Weighted RAIM for Precision Approach

Todd Walter and Per Enge
Stanford University

Abstract

The use of differential GPS is becoming increasingly popular for real-time navigation systems. As these systems migrate to safety-of-life applications (e.g. precision approach and landing), their integrity becomes more important than their accuracy. One method for increasing both accuracy and integrity is the use of weighting in the navigation solution. This method uses *a priori* information to weight certain satellites (e.g. those at higher elevation) over other satellites. The accuracy increases because we better use the information available. The integrity increases because satellites that are more likely to introduce error contribute less to the solution.

A weighted position solution by itself does not provide sufficient integrity to support precision approach. However, this method can be combined with a weighted form of Receiver Autonomous Integrity Monitoring (RAIM) to increase the level of integrity. RAIM uses redundant measurements to check the consistency of an overdetermined solution. This check is crucial because only a user can detect certain error types (e.g. severe airframe multipath or local interference). A differential reference station can detect many types of errors. However, it is only at the user where all the information is combined. The use of RAIM (or some form of integrity at the user) must be combined together with integrity checking at the reference station to provide the overall safety of the system.

Weighted RAIM is investigated for application to Category I precision approach as supported by a Wide Area Augmentation System (WAAS). This paper details how to implement weighted RAIM and how to use geometry selection to guarantee a certain level of protection. Also, we provide information on the availability of these geometries. The results are based upon analysis, Monte Carlo simulation and actual data collected from Stanford University's wide-area differential GPS network.

1.0 Introduction

The use of differential GPS in a safety-of-life application, such as precision approach, requires that the system provide accurate navigational information with a high degree of integrity. Any potentially hazardous or misleading piece of information must be flagged as such, *before* it leads to a positioning error. The most robust way to ensure that the information is valid is to have a multi-layer series of checks with as much independence between the layers as possible. Ideally the validity of the corrections would be verified both on the ground and in the air. Receiver Autonomous Integrity Monitoring (RAIM) is one possible method to check the validity of the corrections in the air.

RAIM is a simple yet powerful technique to check the consistency of the navigation solution. This paper presents the weighted form of RAIM and how it can be applied to precision approach. Specifically, we will concentrate on Category I precision approach as will be supported by the Wide-Area Augmentation System (WAAS). WAAS is the FAA's implementation of Wide-Area Differential GPS (WADGPS) coupled with ground integrity monitoring and supplemental ranging signals.

Although the availability results presented in this paper are specific to WAAS, the application of weighted RAIM can readily accommodate many different systems. This method can be applied to local-area differential GPS and it can also be applied to the use of additional augmentations such as GLONASS satellites, barometric altimeters or precise user clocks, to name a few. The main reason to use the weighted form of RAIM is that some satellites (chiefly those nearer to the horizon) are more likely to suffer from greater errors than other satellites. Instead of treating all satellites equally, we can reduce the contribution from satellites likely to be "noisier" by reducing their weighting. Which satellites should be de-weighted can be determined through *a priori* information or through use of weights broadcast from the source of the differential corrections. The GPS satellites currently broadcast weighting information in the User Ranging Accuracy (URA) parameters.

This paper begins with a brief overview of WAAS and of the integrity that will be provided by its ground network. Next we will discuss using weighting information to improve the position solution and how to implement RAIM in its weighted form. Finally, we will present the results of our analysis and of our Monte Carlo simulations followed by our conclusions.

2.0 WAAS Ground Network Integrity

The Federal Aviation Administration (FAA) of the United States is rapidly developing a system to incorporate WADGPS [1] with other augmentations to create a system that is capable of providing navigation information with sufficient accuracy, integrity, availability and continuity to be used to support Category I precision landings. This Wide Area Augmentation System (WAAS) will initially be available as a supplemental means of navigation by the end of 1997.

The WAAS ground network contains three components: Wide Area Reference Stations (WRSs), Wide Area Master Stations (WMSs) and Ground Earth uplink Stations (GESs). The reference stations consist of dual frequency GPS receivers with antennae located at surveyed sites which also have clear visibility to the horizon. Also attached to each WRS receiver is an atomic clock, a meteorological station and a datalink to the master station. The WRS sends back to the WMS the raw GPS observables, the meteorological measurements and the broadcast ephemeris and GPS data.

The master station uses the information to separate the observed pseudorange errors into contributions from satellite clock error (including SA), satellite ephemeris error, ionospheric delay and tropospheric delay. The first three components are broadcast to the user. This information will be incorporated with a standard model of tropospheric delay to construct the scalar pseudorange correction valid at the user location.

The data is broadcast to the user via the GES and a geostationary satellite. The geostationary satellite will broadcast the WAAS data together with a ranging signal. This satellite signal will appear very similar to a GPS satellite signal except that it will from a geostationary source and it will send information at a 250 rather than 50 bit per second rate. All of the WADGPS corrections in addition to integrity information will be contained in these 250 bps messages.

This system has been described in more detail elsewhere [1-3] [11][12]. Previous attention has been paid principally to the accuracy achievable with WADGPS. This paper instead concentrates on the integrity aspects of WAAS and will focus primarily on user algorithms. The following subsections will describe how integrity is built into the ground network of WAAS to insure the quality of the broadcast messages. In addition, they will outline how an independent set of ground monitors can be used to augment integrity in local areas.

2.1 Reference Station Integrity

Each WRS can separate its measurements satellite by satellite because each WRS antenna is at a surveyed location and because we are using very stable atomic clocks. Each measurement of pseudorange, from the WRS to a satellite, is almost entirely independent of all of the other satellites. Therefore it is possible to detect errors in the pseudorange nearly instantly. However, the WRS does not necessarily have enough information to separate correctable errors (such as SA or small ephemeris offsets) from uncorrectable errors (large satellite clock or ephemeris errors outside of the field of the 250 bit correction message). Additionally, a single receiver may not be able to distinguish between errors originating at the satellite and errors originating at the receiver (channel bias or excessive multipath). Therefore the best location to make integrity decisions is at the WMS.

2.2 Master Station Integrity

Each WMS concentrates the information from multiple WRSs. Here it can perform a consistency check to determine the likelihood of an error. Each WRS has redundant receivers and each satellite is often observed by multiple WRSs. Therefore, in the event of an error, the WMS should be able to isolate the faulty component and take appropriate actions. If one receiver's measurements do not agree with several others (including one or two others at the same location) then the receiver can be declared failed and removed from the solution. If, on the other hand, all of the receiver measurements are consistent but there is a problem with the estimated satellite clock or ephemeris errors, then the WMS would broadcast a message not to use that satellite (or at least to give it a lower weighting). Examples of observable problems would include: errors beyond the maximum range permitted by the 250 bit messages, a lack of consistency between current and previous estimates or a disagreement between new measurements and the derived model.

As a further measure of integrity the WMS can compute an expected user accuracy model for its service region. This model would combine information about the confidences of the measurements, the locations of the measurements and reasonable uncertainties in the models used to form the corrections. Thus, the WMS could determine whether certain geographical locations might suffer from inaccurate corrections due to poor observing conditions.

The final check of the corrections can be performed either at the WMS or at the GES. This check compares the broadcast corrections (and warnings) generated at one WMS to those generated by another. These WMSs could share all of the same WRSs or use different subsets. This comparison would help to mitigate undetected errors in any one WMS. This check is different from the pseudorange comparison because it compares the final calculations. Errors introduced by the WMS itself in addition to measurement errors propagated through the system may be found in this manner.

To guard against data corruption, the WMS 250 bit message includes 24 bits of Cyclic Redundancy Checks (CRCs). The 24 CRC parity bits provide protection against both burst and random errors with a probability of undetected error being less than 6×10^{-8} [4]. This measure reduces the likelihood that valid WMS data would be wrongly interpreted by the user resulting in an erroneous correction.

2.3 Monitor Station Integrity

While the above integrity algorithms should protect against a wide variety of error modes, it is still conceivable that some error modes might pass undetected through the system. Most notably, local disturbances in the ionosphere or troposphere may not be detected by distant reference stations, and yet, through some extremely unlikely combination, create a significant positioning error.

One possibility to detect such effects is to co-locate WAAS capable receivers at the local Air Traffic Control (ATC) centers. Because these receivers would have their antennae positions at known pre-surveyed locations it would be possible to determine if the position solution obtained is of sufficient accuracy to support Category I precision approach. A pilot desiring to make such an approach must contact ATC, who would then check if their local WAAS monitor indicates sufficient system accuracy. Because the WAAS Monitor is close to the region of operation of the airplane a great many possible errors should be common to the monitor and the airplane.

This is a powerful cross-check of the WAAS because the monitors are independent of the WAAS ground network. It is feasible (although extremely unlikely) that some form of common-mode error could propagate through the WAAS ground network undetected, due to its interconnectedness. However, these independent ground stations should be able to detect any significant error of this kind, in addition to flagging errors caused by local ionospheric or tropospheric disturbances.

A very attractive feature is that the monitor station directly measures any error in the positioning domain (as opposed to the pseudorange or ionospheric delay domain). Therefore, if the error is not detected by the monitor station, it is unlikely to lead to a positioning error at the user (provided the same set of satellites is used both on the ground and in the air). Additionally this receiver need not be expensive. A single frequency receiver of the similar to those used in the airplanes would suffice. The main differences would be in the software and the surveying of the antenna.

The best method to protect against potential errors (both foreseen and unforeseen) is to have multiple, independent, redundant systems. The level of redundancy built into the WAAS ground network provides a high degree of integrity. The addition of a “shadow” set of independent monitor stations offers an even greater level of protection against a wider class of possible failure modes.

3.0 Weighted Position Solution and RAIM

For stand alone GPS, the dominant error is caused by Selective Availability (SA). Consequently, there is little advantage in weighting one satellite over another. Errors are separated into geometrical factors (DOPs) and User Ranging Accuracies (URAs). However, with differential GPS, some satellites may have predictably larger errors than others. As an example, it may be desirable to give higher elevation satellites more weighting when performing a position solution. Low elevation satellites suffer from greater multipath effects, increased tropospheric delay uncertainty and usually have a lower Signal-to-Noise Ratio (SNR).

There are several advantages to weighted position solutions: the position fix is more accurate; it is more robust because satellites that are more likely to have errors contribute less to the solution; and the discontinuities in the position fix caused by rising and setting satellites are greatly reduced. While many readers may be familiar with

DOPs and URAs, there is less familiarity with the weighted position solution. This section presents a brief derivation of the weighted form of the Receiver Autonomous Integrity Monitoring (RAIM). We begin by first showing the weighted least-squares navigation solution. Next we show how the consistency of the redundant measurements may be used to generate a test statistic. Finally, we present necessary conditions for this consistency check to gauge the accuracy of the solution.

It is important to have integrity checking at the user because this is the only place where all information used to form the position solution is present. There are many possible error modes that may only affect the user. These include: excessive multipath, receiver error, poor differential corrections resulting from data drop-out and localized ionospheric or tropospheric effects. While these error modes are extremely unlikely, they may not be detectable by either the WAAS ground network or the local monitor station. Therefore, some form of integrity checking must take place within the user's equipment. RAIM is easily implemented, requires no additional hardware and is capable of providing this final layer of integrity.

3.1 Weighted Position Solution

The basic linearized GPS measurement equation is

$$\boldsymbol{y} = \mathbf{G} \cdot \boldsymbol{x} + \varepsilon$$

where $\boldsymbol{x}$ is the four dimensional position vector (north, east, up and clock) about which the linearization has been made, $\boldsymbol{y}$ is an N dimensional vector containing the raw pseudorange measurements minus the expected ranging values based on the location of the satellites and the location of the user ($\boldsymbol{x}$), $\mathbf{G}$ is the observation matrix and ε is an N dimensional vector containing the errors in $\boldsymbol{y}$.

The weighted least squares solution for $\boldsymbol{x}$ can be found by

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

where the definition has been made for $\mathbf{K}$ (the weighted pseudo-inverse of $\mathbf{G}$) and where $\mathbf{W}$ is the inverse of the covariance matrix. For simplification we will assume that the error sources for each satellite are uncorrelated with the error sources for any other satellite. Therefore, all off-diagonal elements are set to zero. The diagonal elements are the inverses of the variances (σ^2s) corresponding to each satellite. While this assumption may not be strictly true, it should be a reasonably good approximation. The equations subsequently derived do not depend on this assumption. It only makes them easier to implement in practice.

Because the satellites are weighted unequally, we can no longer separate the expected positioning errors into a geometrical factor (DOP) and a user ranging accuracy (URA, or σ, common to all satellites). Instead these values are combined into expected positioning confidences. Instead of VDOP given by

$$VDOP \equiv \sqrt{\left[\left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{G}\right)^{-1}\right]_{3\,3}}$$

we now have σ_V given by

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

as a measure of the confidence of the vertical accuracy. In a similar manner the horizontal confidence HRMS can be given by

$$\mathrm{HRMS} \equiv \sqrt{\left[\left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{W} \cdot \mathbf{G}\right)^{-1}\right]_{1\,1} + \left[\left(\mathbf{G}^{\mathrm{T}} \cdot \mathbf{W} \cdot \mathbf{G}\right)^{-1}\right]_{2\,2}}$$

These measures give the 1-sigma expected accuracy in the vertical dimension and the 2-dimensional RMS expected accuracy in the horizontal dimensions respectively. The accuracies of these measures depend on the accuracies of the satellite covariances in the $\mathbf{W}$ matrix.

3.2 Weighted RAIM

So far we have only presented the weighted position solution. Now we wish to assess the accuracy of the least squares fit to the data. The quantity we are most interested in is the positioning error ($\hat{\boldsymbol{x}} - \boldsymbol{x}$). Unfortunately it is not possible to obtain a direct measurement of this quantity, unless we were to have access to an independent, more accurate positioning system. Instead, we can examine the overall consistency of the solution. Provided we have more than four measurements, the system is overdetermined and cannot be solved exactly. This is why a least squares solution is performed in the first place. Since all of the conditions realistically cannot be met exactly, there is a remaining error residual to the fit. By quantifying how closely we were able to make all the observations agree, we can get an estimate of the goodness of the fit. Then we make the assumption that if the fit was good, the error in position is most likely small. This is the foundation for RAIM.

We can get an estimate of the ranging errors from the least squares fit and the basic measurement equation

$$\hat{\varepsilon} = \boldsymbol{y} - \mathbf{G}\cdot\hat{\boldsymbol{x}} = (\mathbf{I} - \mathbf{G}\cdot\mathbf{K})\cdot\boldsymbol{y} \equiv (\mathbf{I} - \mathbf{P})\cdot\boldsymbol{y}$$

where the definition has been made

$$\mathbf{P} \equiv \mathbf{G}\cdot\mathbf{K} = \mathbf{G}\cdot\left(\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}\cdot\mathbf{G}\right)^{-1}\cdot\mathbf{G}^{\mathrm{T}}\cdot\mathbf{W}$$

From these error estimates we can define a scalar measure defined as the Weighted Sum of the Squared Errors

$$WSSE = \hat{\varepsilon}^{\mathrm{T}}\cdot\mathbf{W}\cdot\hat{\varepsilon} = \left[(\mathbf{I}-\mathbf{P})\cdot\boldsymbol{y}\right]^{\mathrm{T}}\cdot\mathbf{W}\cdot\left[(\mathbf{I}-\mathbf{P})\cdot\boldsymbol{y}\right]$$

which is equivalent to

$$WSSE = \boldsymbol{y}^{\mathrm{T}}\cdot\mathbf{W}\cdot(\mathbf{I}-\mathbf{P})\cdot\boldsymbol{y}$$

We use $\sqrt{WSSE}$ as our test statistic in order to judge the goodness of the least squares fit. This statistic is observable whereas the positioning error of the least squares solution ($\hat{\boldsymbol{x}} - \boldsymbol{x}$) is not. Therefore, for integrity purposes we want to use the statistic to flag bad position solutions. Typically, a certain threshold is selected. If the statistic exceeds that threshold the position fix is assumed to be unsafe. However if the statistic is below the threshold, then the position fix is assumed to be valid. Thus, the statistic-vertical error plane is broken up into four regions consisting of: normal operation points, missed detections, successful detections and false alarms (See Figure 1). Ideally, there would never be any missed detections or false alarms.

The threshold T, is chosen such that the probability of false alarm is commensurate with the continuity requirement for precision approach. Under normal conditions, if we assume ε_i is a normally distributed zero mean random variable with a standard deviation of σ_i for all N satellites in view, then the statistic is a chi-square distributed variable with N-4 degrees of freedom. Therefore the threshold T can be selected analytically. $T(N, P_{FA})$ will only be a function of the number of satellites (N) and the desired probability of false alarms (P_{FA}). By examining the distribution it is possible to find the value $T(N, P_{FA})$ such that, for normal conditions, the statistic only has a probability of P_{FA} of exceeding it. Given the probability of false alarms, the threshold is found by inverting the incomplete gamma function [5] [9] [10]

$$1 - P_{FA} = \frac{1}{\Gamma(a)}\int_0^{T^2} e^{-s}s^{a-1}ds$$

where a is the number of degrees of freedom divided by two, or in terms of the number of measurements N

$$a = \frac{N-4}{2}$$

The best way to find the values for $T(N, P_{FA})$ is through an iterative root finding process. Note that these values can be easily computed beforehand and stored for use later in a RAIM algorithm. For convenience several such values are listed in Table 1.

3.3 Protection Levels

Unfortunately, the N errors in the vector ε are mapped into two orthogonal spaces; one of dimension 4 corresponding to the position solution error and one of dimension N-4 corresponding to our statistic. Thus, in the most general case, the statistic cannot be used absolutely to indicate a bad or a good position solution. However, in the case of a single satellite failure, it is possible to restrict the satellite geometries such that a large bias that is mapped into a position error is also mapped into the statistic with certainty. Thus, for this failure mode we can guarantee that the position error will not grow too large without a corresponding growth in the statistic.

This restriction is not necessarily unreasonable because it is assumed that ground monitoring will pick up and isolate any faulty satellite within a relatively short period of time. Thus the likelihood of multiple satellite errors not detected by the ground monitoring network are

N \ P_{FA}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}	10^{-7}	10^{-8}	10^{-9}
5	2.576	3.291	3.891	4.417	4.892	5.327	5.731	6.109
6	3.035	3.717	4.292	4.798	5.257	5.678	6.070	6.438
7	3.368	4.033	4.594	5.089	5.538	5.950	6.335	6.694
8	3.644	4.297	4.849	5.336	5.777	6.184	6.563	6.920
9	3.884	4.529	5.074	5.555	5.991	6.392	6.767	7.120
10	4.100	4.739	5.278	5.754	6.185	6.583	6.954	7.304
11	4.298	4.932	5.466	5.938	6.366	6.760	7.128	7.475
12	4.482	5.111	5.612	6.110	6.535	6.926	7.292	7.636

Table 1. *Values of Threshold (T) for given probabilities of false alarm and number of satellites.*

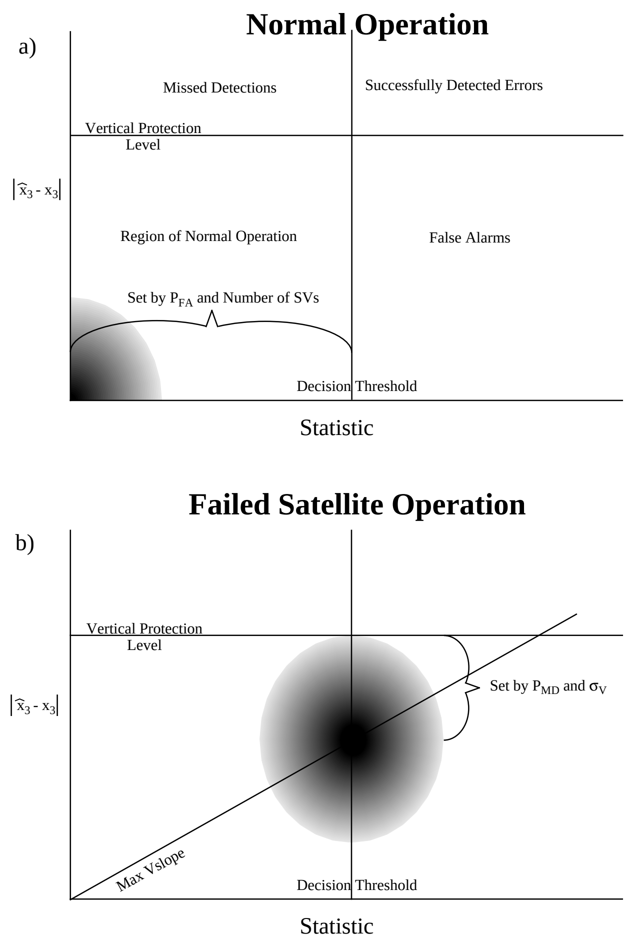

Figure 1. *The distribution of vertical errors and the RAIM statistic are shown here for both normal operation and in the case of a failed satellite.*

expected to be exceedingly small.

We can use the method developed by Brown [5] to guarantee integrity by only accepting geometries which provide adequate redundancy to determine if there is an error on any one channel of the receiver. This method trades availability for integrity. In the vertical dimension, this method requires that the vertical slope for each satellite (*i*), given by [6]

$$Vslope_i \equiv \frac{|K_{3\,i}|\sigma_i}{\sqrt{1-P_{i\,i}}}$$

be less than some maximum allowable slope.

If there is a failure of a single satellite, the expected distribution of operation points in the statistic-vertical error plane is still an ellipse with roughly the same contours as in the absence of failures. The difference is that now the ellipse is no longer centered near the origin. Instead, its center has moved out along the line with the corresponding Vslope for the failed satellite (see Figure 1b). How far it moves along the line depends on the magnitude of the bias. A valid integrity algorithm should alert the pilot to this failure before the vertical error exceeds the desired vertical integrity limit (VIL), thus keeping all points in this ellipse out of region of missed detections.

From Figure 1b we can see that the maximum allowable slope is a function of the desired probability of false alarms, the acceptable probability of missed detection, the vertical error we are trying to protect and σ_V [6]. Integrity is only declared available if each Vslope is less than

$$\frac{\mathrm{VIL} - k(P_{MD})\sigma_V}{T\left(N, P_{FA}\right)}$$

where $k(P_{MD})$ is the number of standard deviations corresponding to the specified P_{MD}.

Another point of view is to assume the WAAS ground network will provide a sufficient level of integrity. Instead we now wish to determine how much additional integrity RAIM can provide. The equation above can be rearranged to give the vertical protection level that the fault detection algorithm is capable of protecting (VPL_{FD} or HPL_{FD} in the horizontal plane). These values are given by

$$\mathrm{VPL}_{\mathrm{FD}} \equiv \max[Vslope]T\left(N, P_{FA}\right) + k(P_{MD})\sigma_V$$

$$HPL_{FD} \equiv \max[Hslope]T\left(N, P_{FA}\right) + k\left(P_{MD}\right)HRMS$$

While this analysis determines the capability of RAIM in the presence of a single satellite failure, this form of RAIM cannot be made robust against any type of multiple satellite error. One can always conceive of a satellite pair failure that would yield zero contribution to the statistic and yet result in a large positioning error. However, RAIM is one layer of a multi-layer integrity structure. It cannot guarantee catching all errors, but no system can. Despite these limitations, RAIM would very likely detect a random multi-satellite failure. In addition, the probability of this failure mode occurring and escaping the detection of both the WAAS ground network and the local monitor is extremely remote.

Results

In order to investigate the availability of RAIM, we ran simulations which made assumptions about the expected satellite variances, protected vertical error and probabilities of missed detection and false alarms. The probabilities that we used were: a probability of false alarms no greater than 10^{-5} per approach (5 independent samples per approach) and a probability of missed detections less than 10^{-3} per independent sample. The overall probability of hazardous or misleading navigation information is required to be below 10^{-7} per approach. This last requirement adds a restriction on σ_V. Even if there were no errors, normal operating conditions might not sufficiently protect the desired vertical error. For a true Gaussian distribution, 10^{-7} corresponds to a 5.33 σ error (see Table 1). Realistic distributions usually have broader tails than a true Gaussian distribution. However the WAAS ground network is designed to prevent long tails. The moderately conservative approach taken here is to require that the 5.5 σ_V error be within the desired protected vertical error. As the actual distribution of errors becomes better characterized over time this value will be subject to change. For a vertical integrity limit of 19 meters, this requirement is equivalent to restricting σ_V to ~ 3.5 meters or below. If that condition is not met, the landing may not safely proceed, regardless of the integrity conditions of the WAAS ground network and/or the local WAAS monitor.

The variances that are assumed for each satellite depend strongly upon elevation angle and take the following form [6]

$$\sigma_i^2 = \sigma_{UDRE\ i}^2 + \mathrm{F}^2(El_i)\,\sigma_{UIVE\ i}^2 + \sigma_{SNR\ i}^2 + \frac{\sigma_{m45}^2}{\tan^2 El_i} + \frac{\sigma_{trv}^2}{\sin^2 El_i}$$

The following definitions have been made: σ_i^2 is the total variance of the i^th satellite, $\sigma_{UDRE\ i}^2$ is the variance of the supplied tropo-free iono-free pseudorange correction, $\sigma_{UIVE\ i}^2$ is the variance of the vertical ionosphere correction, $\mathrm{F}(El_i)$ is the obliquity factor converting vertical measurements into slant, $\sigma_{SNR\ i}^2$ is the receiver noise variance and can be related to signal-to-noise ratio (C/N_0), σ_{m45}^2 is the variance of the multipath contribution at 45 degrees and σ_{trv}^2 is the variance of the vertical tropospheric delay estimate. The values used in our simulation are based on values we observe regularly with the Stanford WAAS network[1] and are given by:

$$\sigma_{UDRE} = 0.5\text{m}$$
$$\sigma_{UIVE} = 0.5\text{m}$$
$$\sigma_{SNR} = 0.22\text{m}$$
$$\sigma_{m45} = 0.22\text{m}$$
$$\sigma_{trv} = 0.15\text{m}$$

Figure 2. shows the weighting curve that results from this model and these values.

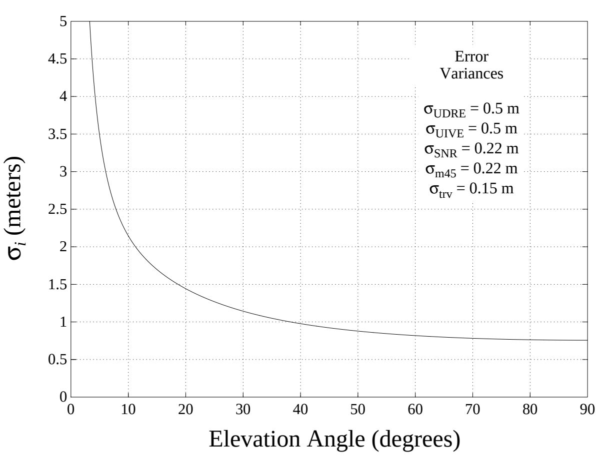

Figure 2. *Here pseudorange uncertainty is plotted as a function of elevation angle.*

In order to calculate availability we combined these error variances with realistic failure models for the GPS satellites and for geostationary satellites [7] [8] to calculate the percentage of the time the accuracy would be sufficient to support Category I or near Category I landings and the percentage of time that the geometry (and variances) would support RAIM.

The Monte Carlo results presented here are based on 10^7 simulated geometries. We computed the results for three different satellite constellations. For reference we calculated a GPS only case, although in reality geostationary satellites are required in order to transmit the differential corrections. The other two cases augment the GPS constellation with INMARSAT and with INMARSAT plus four additional geostationary satellites selected to provide good coverage over the continental United States (CONUS). In all cases the user mask angle was set to 5°. Figure 2 demonstrates that the results should not depend too strongly on user mask angle as satellites are heavily de-weighted below ~ 10°.

All cases were computed over the CONUS region. For the cases using geostationary satellites, we assumed that these satellites provided additional ranging signals with accuracies equivalent to the differential GPS values. The three visible INMARSAT satellites in the CONUS region are; the Atlantic Ocean Region East and West (18.5° W and 55° W) and the Pacific Ocean Region (180°). The four additional geostationary satellites were located at 85° W, 110° W, 140° W and 165° W.

Six different availabilities were investigated and summarized in Table 2. At least four satellites are required in order to obtain navigation solutions. In addition, a minimum of five satellites must be in view to be able to implement RAIM. These raw availabilities are listed in the first two columns of Table 2. The next two

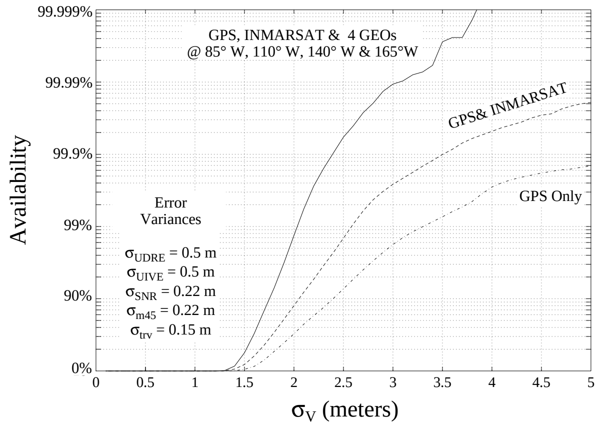

Figure 3. *This graph shows the availability of σ_V for three different satellite constellations.*

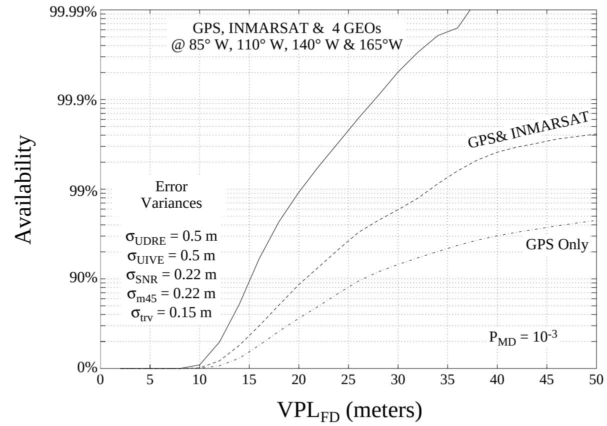

Figure 4. *This plot shows the availability of specific vertical protection levels provided by the fault detection algorithm.*

columns show the availabilities of σ_V and of RAIM that would meet the Category I requirements (for a 200 foot decision height the vertical integrity limit is 19 meters). While the availability of the system is good, particularly with augmentations, the availability of RAIM is low. The last two columns show the availability of near Category I approaches (using a decision height of 400 feet, the vertical integrity limit is ~26 meters which corresponds to an upper bound on σ_V of ~ 4.7 meters). The availability of RAIM markedly improves.

Figures 3 and 4 show the overall availability of σ_V and of the vertical protection limit provided by RAIM respectively. These results clearly show the benefit of having additional ranging signals. GPS with just INMARSAT can very nearly meet the desired availability for Category I landings of 99.9% if RAIM were not a requirement. Unfortunately none of the augmentations tried for this paper can provide RAIM to that same degree of availability.

If the accuracy of the WAAS were to improve there would be a corresponding increase in availability. Smaller values of σ_i lead to smaller values for σ_V and VPL_{FD}. This in turn will increase both the availabilities of the system and of RAIM. Figure 4 clearly shows that only the case with GPS, INMARSAT and the 4 Geos has any chance of making RAIM available 99.9% of the time. Unfortunately it would take a 40% reduction in the values of σ_i in this fully augmented case to make RAIM available to the desired level. This reduction may be overly optimistic, but a modest improvement coupled with alternate augmentations or with higher decision thresholds could sufficiently increase the availability of RAIM.

It should be noted that RAIM should be applied even if the vertical protection level exceeds the desired vertical integrity limit. RAIM would still catch large errors and errors on satellites whose Vslopes were below the maximum allowable slope. The only drawback is that there is no guarantee that RAIM would catch the error with the specified probability of missed detections. As long as RAIM is one part of an overall integrity scheme then it should be applied regardless of the vertical protection level.

Figure 5 shows the distribution of the RAIM statistic for some data collected at Stanford's WAAS network [1][11]. This data represents over 14,000 points in which seven satellites were in view of a statically surveyed passive user. For reference the expected chi-square distribution is also shown. Although the sample size is not large enough to draw any definitive conclusions, it does appear that our estimates of the satellite covariances are too conservative. A reduction of all the sigma values by roughly 30% would bring the actual data more in line with the theoretical curve. We need to perform a more careful evaluation of how we

	$N \geq 4$	$N \geq 5$	$\sigma_V < 3.45$	$VPL_{FD} < 19$m	$\sigma_V < 4.7$	$VPL_{FD} < 26$m
GPS Only	99.996%	99.95%	99.2%	67.6%	99.8%	89.0%
GPS & INMARSAT	> 99.999%	> 99.999%	99.89%	85.0%	99.97%	96.9%
GPS, INMRST & Geos	> 99.999%	> 99.999%	99.996%	98.4%	> 99.999%	99.8%

Table 2. *Availabilities for three satellite constellations.*

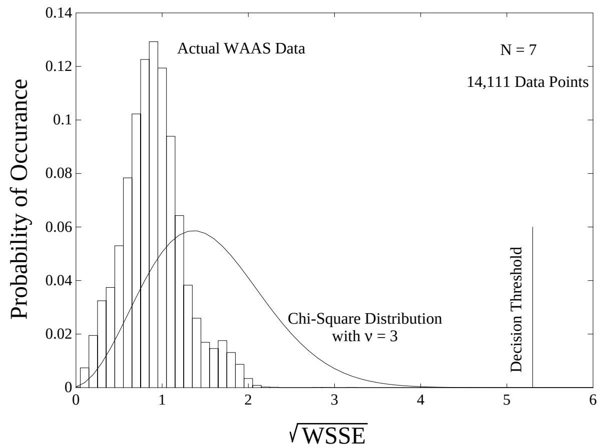

Figure 5. *The distribution of the RAIM statistic from actual WAAS data is compared to theoretical expectations.*

generate the covariances, and collect much more data before we can safely begin to adjust these values.

Conclusions

We have presented the equations for the weighted form of GPS positioning and for RAIM. These equations are simple extensions of the familiar unweighted methods. We argue that correct weighting will improve both the accuracy and the integrity of the navigation solution. As the use of differential GPS (as well as additional augmentations) become more common, weighted position fixes should replace non-weighted positioning as the standard.

We feel that a multiple layer approach, with three independent checks (WAAS ground network, WAAS ground monitors and RAIM) is capable of providing sufficient integrity to support precision approach. RAIM is an important element in providing complete system integrity. Only at the airplane is all of the information present and only at the airplane can certain errors be detected. It is therefore extremely important to verify the integrity of the corrections at the airplane. RAIM can provide this integrity check. RAIM can be implemented without any hardware modifications to the airplane. It can be easily performed using information already available in the receiver. It should be implemented whenever redundant measurements are available.

It seems quite likely that WAAS will be capable of supporting Category I (or near Cat I) precision approaches with availabilities of 99.9%. RAIM should be applied in order to catch errors larger than the VPL_{FD} (RAIM still has a lesser chance than P_{MD} to catch errors below this level). The probability of encountering errors which can escape all ground monitoring and not lead to a large RAIM statistic is likely to be well below 10^{-7}.

In the early stages of WAAS it may be necessary to rely primarily on the WAAS ground network and the monitor stations to provide integrity when RAIM cannot meet the desired protection level. However this should be viewed as a temporary solution. As the accuracies improve and as additional measurements become available, so will the protection level of RAIM improve.

Our results show that the WAAS will provide sufficient accuracy to support near Category I landing requirements. Our experience with our own test-bed supports these results. WAAS will become an invaluable navigation provider for Category I precision approaches.

Acknowledgments

The Authors would like to thank Boris Pervan and Sam Pullen for numerous discussions about integrity. We also gratefully acknowledge the support and assistance of AND-510 (FAA satellite program office).

References

[1] *Flight Trials of the Wide Area Augmentation System (WAAS)*. T. Walter, C. Kee, Y. C. Chao, Y. J. Tsai, U. Peled, J. Ceva, A. K. Barrows, E. Abbott, D. Powell, P. Enge and B. Parkinson, Proceedings of ION-GPS 94, Salt Lake City, Sept 22-24, 1994, pp. 1245-54.

[2] *Evaluation of Orbit and Clock Models for Real Time WAAS*. Y. J. Tsai, Y. C. Chao, T. Walter, C. Kee, D. Powell, P. Enge and B. Parkinson. Proceedings of the National Technical Meeting, Anaheim, Jan 18-20 1995.

[3] *Ionospheric Delay Model Improvement for Stanford WAAS Network*. Y. C. Chao, Y. J. Tsai, T. Walter, C. Kee, D. Powell, P. Enge and B. Parkinson. Proceedings of the National Technical Meeting, Anaheim, Jan 18-20 1995.

[4] *Wide Area Augmentation System Signal Specification.* A. J. Van Dierendonck, RTCA Publication, January 1994.

[5] *A Baseline GPS RAIM Scheme and a Note on the Equivalence of Three RAIM Methods*. R. G. Brown, **Navigation**, Vol. 39, No. 3, Fall 1992, pp. 301-316.

[6] *Baseline Weighted Navigation Solution, Navigation Sensor Error and Total System Error Algorithms for*

Precision Approach. T. Walter, P. Enge and A. J. Van Dierendonck, Appendix E to Draft RTCA WAAS MOPS August 1995.

[7] *GPS Availability, Part II: Evaluation of State Probabilities for 21 and 24 Satellite Constellations.* J. M. Durand and A. Caseeau, **Navigation**, Vol. 37, No. 3, Fall 1990, pp. 285-96.

[8] *Availability Characteristics of GPS and Augmentation Alternatives*. W. S. Phlong and B. D. Elrod, **Navigation**, Vol. 40, No. 4, Winter 1993-4, pp. 409-28.

[9] *Autonomous GPS Integrity Monitoring Using the Pseudorange Residual*. Parkinson, B. W. and Axelrad, P., **Navigation**, Vol. 35, No. 2, 1988, pp 255-74.

[10] **Handbook of Mathematical Functions**. Abramowitz, M. and Stegun, I. A., New York, Dover Pub. Inc., 1972.

[11] *An Algorithm for Inter-frequency Bias Calibration and Application to WAAS Ionosphere Modeling*. Y. C. Chao, Y. J. Tsai, T. Walter, C. Kee, P. Enge and B. Parkinson, Proceedings of ION GPS 95, PAlm Springs, Sept. 12-15, 1995.

[12] *Validation of the RTCA Message Format for WAAS.* Y. J. Tsai, P. Enge, Y. C. Chao, T. Walter, C. Kee, J. Evans, A. K. Barrows, D. Powell and B. Parkinson, Proceedings of ION GPS 95, PAlm Springs, Sept. 12-15, 1995.