Examination of a simple pulse blanking technique for RFI mitigation

Noppasin Niamsuwan

Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH, USA

Joel T. Johnson

Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH, USA

Steven W. Ellingson

Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA, USA

Spectroscopy at L-band can be adversely impacted by radio frequency interference (RFI) due to the presence of numerous sources, especially pulsed RFI from radars operating below 1400 MHz. RFI mitigation is very important to deal with this problem. A simple strategy for reducing pulsed RFI termed "asynchronous pulse blanking" has been implemented in a digital receiver developed at The Ohio State University. This paper presents results from a simulation of the APB algorithm. Several aspects of algorithm use and performance are reported, including means for choosing the algorithm's parameters and the robustness of the method in a realistic RFI environment. Effects of the blanking process on the final output are also examined.

1. Introduction

Radio frequency interference (RFI) mitigation is very important for spectroscopy at L-band due to the presence of numerous interference sources, including strong pulsed RFI from ground-based aviation radars (GBARs) [Ellingson 2003, Ellingson et al 2003a]. For radiometers operating at a sufficiently high temporal sample rate, a simple strategy for reducing pulsed RFI is to remove incoming data whose power exceeds the mean power by a specified number of standard deviations. It may also be advantageous to remove data within a specified time region before and after this "trigger" data, to ensure that any pre- and post-detection "pulse" information is successfully removed. Such an algorithm has been implemented in a digital receiver-based radiometer developed at the Ohio State University ElectroScience Laboratory [Ellingson et al 2003b, ElectroScience Lab]; the process is termed "asynchronous pulse blanking" (APB) because no periodic properties of the interference source are assumed. Although successful performance of this algorithm has been qualitatively demonstrated through local experiments with the digital radiometer [Hampson et al 2004], a detailed quantitative assessment of its performance in a variety of RFI environments has not been reported.

Copyright 2004 by the American Geophysical Union.
0048-6604/04/$11.00

To address this issue, a simulation study of the APB algorithm was initiated using data obtained from the L-Band Interference Surveyor/Analyzer (LISA), an airborne instrument developed for observing the RFI in the region 1200-1700 MHz [Ellingson et al 2003a, Ellingson and Johnson 2003, Johnson and Ellingson 2003]. The LISA instrument included a digital receiver capable of capturing 14 MHz of incoming data; this 14 MHz channel was tuned through the 1200-1700 MHz band throughout system flights. During Jan–Feb 2003, LISA was deployed on NASA's P3-B aircraft to observe RFI in flights in the US and Japan. This data set is very useful for assessing the APB algorithm, since many RFI environments were observed that include multiple sources of interference.

The next section provides a basic description of the digital receiver based radiometer, and a detailed summary of the APB algorithm follows in Section 3. Information on the LISA instrument and dataset is then provided in Section 4, and simulation results presented in Section 5.

2. Interference Suppressing Microwave Radiometer

Under the support of the NASA "Instrument Incubator Program" (IIP), development of a digital-receiver based L-band radiometer was initiated in Dec 2001. Although the project is focused on passive microwave observation of the Earth for L-band remote sensing, the radiometer can be applied for L-band radio astronomy as well, and both applications are severely impacted by RFI.

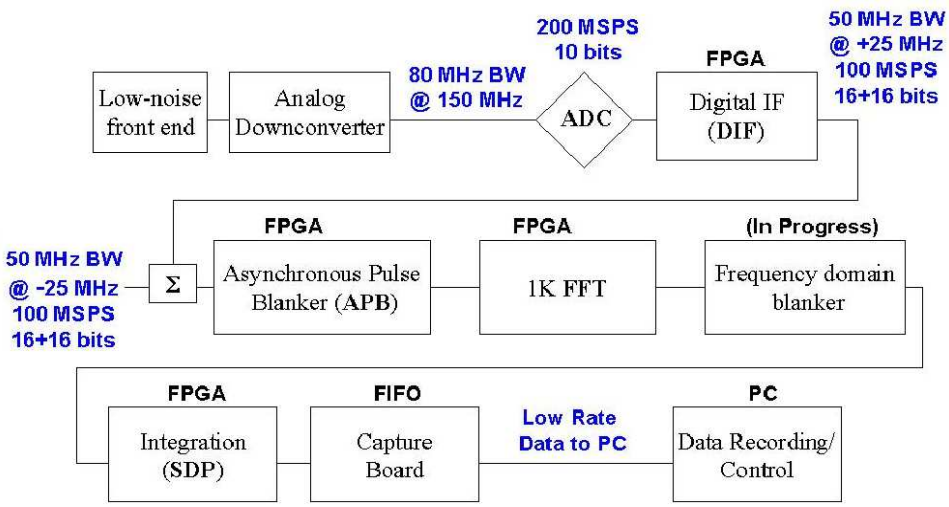

Figure 1. Block diagram of radiometer

A block diagram of the radiometer is shown in Figure 1. The analog front end downconverts an 80 MHz swath of spectrum from L-band to 150 MHz, and samples this signal at 200 MSPS using 10 bits. Note the system can also be operated at frequencies other than L-band simply by modifying the analog front end and downconverter sections. Because the analog IF is in the second Nyquist zone of the A/D, the digital passband is centered at 50 MHz and is spectrally reversed. The "Digital IF" (DIF) FPGA module downconverts this to 0 Hz (so now the samples are complex-valued), filters to 50 MHz bandwidth, decimates by 2, and then upconverts to a center frequency of +25 MHz (still complex). The data emerges from the DIF module in 16-bit "I"+16-bit "Q" format at 100 MSPS. The sample process is applied to a separate, independently-tunable 50-MHz swath at L-Band, with the difference that the digital output is centered at -25 MHz. The two 50 MHz bands are simply added together to form single 100 MHz bandwidth signal.

Following the DIF output is a cascade of FPGA modules which can be programmed to perform a variety of functions. Our favored strategy currently is as shown in Figure 1: mitigation of radar pulses using APB, channelization into 100-kHz bins using a 1K FFT, frequency domain blanking, and integration to generate power spectra.

The APB is designed to detect and blank radar pulses, which typically are the dominant source of external L-Band RFI below 1400 MHz. Radar pulses range from 2-400 μs in length and occur 1-75 ms apart [Ellingson 2003]. To detect these pulses, the APB maintains a running estimate of the mean and variance of the sample magnitudes. Whenever a sample magnitude greater than a threshold number of standard deviations from the mean is detected, the APB blanks (sets to zero) a block of samples beginning from a predetermined period before the triggering sample, through and hopefully including any multi-path components associated with the detected pulses. APB operating parameters are adjustable and can be set by the user.

Following the APB is a length-1K complex FFT, which achieves approximately 98% duty cycle in performing the FFT computations. A triangular window is applied before the FFT. Planned but not yet implemented is a frequency domain blanking module, which is similar in concept to the APB, except applied independently to each frequency bin. The purpose of this module will be to exploit the processing gain achieved through channelization to detect and excise weak, relatively narrowband RFI. The FFT output is processed through a "spectral domain processor" (SDP) module which computes magnitude-squared for each frequency bin and computes a linear power average over many FFT outputs. These results are passed at a relatively low rate to a PC via a capture board. Total power can be computed by summation of frequency bins within the digital hardware, or the same process can be implemented within the PC for increased flexibility in monitoring RFI, selecting subbands, and so on.

This radiometer has been used in both radio astronomy and remote sensing experiments locally at the ElectroScience Laboratory [Hampson et al 2004] and at the Arecibo observatory in Puerto Rico [Ellingson and Hampson, 2003a, 2003b]. The results qualitatively show the APB approach to be successful for pulsed RFI mitigation, although a detailed performance assessment has not yet been achieved.

3. APB Algorithm

3.1. Basic Concept

The APB algorithm is intentionally simple to allow implementation in hardware (FPGA's). For each sample x in the data stream, $||x||^2$ (the squared-magnitude of the sample) is computed and compared to a threshold level δ. *Detection* is declared if $||x||^2$ exceeds this threshold. The threshold level is defined to be $\beta\sigma + m$, where σ and m are the standard deviation and mean of $||x||^2$, respectively. The parameter β sets the "aggressiveness" of the blanker; for small β, the blanker may trigger on larger noise peaks, whereas a high value may allow weak pulses to pass through. Consequently, the β parameter should be optimized to trade-off loss of noise power integration time versus blanking performance.

Because "pulses" detected by the algorithm typically begin before a detection is obtained, it is useful for the APB algorithm to have the capability of blanking information prior to a detection. To blank pre-detection samples, a buffer (memory) is required in the system. The detection operation is then performed on samples entering the buffer, while the blanking operation and all successive operations are performed on samples exiting the buffer. A simple first-in-first-out (FIFO) buffer suffices, and the length of this FIFO (`FIFO_LENGTH`) determines the maximum period that can be blanked before the detected sample.

Even though the APB processor is capable of blanking all the samples in the buffer when a pulse is detected, this is not always necessary. Blanking the whole buffer may remove desired noise information and should be avoided if possible. To control the length of the pre-pulse blanking region, the parameter N_{wait} is introduced. When *detection* is asserted, a process is initiated to wait for a N_{wait} samples before the FIFO output is set to zero. The number of pre-detection samples remaining in the buffer after N_{wait} clock cycles have elapsed is then `FIFO_LENGTH` $- N_{wait}$. The FIFO output is then continually forced to *zero* for a specified number of clock cycles. The length of this blanking "window" is determined by the final APB parameter, N_{blank}. The set of parameters β, N_{wait}, N_{blank}, and `FIFO_LENGTH` are the basic conceptual quantities that control the APB process.

Figure 2 illustrates simulated incoming signal power, and includes simulated RFI pulses. The blanking regions in the plot clarify how the N_{blank} and N_{wait} parameters are defined. It is clear from this Figure that N_{blank} should be chosen large enough to remove a complete pulse and any possible multi-path contributions, while the combination `FIFO_LENGTH` $- N_{wait}$ should be chosen sufficiently large to remove pre-detection contributions. The studies of Section 5 investigate methods for choosing these parameters.

3.2. Hardware Issues

3.2.1. `FIFO_LENGTH`

Ideally, the FIFO should be made as large as possible in order to allow the largest pre-detection period. However, hardware issues (i.e. size of the FPGA used) limit the maximum FIFO size which can be included. For the IIP radiometer and for the remainder of the simulations performed in this paper, `FIFO_LENGTH` was set to 1024 due to size limitations of the FPGA components used in the prototype system.

3.2.2. Running Filter

The APB algorithm requires a running estimate of the mean (m) and variance (σ^2) of the incoming signal power. These estimates are computed as:

$$m = \mu_{mean} \times m_{t-1} + (1 - \mu_{mean}) \times ||x||^2 \quad (1)$$

$$\sigma^2 = \mu_{var} \times \sigma^2_{t-1} + (1 - \mu_{var}) \times (||x||^2 - m)^2 \quad (2)$$

where m_{t-1} and σ^2_{t-1} are the mean and variance generated on the previous clock cycle, and $\mu_{mean/var} \approx 1$ are constants that control the response rate of these simple averaging filters. In the APB hardware, the above computations are accomplished by groups of fixed-point multipliers, adders, and delay registers. In addition, a state machine controls these filters so that contributions from incoming post-detection data that will be blanked are not included in computing m and σ^2.

3.2.3. Large Bit-width

Due to the squaring operation in computing $||x||^2$ (represented originally as 16 bit I and 16 bit Q), a wide datapath of 32 bits is required to preserve accuracy for this quantity. The estimate of σ^2 further computes the square of $||x||^2$, requiring 64 bits. The high dynamic range of the filter constants $\mu \approx 1$ and $(1-\mu) \approx 0$ further increases the number of bits required in estimating the running mean and variance of $||x||^2$. Use of 12 bits to capture both μ and $(1 - \mu)$ was found desirable in setting the filter time constants. The resulting arithmetic operations grow to a maximum size of 76 bits, and severely limit the computation speed of the detector due to this large size.

3.2.4. Processor Speed

Because of the large bit growth of the APB processor and speed limitations of the FPGA components currently used, the detection operation can not be performed on 100% of the incoming data. The current implementation of the design can process only one in four samples (25%) of the data for detection. A decimation controller is introduced to sub-sample the input data stream to every 4 samples for detection operations. However, since radar pulses are typically longer than this sub-sampling period, they are still detectable even under this limitation. This

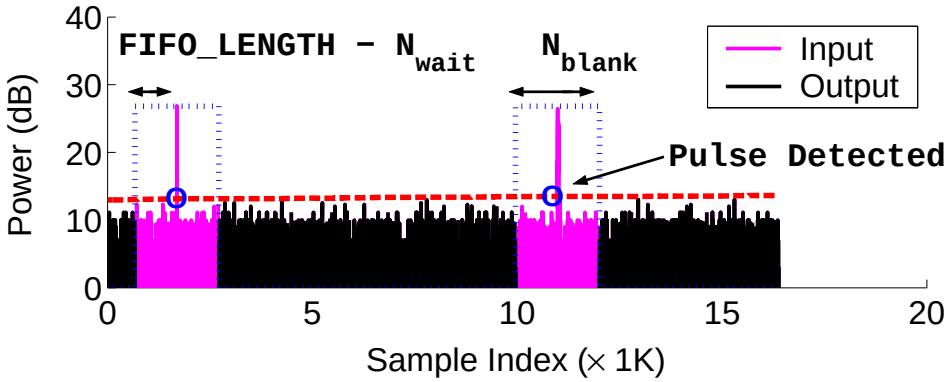

Figure 2. Definition of N_{blank} and N_{wait}: note time is increasing from left to right

sub-sampling operation refers only to samples analyzed for detection, the remainder of the system, including the blanking operation, still operates on 100% of the incoming data.

3.2.5. Control of Blanking Process

When a detection is obtained, a simple state machine is initiated to wait for N_{wait} clock cycles, then blank the output of the FIFO for N_{blank} clock cycles. The state machine controlling this process is called a "blanking timing register" (BTR). Although the BTR is simple to implement in hardware, it still occupies a non-negligible portion of the available FPGA logic. Therefore it is not possible to have an unlimited number of BTR's in the system. Because a BTR is unavailable to begin a blanking process for $N_{wait} + N_{blank}$ cycles after it has been "triggered", the possibility of having no BTR's available to blank a detected pulse must be considered if only a finite number of BTR's exist.

A limitation in the number of BTR's also raises another issue: a large pulsed interferer presumably could trigger all available BTR's on successive clock cycles, with the resulting blanking windows overlapped almost completely at the FIFO output. The ability of the APB to blank detected pulses would then be compromised for approximately $N_{wait} + N_{blank}$ clock cycles when the initial BTR would become available again. This problem motivates the introduction of a parameter N_{sep}, defined to be a number of clock cycles after a detection during which no further BTR's will be triggered. If N BTR's are available, then setting

$$N_{sep} = \frac{1}{N}\left(N_{wait} + N_{blank}\right) \tag{3}$$

will ensure that a BTR is always available for triggering once N_{sep} has elapsed. Although N_{sep} values larger than the above value would also satisfy this requirement, minimizing N_{sep} is desirable to ensure maximal pulse detection.

While introduction of N_{sep} manages the use of a finite number of BTR's, it also may compromise the ability of the APB to remove detected pulses, if a detection occurs within N_{sep} clock cycles of an earlier detection. Although a new BTR process will not be initiated for N_{sep} clock cycles, the APB algorithm can still be designed so that all post-detection signal information within N_{sep} clock cycles will be blanked. Referring to Figure 2, the post-detection blanking region contains $N_{blank} + N_{wait} - \texttt{FIFO_LENGTH}$ samples for a fixed N_{blank}. Choosing

$$N_{sep} \leq N_{blank} + N_{wait} - \texttt{FIFO_LENGTH} \tag{4}$$

will ensure that all signal information within N_{sep} samples after a detection will be blanked.

Combining equations (3) and (5) results in the condition

$$N_{blank} \geq \frac{N}{N-1}\texttt{FIFO_LENGTH} - N_{wait} \tag{5}$$

to ensure that all samples exceeding the specified treshold will be blanked. Although this equation appears to require that N_{blank} increase with FIFO_LENGTH, the associated N_{wait} would also likely increase with FIFO_LENGTH so that N_{blank} would not necessarily increase. The digital radiometer prototype at present contains 4 BTRs; this number will be assumed in the simulations of Section 5. For this case with $\texttt{FIFO_LENGTH} = 1024$, N_{blank} should be chosen $\geq 1366 - N_{wait}$.

4. L-band Interference Surveyor/Analyzer (LISA)

In order to perform a more complete assessment of the APB algorithm, it is desirable to study performance in a complex and varying RFI environment. However, experiments with the radiometer prototype to date have occurred only in the local environment and at the Arecibo

facility. Furthermore, it is desirable to perform these assessments in a software rather than hardware environment, so that quantitative information can be obtained as APB parameters are varied for a fixed specific dataset. The dataset from the LISA sensor provides an excellent resource for these studies.

LISA, completed in September 2002, was developed as a means to observe sources of radio frequency interference (RFI) in the region of 1200-1700 MHz from an airborne platform [Ellingson et al 2003a, Ellingson and Johnson 2003, Johnson and Ellingson 2003]. Physically, LISA is comprised of a small antenna/ front-end unit (AFEU) and an equipment rack. The antenna unit consists of a nadir-facing cavity-backed planar spiral antenna with an integrated custom RF front end including filtering and calibration circuitry. The antenna has a very broad pattern (approximately "cos θ") and is reasonably well-matched over the span of the observations reported here. The antenna unit is connected to an equipment rack mounted in the aircraft cabin by a long and fairly lossy coaxial cable. Although the cable loss degraded the sensitivity of the instrument, the resulting gain profile was an important factor in preserving the linearity of the system while observing strong RFI waveforms, this consideration was paramount, but comes at the expense of the system's ability to detect weak RFI.

Inside the equipment rack, the signal is delivered to a custom-designed coherent sampling subsystem. This subsystem uses a direct-conversion receiver to tune (under PC control) anywhere between 1200 MHz and 1700 MHz. "I" and "Q" signals at baseband are low-pass filtered with ~ 7 MHz cutoff and sampled at 20 MSPS, yielding a digitized bandwidth of ~ 14 MHz. The output samples are queued in a 16K-sample-long first-in-first-out (FIFO), whose contents are acquired by means of the PC parallel port. These 16K "captures" represent an 819.2 μsec sample of the received field at the current center frequency, and require approximately 1 sec to be transferred and recorded on the computer. During flight, the coherent sampling system was successively tuned through center frequencies of 1250, 1264, ..., 1698 MHz; these distinct bands are called channel 1, 2, etc. in what follows. For each channel, 5 819.2 μsec subsequent captures are performed before tuning to the next center frequency. After completing a sweep of all channels, a backup operation was performed to a separate storage component in the computer system. The time required for this operation caused a 10 to 15 minute delay between "sweeps". The final dataset provides a high temporal resolution but very low duty cycle observation of 4.1 msec captured per LISA channel per 12 to 15 minutes.

For the datasets considered here, LISA was installed in NASA's P-3 research aircraft, which is based at the Wallops Flight Facility (WFF) located at Wallops Island, VA. The LISA antenna unit was mounted in the tail radome. The loss due to transmission through the radome is unknown and may be another factor degrading sensitivity. The data used in the simulations to follow were collected during a single flight on January 3, 2003 along an east-west track from Wallops Island, VA to Monterey, CA at approximately 22,000 ft. These data obviously include observation in a variety of RFI environments. Data from LISA channels 6-13 ($1310 - 1428$ MHz) is used in the simulations; for each of these channels, 145 captures (5 captures $\times$ 29 sweeps) were available. LISA was also deployed in the "Wakasa Bay" remote sensing campaign [Colorado State Univ.] in Japan through Feb of 2003; data on the transit flights from Monterey to Japan were also collected [Johnson and Ellingson, 2003]. However, the RFI environments in these flights show less variation than that of the January 3rd flight, and so these data are not considered further in this paper.

5. Simulations and Results

5.1. Choosing β^2, N_{blank}, and N_{wait}

To investigate choice of β^2, data from LISA channel 7 was examined to determine the percent of samples exceeding a specified threshold determined by β. Because the mean measured noise power can be expected to vary significantly throughout a cross-US flight, the mean and variance estimates used to determine the threshold are computed by the running average described in Section 3. This running average used $N_{blank} = 2048$ and $N_{wait} = 0$ in removing pulse contributions from the mean and variance computations, as described in Section 3.2.2. Figure 3 illustrates the percent of samples out of the 145 16K captures that exceed the threshhold specified by β^2 on the horizontal axis. The dashed line included in the Figure represents the percent of exponentially distributed noise that would exceed the same threshold. An exponential, rather than Gaussian, noise distribution is more appropriate here due to the very short integration time of the incoming power for this system. In this case, the presence of RFI in the dataset causes the percent blanked to exceed that of the simulated noise for large β^2 values; however as β^2 is reduced, the two curves become more identical because the blanking of noise dominates both processes. In this case, the turning point with the LISA data appears to be in the range $\beta^2 \approx 30$ to 40. Further simulations with other LISA channels showed β^2 values ranging from 40 to 90 to be reasonable; smaller β^2 values clearly led to excessive blanking as should be expected.

As mentioned previously, the parameter `FIFO_LENGTH` is fixed to 1024 (i.e. 51.2 μsec of LISA data) in order to model the existing radiometer prototype. Due to this somewhat small size, tests using LISA channels 7-13 indi

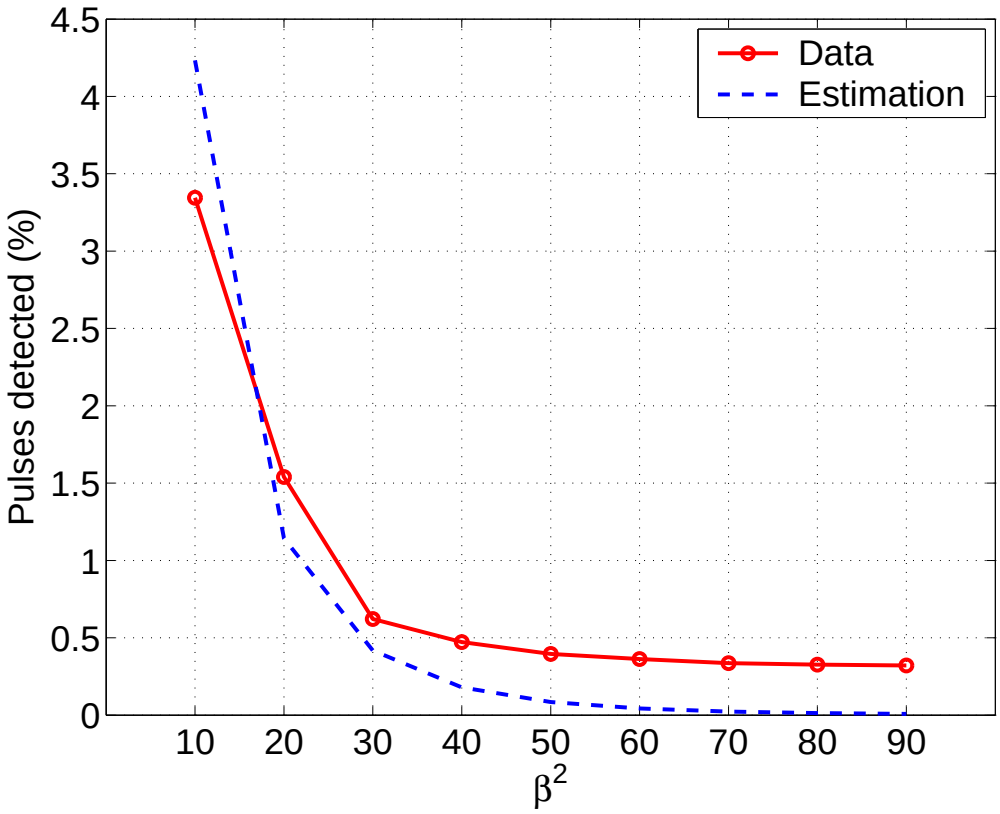

Figure 3. Percentage of samples exceeding threshold for given β^2. All 145 captures of LISA's channel 7 data are used.

cate that immediate blanking of the FIFO is preferable. For this reason, N_{wait} is set to 0 in the remaining simulations. Future work will explore a preferred value of N_{wait} for larger `FIFO_LENGTH` parameters.

A simulation investigating the effect of N_{blank} has also been performed. In this study, two threshold levels are defined: a higher "detection threshold"(β^2=90) used for pulse detection as usual while N_{blank} is varied. The output of the blanker is then examined to determine the number of samples remaining that exceed a lower "reference threshold"(β^2=30 using the same mean and variance computations as those for the detection threshold). This quantity is labeled P_{out} in what follows. The total number of samples exceeding the reference threshold in the original data is also determined, and labeled P_{in}. The ratio $(P_{in} - P_{out})/P_{in}$ then provides information on the effectiveness of the higher threshold blanker with a given N_{blank} at removing lower level pulse contributions. Figure 4 illustrates these quantities for a single LISA capture. Note in some cases, samples exceeding the higher threshold are not blanked in this dataset; this is because such "trigger" samples may be missed by the subsampled detector if they are less than 4 samples long.

Figure 5 plots $(P_{in} - P_{out})/P_{in}$ as a percentage (solid curve) from the entire channel 7 dataset, with $N_{blank} \geq 1536$ to satisfy equation (5). The curve shows only a modest variation with N_{blank}, and an N_{blank} value in the range 1536 to 2048 appears appropriate in this case. This can also be interpreted as indicating that pulsed interference longer than 76.8μsec (i.e. 1536 samples) is not significant in this data. Of course, this N_{blank} parameter will vary for RFI environments dominated by different types of sources, but the cross-US flight considered here should be fairly representative of other datasets. Results for smaller N_{blank} values may show more sensitivity, but are not explored here due to the limitation of equation (5). Data for smaller N_{blank} values cases could be generated by varying the `FIFO_LENGTH` and N_{wait} parameters.

The saturation of the solid curve at a maximum value of approximately 80% is caused by the presence of spurious internally generated interference in the LISA dataset [Johnson and Ellingson, 2003]. The dashed curve is computed by redefining P_{in} and P_{out} to included only consecutive sets of 2 or more samples that exceed the reference threshold. In this case, a removal of approximately 98% of the "low-level" pulse components is achieved.

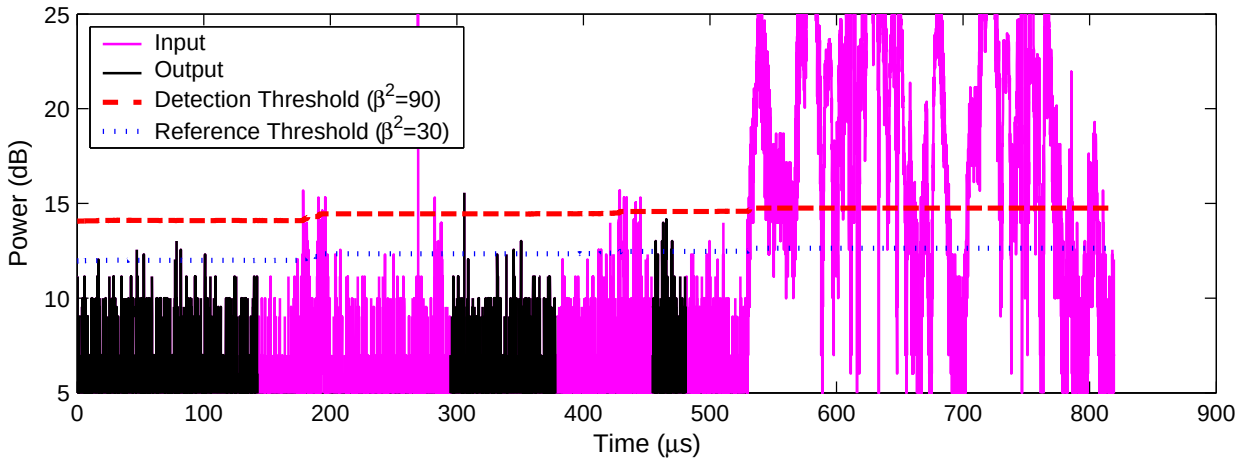

Figure 4. *Detection Threshold* (for detecting pulses) and *Reference Threshold* (for estimating number of pulses). N_{wait}=0, N_{blank}=1536 are used in the APB output results shown

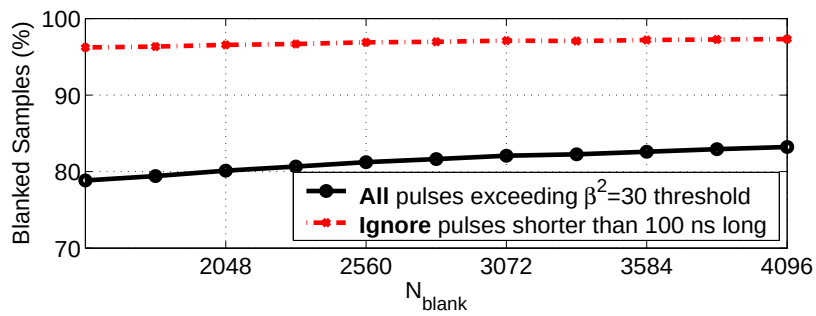

Figure 5. $(P_{in} - P_{out})/P_{in} \times 100\%$, where P_{in} and P_{out} are the number of samples exceeding the "reference threshold" at *input* and *output* of the APB algorithm, respectively.

5.2. Output χ^2 test

Now that means for choosing the APB parameters have been established, it is of interest to quantify the quality of the output data. Because thermal noise power should approach a Gaussian distribution when integrated sufficiently, the χ^2 test against a Gaussian distribution can be utilized to evaluate if the output of the blanker satisfies this expectation. However, the entire dataset for a specific channel cannot be applied in this test, as the mean noise power in a channel will vary as different locations are observed. The test was instead performed using sets of consecutive 5 16K captures, all of which were measured within 5 seconds. The χ^2 statistic using 4 degrees of freedom was computed using data power integrated over 128 samples in order to approach the Gaussian distribution.

Original data from one set of LISA channel 7 data is illustrated in the upper plot of Figure 6 before and after the APB algorithm is applied with $\beta^2 = 40$ and $N_{blank} = 2048$. The presence of interference in the original dataset results in a high χ^2 value of 262.53, indicating that the data are not likely to be from a Gaussian distribution. The lower plot illustrates the χ^2 statistic of the non-blanked data after blanking with $\beta^2 = 40$ versus N_{blank}, and shows a greatly reduced value compared to the pre-blanking case. Critical values based on $\alpha = 1\%$ and 10% (the probability of incorrectly classifying a true Gaussian distribution as non-Gaussian) are also illustrated in the figure. Clearly for this example, the APB output data is much more Gaussian than the input data, particularly for N_{blank} exceeding 1024.

Simulations from other LISA data subsets show similar results, with a few exceptions. In particular, LISA channel 6 (1310 MHz- 1330 MHz) sometimes contains very strong interference from multiple aviation radars, and a large value of χ^2 remains even after blanking with N_{blank} up to 4096. The limitation of the fixed `FIFO_LENGTH` parameter is an issue here, and future work will examine if increasing this parameter can improve these problematic datasets. However it should certainly be expected that there are cases with exceptional RFI corruption for which the APB algorithm cannot retrieve the original noise power.

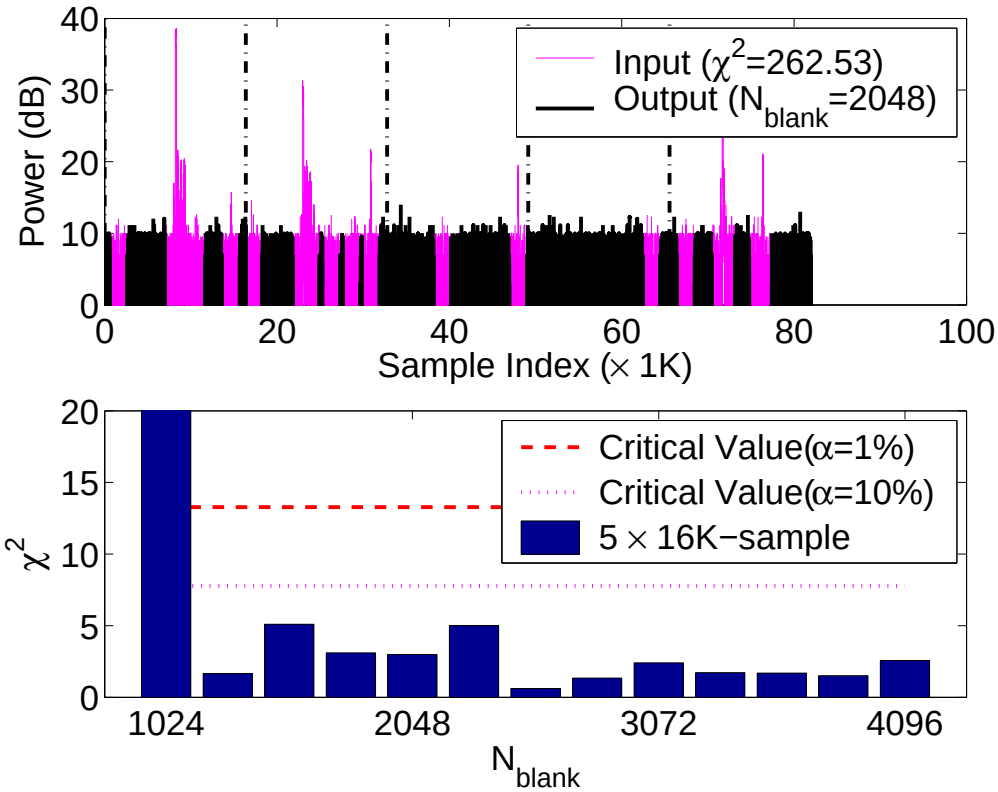

Figure 6. *Top*: Noise power from 5 captures sucessively taken in the vicinity of Wheeling, IL (approx. 39.87° N lat, 80.67° W long) at 22,000 ft. LISA is tuned to channel 7 (1324 - 1344 MHz). Several examples of pulsed RFI are observed. *Bottom*: χ^2 of APB output as N_{blank} is varied (β^2=40, N_{wait}=0).

5.3. Effect of blanking on integrated spectra

The ideal pulse-blanking algorithm would remove only RFI information, without changing properties (particularly the mean power level) of the remaining noise information. One questionable issue of the APB algorithm is the impact of forcing data to zero when pulses are detected. This introduces discontinuities into the signal which may lead to undesired effects on the final output, as well as calibration uncertainties. Note after the APB operation an FFT is performed in the interference suppressing radiometer of Section 2; clearly the impact of blanked samples on the FFT output should be investigated.

To examine these effects, APB outputs were processed through FFT and integration operations. Each 16K LISA capture after blanking was first separated into 32 512 sample "frames" (i.e. a 512 point FFT operation is used). Prior to the FFT, each frame can be categorized as either `BLANK` (contains no non-zero samples), `NO BLANK` (contains no blanked samples), or `PARTIAL BLANK` (some samples are blanked), as shown in Figure 7. The FFT is performed on each frame, the power computed in each FFT bin, and all results in each FFT bin are averaged.

It is clear that the only effect of the `BLANK` category is to decrease the noise power level of the final average. It is trivial to correct for this effect simply by counting the total number of frames and the number of `BLANK` frames. However, the effect of the `PARTIAL BLANK` frames, which contain discontinuities, is more complex. An FFT operation on such a frame clearly will produce a distorted spectrum and a reduced noise power level, with the degree of distortion and power reduction related to the number of blanked samples within the frame. An example `PARTIAL BLANK` spectrum is compared to the corresponding `NO BLANK` average spectrum in Figure 8. The reduction in power level is clearly visible, along with a moderate distortion in the overall shape of the spectrum.

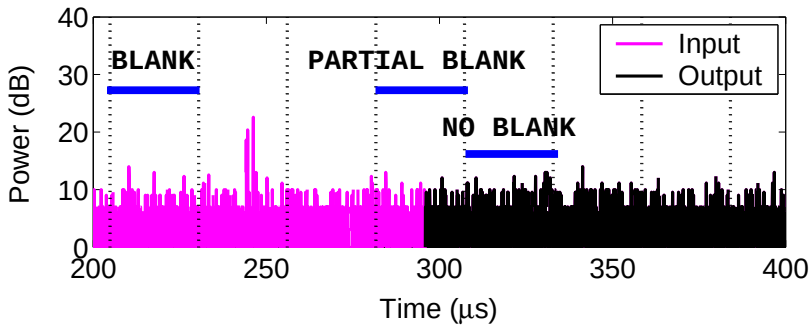

Figure 7. Definition of `BLANK`, `PARTIAL BLANK` and `NO BLANK` frames

Various means for coping with the `PARTIAL BLANK` issue can be conceived. A simple strategy (called method one) is to eliminate such frames from the averaging operation; however this approach may also eliminate a large fraction of the incoming data in high RFI environments. A second approach to retain these frames while correcting for the power level reduction is called "method 2: instantaneous scaling." By Parseval's theorem, the effect of blanking on total average power of the frame can be corrected simply by increasing the power of the computed spectrum by N/N_{rem}, where N_{rem} is the number of non-blanked samples in the frame. This correction is applied to the power level of each frame before including the frame in the average computation. A final approach to is to included all (unscaled) frames in the spectral average operation, and to maintain a separate count of the total number of non-zero time-domain samples included in the average, labeled $N_{tot,\neq 0}$. Only the final average power is scaled by $N_{tot}/N_{tot,\neq 0}$, where N_{tot} is the total number of time domain samples that make up the average operation. This approach is termed "method 3: slow scaling".

The upper plot of Figure 9 compares average spectra from a single Channel 6 LISA 16K capture before any blanking operations with the averages obtained from methods 1 and 2. APB parameters $\beta^2 = 90$ and $N_{blank} = 2048$ were used in the APB algorithm. Both blanking methods are seen to be highly effective in removing contributions from a pulsed interferer centered at approximately 1315 MHz in this example. Results from Methods 1 and 2 are also observed to be very similar, demonstrating that the instantaneous scaling approach is correcting for the reduced noise power level due to blanking effects. To highlight the differences between methods one and two, the lower plot of the figure illustrates the difference (subtracted decibel values) between the method 2 and 1 average spectra. Differences are generally within 1.5 dB in all cases, and appear noiselike, indicating that further averaging would likely make these differences less significant.

The difference between averaged spectra for methods 3 and 1 is also illustrated in the lower plot; errors from method 3 are observed to be somewhat smaller than those from method 2 on average. Clearly method 3 is a simpler operation than that of method 2 (favorable for hardware implementation) since corrections are required at a much slower rate. In addition, method 3 should be preferable to method 2 because method 2 allows `PARTIAL BLANK` frames with a great deal of blanking to be weighted equally in terms of the power averaging computation. However, these frames also have the largest degree of spectral distortion, so reducing their weight should be advantageous. Figure 10 illustrates the mean error for methods 2 and 3 (compared to method 1) for LISA's channels 7 through 13 when averaging results over the entire dataset and the entire frequency spectrum. The mean error is of interest because it indicates the degree to which the average power level is not being corrected properly. Results clearly show the method 3 error (slow scaling) generally to be smaller than that of method 2. A hardware implementation of method 3 is currently in progress for the digital radiometer prototype.

5.4. Frequency domain blanking

Although the digital radiometer prototype does not implement RFI mitigation strategies in hardware after the FFT operation at present, it is of interest to simulate the

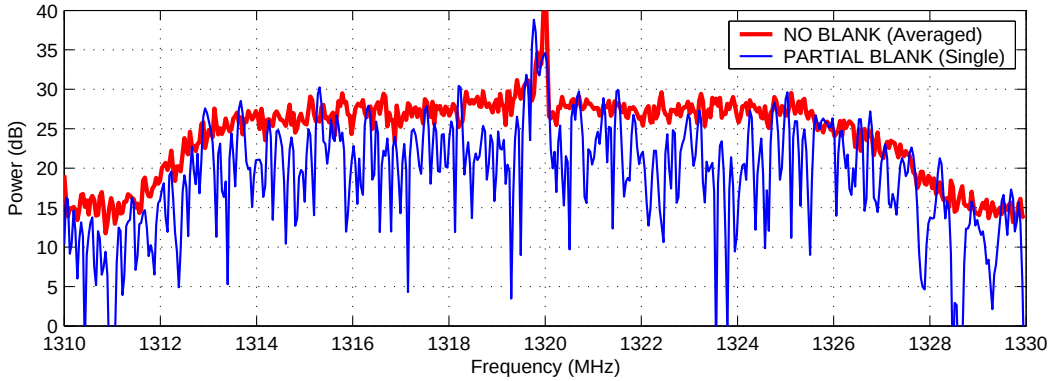

Figure 8. Frequency spectrum of an example `PARTIAL BLANK` frame, compared to the entire `NO BLANK` average for the corresponding capture

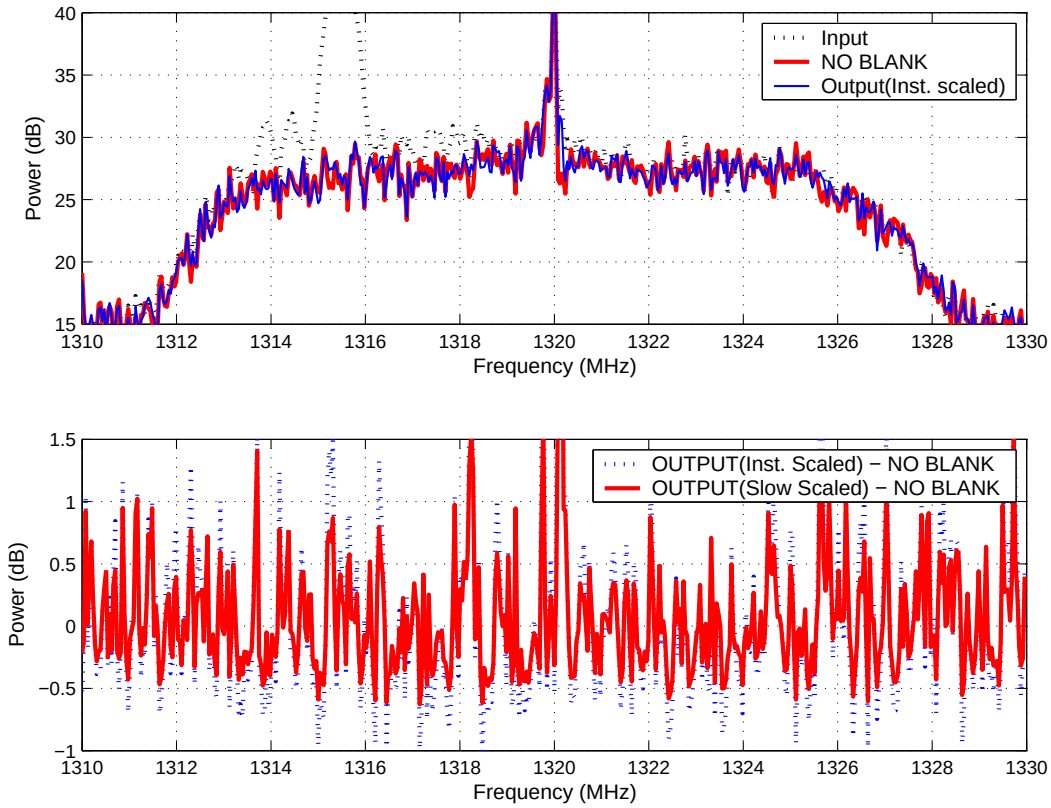

Figure 9. Impact of `PARTIAL BLANK` on the final output. *Top*: Method 2: *Instantaneous Scaled* averages compared to method 1. *Bottom*: The error of methods 2 (inst. scaled) and 3 (slow scaled)

expected performance of such approaches. The "channelization" of the FFT should allow an improved signal-to-noise ratio in detecting pulsed interference within a single FFT bin. A hardware blanking algorithm could conceivably operate on each FFT bin in real time by using a strategy identical to that of the APB processor. Such an approach would allow lower level, rapidly pulsed RFI to be removed if missed by the original APB.

This algorithm was simulated using the LISA data of 145 captures in channel 6. After passing each capture through the time-domain APB algorithm with $\beta^2 = 40$, and $N_{blank} = 4028$, each 16K capture was split into 32 512-point frames. An FFT operation was then applied to each frame, resulting in a total of $32 \times 145 = 4640$ temporal samples for each FFT bin. For each bin, a

second APB algorithm with $\beta^2 = 90$, $N_{blank} = 4$, and $\texttt{FIFO_LENGTH} - N_{wait} = 2$ was then applied to these 4640 samples, with the mean and variance computed by the averaging filter process used in the original APB algorithm. Figure 11 illustrates the average spectra before and after the frequency domain blanking operation, and shows a slight change in results near the center of the spectrum. The effect of the blanker is more obvious in the "max-hold" spectra also illustrated in the plot; "max-hold" refers to the maximum value of the 4640 temporal samples. Clearly a significant degree of RFI is included in this dataset near the center frequency 1320 MHz; the frequency domain blanking operation reduces this interference so that corruption of the average spectrum is less significant.

An alternative approach, "χ^2 blanking" can also be considered if blanking at a slower temporal rate is deemed acceptable. In this method, the χ^2 test is performed on data from each bin. If the χ^2 value exceeds a specified critical value, samples in the dataset exceeding a power threshold are removed and and χ^2 then re-evaluated. This iteration is repeated until the distribution satisfies the χ^2 test for a specified critical value. The requirement for a slower temporal rate here is due to the complexity and iterative nature of this algorithm, which is not suited for integration in hardware. However the algorithm could be applied to already-integrated data as a post-processing step in software.

6. Conclusion and Discussion

Results of the study show the APB approach generally to be effective in reducing corruption from temporally localized RFI. The simulations performed on the LISA dataset, while not completely general, arguably should be representative of a wide range of RFI environments. For these simulations, use of β^2 values ranging from 40 to 90 appears to effective, along with $N_{blank} > 1280$ ($\approx 64\mu$sec blanking window). Use of "slow-scaling" of the output power after blanking was found preferable for correcting errors in the estimate of mean power due to blanking, and averaged spectra after blanking showed only a modest distortion of the underlying noise power. Further studies will continue to explore the effect of other parameters in the APB algorithm (particularly `FIFO_LENGTH`), as well as studies with other data sets. In particular, an airborne version of the digital receiver prototype is currently in progress for RFI observations at C-band [Johnson et al, 2004]. Results from flights of this system will provide a substantial set of RFI observations for further evaluating APB performance in varying RFI environments.

References

Colorado State Univ., "Wakasa Bay Field Experiment Description," http://rain.atmos.colostate.edu/Wakasa.

ElectroScience Laboratory, OSU IIP project document server, *http://esl.eng.ohio-state.edu/∼swe/iip/docserv.html*, 2004.

Ellingson, S. W., *Characterization of Some L-Band Signals Visible at Arecibo*, Technical Report 743467-2,The Ohio State University ElectroScience Laboratory, February 2003.

Ellingson, S. W., and G. A. Hampson, *RFI and Asynchronous Pulse Blanking in the 1230-1375 MHz Band at Arecibo*, The Ohio State University ElectroScience Laboratory Technical Report 743467-3, Feb 2003a.

Ellingson, S. W., and G. A. Hampson, "Mitigation of Radar Interference in L-Band Radio Astronomy," *Astrophysical J. Supp.*, vol. 147 (1), pp. 167–176, 2003b.

Ellingson, S. W., G. A. Hampson, and J. T. Johnson, "Characterization of L-band RFI and Implications for Mitigation Techniques," *Proc. 2003 IGARSS*, Toulouse, July 2003a.

Ellingson, S. W., G. A. Hampson, and J.T.Johnson, *Design of and L-Band Microwave Radiometer with Active Mitigation of Interference*, *Proc. 2003 IGARSS*, Toulouse, July 2003b.

Ellingson, S. W., and J. T. Johnson, *Airborne RFI Measurements over the Mid-Atlantic Coast using LISA,* Technical Report, The Ohio State University ElectroScience Laboratory, January 2003.

Hampson, G. A., J. T. Johnson, and S. W. Ellingson, "Design and demonstration of an interference suppressing microwave radiometer," *IEEE Aerospace Conference 2004*, conf. proc., 2004.

Johnson, J. T., and S. W. Ellingson, *Airborne RFI Measurements Using LISA During Transit to and in the Wakasa Bay Campaign,* Technical Report, The Ohio State University ElectroScience Laboratory, July 2003.

Johnson, J. T., A. J. Gasiewski, G. A. Hampson, S. W. Ellingson, R. Krishnamarachi, and M. Klein, "Airborne radio frequency interference studies at C-band using a digital receiver," *IGARSS 2004,* conf. proc., 2004.

N. Niamsuwan, ElectroScience Laboratory, The Ohio State University, Department of Electrical and Computer Engineering, 1320 Kinnear Road, Columbus, OH 43212 (niamsuwan.1@osu.edu)

J. T. Johnson, ElectroScience Laboratory, The Ohio State University, Department of Electrical and Computer Engineering, 1320 Kinnear Road, Columbus, OH 43212 (johnson@ee.eng.ohio-state.edu)

S. W. Ellingson, Virgiania Tech, Department of Electrical and Computer Engineering, 340 Whittemore Hall, Blacksburg, VA 24061 (ellingson@vt.edu)

(Received _____.)

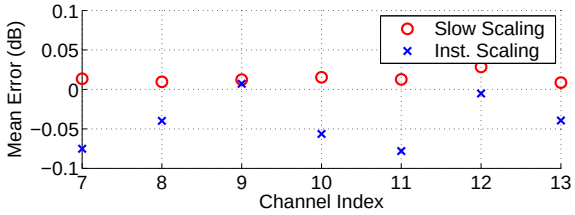

Figure 10. Mean error in total power estimate from methods 2 and 3 in LISA channels 7-13

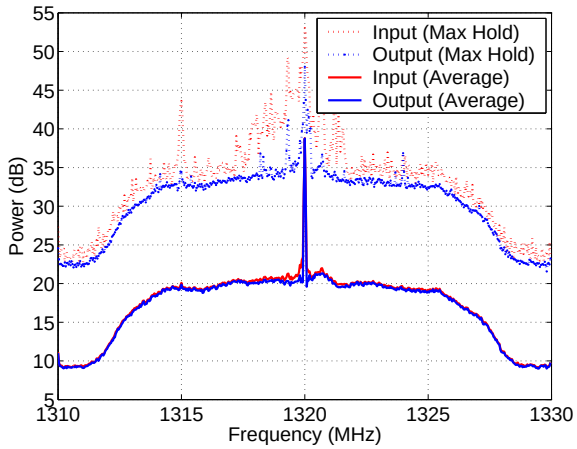

Figure 11. Frequency-domain APB algorithm results.