

Reevaluating the Message Loss Rate of the Wide Area Augmentation System (WAAS) in Flight

Matthew J. Hirschberger, Sherman Lo, Todd Walter
Stanford University

ABSTRACT

The Space-Based Augmentations System (SBAS) Minimum Operational Performance Standards (MOPS) currently specifies that the message loss rate must not exceed 10^{-3} per message in order to meet aviation requirements for the availability of SBAS solutions. There has been little public reporting examining message loss performance in flight. The most recent published results were conducted in 1998 with the prototype version of WAAS known as the National Satellite Testbed (NSTB). Given that the capabilities and performance of WAAS have advanced beyond this prototype in its 18 years of service, a fresh analysis of the message integrity in this new system is necessary. The focus of this work is to examine the loss rate for WAAS reception using in flight data. We cross-referenced message data obtained from on-board receivers with the known broadcast history of each GEO to determine the loss rate. We mapped missed messages to timestamps along the flight path to show where messages are lost and to determine where messages are lost due to coordinated turns or low carrier-to-noise ratio (C/No). From this, we estimate the probability of message loss in nominal level flight as well as in other conditions such as turns and degraded C/No.

INTRODUCTION

One of the important assumptions in the operations of the Wide Area Augmentation System (WAAS) is a low message loss rate (0.1%). Having a low message loss rate is important in order to receive the necessary integrity information and warnings in time as well as meeting the high availability and continuity requirements. While the system design has been analyzed to meet the targeted loss rate [1], there is little publicly available data published to support this analysis. The last published paper that studied the topic in detail was released in 2000, several years before WAAS became operational. The data evaluated in that work also encompassed just the west coast of the United States. So, we embarked on an effort make a more comprehensive evaluation by gathering and assessing data from geographically diverse parts of the WAAS coverage area. In addition to nominal loss rate, we also analyze the data to understand conditions having higher loss rates such as banking.

BACKGROUND

The first published examination of the WAAS reception in flight was published by Fuller in [1] [2] based on flight tests conducted in 1998. While WAAS was not operational in 1998, results could be gathered using the FAA prototype WAAS known as the National Satellite Testbed (NSTB). The NSTB broadcast its transmissions on two geostationary (GEO) satellites: Pacific Ocean Region (POR) and Atlantic Ocean Region West (AOR-W). The analysis conducted by Fuller used a series flight tests conducted in 1998 around Palo Alto, California and in June 1998 around Juneau, Alaska. This data collection was based on the NSTB signal from POR, as AOR-W was not receivable in that part of Alaska and poorly received in Northern California. The Alaska data collected consisted of 12 days of flight testing which demonstrated the effectiveness of Stanford's algorithms and provided insight into the integrity of WAAS message reception. This message reception was examined in more detail [1] and [2]. Fuller determined from the data that the message loss rate met the

established performance requirement for WAAS of 10^{-3} . As the flights included steep bank angles which caused many of the message losses. Fuller also went on to propose a Markov transition model to model the burst-mode losses commonly experienced.

Fuller's conclusions are based on data from the NSTB, but similar analyses have not been publicly presented using flight data from the operational WAAS system. There are substantial differences between the two systems that warrant this new analysis. WAAS has been expanded to the point that it now provides nearly full coverage of North America (CONUS), while the NSTB provided limited coverage of the United States. In addition, Fuller's results used measurements from areas where GEO reception was compromised. AOR-W is at low elevation in California and was not receivable in Alaska, while POR was at low elevation in Alaska. In this work we will show that being at low elevation makes losses due to banking, even slight banks, more likely.

ANALYSIS METHODOLOGY

To conduct our analysis, we collected data from several flight tests. We were able to field a Trimble BX935 receiver on an FAA test aircraft, a Global 5000. This FAA aircraft was flown around the FAA technical center in New Jersey as well as several flights to various FAA facilities in the United States (U.S.) including Hawaii. Additionally, we fielded a similar unit on a U.S. Air Force (USAF) C-12J for some flight tests over the Mojave Desert. Some flights occurred during interference testing. Between these two aircraft and their flights, we have a reasonable set of data collected from a broad swathe of the U.S. One important caveat about this test is that the Trimble receiver is not a flight certified receiver.

The Trimble receiver on the test aircraft records all useful WAAS messages received that passed the cyclic redundancy check (CRC). Ideally that would mean losses from a GEO are time epochs (seconds) where we do not have a record of a received message for that GEO from the receiver. However, since non-useful messages are not recorded, reception of WAAS Message Type (MT) 63, a null message that serves a filler if there are no necessary messages to transmit, is ignored and not recorded by the receiver. So, from the Trimble receiver alone, we cannot tell if a message was lost or if it correctly received MT 63 which had been discarded. Hence to get the full message broadcast history of a GEO, we cross-referenced the Trimble data with a static reference source of transmitted messages for each satellite. This latter set of data is pulled from the database of the Centre National d'Etudes Spatiales (CNES) in France. This provides knowledge of when MT 63 is transmitted so we can properly account from them in our message counts. We pull the "CNES file" for each GEO PRN from the directory that matches the same day as the flight we want to analyze. A MATLAB script reads in both sets of data and parses each line to record the time epoch and message type. For this work, we only consider the L1 frequency, so any L5 messages collected by the Trimble receiver are stored but not analyzed in this work.

There is one strange quirk with harmonizing data from the two sources. A side-by-side view of the messages from the Trimble receiver and those from CNES shows that the message types on each timestamp differ by one second. Our best understanding is that the Trimble receiver marks the message type differently. Regardless, this issue is easily fixed by either shifting the message types or applying a correlation function to two vectors of the message types.

To visualize where the missed messages occurred along the flight path, we aligned the calculated latitude, longitude, and height (LLH) with the message received (or not received) for each epoch. Additionally, velocity was also incorporated to estimate aircraft attitude. This was stored in a MATLAB structure and later plots will show the flight path along with locations with missed messages. From attitude estimates and knowledge of aircraft and GEO location, we can calculate the apparent elevation of the GEO to the aircraft. For the sake of clarity, we define apparent elevation as the elevation of the satellite relative to the horizontal plane of the aircraft (which can change during maneuvers) and actual elevation being the elevation of the satellite relative to the horizon.

DETERMINING THE LOSS RATE

The most basic calculation we must make is the message loss rate. We define the average message loss rate π_L as:

$$\pi_L = \frac{N_{lost}}{N_{total\ expected}}$$

Where N_{lost} is the counted number of lost messages in the time frame and $N_{total\ expected}$ is the total number of messages that should have been received in the time frame. Because we cannot distinguish between the scenarios of losing or receiving a MT 63, epochs with MT 63 are not included in either of these totals.

To generate these counts stated above, we keep a vector of message types for each GEO PRN from each data source. Table 1 shows a vector of data received from one WAAS satellite during one of our flights. When we read each line of our Trimble data, we get a timestamp for each message we receive, as well as the message type of that message. We notice that we are missing times 314927 and 314928 in this set. The absence of timestamps indicates these messages were either lost or received but discarded due to it being a MT 63. We use knowledge of transmission to remove cases of MT 63.

Table 1 – Timeseries from Trimble data

Time	314926	314929	314930	314931	314932	314933	...
MT	28	4	3	2	28	26	...

Time	314926	314927	314928	314929	314930	314931	314932	314933	...
MT	2	28	63	1	4	3	2	28	...

In Table 2 we read in our CNES data to gain information about our lost messages. The one- second shift between message types may throw off the reader here. The best way to understand what happens here is to not focus too much on the timestamps, but rather the order of the message types. We can correctly map the message type to timestamps after. At time 314926, a message was received that was marked MT 28. The following epoch, the receiver should have recorded a message marked with MT 63, and after that one with MT 1. It is unknown whether the first missing message was lost or simply not recorded because it is an MT 63.

Table 3 shows the correct message type-to time mappings and marks all missed messages with a message type of -1. The loss rate can then easily be counted using the message types. The number of lost messages is 1. The number of expected messages in this section of the data is 8 minus 1, since there is one MT 63. Therefore, the loss rate for this example section is 1/7.

Table 3 – Timeseries with missed messages marked and correct MT mappings

Time	314926	314927	314928	314929	314930	314931	314932	314933	...
MT	2	28	63	-1	4	3	2	28	...


ESTIMATED VALUES

Next, we calculate the aircraft state: position, velocity, and attitude. This information will help us confirm or disprove our intuition on which flight maneuvers may significantly impact the loss rate. The position in LLH is calculated by our ARAIM MAAST software using pseudo ranges from the receiver at each given time epoch. From here we can immediately estimate the velocity and acceleration of the plane at all times by taking the respective derivatives of this position data.

The next thing we are interested in is the attitude of the plane, and in particular the bank angle of the plane along the flight path. Unfortunately, we did not have precise attitude for all flights and so we developed a generic method to estimate attitude based on some basic assumptions. This is similar to what we did in [6] though we have refined the estimation method. To estimate attitude, we used our velocity and derived accelerations and an assumption that any turns along the flight path are coordinated turns. Coordinated turns are turns that are conducted where there is no sideslip – meaning the aircraft nose is pointed to direction it is traveling. For the mathematically minded, it means that combination of the gravity (weight) and centrifugal force vector is in the down direction relative to the aircraft platform. This is balanced by the lift force which results in a net zero acceleration in the aircraft down direction. Hence there is no acceleration component to the side of the aircraft. As a result, there is only acceleration primarily perpendicular to the path of travel. A related assumption is that the velocity vector matches the direction of travel.

To calculate attitude, we start by using the GNSS velocity vector to calculate the direction vector, e_{dir} , and the acceleration vector. Based on the assumptions of coordinated turn, GNSS only should measure acceleration in the cross and along track direction – that is no acceleration in the up/down direction relative to the body of the aircraft. Then we can determine the along track acceleration by finding the component of acceleration in the direction of travel from the direction vector. The cross-track acceleration is the remaining acceleration as GNSS should not measure acceleration in other components. We can calculate the apparent acceleration (resultant vector shown below) by adding the cross-track acceleration (i.e. centrifugal in the figure) and gravitational acceleration (not measured by GNSS). The other components of attitude are more straightforward to estimate as they do not require acceleration. Assume the aircraft is pointed in the direction of travel, the horizontal direction vector on the horizontal plane provides the estimated heading. The arc tangent of the vertical velocity divided by the horizontal velocity.

We can calculate the bank angle as the arc cosine of the gravitational acceleration divided by the resultant. The coordinated bank angle assumption is fairly reasonable – this can be seen in our flight test data where the Global 5000 banks are nearly perfectly coordinated turns which is not surprising as there is instrumentation onboard many aircraft, including the Global 5000, to do just that.



Equations for calculating bank angle of the aircraft with the above diagram.

$$e_{dir} = \frac{\vec{v}}{\|\vec{v}\|} \quad (1)$$

$$a_{alongtrack} = e_{dir} \cdot a_{GNSS} \quad (2)$$

$$a_{crosstrack} = a_{GNSS} - e_{dir} * a_{alongtrack} \quad (3)$$

$$a_{resultant} = a_{crosstrack} + [0 \ 0 \ g] \quad (4)$$

$$\theta_{bank} = \arccos\left(\frac{g}{\|a_{resultant}\|}\right) \quad (5)$$

After calculating the bank angle and heading, we can determine whether the aircraft is banking towards or away from a given WAAS satellite, as well as the “apparent elevation” of that GEO to the aircraft. Here we define the apparent elevation as the elevation of the satellite with respect to the horizontal plane of the aircraft.

FLIGHT PATH AND MANEUVERING ERRORS

Plotting the message losses along the flight path and relative to metrics such as carrier to noise ratio (C/N_0) helps visualize some potential sources of WAAS message loss due to conditions other than nominal level flight. In this section, we examine each of these issues and outline a process for excluding data points from our set. Data points that are excluded can be deleted from our time vectors, but a simpler approach would be to mark points for use with a logical vector.

GROUND DATA

In all of our flights there is data collected while the plane is on the runway. At the starting of the flight, the Trimble receiver activates and begins logging shortly after the plane is powered on. Plots, such as the one below, shows that message losses occur often while idling on the runway. The reasons behind this are not investigated, but we assume the losses are due to obstructions by the terminal building or large multipath. We would not want to include this data regardless of how many messages are lost in these circumstances.

While using an attitude based cut off may be sufficient to eliminate these cases – i.e. WAAS correction data is only needed for use when at or over 200 ft above the ground as the applications it supports terminate at that flight level. The loss rate below this altitude is not applicable to the MOPS requirement. To conservatively exclude this ground data from our whole set, we choose to only use data points if they were recorded while the plane was at or above its cruising speed. We know the speed of the plane at a given time by mapping the speeds estimated above to the timestamps in our message history.

AIRCRAFT BANK

Unusually high bank angles along some of our flight paths were often discovered to lead to WAAS message loss rate, especially when the aircraft banks away from the GEO. Figure 1 shows the missed messages from PRN 135 plotted along one of our “racetrack” flights.

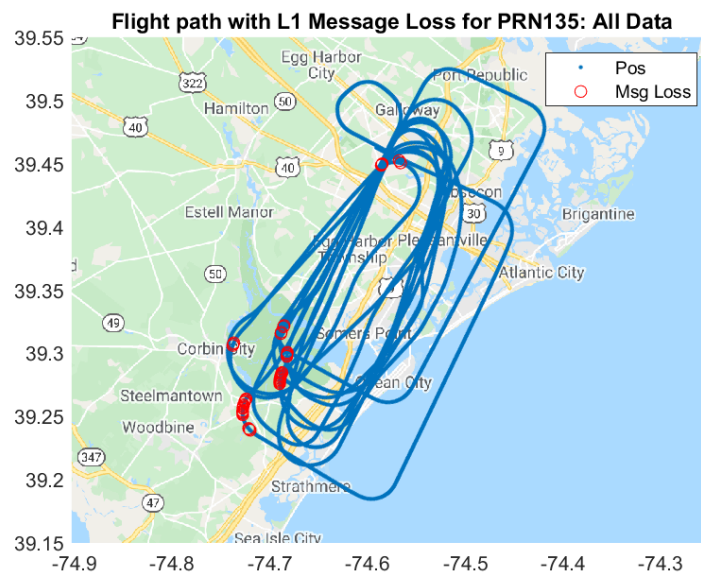


Figure 1 – Flight path showing messages lost from PRN 135




Figure 2 – Flight path where black marks points where the plane banks away from PRN 135

We notice patches where many messages are lost, particularly along turns facing to the southwest. If we add in the fact that the flight is occurring at about 74°W and that PRN 135 orbits above 107.3°W , the pattern of these losses suggests that the line of sight between the receiver and satellite is obstructed during these turns. To test this hypothesis, we estimate the roll angle and heading of the plane along the flight path as described in the last section. Turning points along the path are those where the roll angle is above 5° , and the heading rate of the aircraft is above a minimum threshold value. Combining this information with the known WAAS satellite positions allow for the identification of points where the plane is banking away from the satellite. The result of this process is shown in Figure 2. Figure 3 shows that these coordinated turns do not guarantee message loss, but certainly increase the chance.

Our basis for excluding data collected during these turns lies in a closer reading of the MOPS requirement. The actual statement in the WAAPS MOPS says that the 10^{-3} loss rate requirement applies under “interference conditions described in Appendix C and under minimum signal conditions defined in Section 2.1.1.10” [5]. An important point to consider is that the WAAS messages are most needed during the final approach segment where there should be no significant bank angles. The intention behind the actual MOPS requirement is for the message loss rate to apply to when aircraft is in nominal operations (approach or cruise) where the flight is typically level. From here we proceed with the interpretation that the loss rate specified in the MOPS applies to “straight and level” flight. We set the maximum roll angle to meet this condition at 5 degrees, and any points where the plane rolls beyond this limit is excluded from the data.

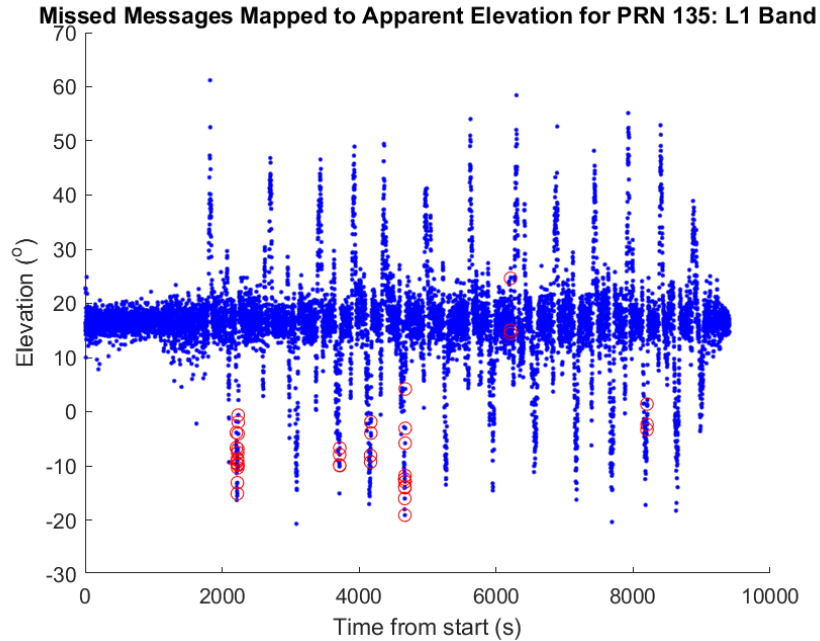


Figure 3 – Apparent elevation along the flight path shown in Figures 3 and 4. Red circles indicate points where a message was lost.

INTERFERENCE CONDITIONS

During some of our flights we noticed frequent message losses at low C/N_0 levels. C/N_0 typically remained between 30 to mid-40 dB-Hz, and losses occurred during sudden dips below 30. Eliminating abnormal C/N_0 helps eliminate effects of possible interference. For example, on a November 2017 flight near Hawaii, several instances of low C/N_0 and message losses occurred on all GEOs. While we could not find a GNSS interference events at the time of the flight, FAA notice to airmen (NOTAMs) had indicated that there was potential GNSS interference testing in the affected area a few days prior to the flight as well as testing in a nearby area a few days after.

Because the C/N_0 values in our data are lower than expected, we do not exclude data points with just a simple minimum value. Another reason we choose not to do this is because of concerns that the C/N_0 varies according to the apparent elevation of the satellite and the satellite itself (they transmit at different power levels). Figure 4 shows the median C/N_0 with satellite elevation. To determine the median value as shown, we organize received messages from each satellite into “bins” of apparent elevation values. All messages, both lost and received (as long as we can track the signal), are included in these bins. Once we have determined the median value for the bin, we can go back and look at the C/N_0 value of each received message and exclude any received below the threshold we set. We set our threshold to be five dB-Hz less than the median (i.e. median – 5).




Figure 4 – Median C/No value according to the apparent elevation of PRN 135.

RESULTS

The WAAS message loss rate can be examined under flight conditions as a function of the parameters calculated in the prior section, particularly apparent elevation, and C/No. The analysis allows us to examine performance under conditions targeted by the MOPS as well as more degraded conditions. Additionally, we will calculate the message loss rate as a function of satellite as WAAS employs three different satellites since they have different characteristics such as different broadcast power levels. Segregating out PRN 133, which operated until November 2017, is of particular interest as it has a narrower signal bandwidth than the other two geostationary satellites examined.

We plot loss rates for each satellite as a function of the apparent elevation. For the calculation, we first create bins of apparent elevation value ranges similar to what was done in the last section while finding the mean C/No value as a function of apparent elevation. Then all the message results for a given satellite are gathered and organized in their respective bins. Finally, we calculate loss rates for all the bins using the number of lost and expected messages in each bin. Bins with less than 100 total messages were ignored. These loss rates are shown in Figure 5 with curves for three different GEOs in different colors. The plot shows the rate if greater than zero. For a given GEO, bins with no message losses are not plotted but indicated with a zero in the color of the GEO plot line. A similar procedure was conducted for loss rates at various C/No bins and the resultant plot is shown in Figure 6.

Plots of the loss rate according to each variable show that the loss rate requirement is met under nominal signal conditions and for high apparent satellite elevations. Exceeding the requirement is not a concern it seems until the apparent elevation of the satellite dips below ~20 degrees, or the C/No value drops below 40. If we are confident that a typical commercial flight only sees shallow bank angles and C/No values above 40, then message loss is not a significant issue. There is concern about the loss rates found around 20 degrees of apparent elevation. Although they do not exceed the loss rate limit, the values here are a bit higher than expected. Of course, as long as another WAAS satellite is available at a higher apparent elevation, losses from the lower satellite pose less of a problem. We also have concerns about the trend of loss rate by C/No value. A theoretical plot of the loss rate by C/No value presented in [4] suggests that we should maintain message integrity in the low 30's dB-Hz, while our plot shows integrity wavering much earlier.

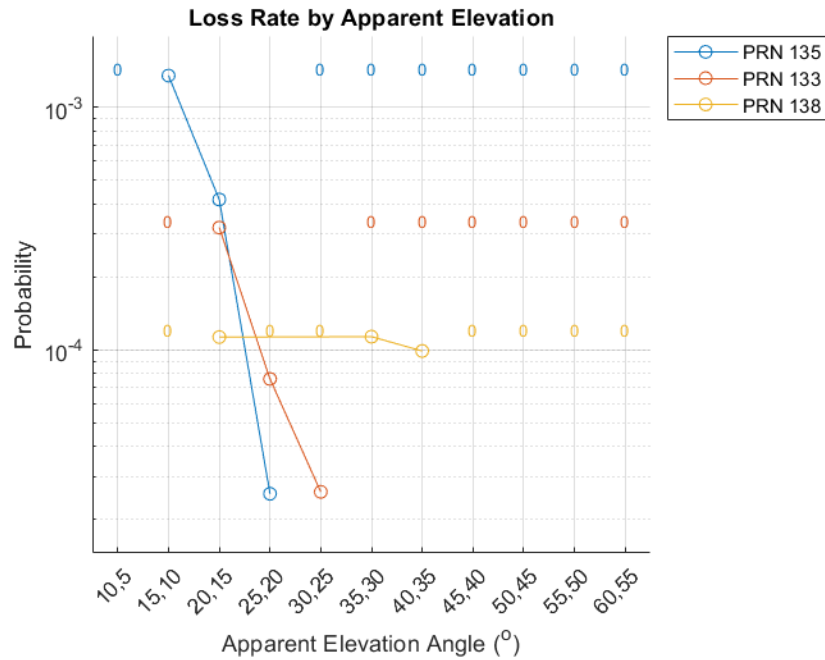


Figure 5 – Loss rate according to the apparent elevation of the WAAS satellite.

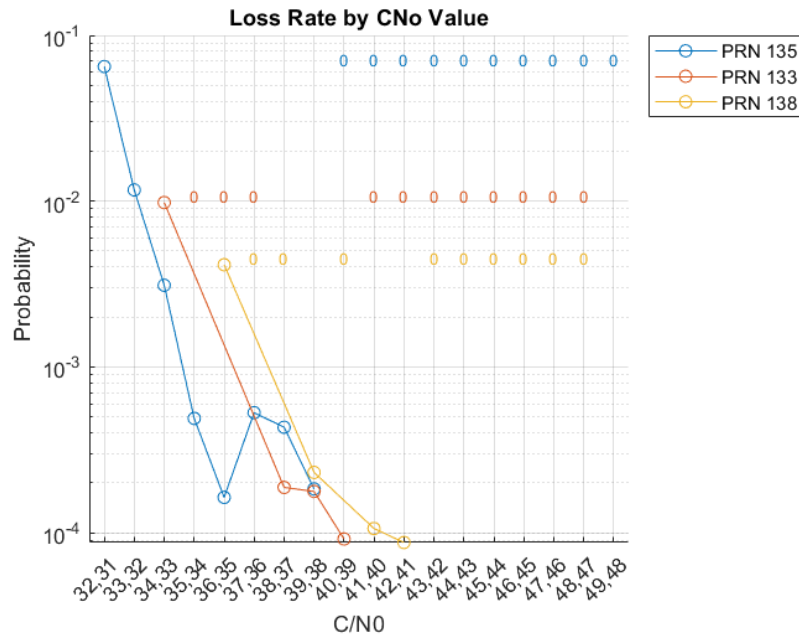


Figure 6 – Loss rate according to C/No.

We are satisfied in confirming that the loss rate meets the SBAS requirement under most conditions, but we also want to consider limitations in this work that may have led to unexpected values or trends in our plots. A detail worth restating is that the Trimble receiver we used is not developed for flight use and is not aviation certified. On one hand, it is encouraging that we still got the results we did with such a receiver. However, it is obviously important that these results also hold for receivers that will actually be used in WAAS approaches. Another factor that limits our result is simply the amount of data we have. Figure 7 shows all the messages received from PRN 135 under our set conditions. One feature that stands out immediately is that the number of messages received in the 5-10-degree bin hardly supports our confidence for the loss rate calculated for

that bin. Of course, we can change the minimum number of messages required for plotting, but this still illustrates the fact that we have a limited number of samples for each bin. If we had more samples for the 5-10 bin, we would certainly expect the number of losses to increase.




Figure 7 – Bar graph of number of messages received from PRN 135 under certain apparent elevation values from all the flights in our data set.

CONCLUSIONS

Examination of our flight data shows that the current iteration of WAAS satisfies operational requirements for nominal conditions. We have identified several conditions and flight patterns where the message loss is impacted. These instances are not of great concern as we believe that these situations fall outside of the conditions stated in the SBAS MOPS. However, it is important to understand and quantify these limitations. Apparent elevation in particular has a noticeable impact on the message loss rate. Avoiding high turning angles minimizes this impact, but if the plane is flying far away from the satellite's geostationary position, this low elevation may still pose an issue. While our results support the claim that WAAS meets the operational requirements, it is clear that more data is needed in certain areas. In particular collecting data with a WAAS/SBAS certified receiver would strengthen our confidence in the message integrity observed in this research.

ACKNOWLEDGEMENTS

The authors thank the Federal Aviation Administration (FAA) and the Stanford Research Experience for Undergraduates (REU) program for sponsoring this research. The authors thank Trimble Inc, particularly Stuart Riley and David De Lorenzo for providing receivers and valuable insights and tools for this research. The authors thank the Centre National D'Etudes Spatiales in France for making publicly available the data necessary for this work to succeed. The authors also thank Edwards Air Force Base and the US Air Force (USAF) Test Pilot School providing us with an opportunity to test.

The views expressed herein are those of the authors only and are not to be construed as official or those of any other person or organization.

REFERENCES

- [1] Fuller, Richard A., Walter, Todd, and Enge, Per K, "Burst Mode Message Loss Effects On WAAS Availability," Presented September 2000 at the Institute of Navigation's GPS conference, Salt Lake City, Utah.
- [2] Fuller, Richard A., Walter, Todd, Houck, Sharon W., and Enge, Per K, "Flight Trials of a Geostationary Satellite Based Augmentation System at High Latitudes and for Dual Satellite Coverage," Presented January 1999 at ION National Technical Meeting, San Diego, California.
- [3] Presented January 1999 at ION National Technical Meeting, San Diego, California.
Fuller, Richard, "Aviation Utilization of Geostationary Satellites for the Augmentation of GPS: Ranging and Data Link," Ph.D. Dissertation, Stanford University, May 2000.
- [4] Enge, P., "WAAS Messaging System: Data Rate, Capacity, and Forward Error Correction" , *Navigation, Journal of The Institute of Navigation*, Vol. 44, No. 1, Spring 1997
- [5] RTCA SC-159, "Minimum Operational Performance Standards for Global Positioning System/Satellite Based Augmentation System Airborne Equipment," RTCA DO-229D, Change 1, Feb 1 2013
- [6] Lo, Sherman, Chen, Yu Hsuan, Reid, Tyler, Perkins, Adrien, Walter, Todd, and Enge, Per, "The Benefit of Low Cost Accelerometers for GNSS Anti-Spoofing," Presented May 2017 at the Institute of Navigation (ION) Positioning, Navigation and Timing Conference, Honolulu, Hawaii.