A METHODOLOGY FOR THE DESIGN AND OPERATIONAL SAFETY ASSESSMENT OF UNMANNED AERIAL SYSTEMS

A Dissertation
Presented to
The Academic Faculty

By

Andrew Kendall

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Daniel Guggenheim School of Aerospace Engineering

Georgia Institute of Technology

December 2022

© Andrew Kendall 2022

A METHODOLOGY FOR THE DESIGN AND OPERATIONAL SAFETY ASSESSMENT OF UNMANNED AERIAL SYSTEMS

Thesis committee:

Dr. John-Paul Clarke
Daniel Guggenheim School of Aerospace Engineering
Georgia Institute of Technology

Dr. Brian J German
Daniel Guggenheim School of Aerospace Engineering
Georgia Institute of Technology

Dr. David Goldsman
H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology

Dr. Karen M Feigh
Daniel Guggenheim School of Aerospace Engineering
Georgia Institute of Technology

Dr. Husni R Idris
Aviation Systems Division
NASA Ames Research Center

Date approved: December 12, 2022

ACKNOWLEDGMENTS

I would like to thank the members of my thesis committee for their help in preparation of this work. As my USRA NAMS internship mentor, Dr. Husni Idris introduced me to the pressing challenges of Remotely Piloted Aerial Systems and gave me the context needed to further explore the problem. My advisor, Dr. John-Paul Clarke, has my gratitude for inviting me into research and giving me the opportunities to explore the topics that interest me.

Special thanks are due to my family and friends who supported me through this process, my bandmates who gave me a creative outlet outside of my studies, and my colleagues who helped me to develop and solidify my ideas.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ACRONYMS

AH Alert Height

ATC Air Traffic Control

ATC Air Traffic Control

C2 Command and Control

CDA Continuous Descent Arrival/Approach

CE Cross-Entropy

DCT Discrete Cosine Transform

DH Decision Height

DOP Dilution of Precision

EFVS Enhanced Flight Vision System

ETOPS Extended-range Twin-engine Operations Performance Standards

FMS Flight Management System

FO Fail Operational

FP Fail Passive

GBAS Ground Based Augmentation System

GCS Ground Control Station

GLS GBAS Landing System

GNSS Global Navigation Satellite System

GPS Global Positioning System

HUD Head Up Display

ICAO International Civil Aviation Organization

IFR Instrument Flight Rules

ILS Instrument Landing System

IMC Instrument Meteorological Conditions

IMU Inertial Measurement Unit

IS Importance Sampling

K-L Karhunen-Loeve

LNAV Lateral Navigation

MA Missed Approach

MLS Microwave Landing System

MLW Maximum Landing Weight

NAS National Airspace System

OODA Observe-Orient-Decide-Act

PDF Probability Density Function

RNAV Area Navigation

RNP Required Navigation Performance

RP Remote Pilot

RPAS Remotely Piloted Aerial System

RVR Runway Visual Range

SATCOM Satellite Communication

SBAS Satellite Based Augmentation System

SDE Stochastic Differential Equation

SHS Stochastic Hybrid System

SVGS Synthetic Vision Guidance System

TASAT Tool for Analysis of Separation and Throughput

TCAS Traffic Collision Avoidance System

TLS Target Level of Safety

UAS Unmanned Aerial System

UAV Unmanned Aerial Vehicle

VMC Visual Meteorological Conditions

VNAV Vertical Navigation

WAAS Wide Area Augmentation System

SUMMARY

Efforts are underway to introduce Unmanned Aerial Systems (UAS) into routine cargo operations within the National Airspace System (NAS). Such systems have the potential to increase transport system flexibility by mitigating crew scheduling constraints and extending operations to remote locations. It is expected that any large UAS operating in the transport category must comply with Federal Aviation Regulations to achieve airworthiness certification for routine operations within the NAS. Regulations on the safety of equipment, systems, and installations require all failure conditions due to malfunctions, environmental events, and inadequate corrective action to be mitigated and shown to be extremely improbable. These system safety requirements are particularly relevant for a UAS as the ability of a Remote Pilot (RP) to detect and respond to risks is dependent on a Command and Control (C2) link. Failure conditions associated with the C2 link system require autonomy onboard the aircraft to supplement the RP in order to mitigate risk. A method for assessing the performance required from automation when the RP cannot adequately mitigate risks is needed to allow routine UAS operations.

The problem of ensuring autonomous UAS safety requirements is addressed in this thesis through the development of a safety assessment methodology that can be applied during both system design and online operations. The contributions are as follows:

- Safety Regulations are formulated as a chance-constraint satisfaction problem, requiring safety on the order of 1 accident per billion operations. Rare event estimation techniques based on Importance Sampling are proposed to assess safety subject to various sources of uncertainty.

- Failure conditions can be due to both discrete events, such as system failures, and continuous state uncertainties, such as navigation errors and turbulence. A stochastic hybrid system model is proposed to handle the coupling between discrete and continuous states and estimate the distribution of aircraft trajectories that may result from

a given set of system parameters, operational conditions, and decision parameters.

- The final approach and landing phase of flight serves as a use case for the methodology. The safety assessment is applied to determine system design parameters required to passively mitigate risks. The methodology is extended to active risk mitigation during operations, where online safety assessments using updated observations are used to ensure decision options always exist that will satisfy safety requirements.

CHAPTER 1
INTRODUCTION AND BACKGROUND

1.1 Introduction

Aviation is moving in the direction of increasing autonomy, reducing reliance on humans in the cockpit and shifting tasks to computers. Since the early days of flight, automation has been introduced to augment pilot senses in low visibility conditions, eliminate distinct human roles such as flight engineer and navigator, and automatically fly procedures and perform landings with minimal human intervention. A pressing goal in aircraft automation is to introduce routine cargo transport operations using Unmanned Aerial System (UAS) into the National Airspace System (NAS) [1]. Concepts for such unmanned operations range from a single Remote Pilot (RP) operating a single Unmanned Aerial Vehicle (UAV) via a Ground Control Station (GCS) and Command and Control (C2) link, to M:N operations in which a small team of RPs oversee a larger fleet of UAVs.

Remotely Piloted Aerial Systems (RPASs) have been in regular use in military operations for decades, however, large RPAS have not yet been shown to meet the standards required for routine civilian operations in the NAS. Several potential difficulties exist due to the latency and reliability of the C2 link used to facilitate communication between the RP and the UAV, and the inherent change in situational awareness due to the RP not being physically situated on the UAV [2]. C2 links are used to uplink commands such as flight controls, chosen procedures, and configuration changes from the GCS to the RPA and downlink avionics data, camera feeds, and alerts back down to the GCS. C2 links can have terrestrial radio line-of-sight communication latencies as low as 10s of milliseconds or Satellite Communication (SATCOM) round trip latencies up to nearly 2 seconds [3]. Availability of C2 links depends on the presence of buildings, infrastructure, and terrain

between the UAV and the ground antenna.

The impact of C2 latencies and lost links on terminal operations is more significant compared to enroute operations. Approach and landing operations present significant challenges due to the precision required to touchdown safely, sensitivity to faults and weather events, and the decrease in number of alternative options available as the ground gets closer. Although taking up only several minutes of flight time, approach and landing accidents consistently contribute to between 40% and 50% of all manned aviation accidents [4]. Significant effort has been made to increase the safety of approach and landing operations via operational procedures, instrument aids, and automatic control.

Landings in Instrument Meteorological Conditions (IMC) have many parallels to RPAS landings, as the pilot has limited access to direct visual cues and must instead rely on onboard instruments to maintain situational awareness. Sensor noise, bias, errors, and failures introduce navigation system errors, possibly resulting in incorrect readings on the instruments. Due to these reasons, instruments can only be certified for a narrow set of conditions. The widely implemented Instrument Landing System (ILS) has been certified to perform in a range of low visibility conditions but it does not address the specific complications associated with RPA landings. The methodology used to certify ILS is used as guidance to generate safety requirements formulation for certifiably safe RPA approach and landing procedures. ILS approaches have a minimum Decision Height (DH), the lowest height to which the aircraft may descend without direct visual contact with the runway lights [5]. If no visual contact is made, position uncertainty is too large to proceed safely, and a Missed Approach (MA) must be executed. The premise behind the decision to continue landing or execute a MA is based around maintaining an option that satisfies a Target Level of Safety (TLS). Concepts used to certify Required Navigation Performance (RNP) procedures are also considered. Rather than designing a one size fits all procedure for any aircraft, RNP only allows aircraft equipped with navigation and control systems that meet minimum performance requirements to fly particular procedures [6].

A methodology is proposed demonstrating the safety of unmanned procedures, with emphasis on approach and landing procedures, by directly estimating the probability of accidents. The framework is capable of evaluating the set of equipment performance requirements that will satisfy the TLS for a given approach and landing procedure as well as updating the safe decision minima as new information is gathered during online operation. Several components contribute to the framework including an Autonomous Agent model, a Stochastic Hybrid System environment model, and a Rare Event Estimation methodology for quantifying risk.

1.2 Approach and Landing Automation

Approach and landing operations have long been the subject of efforts to reduce reliance on human pilots. The precursor to the modern ILS stems from the need to conduct military operations during dense foggy conditions in Britain during World War II [7]. The SCS-51 landing system was used to perform the first completely blind landing in January 1945. The system used pairs of overlapping radio beams to provide pilots with vertical glide slope error and lateral localizer error, as well as marker beacons indicating progress along the approach. ICAO adopted this system as the international standard for ILS in 1948, though it could only provide accurate flight director guidance down to 200 ft above the runway. Further research during the 1950's focused on improving signal accuracy close to the ground, implementing autopilot control laws to account for increasing sensitivity close to the antennas, the transition from ILS glideslope input to pitch and radio altimeter inputs, and control of the landing flare, drift kick, and rollout. High intensity approach lights in standardized patterns were also developed during this time. The ILS system entered service with military aircraft in 1961, though civilian operations required a higher degree of reliability. In order to meet a TLS of 1 accident in 10 million automatic landings, triple redundant systems and risk assessment methods were developed. Several categories of ILS were established, providing a tradeoff between system requirements and the mini-

mum allowable operational conditions. These minimums are the DH, the minimum height the pilots may descend to without visual contact with the runway or approach lights, and Runway Visual Range (RVR), the minimum visibility required for viewing approach and runway lights. The requirements associated with each ILS category are listed in Table 1.1. These were accepted by ICAO in 1965 and the first certified ILS CAT I passenger landings were performed later that year. Development of the higher categories progressed over the next decade until the first ILS CAT III without a Decision Height was first certified in 1979. It is important to note that ILS certification doesn't just depend on the aircraft and ground based equipment, but also on the ability of pilots to maintain situational awareness and quickly react to failures and other events. Pilots must go through advanced training to be certified for ILS approaches. ILS also requires a substantial ground footprint, consisting of an obstacle clearance zone to mitigate collision risks and an ILS critical area that must be kept clear of aircraft, vehicles, and obstructions to prevent signal disruptions [5]. The vertical components of a typical ILS approach procedure are depicted in Figure 1.1.

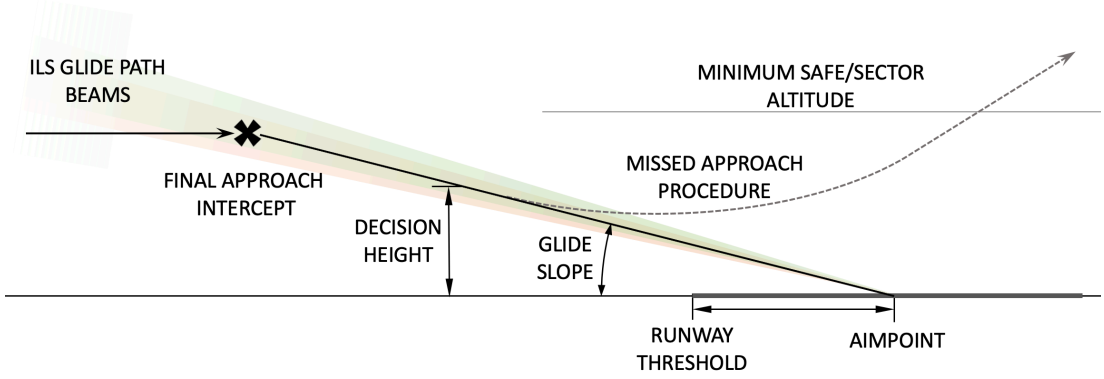

Figure 1.1: ILS Procedure Diagram

CAT III approaches were formerly sub-categorized into CAT IIIa, IIIb, and IIIc, corresponding to decreasing minima. Newer methodologies adopted by the FAA sub-categorize CAT III by redundancy of the autoland and autorollout systems on the aircraft [8]. Several new concepts are introduced. Alert Height (AH) is the height above which a missed approach must be executed if a system failure is detected, while the approach and landing

Table 1.1: ILS Category Requirements, ✓ indicates presence of at least 1 system, ✓ ✓ indicates presence of 2 systems, * indicates independent systems [5, 8]

Requirement	CAT I	CAT II	CAT III
Minimum Decision Height (ft)	200	100	0
Minimum Runway Visual Range (ft)	1800	1200	300
Navigation Receivers	✓✓	✓✓*	✓✓*
Flight Management System	✓✓	✓✓*	✓✓*
Failure Annunciators	✓	✓	✓✓*
Pilot Displays	✓	✓	✓
Radar Altimeter		✓	✓
Flight Director or Autopilot		✓	✓
Autothrottle		✓	✓
Rain Removal Equipment		✓	✓
Missed Approach Automation			✓
Autoland			✓
Rollout Control			✓

may proceed if the failure occurs below the AH. Fail Operational (FO) systems have the required redundancy and monitoring to continue performing without interruption if one sub system fails. Fail Passive (FP) systems are unable to safely continue when a single failure occurs and control is handed back over to the pilot before large trajectory deviations can occur. The relationship between CAT III minima and equipment redundancy is depicted in Table 1.2.

Table 1.2: ILS CAT III Requirements [8]

Landing System Type	Rollout System Type	Minima Height (ft)	RVR (ft)
FP	None	50 DH	600
FO	None	50 AH	600
FP	FP or FO	50 DH	600
FO	FP	50 AH	400
FO	FO	50 AH	300

Further developments to autonomous landing systems attempt to replace the ILS equipment and information sources with alternative technologies that meet the same requirements. The Microwave Landing System (MLS) was developed to provide guidance information using a microwave antenna with a smaller footprint, larger capture angle, and more

channel options. The introduction of Global Navigation Satellite System (GNSS) such as the Global Positioning System (GPS), has allowed precision navigation to be available in many previously unserved areas. Satellite Based Augmentation System (SBAS) such as the Wide Area Augmentation System (WAAS) increase the performance of standard GPS service by using a number of ground reference stations across the country to compute error corrections which are then transmitted to users by satellite. Recent research has focused on using Ground Based Augmentation System (GBAS) to improve GPS navigation performance to the equivalent level required by ILS. Ground stations at surveyed locations near the runway provide pseudorange corrections and signal monitoring to properly equipped aircraft. A GPS based procedure is typically parameterized as a RNP procedure, with accuracy, integrity, continuity, and availability requirements [9, 6]. Accuracy describes the value that bounds navigation error 95% of the time. Integrity describes the probability that the navigation system produces hazardously misleading information without producing an alert within a defined alert time during an operation. This is often accomplished by estimating a protection level based on GPS constellation geometry and error status, and triggering an alert when the protection level surpasses an alert limit. Continuity describes the rate at which alerts are produced during nominal operations, rendering the system unavailable during use requiring a missed approach, while availability describes the probability all required systems meet their requirements at the initiation of operation. The primary obstacles to implementing GBAS Landing System (GLS) are ensuring signal integrity in the face of satellite faults, ground station faults, and rare ionosphere gradient phenomena. CAT I GLS is already in operation at a number of airports, though CAT II/III GLS is still in development and certification. GAST-D attempts to provide the necessary reliability for CAT II/III GLS using multiple ground stations and monitors to catch errors before they compromise safety. The vertical components of a typical GLS approach procedure are depicted in Figure 1.2.

In 2020, the Garmin Emergency Autoland system was certified on the G3000 integrated

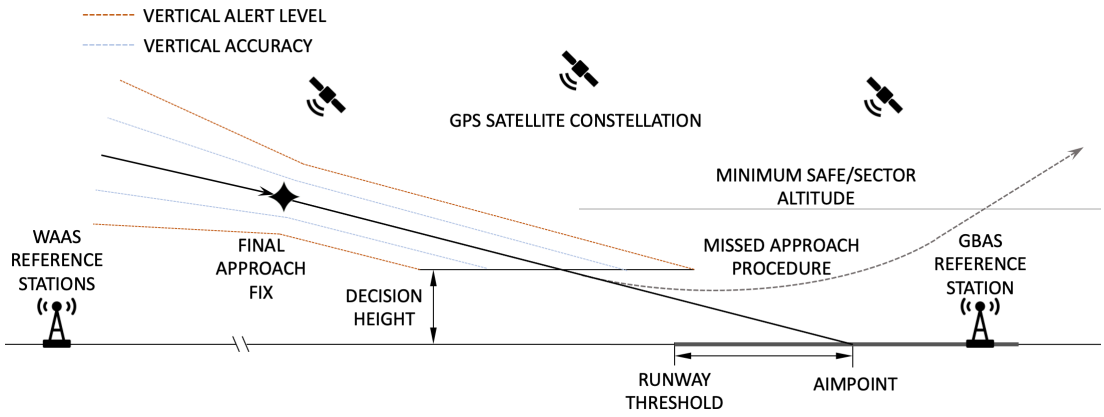

Figure 1.2: GLS Procedure Diagram

flight deck in the Piper M600 [10]. The system detects or is alerted to pilot incapacitation and selects a suitable runway based on range, weather, and availability of a GPS approach with vertical guidance. A route is generated accounting for obstacles, terrain, and weather, and intent is automatically broadcast to Air Traffic Control (ATC). The aircraft initiates the GPS approach, enters a holding pattern at the Final Approach Fix if stabilized approach criteria are not yet met, performs the flare, touchdown, and rollout, and finally shuts the engine off upon coming to a stop on the runway. The system is only certified for emergency situations as it lacks much of the redundancy and monitoring that would be required for routine operation.

Another method for operating in low visibility conditions is augmenting pilot visual information using technologies such as Head Up Displays (HUDs), Enhanced Flight Vision System (EFVS), and Synthetic Vision Guidance System (SVGS). HUDs allow the pilot to fly an instrument procedure while maintaining a constant view out the cockpit window, decreasing reaction time to visual cues. EFVS augments a visual display with an infrared image from a sensor on the nose of the aircraft, allowing a minimum DH of 100 ft [5]. SVGS generates a simulated view of the approach using the estimated state of the aircraft and a database containing terrain, obstacles, and the runway geometry. No augmented vision systems are yet certified for CAT III landings on their own.

True zero visibility CAT III landings require operations not just through approach,

touchdown, and rollout, but also the remainder of taxing and surface operations. Automation of surface operations has been the subject of many new technologies focusing on situational awareness, navigation, and control in low visibility conditions. Systematic lighting such as Rapid Exit Taxiway Indicator Lights are basic visual systems used to assist in low visibility surface operations. Airport Surface Detection Equipment, Model X (ASDE-X) fuses information from sources such as Surface Movement Radar and Automatic Dependent Surveillance — Broadcast (ADS-B) to provide controllers with situational awareness of surface movements [11]. Advanced Surface Movement Guidance and Control System (A-SMGCS) give pilots situational awareness of surface movements and guidance for preventing or resolving conflicts [12]. Despite significant research interest, no technologies have been certified to allow general surface operations in true zero visibility conditions.

Wind shear and microburst events are particularly hazardous during final approaches. These events used to only be detectable through pilot senses, though several technologies in use now provide some degree of autonomous detection and warning [13]. Predictive Wind Shear (PWS) warning systems use aircraft based doppler radar or ground based wind and radar measurements to give pilots advanced notice to perform a missed approach and avoid wind shear events. Reactive Wind Shear (RWS) detection systems use a combination of onboard flight data to detect changes in airspeed associated associated with wind shear events and provide pilots with wind shear escape guidance.

RPAS approach and landing operations have many similarities to conventionally piloted instrument approaches and landings as both heavily depend on systems and equipment to provide reliable information and control the trajectory. The only fundamental difference is an RP never has direct access to visual or haptic information and relies on a C2 link to receive information and transmit commands. To meet safety requirements equivalent to manned operations, additional failure conditions due to the C2 link and altered RP response must be accounted for and mitigated. This will ultimately affect the safe decision minima and introduce requirements on C2 link latency and reliability, autonomous monitoring per-

formance, and RP response.

1.3 Unmanned Aircraft System Automation and Safety

Unmanned and Remotely Piloted Aircraft systems were developed in parallel to low visibility landing capabilities, again starting in the early days of aviation. The first unmanned aerial vehicles developed during WWI were essentially gyro-stabilized aerial torpedoes with simple barometric altitude control [14]. Use of UAVs as gunnery training targets was also of interest. Programs in the 1920's developed simple radio control using discrete commands to adjust heading, altitude, and throttle setting. The pilot-in-the-loop paradigm was introduced in the 1930's with continuous radio controls analogous to manual stick and rudder controls, however, basic flight stabilization was generally performed by an autopilot. Control via direct visual-line-of-sight observation was replaced by radio transmitted avionics signals and eventually supplemented in 1941 by television transmissions with a range up to 30 miles. Beyond-visual-line-of-sight control was also be accomplished using radar tracking. Remote sensing capabilities led to RPAS use in reconnaissance roles. The experimental SD-2 Overseer program of the late 1950's was equipped with many data links and was equipped with a pre-programmed navigation route using a ground based precursor to GPS, which could demonstrate an accuracy of 5 ft at up to 50 miles range [14]. These early RPAS were either recovered through conventional direct visual-line-of-sight landing, recovered via parachute and mid-air helicopter pickup as was the AQM-34, or disposable. Research RPAS such as the F-15 Remotely Piloted Research Vehicle [15] were notoriously demanding for test pilots to remotely control, due to lack of haptic feedback and altered situational awareness. Automatic low-level control is generally necessary to reduce remote pilot workload. Modern military RPAS such as the MQ-1 Predator, MQ-9 Reaper, or RQ-4 Global Hawk are equipped with GPS and the automation necessary to perform landings. Automation also performs basic risk mitigating procedures under lost-link conditions such as entering holding patterns or returning to base.

While military RPAS have been in operation for many years, civilian and commercial RPAS applications are still in their infancy. Safety constitutes a major barrier to routine RPAS operations in the NAS [16]. RPAS can be split into several categories, mainly small vehicles with a weight below 55 lbs and larger vehicles above this weight. Safety analyses have been performed for small RPAS [17], and federal regulations defining the acceptable operational domain and safety requirements exist under 14 CFR Part 107. Large RPAS still have no routine certification process and require Special Airworthiness Certificates or Grants of Exemption to operate [18]. Separation violations and mid-air collisions are major hazards, especially in non-segregated airspace with both manned and unmanned aircraft [19]. The latency and possible failure associated with C2 links reduces the effectiveness of risk mitigation procedures [2]. The Required Link Performance (RLP) concept places minimum performance requirements on the link performance with parameters such as link continuity, integrity, availability and latency [20]. Onboard collision risk mitigation may be performed by Detect and Avoid (DAA) automation, which maintains separation without pilot intervention [21]. On top of the mid-air hazards, RPAS must still maintain safety during final approach and landing, the riskiest flight segments in manned aviation. CAT III landing systems, described in the previous section, could potentially provide the required level of safety, however, installation are rare and would limit the use of RPAS for commercial application. The use of computer vision has been proposed to bridge the decision height gap between CAT I approaches and CAT III landings [22]. A variety of landing control system architectures have been proposed to land safely in the face of disturbances [23], however, navigation system errors still present a hazard that must be accounted for. Methodologies for certification of conventional instrument landing systems are suited towards the analogous problem RPAS final approaches and landings.

1.4 Safety Requirements Formulation

The highest level regulations on safety critical systems for transport category aircraft can be found in FAR § 25.1309. These regulations require all failures conditions and their joint combinations which are not extremely improbable to be mitigated to an acceptable level of safety. AC 25.1309 provides a quantitative interpretation of these requirements. Extremely improbable is defined as 1 catastrophic accident per 1 billion operations or flight hours. These requirements apply to each failure condition and associated mitigation. Safety requirements used in practice have differed from those interpreted by AC 25.1309. The failure condition level safety requirement is less complex to decompose during design as it does away with the higher level risk allocation problems and splits it into many independent sub-problems. However, this requirement does not consider the total accident risk due to any failure condition occurring. If the number of failure conditions increases due to system complexity, the total probability of an accident is also allowed to increase even though each individual failure condition is mitigated according to the requirements. If the outcome regulators wish to prevent is the occurrence of any accident, regardless of failure condition, a different safety requirement must be used. The aircraft Target Level of Safety (TLS) risk metric considers the probability of an accident due to any failure condition, without explicitly defining requirements for individual failure conditions. While this formulation directly considers the total probability of an accident occurring, it introduces the problem of risk allocation.

Both failure condition safety requirements and aircraft TLS requirements can be expressed using a common notation based on Bayesian probability. Let S be the power set of all possible discrete failure conditions and θ be the set of system design parameters and operational decision parameters to be considered. μ contains the set of static and prescribed system parameters. The event acc indicates that the system has experienced an accident during the operational risk exposure time.

The failure condition level safety requirement is with respect to parameter P_{EI}, quantifying the extremely improbable accident probability allocated to an individual failure condition, taken to be 1×10^{-9}.

$$P(s|\theta,\mu)P(acc|s,\theta,\mu) \leq P_{EI} \quad \forall\ s \in S \tag{1.1}$$

The aircraft TLS requirement is with respect to parameter P_{TLS}, quantifying the allowable accident probability due to any failure conditions. ILS certification assumes this probability to be 1×10^{-7}. This values comes from taking the value of 1×10^{-9} used for individual failure conditions, assuming approximately 100 possible independent failure conditions, and lumping them together to produce a safety requirement at a higher level.

$$\sum_{s \in S} P(s|\theta,\mu)P(acc|s,\theta,\mu) \leq P_{TLS} \tag{1.2}$$

An implicit assumption in these formulations is that any uncertainty contributing to $P(acc|s,\theta,\mu)$ is effectively another failure condition considered jointly with s. Further safety requirements not explored here could be formulated. The aircraft TLS requirements could be extended to a fleet or airspace TLS requirement which would need to account for the number of individual aircraft operating in the system. Alternatively, a utility based formulation could consider the expected cost or reward of possible outcomes as a way to set the allowable accident probability.

A methodology for assessing the compliance of a system design with airworthiness safety regulations within this probabilistic framework must address the following:

i) Enumerate possible failure conditions

ii) Determine the probability of each failure condition

iii) Compute the probability of an accident given each failure condition

The first task requires a comprehensive understanding of the full system, including all

system components, possible failure modes, and contributing weather events that may be considered failure conditions. Methodologies such as fault tree analysis [24] or Bowtie analysis [25] are suited to this task. The second task requires acquiring the prior probability of each failure condition, considering the reliability of system components [26] and empirical data on weather events [27]. The third task, the subject of this work, is concerned with determining the safety of an operation given that a particular failure condition has occurred.

1.5 Thesis Outline

The topics covered in this thesis are outlined as follows. A fast-time simulation model for a large fixed-wing RPAS on final approach is developed in Chapter II. Dryden gusts, vertical wind shear profile, and microburst events are accounted for in the disturbance model. A navigation system model composed of GPS pseudorange, radar altimeter, and IMU sensors is outlined as well as a method for fusing observation using an Extended Kalman filter and detecting errors with Bayesian likelihood ratios. Additionally, Chapter II presents practical methods for sampling sources of uncertainty, including multivariate normal distribution, process noise, and the timing of a sequence of discrete events. A methodology for estimating the probability of rare events is presented in Chapter III. A black-box simulation model produces safety distance metric outputs from randomly sampled inputs. Importance sampling is used to efficiently estimate exceedingly rare events and a technique utilizing cross-entropy minimization and stochastic differential equations is presented to acquire optimal importance sampling distributions. The safety assessment methodology is applied in Chapter IV to ensure passive risk mitigation in an approach and landing scenario by setting procedure and system design parameters. An online safety assessment methodology is developed in Chapter V to achieve active risk mitigation that ensures a UAS always has procedure options remaining that satisfy safety requirements. The performance benefits of the online safety assessment in off-nominal situations are demonstrated by comparison to

the baseline system and procedure designed in the previous chapter.

CHAPTER 2
SYSTEM MODELLING

The probability of accidents due to failure conditions and enumeration of failure conditions must be demonstrated through some means to satisfy safety requirements. AC 25.1309 suggests several means for demonstrating validating safety requirements. Sufficiently simple systems may demonstrate safety by qualitative methodologies, relying on judgement of subject matter experts. Systems of higher complexity with many failure conditions and dependencies generally require quantitative methodologies to demonstrate safety. Quantitative model-based methodologies are practically necessary in systems with high degrees of automation and numerous subsystem. Modelling of failure conditions in the RPAS approach and landing operations should consider both the discrete states of the system, such as system faults and monitor alerts, and continuous states of the system, such as flight dynamical states and the wind field. Uncertainty present in the system, such as discrete state transitions due to failure events or path following errors due to gusts and navigation sensor noise, should be modelled. Accidents and hazardous states must be defined within the model as well. Several modelling techniques are proposed to handle the complexity present in the RPAS approach and landing operation.

2.1 Modelling Overview

2.1.1 Stochastic Hybrid System model

It is assumed that the system can be modelled as a Stochastic Hybrid System (SHS) [28], consisting of continuous states $x(t) \in \mathbb{R}^l$ and discrete state $q(t) \in Q$. The dynamics of the continuous states are described by a Stochastic Differential Equation (SDE).

$$dx(t) = f(x(t), q(t))dt + g(x(t), q(t))dw_t \tag{2.1}$$

Where $f : \mathbb{R}^l \times Q \to \mathbb{R}^l$, $w_t : \mathbb{R}^+ \to \mathbb{R}^k$ is a k-dimensional Wiener process, and $g : \mathbb{R}^l \times Q \to \mathbb{R}^{l \times k}$. Dynamics of discrete states are described by transition rates $\lambda_{i,j}(x(t), q(t)) : \mathbb{R}^l \times Q \to \mathbb{R}^+$, where $i, j \in Q$. Discrete state transitions may produce an instantaneous remapping of the continuous states such that $x^+(t) = \phi(x^-(t), q^+(t), q^-(t))$, where $\phi : \mathbb{R}^l \times Q \times Q \to \mathbb{R}^l$. The SHS model can capture the continuous state dynamics of the aircraft flight mechanics, turbulence, and noise processes, coupled with the discrete state dynamics of failure conditions such as sudden navigation system errors, and weather events.

2.1.2 Autonomous Agent Model

Due to the possibility of lost link failure conditions, cases must be considered in which the UAV is operating autonomously. Any autonomous systems onboard the UAV capable of making decisions must be modelled in order to quantify safety. An agent-based framework is adopted to model the autonomous systems. This framework explicitly considers the fact that an autonomous agent does not have direct access to the true state of the system. Information about the state of the system is only available via observations from sensors, monitors, and inbound communications. Likewise, the ability of an agent to change the state of the system depends on available actuators, commands, and outbound communications. Various levels of autonomy can exist to tie the observations to actuation. A basic agent may be purely reactive with limited processing between input and output. The Observe-Orient-Decide-Act (OODA) agent framework provides a higher level of autonomy more applicable to an RPAS. The Orient stage updates an internal model of the relevant system states using observations. This stage may use state estimation and sensor fusion techniques to produce a probabilistic internal model representing the agent's internal belief state. The internal model can be used during the Decide stage to generate the next set of actions. A flowchart of the abstracted structure of an autonomous decision-making agent is depicted

in Figure 2.1.

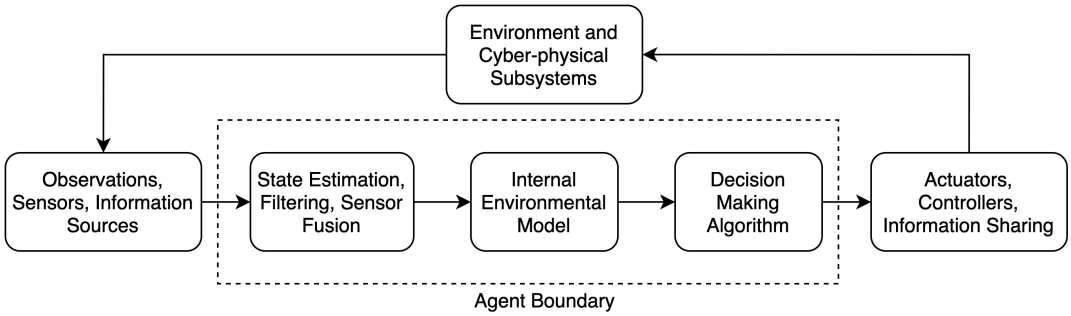

Figure 2.1: Autonomous Agent Model

Observations available to the agent may be describes within the SHS model by $y(t) \in \mathbb{R}^m$.

$$y(t) = h(x(t), q(t)) + r(x(t), q(t))dv_t \tag{2.2}$$

where $h : \mathbb{R}^l \times Q \to \mathbb{R}^m$, $v_t : \mathbb{R}^+ \to \mathbb{R}^p$ is a p-dimensional Wiener process, and $r : \mathbb{R}^l \times Q \to \mathbb{R}^{m \times p}$. The agent uses these observations and system model to compute the belief state estimate as the random variables $\hat{x}(t)$ and $\hat{q}(t)$. An Extended Kalman filter is adopted to estimate the distribution of continuous system states, which are then used as inputs for aircraft guidance, navigation, and control systems. The Kalman filtering framework is also used to detect discrete state changes from nominal conditions to failure mode conditions using likelihood ratio testing.

The belief state distribution may be used for safety assessments by propagating the distribution forward using the stochastic dynamics to estimate the state distribution at future times and estimate accident probabilities. Specifics of the decision making algorithm and how internal states are used to maintain safety in off-nominal conditions are presented and discussed in future chapters.

2.1.3 Model Uncertainty Representation

Various sources of uncertainty contribute to the model and practical methods for representing and sampling random variables are necessary for simulation. Many continuous states are modelled as multivariate random variables, which may be readily sampled. Process noise must be sampled at every time step of the simulation for each white noise channel, resulting in a high dimensional sampling space. Dimensionality reduction is performed using Karhunen-Loeve expansion, producing a suitable approximation of process noise using significantly fewer samples. The sampling of discrete state transition event times is accomplished using an overbounding approximation utilizing the Dirichlet distribution. A sequence of event times can then be obtained using Gamma distributed samples.

2.2 Flight Dynamics Model

The aircraft is modelled using a 6-Degree of Freedom flight dynamics model. The position $[x, y, z]$ is defined in the inertial reference frame attached to the center of the runway threshold with the x axis aligned with the runway center-line and the z axis aligned downwards towards the center of the Earth. The center of mass velocity in the vehicle-fixed frame $[v_x, v_y, v_z]$ is related the velocity in the body-fixed frame $[u, v, w]$ by rotations about the standard Euler angles $[\phi, \theta, \psi]$ as shown in Equation 2.3. s, c, and t respectfully denote the sine, cosine, and tangent of the Euler angle in the subscript.

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix} = \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} = \begin{pmatrix} c_\theta c_\psi & s_\phi s_\theta c_\psi - c_\phi s_\psi & c_\phi s_\theta c_\psi + s_\phi s_\psi \\ c_\theta s_\psi & s_\phi s_\theta s_\psi + c_\phi c_\psi & c_\phi s_\theta s_\psi - s_\phi c_\psi \\ -s_\theta & s_\phi c_\theta & c_\phi c_\theta \end{pmatrix} \begin{bmatrix} u \\ v \\ w \end{bmatrix} \tag{2.3}$$

The rate of change of the Euler angles is related to the body-fixed frame rotation rates

$[p, q, r]$ in Equation 2.4.

$$\begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix} = \begin{pmatrix} 1 & s_\theta t_\theta & c_\phi t_\theta \\ 0 & c_\phi & -s_\phi \\ 0 & \frac{s_\phi}{c_\theta} & \frac{c_\phi}{c_\theta} \end{pmatrix} \begin{bmatrix} p \\ q \\ r \end{bmatrix} \tag{2.4}$$

The rate of change of the velocity in the body-fixed frame is given in Equation 2.5, where m is the aircraft mass, g is the gravitational acceleration, and $[f_x, f_y, f_z]$ are the force components in the body-fixed frame.

$$\begin{bmatrix} \dot{u} \\ \dot{v} \\ \dot{w} \end{bmatrix} = \begin{bmatrix} rv - qw \\ pw - ru \\ qu - pv \end{bmatrix} + \begin{bmatrix} -gs_\theta \\ gc_\theta s_\phi \\ gc_\theta c_\phi \end{bmatrix} + \frac{1}{m} \begin{bmatrix} f_x \\ f_y \\ f_z \end{bmatrix} \tag{2.5}$$

The force components encapsulate aerodynamic and thrust forces. The angular rate dynamics are nominally described by the moment equation, however, assumptions validated in the development of the Tool for Analysis of Separation and Throughput (TASAT) flight dynamics model [29][30] can be applied to simplify the dynamics, allowing larger timesteps to be used in fast-time simulation. It is assumed that low-level flight control of angular rates is sufficiently fast to track reference rates $[p_{ref}, q_{ref}, r_{ref}]$ provided by the higher-level controller on a timescale of τ_{rate}.

$$\begin{bmatrix} \dot{p} \\ \dot{q} \\ \dot{r} \end{bmatrix} \approx \frac{1}{\tau_{rate}} \begin{bmatrix} p_{ref} - p \\ q_{ref} - q \\ r_{ref} - r \end{bmatrix} \tag{2.6}$$

This simplification aids in allowing fast-time simulation of flight dynamics, however, it may be dropped and replaced by the underlying moment equations if fast time-scale angular rate dynamics are required to validate the flight dynamics model of a particular aircraft.

The aircraft flight configuration accounts for the state of the landing gear δ_G, speedbrake/spoiler δ_S, flap angle δ_F, and thrust level δ_T. Configuration is controlled by setting

reference values which the system approaches with first order dynamics defined by configuration change timescales.

The aircraft is controlled by an autopilot and Flight Management System (FMS) which tracks a procedure profile defined analogously to Area Navigation (RNAV) procedures. A lateral controller modelling Lateral Navigation (LNAV) tracks a lateral procedure defined by straight-line path segments linked by constant radius turns, and a longitudinal controller modelling Vertical Navigation (VNAV) tracks a vertical procedure defined by a piece-wise linear altitude profile and calibrated airspeed profile defined with respect to lateral path distance to the runway threshold.

2.3 Wind Model

Much of flight technical error on final approach is due to disturbances from winds and gust acting on the aircraft. FAA AC 120-28D[31] defines several wind models suitable for simulation of final approach and landing scenarios. The wind model accounts for the mean vertical wind profile, stochastic gusts due to Dryden turbulence, and rare events causing severe wind shear.

2.3.1 Mean Wind Profile

The wind profile is defined by a vertical wind shear with a logarithmic profile near the ground and a piecewise linear interpolation for winds aloft. Height above ground level h is expressed in feet and $\bar{V}_{20}$ is the mean wind speed magnitude at 20 ft above ground level.

$$\bar{V}_{wind}(h) = 0.20407\bar{V}_{20} \ln\left(\frac{h+0.15}{0.15}\right) \text{ if } h < 1000 \text{ ft} \tag{2.7}$$

Tailwind/headwind and crosswind components, $[V_{wind,x}, V_{wind,y}]$, are defined with respect to the runway reference frame. Winds above 1000 ft are linearly interpolated from mean wind measurements at several reference heights.

2.3.2 Dryden Turbulence

Turbulence due to various non-uniformities and instabilities of the atmosphere is an important contribution to flight technical error. The Dryden Turbulence model is chosen to simulate gusts acting on the aircraft due to its rational power spectral density, which lends itself to straightforward simulation by filtering white noise inputs [31]. Gust components $[u_g, v_g, w_g]$ are computed in the wind frame and are dependent on airspeed V_∞, height dependent turbulence length scales, L_u, L_v, L_w, and turbulence intensities σ_u, σ_v, σ_w. The time evolution of the gusts may described by SDEs as follows.

$$du_g = -\frac{V_\infty}{L_u} u_g dt + \sqrt{2\frac{V_\infty}{L_u}}\sigma_u dw_u \tag{2.8}$$

$$dv_g = -\frac{V_\infty}{L_v} v_g dt + \sqrt{2\frac{V_\infty}{L_v}}\sigma_v dw_v \tag{2.9}$$

$$dw_g = -\frac{V_\infty}{L_w} w_g dt + \sqrt{2\frac{V_\infty}{L_w}}\sigma_w dw_w \tag{2.10}$$

Independent Wiener processes w_u, w_v, and w_w contribute to the gust processes. The length scale and intensity parameters are described using a piece function of altitude h.

$$L_w = \begin{cases} h & \text{if } h < 1000 \text{ ft} \\ 1000 + 0.75(h - 1000) & \text{if } 1000 \text{ ft} \leq h < 2000 \text{ ft} \\ 1750 \text{ ft} & \text{if } > 2000 \text{ ft} \end{cases} \tag{2.11}$$

$$L_u = L_v = \begin{cases} \frac{h}{(0.177+0.00823h)^{1.2}} & \text{if } h < 1000 \text{ ft} \\ 1000 + 0.75(h - 1000) & \text{if } 1000 \text{ ft} \leq h < 2000 \text{ ft} \\ 1750 \text{ ft} & \text{if } h \geq 2000 \text{ ft} \end{cases} \tag{2.12}$$

The vertical RMS turbulence intensity σ_w is defined in relation to $\bar{V}_{20}$.

$$\sigma_w = 0.1\bar{V}_{20} \tag{2.13}$$

Horizontal turbulence intensities depend on height above ground level and σ_w.

$$\sigma_u = \sigma_v = \begin{cases} \sigma_w(0.177 + 0.00823h)^{0.4} & \text{if } h < 1000 \text{ ft} \\ \sigma_w & \text{if } h \geq 1000 \text{ ft} \end{cases} \tag{2.14}$$

2.3.3 Wind Shear

Exceptional weather events such as microbursts may produce severe wind shear that varies over horizontal distance. The onset wind shear begins at time t_{ws}, at which point $d = 0$, and ramps proportionally with ground distance traveled by γ_{ws} until d_{max} is reached and wind shear reaches a maximum magnitude [32]. An expression for the change in tailwind or headwind due to wind shear, Δu_{ws}, is given below.

$$\Delta u_{ws} = \begin{cases} 0 & \text{if } t < 0 \\ \gamma_{ws} d & \text{if } t \geq 0 \text{ and } d < d_{max} \\ \gamma_{ws} d_{max} & \text{if } t \geq 0 \text{ and } d \geq d_{max} \end{cases} \tag{2.15}$$

More complicated wind shear models may be found in [32], however, a simple model using a tail wind ramp demonstrates a worst case scenario for final approaches.

2.4 Navigation System Model

Position and velocity estimation is accomplished using a navigation system comprised of several components. Primary navigation is provided is by Wide Area Augmentation System (WAAS) Global Positioning System (GPS). This produces position estimates that may be filtered to produce velocity estimates. Vertical position accuracy may be enhanced using

radio altimeters which become more accurate as height above ground level decreases. A model for worst case scenario drift errors affecting vertical position estimates is presented. Further accuracy increases may be obtained by integrating an Inertial Measurement Unit (IMU) accelerometer which acts to smooth the position and velocity estimate. Each of these measurements are integrated using a Kalman filter to produce a state estimate with improved accuracy and a multivariate Normal uncertainty estimate.

2.4.1 GPS Model

Primary navigation is accomplished using a Global Navigation Satellite System, in particular, the Global Positioning System (GPS). A constellation of satellites each send signals at times t_i from individual positions $[x_i, y_i, z_i]$. Initial corrections for clock time and satellite ephemeris errors may already be applied. The elapsed time between when the signal was transmitted and received can be used to compute pseudorange estimates $\rho_i(t)$.

$$\rho_i = |c(t_i - t)| = |\tau_i - \tau| = R_i + \epsilon_i \tag{2.16}$$

$$R_i = R(\mathbf{x} - \mathbf{x}_i) = \sqrt{(x_i - x)^2 + (y_i - y)^2 + (z_i - z)^2} \tag{2.17}$$

Scaled time values τ are defined by the corresponding time t multiplied by the speed of light c. The error between the pseudorange measurement ρ_i and the true range R_i is described by ϵ_i. GPS pseudorange errors are due to many sources including clock errors, ephemeris errors, troposphere errors, ionosphere errors, multipath errors, and measurements errors [33]. Certain error components such as clock and ephemeris errors may be corrected through the use of satellite based augmentation, in particular the Wide Area Augmentation System (WAAS) [34]. Spatially correlated errors due to ionospheric weather may also be corrected by a network of ground receiver stations, however, spatial decorrelations on a length scale smaller than the ground station baseline distance will remain

uncorrected. A general model for pseudorange errors will include a measurement noise term σ_m and a time correlated error term $\epsilon_{GM,i}$.

$$\epsilon_i(t) = \epsilon_{GM,i}(t) + \sigma_m dv_t^i \tag{2.18}$$

The time correlated error is modelled using a 1st order stationary Gauss-Markov Process, also known as an Ornstein–Uhlenbeck process. The error is correlated with time constant τ_{GM} and has a variance of σ_{GM}^2, thus the process has an autocovariance between times t and s of $\sigma_{GM}^2 e^{-|t-s|/\tau_{GM}}$ [33].

$$d\epsilon_{GM,i}(t) = -\frac{1}{\tau_{GM}}\epsilon_{GM,i}(t)dt + \sqrt{\frac{2}{\tau_{GM}}}\sigma_{GM}dw_t^i \tag{2.19}$$

Estimation of position and time using GPS pseudoranges usually requires at least 5 usable satellite signals[35] to produce a sufficiently accurate estimate, thus suggesting a simulation would require modelling many simultaneous Gauss-Markov error processes. By exploiting locally linear estimation solutions and identical error statistics, the position estimation error may be modelled using only 3 error processes and a Dilution of Precision (DOP) matrix. The receiver clock time is also subject to errors, thus it is also estimated alongside the position. We start by linearizing the pseudorange measurements about an initial estimate $\bar{\mathbf{x}} = [\bar{x}, \bar{y}, \bar{z}, \bar{\tau}]$ and solving for position/time update $\Delta\mathbf{x} = [\Delta x, \Delta y, \Delta z, \Delta\tau]$ which minimizes the total square pseudorange error.

$$\Delta\rho = \begin{bmatrix} \Delta\rho_1 \\ \Delta\rho_2 \\ \vdots \\ \Delta\rho_N \end{bmatrix} = \begin{bmatrix} R(\mathbf{x}_1 - (\bar{\mathbf{x}} + \Delta\mathbf{x})) - \sqrt{(\tau_1 - (\bar{\tau} + \Delta\tau))^2} \\ R(\mathbf{x}_2 - (\bar{\mathbf{x}} + \Delta\mathbf{x})) - \sqrt{(\tau_2 - (\bar{\tau} + \Delta\tau))^2} \\ \vdots \\ R(\mathbf{x}_N - (\bar{\mathbf{x}} + \Delta\mathbf{x})) - \sqrt{(\tau_N - (\bar{\tau} + \Delta\tau))^2} \end{bmatrix} \approx \begin{bmatrix} \Delta\bar{\rho}_1 \\ \Delta\bar{\rho}_2 \\ \vdots \\ \Delta\bar{\rho}_N \end{bmatrix} + \mathbf{A} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \\ \Delta\tau \end{bmatrix} \tag{2.20}$$

The Jacobian of the pseudorange error with respect to the position/time updates about the initial estimates is described by matrix $\mathbf{A}$. The initial range estimates are given by $\bar{R}_i$.

$$\mathbf{A} = \begin{pmatrix} (x_1 - \bar{x})/\bar{R}_1 & (y_1 - \bar{y})/\bar{R}_1 & (z_1 - \bar{z})/\bar{R}_1 & -1 \\ (x_2 - \bar{x})/\bar{R}_2 & (y_2 - \bar{y})/\bar{R}_2 & (z_2 - \bar{z})/\bar{R}_2 & -1 \\ \vdots & \vdots & \vdots & \vdots \\ (x_N - \bar{x})/\bar{R}_N & (y_N - \bar{y})/\bar{R}_N & (z_N - \bar{z})/\bar{R}_N & -1 \end{pmatrix} \tag{2.21}$$

The position/time updates which minimize the total square error can be solved for using the linearized pseudorange measurement error.

$$\begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \\ \Delta\tau \end{bmatrix} = \operatorname{argmin} \Delta\rho^T \Delta\rho \approx \operatorname{argmin}(\Delta\bar{\rho} + \mathbf{A}\Delta\mathbf{x})^T(\Delta\bar{\rho} + \mathbf{A}\Delta\mathbf{x}) \tag{2.22}$$

The solution to this ordinary least squares problem is solved using the pseudoinverse of the Jacobian of the pseudorange.

$$\mathbf{A}^T\mathbf{A}\Delta\mathbf{x} = \mathbf{A}^T\Delta\bar{\rho} \tag{2.23}$$

$$\Delta \mathbf{x} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\Delta\bar{\rho} \tag{2.24}$$

The position/time update is a linear combination of the pseudorange measurement errors, thus the position estimation error will have the same statistics as a weighted sum of the pseudorange measurement errors ϵ_i. The covariance between independent pseudorange measurement errors at time t and s can be expressed using the combination of the measurement noise and the time correlated Gauss-Markov process.

$$\begin{aligned} \text{cov}\Delta\bar{\rho} &= \mathbb{E}[\Delta\bar{\rho}\Delta\bar{\rho}^T] \\ &= (\sigma_m^2\delta(t-s) + \sigma_{GM}^2 e^{-|t-s|/\tau_{GM}})\mathbf{I}_N \end{aligned} \tag{2.25}$$

The covariance matrix of the position/time estimation error can be expressed using the DOP matrix $\mathbf{Q}_{DOP} = (\mathbf{A}^T\mathbf{A})^{-1}$, and the pseudorange measurement error covariance.

$$\begin{aligned} \text{cov}\Delta\mathbf{x} &= \mathbb{E}[\Delta\mathbf{x}\Delta\mathbf{x}^T] \\ &= \mathbb{E}[((\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\Delta\bar{\rho})((\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\Delta\bar{\rho})^T] \\ &= (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbb{E}[\Delta\bar{\rho}\Delta\bar{\rho}^T]\mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1} \\ &= (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T(\sigma_m^2\delta(t-s) + \sigma_{GM}^2 e^{-|t-s|/\tau_{GM}})\mathbf{I}_N\mathbf{A}(\mathbf{A}^T\mathbf{A})^{-1} \\ &= (\mathbf{A}^T\mathbf{A})^{-1}(\sigma_m^2\delta(t-s) + \sigma_{GM}^2 e^{-|t-s|/\tau_{GM}}) \\ &= \mathbf{Q}_{DOP}(\sigma_m^2\delta(t-s) + \sigma_{GM}^2 e^{-|t-s|/\tau_{GM}}) \end{aligned} \tag{2.26}$$

The DOP matrix is only dependent on the number and geometry of the satellites used in the position calculation and essentially amplifies the pseudorange measurement errors. DOP is computed during operation and may be used to guide geometry screening for optimizing accuracy. Large values in DOP matrix result in larger estimation errors while small values are desirable for high accuracy. The first three entries along the diagonal of $\mathbf{Q}_{DOP}$, $[\sigma_x^2, \sigma_y^2, \sigma_z^2]$, are of primary interest for describing the navigation accuracy, alongside σ_m,

Table 2.1: WAAS GPS Model Parameters

Parameter	Value
σ_x	2.0
σ_y	2.0
σ_z	2.4
σ_m	0.1 m
σ_{GM}	1.0 m
τ_{GM}	600 s

σ_{GM}, and τ_{GM}. The observed performance of WAAS GPS is used to guide the choice of parameters. 4 m is the recommended conservative 95% lower limit for both horizontal and vertical accuracy as it exceeds the observed 95% accuracy bound and the extrapolated $1 - 1 \times 10^{-9}$ limit assuming a Normally distributed error is equal to the maximum observed error of 12 m [34]. The geometry of a satellite constellation usually results in σ_z larger than the horizontal DOP components, thus an adjustment factor of 1.2 is added to conservatively reflect the reduced vertical accuracy [35]. pseudorange error for individual satellites has been observed to have a standard deviation around 1.0 m and is treated as a Gauss-Markov process with a time constant of 600 s [33]. Measurement noise is estimated to have a standard deviation of 0.1 m [33]. The WAAS GPS model parameters are listed in Table 2.1.

2.4.2 Radar Altimeter Model

Operations very close to the ground require a high degree of navigation accuracy, especially on the vertical axis. This may be provided by Differential GPS, which corrects for navigation errors using an independent ground based reference receiver, however independent height measurements using an aircraft based radar altimeter may also provided the vertical navigation precision required for landing operations [36]. An observation model for radar altimeter measurements returns the height above terrain plus a constant measurement noise term σ_{RALT} and an additional noise term proportional to the height above terrain by a fac-

Table 2.2: Radio Altimeter Model Parameters

Parameter	Value
σ_{RALT}	1.0 m
γ_{RALT}	0.05

tor γ_{RALT}. This accounts for the increase in measurement uncertainty with height due to spurious radar returns from possibly uneven terrain.

$$h(t) = |z(t) - z_0(t)| + \sqrt{\gamma_{RALT}^2 (z(t) - z_0(t))^2 + \sigma_{RALT}^2} dv_t \tag{2.27}$$

The terrain elevation below the aircraft at time is denoted by $z_0(t)$. This term may be subject to spatially correlated errors in elevation surveying and lateral navigation errors, however, these errors are neglected for the approach and landing scenario by assuming the terrain immediately surrounding the runway is level. Conservative parameter values chosen for the radio altimeter are listed in Table 2.2 [37].

2.4.3 Navigation Drift Errors

The error models presented for WAAS GPS and radar altimeters are applicable to nominal operation of the respective system. Off-nominal operating conditions may introduce new errors which are unaccounted for by the nominal models. These off-nominal conditions may cause sudden jumps or bias errors to the measurements due to events such as GPS satellite outages [35]. Drift errors induced by events with a slow onset may be harder to detect and constitute worst case scenarios. GPS errors induced by ionosphere gradients are of particular interest as they slowly decorrelate aircraft pseudorange errors from those measured at ground reference stations, thus these errors are unaccounted for by WAAS corrections. Detailed models account for pseudorange errors for individual satellite signals due to a moving gradient in ionosphere delay, parameterized by orientation, velocity, gradient, and maximum [38]. The vertical axis is most sensitive to errors during final ap-

proach, as runway undershoots may result from vertical errors amplified proportionally to the cotangent of the glide slope. A simplified worst-case model should account for ramp errors on the vertical position estimate, which begin at some event time and grow proportionally to gradient γ_z and distance d travelled until a maximum is reached at $d = d_{max}$. This model is described by Equation 2.28.

$$\Delta z = \begin{cases} 0 & \text{if } t < 0 \\ \gamma_z d & \text{if } t \geq 0 \text{ and } d < d_{max} \\ \gamma_z d_{max} & \text{if } t \geq 0 \text{ and } d \geq d_{max} \end{cases} \tag{2.28}$$

The intent of this error model is to trick the navigation system into slowly overestimating altitude, resulting in an increased sink rate and glideslope undershoot to compensate. A worst case ionosphere induced GPS error gradient of 0.5 m/km has been observed [38]. Off-nominal errors also exist for radar altimeter measurements [39], however they are not well characterized. This same error model is adopted for worst case radar altimeter drift errors. Specific parameter values are left as a variable for analysis of system requirements.

2.4.4 Inertial Measurement Unit Model

Navigation accuracy can be improved further using observations from an Inertial Measurement Unit (IMU) [40]. A general IMU produces measurements of acceleration, angular rates, and attitude from an accelerometer and gyroscope mounted at some location on the aircraft. Some IMUs designs isolate the accelerometer from rotations by mounting it within a gyroscope, while strap down IMUs fix the accelerometer to the aircraft body reference frame. Several assumptions are made to simplify IMU modelling, in particular, angular rate and orientation estimates have small enough error such that the transformation of accelerometer measurements from the body fixed reference frame at the strap down location to the vehicle fixed reference frame at center of mass introduces negligible additional error. A conservative accelerometer noise standard deviation of $\sigma_{accel} = 0.01$ m/s^2 is chosen for

each axis [40]. The integrated velocity error Δv_i for each axis i over time step dt is given in Equation 2.29.

$$\Delta v_i = \sigma_{accel}\sqrt{dt}\; dw_t^i \tag{2.29}$$

2.5 State Estimation and Fault Monitoring Model

Systems onboard the aircraft are tasked with processing incoming observation data from various sensors and information sources to produce an estimate of the true state of the environment. The state estimate is used for aircraft guidance, navigation, control, and decision making, thus the accuracy and timeliness of state estimates are critical for maintaining the situational awareness required for safe operation.

A Bayesian model-based methodology is adopted for both state estimation and fault monitoring. The Kalman filter is used to estimate the state of the system conditional on a nominal fault-free model and likelihood tests are applied to trigger alerts when statistically significant deviations from the fault-free model are observed. A bank of Kalman filters assuming various fault states is then applied to diagnose and exclude the fault.

2.5.1 Kalman Filtering

It is assumed that the aircraft environment is composed of continuous states (such as position, velocity, and time correlated sensor errors) and discrete states (such as fault conditions and aircraft configurations) that can be generally modelled as a Stochastic Hybrid System (SHS). The SHS models the continuous state dynamics using stochastic differential equations with parameters dependent on the discrete states. Discrete state transitions are modelled using Markov transition rates possibly dependent on the continuous states. Discrete state transition events may also be accompanied by a stochastic remapping of the continuous states, including the sampling of new dynamics and observation parameters for the faulted state. The system may often be simplified by assuming the system dynamics are approximately linear, noise terms are Gaussian, and transition rates are constant. Given that

the real observations will be made at a finite sampling rate, the dynamics and observations are discretized in time and presented below.

$$\mathbf{x}_t = \mathbf{F}_t(q)\mathbf{x}_{t-1} + \mathbf{w}_t \tag{2.30}$$

$$\mathbf{z}_t = \mathbf{H}_t(q)\mathbf{x}_t + \mathbf{v}_t \tag{2.31}$$

$\mathbf{x}_t \in \mathbb{R}^n$ is the continuous state vector, $\mathbf{z}_t \in \mathbb{R}^m$ is the observation vector, $\mathbf{F}_t(q) \in \mathbb{R}^{n\times n}$ is the dynamics matrix, and $\mathbf{H}_t(q) \in \mathbb{R}^{m\times n}$ is the observation matrix. Independent noise vectors $\mathbf{w}_t$ and $\mathbf{v}_t$ are drawn from independent multivariate normal distributions $N(\mathbf{0}, \mathbf{Q}_t(q))$ and $N(\mathbf{0}, \mathbf{R}_t(q))$ respectively. $q \in Q$ represents the discrete state of the system, for example, $q = 0$ is the nominal state and $q = 1$ is a particular fault state. Extended Kalman filtering can extend the Kalman filtering framework to nonlinear dynamics, observations, and noise functions by taking Jacobians of the respective functions about the mean state estimate.

The continuous time nonlinear dynamics are restated in the SDE below.

$$dx(t) = f(x(t), q(t))dt + g(x(t), q(t))dw_t \tag{2.32}$$

Linearized discrete time dynamics matrix $\mathbf{F}_t(q)$ is derived by taking the Jacobian of the dynamics function over time step Δt.

$$\mathbf{F}_t(q) = \mathbf{I}_n + \begin{bmatrix} \nabla_{\mathbf{x}}^{\mathbf{T}} f_1(\mathbf{x}_{t-1}, q(t-1)) \\ \vdots \\ \nabla_{\mathbf{x}}^{\mathbf{T}} f_n(\mathbf{x}_{t-1}, q(t-1)) \end{bmatrix} \Delta t \tag{2.33}$$

The dynamics noise covariance matrix $\mathbf{Q}_t(q)$ is obtained from the nonlinear dynamics noise coefficient matrix g.

$$\mathbf{Q}_t(q) = g(\mathbf{x}_{t-1}, q(t-1))g(\mathbf{x}_{t-1}, q(t-1))^T \Delta t \tag{2.34}$$

The nonlinear observation is described below.

$$y(t) = h(x(t), q(t)) + r(x(t), q(t))dv_t \tag{2.35}$$

Linearized observation matrix $\mathbf{H}_t(q)$ is derived by taking the Jacobian of the observation function.

$$\mathbf{H}_t(q) = \begin{bmatrix} \nabla_{\mathbf{x}}^{\mathbf{T}} h_1(\mathbf{x}_t, q(t)) \\ \vdots \\ \nabla_{\mathbf{x}}^{\mathbf{T}} h_m(\mathbf{x}_t, q(t)) \end{bmatrix} \tag{2.36}$$

The observation noise covariance matrix $\mathbf{R}_t(q)$ is obtained from the observation noise coefficient matrix r.

$$\mathbf{R}_t(q) = r(\mathbf{x}_t, q(t))r(\mathbf{x}_t, q(t))^T \tag{2.37}$$

Transitions from the nominal state to a fault state result in a change in the dynamics, observation, and noise parameters. The parameters in the faulted state are not necessarily known beforehand and may be drawn from a distribution. The past observations up to time step t are together represented as $\mathbf{Z}_t$. Observations may also be truncated to length k as $\mathbf{Z}_{t,k}$. Kalman filtering using the observations and assumed fault-free model of the system produces an estimate of the true state position as a multivariate normal distribution. Dependencies on q are dropped from notation in the following sections where the fault-free model is assumed.

$$\mathbf{x}_t | \mathbf{Z}_t \sim \mathcal{N}(\hat{\mathbf{x}}_t, \mathbf{P}_t) \tag{2.38}$$

The mean and covariance of the state estimate can be predicted for the next time step. Superscript $^-$ indicates the prior estimate before the correction step is made.

$$\hat{\mathbf{x}}_t^- = \mathbf{F}_t \hat{\mathbf{x}}_{t-1} \tag{2.39}$$

$$\mathbf{P}_t^- = \mathbf{F}_t \mathbf{P}_{t-1} \mathbf{F}_t^T + \mathbf{Q}_t \tag{2.40}$$

The residuals $\mathbf{y}_t$ between the new observations and expected observations are computed, as well as the anticipated observation covariance $\mathbf{S}_t$.

$$\mathbf{y}_t = \mathbf{z}_t - \mathbf{H}_t \hat{\mathbf{x}}_t^- \tag{2.41}$$

$$\mathbf{S}_t = \mathbf{H}_t \mathbf{P}_t^- \mathbf{H}_t^T + \mathbf{R}_t \tag{2.42}$$

The Kalman gain $\mathbf{K}_t$ is computed and the corrections are applied to the mean and covariance of the state estimate.

$$\mathbf{K}_t = \mathbf{P}_t^- \mathbf{H}_t^T \mathbf{S}_t^{-1} \tag{2.43}$$

$$\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_t^- + \mathbf{K}_t \mathbf{y}_t \tag{2.44}$$

$$\mathbf{P}_t = (\mathbf{I} - \mathbf{K}_t \mathbf{H}_t) \mathbf{P}_t^- \tag{2.45}$$

The state estimate $\hat{\mathbf{x}}_t$ is used in feedback control of the aircraft. When the variance of the state estimate is too large, the state estimate is unsuitable for completing a safe landing. The basic Kalman filter does not consider the occurrence of a fault or how it might change model parameters to produce hazardously misleading information. The structure of the

Kalman filter and system model lends themselves to the application of likelihood-ratio tests for detecting deviations from the assumed fault-free model.

2.5.2 Fault Detection

Beyond producing an estimation of the continuous system state, the Kalman Filtering framework may be used to detect deviations of the system from an assumed model by comparing observations to their nominal expected distributions. Fault events that produce hazardously misleading information through navigation drift errors can be described as deviations from a nominal model. These faults must be detected with high probability within a critical amount of time in order to acceptably mitigate risk. False alerts that would trigger an unnecessary missed approach must also be limited. The monitor responsible for detecting the fault must be designed such that it satisfies these requirements, involving a trade-off between the parameters defining the monitor, the probability of missed and false detection, the time to alert, and the parameters describing the nominal and faulted system states. Requirements concerning system safety and false alerts are summarized as follows [9].

$$P(\text{accident}) \leq P_{integrity} \approx 2 \times 10^{-7} \text{ per approach} \tag{2.46}$$

$$P(\text{false alert}) \leq P_{continuity} \approx 1 \times 10^{-5} \text{ per 15 seconds} \tag{2.47}$$

Accident probability depends on the time at which the fault occurs during the approach and landing, the severity of the fault, and the time between the occurrence of the fault and execution of a missed approach. Longer alert times may detect a fault with higher confidence, but reduce safety due to increased undetected risk exposure. The International Civil Aviation Organization (ICAO) recommends an alert time of 6 seconds for navigation systems during final approach [9].

This trade-off between false alert rate (Type I errors) and missed detection rate (Type II errors) requires an efficient monitor test statistic to meet the continuity and integrity requirements with minimal alert time. Likelihood-ratio tests have been shown to be the most powerful tests by the Neyman-Pearson Lemma [41] Likelihood-ratio tests compare the likelihood of observations given hypothesized models. Detection of any statistically significant deviation from a null hypothesised model can be accomplished by comparing observations to those expected from the nominal fault-free model. The log-likelihood of the observation at time t, d_t, is computed using the pre-correction measurement residual $\mathbf{y}_t$, and the measurement covariance $\mathbf{S}_t$.

$$d_t = \ln f(\mathbf{y}_t|\mathbf{S}_t) = -\frac{m}{2}\ln 2\pi - \frac{1}{2}\ln\det(\mathbf{S}_t) - \frac{1}{2}\mathbf{y}_t^T\mathbf{S}_t^{-1}\mathbf{y}_t \tag{2.48}$$

The log-likelihood of the past k observations up to time t can be computed as the sum of the individual observation log-likelihoods. The case $k = 1$ corresponds to considering only the observation at time t.

$$e_{t,k} = \sum_{s=t-k+1}^{t} d_s = -\frac{km}{2}\ln 2\pi - \frac{1}{2}\sum_{s=t-k+1}^{t}\ln\det(\mathbf{S}_s) - \frac{1}{2}\sum_{s=t-k+1}^{t}\mathbf{y}_s^T\mathbf{S}_s^{-1}\mathbf{y}_s \tag{2.49}$$

The observed log-likelihood $e_{t,k}$ is chosen to be the test statistic and is compared to a random variable representing the nominal distribution of observation log-likelihood $\bar{e}_{t,k}$, which assumes that the fault-free model is generating the observations. It is hypothesized that observations $\mathbf{y}_s$ are multivariate Gaussian random variables with zero mean and covariance $\mathbf{S}_s$. Assuming the sequence of observations are independent, the sum of the quadratic terms is a chi-squared random variable with km degrees of freedom.

$$\bar{e}_{t,k} \sim -\frac{km}{2}\ln 2\pi - \frac{1}{2}\sum_{s=t-k+1}^{t}\ln\det(\mathbf{S}_s) - \frac{1}{2}\chi^2(km) \tag{2.50}$$

We test the hypothesis that the fault-free model will produce observations with a like-

lihood less than or equal to the observed likelihood. Constant terms are cancelled out leaving the quadratic sum and the chi-squared random variable. This reduces to a one sided chi-square test with km degrees of freedom.

$$P(\bar{e}_{t,k} \leq e_{t,k} | q = 0) = P\big(\chi^2(km) \geq \sum_{s=t-k+1}^{t} \mathbf{y}_s^T \mathbf{S}_s^{-1} \mathbf{y}_s\big) \tag{2.51}$$

The threshold $g_{\alpha,km}$ for triggering an alert is chosen by setting a p-value α corresponding to the target false-alert probability. This may be obtained from the quantile function of the chi-squared distribution.

$$g_{\alpha,km} = Q_{\chi^2}(1 - \alpha; km) \tag{2.52}$$

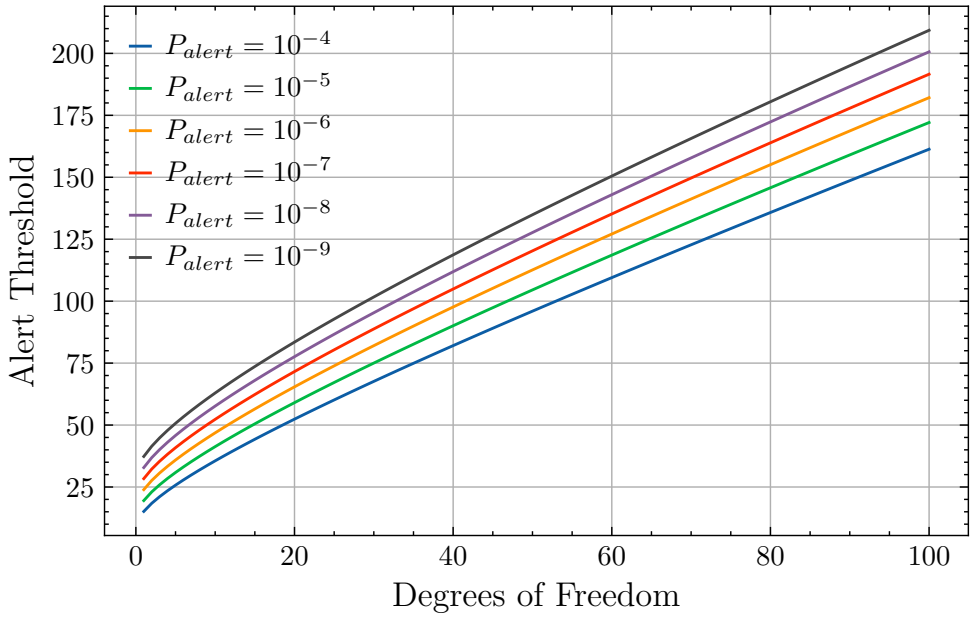

Figure 2.2: Alert Threshold vs Observed Degrees of Freedom for several significance levels

The dependence of the alert limit on degrees of freedom and significance level is depicted in Figure 2.2. In summary, an alert suggesting an α significant deviation from the fault-free model is triggered when the following condition is met:

$$\sum_{s=t-k+1}^{t} \mathbf{y}_s^T \mathbf{S}_s^{-1} \mathbf{y}_s \geq g_{\alpha,km} \tag{2.53}$$

The parameter k is a free variable in the design of the monitor. Small values of k might produce fast alerts, but the high alert threshold required to limit false-alerts may result in a reduced level of protection. Large values of k will reduce the false-alert rate, but time to alert will be larger, potentially increasing the accident rate. The choice of k must meet safety requirements considering both time to alert and error protection level.

2.5.3 Fault Diagnosis and Exclusion

The likelihood-ratio test methodology used for general fault detection may be extended to the diagnosis and exclusion of particular faults that might occur. Examples may include exclusion of faulty radar altimeters or satellites with unacceptably large pseudorange errors from the state estimation. Correct exclusion of faults may allow the system to maintain the performance required for a safe landing or missed approach. Likelihood-ratio tests require hypothesized models for each individual fault state under consideration. Fault models may alter the dynamics, observations, and noise distributions from those of the fault-free model. The Kalman filter effectively 'excludes' measurements from the state estimate solution when their modelled noise magnitude is significantly larger than other available measurements. Separate Kalman filters must be applied in parallel for each hypothesized fault model, resulting in an increase in computational complexity.

The probability of fault model q given the past k observation $Z_{t,k}$ can be computed using Bayes' rule.

$$P(q|Z_{t,k}) = \frac{p(Z_{t,k}|q)P(q)}{p(Z_{t,k})} \tag{2.54}$$

$$p(Z_{t,k}) = \sum_{q \in Q} p(Z_{t,k}|q)P(q) \tag{2.55}$$

The system is assumed to occupy the fault-mode maximizing $P(q|Z_{t,k})$. In addition to k, prior probabilities $P(q)$ are an additional set of parameters in the model. These may be set according to empirical data on fault-rates, or may be free parameters in the design of the monitor. The maximum probability fault-mode will not change under application of the logarithm. The log-likelihood of the past k observations conditional on fault model q is indicated by $e^q_{t,k}$. The most likely fault-mode $\hat{q}$ is thus expressed as follows.

$$\hat{q} = \underset{q \in Q}{\operatorname{argmax}} \left(e^q_{t,k} + \ln P(q) \right) \tag{2.56}$$

2.6 Equipment Performance Model

Equipment performance can be treated as dynamical states analogously to the position and velocity states of the aircraft. Equipment can be split into 2 broad categories, those in which the functional state can be considered as discrete states corresponding to nominal performance and various failure modes, and those in which functional states are continuous, representing gradual failures and degradation of performance.

The dynamics of discrete failures are modelled using a continuous time Markov model within the general SHS. A Markov model consists of a set of discrete states that a system may occupy and a transition matrix containing the rate at which the system transitions between states. As the state of the system is not directly observable, we model the probability that the system is in a particular state conditioned on the available observations. The belief state dynamics of a basic continuous time Markov model with 2 discrete states is given in Equation 2.57. π_0 represents the belief probability that the system occupies the nominal state and π_1 represents the belief probability that the system is occupying the failure state. λ_f is the rate at which the system experiences a failure and transitions from the nominal state to the failure state, while λ_r is the rate at which the system recovers from the failure state to the nominal state.

$$\frac{d}{dt}\begin{bmatrix}\pi_0\\ \pi_1\end{bmatrix} = \begin{bmatrix}-\lambda_f & \lambda_r\\ \lambda_f & -\lambda_r\end{bmatrix}\begin{bmatrix}\pi_0\\ \pi_1\end{bmatrix} \tag{2.57}$$

This model can be used for failure prognostics by instantiating with an initial belief distribution and propagating forward in time. The probability a failure has occurred before the end of an exposure time is equivalent to the continuity parameter in reliability theory. The system reaches an equilibrium failure state belief probability equivalent to the availability parameter P_A in reliability theory.

$$\begin{aligned} 0 &= -\lambda_f\pi_0 + \lambda_r\pi_1 \\ &= -\lambda_f(1-\pi_1) + \lambda_r\pi_1 \end{aligned} \tag{2.58}$$

$$P_A \equiv \pi_{1_{eq}} = \frac{\lambda_f}{\lambda_f + \lambda_r} \tag{2.59}$$

Estimation of the state of the system is accomplished using a monitor. Monitors use observations associated with the functional state of system to diagnose whether a failure has occurred. Common requirements on monitors include alert time, the time τ_A between the occurrence of the failure transition and the monitor failure alert, and integrity, the probability P_I that a failure alert is not produced by the monitor within the alert time. Probability of a false alert may also be a consideration when nuisance is a significant concern. A monitor effectively changes our belief of the failure rate in the past. For times before $t - \tau_A$ the failure rate is $P_I\lambda_f$, while the failure rate stays at the unadjusted value otherwise. The resulting failure belief probability P_{fail} given no alert from the monitor is given in equation Equation 2.60. An approximation is given for the case of $P_A \approx 1$, $P_I \ll 1$, and $\frac{1}{\lambda_f} \gg \tau_A$

$$\begin{aligned} P_{fail} &= \frac{P_I\lambda_f}{P_I\lambda_f + \lambda_r}\, e^{-\tau_A(\lambda_f+\lambda_r)} + P_A(1 - e^{-\tau_A(\lambda_f+\lambda_r)}) \\ &\approx P_I P_A + \lambda_f\tau_A \end{aligned} \tag{2.60}$$

This failure model can be used for failure prediction within a given exposure time T_E.

The probability of a failure occurring before the end of the exposure time consists of monitor integrity failures, undetected failures occurring within the alert time, and failures occurring between the current time and the exposure end time. Assuming $\frac{1}{\lambda_f} \gg \tau_E$ along with the previous assumptions, we compute the approximate exposure time failure probability in equation Equation 2.61.

$$P_{fail} \approx P_I P_A + \lambda_f(\tau_A + T_E) \tag{2.61}$$

When a particular unmitigated failure probability P_T is targeted, the requirements on the various reliability parameters can be expressed by the constraint in equation Equation 2.62.

$$P_I P_A + \lambda_f(\tau_A + T_E) \leq P_T \tag{2.62}$$

Some monitors could rely on several sub-components to accomplish their task, a major example being reliance on the RP to monitor for faults. The resulting alert time and missed detection probability depends on the pilot response, C2 downlink, C2 uplink, and other intermediate equipment between onboard observations and execution of the risk mitigating procedure.

The total time between onboard observation and option execution is the sum of the time constants of each individual sub-component shown in equation Equation 2.63. Likewise, the total detection integrity probability is the product of the individual integrity and detection probabilities according to equation Equation 2.64. A single break in the system such as a sensor, transmitter, or display failure is enough to cause a failure of the entire RP monitor system if there is insufficient redundancy.

$$\tau_A = \sum_i \tau_A^i \tag{2.63}$$

$$P_I = 1 - \prod_i 1 - P_I^i \tag{2.64}$$

Lumping the performance of individual subsystems into the total performance of C2 link dependent monitor allows performance requirements to be set for the total system before being allocated to the individual systems.

Although the autonomous agent will often be interacting with equipment and trying to mitigate equipment failure conditions, the failure of other agents in the system to perform to their required standards can be treated equivalently using the concept of trust. Trust can be divided into deciding whether agents are acting with good intent, which is out of the scope of our safety analysis, and the reliability of agents with good intent to perform to expected standards. For example, in IMC, tower control is trusted to correctly communicate runway clearance to approaching aircraft, aircraft are trusted to follow Instrument Flight Rules (IFR) procedures and comply with Air Traffic Control (ATC) instructions, and regulators are trusted to keep obstacle clearance zones free of hazards. As in the case of equipment failures conditions, some type of monitor is required when the failure of agents to meet their entrusted performance standards results in significant risk. In Visual Meteorological Conditions (VMC), pilots are trusted to use line of sight visual contact to detect other aircraft and maintain separation. When IFR conditions reduce visual range, the responsibility for detecting and mitigating separation conflicts is entrusted to ATC. Systems such as the Traffic Collision Avoidance System (TCAS) may fill the role of a monitor for the event that ATC fails to detect conflicts or other aircraft fail to follow ATC instructions, though aircraft must be properly equipped, and pilots must be properly trained to respond to resolution advisories. Ultimately, both humans and equipment go through a certification process to prove they will meet the reliability and trustworthiness required for particular operations. The certification of humans and equipment sets failure distribution priors that decision making agents can use in their safety assessment. Trust in other agents can be updated when new information is acquired and deviations from their expected nominal

performance can be observed.

2.7 Multivariate Normal Sampling

Samples from a univariate normal distribution can be generated from using basic functions. Sampling from a multivariate normal distribution however requires more computation.

We would like to be able to express our sample X in the following form:

$$X = \Theta Y + \mu \tag{2.65}$$

Where $X \in \mathbb{R}^q$, $Y \in \mathbb{R}^q$, $\mu \in \mathbb{R}^q$, and $\Theta \in \mathbb{R}^{q \times q}$. Y is sampled from $\mathcal{N}(0, I)$, so that each element of Y is independently sampled from univariate normal $\mathcal{N}(0, 1)$. We want to simulate sampling X from $\mathcal{N}(\mu, \Sigma)$, where there is covariance and dependence between variables .

$$\mathbb{E}\big[(Y - \mathbb{E}[Y])\big] = 0 \tag{2.66}$$

$$\mathbb{E}\big[(Y - \mathbb{E}[Y])(Y - \mathbb{E}[Y])^T\big] = I \tag{2.67}$$

$$\begin{aligned}
\Sigma &= \mathbb{E}\big[(X - \mathbb{E}[X])(X - \mathbb{E}[X])^T\big] \\
&= \mathbb{E}\big[(\Theta Y + \mu - \mathbb{E}[\Theta Y + \mu])(\Theta Y + \mu - \mathbb{E}[\Theta Y + \mu])^T\big] \\
&= \mathbb{E}\big[(\Theta Y + \mu - \mu)(\Theta Y + \mu - \mu)^T\big] \\
&= \mathbb{E}\big[(\Theta Y)(\Theta Y)^T\big] \\
&= \mathbb{E}\big[\Theta Y Y^T \Theta^T\big] \\
&= \Theta \mathbb{E}\big[Y Y^T\big] \Theta^T \\
&= \Theta I \Theta^T \\
&= \Theta \Theta^T
\end{aligned} \tag{2.68}$$

Θ can be found using the Cholesky Decomposition of Σ. The Cholesky Decomposition of a Hermitian, positive-definite matrix A (as is Σ), produces lower triangular matrix L such that:

$$A = LL^* \tag{2.69}$$

When A contains only real values (as does Σ), $L^* = L^T$, thus the Cholesky Decomposition of Σ produces $\Theta = L$.

In order to compute the likelihood ratio of a set of parameters, Σ^{-1} must be computed. The inverse of a matrix requires less computation when it is in triangular form. The inverse of lower triangular Θ can be taken to simplify inverting Σ.

$$\Sigma^{-1} = (\Theta^{-1})^T(\Theta^{-1}) \tag{2.70}$$

2.8 Process Noise Sampling

The Stochastic Differential Equations modelling the state evolution and the observer functions modelling sensors both require Wiener process samples w_t and v_t.

$$dx(t) = f(x(t), q(t))dt + g(q(t), x(t))dw_t \tag{2.71}$$

$$y(t) = h(x(t), q(t)) + r(q(t), x(t))dv_t \tag{2.72}$$

When these equations are discretized in time and evaluated over a finite number of timesteps N, the continuous time Wiener process increments are approximated by m independent white noise time series of length N normalized by $\sqrt{dt}$.

$$\begin{aligned} \mathbb{E}[W_i^k] &= 0 \\ \mathbb{E}[W_i^k W_j^l] &= \delta(i-j)\delta(k-l) \qquad i,j = 0,\ldots,N-1 \qquad k,l = 0,\ldots,m-1 \end{aligned} \tag{2.73}$$

Each simulation run requires sampling $N \times m$ standard normal random variables for m white noise processes of length N. A large number of random variables may efficient rare event estimation infeasible, thus a method for dimensionality reduction is required.

2.8.1 Karhunen-Loeve Expansion

The Karhunen-Loeve (K-L) expansion can be used to represent a stochastic process as a linear combination of orthogonal functions $e_k(t)$, where the coefficients Z_k^l are independent standard normal random variables [42]. The general expression for a K-L expansion is given in Equation 2.74.

$$W_t^l = \sum_{k=0}^{N-1} Z_k^l e_k(t) \tag{2.74}$$

The Discrete Cosine Transform (DCT) is used to represent a finite series of data points as a sum of orthogonal cosine functions [43]. Coefficient represent the contribution of various frequency modes to the time series, from the 0-frequency mean component to the higher frequency modes. Slightly modified weighting results in orthonormal basis functions satisfing the requirements of the K-L expansion. The orthonormal weighted DCT-III is chosen for the K-L expansion of the white noise time series and is given in Equation 2.75. The first several modes of the orthonormal weighted DCT-III are depicted in Figure 2.3 for a time series of length 120.

$$W_t^l = \frac{1}{\sqrt{N}} Z_0^l + \sum_{k=1}^{N-1} Z_k^l \sqrt{\frac{2}{N}} \cos\left[\frac{\pi}{N}\left(t+\frac{1}{2}\right)k\right] \tag{2.75}$$

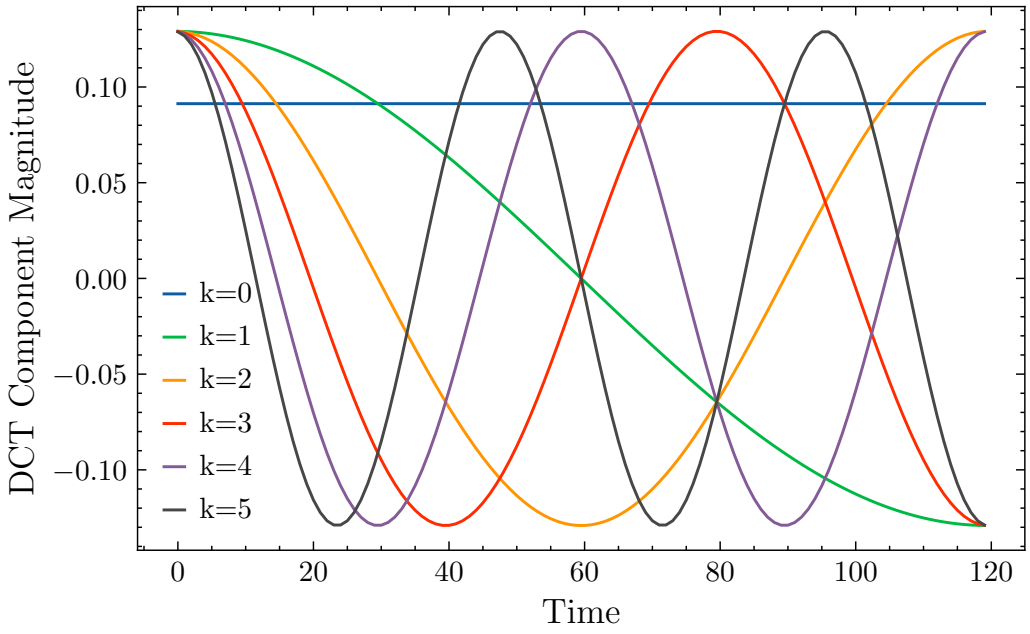

Figure 2.3: Component Magnitudes for the first 6 modes of DCT-III

2.8.2 Dimensionality Reduction

The $N \times m$ random variables required to describe the white noise processes in simulation may be reduced by instead sampling from the subset of K-L expansion coefficients Z_k^l most correlated with the safety metric. For example, high frequency observation noise components or high frequency gust components may have little effect on safety due to filtering and inertia, leaving the low frequency components or components at resonant frequencies primarily responsible for unsafe outcomes. To determine which K-L expansion modes explain the most variance in the safety metric, $1,000,000$ final approach simulations were run, each sampling 1920 independent standard random variables for the K-L expansion of 16 white noise time series of length 120. The output safety metric, describing the distance between the touchdown point and the runway threshold, was normalized and the covariance between K-L expansion coefficients and normalized safety metric was computed. The K-L expansion modes were sorted in descending order by squared covariance magnitude relative variance was computed by the cumulative sum of squared covariance magnitudes divided by the total sum of squared covariance magnitudes.

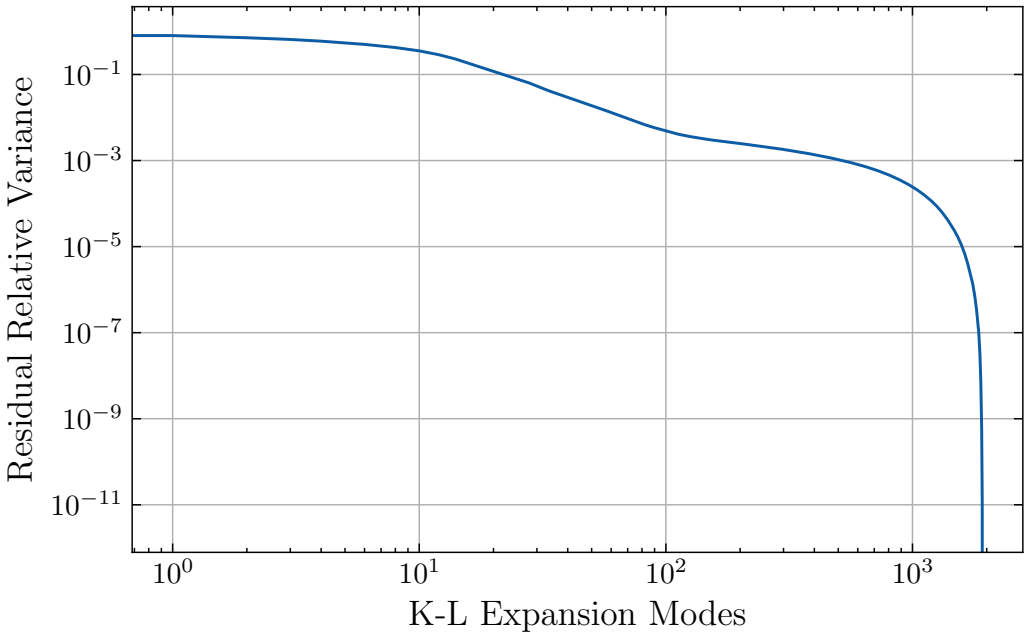

Figure 2.4: Residual Relative Variance vs Top K-L Expansion Modes considered

The relation between relative residual variance and the number of K-L expansion modes considered is depicted in Figure 2.4. 90% of the total variance can be explained with the top 22 modes and 99% of the total variance can be explained with the top 68 modes.

An example of the Safety Metric covariances for the K-L Expansion modes of the white noise contributing to the x-axis Dryden gust process is depicted in Figure 2.5. The $k = 3$ mode has an especially prominent covariance and modes above 40 have a negligible impact in comparison.

2.9 Event Time Sampling

Estimating accident probability using sampling based methods requires considering both the continuous state trajectory as well as the discrete state transition event times. Sampling purely continuous state trajectories can be accomplished by sampling from the initial belief state distribution and propagating the dynamics forward until the end of the risk exposure time using suitably sampled process noise. Including discrete state transition event times complicates this procedure. When the mean time to transition is much larger than the risk

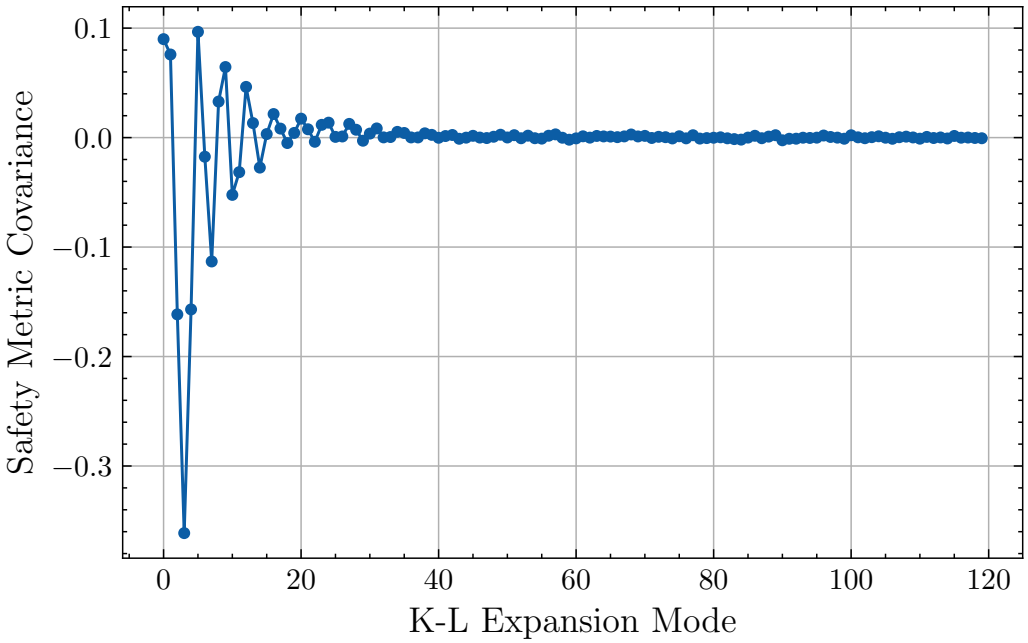

Figure 2.5: Safety Metric covariance vs K-L Expansion mode for x-axis Dryden gust process White noise

exposure time, an infeasible number of samples from the original transition time distribution will be required to observe a single event during the risk exposure time. Approximating the mapping between transition time and the hazard metric becomes difficult unless the sampling distribution is conditioned to ensure transitions happen within the exposure time. Sequences of transition events with state dependent transition rates may also need to be considered, complicating the conditional sampling problem. Discrete state spaces of even a moderate dimension may result in a combinatorial explosion as longer sequences of events are considered. A methodology is required to sample event sequences such that they occur during the exposure time while also addressing the tractability of searching the space of sequences.

For example, consider a final approach trajectory that starts out in a nominal state, but experiences a loss of primary navigation system integrity, engine failure, or wind shear event before touchdown. What sequences of events should be considered in the risk analysis? FAR § 25.1309 allows individual events or combination of events to be neglected if

their probability of occurrence during the exposure time is less than 1 in a billion [44]. This cutoff can be used to limit the set of event sequences that must be considered to demonstrate safety. The intent is to compute an upper bound on the probability that a sequence of discrete state transitions will occur during the exposure time as well as the distribution of transition times conditional on the sequence being observed within the exposure time. The transition event times can be equivalently described by occupation times, describing the elapsed time between transitioning into a discrete state and the first transition out of the discrete state. We first compute the cumulative distribution of discrete state occupation times in Equation 2.76. $\boldsymbol{\tau}$ indicates the vector of occupation times $[\tau_1, \cdots, \tau_K]$ for a sequence of length $K \geq 1$. $\mathbf{q}$ indicates the sequence of discrete states $[q_1, \cdots, q_K]$. The cumulative distribution gives the probability that the occupation time is less than τ_i for states 1 through $K-1$, while the occupation time of the final state is at least τ_K. For the case $K = 1$, the initial state is also the final state and the product is 1. Additionally, we have the probability that the state at time 0 is q_1, conditioned on the discrete belief state distribution $\hat{q}$ of the internal model.

$$P(\boldsymbol{\tau}) = P(q(\tau_K) = q_K | q(0) = q_K) P(q(0) = q_1 | \hat{q}) \prod_{i=1}^{K-1} P(q(\tau_i) = q_{i+1} | q(0) = q_i) \quad (2.76)$$

These probabilities may by obtained analytically if it is assumed discrete state transitions can be modelled by a continuous time Markov model. This requires transition rates to be constant with time. For transitions rates that are sufficiently small with respect to an exposure time, independent of continuous state, or slowly changing, assuming an upper bound on the transition rate is sufficient to ensure a conservative estimate of safety is obtained. $\mathbf{Q}$ is the transition rate matrix for the sequence of discrete states, a subset of the full Markov model. This is augmented with an additional state that absorbs any transitions that would deviate from the desired sequence. For $K \geq 2$, the state initializes in q_1 and transi-

tions into either q_2 or q_{K+1}, representing all other possible transitions that do not go to q_2. From the final state q_K, the state may only transition into q_{K+1}, additionally representing any transition out of q_K.

$$\mathbf{Q} = \begin{bmatrix} \mathbf{S} & \mathbf{s} \\ \mathbf{0} & 0 \end{bmatrix}^T \tag{2.77}$$

$$\mathbf{S} = \begin{bmatrix} \lambda_{q_1,q_1} & \lambda_{q_1,q_2} & & 0 \\ & \ddots & \ddots & \\ & & \lambda_{q_{K-1},q_{K-1}} & \lambda_{q_{K-1},q_K} \\ 0 & & & \lambda_{q_K,q_K} \end{bmatrix} \tag{2.78}$$

Where $\mathbf{s}_i = -\sum_{j=1}^{K} \mathbf{S}_{i,j}$. The ODE describing the evolution of the state occupation probabilities can be solved starting from an initial state. When τ is relatively small, the matrix exponential is approximately linear.

$$P(q(\tau) = q_j | q(0) = q_i) = \boldsymbol{\epsilon}_j^T e^{\mathbf{Q}\tau} \boldsymbol{\epsilon}_i \approx \boldsymbol{\epsilon}_j^T (\mathbf{I} + \mathbf{Q}\tau) \boldsymbol{\epsilon}_i \tag{2.79}$$

Where $\boldsymbol{\epsilon}_i$ is an indicator vector with 1 in index i and zeros elsewhere. $F(\boldsymbol{\tau})$ can be approximated by plugging Equation 2.79 into Equation 2.76.

$$P(\boldsymbol{\tau}) = \boldsymbol{\epsilon}_K^T e^{\mathbf{Q}\tau_K} \boldsymbol{\epsilon}_K P(q(0) = q_1 | \hat{q}) \prod_{i=1}^{K-1} \boldsymbol{\epsilon}_{i+1}^T e^{\mathbf{Q}\tau_i} \boldsymbol{\epsilon}_i \tag{2.80}$$

We would like to sample from the marginal probability density function of the first $K-1$ occupation times with the condition that state q_K is occupied for at least τ_K. The marginal distribution of occupation times is computed by taking the partial derivative of $P(\boldsymbol{\tau})$ with respect to τ_1 through τ_{K-1}.

$$f(\boldsymbol{\tau}) = \boldsymbol{\epsilon}_K^T e^{\mathbf{Q}\tau_K} \boldsymbol{\epsilon}_K P(q(0) = q_1 | \hat{q}) \prod_{i=1}^{K-1} \boldsymbol{\epsilon}_{i+1}^T \mathbf{Q} e^{\mathbf{Q}\tau_i} \boldsymbol{\epsilon}_i \tag{2.81}$$

For relatively small τ, $f(\boldsymbol{\tau})$ can be approximated with a first order expansion.

$$f(\boldsymbol{\tau}) \approx \boldsymbol{\epsilon}_K^T(\mathbf{I} + \mathbf{Q}\tau_K)\boldsymbol{\epsilon}_K P(q(0) = q_1|\hat{q}) \prod_{i=1}^{K-1} \boldsymbol{\epsilon}_{i+1}^T \mathbf{Q}(\mathbf{I} + \mathbf{Q}\tau_i)\boldsymbol{\epsilon}_i \tag{2.82}$$

To compute the probability of observing the sequence of discrete states during exposure time T_e, We would like to compute $f(\boldsymbol{\tau})$ conditioned on the following constraint:

$$\sum_{i=1}^{K} \tau_i = T_e \tag{2.83}$$

$$\sum_{i=1}^{K} x_i = \sum_{i=1}^{K} \frac{\tau_i}{T_e} = 1 \tag{2.84}$$

Each τ_i can be divided by T_e to phrase the constraint and $f(\boldsymbol{\tau})$ in terms of the normalized variables x_i, termed occupation proportions. A diagram of how a sequence of event times may be converted into occupation times and normalized by T_e to produce the occupation proportions is depicted in Figure 2.6.

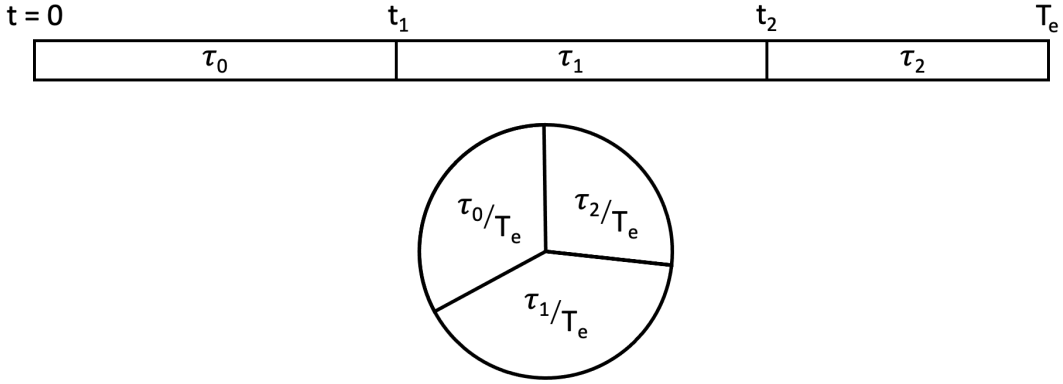

Figure 2.6: Event Time Normalization Procedure, event times t_i are converted into discrete state occupation times τ_i and normalized by T_e to produce occupation proportions that sum to 1.

$$P(\mathbf{x}) = \boldsymbol{\epsilon}_K^T e^{\mathbf{Q}T_e x_K} \boldsymbol{\epsilon}_K P(q(0) = q_1|\hat{q}) \prod_{i=1}^{K-1} \boldsymbol{\epsilon}_{i+1}^T e^{\mathbf{Q}T_e x_i} \boldsymbol{\epsilon}_i \tag{2.85}$$

$$f(\mathbf{x}) = \boldsymbol{\epsilon}_K^T e^{\mathbf{Q}T_e x_K} \boldsymbol{\epsilon}_K P(q(0) = q_1|\hat{q}) \prod_{i=1}^{K-1} \boldsymbol{\epsilon}_{i+1}^T \mathbf{Q} T_e e^{\mathbf{Q}T_e x_i} \boldsymbol{\epsilon}_i \tag{2.86}$$

For relatively small T_e, $f(\mathbf{x})$ can be approximated with a first order expansion.

$$f(\mathbf{x}) \approx \boldsymbol{\epsilon}_K^T (\mathbf{I} + \mathbf{Q}T_e x_k) \boldsymbol{\epsilon}_K P(q(0) = q_1|\hat{q}) \prod_{i=1}^{K-1} \boldsymbol{\epsilon}_{i+1}^T \mathbf{Q} T_e (\mathbf{I} + \mathbf{Q}T_e x_i) \boldsymbol{\epsilon}_i \tag{2.87}$$

By neglecting higher order terms and introducing a function of x_i, an upper bound on the distribution of $\mathbf{x}$ can be obtained.

$$f(\mathbf{x}) \leq P(q(0) = q_1|\hat{q}) \prod_{i=1}^{K-1} \lambda_{q_i,q_{i+1}} T_e \prod_{i=1}^{K} x_i^{\alpha_i - 1} \tag{2.88}$$

Where $\alpha_i = 1$. When conditioning this distribution on the constraint in Equation 2.84, it is apparent that $\mathbf{x}$ is drawn from a Dirichlet distribution with concentration parameters $\boldsymbol{\alpha}$. The occupation proportion of state q_K is defined as $x_K = 1 - \sum_{i=1}^{K-1} x_i$ and as $0 \leq x_i \leq 1$, $\mathbf{x}$ is a point sampled from the $K-1$ simplex.

$$\mathbf{x} \sim \text{Dir}(\boldsymbol{\alpha}) \tag{2.89}$$

When it is assumed that T_e is small, the transition rates are approximately constant during the exposure time and each $\alpha_i = 1$.

$$P\Big(\sum_{i=1}^{K-1} x_i \leq 1 \Big) \leq P(q(0) = q_1|\hat{q}) \prod_{i=1}^{K-1} \lambda_{q_i,q_{i+1}} T_e \frac{\prod_{i=1}^{K} \Gamma(\alpha_i)}{\Gamma\big(\sum_{i=1}^{K} \alpha_i\big)} \tag{2.90}$$

$$f(\mathbf{x}|\boldsymbol{\alpha}) \leq P(q(0) = q_1|\hat{q}) \prod_{i=1}^{K-1} \lambda_{q_i,q_{i+1}} T_e \frac{\prod_{i=1}^{K} \Gamma(\alpha_i)}{\Gamma\big(\sum_{i=1}^{K} \alpha_i\big)} \prod_{i=1}^{K} x_i^{\alpha_i - 1} \tag{2.91}$$

The Dirichlet distribution can be sampled using K gamma distributed random variables with variances α_i. Sampling these random variables G_i and normalizing by their sum gives

the Dirichlet distributed occupation proportion samples X_i for each discrete state i. Gamma distributed random variables may be sampled using acceptance-rejection sampling [45].

$$G_i \sim \text{Gamma}(\alpha_i, 1) \tag{2.92}$$

$$X_i = \frac{G_i}{\sum_{i=1}^{K} G_i} \tag{2.93}$$

CHAPTER 3

RARE EVENT ESTIMATION METHODOLOGY

The development of a suitable simulation model for a final approach scenario and a method for representing and sampling random input variables allows us to estimate the probability of unsafe outcomes. Accident event probabilities may be acquired analytically for sufficiently simple models. However, complex models such as the one developed in the previous chapter require sampling based Monte Carlo methods to estimate probabilities. Equation 3.1 gives the basic form of the Monte Carlo estimation problem, where the accident probability ℓ is estimated by sampling multivariate random variable X.

$$\ell = P(d(X) \leq 0) = \mathbb{E}[\mathbb{1}_{\{d(X)\leq 0\}}] \approx \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} \tag{3.1}$$

An accident event is given by $\{d(X) \leq 0\}$, meaning that the minimum distance d between the state trajectory and the hazard set is less then or equal to 0. $d(X)$ can be considered as a black-box function mapping the input sample to some output value of interest, though for our particular case we will call it the distance function. Quantifying the error of the Monte Carlo estimate is especially important when dealing with safety critical systems. The ratio of the error standard deviation to the probability estimate, called relative error, is often used to quantify the uncertainty in a Monte Carlo estimate. The relative error of the estimate is dependent on the sample variance and the number of samples. Lowering the relative error can be achieved by increasing the number of samples, or decreasing the variance of the sample measurements.

$$\text{Relative Error} = \frac{1}{\ell}\left(\frac{\mathbb{E}[(\mathbb{1}_{\{d(X)\leq 0\}} - \ell)^2]}{N-1}\right)^{\frac{1}{2}} \tag{3.2}$$

Difficulty arises when a rare event is being estimated, as is the case for FAA AC 25.1309

[44] safety requirements in which accident probabilities must be less than 1×10^{-9}. Thus, rare event estimation using basic Monte Carlo requires on the order of 1×10^{9} samples to observe an accident event once and several orders of magnitude more to produce an estimate with low relative error. This many samples is practically infeasible, even with a fast-time aircraft simulation. Efficient rare event estimation is a topic of interest in many fields, such as physics [46], weather forecasting [47], and air traffic systems [48]. Several techniques have been applied to the problem, including importance splitting [49] and sequential Monte Carlo [50], suited towards stationary dynamical systems. The technique of Importance Sampling [46] is suited for the final approach safety assessment application, where trajectories start from an initial position distribution.

3.1 Importance Sampling

Importance Sampling (IS) can be used to draw samples from a new distribution such that the resultant sample variance is reduced significantly. The number of samples required to obtain a probability estimate with equivalent relative variance may be reduced by many orders of magnitude. IS essentially involves shifting the sampling distribution to a new distribution such that the event is observed with a much higher probability on the order of 1×10^{-1}. The likelihood ratio of drawing a sample from the original distribution f vs the new distribution g is used as a weighting to estimate the probability of the rare event using fewer samples while maintaining the same relative error.

$$\hat{\ell} = \mathbb{E}_g\left[\mathbb{1}_{\{d(X)\leq 0\}}\frac{f(X)}{g(X)}\right] = \frac{1}{N}\sum_{i=1}^{N}\mathbb{1}_{\{d(X_i)\leq 0\}}\frac{f(X_i)}{g(X_i)} \tag{3.3}$$

Ideally, the IS distribution will have minimal relative variance and reliably estimate the event probability with relatively few samples. It can be shown that an IS distribution g^* exists which minimizes the variance in the probability estimate [51]. This optimal distribution is proportional to the original distribution f multiplied by an indicator function for the

event occurrence. g^* is normalized by the event probability ℓ.

$$g^*(x) = \frac{\mathbb{1}_{\{d(X)\leq 0\}} f(x)}{\ell} \tag{3.4}$$

The optimal IS distribution has minimal variance and can theoretically estimate the rare event probability with a single sample.

$$\ell = \frac{\mathbb{1}_{\{d(X_i)\leq 0\}} f(X_i)}{g^*(X_i)} \tag{3.5}$$

g^* is often difficult to directly sample from and must be approximated. A common technique for approximating g^* involves minimizing the cross-entropy between a parameterized distribution in the same family as f and a sampled estimate of g^*.

3.2 Cross-Entropy Method

The Kullback-Leibler Divergence, or Cross-Entropy (CE), measures the information theoretic 'distance' between two distributions [51]. Although not a true distance metric, it has properties that are useful for acquiring analytical results for distributions in the natural exponential family.

$$D_{KL}(p||q) = \mathbb{E}_p\left[\ln \frac{p(x)}{q(x)}\right] = \int_{-\infty}^{\infty} p(x) \ln \frac{p(x)}{q(x)} dx \tag{3.6}$$

For our purposes, we wish to minimize the cross-entropy between our IS distribution and the optimum density g^*.

We will assume that the original sampling function belongs to a family of Probability Density Functions (PDFs) denoted by $f(\cdot; u)$ with initial parameters u. The CE between g^* and $f(\cdot; v)$, for some parameters v, is defined as follows:

$$
\begin{aligned}
D_{KL}(g^*||f(\cdot;v)) &= \mathbb{E}_{g^*}\left[\ln \frac{g^*(X)}{f(X;v)}\right] \\
&= \int_{-\infty}^{\infty} g^*(x) \ln \frac{g^*(x)}{f(x;v)} dx \\
&= \int_{-\infty}^{\infty} g^*(x) \ln g^*(x) dx - \int_{-\infty}^{\infty} g^*(x) \ln f(x;v) dx
\end{aligned} \tag{3.7}
$$

$D_{KL}(g^*||f(\cdot;v))$ can be minimized with respect to v to find the parameters which most closely match the optimal IS distribution.

$$
\begin{aligned}
v = \underset{v}{\operatorname{argmin}}\, D_{KL}(g^*||f(\cdot;v)) &= \underset{v}{\operatorname{argmin}} - \int_{-\infty}^{\infty} g^*(x) \ln f(x;v) dx \\
&= \underset{v}{\operatorname{argmax}} \int_{-\infty}^{\infty} g^*(x) \ln f(x;v) dx
\end{aligned} \tag{3.8}
$$

The definition of the optimal IS distribution, dependent on initial distribution parameters u, may be substituted in for g^*.

$$
v = \underset{v}{\operatorname{argmax}} \int_{-\infty}^{\infty} \frac{\mathbb{1}_{\{d(x)\leq 0\}} f(x;u)}{\ell} \ln f(x;v) dx \tag{3.9}
$$

This is equivalent to maximizing an expectation sampled from $f(\cdot;u)$.

$$
v = \underset{v}{\operatorname{argmax}}\, \mathbb{E}_{f(\cdot;u)}[\mathbb{1}_{\{d(X)\leq 0\}} \ln f(X;v)] \tag{3.10}
$$

For a general case where samples are drawn from some intermediate distribution $f(\cdot;w)$ parameterized by w, IS can be used to obtain the following equivalent expectation.

$$
v = \underset{v}{\operatorname{argmax}}\, \mathbb{E}_{f(\cdot;w)}[\mathbb{1}_{\{d(X)\leq 0\}} \frac{f(X;u)}{f(X;w)} \ln f(X;v)] \tag{3.11}
$$

This can be estimated using Monte Carlo samples from $f(\cdot;w)$.

$$
v \approx \underset{v}{\operatorname{argmax}} \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w) \ln f(X_i;v) \tag{3.12}
$$

The likelihood ratio W at x is defined as follows.

$$W(x;u,w) = \frac{f(x;u)}{f(x;w)} \tag{3.13}$$

The gradient condition for minimizing $D_{KL}(g^*||f(\cdot;v))$ is given in Equation 3.14. v can be solved for when $f(\cdot;v)$ belongs to the natural exponential family, as is the case for the Multivariate Normal and Dirichlet distributions sampled for the simulation model. If $f(\cdot;v)$ is composed of the product of multiple independent distributions, it may be separated such that the CE minimization may be performed independently for the respective distribution parameters. This allows the Multivariate Normal parameters to be solved for independently from the Dirichlet parameters.

$$\begin{aligned} 0 &= \nabla_v \left[\frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w) \ln f(X_i;v) \right] \\ &= \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{\{S(X_i)\geq \gamma\}} W(X_i;u,w) \nabla_v \ln f(X_i;v) \end{aligned} \tag{3.14}$$

3.2.1 Cross-Entropy Minimization for Multivariate Normal Distribution

Equation 3.14 has an analytical solution when $f(\cdot;v)$ belongs to a Multivariate Normal distribution, the PDF of which is given in Equation 3.15.

$$f(x;v) = \frac{1}{\sqrt{(2\pi)^q \det \boldsymbol{\Sigma}}} \exp\left((x-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (x-\boldsymbol{\mu}) \right) \tag{3.15}$$

μ is the mean column vector belonging to $\mathbb{R}^q$ and Σ is the covariance matrix belonging to $\mathbb{R}^{q\times q}$. v can be expressed as $[\boldsymbol{\mu}, \boldsymbol{\Sigma}]^T$. Equation 3.14 involves the gradient of $f(\cdot;v)$ with respect to its parameters, which can be broken up as follows:

$$\nabla_v \ln f(X_i;v) = [\nabla_{\boldsymbol{\mu}} \ln f(X_i;v), \nabla_{\boldsymbol{\Sigma}} \ln f(X_i;v)]^T \tag{3.16}$$

These gradients have solutions as follows.

$$\nabla_{\boldsymbol{\mu}} \ln f(X_i; v) = \boldsymbol{\Sigma}^{-1}(X_i - \boldsymbol{\mu}) \tag{3.17}$$

$$\nabla_{\boldsymbol{\Sigma}} \ln f(X_i; v) = \frac{1}{2}\Big[-\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Sigma}^{-1}(X_i - \boldsymbol{\mu})(X_i - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}\Big] \tag{3.18}$$

With the proper substitutions, solutions for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ can be found.

$$\begin{aligned} 0 &= \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w) \nabla_{\boldsymbol{\mu}} \ln f(X_i; v) \\ &= \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w) \boldsymbol{\Sigma}^{-1}(X_i - \boldsymbol{\mu}) \\ &= \frac{\boldsymbol{\Sigma}^{-1}}{N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w)(X_i - \boldsymbol{\mu}) \\ &= \sum_{i=1}^{N} \mathbb{1}_{d(X_i)\leq 0\}} W(X_i; u, w)(X_i - \boldsymbol{\mu}) \end{aligned} \tag{3.19}$$

The means are solved for first.

$$\begin{aligned} \sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w) X_i &= \sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w) \boldsymbol{\mu} \\ &= \boldsymbol{\mu} \sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w) \end{aligned} \tag{3.20}$$

The means are simply the weighted average of the samples, where the weights are the likelihood ratios.

$$\boldsymbol{\mu} = \frac{\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w) X_i}{\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i; u, w)} \tag{3.21}$$

The covariance matrix is solved for next:

$$\begin{aligned}
0 &= \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w)\nabla_{\boldsymbol{\Sigma}} \ln f(X_i;v) \\
&= \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w)\frac{1}{2}\Big[-\boldsymbol{\Sigma}^{-1} + \boldsymbol{\Sigma}^{-1}(X_i-\boldsymbol{\mu}_k)(X_i-\boldsymbol{\mu}_k)^T\boldsymbol{\Sigma}^{-1}\Big] \\
&= \frac{1}{2N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w)\Big[-\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma} + \boldsymbol{\Sigma}\boldsymbol{\Sigma}^{-1}(X_i-\boldsymbol{\mu})(X_i-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}\Big] \\
&= \frac{1}{2N}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq \gamma_k\}} W(X_i;u,w)\Big[-\boldsymbol{\Sigma} + (X_i-\boldsymbol{\mu})(X_i-\boldsymbol{\mu})^T\Big]
\end{aligned} \tag{3.22}$$

$(X_i-\boldsymbol{\mu}_k)(X_i-\boldsymbol{\mu}_k)^T$ is effectively the covariance matrix of X_i with mean $\boldsymbol{\mu}_k$

$$\begin{aligned}
\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w)\text{cov}(X_i-\boldsymbol{\mu}) &= \sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w)\boldsymbol{\Sigma} \\
&= \boldsymbol{\Sigma}\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w)
\end{aligned} \tag{3.23}$$

The covariance matrix of the parameters is the covariance of the samples weighted by likelihood ratios.

$$\boldsymbol{\Sigma} = \frac{\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0n\}} W(X_i;u,w)\text{cov}(X_i-\boldsymbol{\mu})}{\sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq 0\}} W(X_i;u,w)} \tag{3.24}$$

The likelihood ratio is defined for the Multivariate Normal distribution.

$$\begin{aligned}
W(X_i;u,w) &= W(X_i;v_0,v_k) \\
&= W_{0,k}(X_i) \\
&= \left(\frac{\det\Sigma_0}{\det\Sigma_k}\right)^{\frac{1}{2}} \exp\Bigg((X_i-\boldsymbol{\mu}_k)^T\Sigma_k^{-1}(X_i-\boldsymbol{\mu}_k) \\
&\qquad - (X_i-\boldsymbol{\mu}_0)^T\Sigma_0^{-1}(X_i-\boldsymbol{\mu}_0)\Bigg)
\end{aligned} \tag{3.25}$$

3.2.2 Cross-Entropy Minimization for Dirichlet Distribution

The Cross-Entropy entropy minimization to solve for the optimal importance sampling distribution must also be applied to the Dirichlet distribution used to sample event times. The probability density of the Dirichlet distribution is given in Equation 3.26.

$$f(\mathbf{x}_i;\boldsymbol{\alpha}) = \frac{\Gamma\big(\sum_{j=1}^{K}\alpha_j\big)}{\prod_{j=1}^{K}\Gamma(\alpha_j)}\prod_{j=1}^{K}x_j^{(\alpha_j-1)} \tag{3.26}$$

The parameters which minimize the Cross-Entropy with an optimal sample population are obtained by the maximization problem in Equation 3.27.

$$\begin{aligned}
\boldsymbol{\alpha}^* &= \underset{\boldsymbol{\alpha}}{\operatorname{argmax}}\,\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}_{S(X_i)\geq\gamma}W(X_i;u,w)\ln f(\mathbf{x}_i;\boldsymbol{\alpha}) \\
&= \underset{\boldsymbol{\alpha}}{\operatorname{argmax}}\,\frac{1}{N}\sum_{i=1}^{N}\mathbb{1}_{S(X_i)\geq\gamma}W(X_i;u,w)\bigg(\ln\Gamma\big(\sum_{j=1}^{K}\alpha_j\big)-\sum_{j=1}^{K}\ln\Gamma(\alpha_j)+\sum_{j=1}^{K}(\alpha_j-1)\ln x_{i,j}\bigg) \\
&= \underset{\boldsymbol{\alpha}}{\operatorname{argmax}}\,\frac{\bar{W}}{N}\bigg(\ln\Gamma\big(\sum_{j=1}^{K}\alpha_j\big)-\sum_{j=1}^{K}\ln\Gamma(\alpha_j)+\sum_{j=1}^{K}(\alpha_j-1)\ln\bar{x}_j\bigg)
\end{aligned} \tag{3.27}$$

$\bar{W} = \sum_{i=1}^{N} W(X_i|u,w)$ and $\ln\bar{x}_j = \sum_{i=1}^{N}\mathbb{1}_{S(X_i)\geq\gamma}\frac{W(X_i|u,w)}{\bar{W}}\ln x_{i,j}$. This problem may be solved iteratively using the Newton-Raphson method [52].

$$\boldsymbol{\alpha}^{new} = \boldsymbol{\alpha}^{old} - \mathbf{H}^{-1}(F)\nabla F \tag{3.28}$$

$$F = \ln\Gamma\big(\sum_{j=1}^{K}\alpha_j\big)-\sum_{j=1}^{K}\ln\Gamma(\alpha_j)+\sum_{j=1}^{K}(\alpha_j-1)\ln\bar{x}_j \tag{3.29}$$

$\mathbf{H}$ indicates taking the Hessian matrix of a function. Partial derivatives of F involve the Digamma function Ψ and its derivative Ψ'. The gradient of F and the diagonal and off-diagonal Hessian entries are given below.

$$\frac{\partial F}{\partial \alpha_k} = \Psi\big(\sum_{j=1}^{K} \alpha_j\big) - \Psi(\alpha_k) + \ln \bar{x}_k \tag{3.30}$$

$$\frac{\partial^2 F}{\partial \alpha_k^2} = \Psi'\big(\sum_{j=1}^{K} \alpha_j\big) - \Psi'(\alpha_k) \tag{3.31}$$

$$\frac{\partial^2 F}{\partial \alpha_k \partial \alpha_l} = \Psi'\big(\sum_{j=1}^{K} \alpha_k\big) \tag{3.32}$$

Equation 3.28 is iterated until the size of the update is approximately 0.

3.3 Optimal Importance Sampling

In order to estimate the optimal IS distribution parameters using a reasonable number of samples, the probability l must be relatively large ($l \geq 10^{-2}$). For rare events ($l \leq 10^{-5}$), the IS parameters must be estimated using an iterative approach. When enough samples demonstrating an accident event are observed with $d(X_i) \leq 0$, CE minimization may be performed on the final sample population to obtain an approximation of the optimal IS distribution g^*.

3.3.1 Quantile Cross-Entropy Rare Event Importance Sampling

The quantile Cross-Entropy Importance Sampling algorithm iteratively adjusts the sample population by shifting the distance function by a factor γ_k, such that a target quantile of samples experience a redefined accident. At iteration k, $d_k(x) = d(x) - \gamma_k$ is used in place of $d(x)$, such that $\mathbb{E}_{f(\cdot;v_{k-1})}[\mathbb{1}_{\{d(X)\leq\gamma_k\}}] \geq 10^{-2}$. A suitable γ_k can be chosen from a sample of size N from $f(\cdot; v_{k-1})$. The best performing ρ quantile ($\rho \geq 10^{-2}$) is chosen to be greater than $d(X_\rho)) \leq \gamma_k$. The next IS parameter, v_k, is calculated using Equation 3.14.

$$0 = \frac{1}{N} \sum_{i=1}^{N} \mathbb{1}_{\{d(X_i)\leq\gamma_k\}} W(X_i; u, v_{k-1}) \nabla_{v_k} \ln f(X_i; v_k) \tag{3.33}$$

When $\gamma_k \approx 0$, the final IS parameters $v_{final} = v_{k-1}$ have been found. An IS Monte-Carlo simulation sampled using v_final can be used to achieve a rare event estimation with less variance using fewer samples.

$$\hat{\ell} = \frac{1}{N} \sum_{i=1}^{N_{MC}} \mathbb{1}_{\{d(X_i) \leq 0\}} W(X_i; u, v_{final}) \tag{3.34}$$

The Quantile Cross-Entropy Importance Sampling algorithm is summarized in Algorithm 1.

Algorithm 1 Quantile Cross-Entropy Rare Event Importance Sampling

$v_0 = u$ ▷ Initial distribution parameters
$k = 1$ ▷ Initial iteration
while $k \leq k_{max}$ **do**
  $\mathbf{X}_i \leftarrow f(\cdot; v_{k-1})$ ▷ N samples
  $\mathbf{d}_i \leftarrow d(\mathbf{X}_i)$ ▷ Simulate samples
  Sort $\mathbf{X}_i$ and $\mathbf{d}_i$ by $\mathbf{d}_i$ ascending ▷ Sort by distance metric
  $\gamma_k \leftarrow \mathbf{d}_{\lfloor \rho N \rfloor}$ ▷ ρ quantile distance metric
  if $\gamma_k \leq 0$ **then** ▷ Until accident ratio greater than ρ
    $v_{final} = v_{k-1}$
    break
  end if
  $v_k \leftarrow \text{argmax}_v \sum_i \mathbb{1}_{\{d(\mathbf{X}_i) \leq \gamma_k\}} W(\mathbf{X}_i; u, v_{k-1}) \ln f(\mathbf{X}_i; v)$ ▷ Optimize IS parameters
  $k+ = 1$ ▷ Iterate
end while
$\mathbf{X}_i \leftarrow f(\cdot; v_{final})$ ▷ Final N_{final} samples
$\mathbf{d}_i \leftarrow d(\mathbf{X}_i)$ ▷ Simulate final samples
$\hat{\ell} = \frac{1}{N_{final}} \sum_{i=1}^{N_{final}} \mathbb{1}_{\{\mathbf{d}_i \leq 0\}} W(\mathbf{X}_i; u, v_{final})$ ▷ Estimate Probability

3.3.2 Stochastic Differential Equation Rare Event Importance Sampling

The quantile based cross entropy method for finding the optimal Importance Sampling distribution is simple to implement but is prone to poor and inconsistent convergence behavior in practice. Other methods exist for finding optimal importance sampling distributions including use of Covariance Matrix Adaptation Evolution Strategy [53], Sequential Importance Sampling [54], and Ensemble Kalman Filtering [55]. The iterative estimation of the optimal importance sampling distribution can be treated as a stochastic drift-diffusion pro-

cess driving the sample distribution variable $\mathbf{X}_t$ from the initial distribution $f(\mathbf{x})$ towards an optimal equilibrium distribution $g^*(\mathbf{x})$. In practice, this process will take place in discrete time increments, however, it can be approximated in the continuous time limit by a Stochastic Differential Equation (SDE).

$$d\mathbf{X}_t = \boldsymbol{\mu}(\mathbf{X}_t, t)dt + \boldsymbol{\sigma}(\mathbf{X}_t, t)d\mathbf{W}_t \tag{3.35}$$

Where $\boldsymbol{\mu} \in \mathbb{R}^N$ describes the drift and $\boldsymbol{\sigma} \in \mathbb{R}^{N\times M}$ describes the diffusion with respect to an M-dimensional Wiener process $\mathbf{W}_t$. The Fokker-Planck equation for this process can be obtained, describing the evolution of the probability density $p(\mathbf{x}, t)$ for $\mathbf{X}_t$.

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} = -\sum_{i=1}^{N} \frac{\partial}{\partial x_i}[\mu_i(\mathbf{x}, t)p(\mathbf{x}, t)] + \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{\partial^2}{\partial x_i \partial x_j}[D_{ij}(\mathbf{x}, t)p(\mathbf{x}, t)] \tag{3.36}$$

Where $D = \frac{1}{2}\boldsymbol{\sigma}\boldsymbol{\sigma}^T$ is the diffusion tensor. At equilibrium, $p(\mathbf{x}, t)$ is stationary and the following equation holds.

$$\sum_{i=1}^{N} \frac{\partial}{\partial x_i}[\mu_i(\mathbf{x}, t)p(\mathbf{x}, t)] = \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{\partial^2}{\partial x_i \partial x_j}[D_{ij}(\mathbf{x}, t)p(\mathbf{x}, t)] \tag{3.37}$$

The optimal importance sampling distribution for observing the event $\{d(\mathbf{X}) \leq 0\}$ is defined by multiplying the initial distribution $f(\mathbf{x})$ by the event indicator function and normalizing by the event probability ℓ.

$$g^*(\mathbf{x}) = \frac{\mathbb{1}_{\{d(\mathbf{x})\leq 0\}} f(\mathbf{x})}{\ell} \tag{3.38}$$

To overcome the discontinuous indicator function, an approximate indicator function parameterized by width ϵ is assumed, which converges point-wise to the indicator function as $\epsilon \to 0$. The approximate indicator function is depicted in Figure 3.1.

$$G(\mathbf{x}, \epsilon) = \exp\left(-\frac{\max(d(\mathbf{x}), 0)^2}{2\epsilon^2}\right)$$
$$\lim_{\epsilon \to 0} G(\mathbf{x}, \epsilon) = \mathbb{1}_{\{d(\mathbf{x}) \leq 0\}} \tag{3.39}$$

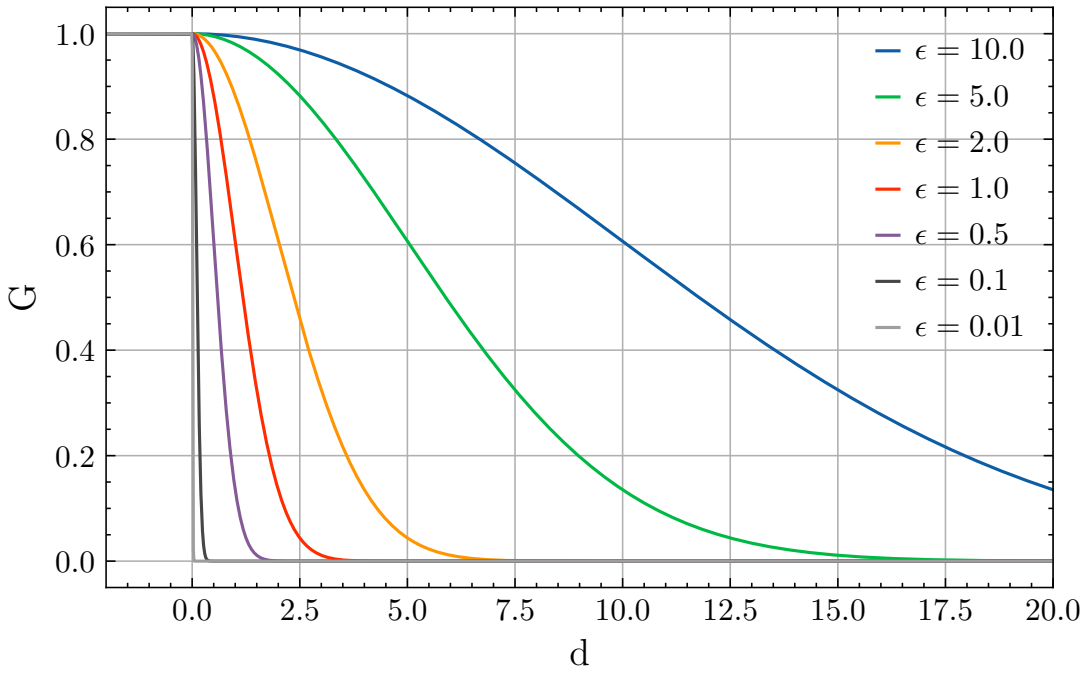

Figure 3.1: Indicator Function Approximation for several values of ϵ

We can solve for equilibrium drift and diffusion coefficients by substituting $\frac{G(\mathbf{x},\epsilon)f(\mathbf{x})}{\ell}$ for $p(\mathbf{x}, t)$ into Eq. Equation 3.37. The constant coefficient ℓ appears on both sides and can be removed.

$$\sum_{i=1}^{N} \frac{\partial}{\partial x_i}[\mu_i(\mathbf{x}, t)G(\mathbf{x}, \epsilon)f(\mathbf{x})] = \sum_{i=1}^{N}\sum_{j=1}^{N} \frac{\partial^2}{\partial x_i \partial x_j}[D_{ij}(\mathbf{x}, t)G(\mathbf{x}, \epsilon)f(\mathbf{x})] \tag{3.40}$$

Drift and diffusion terms may be found to satisfy this equation for certain classes of distributions. Multivariate normal distributions parameterized by a mean $\boldsymbol{\mu_0}$ and covariance matrix $\boldsymbol{\Sigma}$ are obtained if drift and diffusion terms satisfy the following.

$$\begin{aligned} \mathbf{X}_t &\sim \mathcal{N}(\boldsymbol{\mu_0}, \boldsymbol{\Sigma}) \\ \boldsymbol{\mu} &= \frac{\boldsymbol{\mu_0} - \mathbf{X}_t}{2} \\ \boldsymbol{\sigma}\boldsymbol{\sigma}^T &= \boldsymbol{\Sigma} \end{aligned} \tag{3.41}$$

$\boldsymbol{\sigma}$ can be defined using the Cholesky decomposition of $\boldsymbol{\Sigma}$. The Dirichlet distribution parameterized by $\boldsymbol{\alpha}$ can be obtained with drift and diffusion terms that satisfy the following constraints [56].

$$\begin{aligned} \mathbf{X}_t &\sim \text{Dir}(\boldsymbol{\alpha}) \\ \mu_i &= \frac{b_i}{2}[S_i X_{Nt} - (1 - S_i)X_{it}] \\ \sigma_{i,i} &= \sqrt{\kappa_i X_{it} X_{Nt}} \\ \alpha_i &= \frac{b_i}{\kappa_i} S_i \qquad i = 1, \ldots, N-1 \\ \alpha_N &= \frac{b_1}{\kappa_1}(1 - S_1) = \cdots = \frac{b_{N-1}}{\kappa_{N-1}}(1 - S_{N-1}) \end{aligned} \tag{3.42}$$

Where $b_i > 0$,$\kappa_i > 0$, and $0 < S_i < 1$. The final variable X_{Nt} is constrained by $X_{Nt} = 1 - \sum_{i=1}^{N-1} X_{it}$. Off-diagonal entries of $\boldsymbol{\sigma}$ are set to 0.

When $f(\mathbf{x})$ is a Multivariate Normal distribution, the indicator function biases the drift with the following term.

$$\Delta\boldsymbol{\mu}(\mathbf{x}, \epsilon) = \frac{\Delta\hat{\boldsymbol{\mu}}(\mathbf{x})}{\epsilon^2} = -\boldsymbol{\Sigma}\frac{\max(d(\mathbf{x}), 0)}{\epsilon^2}\nabla_{\mathbf{x}} d(\mathbf{x}) \tag{3.43}$$

Where $\nabla_{\mathbf{x}} d(\mathbf{x})$ indicates the Jacobian of $d(\mathbf{x})$. When $f(\mathbf{x})$ is a Dirichlet distribution, the indicator function biases the drift with the following term.

$$\Delta\mu_i = \frac{\kappa_i \mathbf{x}_i \mathbf{x}_N \max(d(\mathbf{x}), 0)(\partial_N d(\mathbf{x}) - \partial_i d(\mathbf{x}))}{\epsilon^2} \qquad i = 1, \ldots, N-1 \tag{3.44}$$

Where ∂_i indicates the partial derivative with respect to $\mathbf{x}_i$. The differentiable distance function required for computing the drift bias is generally unavailable and may be approximated

using a surrogate model $\tilde{d}(\mathbf{x}, \boldsymbol{\theta})$ parameterized by $\boldsymbol{\theta}$. A quadratic model is satisfactory for sufficiently smooth and unimodal distance functions. $\boldsymbol{\theta}$ may be learned via recursive least squares regression on the ensemble of distance function samples [57].

$$\begin{aligned} d(\mathbf{x}) \approx \tilde{d}(\mathbf{x}, \boldsymbol{\theta}) &= [1 \;\; \mathbf{x}_1 \;\; \dots \;\; \mathbf{x}_N \;\; \mathbf{x}_1^2 \dots \;\; \mathbf{x}_N^2]\boldsymbol{\theta} \\ \boldsymbol{\theta} &= \underset{\theta}{\operatorname{argmin}} \sum_{j=1}^{N} (\tilde{d}(\mathbf{X}^j, \boldsymbol{\theta}) - \mathbf{d}^j)^2 \end{aligned} \tag{3.45}$$

The resulting SDE may be discreteized in time using the Euler–Maruyama method [58]. Indicator width ϵ and time step Δt and are the remaining free parameters when the unbiased drift and diffusion process parameters are defined. In practice, both ϵ and Δt should be initialized large and decrease until convergence is observed.

The choice of these parameters may be guided by simultaneously controlling the expected change in mean and variance of the distance metric.

$$\begin{aligned} \Delta \mathbf{x} &= (\boldsymbol{\mu}(\mathbf{x}) + \Delta\boldsymbol{\mu}(\mathbf{x}, \epsilon))\Delta t + \boldsymbol{\sigma}(\mathbf{x})\mathbf{w}\sqrt{\Delta t} \\ &= (\boldsymbol{\mu}(\mathbf{x}) + \frac{\Delta\hat{\boldsymbol{\mu}}(\mathbf{x})}{\epsilon^2})\Delta t + \boldsymbol{\sigma}(\mathbf{x})\mathbf{w}\sqrt{\Delta t} \end{aligned} \tag{3.46}$$

$$\begin{aligned} \Delta d(\mathbf{x}) &= \nabla_{\mathbf{x}} d(\mathbf{x})^T \Delta \mathbf{x} \\ &= \nabla_{\mathbf{x}} d(\mathbf{x})^T (\boldsymbol{\mu}(\mathbf{x}) + \Delta\boldsymbol{\mu}(\mathbf{x}, \epsilon))\Delta t + \nabla_{\mathbf{x}} d(\mathbf{x})^T \boldsymbol{\sigma}(\mathbf{x})\mathbf{w}\sqrt{\Delta t} \end{aligned} \tag{3.47}$$

The first constraint places the variance of the change in distance metric due to diffusion on the same scale as the ϵ^2. This ensures that the sample distances stay on the order of ϵ during each iteration without excessive deviations.

$$\mathbb{E}[\left(\sqrt{\Delta t}\mathcal{N}\left(0, \nabla_{\mathbf{x}} d(\mathbf{x})^T \Sigma(\mathbf{x}) \nabla_{\mathbf{x}} d(\mathbf{x})\right)\right)^2] \approx \epsilon^2 \tag{3.48}$$

$$\Delta t \approx \frac{\epsilon^2}{\mathbb{E}[\nabla_{\mathbf{x}} d(\mathbf{x})^T \Sigma(\mathbf{x}) \nabla_{\mathbf{x}} d(\mathbf{x})]} \tag{3.49}$$

The second constraint requires a reduction in mean distance metric by a factor $0 < \gamma <$

1 during each iteration.

$$\mathbb{E}[\nabla_{\mathbf{x}} d(\mathbf{x})^T(\boldsymbol{\mu}(\mathbf{x}) + \frac{\Delta\hat{\boldsymbol{\mu}}(\mathbf{x})}{\epsilon^2})\Delta t] \approx (1-\gamma)\mathbb{E}[d(\mathbf{x})] \tag{3.50}$$

The SDE propagation and sample simulation are iterated until ρ fraction of samples experience an accident. At this point, we treat the sample population as the optimal IS distribution and learn final distribution parameters v_{final} to minimize the CE with the optimal distribution. A final set of samples is simulated and the accident probability is evaluated. Algorithm 2 summarizes the iterative sampling process.

Algorithm 2 SDE Rare Event Importance Sampling

$\mathbf{X}_i \leftarrow f(\cdot; u)$ ▷ Initial N samples
$\mathbf{d}_i \leftarrow d(\mathbf{X}_i)$ ▷ Simulate initial samples
while $\sum(\mathbf{d}_i \leq 0) < \rho N$ **do** ▷ Until accident ratio greater than ρ
  $\theta \leftarrow \text{argmin}_\theta \sum(\tilde{d}(\mathbf{X}_i, \theta) - \mathbf{d}_i)^2$ ▷ Learn surrogate model
  $\epsilon, dt \leftarrow \epsilon(\mathbf{d}), dt(\mathbf{d})$ ▷ Update parameters
  $\mathbf{X}_i \leftarrow \mathbf{X}_i + (\mu(\mathbf{X}_i) + \Delta\mu(\mathbf{X}_i, \epsilon, \theta))dt + \sigma(\mathbf{X}_i)\mathcal{N}(\mathbf{0}, \mathbb{I})\sqrt{dt}$ ▷ Propagate SDE
  $\mathbf{d}_i \leftarrow d(\mathbf{X}_i)$ ▷ Simulate samples
end while
$v_{final} \leftarrow \text{argmax}_v \sum_i \mathbb{1}_{\{d(\mathbf{X}_i)\leq 0\}} W(\mathbf{X}_i; u, u) \ln f(\mathbf{X}_i; v)$ ▷ Optimize IS parameters
$\mathbf{X}_i \leftarrow f(\cdot; v_{final})$ ▷ Final N_{final} samples
$\mathbf{d}_i \leftarrow d(\mathbf{X}_i)$ ▷ Simulate final samples
$\hat{\ell} = \frac{1}{N_{final}} \sum_{i=1}^{N_{final}} \mathbb{1}_{\{\mathbf{d}_i \leq 0\}} W(\mathbf{X}_i; u, v_{final})$ ▷ Estimate Probability

CHAPTER 4
DESIGN PARAMETER SAFETY ASSESSMENT FOR PASSIVE RISK MITIGATION

The development of a parameterized model of an unmanned aircraft on final approach and a method for evaluating the risk associated with a given set of design parameters allows the evaluation of the range of parameters which will satisfy the imposed safety requirements and ensure risk is passively mitigated. This is in contrast to active risk mitigation which requires decisions to be made during operation.

The aim of this chapter is to fill in the variable parameters not yet fixed or assumed in the system design. This can be decomposed into safety evaluations on the various failure conditions, beginning with the nominal operating case.

4.1 Safety Assessment Methodology

Free parameters of the system and procedure design must be set such that safety requirements are satisfied. A full safety assessment involves three broad tasks, of which the last is primarily addressed here.

i) Enumerate possible failure conditions

ii) Determine the probability of each failure condition

iii) Compute the probability of an accident given each failure condition

Techniques for enumerating failure condition of a detailed design, such as Fault Tree Analysis, start with many individual equipment failure events or environmental events specific to the system in questions. The tree structure of these methodologies computes the probability of high-level system failures dependent on low-level failure conditions. This

allows a multitude of specific system details to be abstracted into a reduced number of high level failure conditions. The basic ways in which low-level failure conditions can be composed to produce higher-level failure conditions may be modelled using AND gates, which require a combination of low-level failure events to trigger a higher-level failure, and OR gates, which require only one of a set of low-level failure events to trigger a higher-level failure [9].

An AND gate is depicted in Figure 4.1, showing individual events e_i with probabilities P_i combining to form a higher-level failure e_{AND} with probability P_{AND}. Equation 4.1 computes this probability with e_{AND} as the intersection of events e_i. This is approximated for the case that each low-level event is independent, however, this is not necessarily the case.

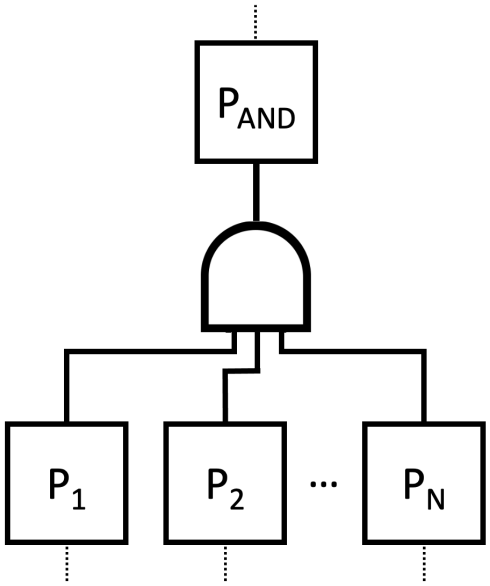

Figure 4.1: Fault Tree AND Gate

$$P_{\text{AND}} = P\left(\bigcap_{i=1}^{N} e_i\right) \approx \prod_{i=1}^{N} P_i \tag{4.1}$$

An OR gate is depicted in Figure 4.2, showing individual events e_i with probabilities P_i

combining to form a higher-level failure e_{OR} with probability P_{OR}. Equation 4.2 computes this probability with e_{OR} as the union of events e_i. The Bonferroni inequality is used to provide an upper bound for the union probability [59].

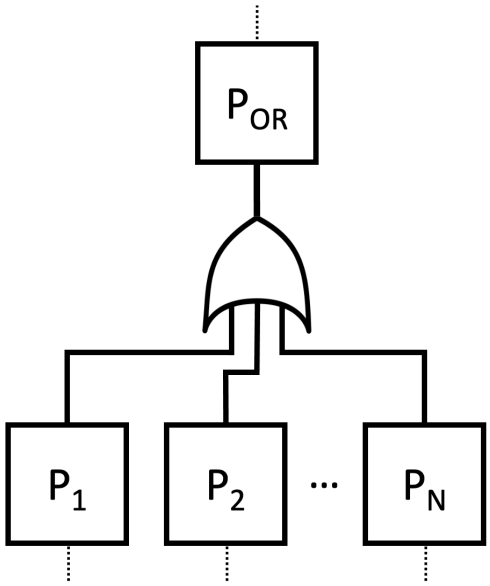

Figure 4.2: Fault Tree OR Gate

$$P_{\text{OR}} = P\left(\bigcup_{i=1}^{N} e_i\right) \leq \sum_{i=1}^{N} P_i \tag{4.2}$$

AC 120-28D [31] requires several high-level failure cases to consider for the approach and landing safety assessment, namely hazardously misleading navigation system errors, severe wind shear events, and critical engine failures. Other possible failure conditions include control system failures and runway incursions. We will assume control systems have adequate redundancy such that they remain fail operational. Runway incursion are assumed to be adequately mitigated by surface control procedures such that they are unlikely to occur while the aircraft is below a minimum decision height. Runway incursion could potentially be detected by computer vision onboard the aircraft, but this requires sufficient integrity from the computer vision system to minimize missed detections.

Failure conditions may be combined to create further failure conditions composed of a sequence of events. This can be treated in the discrete state transition event sequence framework developed in Chapter II. The system starts in an initial discrete state, such as the nominal fault free state, then transitions into further discrete states as failure condition events occur. The safety assessment is confined to a finite risk exposure time, so only a finite number of events are considered in the sequence. The branching tree structure of the event sequence is depicted in Figure 4.3. Events occur with rate λ_i until the sequence is terminated.

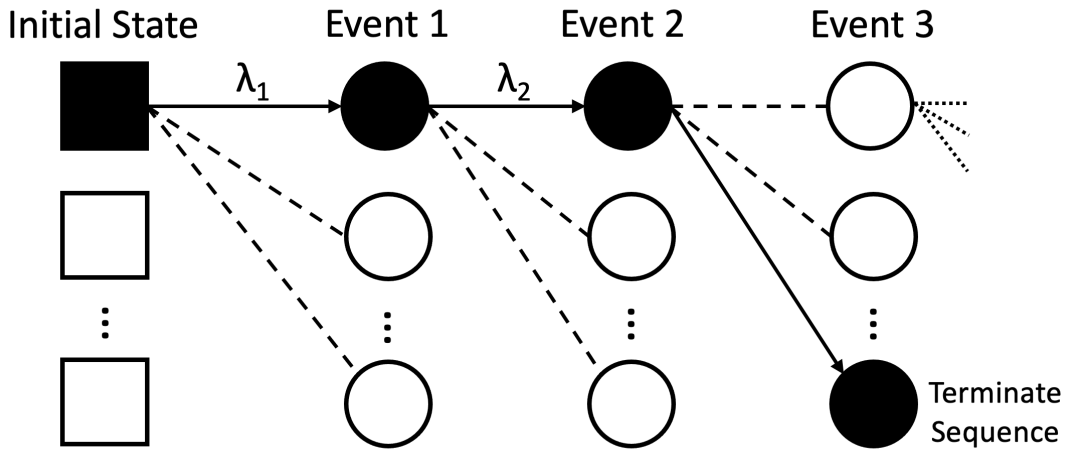

Figure 4.3: Failure Condition Event Tree

The probability of a given sequence of events transitioning through a sequence of discrete states is given in Equation 4.4 using results derived Chapter II. $\mathbf{q}$ indicates the sequence of discrete states $[q_1, \cdots, q_K]$. $P(\mathbf{q}|T_e)$ is the probability that the sequence of discrete states $\mathbf{q}$ occurs within the exposure time T_e.

$$P(\mathbf{q}|T_e) \leq P(q(0) = q_1)\frac{T_e^{K-1}}{\Gamma(K)}\prod_{i=1}^{K-1}\lambda_{q_i,q_{i+1}} \tag{4.3}$$

Upper bounds on the event rates and initial state probability are needed to evaluate this probability. In the absence of any information, the probability should conservatively be set to 1. Conversely, system requirements may be set choosing values of λ that limit the

product of the event sequence probability and the conditional accident probability to the desired level of safety.

$$P(acc|\mathbf{q}, T_e)P(\mathbf{q}|T_e) \leq 1 \times 10^{-9} \tag{4.4}$$

This content of this Chapter directly addresses estimation of accident probability using the rare event estimation methodology presented in Chapter III. The Importance Sampling Monte Carlo method used to compute accident probabilities is a fundamentally stochastic method and has an estimated relative error. In the large sample size limit, this error is approximately normal. The problem of errors in the safety critical accident probability estimate may be resolved by treating this uncertainty as another failure condition. The 1×10^{-9} quantile error upper bound, corresponding to approximately 6 standard deviations of the estimation error, may be assumed for the accident probability to account for this failure condition.

4.2 Arrival and Approach Scenario

The parameters feeding into a particular safety assessment case can be split into several types depending on their role in the scenario. Parameter type also informs how they are instantiated for a particular safety assessment, whether they are fixed into the system or procedure design, assumed for the purpose of a conservative safety assessment, variable system design parameters, or state variables sampled in each simulation. The different simulation parameter types and parameter sources are listed in Table 4.1.

The methodology used for the design safety assessment initializes the simulation the during the descent and arrival phase of flight and stops the simulation before the final 200 ft of the final approach, the point at which safety becomes critical for instrument approach procedures. This initial simulation serves to 'warm-up' the internal states of the guidance, navigation, and control system, and damp out transients in the aircraft dynamics. A Boeing

Table 4.1: Scenario Parameters

Parameter Type	Parameter Name	Parameter Source
Initial State	x, y, z	Sampled from belief state
	v_x, v_y, v_z	
	ϕ, θ, ψ	
	p, q, r	
	configuration	Landing Configuration
	mass	737-800 MLW, 66,350 kg
Belief State	x, y, z	warm up simulation
	v_x, v_y, v_z	
	$\epsilon_{GM,x}, \epsilon_{GM,y}, \epsilon_{GM,z}$	
Process Noise	Gust Processes	K-L Expansion Sampling
	GPS GM Process	
	Measurement Noise	
Discrete State Transitions	Event Sequence	failure condition tree
	Event Probability	conservatively assume 1
	Event Magnitudes	overbound or variable
Procedure Parameters	Runway Aimpoint	variable
	Glide Slope	3^o
	Approach Speed	1.3 Stall Speed
	Runway Elevation	304.8 m
	Runway Size	150 ft by 10000 ft
	Terrain Profile	assumed flat
Operational Conditions	Wind Profile	15 kt cross by 10 kt tail [31]
	Gust Strength	10% mean wind [31]
Design Requirements	GPS Accuracy	WAAS GPS observed performance
	Radar Altimeter Accuracy	5% height + 1m
	IMU Accuracy	0.01 m/s^2
Fixed Design Parameters	Aircraft Aerodynamics	737-800 TASAT model [30][29]
	Engine Characteristics	
	Control Characteristics	

737-800 at Maximum Landing Weight (MLW) is used as the subject of the safety assessments. Aerodynamic, engine, and control parameters are defined using the parameters verified in the TASAT model [29]. The initial states at the start of the safety assessment are sampled from the belief state distribution. AC 120-28D requires safety to be ensured in a defined range of operational conditions, specifically mean head winds of 25 knots, crosswinds of 15 knots, and tail winds of 10 knots [31], with a gust intensity of 10% mean wind at 20 feet. A worst case mean wind with 15 knot crosswind and 10 knot tailwind is used to provide a conservative safety assessment. A runway of standard size at 1000 feet above mean sea level on flat terrain is used for the safety assessment. The runway aimpoint location relative to the runway threshold is used as a variable in the procedure design to set the baseline undershoot risk level. Variable aimpoint and glide slope are features of GPS based procedures as they are not determined by a fixed antenna location. Glide slope, however, is fixed at 3^o for the safety assessment. Process noise is sampled using the Karhunen-Loeve (K-L) expansion technique and event times are sampled using the Dirichlet distribution technique, both presented in Chapter II.

Although the last 200 feet of the final approach are the primary interest of the safety assessment, a full arrival and approach procedure is simulated beforehand to 'warm-up' the internal states of the aircraft simulation. A Continuous Descent Arrival/Approach (CDA) procedure is defined, with a lateral path and vertical profile similar to an Area Navigation (RNAV) procedure. CDA procedures reduce fuel burn, emissions, and noise exposure by reducing thrust requirement with a continuous rate of descent [60][61]. The lateral path is depicted in Figure 4.4, featuring a 45^o turn onto what would nominally be the downwind approach segment followed by a 180^o turn to intercept the final approach. It is worth reiterating that the aircraft is experiencing a tail wind during final approach as opposed to the nominally desired head wind to serve as a worst case scenario.

The altitude profile of the arrival and approach, starting partway through the arrival descent is depicted in Figure 4.5. The vertical profile has constant 3^o glide slope except

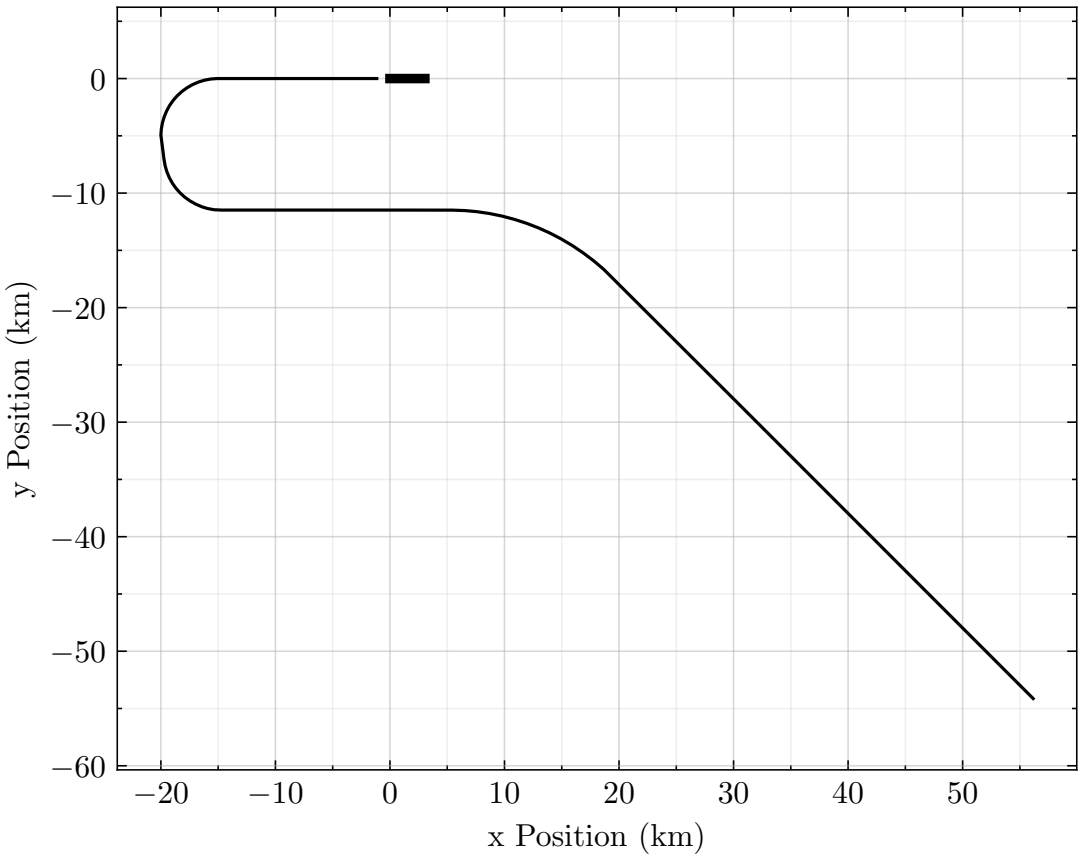

Figure 4.4: Arrival and Approach Lateral Path

for a brief segment of reduced glide slope that serves to help decelerate the aircraft to meet the 250 knot (128 m/s) calibrated air speed limitation imposed by air traffic control below 10000 feet (3048 m). The corresponding calibrated airspeed profile is depicted in Figure 4.6, with numerous steps down in airspeed as flaps are deployed and the aircraft stabilizes on the final approach at the reference approach speed.

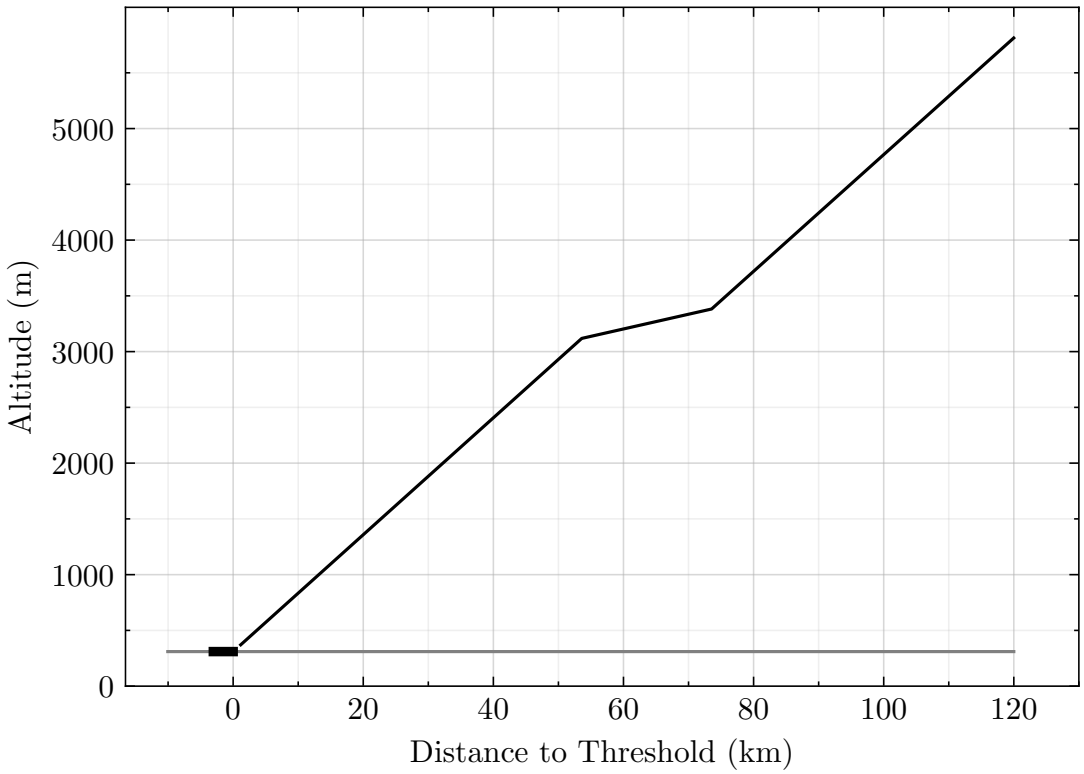

Figure 4.5: Arrival and Approach Altitude Profile

4.3 Nominal Condition Safety Assessment

The first safety assessment analyzes the nominal system condition to set free parameters and passively mitigate risk by ensuring the 1×10^{-9} level of safety when there are no off-nominal failure conditions. Mean accident probability estimates and 6σ upper bounds are plotted. A risk exposure time of 30 s is used with a simulation time step of 0.25 s.

Safety assessment plots should be read by looking for the point where the 6σ accident upper bound bar crosses 1×10^{-9}. This is the limiting parameter value for which safety is

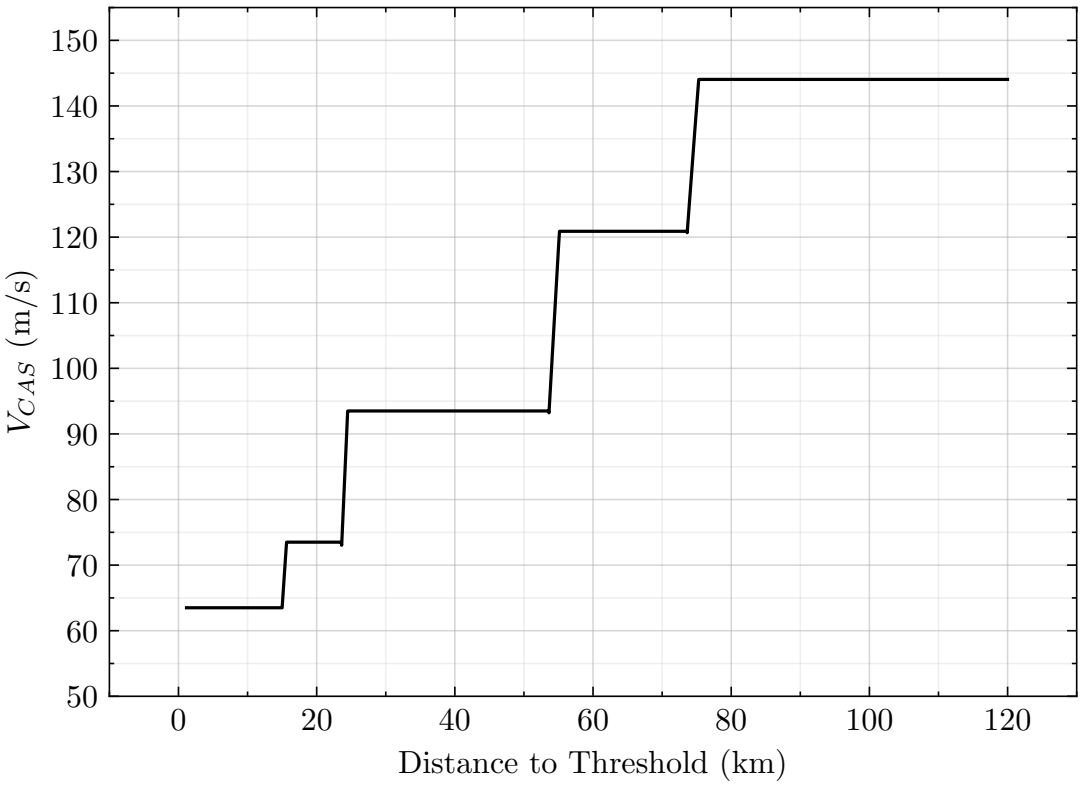

Figure 4.6: Arrival and Approach Calibrated Airspeed Profile

passively maintained for the particular system configuration and failure condition.

4.3.1 Runway Aimpoint

The remaining free parameter in the procedure design, the runway aimpoint, should be minimized relative to the runway threshold to maximize remaining runway distance, but kept large enough that the probability of an undershoot accident is kept below 1×10^{-9}. The runway undershoot accident probability is computed across a range of aimpoint locations, using WAAS GPS navigation and the full wind model.

Results plotted in Figure 4.7 suggest that a runway aimpoint 210 m or greater is required under nominal conditions. Major sources of uncertainty in this scenario include turbulence and vertical navigation uncertainty.

Results plotted in Figure 4.8 suggest that under conditions without gust turbulence, runway aimpoint could potentially be reduced to 185 m while maintaining safety.

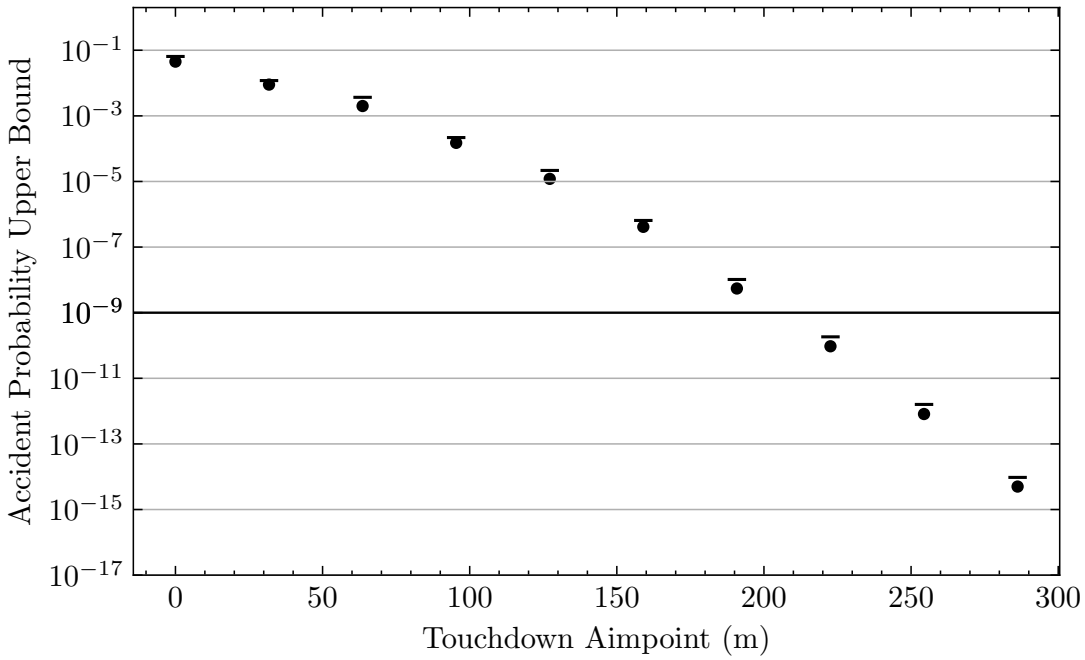

Figure 4.7: Accident Risk vs Runway Aimpoint, WAAS GPS and nominal wind model

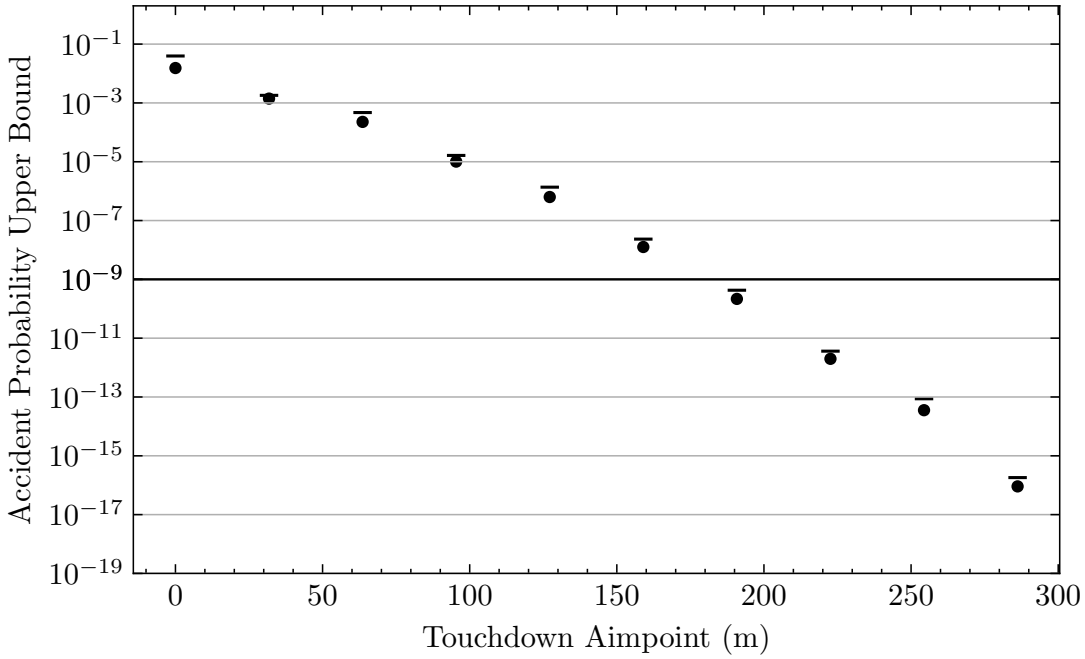

Figure 4.8: Accident Risk vs Runway Aimpoint, WAAS GPS and no turbulence

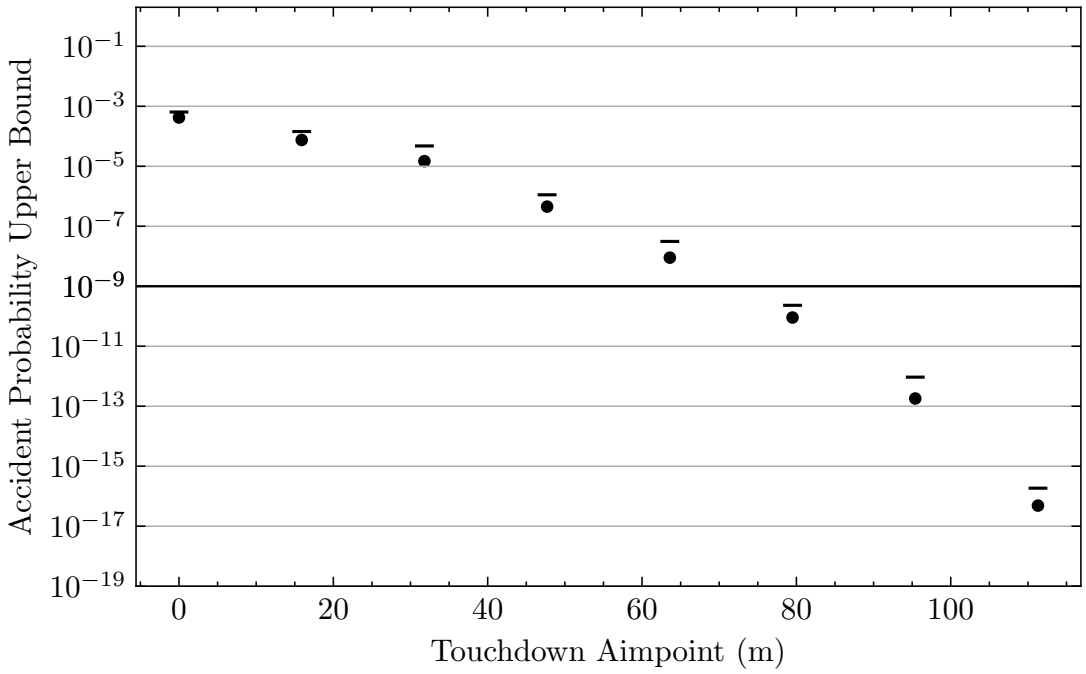

Figure 4.9: Accident Risk vs Runway Aimpoint, Radar Altimeter augmented WAAS GPS and nominal wind model

Vertical navigation error can be reduced by integrating a radar altimeter into the state estimation. Results plotted in Figure 4.9 suggest that the runway aimpoint could potentially be reduced to 75 m if the aircraft is equipped with a radar altimeter.

4.3.2 Decision Height

Decision Height is defined by the minimum height at which a missed approach may be initiated before the probability of touching the ground exceeds 1×10^{-9}. Although merely touching the ground does not necessarily constitute an accident, it is assumed that it is an accident in this case due to some rare event such as a runway incursion.

Results plotted in Figure 4.10 suggest the minimum DH with WAAS GPS and the nominal wind model to be 26 m, or just about 85 ft, corresponding to a CAT II ILS approach. As in the aimpoint analysis, two major sources of uncertainty contribute to the accident probability, turbulence and vertical navigation error.

Figure 4.11 shows that the minimum DH is reduced slightly when gust turbulence is

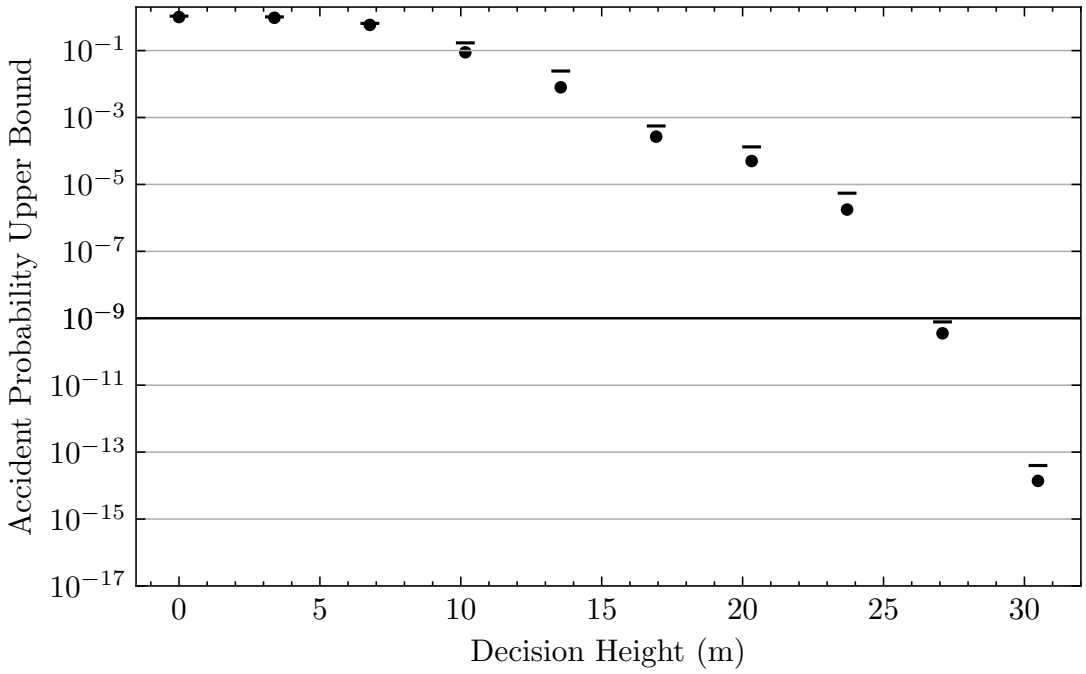

Figure 4.10: Accident Risk vs Decision Height, WAAS GPS and nominal wind model

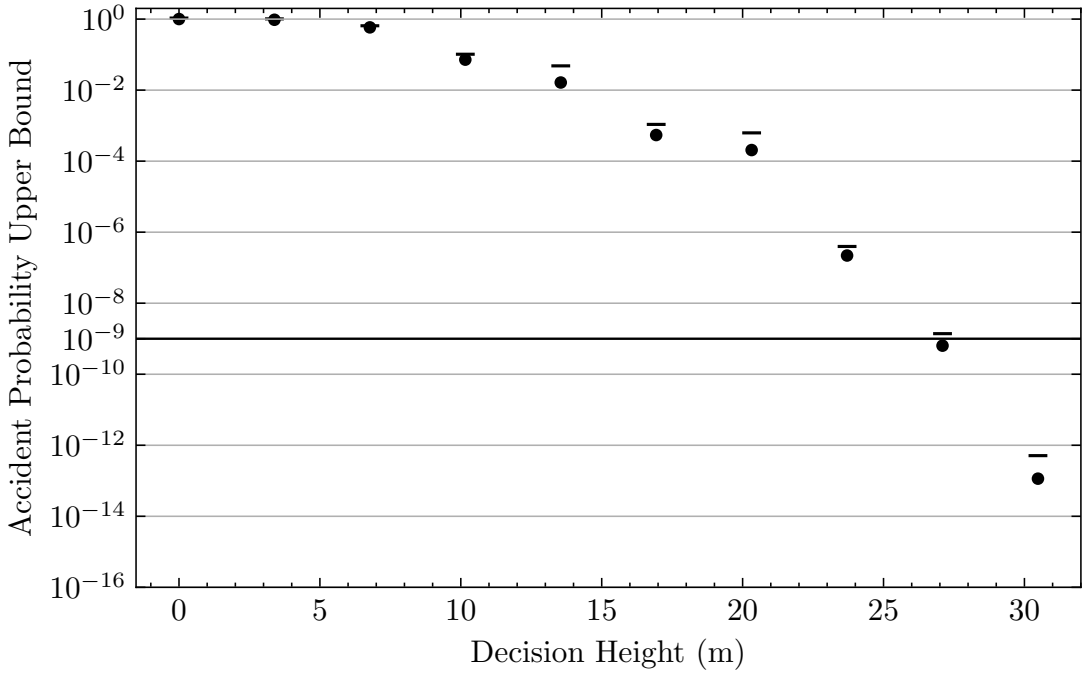

Figure 4.11: Accident Risk vs Decision Height, WAAS GPS and no turbulence

absent.

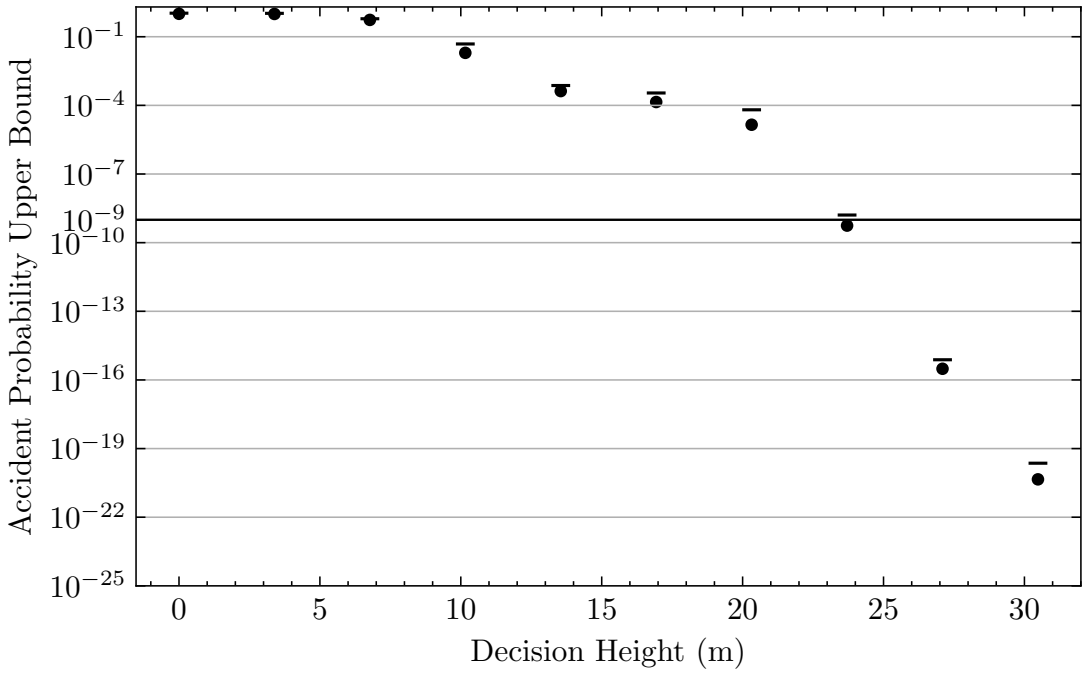

Figure 4.12: Accident Risk vs Decision Height, Radar Altimeter augmented WAAS GPS and nominal wind model

Results plotted in Figure 4.12 suggest that the minimum DH may reduced to 24 m when a radar altimeter is integrated.

4.4 Off-nominal Failure Condition Assessment

Several off-nominal failure conditions are considered in the approach and landing safety assessment, namely engine failures, vertical navigation drift errors, and severe wind shear events.

4.4.1 Engine Failure

FAA regulations require the approach and landing to be safe with failure of the critical engine. The aimpoint analysis is repeated again with an engine failure occurrence.

A runway aimpoint of 240 m is required to maintain safety when an engine failure occurs, 30 m greater than in the nominal case.

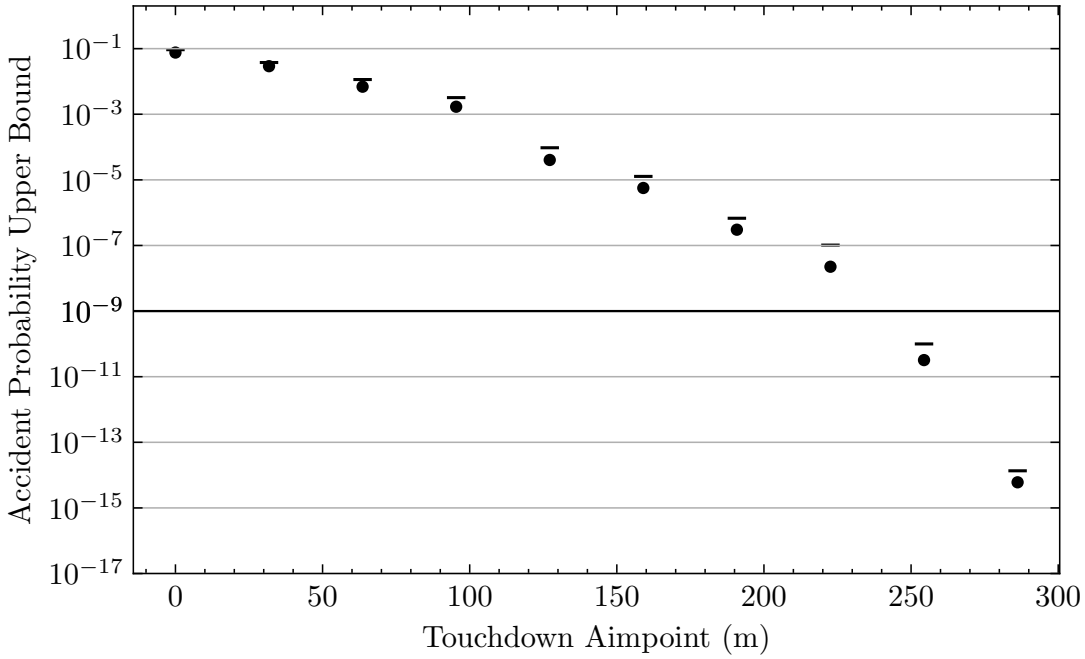

Figure 4.13: Accident Risk vs Runway Aimpoint, Engine Failure, WAAS GPS, and nominal wind model

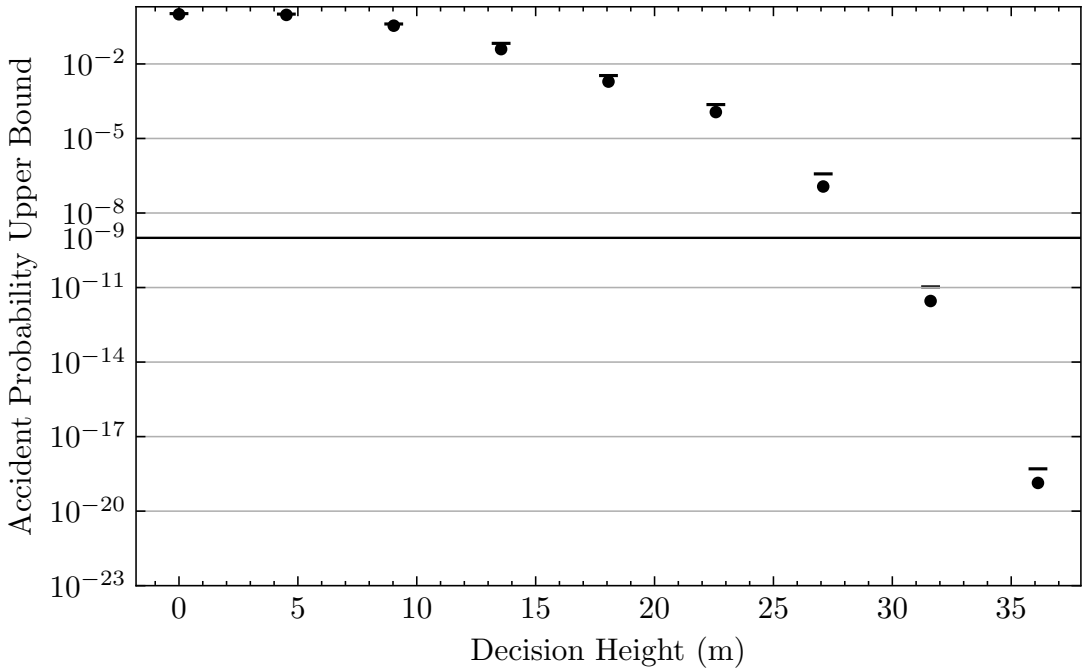

Figure 4.14: Accident Risk vs Decision Height, Engine Failure, WAAS GPS, and nominal wind model

The results plotted in Figure 4.14 suggest decision height must be raised to 30 m to maintain safety when an engine failure occurs, 4 m greater than in the nominal case.

4.4.2 Navigation Drift Error

The free parameter in the navigation drift error safety analysis is the maximum safe vertical error drift rate γ_z. This effectively increases the glide slope of the final approach thus increasing the risk of a runway undershoot. Maximum allowable γ_z is analyzed for both the WAAS GPS, and radar altimeter augmented WAAS GPS, considering failures in each separate system.

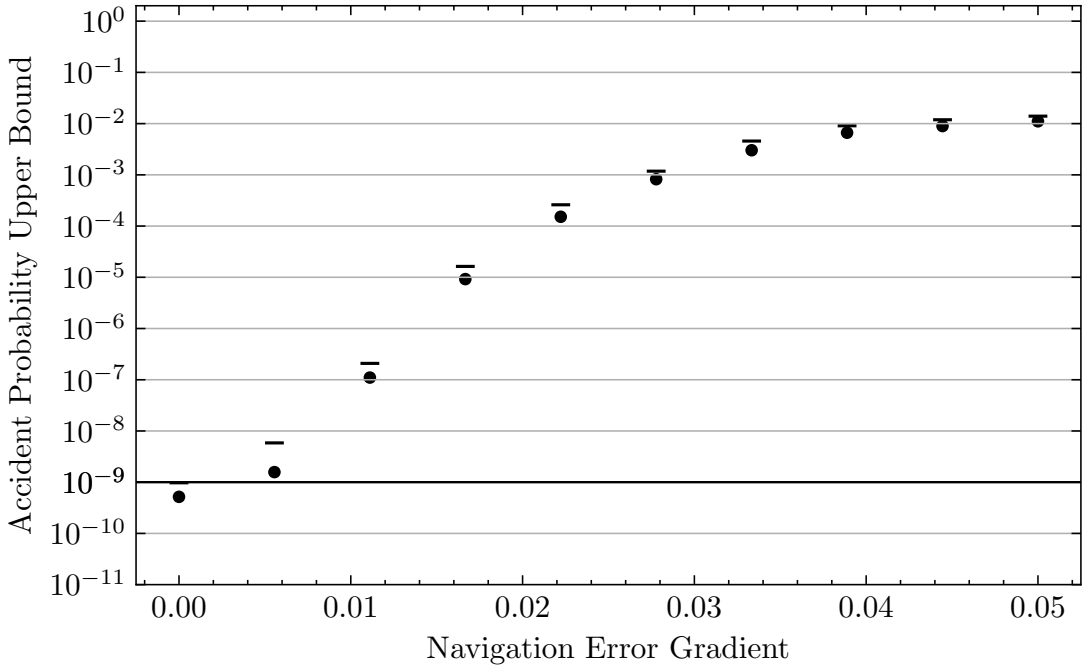

Figure 4.15: Accident Risk vs WAAS GPS Error Gradient, WAAS GPS, and nominal wind model

Results plotted in Figure 4.15 suggest that safety is extremely sensitive to WAAS GPS errors. Any vertical drift error will increase undershoot accident probability above 1×10^{-9}.

Results plotted in Figure 4.16 suggest that the addition of a properly functioning radar altimeter can sufficiently mitigate risk of WAAS GPS error gradients up to 0.04, which is approximately a change in glide slope of 2.3^{o}. This exceeds the largest ionosphere induced

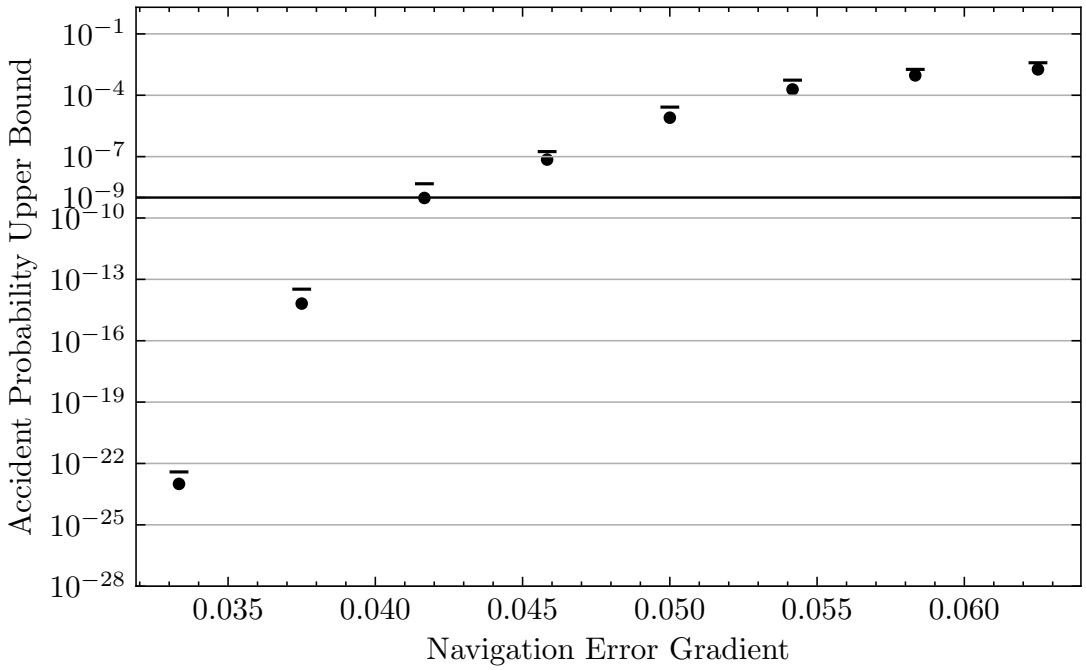

Figure 4.16: Accident Risk vs WAAS GPS Error Gradient, Radar Altimeter Augmented WAAS GPS, and nominal wind model

pseudorange error gradients of 0.0005 (0.5 m per km) that have been observed [62].

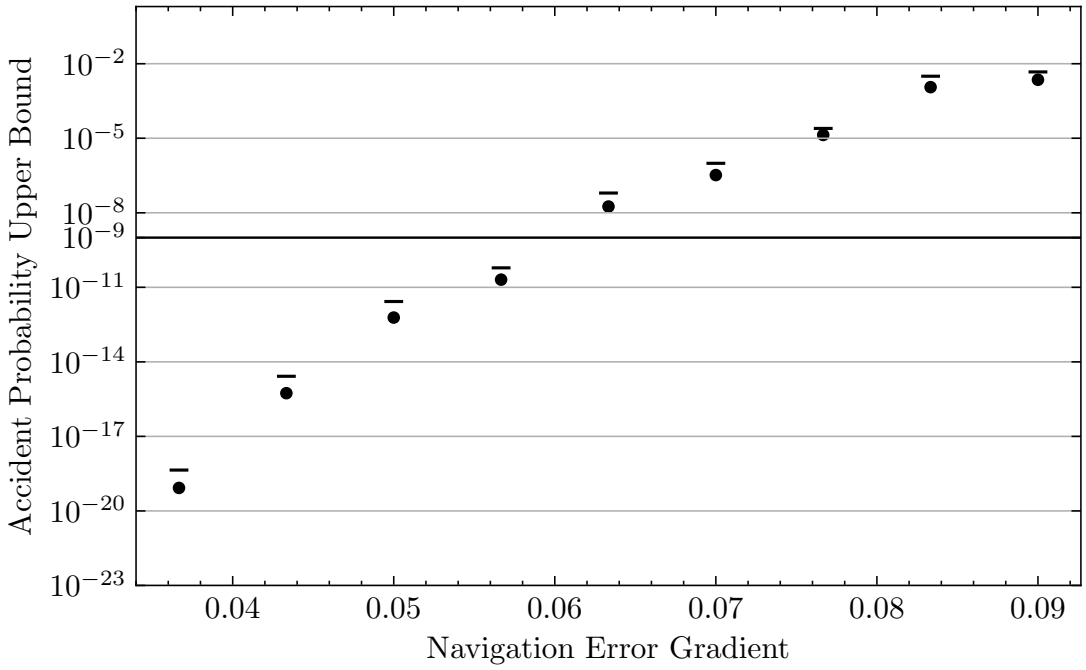

Figure 4.17: Accident Risk vs Radar Altimeter Error Gradient, Radar Altimeter Augmented WAAS GPS, and nominal wind model

The failure condition may be switched such that the radar altimeter fails while the WAAS GPS remains nominal. Results plotted in Figure 4.17 suggest that radar altimeter error gradients up to 0.06 are acceptable when the WAAS GPS is functioning nominally.

4.4.3 Wind shear Response

The free parameter in the wind shear safety analysis is the maximum wind shear gradient γ_{ws}. Any wind shear above this maximum allowable value would need to be mitigated.

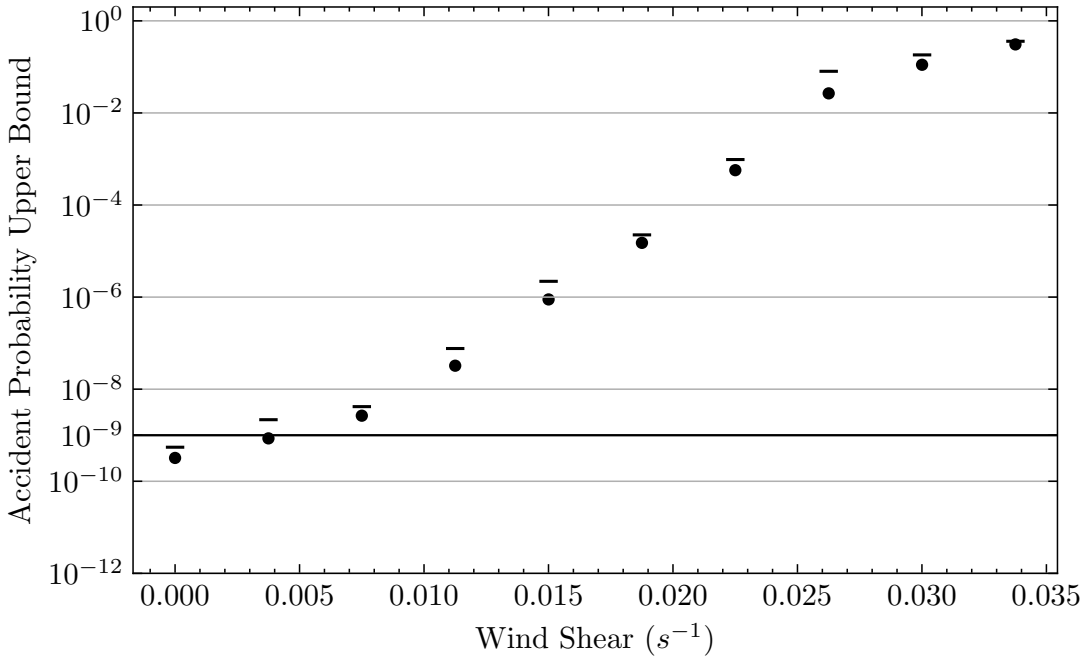

Figure 4.18: Accident Risk vs Wind shear Gradient, WAAS GPS and nominal wind model

Results plotted in Figure 4.18 suggest that wind shear gradients over approximately 0.002 s^{-1} are unsafe with nominal WAAS GPS performance. This is below the maximum wind shear gradient of 0.017 s^{-1} (50 knots per 5000 feet) recommended in the Wind shear Training Aid [32].

Results plotted in Figure 4.19 suggest that with the addition of a radar altimeter, wind shear gradients below approximately 0.02 s^{-1} are safe, exceeding the suggested wind shear gradient.

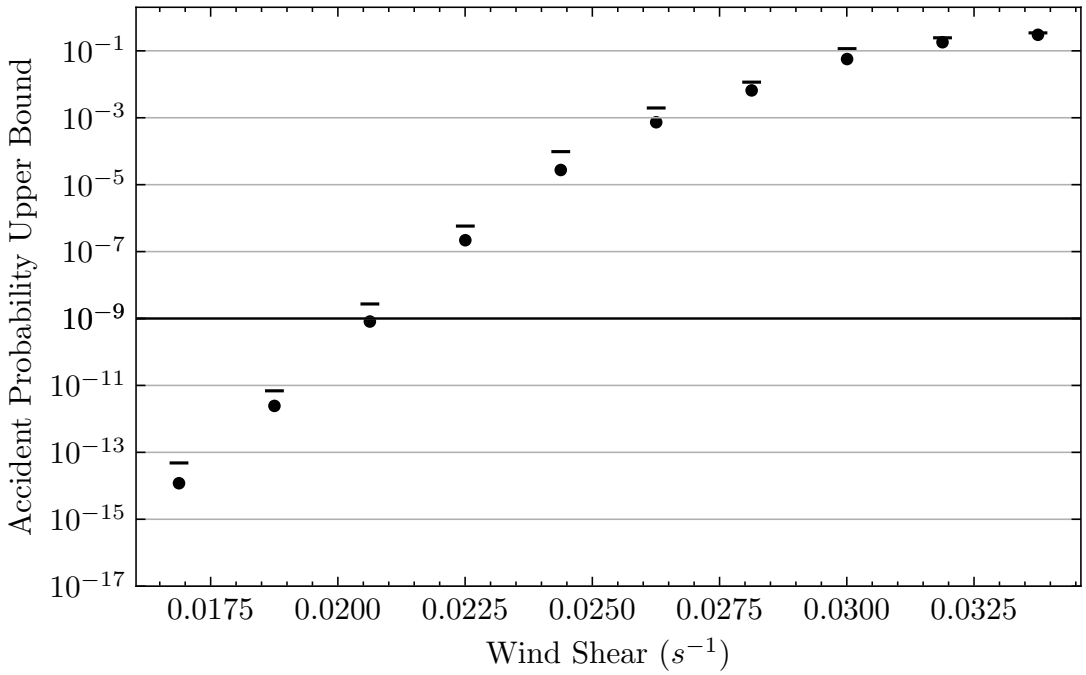

Figure 4.19: Accident Risk vs Wind shear Gradient, Radar Altimeter Augmented WAAS GPS, and nominal wind model

4.5 Discussion

Results from the safety assessments are listed in Table 4.2. The particular navigation system configuration and the associated limiting free parameter values providing passive risk mitigation for individual failure conditions are given.

Table 4.2: Passive Risk Mitigation Parameters

Configuration	Failure Condition	Parameter	Parameter Limit
WAAS GPS	Nominal	Aimpoint	210 m
	Nominal	Decision Height	26 m
	Engine Failure	Aimpoint	240 m
	Engine Failure	Decision Height	30 m
	WAAS GPS Error	Error Gradient	0.002
	Wind shear	Wind shear Gradient	0.002 s^{-1}
Radar Altimeter -Augmented WAAS GPS	Nominal	Aimpoint	75 m
	Nominal	Decision Height	24 m
	WAAS GPS Error	Error Gradient	0.04
	Radar Altimeter Error	Error Gradient	0.06
	Wind shear	Wind shear Gradient	0.02 s^{-1}

These limiting parameters may be used to set the requirements of an autonomous risk mitigation system. If the probability of a failure conditions occurring with a magnitude exceeding the limiting parameter is shown to be below 1×10^{-9}, no monitoring for the failure condition is required. When the probability of exceeding the limiting parameter cannot be neglected, monitoring must be put in place to detect the failure condition and maintain the target level of safety. This reintroduces the trade-off between missed detection probability and false alert probability examined in Chapter II. For example, if the probability of an unmitigated accident, $P(acc|\theta)$, at some failure condition magnitude θ is 1×10^{-8}, the allowable missed detection probability $P(md|\theta)$ is 1×10^{-1}. Further knowledge of the failure magnitude distribution $f(\theta)$ may be used to relax the missed detection requirements further. This requires designing the monitor such that its missed detection distribution satisfies Equation 4.5.

$$\int_0^\infty P(acc|\theta)P(md|\theta)f(\theta)d\theta \leq 1 \times 10^{-9} \tag{4.5}$$

If $P(md|\theta)$ and $f(\theta)$ are unavailable, we can make the conservative assumptions that $P(md|\theta) = 1$ and $P(acc|\theta)$ is monotonic with θ, then find the cutoff θ_{max} such that Equation 4.6 is satisfied.

$$\begin{aligned}\int_0^\infty P(acc|\theta)P(md|\theta)f(\theta)d\theta &\leq \int_0^\infty P(acc|\theta)f(\theta)d\theta \\ &\leq \int_0^{\theta_{max}} P(acc|\theta)f(\theta)d\theta + \int_{\theta_{max}}^\infty P(acc|\theta)f(\theta)d\theta \\ &\leq P(\theta \leq \theta_{max})P(acc|\theta_{max}) + P(\theta > \theta_{max})\end{aligned} \tag{4.6}$$

The $P(acc|\theta)$ curves computed in these safety assessments may be used to place conservative requirements on the prior probability density of failure magnitude. Without further information or monitoring, failure magnitudes exceeding θ_{max} will render the autonomous landing system unavailable.

CHAPTER 5

ONLINE SAFETY ASSESSMENT FOR ACTIVE RISK MITIGATION

The failure condition safety assessments performed in Chapter IV considered an autonomous aircraft functioning with only low-level automation. This automation includes flight control, navigation, state estimation, and procedure path following. Any failure-condition events that occurred during final approach and landing were not actively mitigated, but rather passively mitigated by procedure design, limits on allowable operating conditions, and limits on failure condition magnitudes. Observations made by the aircraft are disregarded when risk mitigation is purely passive. Airspeed deviations, large navigation error residuals, and loss of thrust can be used to make active risk mitigation decisions and execute a missed approach when safety is threatened. Active risk mitigation is performed by a decision making system onboard the autonomous aircraft following a decision making algorithm. This algorithm can take on various levels of sophistication. Human decision making has been classified according to performance level, an notable example being the Skill, Rule, Knowledge framework [63]. These decision making levels may compared to autonomous systems onboard the aircraft. Skill-based behaviors generally convert sensory input directly into outputs through learned sensory-motor mappings, analogously to low-level feedback control and stabilization. Rule-based behaviors generally perform rudimentary processing on inputs and generate outputs through a stored rule or procedure. This could be compared to a checklist based procedure for aborting a landing, where exceeding thresholds of airspeed, flight path deviations, or other system configurations trigger a missed approach. The highest level of decision making behavior is knowledge-based, where a knowledge of system behavior stored in a model and performance goals are used to solve a problem. Model predictive control frameworks fit this category, as a dynamics model and cost function are used to solve a cost minimization problem to generate optimal

flight controls [64]. This knowledge or model-based decision making paradigm is proposed to perform active risk mitigation. The safety assessment procedure used in Chapter IV may be applied in an online setting, using recent observations to update the internal system model and estimate the accident probability associated with the available procedure options. At all times, a procedure option must be available which satisfies safety requirements. Anticipated loss of safety for the primary landing procedure option triggers the execution of the secondary missed approach procedure option. This active risk mitigation framework utilizing an online safety assessment is outlined and compared to passive risk mitigation in scenarios where passive risk mitigation fails to prevent an accident.

5.1 Online Safety Assessment Methodology

Active risk mitigation during final approach and landing is accomplished by using online safety assessments to make decisions that satisfy safety requirements while making progress towards the primary goal of completing the landing. This may be framed generally as a constrained optimization problem that may apply to decision beyond the final approach and landing scenario. A number of procedure options may exist that can be ordered by some measure of optimality. The optimal procedure may accomplish the primary goal, such as landing, or minimize some metric such as flight time, fuel burn, emissions, or noise exposure. This desired procedure option will be termed the primary option. While continuing along the primary option, other alternative procedure options are still available. These include risk mitigating procedures such as missed approach, collision avoidance, or diversion to an alternate airport. At least one of the available procedure options must satisfy the safety requirements, which act as a constraint on accident probability. We will reduce our situation to a case in which only 1 alternative option is available, termed the secondary option. The decision to execute the secondary option is made when the primary option is unsafe and delaying execution of the secondary option any further will render it unsafe as well. If it is possible to continue along the primary option further and still safely

execute the secondary option, continuing the primary option is still deemed safe, even if completing the primary option itself unsafe. In practicality, observations used to update the internal model are made periodically and the online safety analysis takes time to complete. Decisions are made periodically at the end of every update step.

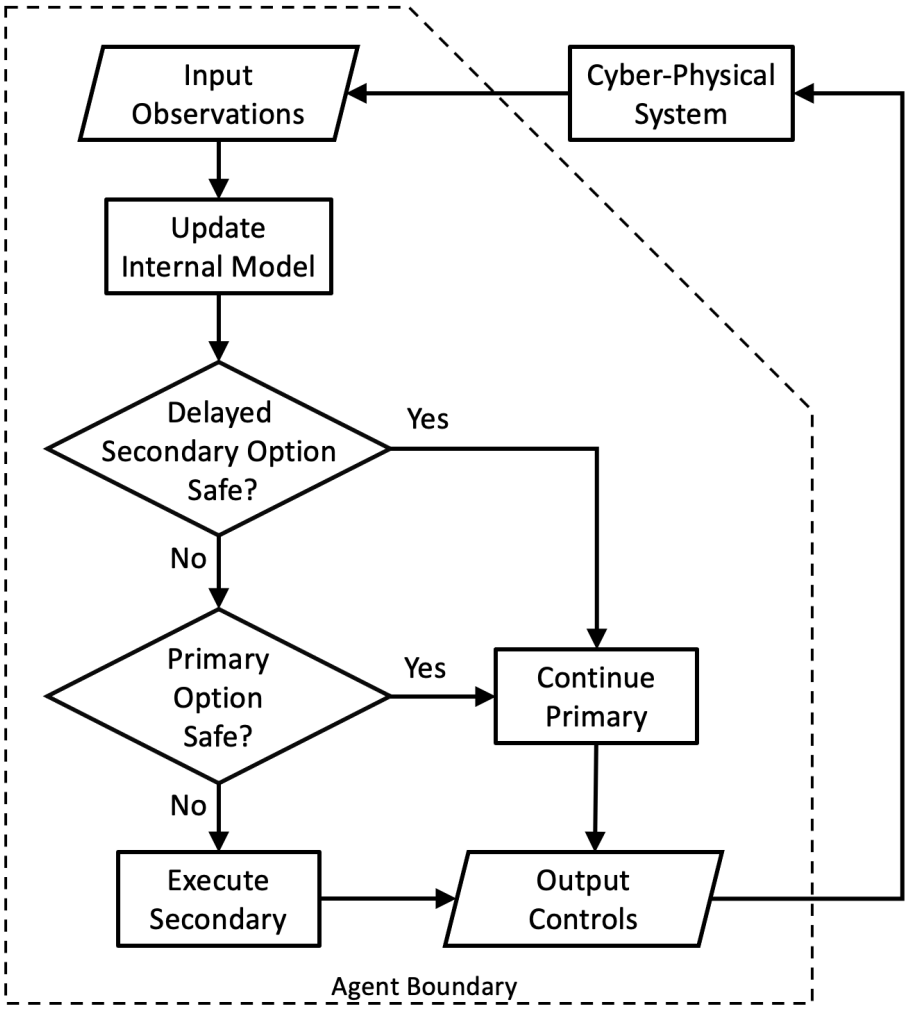

Figure 5.1: Active Risk Mitigation Decision Flowchart

A flowchart contained in Figure 5.1 depicts the active risk mitigation decision making in an OODA Agent framework [65]. The agent takes in observations from the external system and uses them to update the internal model, based around a Kalman filter framework. The periodic decision making occurs after observations have been gathered each update

step with time interval ΔT. The online safety assessments occur in the decision blocks and evaluate the accident probabilities for the primary and secondary options. The secondary option safety assessment assumes we continue along the primary option for the next update step and then execute the secondary option at the end of the next safety assessment. A safe result from this secondary option assessment means we may delay the decision to execute the secondary option another update step, an effective delay of $2\Delta T$, without compromising safety. However, at the point where delaying the secondary option is no longer safe and the primary option cannot be shown to be safe, the secondary option must be executed to maintain safety and mitigate risk. The secondary option is delayed as long as possible in the hope that information gathered from observations will be able to rule out an accident if the primary option is continued to completion. When applied to an instrument approach procedure, the cutoff where the secondary option becomes unsafe is directly analogous to the Decision Height, below which a missed approach is unsafe without improved navigation information. Likewise, during take-off rolls, the cutoff is analogous to the V_1 speed, above which the aircraft will be unable to safely come to a stop within the length of runway remaining. Extended-range Twin-engine Operations Performance Standards (ETOPS) allows twin engine aircraft to operate over 180 minutes flight time from a diversion airport. This ETOPS time serves an analogous purpose to the secondary option cutoff, balancing the risk of engine failure with the single engine flight time.

This decision making methodology can be extended to the more general case in which a preferred primary option and many alternative options are available, such as choosing an ideal diversion airport. In the case that none of the available options satisfy the safety requirements, the safest of the available options should be chosen instead. This would constitute operations outside of the regular operational domain where emergency decision making criteria are required.

5.2 Active Risk Mitigation Scenarios

The passive risk mitigation system parameters identified in Chapter 4 are used as a baseline for comparing the performance gained with active risk mitigation using online safety assessments. The parameters which passively mitigate risk to a level of 1×10^{-9} are used to acquire simulation samples that result in an accident. The rare event estimation methodology is used to search for the simulation trajectory samples with the highest likelihood of causing an accident. This baseline trajectory resulting in an accident is used to demonstrate active risk mitigation using online safety assessment. The risk associated with the delayed missed approach option and continue approach option are computed along the baseline trajectory and a missed approach is triggered by the decision making algorithm when the proper criteria are met. It is expected that when observations suggesting the increased probability of an accident are acquired, a missed approach will be triggered before an accident occurs during the final approach and landing. A crucial parameter in the active risk mitigation methodology is time between safety assessments ΔT. A larger ΔT allows more time for computing the safety assessment when computational limitations exist, however, it may compromise safety due to the delay in processing observations and detecting accident conditions. The arrival and approach procedure and scenario parameter set used in Chapter IV are also adopted for demonstrating active risk mitigation. The internal states of the aircraft are warmed up by an initial simulation down to 200 feet above ground level. At this point, rare event estimation is used to generate a simulation trajectory sample resulting in a runway undershoot accident using the initial belief state distribution estimated by the Kalman filter. Navigation drift error, wind shear, and engine failure conditions are chosen to demonstrate active risk mitigation.

5.2.1 Online Safety Assessment Time Interval Dependence

The first set of scenarios demonstrate how time interval ΔT affects the ability of the active risk mitigation to detect unsafe conditions before the missed approach option becomes unsafe. At each update step, we compute the continue approach risk, plotted in black with a 6σ upper bound error bar, and $2\Delta T$ delayed missed approach risk, plotted in blue with a 6σ upper bound error bar. When both of these accident probabilities are greater than 1×10^{-9}, a missed approach is triggered and the immediate missed approach probability is computed and plotted in red. Down arrows indicate an accident probability estimate below 1×10^{-20}. The immediate missed approach is actually executed at the end of the current update step, so it is in reality delayed by ΔT, as opposed to the delayed missed approach which is executed at the end of the next update step delayed by $2\Delta T$. To maintain safety, the immediate missed approach risk must be safe while the other 2 options are both unsafe.

The online safety assessment plots should be read left to right, with attention on which options safety assessment upper bound bars are above or below 1×10^{-9}. When both the black bar and blue bar are above this value, a missed approach should be performed. The immediate missed approach safety assessment given in red serves to show the validity of the delayed missed approach safety assessment in the previous update step. Fast acting failure conditions will result in sudden jumps in accident probability while slower acting failure conditions will result in gradual upward tendencies in accident probability.

Results plotted in Figure 5.2 depict the safety assessments over time for the WAAS GPS system with runway aimpoint 210 m and time interval ΔT of 5.0 s. The missed approach is triggered at 5.0 s and risk is kept sufficiently low. A delay of another 5.0 s would result in an unsafe missed approach.

Results plotted in Figure 5.3 depict the safety assessments over time for the WAAS GPS system with runway aimpoint 210 m and time interval ΔT of 2.0 s. The missed approach is triggered at 10.0 s and risk is again kept sufficiently low. A delay of another 2.0 s would result in an unsafe missed approach. Missed approach safety assessments too close to the

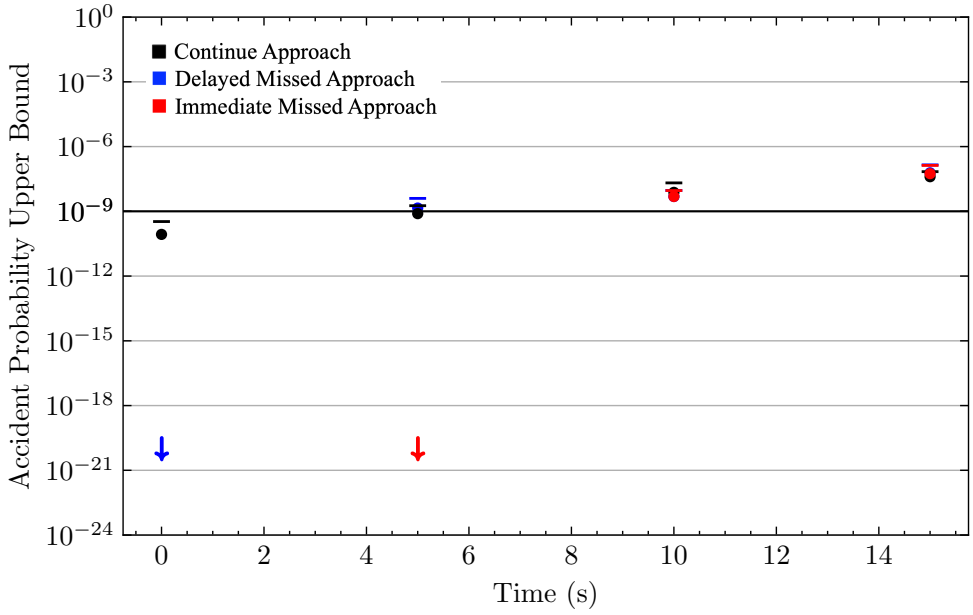

Figure 5.2: Online Safety Assessment vs Time, $\Delta T = 5.0s$, WAAS GPS and nominal wind model

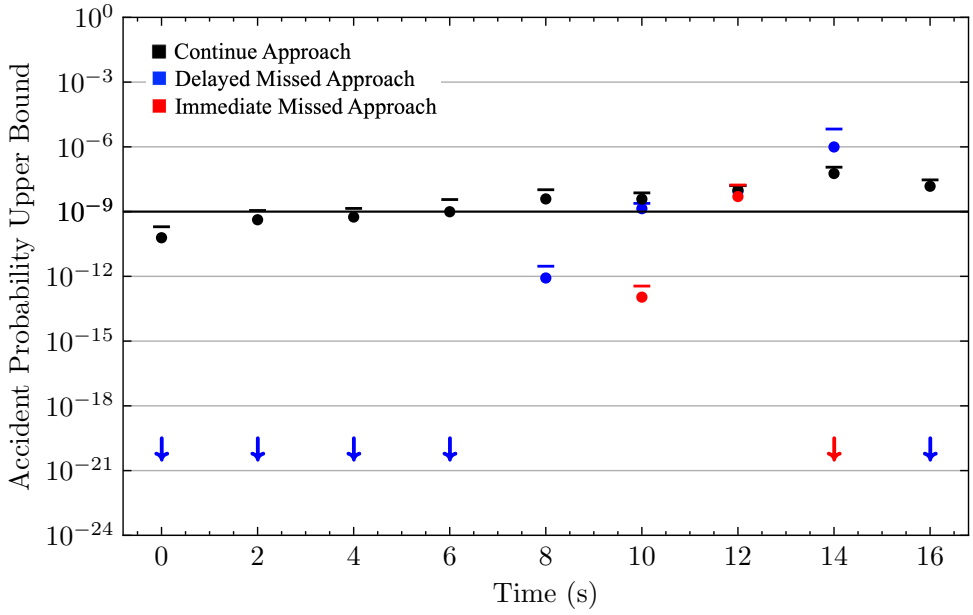

Figure 5.3: Online Safety Assessment vs Time, $\Delta T = 2.0s$, WAAS GPS and nominal wind model

landing are mislead into appearing safe due to large unobserved navigation errors.

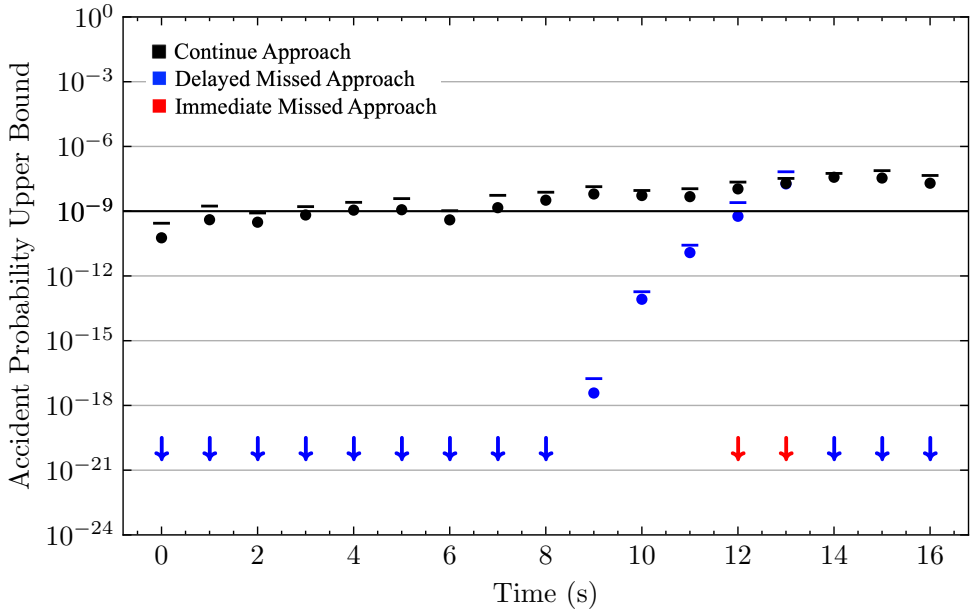

Figure 5.4: Online Safety Assessment vs Time, $\Delta T = 1.0s$, WAAS GPS and nominal wind model

Results plotted in Figure 5.4 depict the safety assessments over time for the WAAS GPS system with runway aimpoint 210 m and time interval ΔT of 1.0 s. The missed approach is triggered at 12.0 s and risk is again kept sufficiently low. A delay of another 1.0 s would still be maintain safety. Again, missed approach safety assessments too close to the landing are mislead into appearing safe due to large unobserved navigation errors. This may be mitigated by augmenting the WAAS GPS with a radar altimeter.

The second set of scenarios demonstrate how time interval ΔT affects the ability of the active risk mitigation to detect unsafe conditions with radar altimeter augmented WAAS GPS.

Results plotted in Figure 5.5 depict the safety assessments over time for the radar altimeter augmented WAAS GPS. system with runway aimpoint 75 m and time interval ΔT of 5.0 s. The missed approach is triggered at 5.0 s and risk is kept sufficiently low. A delay of another 5.0 s would result in an unsafe missed approach.

Results plotted in Figure 5.6 depict the safety assessments over time for the radar al-

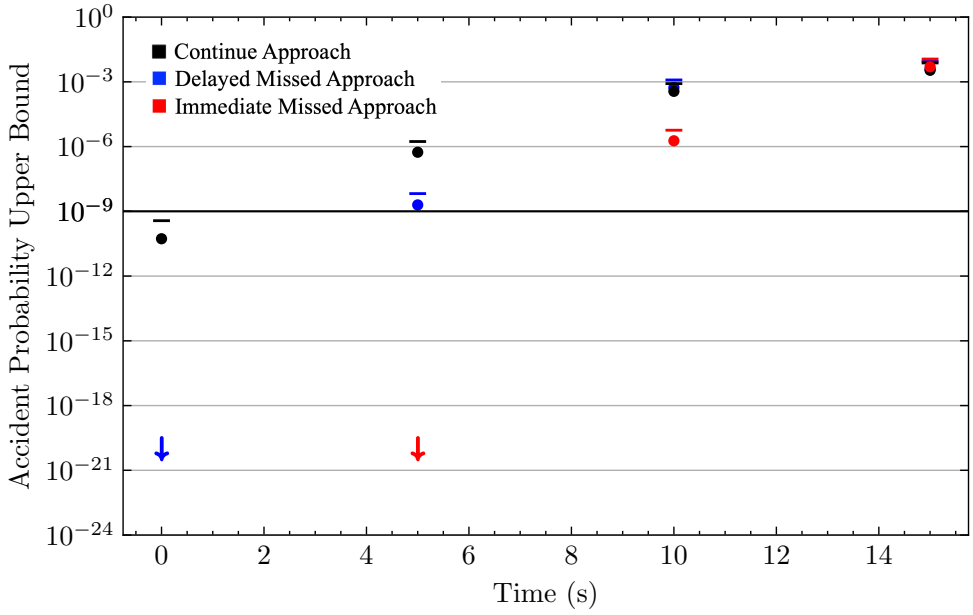

Figure 5.5: Online Safety Assessment vs Time, $\Delta T = 5.0s$, Radar Altimeter Augmented WAAS GPS and nominal wind model

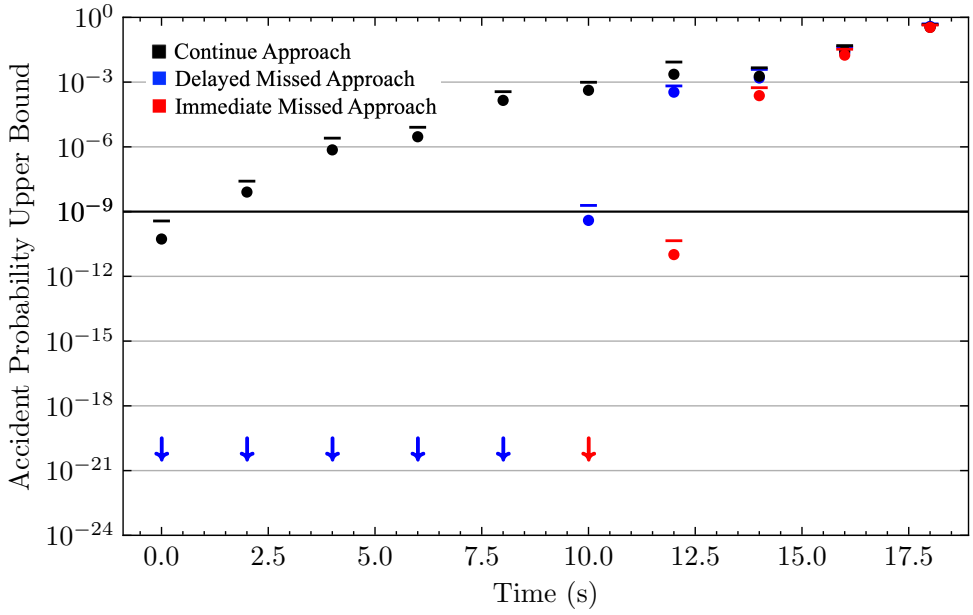

Figure 5.6: Online Safety Assessment vs Time, $\Delta T = 2.0s$, Radar Altimeter Augmented WAAS GPS and nominal wind model

timeter augmented WAAS GPS system with runway aimpoint 75 m and time interval ΔT of 2.0 s. The missed approach is triggered at 10.0 s and risk is again kept sufficiently low. A delay of another 2.0 s would result in an unsafe missed approach. In contrast to the pure WAAS GPS case, missed approach safety assessments close to landing are more accurate because of the radar altimeter corrections.

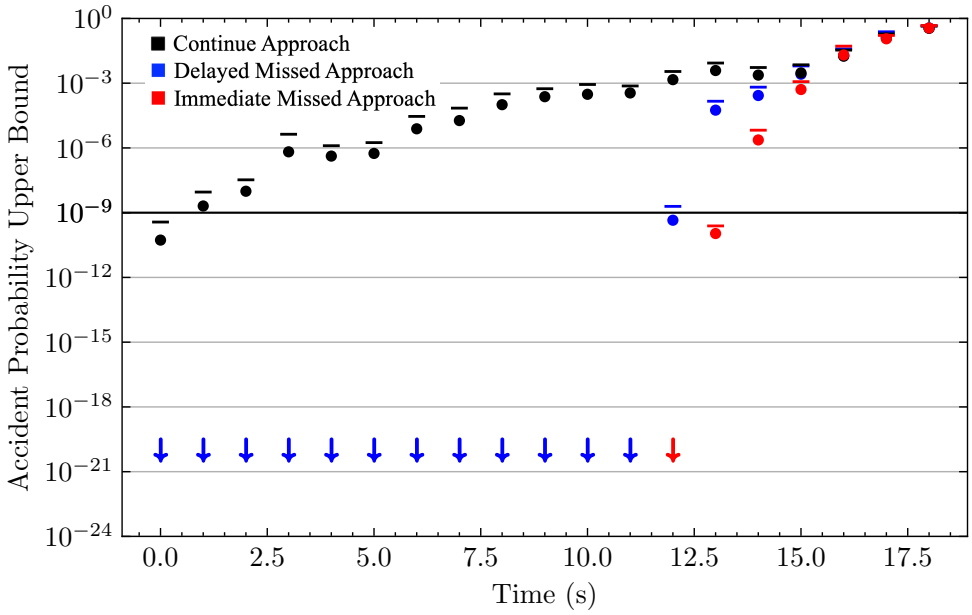

Figure 5.7: Online Safety Assessment vs Time, $\Delta T = 1.0s$, Radar Altimeter Augmented WAAS GPS and nominal wind model

Results plotted in Figure 5.7 depict the safety assessments over time for the radar altimeter augmented WAAS GPS system with runway aimpoint 75 m and time interval ΔT of 1.0 s. The missed approach is triggered at 12.0 s and risk is again kept sufficiently low. A delay of another 1.0 s would still maintain safety.

It is apparent from these experiments that smaller time intervals delay the decision to execute a missed approach. The increased rate of observations updates and reduced missed approach execution delay effectively decrease decision height and allow more time to gather information that may demonstrate the landing is safe. ΔT is set to 1.0 s for the remainder of active risk mitigation experiments.

5.2.2 Engine Failure Online Safety Assessment

The first failure condition considers an engine failure during final approach. It is assumed that the engine failure and subsequent loss of thrust can be accurately detected by the engine control system and used in the online safety assessment.

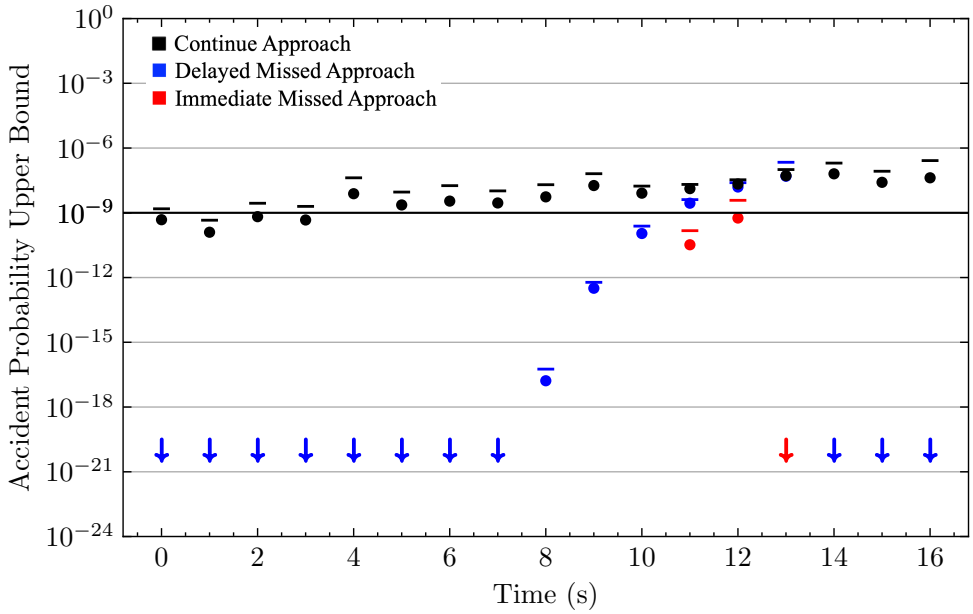

Figure 5.8: Online Safety Assessment vs Time, $\Delta T = 1.0s$, Engine Failure, WAAS GPS, and nominal wind model

Results plotted in Figure 5.8 depict the safety assessments over time for the WAAS GPS system with runway aimpoint 240 m and time interval ΔT of 1.0 s. The missed approach is triggered at 11.0 s and missed approach risk is kept sufficiently low for another update step. The problem of hazardously misleading risk estimates is still apparent near the touchdown due to uncorrected WAAS GPS errors.

Results plotted in Figure 5.9 depict the safety assessments over time for the radar altimeter augmented WAAS GPS system with runway aimpoint 75 m and time interval ΔT of 1.0 s. The missed approach is triggered at 11.0 s and cannot be safely delayed another update step. The detection of a loss of thrust quickly increases the risk estimate and the radar altimeter navigation error corrections produce more consistent risk estimates near

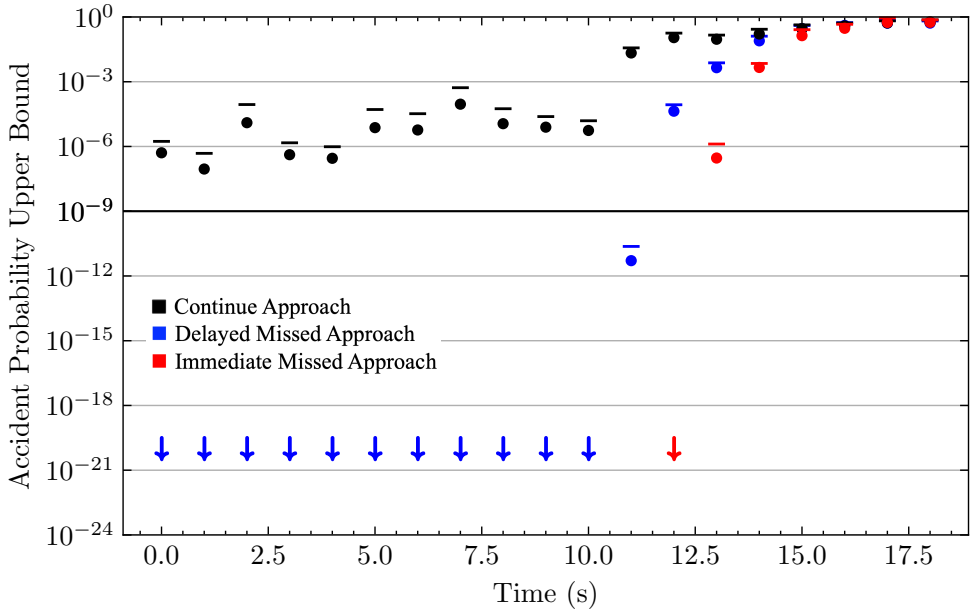

Figure 5.9: Online Safety Assessment vs Time, $\Delta T = 1.0s$, Engine Failure, Radar Altimeter Augmented WAAS GPS, and nominal wind model

touchdown.

5.2.3 Wind Shear Online Safety Assessment

The second failure condition considers an onset of severe wind shear during final approach. It is assumed that a decrease in airspeed can be detected by the air data system and considered in the safety assessment as an increase in tailwind.

Results plotted in Figure 5.10 depict the safety assessments over time for the WAAS GPS system with a wind shear gradient of 0.002 s^{-1}, runway aimpoint 210 m, and time interval ΔT of 1.0 s. The missed approach is triggered at 11.0 s and missed approach risk and cannot be delayed another update step. The problem of hazardously misleading risk estimates is still apparent near the touchdown due to uncorrected WAAS GPS errors.

Results plotted in Figure 5.11 depict the safety assessments over time for the radar altimeter augmented WAAS GPS system with a wind shear gradient of 0.002 s^{-1}, runway aimpoint 75 m, and time interval ΔT of 1.0 s. The missed approach is triggered at 12.0 s and may be safely delayed another update step. The detection of wind shear gradually

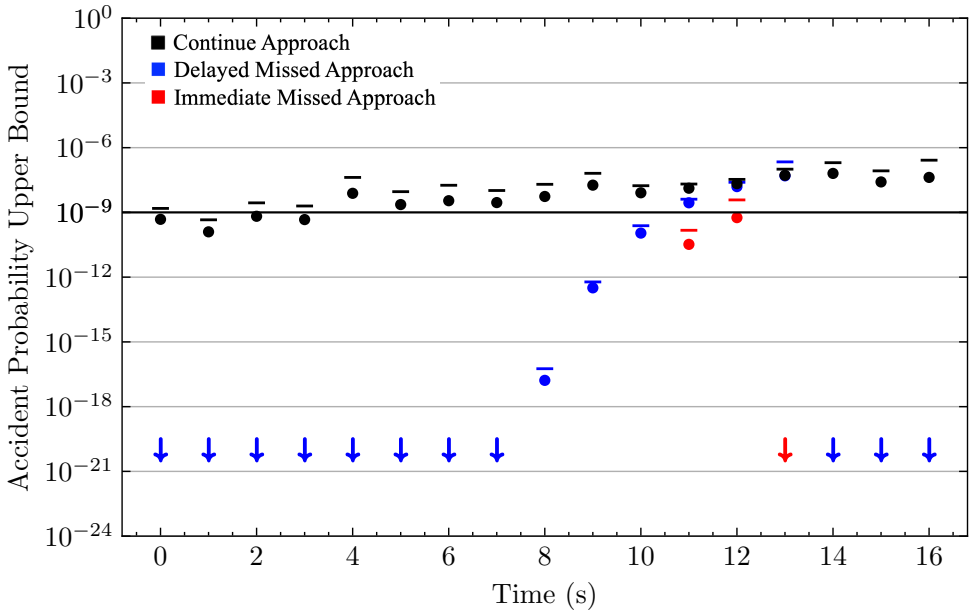

Figure 5.10: Online Safety Assessment vs Time, $\Delta T = 1.0s$, Wind shear, WAAS GPS

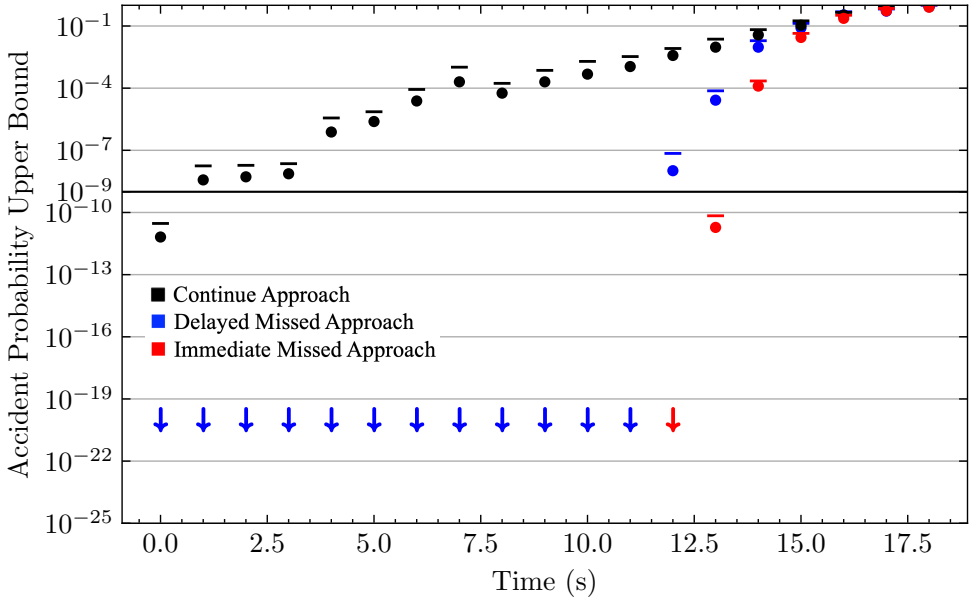

Figure 5.11: Online Safety Assessment vs Time, $\Delta T = 1.0s$, Wind shear, Radar Altimeter Augmented WAAS GPS

increases the risk estimate and the radar altimeter navigation error corrections produce more consistent risk estimates near touchdown.

5.2.4 Navigation Drift Error Online Safety Assessment

The third failure condition considers a vertical navigation drift error during final approach. Navigation errors cannot be detected without independent measurements.

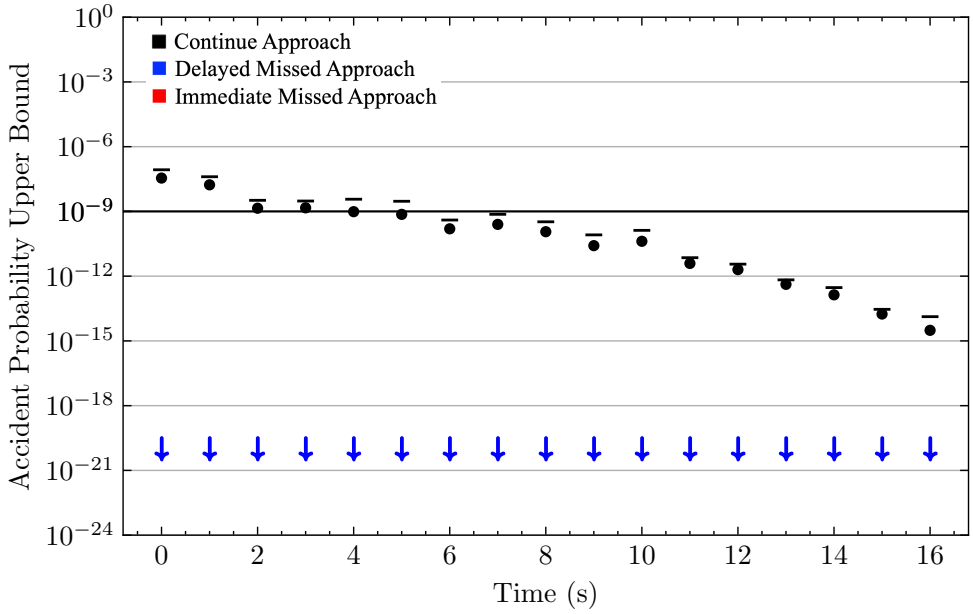

Figure 5.12: Online Safety Assessment vs Time, $\Delta T = 1.0s$, WAAS GPS Drift Error, WAAS GPS, and nominal wind model

Results plotted in Figure 5.12 depict the safety assessments over time for the WAAS GPS system with a vertical drift error gradient of 0.008, runway aimpoint of 210 m, and time interval ΔT of 1.0 s. The risk starts out large due to large risk exposure time before touchdown and decreases as this exposure time decreases. Due to a lack of independent observations to detect or correct navigation errors, the accident precursors go undetected and a missed approach is not executed, resulting in an accident.

Results plotted in Figure 5.13 depict the safety assessments over time for the radar altimeter augmented WAAS GPS system with a GPS vertical drift error gradient of 0.008, runway aimpoint of 75 m, and time interval ΔT of 1.0 s. The missed approach is triggered

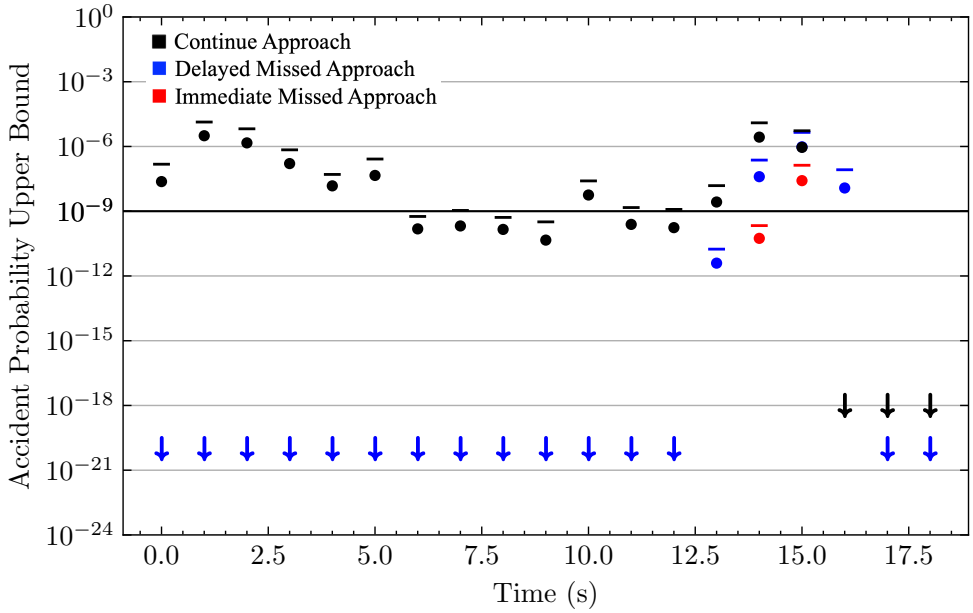

Figure 5.13: Online Safety Assessment vs Time, $\Delta T = 1.0s$, WAAS GPS Drift Error, Radar Altimeter Augmented WAAS GPS, and nominal wind model

at 14.0 s and cannot be safely delayed another update step. The radar altimeter provides the necessary independent information to mitigate the risk, but remaining navigation errors still result in hazardously misleading safety assessments close to touchdown.

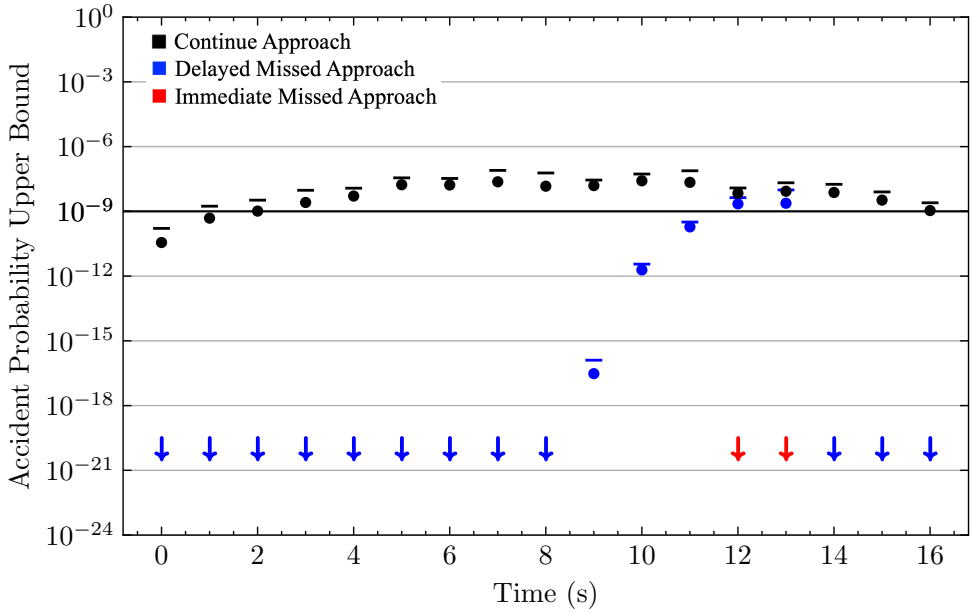

Figure 5.14: Online Safety Assessment vs Time, $\Delta T = 1.0s$, Radar Altimeter Drift Error, Dual Radar Altimeter Augmented WAAS GPS, and nominal wind model

The possibility of a radar altimeter drift error requires a further source of independent observations more accurate than WAAS GPS to provide a correction. This is accomplished by adding a second independent radar altimeter to the aircraft such that navigation is provided by redundant dual radar altimeter augmented WAAS GPS. Results plotted in Figure 5.14 depict the safety assessments over time for the dual radar altimeter augmented WAAS GPS system with a drift error gradient of 0.008 applied to a single radar altimeter, runway aimpoint of 75 m, and time interval ΔT of 1.0 s. The missed approach is triggered at 12.0 s and can be safely delayed another update step. However, remaining navigation errors still result in hazardously misleading safety assessments close to touchdown.

5.3 Discussion

While passive risk mitigation can ensure safety when failure condition magnitudes stay within limits, active risk mitigation can detect deviations from safe conditions while operations are underway. A variety of observation inputs such as air data sensors, engine sensors, and redundant navigation systems can be used to update the internal state estimate of the computer agent onboard the unmanned aircraft and perform online safety assessments to estimate the probability of an accident during operation. The comparison of accident probabilities for the available procedure options allows the agent to always maintain an option that satisfies safety requirements.

The period of the update steps, time interval ΔT, has a significant impact on how far along the final approach and aircraft may proceed before risk grows to large. Small ΔT allows decisions to be made much closer to touchdown, however, it requires safety assessments to be completed in less time. The majority of the time consuming computation during rare event estimation is due to running batches of fast-time aircraft simulations. While large batches consisting of several hundred simulations are too slow to run serially in real-time by a factor of 10, the methodology is readily parallelizable. Utilization of basic parallel computing resources could be applied to speedup the batch simulations by

the required magnitude, allowing real-time online safety assessments

The findings relating to the navigation system redundancy needed to actively mitigate risk ties back to the requirements of conventionally piloted ILS procedures presented in Chapter I. Operations with lower minima require more accurate equipment, redundant systems, and active monitoring to detect and mitigate errors. Augmenting existing navigation systems with independent sources of information is a viable path towards providing the performance required for both zero-visibility conventional landings and autonomous landings of unmanned aerial systems.

CHAPTER 6
CONCLUSIONS

The operation of Remotely Piloted Aerial Systems to the National Air Space in routine cargo transport roles requires systems to meet stringent safety requirements defined in Federal Aviation Regulations. This includes autonomous systems put in place to mitigate risks when the Remote Pilot cannot effectively intervene due to Command and Control link failures or latency. A methodology has been proposed to perform safety assessments on an aircraft in a final approach scenario and estimate the accident probability in several failure conditions.

6.1 Contributions

The primary contribution of this thesis are as follows:

- Safety Regulations are formulated as a chance-constraint satisfaction problem, requiring safety on the order of 1 accident per billion operations for each failure condition.
- A Stochastic Hybrid System model of an Unmanned Aerial Vehicle is proposed to handle the coupling between discrete and continuous system states and estimate the distribution of aircraft trajectories that may result from a given set of system parameters, operational conditions, and decision parameters.
- Autonomy onboard the Unmanned Aerial Vehicle is modelled as an Agent in an Observe-Orient-Decide-Act based Framework. An Extended Kalman Filter updates an internal model which is used for guidance, navigation, control, decision making, and detection of faults using likelihood ratio tests.

- White noise processes contributing to disturbances and measurement noise are reduced in dimensionality using the Karhunen-Loeve Expansion with Discrete-Cosine Transform bases. This allows efficient sampling and rare event estimation.

- Transitions between Discrete States are described using a Markov model. It is shown that a sequence of transition event times may be sampled using a Dirichlet distribution when transition rates are small with respect to a risk exposure time interval.

- Rare event estimation techniques based on Importance Sampling are used to assess accident probability subject to various sources of uncertainty modelled by Multivariate Normal and Dirichlet distributions. A Stochastic Differential Equation based algorithm is proposed to quickly find an optimal Importance Sampling distribution.

- The final approach and landing phase of flight serves as a use case for the methodology. The safety assessment is applied to determine system design parameters required to passively mitigate risks.

- The methodology is extended to active risk mitigation during operations, where online safety assessments using updated observations are used to ensure decision options always exist that will satisfy safety requirements.

6.2 Significance and Extensions

The proposed safety assessment methodology allows efficient evaluation of a Unmanned Aerial System's compliance with safety regulations in uncertain operational conditions. Integration of new systems and subsystems into routine operations may be aided by modelling and parameterizing prospective systems in accordance with the Stochastic Hybrid System model and performing safety assessments to determine the system requirements needed to adequately mitigate risk. The methodology invites extensions beyond the single agent model, including systems with a Remote Pilot (RP) needed to make decisions

Table 6.1: Ground Based Remote Pilot Monitoring Parameters

Subsystem	Time Constant	Integrity Probability
Sensor	Sample Rate	Sensor Integrity
Downlink Transmission	Downlink Latency	Downlink Integrity
GCS Display	Update Rate	Display Integrity
Pilot Response	Response Time	Detection Probability
GCS Interface	Read Rate	Interface Integrity
Uplink Transmission	Uplink Latency	Uplink Integrity
Command Execution	Execution Lag	Actuator Integrity

and monitor for risks. Detailed modelling of a human RP is infeasible for a general decision making task, however, simplified modelling of RP decision making performance for a specific task may be considered in an availability, continuity, integrity, and alert/response time requirements framework. Whether risk mitigation is performed by a human or machine, we still require timely responses with a low missed detection rate to maintain safety, and high availability with a low false alert rate to maintain operational efficiency. Alert time and missed detection probability for ground based a ground based RP in a monitoring role depends not just on the pilot response, but also C2 downlink, C2 uplink, and other intermediate equipment between onboard observations and execution of a risk mitigating procedure. A breakdown of the various subsystem parameters that may contribute to the total ground based RP monitoring performance is listed in Table 6.1. The required total system performance may be determined using the safety assessment methodology and broken down and allocated in subsequent detailed design.

The methodology allows extensions to scenarios involving other aircraft treated as agents. Such scenarios could include merging of flows in terminal airspace, approach sequencing, or autonomous detect and avoid for separation maintenance. Beyond the sources of uncertainty considered in the single agent scenario, multi agent scenarios may require agents to predict the intent and trajectories of other agents. State estimation techniques may be applied to tracking the state of other aircraft using noisy observations and estimation of intent using a sequence of flight segments [66]. Sharing of intent information

between aircraft can provide prior beliefs on trajectory intent that may be updated with observations. The safety assessment may be applied to stochastic models of detect and avoid systems to determine the performance required from such systems. This is needed to determine if UAS must stay confined to airspace corridors where all aircraft are required to be equipped with transponder and intent sharing systems, or whether UAS may operate in uncontrolled airspace with aircraft following Visual Flight Rules. The safety assessment methodology is suitable for quantifying the trade-off between systems requirements and allowable operational domains.

The safety assessment methodology may be applied in an online setting, allowing the possibility of active risk mitigation as demonstrated in Chapter V. Routine implementation of such an active risk mitigation system requires certification and demonstration that failure conditions producing hazardously misleading information are extremely improbable. Safety-critical software must comply with standards set in DO-178C [67], which aims to mitigate risks due to failure conditions arising from software implementations. The iterative nature of the active risk mitigation requires guarantees on the amount of time a safety assessment will take to complete. Safety assessments that take too much time may compromise safety and render the system unavailable. This may be mitigated by ensuring adequate parallel computing resources are available or demonstration of bounded convergence times for the rare event estimation algorithm. A major impediment to certification is the fact that the safety assessment methodology utilizes randomized or stochastic algorithms to demonstrate safety. Stochastic algorithms have typically not been used in safety critical systems and procedures for certifying such algorithms have not been set in stone. The prospect of treating the output of a stochastic algorithm as random variable, alongside other random variables contributing to system uncertainty, is a promising path towards certification. Validating the estimated distribution of the stochastic algorithm output to the 1×10^{-9} quantile should be a requirement for certifying such algorithms. A precursor to full reliance on the safety assessment methodology in autonomous operations may be

implementation as decision support tool for manned or remotely piloted operations. The performance of the online safety assessment may be evaluated with a human available to take control and mitigate risks neglected by the algorithm in failure cases. By accumulating flight hours and comparing safety assessment predictions to actual outcomes, a statistical case may be made for the integrity of the safety assessment, leading to certification for routine UAS operations.

REFERENCES

[1] J. Sakakeeny, N. Dimitrova, and H. R. Idris, "Preliminary characterization of unmanned air cargo routes using current cargo operations survey," in *AIAA AVIATION 2022 Forum*, 2022, p. 3701.

[2] V. Bulusu, G. B. Chatterji, T. A. Lauderdale, J. Sakakeeny, and H. R. Idris, "Impact of latency and reliability on separation assurance with remotely piloted aircraft in terminal operations," in *AIAA AVIATION 2022 Forum*, 2022, p. 3704.

[3] M. Zolanvari, R. Jain, and T. Salman, "Potential data link candidates for civilian unmanned aircraft systems: A survey," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 1, pp. 292–319, 2019.

[4] B. C. Airplanes, "Statistical summary of commercial jet airplane accidents," *Worldwide Operations 1959 - 2019*, 2020.

[5] Federal Aviation Administration, *Ac 120-29a. advisory circular - criteria for approval of category i and category ii weather minima for approach*, 2002.

[6] R. Kelly and J. Davis, "Required navigation performance (rnp) for precision approach and landing with gnss application," *Navigation*, vol. 41, no. 1, pp. 1–30, 1994.

[7] J. Charnley, "The rae contribution to all-weather landing," *Journal of Aeronautical History*, vol. 1, no. 1, 2011.

[8] Federal Aviation Administration, *Ac 120-18. advisory circular - criteria for approval/authorization of all weather operations (awo) for takeoff, landing, and rollout*, 2018.

[9] ICAO, *International standards and recommended practices, annex 10 to the convention on international civil aviation, volume 1, radio navigation aids*, 2006.

[10] F. Wolfe, "Faa certifies garmin autoland for piper m600/sls aircraft," *Aviation Today*, 2020.

[11] E. Perl, "Review of airport surface movement radar technology," in *2006 IEEE Conference on Radar*, IEEE, 2006, 4–pp.

[12] J. Jakobi, D. Teotino, and P. Montebello, "Towards higher-level services of an advanced surface movement guidance and control system (a-smgcs)," *Air Traffic Control Quarterly*, vol. 18, no. 2, pp. 143–175, 2010.

[13] C. A. Authority, *Safety notice sn–2013/010, missed approaches in response to on-board windshear alearts*, 2013.

[14] L. R. Newcome, "Unmanned aviation: A brief history of unmanned aerial vehicles; american institute of aeronautics and astronautics," *Inc.: Reston, VA, USA*, pp. 45–50, 2004.

[15] P. W. Merlin, *Crash Course: lessons learned from accidents involving remotely piloted and autonomous aircraft*, AFRC-E-DAA-TN5128. 2013.

[16] D. Couto, K. Delmas, and X. Pucel, "On the safety assessment of rpas safety policy," 2020.

[17] G. Guglieri and G. Ristorto, "Safety assessment for light remotely piloted aircraft systems," 2016.

[18] J. Malecha and A. S. Inspector, "Drone operations over 55 pounds," in *FAA UAS Symposium*, 2019.

[19] R. B. Ferreira *et al.*, "A risk analysis of unmanned aircraft systems (uas) integration into non-segregate airspace," in *2018 International Conference on Unmanned Aircraft Systems (ICUAS)*, IEEE, 2018, pp. 42–51.

[20] G. Serafino, D. Derin, F. Babich, E. Pietrosemoli, and M. Goiak, "Link performance evaluation procedure for the introduction of unmanned air vehicles in civil airspace," in *2019 IEEE 5th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, IEEE, 2019, pp. 182–186.

[21] R. Sabatini, T. Moore, and C. Hill, "Gnss avionics-based integrity augmentation for rpas detect-and-avoid applications," 2014.

[22] B. Blom, L. Bretschneider, and P. Hecker, "Rpas automatic take-off and landing operations using computer vision," in *2017 Integrated Communications, Navigation and Surveillance Conference (ICNS)*, 2017, 5B2-1-5B2–8.

[23] A. Gautam, P. B. Sujit, and S. Saripalli, "A survey of autonomous landing techniques for uavs," *2014 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 1210–1218, 2014.

[24] A. J. Kornecki and M. Liu, "Fault tree analysis for safety/security verification in aviation software," *Electronics*, vol. 2, no. 1, pp. 41–56, 2013.

[25] A. de Ruijter and F. Guldenmund, "The bowtie method: A review," *Safety science*, vol. 88, pp. 211–218, 2016.

[26] M. G. Pecht and F. R. Nash, "Predicting the reliability of electronic equipment," *Proceedings of the IEEE*, vol. 82, no. 7, pp. 992–1004, 1994.

[27] S. Pullen, J. Rife, and P. Enge, "Prior probability model development to support system safety verification in the presence of anomalies," in *Proceedings of IEEE/ION PLANS 2006*, 2006, pp. 1127–1136.

[28] M. Fan, Z. Zeng, E. Zio, R. Kang, and Y. Chen, "A stochastic hybrid systems based framework for modeling dependent failure processes," *PloS one*, vol. 12, no. 2, e0172680, 2017.

[29] L. Ren and J.-P. B. Clarke, "Flight-test evaluation of the tool for analysis of separation and throughput," *Journal of Aircraft*, vol. 45, no. 1, pp. 323–332, 2008.

[30] G. Nagle, *Tool for analysis and separation (tasat) users manual*, 2009.

[31] Federal Aviation Administration, *Ac 120-28d. advisory circular - criteria for approval of category iii weather minima for takeoff, landing, and rollout*, 1999.

[32] ——, *Windshear training aid*, 1987.

[33] J. Rankin, "Gps and differential gps: An error model for sensor simulation," in *Position Location and Navigation Symposium*, INSTITUTE OF ELECTRICAL & ELECTRONICS ENGINEERS INC, 1994, pp. 260–260.

[34] Federal Aviation Administration, "Global positioning system wide area augmentation system (waas) performance standard," U.S. Department of Transportation, Tech. Rep., 2008.

[35] P. Neri, C. Macabiau, L. Azoulai, and J. Muller, "Gbas nse model for cat ii/iii autoland simulations," in *IEEE/ION Position, Location and Navigation Symposium*, IEEE, 2010, pp. 694–707.

[36] A. Videmsek, M. U. de Haag, and T. Bleakley, "Radar altimeter aiding of gnss for precision approach and landing of rpa," in *2019 Integrated Communications, Navigation and Surveillance Conference (ICNS)*, IEEE, 2019, pp. 1–16.

[37] H. Park, J. Lee, and J. Lee, "Error modelling method of extended kalman filter-based terrain referenced navigation system for integrity assurance under nominal conditions," *IET Radar, Sonar & Navigation*, 2022.

[38] L. Sparks, A. Komjathy, and A. Mannucci, "Sudden ionospheric delay decorrelation and its impact on the wide area augmentation system (waas)," *Radio Science*, vol. 39, no. 1, pp. 1–8, 2004.

[39] C. Hajiyev and R. Saltoglu, "Robust integrated ins/radar altimeter accounting faults at the measurement channels," in *Proceedings of the 23rd ICAS (The International Council of the Aeronautical Sciences) Congress, Toronto, Canada*, vol. 6, 2002.

[40] N. C. Yadav, A. Shanmukha, B. Amruth, *et al.*, "Development of gps/ins integration module using kalman filter," in *2017 International Conference on Algorithms, Methodology, Models and Applications in Emerging Technologies (ICAMMAET)*, IEEE, 2017, pp. 1–5.

[41] J. Neyman and E. S. Pearson, "Ix. on the problem of the most efficient tests of statistical hypotheses," *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 231, no. 694-706, pp. 289–337, 1933.

[42] Z. Liu, Z. Liu, and Y. Peng, "Dimension reduction of karhunen-loeve expansion for simulation of stochastic processes," *Journal of Sound and Vibration*, vol. 408, pp. 168–189, 2017.

[43] N. Ahmed, T. Natarajan, and K. R. Rao, "Discrete cosine transform," *IEEE transactions on Computers*, vol. 100, no. 1, pp. 90–93, 1974.

[44] Federal Aviation Administration, *Ac 15.1309-1a. advisory circular - system design and analysis*, 1988.

[45] D. Kundu and R. D. Gupta, "A convenient way of generating gamma random variables using generalized exponential distribution," *Computational Statistics & Data Analysis*, vol. 51, no. 6, pp. 2796–2802, 2007.

[46] M. Denny, "Introduction to importance sampling in rare-event simulations," *European Journal of Physics*, vol. 22, no. 4, p. 403, 2001.

[47] R. J. Webber, D. A. Plotkin, M. E. O'Neill, D. S. Abbot, and J. Weare, "Practical rare event sampling for extreme mesoscale weather," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 29, no. 5, p. 053 109, 2019.

[48] J. Morio and M. Balesdent, *Estimation of rare event probabilities in complex aerospace and other systems: a practical approach*. Woodhead publishing, 2015.

[49] J. Morio, R. Pastel, and F. Le Gland, "An overview of importance splitting for rare event simulation," *European Journal of Physics*, vol. 31, no. 5, p. 1295, 2010.

[50] F. Cérou, P. Del Moral, T. Furon, and A. Guyader, "Sequential monte carlo for rare event estimation," *Statistics and computing*, vol. 22, no. 3, pp. 795–808, 2012.

[51] P.-T. De Boer, D. P. Kroese, S. Mannor, and R. Y. Rubinstein, "A tutorial on the cross-entropy method," *Annals of operations research*, vol. 134, no. 1, pp. 19–67, 2005.

[52] J. Huang, "Maximum likelihood estimation of dirichlet distribution parameters," *CMU Technique report*, pp. 1–9, 2005.

[53] M. Balesdent, J. Morio, and L. Brevault, "Rare event probability estimation in the presence of epistemic uncertainty on input probability distribution parameters," *Methodology and Computing in Applied Probability*, vol. 18, no. 1, pp. 197–216, 2016.

[54] I. Papaioannou, C. Papadimitriou, and D. Straub, "Sequential importance sampling for structural reliability analysis," *Structural safety*, vol. 62, pp. 66–75, 2016.

[55] F. Wagner, I. Papaioannou, and E. Ullmann, "The ensemble kalman filter for rare event estimation," *SIAM/ASA Journal on Uncertainty Quantification*, vol. 10, no. 1, pp. 317–349, 2022.

[56] J. Bakosi, J. Ristorcelli, *et al.*, "A stochastic diffusion process for the dirichlet distribution," *International Journal of Stochastic Analysis*, vol. 2013, pp. 1–7, 2013.

[57] S. A. U. Islam and D. S. Bernstein, "Recursive least squares for real-time implementation [lecture notes]," *IEEE Control Systems Magazine*, vol. 39, no. 3, pp. 82–85, 2019.

[58] H. Lamba, J. C. Mattingly, and A. M. Stuart, "An adaptive euler–maruyama scheme for sdes: Convergence and stability," *IMA journal of numerical analysis*, vol. 27, no. 3, pp. 479–506, 2007.

[59] W. Schneeweiss, "Approximate fault-tree analysis with prescribed accuracy," *IEEE transactions on reliability*, vol. 36, no. 2, pp. 250–254, 1987.

[60] L. Ren and J.-P. Clarke, "Separation analysis methodology for designing area navigation arrival procedures," *AIAA Journal of Guidance, Control, And Dynamics*, vol. 30, no. 5, pp. 1319–1330, 2007.

[61] A. P. Kendall and J.-P. Clarke, "Stochastic optimization of area navigation noise abatement arrival and approach procedures," *Journal of Guidance, Control, and Dynamics*, vol. 43, no. 4, pp. 863–869, 2020.

[62] S. Pullen *et al.*, "Impact of ionospheric anomalies on gbas gast d service and validation of relevant icao sarps requirements," in *Proceedings of the 30th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2017)*, 2017, pp. 2085–2105.

[63] J. Rasmussen, "Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models," *IEEE transactions on systems, man, and cybernetics*, no. 3, pp. 257–266, 1983.

[64] M. Gopinathan, J. D. Boskovic, R. K. Mehra, and C. Rago, "A multiple model predictive scheme for fault-tolerant flight control design," in *Proceedings of the 37th IEEE Conference on Decision and Control (Cat. No. 98CH36171)*, IEEE, vol. 2, 1998, pp. 1376–1381.

[65] H. R. Idris, Q. Dao, R. C. Rorie, and K. Hashemi, "A framework for assessment of autonomy challenges in air traffic management," in *AIAA AVIATION 2020 FORUM*, 2020, p. 3248.

[66] I. Hwang, H. Balakrishnan, and C. Tomlin, "State estimation for hybrid systems: Applications to aircraft tracking," *IEE Proceedings-Control Theory and Applications*, vol. 153, no. 5, pp. 556–566, 2006.

[67] RTCA, *Do-178c software considerations in airborne systems and equipment certification*, 2011.