Department of Electrical and Computer Engineering

McMaster University

Experiment #1

Part (a b)

Linear phase lowpass filters (w_c = pi/3) designed via truncation (M = 20, 50, 150) clearly exhibited the Gibbs phenomenon. While increasing the filter order M significantly narrowed the transition band, the oscillatory overshoot remained constant at approximately 9%. This confirms that higher orders improve frequency selectivity but do not mitigate the ripples caused by the abrupt truncation of the ideal impulse response.

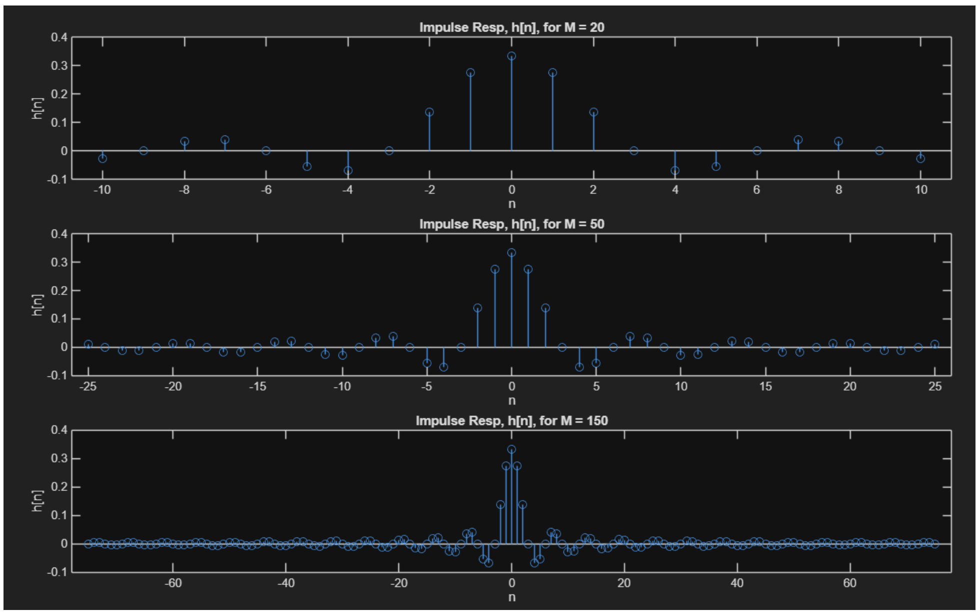

Figure 1: Impulse Responses for Different Values of M, Displaying Gibbs Phenomenon

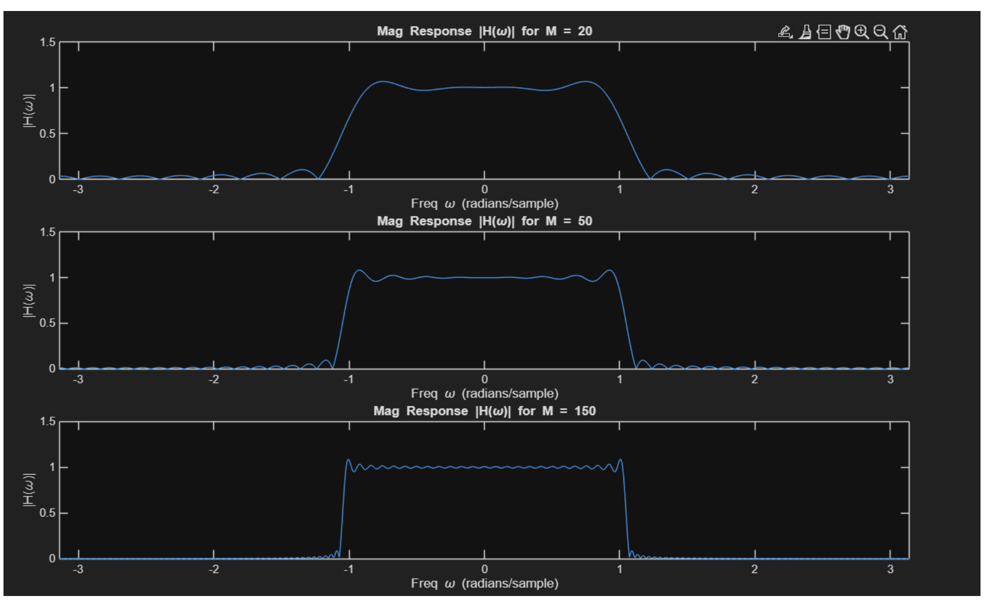

Figure 2: Mag Response for Different Values of M, Displaying Gibbs Phenomenon

Part (c)

The phase response H(w) is strictly linear across the passband for all orders, confirming a constant group delay of Tg = M/2 samples. The visible discontinuities in the plots are solely attributable to standard 2/pi phase wrapping. As expected, the phase slope steepens with increasing M, reflecting the larger delay inherent to longer filter kernels.

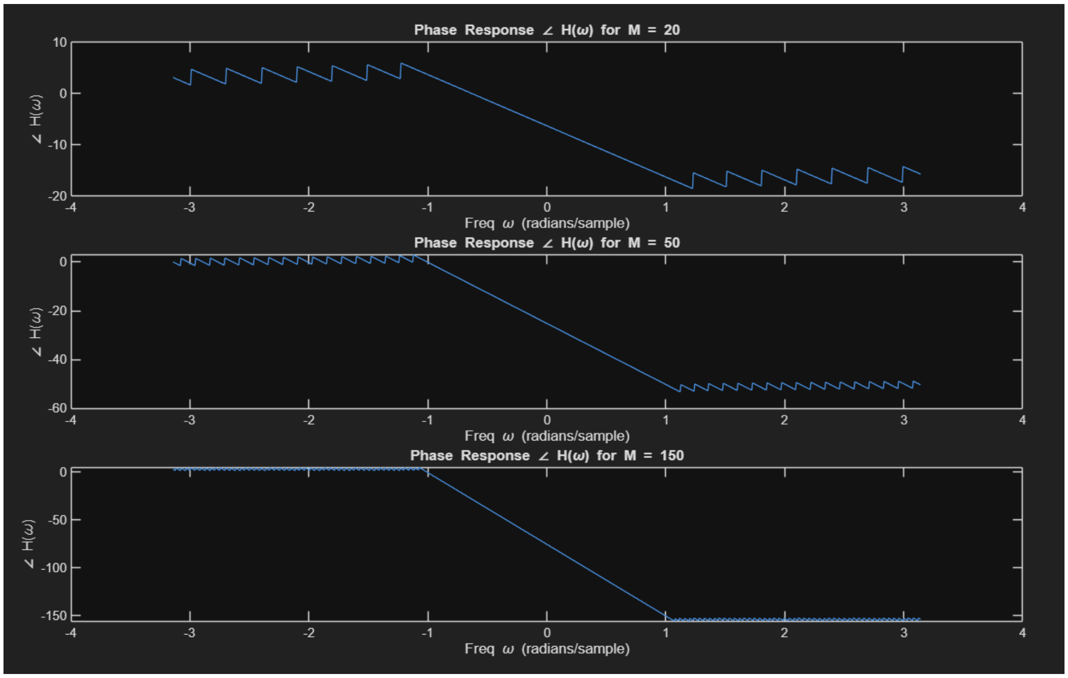

Figure 3: Phase Response Displaying Linear Phase

Part (d)

The impulse responses generated by the truncation method matched those from Matlab firls() function exactly. This empirical result verifies theoretical expectations as truncating the ideal Fourier series with a rectangular window provides the optimal approximation of the frequency response in the least mean square error sense.

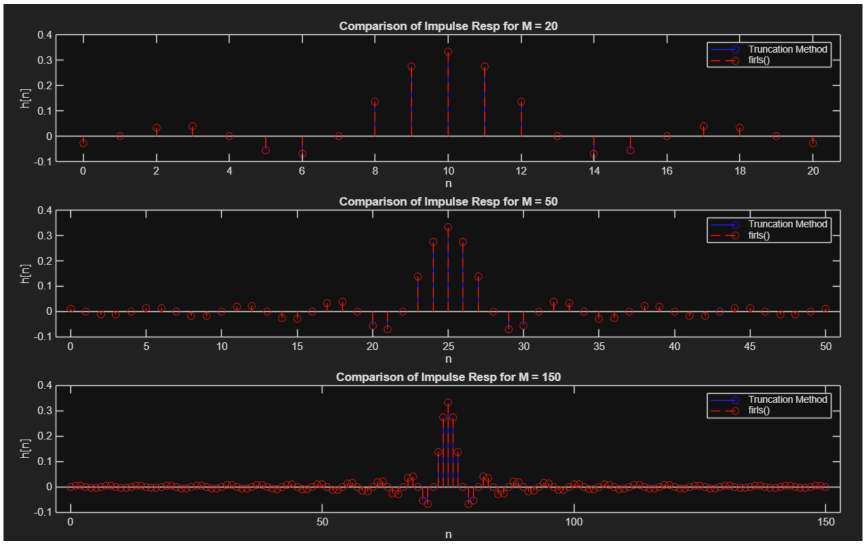

Figure 4: Identical Impulse Responses Using Truncation and firls()

Part (e)

Comparing conv() and filter() with white Gaussian noise reveals identical steady-state values but differing output dimensions. conv() computes the full convolution, yielding a length of Nsig + M that includes the final transient decay. In contrast, filter() implements the difference equation and returns a vector matching the input length (Nsig), effectively truncating the trailing transient response.

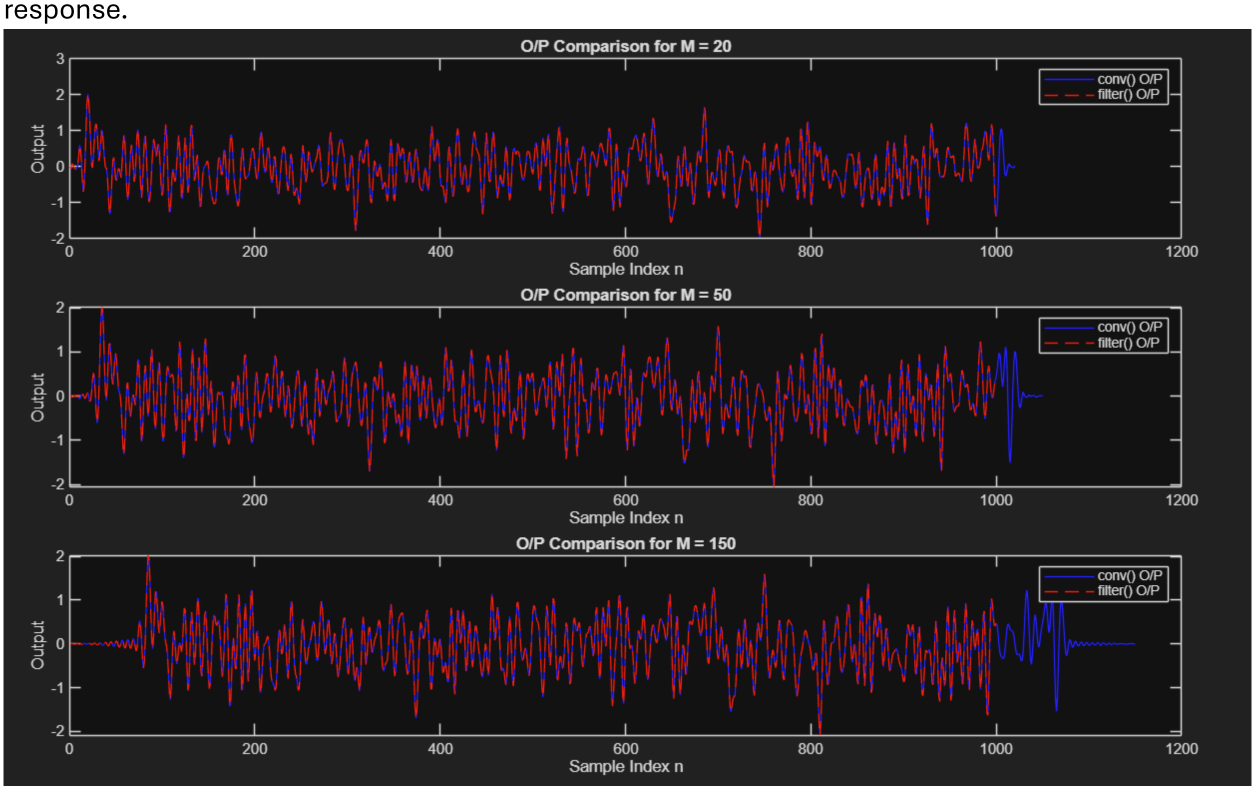

Figure 5: Comparison of Outputs of conv() and filter() Functions

Experiment #2

Part (a)

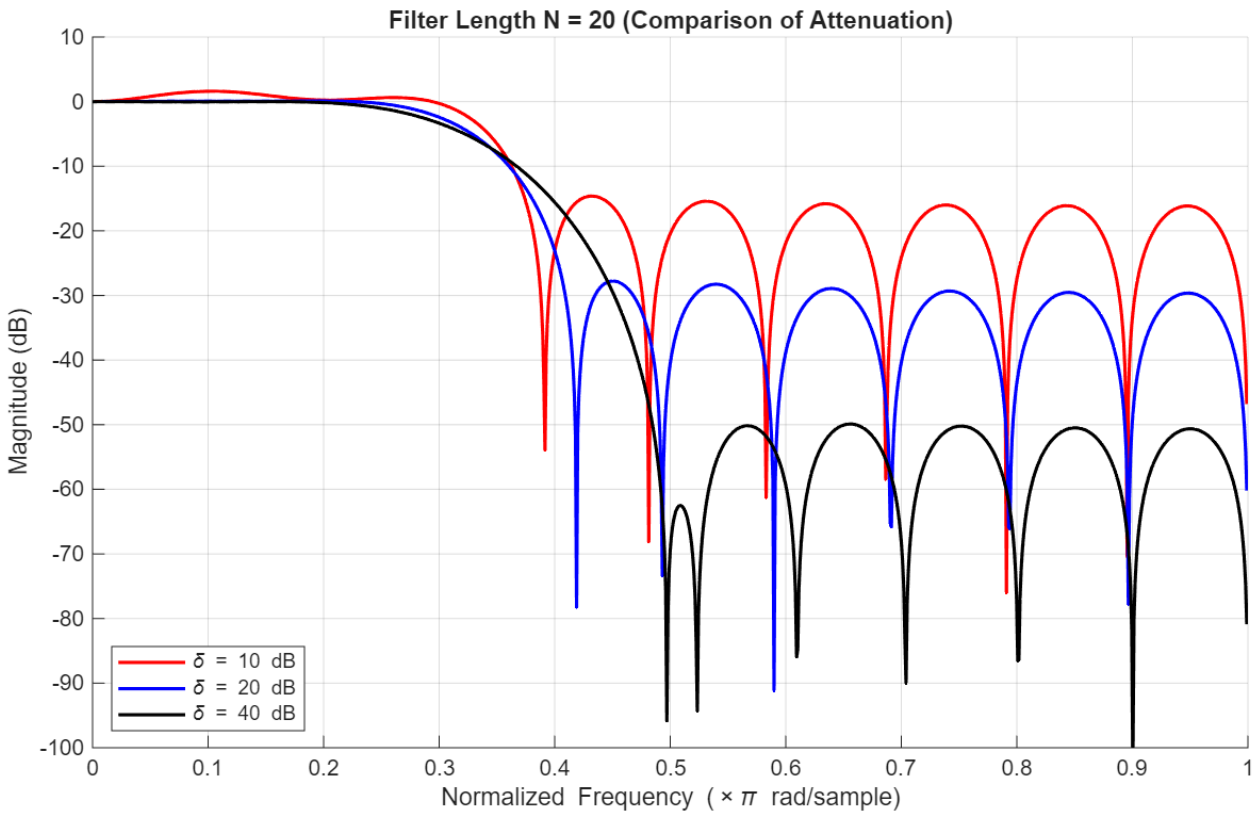

Figure 1: Plot shows that with a short filter length of 20, the transition from passband to stopband is quite gradual, and increasing the attenuation causes the transition width to widen significantly.

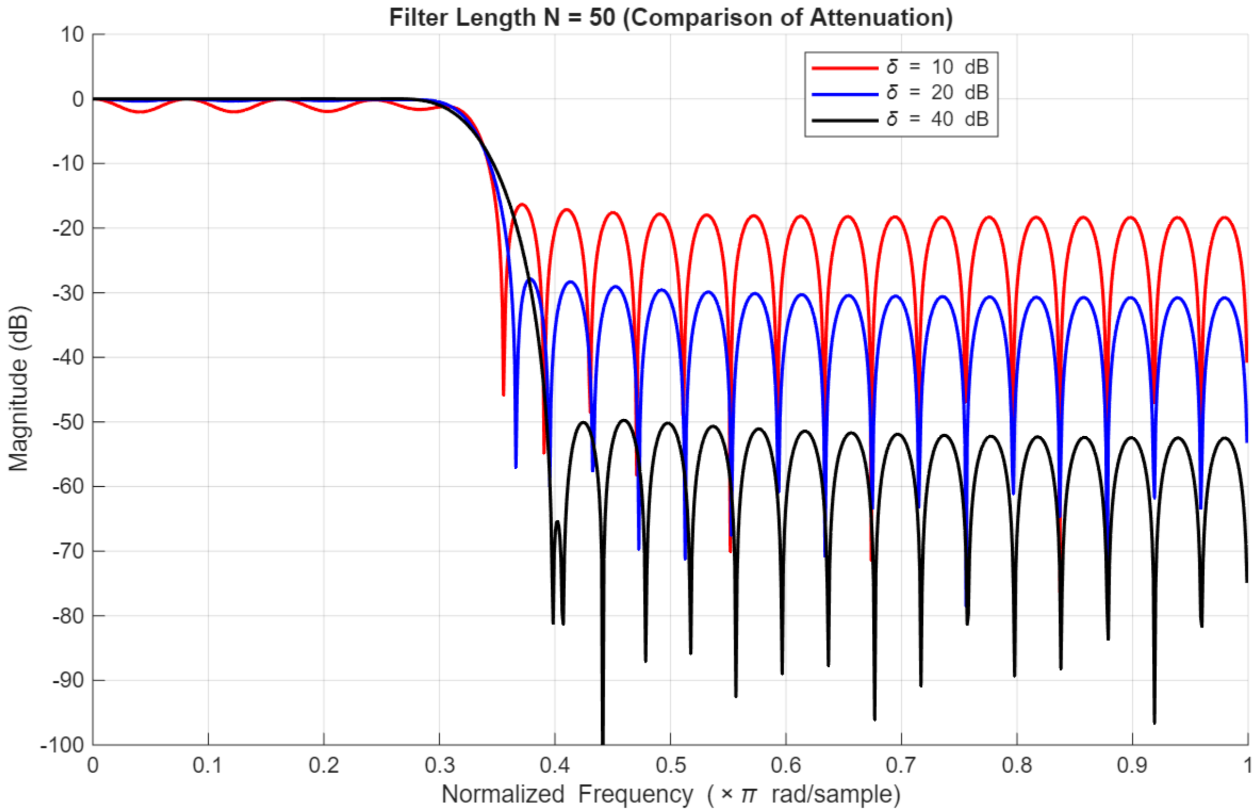

Figure 2: Increasing the filter length to 50 results in a noticeably steeper cutoff slope, though the trade off between deeper attenuation and a wider transition region is still clearly visible.

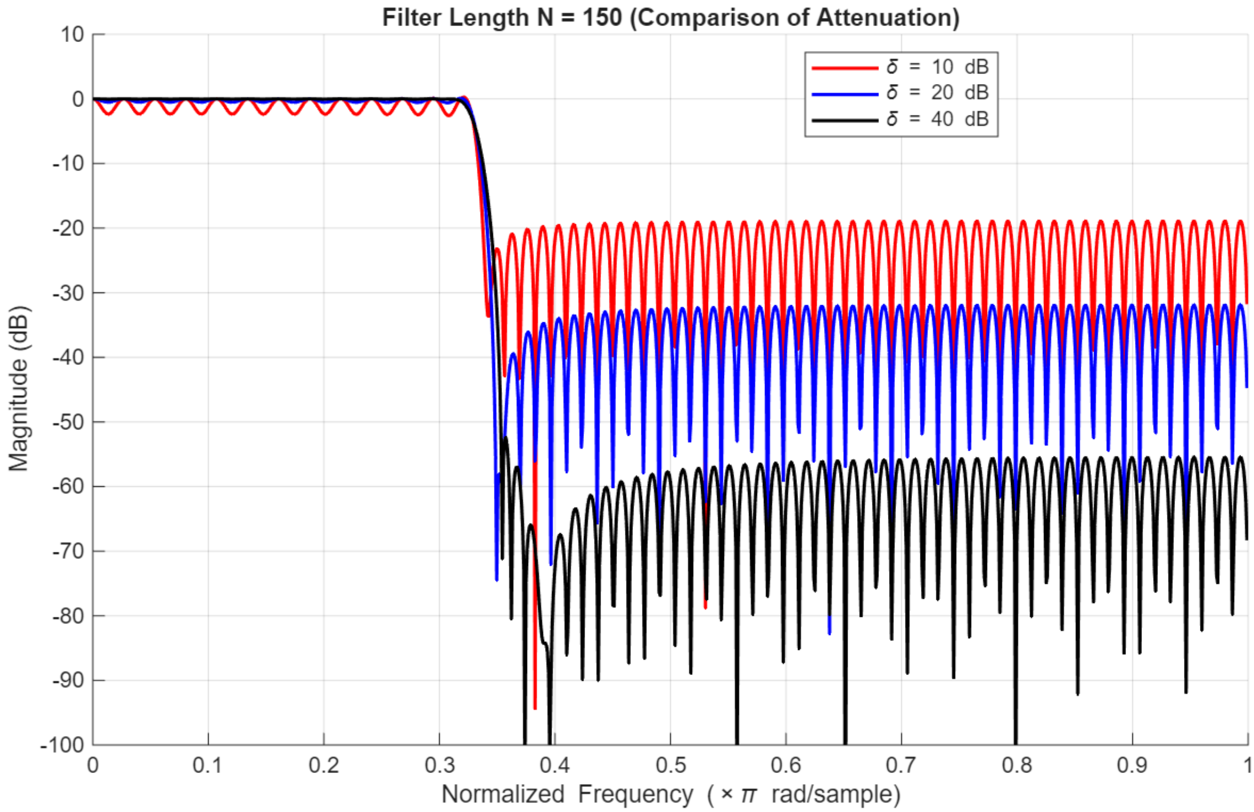

Figure 3: At a filter length of 150, the frequency response approaches an ideal shape with a very sharp cutoff, but higher attenuation settings still result in a slightly slower roll off compared to the lower attenuation settings.

i) The passband ripple and stopband attenuation: When we increase the desired sidelobe attenuation, the filter does a much better job at rejecting frequencies in the stopband and the ripples in the passband become barely noticeable. Essentially, asking for higher attenuation suppresses the "noise" in both the stopband and the passband, making the filter smoother.

ii) The slope of the transition region: The downside to increasing the attenuation is that the slope of the transition region becomes less steep. You can see in the plots that as we push the stopband lower, the curve takes longer to drop from the passband to the stopband, resulting in a wider transition width.

Part (b)

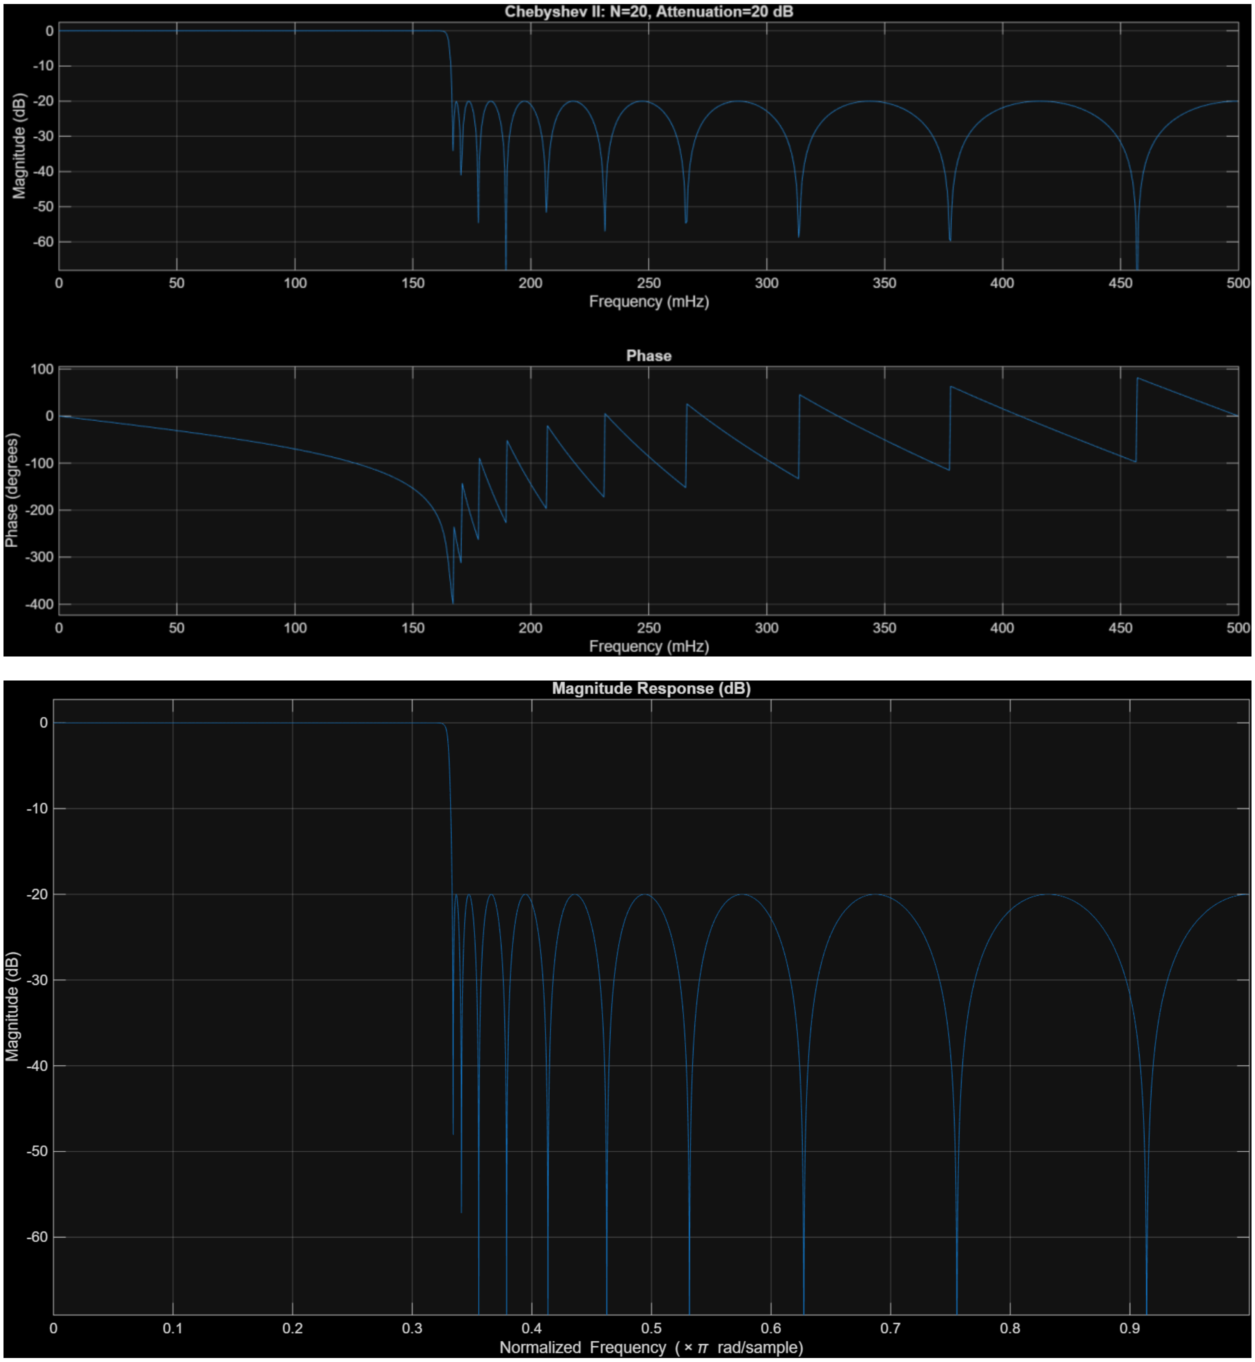

Figure 4: N=20. This filter exhibits a relatively gradual transition from the passband to the stopband, resulting in the widest transition width among the three designs.

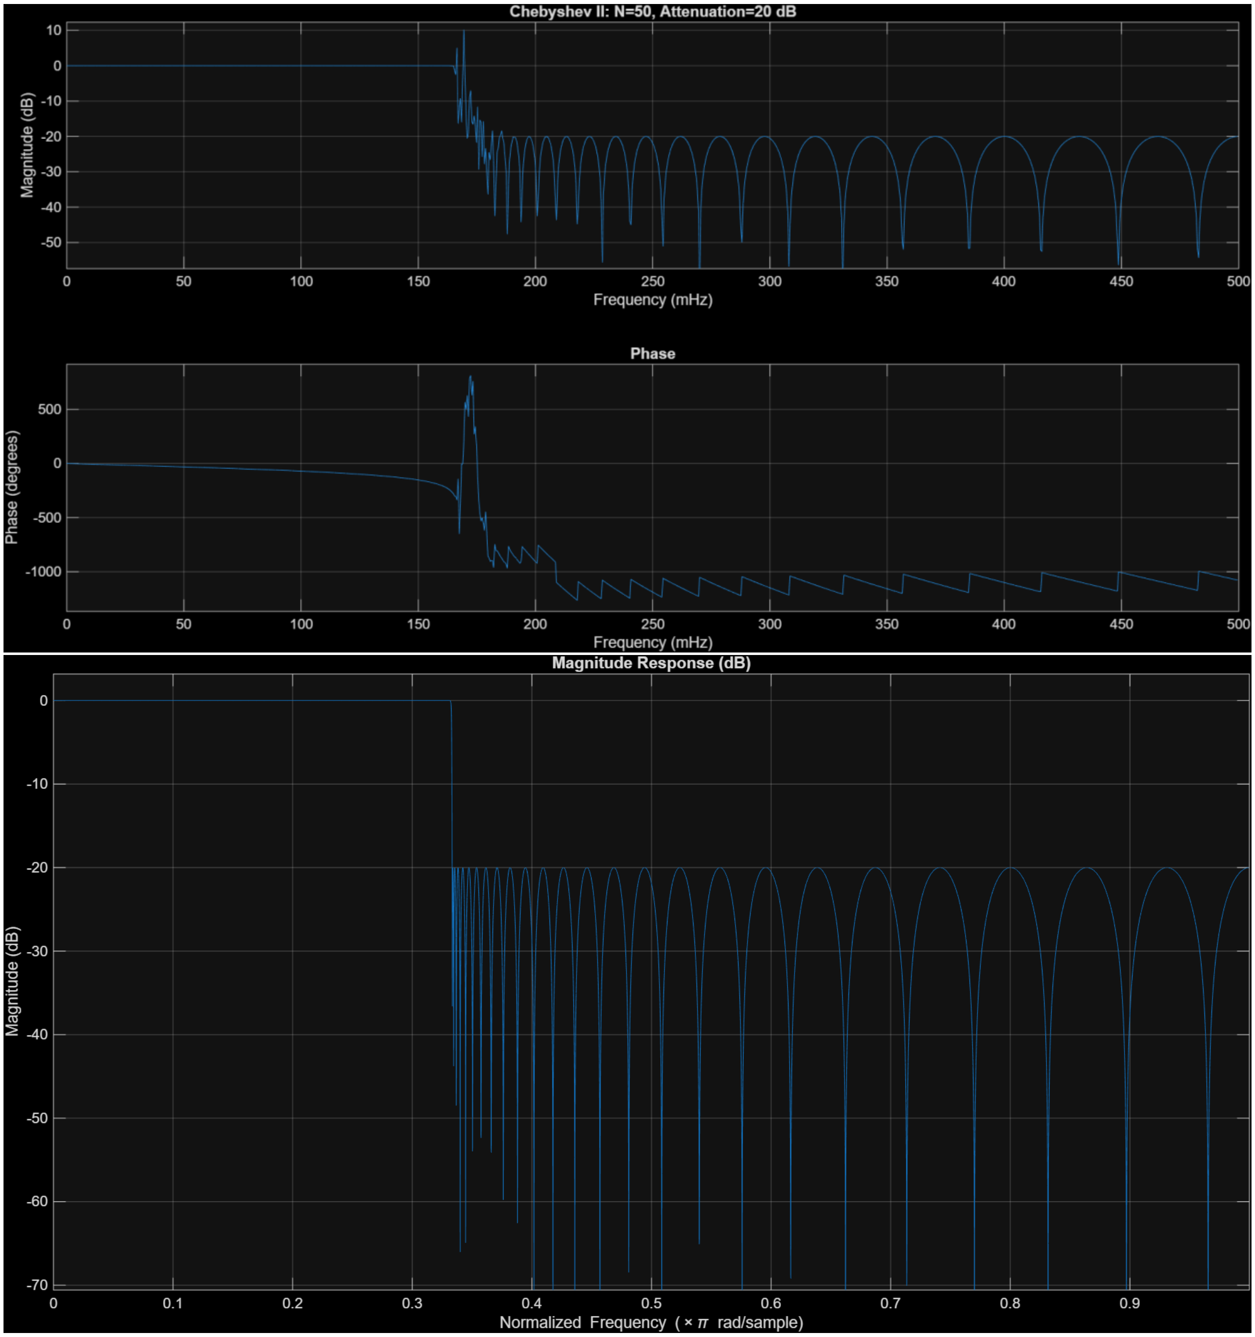

Figure 5: N=50. Increasing the filter order to 50 results in a significantly steeper cutoff slope, noticeably narrowing the transition region compared to the lower order case.

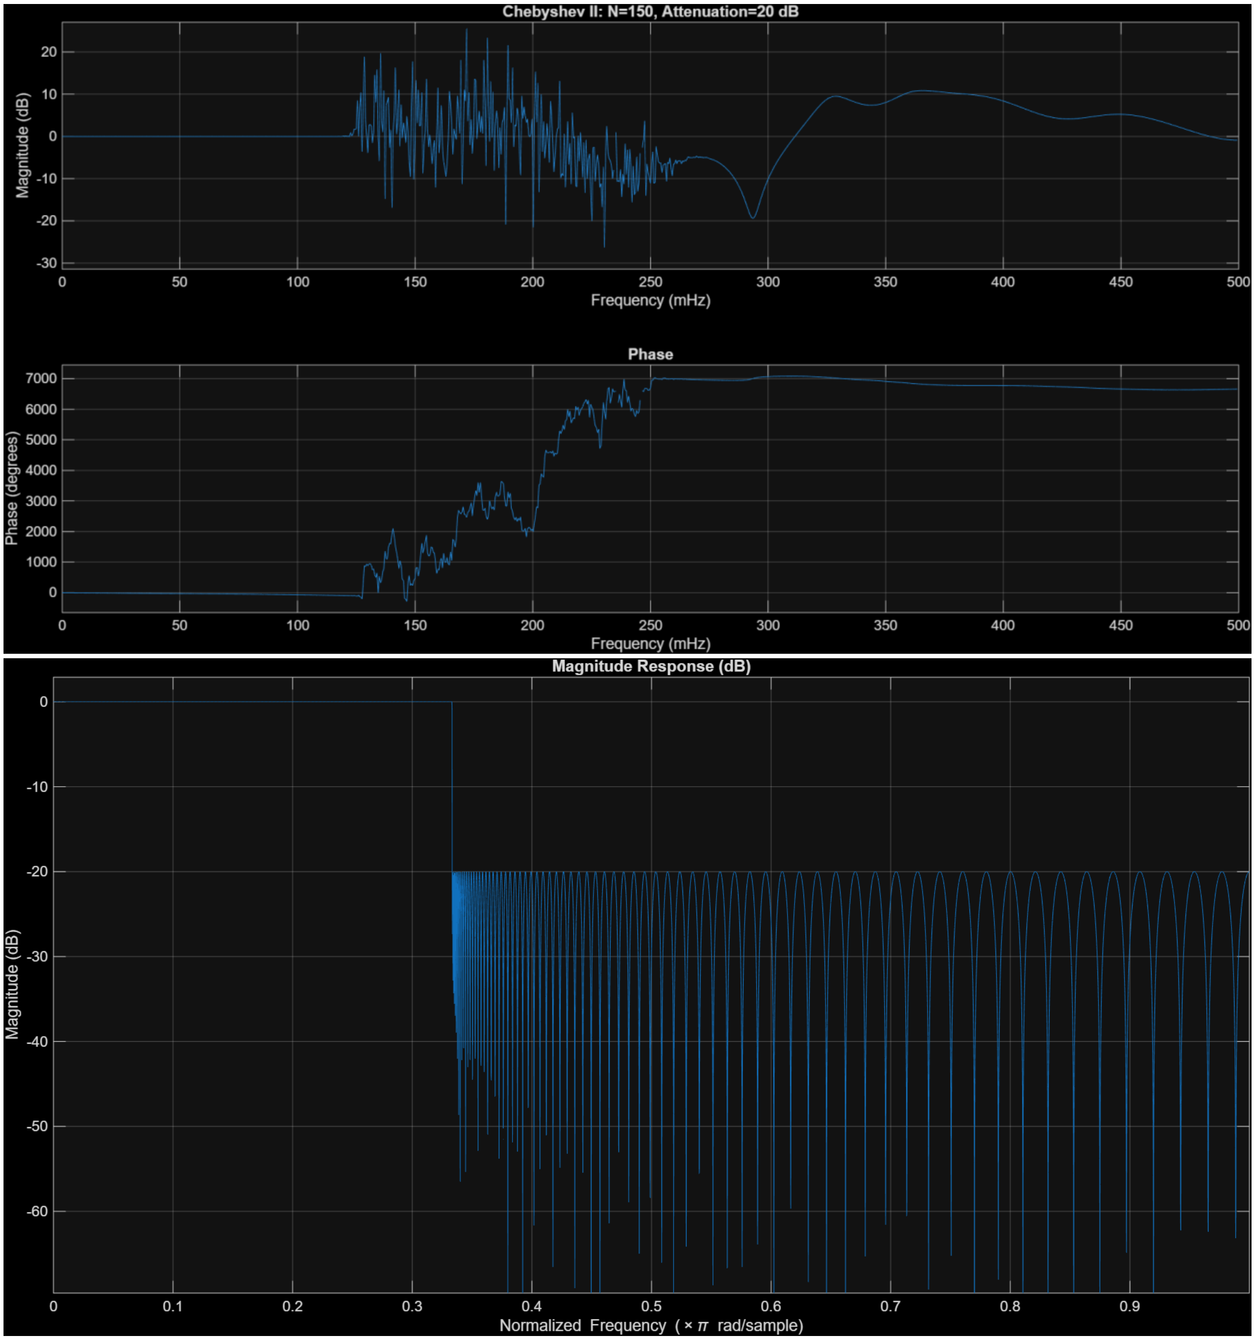

Figure 6: N=150. At this high order, the filter achieves a near-ideal "brick wall" frequency response with an extremely sharp vertical drop off at the cutoff frequency.

Part (c)

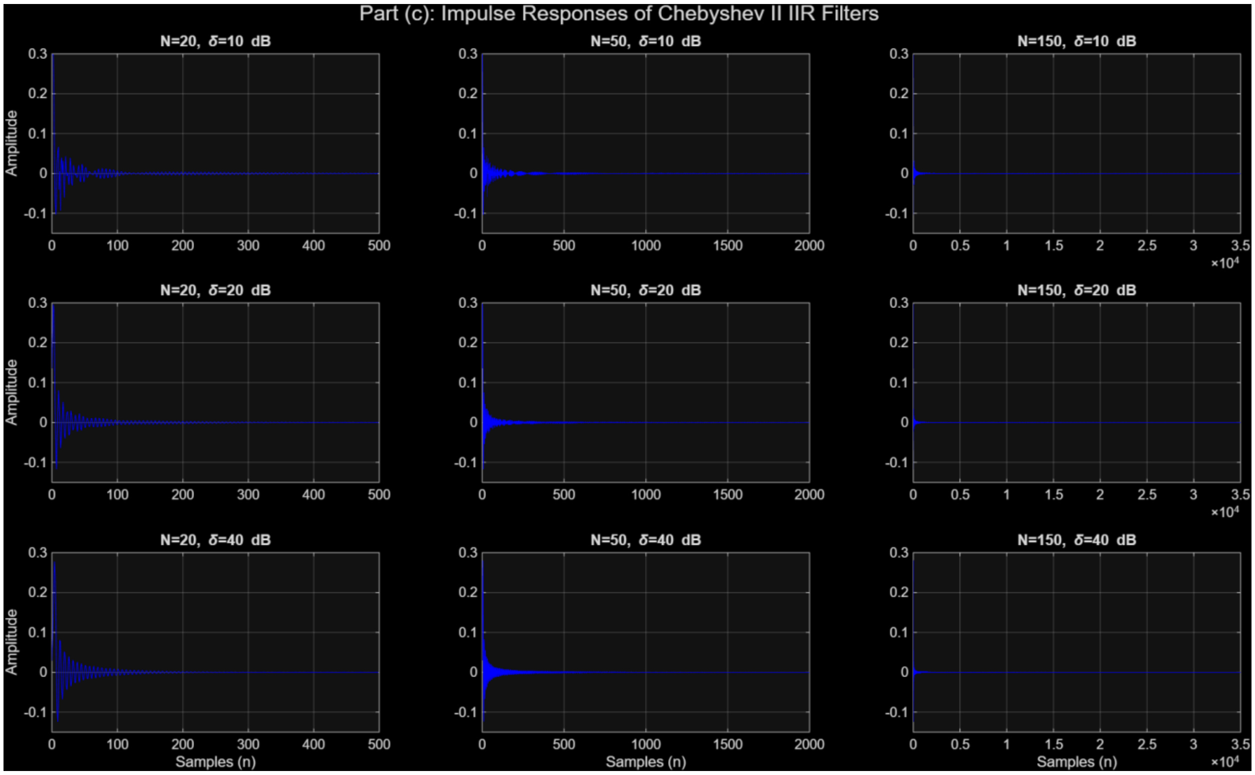

Figure 7: Compares the impulse responses of the designed Chebyshev Type II filters, with the x-axis scale adjusted in each column to capture the full settling time.

Yes, the impulse responses appear to be infinite. Unlike the FIR filters from part (a) which would drop to exactly zero after N samples, these signals show a continuous oscillation that gradually decays asymptotically toward zero but never abruptly stops. This behavior is especially obvious in the N=150 plots, where the signal is still non zero and oscillating long after 30,000 samples, confirming the "Infinite" nature of the IIR design.

Part (d)

When comparing the magnitude responses, the IIR filters are much more efficient than the FIR filters. For the same filter order, the Chebyshev II IIR design produces a significantly sharper transition from the passband to the stopband. You can see this clearly at order 50, where the IIR filter looks almost like a vertical wall, while the FIR filter still has a gradual slope.

The major difference lies in the phase response. The FIR filters show a perfectly straight line, which indicates linear phase. This is important because it means all frequencies are delayed by the same amount, keeping the signal shape intact. The IIR filters, however, show a curved phase response, which introduces phase distortion. This aligns with filter theory, which states that causal IIR filters cannot have linear phase, whereas FIR filters can easily achieve it by having symmetric coefficients.

Part (e)

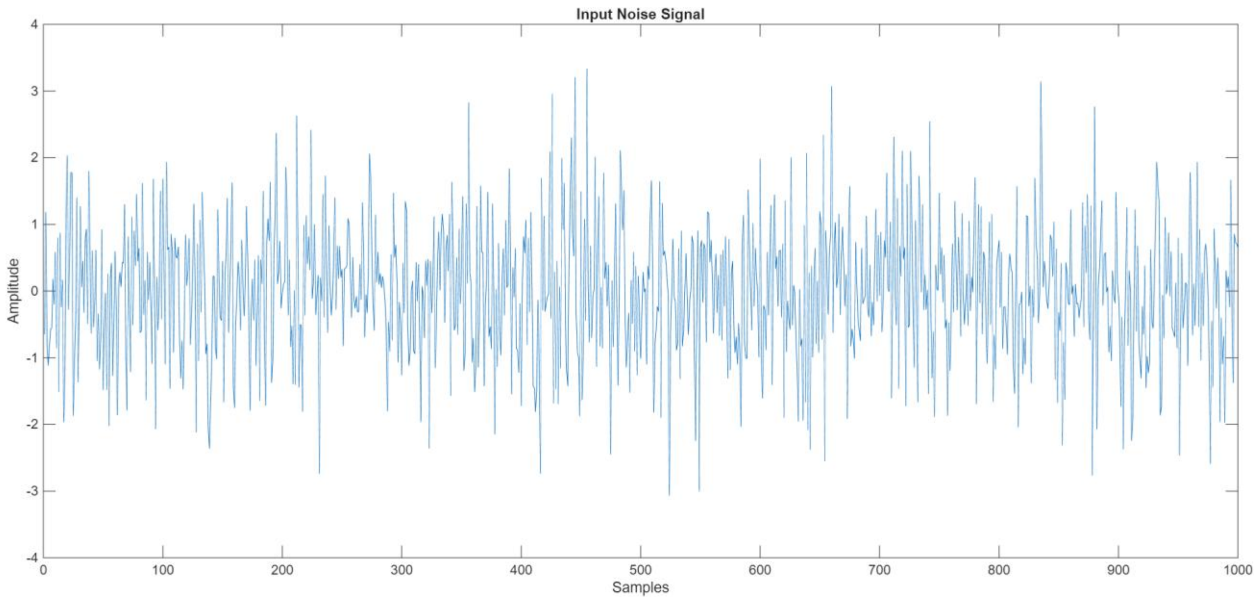

Figure 8: This plot displays the 1000 sample sequence of white Gaussian noise used as the input signal for both filters.

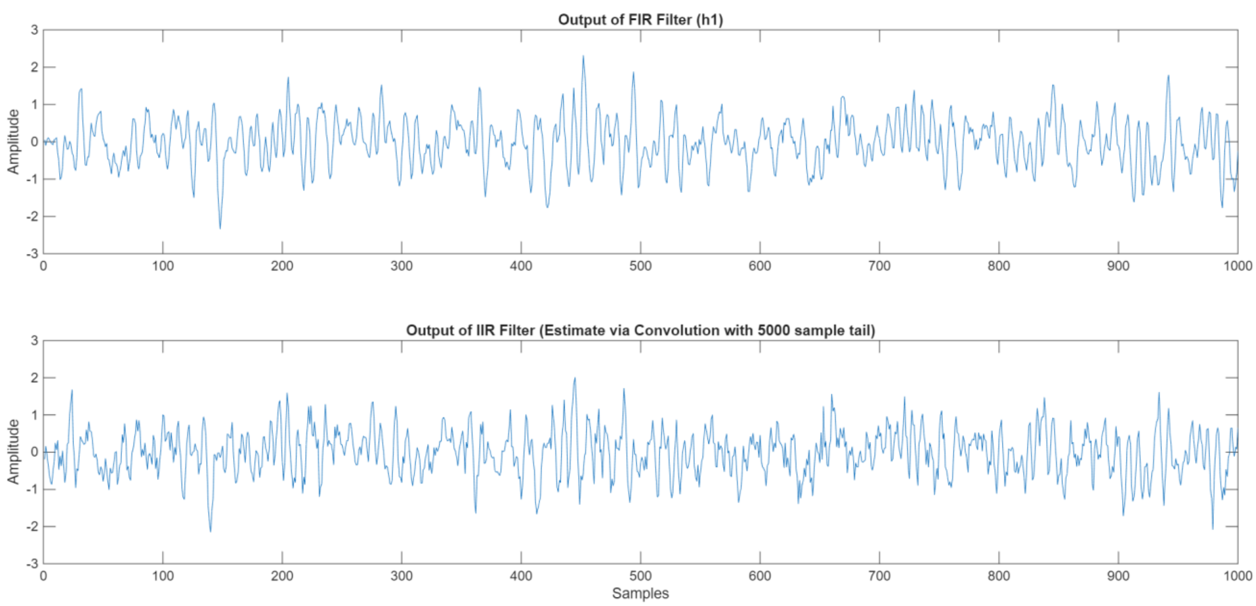

Figure 9: The top subplot shows the output of the FIR filter designed in part (a). The bottom subplot displays the estimated output of the Chebyshev II IIR filter, calculated by convolving the input noise with a truncated 5000 sample impulse response to approximate the filter's infinite nature.

The estimate is excellent and virtually indistinguishable from the true IIR response. Since the impulse response of a stable IIR filter decays over time, utilizing a "long tail" of 5000 samples captures nearly all of the filter's significant energy. Consequently, the convolution result provides a highly accurate approximation, with the error due to truncation being negligible.

Experiment #3

Part (iD)

The oscilloscope shows a massive, high amplitude signal caused by reflections from the stationary metal barriers. Because these barriers are large and extend across the entire detection range, they reflect significantly more radar energy than the targets. This creates an extremely low Signal to Clutter Ratio (SCR), effectively burying the weaker signals from the moving vehicles under the dominant clutter, making them impossible to distinguish without filtering.

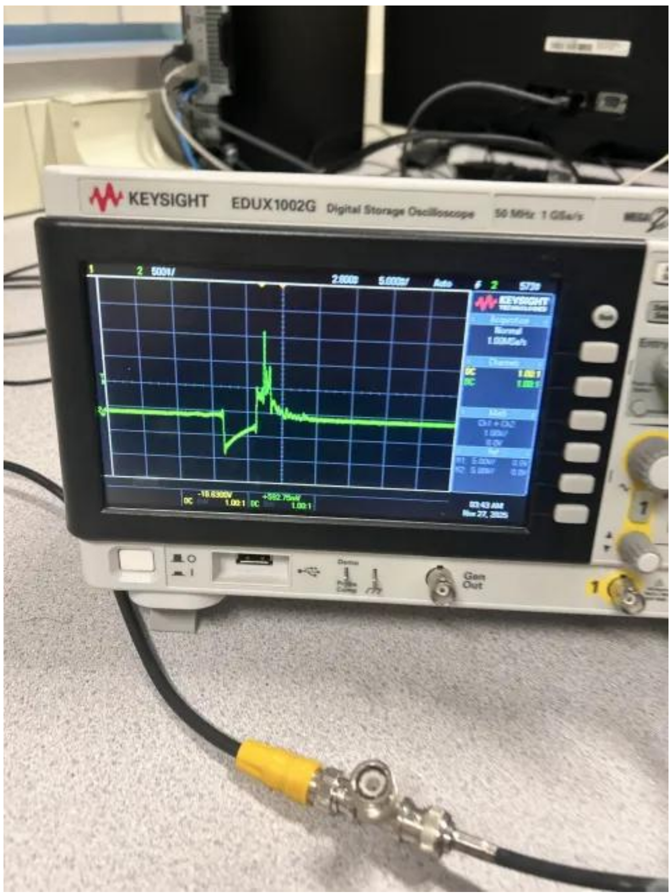

Figure 1: Oscilloscope plot from Lab 4 code

Part (iii)

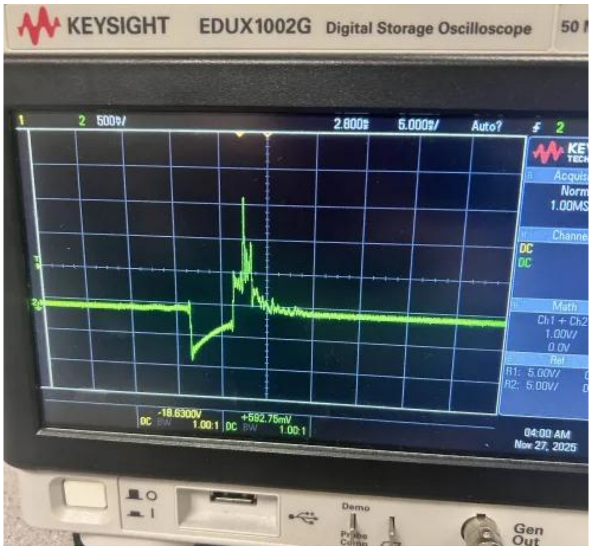

Figure 2: Oscilloscope plot with USE_MTI = 0

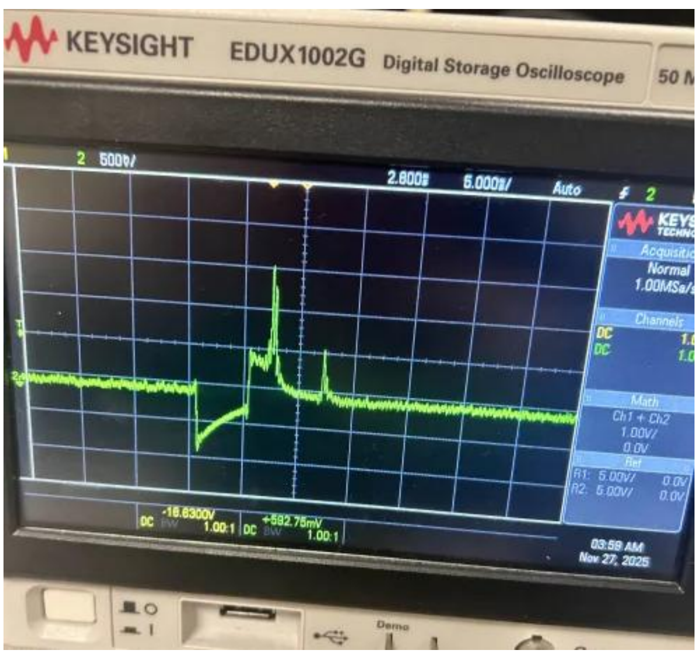

Figure 3: Oscilloscope plot with USE_MTI = 1

When USE_MTI is set to 0, the oscilloscope displays a high amplitude, jagged waveform dominated by strong reflections from the stationary barriers. Switching USE_MTI to 1 activates the Two Pulse Canceller, which drastically reduces the overall signal amplitude by removing these static peaks. This occurs because the canceller acts as a high pass filter, performing the operation $y[n] = x[n] - x[n-1]$. Since the barriers are stationary (0 Hz Doppler shift), their returns are identical in consecutive sweeps and effectively cancel each other out, revealing the moving targets that were previously masked.

Part (iv)

8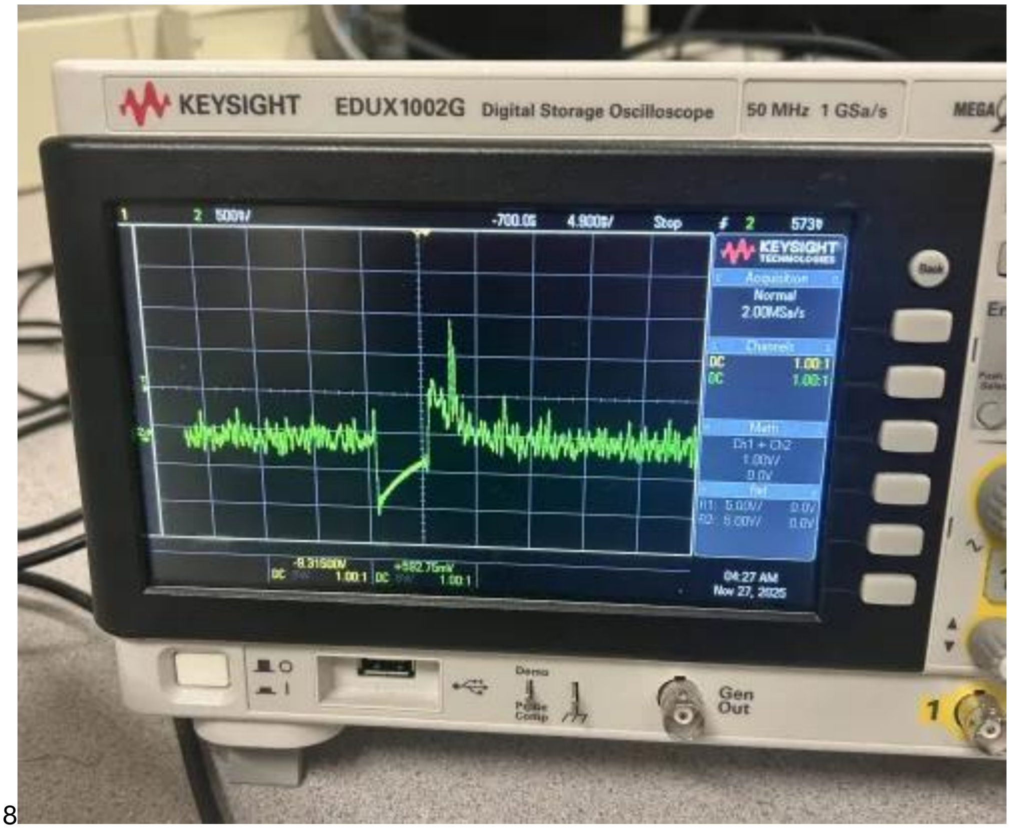

Figure 4: Three Pulse Canceller

Part (v)

The MATLAB simulation compares the frequency response of both filters. While the Two Pulse Canceller has a sharp V shaped notch at 0 Hz, the Three Pulse Canceller exhibits a wider, U-shaped notch. In this scenario, the Three Pulse Canceller proved more effective. Real world clutter often has a slight spectral spread around 0 Hz rather than being a single DC value. The wider stopband of the Three Pulse filter suppresses these low frequency components more thoroughly. This was verified on the oscilloscope in part (iv), which showed a cleaner signal with less residual clutter noise compared to the Two Pulse result.

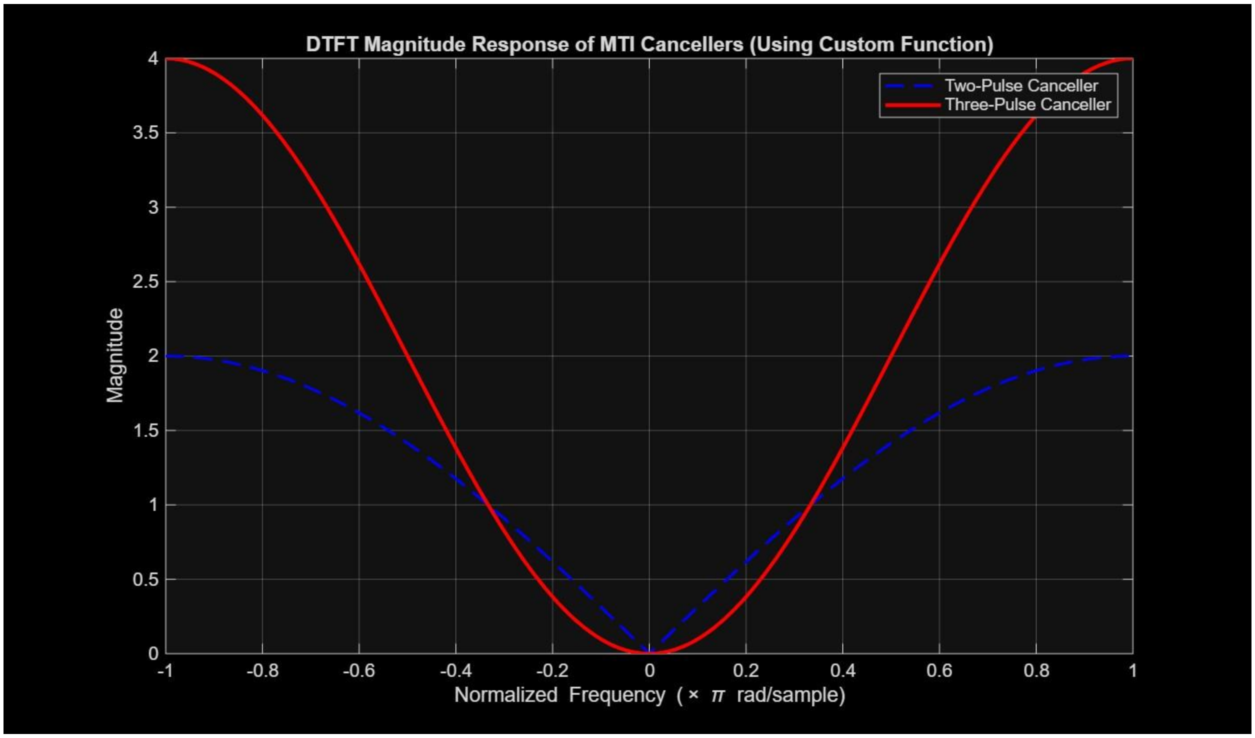

Figure 5: Two Pulse and Three Pulse Canceller Responses

Part (vi)

We modified the prepare_fft_data function to apply a Hanning window to the signal before FFT processing. This resulted in sharper, more distinct target peaks and reduced ripples, or side lobes, in the noise floor. The Hanning window works by tapering the signal at the edges, which minimizes spectral leakage. Without this window, energy from strong signals leaks into adjacent frequencies, potentially masking weaker targets or creating false detections.

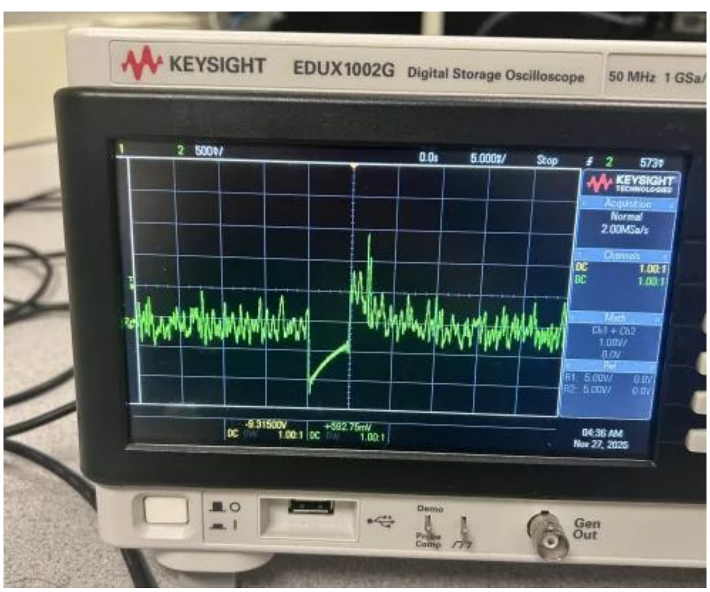

Figure 6: Oscilloscope Plot for Hanning Window

Part (vii)

If the radar platform were moving, the stationary barriers would appear to move relative to the radar, causing their clutter signal to shift away from 0 Hz (DC). A standard MTI filter, which places a notch at 0 Hz, would fail to filter this moving clutter. To fix this, we would do the following: calculate the expected Doppler shift caused by the platform's velocity and shift the center frequency of the MTI filter to match it. This ensures the clutter falls into the filter's stopband. This is a band stop filter.