


Spring 1-1-2019

Automatic Gain Control Measurements as a Gps L1 Interference Detection Metric

Nathan Stephen Levigne

University of Colorado at Boulder, nale6729@colorado.edu

Follow this and additional works at: https://scholar.colorado.edu/asen_gradetds

 Part of the [Aerospace Engineering Commons](#), and the [Electrical and Computer Engineering Commons](#)

Recommended Citation

Levigne, Nathan Stephen, "Automatic Gain Control Measurements as a Gps L1 Interference Detection Metric" (2019). *Aerospace Engineering Sciences Graduate Theses & Dissertations*. 245.
https://scholar.colorado.edu/asen_gradetds/245

This Thesis is brought to you for free and open access by Aerospace Engineering Sciences at CU Scholar. It has been accepted for inclusion in Aerospace Engineering Sciences Graduate Theses & Dissertations by an authorized administrator of CU Scholar. For more information, please contact cuscholaradmin@colorado.edu.

**Automatic Gain Control Measurements as a GPS L1
Interference Detection Metric**

by

Nathan S. Levigne

B.A., University of Colorado Boulder, 2018

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Masters of Science
Department of Aerospace Engineering
2019

This thesis entitled:
Automatic Gain Control Measurements as a GPS L1 Interference Detection Metric
written by Nathan S. Levigne
has been approved for the Department of Aerospace Engineering

Prof. Dennis M. Akos

Dr. Nagaraj C. Shivaramaiah

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Levigne, Nathan S. (M.S., Aerospace Engineering)

Automatic Gain Control Measurements as a GPS L1 Interference Detection Metric

Thesis directed by Prof. Dennis M. Akos

Throughout the course of this research the validity and effectiveness of using automatic gain control measurements as a means of interference detection is explored. Global Navigation Satellite Systems are a crucial part of the global infrastructure and thus it is paramount that the signal coming from these systems is available to all users at all times. The original proclamation arises from the need to create methods of detection that are universally applicable to the vast spectrum of commercial off-the-shelf GNSS receivers in use today. Current forms of detection are found to be inadequate as they require additional hardware or extensive requirements in computation power. AGC is purposed as an alternative, therefore this paper first proves that there exists an inverse proportionality relationship between the AGC metric and the power level of the received signal. This is done through extensive experimentation of the AGC response characteristics of three COTS receivers with respect to varying types of input. With the relationship proven the paper then provides a set of GPS L1 RFI detection applications which take advantage of these findings. Finally, conclusions on the effectiveness of AGC and future applications are provided.

Dedication

This thesis is dedicated to mi abuelo, Joe Sisneros, who has believed in me since the day I was born. Despite being a man of few words, his unconditional love for me, checkers, circus peanuts, and puzzles ignited my personal passion for knowledge. He was without a doubt my first teacher and I deeply resent the fact that he was unable to share this moment of triumph with me, but I know that he is smiling somewhere right now and muttering that timeless phrase, "Good job, hito!".

Acknowledgements

I would like to put forth my deepest appreciation to my advisor Prof. Dennis M. Akos, for providing me with the opportunity to study alongside him at the University of Colorado Boulder. His constant guidance and tendency to operate at breakneck speeds has greatly influenced me and has allowed me to achieve goals in my career that I never thought possible. I also wish to express my gratitude to the other members of my committee Nagaraj C. Shivaramaiah and Nisar R. Ahmed for their patience and support. I would like to extend my appreciation and sincere thanks to every professor that I had the wonderful opportunity to learn from at the University of Colorado Boulder.

Lastly, I wanted to acknowledge the agencies which have funded me throughout my academic career: Colorado Space Grant Consortium, Engineering Excellence Fund, University of Alabama in Huntsville, U.S. Army Space and Missile Defence Command, and the Colorado Center for Astrodynamics Research.

Contents

Chapter	
1	Introduction 1
1.1	Applications of Automatic Gain Control in GNSS 1
1.2	Thesis Objectives and Problem Statement 3
1.3	Thesis Overview 4
2	Background 5
2.1	GNSS Overview 5
2.1.1	GNSS Methodology 5
2.1.2	Civilian Applications of GNSS 7
2.2	GNSS Intentional Interference Overview 9
2.2.1	GNSS Jamming and Spoofing 9
2.2.2	Current Methods of Detection 11
2.3	Automatic Gain Control Overview 13
2.3.1	General Automatic Gain Control Background and Methodology 13
2.3.2	Automatic Gain Control Implementation in GNSS [3, 1] 17
2.4	Previous Work 21
2.5	Power Measurement 23
3	AGC Power Characterization 26
3.1	AGC Power Detection Theory 26

3.2	Receiver Overview	27
3.2.1	SiGe Receiver	27
3.2.2	NT1065 Receiver	28
3.2.3	UBLOX M8 Receiver	30
3.3	AGC Power Relationship Experimentation and Validation	31
3.3.1	White Noise Response	31
3.3.2	Non-White Noise Response	32
3.3.3	Receiver Specific Responses	34
3.3.4	Near Spectrum Response	37
3.3.5	Miscellaneous	40
3.4	Discussion	44
4	Applications and Future Work	45
4.1	Jammer Detection Applications	45
4.1.1	Automotive Jammer Detection Northwest Parkway	46
4.2	Jammer and Spoofing Detection Applications	52
4.2.1	Mobile Applications	53
4.3	Future Work	55
5	Conclusions	57
	Bibliography	58
	Appendix	

Figures

Figure

2.1	GNSS architecture	6
2.2	Trilateration	7
2.4	Typical AGC Block Diagram	14
2.5	Ideal AGC output	15
2.6	Simulated response of AGC loop to large amplitude steps for various detectors . . .	16
2.7	Power monitoring of noise floor. <i>Source: Akos (2015).</i>	18
2.8	Optimal AGC Bin Distributions for a 2-Bit ADC	19
2.9	Left: Incoming signal, Right: ADC Bin Distribution	20
2.10	Raw AGC data	22
2.11	Example of different power measurements	25
3.1	SiGe receiver	28
3.2	NT1065 receiver	29
3.3	UBLOX M8 receiver	30
3.4	White Noise Narrowband AGC Response	31
3.5	White Noise Wideband AGC Response	32
3.6	Constant Wave AGC Response	33
3.7	Chirp AGC Response	34
3.8	SiGe AGC Response	35

3.9	NT1065 AGC Response	36
3.10	UBLOX AGC Response	37
3.11	SiGe high power frequency AGC response	38
3.12	SiGe low power frequency AGC response	40
3.13	SiGe hysteresis results	41
3.14	SiGe hysteresis mean and third deviation	42
3.15	Multiple SiGe module AWGN response	43
3.16	SiGe AGC power calibration curve	44
4.1	Toll road GNSS jammer detection concept	47
4.2	AGC monitoring equipment	48
4.3	Deployed monitor setup	49
4.4	Northwest Parkway AGC timeline	49
4.5	Northwest Parkway trigger daily frequency	50
4.6	Northwest Parkway trigger timeline	51
4.7	C/N_0 vs AGC for TEXTBAT and WAAS station data	53
4.8	Phone jamming/spoofing experiment setup	54
4.9	Phone: C/N_0 vs AGC for TEXTBAT and WAAS station data	55

Chapter 1

Introduction

In the past couple of decades global navigation satellite systems (GNSS) have become an integral part of modern society. Their ability to provide users with accurate navigation and time information anywhere on the surface of Earth has not only allowed for precise positioning, but has also enabled a plethora of dependent technologies and services. Whether it's circumnavigating a Boeing 737 about the Pacific or finding the quickest route to you're child's soccer game, contemporary users have become accustomed to having uninterrupted access to this invaluable resource. It is for these reasons that it becomes paramount to ensure that GNSS spectrums' remain clear of any interference, a denial of service could result in the harm of materials and/or personnel. This simple fact has led to the development of higher performance antenna, receivers, and satellites, which work to counteract all types of interference be it environmental or artificial. The latter is quickly becoming a larger source of worry for users as the increased availability and simplification of intentional interference devices such as jammers and spoofers is also occurring. In addition to this there have been a number of globally recorded events where GNSS has been intentionally denied or altered. Therefore there is an ever present need to detect, localize, and counter these nefarious sources of interference.

1.1 Applications of Automatic Gain Control in GNSS

As stated previously the first step to countering any sort of intentional interference is to detect the presence of such a device. This warning would allow the system or user to make an

informed decision on whether or not the position and time solutions produced by the receiver can be trusted. Engineers from across the globe have created a wide spectrum of potential solutions to this particular problem. One such set of solutions is to add additional hardware to the receiver, an example of this would be the addition of a power meter which can monitor the desired GNSS spectrum for any abnormalities in the power of the band. There is also the idea of using software solutions such as an analysis of the signal in the frequency domain using fast Fourier transforms, or the evaluation of common solution metrics such as Carrier to Noise density ratio (C/N_0). Each set of solutions has pros and cons which can make them individually more applicable in certain scenarios, yet they all share one or two common problems. First, they either require some physical or computational addition to the receiver which can potentially have the effect of abandoning pre-existing receivers. Second, the complexity of the solutions is high thus requiring more time to implement globally.

These issues have led a number of researchers to examine alternative methods of detection which can take advantage of the hardware found in most GNSS receivers. One particularly interesting solution is the use of the automatic gain control (AGC) circuit in the front-end of the receiver. This circuit has been a historical staple in all types of radio wave receivers since its purpose is to dynamical adjust the incoming signal gain such that the proceeding front-end components can work with a stable output despite how the initial signal may have been attenuated by the environment. Since GNSS receivers suffer from the same issue they have also been equipped with this circuit in some fashion which solves one of the problems given previously. The question is how can AGC be utilized to detect the presence of a jammer?

The idea presented by researchers is to use the feed-back voltage of the AGC circuit as a metric to determine the power of the incoming signal. Since GNSS signals are below the noise floor the overall incoming signal power should be relatively low and constant. When a jammer is used to interfere with the receiver it outputs a signal in the band which is more powerful than the original thus masking it. If the AGC were able to showcase this increase in power than a flag could be posted and the jammer would have been successfully detected. From the surface this solution

seems promising since it doesn't require additional hardware and it only requires the comparison of the real-time AGC voltage and some predetermined threshold values. The current catalog of academic papers on the subject have mostly been focused on showcasing the effectiveness of AGC in its anticipated solution space, this being a receiver exposed to a known jamming source. The results suggest that AGC behaves in a predictable one-to-one manner. Specifically if the incoming jamming power is increased then there is an decrease in the AGC metric and inversely if the power is decreased then there is an increase in the AGC metric. The exact shape and magnitude of the relationship between the power and AGC metrics can vary between the receivers but the overall inverse nature of the two is present in all.

1.2 Thesis Objectives and Problem Statement

With this in mind the research for which I am a part of has been attempting to implement AGC metrics into various counter interference systems as a method for detection. The data collected from the testing of these systems have shown results which conflict with the previous conclusions. For example there have been incidents of the AGC metric increasing randomly in nominal settings, which has never been presented in academic papers. After further literature research it became apparent that there is a lack of an in-depth characterization of the AGC metric with respect to general electromagnetic signals. Perhaps there are specific signal characteristics which can create previously unknown responses, and if so what is the effect and to what extent does it compromise the capability of the AGC metric to be used as a way to measure the received power?

This question gets to the primary objective of this research, which is to evaluate the validity of the claim that the AGC metric is essentially a psuedo-measurement of the incoming signal power. By answering this question it will be possible to state in what situations can AGC be used to detect the presence of a jammer or spoofer. Therefore this research will be performing the previously mentioned in-depth characterization of AGC by examining how a set of receivers responds to various types of input. Also there will be a section devoted to the application of the power findings with the intent of detecting spoofing and jamming attacks. The results of such test

will also help verify the preceding concept. Lastly, I will recommend to the GNSS community that there should be an effort made to standardize the AGC metric.

1.3 Thesis Overview

The thesis is organized as follows:

In Chapter 2 an overview of GNSS, GNSS intentional interference, and AGC will be given first. The basic principles and the architecture behind GNSS will be given in Section 2.1 and this will be followed by the methodology used to interfere with GNSS signals and how we counter them in Section 2.2. In Section 2.3 an explanation will be provided on the general purpose and operation of AGC and how it has been implemented in GNSS receivers. This is followed by Section 2.3 where we review some of the previous work done on the subject of using AGC as a method for power determination and even more specifically as a method for detection of intentional interference. Finally, Section 2.5 briefly discusses how a true power measurement is performed and what it represents.

Chapter 3 forms the first main section of the research. Here the results from a number of experiments is examined to complete the primary objective of characterizing the AGC metric in GNSS applications. This begins with a quick overview of the AGC power detection theory in section 3.1. Section 3.2 gives a brief overview of the receivers used throughout the research. Section 3.3 provides the results of the AGC response experiments and this is followed up by section 3.4 which summarizes the results of the experiments with respect to the original goal of this research.

Chapter 4 serves as the second main half of this research. This chapter takes the lessons learned from chapter 3 and showcases how they can be engineered to serve as detectors of GNSS jammers and spoofers. Section 4.1 goes over an automotive jammer detection application while section 4.2 covers a mobile phone method of detecting both jamming and spoofing attacks using AGC and one additional metric.

Chapter 5 summarizes the findings of the research and provides final conclusions.

Chapter 2

Background

2.1 GNSS Overview

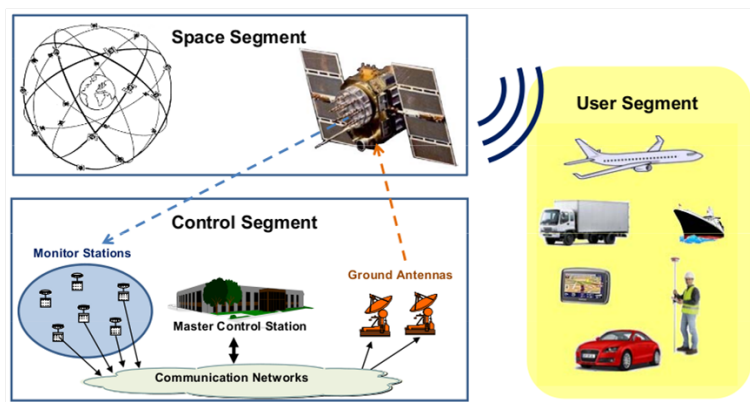
Global Navigation Satellite System (GNSS) is the generic term for a satellite constellation which provides users with accurate, continuous, world-wide, three-dimensional position and velocity information. As of writing this research there are four operational systems: United States' Global Positioning System (GPS), Russia's Global'naya Navigatsionnaya Sputnikovaya Sistema (GLONASS), European Space Agency's Galileo positioning system, and China's BeiDou Navigation Satellite System (BDS)[7]. Each of these systems have different hardware and software specifications, but they all utilize the same underlying principles which have defined radionavigation for the better half of a century. For the sake of a brief and concise explanation we will only be discussing GPS in order to demonstrate how GNSS works, but it should be noted that the other three systems operate in similar manners.

2.1.1 GNSS Methodology

To begin GPS has an architecture which is pieced into three segments: ground, space, and user. The space segment nominally consists of a 24 satellite constellation which is arranged in six orbital planes with four satellites in each plane[7]. The ground segment is made up of a worldwide network of ground control stations that monitor the health and status of the satellites. In addition to this the stations are also responsible for uploading commands, navigation, and ephemeris data to the satellites. Finally, the user segment is simply a receiver which is responsible for capturing


the incoming GPS signal and determining a position solution.

Figure 2.1: GNSS architecture Adapted from eesa



One of the most important requirements of any GNSS is that it has to be able to service an infinite number of users with only 24 satellites. This is accomplished by utilizing the concept of one-way time of arrival ranging. This process starts with the satellites which continuously transmit a carrier signal that has been modulated with navigation data and satellite specific pseudo-random ranging code. This code was specifically chosen due to its good correlation properties thus making it easier to acquire a signal. This signal is then picked up by the users receiver which decodes the navigation data in order to determine the location of the satellite at the time of transmission. The ranging code is used to find the travel time of the signal. The product of this time offset with the speed of light provides a range from the receiver to the satellite. In a three-dimensional space this creates a sphere of potential position solutions. Due to the fact that the receiver and satellite clocks are not synchronized a total of four satellites is needed to solve for a three dimensional position solution via trilateration[7]. The figure below showcases the geometry of this problem,

Figure 2.2: Trilateration, Adapted from eesa



2.1.2 Civilian Applications of GNSS

GPS itself was originally developed for the United States military with the intent to provide its service members with accurate uninterrupted global navigation data[7]. Though the ease at which the user can implement GNSS into various applications has made it a valuable global asset to multiple civilian sectors. In 2018 product sales and services for GNSS were \$290 billion and current estimates expect that number to grow to \$310 billion. This enormous market is a product of the large number of diverse GNSS applications that have been developed over the past couple of decades. In this section we will briefly go over some of these applications in order to get a better appreciation of how important GNSS is to the modern World. One of the most common applications of GNSS is a location-based services (LBS). These services take in the user's location data and outputs information that is applicable to the current location. Some large scale commercial examples of this would be General Motors' OnStar service, Google's Maps, and Uber. Anyone from a soldier to a grade school student can now Emergency service providers have benefited greatly with the integration of GPS. In the United States dispatchers can use the forwarded information of an emergency call from a GNSS enabled vehicle to accurately deploy resources. Once on the road automatic vehicle location systems (AVLS) can be used to manage fleets of vehicles in real-time. This idea has been expanded to commercial maritime communities which can use GNSS to greatly

decrease their transit time and thus lower operational costs. The agriculture sector has implemented GNSS into their architecture in order to cultivate large fields of crops via semi-autonomous harvesters. One of the more unexpected uses of GNSS has been recreational. Outdoorsmen such as hunters, hikers, skiers, climbers, fishermen, etc use handheld GNSS receivers to monitor their position when participating in their respective activity. American software development company Niantic has developed multiple games that utilize a mechanic called *geocaching*, which maps real world locations to virtual objects in the game. Their most popular game called Pokemon Go brings in around \$2 million a day from its user base and this is all possible due to the highly accurate real-time tracking capabilities of GNSS. Without GNSS an idea such as this would almost surely never make it past the drawing board since the burden of localizing the user base would fall entirely on the company itself.

As mentioned in the previous section a receiver requires four satellites to calculate a position solution, because the inexpensive quartz clock used most often in receivers will have an offset with respect to the atomic clock present on each satellite. This time solution gives the user access to precise time measurements anywhere in the World. This capability is applied in the banking sector where transactions are timestamped with the highly accurate GNSS solution. The energy sector also uses a similar method in order to transfer power between energy plants and to perform phase synchronization throughout the entire power grid. Countless number of networks rely on time synchronization protocols which pull time from GNSS and land based atomic clocks.

Perhaps the single largest and most important application of GNSS in the modern World is in the Aviation sector[7]. With the continuous global coverage of GNSS aircraft are able to fly directly from any single destination to another. Prior to GNSS pilots estimate their position by propagating their velocity with a known initial state a method often called dead reckoning, or by using radionavigation offered by land based stations. This limitation was very apparent in trans-oceanic flights which often found themselves in areas that were out of range of these stations. This change in the aviation paradigm greatly increased crew and passengers safety while also giving the industry access to new routes for either the purpose of increased coverage or decreased fuel cost.

Another under appreciated change GNSS brought to the aviation sector was the advent of the Wide Area Augmentation System (WAAS). WAAS basically collects GPS data from a number of ground based stations which then send their data to a master station that then creates an augmentation message which contains a correction for the GPS signal. This message is then broadcasted on the GPS L1 band via communication satellites so that software compatible receivers are able to apply these corrections to their GPS navigation solutions thus providing the user with solutions that accurate to a meter. Using this method it is possible to fly aircraft completely autonomously from takeoff to landing without the need for an additional DGPS connection.

2.2 GNSS Intentional Interference Overview

The big take away from the previous section is that nearly all aspects of modern society have some sort of dependency on GNSS. With this statement in mind it follows that any sort of intentional denial of the service could have major consequences. There are a large number of reasons for why an individual may want to alter or deny the GNSS solutions provided to a user: manipulation/falsification of transaction timestamps, privacy and autoimmunity, sabotage, etc. Therefore it makes sense to look at what are the possible weaknesses in GNSS and how do perpetrators exploit them. At first glance the obvious potential problems are physical failures of the hardware and corruption of the transmitted signal. The first problem is often not an issue since the satellite constellations and ground segments are extremely difficult to access. The latter suggestion is far easier to pull off since the culprit only has to influence the signal captured by the victim's receiver. This research will examine the intentional interference methods of jamming and spoofing.

2.2.1 GNSS Jamming and Spoofing

Jamming is the act of directing an electromagnetic wave at a user's receiver. This is often done in a brute-force manner where the magnitude of the false signal is many magnitudes more powerful than the true signal. This effectively masks the true signal and this can either degrade the

receivers performance or prevent it from ever acquiring one. The act of jamming is nothing new and has been around since the advent of wireless communications. Unfortunately due to the fact that GNSS signals are transmitted from Medium Earth Orbiting (MEO) satellites which causes the signal to lose power from both free space and atmospheric losses. For example the GPS L1 (SPS) signal strength can often be as low as -130.00 dBm when it reaches a receiver. For comparison a weak Wifi signal strength would be classified as -70 dBm. It is this low terrestrial power value of GNSS systems that make it particularly easy to jam. The figure below shows an example of two GPS jammers which were purchasable online at the time of writing this paper.



(a) GPS jammer for 12V car outlet



(b) GPS jammer for 5V USB port

The first jammer plugs directly into the 12V outlet of a vehicle, while the second utilizes a 5V USB port. Both of these run on relatively low outputting power sources, but offer up to 10 m of coverage according to their seller. There have been multiple instances of citizens using these exact types of jammers, most notably was a truck driver in New Jersey. From 2010 to mid-2012 a New Jersey truck driver had operated one of the car compatible jammers in order to block the GPS receiver on his company owned truck. But unknown to the driver his jammer was also causing issues

for new ground-based augmentation system (GBAS) in the Newark Liberty International Airport. On August 4, 2012, Mr. Bojczak's vehicle was identified as the source of GPS interference for the Newark airport and was fined \$31,875. It took the Federal Aviation Administration (FAA) and Federal Communications Commission (FCC) from March 2009 until April 2011 to locate another trucker who used a similar device on the New Jersey Turnpike.

Another, more nefarious form of GNSS intentional interference is a method called Spoofing. The act of spoofing consists of tricking a receiver into outputting a false estimate of time and position by transmitting a GNSS signal which is identical to the ones coming from the satellites. The most simplistic method of spoofing employs a GPS repeater and is called meaconing. This approach utilizes a cable, amplifier, and two antenna. The first antenna is setup in the open sky and is tasked with capturing the true signals for that particular position. The signal is then carried via cable to the second antenna and amplifier which output a more powerful version of the signal towards a user's receiver. The more powerful re-transmitted signal is picked up by the receiver which then outputs a position estimate that is of the first antenna of the meacon. More complex forms of spoofing can replicate the GPS signal profile that a specific receiver is experiencing and slowly surpass that signal with simulated signals thus controlling the time and position solution given by the victim's receiver. Despite being more complicated the consequences of spoofing could be far more serious than jamming. For example in 2017 at least 20 commercial ships sailing the Black Sea were reporting that their GPS receivers were stating that their ships current positions were more than 32 kilometers inland. This incident caused prevented all of the affected ships from effectively navigating away from the port, and the cause for the incident was speculated to be a spoofing attack from the Russian Federation.

2.2.2 Current Methods of Detection

In order to counter the effects of GNSS jamming and spoofing it becomes necessary to have the capability to detect whether or not the incoming signal is being jammed or spoofed. This knowledge can then be used to either notify the user not to trust the time and position estimates

or to take a specific action to negate the effects. One form of jamming detection is through the use of power measurements either through a complementary power meter, C/N_0 values, or frequency domain analysis using fast Fourier transforms performed on the IF data[11, 6]. The first is often the one most commonly used when it comes to modern day enforcement. A qualified technician or automated system can evaluate the incoming power on the GNSS band and determine if the incoming signal is interference by employing a threshold. The monitoring of C/N_0 values is done in the same way, when a jammer is turned on the noise value increases which in turn decreases C/N_0 [11, 6]. Though the same phenomenon can occur if the incoming signal simply gets weaker, this isn't a direct indication of power in the band. The final option of frequency domain analysis follows the same principles as the power meter, but it acquires its information from the evaluation of the digitized IF data used by the receiver.

When it comes to spoofing the solutions are often more complex as the perpetrator can match power with the true signal and still successfully trick the victim's receiver. Current forms of spoofing detection include the use of additional sensors such as IMUs. The two velocity solutions from the GNSS receiver and IMU can be compared to determine if the GNSS solution is incorrect. Another approach is to implement a Kalman filter which can fuse the data from a number of sensors not necessarily measuring the position and velocity in order to determine whether the GNSS receiver measurements are outside of the current state covariance. Finally, a analysis of the received signal in the correlation domain can be used to determine the authenticity[11, 6, 3, 1]. But once again all of these solutions require some form of additional hardware and/or computation power which is not native to modern receivers.

It is for this reason that numerous papers on the subject began to focus on the utilization of AGC to detect various forms of interference. The basic approach is to essentially assume that the AGC can be treated as a power meter[3, 6, 8]. Then by knowing the power of the incoming signal it is possible to once again set a threshold at which the receiver would stop trusting its own estimates. This methodology appears to be an ideal candidate for detecting jammers and non-power matching spoofers since a great number of receivers already include an AGC thus no

additional hardware and the computation requirements would be virtually non-existent since only a single value read and comparison are needed. Despite the increased interest there has been little to none in depth characterization of AGC when it comes to measuring the power of an incoming signal. As stated previously that is the purpose of this research and to begin the characterization of AGC an overview will be provided in the next section.

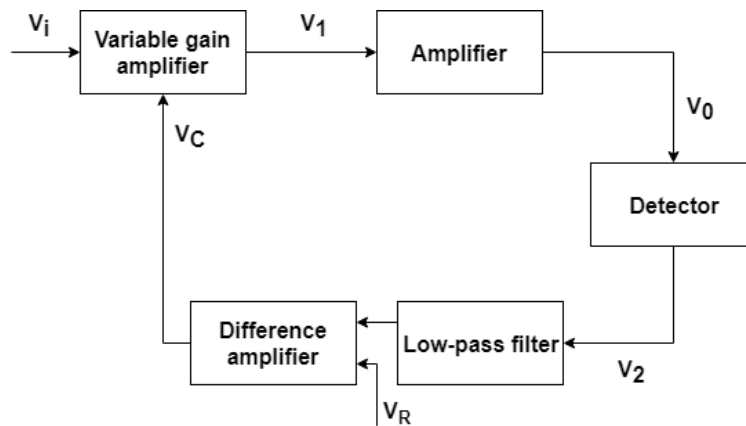
2.3 Automatic Gain Control Overview

As discussed earlier the detection metric being proposed in this research is the relative value of the Automatic gain control (AGC) circuit present in the front-end of most GNSS receivers. To begin let's discuss the purpose of AGC and how it works, then we will discuss how it is implemented in GNSS receivers and how it can be used to detect interference.

2.3.1 General Automatic Gain Control Background and Methodology

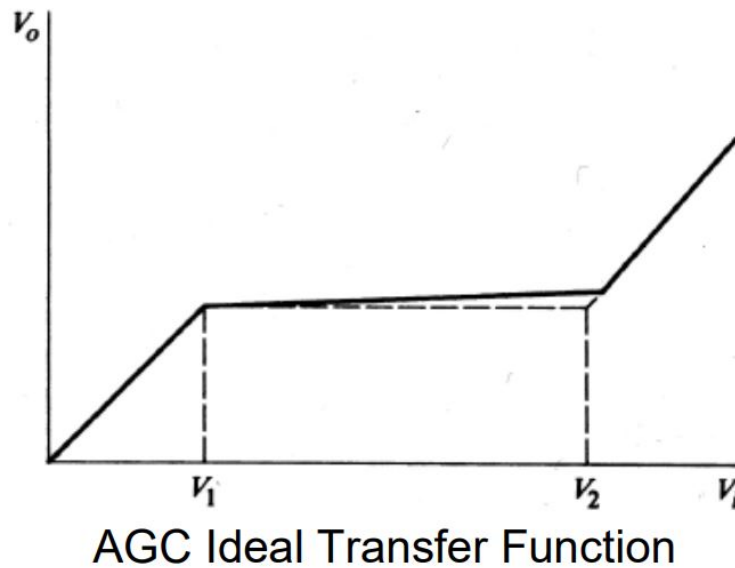
AGC was originally implemented in early radios for the purposing of countering fading propagation a phenomenon defined as slowly occurring variations in the amplitude of the received signal[3]. This constant change in outputted signal can cause large amounts of degradation in the performance of the front-end. Thus the purpose of AGC is to continuously create adjustments in the receiver's gain in order to maintain a relative constant output signal for the circuits which follow the AGC itself. Due to this capability AGC circuits can be found in a number of systems where the amplitude of the input signal varies over a wide dynamic range e.g. AM radio receivers, radar, VCRs, voice-operated gain-adjusting devices (Vogad), telephone communications recorders, and tuners. One of the simplest implementations of AGC would be to use a variable attenuator between the input and output, though this method is often avoided since it decreases the signal to noise ratio (S/N). Therefore a more typical AGC circuit is shown in the diagram below,

Figure 2.4: Typical AGC Block Diagram sourced from Iulian Rosu



as you can see the circuit has a loop or feedback system which is comprised of a forward controlled gain stage, feedback gain stage, and a signal comparison stage which is responsible for generating the differential error signal. This distribution of gain control over multiple stages is done such that the gain in the later stages is reduced first rather than the ones closest to the input. This ensures a high S/N because the later amplification stages are for the final IF while the earlier ones amplify the RF and initial IF signals. Looking back at the diagram we see that the input signal is first amplified by a Variable Gain Amplifier (VGA), which is a circuit that has a variable gain that is controlled by an external signal V_C . The amplified signal V_1 is then further amplified by the second stage in order to have an adequate level for the detector. This signal now dubbed V_0 passes into the detector which identifies various parameters about the signal such as the amplitude, carrier frequency, and index of modulation. The outputted signal V_2 is then filtered and then compared with a reference signal V_R . This comparison generates the control voltage V_C which is responsible for adjusting the gain of the VGA. The ideal AGC output for a linearly increasing input signal is given in the figure below,

Figure 2.5: Ideal AGC output sourced from Iulian Rosu




At low input signals the AGC is often disabled or limited at some maximum gain, this produces a linear response to the input. At some point the input signal will reach a threshold value designated V_1 at which point the AGC will begin to operate and will maintain a constant output level, shown as the dashed line. In reality there is a small linear rise in the output which is determined by the quality of the detector and is called the AGC slope. Finally at some upper threshold value the AGC is saturated and can no longer decrease in gain value, thus the output returns to the initial linear output relationship.[7]

In reality the AGC cannot instantaneously adapt to immediate changes of the amplitude of the input signal due to the delays created in processing. This limitation is often a positive effect since a fast response time is not desirable because it would make the AGC circuit overly sensitive to noise in the signal. There are two terms which define this response "attack time" and "decay time", which are defined as the AGC circuit's response to increases and decreases in input amplitude. For modern receivers there are often four types of detector used for AGC: envelope, square-law, true-RMS, and LOG[7]. An envelope detector has an output voltage that is

proportional to the magnitude of the instantaneous RF input voltage. Given a properly filtered signal with a tight bandwidth this method provides a simplistic solution, though this condition is not applicable to all applications. Square-law detectors output a signal that is proportional to the square of the instantaneous RF input voltage. The result of this method is that the loop's equilibrium average output power independent of the input waveform, but this also results in large responses to large changes in input amplitude. True-RMS detectors combine a square-law detector, low-pass filter, and square-root function. Finally, the LOG detector creates an output proportional to the logarithm of the RF input voltage. Each detector has pros and cons when it comes to various types of input signals. The following figure showcases the response times for the four detector types when given a large amplitude variation (Note that LINBNV is the envelope detector).

Figure 2.6: Simulated response of AGC loop to large amplitude steps for various detectors. Courtesy of Analog Devices, sourced from Iulian Rosu



These simulated results are for clean signals only and thus do not show the real life per-

formance of each detector given noise. Despite this it gives the reader a general idea of where each method is strong, particularly the the log detector gives the quickest response to large abrupt decreases in input level because the logarithmic curve has a very steep slope for low inputs, while abrupt increases create very slow responses. The opposite is true for the square-law detector who utilizes a near inverse[7]. With the following information in mind let's now move on to why and how AGC is implemented in GNSS receivers as that is the specific application outlined in this paper.

2.3.2 Automatic Gain Control Implementation in GNSS [3, 1]

As discussed in the previous section AGC is employed in circuits which desire a constant amplitude input signal but have an input signal that is varying in amplitude over time. GNSS receivers fall well within this subset due to the unique terrestrial signal characteristics of GNSS signals. For example the received GPS signal power on the surface of Earth using a traditional hemispherical RHCP antenna is well below the thermal noise floor P_N . This value can be calculated using the following equation,

$$P_N = kT_A BW \quad (2.1)$$

where k is Boltzmann's constant, T_A is the effective antenna temperature, and BW is the bandwidth[3, 1]. In addition to this there is noise generated by the first stage front-end components of the receiver which can be found using Eqn. 2.2,

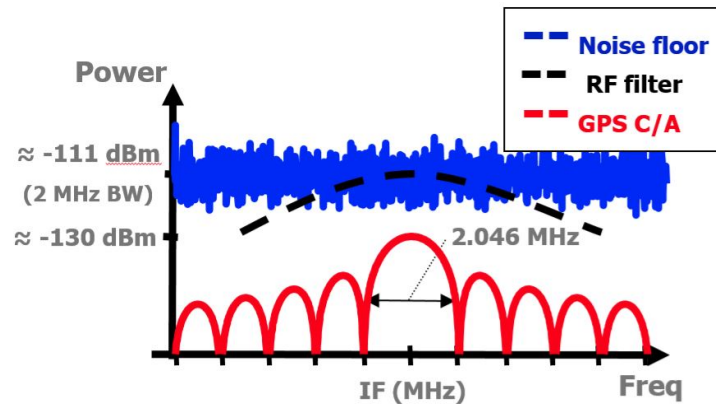
$$P_{N,R} = kT_R BW \quad (2.2)$$

the only difference is the effective temperature is swapped for that of the receiver first stage front-end components which is derived from Friis's formula and the stacking of the front-end components[3, 1]. By summing these sources of noise we get the total noise present in the input signal,

$$P_{N,Total} = k(T_A + T_R) BW \quad (2.3)$$

the result of this analysis is that the total noise comes out to be around -110 dBm which is greater than the -130 dBm received power of GPS L1 C/A[7]. In technical terms this means that the received signal is below the noise floor, meaning that noise is the dominant signal source in the input signal. A visual representation of this is given in Fig. 2.7.

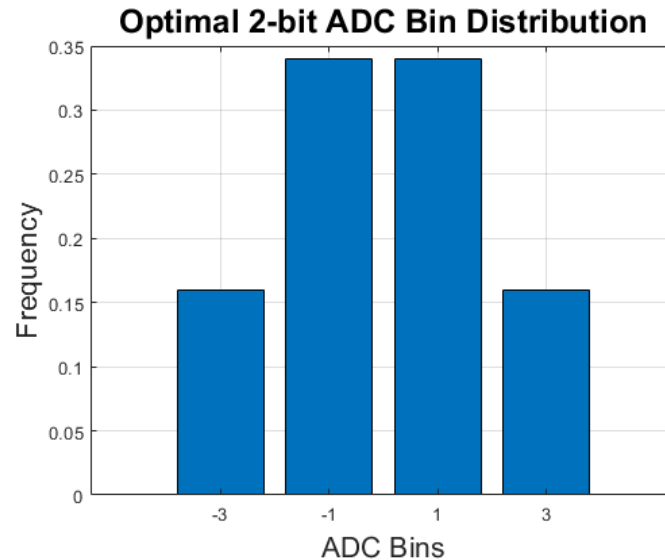
Figure 2.7: Power monitoring of noise floor. *Source: Akos (2015).*



With such input the total noise and sources of interference cause the amplitude of the signal to vary greatly overtime w.r.t. the GNSS signal dynamic range. Apart from the noise factor, commercial GNSS receivers are often sold with the intent that any number of antenna, amplifiers/attenuators, and cables could be used, which all result in a varying levels of gain provided to the input signal. Therefore an AGC circuit is placed in the front-end of the receiver in order to optimize the gain of the front-end such that the amplitude of the incoming signal utilizes the entire range of the analog-to-digital converter (ADC). Now since the input signal is expected to be thermal noise a typical detector like the ones discussed in the previous section cannot be used, since they would be acting on the amplitude of the noise thus potentially decreasing SNR. Therefore the AGC circuits found in GNSS are driven by the noise environment rather than the GNSS signal power itself. The basic principle to how this works is to go forward with the assumption that the incoming noise is white noise and therefore its incoming signal will produce a Gaussian distribution

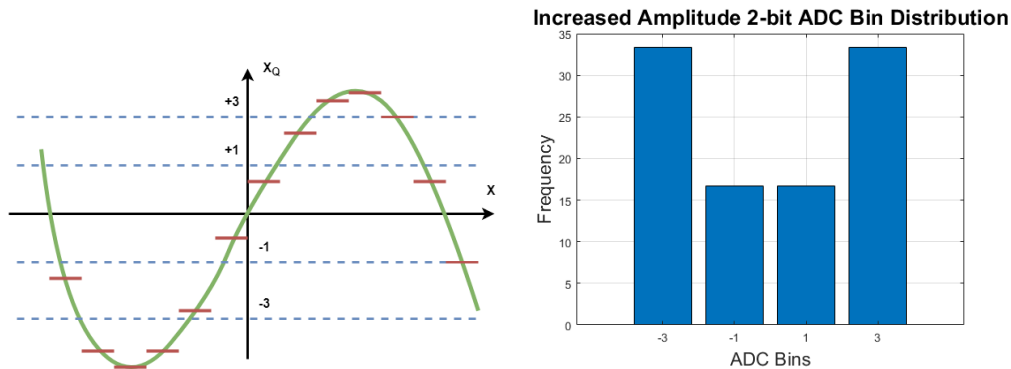
when quantized[3, 1]. Therefore the detector in a multi-bit receiver's AGC attempts to adjust the gain such that the sample distribution is Gaussian thus minimizing digitization losses. An example of an optimal AGC with 2-bit ADC bin distribution is given in Fig. 2.8,

Figure 2.8: Optimal AGC Bin Distributions for a 2-Bit ADC

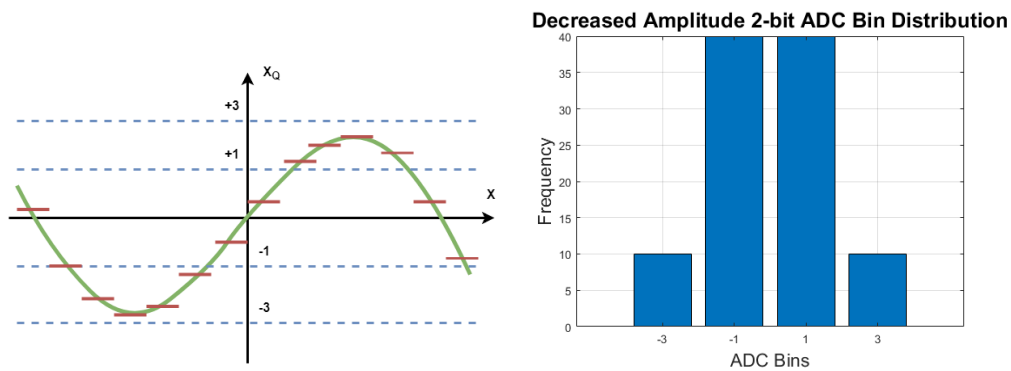


here there are a total of four bins which are uniformly distributed and non-centered. Such a configuration operating at optimal conditions produces two center bins with 67.3% of the total bin count, while the outer bins contain the remaining 32.7% of samples[3, 1]. Now ask what would happen if the incoming signal were to change in amplitude for what ever reason e.g. interference, different antenna, etc. The change in bin distribution is shown in the following figure,

Figure 2.9: Left: Incoming signal, Right: ADC Bin Distribution



(a) Oversaturated ADC



(b) Undersaturated ADC

here two cases are provided. The first showcases an example of when the ADC is oversaturated (the input voltage of the incoming signal exceeds the dynamic range of the ADC) which results in an inverted distribution. This saddle-like shape is a result of all of the quantization samples taken from the input signal, which in the figure are shown as red lines. Remember these are discrete and thus the implementation losses in the receiver are dependent on the sampling rate, precorrelation bandwidth, and the method chosen for quantization. In this case the AGC would decrease its gain to bring the signal back into the dynamic range of the ADC. The second case demonstrates the exact opposite, here an ADC that is undersaturated meaning that it is not

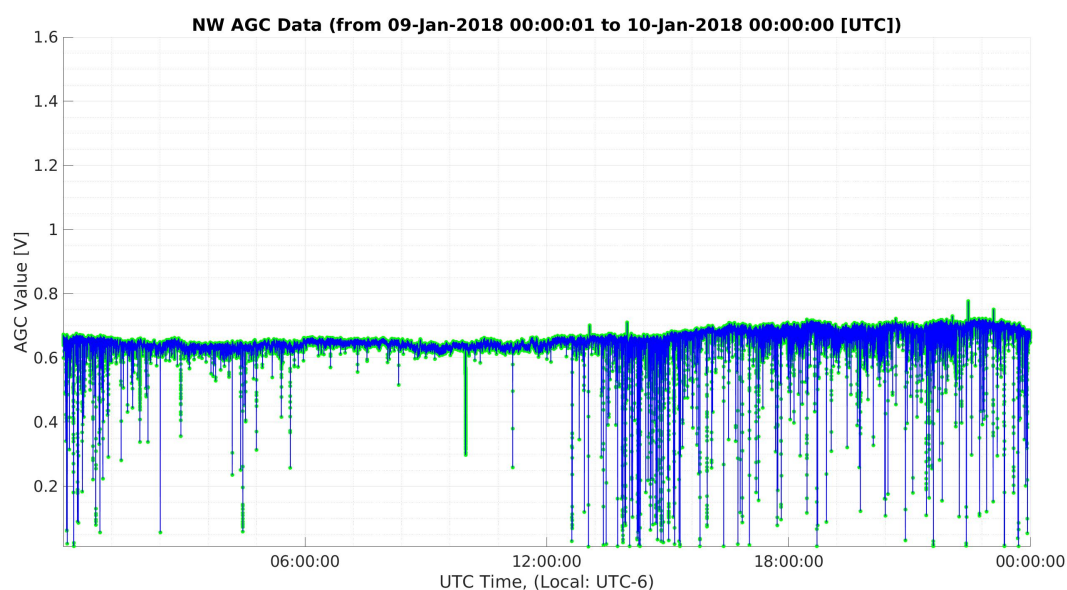
utilizing all of its dynamic range. This scenario would create losses in SNR and therefore the AGC would increase the gain in order to return the system to an optimal quantization ratio. It should be noted that there are analog methods of adjusting the AGC gain, such methods typically estimate the standard deviation of the noise and adjust the power of the VGA to match with that of a Gaussian. Though such methods do have issues when it comes to blanking, regardless this paper solely focuses on the digital implementation.

2.4 Previous Work

As stated previously the idea of using the values from an AGC circuit to detect whether or not a signal is experiencing interference is not an original idea of this paper. The concept was first brought up in the 2003 paper titled "Automatic Gain Control (AGC) as an Interference Assessment Tool" written by Frederic Bastide and others. The purpose of the paper was to evaluate the functionality of AGC for detecting interference in the GPS L1 and L5 band. The authors examined how a Novatel OEM4 receiver's AGC gain responded to an ever increasing AWGN and true GPS L1 signal. The experiment resulted in a large correlation between the AWGN power increase and the decrease in AGC gain. Another more in-depth experiment was performed to evaluate potential of using a Chi-Square test to see if the initial and current distribution are consistent and if not then interference would be prevalent. Another large finding that the paper made was that throughout there multi-day outdoor control testing they would see a daily 1 dB variation in the AGC gain. The source was originally attributed to some natural fluctuation of the noise floor caused by the daily change in temperature, but upon further investigation the source was found out to be the decreasing efficiency of the AGCs amplifiers caused by the increasing ambient temperature. The conclusions of the paper were: the AGC system is an accurate indicator of the noise environment of the receiver whose gain varies with respect to the present interference power and therefore can be a valuable tool to detect interference, with more bins it is possible to apply a Chi-square test on the bins to detect interference, and overall the equipment works for L1 and L5. Apart from this the authors also made an interesting note that during their testing they noticed strange short pulsed

interference signals in about 0.5% of their data. The figure below displays one of these events, note how the interference caused short peaks and valleys in a data set that I collected during on of my own experiments.

Figure 2.10: Raw AGC data



The previous paper has become the basis for the proceeding papers which have attempted to implement the AGC metric into various counter interference systems. The first example of this was the 2012 paper titled "A Low-Cost Monitoring Station for Detection Localization of Interference in GPS L1 Band". This paper outlines a simple monitoring station which comprised of a receiver, camera, and computer. The entire station was placed on the side of road such that when a vehicle would pass with a jammer the receiver's AGC value would drop triggering the camera in an attempt to capture the offender. A known vehicle with a jammer on a circular test was used to test the system and in the end the results showed that the AGC metric was able to detect the incoming vehicle and thus was able to trigger the entire system.

Another brilliant paper was written by Dimc and Bazec (2017) titled "An Experimental

Evaluation of Low-Cost GNSS Jamming Sensors”. This paper compared the effectiveness of multiple metrics when used as jamming detectors. Specifically the paper examined two main types of metrics for detecting jammers. The first was based on the measurements provided by a COTS GNSS receiver, the second used the signal samples provided by a low-cost SDR front-end. The experiment performed had each metric exposed to a jammer and the resulting values were correlated to the true power of the incoming signal. The paper concluded that, ”Among the metrics based on measurements from a standard GNSS receiver, the average C/N_0 and the AGC count seem to be the most effective detection statistics since they directly depend on the distance between the hammer and victim receiver”. Further on when comparing the two the paper states that the AGC is superior since it does not suffer from the ”power ambiguity problem” that C/N_0 does not. The problem the paper was referencing has to do with propagation effects such as signal attenuation.

2.5 Power Measurement

For the final section I wanted to provide a quick overview of how and why we measure the power of an incoming RF signal. This is important since this paper intends to validate the idea that AGC metrics can be used to determine the power of the incoming interference signal. First power is defined as the work done power unit time and is typically given in watts. In the field of electronics a single watt corresponds to a single ampere of current across a potential difference (voltage) of one volt. To calculate power the following formula is used,

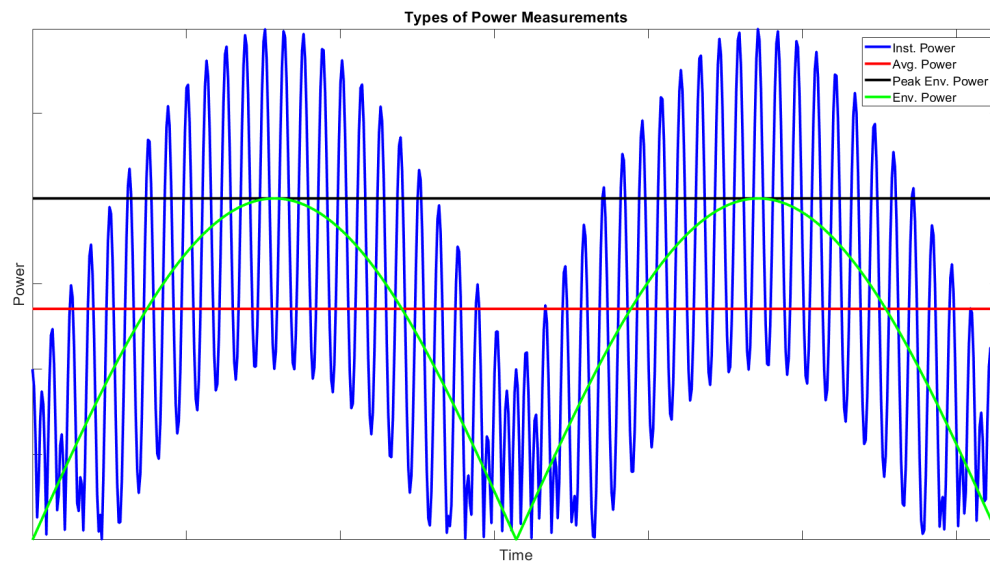
$$P = IV = \frac{V^2}{R} = I^2R \quad (2.4)$$

where P is the power measurement of the signal, I is the current, V is the voltage, and R is the resistance. With this equation in mind it seems obvious that if you want the power of an incoming signal then all you have to do is measure the voltage if you know the load of the circuit. You’d be right about that and this is entirely possible to perform on a pure DC circuit. The problems start to arise when you want to know the power of an analog signal such as an RF signal.

For signals such as these the voltage and current can vary with position along a lossless wire because of standing waves produced by the collisions of incident and reflected traveling waves. Also with signals that have high frequency components such as GNSS which are in the neighborhood of 1 GHz then the idea of measuring the instantaneous power becomes meaningless because of how quickly this metric could change. Therefore the best thing to do is to make a direct power measurement rather than measure the raw voltage of the incoming signal. To do this often thermocouples or a diode detector are used[7]. Thermocouples operate on the principle that the voltage generated between two dissimilar metals generates a voltage that is a function of their temperature. There is a direct relationship between the temperature and the RF power incident on it, therefore it is possible to determine the input power level from the voltage created by the thermocouples. On the other hand diode detectors rectify the signal, meaning they convert the AC signal to DC which can then be measured as it will have a constant voltage though changing current.

Both of the methods described above will provide the user with the instantaneous power measurement of the signal. As stated earlier this measurement is not very useful when trying to characterize RF signals. Therefore other measurements are used, the most commonly used power measurement is the average power. This method averages the power measurements over multiple modulations of the lowest frequency component in the signal. There is also the envelope power which is measured by averaging the power over a period of time that is large in comparison to the period of the highest frequency component, yet short enough to capture a period of the largest. The last measurement we will go over is the peak power measurement which is the largest average value over the measurement period. This can be applied to both the instantaneous and envelope measurements. The figure below showcases how this might look for a given set of power measurements.

Figure 2.11: Example of different power measurements



For the purpose of this paper we will be utilizing the average power measurement, therefore whenever making future references to a "power" measurement I will mean the average power measurement.

Chapter 3

AGC Power Characterization

As stated in chapter 1 the main objective of this research is to validate the claim that the AGC metric can be used to determine the presence of intentional interference and it does so by acting as pseudo-power measurement. The purpose of this chapter is to first explain how theoretically AGC can be used to detect interference and then second provide experimental data which showcases the actual response to various forms of interference. To begin let's discuss how the AGC circuit of a GNSS receiver can be used to detect the power level of the incoming signal.

3.1 AGC Power Detection Theory

The basic idea of AGC GNSS power detection is to relate the power of the input signal with the AGC metric. Going back to the AGC overview we learned that there are various ways to design the feedback loop. The simplest model creates a gain that is proportional to the incoming voltage and thus the incoming power of the signal with respect to the noise floor is given by the following equation,

$$P = \frac{\gamma}{G_{max} - G_{AGC}} \quad (3.1)$$

where P is the relative power of the input signal, G_{AGC} is the current AGC gain, G_{max} is the maximum AGC gain, and γ is the proportionality constant. Applying this to interference detection is trivial given that the noise floor will have a relatively constant power it is possible to detect if a drop in the relative power is seen or to say in a different way the AGC gain increases. The second

implementation attempted to drive the ADC output to a Gaussian distribution by adjusting the variable gain of the AGC. Given that this method could potential produce any relationship the formula used to express it is,

$$P = f(X_{AGC}) \quad (3.2)$$

the assumption here is that if the function $f(X_{AGC})$ is a one-to-one function then it is possible to calculate the incoming power by corresponding the given AGC metric X_{AGC} to an input signal power measurement. Looking back at Fig. 2.9 we noted that the theoretical relationship between the AGC gain and the incoming signal power should be inversely proportional.

3.2 Receiver Overview


One of the advantages of using AGC as a metric for interference detection is that it is readily available in a large number of COTS GNSS receivers, thus no additional receiver hardware is needed. Therefore this paper honors that idea by using the AGC measurements from three COST GNSS receivers and the purpose of this section is to provide a brief overview of each one and to highlight their differences.

3.2.1 SiGe Receiver

The first receiver is the SiGe GN3S Sampler (SiGe) which was co-developed by the GNSS Lab at the University of Colorado Boulder and the SiGe Company. It uses the SE4120 GPS L1 front-end chipset and can output samples over USB via the Cypress USB 2.0 controller. The SiGe has configurable settings including the sampling rate and data type (2-bit real or complex) GPS L1. As for the AGC metric it outputs a byte of data which represents the current AGC voltage. This receiver only outputs the raw IF and AGC samples and therefore does not produce navigation solutions. All post processing must be done using another front-end module or an SDR. The driver used to stream the data from the receiver to the PC was developed in the GNSS Lab and I was personally involved in adding new features and fixing bugs. the program itself is a multi-threaded

C++ program. One of the threads continuously requests and stores the data in a circular buffer on RAM. This is crucial since it allows the PC to quickly access historical data if needed for output to a file. Another thread is responsible for writing the data to disk, while all the other ones handle the times at which the data is written to disk. For example if the user wanted a certain size of data every X number of seconds then the program would have a thread created to manage that request. Finally, there is a trigger thread which is responsible for automatically writing the buffer to disk when the AGC drops below a user specified threshold.

Figure 3.1: SiGe receiver, source sparkfun



3.2.2 NT1065 Receiver

The second receiver is designated as the NT1065. This board was created by the company NT1065 and is capable of processing 4-channels in the bands of GPS/GLONASS/Galileo/BeiDou/NavIC/QZSS L1, L2, L3, L5, E1, E5a, E5b, E6, B1, B2, and B3. Similar to the SiGe receiver the

NT1065 only outputs the raw IF and AGC samples, thus no navigation solutions are generated. The board itself consists of the NT1065 chip which is responsible for outputting the 4-channel IF data. This data is given to a microcontroller that manages the USB 3.0 controller which finally outputs the samples to the PC. There is a large set of configurable options which are set by the uploaded register values. There is commercial software available to interface with the receiver over USB, but this software wasn't used since it did output the AGC data. In order to fix this an in-house program was created to output both the raw IF and AGC data to the PC's disk in a manner that was identical to the SiGe. Given the multi-band capabilities of the receiver there are multiple AGCs present in the receiver which are assigned to the GPS and Glonass L1/L2 spectrum. The AGC metric produced is a float which represents the current AGC level which can be converted to the power gain in dB by using the table given in the datasheet.

Figure 3.2: NT1065 receiver, source NTLabs



3.2.3 UBLOX M8 Receiver

The final receiver being used is ublox's M8. This multi-constellation receiver was developed by the company ublox and is capable of tracking three concurrent GNSS signals of the following types: BeiDou, Galileo, GLONASS, and GPS/QZSS. This receiver is one of the most popular GNSS receivers on the market and is used in a large variety of applications. In order to use the chipset an integration board is needed and for this paper the EVK-M8T module was selected. This evaluation kit allows the user to interface with the chip using I2C, SPI, and RS232 connections. Unlike the last two receivers the software used to read out the AGC metric was developed outside the lab by ublox themselves. The AGC metric in the case of this receiver is given as a percentage of the total available gain of the receiver. Although it should be noted that nowhere in the receiver's documentation is a conversion provided. Therefore this particular receiver is a bit of a black box.

Figure 3.3: UBLOX M8 receiver, source ublox



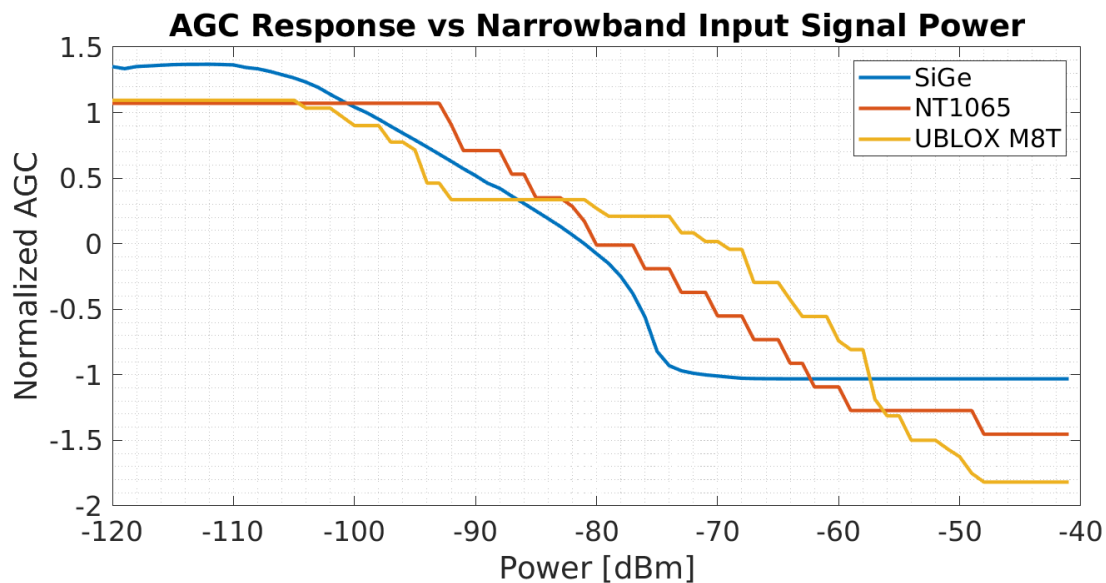
3.3 AGC Power Relationship Experimentation and Validation

Now with the four receivers chosen it is time to move onto the experiment performed to validate the AGC power theory.

3.3.1 White Noise Response

The first type of signal responses examined were for white noise signals. This type of signal is representative of the noise floor that a receiver would expect to see in nominal operation. When it comes to the NEAT there are two options for generating this type of output and that is narrowband and wideband. The difference between the modes is that wideband's signal contains a larger frequency range of components that make up the total signal. Now, let's discuss the AGC response to these white noise signal modes which is displayed in the two figures below,

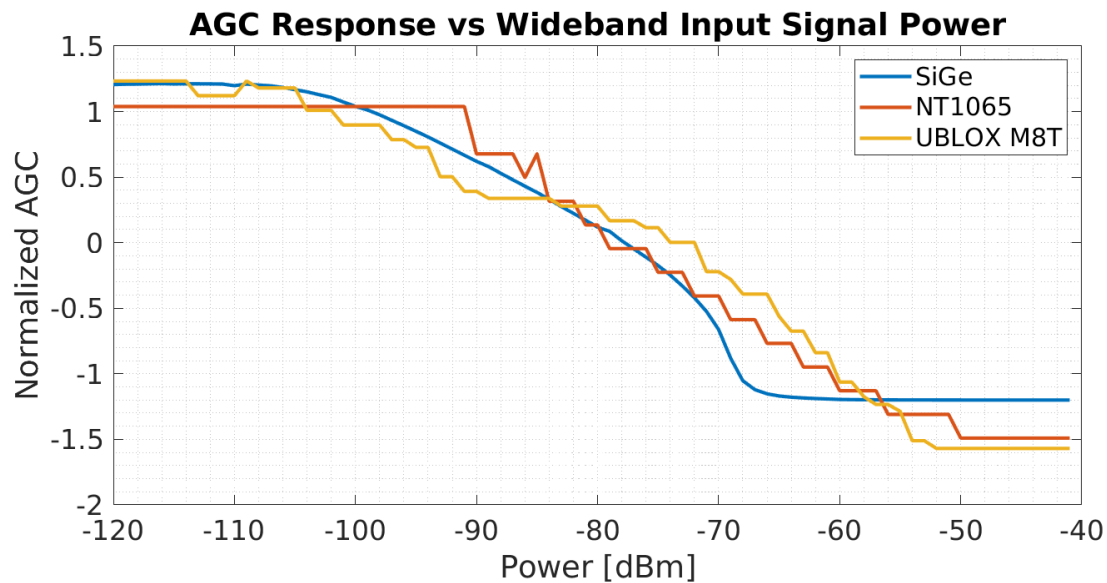
Figure 3.4: White Noise Narrowband AGC Response



here it can be seen that all three receivers followed the theoretical inverse relationship outlined in the previous section. Though it should be noted that the both lower AGC metric resolution of

the UBLOX and NT1065 are very present in this plot. Now let's look at the wideband results given in Fig. 3.5,

Figure 3.5: White Noise Wideband AGC Response



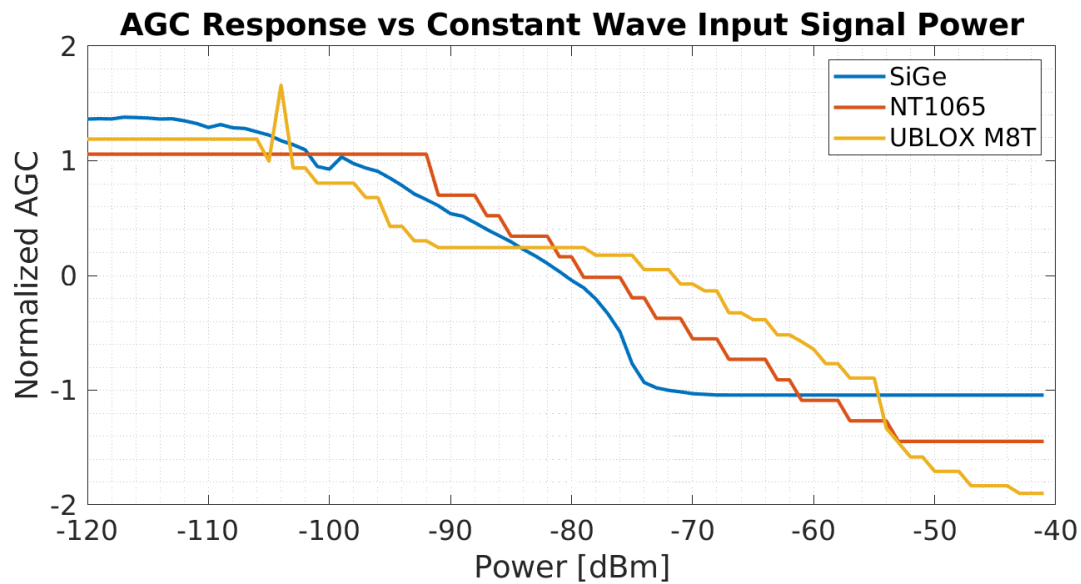
the first thing to note is that this plot shows a greater amount of agreement of AGC metric and the power of the incoming signal. The differences in the receivers seems to come from differences in resolution though once again the overall inverse proportionality holds.

3.3.2 Non-White Noise Response

The second set of signal types are that of the non-white noise signals. These signals are characterized by their deterministic output, meaning that the value of the signal is a function of time and thus a dependent variable. An example of this would be a square-wave where the current value of the signal is always one of two values and is only changed through the duty cycle. The NEAT can generate two of these signals. The first is called constant wave (CW), this signal is simple cosine output and can be thought of as an unmodulated carrier wave whos frequency matches that of the target's center frequency. The second wave form is called a chirp signal (aka sweep signal).

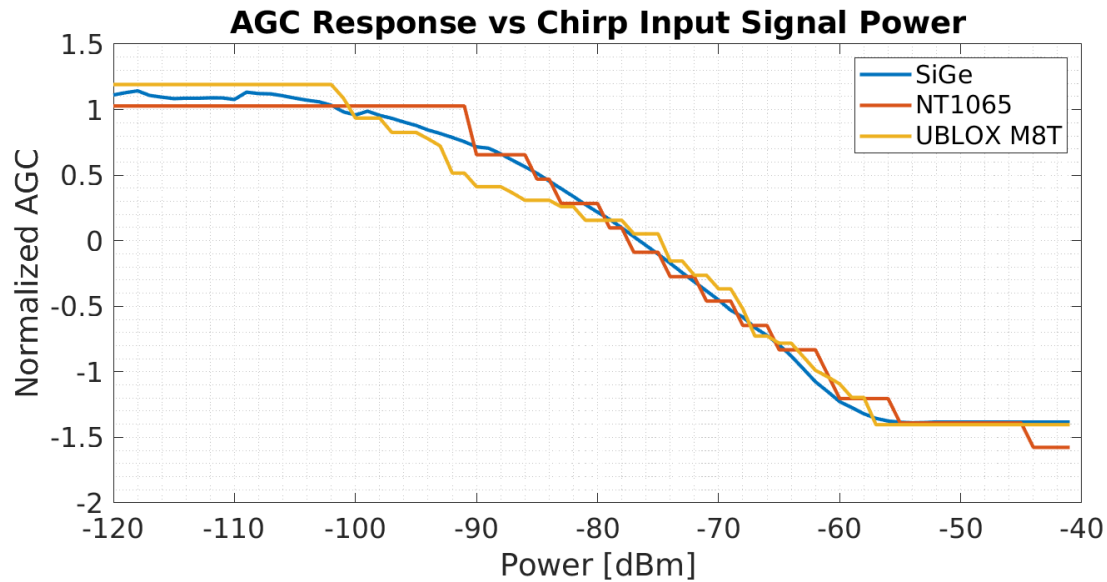
A chirp wave form is a signal in which its frequency changes over time. For this experiment the NEAT goes up and down a range of frequencies centered at the L1 carrier frequency. First, the results of the CW test are shown in the figure below,

Figure 3.6: Constant Wave Wideband AGC Response



here once again we see that overall the inverse AGC and incoming signal power relationship is still present in all three receivers. Next, let's examine the chirp response given in the figure below,

Figure 3.7: Chirp AGC Response

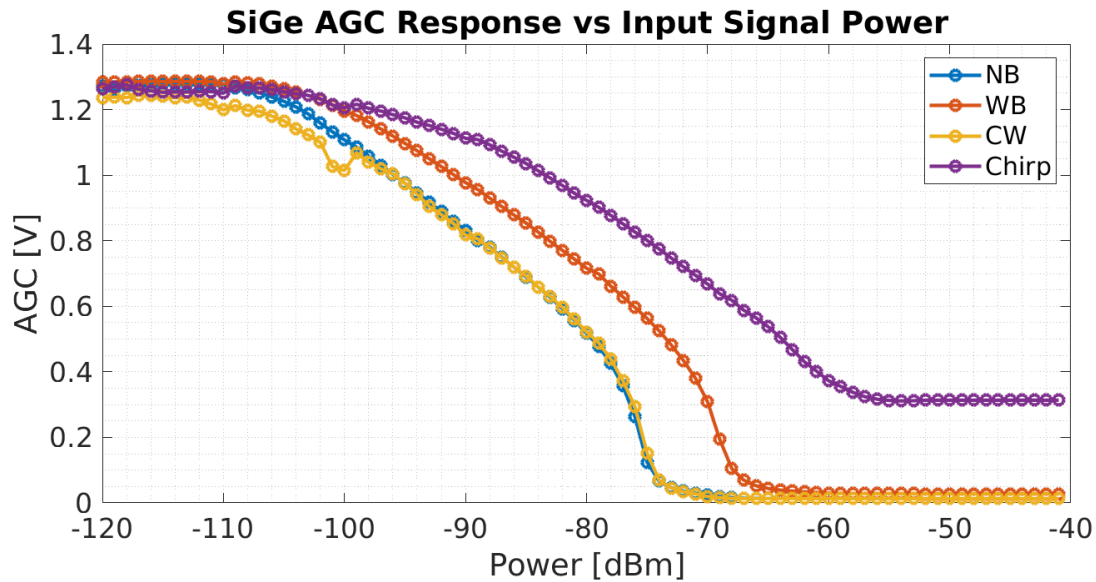


here the relationship still held for each receiver. Though it should be noted that this test showed the greatest agreement of normalized AGC metric across the three receivers. The vast majority of difference appears to come from the NT1065 and M8T's lack of resolution.

3.3.3 Receiver Specific Responses

From the previous section we were able to see that there does exist a strong inverse relationship between the incoming signal power and the AGC metric regardless of the signals characteristics. Now let's reexamine the results and compare how the AGC response of the individual receivers changed with the different types of input signals. To begin let's look again at the total SiGe receiver results,


Figure 3.8: SiGe AGC Response



note here that the true AGC metric in the case of the SiGe is the current control voltage outputted by the AGC itself. Now a few things are very apparent when looking at the results and the most obvious is that only two out of the four total signals gave near identical responses and these were the CW and white-noise nearband. These two showcased a greater sensitivity to the power of the incoming signal and thus hit the minimum gain about 10 dBm quicker than the wideband signal. The reason for this is obvious when it comes to the CW signal, which follows the structure of a cosine wave. Such a wave form produces a bin distribution that is fairly uniform when the peak of the wave corresponds to the maximum dynamic range of the ADC. From theory we determined that this would cause a drop in the AGC's metric and therefore as far as the AGC's detector is concerned the incoming signal appears to be more powerful than it actually is, running a quick simulation of this shows that this effect results in a 9% decrease in response (that is to say that for the same distribution or AGC response the CW signal is 91% of the power of a AWGN signal that gives the same response). The final thing to note from this plot is that the SiGe's AGC proved to be significantly less sensitive to the chirp signal. In fact the AGC didn't even hit

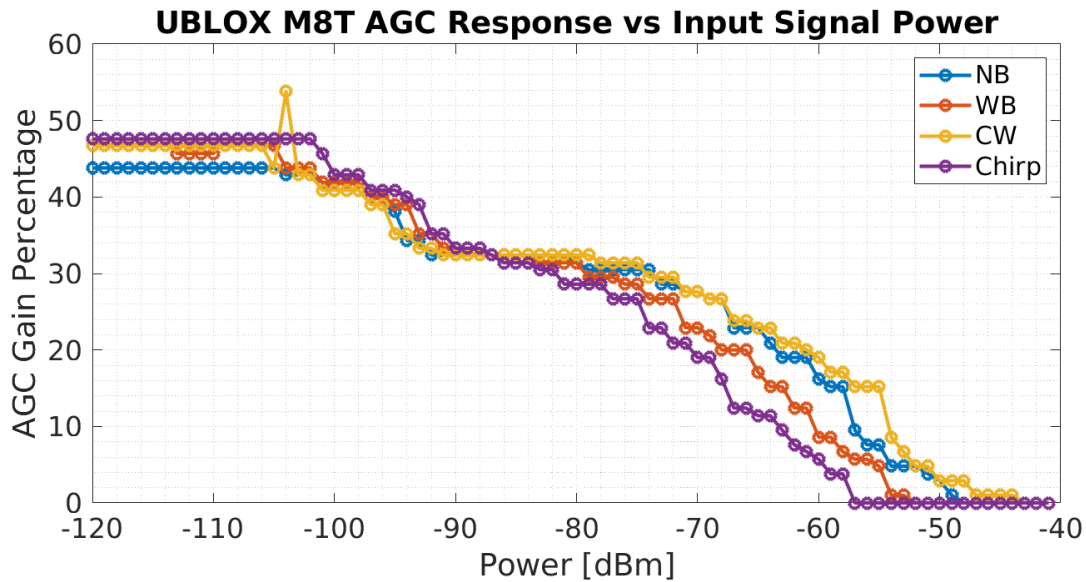
its minimum gain value in the range of power tested. The reason for this behavior is harder to explain as the specifics of the chirp signal outputted by the NEAT were not provided. But it is possible that the rate at which the NEAT swept across a frequency that wasn't filtered out prior to hitting the ADC was less than that of the ADC's sampling rate. Such an occurrence would create essential dead zones in the signal which resulted in more center bins and thus a lower AGC gain for that particular power output when compared to the other signals. Next let's take a look at the individual results of the NT1065.

Figure 3.9: [NT1065 AGC Response



The NT1065 displayed the greatest amount of uniformity in its response to signals. The reason for this could be either the resolution of the AGC was too low to capture the differences and/or the filtering on this receiver is less since it is a multi-GNSS receiver and therefore accepts a broader range of carrier frequencies. For whatever reason this result beautifully demonstrates a near linear relationship between the incoming power and the AGC metric. Finally, let's go over the M8T results shown below.

Figure 3.10: UBLOX AGC Response



The M8T had interesting results in that it showed characteristics that were present in the two previous receivers. To begin the inverse relationship is present throughout all of the signals, though the response is not identical for all four. Everything below -80 dBm is fairly similar but afterwards there is a large amount of variation in the AGC metric when it comes to the chirp and wideband signals. This is similar to the SiGe but the trend of the signals is reversed in this case with the chirp being the most sensitive and the CW being the least. The reason for this will remain a mystery given that this receiver is highly complex and thus it takes into account a vast range of other metrics to improve its performance.


3.3.4 Near Spectrum Response

Continuing the characterization of AGC, we will now perform a preliminary investigation into the response of AGC with respect to frequency. As explained in the previous experiment there is evidence which suggests that AGC metrics are dependent on the frequencies of the incoming signal. To evaluate this claim the following experiment was performed.

A CW signal generator was placed in line with the SiGe receiver and was set to output a

constant power signal. The frequency of the signal was then varied over a range of values which were uniformly distributed about the L1 carrier frequency of 1575.42 MHz. Two sets of data were taken and the first is shown in the figure below,

Figure 3.11: SiGe high power frequency AGC response

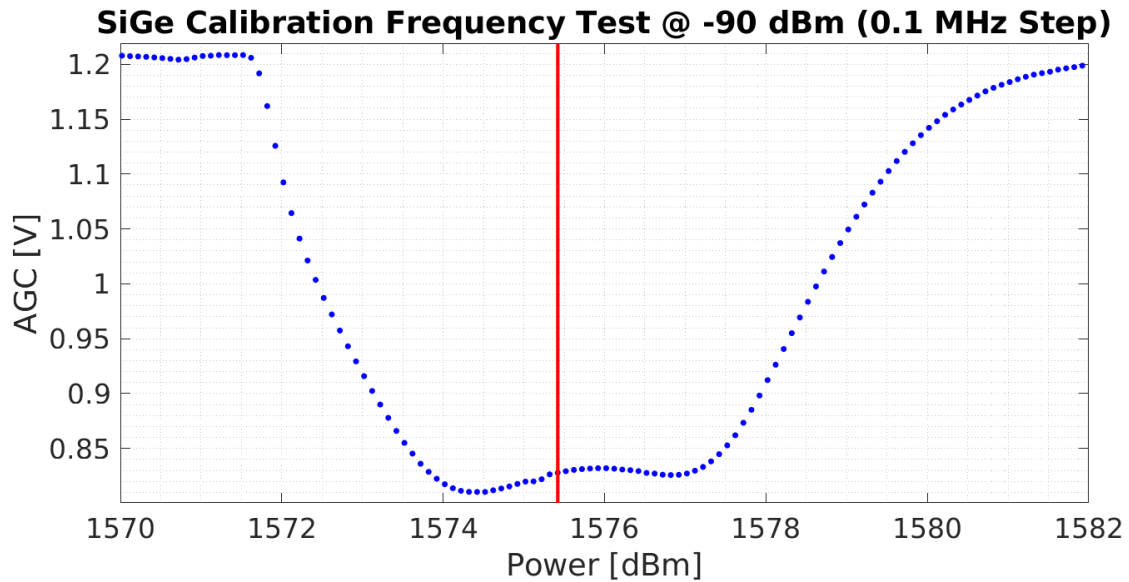


this data set had the power output set to -30 dBm and covered frequencies from 1500.42 MHz to 1650.42 MHz. The choice of power level was determined by the NEAT jammer specifications which had the capability of outputting 10 dBm, thus the power level here would represent a signal coming from the NEAT at a distance and thus is likely to occur in the field. So looking at the plot we can see that the change in frequency has a considerable effect on the AGC. With this amount of power we would expect the signal to always floor the AGC measurement as shown in the previous section, but here the response varies both up and down and is not symmetric about the center frequency. This result is surprising and very important since it could potentially result in the falsification of the inverse proportionality theorem. Though to confirm this another experiment should be performed where the power is varied over a specific off carrier frequency, but for now I will provide a hypothesis as to why this occurs.

I believe the reason for such a result has to do with the both the band-pass filter and time constant of the SiGe's AGC. To begin we know that the band-pass filter is designed such that the it attenuates any signals outside of the specified band. Therefore when the power of the incoming off frequency signal is low it causes the noise to once again be dominate and the AGC response is what we expect, but as we increase the power the filter does not attenuate enough and there is still a small presence of the off frequency signal in the incoming signal. This component I would imagine gets to a point where it now mostly fits in the center bin's voltage therefore a greater frequency of center bins is counted and this causes the AGC to increase. As we go further out the filter attenuates even more and eventually filters out everything except the noise. Now why the response is asymmetrical and goes both up and down probably has to do with the time constant of the AGC. This specific time constant could be set that it is only encountering certain portions of the CW signal such that it always appears to be adding high or low voltages, it could also just be due to the band-pass filter design.

The second set of data used a constant power level of -90 dBm and only varied over the frequencies of 1570.42 MHz to 1582.42 MHz. This scenario is more representative of what a receiver might experience in everyday operation since there are a number of devices which output frequencies that have harmonics near the L1 spectrum or may accidentally bleed into it with poor circuit design. The results are given below,

Figure 3.12: SiGe low power frequency AGC response

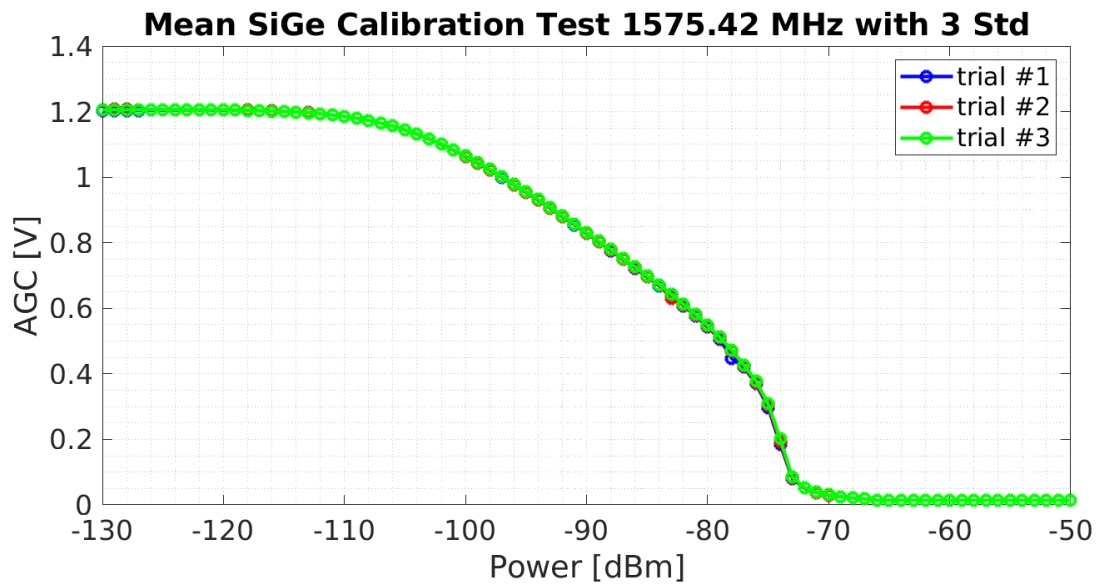


here the AGC response is more in line with what we would expect. The AGC metric creates a trough about the center frequency as a result of the filtering being performed in the front-end of the receiver. This relationship continues out to an asymptote though it should be noted that this behavior is also not symmetric.

3.3.5 Miscellaneous


This final experiment section looks at the effects of repeated measurements and different receivers of the same type on the AGC response. The first experiment had data repetitively and continuously taken from the same receiver in the same manner as the very first experiments. This was done to see if the AGC metric was dependent on the previous measurements or initial state. More concisely we were looking for signs of hysteresis. Once again this experiment was only performed with the SiGe as it was shown in the previous section that the overall performance of the individual receivers varied on a small scale. Below the results of the hysteresis test are shown using two plots,

Figure 3.13: SiGe hysteresis results




this first plot shows the AGC measurements given for the power calibration test after multiple runthroughs. As you can see the grouping of the measurements is very tight and the below graph demonstrates this,

Figure 3.14: SiGe hysteresis mean and third deviation



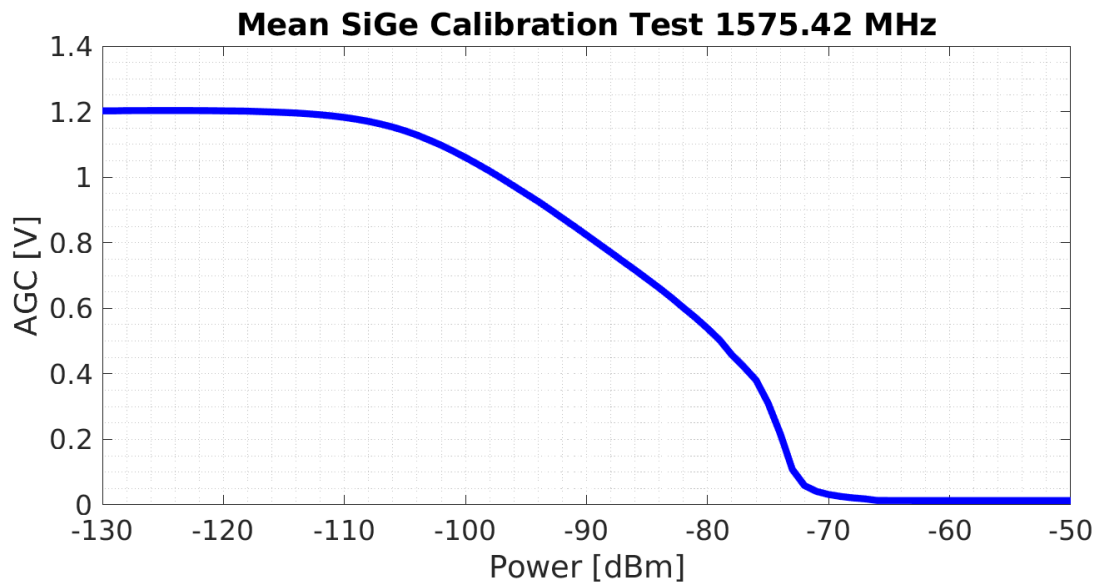
here the average AGC and three times the standard deviation is shown as the dots and corresponding error bars. The greatest third deviation was 0.0426 V while the average was 0.0037 V. Thus it is safe to say that AGC does not suffer from hysteresis and therefore it is possible to relate specific signal type power levels to a discrete AGC value. The last test performed in this experiment was simply a repeat of the original white noise tests but this time a different receiver was used. The purpose of this was to see if there is uniformity in the AGC metric with modules of the same type. The results are given in the figure below,

Figure 3.15: Multiple SiGe module AWGN response



here the results of the three new modules are shown. From the plot it can be seen that the AGC does vary when using a different module and this result isn't surprising since the physical components that make up the receiver all have inaccuracies in which they are built. Though despite this the overall trend remains and is repeatable thus it should be noted that the user or manufacturer should create an estimate of the AGC value based off a large number of tested modules. The last plot of this section Fig. 3.16 showcases what I call a "power calibration curve". This plot is an average of all the previous white noise data and acts as a potential reference tool for estimating the incoming power of a white noise signal. It can also be used to set a threshold for a jammer detector, this will be discussed in the next section, but a quick example would be to say that your SiGe module is returning an AGC value of 1 V then you could interpolate from this curve and determine that the input power is -95 dBm.

Figure 3.16: SiGe AGC power calibration curve



3.4 Discussion

Let's now try to summarize the results of the showcased experiments with respect to the original goal. Now the original goal was to verify the existence of a one-to-one inverse relationship between the incoming signal power and the AGC metric. The results of the experiments performed did back up this original claim though they only did so within the context of a specific signal wave form. Given that all three receivers outputted different AGC values for the same input power but only different wave form disproves the idea that the AGC metric provides a true power measurement. Therefore it seems highly unlikely to design a lookup table for a specific receiver for the purposes of determining the power of the input signal. Instead the AGC should be used as a measurement of relative power trends in the spectrum. Meaning that it can accurately inform a user of whether or not more or less power has been placed in the input signal.

Chapter 4

Applications and Future Work

The second major objective of this paper was to design applications which take advantage of the findings in chapter 3. To remind you the conclusion of the last section was that the AGC metric can not be directly related to a true power measurement of the incoming signal, it can instead inform the user that a change in the power of the spectrum has occurred within a range of values that is specific to the receiver being used. An example of this is that if you were to turn on a jammer with an unknown wave form next to a receiver then you cannot confidently state the exact power being received in dBm by looking at the current AGC metric. Instead you can accurately tell if the overall received power has increased relative to some historical AGC measurement. This is true for all except of the most extreme cases where a high powered signal with a frequency near the target signal frequency is introduced. Though given the isolation of the GNSS spectrum this occurrence is unlikely to occur. Therefore let's now examine some of the applications proposed in this paper.

4.1 Jammer Detection Applications


The first instance of application was inspired by the vast majority of references used as an outline for this paper and that has to do with the detection of GNSS jammers. The idea behind this comes from the knowledge of how a typical jammer operates which was explained in chapter 2. Though as a reminder a jammer simply masks the true GNSS signal with a more powerful signal in the same spectrum[3, 1]. Thus the idea behind GNSS jammer detection with AGC is that if the

AGC metric were to decrease in value significantly then we would know that a jammer is present since the AGC value should remain relatively constant given that it should only be measuring the noise floor which only changes gradually over time and in small magnitudes. This basic idea is what stands as the foundation for the applications given below.

4.1.1 Automotive Jammer Detection Northwest Parkway

The first application was invented to solve the reoccurring problem of GNSS jammer usage in automotive vehicles. As stated in the background section there have been a large number of instances of drivers installing GNSS jammers into their cars for the purpose of maintaining privacy. This specific example is ideal for utilizing AGC since the problem is massive in scope since any vehicle on the road could potentially be hiding one of these jammers. With such a large number of areas to cover it becomes nearly impossible to deploy current detection methods as they are often materially and/or computationally expensive. With AGC it becomes possible to build a jammer detector with just a COTS receiver. The only remaining question is where to place such detectors and how can they be used to potentially capture perpetrators? The answer proposed in this section is to add a GNSS AGC monitoring system into toll road gates. A diagram of the concept is shown in the figure below,

Figure 4.1: Toll road GNSS jammer detection concept



in the figure we can see that the toll road monitor system would consist of a antenna and receiver which would report the current AGC metrics of the road to the main server. This server would then determine if the incoming AGC values are below a predetermined nominal threshold and if so it would trigger the toll road cameras to take a picture of the vehicles near the monitor station. The only additional equipment here is the antenna, receiver, and software used to incorporate the AGC readings with the cameras. Now the idea here is not to capture an offender after one pass but to instead gather a history of repeat offenses such that the offender can be clearly identified in the high traffic environment. Now to test this concept I contacted a toll road here in Colorado called Northwest Parkway. Under their supervision I was permitted to place a simple prototype of the monitoring system without the cameras on the side of their toll gate. The equipment used is shown here,

Figure 4.2: AGC monitoring equipment



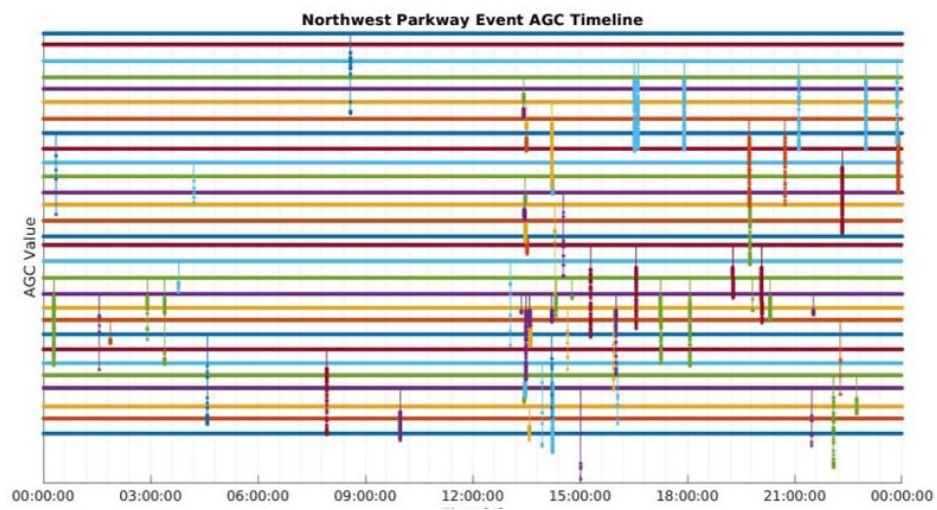
it consists of the SiGe receiver, L1 patch antenna, and laptop. In this case the laptop plays the roll of the server which runs a piece of software that is responsible for monitoring the AGC values. The next photo shows where and how the setup was deployed at the toll road gate.

Figure 4.3: Deployed monitor setup



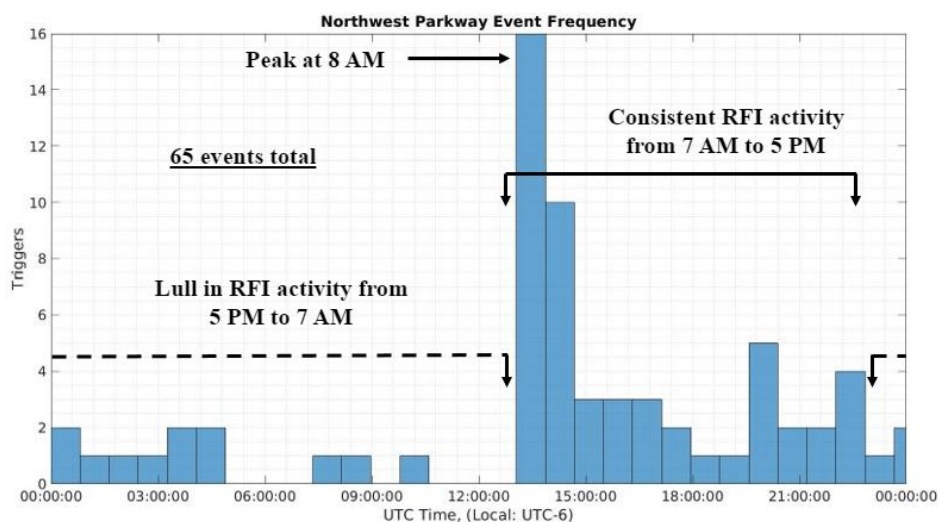
With this system in place I recorded months worth of data in the hopes of showcasing AGC value statistics that pointed towards the idea that some sort of interference was coming from the vehicles. The AGC results for an entire month are given in Fig. 4.4.

Figure 4.4: Northwest Parkway AGC timeline



From Fig. 4.4 we can clearly see that there are instances where the AGC value rapidly dropped and would've resulted in the monitoring system taking a picture if it were actually operational. Though let's take a second to understand how this specific detector operated. First we would measure a day's worth of AGC and come up with an estimate for what the nominal AGC value is (the AGC corresponding to the noise floor of this specific setup), we would then look back to the power curve presented in Fig. 3.16 and ask the question of what difference in AGC value w.r.t nominal represents a significant change in the input power such that we are likely seeing interference? For this particular data set we see that the nominal AGC value is around 0.6 V and thus 0.5 V would represent a 6 dB drop in power from the nominal. With this we would set a threshold value in the software which is compared to every incoming AGC value and if the value drops below this threshold then the program would save a certain amount of historical IF and AGC data for analysis and also trigger the cameras to take a picture. Though this particular figure is difficult to understand so let's look at the next one Fig. 4.5.

Figure 4.5: Northwest Parkway trigger daily frequency



Here each continuous drop in AGC from the last plot was mapped onto a histogram. This

makes it easier to see at what times the interference occurs most often. Looking at the figure it can be seen that during high traffic hours of 7 AM to 5 PM there is a significantly higher number of occurrences than the after work hours. This would lead us to believe that the reported interference is coming from the vehicles themselves and thus a system such as the one purposed in this paper could successfully identify an offender in an efficient and automated manner. To further drive this home let's look at the data in one more way.

Figure 4.6: Northwest Parkway trigger timeline

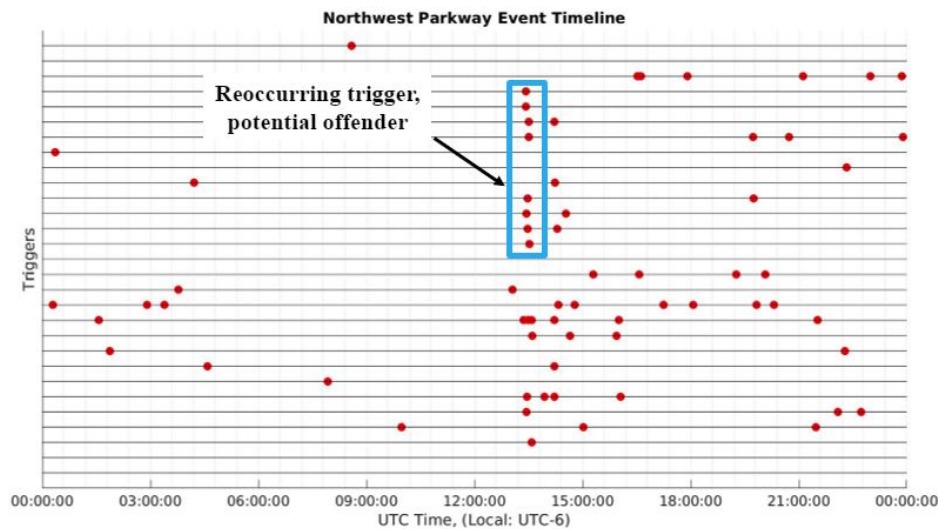
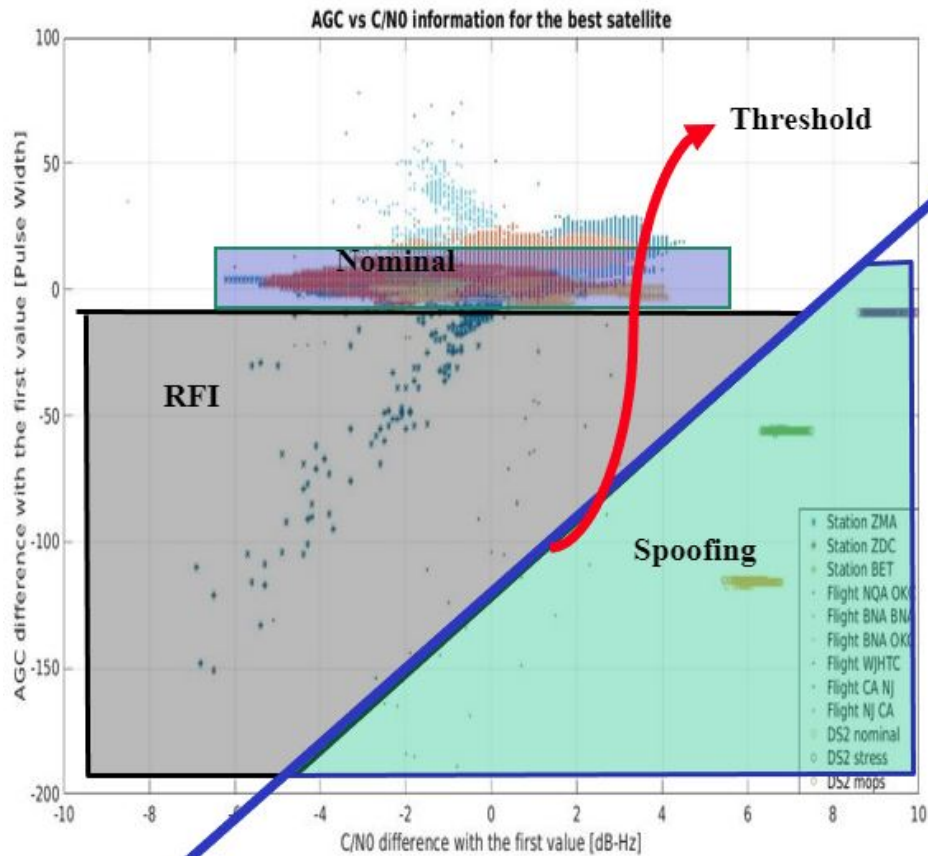


Figure 4.6 displays dots at points in time which correspond to a triggering of the system. From here we can once again see how there is a large concentration of these dots during the higher traffic periods. But more interestingly we can see that for two weeks there was a repeat triggering at roughly the same time during the day. It can be speculated that this is an indicator of a repeat offender who is using this toll road on a daily basis. Without ways to identify the specific vehicle with this setup it is impossible to say. Though hopefully this small example helps demonstrate just how effective and easy the AGC detection method can be incorporated into preexisting infrastructure.

4.2 Jammer and Spoofing Detection Applications

The next set of applications combine the previously mentioned jammer AGC detection method with C/N_0 to detect and distinguish between jammer and spoofer attacks[1]. This is important since spoofers are often harder to detect since they produce good position solutions and replicate the true signals physical characteristics. The consequences of spoofing are often far greater than jamming since it can cause GNSS reliant systems to operate on false data rather than simply preventing them from working. The method used in this paper for accomplishing such goals is the combined observation of both the current AGC and C/N_0 value. The idea here is that all types of radio frequency interference (RFI) must first overpower the original signal in order to mask the true data with either noise or false data. This change in power is easily detected by the AGC but the AGC values themselves are unable to distinguish whether the attack is a jam or a spoof. To know this we look at the C/N_0 value outputted by the receiver, C/N_0 to put simply is a type of SNR and tells us about the current condition of the GNSS signal. A high C/N_0 value corresponds to a high SNR and thus our receiver believes that it is receiving strong GNSS signals with respect to the noise. Now we know that a jammer adds noise to the spectrum and thus causes the C/N_0 to drop while a spoofer increases the power of a simulated GNSS signal thus the C/N_0 would actually increase in this case. Alone this means that C/N_0 cannot be used to detect RFI but when combined with AGC we can create "zones" of C/N_0 and AGC values that correspond to a jammer or spoofer. For example if we use a jammer we would expect both the C/N_0 and AGC values to decrease w.r.t some nominal values. On the other hand if we were to activate a spoofer near the receiver we would expect the AGC value to decrease due to the additional power and the C/N_0 to increase as the new signal introduces only higher power GNSS signals.

Figure 4.7: C/N_0 vs AGC for TEXTBAT and WAAS station data from Miralles (2019)



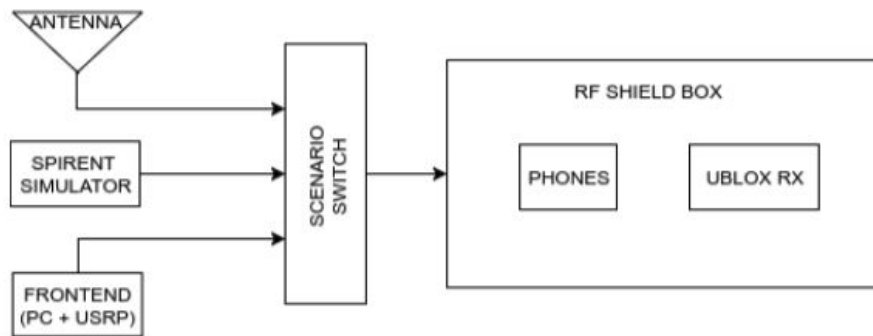
The figure above was produced by running simulated data from the TEXTBAT and in-lab WAAS datasets through a Novatel receiver. Here we can clearly create zones that correlate a range of C/N_0 and AGC values to either a nominal, jammed, or spoofed signal.

4.2.1 Mobile Applications

One way to apply such methods is in the mobile cellular phone industry. This is a vital piece of modern society that is still highly vulnerable to both types of RFI. My colleague and I attempted to tackle this problem for Google's Android operating system by seeing if it was possible to jam/spoof the position solution provided by the OS. This is slightly more difficult than it sounds since modern phone OSs use assisted GPS (AGPS) which is the idea of using navigation data

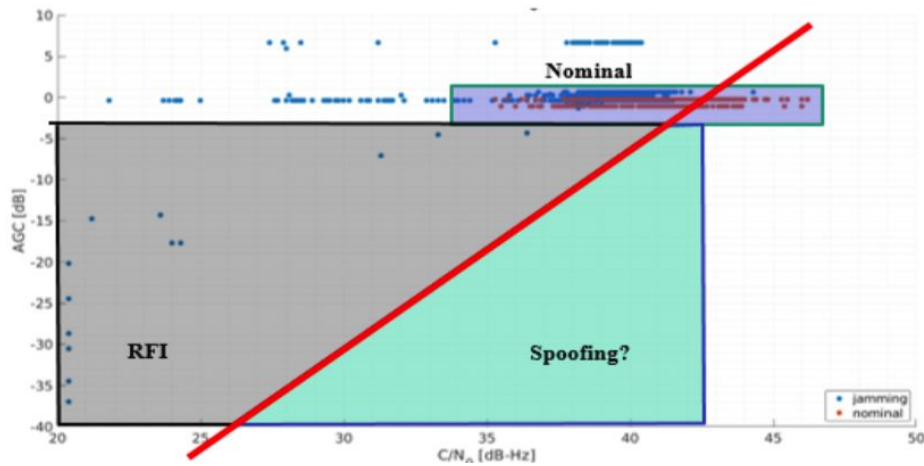
from other sources in order to improve or speed up the classic solution. Also it wasn't till 2017 that Google provided the tools to access the AGC and thus a custom app had to be made for the purposes of logging this data along with the C/N_0 . In order to test this we developed the following setup,

Figure 4.8: Phone jamming/spoofing experiment setup from Miralles (2019)



here we basically have an RF shielded box that contains both an Android phone and a UBLOX receiver. This box had a pair of antennas inside of it that were connected to three outside sources: an outdoor antenna for meaconing (this allows us to get a true signal into the box), a spirent simulator (a GNSS simulator that is acting as a spoofer), and a front-end setup. The idea was to pump in live sky signals from the roof of the building into the box and then attempt to spoof/jam these signals with the simulator with the phone being the true target and the UBLOX receiver acting as a control. During the experiment we were able to spoof the phone position and time solution. Though when it came to detecting the spoofing we were not as lucky.

Figure 4.9: Phone: C/N_0 vs AGC for TEXTBAT and WAAS station data from Miralles (2019)



The results of the experiment are shown in Fig. 4.9 and here we can see that the AGC did drop at points where we attempted to inject RFI, but as you can see there are no data points in the predicted spoofing section of the plot. Unfortunately, it seems that my colleague and I either didn't adequately attenuate the sky signal or perhaps there was an error in the positioning of the phone that lead to one signal being received better than the other. Regardless this test showed that it is possible to spoof an Android phone and it is possible to detect such RFI using the AGC which is a great idea since the processing power on phones is a limited resource given that the platform is mobile and thus battery reliant.

4.3 Future Work

There are two main applications should be worked on in the future. The first is to incorporate RFI detection into the flight stack of open source autopilots. One of the most vulnerable systems to GNSS RFI are unmanned aerial vehicle systems. This is due to the fact that they rely heavily on GNSS position solutions for navigation. It is often the job of GNSS to act as a truth value in which other forms of inertial guidance are aligned to, and the fast moving nature of such systems means that even a temporary disturbance in service could result in catastrophic failure. This

problem isn't new and there have been many developments in the field of jamming detection for UAVs but unfortunately nearly all of them are tied to the solution set that the receiver outputs with each position solution. This message often comes in single digit frequencies and thus there is little time for the autopilot system to react to any jammer warning that is presented along with the position solution such as C/N_0 . AGC values are independent of the navigation engine and therefore can be sent at higher rates they are also smaller in size and thus don't require large amounts of computation power. The proposed change to the flight stack of autopilot software such as PX4 would be to include a plugin which requests AGC values from the onboard GNSS receiver and then compares this value to a nominal threshold in order to warn the autopilot that the GNSS position solutions have been compromised.

The second future application I would like to perform is the use of AGC values to not only detect but localize sources of interference. The idea works like this, you would have either a hand held or vehicle mounted GNSS receiver, IMU and/or SLAM system, and computer that would continuously record the geographical position and corresponding AGC value. That way it would be possible to make a map of an area that showcases the AGC value at every point. Using this method you would be able to determine the likely position of the interference source by looking at what points on the map exhibit the lowest AGC value. This method would really only work for stationary sources but it would small enough to fit on a UAV which could quickly traverse an area.

Finally, one major piece to complete is the analysis of how a change in AGC impacts the receivers performance. That is to say if we see a decrease in AGC would that entail that the performance of the receiver would also decrease proportionally? This is very important because even though a receiver may be receiving some interference doesn't necessarily mean it is being jammed in the sense that it is suffering degradation in its performance. Also the detector mentioned in the automotive section could also incorporate the IF data such that it could identify the jammer's wave form and thus use a power curve to determine the incoming power.

Chapter 5

Conclusions

Through the course of this research we have achieved our main objectives that were laid out at the start. The first was to verify or refute the claim that the AGC metric coming from GNSS receivers could be used to determine the power level of the incoming signal. The results of our experiments showed that that AGC could not be used to calculate the power level of an input signal unless the specific signal wave form had been evaluated previously otherwise the ambiguity of AGC value for general signal types was far too great. Despite this we were able to verify that there did exist an inverse relationship between the AGC metric and a general input signal. Therefore it is possible to determine if a historical input signal's power was greater or less than it had been at some nominal operating value. This finding allowed us to accomplish the second objective of applying the AGC metric as a means of detecting RFI in the GNSS band. Two specific cases were outlined and both looked at GPS L1: automotive GPS L1 spectrum enforcement using AGC as a means of detection, and jamming/spoofing detection using both AGC and C/N_0 values from a cellular phone. Both of these examples purposed RFI detection solutions that were easy to implement into preexisting software yet also effective. Therefore it is my conclusion that the AGC metric can be a powerful addition to the current arsenal of RFI detection methods. Lastly, I will recommend to the GNSS community that there should be an effort made to standardize the AGC metric. As this would allow for easy use across all receivers.

Bibliography

- [1] Dennis Akos. Whos afraid of the spoofer? gps/gnss spoofing detection via automatic gain control (agc). Navigation, 59, 12 2012.
- [2] Erik Axell, Fredrik M. Eklf, Peter Johansson, Mikael Alexandersson, and Dennis Akos. Jamming detection in gnss receivers: Performance evaluation of field trials, 03 2015.
- [3] Frédéric Bastide. Automatic gain control (agc) as an interference assessment tool. 2003.
- [4] J. A. Bhatti, T. E. Humphreys, and B. M. Ledvina. Development and demonstration of a tdoa-based gnss interference signal localization system. In Proceedings of the 2012 IEEE/ION Position, Location and Navigation Symposium, pages 455–469, April 2012.
- [5] Mohammad Zahidul H. Bhuiyan, Heidi Kuusniemi, Stefan Sderholm, and Esa Airos. The impact of interference on gnss receiver observables a running digital sum based simple jammer detector. Radioengineering, 23:898–906, 09 2014.
- [6] Franc Dimc, Matej Baec, Daniele Borio, Ciro Gioia, Gianmarco Baldini, and Marco Basso. An experimental evaluation of low-cost gnss jamming sensors. Navigation, 64:93–109, 03 2017.
- [7] Bernhard Hofmann-Wellenhof, Herbert Lichtenegger, and Elmar Wasle. GNSS - global navigation satellite systems: GPS, GLONASS, Galileo, and more. Springer, 2008.
- [8] Jeong Hwan Yang, Chang Ho Kang, Sun Kim, and Chan Gook Park. Intentional gnss interference detection and characterization algorithm using agc and adaptive iir notch filter. International Journal of Aeronautical and Space Sciences, 13:491–498, 12 2012.
- [9] Jonas Lindstrm, Dennis Akos, Oscar Isoz, and Marcus Junered. Gnss interference detection and localization using a network of low cost front-end modules. 04 2019.
- [10] J. Poncelet and D. M. Akos. A low-cost monitoring station for detection amp; localization of interference in gps l1 band. In 2012 6th ESA Workshop on Satellite Navigation Technologies (Navitec 2012) European Workshop on GNSS Signals and Signal Processing, pages 1–6, Dec 2012.
- [11] Jon S Warner and Roger G Johnston. Gps spoofing countermeasures. 25, 01 2003.