

LAAS Ground Facility Design Improvements to Meet Proposed Requirements for Category II/III Operations

Sam Pullen, Ming Luo, Gang Xie, Jiyun Lee, R. Eric Phelts, Dennis Akos, and Per Enge

Department of Aeronautics and Astronautics, Stanford University

ABSTRACT

Stanford University has developed a Local Area Augmentation System (LAAS) ground facility prototype known as the Integrity Monitor Testbed (IMT) to demonstrate the feasibility of LAAS precision approaches under Category I conditions. While the Category I IMT is essentially complete, research on IMT algorithms continues to improve its performance so that it can eventually meet Category II/III approach requirements. To the extent possible, it is desirable to satisfy Category II/III requirements with modifications to the existing single-frequency (L1) LAAS architecture in order to provide Category II/III initial operational capability (IOC) before the second civil frequency (L5) is present on a sufficient number of GPS satellites. This will also provide a backup operational mode in a future dual-frequency LAAS if either L1 or L5 is interfered with.

This paper addresses IMT improvements to detection of satellite signal deformation, code-carrier divergence monitoring of both potential satellite failures and ionosphere spatial anomalies, and position-error monitoring at a “remote” monitor receiver that is some distance away from the existing reference receiver antennas. With these relatively-limited modifications to the existing Category I LGF architecture, significant performance improvements are demonstrated. While the degree to which these improvements are sufficient depends on changes now being considered to the Category II/III requirements, we believe that, with further refinement, they will be sufficient to provide acceptable IOC and dual-frequency backup availability.

1.0 Introduction

The Stanford University Integrity Monitor Testbed (IMT) is a Local Area Augmentation System (LAAS) Ground Facility (LGF) prototype that meets the requirements to support civil aviation operations up to and including Category I precision approach. The IMT

includes monitor algorithms designed to detect all failure modes of concern to the LAAS Signal-in-Space along with fault-handling logic known as Executive Monitoring (EXM). The IMT has been tested under both nominal conditions and with injected failures to verify its ability to meet the Category I requirements, as specified in the FAA LGF Specification [1]. Details of the existing Category I IMT design and performance can be found in previous papers [2,3,4,8].

While the IMT is designed to maximize integrity and continuity performance for Category I, significant enhancements will be needed to improve performance sufficiently to meet the tighter Category II/III precision approach and landing requirements, which are now being revised by RTCA SC-159 WG-4. The existing LAAS MASPS [5] requirements for Category II/III specify a level of integrity risk that is 20 – 40 times tighter than that for Category I, while the continuity risk requirement is tighter by a factor of 4 – 8. Meeting both of these tightened requirements simultaneously is a challenge for integrity monitoring. In addition, the LGF time-to-alert when hazardous failures occur is reduced from three seconds to approximately one second, which mandates faster monitor response without harming continuity.

This paper describes three enhancements to the IMT that are being designed and tested to see if they can meet the projected Category II/III LGF requirements. Section 2.0 discusses monitoring of satellite signal deformation, which has been enhanced by adding a fourth “ultra-narrow” (0.05 chip spacing) correlator pair to the three included in the Category I baseline “SQM 2b” monitor [7]. Another element of SQM, the code-carrier divergence monitor, is the focus of Section 3.0. This monitor has been significantly enhanced by implementing a Cumulative Sum or CUSUM algorithm, which significantly speeds up detection of anomalies that take time to build up in smoothed pseudoranges before they threaten users. Section 4.0 discusses the use of a remote monitor receiver to more tightly bound the user protection

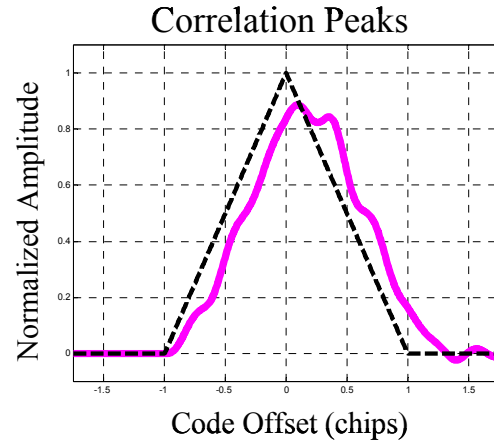
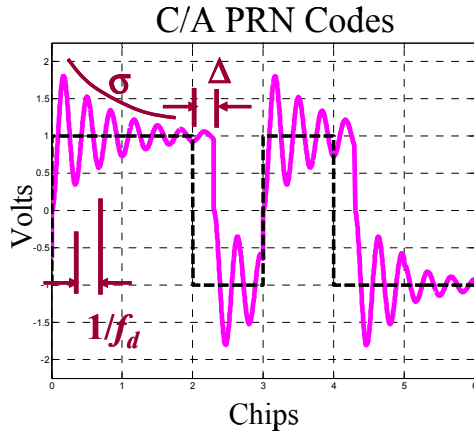


Figure 1: “2nd-Order Step” Signal Deformation Example and Corresponding Correlation Peak

levels assumed in the LGF to enhance overall continuity risk and enhance EXM. Taken together, these improvements demonstrate that a significant fraction of the gap between Category I and Category II/III requirements can be satisfied by relatively minor upgrades to the existing LAAS ground system.

2.0 Enhanced Signal Deformation Monitoring

One potential failure mode of concern to LAAS is a subtle failure of the signal generating hardware onboard the satellite may distort the incoming signal. These anomalous or “evil” waveforms (EWFs, otherwise known as “signal deformations”) distort the correlation function generated within a GPS receiver. This affects code-tracking loops and leads to erroneous pseudorange measurements. Furthermore, for receivers of different configurations (i.e., discriminator types, correlator spacings, and front end bandwidths) these correlation peak distortions result in different pseudorange errors. Since user receivers vary and differ from the reference receivers, these errors cannot, in general, be differentially corrected.

The ICAO EWF threat models this class of integrity faults as a combination of both digital and analog failure modes on the satellite signal-generating hardware [6]. The digital parameter Δ models a lead (or lag) of the falling edge of the C/A code chip transition. The parameters f_d and σ model the frequency and damping of a (2nd-order) failed, analog filter response. The 2nd-order response is given by:

$$e(t) = \begin{cases} 0 & t \leq 0 \\ 1 - \exp(-\sigma t) \left[\cos \omega_d t + \frac{\sigma}{\omega_d} \sin \omega_d t \right] & t \geq 0 \end{cases}$$

$$\omega_d = 2\pi f_d$$

Figure 1 illustrates an example of these waveforms for $f_d = 3$ MHz, $\sigma = 0.8$ MNepers/sec, and $\Delta = 0.3$.

Figure 2 depicts, to scale, the correlator spacings used for SQM. For SQM2b, which has been validated as sufficient to protect all airborne receiver designs permitted in the LAAS MOPS [8], the early-to-late spacings are 0.1, 0.15, and 0.2 chips wide; the code tracking loop employs the pair at 0.1 chips. To increase detection capability, Cat II/III SQM analysis will employ an additional correlator spacing pair (at 0.05 chips). This is made possible by the availability of receivers such as the NovAtel OEM-4 receiver, which can (with specially-designed firmware) provide 8 correlator outputs for each of 12 different satellites.

Nominal thermal noise and multipath at LAAS installations cause distortions of the correlation peak which can conceal the presence of an EWF. Figure 3 illustrates this nominal distortion of the correlation peak for a single satellite pass. For this figure, actual correlation peak measurements were taken using a single-channel, 48-correlator receiver with the full set of spacings shown in Figure 2.

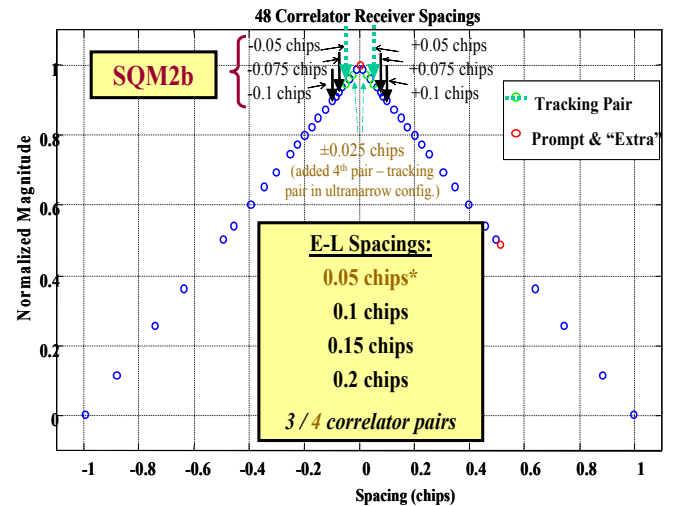


Figure 2: “SQM 2b” Multicorrelator Receiver and Addition of 4th Correlator Pair

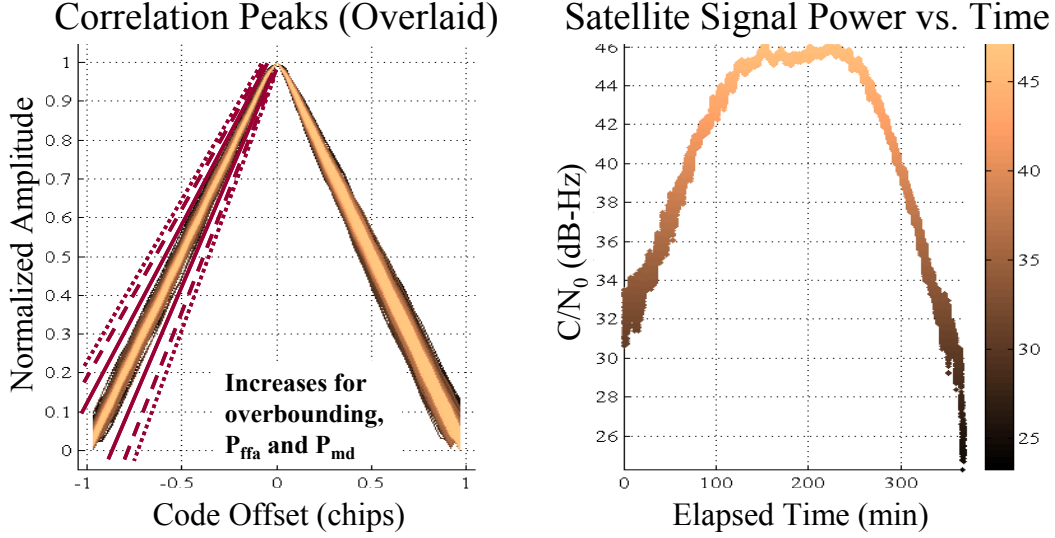


Figure 3: Nominal Distortion of Correlation Peak Due to Multipath and Thermal Noise

Because of the nominal errors shown in Figure 3, detection tests formed using algebraic combinations of these correlator outputs have minimum values beneath which they cannot detect EWFs with sufficiently small fault-free alert (to protect continuity) and missed detection (to protect integrity) probability. These values are known as minimum detectable error (MDEs). MDEs are site-specific and must generally be pre-computed and stored during LGF site installation. They effectively represent the detection capability of each test.

The MDE data collection and computation process has been presented in detail in [6,7]. Here is a summary:

- 1) Record nominal multipath and noise measurements over multiple satellite passes.
- 2) Form each test statistic (delta, ratio, fit, delta-squared, etc.) using the multiple correlator outputs (filtered by 100 seconds of averaging).
- 3) Bin the data according to elevation angle (note that, for the current IMT survey-grade antennas, higher multipath distortion occurs for lower elevation angles).
- 4) For each test statistic, compute standard deviations in each elevation bin. This is known as σ_{test} .
- 5) Multiply by the sum ($K_{ffd} + K_{md}$) of the fault-free detection and missed-detection multipliers needed to achieve the allocated continuity and integrity probabilities based on a standard Gaussian distribution with zero-mean. To the extent that the actual test statistic distribution violates this Gaussian assumption at the required probabilities, σ_{test} must be inflated such that the probability derived from the Gaussian distribution “overbounds” the true probability.

- 6) Fit a 3rd-order polynomial to the MDEs as a function of elevation angle.

Figure 4 summarizes the detection performance for the 3-correlator pair SQM configuration (SQM 2b). The MDEs used were taken using four satellite passes and smoothed using a 100-second moving average filter. No overbounding of non-Gaussian errors was performed, but a margin of 20% (20% inflation of the measured σ_{test}) was applied to provide some protection against test statistic error variation. Also, several additional detection tests were used (as explained in [7]) compared to the original definition of SQM 2b in [6]. The maximum user pseudorange errors consistently remain below the maximum errors allowable (MERR) for Category I LAAS approaches, shown by the top two solid curves. The lower dashed curves are based on Multipath Limiting Antennas (MLAs) and roughly correspond to the MERRs for Category II/III performance if the 10-meter vertical alert limit (VAL) for Category I is retained for Category II/III. Note that only 667 of approximately 15,000 EWFs remained undetected at a 5° elevation angle for both $\Delta\Delta$ and Early-Late user receivers. Because the 20% inflation applied may not be sufficient for overbounding at the required probabilities, the actual margin between MERRs and MDEs is likely smaller than what is shown in these results.

Figure 5 summarizes the detection performance for the four-correlator-pair SQM configuration, which employs “ultranarrow” DLL-tracking at a 0.05=chip spacing. As for SQM 2b, these MDEs used were taken using four satellite passes and smoothed using a 100-second moving average filter, with no sigma overbounding beyond the 20% inflation of each σ_{test} . The improved theoretical

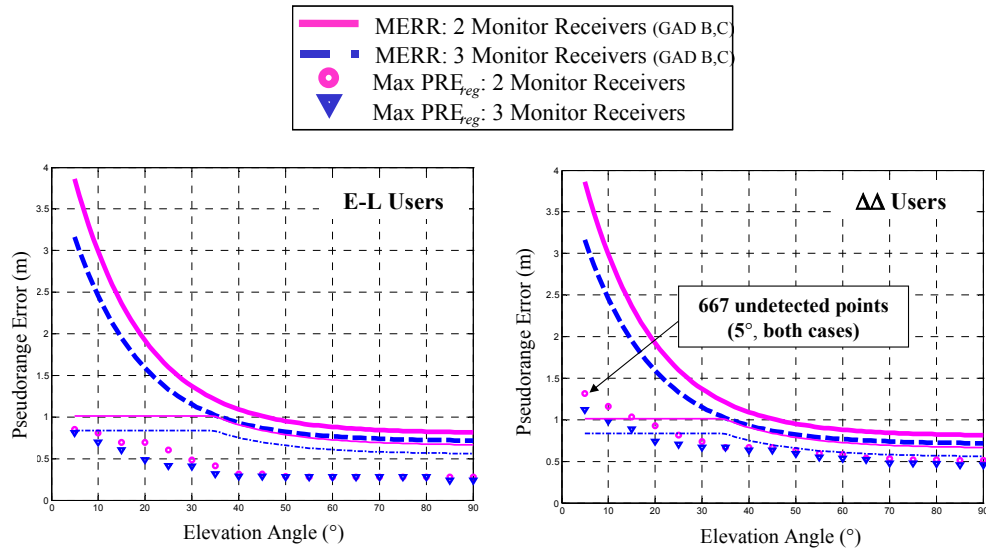


Figure 4: Theoretical Performance of *SQM 2b* (Narrow Tracking, 3 Correlator Pairs) for E-L and $\Delta\Delta$ User Receivers




Figure 5: Theoretical Performance with Ultranarrow Tracking, 4 Correlator Pairs

performance shown here is primarily due to the additional correlator pair. It enables the use of more numerous (and more sensitive) detection tests. Again, a dashed line is drawn to indicate the expected SQM performance given LGF use of MLAs. Note that the MDEs for the MLA should be no worse than those at approximately 35° for a traditional antenna [5,14].

Again, while overbounding would reduce the error margin shown in these results, what is significant is the factor of approximately 2 – 4 of margin that exists with the addition of the fourth correlator pair. As long as the combination of MERR “tightening” to satisfy the final Category II/III Vertical Alert Limit (VAL) and loss-of-integrity probability and required inflation of σ_{test} for overbounding does not exceed this theoretical gap, the

Category II/III evil-waveform detection requirements will be satisfied without requiring major changes to the LGF or to the allowed airborne receiver design space.

3.0 Enhanced Monitoring of Code-Carrier Anomalies

3.1 Introduction

The code-carrier divergence monitor within the SQM function was implemented to detect the remote (and never before observed, to the authors’ knowledge) possibility of satellite generated code-carrier divergence that would impact users whose carrier smoothing filters differ from that of the LGF. More recently, a handful of very large spatial ionosphere gradients have been discovered in WAAS “supertruth” data during ionosphere storms [15].

These gradients are large enough to be hazardous to LAAS users under worst-case ionosphere approach geometries [10]. However, once an ionosphere “wave front” of this type affects an LGF ionosphere pierce point, it can be detected by the code-carrier divergence monitor. While the existing monitor, in combination with other monitors, appears sufficient for Category I [10], improving the time-to-detect of this monitor will likely be needed to meet the Category II/III integrity requirement.

The ionosphere affects GPS signal propagation by making code-phase measurements ρ longer while making carrier phase measurements ϕ shorter. Thus, opposite signs appear in front of the ionosphere error components I in the GPS navigation equations (1) and (2) [16]:

$$\rho = r + b_u - B^s + I + T + M + \varepsilon_\rho \quad (1)$$

$$\phi = r + b_u - B^s - I + T + N\lambda + \varepsilon_\phi \quad (2)$$

where M represents multipath error and ε_ρ and ε_ϕ represent code-phase and carrier-phase measurement errors. Since anomalous code-carrier divergence undermines the accuracy and integrity of carrier-smoothed code (CSC) and may be non-common between the LGF and aircraft (particularly in the case of ionospheric spatial gradients), the LGF must detect such anomalies before the resulting differential CSC errors become hazardous.

The current IMT uses a geometric moving averaging (GMA) method to estimate the divergence d [10,17]:

$$d(k) = \frac{\tau - T_s}{\tau} d(k-1) + \frac{1}{\tau} (z(k) - z(k-1)) \quad (3)$$

where k is the epoch index, $T_s = 0.5$ seconds is the IMT GPS measurement update rate, τ is the time constant of averaging, and z equals the raw code-minus-carrier:

$$z(k) = \rho(k) - \phi(k) = 2I(k) + M(k) - N\lambda \quad (4)$$

$$z(k) - z(k-1) = 2(I(k) - I(k-1)) + (M(k) - M(k-1)) \quad (5)$$

The GMA monitor in (3) gives low-noise divergence estimates by averaging multipath differences $M(k) - M(k-1)$. Though this method is very stable, it is slow to detect changes in $I(k) - I(k-1)$ since they are also averaged.

When a new satellite rises in view or a filter reset occurs after a satellite is temporarily excluded, a variable time constant $\tau(k)$ is used for the first $\tau_{\max} = 200$ sec:

$$d(k) = \frac{\tau(k) - T_s}{\tau(k)} d(k-1) + \frac{1}{\tau(k)} (z(k) - z(k-1)) \quad (6)$$

where

$$\tau(k) = \begin{cases} kT_s, & kT_s < \tau_{\max} \\ \tau_{\max}, & \text{else} \end{cases} \quad (7)$$

Figure 6 shows the results of a failure test using the GMA method. An ionosphere gradient with a magnitude of 0.04 m/s and duration of 173 seconds (representing a severe ionospheric spatial gradient based on the model in [10]) was added to the IMT measurements for one satellite (PRN 2) at an elevation angle of about 70 degrees and a GPS time of 215,000 seconds. The original IMT measurements are from a 12-hour data set collected on February 25, 2002. The IMT post-processed the failure-injected GPS measurements with executive monitoring disabled so that the times-to-detect from all relevant monitors can be observed. In the receiver channel shown here (RR 0 and PRN 2), detection occurs 58.0 seconds after the ionospheric gradient is injected. Note that this does not represent the “time to alert” since time to alert is measured from the time that the error becomes hazardous, not from the time that the anomaly begins.

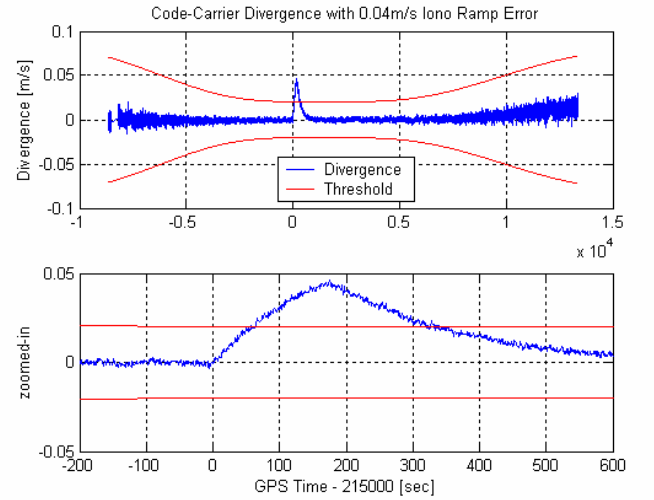


Figure 6: GMA Divergence Failure Test for PRN 2

Other monitors in the IMT, such as MQM acceleration-ramp-step test and innovation test, can also detect ionospheric gradients. The acceleration-ramp-step test can detect very large ionospheric gradients in seconds, but it usually fails to detect smaller (but still hazardous, and presumably less rare) ones. Furthermore, the GMA divergence test is very similar to the innovation test in practice, and the innovation test sometimes detects slightly faster than the GMA divergence method mostly due to its smaller time constant. Thus, our goal is to develop a divergence test that detects smaller anomalies significantly faster than the GMA and innovation tests. Aside from the GMA, two different methods for detecting ionospheric gradients are discussed in this paper. Section 3.2 explores a stochastic predictor using adaptive filtering, while Section 3.3 focuses on the second method – the Cumulative Sum (CUSUM) method.

3.2 Adaptive Filtering of Divergence

Adaptive filtering is widely used in signal processing. We use a method similar to that developed for a

completely different application in [18] to adaptively “predict” divergence.

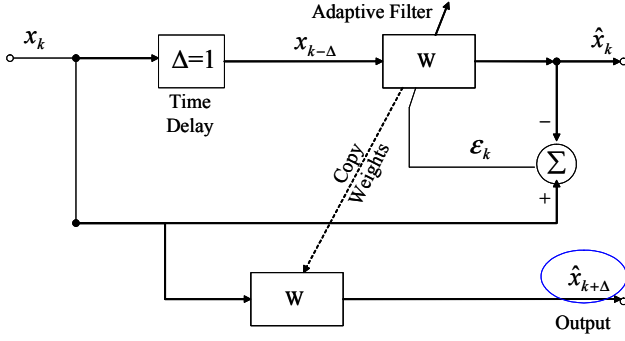


Figure 7: Adaptive-Filter-Based Stochastic Predictor

Figure 7 is a block diagram of an adaptive predictor [19]. It consists of two parts: adaptive filtering and prediction. There are $L + 1$ weights ($\omega_{k0}, \omega_{k1}, \dots, \omega_{kL}$) for the corresponding inputs to the adaptive filter ($x_{k-1}, x_{k-2}, \dots, x_{k-L-1}$). The output \hat{x}_k is compared to a desired response that is equal to x_k , and the error ε_k is used to adjust the weights. The linear combination of $x_{k-1}, x_{k-2}, \dots, x_{k-L-1}$ matches x_k with the minimum error in the least-squares sense, while the linear combination of $x_k, x_{k-1}, \dots, x_{k-L}$, with the same weights copied from the adaptive filter, represents the prediction of x_{k+1} . The prediction \hat{x}_{k+1} is then regarded as the divergence estimate for the current epoch k .

In our adaptive filtering of divergence, raw divergence $dz(k)$ normalized by its standard deviation at the corresponding elevation angle is the input x_k . The normalization helps the adaptive filter converge quickly. We define a general form of the raw divergence $dz(k)$ as:

$$\begin{aligned} dz(k) &= \frac{1}{2T_s k_0} (z(k) - z(k - k_0)) \\ &= \frac{I(k) - I(k - k_0)}{T_s k_0} + \frac{1}{2T_s k_0} (M(k) - M(k - k_0)) \\ &= \dot{I}(k) + \frac{1}{2T_s k_0} (M(k) - M(k - k_0)) \end{aligned} \quad (8)$$

A normalized least-mean-square (NLMS) method is used in the adaptive filter. With a vector representation, it is expressed as [19]:

$$\begin{cases} \varepsilon_k = x_k - W_k^T X_{k-1} \\ W_{k+1} = W_k + \frac{\alpha \varepsilon_k}{\|X_{k-1}\|^2} X_{k-1} \end{cases} \quad (9)$$

where $W_k = [\omega_{0,k}, \omega_{1,k}, \dots, \omega_{L,k}]^T$ and $X_k = [x_k, x_{k-1}, \dots, x_{k-L}]^T$.

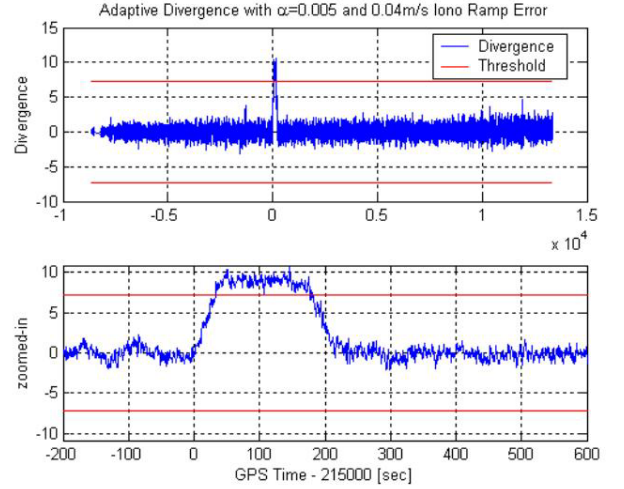


Figure 8: Adaptive Divergence Failure Test for PRN 2

To compare the performance of the GMA and adaptive filter methods, a failure test identical to that shown in Figure 6 was conducted on the adaptive predictor with $L = 10$, $k_0 = 80$, and $\alpha = 0.005$. Figure 8 shows that it only takes 32.5 seconds to detect the gradient, which is much faster than the GMA method. However, this degree of improvement is unusual. For most anomaly scenarios tested, the adaptive method reduces time to detect by less than 10 – 20%. Because the adaptive predictor needs a high signal-to-noise ratio (SNR) to be of any benefit, the epoch delay k_0 is set to 80 to provide significant noise reduction (in a manner analogous to smoothing), but this lengthy delay produces a detection lag similar to that of the GMA method. While the adaptive weights do provide a small improvement over the fixed GMA gains, this improvement does not appear to justify the complexity of this approach.

3.3 Divergence Cumulative Sum (CUSUM)

A variant of the CUSUM method is used for sigma-mean monitoring in the IMT, as described in [9]. In this paper we use a similar CUSUM approach to monitor the code-carrier divergence. The CUSUM approach has the advantage of having the theoretically minimal time-to-detect of a change in the parameters of a random process under certain assumptions, and it has been shown to work well in practice [20,21]. Because anomalous code-carrier divergence takes time to affect users and cannot be detected immediately in any case, the CUSUM approach is applicable, although it is complicated by the fact that nominal code-carrier divergence is not zero-mean at any given time (it is always slightly positive or negative depending on the nominal ionosphere delay trend).

The CUSUM algorithm to detect a positive mean jump in a discrete signal $Y(k)$ is expressed as [20]:

$$C^+(k) = \left[C^+(k-1) + \frac{\mu_1(k) - \mu_0(k)}{\sigma^2(k)} (Y(k) - K(k)) \right]^+ \quad (10)$$

where the superscript “+” means that $C^+(k)$ is reset to a non-negative value if it is less than zero. In this case, the input $Y(k)$ has a Gaussian distribution, i.e., $Y \sim N(\mu_0, \sigma^2)$, although this is not required in general. We will explain the other parameters in this equation shortly. Note that the parallel CUSUM C^- , used to detect negative mean jumps is completely analogous to (10), so we will limit our discussion to positive mean jumps and C^+ .




Figure 9: Raw Divergence dz for all Channels ($k_0 = 1$)

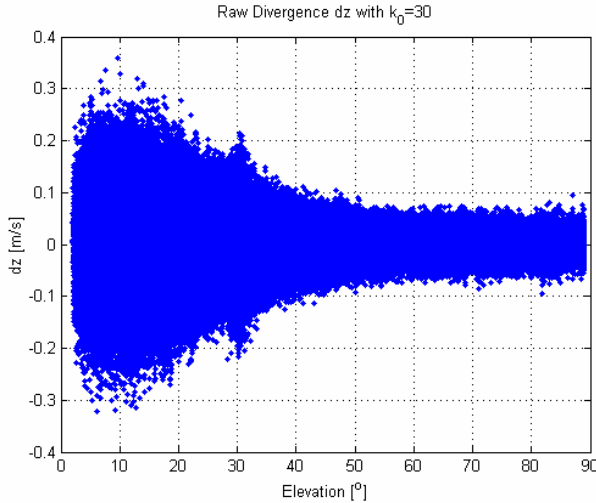


Figure 10: Raw Divergence dz for all Channels ($k_0 = 30$)

As for the adaptive predictor, the input to the CUSUM is also the raw divergence $dz(k)$, i.e.:

$$Y(k) = dz(k) \quad (11)$$

where $dz(k)$, from (8), has a high SNR with a large k_0 , although a larger k_0 implies a longer monitor lag. This advantage is shown in Figures 9 and 10, which demonstrate that the noise in dz is reduced by about a factor of 10 for $k_0 = 30$ compared to $k_0 = 1$. The additional lag is of lesser consequence because we accept that no algorithm can flag small anomalies in less than 15 – 30 seconds. Based on trial and error with nominal IMT datasets, the value of k_0 is selected as 30 in the code-carrier divergence CUSUM.

The “constant” windowing parameter $K(k)$ is calculated by the following formula [20]:

$$K(k) = \frac{\mu_0(k) + \mu_1(k)}{2} = \mu_0(k) + \frac{\nu(k)}{2} \quad (12)$$

where $\mu_1(k)$ is the out-of-control mean of $Y(k)$, and

$$\mu_1(k) = \mu_0(k) + \nu(k) \quad (13)$$

$\nu(k)$ can be regarded as the target magnitude of the ionospheric gradients that we are trying to detect. It is defined as:

$$\nu(k) = \dot{I}_{90} \cdot OF(k) \quad (14)$$

where \dot{I}_{90} is fixed at a 90-degree-elevation ionospheric rate of 0.01 m/s. At other elevation angles, $\nu(k)$ is magnified by the obliquity factor $OF(k)$ that equals [16]:

$$OF(k) = \left(1 - \left(\frac{R_E \cos(\theta_{elev}(k))}{R_E + h_i} \right)^2 \right)^{-1/2} \quad (15)$$

\dot{I}_{90} is set at this relatively small value (for ionospheric anomalies – it is approximately the smallest anomaly that could be hazardous to LAAS users) because our divergence CUSUM monitor is designed to detect small ionospheric gradients. It performs optimally for gradients of this magnitude, and it will also detect gradients with larger magnitudes with near-optimal speed.

Returning to (10), $\mu_0(k)$ is the mean of the input $Y(k)$. Because it is non-zero and changes with elevation angle and ionosphere state, it is estimated in real time by the GMA method (see Section 3.1):

$$\mu_0(k) = \frac{\tau(k)-1}{\tau(k)} \mu_0(k-1) + \frac{1}{\tau(k)} dz(k) \quad (16)$$

where

$$\tau(k) = \begin{cases} k, & k < \tau_{\max} \\ \tau_{\max}, & \text{else} \end{cases} \quad (17)$$

The CUSUM starts to be updated after $2\tau_{\max}$, when $\mu_0(k)$ has nearly converged. For this use, τ_{\max} is 800 epochs, or

400 seconds. Since $\mu_0(k)$ is the mean “in front” of an ionospheric gradient, it takes the value k_1 epochs before in order not to be polluted by the gradients as they appear:

$$u_0(k) \leftarrow u_0(k - k_1) \quad (18)$$

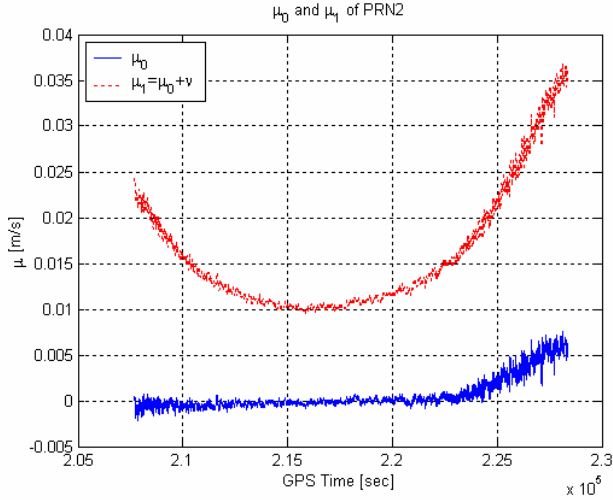


Figure 11: $\mu_0(k)$ and $\mu_1(k)$ for PRN 2 on Receiver 0

Considering both the duration of ionospheric gradients from the model in [10] and typical CUSUM detection times, k_1 is set at 500 epochs. Figure 11 shows an example of $\mu_0(k)$ (nominal mean) and $\mu_1(k)$ (target anomalous mean) on Channel (RR 0, PRN 2). The mean $\mu_0(k)$ changes significantly with time, illustrating the need for real-time mean estimation. Though there is noise in $\mu_0(k)$, variations in $\mu_0(k)$ are very small compared to $v(k)$. The large gap between $\mu_0(k)$ and $\mu_1(k)$ indicates that this real-time mean estimation achieves its purpose.

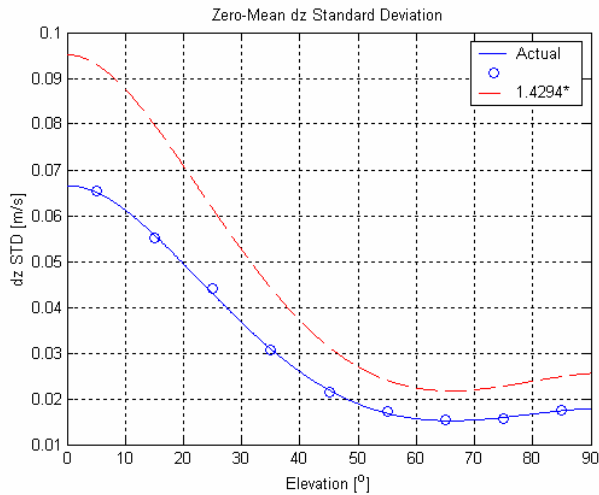


Figure 12 Sample and Inflated Standard Deviation of Raw Divergence dz

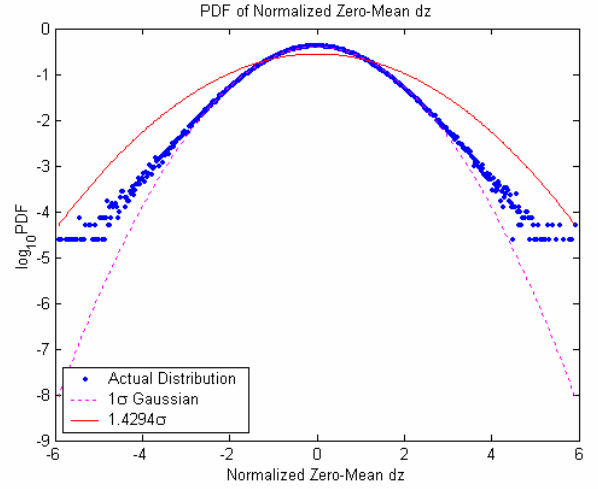


Figure 13: PDF of Normalized Zero-Mean Raw Divergence dz

Unlike $\mu_0(k)$, the variance $\sigma^2(k)$ of the input $Y(k)$ does not require real-time estimation. Instead, it is pre-computed based on statistical analysis of nominal data sets. Starting from zero-mean raw divergence dz over time (almost the same as shown in Figure 10), we calculate standard deviations in nine elevation bins of 10 degrees each and linearly interpolate for other elevations, as shown by the solid blue line in Figure 12. The zero-mean dz is then normalized by its standard deviation at the corresponding elevation. A zero-mean Gaussian distribution with an inflated standard deviation is used to “overbound” the tails of the normalized zero-mean probability density function (PDF) of dz , as shown in Figure 13. The assumed standard deviation is inflated to 1.4294 times the original measured σ . The final standard deviation σ of dz is shown as the red dash line in Figure 12. Inflating σ raises the CUSUM thresholds but is necessary to maintain the required low probability of false free detections (sub-allocated from the assumed continuity requirement).

The CUSUM threshold $h(k)$ is a function of $\mu_0(k)$, $\mu_1(k)$ (or $v(k)$), and $\sigma(k)$ that all change with time or elevation angle. In order to simplify the threshold computations, a normalization process is applied before the input to the CUSUM as follows:

$$\begin{cases} Y(k) \leftarrow \frac{Y(k) - \mu_0(k)}{\sigma(k)} \\ v(k) \leftarrow \frac{v(k)}{\sigma(k)} \end{cases} \quad (19)$$

and then (10) becomes:

$$C^+(k) = \left[C^+(k-1) + v(k) \left(Y(k) - \frac{v(k)}{2} \right) \right]^+ \quad (20)$$

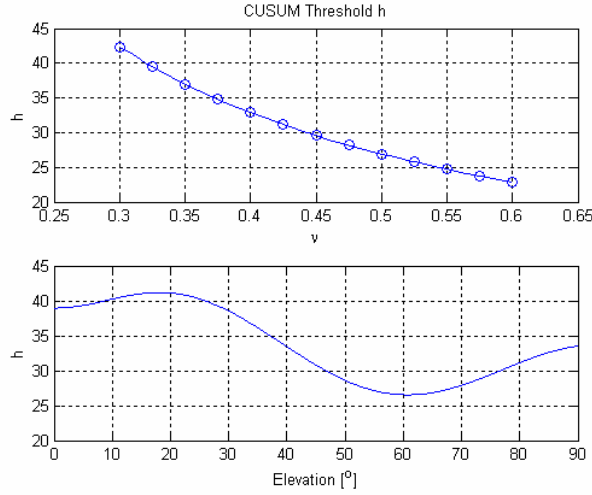


Figure 14: CUSUM Threshold as Function of ν Elevation.

After this normalization, the distribution of the input $Y(k)$ is overbounded by a standard Gaussian distribution. Thus, since $\mu_0(k) = 0$ and $\sigma(k) = 1$ after normalization, the CUSUM threshold $h(k)$ is now only a function of $\nu(k)$. Based on the desired average run length ($ARL \cong 10^7$ epochs) set by the continuity sub-allocation, the threshold $h(k)$ is pre-computed for a range of values of ν using the Markov Chain method and numerical search, as shown in the top plot of Figure 14 (the technique used to model the CUSUM as a Markov chain and iteratively search for the threshold that achieves the required ARL is described in [9,21]). Linear interpolation is used between these sampled points. Since $\nu(k)$ is elevation-dependent, $h(k)$ is also only a function of elevation angle, as shown in the bottom plot of Figure 14. The $h/2$ Fast Impulse Response (FIR) CUSUM is used, in which the value of $C^+(k)$ is reset to $h(k)/2$ if $C^+(k)$ is less than zero on any epoch as computed from (20). Because of the windowing factor $K(k)$, resets like this are common under nominal conditions and can be observed in Figure 16 later on.

Figure 15 shows nominal $C^+(k)$ test results for all channels (all receivers and all satellites). As expected, all points are below the thresholds shown as solid red lines. The amount of margin between the highest CUSUM points and the threshold is not large for some elevation angles, suggesting a possible need to raise the threshold slightly. This will be addressed in Section 3.4.

Failure testing based on the technique used for the GMA and adaptive methods is used to examine typical CUSUM detection times (theoretical predictions of mean ARL under anomaly conditions can be obtained from the Markov Chain model [9,21]). The same satellite used before, PRN 2, is used for these tests since it makes a lengthy pass in view of the IMT and reaches an elevation of greater than 70 degrees. Failures are injected at every

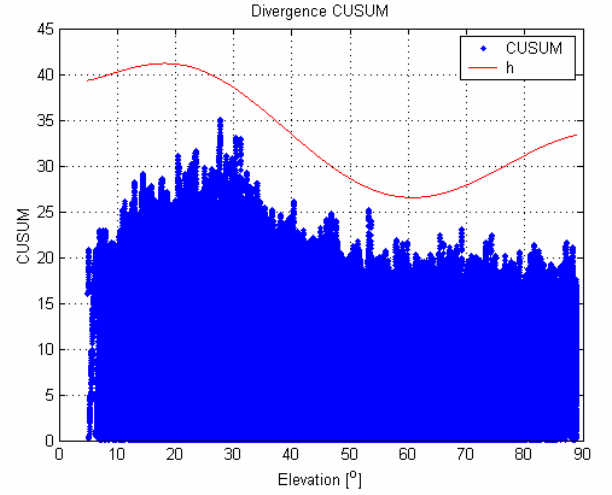


Figure 15: Nominal CUSUM Results (All SV's, RR's)

10 degrees of elevation as PRN 2 rises and sets. Each failure test is independent of the others, as it acts upon a different swath of nominal data.

Elev. [deg]	i [m/s]	Time to Flag [sec]			
		Acc, Ramp, Step	GMA Dvgc.	Inno.	Dvgc. CUSUM
30	0.017514	∞	∞	∞	157.5
40	0.014546	∞	∞	∞	134.5
50	0.012612	∞	∞	∞	97.0
60	0.011357	∞	∞	∞	53.5
70	0.010571	∞	∞	∞	100.5
70	0.010571	∞	∞	∞	99.5
60	0.011357	∞	∞	155.5	84.5
50	0.012612	∞	∞	∞	118.0
40	0.014546	∞	∞	∞	172.5
30	0.017514	174.5	∞	∞	146.5

Table 1: Failure test results for PRN 2 (Inject $i = 0.01 \times OF$ m/s for 173 Seconds)

The first set of failure tests is to check the performance of the CUSUM method at its targeted fault condition, i.e., when the divergence mean jump ν is equal to 0.01 m/s magnified by the corresponding obliquity factor. Table 1 shows the time needed for the acceleration-ramp-step, innovation, GMA divergence, and divergence CUSUM monitors to detect. The symbol " ∞ " indicates that the corresponding method fails to detect that gradient in this test. The table shows that the CUSUM method always detects this gradient at all elevation samples as expected; while the other three monitors fail to with two exceptions. These exceptions are usually due to "fortunate" measurement noise, meaning that the mean response of the test statistic is below its threshold, but 2σ or greater noise in the right direction on a given epoch is sufficient to push the statistic above the threshold briefly.

Elev. [deg]	Time to Flag [sec]			
	Acc, Ramp, Step	GMA Dvgc.	Inno.	Dvgc. CUSUM
30	∞	∞	∞	143.0
40	174.5	∞	153.0	77.0
50	67.0	∞	91.5	59.5
60	174.0	150.0	55.5	25.0
70	174.0	136.0	56.5	36.0
70	81.5	128.0	∞	30.0
60	17.5	153.5	66.0	29.5
50	16.0	163.5	88.0	67.0
40	174.0	∞	173.5	97.5
30	173.5	∞	∞	117.5

Table 2: Failure Test Results for PRN 2 (Inject $\dot{i} = 0.02$ m/s at all Elevations for 173 Seconds)

The second set of failure tests always injects 0.02 m/s at all elevations (i.e., without multiplying by the obliquity factor). As shown in Table 2, the CUSUM method first detects the gradient in eight of the ten cases. The two cases where the acceleration-ramp-step monitor is fastest are again due to “fortunate” noise – when this monitor flags reliably, it does so almost immediately. Unlike the other three monitors, the CUSUM method consistently succeeds in detecting the gradients. The table also shows that the innovation monitor can detect faster than the GMA divergence monitor, as mentioned previously.

Figure 16 shows the CUSUM response over time for the rising 70-degree failure test in Table 2. Before the fault affects IMT measurements (at time 0 on the x-axis), $C^+(k)$ is repeatedly reset to $h(k)/2$ as its value becomes negative. Once the anomaly is injected, the CUSUM rapidly rises and reaches the threshold 36.0 seconds later.

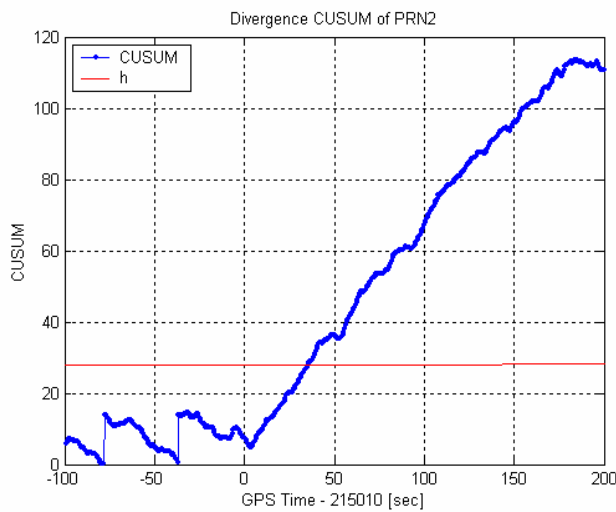


Figure 16: Failure Test Results (Inject Anomaly with $\dot{i} = 0.02$ m/s on PRN 2 at 70°)

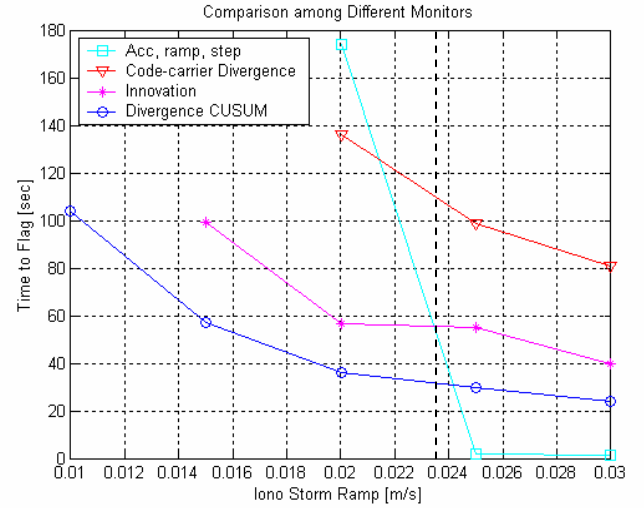


Figure 17: Failure Test Results for PRN 2 (Inject different magnitudes of \dot{i} at 70°)

The final set of failure tests is carried out by injecting ionospheric gradients of 0.01, 0.015, 0.02, 0.025, and 0.03 m/s onto PRN 2 measurements at the same elevation of 70 degrees (rising). Figure 17 shows that the divergence CUSUM consistently detects smaller anomalies faster and more reliably than the other tests in the IMT. Anomalies with gradients larger than about 0.024 m/s at this elevation, the MQM acceleration-ramp-step reliably alerts in less than 5 seconds, with the CUSUM taking 20 – 30 seconds. These results suggest that the combination of the divergence CUSUM and the acceleration-ramp-step monitor provide optimal coverage against all anomalous ionosphere divergences that could threaten LAAS users.

3.4 Divergence Conclusions and Ongoing Work

This results reported in this section demonstrate that the divergence CUSUM has the potential to significantly reduce detection times for relatively small, but still hazardous, ionospheric spatial anomalies that are below the MDE of the MQM carrier-phase acceleration-ramp-step monitor. Adaptive filtering also has potential, but its benefit for divergence monitoring is limited because of the relatively low SNR of the ionosphere divergence signal in the presence of nominal multipath (this SNR may improve when outputs from the FAA Technical Center LAAS Test Prototype (LTP) with Multipath Limiting Antennas are used in the IMT [22]).

One concern regarding the CUSUM’s performance is that it may be too effective in detecting small anomalies that may occur more often than the LGF continuity sub-allocation allows. The targeted CUSUM divergence of 0.01 m/s is anomalous, but it is probably no more than a 4 – 5 σ event at or near the peak of the ionosphere solar cycle. Furthermore, anomalies of this magnitude do not appear to pose any significant threat to LAAS users based

on the analysis in [10] (the baseline threat model in [10] has a divergence rate of about 0.035 m/s).

Thus, while targeting the CUSUM at 0.01 m/s of divergence demonstrates its ability to detect relatively small anomalies in the presence of noise and multipath, it is probably better not to alert at 0.01 m/s and instead target the CUSUM at a higher value (e.g., 0.02 m/s). If necessary to avoid alerts at 0.01 m/s, this target value can be made higher than the value that MQM can reliably detect (about 0.025 m/s), since CUSUM performance in practice is not strongly dependent on the exact target failure value. This will have the effect of increasing the windowing factor $K(k)$ and raising the effective thresholds above the values shown in Figure 15. The detection time shown in Figure 16 for a divergence of 0.02 m/s will likely degrade somewhat but should still be significantly better than the existing GMA and innovation tests. Work on optimizing the CUSUM along these lines is continuing and will incorporate LTP data in the near future.

4.0 Use of Position Domain Remote Monitoring

One of the consequences of the tighter integrity requirements for Category II/III LAAS is that failure modes determined to be too improbable to require mitigation for Category I must be re-evaluated. One example is the so-called “Type A” ephemeris failure, in which a GPS satellite moves far away from its broadcast ephemeris without being alerted or corrected by the GPS Operational Control Segment [11]. A carrier-phase based algorithm to provide guaranteed detection of this class of failure has been proposed in [23]. It requires a baseline between LGF receiver antennas of at least several hundred meters, which is well beyond the typical 100-meter spacing of a Category I LGF.

If this long baseline were created by adding a redundant “remote monitor” receiver (RMR) to the LGF architecture (rather than spreading out all LGF reference receiver antennas and thereby making LGF siting much more

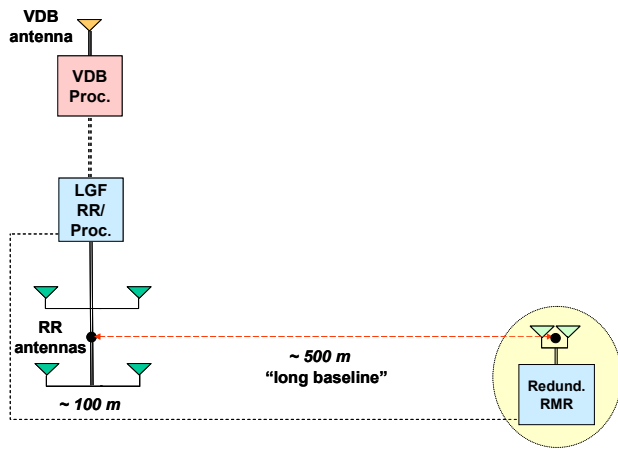


Figure 18: Remote Monitor Receiver Concept




Figure 19: Typical PDM VPL_{H10} at Chicago, Illinois (LGF Contribution Included)

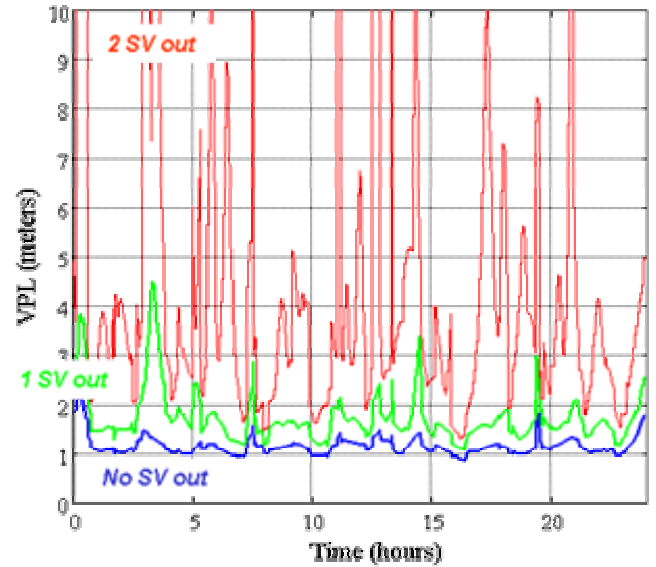


Figure 20: Typical PDM VPL_{H10} at Chicago, Illinois (LGF Contribution Not Included)

difficult), as shown in Figure 18, it could also support the Position Domain Monitor (PDM) concept proposed in [12]. In this concept, the RMR hosting the PDM derives position solutions from the current LGF corrections using all visible satellites approved by the LGF and all reasonable subsets of these satellites that an aircraft may be limited to. These position solutions are compared to the known (surveyed) location of the RMR antenna, and errors exceeding the Vertical or Lateral Alert Limits would be alerted (with the subset position solutions being used to isolate the satellite with faulty measurements). This approach would improve upon the existing sigma-mean monitoring and would support the smaller σ_{pr_gnd} inflation factors needed for high Category II/III

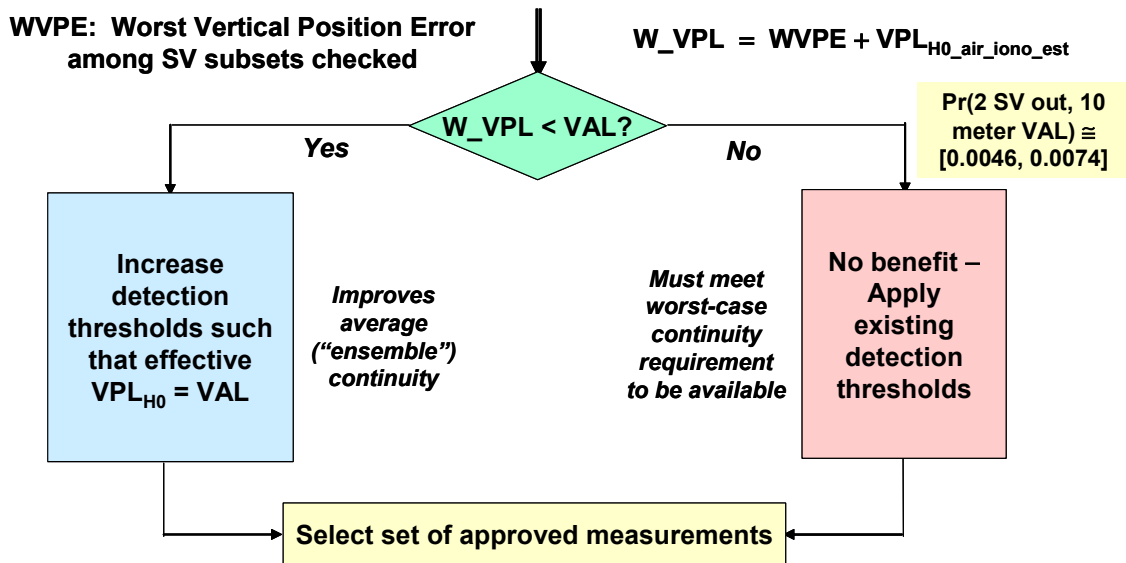


Figure 21: Use of Real-Time PDM to Enhance Average Continuity and Executive Monitoring

availability. Combining RMR/PDM with the CUSUM approach of [9] would improve sigma-mean monitoring further and is a subject of ongoing work.

Given that an RMR is in place to perform PDM, the PDM outputs could potentially be used to relax a key assumption of Category I LGF monitoring, which is that all airborne users have vertical protection levels (VPLs) right at the 10-meter maximum imposed by the VAL [24]. In practice, the truth is almost certainly much better, as shown in Figures 19 and 20. These plots show typical VPLs using the RTCA standard 24-satellite GPS constellation [25] over a 24-hour period in Chicago, Illinois, a 5-degree LGF elevation mask angle, and the LGF error model given in [26] with three exceptions:

- 1) the conservative airborne error variance σ_{air}^2 is reduced by a factor of 3.0;
- 2) σ_{vig} in the ionosphere error model is 0.001 rather than 0.004 m/km (0.001 m/km is more typical of nominal ionosphere spatial gradients);
- 3) the approach speed v (or v_{air}) is 0.070 rather than 0.129 km/s to correspond to typical passenger jet aircraft speeds at Category II/III decision heights.

Figure 19 includes all contributions to nominal VPL (or VPL_{H0}), whereas Figure 20 excludes the LGF contribution. Since the PDM could apply a function of its measured position error in place of the LGF contribution to VPL_{H0} (which is a rare-event upper bound on LGF-induced error), the VPLs used by the PDM are likely to fall in between the results shown in these two figures. Note that the “1 out” and “2 out” results represent the worst single satellite or pair of satellites unavailable to user aircraft, so they are conservative. In both figures, all “no out” and “1 out” cases have VPLs well below the 10 meter Category I VAL, and these represent the vast

majority of cases in practice. Among the worst “2 out” cases, the probability that VPL exceeds 10 meters is 0.0074 with the LGF contribution included and 0.0046 without. For the current Category II/III VAL of 5.3 meters [5], these probabilities increase to 0.0266 and 0.0130, respectively (the eventual Category II/III VAL is likely to be revised to an intermediate value that is likely to be closer to 10 meters [27]).

Figure 21 shows how this information could be used in real time to improve LGF performance. For the vast majority of cases in which the worst computed VPL from the PDM outputs (denoted as W_VPL) is less than VAL, as shown in the left-hand fork, the effective MDEs can be increased by loosening the integrity monitor detection thresholds such that the effective VPL_{H0} equals the VAL. This maintains the required integrity while significantly lowering continuity risk. Aircraft that, on rare occasions, happen to “see” a subset of GPS satellites that was not directly checked by the PDM are still protected by their own VPL calculations – they only suffer a slight increase in integrity risk if their VPL exceeds the W_VPL but is still below VAL (if it were above VAL, the aircraft could not conduct a Category II/III approach). This limited integrity risk increase is deemed acceptable for sigma-mean monitoring if it is sufficiently rare and implies no greater than one order of magnitude of increased overall system risk [1].

The limitation of this concept is that, to achieve a required availability of 0.999 or higher, the LGF must still meet the integrity and continuity requirements when the worst-case PDM VPL exceeds VAL (the right-hand fork of Figure 21). Since no threshold increase is possible in this case, the baseline thresholds must meet the worst-case continuity requirement. However, the improvement in average continuity is significant by itself and is of

substantial benefit to Executive Monitoring (EXM). The most difficult task of EXM is to distinguish between different failure classes and to separate hazardous anomalies from fault-free alerts [3,28]. As thresholds are pushed lower to satisfy tighter Category II/III integrity requirements, smaller off-nominal conditions that are not hazardous to LAAS users are more likely to be flagged, making it more difficult for EXM to distinguish “real” failures. Reducing the rate of off-nominal but non-hazardous exclusions will reduce the conservatism of EXM fault exclusion and recovery and thereby improve overall satellite availability for LAAS in addition to lowering average continuity risk.

5.0 Conclusions and Future Work

This paper has demonstrated several means by which the existing Category I LGF architecture can be improved to at least partially cover the gap between the Category I and (evolving) Category II/III requirements. Monitoring of satellite signal deformation (the key component of Signal Quality Monitoring, or SQM) has been enhanced by adding a fourth “ultra-narrow” (0.05 chip spacing) correlator pair to the three included in the Category I baseline “SQM 2b” monitor. Preliminary data from NovAtel OEM-4 receivers modified to output four correlator pairs for each of 12 satellites shows that this, combined with additional monitor test statistics, significantly improves the minimum detectable errors (MDE’s) that can be supported.

Another element of SQM, the code-carrier divergence monitor, has been enhanced by implementing a Cumulative Sum or CUSUM algorithm. This approach, which has also been used for LGF error mean-sigma monitoring [9], significantly speeds up detection of anomalies that take time to build up in smoothed pseudo-ranges before they threaten users. While satellite-based divergence failures have never been reported to the authors’ knowledge, severe ionosphere spatial anomalies were noted during the last solar maximum, and this improvement is significant in terms of reducing the risk that they pose [10]. Algorithms that attempt to improve performance in real-time by adapting the parameters of the existing code-carrier divergence filter have also been tested but do not show the same level of improvement.

Finally, because a remote monitor receiver some distance (at least several hundred meters) from the centroid of the LGF antennas is needed to detect all possible threatening ephemeris failures [10], an additional tool exists for monitoring the LGF corrections. Using this remote receiver as a position-domain monitor (PDM) has been studied in detail by MITRE/CAASD [12], and its performance could be improved by applying a CUSUM approach similar to that used in [9] and Section 3.3 of this paper. In addition, it has been proposed to use the PDM position error outputs for subsets of satellites in view to

produce an upper bound on airborne protection levels that is less conservative than the worst case value assumed for Category I. This concept can lower average continuity loss risk over all approaches while enhancing the ability of EXM to separate real anomalies from false alarms.

ACKNOWLEDGMENTS

The authors would like to thank Boris Pervan (IIT), Ron Braff and Curt Shively (MITRE/CAASD), Ted Urda (FAA AND-710), Barbara Clark (FAA AIR-130), Victor Wullschlegler and John Warburton (FAA ACT-360), Navin Mathur (AMTI), and Frank Van Graas (Ohio University) for their help during this research. The advice and interest of many other people in the Stanford GPS research group is appreciated, as is funding support from the FAA LAAS Program Office (AND-710). The opinions discussed here are those of the authors and do not necessarily represent those of the FAA or other affiliated agencies.

REFERENCES

- [1] *Specification: Performance Type One Local Area Augmentation System Ground Facility*. U.S. Federal Aviation Administration, Washington, D.C., FAA-E-2937A, Apr. 17, 2002.
- [2] M. Luo, S. Pullen, *et.al.*, “Development and Testing of the Stanford LAAS Ground Facility Prototype,” *Proceedings of the ION 2000 National Technical Meeting*. Anaheim, CA., Jan. 26-28, 2000, pp. 210-219.
- [3] S. Pullen, M. Luo, *et.al.*, “GBAS Validation Methodology and Test Results from the Stanford LAAS Integrity Monitor Testbed,” *Proceedings of ION GPS 2000*, Salt Lake City, UT, Sept. 19-22, 2000, pp. 1191-1201.
- [4] G. Xie, S. Pullen, *et.al.*, “Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed,” *Proceedings of the ION 57th Annual Meeting*, Albuquerque, NM, June 11-13, 2001, pp. 681-693.
- [5] *Minimum Aviation System Performance Standards for the Local Area Augmentation System*. Washington, D.C.: RTCA SC-159, WG-4A, DO-245, September 28, 1998.
- [6] R.E. Phelts, D. Akos, P. Enge, “Robust Signal Quality Monitoring and Detection of Evil Waveforms,” *Proceedings of ION GPS 2000*, Salt Lake City, UT., Sept. 19-22, 2000, pp. 1180-1190.
- [7] R.E. Phelts, A. Mitelman, *et.al.*, “Transient Performance Analysis of a Multicorrelator Signal Quality Monitor,” *Proceedings of ION GPS 2001*, Salt Lake City, UT., Sept. 11-14, 2001, pp. 1700-1710.

- [8] *Minimum Operational Performance Standards for GPS Local Area Augmentation System Airborne Equipment*. Washington, D.C.: RTCA SC-159, WG-4A, DO-253A, Nov. 28, 2001.
- [9] J. Lee, S. Pullen, *et.al.*, "LAAS Sigma-Mean Monitor Analysis and Failure-Test Verification," *Proceedings of the ION 57th Annual Meeting*, Albuquerque, NM, June 11-13, 2001, pp. 694-704.
- [10] M. Luo, S. Pullen, *et.al.*, "Assessment of Ionospheric Impact on LAAS Using WAAS Supertruth Data," *Proceedings of the ION 58th Annual Meeting*, Albuquerque, NM, June 24-26, 2002, pp. 175-186.
- [11] S. Pullen, B. Pervan, *et.al.*, "Ephemeris Protection Level Equations and Monitor Algorithms for GBAS," *Proceedings of ION GPS 2001*, Salt Lake City, UT., Sept. 11-14, 2001, pp. 1738-1749.
- [12] R. Braff, "Position Domain Monitor (PDM) Performance Analysis for CAT III," McLean, VA., MITRE/CAASD, October 1, 2002.
- [13] R. Braff, "Description of the FAA's Local Area Augmentation System (LAAS)," *Navigation*. Vol. 44, No. 4, Winter 1997-98, pp. 411-424.
- [14] G. McGraw, T. Murphy, *et.al.*, "Development of the LAAS Accuracy Models," *Proceedings of ION GPS 2000*, Salt Lake City, UT., Sept. 19-22, 2001, pp. 1212-1223.
- [15] S. Datta-Barua, *et.al.*, "Assessment of Ionospheric Impact on LAAS Using WAAS Supertruth Data," *Proceedings of the ION NTM 2002*, Anaheim, CA., January 24-26, 2002, pp. 175-186.
- [16] P. Misra, P. Enge, *Global Positioning System: Signals, Measurements, and Performance*. Lincoln, MA.: Ganga-Jamuna Press, 2001.
- [17] Boris Pervan, "IIT Research Update," Illinois Institute of Technology, June 28, 2001.
- [18] P. DeLeon, C. Screenan, "An Adaptive Predictor for Media Payout Buffering," *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Vol. 6, March 1999, pp. 3097-3100.
- [19] B. Widrow, S. Stearns, *Adaptive Signal Processing*. Englewood Cliffs, N.J.: Prentice-Hall, 2001.
- [20] M. Basseville and I. Nikiforov, *Detection of Abrupt Changes – Theory and Application*. Englewood Cliffs, N.J.: Prentice-Hall, 1993. Internet URL: <http://www.irisa.fr/sigma2/kniga/>
- [21] D.M. Hawkins, D.H. Olwell, *Cumulative Sum Charts and Charting for Quality Improvement*. New York: Springer-Verlag, Inc., 1998.
- [22] P.L. Normark, D. Akos, *et.al.*, "Integration of the Multipath Limiting Antenna with the LAAS Integrity Monitor Testbed," *Proceedings of ION GPS 2002*, Portland, OR., Sept. 24-27, 2002.
- [23] B. Pervan, F. Chang, "Detecting Global Positioning Satellite Orbit Errors Using Short-Baseline Carrier Phase Measurements," *AIAA Journal of Guidance, Control, and Dynamics*. Vol. 25, No. 6, Nov.-Dec. 2002.
- [24] F. Van Graas, Informal Discussion during Ohio University/Stanford University LAAS Review Meeting, Stanford University, June 21, 2002.
- [25] *Minimum Operational Performance Standards for Global Positioning System/Wide Area Augmentation System Airborne Equipment*. Washington, D.C.: RTCA SC-159, WG-2, DO-229C, Nov. 28, 2001.
- [26] C. Shively, "LAAS CAT III Potential Availability Benefit from Dual Frequency (Preliminary Results)," Revision 1. McLean, VA., MITRE/CAASD, May 22, 2002.
- [27] T. Murphy, "Development of Signal in Space Performance Requirements for GBAS to Support CAT II/III Landing Operations," *Proceedings of ION GPS 2002*, Portland, OR., Sept. 24-27, 2002.
- [28] B. Pervan, S. Pullen, M. Luo, "Navigation Integrity Monitoring for DGPS-Based Precision Landing of Aircraft," *Proceedings of the 1999 American Control Conference*, San Diego, CA., June 2-4, 1999.