THE JOURNAL OF NAVIGATION (2006), **59**, 403–422. © The Royal Institute of Navigation
doi:10.1017/S0373463306003882 Printed in the United Kingdom

Detecting RFI Through Integrity Monitoring at a DGPS Reference Station

Youngsun Yun and Changdon Kee

(*Seoul National University*, *Korea*)
(Email: zoro@snu.ac.kr)

Jason Rife, Ming Luo, Sam Pullen and Per Enge

(*Stanford University*, *USA*)

Because GPS is a radio navigation system which has a very low power level, it is vulnerable to RFI. Excessive RFI could cause receiver performance degradation, such as degradation of position accuracy, loss of lock and increased acquisition time. After GPS modernization plans introduce dual-frequency civil signals to mitigate ionospheric errors, RFI will remain as one of the dominant threats for differential GPS navigation systems. Examples of safety-critical civil aviation and military missions threatened by RFI include the Local Area Augmentation System (LAAS) and the Joint Precision Approach and Landing System (JPALS). This paper focuses on RFI mitigation through integrity monitoring for a DGPS system like LAAS or JPALS. The mitigation strategy consists of two parts. First, the paper develops a new RFI detection method, using a raw divergence statistic. Second, the paper investigates strategies for maintaining integrity in the case that RFI is detected.

To validate the utility of the divergence-based RFI monitor, this paper takes an experimental approach. The experiments assess the performance of the divergence metric and compare it to existing alternatives for RFI detection, such as metrics for Automatic Gain Control (AGC) and carrier-to-noise ratio (C/N0). Generating a monitoring threshold for these statistics proves challenging, because the threshold depends both on the type of RFI threat (e.g. continuous wave, narrow band, wideband, pulsed) and on environmental conditions, such as temperature. As experiments illustrate, the divergence statistic resolves these limitations, as divergence directly estimates ranging source error, independent of the type of RFI threat or the environmental conditions.

KEY WORDS

1. DGPS. 2. Radio Frequency Interference. 3. Integrity Monitoring.

1. INTRODUCTION. The Local Area Augmentation System (LAAS) and the Joint Precision Approach and Landing System (JPALS) are aircraft landing systems based on the Differential GPS (DGPS) concept. These systems consist of a reference station, which provides accurate correction information with integrity,

and users, who utilize the broadcast information to estimate their own positions and navigation error bounds. To develop and evaluate these functions, the Stanford GPS laboratory built an integrity monitor testbed (IMT) for LAAS. It monitors the quality of the GPS signal, navigation data and measurements and generates pseudorange corrections and standard deviations of the correction error suitable for transmission to airborne users (Luo, 2000) (Pullen, 2000) (Xie, 2001) (Normark, 2001).

A significant threat for DGPS that can be investigated using the IMT is excessive radio frequency interference (RFI). GPS has vulnerability to RFI because of its low signal power. RFI causes increase of measurement noise, loss of lock, cycle slip or false lock. Because RFI-induced measurement noise causes deterioration of GPS performance, RFI detection and isolation is essential to GPS related applications, especially safety-critical systems.

To detect the RFI effect on GPS signals, other researchers, such as Ndili (Ndili, 1998) and Luo (Luo, 2003), have focused primarily on automatic gain control (AGC) and carrier to noise ratio (C/N0) statistics. This research has shown that AGC, in particular, is very sensitive to RFI. In isolation, however, the AGC can encounter two problems when used to detect RFI. First, because AGC outputs depend on the nature of the RFI signal, as well as on receiver design and ambient temperature (Bastide, 2003), it is difficult to predict AGC response and to design a threshold for an AGC monitor. Second, AGC monitor outputs do not directly reflect navigation error of the system. Even under weak RFI changes that do not degrade range measurements severely, an AGC monitor may produce a significant response that could affect the continuity performance of the system.

Because of these limitations, it is desirable to define a detection metric that estimates the actual pseudorange error level experienced by the differential GPS ground station. One candidate metric in the IMT is the B-value statistic, used to detect hazardous range corrections or violations of the broadcast error variance bound (Lee, 2001). Because the B-value is based on 100 second smoothing, however, the detection time for RFI threats is too long.

The divergence statistic is a useful complement to other system parameters that might be employed for RFI monitoring. In contrast with the AGC measurement, the divergence statistic provides a direct indication of ranging errors. In this sense, thresholds are more easily defined for a raw divergence monitor than for an AGC monitor, since they are less dependent on the nature of the noise source or the ambient RFI environment. In contrast with the B-value statistic, the raw divergence statistic provides an independent sample from only two sequential epochs and, consequentially, enables rapid detection of RFI threats.

In this paper, AGC, SNR and raw divergence statistics are experimentally compared for various RFI sources. These results motivate an operational strategy to detect and deal with RFI. For civil applications, like LAAS, if variance violation is detected and exceeds our threshold, the system stops operation. For military applications, like JPALS, which may require operation even under hazardous RFI conditions, the detection of RFI results in inflation of thresholds for all monitors and simultaneous inflation of broadcast sigmas, so that users can decide whether or not to continue operating. This strategy can thus provide continuous system operation under moderate RFI conditions without losing integrity.

2. DESIRABLE METRIC.

2.1. *Raw Divergence*. Divergence is a desirable RFI metric. Divergence estimates current error levels because it is a time difference of pseudorange minus carrier phase as follows.

$$\begin{aligned} dz_n(k) &= [\rho_n(k) - \phi_n(k)] - [\rho_n(k-1) - \phi_n(k-1)] \\ &= 2[I_n(k) - I_n(k-1)] + [\varepsilon_{n,\rho}(k) - \varepsilon_{n,\phi}(k) - \varepsilon_{n,\rho}(k-1) - \varepsilon_{n,\phi}(k-1)] \end{aligned} \quad (1)$$

Where,

dz: *raw divergence*; k: k^{th} *epoch*; n: n^{th} *satellite*;
ρ: *pseudorange*; ϕ: *carrier phase*; I: *ionospheric delay*;
ε_ρ: *pseudorange measurement noise*
ε_ϕ: *carrier phase measurement noise*

Assuming the multipath effect is tightly correlated in one epoch difference, the raw divergence reflects errors caused by ionospheric divergence and measurement noise terms, as shown in the second line of Equation (1). If the ionospheric bias term can be estimated, the raw divergence is dominated by contributions from thermal noise and RFI. The raw divergence statistic is useful for RFI detection because multiple independent measurements occur quickly, allowing for rapid detection. Independent samples of raw divergence are available once every two epochs for each channel. Assuming a LAAS or a JPALS reference station with at least three GPS receivers and an average of six satellites in view, there are as many as 18 independent samples available per epoch. Assuming Gaussian statistics, a combination of the independent samples gives a decision variable which approximately follows a χ^2 distribution. A monitor can be implemented based on a χ^2. Raw divergence statistics can also be combined with AGC statistics to improve performance.

2.2. *Obtaining Zero Mean Raw Divergence*. As described by the prior section, Equation (1) includes ionospheric delay terms which should be eliminated if the RFI and thermal noise levels are to be estimated. For the purposes of this work, nominal divergence (such as that which results from the change in ionosphere obliquity factor as GPS satellites rise through the sky) is assumed to change over an extended duration. Rapid ionosphere changes, such as scintillation effects, are assumed to be additional threats that the divergence metric would monitor. Scintillation effects are not discussed in this paper, however. In order to remove low-frequency ionospheric changes from the raw divergence statistic, a Geometric Moving Average (GMA) of the divergence is used (Xie, 2004).

$$d\bar{z}(k) = \frac{\tau(k) - T_S}{\tau(k)} d\bar{z}(k-1) + \frac{1}{\tau(k)} dz(k) \quad (2)$$

Where,

$d\bar{z}$:mean of raw divergence

$$\tau(k) = \begin{cases} kT_s, k < \tau_d / T_s \\ \tau_d, else \end{cases}, \quad \tau_d = 200 \text{ sec}, T_s = 0{\cdot}5 \text{ sec}$$

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

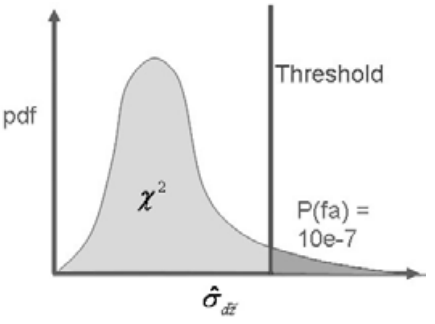

Figure 1. Sigma-estimation threshold.

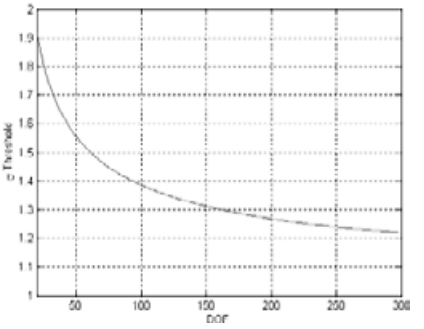

Figure 2. Sigma-estimation threshold vs. DOF.

Removal of the estimated mean from the raw divergence makes the zero-mean divergence ($d\tilde{z}$), which has only noise components.

$$d\tilde{z}(k) = dz(k) - d\bar{z}(k) \tag{3}$$

It is this zero-mean divergence statistic which will be employed for RFI monitoring.

3. ZERO-MEAN DIVERGENCE MONITORS. In this paper, two kinds of methods based on zero-mean divergence are used for RFI detection. They are 'Sigma-estimation' and 'CUSUM' methods, which have previously been employed for detecting small and large sigma violations, respectively, in the Stanford IMT (Pullen, 2003).

3.1. *Sigma-Estimation.* The sigma-estimation method (Lee, 2001) simply estimates the standard deviation of normalized raw divergence in real time.

$$\hat{\sigma}_{d\tilde{z}}(k) = \sqrt{\frac{1}{A-1}\sum_{i=k-A+1}^{k}\left[\frac{d\tilde{z}_i(i)}{f\sigma_{d\tilde{z}}(el(i))}\right]^2} \tag{4}$$

Where,

$\hat{\sigma}_{d\tilde{z}}$: estimated sigma
A: number of independent samples
$d\tilde{z}$: zero-mean divergence
$\sigma_{d\tilde{z}}$: standard deviation of normalized raw divergence in nominal condition
f: inflation factor
el: elevation angle

The threshold for this monitor is determined by the false alarm rate and by the number of available independent measurements, also called degrees of freedom. Figure 1 shows the probability density function of a χ^2 distribution to illustrate the sensitivity of sigma-estimation to false alarm rate. Here, the false alarm rate is 10^{-7}. In this case, the threshold is set as a value which makes the probability beyond the right tail equal to 10^{-7}. The sensitivity of the sigma-estimation method to the

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

degrees of freedom (DOF) of the χ^2 distribution is illustrated in Figure 2. Here the acronym DOF is used as an abbreviation that means "number of independent samples." If the monitor has more independent samples, it can detect smaller violations.

3.2. *CUSUM* (*Cumulative Sum*) In contrast with the sigma-estimation monitor, which examines the magnitude of divergence errors, a second monitor type, called CUSUM, examines the rate at which divergence errors accumulate. The sigma CUSUM takes squared normalized zero-mean divergence as the monitor input as shown in Equation (5). This input is normalized by the nominal standard deviation of the divergence, σ_0, which depends on the elevation angle of the corresponding satellite. Note that this monitor uses all the measurements from all reference receivers at one epoch, therefore M in Equation (5) is the number of all the measurements. The rate at which the CUSUM allows errors to accumulate is called the 'failure slope,' k_{sigma}. The failure slope term is a function of the nominal sigma, σ_0, and the lower-bound faulted sigma, σ_1. This faulted sigma is also known as the 'out-of-control' sigma. In this paper, the monitor uses a particular variant of CUSUM known as the FIR (Fast Initial Response) CUSUM, which has a "head-start" value of $3h/4$ for fast detection. The threshold h is determined to achieve a fault-free alert rate of 10^{-7}.

$$
\begin{aligned}
C(0) &= \frac{3h}{4} \\
C(k) &= \max\left(0,\ C(k-1) + \sum_{i=1}^{M}\left(\frac{d\tilde{z}_i(k)}{f\sigma_{d\tilde{z}}(el_i(k))}\right)^2 - M\cdot k_{sigma}\right) \\
k_{sigma} &= -\frac{\ln(\sigma_1) - \ln(\sigma_0)}{(2\sigma_1^2)^{-1} - (2\sigma_0^2)^{-1}}
\end{aligned}
\tag{5}
$$

Where,

f: inflation factor
M: number of measurements at k-th epoch
σ_0: nominal (in-control) sigma
σ_1: target (out-of-control) sigma
h: threshold

3.3. *Accounting For Non-Gaussian Statistics.* Both monitors assume Gaussian statistics for zero-mean divergence. In fact, divergence statistics typically display heavier than Gaussian tails (Braff, 2005). In order to ensure continuity, given this heavy tailed distribution, the actual divergence distribution can be modelled as a Gaussian with an inflated sigma. The monitor threshold is defined conservatively, in terms of this inflated sigma. As described in the following section, an appropriate value of the sigma-inflation, f, for the experiments described in this paper is 2·45, at low elevations below 20°, and 1·6 at higher elevations.

4. EXPERIMENTAL DATA.

4.1. *Experimental Data Collection.* The data used in this divergence monitor validation study was previously collected by Luo (Luo, 2003). The data set includes

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

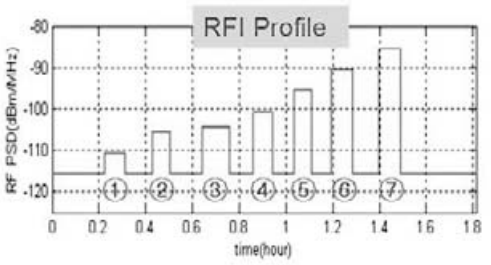

Figure 3. Injected RFI power profile.

two kinds of RFI experiments – a single-channel bench test and a three receiver live GPS test. In this paper, the saved data in those tests will be analyzed to show the raw divergence is an appropriate metric for RFI detection.

4.1.1. *Single-channel Bench Test Setup*. A single-receiver, single-satellite bench test demonstrated strong correlation between the divergence metric and RFI-induced errors. In this test, a GPS simulator produced a synthetic single-satellite signal. This simulated GPS signal was combined with an RFI signal and input to a NovAtel GPS receiver. The RFI generator was tuned to produce different types and power levels of RFI. For each type of RFI, the power level was increased by 1 dB every hour. Five types of RFI were injected into the received GPS signal including (1) wideband, (2) continuous wave at L1, (3) 10 MHz narrow band, (4) 30 MHz narrow band, and (5) pulsed continuous wave at L1 with a 5% duty cycle. These five RFI signals are abbreviated as WB, CW, NB10, NB30, and PULSE, respectively.

4.1.2. *Three-Receiver Live GPS Test Setup*. In order to assess the response time and sensitivity of the sigma-estimation monitor and the CUSUM monitor based on the divergence statistic, a multi-receiver experimental configuration was used. The configuration consisted of three NovAtel GPS receivers installed in the Stanford IMT. Two of them operated under nominal conditions, receiving live GPS signals from roof-mounted antennas. The input to the third receiver was combined with a signal from the RFI generator used in the prior experiment. The data from nominal condition receivers were processed for threshold determination. The data from the RFI-corrupted receiver were processed to evaluate monitor performance. In contrast with the single-channel test, which employed five types of RFI, the three-receiver test employed only a single type of RFI: wideband. The RFI generator was configured to produce a wideband RFI condition with a time varying power level, as shown in Figure 3. The nominal RFI power level is $-110{\cdot}5$ dBm/MHz. The test included seven periods of elevated RFI power, above the nominal level. Each of these seven periods of elevated power, numbered as shown in the figure, will be referred to as a 'section.' Each section lasted one hour, with data collected at a 2 Hz sample rate.

4.2. *Pre-Processing Of Experimental Data*. This section provides an overview of the pre-processing steps employed in both the single-channel and three-receiver tests. An overview of the analysis for each test provides a context for the pre-processing required in each case.

The single-channel test enables a comparison of the divergence statistic to AGC and C/N0 as estimates of RFI induced error. To assess the correlation between

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

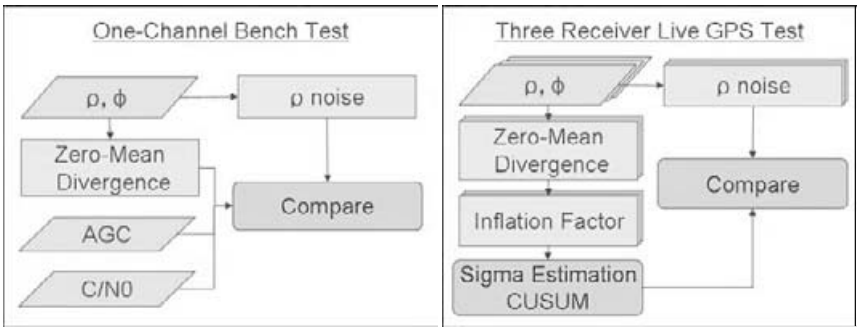

Figure 4. Data analysis overview.

these metrics and actual error levels, the raw divergence, AGC and C/N0 metrics are compared with an offline estimate of the pseudorange noise. The comparison is repeated under various types of excessive RFI, including continuous wave, narrowband and wideband threats. In each case, the raw data files include AGC and C/N0 statistics. Pre-processing is required, however, to transform the carrier and code measurements stored in the data file into zero-mean divergence and pseudorange noise estimate statistics.

The three-receiver live-GPS test enables an assessment of RFI detection performance possible using two monitor methods based on the zero-mean divergence statistic: 'Sigma-estimation' and 'CUSUM', which were described previously. The result is compared with the offline estimate of actual pseudorange noise to show that the monitor can successfully detect severe RFI threats within a reasonable time-to-alert. Again, pre-processing is required to generate the zero-mean divergence and psuedorange noise estimate. Additionally, pre-processing determines the inflation factor, f. This inflation factor is used in defining monitor thresholds in order to prevent frequent false alarms due to heavy tails in the divergence distribution.

Figure 4 summarizes these pre-processing steps for both the single-channel and the three-receiver tests. The procedures for determining the zero-mean divergence inflation factor and the offline estimate of pseudorange noise are described in the subsequent sections.

4.3. *Determining Heavy-Tail Inflation Factor*. As described previously, the divergence monitors presume Gaussian statistics. To provide a conservative bound for actual, heavy-tailed statistics, a sigma-inflation factor is used. The sigma-inflation factor is selected so that the inflated Gaussian provides a conservative approximation of the actual heavy-tailed error distribution.

The following figures illustrate the determination of the sigma-inflation factor for the data used in the experiment. The dots in Figure 5 depict the normalized probability density function for the zero-mean divergence measurements. The figure illustrates data for satellites whose elevation angles were between 10 and 20°. The line 1 illustrates the normal distribution with standard deviation equal to that of the raw divergence measurements. The other lines are the inflated distributions with inflation factors of 1·6, 2·0, and 2·45. Theoretically, if the divergence metric were Gaussian, line 1 would overlay the dots. However, in the real case, the theoretical

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

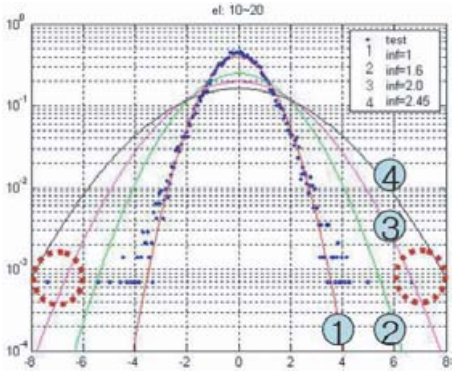

Figure 5. Inflation factor (El: 10 ~ 20°).

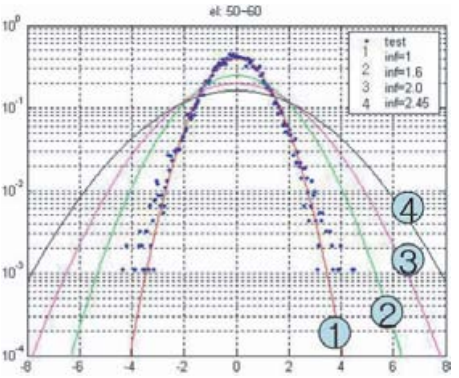

Figure 6. Inflation factor (El: 50 ~ 60°).

standard deviation does not bound all the measurements, especially near tail. For low elevation angles (10–20°), the inflation factor should be at least 2·45 (the widest line) to bound the measurements. For elevation angles above 20°, an inflation factor of 1·6 is sufficient to bound the data. As an illustration, Figure 6 plots the data in the elevation bin between 50 and 60°.

4.4. *Estimating Actual Pseudorange Noise*. In order to evaluate the relationship between the divergence metric and the actual error levels induced by RFI, it is necessary to provide a "truth" measurement that describes the RFI noise contribution to pseudorange measurements. For this work, noise contributions were estimated by an offline procedure. The offline procedure permits a batch analysis of the entire dataset available. The offline procedure bases its error estimate on a code-minus-carrier statistic. The primary error sources in the code-minus-carrier signal are the ionosphere, the carrier-phase integer ambiguity, the multipath error and the thermal/RFI noise contribution as shown in the third line of Equation (6).

$$\begin{aligned} \rho &= d + dR - b + I + T + B + M_\rho + \varepsilon_\rho \\ \phi &= d + dR - b - I + T + B + \lambda N + M_\phi + \varepsilon_\phi \\ \rho - \phi &= 21 - \lambda N + M_\rho - M_\phi + \varepsilon_\rho - \varepsilon_\phi \end{aligned} \tag{6}$$

Where,

d: true distance from satellite to receiver; dR: ephemeris error;
b: satellite clock bias; T: tropospheric delay; B; receiver clock bias;
M_ρ: pseudorange multipath: M_ϕ: carrier phase multipath;
λ: wavelength of L1 signal; N: integer ambiguity

For the single-channel bench test, ionosphere and multipath errors are not present. The offline pseudorange error estimate is established by subtracting the mean value from the code-minus-carrier statistic, in order to remove the integer ambiguity.

For the three-channel live GPS test, ionosphere divergence, integer ambiguity and multipath errors are present. In order to estimate the RFI-induced error component in the code-minus-carrier signal, the ionosphere delay, integer ambiguity and

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

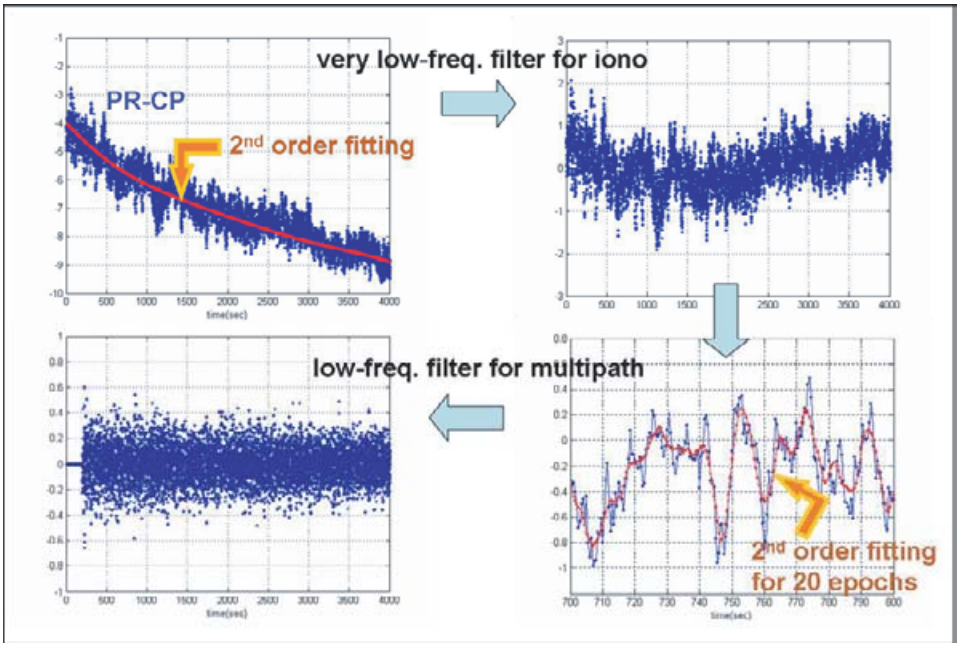

Figure 7. Estimating actual pseudorange noise.

multipath noise terms must be removed. As a first step, the ionospheric delay and integer ambiguity terms are eliminated by a second-order fitting for all epochs in the data set. In effect, this process acts as a very low frequency filter. As a second step, the instantaneous multipath is estimated by a second order fitting over a 10 second window. In effect, this process acts as a low frequency filter to eliminate multipath (since the time-constant for multipath is generally significantly longer than this 10 second window). The pseudorange error estimate consists of the remaining high-frequency noise. Figure 7 depicts this estimation process.

5. EXPERIMENTAL RESULTS.

5.1. *Test 1: Single-Channel Bench Test.* The single-channel bench test enables a comparison of the correlation between the estimated "true" value of the pseudorange noise and each of the possible monitor inputs – divergence, AGC and C/N0. Results show that the divergence metric, in contrast with the other test statistics, correlates strongly with pseudorange noise, independent of the nature of the RFI threat.

Plots of pre-processed data clearly show the correlation between each detection metric and the magnitude of the offline error estimate. The experimental data are plotted in Figures 8–10. Measurements were collected once every minute. The x-axis indicates the offline estimates for the pseudorange noise sigma, computed as the standard deviation of the error estimate over each one hour "section" during which the RFI noise level is constant. The y-axis represents values of the test statistic. The test statistic is the standard deviation of the divergence for Figure 8, the mean AGC for Figure 9, and the mean C/N0 for Figure 10. Note that the standard deviation of

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

the divergence should get larger as the pseudorange noise increases while the mean AGC and the mean C/N0 should get smaller.

In order to establish correlation, the test statistic estimates should vary linearly with the offline pseudorange error estimate. Ideally, the correlation should become increasingly apparent as more independent samples are used. This convergence to a linear relationship occurs because the increased number of independent samples reduces random measurement noise. If, however, the correlation is weak, the detection statistic will not display linear variation with the error estimate, even when the number of independent samples is high.

Figures 8 through 10 indicate strong correlation for the divergence metric but not for the AGC or C/N0 metrics. For each figure, plots are shown for two cases with varying independent samples used in computing the detection metric (DOF of 5 and 30). The number of estimates plotted is inversely proportional to the DOF, since data was compiled for each RFI case over a 60 minute interval.

In general, an integrity monitoring system does not have prior knowledge of the type of RFI that might corrupt its measurements. For this reason, it is necessary to consider the correlation between the detection metric and the estimated error across a wide range of RFI conditions (wideband, continuous wave, narrowband and pulsed). For each individual RFI condition, AGC and C/N0 metrics display a roughly linear trend as a function of estimated error. As shown in Figures 9 and 10, however, the correlation breaks down when AGC and C/N0 data for a variety of different RFI conditions are lumped together.

By comparison, the divergence metric is proportional to the estimated error regardless of the classification of the RFI source. This linear trend is clearly evident in the 30 DOF plot of Figure 8. The slope of the proportionality for the line is approximately $\sqrt{2}$. This slope is explained by noting that the divergence error is dominated by the difference of two pseudoranges and that the sigma for a difference of two independent, identically distributed Gaussian random variables is $\sqrt{2}$ larger than the sigma of either individual random variable.

The strong linearity of the divergence metric data, in comparison to the AGC and C/N0 data, is further emphasized by computing the residuals of the data relative to a linear regression. The relationship between the linear regression residuals and the DOF is shown for each metric in Figure 11. The regression is computed over a data set that combines all RFI types evaluated in the single-channel test. While the mean AGC and the mean C/N0 residuals are not reduced as DOF increases, the standard deviation of the divergence does converge toward zero as DOF increases. This data clearly demonstrates the direct linear relationship between the divergence metric and the estimated error, independent of the type of RFI present. This is a highly desirable characteristic for the RFI monitor statistic.

5.2. *Test 2: Three-Receiver Live GPS Test.* The three-receiver live GPS test enables performance assessment of RFI monitors based on the zero-mean divergence statistic. In this test, wideband RFI of varying power levels is injected into the live GPS signal. The sigma-estimation and CUSUM monitors process this data to detect integrity violations caused by the injected RFI.

5.2.1. *RFI Mask.* Not all RFI is considered harmful to the operation of a differential augmentation system. In fact, systems like LAAS and JPALS are designed to operate with full integrity even under slightly elevated RFI conditions. The maximum level of RFI in which an augmentation can safely operate is called

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

the RFI mask. The threshold for RFI monitoring must be defined in terms of the RFI mask. Previous research efforts in RFI monitoring and detection for GPS augmentation systems have not focused, however, on the challenge of defining an appropriate monitoring threshold based on the RFI mask. For instance, Luo showed that an AGC monitor could be designed to detect all of the elevated RFI sections in Figure 3 and that a C/N0 monitor detected sections 5 and 6 (Luo, 2003). This analysis focused only on whether RFI conditions were more severe than nominal. For operational system integrity, it would be necessary to extend this analysis to distinguish between RFI conditions more severe than nominal but nonetheless below the RFI mask. To reflect the impact of the mask on an operational system, this work derives the thresholds for the divergence-based sigma-estimation and CUSUM monitors using the RFI mask.

Augmentation systems automatically account for the noise level at the RFI mask in their broadcast error bounds. In LAAS, for example, the RFI mask is embedded in the broadcast parameter, σ_{pr_gnd} which represents the error level of the pseudorange correction measured at the ground station. If the true correction error is not bounded by σ_{pr_gnd}, an integrity violation occurs. The LAAS σ_{pr_gnd} values are generally described as a function of elevation using Ground Accuracy Designator (GAD) curves. GAD curves represent the performance of the reference receivers and antennas at three levels of quality, labelled A, B and C. The Stanford IMT antennas, for example, support a GAD-B curve for σ_{pr_gnd}. For analysis of IMT data, the level of RFI noise at the RFI mask can thus be derived from the GAD-B standard.

The RFI noise contribution is only one component of the GAD-B curve. The GAD-B curve, modelled by Equation (7), has both a RFI/thermal noise contribution and a multipath contribution as shown in Table 1 (McGraw, 2000).

$$\sigma_{PR}(\theta) = a_0 + a_1 e^{-\theta/\theta_c} \tag{7}$$

The RFI noise level defined in the GAD-B model assumes 100 second smoothing for data acquired at 2 Hz. In order to compare this smoothed pseudorange sigma at the RFI mask to the estimated actual pseudorange noise sigma, the smoothed values must be scaled by a factor of approximately 19, because the smoothed pseudorange sigma is about 20 times that of a raw pseudorange sigma (McGraw, 2000) and the estimated actual pseudorange sigma is about 1·05 times of that due to the low-frequency filter. When the estimated actual pseudorange data from the IMT are normalized by the scaled sigma at the GAD-B RFI mask, the data can be plotted as shown in Figure 12 (dots). For nominal operating conditions, the normalized distribution is tighter than the GAD-B level representing the noise at the RFI mask. In the figure, the GAD-B limiting error is plotted as a Gaussian curve with unity (normalized) sigma.

As the RFI power level increases, the pseudorange error distribution eventually grows wider than the GAD-B error bound which is about three times wider than the nominal distribution. Therefore, the target out-of-control sigma of the divergence monitors is set to 3·0. Figure 13 compares the pseudorange error distribution to the bounding GAD-B curve for the four different power levels associated with sections $3 \sim 6$. The third and the fourth sections' pseudorange error distribution are bounded, despite RFI injection, so the monitor need not issue an alert to protect integrity.

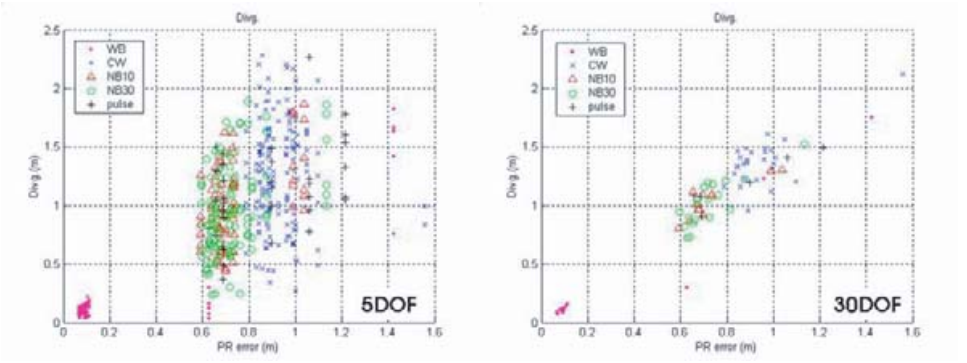

Figure 8. STD of raw divergence vs. STD of PR.

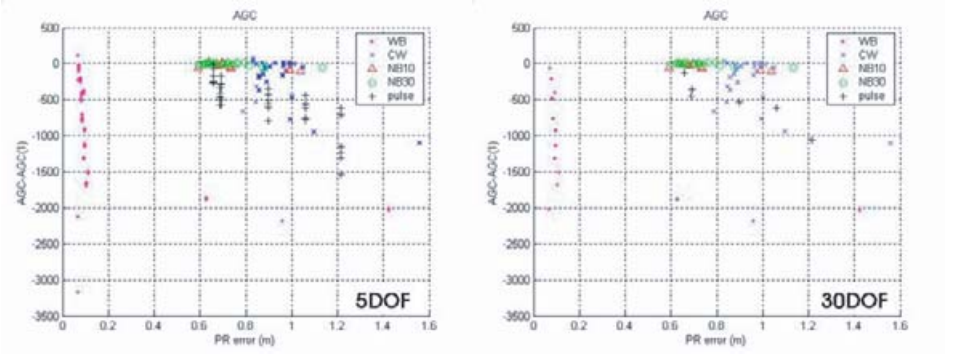

Figure 9. Mean of AGC vs. STD of PR noise.

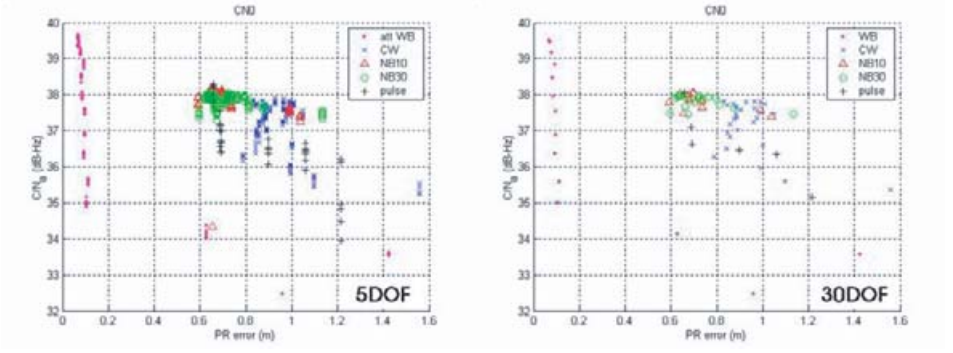

Figure 10. Mean of C/N0 vs. STD of PR noise.

In sections 5 and 6, however, the GAD-B curve does not bound the pseudorange noise. Consequently, RFI in the sections 5 and 6 must be detected. Therefore the set of "target" sections, for which RFI exceeds the allowed level, consists of sections 5 and 6 (Figure 14). Section 7 is discounted, since the receiver did not track the GPS signals in this section. Because, from the actual pseudorange noise estimation, the estimated noise has an independent sample every 21 epochs (see 4.4), we

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

Table 1. GAD-B curve model.

Contribution	a_0 (m)	a_1 (m)	θ_C (deg)
Multipath	0·15	1·06	15
RFI/noise	0·029	0·081	16·4
Total	0·16	1·07	15·5

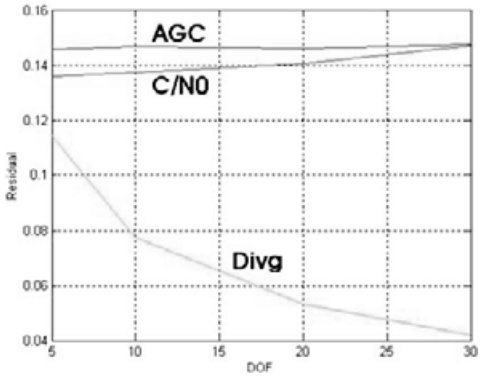

Figure 11. Linear regression residuals for each metric.

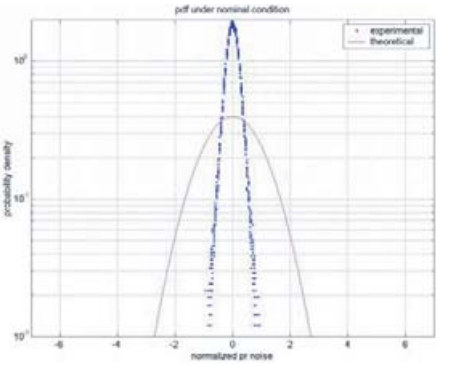

Figure 12. PDF of PR noise under nominal condition.

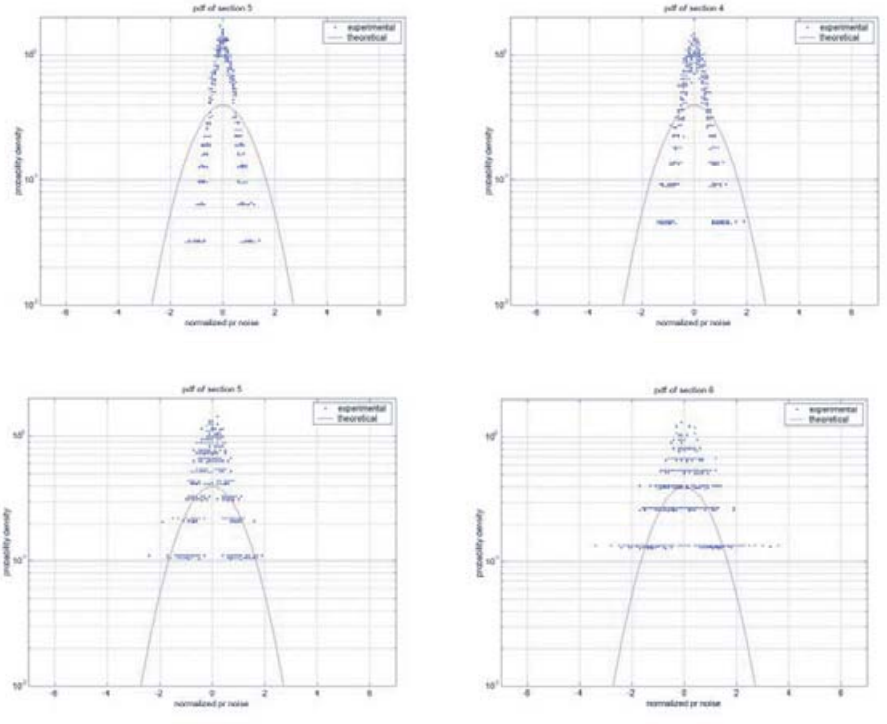

Figure 13. PDF of PR noise in each section (3,4,5,6).

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

can have 20 different probability density functions (PDF) for a section. The plots in Figure 13 show the multiple PDFs for each section together. The two monitor methods, sigma-estimation and CUSUM, need only detect RFI in the target sections. Within these sections, moreover, it is desirable that the monitors respond with a rapid time-to-alert.

5.2.2. *Sigma-Estimation Result*. The sigma-estimation monitor can be designed either to respond rapidly to major anomalies, or to detect minor anomalies with a slower alert time. The key parameter that establishes the tradeoff between detection sensitivity and detection time is the number of epochs (DOFs) used in the sigma-estimation. Figure 15 represents the relationship between sigma-estimation DOFs and detectable out-of-control sigmas when the false alarms rate is 10^{-7} and a probability of missed detections is 10^{-3}. Because the target out-of-control sigma is set to 3·0 in the previous paragraph, the sigma-estimation monitor needs at least 25 independent samples to detect the RFI effect over the mask.

A larger number of DOFs slows the monitor response, but also improves sensitivity by lowering the monitor threshold (given a fixed probability of false alarms, 10^{-7}. See Figure 2). Five test cases of varying DOF were investigated in this study: 25, 50, 100, 150 and 200.

Figures 16–18 display the results of this comparison. In each figure, the left plot depicts the estimated sigma (dot) and the threshold (line). The right upper plot shows the alarm flag issued by the monitor. The right lower plot depicts the injected RFI power profile. This last plot indicates the times at which the excessive RFI signal switched on and off.

Because the sigma-estimation collects only as many DOFs as there are visible satellites at any given epoch, the time to first detection is delayed as the number of DOFs grows. This time to first detection is called 'detection time'. When the RFI condition disappears, this latency effect delays the switching of the monitor flag back to its neutral (zero) state. In this paper, the elapsed time between the end of the section and time when the monitor flag returns to zero is referred to as 'reset time'. A long reset time is less harmful than a long detection time, since the reset time only impacts system availability, whereas the detection time impacts integrity.

The 25-DOF sigma-estimation detects the sigma violation in sections 5 and 6 within 7 and 2 seconds, respectively. The reset times for the sections are 0 and 1 second in this case. In section 5, because the estimated sigma is noisy and not large enough to exceed the threshold consistently, the monitor does not remain flagged throughout the entire section. After the first detection in section 5, the sigma-estimation monitor with 25 DOFs flags only 15·56% of the time, as compared to 93·56% in section 6.

In order to increase the likelihood of detecting the RFI event in section 5, the threshold must be decreased by considering a larger number of DOFs. Accordingly, the number of DOFs was changed to 50, 100, 150 and 200. Figure 17 and 18 are the results of 100 and 200 DOFs cases. As the DOFs increase, the estimated sigma becomes better resolved and the detection probability for test section 5 increases to 53·33%, 93·23%, 98·85%, and 100% as shown in Figure 19. However with a larger number of DOFs, the monitor needs more time to collect independent samples, so the detection and the reset times grow.

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

Figure 20 summarizes the response-time results for the different DOF cases. The upper two lines are the reset times versus DOF and the lower two are the detection times. The figure implies that the detection time is proportional to the DOF. Section 5 has lower RFI power, therefore it has longer detection time and shorter reset time than section 6. The figure suggests that a sigma-estimation monitor based on the divergence statistic can detect the target sections within 20 seconds. Because this test used only one GPS receiver under RFI conditions, a reference station that had three or four reference receivers might detect RFI about three or four times faster than these results.

5.2.3. *CUSUM Result*. In statistical process control, the CUSUM monitor is frequently used because it tends to respond to faults more quickly than the sigma-estimation monitor. A significant parameter in the design of a CUSUM monitor is the out-of-control sigma, σ_1, which is set to 3·0 in the test, or 3 times of a nominal standard deviation. To compare the performance of the CUSUM monitor with the sigma-estimation, the 25 DOF result was selected because in Figure 2, a 25 DOF sigma-estimation monitor can detect about 3·0 times of a nominal standard deviation. The left plot of Figure 21 depicts the calculated CUSUM values and thresholds, which shows the CUSUM values quickly exceed their thresholds in sections 5 and 6. The detection times in sections 5 and 6 are 4 seconds and 1 second, respectively. As such, this monitor detects the violation much faster than the sigma-estimation monitor whose detection times were 7 and 2 seconds. Again, the detection times could get shorter if the monitor used three or four reference receivers.

In the CUSUM method, the CUSUM value, C, is reset to the start value ($3h/4$ in this case) after it exceeds the threshold. Thus, there is no delay associated with reset because the CUSUM monitor resets the output value to the start value immediately following detection. Also, the fraction of time the monitor remains flagged after the first detection (as illustrated for the sigma-estimation method in Figure 19) is not relevant for the CUSUM monitor. The CUSUM monitor would remain flagged 100% of the time if it were not reset to the start value following each detection.

6. CONCLUSION AND FUTURE WORK.

This paper developed divergence-based monitors that directly detect the RFI contribution to the navigation error. A single-channel bench test showed that the zero-mean divergence metric is directly proportional to high-frequency pseudorange errors. For the divergence metric, unlike AGC or C/N_0, this proportionality is insensitive to RFI type. A three-receiver live GPS test showed that monitors based on the zero-mean divergence metric can detect RFI threats within a reasonable time. Two monitors were tested: a sigma-estimation method and a CUSUM method. Of the two monitors, the CUSUM method detects hazardous RFI much faster than the sigma-estimation method.

Future work will integrate the new divergence-based monitor with previously proposed RFI detection methods to provide more robust ground-based integrity monitoring. The AGC monitor is of particular interest because of its fast response. This fast response means that the AGC monitor may play a useful role in detecting RFI for an augmentation system. However, because the AGC does not provide

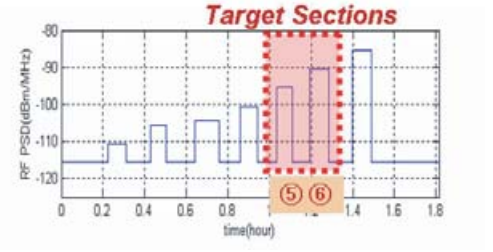

Figure 14. Target sections for RFI detection.

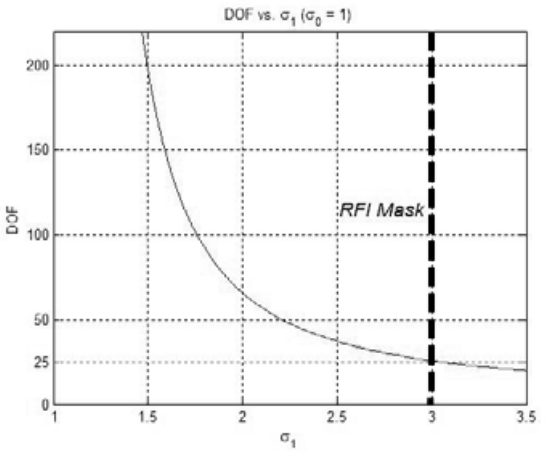

Figure 15. Sigma-estimation DOF vs. Detectable out-of-control sigma.

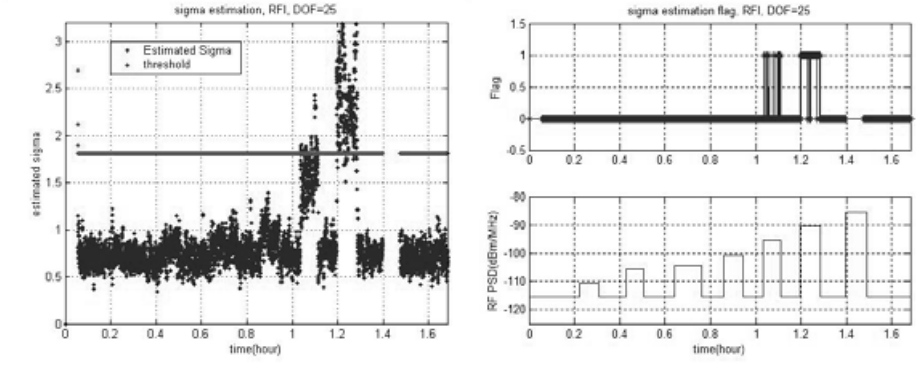

Figure 16. Sigma-estimation result (DOF = 25).

enough information to determine the navigation error level caused by an arbitrary RFI threat, the divergence-based monitor serves as a useful complement.

Figure 22 illustrates a recommended strategy to deal with RFI threats using these complementary monitors. The proposed strategy uses a combination of AGC and zero-mean divergence monitors for detection of RFI above the mask

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

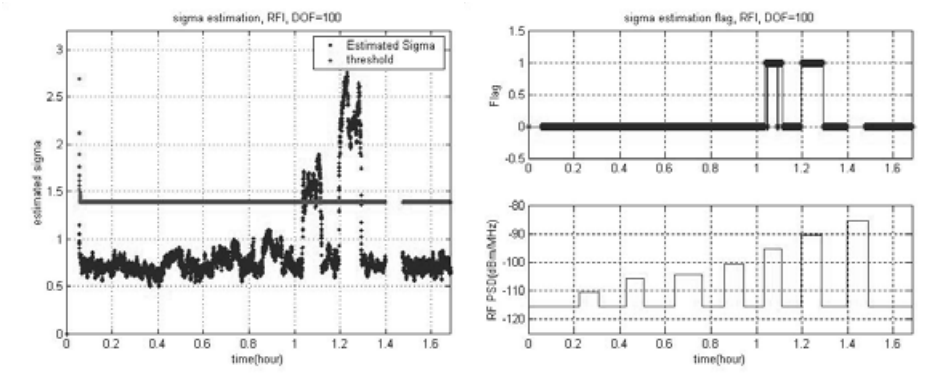

Figure 17. Sigma-estimation result (DOF = 100).

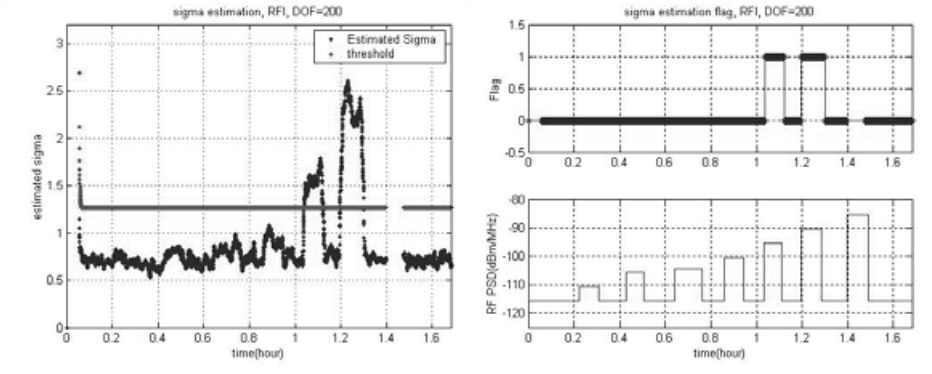

Figure 18. Sigma-estimation result (DOF = 200).

level. The divergence monitor, most likely a CUSUM monitor on account of the CUSUM's rapid performance, determines whether the system is operational or not. If the system needs to keep operating in spite of the RFI threat, it may continue to operate normally by adjusting the integrity monitor thresholds and the broadcast sigma based on the RFI level detected by the divergence-based monitor. By incorporating both AGC and divergence-based RFI monitoring, a civilian differential augmentation system, such as LAAS, can provide better availability and continuity performance under low power RFI conditions, since the system will only shut itself down under truly threatening high-RFI conditions. Likewise, a military system, such as JPALS, could justify continued operations under elevated RFI conditions without loss of integrity by broadcasting scaled sigma error bounds.

ACKNOWLEDGEMENT

This research was supplied by Brain Korea 21 project. And the authors are thankful for Godwin Zhang's help for tests. Parts of the paper were presented at the Institute of Navigation 61st Annual Meeting in Cambridge, Massachusetts, USA, 27–29 June 2005.

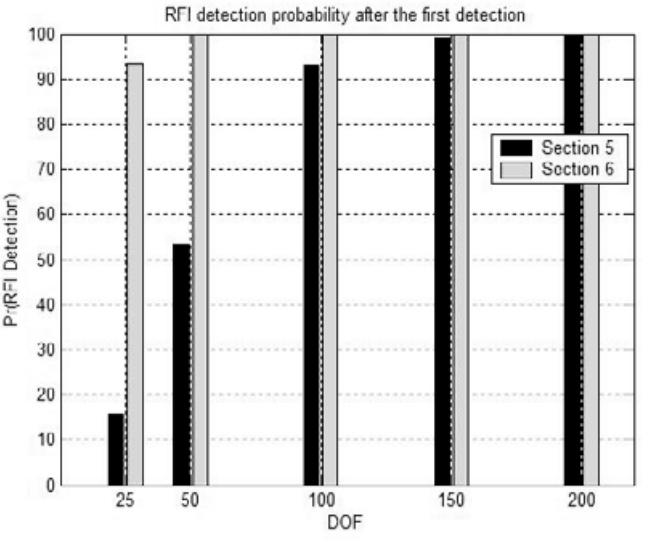

Figure 19. Pr (RFI detection after the first detection).

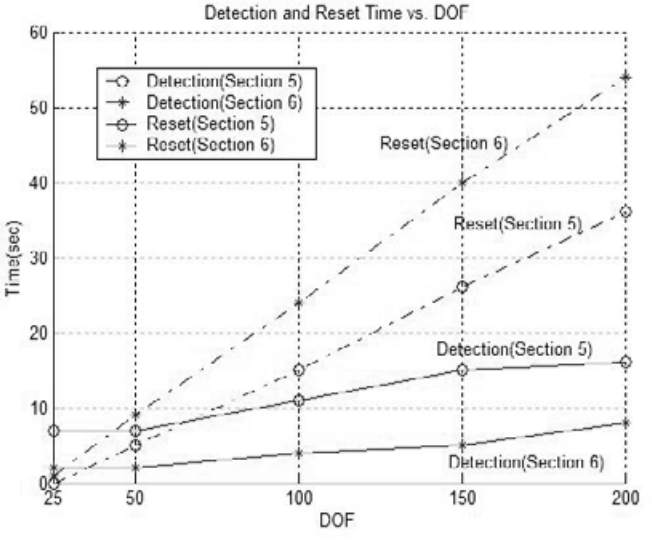

Figure 20. Detection and Reset time.

REFERENCES

Bastide, F., Akos, D., Macabiau, C. and Roturier, B. (2003). Automatic Gain Control (AGC) as an Interference Assessment Tool. *Proceedings of the Institute Of Navigation GPS 2003*, Portland, OR.

Braff, R. and Shively, C. (2005). A Method of Over Bounding Ground Based Augmentation System (GBAS) Heavy Tail Error Distributions. *The Journal of Navigation*, **58**, 83–103.

Lee, J., Pullen, S., Xie, G. and Enge, P. (2001). LAAS Sigma-Mean Monitor Analysis and Failure-Test Verification. *Proceedings of the Institute Of Navigation 57TH Annual Meeting*, Albuquerque, NM.777

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press

Luo, M., Pullen, S., Zhang, J., Gleason, S., Xie, G., Yang, J., Akos, D. and Enge, P. (2000). Development and Testing of the Stanford LAAS Ground Facility Prototype. *Proceedings of the Institute Of Navigation National Technical Meeting*, Anaheim, CA.

Luo, M., Xie, G., Akos, D., Pullen, S. and Enge, P. (2003). Radio Frequency Interference Validation Testing for LAAS using the Stanford Integrity Monitor Testbed. *Proceedings of the Institute Of Navigation National Technical Meeting*, Anaheim, CA.

McGraw, G. A., Murphy, T., Brenner, M., Pullen, S. and Van Dierendonck, A. J. (2000). Development of the LAAS Accuracy Models. *Proceedings of the Institute Of Navigation GPS 2000*, Salt Lake City, UT.

Ndili, A. and Enge, P. (1998). GPS Receiver Autonomous Interference Detection. *Proceedings of the Institute of Electrical and Electronics Engineers Position Location and Navigation Symposium '98*, Palm Springs, CA.

Normark, P. L., Xie, G., Akos, D., Pullen, S., Luo, M., Lee, J. and Enge, P. (2001). The Next Generation Integrity Monitor Testbed (IMT) for Ground System Development and Validation Testing. *Proceedings of the Institute Of Navigation GPS 2001*, Salt Lake City, UT.

Pullen, S., Luo, M., Gleason, S., Xie, G., Lee, J., Akos, D., Enge, P. and Pervan, B. (2000). GBAS Validation Methodology and Test Results from the Stanford LAAS Integrity Monitor Testbed. *Proceedings of the Institute Of Navigation GPS 2000*, Salt Lake City, UT.

Pullen, S., Lee, J., Xie, G. and Enge, P. (2003). CUSUM-Based Real-Time Risk Metrics for Augmented GPS and GNSS. *Proceedings of the Institute Of Navigation GPS 2003*, Portland, OR.

Xie, G., Pullen, S., Luo, M., Normark, P. L., Akos, D., Lee, J. and Enge, P. (2001). Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed. *Proceedings of the Institute Of Navigation 57TH Annual Meeting*, Albuquerque, NM.

Xie, G. (2004). Optimal On-Airport Monitoring of the Integrity of GPS-based Landing Systems. *A Stanford University Ph.D. Dissertation*.

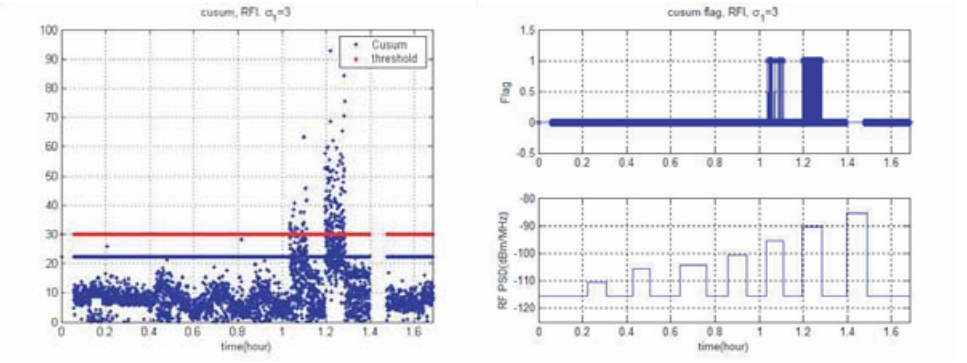

Figure 21. CUSUM Result (target out-of-control sigma = 3·0)

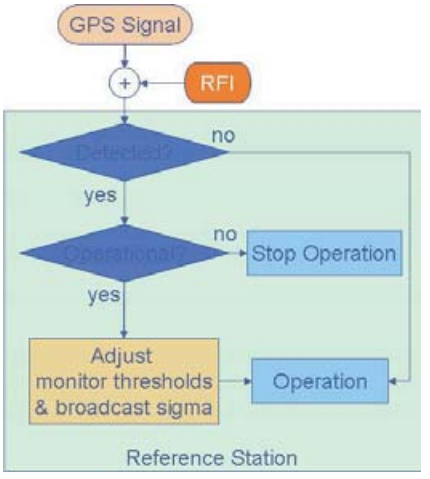

Figure 22. RFI detection strategy.

https://doi.org/10.1017/S0373463306003882 Published online by Cambridge University Press