risk comes when only one of the two main touchdown points makes contact with a surface and becomes a pivot point for the vehicle. Standard operating procedures list a limit for slope based on vehicle configuration and environmental conditions. HACs are expected to evaluate the slope for suitability from altitude, and they continually evaluate the LZ through touchdown.

3) The W in SWEEP stands for wind. Unlike their fixed-wing counterparts, helicopters normally do not land with a forward velocity that dominates the local wind during landing. A fixed-wing aircraft may be able to withstand crosswinds of 30+ kt due to its forward velocity of 100+ kt. A helicopter may have crosswind limits of 5–10 kt while landing. A HAC is expected to evaluate the landing area before approach and continuously during approach to ensure the aircraft can complete a safe landing. In a CAL/LZ, when an aircraft gets near the ground, the wind has a tendency to shift greatly due to local conditions. These shifts may be difficult for the HAC to anticipate from altitude. The HAC is expected to abort a landing if an unsafe wind condition is present.

4) The first E in SWEEP stands for elevation. Tactical helicopters are historically underpowered due to their weight. The closer to sea level, the better the performance of the engines on the aircraft. As altitude increases, the performance of the engines is reduced. The USN trains selected naval aviators at the mountain training school in Fallon, Nevada. There, pilots learn how to control their aircraft when its performance is limited due to elevation. A HAC is expected to be able to accurately evaluate the vehicles performance based on the altitude of the LZ. They are also expected to abort the landing if an unsafe condition exists.

5) The second E in SWEEP stands for the escape route. When evaluating an unprepared LZ, HACs are expected to be able to find a way out (if one exists). The way out is used as an escape route when aborting a landing/approach. This route may be used when an unexpected unsafe condition develops. One example would be if the LZ becomes fouled by an interloper (such as a moving vehicle or wildlife). On this step of the SWEEP procedure, the HAC must select their escape route if a safe landing can no longer be executed. If any escape route does not exist, some low-priority missions will be aborted because the extra risk associated with the mission is not acceptable based on the priority level.

6) The P in SWEEP stands for power. As with all aspects of vertical lift aviation, power is the most critical part of aircraft performance. The two main expressions are the hover in ground effect (known as HIGE) and the hover out of ground effect (known as HOGE). These values define the power margin available to the pilot on the day in question and are constantly evaluated during flight as conditions change. Environmental factors, such as temperature and density altitude, combined with mechanical factors (the actual performance of the engines installed on the vehicle) define the power available to the pilot for use. A HAC is expected to be able to evaluate the power they have available for approach to determine suitability.

Regarding the application of SWEEP for certification, this paper proposes a clearance envelope where the decision engine can exhibit nondeterministic behavior. If the vehicle reaches one of the edges, it will abort the approach and proceed to a predetermined point. The question is how to define the edges. Using SWEEP as an outline, a protocol can be developed based on a specification for keeping a vehicle within the clearance envelope. We then systematically examine the specification in an effort to ensure it satisfies the requirements in Sec. IV. This will serve as an artifact for flight clearance officials to accept the risk of allowing a decision engine to make a decision (landing) normally reserved for a qualified HAC.

B. Specification

For the limited purpose of defining a specification for the landing of a large rotorcraft in a CAL/LZ using guidance and control from an onboard decision engine, we elected to use a state machine specification [69] (Fig. 1). The state machine specification follows the various states required for the vehicle to transition through from the initial (or reset) point and being safe on deck. Table 1 details the various events that happen as the specification transfers from one state to another.

The transition states can be summarized as follows:

State A is the "initial/reset" state. At this point, the decision engine is at the start of the loop. Following a fuel check [to determine if the current state is above a predetermined bingo fuel (fuel required to return to a safe landing field)], it will begin the process of selecting a LZ and evaluating it against the SWEEP checklist. If the vehicle is below the predetermined bingo fuel, the decision engine reverts to the return to base (RTB) state, and it returns to base for more fuel before it attempts the find a valid LZ.

State B is the "conduct SWEEP checks to determine if selected LZ is a valid LZ" state. In this state, the decision engine selects a possible LZ and evaluates the SWEEP checks. If the selected LZ has a valid SWEEP check, the decision engine can then proceed to state C ("build ingress route"). If not, the decision engine retrogrades to state A (initial/reset).

State C is the build ingress route state. In this state, the decision engine builds an ingress route from the start point to a HOGE point. Providing it can be completed with the remaining fuel on board, avoid obstructions/traffic, and remain within the performance envelope of the vehicle, the ingress route is considered valid and the decision engine can proceed to state D ("monitor ingress"). If not, the decision engine retrogrades to state A (initial/reset).

State D is the monitor ingress state. In this state, the decision engine monitors the LZ and the performance parameters of the vehicle to ensure that SWEEP remains valid while the vehicle is transitioning from the start point to the HOGE point. Once the vehicle reaches the HOGE point, the decision engine shifts to state E ("HOGE over spot to LZ transition"). If SWEEP was to become invalid before the vehicle reached the HOGE point, the vehicle would execute the escape route, return to the initial/reset point, and retrograde to state A (initial/reset).

State E is the HOGE over spot to LZ transition state. In this state, the decision engine monitors the LZ and the performance parameters of the vehicle to ensure that SWEEP remains valid from HOGE to touchdown. If SWEEP remains valid, the vehicle will complete the mission (land safely). If SWEEP were to become invalid before touchdown, the vehicle would execute the escape route, return to the initial/reset point, and retrograde to state A (initial/reset).

This state machine specification can be considered a top level. Each of the events described in Table 1 has conditions and assumptions built into them. Some examples of the assumptions are the environmental conditions (weather, atmospheric conditions) and vehicle limitations (actual limits of the air vehicle). These conditions and assumptions must be valid for Fig. 1 to be a valid flight clearance artifact. Top-level assumptions become lower requirements.

As the specification in Fig. 1 represents a subset of the overall functionality of the aircraft, it has one defined start point (initial/reset

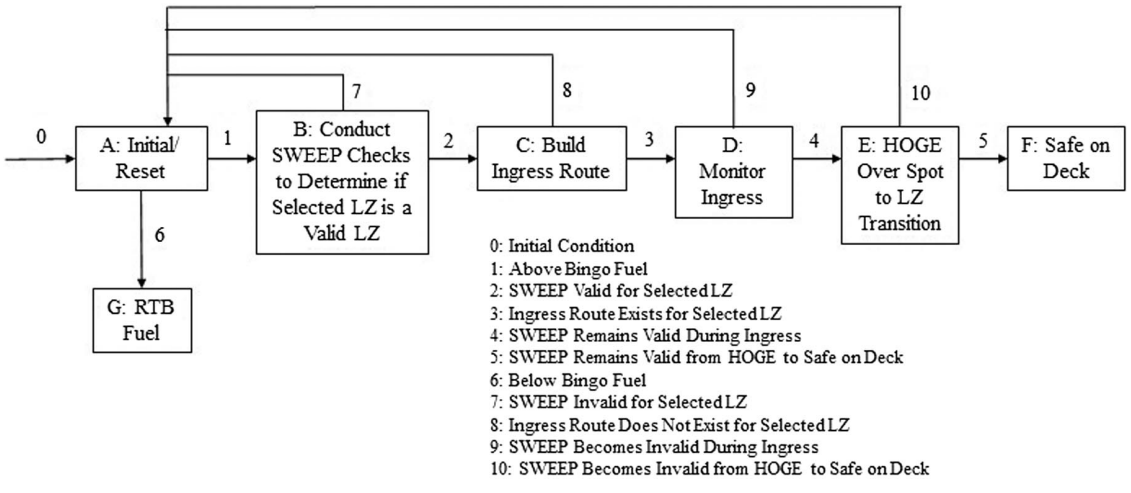

Fig. 1 State machine specification that details the decision process for an unmanned system to make a decision currently reserved for a qualified pilot.

Table 1 Event description for state machine specification that details the decision process for an unmanned system to make a decision currently reserved for a qualified pilot

Identification	From state	Events	To state
1	A	Above bingo fuel	B
2	A	Below bingo fuel	G
3	B	SWEEP valid for LZ	C
4	C	Ingress route exists for selected LZ	D
5	D	SWEEP remains valid during ingress	E
6	E	SWEEP remains valid from HOGE to safe on deck	F
7	B	SWEEP invalid for selected LZ	A
8	C	Ingress route does not exist for selected LZ	A
9	D	SWEEP becomes invalid during ingress	A
10	E	SWEEP becomes invalid from HOGE to safe on deck	A

state). From there, the decision engine executes the evaluation of possible landing locations until it either completes a safe landing ("safe on deck" state) or is forced to abandon the task due to fuel constraints (RTB state).

IV. Analysis of the Specification

In this section, we will begin with the state machine specification as it relates to controlling the unmanned system in its decision process. We will show consistency and completeness via an operational procedure table. We will then break down the various processes within the specification into propositions that must be held valid for the specification to be valid. The propositions will then be tracked and analyzed by a theory-proving software package to complete the analysis of the specification, detailing the decision process for an unmanned system to make a decision currently reserved for a qualified pilot.

Formal methods have been used for aircraft software verification and ultimately certification of aerospace software [16]. The power of formal methods lies in providing precise and unambiguous descriptions and mechanisms that facilitate the development of safety-critical systems in a more robust fashion [70]. By first developing a specification that tracks the various states for landing, then completing the formal methods activities (analyze specification for consistency/completeness, prove the behavior will satisfy the requirements (with assumptions), prove that a more detailed design implements a more abstract one [71]), TAEs can use the results as artifacts for certifying an autonomous controller to complete the CAL/LZ mission. The analysis in this section uses PVS, a theorem proving tool, to examine a high-level specification for an autonomous system in an attempt to certify that the system can complete tasks currently reserved for qualified pilots. This analysis is not a formal verification of the software but is rather a preliminary example of a path toward formal verification of such systems.

A. Operational Procedure Table

An operational procedure table was used to begin the analysis of the specification (Fig. 2). The variables along the top row represent the requirements for each associated landing segment (of flight) task (left column) required for the CAL/LZ mission. Each variable has its own assumptions (which would translate to requirements at lower levels). Each task is performed sequentially (top to bottom). Each variable is unknown until the associated segment is complete (changing the variable to a one or a zero). A common underlying assumption for all the variables is that the situational awareness

	VARIABLE								
		Initial SWEEP Checks for LZ							
LANDING SEGMENT TASK	Above Bingo Fuel	Suitable LZ (Size/Slope)	Winds Within Limits	Valid Elevation Data	Valid Escape Route	Favorable Power Margin	Valid Ingress Route	SWEEP Valid on Ingress to HOGE Point	SWEEP Valid HOGE to Land
Initial/Reset	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
RTB	0	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
(Return to Hold/Start if a 0 Exists in the Following		*	*	*	*	*	*	*	*
Conduct SWEEP Check to Determine if Selected LZ is a Valid LZ									
- Size/Slope	1	1	U	U	U	U	U	U	U
- Wind	1	1	1	U	U	U	U	U	U
- Elevation	1	1	1	1	U	U	U	U	U
- Escape Route	1	1	1	1	1	U	U	U	U
- Power Margin	1	1	1	1	1	1	U	U	U
Build Ingress Route	1	1	1	1	1	1	1	U	U
Monitor Ingress	1	1	1	1	1	1	1	1	U
HOGE Over Spot to LZ Transition (Safe on Deck)	1	1	1	1	1	1	1	1	1

* Return to initial/reset segment if before ingress, execute escape route to start point and then initial/reset segment, and clear variables

Possible Variable Values
1 Positive
0 Negative
U Has Not Reached Evaluation Yet
N/A Not Applicable

Fig. 2 Operational procedure table converting the state machine specification into the various tasks required for an unmanned system to make a decision currently reserved for a qualified pilot.

provided by the vehicle's sensors to the decision engine is adequate for the current conditions (not degraded to an unsatisfactory level by weather or malfunction).

The following are the variables and their underlying assumptions:

1) The first assumption is above bingo fuel, in which the vehicle is above the amount of fuel required to return to a safe landing area. This assumes the fuel management system is functioning properly and the decision engine is able to accurately measure the value.

2) The second assumption is suitable LZ (size/slope): The decision engine is able to choose a LZ that is suitable for the vehicle. It assumes the LZ requirements are programmed properly (size and slope) and can properly classify obstructions as threat or no threat.

3) The third assumption is winds within limits: The decision engine is able to compare the current wind conditions to the programmed limits for the vehicle. It assumes the wind limits are programmed properly (headwind and crosswind).

4) The fourth assumption is valid elevation data: The decision engine is able to determine its current mean sea level (MSL) altitude from its internal systems (some combination of the Global Positioning System, the inertial navigation system, and the internal pitot static system).

5) The fifth assumption is a valid escape route: The decision engine has developed an escape route that will return the vehicle to the start point and remain within safety limits. It assumes the safety limits are developed and defined within the programming of the decision engine.

6) The sixth assumption is a favorable power margin: The decision engine has defined the power margin (power required/power available) to be adequate for the LZ. It assumes the margin has been defined and programmed into the decision engine.

7) The seventh assumption is a valid ingress route: The decision engine is able to build an ingress route that will keep the vehicle free from collision and within the flight limits of the vehicle. It assumes the limits of the vehicle are programmed into the decision engine.

8) The eighth assumption is that SWEEP is valid on ingress to the HOGE point: The decision engine is able to continuously monitor the LZ during the approach to its HOGE point. Should the status of SWEEP change to invalid, the vehicle would need to abort the approach, execute the escape route, and transition to the reset point.

9) The ninth assumption is that SWEEP is valid from HOGE to land: The decision engine is able to continuously monitor the LZ during its landing through touchdown. Should the status of SWEEP change to invalid, the vehicle would need to abort the landing, execute the escape route, and transition to the reset point.

B. Consistency and Completeness

The operational procedure table [which contains cell values (1, 0, U, or N/A) of each requirement] was used to help define consistency and completeness. The table shows consistency by the fact that no two columns are operational for any combination of values for the variables because no two columns have the same cell values (at most, one outcome assigned under each possible scenario). The table shows completeness by the fact that for all values of variables, only one column is operational because all possible combinations of the variables are listed within the table, and no two columns are equal (some outcome assigned to every possible scenario) [72].

C. Theorem Proving Model

To prove that the system will complete the task, and show what the system will not do, the top-level requirements outlined in Fig. 2 were separated into three propositions [each of which have supporting propositions (e.g. Proposition 1.1 and Proposition 1.2 and Proposition 1.3 imply Proposition 1.0 is true)], which must remain true for the overall model of a successful landing to be valid. These propositions alone would not satisfy formally verifying the specification. That would require detailed formal analysis of the specification. This analysis would include validating all of the assumptions underneath the top level specification presented in this research. Which in turn would would require more explicit definitions than the Booleans presented and is beyond the scope of this research.

Proposition 1.0: The LZ is suitable for landing (all of the supporting propositions are true).

Proposition 1.1: The size of the LZ is adequate for the vehicle.

Proposition 1.2: The slope of the LZ is adequate for the vehicle.

Proposition 1.3: The LZ is clear of obstructions.

Proposition 2.0: The conditions for landing are suitable (all of the supporting propositions are true).

Proposition 2.1: The altitude of the LZ is within the envelope of the vehicle.

Proposition 2.2: The local wind conditions are within the envelope of the vehicle.

Proposition 2.3: The power margin is within acceptable parameters (nominally +10%).

Proposition 2.4: The decision engine can define a valid ingress route.

Proposition 2.5: The decision engine can define a valid egress/abort route.

Proposition 3.0: The approach and landing can be completed while maintaining suitable conditions (all of the supporting propositions are true).

Proposition 3.1: SWEEP can remain valid during the approach phase of the vehicle (from start to HOGE).

Proposition 3.2: SWEEP can remain valid from HOGE to landing.

D. PVS Model

After establishing the top-level propositions, we translated them into the theorem proving software package Prototype Verification System. PVS is a computer program that contains a theorem prover (symbolic engine that implements the deductive rules of a logic system). It allows us to use precise statements of logic such as lemmas and theorems. Proofs of logic formulas can be mechanically proven using the PVS theorem prover, which guarantees that every proof step is correct and that all possible cases of a proof are covered. Similar to the work performed by Narkawicz and Muñoz [73], all propositions presented were mechanically checked in PVS for logical correctness.

PVS has been used by NASA and other organizations for documentation of requirements for autonomous behavior for FAA certification [63]. The PVS specification (Fig. 3) is broken down into three sections (similar to the three main propositions). The first deals with the physical size of the LZ (Proposition 1.0). The second deals with the environmental conditions of the LZ (Proposition 2.0). The third deals with SWEEP remaining valid during approach to landing (Proposition 3.0). Using theorem proving software provides a repeatable, traceable model of the system's behavior that satisfies the specification. Figure 3 is a PVS top-level specification that illustrates the requirements for completing the initial SWEEP checks by a decision engine. While this model is not sufficient for formally verifying the specification, we use the model to illustrate how documenting the requirements through a formal process can provide TAEs with artifacts. These artifacts can be additional risk mitigation measures during the certification process for allowing an autonomous system to complete a task currently reserved for qualified pilots.

PVS offers the ability to analyze the propositions listed in Sec. IV.C within the interactive proving environment. While using the interactive environment, lemmas can be defined from sections of a PVS specification. An example of this would be an evaluation of the environmental condition of the LZ (wind and elevation). If either were outside of the defined parameters of a valid LZ, the selected LZ would be unsuitable due to conditions. An example of this lemma in PVS can be found in Fig. 4. For further details on the functionality and utility of PVS, we refer the reader to Ref. [74].

Theorem provers provide an analytical framework that can completely define the environment the vehicle will be operating in. While the model that is defined is a simplified model of the real world, it is robust enough that flight certification officials can use it to justify what the decision engine will not allow the vehicle to do. Thus, allowing the officials to approve the decision engine to exhibit nondeterministic behavior provided the behavior remains within the limits of its clearance envelope.

```
sweep_to_land: THEORY
BEGIN

suitable:       bool    % LZ is good for size, slope, and is obstruction free
too_small:      bool    % LZ is too small for the vehicle
bad_slope:      bool    % Slope of LZ is unsat for vehicle limits
obstructed:     bool    % LZ has too many obstructions
unsuitable:     bool =  NOT suitable    % LZ is unsuitable

good_conditions:  bool   % Condiitions work for elevaton, wind, Power margin,
                         % ingress route and egress route
too_high:         bool   % Elevation is to to high for vehicle
too_windy:        bool   % Wind exceeds vehicle limits
bad_power_margin: bool   % Less than a 10% power margin
no_ingress:       bool   % no valid ingress route
no_egress:        bool   % no valid egress route
bad_conditions:   bool = NOT good_conditions   % conditions not valid for approach

good_landing:     bool   % successful landing
abt_approach:     bool   % SWEEP goes invalid on approach
abt_landing:      bool   % SWEEP goes invalid on landing
abort:            bool = NOT good_landing   % vehicle aborts on approach

approach: bool = NOT unsuitable AND NOT bad_conditions
valid_landing: bool = approach AND NOT abort

cond_ax1: AXIOM too_small => unsuitable
cond_ax2: AXIOM bad_slope => unsuitable
cond_ax3: AXIOM obstructed => unsuitable
cond_ax5: AXIOM too_high => bad_conditions
cond_ax6: AXIOM too_windy => bad_conditions
cond_ax7: AXIOM bad_power_margin => bad_conditions
cond_ax8: AXIOM no_ingress => bad_conditions
cond_ax9: AXIOM no_egress => bad_conditions
```

Fig. 3 PVS specification for SWEEP checks to landing, detailing the decision process for an unmanned system to make a decision currently reserved for a qualified pilot.

```
%%%% LEMMA 3: Deals with Environmentals %%%%%
Lemma_3: LEMMA
(too_high OR too_windy)
IMPLIES
bad_conditions
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

Fig. 4 Lemma 3 deals with the environmental conditions of the LZ: If the elevation or the winds are out of limits, the LZ is not valid due to bad conditions.

For theorem provers, assumptions at a top level become requirements at lower levels. The specification outlined in Fig. 3 has a number of requirements embedded in the assumptions and can be broken up into three categories: LZ suitability, environmental conditions, and status during movement. Providing all three are satisfied, the specification would be valid and verified, and thus provide certification officials evidence of what the system would not do. Therefore, it can be used to prove the specified behavior will satisfy the requirements, given the assumptions.

For the PVS model to be a valid artifact for certification officials, it must be representative of actual conditions a vehicle would be faced with. To accomplish this, the assumptions built into the top level must be valid. These assumptions are what would define the real world situation. Weather and atmospheric conditions are built into the various states of the model as assumptions. Aircraft procedures and mechanics (such as aircraft size and operational limitations) are also built into the assumptions. Provided the assumptions are valid, a more detailed design implementation is implemented by a more abstract one (the PVS model in Fig. 3).

Figure 5 depicts the results of the PVS model against 11 separate hypothetical LZs. Of the 11 LZs, only one is acceptable for landing. LZ 1 is an ideal LZ because all 10 supporting propositions remain true. LZs 2 through 11 all have one supporting proposition that is false. The PVS specification shows that the final 10 LZs are not acceptable for landing.

V. Protocol

We used the analyzed specification as a baseline for the requirements the decision engine will need to fulfill in executing the CAL/LZ mission. By translating the state machine specification into a flowchart protocol, software designers can develop code based on the analyzed specification. The protocol has been broken into several steps that mirror what a qualified pilot would do while completing the CAL/LZ mission. The protocol translates the propositions into assessments. These steps can be traced directly to the supporting propositions presented in Sec. IV.C:

Size assessment can be traced to Proposition 1.1
Slope assessment can be traced to Proposition 1.2
Obstruction assessment can be traced to Proposition 1.3
Wind assessment can be traced to Proposition 2.2
Power margin assessment can be traced to Proposition 2.3
Elevation assessment can be traced to Proposition 2.1
Ingress assessment can be traced to Proposition 2.4
Escape route assessment can be traced to Proposition 2.5
Sweep valid ingress to HOGE can be traced to Proposition 3.1
Sweep valid HOGE to touchdown can be traced to Proposition 3.2

The protocol depicted in Fig. 6 satisfies the specification. It serves as an artifact for flight clearance officials when certifying a decision engine to make the decision on where to land a large rotorcraft (a task normally reserved for a fully qualified HAC). The various steps of the protocol can be completed autonomously using current day technology. Size, slope, and obstruction assessment can be accomplished via Light Detection and Ranging and Electro-Optical/InfarRed vision systems under challenging environmental condition to include degraded visual environments. Wind assessment can be accomplished by comparing the rotorcraft ground track against the current control inputs of the vehicle [75]. Onboard health monitoring systems can be programmed to assess the vehicle performance under all known operating conditions (to include degraded modes possible during a malfunction or

	Proposition 1.1 (Size)	Proposition 1.2 (Slope)	Proposition 1.3 (Obstructions)	Proposition 2.1 (Elevation)	Proposition 2.2 (Wind)	Proposition 2.3 (Power Margin)	Proposition 2.4 (Ingress)	Proposition 2 (Egress/Abort)	Proposition 3.1 (SWEEP Approach)	Proposition 3.2 (SWEEP HOGE to Land)	Proposition 1.0 (Suitability for Landing)	Proposition 2.0 (Conditions Suitability)	Proposition 3.0 (Movement)	Successful Landing
LZ 1	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	YES
LZ 2	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	NO
LZ 3	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	NO
LZ 4	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	NO
LZ 5	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	NO
LZ 6	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	NO
LZ 7	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	NO
LZ 8	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	NO
LZ 9	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE	FALSE	TRUE	NO
LZ 10	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE	FALSE	NO
LZ 11	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	TRUE	FALSE	TRUE	TRUE	FALSE	NO

Fig. 5 Depiction of 11 hypothetical LZs against the propositions listed in Sec. III.C and later detailed in the PVS model.

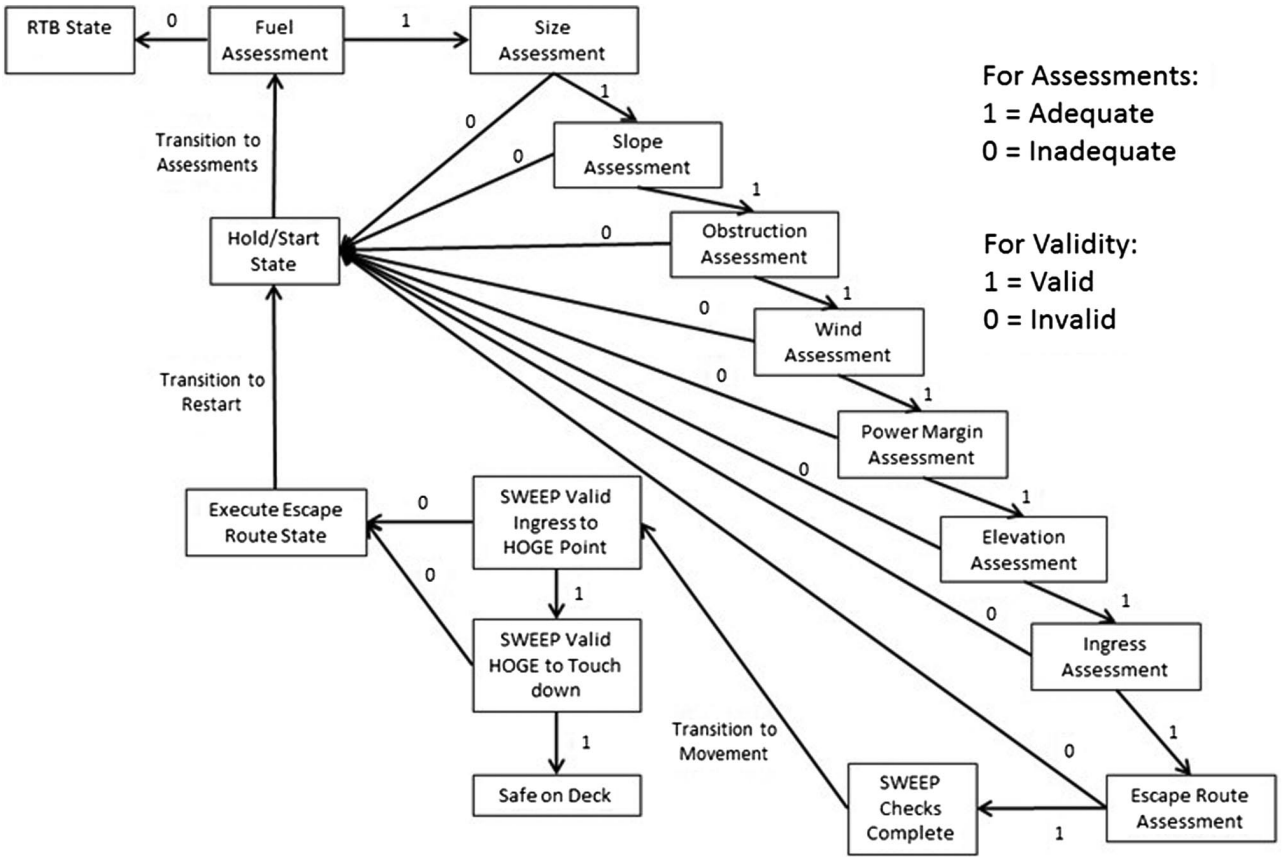

Fig. 6 Protocol that meets the requirements of the specification detailing the decision process for an unmanned system to make a decision currently reserved for a qualified pilot.

emergency situation). The performance characterization can be used during elevation, ingress, and escape route assessment.

As stated in earlier sections, this research focused on defining an envelope where the system can exhibit nondeterministic behavior. In the event that the LZ under evaluation does not pass all eight assessments (or SWEEP becomes invalid before touchdown), the system would return to the hold/start point and evaluate other possible LZs, in an attempt to find a valid LZ, until it no longer has enough fuel to complete the mission. Provided the LZ in question is within the limits established by the protocol (which defines the envelope where a system can exhibit nondeterministic behavior), it can land autonomously. This can be demonstrated by the system attempting to execute a landing on an empty football field, at sea level, in calm wind conditions. Assuming there were no stands or benches adjacent to the field, SWEEP would easily be valid between the 15 yd lines (the goal posts would obstruct from approximately the 15 yd line back to the end of each end zone). When executing the landing, the input conditions cannot guarantee the system would choose one landing spot on the field (as there will be multiple that satisfy the protocol). Under our methodology, the system would be certified to choose its landing point autonomously (cleared to land anywhere on the field that satisfies SWEEP).

This would allow the system to exhibit nondeterministic behavior, provided SWEEP is valid.

When assessing various LZs, the protocol performs eight assessments: each with a binary outcome. These eight binary outcomes translate into 256 possible combinations for each evaluated LZ. A LZ may be large enough for the vehicle in question (so, the first value would be a one) or it may not be large enough (so, the first value would be a zero). Of the 256 possibilities, only a LZ that passes all of the assessments (size, slope, obstruction, wind, power margin, elevation, ingress, and escape route) is conserved to be a valid LZ that the decision engine can select for landing. The assessments can be linked directly to SWEEP (and the specification) and a limited H-60 clearance envelope:

Assessment 1, size: It assumes H-60 and requires a 1.5 rotor arch (75-ft-diameter circle).

Assessment 2, slope: It assumes a limited H-60 slope envelope (5 deg forward/aft, and 2 deg port/starboard).

Assessment 3, obstruction: Within the circle defined in Assessment 1, there are no obstructions that would hinder a safe landing.

Assessment 4, wind: Assuming a limited H-60 wind envelope, it requires between 2 and 20 kt of headwind and less than 5 kt of crosswind.

Assessment 5, power margin: A positive 10% power margin can be maintained through approach to landing.

Assessment 6, elevation The LZ elevation is within the operating envelope of the vehicle (below 3000 ft MSL).

Assessment 7, ingress: A valid ingress route exists from the start point to the HOGE point.

Assessment 8, escape route: A valid escape route exists along the ingress route (to the HOGE point) that returns the vehicle to the reset point.

The results of the eight assessments can be displayed as a binary output. A subset of the 256 possible outcomes of the eight assessments is detailed in Table 2. If a LZ fails all eight assessments, its output would be 00000000 (outcome 1 in Table 2). If it only fails the wind (Assessment 4), its output would be 11101111 (outcome 240 in Table 2). If it only fails the size assessment (Assessment 1), its output would be 01111111 (outcome 128 in Table 2). If it only fails the obstruction and power margin assessments (Assessments 3 and 5), its output would be 11010111 (outcome 216 in Table 2). Only an LZ that passes all eight assessments with an output of 11111111 (outcome 256 in Table 2) would be valid for an attempted landing. After the decision engine chooses a LZ, it would then continuously assess SWEEP until it is safe on deck. While Table 2 may seem a trivial contribution, it is in fact considered an artifact that a TAE would use when accepting risk during the flight clearance process [8].

While analytically this appears to be a valid protocol for allowing a decision engine to make the decision currently reserved for HACs consistently, the following question remains: How can certification officials, within NAVAIR 4.0P, negate the current approved process (CNAF process for naval aviation) where a CO determines they have adequate trust in the HAC before full qualification? As a first step, we propose current senior officers become involved early in the process. These officers need to have, or have had, the authority to designate naval aviators as HACs. This is crucial for this effort because it can be used as an additional risk mitigation step to have qualified officers involved in the process.

Table 2 Depiction of five of the 256 possible outcomes of the eight protocol assessments

	Assessment number							
Outcome	1	2	3	4	5	6	7	8
1	0	0	0	0	0	0	0	0
128	0	1	1	1	1	1	1	1
216	1	1	0	1	0	1	1	1
240	1	1	1	0	1	1	1	1
256	1	1	1	1	1	1	1	1

The protocol (and related artifacts) was also shown to four naval commanders: all of which have been granted the authority by the CNAF for determining when a naval aviator can be qualified as a HAC. All agreed that, assuming the assumptions were valid, the assessments provided would be sufficient to qualify the decision engine to complete the task of landing in a CAL/LZ (a task that currently requires a HAC) safely.

Currently, all flight clearances for naval aircraft and subsystems are processed by the airworthiness process using approved V&V techniques/metrics detailed in NAVAIR Manual M-13034.1 [8]. While the evidence presented in this paper is not currently detailed in that manual, it has been submitted to flight clearance officials for consideration in the next revision of the naval airworthiness process. This may lead to a new process for clearing autonomous behavior under limited circumstances.

VI. Conclusions

To facilitate a flight clearance for a software intensive system, a clear definition of the requirements needs to be agreed upon before software development. This paper presented artifacts for a safety of flight certification in support of an autonomous controller that is designed to complete the unprepared CAL/LZ mission in a large rotorcraft. The actual path toward this certification does not currently exist.

This paper was a first step toward a methodology for clearing autonomous behavior to complete the CAL/LZ mission. The requirements normally reserved for a pilot to execute a safe landing on an unprepared CAL/LZ were defined. These requirements were developed through coordination with safety of flight clearance officials, the naval test and evaluation community, and fleet officials who currently certify pilots as fully qualified. A specification was developed. Then, the specification was systematically examined in an effort to ensure it satisfies the requirements. Finally, the analyzed specification was translated into a protocol and evaluated against all possible combinations of the conditions of a LZ. The protocol can then be used by software designers when developing the decision engine of the autonomous vehicle. All of the artifacts developed in this paper can be used as certification evidence for a safety of flight clearance of autonomous behavior.

The logical next step is a limited M&S of the protocols/control laws to insure that the vehicle will function as intended. This will be in an attempt to show the system will only display nondeterministic behavior while it is within the clearance envelope, and it would serve as a risk mitigation step before the actual flight test. Following M&S, a design of experiments for flight tests needs to be developed (flight-test plan). Most conventional developmental flight-test techniques are designed for a pilot to test an unproven system. In this case, test points will need to be developed that demonstrate in an operational-relevant environment that the test vehicle can produce results similar to the demonstrated behaviors in modeling and simulation. This is the test plan of the specification. Following the flight test, a summary of the test results would be the final piece of data that flight clearance officials would need to certify a machine to make decisions currently reserved for qualified pilots.

This paper has presented top-level evidence from the requirements definition to the protocol definition. Future work that focuses on each of the subtasks would be invaluable to certification officials. They could use this to define/refine an envelope by which to certify what a decision engine would not do. Future works to concentrate on developing tools to complete the remaining steps of the proposed methodology are critical before an actual autonomous controller's development. Without a clear path, and mature flight clearance tools, certification of controllers to accomplish tasks currently reserves for pilots will continue to be limited in their scope and limited in their real world application.

References

[1] "Teal Group Releases Study, 'World Unmanned Aerial Vehicle Systems, Market Profile and Forecast 2015'," Manufacturing Close-Up, Aug. 2015.

[2] "FAA Creates Commercial Drone Pilot's License (Part 107): American Drone Operators Will Be Able to Apply for a Drone Pilot's License Starting in August 2016," Marketwired, June 2016.

[3] Weibel, R. E., and Hansman, R. J., "Safety Considerations for Operation on Unmanned Aerial Vehicles in the National Airspace System," Massachusetts Inst. of Technology International Center for Air Transportation TR-ICAT 2005-1, Cambridge, MA, 2005.

[4] "BCAR Section A: Airworthiness Procedures Where the CAA has Primary Responsibility for Type Approval of a Product," Civil Aviation Authority Civil Aviation Publication 553, London, Oct. 2011.

[5] "Naval Air Training and Operations Procedures Standardization (NATOPS) General Flight and Operating Instructions," Commander Naval Air Forces, North Island, CA, 2016.

[6] "Airborne Software Assurance," U.S. Dept. of Transportation, Federal Aviation Administration Advisory Circular 20-115C, June 2013.

[7] Webster, M., Cameron, N., Fisher, M., and Jump, M., "Generating Certification Evidence for Autonomous Unmanned Aircraft Using Model Checking and Simulation," *Journal of Aerospace Information Systems*, Vol. 11, No. 5, 2014, pp. 258–279. https://doi.org/10.2514/1.I010096

[8] "NAVAIR Airworthiness and Cybersafe Process Manual, NAVAIR Manual M-13034.1," Naval Air Systems Command, 2016.

[9] "NATOPS Flight Manual, MH-60R Helicopter," Chief of Naval Operations and Under the Direction of the Commander, Naval Air Systems Command, 2014.

[10] Eldridge, C., "Electronic Eyes for the Allies: Anglo-American Cooperation on Radar Development During World War II," *History and Technology*, Vol. 17, No. 1, 2000, pp. 1–20. https://doi.org/10.1080/07341510008581980

[11] Bradley, M., Nealiegh, M., Oh, J. S, Rothberg, P., Rodd, S. A., and Rodd, S. A., "Combat Casualty Care and Lessons Learned From the Last 100 Years of War," *Current Problems in Surgery*, Vol. 54, No. 6, 2017, pp. 315–351. https://doi.org/10.1067/j.cpsurg.2017.02.004.

[12] Giffard, H., "Engines of Desperation: Jet Engines, Production and New Weapons in the Third Reich," *Journal of Contemporary History*, Vol. 48, No. 4, 2013, pp. 821–844. https://doi.org/10.1177/0022009413493943.

[13] Rodd, S. A., "Government Laboratory Technology Transfer: Process and Impact Assessment," Ph.D. Thesis, Virginia Polytechnic Inst. and State Univ., Blacksburg, VA, 1998.

[14] Alexander, C., "Development of the Combiner-Eyepiece Night-Vision Goggle," *Proceedings of SPIE 1290*, Helmet-Mounted Displays II, Oct. 1990. https://doi.org/10.1117/12.20951

[15] Clark, M., Alley, J., Deal, P. J., Depriest, J. C., Hansen, E., Heitmeyer, C., Nameth, R., Steinberg, M., Turner, C., Young, S., Ahner, D., Alonzo, K., Bodt, B. A., Friesen, P. F., Horris, J., Hoffman, J. A., Gross, K. H., Humphrey, L., Childers, M., and Corey, M., "Autonomy Community of Interest (COI) Test and Evaluation, Verification and Validation (TEVV) Working Group: Technology Investment Strategy 2015-2018," Office of the Assistant Secretary of Defense for Research and Engineering (USD (R&E)), The Pentagon, Washington, D.C., 2015.

[16] Gross, K. H., Clark, M. A., Hoffman, J. A., Swenson, E. D., and Fifarek, A. W., "Run-Time Assurance and Formal Methods Analysis Nonlinear System Applied to Nonlinear System Control," *Journal of Aerospace Information Systems*, Vol. 14, No. 4, 2017, pp. 232–246. https://doi.org/10.2514/1.I010471

[17] Ashokkumar, C. R., and York, G. W., "Trajectory Transcriptions for Potential Autonomy Features in UAV Maneuvers," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2016-0380, 2016. https://doi.org/10.2514/6.2016-0380

[18] Humphreys, C. J., Cobb, R., Jacques, D. R., and Reeger, J. A., "Dynamic Re-Plan of the Loyal Wingman Optimal Control Problem," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2017-1744, 2016. https://doi.org/10.2514/6.2017-1744

[19] Avram, R., Zhang, X., Muse, J. A., and Clark, M., "Nonlinear Adaptive Control of Quadrotor UAVs with Run-Time Safety Assurance," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2017-1896, 2017. https://doi.org/10.2514/6.2017-1896

[20] Gross, K. H., Fifarek, A. W., and Hoffman, J. A., "Incremental Formal Methods Based Design Approach Demonstrated on a Coupled Tanks Control System," *2016 IEEE 17th International Symposium on High Assurance Systems Engineering (HASE)*, Inst. of Electrical and Electronics Engineers, New York, 2016, pp. 181–188. https://doi.org/10.1109/HASE.2016.16

[21] Abraham, J., "Verification and Validation Spanning Models to Code," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2015-1593, 2015. https://doi.org/10.2514/6.2015-1593

[22] Blanchette, S., "Giant Slayer: Will You Let Software Be David to Your Goliath System?" *Journal of Aerospace Information Systems*, Vol. 13, No. 10, 2016, pp. 407–417. https://doi.org/10.2514/1.I010448

[23] Muli, M., Moudgal, V., and Allen, J., "Best Practices and Recommendations for Model-Based Development Process," SAE TP 2015-01-2529, Warrendale, PA, 2015. https://doi.org/10.4271/2015-01-2529

[24] Israelsen, B. W., Ahmed, N. R., Center, K., Green, R., and Bennett, W., "Towards Adaptive Training of Agent-Based Sparring Partners for Fighter Pilots," *AIAA Information Systems-AIAA Infotech @ Aerospace*, AIAA Paper 2017-0343, 2017. https://doi.org/10.2514/6.2017-0343

[25] Fisher, M., Dennis, L., and Webster, M., "Verifying Autonomous Systems," *Communications of the Association for Computing Machinery*, Vol. 56, No. 9, 2013, pp. 84–93. https://doi.org/10.1145/2494558

[26] Tobias, E., Tischler, M., Berger, T., and Hagerott, S. G., "Full Flight-Envelope Simulation and Piloted Fidelity Assessment of a Business Jet Using a Model Stitching Architecture," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2015-1594, 2015. https://doi.org/10.2514/6.2015-1594

[27] Eisemann, U., and Allen, J. L., "New Requirement-Definition and Verification Techniques According to DO-178C, DO-331, and DO-333," *AIAA Infotech @ Aerospace*, AIAA Paper 2016-0223, 2016. https://doi.org/10.2514/6.2016-0223

[28] Berger, T., Tischler, M., Hagerott, S. G., Cotting, M. C., Gray, W. R., Gresham, J., George, J., Krogh, K., D'Argenio, A., and Howland, J., "Development and Validation of a Flight-Identified Full-Envelope Business Jet Simulation Model Using a Stitching Architecture," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2018-0525, 2017. https://doi.org/10.2514/6.2018-0525

[29] Walker, S. M., Shan, J., and Liu, L., "DIMMACSS-Stage: A Distributed Intelligence Model for a Multi-Agent Control System Using Simulink and the Stage Robotic Simulator," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2014-0639, 2014. https://doi.org/10.2514/6.2014-0639

[30] Hangal, S. A., Tak, B., and Arya, H., "Distributed Hardware-in-Loop Simulations for Multiple Autonomous Aerial Vehicles," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2015-0151, 2015. https://doi.org/10.2514/6.2015-0151

[31] Mohamedy, A. E.-D. S., Aly, A. M., and Elnashar, A. H., "Modeling and Simulation Hardware-in-the-Loop for Unmanned Aerial Vehicle," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2016-1428, 2016. https://doi.org/10.2514/6.2016-1428

[32] Baier, C., and Katoen, J., *Principles of Model Checking*, MIT Press, Cambridge, MA, 2008.

[33] Kane, A., "Runtime Monitoring for Safety-Critical Embedded Systems," Ph.D. Dissertation, Electrical and Computer Engineering, Carnegie–Mellon Univ., Pittsburgh, PA, 2015.

[34] Coombes, M., Chen, W.-H., and Render, P., "Site Selection During Unmanned Aerial System Forced Landings Using Decision-Making Bayesian Networks," *Journal of Aerospace Information Systems*, Vol. 13, No. 12, 2016, pp. 491–495. https://doi.org/10.2514/1.I010432

[35] Gross, K. H., Clark, M., Hoffman, J. A., Fifarek, A., Rattan, K., Swenson, E., Whalen, M., and Wagner, L., "Formally Verified Run Time Assurance Architecture of a 6U CubeSat Attitude Control System," *AIAA Infotech @ Aerospace*, AIAA Paper 2016-0222, 2016. https://doi.org/10.2514/6.2016-0222

[36] Torens, C., Adolf, F., Faymonville, P., and Schirmer, S., "Towards Intelligent System Health Management Using Runtime Monitoring," *AIAA Information Systems-AIAA Infotech @ Aerospace*, AIAA Paper 2017-0419, 2017. https://doi.org/10.2514/6.2017-0419

[37] Schierman, J., Ward, D., Dutoi, B., Aiello, A., Berryman, J., DeVore, M., Storm, W., and Wadley, J., "Run-Time Verification and Validation for Safety-Critical Flight Control Systems," *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA Paper 2008-6338, 2008. https://doi.org/10.2514/6.2008-6338

[38] Lichter, M., Bateman, A., and Balas, G., "Flight Test Evaluation of a Run-Time Stability Margin Estimation Tool," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2009-6257, 2009. https://doi.org/10.2514/6.2009-6257

[39] Rabideau, G., Chien, S., and Mclaren, D., "Onboard Run-Time Goal Selection for Autonomous Operations," *SpaceOps 2010 Conference*,

AIAA Paper 2010-2203, 2010.
https://doi.org/10.2514/6.2010-2203
[40] Sababha, B., Yang, H. C., and Rawashdeh, O., "An RTOS-Based Run-Time Reconfigurable Avionics System for UAVs," *AIAA Infotech@Aerospace 2010*, AIAA Paper 2010-3414, 2010.
https://doi.org/10.2514/6.2010-3414
[41] Aiello, M., Berryman, J., Grohs, J., and Schierman, J., "Run-Time Assurance for Advanced Flight-Critical Control Systems," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2010-8041, 2010.
https://doi.org/10.2514/6.2010-8041
[42] Wong, E., Schierman, J. D., Schlapkohl, T., and Chicatelli, A., "Towards Run-Time Assurance of Advanced Propulsion Algorithms," *50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA Paper 2014-3636, 2014.
https://doi.org/10.2514/6.2014-3636
[43] Berezin, S., "Model Checking and Theorem Proving: A Unified Framework," Ph.D. Dissertation, School of Computer Science, Carnegie–Mellon Univ., Pittsburgh, PA, 2002.
[44] Webster, M., Cameron, N., Jump, M., and Fisher, M., "Towards Certification of Autonomous Unmanned Aircraft Using Formal Model Checking and Simulation," *Infotech@Aerospace 2012*, AIAA Paper 2012-2573, 2012.
https://doi.org/10.2514/6.2012-2573
[45] Good, N., Aboutalib, O., Thai, B., Yamaoka, N., Kim, C., Wilkinson, C., and Findlay, D., "Validation Process of the Physics-Based Modeling of Navigation Sensors for Sea-Based Aviation Automated Landing," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2016-1919, 2016.
https://doi.org/10.2514/6.2016-1919
[46] Bakera, M., Margaria, T., Renner, C. D., and Steffen, B., "Game-Based Model Checking for Reliable Autonomy in Space," *Journal of Aerospace Computing, Information, and Communication*, Vol. 8, No. 4, 2011, pp. 100–114.
https://doi.org/10.2514/1.32013
[47] Sirigineedi, G., Tsourdos, A., White, B., and Zbikowski, R., "Kripke Modelling and Model Checking of a Multiple UAV System Monitoring Road Network," *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper 2010-7582, 2010.
https://doi.org/10.2514/6.2010-7582
[48] Humphrey, L., "Model Checking UAV Mission Plans," *AIAA Modeling and Simulation Technologies Conference*, AIAA Paper 2012-4723, 2012.
https://doi.org/10.2514/6.2012-4723
[49] Verzino, G., "Model Checking Driven Simulation of Sat Procedures," *AIAA SpaceOps 2012 Conference*, 2012.
https://doi.org/10.2514/6.2012-1275611
[50] Humphrey, L., and Patzek, M., "Model Checking Human UAV Mission Plans," *AIAA Guidance, Navigation, and Control (GNC) Conference*, AIAA Paper 2013-5183, 2013.
https://doi.org/10.2514/6.2013-5183
[51] Torens, C., and Adolf, F., "Using Formal Requirements and Model-Checking for Verification and Validation of an Unmanned Rotorcraft," *AIAA Infotech @ Aerospace*, AIAA Paper 2015-1645, 2015.
https://doi.org/10.2514/6.2015-1645
[52] Hansen, J. P., and Wrage, L., "Verification of Real-Time Systems Using Statistical Model Checking," *AIAA Infotech @ Aerospace*, AIAA Paper 2015-1866, 2015.
https://doi.org/10.2514/6.2015-1866
[53] Nandiganahalli, J. S., Lee, S., and Hwang, I., "Flight Deck Mode Confusion Detection Using Intent-Based Probabilistic Model Checking," *AIAA Information Systems-AIAA Infotech @ Aerospace*, AIAA Paper 2017-0345, 2017.
https://doi.org/10.2514/6.2017-0345
[54] Ouimet, M., "Formal Software Verification: Model Checking and Theorem Proving," Embedded Systems Lab., Massachusetts Inst. of Technology TR ESL-TIK-00214, Cambridge, MA, 2008, http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.90.5587&;=rep1&=pdf [retrieved 11 Oct. 2020].
[55] Sutcliffe, G., Denney, E., and Fischer, B., "Practical Proof Checking for Program Certification," NASA Arc., 2005, https://ti.arc.nasa.gov/m/profile/edenney/papers/escar05.pdf [retrieved 19 Dec. 2017].
[56] Muñoz, C., Dutle, A., Narkawicz, A., and Upchurch, J., "Unmanned Aircraft Aystems in the National Airspace System: A Formal Methods Perspective," *ACM SIGLOG News*, Vol. 3, No. 3, 2016, pp. 67–76.
https://doi.org/10.1145/2984450.2984459
[57] Goodloe, A., Gunter, C. A., and Stehr, M.-O., "Formal Prototyping in Early Stages of Protocol Design," *Proceedings of the 2005 Workshop on Issues in the Theory of Security (WITS '05)*, Association for Computing Machinery, New York, 2005, pp. 67–80.
https://doi.org/10.1145/1045405.1045413
[58] Jiang, Y., Liu, J., Dowek, G., and Ji, K., "SCTL: Towards Combining Model Checking and Proof Checking," 2016, https://www.researchgate.net/publication/304547954_SCTL_Towards_Combining_Model_Checking_and_Proof_Checking.
[59] Asokan, S., Kumar, G. S., and Lal, N. J., "Modeling of ALFA Programs Using PVS Theorem Prover," *International Conference on Advances in Recent Technologies in Communication and Computing, ARTCom*, Inst. of Electrical and Electronics Engineers, New York, 2009, pp. 373–375.
https://doi.org/10.1109/ARTCom.2009.134
[60] Muñoz, C., "Formal Methods in Air Traffic Management: The Case of Unmanned Aircraft Systems (Invited Lecture), Vol. 9399, Lecture Notes in Computer Science," *Proceedings of the 12th International Colloquium on Theoretical Aspects of Computing*, Springer, New York, 2015, pp. 58–62.
[61] Jenie, Y. I., Kampen, E.-J. V., Ellerbroek, J., and Hoekstra, J. M., "Safety Assessment of Unmanned Aerial Vehicle Operations in an Integrated Airspace," *AIAA Infotech @ Aerospace*, AIAA Paper 2016-1000, 2016.
https://doi.org/10.2514/6.2016-1000
[62] Guarro, S., Yau, M. K., Ozguner, U., Aldemir, T., Kurt, A., Hejase, M., and Knudson, M., "Formal Framework and Models for Validation and Verification of Software-Intensive Aerospace Systems," *AIAA Information Systems-AIAA Infotech @ Aerospace*, AIAA Paper 2017-0418, 2017.
https://doi.org/10.2514/6.2017-0418
[63] Muñoz, C., and Narkawicz, A., "Formal Analysis of Extended Well-Clear Boundaries for Unmanned Aircraft," *Conference Proceedings of NASA Center for AeroSpace Information (CASI)*, 2016.
[64] Ghatas, R. W., Jack, D. P., Tsakpinis, D., Vincent, M. J., Sturdy, J. L., Muñoz, C. A., Hoffler, K. D., Dutle, A. M., Myer, R., DeHaven, A. M., Lewis, T., and Arthur, K. E., "Unmanned Aircraft Systems Minimum Operational Performance Standards End-to-End Verification and Validation (E2-V2) Simulation," NASA TM-2017-20780, 2017.
[65] Narkawicz, A., Munoz, C., and Dutle, A., "Coordination Logic for Repulsive Resolution Maneuvers," *16th AIAA Aviation Technology, Integration, and Operations Conference*, AIAA Paper 2016-3156, 2016.
https://doi.org/10.2514/6.2016-3156
[66] "Preliminary NATOPS Flight Manual, MQ-8B/C Unmanned Aircraft System," Chief of Naval Operations and Under the Direction of the Commander, Naval Air Systems Command, Patuxent River Maryland, 2019.
[67] Donley, M. B., and Schwartz, N. A., "Technology Horizons, A Vision for Air Force Science and Technology 2010–30," Office of the U.S. Air Force Chief Scientist, Air Univ. Press, Air Force Research Inst., Maxwell AFB, AL, 2010.
[68] Tate, D., Grier, R., Martin, C., Moses, F. L., Sparrow, D. A., Edmonson, J. R., Chaki, S., Scheidt, D. H., Scheidt, D. H., Piatko, C. D., Davis, D., and Stausberger, D., "A Framework for Evidence-Based Licensure of Adaptive Autonomous Systems: Technical Areas," Inst. for Defense Analyses, IDA Paper P-5325, Log H 16-000680, Alexandria, VA, 2016.
[69] Domínguez, E., Pérez, B., Rubio, Á. L., and Zapata, M. A., "A Systematic Review of Code Generation Proposals from State Machine Specifications," *Information and Software Technology*, Vol. 54, No. 10, 2012, pp. 1045–1066.
https://doi.org/10.1016/j.infsof.2012.04.008
[70] Blooshi, M. A., Jafer, S., and Patel, K., "Review of Formal Agile Methods as Cost-Effective Airworthiness Certification Processes," *Journal of Aerospace Information Systems*, Vol. 15, No. 8, 2018, pp. 471–484.
https://doi.org/10.2514/1.I010601
[71] Kumar, S., Suryavanshi, R. S., and Chandra, G., "Formal Methods: Techniques and Languages for Software Development," *International Journal of Engineering Science and Advanced Research*, Vol. 1, No. 1, 2015, pp. 35–42.
[72] Hoover, D., Gauspari, D., and Humenn, P., "Applications of Formal Methods to Specification and Safety of Avionics Software," NASA CR, Odyssey Research Assoc., Ithaca, NY, 1996, https://ntrs.nasa.gov/citations/19960023949 [retrieved 19 Dec. 2017].
[73] Narkawicz, A., and Muñoz, C. A., "Formal Verification of Conflict Detection Algorithms for Arbitray Trajectories," *Reliable Computing*, Vol. 17, June 2012, pp. 209–237, https://interval.louisiana.edu/reliable-computing-journal/volume-17/reliable-computing-17-pp-209-237.pdf [retrieved 11 Oct. 2020].
[74] Crow, J., Owre, S., Rushby, J., Shankar, N., and Srivas, M., "A Tutorial Introduction to PVS," *WIFT '95: Workshop on Industrial-Strength Formal Specification Techniques*, Boca Raton, FL, April 1995.
[75] "Deliverables for Contract N00014-12-C-0671, Aurora Flight Sciences," Releases by NAVAIR Center for Autonomy, 2018.

E. Atkins
Associate Editor