Mitigation of Radar Interference in L-Band Radio Astronomy

S.W. Ellingson & G.A. Hampson

The Ohio State University, ElectroScience Laboratory, 1320 Kinnear Road, Columbus OH 43212 USA; ellingson.1@osu.edu

ABSTRACT

The 1215–1400 MHz band is important for spectroscopy of HI at high redshift, pulsar work, and SETI. Observations at these frequencies are complicated by pulsed interference from ground-based aviation radars. In this paper, we characterize one such radar received at the Arecibo Observatory using coherently-sampled datasets obtained during a recent observation. Using this data, we demonstrate some simple methods for detection and removal of the radar pulses. One of these uses a coherent subtraction technique that has not previously been applied to the radar problem. This new technique provides an alternative to blanking, which is undesirable in pulsar and SETI work. We demonstrate that the radar studied in this paper can be suppressed by at least 16 dB in integrated spectra using the coherent subtraction technique. The maximum single-pulse power observed at the output of the canceler is ~ 15 dB less. The primary limitation appears to be the detector performance; as a result, the performance using blanking is about the same. Also, we demonstrate that the matched detector for pulses from this radar is relatively insensitive to astronomical transients (e.g., giant pulses), and quantify the risk of such transients being falsely identified as radar pulses. The techniques described in this paper can easily be adapted to radar waveforms other than the one examined here.

Subject headings: instrumentation: detectors — Methods: analytical – radio continuum: general techniques: radio astronomy — radio lines: gerneral

1. Introduction

The 1215–1400 MHz band is important for spectroscopy of HI at high redshift, pulsar work, and SETI. Observations at these frequencies are complicated by interference from the legitimate transmissions of ground-based aviation radars (GBARs). These radars typically transmit pulsed fixed-frequency or chirped sinusoidal waveforms with pulse lengths

of 2–400 μs, 1–27 ms between transmitted pulses, and bandwidths on the order of 1 MHz (Ellingson 2002c)[1]. Transmit powers range from 10^3 W to 10^6 W into a highly-directional antenna which rotates in azimuth with period on the order of 10 s. Pulses are easily detected through the sidelobes of radio telescopes at least 10^2 km away, and further when the radar beam is directly pointed at the receiving site or a reflecting object, such as an aircraft. The carrier frequency separation between radars detected at any given site can be 10 MHz or less, making observations in larger bandwidths (desirable for pulsar work, for example) a challenge. The spectrum between carrier frequencies is also corrupted to some degree due to the extended sidelobes associated with the pulse edges, which distorts noise baselines and overwrites spectral features. Pulsar surveys are affected because GBAR pulses tend to saturate dedispersion/periodicity searches (especially at low dispersion measures (DMs)), masking pulsar candidates.

Some unique properties of GBARs relative to other forms of interference make them especially difficult to mitigate. Although the transmission duty cycle is relatively low (typically less than 0.1%), data editing is complicated by the persistent nature of the transmission coupled with the short period between pulses. To be practical, the data editing should be automated. Unfortunately, most present-day instrumental configurations are not suitable for post-observation editing of GBAR pulses; for example, most spectrometers and pulsar machines output partially-reduced data at intervals that are often many times longer than a radar pulse. Although faster recording is possible (e.g., coherent sampling of bandwidths greater than 10 MHz or so, with recording directly to tapes or hard drives), this yields quantities of data which may be too large to manage. In this situation, automated, real-time mitigation of GBARs concurrent with the observation may be desireable.

A number of real-time techniques have been proposed and developed to various degrees. Friedman (1996), Weber *et al.* (1997), and Leshem, van der Veen & Boonstra (2000) each describe various methods for detecting impulsive interference (not specifically GBAR pulses) and blanking the output accordingly. The National Astronomy and Ionosphere Center (NAIC) developed a device for real-time mitigation of a strong local radar (the same one examined in this paper, in fact) at the Arecibo Observatory (Puerto Rico).[2] This device works by tracking the known pattern of the timing between pulses for this particular radar, and then blanking the output of the receiver in a time window around the expected pulse arrival times. Reflections from terrain features and aircraft generate additional copies of the

[1]See also `http://www.gb.nrao.edu/~rfisher/Radar/analysis.html`, `http://esl.eng.ohio-state.edu/~swe/ska1/rparadar.pdf`, and `http://esl.eng.ohio-state.edu/~swe/docs/ar_char.pdf`.

[2]`http://www.naic.edu/~astro/spectral_line/handbook/node16.html`

pulse which arrive after the "direct path" pulse (as is demonstrated in the Appendix); thus the blanking window must be many times longer than the pulse length to be effective. For this reason, the NAIC blanker is typically configured to blank a window 400 μs long ($\sim$67 times longer than the pulse itself). This device is occasionally employed in spectroscopy at Arecibo (P. Perillat 2002, private communication).

For future backends and new radio telescopes, there is much interest in devices of a similar nature, but with some improvements. First, in order to mitigate weak radars, it is desirable to optimize detection performance. This infers matched filtering, which requires detailed knowledge of the pulse waveform. Second, it is desired to have the ability to detect pulses singly, as opposed to relying on detections of pulse patterns, which may be initially unknown or subject to change. Third, it would be useful to be able to have the option to coherently subtract radar pulses from the telescope output, as opposed to simply blanking the output. This feature would be useful for pulsar work, since blanking may introduce undesirable artifacts in dedispersion and periodicity measurements; and for SETI, where blanking complicates signal detection processes which are already tedious and complex.

A special case of problems in pulsar astronomy and SETI is the detection of intermittent or transient events; e.g. giant pulses from pulsars (Staelin & Reifenstein 1968[3], Cognard *et al.* 1996). No matter how well designed the radar pulse detector is, it is possible that astronomical transients could be mistaken for radar pulses. Thus, another consideration in the development of new GBAR mitigation techniques is the sensitivity of the proposed detectors to various astronomical transients. The goal is obviously to develop detectors with the best possible sensitivity to GBAR pulses subject to the constraint that the likelihood of false alarms associated with interesting astronomical transients is acceptable.

In this paper, we characterize a commonly-seen GBAR waveform and it's associated propagation characteristics. Based on this information, we develop a simple technique for detection of single radar pulses and characterize it's performance. We then test this detector with two different pulse removal techniques: pulse blanking, and a coherent subtraction (canceling) technique that exploits *a priori* knowledge of the pulse waveform. The canceling strategy has not previously been applied to GBAR mitigation, and is attractive because it preserves integration time and avoids the problem of introducing blanking artifacts in time-domain processing.

This paper is organized as follows. In Section 2, we describe the collection and analysis of data pertaining to the GBAR of interest in this paper, including the development of a suitable single-pulse detection technique. Section 3 discusses various methods for removing detected

[3]The Crab Pulsar was discovered via detection of it's giant pulses

pulses, including blanking and canceling. Section 4 summarizes the performance of these methods when applied to the data set presented in Section 2. Section 5 discusses the problem of falsely detecting interesting astronomical transients as radar pulses. Section 6 discusses some ideas for improving performance and the applicability of this work to other GBARs. Finally, a brief analysis of the GBAR multipath propagation characteristics encountered at Arecibo – which has implications in the design of mitigation techniques – is presented as an Appendix.

2. Characterization of the 1350 MHz GBAR at Arecibo

Very little technical information on GBARs is publicly available. However, much has been learned from a number of monitoring programs and experiments over recent years. Prominent among these is the interference monitoring program at the Arecibo Observatory, which has gone to great lengths to identify and characterize local radar transmitters and coordinate with them when possible.[4]

In this paper, we focus on the notorious 1350 MHz GBAR received at Arecibo. This radar is located at Pico del Este, PR (near San Juan), about 85 km to the east over rugged terrain. Previous observations have determined that this radar uses fixed-frequency CW modulation, a pulse length ~ 6 μs, and inter-pulse spacings following a repeating pattern of discrete values between 2.6 ms and 3.3 ms (P. Perillat 2002, private communication; Ellingson 2002c). In this section, we report on a new observation of this radar with the aim of better understanding the waveform, which is important in the design of mitigation techniques.

2.1. Data Collection

The data were collected on 2002 November 3 at Arecibo using the "L-Wide" receiver with a custom back end developed at the Ohio State University (Ellingson, Johnson & Hampson 2002)[5]. The back end shifts the desired center frequency to 150 MHz and samples at 200 MSPS using 10 bits. 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 A/D is followed by digital hardware that downconverts the desired center frequency to 0 Hz (so now the samples are complex-valued), filters to 50 MHz bandwidth, reduces the sample rate by

[4]NAIC Spectrum Management web page: `http://www.naic.edu/~rfiuser/`

[5]See also `http://esl.eng.ohio-state.edu/~swe/ska1/naic0211.pdf`

a factor of 2, and then upconverts to a center frequency of +25 MHz (the last step was a legacy function which was easier to leave in place than to modify for this experiment). The data emerges from this process in 16-bit "I" + 16-bit "Q" format at 100 MSPS. The data are collected in a first-in first-out (FIFO) buffer of length 256K samples. When the buffer fills (in 2.6 ms), the acquisition is halted while the samples are transferred to a PC, which requires about 1 s. When the transfer is complete, the buffer is triggered and another 256K samples are collected. This process repeats continuously.

The dataset examined in this report consists of a set of 400 data blocks collected at a center frequency of 1350 MHz, representing about 1 s of observation collected over about 13 min in real time.[6] This particular radar is known to drive the L-Wide receiver into compression when the radar is pointed directly at Arecibo. However, this condition exists only for a few milliseconds, once per $\sim$ 10 s revolution. Presumably for this reason, five data blocks were found to contain clipping (and probably also front-end compression, although this is difficult to verify) and were discarded. It is important to note that only blanking is reasonable when the data contains clipping or compression. Coherent subtraction – as proposed later in this paper – is not effective for the $\sim$ 1% of the time that the radar is pointed directly at Arecibo.

The remaining blocks were processed as follows: The center frequency was shifted to zero and a bandlimiting filter was applied by convolution in the time domain. The filter was a 119-coefficient Kaiser window with a shape parameter of 0.4, designed to have a bandpass which is flat to within a few tenths of a dB inside $\pm$2.5 MHz and which cuts off ($>$ 50 dB attenuation) beyond $\pm$5 MHz. The resulting noise bandwidth for the purposes of this study was about 5 MHz, and there appear to be no other signals present in the passband.

After each sample block is filtered as described above, the pulse modulation can be explored. Figures 1 and 2 show an example of a single strong pulse with relatively little multipath. The pulse width is 5.8 μs between half-power ($-$3 dB) points and 6.5 μs between $-$20 dB points. The nearly-constant magnitude and phase over the pulse period confirm that this radar is of the "pulsed CW" variety. The small variations in the magnitude and phase during the pulse period could be due to either transmitter-induced artifacts or scattering within $\sim$300 m of the transmitter or receiver.

[6]All data collected is freely available; contact the authors for distribution information.

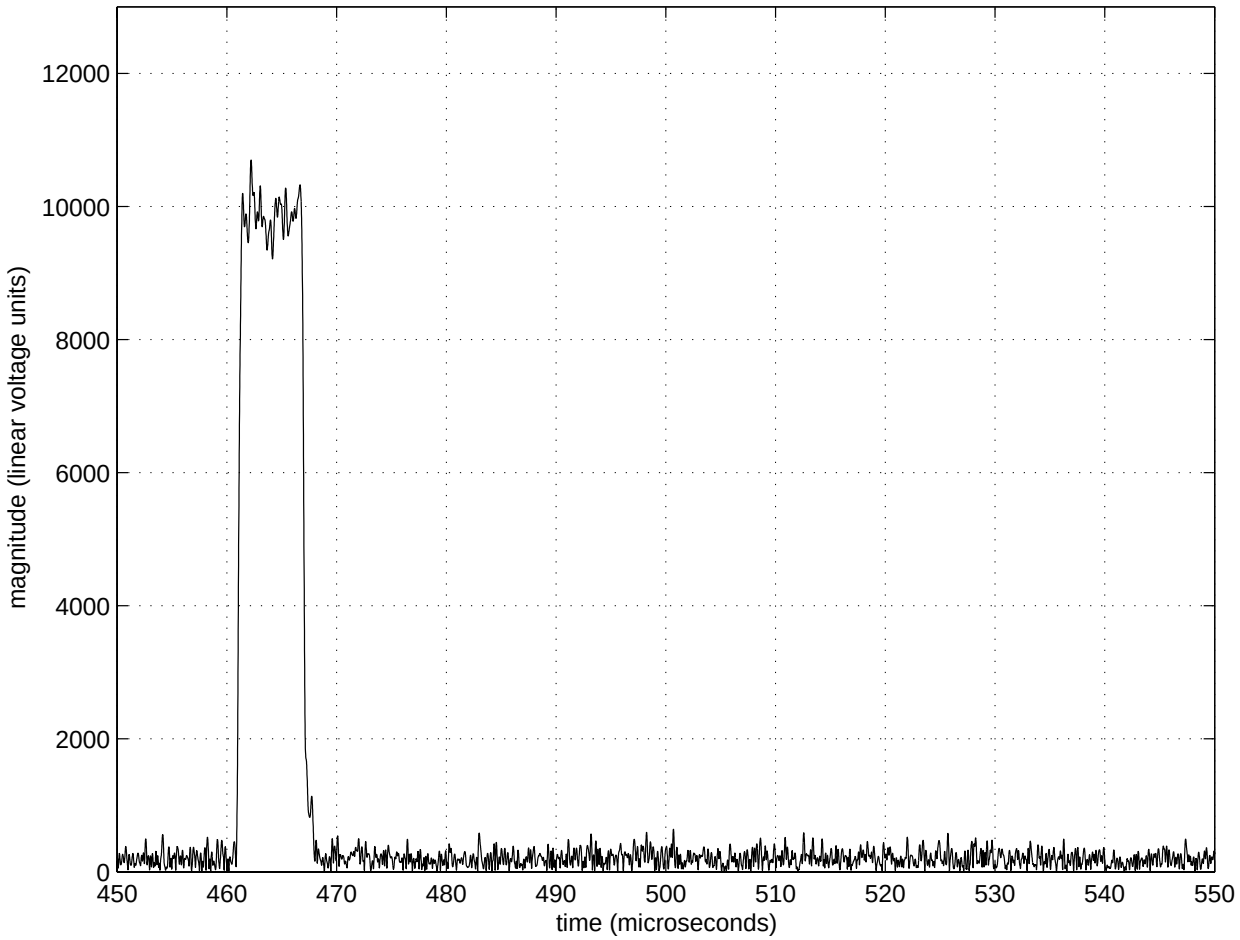

Fig. 1.— An example of a single strong pulse with relatively little multipath.

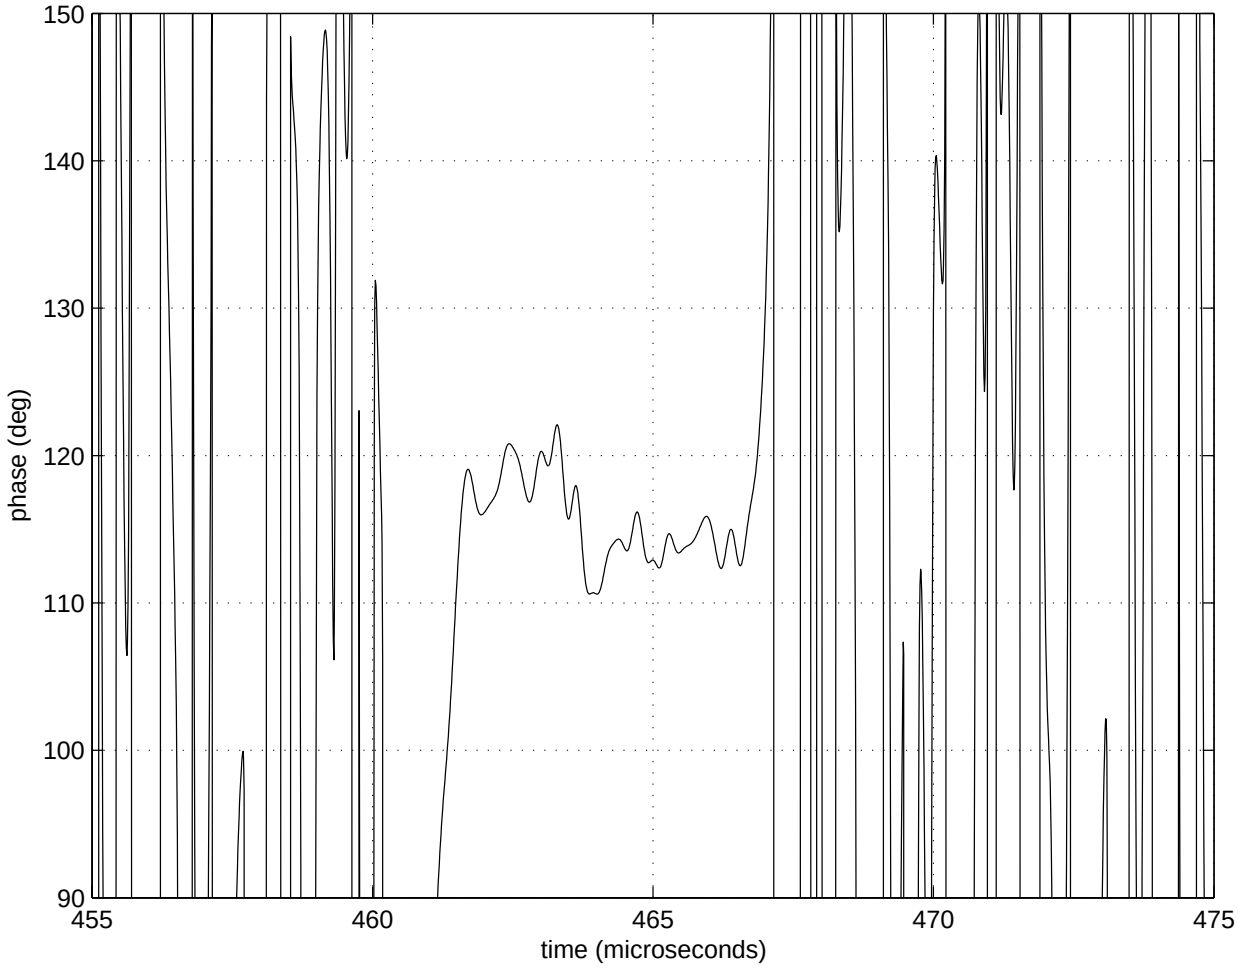

Fig. 2.— Phase of the pulse shown in Figure 1.

2.2. Pulse Detection

To proceed further, a scheme for detecting individual radar pulses is required. In the absence of multipath and assuming additive spectrally-white Gaussian noise, the optimum detector of a single pulse is a filter matched to the transmitted pulse waveform, followed by a threshold test (Kay 1998). The threshold is set to the smallest value that yields a false-alarm rate (FAR) that is acceptable to the user. Let us initially assume the following: (1) A simple model of the pulse as 6 μs of a sinusoid is a reasonable approximation to the transmitted pulse, (2) Time-variable multipath contributions do not significantly degrade the pulse before it is received, and (3) Pulses do not overlap. Upon inspection, we found that many of the pulses in the data set are spectacularly distorted by multipath; i.e., assumptions (2) and (3) are often violated. Nevertheless, assumptions (1)–(3) do seem to yield a useful detector, as will be shown below.

The proposed detection algorithm is as follows:

1. Compute the matched filter output $y(t) = \|q^*(t) * x(t)\|$, where $q(t)$ is the waveform of the model pulse described above, "$*$" denotes convolution, and the superscript "*" denotes conjugation.

2. Compute the "local" (to be quantified shortly) mean and standard deviation of $y(t)$ in the absence of radar pulses. Let these quantities be $m(t)$ and $\sigma(t)$, respectively.

3. A detection is declared when $y(t) - m(t) \geq \beta\sigma(t)$, where β is the user-selected threshold which sets the FAR.

In Step 2, we propose calculating $m(t)$ and $\sigma(t)$ from the samples of $y(t)$ within a 2.6 ms (equal to the length of a block in the data set) window, after removing the largest 10% of the samples. This works because the transmit duty cycle of most GBARs is usually much less than 1%; therefore, pulse signatures are effectively eliminated from $y(t)$ using this procedure. It should be noted, however, that this method may also discard outliers in the noise distribution, which will tend to artificially reduce the measured value of $\sigma(t)$ and thereby increase the FAR. In particular, if this method is used and the largest pulse signal-to-noise ratios (SNRs) in the window approach zero (or if radar pulses are not present), then β must be increased to maintain an acceptable FAR. In the present data set, we see large-SNR pulses in nearly every block; thus we shall assume the use of the adaptive method from this point forward.

It should also be pointed out that adaptive estimation of $m(t)$ and $\sigma(t)$ is not strictly necessary. For example, if the system temperature does not vary significantly with time,

then this measurement could be done just once. Even if the system temperature does vary significantly, one can simply to set β sufficiently high to ensure that the maximum acceptable FAR is never exceeded. However, this latter approach will certainly degrade detection sensitivity.

It is not strictly necessary to perform the detection in the time domain, as implied above. An alternative is to operate on the Fourier-transformed input, in which case the matched filter can be implemented as a multiplication of spectra, as opposed to a convolution of time series. This approach is theoretically equivalent and might have advantages in certain implementations.

The detector defined by steps 1–3 above only detects the onset of the pulse; it does not precisely indicate it's position in time. With no other information, the timing of strong pulses will tend to be biased earlier in time with respect to that of weaker pulses. This may be a problem for certain pulse removal techniques, such as pulse canceling. To overcome this problem, it is recommended to calculate the pulse position using the zero-crossing of the first derivative $dy(t)/dt$ of the matched filter output. In practice, a sufficient approximation of $dy(t)/dt$ can be obtained from the slope of $y(t)$ calculated between discrete samples separated by about one-fourth of the pulse length.

Finally, it should be noted that alternative, suboptimal but potentially useful detection algorithms exist. Perhaps the simplest is to assume the pulse has the form of an impulse function (equal to 1 at $t = 0$, equal to 0 otherwise). This is equivalent to having no matched filter; thus, all signals in the passband are detected with equal, albeit poor, sensitivity. Nevertheless, blankers using this detection strategy have been shown to be surprisingly effective (Ellingson *et al.* 2002, Ellingson & Hampson 2002).[7]

2.3. A Refined Model Pulse

To develop the best possible detector, the original transmitted pulse shape – free of the time-variable distortions associated with multipath – is desired. A reasonable approximation is obtained from the mean over many strong (i.e., high SNR) pulses, under the assumption that the contributions of the time-variable multipaths average to zero. To calculate this, we first identified all of the pulses in the dataset which were 10σ or better (i.e., very high SNR) at the output of the $q(t)$-matched detector described above. The detected pulses were then coherently averaged using the following procedure: (1) The pulses were aligned in time; (2)

[7]See also `http://esl.eng.ohio-state.edu/~swe/docs/ap021109a.pdf`.

The mean magnitude and phase (a single complex constant) over the duration of each pulse was determined; (3) Each pulse was divided by the value from (2); and then finally (4) The mean waveform was obtained from the mean across the pulses from (3) for each time sample. The result is shown in Figures 3, 4, and 5. Henceforth, we shall refer to this waveform as $p(t)$.

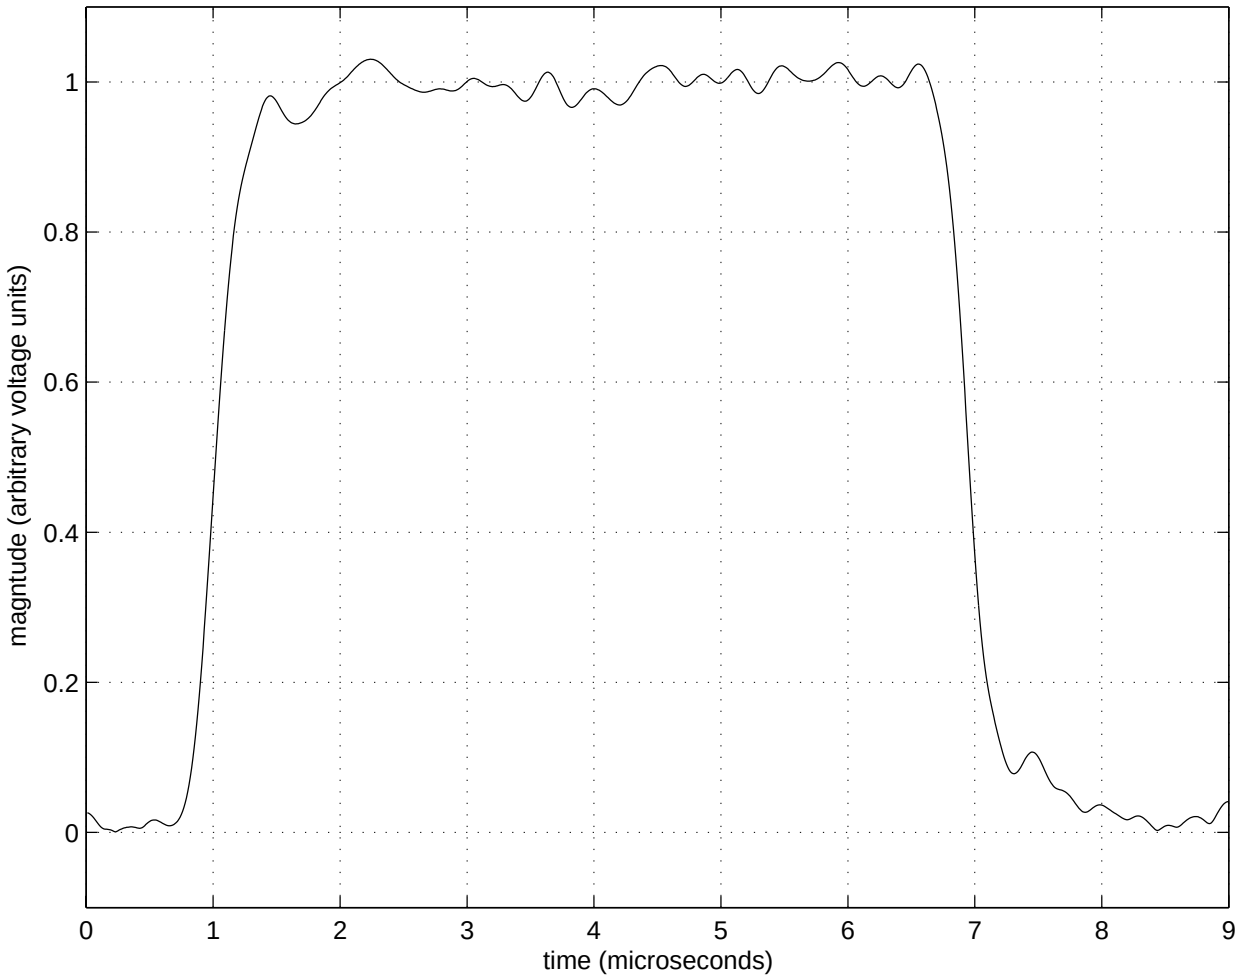

Fig. 3.— Magnitude of the refined model pulse, $p(t)$.

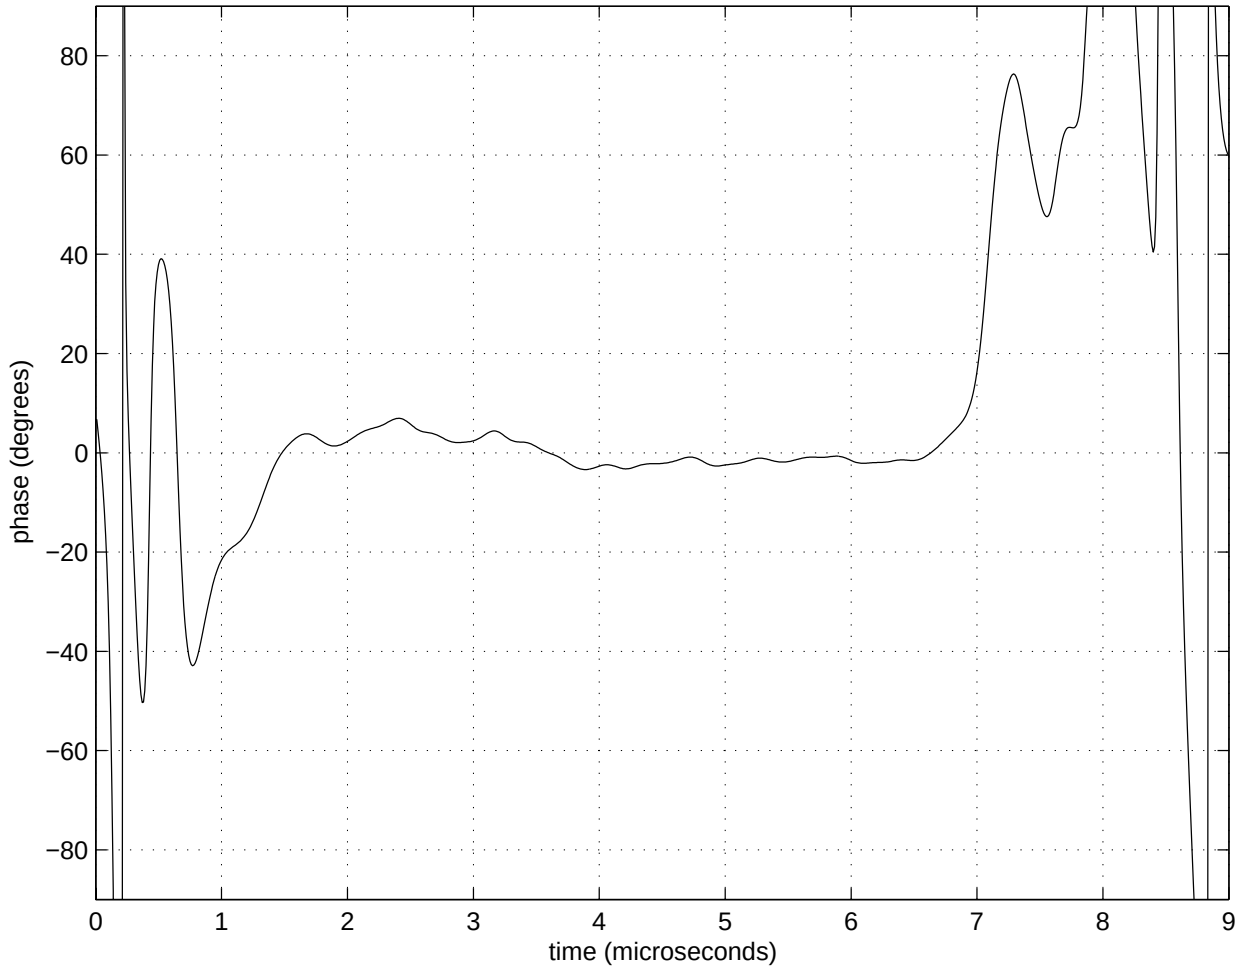

Fig. 4.— Phase of the refined model pulse, $p(t)$.

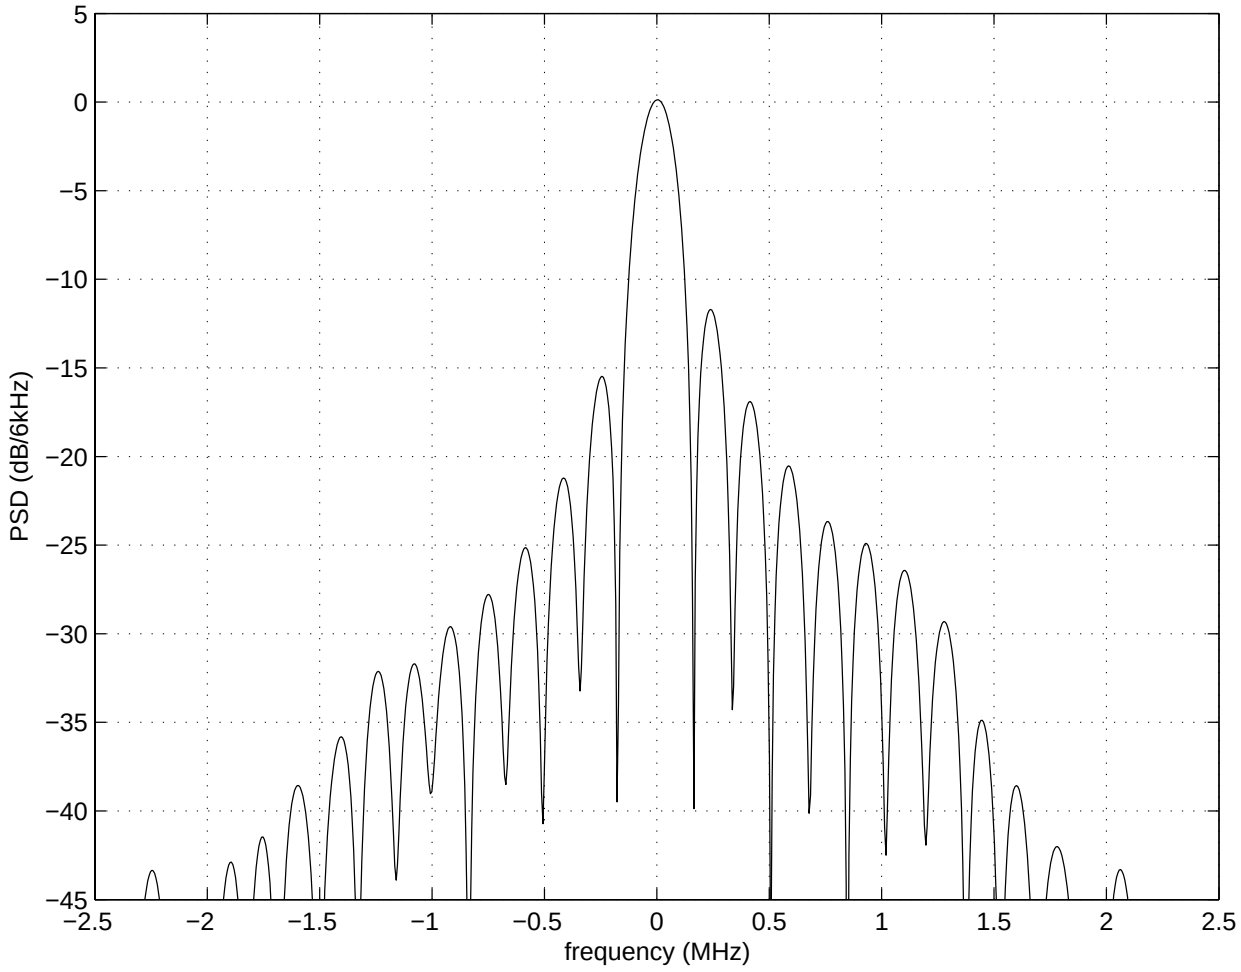

Fig. 5.— Power spectral density of refined model pulse $p(t)$.

2.4. Pulse Magnitude Distribution and Threshold Selection

Using the same detector except with $p(t)$ in lieu of $q(t)$, we recorded the magnitude of all pulses detected at 3σ or better. The result, shown in Figure 6, suggests a choice of $\beta \sim 10$ for unambiguous detection (negligible FAR), or $\beta \sim 5$ for aggressive detection (high sensitivity, albeit at an elevated FAR).

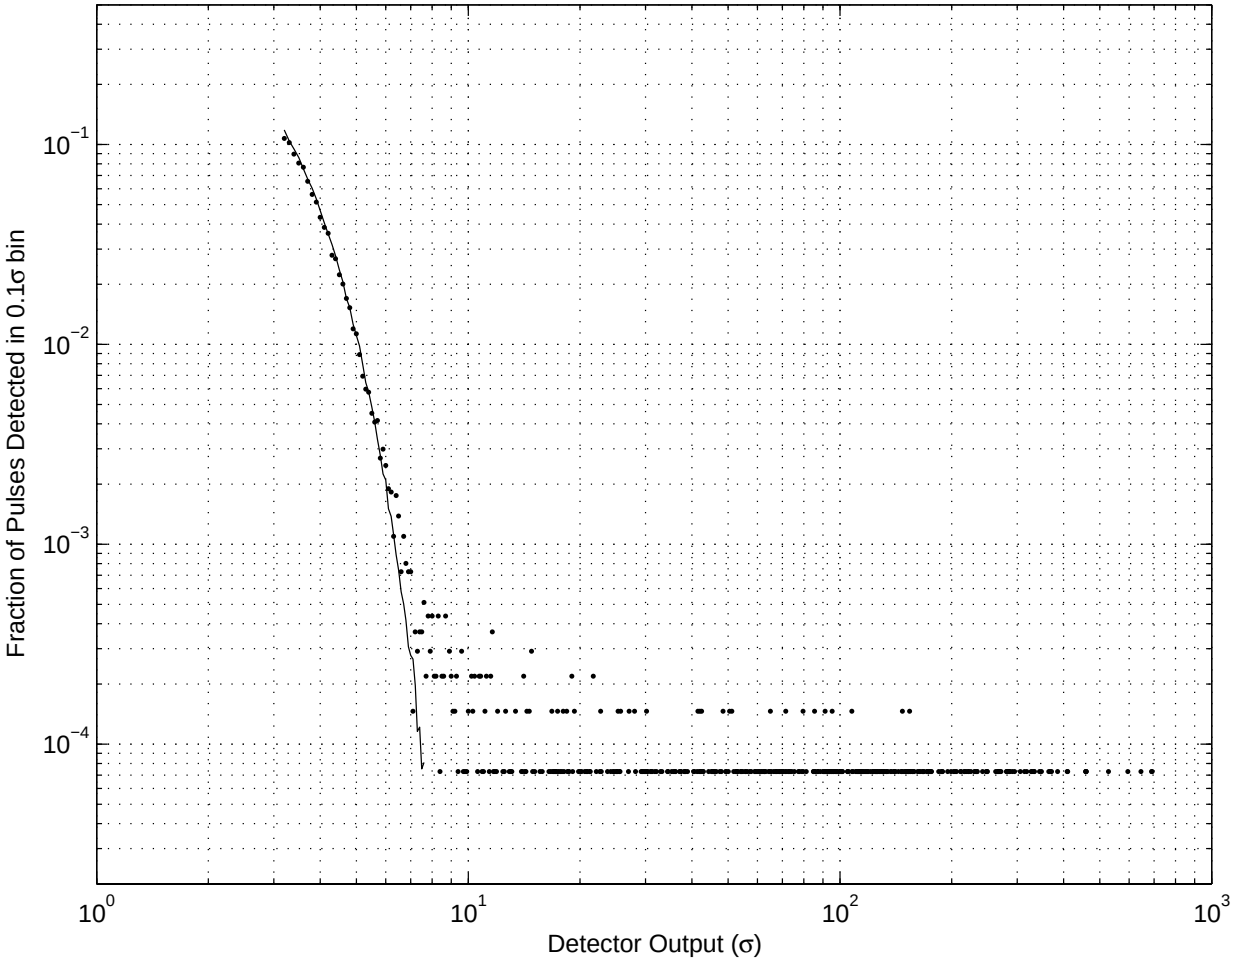

Fig. 6.— *Dots:* Fraction of all pulses detected over 3σ at the output of the $p(t)$-matched detector in bins 0.1σ wide. *Line:* Same result for perfect Gaussian white noise input. Note that detections become unambiguous (distinct from spurious noise detections) between 5σ and 10σ.

Figure 7 shows the detection sensitivity for $\beta = 5$ and $\beta = 10$. Note that the detector is able to detect pulses that are much weaker than the noise due to the "processing gain" associated with the matched filter.

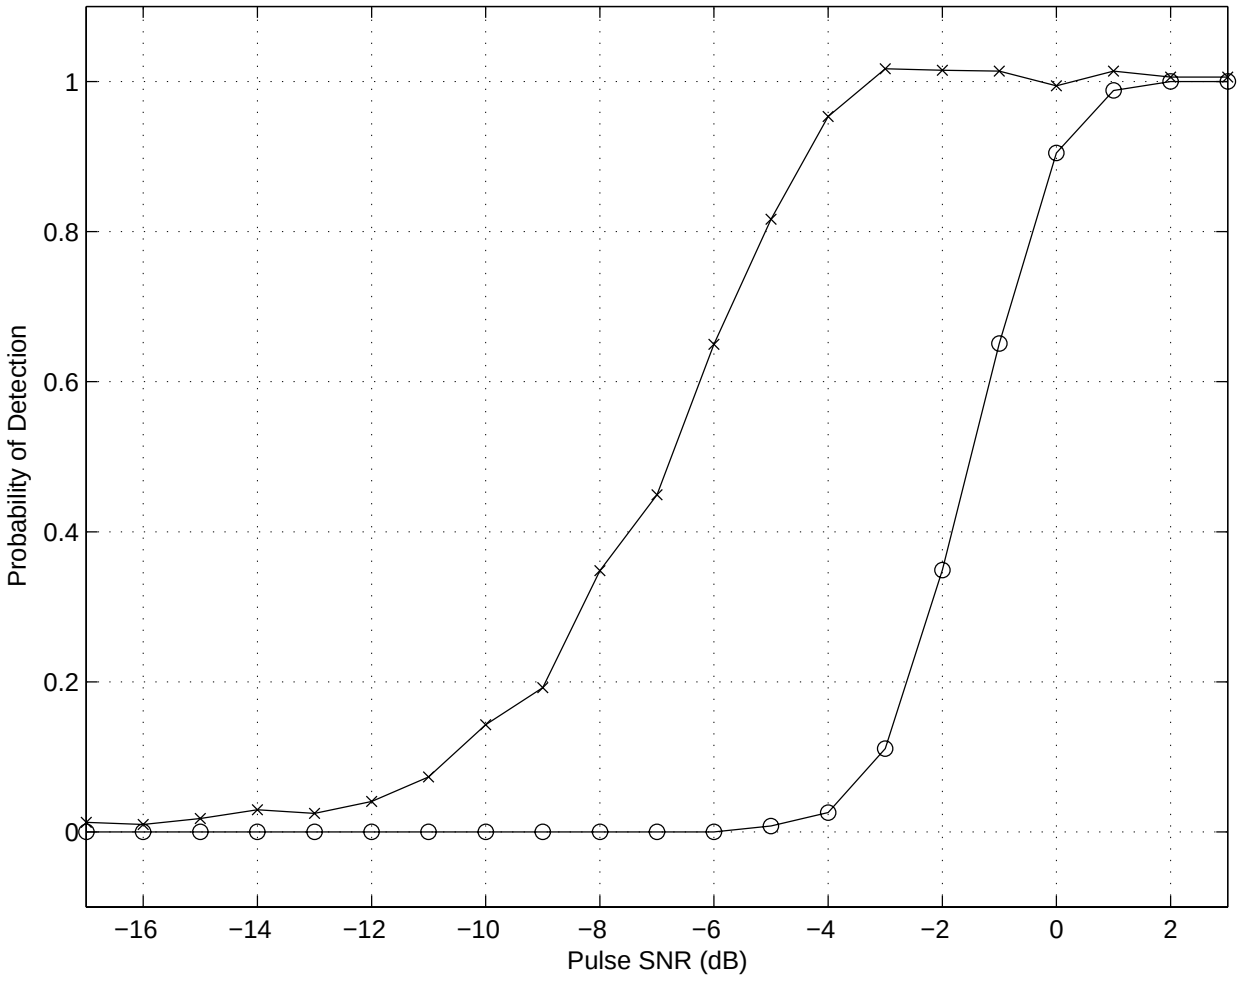

Fig. 7.— Pulse detection performance for $\beta = 5$ ("×") and $\beta = 10$ ("○"). "Pulse SNR" refers to the instantaneous SNR (when the pulse is fully "on") in a 5 MHz noise bandwidth. To obtain constant FAR, these results are generated with $m(t)$ and $\sigma(t)$ held fixed at their nominal "noise only" values (see text). Note that the $\beta = 5$ results can be slightly greater than one due to false alarms.

3. Pulse Removal Techniques

In this paper, we focus on methods of mitigating GBAR pulses which follow the two-step process of (1) detection, followed by (2) removal. For the remainder of this paper, we assume the use of the $p(t)$-matched detector described in the previous section. Now, we consider two classes of techniques for pulse removal: blanking, and canceling.

3.1. Blanking

There are at least two different strategies for blanking detected pulses. We refer to the technique used by NAIC's blanker (described in Section 1) as "delay window blanking".[8] As discussed in the Appendix (see Figure 16), blanking a window beginning about 10 μs before the largest pulse copy and extending out to 300 μs (perhaps more) eliminates all multipath components as long as at the strongest single component is detected. Thus, delay window blanking is useful when it is important to eliminate as much interference as possible, and long gaps in the data can be tolerated. Such an application is HI spectroscopy. The disadvantage of delay window blanking for spectroscopy is the loss of integration time: For the radar studied in this paper, a blanking window 400 μs long will result in an increase of about $\sim$14% in observating time to acheive the same sensitivity.

An alternative is pulse blanking; i.e., blanking only for a period equal to the length of the pulse, and only when a pulse is detected. With pulse blanking, one runs the risk of failing to blank pulses which are too weak to detect, but are collectively strong enough to degrade the quality of the observation. The advantage is that the loss of integration time is much less.

3.2. Canceling

In this paper, we propose pulse canceling, a technique for removing pulses using coherent subtraction, exploiting our knowledge of the pulse waveform as embodied by $p(t)$. Similar concepts have been proposed by Miller, Potter, & McCorkle (1997) for the mitigation of strong narrowband communications signals in ultrawideband radar applications, and also by Ellingson, Bunton, & Bell (2001) to suppress interference from GLONASS satellites in

[8] A similar technique is also discussed in `http://www.gb.nrao.edu/~rfisher/Radar/analysis.html` as "time window blanking".

L-Band OH spectroscopy. To apply this strategy to GBAR pulse mitigation, we assume that the propagation channel can be modeled as a complex-valued constant α over the period of a single pulse. In other words, it is assumed that multipath-generated copies of the pulse do not overlap. As mentioned above, we are certain that this assumption is often violated; however, it appears that most of the violations are due to local scattering, resulting in combinations of one strong pulse with many weaker copies. Therefore, this assumption yields a useful model, even if it is not literally true. Since the radar of interest uses a CW modulation, estimating α turns out to be a matter of estimating the magnitude and phase of a sinusoid in noise. If we assume that the noise is additive white gaussian noise, then this becomes one of the canonical problems of estimation theory: $\hat{\alpha}$, the optimum estimate of α, is simply the mean value of the noisy received signal $x(t)$ over the period of the pulse, after shifting the center frequency to zero (Kay 1993). Then, pulse removal is effected by subtracting $\hat{\alpha}p(t)$ from $x(t)$ for each pulse detected.

It should be noted that it is important use $p(t)$ in this technique, as opposed to a simple truncated sinusoid ($q(t)$). Although there is not much difference in the *detection* performance, there is a dramatic difference in the canceling performance because $p(t)$ has the correct asymmetric sidelobes (as shown in Figure 5), whereas $q(t)$ does not.

4. Performance of GBAR Mitigation

In this section, we examine the performance of the $p(t)$-matched detector combined with pulse removal, both by pulse blanking and pulse canceling. The results are shown in Figure 8–11. For $\beta = 10$, pulse canceling achieves greater than 16 dB of suppression in the integrated spectrum. Pulse blanking performs marginally, but not significantly, better. The dataset used here is not suitable for the evaluation of delay window blanking, since blanking windows might extend beyond the leading and trailing edges of individual data blocks. However, we anticipate that delay window blanking would have the best overall suppression performance (since it is not limited by the need to detect weak multipth pulses), albeit with the relatively large gaps and reduction in effective observation time noted earlier.

Reducing β to 5 results in a slight improvement for both techniques. Since $\beta = 5$ detects pulses well below the noise level, the limited suppression is probably due to pulses missed due to various model violations; e.g., overlapping pulses. The max hold spectra shown in Figure 12 suggest that the principle weakness of the technique lies with the detector: clearly, some strong pulses are not being detected. However, it is interesting to note that both techniques are equally effective in reducing the power of the maximum single pulse by ~ 15 dB.

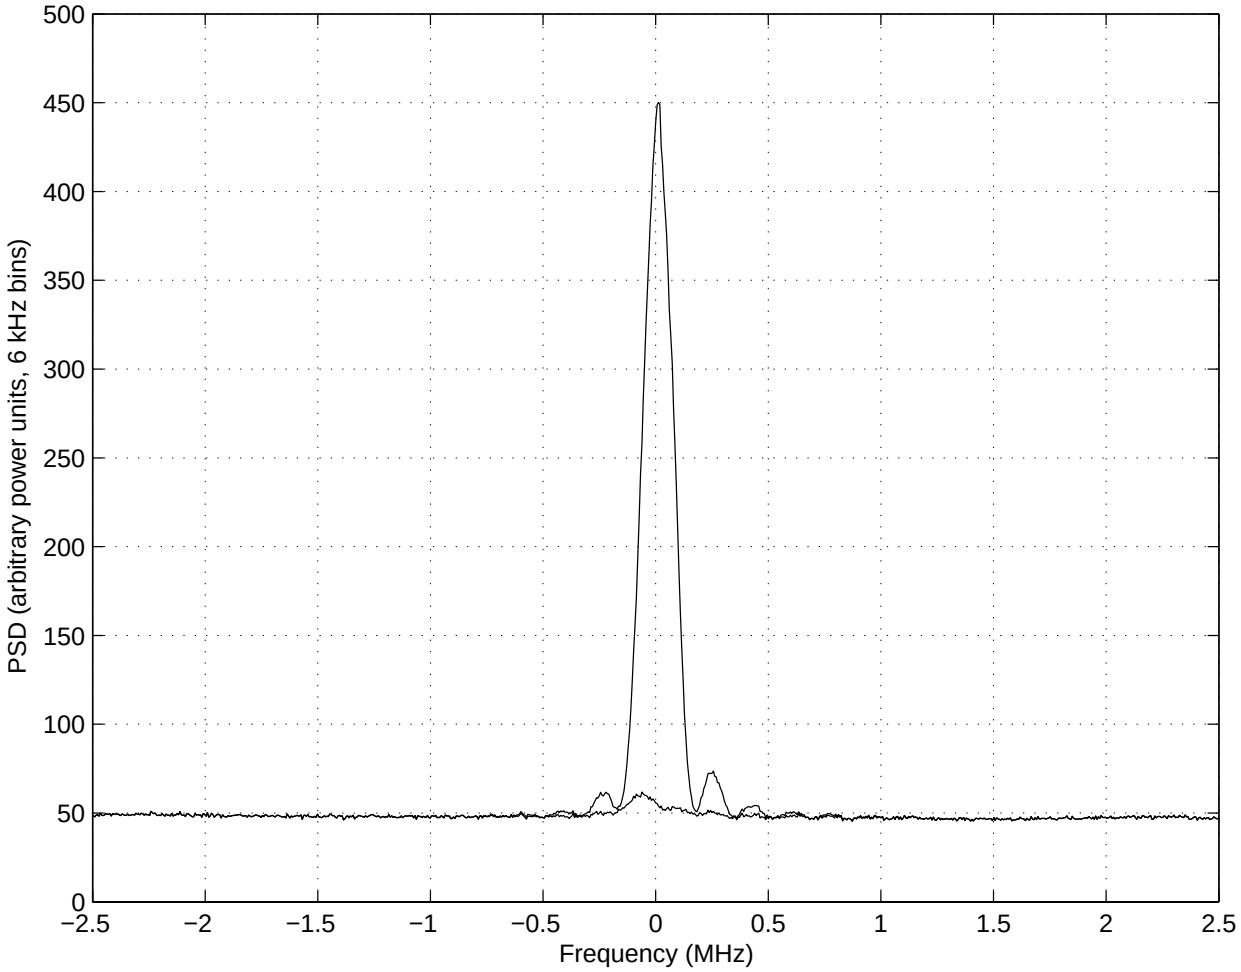

Fig. 8.— Integrated spectra (linear power average). *Top:* Before, *Bottom:* After pulse canceling with $\beta = 10$.

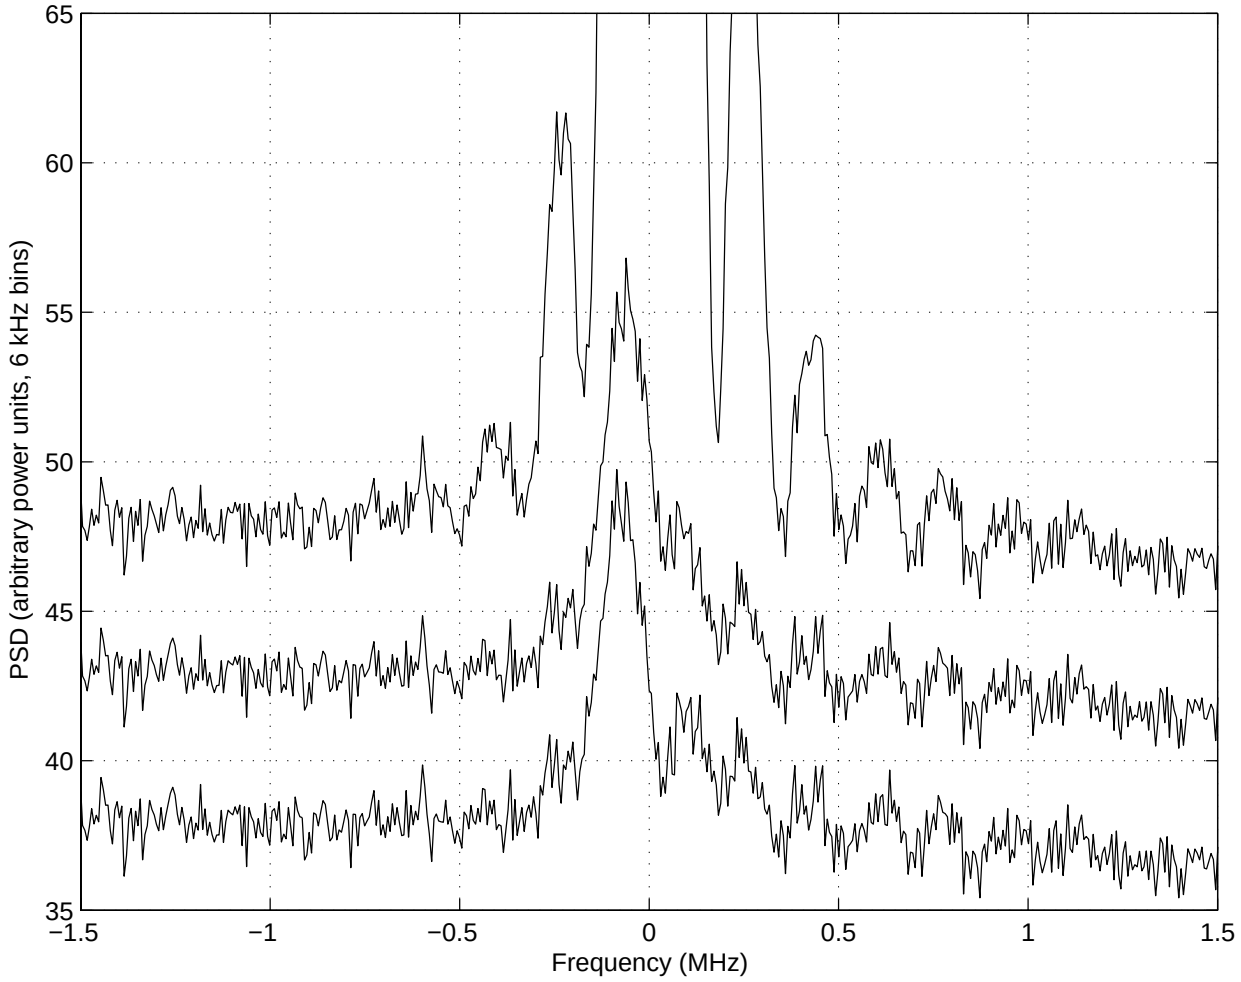

Fig. 9.— Integrated spectra (linear power average). *Top:* Before, *Middle:* After pulse canceling with $\beta = 10$ (shifted down 5 units for clarity) and *Bottom:* $\beta = 5$ (shifted down 10 units for clarity). Note that the difference between the $\beta = 10$ and $\beta = 5$ results is barely noticeable.

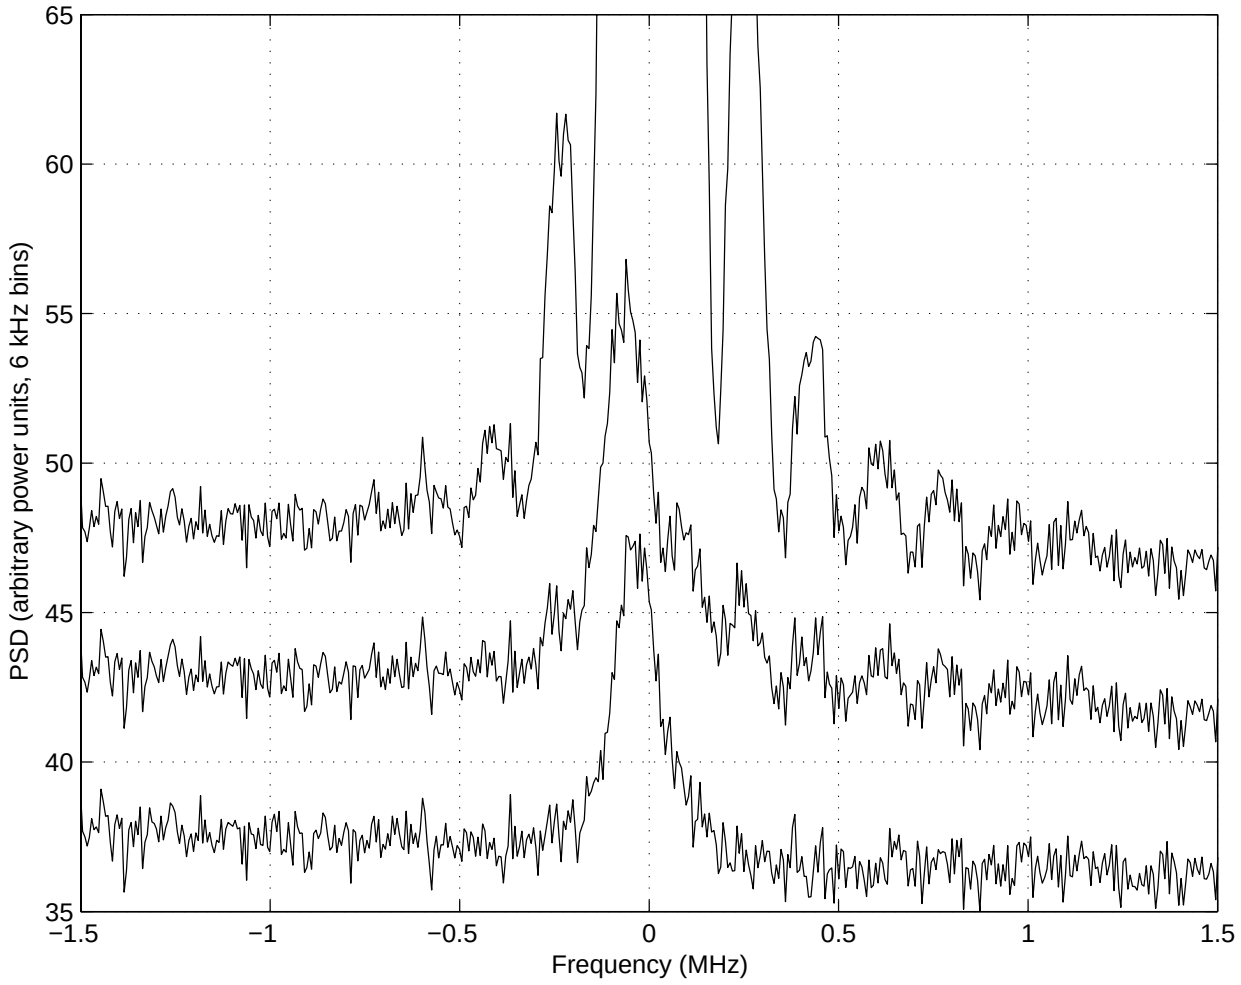

Fig. 10.— Integrated spectra (linear power average) for various pulse removal techniques with $\beta = 10$. From top to bottom: (1) Before, (2) Pulse canceling (shifted down 5 units for clarity), (3) Pulse blanking (shifted down 10 units for clarity).

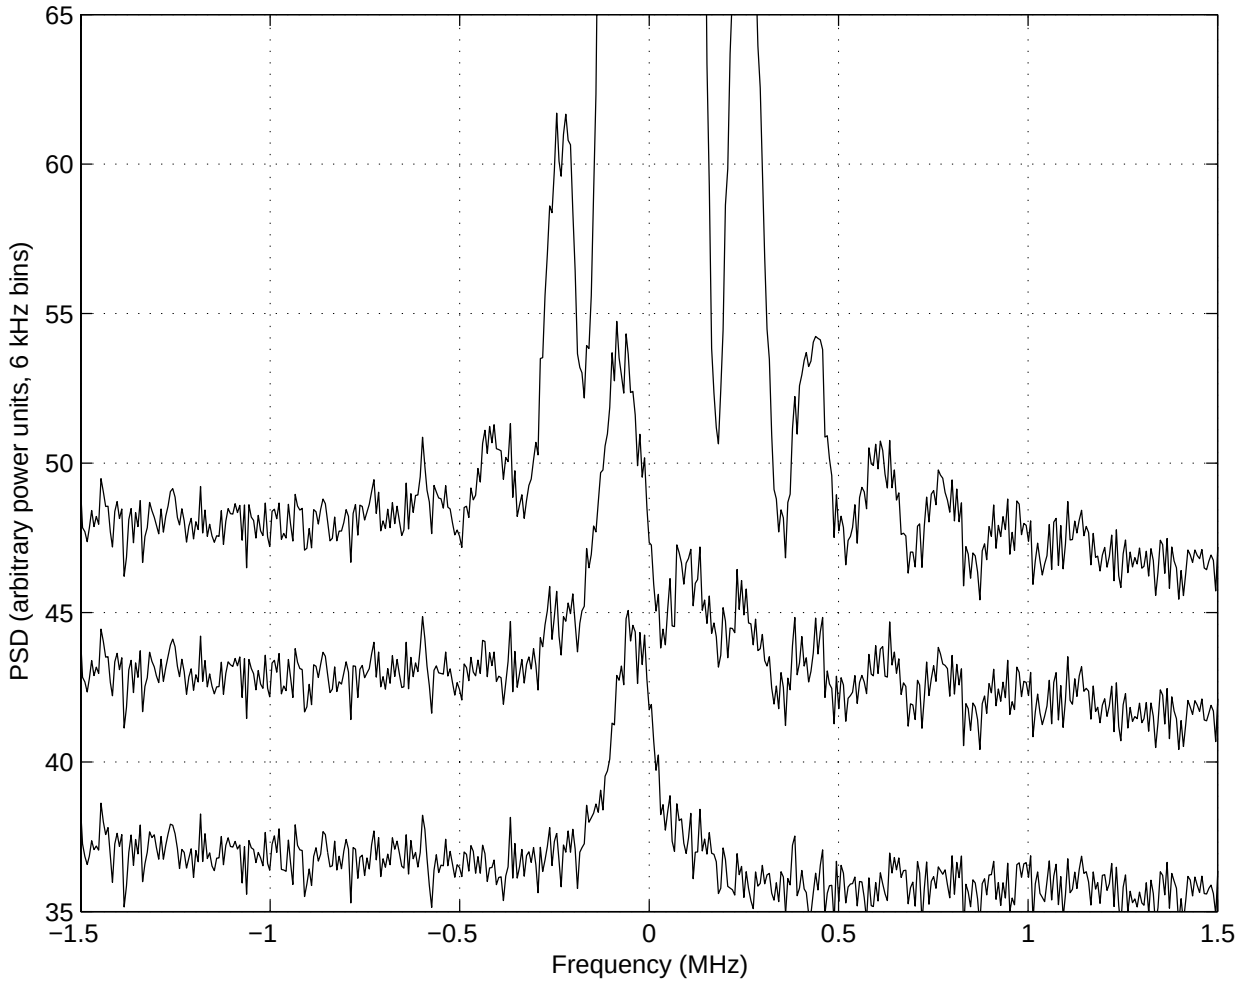

Fig. 11.— Integrated spectra (linear power average) for various pulse removal techniques with $\beta = 5$. From top to bottom: (1) Before, (2) Pulse canceling (shifted down 5 units for clarity), (3) Pulse blanking (shifted down 10 units for clarity).

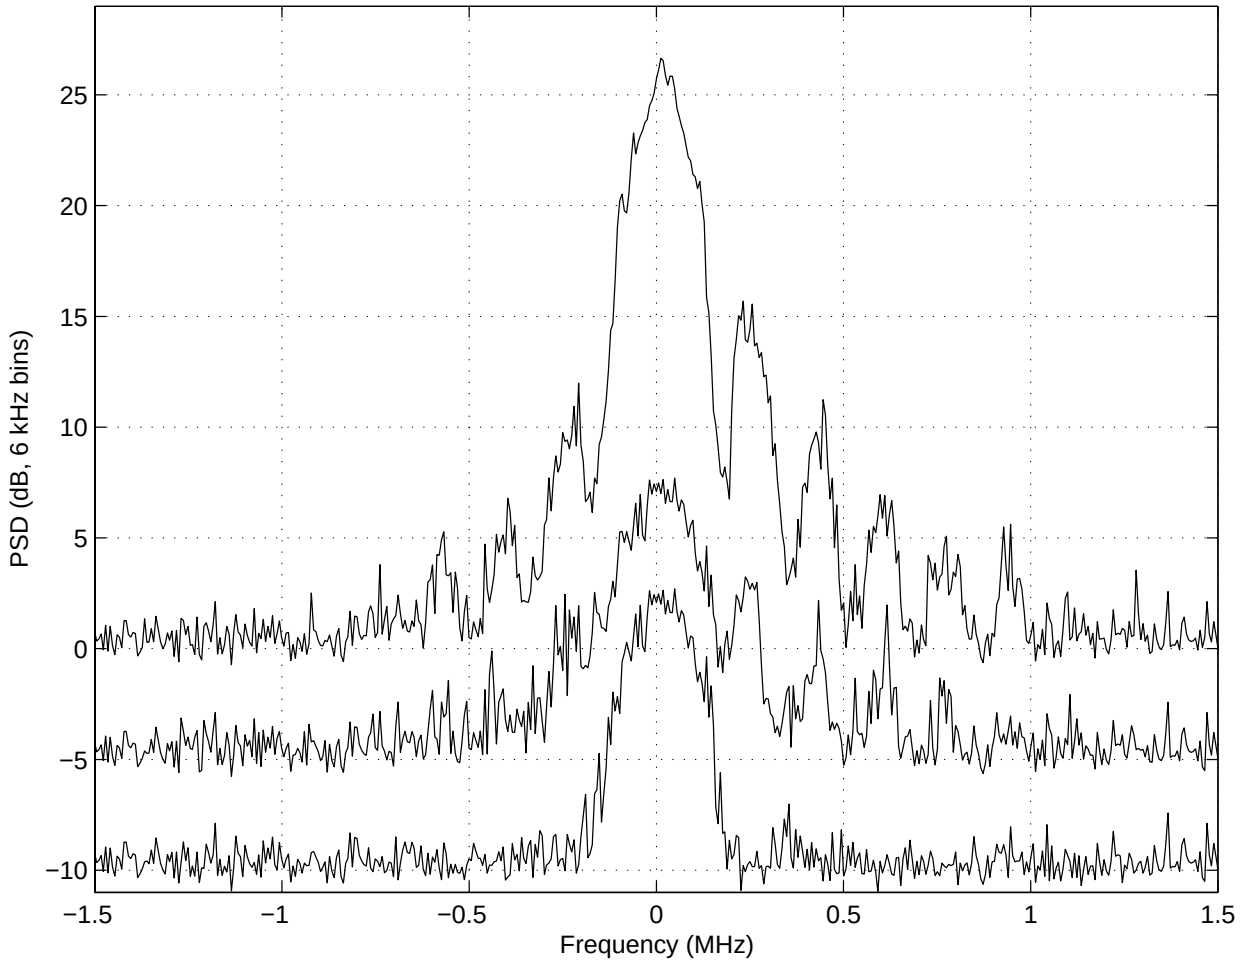

Fig. 12.— "Max hold" spectra for various pulse removal techniques with $\beta = 10$. *Note: The vertical axis for this plot is in dB.* A max hold spectrum is obtained by taking the maximum, as opposed to the mean, over the individual spectra. *Top to Bottom:* Before, Pulse canceling (shifted down 5 dB for clarity), Pulse blanking (shifted down 10 dB for clarity). The result for $\beta = 5$ is not significantly different.

Despite the disappointing performance of the detector, it is interesting to find that the suppression performance of pulse canceling is competitive with that of pulse blanking. This raises question of the actual suppression achieved by pulse canceling on a pulse-by-pulse basis. The performance of sinusoidal waveform cancelers was considered in Ellingson (2002a) and Ellingson (2002b). It was found that the suppression was upper-bounded at $\mathrm{SNR} \cdot L$, where L is the number of samples used to estimate the waveform parameters assuming noise bandwidth equal to the Nyquist bandwidth. In this case, L is about 1/20 of the actual number of samples used ($\sim 6\ \mu\mathrm{s} \cdot 100\ \mathrm{MSPS} = 600$), since the noise bandwidth at the output of the 5 MHz bandlimiting filter is 1/20 of the complex-valued sample rate (100 MSPS). This gives $L \sim 30$, so the predetection SNR at which one should theoretically achieve the observed suppression of ~ 16 dB is about 2 dB. Since the suppression degrades linearly with the SNR, we expect pulse canceling to become completely ineffective for SNR ≤ -14 dB even if every pulse is detected. However, referring to Figure 7, note that such pulses are below the sensitivity of the current detector. Thus, for this particular detector and dataset, it appears that (1) The mitigation performance is limited by the detector, as opposed to the canceler; and (2) Even if the detector was improved, the mitigation performance would then be limited by the canceler, as indicated above.

It is important to note, however, that the above analysis assumes that the likelihood and extent of intra-pulse multipath distortions are independent of pulse strength. If, on the other hand, weaker pulses are more likely to be distorted by multipath than stronger pulses, then it is less certain whether it is detection or canceling that is limiting performance. Because weaker pulses have lower SNR, a compelling quantitative comparison of the "integrity" of weak vs. strong pulses is difficult. However, a casual review of the data does not leave the impression that weak pulses show greater distortion than strong pulses.

5. Sensitivity to Astronomical Transients

A potential application of GBAR mitigation techniques is to improve the productivity of searches for astronomical transients, such as giant pulses, at L-band. To be useful in pulsar and other transient applications, GBAR mitigation techniques should not mistakenly detect astronomical transients as radar pulses. For pulsar surveys, the impact of doing so depends on the bandwidth of observation. For very wideband (e.g., 100 MHz) pulsar observations, deleting the portion of every pulse appearing within a single 1 MHz channel – whether from a radar or pulsar – is perhaps not a problem. However, for observations using less bandwidth, or in the presence of multiple radars, this problem may result in a significant loss of SNR, and perhaps also might introduce undesirable artifacts in the dedispersed data. Thus, it is

useful to note under what conditions astronomical pulses might be mistaken for radar pulses using the detector described earlier.

Consider the case of an impulse which is received after being subject to dispersion due to the interstellar medium. The dispersion broadens the impulse into a tone which drifts in frequency such that

$$\Delta t \approx (8.3 \times 10^{15})\ (\mathrm{DM})\ (\Delta\nu)\ \nu^{-3} \tag{1}$$

(Lyne & Graham-Smith, 1998), where Δt is the time required for the frequency of the dispersed pulse to drift across the bandwidth $\Delta\nu$ at center frequency ν, and DM is in pc cm^{-3}. Figure 13 shows shows that for DM > 25, the matched filter output is significantly longer than that associated with a radar pulse. Therefore, detected astronomical pulses can be discriminated from radar pulses simply by observing the length of time over which the matched filter output exceeds the detection threshold.

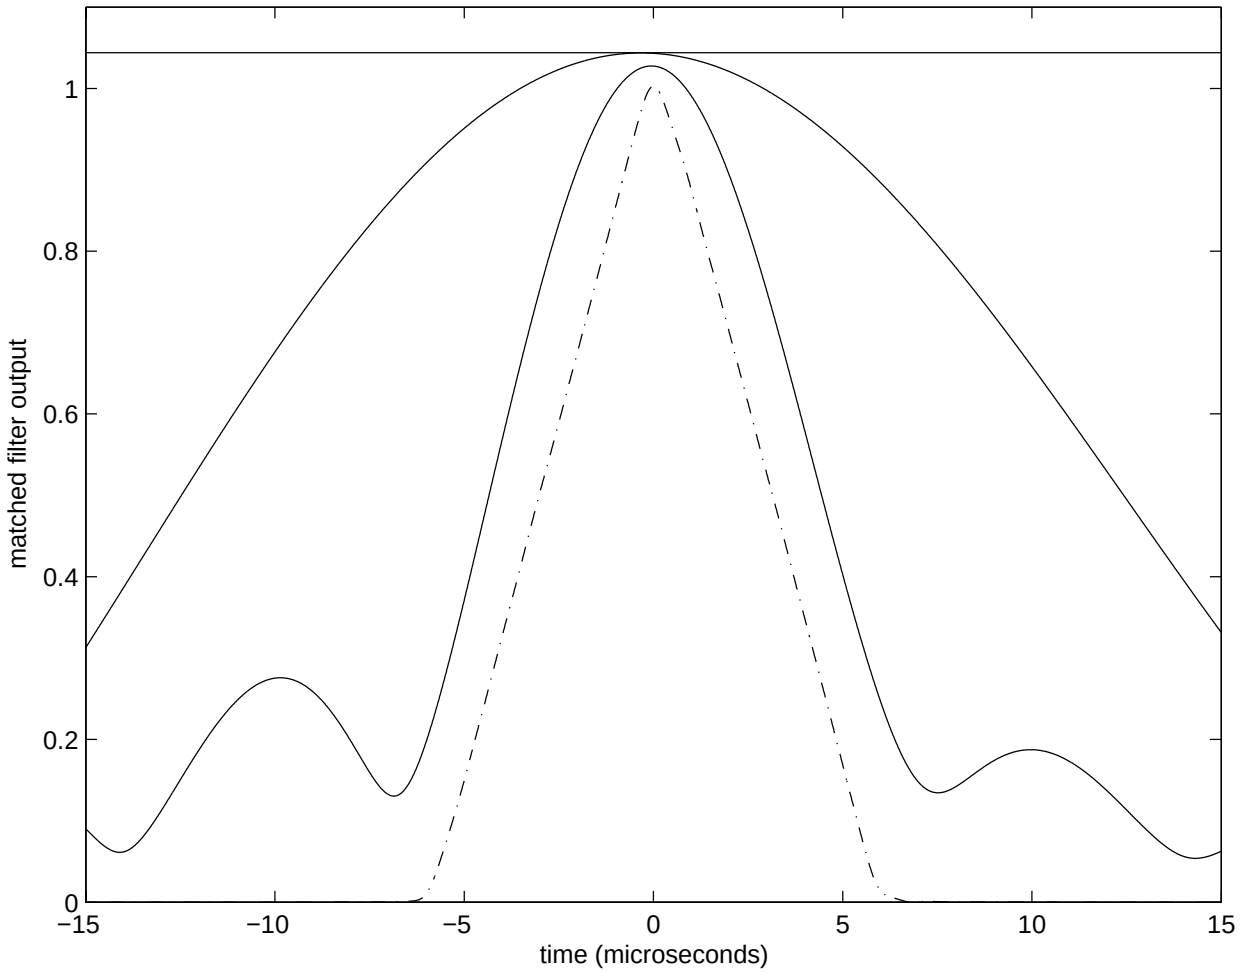

Fig. 13.— Matched filter output $y(t)$ for impulses with DM (pc cm^{-3}) from (*top to bottom*) ∞, 71, and 25. The dashed line is $y(t)$ for the radar pulse. The instantaneous SNR for each case is the same.

Below 25 pc cm^{-3}, the resulting matched filter output is about the same length as that expected for a radar pulse, so this method cannot reliably discriminate. This is shown in Figure 14. However, note that it is possible to determine a conservative lower bound on the strength of a pulse that can be falsely detected as a radar pulse, given the DM. Since the matched filter output for a DM=15 pulse is nearly identical to that of the radar pulse (and therefore also the worse case), the associated detection "performance" is given by Figure 7. At $\beta = 5$, for example, the probability of detection becomes significant for an instantaneous SNR of about -15 dB. For Arecibo at L-Band, this corresponds to about 79 mJy assuming an effective aperture of 10 K/Jy and $T_{sys} = 25$ K. This is can be improved to about 790 mJy using $\beta = 10$. Also, the situation improves considerably for DM < 15, as shown in Figure 14, since the matched filter output magnitude decreases with decreasing DM.

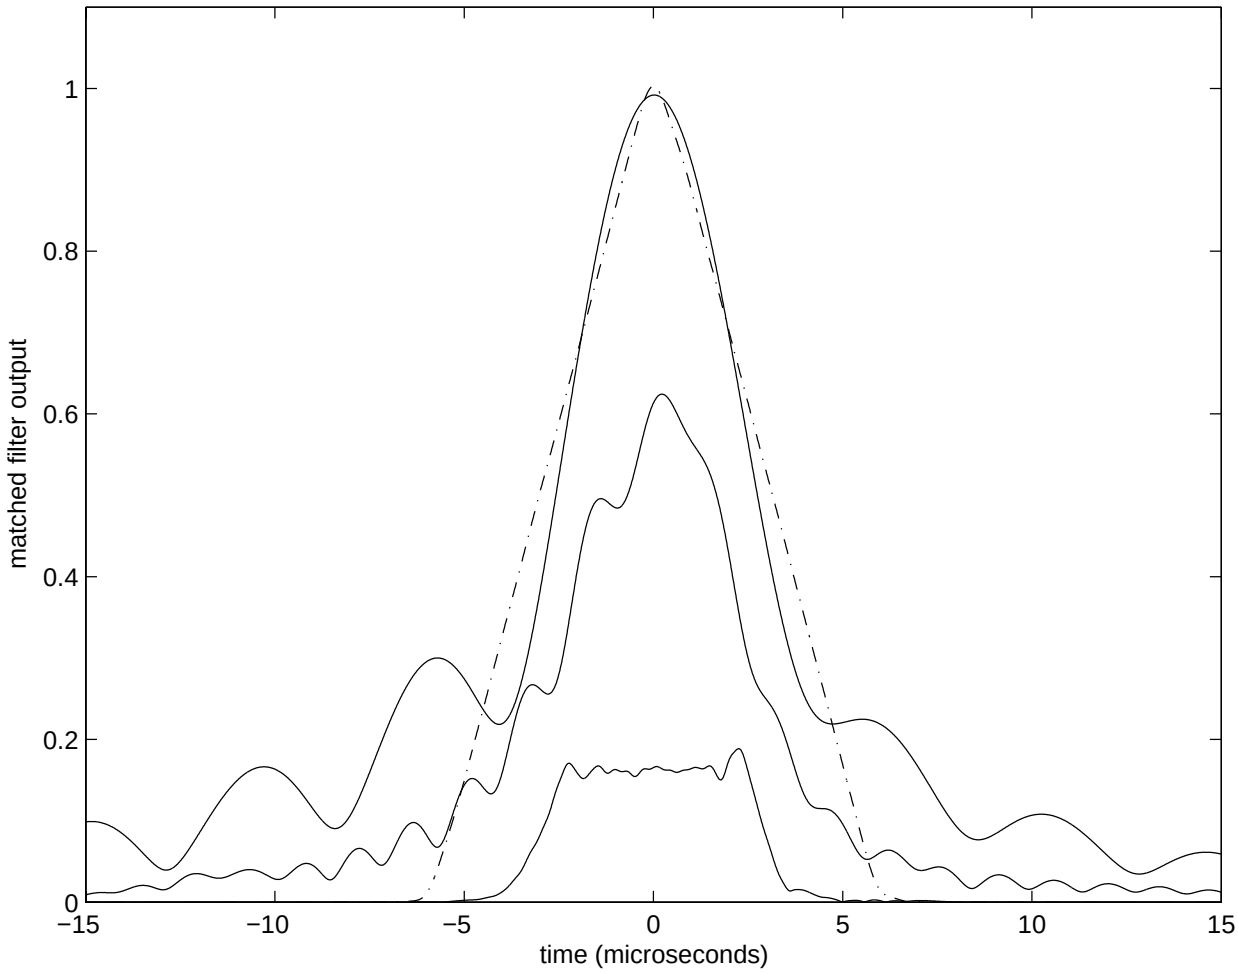

Fig. 14.— Matched filter output $y(t)$ for impulses with DM (pc cm^{-3}) from (*top to bottom*) 15, 5, and 0.5. The dashed line is $y(t)$ for the radar pulse. The instantaneous SNR for each case is the same.

It should be noted that the worst case DM of ~ 15 corresponds to the case in which the apparent frequency drift covers the ~ 150 kHz main lobe bandwidth of the matched filter in ~ 6 μs, equal to the length of the matched filter. This criterion can presumably used to estimate the worst case DM for other matched filters as well.

In summary, the risk of false detection of naturally-occurring pulses as GBAR pulses is significant only for pulses > 79 mJy (assuming $\beta = 5$) at DM < 25 or so, decreases with decreasing DM, and is negligible for DM > 25 assuming the duration of the detection in $y(t)$ is monitored.

6. Enhancements and Applications to Other GBARs

Although the simple techniques described above achieved significant (16 dB, in this case) suppression of interference from a GBAR, this may not be adequate for some applications – especially spectroscopy. However, various improvements are possible. Perhaps the greatest improvement would be achieved by developing a detector that can sense overlapping pulses, combined with modifications to the pulse blanking and canceling schemes to suppress overlapping pulses. Another strategy is to use "multichannel" detectors: Since most radio telescopes employ multiple polarizations – and increasingly, array feeds – the use of these concurrent observations can be exploited to improve the sensitivity of the detector. The relevant theory is addressed in Leshem, van der Veen, & Boonstra (2000).

The focus of this paper was on the 1350 MHz GBAR received at Arecibo. Although the waveform used by this radar is very common, it is not the only GBAR waveform received at Arecibo or at other locations. However, most other waveforms observed are very similar, or only slightly more complex. (The reason for this is simply that it is relatively difficult to generate complex waveforms at the high power desired for GBAR applications.) The most common variation to the simple CW pulse waveform is the so-called "Linear FM" (LFM) waveform. LFM is identical to pulsed CW, with the exception that the frequency is swept (or "chirped") as linear function of time as the pulse is transmitted. As an example, Figure 15 shows a time-frequency plot of another GBAR commonly observed at Arecibo, the frequency-hopping "spread spectrum" radar located at Punta Salinas, PR. To detect this pulse, one can simply "dechirp" the signal using the known (fixed) frequency sweep rate – thereby converting it to a pulsed CW waveform – before using a pulse-matched detector of the appropriate length. Pulse canceling is a simple matter of computing the same coefficient α, except now for a model of the dechirped LFM pulse. Other GBAR pulse waveforms, despite their apparent complexity, might similarly yield to simple processing techniques.

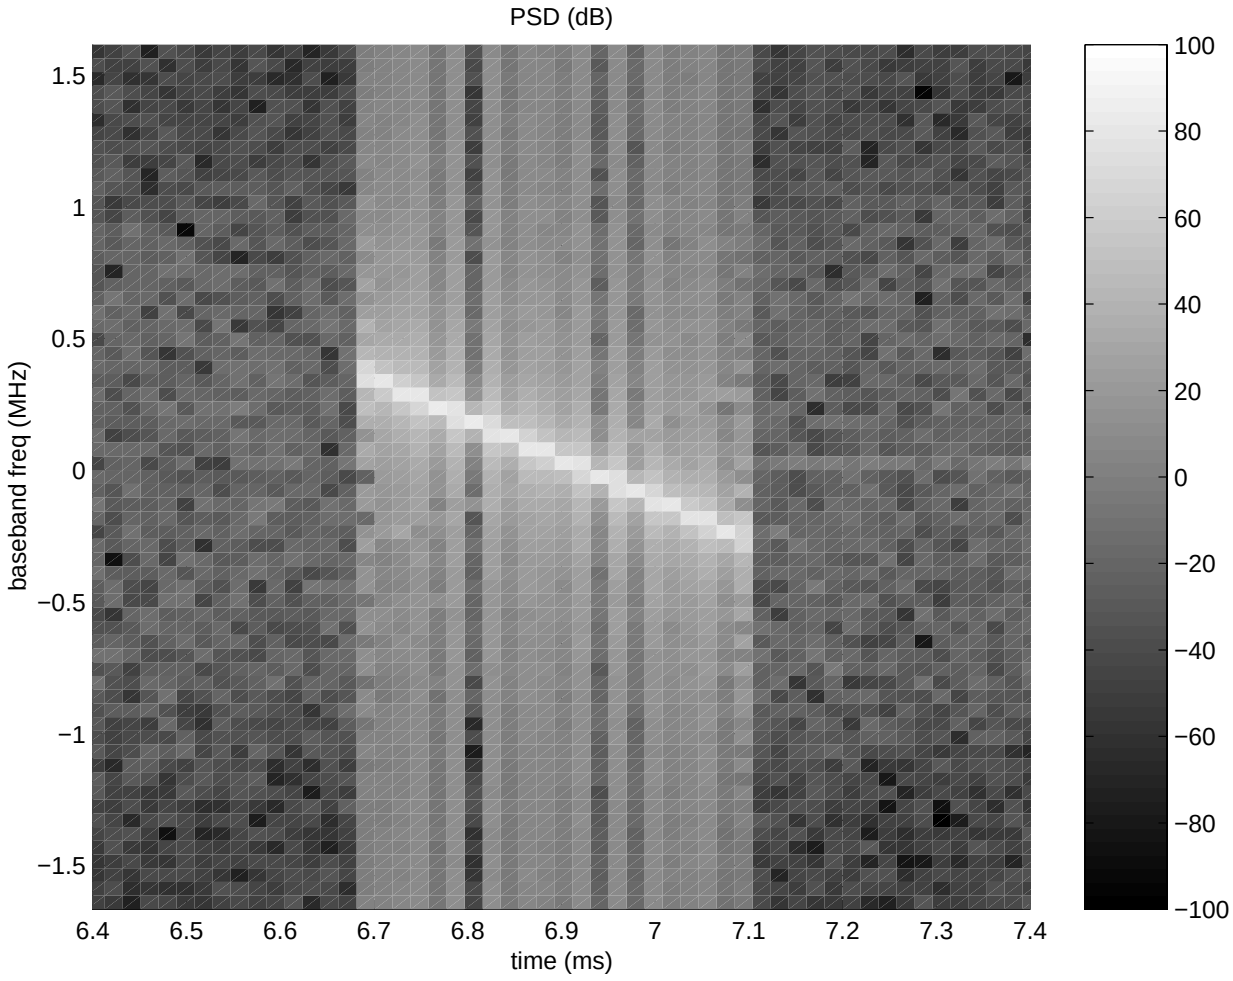

Fig. 15.— Time-frequency plot of power spectral density (PSD) for one of two types of pulse waveforms received from the FPS-117 radar located at Punta Salinas, PR (Ellingson 2002c). The other pulse waveform differs only in length.

A. Impulse Response of the Propagation Channel

In the development of mitigation techniques, it is useful to know how radar pulses are dispersed as they propagate from the transmitter to the receiver. An appropriate characterization is the impulse response $h(t)$, which can be described in terms of the measured signal $x(t)$, the transmitted pulse $p(t)$ (as given in Section 2.3), and the measurement noise $n(t)$ as follows:

$$x(t) = h(t) * p(t) + n(t) \tag{A1}$$

The deconvolution of $p(t)$ and $h(t)$ tends to be ill-conditioned; however, it is possible to obtain a useful *modified* impulse response $g(t)$, derived as follows. In the frequency domain, we have:

$$X(\nu) = H(\nu)P(\nu) + N(\nu) \tag{A2}$$

where the upper case quantities denote the Fourier transforms of the associated lower case quantities. Multiplying both sides by $P^*(\nu)$:

$$P^*(\nu)X(\nu) = \|P(\nu)\|^2 H(\nu) + P^*(\nu)N(\nu) \tag{A3}$$

Note that if the pulse SNR is large, the first (signal) term on the right-hand side dominates over the second (noise) term. Neglecting the noise term and transforming back into the time domain we have

$$p^*(t) * x(t) \sim F^{-1}\{\|P(\nu)\|^2\} * h(t) \tag{A4}$$

Since $p(t)$ is approximately a rectangular pulse, $F^{-1}\{\|P(\nu)\|^2\}$ is approximately a triangular pulse, and each side of Equation A4 is the same as $h(t)$ but "smeared" in the time domain due to the convolution with a triangular pulse function. We shall call this modified version of the impulse response $g(t)$. $g(t)$ is a more useful estimate of the impulse response than Equation A1 because it improves the SNR of the result and shows sharp, clearly defined peaks.

For each block in the data set described in Section 2.1, we calculated $g(t)$. Since the inter-pulse spacing is between 2.6 ms and 3.3 ms, there is at most one transmitted pulse per block. All detections of 10σ or greater at the output of the $p(t)$-matched detector were identified. For each of these, the peak of $g(t)$ was aligned with an arbitrary time reference $\tau = 0$. The "excess delay" τ refers to the time relative to the arrival of the detected peak, facilitating comparisons between different pulses. Figure 16 shows the results. In addition to intermittent strong discrete multipaths, note that the average curve shows a continuum of multipath that can be observed out to about 100 μs, where it is about -36 dB relative to the main peak and merges into the noise. Note also that individual multipaths are detected as long as 200 μs after the main pulse, and at least one multipath component is detected about

10 μs *before* the triggering pulse. The latter demonstrates that the alignment of received pulse events is subject to bias due to the possibility of multipath (i.e., delayed) pulses which are stronger than direct path pulses.

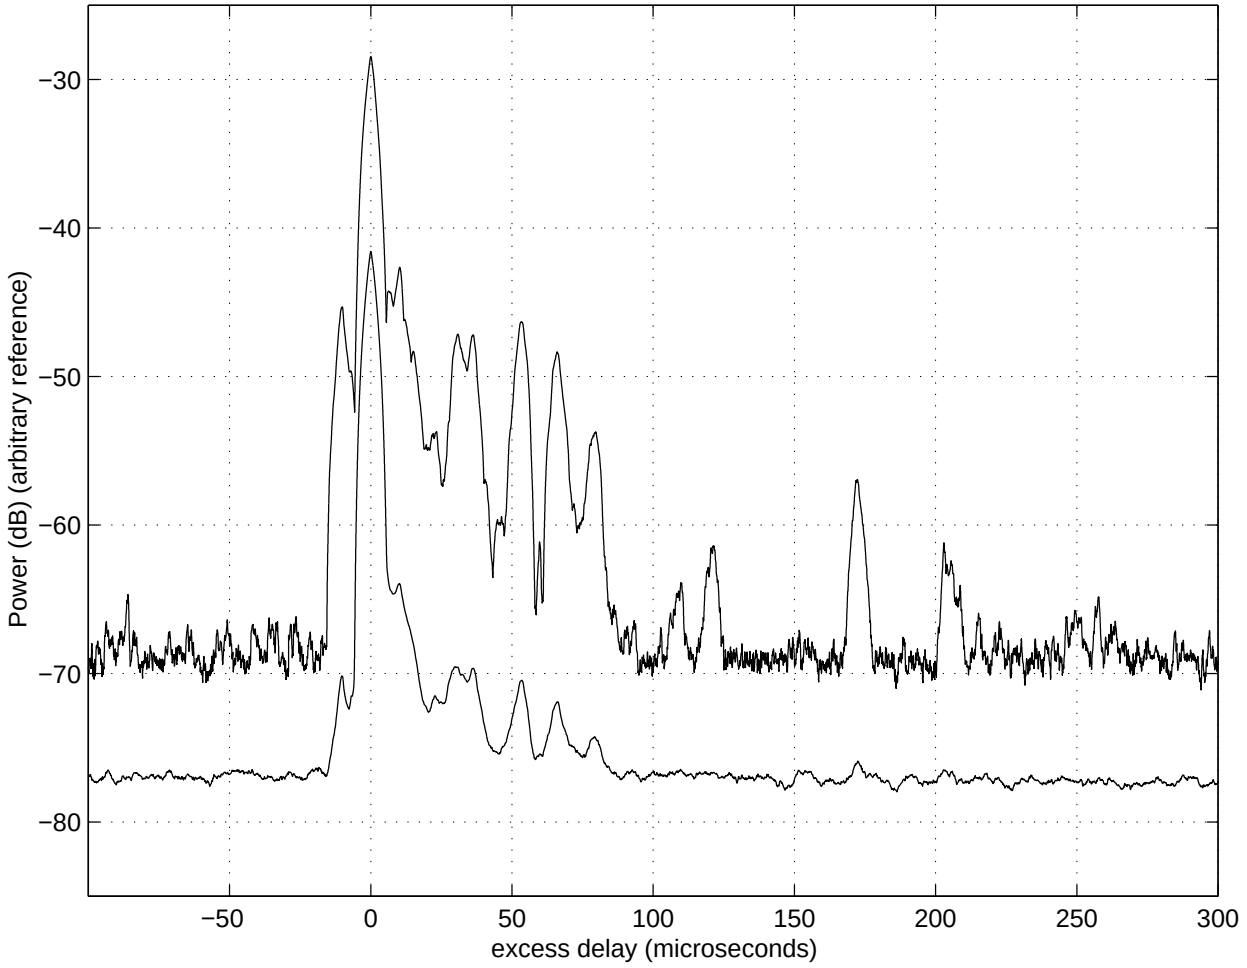

Fig. 16.— Max hold (top) and linear power average (bottom) of $g(\tau)$. The "max hold" measurement is computed as follows: For each value of τ, one uses the maximum value of $g(\tau)$ observed from the available pulses.

Acknowledgments

The authors gratefully acknowledge the assistance and advice of R. Bhat (Haystack), J. Cordes (Cornell), R. Fisher (NRAO), J. Johnson (OSU), P. Perillat (NAIC), and L. Wray (NAIC). This work was supported in part by the National Science Foundation (NSF) under Award No. AST-0138263 via subcontract to Cornell University. The Arecibo Observatory is operated by Cornell University under a Cooperative Agreement with the NSF.

REFERENCES

Cognard, I., Shrauner, J.A., Taylor, J.H., & Thorsett, S.E. 1996, ApJ, 457, L81

Ellingson, S.W. 2002a, in Int'l Geoscience & Remote Sensing Sym. (IEEE: Toronto), 1685.

Ellingson, S.W. 2002b, A Study of Adaptive Canceling for Microwave Radiometry and Spectrometry, Tech. Rep. 743467-1, The Ohio State University ElectroScience Lab.

Ellingson, S.W. 2002c, Characterization of Some L-Band Signals Visible at Arecibo, Tech. Rep. 743467-2, The Ohio State University ElectroScience Lab.

Ellingson, S.W., Bunton, J., & Bell, J.F. 2001, ApJS, 135, 87

Ellingson, S.W. & Hampson, G.A. 2002, RFI and Asynchronous Pulse Blanking at Arecibo, Tech. Rep. 743467-3, The Ohio State University ElectroScience Lab.

Ellingson, S.W., Johnson, J.T. & Hampson, G.A. 2002, Development of an RFI-Robust Receiver for Microwave Radiometry (Year 1), Tech. Rep. 742068-1, The Ohio State University ElectroScience Lab.

Friedman, P. 1996, in 8th IEEE Signal Processing Workshop on Statistical Signal and Array Processing (Los Alamitos: IEEE), 264

Kay, S.M. 1993, Fundamentals of Statistical Signal Processing: Estimation Theory (Upper Saddle River, NJ: Prentice-Hall PTR)

Kay, S.M. 1998, Fundamentals of Statistical Signal Processing: Detection Theory (Upper Saddle River, NJ: Prentice-Hall PTR)

Leshem, A., van der Veen, A.-J., & Boonstra, A.-J. 2000, ApJS, 131, 355

Lyne, A. & Graham-Smith, F. 1998, Pulsar Astronomy (Cambridge, UK: Cambridge Univ. Press)

Miller, T., Potter, L., & McCorkle, J. 1997, IEEE Trans. Aerospace & Electronic Systems, 33, 1142

Staelin, D.H. & Reifenstein, E.C. 1968, Science, 162, 1481

Weber, R., Faye, C., Biraud, F., & Dansou, J. 1997, A&AS, 126, 161

This preprint was prepared with the AAS LaTeX macros v5.0.