Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed

Gang Xie, Sam Pullen, Ming Luo, Per-Ludvig Normark, Dennis Akos, Jiyun Lee, and Per Enge
Department of Aeronautics and Astronautics, Stanford University

Boris Pervan
Department of Mechanical, Materials, and Aerospace Engineering, Illinois Institute of Technology

ABSTRACT

The Stanford University Integrity Monitor Testbed (IMT) is a prototype of the Local Area Augmentation System (LAAS) Ground Facility (LGF) that is used to evaluate whether the LGF can meet the integrity requirements that apply to Category I aircraft precision approach. Because of the complexity of the monitoring algorithms, it is necessary to show that these requirements are met under a variety of possible failure situations.

This paper explains the integrity monitor algorithms and fault-handling logic implemented in the IMT and reports the most recent nominal and failure test results with upgraded hardware. A significant fraction of the IMT code is dedicated to processing alert messages generated by different IMT monitor algorithms. Once these monitors begin flagging questionable measurements, several steps of logical reasoning and trial removals are required to determine which failed system elements are the source of the problem. This LGF function is known as "Executive Monitoring" (EXM). This paper describes in detail the EXM procedures implemented in the IMT and their role in meeting the LGF integrity and continuity requirements.

For failure testing, two methods are used to inject failures into the IMT. One is to program a WelNavigate 40-channel GPS signal simulator to generate failed RF signals, which are then fed into the IMT. The other is to modify stored IMT receiver measurements collected under nominal conditions to inject failures after the fact and then have the IMT post-process these packets. Several recent results of IMT failure testing are reported.

The results of these tests demonstrate that the IMT is capable of detecting and excluding failures in compliance with the LGF integrity requirements in the current LAAS Ground Facility Specification. Improved algorithms are under study to meet the much tighter requirements for Category III precision landings.

1.0 INTRODUCTION

The U.S. Federal Aviation Administration (FAA) is developing the Local Area Augmentation System (LAAS) to support aircraft precision approach. This local-area differential GPS ground-based system places the responsibility for detecting and alarming space-segment and ground-segment failures on the LAAS Ground Facility (LGF), which is also responsible for generating and broadcasting carrier-smoothed code differential conditions and approach-path information to user aircraft [1]. The LGF must insure that all ranging sources for which LAAS corrections are broadcast are safe to use. If a failure occurs that threatens user safety, the LGF must detect and alert users (by not broadcasting corrections for the affected ranging source) within a 3-second time-to-alarm for Category I precision approaches.

The LGF must apply several different types of monitoring algorithms to detect a varied array of possible failures. In order to coordinate the LGF response to detected failures (some of which may trigger more than one monitoring algorithm), complex failure-handling logic must be included in the LGF. Stanford University researchers have developed an LGF prototype known as the Integrity Monitor Testbed (IMT) that includes a comprehensive set of monitoring algorithms and Executive Monitoring (EXM) logic to isolate failed measurements and reintroduce these measurements after the failure is clearly determined to be over.

To show that the requirements in the LGF Specification are achievable, it is important to verify that the IMT works as designed and meets the LGF integrity requirements. The LGF specification defines six ranging source failure

classes (see Section 3.2.1.2.7.3.1 of [1]). Two methods are used to inject failures into the IMT for failure testing. Failed RF signals can be generated by a WelNavigate 40-channel GPS signal simulator and then fed into the IMT. The second method, generally more flexible than the first one, is to inject failures in software by modifying stored IMT reference receiver packets (from measurements taken under nominal conditions) to reflect the desired failures and let the IMT post-process the altered measurements.

This paper gives an overall picture of the IMT system architecture. It explains integrity monitor algorithm components with their nominal test results and describes the fault-handling logic (Executive Monitoring, or EXM) implemented in the IMT. The most recent results of IMT failure testing are also reported. Finally, the plans for the next phase of IMT development and testing are discussed. Note that an updated version of the LGF Specification for Non-Federal procurement has just been released [16]. The IMT is designed to satisfy the more specific integrity requirements of the previous specification and, with minor exceptions (several new algorithms are now being implemented), it meets the revised specification as well.

2.0 SUMMARY OF IMT SYSTEM ARCHITECTURE

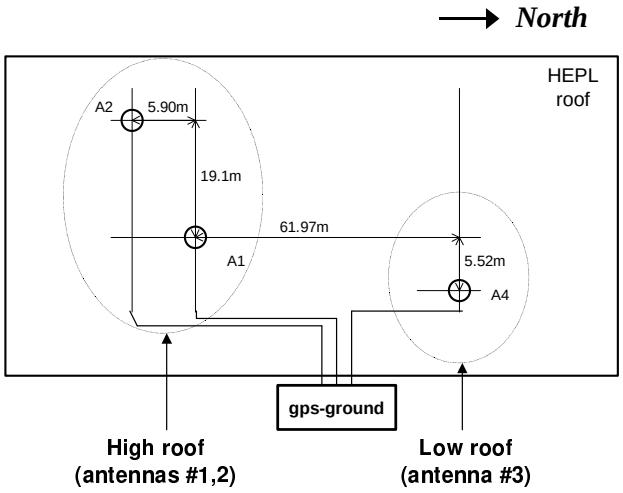

Figure 1: IMT Hardware Configuration

The LAAS Ground Facility requires redundant DGPS reference receivers to be able to detect and exclude failures of individual receivers. Figure 1 shows the configuration of the three IMT antennas on the Stanford HEPL laboratory rooftop. Although limited by the space available on the rooftop, these antennas are sufficiently separated to minimize correlation between individual receiver multipath errors [5]. Recently, the IMT has been updated with new GPS reference receivers and antennas. Three NovAtel OEM 4 receivers are now connected to three separated GPS Pinwheel antennas. Each receiver can track as many as 12 satellites simultaneously. The receivers sample GPS signals and provide receiver measurements, or packets, every 0.5 seconds, and these measurements are fed into the IMT at that rate.

Figure 2 shows the IMT's functional blocks. The dataflow and chronology of the IMT processing are shown in Figure 3. The RF conditioning and Signal-in-Space Receive and Decode (SISRAD) functions provide the pseudorange measurements, carrier-phase measurements, and navigation data messages that are the core of IMT processing and enable the generation of differential corrections.

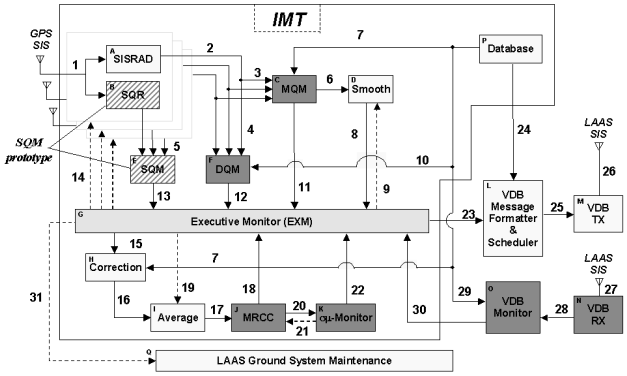

Figure 2: IMT Functional Flow Diagram

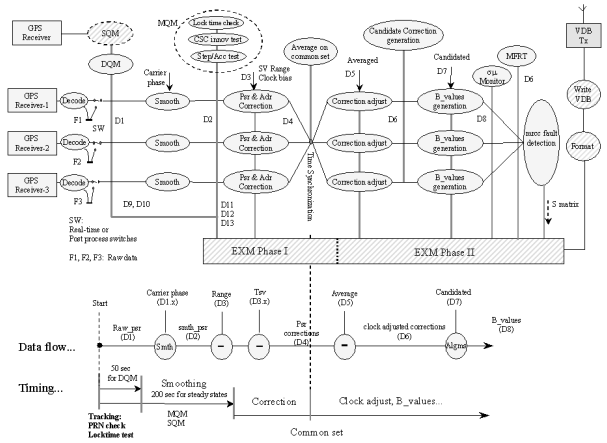

Figure 3: IMT Processing Chronology

The IMT processing core is divided into three parts: nominal processing (carrier smoothing and calculation of differential corrections), integrity monitor algorithms, and executive monitoring (EXM). The monitor algorithms are grouped into SQM, DQM, MQM, MRCC, σμ–monitor, and MFRT classes (these are defined below). They are targeted at detecting a wide range of possible failures in the GPS Signal in Space (SIS) or in the LGF itself. EXM

consists of two phases of logic reasoning. It determines, isolates and removes error sources as needed to meet the integrity requirements while minimizing false exclusions (threatening continuity) to the degree possible. In addition to the IMT processing core, the LGF also contains other essential components, such as VHF Data Broadcast (VDB), through which the LAAS corrections will be broadcast to users.

Note there is a switch after the Decode function in Figure 3. It represents the concept presented in Figure 4. When Point 1 is closed, GPS receiver packets are directly fed into the IMT processing. This is referred to as "real-time" IMT operation. If Point 2 is closed, then the packets can be saved into a non-volatile storage device. Depending on whether Point A or B is closed, the saved data packets can be used to perform nominal or failure testing. This is called "post-processing" IMT operation. The post-processing version of IMT is very important and useful in debugging and testing, since one fixed dataset of receiver packets can be processed repeatedly.

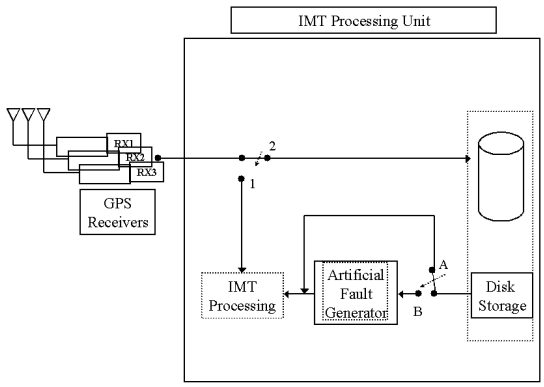

Figure 4: Real-Time vs. Post-Processing IMT

In this paper, post-processing is used to do testing on two datasets. One was collected on October 17, 1999, when Selective Availability (SA) (the deliberate degradation of GPS satellite clock accuracy) was still active. The other was collected on March 16, 2001, after SA had been permanently deactivated and after the recent IMT hardware upgrade. These will be designated "SA Off" and "SA On" to distinguish them. Both datasets are about four and a half hours long.

3.0 INTEGRITY MONITOR ALGORITHMS

LGF integrity monitor algorithms are targeted at a wide range of possible failures in the GPS SIS or in the LGF itself and are designed to insure that threatening failures are detected within three seconds at least 99.9% of the time. In this section, the SQM, DQM, MQM, MRCC, σμ–monitor, and MFRT monitor algorithms implemented in the current IMT are discussed in approximately the chronological order of processing. Each algorithm may generate one flag (failure) per channel (a given satellite tracked on a given receiver) or per satellite. The coordination of these monitors and combination of results will be discussed in Section 4.

3.1 SIGNAL QUALITY MONITORING (SQM)

The SQM prototype includes the distinct Signal Quality Receivers (SQR) and was developed separately from the IMT [14]. SQR consists of modified reference receivers that report C/A-code correlation measurements at several different correlator spacings. These measurements are processed by the SQM algorithms to determine if the ideal triangular C/A-code correlation shape has been appreciably altered by the presence of signal-deformation failures; otherwise known as "evil waveforms" [6,7,8]. The SQM prototype is being failure-tested separately and will soon be integrated into the IMT system.

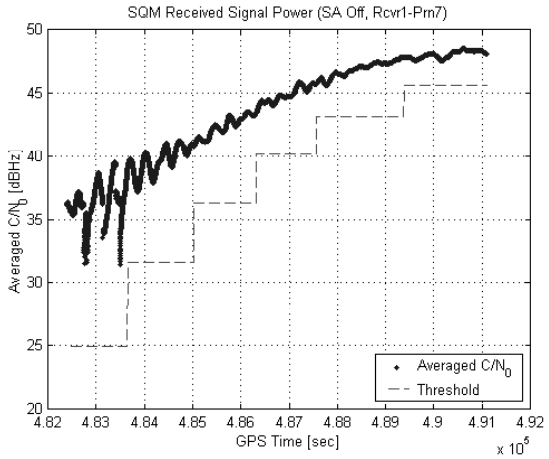

Figure 5: Nominal SQM Received Signal Power Test Results

SQM also has separate algorithms to confirm that received satellite signal power is within specifications and that the received signals do not include an anomalous amount of code-carrier divergence, which would interfere with LGF and user carrier smoothing and could lead to erroneous navigation solutions (see [17,18] for details of the code-carrier divergence monitor). The signal-power monitor takes an average of reported receiver C/N_0 for each channel at the current epoch k and the one at the previous epoch (k-1):

$$C/N_{o\_Avg}(k) = \frac{1}{2}\left(C/N_o(k-1) + C/N_o(k)\right) \qquad (1)$$

The averaged C/N_0 is compared with a threshold value. If the average is larger than the threshold, it passes this signal power test. In the IMT, the threshold is receiver-specific and is computed by the following steps. Under

nominal operation, the C/N_0 averages are recorded into 7 groups according to their corresponding elevation angles, i.e., less than 5°, 5° to 15°, 15° to 25°, 25° to 35°, 35° to 45°, 45° to 60°, and larger than 60°. Then, each group's mean (μ) and standard deviation (σ) are calculated. The threshold in each group is set to (μ–6σ). Though the distribution of signal power C/N_0 [dB] is still unknown (but probably not Gaussian), this threshold appears sufficiently tight to ensure that the relative aircraft/reference ranging error is not hazardous.

Figure 5 shows typical nominal results for the C/N_0 test. Data for GPS Satellite 7 (which is rising during this period) is shown while tracked on IMT Receiver 1. As can be seen in the figure, when the satellite elevation is low, the averaged C/N_0 received is also low and has a large variance, which makes it difficult to tightly monitor the signal power of satellites that first rise in view of the LGF. Averaging C/N_0 measurements over a 100-second interval, as suggested in [19], has been tried in the IMT, but because of the dominance of multipath on the IMT pinwheel antennas at low elevation angles, the C/N_0 variance reduction is small.

3.2 DATA QUALITY MONITORING (DQM)

The purpose of DQM is to verify that the satellite navigation data is of sufficient fidelity. DQM continually checks the validity of the GPS ephemeris and clock data for each satellite that rises in view of the LGF [9]. For a newly-risen satellite, the DQM compares the satellite orbit locations over the next 6 hours at 5-minute intervals based on the received new ephemeris message to those generated from the most recent almanac message and insures that the satellite positions always agree to within 7000 meters (limited by almanac accuracy). When navigation messages are updated (typically every two hours), DQM compares satellite positions based on old and new ephemeris to insure that the new ephemeris is consistent with the old, validated parameters to within 250 meters over the past 2 hours and the upcoming 2 hours.

DQM works in tandem with the MFRT function [1,9]. Measurements that pass all other IMT tests but cause LAAS range or range-rate corrections to exceed reasonable SPS error bounds could be affected by large ephemeris errors. MFRT is included in EXM-II and will be discussed in Section 3.6.

Figure 6 shows how the health status of a satellite's ephemeris evolves. When a satellite rises and becomes visible to the IMT (with status value 4), there is no old ephemeris, and DQM begins to validate the newly received ephemeris. DQM begins validation when at least two IMT reference receivers obtain identical new navigation data. After the initial ephemeris is validated (status 0), the navigation data is updated twice (and the status value jumps to 3), after which re-validation with the prior ephemeris occurs. Note that if the status is 4 or 5, it means that there is no validated ephemeris for a satellite. Status 1 implies a bad ephemeris, and the corresponding satellite should be excluded. In other cases (0, 2 and 3), there is always a validated ephemeris, and the messages from this satellite are safe to use.

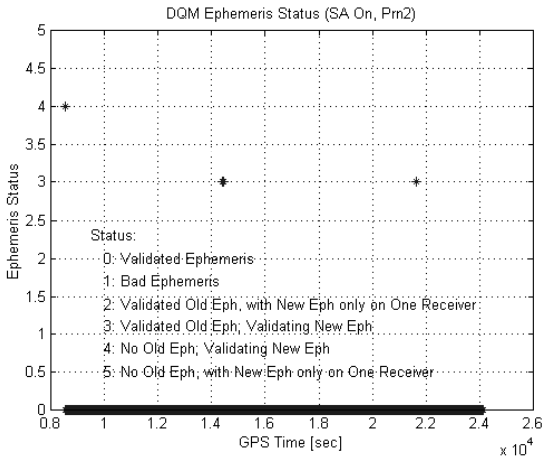

Figure 6: DQM Ephemeris Status of a Typical Satellite

The newly-released Non-Federal LGF Specification [16] includes more stringent requirements on DQM which require that additional monitors be added. These monitors include a more precise validation of ephemerides for newly-risen satellites (checking against the previous day's ephemeris in addition to the almanac) will be described in detail in an upcoming paper [20].

3.3 MEASUREMENT QUALITY MONITORING (MQM)

MQM confirms the consistency of the pseudorange and carrier-phase measurements over the last few epochs to detect sudden step and impulsive errors due to GPS SIS clock anomalies or LGF receiver failures. It consists of three monitors: the Receiver Lock Time check, the Carrier Acceleration-Step test, and the Carrier-Smoothed Code (CSC) Innovation test. Together, they generate one flag per channel if any of these three tests fails.

3.3.1 Receiver Lock Time Check

This check ensures continuous receiver phase lock by computing the numerical difference of the lock times reported by each reference receiver. Since this type of failure is most likely not hazardous (if it were, the acceleration-step test would catch it), the current IMT is implemented in such a way that it will not reject this measurement as hazardous even if this check fails; instead, the IMT initializes and resets several memory buffers

properly before using it, including re-initializing its carrier smoothing filter (see Section 3.3.2).

3.3.2 Carrier Smoothing and Pseudorange Corrections

Carrier smoothing and generation of differential corrections are the backbone of LGF processing under nominal conditions. For each channel, a pair of Receiver *m*-Satellite *n* at epoch *k*, the IMT applies carrier smoothing to reduce raw pseudorange measurement errors by using the following filter [1,2]:

$$PR_s(k) = \frac{1}{N_s} PR(k) + \frac{N_s - 1}{N_s} PR_{proj}(k) \quad (2)$$

where

$$PR_{proj}(k) = PR_s(k-1) + \phi(k) - \phi(k-1) \quad (3)$$

$$N_s = \tau_s / T_s \quad (4)$$

and PR and ϕ are the pseudorange and carrier-phase measurements, respectively. The smoothing filter uses a time constant τ_s of 100 seconds. As described earlier, the sample interval T_s is 0.5 seconds; thus N_s is equal to 200.

In the IMT, corrections are done on both pseudorange and carrier-phase measurements (note that only pseudorange corrections are required of the LGF):

$$PR_{sc}(k) = PR_s(k) - R(k) - \tau(k) \quad (5)$$

$$\phi_c(k) = \phi(k) - R(k) - \tau(k) - \phi(0) \quad (6)$$

where R and τ are the antenna-to-satellite range and satellite clock corrections computed from DQM-approved satellite navigation data.

3.3.3 Carrier Acceleration-Step Test

The carrier acceleration-step test is designed to detect an impulse, step, excessive acceleration, or other rapid changes in carrier phase measurements that could cause errors in the pseudorange corrections or in linear user latency corrections (to update corrections computed before user measurements are taken). On Receiver *m* at each epoch *k*, Satellite *n*'s last 10 continuous epochs (i.e., from epoch *k*-10 to epoch *k*-1) of ϕ^*, written as $\phi^*(n)$, is used to fit the following quadratic model:

$$\phi^*(n) = \phi_0^*(n) + \frac{d\phi^*(n)}{dt} t + \frac{d^2\phi^*(n)}{dt^2} \frac{t^2}{2} \quad (7)$$

where at each given fitting point/epoch (from *k*-1 to *k*-10),

$$\phi^*(n) = \phi_c(n) - \frac{1}{N_m} \sum_{i \in S_m} \phi_c(i) \quad (8)$$

and where S_m is the set of N_m satellites tracked on Receiver *m*. These satellites are assumed to be failure-free so far because they have passed all tests on previous epochs.

The least-squares method is used to solve for coefficients ϕ_0^*, $\frac{d\phi^*(n)}{dt}$, and $\frac{d^2\phi^*(n)}{dt^2}$ in the model, one of which expresses range acceleration:

$$Acceleration(n) \equiv \frac{d^2\phi^*(n)}{dt^2} \quad (9)$$

A second test statistic expresses the apparent change (or "step") in the latest measurement and is defined as [2]:

$$Step(n) \equiv \phi^*_{meas}(n) - \phi^*_{pred}(n) \quad (10)$$

where $\phi^*_{meas}(n)$ is the computed value at the current epoch *k* from (8), and $\phi^*_{pred}(n)$ is the predicted value from (7) at epoch *k*.

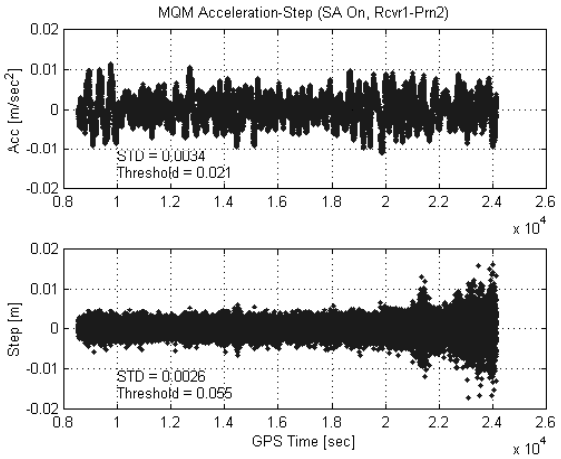

Figure 7: Nominal MQM Acceleration-Step Test Results with SA On

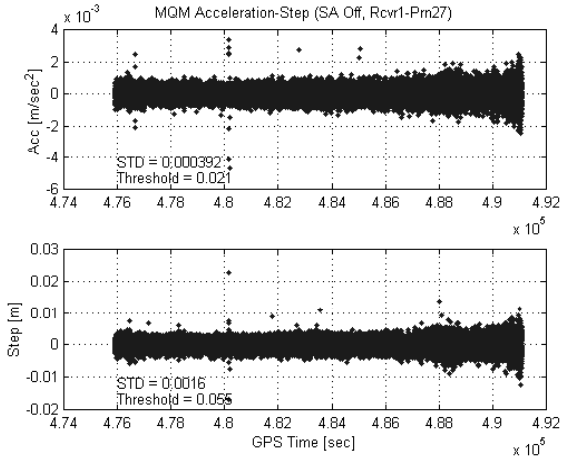

Figure 8: Nominal MQM Acceleration-Step Test Results with SA Off

For comparison, nominal acceleration-step test results with SA on and off are shown in Figure 7 and Figure 8, respectively. With SA on, the acceleration estimate oscillates due to the satellite clock dither being driven by SA, and its standard deviation is several times larger than

the one with SA off, which appears more or less as white noise. Both figures also show the trends of acceleration and step magnitudes versus satellite elevation angle – both grow as the elevation angle decreases.

Based on theoretical analysis and nominal testing, acceleration and step thresholds, with SA on, are set to be seven times their standard deviations in order to limit the probability of flags under fault-free conditions to levels low enough to meet the LGF continuity requirements [1,2]. If either an acceleration or step test statistic exceeds the corresponding threshold, the channel fails and is flagged for later exclusion by EXM-I (see Section 4.1).

3.3.4 CSC Innovation Test

This test is used to detect impulse and step errors on raw pseudorange measurements. The innovation test statistic is defined as:

$$Inno(k) \equiv PR(k) - \left(PR_s(k-1) + \phi(k) - \phi(k-1)\right) \quad (11)$$

To limit fault-free alarms, the innovation threshold is set as 6.82 times its theoretical standard deviation, i.e.,

$$\begin{aligned} Threshold_{inno}(k) &\equiv 6.82\sigma_{theoretical}(k) \\ &= 6.82\times(0.1+2.0e^{-3.0\alpha(k)}) \end{aligned} \quad (12)$$

where $\alpha(k)$ is the satellite elevation angle in radians at epoch k. If two or all of innovations in three successive epochs are over thresholds, a real flag will be generated from this CSC innovation test, and the channel will be excluded by EXM-I. If not, the measurement can still be used, but the smoothing-filter update for that epoch does not use the flagged raw pseudorange. Instead, the weight on the *PR* term in (2) is set to zero, and the weight on the carrier-phase update term in (2) is set to one for this epoch.

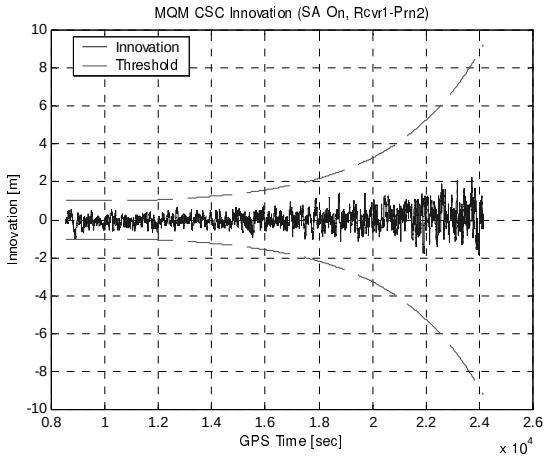

Figure 9: Nominal MQM CSC Innovation Test Results with SA On

Nominal innovation test results are shown in Figure 9 and Figure 10. In both plots, the innovation test statistics appear as white noise (as would be expected of filter residuals under nominal conditions), and the test statistic variances grow as the elevation angle decreases (this is reflected in the detection thresholds). With SA off, the raw pseudorange inputs to the smoothing filter are no longer driven by SA clock dither and are more consistent from epoch to epoch (multipath and ionosphere errors remain but are highly time-correlated); thus the innovation test statistic variance is much smaller than that of the SA-on case.

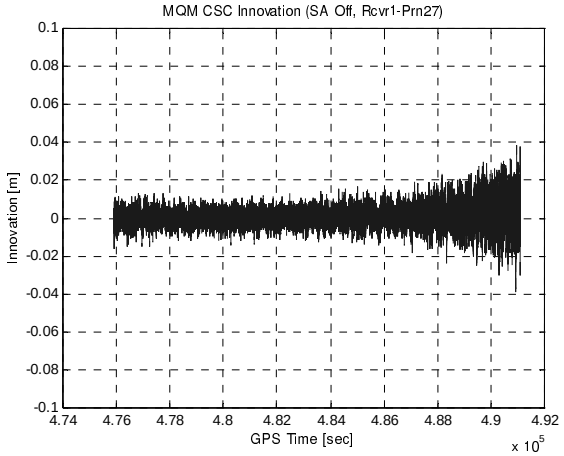

Figure 10: Nominal MQM CSC Innovation Test Results with SA Off

3.4 MULTIPLE REFERENCE CONSISTENCY CHECK (MRCC)

After EXM-I (see Section 4.1), the remaining measurements are used to compute candidate corrections; then B-values are computed to express the consistency of these corrections for each satellite across all reference receivers. B-values are best thought of as estimates of pseudorange error under the hypothesis that a given reference receiver has failed. MRCC is a process of computing and checking these B-values to be able to broadcast them to users (they are needed for users to compute "H1" protection levels) and to isolate any receivers or receiver channels that create anomalously large errors in the corrections.

After computing the individual corrections for pseudorange and carrier phase for each Receiver *m*, Satellite *n* pair in (5) and (6), a receiver clock correction is computed as follows to allow measurements to be compared across receivers [1,2]:

$$PR_{sca}(m,n) = PR_{sc}(m,n) - \frac{1}{N_c}\sum_{j\in S_c} PR_{sc}(m,j) \quad (13)$$

$$\phi_{ca}(m,n) = \phi_c(m,n) - \frac{1}{N_c}\sum_{j\in S_c} \phi_c(m,j) \quad (14)$$

where S_c denotes a set of N_c ranging sources in the maximum common set (see Section 4.1). Then the averaged, clock-adjusted corrections for Satellite n are:

$$PR_{corr}(n) = \frac{1}{M_n} \sum_{i \in S_n} PR_{sca}(i,n) \quad (15)$$

$$\phi_{corr}(n) = \frac{1}{M_n} \sum_{i \in S_n} \phi_{ca}(i,n) \quad (16)$$

where S_n is a set of M_n reference receivers, each of which reports an EXM-valid measurement for Satellite n. B-values are generated by the following equations [1]:

$$B_{PR}(m,n) = PR_{corr}(n) - \frac{1}{M_n - 1} \sum_{\substack{i \in S_n \\ i \neq m}} PR_{sca}(i,n) \quad (17)$$

$$B_{\phi}(m,n) = \frac{1}{M_n} \sum_{i \in S_n} \Delta\phi_{ca}(i,n) - \frac{1}{M_n - 1} \sum_{\substack{i \in S_n \\ i \neq m}} \Delta\phi_{ca}(i,n) \quad (18)$$

where

$$\Delta\phi_{ca}(i,n) = \phi_{ca}(i,n) - \phi_{ca0}(i,n) \quad (19)$$

Note that $\phi_{ca0}(m,n)$ is equal to $\phi_{ca}(m,n)$ evaluated at the first measurement epoch for the pair (m,n). $\phi_{ca0}(m,n)$ is reset after any change in the common set of satellites used to generate the clock adjusted carrier phase $\phi_{ca}(m,n)$ (this is a fairly common event).

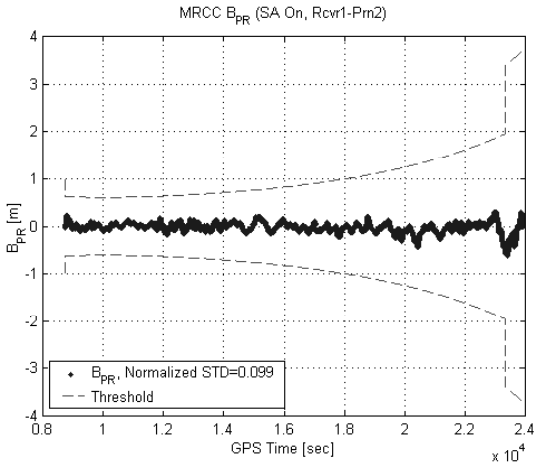

Figure 11: Nominal MRCC B-Value (B_{PR}) Test Results with SA On

Nominal test results are shown in Figure 11 and Figure 12. These sample B-values show that, as with the previous test statistics, the standard deviation of B_{PR} is larger for lower elevation satellites. Note that B_{PR} is time-correlated over 100 seconds due to the carrier smoothing. The B-values are compared to an elevation-dependent MRCC threshold. The IMT uses the following pseudorange B-value threshold [1,2]:

$$D_{PR} = 5.6 \frac{\sigma_{pr\_lgf}}{\sqrt{M_n(M_n - 1)}} \quad (20)$$

Note that no B-value can be generated when a satellite is tracked only on one receiver. The following formula is used to calculate the elevation-dependent σ_{pr_lgf}, which is the assumed PR_{corr} (from (15)) error standard deviation:

$$\sigma_{pr\_lgf} = 0.16 + 1.07 e^{-\alpha \times 90.0/\pi/15.5} \quad (21)$$

where α is the satellite elevation angle in radians. The threshold changes both in Figure 11 and in Figure 12 reflect the effects of changes in both M_n and α. The slow trend is due to changes in α, while the sudden jump near the end is due to one of the three reference receivers losing lock on the satellite (M_n decreases from 3 to 2).

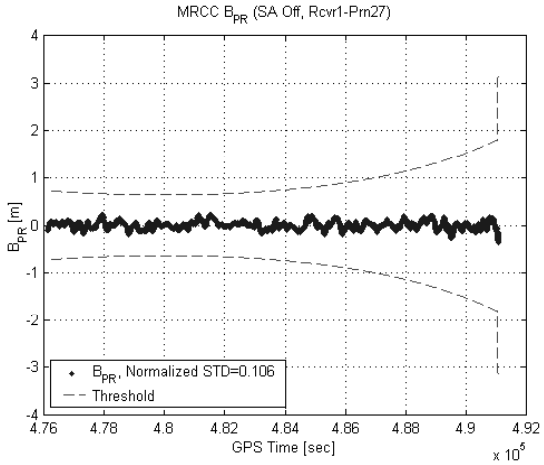

Figure 12: Nominal MRCC B-Value (B_{PR}) Test Results with SA Off

After B-values for both pseudorange and carrier-phase measurements are computed, the first step of MRCC is to compare each B-value with its threshold and, if any B-values exceed their thresholds, find the largest B_{PR} and B_{ϕ} that are over the thresholds. Figure 13 shows the diagram of the MRCC fault detection process as designed (the current IMT is very similar to but may not exactly follow this flowchart). This process determines whether there are any B-values over threshold, whether they are large B-values of different types (B_{PR} and B_{ϕ}), and whether the different types of the largest B-values are on the same channel.

The MRCC isolation procedure, which is also called the EXM-II "pre-screen", picks up the detection status from Figure 13 and continues to isolate errors [2]. Figure 14 shows this procedure. If MRCC fault detection returns a status of 0, i.e., no B-values are over threshold, and the MRCC isolation procedure does nothing. If the status is 2, i.e., the largest B_{PR} and B_{ϕ} are not on the same channel, and the pre-screen is not able to handle this complex situation. In other cases, it removes the channel with the largest B-value and then re-computes B-values with a reduced common set. If no MRCC flags remain, the pre-screen was successful (the reason for the pre-screen is that

a single measurement fault will affect multiple B-values due to clock adjustment and correction averaging). If MRCC flags still exist, it means that the pre-screen may have mishandled the failure, so the original B-values are retrieved and MRCC passes the task to EXM-II.

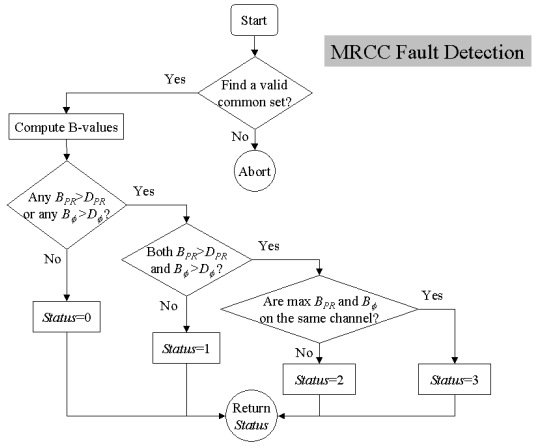

Figure 13: MRCC Fault Detection Flowchart

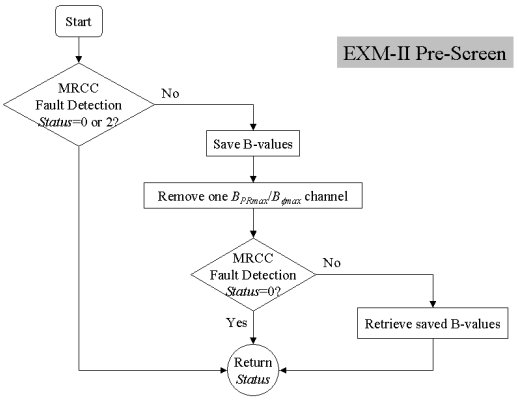

Figure 14: EXM-II Pre-Screen Flowchart

3.5 SIGMA-MEAN ($\sigma\mu$) MONITOR

B-values from MRCC are used as input to this $\sigma\mu$ monitor. This monitor is designed to insure that the pseudorange correction error distribution is bounded by a zero-mean Gaussian distribution with the broadcast σ_{pr_lgf} value. Two different monitors, the estimation and Cumulative Sum (CUSUM) algorithms, have been implemented for both mean and sigma violations. The details of these monitors are covered in [10,11,12].

3.6 MESSAGE FIELD RANGE TEST (MFRT)

MFRT is actually executed as the last step in EXM-II before measurements are approved for broadcast. It confirms that, with SA on, the computed average pseudorange corrections PR_{corr} from (15) are within a bound of ± 327.67 meters, and that the correction rates RR_{corr} are also within a threshold of ± 3.4 m/sec. With SA off, the bounds are ± 75 meters and ± 0.8 m/sec [13,16]. Note that MFRT operates on the averaged corrections and thus generates only one flag per satellite. RR_{corr} at epoch k is computed by:

$$RR_{corr}(k) = \frac{PR_{corr}(k) - PR_{corr}(k-1)}{T_s} \qquad (22)$$

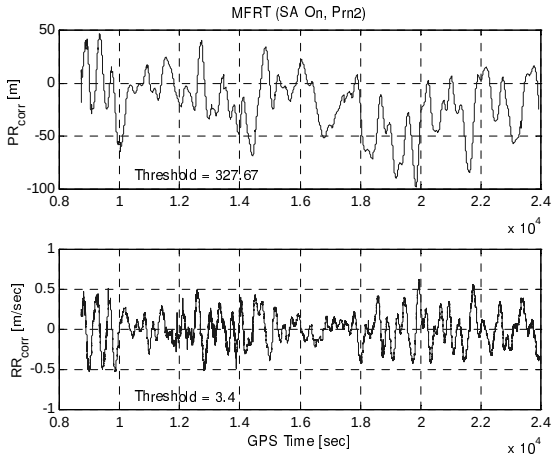

Figure 15: Nominal MFRT PR_{corr} and RR_{corr} Test Results with SA On

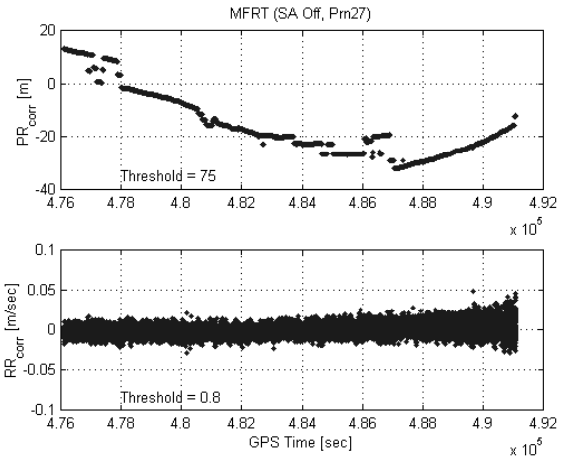

Figure 16: Nominal MFRT PR_{corr} and RR_{corr} Test Results with SA Off

Examples of nominal MFRT test results with SA on and SA off, respectively, are shown in Figure 15 and Figure 16. The characteristic oscillatory clock dithering of SA is clearly visible in Figure 15, as the *PR* and *RR* corrections trace this dominant error so that users can remove it. Figure 16, in contrast, shows gradual changes in the actual satellite clock, ionosphere, troposphere (no ionosphere or troposphere corrections are applied), and IMT multipath punctuated by occasional jumps (a few of which are also visible on the SA-on plot after zooming in). Small jumps

(typically under one meter) occur when the satellite navigation data changes and is revalidated by DQM. These jumps occur on the aircraft as well (it switches ephemerides on the same epoch) and are thus not a concern. However, several back-and-forth jumps appear to be due to changes in the common set used to perform clock corrections. When a low-elevation satellite is tracked on 2 of 3 IMT receivers but loses and regains lock on the third receiver, it comes in and out of the common set as currently constructed, which means that the ionosphere and troposphere errors on the low-elevation satellite (which could be tens of meters near solar maximum) come in and out of the satellite clock correction (the far-right-hand term in (13)), causing PR_{corr} jumps of as high as several meters.

Common-set-driven jumps in PR_{corr} do not threaten users because the corrections for all approved satellites jump by the same amount; thus the common change in user pseudoranges after applying LGF corrections goes entirely into the user clock bias solution and does not affect position solutions. However, frequent common-set changes stress EXM processing and should be reduced to the extent possible. The next IMT update will limit the common set to satellites above 10^{o} elevation in order to minimize the in-and-out common set issue visible in Figure 16 (LGF sites with Multipath Limiting Antennas reliably track down to 5^{o} and do not require this limitation).

4.0 EXECUTIVE MONITORING (EXM)

In the previous sections, several major integrity monitor algorithms were discussed. Each monitor algorithm is targeted to detect certain failures and may generate one flag per receiver-satellite pair or per satellite. This section will explain how EXM coordinates the results of these monitor algorithms and isolates faulty measurements when necessary. EXM prevents faulty measurements from being used in the LAAS VHF data broadcast.

EXM consists of two stages, EXM-I and EXM-II, as shown in Figure 3.

4.1 EXM-I

This first stage of EXM acts on any quality monitor (QM) flag and excludes measurements with at least one flag (value of 1) from any of SQM, DQM, and MQM. In the current IMT implementation, two matrices, called tracking (**T**) and decision (**D**) matrices, are created to support EXM. Each entry in the matrices is for a single channel or receiver-satellite pair. The **T**-matrix records which satellites are physically tracked by each receiver. The **D**-matrix is constructed by combining the three QM flags in a logical-OR operation (i.e., a flag in any QM algorithm leads to a flag in **D**) [2].

The details of EXM-I fault handing logic and its 11 isolation action cases are given in [2,15]. All but two of these cases can be regarded as combinations of three fundamental situations:

a) Flag on a single satellite on a single reference receiver;
b) Flags on a single satellite and multiple reference receivers;
c) Flags on multiple satellites on a single reference receiver.

In case (a), resolution is straightforward: the single flagged measurement is excluded. Cases (b) and (c), where multiple flags exist on a single satellite or reference receiver, lead to the exclusion of that satellite or receiver entirely. When combinations of cases (b) and (c) occur, all suspect satellites and receivers are excluded – no attempt is made to guess whether a satellite or receiver is the cause. These rules are conservative but make sense when not enough redundancy is present to clearly determine which elements have failed. Note that the LGF continuity risk probabilities allocated to satellite or LGF hardware failures are on the order of 10^{-6} per 15 seconds. The EXM exclusion logic is based on the assumption that failures and fault-free alarms are truly rare; thus any cases with multiple flags should be treated very cautiously.

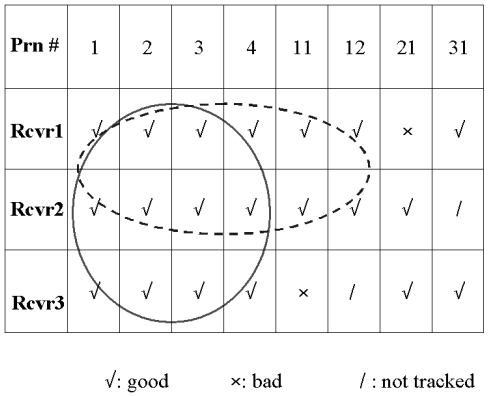

Figure 17: An Example of Selecting a Common Set

A common set of visible satellites that survive EXM-I is selected so that the receiver clock adjustment calculation in (13,14) uses the same set of satellites. It consists of at least four satellites that are all tracked on at least two common receivers. A common set that consists of all three receivers has higher priority than one consisting of two receivers. For example, Figure 17 shows a table similar to the form of the **D** and **T**-matrices. The common set that the IMT will select is the one circled by the red solid line

instead of the blue dashed line. If fewer than four satellites are tracked by all three receivers, the largest common set that can be formed from any two receivers is chosen. If a common set of at least four satellites cannot be approved, and this event was not predicted by the LGF "constellation alert" algorithm (which considers known satellite outages), all measurements are excluded, and the IMT is reset.

4.2 EXM-II

Once corrections, B-values, and MRCC flags have been generated, a second phase of executive monitoring is performed. The logic used to handle MRCC flags is complex because each attempted measurement exclusion usually requires the IMT to reduce the common set S_c. When this happens, corrections and B-values must be re-computed for all remaining measurements, and then EXM-II must confirm that this new, reduced set of measurements passes MRCC.

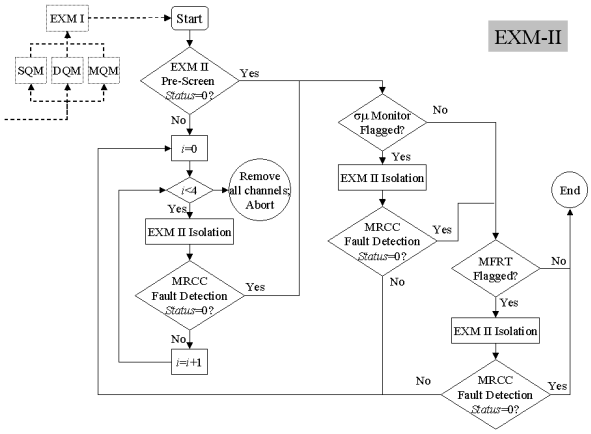

Figure 18: EXM-II Flowchart

Figure 18 shows a diagram of the EXM-II flowchart. In this diagram, the block named "EXM-II Pre-Screen" has been explained in Figure 14. The block named "EXM-II isolation" operates using logic similar to EXM-I. In other words, single B-value flags on individual channels are removed, but if more than one B-value is flagged on a given satellite or receiver, the entire satellite or receiver is excluded. In order to bound the processing time required by EXM-II, the IMT places a limit of four attempts to exclude measurements and repeat the EXM-II isolation-MRCC fault detection process before it "gives up" and excludes all measurements. It should be very rare for EXM-II to need to make more than one or two exclusions. Note that all measurements must also pass the σμ monitor and MFRT tests – any flags from these tests also necessitate EXM-II exclusions.

Finally, EXM is responsible for managing excluded ranging sources and/or receiver channels. If measurements are excluded due to failures detected by the σμ monitor or MFRT, they will be excluded for the current pass because, for the faults targeted by these monitors, it is difficult to verify during the current pass that a detected failure is no longer present. In most other cases, excluded measurements enter "self-recovery" mode, in which the carrier smoothing filters are re-initialized. During self-recovery, if the re-smoothed measurement passes all tests with thresholds tightened to 3-sigma levels, the measurement is declared safe for use again and is reintegrated into the set of usable measurements [2,3]. Tightened thresholds are also applied to those measurements that were excluded because of σμ monitor or MFRT flags in the previous pass. Tightened thresholds help insure that the probability of reintroducing a measurement that is still failed is below 1.94×10^{-9} [1]. If a failed channel does not pass the tightened thresholds, it enters self-recovery mode again and repeats the recovery process. If self-recovery is failed a second time, the affected channel enters "external maintenance" mode and cannot be used again until the affected equipment is certified to be health by an external agency, such as FAA maintenance personnel [2].

5.0 FAILURE TEST RESULTS

LGF verification testing can be divided into two components: nominal testing and testing with simulated failures. Many nominal test results have been shown in Section 3. As is expected, no flags are generated under nominal conditions in these plots (however, flags representative of real measurement faults do occur once in a while, indicating that the current IMT hardware falls short of meeting the Category I LGF continuity allocation). In order to test the IMT response to actual failures, "failure testing" is conducted in which deliberate errors are injected into the system.

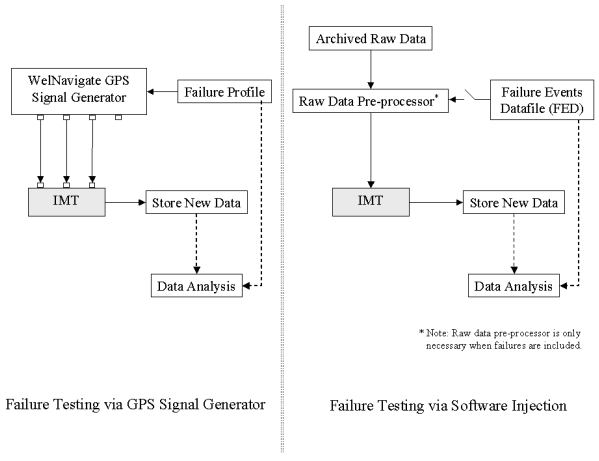

Figure 19: Two Failure Testing Methods: via GPS Signal Generator and via Software Injection

Figure 19 illustrates two possible failure-testing methods. One is to program our WelNavigate 40-channel GPS signal simulator to generate failed RF signals, which are then fed into the IMT. The other is to modify stored IMT data collected under nominal conditions to inject failures and then to let the IMT post-process these packets. Each method has its pros and cons with respect to the different failure scenarios. Some failure modes, such as range jumps or satellite ephemeris errors, can be implemented in either way, but it is usually much easier and faster via the post-processing method.

A sample of our recent failure test results is included here. In these plots, errors are injected into the dataset with SA on by the post-processing software method.

5.1 STEP ERRORS ON PSEUDORANGE MEASUREMENT

On the raw pseudorange measurement from Receiver 1 and Satellite 2, two step errors are injected: a +2.0-meter step from GPS time 10000 [sec] to 10020, and a –2.0-meter step from 10040 to 10060. These steps are shown in the top plot in Figure 20. This scenario simulates a sudden dramatic change in the C/A-code due to either ground multipath or satellite signal deformation.

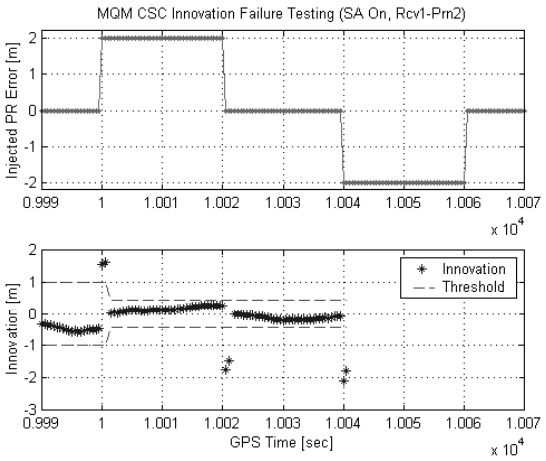

Figure 20: Failure Test Results with Injected Step Errors on Pseudorange Measurement

As stated in Section 3.3.4, the MQM CSC innovation test can catch step or jump errors on raw pseudorange measurements. The bottom plot in Figure 20 demonstrates this. When the first jump occurs, two successive innovations are larger than the threshold; thus an MQM flag occurs and the measurement is excluded by EXM-I at 10000.5 [sec], i.e., 0.5 seconds after the step jump. During self-recovery, the innovation threshold is reduced. After two chances of flagging during self-recovery periods are used up, the satellite is removed and goes into the external maintenance mode. As shown here, the CSC innovation test is fast and effective in responding to sudden, significant errors in pseudorange measurements.

5.2 STEP ERRORS ON CARRIER-PHASE MEASUREMENT

Errors with the same pattern as in Section 5.1 are injected on the carrier-phase measurement from Satellite 2. The magnitude of both steps is 3 cycles, or 0.382 meters, as shown in the top plot in Figure 21. Note that this test is independent of the previous one. It is more typical of a satellite clock jump or receiver cycle slip. In the satellite-clock-failure case, the jump will be seen on pseudorange as well, but the fault magnitude in this case is too small to be reliably detected by the CSC innovation test. However, the MQM acceleration-step test is very sensitive to carrier-phase jumps of this type.

The middle plot of Figure 21 is similar to the plot in Figure 20 and shows the acceleration-step test responses on Satellite 2. It shows that the errors on carrier phase cause large jumps in the step test result such that it exceeds the threshold. After Satellite 2 is excluded, 5.5 seconds are required to collect data and re-calculate acceleration-step test statistics, which is required before carrier smoothing can safely begin again. If 200 seconds of smoothing are completed without further violations of the tightened self-recovery thresholds, the measurement is reset to be healthy and can be used again.

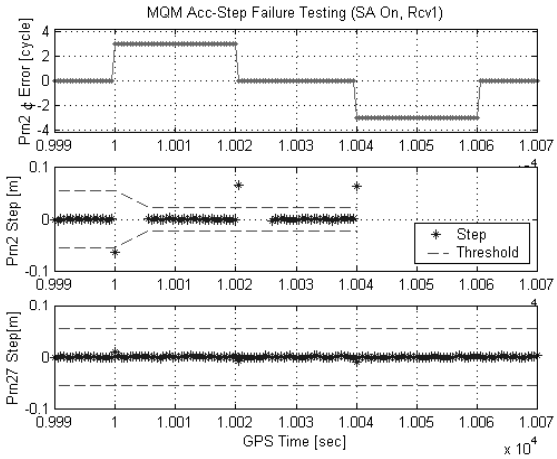

Figure 21: Failure Test Results with Injected Step Errors on Carrier Phase Measurement

The bottom plot of Figure 21 shows how the test responses on other satellites that are also tracked on Receiver 1 are impacted by the fault on Satellite 2. Equation (8) tells us that if there is an error in Satellite 2, other error-free satellites (like Satellite 27 in this test) may have slight jumps in their acceleration-step results due to the clock adjustment included in ϕ^*. This plot shows that jumps on non-failed satellites do occur but are very small

and are far below the threshold. If the step error injected in Satellite 2 were large enough, it could also cause other satellites' step-test results to exceed the threshold. In our implementation of the acceleration-step test, a simple isolation function, similar to the MRCC pre-screen isolation procedure, prohibits mis-exclusion in this scenario. Only the satellite with the largest test statistic over the threshold is excluded at first, and if only this satellite has failed, the jumps on the other satellites will disappear.

5.3 EXM-II FAILURE TESTING

Because of its complexity and the fact that it follows EXM-I, failure tests of EXM-II are often more complicated than those above. This section gives a relatively simple example. Independent of the two previous failure tests, a jump error of 15 meters at GPS time 10000 [sec] is injected on raw pseudorange measurements from the channel of Receiver 1 and Satellite 2, as shown in the top plot in Figure 22. Section 5.1 demonstrates that MQM would issue flags due to the injected error before EXM-II starts to response and then the channel would go into self-recovery period. During the self-recovery period, no LAAS corrections or MRCC B-values would be generated and broadcast. Thus, this failure test requires that the MQM be disabled in order to let the fault appear in the B-values and in the EXM-II logic.

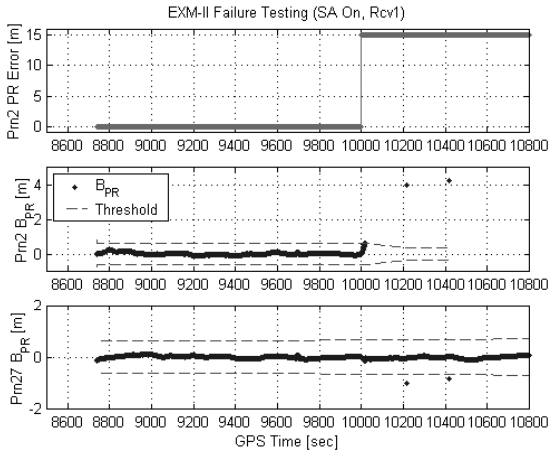

Figure 22: EXM-II Failure Test Results

The MRCC response on the failed channel is shown in the middle plot of Figure 22. Its affected B_{PR} starts to increase at time 10000 and rises over the threshold at time 10017. After two failed self-recovery periods (note that the original fault is still present), the failed channel (Receiver 1, Satellite 2) is placed into external maintenance mode.

Notice that two MRCC B-value flags are also issued on another channel (Receiver 1, Satellite 27), due to the magnitude of the fault and its impact on the clock-adjust and averaging functions, as shown in the bottom plot in Figure 22. However, MRCC and EXM-II logic isolates the actual failures without removing these indirectly-affected measurements. EXM does not remove the channel (Receiver 1, Satellite 27) but only the channel (Receiver 1, Satellite 2) which contains the actual fault. Once the failed channel (only) is excluded, the B-values on the non-failed channels return to normal.

6.0 CONCLUSIONS AND FUTURE WORK

This paper has demonstrated that Stanford's LAAS Integrity Monitor Testbed (IMT) can detect and exclude failures such that it meets the LGF integrity requirements in the Category I LAAS Ground Facility Specification. Based on its implementation of Executive Monitoring, the IMT continues to support LAAS operations after detecting and excluding failed satellites and/or receiver channels. Stanford has also developed a failure test methodology to validate that the IMT meets the LAAS Ground Facility (LGF) integrity requirements and that these requirements can be met. These methods have been adapted to form a verification test methodology for FAA Type Acceptance of future LGF equipment [4,15].

The test results shown here demonstrate that the IMT functions well under both nominal and failure situations. The IMT software has recently been upgraded to be compatible with the new NovAtel OEM-4 GPS receivers and a new faster Linux PC. The next update to the IMT will include updated DQM ephemeris monitors [20] and common-set logic improvements to reduce the number of satellite common-set changes. In addition, the SQM prototype will soon ready to be integrated into the IMT. We expect that testing of the integrated Category I IMT will be completed by the end of this year or early in 2002. New algorithms, including enhanced "soft error" detection algorithms and isolation logic and improved ephemeris-error monitors, are now under study to meet the more stringent Category II/III integrity and continuity requirements.

ACKNOWLEDGEMENTS

The authors would like to give thanks to many other people in the Stanford GPS research group for their advice and interest. Funding support from the FAA Satellite Navigation LAAS Program Office (AND-710) is appreciated. The opinions discussed here are those of the authors and do not necessarily represent those of the FAA and 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-2937, Sept. 21, 1999. Internet: http://gps.faa.gov/Library/Documents/laas_faa2937.pdf

[2] "FAA LAAS Ground Facility (LGF) Functions," Version 2.4. LAAS KTA Group, Unpublished Manuscript, September 9, 1998.

[3] 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.

[4] S. Pullen, M. Luo, *et.al.*, "GBAS Validation Methodology and Test Results from the Stanford LAAS Integrity Monitor Testbed,", *ION GPS 2000*, 19-22 September 2000, Salt Lake City, UT, pp. 1191-1201.

[5] M. Luo, S. Pullen, "LAAS Reference Receiver Correlation Analysis," Stanford University, Unpublished Manuscript, Feb. 17, 1999.

[6] P. Enge, E. Phelts, D. Akos, "Detecting Anomalous Signals from GPS Satellites", Stanford University, Draft Manuscript, Oct. 12, 1999.

[7] R.E. Phelts, D. Akos, P. Enge, "SQM Validation Report," ICAO GNSSP WG-B Meeting, Seattle, WA., May 30, 2000.

[8] D. Akos, *et.al.*, "GPS Signal Quality Monitoring: Test Results," *Proc. of the ION 2000 National Technical Meeting*. Anaheim, CA., Jan. 26-28, 2000, pp. 536-541.

[9] S. Matsumoto, *et.al.*, "GPS Ephemeris Verification for Local Area Augmentation System (LAAS) Ground Stations," *Proceedings of ION GPS '99*. Nashville, TN., Sept. 14-17, 1999, pp. 691-704.

[10] S. Pullen, *et.al.*, "The Use of CUSUMs to Validate Protection Level Overbounds for Ground-Based and Space-Based Augmentation Systems," *Proceedings of ISPA 2000*. Munich, Germany, July 18-20, 2000.

[11] B. Pervan, *et.al.*, "Sigma Estimation, Inflation, and Monitoring in the LAAS Ground System," *Proceedings of ION GPS 2000*. Salt Lake City, UT., Sept. 19-22, 2000, pp. 1234-1244.

[12] J. Lee, S. Pullen, *et.al.*, "LAAS Sigma Monitor Analysis and Failure-Test Verification,", ION 57th Annual Meeting, Albuquerque, NM, June 11-13, 2001 (forthcoming).

[13] S. Pullen, "Performance of LAAS Ephemeris Monitoring with Selective Availability Deactivated," Stanford University, Unpub. Presentation, June 7, 2000.

[14] A. Mitelman, *et.al.*, "A Real Time Signal Quality Monitor for GPS Augmentation Systems," *Proceedings of ION GPS 2000*. Salt Lake City, UT., Sept. 19-22, 2000, pp. 862-871.

[15] R. Key, *et.al.*, "LAAS Verification Methodologies and Cases," FAA AOS-240, Oklahoma City, OK., Unpublished Draft, May 20, 2000.

[16] *Specification: Category I Local Area Augmentation System Non-Federal Ground Facility*. U.S. Federal Aviation Administration, Washington, D.C., FAA/AND710-2937, May 31, 2001.

[17] B. Pervan, S. Pullen, J. Andreacchi, P. Enge, "Characterizing the Effects of Ionospheric Divergence and Decorrelation on LAAS," *Proceedings of ION GPS 2000*. Salt Lake City, UT., Sept. 19-22, 2000, pp. 653-661.

[18] B. Pervan, "A Review of LGF Code-Carrier Divergence Issues." Illinois Institute of Technology, Draft Revision 2, May 29, 2001.

[19] F. van Graas, "Detection of Satellite Low Signal Power." Ohio University, Revised Draft, April 30, 2001.

[20] S. Pullen, M. Luo, B. Pervan, F.C. Chang, "Ephemeris Protection Level Equations and Monitor Algorithms for GBAS," *Proceedings of ION GPS 2001*. Salt Lake City, UT., Sept. 11-14, 2001 (forthcoming).