Comparison of FIR and IIR Filters in Coherent Lidar Processing

Dr. Scott Shald
Lockheed Martin Coherent Technologies
135 South Taylor Ave.
Louisville, CO 80027
scott.shald@lmco.com

Introduction

The architect of a remote sensing system is always faced with the question of how to deal with sensor noise. The architect's goal is to produce a system that reliably detects a target signal even when it is relatively weak, while avoiding the generation of an overwhelming number of false alarms. Reliable detection and minimal false-alarms are competing interests which the system architect has to balance. This is just as true in coherent detection systems as it is in any other remote sensing system.

A common method for dealing with noise involves two steps: 1) averaging, and 2) threshold detection. In the first step, intermediate signal data are averaged over multiple realizations to reduce the fluctuations caused by noise. In a pulsed coherent sensor, averaging is often performed on the power spectra that are generated from each transmit pulse. In the second step, the averaged data are compared to a threshold, and any values that exceed the threshold are declared detections.

The type of averaging used in the above process will impact the detection performance. This work compares two types of averaging, one (a 'boxcar' finite impulse response (FIR) filter) that has been used often in the past and is well characterized and another (a first-order infinite impulse response (IIR) filter) that is simple to implement but has not been well characterized. The goal of this work is to characterize the detection performance when using an IIR filter for averaging and to determine the advantages and disadvantages of this approach in comparison to an FIR filter.

Signal Averaging

Two types of averaging are examined here. The first is a simple 'boxcar' FIR filter, which is described by

$$y_n = \frac{1}{N}\sum_{k=0}^{N-1} x_{n-k} \ . \tag{1}$$

y_n is the simple average of the last N values of the sequence $\langle x_k \rangle$. The level of averaging, N, is an important parameter and is usually selected to achieve a desired performance level. The resulting detection performance when using an FIR filter in a coherent lidar has been well characterized in prior works and will be summarized below.

The other type of averaging examined in this work is a simple first-order IIR filter, described by

$$y_n = (1-\alpha)x_n + \alpha y_{n-1}, \tag{2}$$

where α is a weighting parameter between 0 and 1. This filter calculates a new value for y as the weighted sum of the newest value of x and the most recent value of y. It is simple to implement because it is completely determined by the current measurement, the previous average, and the weighting parameter.

Detection Performance Characterization

This work compares the detection performance in a coherent lidar between an IIR filter and an FIR filter. Detection performance is characterized by probability of detection (P_D), probability of false alarm (P_{FA}), and CNR (and the relation among these quantities). A clairvoyant detection process is used for the comparison, in which the detection threshold is determined from the a priori knowledge of the noise and signal distributions. In particular, this work addresses the following questions:

1. How are the averaging parameters (N for FIR, α for IIR) for the two averaging techniques related to each other?
2. How is the threshold calculated for each averaging approach?
3. How do the IIR filter and FIR filter compare in detection performance? What are the advantages of each approach?

The detection statistics for the case of the FIR filter are well known. However, the detection statistics for the IIR filter are not well known and are derived below. Once the statistical descriptions of each approach are obtained, the detection performance easily follows.

Assumptions

This work examines a statistical question arising when averaging signal power spectra in a pulsed coherent lidar. When shot noise is dominant, each noise bin (a bin containing only shot noise) of such a spectrum is exponentially distributed. This work assumes in addition that the data are normalized so that the mean value of a noise bin is 1.

In addition, this works assumes that the power spectrum of a signal bin (a bin containing the target signal and shot noise) is also exponentially distributed (with a mean of *CNR*+1) and independent from one pulse to the next. This type of behavior occurs for some targets, but not for all. Still, it is expected that the lessons learned in this work will also apply to other cases.

Summary of FIR Filter Results

As mentioned above, analysis of the FIR case has already been published[1]. The behavior of an averaged noise bin and an averaged signal bin are well known, and are briefly reviewed here.

A noise bin of the spectrum contains only shot noise. The averaged noise spectrum is just the average of *N* exponentially distributed values, each with mean of 1. Therefore, the averaged value is gamma distributed with mean 1 and diversity *N*.

The average of a signal bin behaves much like the average of a noise bin. The averaged signal, when using an FIR filter, is gamma distributed with mean of *CNR*+1 and diversity of *N*.

IIR Filter Results

Since the IIR filter case has not been previously examined, the distribution of noise and signal bins must be derived here. Then the detection performance can be evaluated.

Distribution of Averaged Power Spectrum

The distribution of the average is derived by examining Equation 2, the averaging equation for the IIR filter. Two approaches are taken to examine the resulting probability distribution: 1) the characteristic function, and 2) the first few moments. Neither approach produces a simple analytic solution for the distribution, but an approximation is found.

Characteristic Function

Assume that enough averaging has been performed so that y_n has converged to a stationary distribution (y_n has the same distribution as y_{n-1}). Using the properties of characteristic functions, we find

$$\phi_y(t) = \phi_x\big((1-\alpha)t\big)\phi_y(\alpha t), \tag{3}$$

where ϕ_x and ϕ_y are the characteristic functions of *x* and *y*. When *x* represents a noise bin, it will be exponentially distributed with mean of 1, so $\phi_x(t) = 1/(1-it)$ and

$$\phi_y(t) = \frac{1}{1-i(1-\alpha)t}\phi_y(\alpha t). \tag{4}$$

No simple analytic solution for this equation has been found. However, it quickly becomes clear that *y* cannot be gamma distributed, since the characteristic function for a gamma distribution does not fit the above equation. Beyond this realization, there is little more to learn about the distribution of *y* from this approach.

Moments

More can be learned about *y* by examining the first few moments. The first moment of *y* can be found by calculating the expectation of Equation 2, again assuming that the averaging has converged to a stationary distribution for *y*:

$$\begin{aligned} \mu_{y,1} &= (1-\alpha)\mu_{x,1} + \alpha\mu_{y,1} \\ \Rightarrow \mu_{y,1} &= \mu_{x,1} \end{aligned} \tag{5}$$

where $\mu_{y,n}$ is used to represent the *n*-th non-central moment of *y*. This result shows that the mean of *y* is equal to the mean of *x*. This occurs because of the choice of coefficients in (2). The second moment can be found by taking the expectation of the square of Equation 2. The details are omitted, but the result is

$$\mu_{y,2} = \frac{2}{1+\alpha}\mu_{x,1}^2. \tag{6}$$

The variance of *y* is found to be

$$\sigma_y^2 = \frac{1-\alpha}{1+\alpha}\mu_{x,1}^2. \tag{7}$$

The third moment is found in the same manner:

$$\mu_{y,3} = \frac{6\mu_{x,1}^3}{(1+\alpha)(1+\alpha+\alpha^2)} \tag{8}$$

Fit to Gamma Distribution

Although it was shown that *y* cannot have a gamma distribution, it is useful to see how close *y* is to a gamma random variable. Fitting a gamma distribution to the first two moments of *y* results in

$$\begin{aligned} \mu_z &= \mu_{x,1} \\ N_z &= (1+\alpha)/(1-\alpha), \end{aligned} \tag{9}$$

where μ_z and N_z are the mean and diversity of the gamma distribution.

The next question is whether the third moment of this gamma distribution matches the third moment calculated for *y*. Unfortunately, the answer is no. The third moment of a gamma distribution with the above parameters is

$$\mu_{z,3} = \frac{6-2\alpha}{(1+\alpha)^2}\mu_{x,1}^3 \qquad (10)$$

This is clearly not the same as (8). However, it is quite close. Figure 1 shows the normalized third moment (normalized to $\mu_{x,1}^3$) of the two distributions. A value of $\alpha = 0$ corresponds to the case $y = x$, which is an exponential distribution (or a gamma distribution with diversity of 1). The case of $\alpha = 1$ is not well defined, but if we initialize the averaging with $y_0 = 1$, then y is equal to 1 for all time (a 'gamma' distribution with infinite diversity) and the third moment is equal to 1, as the figure shows. In these two extreme cases, *y* does have a gamma distribution (although somewhat degenerate), and in between it clearly does not, as shown by the figure.

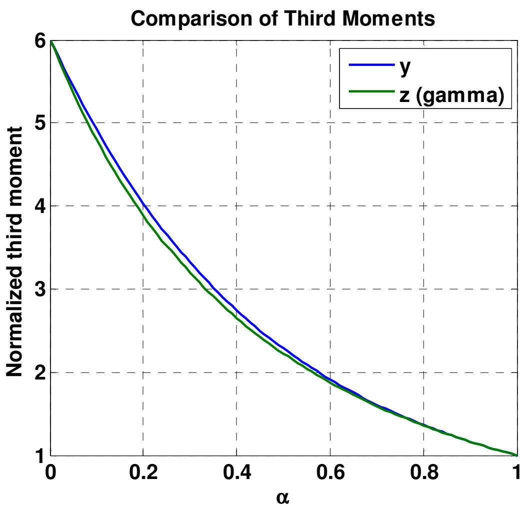

Figure 1. Comparison of the normalized third moments of *y* and *z* (a gamma distribution).

The above results show that the IIR average cannot converge to a gamma distributed random variable. However, the first three moments agree quite well, and it appears that a gamma distribution can be used as a good approximation. This is further supported through simulation.

Simulation

Simulated realizations of *y* were generated to examine its behavior in comparison to a gamma distribution. Figure 2 shows the pdf of *y* along with the pdf of a gamma distribution that has the same mean and variance. In these results, 10^7 realizations were generated. Two cases are shown: diversity of 10 ($\alpha = 9/11$) and 100 ($\alpha = 99/101$). There is some noticeable difference with a diversity of 10, but *y* is still 'close' to a gamma distribution. With a diversity of 100, the two pdf's are almost indistinguishable.

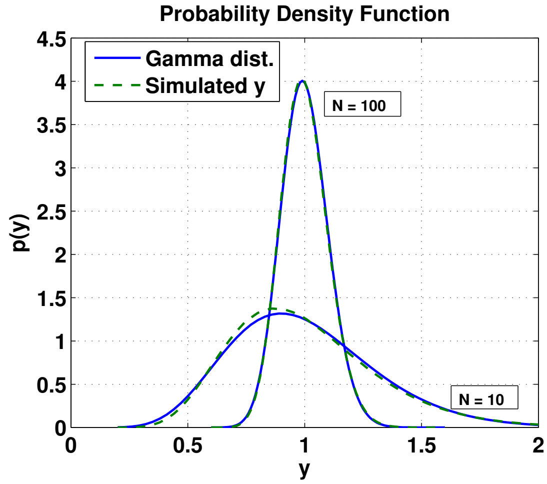

Figure 2. Comparison of probability density functions of *y* (simulated) and a gamma distribution for different diversity levels (10 and 100).

Summary of IIR Filter Results

The above results show that the averaged power spectrum, when using an IIR filter, behaves very much like a gamma distribution. In fact, we can define an effective diversity level in terms of the weighting parameters α:

$$N_{eff} = \frac{1+\alpha}{1-\alpha}. \qquad (11)$$

If we accept the approximation that the averaged power spectrum is gamma distributed, then both types of filters produce the same behavior (asymptotically). The detection threshold and asymptotic probability of detection will be calculated in the same manner for both filters. The only differences between the two would be the transient behavior before they have converged to a stationary distribution.

Transient Probability of Detection

Since the two filters perform essentially the same asymptotically, the only significant difference between the two filters is the transient behavior. Consider this example: a system is scanning a search region trying to detect a target. The system slowly scans the region resulting in a large degree of overlap of the lidar beams. The lidar updates its signal power average with every pulse and declares a detection if a signal exceeds the pre-determined threshold.

As the scan approaches the target, the received signal will transition from containing only shot noise to also containing the target signal. For the sake of analysis, we model this as a binary (or step function) transition between shot noise with mean of 1 to target plus shot noise with mean *CNR*+1.

This sudden transition is not entirely accurate for the described scenario, since the transition is more gradual, but this assumption makes the calculations more straightforward.

Figure 3 shows the resulting transient probability of detection in an example case, where the diversity is 10 and the P_{FA} is 10^{-4}, which leads to a detection threshold of 3.58. The target *CNR* is assumed to be 4.75, which leads to an asymptotic probability of detection of 90%. However, when only a fraction of the averaged pulses contain the averaged signal, the probability of detection is decreased.

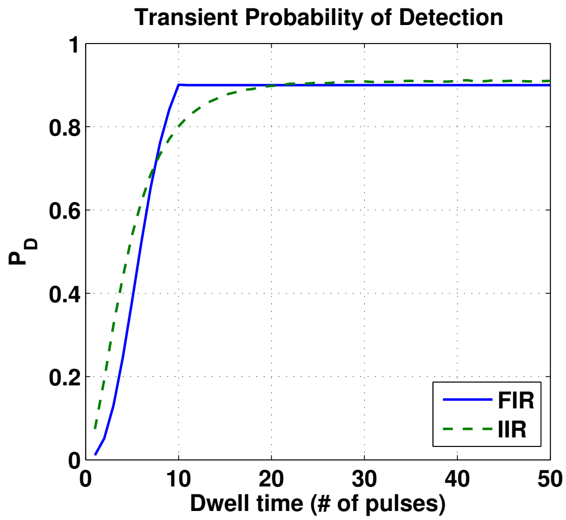

Figure 3. Probability of detection vs. dwell time. This case uses a threshold of 3.58, *CNR* of 4.75, and diversity of 10, which results in asymptotic P_D of 0.9. The transient P_D shows some difference between the FIR and IIR filters.

The plot shows the performance for both an FIR and IIR filter. The FIR filter has converged to its ultimate behavior after 10 pulses. From this point on, the most recent 10 pulses all contain the target signal so the performance does not change. The IIR filter, however, takes longer to converge. Initially, it performs better than the FIR filter, where the probability of detection is about twice as high over the first 3 or 4 pulses. This occurs because the IIR filter more heavily weights (by roughly a factor of 2) recent measurements. After 10 pulses, the probability of detection reaches about 80% with the IIR filter. It appears to take about 20 pulses, or twice the diversity, for the IIR filter to converge to its asymptotic behavior.

These two filters, therefore, do show some differences in their transient behaviors. The IIR filter is more likely to detect a target early in the transition. In the above example, the probability of detection is nearly twice as high for the IIR filter in the first few pulses. However, the IIR filter takes longer, by about a factor of two, to converge to the asymptotic behavior.

Conclusions

Two types of averaging techniques for target detection in a pulsed coherent lidar system were examined. The first, an FIR filter, results in a gamma distribution when used to average power spectra. The detection performance follows from the statistics of a gamma distribution. The second type of averaging technique, a first order IIR filter, is even simpler to implement and produces many of the same effects that the FIR filter does. While the distribution it produces does not precisely fit a gamma distribution, it can be approximated by a gamma distribution if the effective diversity is defined by

$$N_{eff} = \frac{1+\alpha}{1-\alpha}. \qquad (12)$$

If this approximation is used, then the detection threshold and the asymptotic probability of detection are calculated in the same manner as with the FIR filter.

The transient behavior of the IIR filter shows some differences from the transient behavior of the FIR filter. The IIR filter produces a higher probability of detection initially, but takes about twice as long to converge.

In many lidar applications, there will be very little difference in performance between the FIR and IIR filter. The calculation of the detection threshold and the asymptotic probability of detection are essentially the same with either filter. The transient behavior shows some difference between the two but does not strongly favor one or the other. The IIR filter may be a more desirable choice (depending upon the application), simply because it is easier to implement.

References

[1] Gatt and Henderson, "Laser Radar Detection Statistics, ...", *Proceedings of the SPIE*, Vol. 4377, pp. 251-262, (2001).