

Finite Impulse Response (FIR) Filters
to Simulate Response of an Antenna

Thomas M. Sokol, Dr. Inder J. Gupta and Prof. Edward H. Newman

Honors Thesis for Graduation with Distinction
Submitted by Thomas M. Sokol

The Ohio State University
Department of Electrical Engineering
2015 Neil Avenue
Columbus, OH 43210

Honors Committee:

Prof. Edward Newman, Advisor

Dr. Inder J. Gupta, Senior Research Scientist, Advisor

Approved by:

Advisor

Acknowledgements

I would like to start out by thanking everyone who helped me with this project. First, I would like to thank my mom for encouraging me to challenge myself. I would also like to thank Dr. Gupta for his patience and helping me understand the concepts in this paper intuitively. I would also like to thank Prof. Newman for co-advising me on this paper as well as for teaching me a lot about radar systems. Finally, I would like to thank my girlfriend Genie for giving me the motivation to work hard since I know she is always there to make my days better. I would also like to thank the graduate students Andy, Chad, Erik, and Jon at the ElectroScience Lab for making it a fun place to be.

Abstract

In performance analysis of radar/communications systems using computer simulations or physical wave-front generator, it is often desired to include in the simulation the response of the antenna. Incoming signals are distorted by the antenna as they are received, and therefore modeling this effect is crucial to accurate signal processing and analysis. In general, the impulse response of an antenna is not available. However, one can measure or otherwise acquire the gain and phase response of the antenna over the bandwidth of interest at several discrete frequencies. Using these discrete frequency-domain data points, one is interested in obtaining a finite-length filter to represent the antenna response. We have studied several methods using an inverse Fourier transform approach to obtain the filter coefficients. We have shown that if the bandwidth of the filter is somewhat larger than the bandwidth of the incoming signal, that various approaches involving the inverse Fourier transform to obtain the coefficients can be used. The simplest case involves taking the inverse Fourier transform of the frequency data and using the result as the filter coefficients. However, this simple approach may not create a filter that models the antenna within an acceptable error tolerance over the band of interest. The methods studied significantly improved upon the basic Fourier method for determining the filter coefficients. One approach involves windowing the coefficients to further reduce the error. The second approach studied involves windowing the frequency-domain samples. The final approach involves interpolating the data and solving for the coefficients in the least squares sense to obtain further improvements in modeling the antenna.

Table of Contents

ACKNOWLEDGEMENTS	2
ABSTRACT	3
TABLE OF CONTENTS	4
TABLE OF FIGURES	5
INTRODUCTION	6
SIMPLE FOURIER TRANSFORM	8
OVERVIEW.....	8
ASSUMPTIONS	8
PROCEDURE.....	10
RESULTS & DISCUSSION.....	13
TIME-DOMAIN WINDOWING	17
OVERVIEW.....	17
TRIANGULAR WINDOW FUNCTION	17
HIGHER-ORDER WINDOW FUNCTIONS	20
FREQUENCY DOMAIN WINDOWING	24
OVERVIEW.....	24
PROCEDURE.....	24
LINEAR WINDOW FUNCTION	25
SINE SQUARED WINDOW FUNCTION.....	29
LEAST SQUARES DESIGN APPROACH	33
OVERVIEW.....	33
PROCEDURE.....	33
RESULTS.....	35
SUMMARY	39
SUMMARY	42

Table of Figures

FIGURE 1: ANTENNA GAIN AND PHASE RESPONSE.....	9
FIGURE 2: INCIDENT SIGNAL POWER SPECTRUM.....	10
FIGURE 3: ANTENNA GAIN AND PHASE SAMPLES, 2 MHz SPACING.....	11
FIGURE 4: REAL & IMAGINARY COMPONENTS OF IDFT FILTER COEFFICIENTS	12
FIGURE 5: IDFT FILTER GAIN RESPONSE COMPARED TO ACTUAL ANTENNA GAIN RESPONSE	13
FIGURE 6: IDFT FILTER PHASE RESPONSE COMPARED TO ACTUAL ANTENNA PHASE RESPONSE	14
FIGURE 7: NORMALIZED ERROR FUNCTION OF IDFT FILTER	15
FIGURE 8: TRIANGULAR WINDOW FUNCTION	18
FIGURE 9: RECTANGULAR & TRIANGULAR WINDOW TECHNIQUE GAIN RESPONSE.....	18
FIGURE 10: RECTANGULAR & TRIANGULAR TECHNIQUE PHASE RESPONSE	19
FIGURE 11: RECTANGULAR & TRIANGULAR WINDOW TECHNIQUE NORMALIZED ERROR FUNCTIONS	19
FIGURE 12: HAMMING & KAISER WINDOW FUNCTIONS.....	21
FIGURE 13: HAMMING & KAISER WINDOW TECHNIQUES GAIN RESPONSE.....	21
FIGURE 14: HAMMING & KAISER TECHNIQUES PHASE RESPONSE	22
FIGURE 15: NORMALIZED ERROR FUNCTIONS FOR ALL WINDOW TYPES.....	22
FIGURE 16: LINEAR WINDOW TAPER	25
FIGURE 17: LINEAR FREQUENCY WINDOW SAMPLES.....	26
FIGURE 18: LINEAR FREQUENCY WINDOW GAIN RESPONSE COMPARED TO KAISER & ACTUAL.....	27
FIGURE 19: LINEAR FREQUENCY WINDOW PHASE RESPONSE COMPARED TO KAISER & ACTUAL.....	27
FIGURE 20: LINEAR FREQUENCY WINDOW NORMALIZED ERROR FUNCTION COMPARED TO KAISER & ACTUAL.....	28
FIGURE 21: SINE SQUARED FREQUENCY WINDOW.....	30
FIGURE 22: SINE SQUARED WINDOWED FREQUENCY SAMPLES	30
FIGURE 23: GAIN RESPONSE OF SINE SQUARED FREQUENCY TECHNIQUE	31
FIGURE 24: PHASE RESPONSE OF SINE SQUARED FREQUENCY WINDOW TECHNIQUE.....	31
FIGURE 25: NORMALIZED ERROR FUNCTION OF SINE SQUARED FREQUENCY WINDOWED TECHNIQUE.....	32
FIGURE 26: N=100 FREQUENCY SAMPLES	35
FIGURE 27: WEIGHTING MATRIX DIAGONAL ELEMENTS	36
FIGURE 28: LEAST SQUARES GAIN RESPONSE.....	37
FIGURE 29: LEAST SQUARES PHASE RESPONSE.....	37
FIGURE 30: LEAST SQUARES NORMALIZED ERROR FUNCTION.....	38

Introduction

In order to find a good method for modeling an arbitrary antenna response as a Finite Impulse Response (FIR) filter, several techniques were studied or devised to improve the filter reconstruction of the antenna response. In each case, it was assumed that N equally spaced frequency samples were available of the antenna response in the given direction. It was also assumed that the antenna bandwidth was greater than the bandwidth of the incident signal. The goal was to reconstruct the antenna response over the portion of the antenna bandwidth containing the desired signal using a FIR filter.

The first case studied was a simple inverse Fourier approach, using the transform of frequency samples of the antenna response as the filter coefficients. This technique had a problem of creating large ripples and oscillations in the reconstructed frequency response.

The second technique studied sought to correct the oscillations and ripples caused by the finite nature of the IFT approach. A higher-order window, such as a Hamming or Kaiser window, was applied to the filter coefficients obtained in the simple IFT case in order to significantly reduce the ripple. However, this technique had a problem with energy loss in the desired portion of the bandwidth.

The third technique involved windowing the frequency samples before the Fourier operation was performed. By windowing the frequency samples, one has control over which portion of the frequency spectrum has energy loss, thus eliminating the energy loss problem.

The fourth technique was an improvement upon the other techniques using a least squares approach to solve for the filter coefficients. In order to solve for the filter coefficients using the Fourier operation in the least squares sense, one is required to have more frequency samples than desired filter coefficients. This condition can be met by either obtaining for frequency samples or interpolating the ones available.

A summary of the results explains when to use each technique outlined in this study. Depending on the scenario, constraints, and assumptions, different techniques are suitable for different applications.

Simple Fourier Transform

Overview

The study began with a review of the inverse Fourier transform. In general, in order to find the discrete impulse response of a discrete series of frequency samples, one applies the Inverse Discrete Fourier Transform (IDFT) to the frequency data. Using the IDFT as the antenna filter coefficients, an analysis was performed on the effectiveness of this method. Finally, the errors and problem that result from using this method are discussed.

Assumptions

When designing a FIR filter, there are many well-known techniques for determining the filter coefficients assuming the frequency response is linear in magnitude or phase. This is almost never the case for a real antenna response, therefore none of the aforementioned techniques can be used for this study. In general, an antenna response will be well-behaved but non-uniform and non-linear phase over its bandwidth. This means that there are no sudden discontinuities in either the gain or phase response, and neither the gain and phase response can be modeled by a low-order polynomial. This study assumes that the antenna's gain and phase responses are available over its bandwidth. In general, antennas are designed to receive incident signals over a certain bandwidth B around a center frequency, f_0 . In this study, we are not concerned about the carrier frequency aspect of the antenna, only the bandwidth itself. Therefore, this study set the center frequency to 0 MHz for simplicity and assumed the antenna has a 50 MHz bandwidth. Figure 1 illustrates a typical gain and phase response for an antenna, and this response is used throughout the study. In this study, the actual antenna response will be

denoted $A(f)$. The spectrum containing the incident signal, $0.64B$ (32 MHz) is outlined with dashed lines and will be so in subsequent figures.

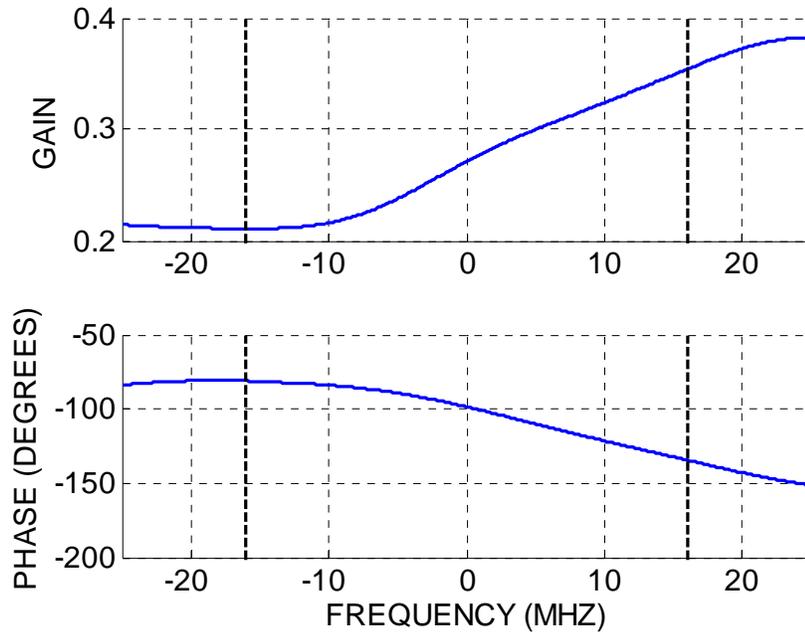


Figure 1: Antenna Gain and Phase Response

In general, hardware FIR filters operating at high frequencies do not have extremely long filter lengths. For the purpose of this study, we have chosen a 25-tap filter length (N) for suitability and relevance to applications. All filters designed and plotted in this study are length 25. However, the methods used in this study are applicable to any arbitrary length filter.

Typically, the bandwidth of the antenna is greater than the bandwidth of the incident signal. In this study, the incident signal is assumed to have a bandwidth of 32 MHz ($0.64B$) around the center frequency, and a uniform power spectrum over the 32 MHz, as illustrated in Figure 2. The figure shows that the incident signal has no frequency component outside the $0.64B$ of the total bandwidth. Thus, when modeling the antenna response, one only needs to consider the portion of filter frequency response that contains the desired signal.

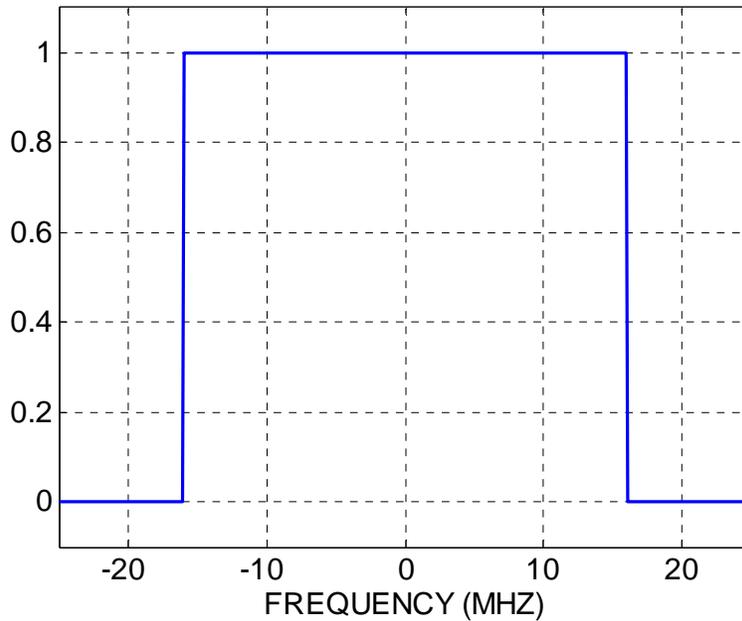


Figure 2: Incident Signal Power Spectrum

Procedure

In order to generate the antenna filter coefficients using the Simple Fourier Transform method, one first starts with the N Δf -spaced antenna frequency samples, where N times the frequency spacing Δf is equal to the antenna bandwidth. Therefore, the first step is to measure or acquire the frequency samples. For this study, the 25 frequency samples with 2 MHz spacing are assumed to be available; the methods for acquiring or calculating them are numerous and unimportant for this study. Figure 3 shows the gain and phase of each of these uniformly spaced antenna frequency samples, spanning the 50 MHz bandwidth. The frequency samples are denoted $E[k]$ and have a corresponding frequency f_k .

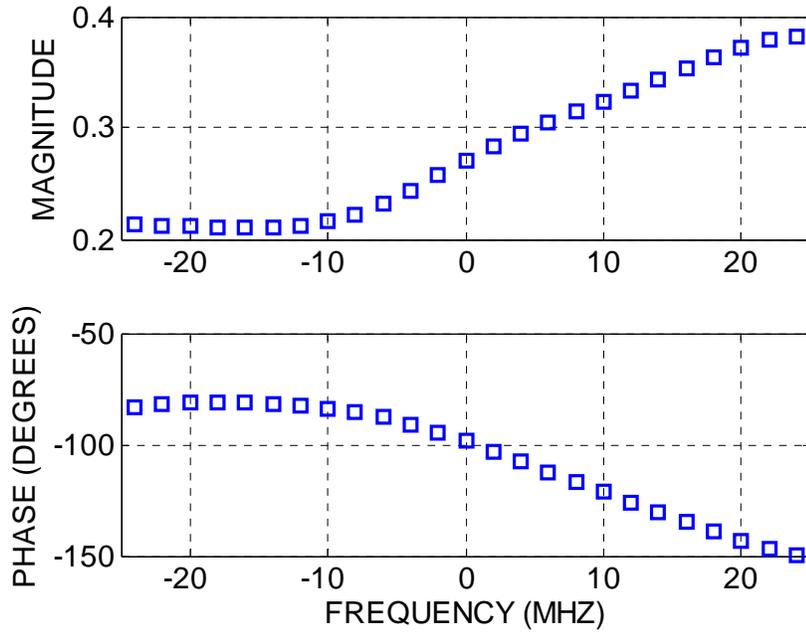


Figure 3: Antenna Gain and Phase Samples, 2 MHz Spacing

The Fourier Transform and its' inverse are widely understood, so the equations relating the frequency samples to the time domain samples are simply used with no derivations in this paper. When one has a finite series of frequency samples, denoted $E[f]$, their impulse response, $e(t)$, can be computed at any time t as follows.

$$e(t) = \frac{1}{N} * \sum_{k=0}^{N-1} E[k] * e^{j2\pi f_k t} \quad (1)$$

Using simple extensions of this formula, one can obtain a series of equally spaced time samples representing the impulse response. In general, when one has $N \Delta f$ -spaced frequency samples, one will have N coefficients in the time domain with $1/N\Delta f$ spacing after applying a simple IDFT.

$$e[m] = \frac{1}{N} * \sum_{k=0}^{N-1} E[k] * e^{j2\pi f_k t_m} \quad (2)$$

$e[m]$ denotes the m -th time coefficient. There are N elements in $e[m]$, and t_m is the time value of the m -th time sample in $e[m]$. In our example, the frequency spacing Δf was 2 MHz and the time spacing is $1/25\Delta f$ (20 nanoseconds). The resulting coefficients are used to model the antenna response. Figure 4 shows the real and imaginary components of the filter coefficients calculated using this method.

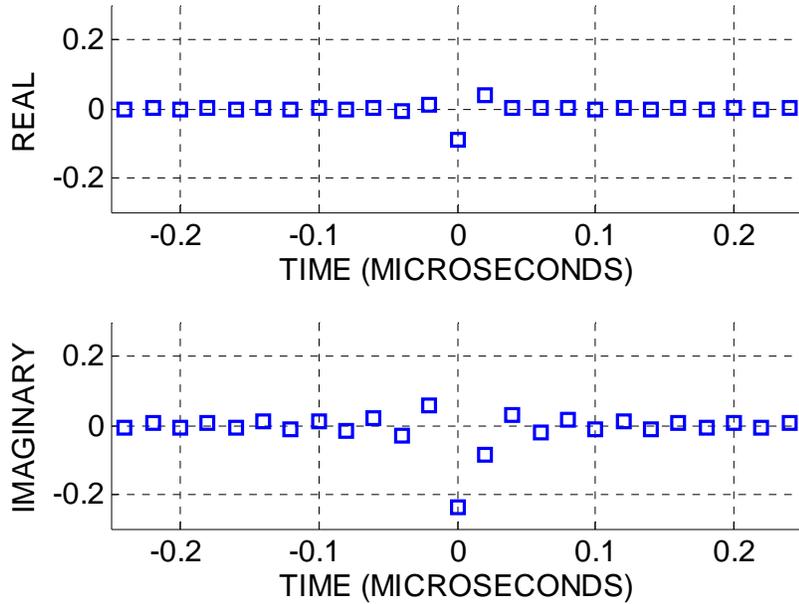


Figure 4: Real & Imaginary Components of IDFT Filter Coefficients

Using the filter coefficients, one can find the frequency response of the filter for any frequency using the following inverse operation.

$$\alpha(f) = \sum_{m=0}^{N-1} e[m] * e^{-j2\pi f t_m} \quad (3)$$

$\alpha(f)$ denotes the frequency response of the filter coefficients. In general, this identity can be used to find the frequency response at all frequencies. One can calculate the gain and phase response for all values in between the original Δf frequency samples and plot the

result vs. the actual antenna response to analyze how good of a job the filter is doing representing the original antenna response over the band of the incident signal.

Results & Discussion

Using the simple Fourier technique, the filter coefficients were determined and the frequency response, $\alpha(f)$ was calculated and plotted over the antenna bandwidth. Figure 4 shows the gain response of the filter designed using the IDFT technique, $\alpha(f)$ compared with the original antenna gain response, and Figure 5 shows the phase response of the IDFT filter compared to the phase response of the antenna. In both figures, the bandwidth of the incident is outlined with dashed lines.

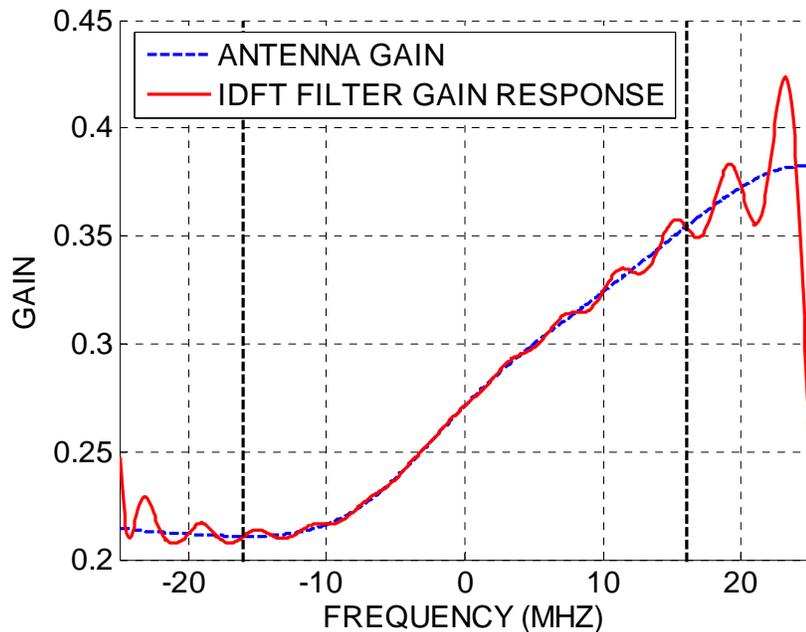


Figure 5: IDFT Filter Gain Response Compared to Actual Antenna Gain Response

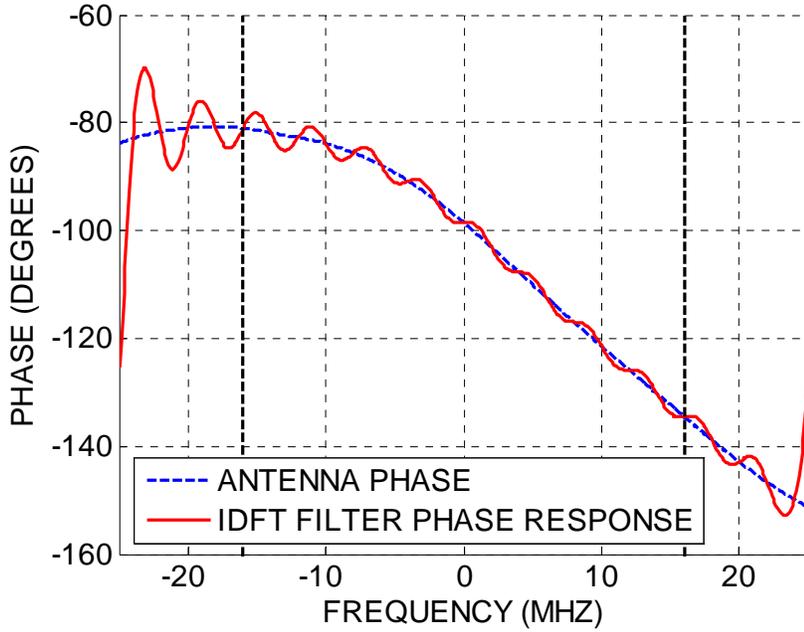


Figure 6: IDFT Filter Phase Response Compared to Actual Antenna Phase Response

As shown in Figures 5 & 6, the filter does a reasonable job approximating both the gain and phase response of the actual antenna. In fact, at the frequencies corresponding to the original frequency samples f_k , the filter perfectly reconstructs the antenna pattern. This result should be intuitive because the only information available was the gain and phase response at the frequencies f_k . However, in between the frequencies corresponding to the original samples, a significant ripple appears in both the gain and phase response.

The ripple is the significant challenge to overcome in this study, because the ripple introduces some error into the simulation if left uncorrected. In order to better analyze this ripple, and error function $R(f)$ was defined as follows:

$$R(f) = \left| \frac{A(f) - \alpha(f)}{A(f)} \right| \quad (4)$$

$R(f)$ will be referred to as the normalized error of the filter. This metric is defined as the magnitude of the following expression: the complex difference between the actual antenna frequency response and the filter frequency response divided by the actual

antenna frequency response. The normalized error function is typically plotted in dB by taking the base ten logarithm of the normalized error function and multiplying the result by 20. Figure 7 shows the normalized error function $R(f)$ of the IDFT filter.

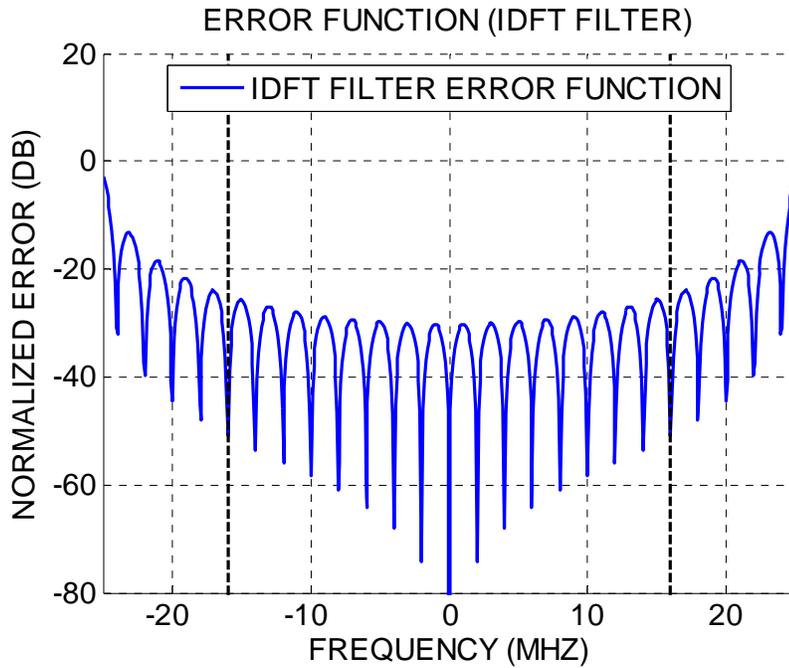


Figure 7: Normalized Error Function of IDFT Filter

This plot of the normalized error of the IDFT filter illustrates why this metric was defined. Normalizing the error function allows the analysis to properly account for the relative severity of the ripple over certain portions of the band. Figure 7 shows several effects which were observed from the gain and phase response plots of the IDFT filter response (Figures 5 & 6). Namely, the ripple is more severe near the fringes of the band and less severe in the portion of the band containing the desired signal. The dashed lines outline the 32 MHz portion containing the incident signal, and the worst case error is approximately -26 dB in this region, which corresponds to a 5% worst case error. Additionally, the error function shows that the IDFT filter frequency response matches the actual antenna pattern exactly at the frequencies f_k , shown on the figure as the deep

nulls in the error function every 2 MHz. For many applications, a 5% error is simply unacceptable.

Determining why the IDFT method produces such a large ripple is the key to figuring out how to reduce the problem. Consider for a moment that anything which is finite duration in the frequency domain will have infinite duration in the time domain. Thus, a finite-bandwidth antenna response requires an infinite duration time series representation in order to achieve perfect reconstruction. However, the time series implemented by the IDFT method was only 25 coefficients spaced 20 nanoseconds apart. This method is equivalent to calculating the infinite duration time filter and applying a rectangular window function to the center 25 filter coefficients. It is well known that a rectangular window function in the time domain causes large ripples in the frequency domain. The second filter design approach considers whether different windowing functions other than the rectangular method can improve performance.

Time-Domain Windowing

Overview

As discussed in the previous section, a finite-bandwidth antenna response requires an infinite time duration filter for perfect reconstruction. However, the simple IDFT approach to designing the filter produces an N-length finite duration filter. This filter is equivalent to taking the infinite duration filter required for perfect reconstruction and applying a rectangular window to the center N elements. A rectangular windowing operation in the time domain results in a fast ripple in the frequency domain. This second method for designing the filter studies various types of windowing functions in addition to a rectangular window to determine what performance improvements can be achieved. The premise is the ripple can be reduced with higher order windowing functions, and to compare several windowing functions and study the performance tradeoffs.

Triangular Window Function

The method for creating this type of filter is step-for-step identical to the simple inverse Fourier transform approach except for one step. Previously, the antenna frequency samples $E[k]$ were operated on by the IDFT producing the time domain filter coefficients $e[m]$. The frequency response $\alpha(f)$ was then computed from the antenna filter coefficients. This second method varies in that it applies a window function to the filter coefficients $e[m]$. Several types of windowing functions were compared, the first is the simple rectangular window function, not illustrated, which is equivalent to the IDFT method. The second type is a triangular window function, which is term-by-term multiplied with the filter coefficients $e[m]$ to produce $e_{tri}[m]$, a corresponding filter response $\alpha_{tri}(f)$, and a corresponding normalized error function $R_{tri}(f)$. Figure 8 shows the triangular window function, Figure 9 shows the gain response of $a_{tri}(f)$ and Figure 10

shows the phase response using $a_{tri}(f)$. Figure 11 shows the corresponding normalized error functions for both the rectangular and triangular window methods.

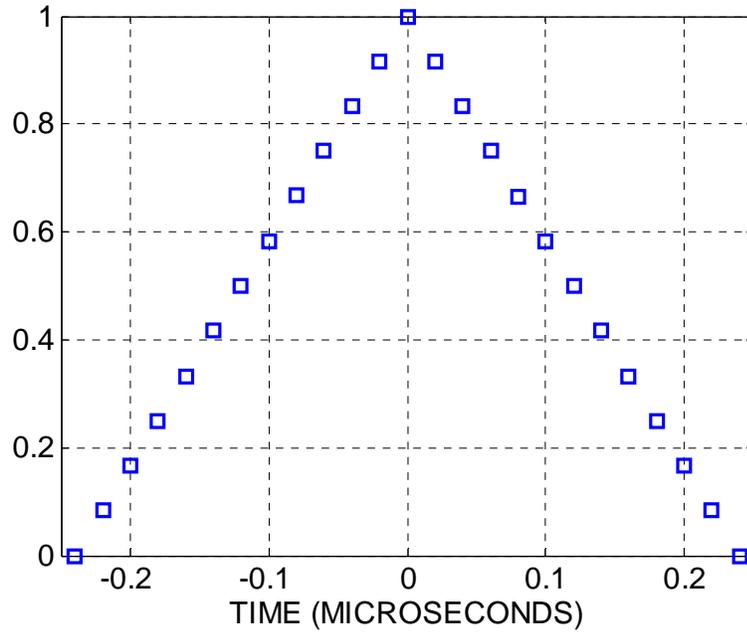


Figure 8: Triangular Window Function

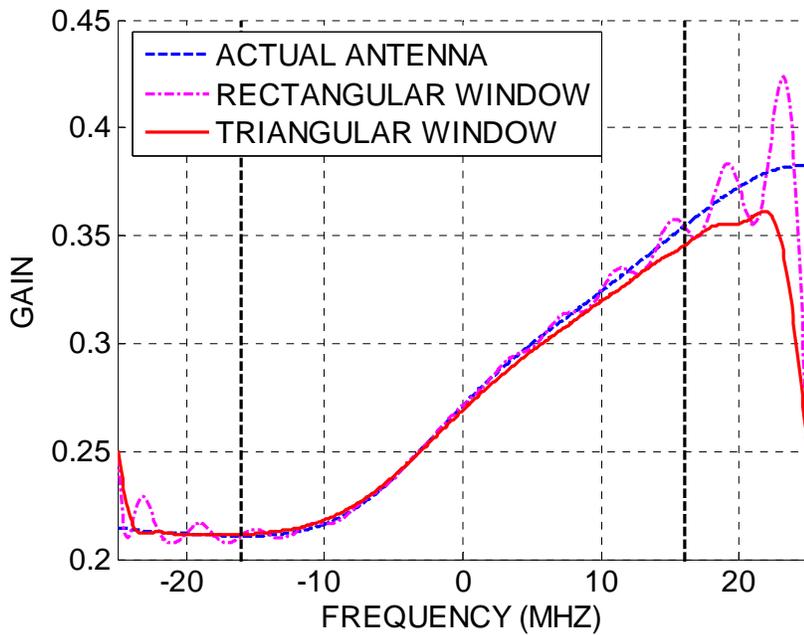


Figure 9: Rectangular & Triangular Window Technique Gain Response

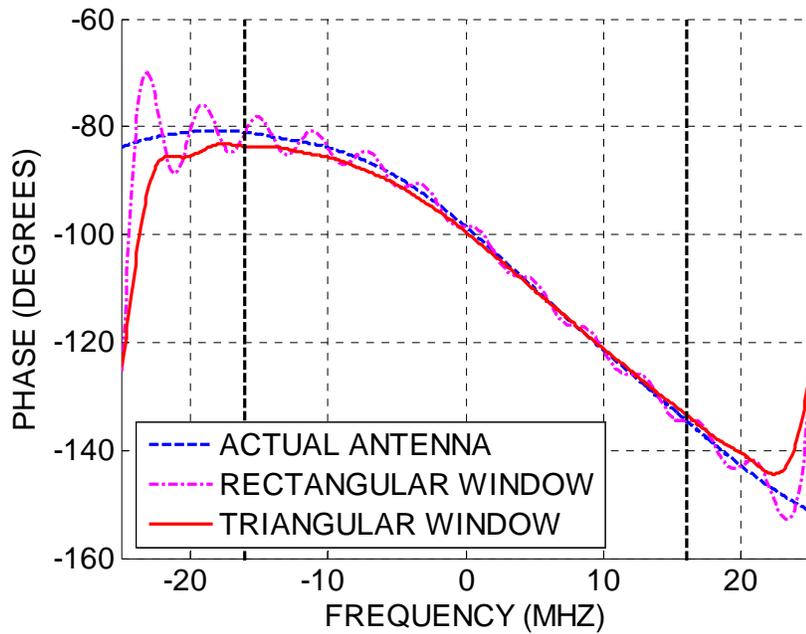


Figure 10: Rectangular & Triangular Technique Phase Response

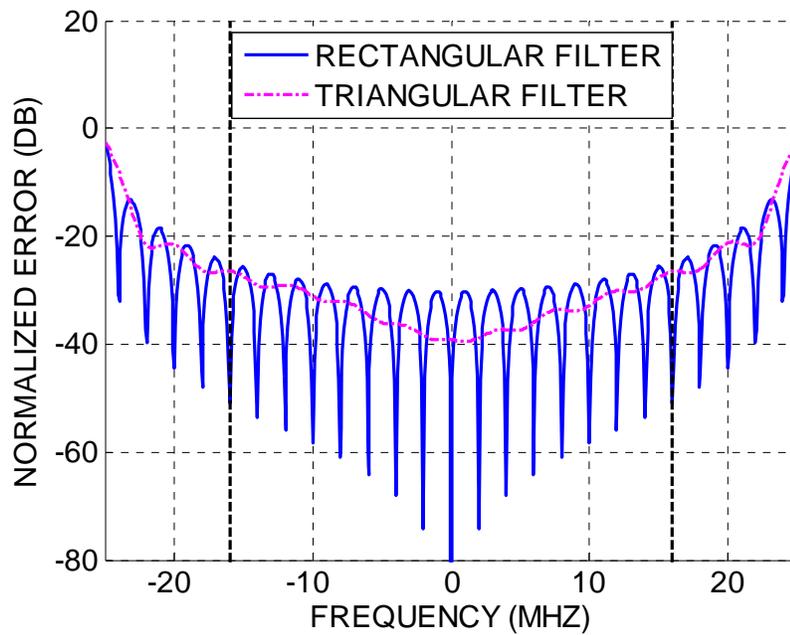


Figure 11: Rectangular & Triangular Window Technique Normalized Error Functions

A comparison of the two normalized error functions reveals that the triangular technique marginally improves performance compared to the rectangular technique. The worst case error improves to -27 dB from -26 dB. However, there is some “bias” to the frequency response, as it no longer matches exactly with the actual antenna response at

the frequency points f_k . Figures 9 & 10 show that the gain and phase response using the triangular window no longer exactly tracks the magnitude and phase response of the actual antenna. This can be understood simply as a tradeoff, reducing the ripple by altering the filter coefficients causes a less accurate tracking of the actual antenna response. Nonetheless, for some applications a “bias” in the antenna response may be preferable to a larger ripple. Because of this “bias”, one of the problems with the time-domain windowing technique is that there is some energy loss in the portion of the band corresponding to the incident signal. This is because the windowing function applied to the filter coefficients affects the entire bandwidth of the filter, not simple one portion of it.

Higher-Order Window Functions

The triangular window function showed marginally improved performance over the rectangular window function. Higher order windowing functions should show larger performance improvements than the triangular approach. Both a Hamming window and a Kaiser-Bessel window (Beta of 6) are also compared to see how much better higher-order windowing functions improve performance over the rectangular and triangular approaches. The Hamming & Kaiser approaches have corresponding coefficients $e_{\text{hamming}}[m]$ and $e_{\text{kaiser}}[m]$, corresponding frequency responses $a_{\text{hamming}}(f)$ and $a_{\text{kaiser}}(f)$ respectively, and corresponding error function $R_{\text{hamming}}(f)$ and $R_{\text{kaiser}}(f)$ respectively. Figure 12 shows the Hamming & Kaiser windows, Figure 13 compares both new filter gain responses, while Figure 14 compares their phase responses. Figure 15 compares the error functions for all four windowing techniques.

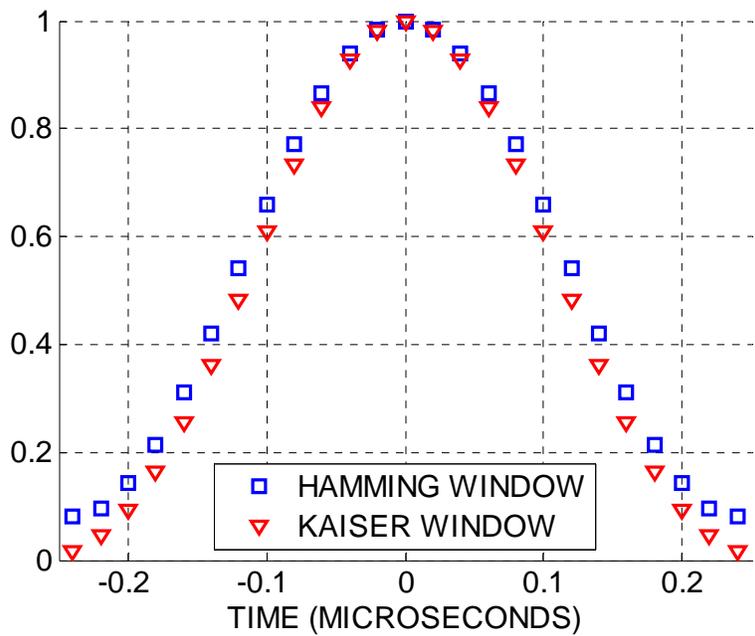


Figure 12: Hamming & Kaiser Window Functions

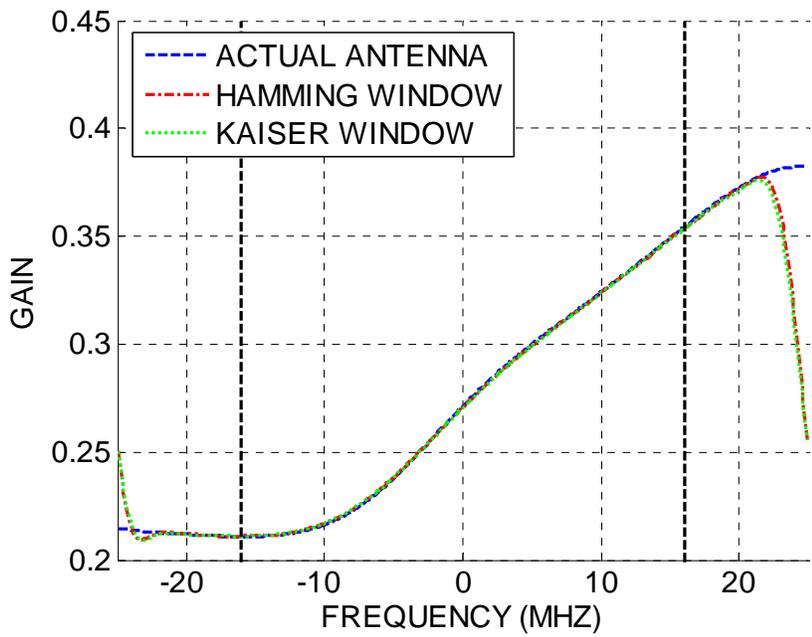


Figure 13: Hamming & Kaiser Window Techniques Gain Response

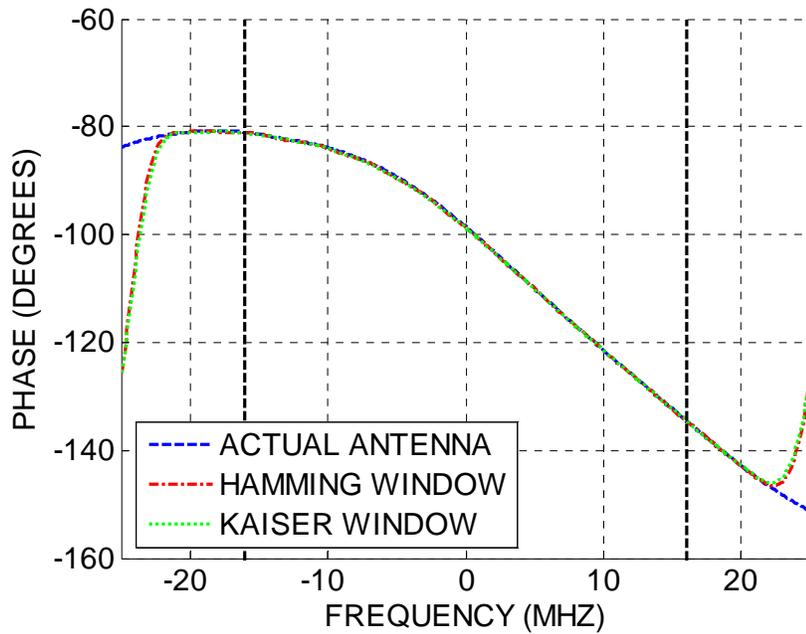


Figure 14: Hamming & Kaiser Techniques Phase Response

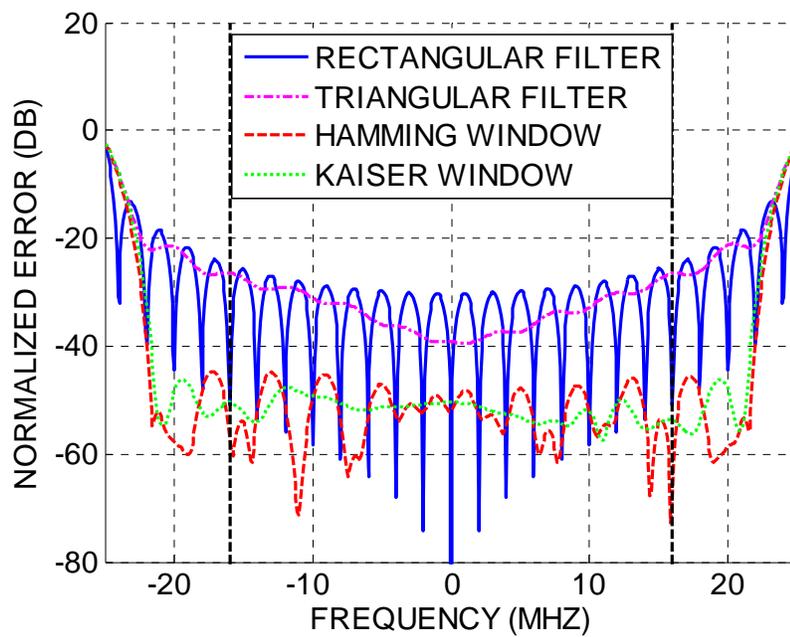


Figure 15: Normalized Error Functions for All Window Types

As expected, the higher order windowing functions improved performance in the desired band further than the lower order windowing techniques. The Hamming window approach resulted in a worst case error of -44 dB and the Kaiser window approach resulted in a worst case error of -48 dB, or maximum errors of 0.63% and 0.44%

respectively. Additionally, for the higher order techniques, the “bias” effect discussed previously becomes less pronounced. The ripple is almost complete non-existent in the portion of the band corresponding to the incident signal for the Kaiser technique, despite the small “bias”.

This second technique demonstrated that significant performance gains in the portion of the antenna bandwidth corresponding to the incident signal can be achieved by using a high order window function. Lower order window functions such as a rectangular or triangular technique exhibited significantly higher normalized error than the higher-order techniques. The best technique studied was the Kaiser-Bessel window with a Beta of 6.

Frequency Domain Windowing

Overview

The time domain windowing technique demonstrated that windowing the antenna filter coefficients after taking the IDFT of the frequency samples can improve the performance of the filter over the band of interest (center 32 MHz). When no window was applied to the IDFT coefficients, there was a significant ripple in the gain and phase response of the filter. The ripple effect was significantly reduced when a high order window was applied to the coefficients. However, the tradeoff for the reduced ripple was a “bias” effect: the distortion of the magnitude and phase response caused by a loss of energy in the desired band. Perhaps for some simulations, the bias is less of a problem than a fast ripple. In others, perhaps the bias is a worse problem to have. Nonetheless, this third technique was motivated by attempting to reduce the bias while still using a windowing technique to improve performance. Instead of windowing after the IDFT operation is performed on the frequency samples, $E[k]$, a window is applied before the IDFT operation is performed. These frequency domain windows must be more carefully designed than the time domain windows, so as to not cause energy loss in the $0.64B$ (32 MHz) of the spectrum containing the incident signal.

Procedure

The steps for this technique are once again similar to the simple IDFT approach. The N frequency samples, $E[k]$, are term-by-term multiplied by a frequency window which preserves the samples within the incident band ($0.64B$) while tapering to 0 those outside the incident band resulting in $E'[k]$. After this windowing is applied, the steps are identical to the simple IDFT technique to produce the filter: the IDFT operation is

performed on the $E'[k]$ samples and the filter's gain and phase response are calculated, as well the new error function.

Linear Window Function

The first frequency window tested was a simple linear taper, shown in Figure 16. The window is term by term multiplied with the original antenna samples, $E[k]$, resulting in the new samples $E_{\text{linear}}[k]$, whose magnitudes are shown in Figures 17 compared with the original samples, $E[k]$. The windowing operation does not affect the phase of the samples, so a phase comparison is not shown because they are identical.

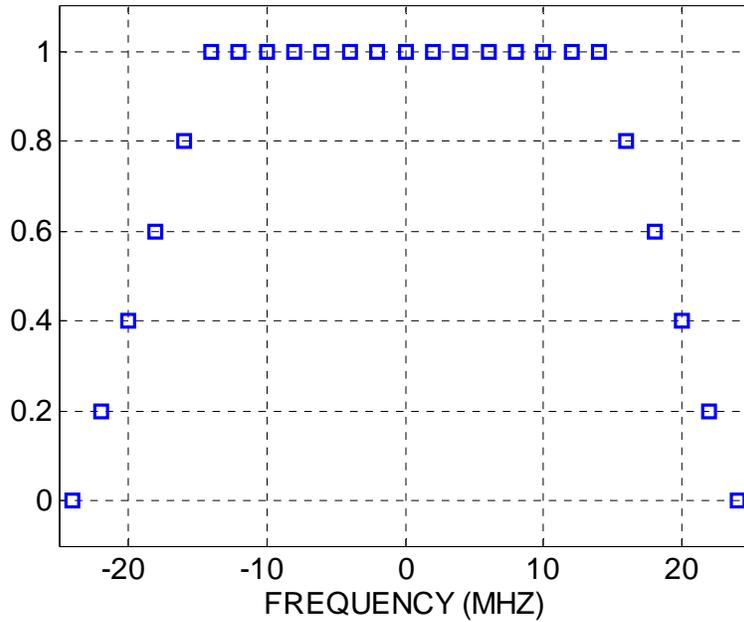


Figure 16: Linear Window Taper

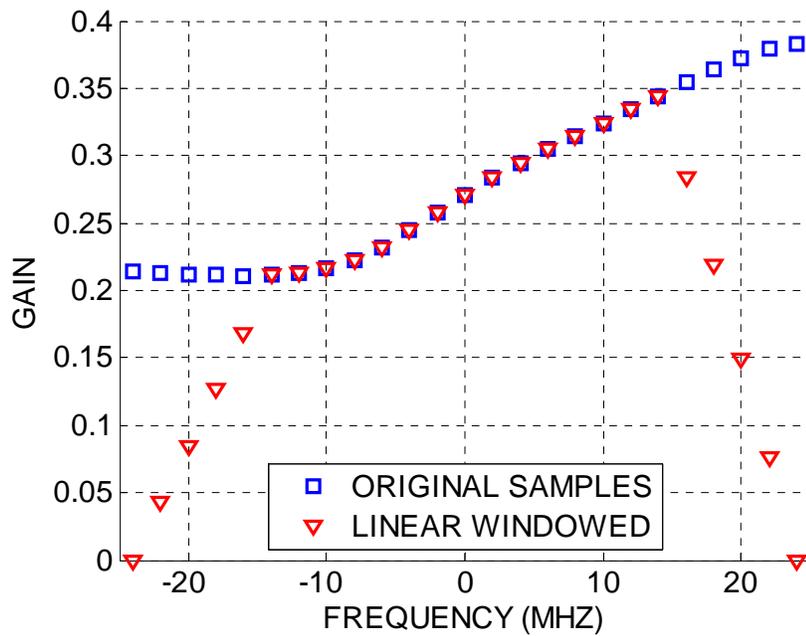


Figure 17: Linear Frequency Window Samples

After the IDFT operation is performed on the frequency windowed samples and the new filter coefficients are obtained, the filter gain and phase response, $\alpha_{linear}(f)$ can be calculated as well as the normalized error function, $D_{linear}(f)$. Figure 18 shows the filter gain response compared to the Rectangular and Kaiser window techniques discussed previously. Figure 19 shows their phase responses. Figure 20 shows the corresponding error functions compared.

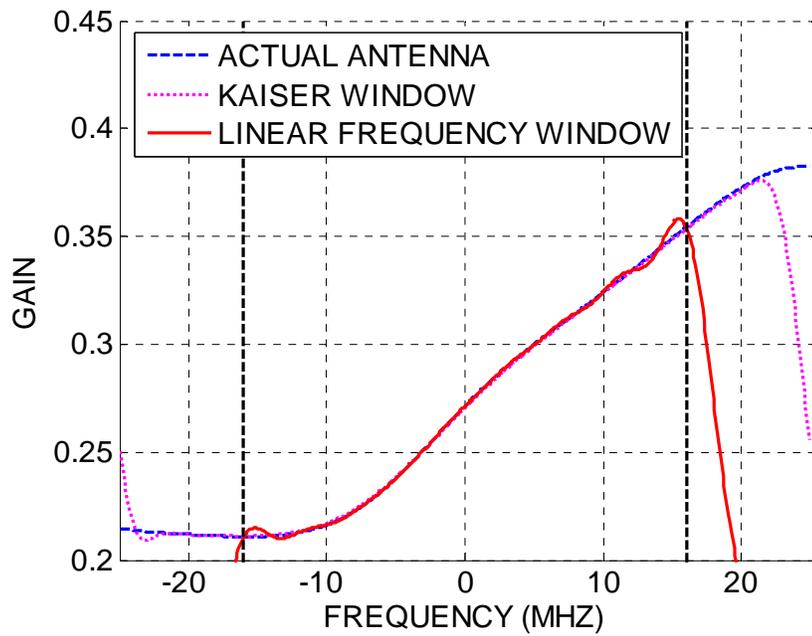


Figure 18: Linear Frequency Window Gain Response Compared to Kaiser & Actual

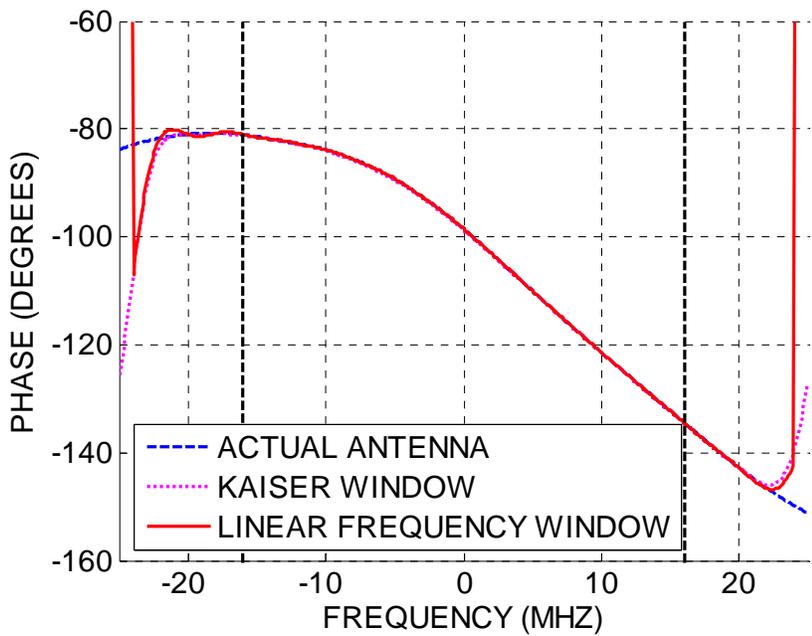


Figure 19: Linear Frequency Window Phase Response Compared to Kaiser & Actual

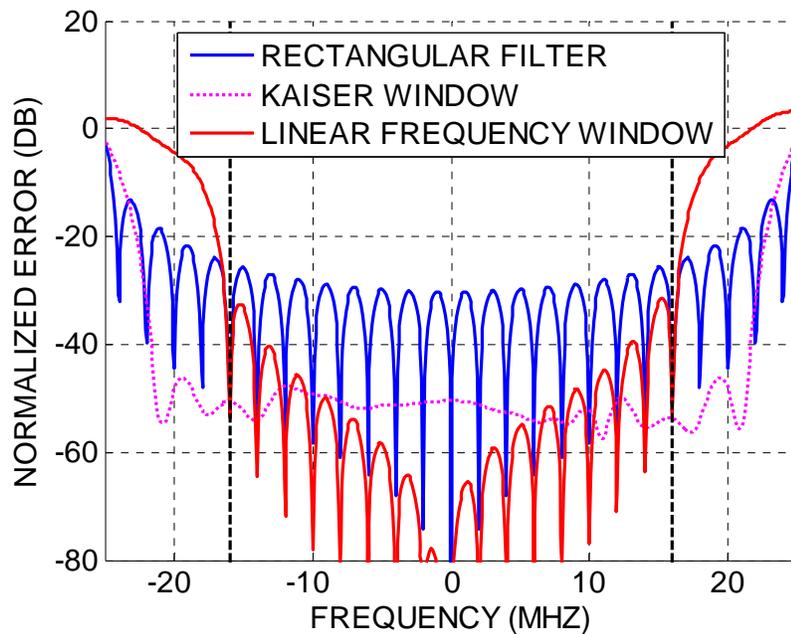


Figure 20: Linear Frequency Window Normalized Error Function Compared to Kaiser & Actual

Based on Figures 18 through 20, we can see that this technique produces fairly good results compared with the simple IDFT filter (Rectangular Filter). There is no evidence of bias, as shown in the error function by the deep nulls corresponding to the original frequencies f_k . The worst case error in the desired band (center 32 MHz denoted by hashed black lines) for this technique is -33 dB compared to -26 dB for the simple IDFT technique. This represents a maximum error of 2.2% vs. 5% for the simple IDFT approach. However, based on the shape of the error function, one can observe that this technique is actually doing much better in the center part of the desired band than the other techniques. This technique actually outperforms the Kaiser technique in the center half of the desired band. For applications where the bias is a more significant problem than the ripple, this technique is a significant improvement over the standard IDFT filter approach. Changing the type of frequency window from linear to something else less abrupt would be expected to improve performance further.

Sine Squared Window Function

An obvious improvement upon the linear frequency windowing technique is a window which tapers from 0 to 1 much more smoothly. A sine squared window tapering was chosen for the portion of the band outside the spectrum of the incident signal. This window is shown in Figure 21, and clearly is less abrupt than the linear window technique. As with the linear frequency window, this window is term by term multiplied with the original frequency samples, $E[k]$, to produce a new set of samples, $E_{ss}[k]$. Figure 22 shows the gain of both $E[k]$ and $E_{ss}[k]$ for comparison, they have identical phase. Once again, the antenna filter coefficients using this method are calculated by taking the inverse Fourier transform of the weighted frequency samples. The same procedure is followed for determining the gain and phase response as well as the normalized error function of the sine squared window technique. Figures 23 & 24 show the gain and phase responses of the new sine squared window technique compared with previous techniques, and Figure 25 compares their respective error functions.

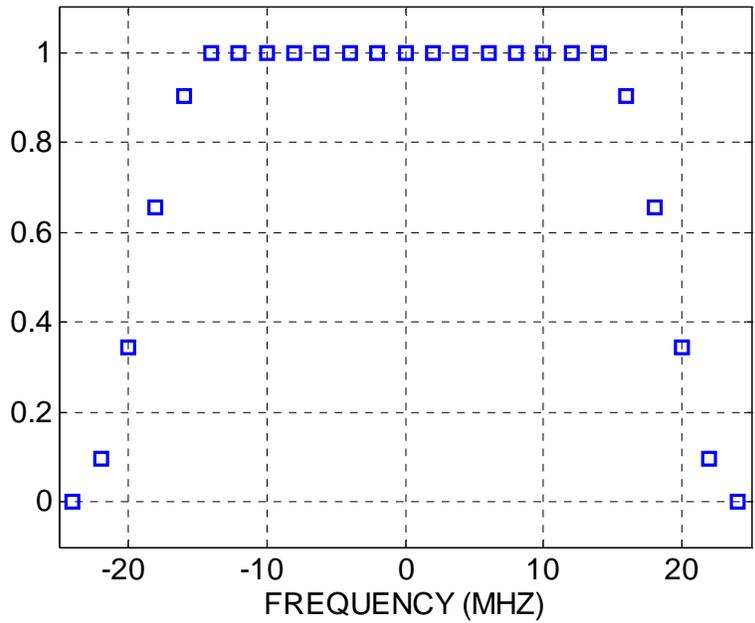


Figure 21: Sine Squared Frequency Window

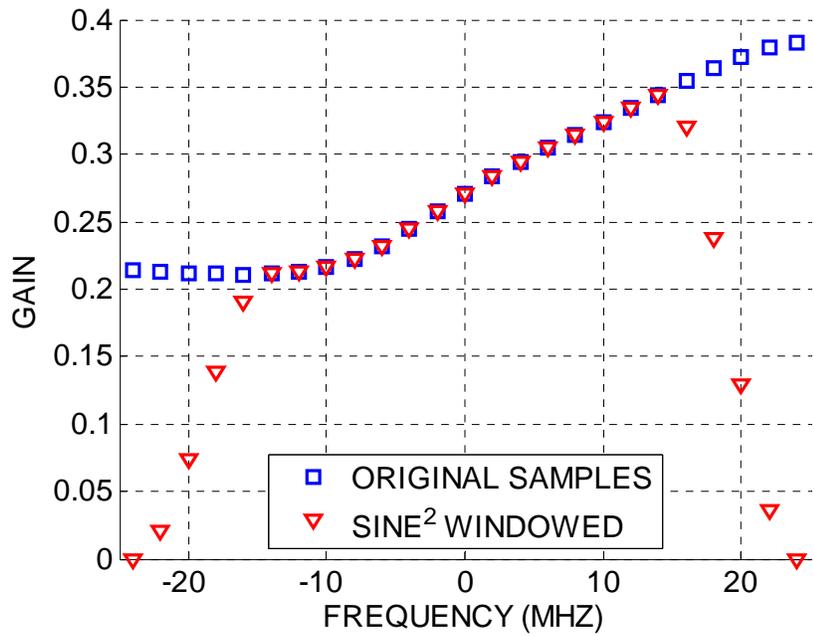


Figure 22: Sine Squared Windowed Frequency Samples

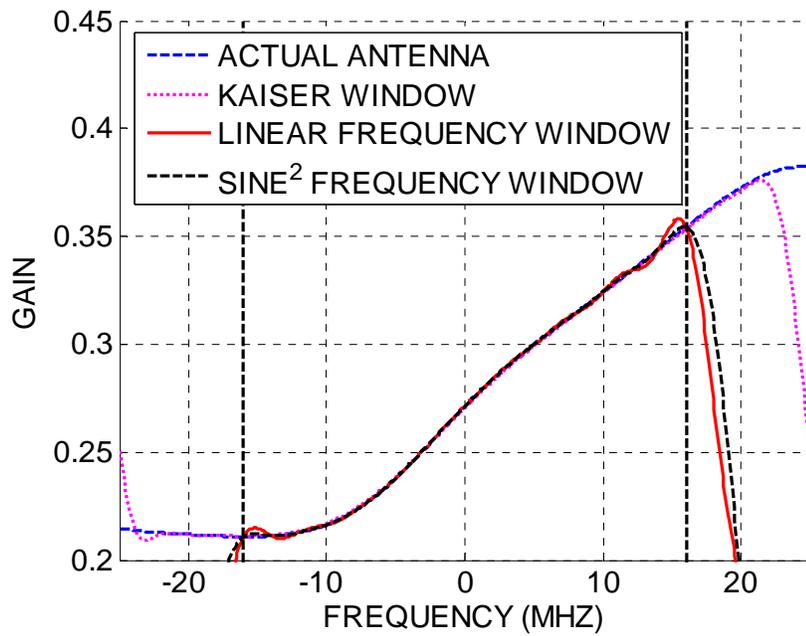


Figure 23: Gain Response of Sine Squared Frequency Technique

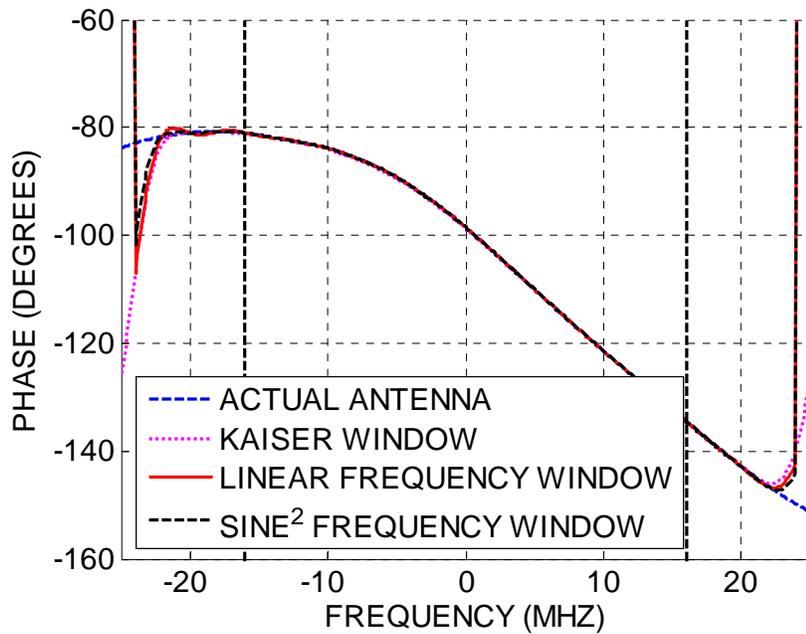


Figure 24: Phase Response of Sine Squared Frequency Window Technique

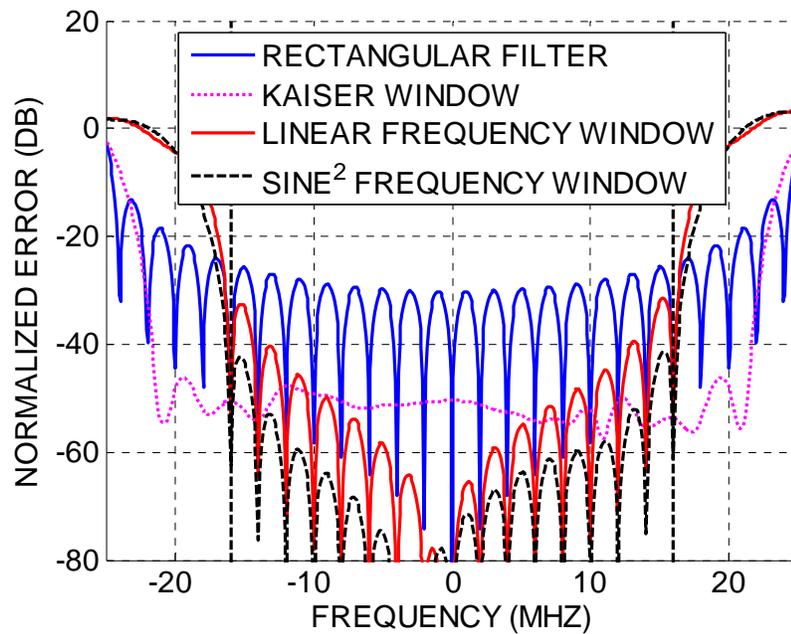


Figure 25: Normalized Error Function of Sine Squared Frequency Windowed Technique

As expected, the sine squared frequency window improved performance notably compared to the linear frequency window technique. The worst case error in the portion of the band corresponding to the incident signal is -42 dB, or 0.79%, close to the worst case -48 dB error seen by the Kaiser approach. A few observations are noteworthy based on the results in Figure 25. First, the sine squared technique significantly outperforms the Kaiser technique over the center 28 MHz of the 32 MHz band of interest. Furthermore, there is no evidence of bias using the sine squared technique, because the nulls in the error function once again coincide with the original frequency samples at f_k .

Least Squares Design Approach

Overview

The previous methods described for finding the finite impulse response (FIR) filter to represent an antenna all rely on the inverse Fourier transform approach when the number of known frequency samples is equal to the number of desired filter coefficients. However, in many instances, one may have available more frequency samples over the antenna bandwidth than number of desired filter coefficients. These additional samples can either be obtained through measurements, calculations, or interpolation. Suppose one has available M frequency samples on an antenna response over bandwidth B and desired an N length FIR filter to model the antenna response. If $M > N$, it is possible to set up M equations with N unknown filter coefficients. Furthermore, a constraint vector can be applied if one is only interested in matching the antenna frequency response over a portion of the total bandwidth, for example $0.64B$. It is well known that an efficient way to solve a system of equations with more equations than unknowns is the least squares technique. This approach was studied to determine if the filter performance can be improved over the previous techniques.

Procedure

We first must have available to us N frequency samples across the bandwidth B , and desire an N length filter, where $M > N$. In general, M should be at least twice N in order for this technique to work well in any situation. The frequency samples are stored in the vector \mathbf{H} , while the N unknown filter coefficients are denoted \mathbf{h} .

First, set up a Fourier Transform matrix relating M frequency samples to N filter coefficients. This M -by- N matrix is denoted, \mathbf{W} . The value of the element in the m^{th} row and n^{th} column is simply:

$$W[m,n] = e^{\frac{j2\pi mn}{N}} \quad (5)$$

This defines the Fourier matrix which relates N antenna frequency samples to M filter coefficients by the equation

$$H = W * h \quad (6)$$

Where H is an M -by-1 vector containing the M known frequency samples, and h is an N -by-1 vector containing the N unknown filter coefficients. This relationship relates N unknown filter coefficients to M known frequency samples across the antenna bandwidth. This means we can solve for h , the filter coefficients, in the least squares sense. Before solving this system a weighting vector is added that accounts for the interest in only a portion of the total antenna bandwidth. Namely, the center $0.64B$ (32 MHz) of the total system bandwidth (50 MHz) is the portion of the bandwidth where the antenna response must match as closely as possible. Since the incident signal is not present in the remaining 18 MHz of the antenna bandwidth, a weighting vector can be added which ignores the portion of the band where there is no incident signal. The weighting vector, \mathbf{L} , is an M -by- M diagonal matrix. All non-diagonal elements in \mathbf{L} are 0. The diagonal elements of \mathbf{L} are weighted 0 or 1 corresponding to either the undesired or desired portions of the band. The new equation becomes

$$L * H = (L * W) * h \quad (7)$$

Solving for h ,

$$h = (L * W)^+ * (L * H) \quad (8)$$

Where the $^+$ operation denotes the pseudo-inverse of a matrix. In this way, we can solve for the filter coefficients using a least squares approach.

Results

For the least squares technique for setting up the least squares approach, we start with 100 frequency samples over the same bandwidth B , with spacing Δf of 0.5 MHz, shown in Figure 26. Since there are 100 frequency samples, N is 100.

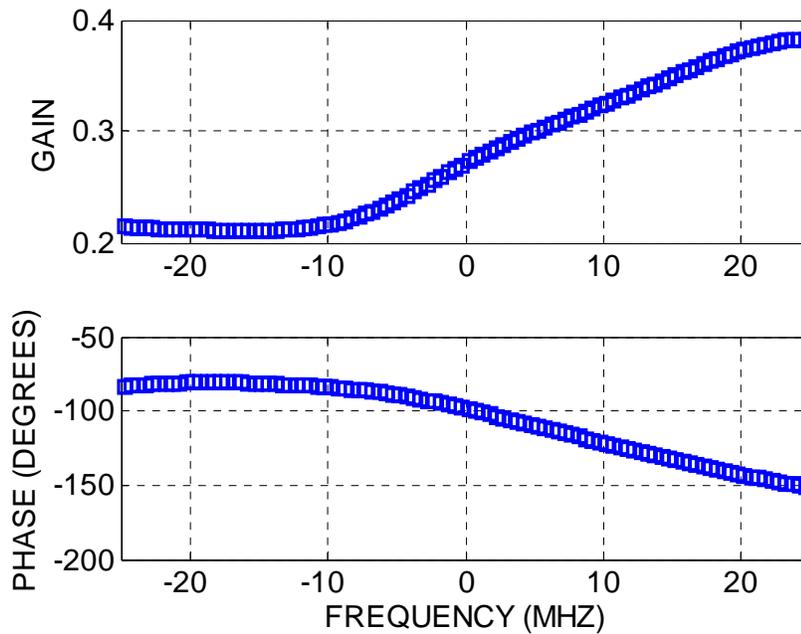


Figure 26: N=100 Frequency Samples

In the weighting matrix \mathbf{W} , all non-diagonal entries are 0. The diagonal entries correspond to a weighting vector which establishes the relative importance of certain portions of the bandwidth. Figure 27 shows the diagonal entries for the weighting vector \mathbf{W} assuming we have $N = 100$ frequency elements, $B = 50$ MHz, and the incident signal spectrum of $0.64B = 32$ MHz. Since only the portion of the filter bandwidth containing the desired signal needs to be matched, the weights outside of this region are set to zero.

In Figure 27, each weight in the center 32 MHz (the 0.64B) is set to 1, while those corresponding to the remaining portion of the band are set to 0. In this way, the diagonal entries of the weighting vector can be constructed. More elaborate scenarios involving multiple desired bands and various relative weightings can be devised depending on the application and scenario.

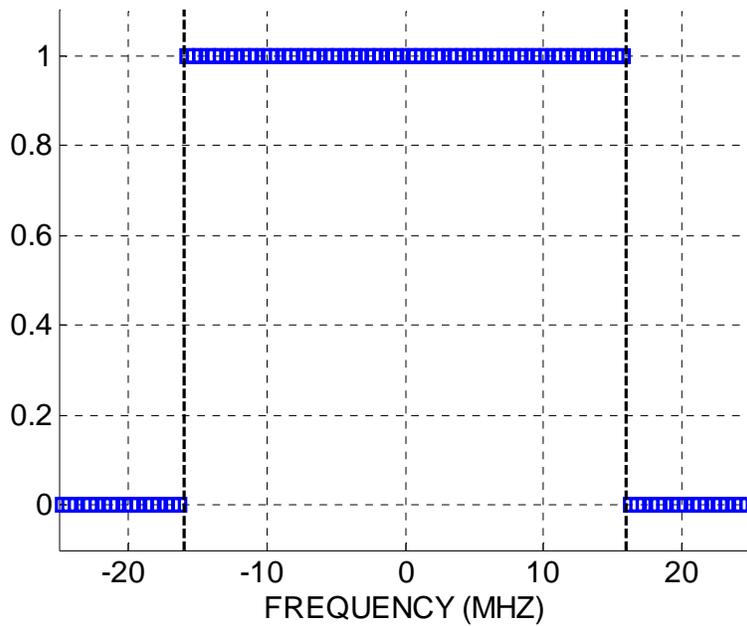


Figure 27: Weighting Matrix Diagonal Elements

After solving for the filter coefficients using the least squares approach, the gain and phase response are calculated, as well as the error function. The gain and phase responses are shown in Figures 28 and 29, while the error function is shown in Figure 30.

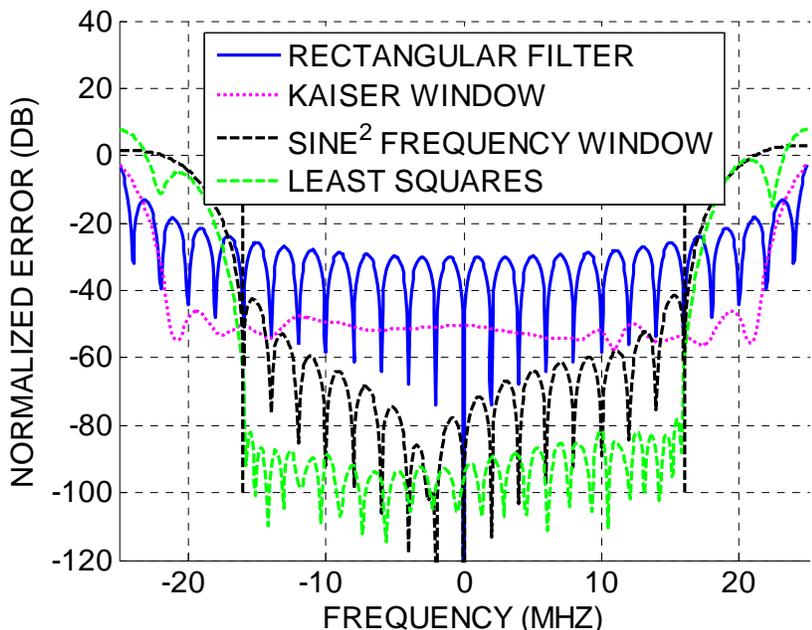


Figure 28: Least Squares Gain Response

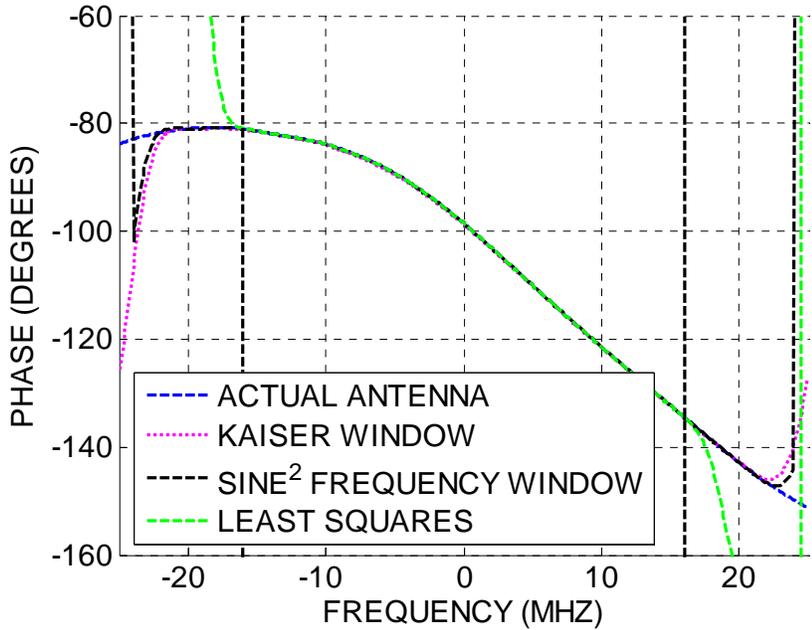


Figure 29: Least Squares Phase Response

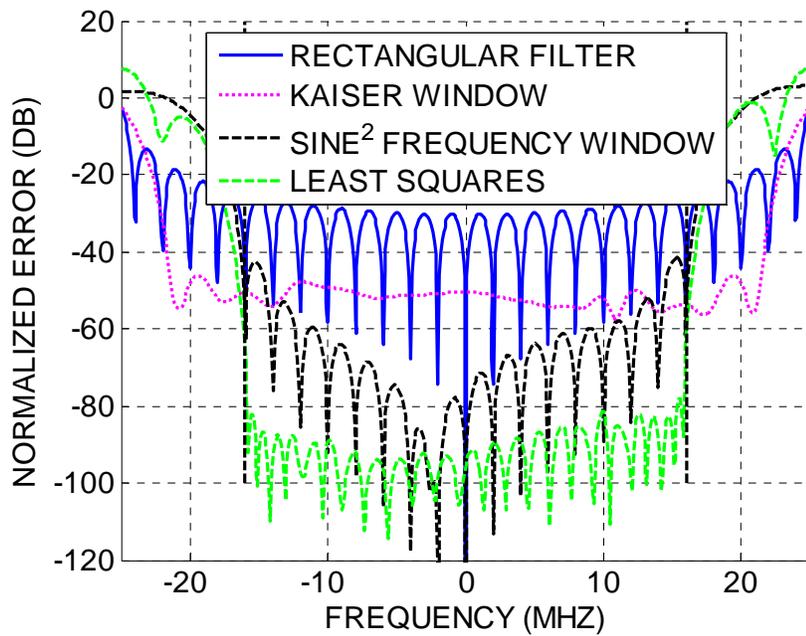


Figure 30: Least Squares Normalized Error Function

The least squares technique performs significantly better than the previous techniques. From Figure 30, the worst case error inside the desired band is approximately -78 dB, or 0.013%. In general, this technique is superior to both types of windowing techniques (time domain windowing and frequency domain windowing).

Summary

When simulating a radar/ communication system, the distortion caused by the antenna to the incoming signal must be accounted for in order to achieve accurate performance analysis. A common way to account for the antenna response in hardware-in-the-loop simulations or computer simulations is to model the antenna response as a finite impulse response (FIR) filter. For each angle of arrival of incident signals, the antenna response is different and requires a different set of FIR filter coefficients to model the antenna. Various techniques were studied and compared for generating these FIR filters to model the antenna responses.

In this study, the incident signal bandwidth was only $0.64B$ of the total antenna bandwidth B . In most applications, the antenna bandwidth exceeds the incident signal bandwidth. The filter is only required to match the response of the antenna over this portion of the total bandwidth because there is no incoming signal containing in the remaining frequency spectrum, only noise. The antenna responses were assumed to be well behaved with non-linear phase and non-uniform magnitude.

The simplest case studied was a simple Inverse Fourier Transform (IFT) approach. This approach results in a significant ripple in the frequency response of the filter when compared to the actual antenna response. The reason there was a ripple has to do with the properties of the Fourier transform. The Fourier transform of a finite-bandwidth antenna response will have an infinite time series filter for perfect reconstruction. The simple IFT approach is equivalent to truncating this infinite-length filter with a rectangular window. It is well known that rectangular windows in the time domain cause ripples and oscillations in the frequency domain.

The second technique used the observation that the simple IFT approach is equivalent to a rectangular window applied to the infinite-length filter. The motivation was that a smoother windowing operation should reduce the ripple seen in the filter frequency response. Instead of using a rectangular truncation operation, higher order windows were considered. These included a triangular window, a Hamming window, and a Kaiser window. As expected, the ripple was reduced for each of these cases, with the best case being the Kaiser window. However, a different type of error was introduced in the desired band: a so-called bias caused by a loss of energy from the filtering operation. So while the performance of the filters created with this technique was superior to the simple IFT approach, a different type of error was introduced, the “bias,” that may be less desirable than the ripple for certain applications.

The third technique was motivated by the desire to eliminate the “bias” yet still improve the filter performance at matching the antenna response over the desired band. Instead of windowing the filter coefficients, a window function was applied to the original frequency samples. When designing the frequency windows, the weighting were constructed such that only frequency samples which lay outside the desired band were tapered. Two types of frequency windows studied were a linear and sine squared technique. The windowed frequency samples were then operated on by the Fourier transform to produce filter coefficients. Once again, the filter performance in modeling the antenna was improved over the standard IFT approach. And as expected, the “bias” problem was not evident in the desired band. The ripple was reduced significantly using a frequency window as well, but the result was non-uniform over the antenna bandwidth. The performance was best closest to the center of the desired band.

The final technique considered assumed instead that one has more antenna frequency samples than desired filter coefficients. Also, if one did not have more frequency samples than filter coefficients, it was suggested that a cubic spline interpolation be used on the available frequency samples to create more. Since one can write a system of equations relating M frequency samples to N filter coefficients by the Fourier operation, this system can be solved in the least squares sense when $M > N$. Since only the desired band, $0.64B$ of the total bandwidth B , is required to match the antenna response, a weighting vector was applied to set the error tolerance outside the desired band to infinity. In other words, no consideration was given to the portion of the bandwidth which did not contain the incident signal. Using the least squares technique, further improvements in the filter performance at modeling the antenna were obtained. The worst case error in the desired band was over an order of magnitude better than any other technique.

The recommendations for the use of these techniques depend on the application and assumptions. If one has available more frequency samples than the desired number of filter coefficients, the least squares technique is recommended. If one does not have more frequency samples than filter coefficients, interpolation is recommended to generate additional samples such that this method can be used. If one is unable or unwilling to interpolate, then there exists a choice between the other two methods studied in this paper. Depending on the type of error that will most affect the simulation, choose the time domain windowing technique if ripple is a concern (convolution sensitive, adaptive processing), and choose the frequency domain technique if "bias" is a concern (received power sensitive).

Sources

- [1] Porat, Boaz. *A Course in Digital Signal Processing*. 1997.
(Equations 1-8 directly or derived from this book)