

A NEW APPROACH TO INTERFERENCE EXCISION IN RADIO ASTRONOMY: REAL-TIME ADAPTIVE CANCELLATION

CECILIA BARNBAUM

Space Telescope Science Institute,¹ 3700 San Martin Drive, Baltimore, MD 21218; barnbaum@stsci.edu

AND

RICHARD F. BRADLEY

National Radio Astronomy Observatory,² 520 Edgemont Road, Charlottesville, VA 22903; rbradley@nrao.edu

Received 1998 May 20; revised 1998 August 4

ABSTRACT

Every year, an increasing amount of radio-frequency (RF) spectrum in the VHF, UHF, and microwave bands is being utilized to support new commercial and military ventures, and all have the potential to interfere with radio astronomy observations. Such services already cause problems for radio astronomy even in very remote observing sites, and the potential for this form of light pollution to grow is alarming. Preventive measures to eliminate interference through FCC legislation and ITU agreements can be effective; however, many times this approach is inadequate and interference excision at the receiver is necessary. Conventional techniques such as RF filters, RF shielding, and postprocessing of data have been only somewhat successful, but none has been sufficient. Adaptive interference cancellation is a real-time approach to interference excision that has not been used before in radio astronomy. We describe here, for the first time, adaptive interference cancellation in the context of radio astronomy instrumentation, and we present initial results for our prototype receiver.

In the 1960s, analog adaptive interference cancelers were developed that obtain a high degree of cancellation in problems of radio communications and radar. However, analog systems lack the dynamic range, noised performance, and versatility required by radio astronomy. The concept of digital adaptive interference cancellation was introduced in the mid-1960s as a way to reduce unwanted noise in low-frequency (audio) systems. Examples of such systems include the canceling of maternal ECG in fetal electrocardiography and the reduction of engine noise in the passenger compartments of automobiles. These audio-frequency applications require bandwidths of only a few tens of kilohertz. Only recently has high-speed digital filter technology made high dynamic range adaptive canceling possible in a bandwidth as large as a few megahertz, finally opening the door to application in radio astronomy.

We have built a prototype adaptive canceler that consists of two receivers: the primary channel (input from the main beam of the telescope) and a separate reference channel. The primary channel receives the desired astronomical signal corrupted by RFI (radio-frequency interference) coming in the sidelobes of the main beam. A separate reference antenna is designed to receive only the RFI. The reference channel input is processed using a digital adaptive filter and then subtracted from the primary channel input, producing the system output. The weighting coefficients of the digital filter are adjusted by way of an algorithm that minimizes, in a least-squares sense, the power output of the system. Through an adaptive-iterative process, the canceler locks onto the RFI, and the filter adjusts itself to minimize the effect of the RFI at the system output. We have designed the adaptive canceler with an intermediate frequency (IF) of 40 MHz. This prototype system will ultimately be functional with a variety of radio astronomy receivers in the microwave band. We have also built a prototype receiver centered at 100 MHz (in the FM broadcast band) to test the adaptive canceler with actual interferers, which are well characterized. The initial laboratory tests of the adaptive canceler are encouraging, with attenuation of strong frequency-modulated (FM) interference to 72 dB (a factor of more than 10 million), which is at the performance limit of our measurements. We also consider requirements of the system and the RFI environment for effective adaptive canceling.

Key words: instrumentation: detectors — methods: analytical

1. INTRODUCTION

Just as optical astronomy faces serious problems with light pollution, radio astronomy is similarly plagued by its own brand of light pollution: radio-frequency interference

(RFI). Growing technological resources usurp more and more radio-frequency spectrum in the VHF, UHF, and microwave bands. Advances in both microprocessor and monolithic microwave integrated circuits have spurred a plethora of new applications, including improved point-to-point communications, wireless computer communications, and a growing number of cellular telephones; all have potential to interfere with radio astronomy observations. Signals from some existing satellites leak into protected radio astronomy bands, and the problem will likely increase as more satellites are put into orbit. Low-Earth-orbit satel-

¹ The Space Telescope Science Institute is a facility of the National Aeronautics and Space Administration operated under contract by the Association of Universities for Research in Astronomy, Inc.

² The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

lites (LEOs) create an RFI problem no matter where on Earth a radio telescope is located, even in remote locations such as Tasmania and the Antarctic. The increasing congestion of the radio spectrum is making astrophysical research in the designated radio and microwave bands ever more difficult. Outside the bands protected for astronomy the situation is quite a bit worse.

Preventive measures through FCC legislation and ITU agreements can be effective, such as protected radio astronomy bands and designated radio-quiet areas like the National Radio Quiet Zone around Green Bank Observatory, in West Virginia (Sizemore 1991). When prevention is not possible, various methods to reduce or eliminate interference have been employed; conventional approaches include (1) blanking techniques to remove pulse-type signals from the data stream (Gerard 1983; Fisher 1983), (2) filtering techniques such as superconducting notch filters to remove fixed-frequency interference (e.g., Moffet 1982), (3) radio-frequency (RF) shielding to suppress spurious digital signals and local oscillator signals from adjacent electronic equipment or communication systems (Morrison 1998), (4) postprocessing techniques on array systems, such as sidelobe-beam nulling to remove fixed-frequency signals (Erickson 1983), and (5) adaptive beam-forming techniques (Goris 1997). All these schemes are successful to some degree, yet each suffers from either insufficient interference cancellation, inability to adapt to changing statistics of the interference signal, partial removal of wanted data, or excessively large amounts of postprocessing of the accumulated data. Clearly, we need to investigate new approaches to interference excision that have potential to improve upon the shortcomings of conventional techniques.

Even with all the efforts described above, removing RFI in a radiometer remains a very difficult exercise. Astronomical signals are weak compared with ground (or satellite) signals, and so large attenuation capability is necessary. Often, RFI and the astronomical signal occur at the same frequency, and conventional rejection schemes, such as notch filters, matched filters, and beam nulling, remove some or all of the desired signal along with the RFI. A notch filter is a very sharp band-reject filter that can be quite useful in removing strong interference which can cause saturation to occur in the amplifier and correlator. However, this type of filter is not easily adjustable in frequency and removes some of the desired astronomical signal. A matched filter is a digital filter whose characteristics "match" that of the interference and remove it from the radiometer data. However, the spectrum of the interference signal must be known a priori, and in some cases the filter characteristics are not physically realizable. The matched filter can also cause distortion of the desired signal. In array systems, beam nulling has been tried through extensive postprocessing of the data; although somewhat successful, this approach is very time-consuming and does not attenuate enough for astronomy observations at frequencies that are plagued by RFI. An adaptive beam-former is being designed for the Square Kilometer Array Interferometer (SKAI) (Goris 1997), but the problem here is that the weighting coefficients that create the null can also cause small distortions in the main telescope beam.

At first, one might think that a simple correlation technique could work well, given reference input from a completely separate radiometer that monitors the interference but not the astronomical source. With this approach, the

reference input is cross-correlated with the signal in the main feed, and, after correcting for differences in phase delay and amplitude of the RFI in the two inputs by way of weighting coefficients, the reference input is subtracted from that of the main feed. This approach would work as long as the main and reference inputs continued to see the same magnitude of *change* in phase delay or amplitude; if the *relative difference* in the characteristics of the RFI in the reference and primary inputs does not change, the situation is stationary, and the weighting coefficients do not change over an integration time. However, in the real world conditions change as the telescope tracks or as propagation phenomena, such as reflection and multipath, affect the RFI at the reference differently than at the main feed. In these non-stationary conditions, correcting coefficients must be able to update in real time as the statistics change. To cancel RFI well enough to retrieve astronomical signals, the canceling scheme must be able to *adapt* to changing statistics.

For the simple case of stationary interference, a linear filter is used to minimize the least mean square value of the difference between the desired response and the actual output. The resulting solution is commonly known as the Wiener solution. For the more realistic case of nonstationary conditions, the Wiener solution is inadequate since it is nonadaptive. To be effective in a nonstationary environment, the Wiener solution must be updated as quickly as conditions change. An adaptive canceler initially finds the Wiener solution and continually updates the solution by using a recursive algorithm to track the changing statistics, providing that these variations are sufficiently slow. This is the essential difference between a stationary filter (such as the Wiener variety) and the nonstationary *adaptive* filter we describe here.

Adaptive canceling has not been used before in radio astronomy, but the concept is not a new one. Since the 1960s, analog adaptive canceling has been used successfully for antijamming in radar and interference excision in communication systems (Ghose 1996). Analog adaptive canceling is based on hardware techniques that are inherently noisy and lack the dynamic range needed for radio astronomy applications. Digital techniques solve these problems but until recently have been limited to audio applications because of their slow processor speed. In 1965, an adaptive echo canceler for telephone lines was developed (Sondhi 1967). In the early 1970s Stanford University students used adaptive filtering to cancel the maternal signal in fetal electrocardiography, where the electrical signature of the mother's heart has an amplitude from 2 to 10 times stronger than that of the fetal heartbeat (Widrow et al. 1975). Since then, digital audio interference cancellation systems have been developed for many diverse applications from speech enhancement in noisy environments to the reduction of harmful noise in factories. Audio applications require bandwidths of only a few tens of kilohertz, which do not require especially fast digital hardware. Fast digital technology for adaptive interference cancellation over the wide bandwidths necessary for astronomy has become available only relatively recently. Operation up to a few megahertz can now be performed using modern digital signal processing chips at reasonable cost, such as the Logic Devices Inc. LMA1009 12 × 12 bit multiplier-accumulator (Logic Devices 1995, p. 4). Interference canceling by an adaptive filter in the context of radio astronomy requires a low-gain reference antenna that monitors the RFI. The

adaptive filter then tracks the two inputs and removes the RFI, leaving the desired astronomical signal. The canceling is accomplished in real time, and postprocessing of data is not required.

Given the serious, ever-growing problems that RFI poses to radio astronomy, new and sometimes radical approaches for interference cancellation must be examined. If the adaptive interference-canceling scheme is shown to be viable for use in radio astronomy, the potential would be significant. Viability will depend on how well it attenuates RFI and how much noise is added to the system output as a result. The prototype adaptive-canceling system we have designed can be used with any radiometer accepting a 40 MHz IF. We have characterized the performance of the canceler in the laboratory and present the results here. The next step will be to field-test the system using actual interference signals. We chose a portion of the FM broadcast band near 100 MHz, certainly a worst-case frequency location for astronomical observations, and have developed appropriate low-noise receivers for these tests. The FM broadcast band spans 88–108 MHz and contains channels that are spaced 200 kHz apart. The signals are relatively narrow in bandwidth (< 100 kHz at the 3 dB level) with nonvarying central frequencies, fixed power levels, and known transmit antenna polarizations. Multiple FM signals in the filter bandpass will allow us to test fully the capability of the adaptive canceler. Even at a remote site such as at NRAO's Green Bank Observatory in West Virginia, the entire FM band is occupied (Fig. 1). The spectrum shown in Figure 1 makes it quite obvious that it would be impossible to observe in this band without highly effective interference excision. The initial bench tests with our prototype system are encouraging, with attenuation of strong frequency-modulated interference down to at least 70 dB (a factor of 10 million), at the limit of the capabilities of our measurements.

In this paper we describe, for the first time, adaptive interference canceling in the context of radio astronomy instrumentation (see also Bradley et al. 1996, 1997). The basic theory of the adaptive-canceling system is described in § 2. The special case of the stationary solution and expected performance are described in § 3, and the nonstationary solution is discussed in § 4. The design of our prototype

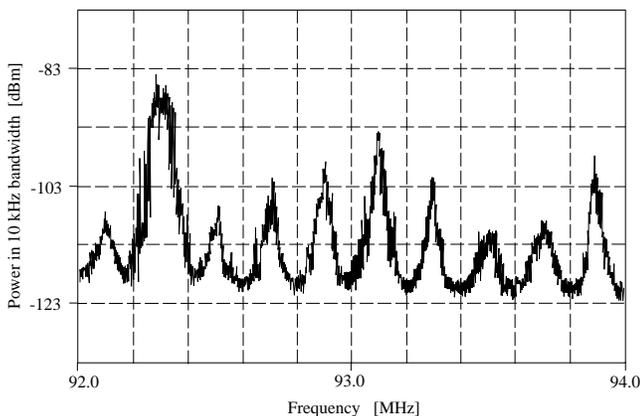


FIG. 1.—Spectrum of the FM broadcast band near 93 MHz. The data were taken using the 50–500 MHz receiver and cross-dipole feed on the 140 foot (43 m) telescope at Green Bank Observatory. The receiver gain was 44.5 dB and the system temperature was 750 K, with a bandwidth resolution of 10 kHz.

adaptive canceler and results of initial trials are presented in § 5.

2. BASIC CONCEPTS OF ADAPTIVE INTERFERENCE CANCELING

An adaptive interference-canceling system for use on a radio telescope is illustrated in Figure 2. All signals are digitized and have a constant sampling period with discrete time sequences indexed by n . The telescope radiometer (or *primary channel*) receives both the astronomical signal $s(n)$ entering the main beam and the interference $i_p(n)$ entering the sidelobes. The primary input is the sum of these two signals, $s(n) + i_p(n)$. A low-gain antenna is connected to a second receiver (the *reference channel*), whose input is only the interference $i_r(n)$. The reference antenna is pointed toward the interferer (at the satellite or toward the horizon). The collecting area of the reference antenna is much smaller than the primary telescope, and therefore the astronomical signal would be essentially absent in the reference input on the filter's adaptation timescale. The interference in the reference channel $i_r(n)$ is correlated in some unknown way with the interference in the primary channel $i_p(n)$, and the job of the adaptive filter is to estimate this correlation as a function of time. The adaptive algorithm compares the previous solution to current information and sends updated coefficients to the digital filter. The digital filter uses these coefficients to alter $i_r(n)$, producing $y(n)$, which closely resembles the interference in the primary channel; $y(n)$ is subtracted from the primary input to produce the system output $\epsilon(n)$:

$$\epsilon(n) = s(n) + i_p(n) - y(n), \quad (1)$$

and then $\epsilon(n)$ is sent to the telescope's spectrometer. No prior knowledge is required of $s(n)$, $i_p(n)$, $i_r(n)$, or of their interrelationships.

The signal path through the adaptive filter shown in Figure 2 illustrates the iterative nature of the system. The adaptive algorithm finds new coefficients by comparing $\epsilon(n)$ with $\epsilon(n-1)$ using a least mean squares (LMS) algorithm that minimizes the total power. The power is the square the system output:

$$\epsilon^2(n) = s^2(n) + [i_p(n) - y(n)]^2 + 2s(n)[i_p(n) - y(n)]. \quad (2)$$

Since $s(n)$ is uncorrelated with the interference in the primary and reference channels, the cross terms vanish, and so the expectation value (time averaged) of the system output is

$$E\{\epsilon^2(n)\} = E\{s^2(n)\} + E\{[i_p(n) - y(n)]^2\}. \quad (3)$$

As the filter adjusts the coefficients to minimize $E\{\epsilon^2(n)\}$, the power in the astronomical signal $E\{s^2(n)\}$ is unaffected, and so $E\{[i_p(n) - y(n)]^2\}$ reaches a power minimum:

$$E_{\min}\{\epsilon^2(n)\} = E\{s^2(n)\} + E_{\min}\{[i_p(n) - y(n)]^2\}. \quad (4)$$

Since the astronomical signal is constant, *minimizing the total output power minimizes the output interference power and therefore maximizes the output signal-to-interference ratio.*

In a stationary environment, once the filter finds the weighting coefficients so that $y(n)$ is a best least-squares estimate of the interference in the primary channel, those coefficients are fixed. More specifically, even though the interference signal will have a nonzero bandwidth due to

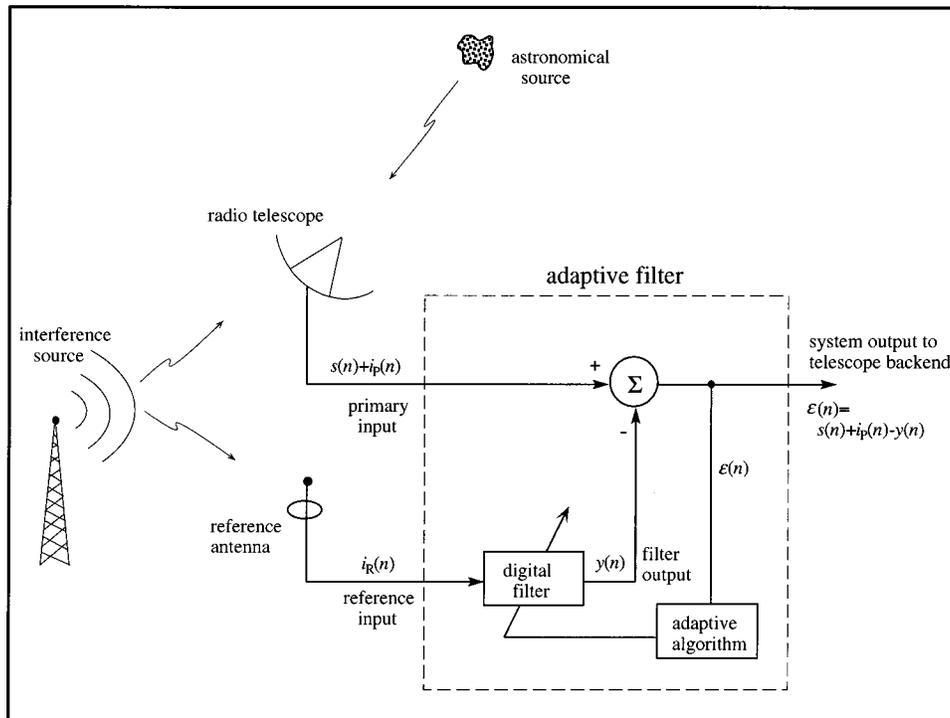


FIG. 2.—Schematic of the adaptive interference canceler. The telescope receiver, or *primary channel*, is located at prime focus and receives the primary input, which consists of the astronomical signal $s(n)$ in the main beam and the interference $i_p(n)$ entering the sidelobes. A low-gain antenna is connected to a second receiver, the *reference channel*, whose input is the interference $i_R(n)$, which is correlated in some unknown way with $i_p(n)$. The adaptive algorithm sends updated weighting coefficients to the digital filter, which are found by comparing the previous solution to the current information. The digital filter uses the weighting coefficients to alter $i_R(n)$ thus producing $y(n)$, which closely resembles $i_p(n)$. Subtracting $y(n)$ from the primary input produces the system output $\epsilon(n)$, which is then sent to the spectral processor. The signal path through the adaptive filter illustrates the iterative nature of this system. Note that no prior knowledge of $s(n)$, $i_p(n)$, or $i_R(n)$ is required.

modulation, as long as the statistics of the waveform (i.e., mean variance and autocorrelation) in both the primary and reference channels remain the same, the coefficients found initially will suffice. However, in more realistic conditions with a radio telescope, the weighting coefficients will quickly become obsolete if propagation effects cause a *relative* change in what the reference and primary channels see. Significant relative changes that require updated coefficients can be caused by reflection, dispersion, and telescope slewing; the timescale of these effects is on the order of hundreds of milliseconds. In an adaptive scheme, the coefficients used to weight $i_R(n)$ are updated; in our prototype receiver, coefficients are updated every $2 \mu\text{s}$.

Two special cases of interest can arise. First, if the reference input is perfectly correlated with the interference in the primary channel, the output signal will be completely free of interference, since $E\{[i_p(n) - y(n)]^2\} = 0$. In the second case, if the reference channel is completely *uncorrelated* with the interference in the primary channel (e.g., if for some reason the RFI does not appear in the primary channel), $E_{\min}\{y^2(n)\}$ goes to zero then equation (3) becomes

$$E\{\epsilon^2(n)\} = E\{s(n)^2\}, \quad (5)$$

and the filter turns off (in practical terms, the weighting coefficients are set to zero).

In summary, in an adaptive interference-canceling scheme, the system output is fed back into the adaptive filter, and then the adaptive filter adjusts itself to minimize the total system output power, updating the weighting coefficients as the need arises. The advantage of this system is that it works well in nonstationary conditions, i.e., when the

relative difference in the characteristics of the RFI in the primary and reference channels change with time. Refer to Widrow & Stearns (1985) and Haykin (1996) for a complete discussion of digital adaptive filter theory and techniques.

3. CHARACTERISTICS OF ADAPTIVE INTERFERENCE CANCELING IN THE FRAMEWORK OF THE WIENER SOLUTION

In stationary conditions, the adaptive system converges to the optimal Wiener filter. To minimize the output power, the algorithm needs to find the minimum of an error surface, defined in the next section. This surface is a multidimensional hyperboloid bowl with a unique minimum. Once the adaptive filter finds the set of coefficients corresponding to the bottom of the bowl, those coefficients are fixed and are not updated or changed (the Wiener solution). However, the adaptive version continuously tracks changing statistics between the reference and primary channels and adjusts the coefficients in real time. In a realistic, nonstationary environment, the bottom of the hyperboloid bowl is slowly moving around, not by large amounts but significantly, as the sides of the bowl change shape. The adaptive system finds the Wiener solution by locking onto the minimum point and then tracking the minimum as the bowl moves around in multidimensional space. Since the basis of adaptive interference canceling is the Wiener solution, we will start by examining the framework of the Wiener filter. In this section, we discuss the characteristics of a Wiener filter in the presence of random noise and formulate performance expectations. In § 4, we discuss methods of tracking and the

design of an adaptive interference-canceling receiver for radio astronomy.

3.1. *The Optimal Transfer Function: Finding the Bottom of the Bowl*

Since $i_R(n)$ is not an exact duplicate of $i_p(n)$, we process $i_R(n)$ with an adjustable-weight filter with coefficients $w(n)$ to produce $y(n)$, which is a close replica of $i_p(n)$. A schematic of the filter is shown in Figure 3; it is constructed using a tapped-delay line (or transversal filter) in the reference channel and a linear combiner. The unit delay z^{-1} is one sample-time (the unit delay is the generalized discrete Fourier transform, or z -transform, of the unit sequence delayed by one sample; see Oppenheim & Schaffer 1989).

The tapped-delay line is a finite impulse response (FIR) filter, used in nearly all digital interference-canceling applications because of its inherent stability. The taps are scaled by weighting coefficients and then summed to form the FIR filter whose output is $y(n)$. Finally, $y(n)$ is subtracted from the primary input to form the combiner output, $\epsilon(n)$. As described by equation (1), $\epsilon(n)$ is the difference between the primary output and the processed output of the reference channel. Note that no filtering is done in the primary channel.

In analogy to equation (1), let $I(n)$ be the vector of delayed versions of $i_R(n)$ and $W(n)$ be the vector of tap weights containing the set of weights $w(n)$. Then equation (1) becomes

$$\epsilon(n) = [s(n) + i_p(n)] - I^T(n) \cdot W(n), \quad (6)$$

where T indicates the transpose. The output power is $\epsilon^2(n)$. The filter finds $W(n)$ by minimizing the total output power, so $E\{\epsilon^2(n)\}$ can also be considered as an estimation of how well the system is working; in control system theory, $\epsilon(n)$ is called the error signal. $E\{\epsilon^2(n)\}$ is the error performance surface, which is a multidimensional quadratic hyperboloid and has a unique minimum. In stationary conditions, this minimum is fixed and is described by the optimum weight vector W_{opt} forming the Wiener filter response. The values of W_{opt} are found by setting the derivative of $E\{\epsilon^2(n)\}$ with respect to the weights equal to zero. The expectation value for a stationary process is equivalent to the autocorrelation

function (power spectrum), and so after matrix operations we arrive at

$$\phi_{RP}(n) = \sum_{l=-\infty}^{\infty} w_{opt,l} \phi_{RR}(n-l), \quad (7)$$

where ϕ_{RR} is the autocorrelation of the reference input and ϕ_{RP} is the cross-correlation function of the reference with the primary input (see Widrow & Stearns 1985 for a derivation). Taking the z -transform (the generalized discrete Fourier transform) of equation (7), we have

$$W_{opt}(z) = \Phi_{RP}(z)/\Phi_{RR}(z). \quad (8)$$

This result represents the unconstrained, noncausal Wiener solution (see, e.g., Oppenheim & Schaffer 1989 for details). However, any realizable physical system must be causal. In order for the performance to approach the ideal noncausal filter, a delay must be inserted in the primary input. This forces an equal delay in the response of the filter. The length of the delay is chosen to cause the peak of the impulse response to be centered along the tapped-delay line. This causal system can behave in a noncausal manner for a limited time frame, since the solution will depend on samples at $n-1$, n , and $n+1$. A real filter has a finite number of taps, but the more taps, the closer the impulse response will be to the ideal, infinitely long filter. The number of taps for a particular digital signal processor is a cost-performance trade-off.

3.2. *Random Noise and Propagation Path*

In this section we include the effects of random noise and propagation paths on the system output by rewriting W_{opt} in terms of the interference power spectra in each channel. A power spectrum of the interference will depend upon the noise and transfer function through the system. This discussion is for a Wiener solution but is equivalent for the adaptive system. In this paper, "noise" always refers to random noise unless otherwise specified.

Noise temperatures for the primary and reference receivers are the uncorrelated noise components $m_p(n)$ and $m_R(n)$, respectively. However, the interference in both channels

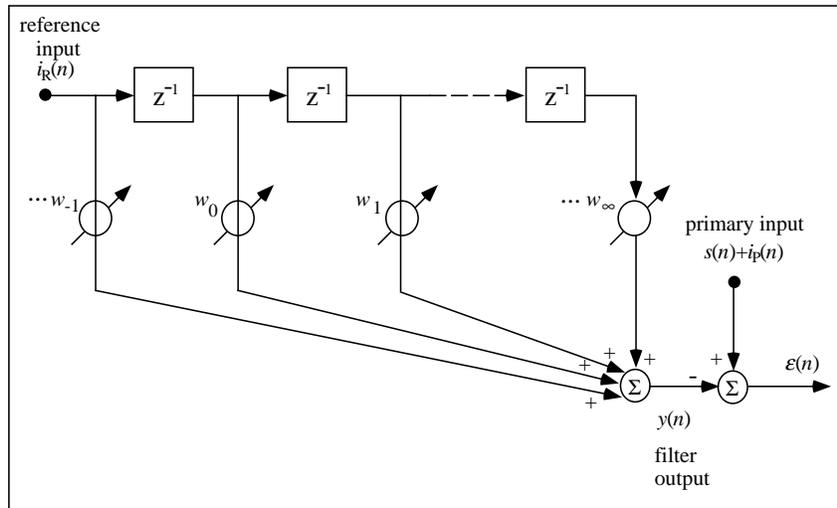


FIG. 3.—An ideal Wiener filter with an infinitely long tapped-delay line. The filter tap weights, w_i (or weighting coefficients), are adjusted to yield optimal filter performance for the case of stationary conditions (the Wiener solution). z^{-1} is one sample-time.

comes from the same source, so $i_p(n)$ and $i_R(n)$ are correlated with each other but are uncorrelated with either $m_p(n)$ or $m_R(n)$. Therefore, the reference input to the filter is $i_R(n) + m_R(n)$, and likewise, the interference in the primary input is $i_p(n) + m_p(n)$. Propagation paths through the system also affect the output. The interference signal that arrives at each channel is the original interference signal convolved with some impulse response function for each path, $h_R(n)$ and $h_p(n)$, for the reference and primary inputs, respectively. The transform of each, called the transfer function, describes the characteristics of each propagation path. Since $i_p(n)$ and $i_R(n)$ are correlated, we are only interested in the ratio of their transfer functions. So we can define $H(z)$ to be normalized to the transfer function in the primary channel and write the interference power spectra of the primary and reference channels in terms of each other:

$$\Phi_{i_R i_R}(z) = \Phi_{i_p i_p}(z) |H(z)|^2. \quad (9)$$

The autocorrelation function of the reference input becomes

$$\Phi_{RR}(z) = \Phi_{m_R m_R}(z) + \Phi_{i_p i_p}(z) |H(z)|^2, \quad (10)$$

where $\Phi_{m_R m_R}$ is the power spectrum of the *noise* in the reference channel and $\Phi_{i_p i_p}$ is the power spectrum of the *interference* in the primary channel. Since the cross-correlation of the reference and primary inputs depends only on the correlated components of each, we have

$$\Phi_{RP}(z) = \Phi_{i_p i_p}(z) H(z). \quad (11)$$

Substituting equations (10) and (11) into equation (8), we obtain an expression for the optimal transfer function that includes noise and propagation paths:

$$W_{\text{opt}}(z) = \frac{\Phi_{i_p i_p}(z) H(z)}{\Phi_{m_R m_R}(z) + \Phi_{i_p i_p}(z) |H(z)|^2}. \quad (12)$$

Note that $W_{\text{opt}}(z)$ is independent of both $\Phi_{m_p m_p}(z)$ and the power spectrum of the astronomical signal of interest.

3.3. Performance Expectations

The goal of any interference excision scheme is to recover the astronomical signal without distortion by the filter. To this end, the interference at the output of the canceler must be reduced down to or below the rms noise over the integration time needed for the science, and the averaged baseline noise should not be altered by the presence of the canceler. The canceler's performance depends on a number of factors, including: random noise in each channel, quantization noise, type of algorithm used to find the optimal weights, and, in the case of the adaptive system, tracking ability. One of the unique characteristics of adaptive interference cancellation is that the filtering process occurs in the reference channel, and so the astronomical signal coming through the primary channel is not distorted by the canceler. As a result, the filter is linear, and so the attenuation achieved by the filter will be entirely independent of the astronomical and interference signals in the primary channel. Linearity is preserved as long as the system is operated within the dynamic range set by the RF and IF amplifiers, quantization of the A/D converter, etc. We have used 12 bit converters in our prototype system, resulting in a dynamic range (see Ifeachor & Jervis 1993) of 72 dB.

In this section, we will describe the performance of an adaptive filter in a stationary environment as modeled by a Wiener filter; the tracking of the adaptive system in the case

of nonstationary conditions will be considered separately in § 4.

3.3.1. Attenuation of RFI

The canceler's performance over a given integration time is measured by how well the RFI signal is attenuated, i.e., whether the attenuation reached the rms noise, and how long one could integrate before the rms noise level would fall below the residual RFI. Quantities of interest are (1) theoretical interference attenuation A_T , (2) measured attenuation $A(\tau)$, where τ is the integration time, and (3) integration time for the rms noise to fall below the residual RFI τ_A . For an ideal filter $A(\tau) = A_T$, but the dynamic range of a real filter is limited by the resolution of its digital processor, so $A(\tau)$ will always be less than or equal to A_T , assuming that no saturation occurs at any point in the radiometer or canceler. Conditions at the telescope will not always push the system to the edge of possible performance; if the RFI is weak to start with, a low value of $A(\tau)$ might still result in attenuating RFI below the level of rms noise and yield a successful observation. It is important to note that an RFI signal is never completely excised; there is always some residual RFI even when the attenuated RFI is buried in the rms noise. As the integration time gets longer, the rms noise gets smaller, and eventually the residual RFI signal could be recovered. Using A_T we can calculate the longest integration time before the residual RFI ruins the astronomical observation. These quantities, A_T , $A(\tau)$, and τ_A are formulated below.

The theoretical attenuation is the ratio of the interference power spectra in the system output to primary input:

$$A_T(z) = \Phi_{i_{\text{output}}}(z) / \Phi_{i_p i_p}(z). \quad (13)$$

To express A_T in terms of measurable quantities, it can be written as a function of the interference-to-noise ratios in the primary and reference channels, IN_P and IN_R , respectively, where

$$\text{IN}_P(z) \simeq \Phi_{i_p i_p}(z) / \Phi_{m_p m_p}(z), \quad (14)$$

$$\text{IN}_R(z) \simeq \Phi_{i_p i_p}(z) |H(z)|^2 / \Phi_{m_R m_R}(z). \quad (15)$$

Substituting equations (14) and (15) into the optimal transfer function, equation (8) becomes

$$W_{\text{opt}}(z) = \frac{\text{IN}_R(z)}{H(z)[\text{IN}_R(z) + 1]}, \quad (16)$$

and after some manipulation, these expressions combine to give

$$A_T(z) = [\text{IN}_R(z) + 1]^2. \quad (17)$$

(Note that all these quantities, A_T , $A(\tau)$, τ_A , IN_R , and IN_P in the discussion that follows are explicit functions of z .) A plot of A_T versus IN_R is shown in Figure 4 (*top*). Both axes are in units of logarithmic relative power, decibels, defined as $10 \log (P_1/P_2) = 0$ dB, if $P_1 = P_2$. Equation (17) gives the somewhat nonintuitive result that the theoretical attenuation depends on IN_R and is independent of IN_P . Yet, as mentioned above, this is consistent with the system's linearity; the filtering process occurs in the reference channel, and so the attenuation in the system output is independent of both the astronomical signal and interference in the primary channel. Also, since A_T varies as IN_R^2 , moderately good interference-to-noise in the reference

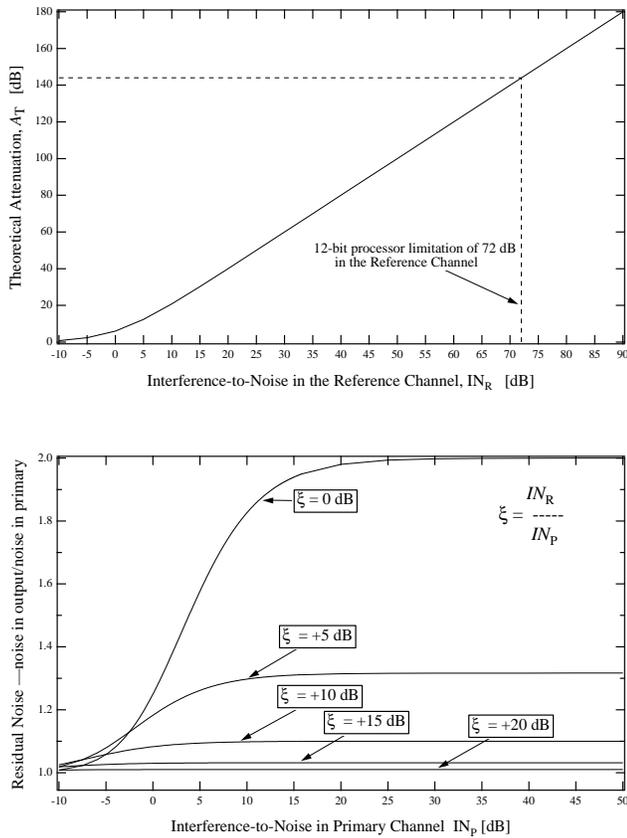


FIG. 4.—*Top*: Theoretical interference attenuation A_T as a function of the interference-to-noise in the reference channel IN_R . Both axes are in units of logarithmic relative power, dB, defined to be $10 \log(P_1/P_2)$. Since A_T goes as IN_R^2 , moderately good interference-to-noise in the reference channel produces significant attenuation of interference in the primary channel. The maximum achievable attenuation depends on the resolution of the digital processor in the reference channel. The dynamic range for the 12 bit processor is 72 dB. *Bottom*: The residual noise ratio RNR as a function of the relative interference-to-noise in the primary and reference channels. RNR is a measure of the noise introduced into the system output by the reference channel. Several curves are shown, each for a different ratio of IN_R to IN_p , such that $\xi = IN_R/IN_p$. ξ is given in dB, so that if $\xi = 0$ dB, the interference-to-noise is the same in the reference and primary channels.

channel should produce significant attenuation of interference in the system output. For example, our prototype system has a maximum dynamic range in the reference channel of 72 dB; however, with IN_R of 72 dB, Figure 4 (*top*) shows an achievable attenuation of 144 dB. This is achievable in our prototype system since calculations within the filter are represented by 27 bits (giving a limit of 162 dB), yet our output D/A converter places an upper limit of 72 dB. This means that although the attenuation achieved is much better than 72 dB, it is only measurable in the system output to that limit.

A useful way to measure the achieved attenuation in a given observation is

$$A(\tau) = I_{\text{unfiltered}}/I_{\text{filtered}}, \quad (18)$$

where $I_{\text{unfiltered}}$ and I_{filtered} are the interference peaks in the system output before and after filtering, respectively. If there is no residual interference peak, i.e., the RFI is attenuated down to or below the noise, then we take I_{filtered} to be the peak-to-peak noise σ_p in the baseline.

The residual of the interference can be related to the radiometer equation by comparing the power spectra of the interference and average noise at the system output:

$$\begin{aligned} IN_{\text{output}}(z) &= \frac{\Phi_{iip}(z)}{\Phi_{mpmp}(z)A_T(z)} = \frac{IN_p(z)}{A_T(z)} \\ &= \frac{\sigma_i}{T_{\text{sys}}} = \frac{1}{\sqrt{B\tau}}. \end{aligned} \quad (19)$$

IN_{output} is the interference-to-noise of the output, where B is the bandwidth resolution in hertz, τ is the integration time, T_{sys} is the system temperature, and σ_i is the rms noise. We are assuming here that the canceler itself does not introduce additional noise. We can use this expression for IN_{output} to find the maximum integration time before the rms noise would fall below the residual RFI and ruin the observation. If we define τ_A to be the integration time required for RFI to appear as a 3σ residual above the noise, then equation (19) gives

$$\tau_A = \frac{9}{B} \left(\frac{A_T}{IN_p} \right)^2, \quad (20)$$

which is valid for IN_R up to the digitization limit of 72 dB for our converters. For example, this expression shows that for moderately good IN_R and IN_p of 30 and 20 dB, respectively, A_T is 60 dB, and so the integration time could be as long as 3 weeks before a 3σ residual would pop up above the rms noise.

3.3.2. Random Noise Contribution by the Reference Channel

Although noise in the adaptive canceler will not distort the astronomical signal while attenuating the interference significantly, the contribution of noise from the reference channel is an important consideration. Noise added by the canceler is a result of three factors: (1) quantization noise caused by digitization, (2) residual noise within the bandwidth of the interference, and (3) residual noise outside the interference bandpass caused by the tapped-delay line.

Quantization noise is a result of the digitization of the analog signal and is assumed to be uniformly distributed over the quantization step size, leading to a signal-to-quantization noise power ratio of approximately 74 dB for the 12 bit digital processor of our prototype system. The quantization noise contribution can be extremely small with appropriate signal adjustments, that is, gains in the reference and primary channels can be adjusted so that T_{sys} gives a noise floor significantly larger than the quantum noise floor.

Noise in the system output introduced by the reference channel is a more critical issue. Random noise in the reference channel is never zero, and so incorporating a reference channel necessarily injects noise into the output spectrum. This noise will have structure in the frequency domain. A measure of this effect is the residual noise ratio (RNR), defined as the ratio of the noise power spectra at the system output ($\Phi_{mm_{\text{output}}}$) to that in the primary channel input (Φ_{mpmp}):

$$\begin{aligned} \text{RNR}(z) &= \frac{\Phi_{mm_{\text{output}}}(z)}{\Phi_{mpmp}(z)} \\ &= \frac{\Phi_{mpmp}(z) + \Phi_{mrmr}(z)W_{\text{opt}}^2(z)}{\Phi_{mpmp}(z)}. \end{aligned} \quad (21)$$

This expression, after some algebra, can be put in terms of interference-to-noise in the primary and reference channels, which are measurable quantities, and so equation (21) becomes

$$\text{RNR}(z) = \frac{\text{IN}_R(z)\text{IN}_P(z)}{[\text{IN}_R(z) + 1]^2} + 1. \quad (22)$$

For good performance, we want this residual noise ratio to be as close to 1 as possible; that is, we want the noise spectrum in the system output to be no greater than the noise spectrum in the primary input channel. RNR is unity if there is no reference channel in the system. It is useful to consider equation (22) in terms of the *relative* interference-to-noise in the reference and primary channels; in general, we expect better performance when $\text{IN}_R > \text{IN}_P$. We want to know how much noise is introduced into the system output for different ratios of IN_R to IN_P . If we write IN_R in terms of IN_P , so that $\text{IN}_R(z) = \xi \text{IN}_P(z)$, then equation (22) becomes

$$\text{RNR}(z) = \frac{\xi \text{IN}_P^2(z)}{[\xi \text{IN}_P(z) + 1]^2} + 1. \quad (23)$$

Figure 4 (*bottom*) shows the residual noise as a function of IN_P for different values of ξ . If the interference-to-noise in the reference and primary channels are equal ($\xi = 0$ dB), the noise in the system output (at the frequencies where the interference signal was located) will be almost twice as high as in the primary input before filtering. This is a result of a noisy interference signal in the reference channel being subtracted from the interference signal in the primary; since the noise is uncorrelated, this operation adds significant noise power to the system output. Therefore, the higher the interference-to-noise in the reference channel relative to that in the primary, the lower the residual noise in the system output. In the case of a radio telescope, the primary channel receives RFI in the relatively weak sidelobes of the beam. The low-gain antenna for the reference channel is pointed toward the horizon in the direction of the RFI, so interference-to-noise in the reference channel can easily be higher than in the primary channel. The curves of Figure 4 indicate that over a wide range of IN_P (especially for $\xi \geq 10$ dB) the injected noise is nearly constant. This implies that even though the interference signal level might fluctuate in the primary channel due to the telescope's slewing or to propagation effects, the amount of noise injected in place of the interference will be constant.

Another noise component introduced by the reference channel occurs outside the bandwidth of the interference and can result in baseline ripple. Ideally, the digital filter in the reference channel should have a very sharp frequency response, such that it adjusts to the bandwidth of the interference but blocks reference channel noise outside this bandwidth. However, the sharpness of the digital filter is proportional to the number of taps. If the filter contains too few taps, then excess noise from outside the interference bandwidth will enter the primary channel, causing a ripple in the baseline power. Therefore, as large a number of taps as possible should be used to minimize this effect. However, to choose the optimal number of taps, many factors must be considered. For a given filter passband ripple, stop-band attenuation, and transition width between passband and stopband, an estimate of the number of taps required to meet the specification can be obtained from optimal FIR filter theory (Mintzer & Liu 1979).

In summary, the adaptive canceler will not distort the spectrum of the astronomical signal, yet it will provide a high degree of interference attenuation. Noise injected by the reference channel can be minimized if $\text{IN}_R > \text{IN}_P$ and a large number of filter taps are used.

4. ADAPTIVE INTERFERENCE CANCELING IN A NONSTATIONARY ENVIRONMENT

The results of the previous section for a Wiener filter can be applied to the adaptive filter system. In addition to the considerations for a stationary filter, an adaptive process introduces other sources of error and noise, specifically from the tracking capability and multiple reference channels.

4.1. The Least Mean Square Algorithm and Tracking Concerns and Constraints

The basic algorithm used to find the minimum of the error surface for the Wiener solution is also used to find the minimum in the case of an adaptive system, with the added complication of tracking the minimum as it changes. The least mean square (LMS) algorithm uses an estimate of the error surface gradient that is closely tied with the structure of the tapped-delay line and requires a minimal amount of computing. There are other algorithms that offer improvements over LMS that would increase the performance of an adaptive system; examples include recursive least squares (Haykin 1996), which uses off-line gradient estimations, and higher-order statistics (Shin & Nikias 1994), which is computationally complex. For our prototype receiver, we have implemented the LMS algorithm for its computational simplicity.

As in equation (6), $\mathbf{I}(n)$ is the vector of delayed versions of $i_R(n)$, and $\mathbf{W}(n)$ is the vector of tap weights containing the coefficients $w(n)$. For the Wiener solution and for each iteration in the adaptive system, the gradient of the error surface can be estimated from

$$\nabla \epsilon^2(n) = \begin{bmatrix} \frac{\partial[\epsilon^2(n)]}{\partial w_0(n)} \\ \vdots \\ \frac{\partial[\epsilon^2(n)]}{\partial w_L(n)} \end{bmatrix} = 2\epsilon(n) \begin{bmatrix} \frac{\partial[\epsilon(n)]}{\partial w_0(n)} \\ \vdots \\ \frac{\partial[\epsilon(n)]}{\partial w_L(n)} \end{bmatrix} = -2\epsilon(n)\mathbf{I}(n), \quad (24)$$

where L is the number of filter tap weights, defining a direction in error space. By starting with this estimate of the gradient and using the method of gradients (see Widrow & Stearns 1985):

$$\mathbf{W}(n+1) = \mathbf{W}(n) - \mu \nabla \epsilon^2(n) = \mathbf{W}(n) + 2\mu \epsilon(n)\mathbf{I}(n). \quad (25)$$

$\mathbf{W}(n+1)$ is found by tweaking $\mathbf{W}(n)$. The parameter μ is the gain constant and is related to the step size in the search for minimum as the system tracks. The smaller the step size, the longer it takes to find the bottom of the error surface bowl. This is especially important in an adaptive filter, since the tracking time must be able to keep up with the timescale of the changing statistics, i.e., the ongoing movement of the bottom of the bowl. Speed and stability of adaptation, as well as noise in the weight vector solution, are determined by the size of μ ; the smaller μ is, the smaller the error is in $\mathbf{W}(n)$, but the longer it takes for the solution to converge. A

compromise between speed and introduced error is made in choosing μ . The weight vector will converge to an optimal solution when

$$0 < \mu < \frac{1}{(L + 1)E\{i_R^2\}} \quad (26)$$

(see Widrow & Stearns 1985). Note that the optimum value of μ is a function of the interference power in the reference channel. The optimal value of μ is a trade-off between better adaptability and time for convergence. In our prototype system, the value of μ is manually adjustable. For a future system, we plan to make μ self-adjustable, that is, capable of responding to the interference power in the reference channel (S. Wilson 1996, private communication).

4.2. Adaptive Interference Canceler with Multiple Reference Channels

In theory, any number of reference channels can be used to cancel interference in the primary channel. Multiple reference channels are necessary if there is more than one interferer in the passband, since a single reference channel does not have the necessary degrees of freedom to eliminate more than one. Additional degrees of freedom are required whenever (1) there are several uncorrelated interference signals in the filter bandpass, (2) a single interference signal encounters severe multipath propagation and appears as several signals, or (3) the spatial polarization of a single interference signal differs significantly between the primary and reference inputs. With two or more uncorrelated sources of interference, the synthesis of the transfer functions, and hence the set of optimal weight vectors, becomes more complicated. Not only are there transfer functions that describe the propagation paths from the sources through the primary input, but there are other transfer functions that represent propagation paths through all the reference inputs with allowance for cross-coupling.

Additional degrees of freedom can be achieved by increasing the number of reference channels processed by the canceler. In the case of multiple interference signals in the filter passband (in our case, more than one radio station), uncompromised performance occurs when the number of reference channels is equal to or greater than the number of interferers. This is also relevant under severe multipath conditions where time delays cause the same interference source to appear multiple times (similar to ghost images on a television receiver). Note that a one-to-one correspondence between an interferer and a given reference channel is not a requirement; the adaptive filter will synthesize the individual transfer functions as linear combinations of the reference channel inputs.

Another degree of freedom is needed when the spatial polarization of the reference antenna differs from that of the sidelobe of the telescope antenna in the direction of the interference source, as when the telescope slews across the sky during any given observation. As an example, the RFI in our bandpass is from broadcast stations in the FM band that transmit with circular polarization. Yet propagation effects between the transmitter and the telescope can cause the polarization of the arriving RFI to change from its initial state or even to fluctuate randomly with time. If the polarization response in the reference and primary antennas match that of the RFI, then there is no problem and another degree of freedom is not needed. However, if there

is a spatial polarization difference between the reference and primary antennas, then one or the other will see a stronger signal (even a null at the reference antenna, which would result in the adaptive filter turning off) and possibly a phase delay. If the difference in antenna orientations causes IN_p to be higher than IN_R , the attenuation of RFI in the output will be less than optimal. It is unlikely that a single reference antenna will always be oriented to match the telescope sidelobe polarization, especially as the telescope slews. The addition of a second reference antenna that responds to the orthogonal polarization sense with respect to the first reference antenna should be adequate to achieve optimal attenuation of RFI. The transfer function is then synthesized as a linear combination of the signals from the two orthogonal reference inputs.

4.3. Practical Considerations for an Adaptive-canceling System on Radio Telescopes

Adaptive canceling shows promise as an effective means to attenuate interference in both single-dish and interferometer radio telescope systems. However, there are three basic requirements of the system and the RFI environment for success with adaptive canceling: (1) the receiver must always operate in the linear regime, (2) the adaptive convergence time must be finite, on the order of a few seconds, and (3) interference-to-noise in the reference channel must be greater than that in primary channel.

A linear operating regime ensures that the RFI will not overload the receiver front end, since distorted interference cannot be removed through a linear adaptive filter system. A finite adaptive convergence time places practical limits on the type of RFI that can be canceled effectively by this method. Signals of the continuous-wave variety such as from broadcast stations or satellite downlinks (regardless of the modulation type) and moderately long duration signals (greater than a few seconds) from personal communication systems can be attenuated effectively since the canceler has the necessary time required to lock on and adapt. However, short bursts of interference from systems such as aviation or shipboard radar make canceling difficult without additional processing to "remember" the filter parameters from burst to burst yet turn the filter off rapidly between bursts to eliminate unwanted noise injection. Therefore, quasi-random frequency-hopping signals from certain spread-spectrum systems simply cannot be canceled using this adaptive design.

Maintaining a greater interference-to-noise in the reference channel than that in the primary channel places limits on where the interference source is located with respect to the main beam of the telescope. For example, if the main beam happens to be pointing directly at a satellite producing RFI, it will be impossible for the reference channel IN_R to be greater than that in the primary channel, since the gain of the main beam of the telescope will be greater than the gain of the reference antenna. In contrast, if it is the sidelobes that pick up the interference, IN_R can easily be greater than IN_p since the gain of the telescope in the direction of the interference can be quite low. Note that when the interference source is moving, e.g., LEO satellites such as the *Iridium* series, the reference antenna must track the satellite across the sky to achieve good IN_R . Finally, in the case of interference arriving at the telescope from an over-the-horizon source, the adaptive filter performance is reduced as the telescope elevation is decreased in the direction of the

interference, but the larger the telescope, the greater the tolerable elevation angle.

Even with these practical restrictions, however, adaptive canceling can greatly improve observing in many current and future situations where RFI poses serious problems for radio astronomy.

5. THE ADAPTIVE INTERFERENCE-CANCELING RECEIVER: THE PROTOTYPE

Ultimately, we hope to build adaptive interference cancelers for bands of interest in radio astronomy that have serious RFI problems, e.g., the *L* band. As a first step toward this goal, we have built a prototype adaptive canceler with a single reference channel, which we will use with a receiver in the FM band. Our next effort will be to complete the system with four pairs of reference channels to test on the 140 foot (43 m) radio telescope at Green Bank Observatory. Here we report the results of our single reference channel prototype tested under controlled laboratory conditions.

5.1. Design

The primary and reference inputs to this canceler are centered at 40 MHz and can be used with a variety of microwave receiver front ends that have a compatible IF bandpass. All of our laboratory measurements were performed by injecting controlled signals directly into the IF band. Initial system tests were performed using a front end designed for 100 MHz. The canceler design and measurements are described below.

5.1.1. The Adaptive Canceler

A block diagram of the prototype adaptive canceler is shown in Figure 5. The primary and reference inputs, both centered at 40 MHz and assumed band-limited to 500 kHz, are down-converted to baseband using a common local oscillator. At baseband, both signals are prefiltered, amplified, and then digitized using a 12 bit A/D converter sampling at a 4 MHz rate. To achieve the Nyquist-sampled 0–1 MHz baseband, the A/D outputs are decimated by 2 (i.e., every other sample is used). The primary channel input is directed to the summing junction. The reference channel input is directed to a nine-tap FIR adjustable-weight filter whose coefficients are determined by the microprocessor that implements the LMS algorithm. The filter output is directed to the summing junction, where it is subtracted from the primary data and forms the system output (which is also the error signal input to the LMS algorithm). The microprocessor, operating at 22 MHz, is capable of handling up to eight reference channels for future expansion.

5.1.2. The 100 MHz Front End

We chose the FM band for our first effort to build an adaptive filter receiver for radio astronomy mainly because the characteristics of the broadcast signals are well known: (1) each station occupies a relatively narrow bandwidth (<100 kHz at the defined level of 3 dB), (2) each station is well characterized in frequency (spaced at 200 kHz intervals), (3) each station is in continuous operation at a fixed power level, (4) most stations have a known transmit polarization, (5) the filter bandpass includes several FM channels to gauge the effectiveness of multi-reference

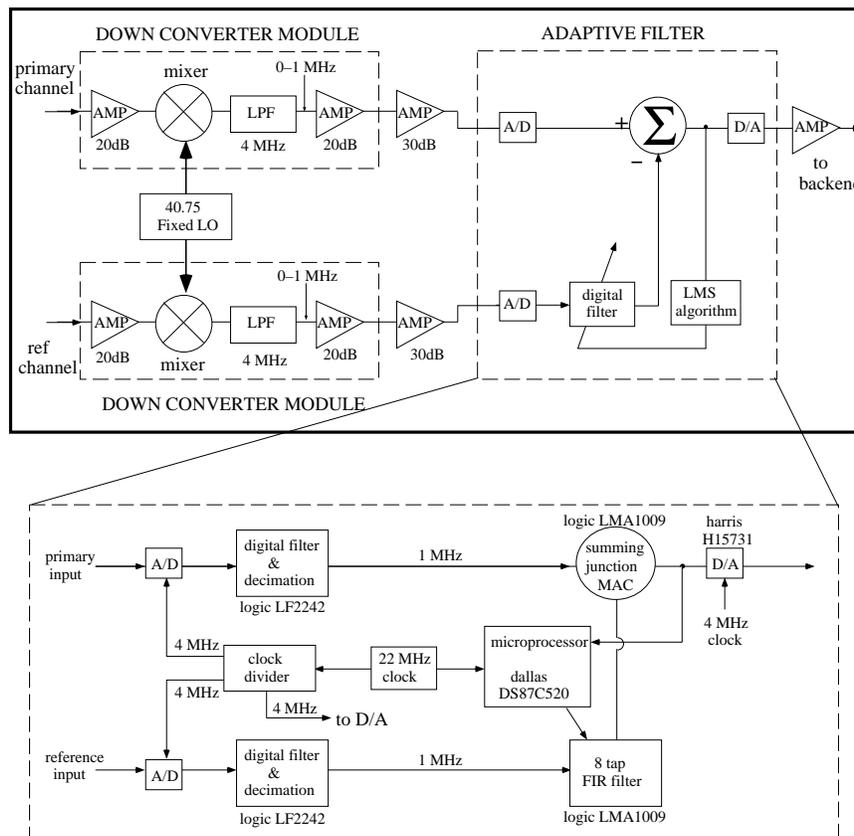


FIG. 5.—Block diagram of the intermediate-frequency (IF) downconverter and adaptive filter for a system with a single reference channel

channel system, and (6) each FM channel can have multiple interference signals on the same frequency. Green Bank Observatory is centrally located in the radio-quiet zone, and although broadcast stations are spaced 200 kHz apart, it is possible to receive multiple stations per channel at this location. Cautiously, we expect four radio stations broadcasting in our passband near 100 MHz. Although our prototype system currently has only one reference channel, there will be eight in the final system, one for each polarization for the four interferers.

A block diagram of the RF section is shown in Figure 6. The primary channel input is half of a cross dipole feed located at the prime focus of the radio telescope. Four pairs of orthogonal Yagi antennas will be the inputs to the eight reference channels (only two are shown in the figure). Each of the five RF inputs have a room temperature, low-noise amplifier ($T = 30$ K), followed by a bandpass filter and additional amplification. The inputs are translated to the 40 MHz IF bands using mixers that share a common local oscillator, and they are then directed to the adaptive canceler.

5.2. Bench Tests

We have completed a first-phase performance characterization of our prototype receiver in the laboratory. These

initial tests were performed at NRAO Headquarters in Charlottesville, Virginia, in the Central Development Laboratory. The goal was to measure the dynamic range of the adaptive filter and gauge the filter's performance under stationary and nonstationary conditions.

Figure 7 shows a block diagram of the bench test configuration. We created an interference signal using a signal generator with the option of single-tone and/or random frequency modulation. The output of the signal generator was split and directed to the reference and primary inputs. Both inputs were prefiltered using a 500 kHz bandpass filter centered at 40 MHz, with random noise power from two independent noise sources coupled into each input. In some experiments, we injected a test source into the primary input to evaluate the canceler's ability to recover the test signal in the system output. The test source was used to simulate an astronomical signal, that is, a signal that does not vary in strength or frequency over the course of the observations. To simulate nonstationary conditions, we used a voltage-controlled amplitude and phase shifter driven by a sine wave generator that varied the amplitude and phase of the primary input interference signal relative to that present at the reference input. The output of the adaptive filter was directed to a Hewlett-Packard spectrum analyzer for processing. The bandwidth resolution was 1 kHz and integra-

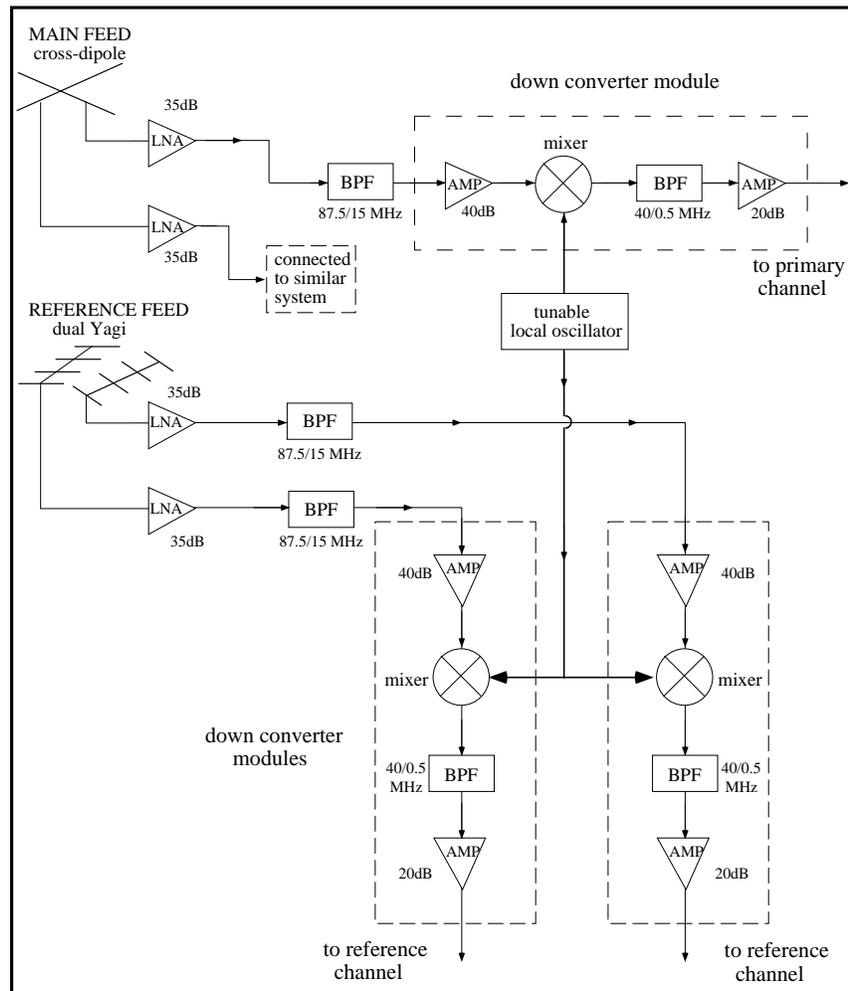


FIG. 6.—Block diagram of the radio-frequency (RF) and intermediate-frequency (IF) downconverter for a system with a two reference channels. The reference feeds are Yagis. In this diagram, both polarizations of the primary input are power-combined.

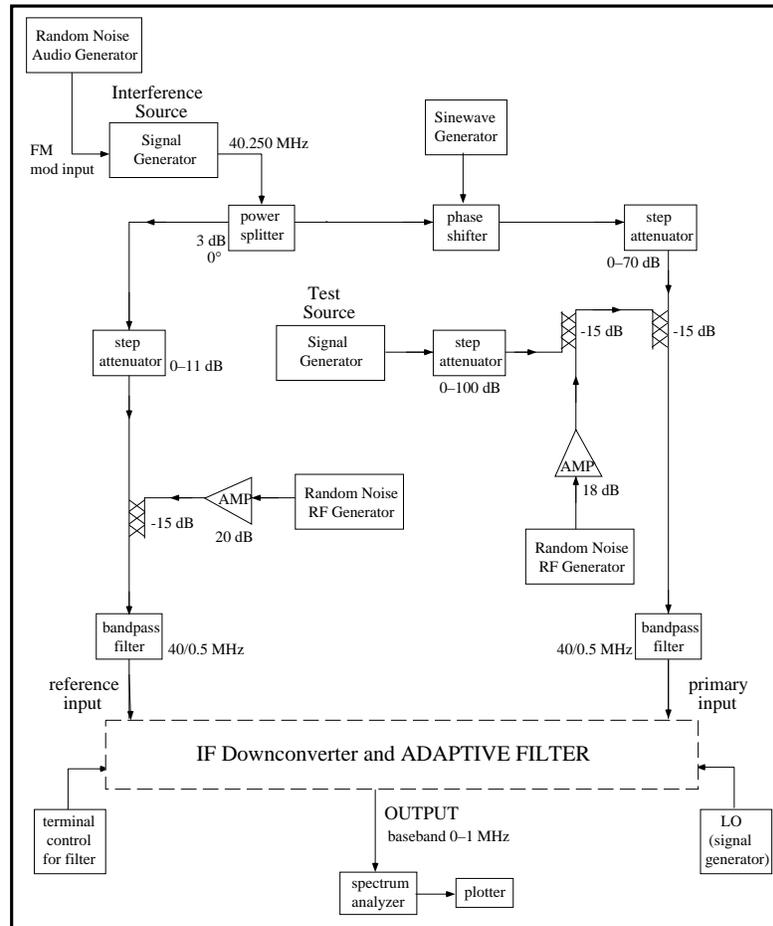


FIG. 7.—Block diagram of the experimental setup to test the prototype in the laboratory. The test source, imitating an astronomical signal, and the interference signal are produced by a separate signal generators. The IF downconverter and adaptive filter box referred to in this figure are shown in detail in Fig. 5.

tion time was typically 30 s. There was no capability to measure rms noise, so we quote the peak-to-peak noise in microwatts for each test.

Initial tests confirmed system linearity. The first experiments described below characterize the system's dynamic range under stationary conditions, as well as its ability to recover a weak test signal buried in the interference. The second set of experiments tests the adaptive tracking capability in nonstationary conditions. In these experiments the term "deviation of modulation" refers to the spread of the RFI signal on either side of the central frequency, and the term "modulation function" designates the band-limited signal that carries the RFI about the central frequency. Modulation signals can be random (Gaussian) and/or single-tone sine wave.

5.2.1. Performance for Stationary Conditions

To establish performance at or near maximum dynamic range, the first tests address the ability of the adaptive filter to attenuate very strong interference in the primary channel for different values of $\xi = IN_R/IN_P$, without tracking (stationary solution) and without the presence of a simulated astronomical signal. The results are shown in Table 1. The integration time for all measurements was 30 s; once the canceler was activated, it took less than 1 s for the canceler to lock up on the signal and typically less than 20 s to extinguish the RFI below the rms noise. The RFI is

successfully attenuated below the noise for all the experiments in Table 1. The RFI in each of these experiments had a sine wave modulation function of 500 Hz with a deviation of 5 kHz. The measured interference attenuation $A(\tau)$ in all cases goes up to or near the theoretical limit of the 12 bit digitization (72 dB). Based on the error in our measurements of the power, we find the error in $A(\tau)$ is ± 2 dB.

Figure 8 shows spectra from column (2) in Table 1. For all spectra in this and subsequent figures, the data were recorded in volts. The right axis shows the linear scale in volts, and the left axis shows the quantity of interest, power, which is directly proportional to the antenna temperature; note that the left axis is not linear since power is proportional to the square of the voltage. The top plot in Figure 8 is the RFI at the system output before the canceler is activated (this spectrum shows a 1 s integration; after 30 s the peak of the RFI integrated down to $0.071 \mu\text{W}$, which is the value we use in the table). The filtered spectrum appears below. Note that the RFI in the top plot is a factor of 10^3 times stronger than the strongest FM station shown in Figure 1 (after correcting for the gain and bandwidth resolution of the systems).

The next set of experiments tested the ability of the filter to recover a weak test signal buried in relatively strong RFI. The results for four trials with different values of ξ are in Table 2. Figure 9 shows spectra for the first trial in Table 2. The RFI signal was simulated with a deviation of 5 kHz and

TABLE 1
HIGH DYNAMIC RANGE TESTS: STATIONARY CONDITIONS

Parameter	Trial 1	Trial 2	Trial 3	Trial 4
$\xi = IN_R/IN_P$ (dB)	10	10	20	20
IN_P (dB)	62	62	52	48
Deviation of modulation (kHz)	5	5	5	5
Modulation function (1 frequency) (kHz)	5	5	4	5
RFI before filtering ^a ($10^{-6} \mu\text{W}$)	157000	71000	9380	12480
Output peak-to-peak noise ($10^{-6} \mu\text{W}$)	0.008	0.003	0.008	0.01
$A(T)^b$ (dB)	72	72	61	61

NOTE.—The errors are $\pm 2\%$ in IN_P and the RFI peak, ± 0.002 in the output peak-to-peak noise, and ± 2 dB in the attenuation $A(\tau)$.

^a Note that there was no measurable residual RFI after filtering.

^b $A(\tau)$ is the measured attenuation, achieved, to a limit here of 72 dB. It is defined to be the peak of the RF signal before filtering divided by the peak-to-peak noise after filtering. It is the measured attenuation.

a random modulation function band-limited to 20 kHz. The middle panel shows a $\sim 5 \sigma$ test signal before the RFI was turned on, and the canceler’s successful recovery of the test signal is in the bottom panel. The RFI signal disappears within the noise, and the test signal is recovered. Figure 9 shows that the random noise in the spectrum before and after filtering has not changed significantly, so the presence of the reference channel does not seriously degrade the output. However, a slight structure to the baseline is introduced by the reference channel. Note that a second, very weak RFI signal centered at -13 kHz (coming from some unidentified RFI source in the room) appears in the middle panel and is also successfully removed by the canceler.

Figures 10, 11, and 12 show spectra from the experiments of the last three columns of Table 2. The top panel of each

figure displays the RFI alone coming through the system output. The test source alone in the system output is shown in the second panel. The canceler is not turned on for either. The third panel in each figure shows the “OFF position” obtained by turning off the test source and turning on the RFI and canceler. This approximates beam switching off-source in astronomical observations. The OFF position spectra in Figures 10 and 12 show no measurable sign of residual RFI, and the bottom panels (with the OFF posi-

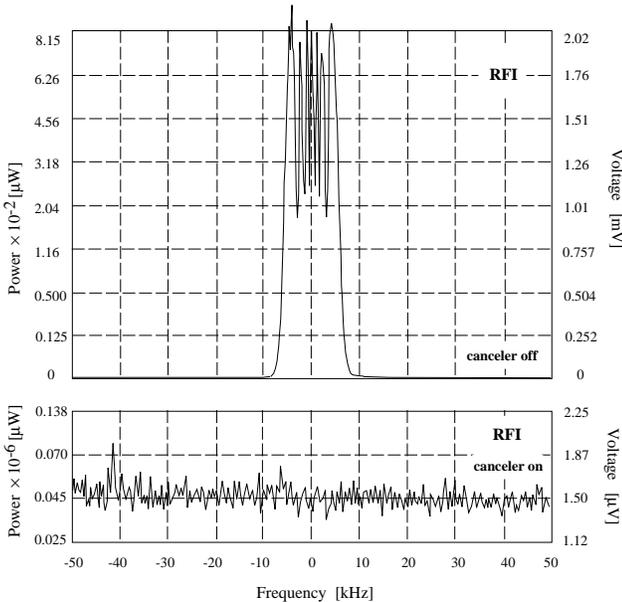


FIG. 8.—High dynamic range test of the adaptive-canceling system under stationary conditions. The data are shown in col. (2) of Table 1. The spectra in this and subsequent figures were recorded in volts. The right axis shows the linear scale in volts, and the left axis is power. Note that the left axis is not linear since power is proportional to the square of the voltage. The top panel shows the RFI at the system output before the adaptive canceler is activated (this spectrum shows a 1 s integration; with a 30 s, the peak of the RFI integrated down to $0.071 \mu\text{W}$, which is the value we use in col. [2] of Table 1). The bottom panel shows the spectrum after activating the adaptive canceler and integrating for 30 s.

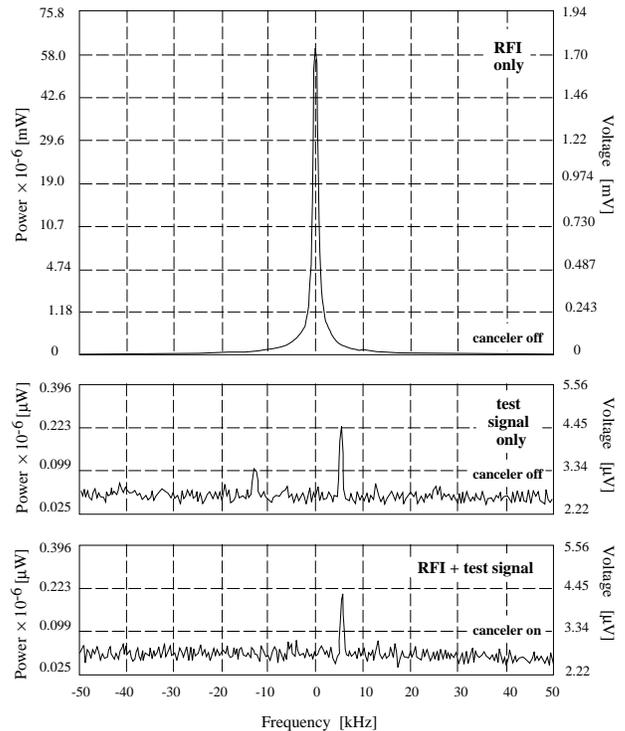


FIG. 9.—Simulated astronomical signal $s(n)$ (or “test signal”) buried in strong RFI with stationary conditions. The top panel shows the RFI alone in the system output before the filter is activated. In the middle panel, the test signal alone is shown (RFI and canceler are off). Note that the test signal is offset by about 5 kHz from the central frequency. The spectrum in the bottom panel results when the RFI and test signal are on and the canceler activated. The test signal that was buried in the RFI is recovered. Note that in this figure, the baseline has not been subtracted. Also note that at -13 kHz a very weak RFI signal leaked into the primary channel (coming from some extraneous source in the laboratory); the canceler was able to excise this as well as the much stronger RFI signal. The data are displayed in col. (1) of Table 2. See Fig. 8 for an explanation of the axes.

TABLE 2
RECOVERY OF A TEST SOURCE BURIED IN THE RFI: STATIONARY CONDITIONS

Parameter	Trial 1	Trial 2	Trial 3	Trial 4
$\xi = IN_R/IN_P$ (dB).....	10	20	15	15
IN_P (dB)	62	20	11	16
Deviation of modulation (kHz)	5	5	15	30
Modulation function ^a (random) (kHz)	5	4	4	4
Peak of test signal (10^{-6} μ W).....	0.22	0.08	0.12	0.08
RFI signal before filtering (10^{-6} μ W)	58000	48.0	6.48	20.4
RFI signal after filtering (10^{-6} μ W).....	0.09	...
Output peak-to-peak noise (10^{-6} μ W).....	0.01	0.001	0.001 ^b	0.002
$A(T)^c$ (dB)	68	47	38 (19) ^d	40

NOTE.—The errors are $\pm 2\%$ in IN_P and the RFI peak, ± 0.002 in the output peak-to-peak noise, and ± 2 dB in the attenuation $A(\tau)$.

^a The frequency function for trial 3 was random plus single tone.

^b This is the peak-to-peak noise measured in the filtered spectrum after subtracting the OFF position.

^c $A(\tau)$ as in Table 1.

^d The numbers in parentheses are the results in the filtered spectrum before subtraction by the OFF position. The other values are the results after subtraction by the OFF position (see Fig. 12).

tion subtracted) show a good recovery of the test signal with no significant increase in rms noise. Even with a 30 kHz deviation for the RFI in Figure 12, the canceler performs well. The experiment shown in Figure 11, however, was

performed with a random plus single-tone modulation function and a 15 kHz deviation. As can be seen from the OFF position spectrum in the third panel, the canceler was not able to extinguish the RFI to the level of the rms noise. Note that the residual RFI is double peaked. It is as if the canceler is seeing two separate RFI signals and cannot eliminate both at once. Yet when the OFF position is subtracted from the ON source spectrum, the test signal is recovered (bottom panel). A single-tone modulation is atypical in the

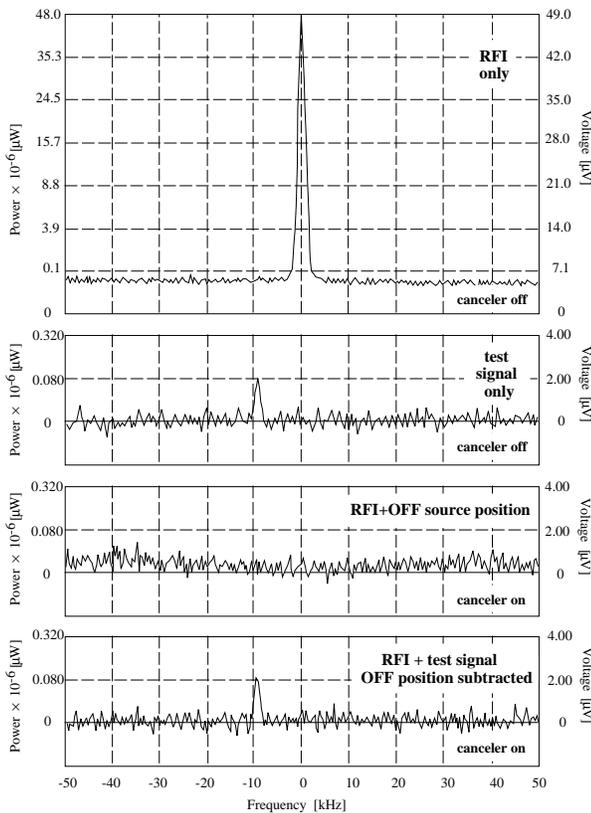


FIG. 10.—Results of adaptive cancellation with a test signal buried in strong RFI, random modulation, “OFF source position” subtraction, and with stationary conditions. The experimental procedure is the same for Figs. 11 and 12. The first panel shows the RFI alone in the system output before the canceler has been activated. The second panel shows the test signal before the canceler is activated. To obtain an “OFF source position,” we turned off the test signal generator but kept the canceler on. This is equivalent to beam switching during an astronomical observation, and the resulting spectrum is shown in the third panel. The bottom panel shows the ultimate result, with test source, RFI, and canceler operating, and subtracted by OFF position spectrum. See col. (2) of Table 2. See Fig. 8 for an explanation of the axes.

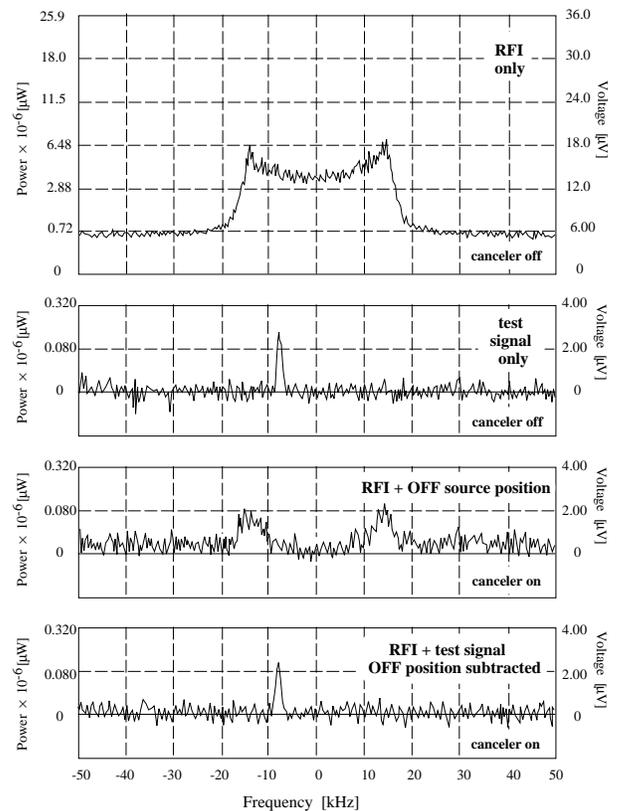


FIG. 11.—Results of adaptive cancellation with a test signal buried in a broad (15 kHz) RFI signal, random and single-tone modulation function, “OFF source position” subtraction, and with stationary conditions. The legend for Fig. 10 describes the basic experimental setup. See col. (3) of Table 2 for measurements.

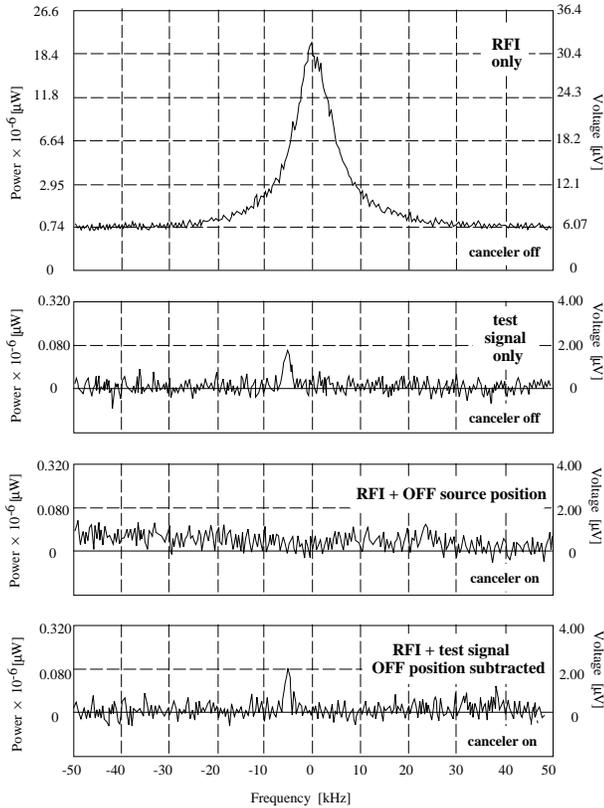


FIG. 12.—Results of adaptive cancellation with a test signal buried in a very broad (30 kHz) RFI signal, random modulation function, “OFF source position” subtraction, and with stationary conditions. The legend for Fig. 10 describes the basic experimental setup. See col. (3) of Table 2 for measurements.

real world, but this setup pushes the system to its limit and these results will be useful when the prototype is being tested at the telescope.

5.2.2. Performance for Nonstationary Conditions: The Adaptive Solution

Given the success with stationary conditions in the previous trials, we simulated nonstationary conditions. We injected input that continuously varied the difference

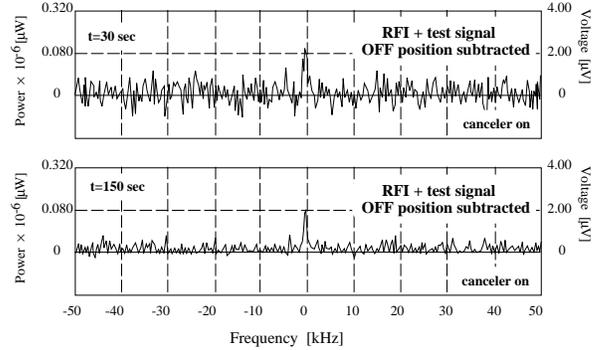


FIG. 13.—Same experiment as in Fig. 12, but with nonstationary conditions. A spectrum of the RFI in the system output, and OFF source position are shown in Fig. 12. The test signal has about the same strength as in Fig. 12 but is displaced in frequency. Both panels here show the recovered test source, with RFI and canceler operating and with the OFF position spectrum subtracted, but differ in integration time. The integration time is 30 s in the top panel and 150 s in the bottom panel. See col. (4) of Table 2 for measurements.

between the reference and primary channels in both amplitude and phase. Limitations of our equipment determined the magnitude of variations in amplitude and phase between the reference and primary inputs. We expect that these values are quite large compared to realistic conditions at the telescope, so these experiments pushed the filter beyond the edge of best performance. The four experiments in Table 3 were performed with sinusoidal amplitude and phase variations of $2 \mu\text{W}$ and 15° , respectively, between the reference and primary inputs. The first trial cycled the amplitude and phase and once every second, and the last three cycled every 10 s.

Spectra for trial 1 are shown in Figure 13. The test source was recovered and the RFI disappeared into the noise. This experiment had the same RFI of Figure 12 but this time with nonstationary conditions. The top and bottom panels show the recovered test signal after 30 and 150 s, respectively. The noise values in the table are for the 30 s integration. The positive baseline in the bottom panel is most likely due to the finite nature of our nine-tap system and not to the canceler. Trial 2 used a very strong RFI signal and weak test source, and again, good attenuation was achieved.

TABLE 3
HIGH DYNAMIC RANGE TESTS: NONSTATIONARY CONDITIONS

PARAMETER	TEST SOURCE		NO TEST SOURCE	
	Trial 1	Trial 2	Trial 3	Trial 4
$\xi = \text{IN}_R/\text{IN}_P$ (dB)	15	10	10	10
IN_P (dB)	16	61	62	62
Deviation of modulation (kHz)	30	10	10	10
Modulation function (random) (kHz)	4	5	20	20
Peak of test signal ($10^{-6} \mu\text{W}$)	0.10	0.56
RFI before filtering ($10^{-6} \mu\text{W}$)	20.4	45000	95050	95050
Residual RFI after filtering ($10^{-6} \mu\text{W}$)	0.10	0.12
Output peak-to-peak noise ($10^{-6} \mu\text{W}$)	0.003 (0.001) ^a	0.008	0.02	0.02
$A(T)$ ^b (dB)	40	67	60	59

NOTES.—The errors are $\pm 2\%$ in IN_P and the RFI peak, ± 0.002 in the output peak-to-peak noise, and ± 2 dB in the attenuation $A(\tau)$. Nonstationary conditions: for all trials, Δ amplitude and Δ phase are $2 \mu\text{W}$ and 15° , respectively; for trial 1, statistics cycle once a second. For trials 2, 3, and 4, the statistics cycle once every 10 s.

^a The number in parentheses is the peak-to-peak noise after 150 s integration (see Fig. 13).

^b $A(\tau)$ as in Table 1.



FIG. 14.—The 140 foot telescope at Green Bank Observatory with the prototype adaptive interference-canceling receiver for 100 MHz mounted at prime focus. The orthogonal Yagi antennas mounted at the top of the receiver box are pointed toward the horizon and serve as the reference feed.

Trials 3 and 4 were the only cases where the RFI was not attenuated down to the level of the noise. However, the canceler was still able to bring down the RFI by at least a factor of one hundred thousand [$A(\tau) = 50-60$]. In this

prototype, we have used a fixed step size, μ , and the most simplistic algorithm available (LMS) for finding and tracking the solution for the weighting coefficients. Since the tracking success depends on the algorithm and on μ , future

design improvements will use a more sophisticated algorithm and allow μ vary as necessary.

5.3. System Compatibility Tests on the 140 Foot Telescope

In 1997 June we had an engineering test of the single reference channel prototype receiver on the 140 foot telescope at Green Bank Observatory. The primary feed was a cross dipole mounted at prime focus. For the reference channel, we used two FM antennas mounted on the top of the receiver box (see Fig. 14). The input from the main feed was mixed before arriving at the IF. There were four interferers in the passband. We also found RFI leaking into our system from the control room and from the spectral processor, so a design improvement will be to put the adaptive canceler in a shielded box. During this shakedown of the system, we found that the anti-imaging filters need sharpening to confine the bandwidth and that a better interface to the spectral processor is necessary. We also verified that a separate canceler is needed for each polarization.

Even in the face of multiple interferers in our passband, polarization effects, and having only one reference channel, the system still attenuated radio stations to better than 25 dB. In light of the design limitations of this first prototype, we find these results to be very encouraging. The next step will be to implement the changes mentioned above and build up the system to have four reference channels with a canceler for each polarization. We plan to begin this work during the summer of 1998.

6. CONCLUSIONS

Laboratory experiments with a prototype adaptive-canceling receiver have shown that with one interferer in the passband, under stationary or nonstationary conditions,

the system can attenuate RFI to the maximum performance limited by its digital processor, 72 dB. The canceler locks up on the interference signal in less than 1 s and attenuates down to the rms noise in typically less than 20 s. The noise added by one reference channel is minimal for the experiments we performed. Since the interference-to-noise in the reference and primary inputs is a function of frequency, a slight ripple is introduced into the baseline when using the adaptive-canceling system. Tests that will be performed under more realistic conditions will allow a good characterization of this problem. We have verified that a separate adaptive-canceling filter is necessary for each polarization, and we will implement this in the next system we build. We also expect that a more sophisticated algorithm than LMS and the ability to vary the step size parameter μ will also improve the results. The attenuation performance is related to a higher interference-to-noise in the reference than primary channel, and if the RFI signal can be kept in the sidelobes of the main beam, IN_R can be higher than IN_P . To keep the RFI signal in the sidelobes, there will be an elevation angle limit of the telescope (e.g., the main beam should not be pointed at the horizon in the direction of the RFI source). This is an important consideration if the interference is coming from a satellite. The limit on the elevation angle will, of course, depend on the telescope and the nature of the interference sources.

We conclude that these initial results for the prototype adaptive interference-canceling receiver are promising.

We wish to thank R. Escoffier and R. Fisher for invaluable contributions to the development and implementation of the digital hardware. We are also indebted to S. Wilson for directing much of the theoretical design.

REFERENCES

- Bradley, R., Wilson, S., Barnbaum, C., & Wang, B. 1996, An Adaptive Interference Canceling Receiver for Radio Astronomy—Theory (NRAO Electron. Div. Internal Rep. 305) (Charlottesville, VA: NRAO)
- Bradley, R., Wilson, S., Escoffier, R., & Barnbaum, C. 1997, in National Radio Science Meeting Digest (Montreal: Intl. Union Radio Sci.), 609
- Erickson, W. 1983, in Interference Identification and Excision, ed. J. R. Fisher (Green Bank, WV: NRAO), 78
- Fisher, J. R. 1983, in Interference Identification and Excision, ed. J. R. Fisher (Green Bank, WV: NRAO), 10
- Gerard, E. 1983, in Interference Identification and Excision, ed. J. R. Fisher (Green Bank, WV: NRAO), 66
- Ghose, R. N. 1996, Interference Mitigation: Theory and Application (New York: IEEE Press)
- Goris, M. 1997, The Adaptive Beamformer of the SKAI Adaptive Antenna Demonstrator (NFRA Tech. Rep. 459/MG/CB/V2.2) (Dwingeloo: Netherlands Found. Res. Astron.)
- Haykin, S. S. 1996, Adaptive Filter Theory (3d ed.; Upper Saddle River, NJ: Prentice Hall)
- Ifeachor, E. C., & Jervis, B. W. 1993, Digital Signal Processing: A Practical Approach (Reading, MA: Addison-Wesley)
- Logic Devices Incorporated. 1995, Digital Signal Processing Data Book (Sunnyvale, CA: Logic Devices Inc.)
- Mintzer, F., & Liu, B. 1979, IEEE Trans. Acoustics, Speech, Signal Processing, 27, 204
- Moffet, A. 1982, in Interference Identification and Excision, ed. J. R. Fisher (Green Bank, WV: NRAO), 91
- Morrison, R. 1998, Grounding and Shielding Techniques (4th ed.; New York: Wiley)
- Oppenheim, A. V., & Schafer, R. W. 1989, Discrete-Time Signal Processing (Englewood Cliffs, NJ: Prentice Hall)
- Shin, D. C., & Nikias, C. L. 1994, IEEE Trans. Signal Processing, 42, 2715
- Sizemore, W. A. 1991, in IAU Colloq. 112, Light Pollution, Radio Interference, and Space Debris, ed. D. L. Crawford (ASP Conf. Ser. 17) (San Francisco: ASP), 176
- Sondhi, M. M. 1967, Bell Syst. Tech. J., 46, 497
- Widrow, B., et al. 1975, Proc. IEEE, 63, 1692
- Widrow, B., & Stearns, S. D. 1985, Adaptive Signal Processing (Englewood Cliffs, NJ: Prentice Hall)