Maximizing Aviation Benefits from Satellite Navigation
(FAA Award No. 08-G-07)
Technical Description of Project and Results
Stanford University
November 2012

1.0 Introduction

This report describes the key elements and results of Stanford University's contribution to the development of the Wide Area Augmentation System (WAAS), its underlying technologies, and alternative methods to provide global vertical guidance to aircraft. WAAS became the first operational space based augmentation system (SBAS) in July 2003. WAAS augments the Global Positioning System (GPS) with the following three services: integrity monitoring to improve safety; a ranging function to improve availability and continuity; and differential GPS corrections to improve accuracy. WAAS provides the continuity required for enroute and terminal area flight. It also has the integrity required for vertical guidance during airport approach. WAAS also protects the aviation community from uncertainties in GPS satellite replenishment.

Stanford University played a key role in the early development and prototyping of WAAS. Stanford then became actively involved in implementing WAAS and establishing its initial safety certification. Later, Stanford played a significant role in the improvement of WAAS to the point that today it provides LPV service to 100% of CONUS and more than 95% of Alaska. Stanford has played and continues to play a very active role in the modernization of WAAS to incorporate the future GPS L5 signal and new GNSS constellations such as the European Galileo system.

Stanford University has also taken a leadership role in investigating alternative methods to exploit signals from other core GNSS constellations. Stanford pioneered the development of Advanced Receiver Autonomous Integrity Monitoring (ARAIM) to compare signals from many satellites belonging to different constellations in order to achieve the required levels of accuracy, integrity, continuity, and availability in order to provide vertical guidance. The goal of this research path is to find lower cost methods to provide service while exploiting new constellations.

2.0 Optimize Performance for Legacy Users of WAAS

From 2009 through 2012, there one major (Release 3 in October 2011) and four minor software releases that significantly improved WAAS performance. The major release changed the ionospheric estimation process from the planar fit method to a new algorithm called kriging. The Kriging algorithm was developed at Stanford and improves upon the older Stanford-developed method by being better able to model more ionospheric conditions. Under kriging, the ionospheric estimator is less likely to encounter incompatible ionospheric conditions that can not be accurately estimated. Further, the overall uncertainty may be reduced and availability is improved. Figures 2-1 and 2-2 show an example of the improvement during a minor storm in 2010. This level of

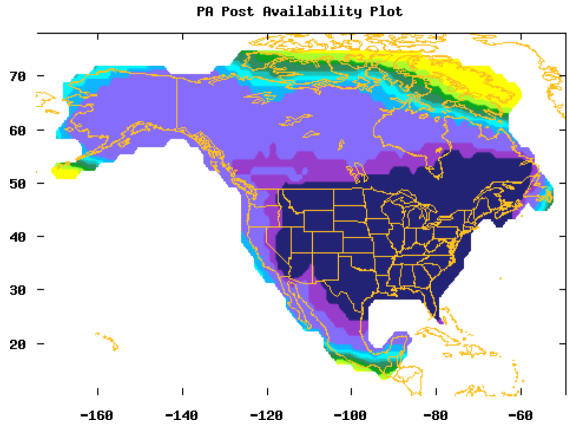

Figure 2-1. *This map shows availability for the planar fit algorithm during a minor ionospheric storm on April 5, 2010. Service is briefly lost to Alaska. 82.8% of Alaska achieved 95% availability, but none of Alaska achieved 99% or higher.*

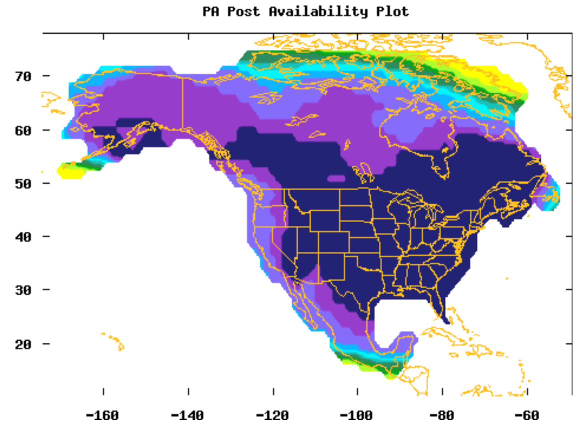

Figure 2-2. *This map shows what availability for the kriging algorithm would have been for the same storm. Most of Alaska has a much smaller loss of service. 89.5% of Alaska achieved 95% availability, 78.5% of Alaska achieved 99%, and 14.9% maintained 100% availability.*

improvement has been seen on many more days as the level of solar activity increases. Stanford also studied the observed behavior of the major storms from 2003 to better understand their spatial and temporal characteristics and incorporated this knowledge into the algorithm development.

Stanford University also improved the algorithms in the Signal Quality Monitor (SQM) that look for deformations in the broadcast signals from the satellites. There had been a few instances where multipath at a single reference station was large enough to affect the average value under certain circumstances. The SQM algorithm was changed to be more robust in these rising satellite situations, where only a few stations can see the satellite. This has resulted in a reduction in false alerts, and a small improvement in system continuity.

The geostationary (GEO) satellites used to broadcast the WAAS messages also provide ranging measurements to the user. Unfortunately, the control loops for the GEOs are not perfect and the code and carrier components of the signals are not fully coherent. This offset leads to bias errors on the users ranging that can grow up to a few meters. As this incoherency has changed over time and with equipment, bigger and bigger values need to be protected. Measurements made for Release 3 indicated that the bias error could be large enough to require that a larger minimum User Differential Range Error (UDRE) be broadcast. This larger value would lead to a loss of availability, particularly over Alaska. Stanford developed a new analysis tool to better account for the combination of larger GEO biases, small GPS biases, and other random error components, to demonstrate that the existing minimum UDRE values were adequate. This new tool, the GEO Code Minus Carrier Incoherence (CMCI) tool better accounts for all the errors so it simultaneously strengthens the integrity argument and increases the maximum tolerable bias error. Stanford developed the safety analysis and has provided the tool and accompanying documentation. Further, application of this tool may even allow for lowering the minimum UDRE value, which would lead to an improvement in availability and

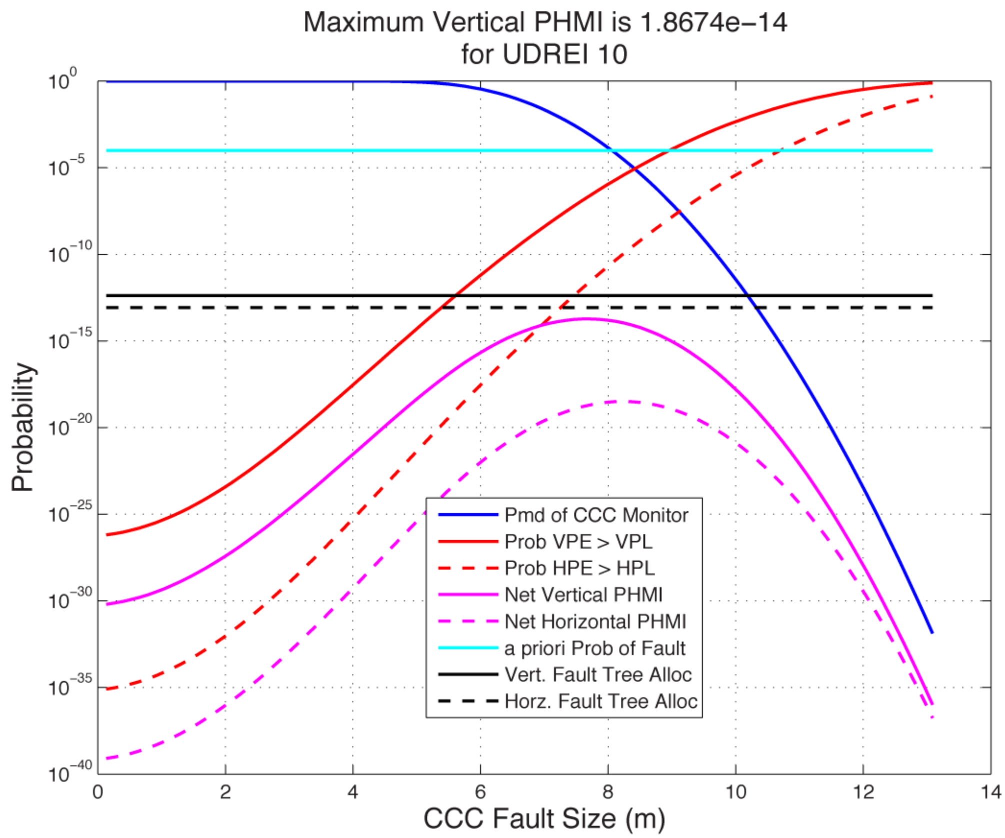

Figure 2-3. *This is an example of a specific GEO CMCI analysis. The cyan line shows the prior probability of there being a fault. The blue line shows the probability of the CCC monitor missing a certain size fault given a specific UDRE value. The red line shows the probability of the same fault leading to a hazardous situation. The product is the magenta line, which remains below the allocation (black line) for all possible values of the bias. The analysis is conducted separately for vertical position errors (solid lines) and horizontal position errors (dashed lines).*

continuity. Figure 2-3 shows an example of one of the many components of the GEO CMCI analysis tool.

In addition to system performance improvements, the Release 3 analysis had other significant goals. Prior releases had major analysis tools that were under control of the WAAS prime contractor, Raytheon. While the FAA, Stanford University, and others had reviewed the outputs of the tools, they had not been independently run and confirmed. For Release 3, the FAA independently replicated all major safety analyses. Stanford had a lead position in identifying all of the analyses and assigning teams to transfer tools and knowledge to the FAA. This was a significant effort on the part of all parties. Many small details had not previously been documented and the analyses were not all initially repeatable. By ensuring constant communication and working to resolve each issue, the safety analysis for Release 3 was finally fully reproduced and captured outside of Raytheon. This was a major step to raising visibility into all of the analyses and to transitioning the FAA towards an "organic" or in-house capability to maintain the WAAS software. Stanford University and other FAA support contractors continue to assist the FAA in maintaining and updating these safety analyses and tools.

One part of this move towards organic capability uncovered issues with how the ionospheric threats were identified. To study the ionosphere, measurements from all of the reference stations are processed to remove outliers. This process combines code and carrier information to reduce multipath and votes between co-located receivers to remove artifacts. The resultant output is called "supertruth." It was discovered that this process was not repeatable and did not lead to consistent results. A significant amount of effort was put into identifying performance differences and ultimately discovering sources of discrepancy. As a result, the overall supertruth process was improved and now it is fully captured and repeatable.

In 2010, the solar cycle started increasing towards it maximum phase which is expected to peak in mid-2013. WAAS began observing large ionospheric delays in southern Mexico from newer reference stations that had not been in place during the previous solar maximum. This increased ionospheric error tripped the range domain monitor (RDM) due to its having an overly conservative bound on the possible magnitude of the reference station clock error. Unfortunately the RDM trip affected all satellites that could be seen by the Mexico reference station and not just those with larger ionospheric delays. Stanford University analyzed the effect of the RDM clock bound and proposed a solution that sufficiently bounded the reference station clock error, but did not negatively affect all satellites in view. This new clock bound is much more resistant to large individual errors, but still fully protects the user. This fix has been tested and will be put to field in 2013. In early 2012, a stopgap measure was taken to remove the southernmost grid points from the ionospheric mask. This action prevents WAAS from processing measurements associated with the equatorial anomaly crest. Stanford is working with the FAA to look at restoring these grid points after the RDM fix is implemented.

Stanford has also actively been working with the international community to address issues with how the avionics process data from WAAS and similar systems. The European counterpart to WAAS is called the European Geostationary Navigation Overlay System (EGNOS). EGNOS does not implement ranging with its GEOs and as a result had some differences with its GEO ephemeris and almanac messages. Unfortunately, these messages are used to identify which service provider is using which GEO, but this usage had not been made clear in the Minimum Operational Performance Standards (MOPS), the document that describes the use of the data. Stanford worked with the FAA, the Europeans, and the avionics manufacturers to resolve this issue and help to create an update to the MOPS to clarify this issue. EGNOS has also changed their operation to be consistent with WAAS.

Stanford continues to work to improve performance. We have analyzed a new method for generating and assuring the broadcast UDRE that could result in safely broadcasting much smaller values. This same philosophy can be applied to the RDM. We are looking at using the receiver measurement data more efficiently to improve both accuracy and availability. These could either be put in place for the legacy L1-only user or they may enhance operation for a future L1/L5 user. Stanford continues to work closely with the FAA to evaluate these options and determine whether they should be implemented and when to put them into place.

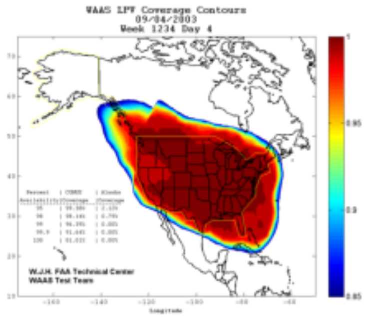

Figure 2-4. *This map shows availability for WAAS at IOC in 2003*

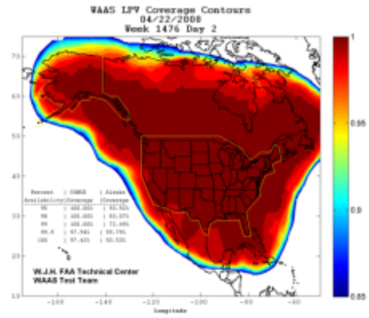

Figure 2-5. *This map shows availability for WAAS after adding reference stations and other improvements in 2008*

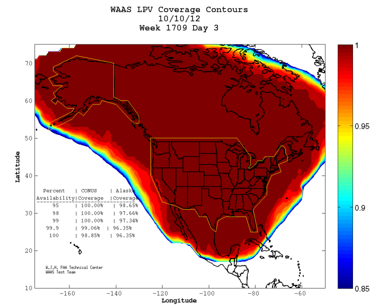

Figure 2-6. *This map shows availability for WAAS after adding kriging in 2011*

WAAS service has continued to improve since its Initial Operating Capability (IOC) at commissioning in 2003. By 2008 there had been many improvements including better handling of ionospheric threats, and new reference stations in Alaska, Canada, and Mexico. Figures 2-4 through 2-6 show this progression from CONUS coverage only to nearly full North American coverage. Stanford University researchers, funded by this FAA cooperative agreement, were the primary designers of the majority of the algorithms that contributed most significantly to this improvement.

3.0 Protect WAAS Reference Stations and Airports from Personal Privacy Devices

Personal privacy devices (PPDs) are low-cost jammers that mask GPS signals, so that the location of the host vehicle is not revealed to other parties. Although it is illegal to use PPDs in the United States, they are being used and have caused problems for GPS users. PPDs have affected operation of the Local Area Augmentation System (LAAS) and have affected measurements at individual WAAS reference stations. The FAA and the Zeta Corporation collected statistics on the frequency of observed PPD interference and provided this information to Stanford University so that we could model the potential impact of these devices on performance. Stanford had previously developed its Matlab Algorithm Availability Simulation Toolset (MAAST) that could predict availability under different conditions. We modified our MAAST code to conduct a Monte Carlo analysis with simulated PPDs affecting different reference stations. Conservatively following the observed pattern of PPD interference, a simulated PPD caused a loss of tracking for one minute and afterwards the code noise and multipath (CNMP) error bound was reset to maximum and slowly decreased over time as would occur in the operational system.

The FAA identified 16 stations most at risk for interference from PPDs given their proximity to larger populations and busy highways. They also identified 10 events a day, where all signals were lost, as a conservative upper bound on expected behavior. Currently, there are far fewer than 10 events a day and only low elevation satellites are typically affected. MAAST was modified to implement 10 random events concentrated during daytime commute hours and at these 16 most at risk stations.

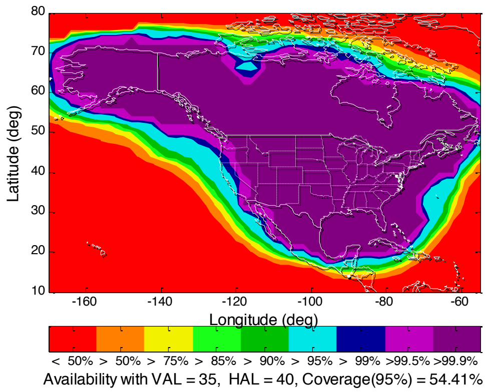

Figure 3-1. WAAS availability coverage for current constellation, baseline No PPD Case

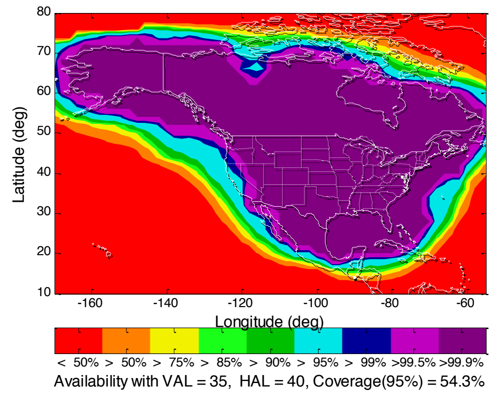

Figure 3-2. Average availability coverage for 10 outages in 24 hours for current constellation, all satellites experience outages

Figures 3-1 and 3-2 show results from this analysis. As can be seen in the figures, there is very little impact from PPDs to WAAS at the simulated levels. WAAS has 38 reference stations and this tremendous amount of redundancy means that WAAS is very robust to interruptions at individual stations. Also the 16 most at risk station are interior stations that have less impact on the outer boundary of coverage. Had more of the outermost stations been at risk, the effects could be larger. As a sensitivity analysis, we increased the number of events from 10 a day to 100 a day and saw very little further degradation.

Stanford also analyzed the more realistic scenario of having PPDs only blank satellites below 35 degrees elevation. As expected this reduced the impact even further. Even for the extremely pessimistic scenario of 100 events per day, there was only a 0.28% reduction in coverage over the baseline scenario. Thus, we do not see PPDs as an imminent threat to WAAS performance, although we should continue to monitor for their impact and be prepared to take action in the event of unforeseen growth in their use. In previous years, Stanford developed an antenna that had a much sharper cutoff for low elevation emitters. Assuming PPDs are primarily on the ground, such an antenna could help the reference station be relatively immune to their interference. The cost of these antennas is that they also mask out low elevation satellites as well, so a reference station with such an antenna may not be able to reliably track a satellite below 10 degrees. At the moment, this cost does not seem justified, but we could apply it in the future if a particular reference station is exposed to numerous PPD events.

Stanford has also developed control radiation pattern antennas that can steer nulls toward interference sources and gains towards GPS satellites. This work is being done as part of our efforts to support Alternate Positioning, Navigation, and Timing (APNT) and is described separately. Again this does not seem necessary, but could be considered in the future. Such antennas can also be useful in identifying the origin of the interference and could assist in finding the PPDs and shutting them down.

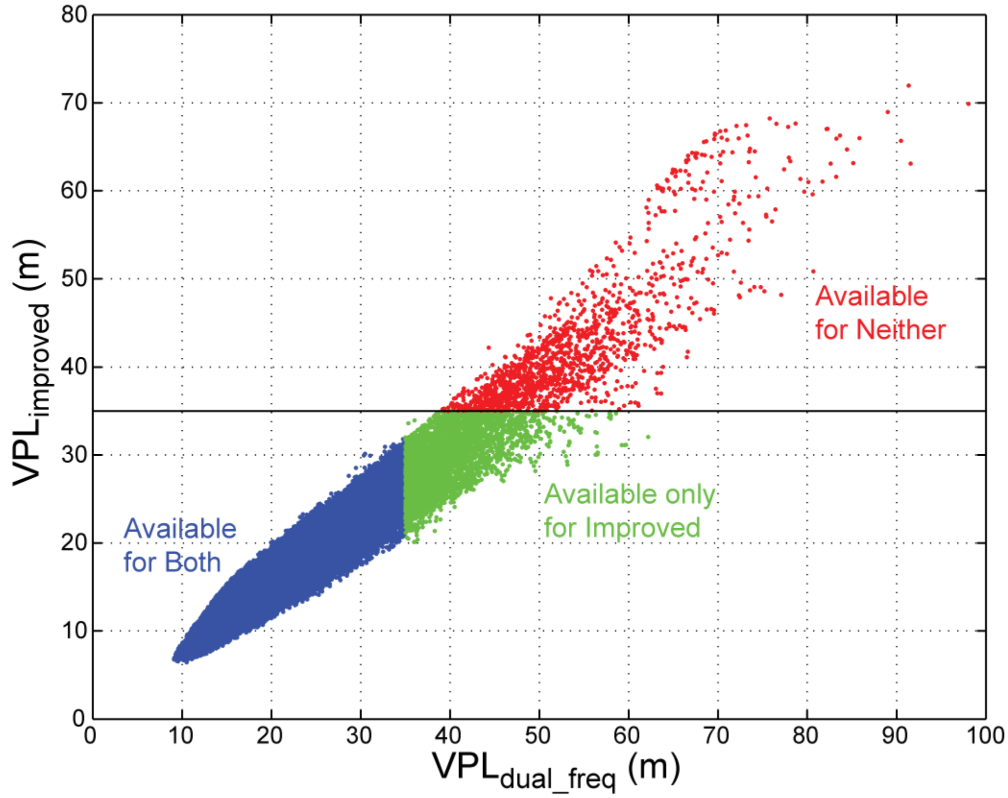

Figure 4-1. *The improved VPL is compared against a straight-forward update of the VPL for a days worth of geometries. Anything below 35 m would be declared available. The blue points are below 35 m for either VPL. The green points are available only for the improved VPL. As can be seen, the improved VPL makes many more geometries available while maintaining safety.*

4.0 Dual Frequency SBAS

In addition to improving performance for current users of WAAS, Stanford is actively engaged in preparing for the next phase of WAAS that will utilize both the GPS L1 and L5 signals. This second signal at L5 will allow aviation users to directly estimate and eliminate the ionospheric delay affecting their ranging measurements. This removal both eliminates the possibility of losing vertical guidance service due to ionospheric storms and dramatically increases the coverage area of the vertical guidance service. However, dual frequency WAAS service is contingent upon having a sufficient quantity of GPS satellites with L5 capability. Initially the Air Force intended to have a full constellation of such satellites by 2018, but the extended life of the existing satellites has slowed down the launching of this new capability.

Stanford University has proposed a new position error bounding approach. This bound, termed the Vertical Protection Level (VPL), addresses some shortcomings of the current approach. In particular, it allows for the explicit handling of biases and single satellite faults. These changes allow for more geometries to be declared available because the new VPL more accurately reflects the effects of potential fault modes. Figure 4-1 shows an example result from this VPL change. This approach better matches the integrity approach taken by the WAAS monitors and increases availability. Figures 4-2 and 4.3 show the improvement in coverage resulting from the change.

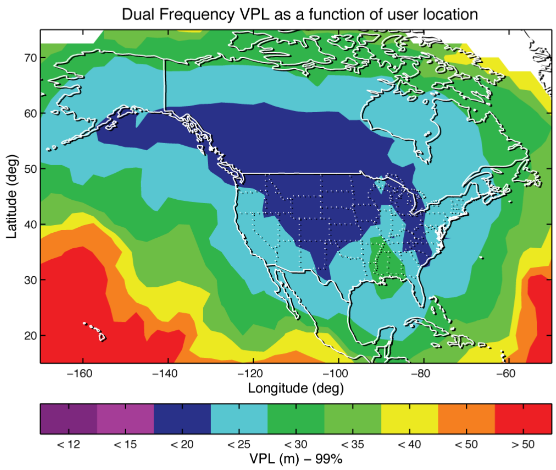

Figure 4-2. WAAS for nominal dual frequency VPL

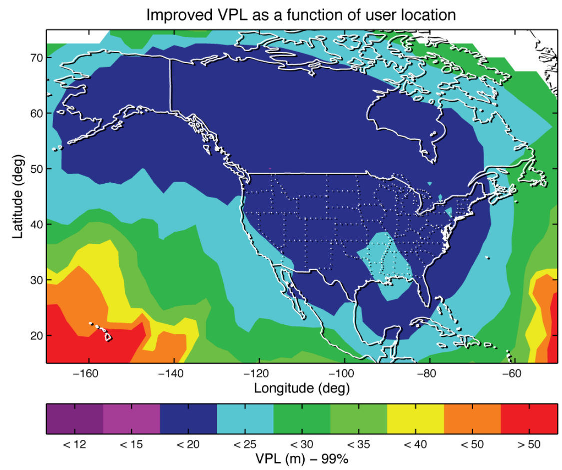

Figure 4-3. WAAS for improved dual frequency VPL

Given the overall integrity bounding approach that is implied by the VPL, we investigated what changes are required to the ground system to support it. Stanford has undertaken significant activity in this area in recent years. We evaluated all of the existing Algorithm Description Documents (ADDs) associated with the current WAAS integrity monitors. We then updated them to be consistent with the replacement of the L2 frequency by the L5 and to monitor threats consistent with the proposed dual frequency VPL approach. These updated ADDs were distributed to the FAA and their supporting contractors and updated with their comments. These ADDs will be distributed to potential contractors to describe the safety software that will be implemented in the dual frequency update to WAAS. Stanford also investigated the overall architecture approach to the integrity monitors and identified which ones would no longer be needed and which needed modifications. We also investigated the safety analysis and identified fault tree allocation changes that would help ensure that the targeted levels of performance would be met. We used our new GEO CMCI tool to analyze the performance of the SQM and CCC monitors.

Another significant effort towards moving to L5 usage is the development of the L5 MOPS that describes the messages broadcast on the GEO L5 signal that support users of the L1/L5 iono-free frequency combination. We have led the development of the MOPS and have proposed new messaging schemes that make more efficient use of the available message bandwidth. The current MOPS support a single constellation of single frequency corrections. Our proposed design supports up to four constellations of dual frequency corrections. We also improve upon the existing MOPS in several key ways. We improve accuracy by reducing the quantization error in the messages. We combine all satellite corrections into a single message greatly simplifying the overall message structure. We support flexible update rates and bounding values to allow service providers to optimize their message streams. Our proposal supports many different types of SBAS satellite orbits, not just geostationary. We robustly identify the PRN and service provider ID of the broadcasting satellite. We have coordinated our proposal with international SBAS service providers and with avionics manufacturers. We have

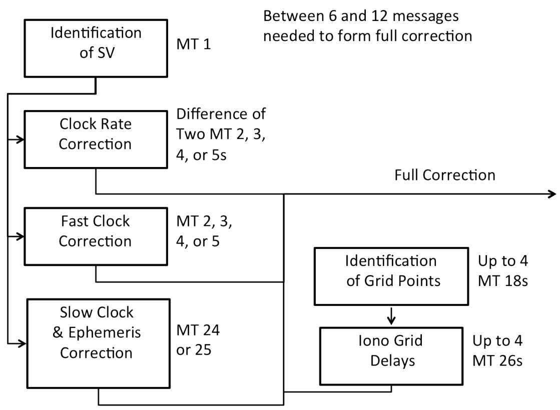

Figure 4-4. *A high-level schematic of the messages and components required to correct a single satellite under the current L1 SBAS MOPS.*

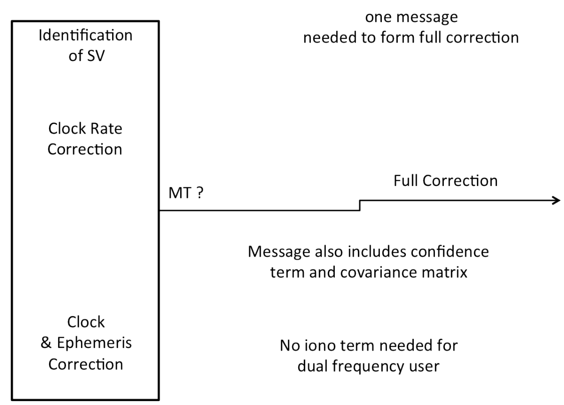

Figure 4-5. *The corresponding schematic of the proposed message structure for the L5 SBAS MOPS. All corrections are contained in a single message.*

received feedback and are ready to draft firm proposals that can be made into official updates. Figures 4-4 and 4-5 highlight the advantages of combining the corrections into a single message. Instead of combining information across up to 12 different messages, only a single message's contents are required.

Another change being made to GPS is the usability of the L2 P(Y) signal for codeless and semi-codeless receivers. After 2020, the Air Force no longer guarantees that these signals will function as they have in the past. WAAS had formulated a plan to switch to L5 before that time. However, it is becoming clearer that L5 will not be available in time. Therefore, we have begun to study our back up plan of first switching to L2C before moving to L5. At first glance it appears that L2C should be a relatively simple switch. However, we are still in the process of evaluating the impact on the estimation of the receiver inter-frequency bias estimates. There may be some minor modifications required to our estimator and the associated confidence bounds.

L1/L5 use still has some ionospheric concerns. Although the first order ionospheric delays are removed, higher order terms remain. These could cause a small amount of error and need to be bounded. Stanford has analyzed these terms and found them to be small for most times and places. It is possible to monitor the total delay observed and bound a small fraction of that observed value. Typically this is less than 0.2% of the total delay.

Another concern is scintillation. This effect can cause loss of lock on the signal to the satellite. It is not a significant problem at the mid-latitudes where WAAS principally operates, but it can be a problem near equatorial latitudes and at polar latitudes. WAAS does experience some scintillation effects at its northernmost stations, however, it does not have a very significant impact on performance. The main concern is for equatorial regions. Stanford has collected data and identified important characteristics of scintillation behavior such as fade duration and mean time between fades. We then conducted a sensitivity analysis and determined that the most important parameter was the time it takes the receiver to reintroduce a satellite into the navigation solution after a

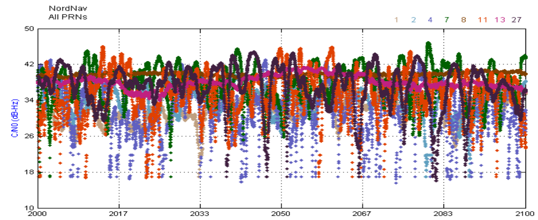

Figure 4-6. *A 100 second example of signal to noise ratios for some of the most severe scintillation observed. Each color represents a different satellite. Seven of the eight satellites in view are affected.*

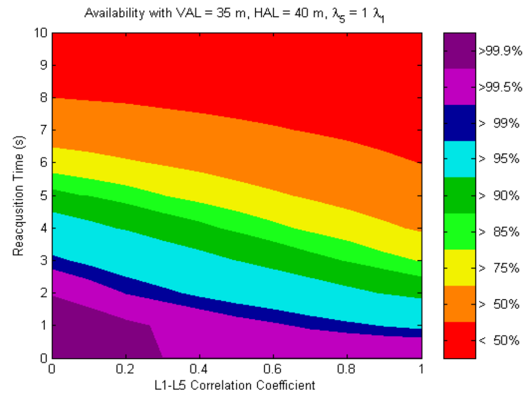

Figure 4-7. *The availability as a function of reacquisition time and correlation between L1 and L5 fades. For severe scintillation this correlation is expected to be low.*

deep fade. We found that as long as the initial geometry is strong, if the receiver can recover from the deep fade within two seconds, availability can be maintained in even the strongest scintillation environments. Figure 4-6 shows the signal to noise ratio for eight satellites that were in view of a station experiencing very severe equatorial scintillation. Figure 4-7 shows that even during this extreme period, availability can be kept high provided that the satellite reacquisition time is short and the correlation between fading on the L1 and L5 frequencies is low. Our best understanding is that correlation is high for moderate to weak scintillation, but low for very strong scintillation. We continue to collect data in equatorial areas and as more satellites broadcast L5 signals we will gain a better understanding of the risk posed by scintillation.

Stanford has also conducted research into making receiver tracking more robust against scintillation. We investigated the performance of aiding receiver tracking loops with inputs from inertial measurements and from the other GPS tracking channels. We used the scintillation data together with clock noise models, aircraft vibration models, and an RF GPS simulator to create a very realistic channel model of the effects of severe scintillation on a receiver in an aircraft. This signal was fed into an actual receiver that implemented the aided tracking loops. We found that the aiding could provide an extra 6 dB of fade resistance. This means that signal fades need to be at least four times stronger to cause a loss of lock and therefore fewer instances of deep fading will affect tracking. Combined with a faster reacquisition time and the ability to track satellites from as many constellations as possible, we are very optimistic for the ability of future receivers to operate in regions that experience strong scintillation.

Another concern over dual frequency operation is the increase of small error sources due to the iono-free combination. Certain error sources can be inflated more than 2.5 times their L1 only value when the L1 and L5 signals are combined in a way to cancel the ionospheric delay. One of these error sources is signal deformation. Because each satellite signal is not ideal and not identical to the signals from other satellites, there are small differences created when tracking these satellites. These nominal biases are essentially constant, although there can be changes with aging or component swapping.

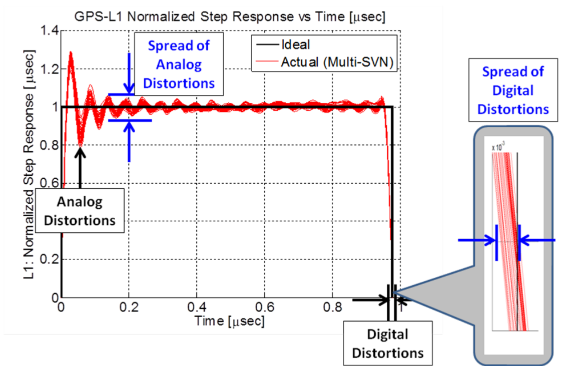

Figure 4-8. *Nominal Analog and Digital Distortions as measured for different GPS satellites. Data collected with SRI 46m-Dish: Aug 2008, Jul 2009, Aug 2010.*

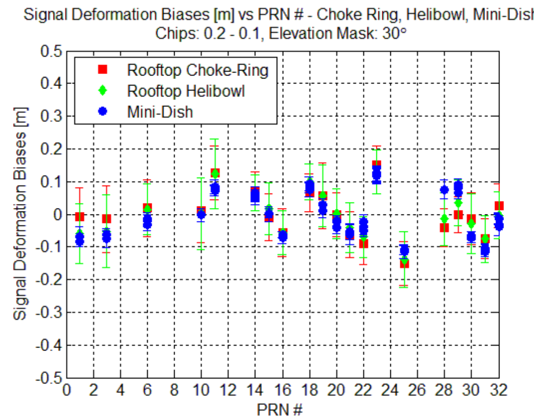

Figure 4-9. *Signal deformation biases and standard deviations distributions for individual satellites.*

These biases can be of order 30 cm for the current L1 system. However, for the iono-free combination they could approach 1 m. Stanford has carefully analyzed these biases using our 47 m big dish antenna, our small 1.8 m dish antenna and a variety of other multipath limiting antennas and receivers. We have observed that if the aviation receivers can be made more similar to one another, the L1 only bias could be reduced to less than 10 cm and the iono-free error to less than 15 cm. Figure 4-8 illustrates nominal signal variations that exist on the satellites. Figure 4-9 shows measured biases using a several different methods. All methods are in good agreement and show that for small correlator differences the bias error can be small.

5.0 Multi-Constellation Navigation

In addition to dual frequency WAAS, Stanford is actively investigating the optimal methods to make use of the new signals that will be available from other constellations. This includes incorporating additional constellations into WAAS and alternate methods to provide vertical guidance. As mentioned above, the dual frequency SBAS MOPS are being developed to support multiple constellations. The alternative method is through Advanced Receiver Autonomous Integrity Monitoring (ARAIM), a method that was developed as part of the GNSS Evolutionary Architecture Study (GEAS). Stanford acted as a co-chair of the GEAS and was one of the principle contributors to the algorithms and architectures developed under this study. Stanford also acted as the lead author on the GEAS reports that documented the progress and final outcomes of the GEAS.

Later, the FAA partnered with the European Union (EU), specifically the European Commission (EC) and the European Space Agency (ESA), to create a bi-lateral group to continue the investigation into ARAIM. Stanford has participated in this new US-EU ARAIM subgroup in the roles of co-chair and principle algorithm developers. Stanford also took the lead in writing two joint papers and the Milestone 1 report which is the first comprehensive documentation of the subgroup's activities and outcomes. As part of this subgroup Stanford helped to establish a parametric study of ARAIM performance and ran the analyses to evaluate performance for each scenario. Figure 5-1 shows a summary of

$P_{const} = 10^{-8}$

P_{sat}/URA	**.5 m**	**1 m**	**1.5m**	**2 m**	**3 m**	**3.5 m**	**4 m**
10^{-5}	100%	100%	100%	100%	100%	42.9%	3.4%
10^{-4}	100%	100%	100%	100%	100%	0	0
10^{-3}	100%	100%	100%	99.6%	6.6%	0	0
$P_{const} = 10^{-6}$							
10^{-5}	100%	100%	95.0%	51.5%	0	0	0
10^{-4}	100%	100%	95.0%	51.5%	0	0	0
10^{-3}	100%	100%	95.0%	51.3%	0	0	0
$P_{const} = 10^{-4}$							
10^{-5}	100%	98.5%	79.2%	.1%	0	0	0
10^{-4}	100%	98.5%	79.2%	.1%	0	0	0
10^{-3}	100%	98.5%	79.2%	.1%	0	0	0

Figure 5-1. *ARAIM availability is shown for three critical parameters: probability of constellation fault, probability of satellite fault, and user range accuracy (URA). As can be seen, URA is the most critical parameter, P_{const} is the next most critical and P_{sat} is the least critical.*

those results. This figure shows three tables corresponding to three different probabilities of faults that can affect multiple satellites within a constellation (P_{const}), faults that can affect a single satellite within a constellation (P_{sat}), and the expected one-sigma User Range Accuracy (URA) of each ranging signal. As can be seen, the URA is the most important factor and it should ideally be kept below 1 m (depending on P_{const}). P_{const} is the next most important parameter as it can also strongly influence availability. P_{sat} has little influence over availability, but it does change the amount of work required at the aircraft. The larger this number, the more failed satellites the avionics has to consider, meaning that it has to evaluate more subsets. Although it is too early to know how other constellations will perform, these results demonstrate what level will be required to make ARAIM successful.

Stanford has also been very active in the use of other satellites to prototype ARAIM and other safety critical uses of multiple constellations. When China launched the first of their medium earth orbiting (MEO) navigation satellites, Stanford was able to identify the code and demonstrate that it was in the Gold code family. China had not published this information, but by providing these codes Stanford enabled receiver manufactures around the world to track and use this first Chinese satellite. Stanford also characterized the signals from the very first GPS satellite to broadcast a true L5 signal. As did other researchers, we were able to identify small variations in the L5 signal with respect to the L1 and L2 signals.

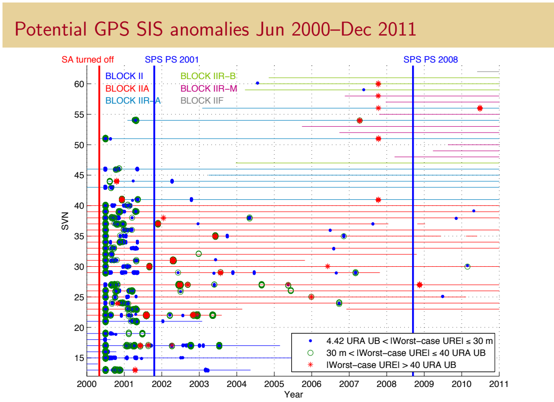

Figure 5-2. *Potential GPS anomalies over the last ten years.*

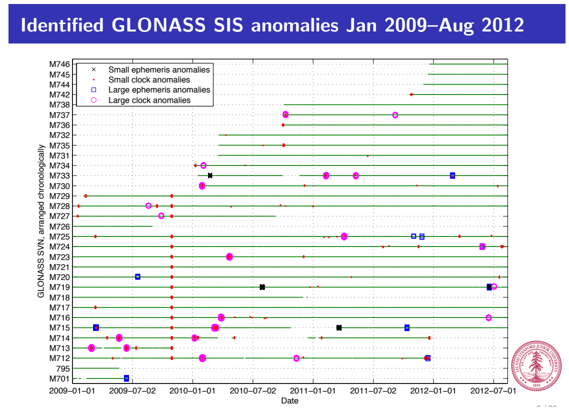

Figure 5-3. *Potential GLONASS anomalies over the last three years.*

Stanford also studied the historical performance of the GPS and GLONASS constellations. GPS was studied in depth for its performance over the last ten years, while GLONASS, which is only recently resurgent, had its performance evaluated over the last three years. It has been observed that GPS nominal performance is very good: the stated URA is usually at the minimum current value of 2.4 m while the in fact performance is typically better than 1 m. GLONASS performance is also quite good. The stated URA unfortunately is not always available to us but the in fact performance ranges from 1 to 5 m with the average close to 2 m. Both have distributions with more likelihood of larger errors than would be expected for a Gaussian. Thus, when bounded in an integrity sense, larger URA values are required.

Both systems also experience faults where the error is much larger than would otherwise be expected. GPS faults are defined as errors larger than 4.42xURA or most commonly 10.6 m. Since the URA is not always available for GLONASS and there is no official definition of a fault, nor is there a performance commitment, we declared a signal to be anomalous if the error is greater than 50 m. Figure 5-2 shows all of the anomalies found on GPS from 2000 – 2011. We found that GPS meets its current performance commitment of no more than three satellite faults a year. This supports a P_{sat} value of 10^{-5}. GPS also averages close to an hour in duration per anomaly, which is much better than the committed value of six hours. Figure 5-3 shows all of the anomalies observed for GLONASS from 2009 onwards. In 2011, we observed 30 GLONASS anomalies indicating that its fault rate is approximately ten times higher than GPS. P_{sat} for GLONASS should thus be at least 10^{-4} or larger. However, we notice that the trend is very favorable and the number of observed anomalies has decreased in each of the last three years. Therefore, GLONASS may also be able to achieve 10^{-5} or better in future years. The duration per outage has also improved over the last three years. Initially it averaged closer to ten hours per anomaly and now the average is closer to two hours.

We have also proposed additional monitoring constraints that could be placed on any of the constellations. These would monitor performance over multiple time-scales (from hours to years) to evaluate if the actual performance is bounded by the specified URA.

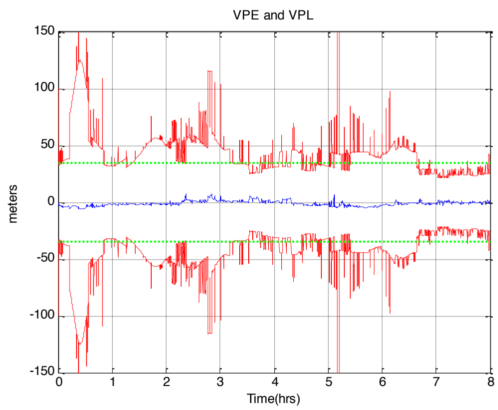

Figure 5-4. *ARAIM vertical position error and protection level using GPS only, using data collected at Stanford University.*

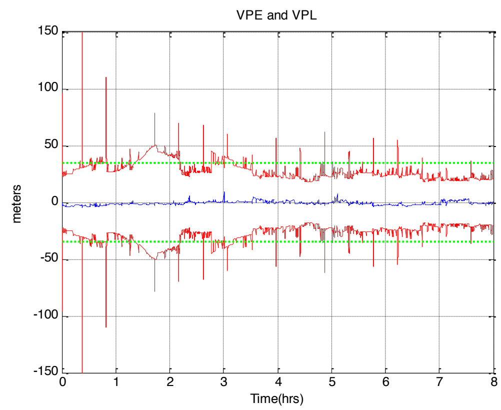

Figure 5-5. *ARAIM vertical position error and protection level using GPS and GLONASS together, using data collected at Stanford University.*

Further, chi-square checks were proposed to evaluate if the errors were correlated over multiple satellites. These checks could be part of an off-line monitoring system that could be used to improve confidence in the use of satellites for safety-of-life applications such as vertical navigation for aircraft. There is a larger effort being led by Stanford within the US-EU ARAIM subgroup to determine the exact monitoring architecture required to support ARAIM. When such an architecture is agreed to, it will become practical to directly compare ARAIM and SBAS in terms of performance and cost of maintenance.

Stanford has used the knowledge gained from monitoring the actual constellation performance and anomaly rates to determine the parameters that should be put into an actual evaluation of ARAIM using measured data from the GPS and GLONASS constellations. This is an ideal test platform for ARAIM as we already have two functioning constellations operating on two frequencies. Although the second frequency is not L5, much can be learned through practical application of the algorithm. We have observed that accuracy and availability can be enhanced by combining measurements from both constellations compared against using just GPS by itself. Figure 5-4 shows the vertical position error (blue) over an eight hour period using just GPS data. Also shown is the ARAIM vertical position level (red). Figure 5-5 shows the same period of time but now including the GLONASS measurements as well as the GPS. As can be seen, both the position error and the protection level are reduced by the combination.

A related activity is the application of operational requirements on ARAIM. WAAS has an approach service called LPV-200 that allows it to guide aircraft to within 200 feet above the ground. This procedure is based upon an understanding of the error characteristics of WAAS. In order to have ARAIM support the same procedure it is important to ensure that its error characteristics also support the requirements of the procedure. Stanford led an effort to understand the intention of the SBAS requirements as specified by the International Civil Aviation Organization (ICAO) and translate them into tests appropriate to ARAIM. There is concern that uncorrected GNSS signals will have less accuracy than SBAS and that ARAIM may not flag errors as quickly as SBAS can. Therefore, Stanford proposed tests to be performed in the aircraft to help ensure that when all such evaluations pass, the expected error distribution will meet the ICAO requirements. Stanford documented its interpretation of the ICAO requirements and its recommended approach. This was coordinated within the US-EU subgroup and brought back to ICAO. There is still ongoing discussion at ICAO as the requirements specified in their documentation have some ambiguities and there is a desire to revisit and revise the requirements at the ICAO level. Fortunately, there is sufficient agreement that the evaluation of ARAIM can proceed with the Stanford recommended evaluations.

6.0 International Outreach

Stanford is actively engaged in promoting and developing SBAS and ARAIM internationally. Stanford is strongly involved in the development of international standards that describe these systems. As previously mentioned, Stanford is the principle developer of the dual-frequency, multi-constellation SBAS MOPS. This is being done at RTCA, which has international participation. Stanford is also coordinating this activity at EUROCAE, which is primarily focused on the development of Galileo in Europe.

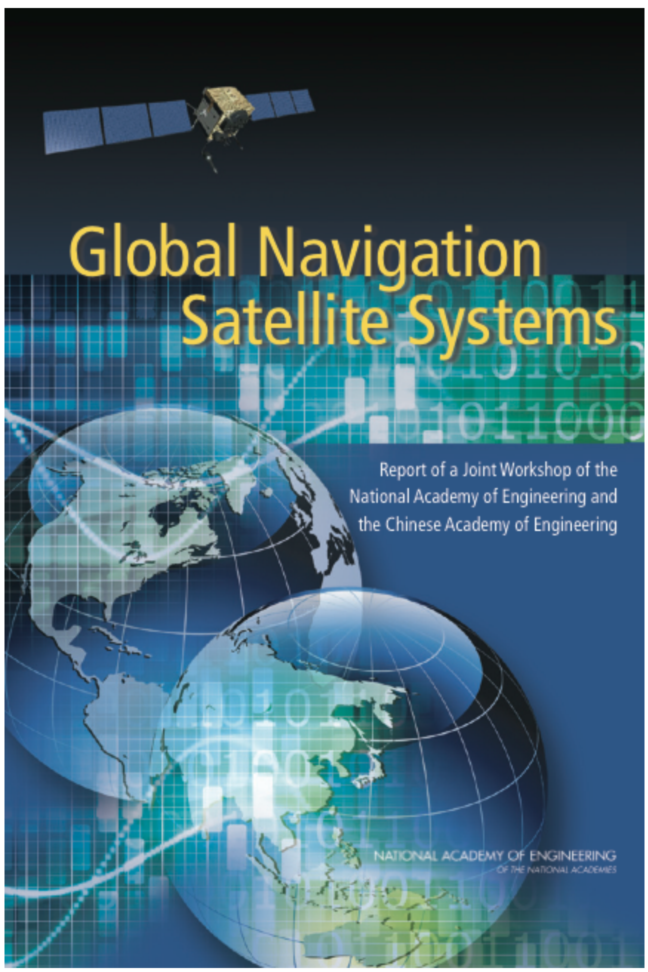

Figure 6-1. *The output report from the joint meeting of the national academy of engineering and the Chinese academy of engineering on the use of GNSS.*

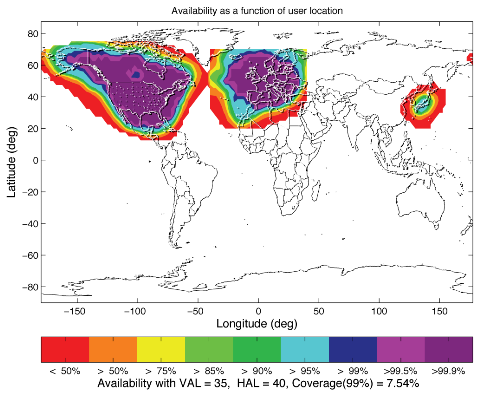

Figure 6-2. *Current, global availability of LPV-200 as provided by WAAS, EGNOS, and MSAS.*

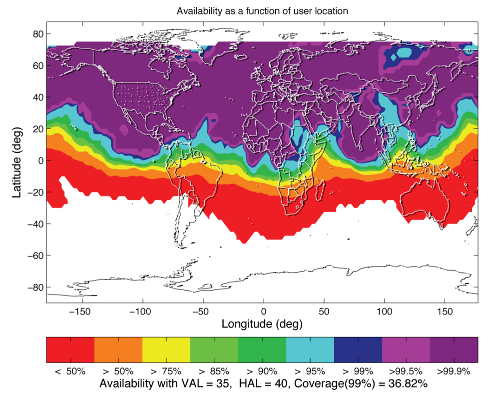

Figure 6-3. *Future availability assuming WAAS, EGNOS, and MSAS upgrade to dual frequency service and that GAGAN and SDCM also offer dual frequency service.*

Further, Stanford is a principle contributor to the Interoperability Working Group (IWG) a consortium of SBAS service providers that seeks to coordinate their services and harmonize their future plans. Both EUROCAE and IWG have had some limited participation from Russia, which provides an opportunity to coordinate the usage of GLONASS.

Stanford University has also reached out to China to coordinate the incorporation of its GNSS constellation, Compass. Stanford helped to coordinate a meeting between the US National Academy of Engineering (NAE) and the Chinese Academy of Engineering (CAE). This NAE-CAE meeting brought together leading developers of the US GPS system with leading developers of Compass. Stanford brought four speakers to this two-day meeting and coordinated the meeting minutes that were published with papers provided by the meeting participants into a joint monograph, see Figure 6-1.

Stanford also participates in the evaluation and certification of the Japanese SBAS: MSAS, and the Indian SBAS: GAGAN. Stanford has also met with interested researchers in Brazil and South Korea to discuss the possibility of developing SBAS systems in those regions. The net goal is to achieve seamless global coverage of SBAS. Stanford published a seminal article on the potential evolution of SBAS as it grows from today's set of three single frequency systems, see Figure 6-2. Figure 6-3 shows the availability if the three single frequency services upgrade to dual frequency service and that GAGAN and a Russian SBAS, called System for Differential Corrections and Measurements (SDCM), also offer dual frequency service.

Figure 6-4 shows the improvement in performance if either of the existing SBAS providers expand their networks into the southern hemisphere or new SBASs are fielded in this hemisphere. Now nearly all landmasses would have access to a high availability of LPV-200 service. The few weaker areas can be covered by either adding additional reference stations or by adding another constellation. Figure 6-5 shows coverage when all of the SBASs are upgraded to include both dual frequency and a second constellation,

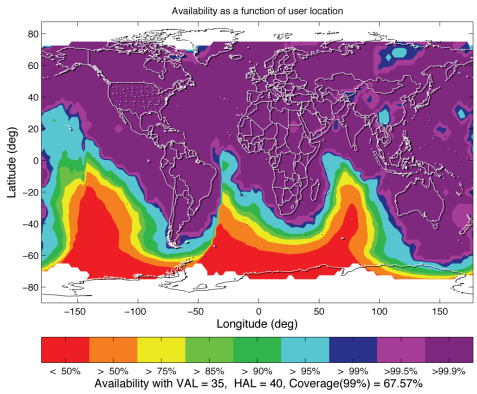

Figure 6-4. *Future availability assuming WAAS, EGNOS, MSAS, GAGAN, and SDCM are dual frequency and add some reference station in the southern hemisphere.*

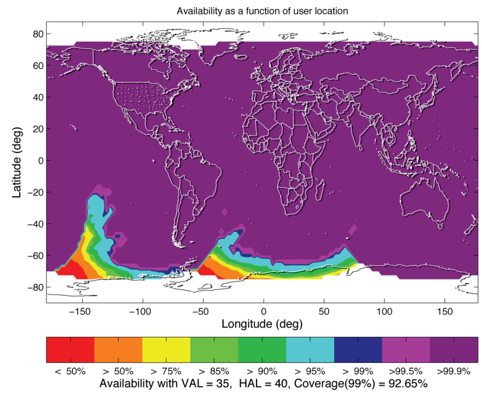

Figure 6-5. *Future availability assuming the same conditions as Figure 6-4, but now including SBAS corrections for both GPS and Galileo.*

modeled in this figure by the inclusion of GPS and Galileo. We can see that with these improvements the targeted goal of seamless global coverage would be achieved.

7.0 Summary

Stanford has played a critical role in the development, certification, and evolution of WAAS. Under this cooperative agreement, Stanford has worked with the FAA and Raytheon to implement significant algorithm improvements that resulted in a dramatic increase in the WAAS service region. Stanford has also worked to resolve anomalous performance characteristics of WAAS and has proposed several solutions that have led to improved continuity and better confidence in system performance.

Stanford is leading the effort to develop L5 WAAS algorithms and the L5 SBAS MOPS. Already we have recommended an architecture for the dual frequency safety monitors and drafted first versions of the associated algorithm description documents. We have also drafted the first version of the messages to be broadcast to the user and how the information is to be formatted by the ground and applied by the user. Ultimately, this will eliminate the current system's vulnerability to ionospheric disturbances and will allow the coverage region to be dramatically expanded.

There is widespread interest in utilizing new constellations for aviation as they become populated. The FAA has asked Stanford to evaluate alternate methods of using these new satellites that might be more cost effective than WAAS or LAAS. Stanford has taken a leadership role in the GEAS and the EU-US ARAIM subgroup in proposing and evaluating methods to compare the satellites against one another to identify potential errors. Stanford has investigated the ability of the aircraft to detect errors and what would be required for the FAA to certify such an approach. Much work has been done in defining the airborne algorithm. Stanford has written several detailed descriptions of this algorithm and provided these to avionics manufacturers and service providers. This method has great promise, but still has to be evaluated for the ground requirements.

Stanford has worked very actively with the FAA and its contractors to ensure the success of satellite navigation for aircraft. Much progress has been made, both to the operational WAAS system and to defining its forward evolutionary paths. Stanford has also made much progress in examining alternate methods of providing satellite navigation that may one day lead to more efficient means to provide vertical navigation.

This report highlights but a few of the contributions Stanford has made toward ensuring that aviation users take maximum benefit from satellite navigation. A great deal more work was actually performed. The following list of publications provides many more topics and details.