

GPS-BASED AIRCRAFT LANDING SYSTEMS
WITH ENHANCED PERFORMANCE:
BEYOND ACCURACY

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Jiyun Lee

March 2005

© Copyright by Jiyeun Lee 2005
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and in quality as a dissertation for the degree of Doctor of Philosophy.

Per K. Enge (Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and in quality as a dissertation for the degree of Doctor of Philosophy.

Sam Pullen

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and in quality as a dissertation for the degree of Doctor of Philosophy.

Steve Rock

Approved for the University Committee on Graduate Studies.

Abstract

The Local Area Augmentation System (LAAS) is a differential GPS navigation system being developed to support aircraft precision approach and landing with guaranteed accuracy, integrity, continuity, and availability. To quantitatively appraise navigation integrity, an aircraft computes vertical and lateral protection levels using the standard deviations (σ) of pseudorange correction errors broadcast by the LAAS Ground Facility (LGF). Thus, one significant integrity risk is that the true standard deviation of the pseudorange correction error distribution may grow to exceed the broadcast correction error σ during LAAS operation. This event may occur due to unexpected anomalies of GPS measurements or the statistical uncertainty of the true error distribution.

This thesis presents two approaches to ensure that the error distribution based on the broadcast σ overbounds the true error distribution for a LAAS Category I (CAT I) precision approach. First, real-time σ monitoring is needed to detect violations due to unexpected anomalies with acceptable residual integrity risk. Both the statistical σ estimation method and Cumulative Sum (CUSUM) method are useful in this respect. σ estimation more rapidly detects small σ violations, while the CUSUM variant more quickly detects significant violation that would pose a larger threat to user integrity. The thesis demonstrates that these two different σ -monitoring algorithms together are capable of detecting any size of σ violations that is hazardous to users. Second, σ inflation is necessary to account for imperfect knowledge of the true error distribution. The main sources of the uncertainty are statistical estimation error during site installation and

non-stationary error distributions caused by environmental changes that affect multipath, as well as the fact that the tails of the true error distribution may not be Gaussian. A new and detailed method of sigma inflation factor determination was created and validated with test results using the Stanford LGF prototype and a “pseudo-user” receiver. This test demonstrated that sigma overbounding with the resulting inflation factor is sufficient to support LAAS CAT I operation.

Another concern related to sigma overbounding is that the conservatism applied to LAAS CAT I is no longer feasible if a navigation system requires higher performance. Thus LAAS CAT II/III precision approaches, which may need to meet tightened Vertical Alert Limit and higher availability requirements, cannot tolerate high levels of sigma inflation. This thesis describes how Position Domain Monitoring (PDM) may be used to improve system availability by reducing the inflation factor for the standard deviation of pseudorange correction errors. LAAS prototype testing using both a PDM receiver and a “pseudo-user” receiver verified the utility of PDM to enhance CAT II/III user availability. In addition, PDM helps to mitigate continuity risk using outputs of subsets of satellites in view while maintaining the required integrity. When combined with a CUSUM approach, PDM provides extra navigation integrity to users.

Acknowledgements

First and foremost, I would like to thank my advisor, Professor Per Enge, for giving me the opportunity to work on this project and for guiding me throughout my graduate career. His expertise, understanding, and continuous encouragement have made it possible for me to complete my Ph.D. degree. I appreciate his vast knowledge in many areas and his assistance in writing this dissertation. Working with him has been an honor. I am also extremely grateful to the members of my committee, Professor Donald Cox and Professor Sanjay Lall, for their input and interest in this research. My special gratitude goes to Professor Steve Rock for taking his time to read this dissertation.

I cannot fully express my thanks to Dr. Sam Pullen, director of the LAAS Laboratory, for his patience, faith, and superb guidance. He has helped me, at all levels of my development, choose research directions, extend my research strengths, and improve in my weaker areas, and thus gain confidence in my abilities. I sincerely hope I continue to learn from him during my research career.

Very special thanks go to Professor Penina Axelrad, my former advisor. Without her encouragement and endless support, I could not have pursued a graduate career in GPS research. She believed in me without question and put my interests as a student ahead of her own. She is the one professor who truly made a difference in my life.

I appreciate Professor Kyuhong Choi for helping me complete my undergraduate degree and for encouraging me to apply for graduate training. I am also grateful to Professor Sangyoung Park for his support and interest in my research.

I would like to thank all my colleagues in the Stanford GPS laboratory for their help and friendship, specially the LAAS laboratory members: Ming Luo, Dennis Akos, Jason Rife, Gang Xie, Hiroyuki Konno, Sasha Mitelman, and Per-Ludvig Normark. My thanks also go to Fiona Walter for proofreading this dissertation. I would like to acknowledge the Federal Aviation Administration for its financial support of this project. Appreciation also goes to Jeff Wade and Tim Huynh for technical assistance and to Lynn Kaiser, Robin Heinen, Sherann Ellsworth, and Dana Parga for their administrative help.

Finally, I would like to thank my parents – Wonchun Lee, the most extraordinary businessman and father, and Jungrye Seo, the most astonishingly talented ceramic artist and mother, I have ever known – for their endless love, support, and encouragement.

Table of Contents

GPS-BASED AIRCRAFT LANDING SYSTEMS WITH ENHANCED PERFORMANCE: BEYOND ACCURACY	I
ABSTRACT	IV
ACKNOWLEDGEMENTS	VI
TABLE OF CONTENTS	VIII
LIST OF FIGURES	XI
LIST OF TABLES	XIV
GLOSSARY OF ACRONYMS	XV
CHAPTER 1	1
INTRODUCTION	1
1.1 THE GLOBAL POSITIONING SYSTEM (GPS)	2
1.1.1 GPS SYSTEM SEGMENTS	2
1.1.2 SIGNALS	4
1.1.3 MEASUREMENTS AND ERROR SOURCES	5
1.2 GPS AND AVIATION NAVIGATION	9
1.2.1 DIFFERENTIAL GPS	10
1.2.2 CIVIL AIRCRAFT NAVIGATION	12
1.3 AUGMENTATION SYSTEMS	13
1.3.1 LOCAL AREA AUGMENTATION SYSTEM (LAAS)	14
1.3.2 WIDE AREA AUGMENTATION SYSTEM (WAAS)	15
1.4 PROTECTION LEVEL CALCULATION	17
1.5 REAL TIME ERROR BOUNDING	20
1.5.1 CHALLENGES IN ERROR BOUNDING	20
1.5.2 PREVIOUS WORK	21

1.6 OUTLINE AND CONTRIBUTIONS	23
CHAPTER 2	25
THE LOCAL AREA AUGMENTATION SYSTEM	25
2.1 INTRODUCTION	25
2.2 LAAS ARCHITECTURE OVERVIEW	26
2.3 LAAS REQUIREMENTS	27
2.4 LAAS GROUND FACILITY PROCESSING ALGORITHMS	31
2.4.1 CARRIER SMOOTHING AND PSEUDORANGE CORRECTIONS	31
2.4.2 GROUND FACILITY ERROR STANDARD DEVIATION	33
2.5 STANFORD LAAS INTEGRITY MONITOR TEST-BED (IMT)	34
2.5.1 IMT HARDWARE CONFIGURATION	35
2.5.2 IMT FUNCTIONS	36
2.5.2.1 SIGNAL QUALITY MONITORING (SQM)	37
2.5.2.2 DATA QUALITY MONITORING (DQM)	38
2.5.2.3 MEASUREMENT QUALITY MONITORING (MQM)	39
2.5.2.4 PHASE ONE OF EXECUTIVE MONITORING (EXM-I)	41
2.5.2.5 MULTIPLE REFERENCE CONSISTENCY CHECK (MRCC)	43
2.5.2.6 SIGMA-MEAN ($\sigma\mu$) MONITORING	46
2.5.2.7 MESSAGE FIELD RANGE TEST (MFRT)	46
2.5.2.8 PHASE TWO OF EXECUTIVE MONITORING (EXM-II)	46
2.6 CONCLUSION	48
CHAPTER 3	50
SIGMA-MEAN MONITORING	50
3.1 INTRODUCTION	50
3.2 THREAT SPACE	51
3.3 SIGMA MONITORING	52
3.3.1 SIGMA ESTIMATION METHOD	52
3.3.1.1 ALGORITHM	52
3.3.1.2 THEORETICAL ANALYSIS	55
3.3.2 SIGMA CUMULATIVE SUM METHOD	56
3.3.2.1 ALGORITHM	56
3.3.2.2 THEORETICAL ANALYSIS	60
3.3.3 IMT TEST RESULTS	62
3.3.3.1 NOMINAL TESTING	62
3.3.3.2 FAILURE TESTING	65
3.4 MEAN MONITORING	70
3.4.1 MEAN ESTIMATION METHOD	70
3.4.2 MEAN CUMULATIVE SUM METHOD	70
3.4.2.1 ALGORITHM	71
3.4.2.2 THEORETICAL ANALYSIS	71
3.4.3 IMT TEST RESULTS	72
3.4.3.1 NOMINAL TESTING	72
3.4.3.2 FAILURE TESTING	74
3.5 COMPARISON OF ESTIMATION AND CUSUM RESULTS	77
3.6 CONCLUSION	80
CHAPTER 4	82
SIGMA INFLATION	82
4.1 INTRODUCTION	82
4.2 SIGMA INFLATION FACTOR DETERMINATION METHOD	83
4.2.1 FINITE SAMPLE SIZE	83
4.2.2 PROCESS MIXING	84

4.2.3 LIMITATION OF SIGMA MONITORS.....	86
4.2.3.1 GAUSSIAN ASSUMPTION ON ERROR MODEL	86
4.2.3.2 NON-GAUSSIAN ASSUMPTION ON ERROR MODEL	87
4.2.4 TOTAL INFLATION FACTOR	88
4.3 PERFORMANCE ANALYSIS	90
4.3.1 STANFORD LAAS PERFORMANCE TEST-BED	90
4.3.2 PERFORMANCE TEST RESULTS	91
4.4 CONCLUSION	93
CHAPTER 5	94
POSITION DOMAIN MONITORING.....	94
5.1 INTRODUCTION.....	94
5.2 POSITION DOMAIN MONITORING (PDM).....	96
5.2.1 PDM HARDWARE CONFIGURATION	97
5.2.2 PDM ALGORITHM.....	98
5.2.3 THRESHOLD DERIVATION	101
5.2.4 NOMINAL TESTING.....	104
5.2.5 FAILURE TESTING	106
5.3 SIGMA INFLATION IN POSITION DOMAIN	109
5.3.1 ERROR DISTRIBUTIONS	109
5.3.2 SIGMA INFLATION FACTOR.....	111
5.4 PERFORMANCE ANALYSIS	113
5.5 USE OF POSITION DOMAIN MONITOR MEASUREMENTS.....	117
5.5.1 PDM CUMULATIVE SUM (CUSUM) MONITORING.....	117
5.5.2 SCREENING PROCESS	120
5.6 CONCLUSION	124
CHAPTER 6	125
CONCLUSION.....	125
6.1 SUMMARY OF CONTRIBUTIONS	125
6.1.1 SIGMA-MEAN ESTIMATION AND MONITORING	126
6.1.2 SIGMA INFLATION AND PERFORMANCE	128
6.1.3 POSITION DOMAIN MONITORING	128
6.2 SUGGESTIONS FOR FUTURE WORK.....	130
APPENDIX A	132
LAAS ERROR MODELS FOR NOISE, MULTIPATH, TROPOSPHERE, AND IONOSPHERE.....	132
A.1 MODEL OF AIRBORNE PSEUDORANGE PERFORMANCE.....	132
A.2 MODEL OF TROPOSPHERIC RESIDUAL UNCERTAINTY.....	133
A.3 MODEL OF IONOSPHERIC RESIDUAL UNCERTAINTY.....	134
APPENDIX B.....	136
CUMULATIVE SUM (CUSUM) DESIGN.....	136
B.1 THE GENERAL EXPONENTIAL FAMILY	136
B.2 DERIVATION OF CUSUM FOR A NORMAL SIGMA-MEAN SHIFT	138
B.3 MARKOV CHAIN ANALYSIS OF TRANSITION PROBABILITY.....	139
BIBLIOGRAPHY	141

List of Figures

Figure 1.1: GPS Space Segment (Courtesy: FAA)	2
Figure 1.2: GPS Operational Control Segment Facilities (The MCS located at Colorado Springs; USAF monitor stations at Colorado Springs, Ascension Island, Diego Garcia, Kwajalein, Hawaii).....	3
Figure 1.3: GPS Signal Structure Showing Relations Between the Carrier, Code, and Navigation Data. The C/A coded signal on 1575.42 MHz is used as an example.	5
Figure 1.4: Error Sources in GPS Measurements.....	8
Figure 1.5: Differential GPS	11
Figure 1.6: Vertical Alert Limit and Horizontal Alert Limit	13
Figure 1.7: Local Area Augmentation System (LAAS) Overview	14
Figure 1.8: Wide Area Augmentation System (WAAS) architecture	16
 Figure 2.1: Local Area Augmentation System (LAAS) Architecture	 26
Figure 2.2: Aviation Navigation Requirements.....	28
Figure 2.3: Precision Approach and Landing Categories [27, 28].....	29
Figure 2.4: IMT Hardware Configuration	35
Figure 2.5: IMT Hardware Configuration	36
Figure 2.6: An Example of Selecting a Common Set [30]	42
Figure 2.7: MRCC Fault Detection Flowchart [38]	44
Figure 2.8: EXM-II Pre-Screen Flowchart [38]	45
Figure 2.9: EXM-II Flowchart [38]	47

Figure 3.1: Chi-Square Distribution of Sigma Estimate	54
Figure 3.2: Performance of Sigma Estimation Method and MRCC Test.....	55
Figure 3.3: CUSUM Performance Modeling with Markov Chains.....	59
Figure 3.4: Failure-State ARLs for Sigma CUSUM Method	61
Figure 3.5: Threshold for Sigma CUSUM Method.....	61
Figure 3.6: Sigma Estimation Results from IMT Nominal Data	63
Figure 3.7: Zero-Start CUSUM Result from IMT Nominal Data	64
Figure 3.8: FIR CUSUM Result from IMT Nominal Data.....	64
Figure 3.9: Sigma Estimation Results from Failure Test	66
Figure 3.10: FIR CUSUM Results from IMT Failure Test.....	68
Figure 3.11: FIR CUSUM Results of Nominal RR from IMT Failure Test	68
Figure 3.12: Thresholds and Failure-State ARLs for Mean CUSUM Monitor.....	72
Figure 3.13: Mean Estimation Results from IMT Nominal Data	73
Figure 3.14: Mean FIR CUSUM Results from IMT Nominal Data.....	74
Figure 3.15: Mean Estimation Results from IMT Failure Test with L=0.8 Injected on Channel (RR2, SV 2)	75
Figure 3.16: Mean FIR CUSUM Results from IMT Failure Test with L=0.8 Injected on Channel (RR2, SV 2)	76
Figure 3.17: Time-to-Alert for Sigma CUSUM and Sigma Estimation Monitors	78
Figure 3.18: Time-to-Alert with $P_{MD}<0.001$ for Mean Estimation and Mean CUSUM Monitor Performance	79
Figure 4.1: Probability Density Function of the Normalized B -values	85
Figure 4.2: Failure-State Average Run Lengths for CUSUM and Sigma Estimation Monitors	86
Figure 4.3: Probability Density Function of Sample Standard Deviation	88
Figure 4.4: Inflation Factor Determination Method for Broadcast σ_{pr_gnd}	89
Figure 4.5: Stanford LAAS Performance Test-bed Hardware Configuration.....	91
Figure 4.6: Stanford LAAS Pseudo-User Performance	93
Figure 5.1: Stanford LAAS Pseudo-User Performance	95
Figure 5.2: IMT-PDM Hardware Configuration	98
Figure 5.3: Probability Density Function of Normalized Vertical Position Errors	103
Figure 5.4: Cumulative Distribution Function of Normalized Vertical Position Errors	103

Figure 5.5: Normalized Vertical Position Errors and Detection Thresholds from IMT-PDM Nominal Data (All Approved SVs in View).....	104
Figure 5.6: Normalized Vertical Position Errors and Detection Thresholds from IMT-PDM Nominal Data (all “one-SV-out” combinations)	105
Figure 5.7: Normalized Vertical Position Errors and Detection Thresholds from IMT-PDM Nominal Data (“two-SV-out” combinations)	105
Figure 5.8: IMT-PDM Sigma Failure Test with L=3 (All Approved SVs in view).....	107
Figure 5.9: IMT-PDM Sigma Failure Test with L=8 (All Approved SVs in view).....	107
Figure 5.10: Error Distributions in Position Domain and in Range Domain	110
Figure 5.11: Probability Density Function of the Normalized Vertical Position Errors (Error Distribution in Position Domain)	111
Figure 5.12: Inflation Factors for Broadcast σ_{pr_gnd} with RDM Only and RDM+PDM.....	112
Figure 5.13: Stanford LAAS Performance Test-bed IMT-PDM-User Hardware Configuration....	113
Figure 5.14: System Performance in Vertical Direction with RDM and PDM.....	115
Figure 5.15: System Performance in Vertical Direction with RDM only	115
Figure 5.16: PDM-CUSUM Results from Nominal Data.....	118
Figure 5.17: PDM-CUSUM Results from Failure Test (3 x Error Sigma on All SV and All RR) .	119
Figure 5.18: Use of PDM Screening Process Outputs to Enhance Average Continuity.....	121
Figure 5.19: The Worst-case VPL_{H0} Out Of All "Two-SV-Out" combinations.....	122
Figure 5.20: Increase Detection Thresholds Such That Effective $VPL_{H0} = VAL$	122
Figure 5.21: Prior Probability Density Function for Out-Of-Control σ	123
 Figure 6.1: Time-to-Alert for Sigma CUSUM and Sigma Estimation Monitors	127
Figure 6.2: Time-to-Alert with $P_{MD}<0.001$ for Mean Estimation and Mean CUSUM Monitor Performance	127
Figure 6.3: Stanford LAAS Pseudo-User Performance	129
Figure 6.4: Real-Time Use of Bayesian CUSUM Outputs to Preserve Continuity and Availability while Protecting Integrity	131
 Figure B.1: CUSUM Performance Modeling with Markov Chains	140

List of Tables

Table 1-1: A Summary of Error Size in GPS Measurements	11
Table 2-1: Requirements for Precision Approach and Landing	30
Table 2-2: Ground Facility Error Allocation Model	34
Table A-1: Airborne Error Model Parameters.....	133

Glossary of Acronyms

Local-Area DGPS	
AAD	Airborne Accuracy Designator
CAT I/II/III	Different categories of the Local-Area Augmentation System (LAAS) precision approaches
GAD	Ground Accuracy Designator
GBAS	Ground Based Augmentation System
HAL	Horizontal Alert Limit
HMI	Hazardously Misleading Information
HPL	Horizontal Protection Level
IMT	Integrity Monitor Test Bed
LAAS	Local Area Augmentation System
LGF	LAAS Ground Facility
MOPS	Minimum Operational Performance Standards. Refers to RTCA159 specification for the Local Area Augmentation System (LAAS).
NAVSTAR	Navigation Satellite Timing and Ranging
PDM	Position Domain Monitoring
Protection Level	Broadcast indication of the bound on the accuracy of the state. This value is compared to the Alert Limit to determine if a flight operation can begin or continue.
RDM	Range Domain Monitoring

VAL	Vertical Alert Limit
VPE	Vertical Position Error
VPL	Vertical Protection Level

Chapter 1

Introduction

Local Area Augmentation of GPS is being developed to become the primary navigational aid in civil aircraft precision approach and landing. While the system promises great performance, a number of technical obstacles have been encountered in meeting aviation requirements. These obstacles include statistical uncertainties in the knowledge of the pseudorange correction error standard deviation (sigma) and potential changes of these “sigmas”. The broadcast sigmas are used by the aircraft to compute their position bounds. If the true sigma exceeds the broadcast sigma, increased integrity risk results. In this thesis, two approaches are presented to ensure that the error distribution based on the broadcast sigmas overbounds the true error distribution. The first method is real time sigma monitoring, based on the measurements of pseudorange correction error, which estimates sigma and detects anomalies. The second method is sigma inflation [1], which compensates for the uncertainty of the true error distribution. In addition, the thesis describes how position-domain monitoring may be used to support precision approaches with more stringent requirements.

In this chapter, we first present some background on GPS and explain how to enhance GPS with differential techniques to become an aviation navigation aid. This is followed by a description of GPS augmentation systems and how they provide error bounds in real time and consequently guarantee flight safety. We then focus on error bounding using the sigma

values broadcast by GPS augmentation systems and present the motivation for this thesis. Next, previous work in related fields is presented. Finally, contributions are given along with an outline of the thesis.

1.1 THE GLOBAL POSITIONING SYSTEM (GPS)

The NAVSTAR Global Positioning System is a space-based radio-navigation system. This satellite system is deployed and managed by the U.S Department of Defense (DoD) originally to provide accurate information of position, velocity and time to military forces. However, GPS also provides significant benefits to civil users. The civil community has developed an increasingly large variety of applications in space and marine navigation, vehicle transportation, civil aviation, auto-farming, surveying and mapping, telecommunications, public safety, and outdoor leisure activities. Today, GPS serves nearly 20 million users worldwide [2], and the vast majority are civilians.

1.1.1 GPS SYSTEM SEGMENTS




Figure 1.1: GPS Space Segment (Courtesy: FAA)

GPS is comprised of three segments: the Space Segment, the Control Segment, and the User Segment. The space segment consists of at least 24 nominal satellites which are positioned in six nearly circular orbital planes with an orbital radius of 26,560 km and a

period of 11 hr 58 min or one-half of a sidereal day (After two rotations, each satellite rises at the same spot, but four minutes earlier than the day before [3]). These satellites provide the ranging signals and data messages to the user's equipment.




Figure 1.2: GPS Operational Control Segment Facilities (The MCS located at Colorado Springs; USAF monitor stations at Colorado Springs, Ascension Island, Diego Garcia, Kwajalein, Hawaii)

The Operational Control Segment (OCS) operates the system and maintains the satellites in space. It monitors satellite orbits and satellite health and maintains GPS time. There are five monitor stations spread around the world, as shown in Figure 1.2. These stations passively track the satellites and transmit raw data and the received navigation message to the Master Control Station (MCS) located at Colorado Springs. The MCS then predicts satellite ephemerides and clock corrections and updates satellite navigation messages which are essential for users to estimate position, velocity and time. The user segment (i.e., GPS receivers) processes ranging signals transmitted from the satellites and performs the navigation. A GPS receiver acquires the locations of satellites based on the received navigation messages and measures the distance between the user and satellites in terms of transit time of the signal from satellites to users. To estimate position precisely using trilateration, accurate timing is essential. This is accomplished by synchronizing satellite atomic clocks very accurately. Although the clocks in the satellite and the receiver also must be synchronized to measure the true transit time of signals, this condition is generally

not met by the inexpensive quartz oscillators in most GPS receivers. GPS receivers, therefore, need at least four satellites in view to solve for the three-dimensional user position and receiver clock bias. In other words, four observation equations are needed to solve for four unknowns (x, y, z, f) .

1.1.2 SIGNALS

The GPS satellites transmit two radio frequencies: L1 centered at 1575.42 MHz and L2 centered at 1227.60 MHz. These carriers are modulated with two types of codes and a navigation message. The two types of codes are the coarse/acquisition pseudorandom noise (PRN) code (C/A-code) on L1 carrier phase and the precision (encrypted) code (P(Y)-code) on both L1 and L2. P(Y)-code is accessible only to authorized users, and while C/A-code is provided for all users. Though current civil users can only access L1 C/A-code, there are receiver variations types such as codeless L2 tracking receivers that enable users to obtain centimeter-level measurement accuracy by utilizing the carrier phases of both the L1 and L2 frequencies.

The GPS C/A-code is a Gold code [4] with a unique sequence length of 1023 bits, called “chips.” Since the chipping rate of the C/A-code is 1.023 MHz, the C/A-code is repeated each millisecond. The duration of each C/A-code chip is about $1\ \mu\text{s}$ as shown in Figure 1.3, and the corresponding distance is about 300 m. The sequence length of P(Y)-code is extremely long (about 10^{14} chips) [3] and the repetition period is one week. Since the P(Y) code has a smaller “wavelength” of 30 m and, equivalently, a chipping rate of 10.23 MHz, the precision in range measurements is much greater than that for the C/A-code. The spread spectrum codes are designed to provide range measurements by having “peaked” auto-correlation functions. In addition, the unique PRN sequences associated with each satellite are nearly uncorrelated with respect to each other. This property allows all satellites to transmit at the same frequency without any time-sharing. This modulation technique is called code division multiple access (CDMA) and is used for separating and detecting the GPS signals [5].

Each PRN code is modulated with navigation data, which is a binary code message transmitted at 50 bits per second (bps) [3]. The bit duration of the navigation message is 20 ms as shown in Figure 1.3. The information contents of the message are satellite clock corrections, health status, ephemeris parameters, and almanac. This combined binary signal using modulo-2 addition then modulates the carrier using a specific technique, called binary phase shift keying (BPSK) [3].

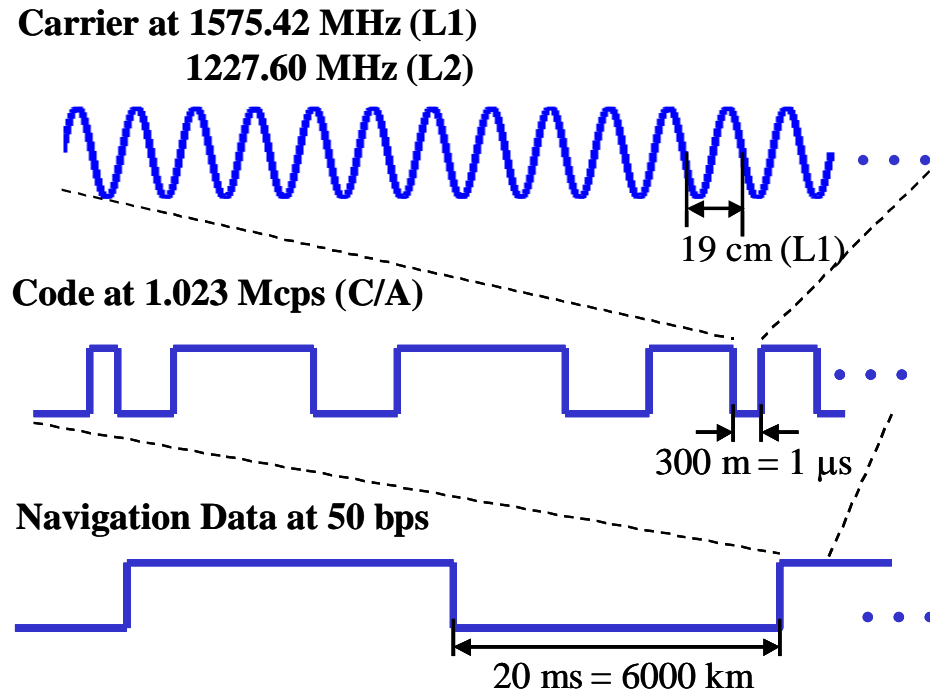


Figure 1.3: GPS Signal Structure Showing Relations Between the Carrier, Code, and Navigation Data. The C/A coded signal on 1575.42 MHz is used as an example.

1.1.3 MEASUREMENTS AND ERROR SOURCES

Two types of measurements are of interest to GPS users. One is the pseudorange, which is the distance between the satellite and the receiver plus a bias due to the difference in the user clock from the GPS clock. Pseudoranges are a measure of the travel time of the PRN codes. To acquire a signal, first the receiver replicates the PRN code that is transmitted by the satellite. Then it attempts to shift the replica in time until it is aligned with the incoming PRN code. When the code replica matches the incoming code, the correlation is maximized.

At that point, the time shift required to achieve the maximum correlation is the transit time of the signal modulo 1 ms. The transmission time is “marked” on the signal with the satellite clock, and the reception time can be read from the receiver clock. The pseudorange, ρ , is determined from multiplying the transit time by the speed of light. Code lock is maintained by a feedback control loop, called a delay lock loop (DLL), which continuously aligns the replica code with the incoming signal. Within the DLL, the PRN code is removed from the signal, and the carrier (modulated by the navigation message) is available for further processing.

The second measurement, the carrier phase, is the difference between the received phase and the phase of a receiver oscillator at the epoch of measurement. The receiver continues tracking the carrier modulated by the navigation data with a phase lock loop (PLL). The PLL attempts to match the phase of the receiver-generated signal to that of the incoming signal. With the PLL, the receiver can measure only a partial cycle. However, this partial cycle, when combined with an initial unknown number of whole cycles, also indicates the range to the satellite. In order to take full advantage of the carrier phase measurements, ϕ , we need to resolve this unknown number of whole cycles, called the integer ambiguity [3]. The PLL also measures the Doppler shift, which can be converted into a pseudorange rate (this measurement is used for ultra-precise static and kinematic surveying or for attitude determination). After the phase lock is accomplished, the navigation message is extracted.

The GPS observation equations for code and carrier phase measurements are:

$$\rho = R_m^n + c(b_m - B^n) + I_m^n + T_m^n + M_m^n + v_m^n \quad (1-1)$$

$$\phi = R_m^n + c(b_m - B^n) - I_m^n + T_m^n + N_m^n \lambda + p_m^n + \varepsilon_m^n \quad (1-2)$$

where,

ρ is the measured code phase measurement, or pseudorange,

ϕ is the measured carrier phase measurement,

R_m^n is true range from satellite n to receiver m ,

b_m is the receiver clock bias (offset from GPS time),

B^n is the satellite clock bias (offset from GPS time),

I is the ionospheric delay,

T is the tropospheric delay,

M, p are multipath errors,

N is the integer ambiguity,

λ is the carrier wavelength (for L1 frequency, $\lambda_{L1} = \frac{c}{f_{L1}} \approx 19 \text{ cm}$),

v represents other code phase measurement errors, and

ε represents other carrier phase measurement errors.

As shown in these observation equations, GPS measurements are subject to various errors. It is important to understand the effects of the measurement errors, since the quality of PVT estimates depends on the quality of the range and range-rate measurements.




Figure 1.4: Error Sources in GPS Measurements

The primary GPS error sources are illustrated in Figure 1.4. These errors can be grouped into three categories [3]. The first set is due to control segment imperfections. The satellite ephemeris and clock parameters estimated by the control segment are broadcast to the user receiver. The satellite ephemeris error is the difference between the actual position and velocity of a satellite and those predicted by the broadcast ephemeris model. This error is typically 1-2 m in the root mean square (rms) sense. The satellite clock bias, the difference between the true clock and the satellite clock, also introduces about a 1-2 m range error in the rms sense. Civil users were also compromised by Selective Availability (SA), which intentionally dithered the clock to cause about 22 m error in rms [6] until it was deactivated on May 2, 2000 by Presidential decision [3].

The second set of errors is introduced by uncertainties in the propagation mediums: the ionosphere and troposphere. The ionosphere is a region of ionized gases which affect the speed of GPS signal propagation from a satellite to a receiver. The code phase

measurements are delayed while the carrier phase measurements are advanced, as shown in Equations (1-1) and (1-2). Since this delay is inversely proportional to the signal frequency, dual frequency users can remove this error by themselves. Single frequency users can reduce this delay by approximately 50% after utilizing the Klobuchar ionospheric model broadcast in GPS navigation data [7]. The resulting ranging error, proportional to the total electron content (TEC) in the ionosphere, is about 1-5 m. The dry gases and water vapor composing the troposphere refract GPS signals and introduce an additional delay. The delay is small for satellites directly overhead and larger for low-elevation satellites. This tropospheric delay can be corrected using atmospheric models [8]. If corrected based on average meteorological conditions, the resulting error is about 0.1-1 m.

The remaining errors are multipath and receiver noise. Multipath errors are caused by the interfering signals reflected from surfaces. Since the code and carrier measurements are based on the sum of the direct and reflected signals, the ranging error depends on the strength of the reflected signal and the delay between direct and reflected signals [3]. Multipath affects code measurements with a 1-5 m error and carrier measurements with a 1-5 cm rms error. Adopting a multipath-limiting antenna, a narrow correlator receiver or carefully choosing an installation site for the antenna can reduce these errors. Finally, receiver noise errors are due to thermal noise in the receiver front end, multi-access interference, and signal quantization noise. The receiver noise introduces less than 0.5 m of code measurement error and about 1-2 mm of carrier phase measurement error.

1.2 GPS AND AVIATION NAVIGATION

In an effort to make GPS service available to commercial, national and international civil users while maintaining the original U.S military function, two GPS services are provided. DoD authorized users have access to the Precise Positioning Service (PPS), which provides full system accuracy by utilizing extremely long and fast P(Y)-code (detailed in Section 1.1.2). Access to PPS is restricted by cryptographic techniques, and users must be equipped with a decryption device to lock onto the encrypted P-code, referred to as the Y-code. This feature is called Anti-Spoofing (AS). The Standard Positioning Service (SPS) is provided to civilian and all other users throughout the world with a less accurate positioning

capability than PPS. Without SA, current GPS/SPS provides position accuracy of approximately 10 m (with 95% confidence) in the horizontal direction and 15 m (95%) in the vertical direction.

A significant civil application of GPS is aviation navigation. With air travel doubling in the 21st century, the aviation community is already relying extensively on GPS. The economy and safety of aircraft navigation, supported by on-board inertial navigation systems and ground-based radionavigation aids in the past, are now greatly enhanced with GPS. However, civil aviation requires greater accuracy than what GPS alone can provide. For instance, Category I Precision Approaches requires navigation sensor errors below 1 meter. The required accuracy of Category III precision approaches is even higher: the ranging error is restricted to decimeter-level [3]. These precision approach operations will be described in Chapter 2, which will also give the performance requirements in detail. This thesis focuses on how to augment GPS to be a primary system for precision approaches and auto-landing as well as en route and surface traffic surveillance.

1.2.1 DIFFERENTIAL GPS

As addressed earlier, standalone GPS is not capable of supporting all phases of flight from cruise to landing due to insufficient accuracy. In this context, the use of Differential GPS (DGPS) enhances standalone GPS accuracy. The basic concept of DGPS, shown in Figure 1.5, lies in the mitigation of measurement errors with one or more stationary reference receivers viewing the same satellite as the roving users. DGPS places reference receivers at precisely surveyed locations. The biases associated with the worst error sources are similar if a user receiver is close to the reference receivers. DGPS estimates the errors in the reference measurements and broadcast these errors as correction. All users in the coverage area can then use differential corrections to improve their navigation accuracy. Since most of the ionospheric, tropospheric, satellite ephemeris and clock errors are correlated between receivers spatially and temporally, residual correction errors are small, as shown in Table 1-1 [3]. On the other hand, multipath and receiver noise errors are uncorrelated between reference and roving receivers, and cannot be corrected by DGPS. However, these types of errors can be mitigated through receiver design, antenna design, and siting.




Figure 1.5: Differential GPS

Source	Error Size (GPS/SPS)	Residual Error (DGPS)
Satellite Clock Model	1 – 2 m (rms)	0.0 m
Satellite Ephemeris Prediction	1 – 2 m (rms)	0.1 m (rms)
Ionospheric Delay	2 – 10 m in zenith direction	0.2 m (rms)
Tropospheric Delay	2.3 – 2.5 m in zenith direction at sea level	0.2 m (rms) plus altitude effect
Multipath	Code: 1 – 5 m Carrier: 1 – 5 cm	Uncorrelated between reference and rover
Receiver Noise	Code: 0.5 m (rms) Carrier 1 – 2 mm (rms)	Uncorrelated between reference and rover

Table 1-1: A Summary of Error Size in GPS Measurements

1.2.2 CIVIL AIRCRAFT NAVIGATION

Improving accuracy with DGPS is not enough to support aircraft operations. These operations also demand flight safety and reliability. For a better understanding of these demands, Figure 1.6 illustrates two concepts that are commonly used to describe aviation navigation systems. In order to conduct a safe flight, the pilot or aircraft guidance system should be alerted if the position error exceeds a certain bound. This bound (an outer red box in Figure 1.6) is defined as the Alert Limits: the Vertical Alert Limit (VAL) in the vertical direction and the Horizontal Alert Limit (HAL) in the horizontal direction. If the aircraft lies outside the box without any alarm, it may crash into an obstacle. For that reason, the pilot needs an error bound in real time. This error bound, indicating how poor the position fixes can be, is called the *protection level*: the Vertical Protection Level (VPL) in the vertical direction and Horizontal Protection Level (HPL) in the horizontal direction. As shown in Figure 1.6, for safety we need the protection levels always to be smaller than the alert limits corresponding to the current phase of flight. The computation of protection levels will be discussed in Section 1.4.

Unlike land-based equipment, GPS accuracy varies significantly with time. As described earlier, several sources of error can corrupt the pseudorange measurement, and the position error thus varies. In view of this, computing protection bounds is necessary to obtain an assurance for the position solution at a certain level. In addition, for the safety of aircraft guidance, the system needs to provide warnings quickly enough for an aircraft to act when its position error exceeds alert limits. To fulfill these requirements, the concept of the augmentation system was introduced as an application of DGPS. Ground-based augmentation systems (GBAS) and Space-based Augmentation Systems (SBAS) are two major categories. They will be introduced in Section 1.3.




Figure 1.6: Vertical Alert Limit and Horizontal Alert Limit

1.3 AUGMENTATION SYSTEMS

Augmentation systems enhance GPS position estimates by sending differential corrections to a user and by improving satellite geometry. The reference stations also broadcast warnings of any system malfunctions and the quality of the corrections to the user in such a way that the system helps insure flight safety. These systems are categorized as “ground-based” or “space-based” augmentation systems depending on the coverage area and how they improve the geometry. GBAS is designed to provide service in a local area (within several kilometers to tens of kilometers). Reference receivers are placed close to each other and determine the measurement errors at their locations. This system may include pseudolites (to augment the geometry), which are GPS-like ranging signals radiated from the ground. With these pseudolite signals, the system accuracy improves and sensitivity to the failure of any GPS signals is reduced. In contrast, SBAS operates in a wide area up to continental coverage [9]. A network of ground receivers at precisely known locations continually updates its error estimates and makes the correction available for each

monitored satellite [5]. This system requires geostationary satellites, which broadcast the correction message over the entire region of coverage and also augment the geometry with GPS-like signals. The Federal Aviation Administration (FAA) is developing a GBAS called the Local Area Augmentation System (LAAS) and an SBAS known as the Wide Area Augmentation System. Although this thesis concentrates on how to improve the performance of LAAS, my work is also relevant to WAAS. Both LAAS and WAAS will be described in Sections 1.3.1 and 1.3.2 respectively.

1.3.1 LOCAL AREA AUGMENTATION SYSTEM (LAAS)




Figure 1.7: Local Area Augmentation System (LAAS) Overview

LAAS is a local area differential GPS system because it typically serves receivers close to the reference station. Differential corrections are computed based on the surveyed location of multiple nearby reference receivers and broadcast to an approaching airplane using a

VHF data link. The broadcast data also contains any alerts on system failures and error bounds on corrections. A LAAS user first measures the pseudoranges to the GPS satellites, and he/she then determines which satellites can be used safely based on the LAAS message and corrects their ranging measurements. The user also computes a VPL and an HPL in real time using the information on error bounds. These are then compared to the VAL and HAL, respectively, to determine whether the system provides safety to the user.

Due to the proximity between the reference station and LAAS users, the pseudorange error components, which are common to all receivers within the local geographical area, can be nearly cancelled, and thereby sub-meter accuracy is achieved. The spatially correlated errors increase as the separation of the user from the reference station increases, and accordingly LAAS performance degrades. In general, LAAS is more accurate than SBAS if the user is within 100 km or so of the reference receiver.

LAAS will provide many benefits for all users. It is capable of supporting Category I and II/III precision approaches, as will be explained in depth in Chapter 2. With LAAS, curved precision approaches will also be possible, while these approaches cannot be conducted using current instrument landing systems (ILS) [10]. Unlike ILS, which requires multiple installations to serve multiple runways, a single LAAS reference station will typically provide precision approach capability to all runways at an airport [10].

1.3.2 WIDE AREA AUGMENTATION SYSTEM (WAAS)

In contrast to LAAS, WAAS offers coverage over a continent-wide area based on the concept of wide-area DGPS [11]. WAAS was made operational over the Conterminous United States (CONUS) by the FAA on July 10, 2003. There are several WAAS-like GPS augmentations under development: the European Geostationary Navigation Overlay System (EGNOS) in Europe [12], the Multifunction Transportation Satellite (MTSAT)-Based Satellite Augmentation System (MSAS) in Japan [13], and the GPS and GEO Augmented Navigation (GAGAN) in India.

The architecture of WAAS is illustrated in Figure 1.8. The master station collects observation data from about 25 WAAS reference stations (WRS) distributed over the

CONUS and in neighboring regions. It then generates two corrections for each satellite: one for the satellite clock and the other for the three dimensional location of the satellite. Because dual frequency (L1-L2) measurements are available, the master station also estimates a set of corrections for the ionospheric delay. The WAAS data, which are the differential corrections and their error bounds, are coded in a 250-bps navigation message of GPS-like signals at L1. The message is uploaded to a geostationary satellite and transmitted back to users. The geostationary satellite not only serves as a data link but also as a potentially valuable source of ranging.

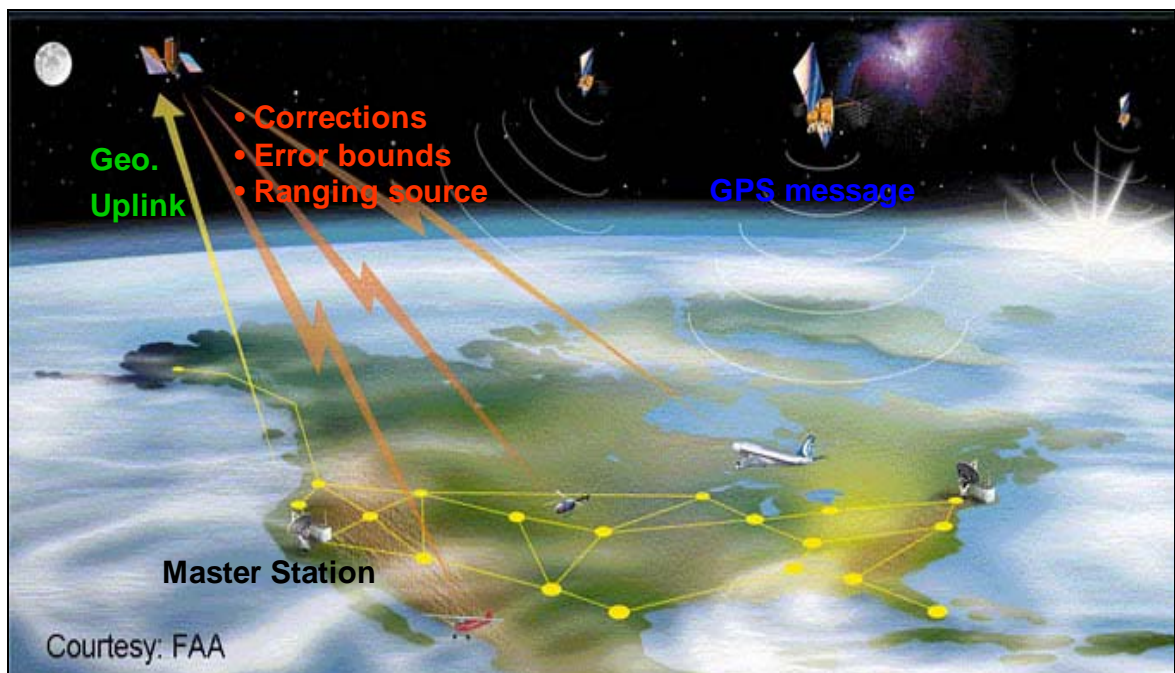


Figure 1.8: Wide Area Augmentation System (WAAS) architecture

1.4 PROTECTION LEVEL CALCULATION

Most importantly, augmentation systems provide real-time error bounds. As noted in Section 1.2.2, these bounds are called protection levels (PLs). They are defined to meet the following requirement:

$$\text{Prob}(\text{error} > PL) \leq \gamma \quad (1-3)$$

Namely, the protection level (PL) must overbound the true position error, which is unknown in real time, with a probability of one minus γ (γ differs by application and is on the order of $10^{-7} \sim 10^{-10}$). In this section we describe how PLs are calculated based on the information broadcast through the VHF data link by a user of LAAS, which is the primary system addressed in this thesis. (Refer to [10] for PL computation in WAAS.)

Let us first build a simple model for the corrected pseudorange measurement from the i th satellite,

$$\rho_c^{(n)} = \sqrt{(x^{(n)} - x)^2 + (y^{(n)} - y)^2 + (z^{(n)} - z)^2} + b + \tilde{\epsilon}^{(n)}; \quad n = 1, 2, \dots, N \quad (1-4)$$

where the position of the n th satellite, $(x^{(n)}, y^{(n)}, z^{(n)})$, is computed based on the navigation message, and the user position, (x, y, z) , is to be determined. b is the unknown user clock bias and $\tilde{\epsilon}^{(n)}$ are the errors that remain after applying the LAAS correction for the measurement errors discussed in Section 1.1.3. We solve the N equations by linearizing them about initial estimates of the user position and the clock bias: x_0, y_0, z_0 and b_0 . Let the expected ranging value based on these initial guesses be:

$$\rho_0^{(n)} = \sqrt{(x^{(n)} - x_0)^2 + (y^{(n)} - y_0)^2 + (z^{(n)} - z_0)^2} + b_0 \quad (1-5)$$

We now develop the linearized equation in which $\delta x, \delta y, \delta z$ and δb are the unknowns to be solved.

$$\rho_c - \rho_0 \doteq \delta\rho = G \begin{bmatrix} x - x_0 \\ y - y_0 \\ z - z_0 \\ b - b_0 \end{bmatrix} + \tilde{\varepsilon} = G \begin{bmatrix} \delta x \\ \delta y \\ \delta z \\ \delta b \end{bmatrix} + \tilde{\varepsilon} \quad (1-6)$$

where $\delta\rho$ is an N dimensional vector containing the differentially corrected pseudorange measurement (ρ_c) minus the expected ranging value (ρ_0). G is the user-satellite geometry matrix consisting of N rows of line of sight vectors ($1^{(N)}$), augmented by a “1” for the clock.

$$G = \begin{bmatrix} \left(-1^{(1)}\right)^T & 1 \\ \left(-1^{(2)}\right)^T & 1 \\ \vdots & \vdots \\ \left(-1^{(N)}\right)^T & 1 \end{bmatrix} \quad (1-7)$$

The next step is to obtain the optimal solution by the least-squares method iteratively until the change in the estimates is sufficiently small. The weighted least-squares solution for the corrections to the estimates of states can be written as [14]

$$\begin{bmatrix} \delta\hat{x} \\ \delta\hat{y} \\ \delta\hat{z} \\ \delta\hat{b} \end{bmatrix} = \left(G^T W G\right)^{-1} G^T W \delta\rho \quad (1-8)$$

For simplicity, let us define the weighted least-squares projection matrix as

$$S \equiv \left(G^T W G\right)^{-1} G^T W \quad (1-9)$$

To account for unequal measurement quality, the measurement residuals are weighted with a covariance matrix, which characterizes the errors, $\tilde{\varepsilon}$, in the pseudorange measurements. This covariance matrix is called the least-squares weighting matrix and the inverse of it is

$$W^{-1} = \begin{bmatrix} \sigma_{PR,1}^2 & 0 & \cdots & 0 \\ 0 & \sigma_{PR,2}^2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \sigma_{PR,N}^2 \end{bmatrix} \quad (1-10)$$

The weighting matrix treats the measurement errors from different satellites as zero-mean, uncorrelated, and Gaussian distributed; otherwise such a characterization is very difficult in general. For each measurement we have an error model given by

$$\sigma_{PR,n}^2 = \sigma_{air,n}^2 + \sigma_{tropo,n}^2 + \sigma_{iono,n}^2 + \sigma_{pr_gnd,n}^2 \quad (1-11)$$

The airborne error, σ_{air} , is determined from the receiver noise estimate and the specified multipath model. The second and the third terms are introduced by the residual tropospheric and ionospheric errors respectively (see Appendix A for details). The ground error, σ_{pr_gnd} , includes the ground station receiver noise and multipath error. The LAAS message broadcasts this fourth term for each satellite [15]. The vertical position error is then characterized by its standard deviation,

$$\sigma_{VerticalPositionError} = \sqrt{\sum_{n=1}^N S_{vertical,n}^2 \sigma_{PR,n}^2} \quad (1-12)$$

where $S_{vertical,n}$ is the projection of the local vertical component for the n th ranging source. Assuming that vertical position errors are Gaussian distributed, the vertical protection levels (VPL) can be computed as

$$VPL = K_{ffmd} \sigma_{VerticalPositionError} \quad (1-13)$$

where K_{ffmd} is the quantile of a unit Gaussian distribution corresponding to γ . The computation of horizontal protection levels (HPL) is essentially the same except to project pseudorange errors onto the horizontal direction. Since the vertical direction requirement is the most stringent and errors in this direction are the largest, we will only focus on the VPL in this thesis. Note that the VPL in Equation (1-13) is computed under the hypothesis of

“fault-free” conditions (H_0). The VPLs computed under different operational hypotheses—single reference-receiver failure or single satellite ephemeris fault [14] – at an aircraft are out of scope for this work.

1.5 REAL TIME ERROR BOUNDING

LAAS avionics use the computed PL in real time to determine whether an operation is safe. As stated earlier, the protection level needs to be smaller than the required alert limit (AL) in order for the operation to be conducted. For this reason, the computed PL must be credible. If the PL fails to bound the true position error, then the pilot may attempt a flight that is not safe. However, we encounter some technical challenges in the error bounding. In Section 1.5.1, three problems concerning the error bounding are explained. The previous work on error bounding is then discussed in Section 1.5.2.

1.5.1 CHALLENGES IN ERROR BOUNDING

The first problem is that the error model of differentially corrected pseudorange measurements may not be accurate enough to be used for error bounding. As described in the previous section, the algorithms for the generation of PL assume a zero-mean and normally distributed error model for corrected measurements. Yet the errors are neither necessarily zero-mean nor Gaussian. Since an accurate characterization of the correlation across errors is very difficult, we assume the errors for each satellite are uncorrelated. However, such an assumption may be unjustified. The standard deviation of correction error, “sigma,” is further assumed to be equal to the broadcast value determined with the error model. Because the computation of PL is based on these broadcast values of standard deviations, as shown in Equations (1-11), (1-12) and (1-13), special care must be taken with these assumptions. If the error model does not overbound the true error distribution, it may cause a serious threat to the aircraft. The approach taken in LAAS to ensure the error model overbounds the true distribution is, for each satellite, to transmit an inflated value of the standard deviations. Previous research on this subject will be discussed in the next section, and a new approach to determine how much we should inflate the broadcast sigma will be presented in Chapter 4.

The second problem arises in abnormal situations. Let us suppose that the broadcast sigma is magnified enough so that the true error distribution is bounded by a zero-mean Gaussian distribution defined with that inflated sigma. This technique enables the aircraft to compute the PL that can ensure an acceptable level of risk. However, this may not be true in all conditions. There may be unexpected anomalies that cause the true sigma to exceed the broadcast sigma. The source of such anomalies in corrected pseudoranges can be: multipath error increases when environmental conditions vary, receiver noise error amplifications due to a receiver or antenna failure, or all other possible malfunctions. In order to provide reliable PL, LAAS is required to detect these abnormalities and send an alarm to users within the required time-to-alert. This task can be accomplished by monitoring the standard deviation of pseudorange correction error in real time. This “sigma monitoring” is the subject of Chapter 3.

The LAAS sigma-overbounding issue is especially difficult for Category II/III operations (which will be further discussed in Chapter 2), because the Alert Limits (ALs) are small. In other words, because of the tightened Vertical Alert Limit (VAL), the VPLs cannot be overly conservative, and high levels of sigma inflation cannot be tolerated. Otherwise, the full capacity of the system will not be utilized. Consequently, the goal is to make the sigma inflation as small as possible while maintaining the reliability of VPLs, which are driven by the broadcast sigmas. In order to accomplish this goal, the Position Domain Monitor (PDM) concept has been proposed. Previous work on PDM will be discussed in the next section. In Chapter 5, we will describe a new use of PDM to support smaller inflation of the broadcast sigma.

1.5.2 PREVIOUS WORK

The first group of previous studies deals with sigma establishment and estimation. The goal of this research is to characterize the broadcast standard deviation well enough to be used for the computation of protection levels. A detailed approach to establishing the broadcast sigma was suggested by Pervan and Sayim [16]. This suggested method computes the root-sum-square (RSS) of the standard deviation due to receiver noise, diffuse multipath, and ground reflection multipath. The implicitly assumed zero-mean

Gaussian error model appears to be a consistent one for receiver noise and diffuse multipath errors. Thus, the standard deviations of those receiver noise and diffuse multipath error models can be obtained from experimental data for a specified LAAS installation. In contrast, ground reflection multipath can slowly vary with environmental conditions and, consequently, it is impractical to characterize the underlying distribution with experimental data alone. The sigma for ground reflection multipath is established by first obtaining a theoretical model and then validating the model with empirical data. This approach has been applied to define standardized standard deviations based on currently available GPS receiver/antenna technology by McGraw et al. [17].

The second group of prior work addresses statistical uncertainty in the definition of the broadcast sigma. The estimated sigma from the error model may not be a good representative of the true sigma, because of the finite sample sizes used to generate the model and error correlation across multiple reference receivers. The sigma needs to be inflated to account for this statistical uncertainty. In this regard, the minimum acceptable inflation parameters for the value of the broadcast sigma have been derived by Pervan and Sayim in [1]. This work implicitly assumed the zero-mean Gaussian error model associated with thermal noise and diffuse multipath. However, the authors acknowledged that other error sources, such as ground reflection multipath and systematic reference receiver/antenna errors, may not be zero-mean Gaussian distributed. Because a user computes PLs assuming a zero-mean Gaussian distributed error model based on the broadcast sigma, an overbound concept needs to be applied if the error distribution is not Gaussian. Shively and Braff [18] derived inflation factors to deal with this non-Gaussian effect using a synthetic model of a Gaussian core and Laplacian tails. Despite all this progress, the study of sigma overbounding and inflation is not complete. To establish a reliable model of the error distribution, more strong physical bases are needed. The resulting inflation factors from the model of Laplacian tails may significantly be larger than would be needed if the tails were known to be nearly Gaussian. In contrast, even with tens of thousands of samples, the resulting inflation factors may provide a limited confidence without a dependable overbounding method. Rife [19] introduced a modified overbounding technique, called core overbounding. His Gaussian Core with Gaussian Sidelobe (GCGS)

approach mitigates over-conservatism associated with bounding heavy tails by providing an allowable envelope of tail distributions. Many researchers actively work on these issues.

Another concern of sigma overbounding is that the same conservatism cannot be adopted for Category II/III approaches. The concern led to the application of a position domain method, which was originally introduced by Markin and Shively in [20] as an alternative to a range domain method. The computed protection levels in the range domain may be conservative, since the range domain method requires a transformation from pseudorange correction errors to position error estimates. In contrast, this alternative technique avoids the conservatism by performing a safety check directly in the position domain. An extended benefit from the position domain method has been shown by Braff [21]. In this work, the method was found to be effective to reduce any inflation factor that was applied to protect against the event that the pseudorange correction error distribution was not modeled properly.

Lastly, the previous work on the Cumulative Sum (CUSUM) method is essential, since this method is directly applied to the sigma monitoring in real time. The CUSUM method was originally invented by Page in 1954 [22]. The principles and applications of this method were analyzed in depth by Hawkins and Olwell [23]. Pullen first considered adapting the CUSUM method to LAAS and used the method to validate protection level overbounds for ground-based and space-based augmentation systems [24].

1.6 OUTLINE AND CONTRIBUTIONS

Since major contributions will be described thoroughly in the thesis, only a brief summary of those efforts is given here with an outline of this thesis. The second chapter gives a more detailed picture of the LAAS architecture and of the requirements for aviation navigation. Chapter 2 also describes the Stanford University Integrity Monitor Testbed (IMT), which is a prototype of the LAAS ground facility. It is not a contribution in this thesis, but it needs to be explained since it is the basis upon which the new monitors are implemented.

The third chapter explains how to estimate and monitor standard deviations of differentially corrected pseudorange errors in real time. This chapter also shows the responses of sigma

monitors under the failure modes. In order to detect and remove abnormal behaviors of the pseudorange correction error distribution, two different sigma-monitoring algorithms were developed: the sigma estimation method and Cumulative Sum (CUSUM) method. The former detects relatively smaller violations faster, while the latter detects larger violations faster. This is one of the most valuable contributions of this thesis, as the two sigma-monitoring algorithms together are able to detect any size of sigma violation that is hazardous to users. At the end of Chapter 3 we include an analysis of mean monitors. The “head-start” CUSUMs are superior to the mean estimation method and sufficient to detect mean anomalies during LAAS operations.

The fourth chapter creates the new inflation factor determination method and analytically derives the inflation factor for the broadcast sigma based on both experimental and theoretical data. To do that, the effect of sigma monitor performance on the determination of the inflation factor was evaluated.

The fifth chapter theoretically demonstrates the advantages of position-domain monitors and shows how to improve the overall system performance based on empirical tests. The last important contribution resides in the fact that the position-domain monitoring algorithm added to the current algorithm can support a reduced inflation factor. Used in combination, a 25% reduction in VPL is achieved with the same safety standard.

Finally, in Chapter 6, the accomplished work is summarized, and directions for future research are suggested.

Chapter 2

The Local Area Augmentation System

2.1 INTRODUCTION

GPS is already used for many different types of aircraft navigation: cruising through oceanic routes, flight over continents, etc. In time, GPS will also be used for the final approach to airports and auto-landings. Among these operations, precision approach and landing navigation demand the greatest safety and reliability. To serve these applications, the Federal Aviation Administration (FAA) has been developing a ground-based system to augment GPS. This augmented system is known as the Local Area Augmentation System (LAAS) [25] because it locates a reference station on the ground at an airport to improve the performance of airborne GPS receivers approaching that airport (over approximately a 20-30 mile radius). This chapter starts by giving a brief description of the LAAS architecture. The completed augmentation system will meet stringent requirements with respect to accuracy, integrity, continuity, and availability. Chapter 2 introduces these technical terms which will be used to describe the performance of LAAS. It also defines categories of LAAS precision approach and landing based on the level of these requirements. The chapter then reviews how the GPS measurements are processed to generate differential corrections, and how their residual errors are characterized in LAAS. Finally, its attention turns to the Stanford Integrity Monitor Testbed (IMT), which has been developed to evaluate whether LAAS can meet the defined integrity requirements.

2.2 LAAS ARCHITECTURE OVERVIEW

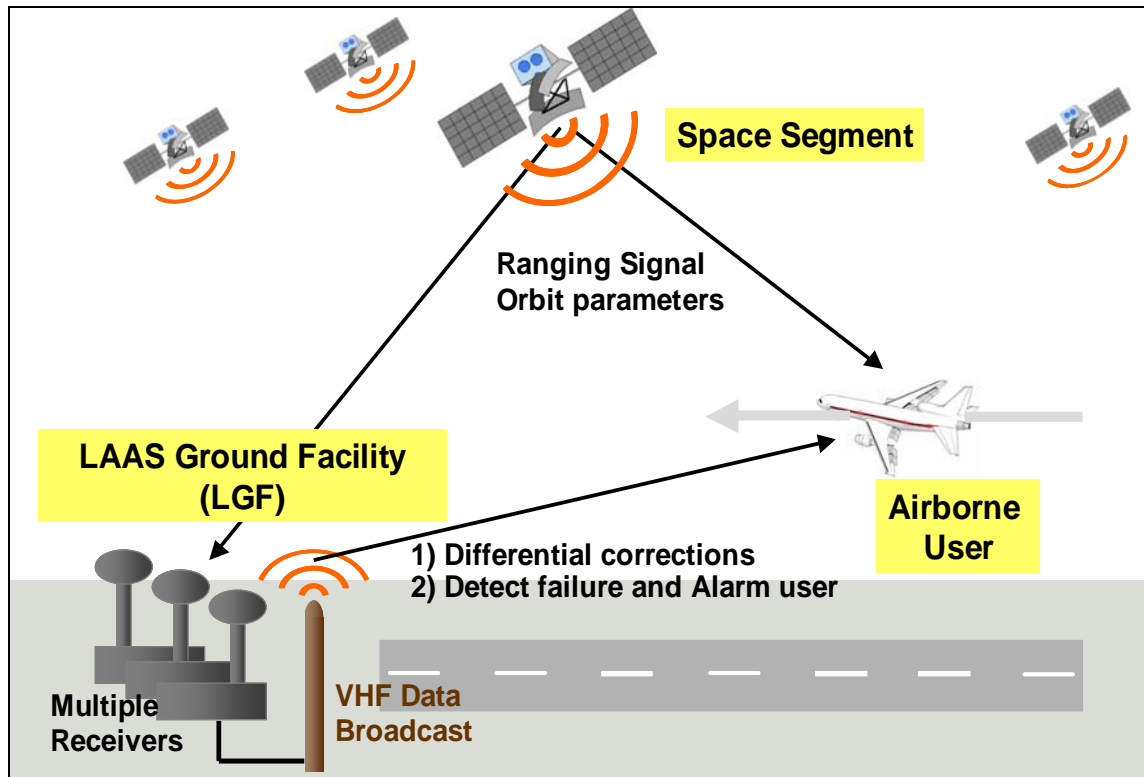


Figure 2.1: Local Area Augmentation System (LAAS) Architecture

As shown in Figure 2.1, LAAS consists of three segments: the space segment, the LAAS Ground Facility (LGF), and the airborne user segment. The space segment provides ranging signals and orbit parameters to the LGF and to users. The LGF includes a small collection (typically 3 or 4) of GPS reference receivers and antennas placed at precisely known locations. With this set of receivers, the LGF continuously tracks, decodes, and monitors GPS signals and generates differential corrections. To help insure flight safety, it is also responsible for detecting both space segment and ground segment failures and rapidly warning users. The corrections, along with integrity parameters and approach-path information, are broadcast to approaching aircraft via a Very High Frequency (VHF) data link. The airborne GPS receivers use this information to correct their own ranging measurements, obtain the required accuracy, and verify the required integrity.

2.3 LAAS REQUIREMENTS

In order to specify requirements for precision approaches using LAAS, we must first introduce the terminology that is used to describe all aircraft navigation applications. The four criteria to evaluate the performance of air navigation systems are defined as follows [25, 26] and illustrated in Figure 2.2.

- *Accuracy*: A measure of the difference between the estimated position and the true aircraft position under nominal fault-free conditions. It is typically a 95% bound on *navigation sensor error* (NSE).
- *Integrity*: The ability of a navigation system to detect anomalies and provide warnings to users in a timely fashion.
- *Continuity*: The probability that the system supports Accuracy and Integrity requirements throughout a flight operation without interruption.
- *Availability*: The percentage of time for which the system is operational and the Accuracy, Integrity and Continuity requirements are met.

Figure 2.2 illustrates this terminology in a two-dimensional plane for simplification. The origin indicates the true position, and a dot specifies the estimated position. Integrity fails when the position error (the dot) exceeds a certain AL (alert limit) (the outer circle) and this event is not notified to a pilot within a specified time-to-alarm. Thus *integrity risk* is defined as the probability that no alert is issued while the position error exceeds the AL for a time longer than the time-to-alarm. However, the true position, and consequently, the position error are not quantities that can be known in real-time. For that reason, protection bounds (the inner circle in Figure 2.2), defined as VPL and HPL in Section 1.4, need to be computed with respect to the acceptable level of integrity risk.

In contrast to integrity risk, continuity risk is defined as how often the system fails during a specified time interval. Continuity and integrity are competing requirements. If integrity algorithms are overly sensitive, too many “false alarms” will be sent to the user, and the

system continuity will decrease. Lastly, if the protection bound (the inner circle) exceeds the AL (the outer circle), the system is no longer available. If this occurs before an approach, there will be only availability loss. On the other hand, if the system fails during the approach, continuity risk will increase along with availability loss.

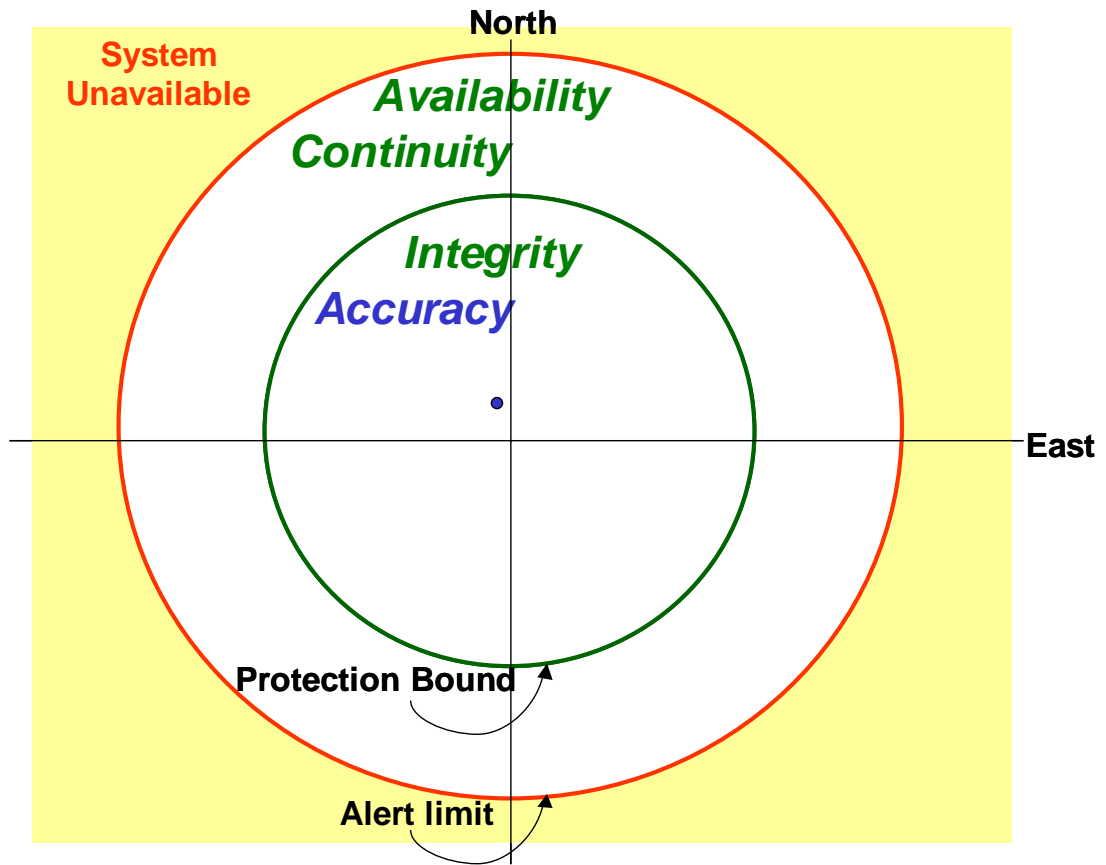


Figure 2.2: Aviation Navigation Requirements

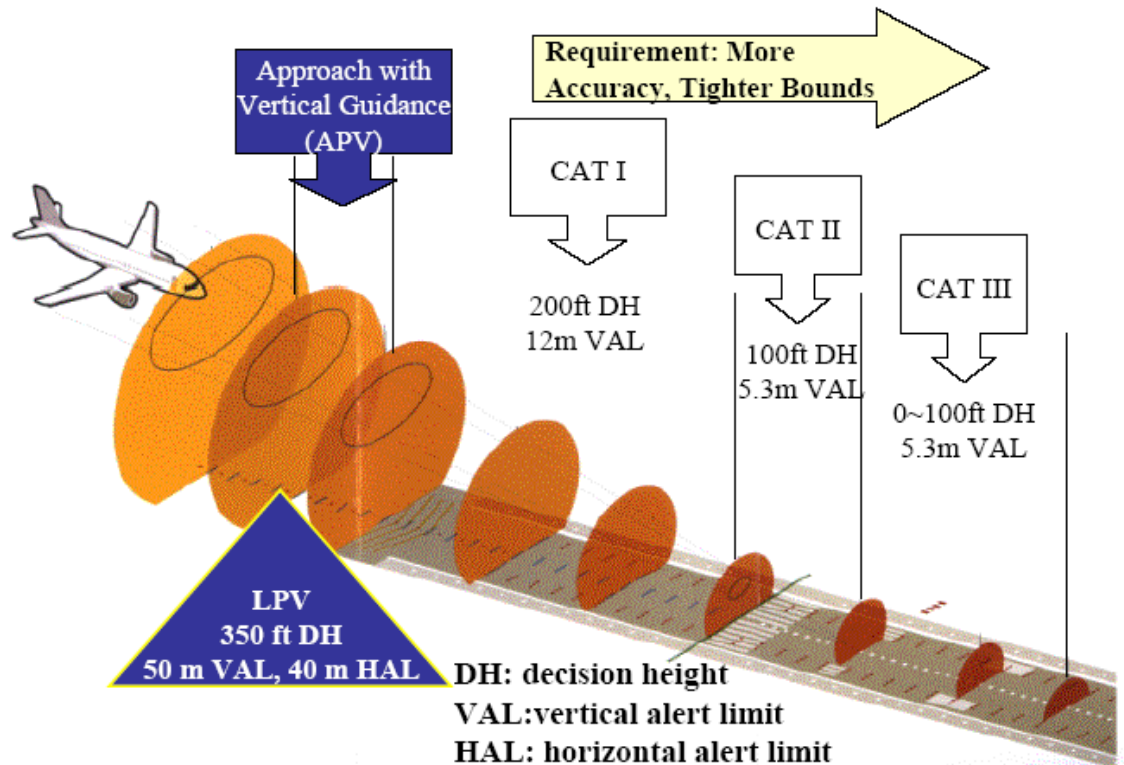


Figure 2.3: Precision Approach and Landing Categories [27, 28]

We now turn our attention to categorizing precision approaches based on their capability to provide various levels of accuracy, integrity, continuity, and availability. As shown in Figure 2.3 [27, 28], Lateral Precision Approach with vertical guidance (LPV) and Approach with Vertical guidance (APV-2) guide an aircraft to a minimum altitude known as the Decision Height (DH) after which the aircraft can proceed only if the runway is visible. LAAS Category I (CAT I), CAT II and CAT III Precision Approaches (PA) are the subjects of interest in this thesis. CAT II and CAT III involve more stringent requirements that allow users to operate at lower DHs. The current requirements and VAL/HAL – a bound on maximum tolerable VPL/HPL – for these precision approaches are summarized in Table 2-1 [15, 25]. As an example, if Hazardously Misleading Information (HMI) causes a CAT I user’s vertical position error to exceed 10 meters, the LGF must detect the event and alert the user within a 6-second time-to-alarm. The probability of the LGF failing in this task, $\text{Pr}(\text{HMI})$, should be no greater than 2×10^{-7} per approach.

Phase of Flight	Accuracy (95% error)	Integrity		Alert Limit (H: Horizontal V: Vertical)	Continuity	Availability
		Time to Alert	Pr(HMI)			
LPV (APV1.5)	H: 16 m V: 20 m	10 sec	2×10^{-7} / approach	H: 40 m V: 50 m	5.5×10^{-5} / approach	0.99 to 0.99999
APV-2	H: 16 m V: 7.6 m	6 sec	2×10^{-7} / approach	H: 40 m V: 20 m	5.5×10^{-5} / approach	0.99 to 0.99999
CAT I	H: 16 m V: 4 to 7.6 m	6 sec	2×10^{-7} / approach	H: 40 m V: 10 to 12 m	5.5×10^{-5} / approach	0.99 to 0.99999
CAT II	H: 6.9 m V: 2.0 m	2 sec	2×10^{-9} / approach	H: 17.4 m V: 5.3 m	4×10^{-6} / 15 sec	0.99 to 0.99999
CAT III	H: 6.1 m V: 2.0 m	1 to 2 sec	2×10^{-9} / approach	H: 15.5 m V: 5.3 m	H: 2×10^{-6} / 30 sec V: 2×10^{-6} / 15 sec	0.99 to 0.99999

Table 2-1: Requirements for Precision Approach and Landing

Although the FAA originally proposed a VAL of 5.3 meters for CAT II/III, the appropriate value is being reconsidered by the Radio Technical Commission for Aeronautics (RTCA). Tim Murphy of Boeing developed an alternative methodology to determine the alert limits and showed that the probability of unsuccessful landing is still on the order of 10^{-7} or less for alert limits up to 10 meters [29].

2.4 LAAS GROUND FACILITY PROCESSING ALGORITHMS

The LGF is responsible for generating and broadcasting carrier-smoothed code differential corrections to users. The processing algorithms, including carrier smoothing and the computation of differential corrections, will be explained in Section 2.4.1. As addressed earlier, the broadcast also carries integrity data which estimates the quality of the correction. We will define a ground facility error standard deviation, σ_{pr_gnd} , and present the error model in Section 2.4.2. The last important role of the LGF is to insure that all ranging sources, for which LAAS corrections are broadcast, are safe to use by detecting faulty measurements. Since fault detection involves highly complex mechanisms and requires a lengthy explanation, LGF monitoring algorithms will be presented separately in Section 2.5.

2.4.1 CARRIER SMOOTHING AND PSEUDORANGE CORRECTIONS

The first step to compute a differential correction is the smoothing of pseudorange measurements, ρ , with the change in carrier phase measurements, ϕ . Carrier smoothing reduces rapidly changing errors in raw pseudorange measurements such as high frequency errors due to receiver noise. For each channel, a pair consisting of Receiver m - Satellite n at epoch k , the following filter is applied with a time constant, τ_s , of 100 seconds. The sample interval, T_s , is 0.5 seconds; thus N_s is equal to 200 [15, 30].

$$\rho_{s,m,n}(k) = \frac{1}{N_s} \rho_{m,n}(k) + \frac{N_s - 1}{N_s} (\rho_{s,m,n}(k-1) + \phi_{m,n}(k) - \phi_{m,n}(k-1)) \quad (2-1)$$

where

$$N_s = \tau_s / T_s = 100 / 0.5 = 200 \quad (2-2)$$

We then compute corrections for both pseudorange and carrier phase measurements [31]. The smoothed pseudorange correction, ρ_{sc} , and the raw carrier phase correction, ϕ_c , for each channel (m, n) are:

$$\rho_{sc,m,n}(k) = \rho_{s,m,n}(k) - R_{m,n}(k) + \tau_{m,n}(k) \quad (2-3)$$

$$\phi_{c,m,n}(k) = \phi_{m,n}(k) - R_{m,n}(k) + \tau_{m,n}(k) - \phi_{ci,m,n}(0) \quad (2-4)$$

where $R_{m,n}$ is the range from reference antenna to satellite (computed based on the known location of the reference antenna and the broadcast ephemeris), and $\tau_{m,n}$ is the satellite clock correction. The initial carrier phase correction, $\phi_{ci,m,n}(0)$, is equal to $\phi_{m,n}(0) - R_{m,n}(0) + \tau_{m,n}(0)$. After generating the individual corrections for each channel using Equations (2-3) and (2-4), we adjust receiver clock biases as follows to allow measurements to be compared across receivers [31]:

$$\rho_{sca,m,n}(k) = \rho_{sc,m,n}(k) - \frac{1}{N_c(k)} \sum_{j \in S_c(k)} \rho_{sc,m,j}(k) \quad (2-5)$$

$$\phi_{ca,m,n}(k) = \phi_{c,m,n}(k) - \frac{1}{N_c(k)} \sum_{j \in S_c(k)} \phi_{c,m,j}(k) \quad (2-6)$$

Here, S_c designates a set of N_c ranging sources in the maximum common set, which will be detailed in Section 2.5.2.4. Lastly, we compute the averaged corrections for each satellite using the following equations.

$$\rho_{corr,n}(k) = \frac{1}{M_n(k)} \sum_{i \in S_n(k)} \rho_{sca,i,n}(k) \quad (2-7)$$

$$\phi_{corr,n}(k) = \frac{1}{M_n(k)} \sum_{i \in S_n(k)} (\phi_{ca,i,n}(k) - \phi_{ca,i,n}(0)) \quad (2-8)$$

In these equations, S_n is a set of M_n reference receivers, and $\phi_{ca,m,n}(0)$ is evaluated at the first measurement epoch for channel (m, n) . This computed ρ_{corr} is broadcast to all LAAS users. The carrier-phase corrections, ϕ_{corr} , are not needed for CAT I but may be needed for CAT II/III users on account of their tightened VAL of 5.3 meters.

2.4.2 GROUND FACILITY ERROR STANDARD DEVIATION

In the previous chapter, we saw how user aircraft compute their protection levels using the standard deviation of differentially corrected pseudorange measurements. This section reminds the reader of the different components of the standard deviation.

$$\sigma_{PR,n}(k) = \sqrt{\sigma_{air,n}^2(k) + \sigma_{tropo,n}^2(k) + \sigma_{iono,n}^2(k) + \sigma_{pr_gnd,n}^2(k)} \quad (2-9)$$

The first three terms are the airborne RMS pseudorange error, the residual tropospheric error and the residual ionospheric error, respectively. The error models corresponding to these terms are specified in Appendix A. Let us now turn our attention to the fourth term, the ground facility pseudorange error, which is broadcast to users by the LGF. When combined with the differential correction methodology, this broadcast σ_{pr_gnd} is a critical factor for airborne users to compute their position and integrity protection levels (VPL and HPL) as shown in Section 1.4. This error standard deviation for each ranging source should account for all equipment and environmental effects, including receiver noise, local interference, and ground station multipath.

Working Group-4 of RTCA Special Committee 159 developed standard error models for LAAS differential processing [15]. The standard GPS interference environmental conditions – the RF interference environment, at and around L-band frequencies, for LAAS airborne receivers – assumed in the models are defined in Appendix D of the LAAS Minimum Operational Performance Standards (RTCA/DO-253A [32]). The group defined Ground Accuracy Designators (GADs) that reflect different performance levels of GPS receiver technologies [17]. GAD-A represents a level of performance achievable with early and low-cost LAAS installations using a standard correlator receiver and a single-aperture antenna. GAD-C was defined to characterize the performance realizable with a narrow correlator receiver and a multipath limiting antenna (MLA). GAD-C performance is expected to be able to support LAAS CAT II/III precision approaches. GAD-B represents an intermediate level of performance between GAD-A and GAD-C. The performance of GAD-B is attainable with advanced receiver technologies similar to GAD-C but with a single-aperture antenna instead of an MLA.

The standard deviation of the ground facility error is [15]:

$$\sigma_{pr_gnd,n}(k) = \sqrt{\frac{\left(a_0 + a_1 e^{-\theta_n(k)/\theta_0}\right)^2}{M_n(k)} + (a_2)^2} \quad (2-10)$$

where M is the number of reference receivers that are averaged to obtain a differential correction, θ_n is the n th ranging source elevation angle, and a_0, a_1, a_2 and θ_0 for the applicable Ground Accuracy Designators (GAD) are defined in Table 2-2.

Ground Accuracy Designator		a_0 meters	a_1 meters	a_2 meters	θ_0 degrees
GAD-A		0.50	1.65	0.08	14.3
GAD-B		0.16	1.07	0.08	15.5
GAD-C	$\theta_n \geq 35^\circ$	0.15	0.84	0.04	15.5
	$\theta_n < 35^\circ$	0.24	0	0.04	-

Table 2-2: Ground Facility Error Allocation Model

2.5 STANFORD LAAS INTEGRITY MONITOR TEST-BED (IMT)

As addressed earlier in this chapter, the LGF guarantees users that it is safe to use all ranging sources for which LAAS corrections are broadcast. If a failure threatening user safety occurs, the LGF must detect and alert users – by not broadcasting the correction for the affected measurement – within a 3-second time-to-alarm (out of the 6-second requirement) for CAT I precision approaches. Furthermore, the probability of failing this task should be less than 2×10^{-7} per approach as shown in Section 2.3. In order to evaluate whether the LGF can meet these integrity requirements, Stanford University researchers have developed an LGF prototype known as the Integrity Monitor Testbed (IMT) [30, 33-37].

Section 2.5.1 will first give an overall picture of the IMT system architecture. Section 2.5.2 will then briefly present the integrity monitor algorithms and failure-handling logic implemented in IMT. For a full description of the IMT, please refer to [38].

2.5.1 IMT HARDWARE CONFIGURATION

The LGF requires redundant DGPS reference receivers to be able to detect and exclude failures of individual receivers. Figure 2.4 shows the configuration of the three IMT antennas on the Stanford HEPL (Hansen Experimental Physics Laboratory) rooftop. The existing IMT antennas are connected to three NovAtel OEM-4 reference receivers which are connected to the IMT computer by a multiport serial board. The separations between these three NovAtel Pinwheel (survey grade) antennas are limited to 20 – 65 meters by the size of the HEPL rooftop but are sufficiently separated to minimize the correlations between individual reference receiver multipath errors (this has been demonstrated by previous work) [33, 39].

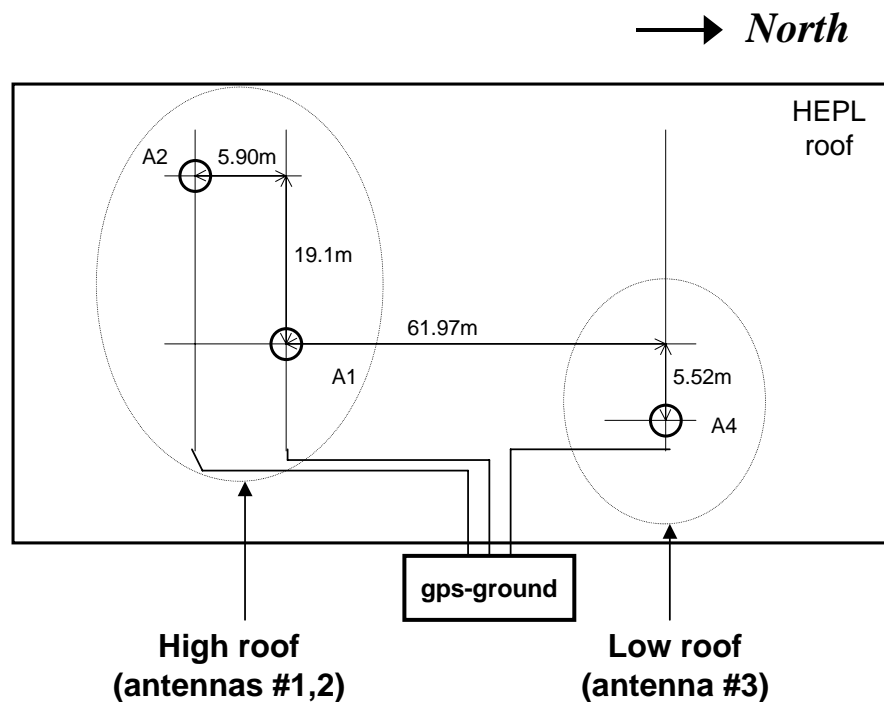


Figure 2.4: IMT Hardware Configuration

Each receiver can track as many as 12 satellites simultaneously. Each receiver samples GPS signals every 0.5 seconds and provides receiver measurement packets, which contain pseudorange measurements, carrier-phase measurements, and navigation messages. These GPS measurements are fed into the IMT processor for further calculations.

2.5.2 IMT FUNCTIONS

The LGF must apply a comprehensive set of monitoring algorithms to detect a varied array of possible failures in the GPS Signal in Space (SIS) or in the LGF itself. In order to coordinate the LGF response to detected failures (some of which may trigger more than one monitoring algorithm), complex fault-handling logic is included in the LGF. This logic is called Executive Monitoring (EXM), and it isolates failed measurements and reintroduces these measurements only after the failure is clearly determined to have been corrected.

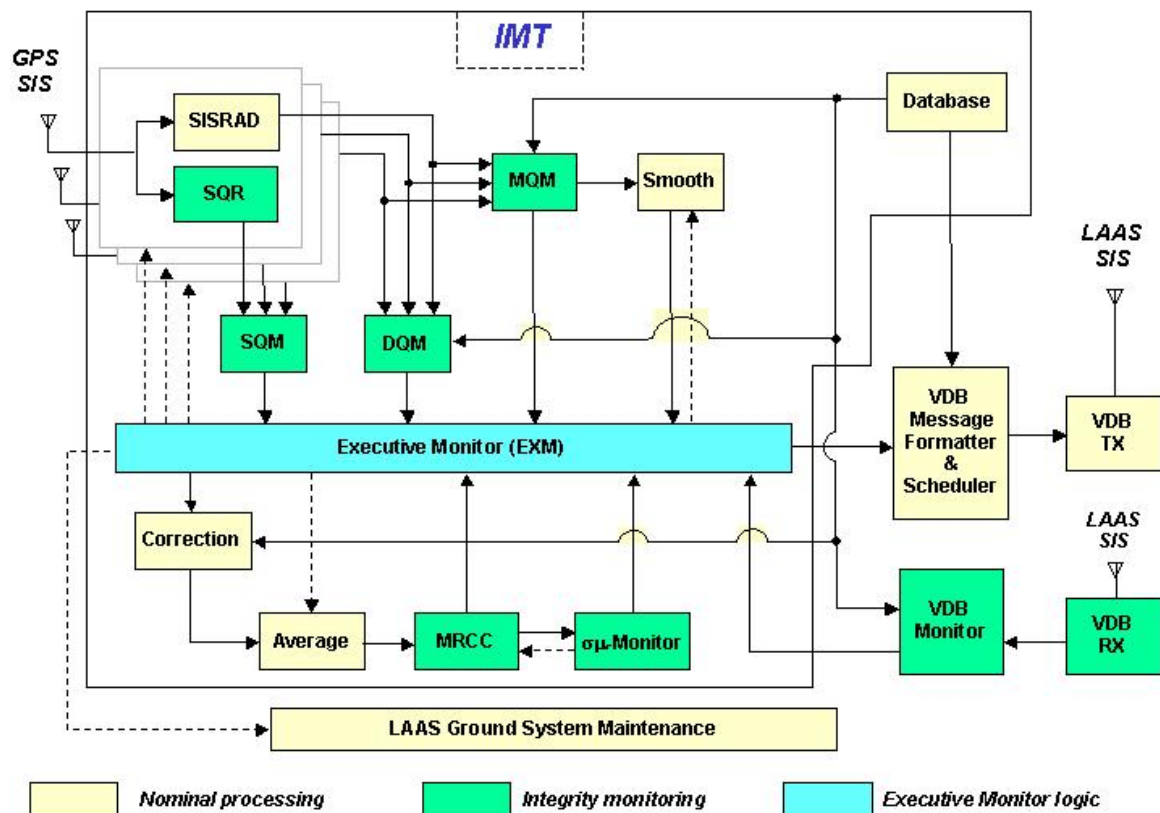


Figure 2.5: IMT Hardware Configuration

Figure 2.5 shows the IMT functional blocks. The preliminary functions are RF conditioning and Signal-in-Space Receive and Decode (SISRAD) functions. The role of these functions is to provide pseudorange measurements, carrier-phase measurements, and navigation data messages in order to enable the generation of differential corrections. The core of IMT processing consists of three parts: nominal processing (carrier smoothing and

calculation of differential corrections), integrity monitor algorithms, and executive monitoring (EXM). We have already discussed nominal processing in Section 2.4.1. We can divide the integrity monitor algorithms into Signal Quality Monitoring (SQM), Data Quality Monitoring (DQM), Measurement Quality Monitoring (MQM), Multiple Reference Consistency Check (MRCC), Sigma-Mean ($\sigma\mu$) monitoring, and Message Field Range Test (MFRT) classes. In addition to the IMT processing core, the LGF also contains VHF Data Broadcast (VDB). It is an essential component through which the LAAS corrections and integrity data will be broadcast to users.

2.5.2.1 SIGNAL QUALITY MONITORING (SQM)

SQM identifies deformations of the C/A-code broadcast by GPS satellites [40-42]. The IMT SQM component comprises distinct Signal Quality Receivers (SQR) which report C/A-code correlation measurements at several different correlator spacings. These measurements are processed to determine whether signal-deformations have altered the ideal triangular C/A-code correlation shape significantly. The details of its algorithms and performance are demonstrated in [43, 44].

SQM also confirms that received satellite signal power is within SPS specifications [45] by averaging the reported receiver carrier-to-noise power density ratio, C/N_0 . If the signal power is significantly lower than the specification, ranging errors may increase and present an integrity risk. The averaged C/N_0 for each channel (m, n) at the current epoch k and the one at the previous epoch ($k-1$) is:

$$C / N_{0_Avg,m,n}(k) = \frac{1}{2} \left(C / N_{0,m,n}(k-1) + C / N_{0,m,n}(k) \right) \quad (2-11)$$

This value is compared with a threshold value that is predetermined based on the hardware configuration, antenna sites, antenna gain patterns, and cable losses. The general method to determine thresholds utilized in the IMT is proposed in [38]. If the averaged C/N_0 is less than the threshold, a flag is issued for this channel.

Separate SQM algorithms are used to confirm that the received signals do not include an anomalous amount of code-carrier divergence, which would interfere with carrier smoothing and could lead to larger errors. A geometric moving averaging (GMA) method and a divergence cumulative sum (CUSUM) method are applied to estimate the code-carrier divergence (see [38, 46] for details of the code-carrier divergence monitor).

2.5.2.2 DATA QUALITY MONITORING (DQM)

The role of DQM is to verify that the received satellite navigation data are sufficiently reliable. Several DQM methods have been developed to validate the GPS ephemeris and clock data of each satellite for two different situations: when a satellite first rises into view of the LGF and when navigation data messages are updated [47, 48].

For a newly-risen satellite, DQM compares the satellite positions over the next six hours at 5-minute intervals based on the broadcast ephemeris message to those generated from the most recent almanac message. Ephemeris parameters are updated by a satellite every two hours, while the almanac is a subset of the clock and ephemeris data of all satellites and is updated much less frequently with reduced precision [3]. DQM then insures that the ephemeris-based satellite positions always agree with the almanac-based positions to within 7000 meters, where this threshold is set by the accuracy of the almanac [38]. When navigation messages are updated, DQM computes satellite positions based on both old and new ephemerides and compares them to insure that the new ephemeris is consistent with the old and validated ephemeris to within 250 meters.

DQM works in cooperation with Message Field Range Test (MFRT) functions (see Section 2.5.2.7). Although most large errors will be detected by the MFRT test, ephemeris errors orthogonal to the line-of-sight between a failed satellite and an LGF are not detectable by this test. To account for this possibility, a monitor concept has been developed that is known as the “Yesterday Ephemeris-minus-Today Ephemeris (YE-TE) test.” This test is a more precise validation of ephemerides for newly risen satellites than the ephemeris-minus-almanac test. The concept of the YE-TE test is to confirm that today’s broadcast ephemeris

data for each GPS satellite is correct by comparison with the most recently validated ephemeris data. The detailed algorithms and performance of this test are presented in [48].

2.5.2.3 MEASUREMENT QUALITY MONITORING (MQM)

MQM is designed to detect sudden step errors and any other rapidly changing errors due to GPS clock anomalies or LGF receiver failures by verifying the consistency of the pseudorange and carrier-phase measurements over the last few epochs [30, 31]. It consists of three monitors: the Receiver Lock Time check, the Carrier Acceleration-Step test, and the Carrier-Smoothed Code (CSC) Innovation test. MQM generates one flag per channel if it detects any failure of the channel from these three tests.

First, the Receiver Lock Time Check ensures continuous receiver phase lock by computing the numerical difference of the lock times reported by each reference receiver. Since this type of failure is most likely not hazardous and frequently occurs when a receiver tracks satellites at low elevation angles, this test is implemented in such a way that EXM will not declare the channel to be failed. Instead, the IMT initializes and resets several memory buffers properly before using them, including re-initializing its carrier smoothing filter discussed in Section 2.4.1. This prevents the system from a loss of continuity.

Second, the purpose of the carrier acceleration-step test is to detect an impulse, step, excessive acceleration, or other rapid changes in carrier phase measurements. Such anomalies could cause errors in the pseudorange corrections or in the carrier-phase corrections. The last 10 continuous epochs (i.e., from epoch $k-9$ to epoch k) of ϕ^* are calculated for each channel (m, n) at each epoch k by this equation:

$$\phi_{m,n}^*(k) = \phi_{c,m,n}(k) - \frac{1}{N_m} \sum_{j \in S_m(k)} \phi_{c,m,j}(k) \quad (2-12)$$

where $\phi_{c,m,n}$ is computed by Equation (2-4), and S_m is the set of N_m satellites tracked on receiver m . These ten fitting points are used to fit the following quadratic model:

$$\phi_{m,n}^*(k,t) = \phi_{0,m,n}^*(k) + \frac{d\phi_{m,n}^*(k,t)}{dt}t + \frac{d^2\phi_{m,n}^*(k,t)}{dt^2}\frac{t^2}{2} \quad (2-13)$$

The least-squares method is used to solve for the coefficients of this model [15]. The acceleration and ramp are defined as:

$$Acceleration_{m,n}(k) \equiv \frac{d^2\phi_{m,n}^*(k,t)}{dt^2} \quad (2-14)$$

$$Ramp_{m,n}(k) \equiv \frac{d\phi_{m,n}^*(k,t)}{dt} \quad (2-15)$$

The third test statistic expresses the apparent change (or “step”) in the latest measurement and is defined as [31]:

$$Step_{m,n}(k) \equiv \phi_{meas,m,n}^*(k) - \phi_{pred,m,n}^*(k) \quad (2-16)$$

where $\phi_{meas,m,n}^*$ is the actual value at the current epoch k from Equation (2-14), and $\phi_{pred,m,n}^*$ is the predicted value from Equation (2-13). If any of the acceleration, ramp, and step test statistics on a channel exceed their thresholds, the channel fails and is flagged for later exclusion by EXM-I (see Section 2.5.2.4).

Lastly, the Carrier-Smoothed Code (CSC) Innovation Test is designed to detect impulse and step errors on raw pseudorange measurements. The innovation test statistic is defined as [31]:

$$Inno_{m,n}(k) \equiv \rho_{m,n}(k) - (\rho_{s,m,n}(k-1) + \phi_{m,n}(k) - \phi_{m,n}(k-1)) \quad (2-17)$$

where $\rho_{s,m,n}$ is the output of the carrier smoothing computed by Equation (2-1). If two or all of innovations at three successive epochs exceed thresholds, a flag will be issued from this CSC innovation test, and the channel will be excluded by EXM-I. If only the innovation at the current epoch is over its threshold, the smoothing filter does not use the raw pseudorange for that epoch. Instead, the carrier-smoothed code is updated based only

on the carrier phase measurement (if the carrier measurement fails the acceleration-step test, EXM-I will treat this as a failed channel).

2.5.2.4 PHASE ONE OF EXECUTIVE MONITORING (EXM-I)

In the previous sections, we described each of the LGF quality monitoring (QM) algorithms. Each integrity monitor is targeted to detect certain failures and may generate one flag per receiver-satellite pair. Once these monitors begin flagging questionable measurements, several steps of logical reasoning and trial removals are required to determine which failed system elements are the source of the problem. This LGF function is known as "Executive Monitoring" (EXM). In this section, we will explain the first phase of EXM (EXM-I), which is designed to exclude measurements flagged by any of the SQM, DQM, and MQM algorithms. The measurements that survive EXM-I can then enter the second phase, which will be described in Section 2.5.2.8.

The details of EXM-I fault handling logic are given in [31]. The first step is to build two matrices, called tracking (**T**) and decision (**D**) matrices, to support EXM. Each entry in the matrices is for a single channel or receiver-satellite pair. We record which satellites are physically tracked on each receiver in the **T**-matrix. We construct the **D**-matrix by combining the three QM flags in a logical-OR operation (a flag from any QM algorithm causes a flag in **D**). We then match these matrices to a set of 11 generic failure cases to determine which measurements should be excluded from correction generation. We can consider all but two of these cases as combinations of three fundamental situations:

- a) Flag on a single satellite on a single reference receiver;
- b) Flags on a single satellite and multiple reference receivers;
- c) Flags on multiple satellites on a single reference receiver.

If a flag is found to meet case (a), a single flagged measurement cannot be used. Cases (b) and (c), where multiple flags exist on a single satellite or reference receiver, lead to the exclusion of that satellite or receiver entirely. When flags meet both case (b) and (c), all suspicious satellites and receivers are excluded – no attempt is made to determine whether a satellite or receiver is the cause of the failure, because there is no particular reason to

believe that one is more likely to fail than the other. These rules are conservative because the EXM exclusion logic is based on the assumption that failures and fault-free alarms are truly rare. Thus, any cases with multiple flags should be treated very cautiously.

Prn #	1	2	3	4	11	12	21	31
Rcvr1	✓	✓	✓	✓	✓	✓	×	✓
Rcvr2	✓	✓	✓	✓	✓	✓	✓	/
Rcvr3	✓	✓	✓	✓	×	/	✓	✓

✓: good ×: bad / : not tracked

Figure 2.6: An Example of Selecting a Common Set [30]

Once we determine which measurements should be excluded, we select a common set, S_c , of visible satellites that survive EXM-I. This set of satellites is used to calculate the receiver clock adjustment as shown in Equations (2-5) and (2-6). The measurements of this set are also used to compute candidate corrections as shown in Equations (2-7) and (2-8). We form the common set based on two principles [31]:

1. If all three reference receivers track at least four satellites, the common set comprises all satellites tracked by all three receivers;
2. Otherwise the common set is the largest set of satellites tracked by any two of the three receivers.

As an example of selecting the common set, a table similar to the form of the **D** and **T**-matrices is shown in Figure 2.6. In this case, the common set that we will select is the one

circled by the solid line instead of the dashed line. Let us assume that we cannot approve a common set of at least four satellites, and it was not predicted by an LGF “constellation alert” algorithm that takes known satellite outages into account. If such an event occurs, we should exclude all measurements and reset the IMT.

2.5.2.5 MULTIPLE REFERENCE CONSISTENCY CHECK (MRCC)

MRCC is designed to isolate an anomalous receiver that creates large errors in the candidate corrections. To accomplish this task, we examine the consistency of corrections for each satellite across all reference receivers by computing B -values [30, 31, 49]. These B -values are best thought of as the estimates of pseudorange errors under the hypothesis that a given reference receiver has failed. The B -values are generated by the following equations:

$$B_{\rho,m,n}(k) = \rho_{corr,n}(k) - \frac{1}{M_n(k) - 1} \sum_{\substack{i \in S_n(k) \\ i \neq m}} \rho_{sca,i,n}(k) \quad (2-18)$$

$$B_{\phi,m,n}(k) = \phi_{corr,n}(k) - \frac{1}{M_n(k) - 1} \sum_{\substack{i \in S_n(k) \\ i \neq m}} (\phi_{ca,i,n}(k) - \phi_{ca,i,n}(0)) \quad (2-19)$$

In these equations, ρ_{corr} and ϕ_{corr} are the pseudorange correction and the carrier-phase correction, respectively (see Equations (2-7) and (2-8)), ρ_{sca} and ϕ_{ca} are clock-adjusted corrections (see Equations (2-5) and (2-6)), and S_n denotes a set of M_n reference receivers as explained in Section 2.4.1. Let us assume that a given reference receiver, m , has failed. In that case, the best estimate of the true pseudorange correction is the correction computed by taking the average of all receivers except the one hypothesized to have failed. This is the second term in the right side of Equation (2-18). We can see from Equation (2-7) that the first term, $\rho_{corr,n}(k)$, is computed by averaging all receivers including the hypothetically failed receiver. The difference between these two terms, the B -value, thus expresses the resulting error if the failed receiver were not excluded. It is broadcast to users so that they

can compute an “H1” protection level, which is the position error bound under this hypothesis [33, 49].

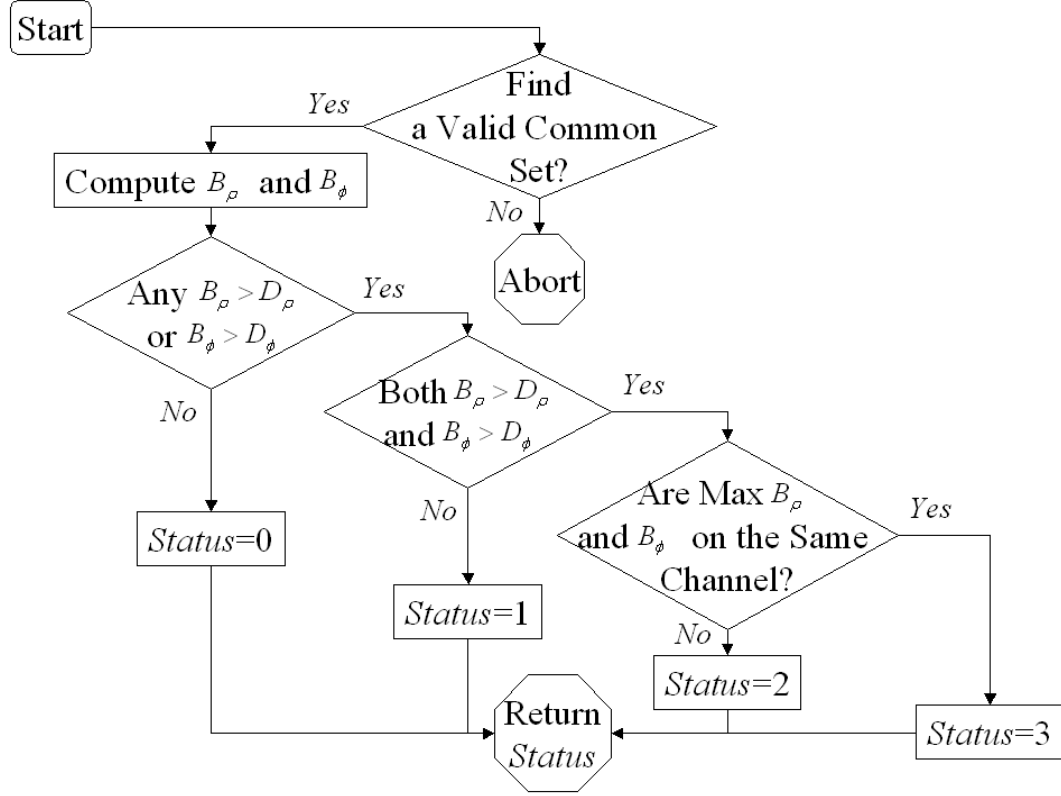


Figure 2.7: MRCC Fault Detection Flowchart [38]

After we compute B -values for pseudorange and carrier-phase measurements, we compare each B -value with its threshold. Then we follow the logic of the MRCC fault detection process as shown in Figure 2.7. If there are any B -values that exceed their thresholds, we find the largest B_ρ and B_ϕ over the thresholds. Next, we determine whether they are B -values of different types (B_ρ and B_ϕ), and whether the different types of the largest B -values are on the same channel.

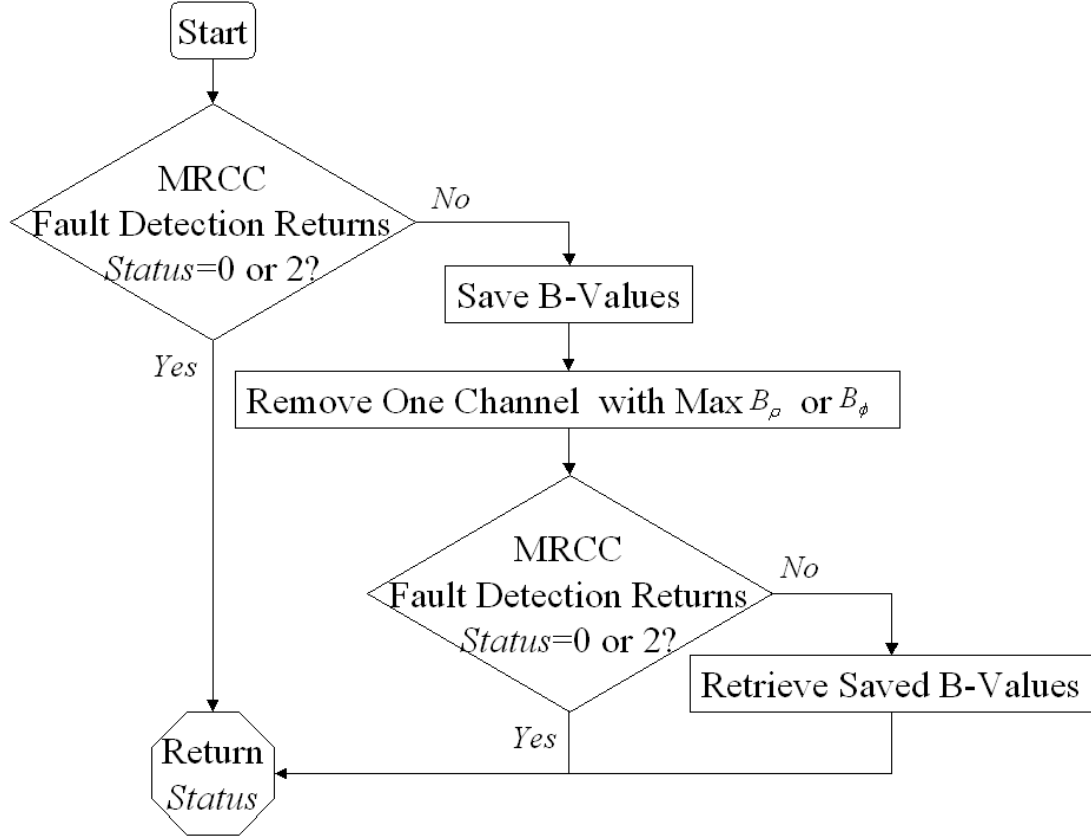


Figure 2.8: EXM-II Pre-Screen Flowchart [38]

The final procedure in MRCC is to isolate any receiver channels that create anomalously large errors in the corrections. The MRCC isolation procedure, which is also referred to as the “EXM-II pre-screen,” is shown in Figure 2.8. We pick up the detection status from Figure 2.7 and continue to isolate errors [31]. If the MRCC fault detection returns a status of 0, which means that no B -values are over threshold, then we do not isolate any channel. If the status is 2, which implies that the largest B_ρ and B_ϕ are not on the same channel, the pre-screen is not able to handle this complex situation. Thus we pass this task to the similar, but more comprehensive, logic of EXM-II (see Section 2.5.2.8). For all other cases, we remove the channel with the largest B -value and then re-compute B -values using a reduced common set. If we see no remaining MRCC flags, we can conclude that the pre-screen was successful. Otherwise, it means that the pre-screen may have mishandled the failure. Therefore we retrieve the original B -values and pass the task from MRCC to EXM-II.

2.5.2.6 SIGMA-MEAN ($\sigma\mu$) MONITORING

The purpose of sigma-mean monitoring is to ensure that the zero-mean Gaussian distribution defined with the broadcast σ_{pr_gnd} overbounds the true error distribution of broadcast differential corrections [16, 50]. We have already described the broadcast σ_{pr_gnd} in Section 2.4.2. We have also derived B -values, which are used as input to this sigma-mean monitor, in the previous section. In order to detect both mean and sigma violations, we implement two different monitoring algorithms: the estimation method and the Cumulative Sum (CUSUM) algorithm [24]. The details of these monitors are a key focus of this thesis and are covered in Chapter 3.

2.5.2.7 MESSAGE FIELD RANGE TEST (MFRT)

The function of MFRT is to confirm that the computed average pseudorange corrections and correction rates are within the confidence bounds [31]. This test is the last step of EXM-II before approving corrections to be broadcast. The computed average pseudorange corrections, ρ_{corr} , from Equation (2-7) should be within a bound of ± 125 meters, and the correction rates, $R_{\rho_{corr}}$, need to be within a threshold of ± 0.8 m/sec [51]. These thresholds are determined based on nominal data (the detailed method is described in [38]). $R_{\rho_{corr}}$ is the rate of change of ρ_{corr} and is equal to:

$$R_{\rho_{corr},n}(k) = \frac{\rho_{corr,n}(k) - \rho_{corr,n}(k-1)}{T_s} \quad (2-20)$$

where the sample interval, T_s , is 0.5 seconds.

2.5.2.8 PHASE TWO OF EXECUTIVE MONITORING (EXM-II)

A significant fraction of the IMT code is dedicated to processing alert messages generated by different IMT monitor algorithms. EXM isolates faulty measurements based on those alerts as needed to meet the integrity requirements while minimizing incorrect exclusions

to the degree possible. In Section 2.5.2.4, it was shown how to execute the first phase of EXM (EXM-I) and remove measurements flagged by the various QMs. This section describes the second phase of EXM (EXM-II), which tries a series of exclusion steps based on MRCC, the $\sigma\mu$ monitor, and MFRT flags [30, 36].

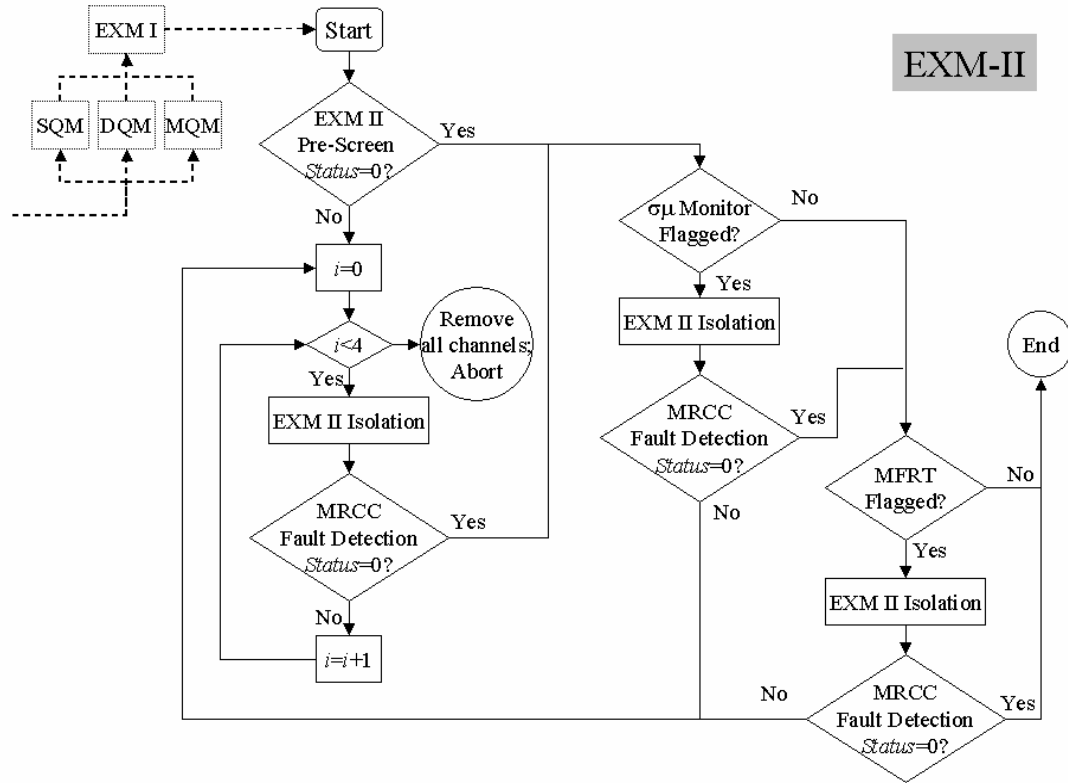


Figure 2.9: EXM-II Flowchart [38]

It is not simple to handle MRCC flags because we usually need to reduce the common set S_c when we attempt to exclude measurements. When this happens, we must re-compute corrections and B -values for all remaining measurements and confirm that this new, reduced set of measurements passes MRCC. This procedure has been named “EXM-II Pre-Screen”, and is shown in Figure 2.8. Including this block, Figure 2.9 shows the whole diagram of the EXM-II flowchart. In this diagram, “EXM-II isolation” operates based on logic similar to that of EXM-I. In other words, single B -value flags on individual channels are removed, but if more than one B -value is flagged on a given satellite or receiver, the

entire satellite or receiver is excluded. In order to complete EXM-II before the next epoch begins, we place a limit of four attempts to exclude measurements and repeat the EXM-II isolation-MRCC fault detection process before we “give up” and exclude all measurements. In addition, we must confirm that all measurements pass the $\sigma\mu$ monitor and MFRT. Any flags from these tests also demand EXM-II exclusions.

Lastly, we complete the EXM-II process by managing excluded ranging sources and/or receiver channels. Once we exclude faulty measurements detected by the $\sigma\mu$ monitor or MFRT, they are excluded for the entire pass of the satellite in question. This is because it is difficult to verify that the faults detected by these monitors are no longer present during the current pass. In most other cases, we put excluded measurements into a “self-recovery” mode, in which the carrier smoothing filters are re-initialized. During self-recovery, we compare the re-smoothed measurement with thresholds tightened to 3-sigma levels – the original thresholds are approximately at 6-sigma levels. If this measurement passes all tests, it is declared safe for use again and is reintegrated into the set of usable measurements [31, 33]. We also apply the tightened thresholds to measurements excluded due to $\sigma\mu$ monitor or MFRT flags in the previous pass. These stringent thresholds are used to insure that the probability of reintroducing a measurement that still fails is below 1.94×10^{-9} [52]. If a rejected channel does not pass the tests with the tightened thresholds, it enters the self-recovery mode again and repeats the recovery process. If self-recovery fails a second time, the defective channel enters an “external maintenance” mode and cannot be used again until the affected equipment is certified to be healthy by an external agency, such as FAA maintenance personnel [31].

2.6 CONCLUSION

This chapter provided the background information on LAAS: the description of the LAAS architecture, the requirements of the LAAS performance (accuracy, integrity, continuity, and time availability), and the various types of LAAS precision approaches and landings. We have also seen how the LGF computes and broadcasts differential corrections to users in real time. However, users cannot eliminate the errors of their pseudorange measurements entirely even after applying these differential corrections. What is important here is that the

residual error bounds should also be sent, because this information is essential for computing position bounds within the specified integrity level. The chapter described the broadcast sigma of the ground facility error representing various allocation models.

The last part of this chapter presented the functions of the IMT integrity monitoring algorithms. Each monitor is designed to detect and exclude its targeted faulty measurements before those faults influence the broadcast correction. The sigma-mean monitor is one of the key integrity monitors, since the broadcast sigmas are used in the calculation of the protection levels. In Chapter 3, we will see how the sigma-mean monitoring algorithms were incorporated to the IMT.

Chapter 3

Sigma-Mean Monitoring

3.1 INTRODUCTION

In the previous chapter, we saw how GPS measurements are processed to obtain differential corrections and integrity parameters in real time. We also investigated how the LGF integrity monitors are designed to detect and remove anomalies such as ground-based or space-based system failures. This chapter focuses on monitoring the standard deviations of pseudorange correction errors, σ_{pr_gnd} , which are one of the key integrity parameters broadcast to users. As explained in Chapter 1, aircraft use these σ_{pr_gnd} values to compute the Vertical Protection Level (VPL) and the Lateral Protection Level (LPL) as position bounds. Since user navigation integrity is quantitatively appraised by the position bounds and these bounds are based on σ_{pr_gnd} , user integrity relies on these “sigmas”. Because of the direct connection between the broadcast σ_{pr_gnd} and user integrity, real-time sigma monitoring is necessary to detect anomalies or other events in which the true sigma exceeds the broadcast σ_{pr_gnd} during LAAS operations.

First, we need to define the “threat space” (or integrity risk space) for sigma and mean monitoring. Section 3.2 establishes the territory of unexpected anomalies which the sigma-mean monitoring is targeted to resolve. Note that these anomalies are not large enough to

represent immediate hazards; failures resulting in immediate hazards must be detected within the time-to-alert by the other monitors described in Chapter 2. Second, Section 3.3 presents sigma-monitoring methods to detect violations with acceptable residual integrity risk. Both the sigma estimation and the Cumulative Sum (CUSUM) methods are useful in this respect. These algorithms have been implemented in the IMT and have been tested under both nominal and failure conditions. In Section 3.4, we turn our attention to monitoring the mean. Real-time monitoring of the mean of pseudorange correction errors is needed to detect unforeseen conditions that cause the true mean to become substantially non-zero. We will also describe the mean estimation and CUSUM algorithms implemented in the IMT and report their test results. The results of both nominal and failure tests demonstrate that the sigma-mean monitoring algorithms can detect failures large enough to threaten user integrity and are integrated smoothly with the Executive Monitoring (EXM) Phase II logic in the IMT [38]. We will then compare the estimation and CUSUM methods analytically in Section 3.5. Overall, the estimation method more rapidly detects small violations of σ_{pr_gnd} , but the “fast-impulse-response” (FIR) CUSUM variant more promptly detects significant violations that would pose a larger threat to user integrity.

3.2 THREAT SPACE

There are two major assumptions in the calculation of protection level (PL). One is that the error distribution of the broadcast pseudorange corrections can be characterized with a zero-mean, normally distributed fault-free error model. The other is that the error distribution based on the broadcast σ_{pr_gnd} overbounds the “true” error distribution so that the computed PL also bounds the true position error. However, these assumptions may not hold when man-made or natural system failures are experienced. In particular, a sudden change in environmental conditions – as an example, a fire truck enters the protected zone of the LGF – may cause multipath errors to increase significantly. In such cases, one significant integrity risk is that the mean of the pseudorange correction error distribution becomes non-zero. To make matters worse, the standard deviation (sigma) for the correction error may grow to exceed the broadcast correction error sigma, σ_{pr_gnd} . We therefore need real-time mean and sigma monitoring to help insure that the true error

distribution is bounded by the zero-mean Gaussian distribution defined by the broadcast sigma value.

3.3 SIGMA MONITORING

Sigma monitoring plays an important role in ensuring that the possibility of the true sigma exceeding the broadcast sigma poses no significant integrity risk to LAAS users. Among many proposed algorithms for monitoring sigmas in real-time, the following two algorithms – the estimation method and the CUSUM – have been implemented in the Stanford IMT. The detailed algorithms are described in Sections 3.3.1 and 3.3.2, respectively. Next we analyze the expected performance of the methods theoretically. To show that the requirements in the LGF specification are achievable [15], it is important to verify that the monitor works as designed. Thus, in Section 3.3.3, we will explain how to conduct verification testing under both nominal and failure conditions, as well as show some test results.

3.3.1 SIGMA ESTIMATION METHOD

The real-time sigma estimation method estimates sample standard deviations of the pseudorange correction error from LGF B -values, B_ρ , computed in the Multiple Reference Consistency Check (MRCC) for each visible satellite. As explained in Chapter 2, since the B -values represent pseudorange correction differences across reference receivers – ideally, the pseudorange corrections from all reference receivers should be the same for a given satellite – the B -values represent pseudorange correction errors that would exist if a given reference receiver has failed [30, 31].

3.3.1.1 ALGORITHM

The normalized values of B_ρ (i.e., B -values divided by their theoretical sigmas, σ_{B_ρ}) are the inputs to sigma estimation:

$$B_{\rho_normal,m,n}(k) = \frac{B_{\rho,m,n}(k) - \mu_{B_{\rho,n}}(k)}{\sigma_{B_{\rho,n}}(k)} \quad (3-1)$$

$$\sigma_{B_{\rho,n}}(k) = \frac{\sigma_{pr_gnd,n}(k)}{\sqrt{M_n(k) - 1}} ; \quad \mu_{B_{\rho,n}}(k) = 0 \quad (3-2)$$

where we can compute $\sigma_{pr_gnd,n}(k)$ using Equation (2-10) at the epoch k , and M_n is the number of reference receivers for the n th ranging source. Now that we have the normalized B -values, we can compute the estimated sigma, $\hat{\sigma}_{B_{\rho_normal}}$:

$$\hat{\sigma}_{B_{\rho_normal,m,n}}(k) = \sqrt{\frac{1}{k-1} \sum_{i=1}^k [B_{\rho_normal,m,n}(k) - \mu_{B_{\rho_normal,n}}(k)]^2} \quad (3-3)$$

According to the Gaussian error model, the estimated sigma has a chi-square distribution, given by:

$$(N(k)-1) \frac{\hat{\sigma}_{B_{\rho_normal,m,n}}^2(k)}{\sigma_{B_{\rho_normal,m,n}}^2(k)} \sim \chi^2 \{N(k)-1\} \quad (3-4)$$

with $N(k)-1$ degrees of freedom, where $N(k)$ is the number of independent samples used to derive the estimate at the epoch k . Note that the interval between independent B -values is expected to be equal to 200 seconds, which is twice as long as the time constant of the carrier-smoothing filter. In other words, the relation between the number of independent samples and k is:

$$N(k) = k / (T_s \times 200) = k / 400 \quad (3-5)$$

Typically the sample interval, T_s , is 0.5 seconds. Regardless of T_s , it takes at least one hour to collect 18 independent B -values.

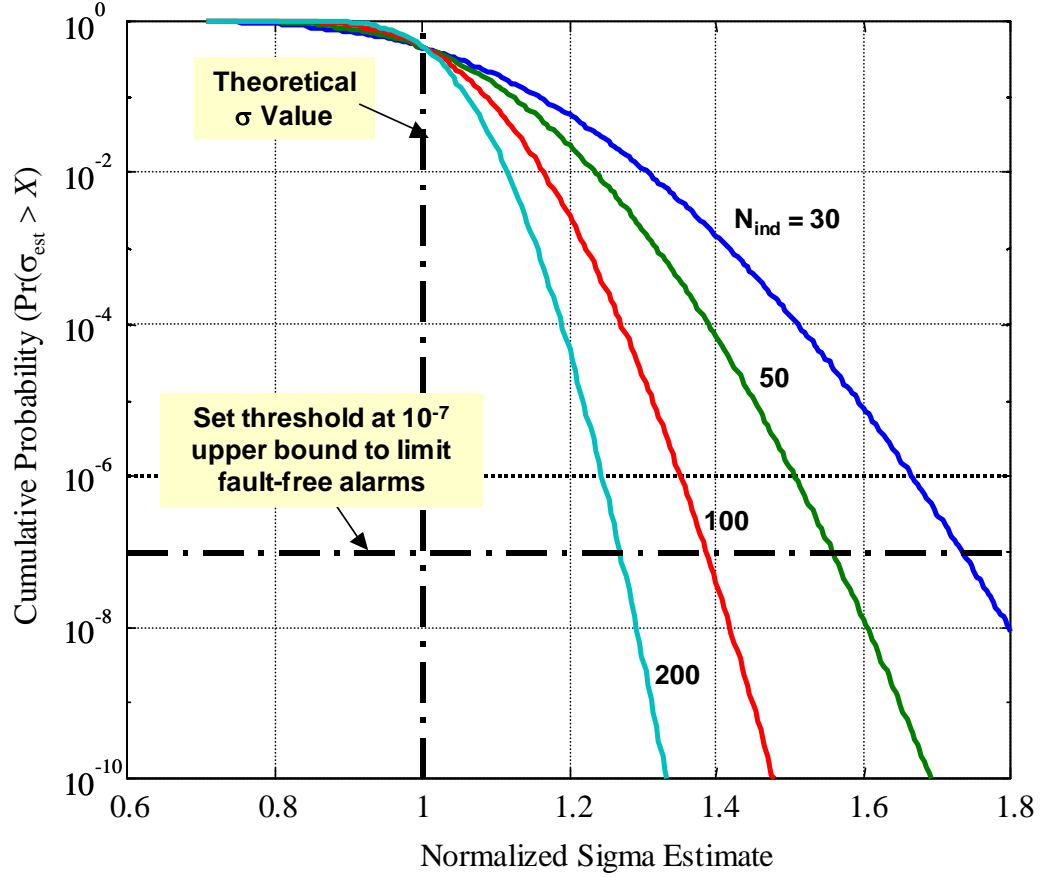


Figure 3.1: Chi-Square Distribution of Sigma Estimate

Figure 3.1 shows the resulting cumulative distributions of the sigma estimates for varying numbers of independent samples. As expected, more independent samples provide tighter distributions on $\hat{\sigma}_{B_{\rho_normal}}$. Based on this chi-square distribution, the detection threshold is set to provide an acceptably low fault-free alarm rate (10^{-7} , based on a sub-allocation of the specified Category I continuity risk allowed per 15-second interval [15]). The estimated sigma is then compared to this time-dependent threshold, which is lowered as more independent samples are collected. Any alerts generated from this monitor are passed on to the second phase of EXM (EXM-II) for resolution, as described in Section 2.5.2.8.

3.3.1.2 THEORETICAL ANALYSIS

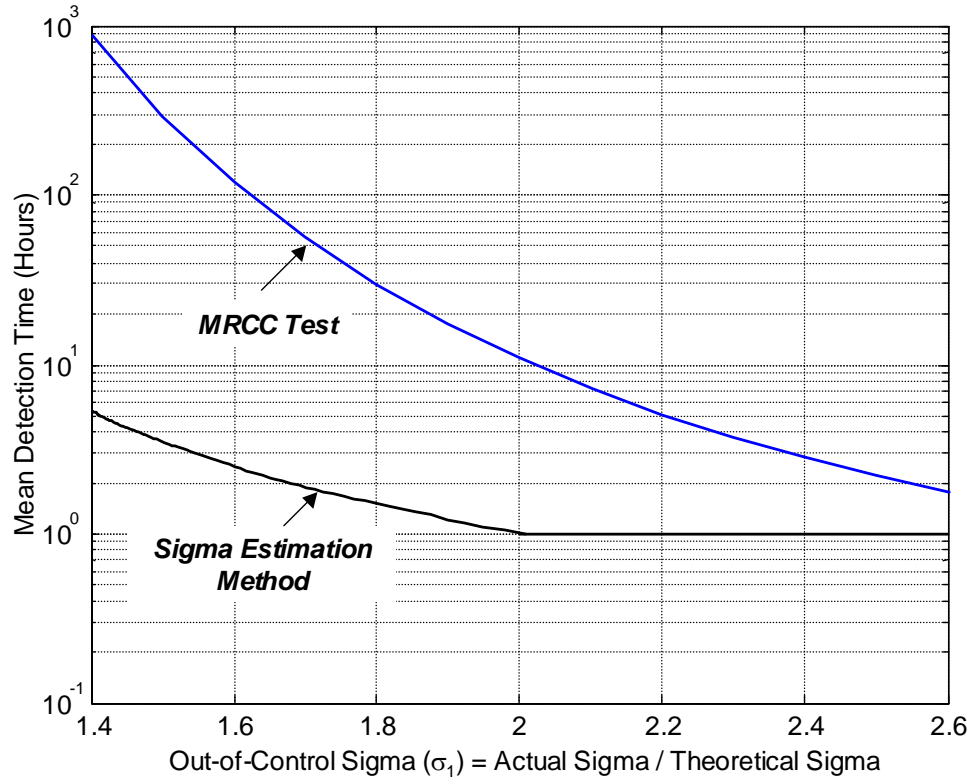


Figure 3.2: Performance of Sigma Estimation Method and MRCC Test

As introduced in Chapter 2, the existing Multiple Reference Consistency Check (MRCC) has some utility to screen IMT B -values on an epoch-by-epoch basis as a sigma monitor, but its times-to-alert are much longer than that for sigma estimation. Figure 3.2 compares the theoretical performance of the MRCC test to the sigma estimation method. We compute the mean number of independent samples prior to detection based on the probability of the normalized sigma exceeding thresholds and convert it into detection time. The out-of-control sigma, σ_1 , is a potential sigma violation as the ratio of an actual sigma over a theoretical sigma. We can see that the MRCC test requires much more time for detection than the sigma estimation method does. This is because the MRCC test checks only the latest B -value, while the sigma estimation method accumulates prior information. However, the sigma estimation method has a limitation of one hour to detect a sigma ratio

greater than 2 because we need at least 18 independent samples (one hour) to estimate the sigma in order to rely on this sigma estimation method. As shown in Figure 3.1, the thresholds for this method are derived from the chi-square distribution assuming that random samples of size N are taken from a population having a normal distribution and that the sample variance has a chi-square distribution with $N-1$ degrees of freedom. Yet, we cannot assure that the population (i.e., the error distribution) is a normal distribution. In order to reasonably use the chi-square distribution, we need a large N , so that the distribution of the sum of N independent variables approaches a normal Gaussian distribution by the Central Limit Theorem. In this sense, $N=18$ represents an assessment of the minimum allowable number for this application.

3.3.2 SIGMA CUMULATIVE SUM METHOD

Cumulative Sum (CUSUM) monitoring is very simple and is relatively easy to analyze. It can be shown to be "optimal" in terms of minimizing time-to-alert under specified failure conditions [23]. It is thus commonly used in manufacturing, where the goal is to detect poor-quality products with reasonably low missed-detection and false-alarm rates (but nowhere near as low as required by the LGF). The CUSUM is theoretically the most expeditious tool to detect small but persistent shifts of random process parameters under two assumptions [23]: first, the input to the CUSUM can be characterized with a Gaussian distribution under nominal conditions and is statistically independent at each epoch; second, we know the "true" mean and standard deviation of the parameter of interest under nominal conditions.

3.3.2.1 ALGORITHM

The basic idea of the CUSUM is to maintain running sums of statistically independent quality metrics (see Appendix B.1). A "windowing factor" k is subtracted from the running sum at each update. This factor is chosen to minimize the time-to-alert for a particular failure case with a specific out-of-control distribution. If the targeted fault case is a large deviation from nominal, k will be large as well to reduce the sensitivity of the CUSUM to smaller anomalies. If the targeted fault is closer to nominal performance, k gets smaller, but

the price is more fault-free alerts unless the CUSUM threshold, h , is increased to compensate.

As implemented in the IMT, the CUSUM method collects cumulative sums ($C_{m,n}^+$) of squared and normalized B -values ($Y_{m,n}$) for each receiver channel (m, n) tracking a GPS satellite and is updated every 200 seconds. Note that updates must be statistically independent in time. If we increment a CUSUM with highly correlated inputs greater than the k factor, it will quickly exceed the threshold as similar values are added one after another. In this case, each independent epoch, N , corresponds to two carrier-smoothing time intervals (see Equation (3-5)). We start the CUSUM at zero or a head-start value of $H^+ > 0$ and then increment each epoch by the size of the monitored input, Y , minus the windowing factor, k .

$$\begin{aligned} C_{m,n}^+(0) &= 0 \text{ or } H^+ \\ C_{m,n}^+(N) &= \max \left(0, C_{m,n}^+(N-1) + Y_{m,n}(N) - k_\sigma^+ \right) \end{aligned} \quad (3-6)$$

$$Y_{m,n}(N) = \left(\frac{B_{\rho,m,n}(N) - \mu_{B_{\rho,n}}(N)}{\sigma_{B_{\rho,n}}(N)} \right)^2 \quad (3-7)$$

$$k_\sigma^+ = - \frac{\ln(\sigma_1) - \ln(\sigma_0)}{(2\sigma_1^2)^{-1} - (2\sigma_0^2)^{-1}} \quad (3-8)$$

This desired failure slope, k , is based on the targeted out-of-control sigma, σ_1 , that represents “failed” performance (the derivation of the k factor is shown in Appendix B.2). The sigma in a nominal case is defined as σ_0 and is called the in-control sigma. Since the CUSUM is sensitive to only one direction, separate upward “+” and downward “-” CUSUMs are used. However, this is not true if sigma is the input, as decreasing sigma is not a concern. As shown in Equation (3-6), if the CUSUM falls below zero on a given epoch, it is reset to zero. If the sum is above zero at any update epoch, the CUSUM is

compared to a threshold, h , which does not vary with time. If it accumulates to above the threshold, an alert is issued.

While CUSUM behavior is more complicated than sigma estimation, it can be analyzed via Markov Chains (MCs) [23]. CUSUMs follow the Markov property, since each CUSUM belongs to a finite set of states, and the distribution of the CUSUM state at epoch N depends only on its state at the previous epoch $N-1$ and the distribution of the incrementing value at epoch N . In Figure 3.3, we introduce the MC model for CUSUMs. Using a mesh of width $\Delta = h / M$, we discretize the range of possible C^+ values into the states:

$$\begin{array}{ll} \text{State } 0 & C^+ = 0 \\ \text{State } i & C^+ \in [(i-1)\Delta, i\Delta] \quad i = 1, 2, \dots, M \\ \text{State } M+1 & C^+ > h \end{array} \quad (3-9)$$

Under either nominal or specified failure conditions, we can derive a "one-step" MC transition matrix, \mathbf{P} , with elements:

$$P_{i,j} = \Pr(C^+(N+1) \in \text{State } j \mid C^+(N) \in \text{State } i) \quad i, j = 0, 1, \dots, M+1 \quad (3-10)$$

This transition matrix gives the probability of going from each discretized CUSUM state between zero and the threshold h on epoch N to each possible state on epoch $N+1$. We can determine each element by using the known distribution function of the input, Y , and using Simpson's rule for a more accurate approximation of the transition probabilities (see Appendix B.3). A state is called "absorbing" if the system remains in the state once it enters there, as shown in the right-bottom corner of Figure 3.3. In other words, when the CUSUM exceeds its threshold, it will be restarted.

From this \mathbf{P} , we can compute the steady-state distribution of the MC and thus determine how long, on average, it takes for the MC to exceed a given value [23]. If this value is the threshold h , this gives the average run length (ARL). We want long runs before false alarms occur but short runs after the parameters actually shift to insure low integrity risk. The nominal (or in-control) ARL should be sufficiently long in order to result in a low

probability of false alarms. In contrast, a short out-of-control ARL is desired to detect failures as quickly as possible.

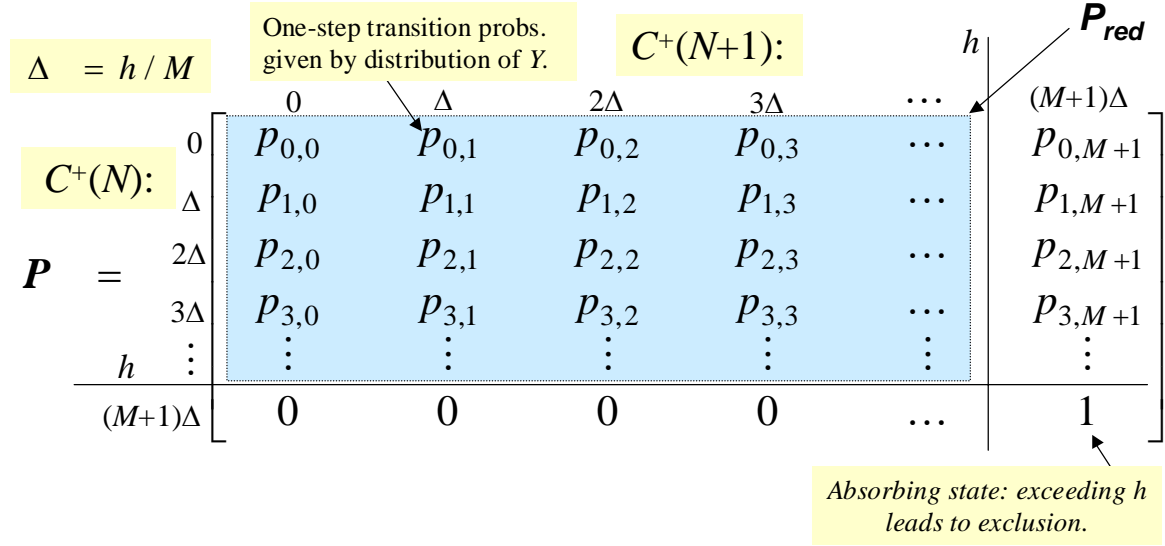


Figure 3.3: CUSUM Performance Modeling with Markov Chains

An alternate approach to finding ARLs is to solve the following matrix representation:

$$(\mathbf{I} - P_{red})\lambda = \mathbf{1} \quad (3-11)$$

We use the matrix, P_{red} , which disregards the transitions to and from the last state, $M+1$. λ is the vector of ARL values with a length of $M+1$, and each component represents the ARL starting in the corresponding states $0, 1, \dots, M$. \mathbf{I} is the identity matrix, and $\mathbf{1}$ is a vector of length $M+1$ all of whose elements are 1. The normal solution process is to guess a threshold value h , form the nominal MC, and then solve for the in-control ARL, iterating on h until the resulting in-control ARL equals the inverse of the desired fault-free alert rate [24].

In addition, multiplying \mathbf{P} by itself d times (e.g., computing \mathbf{P}^d) gives the transition probabilities between epoch N and epoch $N+d$. This allows us to determine, from the failure-state MC, the number of epochs required to exceed the threshold with a given missed-detection probability. Thus, we change the MC to represent the failed or "out-of-

control" state, compute the out-of-control ARL (or mean time to detect), and successively multiply the \mathbf{P} matrix by itself until it shows a probability of exceeding the threshold that is one minus the desired missed-detection probability [23, 24].

3.3.2.2 THEORETICAL ANALYSIS

As a function of σ_1 , Figure 3.4 shows the resulting failure-state average run lengths for the various CUSUMs: zero-start ($H^+ = 0$) and "head-start" ($H^+ = h/2$ or $3h/4$) CUSUMs. The "head start" or "Fast Initial Response (FIR)" CUSUM improves its performance if the process is already out of control when the CUSUM charting begins. We can see that the FIR CUSUMs result in much faster detection because starting the CUSUM part way toward the threshold will hasten detection compare to starting it at zero. However, we have to pay for this improved initial response somewhere. To maintain the same in-control ARL, it may be necessary to increase the threshold, resulting in slightly slower detection if the process starts in control.

Figure 3.5 shows the CUSUM threshold (h), which is set to achieve the desired average run length ($ARL = 10^7$ independent epochs) under nominal conditions based on the sub-allocated LGF continuity requirement, and the optimal k given by the target out-of-control sigma (σ_1). Because the threshold must be very large to achieve an ARL of 10^7 epochs, the ARL under nominal conditions is practically the same for the zero-start ($H^+ = 0$) and 'fast-impulse-response' (FIR; $H^+ = h/2$ or $3h/4$) CUSUMs. Thus, the thresholds are almost the same, and there is little continuity penalty for using the FIR CUSUM. Under these conditions, the CUSUM with the highest H^+ that does not give a significant increase in the fault-free alarm probability is optimal.

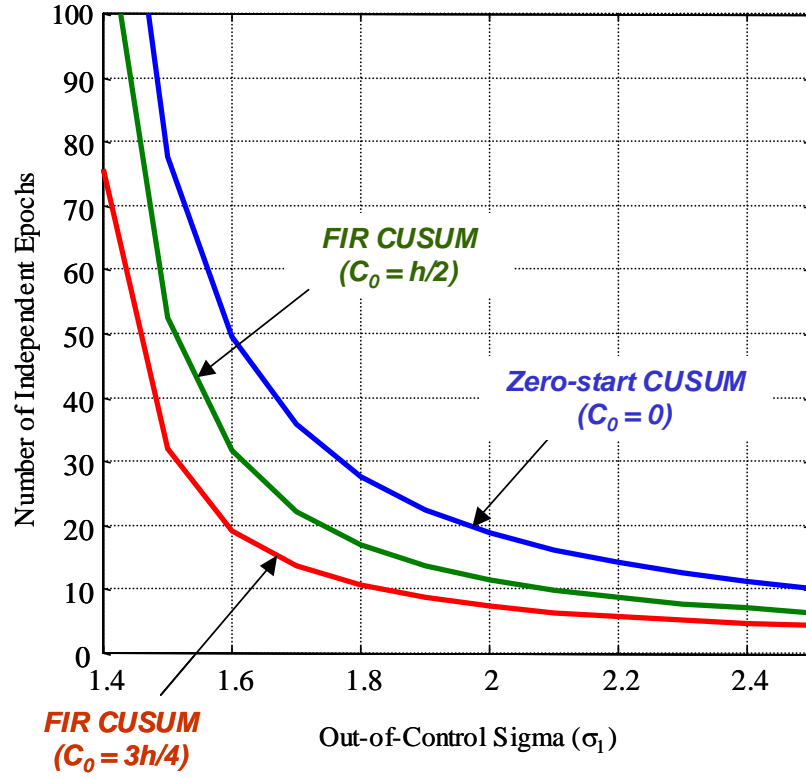


Figure 3.4: Failure-State ARLs for Sigma CUSUM Method

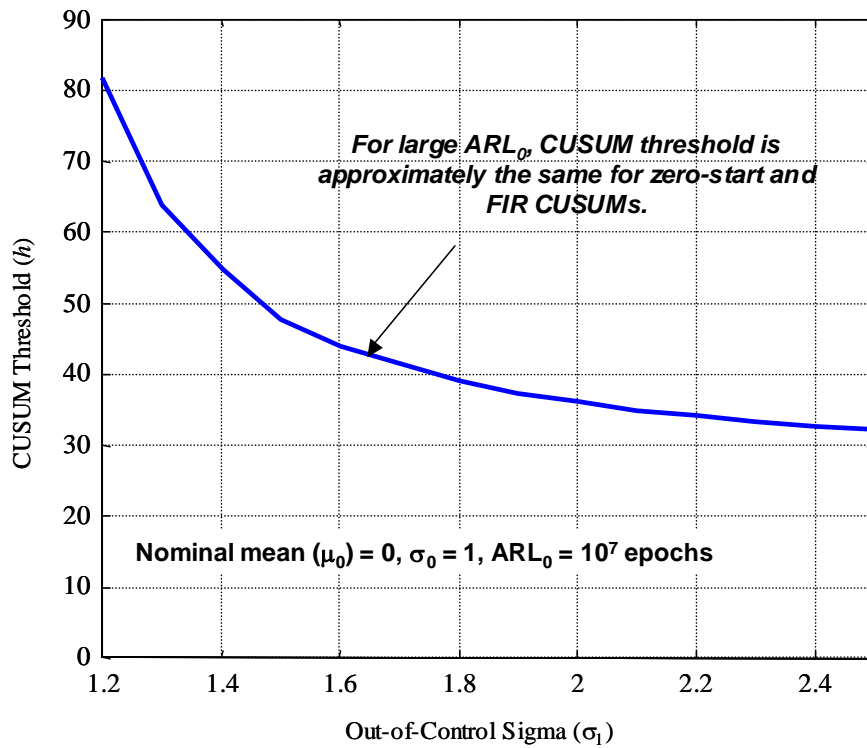


Figure 3.5: Threshold for Sigma CUSUM Method

3.3.3 IMT TEST RESULTS

Based on analytical results shown in Sections 3.3.1.2 and 3.3.2.2, we have implemented and tested the sigma estimation and CUSUM methods in the IMT under both nominal and failure conditions. Under nominal conditions, both sigma estimates and CUSUMs stay below the relevant detection thresholds for all visible satellites in the IMT dataset (a 4 hour dataset collected on March 20, 2000) we have tested. In failure testing, we induce sigma violations by modifying stored IMT receiver packets collected under nominal conditions to represent sigma anomalies, and inject the altered measurements back into the IMT in a post-processing mode [36]. Both sigma estimation and CUSUM methods reliably detect injected sigma violations, although both methods are limited by the 200-second interval between independent B -values.

3.3.3.1 NOMINAL TESTING

Figure 3.6 shows the results of applying the sigma estimation algorithm to the IMT under nominal conditions. The darker (blue) curves show the normalized sigma estimate of a satellite (SV 9) on three reference receivers, and the lighter (green) curves show the detection thresholds, which get smaller and converge to one over time as the number of independent samples increases. As mentioned earlier, monitoring starts after 18 independent B -values have been collected, which corresponds to one hour with a 200-second interval between independent updates. The normalized sigma estimates stay well below the detection thresholds and converge toward one over time. Thus, the theoretical sigma, σ_{B_p} (see Equation (3-2)), which depends on satellite elevation angles, appears to be a good estimate. Very similar results have been obtained from other satellites in this and other IMT datasets.

The ‘zero-start’ CUSUM and FIR CUSUM variants have also been tested with the same IMT data under nominal conditions. The top plot in Figure 3.7 presents the zero-start CUSUM result for Satellite 2 and IMT Reference Receiver (RR) 2, and the lower plot shows the normalized B -values from Equation (3-1) that fed the CUSUM. The CUSUM in this case is targeted at an out-of-control sigma twice that of the theoretical sigma ($\sigma_1 = 2$),

which gives a high windowing factor ($k = 1.848$). The CUSUM rarely departs far from zero due to subtraction of k at each independent B^2 update and stays well below the threshold ($h = 36$).

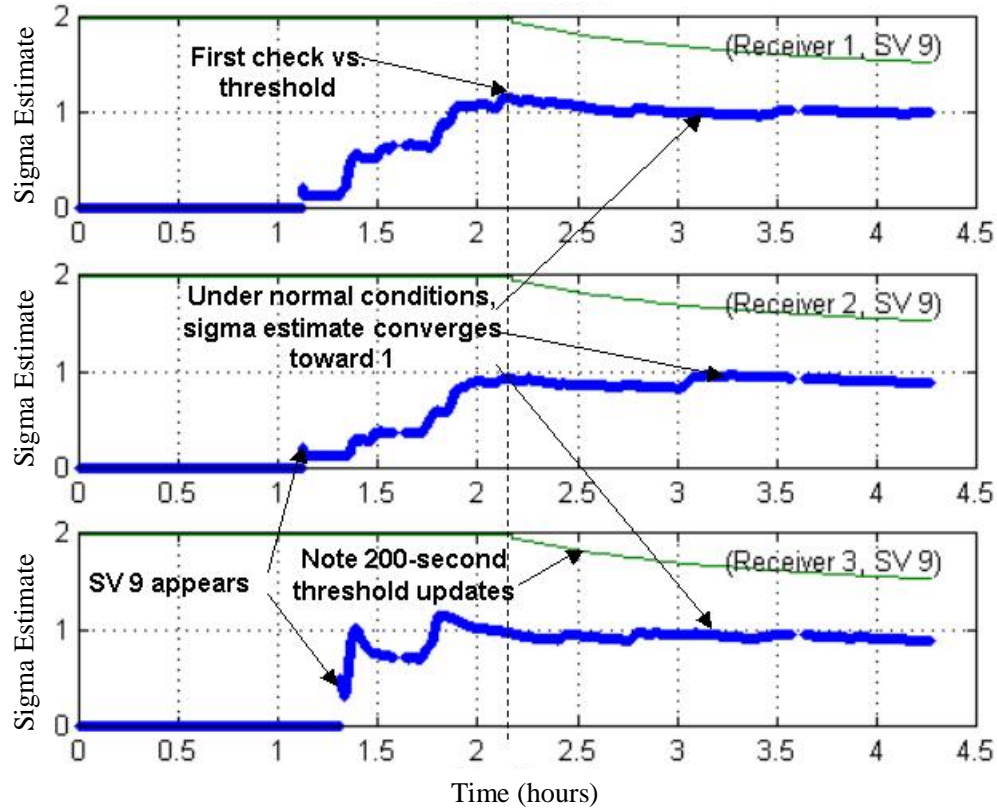


Figure 3.6: Sigma Estimation Results from IMT Nominal Data (the thick blue curves show the sigma estimates and the light green curves show the detection thresholds)

The FIR CUSUM result of Satellite 7 and IMT Receiver 1 is shown on the top plot of Figure 3.8. The CUSUM is initialized at $h/2 = 18$ and is reset there every time the CUSUM falls below zero. Recall that the CUSUM is updated every 200 seconds so that successive updates are statistically independent. Under nominal conditions, the CUSUM slowly falls toward zero and is then reset, since the normalized B^2 is usually below k and k is subtracted off at each epoch. The other satellites tracked by this IMT dataset show very similar patterns for both the zero-start CUSUM and the FIR CUSUMs. The threshold of 36 is never threatened, and no flags are observed at all.

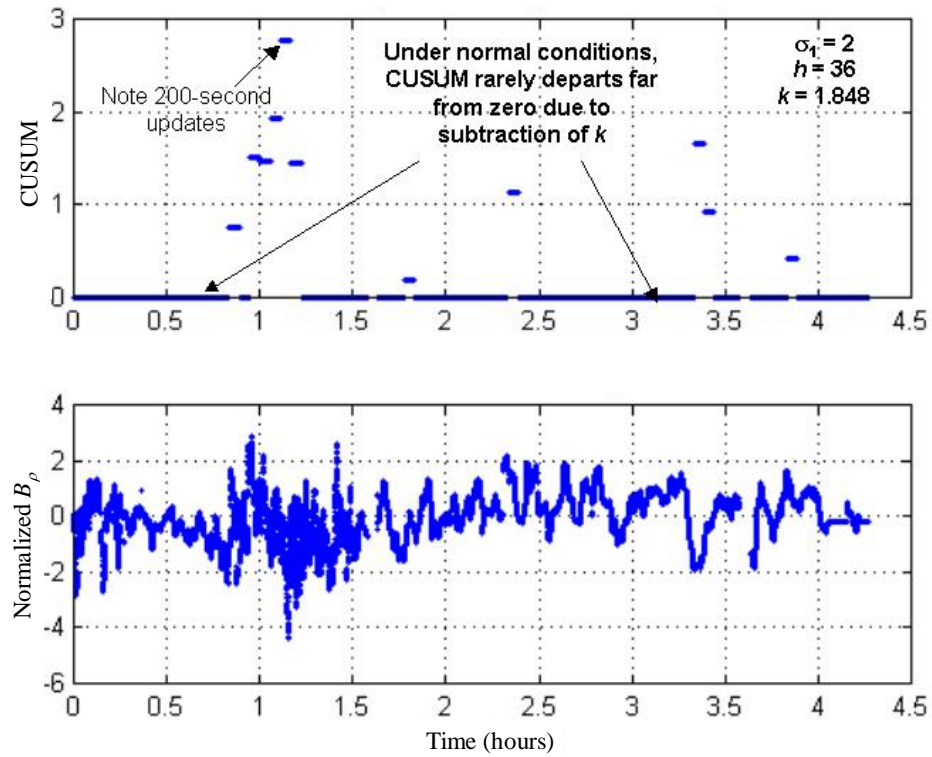


Figure 3.7: Zero-Start CUSUM Result from IMT Nominal Data

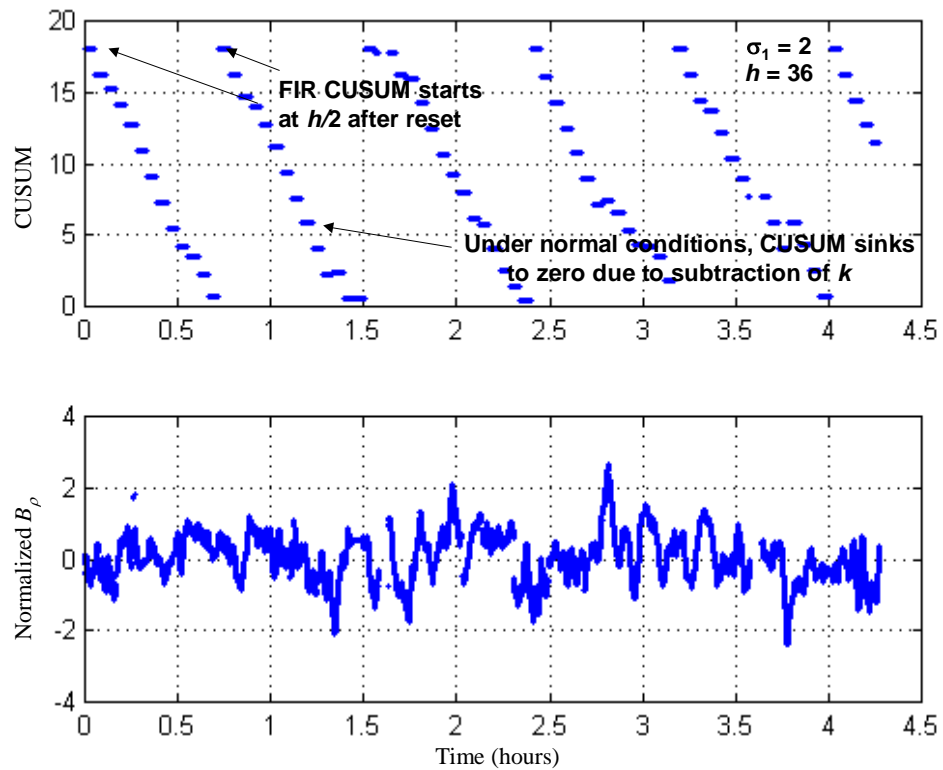


Figure 3.8: FIR CUSUM Result from IMT Nominal Data

3.3.3.2 FAILURE TESTING

In order to conduct failure testing, we inject controlled errors into the IMT and test the detection of anomalies with the current sigma monitoring algorithms. We first induce sigma violations by inserting errors into stored nominal receiver packets. To roughly estimate the nominal error in ρ_{raw} , we use pseudorange-minus-carrier measurements with a polynomial fit to remove ionosphere divergence (multipath and receiver noise errors are left afterward). The estimated error for Channel (m, n) at the epoch k is:

$$error_{estimate,m,n}(k) = [\rho_{raw,m,n}(k) - \phi_{raw,m,n}(k)] - Polyfit[\rho_{raw,m,n}(k) - \phi_{raw,m,n}(k)] \quad (3-12)$$

where *Polyfit* finds the coefficients of a polynomial curve that fits the data in a least square sense. The sigma of the nominal pseudorange error is approximately increased to L times the previous value by adding $(L-1)$ times this error estimate to the nominal measured ρ_{raw} :

$$\rho_{failed,raw,m,n}(k) = [error_{estimate,m,n}(k)] \times (L-1) + \rho_{raw,m,n}(k) \quad (3-13)$$

We then put these modified pseudorange measurements back into the IMT in a “post-processing” mode (In this mode, the data packets stored in a non-volatile storage device can be processed repeatedly [36].)

Figure 3.9 shows the results of applying the sigma estimation algorithm under failure conditions. The pseudorange error on Channel (RR 2, SV 2) is increased to $L = 3$ times the nominal error. The dark (blue) line of the second plot shows the normalized sigma estimate, which exceeds the detection threshold through the whole IMT run. For the purpose of this test, the Executive Monitoring (EXM) logic is not active for sigma monitoring flags, such that the sigma values are estimated over time without a reset. After integrating sigma monitor flags into EXM, the flagged measurement will be excluded by EXM, and its sigma value will be reset upon recovery of the measurement (after the failure goes away). The EXM fault-isolation logic has been tested in prior work [36, 38] and is summarized in Chapter 2. Since threshold checks must wait until enough independent samples have been collected for the sigma estimates to be reliable, the initial transient is

ignored – the first check starts when 18 independent B -values have been collected (corresponding to 1 hour) as shown in Figure 3.9 – and sigma estimates for Channel (RR 1, SV 2) and Channel (RR 3, SV 2) do not exceed the threshold over time. However, the sigma estimates for those channels also converge to values over 1 due to the fact that the B -values are correlated across the three receivers (recall that the B -values are the difference between the correction computed by taking the average of all receivers except the one hypothesized to have failed and the correction computed by averaging all receivers including the hypothetically failed receiver).

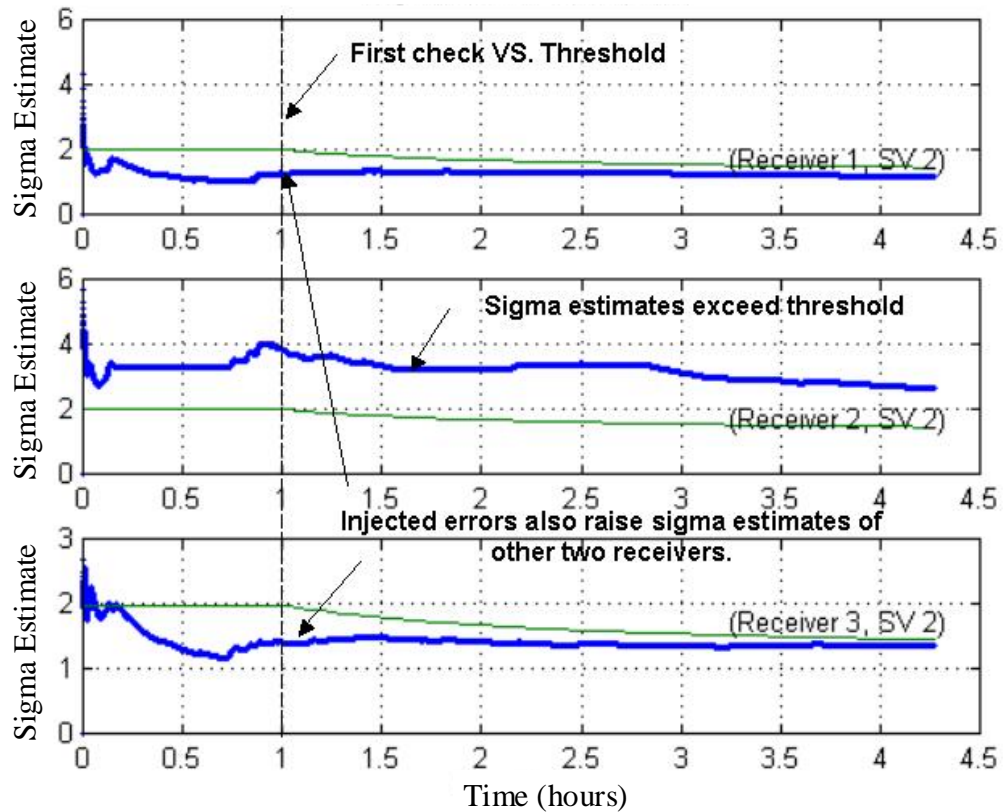


Figure 3.9: Sigma Estimation Results from Failure Test (the thick blue curves show the sigma estimates and the light green curves show the detection thresholds)

We have performed the failure test of IMT sigma monitoring with various L factors. With $L = 1.7$ on the same Channel (RR 2, SV 2), the sigma estimates remain just under thresholds, meaning that no flag was issued before the end of the run. The sigma estimates of the other two receivers and this satellite appear essentially nominal.

Figure 3.10 shows the result of applying the FIR CUSUM variant to failure-injected IMT data for the same Channel (RR 2, SV 2) shown in the previous plot. The CUSUM is tuned to be optimal at an out-of-control sigma twice that of the theoretical sigma ($\sigma_1 = 2$), which gives the windowing factor ($k = 1.848$). Based on this windowing factor and the ARL, the detection threshold (h) is determined to be 36. The FIR CUSUM, initialized at $h/2 = 18$, adds up the increased normalized B -values due to the severe error factor ($L = 3$). The CUSUM crosses the threshold at the third independent epoch – which corresponds to 600 seconds (10 minutes) – after the fault is injected. This is much faster than the sigma estimation method for a newly risen satellite because of the 1-hour delay before sigma-estimate threshold checks can be made. Recall that the sigma estimation method requires this 1-hour delay for applying the Central Limit Theorem to the sample distribution, which the CUSUM method does not need. Since EXM is not active in order to demonstrate how CUSUM responds with respect to injected failures, Channel (RR 2, SV 2) is not excluded, and the CUSUM continually grows regardless of subtraction of the windowing factor at each independent B^2 update. The flat line on the lower plot indicates that the Multiple Reference Consistency Check (MRCC) isolates normalized B -values at this point. In failure tests like this one where severe errors are injected, very similar results have been obtained from the other satellites tracked by the IMT in this dataset.

The top plot in Figure 3.11 shows the FIR CUSUM result of the nominal reference receiver (RR 3) and the same satellite (SV 2) affected by injected pseudorange errors on Channel (RR 2, SV 2). We can see that the FIR CUSUM slightly exceeds the threshold at epoch 21 due to the B -value correlation among three reference receivers and then decreases towards zero. However, when EXM is active in a real operation, we see a different result: Because the source of the failure is excluded after 10 minutes as expected from Figure 3.10, the B -values on Channel (RR 1, SV 2) and Channel (RR 3, SV 2) would be protected from the RR 2 error and would return to normal.

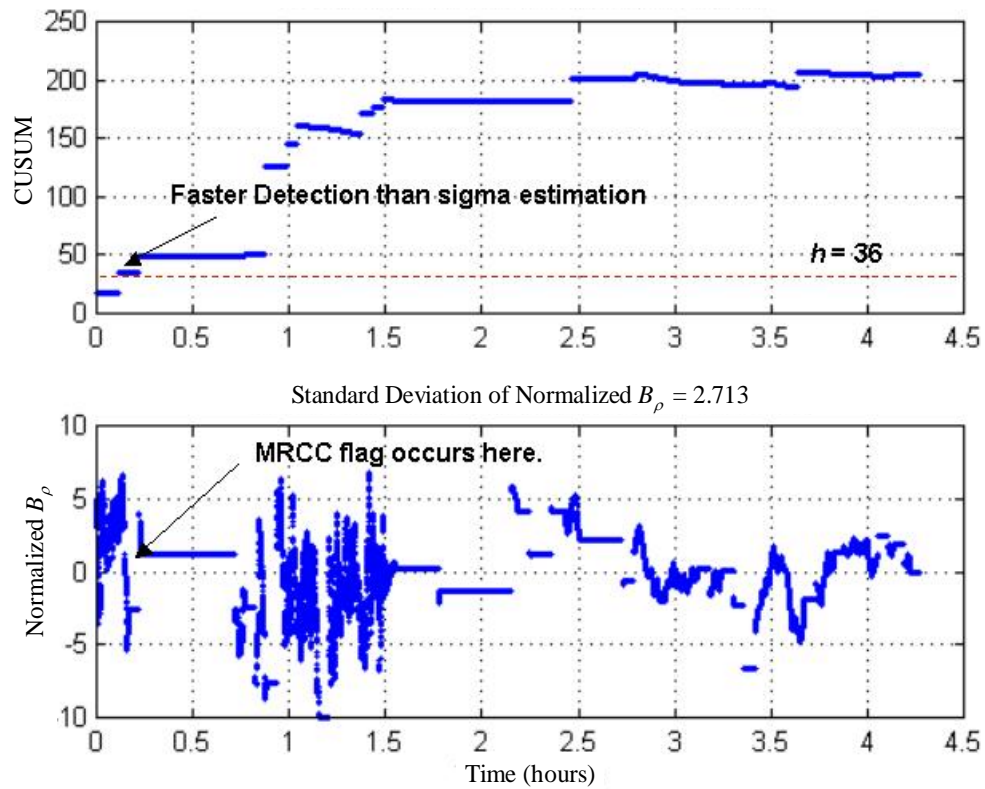


Figure 3.10: FIR CUSUM Results from IMT Failure Test

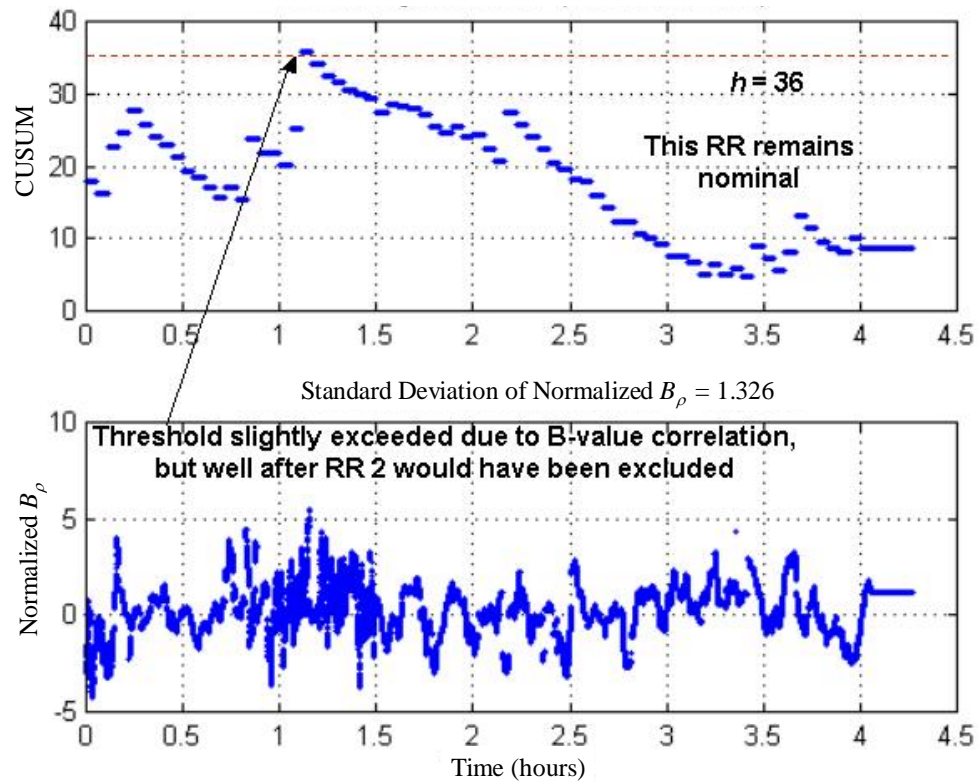


Figure 3.11: FIR CUSUM Results of Nominal RR from IMT Failure Test

We have also conducted the failure test of IMT sigma CUSUM monitoring with various L factors. With a moderate error factor of $L = 1.7$ on the same Channel (RR 2, SV 2), the FIR CUSUM exceeds the threshold at the 21st epoch (1.1 hours after error injection). Neither CUSUM nor MRCC flags appear on the non-failed receivers (RR 1 and RR 3). In the case of $L = 1.4$, no violation is detected. This fault is too small to be reliably detected during the 4-hour IMT run used here, as predicted by the theoretical result in Figure 3.4.

Overall, the CUSUM times-to-detect are much shorter (typically well under one hour) for large anomalies than those of the sigma estimation method, which requires waiting one hour until 18 independent samples are collected. Moreover, a FIR CUSUM with a “head start” at $3h/4$ would detect violations quicker than a FIR CUSUM with initialization at $h/2$ under the same failure conditions, but it has slightly higher continuity risk under fault-free conditions. We have also tested multiple CUSUM monitors tuned to target different values of the out-of-control sigma ($\sigma_1 = 1.7$ and 2.3), but these do not improve the time-to-detect measurably over a single CUSUM with $\sigma_1 = 2$. A subset of these failure tests has been rerun after integration with the EXM-II logic, and these tests confirmed that the IMT properly excludes measurements that triggered CUSUM and/or sigma estimation alerts.

3.4 MEAN MONITORING

As addressed in Section 3.2, real-time mean monitoring is required to detect possible protection-level violations due to non-zero means of the true pseudorange-correction errors. Similar to sigma monitoring, a common approach is to estimate real-time sample means. In Section 3.4.1, we present the mean estimation method. A mean CUSUM algorithm is described in Section 3.4.2. Lastly, Section 3.4.3 analyzes two methods with IMT test results.

3.4.1 MEAN ESTIMATION METHOD

The mean estimation method derives sample means from LGF B -values for each visible satellite. As with the sigma estimation method, normalized B -values from Equation (3-1) are the inputs to mean estimation. The estimates are compared to time-dependent thresholds, which are set based on the normal distribution of the sample mean, $\hat{\mu}_{B_{\rho_normal,m,n}}$.

The normal distribution as a function of the number of independent measurements, N , is:

$$\hat{\mu}_{B_{\rho_normal,m,n}}(k) \sim Normal\left(\mu_{B_{\rho_normal,m,n}}(k), \frac{\sigma_{B_{\rho_normal,m,n}}(k)}{\sqrt{N(k)}}\right) \quad (3-14)$$

Note that using the B -values as inputs to mean monitoring – both for the estimation and CUSUM methods – limits the observability of non-zero means to cases where mean violations occur on only one reference receiver. A common-mode failure that causes the same non-zero mean to occur on all three receivers is not observable from B -values and must be made extremely improbable to meet the LGF integrity allocation to multiple-receiver failures [51].

3.4.2 MEAN CUMULATIVE SUM METHOD

CUSUM applied to mean monitoring is essentially the same as CUSUM for sigma monitoring. Thus, this section mainly explains the differences to be made with the purpose of adapting the CUSUM for mean monitoring.

3.4.2.1 ALGORITHM

The input Y for the mean CUSUM method is the normalized B_ρ , whereas the square of the normalized B_ρ is the input for the sigma CUSUM method:

$$Y_{m,n}(N) = \frac{B_{\rho,m,n}(N) - \mu_{B_\rho,n}(N)}{\sigma_{B_\rho,n}(N)} \quad (3-15)$$

We choose a windowing parameter (k) based on a target out-of-control mean (μ_1), which represents “failed” performance (see Appendix B for derivation):

$$k_\mu^+ = \frac{\mu_0 + \mu_1}{2} \quad (3-16)$$

where the mean in a nominal case is defined as μ_0 , and referred to as the in-control mean. Note that, unlike sigma violations, threatening mean violations can be either positive or negative; thus two parallel CUSUMs (C^+ and C^-) are needed for each measurement so that violations in either direction will be detected.

3.4.2.2 THEORETICAL ANALYSIS

The CUSUM threshold (h) is found by numerical search to match the desired average run length (ARL) and k value, which is derived in Equation (3-16) given the target out-of-control mean (μ_1). As a function of μ_1 , Figure 3.12 shows the resulting thresholds to achieve an ARL of 10^7 independent epochs for the mean CUSUM monitor on the left and the resulting failure-state ARLs for the zero-start ($H^+ = 0$) and FIR ($H^+ = h/2$) CUSUMs on the right. The out-of-control ARL for the FIR CUSUM is significantly better than the zero-start CUSUM; and as with the sigma CUSUM method, it is possible to increase H^+ beyond $h/2$ to decrease detection time further with little nominal false alarm penalty. Again, the nominal thresholds for these two CUSUMs are essentially the same due to the fact that very large thresholds are needed to achieve an ARL of 10^7 under nominal conditions.




Figure 3.12: Thresholds and Failure-State ARLs for Mean CUSUM Monitor

3.4.3 IMT TEST RESULTS

As with sigma monitoring, both the estimation and CUSUM algorithms are implemented in the Stanford IMT for mean monitoring. This section discusses their test results under both nominal and failure conditions.

3.4.3.1 NOMINAL TESTING

Figure 3.13 shows the results of applying the mean estimation algorithm to the IMT under nominal conditions. The mean estimates of a single satellite (SV 2) on all three IMT reference receivers are shown here. The dark (blue) curves show the mean estimates of the normalized B -values, and the light (green) curves show the detection thresholds. For the same reason explained in sigma monitoring, mean monitoring starts after 6 independent B -values have been collected, or after 20 minutes. Note that the mean estimates stay well below the detection thresholds – which get smaller over time as the number of independent samples increases – and converge over time toward zero. We restart the mean monitor

when there is no B -value (if the GPS satellite sets and its elevation angle drops too low, then the receiver loses lock and no B -value is generated).

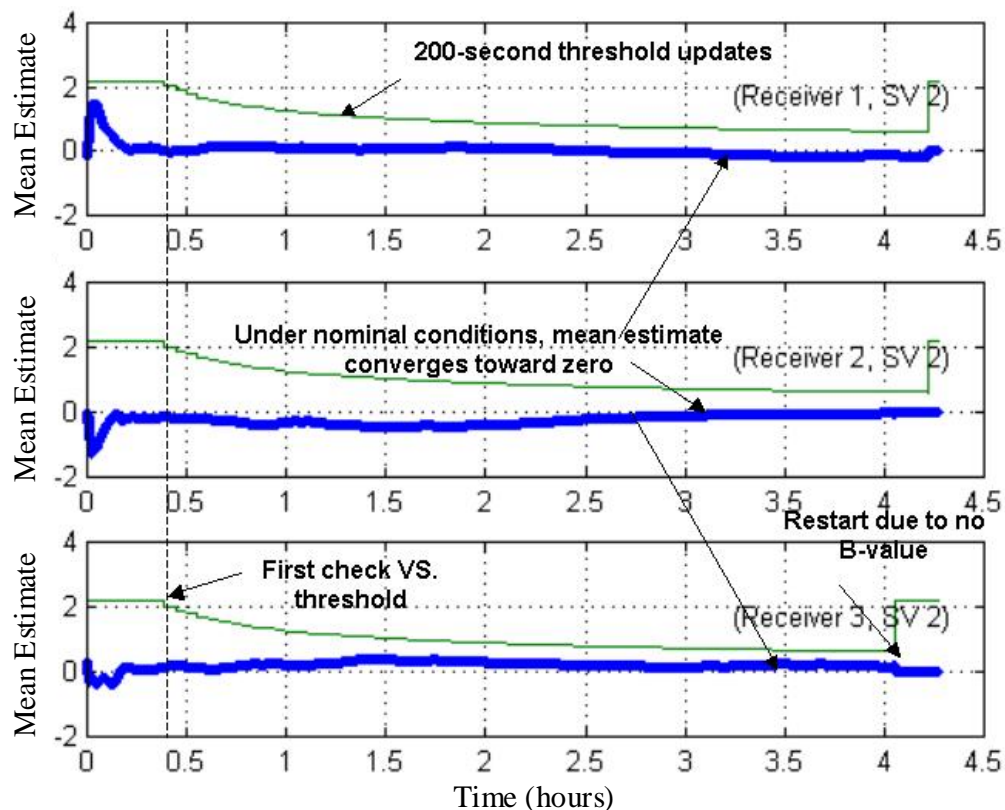


Figure 3.13: Mean Estimation Results from IMT Nominal Data (the thick blue curves show the mean estimates and the light green curves show the detection thresholds)

The mean FIR CUSUM has also been tested with the same IMT data shown in the previous plot. The top plot in Figure 3.14 presents the FIR CUSUM result for Channel (RR 1, SV 2), and the lower plot shows the normalized B -values that fed the CUSUM. The Mean CUSUM method is tuned to an out-of-control mean, $\mu_1 = 0.4$, which gives a windowing factor, $k = 0.2$. Based on the desired average run length (ARL) and k value, we determine the Mean CUSUM threshold ($h = 32.85$) using numerical search, since we can consider the CUSUM as a Markov Chain as described in Section 3.3.2.1. Again, the CUSUM is updated every 200 seconds, which makes each update statistically independent. We initialize the CUSUM at $h/2 = 16.4$ and reset it there every time the CUSUM falls below zero. The normalized B_ρ is usually below k ; thus the CUSUM slowly falls toward zero and is then

reset. We see the same patterns on the negative CUSUM, which we skip here. Again, we get very similar results for the other satellites included in this IMT dataset, and no flags are generated by either the mean estimation or mean CUSUM methods.




Figure 3.14: Mean FIR CUSUM Results from IMT Nominal Data

3.4.3.2 FAILURE TESTING

Both the mean estimation and mean CUSUM methods have been tested under failure conditions to verify the capability of mean monitoring to detect threatening anomalies. In order to induce mean violations, we insert controlled bias errors into stored nominal receiver packets previously collected through the IMT antennas. The bias added to ρ_{raw} at each epoch is pre-selected to be L times the nominal sigma of the error in ρ_{raw} . Then we put this modified ρ_{raw} back into the IMT in post-processing mode.

We have tested IMT mean monitoring with three different non-zero mean values ($L = 0.4$, 0.8 , and 1.2). Unlike the failure test of sigma monitoring, EXM-II is active with the mean monitors such that flagged measurements are removed. The results of the mean estimation method are presented in Figure 3.15. The mean of pseudorange errors on Channel (RR 2, SV 2) is increased to 0.8 times the nominal error. The lower plot shows that the mean estimate of the normalized B -values (the dark blue line) on Channel (RR 2, SV 2) is reset to zero at 2.72 hours. This is due to the fact that the mean FIR CUSUM ($C = 33.17$) exceeded its threshold ($h = 32.85$) at 2.72 hours, as shown in the lower plot of Figure 3.16. Since the EXM logic is active, the flagged Channel (RR 2, SV 2) is excluded, and the mean estimate and CUSUM (and all other measurements of that channel) are reinitialized at the same time. The top plots in Figures 3.15 and 3.16 show that Channel (RR 1, SV 2) is not affected much by the injected errors on Channel (RR 2, SV 2); thus the mean estimate and CUSUM of Channel (RR 1, SV 2) remain nominal. Since the bias is in the positive direction, no flags are seen in the negative CUSUM.

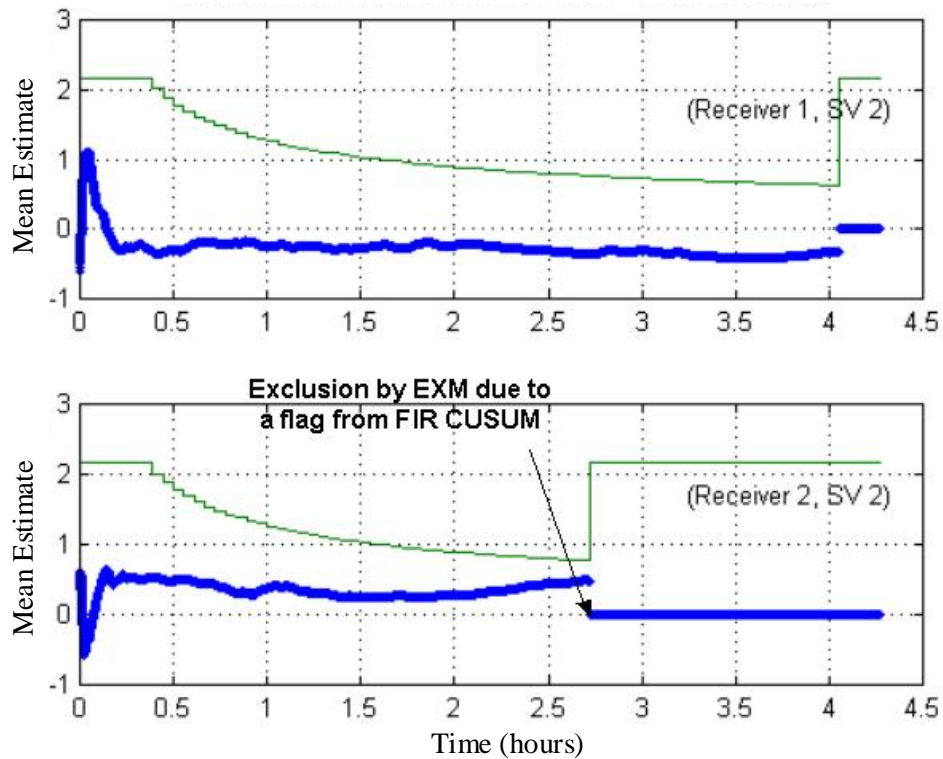


Figure 3.15: Mean Estimation Results from IMT Failure Test with $L=0.8$ Injected on Channel (RR2, SV 2) (the thick blue curves show the mean estimates and the light green curves show the detection thresholds)

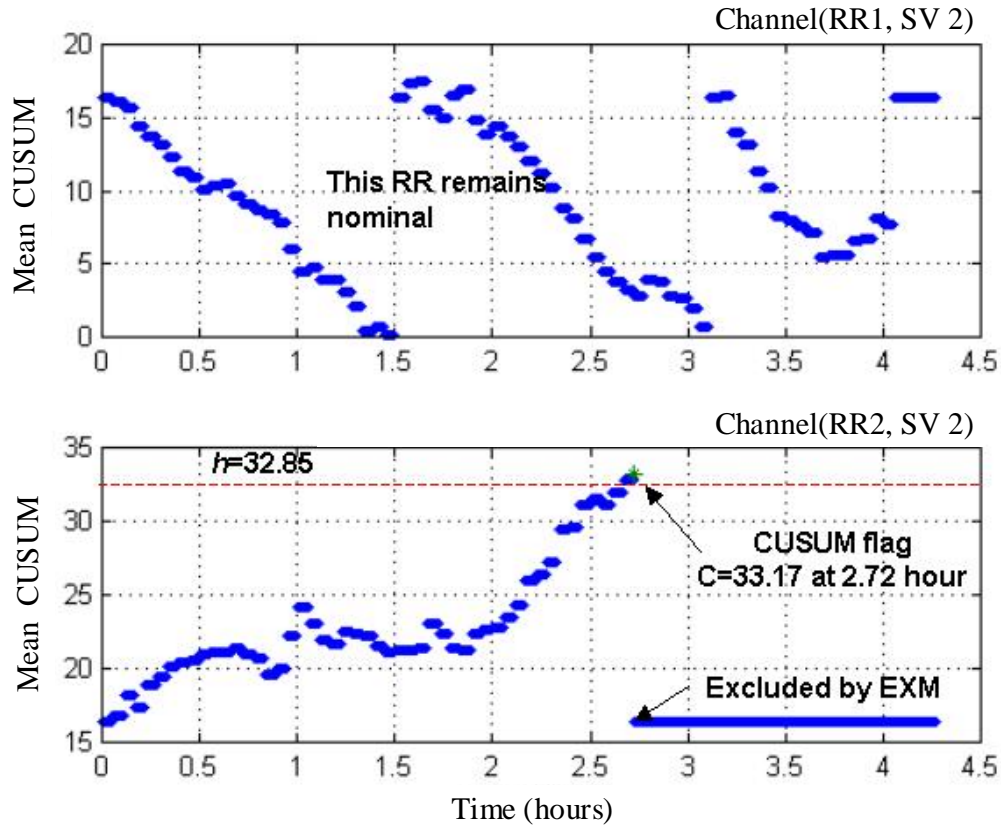


Figure 3.16: Mean FIR CUSUM Results from IMT Failure Test with $L=0.8$ Injected on Channel (RR2, SV 2)

When L is lowered to 0.4 on the same Channel (RR 2, SV 2), no flag is generated by either mean estimation or CUSUM methods, which is predicted by the theoretical performance analysis in Figure 3.12. With a severe mean error with $L = 1.2$ on the same Channel (RR 2, SV 2), the positive FIR CUSUM exceeds its threshold at 1.65 hours after the fault was injected. Again the FIR CUSUM method gives a faster detection than the mean estimation method in this case, and thus the mean estimate is reset to zero at 1.65 hours. Neither CUSUM nor mean flags appear on the non-failed receivers (RR 1 and RR 3).

3.5 COMPARISON OF ESTIMATION AND CUSUM RESULTS

So far, we have analyzed and tested two sigma-mean monitoring algorithms: the estimation and CUSUM methods. In this section, we compare and summarize the performance of these methods.

Figure 3.17 compares the average times-to-detect for the sigma estimation monitor and three CUSUMs based on potential sigma violations as a function of the out-of-control sigma (these are combined results from Figures 3.2 and 3.4). Here we can see that the sigma estimation method is best for $\sigma_1 < 1.45$, but the FIR CUSUM ($H^+ = h/2$) is superior for higher sigma values, which is important because larger sigma violations lead to more severe integrity threats. As explained earlier, the FIR (head-start) CUSUM achieves faster detection by initializing the CUSUM to a non-zero value closer to the threshold every time the CUSUM resets while having little continuity penalty (see Section 3.3.2.2). The CUSUM has an additional advantage: CUSUM monitoring begins immediately, while sigma estimation requires that 18 independent epochs (1 hour) be observed before threshold checks can begin (Prior to one hour, the sigma estimate is too unreliable to be compared to a chi-square-based threshold). Thus, we can improve sigma monitoring further, especially when σ_1 is greater than 2, by adding CUSUM algorithms. Overall, the IMT is sufficient to detect any threatening size of sigma violation using sigma estimation and sigma CUSUM methods.




Figure 3.17: Time-to-Alert for Sigma CUSUM and Sigma Estimation Monitors

As explained earlier, modeling the CUSUM as a Markov chain allows us to determine the probability of exceeding the threshold h under failure conditions at any future epoch. Thus we can determine how soon the failure is detected with a probability of 0.999, or, conversely, a missed-detection probability ($P_{MD} \leq 0.001$). Figure 3.18 compares the times-to-detect with $P_{MD} \leq 0.001$ for three CUSUMs and the mean estimation method based on potential non-zero mean violations as a function of the out-of-control mean. The results show that FIR CUSUM methods are superior to the mean estimation method in detecting any mean violations. Note that the zero-start CUSUM is slightly worse than mean estimation, but the $h/2$ FIR CUSUM is significantly better for all μ_1 . This can be further improved – with slight loss of continuity under fault-free conditions – by increasing H^+ above $h/2$. Both CUSUM and mean estimation methods detect larger mean violations

almost simultaneously, though the FIR CUSUM with a “head start” of $3h/4$ achieves faster detection.




Figure 3.18: Time-to-Alert with $P_{MD} < 0.001$ for Mean Estimation and Mean CUSUM Monitor Performance

The CUSUM method is often compared to the Exponentially Weighted Moving Average (EWMA) method, which may also be called the geometric moving average [23]. The EWMA is a method that can be nearly as fast as the CUSUM in detecting step changes. Another attraction of this method is that its value at any time gives an immediate estimate of the current process mean, something that the CUSUM provides only after measuring the slope of the latest segment and adding in the reference value. On the other hand, the EWMA is not as fast as the CUSUM in detecting step shifts and is not as good as the CUSUM for estimating when the step changes or shifts occurred.

3.6 CONCLUSION

This chapter has summarized the direct estimation algorithms and analyzed the application of CUSUM algorithms for mean and sigma monitoring. The Sigma-Mean monitor has been successfully implemented in the Stanford LAAS Integrity Monitor Testbed and tested under both nominal and failure conditions. We have seen that the test results from both estimation and CUSUM methods generally agree with analytical predictions.

FIR CUSUMs are superior to mean and sigma estimation in most cases, although sigma estimation should still be used to detect relatively small sigma violations. Further improvement of FIR CUSUM performance is possible with higher “head-start” H^+ . However this causes the fault-free alarm rate to increase. Given this trade-off, the optimum head start is yet to be determined, but the $h/2$ head start implemented in the IMT is a reasonable compromise; in fact, the choice $H^+ = h/2$ is recommended for general use by Lucas and Crosier [53]. While CUSUM *mean* times-to-detect are well under one hour for large violations, the time-to-detection with $P_{MD} \leq 0.001$ is somewhat longer.

Since small mean and sigma violations are not detectable, some inflation of the nominal sigma (roughly a factor of 1.4 – 1.5) is needed to provide margin for sigma-mean monitoring so that anomalies too small to be detected are not hazardous to users. With this amount of inflation, the performance achieved by the CUSUMs appears good enough to meet the LGF specification requirement [15, 51]. In other words, “non-minimal-risk” anomalies – which cause the computed protection levels to be well below reality – can be detected within one to three hours with 99.9% reliability, while the mean detection times will be typically under one hour. If faster times-to-detect are desired, additional sigma inflation could be implemented, but “diminishing returns” applies above an inflation factor of 2.0 because the CUSUM time-to-detect does not improve much further. Further analysis will come when we investigate sigma inflation in Chapter 4.

The Sigma-Mean monitor has been smoothly integrated with the EXM logic within the Stanford IMT and clearly accomplishes removal of single-channel anomalies, allowing other nominal measurements to continue to be used. Overall, this monitor is sufficient to

detect anomalies that cause the true sigma to exceed the broadcast sigma or the true mean to become non-zero during LAAS operations; thus, it helps to provide navigation integrity to airborne users.

Chapter 4

Sigma Inflation

4.1 INTRODUCTION

In Chapter 3, the Sigma-Mean monitor was developed to guarantee the safety of the broadcast σ_{pr_gnd} when measurements are corrupted by unexpected anomalies. However, monitoring alone is not sufficient to ensure that a zero-mean Gaussian distribution with the broadcast sigma “overbounds” the true (unknown) error distribution. Even in nominal conditions, the true sigma may exceed the broadcast σ_{pr_gnd} due to the uncertainty of the true error distribution. The main sources of this uncertainty are mean and sigma estimation error during site installation and non-stationary error distributions caused by environmental changes that affect multipath. In order to deal with this statistical uncertainty, we need to broadcast an inflated σ_{pr_gnd} such that the broadcast distribution overbounds all reasonable error distributions out to the probabilities assumed in the computation of the protection levels (PLs).

A great deal of prior work has been done regarding sigma inflation that accounts for each individual cause of the uncertainty [1, 18]. However, an inflation factor that copes with all of these uncertainties at once has not yet been investigated. This chapter will introduce a comprehensive method of determining the inflation factor to insure that the zero-mean

Gaussian distribution implied by the broadcast sigma values overbounds the tails of the true distribution (which is possibly non-Gaussian and non-zero-mean). We will continue by deriving the inflation factor for the broadcast σ_{pr_gnd} with this method. In the last section we will evaluate the induced inflation factor by computing PLs and quantitatively appraise LAAS navigation integrity with these position bounds.

4.2 SIGMA INFLATION FACTOR DETERMINATION METHOD

As addressed earlier, sigma inflation is needed to provide the safety margin on protection bounds, since the error distribution of the differentially corrected pseudorange measurements is subject to the following sources of uncertainty: finite sample size, process mixing, and the limitation of the sigma monitor. In Section 4.2.1, we consider the effect of sigma estimation error due to the limited number of samples. As mentioned in Chapter 1, a basic assumption in PL computing is that correction errors are zero-mean Gaussian distributed. However, in practice, the tails of the true error distribution may not be exactly Gaussian due to time-varying environmental conditions. In addition, even though we assume a stationary condition, mixing of Gaussian errors with different sigmas may cause the fattened tails. In Section 4.2.2, we deal with mixing of time-varying errors – such as ground reflection multipath – and mixing of different Gaussian distributions. A buffer parameter is derived for the non-Gaussian tails. Then we review the performance of sigma monitoring and provide a factor to overcome its limitations in Section 4.2.3. Lastly we present a way to combine all factors and determine the inflation factor for the broadcast

$$\sigma_{pr_gnd} \cdot$$

4.2.1 FINITE SAMPLE SIZE

When we determine the broadcast sigma of the ground facility error, we account for specific environmental conditions – antenna siting, gain patterns, and system configuration – of each LGF site. Even though these conditions are factored very accurately into sigma estimation and the environment is assumed to be stationary, the estimated sigma may have a statistical noise due to finite sample size. Previous research on this subject has been done by Pervan and Sayim in [1]. They investigated the sensitivity of integrity risk to statistical

uncertainties to which the correction error standard deviation and error correlation between multiple reference receivers are susceptible. Based on their work, the minimum acceptable buffer parameter for the broadcast σ_{pr_gnd} is approximately 1.2 (we skip the details in this thesis).

4.2.2 PROCESS MIXING

As noted above, the true LGF error distribution may change with time, as the environment condition varies. In addition, mixing of the time-varying errors – ground reflection multipath and systematic receiver/antenna errors – makes the characteristics of the error distribution to be complex. Even if stationary Gaussian error distributions are assumed, some degree of the mixture problem is expected. The standard deviation of the true error distribution varies as a function of the elevation angle of each satellite. If pseudorange correction errors are normalized by the theoretical sigma which depends on the ranging source elevation angle but which is not perfect, this normalization cause to mix Gaussian distributions with different sigmas. The process mixing may result in non-Gaussian tails.

Figure 4.1 shows the distribution of the LGF B -values collected by the Stanford IMT for a period of five hours. Recall that the B -values represent the correction errors of pseudorange measurements. Here we can clearly see that the correction error distribution (the dotted curve) has non-Gaussian tails (note that the scale of its vertical axis is logarithmic). Thus we should inflate the nominal 1σ Gaussian distribution – shown as the dashed curve – to overbound the error distribution with non-Gaussian tails. However, this error distribution modeled with experimental data is not sufficient to represent the true error distribution. In other words, the reliable estimation of the tail probabilities is impossible because their small magnitude (on the order of 10^{-10}) requires a huge sample size (greater than 10^{10}) that cannot be collected in a realistic time frame. Hence the limited number of samples makes a theoretical model necessary for estimating the error distribution. We use the following Gaussian-Mixture distribution as the theoretical model

$$\begin{aligned} f_{GM} &= (1 - \varepsilon) * N(\mu_0, \sigma_0) + \varepsilon * N(\mu_1, \sigma_1); \\ \varepsilon &= 0.15, \mu_0 = \mu_1 = 0, \sigma_0 = 0.75, \sigma_1 = 1.82 \end{aligned} \tag{4-1}$$

where $N(\mu, \sigma)$ is a normal distribution with mean, μ , and sigma, σ . Thus this function is the weighted sum of two normal distributions with one nominal sigma and one relatively large sigma. This model (the solid green curve) shown in Figure 4.1 well characterizes the actual distribution of Gaussian-core and non-Gaussian tails. Again, the nominal sigma should be inflated in order to cover the non-Gaussian tails of the actual distribution with a normal distribution. For CAT I approaches, the tails need to be overbounded so that the probability of the error exceeding protection levels is less than or equal to 6×10^{-9} under the hypothesis of fault-free conditions (H_0); for CAT II/III approaches the required probability is 1.2×10^{-10} . To meet the integrity requirement, we need to inflate the sigma by a factor of 2.32 or greater. We can see that the 2.32σ Gaussian distribution (the solid red curve) well overbounds the theoretical model (the solid green curve). As a result the minimum tolerable inflation factor is 2.32. Note that test statistics highly depend on system configurations; thus this analysis should be conducted at each LGF site.




Figure 4.1: Probability Density Function of the Normalized B -values

4.2.3 LIMITATION OF SIGMA MONITORS

The possibility of sigma violations exists because of not only the nominal sigma uncertainty but also unexpected anomalies. As explained in Chapter 3, the sigma monitor is designed to provide necessary integrity in the event that the true sigma significantly exceeds the broadcast sigma [50]. However, the current sigma monitor has a limitation on mean time-to-detect which must be overcome with an additional inflation factor. First we derive the additional parameter assuming that the error distribution is Gaussian in Section 4.2.3.1. Second we assume a specific non-Gaussian error distribution and then derive the buffering parameter in Section 4.2.3.2.

4.2.3.1 GAUSSIAN ASSUMPTION ON ERROR MODEL




Figure 4.2: Failure-State Average Run Lengths for CUSUM and Sigma Estimation Monitors

Let us review the performance of the sigma monitor implemented in the IMT. Figure 4.2 shows the Average Run Lengths (ARL) for different sigma monitoring methods to detect certain failure-states – out-of-control sigmas (σ_1) – given the condition that the error distribution is Gaussian. We now turn our attention to the LGF requirements specified in [15] and reexamine the capability of this monitor. Based on the time-to-alert requirements, if the actual integrity risk is greater than the total allocation but the resulting risk increase is minimal (i.e., is no greater than one order of magnitude), it is defined as “minimal-risk-increase.” Since degraded performance due to such sigma failure is minimal, we need not detect it immediately but instead within a day. Note that if sigma failure causes “non-minimal-risk-increase” (i.e., the integrity risk is increased by more than one order of magnitude from the total allocation), it should be detected within an hour. The limitation of the sigma monitor is now defined: Assuming that we can continuously collect data in one satellite pass for five hours on average, the minimum out-of-control sigma detectable within a day is 1.41. Since out-of-control sigmas (σ_1) greater than the inflation factor ($f_{inflation}$) are categorized as “minimal-risk-increase” (i.e. the actual sigma exceeds the broadcast sigma), by definition:

$$\sigma_1 = \frac{\sigma_{Actual}}{\sigma_{Theoretical}}; \quad f_{inflation} = \frac{\sigma_{Broadcast}}{\sigma_{Theoretical}} \quad (4-2)$$

if the inflation factor is less than 1.41, sigma violations with minimal risk – between the inflation factor and 1.41 – cannot be alarmed within a day. Accordingly, in order to meet the LGF requirements, the inflation factor should be at least 1.41.

4.2.3.2 NON-GAUSSIAN ASSUMPTION ON ERROR MODEL

As we pointed out in Section 4.2.2, the error distribution may not be precisely Gaussian. Thus, we also need to consider the restriction of the sigma monitor given the assumption that the errors are non-Gaussian distributed. Results corresponding to Figure 4.2 are generated using the non-Gaussian model described in Equation (4-1). This model is an example to represent the actual distribution. From this distribution we collect 90 independent samples to compute a sample standard deviation. Note that 90 is the number of

independent samples that we can collect continuously in five hours, since 18 independent samples correspond to one hour. Repeating this process randomly, we then generate 100000 sample sigmas. As a result, the probability density function (PDF) of sample sigma is shown in Figure 4.3. Based on the specified fault-free alarm rate, 10^{-7} (a sub-allocation of Category I continuity risk allowed per 15-second interval [15]), the minimum out-of-control sigma detectable within five hours is now 1.58. Consequently, to protect this particular non-Gaussian error model, the inflation factor should be at least 1.58 for the same reason explained in Section 4.2.3.1.




Figure 4.3: Probability Density Function of Sample Standard Deviation

4.2.4 TOTAL INFLATION FACTOR

So far we have investigated three sources of sigma uncertainty and derived a buffering parameter for each source. The final step is to determine the total inflation factor for the broadcast σ_{pr_gnd} considering all conditions discussed in Sections 4.2.1, 4.2.2 and 4.2.3:

- The theoretical (or pre-estimated) sigma is to be inflated by a factor of 1.2 to account for finite sample size (Section 4.2.1).
- The buffering parameter to overbound the tails of the non-Gaussian distribution derived from IMT data is 2.32 (Section 4.2.2).
- The inflation factor should be at least 1.58 to overcome limitations of the existing sigma monitor (Section 4.2.3).

In Figure 4.4, we present the inflation factor determination method for the broadcast σ_{pr_gnd} . Since the conditions described in Sections 4.2.1 and 4.2.2 are independent, we multiply the two parameters (1.2 and 2.32). The resulting factor (2.78) already exceeds what is required by the sigma monitor (which is 1.58) satisfying the third condition. Thus, the total inflation factor is 2.78.

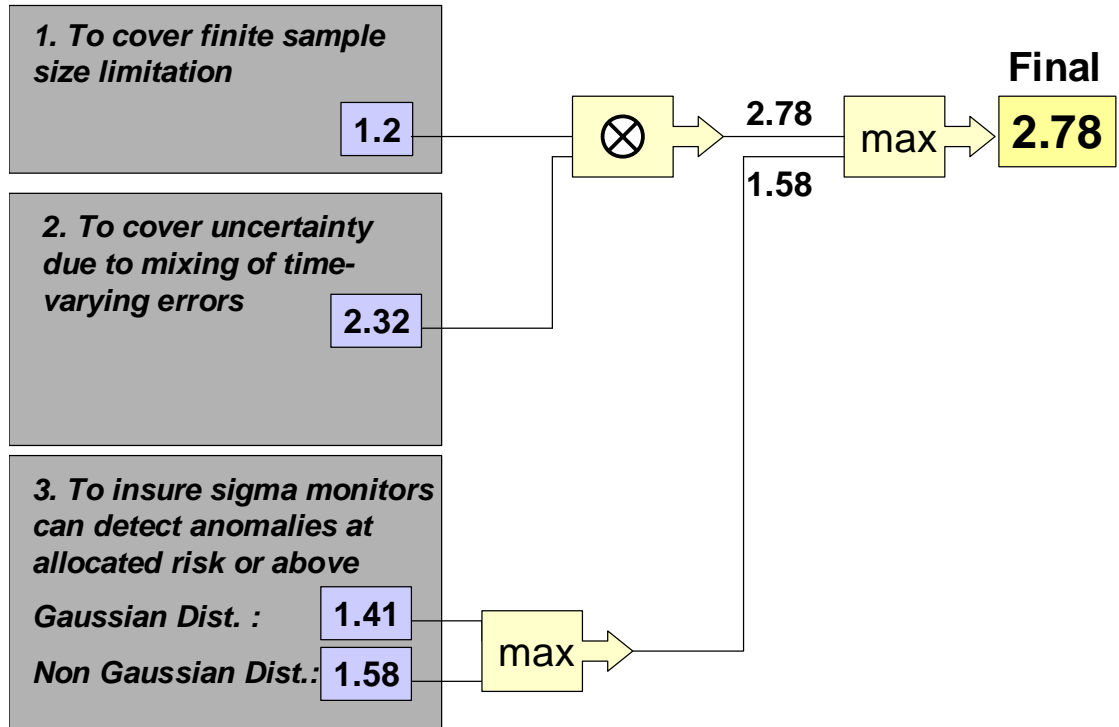


Figure 4.4: Inflation Factor Determination Method for Broadcast σ_{pr_gnd}

4.3 PERFORMANCE ANALYSIS

For a safety-critical system like LAAS whose main purpose is to provide integrity to users, a key requirement is the capability to provide a hard bound on the position error. For this reason, the fundamental requirement on sigma inflation is that the computed error bounds with the inflated sigma are really bounding the possible position error. In this section, we compute protection levels by applying the total inflation factor determined in Section 4.2.4 to the broadcast σ_{pr_gnd} and evaluate the performance of the LGF with sigma inflation.

4.3.1 STANFORD LAAS PERFORMANCE TEST-BED

In order to test the ability to meet the LAAS precision approach requirements, we have installed a static “pseudo-user” receiver in addition to the existing Stanford IMT architecture (see Chapter 2). In Figure 4.5 we show the configuration of the three IMT antennas on the Stanford HEPL laboratory rooftop as well as the “pseudo-user” antenna/receiver on top of the nearby parking structure. The NovAtel Pinwheel antenna (“pseudo-user” antenna) and the center of the IMT are approximately 230 meters apart. The NovAtel OEM-4 receiver connected to the “pseudo-user” antenna collects pseudorange measurements, carrier-phase measurements, and navigation messages of GPS satellites (“pseudo-user” and IMT receivers are set up to collect measurements simultaneously). These measurements are post-processed in a single computer where the processing algorithms have been developed and tested.

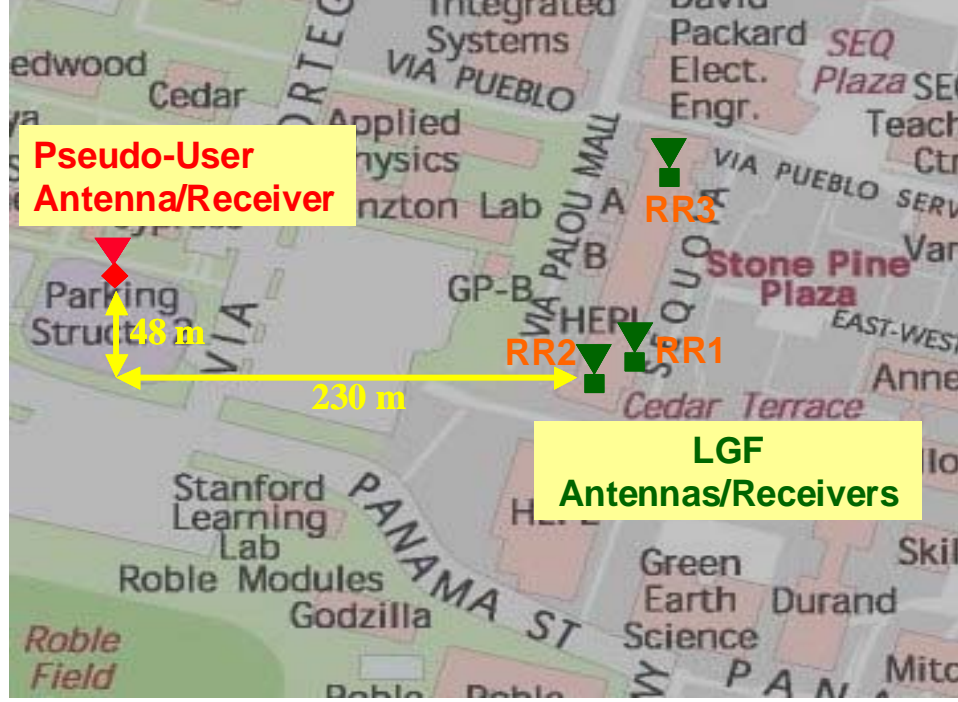


Figure 4.5: Stanford LAAS Performance Test-bed Hardware Configuration

4.3.2 PERFORMANCE TEST RESULTS

In order to evaluate IMT system performance with sigma inflation, we first compute position errors, which are obtained by comparing the surveyed location of the pseudo-user's antenna to position solutions. Pseudo-user position solutions are computed in the manner required of the LAAS airborne receivers to mirror LAAS aircraft operations to the degree possible (the detailed algorithm is specified in the RTCA LAAS *Minimum Operational Performance Standards* (MOPS)[14] and in Section 5.2.2). In this analysis, Accuracy Designator C (AD-C) is applied to the pseudorange error model [54], as the upcoming hardware installation for CAT I will most likely be similar to those of CAT II/III. Second, we compute protection levels, through which the final quantitative appraisal of the navigation performance is realized. The Vertical Protection Level (VPL) under the hypothesis of fault-free conditions (H0) is (from Equations (1-11), (1-12) and (1-13)):

$$VPL_{H0} = K_{ffmd} \sqrt{\sum_{n=1}^N S_{vertical,n}^2 \cdot \left[\sigma_{air,n}^2 + \sigma_{tropo,n}^2 + \sigma_{iono,n}^2 + \left(f_{inflation} \cdot \sigma_{pr_gnd,n} \right)^2 \right]} \quad (4-3)$$

where $f_{inflation}$ is the inflation factor and K_{ffmd} is a specified multiplier that determines the probability of fault-free missed detections [14]. Again, since the vertical direction is the most stringent one and errors in this direction are generally worse than those in the lateral direction, we only pay attention to the vertical direction.

Figure 4.6 shows the results of applying sigma inflation to the broadcast σ_{pr_gnd} . The pseudo-user's Vertical Position Errors (VPEs), the Vertical Protection Levels (VPLs), and the Vertical Alert Limit (VAL) for CAT I are plotted in this figure. First, we can see that the VPEs are well within ± 2 meters; thus the accuracy requirements for LAAS precision approaches are met (see Table 2.1). The next question is how well system integrity is guaranteed; i.e. whether the computed error bounds are really bounding the position errors. The results show that the VPEs are well below the VPLs without sigma inflation ($f_{inflation} = 1$). However, this is not enough for us to be confident that the error bounds are always an upper bound of the position errors; note that the requirement is $\text{Prob}(error > PL) \leq 10^{-7}$. Thus the VPLs with sigma inflation ($f_{inflation} = 2.78$) provide better safety margin on integrity than those without sigma inflation. Lastly, the VPLs never exceed a VAL of 10 meters. As a consequence, the continuity and availability requirements are also met for CAT I precision approaches in this period of time.




Figure 4.6: Stanford LAAS Pseudo-User Performance

4.4 CONCLUSION

This chapter provides a comprehensive method to determine the sigma inflation factor. The derived inflation factor includes partial parameters for all sources of the sigma uncertainty and for the limitation of the current sigma monitor. Pseudo-user performance test has demonstrated that navigation integrity can be improved by applying sigma inflation.

Chapter 5

Position Domain Monitoring

5.1 INTRODUCTION

Chapters 3 and 4 found a way to insure that the zero-mean Gaussian distribution based on the broadcast σ_{pr_gnd} overbounds the true distribution, which may be non-Gaussian and non-zero-mean. This is done by broadcasting an inflated σ_{pr_gnd} and detecting violations of the resulting overbound using the sigma-mean monitor. However, the LAAS sigma overbounding issue persists because high levels of sigma inflation cannot be tolerated for CAT II/III precision approaches. In Figure 5.1 we review the performance of the Stanford LGF prototype shown in Chapter 4. Here we can see that the VPLs with sigma inflation are too conservative to meet CAT II/III requirements (when the VPLs exceed the VAL, LAAS is not available). While the system promises to support CAT I operations, significant technical challenges are encountered in supporting CAT II/III operations on account of the tightened VAL of 5.3 meters and similarly high availability requirements (0.999 or higher, depending on the airport). For this reason, we introduce Position Domain Monitoring (PDM) and investigate how PDM may be used to improve system availability by reducing the inflation factor for the standard deviation of pseudorange correction errors.




Figure 5.1: Stanford LAAS Pseudo-User Performance

This chapter starts by introducing the basic function of PDM and a proposed algorithm in Section 5.2. Once we have gotten position solutions from this algorithm, we discuss the characteristics of error distributions in the position domain. Then we demonstrate that PDM supports a smaller σ_{pr_gnd} inflation factor needed for CAT II/III operations in Section 5.3. For this purpose the inflation factor determination method (developed in Chapter 4) is used to derive the new inflation factor in the position domain. We will continue by field-testing the improved performance of the Stanford IMT in Section 5.4, demonstrating that PDM helps meet the availability requirement of CAT II/III once it is integrated in the LGF architecture. Furthermore, in Section 5.5 we examine different methodologies to enhance system integrity and continuity using PDM outputs.

5.2 POSITION DOMAIN MONITORING (PDM)

Position domain monitoring (PDM) was introduced in the mid-1990 by Markin and Shively in [20]. PDM performs the integrity check by monitoring position solutions computed based on LGF differential corrections. This method was considered an alternative to range domain monitoring (RDM), which monitors each pseudorange measurement individually and approves each satellite for the aircraft to use. Relying on only PDM turned out to be impractical due to limited data-link capacity and flexibility because it was thought that PDM required generating every possible combination of satellites that may be used to compute the position solution and approving usable sets on a combination-by-combination basis. Thus, the current LGF is based on RDM, as shown in Chapter 2.

Given that an enhanced LGF architecture is required to meet CAT II/III requirements, the PDM concept has been reconsidered in [21, 55]. In this concept, PDM collects measurements with a remote receiver and derives position solutions by applying LGF corrections to all visible satellites approved by the LGF and all possible subsets of satellites. The position solutions are then compared to the known (surveyed) location of the PDM antenna. By performing the integrity check directly in the position domain, this method avoids the conservatism that prevails in the range domain (RDM requires a transformation from an estimate of the pseudorange correction errors to a bound on user position error, and the resulting position error estimate may be conservative). Thus, an inflation factor for the broadcast σ_{pr_gnd} can be reduced in the position domain, as will be discussed further in Section 5.3. Consequently, the use of the position domain concept will provide an availability advantage on the LAAS.

PDM could play an important role in providing extra integrity. It would augment sigma-mean monitoring (discussed in Chapter 3) and help detect when the true sigma exceeds the broadcast sigma. The possibility of sigma violation exists because of sigma anomalies (caused by man-made or natural system failures) and because of the limited number of independent samples (limited by the 200 second interval between independent B -values)

[50]. Thus, sigma monitoring cannot always detect when the broadcast σ_{pr_gnd} fails to bound the true sigma. Moreover B -values are sensitive only to multipath errors, while position errors are sensitive to all error sources, including residual tropospheric and ionospheric errors, and ephemeris errors. Therefore, this position-domain approach can improve upon the existing range-domain sigma monitoring.

In this section, we describe the Stanford PDM architecture and present the PDM algorithm to compute position error estimates and perform a safety check in the position domain. We also explain how to derive detection thresholds for PDM and show test results under both nominal and failure conditions.

5.2.1 PDM HARDWARE CONFIGURATION

A prototype of PDM has been implemented as a LAAS “pseudo-user” receiver augmentation to the Stanford IMT shown in Chapter 2. Figure 5.2 gives the configuration of the three IMT antennas on the Stanford HEPL laboratory rooftop and a PDM antenna on the Stanford Durand building. We use the existing Stanford WAAS Reference Station antenna for the PDM antenna, which is separated by approximately 145 meters from the IMT antennas. The NovAtel OEM-4 receiver connected to the PDM antenna collects pseudorange measurements, carrier-phase measurements, and navigation messages. The measurements are processed in a single computer where the PDM algorithm is developed and tested. To compare performance, one post-processing run is conducted with existing IMT measurements only, and a second run is executed using the combined IMT-PDM algorithms.

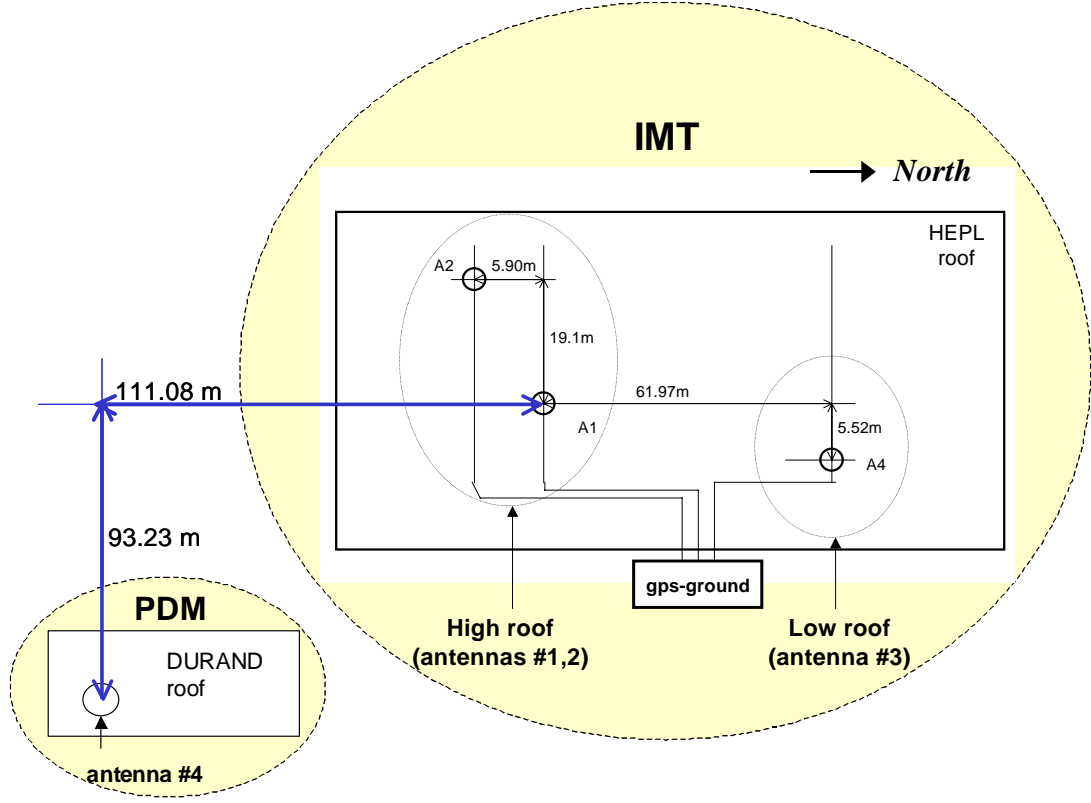


Figure 5.2: IMT-PDM Hardware Configuration

5.2.2 PDM ALGORITHM

PDM position solutions are computed using the approach required of LAAS airborne receivers – as specified in the LAAS *Minimum Operational Performance Standards* (MOPS) [14] – to emulate LAAS aircraft conditions as much as possible (the same method has been used to obtain pseudo-user position estimates and evaluate performance in Chapter 4). In order to reduce errors in raw pseudorange measurements, we first apply the following carrier-smoothing filter in the same manner executed in the IMT [14, 15, 31, 56]. The smoothed pseudorange for satellite n at epoch k is:

$$\rho_{s,n}(k) = \frac{1}{N_s} \rho_n(k) + \frac{N_s - 1}{N_s} (\rho_{s,n}(k-1) + \phi_n(k) - \phi_n(k-1)); \quad n = 1, 2, \dots, N \quad (5-1)$$

where ϕ is the carrier phase measurement, and N_s is equal to 200 since this filter uses a time constant, τ_s , of 100 seconds and a sampling interval, T_s , of 0.5 seconds.

$$N_s = \tau_s / T_s = 100 / 0.5 = 200 \quad (5-2)$$

Next we apply the set of LGF differential corrections to these carrier-smoothed code measurements [14]. The corrected pseudorange measurements are:

$$\rho_{c,n}(k) = \rho_{s,n}(k) - \rho_{corr,n}(k-1) - R_{\rho_{corr},n}(k) \cdot T_s + TC(k) + c \cdot \Delta t_n(k) \quad (5-3)$$

where ρ_{corr} and $R_{\rho_{corr}}$ are the pseudorange correction and the range rate correction from the IMT-approved message (see Equations (2-7) and (2-20)). TC is the tropospheric correction and is small enough to be neglected in this application (see Appendix A). The parameter c represents the vacuum speed of light, and Δt_n is the satellite clock correction computed using clock parameters in Sub-frame 1 of the GPS navigation message. Based on this set of differentially corrected measurements, we compute three-dimensional positions using a linearized, weighted least-squares estimation method. The linearized measurement model is (see Section 1.4):

$$\underbrace{\rho_{c,n}(k) - \rho_{0,n}(k)}_{\Delta y} = G \underbrace{\begin{bmatrix} \delta x(k) \\ \delta y(k) \\ \delta z(k) \\ \delta b(k) \end{bmatrix}}_{\Delta x} + \tilde{\epsilon} \quad (5-4)$$

where Δx is the four dimensional position/clock vector, Δy is an N dimensional vector containing the corrected pseudorange measurements minus the expected ranging values based on the location of the PDM antenna and satellites, $\tilde{\epsilon}$ is an N dimensional vector containing the errors in the corrected measurements (y), and G is the observation matrix consisting of N rows of line-of-sight vectors from each satellite to the PDM antenna, augmented by a “1” for the clock. Thus, the n^{th} row of G corresponds to the n^{th} satellite in

view and can be written in terms of the azimuth angle (Az_n) and the elevation angle (El_n).

This matrix is unitless and is defined as:

$$G = \begin{bmatrix} -\cos El_1(k) \cos Az_1(k) & -\cos El_1(k) \sin Az_1(k) & -\sin El_1(k) & 1 \\ -\cos El_2(k) \cos Az_2(k) & -\cos El_2(k) \sin Az_2(k) & -\sin El_2(k) & 1 \\ \vdots & \vdots & \vdots & \vdots \\ -\cos El_N(k) \cos Az_N(k) & -\cos El_N(k) \sin Az_N(k) & -\sin El_N(k) & 1 \end{bmatrix} \quad (5-5)$$

We find the weighted least-squares solutions for the estimate of the states by:

$$\Delta \hat{x} = S \cdot \Delta y; \quad S \equiv (G^T W G)^{-1} G^T W \quad (5-6)$$

where S is the weighted least square projection matrix, and the inverse of the least-squares weighting matrix is:

$$W^{-1} = \begin{bmatrix} \sigma_{PR,1}^2 & 0 & \cdots & 0 \\ 0 & \sigma_{PR,2}^2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \sigma_{PR,N}^2 \end{bmatrix} \quad (5-7)$$

Here, $\sigma_{PR,n}$ is the fault-free error model associated with satellite n :

$$\sigma_{PR,n}^2 = \sigma_{air,n}^2 + \sigma_{tropo,n}^2 + \sigma_{iono,n}^2 + \sigma_{pr_gnd,n}^2 \quad (5-8)$$

We describe the details of the first three terms in Appendix A (or see Section 1.4). For the fourth term, σ_{pr_gnd} , we apply the Ground Accuracy Designator C (GAD-C) model explained in Section 2.4.2. Although the purpose of PDM is to imitate aircraft operations, the PDM is still a ground-based system with ground-reflection multipath. Thus, we need to replace the airborne error sigma, σ_{air} , with the ground facility error sigma, σ_{pr_gnd} . Comparing the position solutions (\hat{x}) to the known location of the PDM antenna ($x_{surveyed}$), we have:

$$\chi_{VPE} = ENU \left[\hat{x} - x_{surveyed} \right]_{vertical}, \quad \chi_{HPE} = ENU \left[\hat{x} - x_{surveyed} \right]_{horizontal} \quad (5-9)$$

Here, χ_{VPE} is the vertical component and χ_{HPE} is the horizontal component of the position error in an east, north, up (ENU) coordinate system. These position errors are compared to a fixed threshold which is derived in the next section. Finally, if the position error exceeds the detection threshold, a flag is issued and sent to Executive Monitoring (EXM), which excludes the underlying faults.

Integrity checks with satellite outages are needed to anticipate cases where the aircraft is not tracking all satellites approved by the LGF. Thus, we compute position solutions from the LGF corrections using not only all visible satellites approved by the LGF but also all reasonable subsets of these satellites to which an aircraft may be limited. These subsets include the “all approved SVs in view” case (approved by the LGF prior to the PDM taking action), all “one-SV-out” combinations, and all “two-SV-out” combinations. Again PDM compares the position errors of each subset to a fixed threshold, and an alarm is issued if the position error exceeds the threshold. If all alerts include a single common satellite, that satellite can be excluded (by EXM) from the broadcast correction. If all alerts include two common satellites, it is acceptable for EXM to exclude one of them and recheck on the next epoch provided that the LGF time-to-alert requirement is still met. If neither one nor two common satellites exist, the system must exclude all corrections and generate empty Type 1 messages. This should be rare because the majority of failures will be limited to individual ranging sources.

5.2.3 THRESHOLD DERIVATION

The position-domain monitor is designed to detect anomalous behavior in a satellite or reference receiver that is too small to be detected by the existing LGF range-domain monitors. The key in monitoring is to determine what anomalous behavior of the system would result in an integrity or safety risk. For this reason, thresholds for the monitors are derived and verified using real data. A threshold may be determined theoretically. However, in most cases, it is hard to find a theoretical model for an actual distribution that is accurate beyond two to three standard deviations. In addition, the theoretical bounds are

not reliable since test statistics are highly dependent on a system configuration such as antenna siting, gain pattern, and operation environment. Thus we calculate the threshold from a Gaussian distribution that overbounds the two tails of the apparent – observed from data – distribution. This overbounding procedure is applied to most of the integrity monitor test statistics [38].

Here we determine the threshold based on empirical data. The vertical position errors, χ_{VPE_normal} , are normalized by their theoretical sigma, $\sigma_{VerticalPositionError}$, which is projected into the position domain. They are used as inputs for deriving the threshold.

$$\chi_{VPE_normal} = \frac{\chi_{VPE}}{\sigma_{VerticalPositionError}}; \quad \sigma_{VerticalPositionError} = \sqrt{\sum_{n=1}^N S_{vertical,n}^2 \sigma_{PR,n}^2} \quad (5-10)$$

Based on the distribution of χ_{VPE_normal} , a sigma inflation factor is determined so that a zero-mean Gaussian distribution overbounds the apparent tails of the measured distribution. In order to limit fault-free alarms, the PDM threshold is set as six times its theoretical standard deviation based on the continuity sub-allocation [15], that is,

$$Threshold_{PDM} \equiv \mu_{VPE_normal} \pm 6 \cdot f \cdot \sigma_{VPE_normal} \quad (5-11)$$

where μ_{VPE_normal} and σ_{VPE_normal} are the sample mean and standard deviation of the test statistics, respectively, and f is the sigma inflation factor.

The threshold for the IMT-PDM is established based on real data collected on 25 November 2002. Figures 5.3 and 5.4 show the probability density function and the cumulative distribution function of the normalized vertical position errors on a log scale. For this nominal IMT-PDM dataset, the required sigma inflation factor is 1.56, and the resulting upper and lower thresholds are ± 6.2531 .




Figure 5.3: Probability Density Function of Normalized Vertical Position Errors




Figure 5.4: Cumulative Distribution Function of Normalized Vertical Position Errors

5.2.4 NOMINAL TESTING

The position domain monitor has been tested for the cases of “all SVs in view”, all “one-out” SV combinations, and all “two-out” SV combinations under nominal conditions. For these tests, we applied the Ground Accuracy Designator B (GAD-B) model for σ_{pr_gnd} . Figure 5.5 shows the results of applying the PDM algorithm with “all SVs in view” of the IMT. The thick curve displays the normalized Vertical Position Errors (VPE), which stay well between the fixed detection thresholds of ± 6.2531 . Thus, the theoretical sigmas appear to be good estimates.




Figure 5.5: Normalized Vertical Position Errors and Detection Thresholds from IMT-PDM Nominal Data (All Approved SVs in View)




Figure 5.6: Normalized Vertical Position Errors and Detection Thresholds from IMT-PDM Nominal Data (all “one-SV-out” combinations)




Figure 5.7: Normalized Vertical Position Errors and Detection Thresholds from IMT-PDM Nominal Data (“two-SV-out” combinations)

Figures 5.6 and 5.7 show the results obtained from all “one-out” and “two-out” SV combinations in this IMT-PDM dataset. For “one-out” cases, the number of combinations is N , where N is the number of measurements or satellites approved by the IMT. Since the maximum number of measurements is eight in this dataset, the resulting number of cases is eight. For “two-out” cases, the number of combinations is 28, or $N(N-1)/2$ permutations, where the maximum N is eight. Similar to the result of the “all SVs in view” case the normalized VPEs stay well between the fixed detection thresholds for all combinations over time. The results from only four combinations out of eight and 28 permutations are presented and very similar results have been achieved from the remaining combinations.

5.2.5 FAILURE TESTING

PDM has been tested under failure conditions to verify that it can detect threatening anomalies. In failure testing, controlled errors are injected into IMT datasets to test the detection of anomalies. In order to induce sigma violations, we first insert errors into stored nominal receiver packets previously collected by the IMT antennas. We have presented the detailed method to modify nominal pseudorange measurements in Section 3.3.3.2 (see Equations (3-12) and (3-13)). The sigma of the nominal pseudorange error is approximately increased to L times the previous value. Then we put these modified values back into the IMT in a post-processing mode. Lastly, we apply the erroneous pseudorange corrections generated with failure-injected IMT datasets to the PDM algorithms.




Figure 5.8: IMT-PDM Sigma Failure Test with L=3 (All Approved SVs in view)




Figure 5.9: IMT-PDM Sigma Failure Test with L=8 (All Approved SVs in view)

Here we show the results of applying the PDM algorithms under failure conditions. The ρ_{raw} errors on a single satellite (SV 4) and all three reference receivers are increased to $L = 3$ times the nominal error at every epoch. The dotted line in Figure 5.8 shows the normalized vertical position errors, which exceed the detection threshold at 92.5 minutes after the PDM starts its run. The position error monitor may not detect moderate sigma violations faster than the sigma monitor, which normally flags in one hour (when 18 independent B -values have been collected) [50]. However, when the anomalies are not immediately hazardous, the position error monitor is a better discriminator than the sigma monitor from the airborne user's point of view; it has the ability to separate anomalies which can be tolerated from those which must be alerted to protect LAAS user integrity.

In addition to the sigma failure test as described above, another failure test has been done with an increased error factor of $L = 8$, for which an immediate alarm is required. Errors with $L = 8$ are injected on all three receivers and SV 4, which may not be detected by the Multiple Receiver Consistency check (MRCC) in the IMT, since B -values represent pseudorange correction differences across reference receivers. The normalized VPEs – shown in Figure 5.9 as the dotted line – cross the lower threshold (-6.2531) at 9.3 minutes after the fault is injected. Overall these tests suggest that the position error monitor properly detects urgent hazards and provides added integrity to LAAS users.

5.3 SIGMA INFLATION IN POSITION DOMAIN

As mentioned earlier, the PDM can avoid conservatism encountered by range domain monitoring (RDM) when computing protection bounds. This comes from the fact that PDM supports a smaller inflation factor for the broadcast σ_{pr_gnd} . To demonstrate this, we will first examine an error distribution in the position domain, since the inflation factor strongly depends on this error distribution. After that, we will derive a new inflation factor based on both empirical data – the results from Section 5.3.1 – and the inflation factor determination method described in Chapter 4.

5.3.1 ERROR DISTRIBUTIONS

In this section we investigate how range domain error statistics are converted into position domain error statistics. The purpose here is to show that the error distribution has thinner tails than before the conversion, which is a key factor supporting a smaller sigma inflation factor. First let us examine the relationship between pseudorange correction errors and position errors. We have:

$$\chi = S_1(\hat{\varepsilon}_1 - b_1) + S_2(\hat{\varepsilon}_2 - b_2) \dots + S_N(\hat{\varepsilon}_N - b_N) \quad (5-12)$$

In this equation, $\hat{\varepsilon}_n$ is the pseudorange correction error with zero-mean and b_n is the mean bias of the correction errors for each satellite n . The position error, χ , is the sum of mean-biased correction errors, which are also weighted by the coefficients of the projection matrix (S_n). We now develop the connection between error distributions in the range and position domains. By application of the Central Limit Theorem, the probability density function (PDF) of the sum of the weighted and mean-biased independent variables is the convolution of their respective scaled and mean-shifted PDFs [57]. Based on this theorem and Equation (5-12), the probability density function of position errors $f(\chi)$ is:

$$f(\chi) = \frac{1}{S_1} f\left(\frac{\hat{\varepsilon}_1 - b_1}{S_1}\right) * \frac{1}{S_2} f\left(\frac{\hat{\varepsilon}_2 - b_2}{S_2}\right) \dots * \frac{1}{S_N} f\left(\frac{\hat{\varepsilon}_N - b_N}{S_N}\right) \quad (5-13)$$

Since $\hat{\varepsilon}_n$ is weighted by S_n and biased by b_n , the PDF of $\hat{\varepsilon}_n$ is scaled by S_n and shifted by b_n . We take convolutions of these PDFs to obtain $f(\chi)$.

A theoretical model of pseudorange correction errors ($\hat{\varepsilon}_n$) was developed in Section 4.4.2 (see Equation (4-1)). Using this model – shown in Figure 5.10 as an outer (blue) curve – we transform the error distribution in the range domain into the position domain. This is done by convolving the correction error PDFs, which are scaled and mean-shifted. The given weighting parameters, S_n , and mean-bias parameter, b_n , make the established position error model into a good representation of the empirical data (which will be shown in Figure 5.11). It is clear that the tails of the position error distribution (the inner green curve in Figure 5.10) are thinner than those of the individual correction error distributions.




Figure 5.10: Error Distributions in Position Domain and in Range Domain

5.3.2 SIGMA INFLATION FACTOR

The result from the previous section implies that PDM allows us to reduce the inflation factor for broadcast σ_{pr_gnd} because its error density has thinner tails. The purpose of this section is to determine a new inflation factor which will be used for an availability analysis in the next section. For this we will use the inflation factor determination method developed in Chapter 4. We have already discussed three sources of sigma uncertainty: the finite sample size, process mixing, and the limitation of the sigma monitor on mean time-to-detect. Among induced parameters to cope with these sources, only the second one – the buffering parameter to cover process mixing – will change; the others remain the same in the position domain.




Figure 5.11: Probability Density Function of the Normalized Vertical Position Errors (Error Distribution in Position Domain)

In Figure 5.11 we show the distribution of vertical position errors, computed using the PDM algorithm for a period of 5 hours. The actual position error distribution (the dotted curve) is well characterized by the theoretical model (the solid green curve) shown in Figure 5.10. Note that this distribution has a shifted mean due to the mean biases of pseudorange correction errors. Since we assume a zero-mean normally distributed error model in the computation of PLs, we need to inflate the nominal sigma of a zero-mean Gaussian distribution to cover the non-Gaussian tails of the non-zero actual distribution. Thus we inflate the sigma to meet the H0 integrity risk allocation (1.2×10^{-10} for CAT II/III [15]). We see that the 1.56σ Gaussian distribution (the solid red curve) well overbounds the theoretical model (the solid green curve). Consequently the minimum tolerable buffering parameter to mitigate integrity risks due to mixing of process is 1.56, which is identical to the buffering parameter shown in Figure 5.3. We then combine this factor with other partial parameters derived in Chapter 4 and obtain the total inflation factor using the same method presented in Chapter 4. As shown in Figure 5.12, when we add PDM to the current RDM, we have decreased the overall inflation factor from 2.78 to 1.87.

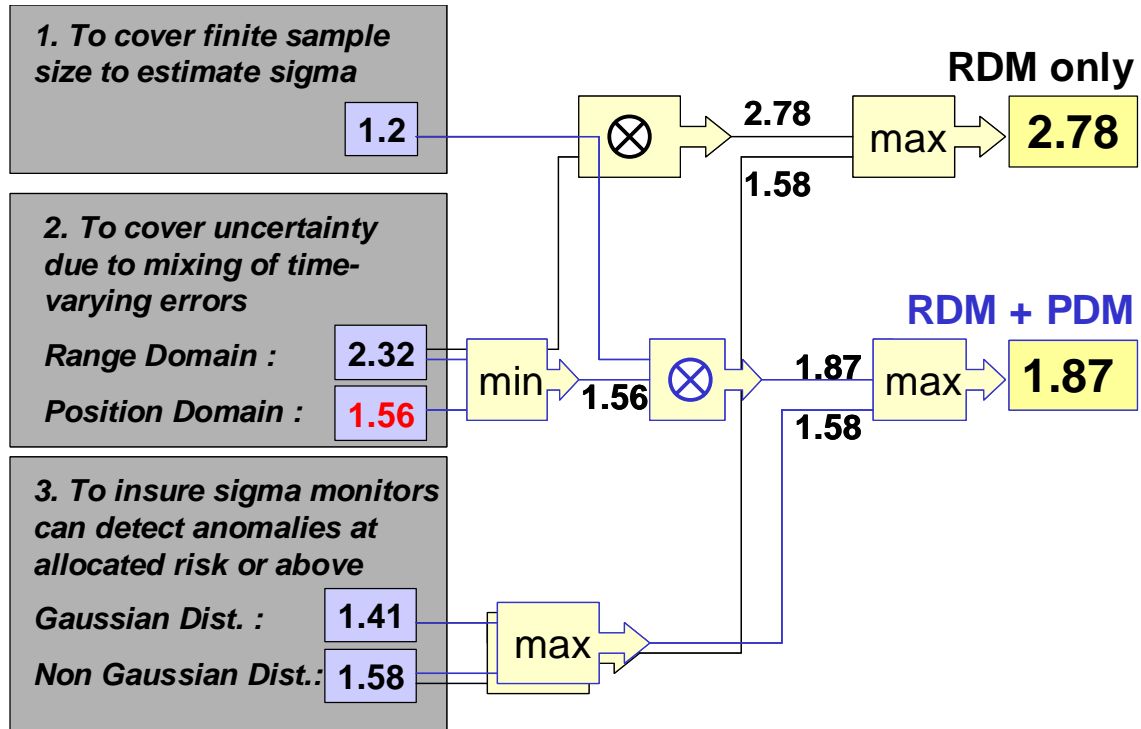


Figure 5.12: Inflation Factors for Broadcast σ_{pr_gnd} with RDM Only and RDM+PDM

5.4 PERFORMANCE ANALYSIS

In order to test the capability to meet the high availability requirement for CAT II/III precision approaches, we have tested LAAS augmented with the PDM algorithm proposed here. We use the same LAAS “pseudo-user” receiver introduced in Chapter 4. In Figure 5.13, we show the configuration of the Stanford LAAS performance test-bed. We skip the details because we have presented “pseudo-user” augmentation in Section 4.3.1 and the PDM setup in Section 5.2.1. The distance between the pseudo-user antenna and the PDM antenna is approximately 360 meters.

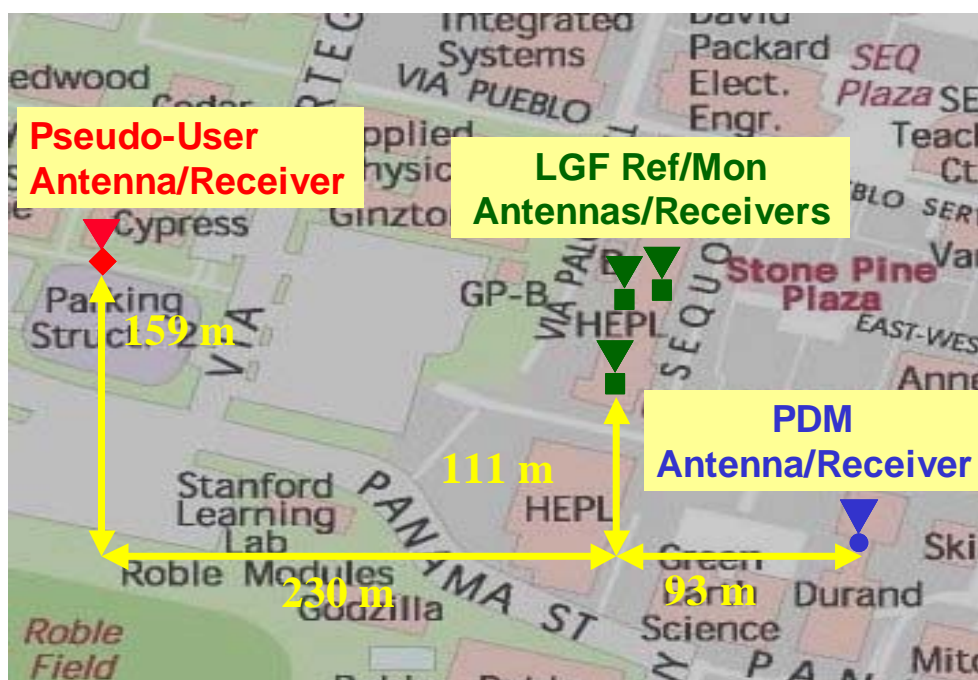


Figure 5.13: Stanford LAAS Performance Test-bed
IMT-PDM-User Hardware Configuration

We have already computed pseudo-user position errors while evaluating the performance of LAAS with only RDM (i.e., IMT) in Section 4.3.2. These quantities are not affected by augmenting the current LGF – which consists of range domain monitors only – with PDM. However, this augmented system provides users with sharper confidence bounds due to the reduced inflation factor. We now focus on computing the new vertical protection levels under the “fault-free” H_0 hypothesis. From Equation (1-13):

$$VPL_{H0} = K_{ffmd} f_{inflation} \sigma_{VerticalPositionError} \quad (5-14)$$

where $\sigma_{VerticalPositionError}$ is the standard deviation of vertical position errors, and $f_{inflation}$ (equal to 1.87) is the inflation factor derived in Section 5.3.2. K_{ffmd} is the quantile of a unit Gaussian distribution corresponding to 10^{-9} for CAT II/III operations and is equal to 6.441 when the number of ground reference receivers is three [14]. We know that

$$\sigma_{VerticalPositionError} = \sqrt{\sum_{n=1}^N S_{vertical,n}^2 \sigma_{PR,n}^2}; \quad \sigma_{PR,n}^2 = \sigma_{air,n}^2 + \sigma_{tropo,n}^2 + \sigma_{iono,n}^2 + \sigma_{pr_gnd,n}^2 \quad (5-15)$$

from Equations (1-11) and (1-12). Because the “pseudo-user” receiver is located on the ground, we replace the airborne error sigma, σ_{air} , with the ground facility error sigma, σ_{pr_gnd} . Then we can write:

$$VPL_{H0} = K_{ffmd} f_{inflation} \sqrt{\sum_{n=1}^N S_{vertical,n}^2 \cdot [3\sigma_{pr_gnd,n}^2 + \sigma_{tropo,n}^2 + \sigma_{iono,n}^2 + \sigma_{pr_gnd,n}^2]} \quad (5-16)$$

Here $\sigma_{air,n} = \sqrt{3}\sigma_{pr_gnd,n}$ since $\sigma_{pr_gnd,n}$ is set based on three reference receivers (see Equation (2-10)). We also use the fact that $\sigma_{tropo,n}$ and $\sigma_{iono,n}$ are negligible because of the short distance (approximately 230 meters) between the IMT and “pseudo-user.” As a result we have:

$$VPL_{H0} = K_{ffmd} \sqrt{\sum_{n=1}^N S_{vertical,n}^2 \cdot [3(f_{inflation} \cdot \sigma_{pr_gnd,n})^2 + (f_{inflation} \cdot \sigma_{pr_gnd,n})^2]} \quad (5-17)$$

This means that we can directly apply the inflation factor derived using the position domain error statistics to the broadcast σ_{pr_gnd} .




Figure 5.14: System Performance in Vertical Direction with RDM and PDM




Figure 5.15: System Performance in Vertical Direction with RDM only

Figure 5.14 shows the performance of the PDM-augmented system. For the purpose of comparison, we reproduce the performance results with RDM only – which have already been shown in Section 4.3.2 – and plot it in Figure 5.15. Horizontal axes indicate the absolute value of Vertical Position Errors ($|VPE|$), and Vertical Protection Levels (VPLs) are plotted in the vertical axes. Each bin represents the number of occurrences of a specific (error, protection level) pair and the color of each grid indicates the total number of epochs when that pair occurred. The $|VPE|$ s are always less than 2 meters, which means that both types of LGF systems meet the accuracy requirement for CAT II/III approaches. As noted earlier, integrity risk is defined as the probability that the position error exceeds the alert limit and no navigation system alert occurs within the time-to-alarm. The event with VPL less than the Vertical Alert Limit (VAL) but error greater than the VAL – which leads to Hazardously Misleading Information (*HMI*) – indicates a violation of integrity. In both plots, the errors are always less than the VPL and also VAL; thus no points constitute integrity failures.

Now we turn our attention to LAAS availability, which is defined as the fraction of time for which the system is providing position fixes to the specified level of accuracy, integrity and continuity (see Section 2.3). If computed protection levels exceed the alert limit, then the system no longer meets the integrity requirement and thus loses availability. As we know, the VAL for CAT II/III precision approaches – indicated with horizontal and vertical lines in Figures 5.14 and 5.15 – is 5.3 meters based on the 1998 RTCA LAAS MASPS [56] (see Table 2.1). Without PDM (and given the condition that available satellites in view are more than five), the system availability achieved in this analysis is only 89.258 %, as shown in Figure 5.15. Thus RDM alone cannot meet the availability requirement of CAT II/III approaches, which is 99.999 %. In contrast, we can see in Figure 5.14 that the system augmented with PDM maintains the availability of 99.999% or greater in vertical positioning (when the same GPS constellation is provided as in the RDM only case). Since PDM supports a smaller inflation factor ($f_{inflation} = 1.87$ versus 2.78) – inserted in Equation (5-17) to compute VPLs – and consequently provides sharp protection bounds, the PDM-augmented system is able to meet the high CAT II/III availability requirements.

5.5 USE OF POSITION DOMAIN MONITOR MEASUREMENTS

So far, we have dealt with PDM for providing sharp error bounds and supporting high availability in the previous section. In this section, we examine two proposed methodologies to further enhance integrity and continuity using PDM: the PDM Cumulative Sum (CUSUM) Monitor and the Screening Process Method. Test results are presented along with the algorithms.

5.5.1 PDM CUMULATIVE SUM (CUSUM) MONITORING

As we saw in Chapter 3, the Cumulative Sum (CUSUM) method effectively detects sigma-mean anomalies. In this section, we apply the CUSUM to PDM error statistics to further improve integrity. The CUSUM algorithm implemented in PDM is essentially the same as that of the range-domain sigma monitoring (refer to Section 3.3.2.1 for the details of this algorithm). The only change is that the input (Y) for each epoch (N) is the squared and normalized values of the vertical position errors (VPE).

$$Y(N) = \left(\frac{\chi_{VPE}(N) - \mu_{VerticalPositionError}(N)}{\sigma_{VerticalPositionError}(N)} \right)^2 \quad (5-18)$$

The Head-Start CUSUM variant has been tested with the IMT-PDM data under nominal conditions. The top plot in Figure 5.16 displays PDM-CUSUMs, and the lower plot shows the normalized VPE that fed the CUSUM. Note that the required inflation factor in the position domain is 1.87 as shown in Section 5.3, and any out-of-control sigma greater than this inflation factor should be detected within a day based on the LGF specifications. Thus, the CUSUM in this case is targeted at an out-of-control sigma 1.87 times that of the theoretical sigma ($\sigma_1 = 1.87$), which gives a high windowing factor ($k = 1.753$). The CUSUM is initialized at $h/2 = 18.9$ and is reset there every time it falls below zero. Under nominal conditions, the CUSUM slowly falls toward zero, since the normalized VPE² is usually below k and k is subtracted off at each epoch. We update the CUSUM every 200 seconds, which correspond to two times the carrier-smoothing intervals, so that successive

updates are statistically independent. The threshold of 37.8 is never approached, and no flags are observed.




Figure 5.16: PDM-CUSUM Results from Nominal Data

Range measurements from all reference receivers could possibly experience errors with equal variances under failure conditions. However, the existing range-domain sigma monitors may not observe common mode failures (for example multipath correlation), since those monitors rely on B -values, which are based on differences between pseudorange corrections across reference receivers. In order to simulate such failure conditions, we injected controlled errors into stored nominal receiver packets using the code-minus-carrier method (refer to Section 3.3.3.2). The ρ_{raw} errors on all satellites in view are increased to three times the nominal error, and these injected errors are exactly the same for all reference receivers.

Figure 5.17 shows the result of applying the Head-Start CUSUM variant to failure-injected IMT-PDM. The Head-Start CUSUM – initialized at $h/2 = 18.9$ – adds up the increased normalized VPE due to the severe errors injected onto range measurements. The PDM-CUSUMs cross the detection threshold at 26.65 minutes after the fault injection. In contrast, these anomalies were not detected by the current range-domain sigma monitoring algorithms because the fault could not be observed in the B -values.




Figure 5.17: PDM-CUSUM Results from Failure Test
(3 x Error Sigma on All SV and All RR)

5.5.2 SCREENING PROCESS

The procedure to determine the satellite subsets to be processed by the PDM is called the “screening process.” Each subset is considered based on its calculated vertical protection level (VPL) at the PDM, and only those with VPLs that are less than VAL are processed further. As explained, PDM derives position solutions by applying LGF corrections to all visible satellites approved by the LGF and all reasonable subsets of these satellites – determined by this screening process – that an aircraft may be limited to using.

If we use the PDM outputs from the screening process, we can improve LGF performance. The idea is to relax a key assumption of Category II/III LGF monitoring, which is that all airborne users have vertical protection levels (VPLs) right at the 5.3-meter maximum imposed by the Vertical Alert Limit (VAL) [58]. In practice, the truth is almost certainly better, as shown Figure 5.15. Figure 5.18 illustrates how this information is used in real time to improve average continuity. If the worst computed VPL from the PDM outputs (denoted as W_VPL) is less than the VAL, as shown in the left-hand fork, the effective VPL_{H0} can be made to be equal to the VAL by loosening the integrity monitor detection thresholds. This increases the effective Minimum Detectable Errors (MDEs). As a result of this process, continuity risk is lowered while integrity is maintained.

On rare occasions, aircraft may happen to see a subset of GPS satellites that was not directly checked by the PDM. These cases are still protected by their own VPL calculations; they only suffer a slight increase in integrity risk if their VPL exceeds W_VPL but is still below VAL (if it were above VAL, the aircraft could not conduct a Category II/III approach). This limited integrity risk increase is deemed acceptable for sigma-mean monitoring if it is sufficiently rare and implies no greater than one order of magnitude of increased overall system risk [15].

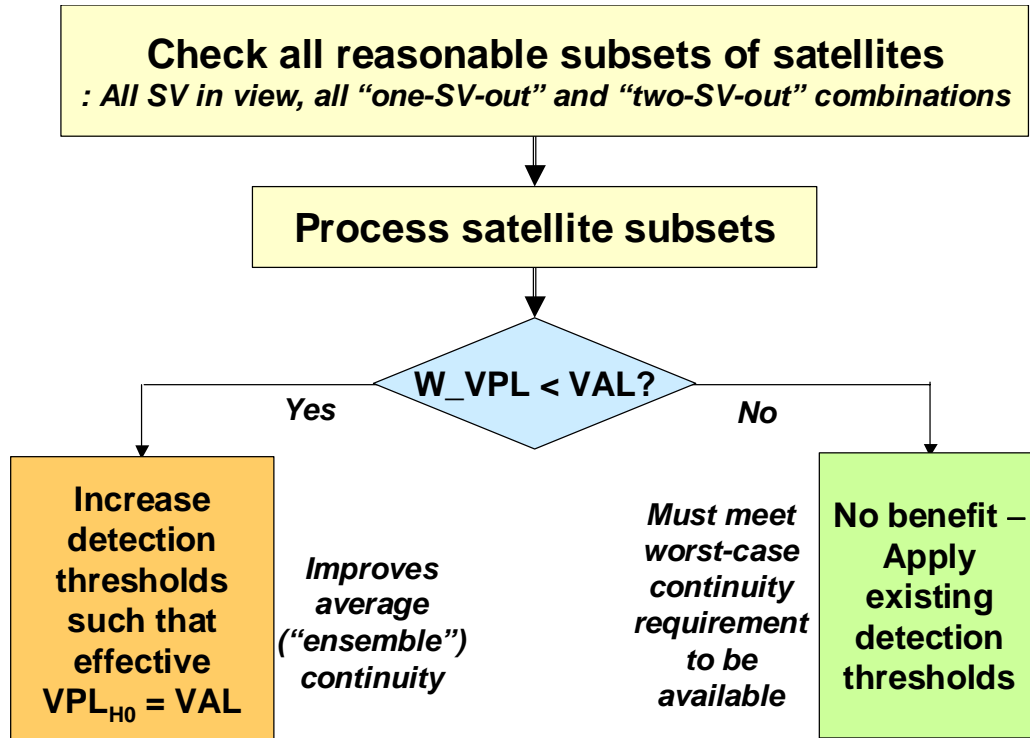


Figure 5.18: Use of PDM Screening Process Outputs to Enhance Average Continuity

The limitation of this concept is that to achieve a required availability of 0.999 or higher the LGF must still meet the integrity and continuity requirements when W_VPL exceeds VAL (the right-hand fork of Figure 5.18). Since no threshold increase is possible in this case, the baseline thresholds must meet the worst-case continuity requirement.

The worst VPL_{H0} obtained from “two-out” SV combinations with the same IMT-PDM dataset are plotted in Figure 5.19. The maximum number of measurements is ten in this dataset, resulting in 45 permutations. Given that all W_VPL s are less than VAL , increased detection thresholds can be applied for sigma monitoring, as shown in Figure 5.20. The existing detection threshold, 1.87, is equal to the inflation factor in position domain, since an out-of-control sigma above the inflation factor is defined as a failure. The increased inflation factor – with which the effective VPL_{H0} would be the same as VAL – becomes the new detection threshold. If W_VPL is greater than VAL , there will be no benefit from using the PDM outputs.




Figure 5.19: The Worst-case VPL_{HO} Out Of All "Two-SV-Out" combinations




Figure 5.20: Increase Detection Thresholds Such That Effective $VPL_{HO} = VAL$

Figure 5.21 presents the prior probability density function of an out-of-control sigma (σ_{ooc}) modeled as a Gamma distribution with parameters $a = 20.5$ and $b = 0.024$. In this case, the probability of σ_{ooc} exceeding the detection threshold, 1.87, is set to 10^{-4} . Based on this prior PDF and the results shown in Figure 5.20, we compute probabilities of σ_{ooc} exceeding new thresholds. The synthetic result demonstrates an improvement on average continuity (*Mean Time Between Failure*) by 27%.




Figure 5.21: Prior Probability Density Function for Out-Of-Control σ

The improvement in average continuity is also a substantial benefit to Executive Monitoring (EXM). The most difficult task of EXM is to distinguish between different failure classes and to separate hazardous anomalies from fault-free alerts [36, 59]. As thresholds are pushed lower to satisfy tighter Category II/III integrity requirements, smaller off-nominal conditions that are not hazardous to LAAS users are more likely to be flagged. As a result, EXM has more trouble distinguishing “real” failures. In contrast, increasing thresholds will reduce the rate of off-nominal but non-hazardous exclusions and lessen the conservatism of EXM fault exclusion and recovery. Thus, this practice improves overall satellite availability for LAAS in addition to lowering average continuity risk.

5.6 CONCLUSION

Position Domain Monitoring (PDM) improves the performance of the existing Category I LGF such that it can support Category II/III operations. We found that the performance achieved by adding PDM aids significantly in meeting the stringent availability requirements of Category II/III operations. This improvement is possible because PDM supports a lower inflation factor for the broadcast σ_{pr_gnd} . We have also seen that the PDM algorithm implemented in the Stanford LAAS Integrity Monitor Testbed detects threatening anomalies, including sigma violations that would not be detected by RDM alone. In addition, the PDM CUSUM approach improves upon the PDM algorithms by providing extra navigation integrity to users. Lastly the PDM screening process – utilizing protection level for subsets of satellites in view – lowers average continuity risk while maintaining the required integrity.

Chapter 6

Conclusion

6.1 SUMMARY OF CONTRIBUTIONS

The objective of this thesis was to design a set of algorithms for ground based augmentation systems that bound the error of differentially corrected measurements and position estimates without excessive conservatism. These algorithms process GPS measurements collected by a set of redundant reference receivers. They estimate and monitor the standard deviation (σ) of differentially corrected pseudorange errors and provide an error bound. Several problems are addressed in this thesis. First, it is difficult to characterize the error distribution of differentially corrected pseudorange measurements precisely. Second, sudden and unexpected measurement anomalies may occur. Third, the estimated σ has statistical uncertainties and must be inflated when broadcast to aircraft users. Finally, the same level of σ inflation applied for LAAS CAT I precision approaches may be not acceptable for CAT II/III operations, which require higher availability. The contributions of this work in solving these problems are summarized in the following sections.

6.1.1 SIGMA-MEAN ESTIMATION AND MONITORING

Two sigma-monitoring algorithms were designed in Chapter 3. The first algorithm computes the sample standard deviations of the pseudorange correction errors in real time. The classical application of the estimation method requires knowledge of the error distribution. Since the error distribution is not known in our problem (we cannot assure that the population is normally distributed), a significant number of independent samples are needed to ensure that the estimated sigma has a chi-square statistic. This results in the time to detection of sigma violations (i.e., the response time) being at least one hour for errors of any size. In contrast, the Cumulative Sum (CUSUM) method does not require any initial waiting duration before the first check, since the Markov property of the CUSUM allows us to determine the threshold regardless of a sample size. Further improvement is possible with “head-start” CUSUMs; these expedite detection by starting the CUSUM closer to the threshold. Figure 3.17 (repeated here as Figure 6.1) shows that the sigma estimation method detects smaller violations faster, and the “head-start” CUSUM is superior for detecting larger sigma anomalies. The successful combination of these two sigma-monitoring algorithms provides the most rapid detection regardless of fault magnitude.

This thesis also applies both direct estimation and CUSUM methods to mean monitoring. Real-time monitors are needed to detect unexpected situations in which the true mean becomes non-zero and, consequently, the position error exceeds the protection bound. It is shown that the performances of the estimation and CUSUM methods are similar when detecting larger mean violations (refer to Figure 6.2, which is a duplicate of Figure 3.18). For smaller mean anomalies, the “head-start” CUSUM achieves faster detection as the head-start values increase.

All sigma-mean monitors have been successfully integrated with the existing IMT fault-exclusion logic and have been demonstrated to detect anomalies that cause the true sigma to exceed the broadcast sigma or the true mean to become non-zero during LAAS operations. Specifically, any out-of-control sigma or mean greater than 1.8 can be detected within one hour. This performance meets the LGF requirements.




Figure 6.1: Time-to-Alert for Sigma CUSUM and Sigma Estimation Monitors




Figure 6.2: Time-to-Alert with $P_{MD} < 0.001$ for Mean Estimation and Mean CUSUM Monitor Performance

6.1.2 SIGMA INFLATION AND PERFORMANCE

The second step of my dissertation (Chapter 4) addressed sigma inflation based on the characterization of the error distribution. This thesis considered three main sources of statistical uncertainties on the estimated sigma. First, the limited size of measurements available to estimate the sigma and mean prior to commissioning introduces the estimation error. Second, time-varying environmental conditions and normalization by imperfect theoretical sigmas introduce mixing and therefore, the Gaussian assumption may not be valid. Lastly, a limited number of samples available in real time for sigma monitors necessitates inflating sigma to meet reasonable time-to-alert requirements. The thesis developed the inflation factor determination method and determined the total inflation factor for the broadcast sigma by combining all necessary buffering parameters. This sigma inflation method, when combined with my monitoring scheme shown in Chapter 3, is sufficient to maintain user integrity under both nominal and failure conditions. This thesis showed that for Category I approaches, the requirements of continuity and availability are also met with sigma inflation.

6.1.3 POSITION DOMAIN MONITORING

Chapter 5 showed that the position-domain algorithm designed in this thesis improves availability relative to the current range-domain algorithm. This position-domain method computes a position estimate by applying LGF corrections to range measurements, and it generates a position error estimate by comparing this position estimate to the known true position. The sigma inflation factor derived from the position error statistics was smaller than that derived from the range correction error statistics. In fact, it is difficult to meet the tightened requirements of Category II/III approaches with the sigma inflation factor derived in the range domain. In this case, the availability of a system was 89% as shown in Figure 6.3. In contrast, the system augmented with the position-domain algorithm provides an availability of 99.999% or higher, because there is no availability penalty due to the conservative inflation factor as there is with the range-domain method.

Further improvement on navigation integrity was achieved by applying the position domain monitor (PDM) and the PDM CUSUM monitor. Both methods perform integrity checks using the position solutions generated from the position-domain algorithm as inputs.

The thesis has also shown that outputs from the PDM screening process can be used to improve continuity. This improvement does not have an impact on the required integrity, since this method takes advantage of the margin between the maximum tolerable position error (i.e., alert limit) and the actual protection levels (all monitors are designed to protect users assuming that their protection levels are right at the alert limit). As an average over all satellite combinations, this algorithm will provide at least a 20% improvement (from 10^{-6} to 0.8×10^{-6}) in average continuity.




Figure 6.3: Stanford LAAS Pseudo-User Performance

6.2 SUGGESTIONS FOR FUTURE WORK

As mentioned in Chapter 4, test statistics for determining a sigma inflation factor depend heavily on antenna sites, gain patterns, and the nearby environment. Thus, the inflation factor derived in Chapter 5 is not applicable to all situations. For example, a challenging airport site or a ground-based system with upgraded hardware would require different inflation factors. Future research should focus on constructing more PDM test statistics with datasets collected from different environmental conditions for a longer period of time (multiple 24-hour datasets). Analyzing these statistics will help better evaluate PDM performance.

Another area of future work is real-time integration of PDM into the Stanford integrity monitor testbed (IMT). This may be necessary to better emulate an operational LGF. Since the PDM computer has access to a T1 network connection, it is possible to relay the processed PDM test statistics (worst VPL, position error, and CUSUM) back to the IMT computer in real-time. This would make it possible for Executive Monitoring (EXM) to better isolate hazardous anomalies from fault-free alerts. Reducing the rate of off-nominal but non-hazardous exclusions will reduce the conservatism of EXM fault exclusion and thereby improve overall LAAS availability in addition to lowering average continuity risk.

Further improvement of PDM-CUSUM performance is possible with the Bayesian CUSUM [60]. By combining the CUSUM with an application of Bayes' rule and a prior probability distribution for the sigma parameter being monitored, we can infer a real-time posterior distribution of sigma based on the current CUSUM state. This gives ground systems more flexibility in responding to sigma increases before an alert (possibly causing loss of user continuity) is required. Figure 6.4 shows how the Bayesian PDM-CUSUM could be used to improve real-time LGF response to anomalies. The Bayesian CUSUM offline analysis results provide a map relating CUSUM states to H0 integrity risk conditions. When the PDM-CUSUM exceeds the state that defines a "minimal" risk increase (C_{min}), no immediate action is required because this condition is allowed to last for one day before being alerted [15]. The sigma estimation monitor will separately detect this condition if it lasts long enough to be of concern.

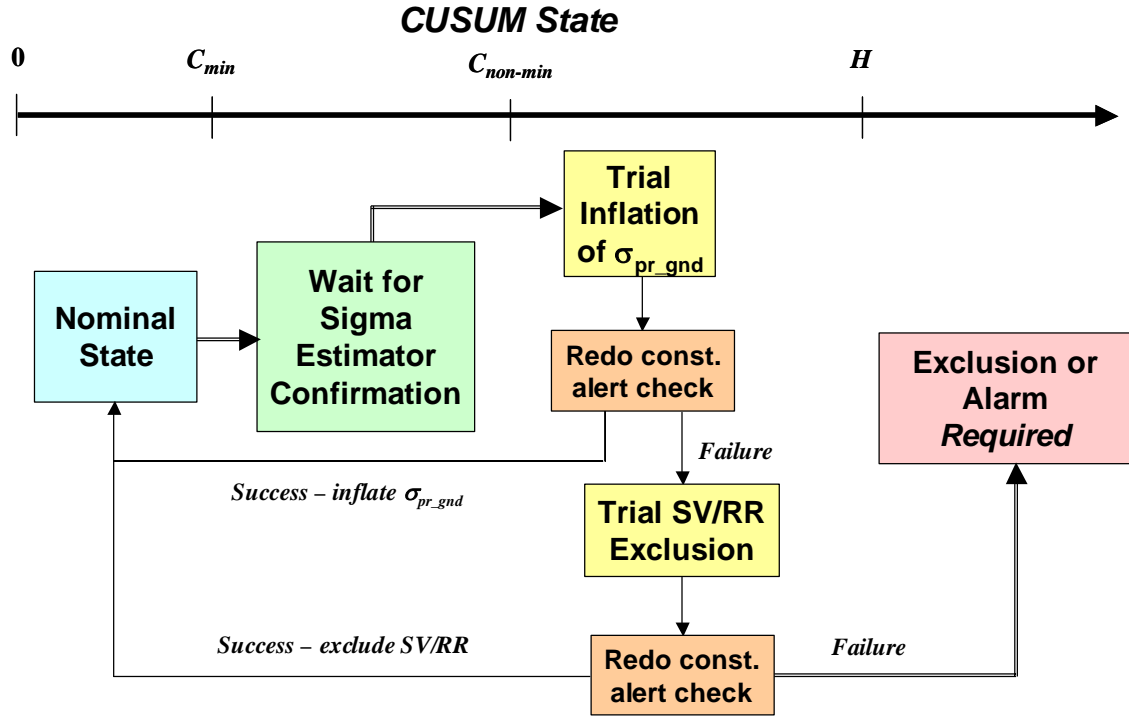


Figure 6.4: Real-Time Use of Bayesian CUSUM Outputs to Preserve Continuity and Availability while Protecting Integrity

When the PDM-CUSUM exceeds the “non-minimal” risk increase state ($C_{non-min}$), action is taken. The broadcast σ_{pr_gnd} is increased such that the H_0 integrity risk becomes acceptable again. In this case, to confirm no loss of continuity, the constellation-alert-check procedure is performed assuming an inflated value of σ_{pr_gnd} . If VPL remains under VAL, neither continuity nor availability will be lost; thus, inflation can be implemented. At the same time, the CUSUM must be recomputed – based on the re-normalized prior sigma distribution and B -value inputs – to confirm that the revised CUSUM state is below $C_{non-min}$. If inflation cannot be implemented without possible continuity loss, exclusion of problematic measurements would be tried next using the same constellation-check and CUSUM-reset process. In this manner, integrity monitoring is potentially enhanced while better preserving continuity and availability.

Appendix A

LAAS ERROR MODELS FOR NOISE, MULTIPATH, TROPOSPHERE, AND IONOSPHERE

This appendix describes the *LAAS Accuracy Models* used in this thesis. These models are required to compute the protection levels.

A.1 MODEL OF AIRBORNE PSEUDORANGE PERFORMANCE

This section defines airborne pseudorange error allocations for carrier-smoothed code processing. Let us consider two components. First, σ_{noise} is the error due to wideband noise and interference including receiver noise, thermal noise, inter-channel biases, extrapolation and processing errors, and is modeled as:

$$\sigma_{noise,n} = a_0 + a_1 e^{-\theta_n/\theta_c}; \quad 5^\circ \leq \theta_n \leq 90^\circ$$

where θ_n is the n th ranging source elevation angle, and a_0, a_1 and θ_c for the applicable Airborne Accuracy Designator (AAD) are defined in Table A-1. The AADs were defined to reflect different performance levels of GPS receiver technologies [17].

Airborne Accuracy Designator	a_0 meters	a_1 meters	θ_c degrees
GAD-A	0.15	0.43	6.9
GAD-B	0.11	0.13	4.0

Table A-1: Airborne Error Model Parameters

Second, $\sigma_{multipath}$ is the error due to airframe multipath and is described by the distribution, $N(0, \sigma_{multipath}^2)$ where:

$$\sigma_{multipath,n} = 0.13 + 0.53e^{-\theta_n/10^\circ}$$

The overall airborne accuracy model is computed as follows:

$$\sigma_{air,n} = \sqrt{\sigma_{noise,n}^2 + \sigma_{multipath,n}^2}$$

A.2 MODEL OF TROPOSPHERIC RESIDUAL UNCERTAINTY

An airborne subsystem needs to correct for a differential tropospheric delay error because the airborne user is at a different altitude than the ground reference receivers. The differential tropospheric delay correction is computed as follows [32]:

$$TC = N_R h_0 \frac{10^{-6}}{\sqrt{0.002 + \sin^2(\theta_n)}} (1 - e^{-\Delta h/h_0})$$

where: N_R = Refractivity index transmitted by the ground subsystem in the LAAS Type 2 message (unitless) [61],

Δh = Height of the aircraft above the GBAS reference point (meters),

θ_n = Elevation angle of the satellite n ,

h_0 = Tropospheric scale height from the LAAS Type 2 message (meters).

In addition to the tropospheric correction, the airborne subsystem must also estimate the standard deviation of the post correction error [32]:

$$\sigma_{tropo,n} = \sigma_N h_0 \frac{10^{-6}}{\sqrt{0.002 + \sin^2(\theta_n)}} (1 - e^{-\Delta h/h_0})$$

where σ_N is the refractivity uncertainty from the LAAS Type 2 message (unitless).

The LAAS Ground Facility (LGF) is responsible for providing broadcast parameters (N_R, σ_N, h_0 , and Δh) that are consistent with prevailing conditions at the site. If this real time estimation process is to be avoided, then these parameters must be set to constant values that cover the worst case conditions expected during operations [17].

A.3 MODEL OF IONOSPHERIC RESIDUAL UNCERTAINTY

Ionospheric temporal and spatial decorrelation can lead to differential LAAS user range errors. The uncertainty of this residual ionospheric error is defined as [32]:

$$\sigma_{iono} = F_{pp} \times \sigma_{vert_iono_gradient} \times (X_{air} + 2 \times \tau \times v_{air})$$

where F_{pp} is the vertical-to-slant obliquity factor (unitless) for a given satellite. It is computed as follows.

$$F_{pp} = \left[1 - \left(\frac{R_e \cos(\theta_n)}{R_e + h_I} \right)^2 \right]^{-\frac{1}{2}}$$

R_e = Radius of the earth = 6378.1363 km

h_I = Ionospheric shell height = 350 km

θ_n = Elevation angle of the satellite n

$\sigma_{vert_iono_gradient}$ is the standard deviation of a normal distribution associated with the residual ionospheric uncertainty due to spatial decorrelation (a parameter provided by the ground subsystem in the LAAS Type 2 message). X_{air} is the distance (slant range) between the aircraft and the GBAS reference point (meters). τ is the time constant of the smoothing filter (100 seconds), and v_{air} is the horizontal speed of the aircraft (meters/sec).

Finally, the total error model is given by (σ_{pr_gnd} is the ground error model provided in Equation (2-10)):

$$\sigma_{PR} = \sqrt{\sigma_{air}^2 + \sigma_{tropo}^2 + \sigma_{iono}^2 + \sigma_{pr_gnd}^2}$$

Appendix B

CUMULATIVE SUM (CUSUM) DESIGN

This appendix contains theoretical foundations of the *CUSUM Design* used in this thesis [23]. In the first section, we will start with the general CUSUM for the exponential family of distributions, which is important in theoretical statistics because of its good inferential properties. Based on this general theory, we will derive the CUSUM for a normal distribution in the second section. The last section deals with the computation of transition probabilities to implement Markov Chain approaches introduced in Chapter 3.

B.1 THE GENERAL EXPONENTIAL FAMILY

Let us consider the probability density for any member of the exponential family written as [62]:

$$f(x|\theta) = \exp\{a(x)b(\theta) + c(x) + d(\theta)\} \quad (\text{B-1})$$

where θ is the single parameter of the distribution and X is the corresponding random variable. This formulation includes the normal distribution with an appropriate choice of a , b , c , and d . The expression for the joint density, where \mathbf{x} is a random vector of independent identically distributed exponential variables, is:

$$f(\mathbf{x}|\theta) = \exp \left\{ \sum_{i=1}^N a(x(i))b(\theta) + \sum_{i=1}^N c(x(i)) + Nd(\theta) \right\}$$

Now suppose we want to test whether the process has gone from an in-control parameter value θ_0 to an out-of-control value θ_1 . Then we have the “score” variables, $Z(i)$, which are logarithms of the likelihood ratio ($f_0(X)$ and $f_1(X)$ are associated with a null (in-control) hypothesis (H_0) and an alternative (out-of-control) hypothesis (H_1), respectively).

$$Z(i) = \ln \left(\frac{f_1(X(i))}{f_0(X(i))} \right) = a(X(i)) \{b(\theta_1) - b(\theta_0)\} + \{d(\theta_1) - d(\theta_0)\}$$

In practice, the sequential probability ratio test (SPRT) works with this log-likelihood, and is based on the cumulative sum of $Z(i)$. If the hypothesis H_1 is true, $Z(i)$ will have a positive expected value and the cumulative sum of $Z(i)$ will tend to drift upward. If the hypothesis H_0 is true, $Z(i)$ will have a negative value and the sum will tend to drift downward. The primary difference between SPRT and CUSUM tests is that, in the CUSUM, we never accept the hypothesis of in-control, H_0 , and stop sampling. Instead, we restart the test – by setting the sum to zero – each time the evidence supports the hypothesis that we are in control (i.e. when the sum is negative).

Let us define a generic decision interval CUSUM with its recursion as:

$$\begin{aligned} D(N) &= \max(0, D(N-1) + Z(N)) \\ &= \max(0, D(N-1) + a(X(N)) \{b(\theta_1) - b(\theta_0)\} + \{d(\theta_1) - d(\theta_0)\}) \end{aligned}$$

If $b(\theta_1) - b(\theta_0) > 0$, we can rescale by dividing the CUSUM in this equation through by $\{b(\theta_1) - b(\theta_0)\}$. The expression above is now:

$$C^+(N) = \max(0, C^+(N-1) + Y(N) - k) \quad (\text{B-2})$$

where we define $Y(N) = a(X(N))$, $C^+(N) = D(N) / \{b(\theta_1) - b(\theta_0)\}$, and

$$k = -\frac{d(\theta_1) - d(\theta_0)}{b(\theta_1) - b(\theta_0)} \quad (\text{B-3})$$

If $b(\theta_1) - b(\theta_0) < 0$, then dividing the CUSUM through by $\{b(\theta_1) - b(\theta_0)\}$ reverses the direction of the CUSUM, giving the recursion:

$$C^-(N) = \min\left(0, C^-(N-1) + Y(N) - k\right) \quad (\text{B-4})$$

The CUSUM signals if its recursion exceeds a detection threshold.

B.2 DERIVATION OF CUSUM FOR A NORMAL SIGMA-MEAN SHIFT

Given the general formulation for the CUSUM shown in the previous section, we can easily derive the CUSUM for a mean or sigma shift of a normal distribution. Let us first regard σ as fixed and known, and consider monitoring for possible changes in μ . The normal density in “exponential family” form is given by:

$$f(y|\mu) = \exp\left(\frac{y\mu}{\sigma^2} - \frac{y^2}{2\sigma^2} - \ln(\sigma\sqrt{2\pi}) - \frac{\mu^2}{2\sigma^2}\right) \quad (\text{B-5})$$

From Equations (B-1), (B-3), and (B-5), it follows that

$$\begin{aligned} a(y) &= y \\ b(\mu) &= \frac{\mu}{\sigma^2} \\ d(\mu) &= -\frac{\mu^2}{2\sigma^2} \\ k &= -\frac{(-\mu_1^2 + \mu_0^2)/(2\sigma^2)}{(\mu_1 - \mu_0)/\sigma^2} = \frac{\mu_0 + \mu_1}{2} \end{aligned}$$

Applying Equations (B-2) and (B-4), we obtain the CUSUM scheme for the mean of a normal distribution as follows:

$$\begin{aligned} C^+(N+1) &= \max\left(0, C^+(N) + Y(N) - k\right) \\ C^-(N+1) &= \min\left(0, C^-(N) + Y(N) - k\right) \end{aligned} \tag{B-6}$$

This scheme signals if $C^+(N) > h^+$ or $C^-(N) < -h^-$, where h^\pm are chosen so that the scheme has a specified in-control ARL.

Now let us turn our attention to monitoring possible shifts in σ with the mean known. We want to rewrite the normal density in the form:

$$f(y | \sigma) = \exp\left(-\frac{(y - \mu)^2}{2\sigma^2} - \ln(\sigma) - \ln(\sqrt{2\pi})\right)$$

If we identify the components, we have:

$$\begin{aligned} a(y) &= (y - \mu)^2 \\ b(\sigma) &= -(2\sigma^2)^{-1} \\ d(\sigma) &= -\ln \sigma \end{aligned}$$

which yields the CUSUM with a reference value:

$$k = -\frac{\ln(\sigma_1) - \ln(\sigma_0)}{(2\sigma_1^2)^{-1} - (2\sigma_0^2)^{-1}}$$

and the CUSUM scheme shown in Equation (B-6).

B.3 MARKOV CHAIN ANALYSIS OF TRANSITION PROBABILITY

In Section 3.3.2.1 we needed to find the transition probabilities for our transition matrices, \mathbf{P} , shown in Figure B.1 (a duplicate of Figure 3.3). In other words, we sought a general way to find:

$$P(a < C^+(N+1) < b | c < C^+(N) < d)$$

Let us define $\eta(y)$ to be the distribution function of $C^+(N)$ conditional on $c < C^+(N) < d$, and the distribution function of $Y = a(X)$ to be $F(y)$. Then we have:

$$P(a < C^+(N+1) < b | c < C^+(N) < d) = \int_c^d \{F(b-c+k) - F(a-c+k)\} d\eta(c)$$

If we use Simpson's rule for the interval $[c, d]$ with midpoint m , we can further approximate the integral. As a result, we obtain the following approximation:

$$\begin{aligned} &P(a < C^+(N+1) < b | c < C^+(N) < d) \\ &= [\{F(b-c+k) + 4F(b-m+k) + F(b-d+k)\} \\ &\quad - \{F(a-c+k) + 4F(a-m+k) + F(a-d+k)\}] / 6 \end{aligned}$$

We apply this equation repeatedly to find accurate values for the transition matrices.

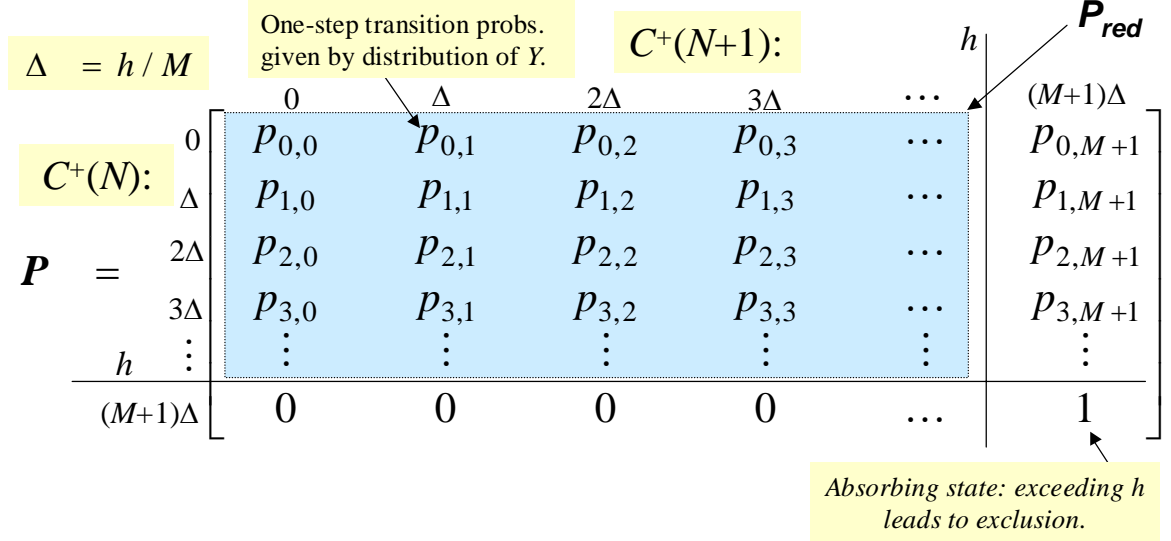


Figure B.1: CUSUM Performance Modeling with Markov Chains

Bibliography

- [1] B. Pervan and I. Sayim, "Sigma Inflation for the Local Area Augmentation of GPS," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 37, October, 2001.
- [2] P. Enge, "GPS Modernization: Capabilities of the New Civil Signals," presented at Australian International Aerospace Congress, Brisbane, Australia, July 29-Aug. 1, 2003.
- [3] P. Enge and P. Misra, *Global Positioning System: Signals, Measurements, and Performance*: Ganga-Jamuna Press, 2001.
- [4] R. Gold, "Optimal Binary Sequences for Spread Spectrum Multiplexing," *IEEE transactions Info. and Theory*, Vol. 33, pp. 619-621, 1967.
- [5] E. Kaplan, *Understanding GPS Principles and Applications*: Artech House, 1996.
- [6] B. W. Parkinson and J. J. Spilker, *Global Positioning System: Theory and Application*: AIAA Publication, 1996.
- [7] J. Klobuchar, "Design and Characteristics of the GPS Ionospheric Time Delay Algorithm for Single Frequency Users," presented at Rec. IEEE 1986 Position Location and Navigation Symp., Las Vegas, NV, 1986.
- [8] J. J. Spilker, "GPS Navigation Data," *Global Positioning System: Theory and Applications*, Vol. 1., Washington D.C.: AIAA, 1996.

- [9] G. Mack, J. G., and B. M., "Skyfix - The Worldwide Differential GPS Systems," presented at DSNS-93, Amsterdam, The Netherlands, 1993.
- [10] T. Walter, P. Enge, and A. J. Hansen, "A Proposed Integrity Equation for WAAS MOPS," presented at ION GPS-97, September 1997.
- [11] C. Kee, B. W. Parkinson, and P. Axelrad, "Wide Area Differential GPS," *Journal of the U.S. Institute of Navigation*, Vol. 38, 1991.
- [12] J. Benedicto, P. Michael, and J. Ventura-Traveset, "EGNOS, the European Regional Augmentation to GPS and GLONASS," presented at the Fifth ESA International Workshop on DSP Techniques Applied to Space Communications, Noorwijk, the Netherlands, 1998.
- [13] A. Shimamura, "MSAS Project Status," *Selected Papers on Satellite Based Augmentation Systems*, Vol. VI, pp. 53-60.
- [14] "Minimum Operational Performance Standards for GPS/Local Area Augmentation System Airborne Equipment," RTCA, Washington, D.C., RTCA SC-159 WG-4A, DO-253A, Nov. 28, 2001.
- [15] "Specification: Performance Type One Local Area Augmentation System Ground Facility," U.S. Federal Aviation Administration, Washington, D.C., FAA-E-2937A, April 17, 2002.
- [16] B. Pervan, S. Pullen, and I. Sayim, "Sigma Estimation, Inflation, and Monitoring in the LAAS Ground System," presented at Proceedings of ION GPS, Salt Lake City, UT, Sept. 19-22, 2000.
- [17] G. McGraw, T. Murphy, M. Brenner, and S. Pullen, "Development of the LAAS Accuracy Models," presented at Proceedings of ION GPS, Salt Lake City, UT, Sept. 19-22, 2000.
- [18] C. Shively and R. Braff, "An Overbound Concept for Pseudorange Error from the LAAS Ground Facility," presented at Proceedings of IAIN World Congress/ION 56th Annual Meeting, San Diego, CA, June 26-28, 2000.
- [19] J. Rife, "Core Overbounding and its Implications for LAAS Integrity," presented at ION GNSS 2004, Long Beach, CA, Sept. 21-24, 2004.

- [20] K. Markin and C. Shively, "A Position-Domain Method for Ensuring Integrity of Local Area Differential GPS (LDGPS)," presented at ION Annual Meeting, 1995.
- [21] R. Braff, "Position Domain Monitor (PDM) Performance Analysis for CAT III," MITRE/CAASD, McLean, VA, October 1 2002.
- [22] E. S. Page, "Continuous Inspection Schemes," *Biometrika*, Vol. 41, pp. 100-115, 1954.
- [23] D. M. Hawkins and D. H. Olwell, "Cumulative Sum Charts and Charting for Quality Improvement." New York: Springer-Verlag, Inc., 1998.
- [24] S. Pullen, M. Luo, B. Pervan, and P. Enge, "The Use of CUSUMs to Validate Protection Level Overbounds for Ground Based and Space Based Augmentations Systems," presented at Proceedings of ISPA 2000, Munich, Germany, 18-20 July 2000.
- [25] P. Enge, "Local Area Augmentation of GPS for the Precision Approach of Aircraft," *IEEE*, Vol. 87, January 1999.
- [26] S. Pullen, T. Walter, and P. Enge, "System Overview, Recent Developments, and Future Outlook for WAAS and LAAS," presented at Tokyo University of Mercantile Marine GPS Symposium, Tokyo, Japan, Nov. 11-13, 2002.
- [27] "Navigation and Landing Transition Strategy," Federal Aviation Administration (FAA), Washington, D.C. August 2002.
- [28] ICAO, GNSS SARPS, ICAO 1999.
- [29] T. Murphy, "Development of Signal in Space Performance Requirements for GBAS to Support CAT II/III Landing Operations," presented at ION, Portland, OR, September 2002.
- [30] G. Xie, S. Pullen, M. Luo, P. Normark, D. Akos, J. Lee, and P. Enge, "Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed," presented at ION NTM, Albuquerque, NM, June 11-13, 2001.
- [31] "FAA LAAS Ground Facility (LGF) Functions," in *Unpublished Manuscript*, Vol. 2.4: LAAS KTA Group, September 9, 1998.

- [32] RTCA, *Minimum Operational Performance Standards for GPS/Local Area Augmentation System Airborne Equipment*. Washington, D.C., Nov. 28, 2001.
- [33] M. Luo, S. Pullen, J. Zhang, S. Gleason, and G. Xie, "Development and Testing of the Stanford LAAS Ground Facility Prototype," presented at Proceedings of ION NTM, Anaheim, CA, 26-28 January 2000.
- [34] P. Normark, D. Akos, S. Pullen, G. Xie, J. Lee, M. Luo, P. Enge, and B. Pervan, "Integration of the Multipath Limiting Antenna with the LAAS Integrity Monitor Testbed," presented at ION GPS 2002, Portland, OR, Sept. 24-27, 2002.
- [35] P. Normark, G. Xie, D. Akos, S. Pullen, M. Luo, J. Lee, P. Enge, and B. Pervan, "The Next Generation Integrity Monitor Testbed (IMT) for Ground System Development and Validation Testing," presented at ION GPS 2001, Salt Lake City, UT, Sept. 11-14, 2001.
- [36] S. Pullen, M. Luo, G. Xie, and P. Enge, "GBAS Validation Methodology and Test Results from the Stanford LAAS Integrity Monitor Testbed," presented at ION GPS 2000, Salt Lake City, UT, 19-22 September 2000.
- [37] S. Pullen, M. Luo, G. Xie, and J. Lee, "LAAS Ground Facility Design Improvements to Meet Proposed Requirements for Category II/III Operations," presented at Proceedings of ION GPS, Portland, OR, September 2002.
- [38] G. Xie, "Optimal On-Airport Monitoring of the Integrity of GPS-Based Landing Systems," in *Electrical Engineering*, Stanford University, 2004.
- [39] M. Luo and S. Pullen, *LAAS Reference Receiver Correlation Analysis*: Stanford University, Feb. 17, 1999.
- [40] D. Akos, R. E. Phelts, S. Pullen, and P. Enge, "GPS Signal Quality Monitoring: Test Results," presented at ION 2000 National Technical Meeting, Anaheim, CA, Jan. 26-28, 2000.
- [41] R. E. Phelts, D. Akos, and P. Enge, "SQM Validation Report," ICAO GNSSP WG-B Meeting, Seattle, WA. May 30, 2000.
- [42] A. J. Van Dierendonck, D. Akos, S. Pullen, R. E. Phelts, and P. Enge, "Practical Implementation Considerations in the Detection of GPS Satellite Signal Failures,"

presented at IAIN World Congress/ION 56th Annual Meeting, San Diego, CA, June 26-28, 2000.

- [43] P. Enge, R. E. Phelts, and D. Akos, "Detecting Anomalous Signals from GPS Satellites," in *Draft Manuscript*. Stanford University, Oct. 12, 1999.
- [44] A. Mitelman, R. E. Phelts, D. Akos, S. Pullen, and P. Enge, "A Real-Time Signal Quality Monitor for GPS Augmentation Systems," presented at ION GPS 2000, Salt Lake City, UT, Sept. 19-22, 2000.
- [45] "Global Positioning System (GPS) Standard Positioning Service (SPS) Performance Standard," U.S Department of Defense, October 2001.
- [46] B. Pervan, "A Review of LGF Code-Carrier Divergence Issues," Illinois Institute of Technology, MMAE Dept., May 29 2001.
- [47] S. Matsumoto, S. Pullen, M. Rotkowitz, and B. Pervan, "GPS Ephemeris Verification for Local Area Augmentation System (LAAS) Ground Stations," presented at ION GPS 1999, Nashville, TN, 1999.
- [48] S. Pullen, J. Lee, M. Luo, B. Pervan, F. C. Chan, and L. Gratton, "Ephemeris Protection Level Equations and Monitor Algorithms for GBAS," presented at ION GPS 2001, Salt Lake City, UT, Sept. 11-14, 2001.
- [49] B. Pervan, S. Pullen, and J. Christie, "A Multiple Hypothesis Approach to Satellite Navigation Integrity," *Journal of the U.S. Institute of Navigation*, Vol. 45, Spring 1998.
- [50] J. Lee, S. Pullen, G. Xie, and P. Enge, "LAAS Sigma Monitor Analysis and Failure-Test Verification," presented at ION NTM 2001, Albuquerque, NM, June 11-13, 2001.
- [51] "Specification: Category I Local Area Augmentation System Non-Federal Ground Facility," U.S. Federal Aviation Administration, Washington, D.C. FAA/AND710-2937, May 31, 2001.
- [52] "Specification: Performance Type One Local Area Augmentation System Ground Facility," U.S. Federal Aviation Administration, Washington, D.C. FAA-E-2937, September 21, 1999.

- [53] J. M. Lucas and R. B. Crosier, "Fast Initial Response for CUSUM Quality Control Schemes: Give your CUSUM a Head Start," *Technometrics*, Vol. 24, pp. 199-205, 1982.
- [54] M. Gary, T. Murphy, M. Brenner, and S. Pullen, "Development of the LAAS Accuracy Models," presented at Proceedings of ION GPS, Salt Lake City, UT, Sept. 19-22, 2000.
- [55] J. Lee, S. Pullen, G. Xie, and P. Enge, "LAAS Position Domain Monitor Analysis and Failure-Test Verification," presented at 21st AIAA ICSSC, Yokohama, Japan, April 2003.
- [56] "Minimum Aviation System Performance Standards for the Local Area Augmentation System (LAAS)," Washington D.C. RTCA SC-159 WG-4A, DO-245, September 1998.
- [57] R. E. Ziemer, *Elements of Engineering Probability & Statistics*, International ed. Upper Saddle River: Prentice-Hall, Inc., 1997.
- [58] F. Van Grass, "Informal Discussion during Ohio University/Stanford University LAAS Review Meeting," Stanford University, June 21, 2002.
- [59] B. Pervan, S. Pullen, and M. Luo, "Navigation Integrity Monitoring for DGPS-Based Precision Landing of Aircraft," presented at Proceedings of the 1999 American Control Conference, San Diego, CA, June 2-4, 1999.
- [60] S. Pullen, J. Lee, G. Xie, and P. Enge, "CUSUM-Based Real-Time Risk Metrics for Augmented GPS and GNSS," presented at ION GPS/GNSS, Portland, Oregon, September 2003.
- [61] RTCA, *GNSS Based Precision Approach Local Area Augmentation System (LAAS) Signal-in-Space Interface Control Document (ICD)*: RTCA/DO-246A, Jan. 11, 2000.
- [62] D. M. Hawkins, "Evaluation of the average run length of cumulative sum charts for an arbitrary data distribution," *Communications in Statistics, series B*, Vol. 21, pp. 1001-1020, 1992.