

PROGRAMMABLE REAL-TIME CANCELLATION OF *GLONASS* INTERFERENCE WITH THE GREEN BANK TELESCOPE

A. J. POULSEN,¹ B. D. JEFFS, AND K. F. WARNICK

Electrical and Computer Engineering Department, Brigham Young University, Provo, UT 84602;
poulsen@mit.edu, bjeffs@ee.byu.edu, warnick@ee.byu.edu

AND

J. R. FISHER

National Radio Astronomy Observatory, Green Bank, WV 24944; rfisher@nrao.edu

Received 2004 December 16; accepted 2005 July 20

ABSTRACT

A programmable, real-time, digital signal processing (DSP), least-mean-squares (LMS) interference canceler was developed and demonstrated as a successful method of excising the *GLONASS* satellite *L*-band downlink signals from the 100 m Green Bank Telescope (GBT). This canceler has the potential for aiding the observation of redshifted OH lines near 1600 MHz in the presence of *GLONASS* signals. Our canceler approach is similar to experiments reported by Barnbaum and Bradley, Baan et al., and Kesteven et al. for mitigating ground-based radio frequency interference (RFI), but we add the new features of a tracking dish antenna to follow the satellite for improved cancellation performance, complex adaptive filter coefficients that reduce the required filter length, and a programmable delay line for bulk time delay correction. An analysis of LMS tracking performance for moving satellites is presented, and effective use of a reference antenna distantly located from the telescope is demonstrated. As is expected from the known properties of LMS cancelers, larger filter orders are shown to yield more reliable cancellation and are less sensitive to data time misalignment. *GLONASS* interference seen by the GBT was suppressed below the noise floor without corrupting a signal that simulated an astronomical source, even with rapid orbital transits and long integration times.

Key words: instrumentation: miscellaneous — methods: numerical — radio continuum: general — radio lines: general — techniques: spectroscopic

Online material: color figures, mpeg animations

1. INTRODUCTION

In recent years, radio astronomy has faced a growing interference problem as radio frequency (RF) bandwidth has become an increasingly scarce commodity. Pertinent frequency band allocation restrictions and designated radio-quiet zones have historically been preserved by the Federal Communications Commission, International Telecommunication Union, and other regulatory agencies. Some of this protection has been undermined as the global community has launched satellite communication systems and the spectral range of interest to astronomy has grown. Many important astronomical sources exist at frequencies outside of the traditionally protected observational bands. In particular, as more distant objects with higher redshifts are observed, no set fixed protection band arrangement is suitable, since sources can be shifted to virtually any frequency. For example, frequency bands of spectral line emissions of neutral hydrogen (1420.4 MHz) and the hydroxyl molecule (1612.2, 1665.4, 1667.4, and 1720.5 MHz) from cosmic sources are protected by international spectrum allocations, but very distant objects are Doppler shifted to much lower frequencies. Some of this radiation is shifted into the lower *L*-band frequency ranges allocated to a variety of radio services, including satellite downlink channels.

This paper concentrates on satellite radio frequency interference (RFI) rather than terrestrial-based emissions. Satellite

sources are particularly problematic because they sweep rapidly across a telescope's field of view, so interference can arrive from any direction, including both the primary antenna deep sidelobes and near the main lobe look direction. Although simple line-of-sight signal paths with little multipath structure are common, signals are relatively strong and the geometry changes rapidly, so an adaptive, rapid-response mitigation technique is required.

The Russian Federation *Global Navigation Satellite System* (*GLONASS*) has been particularly problematic. Its downlink signal consists of two direct sequence digital spread spectrum modulation components with bandwidths of 600 kHz and 6 MHz, respectively. They have equal total transmit powers, so the spectral density is 10 dB higher in the 600 kHz component. Due to an unfortunate initial choice of transmission bands, this system infringes on astronomically important frequencies, including the 1612 MHz hydroxyl (OH) line. More recently, perhaps due to pressure from the radio astronomical community, the *GLONASS* downlink frequency assignments have all been in lower channels, which has helped the situation.² But some spectral sidelobe energy still creeps into critical bands, and moderately redshifted OH lines can be completely obscured. Currently, significant

¹ Current address: Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139.

² As of 2005 September all operational satellites were assigned to frequency channels 1, 2, 3, 5, 6, 10, 11, and 12 (see <http://www.glonass-center.ru/nagu.txt>). The corresponding carrier frequencies are computed as $f = 1602 + n \cdot 5625$ MHz, where n is the channel number. Thus, the current highest carrier frequency is 1608.75 MHz, which is outside of the protected OH band. Historically, channels 14–24 were commonly used, pushing the carriers as high as 1615.500 MHz, which enters the protected band (see http://www.glonass-center.ru/hist_e.html).

TABLE 1
 SUMMARY FOR THE LMS FILTER CONFIGURED AS AN ADAPTIVE CANCELER

Expression	Description
Inputs:	
$d[n]$	Signal received by radio telescope
$s[n]$	Desired astronomical source
$i_p[n]$	Interference at primary channel
$\eta_p[n]$	Noise at primary channel
$x[n]$	Signal received by reference antenna
$i_r[n]$	Interference at reference channel
$\eta_r[n]$	Noise at reference channel
Output:	
$\hat{s}[n]$	Estimate of desired signal
Parameters:	
μ	Filter update parameter
L	Filter order
\mathbf{h}_n	Filter weight vector at time n
Internal Signals:	
$\mathbf{x}_n = [x[n], x[n-1], \dots, x[n-L+1]]^T$	Data vector
$y[n]$	Filtered reference channel
Initialization:	
$\mathbf{h}_0 = \mathbf{0}$	Initial filter tap weights
Computation:	
$y[n] = \mathbf{h}_n^T \mathbf{x}_n$	Filtered reference
$\hat{s}[n] = d[n] - y[n]$	Canceller output update
$\mathbf{h}_{n+1} = \mathbf{h}_n + \mu \hat{s}[n] \mathbf{x}_n^*$	Update filter weights

observation time can be lost due to a lack of effective interference mitigation solutions.

Ellingson et al. (2001) presented a successful mitigation method tailored to the *GLONASS* problem based on a detailed model of the transmission waveform and estimating parameters such as gain and timing from the observed data. The matched model waveform was subtracted for the observation. In contrast, our adaptive canceler approach requires no precise prior knowledge of the waveform and can work with a wide variety of signal types.

The groundbreaking real-time RFI canceler reported by Barnbaum & Bradley (1998) demonstrated concept feasibility with a dedicated digital hardware implementation of an eight-tap, real, least-mean-squares (LMS), finite impulse response (FIR) filter that could support 1 MHz bandwidth operation. Compared to that experiment, our canceler offers the following enhancements: Our platform, like those of Baan et al. (2004a, 2004b), Kesteven et al. (2005), and Kesteven (2005), is programmable rather than fixed dedicated hardware. In addition, we implement a complex valued LMS FIR filter and allow a programmable trade-off between filter length (up to 42 taps) and bandwidth (up to 4 MHz). Complex valued filters often require significantly fewer taps than real filters with comparable phase and frequency response. Our digital receiver front end uses no digital signal processing (DSP) resources, supports sample rates up to 65 megasamples s^{-1} , mixes the signal to a decimated complex baseband for DSP input, and allows the tuning of an arbitrary width receiver bandpass window anywhere from 0 to 25 MHz. Finally, we have used active tracking of the interfering satellite with a high gain reference antenna and demonstrated cancellation of real satellite interference seen in the L band with an operational telescope, the Green Bank Telescope (GBT).

The tracking dish antenna provides a high interference-to-noise ratio (INR) copy of the detrimental signal to the adaptive canceler and thus significantly improves the accuracy of signal subtraction from the primary channel. A suitable tracking dish system can be purchased at very modest cost, and if necessary

or convenient can be located quite distant from the primary antenna (1.4 km in our trials).

2. REAL-TIME MITIGATION WITH AN LMS ADAPTIVE FILTER

The adaptive algorithm used in this study is based on the classical LMS technique, first introduced in Widrow & Hoff (1988) and Widrow (1970). As pointed out by Barnbaum & Bradley (1998), the LMS algorithm is computationally simple, which is very important when the bandwidth that can be processed is limited by computational resources. Although faster converging methods such as the recursive least-squares (RLS) algorithm exist (Haykin 1996), such measures are not called for here. For a filter length of L taps, the LMS computational load is $2L + 1$ multiplications and $2L$ additions per time sample. By comparison, an efficient RLS code requires $2L^2 + 6L$ multiplications, $\frac{1}{2}(3L^2 + 5L)$ additions, and one division per sample, which would drastically limit the supportable sample rate and thus the usable signal bandwidth. The LMS convergence rate proved acceptable with the 1.4 km antenna separation that we used while observing *GLONASS* interference. Faster moving low Earth orbit satellites such as *IRIDIUM* would be a greater challenge, but LMS would work with a smaller reference antenna separation. Section 2.2 presents an analysis of LMS tracking error as a function of antenna separation and the apparent angular velocity of a satellite.

2.1. Cancellation Filter Implementation

A summary of the LMS adaptive canceler equations and important signal descriptions is presented in Table 1 and the corresponding Figure 1. The radio telescope (primary antenna) receives both the desired astronomical signal $s[n]$ seen in the high-gain main lobe and the interference $i_p[n]$ seen in the side-lobes. The net signal at the telescope feed is

$$d[n] = s[n] + i_p[n] + \eta_p[n]. \quad (1)$$

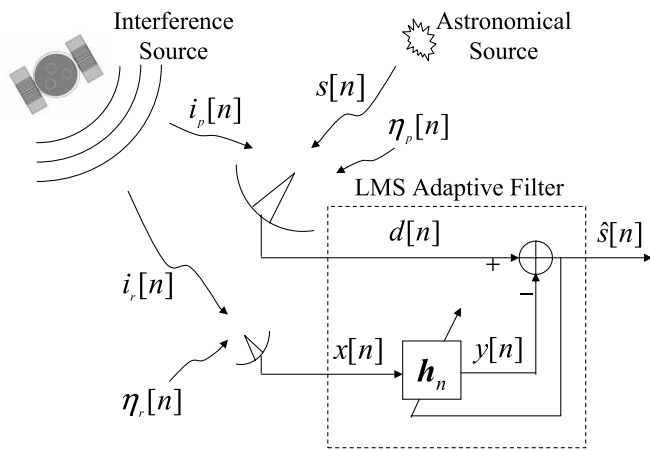


FIG. 1.—Signal flow block diagram for the LMS adaptive canceler used for *GLONASS* RFI mitigation.

Since the reference antenna is aimed at the interferer, $x[n]$ contains virtually no astronomical signal and provides a high-INR copy of $i_r[n]$

$$x[n] = i_r[n] + \eta_r[n]. \quad (2)$$

Subscripts p and r indicate the primary and reference antenna signals, respectively. The noise components $\eta_p[n]$ and $\eta_r[n]$ are mutually uncorrelated.

The L tap filter weight vector, \mathbf{h}_n , adaptively manipulates the reference signal $x[n]$ to closely match the interference incident to the primary channel. The filter output contains both the approximation $\hat{i}_p[n]$ and the noise component $\eta_h[n]$, which is $\eta_r[n]$ colored by \mathbf{h}_n :

$$y[n] = \mathbf{h}_n^T \mathbf{x}_n = \hat{i}_p[n] + \eta_h[n]. \quad (3)$$

At convergence, \mathbf{h}_n approximates the transfer function with respect to $i_p[n]$ from the reference to the primary antenna. Roughly speaking, the filter attempts to make $\hat{i}_p[n]$ match $i_p[n]$ while at the same time minimizing the power in $\eta_h[n]$. Thus, the subtraction block in Figure 1 removes $i_p[n]$ from $d[n]$, leaving $s[n]$ virtually unaffected.

The LMS algorithm has been shown to converge approximately to the optimal minimum mean squared error FIR filter, i.e., the Wiener filter solution (Hayes 1996; Haykin 1996):

$$\mathbf{h}_n \approx \mathbf{h}_{\text{mmse},n} = \mathbf{R}_{x_n}^{-1} \mathbf{r}_{x_d_n}, \quad (4)$$

where

$$\mathbf{R}_{x_n} = E\{\mathbf{x}_n \mathbf{x}_n^H\}, \quad \mathbf{r}_{x_d_n} = E\{\mathbf{x}_n d[n]^*\}, \quad (5)$$

$E\{\dots\}$ denotes the statistical expectation, the superscript “ H ” indicates the complex conjugate transpose, and the asterisk indicates the complex conjugate. The subscript n used throughout indicates that the signal statistics may not be stationary and that a different solution may exist for each time sample. Even this optimal result is only effective as an interference canceler if $i_r[n]$ and $i_p[n]$ are highly correlated over a time lag of not more than L samples. For example, if the two signals are severely misaligned in time, cancellation may not occur (see § 3.1).

2.2. Tracking Performance and Selecting the Update Parameter μ

In this section we analyze the ability of the LMS canceler to adapt to the changing signal geometry due to satellite orbital mo-

tion. A closed-form expression is derived to bound canceler tracking error as a function of antenna geometry and satellite angular velocity. If the filter convergence rate is too slow, then cancellation depth is reduced. The two most significant factors for signal nonstationarity are interferometric fringe (phase) rotation, with a rate proportional to the primary-to-reference baseline, and the fine structure of the telescope beam pattern sidelobes traversed during a satellite overhead pass. Given our long 1.4 km baseline, fringe rotation produced the more rapid channel variations. Assuming the worst-case *GLONASS* geometry leads to approximately 130 full 2π rotations for a 1° change in direction of arrival. A major question addressed by our experiments is whether a practical LMS implementation can track satellite motion without degrading canceler performance.

The Appendix presents an analysis of LMS canceler performance while tracking signal variations due to satellite orbital motion. These results can be used to compute the optimal update parameter μ . The objective is to minimize the misadjustment error, M , given by equation (A12), which is repeated here:

$$M \approx \frac{\mu}{2} L \sigma_{i_r}^2 + \frac{1}{\mu} \frac{\sigma_{i_p}^2}{\sigma_{\eta_p}^2} \left(1 - \cos \frac{2\pi f_0 b \Delta\theta}{c} \right), \quad (6)$$

where $\Delta\theta$ is the apparent satellite angular motion per time sample, b is the baseline length between antennas, f_0 is the signal frequency, c is the speed of light, and $\sigma_{i_r}^2$, $\sigma_{i_p}^2$, and $\sigma_{\eta_p}^2$ are signal power levels for interference in the reference and in the primary and for noise in the primary, respectively. The first term represents the filter error (relative to the optimal Wiener solution), which is proportional to μ , while the second term represents the filter adaptation error due to the changing relative delay for the interfering signal as seen in the primary and reference channels and is inversely proportional to μ . A rapidly changing relative delay argues for a larger, more responsive value for μ . Hence, there is an optimum value for the update parameter determined by these competing requirements.

Figure 2 illustrates equation (A12) for some realistic values encountered in our experiments. A typical *GLONASS* motion of 0.6 minute^{-1} was assumed with a sample rate of 1.0 MHz, yielding




FIG. 2.—LMS filter misadjustment error as a function of the adaptation parameter μ . This plot is an evaluation of eq. (A12) using realistic values encountered in field experiments. The increase in misadjustment for smaller values of μ is due to tracking error.

$\Delta_\theta = 1.75 \times 10^{-10}$, $b = 1.4 \times 10^3$ m, $f_0 = 1.608 \times 10^9$ Hz, and $L = 12$. The power terms depend heavily on gain settings in the signal path, so for this plot we arbitrarily set $\sigma_{i_r}^2 = 1.0$ and then use realistic power ratios with $\sigma_{i_p}^2 = 0.01\sigma_{i_r}^2$ and $\sigma_{\eta_p}^2 = 0.1\sigma_{i_p}^2$. The optimal μ is seen to be near 1×10^{-6} , and low misadjustment error is achieved, indicating successful tracking.

In practice, we were also able to “hand tune” μ in the original LMS vector update equation to achieve acceptable convergence and cancellation performance. Of the hundreds of tests run at the GBT, very few choices for μ led to no RFI attenuation or caused failure of filter convergence. We found the LMS adaptive filter to be robust and effective in a wide range of signal scenarios.

The trade-off between the size of the baseline and system stationarity raises the question of how close to place the primary and reference antennas. The GBT is a huge physical structure. If the 3.6 m reference antenna were placed in close proximity to the GBT, a large portion of the sky would be permanently shielded by the telescope. The location was also dictated by convenience, since an existing experimental station for RFI studies was already located there. One success of our tests at Green Bank is the demonstration that even with the large baseline, the LMS adaptive filter is capable of reliably and effectively canceling a moving interferer in a highly nonstationary environment.

3. EXPERIMENTAL PROCEDURES AND EQUIPMENT

Tests were performed at the GBT in West Virginia to evaluate the effectiveness of *GLONASS* signal cancellation. High levels of cancellation and effective preservation of the desired signal were consistently demonstrated. This section describes the experimental approach and equipment used for these tests.

3.1. Reference Antenna Data Alignment

Since a clean (high-INR) copy of the interference signal is required, a steerable 3.6 m reference antenna was installed at the NRAO Green Bank facility approximately 1.4 km from the GBT (see Fig. 3) and fitted with a relatively wide band, low-noise amplifier (LNA). The reflector is a consumer-grade, aluminum mesh, 3.6 m dish manufactured by Kaultronics, Inc. The azimuth-elevation positioner consists of two Kaultronics HT-100 equatorial mount rotators joined together with a custom adapter mount. This inexpensive system follows the design concept of the “small radio telescope” developed by the Massachusetts Institute of Technology as an educational instrument. Satellite tracking software was purchased commercially (NOVA for Windows) for less than \$100. Antenna gain is approximately 32 dB at 1.6 GHz, and the uncooled LNA and receiver have a system temperature of about 53 K. Extremely low noise design is not required, since the interferer power level is relatively high. The entire reference antenna system was purchased for less than \$8000.

For proper canceler operation, $d[n]$ must be highly correlated with samples in x_n . For the bandwidths used in our experiments (1–4 MHz), delays caused by different propagation paths and cable lengths for the primary and reference channels can completely decorrelate these signals. This must be compensated for in order to align the correlation peak near the center of the filter sample vector, x_n . In our experiments, for example, there was a 10.73 μ s difference between the cable propagation delays from the GBT and the reference antenna to a central data collection location (see Fig. 3). It is also possible that $d[n]$ may be advanced relative to $x[n]$ due to orbital geometry. We aligned the two channels prior to the LMS filter using large circular memory buffers on the digital signal processor. This operation is similar




FIG. 3.—Details of relative positions and cable delays to the GBT control room for the GBT and 3.6 m antennas. The three-dimensional coordinate system origin for computing bulk time of arrival delays was set at the GBT phase center, which is 24 m higher (+z) than the reference antenna.

to the delay correction performed in synthesis array telescopes to position the field of view phase center on the object of interest. However, here delay values must be updated more frequently to accommodate rapid satellite motion. Delay values depend only on interferer position and not on the desired astronomical source, so with available database information, the positions of satellite interferers can be predicted weeks in advance of observations.

3.2. Real-Time DSP Implementation

A major part of the test platform is the real-time processing subsystem, which uses four high-end Texas Instruments floating-point digital signal processors (TMS 320C6701). Two of the processors have associated digital receiver front-end daughter boards, which sample the analog RF receiver intermediate-frequency output signal. These dedicated digital receivers perform many of the high data rate DSP operations, including digitally controlled adjustable gain analog IF amplification, low-pass filtering, analog-to-digital conversion, complex baseband mixing, tuning, band select filtering, and decimation. This frees up the digital signal processors to implement adaptive filtering and power spectral estimation. A custom *L*-band analog receiver front end was developed for use with our small telescope antennas at Brigham Young University, but for experiments at Green Bank the real-time processing subsystem was connected directly to the existing GBT analog receiver IF signal path. The DSP system is capable of up to four channels of 65 megasamples s^{-1} analog-to-digital conversion and uses industrial/military grade, off-the-shelf components from Pentek. For more specific hardware information, please see Poulsen (2003).

In order to evaluate the performance of the real-time adaptive filter, power spectral density (PSD) estimates were computed in real time for $d[n]$, $x[n]$, and $\hat{s}[n]$. It was thus possible to directly observe the depth of interference cancellation, the quality of the reference signal, and other characteristics of the RFI mitigation scenario. It was necessary to perform these functions in the DSP because of the limited transfer rate provided by the DSP-to-host PC interface. Time averaging (integration) performed in the PSD estimator yields significant data reduction, so the resulting spectra could be streamed to the PC host in real time. It is also possible to perform full rate digital-to-analog conversion and output the filtered signal back into the telescope IF signal path, although this was not done in the experiments reported below.

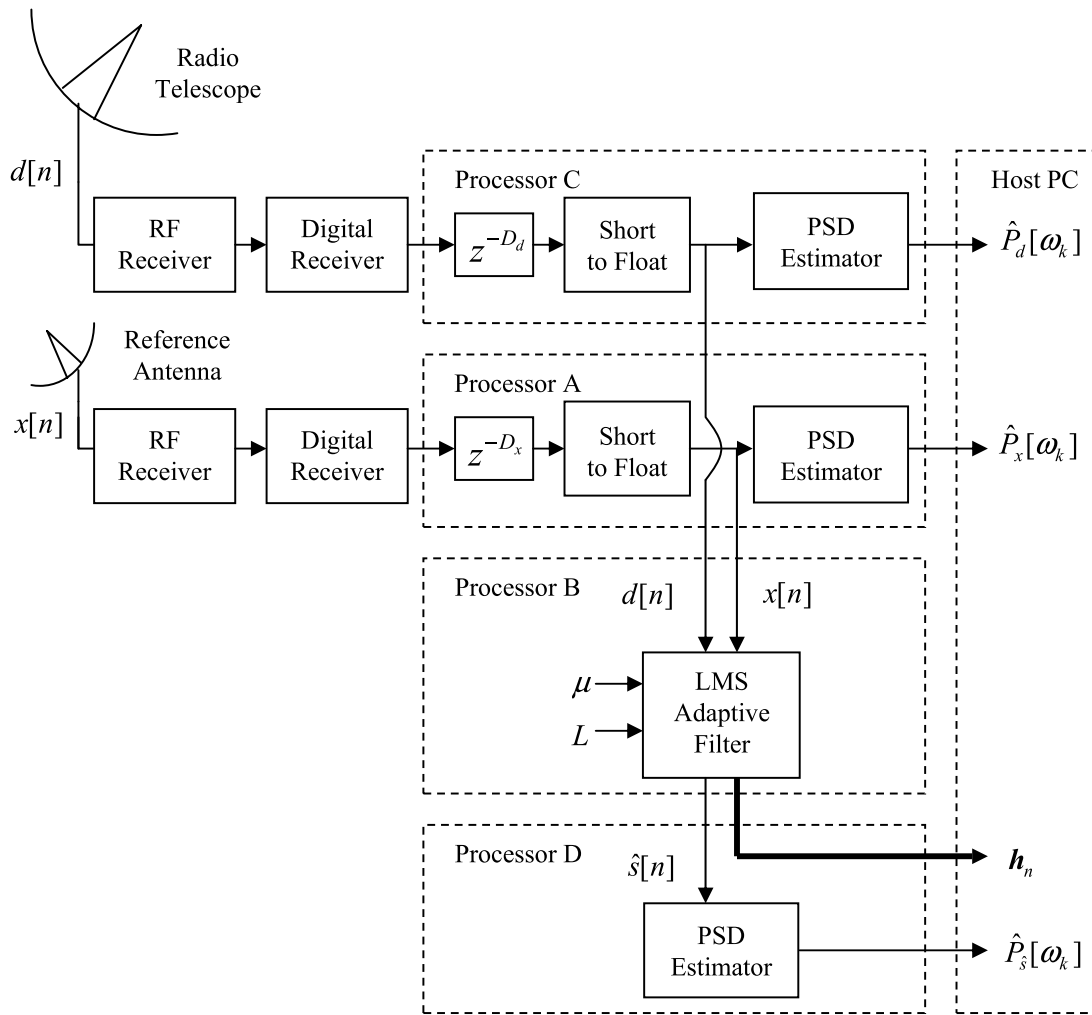


FIG. 4.—DSP implementation of the LMS adaptive filter.

Figure 4 illustrates the general structure of the DSP LMS filtering application. The inputs $d[n]$ and $x[n]$ are sampled by Pentek 6216 digital receivers and sent to processors C and A, respectively, at a complex baseband sample rate consistent with the selected signal bandwidth. If $d[n]$ and $x[n]$ need to be realigned in time, this is accomplished immediately after receiving the data in processors C and A; $d[n]$ is delayed by D_d , and $x[n]$ by D_x . Once the signals are roughly aligned, the signals are

TABLE 2
DSP LMS FILTER PROPERTIES

Signal Bandwidth (MHz)	Filter Order	Real Multiplications (s^{-1})
0.412500	42	173.2500×10^6
0.534375	30	160.3125×10^6
1.006250	12	121.1250×10^6
2.225000	5	111.2500×10^6
2.562500	4	102.5000×10^6
3.250000	3	97.5000×10^6
3.843750	2	76.8750×10^6

NOTES.—Shown are the processing bandwidths, highest respective filter orders, and number of real floating-point multiplications per second used to implement and update the LMS adaptive filter (measured using a 167 MHz processor clock speed).

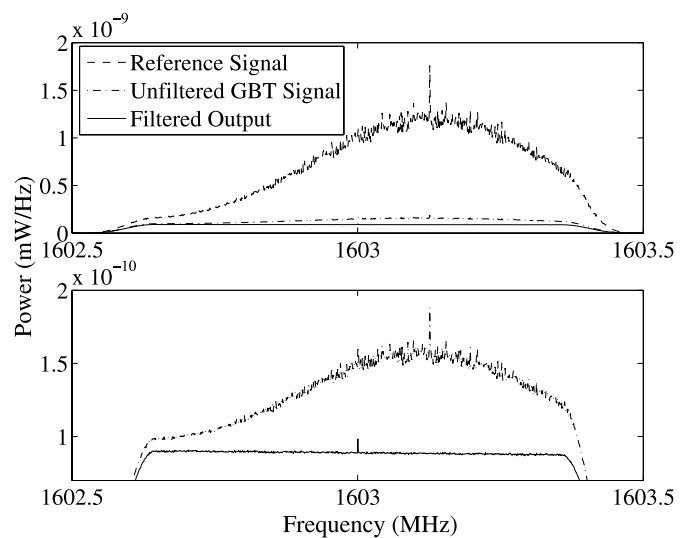


FIG. 5.—Example of real-time cancellation of *GLONASS* interference with the GBT. The lower plot “zooms in” on the filtered and unfiltered GBT signals to emphasize detail. Due to an inability to steer the GBT to an astronomical source, a “fake” source was injected at the feed, in this case at 1603 MHz. Note that the power levels shown in this and other plots are measured at the DSP receiver input and include all analog receiver front-end gains. [See the electronic edition of the Journal for a color version of this figure.]

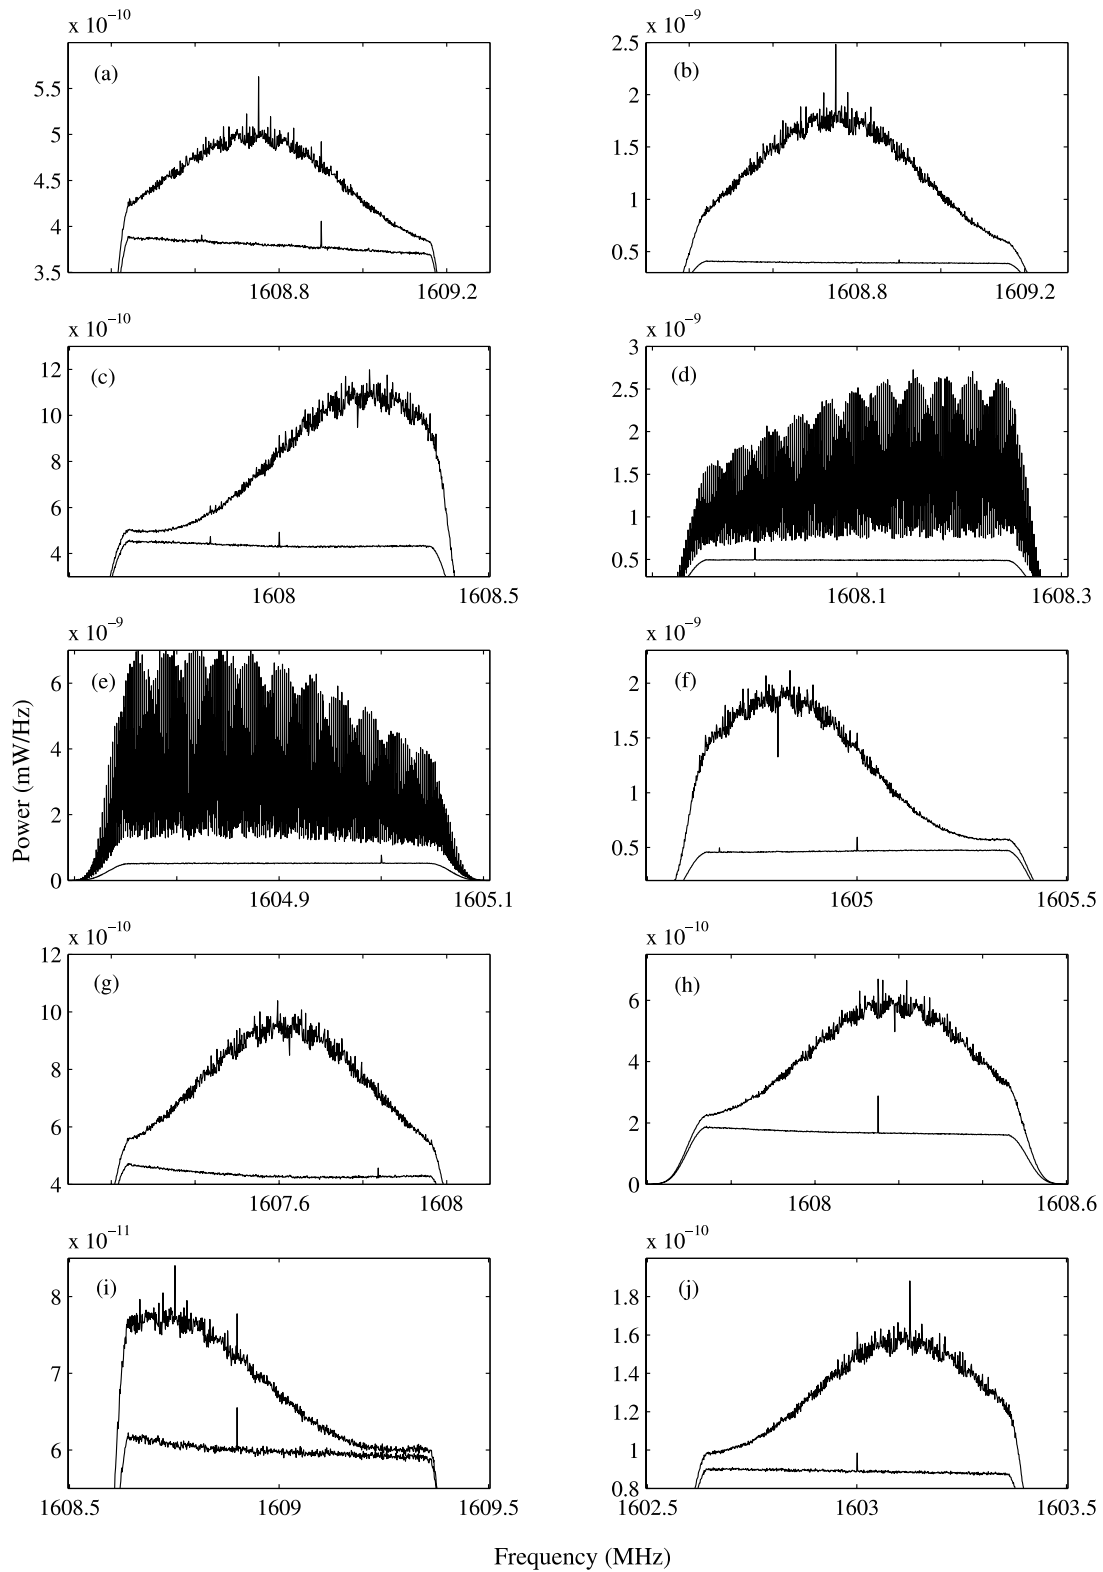


FIG. 6.—Examples of successful cancellation with the GBT. Before filtering, the test tone injected at the feed is hidden in the *GLONASS* spectrum. After cancellation, the desired signal emerges above a clean baseline. The test parameters are (*a-c, f-j*) $L = 12$, bandwidth of 1.00625 MHz and (*d, e*) $L = 42$, bandwidth of 0.4125 MHz. In each case, the upper and lower plots are the measured GBT spectra before and after filtering, respectively. [See the electronic edition of the *Journal* for a color version of this figure.]

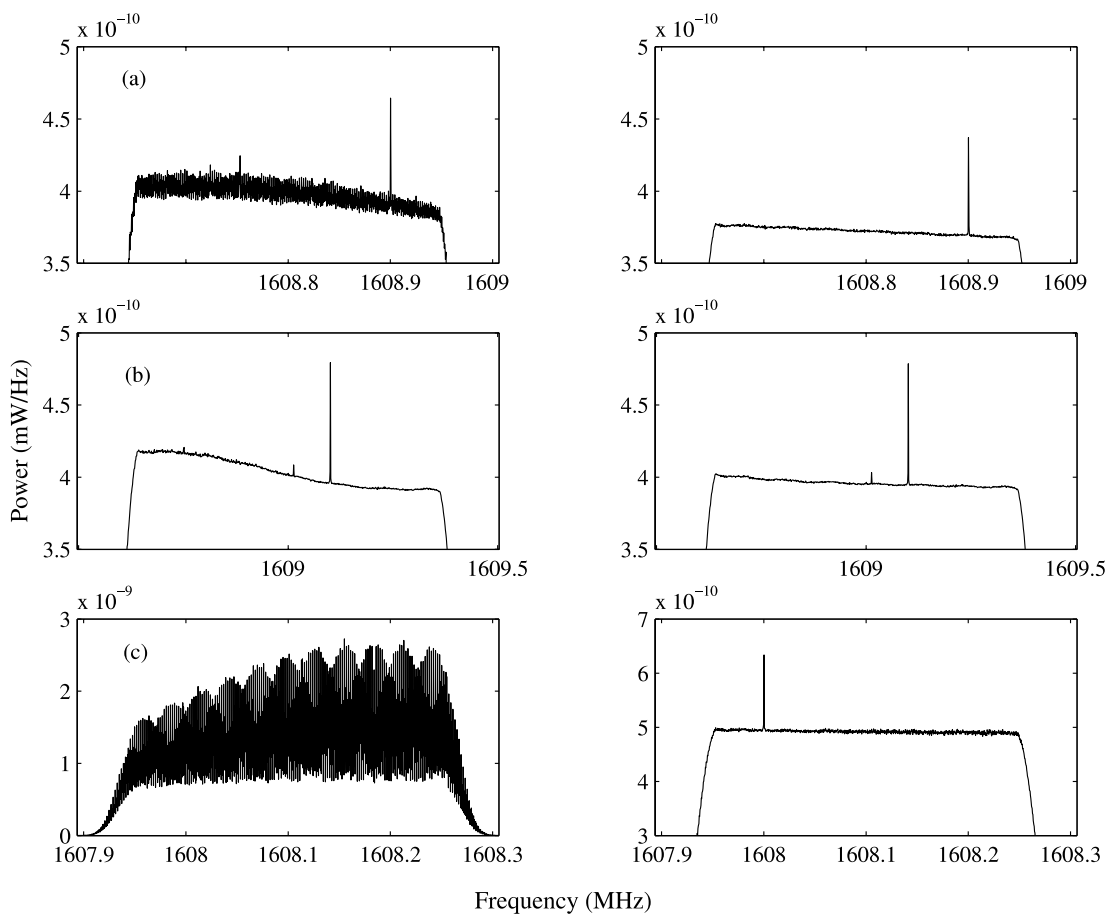


FIG. 7.—GBT signals before (left panels) and after (right panels) filtering. Each test contains a test tone injected at the feed of the GBT. Included are some of the longest integration data sets collected. All powers are at the input to the DSP platform. The test parameters are (a) $L = 42$, 1960 s integration; (b) $L = 12$, 1898 s integration; and (c) $L = 42$, 698 s integration. [See the electronic edition of the *Journal* for a color version of this figure. The electronic edition also contains *mpeg* animations of (a) and (c). The *mpeg* movies display snapshots of the spectra and complex filter coefficients with respect to time.]

converted to a floating-point value in preparation for processing. As shown in Figure 4, the computation of the PSD estimates of $d[n]$ and $x[n]$, $\hat{P}_d[\omega_k]$ and $\hat{P}_x[\omega_k]$, also occurs in processors C and A.

Adaptive filtering occurs in processor B. Since the computational load of the LMS algorithm is quite heavy, processor B is the limiting factor in determining the processing bandwidth, filter order, etc. Therefore, as many calculations as possible were pushed into the other three processors. Table 2 details some of the processing bandwidths and corresponding maximum filter orders for our real-time canceler.

3.3. GBT Test Scenario

In preparation for real-time tests with the GBT, we scheduled several different windows of time that would place 1612 MHz OH spectral line sources in the same passband as *GLONASS* interference. Unfortunately, inspectors found a crack in one of the wear plates of the GBT azimuth track that could not be repaired before the end of the scheduled tests. As a result, we were unable to position the GBT during a majority of the *GLONASS* cancellation experiments. However, we were able to run some convincing *GLONASS* cancellation experiments by injecting a low power signal at the GBT feed horn to simulate an astronomical source. The level was set to be consistent with typical OH sources (near 1 Jy flux) that we had observed earlier with the GBT. Significant interference cancellation was demonstrated

while preserving the desired signal. A remotely controlled signal generator in the GBT receiver room provided the “fake” source (see Fig. 5).

4. RESULTS AND DISCUSSION

This section presents results illustrating the typical interference attenuation and system performance with observations requiring long integration. In addition, we evaluate the required filter order. Finally, we address the difficulties of test scenarios with multiple interferers in the DSP bandwidth. Figure 6 illustrates 10 cases in which the test tone injected at the feed of the GBT is buried or hidden in actual *GLONASS* spectral interference. After cancellation, the test tone can be clearly seen.

4.1. Measuring Interference Attenuation

It is desirable to obtain an estimate of the achieved cancellation depth. Unfortunately, once the canceler pushes interference below the noise baseline estimation error for a given integration time, the final interference amplitude cannot be measured directly. In many cases, both the desired signal and interferer have power levels lower than the noise temperature and are only detected after integration. Successful observations depend on driving the variance of the noise plus residual interference baseline estimate to a level below the perturbation introduced by the desired signal.

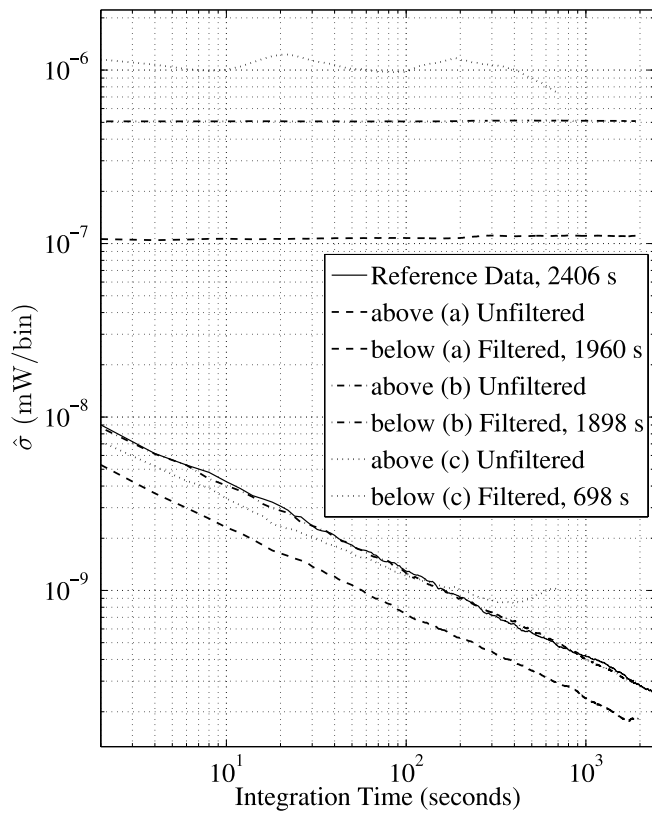


FIG. 8.—Estimates of standard deviation with integration time for those tests shown in Fig. 7. Residual interference is evident in test *c* by the fact that the standard deviation levels out. The reference data set is a 40 minute interference-free sample collected with the GBT. [See the electronic edition of the Journal for a color version of this figure.]

It is assumed that signal detection in frequency bin ω_k occurs when the difference between the current PSD estimate at the canceler output, $\hat{P}_s[\omega_k]$, and a baseline noise estimate, $\hat{P}_\eta[\omega_k]$, exceeds a threshold t ; i.e., there is a detection if

$$V[\omega_k] > t, \text{ where } V[\omega_k] = \hat{P}_s[\omega_k] - \hat{P}_\eta[\omega_k]. \quad (7)$$

The $\hat{P}_s[\omega_k]$ term contains the signal of interest, noise, and residual interference, whereas the $\hat{P}_\eta[\omega_k]$ term is computed from $d[n]$ at a different time when both the signal of interest and interference are absent. The threshold t is set proportional to the standard deviation of the estimation error in $\hat{P}_\eta[\omega_k]$, i.e.,

$$t \propto \sigma(\hat{P}_\eta[\omega_k]) \equiv E\left\{\left(\hat{P}_\eta[\omega_k] - P_\eta[\omega_k]\right)^2\right\}^{1/2}.$$

Any interference component in $\hat{P}_s[\omega_k]$ that is not removed with baseline subtraction appears in $V[\omega_k]$ and, if greater than t , masks

the signal of interest. Since $\sigma(\hat{P}_\eta[\omega_k]) \approx (2/M)^{1/2}P_\eta[\omega_k]$, where M is the number of independent fast Fourier transform averaging windows used in the PSD estimate (proportional to the integration time), the minimum detectable signal power reduces as $1/M^{1/2}$. Since interference terms in $V[\omega_k]$ are not reduced with integration, it is possible to determine the level of residual interference after cancellation by increasing integration time until the $1/M^{1/2}$ reduction in standard deviation of $V[\omega_k]$ is no longer observed. This effect is illustrated by several figures in this section.

An interference-free 40 minute reference data set for computing $\hat{P}_\eta[\omega_k]$ was collected. The standard deviation of this unfiltered data is included as a reference in Figures 8 and 10. Also, a few long integration tests were run in order to evaluate filter stability and robustness, as shown in Figure 7. Tests *a* and *b* contained, on average, a relatively low level of interference, often well below the noise floor, with sporadic spurts of higher interference. Data set *c*, on the other hand, contained average interference power levels 20 times (13 dB above) the integrated noise floor. Note that in the spectral plots of Figure 7, test *c* has visible baseline perturbation due to *GLONASS* interference after filtering.

Figure 8 illustrates the standard deviation and magnitude of cancellation for each of the long integration tests shown in Figure 7. In each case the standard deviation was estimated with frequency bins containing significant levels of interference previous to filtering. In other words, the cleanest regions of the output spectrum were not selected to compute the standard deviation. Since the data sets were taken in widely varying scenarios, the initial standard deviation estimate differs for each case. This is due to changes in the system gain, the size of the PSD bin widths, etc. Note that in Figure 8 the output variance of tests *a* and *b* declined linearly as a function of integration time, with a slope similar to that of the unfiltered interference-free reference data set. In these cases the *GLONASS* interference has been pushed imperceptibly below the integrated noise floor. On the other hand, some residual interference is evident at the longer integration times in test *c*.

In cases in which the variance levels out, one can easily estimate the actual interference attenuation. In other cases in which the interference remains buried in the noise, however, a lower bound on the achieved attenuation can be estimated with the available data as

$$\text{attenuation lower bound} \approx \frac{\min(\hat{\sigma}_s)}{\text{ave}(\hat{\sigma}_d)}, \quad (8)$$

where $\min(\hat{\sigma}_s)$ is the minimum integrated standard deviation of the canceler output obtained through integration and $\text{ave}(\hat{\sigma}_d)$ is the average deviation before filtering (the average interference power in $d[n]$ encountered during the test).

The attenuation bound for the tests of Figures 7 and 8 is presented in Table 3. The variance did not level out for tests *a* and *b*; therefore, the attenuation for these tests is most likely

TABLE 3
ADDITIONAL DETAILS FOR TESTS IN FIGS. 7 AND 8

Test	L	Bandwidth (MHz)	Integration Time (s)	Number of Bins ^a	Frequency Range ^a (MHz)	Minimum Attenuation Power Ratio	Minimum Attenuation (dB)
<i>a</i>	42	0.4125	1960	541	1608.66–1608.88	0.0016	–28.0
<i>b</i>	12	1.00625	1898	301	1608.69–1608.99	0.00056	–32.4
<i>c</i>	42	0.4125	698	401	1608.05–1608.22	0.0013	–29.4

NOTE.—The minimum attenuations in the final column were calculated using eq. (8).

^a Used to estimate $\hat{\sigma}$.

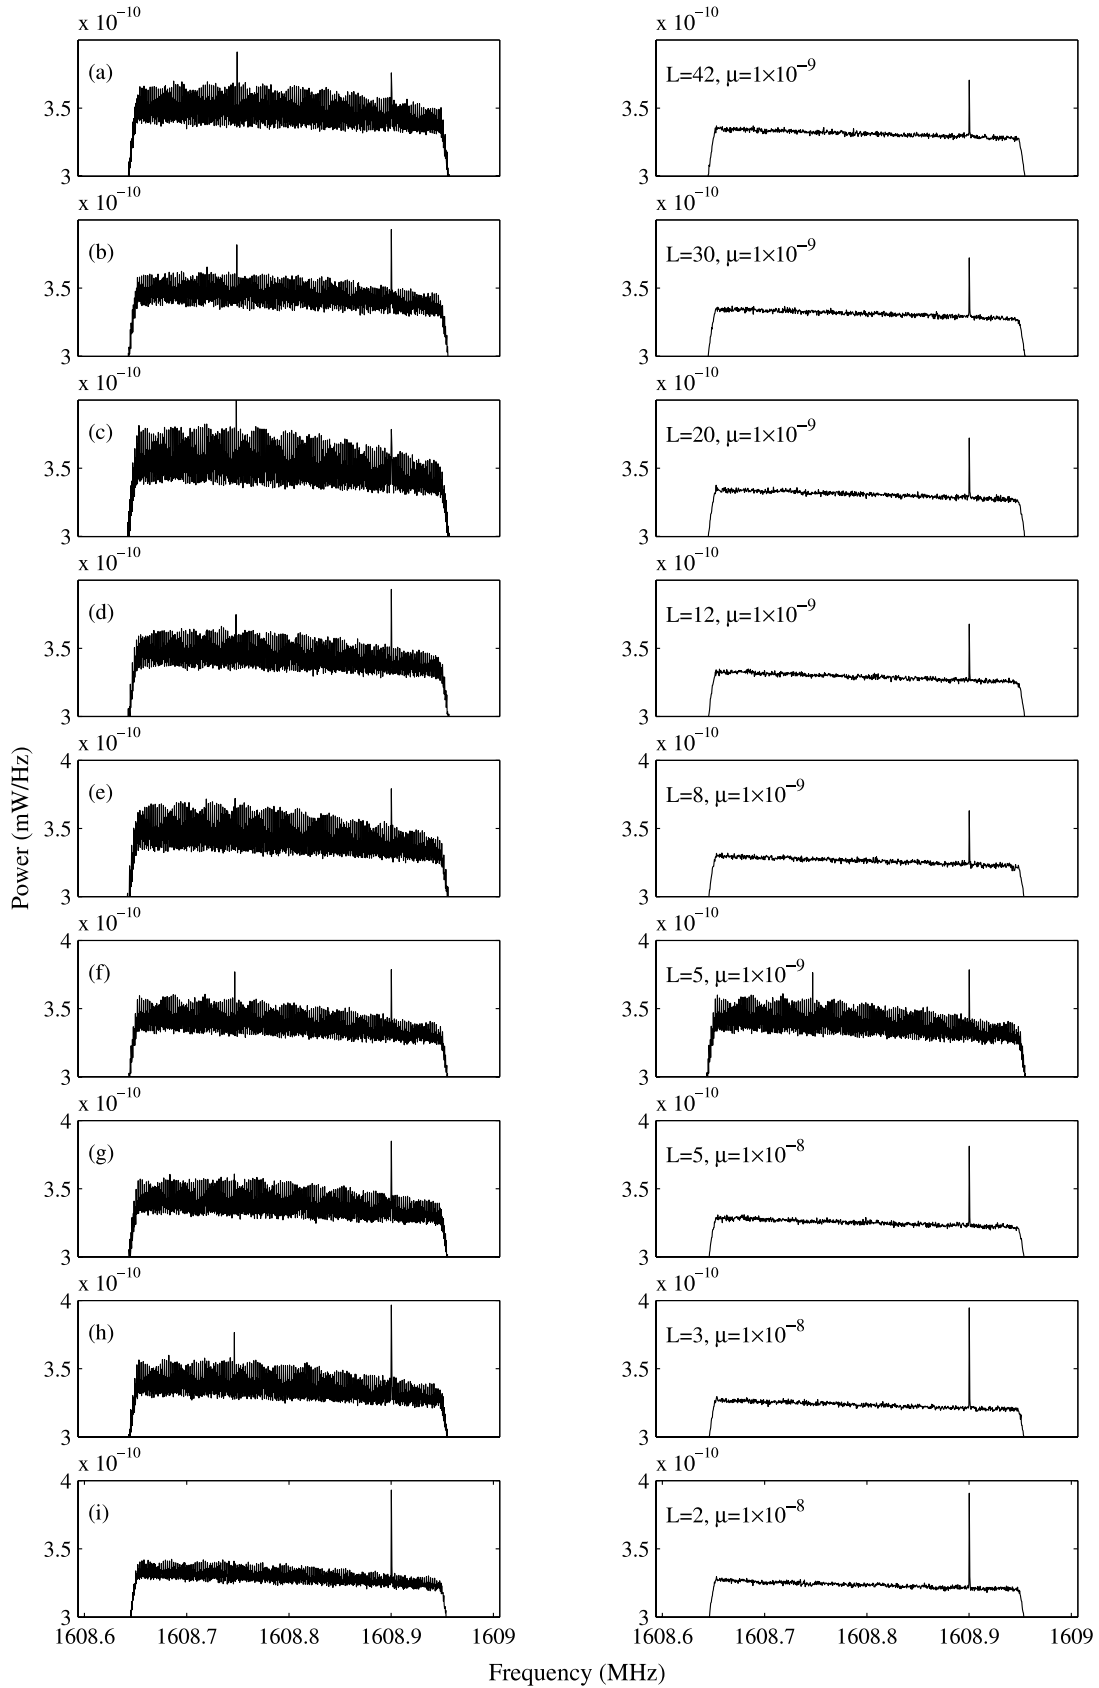


FIG. 9.—GBT signal powers before (*left panels*) and after (*right panels*) filtering. Each real-time cancellation test represents 298 s of integration and a bandwidth of 0.4125 MHz. All tests were run consecutively in order to maintain continuity. The smallest filter orders required a larger adaptation constant in order to track the dynamic system. [See the electronic edition of the *Journal* for a color version of this figure.]

higher than listed. Due to the wide range of interference levels represented by these and other tests, we conclude that a power attenuation factor of 0.001 (−30 dB) is realistically achievable for many test scenarios. This is in addition to any attenuation due to the sidelobe spatial response of the primary antenna.

4.2. Filter Length

An important consideration for canceler performance and computational burden is the minimum filter length required for effective cancellation. The “best” length, L , depends on many factors. These include, but are not limited to, the presence of multipath structure, the length of the baseline between the primary and auxiliary antennas, the data alignment accuracy, the bandwidth of the interferer, and the presence of signals at the auxiliary other than the satellite interferer. It has been shown that for the LMS algorithm, the average convergence time constant, τ_{av} , is proportional to L/μ and that the misadjustment error (the random residual cancellation error at convergence), M , is proportional to μL (Haykin 1996). Thus, using an unnecessarily long filter can adversely affect the ability both to track signal nonstationarities and to achieve high levels of interference cancellation.

With limited time allotted for real-time GBT tests and sporadic availability of interference (*GLONASS* satellites), exhaustive tests were not possible. We did, however, run some experiments by systematically decreasing the filter length while leaving as many other parameters constant as possible. It is important to note that because all experiments were run in real time, each test features a distinct data set. In order to minimize the effect of nonstationary signals, each series of tests was collected as close together in time as was possible.

The complex sample rate/signal bandwidth was 0.4125 MHz, with 5 minutes of integration and various filter orders from 42 down to 2 taps. These results are displayed in Figures 9 and 10. Each filter length tested suppressed interference to undetectable levels after 5 minutes of integration. Discrepancies for the smaller filter orders may have been evident with longer integration. For filter orders of 2–5, the adaptation constant was raised from $\mu = 1 \times 10^{-9}$ to 1×10^{-8} to decrease τ_{av} and ensure convergence while tracking signal nonstationarities.

As demonstrated, successful cancellation can occur with even the smallest filter lengths. There are, however, significant drawbacks to using small filter lengths. First, even tiny errors in data realignment can prevent interference attenuation. With longer filter lengths, one can place the highest correlated sample of the reference signal, $x[n]$, a few lags from the center of the filter and still not overstep the end of the finite-length filter.

Second, any large differences in interference signal characteristics seen in $i_p[n]$ and $i_r[n]$ will require a larger filter order. Differences could include frequency-dependent gain and/or phase differences in the primary and reference channels, multipath structure, frequency misalignment between channels resulting from improper synchronization of local oscillators, the presence of other interferers, etc. Often, several taps are required to adequately adapt to such channel differences.

In summary, there is an unfortunate dilemma of excess available computing power for small bandwidths and a lack thereof for larger bandwidths. Table 2 illustrates this fact. For bandwidths of 0.4125 and 3.84375 MHz, the highest available complex filter orders in our system were 42 and 2 taps, respectively. Ideally, the opposite would be true. For large processing bandwidths with significant channel differences, filter orders such as 42 or larger would be ideal. With the expected continued improvements in computing capacity in DSP platforms, or with a field programmable gate array implementation, real-time cancellation with these




FIG. 10.—Estimates of standard deviation with integration time for those tests shown in Fig. 9. All tests declined at the appropriate slope (except for test f , in which the adaptation constant was too small and cancellation did not occur). The plot of $\hat{\sigma}$ for test f is not included here. [See the electronic edition of the Journal for a color version of this figure.]

filter lengths will be less problematic. For a vast majority of OH maser sources, a bandwidth of 1 MHz is sufficient to include both the redshifted and blueshifted spectral peaks. For a 1 MHz bandwidth, 12 complex taps was sufficient for canceling most *GLONASS* interference scenarios encountered with the GBT. The filter order of 42 used for smaller bandwidths proved to be somewhat of an overkill; seemingly equivalent cancellation was obtainable with shorter filter lengths.

4.3. Multiple Interferers

With multiple interferers simultaneously above the horizon, it is possible, even probable, that each will spill through the telescope sidelobes with varying intensities. Unfortunately, only those interferers seen by the reference antenna will be adaptively subtracted, leaving the others untouched. This is not a problem if the secondary interferers fall outside of the desired signal bandwidth or at power levels low enough to be undetectable with signal integration. In the case of *GLONASS* downlink signals, however, adjacent frequency channels lie 0.5625 MHz apart and can be relatively wide band. With only half of the desired 24 *GLONASS* satellites currently deployed, this problem could worsen in the future, especially if the upper frequency channels coincident with most OH spectral lines are used.

Multiple interferers were encountered on a few occasions while running GBT experiments. Figure 11 illustrates one such case. Two *GLONASS* satellites both lie within the processing bandwidth. The reference antenna collected the interferer seen with greater intensity by the GBT, while the weaker second interferer remained untouched through the filtering process.

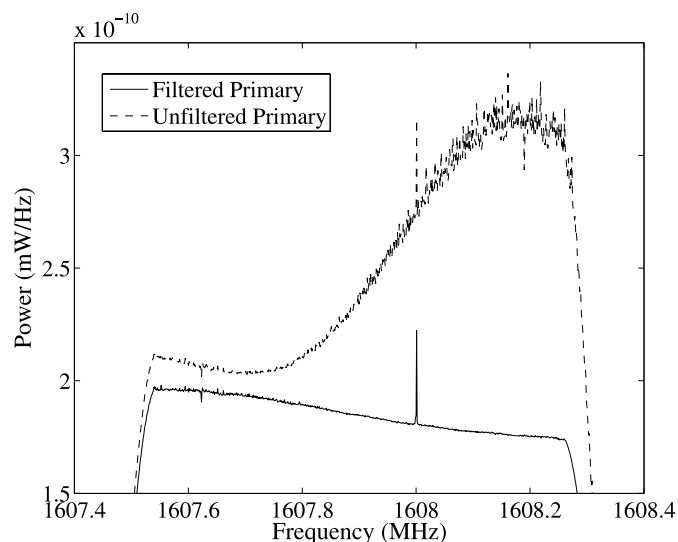


FIG. 11.—GBT spectrum corrupted by two *GLONASS* interferers which were simultaneously in the processing bandwidth, centered at 1607.0625 and 1608.188 MHz. Because the reference antenna only collected the interferer centered at 1608.188 MHz, the signal centered at 1607.0625 MHz remained untouched through the filtering process. The 1608 MHz test tone also remained untouched. The upper and lower curves are the measured spectra for before and after filtering, respectively. This figure along with an associated animation detailing the interferer presence with respect to time clearly indicate the need for a second reference antenna. [See the electronic edition of the *Journal* for a color version and the mpeg animation of this figure.]

An approach to canceling multiple interferers while using the current DSP software would be to cascade LMS filters, one for each interferer, and each with a separate tracking reference antenna. This would require no new code development but just additional DSP boards and associated digital receiver daughter boards (including the analog-to-digital and digital-to-analog functions). Other approaches would likely be more favorable in terms of stability and performance, but this method would be a simple and possibly successful solution. Each signal could be appropriately delayed according to the corresponding interferer's position.

This work was supported by the National Science Foundation, grant AST 99-87339.

APPENDIX

The LMS canceler is structurally in a system identification configuration. The \mathbf{h}_n term estimates, or “identifies,” the true transfer function, \mathbf{h}_n^* , between the reference and primary antennas and must track all changes over n . Tracking performance can be expressed in terms of filter misadjustment M , defined as

$$M = \frac{E\{|\hat{s}[n]|^2\} - J_{\min}}{J_{\min}}, \quad (\text{A1})$$

where

$$J_{\min} = E\{|y[n] - \mathbf{h}_{\text{mmse}, n}^T \mathbf{x}[n]|^2\}. \quad (\text{A2})$$

The first numerator term is the LMS filter mean-squared error, and J_{\min} is the theoretical minimum mean squared error achieved

by the Wiener filter $\mathbf{h}_{\text{mmse}, n}^*$. Low values of M imply good tracking cancellation performance.

Although dynamic changes over n in \mathbf{h}_n^* are due to deterministic orbital motion, we model them as a random process so that the problem allows a statistical analysis. The continuous valued Markov random process model is a powerful but simple tool that provides a good first-order approximation for non-stationarity in \mathbf{h}_n^* (Stark & Woods 2002, pp. 421–430). At each time sample, the new \mathbf{h}_{n+1}^* is represented as the previous value perturbed by a zero-mean, wide-sense, stationary, random-noise vector sequence $\mathbf{q}[n]$ called the Markov process noise:

$$\mathbf{h}_{n+1}^* = a\mathbf{h}_n^* + \mathbf{q}[n]. \quad (\text{A3})$$

Here a determines how much the new \mathbf{h}_{n+1}^* depends on the past and is assumed to be close to 1. System identification misadjustment can then be approximated as (Haykin 1996)

$$M \approx \frac{\mu}{2} \text{tr}\{\mathbf{R}_{i_r, n}\} + \frac{1}{2\mu\sigma_v^2} \text{tr}\{\mathbf{Q}\}, \quad (\text{A4})$$

where $\text{tr}\{\dots\}$ is the matrix trace operator, $\mathbf{R}_{i_r, n} = E\{\mathbf{i}_{r, n} \mathbf{i}_{r, n}^H\}$, $\mathbf{Q} = E\{\mathbf{q}[n] \mathbf{q}^H[n]\}$ is the autocovariance of $\mathbf{q}[n]$, and

$$\sigma_v^2 = E\{|s[n] + \eta_p[n] + \eta_h[n]|^2\} = \sigma_s^2 + \sigma_{\eta_p}^2 + \sigma_{\eta_h}^2.$$

If the reference antenna achieves a high INR and the desired signal $s[n]$ is below the noise floor (the usual situation), then we can substitute $\sigma_v^2 \approx \sigma_{\eta_p}^2$ and $\mathbf{R}_{i_r, n} \approx \mathbf{R}_{x_n}$.

Consider a single interfering satellite following a worst-case ground track parallel to the length b baseline. For an interfering source with no multipath propagation, the true transfer function is a bulk time delay with a single complex gain term. This has the particularly simple form

$$\mathbf{h}_n^* = [0 \dots 0, \alpha_n, 0 \dots 0]^T, \quad (\text{A5})$$

where α_n is a complex scalar whose position in the vector depends on the reference data alignment step of § 3.1. All *GLONASS* signals observed at the GBT had transfer functions of approximately this form. During one time sample interval, the difference between propagation path lengths from the satellite to the primary and reference antennas changes by approximately $b \sin \Delta\theta$. This induces a change in the phase of α_n that can be expressed by the update equation

$$\mathbf{h}_{n+1}^* = \mathbf{h}_n^* \odot \left[1 \dots 1, \exp\left(j \frac{2\pi f_0 b}{c} \sin \Delta\theta\right), 1 \dots 1 \right]^T, \quad (\text{A6})$$

where the circled dot indicates an element-wise Schur matrix product. This ensures no change in the zero-valued elements of \mathbf{h}_n^* while phase shifting α_n to account for propagation path length changes. Using a small-angle approximation, the phase vector multiplication can be approximated by an additive term:

$$\mathbf{h}_{n+1}^* \approx \mathbf{h}_n^* + \mathbf{u}_n, \quad (\text{A7})$$

$$\mathbf{u}_n = \left[0 \dots 0, \alpha_n \left(\exp\left(j \frac{2\pi f_0 b \Delta\theta}{c}\right) - 1 \right), 0 \dots 0 \right]^T. \quad (\text{A8})$$

The Markov model of equation (A3) has the same form as equation (A7) if \mathbf{u}_n is equated with $\mathbf{q}[n]$ and $a \approx 1$. Thus, we interpret \mathbf{u}_n as the random Markov process noise (although, strictly speaking, it is not temporally white, as is usually assumed). Thus, let

$$\mathbf{Q} = E\{\mathbf{u}[n]\mathbf{u}^H[n]\} \quad (\text{A9})$$

$$= E\left\{ \begin{bmatrix} 0 & & \\ & \left| \alpha_n \exp\left(j\frac{2\pi f_0 b \Delta\theta}{c}\right) - 1 \right|^2 & \\ & & 0 \end{bmatrix} \right\}, \quad (\text{A10})$$

which yields

$$\text{tr}\{\mathbf{Q}\} = E\{|\alpha_n|^2\} \left(2 - 2 \cos \frac{2\pi f_0 b \Delta\theta}{c} \right). \quad (\text{A11})$$

Here $E\{|\alpha_n|^2\} \approx \sigma_{i_p}^2 / \sigma_{i_r}^2$ and for high INR will be equal to the magnitude squared of the maximum tap in \mathbf{h}_n . Substitution into equation (A4) produces the final expression for misadjustment:

$$M \approx \frac{\mu}{2} L \sigma_{i_r}^2 + \frac{1}{\mu \sigma_{i_p}^2} \frac{\sigma_{i_p}^2}{\sigma_{i_r}^2} \left(1 - \cos \frac{2\pi f_0 b \Delta\theta}{c} \right). \quad (\text{A12})$$

REFERENCES

- Baan, W., Fridman, P., & Millenaar, R. 2004, in *Mitigation of Radio Frequency Interference in Radio Astronomy* (Blacksburg: Virginia Tech.), <http://www.ece.vt.edu/~swe/RFI2004>
- . 2004, *AJ*, 128, 933
- Barnbaum, C., & Bradley, R. F. 1998, *AJ*, 116, 2598
- Ellingson, S. W., Bunton, J. D., & Bell, J. F. 2001, *ApJS*, 135, 87
- Hayes, M. H. 1996, *Statistical Digital Signal Processing and Modeling* (New York: Wiley)
- Haykin, S. 1996, *Adaptive Filter Theory* (3rd ed.; Englewood: Prentice-Hall)
- Kesteven, M. 2005, in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, Vol. 5 (Los Alamitos: IEEE), 873
- Kesteven, M., Hobbs, G., Clement, R., Dawson, B., Manchester, R., & Uppal, T. 2005, *Radio Sci.*, 40, RS5S06
- Poulsen, A. J. 2003, M.S. thesis, Brigham Young Univ.
- Stark, H., & Woods, J. 2002, *Probability and Random Processes with Applications to Signal Processing* (3rd ed.; Upper Saddle River: Prentice-Hall)
- Widrow, B. 1970, in *Aspects of Network and System Theory*, ed. R. E. Kalman & N. DeClaris (New York: Holt, Rinehart, & Winston)
- Widrow, B., & Hoff, M. E., Jr. 1960, in *Adaptive Switching Circuits*, IRE WESCON Convention Record, Vol. 4 (New York: IRE), 96 [reprinted in *Neurocomputing: Foundations of Research*, 1988, ed. J. A. Anderson & E. Rosenfeld (Cambridge: MIT Press), 126]