International Global Navigation Satellite Systems Association
IGNSS Conference 2016

Colombo Theatres, Kensington Campus, UNSW Australia
6 – 8 December 2016

Improving Sensitivity on Kea CubeSat GPS Receivers

Eamonn P Glennon
The Australian Centre for Space Engineering Research (ACSER)
School of Electrical Engineering and Telecommunications, UNSW, Australia
Phone: +61 2 9385 6702, Email: e.glennon@unsw.edu.au

Andrew G Dempster
The Australian Centre for Space Engineering Research (ACSER)
School of Electrical Engineering and Telecommunications, UNSW, Australia
Phone: +61 2 9385 6890, Email: a.dempster@unsw.edu.au

ABSTRACT

The Namuru and Kea FPGA based GPS receivers have been developed by UNSW and its partners for over a decade. During this period, the hardware and firmware have undergone continuous improvement, with the introduction of L1 carrier phase, inclusion of orbit and high dynamics modes, and disciplined oscillator and timing outputs being cases in point. This paper describes the most recent set of upgrades related to improving receiver sensitivity.

The use of a noise code-phase tap, in addition to the traditional early, prompt and late code-phase taps, is described here, as are the benefits of such a design. Different options concerning the use of the noise tap are discussed. One involves measurement of 1 ms dump statistics and calculation of the statistics of any pre-detection integration (PDI) with additional non-coherent (NC) accumulations based on a mathematical model. Others involve direct measurement of PDI and NC accumulation mean and standard deviations, as well as direct measurement of PDI and NC distributions. It is explained how direct measurement of the noise distribution can be performed efficiently, thereby permitting a threshold to be produced that achieves a given false alarm rate (FAR).

Also described are the modifications to the tracking loops necessary to support the requirement for additional sensitivity, while still maintaining stability. Finally, the need to maintain and achieve bit-synchronisation with reduced signal levels is also discussed. Matlab simulation results comparing the various methods are also presented.

KEYWORDS: GPS, receiver, cubesat, detection, threshold, weak-signal, noise-finger, noise-tap

1 INTRODUCTION

Operating a GPS receiver in space is challenging. Mass and volume constraints imposed by launch, as well as environmental concerns such as vacuum and radiation place demands on the electronics that few terrestrial receivers are required to meet. The dynamic environment of both launch and orbit creates challenges for tracking and acquisition that are an order of magnitude greater than required for most commercial of the shelf (COTS) receivers. Some applications, such as radio-occultation, involve tracking weak signals caused by fading as the signals pass through by the earth's atmosphere and ionosphere. Another application, known as GNSS-Reflectometry, involves detection of weak reflections from the earth's surface. A rarer application involves the use of GPS in high earth orbit (HEO), in which case the receiver operates above the orbits of the GPS satellites themselves. Space based GNSS receivers are also subject to severe power limitations, as is the case for many terrestrial based receivers. In the case of receivers installed on orbiting spacecraft, the receivers are required to run entirely on power generated from solar panels and stored in rechargeable batteries. However, unlike most commercial receivers, the flight models are required to have operate at speeds that exceed the 500 m/s limitations that most commercial receivers are required to enforce in order to satisfy ITAR and Wassenaar regulations. These limitations are one reason that most commercial chipsets are not easily employed for space applications.

UNSW, through the Satellite Navigation And Positioning (SNAP) group, the Australian Centre for Space Engineering Research (ACSER) and its partners have been developing FPGA-based GPS receivers since 2003. This development reached a level of maturity in 2011 that was demonstrated by Defence Science and Technology (DST) Group's request to ACSER to develop, build and supply four flight model GPS receivers for its Biarri project. This resulted in the delivery of four Namuru V3.2R3A receivers, along with several early prototypes used for development (Glennon et al., 2011; Mumford et al., 2012). This work has continued with the development of the more advanced Kea V41SBR3 receiver, which improves on the Namuru V32R3A with higher component integration and lower power consumption, as well as providing support for the CalPoly CubeSat standard. In addition to the improvements made to the hardware, firmware improvements have also been made that further enhance the receiver's ability to operate in space. The most recent of these improvements relate to receiver sensitivity and are the subject of this paper.

Section 2 of this paper provides an overview of the detection and tracking of weak GPS signals. Section 3 reviews the often-neglected aspect of tracking sensitivity, namely the setting and updating of the detection thresholds. Section 4 shows some experimental results showing between each of the methods using a Matlab model, while Section 5 describes some tests with actual Kea V41SBR3 receivers and a Spirent STR8000 GNSS simulator.

2 HIGH SENSITIVITY GPS RECEIVERS

2.1 Overview

GPS C/A code signals are spread spectrum signals transmitted at a frequency of 1.57542 GHz by satellites in medium earth orbit at an altitude of over 20,000 km. The system guarantees a minimum received signal power for the L1 C/A code signal of -160 dBW (-130 dBm) at the input of a 3 dBi linearly polarised antenna at the earth's surface. Such signal levels are easily tracked by a conventional GPS receiver in open sky conditions with a good antenna and a well-designed RF front end. However, in cases where the signal levels are significantly

attenuated, such as for indoor GPS or for GPS receivers in HEO, additional signal processing techniques are required if the signals are to be tracked and used. These signal-processing techniques, which were originally developed to permit the use of GPS in cellular phones located in indoor environments, are described in publications such as (van Diggelen, 2009).

A detailed description of these techniques is beyond the scope of this paper. That said, all techniques involve performing additional integration to increase the signal-to-noise ratio (SNR) of the detection metrics. Additional pre-detection integration (PDI) on the in-phase (I) and quadrature phase (Q) quantities, also known as coherent integration, allows the variance of the I and Q components to be reduced while still maintaining a zero mean for the outputs (assuming zero mean inputs). Such coherent integration is equivalent to a box-car low-pass filter on each component, which reduces the bandwidth of the input signals, but also retains the signal phase-information. Each doubling in the coherent integration period increases the output SNR by 3 dB, while halving the bandwidth. The maximum amount of coherent integration is limited by the bit width of the navigation data, which is 20 ms for the C/A code navigation message. Coherently integrating for longer than 20 ms runs the risk of accumulating across data-bit boundaries, which if accompanied by a change in the data-bit sign will reduce the SNR rather than increase it. If the location of the data-bit boundaries is unknown, then coherently integrating for 10 ms, 5 ms or 4 ms is preferable as this reduces the number of accumulations that may be susceptible to data-bit transition losses to no more than 1 of 2, 1 of 4 or 1 of 5, respectively. Although sensitivity is maximised by increasing the coherent integration period to 20, or possibly more than 20 if data wiping can be performed, this comes with several costs. One of these costs is an increase in the search times that are caused by the reduced bandwidths of each search. However, another significant cost is the reduction in carrier tracking dynamics that are associated with increased PDI. This is especially true for phase-locked-loops (PLLs), which are less robust than frequency-locked-loops (FLLs). The problem of reduced dynamics is of more concern to receivers required to operate in space, as such receivers are subject to acceleration, albeit predictable acceleration.

When sensitivity exceeding the levels achievable using 20 ms coherent integration is required, it is necessary to perform non-coherent integration of the coherently integrated I and Q components. This involves taking the magnitude or magnitude-square of the complex quantity and then accumulating those outputs. When the local PRN code generator is not aligned with the incoming signal, the I & Q components will both be approximately zero mean and normally distributed, but once the magnitude has been taken, the outputs take on a Rayleigh distribution. This has a non-zero mean, which means that each doubling in integration period no longer delivers 3 dB of additional sensitivity. This loss of sensitivity can be attributed to the so-called 'squaring-loss', which is input SNR dependent. Lower input SNRs correspond to larger squaring losses.

2.2 Applicability to the Namuru V32R3A GPS Receiver

The Namuru V3.2 GPS receiver developed for the DST Group Biarri project employs a coherent integration period or PDI of 4 ms when tracking, each of which are bit-aligned. Pairs of PDI I & Q samples are processed by a third order PLL with second order FLL assist (Kaplan & Hegarty, 2006), with the pair sum being used to drive the PLL and the pair difference driving the FLL. Both the PLL and FLL are therefore updated every 8 ms. Detection processing during tracking is performed using 8 rounds of non-coherent integration, thereby resulting in a total integration period of 32 ms. The search algorithm employs a Tong search algorithm that spends at least two PDI intervals per search cell if no non-coherent integration is being employed. This ensures that detection will occur even if the PDI samples

straddle bit boundaries. Some modes required to search extremely wide bandwidths, such as those needed for LEO, employ PDIs of 1 ms thereby increasing the bandwidth of each search cell to 1 kHz. This reduces the TTFF, but at the expense of detection sensitivity.

When these design choices were made, it was for the purposes of allowing the tracking loops to operate with a wide range of dynamics, while supporting a reasonably wide bandwidth of 250 Hz during normal acquisition. However, using a 4 ms PDI with 8 non-coherent rounds of integration does incur a cost of reduced sensitivity. The 8 ms FLL and PLL update rate also causes additional processor load compared to the standard technique of employing 20 ms bit aligned PDI integration during tracking.

2.3 Improving Kea V41SBR3 GPS Receiver Sensitivity

To improve the tracking sensitivity of the Kea receiver, it was first necessary to add support for 20 ms bit aligned PDI. This involved more complexity than might have been expected because in order to allow the acquisition, capture and tracking layer to operate as before, it was necessary to modify the (higher priority) layer performing the PDI such that it could return both the full PDI, as well as the two half period PDIs making up the PDI. For example, selection of a 20 ms PDI would also result in the two half PDIs each of 10 ms being returned. These were necessary to allow the FLL portion of the tracking loop to operate, while the 20 ms PDI output allowed the PLL to update every 20 ms, as well as allowing non-coherent integration of those same 20 ms PDIs.

With these modifications, the SNRs available to both the detection algorithm and the PLL and FLL are improved by 3.98 dB, while the higher SNR present at the start of non-coherent integration also serves to reduce any squaring-losses that might be applicable.

One of the consequences of using a 20 ms PDI is the need to adaptively change the PDI depending on the bit-synchronisation status. To avoid a complete loss of SNR should the PDI straddle a bit-boundary evenly, it becomes necessary to introduce the concept of using different PDIs and non-coherent accumulation during the acquisition and tracking processes, respectively. These values are selected by the satellite selection, depending on the mode of the receiver, and allow the user to tailor the acquisition and tracking sensitivities to each mode. For example, low earth orbit (LEO) mode emphasises search speed over sensitivity, while HEO mode emphasises sensitivity. Allowing for this change was relatively easy, whereby whenever a 20 ms PDI is selected for the tracker, a check to see whether bit-synchronisation has been achieved is also made. In the event that bit-synchronisation has not been achieved, the PDI is halved and the number of associated non-coherent rounds is doubled to partially compensate. Additionally, the receiver then goes back to operating on pairs of 10 ms PDIs in the same way that the Namuru V32R3A does. Once bit-synchronisation has been achieved, the values revert to 20 ms and fewer non-coherent rounds, and the new code to employ these settings is re-activated.

The other issue that needed to be addressed was carrier tracking loop stability when employing the larger PDI values. To account for this, it was necessary for the firmware to check the value of the PDI and use different noise bandwidth values for the FLL and PLL values if 20 ms had been selected.

I/P Power (dBm)	I/P Power (dBHz)	PDI (ms)	NCR	Dwell (ms)	BW (Hz)	SNR (dB)	SNR (ratio)	Remarks
-135	39	4	8	32	250	29.1	815	Lowest sensitivity
-135	39	10	2	20	100	27.6	577	
-135	39	10	4	40	100	30.6	1153	
-135	39	10	8	80	100	33.6	2306	
-135	39	20	2	40	50	30.9	1226	
-135	39	20	4	80	50	33.9	2452	Highest sensitivity
-140	34	4	8	32	250	22.9	197	Lowest sensitivity
-140	34	10	2	20	100	21.9	154	
-140	34	10	4	40	100	24.9	309	
-140	34	10	8	80	100	27.9	618	
-140	34	20	2	40	50	25.4	345	
-140	34	20	4	80	50	28.4	691	Highest sensitivity
-145	29	4	8	32	250	15.8	38	Lowest sensitivity
-145	29	10	2	20	100	15.6	36	
-145	29	10	4	40	100	18.6	72	
-145	29	10	8	80	100	21.6	144	
-145	29	20	2	40	50	19.5	88	
-145	29	20	4	80	50	22.5	177	Highest sensitivity
-150	24	4	8	32	250	7.4	5	Lowest sensitivity
-150	24	10	2	20	100	8.2	7	
-150	24	10	4	40	100	11.2	13	
-150	24	10	8	80	100	14.2	26	
-150	24	20	2	40	50	12.8	19	
-150	24	20	4	80	50	15.8	38	Highest sensitivity

Table 1: Output SNRs for selected input signal power (dBm) processed with coherent integration period PDI (ms) and non-coherent integrations NCR. Note that the input SNRs do not include any receiver implementation losses and as such, approximately 2 dB has to be added in order to see the values that would be reported by the receiver.

2.4 Tuning for Improved Tracking Sensitivity

To determine the improvements expected from these changes, the sensitivity spreadsheets described in (van Diggelen, 2009) were employed. Using this process, sensitivity predictions given combinations of PDI and non-coherent accumulation can be made. Some configurations used in the Aquarius firmware are shown in Table 1, where it can be seen that using a

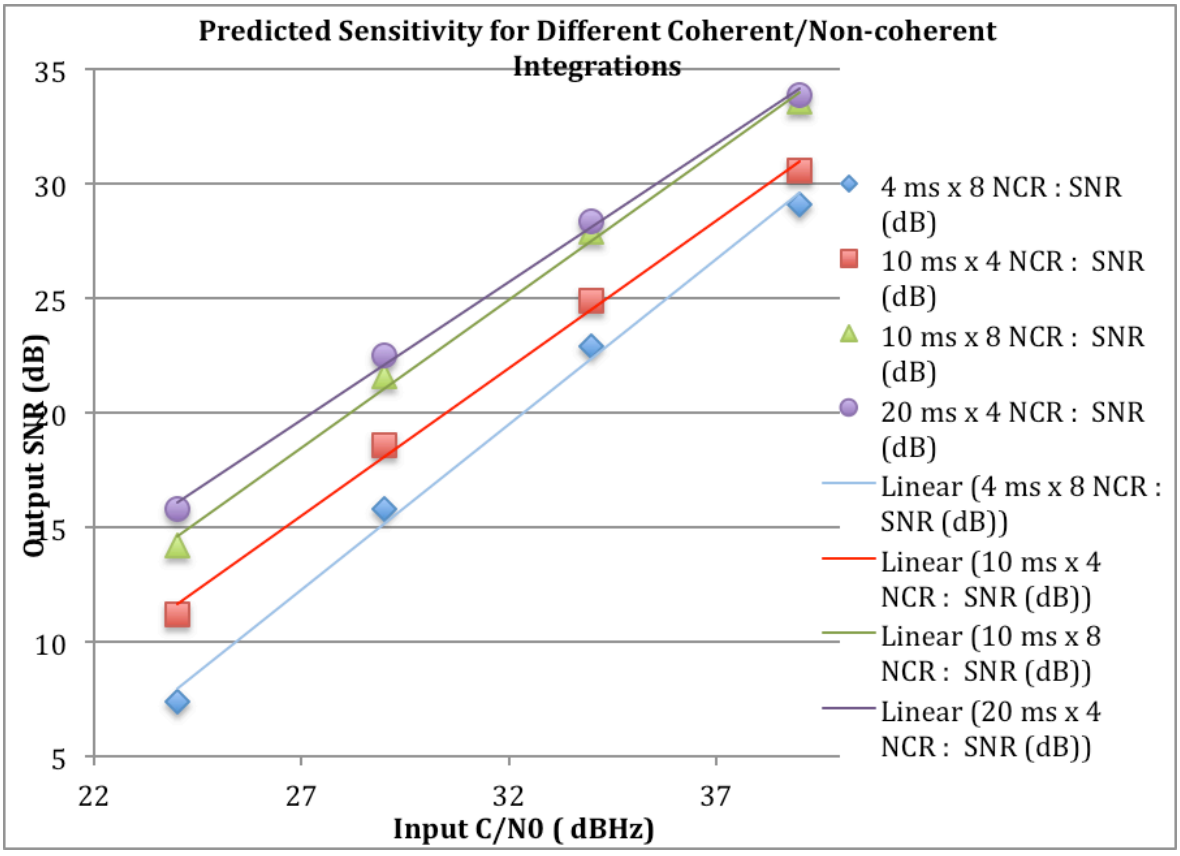

Figure 1: Plot of data from Table 1 showing clearly the improvement to be gained through the use of 20 ms x 4 NCR compared to the legacy setting of 4 ms x 8 NCR.

coherent integration of 20 ms and 2 or 4 non-coherent rounds, tracking of signal levels at -150 dBm should be possible provided the PLLs are able to operate at these levels.

2.5 Bit-Synchronisation at Lower SNR

One difficulty with reliance on 20 ms bit-synchronised coherent integration is the need to first achieve bit-synchronisation. As previously explained, initialising a channel to employ a 20 ms coherent integration period without being properly aligned can result in no detection should a bit transition occur half way through the integration process. One solution is to employ 10 ms coherent integration periods (IP) until bit-synchronisation has been achieved, because as explained by (Psiaki, 2001), that ensures that at least one of the two 10 ms intervals that occur in each 20 ms IP will be free of a data-bit transition. In the Aquarius firmware employed on the Kea receiver, 10 ms coherent with up to 4 rounds of non-coherent integration is employed during acquisition if higher sensitivity is required during acquisition.

Two different bit synchronization methods are employed by Aquarius. The legacy technique implements the histogram method of (van Dierendonck, 1996) in which a count of the number of sign changes in the data-bits are accumulated, modulo 20. This technique works well if the signal levels are strong, but can start to break down as the signal levels reduce. To support correct bit-synchronisation at lower levels, an alternative scheme in which the firmware accumulates bit-matched power across all 20 possible bit boundary locations has been also added (Psiaki, 2001). This metric serves to supplement and possibly supplant the histogram method employed at weaker signal levels.

3 ADAPTIVE THRESHOLDS

3.1 Analysis Limitations

One serious limitation of analysis used to perform sensitivity calculations of this type are the assumptions regarding the signal being processed. In this case, the assumptions are that each in-phase and quadrature channel can be regarded as zero mean Gaussian noise when there is no signal present, or the real and imaginary parts of a rotating phasor with additive Gaussian noise on each component when a signal is present. While such a model is an excellent simplification for a standard sensitivity receiver, the assumptions are not valid when searching for weak signals.

One difficulty, even for strong signals, is that C/A spreading codes employed by the GPS system are not completely orthogonal, either with themselves or with other GPS satellites. This causes artefacts, such as autocorrelation side-lobe peaks and cross correlation with other strong signals, becoming apparent as the receiver sensitivity is increased. Furthermore, any received inadvertent or self-interference of a sufficiently high power will result in 1 kHz line spectra being visible in channel detections, even though these are unwanted and effectively noise.

This becomes important because in order to track a signal with a single channel tracker, it is necessary to compare the outputs of the prompt correlator channel with a noise threshold. The process of coherent and non-coherent integration describes how to maximize the outputs of the prompt correlator channel in order to permit tracking to occur, but it is also necessary to maintain or possess a legitimate value for the detection threshold. For an AGPS application in which massively parallel correlation is performed, such a threshold is moot because the detector is looking for a single peak that is sufficiently above the noise level comprising all the other code phases. A detectability measure can be calculated on the entire correlation output vector, which will also involve calculating a noise level using the non-peak (and surrounds) correlations. However, such a measure cannot be employed on a single channel correlator because even with multiple noise taps, there are insufficient samples to be statistically significant. That said, without a noise tap/finger on each channel it is necessary to hard-code the threshold and that is required to account for the worst case in which very strong signals and possibly some self interference is present. This results in a larger threshold than is required thereby reducing sensitivity.

3.2 Measurement of 1 ms Dump Noise Variance

Having established the need for a correlator noise tap/finger to adaptively set the detection threshold; it remains to determine how such a noise tap/finger should be employed. One approach is to use the noise tap to measure statistics on the 1 ms dump samples. Here 'measuring statistics' amounts to calculating sums and sums of squares of the noise channel samples from which a variance and standard deviation for the noise can be produced. The firmware can then use this estimate to predict the mean and standard deviations applicable after coherent and non-coherent integration (Glennon, 2009). The expected output of the non-coherent accumulations can then be calculated and an appropriate threshold set. The attraction of this technique is that the 1 ms dump samples are output at a 1 kHz rate and statistics can therefore be rapidly accumulated. Any change in conditions, such as the appearance of CWI interference can therefore be rapidly identified and the threshold increased if needs be. The other advantage that this method offers is that the measured statistics can be used to calculate thresholds even if the channel settings for the coherent and non-coherent integration change,

as is the case when transitioning from acquisition to tracking.

A system model of this process is as follows. Consider an in-phase correlator channel *I(n)* and quadrature-phase correlator channel *Q(n)* each outputting independent zero mean Gaussian samples with variance σ^2 every 1 ms.

$$Z(n) = I(n) + j\,Q(n), \quad I(n) \text{ and } Q(n) \sim N\left(0,\sigma^2\right) \qquad ...(1)$$

$$N\left(\mu,\sigma^2\right): \quad p(x) = \frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{1}{2\sigma^2}(x-\mu)^2\right) \qquad ...(2)$$

After *N* ms of normalised pre-detection or coherent integration, the complex outputs are

$$\begin{aligned} Z_{Coh}(m) &= \frac{1}{N}\sum_{n=Nm}^{N(m+1)-1} Z(n) = \frac{1}{N}\left(\sum_{n=Nm}^{N(m+1)-1} I(n)\right) + j\frac{1}{N}\left(\sum_{n=Nm}^{N(m+1)-1} Q(n)\right) \\ &= I_{Coh}(m) + Q_{Coh}(m) \end{aligned} \qquad ...(3)$$

and where the distributions of $I_{Coh}(m)$ and $Q_{Coh}(m)$ are both given by:

$$I_{Coh}(m) \text{ and } Q_{Coh}(m) \sim N\left(0,\frac{\sigma^2}{N}\right). \qquad ...(4)$$

Subjecting $Z_{Coh}(m)$ to a normalised *M* rounds of non-coherent integration gives:

$$W(k) = \frac{1}{M}\sum_{m=Mk}^{M(k+1)-1}\sqrt{I_{Coh}(m)^2 + Q(m)^2} = \frac{1}{M}\sum_{m=Mk}^{M(k+1)-1}\left|Z_{Coh}(m)\right| \qquad ...(5)$$

where the distribution of $|ZCoh(m)|$ is Rayleigh (Kay, 1998) and can be denoted as $R(\sigma^2)$.

$$R\left(\sigma^2\right): \quad p(x) = \frac{x}{\sigma^2}\exp\left(-\frac{x^2}{2\sigma^2}\right) \quad for\ x > 0 \qquad ...(6)$$

The mean and variance of the Rayleigh distributed variable x are given by:

$$E(x) = \sqrt{\frac{\pi}{2}}\sigma \quad and \quad \mathrm{var}(x) = \left(\frac{4-\pi}{2}\right)\sigma^2 \qquad ...(7)$$

Using (4), the central limit theorem and (7), an estimate for mean and variance of the normalised non-coherently integration output can be given as:

$$E\left(w(k)\right) = \sqrt{\frac{\pi}{2}}\left(\frac{\sigma}{\sqrt{N}M}\right) \quad and \quad \mathrm{var}\left(w(k)\right) = \left(\frac{4-\pi}{2}\right)\frac{\sigma^2}{NM} \qquad ...(8)$$

The detection threshold can hence be given as

$$T = \frac{\sigma}{M}\sqrt{\frac{\pi}{2N}} + \kappa\,\sigma\sqrt{\left(\frac{4-\pi}{2MN}\right)} \qquad ...(9)$$

where κ is a constant that defines the FAR, with κ equal to 2 giving a 4.44% FAR and κ equal to 3 giving a 0.27% FAR.

Although this technique can be effective in some cases, it has the problem of not being very effective against cross correlation interference; also known as multiple access interference

(Glennon, 2009). The reason for this is that if a weak signal is to be detected in the presence of another strong signal, the noise variance calculated for a weak signal affected by MAI differs only slightly from the noise variance calculated for a signal in which the MAI has been eliminated. Intuitively, this is to be expected because measurement of variance of 1 ms dump samples is a wide-bandwidth measurement, whereas tracking weak signals is a narrow band measurement. In fact, assuming zero-mean Gaussian noise inputs on each channel, an estimate of the noise power for each channel can be easily made assuming the AGC produces digital samples with given statistics and the mixing signals are sine and cosine signals with other different mixing statistics.

Use of this technique was therefore eliminated from further consideration.

3.3 Measurement of Non-coherent Integration Statistics

An alternative to the previous method is to subject the noise tap/finger to exactly the same processing as the prompt tap/finger. Using this approach has the advantage that the effective bandwidth of the measurement is the same as the actual signal being tracked, so if MAI or CWI is visible in the prompt tap, effects should also be visible in the noise tap. The disadvantage of the method is that with larger integration periods, the rate at which statistics can be measured is significantly slower. Furthermore, it is also necessary to maintain separate statistics for acquisition and capture modes, which typically employ different settings for the coherent and non-coherent integrations.

Several options were considered in the implementation. One was to accumulate statistics on the same non-coherent integrations used in the detection process. By accumulating sum and sum of squares quantities, it is possible to estimate the mean and standard deviation of the values. The firmware can then assume a roughly normal distribution for the processes by application of the central limit theorem and hence estimate a threshold that gives a sufficiently high false alarm rate (FAR). The difficulty is that in reality, the output distributions are not normal, even though they are approximately normal, and therefore the actual FAR ends up being higher than expected.

A solution to this dilemma is straightforward once it is recognised that the values making up the noise lie in narrow range, so rather than calculating sums and sums of squares, the firmware builds up a histogram of the frequency of each of the noise values. Such a histogram directly measures the distribution of the noise and can be used to calculate expected values and variances if needs be. It can also permit direct estimation of a detection threshold satisfying a given FAR. Incorporating new values into the histogram is extremely efficient because the histogram is indexed by the noise values themselves, with each element being incremented as new values are added to the record. As values are added, a check can also be performed to see whether a new maximum has been observed and the maximum value of the histogram updated. This allows thresholds that produce very low FARs to be generated. To allow for changes in the statistics, the maximum count allowed in any of the noise bins is constrained. Whenever this limit is breached, the entire histogram is scaled back by a convenient factor, with a reduction factor of a power of two is particularly convenient as each of the histogram counts can be right shifted by one or more bits.

One danger with maintaining a histogram of this type is ensuring that non-noise values are not added to the histogram. This requires that sanity checks be carried out before any new values are added and values exceeding that limit should be discarded. Another difficulty is the large dwell time per noise sample causing a long delay before a valid histogram is available. This

		SV 1	SV2	SV3	SV4	SV5	SV6	Remark
Scenario 1	C/N0 (dBHz)	50	30	30	30	30	30	1 strong SVs
4 ms, 8 NCR	Doppler (kHz)	1.1	-2.8	0.3	-3.6	-1.5	3.6	
Scenario 2	C/N0 (dBHz)	50	50	50	50	50	50	6 strong SVs
4 ms, 8 NCR	Doppler (kHz)	1.1	-2.8	0.3	-3.6	-1.5	3.6	
Scenario 3	C/N0 (dBHz)	50	50	50	50	50	50	6 strong. SV 2 causes
4 ms, 8 NCR	Doppler (kHz)	1.1	-2.9	0.3	-3.6	-1.5	3.6	MAI on 1
Scenario 4	C/N0 (dBHz)	50	50	50	50	50	50	6 strong. SV 2 & 3
4 ms, 8 NCR	Doppler (kHz)	1.1	-2.9	0.1	-3.6	-1.5	3.6	cause MAI on 1
Scenario 5 (1)	C/N0 (dBHz)	50	30	30	30	30	30	1 strong SVs
20 ms, 4 NCR	Doppler (kHz)	1.1	-2.8	0.3	-3.6	-1.5	3.6	
Scenario 6 (2)	C/N0 (dBHz)	50	50	50	50	50	50	6 strong SVs
20 ms, 4 NCR	Doppler (kHz)	1.1	-2.8	0.3	-3.6	-1.5	3.6	
Scenario 7 (3)	C/N0 (dBHz)	50	50	50	50	50	50	6 strong. SV 2 causes
20 ms, 4 NCR	Doppler (kHz)	1.1	-2.9	0.3	-3.6	-1.5	3.6	MAI on 1
Scenario 8 (4)	C/N0 (dBHz)	50	50	50	50	50	50	6 strong. SV 2 & 3
20 ms, 4 NCR	Doppler (kHz)	1.1	-2.9	0.1	-3.6	-1.5	3.6	cause MAI on 1

Table 2: Threshold calculation scenarios

means that when a channel is started up, it is necessary to create an approximate histogram in order to have valid thresholds that can be output. One way to do this is to copy the histogram of another channel if that channel has the same settings for the coherent and non-coherent integration periods. However, if this is not true, an estimate can always be generated using approximations.

4 MATLAB SIMULATION EXPERIMENT RESULTS

4.1 Experiment Description

To test and demonstrate the concepts described in the previous sections, a Matlab simulation was created to generate IF signals containing GPS C/A code data for a set of satellites with user specified Doppler, SV number and code phases. The signals are produced in selectable batch sizes, with a typical batch size being approximately 1 ms. The generator ensures signal continuity across batches, was well as producing replicas of the signals needed to de-spread code for five code phase fingers and to mix the signals to baseband. The five half-chip separated code phase fingers correspond to very-very early, very early, early, prompt and late, with the very-very early being used as a noise finger.

This Matlab experiment tested a small number of selected scenarios, with the details on the scenarios given in Table 2. The scenarios include cases with a single strong signal, multiple strong signals each at Doppler frequencies that should avoid any cross correlation interference, as well as two cases with one and two satellites that should cause cross correlation interference. Scenarios 1 to 4 employ a coherent integration of 4 ms with 8 NCR, while scenarios 5 to 8 employ a coherent integration of 20 ms with 4 non-coherent rounds. This corresponds to the settings for the least sensitive and most sensitive tracking modes that are typically used. 10 seconds of simulated data was used in each case, although in reality, the histograms used in the Aquarius firmware would employ data accumulated over longer periods thereby improving the statistical significance of the outputs.

This experiment mirrors the processing performed by the Kea GPS receiver, employing a

Scenario	s_{1ms}	$\sigma_{PDI}(s_{1ms})$	$\mu_{NC}(s_{1ms})$	$\sigma_{NC}(s_{1ms})$	σ_{1ms}	$\sigma_{PDI}(\sigma_{1ms})$	$\mu_{NC}(\sigma_{1ms})$	$\sigma_{NC}(\sigma_{1ms})$	μ_{NC}	σ_{NC}
1	373	11.7	14.6	2.7	365	11.4	14.3	2.6	14.4	2.7
2	373	11.7	14.6	2.7	495	15.5	19.4	3.6	16.7	3.5
3	373	11.7	14.6	2.7	499	15.6	19.6	3.6	17.8	4.5
4	373	11.7	14.6	2.7	510	15.9	20	3.7	23.2	6.1
5	373	5.2	6.5	1.7	365	5.1	6.4	1.7	6.5	1.8
6	373	5.2	6.5	1.7	495	6.9	8.7	2.3	6.7	1.9
7	373	5.2	6.5	1.7	499	7.0	8.7	2.3	10.0	4.4
8	373	5.2	6.5	1.7	510	7.1	8.9	2.3	18.6	6.2

Table 3: Fully predicted coherent & non-coherent statistics using estimated 1 ms noise standard deviation, predicted coherent and non-coherent statistics using measured 1 ms noise standard deviation and measured non-coherent noise statistics.

Scenario	$T(2\ s_{1ms})$	$T(3\ s_{1ms})$	$T(2\ \sigma_{1ms})$	$T(3\ \sigma_{1ms})$	$T(2\ \sigma_{NC})$	$T(3\ \sigma_{NC})$	$T_{5\%}$	$T_{1\%}$
1	20	22.7	19.6	22.2	19.8	22.5	18.5	21
2	20	22.7	26.5	30.1	23.6	27.1	22	24.5
3	20	22.7	26.8	30.4	26.8	31.2	24.5	30.5
4	20	22.7	27.4	31.1	35.4	41.6	34.3	38
5	9.9	11.7	9.7	11.4	10.1	11.9	9	11
6	9.9	11.7	13.2	15.5	10.6	12.6	9	11
7	9.9	11.7	13.3	15.6	18.8	23.2	<20	21
8	9.9	11.7	13.6	15.9	31.1	37.3	31	32

Table 4: 2 and 3 sigma thresholds using estimated 1 ms noise standard deviations, 2 and 3 sigma thresholds using measured 1 ms noise standard deviations, and histogram derived thresholds at the 5% and 1% FAR thresholds. Here $T(n\ \sigma)$ is taken to be $\mu + n\ \sigma$, where μ is the mean and σ the standard deviation. $T_{5\%}$ and $T_{1\%}$ are estimates based on a histogram with unit bin sizes and hence have a higher uncertainty.

sampling frequency of 16.368 MHz, an IF frequency of 4.092 MHz and 2-bit sign/magnitude sampling (±1, ±3) from the RF front end. However, unlike the Kea GPS receiver, the digital mixing employs 2-bit signals for the locally generated carrier signals. This was done for simplicity in calculating the predicted noise variance when no signals were present and to simplify the Matlab code. The model also includes a scaling back of the coherent-integrations by a factor of 2^{-4}, which Aquarius does to prevent overflow during other parts of the processing.

4.2 Outcome and Discussion

The results from these experiments are given in Table 3 and 4. The left portion of the table 3 uses a calculated estimate for the noise standard deviation s_{1ms} that is calculated based on the sampling frequency, the distribution of the input sign and magnitude IF samples, as well as the statistics of the locally generated carrier signals used to mix this to baseband. It is then possible to calculate as a function s_{1ms} the standard deviations for the coherent and non-coherent integrations, as well as the mean of the non-coherent integrations. This process is then repeated using a measured value for the noise standard deviation, which is denoted by σ_{1ms}. The last two columns give the measured values for the mean and standard deviations of the non-coherent integrations.

Table 4 continues the results from these experiments, showing the calculated thresholds at the 2 and 3 sigma levels for the values derived from s_{1ms} and σ_{1ms}, as well as the values derived from the histogram at the 5% and 1% FAR.

In general, the results of these experiments confirm much of what has been said regarding the need to avoid using purely calculated results. Although these are acceptable when only one strong signal is present, once other satellites become visible the degradation they cause becomes apparent. Comparing the thresholds calculated from the 1 ms measured statistics σ_{1ms} and statistics measured from the non-coherent outputs also supports the claim that direct measurement of noise statistics from non-coherent results is essential once strong signals with MAI begin to appear. In these experiments, the thresholds measured from the distribution are similar to those obtained by measuring the mean and variance of the non-coherent accumulations, although the former has the advantage of allowing a user specified FAR to be chosen if necessary. Maintaining the histograms to do this do take more memory than maintaining sums and sums of squares, but the flexibility offered more than makes up for this cost.

A measure of the predicted improvements in sensitivity can also be made by comparing the thresholds employed in each of the different scenarios. For example, taking the ratio between the amplitude detection threshold employed in the single strong satellite scenarios 1 and scenario 4, calculated here as $T_1(3\boldsymbol{\sigma}_{\mathbf{NC}})/T_4(3\boldsymbol{\sigma}_{\mathbf{NC}})$, gives a value of 1.97 or 6 dB. This improvement is similar to the improvement implied by Table 1 if the 4 ms and 8 NCR results are compared with the 20 ms and 4 NCR results for an input SNR of -145 dBm, which gives (22.5 - 15.8) dB or 6.7 dB. It is also clear that as the level of MAI increases, the sensitivity predicted by the Matlab simulation reduces, as expected and required by design. Hence $T_4(3\boldsymbol{\sigma}_{\mathbf{NC}})/T_8(3\boldsymbol{\sigma}_{\mathbf{NC}})$ gives 1.11, which is only 0.94 dB.

5 KEA HARDWARE AND SIMULATOR TESTS

The main purpose of this exercise was to improve the sensitivity of the Kea receivers. To prove the efficacy of the measures put in place, a test using the UNSW Spirent STR8000 GNSS simulator was performed. Kea V41SBR3 with serial number 30 was programmed with the legacy QB50/EC0 firmware build and Kea V41SB43 with serial number 32 was programmed with the new high sensitivity firmware. The simulator power was then gradually reduced in 1 dB steps and the outputs of the receiver logged. When the simulator reached a power level of +2, corresponding to a signal level of approximately -135 dBm, the legacy receiver stopped navigating, but continued to track signals for an additional 3 dB of power reduction. The new firmware continued to track and navigate until a power level of -4 was reached. Note that a Spirent power level of 0 corresponds to -130 dBm of output power, but with this test 3.3 dB must be subtracted for the power splitter, and approximately 4 dB subtracted to account for the increased noise temperature of simulator compared to the open sky).

Taking these results at face value gives an improvement of 6 dB, more or less in line with the expected improvement. However, the fact that the legacy code continued to track signals for an additional 3 dB meant that the real improvement was only 3 dB. To fully understand this, the legacy source code was examined and quality thresholds that had been removed from the new firmware, but remained in legacy firmware were identified. The legacy code was rebuilt with those thresholds significantly reduced and a similar test performed.

The revised test eliminated the splitter, choosing instead to connect each receiver to the simulator using exactly the same interface board and cable, with only the receiver swapped over. This eliminated the need to account for the splitter and the additional cable losses, as well as the differences between the interface boards. In the revised test, the new firmware showed marginal navigation at -6 dB and complete loss of all signals at -7 dB, while the revised legacy firmware was perfectly good at -3 dB, but marginal with eventual complete loss of lock at -4 dB. This confirmed the real improvement of 3 dB.

Using inspired guesswork, the high sensitivity mode of the new receiver was then modified from using a coherent integration period of 20 ms and 4 non-coherent rounds of integration to 10 ms and 8 non-coherent rounds, after which the test on the new firmware was repeated. The revised firmware performed significantly better, continuing to navigate down to Spirent power levels of -8 dB, with loss of signals and navigation at -9 dB. This gave an improvement of 5 dB compared to the legacy build and was only 1.7 dB short of the expected 6.7 dB improvement. Comparison of the (10 ms, 8 NCR) and (20 ms, 4 NCR) suggests 1.6 dB squaring losses differences at these levels, so if the PLL and FLL tracking loops can be optimised for operation at 20 ms, an additional 1.6 dB improvement could be expected.

6 CONCLUSIONS

The techniques used to improve the sensitivity of the Kea GPS receiver have been described. These include use of 20 ms bit aligned coherent integration, along with additional non-coherent integration. It was necessary to adaptively change the coherent integration period depending on whether bit-synchronisation had been achieved. It was also necessary actively measure the noise floor appropriate to the selected coherent and non-coherent integration settings. This allows the receiver to operate in standard conditions without getting excessive false alarms or tracking noise, but also to track weak signals when the noise levels have dropped to a sufficiently low level. Some additional work remains in order to get the tracking loops to operate with 20 ms coherent integration, but the results obtained using 10 ms coherent integration are in line with expectations.

REFERENCES

Glennon EP. (2009), Cross Correlation Mitigation for C/A Code GPS Receivers, PhD Thesis, University of New South Wales.

Glennon EP, Parkinson K, Mumford P, Shivaramaiah N, Li Y, Li R, Jiao Y. (2011), A GPS Receiver Designed for Cubesat Operations: Australian Space Science Conference, pp. 213-222.

Kaplan ED, Hegarty CJ. (2006), Understanding GPS: Principles and Applications: 2nd Edition. Artech House.

Kay SM. (1998), Fundamentals of Statistical Signal Processing: Detection Theory, Prentice Hall.

Mumford PJ, Shivaramaiah NC, Glennon EP, Parkinson K. (2012), The Namuru V3.2A Space GNSS receiver: Australian Space Science Conference 2012, Melbourne.

Psiaki ML. (2001), Block Acquisition of Weak GPS Signals in a Software Receiver: ION GPS 2001, Salt Lake City, UT.

van Dierendonck AJ. (1996), GPS Receivers, Parkinson, BW, Spilker, JJ (eds), Global Positiioning System: Theory and Applications. AIAA, Cambridge, MA.

van Diggelen F. (2009), A-GPS: Assisted GPS, GNSS, and SBAS, Artech House.