NASA/TM–20210024234

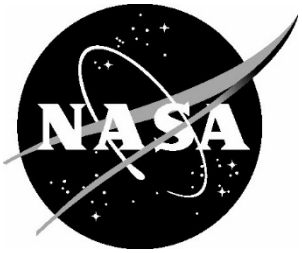

Functional Hazard Assessment for the eVTOL Aircraft Supporting Urban Air Mobility (UAM) Applications: Exploratory Demonstrations

Kim Wasson
Federated Safety, LLC, Crozet, Virginia

Natasha Neogi, Mallory Graydon, Jeffrey Maddalon, andPaul Miner
NASA Langley Research Center, Hampton, Virginia

G. Frank McCormick
Certification Services, Inc., Eastsound, Washingston

April 2022

NASA STI Program . . . in Profile

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

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

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

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

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

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

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

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

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

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

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

- E-mail your question to help@sti.nasa.gov

- Fax your question to the NASA STI Information Desk at 443-757-5803

- Phone the NASA STI Information Desk at 443-757-5802

- Write to:
  STI Information Desk
  NASA Center for AeroSpace Information
  7115 Standard Drive
  Hanover, MD 21076-1320

NASA/TM–20210024234

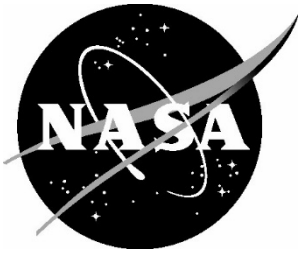

Functional Hazard Assessment for the eVTOL Aircraft Supporting Urban Air Mobility (UAM) Applications: Exploratory Demonstrations

Kim Wasson
Federated Safety, LLC, Crozet, Virginia

Natasha Neogi, Mallory Graydon, Jeffrey Maddalon, and Paul Miner
NASA Langley Research Center, Hampton, Virginia

G. Frank McCormick
Certification Services, Inc., Eastsound, Washingston

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

April 2022

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

February 2022

Nomenclature

AC	advisory circular
AI	artificial intelligence
ASTM	ASTM International
ATC	air traffic control
ATM/UTM	air traffic management / unmanned traffic management
DAL	design assurance level
DEP	distributed electric propulsion
DEP	distributed electric propulsion
eVTOL	electric vertical takeoff and landing
FAA	Federal Aviation Administration
FHA	function hazard assessment
FMEA	failure modes and effects analysis
FTA	fault tree analysis
GA	general aviation
HW	
IFR	instrument flight rules
M&S	modeling and simulation
ML	machine learning
MTA	mission task analysis
NMAC	near midair collision
OEM	original equipment manufacturer
SAE	SAE International
SDO	standards development organization
SOTIF	Safety of the Intended Functionality
SSA	system safety assessment
SSA	System Safety Assessment
SVO	Simplified Vehicle Operations
SW	
TOLA	take-off and landing area
UAM	Urban Air Mobility
UAS	unmanned aircraft systems
UID	universal identifier
V&V	verification and validation
VFR	visual flight rules

VTOL	vertical takeoff and landing

Abstract

The active development community surrounding electric vertical takeoff and landing (eVTOL) aircraft has demonstrated potential to bring new technological capabilities to market, encouraging visions of widespread and diverse Urban Air Mobility (UAM) applications. New capabilities always come with safety considerations, some familiar and others less so. OEMs who intend to obtain FAA type certification for their eVTOL designs must plan for and execute sufficient safety engineering processes to demonstrate that credible hazards associated with these eVTOL designs are adequately mitigated. This work provides an orientation for eVTOL stakeholders, especially new entrants and those in hybrid roles, to the safety and regulatory context for assuring eVTOL aircraft for UAM applications. We further present selections of functional hazard assessment (FHA) performed on a reference eVTOL concept, with process and decision narration. The exercise allows exploration of several safety considerations specific to eVTOL systems and demonstrates the FHA process together with negotiation of some of its options and variations.

1 Introduction

Technological advances in power, propulsion, autonomy, and noise management in air vehicles capable of vertical takeoff and landing (VTOL) have encouraged a diversity of new ideas for how such systems can be used. The impending availability of VTOL aircraft running on hybrid or all-electric power, with distributed propulsion configurations, and with the eventual likelihood of autonomous navigation and control, opens avenues for a number of commercial, medical, and humanitarian applications, especially in dense urban environments. Urban Air Mobility (UAM) refers to the set of capabilities that allow safe and efficient transport of passengers and cargo by air in these urban environments. These capabilities comprise vehicle technologies, airspace management, and infrastructural support, together with their regulatory and social contexts.

Potential UAM applications include medical transport (patients, staff, organs), courier services, search and rescue, and many others. The application receiving the most industry attention at present, due to promising business cases as well as the backing and visibility generated by companies like Uber, is on-demand personal air mobility, with air taxi as a focal example. The very active personal air mobility development community is represented by at least 200 urban VTOL aircraft concepts in design and demonstrator stages, by as many companies, covering the continuum of size and experience [25]. Boeing, Bell, and Airbus are all here, as are many startups with just a handful of staff, and everything in between. Each is bringing novel technical solutions to engineering problems, many of which will find homes in the multiple niches that will likely open in these markets.

1.1 Dependability

However, while these novel technical solutions are necessary to obtain access to this airspace and operate in service to these new applications, they are not sufficient. All functional capability in any system comes with a set of dependability objectives that must be met in order for the functional capability to be relevant. If a system can perform a desirable function, but only with unacceptable levels of safety or security or reliability, for example, then the system is not deployable.

FAA certification of the aircraft design is a main gate to prevent insufficiently dependable aircraft from reaching deployment. As stated in [9], “The legal purpose of avionics certification is to document a regulatory judgment that a device meets all applicable regulatory requirements and

can be manufactured properly." Type design and additional certifications sit on foundations of criteria and assurance designed to maintain acceptable levels of risk to the flying public, as well as to third parties incidentally affected by these operations.

For this work we focus on safety. A main component of achieving certification is demonstration that all of the safety requirements of the design are known, understood, and satisfied. This is conventionally done via a rigorous system safety assessment process (SSA) that is well-documented in the aerospace community and outlined later in this report.

Demonstrating that the safety requirements for these systems are known, understood, and satisfied is a significant challenge due to the novelty and complexity of the systems and their target environments. A number of ongoing community discussions center on safety issues that have new sources, new severities, or otherwise new profiles relative to familiar instances. Some will be explored in this report. Rigorous SSA provides a methodical approach to establishing the safety demonstration, as well as a common conceptual model to support OEM engagement with the regulator.

1.2 Regulatory Challenges

The UAM community faces a number of challenges to commercial electric vertical takeoff and landing (eVTOL) aircraft deployment that set it apart from aviation development practice of the last few decades.

First, many OEMs and individual practitioners in this community are new entrants, as this application space draws on expertise from a variety of sources, and its compelling complex problems attract outside interest and talent. Many of these stakeholders also fill hybrid roles, performing several functions within a busy organization, and learning one or more of these roles on the job. OEMs actively developing aircraft concepts in this space range from large, mature organizations with established safety practice, to small, new manufacturers with promising technical expertise but minimal exposure to aviation safety practice or to the aerospace regulatory environment. Those with less exposure will have learning curves to navigate as they determine ways to identify and satisfy safety requirements consistent with certification expectations.

Second, the regulatory environment itself is also in transition. In the last few years several changes have occurred that provide updated guidance and specific allowances relevant to the increasing diversity of type designs and development methods. While these changes generally increase flexibility for OEMs in showing compliance, they also create another learning curve, both for OEMs and the regulators enforcing them. These challenges will be further discussed in section 2. That said, while some surrounding regulation is in transition, most is stable. Importantly, the safety assessment practice that supports it is also stable and will remain a foundation to safety demonstration regardless of changes to criteria and standards.

Earl Lawrence is Executive Director of the Aircraft Certification Service at FAA. At a recent symposium on the challenges to eVTOL aircraft certification, he offered the following:

> *We really don't care about your business case or how much money you are going to make. We look to make sure all the safety risks are mitigated. When you come to us start with safety. That's our focus, that's what all our regulations are built around. If you always come in with safety, you will keep moving forward.*

One of the best ways to bolster a safety case is to liberally adopt industry standards, Lawrence counseled, citing the "huge benefit" derived from the use of industry consensus standards [17].

That is, despite the ongoing transitions in regulatory frameworks, coupled with the diversity and pace of technical development in this area, the best strategy for OEMs seeking type certification is (1) to focus on safety, and (2) to address safety via established best practice.

In this work, we will overview the regulatory backdrop to eVTOL type certification and demonstrate execution of the process that anchors accepted safety assessment practice in aerospace, functional hazard assessment (FHA). We will perform example FHA activities on elements of a reference system concept consistent with features and operations of many eVTOL aircraft in development and narrate this execution. The examples chosen will allow exploration of some safety considerations specific to eVTOL aircraft, and the FHA narration will provide execution guidance not available in applicable standards.

1.3 Scope

This work defines a reference eVTOL concept and a number of operational scenarios consistent with systems in development in the community and their intended applications. Considerations of the certification basis for this aircraft are also addressed. These entities are defined in enough detail to enable the execution of high-level FHA activities for exploration and demonstration purposes. This work focuses on safety hazards of the aircraft design, with some attention to operation. Other significant considerations for UAM, such as infrastructure and traffic management, are only addressed here as they arise in connection with aircraft design hazards under discussion.

1.4 Purpose

The purpose of this work is to orient UAM stakeholders, especially new entrants and those in hybrid roles, to the safety and regulatory context and responsibilities for eVTOL aircraft safety assurance. Toward this purpose, we conduct the following activities:

1. We introduce FHA and execute selections for a reference eVTOL aircraft, providing demonstration and guidance in the application of this SSA anchor.
2. We examine several high-level eVTOL aircraft hazards in some detail to spotlight safety challenges specific to these target systems in their target environments.

We do this from a system-safety perspective using a hypothetical reference system, to provide a concrete set of examples supporting community orientation and discussion.

2 Regulatory Environment

It is possible that "a new avionics device might be brilliantly conceived and flawlessly designed yet ineligible for certification" [9]. This quote highlights the notion that building a system and demonstrating its readiness for certification are two different things. Subtler is the consideration that a brilliant conception and flawless build are only brilliant and flawless with reference to a set of criteria. An OEM can build an eVTOL aircraft that brilliantly and flawlessly accomplishes a particular function, but certification will require the following:

1. The OEM can back up that achievement with appropriate substantiation accessible to a third party.

2. The risks of *undesired* system behavior have also been identified and mitigated with appropriate and accessible substantiation.

That is, demonstrated functional achievement is necessary but not sufficient for certification.

There is a popular notion in the UAM community that the technology to enable desired capabilities is already here and that the immaturity of regulations is creating a bottleneck to progress [12]. The situation is quite a bit more complicated than this. There have been great advances in distributed electric propulsion (DEP), energy density of batteries, autonomous control, and other eVTOL-enabling technologies. However,

1. vehicle design is only one part of a complex problem; infrastructure such as vertiports and energy grid considerations and air traffic management in new and complex airspaces need also to be developed, each requiring significant systems design and development in addition to regulatory involvement; and
2. advances in functionality require complementary advances in assuring the safety of that new functionality, itself to be deployed in complex new environments. These safety criteria, standards, and methods of compliance are also *developing*. This supports, though is distinct from, regulation. Regulation will look to and assess these advances in order to stabilize regulatory expectations, but the safety analyses and their design ramifications are technical problems largely in immature states of solution.

Further, regardless of level of uncertainty in the development context, "[o]n any new project, it is unwise to presume that all regulatory requirements are known" [9]. That is, any new system development, even for something familiar and well-understood, contains surprises. For these eVTOL systems, there will be many, against an also-moving backdrop. Currently available certification pathways for eVTOL aircraft are *developing*, are not standardized, and are best described on a case-by-case basis. Starting points are described in section 2.1

The implication is that OEMs should plan to work with their regulators, early and often, to establish the certification basis and compliance expectations for their type design and all of its component systems: "applicants are encouraged to start a dialogue with certification authorities as early in the process as possible to reach a common understanding of means of achieving compliance […] This is especially important as new technology is applied to avionics and as new personnel enter the field" [24]. As a part of this, the OEM must demonstrate to the regulator early that it is capable and prepared to properly identify and realize the safety objectives for the system. The OEM must work with the regulator to agree on the plan, and then the OEM must track and fulfill the plan (with periodic regulator oversight). As a prerequisite to eVTOL aircraft certification, the OEM must ultimately substantiate the assertion, to FAA's satisfaction, that the system is adequately safe.

2.1 Candidate Certification Pathways

At the time of this writing, there is much active discussion on possible certification pathways for commercial eVTOL aircraft in the U.S., and no eVTOL design has yet set a precedent [23]. This significantly complicates the OEMs' task of drafting a certification basis to propose to the FAA and to guide system development.

As aircraft, eVTOL vehicles will be subject to the 14 CFR Part 21 baseline regulations for type and airworthiness certification and other approvals for aircraft [2]. This is where the certainty ends. In accordance with the size and weight of these aircraft, Part 23 airworthiness standards for small

airplanes have relevance to the discussion, but do not address all factors. Because many of these systems have rotors in addition to wings or do not have wings, the Part 27 airworthiness standards for rotorcraft also have relevance and also do not address all factors. Other parts similarly have potential application.

Further, Part 21 separates regular (Part 21.17(a)) from special class (Part 21.17(b)) aircraft. A regular aircraft classified as a particular type must meet all of the requirements for the type, plus any special conditions derived through the negotiation between FAA and the OEM. Special conditions are line-item additions to the certification basis to account for a system's divergence from available regulation. A special class aircraft, on the other hand, is one that does not conform satisfactorily to any single type definition, and its certification basis is constructed by parts through custom assignment of applicable criteria drawn from relevant types, for example, from Part 23 for small airplanes and from Part 27 for rotorcraft. A certification basis for a special class aircraft is sometimes viewed as entirely a complex special condition.

While it might seem straightforward to assume that a certification basis in accordance with Part 21.17(b) (special class) makes the most sense because it can cover in a customized way every new configuration, the issue is more complicated. The industry and the public benefit from the deliberate development of well-thought-out regulations with the big picture in mind. If every new aircraft concept gets a completely idiosyncratic certification basis, the means of normalizing expectations for performance and substantiation suffer, and therefore so do the build and evaluation processes. This serves nobody.

Rather, the community is currently in the thick of discussions about what it would mean for an eVTOL aircraft to be classified as

- a Part 23 normal category airplane under Part 21.17(a), with or without special conditions (the latter is difficult to imagine here);
- a Part 27 normal category rotorcraft, with or without special conditions (again, the latter is difficult to imagine here);
- a special class aircraft under Part 21.17(b); or
- any of a number of other less prominent but still possible options.

In addition, the community is discussing

- which criteria and methods of compliance from Parts 23, 27, and others apply in what ways to the configurations in development;
- what kinds of special conditions would make sense for these purposes as extensions to Parts 23 or 27;
- which new considerations are not yet covered at all and require entirely new regulation;
- how many pathways we need, and how few can reasonably fulfill the objectives; and
- how all of these issues and their potential resolutions in the U.S. interact with the same discussions being had in Europe and elsewhere.

The certification pathway(s) for eVTOL aircraft will not stabilize immediately. There appears to be consensus across the FAA and industry right now that the first several eVTOL designs to be type certified will be case-by-case undertakings from which we will learn and gain experience to help make the necessary decisions to normalize the process. Against this backdrop, several factors impact the tasks of identifying the safety requirements for these systems and demonstrating their satisfaction. Some of these are elaborated below.

2.1.1 On the Part 23 Rewrite

14 CFR Part 23 recently underwent a significant rewrite that completed in 2017 [3]. The rewrite shifts compliance requirements for airworthiness standards for small airplanes toward performance-based assessment, focusing more on desired properties to be shown and less on prescribed process. This came partially in response to both the explosion of features and technologies in development and to the multiplication of development methods. Performance-based assessment relieves both the OEM and the regulator from attempting to manage all of that diversity prescriptively, and instead puts the focus on how the resulting system behaves. For aircraft subject to Part 23 criteria, to include most eVTOL aircraft in some way, this update provides some demonstration relief to OEMs with regard to process requirements, and this relief can be significant where certain conventionally-required processes have become incongruent to current needs over time. Makers must still justify claims of performance and safety, but methods of justification are more flexible [14].

With this flexibility comes a new burden: OEMs must now evaluate and decide among options for showing compliance, and new options do not arrive fully formed. With the Part 23 rewrite, the FAA is *allowing* industry some space to develop new means of substantiating performance claims. The FAA is also *expecting* industry to step up with innovation here, proposing and maturing criteria and means of compliance that are sound and convincing and that can begin to populate new certification toolboxes. Several standards development organizations (SDOs) are working on foundations to establish options for acceptable means and methods to meet performance-based requirements (e.g., ASTM, SAE).[1] Safety is a primary topic in these discussions, from substantiation of safety requirements to the assurance levels warranted by different development choices, to which methods are eligible for certification credit toward achieving target assurance levels. For eVTOL aircraft, these discussions further include rationale for the applicability or inapplicability of performance and safety criteria, whether new or sourced from existing standards, as well as new means of substantiation.

In all versions of the most likely certification pathways in consideration for eVTOL aircraft, 21 CFR Part 23 is expected to be directly applicable in whole or in part, including the system safety compliance requirements. While the rewrite offers the potential for flexibility in substantiation data, Part 23 compliance still rests on a foundation of practice for system development, software development, hardware development, and safety. FAA advisory circulars recognize particular standards for each of these areas as acceptable means of compliance. For safety, AC 23.1309.1E [22] provides guidance and also recognizes SAE ARP4761, *Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment* [10].

In other words, a safety requirement provoked by mitigation of a hazard might now be verifiable through potentially new means. However, it is still strongly recommended that identification of the hazard and design of the mitigation in accordance with target assurance levels are accomplished via conventional system safety assessment practice.

[1] More generally there is ongoing work supporting the effective realization of performance-based assessment at various levels, for example, the FAA's Overarching Properties project [1] and increasing activity in structured argument and assurance cases.

2.1.2 Target Safety Levels

In determining whether hazards have been adequately mitigated, the achieved levels of mitigation must be compared to target safety levels defined in accordance with the certification basis established for the aircraft. These target safety levels then map to required design assurance levels (DALs) for system functions and components, that are satisfied through development evidence. At the time of this writing, the target safety levels in discussion for eVTOL aircraft are evolving and derive from a risk-based classification framework in development by FAA. This framework takes into account considerations like the kinetic energy of the aircraft configuration, the passenger capacity, and other contributors to risk estimation. Because these aircraft have significantly lower mass and passenger capacity than conventional airliners, it is possible we will see some certification bases in the U.S. with target safety levels slightly relaxed from those of commercial airliners. This would affect the required substantiation data, *but not the structure of the safety assessment process*.

In parallel, EASA has recently released a document entitled *Special Condition for Small-Category Vertical Take-Off and Landing (VTOL) Aircraft* (SC-VTOL-01) that specifies the current European regulatory requirements for eVTOL designs [19]. In the environments characteristic of air taxi operations, SC-VTOL-01 expects a target level of safety for these aircraft in line with that of commercial airliners. Reasons for and impacts of this difference are beyond the scope of this work, other than to note that this channel also relies foundationally on a rigorous SSA process.

2.1.3 Hazard Assessment Basis

Section 2.1.1 introduced the Part 23 compliance foundations, including specific recognition by advisory circulars of acceptable means of compliance. The industry standards SAE ARP4754A [11] for systems, SAE ARP4761 [10] for safety, RTCA DO-178C [18] for software, and RTCA DO-254 [5] for hardware work together with specific interfaces and shared information to create a combined process that, implemented with fidelity, establishes a strong foundation toward certifiability. While the FAA states in each associated advisory circular that these standards do not represent the only acceptable means, the FAA expectation is that an OEM that intends to deviate from this framework will coordinate this intention with the FAA early in the certification program and obtain acceptance of the alternate plan.

The hazard assessment demonstration documented in the remaining sections is presented within the context of these development and safety process standards, recognizing that ultimate performance requirements and means of compliance (in fact most of the certification basis) will remain in development for the foreseeable future, modulo case-by case evaluations for the first applicants. Early applicants will contribute to precedent. Submissions that successfully demonstrate hazard identification and mitigation for the novel technologies in question and the novel methods supporting their development will help shape follow-on expectations.

3 Reference System and Scenarios

The aircraft FHA begins from an aircraft function list. The function list for this demonstration FHA is derived from the reference system and scenarios described in this section. The FHA reference scenario presented here extends an intra-metro air shuttle mission scenario that has arisen

in community discussion to describe an intermediate state of UAM maturity.[2] This intermediate maturity scenario provides piloted air metro services between high-density ground and air transport hubs; in the mature state, these missions are envisioned to be autonomous. In the air shuttle mission scenario, air vehicles of multiple types serve users who require transport between these hubs and who embark and disembark at designated UAMports and/or take-off and landing areas (TOLAs). Air shuttles fly predefined routes, allowing significant use of strategic conflict management, for example flight planning and procedural separation. Tactical conflict management is additionally employed to address remaining encounters. Both VFR and IFR operations are supported. Air shuttles may be yielded right of way due to limited maneuverability.

The air shuttle scenario further considers how users will locate UAM ports, acquire tickets, progress through security, and be assessed along with their luggage in support of weight and balance calculations for the flight. As we are focused here on aircraft system hazards, these considerations are out of scope for current purposes. The FHA reference scenario specified below is confined to the period from when the vehicle initiates transit under its own power to when it comes to rest, ceasing this control. Pre- and post-flight commercial and related activities are therefore also out of scope for this FHA demonstration activity.

Likewise, since the model FHA is focused on the aircraft system and functions, other systems and functions such as those of the Air Traffic Management system and vertiport infrastructure are beyond current scope, though dependencies between the aircraft and these systems will be referenced later in the discussion.

Performance requirements of the aircraft, levied to enable integration with these and other systems, remain in scope insofar as they become part of the aircraft specification.

Figure 1 illustrates the aircraft system boundary relative to these other systems with which it interacts. The shaded area is the focus for this analysis, in the context of these other systems.

[2] While these informal scenario descriptions frequently anchor stakeholder discussions, they do not yet appear to be publicly documented. As NASA's UAM Grand Challenge program proceeds, scenarios will be refined, and, eventually, made public. For our purposes, the current informal inputs are sufficient.

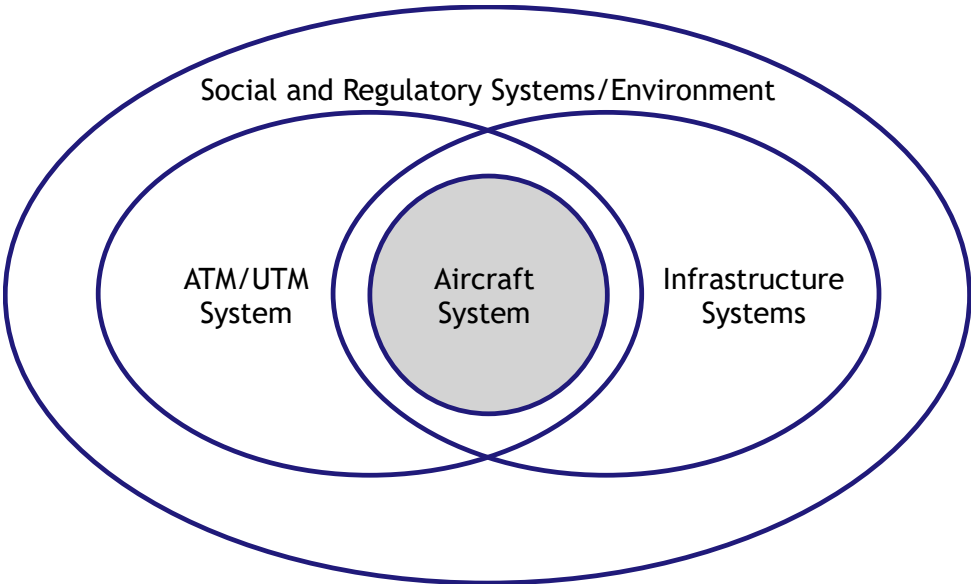

Figure 1: Interacting System Boundaries

For purposes of FHA, note also that the aircraft system design includes all users. Passengers and any pilot(s) or operator(s) are first-person parties to various aircraft functions and therefore also factor into the specifications and mitigations of associated hazards.

The reference scenario defined below further specifies some aspects of a model environment, aircraft, and user interaction (both pilots and passengers) with the vehicle system, in order to allow the identification of functions upon which to conduct the demonstration hazard assessment. We present one of many possible design concepts, and the design space represented currently within the air taxi community is quite large. Some of the prominent design dimensions include the following:

- **Control.** Piloted and autonomous concepts are both represented, as well as combinations thereof.
- **Propulsion.** Fully-electric and hybrid-electric-combustion concepts are both represented.
- **Lift mechanism.** Wing-borne and vertical-lift concepts are both represented, as well as combinations that transition between the two.
- **Configuration.** Fixed and tilt-rotor/-wing concepts are both represented.

Each of these dimensions provokes specific questions in the SSA and hazard assessment. Aircraft concepts that transition from vertical to forward flight and back are particularly interesting in this regard and create new challenges for hazard mitigation that we will discuss later in this work.

We have specified a concept within this space, consistent with the intermediate maturity stage scenario described above. Where some alternative design choices are likely to be salient, we will narrate additional safety considerations within the demonstration. The remainder of section 3 outlines the system concept upon which we will base our assessment.

3.1 Environment Features

The operating environment for this reference scenario mirrors that of the eVTOL concepts in active development for UAM. In particular, we consider operations characterized by the following:

- A service area focused on a large modern city, with features including an urban metropolitan landscape, a variety of building heights including high-rises, a major airport hub, and active road and/or building construction including large equipment and cranes.
- Shared airspace with other manned and unmanned air traffic, VTOL and otherwise, and some standard of ATM/UTM in place.[3]
- A weather continuum spanning northern U.S. winters and southern U.S. summers, with representative temperature ranges, precipitation rates, and wind profiles.
- Infrastructure sufficient to address battery charging, dispatch, passenger management, and other associated needs, which can include TOLAs of various levels of sophistication from cleared fields to purpose-built vertiports.[4]

The development of requisite energy and transportation infrastructure, together with the development of traffic management policy and systems addressing deployment of these systems in these airspaces, can be seen as two additional pillars supporting the goals of UAM. These are as critical to the mission as the aircraft themselves, and though they are not the focus of this exercise, some important dependencies will be noted in the discussion.

3.2 Aircraft Features

The aircraft for this reference scenario represents features common to many of the design concepts in active development in the eVTOL space. Since there are over 200 active concepts as of this writing, and many dimensions in which they vary, no single specification can represent them all. The design used for this activity is based on feature choices that are both well-represented in the active design space and provide opportunities for discussion of hazards arising from some of the differences between these and traditional aircraft.

The aircraft for this reference scenario is characterized by the following:

- 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, representing the intermediate stage of UAM capability maturity.[5]

[3] The development of ATM/UTM support for heterogenous aircraft operations in dense urban environments is, like other components of UAM, developing. While the specification of the eventual systems is as yet unknown, the requirements and hazards of these traffic management solutions are in active analysis in parallel with development of the aircraft systems themselves.

[4] The development of urban infrastructure sufficient to support UAM operations is, like other components of UAM, developing. While the specifications for power provision, dispatch, passenger management, and other features of TOLAs and vertiports are as yet unknown, the requirements and hazards of these entities are in active analysis in parallel with development of the aircraft systems themselves.

[5] Though the mature state for most of the concepts in active development includes significant autonomous capability, technology development and public acceptance drivers indicate an intermediate state where these systems will be controlled by onboard pilots or operators.

- Design/configuration for passenger transit, up to five persons, one of which will be the pilot[6]; this implies, for example, requirements for HVAC support, seating, passenger restraints, ingress/egress, etc.
- Commuter distance flight ranges within and across metropolitan areas, up to 50 nm and FL100; this implies, for example, no requirement for cabin pressurization, but an analog to conventional fuel reserve requirements is in active discussion in the community.
- 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.

3.3 Operational Scenario

We next define a hypothetical operational scenario bounded by vehicle power engagement and disengagement broken into eight flight phases.[7] In this scenario, all phases are executed under pilot control, since this is a piloted scenario representing an intermediate-maturity UAM capability.

There are options for decomposition of the scenario into flight phases. This decomposition was chosen in order to

1. maintain consistency with the takeoff/climb/enroute/approach/land decomposition commonly used for conventional aircraft; and
2. distinguish additional flight phases of potential salience resulting from relative differences in associated risk profiles.

This latter objective derives from direction in SAE ARP4761: "The FHA should identify the failure conditions for each phase of flight when the failure effects and classifications vary from one flight phase to another" [10]. This results in a decomposition with an additional phase change for each of the transitions in lift mechanism for this aircraft, since the risk profile changes at these transitions. In addition, we have for this demonstration identified "avoidance" as an explicit phase, for purposes of forcing specific analytic attention there during the FHA activity, on the premise that the denser airspace environment and novel encounter scenarios suggest potential changes to the risk profile during avoidance activities.[8] The eight flight phases are as follows:

1. **Lift to Hover (Takeoff).** The aircraft departs vertically from the pad at the origin TOLA and attains altitude appropriate to transition.

[6] Concepts in active development tend toward two to five passengers for this air metro scenario, resulting from a tradeoff sweet spot in balancing market drivers with size, weight, and power needs and the associated noise implications. Specifically, any market for larger capacity aircraft is dampened by the undesired technical side effects of the increased aircraft size.

[7] In the baseline scenario specification, these event boundaries refer to engagement and disengagement as *intended* behavior. Unintended power engagement that occurs prior to block-off is beyond the scope of the scenario. Unintended power disengagement that occurs prior to block-on is within scope and is addressed within the FHA as resulting from a functional failure, creating a hazard of loss of power.

[8] Because FHA is iterative, such provisional decisions may be reversed upon availability of analytic results that favor an alternate organization. Given this flexibility, it is recommended to choose to provide for (vs. not provide for) such potentially relevant considerations and reevaluate as results warrant.

2. **Transition to Forward Flight.** The aircraft transitions from reliance on a vertical lift mechanism(s) to reliance on a forward flight lift mechanism(s). Note: this transition might or might not involve a reconfiguration, such as rotor tilt, depending on the concept design.
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 still has the responsibility to see and avoid in accordance with 14 CFR 91.113 [4].
6. **Approach.** The aircraft decreases altitude until reaching the target altitude for transition.
7. **Transition to Hover.** The aircraft transitions from reliance on a forward flight lift mechanism(s) to reliance on a vertical lift mechanism(s). Note: this transition might or might not involve a reconfiguration, such as rotor tilt, depending on the concept design.
8. **Set Down (Land).** The aircraft descends through remaining altitude vertically and lands on the pad at the destination TOLA.

3.4 High-level Safety Considerations

This reference aircraft in this reference environment shares high-level safety considerations common to all passenger-carrying aircraft. In particular, throughout execution of intended functionality, associated risks to persons (both on- and off-board), property, and the environment must all be adequately managed. Traditionally, these needs raise questions regarding the ability to reliably fly the aircraft, and the ability to reliably avoid collisions.

These universal considerations will factor in our analysis from the beginning and be joined by considerations specific to this aircraft and how it can achieve (or fail to achieve) these basic performance requirements.

How we identify, break down, and prepare to address these and other safety needs, is the focus of the next section.

4 FHA Demonstration

FHA is a process to focus performer attention systematically on system and operational conditions that could prove hazardous, and to define associated safety requirements for the system. It is, in effect, structured brainstorming, its effectiveness boosted by ensuring the representation of domain expertise, cross-functional thinking, and raw imagination in the team conducting it.

FHA is the origin process for system safety assessment: it identifies hazards and sets safety objectives for the system. The later, comprehensive system safety assessment (SSA) will then show that the implemented system meets the safety objectives established by the FHA [10]. This makes FHA critical to the eventual safety argument for the delivered system; it is impossible to say the hazards have been adequately mitigated unless the hazards are first identified and assessed.

This section documents how we executed FHA for selected considerations of our aircraft design with reference to its intended operational environment. With this exercise we intend to demonstrate the process for new entrants, with attention to approach and decision considerations not provided

in the standards. In addition, we present examples that illuminate specific features and challenges of the UAM application domain.

Type certification for any aircraft requires demonstration that hazards for the type design in question have been adequately assessed and mitigated. Since there is significant novelty here in both the designs and the environment, the hazard list will be unique, and the mitigations will include new and different means, relative to more conventional applications. The FHA establishes the profiles of the hazards to be mitigated and the bases for their mitigation.

Since FHA and SSA are critical to the acceptability of the delivered system, the development plan for any type design should provide for dedicated resources and personnel for these activities. Someone within the organization must own the safety process, manage its execution, ensure its validation, and sign it off. The new entrant OEM is advised to resource this process appropriately and early.

4.1 FHA Process

FHA discharges its obligations to SSA through an iterative, stepwise process designed to uncover cause-and-effect relationships and prioritize them for response. The analyst performing an FHA considers functions at the most appropriate level and "identifies failure conditions and the associated classifications while considering both loss of functions and malfunctions" [10]. FHA provides for the identification of functions, the independent consideration of each function and its possible failure conditions, identification of potential consequences, and the assignment of associated severity. Functional failures are considered at both the aircraft level and at the level of supporting systems, individually and sometimes in combination [10]. For this effort, we focus on individual failures at the aircraft level.

4.1.1 Standard Practice

FHA is used across safety-critical applications from multiple industries, as well as in both the commercial and defense sectors. Two safety standards in active use that include FHA are the civil standard SAE ARP4761, *Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment* [10], and the defense standard MIL-STD-882E, the *Department of Defense Standard Practice for System Safety* [21]. The FAA's AC 23.1309-1E offers some elaboration on how to comply to SAE ARP4761 and includes some high-level example hazards presented in a representative table format [22]. With some variations, the FHA guidance provided in these and other documents is foundationally consistent. While referencing consensus standards for concreteness, we will focus our presentation on the foundations common across recognized best practice. That is, it is ultimately most important to surface and sort the relevant cause-and-effect relationships sufficiently to enable appropriate management. With this in mind, we will also point out some process variations and options and how they might be negotiated.

Consistent with established practice, we first identify and decompose the aircraft-level functions. This is accomplished via the following process:

1. **Begin with an expected core function list.** This is a high-level function list that enumerates and organizes the core functions of any aircraft. Some generic lists are available in the literature, standards, and other regulatory documents [10, 16, 22]. Others are proprietary to aircraft manufacturers. For this demonstration, we began with the

function list provided in SAE ARP4761 [10] together with the additional organization and analysis provided by Hayhurst et al. [16].

2. **Customize this core function list using inputs specific to the system in development.** These inputs can usually be sourced from early conceptual documents for the system, and should include, but are not limited to
    a. aircraft objectives and customer requirements, for example, passenger-carrying capacity, range, etc.;
    b. initial design decisions, for example, take-off geometry, engine type (combustion vs. electric), etc.; and
    c. results of targeted discovery activities such as mission task analysis (MTA), intended to support early requirements and function identification.

We will elaborate on function identification and decomposition in section 4.1.2.

Then, given the function list as customized to the aircraft in development, we conduct the systematic functional failure analysis. This begins as a top-down process outlined by the following steps:

1. For each function at the target level of abstraction, **consider how the function can fail**. The rigor of this step can be modulated with further prompts, for example, there can be a malfunction, a delay of function, absent function, and potentially other variations. The effects in each of these scenarios can be considered.
2. For each functional failure, **define the resulting failure condition(s), or hazardous state(s)**, if any. Terminology here varies across the community. In either case, we are referring to a state (the hazard) that results from an event (the functional failure).
3. For each identified hazard, **consider possible consequences**. These are the undesired possible outcomes should the hazard be activated. These consequences might be actual losses or might be additional hazardous states.
4. For identified consequences, **assign severity classifications**. This is usually done in accordance with an applicable convention associated with the regulatory pathway.
5. For those hazards posing risk above a threshold (as defined elsewhere in the program's system safety process), **propose mitigations**. Mitigations are design changes that manage the risk exposure presented by a hazard, and they become new safety requirements as they are selected and stabilized.

While the process begins top-down, it will surface additional information that will generate feedback loops and iteration. For example, some hazards will cause consequences that will imply the existence of additional functions not yet captured in the top-down phase. We provide examples of this and other process flow considerations in later sections.

4.1.1.1 On Limitations of Standards

We noted that standards and other available guidance are largely foundationally consistent but do show some variation. In addition, while they describe an assessment process, these sources generally do not provide realistic examples of conceptual options and decisions that must be negotiated by the FHA team during execution of the analysis.

A number of such considerations appear in actual application, and we address some of them explicitly in this work, such as the following:

- **Need for a fault model.** Standards are good at prescribing high-level steps, and less good at conveying a mental model from which the steps can be derived. We anchor the hazard assessment in a system states-and-events model that supports identification and relation of the failures, hazards, and consequences in causal chains and nets.
- **Terminology.** Sources vary in their uses of terms for similar and overlapping concepts. We source and define our terms, and address variation and selection where applicable.
- **Table organization.** We discuss some of the options for what classes of data to track in the table and how we made our decisions.
- **Type hygiene.** Instances of data captured into the table should follow rules of mutual consistency and applicability in order to support valid analysis. We introduce the notion of type hygiene and how we apply it here.

The first item in particular provides guidance when faced with a common hazard assessment pitfall. As analysis proceeds, it will sometimes be the case that a failure appears to produce many hazards. This might be true, or it might require additional examination in order to identify further dependencies among the identified hazards. Some will result only indirectly from the initial failure, realized upon a second or later failure of a different function. Part of the value of using a fault model with FHA is in providing the framework to separate and trace these layered states and events and their dependencies. This separation and tracing enables more direct treatment of individual hazard profiles and surfacing of additional information about the associated functions. We will return to this discussion later in the execution narration.

4.1.2 Narration for this Application

This section describes the workflow by which we executed the FHA together with the conceptual drivers at each step. We begin by establishing a core function list and then conducting a functional decomposition in order to have a starting point for the identification of hazards. We then assess those hazards for potential consequences and associated severity, and then speculate mitigations in accordance with severity and the hazard sources. Once we have this initial top-down draft, we speculate storylines oriented on particular hazards for discussion. Through these storylines, we surface and collect additional considerations, including functions missed in the initial decomposition. Insights and observations from the exercise are elaborated as they occur and generalized to takeaways later in section 4.4.

The narration of this FHA exercise includes details of the actual process as executed by this team, including decision bases, false starts, and dead ends. This is intentional; FHA is a structured but informal process and much of the *noise* of its execution is unaddressed by the standards. Its value is in what it surfaces, but it does not surface anything until the analysis team finds traction with a conceptual organization. We explored several paths into and within the analysis in order to find and maintain traction. Some of these will be discussed later in this report. One takeaway is that there is not a single right way to perform the various conceptual sorting tasks inherent to FHA. The goal is to force out relevant considerations and follow them through for implications.

Notably, some OEMs and other development organizations maintain proprietary baseline function and hazard lists derived and refined over previous developments. To get to this point, an organization must have sufficient experience, both with successful applications as well as with

challenges that led to lessons, to warrant codification in a more permanent reference artifact. These artifacts provide a more advanced starting point for an FHA but must still be customized to the new application. Further, this customization becomes significant if the new application is very different from previous ones.

The current exercise does not begin from such a reference artifact. Rather, this exercise begins from first principles, applying the FHA method to an entirely new system and starting from a blank sheet of paper. We do this to demonstrate for new entrants the process of getting started, and to overview the kinds of considerations that are raised, the decisions that must be made, and their implications. This kind of amplification is not provided in the standards, and where sometimes exercised in academic system safety curricula, it is there generally divorced from the particular challenges of novel systems in evolving regulatory environments. This work demonstrates FHA from first principles on a timely eVTOL concept in the current context of many new entrants and significant regulatory uncertainty.

The remainder of this section describes the application of FHA, foundationally consistent with SAE ARP4761, to the reference eVTOL aircraft described in section 3.

For the orientation purposes of this demonstration, we focus on the aircraft-level hazard assessment. Hazard assessment for supporting systems would proceed once requirements and design decisions for that level had reached suitable definition.

We further address single functional failures, as well as examples of cascading failures and their relation to hazards. Examination of multiple simultaneous failures is beyond the scope of this work, though it must also be completed for production efforts in accordance with standard safety practice.

4.1.2.1 Function List

The aviation community (including the FAA) generally recognizes *Aviate, Navigate,* and *Communicate* as the top-level primary functions of any aircraft, with *Aviate* as the highest priority [7]. This is independent of aircraft type, mission, and other variables such as level of autonomy. For example, whether pilot- or software-controlled, the first priority is to fly the plane, and to do nothing else unless the aircraft is in controlled flight. While under positive control, the other top-level functions are to navigate the aircraft to its destination, and to communicate with bodies such as Air Traffic Control (ATC) in accordance with regulation such that airspace management can be maintained. Most aircraft function lists will include *Aviate*, *Navigate*, and *Communicate* at the top.

To this set, we add *Transport* as this is the primary motivating function for our reference eVTOL aircraft. The eVTOL aircraft described in section 3 exists to transport people and/or cargo commuter distances within an urban environment.

Other top-level functions may be considered in accordance with intended missions and capabilities. For example, one documented application of FHA to unmanned aircraft systems (UAS) includes *Mitigate*, on the premise that the unmanned systems under consideration will perform many monitoring, deconfliction, and other activities traditionally under the responsibility of an onboard pilot [16]. In order to channel focused analytic attention to these activities, *Mitigate*

was identified as a top-level function for that effort.[9] As autonomy enters and takes over more responsibility for control of eVTOL aircraft for UAM, safety engineers should also take into account such considerations. For this effort, and accordance with the presence of an onboard pilot in our reference scenario, we have classed this set of airspace integration functions under *Aviate*.

Another potential top-level function we discarded, but which might be appropriate for other analyses, is a cross-cutting *Manage Systems* function, on the premise that there are many activities that can be batched and addressed for their commonalities this way. In the current case, we found the decomposition more straightforward if we included instances of such potential classes where they arose rather than batching them. For example, there are sub-functions involving various sorts of monitoring under both *Aviate* and *Navigate*.

Thus, for the current effort, we stabilized to the following four top-level functions:

- *Aviate*
- *Navigate*
- *Communicate*
- *Transport*

It is of note that the sources and decisions described above are typical but only make up a small subset of the inputs and negotiation that resulted in our core function list. More accurately, six team members from a range of technical and regulatory backgrounds considered over twenty source documents in order to propose candidate functions and negotiated over several meetings to arrive at these four. While it is reasonable to assume an experienced team working on a more conventional aircraft might take a more direct path, it is strongly recommended that programs involving new entrants and/or novel configurations start from the beginning. It is through this exploration and negotiation that the foundation of the mental model is established.

4.1.2.2 Function Decomposition

We then decomposed these four functions through at least two levels (some further depending on the function and saliency of abstraction). In general, we trialed decomposition choices and levels of abstraction using reference sources as applicable, in combination with several areas of domain expertise represented in our team and negotiated consensus. We proceeded far enough to arrive at a minimum of two levels of decomposition for several interesting paths while ensuring a selection of examples across the main functions for narration purposes.

Note that there are many possible decompositions, and many of them are useful.[10] The idea here is that we are not looking for a single objectively-correct decomposition because it does not exist; there is not one ring to rule them all.[11] Rather, we are looking for decomposition strategies that allow us to identify and force out features of relevance. Consider a simple example of colored shapes. Since the entities have both shape and color, they can be grouped in at least these two dimensions; we can sort the squares separately from the circles or the red ones separately from the blue ones. However, what we're doing next with these colored shapes determines which sort is

[9] The *Mitigate* function identified for that effort is not to be confused with the SSA task of mitigating hazards. For that effort, the *Mitigate* function was subject to its own hazards that still required mitigation.
[10] This is not to be confused with George Box's famous estimation regarding models ("All models are wrong, but some are useful"), though it is of a similar flavor in terms of identifying the value of a tool or method.
[11] Unlike in, say, Middle Earth.

more appropriate. If the objective is to put them through some filter in which the holes are round, the squares won't go, regardless of color. We want to sort by shape here, even though color might be useful in a different context.

With function decomposition, we have many possible options as well as levels of abstraction to contend with. It is quite easy to inject too much foreknowledge and/or domain-specific knowledge into a round of high-level decomposition simply because it is available to capture, but it is not always, or even usually, the right path. A useful rule of thumb is to ask what is the very next thing, and only the next thing, that would be helpful to know here, as opposed to asking what are *all* of the things we know (though many of these can and should be captured to a buffer during the process for consideration downstream).

Further, because FHA is iterative, the decomposition organization will evolve as choices and their relative value surface. At this point, the focus is a top-down breakout of high-level functions, sourced from sensible inputs, oriented toward what will happen next with this information, and refined by available expertise. It will change. The object is to establish a starting point.

Returning to the colored shapes example, it might further be the case that both the shape and color dimensions are relevant, but the dimensions have an ordered priority. For example, we need to find the squares first, but of the squares, our stakeholder then wants the blue ones. In a decomposition model, this means we can break out by shape, and then further break out by color. Any set to be decomposed (colored shapes, system functions) is a multi-dimensional space, and many of these dimensions could be simultaneously relevant for a given need, but order can matter. So, in decomposing system functions in preparation for hazard assessment, if we decide one dimension was not the most important at a given level, it still might arise as relevant at a lower level. If we pushed it on our side stack for later reconsideration, all the better.

Our Level 1 decomposition is presented graphically in figure 2.

For ease of management, our decomposition was stored in its own table during generation, separate from the FHA table. The high-level function decomposition table is provided in the appendix to this document.

These functions and sub-functions, as well as the Level 2 and any additional decompositions we executed, were then entered into the FHA table and assigned numeric identifiers consistent with their decomposition paths and levels of abstraction. Selections will be presented in section 4.2 later in this report.

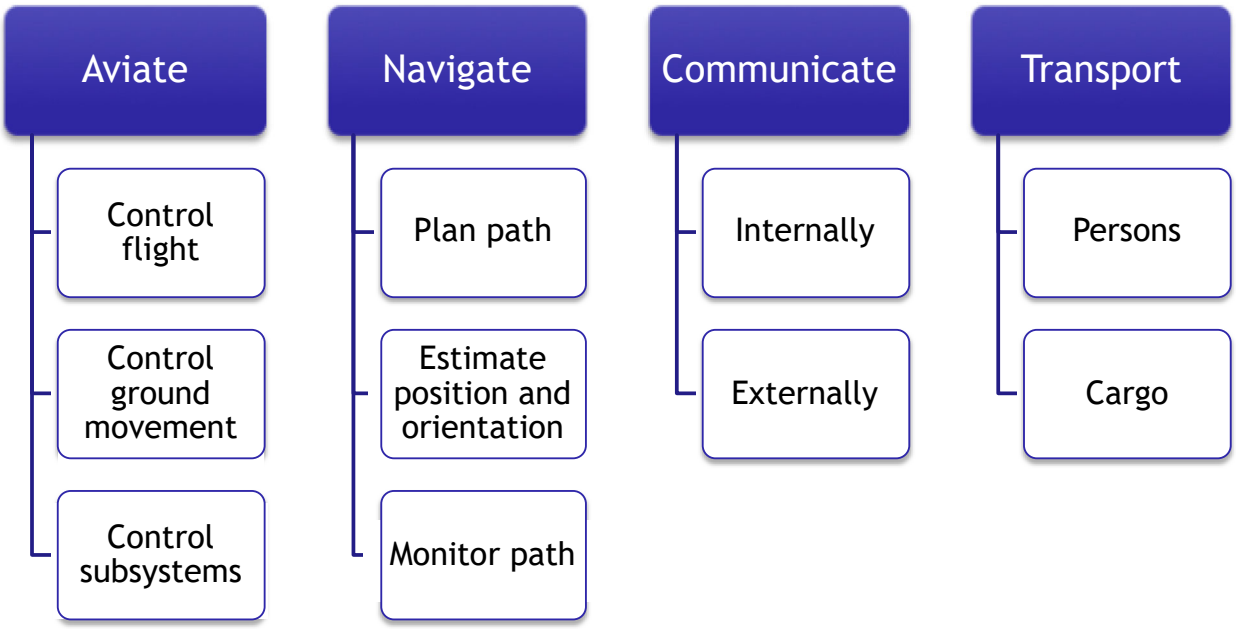

Figure 2: First Level Function Decomposition for Reference eVTOL aircraft

4.1.2.3 Hazard Identification and Analysis

From the draft decomposition, we initiated the hazard identification. At the outset, we ran into several options and challenges that had to be reconciled in order to proceed. These included coping with diversity and information limitations in available guidance, such as the following:

1. **Terminology.** Disagreements such as failure condition vs. hazard, overloaded terms, and conflation of concepts make it difficult for practitioners to apply an approach consistently within and across applications, and therefore difficult also to aggregate results.
2. **Type hygiene.** Example tables in some reference sources were populated inconsistently with regard to the type of data expected or allowed by column; this lack of rigor complicates and potentially invalidates assessment.
3. **Format variations.** FHA table structures (data containers) differed somewhat across sources; this makes it difficult for practitioners to determine exactly which data is most important to collect.

Further, the mental model underlying the prescription contained in the standards and other guidance was not obvious. To orient the demonstration to a set of foundational objectives, we cast the intuitive FHA goals as shown in figure 3. This gave us a reference anchor to aid in decision-making about elements of process, to support resolution to some of the above challenges.

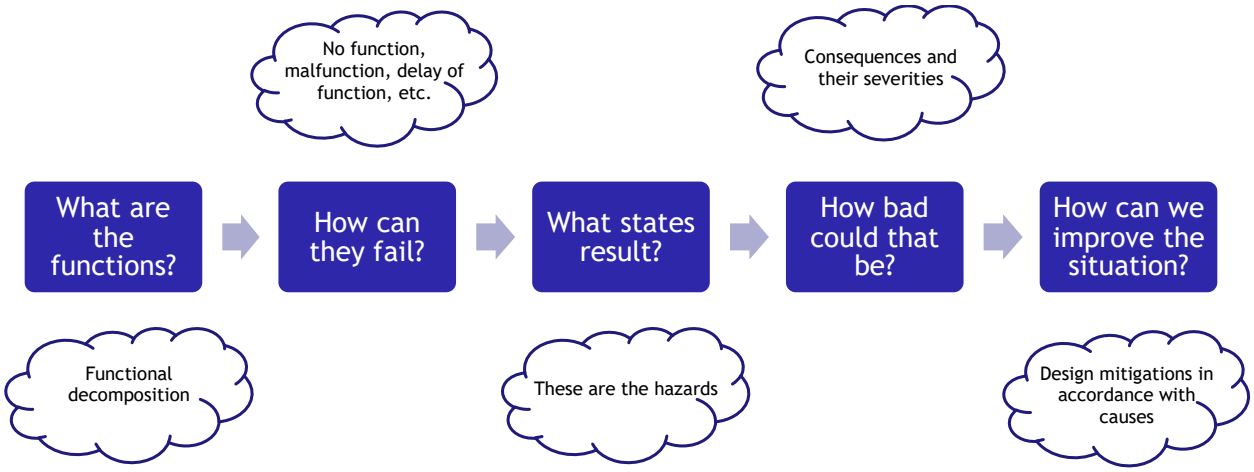

Figure 3: Intuitive FHA Goals

In this figure, we have consolidated a conceptual map of what we are trying to discover and accomplish with the FHA process, regardless of who's conducting it with which tools, expertise, and other resources.[12]

On this basis, we then resolved the initial challenges as follows:

1. **Terminology.** We anchored to consistent terms for salient concepts that support application of the mental model and sourced those terms from established and applicable references (terms to be presented with introduction of the table).
2. **Type hygiene.** We defined and applied consistent rules on the types of information required to fulfill specific roles in the tables and thus the analysis.
3. **Format variations.** We distilled a core FHA table format consistent with the foundations shared by available examples and the mental model and added further detail only where it served specific discussions for purposes of this demonstration.

4.1.2.4 FHA Table Format Used for this Exercise

FHA guides the collection and analysis of a lot of informal data by humans. To manage this data during collection and to organize it for further stakeholder consumption (system developers, project managers, certification engineers, regulators), we require a suitable container. For FHA, this is conventionally done with a table.

Though the use of tables is conventional here, there does exist variation across examples provided by standards and other guidance. Some have more or fewer columns, some use different terms for the same general concepts, and some leave ambiguity regarding the form or type of the data to be collected in various cells.

For this exercise, we present a baseline set of columns broadly consistent with best practice and introduce our negotiated terminology below. None of our choices is unconventional; however,

[12] An even simpler model, courtesy of Frank McCormick, puts it this way:
1. How can you hurt people?
2. What are you going to do about it?
3. How do you know when you're done?

since one challenge for new entrants in implementing established practice is making sense of inconsistent guidance, we label and clarify the data types and rationale for our table structure here.

The "Function" column directly captures the flattened functional decomposition, its hierarchical organization translated to a numbering scheme in the "UID" column (universal identifier). In our exercise, the *Aviate* function has UID 1.0, and all subfunctions of *Aviate* take the form 1.X.

The "Flight Phase" column facilitates consideration of each function and sub-function during each phase of flight. Because some hazards will appear in some phases and not others, or have varying consequences and severities across flight phases, separating in this way provides the structure to surface and document all such differences. Each can then be mitigated in accordance with its accurate profile. The flight phases for this exercise are as presented in section 3.3.

The "Hazard" column captures the descriptions of any immediate system states that are undesirable following associated functional failures. That is, by considering what states might result from the malfunction, delay, or absence of a function, we can systematically enumerate hazards that require mitigation. It can be of further value to consider whether these states are apparent to the operator or not.[13]

When considering these combinations, zero to any number of hazards might be identified per function and captured here. The numbering scheme accommodates unique identification of an arbitrary number of hazards.

A hazardous state is characterized by potential harm not yet realized. Realization occurs when something in the environment actively or passively precipitates the loss event. For example, a software fault might corrupt a critical piece of code in the flight control system. Until that code runs, the system is in a hazardous state, but no loss event has occurred. Once something in the environment (and the software control flow) triggers that code to run, it degrades flight control perhaps causing an unrecoverable loss. In this example, the function is the sub-function of flight control in consideration, the hazard is the erroneous software state deriving from the fault (malfunction), and this hazard has potential consequences including degraded flight control. The consequences are realized in the event the code runs unless adequate mitigations are in place.[14]

The "Severity" assignment is made here based on the standard classification documented in FAA AC 23.1309-1E [22]. This AC provides for the categories *Catastrophic*, *Hazardous*, *Major*, *Minor*, and *No Safety Effect*, and the differences between them. Of note for this demonstration, and for these systems and environments generally, is that familiar hazards that might be assigned one severity level conventionally are more reasonably assigned a higher severity level here in accordance with increased uncertainty and the effects of that uncertainty on possible consequences. For example, hazards that might be classified *Major* for manned commercial transport systems flying in the context of standard and universal operating protocols, might more reasonably be classified *Hazardous* when immature or lacking protocols impose additional workload on an operator responding to off-nominal conditions. Increased workload in particular is one factor that explicitly provokes a higher severity level in accordance with the AC. Since effectively all

[13] In a comprehensive FHA, analysts would in fact consider both cases, as well as any ambiguity between them for a given scenario, and track differences in implications through the analysis. For reasons of scope and focus, we note this dimension but do not follow it further through the discussion. The interested reader is invited to research the area of *failure semantics* and the related notions of design for clarity of system state.

[14] Even with adequate mitigations, the risk is non-zero, but has been definitionally reduced to an acceptable level.

dimensions of this application space, including infrastructure and air traffic management in addition to vehicle design are developing simultaneously, and none is mature, there is insufficient confidence at this stage to assume certain supports such as the ability to rely on universal protocols. Some conventionally *Major* hazards are thus declared *Hazardous* here unless and until sufficient basis exists to reduce this assignment.

We include for this demonstration a "Mitigations" column. This is not universal. We include it because it provides a process trigger to systematically seed and iterate potential responses to hazards in accordance with their causes while the hazard identification discussions are being had. While the activities of hazard identification and design of mitigations can be serialized, time separation, and often team separation, carries significant risks of information loss. Since the process of hazard identification is intrinsically connected to how the hazards arise, causes are already part of the discussion. And since proper mitigation begins by addressing causes of the hazards, there is a valuable opportunity here to capture preliminary mitigation information while it is available.

Given the above, the inclusion of a "Causes" column might also make sense. In fact, we did track causes during the analysis of the functions and their hazards. The choice to leave it out of the presentation is one of convention and space limitations. However, while a Causes column does not generally appear in conventional examples of FHA tables, this does not mean a stakeholder should ignore the option. For the reasons above, tracking causes during this analysis is worthwhile. It is also entirely reasonable to track all potentially relevant information during analysis and then, as we have done, curate the presentation for downstream purposes. Curation can then attend to stakeholders with particular needs, and to submission requirements, so long as the curation does not materially change the reasonable inferences allowable from the data, and the data remain accessible as needed. Further, it provides an interface to cross-validate and enable feedback loops between the FHA and other safety analyses. Fault trees, for example, provide a related but different view on some of the data captured during FHA. Fault trees and other safety analyses will briefly be discussed later in section 5.1.

In general, the structure of the collection container (that is, the working FHA table, as opposed to the final presentation form) should comfortably support the information capture needs of the team conducting the analysis. If something seems potentially relevant, especially in exposing the conceptual relationships addressed by the mental model, the team should establish a way to track that relevant element. As the analysis proceeds, some restructuring decisions will become apparent. Then for presentation, the working table structure can be curated for consumption by the target audience, with attention to that audience's information needs. Keep in mind that the audience, especially in the case of the regulator, is likely to have specific properties and data they're looking for based on protocol and experience, *including the model*. It is in an OEM's best interest both to anticipate these and to be very prepared to justify any deviations from expected practice.

4.2 Selected Discussions

In this section, we address several specific hazards in detail, with reference to their FHA table entries. For coverage, we detail one example functional failure within each of the four main system function hierarchies, and discuss some of the resulting hazards, consequences, and implications. We also include both less- and more-complex examples, illustrate the cascading nature of hazards that arise in some circumstances, and the capture of further data about these cascades.

For each function addressed, we further explore an example failure of that function via exercise within a hypothetical storyline consistent with our reference scenario and mission types. By placing each functional failure within a storyline and speculating events and states preceding and succeeding the failure, we can surface additional considerations relevant to the hazard identification and assessment, and thus to the design of the system to meet overall safety objectives.

Note that the examples as detailed here are *not* comprehensive, nor are they meant to be. They represent, rather, individual examples threaded through the FHA process for demonstration value. A complete FHA for a production system would be a significantly larger undertaking, applied to a more mature system specification, and OEMs new to system safety assessment are advised to plan and resource such activities accordingly.

4.2.1 Navigate Function

The *Navigate* function decomposes in this demonstration into several subfunctions relating to path planning and monitoring. The *Planning* subfunction further decomposes to subfunctions such as dynamic replanning. Dynamic planning is used to update paths to navigate around dynamic obstacles such as other air traffic, construction cranes, and other non-fixed objects in the target environment that create the possibility of collision. (These dynamic obstacles would be detected via subfunctions of path monitoring). For this example, we hypothesize a failure of function 2.1.5, Dynamically Replan, resulting from a hardware, software, or data error. Some considerations of a failure of this function during approach[15] are highlighted in table 1 .

Table 1: Hazard of Insufficient Path Planning

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
2	Navigate					
2.1	Plan path					
…						
2.1.5	Dynamically replan	approach	Insufficient path plan	Misidentified or underspecified waypoints and/or destination conveyed to operator; increased operator workload; impacts to situational awareness	Hazardous	Waypoint and destination confirmation; human factors adjustments to navigation interface; HW/SW/data fault management
2.2	Estimate position and orientation					
2.3	Monitor path					

[15] In doing a comprehensive FHA, the analysis team would consider failures for each function during every flight phase. We present only one example here for each narrated function. Likewise, the remaining discussion of the hazard assessment is demonstrative and not intended to be complete. See section 4.3 for more information.

To characterize implications of a failure of dynamic replanning, we consider effects of its malfunction, delay or complete lack of execution. Any of these can produce the hazardous state of flying with an insufficient path plan, with consequences including underspecified direction provided to the operator, increased operator workload to interpret navigation information, and incomplete situational awareness, that the operator might or might not recognize.

The hazard of flying with an insufficient path plan was declared hazardous in severity based on the uncertainties described in section 4.1.2.4. In particular, impacts to operator workload or situational awareness make it difficult to assign a lower rating, in accordance with AC 23.1309-1E [22].

Mitigations to address operator workload and situational awareness consequences include procedural support, such as protocols for waypoint confirmation, and adjustments to the human factors interface to clarify both navigation variables and system state [26]. Mitigations to address hardware, software, and data faults that contribute to the hazard arising include best practice approaches in fault management [13]. Consider, for example, a software fault as a source of this hazard. A latent software fault, upon execution, might create an erroneous data state. The navigation software then calls this data, for example in creation of an updated path plan. The updated path is invalid as a result (the Replan function has failed), and, depending on how other parts of the system are instrumented, this failure might or might not be apparent to the operator.

Notice that this chain of events includes cascading hazards at various system levels. Essentially, hazards are erroneous system states with potential for harm should a succeeding event externalize them. The erroneous data is inert until a susceptible function executes with it. This execution event then generates another erroneous state (invalid path plan) at the next system level. This invalid path plan might or might not be used. If used, this event extends the chain of potential consequences, and if not used, it still has potential impacts to operator workload in understanding and responding to the invalid state (if annunciated).

Mitigations might eventually also come via Simplified Vehicle Operations (SVO), one instance of which is the EZ-Fly concept [8]. SVO concepts are in development by several organizations to address multiple challenges of fielding eVTOL aircraft and related vehicles, including the recognized pilot shortage. These systems aim to lower the knowledge and experience thresholds required to fly these aircraft to levels sometimes described as similar to driving a car or playing a video game.

While SVO concepts have some promise to simplify tasks related to navigation, it must be noted that new functionality always comes with new hazards. While some hazards might be mitigated, others will not, and still others will arise from the new systems themselves. One such issue is that the question of emergency procedures and the role of the operator in off-nominal situations is not yet resolved. In piloted systems, there are standard emergency response protocols on which pilots are trained. For example, in the event that ATC cannot reach a pilot, every pilot knows what to do in that situation. In operational concepts oriented around far less training by design, significant challenges remain in the design and validation of emergency response protocols. The full system must always be evaluated in the state of its intended deployment.

4.2.1.1 Storyline

We now explore this hazard further via a hypothetical storyline. For this exercise, the aircraft deviates course due to weather on early approach to land. The *Dynamic Replanning* function

instructs the pilot to a new landing site. However, the updated instruction is underspecified relative to the structure of the target landing site; there is more than one landing pad and a plausible ambiguity regarding landing pad identification. The operator's workload is increased by the task of needing to resolve the landing pad ambiguity, and in addition, the operator resolves it incorrectly relative to the computed plan. The operator lands the aircraft on the wrong pad, impacting and damaging equipment and the aircraft. Service personnel conducting maintenance on the landing pad are unharmed, though nearby and otherwise unprotected from the unplanned landing other than by chance. The operator and passengers are unharmed, though also in part by chance as the impact was not significant enough to impart excessive force.

Exercising this storyline around a failure of the dynamic replanning function affords us several insights, such as the following:

- Planning and replanning are critically dependent on data representations, and some of the entities needing representation in this new environment need new and complete data specifications. This should provoke an assessment elsewhere in the analysis of evaluation of software data hazards and related considerations.
- The chance avoidance of harm to persons on the landing pad in this scenario indicates another hazard; any presence of persons or equipment on a landing pad at any time means they are subject to harm from a landing aircraft. *This hazard would be Catastrophic because it could result in multiple fatalities.* Since persons and equipment do need to be present at times on landing pads for legitimate reasons, these occasions require further specification, analysis, and hazard mitigation. A speculated hazard might be failure of appropriate containment of persons, implying the need for a containment function in the *vertiport* specification; see table 2 for elaboration.

Table 2: Failure of Adequate Containment of Persons at Pad

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
TBA	Contain persons at vertiport landing pads	n/a (vertiport vs. aircraft)	Failure to contain persons	Injury or death from landing aircraft	Hazardous or Catastrophic	[Speculate based on how containment failed]

Note the UID is yet to be assigned; we have captured this speculated hazard, but it concerns the vertiport and not the aircraft and so is not recorded in the aircraft FHA; that is, this new requirement is outside the scope of our design authority for this system. Likewise, the mitigations concern the vertiport design and so only a placeholder is recorded. The particular dependencies between aircraft, traffic management, and infrastructure for the UAM capability space are in active discussion in the community; this kind of information should be shared across cooperating development and policy organizations when surfaced.

Finally, note that even though no persons were physically harmed in this scenario, there would be significant costs to the airline in repair of both equipment and infrastructure, in forensic systems analysis and improvement, and perhaps most critically, in public confidence.

4.2.2 Communicate Function

The *Communicate* function decomposes in this demonstration to internal and external communications. External communications concern communications between the aircraft and ATC, the aircraft and other proximate aircraft, or the aircraft and infrastructure such as dispatchers. For this example, we hypothesize a failure of communication with ATC during enroute flight resulting from a hardware or software error. Some considerations for this functional failure are shown in table 3.

Table 3: Hazard of Loss of External Communications

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
3	Communicate					
…	…					
3.2	Communicate externally	Enroute	Loss of external communications	Stale information; increased operator workload; impacts to situational awareness	Hazardous	Data timeouts; Communicaitons redundancy; procedural guidance for lost communications scenarios; HW/SW/data fault management

To characterize implications of a failure of external communication, we consider effects of its malfunction, delay, or complete lack of execution. Any of these results in a state of absent external communications at least temporarily, with consequences including stale information, increased operator workload, and/or decreased operator situational awareness, recognized or not by the operator. Keeping in mind our reference aircraft is piloted, these workload and situational awareness considerations apply to an onboard pilot. For a different aircraft concept, they could likewise apply to a remote pilot, an onboard or offboard non-pilot operator, or the analogous properties of an autonomous controller. That is, the definitions of any function and its associated hazards are always with respect to the particular system as specified. Any "reusable" hazard lists must always be customized accordingly, and the same property applied to different entities can have very different implications.

As in the previous example, we declared loss of external communications hazardous for the reasons previously discussed (see section 4.1.2.4). This is an instance of a familiar hazard (loss of communications) that has a lower severity rating in environments of higher standardization and established protocol; loss of communications is traditionally classified as *Major* if no further complicating factors are present. For this design in this environment, however, a classification of *Major* is not possible, as it rests on assumptions of protocols and practices that do not yet exist. Absent these protocols, concerns about operator workload and situational awareness compel a *Hazardous* severity rating.

Mitigations to address faults in the hardware, software, or data consist of applying best practices in fault management [13], to include architectural redundancy. Mitigations to address stale information include data timeouts where applicable, for example, when delivered visually; this does not work for voice. Mitigations for the operator workload and situational awareness

consequences include development of procedural guidance for lost communications scenarios, that rest additionally on the kinds of assumptions made possible by maturation of pilot/operator training and the traffic management system to be deployed in these environments. Since these are not yet mature for this area, dependent mitigations will be of limited reliability. As the environmental supports mature, so too will our ability to make and use assumptions in designing mitigations.

4.2.2.1 Storyline

We now explore this hazard further via a hypothetical storyline. For this exercise, we hypothesize an eVTOL aircraft enroute with commuters to a TOLA. The eVTOL aircraft encounters a GA aircraft on a collision course, and requests deconfliction support from ATC by voice; this is within the *Communicate Externally* function. The pilot and ATC begin negotiating the deconfliction resolution. External communications are lost during this ATC-supported encounter resolution process. In this case, the pilot has not received complete resolution direction from ATC, which has GA trajectory information not available to the pilot. The pilot and ATC therefore have different situational awareness of the encounter geometry. Unable to recover communications, the pilot makes a maneuver decision completely appropriate to the information available. However, under the circumstances, this maneuver exacerbates the conflict, resulting in a near midair collision (NMAC).

Exercising this storyline around a failure of the external communications function affords us several insights. Our first such insight is that hazards can cascade within system levels and not just through them. In this case, several states and events transpired at the aircraft level; contrast this to the state-and-event chain in the prior discussion in which failures in successive nested subsystems eventually surfaced a failure at the aircraft level.

Our second insight is that, while NMAC is uniformly catastrophic,[16] and the hazard of lost external communications eventually led to an NMAC, this does not render the lost external communications hazard catastrophic in severity. This is directly tied again to the notion of the state-event chain. Severity is assigned based on the possible consequences that are one event away. In this case, several transitions had to occur before the NMAC transpired. In particular, the lost communications led to a loss of situational awareness on the part of the operator (who had incomplete trajectory information regarding the GA aircraft). The operator then *acted* (event) on the information available to him, leading to a state of lost separation. In this state, the event of NMAC was made possible. In this accounting, only loss of separation is catastrophic.

Our third insight is that the fact of intervening states and events raises the consideration of additional functions to capture. That is, each successive hazardous state implies a function that has failed. In theory, we have captured the functions during the initial top-down decomposition. In practice, the initial decomposition creates only a structured starting point for the analysis. In this case, both loss of situational awareness and loss of separation succeed the lost communications. This implies the existence of functions supporting safety requirements of maintaining situational awareness and maintaining separation. Though we have not yet considered where these might

[16] The most severe potential consequence of lost separation is a midair collision (MAC), which can be expected with unacceptable likelihood to result in loss of the aircraft and all persons onboard. Near-midair collisions (NMACs) are considered just as severe for assessment purposes, as the difference between MAC and NMAC as they are defined in terms of airspace volumes is inconsequential to the overall risk exposure. That is, an NMAC is considered a MAC avoided by luck.

integrate into the decomposition, we know they need to be there, and to be likewise followed through in the assessment. Their initial captures might appear as in table 4.

Table 4: Additional Functions

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
TBA	Maintain Situational Awareness	all	Loss of situational awareness	Increased workload, decision error, command error, …	Hazardous	TBD in accordance with causes; this hazard arises many ways
TBA	Maintain Separation	all	Loss of separation	Imposition of response actions on ATC, proximate aircraft, NMAC, MAC	Hazardous, Catastrophic	Strategic and tactical conflict management capabilities, including procedural separation and DAA systems

Maintain Separation can reasonably be integrated within the *Aviate* decomposition in our accounting.[17] *Maintain Situational Awareness*, on the other hand, has potentially several homes, which raises another insight: Some sub-functions cross-cut the decomposition. We first saw this during the initial decomposition round, when we considered support functions such as a *Manage Systems* function, to collect multi-purpose sub-functions like state monitoring and input/output. It is possible to either pull out cross-cutting functions and assign them top-level status, or to leave instances where they occur in the decomposition of other functions. What is most important is to recognize their existence and have a systematic plan for addressing them.

If we consider this particular loss of separation further, we might include this instance of the sub-function under *Navigate*, and specifically under *Dynamic Replanning*, because in this case the generation of a valid updated path plan (the avoidance maneuver) was compromised. This would be in contrast to a valid path plan that was incorrectly executed, perhaps due to a flight control failure, which would indicate an assignment under *Aviate*. Benefits to this inline organization allow the consideration of more specific causes in the design of mitigations. Benefits of batching cross-cutting functions and pulling them up to a higher level include leveraging their commonalities in designing more general mitigation approaches. As with most tradeoffs, a balance must be struck in accordance with the applicable needs. That said, an inline organization, at least to start, will provide a deeper accounting especially when dealing with novelty, as we are here. Abstraction from these details can then be layered into the process as necessary. True mitigation includes the follow-through of adequately mitigating all precipitating hazards (and their antecedents, to origin); any organization that supports this objective can work.

[17] Other options are possible, especially if the top level has been set differently. For example, in decompositions that include *Mitigate* at the top, this might be an appropriate parent function for *Maintain Separation*. Recall that the point is not to find a single correct organization; this does not exist. The point is to enable the surfacing and analysis of hazards and their implications. Many possible organizations support this.

Also of note for the *Maintain Separation* function is that its mitigations include systems associated with the vehicle design, and also systems that characterize the environment. That is, supports like procedural separation rely on design of the ATM/UTM system, itself currently evolving for these environments. There is potential for asymmetry in rules and equipage applicable to these two aircraft. Together, these can result in highly challenging encounter geometries given reduced separation minima in these dense environments, with little time to decide and execute resolution maneuvers, and lack of clarity in making the required decisions. This recalls the discussion of SVO and its challenges in section 4.2.1; any assumptions about pilot and/or operator and/or autonomous control of the vehicle must reflect the environmental systems with which this agent interacts.

4.2.3 Transport Function

The *Transport* function decomposes in this demonstration into transport of persons and transport of cargo. Transport of persons requires subfunctions including support of passenger ingress and egress, as well as support of passenger conveyance. Conveyance of passengers relies on such functions as appropriate cabin HVAC and pressurization, appropriate seating, and passenger restraints both to protect passengers in the case of in-flight turbulence as well as to maintain the weight and balance profile of the loaded aircraft during flight. For this example, we hypothesize a failure of passenger restraints, some considerations of which are highlighted in table 5.

Table 5: Unrestrained Passenger(s) Hazard

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
4	Transport					
4.1	Transport persons					
…						
4.1.2.2	Restrain passengers in flight	enroute	Unrestrained passenger(s)	Aircraft center of gravity out of bounds; reduced flight stability; passenger injury	Hazardous	Procedural passenger management; design to accommodate balance shifts
4.2	Transport cargo					

To characterize implications of a failure of passenger restraints, we consider effects of its malfunction, delay, or complete lack of execution. This example highlights the observation that these failure mechanics are a guide to brainstorming, as opposed to applicable for all functions. A delay in the function of passenger restraints is more difficult to make sense of than a malfunction (unreliable buckle, inconsistent belt tension) or absence of function (a completely broken seatbelt). In this case, either a malfunction or lack of function can produce a state of unrestrained passengers at least temporarily.

Consequences to the aircraft include a shift in the center of gravity, potentially outside of designed bounds, and associated reduced flight stability. Either of these is at least *Hazardous*, because they place the vehicle one event away from a loss of control should the hazard not first be mitigated. It

is of further note that the states of flight instability and out-of-bounds center of gravity are themselves hazardous states that give rise to additional consequences such as increased operator workload. These hazards are of *Hazardous* severity as discussed in sections 4.2.1 and 4.2.2.

Consequences to an unrestrained passenger include potential injury due either to any forces involved in the functional failure itself, or to falls or forces generated for example by an unstable flight path applied when a passenger is unrestrained.

Mitigations to an out-of-bounds center of gravity from this source include procedural passenger management: instructions to remain seated unless otherwise directed. A new consideration for these operating scenarios is that flights will not be crewed to the same levels as current convention provides, if at all, as autonomous control becomes possible. Implementation of passenger management as a mitigation will require attention beyond what is currently practiced.

Mitigations to flight instability include design of the flight envelope to account for off-nominal weight and balance profiles as might occur when a large internal shift relative to the vehicle mass occurs. Derivation of this envelope should take into account the degrees of freedom of unsupervised passengers (and untethered cargo, since those restraints are likewise subject to failure).

Mitigations addressing the initiating failure include appropriate attention to degradation and design faults to which the restraint components are subject.

4.2.3.1 Storyline

We now explore this hazard further via a hypothetical storyline. For this exercise, the aircraft is enroute with passengers from the downtown area of a large metropolitan city to its airport. The flight has an onboard operator and no other crew. Four passengers (a full flight for this configuration) are seated and begin with seatbelts fastened. The flight passes a famous landmark and a passenger releases his belt in order to go across the aisle to point it out through the window. Note that the failure here is of the passenger management mitigation (which implies associated function(s)), not the restraint itself; the passenger released the restraint through its proper mechanical design. A degradation failure of a restraint, on the other hand, might allow a sleeping passenger to fall out of her seat upon an aircraft bank, and lacerate her head during the fall to the floor. Note that either failure can cause overlapping sets of hazards and consequences AND the corollary that identical hazards can come from multiple independent sources.

Continuing with the storyline, the shift (from either source) causes a vehicle imbalance and the operator's attention moves appropriately to stabilizing the flight (first, fly the plane). The abrupt movement of the aircraft jostles another passenger, whose restraint was either malfunctioning or undone, into the aisle. Now perhaps 10 percent or more of the gross takeoff weight is in off-nominal position within the flying vehicle. The center of gravity has shifted and the operator is focused on controlling the aircraft under eroded stability. As remaining passengers react to the lacerated and stumbling passengers, the small cabin becomes further chaotic, and the unmitigated hazardous state(s) of the vehicle now predispose it to catastrophic scenarios.

Exercising a storyline around a failure of the passenger restraint function affords us several insights, such as the following:

- This storyline prompts us to consider refining our decomposition in order to train additional attention on the mechanical and procedural aspects of passenger restraint. We update the

restraint function to include these as subfunctions and document the update as shown in **Error! Reference source not found.**, allowing increased focus and traceability. Note the identical hazards and consequences from different sources and the overlapping but partly distinct mitigations in accordance with different causes.

- This refinement further demonstrates the notion that FHA is a dynamic technique that, done appropriately, updates with successive iterations of the system specification and assessments. Mitigations are responses to hazards, but they contribute design refinements, which create new functions and requirements, which themselves must be assessed for potential hazards. In theory, this can continue ad infinitum. In practice, there are stopping conditions. These will be discussed later in section 4.3

Table 6: Further Decomposition for Passenger Restraint

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
4.1.2.2.1	Apply mechanical restraint	enroute	Unrestrained passenger(s)	Aircraft center of gravity out of bounds; reduced flight stability; passenger injury	Hazardous	Management of design and degradation faults in restraint components; procedural passenger management; design to accommodate balance shifts
4.1.2.2.2	Apply procedural restraint	enroute	Unrestrained passenger(s)	Aircraft center of gravity out of bounds; reduced flight stability; passenger injury	Hazardous	Redesign procedural passenger management; design to accommodate balance shifts

Finally, as in the *Navigate* example, this scenario would likely cause repercussions to public perception and rework to emergency procedures, operator training, and passenger management protocols.

4.2.4 Aviate Function

In this demonstration, the *Aviate* function decomposes into control of flight and ground paths and control of subsystems. Control of the flight path relies on control of several parameters including altitude, attitude, and propulsion, while control of subsystems includes powerplant management and other functions. For this example, we hypothesize a failure of propulsion, some considerations of which are highlighted in table 7. We further address an important common cause failure for eVTOL systems through this example.

Table 7: Hazard of Loss of Propulsion

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
1	Aviate					
1.1	Control flight					
…						
1.1.1.4	Control propulsion	transition	Loss of propulsion	Impacts to aircraft controllability, including loss of control	Catastrophic	Decoupling of propulsion and control; architectural redundancy; glide and other off-nominal landing options
1.2	Control ground movement					
1.3	Control subsystems					

To characterize implications of a failure of propulsion, 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 of circumstances, a loss of propulsion can be hazardous but recoverable. More commonly, and uniformly during the transition phase from vertical to forward flight (or the reverse), this hazard will be catastrophic, as a loss of control from any but a very low altitude will likely result in fatalities.

Mitigations to loss of propulsion address several points in the causal chain of states and events, as well as aspects of the differences among eVTOL designs. Some of the *causes* of propulsion loss include software, hardware, and data faults, for example a premature motor wearout, or a battery controller bug. These are to be addressed via best practices in fault management, including, for example, architectural redundancy. Distributed electric propulsion includes some redundancy by design, and can enable some retention of control with a subset of motor/rotor loss. Decoupling of propulsion and control through mechanisms including actuators not necessary for the propulsion of the vehicle provide another option, allowing another avenue for retained control when propulsion is lost. Note that such actuators add weight as well as failure modes and potential hazards; should they be designed as a mitigation, their own function must be flowed back through the analysis.

Some mitigations target the severity of the consequences once propulsion has already been lost. These include ways to retain some control in a lost propulsion state, as well as ways to reduce impact forces in the event of an uncontrolled landing. For example, a wing-borne eVTOL aircraft in forward flight has some possibility of a controlled glide and landing absent propulsion, assuming

a sufficiently uncluttered glide path and usable landing site. Note carefully that these assumptions must be validated for the target environment, where buildings, ground vehicles, and other structures, and more importantly, people, are present in higher density.

Further, a wingless eVTOL aircraft does not have this glide option, and the rotor complement of a glide, autorotation, will generally not be an available feature for eVTOL designs.[18] At the time of this writing, there is as yet no community consensus on acceptable mitigation to the loss of propulsion and/or control during the transition phases from vertical to forward flight or back again. Much discussion thus far has centered on response-type mitigations to limit harm, such as ballistic parachutes and energy-absorbing materials, but analysis of their effectiveness will be extremely challenging if tractable at all.[19] We revisit this issue later in this section.

4.2.4.1 Storyline

We now explore this loss of propulsion hazard further via a hypothetical storyline. For this exercise, the aircraft is a hexacopter in transition from vertical to forward flight and loses effective propulsion due to a Byzantine fault.[20] In this case, we hypothesize that the six rotor controllers must synchronize on certain data values in order to work together correctly, but a synchronization error arises as a result of data corruption enroute from the master controller to one (or more) of the rotors. The rotors are then working to different ends, unbalancing their respective outputs and distorting the overall propulsion profile. In this flight phase, almost any distortion of propulsion results also in loss of control, since the aircraft is inherently unstable during transition. The consequence is catastrophic, in that the aircraft and all passengers will likely be lost.

This storyline reinforces recognition that control of propulsion relies on subfunctions related to computing, commanding, and actuating the intended rotor behavior. We might then capture these subfunctions in our decomposition and follow them through individually in order to train attention on their specific considerations. Additional table rows are hypothesized in table 8.

The additional table rows show how a cascade of hazards can be traced through several functional failures, and a pattern appears in the data. If the *Compute* sub-function fails, the system enters a hazardous state of not having a correct result of the *Compute* function. It is not until the *Command* function tries to use this data (an event) that a consequence is realized: no command or the wrong command is sent. Again, the system enters a hazardous state of no command or an incorrect command assigned the actuators, which then have undesired effects on propulsion only when they actuate (incorrectly or fail to actuate at all).

[18] Autorotation and related issues are further discussed in section 4.2.4.3.

[19] Accurate reliability estimation will rely on modeling and simulation (M&S) due to limitations of real-world testing, and the value of that M&S will rely on modeling a sufficient variety of scenarios at sufficiently high fidelity to provide useful results. This kind of open-world modeling is proving very difficult for aerospace and automotive applications alike. Some discussions suggest artificial intelligence and machine learning as possibilities to support generation of sufficiently broad and detailed models, but this work is both immature and comes with its own set of new challenges, particularly in verification of AI and ML software and training data, and qualification of associated tools.

[20] A byzantine fault is a fault in a computing system that manifests differently at different system interfaces, making it difficult to determine whether a component has failed and how.

Table 8: Loss of Propulsion Decomposition

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
1.1.1.4	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
1.1.1.4.1	Compute motor/rotor commands	transition	No result or incorrect result produced	No command or wrong command available to send	Hazardous	Proper HW and SW fault management practices
1.1.1.4.2	Command motors/rotors	transition	No command or wrong command sent	No actuation or wrong actuation	Hazardous	Proper HW and SW fault management practices
1.1.1.4.3	Actuate motors/rotors	transition	No actuation or wrong actuation	Impacts to propulsion profile, including loss of effective propulsion	Catastrophic	Proper HW and SW fault management practices

In this example, each of these sub-functions would be engaged in some way via the Byzantine fault. If we hypothesize that it begins via something like an off-specification voltage in a command transmission path, then the command function has already failed, and any directly dependent actuation is also compromised. However, part of coordinating the rotors includes data sharing and voting to enable synchronization. If one rotor now sends back bad data, and this is not detected, then the compute function is now also compromised, as it is computing commands from corrupt information. This insidious property of Byzantine faults can act to corrupt multiple aspects of the system state from a single source, and simultaneously mask that source.

The identification and repair of faulty components in distributed coordinating systems is a main objective of research surrounding the Byzantine Generals Problem and the target of mitigations for this scenario. It is also exceedingly difficult, and it has been recognized that such scenarios are becoming more common due to both technical and social factors of modern systems and development environments [15]. The redundancy inherent in distributed electric propulsion provides some clear benefits to eVTOL systems, but DEP is also subject by design to Byzantine faults. Redundancy alone cannot confer its benefits unless its risks are also systematically managed.

4.2.4.2 On Cascades and Common Causes

The loss-of-control hazard is perhaps the most significant for any aircraft in keeping with the common mandate to first fly the plane. In this storyline, it followed a loss of propulsion in a hazard cascade but arises also from many other sources.

Loss of control through any other cause would also be catastrophic here, as in either case, neither vertical nor forward flight can be recovered from the control loss in transition. Since identical hazards with identical consequences can have multiple sources, however, mitigations must adequately address all of the ways that a hazard can come to pass. It is useful to exercise storylines in part because they reveal these causal chains and cascades. It is further useful to realize these causal chains are more properly thought of as paths through causal nets, where any hazard might have multiple antecedents and consequents in the state-event map.

This mental model also clarifies that since a single hazard can have multiple consequents, then there are distinct hazards that can arise from common causes. Continuing with the propulsion and control examples, loss of each of these functions can be precipitated simultaneously by a loss of power. That is, a failure of the powerplant subfunction (under "Control Subsystems" in our decomposition) during the transition phase of flight will cause loss of both power *and* control. Any version of this scenario is catastrophic. This consideration is outlined in table 9.

Table 9: Loss of Power

UID	Function	Flight Phase	Hazard	Operational Consequence	Severity	Mitigation(s)
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

Possible mitigations to a loss of power include compensatory redundancy in the form of hybrid (electric plus combustion) power and/or battery backup that provides enough energy for a controlled landing. These options of course add weight, change configuration geometry, and contribute other considerations that must be followed through in the analysis. At the time of this writing, hybrid power is expected to be leveraged for some intermediate concepts, but not expected to be a viable long-term solution due to the configuration inefficiency imposed by architectural dissimilarity. Battery technology is expected to continue to improve, but the necessary tradeoffs and new hazards imposed by battery redundancy still require significant evaluation.

4.2.4.3 On New Challenges with Old Hazards

Mitigations to the loss of propulsion, control, or power in the transitions between vertical and forward flight are among the most critical for the entire safety assessment for a passenger-carrying eVTOL aircraft and have complex dependencies among them.

As noted earlier, if in forward flight and under the best of circumstances, a wing-borne eVTOL aircraft has the possibility of a controlled glide to land absent power and propulsion. This option is not available if the aircraft is wingless. The analogous rotorcraft feature, autorotation, uses inertial energy stored in the main rotor of a traditional helicopter to allow a controlled landing in the case of lost power and/or propulsion. However, autorotation relies on a sufficiently high ratio of rotor disc size to aircraft weight to contribute enough energy to allow control of the vehicle. eVTOL aircraft have multiple but smaller rotors; the combined disc areas are generally insufficient

to support the required ratio, and some electric configurations cannot store inertial energy in any case (the rotors simply stop turning when power is lost). Absent autorotation, community discussion on new lost power and propulsion mitigations is active. The inherent redundancy of DEP does allow some retained control if a critical subset of rotors is still functioning; one concept has demonstrated retained control on an 18-rotor eVTOL aircraft with 4 rotors non-operational [6]. Recall, however, that this benefit is only achievable if the attendant hazards are adequately mitigated.

At the current stage of technology maturity, the community does not yet have an answer to sufficient mitigation of lost power and/or propulsion and/or control of an eVTOL aircraft in an urban environment, especially during a transition phase. The difficulty of translating conventional mitigation means into the constraints of this new application have led to valuable discussions that consider options beyond the conventional. One ongoing discussion is centered on crashworthiness and aims to mitigate the harm possible once a loss of one of these critical functions has occurred. It offers ballistic parachutes, energy-absorbing materials, reverse jets/thrusters, and variations of airbags in order to reduce impact forces associated with a falling aircraft.

These discussions should continue and will produce novel partial solutions, but they will not solve the problem. At a minimum, each adds its own set of potential hazards and other considerations that must be followed through; one can only add so much weight or complexity or other property to the design before it fails to meet other requirements. In addition, while these responses can reduce potential harm to passengers, they do not address risk to third parties and property in the crash path, nor the wider effects to public perception and business interests of using parachutes and other recovery systems as a primary line of defense.

This is in general because recovery systems *should not* be used as a primary line of defense. Best practice instructs rather that faults (and thereby the hazards they cause) be managed in proactive priority, first through avoidance, then by elimination of those that could not be avoided, then by tolerance only of those that could not be eliminated, and finally by forecasting, in order to most strategically cope with undesired effects known to be coming [13]. Concepts that focus on crashworthiness over hazard mitigation, and especially positions that argue for certification credit for recovery systems absent the intended means of meeting target assurance levels, fundamentally violate this principle. They mitigate the accident, while abdicating attention to mitigation of the hazard that allowed it.

The FAA will ask hard questions about the assessment and mitigation of these hazards. The successful OEM will have executed and substantiated a safety assessment that is thorough and compelling, convincingly identifying and mitigating all credible hazards of the design configuration in its intended environment. This includes proactive fault management, as well as the consideration of harm to third parties. Crashworthiness as a primary response does not fulfill these properties.

4.3 When Are We Done?

The exercise above is not a complete FHA. Rather, it demonstrates identification, capture, and analysis of FHA data via selected examples. A production FHA will involve significantly more data and several rounds of revision and stabilization. This raises the questions of how much data and how many rounds are needed. The pragmatic answer is that there must be enough to

substantiate, to the FAA's satisfaction, the assertion that all credible hazards have been identified and adequately mitigated.

In practical terms, this is not simply a question of volume, but of coverage and depth. Achieving adequate coverage requires (1) considering each function, in each flight phase, through enough decomposition levels and (2) collecting all credible hazards—potentially numerous—for any given function, together with potential variance in severity and other profile factors by flight phase. Should the causes of the hazards vary, so will the mitigations. The cascades of hazards, and the intersecting state-and-event chains we have discussed in fact create a web vs. a tree, and all relevant paths must be traced and addressed. Further, the process used by the OEM must be adequately documented such that the FAA can make sense of the data and the decisions that led to it. Finally, the mitigations must be validated and the final assessment (with mitigations applied) must demonstrate that overall target levels of safety are achieved.

There are some process indicators to help identify progress short of the final goal. Multiple feedback loops will be active early in the FHA process and begin to settle as the analysis stabilizes. For example, exercising the initial decomposition through storylines will generate rearrangements and new functions as long as there is new space to explore. This might settle as items are captured, only to reactivate as follow-through on those new items guides exploration of further considerations. At some point though, new items stop appearing.

This does not finalize the decomposition or the hazard identification and management. The process of designing and validating mitigations can continue to feed back into the hazard characterizations. For example, a planned mitigation might not pan out in testing, and require redesign, and/or a reconsideration of the cause of the hazard. And importantly, the history of safety-critical systems failures makes clear that hazards can and do continue to be identified during deployed operations. This is of course what we are trying to prevent, but when we fail, our next best steps are to use the new information both to update the assessments and mitigations for the current system, as well as feed the lessons forward into new systems.

Thus, in theory, we are never done.

4.4 Takeaways

Some observations made during execution of this exercise bear repeating for their value to stakeholders new to system safety assessment or to its application under new uncertainties. These observations are collected here:

1. The composition of certification bases for eVTOL aircraft, and some of the means of compliance to those bases, are immature and developing at the time of this writing. However, these uncertainties make no impact to the relevance and applicability of foundational system safety assessment principles and practice. In fact, **the recognized uncertainties in the development and regulatory environments for eVTOL aircraft call for disciplined leveraging of established safety practice to best support identification of, and response to, new hazards and new profiles for familiar hazards.**
2. **There is not a single right way to organize a functional decomposition for an FHA.** Different lenses will net different insights. We are searching for decompositions that offer traction: those that seem to allow placement of all concerns in places that make sense to the analysis team as a whole (and beyond), and for which the decomposition strategies are

relevant to the search for hazards. Thus, a decomposition that only makes sense to one person is probably not viable, and non-trivial workshopping to trial and discuss options is to be expected. Similarly, decomposition steps might be accomplished in many ways, only some of which are useful; whether we choose to sort objects by shape or color depends on what we're doing with them next. In any case, the primary goal of the decomposition activity is to establish a starting point for the hazard assessment. Be sensible in its generation, and also know that it will change.

3. Because we bring in prior knowledge and want to find ways to use what we know, **it can be easy to overcomplicate the artifacts and the workflow.** For example, we might have a preconceived decomposition for a function based on expertise or some other reference. However, it might be more information than is appropriate at a given point in the exercise, and can therefore distract from rather than advance a line of questioning. Further, it might not be the right decomposition for the circumstances, depending on the decomposition steps it uses. Similarly, the numbers of columns in the table, the number of core functions at the top of the tree, and other early decisions should all tend toward the more focused and essential end of the spectrum. We can add features as warranted, but ensure they are warranted. Otherwise, they only serve to bog down the process and muddy the product.
4. Conversely, **we need to be very careful collecting data into that essential framework.** *Type hygiene* is attention to ensuring consistency in the types and relations of data represented in the artifacts. Since a hazard is a state, all descriptions in the hazard column should describe states. Likewise, teasing apart the state-and-event chains requires careful attention to the differences between states and events. Otherwise we cannot properly characterize cascading hazards. This is a topic not well-addressed in the standards or the literature, but good type hygiene is critical to the ability to make subsequent decisions based on the data. Additionally, it helps analysis teams to partition the work; different performers with different practices for collection cannot integrate their parts into a coherent whole.
5. **Familiar hazards might behave differently in new environments.** Conventional severity assignments for common high-level hazards are based on long-standing protocols and practices that allow assumptions about performance and behavior of aircraft, traffic management, and humans acting within associated control loops. If we can't make the same mitigating assumptions that we can make in other environments, then we can't, for example, justify the same severity rankings. We identified some hazards that are traditionally classified as *Major*, but for which that ranking was not possible under the current circumstances of the target environment. These familiar hazards were classified here as *Hazardous* in part because of the effects they implied to operator workload and other factors.
6. More generally, **ALL familiar concepts must be assessed for how they might behave differently under new circumstances.** For example, proprietary hazard lists that have been developed over time through abstraction from experience are necessarily generic and must be customized for the target system. Customization for the specified target system can be significant. The trend from human to autonomous decision-making and control, for instance, raises hard questions about the meanings and roles of concepts like situational awareness.

7. **Identical hazards can come from multiple sources, and individual hazards can give rise to multiple consequences.** Part of FHA is uncovering the state-and-event chains that lead to particular hazards and continue on to cause undesirable consequences. In doing this, we find there are often many ways to end up in an undesirable state. This has some implications for our practice: (1) we need to surface as many realistic causes as possible, because any we miss leave unaddressed risk; and (2) mitigations will sometimes be plural and varied for a single hazard, because good mitigation strategies start with the hazard sources. Further, since multiple consequences can arise, multiple profiles for the hazard are possible, and each must be characterized appropriately to allow for adequate management.
8. ***New technologies and environments can also mean entirely new hazards, and/or new ways of mitigating familiar ones*.** Autorotation as a mitigation for loss of power, for example, does not hold special status here, despite being a fundamental safety requirement for all single-engine helicopters and some other VTOL aircraft. Rather, the focus is on what hazard is being mitigated and whether the mitigations together are adequate to support the target levels of safety. If alternative mitigations to power loss both adequately address the hazard and do not compromise other areas of the safety assessment, then autorotation becomes irrelevant to the application. Note the inclusion that other areas of the safety assessment not be compromised; recall that any new capability, including new mitigations, brings its own set of safety considerations. Mitigating hazards is itself hazardous and the new functionality must be flowed through the same assessment. The driving goal, via whichever set of functionality and dependability attributes are stabilized, is to satisfy the overall safety argument.
9. **FHA of one system can indicate mitigations and/or safety requirements needed in a different system.** We presented examples in which mitigations to eVTOL aircraft hazards included considerations to be designed into other systems, such as vertiports or the ATM/UTM system. Surfacing these mitigation options highlights the dependencies between a system and its environment, including the assumptions we make about these interactions. It also provides insight to be shared among stakeholders and developers of the interacting systems, supporting mutual alignment as these concepts evolve.
10. **Crashworthiness does not appropriately mitigate hazards that can lead to crashes.** Ballistic parachutes, energy-absorbing materials, and other features have been proposed as mitigations to hazards that can lead to crashes (loss of propulsion/control/power). In theory, these features *can* reduce the likelihood and/or severity of harm to passengers onboard an eVTOL aircraft in a crash event. However, crashworthiness as a primary response, and especially as a proposal for certification credit, insufficiently addresses critical factors of these hazards. Among them, risk to third parties in the crash path is unchanged. In the dense urban environments of intended operations, these risks are significant. Further, crashworthiness in fact mitigates the accident but does not address causes of the hazard that allowed it. Fundamental fault management practices conversely are proactive and still apply.
11. **The FHA, and the entire system safety assessment, are living artifacts.** The function decomposition, the hazard identification, the applied mitigations, and the containers that organize them, are all subject to revision and update as warranted throughout the system development as well as during its deployment and operation. Ideally, we want to learn most

of what we need to know sooner rather than later, but we also want the requisite flexibility to integrate impacts and lessons from downstream events. For example, if we learn during V&V that a particular mitigation will not work sufficiently, we need at least to redesign the mitigation, but sometimes we need to redesign the system to otherwise remove sources of the hazard. This can force more significant reassessment and revision to the safety artifacts. Similarly, should an accident happen in an operational system, we want to know whether something was missed or mischaracterized in the assessment, and propagate lessons forward.

5 FHA in the Greater Development and Certification Process

FHA begins early in the development process and continues so long as there is updated information available to it. It connects to other processes via this shared information, in support of delivering a system that demonstrably meets its functional and dependability requirements.

5.1 Relation to Other Safety Processes and to System Development

FHA is the first step in system safety assessment per SAE ARP4761. This standard provides guidelines for further steps in the process, as well as methods for conducting various supporting assessments such as fault tree analysis (FTA) and failure modes and effects analysis (FMEA). Different assessment methods answer different questions about the faults and failures to which a system is subject. For example, while FHA asks what hazards can arise given a set of functions, and characterizes those hazards to support design of mitigations, FTA asks how a given hazard can arise, in terms of events and combinations thereof. FMEA, in contrast, begins from a component failure and asks what consequences are possible as a result [13]. Notice FMEA and FHA both ask about consequences of failures, but FMEA begins from components (i.e., architecture), while FHA begins from functions (i.e., intended behavior). All of these methods and the others referenced in the standard provide ways to expose and clarify the state-and-event chains and nets that lead to undesirable events at the system level, so that they can be managed.[21] This underlying mental model connects all of these analyses.

SAE ARP4761 (on safety) is invoked by SAE ARP4754A (on the system development process more generally). SAE ARP4754A addresses the system development process in support of certification, covering the planning and development of the system at all of its levels, the allocation of functionality to subsystems and components, the assignment of DALs, and validation and verification against requirements. Connections between this process and SSA include aligning the prioritization and rigor of mitigations to the assigned DALs, and feeding safety implications, such as new requirements derived from mitigation needs, back into the development process. Analogous standards exist for defense systems, and for other safety-critical domains such as medical devices and automotive systems. They have much in common, and some interesting differences, which are beyond the scope of this work.

[21] In this work, we have not addressed failures that are not functional. Failures can and do occur that cannot be associated to the incorrect execution of a function. For more information, the interested reader is invited to research Safety of the Intended Functionality (SOTIF).

5.2 Relation to Software Development

This work addresses aircraft-level hazard analysis consistent with SAE ARP4761. From there, an OEM would then proceed to system-level hazard analysis, and on to sub-systems and lower levels, including software. Within the development process running in parallel in accordance with SAE ARP4754A, some system functionality is allocated to software. Once we begin to focus on software in the development and safety processes, another set of concerns comes into play. Since so much of the functionality of eVTOL aircraft and other complex, safety-critical systems is implemented in software, OEMs must prepare accordingly for the software aspects of certification.

Software is fundamentally different from bridges and processors and other concrete physical systems because it is essentially discrete math. Its failures do not fit a continuous reliability curve like a brick, transistor or hinge. Because of this, the identification and management of software faults and hazards imposes additional challenges. RTCA DO-178C, *Software Considerations in Airborne Systems and Equipment Certification* is a software design assurance standard covering the development process and much of the fault management process for airborne software [18]. Many of the system hazards we discussed earlier in this work will have mitigations related to software assurance; DO-178C or similar software design assurance practices will support identification and assurance of adequate mitigations.

Some eVTOL concepts and development methods will run into challenges identifying criteria and means of compliance for some software-supported capabilities. For example, machine-learning-enabled components cannot generally be verified by conventional means required by DO-178C such as demonstrating code coverage and traceability. This means, as a result, that any part of the SSA relying on that software will also face challenges, unless and until the community reaches consensus on certification requirements for ML and other new methods and techniques. The trend is toward performance-based standards as discussed in section 2.1. This means we will likely see proposals that focus on convincing an evaluator that an ML-enabled component possesses a certain property with a certain level of confidence, in some way that achieves this without need for, e.g., traceability. However, at the time of this writing, this consensus is likely several years off, at a minimum.[22]

5.3 Relation to Certification in the Large

FHA initializes SSA, and SSA, together with other processes, supports aircraft type certification. As noted throughout this work, type certification processes are evolving in order to cope with modern challenges associated with development and assurance demonstration of novel capabilities. Certification bases for the first several eVTOL type certifications are expected to be established in case-by-case fashion since eVTOL aircraft do not conform to any single aircraft type currently specified in federal regulation. These certification bases are likely to include various special conditions at a minimum, and perhaps be entirely custom-built through pulling applicable criteria from several types and negotiating additional criteria to fill gaps not addressed at all by any current type. As such, these certification bases will be idiosyncratic and more complex than what might be conventionally expected. Experience gained through these initial submissions, together with criteria and methods of compliance in current development by standards

[22] Production and agreement of standard criteria and methods of compliance for verification of ML-enabled components and other uses of AI is one objective of the G-34 committee of SAE on Artificial Intelligence in Aviation. At the time of this writing, the committee has just kicked off and has a 3-year minimum initial timeline.

development organizations, will provide the foundations for the regulatory and development community to refine and establish a more uniform practice.

Ultimately, the assurance objective is to convince the regulator that the system will predictably behave as intended and within acceptable rates of failure. Recall the Executive Director of the FAA Aircraft Certification Service states, "If you always come in with safety, you will keep moving forward" [17]. Further, *the OEM needs to be able to communicate this safety assertion and its bases in an integrated way that is accessible to the regulator*. That is, submission of a large volume of development data, even if the development data are in fact complete and valid, still places the burden on the regulator to interpret that volume of data into the story and rationale for safety. This is more properly the OEM's responsibility; the OEM must make the case to the regulator.

As such, and in response to the many uncertainties of the modern development and assurance demonstration environment, submission formats, organizations, and presentations for the body of evidence and inference in support of a certification are also evolving. A model seeing increased adoption in recent years in several safety-critical industries is a safety or assurance argument, in which a system is logically argued safe (and/or otherwise assured) for its intended use in its intended environment, often through a decomposition over system requirements satisfaction and hazard mitigation, or another set of properties that covers the intended and excluded behavior of the system. The argument is created and documented explicitly, allowing traceability and audit of the assertions, inferences, and evidence. Textual, graphical, and tabular formats are all possible. The basis, organization, and import of the structured argument should be summarized in an associated report, which can serve as a tool to support communication with the regulator.

The FHA activity demonstrated in this work results in a hazard list that can be input to the construction of a safety argument, and in the context of which an argument for adequate mitigation can be evaluated. Assurance argumentation factors centrally in several active community efforts including the FAA's Overarching Properties initiative for assuring and certifying complex systems [1], and UL 4600, a new cross-industry standard for assuring autonomous systems [20].

6 Summary

Novel eVTOL technologies, and their applications in novel environments, challenge existing foundations and means of aircraft type certification. Many performers in this space are also novel entrants to the aerospace development and regulatory frameworks that provide us the remarkable aircraft safety record from which we all benefit. This work offered an overview of those development and regulatory frameworks, together with an overview of challenges imposed by eVTOL designs. This work then narrated application of an origin step in established safety practice, as applied to a reference eVTOL aircraft, for demonstration and guidance purposes.

We first provided an overview of the regulatory environment within which eVTOL aircraft are emerging and noted some implications for development and certification.

We then hypothesized a reference aircraft design consistent with technologies in development and their intended applications and conducted an FHA demonstration exercise on elements of this design to (1) illustrate the process to new entrants and those in hybrid roles, and (2) further explore some of the design and safety considerations at the forefront of community discussion surrounding certification and deployment of eVTOL systems.

The demonstration offers example consideration of hazards arising from failures across the functional decomposition for this system. It further highlights (1) some options and strategies available in conducting FHA and how to attend to and choose among them with reference to a particular program, and (2) some hazards of new or different salience with regard to eVTOL aircraft as compared to more familiar designs.

We then consolidated and reviewed observations from the exercise and related the FHA activities and outputs to the greater system development and certification process.

7 References

1. C.M. Holloway, "Understanding the Overarching Properties," National Aeronautics and Space Administration (NASA), Langley Research Center, Hampton, VA, NASA/TM-2019-220292, 2019.
2. Code of Federal Regulations, Title 14 Part 21: Certification Procedures for Products and Articles, 2019.
3. Code of Federal Regulations, Title 14 Part 23: Airworthiness Standards: Normal Category Airplanes, 2019.
4. Code of Federal Regulations, Title 14 Part 91.113: Right of way rules: Except water operations, 2019.
5. *Design Assurance Guidance for Airborne Electric Hardware*, RTCA DO-254, RTCA, Inc., Washington, DC, 2000.
6. F. Colucci, "Turning Volts to VTOL," *Vertiflite*, pp. 30-33, Jan/Feb 2018.
7. FAA Aviation Safety, "Fly the Aircraft First," *FAA Safety Briefing GA Safety Enhancement Topic Fact Sheet*, July 2018. Accessed Nov. 29, 2019. [Online]. Available: https://www.faa.gov/news/safety_briefing/2018/media/SE_Topic_18-07.pdf.
8. FAA Aviation Safety, "Regulatory Roadblock Reduction," *FAA Safety Briefing GA Safety Enhancement Topic Fact Sheet*, June 2019. Accessed Nov. 20, 2019. [Online]. Available: https://www.faasafety.gov/files/events/GL/GL03/2019/GL0393086/Regulatory_Roadblock_Reduction.pdf.
9. G.F. McCormick, "Certification of Civil Avionics," in *Digital Avionics Handbook*, 3rd ed., C.R. Spitzer, U. Ferrell, and T. Ferrell, Eds. Boca Raton: CRC Press, 2015, pp. 9-1—9-15.
10. *Guidelines and Methods for Conducting the Safety Assessment Process on Civil Airborne Systems and Equipment*, SAE ARP4761, SAE International, Warrendale, PA, 1996.
11. *Guidelines for Development of Civil Aircraft and Systems*, SAE ARP 4754 Rev. A, SAE International, Warrendale, PA, 2010.
12. I. Smith, "Flying Cars and eVTOL Aircraft," Interview with Scott Drennan, Vice President of Innovation at Bell Flight, *Commercial Drones FM Podcast*, Episode 86, March 13, 2019.
13. J.C. Knight, *Fundamentals of Dependable Computing for Software Engineers*. Boca Raton, FL, USA: CRC Press, 2012.
14. J. Williams, "The Quiet Revolution: What Part 23 Changes Mean for You," *FAA Safety Briefing*, May/June 2019, pp. 8-10. Accessed July 24, 2019. [Online]. Available: https://www.faa.gov/news/safety_briefing/2019/media/mayjun2019.pdf.
15. K. Driscoll, B. Hall, H. Sivencrona, and P. Zumsteg, "Byzantine Fault Tolerance, From Theory to Reality," in *Proc. Computer Safety, Reliability, and Security, 22nd International Conference, SAFECOMP 2003*, Edinburgh, UK, Sept., 2003.

16. K.J. Hayhurst, J.M. Maddalon, P.S. Miner, G.N. Szatkowski, M.L. Ulrey, M.P. Dewalt, and C.R. Spitzer, "Preliminary Considerations for Classifying Hazards of Unmanned Aircraft Systems," National Aeronautics and Space Administration (NASA), Langley Research Center, Hampton, VA, NASA/TM-2007-214539, 2007.
17. M. Huber, "FAA Moving to Smooth Aircraft Certification," *AIN Online*, July 23, 2019. Accessed Aug. 21, 2019. [Online]. Available: https://www.ainonline.com/aviation-news/general-aviation/2019-07-23/faa-moving-smooth-aircraft-certification.
18. *Software Considerations in Airborne Systems and Equipment Certification*, RTCA DO-178C, RTCA, Inc., Washington, DC, 2011.
19. *Special Condition: Vertical Take-Off and Landing Aircraft*, SC-VTOL-01 Issue 1, European Aviation Safety Agency (EASA), Cologne, Germany, 2019.
20. *Standard for Safety for the Evaluation of Autonomous Products*, UL 4600, Underwriters Laboratories, Inc., Northbrook, IL, 2019.
21. *System Safety*, MIL-STD-882E, U.S. Department of Defense, Washington, DC, 2012.
22. *System Safety Analysis and Assessment for Part 23 Airplanes*, AC 23.1309-1E, U.S. Department of Transportation Federal Aviation Administration, Washington, DC, 2011.
23. T. Gunnarson, "Aircraft Type Certification Considerations: Urban Air Mobility," presented at the 5th Annual Transformative Vertical Flight Workshop, San Francisco, CA, Jan. 19, 2018.
24. T.K. Ferrell and U.D. Ferrell, "RTCA DO-178B/EUROCAE ED-12B," in *Digital Avionics Handbook*, 3rd ed., C.R. Spitzer, U. Ferrell, and T. Ferrell, Eds. Boca Raton: CRC Press, 2015, pp. 12-1—12-12.
25. Vertical Flight Society, "Vertical Flight Society Reports More than 200 eVTOL Aircraft Now in Development," *Press Release*, Fairfax, VA, Sept. 6, 2019.
26. N. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*. Cambridge, MA. USA: MIT Press, 2011.

8 Acknowledgements

This report describes work supported by NASA under award No. 80LARC17C0004, awarded to the National Institute of Aerospace.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Aeronautics and Space Administration or the National Institute of Aerospace.

The authors wish to thank the FAA for their input during several discussions over the course of this work. The authors also thank Ken Goodrich and Michael Patterson at NASA for contributions to particular hazard and UAM environment discussions.

9 Appendix

Tables referenced or excerpted throughout the document are provided here. Section 9.1 provides further presentation of the function decomposition, and section 9.2 presents additional detail for the FHA table as it appears during the analysis process.

9.1 Living Function Decomposition Table

This table holds the working function decomposition. This decomposition is not presented as a final artifact; rather, it shows this entity as a living source for the FHA, subject to elaboration and update as the FHA and later activities proceed. Note we include here the candidate fifth core function that we described in section 4.1.2.1. We show it here to indicate an option we trialed and considered before discarding.

1	**Aviate**			
1.1		Control flight		
1.1.1			Control flight path	
1.1.1.1				Control altitude
1.1.1.2				Control attitude
1.1.1.3				Control velocity
1.1.1.4				Control propulsion
1.1.1.5				Manage stability
1.1.2			Control air-to-ground transition	
1.1.2.1				Control engine start and stop
1.1.2.2				Control takeoff to hover, hover to land
1.1.2.3				Stow and deploy landing gear
1.1.2.4				Control emergency land
1.1.2.5				Manage stability
1.1.3			Convey system state	
1.2		Control ground movement		
1.2.1			Control taxi	…
1.3		Control subsystems		
1.3.1			Control powerplant	
1.3.2			Monitor vehicle health	
1.3.2.1				Maintain/protect structural integrity
1.3.2.2				Monitor battery health
2	**Navigate**			
2.1		Plan path		
2.1.1			Defer to pre-loaded plan	
2.1.2			Avoid known obstacles	
2.1.3			Avoid known weather	
2.1.4			Choose low-noise path	
2.1.5			Dynamically replan	
2.2		Estimate position and orientation	…	
2.3		Monitor path	Detect dynamic obstacles	
			Detect dynamic weather	

3	**Communicate**			
3.1		Communicate internally		
3.1.1			Support crew comms with each other	
				Support via voice
				Support via data
3.1.2			Support crew comms with passengers	
				Support via voice
				(Support via data?)
3.2		Communicate externally		
3.2.1			Support comms with ATC	
				Support via voice
				Support via data
3.2.2			Support comms with proximate traffic	
				Support via voice
				Support via data
3.2.3			Support comms with others (airline, vertiports)	
				Support via voice
				Support via data
4	**Transport**			
4.1		Persons		
4.1.1			Support passenger ingress/egress	
4.1.2			Support passenger conveyance	
4.1.2.1				Provide passenger seating
4.1.2.2				Restrain passengers in flight
4.1.2.3				Provide HVAC
4.2		Cargo		
4.2.1			Support cargo loading/unloading	
4.2.2			Support cargo conveyance	
4.2.2.1				Provide cargo hold
4.2.2.2				Provide cargo restraints
[5]	**Manage digital systems**	*[Candidate cross-cutting function eventually discarded]*		
[5.1]		Receive and monitor		
			Accept pilot input	
				Accept manual, voice
			Monitor aircraft systems	
				Monitor structures
				Monitor digital state
			Monitor environment	
				Monitor obstacles, weather, ATC, traffic
[5.2]		Store and compute		
			Provide data storage appropriate to decomposed core functions	
			Provide computing infrastructure appropriate to decomposed core functions	
[5.3]		Produce and provide		
			Generate commands from pilot input	
			Provide system state to pilot	
			Provide environment state to pilot	

9.2 Living FHA Table

This table holds the working FHA. This FHA is not presented as a final artifact; rather, it shows this entity as a living repository for hazard identification and management decisions, subject to elaboration and update as the FHA and later activities proceed. Note that the examples narrated within the body of this report, as well as the function decomposition, have been matured further than the contents of the working table, and themselves represent additional points in a process path. Conducting the assessment creates feedback loops and running them all to ground builds confidence in comprehensive assessment.

Color coding:							
	used in narration						
	surfaced in process						
	cascade groups for further analysis						
	draft for unpacking						
no fill	*otherwise latent draft*						
UID	**Function**	**Flight Phase**	**Hazard**	**Operational Consequence**	**Severity Classification**	**Cause(s)**	**Mitigation(s)**
1	**Aviate**						
1.1	Control flight						
1.1.1	Control flight path						
1.1.1.1	Control altitude						
1.1.1.1.a		lift to hover	Uncommanded ascent	Collision with other traffic or structures	Hazardous	Internal failure of control causes runaway ascent as soon controller has effect	Horizontal control (may have limited effect); disable lift (may result in impact with terrain).
1.1.1.1.b			Failure to ascend	Flight stays on the ground	No safety effect		
1.1.1.1.c			Uncommanded descent	Hard landing, impact with terrain	Minor to catastrophic	Failure of controller or propulsion systems	Land and troubleshoot
1.1.1.1.d			Small deviation from commanded ascent	Probably none	No safety effect	Transient failure of controller	
1.1.1.1.e		transition to forward	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Internal failure of control causes loss of function as soon as transition-related control regime has effect	Not sure, probably depends on the transition mechanism

1.1.1.1.f			Uncommanded ascent / descent	Collision with other traffic or structures	Hazardous or catastrophic	Failure of controller or propulsion systems	Horizontal control (may have limited effect); disable lift (may result in impact with terrain).
1.1.1.1.g			Small deviation from commanded ascent/descent	"Turbulence"	No safety effect	Transient failure of controller	
1.1.1.1.h		climb to enroute	Failure to ascend	Flight cannot continue	Minor	Failure of controller or propulsion systems	Land and troubleshoot
1.1.1.1.i			Failure to descend	No ability to terminate flight if needed	Hazardous	Undetected failure of controller	Troubleshoot once detected
1.1.1.1.j			Uncommanded ascent/descent	Collision with structures, terrain, or other traffic	Hazardous	Failure of controller or propulsion systems	Horizontal control, reconfigure/restart controller, or land and troubleshoot
			Small deviation from commanded ascent/descent	Probably none	No safety effect	Transient failure of controller	
		enroute	Failure to ascend/descend	Unable to change FL	Major	Failure of controller	Declare emergency, troubleshoot
			Uncommanded descent	Collision with terrain	Hazardous	Failure of controller or propulsion systems	Troubleshoot, land (if possible)
			Uncommanded ascent	Collision with other traffic or structures	Major	Failure of controller or propulsion systems	Horizontal control, reconfigure/restart controller
			Small deviation from commanded ascent/descent	Probably none	No safety effect	Transient failure of controller	
		avoidance	Failure to ascend/descend	Collision with structures, terrain, or other traffic; increased workload / distraction	Hazardous	Failure of controller or propulsion systems	Horizontal control
			Uncommanded ascent / descent	Collision with structures, terrain, or other traffic; increased workload / distraction	Hazardous	Failure of controller or propulsion systems	Horizontal control
			Small deviation from commanded ascent/descent	Probably none	No safety effect	Transient failure of controller	

		descent	Failure to ascend	No ability to go around if needed	Major	Undetected failure of controller	Troubleshoot
			Failure to descend	Cannot land	Hazardous	Failure of controller or propulsion systems	Climb and troubleshoot
			Uncommanded ascent / descent	Collision with structures, terrain, or other traffic; increased workload / distraction	Hazardous or catastrophic	Failure of controller or propulsion systems	Horizontal control (may have limited effect)
			Small deviation from commanded ascent/descent	Probably none	No safety effect	Transient failure of controller	
		transition to hover	Uncommanded ascent / descent, failure to ascend/descend, incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Internal failure of control causes loss of function as soon as transition-related control regime has effect	Not sure, probably depends on the transition mechanism
			Small deviation from commanded ascent/descent	"Turbulence"	No safety effect	Transient failure of controller	
		set down	Failure to ascend	No ability to go around if needed	Catastrophic	Undetected failure of controller	Land anyway and troubleshoot
			Failure to descend	Cannot land	Hazardous	Failure of controller or propulsion systems	Climb and troubleshoot
			Uncommanded ascent	Collision with other traffic or structures	Hazardous	Failure of controller	Horizontal control (may have limited effect); disable lift (may result in impact with terrain).
			Uncommanded descent	Hard landing	Catastrophic	Failure of controller or propulsion systems	Aircraft ability to absorb landing energy
			Small deviation from commanded descent	Possible hard landing	Major	Transient failure	Aircraft ability to absorb landing energy
			Failure to descend	Inability to land normally until problem is resolved	Minor	Failure of controller	Reconfigure / retry (possibly after ascending to safe altitude)
1.1.1.2	Control attitude						

1.1.1.2.a		lift to hover	Small deviation from commanded attitude	"Turbulence", possible contact with pad / structure.	Major?	Transient failure of controller	Land and troubleshoot
			Large deviation from commanded attitude	Contact with pad, structure. Loss of stability.	Hazardous	Failure of controller or propulsion systems	Land and troubleshoot (if possible)
		transiition to forward	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Internal failure of control causes loss of function as soon as transition-related control regime has effect	Not sure, probably depends on the transition mechanism
			Small deviation from commanded attitude	"Turbulence"	No safety effect	Transient failure of controller	Land or climb then troubleshoot
			Large deviation from commanded attitude	Collision with other traffic or structures	Hazardous	Failure of controller or propulsion systems	Land or climb then troubleshoot
		climb to enroute	Loss of stability	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Failure of controller or propulsion systems	Reconfigure/restart controller
			Small deviation from commanded attitude	"Turbulence"	No safety effect	Transient failure of controller	Continue climb and troubleshoot
			Large deviation from commanded attitude	Collision with other traffic or structures	Hazardous	Failure of controller or propulsion systems	Mostly other aircraft avoiding yours
		enroute	Loss of stability	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Failure of controller or propulsion systems	Reconfigure/restart controller
			Small deviation from commanded attitude	"Turbulence"	No safety effect	Transient failure of controller	Troubleshoot
			Large deviation from commanded attitude	Collision with other traffic or structures	Hazardous	Failure of controller or propulsion systems	Mostly other aircraft avoiding yours

		avoidance	Loss of stability	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain or other aircraft	Catastrophic	Failure of controller or propulsion systems	Reconfigure/restart controller
			Small deviation from commanded attitude	"Turbulence"	No safety effect	Transient failure of controller	Use vertical control for avoidance; troubleshoot
			Large deviation from commanded attitude	Collision with other traffic or structures	Hazardous	Failure of controller or propulsion systems	Mostly other aircraft avoiding yours
		descent	Loss of stability	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain or other aircraft	Catastrophic	Failure of controller or propulsion systems	Reconfigure/restart controller
			Small deviation from commanded attitude	"Turbulence"	No safety effect	Transient failure of controller	Abort descent; troubleshoot
			Large deviation from commanded attitude	Collision with other traffic or structures	Catastrophic	Failure of controller or propulsion systems	Mostly other aircraft avoiding yours
		transition to hover	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Internal failure of control causes loss of function as soon as transition-related control regime has effect	Not sure, probably depends on the transition mechanism
			Small deviation from commanded attitude	"Turbulence"	No safety effect	Transient failure of controller	Land or climb then troubleshoot
			Large deviation from commanded attitude	Collision with other traffic or structures	Hazardous or catastrophic	Failure of controller or propulsion systems	Land or climb then troubleshoot
		set down	Small deviation from commanded attitude	"Turbulence", possible contact with pad / structure.	Minor?	Transient failure of controller	Land anyway

			Large deviation from commanded attitude	Contact with pad, structure. Loss of stability.	Hazardous	Failure of controller or propulsion systems	Land anyway
1.1.1.3	Control velocity						
1.1.1.3.a		lift to hover	Small deviation from commanded velocity	"Turbulence", possible contact with pad / structure.	Major?	Transient failure of controller	Land anyway
			Large deviation from commanded velocity	Contact with pad, structure.	Hazardous	Failure of controller	Land and troubleshoot
		transition to forward	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Internal failure of control causes loss of function as soon as transition-related control regime has effect	Not sure, probably depends on the transition mechanism
			Small deviation from commanded velocity	"Turbulence"	No safety effect	Transient failure of controller	Land or climb then troubleshoot
			Large deviation from commanded velocity	Collision with other traffic or structures	Hazardous	Failure of controller or propulsion systems	Land or climb then troubleshoot
		climb to enroute	Small deviation from commanded velocity	"Turbulence"	No safety effect	Transient failure of controller	Continue climb and troubleshoot
			Large deviation from commanded velocity	Collision with other traffic or structures	Hazardous	Failure of controller or propulsion systems	Mostly other aircraft avoiding yours
		enroute	Small deviation from commanded velocity	"Turbulence"	No safety effect	Transient failure of controller	Troubleshoot
			Large deviation from commanded velocity	Collision with other traffic or structures	Hazardous	Failure of controller or propulsion systems	Use vertical control for avoidance; troubleshoot; other aircraft avoiding yours
		avoidance	Small deviation from commanded velocity	"Turbulence"	No safety effect	Transient failure of controller	Use vertical control for avoidance; troubleshoot

			Large deviation from commanded velocity	Collision with other traffic or structures	Hazardous or catastrophic	Failure of controller or propulsion systems	Mostly other aircraft avoiding yours
		descent	Small deviation from commanded velocity	"Turbulence"	No safety effect	Transient failure of controller	Climb and troubleshoot
			Large deviation from commanded velocity	Collision with other traffic or structures	Hazardous or catastrophic	Failure of controller or propulsion systems	Climb, troubleshoot, other aircraft try to avoid you
		transition to hover	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Internal failure of control causes loss of function as soon as transition-related control regime has effect	Not sure, probably depends on the transition mechanism
			Small deviation from commanded velocity	"Turbulence"	No safety effect	Transient failure of controller	Land or climb then troubleshoot
			Large deviation from commanded velocity	Collision with other traffic or structures	Hazardous or catastrophic	Failure of controller or propulsion systems	Land or climb then troubleshoot
		set down	Small deviation from commanded velocity	Hard landing	Major?	Transient failure of controller	Land and troubleshoot
			Large deviation from commanded velocity	Collision with other traffic or structures	Hazardous or catastrophic	Failure of controller or propulsion systems	Land and troubleshoot
			Large deviation from commanded velocity	Rollover on landing	Hazardous or catastrophic	Failure of controller or propulsion systems	Land and troubleshoot
1.1.1.4	Control propulsion						
1.1.1.4.a		lift to hover	Failure to control altitude, attitude, or velocity	Contact with pad, structure, or other traffic	Catastrophic	Failure of propulsion	Land and troubleshoot (if possible)
cascade group		transition to forward	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Land and troubleshoot (if possible)
			Failure to control two or more of altitude, attitude, or velocity	Collision with traffic or terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Land and troubleshoot (if possible)

		climb to enroute	Inability to climb	Collision with nearby structures, other traffic, failure to achieve enough altitude to transition back to hover and land at the takeoff point	Hazardous	Loss of power causes loss of ability to control altitude	Land and troubleshoot (if possible)
			Failure to control two or more of altitude, attitude, or velocity	Collision with traffic or terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Land and troubleshoot (if possible)
		enroute	Failure to control two or more of altitude, attitude, or velocity	Collision with traffic or terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Troubleshoot
		avoidance	Failure to control two or more of altitude, attitude, or velocity	Collision with traffic or terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Troubleshoot, let the other aircraft try to avoid you (if it's an AC not an obstacle)
		descent	Failure to control two or more of altitude, attitude, or velocity	Collision with traffic or terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Land and troubleshoot (if possible)
cascade group		transition to hover	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Land and troubleshoot (if possible)
			Failure to control two or more of altitude, attitude, or velocity	Collision with traffic or terrain	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Land and troubleshoot (if possible)
		set down	Failure to control two or more of altitude, attitude, or velocity	Contact with pad, structure, or other traffic	Catastrophic	Loss of power causes insufficient and/or asymmetric lift	Land and troubleshoot (if possible)

1.1.1.4	Control propulsion	transition	Loss of propulsion	Impacts to aircraft controllability, including loss of control	Catastrophic	HW or SW fault	Decoupling of propulsion and control; architectural redundancy; glide and other off-nominal landing options
1.1.1.4.1	Compute motor/rotor commands	transition	No result or incorrect result produced	No command or wrong command available to send	Hazardous	HW or SW fault	Proper HW and SW fault management practices
1.1.1.4.2	Command motors/rotors	transition	No command or wrong command sent	No actuation or wrong actuation	Hazardous	HW or SW fault	Proper HW and SW fault management practices
1.1.1.4.3	Actuate motors/rotors	transition	No actuation or wrong actuation	Impacts to propulsion profile, including loss of effective propulsion	Catastrophic	HW or SW fault	Proper HW and SW fault management practices
1.1.1.5	Manage stability						
1.1.1.5.a		lift to hover	Aircraft is massively and unexpectedly out of balance or trim	Contact with pad, structure, or other traffic, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Land and troubleshoot (if possible)
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible contact with pad / structure.	Minor?	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot
		transition to forward	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Failure of automatic system leads to sudden loss of trim on mode shift	Land and troubleshoot (if possible)
			Aircraft is massively and unexpectedly out of balance or trim	Contact with pad, structure, or other traffic, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Land and troubleshoot (if possible)
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible minor passenger injuries	Minor	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot

		climb to enroute	Aircraft is massively and unexpectedly out of balance or trim	Collision with other traffic or structures, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Troubleshoot
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible minor passenger injuries	Minor	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot
		enroute	Aircraft is massively and unexpectedly out of balance or trim	Collision with other traffic or structures, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Troubleshoot
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible minor passenger injuries	Minor	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot
		avoidance	Aircraft is massively and unexpectedly out of balance or trim	Collision with other traffic or structures, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Mostly other aircraft avoiding yours
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible minor passenger injuries	Minor	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot
		descent	Aircraft is massively and unexpectedly out of balance or trim	Collision with other traffic or structures, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Troubleshoot
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible minor passenger injuries	Minor	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot
		transition to hover	Incomplete/wrong transition	Possible unrecoverable loss of ability to control flight path, resulting in impact with terrain	Catastrophic	Failure of automatic system leads to sudden loss of trim on mode shift	Land and troubleshoot (if possible)

			Aircraft is massively and unexpectedly out of balance or trim	Contact with pad, structure, or other traffic, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Land and troubleshoot (if possible)
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible minor passenger injuries	Minor	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot
		set down	Aircraft is massively and unexpectedly out of balance or trim	Contact with pad, structure, or other traffic, loss of stability	Hazardous	Failure of automatic system, failure of pre-flight planning, etc.	Land and troubleshoot (if possible)
			Aircraft is unexpectedly out of balance or trim to a moderate degree	"Turbulence", possible contact with pad / structure.	Minor?	Failure of automatic system, failure of pre-flight planning, etc., shifting of cargo or passengers	Compensate with attitude control, troubleshoot
1.2	Control ground movement						
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	Design or degradation faults in propulsion system	Power source redundancy
1.3.2	Monitor vehicle health						
2	**Navigate**						
2.1	Plan path						
2.1.1	Defer to pre-loaded plan						
2.1.2	Avoid known obstacles						
2.1.3	Avoid known weather						
2.1.4	Choose low-noise path						

2.1.5	Dynamically replan						
[drafting]		approach	Operator experiencing loss/degradation of situational awareness	Execute controlled flight onto wrong landing pad; possible harm to persons or property in the landing field [unpack; this is not direct consequence]	Catastrophic	[ADD INTERVENING CAUSE re: display of sufficient and current info to allow SA] Failure of dynamic replanning function deriving from, e.g., increased cognitive load from something like weather diversion, novel vertiport, glare, etc.	Pax containment at vertiports; Property/object containment at vertiports; Positive confirmation during descent with e.g., dispatch re: assigned landing pad
2.1.5a		approach	Insufficient path plan	Misidentified or underspecified waypoints and/or destination conveyed to operator; increased operator workload; impacts to situational awareness	Hazardous	HW/SW/data	Waypoint and destination confirmation; human factors adjustments to nav interface; HW/SW/data fault management
2.2	Estimate position and orientation						
2.3	Monitor path						
2.3.1	Detect dynamic obstacles						
2.3.2	Detect dynamic weather						
3	**Communicate**						
3.1	Communicate internally						
3.2	Communicate externally	enroute	Loss of external comms	Stale information; increased operator workload; impacts to situational awareness	Hazardous	HW/SW/data	Data timeouts; Comms redundancy; procedural guidance for lost comms scenarios; HW/SW/data fault management

4	**Transport**						
4.1	Transport persons						
[drafting]		enroute	Aircraft center of gravity is out-of-bounds	Degraded stability, degraded handling qualities, degraded safety margins, increased pilot workload	Hazardous	Failure of passenger restraints, passengers move about, disrupting vehicle balance	Procedural passenger management,
4.1.2.2	Restrain passengers in flight	enroute	Unrestrained passenger(s)	Aircraft center of gravity out of bounds; reduced flight stability; passenger injury	Hazardous	design or degradation faults in belt mechanism; pax not adhering to procedural management	Procedural passenger management; design to accommodate balance shifts
4.1.2.2.1	Apply mechanical restraint	enroute	Unrestrained passenger(s)	Aircraft center of gravity out of bounds; reduced flight stability; passenger injury	Hazardous	design or degradation faults in belt mechanism; tampering	Management of design and degradation faults in restraint components; procedural passenger management; design to accommodate balance shifts
4.1.2.2.2	Apply procedural restraint	enroute	Unrestrained passenger(s)	Aircraft center of gravity out of bounds; reduced flight stability; passenger injury	Hazardous	pax not adhering to procedural management	Redesign procedural passenger management; design to accommodate balance shifts
4.2	Transport cargo						
New/Unhomed Functions							
	Maintain separation		loss of separation	Imposition of response actions on ATC, proximate traffic, NMAC/MAC	Hazardous, Catastrophic	Multiple, e.g., decision error or command error due to loss of situational awareness	Strategic and tactical conflict management capabilities, including procedural separation and DAA systems

	Maintain situational awareness		Loss of situational awareness	Increased operator workload; decision error; command error	Hazardous	Multiple, e.g. loss of (external) comms; failure of dynamic replanning	TBD in accordance with causes; this hazard arises many ways
	Contain persons at vertiport landing pad	n/a (vertiport vs. aircraft)	Failure to contain persons	Injury or death from landing aircraft	Hazardous or catastrophic		[Speculate based on how containment failed]