

Guidance for Designing Safety into Urban Air Mobility: Hazard Analysis Techniques

Mallory S. Graydon* and Natasha A. Neogi†
NASA Langley Research Center, Hampton, VA 23681, USA

Kimberly S. Wasson‡
Federated Safety, LLC, Crozet, VA 22932, USA

The national airspace system is exceptionally safe, specifically in terms of commercial air traffic operations. The introduction of innovative aircraft (such as electric vertical takeoff and landing vehicles) undergoing novel operations (such as for an urban air passenger carrying mission) is a potential disruptor to the current means of regulating and ensuring the safety of air travel. The aviation industry and associated regulatory bodies are adapting their approaches to assuring system safety to enable these new paradigms. However, any system safety analysis requires that a hazard assessment be performed. In this work, we consider the Functional Hazard Assessment (FHA) and Systems Theoretic Process Analysis (STPA) techniques for hazard assessment and evaluate whether they can be used in a complementary fashion for regulatory approval purposes. We perform an FHA and an STPA on an electric vertical takeoff and landing (eVTOL) vehicle undergoing an urban air mobility (UAM) passenger carrying reference scenario and present excerpts of this analysis. We then draw parallels between the techniques and highlight elements where they naturally reinforce to each other's results, specifically in the consideration of hazard severity with respect to flight phases, the design of hazard mitigations, and the applicability of the results to all types of regulatory approvals (e.g., type certification, operational approval, and crew training).

I. Introduction

Urban air mobility (UAM) for passenger carrying operations promises to revolutionize transportation by potentially reducing urban congestion and passenger delay, enabling flexible response to changing transportation needs, and creating new market opportunities. Vehicles capable of vertical takeoff and landing are envisioned to operate from sites that are capable of loading and unloading passengers and handling multiple aircraft. These operations will also require enhanced air traffic management and conflict detection and avoidance capabilities in order to ensure safe, scalable operations.

The UAM community, which includes many new entrants, faces a regulatory environment in flux. The community is presently discussing several options for certifying electric vertical takeoff and landing (eVTOL) aircraft in the USA: (1) as a Part 23 normal category airplane with special conditions; (2) as a Part 27 normal category rotorcraft with special conditions; (3) as a special class aircraft under Part 21.17(b); or (4) as any number of other less prominent but still possible options [1]. However, some safety principles and techniques remain fundamental, regardless of the certification path chosen. Designing a safe UAM operational system begins with a system safety analysis.

In this work, we apply system safety analysis techniques in the context of a fixed intra-metro air shuttle reference mission for UAM. This operation is regarded as being an intermediate state of UAM, whereby passengers are transported between a fixed number of destinations, under medium density and complexity operations. Since operations will need to be conducted under a wide range of weather conditions, low visibility conditions must be accounted for in the required procedures and technological solutions, shaping the need for precise environmental assumptions. Furthermore, these operations will be sharing the skies with diverse entrants with differing levels of performance and equipage, as well as legacy operations, and will pose challenges to infrastructure assumptions. Details about the reference mission are available in a separate technical report [2].

We examine two hazard analysis (HA) techniques as a means of assessing functional and operational hazards in the context of this UAM operation: Functional Hazard Assessment (FHA) and Systems Theoretic Process Analysis (STPA).

*Research Scientist, Safety Critical Avionics Systems Branch.

†Research Scientist, Safety Critical Avionics Systems Branch, Associate Fellow.

‡Principal, Federated Safety, LLC.

Our examination highlights the necessity of considering both the concept of operations, as well as the platform (e.g., vehicle and ground control infrastructure) simultaneously, in order to adequately scope the type and severity of hazards. The hazard analyses can then be used to focus in on a relevant class of hazards (e.g., loss of separation), and refine the chosen hazard down to detailed causes and effects, thereby enabling the development of potential mitigations in order to reduce the overall risk posed by the operation to a potentially acceptable level. We contend that FHA and STPA can be used in a complementary fashion to gain fuller insight into a system's hazards and potential means of addressing them.

In section II, we consider the regulatory aspects involved in enabling a UAM operation, and how this will influence the goals of the hazard analyses performed. In section III, we give some background on system safety and outline the FHA and STPA hazard analysis techniques. In section IV, we describe a reference aircraft and operating scenario and walk through how we applied FHA and STPA to those, giving examples to illustrate the process. In section V, we present observations about the hazard analysis processes we used and discuss the effectiveness of the techniques at meeting certain classes of regulatory requirements and the advantages of employing multiple hazard analysis methods. Finally, we give conclusions and discuss future work in section VI.

II. Regulatory Considerations for Urban Air Mobility

Urban Air Mobility stands to change the way in which people conduct their daily lives, by tightly integrating multiple modes of transport, potentially alleviating burdensome congestion around key locations, and reducing arduous commuting times. However, the class of vehicles required for these operations (e.g., eVTOLs) and the nature of the operation (e.g., continuous overflight of populated areas) raise issues regarding the safety of parties such as passengers onboard the aircraft, other users of the airspace, and third parties such as overflown pedestrians. One of the key hurdles to realizing UAM operations will involve gaining access to the National Airspace System (NAS). This will require establishing the safety of the vehicles and proposed operation to regulators to gain the requisite certification, operational approval and access to air traffic services (if/when needed).

A. Certification Considerations

The Federal Aviation Administration (FAA) is in the process of developing guidance regarding the potential paths to certification for the novel configurations that inhabit the UAM operational space, several of which were outlined in section I.

Some applicants are seeking a path to certification that leverages the new Part 23 rewrite. [3] The traditional, more prescriptive regulations for the certification of normal category aircraft (e.g., general aviation aircraft, small commuter aircraft etc.) were rewritten to embrace a performance based approach, which will hopefully help accelerate certification efforts for small, fixed wing aircraft. Applicants with eVTOL aircraft, which spend the majority of their time in forward, fixed wing flight, are hoping to tailor their certification approach using Part 23 as a starting basis. This approach will require customizing regulations and creating special conditions in order to encompass features that deal with issues related to vertical takeoff and landing, as well as electric propulsion. Similarly, it is possible for applicants to seek certification using Part 27 [4], which is used for normal category rotorcraft, as the starting basis. Again, special conditions and tailored regulations will likely be required with this approach. All of these approaches have merit, and the evolving landscape of vehicle designs and regulatory policy will likely lead to applicants trialing all of these paths to certification. This will enable the FAA to evaluate which means proves most efficacious for specific classes of eVTOLs; this may depend both on the design of the eVTOL, as well as its intended use or means of operation.

Additionally, a natural path to certification for electric vertical takeoff and landing (eVTOL) vehicles would require applicants to initially employ Federal Aviation Regulation (FAR) 21.17(b), which applies to special classes of vehicles. The regulation is as follows:

For special classes of aircraft, including the engines and propellers installed thereon (e.g., gliders, airships, and other nonconventional aircraft), for which airworthiness standards have not been issued under this subchapter, the applicable requirements will be the portions of those other airworthiness requirements contained in Parts 23, 25, 27, 29, 31, 33, and 35 found by the FAA to be appropriate for the aircraft and applicable to a specific type design, or such airworthiness criteria as the FAA may find provide an equivalent level of safety to those parts. [1]

Each vehicle would be handled on a case-by-case basis, and this allows the applicant to tailor the type certification basis of the vehicle. The advantage of the FAR Part 21.17(b) approach is that it can accommodate any particular type design immediately, often leveraging relevant portions of existing standards, without waiting for the standards

development process to take place. However, the disadvantages of this approach are: (1) It is not possible to take advantage of reciprocal airworthiness agreements with other countries, in order to gain certification in, and eventual access to, airspace outside of the US; (2) it is much more labor-intensive for the certifying authority; and (3) since all of the criteria are not known upfront, it is more difficult for the applicant and the regulator to plan for the cost of the certification effort given that there may be a great deal of uncertainty in the outcome. Furthermore, an aircraft which certifies under FAR 21.17(b) does not necessarily have full operational access to the NAS. From a standards perspective, there is a general trade-off between operating limitations and rigor of the certification standard.

B. Operational Considerations

There are two major sets of operational considerations for passenger carrying UAM operations: (1) considerations relating to operations occurring in the shared airspace (Air Operator or Operational Certification) and (2) considerations relating to personnel (e.g., pilots, maintenance personnel, etc.) [5]. While design concerns relating to the safety of the aircraft are considered throughout the aircraft certification process, operational considerations are addressed through the operational approval process, where the type operation the vehicle will perform, as well as the crew who will be responsible for the aircraft, are assessed. For the purposes of the following analyses, we focus primarily on design considerations, due to problem scope and maturity of the models.

For UAM undergoing low or moderate tempo operations, it is generally thought that an Operational Certificate would be sought under conditions similar to today's 14CFR Part 135 Air Carrier and Operator Certification [5]. These certificates are divided into (1) Type, (2) Kind and (3) Scope. Since initial UAM operations will likely be intrastate transportation in an on-demand fashion, the type of certificate sought will be an Operating Certificate, and the kind will be "On Demand". The scope of a Part 135 Operating Certificate is determined by the number of pilots involved in the operation. Due to the potential ambiguity of the nature of this number, it is likely that a Standard scope will need to be sought (though in initial phases, a Single Pilot scope may be sufficient). It should be noted that, for a Standard scope, it is necessary that an Operational Specification be submitted for each type of operation to be pursued (normally accompanied by a Minimum Equipment List). The applicant must apply, qualify, and be granted FAA authorization through these Operational Specifications for each type of operation they wish to conduct (see Figure 1). *This operational approval is generally regarded to be a challenge to the fielding of UAM, as it is unclear as to what level of risk mitigation will be necessary.* Furthermore, standard Part 135 operators are required to develop and maintain manuals, training programs, and have the required management positions.

It is important to note that operational elements, including crew certification, are vital to establishing the safety of the vehicle and operations. Thus, any analysis performed, or argument made in order to establish the overall safety of the system must take into account factors related to vehicle design, as well as crew and operational performance, before a complete assessment of the required emergent properties (e.g., safety) can be made. The development of a safety management system (SMS) for UAM operations would act to enhance safety, and ensure operational integration into the greater NAS, however it is unclear if this will be required.

C. Personnel Considerations

A UAM applicant, other than one who only uses one pilot in their operation (see Part 61 for regulations concerning pilots [6]), is required to submit a training program for their pilot crewmembers and, if applicable, flight attendants. Training curricula must include curriculum segments, as applicable, for each crewmember position, such as (but not limited to): (1) Basic Indoctrination Training, (2) Emergency Training, and (3) Crew Resource Management (CRM) Training. Additionally, the UAM applicant must demonstrate how they will interact with the FAA's Safety Assurance System. The FAA's Safety Assurance System (SAS) policy and procedures provide aviation safety inspectors (ASI) with standardized protocols to evaluate certificate holder programs required by regulations to be approved or accepted. The SAS implements FAA policy by providing safety controls (i.e., regulations and their application), and requires Part 135 Certificate holders to demonstrate: (1) Element Design Assessments (EDAs), which are used to ensure a certificate holder's or applicant's operating systems are designed to comply with the intent of regulations and system safety, and (2) Performance Assessments (PA), which are used to determine if the certificate holder's or applicant's operating systems are producing the intended results, as well as to confirm the certificate holder's or applicant's risk assessment is occurring through their monitoring process and they are taking appropriate corrective action when needed. Again, it is unclear what level of risk will be acceptable under these requirements for UAM operations.

Aviation safety is supported by regulation and policy surrounding aircraft, operations (including crew requirements) and access to airspace. This leads to a holistic approach that allows for the composition of these three elements to yield




Fig. 1 Part 135 Operational Factors

a national airspace system that possesses an unparalleled safety record.

III. System Safety and Hazard Analysis Methods

Safety is an emergent property: it arises from the whole of the system—its components and their interactions. Since safety cannot be inspected, tested or audited into a system, system safety approaches center around building safety into systems, and not attempting to add it on to a completed design. A particular emphasis is placed on using analysis methodologies for system safety, as opposed to relying purely on past experiences or standards; this is invaluable in the design of new systems operating in novel contexts. System safety considers the more general concepts of *hazards*, rather than concentrating on failures, and emphasizes a qualitative approach.

A hazard is a system state or set of conditions that, together with a particular set of worst-case environmental conditions, will lead to an accident (loss) [7]. The primary focus of system safety is hazards: identifying, evaluating, eliminating, controlling and mitigating them through design, analysis and management procedures. It is important to identify hazards early in the system life cycle so corrective action can be taken to minimize the hazards and their effects. Where practical, this corrective action should take the form of intelligent design decisions. Decisions about the type of hazard analysis to perform turn on the purpose or purposes of the analysis: analyses can be used to convince a regulator, to formulate operational policies and procedures or to examine hazards in the design phase to develop control and mitigation strategies.

A hazard analysis provides a means for information to flow both to and from the analysis. Information flowing outwards from the hazard analysis enables designers to evaluate tradeoffs and eliminate or mitigate hazards. Information flowing inwards to the analysis concerns changes in design and enables the overall conceptual system model and safety approach be maintained. Goals of a hazard analysis center around system development, operational management and certification. To support each of these, we must demonstrate that all credible hazards have been adequately mitigated. Identifying and characterizing the hazards will require a complement of analysis techniques applied across the system development lifecycle. Note that while hazard analysis is an essential component of assuring the system safety, it is insufficient by itself: its insights must be acted upon.

In order to use any hazard analysis technique, it is important to define the objectives and scope of the analysis. This requires a clear definition of the system and its boundaries, the information to be used in the analysis, and the customers of the analysis. It is also important to choose the correct level of abstraction in order to define the system. Normally, the level of detail present is constrained by the design phase during which the analysis is performed. However, as hazard analysis is an iterative process, it may be necessary to add detail to characterize potential paths for hazard mitigation or perform a tradeoff analysis.

A. Hazard Analysis Throughout the Life Cycle

There are multiple types of hazard analysis used in aerospace practice [8, 9]. The preliminary hazard analysis (PHA) can be used to identify generalized system hazards and critical system functions in the earliest phases of the design life cycle [8]. The PHA can be used to facilitate design tradeoffs, since it may be started as early as the design exploration phase of the system. Note that this is an iterative process, with updates occurring as the system design matures. Since the PHA occurs early in the design process, the assessment of hazards is limited, and normally qualitative in nature.

A system hazard analysis (SHA) is employed as the design begins to mature [10]. Studies of hazards arising from subsystem interfaces, as well as operational hazards, are a focus of this technique. A SHA examines subsystem interfaces for compliance with safety criteria, degradation of safety from normal operations, and possible combinations of hazardous events (including simultaneous).

A subsystem hazard analysis (SSHA) [10, 11] is used to identify hazards associated with operating modes, critical inputs and outputs and hazards resulting from functional relationships between components and equipment in each subsystem. It is performed once there is sufficient detail concerning subsystem design, is iterated upon as the design changes, and then evaluated to assess the effect on system safety. Each subsystem has hazards identified with its normal operational and failure modes, including performance degradation and inadvertent functioning.

Note that while an SHA and SSHA are performed in much the same way, they have very different goals. The SSHA takes an *individual, component oriented* approach to whether or not normal operation or failures of that component affect system safety. The SHA considers whether normal or failure-oriented modes of the *components operating together* affect the system safety.

Finally, an operating and support hazard analysis (OSHA) identifies hazards, focusing on those created by human-machine interfaces, during all phases of system use and maintenance. All of the mentioned hazard analysis techniques have means to consider risk reduction strategies, or formulate mitigation policies. Note that there are many additional forms of hazard analysis; detailed surveys have been conducted in [10, 11].

IV. Hazard Analysis for UAM

We now examine the application of FHA and STPA to UAM operations, with the eventual goal of enabling certification (and regulatory approval) of these systems. We first perform an FHA to gain an appreciation for the types of hazards illuminated by this analysis. We then perform an STPA to illustrate the efficacy of this approach, specifically in relation to the ability to identify control paths to mitigate hazards. Multiple techniques in hazard analysis may be employed to address different aspects of the safety argument, as complementary techniques may be required to address safety requirements or constraints related to different portions of regulatory approval (e.g., type certification, operational approval, and pilot licensing).

A. Reference Scenario

Our analyses examine a specific scenario for a reference air metro shuttle mission using a hypothesized eVTOL aircraft. In this section, we enumerate our assumptions about the environment, aircraft, and operating scenario. Further details about the reference mission and aircraft are available in a separate technical report [2].

Environment features. We assume operation in an environment characterized by:

- A service area focused on a large modern city with features including an urban metropolitan landscape, high-rise buildings, a major airport hub, and active construction including large equipment and cranes
- Airspace shared with other manned and unmanned air traffic, VTOL and otherwise, and some form of Air Traffic Management (ATM)/UAS Traffic Management (UTM) in place
- A weather continuum spanning northern US winters and southern US summers, with representative temperature ranges, precipitation types and rates, and wind profiles

- Infrastructure sufficient to battery charging, dispatch, passenger management, and other associated needs, which can include take off and landing areas (TOLAs) ranging from cleared fields to purpose-built vertiports

The success of UAM will require the development of requisite energy and transportation infrastructure and the development of traffic management policy and systems. While these matters are as critical as development of the aircraft themselves, we do not address them further in this work.

Aircraft features. We assume an aircraft with the following features:

- Distributed electric propulsion (DEP) realized by 6–8 independently controllable motor/rotor pairs
- Vertical takeoff and landing capability combined with wing-borne forward flight, enabled by a combination of rotors and wings which may be fixed or tilting
- Piloted control with possible limited autonomous capability (i.e., intermediate-stage UAM capability maturity)
- Design/configuration for passenger transit, up to five persons of which one will be the pilot
- Commuter distance flight ranges within and across metropolitan areas, up to 50 nautical miles and altitude of approximately 10 000 ft
- Position reporting and communications appropriate to support ATM/UTM in a form to be determined
- Navigation appropriate to support precision route-following within the reduced separation minima likely to characterize the target urban environments
- Sensing appropriate to support encounter/conflict detection
- Transponding appropriate to support cooperative conflict management

Operational scenario. We will limit our analyses here to a scenario bounded by vehicle power engagement and disengagement and broken into the following eight flight phases:

- 1) *Lift to hover (takeoff).* The aircraft departs vertically from the pad at the origin TOLA and attains altitude appropriate to transition.
- 2) *Transition to forward flight.* The aircraft transitions from reliance on a vertical lift mechanism(s) to reliance on a fixed wing flight lift mechanism(s). This transition might involve a reconfiguration such as rotor or wing tilt.
- 3) *Climb to enroute.* The aircraft continues gaining altitude until reaching the target altitude for enroute flight.
- 4) *Enroute.* The aircraft flies its course at target altitude.
- 5) *Avoidance.* The aircraft maneuvers to deconflict with a detected collision hazard. Note conflict detection can rely on some degree of autonomy, depending on the concept. In this demonstration, the onboard pilot has the responsibility to see and avoid in accordance with 14 CFR §91.113 [12].
- 6) *Approach.* The aircraft decreases altitude until reaching the target altitude for transition.
- 7) *Transition to hover.* The aircraft transitions from reliance on a fixed wing flight lift mechanism(s) to reliance on a vertical lift mechanism(s). This transition might involve a reconfiguration such as rotor or wing tilt.
- 8) *Set down (land).* The aircraft completes a vertical descent to land on the pad at the destination TOLA.

We could have decomposed flight differently so long as we distinguished the flight phases that have distinct risk profiles [8]. We chose to mirror the conventional decomposition commonly used for conventional aircraft to make our analysis more accessible to readers familiar with such aircraft.

B. Functional Hazard Analysis (FHA)

Functional Hazard Analysis (FHA), sometimes called Functional Hazard Assessment, is a means of identifying hazards and an integral part of systems safety activities in civil aviation [8, 9]. FHA is a predictive technique that investigates the effects of failures of the system’s functions. The primary goal of performing a FHA is to identify hazardous failure conditions. As FHA is, in effect, structured brainstorming, its effectiveness is boosted by ensuring the representation of domain expertise, cross-functional thinking, and raw imagination in the team conducting it. In aviation, the FHA is applied first to the aircraft as a whole, that is, it is applied to aircraft functions (before they are assigned to aircraft subsystems).

In this section, we describe the FHA process as applied to the reference scenario and present an example of its output. A more complete writeup of the analysis is given in a separate technical report [2]. The process we use roughly conforms to that described in ARP4761 [8]. The process of performing the FHA comprised the following steps:

- 1) From an appropriate abstraction of the aircraft, create a list of functions.
- 2) Define the purpose and behavior of each function.

- 3) Consider the hypothetical failure modes of each function, and whether or not they are annunciated (e.g., unannunciated loss of function, annunciated incorrect operation of function [high/low] etc.).
 - 4) Determine the effects of the failure mode (e.g., what, if any, are the resultant hazards).
 - 5) Determine, record, and justify associated risk factors (e.g., hazard severity etc.).
- FHA results are usually recorded in a tabular format such as that used in Table 1.

1. Creating the List of Functions and Defining the Purpose and Behavior of Each

For this analysis, we started with an expected list of core functions taken from both ARP4761 and Hayhurst et al. [8, 13]. While there are many function lists we might have used, we started with familiar functions like *aviate*, *navigate*, and *communicate* so that our analysis would be accessible to people familiar with existing aviation hazard analysis practice. Starting from that base, we customized the function list to suit this unique aircraft and operating scenario. For example, we added the *transport* function. We then decomposed functions into subfunctions to achieve a level of abstraction suitable for early-life-cycle analysis: abstract enough to be tractable, with enough detail to yield useful results. The first levels of our function decomposition are given in Figure 2.




Fig. 2 First level function decomposition for the reference eVTOL

2. Considering the Hypothetical Failure Modes of Each Function and Determining Their Effects

After enumerating the system’s functions, we considered each of a subset of functions as it would be executed under each applicable phase of flight. We explored what it might mean for the function to fail in various ways, such as failure to perform the function, performing it when not called for, performing it incorrectly, etc., in both an annunciated and unannunciated matter. For each specific failure considered, we determined whether the state it would put the system into should be considered hazardous, and, if so, defined a hazard. Table 1 gives examples of such hazards, including “loss of propulsion” and “partial or complete loss of power.” (Note that if the vehicle loses all power, it will lose propulsion.)

3. Determining, Recording, and Justifying Associated Risk Factors

In determining the severity of each hazard, we used the standard classification documented in FAA AC 23.1309-1E [14]. This AC provides for the categories *catastrophic*, *hazardous*, *major*, *minor*, and *no safety effect*. In some cases, we have assigned familiar hazards a higher than typical severity level because uncertainty precludes us from assuming that the effects will be the same as for manned commercial air transports. It is usually much easier to decrease severity in the future, when better information justifies that decision, rather than to try to retroactively increase severity. Note that the severity of the hazard can be used as guide in order to help shape the level of effort and critical thinking spent on that topic.

4. Recording the Analysis Results

We recorded the results of our analysis in tables (e.g., Table 1). The *function* column captures the flattened function decomposition; in each row, it indicates which function was examined to produce the result documented by that row. The *flight phase* column repeats the flight phases relevant to the analysis of each function. The *hazard* column captures the descriptions of any undesirable system states that might arise as a result of the failure being analyzed in each row. In considering the combinations of function, flight phase, and manner of failure, zero to any number of hazards might be identified per function and captured in their own rows. The *operational consequence* column identifies the operational consequence(s) that might arise from the system being in each hazardous state during each relevant flight phase for each

Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
1 Aviate	...	...	...	...	...
1.1 Control flight path	...	...	...	...	...
...	...	...	...	...	...
1.1.4 Control propulsion	transition	Loss of propulsion	Impacts to aircraft, controllability, including loss of control	Catastrophic	Decoupling of propulsion and control; redundancy; glide and other off-nominal landing options
...	...	...	...	...	...
1.2 Control ground path	...	...	...	...	...
...	...	...	...	...	...
1.3 Control subsystems	...	...	...	...	...
1.3.1 Control powerplant	transition	Partial or complete loss of power	Impacts to aircraft controllability, including loss of control; impacts to propulsion, including loss of propulsion	Catastrophic	Power source redundancy
...	...	...	...	...	...

Table 1 Illustrative excerpt from the results of our FHA

function. The *severity* column gives the severity assigned to the hazard. The *mitigations* column is a place to record the analysts' insight into potential mitigations.

The use of a table is common and, save for the mitigations column, the structure of our table broadly aligns with those given in popular standards such as ARP4761 [8]. While mitigations will be planned later, it is worth capturing any insights analysts might have, as time and distance might make these insights less visible to designers working later in the life cycle. Analysts might record insights into causes in a separate column. Such non-traditional columns can be elided for presentation to stakeholders with particular formatting needs or preferences so long as such elision does not materially change the reasonable inferences allowable from the data.

5. Example from our FHA

In our analysis, the *aviate* function decomposes into: (1) control of flight path, (2) control of ground path, and (3) control of subsystems as shown in Figure 2. Control of flight path relies on control of several parameters including altitude, attitude, and propulsion, while control of subsystems includes powerplant management. For this example, Table 1 explores the causes and effects of a loss of propulsion.

To characterize implications of a failure of the control propulsion function (1.1.4), we consider effects of its malfunction, delay, or complete absence or unavailability. In this case, the consequences will vary significantly depending on both the design of the aircraft and the phase of flight during which the failure occurs. Generally, loss of propulsion will have implications for control of the vehicle, reducing controllability or removing it altogether. Under the best circumstances, loss of propulsion can be hazardous but recoverable. But during the transition phase, this hazard will be catastrophic: loss of control from any but a negligible altitude will likely result in fatalities.

Some failures might cascade or lead to combinations of distinct hazardous events. For example, the consequences of a loss of power will include the hazards of loss of control and loss of propulsion unless mitigations are provided.

Mitigations to loss of propulsion address several points in the causal chain and turn on differences among eVTOL designs. Some causes of propulsion loss include software, hardware, and data faults, and these should be addressed

via best practices in fault management. For example, the possibility of motor wear-out might be addressed through redundancy, which might be provided by a distributed electric propulsion design. The possibility of defects in battery controller software can be addressed through best practices in dependable software engineering. Other potential mitigations target the consequences of propulsion failure. These include measures to retain some control in a lost propulsion state and means of reducing impact forces in uncontrolled landings. For example, a wing-borne eVTOL might be designed to glide when propulsion is lost during forward flight. Ballistic parachutes and energy-absorbing materials and design might limit impact forces; however these mechanisms act to mitigate the accident (i.e., impact) and not the hazard (i.e., loss of propulsion) and are not traditionally accepted for certification credit.

C. Systems Theoretic Process Analysis (STPA)

The Systems Theoretic Process Analysis (STPA) is a hazard analysis technique that uses control theory as its foundation, rather than reliability theory [7, 15]. Because STPA is based on the System-Theoretic Accident Model and Processes (STAMP), it treats safety as a control problem rather than the result of failures. In STPA, the system is considered to be a collection of interacting control loops; hazard arise due to unsafe control actions. This allows for the consideration of the class of hazards that arise due to adverse interactions, in which no (component) failure may have occurred.

In this section, we apply STPA to the same eVTOL aircraft and UAM mission. We limited the scope of our analysis to safety (ignoring security concerns and loss of service), operation from power-up to power-down (ignoring maintenance), and taking air traffic management as given (rather than part of the system to be analyzed). The process of performing our STPA comprised the following steps:

- 1) Identify losses, hazards, and safety constraints associated with the system a priori.
- 2) Construct the control structure of the system.
- 3) Identify the unsafe control actions (e.g., not providing control action causes hazard, providing control action causes hazard, incorrect timing/order, stopped too soon/applied too long, etc.).
- 4) Identify causal factors and control flaws (e.g., what specific control action causes which particular hazard, analysis of control path and feedback loop etc.).

1. Identifying Losses

A careful scoping of the problem is required in order to conduct an STPA hazard analysis. This begins with definition of the issues of interest (e.g., whether security concerns will be included) and the losses to be avoided. In this analysis, we consider three forms of loss:

- L-1: Death of, or injury to, a human (aircrew, groundcrew, passenger, or third party).
- L-2: Loss of, or damage to, an aircraft (ownership or other).
- L-3: Loss of, or damage to, ground structures.

For the purpose of this analysis, we are ignoring three other potential kinds of losses: (L-4) loss of reputation (of OEM or carrier), (L-5) loss of service (of ownership or other), and (L-6) disclosure of confidential or protected information such as personally identifiable information pertaining to passengers.

2. Identifying Hazards

We then brainstormed nine hazards relating to these losses. As is usual in an STPA, each of these hazards identifies the system, the unsafe condition it can be in, and the related losses [15]. Where we felt it added clarity, we further refined these into sub-hazards. Our complete list of hazards includes:

- H-1: Ownership violates minimum separation standards in flight. [L-1, L-2]
- H-2: Ownership loses its structural integrity. [L-1, L-2, L-3]
 - H-2.1: Ownership is not in airworthy condition. [L-1, L-2, L-3]
 - H-2.2: Ownership has shed debris while in flight. [L-1, L-2, L-3]
- H-3: Ownership is not at a safe distance from terrain or an obstacle. [L-1, L-2, L-3]
 - H-3.1: Ownership does not maintain adequate altitude in flight before landing. [L-1, L-2, L-3]

- H-3.2: Ownship landing at site that is unsuitable for landing. [L-1, L-2, L-3]
- H-3.3: Ownship is too close to infrastructure, vehicles, or humans while taxiing. [L-1, L-2, L-3]
- H-3.4: Ownship is outside of an acceptable taxiing path. [L-1, L-2, L-3]
- H-4: Cargo or passenger is insufficiently restrained. [L-1, L-2, L-3]
- H-5: Ownship motion exceeds limits for occupant health and comfort. [L-1]
- H-6: Ownship parts that a human might contact have high kinetic energy or temperature. [L-1]
 - H-6.1: Ownship is powered up with passengers, aircrew, or other objects in proximity to moving parts (e.g., engines or rotors). [L-1, L-2, L-3]
- H-7: Ownship emissions of radiated energy, toxic or environmentally destructive chemicals, or radioactivity exceed safe limits. [L-1]
 - H-7.1: Ownship cabin environment is toxic. [L-1]
 - H-7.2: Ownship emissions exceed limits within enclosed spaces (e.g., hangars). [L-1]
 - H-7.3: Ownship emissions during service exceed standards. [L-1]
 - H-7.4: Ownship produces rotor / propeller wash in the presence of unprotected persons or unsecured objects. [L-1, L-2, L-3]
- H-8: Ownship is miscommunicating its status or intention to other air traffic or to an air or ground traffic controller. [L-1, L-2, L-3]
- H-9: Ownship is operating in an environment or flight regime (e.g., altitude, airspeed, or weather concerns) that is beyond its operating limits. [L-1, L-2, L-3]

3. Identifying Safety Constraints

From these hazards, we derive system safety constraints by inverting the condition and reformulating the text as a mandate. For example:

- SC-3: Ownship must maintain safe distance from terrain and obstacles. [H-3]
 - SC-3.1: Ownship must maintain adequate altitude in flight before landing. [H-3.1]
 - SC-3.2: Ownship must land only at suitable landing sites. [H-3.2]

4. Constructing the Control Structure of the System

A consideration of the vehicle, environment, operational procedures and human interaction paths is fundamental to being able to construct a model of the system's control structure, including its myriad feedback mechanisms. A systematic consideration of each control variable helps to define its impact on all potential hazards. In STPA, control structures are comprised of *controllers*, *control actions*, and *feedback* [15]. Controllers, which might be human beings, subsystems, or teams of these, are depicted as rectangles. Control actions that controllers make and feedback that they receive are depicted using arrows.

In this analysis, we defined the hierarchical control structure shown in Figure 3. As is usual in STPA, the highest-level controllers appear at the top of the diagram and the lowest-level controllers at the bottom. Text accompanying the figure elaborates on the controllers, control actions, feedback, and the controllers' responsibilities. For example:

Thrust setting control action given to the rotors / propellers / motors. While the exact nature of the aircraft's propulsion and lift mechanisms have yet to be decided, there will be rotors or propellers and their associated motors. In order to make the ownship aircraft fly per the flight control inputs (if being hand-flown) or the given trajectory (otherwise), the flight controller commands the motors and rotors or propellers to produce the thrust needed to accomplish that.

In defining this control structure, we chose a level of abstraction suitable for an analysis that was (a) performed early in the development life-cycle of the eVTOL and (b) focused on the aircraft rather than for managing such air traffic (e.g., onboard functions like collision avoidance are considered, while offboard functions like separation management are abstracted). For example, because system design is as yet immature, we abstracted all of the sensors, hardware, and




Fig. 3 STPA control diagram for the reference UAM aircraft and operation

software responsible for turning pilot or automatic navigation commands into flight surface control actions into a single flight controller. Because we are still exploring options related to the number and placement of rotors or propellers, whether these are fixed or can tilt from vertical to horizontal, whether the aircraft has fixed wings, etc., we abstracted those elements into the controllers labeled rotors / propellers / motors, control effectors, and transition mechanism in Figure 3. Later in the system lifecycle, after related design decisions have been made, we would revisit the analysis and refine the level of abstraction.

Control action	Given to	Kind	Effect
...	...	...	...
Thrust setting	Rotors / propellers / motors	Providing causes hazard	Commanding thrust when the ownship is on the ground and not cleared for takeoff [H-1, H-3, H-5, H-7.4].
(Same)	(Same)	(Same)	Commanding thrust while taxiing when doing so would lead the aircraft out of the cleared taxi path or into a person or object [H-3.3, H-3.4].
(Same)	(Same)	(Same)	Commanding thrust during takeoff, transition, climbout, flight, avoidance, descent, or landing when doing so would cause the ownship to venture unacceptably close to other air traffic [H-1].
(Same)	(Same)	(Same)	Commanding thrust during takeoff, transition, climbout, flight, avoidance, or descent when doing so would cause the ownship to exceed airspeed or altitude limits [H-2, H-9].
Thrust setting	Rotors / propellers / motors	Not providing causes hazard	Not commanding thrust during takeoff, transition, climbout, flight, avoidance, descent, or landing when not commanding thrust would cause impact with terrain, people or objects on the ground, or other air traffic [H-1, H-3.1, H-3.2].
Thrust setting	Rotors / propellers / motors	Incorrect timing / order	Providing too early or too late has generally the same effect as providing or not providing.
...	...	...	...
Selected configuration	Transition mechanism	Providing causes hazard	Commanding the transition mechanism to reconfigure the aircraft when flight condition limits for transition—e.g., on altitude, airspeed, and orientation—are not satisfied [H-2, H-5, H-9].
...	...	...	...

Table 2 Unsafe control actions

5. Identifying Unsafe Control Actions

We next identify unsafe control actions. For each control action given to each controller, we consider four kinds of potentially unsafe control actions: (1) cases where providing the control action causes the hazard; (2) cases where *not* providing the control action causes the hazard; (3) cases where the control action is provided at the wrong time or in the wrong order (vis-à-vis other actions); and (4) cases where the control action is stopped too soon or applied too long [15]. We document our results in a table adapted from that given in the *STPA Handbook*. As shown in Table 2, we adapt the standard table format somewhat: rather than having a column for each of the four kinds of unsafe control actions, we use *kind* and *effect* columns and multiple rows to make the table more readable on a letter-sized page. Rather than one row per combination of control action and recipient, we have four or more: at least one for each of the traditional kind columns, or more if there is more than one unsafe control action of a given kind.

6. Identifying Causal Factors and Control Flaws

For each identified unsafe control action, we then define one or more *loss scenarios*. Each loss scenario explores the potential causes of the unsafe control action. Potential causes include: (a) failures of a physical controller; (b) inadequate control algorithms (e.g., flawed algorithms or defective algorithm implementations); (c) unsafe control inputs (e.g., arising from the failure of another controller); and (d) inadequate process models (e.g., due to bad, missing, or misinterpreted feedback) [15]. For example:

UCA-3: Commanding a high thrust setting during takeoff, transition, climbout, flight, avoidance, or descent when doing so would cause the ownship to exceed airspeed or altitude limits [H-2, H-9]

- *Scenario 1 of UCA-3*: The ownship aircraft is flying on an approved path (in either vertical or horizontal flight mode) when either (a) a failure of the flight controller (e.g., due to hardware failure or a software

defect) or (b) receipt of a spurious flight control input or trajectory control action causes the flight controller to issue a high thrust setting to the rotors / propellers / motors [UCA-3]. As a result, the ownship climbs past its maximum operating altitude [H-9]. This might lead to loss of consciousness of those onboard the ownship aircraft.

- *Scenario 2 of UCA-3:* The ownship aircraft is flying on an approved path in horizontal flight mode when either (a) a failure of the flight controller (e.g., due to hardware failure or a software defect) or (b) receipt of a spurious flight control input or trajectory control action causes the flight controller to issue a high thrust setting to the rotors / propellers / motors [UCA-3]. As a result, the ownship accelerates past its maximum operating airspeed [H-2, H-9]. Thus, the ownship aircraft might suffer damage or become uncontrollable.

UCA-4: Not commanding a sufficient thrust setting during takeoff, transition, climbout, flight, avoidance, descent, or landing when lack of thrust would cause impact with terrain, people or objects on the ground, or other air traffic [H-1, H-3.1, H-3.2]

- *Scenario 1 of UCA-3:* The ownship aircraft is flying (either in vertical flight mode or horizontal flight mode) on an approved path that would take it into terrain, a vertiport, or a ground obstacle when flight control inputs or trajectory control actions dictate that it should slow, stop, or reverse its descent. Instead, either (a) a failure of the flight controller (e.g., due to hardware failure or a software defect) or (b) loss or corruption of the incoming control action causes the flight controller to command little or no thrust [UCA-3]. As a result, the ownship aircraft impacts terrain or lands too hard [H-2.1, H-3.1, H-3.2, H-5].
- *Scenario 2 of UCA-3:* The ownship aircraft is flying in horizontal flight mode on an approved path when it becomes necessary to pull up to avoid either an obstacle (e.g., a crane) or another aircraft. The ownship pilot issues appropriate flight control inputs or the flight planner issues an appropriate trajectory, but either (a) a failure of the flight controller (e.g., due to hardware failure or a software defect) or (b) loss or corruption of the incoming control action causes the flight controller to command little or no thrust [UCA-3]. As a result, the ownship aircraft collides with the obstacle or other aircraft [H-1, H-3.1].

7. Recording the Analysis Results

The results of an STPA include its (a) list of losses, (b) hazard definitions, (c) defined constraints, (d) statements of responsibilities (formalized versions of the required controller behavior we describe above), (e) unsafe control actions, (f) scenarios (with and without UCAs), and (g) defined controller constraints [15]. These outputs are collectively useful to inform safety aspects of system design and operation, e.g., evaluating design concepts for safety impact, defining safety requirements, and planning for the monitoring of leading indicators of an accident. For example, after noting that commanding a transition from vertical to horizontal fixed wing flight might put the aircraft into a hazardous state if it is done outside of airspeed and altitude limits, a designer might consider safety mechanisms to prevent the transition mechanism from acting on such commands.

V. Discussion and Evaluation of Hazard Analysis Techniques

In this section, we present our observations about hazard analysis, FHA, and STPA. We then discuss the utility of FHA and STPA, including areas where these techniques support certification in ways that complement each other.

A. Observations on Hazard Analysis

In this section, we present our observations about hazard analysis generally.

1. There Is Freedom in How the Analyst Models the System and World

Both FHA and STPA require the analyst to construct and analyze an abstract model of the system and its operating environment. For example, in FHA, the analyst must determine what the ‘functions’ to be analyzed are and how far to decompose these. In STPA, the analyst must define a control structure, determining what parts of the system and world to model and what interactions should be modeled. For example, in the control model in Figure 3, we modeled air traffic control as a single, abstract controller. Likewise, we modeled all of the computers, software, interconnections, etc. involved in flight control as a single abstract controller. We modeled all relevant aerodynamic principles as the air data and ownship position and motion ‘controllers.’ If we were attempting to learn more about how details of UAM air traffic management, computer equipment reliability, or specific aspects of aerodynamic performance might contribute to

losses, we could have refined this structure. While detail might be lacking during an analysis done so early in the life cycle, we might even have attempted analysis with details from multiple proposed designs thereby enabling a design space exploration with respect to system safety. The choice to abstract some details in a safety model and elaborate others affects both the tractability of the analysis and its ability to illuminate issues of interest.

2. Knowledge and Skepticism of Familiar Hazards are Helpful

Knowledge of hazards in similar systems is helpful to the analyst performing either FHA or STPA: such knowledge both reduces the likelihood of missing a key insight and helps to frame results in a way that is easily understood by other people familiar with those other systems. But familiar hazards might behave differently in new environments and under new operational assumptions. FHA severity assignments for common high-level aviation hazards are often based on assumptions about the performance and behavior of aircraft, operational models, traffic management, and traditional human-machine role allocations. Even if those assumptions have been borne out in historical practice, they might be undermined by some key difference between a given system and its nearest analogues. In our FHA analysis, such difference caused us to reclassify as hazardous some hazards that are traditionally classified as major [2].

3. Individual Hazards Can Come from Multiple Sources

Identical hazards can arise from multiple sources, and individual hazards can give rise to multiple consequences. In both FHA and STPA, we often found multiple ways to end up in a given undesirable state. This illustrates why several different mitigations will sometimes be needed to address a single hazard: some mitigations will work better for some sources of a given hazard than others. Further, since multiple consequences can arise, multiple profiles for the hazard are possible, and each must be characterized appropriately to allow for adequate management.

4. New Technologies and Environments Can Demand New Ways of Mitigating Familiar Hazards

People familiar with other systems might also be familiar with how those systems' hazards are mitigated. But, just as familiar hazards might differ substantially in new circumstances, familiar mitigations might prove less effective. For example, autorotation is a required mitigation for loss of power in all single-engine helicopters and some other VTOLs. But small rotors do not support autorotation; designers choosing small rotors will need to choose alternative mitigations.

5. Crashworthiness is Not Enough

An aircraft designer might include features like ballistic parachutes and energy-absorbing materials to protect occupants in the case of a crash. But while these features might reduce the likelihood and severity of harm to passengers, they do little to affect risk to other persons, including persons on the ground. Had we split injury to aircraft occupants from injury to people on the ground, rather than combining these into the loss L-1, our STPA's hazard list would show several hazards traceable to both kinds of loss. Mitigations that preclude or reduce the likelihood of entering both of these hazards will still be required.

6. Safety Analyses are Living Artifacts

Our safety analyses are based on knowledge available at the time they were conducted. Functions may be created or changed; new hazards might be identified; and increasing detail about the system's design and construction will become available as system development proceeds. It is necessary to perform analysis early in the lifecycle to gain insights that will lead designers to produce a safe system. But it is also necessary to update the safety assessments as needed.

7. Hazard Analyses are Produced Iteratively

The examples given here and in handbooks, textbooks, and guides might seem to show a one-pass analysis process. But real analysis is not perfect: as our analyses proceeded, we found ourselves reworking earlier parts of the analysis to clarify our meaning or capture previously-overlooked insight. For example, while enumerating the unsafe control actions described in section IV.C.5, we sometimes found ourselves trying to work out which hazard(s) best captured the result of a control action we were considering. When this happened, we clarified or expanded the hazard list to better present our evolved understanding of the system's hazards to the reader.

B. Distinguishing FHA and STPA

In this section, we present observations about how FHA and STPA differ.

1. *The Phase of Operation Is Explicitly Considered in FHA*

Both FHA and STPA can take the phase of operation (e.g., maintenance, taxiing, takeoff, avoidance maneuver, etc.) into account when identifying hazards. But where this is explicit in FHA, it is implicit in STPA. The *STPA Handbook* recognizes the importance of analyzing different phases of operation where these are material, exhorting analysts to encode ‘context’ such as “during normal takeoff” in their descriptions of unsafe control actions [15]. But the STPA process itself does not force the analyst to consider each phase.

2. *Some Important Distinctions Mainly Affect Hazard Severity*

Our FHA revealed that loss of propulsion, control, or both is especially problematic during transition from vertical to forward flight or vice-versa. The distinction it reveals is one of hazard exposure and mitigation efficacy: transition is a flight phase where the envelope is especially small and low altitude limits opportunities to recover before a crash. But STPA is focused on worst-case outcomes, and the worst case for unsafe control of thrust or control effectors is a loss event no matter the phase of flight. As a result, STPA did not highlight this distinction as well as FHA. An analyst can reframe UCAs with context descriptions that highlight such distinctions (although it can be difficult to do that without referencing likelihoods we can only roughly estimate). But we felt as though we were reaching for a way to frame insight that we’d brought to the STPA from the FHA that preceded it rather than being alerted to them by the analysis mechanism.

3. *Scenario Exploration*

While both the FHA and STPA techniques require the analyst to think critically about means by which these hazards manifest, we found that the STPA technique allowed for the exploration of specific scenarios in depth, due to the iterative refinement inherent in considering potentially unsafe actions, as well as causal factors and control flows, in the control loop formalism. While the columns in the FHA format allowed for a progression of events to be captured, the nature of the scenario exploration in STPA allows for an explicit consideration of timing aspects in tandem with event driven exploration. The scenario based exploration explained and contextualized these safety insights in a clear and explicit fashion, as opposed to the more implicit in nature of the FHA analysis.

4. *Hazard Identification*

The FHA and STPA ended up surfacing similar hazards, which is not surprising given that many functions are characterized as aircraft control functions. However, the level of abstraction in which certain aspects of these hazards are treated may differ between the techniques. For instance, the FHA considers a partial or complete loss of power as a hazard, which is a precursor to the loss of propulsion hazard (given the design assumptions). The ability to consider a partial loss of power is important, as the mitigation for a partial loss may require more stringent methods (e.g., pilot training as well as power system redundancy). The STPA explores this mechanism more fully, by examining what are the implications of an incorrect power input being applied, not only in magnitude, but also in time (e.g., too early or too late). This may have profound effects on the strategies used to mitigate this hazard, as a partial loss of power applied before a critical event (e.g., transition to forward flight) may act to negate the severity imposed by that event (e.g., transition is aborted by pilot). These varying levels of abstraction can be used to surface insights that are applicable not only to vehicle certification, but also to operational and pilot training requirements.

5. *Expertise of User*

Both the FHA and STPA techniques require the user to have some knowledge of system safety analysis, as well as a background in the system being analyzed. Inexpert users of either technique tend to delve into increasingly improbable means of manifesting a given hazard, though the mechanism by which this may occur is different. In the case of the FHA, the inexpert user may tend to be sidetracked into causal event chains, where multiple causes may need to occur in order to evince the hazard. For the STPA, the inexpert user may tend to get sidetracked into increasingly more elaborate control flow scenarios, thereby designing the path of control actions that would lead towards an exact accident. Both of these potential pitfalls can be mitigated with repeated use of the technique, as well as circumvented using insight from

more experienced users.

VI. Conclusion and Future Work

In order to make a case to a regulator that a UAM operation with an eVTOL aircraft possesses the level of assurance necessary in order to access the national airspace system, an adherence to the fundamental principles of system safety must be exhibited. The use of sound, repeatable engineering and system safety practices will greatly enhance the ability of an applicant to demonstrate that the system meets its requirements (e.g., certification basis) and mitigates identified hazards to the desired risk level. It is necessary to consider both the vehicle and operation in order to perform a meaningful safety assessment. In this work, we examined two techniques for performing hazard assessment, the Functional Hazard Assessment (FHA) and the Systems Theoretic Process Analysis (STPA), in the context of a novel UAM operation with an eVTOL vehicle. We evaluated these techniques and determined that they could be used in a complementary fashion. The FHA illuminated multiple types of hazards associated with a specific function of the system and allowed for further refinement with respect functional decomposition. The STPA highlighted control paths and specific scenarios that led to identified hazardous states, which could also be used to craft mitigation strategies.

We found that the use of multiple hazard analysis techniques increased the comprehensiveness and depth of our understanding of the hazards to be mitigated. Using both techniques enriched our understanding of all of the means by which a hazard may arise and exposed multiple potential avenues for that hazard's mitigation. The FHA focused on mitigations primarily related to the vehicle and its design, which can be used as a foundation for the generation of the vehicle type certification basis. The STPA allowed for the consideration of operational and crew oriented mitigations, which could inform and impact requirements related to crew training and operational approval. Thus, the use of complementary techniques can help to address different parts of the regulatory process.

The FHA tabular format draws attention to insights related to hazard exposure and severity, which may allow for the identification of design points where the mitigation space may be constrained. The STPA scenarios enable an exploration of detail that allows for the contextualization of safety insight in a clear narrative fashion. For a novel operation with an innovative vehicle design, it is crucial that HA be as thorough as possible; the use of predetermined hazard lists and stock mitigations may mislead stakeholders into confidence in a design that is less safe than they realize. The use of multiple complementary techniques aids in deep exploration, and highlights the issue that a certain level of human judgement and expertise is necessary in the execution of a hazard analysis. No matter the technique that is used, it may be unwise to attempt to substitute a purely mechanized means for this process, as key insights and assumptions in the system design may be missed.

Additionally, the use of both techniques allowed for the consideration of the same hazard through multiple levels of abstraction. This enabled new ways of reasoning about that hazard and exposed different aspects of that hazard (e.g., explicit consideration during flight phase versus general consideration of timing related to specific events). Using two HA techniques may allow for multiple mitigation strategies to be evaluated in the context of a single hazard and enable the generation of vehicle type certification, operational and crew training requirements, which can potentially be traded off during different mission configurations. This is a topic for future work. Further future work centers around investigating other hazard analysis techniques that might be used to further augment the two discussed in this paper. The development of a common vocabulary for hazard assessment techniques (such as a common definition for, or interpretation of the word hazard) would be helpful and is being pursued. Finally, an exploration of the utility of this work to both novel OEMs and regulators would help to validate the approach and is being formulated.

Acknowledgments

The authors would like to thank Mr. Jeffrey Maddalon and Dr. Paul Miner, for all of their invaluable help in conducting the functional hazard analysis, as well as their guidance in scoping the problem. Additionally, the authors would like to thank Mr. Frank McCormick, who provided excellent insight related to the identification of hazards, as well as the regulatory means by which many of these hazards are dealt with under the current certification structure.

References

- [1] 14 CFR Part 21, "Certification Procedures for Products and Articles," *Code of Federal Regulations (CFR)*, Last accessed Jun. 2019. URL <https://www.ecfr.gov/cgi-bin/text-id?node=14%3A1.0.1.3.9>.
- [2] Wasson, K., Neogi, N., Graydon, M., Maddalon, J., Miner, P., and Frank McCormick, G. F., "Functional Hazard Assessment for

an eVTOL Aircraft Supporting Urban Air Mobility (UAM) Applications,” Technical memorandum (draft), National Aeronautics and Space Administration, 2020.

- [3] 14 CFR Part 23, “Airworthiness Standards: Normal Category Airplanes,” *Code of Federal Regulations (CFR)*, Last accessed Dec. 2019. URL <https://www.ecfr.gov/cgi-bin/text-idx?node=pt14.1.23>.
- [4] 14 CFR Part 27, “Airworthiness Standards: Normal Category Rotorcraft,” *Code of Federal Regulations (CFR)*, Last accessed Dec. 2019. URL <https://www.ecfr.gov/cgi-bin/text-idx?node=pt14.1.27>.
- [5] 14 CFR Part 135, “Operating Requirements: Commuter and On Demand Operations and Rules Governing Persons On Board such Aircraft,” *Code of Federal Regulations (CFR)*, Last accessed Jun. 2019. URL <https://www.ecfr.gov/cgi-bin/text-idx?node=14:3.0.1.1.11>.
- [6] 14 CFR Part 61, “Certification: Pilots, Flight Instructors, and Ground Instructors,” *Code of Federal Regulations (CFR)*, Last accessed Jun. 2019. URL <https://www.ecfr.gov/cgi-bin/retrieveECFR?n=14y2.0.1.1.2>.
- [7] Leveson, N. G., *Engineering a Safer World: Systems Thinking Applied to Safety*, The MIT Press, Cambridge, MA, USA, 2011.
- [8] ARP4761, *Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment*, Society of Automotive Engineers, 1996.
- [9] ARP4754A, *Guidelines for Development of Civil Aircraft and Systems*, SAE, 2010.
- [10] Ericson, C. A., II, *Hazard Analysis Techniques for System Safety*, 2nd ed., John Wiley & Sons, Hoboken, NJ, USA, 2016.
- [11] Leveson, N. G., *Safeware: System Safety and Computers*, Addison Wesley, 1995.
- [12] 14 CFR Part 91, “Certification: General Operating and Flight Rules,” *Code of Federal Regulations (CFR)*, Last accessed Nov. 2019. URL <https://www.ecfr.gov/cgi-bin/text-idx?node=pt14.2.91>.
- [13] Hayhurst, K. J., Maddalon, J. M., Miner, P. S., Szatkowski, G. N., Ulrey, M. L., DeWalt, M. P., and Spitzer, C. R., “Preliminary Considerations for Classifying Hazards of Unmanned Aircraft Systems,” Technical Memorandum NASA/TM-007-214539, National Aeronautics and Space Administration, Hampton, VA, USA, Feb. 2007. URL <http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20070008225.pdf>.
- [14] FAA, “System Safety Analysis and Assessment for Part 23 Airplanes,” Advisory Circular AC 23.1309-1E, Federal Aviation Administration, Washington, DC, USA, Nov. 2011. URL http://www.faa.gov/documentLibrary/media/Advisory_Circular/AC%2023.1309-1E.pdf.
- [15] Leveson, N. G., and Thomas, J. P., “STPA Handbook,” Electronic document, Mar. 2018. URL http://psas.scripts.mit.edu/home/get_file.php?name=STPA_handbook.pdf.