Practical Signal Quality Monitoring for Augmentation Systems

R. Eric Phelts, Todd Walter

Department of Aeronautics and Astronautics
Stanford University, Stanford, California
pheltsre@stanford.edu, twalter@stanford.edu

BIOGRAPHY

R. Eric Phelts is a Research Associate in the Department of Aeronautics and Astronautics at Stanford University. He received his B.S. in Mechanical Engineering from Georgia Institute of Technology in 1995, and Ph.D. from Stanford University in 2001. His research involves multipath mitigation techniques and satellite signal anomalies.

Todd Walter received his B. S. in physics from Rensselaer Polytechnic Institute and his Ph.D. in 1993 from Stanford University. He is currently a Senior Research Engineer at Stanford University where his research focuses on algorithms that provide provable integrity for WAAS.

ABSTRACT

Signal Quality Monitoring (SQM) algorithms propose to detect anomalous GPS signal distortions primarily through the use of multicorrelator SQM receivers. However, currently no augmentation systems (e.g., WAAS or LAAS) have true, real-time SQM capabilities. Still, these systems introduce complexities such as multiple-receiver processing, stringent detection requirements, and potential hardware failures.

Among the more subtle complexities are so-called inter-receiver biases (IRBs). IRBs are a product of receiver-induced correlation peak distortions. They result primarily from analog component mismatch in each receiver. In addition, temperature variations may cause these biases to drift slightly over time. Careful monitoring of these biases, however, will permit the system to screen for faulty receivers, to detect small, slowly-varying distortions over time, and also to observe the different levels of correlation distortion corresponding to each SV.

This paper explores a practical SQM processing algorithm in the context of the Wide Area Augmentation System (WAAS). Further, it addressees several issues related to IRBs and describes techniques for mitigating their effects to improve detection sensitivity. Finally, the paper provides sample results using data akin to that which will be observed by the WAAS Offline Monitoring stations operated by the FAA Technical Center.

INTRODUCTION

The Wide Area Augmentation System (WAAS) was commissioned for Initial Operational Capability (IOC) in July of 2003. This first build offers full protection against the "Most Likely Subset"(MLS) signal deformation threat model, but it must equip its reference stations with SQM receivers and protect against all signal deformation threats to achieve compliance for Final Operational Capability (FOC) [6].

WAAS currently mitigates the MLS threat model using its Code-Carrier Coherence (CCC) Monitor. [6] This monitor measures the rate of divergence between the code and carrier measurements on each satellite signal. The occurrence of hazardous MLS signal deformations causes this rate to exceed the detection threshold and cause the monitor to flag the SV as "Not Monitored."

To guard against the complete set of signal deformations defined by ICAO, WAAS currently requires offline monitoring. The Technical Center of the Federal Aviation Administration (FAA Tech

Center) leads this effort for WAAS and will employ true, multicorrelator SQM receivers to mitigate the full ICAO threat. Still. Although the core approaches for SQM have been known for some time, few details have been published detailing all the steps required to accomplish it. [1],[2],[3],[6]

A practical SQM algorithm includes all the steps required to implement SQM from beginning to end. These include such things as normalization, smoothing, inter-receiver bias removal, receiver averaging and threshold computation. This paper provides a step-by-step analysis of implementing SQM and attempts to provide the details required to perform it.

This algorithm presumes an installed network of N stationary SQM-capable receivers at well-placed locations. In addition it requires at a minimum that the following SQM measurements be nominally available (i.e., save for infrequent, random data-dropouts) from all receivers at a 1-Hz rate at a central processing facility or master station:

- All relevant, in-phase correlator measurements (e.g., d_{sqm} = -0.1, -0.075, -0.05, -0.025, +0.025, +0.05, +0.075, +0.1)
- Satellite elevation angles
- Broadcast UDREs (where appropriate)
- Receiver lock status flags/carrier phase lock status

The following procedure will refer to an example case conducted using Stanford University's Integrity Monitoring Testbed (IMT) and utilizes three (relatively closely-spaced) antennas (as shown in Figure 1) each connected to a separate SQM receiver. [5] (Note that this LAAS-prototype setup is analogous to a single WRS simultaneously processing all three threads of data.) Each of the receivers had a 16MHz bandwidth and could provide eight in-phase correlator measurements at the following (approximate) spacings/offsets: d_{sqm} = -0.1, -0.075, -0.05, -0.025, +0.025, +0.05, +0.075, +0.1; due to a firmware limitation the receivers tracked the satellite signals using the "ultranarrow" correlator pair (at ±0.025 chips). Figure 2 shows satellites tracked by each of the three receivers.

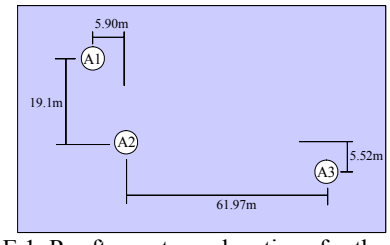

FIGURE 1. Rooftop antenna locations for the SU offline monitoring example dataset; Antenna Heights: 9.434m (A1); 9.459m (A2); 14.1827m (A3)

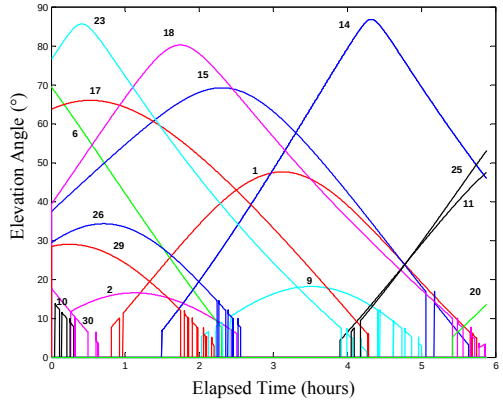

FIGURE 2. Elevation angles for all SVs tracked by the three receivers.

STEP 1: CORRELATION PEAK NORMALIZATION

Conventional normalization involves dividing each correlator measurement from a given receiver channel and by the largest accompanying correlator measurement in that receiver channel. Usually this is the Prompt correlator, P, located at an offset of zero chips. However, if a virtual prompt is used, it may be located at a nonzero offset.

A normalized measured correlation function at any offset, x, can is given by Equation 1. Note that for the example given in this paper, a virtual prompt, $P_{virtual}$ is used by averaging the in-phase (I) measurements at offsets of ±0.025 chips. This implies a correlator spacing, d, of 0.05 chips. (Figures 3 and 4 show the pre-normalized and post-normalized measurements for two correlator outputs.)

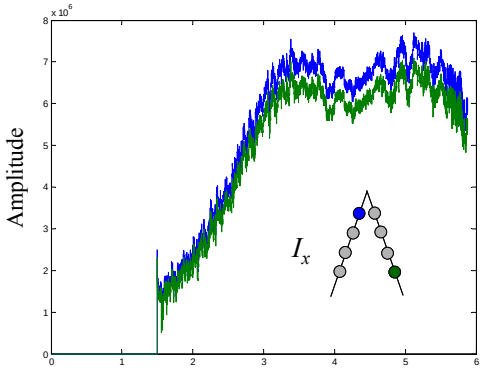

FIGURE 3. Raw (in-phase) correlation amplitude measurements, I_x, for two correlator outputs (at -0.025 and +0.1023 chips) for a single receiver tracking SV14.

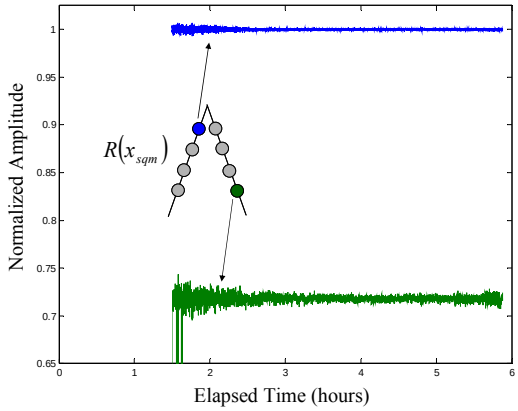

FIGURE 4. Conventionally-normalized correlator (i.e., normalized ratio) measurements for SV14.

$$R(x) = \frac{I_x}{I_0} = \frac{I_x}{P}$$

or (1)

$$R(x) = \frac{I_x}{\left(I_{-d/2} + I_{-d/2}\right)/2} = \frac{I_x}{P_{virtual}}$$

Type-normalization refers to converting each (conventionally normalized) correlation peak type to have dimensions consistent with a "nominal" peak. A normalized, "nominal" peak—conventionally normalized by a true prompt, P—is a triangle with vertices at [-1/1023 1 +1/1023]. (Note that all correlation peaks, regardless of normalization have a base width of two code chips.)

An ideal, normalized and unfiltered correlation function, $\mathsf{R}(x)$, is a triangle consisting of three types: small ("skinny"), medium ("nominal"), and large ("broad"). In general, each measurement may be type-normalized to a single type, ς—where ς equals 0, 1, or 2 for small (e.g, PRN 8 and 22), medium (e.g, PRN 1 and 3), and large (e.g, PRN 7 and 15), respectively)—according to the following:

$${}_{\varsigma\to 1}R(x_{sqm}) = {}_{\varsigma}R(x_{sqm}) \cdot \frac{{}_{1}\mathsf{R}(x_{sqm})}{{}_{\varsigma}\mathsf{R}(x_{sqm})}$$

where

$${}_{1}\mathsf{R}(x_{sqm}) = c_{\varsigma} + \frac{(1+c_{\varsigma})}{1-|x_{prompt}|}\left(1-|x_{sqm}|\right) \quad (2)$$

and

$c_0 = -65/1023$

$c_1 = -1/1023$

$c_2 = 64/1023$

Note that if a "virtual" Prompt ($P_{virtual}$) measurement is used, measurements at $P_{virtual}$ (or P) will equal unity. Accordingly, x_{sqm} in Eq. (2) refers to all receiver correlator pair spacings other than (i.e., wider than) those used in the conventional peak normalization process.

STEP 2: TIME-SMOOTHING OF METRICS

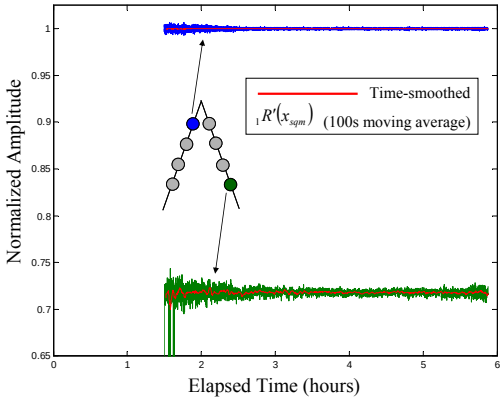

FIGURE 5. Time-averaged single-receiver (red) , normalized correlator (i.e., normalized ratio) measurements for SV14.

Short-term time averaging of the normalized measurements should be performed to reduce the variations due to thermal noise and multipath. However, the time constant (or filter length) should remain small to minimize the filter delay. Equation 3 describes this process. Traditionally, a rectangular FIR window (i.e., a moving average) filter is used with the filter length, F+1, equal to 100 seconds. Optionally, a shaped filter (or weighted average) may be applied based on WAAS code noise and multipath (CNMP) measurements. (See Figure 5.)

$${}_{1}R'(x_{sqm}) = {}_{1}R(x_{sqm}) * h_S(t)$$

where (3)

$$h_S(t) \Rightarrow h_S[n] = \begin{cases} 1, & 0 \le n \le F, \\ 0, & \text{otherwise} \end{cases}$$

STEP 3: COMPUTE *M* DETECTION METRICS FOR ALL RECEIVERS

The detection metrics, ${}_{m}D_i^j(t)$, are algebraic combinations of the correlator measurements for each receiver. There may be as many as M such tests in a given SQM implementation. The primary detection metrics are receiver configuration-specific (i.e., discriminator type, pre-correlation bandwidth and all

correlator spacings of interest). In addition, they are dependent on the nominal noise (and multipath) present in the normalized SQM correlator measurements.

One such metric—the "alpha" metric—is designed from a linear combination of the correlator based on the signal deformation (ICAO) threat model. For a 16MHz receiver having in-phase correlator measurements at offsets of approximately ±0.025Tc (tracking pair), ±0.05Tc, ±0.075Tc, and ±0.1Tc, the SQM algorithm minimally calls for the following (generalized) detection metric:

$$ {}_{m}D_i^j(t) \Rightarrow {}_{1}D_i^j(t) = \square^T \square \quad \text{(for } m = M = 1\text{)} \tag{4} $$

where

$$ \square = \begin{bmatrix} \alpha_{-0.1} \\ \alpha_{-0.075} \\ \alpha_{-0.05} \\ \alpha_{-0.025} \\ \alpha_{+0.025} \\ \alpha_{+0.05} \\ \alpha_{+0.075} \\ \alpha_{+0.1} \end{bmatrix} \text{ and } \square = \begin{bmatrix} {}_{1}R'(-0.1) \\ {}_{1}R'(-0.075) \\ {}_{1}R'(-0.05) \\ {}_{1}R'(-0.025) \\ {}_{1}R'(+0.025) \\ {}_{1}R'(+0.05) \\ {}_{1}R'(+0.075) \\ {}_{1}R'(+0.1) \end{bmatrix} $$

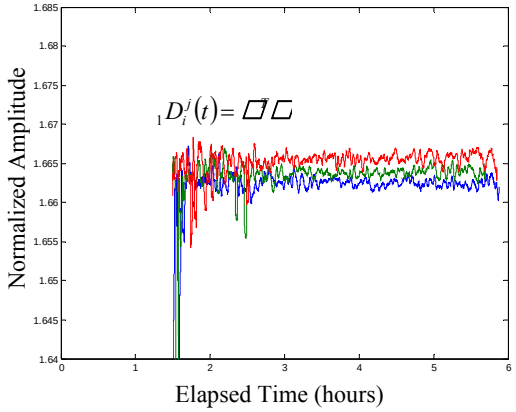

FIGURE 6. Three detection metrics ("alpha" metrics) for each of 3 receivers tracking SV14. Note the offsets between each metric occur primarily due to analog component mismatch and are referred to as inter-receiver biases (IRBs). (See Inter-receiver Bias Section.)

Figure 6 plots the resulting metric (determined as in [4]) for the three SQM receivers. Note: No specific transient (i.e., rapid) detection metrics have been developed at this time for this receiver configuration.

STEP 4: COMPUTE INTER-RECEIVER BIASES (IRBs)

The instantaneous inter-receiver ratio bias (or detection metric bias) is measured relative to a model in analysis only. In practice it can only be measured relative to another receiver. Offsets between receivers may be caused by a combination of the following:

- receiver filter tolerances/precision
- temperature variations
- thermal noise
- multipath (for receivers using non-collocated antennas)
- antenna filtering affects/variations (for receivers using non-collocated antennas)

Assertions

- The IRB estimate approximately equals the true IRB when t exceeds four times the long-term smoothing time constant, L_b; (i.e., $\hat{b}_i^j(t > 4 \cdot L_b) \cong b_i^j(t)$)
- The median ratio is not faulted and represents a good correlation measurement.
- The instantaneous bias is upper-bounded.
- The IRBs for a given receiver are the same across all received SV signals and may therefore be used for cross-comparisons and/or noise mitigation purposes. This is true provided under the following conditions:
  - thermal noise and multipath are averaged down to negligible levels
  - satellite correlation types (S,M,L) are taken into account (e.g., through type-normalization)
- Provided a bias is screened, computed, and measured for a given receiver it will not change unless a corresponding change occurs to one of the following:
  - the receiver (e.g., a malfunction, configuration change, or hardware replacement),
  - the receiver operating environment (e.g., ambient temperature),
  - the multipath environment, and/or
  - the satellite signal itself

The IRBs may be computed according to the following procedure:

1) Smooth each detection metric (for each receiver and SV) using a first-order filter with a *long* time constant.

$$ \hat{b}_i^j(t) = \frac{1}{L_b} b_i^j(t) + \frac{L_b - 1}{L_b} b_i^j(t-1) \tag{5} $$

$$b_i^j(t) = \left\{ \begin{array}{l} {}_m D_i^j(t) - \underset{i}{\text{median}}\left[{}_m D_i^j(t) \right] \\ \text{or} \\ {}_m D_i^j(t) - \dfrac{1}{N_v} \sum_{i=1}^{N_v} {}_m D_i^j(t) \end{array} \right\}$$

$$b_i^j(0) = {}_m \mathrm{D}^j = \text{constant} \qquad (6)$$

$$L_b = 1000$$

where

${}_m D_i^j(t)$ is the m^{th} instantaneous detection metric result of any correlator measurement to the prompt correlator on the peak corresponding to satellite, j and measured by receiver i at time t.

${}_m \mathrm{D}^j$ is the m^{th} analytical (modeled) detection metric result of any correlator measurement to the prompt correlator on the peak corresponding to satellite j

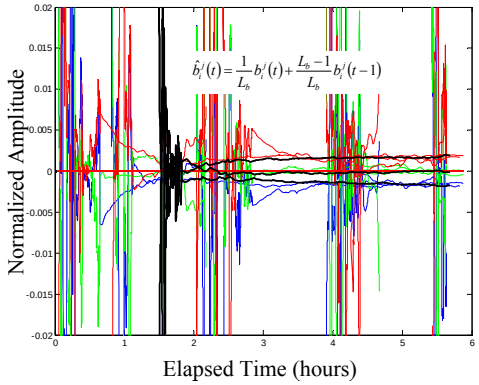

FIGURE 7. IRBs—as described by Equation (5)—for each satellite color-coded (red, green, and blue) according to receiver number. (The IRB traces for SV14 are in black for reference.) Note that each group defines a (receiver-specific) bound on the nominal signal deformations for satellites. However, more data (i.e., more satellite passes) are needed to smooth these traces so the bounds can be estimated.

$b_i^j(t)$ is the instantaneous ratio bias as measured by receiver i corresponding to the correlation peak of satellite, j. (NOTE: For robustness to receiver variability, etc., the median should be computed for large number of receivers, N (e.g., N>10). In WAAS, this restriction implies that only certain SVs may be used for this computation, since many may not be in view of a sufficient number or WREs.)

$\hat{b}_i^j(t)$ is the estimated (i.e., filtered using a first-order filter of time constant L_b) ratio bias as measured by receiver i corresponding to the correlation peak of satellite j.

L_b is the smoothing time parameter for first order filter (>1000)

Figure 5 plots the IRBs—for each SV—for all three SQM receivers of interest. However, each receiver may be viewed as possessing a single IRB value, $\hat{b}_i(t)$ which may be estimated from the correlation peak measurements made from each SV it tracks. (See Figure 8.) Equation (5) provides the basic filter equation for a single tracked SV. Measurements from additional SVs may be weighted and combined into a single estimate according to the following equation:

$$\hat{b}_i(t) = \left[\frac{1}{L_b} \sum_{j=1}^{J} \left[w_i^j b_i^j(t) \right] + \frac{L_b - 1}{L_b} \hat{b}_i(t-1) \right] \qquad (7)$$

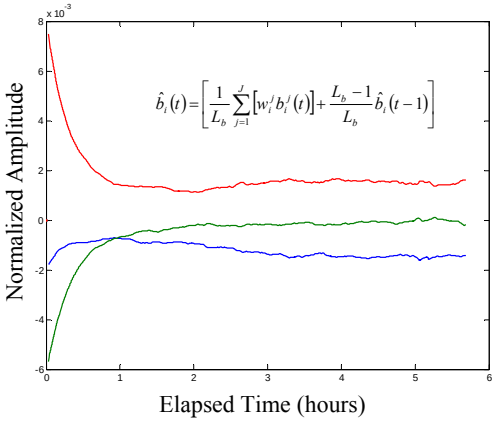

FIGURE 8. "Alpha" metric IRBs as described by Equation (7). Each satellite tracked by a given receiver was used in a weighted solution to form a single trace. (See Inter-receiver Bias Section.)

where

$$w_i^j = \frac{\dfrac{1}{\left[\sigma_i^j(\theta)\right]^2}}{\displaystyle\sum_{j=1}^{J} \frac{1}{\left[\sigma_i^j(\theta)\right]^2}} \qquad (8)$$

The standard deviation, σ_i^j, corresponds to the a priori standard deviation of the detection metric of interest

for SV j and receiver i at elevation angle θ (or WAAS UDRE). Figure 9 plots the weight factors, w_i^j.

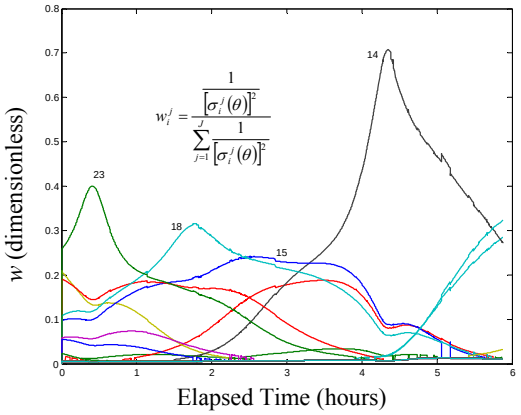

FIGURE 9. IRB weighting factors, w_i^j (Eq. 6) for all SVs tracked by the receivers and used to produce Figure 8.

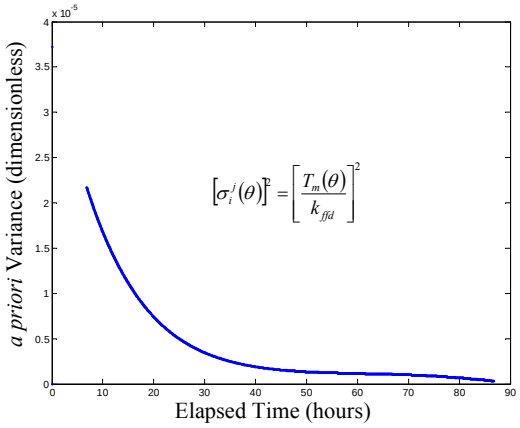

FIGURE 10. A priori variances $[\sigma_i^j(\theta)]^2$ for SV14 used to generate weighting factors, w, described above. For a given elevation angle, the variance profiles are the same for all SVs. (Note, that the elevation angle (θ) is also a function of time for GPS SVs.)

Note that $\sigma_i^j(\theta)$ results from the data used to derive the detection thresholds and is equivalent to $\sigma_i^j(UDRE)$ for WAAS. (Refer to Step 5a.) The equation is given by

$$[\sigma_i^j(\theta)]^2 = \left[\frac{T_m(\theta)}{k_{ffd}}\right]^2$$

or (9)

$$[\sigma_i^j(UDRE)]^2 = \left[\frac{T_m(UDRE)}{k_{ffd}}\right]^2$$

where T_m is the a priori detection threshold for metric, m; it may be a function of either SV elevation angle, θ (as in Figure 10.) or UDRE.

2) Screen (i.e., remove) receivers/IRB traces with divergent or outlying IRBs then find average. This screening may be performed in the execution of Equation 6 for the following cases:
   a. Non-consistent/non-smooth biases
   b. Removal of bias(es) furthest from median (i.e., assume those are always faulted), then use mean of remaining IRBs

Outliers may result from of excessive multipath, receiver hardware faults. They affect only one or two receivers at any given time and are confined to at most at one location/reference site. These biases may affect the IRBs and, in turn, the instantaneous detection metric results themselves. However, their effects may be mitigated as follows:

IRBs

The magnitudes of the IRBs are not as important as their distribution and their smoothness (or continuity) and consistency. The detection metric (or per-correlator) IRBs should result from the mean of the aggregate of screened IRBs from each receiver. A proposed screening algorithm would exclude the two receiver IRBs farthest from the median IRB at all points in time. (Note that for robustness, this median should be applied only over SVs being viewed by a large number of receivers, N (e.g., N>10).) This exclusion would not result in a rejection of any given receiver. However, the number or exclusions per receiver should be logged continuously for maintenance and tracking purposes. The mean IRB of N-2 receivers (for large N) should be relatively smooth if properly screened and should not add significant additional variation to the final detection metric computation.

Weighted Averages

The effects of excessive multipath may be reduced by weighting each receiver metric computations with an independent measure of the multipath. In WAAS, the CNMP monitor may provide such a measure. (Details of this implementation are not yet specified.)

STEP 5: "LEVEL" METRICS TO MEAN INTER-RECEIVER BIAS AND AVERAGE ACROSS RECEIVERS

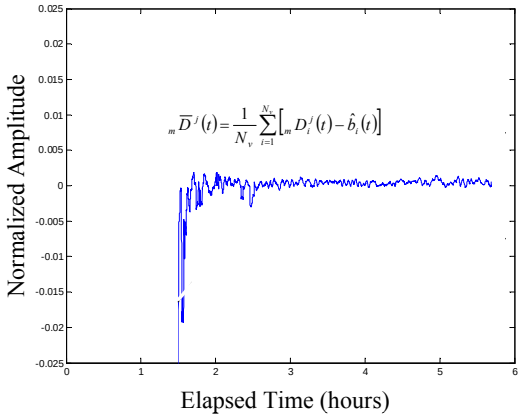

FIGURE 12. Receiver-averaged, "alpha" metric (after IRB removal) for SV14. (Figures 6 and 8 combine to produce this plot.)

After removing the IRBs, a simple, constant-gain average may be taken over the total number of detection metrics (as shown in Figure 12) according to:

$$ {}_m\overline{D}^j(t)=\frac{1}{N_v}\sum_{i=1}^{N_v}\left[{}_mD_i^j(t)-\hat{b}_i(t)\right] \qquad (10)$$

In the above equation, N_v is the number of valid receivers viewing SV j.

A weighted average, based on CNMP information for each receiver, may also be applied here. (The specifics of this approach are TBD.)

STEP 5A: COMPUTE DETECTION THRESHOLDS (IN INITIALIZATION STAGE ONLY)

The detection thresholds are specific to a given monitoring configuration. In other words the number of receivers, N, their antenna locations, the relative paths of the satellites and also the specific receiver configurations and SQM implementations will dictate the threshold magnitudes. Assuming all these conditions are fixed, threshold computation is a one-time calibration/set-up for offline monitoring (or WAAS) that forms thresholds, T_m, for each metric m based on UDRE or, for the purposes of this example, SV elevation angle, θ.

The thresholds are computed by first obtaining elevation angle-dependent (or UDRE-dependent) standard deviations of all detection metrics. (Note that since SV elevation angles vary continuously, the metric standard deviations may be taken from discrete, 5-degree elevation angle "bins.") Next, these standard deviations are appropriately Gaussian-overbounded. (The overbounding process should remove any SV-dependence.) Finally, they are each multiplied by a k-factor (k_{ffd}) corresponding to achieving the desired probability of false-alarm. Equation 11 gives the exact form of these thresholds. (See Figure 13.)

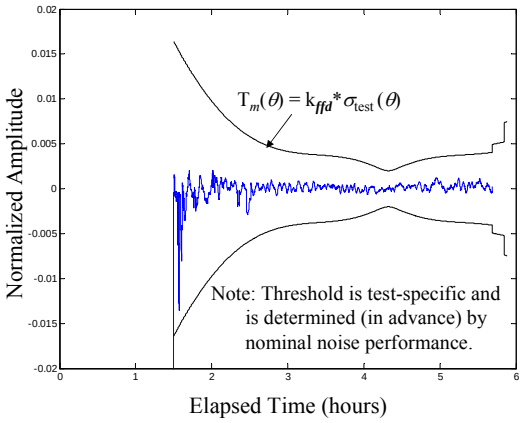

FIGURE 13. Final ("alpha" metric) detection metric and the corresponding threshold for SV14.

$$T_m(\theta)=k_{ffd}\cdot\left[\mathbf{E}\left\{{}_m\overline{D}^j(\theta)\right\}^2\right\}-\left\{\mathbf{E}\left[{}_m\overline{D}^j(\theta)\right]\right\}^2\right]^{1/2}\Rightarrow k_{ffd}\cdot{}_m\sigma_{test}(\theta)$$

or

$$T_m(UDRE)=k_{ffd}\cdot\left[\mathbf{E}\left\{{}_m\overline{D}^j(UDRE)\right\}^2\right\}-\left\{\mathbf{E}\left[{}_m\overline{D}^j(UDRE)\right]\right\}^2\right]^{1/2}\Rightarrow k_{ffd}\cdot{}_m\sigma_{test}(UDRE) \qquad (11)$$

STEP 6: COMPARE METRICS TO THRESHOLDS (IN OPERATING MODE)

Anomalous signal deformation may be present on a signal if the detection metric exceeds the threshold (as shown in Figure 13). This is the final monitoring check which must be performed at each measurement update. It is described in the equation below.

$$\max_m\left[{}_m\overline{D}^j(UDRE)\right]=\begin{cases}\geq 1, & \text{deformation detected, increase UDRE, or}\\ <1, & \text{signal within operating specifications}\end{cases} \qquad (12)$$

CONCLUSIONS

A practical algorithm for SQM includes the following:

- Conventional normalization removes variations due to signal power
- Type-Normalization removes differences across different PRNs
- Time averaging smoothes the noise and multipath
- Inter-receiver bias computations measure the offsets across individual SQM receivers and help enable detection of receiver failures and

small, nominal deformations on individual SV signals
- Receiver averaging permits better averaging of multipath by combining measurements taken from non-collocated antennas

IRBs are critically important since they allow observation of several additional failure modes including receiver malfunctions and slowly time-varying, nominal signal deformations. However, once adequate thresholds are computed for any given siting conditions, SQM can be performed using pre-determined detection tests.

REFERENCES

[1] Brenner, M., Kline, K., Reuter, R., "Performance of a Local Area Augmentation System (LAAS) Ground Installation," *Proceedings of the 15th International Technical Meeting of the Satellite Division of the Institute of Navigation*, ION-GPS-2002, pp. 39-50.

[2] Enge, P. K., Phelts, R. E., Mitelman, A. M., "Detecting Anomalous signals from GPS Satellites," ICAO, GNSS/P, Toulouse, France, 1999.

[3] Phelts, R. E., Akos, D. M., Enge, P. K., "Robust Signal Quality Monitoring and Detection of Evil Waveforms," *Proceedings of the 13th International Technical Meeting of the Satellite Division of the Institute of Navigation*, ION-GPS-2000, pp. 1180-90.

[4] Phelts, R. E., Walter, T., Enge, P., "Toward Real-Time SQM for WAAS: Improved Detection Techniques," *Proceedings of the 16th International Technical Meeting of the Satellite Division of the Institute of Navigation*, ION GPS-2003.

[5] Pullen, S., Luo, M., Gleason, S., Xie, G., Lee, J., Akos, D., Enge, P., Pervan, B., "GBAS Validation Methodology and Test Results from the Stanford LAAS Integrity Monitor Testbed," *Proceedings of the 13th International Technical Meeting of the Satellite Division of the Institute of Navigation*, ION-GPS-2000, pp. 1191-201.

[6] Shloss, P., Phelts, R. E., Walter, T., Enge, P. K., "A Simple Method of Signal Quality Monitoring for WAAS LNAV/VNAV," *Proceedings of the 15th International Technical Meeting of the Satellite Division of the Institute of Navigation*, ION-GPS-2002, pp. 800-8.