Integrity Analysis for Aviation and Automotive Applications

by

Courtney Mario

Submitted in partial fulfillment of
the requirements for the degree of

Master of Science

in

Mechanical Engineering

TUFTS UNIVERSITY

May 2011

ADVISER: Jason Rife

Abstract

As technology advances, automated systems are being considered for many safety applications within the aviation and automotive industries. Such applications include airport surface traffic management systems and automotive lane-keeping systems. Possibly misleading sensor data must be considered when designing these systems since they operate in the absence of a human controller. Currently, no integrity requirements (e.g. no sensor anomaly detection requirements) have been defined for these types of automated airport surface and automotive applications. This thesis both demonstrates an integrity risk assessment for an airport surface application and proposes an integrity monitor for an automated lane-keeping system. Within this work, airport surveillance sensors, such as radar, multilateration, and ADS-B, are evaluated and modeled, with a specific emphasis on the impact of their random noise and off-nominal performance characteristics on the system's integrity. Similarly, for automated lane-keeping, two image processing techniques – gradient detection and optical flow – are applied to perform a cross-check in order to verify measurement quality.

Acknowledgements

Thank you to my advisor Jason Rife for all the help and support he has given me throughout this research. I have gained a lot during the past two years under his guidance and I truly appreciate all that he has done.

Also, thank to you to the members of the ASAR Lab Group at Tufts University. Their feedback and advice during group meetings provided alternative views and helped me to reach new solutions.

Finally, a special thank you to my family and friends. To my parents, Joe and Karen, and brother, Chris, you have offered constant love and support throughout my time at Tufts and I really enjoyed being able to share these experiences with you.

Table of Contents

Figures

Tables

Chapter 1: Introduction

1.1 Background

Automated systems are being designed for both the aviation and automotive industries in order to increase efficiency and improve safety. In order for these systems to operate successfully and avoid additional risk for those involved, safety requirements must be defined and considered from the early stages of system development.

Within the aviation industry, the FAA is currently in the process of implementing the Next Generation Implementation Plan (NextGen) to revamp the entire airspace system. Current estimates show that total delays will be reduced by about 35% from delays projected in the absence of NextGen. This drop will lead to an estimated $23 billion in total benefits to airlines, the public, and the FAA. It will also save about 1.4 billion gallons of fuel, cutting carbon dioxide emissions by about 14 million tons [1]. On the airport surface, this relates specifically to traffic management systems, which are being designed to increase efficiency and ultimately increase aircraft throughput [2].

Under NextGen, automated systems are being designed to perform these required airport surface traffic management functions. This strategy for improvement is based on the idea of trajectory-based operations, and particularly, the use of 4D trajectories, which adds an additional time parameter to the 3D position. One specific decision support tool to ensure that aircraft conform to these trajectories is a Conflict Detection and Resolution (CD&R) algorithm, which both detects non-conforming aircraft and provides necessary resolution steps. In order to ensure that these

automated systems are performing correctly, safety requirements must be defined from the beginning of algorithm design.

Similarly, in the automotive industry, there are currently more than 1.2 million road-departure crashes each year in the United States [3] and studies suggest that lane departure warning systems alone can decrease this for passenger cars by 10%, and for tractor-trailer trucks by as much as 30% [4]. The technology for these lane departure warning systems could be further adapted for automated lane-keeping systems, providing an active safety system to further reduce road-departure accident rates. Other similar active safety systems being designed involve braking or steering cars away from impending collisions [5]. Overall, the implementation of these types of systems will allow for safer and more efficient travel.

In recent years, fully autonomous driving has also been a major focus of research in the automotive industry. The DARPA Grand Challenge is a competition for driverless vehicles and encourages autonomous driving designs for both urban and desert environments [6]. Google has also been focusing on a driverless car of their own and have performed several thousand miles of testing on US roads [7]. However, despite the significant amount of work being done in this field, no safety requirements have been established. In order for automated driving to be successfully implemented, these requirements must be considered to ensure safety of life.

1.2 Motivation for Integrity Requirements

Automated systems for both the aviation and automotive industries must be able to demonstrate safe operations before being approved for implementation. Specifically, these systems must be able to show that they can meet defined safety requirements and prove that they in fact will not

cause additional risks for those involved. Therefore, safety requirements such as integrity must be defined in order to provide a metric that can be used to determine if these systems can be safely implemented.

Integrity refers to a system's ability to detect an error and therefore depends strongly on sensor performance. Related to a system's rate of missed detections, this performance metric ensures that a hazardous condition does not go unnoticed. Aviation traffic management relies on sensors to detect aircraft and ground vehicles on the airport surface, while automotive safety applications rely on sensors to determine a car's position – both absolute and relative to surrounding vehicles and objects. These types of sensors are often characterized by their position accuracies, but it is also necessary to consider sensor anomalies. In such cases, these sensors provide misleading information, which impacts a system's integrity and its ability to accurately detect errors.

Automated systems require a close look at integrity since the human operator is replaced by computerization. As a result, the system must be able to interpret sensor anomalies as a human operator would do naturally. For airport applications, automated systems will soon be performing many tasks that are currently the responsibility of air traffic controllers. These include defining and relaying surface routes to pilots, as well as ensuring that aircraft follow these routes. For an automotive application such as automated lane-keeping, computer algorithms will need to ensure that the car is conforming to the correct lane lines, as there is no longer a driver present to correct for system errors. In both applications, a system error could result in injury or even death. Therefore, automated systems must be able to alert for all system errors in a timely manner, corresponding to strict integrity requirements.

A great deal of work has been done to certify that navigation systems based on GPS positioning are able to meet strict integrity bounds. In order to ensure that hazardous errors do not go unnoticed, integrity risk specifications are very strict for these navigation applications. For example, future Category-III (CAT III) systems rely on a GPS augmentation system called the Ground Based Augmentation System (GBAS) to fully support automated aircraft landing. As a result, the Federal Aviation Administration (FAA) is requiring that CAT III systems meet an integrity risk less than 10^{-9} over a 30 second time period [8]. Yet, the algorithms being considered for airport surface and lane-keeping applications rely on more than just GPS positioning. Therefore, further analysis must be done in order to quantify integrity bounds for these types of systems.

Monitoring algorithms are typically used to ensure that a system meets integrity requirements. In these algorithms, a monitoring metric is defined that will allow the algorithm to measure its performance and safety impact, and an alert is triggered when this metric surpasses a determined threshold. For aviation and automotive applications, this metric could for instance be the separation distance between two aircraft on the same runway or the separation distance between two cars traveling in the same lane. The threshold itself is selected based on a value that will provide for low integrity risk, and ultimately a low probability of missed detections. However, the rate of alarms, which impacts another performance metric called continuity, must also be considered. This metric quantifies the requirement for total alarm rate – including both true and false alarms, and is needed in order to ensure that alarm interruptions do not impact normal operations too frequently. Therefore, when designing monitoring algorithms for automated systems, both integrity and continuity requirements must be considered.

Currently, no established integrity requirements exist for automated systems being designed for both airport surface and automated driving applications. In order to determine these bounds, it is necessary to understand the sensor capabilities for the sensors that these types of systems will rely on. In particular, it is important to understand both the associated positioning errors and fault modes. Therefore, this work proposes an in-depth look at the impact of sensor systems on airport surface traffic management systems as well as an automated lane-keeping system. Specifically, on the airport surface, surveillance systems include radar, multilateration and Automatic Dependent Surveillance–Broadcast (ADS-B), while for automated lane-keeping, sensors include video cameras, which provide information through means of various image processing methods. The findings from this research will impact the designs for both types of automated systems.

1.3 Thesis Contributions

The primary goal of this thesis is to address the need for integrity in automated systems for both aviation and automotive applications and relate sensor performance associated with these types of systems to integrity requirements. The contributions associated with is work are summarized as follows.

1. **Demonstration that airport surface surveillance sensors will impact the integrity and continuity of conflict detection algorithms.** This is related specifically to aviation applications and is done by modeling and analyzing the performance of available airport surveillance sensors, and the resulting impact on CD&R algorithms.
2. **Derivation of an integrity monitor for vision-based lane detection systems.** This is related specifically to automotive

applications and is done by applying two image processing techniques to lane marker identification, and monitoring the integrity of resulting detected lane lines.

1.4 Thesis Overview

The remainder of this thesis is divided into several chapters which explain automated system applications for the aviation and automotive industries and their integrity requirements in further detail.

Chapter 2 first looks at the impact of airport surface surveillance sensors on traffic management systems such as CD&R. Within this chapter, a sample conflict detection algorithm is proposed with accounts for integrity and continuity requirements based on generic surveillance sensor performance. Performance capabilities of specific sensor technologies, such as radar, multilateration, and ADS-B, are then analyzed and related to the derived conflict detection algorithm. Further background information for the sensor systems analyzed in this chapter can be found in the Appendices.

Next, Chapter 3 introduces an integrity monitoring strategy for vision-based automotive lane-departure warning systems. This monitor identifies sensor anomalies by comparing two distinct image processing algorithms. The first is an existing algorithm [9] that fits a linear-parabolic model to lane boundaries while the second is an optical flow based algorithm used to verify these boundaries. Applying this monitor to video data collected on urban roadways, this chapter demonstrates the potential benefits of this method and the impact of the performance of these algorithms on the integrity of detecting lane lines.

Finally, Chapter 4 discusses the overall conclusions of the integrity analysis for both applications and outlines future work.

Chapter 2: Impact of Surface Surveillance Sensors on the Integrity and Continuity of Conflict-Detection Monitoring

Current estimates suggest that by 2018, NextGen traffic management will be able to reduce total delays by about 35% from delays projected in the absence of NextGen. This reduction will provide an estimated $23 billion in total benefits to airlines, the public, and the FAA, and save about 1.4 billion gallons of fuel, cutting carbon dioxide emissions by about 14 million tons [1].

In order to achieve these goals, the FAA is planning to enhance airport surface operations through four-dimensional trajectory-based operations, which allow taxiing aircraft to coordinate in a more efficient manner. This efficiency is expected to result in high levels of throughput, which are in turn anticipated to stress the capabilities of local and ground controllers to identify and resolve surface conflicts. As such, automated Conflict Detection and Resolution (CD&R) algorithms are under development [2]. As the name implies, CD&R algorithms consist of two components: (1) a conflict detection algorithm that identifies potential runway incursions, losses of separation, or other surface movement conflicts, and (2) a conflict resolution algorithm that advises controllers and pilots how to achieve a safe outcome with minimal impact on surface traffic throughput. This chapter focuses specifically on the first of these components, the Conflict Detection (CD) algorithm.

The performance of CD algorithms depends in part on the position information provided by airport surface surveillance sensor systems. Existing and emerging surface surveillance systems include radar, multilateration, and Automatic Dependent Surveillance-Broadcast (ADS-B). To understand how surveillance systems will impact the performance of CD

algorithms, it is important to quantify both the nominal operations and off-nominal behaviors of these sensors.

2.1 Defining Continuity and Integrity

To be useful in an operational setting, the CD algorithm must at a minimum satisfy requirements for two essential performance metrics: continuity and integrity. The main function of the CD algorithm is to produce an alert in the event that an actual conflict should occur between two taxiing aircraft or between an aircraft and a surface vehicle. It is also essential that the CD algorithm avoid issuing false alerts when no actual conflict has occurred. Frequent false alerts degrade user confidence in the automated CD function and also introduce risk each time pilots are forced to react rapidly to a conflict resolution order. In loose terms, the allowed continuity risk is a limit on the false (and real alerts) that could interfere with safe, efficient surface operations. On the other hand, while limiting the occurrence of alerts is needed to ensure continuity, the CD algorithm must be able to alert in the event of a hazardous situation. In loose terms, the integrity requirement specifies how often the CD algorithm may fail to produce the necessary alert in a timely fashion. Figure 1 offers a rough overview of desired (shaded) and undesired (unshaded) outcomes of conflict detection.

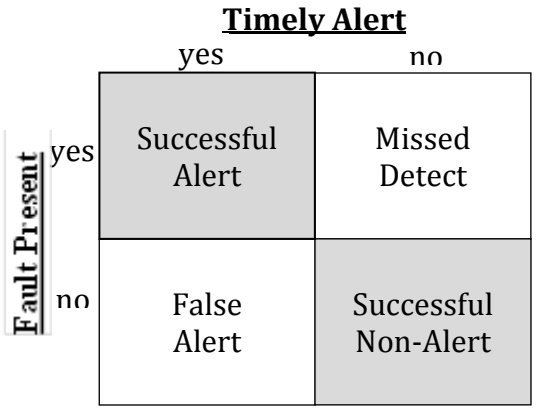

Figure 1. Desired (Shaded) and Undesired (Unshaded) Outcomes for CD Algorithm

CD algorithm continuity and integrity risk depend strongly on surveillance sensor performance. While significant work has been done to quantify continuity and integrity requirements for aviation navigation sensors [10-12], relatively little has been done to relate these metrics to surface surveillance applications [13-14]. Applying continuity and integrity requirements to multi-aircraft surveillance applications like conflict detection is inherently more complex than applying these requirements to single aircraft navigation applications, like precision approach and landing [15-17]. For the purposes of this chapter, we thus focus only on one specific example of conflict detection: the case of a CD algorithm monitoring for the loss of separation between two aircraft. In this example, a loss of continuity would occur if the sensors were to falsely indicate that a separation violation had occurred when the separation minimum were not yet violated. Conversely, a loss of integrity would occur if the sensors were to mask an actual separation violation until there was no longer enough time to properly issue an alert and resolve the conflict.

To compare CD algorithm performance to requirements, a detailed definition of continuity risk for CD algorithms is required. A loss of continuity occurs when normal operations are interrupted, as would occur in the case of a CD

alert. As such, we define continuity risk to be the probability that the CD algorithm issues an alert of any kind. Because true conflicts are anticipated to be very rare occurrences, false alerts are expected to dominate the budget for allowed continuity risk. As such, continuity risk can thus be quantified in terms of two related parameters: a monitor threshold and a continuity risk specification. The first parameter, the threshold, is a bound on the sensor noise level under fault free conditions. The second parameter, the allowed continuity risk, is a requirement on the frequency with which alarms may occur, generally given as a probability per unit time. To prevent overly frequent alerts, the threshold must be set to be sufficiently wide to ensure that the continuity risk requirement is met, given the noise level of available sensor measurements.

A detailed definition of integrity risk for CD algorithms is also needed to evaluate algorithm performance. We define this integrity risk to be the probability that a conflict occurs when no timely alert has been issued by the CD algorithm. In general, integrity is quantified in terms of three parameters: an alert limit, an operation-specific risk specification, and a time-to-alert. For the case of a conflict due to a loss of separation, the first parameter, the alert limit, is a bound on how close the measured separation approaches the separation minimum, given the uncertainty inherent in the surveillance sensor measurements. Second, the integrity risk specification is a parameter that limits the allowed probability that surveillance sensors miss detecting an unsafe condition (such as a loss of separation). The FAA requires this parameter to be a small number, such that any catastrophic event is extremely improbable, with an integrity risk below 10^{-9} per operation [18]. Third, the time-to-alert parameter is a means of quantifying a "timely" alert. An alert is considered to be a missed detection if it is late, meaning that the alert is not issued within the required time-to-alert needed for a successful conflict resolution.

Target values for the five continuity and integrity parameters (continuity threshold and integrity alert limit, continuity and integrity risk specifications, and time-to-alert) have not yet been established for the CD&R application; however, these values should be compatible with other planned and existing surface movement systems, such as the Advanced Surface Movement Guidance and Control System (A-SMGCS). Surveillance requirements for A-SMGCS surveillance in the movement area, under limited visibility conditions down to Runway Visual Range (RVR) of 75 m, are $2 \cdot 10^{-2}$/hour for system continuity and $2 \cdot 10^{-5}$/hour for system integrity (with a 10 second time to alert and a proposed 15 m horizontal alert limit) [13]. The implied threshold (based on a minimum one-sigma horizontal accuracy of 4 m) is approximately 12 m. Requirements in zero visibility conditions would likely be somewhat more strict. In theory, CD&R algorithms could be implemented as a sub-component of A-SMGCS. In this case, a fraction of the total continuity and integrity budgets for A-SMGCS would be sub-allocated to the CD&R function.

2.2 Integrating Continuity and Integrity into CD Algorithm Design

Accounting for Nominal Errors

This section considers a specific implementation of a CD algorithm, for detection of conflicts between two aircraft occupying the same runway, and refines the algorithm to account specifically for continuity and integrity requirements. In detecting conflicts between two aircraft occupying the same runway, there is a required minimum in-trail separation to be maintained between the aircraft. Once the CD algorithm determines that these aircraft no longer meet this separation requirement, an alert is triggered and the CD&R algorithm initiates conflict resolution. An example of

a scenario in which a loss of separation might occur is when an aircraft merges onto a taxiway or runway occupied by a second aircraft, as illustrated in Figure 2.

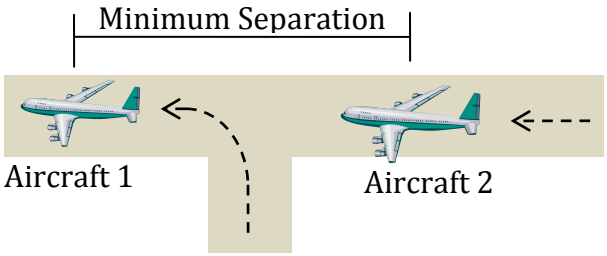

Figure 2. Aircraft Separation Required on Runway

In order to account for continuity and integrity requirements as defined in the preceding section, we must account for the five continuity and integrity parameters. We propose that the CD algorithm should explicitly incorporate the continuity threshold and integrity alert limit parameters in monitoring aircraft separation. The allowed continuity and integrity risk specifications are interpreted as a means to assess whether or not available surveillance sensors comply with the specified alert limit and threshold. In this chapter, we assume that surveillance sensors meet time-to-alert requirements, but this assumption could be relaxed in future research [19].

As a starting point, we identify two distinct separation requirements between aircraft: the minimum required separation S_{tot_req}, which accounts for risk factors, and the reduced minimum separation S, which would be the required separation in the absence of any sensor measurement error. Absent sensor error, S is clearly smaller than S_{tot_req}. The magnitude of S must be set to account for physical separation requirements, such as the dimensions of both aircraft and the buffer needed to allow for nonzero relative speed between aircraft.

We assume that it may be desirable to reduce procedural separation requirements in NextGen, to allow for tighter coordination among aircraft. Hence, we place no restrictions on the magnitudes of S_{tot_req} or S individually. The required difference between these values, however, can be related to our continuity and integrity requirements. In this section, we will show that a suitable relationship between S_{tot_req} and S, which protects continuity and integrity requirements, is the following.

$$S_{tot\_req} = S + T + U \qquad (1)$$

Here the continuity threshold is U and the integrity alert limit is T. The total separation requirement, S_{tot_req}, therefore becomes the sum of three parameters, as illustrated in Figure 3.

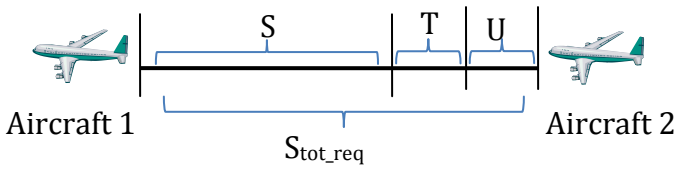

Figure 3. Illustration of Minimum Separation Requirements

Whereas the parameters S and S_{tot_req} refer to required separation minima, the actual separation between an aircraft pair will generally be somewhat larger than this minimum. We define the actual separation between aircraft to be S_{true}. To ensure safe operations,

$$S_{true} \geq S \qquad (2)$$

Of course the CD algorithm is not aware of the true separation distance between aircraft, as sensor measurements are inherently noisy. Accounting for a random sensor error ε, the measured separation distance between aircraft S_{meas} is perturbed from the true separation S_{true} .

$$S_{meas} = S_{true} + \varepsilon \qquad (3)$$

Maintaining a nominal separation of S_{tot_req} between aircraft introduces a margin to ensure that the safety requirement (2) is met, even in the presence

of sensor noise. It is the job of the CD algorithm to verify (2). To do this, the CD algorithm computes a monitor statistic *m*.

$$m = S_{tot\_req} - S_{meas} \tag{4}$$

The monitor statistic is defined to be positive when surveillance sensors detect a loss of separation. Hence the CD algorithm should trigger an alert if the loss of separation grows too large. The alert should not be issued immediately if *m* becomes positive, however, as sensor noise (for an unbiased surveillance sensor) would then have a 50% chance of triggering the alert were the separation at the allowed minimum of S_{tot_req}. Rather, the monitor should trigger at a threshold sufficiently large to limit the probability of false alarms. This operational consideration can be related to the continuity requirement; the alert should trigger if the loss of separation ever exceeds the continuity threshold.

$$\text{alert condition: } m > U \tag{5}$$

The probability of a false alarm P_{FA} is directly related to the distribution of errors ε for a given sensor and to the continuity threshold *U*. Specifically, P_{FA} is the conditional probability that the monitor statistic exceeds the threshold given that the true aircraft separation is satisfactory, meaning that the true separation is greater than the total required separation.

$$P_{FA} = P\left\{m > U \mid S_{true} - S_{tot\_req} \geq 0\right\} \tag{6}$$

Because the continuity requirement must be satisfied for all allowable values of true separation, it is possible to replace (6) with a conservative form that considers only the worst-case allowed true separation. For a unimodal, zero-mean error distribution, the worst case separation, which results in the highest possible value of P_{FA}, occurs when the true separation is a minimum, when $S_{true} - S_{tot\_req} = 0$. Using this limiting case in order to simplify (6), we obtain the following:

$$P_{FA} \leq P\left\{m > U \mid S_{true} = S_{tot\_req}\right\} \tag{7}$$

As long as sensors errors are sufficiently low that the expression on the right side of (7) does not exceed the continuity allocation for nominal random noise C_{NR}, then continuity is assured independent of the true value of aircraft separation.

$$\text{continuity requirement: } P_{FA} \le C_{NR} \qquad (8)$$

Whether or not this requirement can be satisfied depends on the choice of the threshold U and the exact set of surveillance sensors used at the airport facility (see next section).

To help visualize this continuity requirement, the probability distribution for sensor error is illustrated in Figure 4 for the worst-case conforming scenario, where $S_{true} = S_{tot\_req}$. For the purposes of the illustration, the sensor error distribution, $p(\varepsilon)$, is depicted as a Gaussian density function. The integrated probability over the shaded region of the density function is P_{FA}, which represents the total probability that the error magnitude exceeds the threshold U (in the direction of a false alarm.)

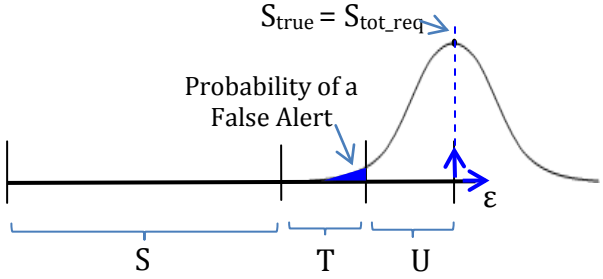

Figure 4. Probability of a False Alert for Worst-Case Conforming Scenario

Our CD algorithm design should also take into account integrity risk requirements. To meet integrity requirements, the probability that the safety criterion (2) is violated must be exceedingly small. In the absence of sensor noise, the safety criterion would always be met because the alert would be triggered at $S_{true} = S + T$, whereas a violation would only occur if S_{true} were to become smaller than S. The margin provided by T is essential to account for

sensor errors, however. A sufficiently large sensor error could cause a missed detection event, in which the CD algorithm assesses an otherwise unsafe situation as safe. To meet integrity requirements, the missed detection probability P_{MD} must be smaller than the integrity allocation for nominal random errors I_{NR}.

$$\text{integrity requirement: } P_{MD} \le I_{NR} \qquad (9)$$

The probability of a missed detection P_{MD} is directly related to the distribution of errors ε for a given sensor and to the alert limit T. Specifically, P_{MD} is the conditional probability that the monitor statistic does not exceed the threshold given that the true aircraft separation is unsafe, meaning that the true separation is smaller than the minimum separation in the absence of a sensor error.

$$P_{MD} = P\{m < U \mid S_{true} - S < 0\} \qquad (10)$$

Because the integrity requirement (9) must be satisfied for all unsafe values of true separation, it is possible to replace (10) with a conservative form that considers only the least-detectable case of a separation violation. For a unimodal, zero-mean error distribution, the least-detectable separation violation occurs when $S_{true} = S$. Using this limiting case in order to simplify (10), we obtain the following:

$$P_{MD} \le P\{m < U \mid S_{true} = S\} \qquad (11)$$

As long as sensors errors are sufficiently low that the expression on the right side does not exceed the integrity allocation I_{NR}, then integrity is assured independent of the true value of aircraft separation. To help visualize this integrity requirement, the probability distribution for sensor error is shown in Figure 5 for the least-detectable unsafe scenario, where $S_{true} = S$. In the figure, the P_{MD} is the integrated probability over the shaded region of the density function, which represents the total probability that the measured

separation is perceived to be safe (loss of separation less than *U*) given that the actual separation is insufficient.

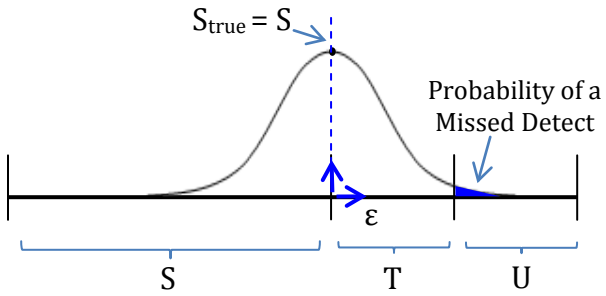

Figure 5. Probability of a Missed Detect for the Least Observable Unsafe Scenario

The salient feature of our proposed CD algorithm design is the significant margin defined between the unsafe state ($S_{true} < S$) and the desired conforming state ($S_{true} > S_{tot\_req}$). The intermediate region ($S_{tot\_req} > S_{true} > S$) is a "gray area," which is neither conforming nor unsafe. The defining characteristic of this region is that continuity and integrity requirements cannot be met within it. In other words, the monitor statistic noise is sufficiently large that the probabilities of a false alarm or of a missed detection occurring in this region are too high to support surface movement operations. The clear implication of the proposed approach is that the aircraft must be required to maintain a somewhat larger minimum separation (S_{tot_req}) than the true safety minimum in the absence of sensor error (S) such that the monitor threshold can be placed at an appropriate location between the two. A corollary is that high quality sensors are needed to permit dense aircraft spacing.

Accounting for Off-Nominal Events

Surveillance sensor performance can be divided into nominal and off-nominal conditions. Nominal performance refers to cases in which the surveillance sensors are operating as intended. The probability density function defined in the previous section (to set *U* and *T*) was assumed to describe sensor errors under nominal conditions.

Off-nominal conditions refer to any non-ideal situation in which the sensor fails to operate as intended. Because off-nominal events are rare exceptions, it can be challenging to account for the specific details of off-nominal events in terms of a probability density function. Hence, even though off-nominal events might be combined into an overall probability distribution for surveillance sensor errors, this distribution would be highly uncertain and difficult to validate rigorously. A conservative alternative is to model all off-nominal events as potential faults causing either a loss of continuity, integrity, or both. In this case, the total continuity budget for the CD algorithm C_{CD} is split between two sub-allocations, one for nominal random events C_{NR} and one for off-nominal events C_{OE}.

$$C_{CD} = C_{NR} + C_{OE} \tag{12}$$

The consequences of this decomposition of nominal and off-nominal events are two-fold. First, the fraction of the continuity risk budget available to bound nominal errors, according to (8), is reduced, placing a somewhat increased demand on sensor accuracy. Second, the probability of all off-nominal events must be summed; this total probability of an off-nominal event must be shown to be smaller than C_{OE}.

Similarly, the total integrity budget for the CD algorithm I_{CD} must be split between an allocation to cover nominal random errors I_{NR} and a second allocation to cover undetected off-nominal events I_{UOE}.

$$I_{CD} = I_{NR} + I_{UOE} \tag{13}$$

It is important to note that only undetected off-nominal events count against the integrity budget. It is possible that some off-nominal events might be detected, and those cases of faulty sensor measurements would be excluded without being used by the CD algorithm. Specific mechanisms to implement such a fault-exclusion logic are not considered in this thesis.

As in the continuity risk case, the consequences of decomposing the nominal and off-nominal integrity risk allocations are two-fold. Again, only a fraction of the total integrity budget is available for bounding nominal errors, according to (9). Additionally, the total probability of all undetected off-nominal events must be summed; this total probability must be shown to be smaller than the allowed risk I_{UOE} in order for integrity requirements to be met.

2.3 Role of Surveillance Sensors in CD&R System Requirements

A detailed analysis of airport surveillance sensors is required to demonstrate whether or not a particular CD algorithm can meet its continuity and integrity requirements. Though such a detailed analysis is beyond the scope of this thesis, we will provide a brief overview of airport surveillance sensors and their performance, in order to outline how their nominal and off-nominal errors might be mapped into continuity and integrity requirements, as described in the preceding section.

Nominal performance is typically described by the accuracy of the reported position (and in a more general CD algorithm, by the accuracy of velocity information), as well as the frequency at which this information is reported. In general, accuracy is the nominal scatter of navigation errors, which for many sensors, is well described by a Gaussian distribution, at least near the distribution core [20]. Accordingly, sensor accuracy is most often specified as a Gaussian sigma value that describes the distribution's standard deviation. Typically, non-Gaussian tail behavior need only be considered in computing the false alert and missed detection probabilities (Figure 4 and

Figure 5, respectively) if the alert limit T and continuity threshold U are much larger than 2-3 times sigma.

Off-nominal performance refers to conditions that can interrupt a sensor's normal function. These occurrences are not well characterized by published accuracy information. In the most extreme cases, off-nominal behavior could result from an unrecoverable hardware failure (continuity loss) or from a partial hardware failure that causes a sensor to produce hazardous misleading information (integrity loss). Other events may be somewhat more common and somewhat less hazardous, such as a one-time sensor transmission or reception error. When considering all possible off-nominal events, it is important to realize that some events will impact both continuity and integrity requirements, while others may only impact one of these two.

In the remainder of the chapter, we will briefly overview current and emerging surveillance sensor technologies, in order to define a rough framework for how to assess the continuity and integrity impact of nominal random errors (C_{NR} and I_{NR}) and off-nominal errors (C_{OE} and I_{UOE}).

Current surveillance systems for detecting aircraft on the airport surface include radar, multilateration, and ADS-B. Radar technology has been around since the 1940s while more recently, multilateration systems have been approved for airport use since 2003 [21]. ADS-B is an emerging surveillance system in the process of being implemented and scheduled for complete implementation by 2020 [22]. This system uses positioning information, typically provided by onboard GPS. The integrity of this GPS sensing capability can be established through Receiver Autonomous Integrity Monitoring (RAIM), the Wide Area Augmentation System (WAAS), or, in the near future, the Ground Based Augmentation Systems (GBAS), which is scheduled to be fully operational sometime after 2016 [23,1].

2.4 Minimum Requirements for Threshold and Alert Limit Based on Sensor Capabilities

Surveillance Sensor Capabilities

It is perhaps most useful to start by characterizing the error distributions for airport surface surveillance systems in terms of their 95% accuracy. The 95% accuracy corresponds to a 2-sigma value for a Gaussian distribution model. Nominal position and velocity measurement error for surface radar, multilateration, and ADS-B are summarized in Table 1 below. These accuracy numbers are obtained from requirements documents and/or field assessments for each system [24-28]. A more detailed description for each system is provided in the appendices.

Surveillance System	Position Accuracy	Velocity Accuracy	Update Rate
Radar	2 m	1 m/s	1 Hz
Multilateration	6 m	0.25 m/s	1 Hz
ADS-B (w/ SBAS)	2 m	0.1 m/s	1 Hz (0.2 Hz if stopped)
ADS-B (w/ GBAS)	1 m	0.1 m/s	1 Hz (0.2 Hz if stopped)

Table 1. Surveillance Sensor Accuracy

For typical airport facilities, more than one of these sensing systems is available (or will be) to a sensor fusion processor. Surface surveillance data at major airports in the Unites States is currently compiled by the Airport Surface Detection Equipment – Model X (ASDE-X), a system produced by Sensis to integrate available sensor data for display to air traffic controllers. The fusion process can improve the overall accuracy of the combined surveillance data.

Threshold and Alert Limit Determination

We can use the accuracies provided in Table 1 as a basis to estimate minimum values for the continuity threshold *U* and integrity alert limit *T* for the proposed CD algorithm. A rigorous analysis would require a detailed characterization of the sensor error probability distribution; such a characterization is beyond the scope of this work, however. Instead, for the purposes of making a rough estimate of *U* and *T*, we will make several assumptions, some which are conservative and some which are not.

Importantly, we assume that the surveillance sensor errors are zero-mean Gaussian and independent for each aircraft. In this case, the separation error (that results from summing two Gaussian random variables) is $\sqrt{2}$ times the error of the aircraft position estimates produced by the sensor fusion system. Rather than model the sensor fusion process in detail, we conservatively estimate that the accuracy of the fused position estimate is equal to the accuracy of the best available sensor (as characterized by Table 1). Noting that the Gaussian sigma value is approximately half the 95% accuracy, the separation error is modeled as zero-mean Gaussian with a sigma value equal to $\sigma_\varepsilon = \sqrt{2}/2\sigma_{table}$, where σ_{table} is the entry in Table 1 associated with the most accurate available sensor technology. This separation error model feeds into computation of the probabilities of a false alarm or missed detection due to nominal random noise, according to (7) and (11). In the limiting cases of (8) and (9), which result in the minimum possible values for *U* and *T*, the false alarm probability P_{FA} can at most be equal to the continuity allocation C_{NR}, and the missed detection probability P_{MD} can at most be equal to the integrity allocation I_{NR}. By evaluating (7) and (11) at these limits we obtain the following equations for *U* and *T*.

$$U = -Q^{-1}(C_{NR}) \cdot \sigma_\varepsilon \qquad (14)$$

$$T = -Q^{-1}(I_{NR}) \cdot \sigma_\varepsilon \qquad (15)$$

Here Q is the Gaussian cumulative density function (CDF) for a unity-variance, zero-mean distribution, and σ_{est} is accuracy of the available sensor, such as the values for position accuracy given in Table 1.

In order to evaluate these equations, values for both nominal continuity and integrity risks, C_{NR} and I_{NR}, are required. The specifications for A-SMGCS (see Section 2.1), $2 \cdot 10^{-2}$/hour for system continuity risk and $2 \cdot 10^{-5}$/hour for system integrity risk, provide representative upper bounds for the requirements of a CD algorithm. These requirements are given on a "per hour" basis, however, and so must be related to the probabilities of a missed detection or a false alert occurring at any particular instant in time. To relate per-hour specifications to the values of C_{NR} and I_{NR} used in (14) and (15), we assume the following. First, we assume that fused sensor errors are highly correlated over the course of any operation with significant risk (such as landing, taking-off, merging, or crossing an active runway). If error values are correlated over an operation, then we can consider missed detection and false alarm probabilities on a *per operation* basis rather than on a *per time-step* basis. Second, we assume that for a high capacity airport there are 60 operations per hour in which aircraft separation approaches close to the minimum (and in which the monitor might flag, either as a false or true alarm). Dividing the total integrity and continuity budget per hour by 60 operations/hour gives the following

$$C_{CD} = \tfrac{1}{60}\ 2 \cdot 10^{-2}/operation \qquad (16)$$

$$I_{CD} = \tfrac{1}{60}\ 2 \cdot 10^{-5}/operation \qquad (17)$$

The total integrity and continuity allocations must be split between nominal and off-nominal events according to (12) and (13). In this light, we make one final important assumption: that C_{OE} and I_{UOE}, continuity and integrity risks for off-nominal events, are allotted half of the overall total continuity and

integrity risk budgets. With this, we obtain the following equation for both C_{NR} and I_{NR}:

$$C_{NR} = 0.5 \cdot C_{CD} \quad (18)$$

$$I_{NR} = 0.5 \cdot I_{CD} \quad (19)$$

Therefore, combining these equations and using the specs given in Table 1, we are able to obtain a very rough estimate of the minimum threshold and alert limit values, U and T. These estimated values, summarized in Table 2, provide a sense of how large the integrity and continuity bounds might be for a practical application.

Best Available Surveillance Technology	Continuity Threshold, U	Integrity Alert Limit, T
Multilateration	15.2 m	21.7 m
Radar	5.1 m	7.2 m
ADS-B with GBAS	2.5 m	3.6 m

Table 2. Threshold and Alert Limit Table

2.5 Contribution of Individual Sensor Systems to Off-Nominal Risks

Contribution of Individual Sensors

Off-nominal sensor performance refers to anomalous sensor behaviors that impact continuity and integrity, either by causing a missed detection or contributing to an alarm. Examples of these events include component failures, signal occlusion, missed messages, and improper time tagging. Here we briefly outline the primary off-nominal behaviors for each sensor of interest (radar, multilateration and ADS-B) in order to construct continuity and integrity risk trees that account for these behaviors. These lists are

preliminary, and are not intended to substitute for a thorough failure modes and effects analysis.

For radar, off-nominal events include signal occlusion, tracking errors, rain and ground clutter interferences, plot extraction errors, and radar tower failures. Signal occlusion refers to instances where the radar signal is blocked by buildings or other aircraft, which often happens in airport areas near the terminal gates. Tracking errors occur since it is possible that surface radar fails to associate targets with the correct track, which is especially true for ground vehicles which often travel out of coverage areas, leading to gaps in the collected data [29]. Additionally, rain and ground clutter can also cause interference for radar, often leading to false alarms for target detection. While surface radar is designed to function in rain falling at a rate of 16 mm/hr, rainfall rates are sometimes significantly higher than this, corresponding to the very cases of extreme weather when radar is needed most [30-31]. Another off-nominal event is an extraction error, in which radar signatures for multiple aircraft are merged into a single track or when signatures for a single aircraft are incorrectly split into multiple tracks [32]. A full description of radar and its impact on CD&R can be found in Appendix A.

With multilateration systems, off-nominal events include anomalous transponder location bias, missed message errors, aircraft transceiver errors, ground equipment failures, signal spoofing, signal jamming, message collisions, and ground clock synchronization error. In terms of the anomalous transponder location bias, multilateration systems solve for the transponder's position on the airport, so there is an additional bias error in the position measurement due to the unknown location of the transponder on the aircraft. In off-nominal cases, the transponder may be located very far from the location assumed by the surface surveillance system. Missed

message errors refer to instances where a message is successfully transmitted by an aircraft but fails to reach the ground station receiver, often due to blockage by objects or buildings. Equipment failures include any potential problem with the aircraft transceiver or with the ground station receiver. Signal spoofing occurs when a fake signal is broadcast by a hostile agent, introducing false position information. Signal jamming can occur when an outside source intentionally or unintentionally interferes with the aircraft's signal transmission, preventing the ground stations from receiving multilateration messages. Similarly, message collisions occur when two transmitters send signals at the same time, a specific case of "unintentional jamming". Finally, ground station clock synchronization errors could occur if the ground station clocks were not properly synchronized, which is an important requirement for multilateration systems, which compute aircraft position based on a time-difference of arrival [33]. A complete description of multilateration can be found in Appendix B.

ADS-B off-nominal events include improper time tagging, missed messages, aircraft transceiver errors, ground equipment failures, signal spoofing, signal jamming, message collisions, and GPS faults. Transponder bias is not a significant issue for ADS-B, since ADS-B requirements place much tighter constraints on transponder location than do multilateration requirements [34]. Improper time tagging could occur if the ADS-B message were broadcast with an imprecise estimate of the time at which GPS position data were acquired, an event which would mean result in the surveillance system measurement being biased by an amount corresponding to the rate (and direction) of aircraft travel. In addition, missed messages, aircraft transceiver errors, ground equipment failures, signal spoofing and jamming, and message collisions can all occur in the same manner as described previously for multilateration. Lastly, GPS faults include code-carrier divergence, excessive satellite clock acceleration, mild ionosphere storms,

and satellite signal deformation, which, while very rare, can have a significant impact on the provided GPS position. A complete description of ADS-B and the GPS systems it relies on for positioning can be found in Appendix C.

These off-nominal events for radar, multilateration, and ADS-B are summarized in the continuity and integrity risk trees shown in Figure 6 and Figure 7. These trees provide a visual aid for how these occurrences can impact continuity and integrity, as the probabilities for each event occurring combine to define the required bounds. From these trees, it is clear that each sensor contributes differently to the cumulative off-nominal integrity and continuity risk. Therefore, it is necessary to understand all aspects of these surveillance systems, in order to establish bounds for continuity and integrity for CD&R.

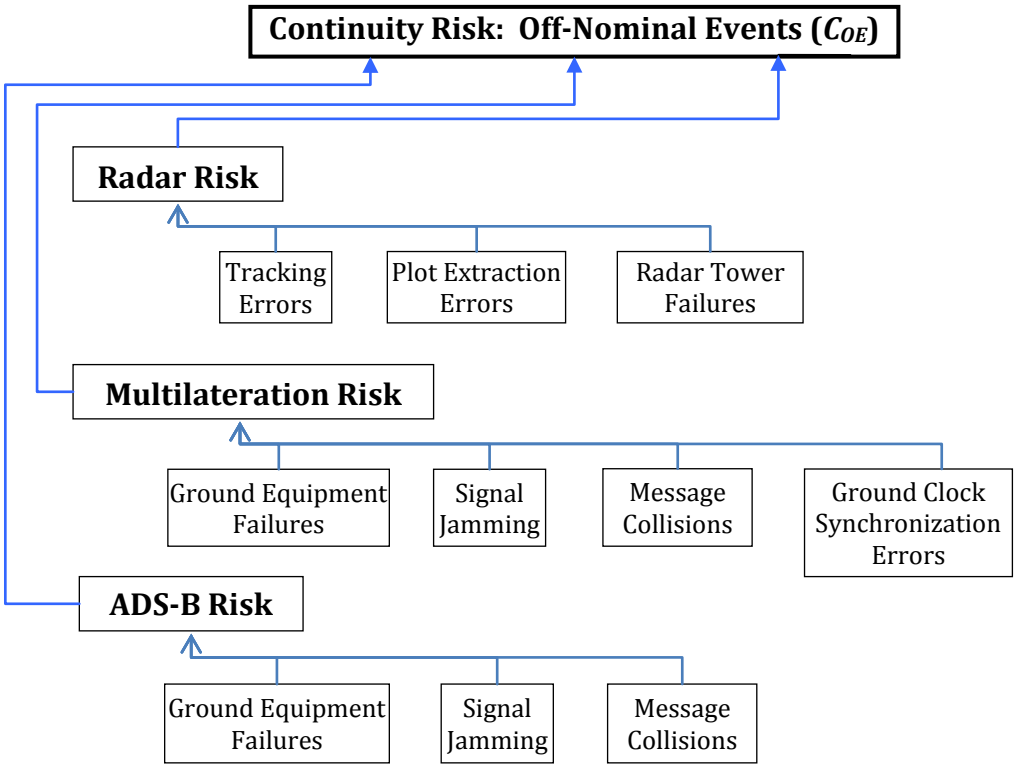

Figure 6. Continuity Risk Tree for Off-Nominal Events of Surveillance Sensors

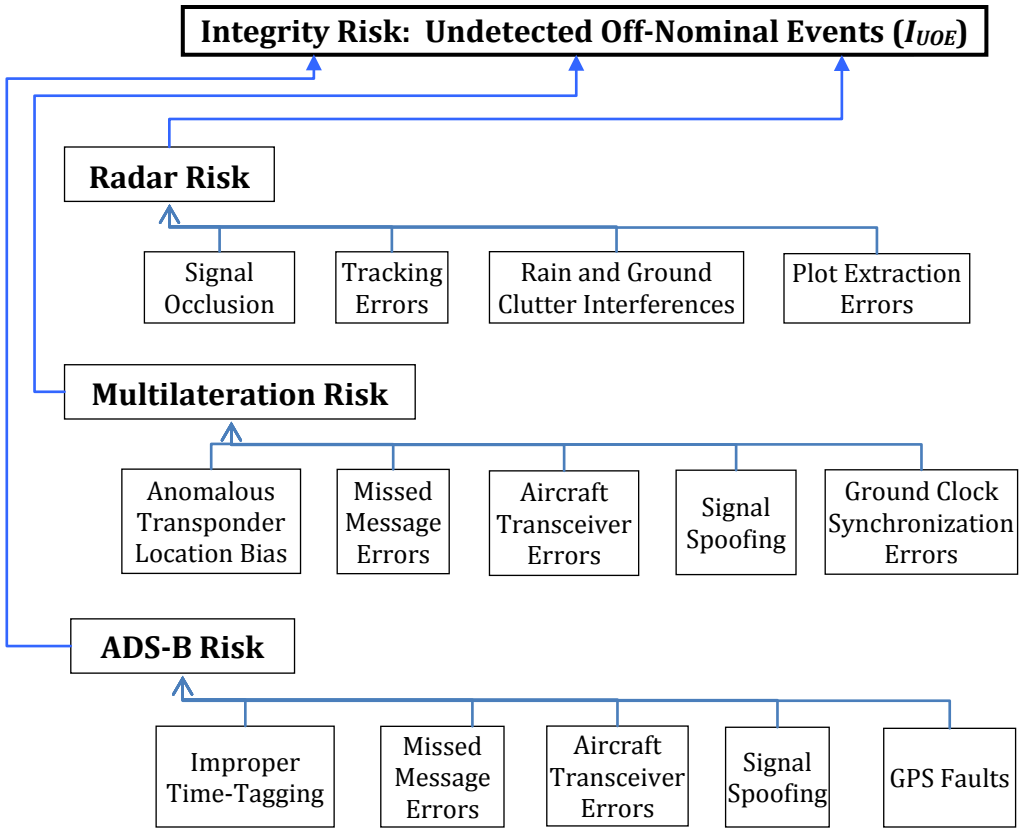

Figure 7. Integrity Risk Tree for Off-Nominal Events of Surveillance Sensors.

Tallying Continuity and Integrity Risks for Off-Nominal Events

Risk trees can be used to determine verify the continuity and integrity risk allocations for off-nominal events, C_{OE} and I_{UOE}, respectively. For the most conservative case, where the CD algorithm has no fault detection and exclusion capabilities, off-nominal continuity and integrity risks would be the sum of the probabilities of the occurrence of each off-nominal event. The total off-nominal integrity risk per operation should not exceed the budget for off-nominal events C_{OE}. This requirement is represented by the following equation, where $C_{i,j}$ denotes the probability per operation of a particular off-nominal event j, for a sensors i used in the sensor fusion system.

$$C_{OE} \geq \sum_i \sum_j C_{ij} \qquad (20)$$

Realistically, the sensor fusion system will incorporate fault detection and exclusion probabilities that reduce the risks associated with any particular fault. Rather than model any particular implementation of a sensor fusion system, we simply represent the sensor fusion operation by a general function *f*.

$$C_{OE} \geq f(C_{1,1}, \quad C_{1,2}, \ldots, C_{i,j}) \qquad (21)$$

By analogy, the combined off-nominal integrity risk for all sensors might also be described in terms of some general function g operating on the per-operation integrity risk probabilities $I_{i,j}$ associated with each fault mode j for any sensor i. The total per-operation integrity risk for off-nominal events, as described by the function g, should not exceed the integrity budget for undetected off-nominal events, I_{UOE}.

$$I_{UOE} \geq g(I_{1,1}, \quad I_{1,2}, \ldots, I_{i,j}) \qquad (22)$$

2.6 Conclusions

There are two key contributions from this chapter. The first is that a margin is needed in between the separation minimum for safety and the nominal required separation among aircraft. The CD threshold should be placed at an appropriate location between these two extremes in order to ensure that continuity and integrity specifications are met. The second contribution is that off-nominal surveillance sensor behaviors should be accounted for when designing a CD&R algorithm. Since this is hard to do in terms of a probability distribution, a suitable conservative alternative to use is risk trees which outline each sensor's impact on continuity and integrity requirements for CD&R.

Specifically, this chapter outlines how continuity and integrity requirements can be met by a sample CD algorithm which monitors the longitudinal

separation distance between two aircraft located on the same runway. However, this algorithm is based on a generic description of sensor performance and, in order to quantify continuity and integrity requirements for CD&R algorithms, it is necessary to understand both nominal and off nominal sensor performance capabilities. Therefore, this chapter also outlines nominal and off-nominal characteristics for surface surveillance sensors such as radar, multilateration, and ADS-B. Overall, when designing CD&R algorithms, continuity and integrity are important specifications which require a complete understanding of off-nominal sensor performance, and which should be considered from the very early stages of CD&R algorithm design.

Chapter 3: Integrity Monitoring of Vision-Based Automotive Lane Detection Methods

According to the Transportation Research Board, there are more than 1.2 million road-departure crashes in the United States each year [3]. Studies suggest that the implementation of LDW systems in the US would lead to a 10% decrease in passenger car road-departure crashes, and a 30% decrease in truck road-departure crashes [4]. In the future, LDW technology could be adapted to enable a new safety feature: fully automated lane-keeping. Such a form of automated driving could allow human operators to perform other tasks than driving, such as texting or reading, without the fear of causing an accident.

A wide range of LDW concepts have been demonstrated in a research setting, but most of these concepts rely on an external infrastructure which would be expensive to deploy and maintain. At the University of California PATH Program, for example, magnetic markers were embedded in the road and used with a magnetometer mounted to the front of the car [35]. While this is a very reliable system, the implementation cost is too high to be practical in all roads. Another system, implemented at the IV Lab at the University of Minnesota, combines differential GPS information with high resolution digital maps to determine lane position and is demonstrated to work in all weather conditions [36]. However, this system requires digital maps with under a meter of accuracy which currently do not exist at wide enough level to make this system practical. Among the technologies which do not require extensive infrastructure are infrared or camera-based sensors. For infrared-based systems, infrared sensors are mounted to the bottom of the car to detect the reflectivity difference between lane lines and bare pavement. While this system has the advantage of being unaffected by poor weather

conditions, it can only detect drift after it has occurred [37]. Lastly, camera-based systems process video data produced by imaging equipment, typically mounted in the crew cabin, to identify lane lines. Unlike infrared-based systems, camera-based systems can predict when a future lane-crossing will occur.

This chapter focuses on camera-based lane-boundary detection systems, because (1) camera-based LDW is the only technology which has been commercialized to date and (2) because camera-based sensors provide an extremely rich stream of data. Due to the density of measurements available – at each camera pixel – not all of the data is required for any single vision processing algorithm. Thus it is possible to implement vision processing algorithms that accomplish the same goal (lane-boundary detection) using fundamentally different signals extracted from the video stream.

By implementing vision processing algorithms based on different components of the video stream, it is possible to perform a cross-check to verify measurement quality. This type of cross-check, which is often called an *integrity monitor* in the navigation community, is essential to ensuring measurement quality in safety-critical automation functions, such as automated driving. Human drivers are highly competent at fusing data from different sources and identifying sensor anomalies; computers used in automated functions, by contrast, do not naturally detect and exclude anomalous sensor data. Sensor anomalies may thus produce unexpected driving behaviors which might quickly result in an accident and possibly, loss of life.

With the goal of ensuring the integrity of video-based lane boundary detection, this chapter presents an integrity monitoring algorithm that compares two largely independent vision processing algorithms. The first

algorithm, referred to as the Lane Boundary (LB) algorithm, was previously defined by Jung and Kelber [9]. The LB algorithm fits a linear-parabolic lane model to high gradient values within regions of interest identified for each video frame. The second algorithm, which we introduce in this chapter, is referred to as the Optical Flow (OF) algorithm. It assumes that all image pixels reside in the plane of the road, and identifies non-road pixels as those whose intensity changes over time are inconsistent with that assumption. The two algorithms are largely independent because they operate on fundamentally different information: the LB method uses spatial gradients while the OF method uses temporal changes.

3.1 Automated Lane-Keeping Architecture

In this chapter, we focus on a particular aspect of automated driving: automated lane-keeping, in which a computer control algorithm attempts to hold a car in the center of its lane. We assume that the lane-keeping algorithm relies on a camera-based sensor for precise lateral positioning. Camera-based sensors are highly appropriate for lane keeping, because they can achieve a 95% accuracy better than 30 cm [38-39], which is commensurate with the requirements for automated lane-keeping. By comparison, automotive navigation packages available today in North America (equipped with WAAS-aided GPS) only achieve a 95% accuracy of approximately 2 m.

Sensor fusion is needed to obtain both the relative and absolute position information needed for automated lane keeping. SBAS-aided GPS can provide absolute information while vision-based algorithms can provide lateral positioning, relative to lane boundaries at the edges of a road. Therefore, we anticipate that future automated driving systems will combine SBAS-aided GPS with camera data in a Kalman Filter (Figure 8).

In automated lane-keeping, lane boundaries determined inaccurately may lead to an unintentional lane departure and possible loss of life. Therefore, it is critical that sensor measurements are verified before being fused. WAAS/SBAS automatically provides verification for GPS measurements, but no such verification is automatically provided by vision-based positioning. The need for integrity monitoring for visual navigation measurements has only recently been considered [40]. The central contribution of this chapter is to identify specific fault modes for a different application (automated lane keeping) and to develop an application-specific approach for verification of vision-based measurements (i.e., *integrity monitoring*).

Specifically, our concept for a vision-based system is as follows: use a primary algorithm (LB) to detect lane position, and a secondary algorithm (OF) to verify measurement quality. The strengths and weaknesses of these two algorithms are complimentary. The LB algorithm is more accurate than the OF algorithm, but requires image features to be tracked over multiple time steps (and may therefore diverge if a false positive occurs). For these reasons, we propose a system-level approach that relies on the LB algorithm for navigation, due to its lower noise, but which exploits the OF algorithm for integrity monitoring, to detect anomalies in the LB output. Such an approach is shown in Figure 8.

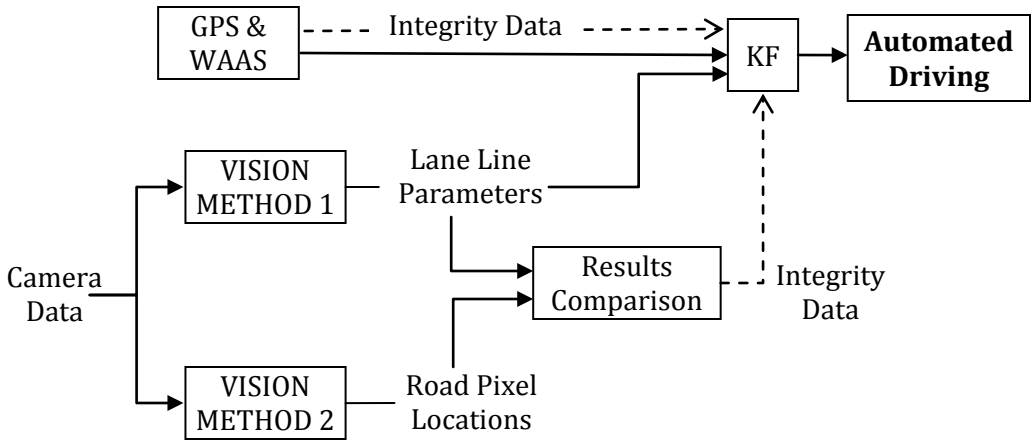

Figure 8. Concept for Automated Lane-Keeping Algorithm. Vision algorithms are needed to augment WAAS aided GPS for precision driving applications, such as lane-keeping. It is likely that data from GPS and vision sensors will be fused in a Kalman Filter (KF); to verify the quality of vision measurements, we recommend that a secondary vision processing algorithm be implemented to cross-check the first.

3.2 Gradient-Based Lane Boundary Detection Algorithm

The LB algorithm is a modified version of a method first introduced by Jung and Kelber [9]. In this algorithm, the lane markers are assumed to lie along a contour that is linear in the region closest to the front of the car and parabolic in the region further away from the car. This model is used because, while a purely linear model would be easier to implement, the combined linear-parabolic model represents curved roads more accurately.

The LB algorithm consists of three sequential processing steps. The first step identifies Regions of Interest (ROI) for subsequent processing. Considering only pixels in the ROI reduces processing costs and sensitivity to false detections of lane-markers in other parts of the image. The second step calculates the gradient values for pixels within the ROI, while the third step applies a Weighted Least-Squares (WLS) fit to this gradient data to

determine the coefficients of the piecewise linear-parabolic contour. These lane line parameters are also then used to build the ROI for processing the next frame. This algorithm is graphically summarized by the block diagram shown in Figure 9.

The LB algorithm makes several assumptions. First, it is assumed that lane markers are present and can be identified by high contrast with the road,

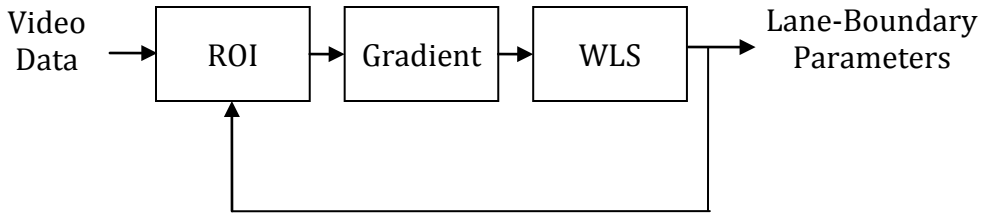

Figure 9. Linear-Parabolic Algorithm. The algorithm extracts Regions of Interest (ROI) from a video frame. Local gradient magnitude is computed at each pixel in the ROI. Subsequently, the best linear-parabolic model of the land boundary is computed using Weighted Least-Squares (WLS) fit to the gradient data in the ROI. These lane-boundary parameters can be used to generate a lane departure warning (and also determine the ROI for the next video frame.

which, in turn, results in high brightness gradients at the edges of the lane markers. Also, it is assumed that lane lines are located in the bottom half of the image and that the lane lines for each new frame will be found in approximately the same position as those from the previous frame.

ROI Identification: For each new image frame, the first processing step is to extract pixels in each of two regions, one ROI associated with the lane markers at the left boundary of the road and a second associated with the lane markers at the right. The centerline of the ROI is based on the lane-marker contour computed from the previous frame. In our algorithm, the ROI consists of a fixed number of pixels evenly distributed to the right and left side of this centerline (20 pixels on either side). The output of the ROI identification is a set of pixel locations in the vicinity of the left and right lane

markers, as shown in Figure 10. As the figure shows, our implementation of the LB algorithm uses an ROI that is approximately two to three times wider than the lane markers. If indeed the ROI contains the lane lines, only a small portion of the image data needs to be processed.

Gradient Filtering: Within each ROI, the gradient magnitude ||g|| is computed for each pixel as the norm of the gradient vector, which consists of the lateral gradient g_x and the vertical gradient g_y, given by equation (23). Throughout

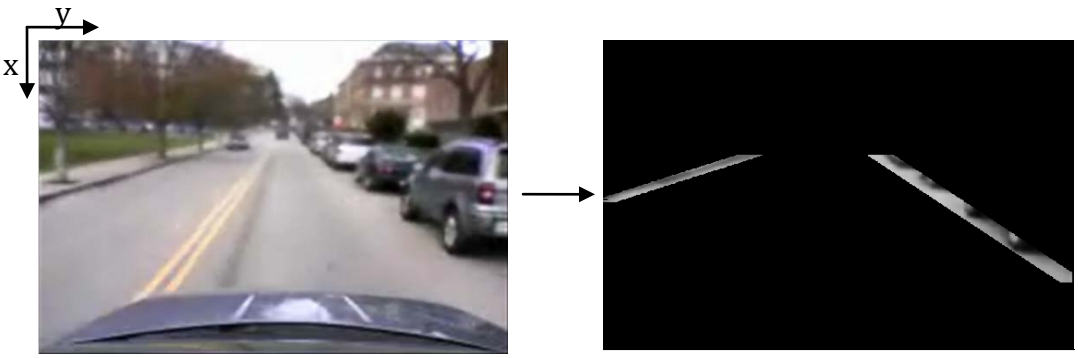

Initial Frame 1

Regions of Interest with Corresponding Pixel Values

Figure 10. Regions of interest for lane detection. Image size is 680x852

the LB algorithm, *x* and *y* are given by the axes shown in the first image of Figure 10.

$$\|g\| = \sqrt{g_x^{\,2} + g_y^{\,2}} \qquad (23)$$

The *x* and *y* components of the gradient vector are calculated by finding a simple pixel difference with lateral and vertical neighbors. Here the variable *p* denotes a pixel brightness value.

$$g_x(x,y) = p(x+1,y) - p(x,y) \qquad (24)$$

$$g_y(x,y) = p(x,y+1) - p(x,y) \qquad (25)$$

It is assumed that the strongest pixel values in the region of interest are associated with the lane boundaries. However, many road features

(oncoming traffic, street lamps, or other structures) introduce mild gradients. For this reason, a threshold is applied to gradient magnitude. Gradient magnitude values less than half the average in the ROI are thrown out. The remaining pixels are called the *high-gradient set*.

WLS Computation: The next step of the LB algorithm involves determining the lane line parameters. The centerline of the lane boundary pixels is modeled as a contour which consists of a straight segment (close to the camera) and a segment which may be curved (in the far field). The centerline coordinates are computed from the high-gradient set using WLS. The straight line segment is a best fit to the locations of the pixels in the high-gradient set below a transition row x_m. The parabolic segment is a best fit to the high-gradient pixel locations above the transition row. For our implementation (with images of size 480x720), the transition row x_m is 300, counting from an origin that lies in the upper left corner of the image. The formula for the linear-parabolic contour is the following.

$$f(x) = \begin{cases} a + bx & x > x_m \\ (1/2)(2a + x_m(b-d)) + dx + \dfrac{(b-d)}{2x_m}x^2 & x \le x_m \end{cases} \qquad (26)$$

Here, a, b, and d are the contour parameters to be determined by the WLS. The first two parameters describe the offset and slope of the linear region. The last parameter determines the curvature of the parabolic region.

The weighted least-square fit minimizes the squared-error function E, which describes the weighted distance between the contour and each pixel in the high-gradient set. Distances are weighted by gradient magnitude, such that stronger gradients (corresponding to the edges of lane markers) exert a greater "pull" on the contour.

$$E = \sum_{k=1}^{K} \|g_k\| \left(y_k - f(x_k)\right)^2 \qquad (27)$$

In this expression for error, the total number of nonzero pixels in the high-gradient set is K; these pixels are stored in a one-dimensional list and referred to by the index k.

A Weighted-Least Squares (WLS) fit minimizes the error expression given by equation (27) [41]. A matrix equation for the WLS solution is:

$$c = \left(\mathrm{A}^{\mathrm{T}}\,\mathrm{W}\mathrm{A}\right)^{-1}\mathrm{A}^{\mathrm{T}}\,\mathrm{W}y \tag{28}$$

Here the column locations of the K pixels in the ROI are listed in the vector y. The parameters describing the lane-boundary contour are components of the vector c.

$$c = \begin{bmatrix} a & b & d \end{bmatrix}^{\mathrm{T}} \tag{29}$$

The A matrix expresses the equation for the linear-parabolic contour, equation (26), for each point in the ROI. In the equation for the A matrix below, the variable n denotes the number of pixels in the near field (with $x > x_m$).

$$A = \begin{bmatrix} 1 & x_1 & 0 \\ \vdots & \vdots & \vdots \\ 1 & x_n & 0 \\ 1 & \frac{1}{2x_m}(x_{n+1}^2 + x_m^2) & \frac{-1}{2x_m}(x_{n+1} - x_m)^2 \\ \vdots & \vdots & \vdots \\ 1 & \frac{1}{2x_m}(x_K{}^2 + x_m{}^2) & \frac{-1}{2x_m}(x_K - x_m)^2 \end{bmatrix} \tag{30}$$

The weighting matrix W is a diagonal matrix, with gradient magnitude values for each of the K pixels in the ROI along the diagonal.

$$W = \begin{bmatrix} \|g_1\| & & 0 \\ & \ddots & \\ 0 & & \|g_K\| \end{bmatrix} \tag{31}$$

By solving (28), it is possible to obtain the lane-boundary parameters. These parameters are used to determine the location of the ROI for the next frame, as well as to obtain the car's position and heading relative to lane lines for the purposes of navigation. Figure 11 shows an example of these lane line results mapped to the original frame.

Figure 11. Lane Boundary Method Results. The calculated lane lines are shown in red.

Initialization: It should be noted that the LB method requires an accurate initialization step in order to define an initial ROI. In our implementation of the algorithm, a starting point is selected based on two assumptions: (1) that the lane lines in the first frame are essentially linear and (2) that these linear parameters can be determined by searching for regions of high contrast in the lowest section (bottom 40%) of the frame. A threshold is applied to this region to find the highest gradient values; subsequently a WLS fit is applied to find an initial set of linear lane-boundary parameters. An example of initialization is shown in Figure 12.

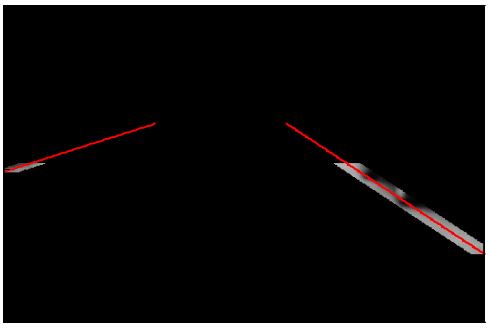

Initial Frame Filtered with Lane Lines

Initial Frame with Lane Lines

Figure 12. LB Initialization. Top: ROI Image with initial linear lane lines in red. Bottom: Initialized lane boundary model (red) superposed on first image.

3.3 Optical Flow Based Lane Detection Method

Optical flow is typically used to compute the flow field between two images. However, the algorithm we present here inverts the conventional optical flow equations to predict the next image in a video sequence. In this case, it is possible to use a known car velocity to advance the previous frame of camera data to compare it with the current video frame. We refer to this prediction as the velocity-warping process. If it is assumed during this warping process that all initial image pixels lie on the road plane, then only the actual road pixels will be accurately warped. Therefore, by comparing the warped image to the actual next image, we will be able to determine which pixels lie in the road. The algorithm for this method is outlined in Figure 13. The output of the OF algorithm is ultimately used to verify the integrity of the first algorithm (Figure 8).

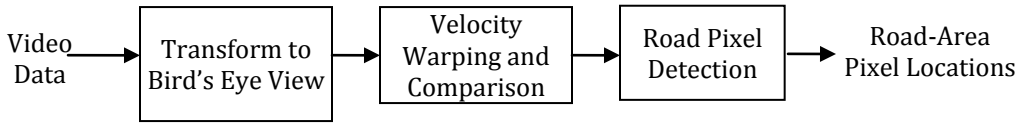

Figure 13. Optical Flow Algorithm. The algorithm converts the camera data to birds eye view, warps this image to a predicted next image using the car's velocity, compares the predicted next image to actual camera data at that time step, and applies threshold values to determine which pixels lie on the road plane.

Bird's Eye View Conversion: To simplify the velocity-warping equations, our OF method transforms forward-looking camera images to a bird's eye view. The bird's-eye transformation treats all points in the image as lying on the ground plane (even though they do not), so that subsequent processing can check the consistency of this assumption. The relationship between the pixels seen in the forward-looking image and the physical locations (x and y) of those pixels, mapped on to the ground plane, is obtained by an inverse perspective mapping operation [42].

$$x(r) = h\left(\frac{1+\left[1\text{-}2\left(\frac{r\text{-}1}{m\text{-}1}\right)\right]\tan(\alpha_v)\tan(\theta_0)}{\tan(\theta_0)-\left[1-2\left(\frac{r-1}{m-1}\right)\right]\tan(\alpha_v)}\right) \quad (32)$$

$$y(r,c) = h\left(\frac{1+\left[1\text{-}2\left(\frac{c\text{-}1}{r\text{-}1}\right)\right]\tan(\alpha_u)\tan(\theta_0)}{\sin(\theta_0)-\left[1-2\left(\frac{r-1}{m-1}\right)\right]\tan(\alpha_v)\cos(\theta_0)}\right) \quad (33)$$

where,

h: height of camera, with respect to ground plane
r: row position for pixel in initial image
c: column position for pixel in initial image
m: row dimension for initial image
n: column dimension for initial image
α_v: vertical camera half angle of view
α_u: horizontal camera half angle of view
θ_0: camera pitch

It is here assumed that *x* axis is in the direction of the camera centerline, projected onto the plane of the ground, and that *y* also lies in the ground plane, perpendicular to *x* (positive to the left). Equations (32) and (33) can be used to determine the ground positions for each pixel value in a given image. Knowing the ground locations for each pixel, the following equations can then be used to calculate the pixel locations within the new bird's eye image.

$$r(x) = \frac{m-1}{2}\left(1 + \frac{h - x \cdot \tan(\theta_0)}{h \cdot \tan(\theta_0) + x}\coth(\alpha_v)\right) + 1 \qquad (34)$$

$$c(x,y) = \frac{n-1}{2}\left(1 - \frac{y}{h \cdot \sin(\theta_0) + x \cdot \cos(\theta_0)}\coth(\alpha_u)\right) + 1 \qquad (35)$$

The images shown in Figure 14 are examples of the results from using these equations. These birds' eye view images are subsequently passed through the velocity-warping process.

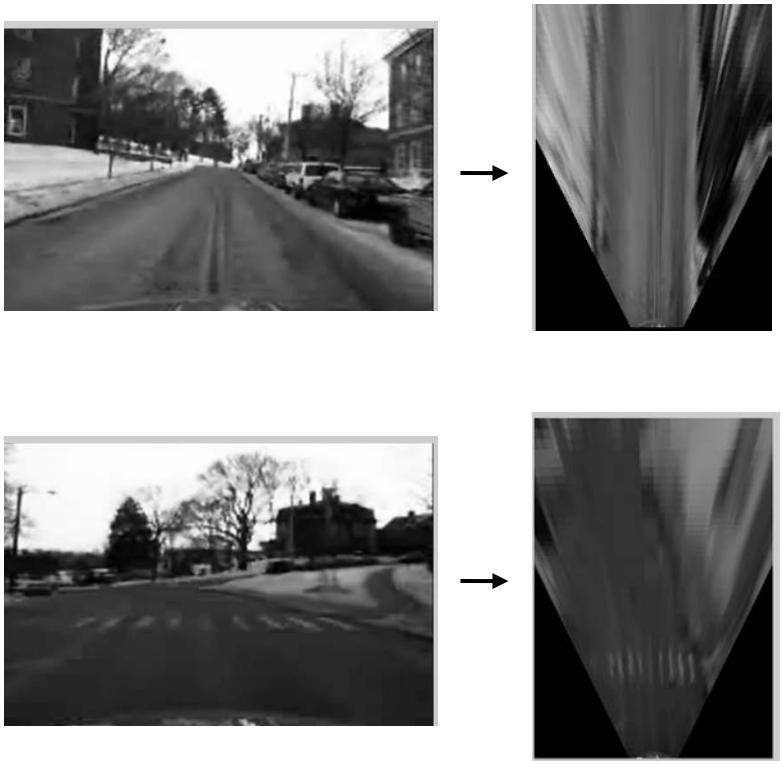

Initial Images Bird's Eye View Images

Figure 14. Bird's Eye View Transformation

Velocity Warping: Knowing the velocity of the car, it is possible to warp an initial image to create a prediction for the image at the next frame. In this way, since the camera data was converted to a bird's eye view with all pixels assumed to lie on the ground plane, only the actual ground pixels are properly warped. Additionally, since the image is in the bird's eye view, the pixel warping can be done by translating each pixel with the direction and speed of the car's velocity. After the warping process, the road pixels in the warped, predictive image should exactly match the road pixels in the actual next image, while all other pixels should differ slightly. Figure 15 shows this warping and comparison progress. In the difference image on the far right, dark regions correspond to small differences between the velocity-warped (predictive) and actual images. Typically, road pixels are dark (small differences) and non-road regions are brighter (larger differences).

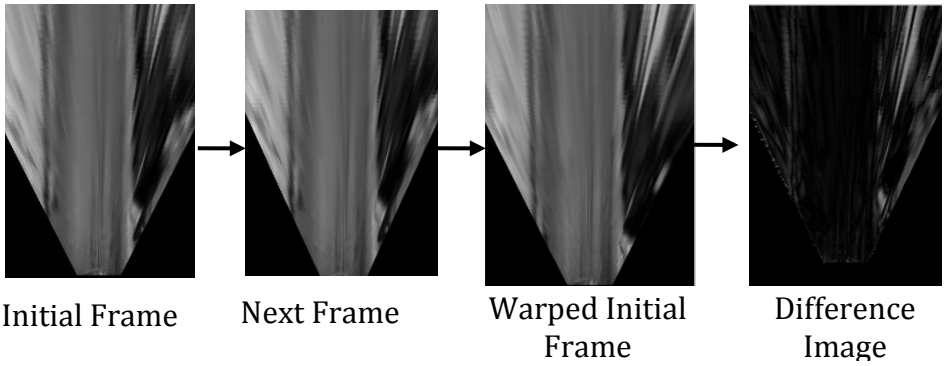

Figure 15. Optical flow results applied to a single frame.

One of the benefits of converting the camera data to birds' eye is that the data can then be combined to show more road history. Therefore, this method can also be applied to build a mosaic image, composed of several sequential bird's eye images, to enhance the sensitivity of road pixel detection. When many bird's-eye images are aligned and combined into a mosaic, any one location on the mosaic may be associated with pixel data from one, two, or many video images. For pixels located in the road, this set of values should be nearly constant; for non-road pixels, these values may vary wildly. Thus it is useful to compute the standard deviation over the set of all pixel values at each mosaic location in order to determine whether that location is (or is not) a road location.

Equations 36 and 37 (below) are used to calculate a running mean and standard deviation. In these equations, n represents the number of pixel values that have been observed at a particular location in the mosaic, p represents a specific pixel value, CA is the computed running average for the pixel value, and σ is the computed standard deviation. Computing a running average and standard deviation means that only the most recent data (and not the full set of historical pixel values) need to be stored in memory.

$$CA_n = CA_{n-1} + \frac{p_i - CA_{n-1}}{n} \tag{36}$$

$$\sigma_n = \sqrt{\frac{(n-2)\cdot\sigma_{n-1}^2 + (p_n - CA_n)(p_n - CA_{n-1})}{n-1}} \tag{37}$$

Figure 16 shows results for running average and standard deviation calculations applied to a typical video sequence. In the standard deviation image (right side), low standard deviation values, expected for road pixels, are dark, and high standard deviation values, expected for non-road pixels, are light.

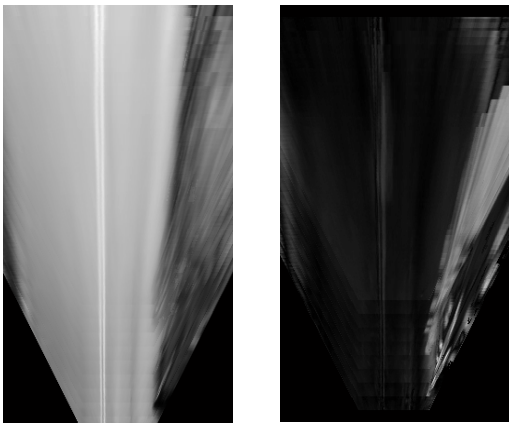

Mean Standard Dev.

Figure 16. Optical flow results applied to mosaic image. The far left image is a mosaic image composed of the mean pixel values while the right image is a mosaic image composed of the standard deviation of pixel values.

Road Pixel Detection: A method can now be applied to the standard deviation image to detect which pixels lie on the road. This method searches the image for pixels with a running standard deviation, computed by (37), that is lower than 8 gray levels. These pixels are declared to be road pixels as long as the local spatial variation of (37) is also small (variance over the nine pixels centered on the pixel of interest is less than 9). Figure 17 shows the results of this algorithm as applied to a typical mosaic image.

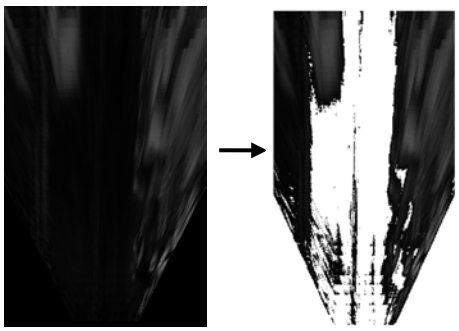

Figure 17. Road Pixel Detection Results. The left image is the mosaic standard deviation; the right images show pixels identified as road pixels (white) and non-road pixels (gray).

3.4 Integrity Monitoring Algorithm

The output of the OF algorithm can be used to verify the lane boundaries identified by the LB algorithms. One strategy for performing this cross-check is outlined in Figure 18. This procedure receives two inputs: lane line parameters and road pixel locations, from the LB and OF methods respectively. The monitoring algorithm then combines this information to compute two monitoring metrics. If these monitoring metrics are below certain threshold values, an integrity alert is triggered.

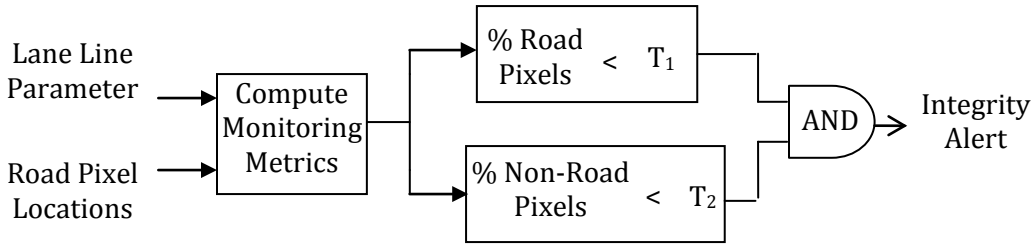

Figure 18. Integrity Monitoring System Algorithm

The monitoring metrics used in this algorithm are the percent of accurately determined road pixels, and the percent of accurately determined non-road pixels, given by equations (38) and (39), respectively. For these metrics, the

road location is estimated by the LB lane model, and the road and non-road pixels are estimated by the OF algorithm.

$$\%\ \text{Road Pixels} = \frac{\text{OF Road pixels between LB lane lines}}{\text{Pixels between LB lines}} \tag{38}$$

$$\%\ \text{Non-Road Pixels} = \frac{\text{OF Non-Road pixels outside LB lane lines}}{\text{Pixels outside LB lines}} \tag{39}$$

These two metrics are compared to threshold values T1 and T2 (as shown in the integrity monitoring algorithm in Figure 18). In order to limit the occurrence of false positives, this implementation only issues an alert if *both* pixel fractions fall below their respective thresholds. The results of applying representative thresholds (see next section) to nominal and off-nominal data are illustrated in Figure 19 – Figure 21.

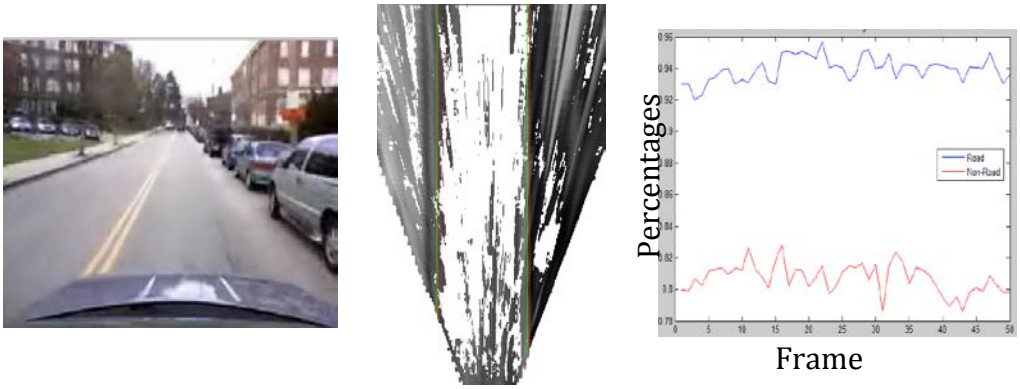

Figure 19: Nominal Method Results. Far Left: Frame taken from movie data. Center: Results from Gradient and Optical Flow Methods where green lines are ground truth and red lines are results from the Gradient method. Far Right: Error percentage results, % road pixels in blue, % non-road pixels in red.

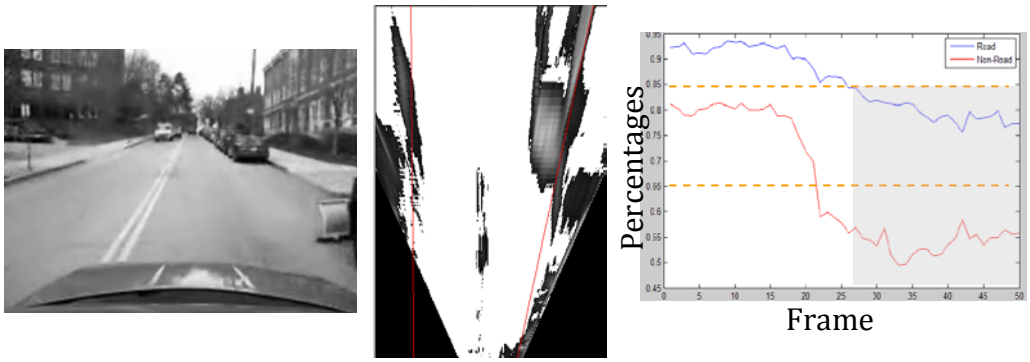

Figure 20: Gradient Method Error Results. Far Left: Frame taken from movie data. Center: Results from Gradient and Optical Flow Methods where red lines are results from the Gradient method. Far Right: Error percentage results, % road pixels in blue, % non-road pixels in red.

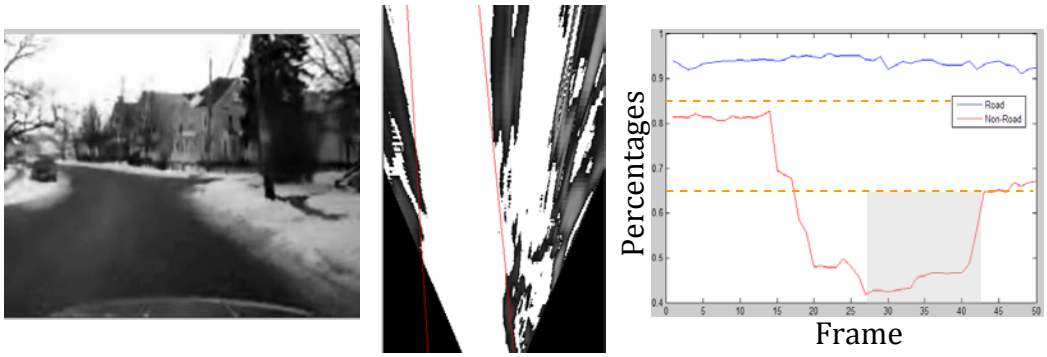

Figure 19: Snow-Error Method Results. Far Left: Frame taken from movie data. Center: Results from Gradient and Optical Flow Methods where red lines are results from the Gradient method. Far Right: Error percentage results, % road pixels in blue, % non-road pixels in red.

3.5 Threshold Selection

With the integrity monitoring algorithm defined, the threshold values for triggering an integrity alert are needed. These values must be determined by the need to identify and exclude known error cases in which the primary vision sensor fails. Once the threshold values are determined (to ensure measurement integrity in the case of known faults), then the rate of false alarms, sometimes called the *continuity risk,* can be assessed.

In determining threshold values, it is instructive to focus on three cases – typical operations (Figure 19), a case in which the primary vision algorithm fails (Figure 20), and a case in which the secondary vision algorithm fails (Figure 21). Figure shows a nominal case where both methods produce accurate results. The forward-looking camera image at one time step is shown on the far left. The OF standard deviation image is shown in the center (with LB data superimposed). The pixel percentage metrics, computed using equations (38) and (39) are shown for 50 time steps, on the far right. It is clear the road-pixel percentages are consistent around 90%,

and the non-road pixel percentages are consistent around 80%. Thus, the thresholds should be set as far below these nominal values as possible.

A case in which the primary vision algorithm fails is shown in Figure 20. In this anomaly case, a gap between cars on the right side of the road causes a tracking error for the LB method. The pixel percentage data (far right) drop significantly in response to this anomaly. If threshold values of 85% and 65% are applied to the road and non-road percentages respectively, an alert is issued around frame 26 (about 10 frames, or 1/3 of a second after the anomaly first occurs).

A case in which the secondary vision algorithm fails is shown in Figure 21. In this case, snow on the right of the road interferes with the optical flow method. (Optical flow looks for changes in pixel values over time. Bright snow causes pixel gray levels to saturate, so that the pixel values are unchanging, always at their maximum value.) Here, it is clear from the middle image that there are several inaccurately detected non-road pixels. The percentage plot on the right indicates that the non-road pixel percentage (red), dropped significantly during this incident, though the road pixel percentage (blue) was relatively unaffected. In this case, it is not necessarily desirable to trigger an alert, since the primary vision algorithm (which is used for navigation) continues to function. The dual thresholds of the integrity algorithm (Figure 18) allow for continued navigation, since the road pixel percentage would not dip below the threshold during this incident. Though it might be acceptable to operate briefly without a secondary vision algorithm to cross-check the primary algorithm, this condition should not be allowed to persist for an extended duration of time. A slight modification to the integrity monitoring algorithm would be required to exclude cases in which the secondary algorithm is persistently unavailable.

Based on these and other cases, we propose a threshold value for road pixels of 85% and for non-road pixels of 65%.

With the thresholds set to detect dangerous anomalies (i.e., cases when the primary vision algorithm reports hazardously misleading information), it is possible to assess the risk of false alarms. To compute the risk of false alarms we consider a set of 4850 frames, or 2.7 minutes of driving time since the frame rate of the data is 30 Hz. These results are shown in Figure 22. This data set was selected because it contained no oncoming traffic (which is likely to cause divergence of the LB algorithm, as it is currently implemented). Over this data set, there was one instance, around frame 2500, where both metrics dropped below their given threshold values. Thus the approximate false alarm rate is 1 per 2.7 minutes. This false alarm rate is clearly too high to deploy the algorithm in an operational automotive system, where false alarms should occur only a handful of times per year (at most). It is evident that further work is needed to refine the proposed system for practical applications. In particular, additional algorithmic development and filtering is required (1) to enable the system to operate robustly in the presence of oncoming traffic and (2) to reduce the rate of false alarms.

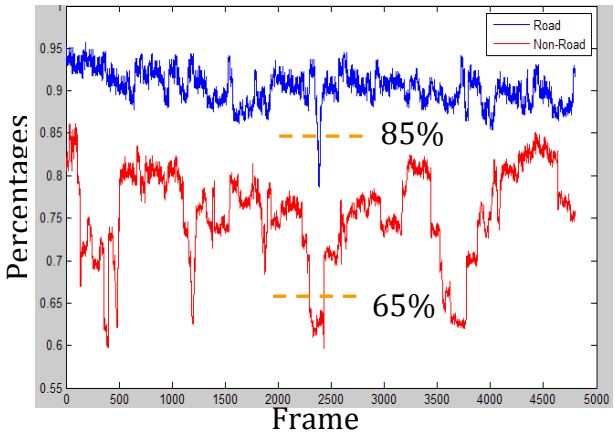

Figure 22. Integrity Monitoring Algorithm applied to nominal case data.

3.6 Summary and Conclusions

The key contributions of this chapter are (1) an identification of the specific integrity risks inherent in camera-based sensors such as those currently deployed in automotive applications like lane-departure warning, and (2) a proposed integrity monitoring approach to identify possible anomalies in visual navigation measurements. This type of integrity monitoring algorithm will be essential in future safety-of-life automotive applications, such as automated lane-keeping.

Specifically, our proposed integrity monitoring system for visual lane detection uses two independent vision processing algorithms. One method detects lane lines, and the second method performs a cross-check to identify instances in which the first method produces an anomalous measurement. If an anomaly is detected, our proposed system would alert the driver of the problem. In its current form, our proposed algorithms do not achieve necessary standards for continuity, as the false alarm rate is too high (approximately one alert per 3 minutes of driving time). Further work is needed to reduce the false alarm rate and to make the algorithms sufficiently robust for commercial application. Despite the current algorithms limitations, the system nonetheless demonstrates a proof of concept for visual integrity monitoring that clearly identifies anomalous vision-based navigation measurements, which would otherwise introduce hazardously misleading information by incorrectly reporting a car's position and orientation relative to the lane boundary.

Chapter 4: Conclusions

4.1 Thesis Contributions

This thesis presented an analysis of sensor performance for integrity requirements for automated aviation and automotive systems. For aviation applications, current airport surveillance sensors were analyzed for their impact on the performance of CD&R algorithms, and specifically, conflict detection monitoring. Similarly, for an automotive application such as automated lane-keeping, an integrity monitoring algorithm was derived to ensure the integrity of lane detection. The main contributions of this research are summarized below.

Demonstration that airport surface surveillance sensors will impact the integrity and continuity of conflict detection algorithms.

This research provided two key results. The first is that for a conflict detection monitoring aircraft separation, a margin is needed between the minimum safety separation and the nominal required separation in order to ensure integrity and continuity. The second result is the significance of off-nominal surveillance sensor behavior. These types of occurrences can impact a CD&R algorithm's ability to detect a hazardous condition, and therefore must be accounted for when designing CD&R systems. The probabilities associated with these off-nominal events feeds into the required probability for both integrity and continuity risk. Therefore, fully understanding these types of sensor systems is critical to developing integrity and continuity requirements for CD&R.

Derivation of an integrity monitor for vision-based lane detection systems.

This research identified specific integrity risks associated with camera-based sensors and derived an integrity monitor to identify these types of anomalies in vision navigation measurements. Specifically, two image processing algorithms were used to detect road lane lines. The first was a gradient-based method, which fit a linear-parabolic lane model to areas of high gradients. The second was an optical-flow based method, which provided additional information to allow for a cross-check of the first method. Both methods were applied to camera data collected in an urban environment, and a monitoring threshold was applied to detect errors. Currently, results have a false alarm rate that is too high to ensure continuity, but the monitor demonstrates that vision-based integrity monitoring can detect sensor anomalies.

4.2 Impact and Future Work

The work presented in this research will impact automated system designs for both the aviation and automotive industries. For the work related to the airport surface, the surveillance sensor findings are currently being included in the design of a CD&R algorithm to be considered by NASA. Similarly, the proposed integrity monitor is a type of algorithm that will be essential in future safety-of-life automotive applications, such as automated lane-keeping.

Future work for CD&R surveillance systems involves further research into the probability of the occurrences of each off-nominal sensor behavior as well as a deeper understanding of how individual surveillance sensors are fused together in systems such as ASDE-X, which incorporates information

from radar, multilateration, and ADS-B. On the other hand, for the lane-detection integrity monitor, further work is needed to reduce the rate of false alarms and improve system continuity. This work includes further developing the image processing methods for robustness, such as by additional filtering to reduce false alarm rate.

Appendix Overview

The key results from Chapter 2 are derived from a more detailed analysis of the sensor systems used for surveillance on the airport surface. This work did not fit into Chapter 2 itself, and therefore is included here in the appendices. Specifically, these surveillance systems include radar, multilateration, and ADS-B. The results of the work presented in these appendices were used to construct Table 1 as well as Figures 6 and 7 of Chapter 2. In particular, Table 1 outlines each system's nominal accuracy capabilities, while Figures 6 and 7 outline each system's off-nominal performance and the resulting impact on continuity and integrity.

Each of the first three appendices is focused on one of the three main surveillance systems. In general, each appendix provides an overview of the sensor technology, equations to model the system's nominal performance, and an analysis of sensor error models and off-nominal performance. The final appendix provides additional information about specific accuracy bounds referred to in the previous appendices. This comprehensive overview of the available surveillance technologies provides the background for how each sensor system will impact CD&R algorithms, as described in Chapter 2.

Appendix A: Radar Performance Model

I. Executive Summary

This appendix provides a performance model for current day airport ground surveillance systems as they impact Conflict Detection and Resolution algorithms (CD&R). These current day systems include primary and secondary surveillance radar, either as standalone systems or as components of multi-sensor systems. The goal of these sensors is to use radar capabilities to accurately detect the position information for all objects, such as aircrafts and vehicles, located on runways and taxiways. While there is ongoing research using new methods for this detection, these radar systems are still the basis of most current day ground surveillance systems.

II. Background

There are two main sensors used in current day ground surveillance systems: primary radar and secondary radar. Primary radar is designed to detect unequipped (passive) targets by broadcasting a signal and waiting for reflections to return. Secondary radar is designed to detect equipped (active) users by broadcasting a signal and waiting for user transponders to broadcast a message in return.

Primary Radar

With primary surveillance radar (PSR), a transmitter sends out an electromagnetic pulse over a narrow rotating beam and waits for its reflection off of an aircraft or another object to return [43]. The signal reflection (echo) is then used to determine the position of the target. Typically, radar updates are provided at 1 Hz. Target position is described in terms of a range distance and azimuth angle. Range is defined as the distance from the antenna to the echo centroid, and is calculated from the total travel

time and the signal speed. Similarly, azimuth is defined as the horizontal angular distance measured clockwise from true North and is measured as the angular position of the rotating beam at the instant when the echo returns. Primary radar is also capable of measuring elevation angle, the angle of the target above the horizon; however the service volume for surface systems only extends to about 200-300 feet above the ground [31].

Since the results of this system are based on the echo alone, one of the drawbacks is that it is not possible to identify the detected object directly. Rather, the identity of the target must be inferred from additional information sources. Furthermore, a data association process is required to relate radar returns acquired at different times with a particular aircraft or vehicle [44]. A set of related radar returns is sometimes called a *track*. Because it is possible that the PSR fails to associate targets with the correct track, the risk of a spurious association must be considered in automating safety-critical air traffic functions, such as CD & R.

Secondary Radar

In addition to determining a target's range and azimuth, Secondary Surveillance Radar (SSR) provides specific identification information about the detected target, as well as additional information (such as barometric altitude). With this technology, the antenna transmits 'interrogation" pulses that are received by transponders located on equipped targets (principally aircraft). These transponders respond with a signal at a different frequency, providing position and identification information [43]. Depending on the specified mode of the interrogation pulse, different aircraft information will be provided. Unlike PSR, SSR signals sent out from the antenna only need to be able to reach the target, and do not need to be strong enough to return to the antenna. Therefore, SSR has a larger coverage area than PSR and less transmission power is required. Likewise, since the initial signal travels a

shorter distance, SSR signals do not lose as much power due to attenuation as PSR signals do, resulting in higher SSR performance in poor weather conditions. Lastly, since the return signal received by the antenna is emitted by the target itself, all targets within the range of the SSR antenna can be heard equally well [45].

Although SSR offers many benefits, it has a significant disadvantage in that all targets must have a transponder in order to be detected. As such SSR is generally reserved for tracking aircraft movement and is not specifically used to track targets on the surface. However, surface surveillance systems often interface with SSR, which is used to track aircraft in the terminal area before they descend into the range of primary surveillance radar.

Multi-Sensor Systems

In surface movement applications, PSR is generally referred to as Surface Movement Radar (SMR). SMR is a key component of Airport Surface Detection Equipment, (ASDE-3 or the newer ASDE-X), and function to detect targets on runways and taxiways. Currently, ASDE-X fuses SMR with other data, such as SSR or multilateration data, if available. ASDE-X displays target position and identification information to air traffic controllers by overlaying this data on a color map of runways, taxiways, and approach corridors [46].

The newest update for ASDE-X includes an allowance for up to 4 SMR systems and is capable of processing three different types of SMR. These SMR systems include legacy SMR, 3X – the SMR system for ASDE-3, and SMR Improved (SMRi). The legacy SMR system operates in the X-Band frequency range (9.0-9.2 GHz) and uses an antenna, transceiver, and Radar Data Processor (RDP). Here, the transceiver and antenna work in parallel to transmit and receive signals. The transceiver generates X and radio frequency pulses that are transmitted by radar, and then converts received

antenna signals to baseband. Finally, the RDP converts these input signals into SMR track reports and performs target detection, association and tracking, as well as clutter rejection, and multipath target elimination. The 3X operates in the KU-Band frequency range (12-18 GHz) and uses a Radar Interface Unit (RIU) to interface to ASDE-3. The RIU receives analog data from the frequency detector and digitizes them, sending the digitized samples to the RDP. Finally, the SMRi system operates in the X band (9.0-9.2 GHz), and uses updated technology to provide more accurate position information [24, 25]. All of these SMR systems use PSR technology to determine an object's range and azimuth angle.

III. Mathematical Model

To obtain an error model for the current day sensor scenario to be used for CD&R, we must look at the individual radar components used in current day multi-sensor systems such as ASDE-X. A representative model has been previously derived for PSR [45]. Though this model applies nominally to fusion of PSR and other systems, it is descriptive of performance for PSR alone. For this model, we will use accuracy values provided by Sensis Corporation for the SMRi system. These accuracy values are 6.6ft (2.012m) RMS for range, and 0.05 degrees for azimuth [24]. Slightly more conservative estimates of azimuth accuracy are stated in product literature [47].

In order to convert positioning accuracy in a Cartesian coordinate system, it is necessary to transform the radar measurements, which are obtained in a polar coordinate system. The following transformation is used to relate a target's range and azimuth position to x-y coordinates, relative to the site of the radar antenna.

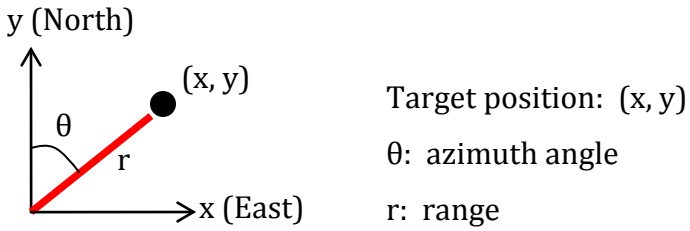

Figure 20. Radar Model Coordinate System

$$x = r \sin\theta \tag{40}$$

$$y = r \cos\theta \tag{41}$$

From these equations, it is clear there is a nonlinear relationship between the range and azimuth error values, and the resulting position error values. As a result, the error distribution produced by a radar return may have an elongated form, where lateral error (in meters) is smaller than ranging error at short distances from the radar, but much greater than ranging errors at long distances. A Monte Carlo simulation of errors at a range of 4km (near the outer range for SMR) is shown in Figure 21.

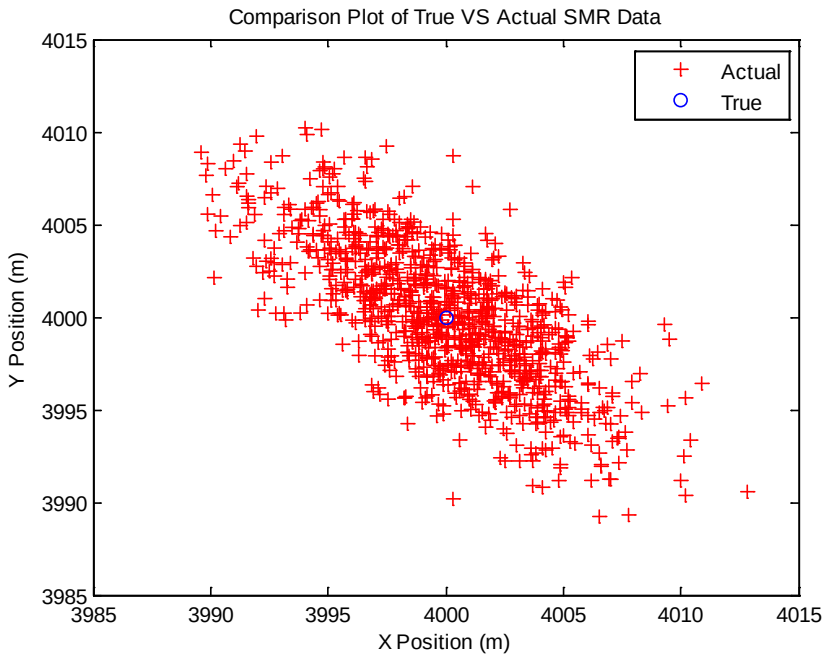

Figure 21. Radar Error Model Results.

In these plots, the blue circle represents the true target position, while the red crosses represent the target's position as seen by the radar sensor. In this simulation, the x and y position covariance matrix is [14.5 -10.1; -10.1 13.8], with corresponding standard deviation values for x-position of 3.8023 m and for y-position of 3.7166 m. (At this range, the principle axis standard deviations are 2.0 m in the ranging direction and 4.9 m in the lateral direction.) These accuracy estimates are in line with measurements made in the field [48]; however, in a sense, they are somewhat optimistic, as they do not account for data association problems, clutter, and multipath.

SMR also provides velocity measurements. Speed accuracies have been reported as 1.6 knots (0.81 m/s) one sigma, subject to a 1.9 knot (0.97 m/s) bias in mean. Heading accuracies are approximately 4.3 degrees (one sigma), subject to a 4.9 degree bias in mean. Sensor fusion processing may cause velocity measurements to lag significantly behind actual changes in velocity, with latency as high as 10 seconds [25].

IV. Additional Implications

In addition to range and azimuth angle detection accuracy, several other factors affect the use of radar for CD&R.

- *Signal Occlusion:* First and foremost, radar is only capable of detecting targets in areas where the signal is not blocked by buildings or aircraft. Therefore, airport areas near the terminal gates will be hard to cover completely using only radar.

- *Ground Vehicle Tracking:* Additionally, for CD&R algorithms, it is necessary to track vehicle trajectories. This is particular difficult for ground vehicles, which are not typically equipped with transponders that transmit identifying information. For these cases, data association issues often arise. This is especially true since ground vehicles often travel out of coverage areas, leading to gaps in the collected data [29].

- *Weather:* Rain and ground clutter can also cause interference with both radar systems, often leading to false alarms for target detection [47,30].

- *Cell Resolution:* Another error with radar is due to resolution errors and differentiating between objects close to each other. Also, a target splitting phenomenon can occur since aircraft targets are large, on the scale of 60-70m, compared to radar resolution cells, on the scale of 6m x 0.4 degrees. Therefore, large targets will occupy hundreds of cells and, as a result, often appear disconnected [32].

- *Incomplete Information:* Lastly, if the altitude is unknown, some bias can occur in the ranging direction due to an incomplete trigonometry model.

V. Conclusion

Surface movement radar provides high accuracy surveillance, with positioning errors on the order of 2-5 m (one sigma) and velocity errors on the order of 1 m/sec (one sigma), both subject to minor biases. Measurement accuracy is sufficiently high to support sophisticated CD & R processing. However, ancillary issues, such as measurement lag or the unavailability of measurements in certain occluded areas of the facility, may place restrictions on how CD & R algorithms are implemented. A major benefit of fusing additional sensor data with primary radar is to provide this additional coverage and to establish the ability to perform cross-checks for anomaly detection and outlier rejection.

Appendix B: Multilateration Performance Model

I. Executive Summary

This appendix provides nominal error models for multilateration, a surveillance technology that is currently being used in the United States in conjunction with radar as a subsystem of ASDE-X. Since it is capable of detecting aircraft in areas otherwise blocked for radar, multilateration is a useful tool for conflict detection & resolution (CD&R). This appendix examines multilateration's typical performance, as well as other off-nominal situations circumstances that can affect CD&R.

II. Multilateration Overview

Multilateration is a sensor system currently used for ground surveillance at airports throughout the United States. It determines an aircraft's position based on the time difference of arrival (TDOA) of signal emitted by the aircraft and received by at least three receivers. This system's biggest advantage is that it is a quickly implemented technology with a relatively low-cost that can provide both position and unique identification information for aircraft on the surface.

The earliest multilateration prototype for airport surveillance dates back to the 1970s, but technology has improved since then and equipment costs have been reduced significantly. [43,49]. There are two classes of multilateration currently in use in the United States. The first form of multilateration is used in conjunction with Surface Movement Radar (SMR) in an airport surface surveillance system produced by Sensis Corporation called Airport Surface Detection Equipment – Model X (ASDE-X). The FAA has identified 35 airports for ASDE-X implementations, and as of October 2010, ASDE-X was reported operational at 32 of these airports [50]. The second application of

multilateration is called Wide Area Multilateration (WAM), and has been implemented in Juneau, Alaska and four airports in Colorado to improve guidance for aircraft approaches in mountainous regions where radar signals are blocked. This appendix will focus only on the first case: multilateration used on the airport surface.

III. System Architecture

A basic multilateration system determines an aircraft's position by having several ground stations listen for a signal emitted by a transponder-equipped aircraft. The targeted aircraft's position can then be computed based on the time differences between the receipts of the signal at the different ground stations. Unlike GPS, where a single receiver processes signals from four or more satellites, here a single transponder is emitting one signal to multiple ground stations with known locations. A central station then processes these signals and computes the transponder's location. For a 2D position calculation, such as the cases where the height of the aircraft is known, a minimum of three ground stations are required, while 3D position calculations require four or more ground stations. A full multilateration system therefore requires a properly equipped aircraft, a set of working ground stations, and a central ground station to compute the aircraft's position. Figure 22 outlines this system, showing an aircraft emitting a signal to four different ground stations (A-D), with travel times t_A - t_D. The central station then computes the aircraft's position based on these different signal travel times.

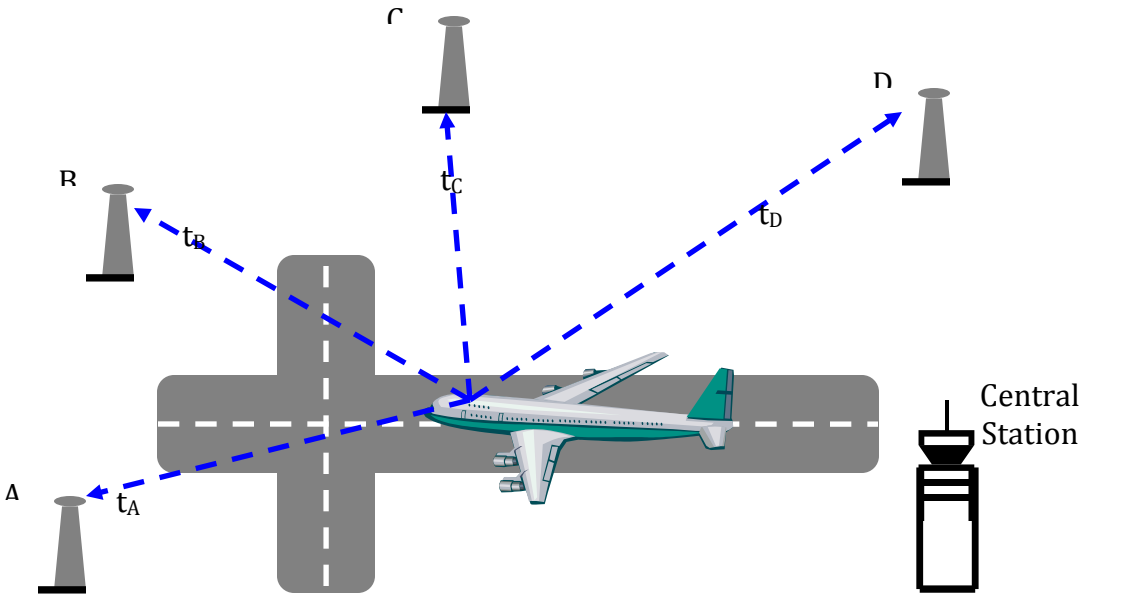

Figure 22. Multilateration System Schematic. Aircraft on the ground and in the air emit signals to ground stations, which use the time delay of arrival (TDOA) to calculate the aircraft's position.

Aircraft Equipage

The first step in multilateration involves an aircraft emitting a signal to then be received by the ground stations. Therefore, aircraft must be properly equipped to perform this action. Current Secondary Surveillance Radar (SSR) and Automatic Dependent Surveillance – Broadcast (ADS-B) technologies use transponders for aircraft communication, consequently providing an existing infrastructure that can be used for multilateration.

Several message structures are broadcast by aircraft equipped with SSR transponders: Mode 3/A, Mode C, or Mode S structures (see Table 3). Mode 3/A and Mode C are both 4 digit octal codes, where Mode 3/A provides an aircraft identification code, assigned by the air traffic controller and not necessarily always unique, while Mode C provides aircraft pressure altitude information. Aircraft transponders emitting either Mode 3/A or Mode C messages are passive and must first be interrogated by an outside source. On the other hand, Mode S is a 24-bit address which combines an aircraft

identification that is fixed for each aircraft (and always unique) with aircraft altitude information. Furthermore, Mode S transponders emit signals automatically, usually at a rate of 1 Hz.

ADS-B transponders emit signals, referred to as Mode S Extended. These signals are automatically generated. They combine a Mode S message structure with additional aircraft state information provided by the aircraft navigation system. Multilateration does not use this position information contained within ADS-B messages; however, this information may later be combined through sensor fusion [51-54].

Available Surveillance Technology	Communication Message Type	Message Structure	Message A/C Information	Automatic/ Interrogated
SSR	Mode 3/A	4 digit octal	ATC provided ID	Interrogated
	Mode C	4 digit octal	Altitude	Interrogated
	Mode S	24-bit	Unique ID and altitude	Automatic (1 Hz)
ADS-B	Mode S Extended	112-bit	Unique ID, altitude, and state	Automatic (1 Hz)

Table 3. Available Surveillance Technologies and Message Structure Information

Ground Station Requirements

In order to process these aircraft signals, a set of functioning ground stations located throughout the airport is also required. For a typical large airport, the full multilateration system includes approximately 14-18 ground stations located on the airport surface and around the perimeter. For example, the installation at Atlanta International Airport includes 16 ground stations [55]. Usually, two of these ground stations are Reference Transmitters, which automatically emit time synchronization signals and monitor the entire

system [56]. Since the calculation of the aircraft's position is dependent on the timing of the signal arrival at each of the ground stations, it is important that the ground station clocks remain synchronized. The rest of the ground stations are one of two different types of Remote Units (RU): Receiver/Transmitter (R/T) and Receiver Only (RO). R/Ts are capable of interrogating transponders, such as those operating with Mode 3/A or Mode C, while ROs are only capable of listening for aircraft emitted signals. Both types of RUs are capable of processing signals from all possible aircraft transponders. In concept, only one R/T is needed to request data from aircraft. However, in practice, multiple R/Ts are needed to reach all users, due to large airport buildings that block transmission to aircraft receivers. Figure 23 illustrates the ground station layout for Atlanta/Hartsfield-Jackson Atlanta International Airport, with red asterisks representing the RefTrans, blue circles representing the R/Ts, and green squares representing the ROs. Lastly, there is a central processor which collects transponder reply information from the RUs, tracks the multilateration targets, monitors the status and health of all pieces of equipment, and schedules interrogations to be emitted by R/Ts as needed [25].

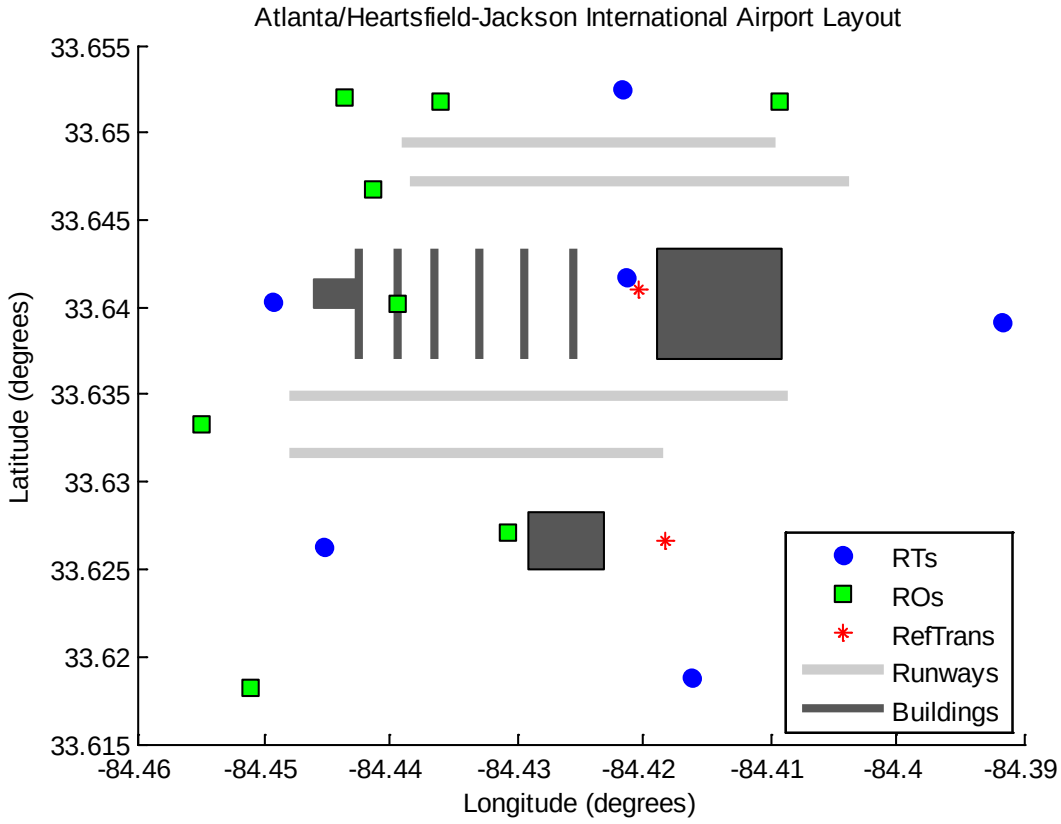

Figure 23. Ground station layout for Atlanta International Airport.

Position Determination

Determining the aircraft's location based on the travel time of the signal emitted by the aircraft's transponder relies on an equation that relates the measured time the signal traveled to the distance between the aircraft and the ground station. Equation (42) below represents the range from the aircraft at the unknown location (x_a, y_a, z_a) to the m^{th} ground station at the known location (X_m, Y_m, Z_m) related to the travel time of the signal sent at an unknown time t_a and received by the ground station at a known time t_m:

$$\sqrt{(x_a - X_m)^2 + (y_a - Y_m)^2 + (z_a - Z_m)^2} = c(t_m - t_a) \qquad (42)$$

Subtracting this range equation for one ground station from the range equation for a second ground station leads to the following equation (43), which can also be referred to as the time difference of arrival (TDOA) [57]:

$$\sqrt{(x_a - X_1)^2 + (y_a - Y_1)^2 + (z_a - Z_1)^2} - \sqrt{(x_a - X_2)^2 + (y_a - Y_2)^2 + (z_a - Z_2)^2} = c(t_1 - t_a) - c(t_2 - t_a) \tag{43}$$

Simplifying (2) eliminates t_a from the equation, which is beneficial because it also eliminates an unknown clock error associated with the aircraft's transponder [57]. For a set of *M* ground stations, the set of nonlinear equations to describe a transponder's location $\vec{x} = [x_a, y_a, z_a]^T$ from a multilateration solution is given by equations (44-46). Note that for this set of equations, the vector *y* has (*N-1*) elements for a set of *N* measurements due to the subtraction used for TDOA.

$$\vec{y} = h(\vec{x}) \tag{44}$$

where,

$$\vec{y} = \begin{bmatrix} c(t_1 - t_2) \\ c(t_2 - t_3) \\ \vdots \\ c(t_{m-1} - t_m) \end{bmatrix} \tag{45}$$

$$h(\vec{x}) = \begin{bmatrix} \sqrt{(x_a - X_1)^2 + (y_a - Y_1)^2 + (z_a - Z_1)^2} - \sqrt{(x_a - X_2)^2 + (y_a - Y_2)^2 + (z_a - Z_2)^2} \\ \sqrt{(x_a - X_2)^2 + (y_a - Y_2)^2 + (z_a - Z_2)^2} - \sqrt{(x_a - X_3)^2 + (y_a - Y_3)^2 + (z_a - Z_3)^2} \\ \vdots \\ \sqrt{(x_a - X_{m-1})^2 + (y_a - Y_{m-1})^2 + (z_a - Z_{m-1})^2} - \sqrt{(x_a - X_m)^2 + (y_a - Y_m)^2 + (z_a - Z_m)^2} \end{bmatrix} \tag{46}$$

In order to solve for the aircraft's position, we must first linearize equation (44). This process yields the following system of equations that can then be used to determine $\hat{\vec{x}}$, which is an estimate of the true state, $\vec{x}$.

$$\vec{y} = h(\hat{\vec{x}}) + H \cdot d\hat{\vec{x}} \tag{47}$$

where,

$$h\left(\hat{\vec{x}}\right)=\begin{bmatrix} \sqrt{\left(\hat{x}_a-X_1\right)^2+\left(\hat{y}_a-Y_1\right)^2+\left(\hat{z}_a-Z_1\right)^2}-\sqrt{\left(\hat{x}_a-X_2\right)^2+\left(\hat{y}_a-Y_2\right)^2+\left(\hat{z}_a-Z_2\right)^2} \\ \sqrt{\left(\hat{x}_a-X_2\right)^2+\left(\hat{y}_a-Y_2\right)^2+\left(\hat{z}_a-Z_2\right)^2}-\sqrt{\left(\hat{x}_a-X_3\right)^2+\left(\hat{y}_a-Y_3\right)^2+\left(\hat{z}_a-Z_3\right)^2} \\ \vdots \\ \sqrt{\left(\hat{x}_a-X_{m-1}\right)^2+\left(\hat{y}_a-Y_{m-1}\right)^2+\left(\hat{z}_a-Z_{m-1}\right)^2}-\sqrt{\left(\hat{x}_a-X_m\right)^2+\left(\hat{y}_a-Y_m\right)^2+\left(\hat{z}_a-Z_m\right)^2} \end{bmatrix} \tag{48}$$

$$H=\begin{bmatrix} \frac{\hat{x}_a-X_1}{\hat{R}_1}-\frac{\hat{x}_a-X_2}{\hat{R}_2} & \frac{\hat{y}_a-Y_1}{\hat{R}_1}-\frac{\hat{y}_a-Y_2}{\hat{R}_2} & \frac{\hat{z}_a-Z_1}{\hat{R}_1}-\frac{\hat{z}_a-Z_2}{\hat{R}_2} \\ \frac{\hat{x}_a-X_2}{\hat{R}_2}-\frac{\hat{x}_a-X_3}{\hat{R}_3} & \frac{\hat{y}_a-Y_2}{\hat{R}_2}-\frac{\hat{y}_a-Y_3}{\hat{R}_3} & \frac{\hat{z}_a-Z_2}{\hat{R}_2}-\frac{\hat{z}_a-Z_3}{\hat{R}_3} \\ \vdots & \vdots & \vdots \\ \frac{\hat{x}_a-X_{m-1}}{\hat{R}_{m-1}}-\frac{\hat{x}_a-X_m}{\hat{R}_m} & \frac{\hat{y}_a-Y_{m-1}}{\hat{R}_{m-1}}-\frac{\hat{y}_a-Y_m}{\hat{R}_m} & \frac{\hat{z}_a-Z_{m-1}}{\hat{R}_{m-1}}-\frac{\hat{z}_a-Z_m}{\hat{R}_m} \end{bmatrix} \tag{49}$$

$$\hat{R}_m=\sqrt{\left(\hat{x}_a-X_m\right)^2+\left(\hat{y}_a-Y_m\right)^2+\left(\hat{z}_a-Z_m\right)^2} \tag{50}$$

With four ground stations, this system can be solved for exactly one 3D aircraft position. However, it is likely there will be more than four ground stations. As a result, this system will be overdetermined, where there are more equations than unknowns. Therefore, a Nonlinear Least Squares method is applied [57]. The goal of Least Squares is to optimize the system solution by minimizing the square of the residuals [41]. The least squares solution for equation (47) then becomes:

$$d\hat{\vec{x}}=\left(H^TH\right)^{-1}H^T\cdot\left(\vec{y}-h\left(\hat{\vec{x}}\right)\right) \tag{51}$$

Since this is a nonlinear system, this computation, based on an initial guess for $\hat{\vec{x}}$, must be iterated several times until it converges on a solution. In each iteration, the previous $\hat{\vec{x}}$ is replaced with a new estimate $\hat{\vec{x}}$.

$$\hat{\vec{x}}=\hat{\vec{x}}+d\hat{\vec{x}} \tag{52}$$

This method will provide a minimum error result for any case where there is more than four ground stations that have received the signal, and as a result, more than three observation equations.

These equations can also be applied to the case where the height of the aircraft is known, such as determining an aircraft's position on the airport surface. In this case, the vertical (z_a) term is known, and so only two states remain unknown (x_a and y_a). For a TDOA solution, three measurements are needed to obtain these two unknown states.

IV. Error Model

Installed Systems

In order to ensure that multilateration systems provide useful surveillance information, the FAA has outlined several system requirements, including specifications for accuracy, detection probability, and coverage zones. For positioning, the FAA requires that multilateration systems, for surface applications within a coverage area that spans all taxiways and runways from the ground to 300ft above, meet a position accuracy of 20ft rms (6.1m) one-sigma. In addition to meeting this position accuracy, the multilateration system reports an associated quality indicator, the Geometric Dilution of Precision (GDOP), a scalar value used to characterize the quality of the position solution [59]. The FAA requires that the GDOP be less than 1.5 for 98% of the entire coverage volume, including larger Wide Area Multilateration regions. Every multilateration measurement also includes a time associated with each aircraft position. This time parameter is required to be accurate to within 10 milliseconds one-sigma.

In order to ensure sufficient operational accuracy, the FAA places several requirements on each multilateration installation. All runways and taxiways

longer than 150ft must be fully supported by multilateration; moreover, if any dead zones (areas shorter than 150ft) do exist, they must not overlap with radar dead zones. Similarly, for detection at 1-Hz, multilateration systems must be able to detect 93% of all possible transponders, averaged over the entire coverage area, including dead zones. For false reporting, multilateration systems must have less than 1 false report per six hours during peak operation hours for Mode S transponders, and less than 2% false plots of all plots during peak hours for Mode 3/A and Mode C transponders. (A false report is defined as a measurement for a non-existent target or as a measurement more than 100 ft from a target's true position.) Lastly, multilateration systems must be able to provide position information for at least 200 aircraft within the coverage area [25].

Evaluations of installed multilateration systems indicate that operational systems meet or exceed these requirements. In the US, multilateration is currently used as a subsystem of ASDE-X, a surveillance system produced by Sensis Corporation. Currently, Sensis reports that their multilateration system meets the FAA outlined specs of 20ft (6.1m) one-sigma position accuracy, with 93% detection probability [24]. A study conducted at Atlanta's airport by MIT Lincoln Lab in 1998 also reported 20ft positioning accuracies [49]. Sensis has also provided documentation of accuracies reported over an entire airport's coverage area for Atlanta and Milwaukee. For both of these airports, the position errors usually remain low, less than 10ft, but with a few spots of errors as high as 22.5ft [55]. These higher errors usually correspond to areas with limited line of sight to the ground stations. However, the fact that ground station layouts are specified for individual airports helps minimize this error.

Simulation

A component of this appendix is a simulation of nominal multilateration positioning performance, for use in CD & R algorithm development. For this simulation, we use the same ground station layout for Atlanta's airport outlined previously. Figure 24 shows the airport layout for only those ground stations used for TDOA. The simulation allows the user to be positioned at any point on the airport surface. The simulation is implemented using metric units.

In implementing the simulation, a nominal value must be assigned to the ranging measurement error for each RU. A nominal value for the ranging error to each RU is 4.1 m. This number can be computed by dividing the target accuracy (6.1 m) by the maximum GDOP (1.5).

In our simulation, we have inflated the nominal ranging measurement error to account for line-of-site blockages. A full simulation environment would account for building locations and would detect which RUs were visible to an aircraft located at a particular location on the airport surface. This ray-tracing feature is not included in the current version of our simulation. As such, we have arbitrarily inflated the ranging accuracies to reflect the worse GDOP values that would result due to signal blockage. The inflation factor (1.5) has been chosen for this specific airport facility as a value that places typical multilateration accuracies into the range observed at airports like Atlanta and Milwaukee (e.g. into the range of 10-22.5 ft). As a result, our simulation uses a modified ranging accuracy of 6.1 m for each RU, such that the ranging and positioning accuracy are equivalent for a case of GDOP equal to one.

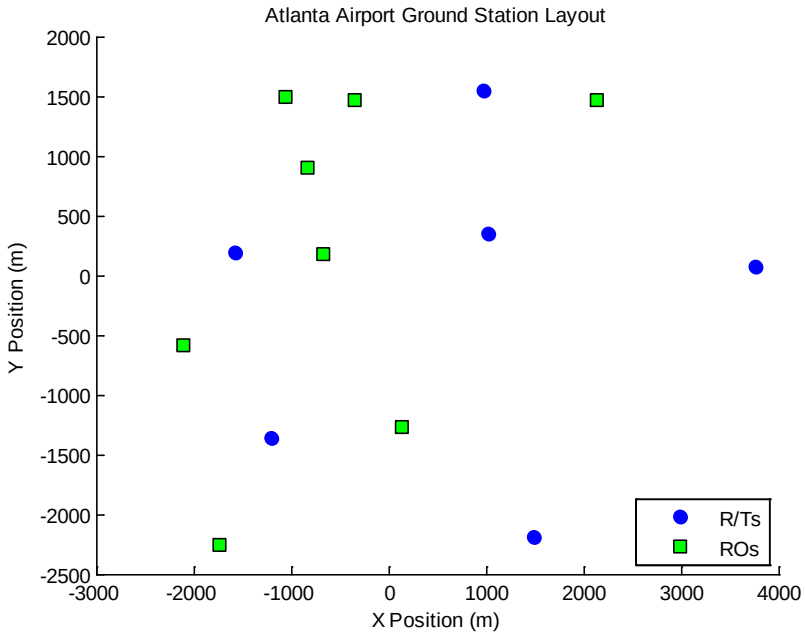

Figure 24. Multilateration Ground Station Layout for Atlanta/Hartsfield-Jackson Atlanta International Airport.

To demonstrate the simulation, we consider a specific instance in which 500 Monte Carlo samples were generated for a user with a true location at (0,0). Figure 25 shows the simulation results. In the simulation shown below, the x and y position covariance matrix is [22.9 1.5; 1.5 17.3], with corresponding standard deviation values for x-position of 4.8 m and for y-position of 4.2 m.

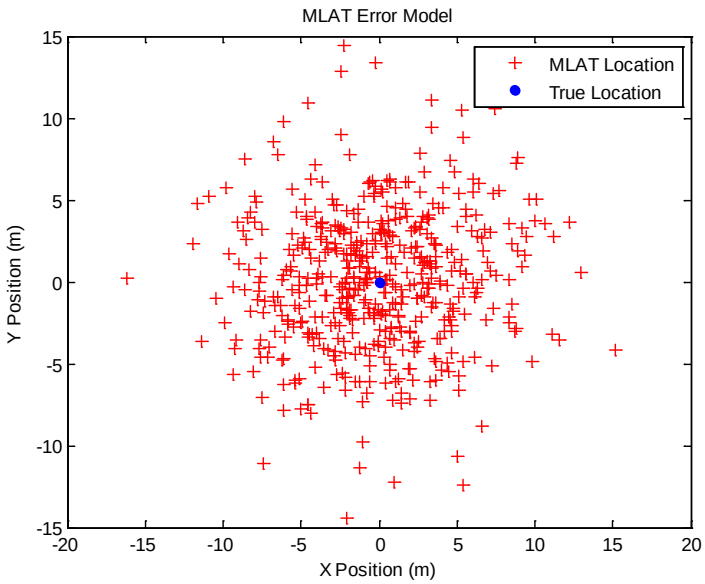

Figure 25. Multilateration Error Model (500 points).

In addition to solving for the transponder's position on the airport, there is also a bias error for where this transponder lies on the aircraft itself. Since the exact transponder location on an aircraft is unknown, there is a considerable amount of uncertainty with regard to where the body of the aircraft lies relative to the transponder. This is especially important for CD&R algorithms. For example, a CD&R algorithm could rely on a detected transponder's location as being centered on the aircraft, when in reality, the transponder is located at the front of the aircraft. In this case, the CD&R algorithm would be incorrectly regarding the rear half of the aircraft as open space, which could inadvertently lead to a collision.

V. Multilateration Fault Modes and Their Impact on CD&R

In addition to multilateration's reported accuracies, there are several possible fault modes which can impact CD&R. These include a range of communication faults which occur between the equipped aircraft and the ground stations, and ground station timing faults which can impact the computed position.

- Missed Message: Occurs when a message is successfully transmitted by an aircraft but fails to reach the ground system receiver. For example, the transmission might be blocked by objects or building along the path between the aircraft and ground system receiver(s). If the signal's path towards a high enough number of ground stations is blocked, it is possible the multilateration system will not be able to compute its location. In this case, since only old position data is available to the CD&R algorithm as a result and also since, under extreme conditions, the existence of an aircraft or a ground vehicle may not be recognized, missed message faults are primarily an integrity issue.

- Aircraft Transceiver Error: Occurs when the transceiver located on the aircraft is inactive, missing, or experiencing failure. Therefore, no position data is provided for that aircraft. Since it is possible for the aircraft to go undetected, this is an integrity issue.

- Ground Station Equipment Failure: Occurs due to an electronic equipment failure at an RU. This fault can largely be mitigated by enabling a cross-check procedure over a redundant set of ground stations. (Then, if the signals into one ground-station antenna are corrupted, the comparison to other ground station antennas will reveal the inconsistency.) Because such monitoring is possible, the threat of a ground-station antenna failure is primarily a continuity issue (i.e. an issue which might briefly interrupt service due to a real or false alarm issued by the consistency cross-check monitor).

- Signal Jamming: Occurs when an outside source intentionally or unintentionally interferes with the aircraft's signal transmission

and prevents ground stations from receiving the messages. Elevated Radio-Frequency Interference is generally easy to detect, so this fault is primarily a continuity issue rather than an integrity issue.

- Message Collision: Occurs when two transmitters send signals at the same time. This can also be considered 'unintentional jamming' and is a continuity or availability issue.

- Ground Station Clock Synchronization: Occurs when the ground station clocks are not properly synchronized. If the clocks are not properly synchronized, then the calculated aircraft position will be incorrect. The FAA requires that all ground stations are synchronized with a time source within an accuracy of 10 milliseconds. Therefore, signals received at the ground stations and used in multilateration calculations are time stamped within an accuracy of 10 milliseconds. If the time source itself fails, a backup timing capability will ensure degraded timing (100 millisecond accuracy) for up to one hour after the fault occurred [25].

VI. Summary

In this appendix we have described the architecture of multilateration, provided an error model for position, and investigated the impact of multilateration on CD&R algorithm design. Multilateration is a useful system to provide aircraft identification information on the airport surface as well as position information for areas not covered by primary radar. It is capable of providing 6.1m accuracy and identification information for all equipped

aircraft and ground vehicles on the airport surface. While this accuracy is currently not as high as the WAAS or GBAS GPS systems used by ADS-B, multilateration can be more easily implemented than ADS-B since there is no need to update aircraft equipment. Ground stations do need to be implemented at all airports, with optimal ground station geometries subject to individual airport layouts. However, once the ground stations have been set up, multilateration is capable of working with new and old transponder systems, making it a valuable tool as the airport industry makes the transition to full ADS-B equipped systems. On the other hand though, since multilateration requires signals emitted by transponders, it cannot detect unequipped vehicles. Currently, ASDE-X fuses multilateration data with radar, which is capable of detecting unequipped vehicles, and ADS-B, which provides higher accuracy position information, to accurately detect all aircraft and vehicles on the airport surface. By using all three systems, ASDE-X can provide position information for unequipped ground vehicles using radar, high accuracy position information for all ADS-B equipped aircraft, and position and identification information for all transponder-equipped aircraft using multilateration. Overall, multilateration, as part of ASDE-X, is currently a relatively low-cost system to provide aircraft positioning information on the airport surface as the FAA transitions to fully equip airports and aircraft for ADS-B by 2020.

Appendix C: ADS-B Performance Model

I. Executive Summary

This appendix provides nominal error models for ADS-B, an emerging surveillance technology that can be used for conflict detection & resolution (CD&R). In addition to characterizing typical performance, we also consider other implications of ADS-B on CD&R. An important benefit is that ADS-B explicitly characterizes integrity by establishing a far-tail bound on sensor error (10^{-7}). A potential liability is that the coarse, discrete accuracy and integrity levels used in ADS-B may limit the sensitivity of our proposed conflict detection algorithms.

II. ADS-B Overview

ADS-B, which stands for Automatic Dependent Surveillance – Broadcast, is a system where an aircraft equipped with an ADS-B transponder periodically broadcasts a message containing its state vector [25]. More specifically, this message includes the following aircraft information: position, altitude, speed, heading, air/ground status, navigation uncertainty, aircraft ICAO address, aircraft type, and flight ID [60]. This message is received by several ground stations, which relay the information to the control tower.

The system is automatic in that nothing elicits the transmission of the aircraft status messages, and is dependent because it relies on on-board equipment to provide its navigation information. This on-board equipment responsible for emitting messages is often referred to as "ADS-B Out". A receiver located at a ground station processes these aircraft messages and sends the information along to the control tower. The control tower can then forward this position information to surrounding aircraft [61]. Additionally, other ADS-B equipped aircraft can choose to enable their ADS-B system to receive

surrounding aircraft position information, eliminating the need for the control tower middle step. This equipment is referred to as "ADS-B In", and provides a cockpit display of traffic information, allowing pilots to be aware of surrounding aircraft, even in low visibility situations, without relying solely on the control tower [62]. Figure 26 outlines a basic ADS-B system.

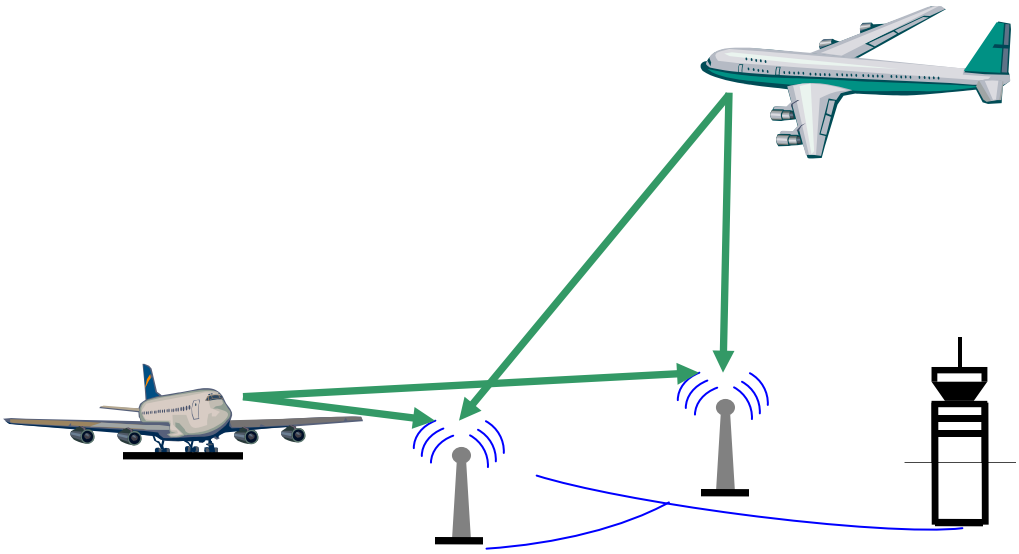

Figure 26. ADS-B Schematic. Aircraft on the ground and in the air emit messages to ground stations, which relay the information to the control tower.

ADS-B is beneficial to CD&R in its ability to quantify integrity bounds, and provide a rigorous bound on rare errors. Additionally, for flight-deck based CD&R, pilots have a direct knowledge of other aircraft position information.

III. Adoption Schedule

The success of ADS-B implementation requires all aircraft to be equipped with, at a minimum, ADS-B Out technology. This is an expensive and therefore slow conversion process. The FAA's current implementation plan

will lead to full aircraft and airport equipage by 2020. Other phases of that implementation plan include the first phase, 2006-2010, where ground stations are in the process of being built at several airports in key areas; the second phase, 2009-2014, where ground station coverage will be completed and aircraft equipage will be increased to 40%; the third phase, 2015-2020, where 100% of aircrafts will be equipped with at least ADS-B out; and finally, the fourth phase, 2020-2025, where old surveillance equipment, especially secondary radar, will be decommissioned [34]. There is no required implementation of ADS-B In technologies, but this implementation is beneficial to aircraft by providing direct information from surrounding aircraft.

Not only does this implementation plan span almost twenty years, but it also is associated with several costs. For initial costs, the FAA will be required to purchase or contract hardware for both the ground stations and air traffic control towers, develop procedures, and train personnel, including controllers, maintenance, and management. Operators' initial costs will go towards new avionics equipment, including ADS-B Out (as a minimum), a position source, antenna, and interface with other aircraft systems, as well as labor costs and any incurred losses for aircraft downtime. Additionally, the FAA will face recurring costs for calibrating and repairing ground stations and other equipment, and continuously reevaluating and developing procedures. Operators will also have to maintain their avionics equipment and may face additionally losses for aircraft downtime for any equipment repairs [62]. However, despite these costs, successfully implemented, fully equipped ADS-B systems will be able to replace secondary radar systems, improve cockpit situational awareness, and provide a more accurately defined layout of the active airport.

IV. System Architecture

Communication

ADS-B relies on the communication between the on-board aircraft equipment, ground stations, and control tower. ADS-B Out transponders located on the aircraft emit messages at a context-dependent rate (between 1 Hz for moving surface vehicles and 0.2 Hz for stationary vehicles) [25]. These messages are picked up by receivers at various locations on the airport surface, which in turn relay the information to the airport control tower. Other aircraft equipped with ADS-B In technology are also capable of receiving and processing these messages. In the long term, the FAA has stated that ground vehicles could also be equipped with ADS-B Out transponders, providing additional traffic information to the control tower [34]. Ground vehicles can also transmit their position using the Traffic Information Service – Broadcast (TIS-B), and ADS-B In technology is capable of receiving and processing these TIS-B messages [63].

Navigation

The position accuracy and integrity of ADS-B is based entirely on the system used to determine the aircraft's position. For ADS-B equipped aircraft, most aircraft rely on a Global Positioning System (GPS) to determine their location. This basic GPS model provides position accuracy on the order of 8m. A differential correction system, such as the Wide Area Augmentation System (WAAS), further improves this accuracy to 1.5-2m [64] [65]. WAAS also provides integrity information for the overall WAAS system, individual GPS satellites, and other error estimates used in the WAAS position correction. As an alternative to WAAS, some aircraft may rely on Receiver Autonomous Integrity Monitoring (RAIM) to supply integrity information for GPS. RAIM operates by monitoring the residuals of the GPS position solution to detect inconsistencies [66] [67] [68].

Global Navigation Satellite Service (GNSS) refers to the generic term for satellite navigation systems, while GPS refers to the US government operated satellite constellation used for positioning. This constellation consists of a minimum of 24 active satellites, with four satellites evenly spaced within one of six different orbital planes [69]. These planes are evenly spaced at 60° around Earth, relative to the Equator, with a 55° angle of inclination above the Equator. Each orbit is positioned at 20,200 km (10,900 NM) above Earth's surface, and operates with a 12 hour period [70]. These satellites emit signals to GPS receivers which then determine their position based on the time difference between when the signal left the satellite and when the signal reached the receiver. Typically, for any given point on Earth, there are four to twelve visible GPS satellites.

In a setting where the satellite clocks are perfectly synchronized, there are four unknowns in determining the receiver's position: three position coordinates and a receiver clock bias. With four unknowns, a minimum of four satellites is needed to calculate the receiver's position. The following equation relates the pseudorange to the four unknowns:

$$\rho = \sqrt{(X-x)^2 + (Y-y)^2 + (Z-z)^2} - c \cdot dT \qquad (53)$$

where (X, Y, Z) represents the satellite's position, dT represents the receiver clock bias, and (x, y, z) is the receiver's position [71]. Linearizing this equation for a set of m satellites yields the following:

$$d\rho = G \cdot ds \qquad (54)$$

where,

$$d\rho = \begin{bmatrix} d\rho_1 \\ d\rho_2 \\ \vdots \\ d\rho_m \end{bmatrix} \qquad (55)$$

$$G = \begin{bmatrix} \frac{\hat{x} - X_1}{\hat{R}_1} & \frac{\hat{y} - Y_1}{\hat{R}_1} & \frac{\hat{z} - Z_1}{\hat{R}_1} & 1 \\ \frac{\hat{x} - X_2}{\hat{R}_2} & \frac{\hat{y} - Y_2}{\hat{R}_2} & \frac{\hat{z} - Z_2}{\hat{R}_2} & 1 \\ \vdots & \vdots & \vdots & 1 \\ \frac{\hat{x} - X_m}{\hat{R}_m} & \frac{\hat{y} - Y_m}{\hat{R}_m} & \frac{\hat{z} - Z_m}{\hat{R}_m} & 1 \end{bmatrix} \tag{56}$$

$$ds = \begin{bmatrix} dx \\ dy \\ dz \\ c \cdot dT \end{bmatrix} \tag{57}$$

$$\hat{R}_m = \sqrt{(X_m - \hat{x})^2 + (Y_m - \hat{y})^2 + (Z_m - \hat{z})^2} \tag{58}$$

Aiming to minimize ds, this set of equations is iterated several times, each time updating the state vector (x, y, z, cdT) with the calculated ds values [72].

To date, the only differential GPS correction service that has been certified to Fully Operational Capability (FOC) by the FAA for aviation applications is the Wide Area Augmentation System (WAAS). WAAS uses a combination of ground and satellite equipment to improve the accuracy of GPS for the continental US and large areas of Alaska, Canada, and Mexico [73]. WAAS achieved FOC status in July 2003 [74] and currently involves a network of ground reference stations, ground master stations, and geostationary satellites. Ground reference stations can determine GPS signal errors by comparing GPS calculated position with the known reference station position. These errors include GPS satellite position and clock errors, ionospheric delays, and ephemeris and clock drift. Once determined, the ground stations forward this error information to WAAS Master Stations via terrestrial communication, where WAAS messages are generated. These messages are uploaded to geostationary satellites, which in turn, broadcast

the information to GPS/WAAS receivers. The GPS/WAAS receivers are then able to incorporate this error information in their position calculation and improve the GPS position accuracy [64,65].

In the future, airports and aircraft will be fully equipped for a differential GPS correction service known as the Ground Based Augmentation System (GBAS), which is also referred to as the Local Area Augmentation System (LAAS) in the United States. This system is capable of further improving position accuracy to less than 1m in both the horizontal and vertical directions [75]. Similarly to WAAS, GBAS uses a set of ground reference stations to generate differential corrections, which are subsequently broadcast to users. However, with GBAS, these GPS reference stations are located on airport grounds and serve an area within a 20-30 mile radius [75]. The proximity of the reference antennas to the user in GBAS significantly enhances the benefits of differential corrections, since certain GPS errors sources (notably delays caused by the Earth's ionosphere and troposphere) are more highly correlated over shorter distances. After comparing GPS calculated positions with known reference station positions, a central ground facility, also located at the airport, uses a VHF data broadcast (VDB) transmitter to distribute differential corrections and integrity data to all equipped aircraft in the GBAS coverage area. [76]

GBAS has been identified by the Joint Program Development Office (JPDO) as a key enabling capability for NextGen [20]. Currently, GBAS is in the initial stages of certification. Honeywell received FAA approval for a Non-Federal CAT I System called Smartpath SLS-4000 in September 2009 [78] and work is still underway to develop the high level requirements needed for CAT II/III GBAS. Ultimately, the goal for GBAS systems is to meet CAT II/III requirements [79-80]. While these requirements are defined for high precision aircraft approaches to an airport, GBAS will also be able to be used

for airport surface movement either through a general Differential Correction Positioning Service (DCPS) or through a new service defined specifically for surface movement [81-83].

V. Error Models

ADS-B Accuracy and Integrity Categories

In order to ensure that ADS-B provides reliable navigation information, the FAA uses a set of standards determined by the RTCA (Radio Technical Commission for Aeronautics) called Minimum Aviation System Performance Standards (MASPS). These standards specify accuracy and integrity performance parameters for ADS-B operations in the U.S. Within these standards, the FAA has mandated that ADS-B equipped aircraft meet certain Navigation Accuracy Category bounds for position and velocity (NACp and NACv, respectively), which specify 95% accuracy bounds on horizontal position. Currently, FAA requirements for these parameters are NACp $\geq$ 9 and NACv $\geq$ 1, corresponding to a horizontal position accuracy within 30m, and a computed aircraft velocity accuracy within 10 m/s [84].

In addition to these accuracy standards, MASPS for ADS-B include integrity standards as well. The Navigation Integrity Category (NIC) specifies an integrity containment radius and is closely associated with another parameter, the Source of Integrity Level (SIL). SIL further defines integrity requirements by representing the probability of the true aircraft position lying outside the NIC containment radius. The FAA currently requires NIC $\geq$ 7, corresponding to an integrity containment radius less than 370.4 m, and SIL = 3, for a probability less than 10^{-7} per flight hour or per sample that the true position is not within the containment radius.

In order to conserve bandwidth by keeping the ADS-B message as short as possible, NACp, NACv, and NIC values are discretized coarsely into a limited

number of integer values, which may be represented by a 4-bit field. These values, specified by the MASPS [34,85], are summarized in the tables in Appendix D. In general, a higher NACp, NACv, or NIC category indicates better performance. The categories are intended to support a range of operations, include en route, terminal area, approach, landing, and surface movement operations. As such, the defined categories span a very wide range of values. When no category exists to precisely describe sensor accuracy and integrity, ADS-B broadcasts NACp, NACv and NIC bounds that are conservatively rounded up. Consider, for instance, the two highest NACp values: a NACp of 10, which corresponds to a 10 m error bound (95%), and a NACp of 11, which corresponds to a 3 m error bound (95%). If the actual sensor error is 4m (95%), ADS-B would round this value up to the next available category, which in this case would be a NACp of 10 (corresponding to a 10 m error bound). The noise levels implied by the 10 m containment radius would clearly be more severe than the actual errors, as illustrated in Figure 27. In this sense, ADS-B presents a challenge for developing CD & R algorithms, since the reported accuracy may be a conservatively inflated representation of actual sensor accuracy.

Low-Fidelity Sensor Error Models

As described in the previous section, ADS-B broadcasts error modeling parameters, or categories, which conservatively describe the actual sensor error. For this reason, we take a two stage approach to modeling ADS-B errors. First, in this subsection, we discuss how a low-fidelity error model can be generated directly from ADS-B broadcast message parameters. Second, in subsequent sections, we provide higher-fidelity models that represent actual sensor performance.

The following equations can be used to relate an ADS-B containment radius to a horizontal position error. These equations assume a nominal Gaussian distribution and provide errors in a local East-North coordinate system.

$$\varepsilon_E = N\left(0, \frac{NAC_p}{2}\right) \tag{59}$$

$$\varepsilon_N = N\left(0, \frac{NAC_p}{2}\right) \tag{60}$$

The low-fidelity errors in both the east and north directions (ε_E and ε_N) are both sampled randomly from a zero-mean, Gaussian distribution N characterized by a standard of $NAC_p/2$. These equations were used to generate the synthetic noise plot shown on the left side of Figure 27.

The low-fidelity model of velocity error, based on ADS-B broadcast parameters, has a related form. The standard deviation of the velocity errors ($\varepsilon_{v,E}$ and $\varepsilon_{v,N}$), however, is obtained from the broadcast NAC_v parameter.

$$\varepsilon_{v,E} = N\left(0, \frac{NAC_v}{2}\right) \tag{61}$$

$$\varepsilon_{v,N} = N\left(0, \frac{NAC_v}{2}\right) \tag{62}$$

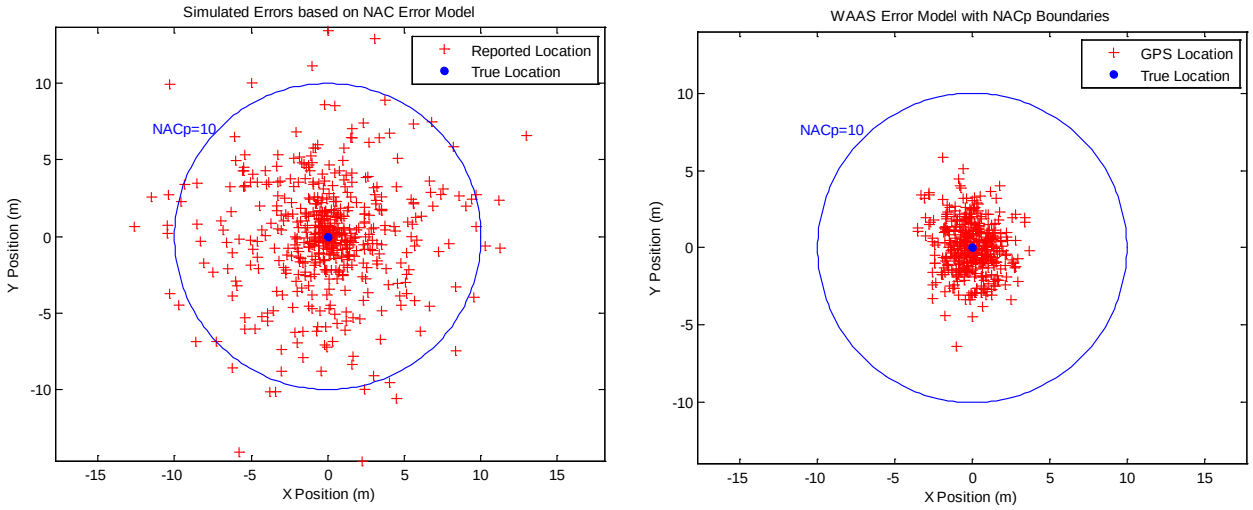

Figure 27. Comparison of Synthetic Noise Generated from NACp bound (left) to Nominal Noise Generated from a Higher Fidelity Error Model (right). Note that the noise implied by the NACp bound is conservative compared to the higher fidelity noise model.

Higher Fidelity Position Error Model

Although in designing a CD & R algorithm it may not be possible to take full credit for a higher fidelity error model, it is still useful to consider such models in order to anticipate what levels of ADS-B performance will be seen over two time frames: a transitional scenario, in which GPS differential corrections (if used) are provided by WAAS, and a modernized scenario, in which GPS differential corrections are provided by GBAS.

In both the transitional and modernized scenarios, we can compute a higher fidelity representation of GPS positioning errors by modeling the geometry of the satellite constellation in orbit around Earth. At any given instant, the constellation geometry can be described by a matrix **G**, which was defined by equation (56). Errors on pseudorange measurements from each satellite contribute to the user position error. If a weighted, iterative least-squares solution to (54) is used to compute the GPS position, then the state error

vector $\boldsymbol{\varepsilon}_s$ is related to the satellite pseudorange error vector $\boldsymbol{\varepsilon}_\rho$ by the following equation.

$$\boldsymbol{\varepsilon}_s = \left(\mathbf{G}^T\mathbf{W}\mathbf{G}\right)^{-1}\left(\mathbf{G}^T\mathbf{W}\right)\boldsymbol{\varepsilon}_\rho \tag{63}$$

The first three elements of the state error vector describe the vector position error. The fourth component of the state error is the clock-correction error. In its most typical form, the weighting matrix **W** used in a GPS solution is the inverse of the covariance matrix for the pseudorange errors. Such a weighting matrix functions by emphasizing low-noise measurements and de-emphasizing higher-noise measurements.

$$\mathbf{W} = \left(E\left[\boldsymbol{\varepsilon}_\rho \boldsymbol{\varepsilon}_\rho^T\right]\right)^{-1} \tag{64}$$

A Monte Carlo simulation of position errors for a given satellite geometry can be generated using equation (63). In such a simulation, a multivariate distribution for the pseudorange error vector $\boldsymbol{\varepsilon}_\rho$ is needed both to compute the weighting matrix **W** and to generate random samples of the pseudorange errors, themselves.

For the transitional scenario, we use a relatively simple model for the error. Specifically, we assume that the errors for all satellites are independent, identically distributed random variables described by a zero-mean Gaussian distribution of standard deviation 1.7 m. In this case, it is easy to show that the weight matrix cancels out of equation (63), such that the state error for the transitional scenario may be simplified as follows.

$$\boldsymbol{\varepsilon}_s = \left(\mathbf{G}^T\mathbf{G}\right)^{-1}\left(\mathbf{G}^T\right)\boldsymbol{\varepsilon}_\rho \tag{65}$$

A somewhat more sophisticated error model is introduced for the modernized scenario. In this case, we model each satellite error as an independent, zero-mean Gaussian distribution; however, the standard deviation of each satellite error has a distinct standard deviation that depends on the elevation of the satellite in the sky. Typically, low-elevation satellites are subject to much larger errors than high-elevation satellites, due to grazing multipath and atmospheric obliquity. Refined models of GBAS errors that take these effects into account have been previously developed to support system design [86,87]. These standard GBAS error models, in fact, consider four contributions to the total error for each satellite *i*: ground receiver noise, airborne receiver noise, nominal ionosphere gradients, and nominal troposphere gradients.

$$\sigma_i^2 = \sigma_{gnd,i}^2 + \sigma_{air,i}^2 + \sigma_{iono,i}^2 + \sigma_{tropo,i}^2 \tag{66}$$

These noise terms all depend on each satellite's elevation, E_i, as described by the following equations. First, the ground receiver noise is modeled by the Ground Accuracy Designator (GAD) "C4" curve [86,87].

$$\sigma_{grnd,i}(E_i) = \sqrt{\frac{1}{4}(\sigma_{RR})^2 + 0.04^2 + \left(\frac{0.01}{\sin(E_i)}\right)^2} \tag{67}$$

$$\sigma_{RR,i} = \begin{cases} 0.15 + 0.84e^{-E_i/15.5^\circ} & E_i \le 35^\circ \\ 0.24 & E_i > 35^\circ \end{cases} \tag{68}$$

Next, the airborne receiver noise is described by the Aircraft Accuracy Designator (AAD) curves. Here, to model the error associated with CAT III-qualified equipment, the "B" AAD curve is used [86].

$$\sigma_{air,i}(E_i) = \sqrt{\sigma^2_{air\_noise,i} + \sigma^2_{air\_multipath,i}} \quad (69)$$

$$\sigma_{air\_noise,i} = 0.11 + 0.13e^{-E_i/4^\circ} \quad (70)$$

$$\sigma_{air\_multipath,i} = 0.13 + 0.53e^{-E_i/10^\circ} \quad (71)$$

The ionosphere contribution is described by the following equation.

$$\sigma_{iono,i}(E_i) = OF(E_i) \cdot \sigma_{vig}(X + 2\tau V_{air}) \quad (72)$$

The ionosphere error equation depends on the following parameters: σ_{vig}, the vertical ionosphere gradient standard deviation (nominally 4 mm/km but here taken as 10 mm/km to bound worst case error as described by [83]); X, the distance between the aircraft and the ground facility antenna centroid (taken as 4 km, assuming both aircraft and antennas are located on the airport surface); τ, the filter smoothing time (taken as 30 s); and V_{air}, the aircraft approach velocity (modeled with a conservative upper bound of 0.130 km/s). The ionosphere error also depends on the distance through the ionosphere that the GPS signal travels, a distance which is greater when the satellite is low in the sky. To account for this elevation-dependent distance, an obliquity factor $OF(E_i)$ is introduced.

$$OF(E_i) = \sqrt{1 - \left(\frac{R_e \cos(E_i)}{R_e + h_I}\right)^2} \quad (73)$$

In this equation, R_e is the approximate radius of the Earth's ellipsoid (6378.1363 km) and h_I is the height of the maximum electron density of the ionosphere (approximately 350 km).

Lastly, the troposphere error is assumed to be much smaller than the other error sources and therefore is taken as zero in this model.

$$\sigma_{tropo,i} = 0 \tag{74}$$

Because all errors are essentially independent, the weighting matrix **W** for GBAS is diagonal, with the inverse of the variance of each satellite's pseudorange error along the diagonal. In short, the weighting matrix for GBAS has the following form.

$$W_{ii} = 1/\sigma_i^2 \tag{75}$$

Comparing both the transitional (WAAS-based) and modernized (GBAS-based) error models demonstrates that GBAS provides a significant benefit. Under nominal conditions, GBAS position errors will fall within the NACp=11 bound, while WAAS nominal errors will only meet the requirements for the wider NACp=10 bound.

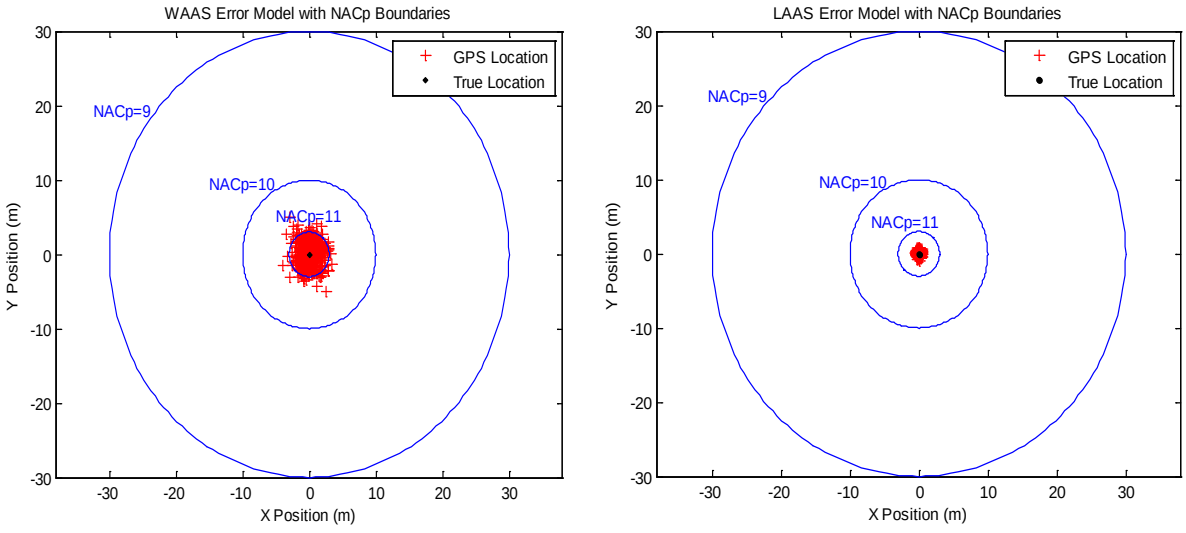

Figure 28. WAAS and LAAS Error Model Results for NACp Bounds

Higher Fidelity Velocity Error Model

For a GPS-equipped user, ADS-B velocity is estimated using GPS-signal Doppler measurements. Doppler shifts reflect changes in the observed satellite signal frequency due to the relative motion of the satellite and the GPS receiver. Because satellite velocities are known, Doppler shifts can be used to estimate the velocity of the user receiver, directly. Doppler components can be measured in the direction of each satellite; hence, it is possible to reconstruct all three components of an aircraft's velocity through GPS Doppler measurements. Augmentation systems do not publish corrections for velocity measurements. Hence, since neither WAAS nor GBAS provides any benefit in estimating velocity, the accuracy of GPS velocity measurements is the same for both the transitional and modernized scenarios.

The equation for user receiver velocity is similar to that for position [88]. Here the variable $\dot{\boldsymbol{\phi}}$ is used to describe the vector of Doppler measurements, the vector $\mathbf{v}$ is the user position and the scalar $\dot{b}$ is the drift of the user receiver clock.

$$\dot{\boldsymbol{\phi}} = \mathbf{G}\begin{bmatrix}\mathbf{v}\\ \dot{b}\end{bmatrix} \tag{76}$$

The structure of this equation resembles that of the linearized position equation (53). Accordingly, the method for characterizing velocity error is similar to that for characterizing position error. In the following equation, the velocity error $\boldsymbol{\varepsilon}_v$ is related to the measurement error $\boldsymbol{\varepsilon}_{\dot{\phi}}$ through the pseudoinverse of the geometry matrix $\mathbf{G}$.

$$\boldsymbol{\varepsilon}_v = \left(\mathbf{G}^T\mathbf{G}\right)^{-1}\left(\mathbf{G}^T\right)\boldsymbol{\varepsilon}_{\dot{\phi}} \tag{77}$$

We model Doppler measurement errors as being independent, Gaussian-distributed errors with a standard deviation set at 0.1 m/s. For our baseline GPS constellation, used to support Monte Carlo simulations, this measurement noise level results in a user velocity error in the East direction of 0.06 m/s and in the North direction of 0.08 m/s. These numbers are a reasonable, if slightly conservative, approximation of typical results for GPS velocity estimation (which, for example, have been reported to be on the order of 0.05, one-sigma by [88]).

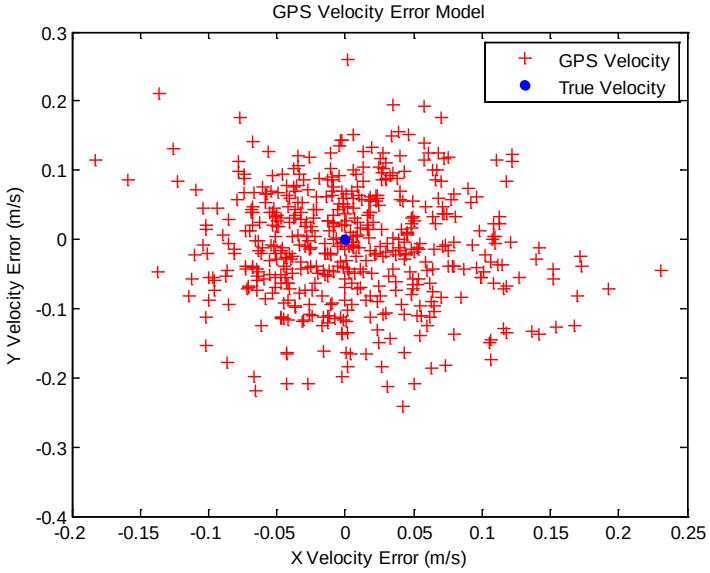

Figure 29. GPS Velocity Error Model Results

This velocity accuracy of GPS meets the requirements for the highest velocity category supported by the ADS-B standard: a NACv equal to 4. This NACv category corresponds to a bound of 0.3 m/s on horizontal speed (95%). The following figure illustrates that simulated GPS velocities (500 points) are contained by accuracy category 4.

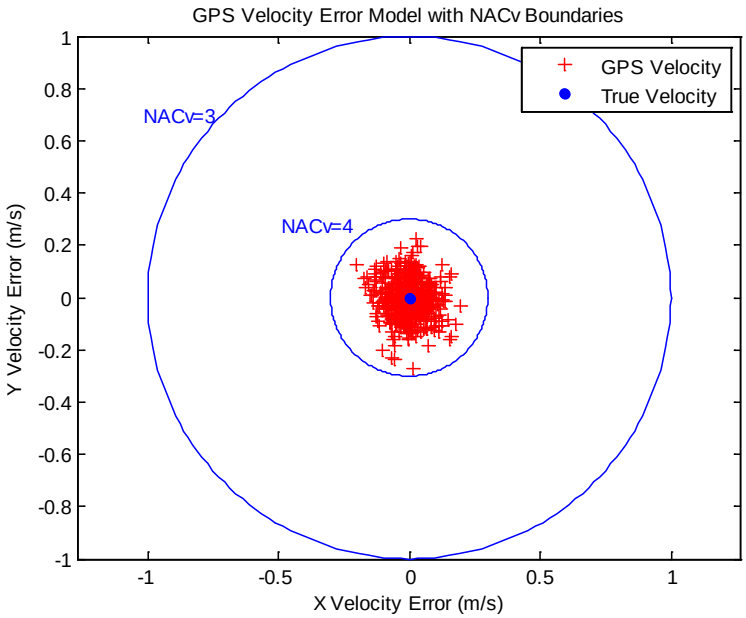

Figure 30. GPS Velocity Error Model Results with NACv Bounds

Far-Tail Error Modeling: Integrity

Although Gaussian distributions are commonly used to describe nominal errors, far-tail sensor errors are not typically Gaussian distributed. As such, it is useful to use a refined error model to describe rare, hazardously large sensor errors. This is the primary distinction between the ADS-B accuracy and integrity categories. The accuracy categories are intended to describe typical, essentially Gaussian sensor error behavior. The integrity category (NIC) is used to bound large errors which may not follow a Gaussian distribution.

Integrity is defined in terms of an integrity risk probability and a containment radius. The integrity risk probability is the chance that an error exceeds the containment radius. RAIM, WAAS and GBAS are all capable of supporting the highest ADS-B integrity risk requirements (SIL = 3). RAIM and WAAS are already capable of achieving an integrity requirement of 10^{-7} per hour. The GBAS DCPS has not yet been certified, but is anticipated to achieve this requirement also [89]. In fact, future Cat III GBAS systems,

which are intended to support fully automated aircraft landing, will require that the combined vertical and lateral integrity risk shall not exceed $1x10^{-9}$, where the vertical risk applies over any 15 second time period and the lateral risk applies over any 30 second time period. For GBAS to ensure integrity over these short windows requires a very rapid alert time (nominally 2 second time-to-alert for GBAS, in contrast with a 10 second time-to-alert for WAAS) [89,90].

Containment radius is a function of the satellite geometry and the types of differential corrections available. In GPS applications, containment radii are sometimes called *protection levels*. Protection levels for RAIM are typically 60 m or more [91], corresponding to a NIC of 9. Protection levels for WAAS users are typically on the order of 15 to 25 m [92], corresponding to a NIC of 10, but occasionally down to a NIC of 9. While there are few available reports with information about GBAS protection levels, these values typically range from 5 to 10 m [93,34], corresponding to a NIC of 10, or occasionally as high as 11.

VI. Nominal Function of ADS-B and Its Impact on CD&R

While ADS-B is useful in providing detailed position and identification information for aircraft on the runway surface, there are also several characteristics of the system that will impact CD&R. These include different sampling times, delayed sensor data, discrete accuracy and integrity levels, and the availability of higher accuracy positioning data needed for CD&R.

Heterogeneous Sampling Times

ADS-B systems rely on an external GPS unit to determine the aircraft's position before emitting a message to surrounding receivers. As a result, it is possible the time at which the position is determined by the GPS differs from the time at which the aircraft position message is transmitted. This timing

issue is referred to as latency, which describes the delay between when the time of the position measurement and the position signal transmission time [34]. To avoid transmitting old and inaccurate position information, the FAA requires that the aircraft transmit its position no later than 2 seconds after the time the measurement was taken. Within this 2 second latency allowance, a maximum of 0.6 seconds of uncompensated latency is allowed. This means that, for all other latency times greater than 0.6 seconds (but within the 2 second maximum), the aircraft must compensate by extrapolating its position to the position at the time of the message transmission [84]. As a result, if the aircraft's transmitted message is not properly time-tagged with position measurement and transmission timing data, there could be an added position error of up to 0.6*velocity. For an aircraft traveling at 10 m/s, this corresponds to an error of 6m, an error large enough to cause a discrepancy between where the CD&R algorithm expects an aircraft to be, and where that aircraft truly lies.

In addition to the latency requirements, the FAA requires that an aircraft transmit its position and velocity at a rate of 1 Hz while in the air or moving on the airport surface. For aircraft stationary on the ground, this transmission frequency is reduced to once per 5 seconds [84]. However, this does not guarantee that messages from different aircraft are received, or that measurements were taken, all at the same instant. As a result, a CD&R algorithm relying on these position messages may need to incorporate a capability to manage asynchronously sampled data.

Pure Delay on Sensor Data Input to CD&R

In addition to timing delays between the position measurement and the message transmission, it is also possible to have sensor data delays due to missed messages at the receiver end. It is possible for a receiver to detect and report an aircraft's position at one time instance, and then miss

subsequent messages due to either a transmission error or signal blockage. As a result, a CD&R algorithm would be forced to rely on that aircraft's older position until its position message was picked up by a receiver again.

Another delay issue occurs when the error model of the aircraft's position measurement changes. The FAA requires that any change in the accuracy parameters, NACp or NACv, be broadcast within 10 seconds [84]. In this case, a CD&R algorithm could base its calculations off of an incorrect position error bound for up to 10 seconds.

Discrete Accuracy and Integrity Levels and Reduced Sensitivity of CD&R

Due to the discrete accuracy and integrity levels provided with aircraft position information in an ADS-B message, threshold used in conflict detection may need to be inflated significantly. Assuming that threshold values are set at a particular multiple of sensor noise (in order to ensure a required false alarm rate), threshold values will be pushed farther into the tails of a distribution since NAC levels are set to be conservative, meaning they are as large or larger than actual accuracy values. In short, since the monitor thresholds are made coarse by the coarse NAC discretization, some monitoring sensitivity is lost. With wider thresholds, more missed detections may occur (reducing the integrity of the CD & R algorithm). Figure 6 shows how the imposed discretized threshold leads to a loss of sensitivity.

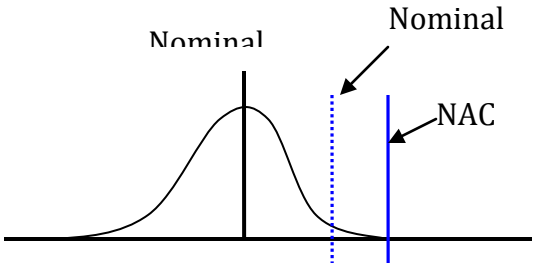

Figure 31. Discretized threshold display.

It is also significant to note that conflict detection thresholds may need to be determined dynamically, to optimize CD & R performance when NIC and NAC levels are high, and to compensate for reduced performance when NIC and NAC levels are low.

Availability of CD&R Function

It is possible that effective CD & R may require higher NIC and NACp levels than the minimums dictated by the FAA. As a consequence, the CD & R function may not be available 100% of the time. If the CD & R function becomes unavailable, then alternative procedures – such as a switch to visual surveillance by controllers in the tower or to an increased aircraft separation – may be required.

As an example, it may be determined that an integrity bound of 25 m or better (NIC = 10) is needed for adequate CD&R function. At the very least, it is unlikely that a circular integrity bound of 370.4 m (NIC = 8) is acceptable for CD &R function, even though this NIC level may be broadcast (per FAA policy). If, indeed, CD & R requirements are only satisfied when NIC > 8, then the availability of the CD & R function (i.e. the percentage of time the CD & R function can be used safely by air traffic controllers) will be reduced by the fraction of time that aircraft broadcast a NIC equal to 8.

GPS Time-to-Alert (TTA) Specifications

Current regulations require that CAT III GBAS systems meet a 2 second TTA constraint while WAAS is only required to meet a 10 second TTA constraint. This constraint specifies the amount of time for an alert to be broadcasted after the navigation system in use detects a failure [89,90]. For GPS systems, these failures could be related to specific satellite failures. Therefore, in the case of GBAS, incorrect position information would be relayed for up to 2 seconds before the system determined the information did not meet performance standards. Furthermore, for WAAS, this error would increase even more, as the incorrect information could be relayed for up to a total of 10 seconds. Overall, this amount of time where incorrect position information is being provided without an alert can have a negative impact on CD&R functions

VII. ADS-B Fault Modes and Their Impact on CD&R

The fault modes of ADS-B fall under two categories, communication faults and navigation faults, and can be mapped into integrity, continuity, or availability issues. Communication faults occur between the equipped aircraft and the ground stations, and include missed messages, transceiver/receiver failures, and message collisions, while navigation faults occur in the sensor's position measurement, which in this case is GPS.

Communication Faults

- Missed Message: Occurs when a message is successfully transmitted by an aircraft but fails to reach the ground system receiver. For example, the transmission might be blocked by objects or building along the path between the aircraft and ground system receiver(s). Since only old position data is available to the CD&R algorithm as a result and also since, under extreme

conditions, the existence of an aircraft or a ground vehicle may not be recognized, missed message faults are primarily an integrity issue.

- Aircraft Transceiver Error: Occurs when the transceiver located on the aircraft is inactive, missing, or experiencing failure. Therefore, no position data is provided for that aircraft. Since it is possible for the aircraft to go undetected, this is an integrity issue.

- Ground Station Equipment Failure: Occurs when position information from aircraft does not properly arrive at sensor fusion processor located at the ground station. This fault can largely be mitigated by enabling a cross-check procedure over a redundant set of ground stations. (Then, if the signals into one ground-station antenna are corrupted, the comparison to other ground station antennas will reveal the inconsistency.) Because such monitoring is possible, the threat of a ground-station antenna failure is primarily a continuity issue (i.e. an issue which might briefly interrupt service due to a real or false alarm issued by the consistency cross-check monitor).

- Signal Spoofing: Occurs when a fake signal is broadcast by a hostile agent, introducing false positions which are harder to detect as inaccurate. As a result, this is an integrity issue.

- Signal Jamming: Occurs when an outside source intentionally or unintentionally interferes with the aircraft's signal transmission and prevents ground stations from receiving the messages. Elevated Radio-Frequency Interference is generally easy to detect,

so this fault is primarily a continuity issue rather than an integrity issue.

- Message Collision: Occurs when two transmitters send signals at the same time. This can also be considered 'unintentional jamming' and is a continuity or availability issue.

- Improper Time Tagging: Occurs when the message lacks information about the time when the aircraft's position was measured. Due to a maximum uncompensated latency time of 0.6 seconds, this fault introduces a maximum bias error of 0.6*velocity. This is an integrity issue and could cause a CD&R missed detect.

Navigation Faults

- GPS Faults: Occurs when there are GPS faults not otherwise detected by WAAS (or by RAIM). Examples of these faults include code-carrier divergence, excessive satellite clock acceleration, mild ionosphere storms, and satellite signal deformation. Such faults have a low probability, on the order of one satellite fault per year, and will typically be detected by WAAS. These detections will interrupt the surveillance data signal, possibly causing a loss of continuity for the CD & R function. The stated continuity risk for a loss of the WAAS signal is 10^{-5} per hour. There is a very small chance that WAAS fails to detect a hazardously large GPS error, which might result in elevated integrity risk for CD & R. The stated risk of an undetected large error (an integrity violation) for WAAS is $1x10^{-7}$.

In the future, as NextGen implements GBAS for use in ADS-B systems, navigation information will meet improved accuracy and integrity requirements. As it stands now, WAAS is the only differential GPS system approved for use on airports. However, as research is done to improve GBAS in order to meet CAT III requirements, CD&R functions in the future will be able to rely on much more precise navigation information.

VIII. Summary

In this appendix we have described the architecture of ADS-B, provided an error model for position, and investigated the impact of ADS-B on CD&R algorithm design. While ADS-B is not currently fully operational at all airports, it has the ability to work with multilateration and primary and secondary radar systems as the airport industry transitions to fully equipped systems, and its adoption is anticipated (for all aircraft) by the year 2025.

Appendix D: Navigation Accuracy Category (NAC) Bounds

NACp	95% Containment Radius
0	EPU ≥ 10 NM
1	EPU < 10 NM
2	EPU < 4 NM
3	EPU < 2 NM
4	EPU < 1 NM
5	EPU < 0.5 NM
6	EPU < 0.3 NM
7	EPU < 0.1 NM
8	EPU < 0.05 NM
9	EPU < 30 m
10	EPU < 10 m
11	EPU < 3 m

NACv	95% Horizontal Velocity Error
0	Unknown or ≥ 10 m/s
1	< 10 m/s
2	< 3 m/s
3	< 1 m/s
4	< 0.3 m/s

NIC	Horizontal Containment Bounds
0	Unknown
1	Rc < 37.04 km (20 NM)
2	Rc < 14.816 km (8 NM)
3	Rc < 7.408 km (4 NM)
4	Rc < 3.704 km (2 NM)
5	Rc < 1852 m (1 NM)
6	Rc < 1111.2 m (0.6 NM)
7	Rc < 370.4 m (0.2 NM)
8	Rc < 185.2 m (0.1 NM)
9	Rc < 75 m
10	Rc < 25 m
11	Rc < 7.5 m

SIL	Probability of Exceeding the Horizontal Integrity Containment Radius (Rc) Without Indication (per flight hour or per sample)
0	Unknown
1	≤1×10-3
2	≤1×10-5
3	≤1×10-7*

*1×10-7 / hr reflects augmented GPS specification for Signal-in-Space integrity

References

1. Federal Aviation Administration. “FAA’s NextGen Implementation Plan 2011.” Washington D.C. March 2011.
2. “Next Generation Air Transportation System Concepts and Technology Development Program.” NASA Airspace Systems Program. May 18, 2010.
3. B. H. Wilson, “Safety Benefits of a Road Departure Crash Warning System,” *Transportation Research Board*. Rep. no. 08-1656. 2008.
4. C. Visvikis, T. L. Smith, M. Pitcher, and R. Smith. “Study on Lane Departure Warning and Lane Change Assistant Systems,” *Transport Research Laboratory.* Rep. no. PPR 374. 2008.
5. A. Vahidi and A. Eskandarian. “Research Advances in Intelligent Collision Avoidance and Adaptive Cruise Control.” *IEEE Transactions on Intelligent Transporation Systems*. Vol. 4. No. 3. September 2003.
6. U. Ozguner, C. Stiller, and K. Redmill. “Systems for Safety and Autonomous Behavior in Cars: The DARPA Grand Challenge Experience.” *Proceedings of the IEEE.* Vol. 95, No. 2. February 2007.
7. Md. Hasan, D. Alam, and S. Huq. Intelligent Car Control for a Smart Car.” *International Journal of Computer Applications.* Vol. 14, No. 3. January 2011.
8. “Minimum Aviation System Performance Standards for the Local Area Augmentation (LAAS).” RTCA, Inc. Washington, DC. Report No RTCA/DO-245A. December 2004.
9. Jung, C.R., Kelber, C.R. "A lane departure warning system using lateral offset with uncalibrated camera," *Intelligent Transportation Systems, 2005. Proceedings. 2005 IEEE*, pp. 102- 107, 13-15 Sept. 2005.
10. Enge, P. Local area augmentation of GPS for the precision approach of aircraft, Proc. IEEE, 87(1):111-132, 1999.
11. ICAO, Performance-based Navigation (PBN) Manual, Doc 9613, 2008. http://www.ecacnav.com/downloads/PBN%20Manual%20-%20Doc%209613%20Final%205.10.08%20with%20bookmarks.pdf
12. Rife, J., and Pullen, S., Aviation applications, GNSS Applications and Methods, S. Gleason and D. Gebre-Egziabher, Eds., Artech House, 2009.
13. “Evaluation of Category I LAAS to Support Airport Surface Operations.” RTCA SC-159 WG-5. May 14, 2003.
14. R. Cassel. “GBAS and Airport Surface Movement.” Era/FAA AJP-652. Reston, VA. January 16, 2009.
15. “Minimum Aviation System Performance Standards for the Local Area Augmentation (LAAS).” RTCA, Inc. Washington, DC. Report No RTCA/DO-245A. December 2004.
16. “Minimum Operation Performance Standards for GPS Local Area Augmentation System Airborne Equipment.” RTCA, Inc. Washington, D.C. Report No RTCA/DO-253C. 2008.

17. FAA. “Specification for Category 1 Local Area Augmentation System Ground Facility.” FAA-E-2937A. October 2005
18. FAA. *Chapter 3: Principles of System Safety*, FAA System Safety Handbook. 2000.
19. Rife, J., and Phelts, R.E., *Formulation of a Time-varying Maximum Allowable Error for Ground-Based Augmentation Systems*, IEEE Transactions on Aerospace and Electronic Systems, 2008, Vol. 44, No. 2, pp. 548-560.
20. Rife, J., Pullen, S., and Pervan, B. Core Overbounding and its Implications for LAAS Integrity. In *Proceedings of the ION-GNSS*, 2004, 2810-2821.
21. FAA. “Fact Sheet – Airport Surface Detection Equipment, Model X.” October 5, 2010. http://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=6296
22. E. A. Lester, R. J. Hansmann. “Benefits and Incentives for ADS-B Equipage in the National Airspace System.” MIT International Center for Air Transportation, Department of Aeronautics & Astronautics, Cambridge MA. Report No. ICAT-2007-2. August 2007.
23. “Wide Area Augmentation System (WAAS) Commissioning Information.” FAA. May 2003.
24. Sensis. “ASDE-X Solution Overview and Specifications,” East Syracuse, NY. 2008.
25. FAA. “ASDE-X Preliminary Report.” Federal Aviation Administration, National Airspace System. August 27, 2009.
26. Eurocontrol. “Eurocontrol Standard Document for Surveillance Data Exchance: Multilateration Target Reports.” European Air Traffic Management. Part 14: Category 20. December 2010.
27. M. S. Grewal, L. R. Weill, A. P. Andrews. *Global positioning systems, inertial navigation, and integration*. Hoboken, NJ. Wiley-Interscience, 2007.
28. P. Misra, P. Enge. Global Positioning System: Signals, Measurements, and Performance. Lincoln, MA. Ganga-Jamuna. 2001.
29. T. P. Waldron. “Detecting Airport Surface Movement Events Using Ground Surveillance.” Sensis Corporation, East Syracuse NY. *28th Digital Avionics System Conference (DASC).* October 2009.
30. Lanzkron, P.; Brookner, E. “Solid state X-band surface surveillance radar,” *IEEE Radar Conference*, 2007.
31. Perl, E. “Review of airport surface movement radar technology,” *IEEE Radar Conference*, 2006.
32. Lukin, K.; Mogila, A.; Galati, G.; Pavan, G. "Novel concepts for Surface Movement Radar design," *Tyrrhenian International Workshop on Digital Communications - Enhanced Surveillance of Aircraft and Vehicles, 2008.* Sept. 2008.

33. M.L. Wood, R.W. Bush. "Multilateration on Mode S and ATCRBS Signals at Atlanta's Hartsfield Airport." MIT Lincoln Laboratory. Project Report ATC-260. January 1998.
34. "Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast (ADS-B)." RTCA, Inc. Washington, DC. Report No RTCA/DO-242A. June 2002.
35. R. Rajamani, H-S Tan, B. K. Law, and W-B Zhang, "Demonstration of integrated longitudinal and lateral control for the operation of automated vehicles in platoons," *IEEE Transactions on Control Systems Technology*, vol.8, no.4, pp.695-708, Jul 2000.
36. Bishop, Richard, "Lateral/Side Sensing and Control Systems," *Intelligent Vehicle Technology and Trends*. Norwood, MA: Artech House, 2005.
37. R. Schulz, and K. Furstenberg. "Laserscanner for Multiple Applications in Passenger Cars and Trucks." *Advanced Microsystems for Automotive Applications 2006*. pp. 129-141. Berlin: Springer Berlin Heidelberg, 2006.
38. M.A. Sotelo, F.J. Rodríquez, L. Magdalena, L. M. Bergasa, and L. Boquete, "A color vision-based lane tracking system for autonomous driving on unmarked roads," *Autonomous Robots*, vol. 16, no. 1, pp. 95-116, October 2004.
39. J.C. McCall and M.M. Trivedi, "Performance evaluation of a vision based lane tracker designed for driver assistance systems," *Intelligent Vehicles Symposium, 2005. Proceedings. IEEE*, pp. 153- 158, 6-8 June 2005.
40. C. Larson, J. Raquet, and M. Veth. "Developing a Framework for Image-Based Integrity." *Proc. Institute of Navigation ION GNSS 2009,* pp. 778-789.
41. D. Simon. *Optimal State Estimation*. Wiley Interscience, 2006.
42. G. Cario, A. Casavola, G. Franze, and M. Lupia. "Predictive Time-to-Lane-Crossing Estimation for Lane Departure Warning Systems," *National Highway Traffic Safety Administration.* Pap. no. 09-0312, 2009.
43. M.L. Wood. "Multilateration System Development history and performance at Dallas/Ft. Worth Airport." *The 19th Digital Avionics Systems Conference.* 2000.
44. Baud, O.; Honore, N.; Taupin, O. "Radar / ADS-B data fusion architecture for experimentation purpose." *IEEE International Conference on Information Fusion*, 2006.
45. S. Kingsley and S. Quegan. *Understanding Radar Systems,* USA: SciTech Publishing, 1999.
46. Capezzuto, V.; Currier, E. "Acquiring new technology for US airports," *Engineering Management Conference, 2002. IEMC '02. 2002 IEEE International* , vol.1, pp. 481- 484, 2002.

47. Sensis. “High performance surface movement radar: SMRi,” East Syracuse, N.Y., 2010.
48. Evers, C., Cassel, R., and Lee, D. “Analysis of ADS-B, ASDE-3, and multilateration surveillance performance – NASA Atlanta demonstration,” *17th Annual Digital Avionics Systems Conference (DASC),* 1998.
49. M.L. Wood, R.W. Bush. "Multilateration on Mode S and ATCRBS Signals at Atlanta’s Hartsfield Airport.” MIT Lincoln Laboratory. Project Report ATC-260. January 1998.
50. Sensis Corporation Press Release. “Sensis Continues On-Time Deployment of ASDE-X – Now Operational at 32 Airports.” East Syracuse, NY. October 25, 2010.
51. Aeronautical Surveillance Panel. “Draft Manual on Multilateration Surveillance.” WP ASP03-11. October 2007.
52. International Civil Aviation Office. “Multilateration (MLAT) Concept of Use.” Edition 1.0. September 2007.
53. M. Sushko. “ADS Mode-S Channel Performance for Aircraft Surveillance Using GPS Positioning ADS.” *48th IEEE Vehicular Technology Conference.* May 1998.
54. V. Orlando. “ADS-B Mode S Extended Squitter Presentation.” MIT Lincoln Laboratory. RTCA SC-186 Working Group #3. December 2001.
55. Sensis Corporation. “Multilateration Systems Implementations.” International Civil Aviation Office Meeting. 2007.
56. H. Miyazaki, T. Koga, E. Ueda, Y. Kakubari, S. Nihei. “Evaluation Results of Airport Surface Multilateration.” EIWAC 2010. Tokyo, Japan.
57. D. Muñoz, C. Vargas. Position Location Techniques and Applications. Elsevier Inc. 2009.
58. M. Geyer, A. Daskalakis. "Solving Passive Multilateration Equations Using Bancroft's Algorithm." *The 17th Digital Avionics Systems Conference.* 1998.
59. C. Yang, E. Blasch, I. Kadar. "Geometric Factors in Target Positioning and Tracking." *12th International Conference on Information Fusion.* July 2009.
60. D.R. Jones, C.C. Quach, S.D. Young. "Runway incursion prevention system-demonstration and testing at the Dallas/Fort Worth International Airport," *20th Digital Avionics Systems Conference*. October 2001.
61. “Multilateration and ADS-B Executive Reference Guide.” ERA Corporation.
62. E. A. Lester, R. J. Hansmann. “Benefits and Incentives for ADS-B Equipage in the National Airspace System.” MIT International Center for Air Transportation, Department of Aeronautics & Astronautics, Cambridge MA. Report No. ICAT-2007-2. August 2007.

63. “Minimum Aviation System Performance Standards for Traffic Information Service – Broadcast (TIS-B).” RTCA, Inc. Washington, DC. Report No RTCA/DO-286B. October 2007.
64. P. Kemppi. “Next Generation Satellite Navigation Systems.” VTT Technical Research Centre of Finland. 2007.
65. FAA Navigation Services - Wide Area Augmentation System. http://www.faa.gov/about/office_org/headquarters_offices/ato/service_units/techops/navservices/gnss/waas/. August 2010.
66. Q. Sun, J. Zhang. "RAIM method for improvement on GNSS reliability and integrity," 2009 *Digital Avionics Systems Conference.* October 2009.
67. Y. Lee. “Receiver Autonomous Integrity Monitoring Availability for GPS Augmented with Barometric Altimeter Aiding and Clock Coasting,” *Global Positioning System: Theory and Applications, Vol. II*, Progress in Astronautics and Aeronautics Series, Vol. 164, Parkinson, B., and Spilker, J., Eds., Washington, DC: AIAA 1996.
68. C. A. Shively, D. G. O'Laughlin. "Detailed Analysis of Raim Performance for ADS-B Separation Error". Institute of Navigation. Vol. 56, No. 4. pp. 261-273. Winter 2009-2010,
69. G. Beutler, M. Rothacher, S. Schaer, T. A. Springer, J. Kouba, R. E. Neilan. **“**The International GPS Service (IGS): An interdisciplinary service in support of Earth sciences.” *Advances in Space Research.* Vol. 23. Issue 4. 1999.
70. R. B. Thompson, “Global positioning system: the mathematics of GPS receivers.” *Mathematics Magazine.* 1998.
71. R. B. Langley. “The Mathematics of GPS.” *GPS World*. 1991.
72. G. Strang, K. Borre. *Linear Algebra, Geodesy, and GPS*. Wellesley-Cambridge Press, USA. 1997.
73. “Maximizing Airport Operations Using Wide Area Augmentation System (WAAS).” FAA. Rep No. HQ-027007 INDD.
74. “Wide Area Augmentation System (WAAS) Commissioning Information.” FAA. May 2003.
75. M. S. Grewal, L. R. Weill, A. P. Andrews. *Global positioning systems, inertial navigation, and integration*. Hoboken, NJ. Wiley-Interscience, 2007.
76. R. Thomas, M. Dibenedetto. "The Local Area Augmentation System: an airport surface guidance application supporting the NASA runway incursion prevention system demonstration at the Dallas/Fort Worth International Airport." *Digital Avionics Systems Conference.* October 2001.
77. JPDO. “NextGen Avionics Roadmap, Version 1.0.” October 24, 2008.
78. “FAA SatNav News”. Fall 2010. Washington, D.C.
79. FAA Navigation Services - Local Area Augmentation System. http://www.faa.gov/about/office_org/headquarters_offices/ato/service_units/techops/navservices/gnss/laas/. August 2010.

80. International Civil Aviation Organization, CAR/SAM Regional Planning Implementation Group. "Fourth Meeting of the GNSS Task Force." September 2009.
81. Y. S. Park, S. Pullen, P. Enge. "Mitigation of Anomalous Ionosphere Threat to Enhance Utility of LAAS Differentially Corrected Positioning Service (DCPS)." Proceedings of IEEE/ION PLANS 2008. Monterey, CA. May 6-8, 2008.
82. Y. S. Park, S. Pullen, P. Enge. "Enabling the LAAS Differentially Corrected Positioning Service (DCPS): Design and Requirements Alternatives." Proceedings of ION GNSS 2009. Savannah, GA. September 22-25, 2009.
83. Y. S. Park, S. Pullen, P. Enge. "Enabling LAAS Airport Surface Movement: Mitigating Anomalous Ionospheric Threat." Proceedings of IEEE/ION Plans 2010. Palm Springs, CA. May 3-6, 2010.
84. "Automatic Dependent Surveillance--Broadcast (ADS-B) Out Performance Requirements To Support Air Traffic Control (ATC) Service." FAA. Docket No. FAA-2007-29305. May 2010.
85. "ADS-B Deployment in the United States". Briefing to Colorado Springs ION. ITT. December 2009.
86. J. Rife, S. Pullen. "The Impact of Measurement Biases on Availability for CAT III LAAS." *Proceedings of ION Annual Meeting.* Cambridge, MA. June 2005.
87. G. McGraw, T Murphy, *et al.* "Development of the LAAS Accuracy Models." *Proceeding of ION GPS 2000*. Salt Lake City, UT. September 2000.
88. P. Misra, P. Enge. Global Positioning System: Signals, Measurements, and Performance. Lincoln, MA. Ganga-Jamuna. 2001.
89. "Minimum Aviation System Performance Standards for the Local Area Augmentation (LAAS)." RTCA, Inc. Washington, DC. Report No RTCA/DO-245A. December 2004.
90. "Global Positioning System Wide Area Augmentation System (WAAS) Performance Standard." US DOT, FAA. October 2008.
91. B. Peterson, S. Lo, P. Enge. "Integrating Loran and GNSS for Safety of Life Applications." Proceedings of the Institute of Navigation GNSS Conference. Savannah, GA. September 2008.
92. FAA, William J. Hughes Technical Center. "Wide Area Augmentation System Performance Analysis Report, Report #32." April 2010.
93. C.A. Shively, T.T. Hsiao. "Availability Enhancements for Cat IIIB LAAS." Journal of the Institute of Navigation. Vol. 51, No. 1. 2004.
94. D. Ivanescu, E. Hoffman, K. Zeghal. "Impact of ADS-B Link Characteristics on the Performances of In-Trail Following Aircraft." Eurocontrol Experimental Centre, France. *American Institute of Aeronautics and Astronautics.* 2002.