Adaptive Noise Canceling in Geophysical Data
EE 373a Final Project

Salman H. Naqvi

Department of Electrical Engineering, Stanford University *

March 23, 2013

Abstract

Very Low Frequency (VLF) radio waves occur naturally due to various geophysical phenomena such as lightning discharges, solar wind interaction with the magnetosphere, and the energetic particle precipitation in the Earth's ionosphere. However, ground measurements of these waves using magnetic field antennas are corrupted by various types of natural and anthropogenic interference. Sources of these interference are primarily lightning, power transmission lines, and naval VLF transmitters, which manifest themselves in the antenna signal as intense impulsive broadband and narrowband harmonic noise respectively. In this paper, we implement the well-known adaptive noise canceling technique to filter the stationary components of noise in VLF signal. The non-stationary nature of the lightning-induced intense impulsive broadband noise makes the reference noise input non-stationary. Thus, we describe and first apply another technique, namely linear predictive coding, to remove this click-like noise. We show then that the application of adaptive filtering on the VLF data efficiently reduces the stationary noise components in the primary signal. We determine the efficiency of the adaptive noise cancellation by computing the results for different parameters and by comparing them with another technique for removing VLF noise recently described in the literature.

Keywords: VLF antenna, stationary noise, adaptive filtering, LMS algorithm, linear predictive coding

1 Introduction

The extremely low frequency / very low frequency (ELF/VLF) electromagnetic waves (defined here as ranging from $0.3-30kHz$ in the electromagnetic spectrum) are highly significant because of their very low-attenuation (on the order of a few dB/Mm) long-range propagation characteristics. For this reason, analysis of naturally occurring ELF/VLF waves is useful in investigating various ionospheric and magnetospheirc phenomena, in long-distance naval communication, in lightning geo-location, in geophysical prospecting, and in global positioning of submarines.

In particular, many geophysical phenomena are the sources of generation of VLF waves in the Earth's atmosphere and the near Earth space. A recent review paper [1] extensively reviews many of these phenomena including, but not limited to, the lightning strikes from cloud to ground and cloud to cloud, sprites generated over the thunderstorms, energetic particles population in Van Allen Radiation Belts in the Earth's magnetic field, interaction of solar wind and cosmic rays with the magnetosphere, and electron precipitation in the auroral regions. Studying these waves either in-situ via antennas on-board the satellites or at ground, is an invaluable way of remote sensing the Earth's ionosphere, and the various layers of the magnetosphere, and can give us information about the interaction of the geomagnetic field with the interplanetary field and the magnetospheric-ionospheric interaction.

Moreover, since ELF/VLF waves can propagate

to global distances (via Earth-Ionosphere waveguide) without attenuating significantly, and can penetrate the deep-sea water, navy utilizes VLF radio spectrum to communicate with their submarines and to have the submarines locate their position via triangulation of the VLF transmitter signals. Henceforth, many countries operate their own naval VLF transmitters operating at a specific VLF frequency precise upto $.1kHz$. Examples include the Cutler Maine VLF transmitter operated by US Navy at $24kHz$, which is the most powerful radio transmitter in the world with a transmission power of $1800kW$. It is formed by two arrays of antenna systems spanning a diameter of approximately $1.2m$. Figure 2 shows a cartoon of this transmitter.

A relatively recent usage of VLF waves generated by lightning, called the radio atmospherics, is in the real-time global geo-location of lightning occurrences for the purposes of safer aviation regulation. This is achieved by simultaneously measuring the radio atmospheric (or simply a "sferic") generated by a lightning strike with three or more ground antennas separated appropriately in azimuth. By evaluating the angle of arrival of the sferic at each antenna, we can triangulate its precise location.

Nevertheless, ELF/VLF radio waves have broad applications but most important of all for the scientific community are related to their ability to remotely sense the Earth's ionosphere and magnetosphere.

Ground-Measurement of VLF waves

The ELF/VLF ground receivers can be electric or magnetic field antennas, which are connected to the amplifier and then a digitizer that digitizes at baseband to acquire a full broadband spectrum. Stanford VLF receivers consist of two orthogonal air-core magnetic loop antennas, pointing North-South and East-West respectively. Stanford VLF group is operating a global constellation of over 80 of these receivers for many years. They operate at a sampling frequency of $100kHz$, measuring a broadband spectrum of $50kHz$, which corresponds to a time resolution of $20msec$. It also has sensitivity to the magnetic field of as low as $3fT$. Their operation is described in detail in [2] and the schematic is given in Figure 1.

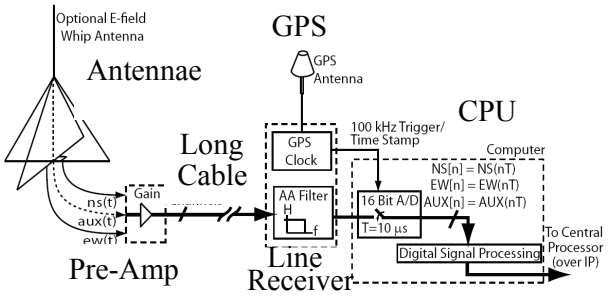

Figure 1: Schematic showing the complete VLF receiver with all the components used to process the antenna signal. Taken from [6]

Noise Sources in Geophysical Remote Sensing Data

For scientific experiments aimed at measuring the natural VLF waves originating at over $10kkm$ altitude in the auroral ionosphere and deep in the magnetosphere, signals from the VLF transmitters, or radio sferics constitute the unwanted noise in the ground antennas. Sferics are usually intense, amplitudes are 1-100 pT, impulsive, 1ms, in time and broadband in spectrum, with a maximal energy concentrated between 9-13 kHz. Signals from the VLF transmitters can also be intense depending on the location of the ground antenna.

Moreover, another significant noise source at most antenna locations is the local power transmission lines. These alternate the voltages at 50 or 60 Hz, but the non-monochromatic nature of these interference result in harmonics of this fundamental $60Hz$ that can extend up to several kHz. Overhead power transmission is the usual way of electric power transmission, and the electromagnetic interference from these lines, or "$60Hz$ hum" can often be significantly strong than the naturally occurring signals. So much so that the relatively "low-power" lines at km distances can present significant interference hazard when compared to natural sources.

Finally, there are high amplitude components of sferics that propagate below $500Hz$ in the transverse electromagnetic mode of the Earth-ionosphere waveguide and distort any VLF signal below this approximate upper cutoff.

Thus, the naval VLF transmitters, electric power transmission lines, and the lightning-induced sferics and slowtails are the major noise sources that need to be mitigated in the VLF antenna signal. Figure 2 shows a cartoon showing the above-mentioned major sources of the natural and anthropogenic interference in VLF antennas.

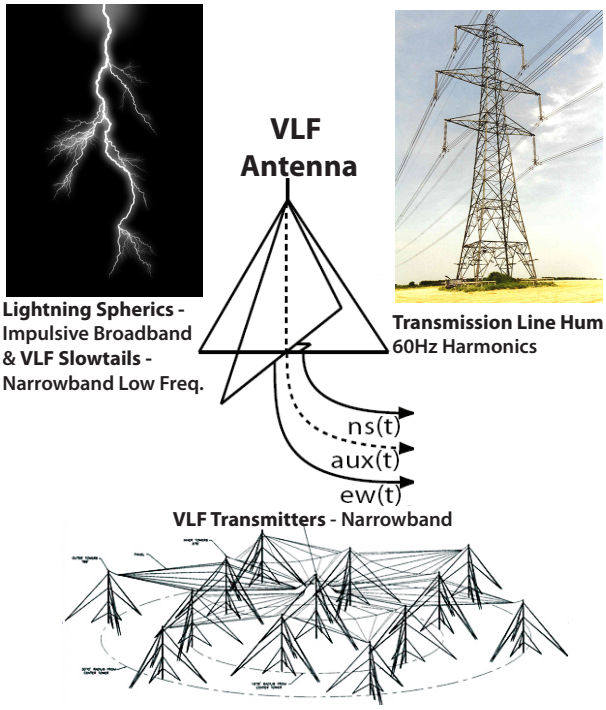

Figure 2: Few key Anthropogenic and Natural sources of Interference/ Noise in the VLF magnetic field signal

Data Acquisition and Representation

Figure 3 shows the time domain signal $25sec$ in duration acquired from the East-West channel of the Stanford VLF antenna located at the South Pole station (90^0S) from a day in June in 2006.

Time domain representation gives very little information about the spectral features of the signal. Hence, to detect particular interference sources present in the data, we evaluate the Short Time Fourier Transform (STFT) of this data, knowing its sampling frequency. The spectorgram, which is the Frequency-Time representation of the signal, is plotted in Figure 4.

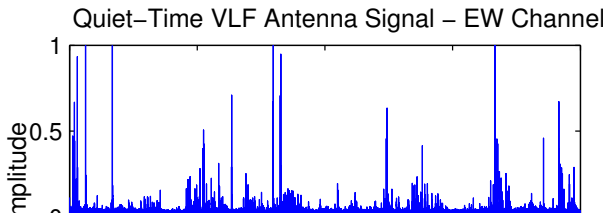

Figure 3: Magnetic Field Amplitude Signal from the VLF Antenna

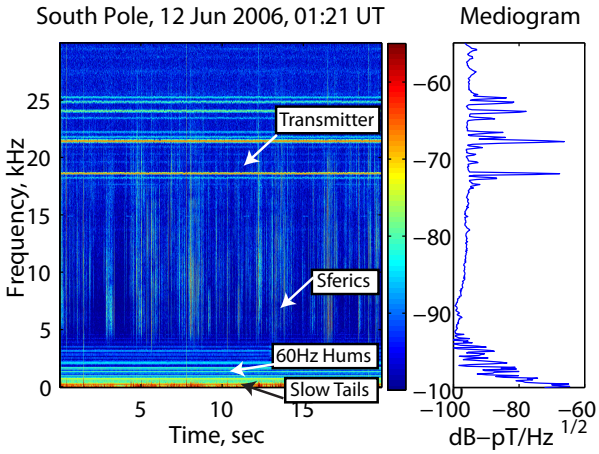

Figure 4: Spectrogram of Quiet-Time signal showing various types of anthropogenic and natural Interference

Plotted next to it is the "Mediogram" plot, which is essentially the median taken over the spectrogram matrix along the time index, thus plotting frequency against the power spectral density

(PSD) of the detected signal.

We can now easily notice the spectral signature of the four key noise sources discussed earlier. We also clearly see in the mediogram, the spikes in (PSD) correspondi6ng to the transmitter interference and $60Hz$ hum. This spectrogram is a typical of the East-West channel of the South Pole antenna during the times of low natural electromagnetic wave activity, which we call from now on as the "Quiet-Time."

During the "Active-Time," we expect to see the naturally occurring VLF waves generated higher in the ionosphere or deep in the magnetosphere. In particular, a natural VLF emission that is prevalent at the South Pole station, because of its strategic location in the auroral zone is the "Auroral Hiss" VLF wave. Figure 5 shows that auroral hiss waves are generated in the auroral ionosphere as a result of the precipitation of energetic electrons from the electron rich Van Allen radiation belts into the ionosphere, causing the auroral nightlights.

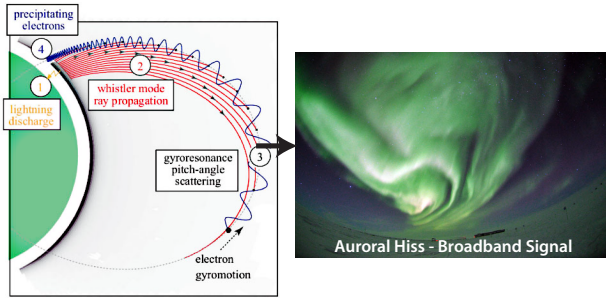

Figure 5: A specific case of a scientifically meaningful signal being generated by the geomagnetic process of energetic electrons precipitation to the auroral ionosphere and the resulting auroral nightlights generation. Auroral Hiss VLF electromagnetic waves are generated in conjunction with this phenomenon

Figure 6 shows the spectrogram of the South Pole VLF signal taken from an "Active-Time." We can approximate describe this situation as having the auroral hiss signal from the "Active-Time" riding on top of the cumulative interference signal from the "Quiet-Time. Here onward, we will call the naturally occurring magnetospheric-ionospheric waves, such as auroral hiss, as the "Meaningful" scientific signal, while terming the VLF signals from other sources as "Noise."

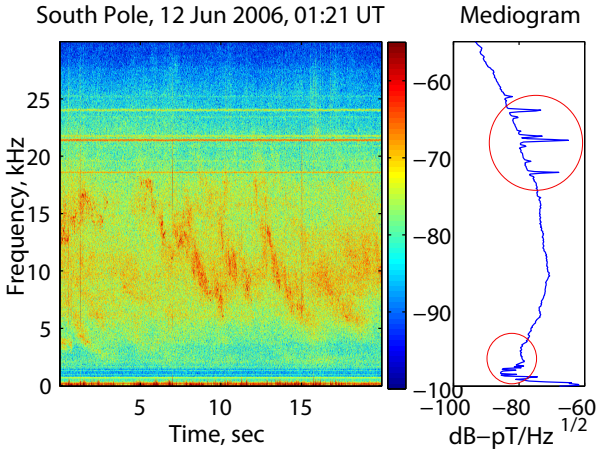

Figure 6: Frequency-Time Plot of 20sec of an Active-Time VLF signal

Therefore, if we can approximate the quiet-time signal as the stationary noise then we can use adaptive filtering technique to remove this stationary noise component from the "active-time" VLF signal. We now discuss our method of noise cancellation before applying it to the VLF data.

2 Method

2.1 Adaptive Noise Canceling

Noise cancellation can be actively achieved in the most fundamental setting by noting the interference signal present in the primary signal and subtracting it from the primary signal to extract the meaningful signal component. Figure 7 shows this setup adapted from the Panphonics Inc. website. Here, the ambient noise in the headphones can be removed by sensing it and subtracting it from the signal.

In reality, the scenario is not so simple, we don't have an idealized sinusoid, rather the interference is a sum of multiple discrete sinusoids of variable

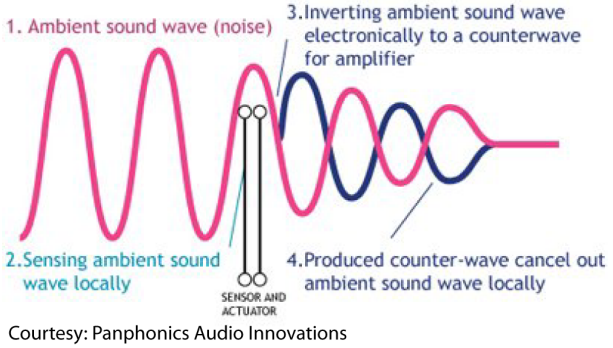

Figure 7: Fundamental concept of noise canceling

amplitudes. A fixed filter is not feasible because the noise mixed with the signal is not exactly the reference noise available to us. This is because the transmission paths of the reference noise to the primary signal have generally unknown characteristics. The scenario is shown in Figure 8, where a signal s is corrupted with noise n_0. n_0 is basically the noise source, n_1, added to the signal via a transmission path of unknown characteristics which has modified n_1 in an unknown way. The assumptions are that the noise n_0 is uncorrelated with the signal s and correlated in some unknown way with the reference noise input n_1.

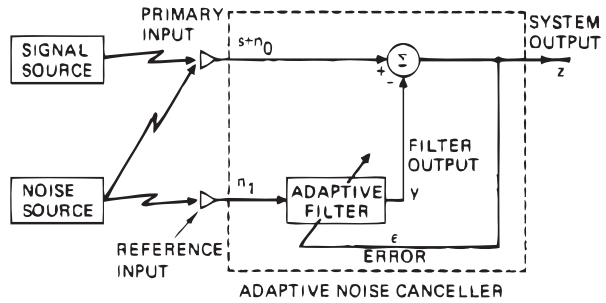

Figure 8: Schematic showing the basic concept of adaptive noise cancellation. Figure 1 of [7]

In this adaptive noise canceling scheme, the reference noise input n_1 is filtered to generate an output y which is as close a replica of the noise component n_0 in primary signal as possible at each time step. This output is then subtracted from the primary input $s + n_0$ to get the noise canceled output, $z = s + n_0 - y$. The filtration is done through an adaptive filter, which adjusts its own impulse response via an algorithm that is responsive to the error signal z. The objective of the adaptive filtering is to have the system output (error signal) be a best fit to the signal s in the least squares sense. In a nutshell:

"Minimizing the total output power,
Minimizes the output noise power, which
Maximizes the output signal-to-noise ratio."

Thus minimizing the output power, $E[z^2]$, makes the output nearly noise-free. Adaptive filters initially "learn" the input noise and signal statistics and then track them if they vary slowly. If the noise's statistics are rapidly varying then it might be difficult for the adaptive filter to learn them, let alone to track its variation, and it may never converge to the ideal solution. For stationary stochastic signals, however, the steady-state performance of the adaptive filters closely approximates that of the fixed filters.

Figure 9 shows an adaptive linear combiner implemented with a tapped delay line. Adaptive linear combiners are the main component of the adaptive filters, which weights and sums a set of delayed inputs to form an output signal. The adaptive filter algorithm we use in this project is the Least Mean Squared (LMS) Algorithm of updating the weights at each iteration or time step based on the error signal. The input signal vector, $\mathbf{X_j}$, and the Weight vector, W, used in Figure 9 are given below, where j is the iteration step:

$$X_j = \begin{bmatrix} x_{0j} \\ x_{j-1} \\ \vdots \\ x_{j-n+1} \end{bmatrix}; W = \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_n \end{bmatrix}$$

And the output of the adaptive filter is given as:

$$y_j = W^T X_j \tag{1}$$

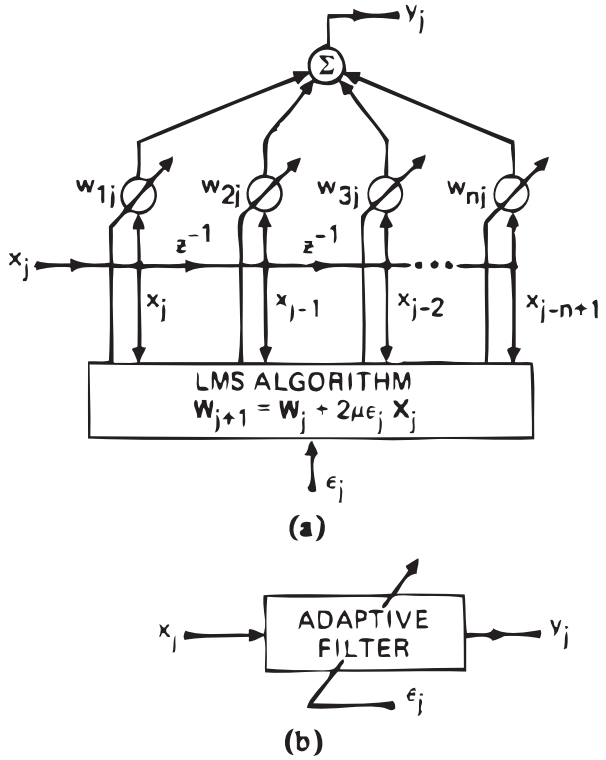

Figure 9: Adaptive Filter Block with the LMS algorithm. Panel (a) shows the Adaptive Linear Combiner with Tapped Delay lines as an expanded view of (b). Figure 29 of [7]

We can then write the error at each j^{th} iteration and its square as:

$$\varepsilon_j = d_j - W^T X_j \tag{2}$$

$$\varepsilon_j^2 = d_j^2 - 2d_j X_j^T W + W^T X_i X_j^T W \tag{3}$$

where, the desired response signal, d_j, is chosen to be the primary input signal.

From here, it is easy to see that the squared error and hence the MSE is a quadratic function of the weights of the adaptive filter, which can be pictures as a concave hyperboloidal surface that is always non-negative. Under this view, the process of adjusting the weights is equivalent to descending along this surface with the objective of getting to the "bottom of the bowl." And for the purpose of this descent, gradients methods are usually used to approximate the next step.

LMS algorithm is the implementation of the steepest descent method of updating the weight vector to arrive at the Wiener solution. Steepest descent method states that "the weights, W_{j+1} at the next iteration is equal to the current weights W_j plus a change that is proportional to the negative gradient, ∇_j, of the MSE" where μ is the factor determining the rate of convergence and stability.

$$W_{j+1} = W_j - \mu\nabla_j \tag{4}$$

LMS algorithm estimates an instantaneous gradient in an efficient but sloppy way by considering ε_j^2 as an estimate of the MSE, $E[\varepsilon_j^2]$, and by differentiating it with respect to W. This vector differentiation gives:

$$\nabla_j = -2\varepsilon_j X_j \tag{5}$$

which is the gradient vector of the MSE surface

Hence, the LMS algorithm for updating the weight vector at each successive iteration is given by:

$$W_{j+1} = W_j + 2\mu\varepsilon_j X_j \tag{6}$$

This algorithm is easily implemented without the need to square, average, or differentiate any vectors. Starting with arbitrary initial weights, LMS algorithm will converge in the mean and will remain stable if the parameter μ is positive, nonzero, and much less than the reciprocal of the largest eigenvalue of the input correlation matrix, R:

$$0 < \mu < 1/\lambda_{max} \equiv trace(R) \tag{7}$$

where, the trace of the matrix R can be found as $E[x^2]/N$, where N is the total length of signal $\mathbf{X}$, if the input is stationary. In common practice, μ is chosen to be atleast $1/10^{th}$ the μ_{max}. Higher μ causes the algorithm to converge hastily with improper filtering as we see in the results section as well.

This technique of adaptive noise-cancellation has been successfully implemented in many areas, especially the ambient noise cancellation in speech signals. Figure 10 shows the adaptive filter for speech application.

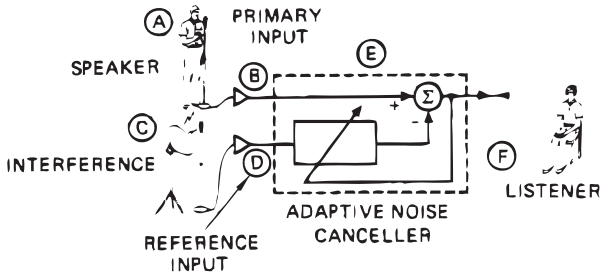

Figure 10: One well-known application of Adaptive Noise-Canceling to remove stationary interference in speech signals. Figure 18 of [7]

2.2 Impulsive Noise Cancellation

As stressed earlier, adaptive noise cancellation can be successfully implemented if the inputs are statistically stationary. Thus only the stationary or very slow varying (in the statistical sense) components of the noise in the primary input can be adaptively filtered out. Thus, before the application of the adaptive filter, we must first remove the intense impulsive broadband lightning-induced non-stationary random noise, termed earlier as sferics, that is present in the reference noise signal. We are assuming here that rest of the noise types are approximately stationary. We will see that this is a good assumption for our data. We achieve this preprocessing noise removal in a two step process, involving sferic detection and removal, followed by data interpolation.

In the first step, we identify the samples as sferics in the reference input using an amplitude threshold and zero them out. In order to speed up this process, we first band-pass the signal between $5 - 15kHz$ in order to focus on the spectrum where the maximum sferic energy is concentrated. We then square the signal to increase the contrast, and label all the samples above the threshold as sferics. A threshold of about $.01\sigma_d$ gives good result [5], where σ_d is the standard deviation of the band-passed and squared signal amplitude. If the detected sferic has a length under $1msec$ then it is extended to a $1msec$ duration, and if two sferics are detected less than $1.5msec$ apart then they are combined.

The detected sferics are then zeroed out in the original reference input. Before interpolation, to speed up the process, we low pass filter the data from $100kHz$ to $25kHz$ to focus on the sferic dominated region of the spectrum. We also high pass filter the data with a lower cutoff of $400Hz$ to filter out the slowtails that are bounded by the upper cutoff of about $500Hz$. We are then ready to interpolate over the zeroed out data samples. Interpolation is necessary in order to maintain the continuity of the data and hence to keep the other spectral features of the reference input intact.

The interpolation is achieved using the linear-predictive coding (LPC) of the signal using the auto-regressive (AR) model as explained in [4]. AR model represents approximately the n^{th} discrete sample of the signal as a linear combination of P previous samples plus a scalar error term:

$$x_n = \sum_{i=0}^{P} a_i x_{n-i} + e_n \tag{8}$$

where, a_i is the i^{th} AR coefficient, and P is the order of the AR process. Define a $(P+2)\times(2P+2)$ matrix A as follows:

$$\mathbf{A} = \begin{bmatrix} -a_P & \cdots & -a_1 & 1 & 0 & \cdots & 0 \\ 0 & -a_P & \cdots & -a_1 & 1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & -a_P & \cdots & 1 & 0 \\ 0 & \cdots & 0 & 0 & -a_P & \cdots & 1 \end{bmatrix} \tag{9}$$

Then the error sequence of length $(P+2)$ in equation (8) can be written in terms of the matrix A and signal vector:

$$\mathbf{e} = \mathbf{A}\mathbf{x} \tag{10}$$

For each detected sferic, we create three segments of the reference input vector $\mathbf{x}$; $\mathbf{x_u}$ is the unknown signal which consists of all zeros representing the sferic. This is preceded by $\mathbf{x_b}$, the known $20msec$

long prior samples of the signal, which is already free of sferics because of sequential cleaning. And it is followed by $\mathbf{x_a}$, which are also known but not sferic-cleared $5msec$ of data samples. Then we write the signal $\mathbf{x}$ as:

$$\mathbf{x} = \begin{bmatrix} \mathbf{x_b} \\ \mathbf{x_u} \\ \mathbf{x_a} \end{bmatrix} \tag{11}$$

We can also combine the known components of the input vector $\mathbf{x}$ in another vector $\mathbf{x_k}$, and correspondingly separate the matrix $\mathbf{A}$, into $\mathbf{A_k}$ and $\mathbf{A_u}$. Thus, re-writing the error $\mathbf{e}$ as:

$$\mathbf{e} = \mathbf{A_k x_k} + \mathbf{A_u x_u} \tag{12}$$

Our goal is to find $\mathbf{x_u}$, given $\mathbf{A_k}$ and $\mathbf{A_u}$ of the known data samples $\mathbf{x_k}$, that will minimize the mean-squared of the error vector, $\mathbf{e}$. By setting $\mathbf{e}$ to zero:

$$\mathbf{A_u x_u} = -\mathbf{A_k x_k} \tag{13}$$

Thus, we have the traditional least squares solution when the matrix $\mathbf{A_u}$ is skinny and full rank, which is a suitable assumption to make here:

$$\mathbf{x_u} = -(\mathbf{A_u^T A_u})^{-1}\mathbf{A_u^T A_k x_k} \tag{14}$$

Hence, $\mathbf{x_u}$ can be found if $\mathbf{A}$ is known. The LPC coefficients used to populate the matrix $\mathbf{A}$ are found using the built-in MATLAB function `lpc`. We use the entire $\mathbf{x}$ vector in estimating these coefficients. This process is repeated sequentially until all the sferics are detected and interpolated over.

3 Implementation and Results

We start by applying the impulsive noise removal technique described in the last section to the reference input noise signal. The algorithm does a good job at reducing the sferics to the level of the background noise. This is shown in Figure 11, where the blue line plot shows the filtered reference input.

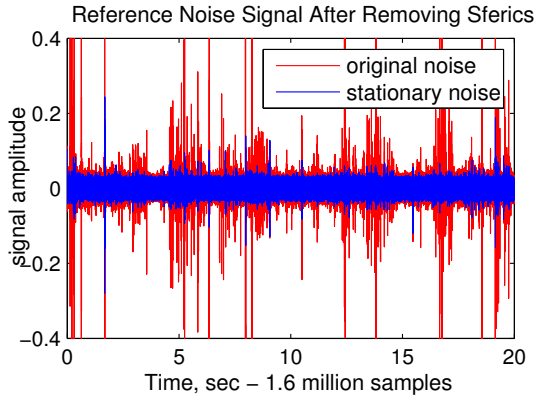

Figure 11: Time Domain noise signal before and after the impulsive noise reduction

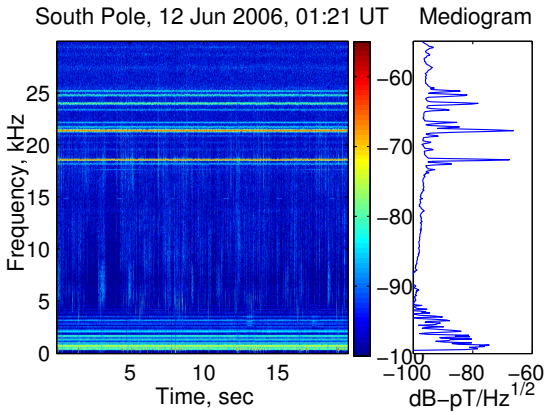

(a) Spectrogram of Sferic-removed noise signal

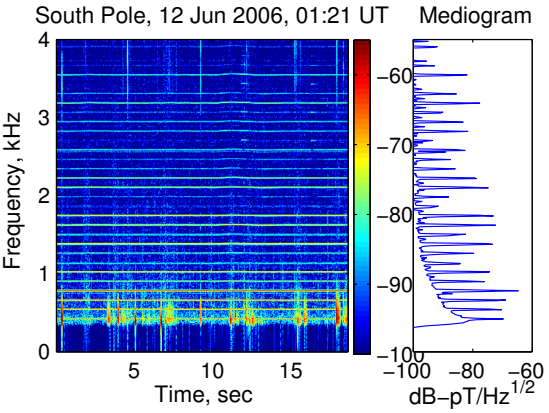

(b) Zoomed in plot emphasizes that the non-stationary slow-tail noise is also removed

Figure 12: After removing the impulsive noise. Spectral features are now approx. stationary

The effect of sferic filtering can best be seen by computing the STFT of the reference input. Spectrograms in Figure 12 show how the sferics and the slowtails have been filtered out to make the reference input a stationary signal.

Now, the reference input is stationary and correlated with the primary input. We can send these two signals as input to the adaptive filter described in Section 2 and shown in Figure 13. We use 10^4 weights and $\mu = .01 * \mu_{max}$.

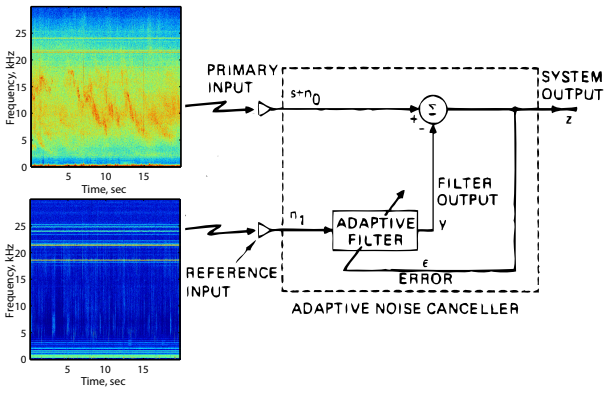

Figure 13: Spectrograms of the two input signals to the adaptive filter

The result of the adaptive filter are shown in time domain representation in Figure 14.

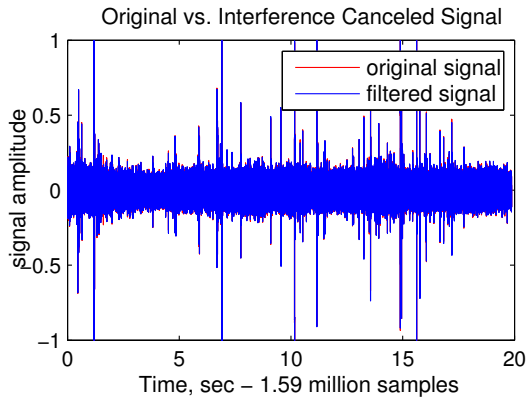

Figure 14: Output of the Adaptive Linear Combiner is overlaid on the Primary Signal

We see that the adaptive process has reduced the amplitude by just the right amount such that the non-stationary components in the primary input are not adversely affected.

In Figure 15b, we compare the spectrograms of the output signal with that of the primary input. The circled boxes emphasize how much are the transmitter noise and the transmission line harmonics reduced by the adaptive process. Note specifically that the transmitter signal at $21.32kHz$ is reduced by $4dB$ while that at $18.62kHz$ reduced by $3.5dB$, which is a significant reduction.

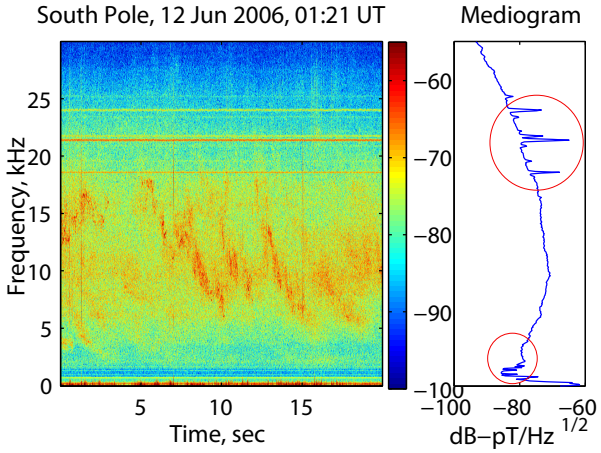

(a) STFT of the Primary Signal (Noisy) before ADF

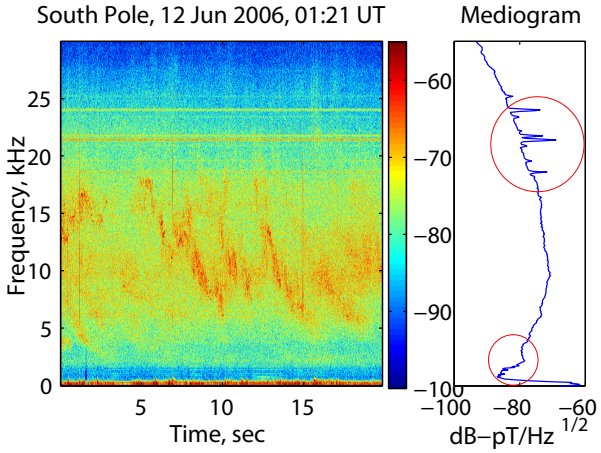

(b) STFT of the Output of the ADF

Figure 15: Primary Signal before and after the Adaptive Noise Cancellation

We low pass filter the output signal to zoom in to

the lower $4kHz$ passband to see the effect of filtering on the power line hums. Figure 16 shows the original input in panel (a) along with the filtered output in (b). We note that some of the lower frequencies at $.8kHz$ are reduced by nearly $19dB$, the higher ones at $1.3kHz$ are reduced by $4dB$, and the ones higher still, which already have an increasingly lower effect on the meaningful signal are lowered by $1dB$ at $2.1kHz$ and so on.

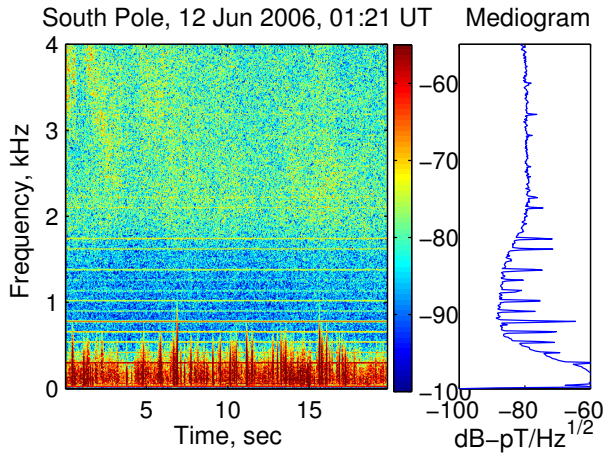

(a) STFT of the Low pass filtered Primary Signal

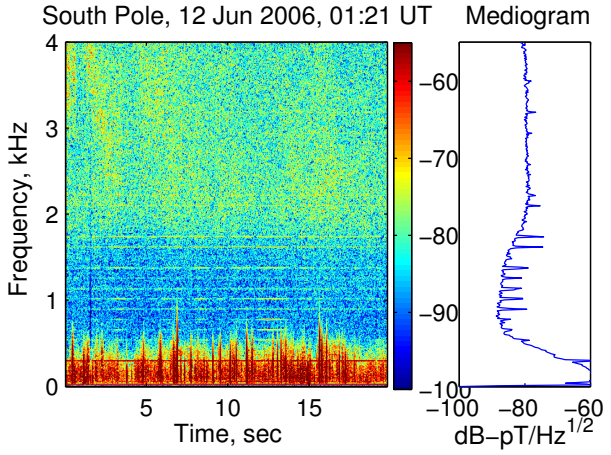

(b) Low pass filtered Noise-Canceled Output

Figure 16: Plots showing that the 60Hz Hums of the Transmission lines are also removed

We also compute the learning curve of this adaptive filtering process and to extract the trend, we applying a running median to the .4 million samples of the MSE. This is shown in Figure 17

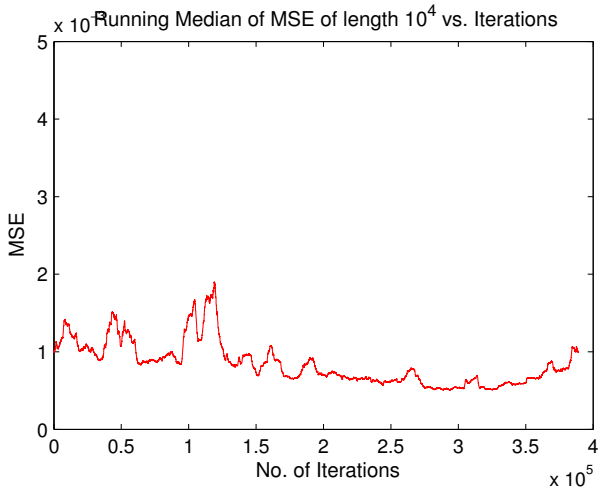

Figure 17: Learning curve of the adaptive process. MSE is computed over .4 million iterations and a running median over the 10000 samples is plotted to see the underlying trend.

4 Discussion

In this example, we used the "Quiet-Time" noise data to be from the same data as the "Active-Time" data just 5 hours apart. In this technique, we can use the noise reference for the same channel and antenna at the South Pole from any other day, month or year as well as long as there was not considerable variation in the noise characteristics. In case of doubt, we can always search for a small segment of "Quiet-Time" noise reference signal in the database of the same day as that of the "Active-Time" data required to be filtered, and repeat the reference input to make it as long as the "Active-Time" data, and then perform the adaptive noise cancellation.

4.1 Parameter Variation

Our adaptive filter has produced commendable results but the process was not so efficient. The adaptive process takes $158sec(2.6min)$ to run on a $20sec(1.6millionsamples)$ long input signal with 10^4 weights. The slow processing speed makes the process inefficient.

So we try using $1/10^{th}$ as many weights, i.e. 10^3. We notice that the processing time is substan-

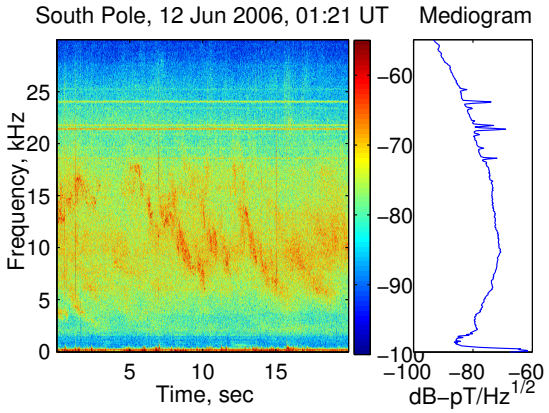

Figure 18: Noise Canceled Output with $1/10^{th}$ as many weights as before

tially reduced to about $22sec$ and the noise reduction is nearly perfectly the same as shown in Figure 18.

We also try varying the parameter μ for faster convergence and better performance at the expense of the algorithm's stability. We use a 10 times larger value of μ, i.e. $\mu = .1 * \mu_{max}$. Figure 19 shows the results.

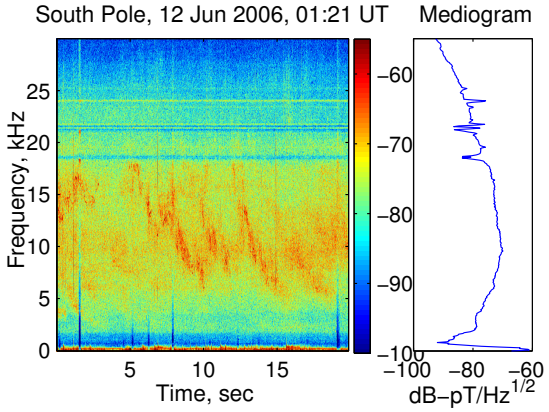

Figure 19: Noise Canceled Output from the Adaptive Filter with 10 times higher convergence factor μ ($\mu = .1 * \mu_{max}$)

Here, we see that the transmitter noise is reduced by a further $3-4dB$ and the hum removal seems to have performed better. But there are clearly visible artifacts in the form of broadband nulls. These might have resulted because the primary input was not stationary because of the presence of sferics, and hence any attempt at removing these sferics would have caused broadband frequency components to be removed.

Also, it is worth pointing out that the filter diverges for high values of μ if we use low number (10^3) of weights. Hence, we have to work with large number of weights at the expense of processing speed.

4.2 Other Techniques for 60Hz Harmonics Removal

We also compare our transmission line hum removal results with another method recently discussed in the literature [3]. In this method, Cohen et al. first estimate the frequencies of all the $60Hz$ harmonics using quadratic interpolation method using odd or even harmonics (whichever are stronger at a given station). Next, hum is removed via the least squares estimation method. They then repeat this process in the intervals of 200msec to account for any slowly varying changes in the hum frequency. Figure 20b show the results achieved using this method.

Comparison with Figure 20a reveal that both the methods perform well. Although the least squares method doesn't do such a good job at lower frequencies, which are equally well filtered by the adaptive noise canceler.

The least squares method is computationally faster though, taking nearly half the processing time than the adaptive filter with 10^3 weights. On the downside, this technique is limited to removing power line hum of $60Hz$ while adaptive filter also reduces the transmitter noise. Moreover, the adaptive filter automatically detects the variation in noise frequency and we don't have to track this variation explicitly.

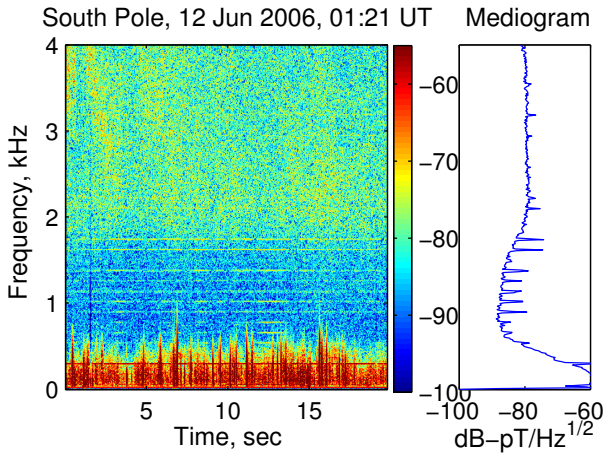

(a) 60Hz Hums removed using ADF

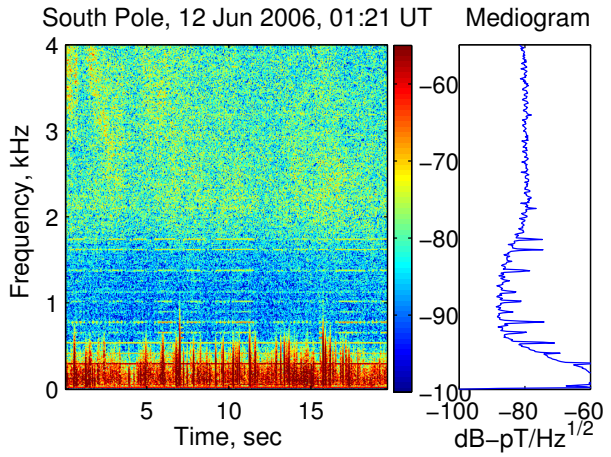

(b) 60Hz Hums removed using Least Squares Method [3]

Figure 20: Comparison of 60Hz Hums removal with a least squares technique, recently described in literature [3]

5 Conclusion

We have successfully implemented adaptive noise canceling algorithm on the VLF magnetic field data. We have shown that if the reference input is stationary then the stationary component of the noise that are correlated with primary input can be successfully removed or attenuated. Our comparison of this method with other methods discussed in the literature shows that our algorithm is more resilient and gives a better performance at lower frequencies. It is computationally less efficient though but only by a scalar factor, not by an order of magnitude.

6 References

[1] R. Barr, D. Jones, and C. Rodger. ELF and VLF radio waves. *Journal of Atmospheric and Solar-Terrestrial Physics*, 62(17-18):1689–1718, Nov. 2000.

[2] M. Cohen, U. Inan, and E. Paschal. Sensitive Broadband ELF/VLF Radio Reception With the AWESOME Instrument. *IEEE Transactions on Geoscience and Remote Sensing*, 48(1):3–17, Jan. 2010.

[3] M. B. Cohen, R. K. Said, and U. S. Inan. Mitigation of 50-60 Hz power line interference in geophysical data. *Radio Science*, 45(6):1–12, Dec. 2010.

[4] S. J. Godsill and P. J. W. Rayner. *Digital Audio Restoration - a statistical model based approach*. Springer, 1 edition, 1998.

[5] D. Golden. *Source Variation and Ground Accessibility of Magnetospheric Mid-Latitude ELF/VLF Chorus and Hiss*. PhD thesis, Stanford University, 2011.

[6] R. Khalil. *Accurate and Efficient Long-Range Lightning and Geo-Location using a VLF Radio Atmoshperic Waveform Bank*. PhD thesis, Stanford University, 2009.

[7] B. Widrow, J. Glover, J. McCool, J. Kaunitz, C. Williams, R. Hearn, J. Zeidler, E. Dong, and R. Goodlin. Adaptive Noise Cancelling : Principles and Applications. *Proceedings of the IEEE*, 63(12):105–112, 1975.