Satellite Autonomous Integrity Monitoring and its Role in Enhancing GPS User Performance

Logi Viðarsson, Sam Pullen, Gaylord Green and Per Enge
Department of Aeronautics and Astronautics, Stanford University and NavAstro Co.

BIOGRAPHY

Logi Viðarsson completed his B.Sc. in Electrical Engineering at the University of Iceland. He now is working towards his Ph.D. degree in Electrical Engineering at Stanford University.

Sam Pullen completed his Ph.D. in Aeronautics and Astronautics at Stanford University, where he is now the Technical Manager of Stanford's Local Area Augmentation System (LAAS) project. His research focuses on system design and integrity algorithms for both the Local Area and Wide Area Augmentation Systems.

Gaylord Green has been an ION member since 1972, past President and Fellow of the ION. Colonel Green retired from the Air Force in 1988, when he was the GPS Joint Program Office (JPO) Program Director. Captain Green was a member of the original GPS Program Office Cadre as the Space Segment IPT Leader. He is in the Space Technology Hall of Fame and the GPS JPO Hall of Fame for his contributions to GPS.

ABSTRACT

One of the difficulties inherent in use of Global Navigation Satellite Systems (GNSS) for safety-critical applications is the need to insure user integrity in the face of an unbounded array of possible failures in the GNSS satellite signals. GPS satellites, for example, are generally operated until failure, and some of those failures will cause ranging errors which impact user integrity. Ground-based satellite-operator solutions to this problem are complex and costly and have difficulty alerting users of the failure quickly. Models for possible satellite failures have been developed but are imperfect because of the limited statistical information available to the satellite operators and to the civil user community.

As a result, it is difficult to support the single-failure assumption made by Receiver Autonomous Integrity Monitoring (RAIM) or to certify that Space Based Augmentation Systems (SBAS) and Ground Based Augmentation Systems (GBAS) safely detect all threatening satellite failures. In addition, when satellite failures occur, they must be detected and sorted out from an array of possible failures in SBAS and/or GBAS ground systems, or else continuity will be sacrificed unnecessarily. These difficulties would be greatly lessened if integrity monitoring were conducted within the satellite constellation itself so that immediate warnings could be transmitted to users.

We have developed and are prototyping a method for Satellite Autonomous Integrity Monitoring (SAIM) that could be applied to future GNSS satellites such as GPS III. The processing demands on SAIM are much lighter than those on a LAAS or GBAS ground system because each satellite's SAIM function monitors the ranging signals of the satellite it is attached to rather than the 12 or more satellites that must be simultaneously handed by SBAS and GBAS.

This paper describes our SAIM concept in detail and presents test results from a SAIM software prototype now under development. This prototype has been tested against nominal satellite signals (to confirm that fault-free alarms are rare enough to support civil aviation continuity requirements) and several classes of failed signals. Practical implementation issues such as satellite multipath and receiver clock calibration will also be addressed.

Introduction

As satellite navigation is used for a wider array of civil and military applications, insuring the safety of the broadcast signals and data is a higher priority than ever. Civil aviation precision approach, for example, has very little room for error, and the integrity requirements for this operation must be very tight to satisfy modern civil aviation safety standards.

One challenge to assuring high integrity is that the systems responsible for detecting failures must alert the user in as little as 1-2 seconds from the time when the failure becomes hazardous. Because this is beyond the capability of the existing GPS constellation, integrity monitoring has been designed into Space Based and Ground Based Augmentation Systems (SBAS and GBAS). Tight integrity requirements have made the design of these systems very complex, and when a ground based station discovers a failure, it still takes several seconds to assure that users receive the alert. In addition, no single ground station can constantly monitor a given satellite for more than several hours at a time.

We suggest a simpler, more cost-efficient implementation by adding Satellite Autonomous Integrity Monitoring (SAIM) onboard navigation satellites themselves. This would significantly simplify the integrity monitoring requirements for augmentation systems and users because each satellite is constantly monitored to a level of integrity sufficient for almost all applications.

As part of its research on the Local Area Augmentation System (LAAS) Ground Facility (LGF) design for GBAS, Stanford has already developed the Integrity Monitor Testbed (IMT), which is a prototype of the integrity monitors required for an operational LGF [1]. The IMT provided a starting point for the design of the SAIM prototype that is described in this paper. Several aspects that are unique to SAIM, such as removal of the monitor receiver clock bias, are discussed in detail, and preliminary test results are shown that demonstrate the potential of this concept.

1 Goals and Motivation

As noted above, the integrity requirements of civil aviation applications can be met by augmentation systems such as SBAS and GBAS, but the design of these systems is complicated by the need to monitor for an array of possible failures that are difficult to even define. This increases the cost of these systems, delays their certification, and restricts their use once they are finally fielded and approved.

For example, to satisfy requirements for Category I precision approaches (instrument guidance down to a 200-foot decision height), the navigation system must be able to detect failures and warn users within 6 seconds after they become a hazard. For CAT II/III approaches to decision heights of 100 feet or less, that number goes down to 2 seconds. Ground-based integrity monitoring, such as that implemented within SBAS and GBAS, has a hard time meeting those specifications, since a ground-based system must first detect the failure and then alert users via its mechanism for broadcasting differential correction data. All of these steps impose time delays, and the possibility of missed correction messages further extends the time within which a warning can be guaranteed.

In addition, "stand-alone" users that operate without augmentation systems have come to expect similar protection against system failures. Some protection for stand-alone users is provided today by Receiver Autonomous Integrity Monitoring (RAIM), but that approach has its weaknesses, as described in Section 2.1 below. It is not possible to rely on today's GPS Operational Control Segment (OCS) for rapid warning of failures. OCS failure alerting is even more daunting than that for SBAS and GBAS, as OCS must contact a GPS uplink station once a failure has been detected and have it send a message to update the signal of the affected satellite so that users stop applying it to their position fixes. It is unreasonable to expect the current OCS to provide warnings in less than 10 - 20 minutes [10]. Future OCS enhancement to provide faster warnings is possible, but given the difficulties mentioned above, it would be almost impossible to provide warnings within 2 seconds to satisfy all civil user needs.

SAIM resolves the problem of rapid failure alerting by integrating detection and alerting within the satellite itself. Once a failure is detected in the satellite, a message is sent to the signal processor to change the transmission message such that it is immediately unusable. This should make it possible to alert failures within 1 second, which meets all currently-envisioned time-to-alert requirements.

To be useful, SAIM should have the ability to detect, at a minimum, the following satellite failure modes:

- **Abnormal signal power levels.** Signal power levels significantly above or below specified levels can disrupt the ability of users to track the affected satellite and could increase the resulting ranging errors.

- **Distortion of pseudorandom code signals.** Deviations in the pseudorandom code signal patterns, such as is believed to have occurred on GPS SVN 19 in 1993 [11], can cause pseudorange errors. Differential users with tight accuracy and integrity limits are threatened if the reference and user

receivers differ in how they receive and track the deformed signals [12].

- **Code/carrier divergence.** Most modern receivers use both the GPS code and carrier signals to provide more accurate ranging measurements. If there is a divergence between these, these range measurements will become erroneous over time.

- **Excessive clock acceleration.** At the heart of each GPS satellite lies an atomic clock that is kept synchronized with GPS system time. However as we will see later, it does not take much of a divergence from GPS time to cause serious pseudorange errors to emerge. Augmented users generally remove these clock errors when they apply differential corrections, but unusual clock dynamics introduce errors into user corrections for the latency, or age, of the corrections.

- **Erroneous navigation data.** Each GPS satellite broadcasts its position via the ephemeris and almanac messages in the GPS navigation data, and satellite clock correction coefficients are provided in another data message. Large errors in these messages will cause errors for stand alone and (to a lesser degree) differential users.

- **Single Event Upset (SEU) navigation signal degradation of digital payloads.** The next generation of GNSS satellites will likely be equipped with more and more digital circuitry for ranging signal generation, but the space environment can be harmful to such equipment and may cause sudden. unpredicted "bit-flips" in digital signals.

Ideally, SAIM would be built into future navigation satellite designs so that signals are monitored as they are generated. This may be technically feasible, but it is more practical to design SAIM as a modular "add-on" to existing GNSS satellites. This would allow a separate agency, such as a Civil Aviation Authority, to procure and certify the SAIM component and then provide it to the GNSS operator. Even if this is not necessary, modular SAIM minimizes the changes that are required to existing navigation satellite designs and limits the interface between SAIM and the satellite to two items: (1) a means to receive the satellite signals, and (2) a means to alert the satellite when a failure is detected. A patent for this implementation of SAIM is now pending.

2 SAIM Integrity Benefits for Stand-alone and Augmented Users

SAIM is designed to significantly enhance the integrity of all GPS users by reducing the probability of a latent space-segment failure being present in user measurements. Figure 1 summarizes these benefits for the two classes of users (stand-alone and augmented) that are discussed below.

2.1 Benefits for Stand-alone Users

The majority of GPS users today do not have access to differential corrections (which can correct for and warn of satellite anomalies) and thus rely on the accuracy and integrity of the transmitted satellite signals themselves. An estimate of the integrity of the existing GPS signals was made in the GPS Standard Positioning Service (SPS) Signal Specification [7]: no more than three (3) "major service failures" should occur per year across the entire constellation, where a "major service failure" is defined as a failure that causes a user pseudo range error of greater than 150 meters. Three such failures in one year across the existing 24-satellite primary constellation combined with an average Operational Control Segment Response time of 6 hours implies a state failure probability of 8.6×10^{-5} per satellite per hour. More typical numbers would be on the order of one such failure per year and a 30-minute response time, which would give a probability of 2.4×10^{-6} per satellite per hour [7].

Because of the conservatism in the failure likelihood and duration used to derive the former probability, the civil community has generally accepted a value of 10^{-4} per satellite per hour as an upper bound [8]. However, it should be noted that some civil GPS applications, such as aircraft precision approach, would be threatened by failures that are not "major" according to the above definition.

SAIM would greatly improve stand-alone user integrity by acting as an additional "screen" on satellite failures before they affect users. The integrity monitors that make up SAIM will be designed to support a missed-detection probability (P_{MD}) of 0.001 or less for faults which cause stand-alone ranging errors of 20 meters or more. Thus, the actual integrity provided to stand-alone users will be at least three orders of magnitude better than it is today and should be good enough to meet the 10^{-7} per hour requirements for civil aviation Signal-in-Space integrity (this number applies to lengthy operations; precision approach exposure times are 150 - 250 seconds)[6]. It will also support the needs of military users who would like to be able to bound the magnitude of errors in munitions guidance to reduce the probability of collateral damage.

Integrity for stand-alone users today is based on Receiver Autonomous Integrity Monitoring (RAIM) [6]. RAIM is a useful technique, but it has limitations that do not apply to SAIM. RAIM requires redundant satellite geometry (at least 5 or 6 satellites in view) to check the health of each satellite measurement against the others, which limits user availability and is less sensitive to failures on multiple satellites at the same time. SAIM, on the other hand, will provide integrity to each satellite individually, so users will generally be able to navigate safely with only 4 visible

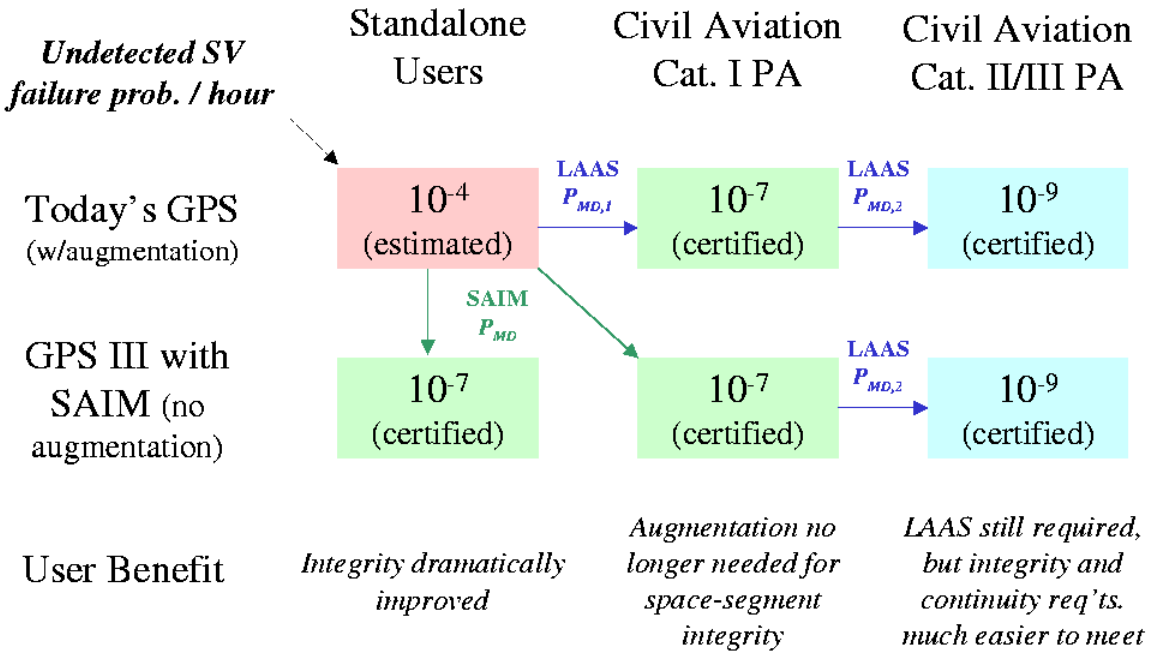

Figure 1: Summary of SAIM Integrity Benefits

satellites. RAIM may still have a role to play in detecting anomalies in the vicinity of a given user, such as RF interference, that will not be covered by SAIM.

2.2 Benefits for Augmented Users

Because the ranging accuracy and integrity guarantees provided for the existing GPS Standard Positioning Service fall short of what is needed for civil aviation navigation, the Federal Aviation Administration is developing Wide Area and Local Area Augmentation Systems (WAAS and LAAS) to provide differential corrections and integrity monitoring to users with tight requirements. One of the roles of WAAS and LAAS is to detect the satellite failures that threaten integrity and alert users accordingly, either by excluding the affected satellite from the set of broadcast corrections or, for less-severe failures, by increasing the broadcast standard deviation of ranging errors for that satellite (after applying differential corrections). With SAIM in place, this burden will be lifted for most WAAS and LAAS installations, which means that these systems either will not be necessary or will have less-stringent requirements, making them easier to field, certify, and afford.

For example, LAAS systems fielded to meet the requirements of Category I precision approach (where guidance is required down to a 200-foot decision height) are required to achieve a Signal-in-Space integrity of 2×10^{-7} per 150-second approach [6]. SAIM will be able to achieve that level of integrity against satellite failures; thus relieving Category I LAAS of that responsibility. Because differential corrections are based on reference receiver measurements, LAAS must also detect failures in the reference receivers as well as any other failures affecting the reference receiver measurements. With SAIM in place, the difficulty in separating satellite failures from reference receiver failures is greatly reduced because receiver failures will be far more probable than satellite failures(roughly 2.4×10^{-4} vs. 10^{-7} per hour). If the ranging accuracy of future GPS satellites is small enough to make differential corrections no longer necessary to achieve the Category I accuracy and alert-limit requirements [9], it may be possible to achieve Category I operations without any augmentation, although this is a far-off prospect.

The tightest civil-aviation requirements apply to Category III precision approaches down to a 50-foot or lower decision height. In this case, the Signal-in-Space integrity requirement is 10^{-9} per 250-second approach [6]. Furthermore, because the accuracy and alert-limit requirements for Category III are much tighter than for Category I, it seems likely that LAAS differential corrections will continue to be necessary regardless of the extent of future GPS modernization. Even so, SAIM would reduce the missed-detection probability against satellite failures required of Category III LAAS by three orders of magnitude. This is not a trivial benefit, as it is much easier to achieve a missed-detection probability of 10^{-3} as opposed to 10^{-6}. Again, because reference receiver failures will be far more common than

satellite failures, the task of diagnosing the causes of failures detected by LAAS would be much easier, and this has significant implications for the degree of redundancy and software required of the LAAS ground system.

3 Satellite Autonomous Integrity Monitoring (SAIM)

As noted in Section 1, it is desirable to design SAIM as an add-on to existing navigation satellites such that SAIM is essentially an "on-board augmentation" to the satellite. This suggests that existing designs for airport-based GBAS ground systems can serve as a model for how to implement SAIM on a satellite. As part of its LAAS research, Stanford has developed the Integrity Monitor Testbed (IMT), which is a prototype of the LAAS Ground Facility (LGF) that includes the integrity monitors needed to meet civil-aviation integrity requirements in the absence of SAIM [1]. While it is not the only means of implementing SAIM, the IMT serves as a good basis for SAIM prototype development.

3.1 Introduction to the IMT

The IMT, which is now in its second version, consists of three main parts. After receiving and decoding the GPS signal on three redundant GPS reference receivers with antennas separated sufficiently to make multipath errors statistically independent, a first phase integrity monitoring is performed separately on each of the satellites tracked by each GPS receiver (one satellite on one receiver is known as a "channel"). The primary intent of the first phase of monitoring is to detect satellite failures, although any failure that affects a given channel will be picked up. In the second phase of monitoring, measurements are combined across receivers on each of the satellites in view in order to identify potential failures on a single receiver. Executive Monitoring (EXM) oversees both phases of integrity monitoring and decides which measurements, if any, are unhealthy and must be excluded from use. A block diagram of the IMT is included in figure 2.

Signal-in-Space Receive and Decode (SISRAD) is the interface between the SAIM monitor receivers (three NovAtel OEM4 L1-only GPS receivers) and the integrity processor. It converts raw receiver data into an internal data structure suitable for our integrity tests. This includes carrier phase and pseudorange measurements from the receivers along with the decoded navigation data [1, 5].

Signal Quality Receiver (SQR) and Signal Quality Monitoring (SQM) monitor signal power levels and check for the existence of evil waveforms on the satellite signal by tracking the C/A code at multiple correlator spacings [12]. In the IMT, the SQR receiver function is similar to SISRAD but is carried out by separate NovAtel Millennium receivers with multiple-correlator firmware. SAIM monitor receivers should include both SISRAD and SQR functions [1, 5].

Smoothing is performed on the raw pseudorange measurements using a modified first-order FIR filter that uses carrier phase to aid the smoothing process. This is often done inside GPS receivers, but we want to have control over the smoothing process along with doing integrity monitoring on raw pseudorange data. The filter applies the following two equations [1]:

$$PR_{sm}(k) = \left[\frac{1}{B}\right] PR_{raw}(k) + \left[\frac{B-1}{B}\right] PR_{proj}(k) \quad (1)$$

$$PR_{proj}(k) = PR_{sm}(k-1) + \theta(k) - \theta(k-1) \quad (2)$$

Here, B is either the number of epochs since the filter was reset or 200 (whichever is lower), since we would like our filter to have reached steady-state after 200 epochs. PR_{sm} is the smoothed pseudo range and PR_{raw} is the raw pseudo range from the receiver. θ is the carrier phase from the receiver whose pseudo range we are smoothing.

Measurement Quality Monitoring (MQM) performs checks on pseudorange and carrier phase separately. For pseudorange we perform what is commonly called an innovation test. That is, we calculate the raw pseudorange measurement for each epoch and subtract the projected pseudorange obtained by projecting the smoothing filter forward from the previous epoch. This difference is compared to a threshold that is driven by noise in the raw pseudoranges.

Clock monitoring is performed using the carrier phase. A second- order model is constructed from the last ten measurements received (with 2 Hz updates, 5 seconds of measurements are used). Step, ramp, and acceleration estimates are computed from the result using an innovation test (latest value minus value predicted from the fit of the previous epoch) for the step test statistic and the 1st and 2nd-order coefficients of the polynomial fit for ramp and acceleration, respectively. Code-carrier divergence is also monitored using a 200-second moving average to estimate the divergence rate.

Data Quality Monitoring (DQM) is responsible for monitoring navigation messages coming from the satellite. Checking the accuracy of the ephemeris and clock-correction messages is particularly important. All decoded navigation data is checked for transmission errors using the

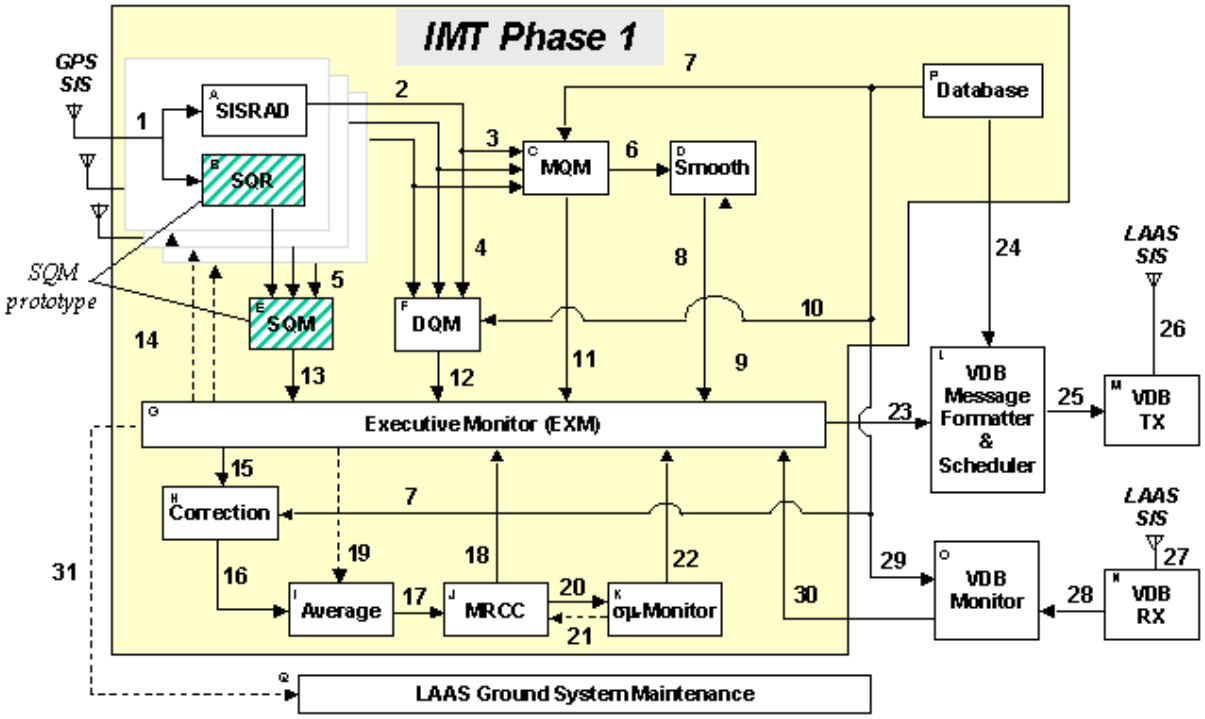

Figure 2: IMT Block diagram

parity checks built into the data words. When a new or updated ephemeris message is broadcast, we compare it to older messages to confirm that the latest data is reasonably close [13].

Executive Monitoring (EXM) uses the results from all of the tests so far and takes the appropriate action if faults are detected. As an example, if a certain failure is discovered on only one of the three reference receivers on more than one satellite, that receiver is excluded from further contribution to differential corrections. Another possibility is a failure that is discovered on a given satellite on more than one reference receiver. In that case, the affected satellite will be excluded (no corrections will be broadcast for it by LAAS).

EXM continues monitoring the excluded measurements and will "self-recover" if the fault is clearly determined to be over. In the IMT, this self-recovery process is attempted no more than two times. If the failure still shows up for that receiver, it is declared non-operational, pending external intervention and maintenance [1].

Once the integrity monitors described above have been acted upon, EXM constructs a common set of satellites. A satellite is in the common set if all receivers are tracking it with healthy measurement (it passes all of the integrity monitors so far). These satellites now enter the next phase of integrity monitoring.

Correction generation. Candidate pseudorange and carrier phase corrections are computed for each satellite and receiver by subtracting the theoretical range (based on the broadcast satellite ephemeris and clock navigation data) from the measured range. This is the central step in Differential GPS – when users apply the corrections, ranging errors that are correlated between reference and user receivers are canceled out.

Reference-receiver clock adjustment is then performed using the candidate corrections of the satellites in the common set. After this step, the corrections are approximately zero-mean. The correction that will be broadcast for a given satellite is the average of the clock-adjusted candidate corrections from each receiver tracking that satellite, pending approval from the two sets of monitor algorithms described next and confirmation that the averaged correction falls within acceptable bounds [1].

Multiple Reference Consistency Check (MRCC) compares the candidate pseudorange and carrier phase corrections across the three different receivers. This is done using "B-values" that give a numeric representation of how well a given receiver's measurement compares to what the other receivers measured for the same satellite.

The following example illustrates how B-values are calculated. Reference receiver 1 provides a candidate clock-adjusted pseudorange correction of 10 m for a given satel-

lite, receiver 2 gives a correction of 11 meters, but receiver 3 gives a correction of 30 meters. In this case, the correction of receiver 3 appears erroneous. If receiver 3 has in fact failed, the correction error that would result is the difference between the averaged correction from all three receivers (17 m) and the averaged correction from (healthy) receivers 1 and 2 (10.5 m), or 6.5 meters. This is defined to be the B-value for receiver 3 on that satellite.

B-values for each satellite and receiver are computed and checked against thresholds driven by code-phase multipath, which should be independent among the reference receivers. If one or more B-values exceeds its threshold, a set of logical steps is undertaken in the second phase of EXM (see below) to determine which measurements are erroneous, and then the second phase of processing must be repeated, starting with "correction generation" [1].

The $\sigma - \mu$ **Monitor** uses B-values as inputs and estimates their means and standard deviations, ensuring that they are within specifications. In addition, Cumulative Sum (CUSUM) monitors are applied, as described in [14]. These monitors are designed to catch small violations of the expected mean error in the corrections (zero) and the error standard deviation broadcast by the LGF [15]. Such violations may not be apparent immediately: the length of time required for detection is generally inversely proportional to the degree of increased integrity risk. Significant increases in integrity risk will almost always be detected within 30 minutes [14].

Executive Monitoring (EXM) now takes the results from all of the second phase of tests and makes a final decision on approving or rejecting each satellite and receiver. A satellite is retained in the common set if it is seen by all receivers and passes all of the integrity tests. Users listening to the LAAS VHF data broadcast can then difference the corrections computed by the IMT from their measured pseudoranges to improve positioning accuracy. Satellites for which corrections are not broadcast are considered to be unhealthy and cannot be used.

3.2 Adapting the IMT for SAIM

The IMT is designed to operate from the ground; thus some modifications are needed to apply it to SAIM. When a monitor receiver antenna is located on a GPS satellite, that satellite will most likely be the only satellite that receiver can track because the received power from that satellite on which the receiver is installed is many orders of magnitude greater than that from other satellites. With only one satellite to work with, the receiver cannot acquire a position fix nor directly solve for the receiver clock offset. Thus, our pseudorange and carrier-phase measurements will have a large user clock drift in them.

In order to closely monitor the received satellite signals, it is necessary to carefully remove this receiver clock drift. We do not want to be too aggressive in doing so, since we want to be still able to catch harmful ramp errors in the satellite itself. Note that, in SAIM, there is no need to generate pseudorange corrections (box H in figure 2) for users to apply. Correction calculations are retained as "by-products" of IMT-like processing.

4 Receiver Clock Drift Removal

The monitor receiver's ability to keep track of time is handicapped by the fact that typical receiver oscillators have a difficult time staying exactly fixed on the desired frequency. However, given stable environmental conditions, the oscillator frequency typically remains slightly off by a nearly-constant amount.

The receiver clock frequency can be written as:

$$f_{Rcvr} = f_{Des} + \Delta f_{Rcvr} \tag{3}$$

where Δf is the receiver's frequency deviation in Hz and f_{Des} is the desired ocillator frequency. Two things make up the essentials of a clock: a frequency oscillator and a counter. Assuming that our counter is ideal, the receiver's estimate of GPS time can be written as [2]:

$$T(t) = \frac{f_{Rcvr}}{f_{Des}} t + T_{Start} = t + \frac{\Delta f_{Rcvr}}{f_{Des}} t + T_{Start} \tag{4}$$

where $T(t)$ is our estimate of GPS time in seconds, and T_{start} is our initial guess of GPS time. We will assume that Δf is fairly constant (for crystal oscillators, this is a valid assumption if the temperature stays constant); thus the time drift will be linear [2].

It is possible to examine this supposition with measured pseudorange drifts from existing receivers. In a simple experiment, a NovAtel Millenium GPS receiver is connected to a WelNavigate single-channel GPS simulator (see figure 8 for a illustration of this setup). The pseudorange and carrier phase reported by the receiver are logged. The results can be seen in figure 4.

As theory predicts, figure 4 is a straight line, and the measured drift is aproximetely 2100 km/hr. If our monitor receiver were to use a crystal oscillator on L1, this would correspont to a frequency deviation of only 3 kHz, which

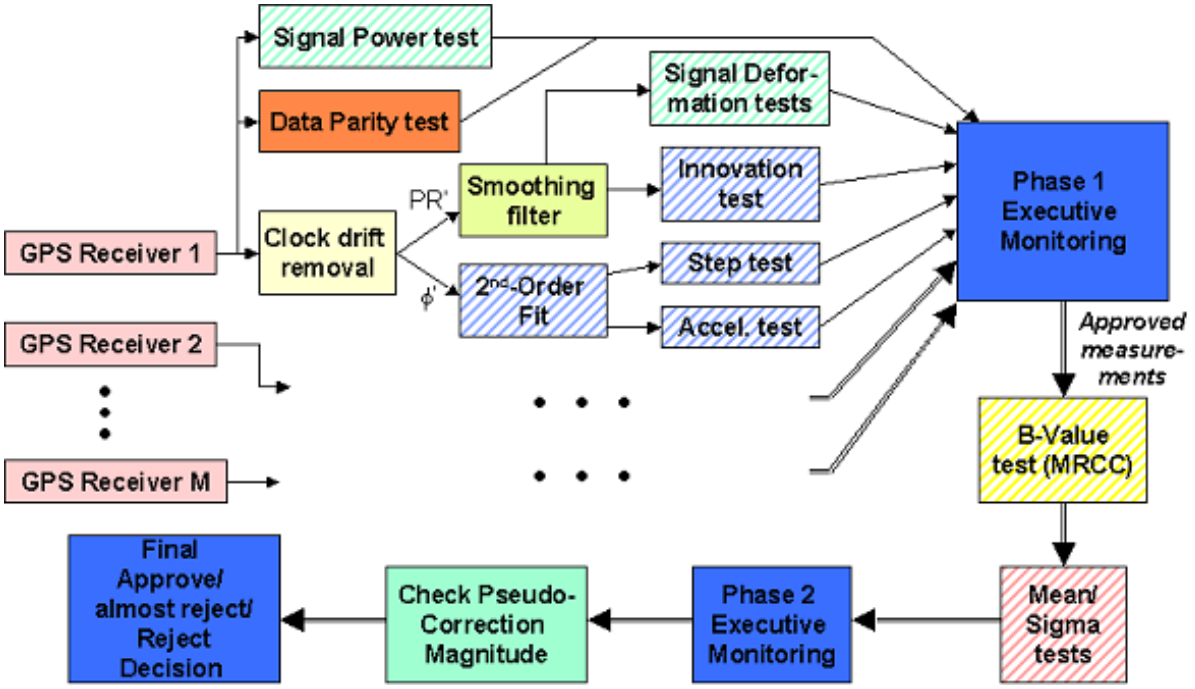

Figure 3: A detailed SAIM Block diagram

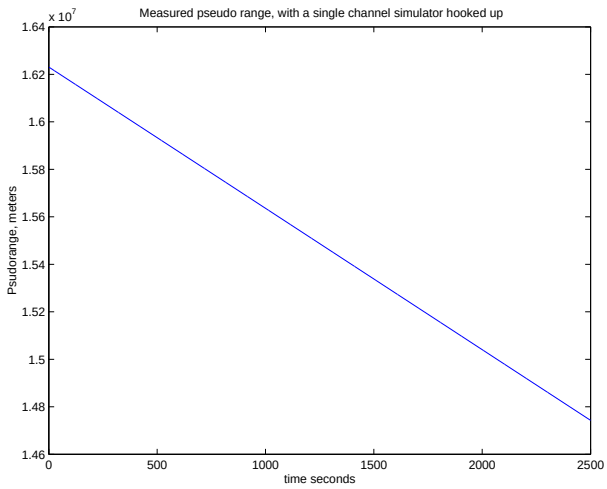

Figure 4: Receiver clock drift: roughly 2100 km/hr. Data is generated using the measurement setup in figure 8. Note that the data was shifted such that the first pseudorange was 0.

is a deviation expected nominally from Doppler effects for a ground- based user. This drift is much larger than our true pseudorange measurement deviations (due to nominal measurement errors for a fixed receiver) of several meters. Even when more-advanced oscillator types are used, the drift rate still stays fairly high. Measurements taken using a rubidium oscillator as a reference oscillator for the NovAtel Millennium GPS receiver revealed that, relative to the GPS constellation time, the rubidium oscillator still created a clock drift of roughly 6.7 $\mu s/h$ that translates into 2 km/hr of pseudorange drift.

Again considering equation 4, the receiver's time estimate error will be:

$$T_{err}(t) = T_{GPS} - T_{rcvr} = \frac{\Delta f_{Rcvr}}{f_{L1}} t + T_{Init} \qquad (5)$$

Converting this into pseudorange error (multiplying with c), the user clock drift is obtained:

$$D(t) = \underbrace{c\frac{\Delta f_{Rcvr}}{f_{L1}}}_{A} t + \underbrace{cT_{Init}}_{B} \qquad (6)$$

Since the true pseudorange is constant for a fixed receiver, it should be included in B. Hence, we can perform "IMT-like" correction calculations while removing the user clock drift. Given the linearity of drift, a least-squares method is suitable. It is known that A and B do not change that much over time and that the drift term will dominate pseudorange measurements. Hence it is possible to take the past values of the measurements and fit them to a least-squares line. This line is projected forward to the current epoch and subtracted from the current epoch'ss measurement, removing the drift and performing the correction calculation. This approach has been tested in Matlab by using the last

10 epochs to construct a least-squares linear fit. The difference between the linear prediction and current pseudo range is shown in the following figure.

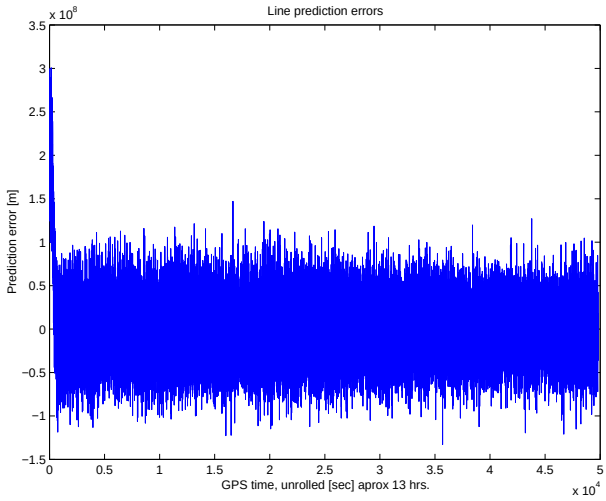

Figure 5: Filter clock prediction errors using the least squares method

One problem with using the least-squares linear fit is the computational cost of redoing it every epoch. However, by using a Kalman filter, it is possible to retain similar clock drift removal properties for a fraction of the computational cost (in fact, Kalman filters may be more accurate since they take all of the past data into consideration).

The Kalman filter is an optimal solution of the general linear estimation problem. Consider a linear process of the form:

$$X(n+1) = AX(n) + BU(n) + W(n) \tag{7}$$

where A, B are known constant matrices, and W is a Gaussian random noise vector with known mean and variance. A Kalman filter provides the best RMS estimate of X(N+1) given all past values of C X(N) where C is a constant matrix.

Transferring this general solution to the drift estimate problem, define a two-state vector X containing the current clock drift and drift rate. This system is assumed to be input-free for the time being, although clock temperature could be added as an input. The A matrix is given by [4]:

$$A = \begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix} \tag{8}$$

The statistical processes of the Gaussian random vector W are modeled in the form of a covariance matrix. The formula (see [4] for details) used for calculating the covariance matrix for $W(N)$ is given by [4]:

$$Q = \begin{bmatrix} s_b \Delta t + s_f \frac{\Delta t^3}{3} & s_f \frac{\Delta t^2}{2} \\ s_f \frac{\Delta t^2}{2} & s_f \Delta t \end{bmatrix} \tag{9}$$

where s_b and s_f are derived from Allan variances h_0 and h_{-2} (properties of the oscillator) by the following formulas [4]:

$$s_f = 2h_0 \tag{10}$$

$$s_b = 8\pi^2 h_{-2} \tag{11}$$

Further details of Kalman filters can be found in [3].

The Kalman estimator can be applied in a similiar fashon as the least-squares fit – the Kalman estimator generates this epoch's pseudorange value instead of the least-squares line. As was done for the least-squares fit, the difference between the Kalman estimate for the current epoch and the actual measured value for this epoch gives the estimation error for the current epoch (presented in figure 6).

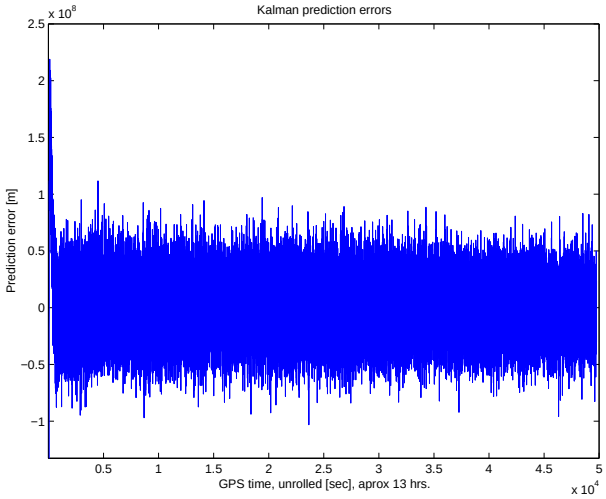

Figure 6: Kalman filter clock prediction errors

By comparing figures 5 and 6, we can see that the performance of the Kalman filter is slightly better than that of our least-squares estimator. The real gain, however, is in the severely reduced computation power. Informal measurements (a tic-toc pair in Matlab) revealed that a dual-Pentium II/400 processor machine required 550 seconds to execute the least squares estimate for a 13-hour dataset

but only 17 seconds to process the Kalman estimate. Reduced computational load provides the ability to have several Kalman filters running in parallel for the same computational cost. This is quite useful, since it is now possible to get broader coverage of the error space.

For example, if we set the feedback frequency to its maximum (feedback every epoch), estimation noise will go down, but the filter will adapt quickly, hence when a small harmful ramp is introduced the MQM part will fail to detect it. Thus, it is a good step error detector but does not detect ramp errors that well. The feedback frequency could also be set to a somewhat slower rate. This increases estimation noise, yielding higher thresholds. However, it does not adapt as quickly, so its strengths and weaknesses are reversed from the previous case (the slower-feedback filter is a better ramp detector but is not as good in detecting step errors).

A solution to this dilemma is to use two or more Kalman filters running in parallel, each removing estimated receiver clock biases feeding the subsequent monitor algorithms. As suggested above, one would have a high feedback frequency (aimed at step errors), and another one would have a lower feedback frequency (aimed at ramp or other slowly-changing errors).

5 Matlab SAIM Prototype

Using the IMT model described above along with receiver clock estimation, a SAIM software prototype has been developed in Matlab. In this first version, the signal quality monitors (SQM and SQR), $\sigma - \mu$ and the DQM monitors are not yet implemented.

As mentioned earlier, the prototype is to have several Kalman-MQM sets running in parallel (looking at the block diagram in figure 3 a Kalman-MQM block is everything but the receiver part and the SQM/SQR blocks). A revised block diagram of the SAIM prototype is shown in figure 7.

To derive detection thresholds for each monitor (after receiver clock removal), we examined the test variable's mean (if non-zero) and standard deviation from several sets of nominal data. After removing any mean bias, multiplying the standard deviation by a "K-value" between 6 and 7 to yield the threshold. Assuming that the test variable is Gaussian or is overbounded by a Gaussian distribution with the measured standard deviation, the probability of a false alarm (a threshold being crossed under nominal conditions) is only 10^{-9} or lower, which is well below existing continuity specifications for civil aviation (note that the overall loss-of-continuity probability, which is on the order of 10^{-5} per operation, must be divided among many monitor algorithms and parallel filters). In the current SAIM prototype, two Kalman-MQM blocks are used. The one aimed at detecting step errors feedbacks every epoch, and the one focused on slower-changing errors feedbacks every 10 epochs.

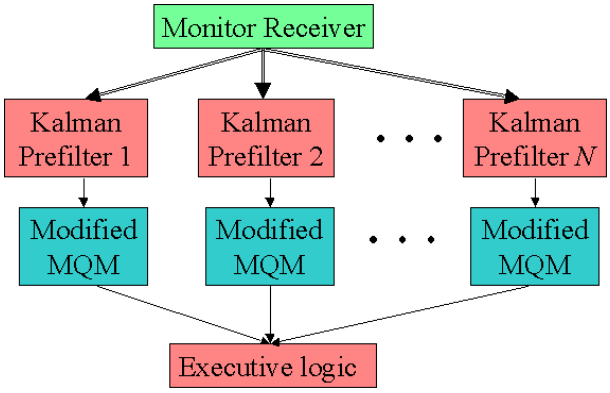

Figure 7: SAIM block diagram as used in the Matlab prototype, see figure 3 for a more detailed illustration of a Kalman-MQM block

In order to test our Matlab implementation and to derive nominal test thresholds as described above, a WelNavigate single-channel GPS simulator (GS-100) was configured to give a nominal GPS signal, which was fed into a NovAtel OEM3/Millennium GPS receiver, logging the carrier phase (in cycles) and pseudorange (in meters) every second approximately 13 hours. A laptop PC with Windows 98 was used to log the data.

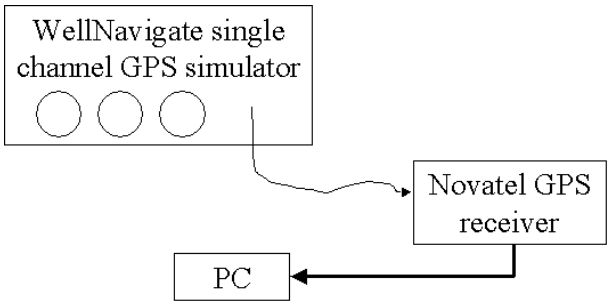

Figure 8: Measurement setup, using a WelNavigate GS-100 single channel simulator feeding a NovAtel Millennium CA/Code receiver

In order to simulate the presence of three receivers in SAIM, three independent 40-minute (2500 seconds) data fragments were extracted from the 13-hour set. To make the test more realistic, the sign of one of the pseudoranges was flipped (to simulate the drift of a different receiver).

5.1 Injecting a step error

The SAIM prototype was failure-tested by injecting a carrier phase step error. This was done purely in software by adding the step to all measurements on all receivers after a given epoch.

Since the Kalman filter requires about 100 epochs to settle, the first 200 epochs of the prefiltered data are removed before passing it to the IMT. The deviation of the current carrier phase measurement from the forward-projected second order model is examined (see the IMT MQM description in Section 3.1). The results can be viewed in figure 9.

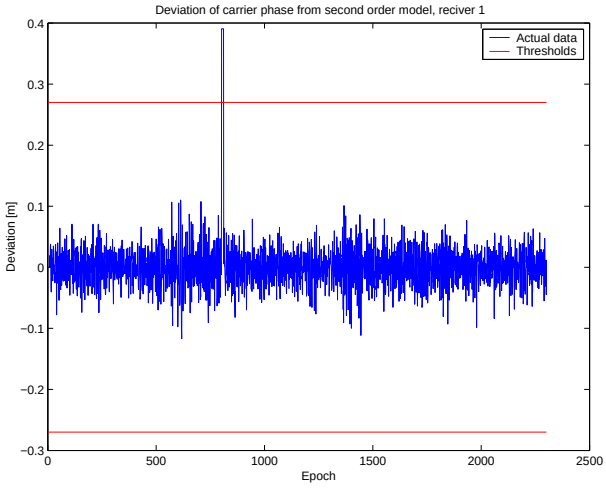

Figure 9: Carrier phase deviation from our second order model, here using the faster Kalman pre-filter. The injected step error should have an effects around the 800-epoch mark.

There is indeed a spike occurring around 800 seconds into the measurements in figure 9. Zooming in shows that the MQM component catches this error immediately. Figure 10 shows the same IMT variable for the slower Kalman filter.

As expected, the high test statistic noise requires a high threshold, thus the step error is not detected. It is possible by zooming in to see a small impact of the step error around the 800 epoch mark; however, the resulting spike is at nominal levels and does not cross our thresholds.

5.2 Injecting a ramp error

In this test, a 0.65 m/sec ramp error is injected into all of the pseudorange measurements after the 1000-epoch mark (one epoch is one second). Again, the Kalman filter needs 100 seconds to settle, so the first 200 seconds of prefiltered data are removed before continuing to the IMT phase. It

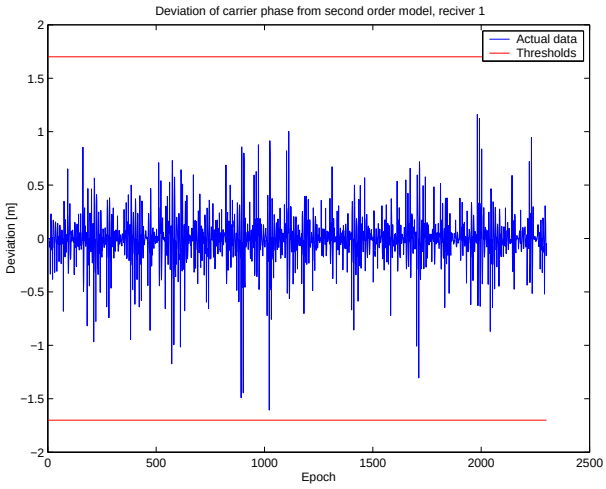

Figure 10: Carrier phase deviation from a second order model (see Section 3.1 for details), here using the slower Kalman prefilter. Again, the injected step error should take effect around the 800 epoch mark.

is interesting to note that this ramp error is significantly smaller than the one caused by a clock anomaly discovered on SVN 22 on July 28, 2001 [16].

The monitor designed to catch large pseudorange errors is the innovation test. It compares the forward-projected smoothed pseudorange to the current raw pseudorange (see Section 3.1 for more details, and also refer to [1] and [5] for a more complete description). For the fast Kalman filter, the results are in figure 11.

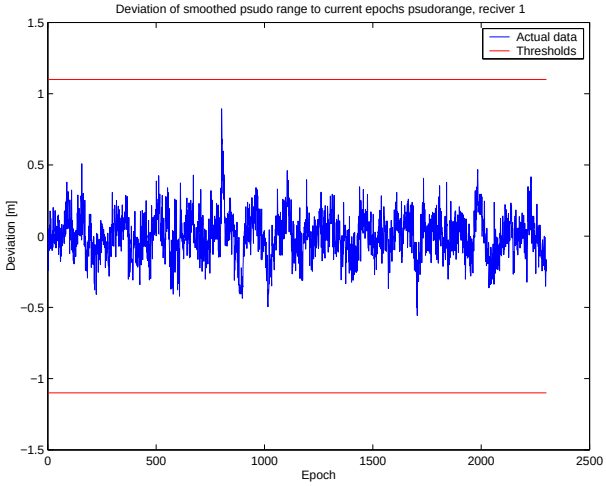

Figure 11: Raw pseudo range deviation from our smoothed pseudorange, here using the faster Kalman prefilter. The injected ramp error should appear around the 800-epoch mark.

A spike around the 800-epoch mark is found as expected. However, that spike is not large enough to trigger the

threshold. The fast Kalman filter simply adapts too fast so that the monitor does not have a chance to detect the ramp error. We now observe the same IMT test variable on the slower Kalman filter. These results are displayed in figure 12.

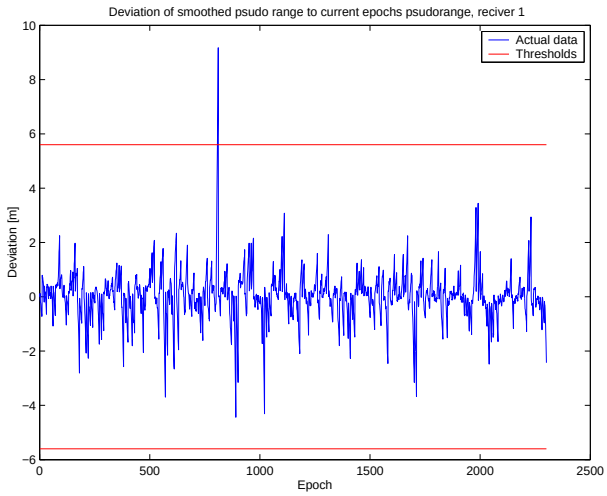

Figure 12: Raw pseudo range deviation from our smoothed pseudorange, here using the slower Kalman prefilter (0.1Hz)

Although the thresholds are about five times higher, the spike caused by the injected ramp is about ten times higher so here IMT has no trouble detecting the error. However it is important to note that the spike does not look like a ramp (it does not continue to infinity). There are two forces at work. First we have the Kalman filter designed to cut out the large ramp caused by the user clock drift via feedback. The Kalman filter will eventually adapt to the change in the slope of the ramp. Second, the pseudoranges are put through a smoothing filter and that filter too will eventually catch up.

6 Conclusions and future work

It has been shown that SAIM is a feasible, cost-effective solution to achieve tightened integrity requirements for all GPS users while reducing the degree of augmentation needed by civil-aviation users. The results of SAIM prototype testing have partially demonstrated its ability to rapidly detect errors that have occurred or could occur onboard GPS satellites.

The Matlab SAIM prototype serves as a tool to aid research in SAIM since one can easily try out different algorithms and combinations of Kalman filters for monitor receiver clock removal and supports extensive failure testing. However, much work still remains to be done. For one thing, an interface between SAIM outputs and satellite onboard systems needs to be developed (remember that the SAIM concept advanced here is a modular add-on to existing GPS satellite designs). In addition, a thorough study is needed into the effects of receiving GPS signals from a receiver located on the side of the transmitting GPS satellite. Local multipath (reflections of other satellite equipment such as solar cells) will affect our received signal quality and thus impact the detection thresholds we can set with acceptable false-alarm rates. In addition, we must confirm that monitoring the sidelobe allows SAIM to detect all possible faults in the main lobe of the signal.

Finally, as mentioned above, for more-thorough coverage of the possible satellite failure space, it is better to have several Kalman filters running in parallel with different settings to remove the receiver clock bias. It still remains to do an optimization of how many Kalman filters are needed and what their parameter settings should be to best cover the failure space.

Acknowledgments

The authors would like to thank all the great people of the Stanford University GPS lab for their professional support and advice during this research, the FAA Satellite Navigation Program Office for funding this research, and NavAstro Co. for its interest and innovation. The opinions discussed here are those of the authors and do not necessarily represent those of the FAA or other affiliated agencies.

References

[1] G. Xie, S. Pullen, M. Luo, et al., "Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed," *Proceedings of ION 57th Annual Meeting and CIGTF 20th Biennial Guidance Test Symposium*. Albuquerque, NM., June 11-13, 2001, pp. 681-693.

[2] D.W. Allan, N. Ashby, C.C. Hodge, *The Science of Timekeeping*. Available from `http://www.allanstime.com/Publications/DWA/Science_Timekeeping/` 1997

[3] G. Welch and G. Bishop, *An Introduction to the Kalman Filter*. Avaliable from `http://www.cs.unc.edu/~welch/publications.html` 1995.

[4] B. Parkinson and J. Spilker, Eds., *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996. Vol. I, Ch. 9, pp. 417- 418.

[5] M. Luo, S. Pullen, B. Pervan, et.al., "GBAS Validation Methodology and Test Results from the Stanford LAAS Integrity Monitor Testbed," *Proceedings of ION GPS 2000*. Salt Lake City, UT., Sept. 14-17, 2000, pp. 1191-1201.

[6] T.M. Corrigan, et.al., *GPS Risk Assessment Study: Final Report*. The Johns Hopkins University Applied Physics Laboratory, VS-99-007, January 1999. Avaliable from `http://www.jhuapl.edu/transportation/aviation/gps`

[7] *Global Positioning System Standard Positioning Service Signal Specification*. U.S. Department of Defense, DOD 4650.5/SPSSP V3, 3rd Edition, August 1, 1998. Annex B, Section 4.0, pp. B-13 - B-16.

[8] K. VanDyke, "Use of Standalone GPS for Approach with Vertical Guidance," *Proceedings of ION 2001 National Technical Meeting*. Long Beach, CA., Jan. 22-24, 2001, pp. 301-309.

[9] D. Loverro, "GPS Modernization," *Proceedings of ION 57th Annual Meeting and CIGTF 20th Biennial Guidance Test Symposium*. Albuquerque, NM., June 11-13, 2001, pp. 1-32.

[10] *GPS OCS Performance Analysis and Reporting (GOSPAR) Project: Final Report*. Colorado Springs, CO.: Overlook Systems Technologies, Inc., Version 2, Sept. 30, 1996.

[11] C. Edgar, F. Czopek, B. Barker, "A Co-operative Anomaly Resolution on PRN-19," *Proceedings of ION GPS 1999*. Nashville, TN., Sept. 14-17, 1999, pp. 2281-2284.

[12] R.E. Phelts, D.M. Akos, P. Enge, "Robust Signal Quality Monitoring and Detection of Evil Waveforms," *Proceedings of ION GPS 2000*. Salt Lake City, UT., Sept. 14-17, 2000, pp. 1180-1190.

[13] S. Pullen, B. Pervan, et.al., "Ephemeris Protection Level Equations and Monitor Algorithms for GBAS," *Proceedings of ION GPS 2001*. Salt Lake City, UT., Sept. 11-14, 2001 (forthcoming).

[14] J. Lee, S. Pullen, G. Xie, P. Enge, "LAAS Sigma-Mean Monitor Analysis and Failure-Test Verification," *Proceedings of ION 57th Annual Meeting and CIGTF 20th Biennial Guidance Test Symposium*. Albuquerque, NM., June 11-13, 2001, pp. 694-704.

[15] B. Pervan, I. Sayim, S. Pullen, "Sigma Estimation, Inflation, and Monitoring in the LAAS Ground System" *Proceedings of ION GPS 2000*. Salt Lake City, UT., Sept. 14-17, 2000, pp. 1234-1244.

[16] S. Pullen, T. Walter, J. Warburton, "A Quick Look at Last Weekend's SVN/PRN-22 Clock Anomaly." Stanford University, Unpublished Presentation, Version 1.1, Aug. 3, 2001.