Validation of the WAAS MOPS Integrity Equation

Todd Walter, Andrew Hansen, and Per Enge
Stanford University

ABSTRACT

There has been widespread growth in the number of differential augmentation systems for GPS under development or in operation. Such systems are being developed by both civil authorities and commercial interests. These systems serve a variety of users and applications including precision approach for aviation, where the system provides vertical and horizontal guidance. Precision approach has very strict requirements for accuracy, integrity, continuity, and availability, and these become more stringent as the decision height decreases. To date, it appears that the Wide Area Augmentation System (WAAS) will be able to meet the accuracy requirements all the way down to a 200 ft decision height. The primary concern for such a system is that it always maintain integrity.

The WAAS Minimum Operational Performance Standards (MOPS) specifies how users combine error confidences from the different sources to form a position bound. The service provider guarantees that the error at any user location is smaller than the respective bound with a sufficiently high confidence. This paper describes the validation of the integrity equation. Actual data from the National Satellite Test Bed (NSTB), a prototype for WAAS, is compared side-by-side to simulated data. The difference between actual and expected performance is investigated in detail. It is shown that compared to the real data, the assumptions used in the integrity equation are conservative. Integrity is maintained both in the simulated data and in the live data. The comparison of the two data sets provides insights as to the actual probability distribution of the errors in the live data and about correlations between different error components. This knowledge helps to ensure that the full integrity requirements are always met. In the future, it may also be possible to utilize this information to increase the availability of the system.

INTRODUCTION

Integrity of a system is often extremely difficult to prove. One must demonstrate safe performance in the past and an expectation of continued safe operation even in the face of potentially unknown threats. Past safe performance can be demonstrated easily, but it may be difficult or impossible to gather enough data to meet stringent requirements at 10^{-7} or 10^{-9} levels. In addition, there is the question of whether all possible fault modes were adequately tested. In fact it would not be possible to prove the integrity of any system to a true skeptic.

In order to gain confidence in any system, one must be able to predict the performance of the system under both nominal and faulted operation. For example, if the errors have a gaussian distribution with certain means and variances under “fault free” and various faulted modes of operation, performance can be predicted if enough data is collected to determine those values. This system must also be robust against general fault modes to ensure safety in the face of unexpected errors.

The Wide Area Augmentation System (WAAS) [1] protects the users of the service by providing timely alarms and bounds on the error in the position solution. These bounds, called protection levels, provide an indication of the quality of service. In order for the system to be usable, the protection levels must be below predefined thresholds known as alert limits. The most challenging aspect of the system is to generate bounds which are large enough to always protect the user but small enough to permit the operation. At the center of this challenge is the integrity equation.

INTEGRITY EQUATION

The WAAS MOPS integrity equation is based on the concept that the actual pseudorange errors can be conservatively bounded at and beyond the 10^{-7} probability

level by a zero mean gaussian [2] [3]. The variance of this gaussian is described in the MOPS [4] and is largely based on information broadcast to the users from the geostationary satellite (GEO). Another principle of the integrity equation is that the pseudorange errors are uncorrelated with one another. This is a conservative assumption for navigation as the actual correlations appear to be negative. That is, they combine to reduce the overall positioning error rather than increase it. For example, errors common to all pseudoranges will affect the clock, but will not influence the navigation solution. Additionally, positive errors in the ionospheric estimation may lead to negative errors in the satellite clock and ephemeris terms. When these are combined, the two tend to cancel and reduce the overall error.

The integrity equation is based on covariance propagation from errors in the pseudorange domain, $\Delta\hat{\boldsymbol{y}}$, to the error in the position domain, $\Delta\hat{\boldsymbol{x}}$. This mapping follows from the navigation solution [2] and can be expressed as

$$\Delta\hat{\boldsymbol{x}} = \left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\Delta\hat{\boldsymbol{y}} \qquad (1)$$

where **G** is the observation matrix and **W** is the weighting matrix for the measurements. The matrix

$$\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1} \qquad (2)$$

is the full position estimate covariance matrix which is available when the position solution is calculated. The variance of the vertical position estimate is given by the third diagonal element of this covariance matrix,

$$\sigma_V \equiv \sqrt{\left[\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\right]_{33}} \qquad (3)$$

For WAAS, **W** is a diagonal matrix and the inverse of the i^{th} diagonal element is given by the variance for the corresponding satellite, σ_i^2, which is defined as [4]

$$\sigma_i^2 = \sigma_{i,\,flt}^2 + \sigma_{i,\,UIRE}^2 + \sigma_{i,\,air}^2 + \sigma_{i,\,tropo}^2 \qquad (4)$$

The four variance terms on the right represent the confidences for the fully degraded [4] [5] [6] clock and ephemeris corrections, $\sigma_{i,\,flt}^2$, the fully degraded ionospheric correction, $\sigma_{i,\,UIRE}^2$, the contribution from the airborne receiver, $\sigma_{i,\,air}^2$, and the tropospheric model correction, $\sigma_{i,\,tropo}^2$.

The Vertical Protection Level (VPL) is given by

$$\mathrm{VPL}_{WAAS} \equiv \kappa(Pr)\cdot\sigma_V \qquad (5)$$

where $\kappa(Pr)$ is a constant defined in the MOPS which depends on the tolerable probability of having an error greater than this value. For a 10^{-7} probability, $\kappa(Pr) = 5.33$. For example, if the UDREs, GIVEs, airborne receiver variance, and geometry combined to create a σ_V of 2 m, then $\mathrm{VPL}_{WAAS} = 10.66$ m. This would indicate that the aircraft would only have one chance in ten million of having a navigation error exceed 10.66 meters.

NSTB DATA

The NSTB data presented here was generated using the Stanford developed master station code [7]. It is an engineering prototype version of the WAAS master station code. It uses the stations shown in Figure 1 to generate differential corrections applicable throughout the United States. The differential corrections are put into the format specified by the WAAS MOPS [4]. Thus all the correction data is passed through the 250 bit per second bandwidth constraint. The system latency is also in effect. Although we are infrequently granted access to the geostationary delivery channel, we always apply the same latency as though we were.

The ability of the integrity equation to protect users has been borne out over time using the NSTB. Tests include nominal fault free operation and faulted modes. These faulted modes include the use of maneuvering GPS satellites which the Master Control Segment (MCS) had declared "unhealthy," but were monitored and differentially corrected anyway [8]. Also included are six second non-

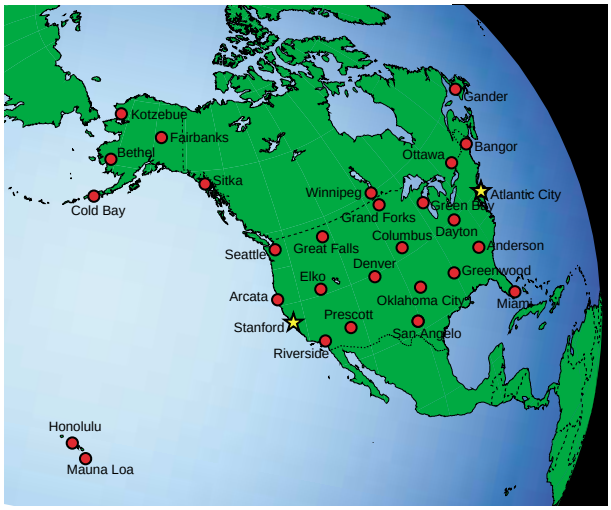

Figure 1. *The NSTB network. Each circle represents the location of a reference station. The master stations at the FAA Technical Center and Stanford University are shown as stars. [11]*

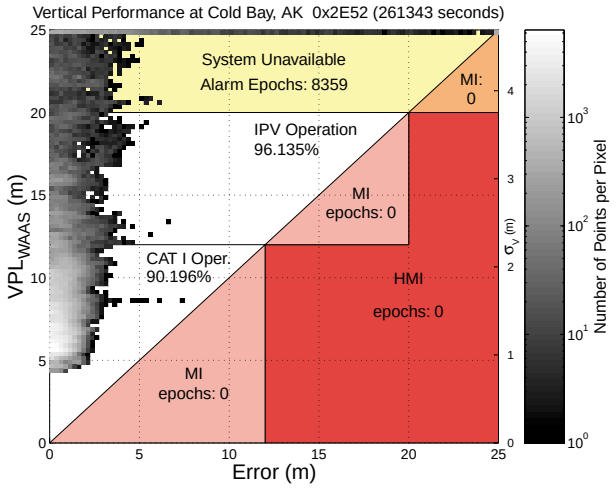

Figure 2. *Actual vertical performance at Cold Bay. This is a combined 2-D histogram of error and VPL showing accuracy, integrity, and availability. The errors caused by PRN 18 can be seen in the CAT I region.*

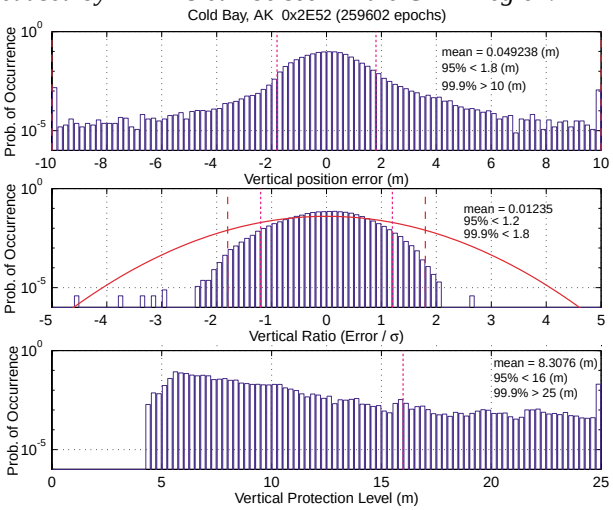

Figure 3. *Actual vertical performance at Cold Bay. These three histograms show the accuracy, integrity, and availability separately. Cold Bay had the worst availability of all our reference stations on these days.*

standard code outages [9] and accelerations larger than the specification in the Standard Positioning Service (SPS) [10]. In addition several days of solar storms have been recorded. Despite these events, the integrity equation was capable of protecting all users investigated.

It is easy to look at a data set after the fact and describe its characteristics. It is more difficult to predict future behavior, particularly the integrity of the system. We have developed tools to help us examine our data and rapidly identify problems. In particular we have different ways of representing the data. Figure 2 shows a result from three days in June of 1998. This figure shows a two dimensional histogram of the vertical navigation data. The error-VPL space is divided into 25 cm by 25 cm bins.

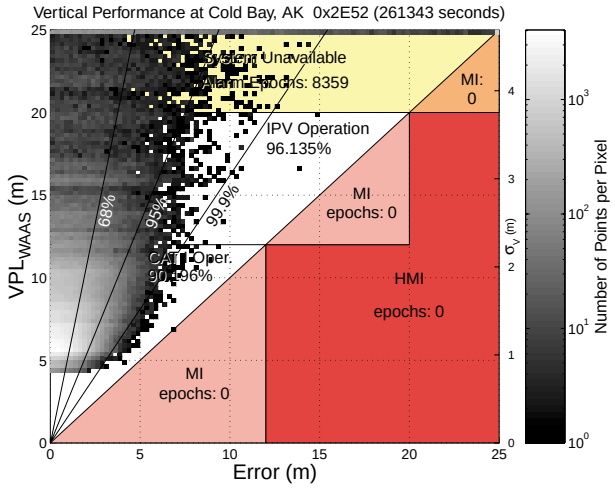

Figure 4. *Simulated vertical performance at Cold Bay. The errors are generated independently for each time and satellite. For reference, the lines of constant expected probability are shown.*

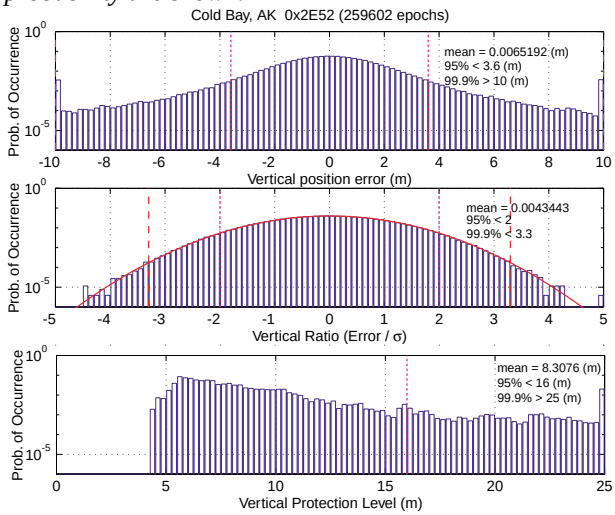

Figure 5. *Simulated vertical performance at Cold Bay. Each point is independently generated following a gaussian distribution. The center histogram shows the close agreement between the simulated values (bars) and the theoretical distribution (solid line).*

For every position solution the error is determined by comparing that solution to the pre-surveyed location of the antenna along with the VPL. These two values are quantized to within 25 cm values and the appropriate bin of the histogram is incremented. The bins which contained one or more data points are shown at their appropriate location in the error-VPL space. Matlab routines for generating these figures as well as exemplar code for formatting the data are available at our web site [12].

We commonly refer to these plots as VPL or “triangle” charts. The chart is broken into three main sections: an unavailable region where the VPL is too large to support

the desired navigation procedure, an unsafe region where the VPL supports the operation but the error is large enough to create Hazardously Misleading Information (HMI), and a usable region where both the VPL and the actual error are below the Vertical Alert Limit (VAL) so the system is usable and safe. However, it should be noted in normal operation mobile users do not have access to the actual error. They are entirely dependent on the accuracy of the VPL. The usable and unavailable regions are further divided. Above the diagonal line in the triangle chart, the VPL is always larger than the actual error which is the desired outcome. In the lower right hand regions the error has exceeded the VPL and provides Misleading Information (MI). Operationally, these regions are not necessarily hazardous. In the unavailable region, the procedure will not be flown since the VPL exceeds the VAL, while in the usable region the error is small enough to keep the aircraft within the obstacle clearance region. Despite these operational considerations, from a systems standpoint, the master station and/or integrity equation have failed to protect the navigation solution if the error becomes larger than the VPL. Thus all points should be above the diagonal line. As can be seen in Figure 2, all points are above the diagonal, so for these days, at Cold Bay, the integrity requirement was met.

The two most stringent applications for WAAS will be Category I precision approach (CAT I) and the still evolving instrument approach with vertical guidance (IPV). These operations have VALs of 12 and 20 meters, respectively. Thus we have further subdivided the usable region by these two values. If the VPL is below 12 meters then the system is usable for CAT I. If it is below 20 meters the system will support IPV. If the VPL is larger than 20 meters the system is unavailable for these stringent applications but may still be usable for non-precision approach, terminal, and en route operations.

This data is typical of the performance we see with our Testbed Master Station (TMS). It was taken at Cold Bay, which was chosen because it had the worst availability of all of our reference stations. This is helpful because a common question on the behavior of our system is the performance when the VPL is large (low availability). On these days, CAT I operations could have proceeded more than 90% of the time and IPV more than 96%. All other stations had higher availability, most better than 99.9% for both CAT I and IPV. Again, Cold Bay was chosen for illustrative purposes.

Figure 3 is another way of representing the data. Here instead of a two-dimensional histogram we have three separate one-dimensional distributions. They show, from top to bottom, the accuracy, integrity, and availability of the system. As with Figure 2 the probabilities are plotted on a logarithmic scale. The logarithmic scales are used to emphasize the tails of each distribution. We already have great confidence in the TMS operation the vast majority of the time. We are now interested in exceptions. The middle histogram is of primary interest. It shows the ratio of actual error to the one sigma value in Equation 3. For reference, a gaussian curve is also shown. This parabola would be applicable if all the errors were gaussian, zero mean, and independent. As can be seen, the errors are more tightly distributed than the gaussian reference. All but seven of the 259,602 data points have values less than 2.5. These seven points can also be seen in Figure 2 with a corresponding VPL of roughly 8.5 meters.

These seven points clearly belong to a different distribution and represent the greatest concern. These points do not follow the same probability distribution as the bulk of the errors. Upon investigation it was determined that these points were caused by excessive phase noise on a satellite. Our investigation is still ongoing, but the preliminary evidence is that on June 25th, 1998 around 5:19 UTC, SVN 18 (also PRN 18) exhibited phase noise some twenty times larger than the nominal value. All 23 of the NSTB receivers that were able to track this satellite saw identical effects on their L1 and L2 carrier phase measurements. These rapid fluctuations caused the clock error to vary more quickly than the fast corrections could track, leading to errors that were several meters in magnitude. Our prototype software does not currently have a trap for error sources of this type. All that is required is a simple clock acceleration check which both the WAAS and LAAS systems will have. However, we have not implemented such traps, in part to determine and characterize the effects of such errors. It should be noted that this error persisted for more than an hour yet only seven points were driven off of the main distribution and no errors were larger than their corresponding VPLs. More detail on this anomaly will be presented at the end of this paper.

SIMULATED DATA

The real NSTB data described above was used to generate a distribution of observation and weighting matrices. By using the same lines of sight, UDREs, GIVEs and airborne variances, we could create simulated errors for identical conditions as experienced with real data. The true range was defined by taking the known antenna locations and corrected satellite locations. To this true range we added simulated noise.

Initially we started with independently distributed zero mean gaussian noise. Each measurement is independent

of all others whether for another satellite or for a different time. This result is not intended to reflect actual performance but to simulate the conditions which are assumed by the integrity equation. As will be seen, this situation is more conservative than the actual performance. These results are shown in Figures 4 and 5.

One of the most common questions regarding the triangle charts is why the data points do not fill in the whole upper left hand triangular region. This pure gaussian data set points to part of the answer. Here, as can be seen in Figure 5, the errors are truly gaussian in distribution and fill in the tails as expected yet the distribution of errors does not seem to get much larger as the VPL gets worse. The reason behind this phenomenon is that there are fewer points sampled at large VPL. When the VPL is between 5 and 11 meters, there are greater than 2,000 data points per row. The number of data points per row rapidly drops down below 300 for VPLs above 15 meters. Thus the reason the tails seem to decrease at higher VPLs is not because the higher VPLs are overly conservative, but rather there are fewer data points in this area. More than 95% of the data has a VPL below 16 meters. The maximum likelihood value for a zero mean gaussian is zero, 68% of the points are contained within one sigma, 95% within two sigma and 99.9% within 3.29 sigma. For reference, these lines are shown in Figure 4.

The reference probability lines further illustrate why the upper points tend to be nearer to zero than intuition expects. These lines all converge at the origin. Since their slopes are greater than the diagonal line, there is more space between the 99.9% line and the 10^{-7} line as the VPL becomes larger. Thus there is a larger region with fewer points going into it resulting in a sparsely filled appearance. In reality, as many points are in the region as would be expected given the distribution.

The empty appearance in the upper right part of the safe region is not due to overcautiousness when the VPL becomes large. We are not discarding availability in the face of uncertainty. Rather, we are undersampling the region, and as always, the errors are more likely to be small than they are to be large. The comparatively smaller number of points we have at these upper VPLs are most likely to be distributed near zero. Even if the VPLs were uniformly distributed from 5 to 25 meters, the bulk of the data points would still appear to pull away from the diagonal line at the top as a consequence of the increasing separation of the lines of constant probability (68%, 95% and 99.9%) from the diagonal line. Only if each discrete VPL level had greater than 10^7 points (an impossibly large number for real data and nearly so for simulated) would the distribution of data appear to fill in the whole of the upper triangular region.

Temporal Correlations

Correlations in the data will also affect the shape of the distribution. The real NSTB data contains correlations across both time and satellite corrections. To better understand the distribution of the correction errors and any correlations between them, we studied the pseudorange errors from our real time master station. One difficulty with this method is the lack of a precise truth source. Although we know the location of our reference antennas very accurately, it is difficult to translate this into precise pseudorange errors. Instead we must contend with errors in the "truth" reference which may be larger than the errors in the differential corrections. We separated the correction errors into two effects: clock and ephemeris correction errors and ionospheric correction errors. The error contribution from the user receiver is included in the "truth" reference.

Clock and Ephemeris Correction Errors

The clock and ephemeris correction errors are covered by the appropriately degraded UDRE terms in addition to the user receiver variance. In our case the user is not an airborne receiver, but a static ground receiver at a known location. Even though the reference pseudorange measurement can be smoothed using dual frequency carrier phase measurements in post processing, there is still a non-negligible error left. This data for one satellite trace is shown in the top part of Figure 6. The data was examined for correlations in time. Common mode errors, as would be expected from the reference clock were removed. Figure 7 shows the autocorrelation for the real data as a solid line. It is clear that the errors are correlated for times less than 150 seconds.

In order to simulate this temporal correlation, we used a first order Gauss-Markov process. The errors were generated using

$$\begin{aligned} \varepsilon_{no\ iono,i,k} &= w_{no\ iono,i,k} + \gamma\ \varepsilon_{no\ iono,i,k-1} \\ \sigma^2_{\varepsilon\_no\ iono,i,k} &= \sigma^2_{no\ iono,i,k} + \gamma^2\ \sigma^2_{\varepsilon\_no\ iono,i,k-1} \end{aligned} \tag{6}$$

where $\sigma^2_{no\ iono,i,k}$ is given by

$$\sigma^2_{no\ iono,i,k} \equiv \sigma^2_{flt,i,k} + \sigma^2_{trop,i,k} + \sigma^2_{air,i,k} \tag{7}$$

and $w_{no\ iono,i,k}$ is zero mean gaussian variable with variance given by $\sigma^2_{no\ iono,i,k}$. The subscript k denotes time. The actual error applied to the satellite, $\delta_{no\ iono,i,k}$, is scaled to have the correct variance by

$$\delta_{no\ iono,i,k} = \alpha \sqrt{\frac{\sigma^2_{no\ iono,i,k}}{\sigma^2_{\varepsilon\_no\ iono,i,k}}} \varepsilon_{no\ iono,i,k} \qquad (8)$$

For uncorrelated white noise, the values of γ and α would be 0 and 1 respectively. To match the amplitude and shape of the autocorrelation of the actual errors in Figure 7, γ and α were set to 0.975 and 0.467, respectively. The resulting autocorrelation of the simulated errors is shown as the dashed line in Figure 7. In this section we wish to investigate only the effects of the temporal correlations. Therefore we used the value of α = 1, rather than matching the amplitude as well. In a later section we will further investigate the effect of variance scaling.

Ionospheric Correction Errors

The ionospheric correction errors are covered by the appropriately degraded GIVEs and then scaled by the obliquity factor to a slant range. Upon examination of these errors they appear to have very long time constants. One such error trace is shown in the bottom part of Figure 6. The errors are more biases than random fluctuations. However, the aggregate of errors, across different satellites, different reference stations, and different times, appear to fill in a gaussian distribution. The dominant error term most likely comes from our local dual frequency reference rather than the ionospheric corrections themselves. For simplicity, we modeled ionospheric errors as constant errors for the entire satellite pass. Thus, when a satellite first rises, it is assigned a random variable, $\delta_{iono,i,0}$, which is zero mean and gaussian in distribution with variance given by $\sigma^2_{i,\,UIRE,\,0}$. Then at each time thereafter the value is scaled by the new variance

$$\delta_{iono,i,k} = \delta_{iono,i,0} \sqrt{\frac{\sigma^2_{i,\,UIRE,\,k}}{\sigma^2_{i,\,UIRE,\,0}}} \qquad (9)$$

so that the distribution of errors always matches the gaussian expectation. The scaling by the variance also incorporates the expected obliquity factor term.

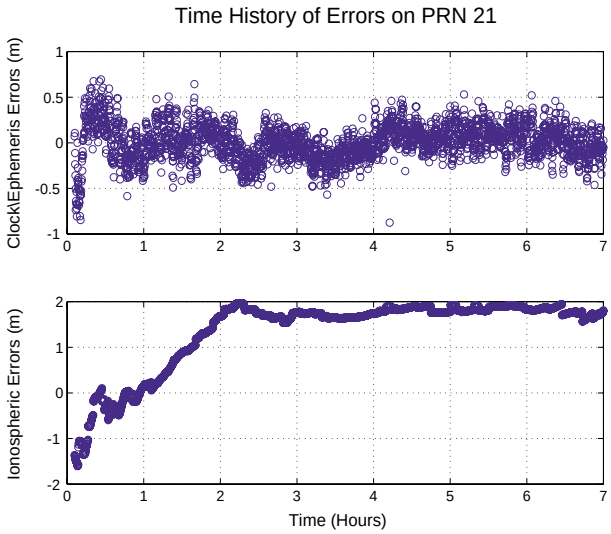

Figure 6. *Time histories of the clock and ephemeris errors and the ionospheric errors. Note that ionospheric errors are dominated by long bias-like errors.*

Temporal Correlation Results

The two simulated, time-correlated error sources are combined to form the total satellite correction error

$$\Delta\hat{y}_{i,k} = \delta_{no\ iono,i,k} + \delta_{iono,i,k} \qquad (10)$$

These individual satellite errors combine as described by Equation 1 to form the position error.

Figure 8 shows the triangle chart for the vertical navigation data simulated with temporal correlations, using the same distribution of VPLs as Figures 2 and 4. Because of the time correlations, there are fewer independent samples and this data fills out less of the histogram than did the uncorrelated simulated data. In Figure 9, the histogram still follows out the theoretical gaussian curve near to the origin. However, as the errors grow larger, the likelihood drops more rapidly. This is due to undersampling. Since the errors are correlated in time there are fewer independent samples and the histogram becomes narrower than the uncorrelated model. Note that the 95% and 99.9% values match the model to within about 5%. It is largely the low probability errors which are under represented. This distribution still does

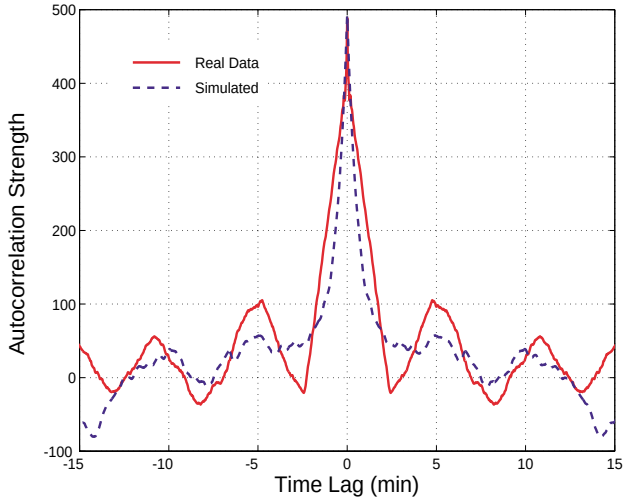

Figure 7. *The autocorrelations for the real and simulated clock and ephemeris corrections.*

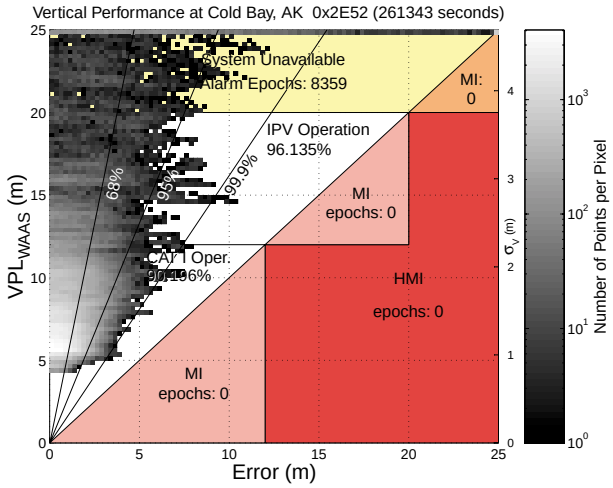

Figure 8. *Simulated vertical performance at Cold Bay. The errors are generated independently for each satellite, but include the time correlations described. For reference, the lines of constant expected probability are shown.*

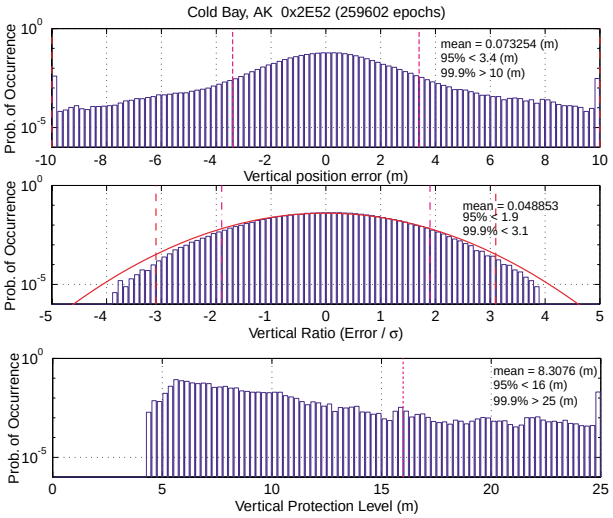

Figure 9. *Simulated vertical performance at Cold Bay. These errors follow the gaussian model close to the origin, but fall off at the tails.*

not match the actual data. Time correlation alone cannot explain the observed behavior. We must invoke another process to accurately simulate the true performance.

Variance Scaling

The simulated errors described above only incorporated correlations in time. The variance of the real data is of a different magnitude than the predicted variances in (4). Recall that the magnitude of the autocorrelation of the actual clock/ephemeris noise is less than half of the predicted magnitude. We now reduce the size of the simulated errors to determine this effect. Here the simulated clock/ephemeris data was generated using the value $\alpha = 0.467$.

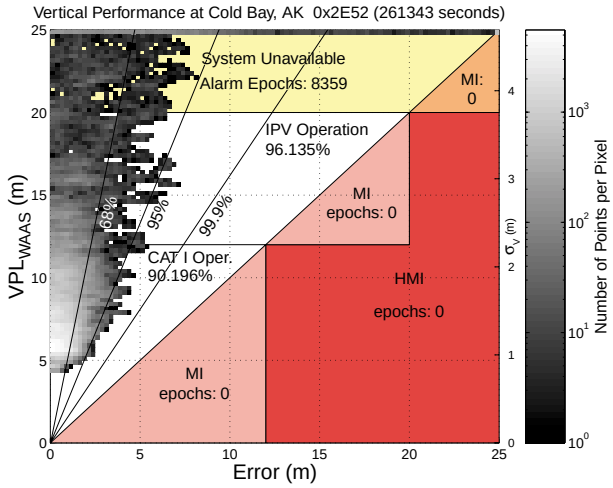

Figure 10. *Simulated vertical performance at Cold Bay. The errors are generated identically to the previous case (Figure 8), except that the clock ephemeris errors have been reduced to 46.7% of their predicted value.*

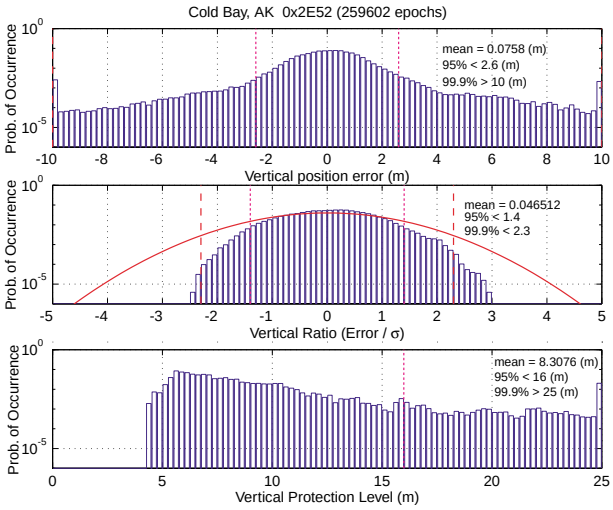

Figure 11. *Corresponding vertical performance histograms. These errors are closer to, but still larger than the actual errors in Figures 2 and 3.*

The actual ionospheric errors observed were about 20% larger than the prediction from the UIRE confidence. It is believed that this discrepancy actually resulted from uncertainty in the truth reference. Therefore, the simulated errors were not increased to match the observed errors, but were left at the lower predicted levels. If the larger values were correct, we would expect to see integrity violations as the GIVEs would not be bounding the actual error. However, this subject does require further investigation.

Figure 10 shows the triangle chart for the vertical navigation data simulated as described above, using the same distribution of **G** and **W**, and therefore VPLs, as Figures 2, 4, and 8. The reduced variance of the simulated

noise decreases the error in the navigation solution. The triangle chart for this data looks much more like the actual data in Figure 2. However, as can be seen in Figure 11, the distribution of errors is still more than 20% larger. The temporal correlations for these simulated data are conservative compared to the actual data. The clock/ephemeris terms match closely, but we assumed worse correlations for the ionospheric corrections. Also, the magnitudes of the errors either match (for clock ephemeris) or are smaller than the real data (for ionospheric corrections). Despite this, the resulting simulated position solution is still worse than the real data! This discrepancy arises because we have only incorporated correlations in time. There are obviously correlations between the different error sources which tend to reduce the overall position error.

SATELLITE ANOMALY

The integrity equation works because the errors are either uncorrelated or beneficially correlated. Thus they do not tend to combine in a worst possible fashion. The anomalous points in the real data shown in Figures 2 and 3 reflect a failure mode on a single satellite. SVN 18 exhibited large accelerations and jerk in its clock signal. As mentioned, all 23 stations viewing the satellite observed identical behavior in their L1 and L2 carrier phases.

Figure 12 shows the time history of clock error as determined by our TMS for both PRN 16 and PRN 18. Notice that PRN 16 exhibits the expected behavior for selective availability. In this figure we have also placed a simulated WAAS message. A fast correction message is generated every 6 seconds and is based on 18 seconds of previous data which is then forward predicted to account for a 4 second latency [5]. Each correction is also discretized by 0.125 meters and then used to extrapolate the simulated WAAS correction shown by the solid line. As can be seen, this process provides a correction for PRN 16 accurate to better than a meter. However, for PRN 18, the large accelerations and sudden sign changes combined with the message rate and latency result in a very poor correction message. Here the errors can be off by as much as 10 meters.

In the region between -5 and 10 seconds, the clock acceleration changes from greater than 200 mm/s^2 to less than -150 mm/s^2 and back greater than 170 mm/s^2. Remember that the Standard Positioning Service (SPS) specification for the magnitude of the clock acceleration is not to exceed 19 mm/s^2 [13]. Thus we have accelerations greater than 10 times the specification and rapid changes as well. This is more than the WAAS messaging system

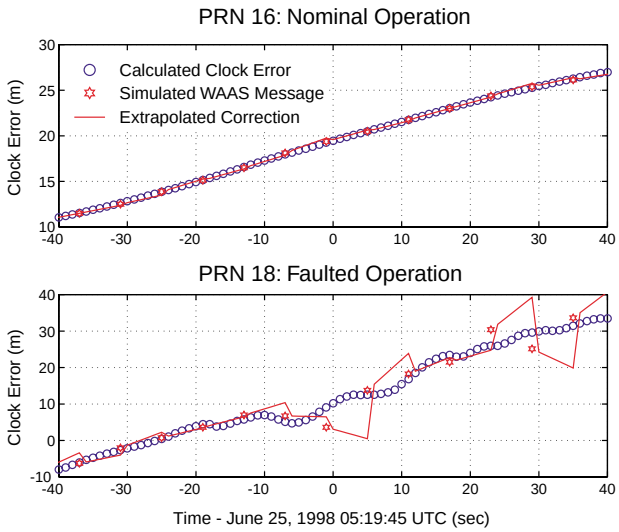

Figure 12. *Normal satellite behavior shown up top for PRN 16 and excessive accelerations shown below for PRN 18. Also shown are the simulated effects on the WAAS fast correction messaging. Not shown is the alarm response of the system.*

can keep up with. If we had implemented a more sophisticated acceleration trap, as will be in place for the operational WAAS and LAAS systems, we could have recognized and flagged this problem before data was transmitted to the user. As it was, our TMS recognized a problem and broadcast fast clock corrections more often than every six seconds to mitigate this problem. Thus, within the required six second time to alarm, when our system recognized that the user would suffer an unacceptably large error, it sent out an emergency correction message (not shown in Figure 12). Operationally there were no ill effects as the VPL always covered the actual error. However, we feel the system can be made to perform better if we installed the kind of error handling that will be present in the operational system. In this case, the seven anomalous points will either be recognized and removed beforehand, or the UDRE will be sufficiently increased to account for this unusual satellite behavior. Again it should be noted that this anomalous behavior persisted for hours yet only seven points departed from the nominal distribution and no points led to integrity violations.

CONCLUSIONS

This paper has addressed two key assumptions of the integrity equation: the position errors can be overbounded at the tails by a zero mean gaussian with specified variance, σ_V, and the errors do not combine in a worst case fashion. The data in Figures 2 and 3 demonstrate that the real errors are always bounded by the VPL. In addition, the data is more tightly distributed than the reference

gaussian curve. We have seen these effects at every reference station, over numerous data sets spanning many years.

The weighting matrix for the integrity equation is diagonal, meaning that any correlations between the satellites are ignored. Also the variances for the individual error components for a single satellite are combined assuming independence (4). We have asserted that this is a conservative assumption because in reality they are correlated in such a manner so as to reduce the overall error. This has been demonstrated via the simulated data. Figures 10 and 11 incorporated conservative assumptions about correlations in time and the magnitudes of the error components yet the resulting distribution is still larger than the actual one. We did not simulate correlations between the satellites. At a minimum such correlations do not appear to inflate the error and all of the evidence to date suggests that they in fact reduce it. This is the expected outcome from analysis as well, given that common errors affect only the clock, and an error in one part of a satellite's correction will tend to create an offsetting error in another part.

The anomalous points in the real data lead to a weakness in the Stanford generated UDREs and not in the integrity equation itself. Although the errors were subsequently corrected within the six second time to alarm, an algorithm exists which could have prevented their transmission altogether. It is important to note that despite having clock errors on PRN 18 much larger than what is protected by the UDRE, these errors combined with other satellites to create a small overall position error. Thus, even in a faulted mode, the errors combined in a beneficial manner.

The real data shows the VPLs to be nearly twice as large as they need to be. Unfortunately we are already using the lowest possible values for UDRE and have reduced the degradation parameters as well. However, the observed variance of the errors appears to be significantly lower than the MOPS allows the master station to predict. Unfortunately the broadcast UDREs are also coarsely discretized and for safety the master station must broadcast the next larger value. Thus the MOPS forces much of this conservatism. However, some of it is a reflection of conservatism on the part of our master station.

Ideally, as these error sources become better understood we will be able to reduce the UDREs and GIVEs and hence the VPLs for the users. This will increase availability at stations such as Cold Bay, our worst station, and increase the margin of availability and continuity at other stations already having high availability.

ACKNOWLEDGMENTS

This work was sponsored by the FAA GPS Product Team (AND-730). We would also like to thank our colleagues at the Wide Area Differential GPS lab at Stanford University for their many contributions to this work.

REFERENCES

[1] Loh, R., Wullschleger, V., Elrod, B., Lage, M., and Haas, F., "The U.S. Wide-Area Augmentation System (WAAS)," NAVIGATION, Fall 1995.

[2] Walter, T., Enge, P., and Hansen, A., "A Proposed Integrity Equation for WAAS MOPS," in proceedings of ION GPS-97, pp. 475-484, Kansas City, MO, September 1997.

[3] Walter, T., Enge, P., and Hansen, A., "Integrity Equations for WAAS MOPS," in Global Positioning System: Papers Published in NAVIGATION Volume VI, to be published 1999.

[4] RTCA Special Committee 159 Working Group 2, "Minimum Operational Performance Standards for Global Positioning System / Wide Area Augmentation System Airborne Equipment," RTCA Document Number DO-229A June 1998.

[5] Walter, T., "WAAS MOPS: Practical Examples," in proceedings of ION National Technical Meeting, San Diego, CA, Janurary 1999.

[6] Slattery, R., Peck, S., Anagnost, J., and Moon, M., "Guaranteeing Integrity for all Users of Active Data," in Global Positioning System: Papers Published in Navigation Volume VI to be published 1999.

[7] Enge, P., Walter, T., Pullen, S., Kee, C., Chao, Y.C., Tsai, Y.J., "Wide Area Augmentation of the Global Positioning System," in Proceedings of the IEEE, Vol., 84, No. 8, August 1996.

[8] Tsai, Y. J., "Wide Area Differential Operation of the Global Positioning System: Ephemeris and Clock Algorithms," Stanford University Ph. D. Dissertation, June 1999.

[9] Cobb, H. S., Lawrence, D., Christie, J., Walter, T., Chao, Y. C., Powell, J. D. and Parkinson, P., "Observed GPS Signal Continuity Interruptions," in Proceedings of ION GPS-95, Palm Springs, CA, September 1995.

[10] Hansen, A., Walter, T., Lawrence, D., and Enge, P., “GPS Satellite Clock Event of SV#27 and Its Impact on Augmented Navigation Systems,” in proceedings of ION GPS-98, Nashville, TN, September 1998.

[11] Wessel, P. and Smith, W. H. F., “New Version of the Generic Mapping Tools Released,” EOS Trans. Amer. Geophys. U., Vol. 76, pp 329, 1995.

[12] http://waas.stanford.edu

[13] Global Positioning System Standard Positioning Service Signal Specification, June 1995.