Sequential test for signal-level integrity monitoring in GNSS receivers

Juan M. Parro-Jiménez[1,2], Rigas T. Ioannides[2], Massimo Crisci[2] and José A. López-Salcedo[1]

[1]Universitat Autònoma de Barcelona (UAB), Bellaterra (Barcelona), Spain
[2]European Space Agency, Noordwijk, The Netherlands
Email: {juan.parro, rigas.ioannides, massimo.crisci}@esa.int; jose.salcedo@uab.es

Abstract—**This paper addresses the problem of integrity monitoring in global navigation satellite system (GNSS) receivers. A new technique is proposed to quickly detect the presence of corrupted measurements caused by either multipath or spoofing threats. To do so, the symmetry of the correlation curve is permanently monitored, and a CUSUM-based sequential test is applied later on. Numerical results are provided to validate the technique and show its effectiveness.**

I. INTRODUCTION

The use of Global Navigation Satellite System (GNSS) receivers has spread in the past years due to the increasing demand of autonomous positioning and location-based services (LBS). Nowadays, GNSS is used as the main tool for localisation in many applications, among which are airplane landing, maritime, railway or road guidance, just to mention a few. The criticality of these activities requires an alert to be raised whenever the integrity of the navigation solution cannot be guaranteed, and this often has to be done in less than 2 seconds [1]. In this context, the concept of *integrity* refers to the level of trust that the user can expect from the results provided by the receiver. In the recent years, integrity has become of great importance due to the increasing demand of reliable navigation for critical applications, and the emergence of commercial services where a given service license agreement (SLA) needs to be guaranteed (e.g. road tolling). Currently, there are different approaches to measure the level of integrity such as the use of Ground or Satellite Based Augmentation Systems (GBAS or SBAS) or the implementation of Receiver Autonomous Integrity Monitoring (RAIM) techniques. However, other alternatives are necessary for standalone receivers, where the need of low complexity techniques is a usual requirement. In this sense, signal-level techniques are a good solution for extracting the information present at the physical-layer, thus providing a faster integrity response without having to resort to the output observables or the navigation solution.

One of the phenomena that can degrade the quality of the navigation signals is the well-known multipath effect. It is caused by the reception of multiple reflections of the authentic signal coming from objects surrounding the user receiver. Multipath represents one of the major sources of error in GNSS introducing errors in the pseudorange and carrier phase measurements, which in turn produce errors in the computation of the Position, Velocity and Time (PVT) solution. Together with multipath, another concerning threat for reliable satellite navigation is the transmission of spoofing signals, i.e. the deliberate replication of GNSS signals. This kind of intentional interference is only feasible in open GNSS where all the information regarding the signals is public, and thus the signal can be easily replicated [2].

In the presence of multipath or spoofing, the effects at the signal level share important similarities. Mainly, the two phenomena produce a distortion of the correlation peak between the received signal and the local replica, and the expected triangular shape of the main correlation lobe is no longer symmetric. This fact was exploited by the authors in [3], [4] for introducing so-called Slope Asymmetry Metric (SAM), with the objective of detecting the presence of multipath and spoofing. By exploiting the geometric properties of the correlation peak, a faulty signal can be easily detected, and the user has the possibility to discard the satellite and avoid the propagation of errors in the measurements to the PVT solution. Nevertheless, the problem of fault *detection* was actually not addressed in previous works on the SAM metric, where the emphasis was placed, instead, on the preliminary characterization of this metric, and the performance validation in realistic working conditions. The purpose of the present work is thus to fill this missing gap. To so do, we will propose a SAM-based sequential test capable of detecting sudden changes in the expected behaviour of the correlation shape. The underlying principle of this test is based on the CUSUM algorithm [5], which is very well-suited for quickest detection problems (as it is the case in integrity monitoring applications) and can be easily implemented with a very low complexity.

The paper is structured as follows: in the Section II, the optimal configuration of the SAM and its statistical behaviour are presented. Later in Section III, the fundamentals of sequential testing are presented and the use of the CUSUM algorithm for detection of integrity anomalies is proposed. Finally, Section IV presents different simulation results in order to show the performance of the proposed method and relevant conclusions about the presented work are drawn in Section V.

II. Slope Asymmetry Metric

As previously introduced, the SAM metric exploits the triangular shape that can be observed in the central part of the autocorrelation function of most GNSS signals. In particular, this metric compares the symmetry of the left and right slopes on both sides around the maximum correlation peak. Ideally, both slopes should be equal but with opposite sign, and thus their sum should equal zero. Based on this observation, the first step for calculating this metric consists on fitting each side of the correlation peak by a straight line. To do so, let us first represent the received signal correlation curve at n-th code period and time instant t, by the following signal model:

$$z_n(t) = \sqrt{\text{SNR}}\exp\{\theta_a\}R(t) + \eta_n(t), \quad (1)$$

where SNR and θ_a are the Signal-to-Noise-Ratio and phase of the authentic signal, $R(\tau)$ is the correlation function of the spreading code, and η_n is the correlated noise component, which has a complex Gaussian distribution with zero mean and variance $\sigma_\eta^2 = 1$.

In order to compute the SAM metric, we will consider a set of L correlation samples on each side of the received signal correlation curve, taken at time instants τ_k equi-spaced by T_Δ. Since we are focusing on GPS L1 C/A signals, we will consider that $LT_\Delta \leq T_c$, with T_c the chip period. The correlation samples from the left and right slopes will be stacked into the $(L \times 1)$ vectors $\mathbf{z}_{l,n} \doteq [z_n(\tau_1), z_n(\tau_2), \ldots, z_n(\tau_L)]^T$ and $\mathbf{z}_{r,n} \doteq [z_n(\tau_{-1}), z_n(\tau_{-2}), \ldots, z_n(\tau_{-L})]^T$, respectively. For any two values of $z_n(\tau_k)$ their covariance is given by:

$$\text{cov}\{z_n(\tau_i), z_n(\tau_j)\} = \sigma_\eta^2 R(\tau_i - \tau_j) \quad (2)$$

where it can be seen that samples with $|\tau_i - \tau_j| < T_c$ turn out to be correlated, and this will have an impact on the behaviour of the SAM metric.

Using the correlation samples in $\mathbf{z}_{l,n}$ and $\mathbf{z}_{r,n}$, we will obtain the least-squares (LS) estimate of the straight line that best fits these measurements. The LS estimates for the slope $\hat{a}$ and offset $\hat{b}$ of each line are obtained as $[\hat{a}, \hat{b}]^T = \mathbf{M}^\#\mathbf{z}$, where $\mathbf{M}^\#$ represents the Moore-Penrose inverse of matrix $\mathbf{M} = (\mathbf{T} \quad \mathbf{1})$ being $\mathbf{T}$ the column vector containing the correlation lags of each correlator value and $\mathbf{1}$ a column vector of L ones. Using either $\mathbf{z}_n = \{\mathbf{z}_{l,n}, \mathbf{z}_{r,n}\}$, we obtain the estimates of the slopes for the left line, $\hat{a}_l(n)$, and for the right line, $\hat{a}_r(n)$, respectively. Then, the SAM metric becomes:

$$\text{SAM}(n) \doteq \hat{a}_l(n) + \hat{a}_r(n). \quad (3)$$

Note that, if any symmetry distortion is produced in the received signal correlation, the metric will exhibit $\mu_{\text{SAM}} \doteq \text{E}[\text{SAM}(n)] \neq 0$. This effect will be used to detect any possible degradation that multipath or spoofing can cause to the received signals. In the presence of either these two threats, the correlation output will be affected by an additional replica according to:

$$\begin{aligned} z_n^{a+r}(\tau_k) &= \sqrt{\text{SNR}}\exp\{\theta_a\}R(\tau_k) \\ &+ \sqrt{\frac{\text{SNR}}{D/U}}\exp\{j\Delta\theta_r\}R(\tau_k - \Delta\tau_r) + \eta_n, (4) \end{aligned}$$

where the subscript $^{a+r}$ indicates *authentic* plus *replica* signals. The power of these two signals is related through $D/U \doteq P_a/P_r$ and $\Delta\theta_r$ and $\Delta\tau_r$ represent the relative phase and code delay of the replica with respect to the authentic/original signal.

A. Optimal configuration for the estimation of the slopes

The estimation of the correlation slopes will determine the performance of the SAM metric for detecting the replica. In order to find optimal configuration, the spacing between correlators T_Δ and their number L will be studied by comparing the performance in terms of noise and integrity detection capabilities. The target is to make the SAM metric very sensitive to anomalies affecting the region close to the prompt correlator, since these are the ones introducing larger errors in the time-delay estimation of the received signal, and subsequently, in the resulting pseudorange. For this reason we would like to use correlators close to the prompt one, when computing the metric.

Figure 1 shows the SAM envelope for three sets of correlators in the range from -0.5 to 0.5 chips. In particular, one set with $L = 3$ correlation samples given by $\mathcal{S}_1 : \tau_{|k|} = \{0.1, 0.2, 0.3\}$, and two different sets with $L = 4$ given by $\mathcal{S}_2 : \tau_{|k|} = \{0, 0.1, 0.2, 0.3\}$ and $\mathcal{S}_3 : \tau_{|k|} = \{0.2, 0.3, 0.4, 0.5\}$. A replica signal with a relative power $D/U = 6$ dB has been added, whose relative delay $\Delta\tau_r$ is varied from -0.5 to 0.5 chips.
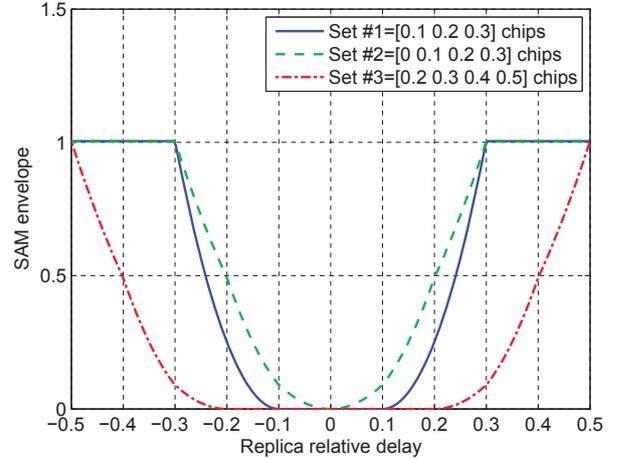
Fig. 1. SAM envelope as a function of the relative delay of a replica signal (i.e. multipath/spoofing) with $D/U = 6$ dB.

We can observe how those sets excluding the prompt correlator (i.e. $\mathcal{S}_1$ and $\mathcal{S}_3$) become very insensitive to the presence of replica signals closely aligned with the authentic one. This effect is totally undesired since the detection of anomalies close to prompt is essential. This blind effect is caused by the fact that both estimated slopes $\hat{a}_l$ and $\hat{a}_r$ are affected by the same bias but opposite sign. Thus, the resulting SAM does not reflect any asymmetry in the signal. In contrast, the set $\mathcal{S}_2$ that uses the prompt correlator, reflects the asymmetry in the signal even for short replica delays. For this reason, the prompt

correlator should always be included in the computation of each slope.

For selecting the number of correlators, two elements will be evaluated, namely the SAM envelope, which indicates the level of degradation suffered by the correlation peak in the presence of a replica, and the variance of the SAM metric. Note that a low variance is desired since that would make easier the detection of an anomaly, reducing at the same time the rate of false alarms.

In order to carry out the estimation of a slope, a minimum of $L = 2$ correlators is needed per side. This configuration has the advantage that requires few correlators, which is always desirable for the sake of reducing the complexity of the receiver processing. More correlators can be added as inputs to the linear regression problem. This would improve the estimation of the slopes if the correlators were uncorrelated. However, for the case under study, pairs of correlators at instants τ_i and τ_j are actually correlated through $R(\tau_i - \tau_j)$.

Figure 2 shows the SAM envelope for different configurations using $\tau_{|k|} = kT_\Delta$ for $k = 0, \ldots, L - 1$ and different values of T_Δ. Black lines represent the asymmetry error when using $L > 2$ correlators per side spaced $T_\Delta = 0.1$ chips respectively. The coloured lines are the SAM envelopes for those cases in which $L = 2$ correlators are selected for each slope. The results show that there is no significant improvement in terms of bias in the peak symmetry when more correlators are added. This is caused by the high correlation between the values used in the regression problem. Consequently, the decision made regarding the number of correlators is that $L = 2$ should be used for computing each slope.

that will determine the detection performance of the metric is its variance. For the specific case of $L = 2$, the variance of the SAM can be calculated from the statistical behaviour of correlators using the propagation of uncertainty [6]. Being the metric SAM a function of the correlators $z(\tau)$, we can obtain its variance with $\sigma_{\mathrm{SAM}}^2 = \mathbf{A} \boldsymbol{\Sigma} \mathbf{A}^{\mathbf{T}}$, where the matrix $\mathbf{A}$ is the Jacobian matrix formed from the partial derivatives of the SAM with respect to each of the correlators and $\boldsymbol{\Sigma}$ is the covariance matrix of the correlators, which can be built from (2). Note that these two elements describe how the input variables (correlators) are related to each other and with the metric SAM. For $L = 2$ the variance of the SAM yields:

$$
\begin{aligned}
\sigma_{\mathrm{SAM}}^2 \;=\; & \frac{2}{T_\Delta^2}(2 - R(\tau_{-2} - \tau_{-1}) + R(\tau_{-2} - \tau_1) \\
& - \; R(\tau_{-2} - \tau_2) - R(\tau_{-1} - \tau_1) \\
& + \; R(\tau_{-1} - \tau_2) - R(\tau_2 - \tau_1)).
\end{aligned}
\tag{5}
$$

Replacing the correlators at τ_{-1} and τ_1 with the prompt correlator (i.e. $\tau_{-1} = \tau_1 = 0$), we can evaluate the variance of the metric for different values of the furthest correlator delay, in this case $\tau_{-2} = -T_\Delta$ and $\tau_2 = T_\Delta$. Figure 3 shows that the variance of the SAM metric decreases inversely quadratically to the spacing between correlators T_Δ. Thus, with the objective of reducing the variability of the metric the two correlators of each side should be as spaced as possible. On the other hand, as shown previously, increasing the space between the two correlators of each side reduces the detection capabilities for short delay replicas. Therefore, the choice of T_Δ will determine the trade-off between the mean value of the SAM metric and its variance.

Fig. 2. SAM envelope as a function of the relative delay of a replica signal (i.e. multipath/spoofing) with $D/U = 6$ dB.
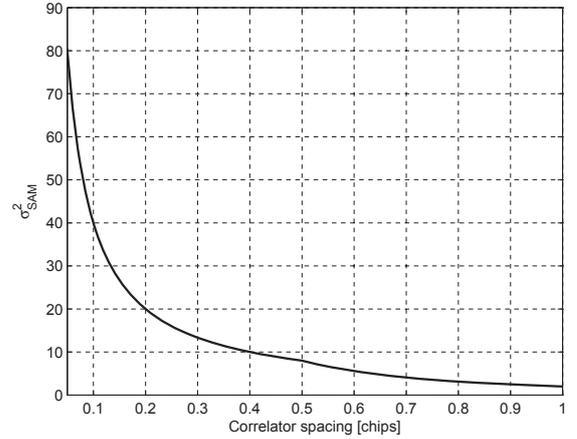
Fig. 3. Evolution of the SAM variance as a function of T_Δ for $L = 2$ and $\tau_{-1} = \tau_1 = 0$.

From all the possible configurations that involve the use of $L = 2$, the best results for detecting close replicas are obtained when the spacing between correlator T_Δ is small. In these cases the level of asymmetry is more significant for replica delays close to $\Delta\tau_r = 0$ but becomes smaller when the correlators are further spaced.

Together with the SAM envelope, another figure of merit

B. SAM statistical characterisation

The statiscal behaviour of the metric will be studied in two different situations: first, when only the authentic signal is present, and second, when a replica of the signal is present. In each case, we will use (1) and (4) to represent the correlator outputs.

For the case of $L = 2$ and $\tau_{-1} = \tau_1 = 0$, we have a total of three correlators distributed along the correlation function. Based on (1), the correlator $z_n(t)$ is a random complex value with a Gaussian distribution according to:

$$z_n(t) \sim \mathcal{N}(\sqrt{\text{SNR}} \exp\{j\theta_a\} R(\tau), \sigma_\eta^2). \tag{6}$$

The estimation of the slopes $\hat{a}$ is also normally distributed since it is a linear combination of the correlators $z_n(t)$. Besides, since they have the same mean but opposite sign, the resulting SAM has zero mean. The metric SAM is thus distributed as:

$$\text{SAM}(n) \sim \mathcal{N}(0, \sigma_{\text{SAM}}^2). \tag{7}$$

In order to carry out a detection process, and since the SAM metric will be complex valued, we take here the squared absolute value of the SAM. As a result, we obtain a random variable that is distributed according to a chi-squared distribution with two degrees of freedom and mean and variance:

$$\text{E}\left[|\text{SAM}(n)|^2\right] = \sigma_{\text{SAM}}^2, \tag{8}$$
$$\text{var}\left[|\text{SAM}(n)|^2\right] = (\sigma_{\text{SAM}}^2)^2. \tag{9}$$

In the hypothetical case that a replica is present in the incoming signal, the correlator values are no longer determined by (1) since the presence of the replica has to be taken into account. In this scenario, the obtained SAM maintains the variance of (5) but its mean value is different from zero. More specifically, it is a value that depends on the power, the phase and the delay of the replica according to:

$$\mu_{\text{SAM}} = \sqrt{\frac{\text{SNR}}{D/U}} \frac{\exp\{j\Delta\theta_r\}}{T_s} \beta, \tag{10}$$

where $\beta \doteq R(\tau_{-1} - \Delta\tau_r) - R(\tau_{-2} - \Delta\tau_r) + R(\tau_2 - \Delta\tau_r) - R(\tau_1 - \Delta\tau_r)$. As can be seen, if the relative delay of the replica is $\Delta\tau_r = 0$, the mean value of the SAM is also equal to 0 because the two signals are aligned in time and no deformation is observed on the final peak. The variance in this scenario is equal to (5).

In this case, when the squared absolute value of the SAM is calculated, we obtain a non-central chi-squared distributed random variable due to the bias present in (10), and mean and variance:

$$\text{E}\left[|\text{SAM}(n)|^2\right] = |\mu_{\text{SAM}}|^2 + \sigma_{\text{SAM}}^2, \tag{11}$$
$$\text{var}\left[|\text{SAM}(n)|^2\right] = 2|\mu_{\text{SAM}}|^2 \sigma_{\text{SAM}}^2 + (\sigma_{\text{SAM}}^2)^2. \tag{12}$$

Note that in the case that $\mu_{\text{SAM}} = 0$, the resulting distribution matches the one obtained in the absence of replica.

The following section will use the statistical analysis presented here for implementating a sequential detection test.

III. SEQUENTIAL PROBABILITY RATIO TESTS

A. Fundamentals

A Sequential Probability Ratio Test (SPRT) is a procedure for quickly detecting "change" events in a signal under analysis. The concept of Sequential Analysis was introduced by A. Wald in the 40's [7]. The benefit of SPRT with respect

to classical hypothesis testing, such as the Likelihood Ratio Test (LRT) is that, in general, it requires an expected number of observations considerably smaller. This means that the sequential test is able to made a decision at an earlier stage than other sorts of testing. This property of the SPRT has attracted the attention of researchers from the field of finance and medicine. In the field of GNSS, some contributions can be found mostly for integrity applications at the observable level [8]. A decade after Wald presented his work on SPRT, Page [5] introduced the Cumulative Sum (CUSUM) algorithm as a specific case of SPRT, which is nowadays widely used due to its simplicity. The work of Page was later extended by others such as [9]–[11].

A general SPRT is defined by the decision rule d that is taken at every stopping time T, i.e. every time the observation made exceeds some predefined upper and lower values A and B. For a given stopping time T, the decision rule indicates whether the null ($\mathcal{H}_0$) or the alternate hypothesis ($\mathcal{H}_1$) are accepted, or whether the confidence reached so far is not enough, and thus a new observation is needed before a decision can be made. In this way, the test is ran in a sequential manner until a decision is made.

In the case of the CUSUM algorithm it is assumed that, initially (i.e. under $\mathcal{H}_0$), the random process y under study follows some statistical distribution $p_0(y)$. Thus, the problem becomes whether or not the observation y changes at some instant t_0, and follows a new distribution $p_1(y)$, corresponding to the alternate hypothesis $\mathcal{H}_1$. The detection problem can thus be expressed as:

$$\begin{cases} \mathcal{H}_0: & y_n \sim p_0(y), \quad \forall n \\ \mathcal{H}_1: & y_n \sim p_0(y), \quad n < t_0, \\ & y_n \sim p_1(y), \quad n \geq t_0. \end{cases} \tag{13}$$

The statistic used in the algorithm is built from the Log-Likelihood Ratio $\text{LLR} = \log \frac{p_1(y)}{p_0(y)}$. The stopping time in this case can be written as:

$$T \doteq \min(n : g(n) \geq h), \tag{14}$$
$$g(n) = [g(n-1) + \text{LLR}(n)]^+, \tag{15}$$

where $[x]^+ \doteq \max(0, x)$ and h is the threshold that guarantees that the time between false alarms ($\overline{T}_0$) is beyond a given false alarm rate α:

$$\overline{T}_0 \geq e^h = \frac{1}{\alpha}. \tag{16}$$

This threshold can be derived from the inequalities of Wald [7] for the selection of the thresholds in a SPRT. From this result, Lorden [10] demonstrated the optimality of the CUSUM algorithm for reducing the worst mean delay for detection $\overline{T}_1$.

One of the drawbacks of the CUSUM algorithm is that its optimality is subject to detect exactly the change for which is designed. Meaning that smaller or larger changes in the actual model may cause the test to perform far from optimal. This lack of optimality is shown in Figure 4 where a CUSUM test is carried out to detect a change in the mean of a process, while the variance σ^2 is maintained. For this test, a value of $\alpha = 0.05$ is selected. It can be seen in the upper figure that

when the ratio between the selected mean μ_{tuned} and the actual mean of the change μ_{real} is different from 1, the CUSUM needs significantly more time to detect the change. For what concerns the false alarm rate, even though the obtained value is below the configured α, an increment of the number of false alarms is observed when the selected mean is far from the actual one.
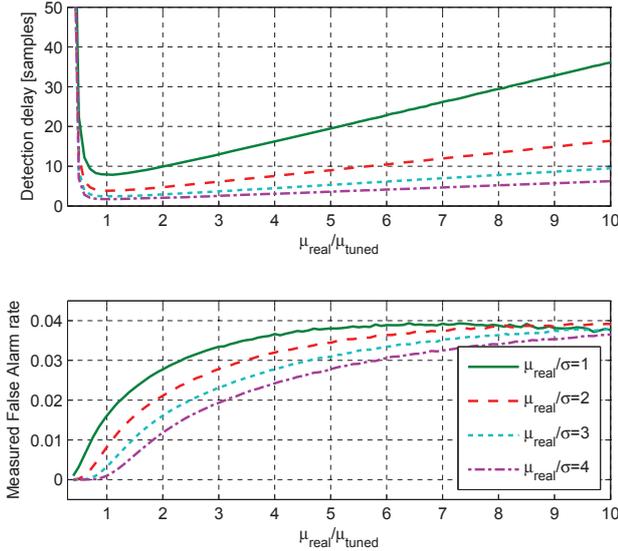
Fig. 4. Evolution of the detection delay and the false alarm rate of the proposed algorithm for different values of real and tuned mean.

B. Sequential test for detection of changes in SAM

The statistical behaviour of the SAM presented in Section II-B will help us in implementing an integrity monitoring technique based on the CUSUM algorithm. As previously seen, the squaring operation on the SAM metric forces us to work with a chi-squared distributed random variable, instead of a conventional Gaussian one (as expected by the CUSUM algorithm). A way to circumvent this problem is to take advantage of the Central Limit Theorem, and let N_{acc} observations to be averaged before using them for detection purposes. By doing so, we know that if N_{acc} is large enough ($N_{\text{acc}} > 50$ according to [12, p. 46]), the resulting averaged metric:

$$\overline{\text{SAM}}(m) \doteq \frac{1}{N_{\text{acc}}} \sum_{n=1}^{N_{\text{acc}}} |\text{SAM}(n - mN_{\text{acc}})|^2 \quad (17)$$

can be approximated by a normal distribution whose mean coincides with the mean of the averaged variable, and whose variance is N_{acc} times lower. Therefore, the resulting distribution in the absence of threats is:

$$p_0(\overline{\text{SAM}}) = \mathcal{N}(\sigma_{\text{SAM}}^2, (\sigma_{\text{SAM}}^2)^2/N_{\text{acc}}), \quad (18)$$

while the presence of a threat modifies the distribution as:

$$p_1(\overline{\text{SAM}}) = $$
$$\mathcal{N}(|\mu_{\text{SAM}}|^2 + \sigma_{\text{SAM}}^2, (2|\mu_{\text{SAM}}|^2\sigma_{\text{SAM}}^2 + (\sigma_{\text{SAM}}^2)^2)/N_{\text{acc}}). \quad (19)$$

In other words, the detection procedure consists on a change in the mean and the variance of a normal distribution. For this case the LLR can be easily derived resulting in:

$$\text{LLR}(\overline{\text{SAM}}(m)) = $$
$$\log\left(\frac{\sigma_0}{\sigma_1}\right) + \frac{(\overline{\text{SAM}}(m) - \mu_0)^2}{2\sigma_0^2} - \frac{(\overline{\text{SAM}}(m) - \mu_1)^2}{2\sigma_1^2}. \quad (20)$$

where μ and σ represent the mean and standard deviation of the probability density functions $p_0(\overline{\text{SAM}})$ and $p_1(\overline{\text{SAM}})$.

IV. SIMULATION RESULTS

In order to assess the performance of the proposed sequential test for the correlation asymmetry detection, we will simulate different GPS L1 C/A signal scenarios. For this evaluation, the correlation values $z_n(t)$ are obtained with an integration time of $T_i = 1$ ms and the squared magnitudes of the SAM metric are accumulated during $N_{\text{acc}} = 100$ samples. The number of correlators to estimate each slope is set to $L = 2$ and a correlator spacing of $T_\Delta = 0.2$ chips is chosen. This decision is based on the detection capabilities and the noise level of the metric as discussed in Section II.

A. Performance of the test based on expected multipath

For this scenario an expected replica of $D/U = 12$ at $\Delta\tau_r = 0.2$ chips is selected. This expected replica power, from now on D/U_{tuned}, is a design parameter that we set in order to detect replicas with larger power values. Since replicas with $D/U < D/U_{\text{tuned}}$ cause a greater deformation in the correlation peak, the obtained μ_{SAM} will also increase, according to (10), hence reducing the detection interval. This effect can be observed in Figure 5 where the mean detection delay is shown for different D/U and different values of C/N_0. The results confirm that, even though the experiment is design to detect a replica of $D/U_{\text{tuned}} = 12$ dB, greater replicas are also detected with even shorter delay. For observed $D/U > 12$ dB, the algorithm detects the change only in rare occasions. This fact can be used to design a replica detection algorithm with the ability of ignoring low power replicas. For instance, if a receiver is design to work in an environment with an expected D/U occasioned by multipath, the tuned value in the CUSUM could be chosen to ignore any replica lower from that value. However, in this case, the algorithm will not work under optimal conditions since the specified change in the mean will not match, in general, the observed one. As explained in the previous section, the optimality of the CUSUM algorithm is subject to the correct modelling of the change. Meaning that for selected D/U_{tuned} the detection of any different replica power will take more time to be detected than if the algorithm was designed exactly to detect that change. This effect can be observed in Figure 6 where for different observed D/U we can see the additional time required to detect due to the

misalignment between the design power and the observed one. Note that the optimal value in each case corresponds to the case in which $D/U_{tuned} = D/U_{real}$. The additional detection delay is a drawback of the proposed algorithm for designing the CUSUM method for a wide number of cases instead of for each specific situation. However, as long as the application allows a certain margin in the detection delay, the designer could tune the algorithm to take into account this extra delay.
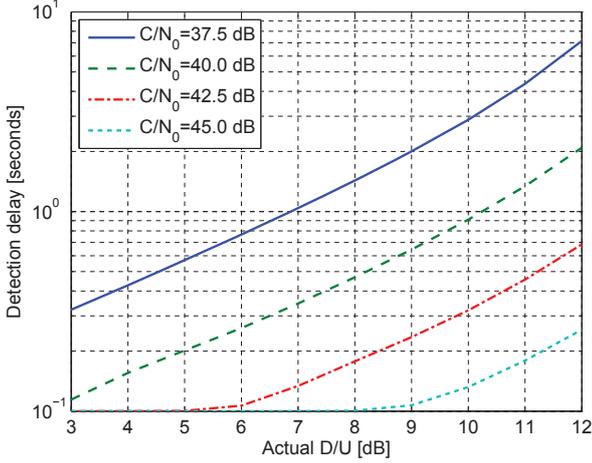
Fig. 5. Mean detection delay for different observed values of D/U and power level C/N_0 when the algorithm is designed for detecting D/U_{tuned}.
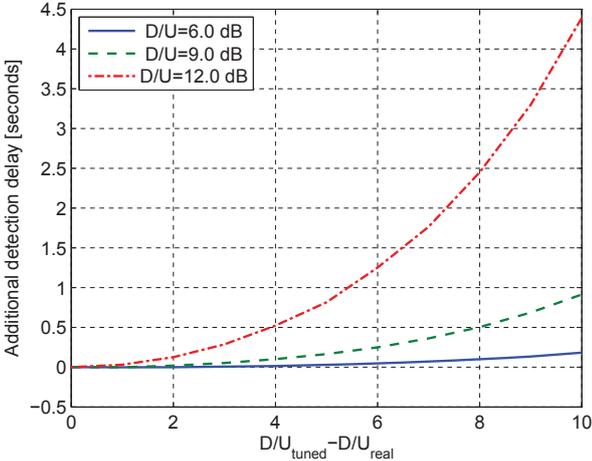
Fig. 6. Additional detection delay due to the misalignment between the observed and the tuned D/U.

B. Effect of the chosen false alarm rate

In this case we analyse the false alarm rate α that has to be specified in the algorithm and which determines the threshold according to equation (16). Note that the obtained threshold will act as an upper bound as demonstrated in Figure 4 where the obtained false alarm never exceeded the selected α. For

this experiment we tune the algorithm to detect a replica of $D/U = 12$ dB and we evaluate the mean detection delay for a range of pre-set false alarm rates α between 1E-5 and 2.5E-3. Figure 7 shows how the mean detection time increases as the rate of false alarms is decreased. In addition, larger detection intervals are required to obtain the same false alarm rates when the power level of the observed signal decreases.
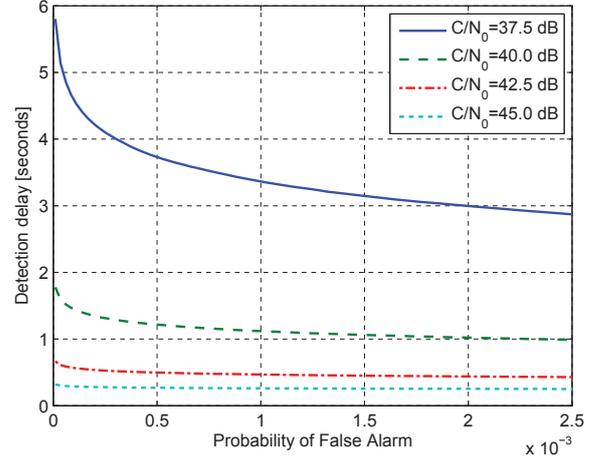
Fig. 7. Evolution of the obtained mean detection delay with the selected false alarm rate.

As explained in the previous section the pre-set false alarm rate α defines the threshold that guarantees the obtained probability of false alarm to be always below α. For the specific case of $D/U = 12$ dB, Figure 8 shows how the observed false alarm increases with α for different C/N_0. Note that rate of false alarms in the experiment grows at a lower rate than the pre-set value α. This means that during the tuning of the proposed algorithm, the designer should take into account that the value α is not the desired false alarm rate but an upper threshold.
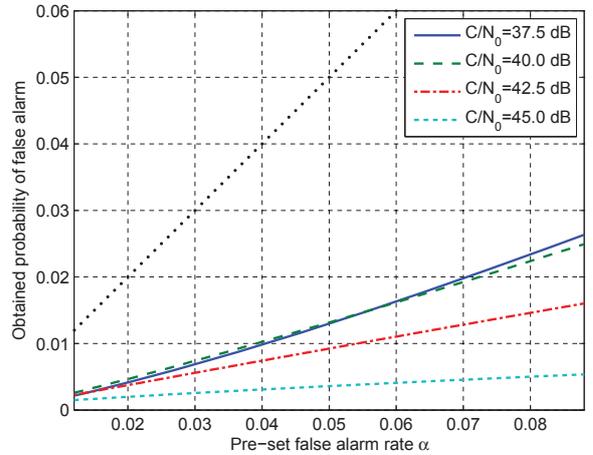
Fig. 8. Obtained false alarm rate as a function of the set upper threshold for the false alarm.

V. Conclusions

The use of of the CUSUM algorithm has been investigated in this paper with the objective of a quick detection of anomalies in the expected shape of satellite navigation signals. The use of this sort of sequential test allows the GNSS receiver to easily and quickly alert the user about the presence of a signal replica. This kind of signal can be present in the received signal either unintentionally, due to environmental reflections of the LOS signal, or intentionality due to a spoofing attack. The presence of a replica is measured with a metric that reflects the asymmetry of the expected correlation peak and which is monitored with the change detection algorithm. An exhaustive analysis of the behaviour of this metric has been presented here which allows to carry out a mechanism for detecting the moment in which the replica appears. The simulation results presented allow to understand how the proposed algorithm should be tuned to exploit its optimal properties in terms of detection delay and false alarm rate. Moreover, these results allow to design the algorithm to be tolerant to certain ranges of replica power at the expense of a slower detection of the actual replica.

References

[1] S. Gleason and D. Gebre-Egziabher, *GNSS Applications and Methods*, ser. GNSS technology and applications series. Artech House, Incorporated, 2009.

[2] R. Di, S. Peng, S. Taylor, and Y. Morton, "A USRP-based GNSS and interference signal generator and playback system," in *Position Location and Navigation Symposium (PLANS), 2012 IEEE/ION*, april 2012, pp. 470 –478.

[3] J. A. Lopez-Salcedo, J. M. Parro-Jimenez, and G. Seco-Granados, "Multipath detection metrics and attenuation analysis using a GPS snapshot receiver in harsh environments," in *3rd European Conference on Antennas and Propagation, 2009. EuCAP 2009.*, march 2009, pp. 3692 –3696.

[4] J. M. Parro-Jimenez, J. A. Lopez-Salcedo, R. T. Ioannides, and M. Crisci, "Signal-level integrity monitoring metric for robust GNSS receivers," in *31st AIAA International Communications Satellite Systems Conference*, 2013.

[5] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. pp. 100–115, 1954.

[6] H. H. Ku, "Notes on the use of propagation of error formulas," *Journal of Research of the National Bureau of Standards. Section C: Engineering and Instrumentation*, vol. 70C, no. 4, pp. 263–273, Oct. 1966.

[7] A. Wald, "Sequential tests of statistical hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117–186, 06 1945.

[8] I. V. Nikiforov, "Sequential, FSS and snapshot approaches to GPS/DGPS integrity monitoring," in *Proceedings of the 10th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 1997) September 16 - 19, 1997 Kansas City, MO*, 1997.

[9] A. Shiryaev, "The problem of the most rapid detection of a disturbance in a stationary process." *Soviet Math. Dokl.*, vol. 2, pp. 795–799, 1961.

[10] G. Lorden, "Procedures for reacting to a change in distribution," *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, 12 1971.

[11] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *The Annals of Statistics*, vol. 14, no. 4, pp. 1379–1387, 12 1986.

[12] G. E. P. Box, W. G. Hunter, and J. S. Hunter, *Statistics for Experimenters*. New York, NY: Wiley, 1978.