Multilateration based ADS-B validation using a Particle Filter

Master Thesis

T.D. Landzaat

Delft University of Technology

TU Delft

LVNL

Multilateration based ADS-B validation using a Particle Filter

Master Thesis

by

T.D. Landzaat

This thesis work is a collaboration between the student, LVNL and the TU Delft.

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday June 30, 2023 at 14:00.

Thesis Committee	
Supervisor	Dr. ir. Hans Driessen
Responsible Full Professor	Prof. Dr. Alexander Yarovoy
Thesis Committee Member	Dr. Geethu Joseph
Company Supervisor	Dr. Hans van Hintum

Project Duration: September, 2022 - June, 2023
Faculty: Faculty of Electrical Engineering, Mathematics and Computer Science, Delft

Cover: LVNL Internal

TU Delft

LVNL

Preface

With great pleasure, I hereby present my thesis, the culmination of an intense period of research and hard work, it was a collective effort made possible only through the invaluable support and guidance of various individuals.

I would like to express my sincere gratitude to my family, my girlfriend and my friends for their unconditional love, constant support, and encouragement that has inspired and motivated me throughout my time at TU Delft. I would also like to thank Professor Hans Driessen, whose knowledge, valuable insights, and stimulating guidance have been invaluable to my academic growth. I would like to express my sincere gratitude to my supervisor at LVNL, Hans van Hintum, and all of my colleges whose expert guidance, and valuable feedback helped me throughout this entire process. Finally, I thank all the professors, teachers and fellow students of the MS3 group for their guidance, and inspiring lectures.

As a born and raised Delftenaar, I am proud to achieve this Masters degree in Science in *my* city, the city who will always be my home. But, my time as a student is coming to an end. New challenges await at my new job at the Netherlands Aerospace Centre in Amsterdam...

T.D. Landzaat
Delft, June 2023

Summary

Automatic Dependent Surveillance Broadcast (ADS-B) allows aircraft to broadcast their own position, speed, altitude, and other information to ground stations and other nearby aircraft. This information is then used by air traffic control for situational awareness, and collision avoidance. ADS-B spoofing is possible due to the lack of authentication and encryption in the ADS-B protocol. This can result in incorrect decision-making and potential safety hazards. Validation of the location of the ADS-B message is required for Luchtverkeersleiding Nederland (LVNL) such that it can maintain separation minima between civil aircrafts, whilst using ADS-B operationally.

Analysis of possible approaches for ADS-B validation has resulted a multilateration (MLAT) based approach. Time of Arrival (TOA) measurements of the ADS-B messages are used to validate the location. For validation, at least two Ground Stations (GS) are required instead of the four GSs required for a MLAT track, allowing for ADS-B validation in a larger area than it is currently used for in a tracking application. If an ADS-B message is considered validated, its content can be used by ATC. Therefore MLAT based validation results in an increased surveillance coverage. Validation is achieved in two steps, first a tracker is used to compute the state of the target using the TOA measurements, secondly this state is compared to the ADS-B location using a likelihood ratio test.

Tracking is done using a Sequential Importance Resampling (SIR) Particle Filter (PF). Classical PF issues as the degeneracy problem and sample impoverishment problem are mitigated by using a novel sampling method that samples directly from the measurement at the initialization of the SIR filter. Without this novel method a traditional SIR filter (where the proposal density is uniformly distributed) requires roughly a million particles to converge on the location of the target. Below this amount of particles the traditional SIR filter fails. The proposed SIR filter can converge on the location of the target using only 1000 particles.

To provide LVNL options and insights, three different likelihood ratio tests are proposed, namely the Minimum Bayes Risk, The Neyman-Pearson and the Minimax Hypothesis test.

Performance of the algorithm is investigated using a case study where data from LVNL's Surveillance Data North Sea (SDNS) MLAT system is used. Results have found that each test is capable of correct ADS-B validation. The limiting factor in the validation algorithm is the quality of the state estimate. At lower altitudes (<FL20) state estimation can fail and therefore also the hypothesis test. Above this altitude, spoofed targets can be detected if the distance between the spoofing transmitter and the location inside the spoofed message is roughly 1000 to 2000 meters depending on the hypothesis test used. Horizontally, this falls within LVNL's separation minima, vertically, this falls outside the separation minima.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
ADS-B	Automatic Dependent Surveillance Broadcasting
ANSP	Air Navigation Service Provider
ARTAS	ATM Surveillance Tracking and Server
ATC	Air Traffic Control
ECEF	Earth Centered Earth Fixed
ESS	Effective Sampling Size
ENU	East, North, Up
FL###	Flight Level ###
GS	Ground Station
GSP	Geographic Positioning System
ICAO	International Civil Aviation Organization
INS	Internal Navigation Sysem
LR	Likelihood Ratio
LRT	Likelihood Ratio Test
NP	Neyman-Pearson
MBR	Minimum Bayes Risk
MLAT	Multilateration
PSR	Primary Surveillance Radar
TDOA	Time Difference of Arrival
TOA	Time of Arrival
SDNS	Surveillance Data North Sea
SSR	Secondary Surveillance Radar
VDOP	Vertical Dilution of Precision
WAM	Wide-Area Multilateration

1

Introduction

In section 1.1 Air Traffic Control The Netherlands (LVNL) will be introduced. After this in section 1.2 ADS-B will be explained including its shortcomings. From this in 1.3 the motivation of the thesis is elaborated leading to the thesis objective in 1.4. The final section of the chapter provides the structure of the thesis.

1.1. Air Traffic Control The Netherlands

Air Traffic Control The Netherlands is the Dutch Air Navigation Service Provider (ANSP), founded in 1923. Currently about 1300 employees work for the company, of which roughly 250 are air traffic controllers. LVNL provides Air Traffic Control (ATC) in the civil airspace, which is required by law. This law specifically states that LVNL must provide communication, navigation and surveillance services (commonly abbreviated by ANSPs as CNS). LVNL also provides other services such as air maps, optimal use of runways and educating new air traffic controllers. Providing CNS services is a complex task. LVNL is expected to keep these services up and running 100% of the time for airports such at Rotterdam The Hague Airport and Schiphol Airport. Fig. 1.1 shows all sites in operation by LVNL and military operated airports.

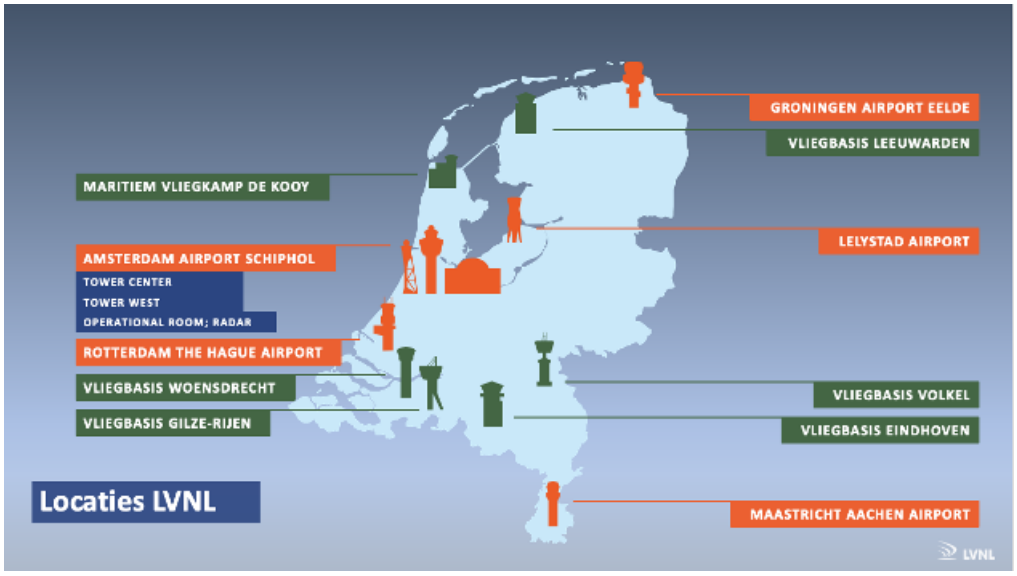

Figure 1.1: Shown is an overview of the sites that LVNL operates in orange. Green are military operated airports. Blue gives all the different locations present at Schiphol Airport

These services are provided using several technical systems. For navigation services Distance Measuring Equipment (DME) is used. This system is used to determine the distance from a plane to the

DME itself. Another example is the Instrument Landing System (ILS). An ILS guides an aircraft on descent to the runway helping the pilots on final approach. DMEs and ILSs are widely used navigation equipment for civil aviation. These are two examples to illustrate the operational equipment of LVNL allowing 61.3 million passengers [7] to travel from and to the Netherlands. Fig. 1.2 shows a schematic overview of the surveillance systems used by LVNL. There are currently four surveillance techniques in operation at LVNL which can be seen in Fig. 1.2.

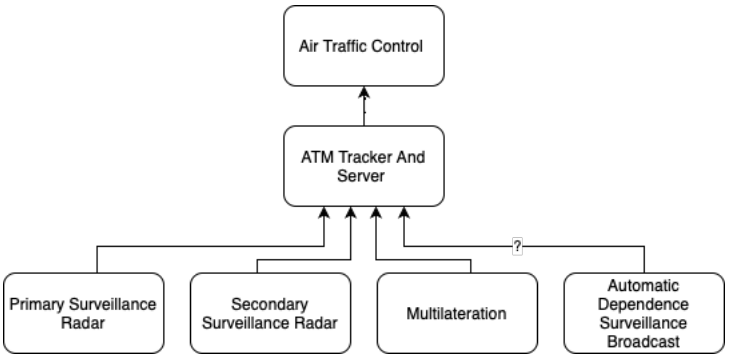

Figure 1.2: Schematic representation of LVNL surveillance systems.Question mark indicating validation of the data not done.

1.1.1. Primary Surveillance Radar

Primary Surveillance Radar (PSR) is a traditional radar where a rotating antenna illuminates a target with electromagnetic waves. The location is determined by the round trip time of the signal, and the incident angle of the waves. PSR is independent and non-cooperative radar. This means that the target does not participate in aiding the detection, and independent because the radar itself determines the location of the target.

1.1.2. Secondary Surveillance Radar

Secondary Surveillance Radar (SSR) / Mode-S Radar is an evolution of PSR. Here the target is equipped with a transmitter and responder, accordingly named a transponder. This transponder can be interrogated by the radar at 1030 MHz, and the target responds at 1090 MHz. With such an interrogation the altitude can be requested among other flight information. A type of interrogation is the All-Call. With an All-Call all the transponder equipped aircrafts responds to interrogation of the radar with its personal Mode-S address. Each aircraft has its own Mode-S address distributed by ICAO[1]. All documentation describing what information can be requested is publicly available. Requesting specific data from an aircraft using an interrogation is what is called interrogator modes. Several iterations of this system have been developed and used operationally. The current system is called Mode-S. Here the S indicates *Select*, illustrating the capability of the SSR to select which target it wants to interrogate. All the information obtained is only received upon interrogation by the SSR.

Secondary Surveillance Radar is a cooperative radar. The radar needs the transponder of the aircraft to respond otherwise it is not able to determine the location of the target. But if the target responds it is able to determine the location itself making it independent. A SSR is not able to detect a target via an echo only.

1.1.3. Multilateration

The third source is Multilateration (MLAT). MLAT has a significantly different working principle than primary or secondary radar. MLAT consists of several (at LVNL up to 37) receivers placed across a certain area. MLAT uses Mode-S interrogations where the target replies to using its transponder. Multiple Ground Stations (GS) in the MLAT system register this reply and assign a Time of Arrival (TOA). The time of transmission is not available at each GS therefore the system uses Time Difference Of Arrival (TDOA) to determine the location of the target. For this at least four GS are needed, which will be explained in detail in chapter 3. Similar to SSR this principle is cooperative and independent. But

[1]International Civil Aviation Organization

MLAT is not limited to Mode-S replies. Any received signal by a MLAT system can be used to compute the target's location. MLAT systems can be divided into two main areas of application, the first one is Wide-Area Multilateration (WAM), such a system is designed to cover a large area. Second application of MLAT is close range application, frequently used for ground surveillance at civil airports.

Automatic Dependent Surveillance Broadcast (ADS-B)
ADS-B will be discussed in further detail in section 1.2.

Multi Sensor Tracker
To combine all these surveillance sensors a multi-sensor tracker is used, namely ATM Tracker and Server (ARTAS). Offered by EUROCONTROL this multi-sensor tracker is used by 23 European ANSPs. The track computed by ARTAS is the track that is shown to ATC.

1.2. Automatic Dependent Surveillance Broadcast

Automatic Dependent Surveillance Broadcast (ADS-B) is a system where an aircraft transmits its location using the Mode-S transponder[2], as is illustrated in Fig. 1.3. This information is broadcast periodically every 0.5 seconds. A fundamental difference between ADS-B and Mode-S is that ADS-B always broadcasts, and Mode-S only replies when it receives an interrogation. Information that is transmitted contains also other information besides location such as heading, speed and other flight data. The aircraft uses GPS and the Internal Navigation System to determine it's own location. Notable difference between ADS-B and SSR is in the way how ATC obtains the location of a target. Using SSR the radar determines the location, and with ADS-B the aircraft determines the location.

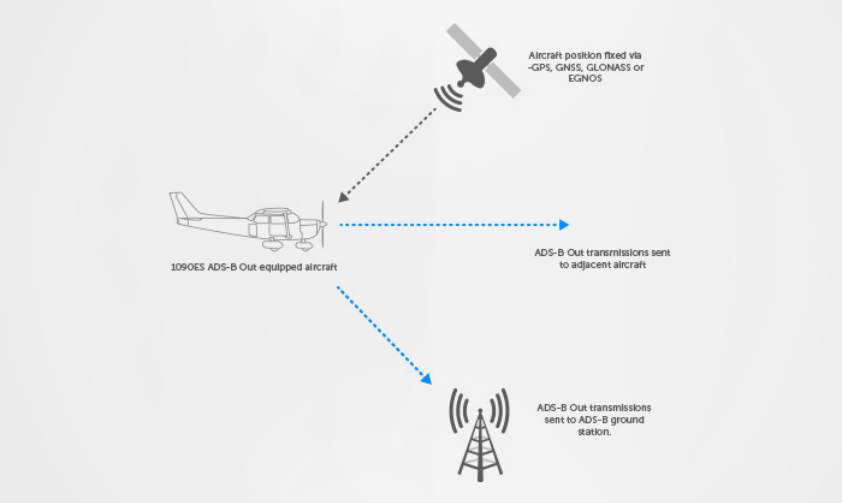

Figure 1.3: Schematic overview of working principal of ADS-B. Illustration from [1].

Advantages of ADS-B
The biggest advantage of ADS-B is that obtaining the location of an aircraft is now done using receivers and the aircraft's GPS or INS system. Full implementation of ADS-B could thus make SRR obsolete. It has been shown in practice by NAV-CANADA's ADS-B system in the Hudson Bay area that the implementation is very beneficial in cost saving, and also showed [8] they were able to optimize flight paths and save on CO2 emissions because of this. In addition to this the accuracy of ADS-B does not decrease when the range increases [9].

Disadvantages ADS-B
The content of an ADS-B message is of a predefined structure [10]. With ADS-B being an extension of Mode-S, the message protocol is not encrypted similar to Mode-S. With regular Mode-S replies this

[2]Specifically the Mode-S Extended Squitter

does not pose a problem. Because Mode-S is always the result of a interrogation, the messages are one to one correlated (interrogation to reply) and by this the origin of the message is determined. In addition to this the location of the target is still determined by the radar itself, ensuring that at least the location of the target is determined independently. ADS-B is a general broadcast of the location and not correlated to an interrogation. This means that received messages are not automatically validated. Validation in this context means that the contents of the ADS-B message is true flight information of the target that transmits the message, i.e. the content of the message is not manually altered due to bad intent.

Received ADS-B messages can be of a real or fake target, due to this there are great concerns about the security of ADS-B. McCallie provided a clear overview of the threats and vulnerabilities [11]. For LVNL the most important threat is,

- **Ghost Target**: By generating a fake ADS-B message one can inject a ghost target into the ATC radar system.

[12] demonstrated there are several types of ghost target attacks and [13] shows how to create such a spoofed ADS-B message, also is concluded that an ADS-B attack is inexpensive and possibly highly successful. Detailed instructions are thus available on how to spoof ADS-B messages indicating that for a willing attacker with some technical background it is possible to inject a ghost target into an ATC system. NAV-CANADA and Air-services Australia, but also LVNL is using ADS-B operationally. Implementation of ADS-B at LVNL is done via whitelisting. a small number of helicopters that supply oil rigs in the North Sea are added to this list and shown to ATC. On the ATC screen it can occur that this target is only detected by ADS-B. If this is the case it is indicated on the radar screen that this target is an ADS-B only target to provide an extra safety measure. These two safety measures where implemented as a quick fix, and have several disadvantages.

- White-listing of aircraft must be done manually and is time consuming as LVNL has the intention of tracking all ADS-B equipped aircraft.
- ICAO addresses are publicly known and if a spoofed ADS-B message uses a valid ICAO address that is added to the white-list, the safety measure is not able to stop the spoofed messages from appearing on the ATC radar screen.
- Showing the possible safety risk to ATC adds complexity to their task. This is undesirable as the task of ATC is to separate aircraft, not to determine and mitigate spoofed ADS-B messages.

1.3. Motivation

Due to the advantages of ADS-B LVNL seeks a proper validation method, that does not use white-listing or ADS-B only indicators. LVNL has installed receivers that are able to receive ADS-B messages in multiple locations throughout The Netherlands. These receivers are the same receivers used in MLAT systems in operation by LVNL. They allow for quick and widespread implementation of ADS-B if the messages can be validated.

Added value to MLAT

LVNL is able to receive ADS-B messages with two surveillance systems. The first one is a MLAT system located at Schiphol Airport named Cooperative Surveillance System (CSS) and is used for observing aircraft on the ground at the airport. CSS has its own tracker and is incorporated into the advanced-surface movement guidance and control system (A-SMGCS). This system combines several important tasks that are needed to operate the airport such as GARDS, a system to notify ATC that an aircraft has executed a go-around, or to control runway crossing lights. This MLAT system enables LVNL to track aircraft with high accuracy on the ground which has significant added value. Ground vehicles are also equipped with ADS-B resulting in a complete overview of all moving targets on the airport. The second system able to receive ADS-B messages is the Surveillance Data North Sea in short SDNS. SDNS is a WAM based surveillance system designed to provide coverage on the North Sea, which is part of the Dutch airspace and can be seen in Fig. 1.4. There is a total of 37 receivers making up the entire surveillance system. These receivers are installed on oil rigs in the North Sea. Each of these receivers in SDNS is also able to receive and decode the contents of the ADS-B message. LVNL is in the

deployment of a third WAM system called Wide-Area Multilateration Nederland (WAM-NL). WAM-NL is designed to provide full multilateration and ADS-B coverage on land. Combining WAM-NL with SDNS the entire Dutch airspace is covered by MLAT and ADS-B.

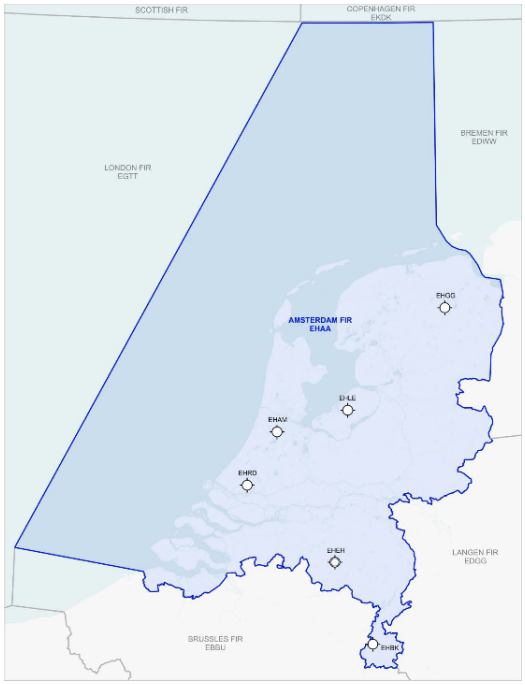

Figure 1.4: Flight Information Region for Dutch airspace. Illustration from [2].

For SDNS, ADS-B validation can be of particular importance. This airspace is used at low altitude to supply oil rigs in the North Sea. In addition to this there is very limited SSR coverage at low altitude due to the curvature of the earth. The placement of the receivers is defined by the location of the oil rigs meaning that ideal placement of all GSs is not possible, resulting in that the coverage is not always sufficient. MLAT requires four receivers to determine the location of a target and ADS-B only requires one receiver. Using ADS-B thus increases the coverage for this WAM system, increasing the desire to validate ADS-B. The coverage is increased because the area that is covered by one receiver is always larger then the area that is covered by the overlapping area of four receives. This will also occur with WAM-NL, or any other MLAT/WAM system where all GS's are able to receive and decode ADS-B messages. Fig. 1.5 schematically shows the surveillance systems including SDNS and WAM-NL. CSS system located at Schiphol Airport is not taken into account, because this is used in LVNL's A-SMGCS.

Added value to ARTAS

As Fig. 1.5 illustrates, ARTAS combines all surveillance sources. ADS-B has been deployed throughout the years in several iterations starting from V0 to V2. First versions where very unreliable and inaccurate. Nowadays the reported location of V2 equipped aircraft is of high quality and reliable. Disadvantage of the ADS-B message protocol is that is does not indicate the version of the system, thus for a received message is it impossible to know if the airplane used V0 or V2, and thus the accuracy of the location report. This introduces another significant benefit to ADS-B validation because incorrect messages can also be detected and mitigated. ARTAS can thus use make use of ADS-B validation to detect spoofing, but also to detect errors in the ADS-B message, that are not the result of spoofing.

Added value to ATC

Providing ATC services is a large and complex task. To achieve this, extensive safety measures are taken to ensure safe civil aviation. Air traffic controllers rely on their radar screen to provide correct and validated radar positions. If ADS-B would be allowed on the same radar screen this introduces

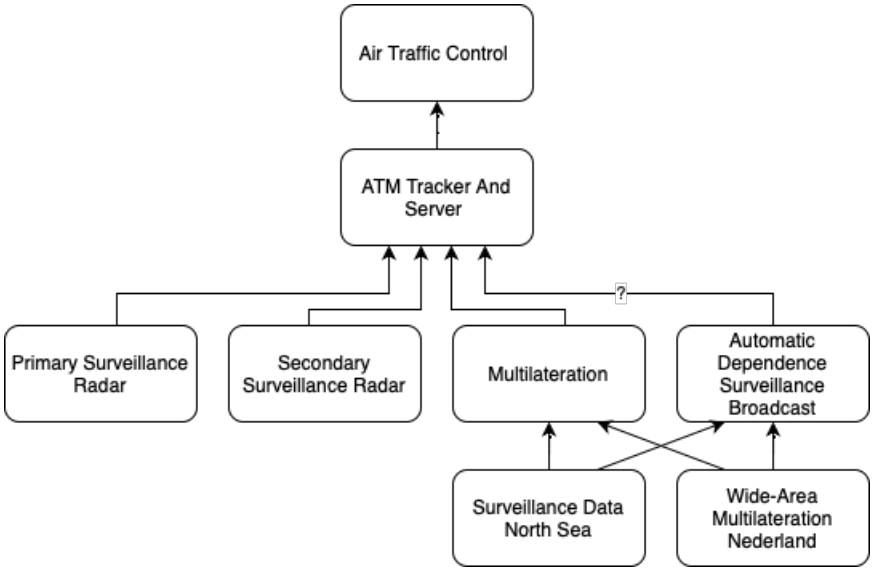

Figure 1.5: Schematic representation of LVNL surveillance systems, including both MLAT systems. Question mark indicating that the implementation of ADS-B validation is still unknown.

additional complexity to this task, as ATC now also has to determine if the message is correct or spoofed. This task of validating the ADS-B location is not the responsibility of ATC and thus must be handled with accordingly prior to showing ADS-B to ATC.

In short the motivation is as shown below in Tab. 1.1

Motivation	Area of application
ADS-B able systems are in place	SDNS (WAM-NL)
Increased coverage	SDNS (WAM-NL)
Added redundancy	ARTAS
Higher data quality	ARTAS
Reducing ATC complexity	ATC

Table 1.1: Motivation for ADS-B validation and which area it applies to

1.4. Thesis Objective

The goal of the thesis is to design a validation algorithm capable of validating the location reported in an ADS-B message and investigate its performance.

The main job of ATC is to separate traffic according to a separation minimum. At the core they need to know the reliable position and velocity of the aircraft for which they are responsible. Position will thus be the parameter validated.To validate ADS-B messages there are numerous approaches possible which will be discussed in chapter 2.

This thesis presents the design of the proposed algorithm, including important design choices. The performance of the algorithm is analyzed using a case study.

1.5. Thesis Structure

After the introduction in chapter 1, the second chapter consist of the literature review. In this literature review several methods are discussed to achieve the goal set out by the thesis. From this review the novelty of the proposed solution follows. Detailed discussion of the designed validation algorithm is done in chapter 3. To obtain results and analyze the performance of the proposed algorithm, a case

study is done. In chapter 4 this is introduced along with pre-processing steps required to obtain the desired results. In chapter 5 the results of the case study is presented and certain spoofing scenarios are investigated. In the final chapter of the thesis recommendations and conclusions are given.

2
Literature Review

In this literature review the goal is to review approaches that can be used to validate ADS-B messages. The findings are presented in three main sections. In section 2.1 encryption as a solution is discussed followed by section 2.2 where Machine Learning is investigated. The final approach is measurement based and discussed in section 2.3. Section 2.4 considers advantages and disadvantages of all proposed methods and presents the chosen solution. In the final section the novelty of this solution is presented.

2.1. Encrypting ADS-B

Proposed by [14] and [15] is to encrypt the ADS-B protocol. One of the methods discussed is Symmetric-Key encryption. With this encryption method the transmitter and receiver share a key which is used to encode and decode the ADS-B message. By distributing these keys between ATC and trusted ADS-B equipped airliners it is impossible to spoof the ADS-B message. Encryption as validation system relies on the trust between ATC and the airliners. Airliners still have the theoretical possibility to transmit spoofed ADS-B messages because they own the keys that allow for the encryption. Disadvantage of this method is that a single key leak renders the system unusable. Asymmetric-Key Cryptography mitigates this problem. Here, sets of keys are distributed where each ADS-B equipped plane has his own private key assigned. If such a private key will be hacked or leaked only a single ADS-B equipped aircraft is compromised. Other approaches have been proposed such as [16] where Identity-Based Encryption is discussed. With this approach less complex key management is achieved. It has been recognized by Yang *et al.* in [17] that proposing key based cryptography is notoriously difficult. Yang *et al.* aims to use the existing ADS-B infrastructure for encryption by proposing an authentication method that makes use of the existing ADS-B fields that are not in use. It has been demonstrated on real ADS-B data from the OpenSky network that their proposed encryption method is compatible with existing ADS-B systems.

Encryption is from a technical point of view a good solution. Key management and acceptation has been found to be a fundamental disadvantage of the solution. Suggestions have been done that ICAO could be responsible for key management, but is it likely that countries want to manage these keys themselves because they form a vital part of their national infrastructure. These management challenges need to be resolved before one can decide on any encryption method. Another disadvantage of encryption of ADS-B is the deployment. Airliners are reluctant to implement changes to the system due to associated costs.The FAA (Federal Aviation Administration) mandated in a regulation published May 2010 that aircraft must be ADS-B equipped by January 2020. This spans one decade indicating that the rate of adaption is slow. Introducing fundamental changes to the system will also likely take at least a decade. For LVNL this time period is to long and another approach should be considered.

2.2. Validation using Machine Learning

Proposed in literature is the Machine Learning (ML) approach. ML has several advantages when used for ADS-B validation. Performance of ML increases when it has more data to learn from. ADS-B messages are already transmitted every day and thus millions of messages are stored which can be used for training. ML is a rapid growing field that is rapidly innovating, indicating that in the

near future performance may increase. One ML method is feature based where a feature could be the limits of the velocity of the target, which must not exceed the physical maximum velocity of an airplane. Features must be hand crafted and thus require extensive development. Another approach is letting a Neural Network (NN) learn features by itself. Multiple options are available as input for a NN such as Ying *et al.* demonstrate in [18] by using the physical data link and [19] shows that long short-term memory (LSTM) provides promising results in detecting spoofed information such as speed, heading, and climb rate. Advantage to this specific proposed approach is that it can detect what parameter of the ADS-B message is fake and when, due to the use of a sliding window. Applied Logistic regression, Naive Bayes and K-Nearest Neighbor are methods explored in [20] to detect fake ADS-B messages on a total of four attack types, namely; false information (location and velocity), false heading and a jumping attack where the location changes randomly. They achieve a very high F1-score for all types of attack. A significant disadvantage to ML is that it cannot detect real ADS-B messages replayed at a later time. In this type of attack a spoofer records ADS-B messages from real targets, and replays these at a later time. Any ML model will not be able to detect such an attack because all properties of the ADS-B messages are real. This introduces a significant vulnerability that cannot be mitigated.

2.3. Validation using Measurements

In this section measurement bases validation methods are discussed. All methods share a common working principle that is that location, time or Doppler shift measurements are made and these measurements must correlate with the reported location or velocity inside the ADS-B message.

2.3.1. Primary / Secondary Radar

PSR and SSR determine the location of a target that is transmitting ADS-B messages. The ADS-B message and the radar measurement can be compared and it can be concluded that the message is spoofed or not. Advantage to this methods is that all hardware is available and operational. Disadvantage is that where the radar has coverage, ADS-B is of limited added value because the location of the target is already determined. In addition, at large distances from the radar the quality of the estimate decreases, which decreases the certainty in the location validation.

2.3.2. Doppler shift

The method of measuring the Doppler shift of an ADS-B message has been proposed in [21]. With this method a GS verifies ADS-B messages by computing the expected Doppler shift based on the reported location and velocity and compares this to the measured Doppler shift. The measured Doppler shift is computed by sampling the frequency of the carrier wave of the received ADS-B message. Using this method the radial velocity is estimated. This estimated radial velocity must be equal to the radial velocity that can be determined from the position and velocity reported by ADS-B. This validation method can mitigated by the attacker if he/ she calculates the relative velocity between the GS receiving the ASD-B message and the ghost aircraft that it wants to inject, and adjust the transmitted frequency such Doppler shift that the GS measures is equal to the Doppler shift the GS computes. In practice it is highly likely that the ADS-B message will be received by two or more GSs, for the spoofer to account for this, he/ she must be able to calculate the Doppler shifts with respect to all GSs that receive the ADS-B message and transmit modulated frequencies accordingly. Remaining undetected thus adds extreme complexity that rapidly increases when the ADS-B message is received by multiple GSs, which is almost always the case in practice. ADS-B validation based on Doppler shift measurements is expected to have very good performance due to the complexity to remain undetected. Disadvantage of the method is that new hardware is required. This hardware must be able to measure the Doppler shift with very high precision.

2.3.3. Angle Of Arrival (AoA)

One can determine the angle of arrival of an ADS-B message using uniform linear array's (ULA).The ULA approach is presented in [22] and [23]. Based on the different phase that is measured at each element of the array the ULA is able to determine if the angle of arrival correlates with the claimed position in the ADS-B message. AoA validation means that the direction of the incident EM wave must

correlate with the relative angle between the reported ADS-B location and the measuring equipment. At large distances from the ULA the quality of the measurement deteriorates, resulting in less accurate AoA estimation. A less accurate AoA measurement increases the possibility of a fake ADS-B message to remain undetected. For AoA validation the area where spoofing remains undetected is determined by the resolution of the measurement and the distance from the ULA. Higher angular resolution reduces the size of this area similar for targets closer to the ULA.Using this validation method multiple ULAs must be installed to cover the entire Dutch airspace which results in high costs. In addition they need to cover large distances with means deteriorating performance.

2.3.4. Time Difference of Arrival (TDOA)

A TDOA measurement is the result of subtracting two Time of Arrival (TOA) measurements. This approach is commonly used to estimate the location of a target when the time of transmission is not known by the receiver. TDOA measurements can be used to validate the received ADS-B location. Determining the location of a target using TDOA has been proposed in [24] and the performance is analyzed and a linearized least squares solution is presented to estimate the location of the target. A more detailed solution including a geometric perspective is given in [25]. A single TDOA measurement creates a surface in space, this surface represents all possible target locations. When three TDOA (i.e. the ADS-B message is received at four GSs) measurements are available all these surfaces intersect in one point. This point of intersection is the location of the target and ADS-B validation can then be performed by comparing this location to the received ADS-B location. When only two or three GSs contribute to the TDOA measurement, validation can still be done as the ADS-B location must still intersect the surfaces.

TDOA measurements are the working principle of multilateration. SDNS and WAM-NL are such WAM systems where the location of a target is computed using TDOA. Using TDOA as validation approach has one significant advantage which is that no new hardware is required as the TDOA measurements are already available in SDNS (and in the future WAM-NL). In SDNS the TOA measurement that are used to compute the TDOA measurements are accurate up to 3 ns. The TDOA measurement are thus of very high precision. Validation is done by verifying if the reported ADS-B location correlates to the TDOA measurements made by the multilateration system.

2.3.5. Frequency Difference of Arrival (FDOA)

Based on relative velocity between the transmitter and receiver a Doppler shift occurs. FDOA is a less popular method for estimating location and velocity then TDOA due to it's complexity. But several papers such as [26] and [27] have presented a solution. It is commonly applied in combination with TDOA. It can be advantageous to only use FDOA measurements when the received signal is narrowband [28].

For validation purposes using FDOA in combination with TDOA will lead to better spoofing detection. This is because the system is able to determine position and velocity with an higher accuracy. But similar to the problem of measuring Doppler shift, to use FDOA the Doppler shift must also be available for this approach to work, which at SDNS or WAM-NL they are not. It is expected that this gain in accuracy will not lead to significantly better spoofing detection compared to TDOA based validation.

2.4. Comparison of Methods

Multiple methods to validate the reported ADS-B location are discussed in this chapter. To choose between the available options a decision matrix is made as seen in Tab. 2.1. Important aspects of the validation approaches are given a score. This score is based on the relative advantages/disadvantages compared to one another.

Encryption based ADS-B is impossible to spoof, but there are significant management challenges that render it impractical.

The ML method arises as a robust solution because it can validate all fields inside an ADS-B message which no other method is able to do, but any ML model must be trained to do so which is a complex

task because there are numerous data registers each with their own characteristics. Validation based on all data fields adds significant complexity and if often not required by ATC. In addition to the disadvantages of ML it is not able to detect replayed ADS-B messages which introduces a significant vulnerability that cannot be mitigated.

The measurement based validation methods are only able to validate position, velocity and heading. From a ATC point of view this is exactly what is needed. Both AoA, Doppler estimation and FDOA require additional hardware to operate. This results in TDOA being the best solution for the LVNL to validate ADS-B messages. Therefore is it decided to design a validation algorithm based on TDOA measurements.

	Expected Performance	Feasibility	Implementability
Encryption	+++	+	- - -
Machine Learning	+	++	++
Angle of Arrival	-	-	-
Doppler Shift	+	+	- -
TDOA	++	+++	+++
FDOA	+++	- -	- - -

Table 2.1: Decision matrix for ADS-B validation methods.

2.5. Novelty

The novelty of this thesis is twofold,

- MLAT systems do not produce an output when a measurement consists of less than four GSs (or three TDOA measurements). This is because four GSs are needed to determine a point estimate of a target. In validation the problem is not to determine the location of a target, but to determine if the TDOA measurement is coherent with the ADS-B location. This means that at least one TDOA measurement is sufficient to validate. A validation algorithm based on this principle is novel.
- The presented Particle Filter in chapter 3 uses a novel initialization method that allows for efficient state estimation at a low computational cost.

3

Validation Algorithm

The goal of this chapter is to present the design of the validation algorithm. The preceding literature review has shown that TDOA based ADS-B validation will be the working principle of the algorithm, therefore in 3.1 the theory behind TDOA localization will be explained in detail. The validation algorithm consist of two main parts, the first part is the filter discussed in section 3.2. The second part is the hypothesis testing used to determine if the ADS-B message under test is real or fake, which is discussed in section 3.3.

The validation algorithm has two data inputs which is schematically shown in Fig. 3.1. The first input is the ADS-B message the algorithm tries to validate. The second input is the TDOA measurement. TDOA measurements and ADS-B messages are coupled to one another based on ICAO the address that is reported inside the ADS-B message.

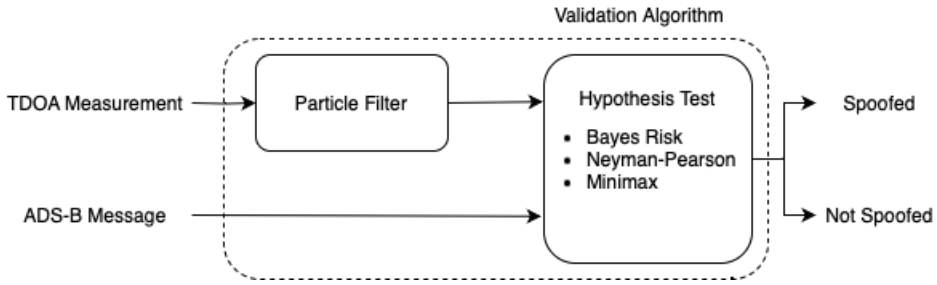

Figure 3.1: Schematic design of the validation algorithm.

There are several advantages in using a filter to determine the state of the target, and not the individual measurements to validate the target. These are:

- Faulty measurement do not directly result in the ADS-B message to be rejected, leading to less false alarms.
- It allows the filter to estimate the velocity of the target. This estimated velocity can also be validated, leading to more secure validation.
- Validating individual measurements only provides the information if the message is spoofed or not, tracking allows estimation of the location of a (possibly) spoofing target. In such a scenario traffic can still be separated from the actual location of the target that is transmitting spoofed messages. This allows LVNL to maintain the separation minima.

Due to these advantages a filter is used in the validation algorithm. The filter used is a Particle Filter which computes the state (location and velocity) of the target that is transmitting the ADS-B message that the algorithm intents to validate. This design choice is made because the filter must deal with the non-linear nature of MLAT localization, and deal with TDOA measurements that contain less than four

GSs. A traditional filter, such as the Kalman Filter is unable to accurately handle measurements that contain less than four GSs, as the measurement is ambiguous in location.

To determine if the ADS-B message under test is real or fake a hypothesis test is done. This test takes into account all the statistical properties of the ADS-B message and the state of the target as determined by the Particle Filter. Three types of hypothesis test are used in the algorithm, in each test the threshold is determined by a different criteria. By using several hypothesis test the performance can be compared and the best solution can be implemented. Choosing the best solution depends largely on the desires and procedures of LVNL. The final implemented validation algorithm can be one of the three, or a combination of the hypothesis test.

At the output each ADS-B message is assigned a label that indicates if the message is real or faked. If a sequence of messages is run through the algorithm this allows for robustness against measurement errors. How to finally decide then a target is actually spoofed or not is considered outside the scope of the thesis. Main reason for this is that the decision must be made in collaboration with Air Traffic Controllers and other stakeholders within LVNL.

3.1. Multilateration Theory

To measure the TDOA of a signal one first must determine the Time of Arrival (TOA), for this several Ground Stations (GS) are used. A GS is a remote system that houses all the hardware and software required to operate. Two GSs together can determine one TDOA measurement as shown in eq. 3.1 and Fig. 3.2 illustrates the situation schematically.

$$TDOA_{21} = TOA_2 - TOA_1 \tag{3.1}$$

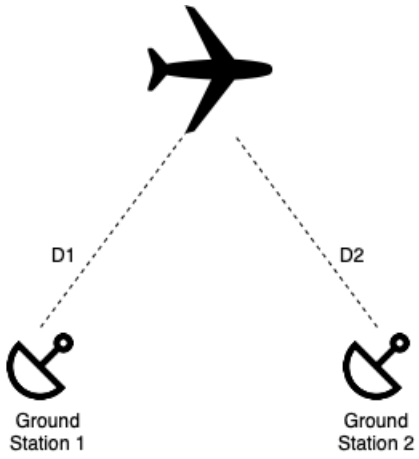

Figure 3.2: Basic working principle of MLAT. An ADS-B message is received at GS 1 and GS 2 and assigned a TOA_1 and TOA_2.

The difference between the measured time at $\mathbf{s}_1$ and $\mathbf{s}_2$ is equal to the difference in distance the signal has to travel divided by the speed of light. By expressing the distances $D1$ and $D2$ shown in Fig. 3.2 in terms of the locations of the GSs and the target, eq. 3.1 can be rewritten to eq. 3.2. $||\cdot||$ denotes the two-norm and $\mathbf{s}_n$ the location of the n'th GS in Cartesian coordinates. $\mathbf{l} = [l_x, l_y, l_z]^T$ where l denotes the location of the target expressed in Cartesian coordinates. T equals the time of transmission by the target.

$$TDOA_{21} = \left(T + \frac{D_1}{c}\right) - \left(T + \frac{D_2}{c}\right) = \frac{||\mathbf{s}_2 - \mathbf{l}|| - ||\mathbf{s}_1 - \mathbf{l}||}{c} \tag{3.2}$$

To estimate the location of the target $\mathbf{l} = [l_x, l_y, l_z]^T$ at least three TDOA measurements are needed, because then three equations are available to solve for three unknown, namely l_x, l_y, l_z.

Eq. 3.2 represents the shape of a two-sided hyperbola, because of this TDOA or multilateration based localization is often called hyperbolic localization. An illustration of a single TDOA measurement

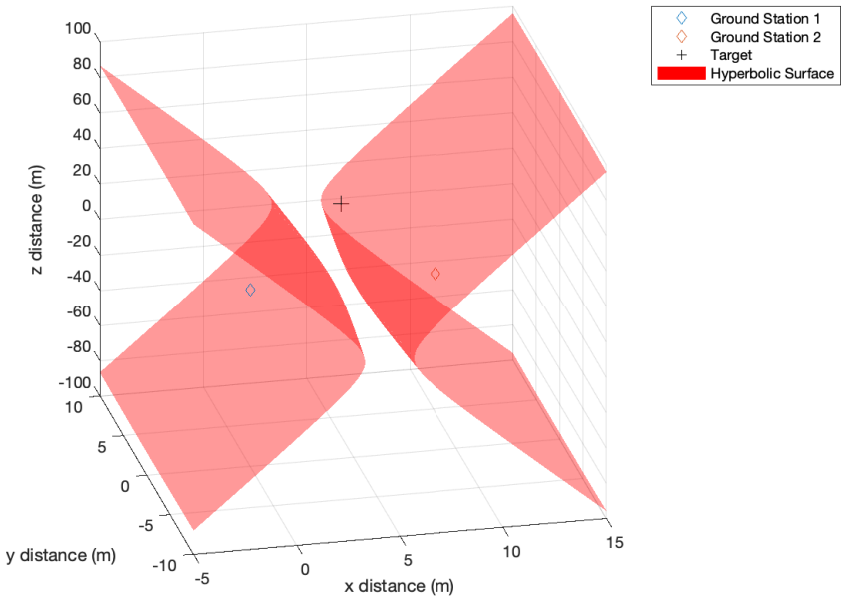

Figure 3.3: Red area represents the solution to the measurement equation. Note : This surface represents an unsigned TDOA measurement, meaning $TDOA = |TOA_1 - TOA_2|$.

is seen in Fig. 3.3. The two GSs receiving the TOA from the target are at height zero, but as can be seen the solution space also extends to negative z values. This means that for a single TDOA measurement the target could theoretically be underground.

When more than three TDOA measurements are available this leads to an overdetermined set of equations as shown in eq. 3.3. To solve for the location of the target several approaches are possible. The most common is the least squares (LS) approach as explained by [29]. LS approach is a method to solve an overdetermined system of equations. [29] Illustrated that Weighted Least Squares (WLS) can be used. A good overview and comparison of two others methods to estimate the location of the target is given by [30] namely, the Taylor Series approach and the Maximum Likelihood approach approach. All these approaches are capable of solving the equations to obtain the location of a target.

$$\mathbf{z} = \begin{cases} TDOA_{21} = \frac{1}{c}\left(\sqrt{(s_{2x}-l_x)^2+(s_{2y}-l_y)^2+(s_{2z}-l_z)^2} - \sqrt{(s_{1x}-l_x)^2+(s_{1y}-l_y)^2+(s_{1z}-l_z)^2}\right) \\ TDOA_{31} = \frac{1}{c}\left(\sqrt{(s_{3x}-l_x)^2+(s_{3y}-l_y)^2+(s_{3z}-l_z)^2} - \sqrt{(s_{1x}-l_x)^2+(s_{1y}-l_y)^2+(s_{1z}-l_z)^2}\right) \\ TDOA_{41} = \frac{1}{c}\left(\sqrt{(s_{4x}-l_x)^2+(s_{4y}-l_y)^2+(s_{4z}-l_z)^2} - \sqrt{(s_{1x}-l_x)^2+(s_{1y}-l_y)^2+(s_{1z}-l_z)^2}\right) \\ \vdots \\ TDOA_{N1} = \frac{1}{c}\left(\sqrt{(s_{Nx}-l_x)^2+(s_{Ny}-l_y)^2+(s_{Nz}-l_z)^2} - \sqrt{(s_{1x}-l_x)^2+(s_{1y}-l_y)^2+(s_{1z}-l_z)^2}\right) \end{cases} \tag{3.3}$$

The problem is shown geometrically in Fig. 3.4, where $N = 4$. When three independent TDOA measurement are available the location of the target can be calculated because eq. 3.3 provides a unique solution.

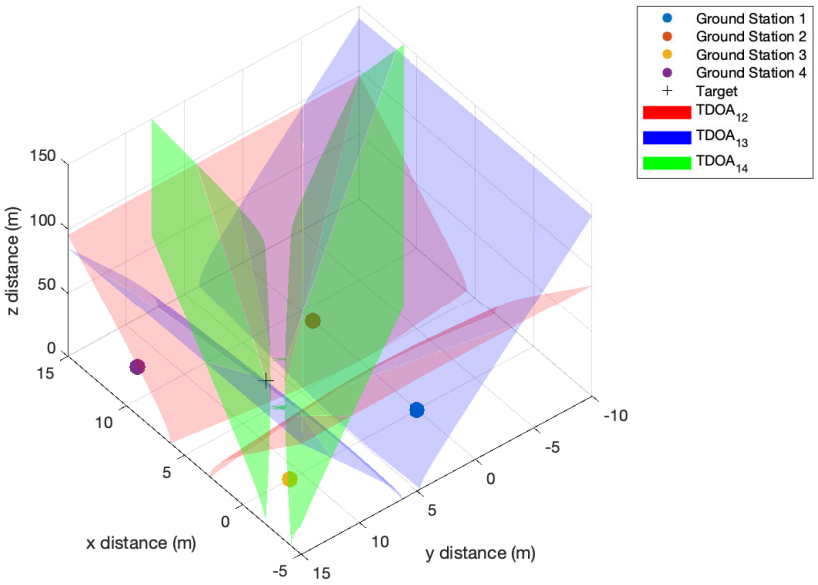

Figure 3.4: TDOA measurement where four GS provide an unique solution. This solution is the location of the target.

3.1.1. Dilution of Precision (DOP)

The generic shape of the hyperbola that results form a TDOA measurement can be seen in Fig. 3.4. The angle at which the hyperbolas intersect has effect on the quality of the location estimate. Every TOA measurement has some uncertainty, consequently the resulting hyperbola also has some uncertainty. To illustrate how this shape effects the uncertainty a simplified image of two hyperbolas is shown in Fig. 3.5. In this figure the blue line represents the measured hyperbola in 2D and the green dotted line its corresponding uncertainty. The red box indicates the associated uncertainty, meaning that somewhere in the red area the hyperbolas intersect. The size of the area depends on the angle of intersection between the hyperbola. A large area means a bigger uncertainty in the location estimate. Fig. 3.5a shows that when the angle of intersection of close to 180 degrees the uncertainty is large, and Fig. 3.5b shows that when the angle of intersection is close to 90 degrees this uncertainty is small. This phenomena is called Dilution of Precision (DOP). GSs in MLAT systems generally have good spread in the horizontal plane, resulting in good Horizontal Dilution of Precision (HDOP). In The Netherlands there is no possibility to put GSs at high altitude due to the absence of mountains. Therefore MLAT systems at LVNL have very little vertical spread of the GSs resulting in a bad Vertical Dilution of Precision (VDOP).

The degree of this effect depends entirely on the location of the target and the locations of all GSs. More TDOA measurements from different locations introduce more hyperbolas, if the geometrical spread of the GSs is favorable, the added hyperbolas diminish this effect. The effect that the GSs have on the location estimate is called Geometric Dilution of Precision (GDOP), where low values means good distributions of GSs in space. Bad GDOP is generally achieved when GSs lie in one line, and a good GDOP when GSs are evenly spread throughout space and thus all GSs have different incident angle w.r.t. to the incoming ADS-B message.

(a) The hyperbolas intersect an angle close to 180 degrees, resulting in a relatively large uncertainty

(b) The hyperbolas intersect an angle close to 90 degrees, resulting in a relatively small uncertainty

Figure 3.5: The blue lines in the figures represent the measured hyperbola in 2D and the dotted green line the according uncertainty. This uncertainty is indicated with the red box. Depending on the angle at which the hyperbolas intersect the area of uncertainty is larger or smaller.

3.1.2. TDOA combinations

When three GSs receive an ADS-B message three different TDOAs can be computed, namely $TDOA_{21}$, $TDOA_{31}$ and $TDOA_{32}$.

$$\begin{cases} TDOA_{21} = \frac{||\mathbf{s}_2-\mathbf{l}||-||\mathbf{s}_1-\mathbf{l}||}{c} \\ TDOA_{31} = \frac{||\mathbf{s}_3-\mathbf{l}||-||\mathbf{s}_1-\mathbf{l}||}{c} \\ TDOA_{32} = \frac{||\mathbf{s}_3-\mathbf{l}||-||\mathbf{s}_2-\mathbf{l}||}{c} \end{cases} \tag{3.4}$$

Geometrically this means that another hyperbola is added to the solution as is illustrated by Fig. 3.6. Fig. 3.6a shows the hyperbolic surfaces corresponding to $TDOA_{21}$ and $TDOA_{31}$, and 3.6b adds the third TDOA options, namely $TDOA_{32}$. This hyperbola intersects on the exact same locations as where $TDOA_{21}$ and $TDOA_{31}$ intersect. Due to this the exact location of the target is still unknown. In Fig. 3.6a the red and blue hyperbolas intersect in a total of four lines. By the addition of $TDOA_{32}$ there only remain two lines as possible solutions. Because of this $TDOA_{32}$ does provide added information on the locations of the target.

Often the situation arises in that an ADS-B message (or any message used by a MLAT system) that is received at 10 GSs. In such an situation there are 45 combinations possible.

$$C(n,r) = \binom{n}{r} = \binom{10}{2} = 45 \tag{3.5}$$

If there are N GSs, which means that there are $N-1$ independent equations that provide independent information on the location of the target. To greatly reduce the computational cost of the multilateration problem only $N-1$ equations are used to compute the location of the target. In theory this set of equations can be reduced further to three, because three equations provide a unique solution. using only three independent measurement equations could result in a bad location estimate if the GDOP is bad. Which combinations one must take to obtain a set of TDOA measurements is usually made by the designer of the system as it is a trade off between efficiency and robustness.

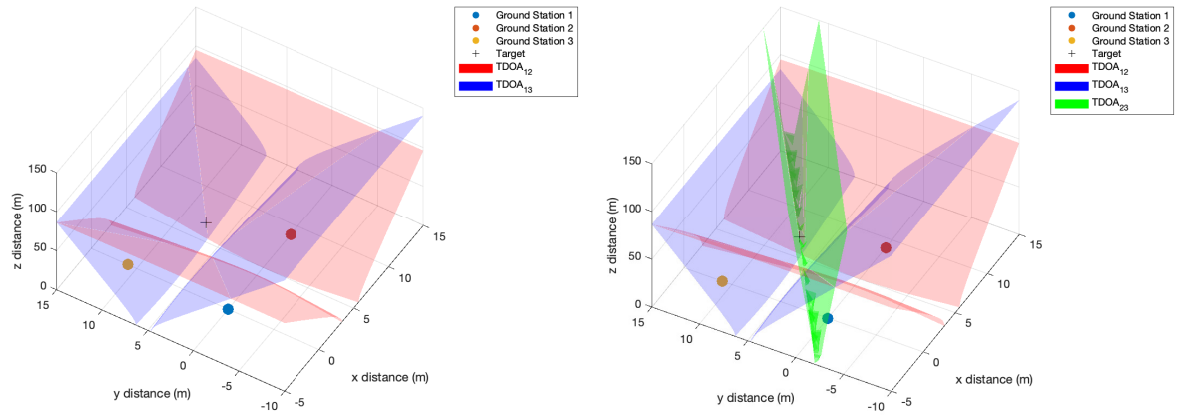

(a) Hyperbolic surfaces resulting of $TDOA_{12}$ and $TDOA_{13}$

(b) Hyperbolic surfaces resulting of $TDOA_{12}$, $TDOA_{13}$ and $TDOA_{23}$

Figure 3.6: In the left figure only $TDOA_{12}$ and $TDOA_{13}$ are plotted, here the intersection point of the two hyperbola creates four lines. In the right figure $TDOA_{23}$ is added. $TDOA_{23}$ hyperbola reduces the amount of solutions (i.e the lower left and upper right intersection lines are no longer a valid solution), but the location of the target is on one of the other remaining intersection lines.

When this validation algorithm will be tested on real WAM data provided by LVNL, it is expected that on average around 4 to 12 GSs will contribute to the TDOA measurements. Due to this, computing all TDOA possibilities requires too much computational power. On the other side, only computing three TDOA measurement is expected to not always be sufficient as the GDOP can result in a bad (or ambiguous) location estimate. Computing only all independent TDOA measurements provides more than sufficient information on the location of the target and acceptable computational cost, therefore the TDOA measurement $\mathbf{z}$ is defines as such;

$$\mathbf{z}_k = \begin{bmatrix} TOA_2 - TOA_1 \\ TOA_3 - TOA_1 \\ \vdots \\ TOA_N - TOA_1 \end{bmatrix} = \begin{bmatrix} TDOA_{21} \\ TDOA_{31} \\ \vdots \\ TDOA_{N1} \end{bmatrix} \tag{3.6}$$

By convention the GS that detects the ADS-B measurement first is assigned TOA_1, and so on till the N'th GS detects TOA_N.

3.2. Particle Filter

In this section background theory on state estimation and Particle Filtering (PF) is reviewed, and the PF design is presented. The background theory is discussed in subsection 3.2.1 by giving an introduction to Bayesian state estimation followed by several PF designs. The Sequential Importance Sampling (SIR) PF is introduced. The SIR filter is expected to perform poorly when there is high uncertainty in the measurement, as is the case with an ambiguous measurement. Therefore the Multiple Importance Filter (MIS) and the Mixture Multiple Importance Sampling (MISS) Filter are investigated as possible solutions.

3.2.1. Particle Filter Theory

Bayesian State Estimation

In many engineering problems estimating the state of a target or process is required. This is done by estimating the state described by the state vector $\mathbf{x}_k$.This state is estimated in two steps, firstly the prediction step where the function $\mathbf{f}_k(\mathbf{x}_{k-1}, \mathbf{w}_{k-1})$ (also called process model) is used to predict the previous state estimate $\mathbf{x}_{k-1}$ to $\mathbf{x}_k$. The second step is called the update step, here the prediction is updated based on evidence. Evidence is considered to be some measurement or observation that provides information on the state $\mathbf{x}_k$. Commonly not the exact state of the target is measured, a radar for instance only measures distance and azimuth, but the state vector of the target commonly contains the

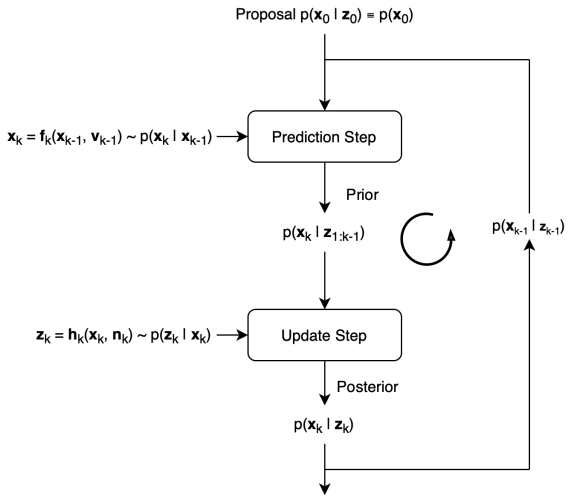

Figure 3.7: Bayesian filtering schematically shown. The filter is initialized by the proposal $p(\mathbf{x}_0)$. Using the process model f_k the state is propagated to the prior belief. This prior belief is updated using the measurement model h_k such that the posterior belief of the state of the target is obtained. For the next iteration this posterior density $p(\mathbf{x}_k|\mathbf{z}_k)$ is considered to be $p(\mathbf{x}_{k-1}|\mathbf{z}_{k-1})$ and the process is repeated.

location of a target described by Cartesian or Geodetic coordinates. Therefore a measurement model $\mathbf{z}_k = \mathbf{h}_k(\mathbf{x}_k, \mathbf{n}_k)$ is used linking the measurements to the state of the target. In this model $\mathbf{n}_k$ describes the uncertainty with respect to the observation. This entire process is repeated when a new measurement is available.

$$\begin{aligned} \mathbf{x}_k &= \mathbf{f}_k\left(\mathbf{x}_{k-1}, \mathbf{w}_{k-1}\right) \\ \mathbf{z}_k &= \mathbf{h}_k(\mathbf{x}_k, \mathbf{n}_k) \end{aligned} \tag{3.7}$$

State estimation is approached from a statistical viewpoint because it involves making predictions and dealing with noisy measurements. The pdf $p(\mathbf{x}_k \mid \mathbf{z}_{1:k})$ describing the state of the target is recursively estimated through prediction and update. In the prediction step the process model is used to propagate the previous posterior density $p(\mathbf{x}_{k-1} \mid \mathbf{z}_{1:k-1})$ to a prior density $p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})$ at the next time index. This is done using the Chapman-Kolmogorov equation shown in eq. 3.8. $p\left(\mathbf{x}_k \mid \mathbf{x}_{k-1}\right)$ is computed using the process model $\mathbf{f}_k$.

$$p\left(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}\right) = \int p\left(\mathbf{x}_k \mid \mathbf{x}_{k-1}\right) p\left(\mathbf{x}_{k-1} \mid \mathbf{z}_{1:k-1}\right) d\mathbf{x}_{k-1} \tag{3.8}$$

In the update step we update the pdf $p\left(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}\right)$ to $p\left(\mathbf{x}_k \mid \mathbf{z}_{1:k}\right)$ using Bayes theorem, and thus the state estimate is obtained. $p\left(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}\right)$ is considered the prior density and $p\left(\mathbf{x}_k \mid \mathbf{z}_{1:k}\right)$ the posterior density. The posterior can then be computed using eq. 3.9. The new measurement is incorporated in the posterior density by the pdf $p\left(\mathbf{z}_k \mid \mathbf{x}_k\right)$.

$$p\left(\mathbf{x}_k \mid \mathbf{z}_{1:k}\right) = \frac{p\left(\mathbf{z}_k \mid \mathbf{x}_k\right) p\left(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}\right)}{p\left(\mathbf{z}_k \mid \mathbf{z}_{1:k-1}\right)} \tag{3.9}$$

The entire process of prediction and update is repeated such that for every measurement available at time k a posterior density is obtained. Fig. 3.7 illustrates this principle. At Time $k = 0$ the filter is initialized by the proposal density $p(\mathbf{x}_0)$, which contains an initial guess of the state of the target. The entire process is repeated by considering that the posterior density at time $k-1$ is the prior density at time k.

Eq. 3.7 can be a linear or a nonlinear function of the state vector $\mathbf{x}_k$. If they are linear functions of the state vector the measurement and process model can be defined as seen in eq. 3.10. For this to hold also the noise model must be additive and Gaussian.

$$\begin{aligned} \mathbf{x}_k &= F_k \mathbf{x}_{k-1} + \mathbf{w}_{k-1} \\ \mathbf{z}_k &= H_k \mathbf{x}_k + \mathbf{n}_k \end{aligned} \tag{3.10}$$

For such models there exists an optimal solution. This solution is obtained using the Kalman Filter [31] and it relies on a linear prediction and measurement model, and additive Gaussian noise. TDOA localization has a nonlinear measurement equation as seen in eq. 3.3. Common solution to nonlinear measurement models is the use of the Extended Kalman Filter (EKF). The EKF is a sub-optimal filtering technique that linearizes the measurement model such that eq. 3.10 can be approximated.

Sequential Importance Resampling (SIR) Filter

The SIR filter was the first working Particle filter since its introduction as the Bootstrap filter by Gordon *et al*. [32]. Particle filtering is a filtering technique that estimates the state of a object of interest by recursive Bayes estimation using Monte-Carlo simulations. The fundamental idea behind the PF is that the posterior density describing the state of the target can be approximated using samples, with a corresponding weight. A particle represents a possible target state and the weight provides a measure of likelihood on this estimate. This set of samples and weight is denoted by $\left\{\mathbf{x}_k^i, \omega_k^i\right\}_{i=1}^{N_s}$, here $\mathbf{x}_k$ is the state vector represented by a set of particles, ω_k the assigned weight and the number of samples N_s. When the sum of the weights is equal to one, i.e. $\sum_i \omega_k^i = 1$, the posterior density can be approximated with eq. 3.11. δ is equal to Dirac's delta function.

$$p(\mathbf{x}_{0:k} \mid \mathbf{z}_{1:k}) \approx \sum_{i=1}^{N_s} \omega_k^i \delta(\mathbf{x}_{0:k} - \mathbf{x}_{0:k}^i) \tag{3.11}$$

At each iteration the weights w_k^i are updated using eq. 3.12. Eq. 3.12 is derived by expressing the posterior density $p\left(\mathbf{x}_{0:k} \mid \mathbf{z}_{1:k}\right)$ by using Bayes rule. A full derivation can be found in [33].

$$\omega_k^i \propto \omega_{k-1}^i \frac{p(\mathbf{z}_k \mid \mathbf{x}_k^i) p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)}{q(\mathbf{x}_k^i)} \tag{3.12}$$

The importance density $q(\mathbf{x}_k^i)$ in literature is often indicated with a q instead of a p. It contains the particles for which the weights are computed by eq. 3.12. In the SIR filter the importance density is set equal to the prediction density $p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)$.

$$\omega_k^i \propto \omega_{k-1}^i p(\mathbf{z}_k \mid \mathbf{x}_k^i) \tag{3.13}$$

In the prediction step particles are predicted to the next possible state of the target, therefore it occurs that particles propagate from a high likelihood region to a low likelihood region (i.e. a bad prediction).

At each iteration of the filter particles are lost to low likelihood regions and because of this there remain less particles in high likelihood regions which represent the state of the target. This phenomenon is called the *Degeneracy Problem*. Eventually almost all particles will have weights close to zero, except for one particle which will be closest to the true state of the target. [34] shows that at every iteration the variance of ω_k^i increases, essentially confirming the degeneracy problem. To solve this problem the SIR filter was proposed in [32] where a resampling step is added. Due to this resampling step the SIR filter was considered to be the first working PF.

The idea behind the resampling step is that particles with a near zero weight will be replaced (resampled) by particles with a high weight. After resampling the weights of the particles are set to $\omega_k^i = 1/N_s$. Several resampling implementations are possible such as residual or stratified sampling [35]. The common implementation is the multinomal resampling approach. With multinomal resampling particles are drawn from $\left\{\mathbf{x}_k^i, \omega_k^i\right\}_{i=1}^{N_s}$ based on the weights ω_k^i. This results in a new set of particles which are assigned equal weights.

Multiple Importance Sampling (MIS) Filter

In the SIR filter the importance density is computed by $p(\mathbf{x}_k \mid \mathbf{x}^i_{k-1})$. This calculation is widely used due to its simplicity and low computational requirements. If possible, one can also draw samples directly from the measurement $p(\mathbf{z}_k \mid \mathbf{x}^i_k)$ as is shown in eq. 3.14.

$$q(\mathbf{x}^i_k) = p(\mathbf{z}_k \mid \mathbf{x}^i_k) \tag{3.14}$$

Computing the weights is done by substitution eq. 3.14 into eq. 3.12. Then the observation density cancels out and the weights update equation is obtain in eq. 3.15. Compared to the SIR filter the observation density and the prediction density are swapped around. Here the measurement is used in the prediction step, and the prediction density is used in the update step.

$$\omega^i_k \propto \omega^i_{k-1} p(\mathbf{x}^i_k \mid \mathbf{x}^i_{k-1}) \tag{3.15}$$

For the problem at hand, namely TDOA localization is it possible to draw location samples from $p(\mathbf{z}_k \mid \mathbf{x}^i_k)$. All hyperbolas can thus be represented by samples and multiple hyperbolas can be used to compute the importance density, hence this implementation of the PF is called the Multiple Importance Sampling (MIS) Filter [36]. Each of the set of samples computed from a single TDOA measurement can be denoted as $q_p(x^i)$.

$$q(\mathbf{x}^i_k) = \prod_{p=1}^{M} q_p(x^i) \tag{3.16}$$

M is the number of densities that are multiplied. Every $q_p(x^i)$ (resulting from one TDOA measurement) can be assigned a mixture weight $\beta(x^i_k)$ [37].

$$\beta_p(x) = \frac{n_p q_p(x)}{\sum_{p=1}^{M} n_p q_p(x)} \tag{3.17}$$

This function can be used to model uncertainties that are present in the each $q_p(x^i)$. $\beta_p(x)$ can be assumed to be equal for all importance densities, this is called the balance heuristic. It is a reasonable assumption because the additive Gaussian noise used to model the measurement is equal for every TOA measurement, resulting in all observation densities to be equally accurate. β_p can be used to tune the proposal density if certain GSs have different noise parameters.

Mixture Multiple Importance Sampling (MMIS) Filter

One can consider to not exclusively sample from $p(\mathbf{z}_k \mid \mathbf{x}^i_k)$ or $p(\mathbf{x}_k \mid \mathbf{x}^i_{k-1})$, but a mixture of the two as is proposed by Kronander and Schön [37] and Zou *et al.* [38]. The importance density is then a combination of the prediction density and the observation density, scaled by a factor α to guarantee that the sum over the density equals one.

In the context of particles the number α defines the ratio of particles between the prediction density and observation density such that $N_s = N_{pred} + N_{meas}$, $N_{pred} = \alpha N_s$ and $N_{meas} = (1-\alpha)N_s$. $p(\mathbf{z}_k \mid \mathbf{x}^i_k)$ denotes the observation density and $p(\mathbf{x}_k \mid \mathbf{x}^i_{k-1})$ the prediction density.

$$q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{z}_k) = \alpha p(\mathbf{z}_k \mid \mathbf{x}^i_k) + (1-\alpha) p(\mathbf{x}_k \mid \mathbf{x}^i_{k-1}) \tag{3.18}$$

With this implementation the particles drawn from the observation density are assigned a weight based on their similarity compared to the prediction density, and vice versa.

3.2.2. Particle Filter Design

In the SIR filter all particles are propagated from the posterior density by prediction density. If the posterior density has a high uncertainty, the prediction density will result in a highly uncertain prediction. In the update step this introduces the degeneracy problem. In anticipation of this problem the MISS filter can be implemented if the degeneracy problem is severe. If this is not the case, the SIR

filter is the final PF design.

The MISS-PF samples from both the observation density and the prediction density. This is expected to solve the degeneracy problem as the samples that are drawn from the observation density have a higher likelihood than samples drawn from a prediction density with high uncertainty. It is not necessary the case that samples from the observation density have a higher likelihood than the prediction density, but when the prediction density has a significantly higher uncertainty (as can be the case) a more robust filter is obtained when samples are drawn from the observation density.

The state vector of the particle filter will contain position and velocity as shown in eq. 3.19. With this state vector all information is available to validate position and velocity. In this state vector the velocity is decomposed into v_x, v_y, v_z. To compute the true velocity the target the two-norm of the velocity components is taken.

$$\mathbf{x}_k = \begin{bmatrix} l_x \\ l_y \\ l_z \\ v_x \\ v_y \\ v_z \end{bmatrix}_k \tag{3.19}$$

In the next subsection the design of the implemented SIR filter is given, and in section 3.2.4 the MISS is presented as a possible solution if the performance of the SIR filter is not sufficient.

3.2.3. Sequential Importance Resampling Filter

At the initialization of the SIR filter samples need to be drawn from a proposal density. The proposal density used at the first iteration of the filter. A proper choice of the proposal density can greatly improve the performance of the filter. It is expected that the degeneracy problem can greatly impact the performance of the filter due to the presence of ambiguous measurements.

The SIR filter has it similarities to Bayesian filtering as seen in Fig. 3.7. In the SIR filter the probability density functions are approximated using samples. This is denoted with the superscript i as seen in Fig. 3.8. The update step is achieved by computing the weights of all particles for which the measurement is used, and multinomal resampling.

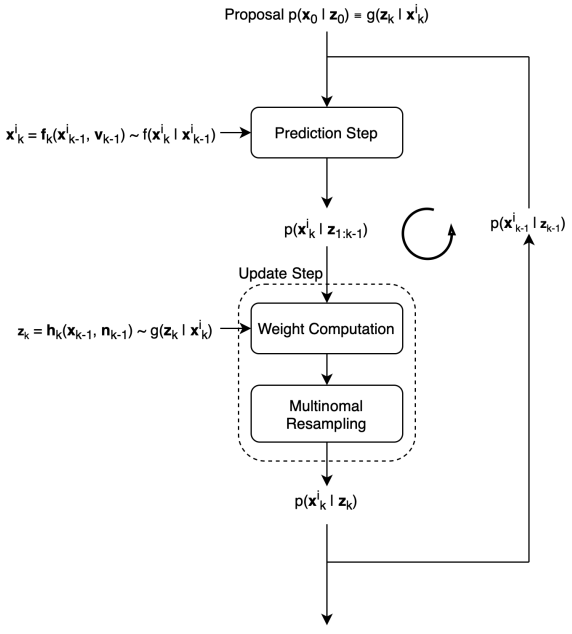

Figure 3.8: Schematic representation of the SIR filter algorithm.

Proposal Density

The proposal density in the SIR filter can be any density in theory. A bad choice can result in the filter not converging to the true target state or have an impact on the severity of the degeneracy problem. In this SIR filter implementation the observation density $p(\mathbf{z}_0|\mathbf{x}_0^i)$ is used as proposal density. A method is proposed to sample from this density that takes into account the noise statistics.

The measurement eq. 3.20 depends on the location of the two corresponding ground stations $\mathbf{s_1}$, $\mathbf{s_2}$, the z_k^j measurement and the location of the target. j equals the j'th element in the measurement vector $\mathbf{z}_k$. To obtain samples from eq. 3.20 Matlab's Symbolic Toolbox is used to rewrite eq. 3.20 such that $l_z = f(l_x, l_y, \mathbf{s}_1, \mathbf{s}_2, z_k^j)$. Using this approach we can input l_x and l_y values and obtain the corresponding l_z value. A sample from the measurement equation is then $\mathbf{l}^i = [l_x^i, l_y^i, l_z^i]^T$. The code that generates the Matlab function is shown in appendix A.

$$z_k^j = \frac{1}{c}\left(\sqrt{(s_{jx} - l_x)^2 + (s_{jy} - l_y)^2 + (s_{jz} - l_z)^2} - \sqrt{(s_{1x} - l_x)^2 + (s_{1y} - l_y)^2 + (s_{1z} - l_z)^2}\right) \tag{3.20}$$

The vectors $\mathbf{l}_x = [l_{x-min}, \ldots, l_{x-max}]$ and $\mathbf{l}_y = [l_{y-min}, \ldots, l_{y-max}]$ are used to compute a meshgrid. This grid contains all coordinates where the observation density will be evaluated. If a target that is spoofing is airborne a large grid size allows for spoofing detection, and localization of the target that is performing the spoofing. A small grid size allows for less computational intense ADS-B validation, but in the event of an airborne spoofer it can be possible the grid size is not large enough to find the location of the spoofing aircraft.

The TDOA measurements contains additive noise represented by n_k^j in eq. 3.21.

$$z_k^j = \frac{||\mathbf{s}_j - \mathbf{l}|| - ||\mathbf{s}_1 - \mathbf{l}||}{c} + n_k^j \quad j = 2, \ldots, N \tag{3.21}$$

The noise measured for each TOA is assumed to be Gaussian distributed $\mathcal{N}(0, \sigma_{toa}^2)$. Because the TDOA measurement is the difference between two normally distributed random variables and the TOA measurements are uncorrelated the noise statistic of the TDOA measurement is equal to,

$$n_k^j \approx \mathcal{N}(0, 2\sigma_{toa}^2) = \mathcal{N}(0, \sigma_{tdoa}^2) \tag{3.22}$$

Realizations of these noise statistics are used to compute the observation density in an efficient manner. In the proposed sampling method, If there are N_s samples, realizations are drawn from the observation density, there are N_s samples drawn from $\mathcal{N}(\mathbf{z}_k, \sigma_{TDOA}^2)$. Each coordinate in the mesh grid represented by an l_x and l_y coordinate is paired with a TDOA sample drawn from $\mathcal{N}(\mathbf{z}_k, \sigma_{TDOA}^2)$. All pairs are then used as input to eq. 3.23.

$$l_z = f(l_x, l_y, \mathbf{s}_1, \mathbf{s}_2, z_k^j) \tag{3.23}$$

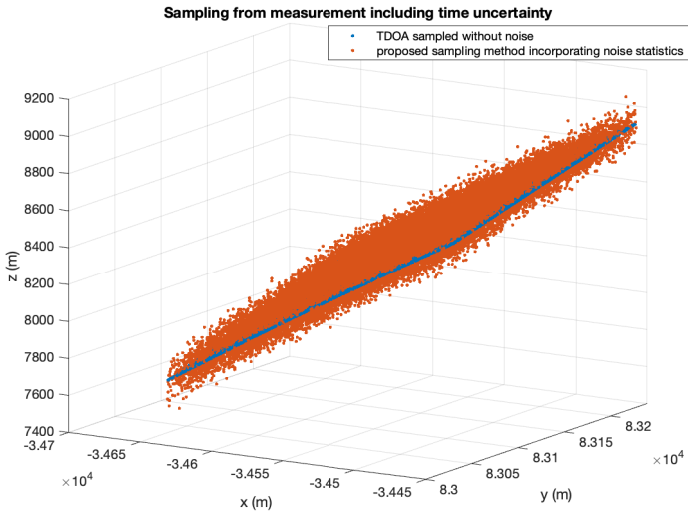

Figure 3.9: Shown in blue are samples drawn from a TDOA measurement with the use of the noise statistics. The orange samples are drawn from the measurement using the proposed method where the noisy realizations of the TDOA measurement are used.

Fig. 3.9 shows this sampling procedure applied to a TDOA measurement. To clearly observe the effect of the proposed sampling method samples are drawn from a TDOA measurement using no noise statistics, and samples are draw including the noise statistics. Samples in blue are drawn without the noisy TDOA realizations, and in orange the noisy realizations are incorporated. The severity of the noise changes w.r.t the location on the hyperbola. Near the vertex point of the hyperbola the noise is small, and as we travel along the asymptotes of the hyperbola the noise increases. Thus the sampling procedure efficiently and accurately computes the effect of a noisy TOA measurement on the observation density.

From the measurement only samples are drawn that provide information on the location of the target. As the state vector also contains velocity, samples describing the initial velocity distribution also need to be drawn. These samples are obtained from the 1st received ADS-B message, and are drawn from $\mathcal{N}(v_{adsb}, \sigma_a^2)$.

Prediction Step

At the first iteration of the filter the proposal density is predicted to a prior belief of the state. This is done using the prediction density $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}^i)$. After the first iteration of the filter, prediction is applied to the posterior density $p(\mathbf{x}_{k-1}^i \mid \mathbf{z}_{k-1})$ of the previous filter iteration.

In general tracking systems make use of several models that describe certain behavior of a target. Three models are used to describe the state of a target, namely Constant Acceleration (CA), Coordinated Turn (CT) and Constant Velocity (CV). If a target is executing a turn, the CT model is able to predict the state of the target with a higher accuracy than the other two models. Similarly if a target is accelerating, the CA is able to predict the state of the target with better accuracy. The filter chooses the model that has the highest likelihood based on the measurement, and as a result the prediction quality improves.

The use of multiple models in a PF is commonly achieved using parallelization, but other methods have been proposed [39]. The problem with running a PF in parallel is the heavy computational load and added complexity. Generally two types of parallel PFs are presented in literature, namely the non-interacting PF and the interacting PF [40]. The non-interacting PF runs several PFs where all PFs have their own process model and own resampling scheme. In an interacting PF the individual PFs communicate with one another, which in general allows for more efficient estimation [40] at the cost of increased computational complexity.

Considering the added complexity, and that the expected traffic in WAM systems at LVNL is not

expected to execute fast maneuvers with rapid direction changes, only the CV model is implemented.

The prediction density $p(\mathbf{x}_k \mid \mathbf{x}_{k-1}^i)$ is thus computed using a CV model as seen in eq. 3.24. In this CV model the process noise is Gaussian distributed with a variance σ_w^2, and the acceleration is modeled as additive noise. The noise is considered to be correlated in each axis, meaning that in the x axis the velocity and position noise is equal at each time index. Therefore w is decomposed in x, y and z.

$$\begin{bmatrix} l_x \\ l_y \\ l_z \\ v_x \\ v_y \\ v_z \end{bmatrix}_k = \begin{bmatrix} 1 & 0 & 0 & T & 0 & 0 \\ 0 & 1 & 0 & 0 & T & 0 \\ 0 & 0 & 1 & 0 & 0 & T \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} l_x \\ l_y \\ l_z \\ v_x \\ v_y \\ v_z \end{bmatrix}_{k-1} + \begin{bmatrix} \frac{T^2}{2} & 0 & 0 \\ 0 & \frac{T^2}{2} & 0 \\ 0 & 0 & \frac{T^2}{2} \\ T & 0 & 0 \\ 0 & T & 0 \\ 0 & 0 & T \end{bmatrix} \begin{bmatrix} w_x \\ w_y \\ w_z \end{bmatrix}_{k-1} \qquad w \sim \mathcal{N}(0, \sigma_w^2) \tag{3.24}$$

Update Step

The update step in the SIR filter consist of two separate steps. Firstly all samples are assigned a weight based on the measurement $\mathbf{z}_k$. The function that assigns a weight to all particles is called the likelihood function. After each particle is assigned a weight the resampling procedure is applied.

Weights Computation

The likelihood function is used to evaluate $\omega_k^i \propto p(\mathbf{z}_k \mid \mathbf{x}_k^i)$. For each particle in the importance density the likelihood function is evaluated and a likelihood λ_k^i is obtained. This likelihood λ_k^i is equal to the weight ω_k^i.

The TDOA measurements are computed from the TOA measurement by the use of a matrix multiplication. This matrix denoted with C is multiplied with the TOA measurement vector such that $\mathbf{z}_k$ is obtained. If we measure a signal at N ground stations $\mathbf{z}_k$will be a $[(N-1)\times 1]$ vector. This results in C being $[(N-1)\times N]$.

$$\mathbf{z}_k = \underbrace{\begin{bmatrix} -1 & 1 & 0 & 0 \\ \vdots & 0 & \ddots & 0 \\ -1 & 0 & 0 & 1 \end{bmatrix}}_{\text{C}} \underbrace{\begin{bmatrix} TOA_1 \\ TOA_2 \\ \vdots \\ TOA_N \end{bmatrix}}_{\mathbf{z}_{toa}} \tag{3.25}$$

Each measured z_k^j has its corresponding measurement equation shown in eq. 3.26.

$$z_k^j = \frac{||\mathbf{s}_j - \mathbf{l}|| - ||\mathbf{s}_1 - \mathbf{l}||}{c} + n_k^j \quad j = 2, \ldots, N \tag{3.26}$$

The TDOA measurement is assumed to contain additive Gaussian noise (eq. 3.22). Thus the likelihood function is equal to eq. 3.27.

$$\omega_k^i = \frac{1}{\sqrt{|2\pi Q|}} \exp\left(-\frac{1}{2}(\mathbf{z}_k - \mathbf{z}_k^i)^T Q^{-1} (\mathbf{z}_k - \mathbf{z}_k^i)\right) \tag{3.27}$$

For each particle $\mathbf{x}_k^i$ a corresponding $\mathbf{z}_k^i$ is computed. This $\mathbf{z}_k^i$ is a hypothetical $\mathbf{z}_k$ measurement computed based on the location of the particle and the contributing GSs to the true measurement $\mathbf{z}_k$. $\mathbf{z}_k^i$ is compared this against the true measurement $\mathbf{z}_k$ according to the likelihood function. $\mathbf{z}_k^i$ is computed using eq. 3.28. When a particle a similar location as the target, the hypothetical measurement $\mathbf{z}_k^i$ is similar to the true measurement $\mathbf{z}_k$, resulting in a high weight.

$$\mathbf{z}_k^i = \begin{bmatrix} \frac{||\mathbf{s}_2 - \mathbf{l}_k^i|| - ||\mathbf{s}_1 - \mathbf{l}_k^i||}{c} \\ \frac{||\mathbf{s}_3 - \mathbf{l}_k^i|| - ||\mathbf{s}_1 - \mathbf{l}_k^i||}{c} \\ \vdots \\ \frac{||\mathbf{s}_N - \mathbf{l}_k^i|| - ||\mathbf{s}_1 - \mathbf{l}_k^i||}{c} \end{bmatrix} \tag{3.28}$$

The covariance matrix Q is equal to $E(\mathbf{n}_k \mathbf{n}_k)$, where $\mathbf{n}_k$ is the noise corresponding to the measurement $\mathbf{z}_k$, the noise is obtained through $\mathbf{n}_k = C\mathbf{n}_{toa}$. The covariance matrix is then obtained by eq. 3.29. $\Sigma = \text{diag}(\sigma_1^2, \ldots, \sigma_N^2)$ is a diagonal matrix containing the variance of each GSs.

$$Q = E(\mathbf{n}_k \mathbf{n}_k^T) = C\Sigma C^T \tag{3.29}$$

When a faulty measurement is received, it is very likely that all particles have a weight of zero. In such scenarios resampling fails. To allow the filter to continue to track the target when faultu measurements are received, ϵ-robustness is applied. In epsilon-robustness an arbitrary small uniform distribution is added to the Gaussian distributed likelihood function according to eq. 3.30

$$w_k^i = \epsilon \frac{1}{\sqrt{|2\pi Q|}} \exp\left(-\frac{1}{2}(\mathbf{z}_k - \mathbf{z}_k^i)^T Q^{-1} (\mathbf{z}_k - \mathbf{z}_k^i)\right) + (1 - \epsilon)\frac{1}{P} \tag{3.30}$$

In the equation above $\epsilon \in [0, 1]$ to allow for the integral over the likelihood function to be equal to one. If a particle is evaluated and assigned a value of zero by the Gaussian part of the likelihood function, the uniform part of the likelihood assures that the final weight of the particle is never zero. Thus when a faulty measurement or spoofed ADS-B message is received all particles are assigned an equally low weight according to $\frac{1}{P}$. The value for P can be set such that $\frac{1}{P}$ is the smallest possible number that the hardware running the filter is possible to compute. For a 64-bit machine this is around 10^{-308}. After resampling all particles are redrawn (due to equal weights) and the posterior density can be used for ADS-B validation.

Multinomal Resampling

Multinomal resampling is the most common resampling method used in particle filtering. In multinomal resampling, particles are drawn based on the weights assigned by the likelihood function. If a particle has a high weight it will be drawn more likely, if is has a low weight it will be drawn less likely. The newly drawn particles are the particles that make up the posterior distribution $p(\mathbf{x}_k^i \mid \mathbf{z}_k)$. Finally, all the weights are considered to be equal. Consequence of the resampling procedure is that particles with a low weight will be lost, and particles with a high weight will be duplicated.

The resampling procedure is as follows, a CDF is calculated based on the particles' weights $\{\omega_k^i\}_{i=0}^{N_s}$. The weights of the particles are normalized to one, thus the range of the CDF is from zero to one. This CDF is visualized as the outer ring in the Fig. 3.10 where the size of the basket represents the weight of a particle and the edges are defined by the CDF. Then, a random uniformly distributed sample is drawn $u \sim U(0, 1]$ if a sample from u falls into the basket of $\{\omega_k^i\}$ that particle is drawn. Particles with a larger weight thus have a higher probability of getting drawn, and particles with a low weight will be drawn less. This procedure is repeated N_s times such that a resampled set of particles is obtained. Multinomal resampling introduces duplicate particles which leads to loss of diversion among the particles. This loss of diversion is called sample impoverishment.

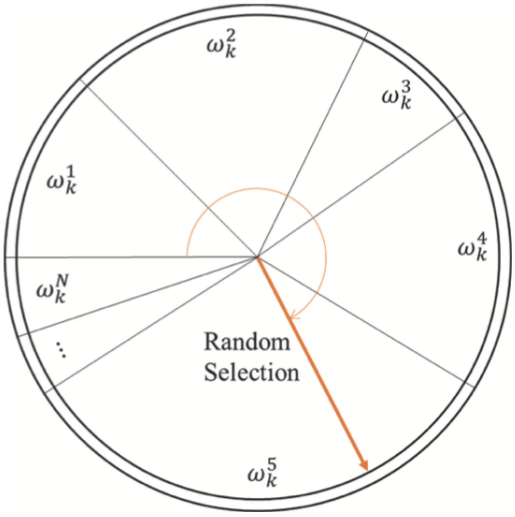

Figure 3.10: The weight of a particle determines the size of the outer ring on the size. Randomly uniform distributed samples are drawn, if this value falls into the basket w_k^i, that particle gets resampled. This procedure is repeated N times till all particles are resampled. Illustration from [3].

SIR Filter Algorithm

Algorithm 1 SIR-Particle Filter Algortihm

Input: $\mathbf{z}_k$
Output: $p(\mathbf{x}_k^i \mid \mathbf{z}_k)$

1: Sample initial particle distribution $p(\mathbf{z}_k \mid \mathbf{x}_k^i)$ for $i = 1, \ldots, N_s$
2: Set $w_k^i = 1/N_s$ for $i = 1, ..., N_s$
3: **for** k = 1,...,K **do**
4: Sample $x_k^i \sim p(\mathbf{x}_k \mid \mathbf{x}_{k-1}^i)$ for $i = 1, \ldots, N_s$
5: Compute $w_k^i = p(\mathbf{z}_k \mid \mathbf{x}_k^i)$ for $i = 1, \ldots, N_s$
6: Resample using multinomal resampling
7: Normalize $w_k^i = \frac{w_k^i}{\sum w_k^i}$
8: **end for**

3.2.4. Mixture Multiple Importance Sampling Filter

The MISS Particle Filter is expected to perform better than the SIR filter due to the degeneracy in problem in the SIR filter. In this section a possible implementation of the MISS filter is given, such that if needed it can be applied to mitigate the problem. Main difference between SIR and MISS is the importance density and the effect that it has on the weight update equation.

Importance density

In the MISS filter the importance density is equal to eq. 3.31. It draws samples from the observation density and the prediction density. Both densities can be sampled from using the method proposed in section 3.2.3. The mixture balance α determines the ratio of how many samples are drawn from both densities.

$$q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{z}_k) = \alpha p(\mathbf{z}_k \mid \mathbf{x}_k^i) + (1-\alpha) p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i) \tag{3.31}$$

$$w_k^i \propto \frac{p(\mathbf{z}_k \mid \mathbf{x}_k^i) p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)}{\alpha p(\mathbf{z}_k \mid \mathbf{x}_k^i) + (1-\alpha) p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)} \tag{3.32}$$

$$\begin{aligned} \widetilde{w}_k^i &= p(\mathbf{z}_k \mid \mathbf{x}_k^i) \text{ for } i = 1, \ldots, N_{pred} \\ \widetilde{w}_k^i &= p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i) \text{ for } i = N_{pred}+1, \ldots, N_s \end{aligned} \tag{3.33}$$

Computing the weights of the particles using the MISS filter is different than the SIR filter. In the MISS filter weights are determined by eq. 3.32. This equation implies that the samples that are drawn from the observation density need to be assigned weights based on the prediction density, and the particles drawn from the prediction density need to be assigned weights based on the observation density. By assuming the mixture weights to be equal, the particle weights can be computed by using eq. 3.33 [37]. Equal mixture weights has shown to be a valid assumption in section 3.2.1.

Sampling Velocity in Observation Density

The state vector contains positions and velocity. When samples are drawn from the measurement as illustrated in section 3.2.3 only a location is sampled. Particles must be assigned a velocity in order to be used in the filter. From the measurement itself it is not possible to obtain a velocity. In the proposed MISS filter implementation the velocity estimate from the previous posterior density is used to obtain velocity samples. The velocity particles from the posterior density are approximated by a Gaussian distribution.

$$\{v_x^i\}_{i=0}^{N_s} \sim \mathcal{N}(\mu_{vx}, \sigma_x^2) \tag{3.34a}$$

$$\{v_y^i\}_{i=0}^{N_s} \sim \mathcal{N}(\mu_{vy}, \sigma_y^2) \tag{3.34b}$$

$$\{v_z^i\}_{i=0}^{N_s} \sim \mathcal{N}(\mu_{vz}, \sigma_z^2) \tag{3.34c}$$

From these approximated Gaussian distributions samples are drawn and added to the samples drawn from eq. 3.23. Using this procedure particles in the observation density contain velocity values which accurately describe the state of the target. Finally the position measurements and the velocity estimates are combined to obtain samples from the observation density $p(\mathbf{z}_k \mid \mathbf{x}_k^i)$ as shown in eq. 3.35.

$$\mathbf{x}_k^i = \begin{bmatrix} l_x^i \\ l_y^i \\ l_z^i \\ v_x^i \\ v_y^i \\ v_z^i \end{bmatrix}_k \tag{3.35}$$

Mixture Balance

The prediction density and observation density are added according to eq. 3.31. The factor α determines the mixture balance. When a TDOA measurement contains four or more ground stations the resulting shape of the estimate is a point in space. When the measurement originates from three GSs the resulting solution has the shape of a curved line created by the intersection on two hyperbolas. For a single TDOA measurement the solution is a hyperbola. The number of GSs thus has a big impact on the estimate the filter, therefore the value of α is set to be a function of the number of GSs present in the measurement $\mathbf{z}_k$.

$$\alpha = f(N_{GS}) \tag{3.36}$$

$$\alpha = \begin{cases} \alpha_{GS \geqslant 4} & \text{if } N_{GS} \geqslant 4 \\ \alpha_{GS < 4} & \text{if } N_{GS} < 4 \end{cases} \tag{3.37}$$

Eq. 3.36 makes the distinction between $N_{GS} \geqslant 4$ and $N_{GS} < 4$. When the measurement $\mathbf{z}_k$ contains four or more GSs the posterior density is non-ambiguous, meaning high certainty in the state estimate. In the following prediction step many particles will be located in high likelihood regions due to the peaked posterior density. In such scenarios drawing samples from the prediction density is more efficient and as accurate as drawing from the observation density. Therefore if four or more GSs contribute to the measurement a low value for α is preferable. Still, it can be helpful to draw samples from the observation density to increase the robustness of the filter.

There are two situations where a high value of α is preferable,

- **Low GS coverage**: If a target is located in an area where it is only detected by two or three GSs, the filter only receives ambiguous position measurement.Which implies a posterior density with high uncertainty, consequently the prediction density is also contains high uncertainty and does not provide an accurate prediction. Then it is preferable to have a high α.
- **Filter startup phase**: When the filter is initialized on a target a prior belief on the location of the target is needed. In such a situation a high value of α is preferable to quickly localize the target. After several iterations the filter the value for α can be decreased if the filter has accurately located the targets' state. This is essentially what the importance density does.

MISS Filter Algorithm

Algorithm 2 MMIS-Particle Filter Algortihm

Input: $\mathbf{z}_k$
Output: $p(\mathbf{x}_k^i \mid \mathbf{z}_k)$

Sample initial particle distribution $p(\mathbf{z}_k \mid \mathbf{x}_k^i)$ for $i = 1, ..., N_s$
Set $w_k^i = 1/N_s$ for $i = 1, ..., N_s$
for k = 1,...,K **do**
  Sample $x_k^i \sim (1-\alpha)p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)$ for $i = 1, ..., N_{pred}$
  Sample $x_k^i \sim \alpha p(\mathbf{z}_k \mid \mathbf{x}_k^i)$ for $i = N_{pred}+1, ..., N_s$
  Set $q(\mathbf{x}_k \mid \mathbf{x}_{k-1}, \mathbf{z}_k) = \alpha p(\mathbf{z}_k \mid \mathbf{x}_k^i) + (1-\alpha) f(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)$
  Compute $\widetilde{w}_k^i = p(\mathbf{z}_k \mid \mathbf{x}_k^i)$ for $i = 1, \ldots, N_{pred}$
  Compute $\widetilde{w}_k^i = p(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i)$ for $i = N_{pred}+1, \ldots, N_s$
  Resample using multinomal resampling
  Normalize $w_k^i = \frac{w_k^i}{\sum w_k^i}$
end for

3.3. Hypothesis Testing

In the first step of the algorithm the posterior density $p(\mathbf{x}_k^i \mid \mathbf{z}_k)$ is computed. In the second step the posterior density is compared with the ADS-B message, and decided if the ADS-B message is real or faked. To do this a hypothesis test is defined. A statistical hypothesis test is a method used to make inferences or decisions based on data. This is achieved by formulating a null hypothesis $\mathcal{H}_0$ and a alternative hypothesis $\mathcal{H}_1$. The null hypothesis is the default assumption that there is no significant difference between variables. The alternative hypothesis is the claim or statement that contradicts the null hypothesis. For spoofing detection the hypothesis are defined as such,

$\mathcal{H}_0$: No spoofing
$\mathcal{H}_1$: Spoofing

In the null hypothesis $\mathcal{H}_0$ no spoofing occurs, thus this is the default assumption. In the alternative hypothesis $\mathcal{H}_1$ *does* spoofing occurs. The goal is to detect spoofing, therefore $p(\text{choose } \mathcal{H}_1|\mathcal{H}_1)$ is defined as P_d. In total four scenarios arise in this binary hypothesis test. Two correct decisions can be made, namely concluding the message is real when it is real, and concluding the message is fake when it is fake. The hypothesis test can make two errors, the first error is concluding the message is real when in fact is it faked. This is considered to be a missed detection and has probability P_m. The second error it can make is a concluding a message is real when it is faked, this probability is denoted by P_{fa}

$$p(\text{choose } \mathcal{H}_1|\mathcal{H}_1) = P_d \tag{3.38a}$$

$$p(\text{choose } \mathcal{H}_0|\mathcal{H}_1) = P_m \tag{3.38b}$$

$$p(\text{choose } \mathcal{H}_1|\mathcal{H}_0) = P_{fa} \tag{3.38c}$$

$$p(\text{choose } \mathcal{H}_0|\mathcal{H}_0) = P_{rej} \tag{3.38d}$$

To determine if either one of the hypotheses occurs a Likelihood Ratio Test (LRT) is applied. In a likelihood ratio test two likelihoods are computed, namely $p\left(\mathbf{x};\mathcal{H}_0\right)$ and $p\left(\mathbf{x};\mathcal{H}_1\right)$. This ratio is compared against a threshold τ, if the LRT is larger than the threshold $\mathcal{H}_0$ is decided, and if the LRT is smaller than the threshold $\mathcal{H}_1$ is decided. Several methods to compute τ can be considered where each method makes some assumption. Choosing between which method to apply is similar to choosing which assumption seems most valid.

$$L(\mathbf{z}^{ADSB}) = \frac{p\left(\mathbf{z}^{ADSB};\mathcal{H}_0\right)}{p\left(\mathbf{z}^{ADSB};\mathcal{H}_1\right)} \lessgtr \tau \tag{3.39}$$

In the hypothesis test for the ADS-B validation algorithm three different hypotheses test are investigated. This stems from the following, the most common hypothesis test one can implement is the Bayesian hypothesis test. The Bayesian test achieves a *minimum probability of error* by use of the prior probability of both hypothesis occurring. This probability denoted with π_0 is often not known or difficult to get an estimate from. Other hypothesis tests rely on a Gaussian assumption of the posterior density $p(\mathbf{x}_k|\mathbf{z}_k)$ or that some other analytically known density fits the data. Most importantly each of these hypothesis tests makes a different assumption. The Bayesian hypothesis test assumes the prior probabilities and associated costs on both hypothesis to be known. Neyman-Pearson and Minimax require the probability distributions of equation 3.38 to be known. Data is often assumed to have some distribution and this assumption is not always accurate. Therefore it is difficult to argue which assumption generally performs the best because it largely depends on the measured data.

Three hypothesis tests are investigated. Mainly to conclude if one performs better, and to provide LVNL with options and insight into approaches that are possible to determine if an ADS-B message is spoofed or not.

- Minimum Bayes Risk: Threshold is set such that the minimum probability of error is obtained
- Neyman-Pearson: Threshold is set such that the tests maximizes P_d for a desired P_{fa}
- Minimax: Threshold is set such that the maximum Bayes Risk is minimized

3.3.1. Hypothesis Definition

In this subsection the two likelihoods $p\left(\mathbf{x};\mathcal{H}_0\right)$ and $p\left(\mathbf{x};\mathcal{H}_1\right)$ are defined. To define the hypothesis test the coordinates and units of the ADS-B message and the state estimate by the SIR filter must be the same. In the next paragraph the chosen coordinate system and units is presented.

Coordinate system & Units

In an ADS-B equipped aircraft the location is reported in latitude, longitude and Flight Level[1] (FL). The velocity reported by the ADS-B message is reported in knots. To combine the measurements from different local coordinate systems the ECEF coordinate system is used. Earth Centered Earth fixed (ECEF) is a Cartesian coordinate system where the zero is at the center of the earth and the x axis extends along the prime meridian and the z axis through the North-Pole. Measurements made in local coordinate systems are combined by transforming them to the ECEF coordinate system. In this global coordinate system tracking and filtering is done. From a human perspective the ECEF coordinate system is difficult to interpret, therefore this coordinate system is converted to a coordinate system of choice, such as geodetic coordinates (latitude, longitude, height) or the the East, North, Up (ENU) system which is easier to interpret. ENU is a Cartesian coordinate system similar to ECEF but with a different zero point and the unit vectors are chosen such that the z vector represents height.

Similar to ARTAS, the validation algorithm will convert the ADS-B messages and MLAT measurement to ECEF where the filtering operations will be performed. This approach allows for expendability and is illustrated by Fig. 3.11.

[1] 1 FL = 100 feet

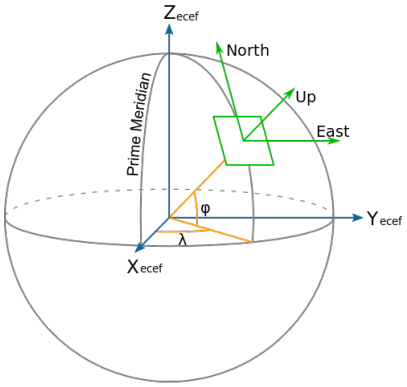

Figure 3.11: Overview of relevant coordinate systems. Illustration from[4].

To compute the ECEF coordinates from geodetic coordinates equation 3.40a is used. ϕ is latitude and λ is longitude, the parameters a, e^2 are used to determine the shape of the earth model that is used. These values are set to be equal with the operational ARTAS settings used at LVNL and are equal to $a = 6378137$ and $e^2 = f(2-f)$ with $f = \frac{1}{298.257223563}$ and correspond to the WGS-84 earth model.

The PF computes the velocity estimate of the target under test in m/s. ADS-B reports ground speed in knots and climb rate (γ) in ft/min. Both quantities are converted to m/s and the true velocity is obtained by equation 3.40b.

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \frac{a}{\sqrt{1 - e^2 \sin^2 \phi}} \begin{pmatrix} \cos\phi \cos\lambda \\ \cos\phi \sin\lambda \\ (1-e^2)\sin\phi \end{pmatrix} + H \begin{pmatrix} \cos\phi \cos\lambda \\ \cos\phi \sin\lambda \\ \sin\phi \end{pmatrix} \tag{3.40a}$$

$$v = \sqrt{(0.514 * v^{ADSB})^2 + (0.00508 * \gamma)^2} \tag{3.40b}$$

After this coordinate change is performed the ADS-B measurement is denoted with,

$$\mathbf{z}_k^{ADSB} = \begin{bmatrix} x \\ y \\ z \\ v \end{bmatrix}_k \tag{3.41}$$

To clearly indicate the difference between the location estimate of the PF and the ADS-B measurement the output of the PF is from now on denoted by $p(\mathbf{x}_k^i \mid \mathbf{z}_k^{MLAT})$

The velocity estimate of the PF is decomposed into v_x,v_y and v_z. These three components of each particle are combined to obtain the true velocity. At the output of the filter this operation is performed and the particles that are used for validation are defines as,

$$\mathbf{x}_{val}^i = \begin{bmatrix} x \\ y \\ z \\ v \end{bmatrix}^i \tag{3.42}$$

To simplify the notation the subscript k is dropped and interchanged with val to indicate that the vector computed by the PF is different from the vector that is used for validation. The difference between the two is that the velocity that is computed in the filter is decomposed in v_x, v_y and v_z, and the velocity in $\mathbf{x}_{val}$ contains the true velocity of the target expressed by v.

Spoofing Hypothesis $\mathcal{H}_0$
The probability distribution of $\mathcal{H}_0$ must represent the likelihood that the ADS-B message is not spoofed. This likelihood is computed by eq. 3.43. The aircraft determines the uncertainty associated with its

location and velocity report and transmits this inside the ADS-B message. It is assumed that the location and velocity are Gaussian distributed with a known covariance matrix R This reported variance is used in the hypothesis test to determine the covariance matrix R. A spoofer is likely to always set the highest possible uncertainty to remain possibly undetected for as long as possible. With the highest possible variance the 95% confidence interval is equal to 20 NM horizontally. This gives a spoofer a significant area to remain undetected, to account for this a upper limit is set to the uncertainty is set.

$$p\left(\mathbf{z}^{ADSB}|\mathcal{H}_0\right) = \int p(\mathbf{z}^{ADSB}|\mathbf{x}_{val};\mathcal{H}_0)p(\mathbf{x}_{val}|\mathbf{z}^{MLAT})d\mathbf{x}_{val} \tag{3.43}$$

Fig. 3.12 illustrates eq. 3.43. The pdf $p(\mathbf{z}^{ADSB}|\mathbf{x}_{val};\mathcal{H}_0)$ is Gaussian distributed but $p(\mathbf{x}_{val}|\mathbf{z}^{MLAT})$ can in principle obtain any distribution, but in the figure it is depicted as a Gaussian distribution. When the pdfs are similar they overlap resulting in a high likelihood.

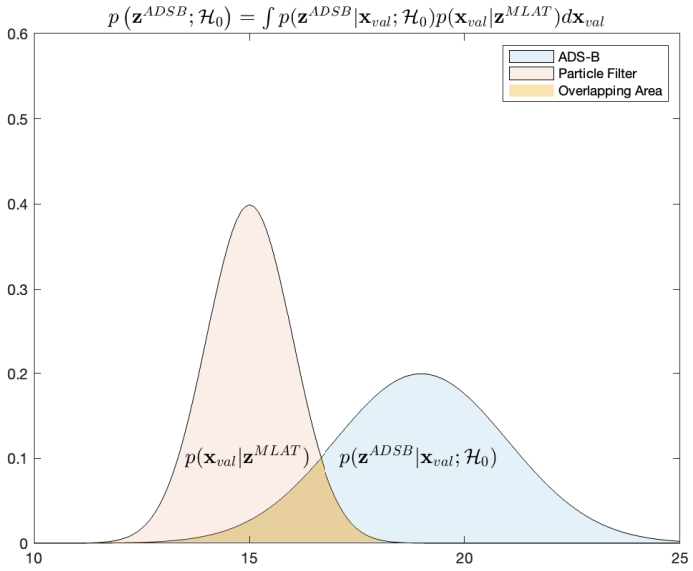

Figure 3.12: 1-dimensional Illustration of eq. 3.43. The two graphs represent the distribution of the PF and the ADS-B message. The overlapping area is a measure of the likelihood between the two distributions.

$p(\mathbf{x}_{val}|\mathbf{z}^{MLAT})$ is approximated by the particle filter with a particle cloud. Therefore the integral cannot be computed analytically, but approximated by the samples that represent $p(\mathbf{x}^i_{val}|\mathbf{z}^{MLAT})$. Each sample of $p(\mathbf{x}^i_{val}|\mathbf{z}^{MLAT})$ has equal density, therefore equation 3.43 can be approximated by eq. 3.44.

$$p\left(\mathbf{z}^{ADSB};\mathcal{H}_0\right) \approx \sum_{i=0}^{Ns} \frac{1}{\sqrt{|2\pi R|}} \exp\left(-\frac{1}{2}\frac{(\mathbf{z}^{ADSB} - \mathbf{x}^i_{val})^2}{R}\right) \tag{3.44}$$

If the ADS-B message is correct, the particle cloud computed by the particle filter will be very similar to the distribution of the ADS-B location and velocity, resulting in a high value.

Spoofing Hypothesis $\mathcal{H}_1$

$p\left(\mathbf{z}^{ADSB};\mathcal{H}_1\right)$ should describe the probability distribution of a spoofed target. The only information on the location of a spoofed target is that is should be within the area where one of the ground stations can receive the message. Then it is reasonable to assume that the location of the target is uniformly distribution within the area where ADS-B can be received. The uniform distribution is defined by eq. 3.45 where the boundaries are set by a and b.

$$\mathcal{U}(a,b) = \begin{cases} \frac{1}{b-a} & \text{for } a \leq x \leq b \\ 0 & \text{for } x < a \text{ or } x > b \end{cases} \tag{3.45}$$

For each dimension of $\mathbf{x}_{val}$ the boundaries a and b must be defined. Therefore each n'th dimension of $\mathbf{x}_{val}$ has its own uniform distribution as shown in eq. 3.46 where $n = 1, \ldots, N$

$$\mathcal{U}_n(a,b) = \begin{cases} \frac{1}{b_n - a_n} & \text{for } a_n \leq x_n \leq b_n \\ 0 & \text{for } x_n < a_n \text{ or } x_n > b_n \end{cases} \tag{3.46}$$

The multivariate uniformly distributed probability density function is obtained through multiplication. Note that the probability that no spoofing occurs is never zero. This is a logical result of the fact that $p\left(\mathbf{z}^{ADSB}; \mathcal{H}_1\right)$ approximates the physical area where ADS-B messages can be received. When an ADS-B messages is received, there is a possibility of spoofing. In the area where no ADS-B messages can be received, there is no possibility of spoofing, hence in that region $p\left(\mathbf{z}^{ADSB}; \mathcal{H}_1\right) = 0$. If the location inside the ADS-B message is falls outside the the area determined by $p\left(\mathbf{z}^{ADSB}; \mathcal{H}_1\right)$ it can immediately be concluded that the message is faked.

$$p\left(\mathbf{z}^{ADSB}; \mathcal{H}_1\right) = \begin{cases} \prod_{n=1}^{N} \frac{1}{b_n - a_n} = \frac{1}{V} & \text{for } a_1 \leq x_1 \leq b_1, \ldots, a_N \leq x_N \leq b_N \\ 0 & \text{for } x_1 < a_1 \text{ or } x_1 > b_1, \ldots, x_n < a_N \text{ or } x_N > b_N \end{cases} \tag{3.47}$$

$p\left(\mathbf{z}^{ADSB}; \mathcal{H}_1\right)$ approximates the physical area area where ADS-B can be received, but also the range of velocities that are acceptable for real ADS-B messages. It is possible for the velocity reported by an ADS-B message to exceed this range of velocities, resulting in $p\left(\mathbf{z}^{ADSB}; \mathcal{H}_1\right) = 0$. Then the LRT is not defined. If the velocity boundaries are set such that all physically possible velocities fall within the boundaries, and the ADS-B message still exceed these values, it can directly be concluded that the message is spoofed. Thus an undefined LRT does not have to be computed if the velocity boundaries are set correctly.

Likelihood Ratio Test

Combining the likelihoods $p(\mathbf{x}_k; \mathcal{H}_0)$ and $p(\mathbf{x}_k; \mathcal{H}_1)$ eq. 3.48 is obtained. For every ADS-B message the LRT is computed and compared against the threshold τ. Then the output of the algorithm is equal to a one if the message is considered *not* spoofed, or a zero is the message is considered to be spoofed.

$$L(\mathbf{z}^{ADSB}) = \frac{p\left(\mathbf{z}^{ADSB}; \mathcal{H}_0\right)}{p\left(\mathbf{z}^{ADSB}; \mathcal{H}_1\right)} = \frac{\sum_{i=0}^{Ns} \frac{1}{\sqrt{|2\pi R|}} \exp\left(-\frac{1}{2} \frac{(\mathbf{z}^{ADSB} - \mathbf{x}_{val}^{i})^2}{R}\right)}{\frac{1}{V}} \lessgtr \tau \tag{3.48}$$

3.3.2. Minimum Bayes Risk

In the Minimum Bayes Risk (MBR) framework, the decision rule is derived by minimizing the expected loss. The loss function quantifies the cost associated with different types of errors, such as the cost of a false alarm or a missed detection. By considering the probabilities of different outcomes and their associated costs, the MBR test provides a decision rule that minimizes the average risk or expected loss.

MBR assigns a cost to each decision as shown in eq. 3.38. In a binary hypothesis test there are two types of errors, a false alarm P (choose $\mathcal{H}_0 \mid \mathcal{H}_1$) or a missed detection P (choose $\mathcal{H}_1 \mid \mathcal{H}_0$). From these expressions a probability of error is defined as seen in eq. 3.49. In this expression for P_e the prior probability of hypothesis $\mathcal{H}_1$ occurring is denoted by π_1, hypothesis $\mathcal{H}_0$ equals π_0.

$$P_e = \pi_1 P\left(\text{choose } \mathcal{H}_0 \mid \mathcal{H}_1\right) + \pi_0 P\left(\text{choose } \mathcal{H}_1 \mid \mathcal{H}_0\right) = \pi_1 P_m + \pi_0 P_{FA} \tag{3.49}$$

Both types of error are assigned a cost C_{ij} where $i = 0, 1$ and $j = 0, 1$.

$$\mathcal{R} = C_{10}\pi_1 P_m + C_{01}\pi_0 P_{FA} \tag{3.50}$$

The threshold that achieves the MBR $\mathcal{R}$ is shown in eq. 3.51. This threshold selects the hypothesis that has the maximum a posteriori probability, thus the MBR estimate is also a MAP estimator.

$$\frac{p\left(\mathbf{x}; \mathcal{H}_0\right)}{p\left(\mathbf{x}; \mathcal{H}_1\right)} \lessgtr \frac{C_{01} - C_{11}}{C_{10} - C_{00}} \frac{\pi_1}{\pi_0} = \tau \tag{3.51}$$

If the probabilities π_0 and π_1 are known, and the associated costs are selected, the threshold is easily computed. In this application estimating π_1 is similar to estimating the probability that spoofing occurs. This is a rather difficult question to answer, and especially to find a mathematical reasoning behind the resulting value. An advantage to the MBR test is that the threshold does not requires the density functions in eq. 3.38 to be known. In many cases is it impossible to find such probability distribution without assuming that some known pdf fits the data.

The costs C_{10} and C_{01} can be used to penalize the different types or error the estimator can make. If it is less preferable for a detector to make a false alarm error instead of a missed detection, the cost C_{01} can be set higher then the cost C_{10}.

3.3.3. Neyman-Pearson Hypothesis Test

The Neyman-Pearson detector is particularly useful in scenarios where the primary objective is to achieve high detection sensitivity, which can be advantageous for spoofing detection. A high detection results in fast spoofing detection, but at the cost of a higher false alarm probability. The Neyman-Pearson detector maximizes P_d for a fixed value of P_{fa}. The threshold is obtained by solving eq. 3.52. Solving this integral is equal to finding an expression of the volume inside a 4-dimensional Gaussian distribution as a function of τ. For this to be solved the full statistics of hypothesis π_0 must be known.

The threshold in this hypothesis test is defined by a chosen value for the false alarm rate, which can be argued is a better tuning parameter than a guess for $P(\mathcal{H}_0)$.

$$P_{fa} = \int_{\{\mathbf{x}_{val}:L(\mathbf{z}_{ADSB})>\tau\}} p\left(\mathbf{x}_{val};\mathcal{H}_0\right) d\mathbf{x}_{val} = \alpha \tag{3.52}$$

If the posterior density of the PF is assumed to be Gaussian distributed, the integral of equation 3.43 can be solved analytically. This means that a Neyman-Pearson test is possible under a Gaussian assumption of $p(\mathbf{x}_{val}|\mathbf{z}^{MLAT})$ then the solution is given by eq. 3.53b.

$$p\left(\mathbf{z}^{ADSB};\mathcal{H}_0\right) = \int \underbrace{p(\mathbf{z}^{ADSB}|\mathbf{x}_{val};\mathcal{H}_0)}_{Gaussian} \underbrace{p(\mathbf{x}_{val}|\mathbf{z}^{MLAT})}_{Gaussian} d\mathbf{x}_{val} \tag{3.53a}$$

$$p\left(\mathbf{z}^{ADSB};\mathcal{H}_0\right) = \mathcal{N}(H\mathbf{x}_k, HPH^T + R) \tag{3.53b}$$

R is the covariance matrix corresponding to the ADS-B message, and P the covariance matrix of $p(\mathbf{x}_{val}|\mathbf{z}^{MLAT})$ which can directly be estimated from the particle cloud. The observation matrix H links the measured Geodetic coordinates to the Cartesian coordinates, this is required because the ADS-B measurements are in latitude, longitude and FL, and $p(\mathbf{x}_k|\mathbf{z}_k^{MLT})$ is defined in ECEF coordinates. The H matrix describes the transformation of coordinate system as described in equation 3.40a.

ARTAS documentation allows for quick conversion of the measurement uncertainty between the measured Geodetic quantities to the required Cartesian uncertainty.

$$S = HPH^T + R = \tilde{P} + R \tag{3.54}$$

The LRT as a result of the Gaussian assumption is then,

$$L(\mathbf{x}) = \frac{p\left(\mathbf{x}_{val};\mathcal{H}_0\right)}{p\left(\mathbf{x}_{val};\mathcal{H}_1\right)} = \frac{\frac{1}{\sqrt{|2\pi S|}} \exp\left(-\frac{1}{2}\frac{(\mathbf{z}^{ADSB}-\mathbf{x}_{val}^i)^2}{S}\right)}{\frac{1}{V}} > \tau \tag{3.55}$$

The false alarm probability is found by using $P_{fa} = 1 - P_{rej}$. The probability of correctly concluding that an ADS-B message is not spoofed is equal to the probability that a valid ADS-B measurement falls within a 4-dimensional Gaussian gate of size G [41] as in eq. 3.56. This approach is inspired the approacht Blackmans *Multiple-Target Tracking with Radar applications* takes in track association problems. Blackman defines a gate as a region a measurement must fall within to be associated with a

corresponding track. The same general approach is applied in the hypothesis test, here the gate is used to determine if the ADS-B message and MLAT measurement originate from the same location. The ADS-B messages must thus fall within the gate defined by the MLAT measurement.

$$P_{rej} = \left(1 - \left(1 + \frac{G}{2}\right)\exp\left(-\frac{G}{2}\right)\right) \tag{3.56}$$

G in eq. 3.56 is obtained by solving eq. 3.55 for all data dependent parameters.

$$G = -2\ln\left(\frac{\tau}{V}\sqrt{2\pi|S|}\right) \tag{3.57}$$

The expression for P_{fa} can be found from eq. 3.56.

$$P_{fa} = \left(1 + \frac{G}{2}\right)\exp\left(-\frac{G}{2}\right) \tag{3.58}$$

Eq. 3.57 is substituted into eq. 3.58 and we obtain a value for P_{fa} that is a function of the threshold τ.

$$P_{fa} = \left(1 + \frac{-2\ln\left(\frac{\tau}{v}\sqrt{2\pi|S|}\right)}{2}\right)\exp\left(-\frac{-2\ln\left(\frac{\tau}{V}\sqrt{2\pi|S|}\right)}{2}\right) \tag{3.59}$$

P_{fa} is a tuning variable which can be set to the preference of the operator of the validation algorithm. For the operator of the validation algorithm it can also be of added value to not only minimize the P_{fa}, but also the other error the detector can make, namely P_m. The Minimax Hypothesis test attempts to achieve this by minimizing both P_{fa} and P_m.

3.3.4. Minimax Hypothesis Test

To construct a minimax test, similar to the MBR test, one defines a loss function that quantifies the cost or penalty associated with the two types of error. This assigned cost to the two types or error reflects the decision maker's preferences and the relative importance of different outcomes. The goal is then to find a decision rule that minimizes the maximum possible loss, considering the underlying probability distributions.

The conditional risks are defined as,

$$R_0(\tau) = C_{10}P_m + C_{00}P_{rej} \tag{3.60a}$$

$$R_1(\tau) = C_{11}P_d + C_{01}P_{fa} \tag{3.60b}$$

These two are combined to obtain the total Bayes Risk,

$$\mathcal{R}(\tau) = \pi_0 R_0(\tau) + \pi_1 R_1(\tau) \tag{3.61}$$

The MBR test assumes π_0 known and the minimum probability of error is then obtained as was shown in equation 3.51. The Minimax test argues that because it can be undesirable to guess some value π_0, then one can attempt to minimize the maximum Bayes Risk. As the Bayes Risk depends on the threshold τ and the prior probability, the Bayes Risk is a function defined as $\mathcal{R}(\pi_0, \tau)$.

$$\mathcal{R}(\pi_0, \tau) = \pi_0 R_0(\tau) + (1 - \pi_0) R_1(\tau) \tag{3.62}$$

$$\min\left(\max_{0\leq\pi_0\leq 1} r(\pi_0, \tau)\right) \tag{3.63}$$

It has been shown in [42] that when the two conditional risks are equal, the maximum Bayes Risk is minimized. The threshold can thus be obtained by eq. 3.64 which is called the equalizer rule.

$$R_0(\tau) = R_1(\tau) \tag{3.64}$$

Under the Gaussian assumption of $p(\mathbf{x}_k|\mathbf{z}_k)$, the conditional risks in eq. 3.60 can be computed analytically. Similar to the MBR hypothesis test the cost associated with a correct decision is zero, meaning $C_{11} = C_{00} = 0$.

$$C_{10}P_m = C_{01}P_{fa} \tag{3.65}$$

From eq. 3.59 the P_{fa} is known. The expression for P_m is shown in eq. 3.66b [41].

$$P_{fa} = \left(1 + \frac{-2 * \ln\left(\frac{\tau}{V}\sqrt{2\pi|S|}\right)}{2}\right)\exp\left(-\frac{-2 * \ln\left(\frac{\tau}{V}\sqrt{2\pi|S|}\right)}{2}\right) \tag{3.66a}$$

$$P_m = \frac{1}{V}\frac{\pi^2}{2}\sqrt{|S|}G^2 \tag{3.66b}$$

Eq. 3.66a and 3.66b are substituted into equation 3.65 and solved for τ. For every ADS-B message that is received a hypothesis test is computed. Only S changes per ADS-B message as it is the covariance matrix of $p(\mathbf{x}; \mathcal{H}_0)$. Hypothesis $\mathcal{H}_1$ remains constant for every hypothesis test, because it is assumed to be uniformly distributed.

4

Case Study

With a case study the performance of the proposed validation algorithm is investigated and analyzed. This is done in this chapter and in chapter 5. In this chapter the case study is introduced, and in chapter 5 the results are presented. In the case study Surveillance Data North Sea (SDNS) is used. This is a operational WAM system in use by the LVNL and is introduced in section 4.1. The MLAT and ADS-B measurements that is provides need to be pre-processed, clustered and extrapolated. These pre-processing steps are discussed in section 4.2. In the presented algorithm several tuning variables need to be set, which is done in section 4.3. Then, in chapter 5 the results of the algorithm are presented.

4.1. Surveillance Data North Sea

SDNS is a WAM system designed to provide surveillance coverage on the North Sea. The system is designed and developed by ERA, a Czech company specialized in civil and military radar solutions. On the North Sea LVNL provides Flight Information Services for helicopters that travel between the oil rigs, and en-route traffic. Before SDNS was operational there was no radar coverage at low altitudes, SDNS solves this problem by making use of the oil rigs located in the North Sea. On these oil rigs the GSs are installed. The WAM system covers an area of 30 thousand square miles from an altitude of 500 ft. This is achieved with 17 remote sites on oil platforms and three onshore sites. Nearly all of the sites have two directional antennas providing full 360 degree coverage. All data is transmitted to LVNL headquarters located at Schiphol-Oost for central processing and tracking.

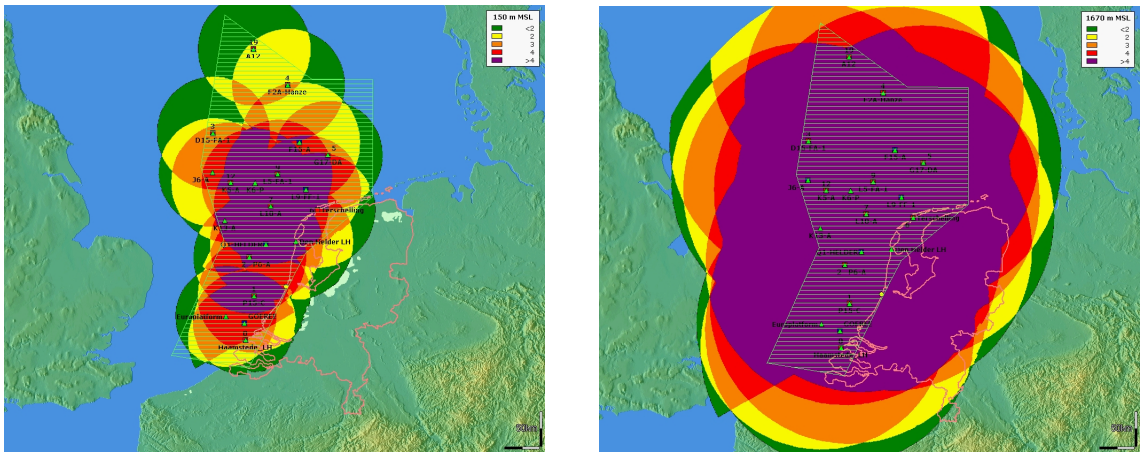

(a) Coverage at 500ft.

(b) Coverage at 5500ft.

Figure 4.1: Coverage map of SDNS, on the left the coverage is shown for 500 ft. where the green with the green stripes indicates the required coverage. The realized coverage is where there are more than 4 GSs, thus the red and purple area. The right image shows the coverage at 5000 ft.

In Fig. 4.1 an overview of SDNS is shown with coverage at 500 ft. and 5500 ft. Four or more GSs are needed to determine the location of a target, the area covered by four GSs is shown in red, and more than four GSs is shown in purple. In Fig. 4.1a the coverage is not sufficient in all areas. At the outer edges of the airspace (indicated with green horizontal stripes) there are large regions that are covered by only one, two of three GSs. Such an area is called an ADS-B only area, because WAM cannot determine the location of the target, but ADS-B can. With the proposed validation method at least two GSs are needed to validate the ADS-B message, then ADS-B messages are validated in all areas except the green area. In the area covered by at least two GSs ADS-B (if validated) can be used by ARTAS, resulting in larger coverage. The strength of this validation method is that by combining the ADS-B and WAM measurements the covered area is increased from the area where four GSs detect the message to the area where two GSs detect the message. This brings offshore sites such as A12 and F2A-Hanze within coverage at low altitudes.

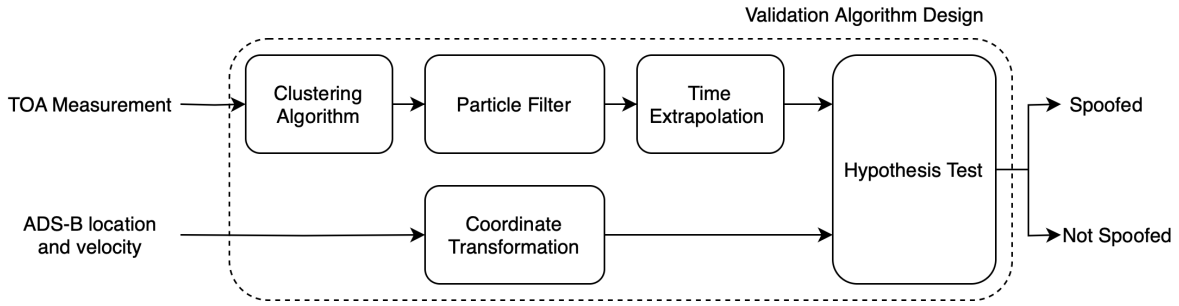

Figure 4.2: Block representation of the designed validation algorithm applied to the case study SDNS

SDNS is the only WAM system operational at LVNL, this means that it is the only option to obtain data and analyze the performance of the validation algorithm. The algorithm is slightly tweaked to process the data provided by SDNS as is shown in Fig. 4.2.

Time Synchronization

The EM waves that are transmitted by the ADS-B transponder propagate at the speed of light. In 1 ns the wave travels 33 cm, this requires timestamping to be done in the range of the tenths of nanoseconds for accurate localization. In addition to this GSs in SDNS must be synchronized. Local oscillators in industrial hardware can achieve the time accuracy, but the main challenge is to synchronize all GSs. Synchronization is required because the oscillator used is able to stay synchronized for only about 30 minutes. After this the TDOA measurements are not accurate enough due to oscillator drift. SDNS solves this problem by using the Common View GNSS synchronization method.

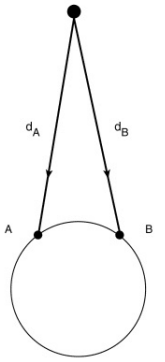

Figure 4.3: GSs A and B have the same satellite in view that is used for synchronization. Both GSs compute the time difference between its own clock and the time reported by the satellite. These time differences are shared between the GSs and the time can be synchronized. Illustration from [5].

Fig. 4.3 shows two GSs we want to synchronize. Both GSs observe the same satellite in the GNSS constellation, the satellite transmits its location and the time measured from its atomic clock. Both GSs A and B measure the same message that the has satellite transmitted, with the location reported by the satellite both GSs can compute the travel time of the message between the satellite and itself. The GSs now measure a time difference between the time in the satellite's message and its own local time. The time differences measured at GS A and B are then shared with each other, with this information the time between the two GSs can be synchronized. This procedure is applied by the system to all GSs.

4.2. Pre-Processing

The data provided by SDNS requires several pre-processing steps. The measurements that are obtained are the TOA measurements and the corresponding pre-processing steps are shown in subsection 4.2.1. These TOA measurements need to be clustered such that the TOA measurements all originate from the same ADS-B broadcast. The clustering method that is used is discussed in section 4.2.2. Some pre-processing is required on the ADS-B data such that fits the input of the validation algorithm, this is discussed in subsection 4.2.3. The MLAT and ADS-B measurements need to be correlated to one another which is elaborated in subsection 4.2.4.

4.2.1. WAM Data Pre-processing

ERA has provided LVNL software that obtains TOA measurements. It provides all TOA measurements of all received messages the GSs are able to decode. Therefore all transponder interrogations replies are inside the data that is received. These types of replies are labeled and called Down-Link Formats (DF) in civil radar surveillance. ADS-B messages are assigned DF-17, therefore we can remove all data that is not DF-17.

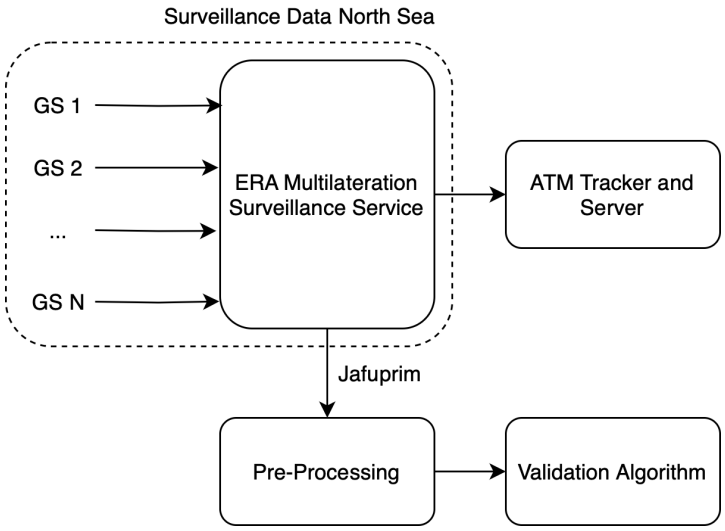

Figure 4.4: Schematic overview of SDNS with corresponding relevant output. Jafuprim is the name of the software ERA has provided to obtain the TOA measurements.

Fig. 4.4 schematically shows the system and how the data it obtained. An example of the data that it provides is shown in Fig. 4.5. The data shown is same data that is used by the central processing unit. This processing is done in ERA's Multilateration Surveillance Service software, from where the output is presented to ARTAS.

```
DMU[ns]=3.125000 DMI[us]=6.400000 DLI[ms]=49.996800 DLI[dmi]=7812 DLI[dmu]=15998976 DMI[dmu]=2048
 block#      byte#   unixtim.nnnn     mitim   ch   toaf1    toaf2                      code  iff spi
 block#      byte#   unixtim.nnnn     mitim   ch   toap1    toap2    toap3    toap4  adress DF parity
     4;s;     84ec;1669798549.0399;1233508016;17;       0;       0;       0;   51028;      ;20;198A0D;
     4;s;     8518;1669798549.0399;1233508016; 2;       0;       0;       0;   60147;      ; 4;D1A3C0;
     4;s;     8534;1669798549.0399;1233508016; 1;       0;       0;       0;   60157;      ; 4;D1AB0D;
     4;s;     8550;1669798549.0399;1233508016;27;       0;       0;       0;   60295;A5ED66;11;0EF7AB;
```

Figure 4.5: Data format

The TOA measurements are stored by making use of three different data fields. By combining these datafields accordingly the TOA measurements accurate up to 3.125 ns can be obtained.

unixtim.nnnn
First data field is called *unixtim.nnnn* and is stored the unix time corresponding to the TOA measurement. This time is used to timestamp the resulting location estimate. The following two data registers are used to determine the TOA measurement.

Time of Arrival Measurement
The TOA measurement is divided into two different counters that together represent the time measurement. One section is called *toaxx* and has 32 bits, the second section *mitim* has 11 bits. Together they construct a virtual 43 bit counter called the MU clock. The smallest time granularity the system is able to measure is 3.125 ns (the *toaxx* counter) and overflows every 13.24 seconds. *mitim* time is higher orders of the *toaxx* clock and has a granularity of 6.4 ms, which is confirmed by $2^{11} * 3.125ns = 6.4ms$. The MU clock overflows every $2^{32} * 6.4ms = 27487.79069s$ or $2^{43} * 3.125ns = 27487.79069s$ which is about 7.6 hours. *toaxx* and *mitim* need to be combined to obtain the TOA measurement which is done using eq. 4.1.

$$TOA = mitim * 6.4 * 10^{-6} + toaxx * 3.125 * 10^{-9} \tag{4.1}$$

When a pulse arrives, it is assigned a TOA depending on the shape of the preamble. For ADS-B messages the preamble has four pulses, where the time between the pulses indicates to the GS that an ADS-B message is received. The preamble for ADS-B is shown in Fig. 4.7. An ADS-B message has 4 pulses in the preamble and the TOA assigned is the average of all these four pulses, and stores this in the *toap4* column.

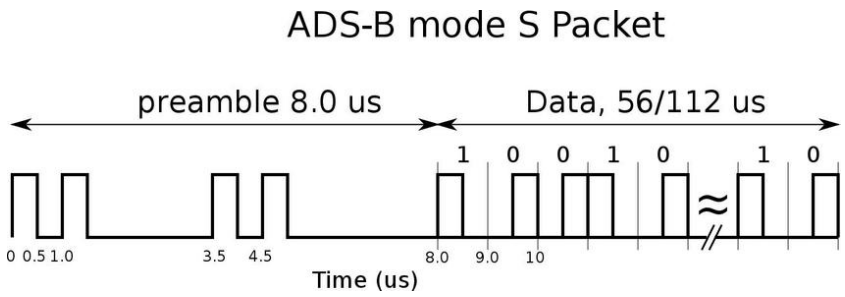

Figure 4.6: ADS-B preable and data (illustration from[6])

From the data shown in Fig. 4.5 the TOA measurements are computed, and the difference is taken with respect to the first received TOA measurement. In Fig. 4.7 the result of this operation is shown. The next step is to compute the TDOA clusters such that every TDOA measurement in a cluster originates from the same ADS-B message.

	1 unix	2 toaHIGH	3 GS	4 toaLOW	5 adress	6 DF	7 time	8 time_norm	9 datetime
1	1.6698e...	1.2728e+09	13	1769158	'484B32'	17	8.1456e...	0	30.11.2022 9:00:...
2	1.6698e...	1.2728e+09	11	1805001	'484B32'	17	8.1456e...	1.1201e-04	30.11.2022 9:00:...
3	1.6698e...	1.2728e+09	25	1810697	'484B32'	17	8.1456e...	1.2981e-04	30.11.2022 9:00:...
4	1.6698e...	1.2728e+09	22	1827035	'484B32'	17	8.1456e...	1.8087e-04	30.11.2022 9:00:...
5	1.6698e...	1.2728e+09	19	12904491	'484B32'	17	8.1459e...	0.3348	30.11.2022 9:00:...
6	1.6698e...	1.2728e+09	27	12920611	'484B32'	17	8.1459e...	0.3348	30.11.2022 9:00:...
7	1.6698e...	1.2728e+09	14	12925905	'484B32'	17	8.1459e...	0.3348	30.11.2022 9:00:...
8	1.6698e...	1.2728e+09	11	12961808	'484B32'	17	8.1459e...	0.3350	30.11.2022 9:00:...
9	1.6698e...	1.2728e+09	25	12967476	'484B32'	17	8.1459e...	0.3350	30.11.2022 9:00:...
10	1.6698e...	1.2728e+09	1	12979525	'484B32'	17	8.1459e...	0.3350	30.11.2022 9:00:...
11	1.6698e...	1.2728e+09	19	13675306	'484B32'	17	8.1460e...	0.4372	30.11.2022 9:00:...
12	1.6698e...	1.2728e+09	27	13691422	'484B32'	17	8.1460e...	0.4372	30.11.2022 9:00:...
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
35585	1.6698e...	1.5253e+09	16	7145242	'484B32'	17	9.7621e...	1.6165e+03	30.11.2022 9:26:...
35586	1.6698e...	1.5253e+09	30	7229358	'484B32'	17	9.7621e...	1.6165e+03	30.11.2022 9:26:...
35587	1.6698e...	1.5253e+09	16	8783833	'484B32'	17	9.7621e...	1.6165e+03	30.11.2022 9:26:...
35588	1.6698e...	1.5254e+09	16	4864345	'484B32'	17	9.7625e...	1.6169e+03	30.11.2022 9:26:...

Figure 4.7: SDNS data format after TOA extraction. The columns *toaHIGH* and *toaLOW* are combined to obtain the true TOA column *time*.

4.2.2. Clustering Algorithm

TOA measurements need to be clustered such that all TOA measurement that are grouped originate from the same ADS-B message. Unfortunately, there is no information the TOA measurement output from which the clustering can be done automatically (i.e. some ADS-B message identifier). Thus in pre-processing the data must be clustered. The proposed clustering algorithm is based on the fact that there is a maximum time difference measurement for which measurements originate from the same ADS-B transmission, and this maximum occurs when the distance difference between the two GSs contributing to the measurement is the largest.

The largest distance difference between any two GSs needs to be obtained. Between any two GSs the maximum distance difference happens when the target is located at one of the GSs, thus the maximum distance difference is equal to the distance between the two GSs that are furthest apart. From Fig. 4.1 it can be seen that the two GSs that are furthest apart are Haams and F2AH, and the distance between them is 364 km. Due to the curvature of the earth the path between the GSs travels through the earth, meaning this is not practically possible to obtain such a distance. Finding the theoretical maximum difference that does not penetrate the earth is a more complex task, but this approximation has been found to work adequately. For the upper bound to be reached Haams and F2AH need to both detect the same ADS-B message which is highly unlikely.

It is found that the maximum time difference within a single TDOA cluster is $TDOA_{max} = \frac{364\text{km}}{c} = 1.2\text{ms}$. Next, to compute the clusters a Matlab script loops through received TDOA measurements such as that in Fig. 4.7. It computes the difference between following TOA measurements and if the difference between two consecutive TOA measurements is lower that 1.2ms they are clustered, if the difference is larger than 1.2 ms the TOA measurement must be resulting from the next ADS-B message and thus a new measurement cluster is created. This process is repeated for all TOA measurements.

4.2.3. ADS-B Data Pre-processing

ADS-B messages are standardized to assure data compatibility between all parties using ADS-B. EUROCONTROL is a pan-European organization managing such standardization tasks. The all-purpose structured EUROCONTROL surveillance information exchange (ASTERIX) defines the data format for all surveillance related data exchanges. The ASTERIX data format helps ANSPs share information in an automated and standardized manner, but it is also used by ANSPs for communications between all its surveillance systems and ARTAS. ADS-B messages are standardized according to the Category 21 (CAT-21) [10] format. This documents describes the data fields and how is should be encoded. This document therefore also illustrates the vulnerability of ADS-B because it effectively tells a spoofer how to forge an ADS-B message. The reader is referred to chapter five of [10] to see all data that can be transmitted using ADS-B.

From the CAT-21 messages received by SDNS the relevant data is selected,

- **Position**: Latitude, Longitude, FL

- **Velocity**: Ground Speed, Climb Rate

Section 3.3.1 eq. 3.40 illustrates the pre-processing steps needed to compute the position and velocity in the units needed by the filter, from the units the ADS-B message is received in. These are the only pre-processing steps required for the ADS-B messages.

4.2.4. Time Extrapolation

It occurs that ADS-B messages received by SDNS contain no location. The GSs still received this message, assigned a TOA and determined based on the preamble that is it in fact an ADS-B message. The clustering algorithm thus can calculate a valid MLAT measurement based on this ADS-B message with no location.

For LVNL such ADS-B messages with no location are of no use and deleted. Consequently, there are more MLAT measurements than there are ADS-B messages at the output of the system. In this case study the MLAT measurements resulting from ADS-B messages without location are still used to improve the validation quality. The state estimate based on an ADS-B message without location is extrapolated in time to the next ADS-B message with a location report. Using time extrapolation the message can still be validated based on the location report of the next ADS-B message.

To illustrate the time extrapolation an ADS-B message without a location report is denoted with m, and the next ADS-B message with a location report is denoted with n. The particle cloud based on the MLAT measurement m is predicted from time T_m to time T_n by eq. 4.2, where $T = T_n - T_m$.

$$\begin{bmatrix} l_x \\ l_y \\ l_z \\ v_x \\ v_y \\ v_z \end{bmatrix}_n^i = \begin{bmatrix} 1 & 0 & 0 & T & 0 & 0 \\ 0 & 1 & 0 & 0 & T & 0 \\ 0 & 0 & 1 & 0 & 0 & T \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} l_x \\ l_y \\ l_z \\ v_x \\ v_y \\ v_z \end{bmatrix}_m^i + \begin{bmatrix} \frac{T^2}{2} & 0 & 0 \\ 0 & \frac{T^2}{2} & 0 \\ 0 & 0 & \frac{T^2}{2} \\ T & 0 & 0 \\ 0 & T & 0 \\ 0 & 0 & T \end{bmatrix} \begin{bmatrix} w_x \\ w_y \\ w_z \end{bmatrix}_m^i \qquad w \sim \mathcal{N}(0, \sigma_w^2) \tag{4.2}$$

Due to the uncertainty in the prediction, the quality of the estimate decreases when compared to an ADS-B message that includes a location report. The time between ADS-B messages with and without a location report is usually not more than one second, (i.e. two consecutive messages) therefore the loss in quality is limited. This proposed time extrapolation allows for the use of all ADS-B messages which increases the quality of the validation algorithm. The increased quality results from the fact that more messages are used. Off course this method only works if there are enough ADS-B messages that do contain location, such that the time extrapolation uncertainty does not become to large.

4.3. Tuning

In the next chapter the result of the validation algorithm are presented. To obtain these results several tuning parameters need to be set. In this section all tuning parameters are discussed and assigned a value.

Process noise Tuning			
location variance	$\sigma_x^2 = 7^2$	$\sigma_y^2 = 7^2$	$\sigma_z^2 = 4^2$
velocity variance	$\sigma_{vx}^2 = 1^2$	$\sigma_{vy}^2 = 1^2$	$\sigma_{vz}^2 = 0.5^2$
$p(\mathbf{x}_k, \mathcal{H}_1)$ **Tuning**			
location	$V_x = 220 * 10^3$	$V_y = 523 * 10^3$	$V_z = 21 * 10^3$
velocity	$V_v = 225 * 10^3$		
Hypothesis Test Tuning	**Bayes Risk**	**Minimax**	**Neyman-Pearson**
π_0	0.75		
π_1	0.25		
C_{00}	0	0	
C_{01}	1	1	
C_{10}	1	1	
C_{11}	0	0	
P_{fa}			10^{-6}

Table 4.1: Tuning variables

Other tuning variables	**Number of particles**	**Epsilon Robustness**
	$N_s = 5 * 10^4$	$\epsilon = 0.99$
		$\frac{1}{P} = 10^{-308}$
	Proposal Density Grid	**ADS-B velocity variance**
	$\mathbf{l}_x = [x_{asdb} - 1e3\ , \ldots , x_{adsb} + 1e3]$	$\sigma_a^2 = 4^2$
	$\mathbf{l}_y = [y_{asdb} - 1e3\ , \ldots , y_{adsb} + 1e3]$	
	measurement std	
	$\sigma_v = 10^{-7}$	

Table 4.2: Tuning variables

The proposal density grid size is a square around the ADS-B location of 2km wide.

Tuning $p(\mathbf{x}_k, \mathcal{H}_1)$

The area used is a rectangular approximation of the area where ADS-B and MLAT measurements can be received. In Fig. 4.8 the approximation can be seen. It is an area 220 km wide and 523 km long. The maximum receivable height is approximated with 16 km, but due to the curvature of the earth the lowest point is 5 km below the land.[1] The expected range of velocities is based maximum expected velocities in the area above the North Sea. Airplanes on cruising speed on average have an maximum speed of 400 knots.

[1]Where a height of zero corresponds to Schiphol Airport

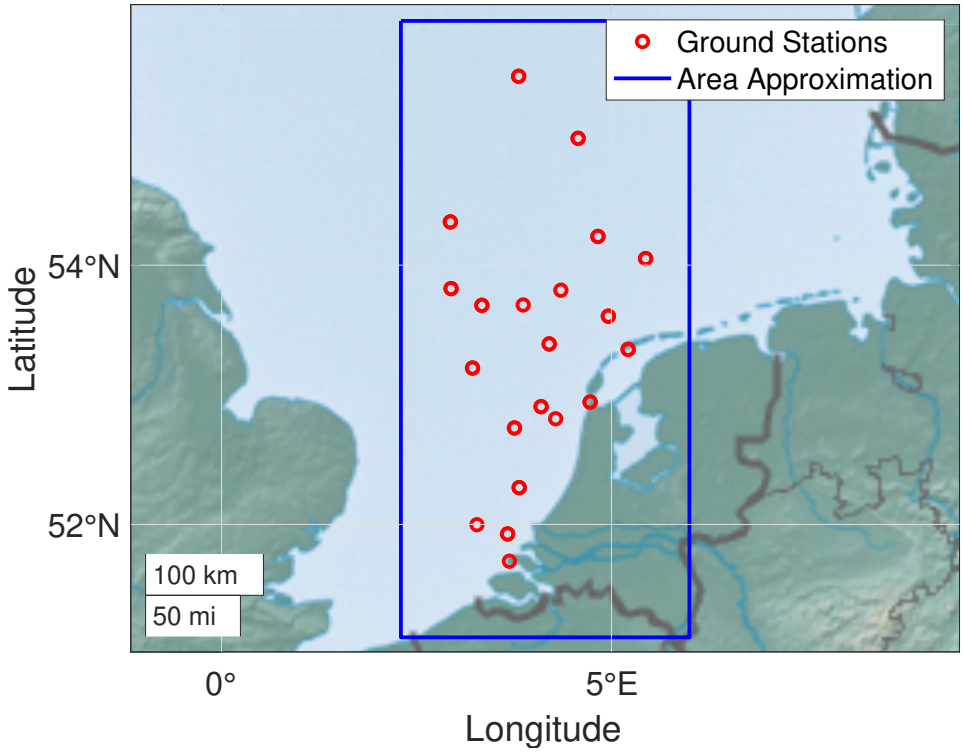

Figure 4.8: Overview of locations of ground stations in the SDNS system. Blue box indicates rough approximation of the area where ADS-B can be received. Red dots are the GSs.

5 Results

In this chapter the results of the case study are presented. Firstly, the filter's performance will be analyzed using several real flights that have flown across the North Sea. In this analysis the state estimation quality is investigated and the effect of parameters. Then in section 5.2 the results of the three investigated hypothesis tests on three flights. In section 5.3 the entire algorithm is tested on spoofed ADS-B messages, and in section 5.4 the chapter is concluded.

5.1. Filter Analysis

The performance analysis of the proposed filtering algorithm is done by analyzing three types of flights. Firstly, an en-route flight that is flying at high altitude in a multiple sensor coverage is investigated. Secondly an approaching flight at low altitude with limited sensor coverage is investigated. Thirdly, an helicopter at low altitude is analyzed. All tracks are selected based on their ICAO address, but are anonymized to flight A, B and C.

The ADS-B tracks used are compared to the ARTAS tracks to know if they are in fact not spoofed. These tracks can be found in appendix A The ARTAS tracks are computed using all available surveillance systems that detect the target. For flight A and B these are several primary and secondary surveillance radars located around The Netherlands and even Belgium. For flight C at low altitude the only surveillance source available in ARTAS is SDNS. SDNS only initiates a track using more than 3 GSs per MLAT measurement, and the SIR filter initiates a track using at least 2 GSs per MLAT measurement. Because the SIR filter can track the target before ARTAS can track the target.

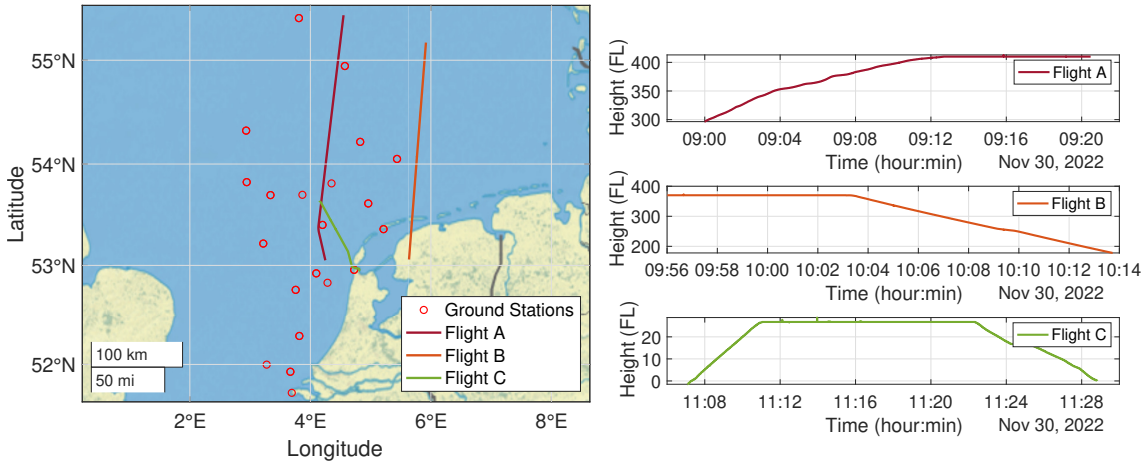

(a) Geographic overview of flights used to investigate the performance and results of the validation algorithm

(b) Altitude plot of flights used to investigate the performance and results of the validation algorithm

Figure 5.1: Overview of flight A, B and C that are used for investigation of the performance of the validation algorithm.

Flight A is a north-bound flight from Schiphol Airport. In this simple scenario the track is initiated after the aircraft has reached FL300. At this altitude the ADS-B messages are received by a large amount of GSs allowing for the best possible tracking of the target.

Flight B is chosen for analysis because it travels along the edges of the covered area which results in a high number of ambiguous measurements. Dealing with ambiguous measurements was an important aspect of the validation algorithm, this flight allows for investigation of this aspect.

Flight C is a helicopter that is traveling between an oil rig and Den Helder. Den Helder Airport is the location from which the helicopters depart to travel between oil rigs located in the North Sea. This type of flight is common in the North Sea at low altitudes. The goal of this analysis is to investigate the performance of the validation algorithm on a helicopter because it has different flight dynamics compared to an fixed wing aircraft, and secondly to investigate the performance at low altitude where vertical dilution of precision can be a significant issue.

5.1.1. SIR Filter Analysis

In this section the performance of the SIR filter is discussed. The quality of the location and velocity estimate is investigated including the Effective Sampling Size (ESS). In addition to discussing the output of the SIR filter also the effect of the tuning variables is discussed, such as the used number of samples, the process and measurement noise, and the impact the proposal distribution has on the performance of the filter. The implemented SIR filter is also compared to a traditional SIR filter where the proposal distribution is uniformly distributed.

Location

The SIR filter initiates a track using the received MLAT measurements of flight A, B and C. These tracks are plotted in Fig. 5.2 along the with ADS-B track. The MLAT track determined by the SIR filter is the mean estimate of the target. In general, the SIR filter is capable of dealing with a considerable amount of ambiguous measurement. Fig 5.3 shows the amount of GSs per measurement for each flight using a histogram. From here it can be observed that around half of the MLAT measurements in flight B are ambiguous measurements. During the design phase is was expected that such amount of ambiguous measurements will cause the SIR filter to fail as the result of the degeneracy or impoverishment problem. This reasoning led to the suggestion of the MISS filter as possible solution. Results now show that the filter is perfectly capable of tracking flight B.

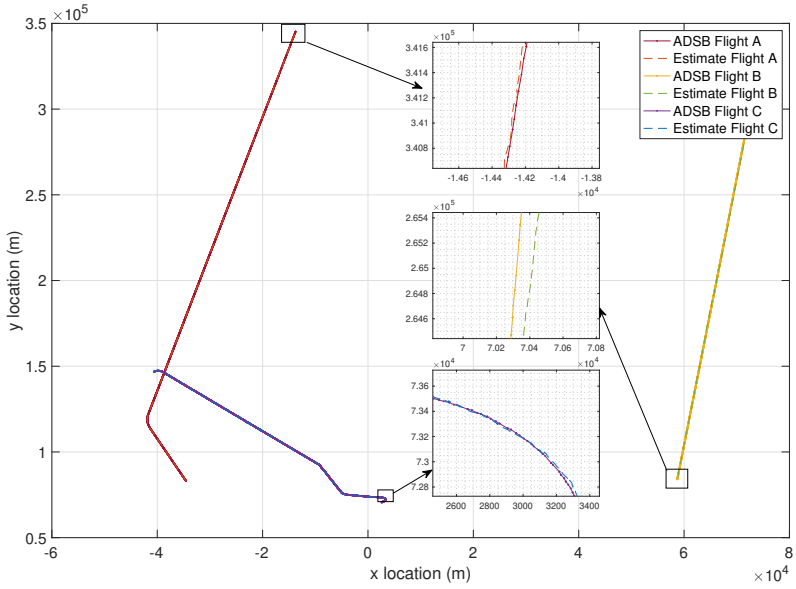

Figure 5.2: Overview of flight A, B and C. Axis of mini-figures are 1 km by 1 km.

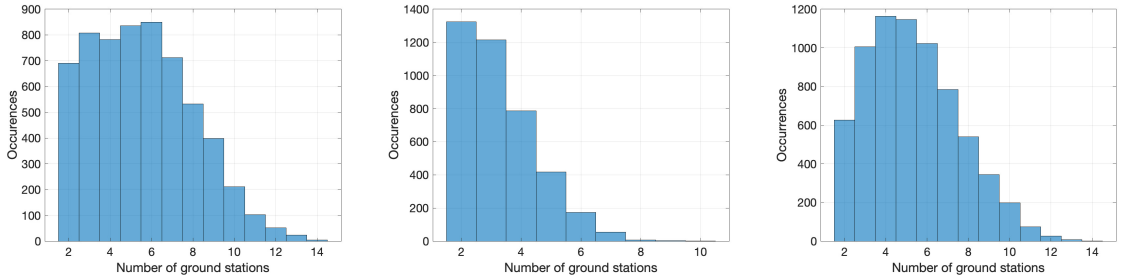

(a) Flight A: Histogram of number of ground stations per measurement

(b) Flight B: Histogram of number of ground stations per measurement

(c) Flight C: Histogram of number of ground stations per measurement

Figure 5.3: Histogram of GS distribution for each flight

The estimate of the altitude the filter provides of each flight is shown in Fig. 5.4. The estimate shown is the weighed mean of the particle cloud, where the weights determined by the likelihood function are used to compute this weighted mean. Both flights A and C need several measurements to accurately compute the altitude of the target. At lower altitude the effect of Vertical Dilution of Precision has a bigger effect on the estimate than at higher altitudes. Although the hypothesis test takes into account the full statistics of the estimate, and in the figure only the weighted mean is shown, the figure indicates that height estimation at lower altitude can be difficult.

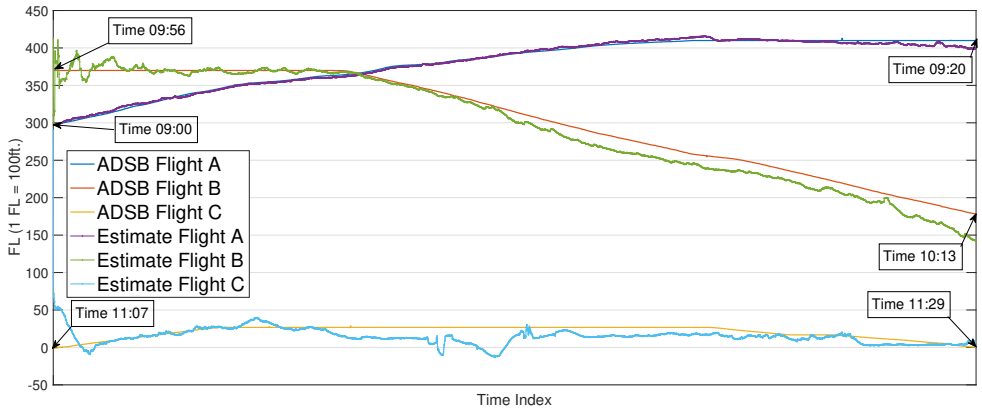

Figure 5.4: Height estimate and ADS-B altitude for flight A, B and C. The estimate is obtained by computing the weighted mean of the particle cloud.

Velocity

In the validation algorithm both position and velocity are used to validate the ADS-B message. Fig. 5.5 shows the velocity estimate by the SIR filter. From flight B it can be seen that the ADS-B velocity lags behind the actual velocity of the target. If this gap between the two gets too large, this can results in a rejected ADS-B message by the hypothesis tests.

At the initialization of flight C the estimate overshoots the ADS-B velocity of the target by around 15 m/s. After this overshoot the SIR filter is able to correct for this overshoot which shows stability. Rapid changes in velocity or location can cause the filter to diverge if the particle cloud is not diverse enough. This result thus shows that the SIR filter is capable of handling these flight dynamics. Similar to the location estimate, the SIR filter provides accurate velocity estimates of the targets, therefore the MISS filter is not required for these three flight. To finally conclude if the MISS filter is not required the SIR filter has to be applied to more flights to know with higher certainty that in fact the MISS filter is not required.

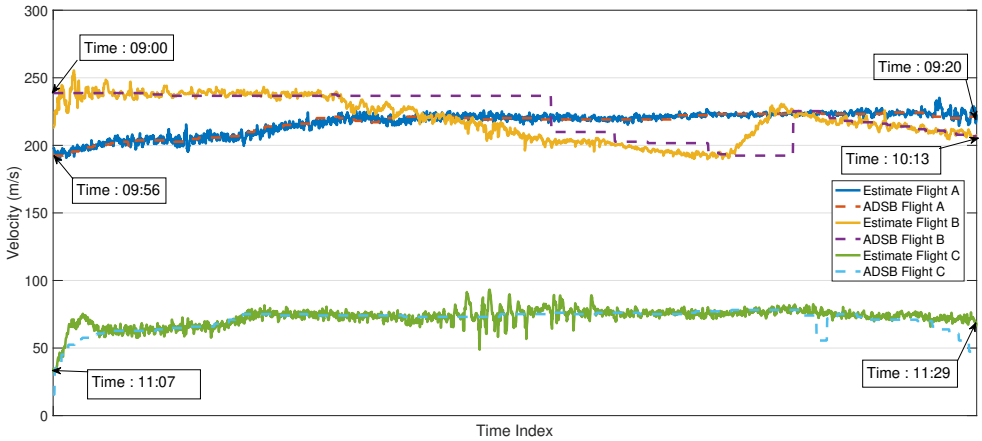

Figure 5.5: Velocity estimate and ADS-B altitude for flight A, B and C. The estimate is obtained by computing the weighted mean of the particle cloud.

Effective Sampling Size

In the analysis the concept of Effective Sampling Size (ESS) is used [43]. The concept of effective sampling size is used to assess the quality of a particle set and determine how well it represents the posterior distribution. The effective sampling size provides an estimate of the number of informative

samples in the particle set. A higher effective sampling size indicates that the particle set is more diverse and representative of the posterior distribution. Conversely, a lower effective sampling size implies that the particle cloud is dominated by a few samples with an extremely high weight, leading to a poorer representation of the posterior. When the effective sampling size is close to the total number of particles, it suggests that the weights are roughly uniform and each particle contributes equally to the approximation. On the other hand, if the effective sampling size is close to one (lowest possible value), the particle set is dominated by a few particles, indicating a degenerate set of particles. The ESS is computes according to;

$$N_{eff} = \frac{N_s}{1 + \text{Var}\left(w_k^{*i}\right)} \tag{5.1}$$

where $w_k^{*i} = p(\mathbf{x}_k^i \mid \mathbf{z}_{1:k})/q(\mathbf{x}_k^i \mid \mathbf{x}_{k-1}^i, \mathbf{z}_k)$ [43]. This expression is impossible to determine exactly, therefore a estimate of N_{eff} is used in practice,

$$\widehat{N_{eff}} = \frac{1}{\sum_{i=1}^{N_s}\left(w_k^i\right)^2} \tag{5.2}$$

Fig. 5.6 shows the ESS for all three flights. All flights have a healthy diverse particle cloud for the majority of the flight. Flights A and C have some drops is ESS which indicate a degenerate set of particles, especially flight A. This result can indicate that in this region more particles are required to represent the posterior density, or that the measurement contains errors. The filter is able to fully recover from this drop in ESS, and the ESS is generally close to the total amount of particles, the results do not indicate that more particles are required by the filter. These dips in ESS cannot be directly related to a degraded quality estimate in location, but the velocity estimate of flight C can be seen to have an increase in uncertainty around the same time that the ESS drops.

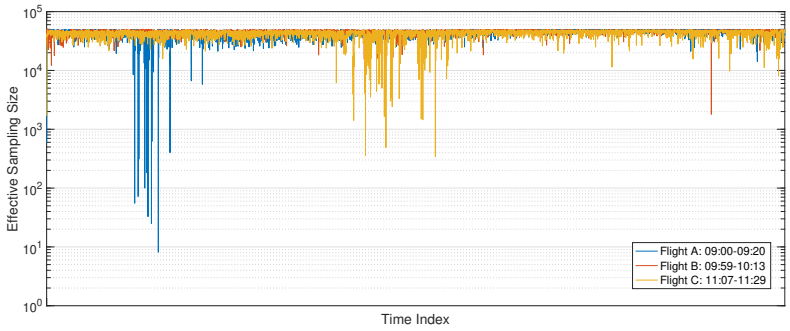

Figure 5.6: Effective Sampling Size for flight A,B and C

5.1.2. Parameter Analysis

The performance of the SIR filter is greatly determined by several important parameters. In this section the effect of the number of samples, the proposal density and the process noise is analyzed. This analysis provides insights on the impact these parameters have on the output the filter produces.

Number of Samples

The number of particles used greatly impacts the performance of the filter. A high number of particles leads to a better approximation of the state of the target. Using high amount of particles is not always an option due to the limited computational time available per iteration. Therefore, in this section the SIR filter is used to track flight B, with various amount of particles. In this analysis the ADS-B report is assumed to be the ground truth to compute a root mean square error (RMSE) that allows for comparison between the used amount of particles.

Fig. 5.7 shows the RMSE error for flight B between the location estimate by the filter and the ADS-B track for various values of N_s. From the figure it can be seen that $N_s = 50$ and $N_s = 100$ have a significant

RMSE which will likely result in the hypothesis tests to reject the ADS-B message. From $N_s = 500$ and higher the RMSE is very similar. Using more samples thus does not lead to a meaningful improvement in performance of the SIR filter.

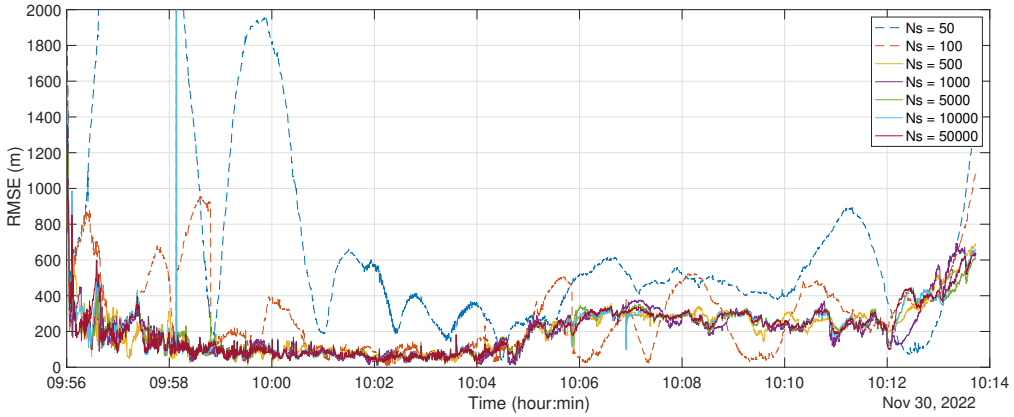

Figure 5.7: Flight B: RMSE of the SIR filter's location estimate with different values of N_s. The RMSE is computed with respect to the reported ADS-B position.

Fig. 5.8 shows the RMSE error of the velocity estimate with respect to the reported ADS-B velocity for different values of N_s. Fig. 5.7 has shown that using more samples than 500 does not lead to a meaningful decrease in RMSE, the same is seen in the RMSE for the velocity estimate.

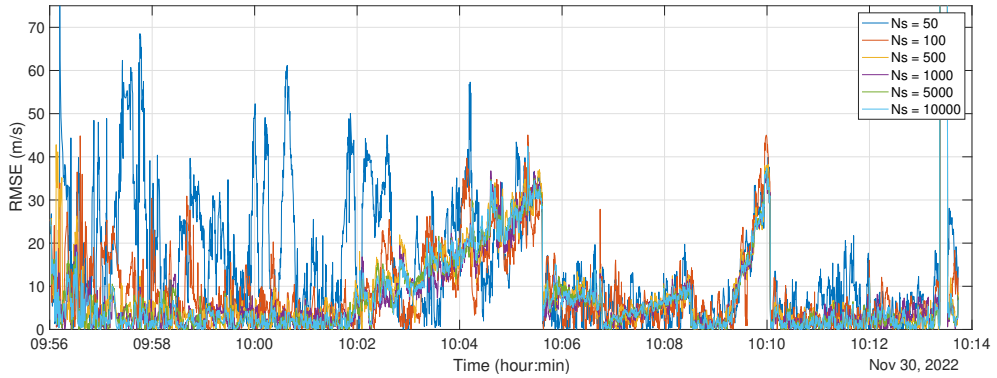

Figure 5.8: Flight B: RMSE of the SIR filter's velocity estimate for different values of N_s. The RMSE is computed with respect to the reported ADS-B velocity

The initial value for the number of particles needed was guessed to be $N_s = 50000$. A relative high number was expected to be required because the filter has to be able to handle ambiguous measurements. The ambiguous measurements require lots of particles to represent the posterior density with good accuracy. These results show that a considerable less amounts of particles are needed to achieve the same RMSE. Using $N_s = 1000$ results in the SIR filter achieving a good accuracy, and the low amount of particles that the filter requires lowers the computational requirement of the filter. In on-line tracking the drawback of the particle filter in general is the amount of particles requires which results in longer computational time when compared to a Kalman or Extended Kalman filter. The choice for the proposal density to sample directly from the measurements allows for the filter to work with such a low amount of particles, without this initialization, the filter becomes unfeasible. The approach also results in a high ESS after the resampling step and limits the impact of the degeneracy and sample impoverishment problem.

The impact of the amount of particles used greatly effects the ESS. The ESS is optimal if it is equal to the total amount of particles. In the case of $N_s = 50, 100$ the drops in relative ESS are larger compared to

the other values shown in Fig. 5.9 which indicate a degenerate set of particles at that iteration, and that this amount of samples is too low. In combination with RMSE from the location estimate is can be seen that this degenerate set of particles goes hand in hand with the RMSE. Low ESS values occurs at the same time a high RMSE is seen. After several minutes the filter's ESS is more stable, which shows that the filter has not lost the track.

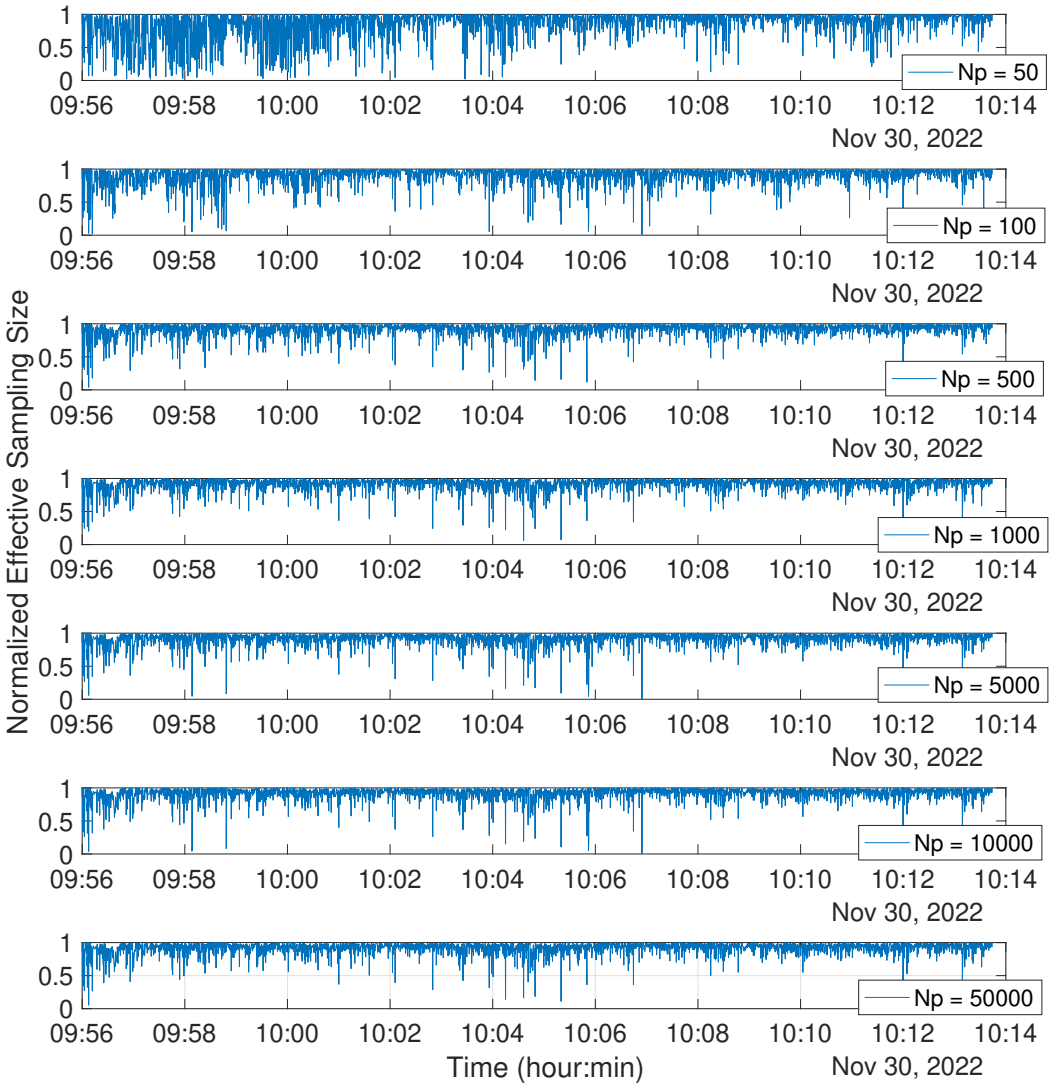

Figure 5.9: Flight B: Relative Effective Sampling size w.r.t total number of samples for different values for N_p

The SIR filter is able to track the target properly using around 1000 particles. The ESS shows that using more particles achieves no substantial increase in performance. The drops in ESS are present at all values which indicates that the drops are the result of faulty or noisy measurements. More importantly, is shows that a higher amount of samples does not mitigate the problem.

Proposal Distribution

The filter is initiated using the proposal density, this density samples directly from the first measurement. This is achieved by using a search area from where the samples are drawn. If the measurement equation

has a solution inside the search area samples are drawn. In case of in-flight spoofing, a scenario where an airborne target alters is ADS-B location, it occurs that the spoofed ASD-B location is inside or outside this search area. In the case that the measurement equation corresponding to the MLAT measurements (from the spoofed ADS-B transmitter) provides no solution around the location of the received ADS-B message, meaning that the transmitter of the spoofed ADS-B messages is outside the search area, no samples are drawn and the filter concludes that the ADS-B message is spoofed. If it occurs that the location of the transmitter that is transmitting the spoofed ADS-B messages is inside the search area, this transmitter can be located whilst concluding that the ADS-B message is spoofed, then safety measures can be taken to mitigate the spoofing. When the search area is large, the probability of finding the location of the transmitter that is transmitting the spoofed ADS-B messages increases. From a spoofing detection point of view it is thus preferable to have a large as possible search area. To sample from this search area the particle density must remain equal to obtain the same performance.

In the previous subsection results were obtained by sampling around the ADS-B location using a search grid of 2km $\times$ 2km, which translates to $2.5 * 10^{-4}$ samples per square meter. The effect of the proposal density is investigated by varying the number of samples per unit area. Fig. 5.10 shows the RMSE with respect to the ADS-B location for varying sample densities. Here it can be seen that for a large grid size the RMSE is large for the first sequence of measurements. After some iterations the RMSE decreases and the filter accurately estimates the location of the target under test. The increased grid size causes the filter to require more measurements to accurately estimate the location of the target, but for the grid sizes shown in the figure, the SIR filter does not diverge.

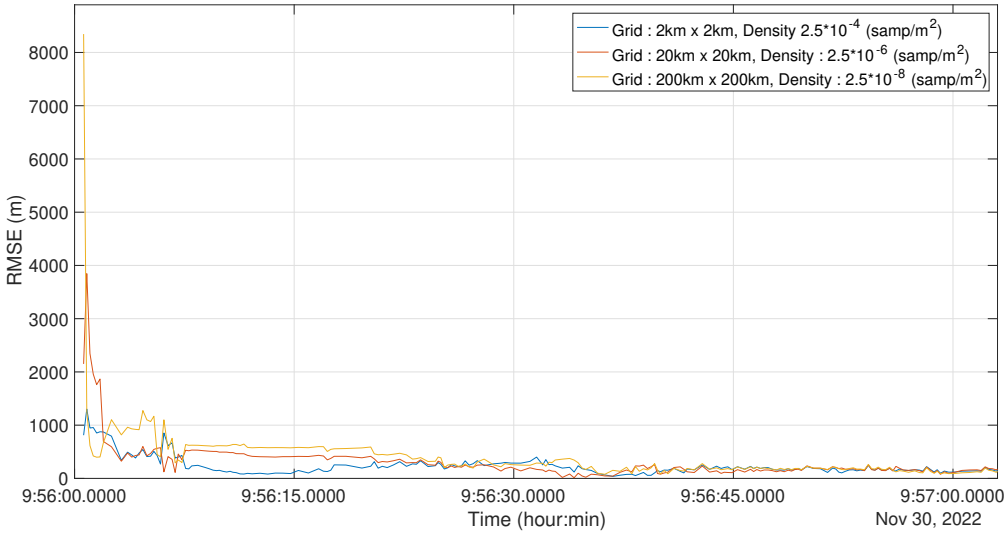

Figure 5.10: Flight B: RMSE of location estimate for various proposal grid densities

Fig. 5.11 shows realization of the proposal density for the different grid sizes that are analyzed in Fig. 5.10. It can be seen that in Fig. 5.11a the search area is very large, but at the cost that the search area is populated by less samples per unit area compared to Fig. 5.11b and 5.11c. It can also be seen that the sampled hyperbola is constant in height, this introduces a disadvantage for the proposed sampling method.

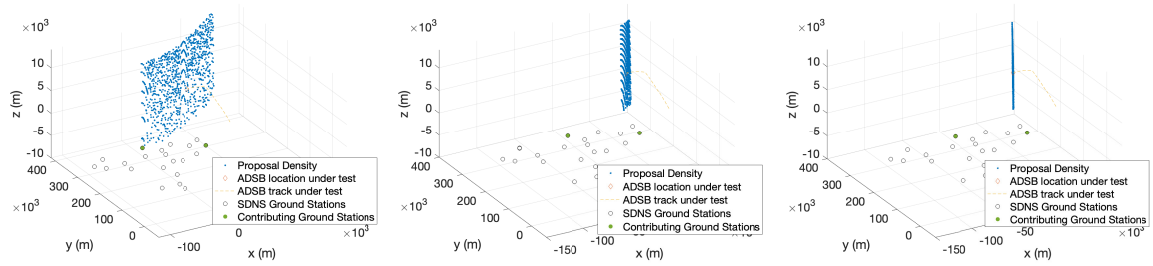

(a) Grid size 200km × 200km **(b)** Grid size 20km × 20km **(c)** Grid size 2km × 2km

Figure 5.11: Flight B: Three different proposal distribution sizes.

When a lower amount of samples per unit area is used the sampling method that draws samples from the measurement fails. The hyperbolas representing the measurement space are sampled using a linear mesh grid in the horizontal plane. If none of the coordinates in the grid intersect the hyperbolas no samples can be drawn. In addition, the general shape of these hyperbolas is vertical and has very little spread in the horizontal plane. Therefore, it occurs when the sampling density becomes too small no samples can be drawn using this sampling method. Although, such a low sample density is detrimental to the performance of the SIR filter which is not preferred and the disadvantage of the sampling method is thus of limited impact.

Eventually it is expected that when the number of samples per unit area becomes too small the filter will not converge on the true location of the target. This because the proposal density is not represented by enough particles such that the degeneracy problem will not effect the state estimate. Similar to previous results such as the location and velocity estimate, the SIR filter performs better than was initially expected. A higher number of samples per unit area helps in the rate of convergence to the true state of the target, but the SIR filter is still able to converge on the location using a low sample density. ($2.5 * 10^{-8}$ samples / m^2)

Comparison w.r.t to traditional SIR filter

The proposed SIR filter samples directly from the measurement at the first iteration, and it has been shown that the filter is able to correctly track the target using around 500 samples, although more samples are preferable for a stable filter. In the traditional SIR filter the proposal density is uniformly distributed in some area where the target could be. Fig 5.12 shows the proposal density in such a filter implementation. In this scenario 2000 particles are used and the ESS after one iteration is $ESS = 1.003$, and the state estimate is wrong as seen in Fig. 5.12b. The traditional SIR filter thus completely fails for this amount of particles.

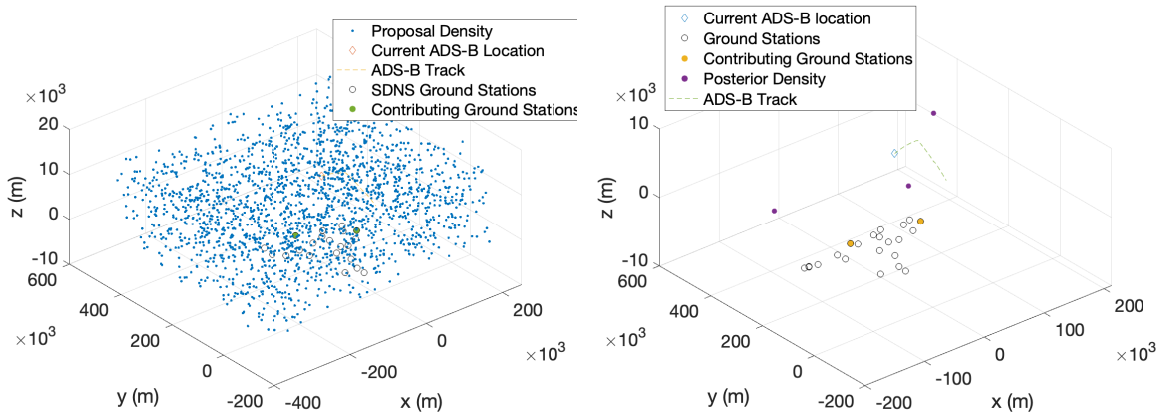

(a) Proposal density as used in a traditional SIR filter where the density is uniformly distributed across the area where the target could be.

(b) Particle distribution after first measurement update. The filter completely diverges and $ESS = 1.003$.

Figure 5.12: Flight B: First filter iteration where the proposal density is that of the traditional SIR filter, where the samples are uniformly distributed across *all* possible target location.

Using an uniform proposal density around 1 million particles are required by the traditional SIR filter to converge on a correct estimation of the location and velocity of the target. Then after the first update the ESS is equal to 757.2. Using the proposed proposal density the filter can be initialized using 500 times less particles compared to the traditional SIR filter, and simultaneously achieve a higher ESS after the initial update. When 2000 samples are used $ESS = 1999.601$, which should be very close the the total amount off samples because the weights are assigned using the same MLAT measurement from which the samples are drawn.

Process Noise

In the SIR filter the process noise must be defined. This process noise quantifies the expected uncertainty in the prediction the filter does at every iteration. In Fig. 5.13 and Fig. 5.14 several process noise configurations the RMSE is shown. This RMSE is computed with respect to the ADS-B track. Tab. 5.1 shows the used noise settings. The figures show that for a large range of noise values the SIR filter converges and in configuration 3 and 4 problems arise. Configuration 3 has a high RMSE at the start of the track but eventually converges, configuration 4 diverges entirely. These results show that when the standard deviation of the process noise becomes too large the filter becomes unstable and can diverge. Noise settings below those of configuration 2 achieve good performance and good stability.

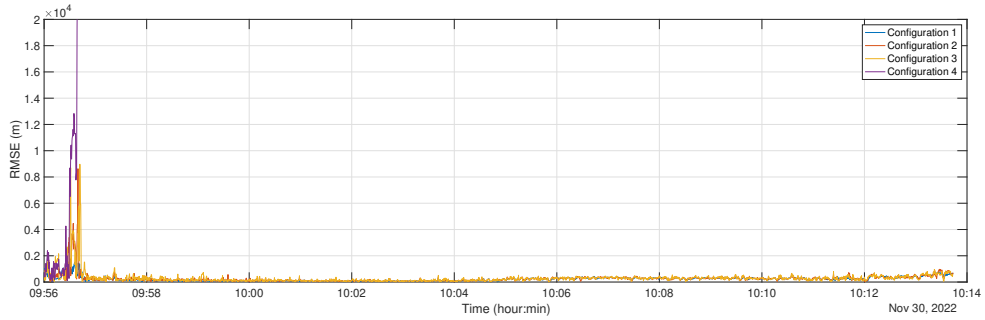

Figure 5.13: Flight B: RMSE error with respect to the ADS-B location for different process noise configurations. Configuration is shown in Tab. 5.1.

The impact of the process noise is less in flight C than in flight B. In flight C the SIR filter manages to converge but configuration 4 does have a significant RMSE during the start-up phase. The main difference between the two flight is that flight B is a fixed-wing aircraft, and flight C a helicopter. Large process noise allows the filter to be better at tracking targets who execute more rapid turns and velocity changes.

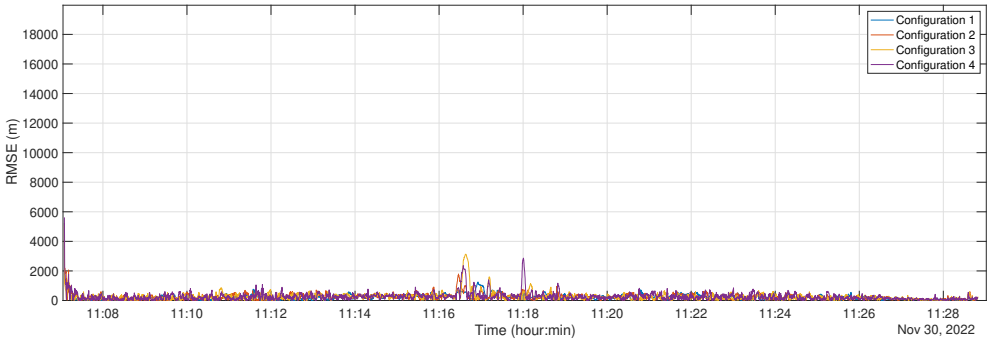

Figure 5.14: Flight C: RMSE error with respect to the ADS-B location for different process noise configurations. Configuration is shown in Tab. 5.1.

Configuration 1	x	y	z
std. position	200	200	100
std. velocity	100	100	50
Configuration 2			
std. position	400	400	300
std. velocity	300	300	200
Configuration 3			
std. position	600	600	500
std. velocity	500	500	400
Configuration 4			
std. position	800	800	700
std. velocity	700	700	600

Table 5.1: Process noise configuration for the results shown in Fig. 5.10.

In the process model used in the SIR filter the acceleration of the target is modeled as additive noise. Therefore the standard deviation is equal to the acceleration by the target. The investigated values thus represent acceleration values target experiences. In configuration 1 the target already experiences G-forces around 20G. Such values are extremely high and not realizable by any target. The SIR filter is thus capable of tracking targets that experience a acceleration up to around 600 m/s^2. Such targets are never tracked in real life but illustrate the capabilities of the filter.

5.2. Hypothesis Test Analysis

In this section the results of the three proposed hypothesis tests are analyzed. In addition to this the Gaussian assumption made in the Neyman-Pearson and Minimax test is investigated, impact of the false alarm rate is investigated and the impact of the assigned cost in the MBR and Minimax hypothesis test.

5.2.1. Gaussian Assumption

In the Minimax and Neyman-Pearson hypothesis test the output of the particle filter is assumed to be Gaussian distributed to compute the threshold. If the assumption is accurate the tests achieves the intended performance. Fig. 5.15a and Fig. 5.15b shows the distribution of each dimension of the state vector at $k = 1,1000$ for flight A. At the first iteration a single measurement is used to update the proposal density. This measurement contains four GSs and thus is not ambiguous in location. From Fig. 5.15a it can be seen that the densities are not smooth estimates, but are somewhat close to a Gaussian distribution. Eventually, as shown in Fig. 5.15b at the 1000th iteration all dimensions of the state vector have a shape that resembles Gaussian distribution. The assumption thus becomes increasingly accurate after several iterations by the filter.

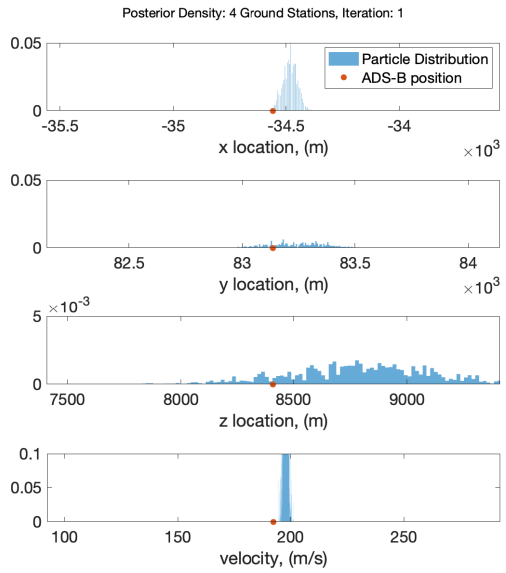

(a) Flight A: Particle distribution and ADS-B report, iteration 1

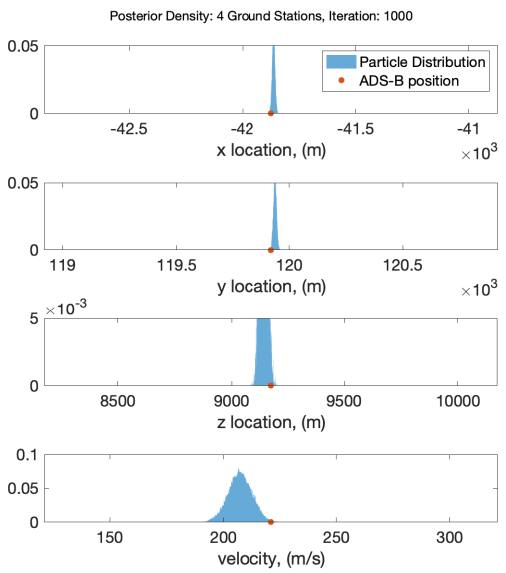

(b) Flight A: Particle distribution and ADS-B report, iteration 1000

Figure 5.15: Flight B: Particle distribution for iteration 1 and 1000. After a considerable amount of iterations the distribution of the particle cloud seem to be Gaussian distributed.

At the initialization of flight A there are only MLAT measurement with more than four GSs per MLAT measurement. To investigate the Gaussian approximation under ambiguous measurements, the algorithm is initialized on a section of flight A that only contains ambiguous measurements. The first iteration of flight A^* can be seen in Fig. 5.16a. In this figure the distribution is close to a uniform distribution, apart from the velocity, which is obtained from the ADS-B message itself. After several iterations as seen in Fig. 5.16b, the distribution concentrates, but there is still a large uncertainty when compared to the an estimate with a non-ambiguous measurement.

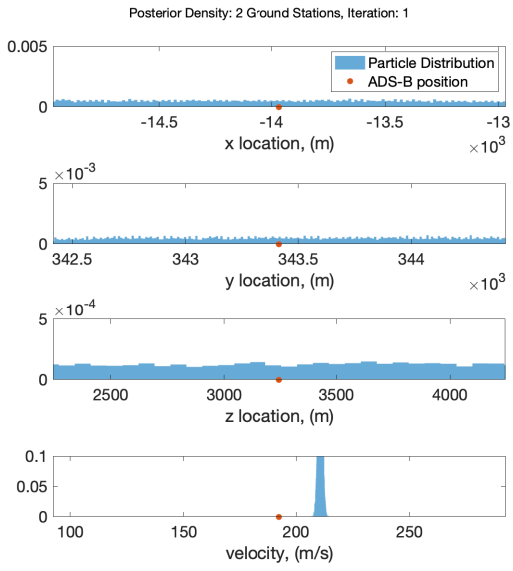

(a) Flight A^*: Particle Distribution and ADS-B report, filter iteration 1

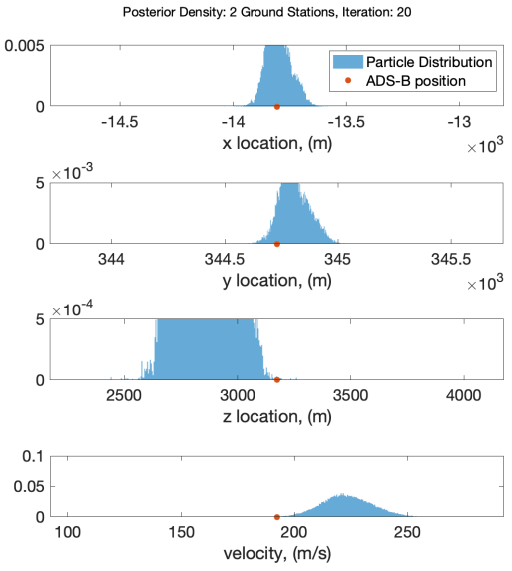

(b) Flight A^*: Particle Distribution and ADS-B report, filter iteration 20

Figure 5.16: Flight A^*: Particle distribution for filter iteration 1 and 20.

The Gaussian assumption on the posterior density is an accurate assumption if some of the measurements are not ambiguous. When ambiguous measurements follow directly after non-ambiguous measurements, the filter is able to hold on to the *Gaussian* distribution. In all three analyzed flights there is always some non-ambiguous measurement thus the probability of only ambiguous measurements for an extended period of time has a very low probability. If this is not the case the densities are close to the densities seen in Fig. 5.16b.

5.2.2. Flight A

Fig 5.17 shows the result of the MBR hypothesis test, Fig. 5.18 the results of the Neyman-Pearson test and Fig. 5.19 the results of the Minimax hypothesis test. All hypothesis tests have determined that all ADS-B messages received are not spoofed, as is the case.

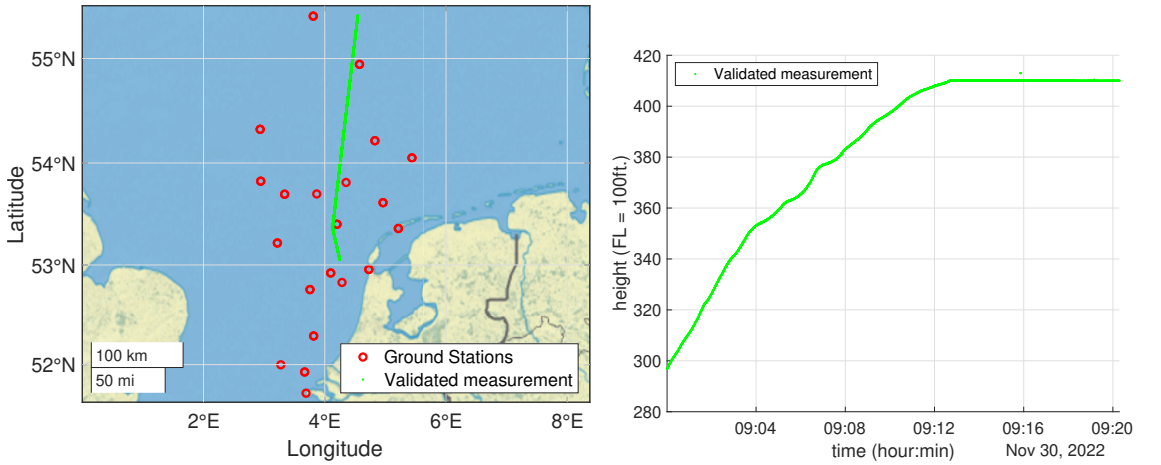

(a) Flight A: Validation Result Minimum Bayes Risk: 100% of messages validated

(b) Flight A: Height Plot, Validation Result MBR: 100% of messages validated

Figure 5.17: Result of the Minimum Bayes Risk for flight A. The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

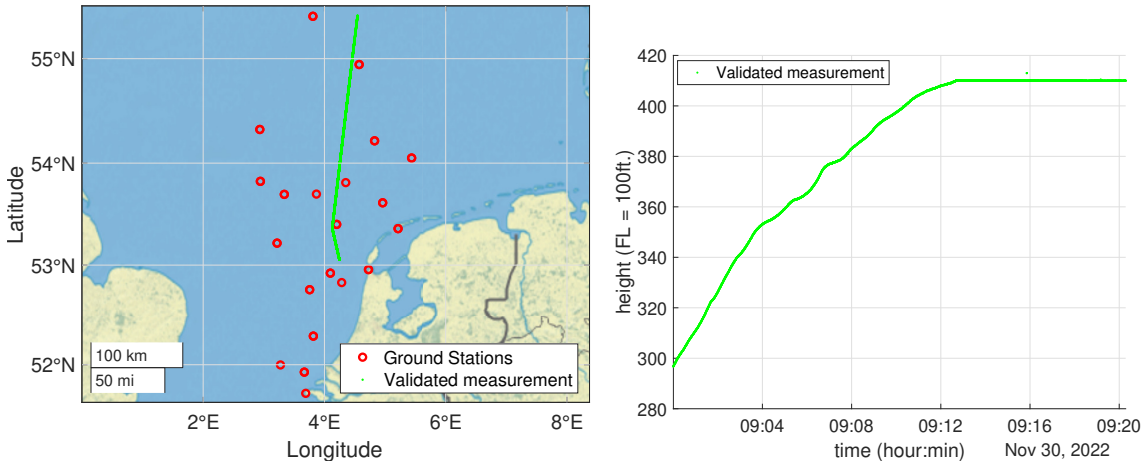

(a) Flight A: Validation Result Neyman-Pearson: 100% of messages validated

(b) Flight A: Height Plot, Validation Result NP: 100% of messages validated

Figure 5.18: Result of Neyman-Pearson for flight A. The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

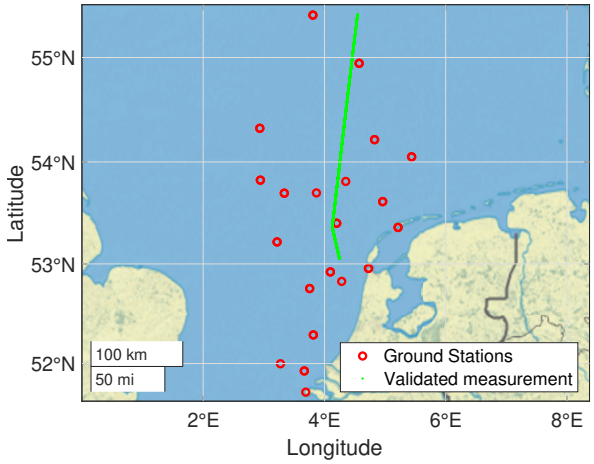

(a) Flight A: Validation Result Minimax: 100% of messages validated

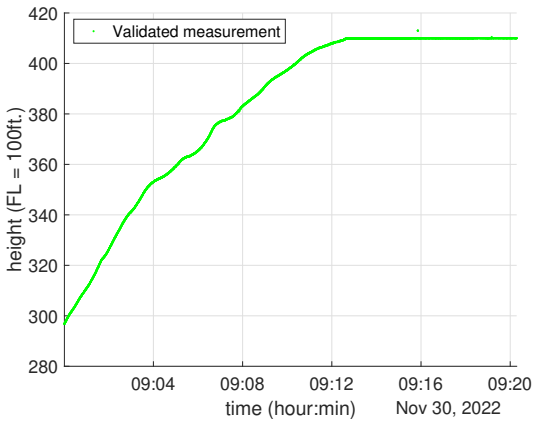

(b) Flight A: Height Plot, Validation Result MM: 100% of messages validated

Figure 5.19: Result of Minimax for flight A. The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

Hypothesis Test

Fig. 5.20a plots the threshold values for all three hypothesis tests over all measurements. The Neyman-Pearson and Minimax threshold have the same general shape, because in both hypothesis tests the threshold's only dynamic variable is the covariance matrix S. From Fig. 5.20a one can also clearly observe the impact that different assumptions on the data or statistics have on the final threshold value. Based on the implemented value of $P_{fa} = 10^{-6}$, the Neyman-Pearson test results in a significantly higher threshold value than the Minimax's threshold. The MBR threshold is constant for all iterations as expected. In a spoofing scenario, the hypothesis test that has the highest threshold value will be the first one to detect spoofing. Therefore the Neyman-Pearson test has the highest spoofing detection sensitivity.

A strong correlation between the threshold value for the Neyman-Pearson and Minimax test can be seen when it is compared to the value of the likelihood ratio. When the likelihood ratio drops, both threshold values also drop. Both are direct result of a increasing uncertainty in the estimate made by the particle filter. As the uncertainty in the estimate grows, the covariance matrix S describing the uncertainty in $p(\mathbf{x}, \mathcal{H}_0)$ also grows. Thus the likelihood and the threshold values simultaneously drop.

Fig. 5.20b plots the values computed by the LRT. In the starting phase of the filter the LRT quickly climbs to around 10^{13}, indicating a very high certainty in the no-spoofing scenario. Eventually the LR value decreases, around the same time the amount of GSs that contribute to the MLAT measurements also drops.

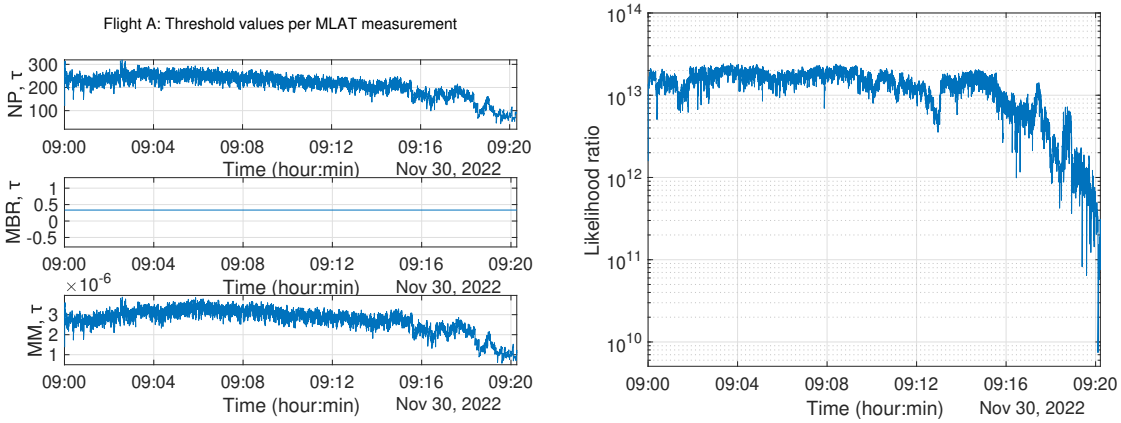

(a) Flight A: Threshold values

(b) Flight A: Likelihood Ratio per MLAT measurement

Figure 5.20: Flight A: Threshold values and Likelihood Ratio

Fig. 5.21 shows the distribution of the amount of GSs per MLAT measurement in time. Each dot represents the amount of GSs that contribute to the measurement at that time index. When the ADS-B messages are detected by less GSs, the LRT drops because when the measurements is ambiguous in location, the uncertainty of the state estimate grows. This drop is only observed when the number of GSs drops below four, because then the measurements are ambiguous. The lowest values of the LRT are still well above the threshold determined by all the three hypothesis tests. Therefore this amount of ambiguous measurements has no meaningful impact on the validation performance.

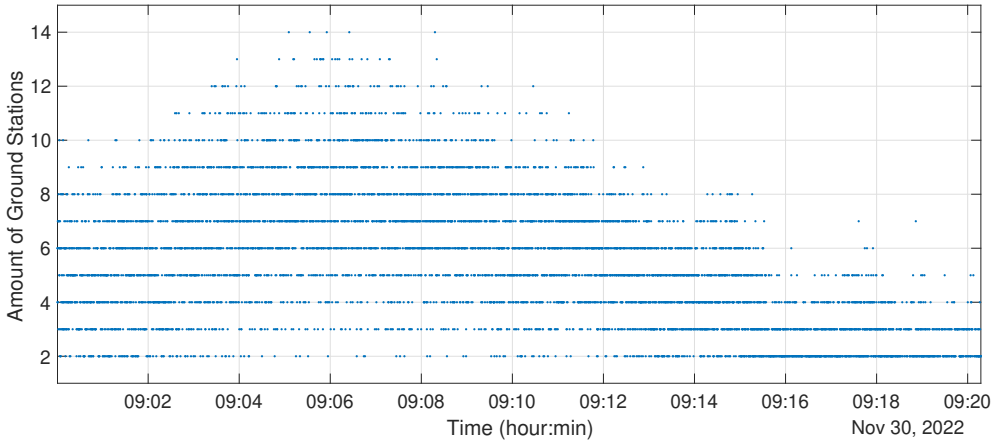

Figure 5.21: Flight A: Number of ground station per MLAT measurement

5.2.3. Flight B

Fig. 5.22, Fig. 5.23 and Fig. 5.24 show the result of the MBR, Neyman-pearson and the Minimax hypothesis tests respectively. Similar to the results of flight A, the ADS-B plot is green if validated, red if spoofed. Again all hypothesis tests classify the messages as valid. The effect of VDOP at this altitude, and a higher amount of ambiguous measurements has no effect on the output of the validation algorithm.

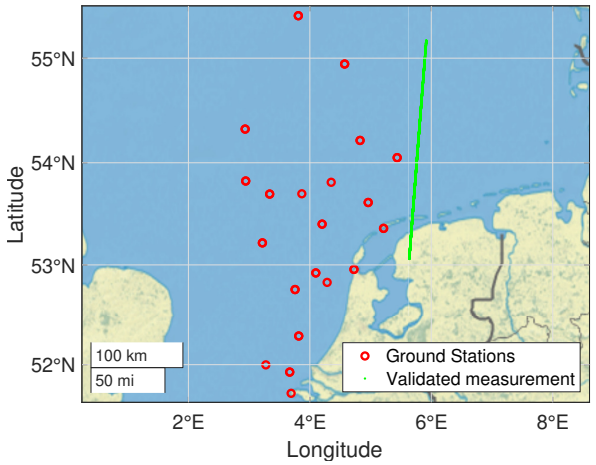

(a) Flight B: Validation Result Minimum Bayes Risk: 100% of messages validated

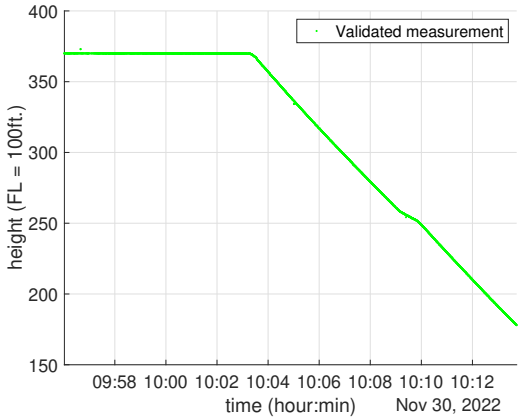

(b) Flight B: Height Plot, Validation Result MBR: 100% of messages validated

Figure 5.22: Result of Minimum Bayes Risk for flight B. The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

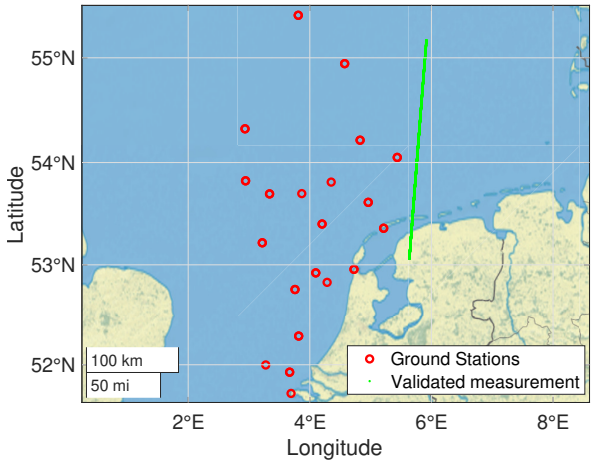

(a) Flight B: Validation Result Neyman-Pearson: 100% of messages validated

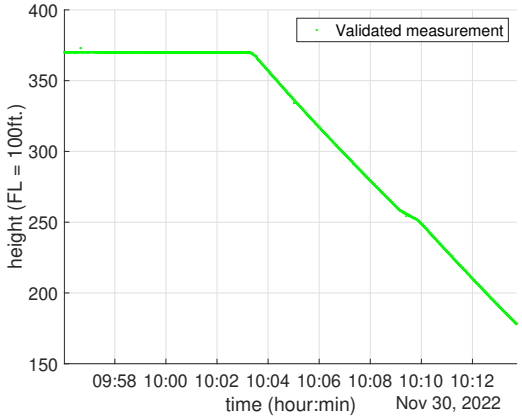

(b) Flight B: Height Plot, Validation Result NP: 100% of messages validated

Figure 5.23: Result of Neyman-Pearson for flight B. The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

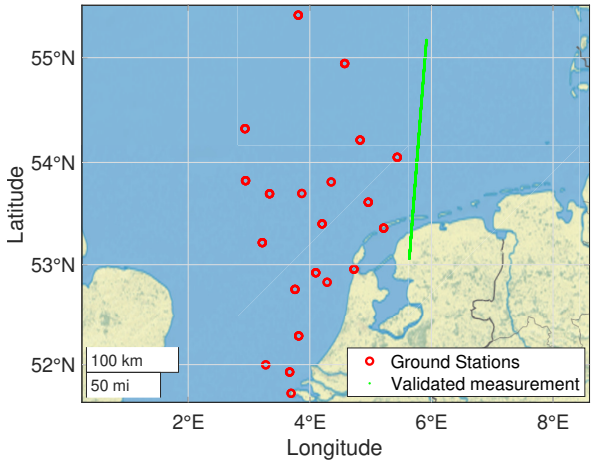

(a) Flight B: Validation Result Minimax: 100% of messages validated

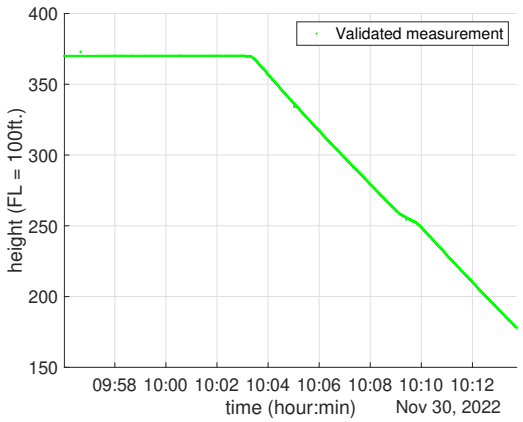

(b) Flight B: Height Plot, Validation Result MM: 100% of messages validated

Figure 5.24: Result of Minimax for flight B: The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

Hypothesis Test

The limited amount of GSs in this scenario has no effect on the output of the validation algorithm. Fig. 5.25b plots the values computed by the LRT. In the initialization phase of the filter the LR quickly climbs to around 10^{12} indicating a very high certainty in the no-spoofing scenario. Eventually the LR starts to drop around the time the aircraft start to descend.

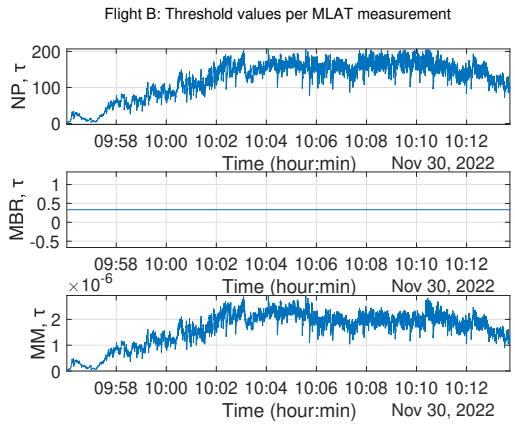

(a) Flight B: Threshold values for the three detectors

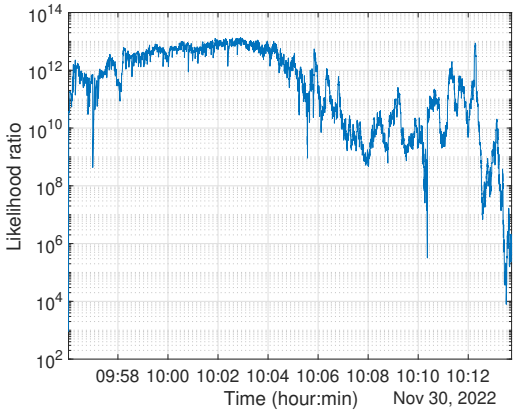

(b) Flight B: Likelihood Ratio per MLAT measurement

Figure 5.25: Flight B: Threshold values and Likelihood Ratio

Fig. 5.26 shows the number of GSs per measurement through time. It can be observed that the value for the LR very roughly correlates with the amount of GSs that contributes to each MLAT measurement. Around 10:06 there is a drop in the number of GSs per measurement, and around the same time the LRT lowers. The same effect can be observed around 10:13 where a rise in number of GSs per measurement occurs at the same time that the LR increases.

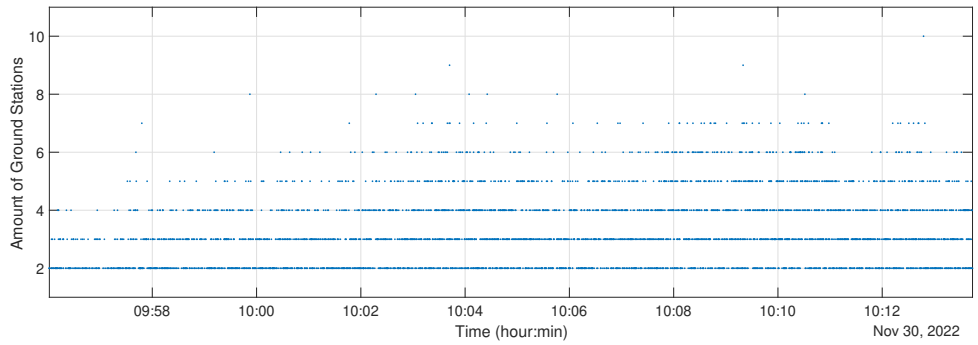

Figure 5.26: Flight B: Number of ground station per MLAT measurement

5.2.4. Flight C

For flight C not all ADS-B messages are validated, although by a small amount. The Neyman-Pearson test has validated 95.27% of all messages, the MBR 96.04%, and the Minimax test 97.07%. The performance of the algorithm decreases at low altitude. Fig 5.27, Fig. 5.28 and Fig. 5.29 show that the all three hypothesis tests reject the first sequence of measurements. In this period the filter is starting and its needs several measurements to accurately determine the state of the target. After this period MBR and Neyman-Pearson reject a sequence of ADS-B messages around 11:17.

In flight C it occurs that not only the filter incorrectly estimates the state of the target, but also the ADS-B location contains errors. These errors occur when the target flies at FL26, here individual messages are out of line with the general track. All three hypothesis test include the statistical properties of the ADS-B distribution and the estimated distribution by the filter, thus the tests are still able to validate the noisy location reports. Depending on the severity of the error ADS-B makes, it can still occur that the three detectors can conclude that the message is spoofed.

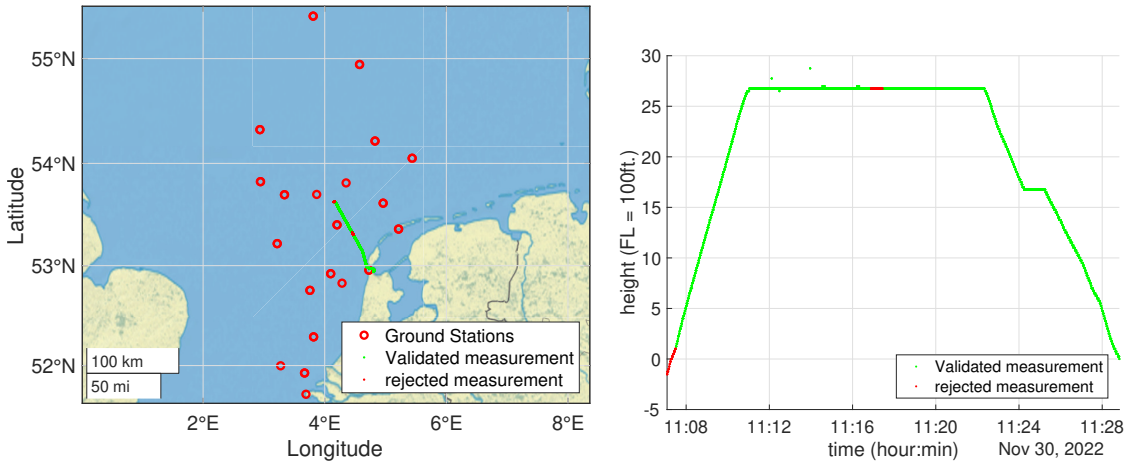

(a) Flight C: Validation Result Minimum Bayes Risk: 96.0375% of messages validated

(b) Flight C: Height Plot, Validation Result MBR: 96.0375% of messages validated

Figure 5.27: Result of Minimum Bayes Risk for flight C : The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

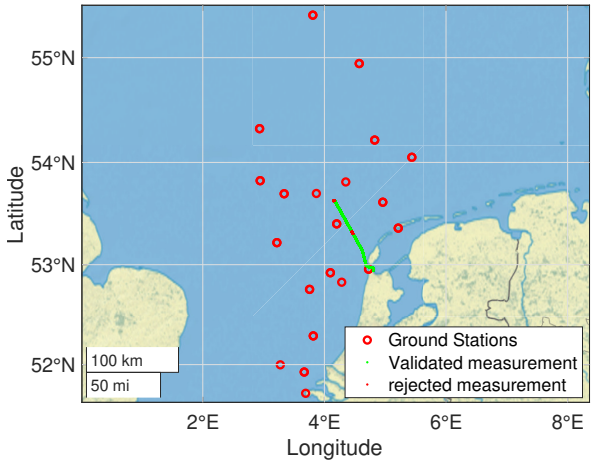

(a) Flight C: Validation Result Neyman-Pearson: 95.2738% of messages validated

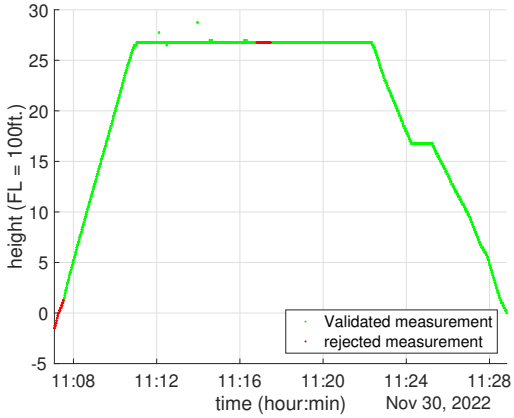

(b) Flight C: Height Plot, Validation Result NP: 95.2738% of messages validated

Figure 5.28: Result of Neyman-Pearson for flight C The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

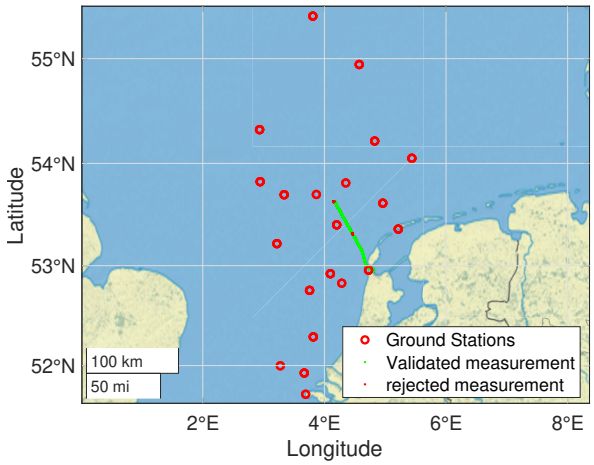

(a) Flight C: Validation Result Minimax: 97.0749% of messages validated

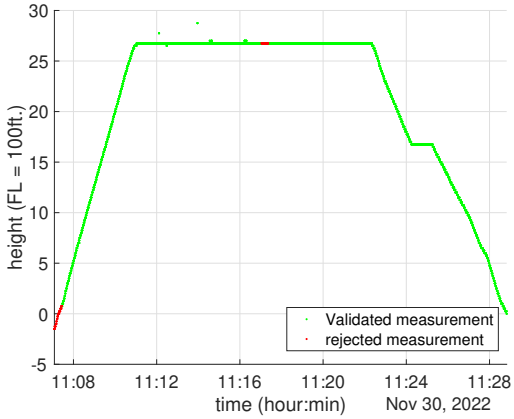

(b) Flight C: Height Plot, Validation Result MM: 97.0749% of messages validated

Figure 5.29: Result of Minimax for flight C The ADS-B message is plotted in green if the message is validated by the corresponding hypothesis test, and red if the message is considered to be spoofed.

Hypothesis Test

Fig. 5.30b shows that in the initialization phase of the filter the likelihood ratio climbs to around 10^{12}. It needs around half a minute to achieve this, which, meanwhile, results in ADS-B messages being determined as spoofed. At 10:17 there is a clear drop in likelihood resulting in messages begin labeled as spoofed.

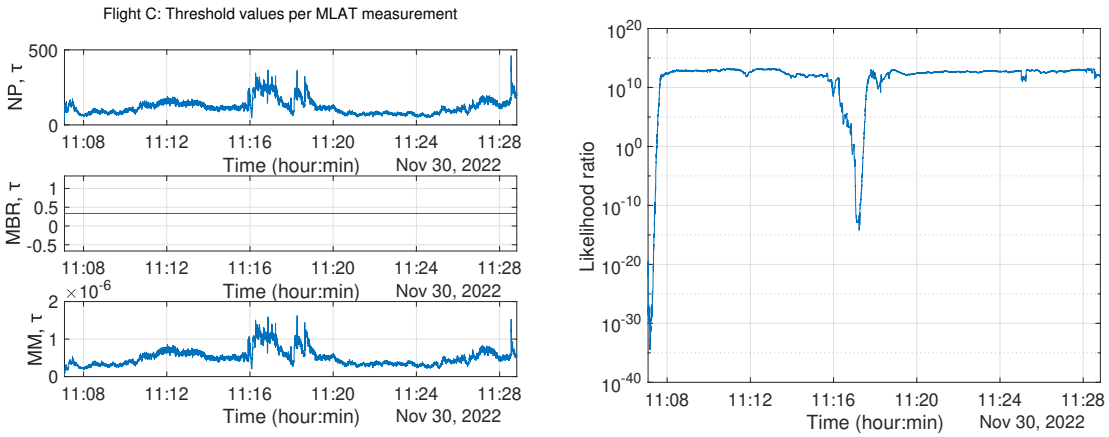

(a) Flight C: Threshold values for the three detectors

(b) Flight C: Likelihood Ratio per MLAT measurement

Figure 5.30: Flight C: Threshold values and Likelihood Ratio

Fig. 5.31 shows the distribution of a message that is considered to be spoofed by all three hypothesis tests at 11:17. From the figure can be seen that the x location , y location and velocity is correctly estimated by the SIR filter, but the z location is off by around 1.2 kilometer. That only the height of the target is incorrectly estimated clearly indicates that at this altitude the effect of VDOP is the cause of the failed validation.

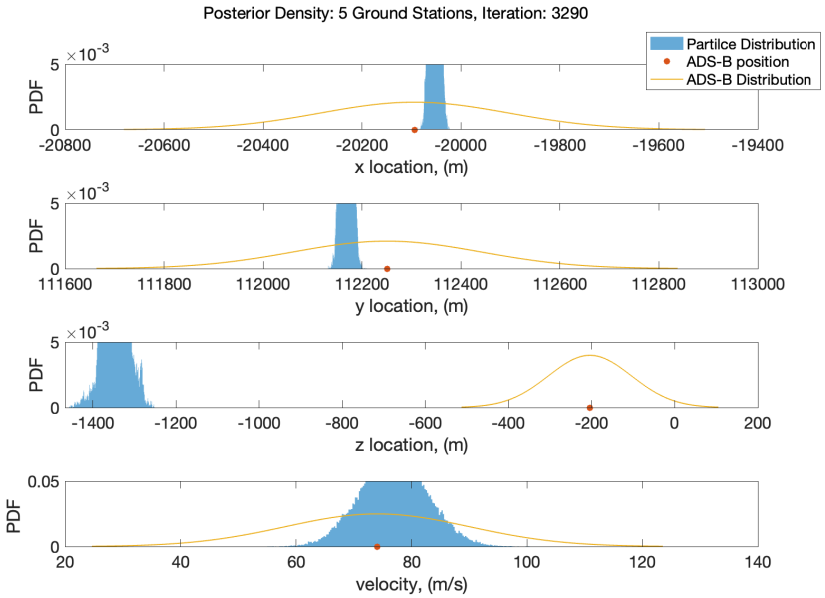

Figure 5.31: Flight C: Distribution of a rejected ADS-B message at 11:17:09 Note: a negative height in this coordinate system means is it below the horizon at Schiphol Airport.

Fig. 5.32 shows the amount of GSs per MLAT measurements in time. Around 11:17 some ADS-B messages are rejected. Previous results show that the likelihood drops when the number of GSs per measurement decreases. In this scenario such effect cannot be seen. In Fig. 5.32 there is no clear drop in GSs per measurement around 11:17 the ADS-B messages are rejected. The cause of the drop in likelihood ratio is due to VDOP as is illustrated by Fig 5.31.

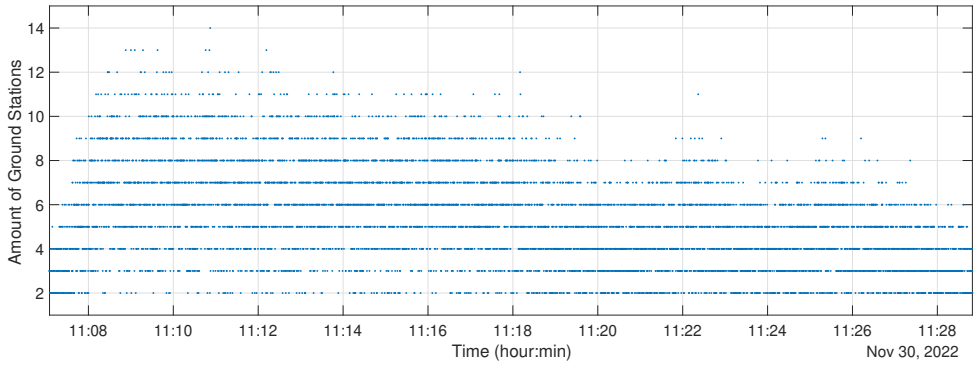

Figure 5.32: Flight C: Number of ground station per MLAT measurement

5.2.5. Parameter Analysis

In all hypothesis test there is some assumptions made on the data or statistics. In this subsection the impact of these variables is analyzed for each hypothesis test.

Neyman-Pearson False Alarm Rate

In the Neyman-Pearson test the threshold is obtained by a pre-selected value for the false alarm rate. As this value is some preferred value, the impact of several values for P_{fa} are investigated and shown in Fig. 5.33 for flight B. The realized values obtained for the false alarm rate are not equal to the theoretically expected values. This is the result of numerous assumptions and approximations made in the design of the filter. The most obvious explanation for the false alarm rate not to be achieved is because this false alarm rate depends on the Gaussian assumption on the output of the SIR filter. The posterior density must be perfectly Gaussian distributed in all 4 dimensions of the state vector to achieve the intended false alarm rate. Although it has been shown that this assumption in general is a relatively good fit, is it not good enough to actually achieve the intended false alarm rate.

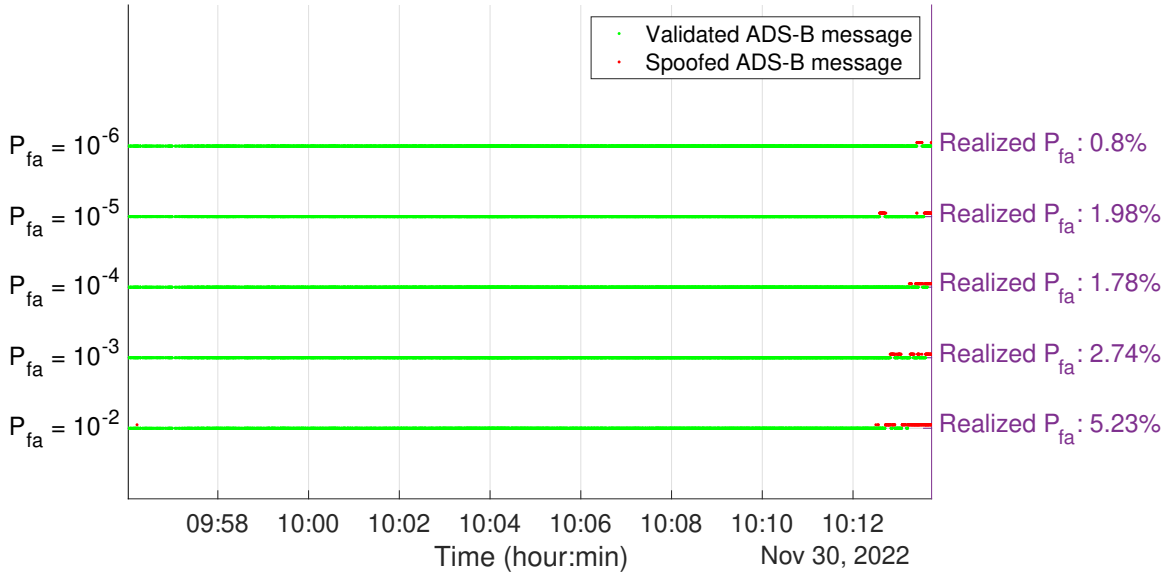

Figure 5.33: Flight B: Validation results for different values for P_{fa}.

Prior Probability for Minimum Bayes Risk

In the Minimum Bayes Risk hypothesis test a value for π_0 is determined. In Fig. 5.34 flight B is validated several values for π_0. The threshold is computed by $\tau = (1 - \pi_0)/\pi_0$, because of this the values for τ for the different values of π_0 are relatively close together, and the results of the validation are quite similar,

compared to the values computed by the LRT. The difference in the result can be explained by the fact that the particle filter is a Sequential Monte Carlo estimation method. This allows for the possibility that each iteration of the filter provides a different realization of the state estimate. Because the threshold of the MBR hypothesis test remains relatively constant the output can be influenced by the stochastic element in the particle filter. This causes the validation result to be different for the analyzed values of π_0.

From the results there appears no best value for π_0, because all values analyzed provide on average the same result. The MBR provides similar performance to the Neyman-Pearson test with a P_{fa}equal to 10^{-6}, the MBR is then preferable due to its simplicity.

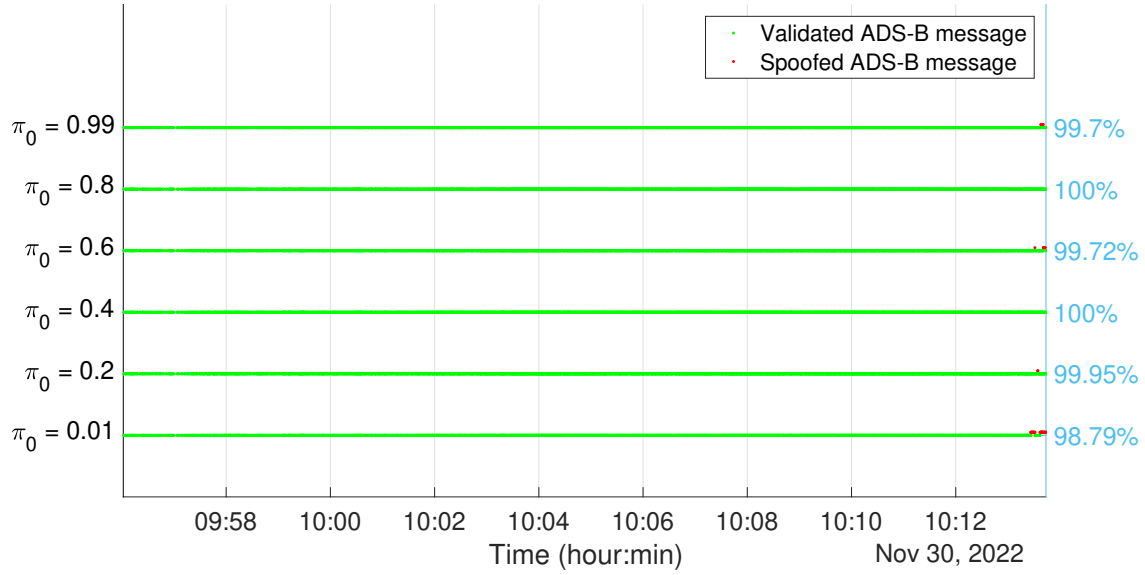

Figure 5.34: Flight B: Validation results for different values for π_0. Left y column represents the set value for π_0 and the right hand column the percentage of validated measurements in the sequence.

Cost Tuning for Minimax

Fig. 5.35 and 5.36 show the threshold values for the Minimax hypothesis test for flights B and C, for three different cost configurations. In blue the cost are set equal to one another, namely $C10 = C01 = 1$, where $C10$ is the cost for a missed detection, and $C01$ the cost for a false alarm. When $C10 = C01 = 1$, the threshold values are very similar to the values where the cost for a false alarm is equal to $C10 = 10^9$. When the cost for a missed detection is high, the corresponding threshold increases as shown in yellow in Fig. 5.35 and 5.36, as a higher threshold leads to more detections it also leads to less missed detections. Conversely, a high cost for the false alarm probability should result in less false alarm, which is achieved by a low threshold.

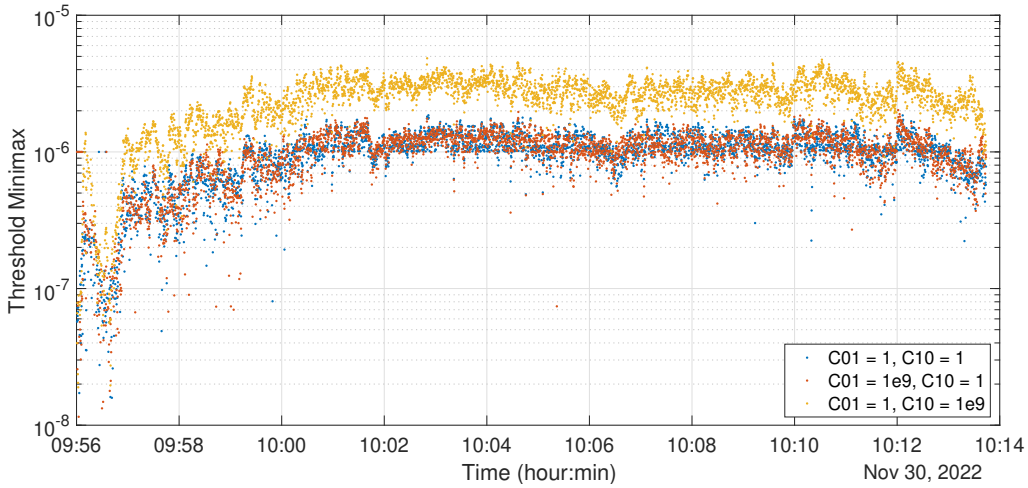

Figure 5.35: Flight B: $C10$ is equal to the cost of a missed detection, $C01$ is equal to the cost of a false alarm. High cost for a missed detection lead to a high threshold. A high cost for a false alarm leads to a threshold that is very similar to equal costs.

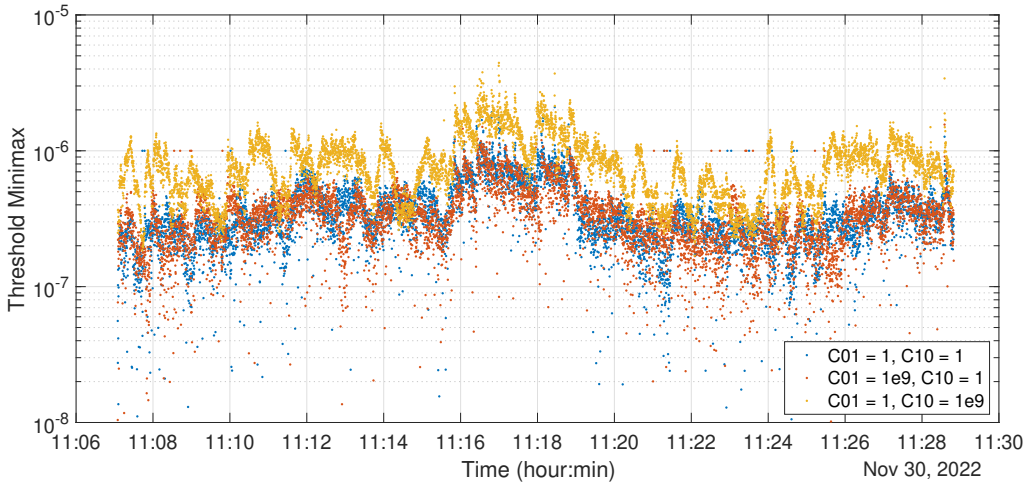

Figure 5.36: Flight C: $C10$ is equal to the cost of a missed detection, $C01$ is equal to the cost of a false alarm. High cost for a missed detection lead to a high threshold. A high cost for a false alarm leads to a threshold that is very similar to equal costs.

5.3. Spoofing Analysis

In this section the performance of the algorithm is investigated when spoofed ADS-B messages are received. Due to the nature of spoofing there are numerous different spoofing scenarios possible. In this thesis, spoofing is categorized into two types, static and dynamic spoofing. In static spoofing the location from where the spoofed ADS-B messages are transmitted does not change, in dynamic spoofing the location does change.

In static spoofing the target that is transmitting the faked ADS-B messages is not moving. If the ADS-B message is received by more than one GSs the algorithm can be initialized and the messages can be validated. As static spoofing will most likely only occur by targets that are not airborne the probability that a track can be initiated is very low. From a validation point of view there is no difference if the target is moving or not. Therefore, only dynamic spoofing scenarios are investigated.

In dynamic spoofing the target that is transmitting the spoofed messages is moving. If the messages are received by sufficient GSs, a track can be computed and the difference in location and velocity can be compared by the hypothesis tests. In this section two different dynamic spoofing scenarios are

investigated based on flights B and C. In the first investigated spoofing scenario an airborne target alters its ADS-B location from beginning to the end of the flight, such that is has a certain offset in the horizontal plane or vertical plane (or both). The spoofed trajectory is thus a parallel track with respect to the real airborne target. The second spoofing scenario is also a dynamic spoofing scenario, but the ADS-B location is altered in-flight. Thus after a certain amount of time, the track of the real target and the ADS-B track diverge from one another. In this scenario it is possible to observe how long it takes for the algorithm to detect this type of spoofing.

5.3.1. Flight B

Horizontal Spoofing

In this spoofing scenario the ADS-B track of flight B is spoofed using a horizontal offset in location as seen in Fig. 5.37b. Fig. 5.37a shows the results of the each hypothesis test on these spoofed tracks. Around a offset of 1250m to 1500m there is a turnover point where the messages are no longer validated, but considered to be spoofed. At distances larger than 2000m all three hypothesis test conclude that all ADS-B messages are spoofed. The Minimax requires a larger offset to determine the messages as spoofed. This results from a lower threshold value that is set in this hypothesis test. Results show that when one wants to obtain the highest spoofing detection the Neyman-Pearson hypothesis test is the best option. Although, this comes at the cost of a high false alarm rate in a no-spoofing scenario.

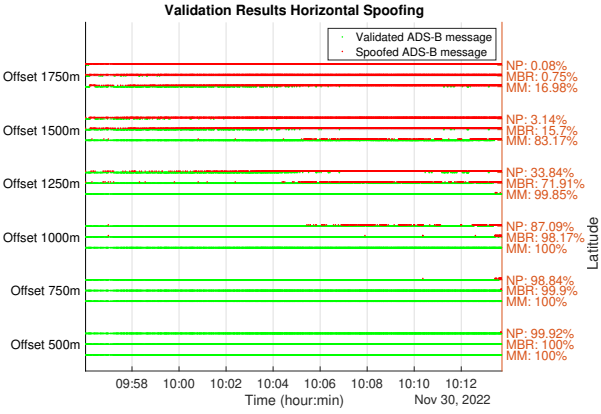

(a) Flight B: Spoofing detection results for horizontal spoofing. Right hand side shows the percentage of messages validated for each hypothesis test. Left hand size shows the offset compared to the true ADS-B track

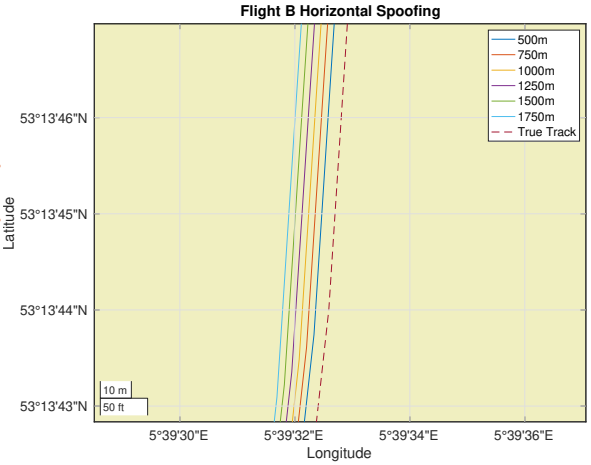

(b) Flight B: Spoofed ADS-B Tracks

Figure 5.37: Flight B: Section of the ADS-B track including spoofed ADS-B trajectories

Vertical Spoofing

In Fig. 5.38b the spoofing scenario is shown where only the height of the ADS-B track is altered. At initialization of the SIR filter considerable amounts of measurements are required to accurately determine that the track is spoofed. When the offset is larger it requires less measurements. The results show that the hypothesis tests are better in detecting spoofing in the vertical plane than in the horizontal plane. This is the result of several factors. It depends on the associated ADS-B uncertainty in height, which is smaller compared to the horizontal uncertainty. Furthermore, it also depends on the quality of the state estimate, which depends completely on the quality of the measurements the filter receives.

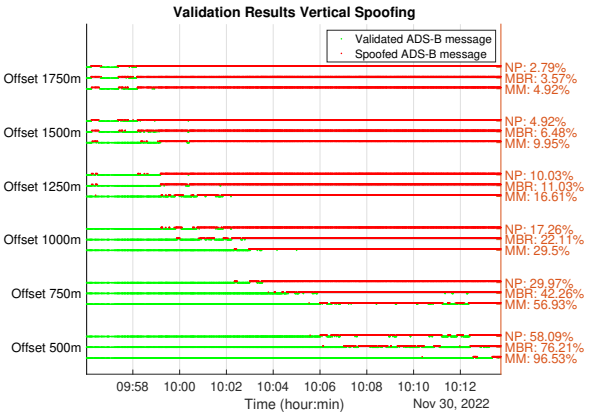

(a) Flight B: Spoofing detection results for vertical spoofing.Right hand side shows the percentage of messages validated for each hypothesis test. Left hand size shows the offset compared to the true ADS-B track

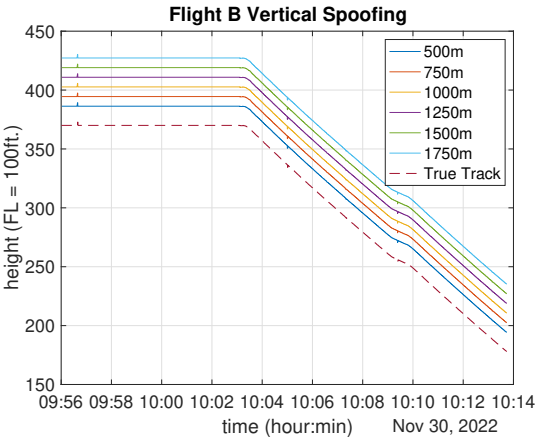

(b) Flight B: Spoofed ADS-B Tracks

Figure 5.38: Flight B: Vertical spoofing scenario

Horizontal and Vertical Spoofing

Fig. 5.39 shows the result of the algorithm when spoofed messages are offset both vertically and horizontally. The spoofed messages are detected earlier as the distance between the target and the spoofed ADS-B location is bigger.

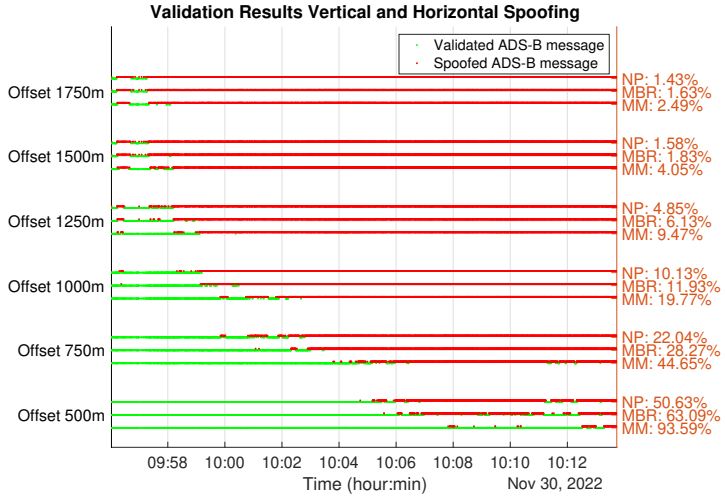

Figure 5.39: Flight B: Spoofed ADS-B Track. Offset indicates that the spoofed messages are off by the indicated amount in the horizontal and vertical plane.

In-flight Spoofing

The second dynamic spoofing scenario investigated is the in-flight spoofing scenario. In the scenario shown in Fig. 5.40, at time 10:03:25 the target starts to descent but the ADS-B reports that the altitude still is FL370. The time is takes for the hypothesis tests to detect the spoofed messages can be seen in Fig. 5.40a. Here, as expected the Neyman-Pearson test is the first test to detect the spoofed messages, after 1 minute the test concludes the messages to be spoofed. The MBR and Minimax need 1 minute and 10 seconds, and 1 minute and 20 seconds respectively.

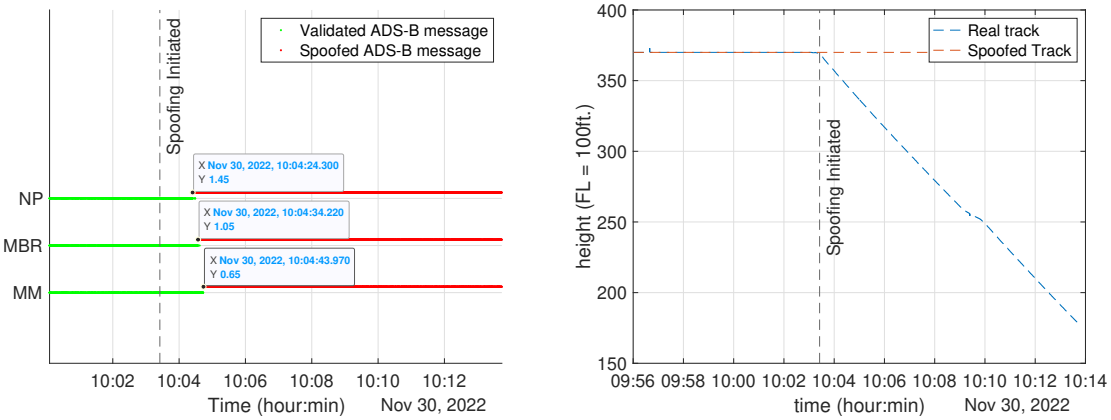

(a) Flight B: Validation result

(b) Flight B: altitude plot of the spoofing scenario

Figure 5.40: Flight B: Vertical In-flight spoofing scenario.

In the horizontal in-flight spoofing scenario the results are similar to the results in the vertical in-flight spoofing scenario, as seen in Fig. 5.41. In this scenario the longitude is spoofed such that the ADS-B deviates from the real track the target flies along.

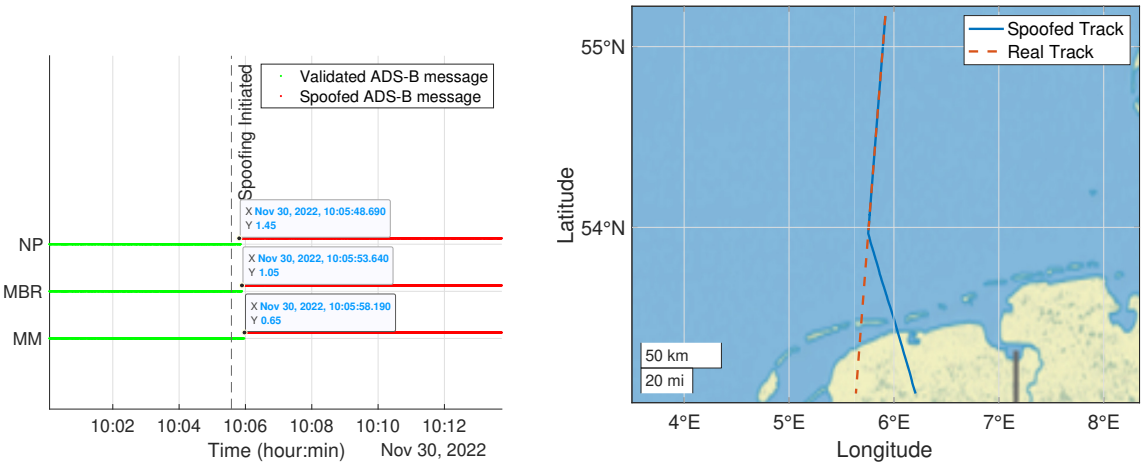

(a) Flight B: Validation result

(b) Flight B: plot of the spoofing scenario

Figure 5.41: Flight B:Horizontal In-flight spoofing scenario

Results have shown that in case of in-flight spoofing the validation algorithm is capable to detect the spoofed message when the distance between the real target and spoofed location is sufficient. Section 5.3 has shown that the distance between the target and ADS-B location must be about 1000 to 1200 meters depending on the test used for the ADS-B message to be rejected. It does take around 1 minute to detect this difference based on the simulated spoofing scenarios. In practice, the time it takes for spoofing to be detected depends largely of the speed by which the target and the spoofed location diverge, and the quality of the location estimate provides by the SIR filter.

5.3.2. Flight C

Fig. 5.42 and 5.43 show the results for the horizontal spoofing scenario in flight C. The scenarios are the same spoofing scenarios as for flight B, and again the results are quite similar. In the vertical plane all three hypothesis test are able to detect the spoofed messages faster than in the horizontal plane. Compared to flight B the validation algorithm obtains similar performance. In vertical spoofing flight B performs better than flight C.

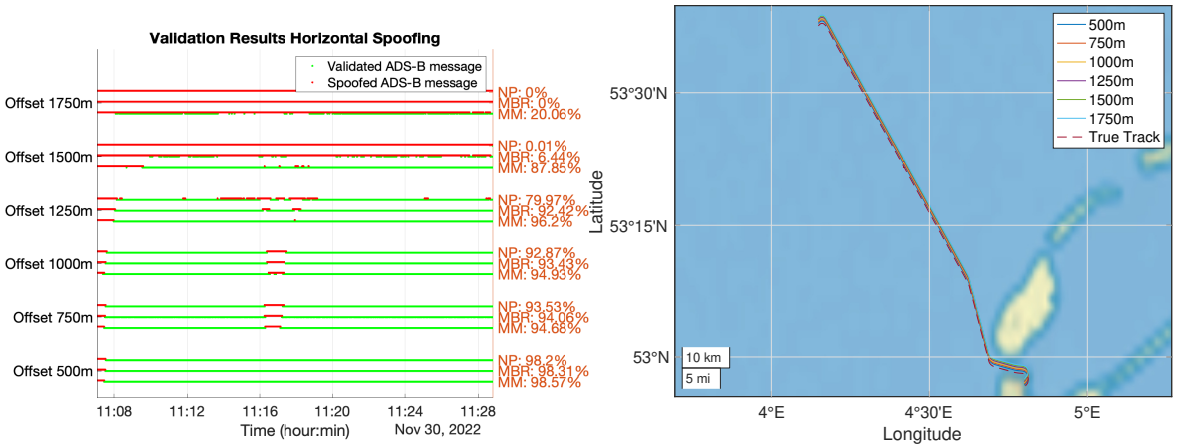

(a) Flight C: Validation result horizontal spoofing

(b) Flight C: plot of the spoofing scenario

Figure 5.42: Flight C: Horizontal spoofing scenario

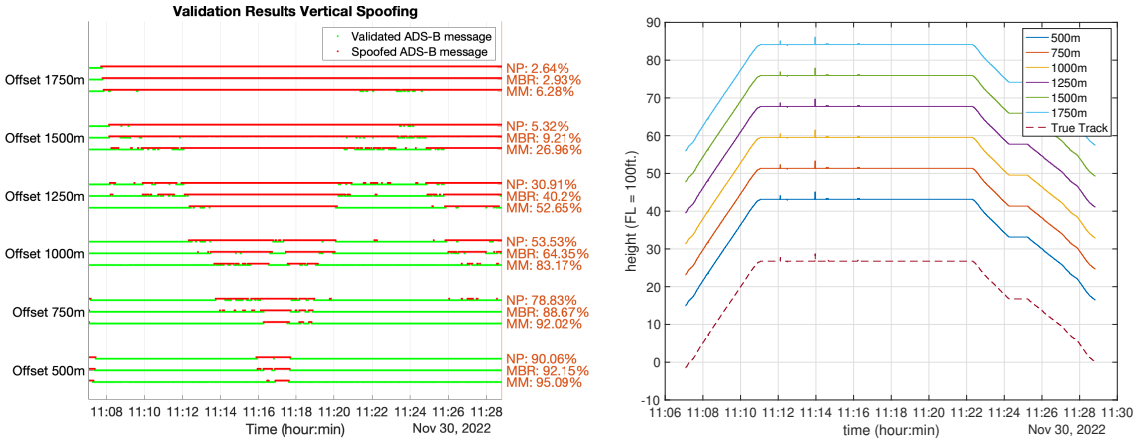

(a) Flight C: Validation result vertical spoofing

(b) Flight C: plot of the spoofing scenario

Figure 5.43: Flight C: Vertical spoofing scenario

5.4. Conclusion

The analysis of the SIR filter, hypothesis test and the investigated spoofing scenarios have provided several insights of the validation algorithm.

SIR Filter

The SIR filter has proven to be capable of proper state estimation in the analyzed cases, provided to proposed initialization method is used. These cases represent the common targets in the airspace and the issues associated with them. By using proposal density, that samples from the measurement, the filter can use significantly less particles compared to the traditional SIR filter where the proposal density is uniformly distributed across a large area. It achieves this performance whilst having a better ESS in the first measurement update.

- **Location**: The location estimate provided by the filter is accurate. In the horizontal plane the accuracy is higher than in the vertical plane. This difference in quality is an intrinsic property of multilateration based localization. At lower altitudes VDOP impacts the performance of the state estimation as is seen in flight C. At such low altitude the quality of the height estimate decreases or the estimate can even be wrong.
- **Velocity**: The ADS-B velocity can lag behind with respect to the SIR filter's velocity estimate because the ADS-B velocity is not updated every measurement, and the velocity estimate by the

SIR filter is updated every measurement. If the delta between the velocity gets too large, the validation will fail. In flight C the track is initiated at take-off of the target. In this flight the estimated velocity overshoots the ADS-B velocity. After several iterations of the filter, the filter corrects this overshoot, showing the filter is able to track such dynamics.

- **Impact of GS**: The impact of less GSs resulting in ambiguous measurements is observable by the value of the LRT, but it does not cause a degenerate set of particles, and the SIR filter does not diverge.
- **Effective Sampling Size**: ESS has shown the particle cloud does not suffer from the degeneracy problem or the sample impoverishment problem. During some iterations the ESS drops very low which indicates that the samples are a poor representation of the posterior distribution, but in all the flights the ESS recovers during next couple measurements. This shows that the drop in ESS is the result of faulty or noisy measurements.
- **Number of Samples**: Analysis of the number of samples required by the filter has shown that when more than 1000 particles are used there is not a significant increase in performance. The ESS shows than when less than 1000 particles are used the degeneracy problem occurs. When the number of samples required by the proposed SIR filter is compared to a traditional SIR filter, significantly less samples are required.
- **Proposal Distribution**: A low particle density per unit area introduces a large error in the state estimate in the first couple iterations of the filter. The rate of convergence of the filter is determined by the amount of samples per unit area.
- **Process Noise**: Investigation of several values for the process noise have shown that the filter converges for very high values, but becomes more unstable. For noise configuration 4 shown in Tab 5.1 the filter converges in flight B. Such values are thus not preferable.

Hypothesis Tests

General performance of all three investigated hypothesis test is good. Results show that there is no preferable test, but also that no test performs significantly worse than the other two. Because of this result the MBR can be preferable due to its simplicity, and it is less computationally intense.

- **Flight A**: Flight A is a very common type of flight in the airspace. All the hypothesis test correcly conclude that all messages are validated.
- **Flight B**: From flight B it can be concluded that the hypothesis tests are able to validate the messages if the state of the target is determined by a large amount of ambiguous measurements.
- **Flight C**: The helicopter flight has shown that validation can work below FL5 if the filter is given enough measurements to converge on the state of the target. Before this moment the quality of the estimate is not good enough for validation
- **Gaussian Assumption for Neyman-Pearson and Minimax**: The Gaussian assumption used in computing the threshold values for the Neyman-Pearson and Minimax hypothesis test is generally a good approximation
- **Neyman-Pearson False Alarm Probability**: The analyzed probability of false alarms do not show a clear preferable false alarm probability, but if high spoofing sensitivity is desired, a higher P_{fa} is advantageous. Analysis of the Neyman-Pearson hypothesis test has shown that the desired false alarm rate is not achieved in a single flight.
- **Minimum Bayes Risk Prior Probability**: The analysis of different prior probabilities in the MBR has shown that the different values have a limited effect on the performance of the algorithm. Any value between 0.01 and 0.99 achieves equal performance. Costs can be assigned for in the MBR to achieve a desired output. Very high cost values must be set such that is has a meaningful impact on the output of the hypothesis test.
- **Minimax Costs**: A high cost for the false alarm rate results in a relatively high threshold, a high cost for a missed detection does not result in a observable difference compared to equal costs. The assigned values to the cost must be of a high order to account for the extremely wide probability distribution of the spoofing-scenario.

Spoofing Analysis
Performance analysis of several spoofing scenarios has provided good results. On average ADS-B messages are validated if the target is roughly 1000 to 2000 meters away from the estimated ADS-B location, provided the velocity of both targets match. This difference in distance depends largely on the hypothesis test in question. All three hypothesis test perform as expected, and there does not appear a clear preferred option for the LVNL.

- **Static Spoofing**: Static spoofing scenarios can be detected using the same approach as dynamic spoofing. The filter must initialize a track on the received messages and the result of the validation logarithm follows.
- **Dynamic Horizontal Spoofing**: An airborne target that alters its ADS-B location such that is creates a parallel track can be detected when the distance is around 1250 to 1500 meters. The Neyman-Pearson hypothesis test will detect the spoofed messages first, next the MBR and finally the Minimax hypothesis test. These results are expected as they appear in order from the highest threshold to the lowest threshold. The horizontal separation minima used by LVNL is 5NM (9260m), this thus falls well within the save region. On approach, the separation minima drops down to 3NM (5556m), which also falls within the save region. Spoofed messages can thus be detected prior the the separation minima being breached.
- **Dynamic Vertical Spoofing**: In vertical spoofing scenarios where an airborne target alters its height only, the distance at which the hypothesis test decide the messages to be spoofed is more gradual, and not a clear cut-off point as is the case of horizontal spoofing. Eventually, if the SIR filter is given enough measurements flight B has shown that messages that are 750 meters away from the real targets can be detected. The time it takes for the hypothesis test to conclude the messages are spoofed is around 6 minutes. During this 6 minutes the spoofed target can only move its height plus or minus 750 meters to remain undetected, and thus the impact of spoofing is thus limited to this region. Vertically the minimum separation norms used by LVNL are 1000 ft (304m). The vertical spoofing detection thus does not fall within this safety regions which can cause possible dangerous situations.
- **In-flight Spoofing**: In this spoofing scenario the detection performance entirely depends on the rate which the targets location and the spoofed ADS-B location deviate from one another. Similar in vertical spoofing and horizontal spoofing, when distance becomes to large or the velocity is wrong, the hypothesis tests determine the ADS-B messages to be spoofed.
- **Detection at low altitudes**: Flight C has shown that the results from the hypothesis test suffer from the degraded state estimation by the SIR filter. At lower altitudes it is harder to detect vertical spoofing, horizontal spoofing performance remains comparable. Below FL5 (SDNS only provides coverage above FL5) the system is able to validate messages of the filter if given enough time to converge on the correct location, but in a spoofing scenario the uncertainty in the state estimate is not sufficient to consider the messages spoofed.

6 Conclusion

In the final chapter of this thesis the conclusions are presented. Section 6.1 gives the general conclusions to the performed work and results of the case study. Section 6.2 states several recommendations regarding the results and possible implementation of the algorithm at LVNL. Section 6.3 gives several aspects of the algorithm that can be explored in further work.

6.1. Conclusion

The advantages of ADS-B are promising, it enables enhanced situational awareness because ADS-B provides accurate information about the position, speed, altitude, and heading for ATC and nearby aircraft who have ADS-B-in capabilities, while requiring only a single receiver, possibly eliminating the traditional SSR. The situation arises that cyber vulnerability limits the full implementation of ADS-B at LVNL, simple mitigating measures have been taken but these are not sufficient for a wide-scale implementation. ADS-B spoofing is possible due to the lack of authentication and encryption in the ADS-B protocol. In this thesis a validation algorithm capable of validating the location reported inside ADS-B messages is designed. To achieve this three approaches where investigated.

- **Encryption**: Encryption adds complexity to the system and may require significant changes to ground-based receivers and other infrastructure. Encryption also introduces additional costs and technical complexities to the ADS-B system. It would require the deployment of encryption mechanisms across a vast number of aircraft and ground stations. This would involve significant investments and potentially impact the affordability and widespread adoption of ADS-B technology.
- **Machine Learning**: A ML approach can be taken where a model is trained using real data and spoofed data, such that it can detect spoofed messages. But the ML decision-making process is not easily interpretable or explainable. This lack of transparency can make it challenging to understand why certain ADS-B messages where validated or not, which is a concern in safety-critical domains like aviation. ML models often require significant computational resources for training and inference. Scaling the model to handle large volumes of ADS-B data in real-time is computationally demanding. Additionally, as new versions of ADS-B are introduced the used model may require frequent retraining or to maintain accurate validation performance.
- **Measurement Based**: The final method investigated makes use of measurements of the target that is transmitting ADS-B messages. PSR, SSR and multilateration are investigated among other methods. Multilateration has been found to be the best option as ADS-B messages can be validated using only two (i.e. ambiguous in location) or more GSs per measurement. This method can validate the ADS-B messages before a multilateration system can even track the target. Validation can thus be done before the location of the target is determined independently. The ADS-B location report can then already be used by ATC and the coverage of the multilateration system is increased from the area where four GSs receive a message, to the area there two GSs receive a message.

Comparison of the three proposed validation methods has found that multilateration based location validation as the best solution. Therefore a filter that can track an ADS-B target based on multilateration

measurements is designed, such that the origin of the ADS-B messages can be validated.

Due to the non-linearity of the multilateration measurement equation, and the added capability of handing ambiguous measurements, a particle filter design is proposed. Classical PF issues as the degeneracy problem and sample impoverishment problem where mitigated by using a novel sampling method that samples directly from the measurement at the initialization of the SIR filter. Without this novel method a traditional SIR filter (where the proposal density is uniformly distributed) requires roughly a million particles to converge on the location of the target. Below this amount of particles the traditional SIR filter fails. The proposed SIR filter can converge on the location of the target using only 1000 particles.

A likelihood ratio test is used to determine if the ADS-B message is spoofed. Determining a threshold value can be somewhat trivial, therefore, three different tests are explored to find which one is best suited. The implemented tests are, the Minimum Bayes Risk, the Neyman-Pearson and the Minimax Hypothesis test. Results have found that each test is capable of correct ADS-B validation. The Neyman-Pearson has the highest threshold generally, followed by the MBR and the Minimax. Determining which test is preferred is left to LVNL's operational experts on ADS-B implementation. From a technical point of view no preferred option arises. For spoofed ADS-B messages the hypothesis test can detect the spoofing if the distance between the transmitter that is transmitting the spoofed message and the location inside the spoofed message is roughly 1000 to 2000 meters, depending on the altitude and hypothesis test used. Horizontally this falls within LVNL's separation minima, vertically this falls outside the separation minima.

Publication FUSION2024
The results of the implemented SIR filter including the sampling method that samples from the measurement is written up as a *draft* paper with the *intention* to submit at the ISIF FUSION2024 congress.

6.2. Recommendations

- **Implementation at LVNL**: For the algorithm to work operationally the surveillance department at LVNL must decide on how to implement the validation tool. Decisions must be made on who will develop the tool for final use, who will be responsible for operational use, and what will be the associated procedures for handling of spoofed ADS-B messages. For now, the algorithm can be used as a standalone tool to analyze incidents regarding ADS-B that occur in the airspace covered by SDNS. To perform such analyses several pre-processing step must be done before the algorithm can be used. These steps require time and are not automated, these processes can be standardized in some way such that the validation tool is easy to use by LVNL.
- **Tuning of filter parameters**: The empirically tuned variables can be tuned differently by LVNL such that the desired performance or stability is obtained. This task can be done by discussing the desires and wishes of the colleagues who work at the operational side of LVNL.
- **Generic Validation**: The implemented hypothesis test used in the algorithm can also be used to validate ADS-B messages based on the track ARTAS provides. Using this approach all surveillance sensors that track and receive ADS-B targets can be used for validation.
- **Spoofing Hypothesis Tuning**: The probability distribution describing a spoofed target is a uniform distribution with the size of the area where ADS-B messages can be received. Currently, this area is approximated using a rectangular approximation. This approximation can be improved such that the quality of the LRT also improves.
- **ADS-B message to MLAT measurement link**: In the case study, and thus SDNS, there is no direct link between the ADS-B messages and its corresponding MLAT measurement which introduces difficulties in data association. This is solved using time extrapolation, but this comes at the cost of accuracy. In future implementation of the algorithm this direct link between the ADS-B message and its MLAT measurement is required for proper one to one validation.
- **When to validate a target based on validated measurements**: The validation algorithm validates individual measurements. There still needs to be some decision moment when the target is validated based on a number of ADS-B messages. This could also be extended to a white-list where

validated ICAO addresses are stored for a certain amount of time to reduce the computational load of validating all aircraft continuously in real-time.

6.3. Future Work

- **Sampling from the MLAT measurement**: Sampling from the measurement has proven to be a crucial component in the performance of the filter, but the used sampling model doesn't take into account which time measurement of the TDOA measurement arrived first. If the sign of the TDOA measurement is taken into account one of the two sheets of the two-sides hyperbola can be neglected as is illustrated in Fig. 6.1. If $T_1 < T_2$, the signal must arrive first at GS1, therefore only the left hyperbola is a valid solution. Drawing samples from the right hyperbola can thus be omitted. Then, more samples are left to populate the remaining hyperbolic sheet resulting in more samples per unit area, which improves the estimation quality of the filter.

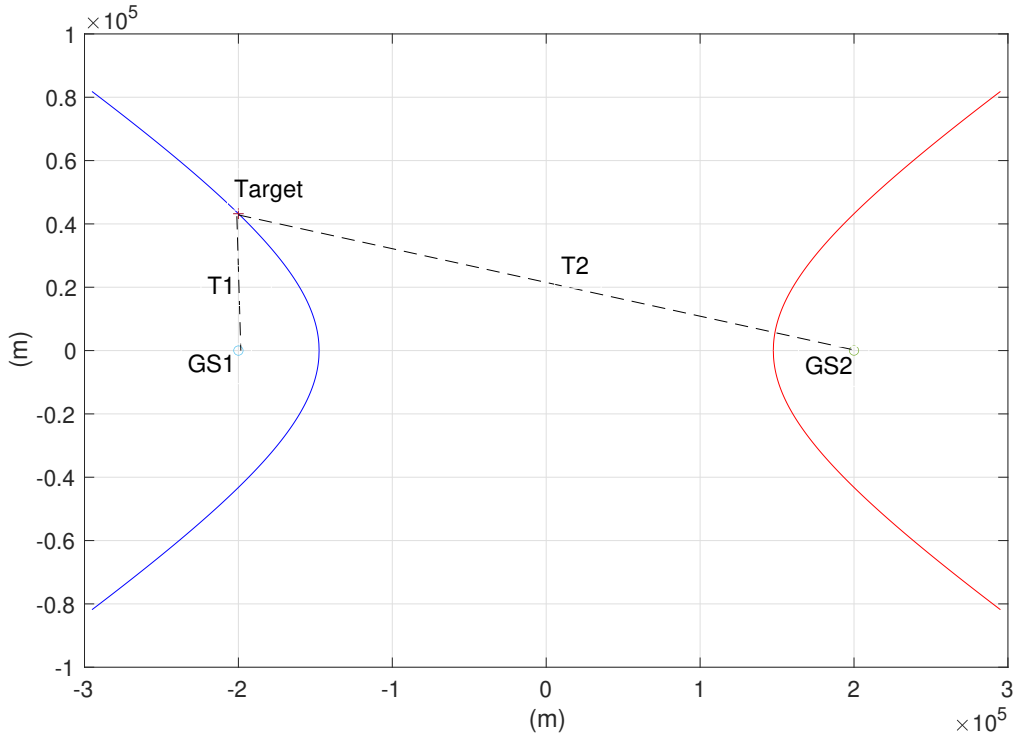

Figure 6.1: 2-dimensional representation of a TDOA measurement. If it is known that $T_1 < T_2$, the right hand size is not a solution to the measurement equation, then drawing samples from is can be neglected.

- **Sampling Method**: The method used to sample from the measurement has difficulties sampling from the measurement when the sample density is to low, and the hyperbolas in question have little horizontal spread. Therefore the method by which the samples are drawn can be improved. One approach could be not sampling based on $z = f(x, y, \ldots)$ but $x = (y, z, \ldots)$ such that samples are drawn based on a vertical grid instead of a horizontal grid. This could possibly solve the problem as there is more vertical spread than horizontal spread in the general shape of the measured hyperbolas.
- **GS selection based on Geometric Dilution of Precision (GDOP)**: In section 3.1 the method for computing TDOA measurements based on TOA measurements is discussed. This method can be interchanged by a method where the TDOA measurements are computed based on four GSs that have the best possible geometric spread. For each MLAT measurement, some of the GSs have good GDOP with respect to one another, and some GSs have a bad GDOP with respect to one another. This GDOP can be computed for each measurement and the GSs with the best GDOP can be selected for the final TDOA measurement. With this approach the computational requirements

of the filter can be decreased. Note that this is only applicable for measurements with four or more GSs.

- **Analyze More Flights**: In the results in chapter 5 the SIR filter and algorithm is only investigated on three characterizing flights. Based on these flights the performance of the algorithm is investigated, but more insights can be obtained when the algorithm is tested on large amount of flights such that new insights or statistical properties can be observed. In addition, when more flights are analyzed it can be concluded with more certainty that the MISS filter is, or is not required in this filtering problem.
- **Interacting Multiple Model Particle Filter**: Targets that LVNL tracks are civil aircraft that can easily be modeled by several dynamic models. In the current implementation only the constant velocity model is used in prediction the state of the target. This prediction quality can be improved by making use of multiple models such a the coordinated turn model and the constant acceleration model.

Bibliography

[1] "Adsb out explanation," 2023, accessed April 14,2023. [Online]. Available: https://www.trig-avionics.com/knowledge-bank/ads-b/ads-b-explanations/

[2] EUROCONTROL, "Lssip 2019 - the netherlands local single sky implementation," p. 9, 2020, accessed April 19,2023. [Online]. Available: https://www.eurocontrol.int/sites/default/files/content/documents/official-documents/reports/LSSIP2017_Netherlands_Released.pdf

[3] Y. Kim, "Ambiguous measurement update in terrain-referenced navigation with particle filtering," Ph.D. dissertation, 11 2017.

[4] "Coordinate systems," 2013, accessed April 14,2023. [Online]. Available: http://dirsig.cis.rit.edu/docs/new/coordinates.html

[5] "Common view gps time transfer," 2016, accessed April 26,2023. [Online]. Available: https://www.nist.gov/pml/time-and-frequency-division/time-services/common-view-gps-time-transfer

[6] M. Leonardi, L. Gregorio, and D. Fausto, "Air traffic security: Aircraft classification using ads-b message?s phase-pattern," *Aerospace*, vol. 4, p. 51, 10 2017.

[7] "Luchtvaart, maandcijfer nederlandse luchthavens van nationaal belang," 2023, accessed April 17,2023. [Online]. Available: https://www.cbs.nl/nl-nl/cijfers/detail/37478hvv

[8] A. A. Barsheshat, "Implementation of ADS-B systems — Benefits and considerations," Tech. Rep., 2011.

[9] L. Purton, P. Hussein Abbass, and S. Alam, "Identification of ADS-B System Vulnerabilities and Threats," Tech. Rep., 2010.

[10] EUROCONTROL, "Eurocontrol specification for surveillance data exchange asterix part 12 category 21 ads-b target reports," Tech. Rep., 2021, accessed April 17,2023. [Online]. Available: https://www.eurocontrol.int/sites/default/files/2021-12/asterix-adsbtr-cat021-part12-v2-6.pdf

[11] D. McCallie, J. Butts, and R. Mills, "Security analysis of the ADS-B implementation in the next generation air transportation system," *International Journal of Critical Infrastructure Protection*, vol. 4, no. 2, pp. 78–87, 8 2011.

[12] M. Schäfer, V. Lenders, and I. Martinovic, "Experimental Analysis of Attacks on Next Generation Air Traffic Communication," Tech. Rep.

[13] W.-P. Air Force Base, "EXPLOITING THE AUTOMATIC DEPENDENT SURVEILLANCE-BROADCAST SYSTEM VIA FALSE TARGET INJECTION AIR FORCE INSTITUTE OF TECHNOLOGY," Tech. Rep.

[14] K. D. Wesson, T. E. Humphreys, and B. L. Evans, "Can Cryptography Secure Next Generation Air Traffic Surveillance?" Tech. Rep.

[15] S. Amin, T. Clark, R. Offutt, and K. Serenko, "Design of a cyber security framework for ads-b based surveillance systems," in *2014 Systems and Information Engineering Design Symposium (SIEDS)*, 2014, pp. 304–309.

[16] J. Baek, E. Hableel, Y.-J. Byon, D. S. Wong, K. Jang, and H. Yeo, "How to protect ads-b: Confidentiality framework and efficient realization based on staged identity-based encryption," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 3, pp. 690–700, 2017.

[17] H. Yang, Q. Zhou, M. Yao, R. Lu, H. Li, and X. Zhang, "A practical and compatible cryptographic solution to ads-b security," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 3322–3334, 2019.

[18] X. Ying, J. Mazer, G. Bernieri, M. Conti, L. Bushnell, and R. Poovendran, "Detecting ads-b spoofing attacks using deep neural networks," in *2019 IEEE Conference on Communications and Network Security (CNS)*, 2019, pp. 187–195.

[19] J. Wang, Y. Zou, and J. Ding, "ADS-B spoofing attack detection method based on LSTM," *Eurasip Journal on Wireless Communications and Networking*, vol. 2020, no. 1, 12 2020.

[20] S. Khan, J. Thorn, A. Wahlgren, and A. Gurtov, "Intrusion Detection in Automatic Dependent Surveillance-Broadcast (ADS-B) with Machine Learning," in *AIAA/IEEE Digital Avionics Systems Conference - Proceedings*, vol. 2021-October. Institute of Electrical and Electronics Engineers Inc., 2021.

[21] M. El-Ghoboushi, A. Ghuniem, A. H. Gaafar, and H. E. D. Abou-Bakr, "Multiple aircrafts tracking in clutter for multilateration air traffic surveillance system," in *Proceedings of 2018 International Conference on Innovative Trends in Computer Engineering, ITCE 2018*, vol. 2018-March. Institute of Electrical and Electronics Engineers Inc., 3 2018, pp. 225–230.

[22] W. Huygen, J. Sun, and J. Hoekstra, "ADS-B Signal Verification Using a Coherent Receiver." MDPI AG, 12 2021, p. 4.

[23] C. Reck, M. S. Reuther, A. Jasch, and L. P. Schmidt, "Verification of ADS-B positioning by direction of arrival estimation," *International Journal of Microwave and Wireless Technologies*, vol. 4, no. 2, pp. 181–186, 4 2012.

[24] D. J. Torrieri, "Statistical Theory of Passive Location Systems," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-20, no. 2, pp. 183–198, 1984.

[25] M. Compagnoni, R. Notari, F. Antonacci, and A. Sarti, "A comprehensive analysis of the geometry of TDOA maps in localization problems," *Inverse Problems*, vol. 30, no. 3, 3 2014.

[26] G. Wang, S. Cai, Y. Li, and N. Ansari, "A Bias-Reduced Nonlinear WLS Method for TDOA/FDOA-Based Source Localization," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 10, pp. 8603–8615, 10 2016.

[27] K. C. Ho and Y. T. Chan, "Geolocation of a Known Altitude Object From TDOA and FDOA Measurements," Tech. Rep.

[28] J. Li, F. Guo, and W. Jiang, "A linear-correction least-squares approach for geolocation using FDOA measurements only," *Chinese Journal of Aeronautics*, vol. 25, no. 5, pp. 709–714, 10 2012.

[29] J. Chen, Y. Zhao, C. Zhao, and Y. Zhao, "Improved two-step weighted least squares algorithm for tdoa-based source localization," in *2018 19th International Radar Symposium (IRS)*, 2018, pp. 1–6.

[30] G. Shen, R. Zetik, and R. S. Thoma, "Performance comparison of toa and tdoa based location estimation algorithms in los environment," in *2008 5th Workshop on Positioning, Navigation and Communication*, 2008, pp. 71–78.

[31] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[32] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," *IEE Proc. F Radar Signal Process. UK*, vol. 140, no. 2, p. 107, 1993. [Online]. Available: http://dx.doi.org/10.1049/ip-f-2.1993.0015

[33] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.

[34] A. Doucet, S. Godsill, and C. Andrieu, "On sequential monte carlo sampling methods for bayesian filtering," *Statistics and Computing*, vol. 10, pp. 197 – 208, 2000.

[35] J. S. Liu and R. Chen, "Sequential monte carlo methods for dynamic systems," *Journal of the American Statistical Association*, vol. 93, no. 443, pp. 1032–1044, 1998. [Online]. Available: http://www.jstor.org/stable/2669847

[36] E. Veach and L. J. Guibas, "Optimally combining sampling techniques for monte carlo rendering," in *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '95. New York, NY, USA: Association for Computing Machinery, 1995, p. 419?428. [Online]. Available: https://doi-org.tudelft.idm.oclc.org/10.1145/218380.218498

[37] J. Kronander and T. B. Schön, "Robust auxiliary particle filters using multiple importance sampling," in *2014 IEEE Workshop on Statistical Signal Processing (SSP)*, 2014, pp. 268–271.

[38] J. Zuo, Y. Liang, Y. Zhang, and Q. Pan, "Particle filter with multimode sampling strategy," *Signal Processing*, vol. 93, no. 11, pp. 3192–3201, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0165168413001667

[39] H. D. Y Boers, "Interacting multiple model particle filter," *IEE Proc. Radar Sonar Navig*, vol. 150, p. 344 – 349, 10 2003.

[40] D. Crisan, J. Miguez, and G. Ríos Muñoz, "On the performance of parallelisation schemes for particle filtering," *EURASIP Journal on Advances in Signal Processing*, vol. 2018, 05 2018.

[41] S. S. Blackman, *Multiple-Target Tracking with Radar Applications*. 610 Washington Street Dedham, MA 02026: Arteck House, Inc, 1986.

[42] H. V. Poor, *An Introduction to Signal Detection and Estimation (2nd Ed.)*. Berlin, Heidelberg: Springer-Verlag, 1994.

[43] N. Bergman, "Recursive bayesian estimation: Navigation and tracking applications," Ph.D. dissertation, Linköping Univ., Linköping, Sweden,, 1999.

A

Matlab Sampling Code

```
syms x y z s1x s1y s1z snx sny snz c1 Z
 f(x,y,z) = sqrt((x - s1x).^2 + (y - s1y).^2 + (z - s1z).^2) - sqrt((x - snx).^2 + (y - sny)
     .^2 + (z - snz).^2) - c1*Z;
 f = solve(f,z);
matlabFunction(f(1,1),'File','meas_eq_z');
```

B

ARTAS Track Flight A,B,C

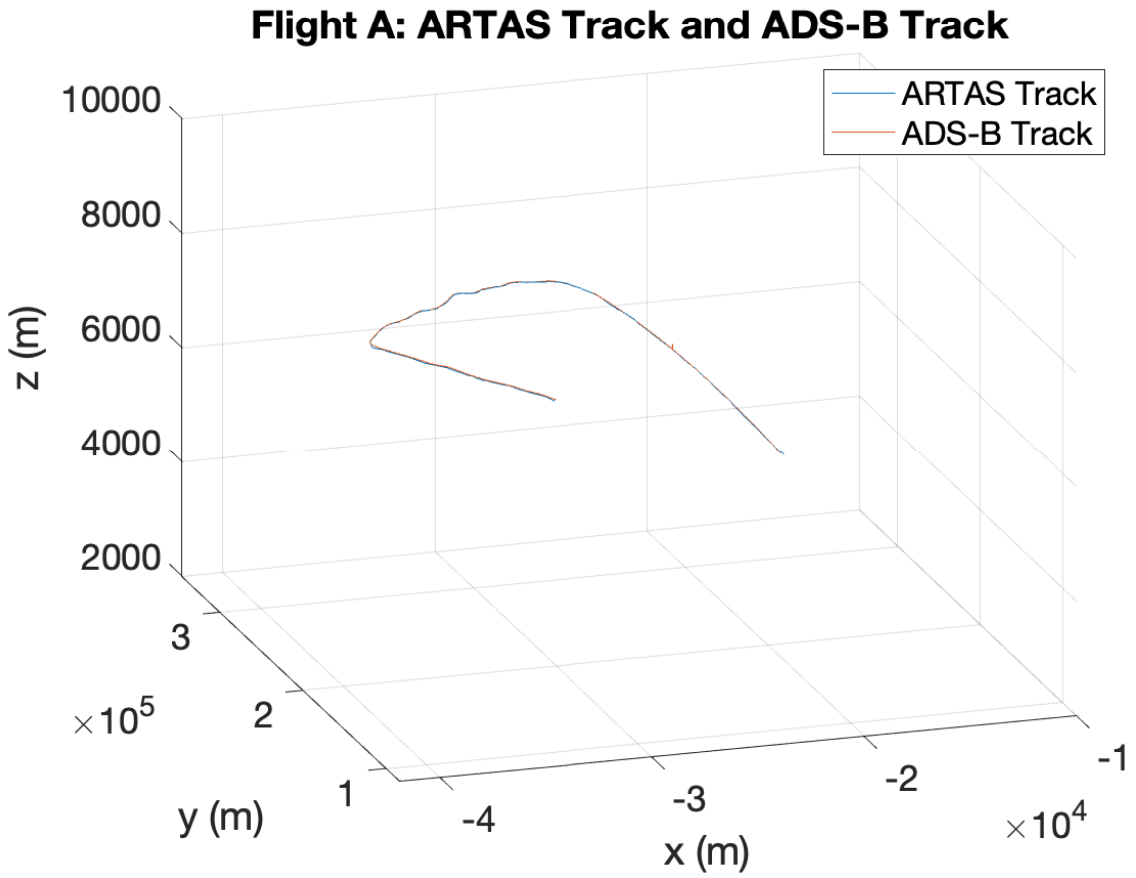

Figure B.1: ARTAS track and ADS-B track for flight A. ARTAS track uses several primary and secondary surveillance radars to determine the track.

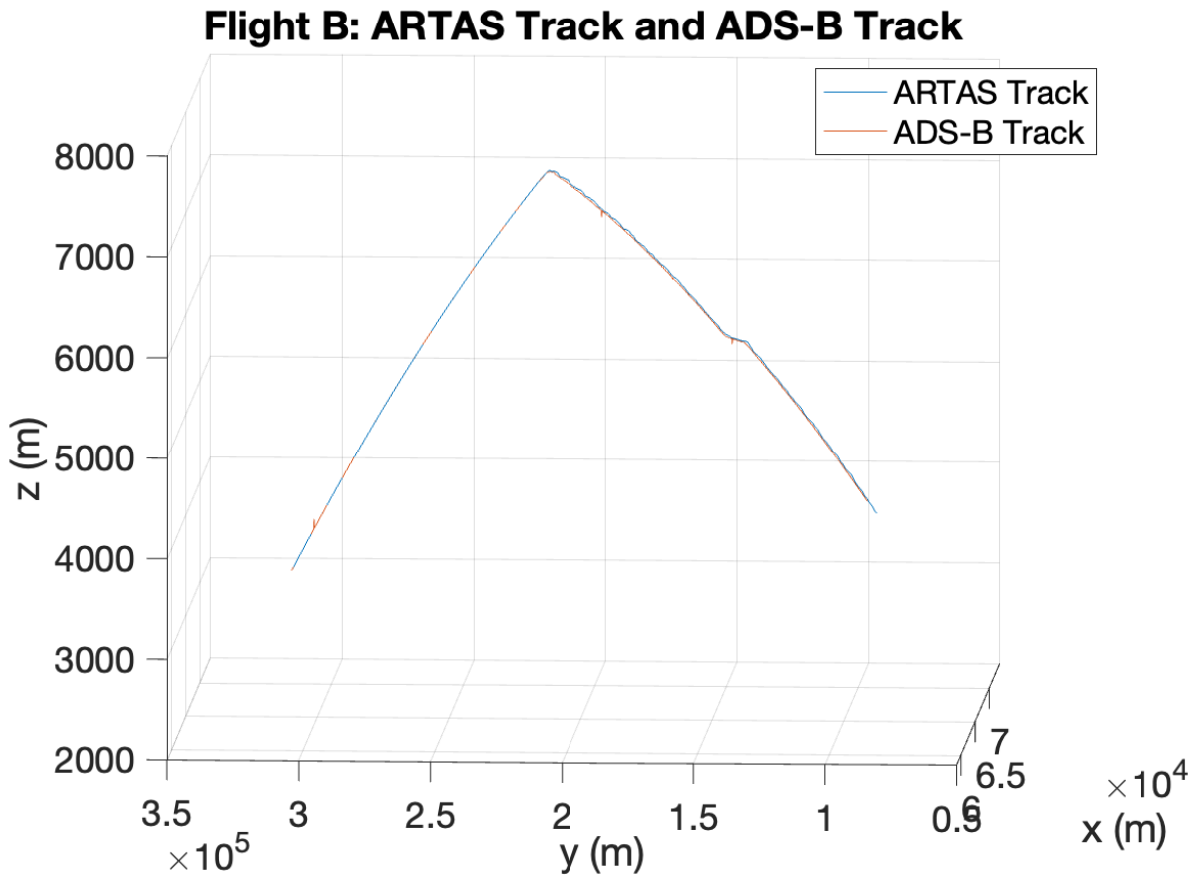

Figure B.2: ARTAS track and ADS-B track for flight C. ARTAS track uses several primary and secondary surveillance radars to determine the track.

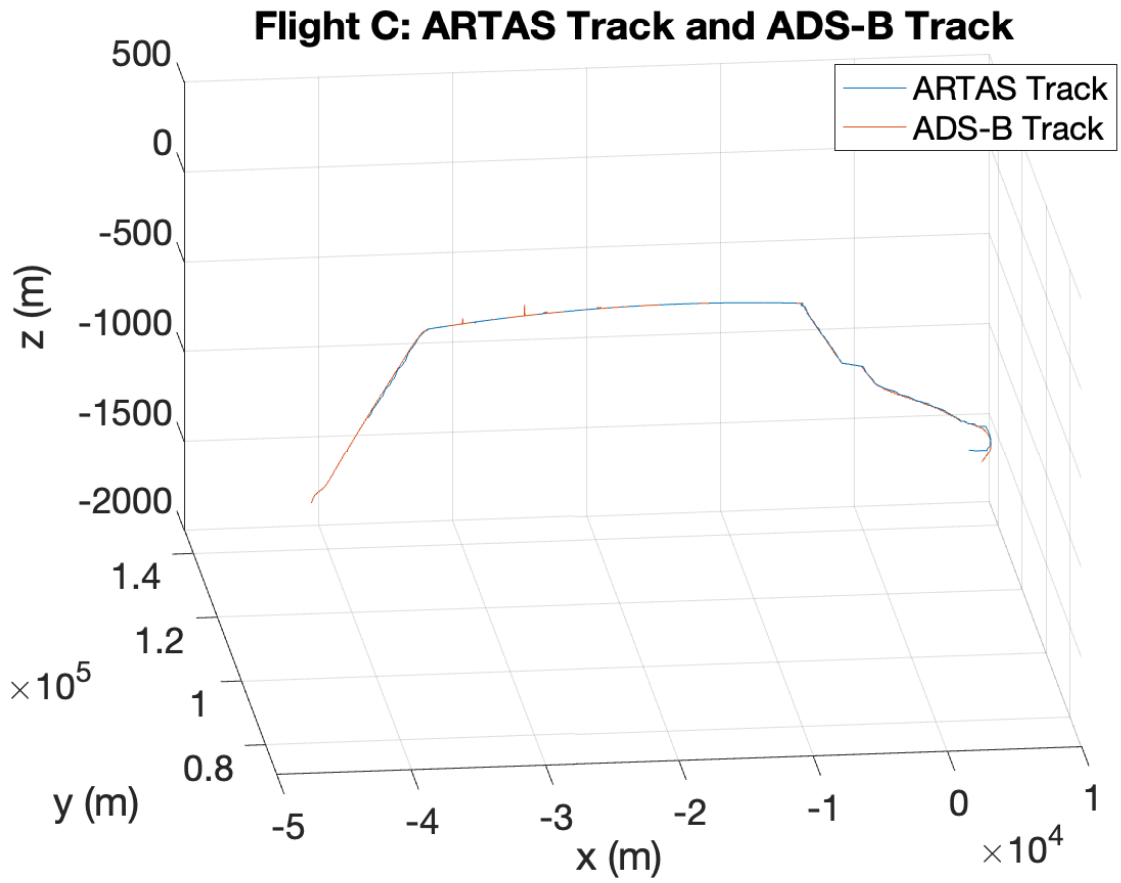

Figure B.3: ARTAS track and ADS-B track for flight C. ARTAS track uses several primary and secondary surveillance radars to determine the track.