Application Note

Interference Mitigation For AWR/IWR Devices

TEXAS INSTRUMENTS

Zigang Yang and Anil Mani

ABSTRACT

Interference between different radars can have disastrous consequences for radar functionality, leading to missed detections, blind-spots, reduced range, and ghost objects. Robust, reliable radar performance requires methods to identify and mitigate interference, or avoid it altogether. This white paper describes the mechanisms of interference and methods to mitigate interference, using algorithms designed for and hardware hooks designed into the TI family of radar devices.

Table of Contents

Trademarks

All trademarks are the property of their respective owners.

1 Introduction

This white paper discusses the problem of radar-to-radar interference and how it can be managed in TI radar devices. Interference is a major issue for reliable radar functioning, as the number of deployed radars has increased in both automotive and industrial contexts. Thus, the likelihood that one radar's transmission is received by another radar has also increased. Interference results in a host of issues, such as a degradation in the noise floor leading to missed detections, or blind spots at certain ranges or directions. It can also create ghost objects in certain cases (ghost targets are targets seen by the radar which do not exist).

This paper is only concerned with cases where FMCW radars interfere with each other.

The information presented here covers the following topics:

- The mechanisms by which interference occurs between FMCW radars and the different types of interferers.
- The methods to avoid interference and control interference, that is, methods to reduce the probability of interference and methods to detect and repair chirps affected by interference.

2 Types of Interference in FMCW Radar

2.1 FMCW Radar

In FMCW radar a chirp, a signal with a linearly ramping frequency, is generated and transmitted (see Figure 2-1).

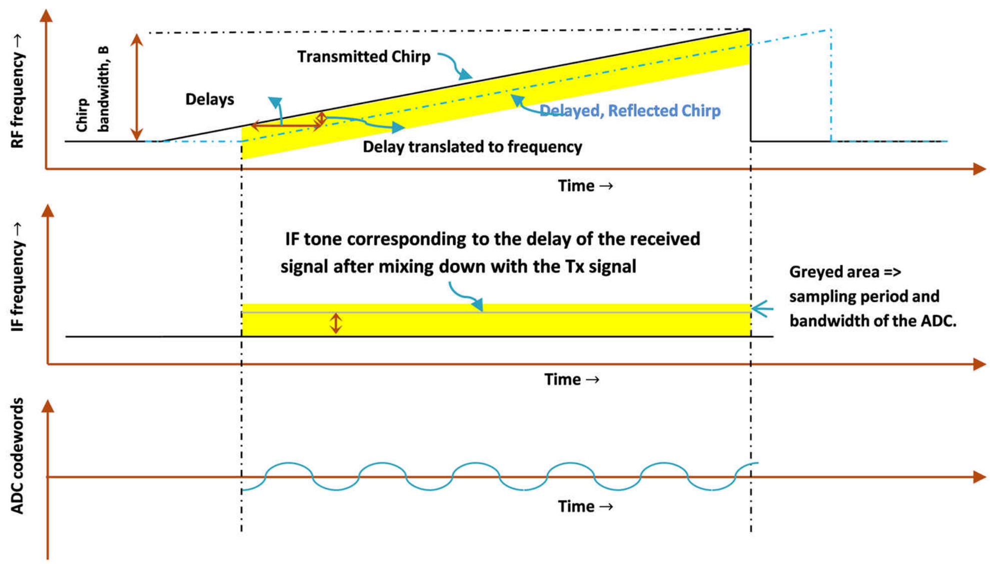

Figure 2-1. FMCW Radar Functionality in Different Domains, RF, IF, and ADC Codewords

This transmitted signal is reflected from targets in its field of view and received at the receiver. The received signal is a delayed copy of the transmitted signal. The signal received is mixed down, using the transmitted signal, and then digitized to create ADC data. Because the reflected signal is a delayed version of the transmitted signal, the mixed down signal corresponds to a sinusoid whose frequency is proportional to this delay. The delay is itself proportional to the distance of the target.

Delays can never be negative. Thus, given a positive slope, all valid objects correspond to positive frequencies. With the tone frequency estimated by a Fourier transform, the delay can be estimated. Using the delay and light speed, the distance to the target can be estimated. Thus, the maximum distance that the receiver can detect is limited by the IF bandwidth. If the target's frequency exceeds the IF bandwidth, it is filtered out.

2.2 The Radar Equation for Interference

First, let us define two terms: a victim and an aggressor. A victim is a radar device whose receiver is affected by interferers. An aggressor is a radar device whose transmit affects the victim's receiver.

The received signal strength in dBm ($P_{Interference}$) of an interfering radar can be computed using Equation 1.

$$P_{Interference} = P_{tx} + txAntGain + rxAntGain - 10log10\left(\frac{4\pi R}{\lambda}\right)^2 \tag{1}$$

where

- P_{tx} is the aggressor radar's transmit power (in dBm)
- txAntGain is the aggressor radar's transmit antenna gain (in dB)
- rxAntGain is the victim radar's receive antenna gain (in dB).

The distance between the aggressor and the victim is R, and the average RF wavelength is λ.

The radar equation for targets is shown in Equation 2.

$$P_r = P_{tx} + txAntGain + rxAntGain + RCS - 10log10\left(\frac{(4\pi)^3 R^4}{\lambda^2}\right) \tag{2}$$

Comparing the two equations shows that the path loss effect (that is, the effect of R) is weaker for interferers than targets. In other words, interference is likely to dominate the received signal even if it is far away.

2.3 Types of Interference

This section introduces two types of interferers: crossing interferers and parallel interferers.

2.3.1 Crossing Interference

If the victim radar and the aggressor radar have different slopes, the two chirps can cross each other. When the crossing happens, the victim observes a crossing interference. The aggressor's transmit signal will mix with victim's transmit signal, and the energy of the aggressor is observable to the victim only if their frequency difference falls into victim's IF bandwidth.

An example is given in Figure 2-2. As the aggressor's chirp crosses the victim's transmitted chirp, the aggressor chirp's energy is observed as a chirp that rapidly moves through the IF bandwidth. It can be a constant slope moving from zero up to IF bandwidth (as shown in this example), or it can be a slope moving from IF down to zero frequency (which happens when the aggressor's slope is bigger than the victim's slope). In time domain, the region affected by interference resembles a glitch.

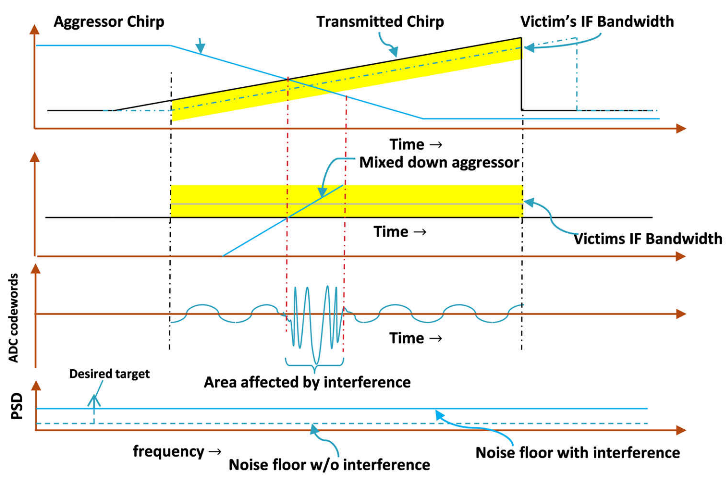

Figure 2-2. Crossing Interference Causing a Glitch in the Time Domain Signal

Finally, after a Fourier transform is applied on the ADC samples in the frequency domain, these crossing Interferers typically increase the noise floor and reduce the SNR of strong targets and bury weak targets, thereby affecting detection and creating momentary blind spots. The glitch duration (τ_{Glitch}) is governed by the victim's IF bandwidth and the slopes of the victim ($slope_{victim}$) and the aggressor ($slope_{aggressor}$). It is given in Equation 3.

$$\tau_{Glitch} = \frac{\text{IF bandwidth}}{|\,slope_{aggressor} - slope_{victim}\,|} \tag{3}$$

The glitch duration is typically small. For example, if the IF bandwidth is 12 MHz and the difference in slopes is 40 MHz/us, approximately 0.3 us, or four samples of the final ADC output, would be affected by interference.

2.3.2 Performance Analysis for Crossing Interference

Based on [1], interference noise level compared to the thermal noise can be calculated as:

$$NoiseIncInDB = P_{Interference} + 10log10\left(\frac{affectedAdcSamp}{totalNumAdcSamp}\right) - \left(-174 + N_F + 10log10\left(BandWidth_{IF}\right)\right) \tag{4}$$

For example, if the aggressor's output power is 10dBm the received power at victim can be calculated as:

$$P_{Interference} = P_{tx} + txAntGain + rxAntGain - 10log10\left(\frac{4\pi R}{\lambda}\right)^2 \quad (5)$$

Assuming a total antenna gain = 14 dB, noise figure = 14 dB. Under this condition, the noise floor increase under different condition is computed as in Table 2-1.

Table 2-1. Noise Floor Increase

Distance of Victim and Aggressor	Percentage of Samples Affected	Noise Floor Increase for 77-GHz System	Noise Floor Increase for 60 GHz
1 m	1%	24 dB	26 dB
5 m	1%	10 dB	12 dB
10 m	1%	4 dB	6 dB
1 m	10%	34 dB	36 dB
5 m	10%	20 dB	22 dB
10 m	10%	14 dB	16 dB

This is the performance degradation, assuming the victim and aggressor are facing each other. As mentioned earlier, when the slope difference gets smaller, the number of samples affected increases, but the probability of crossing interference event reduces. Fewer chirps in the frame are affected, and thus the overall performance is not as bad.

This is the noise floor degradation before any interference mitigation (signal healing techniques) is applied.

2.3.3 Parallel Interference

When the aggressor chirp and the victim chirp have the exact same slope, interference only occurs when the starting time between the victim chirp and the aggressor chirp is so close that the aggressor chirp is within the IF bandwidth of the victim's chirp.

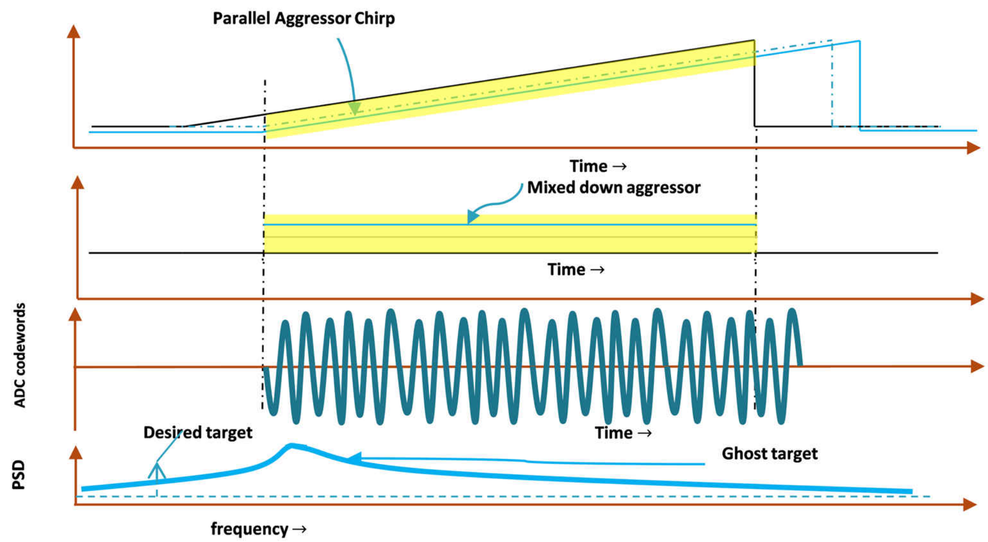

Figure 2-3. Parallel Interferer Causing a Ghost Object

When mixed down with the transmitter chip, the parallel interferer becomes a constant frequency tone in the ADC data. After the Fourier transform has been applied on the ADC data, in the Fourier domain, it becomes a ghost object. That is, it behaves like a target at random distance with a random velocity. This type of interference is called parallel interference. When it happens, the region of interference is almost the entire chirp.

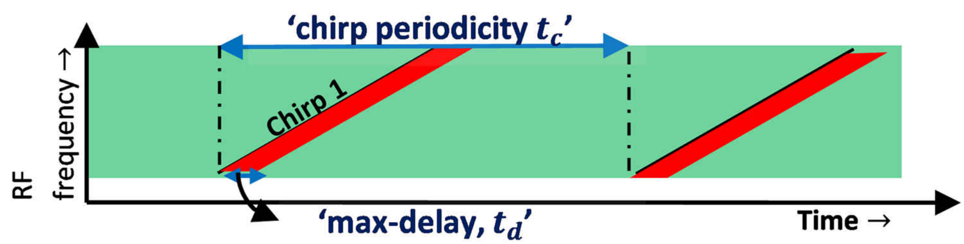

Figure 2-4. Chirps Occupy a Fraction of the Spectrum-Time Space

However, the probability of a parallel interferer is very small (see Figure 2-4). Interference only occurs if the two radars start nearly simultaneously, such that the aggressor's radar signal is present in the victim radar's IF bandwidth. Otherwise, the aggressor radar signal is filtered out by the victim's Rx. The probability of interference (p_{intf}) for a parallel interferer can be calculated using the max-delay (t_d), the chirp repeat periodicity (t_c), and the number of radars present in the scene (N_r), as shown in Equation 6.

$$\mathrm{Pintf} = 1 - \left(1 - \frac{t_d}{t_c}\right)^{Nr-1} \tag{6}$$

For example, in ultra-short range radar, with a max distance of 20 meters, the t_d is 0.13 µs. In this case, interference only occurs if the two radars start within 0.13 µs of each other. Assuming the chirp duration is 100 µs and there are 10 radars operating in the area, the probability of interference is only 1.3% assuming 100% of duty cycle. If the duty cycle of each radar is only 10%, then the probability of interference will be reduced further significantly.

When the victim and aggressor have an independent local oscillator (LO), it is difficult to get them an exact frequency slope even when the user programs them the same slope in the chirp configuration. In that case, the ghost no longer looks like a clean target and the range spectrum and Doppler spectrum all look much noisier, which can be used to identify a parallel interference situation.

2.3.4 Between Crossing and Parallel Interference

Crossing interference happens often, but only affects a small number of samples in the chirp, while parallel interference happens rarely, but affects almost the whole chirp; thus, as the chirp slope difference gets smaller, the glitch gets longer, but the crossing probability decreases.

3 Interference Avoidance

Users should always try to avoid interference if possible. This section describes the methods to avoid interference.

3.1 Standardization: Different Frequency Band and Time Slot for Different Radars

The first method is standardization. Standardization refers to frequency planning and chirp design, as well as time slot management. Frequency planning, based on the resolution requirement, lets different radars coexist in different RF bands. For example, the AWR family of devices has 4 gigahertz of RF bandwidth, which can be divided into 2 gigahertz bands and used simultaneously by two radars.

On the other hand, there is usually silent time between neighboring frames without any active chirping. If the duty cycle for the radar system is 10%, potentially, 10 different radars can separate in time.

Figure 3-1 shows that a radar signal separate in RF frequency and time slot does not experience any interference. The frequency separation is easy to implement. Time slot management must have a common global timing source for all users to synchronize to. In this case, coarse synchronization is sufficient for frames.

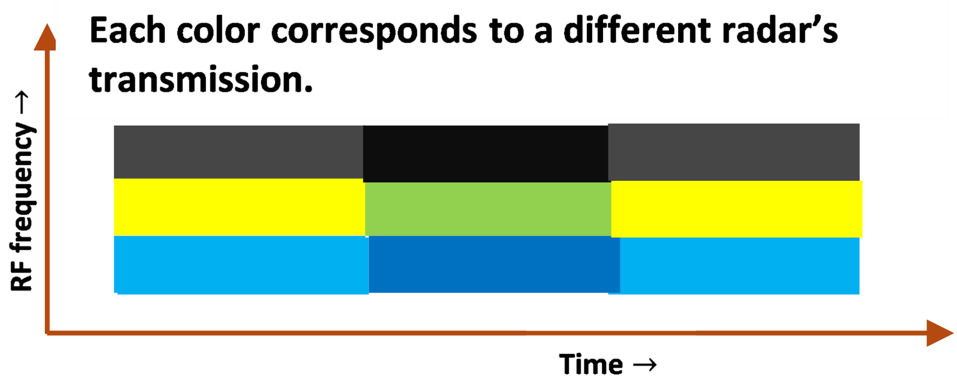

Figure 3-1. Using Different Time Slots and Frequency Slots to Avoid Interference

Another use case can be direction-specific predefined frequency band separation. For example, users can use a separate band for long range radar and short range radar so that they do not interfere with one another. TI also recommends the use of a different band for front-facing radar, and another band for a rear-facing radar.

3.2 Different Starting Time for Parallel Interference

If a single manufacturer is building all the radars, they can be made so as to be synchronized to the same clock, to a global time for that particular factory. If every radar device is then configured with the same kind of chirp and frame, it can result in parallel interferers. However, if every radar's frame is offset, approximately one microsecond or so to the global times so that they do not interfere with other radars, then a large number of radars can coexist in a limited space and in the same bandwidth. For example, if the chirp time is 100 microseconds, and the max distance of interest is 150 meters; that is, the time of transmit is less than one microsecond then approximately 100 such radars can coexist in the same bandwidth. Synchronization also lets frames be stacked one after the other, so as not to interfere with other radars.

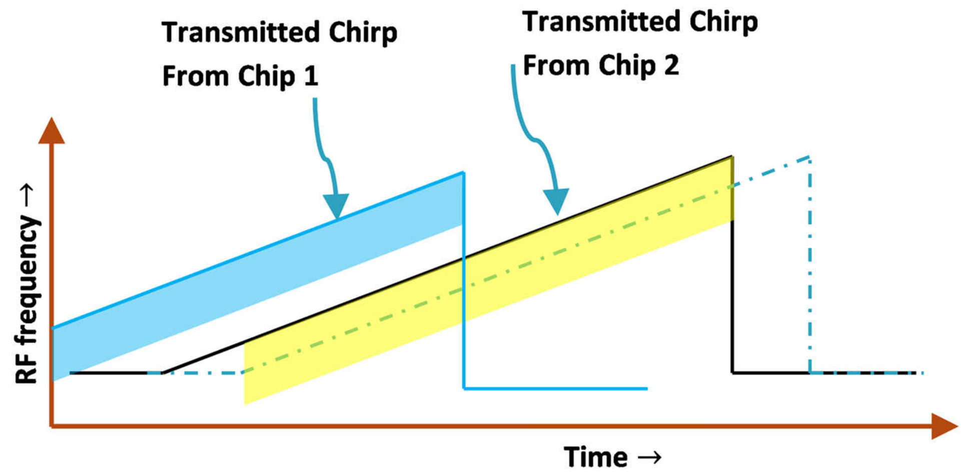

Figure 3-2. Precise Intra-Chirp Time-Slot Planning

A simple method to achieve synchronization between radars that are placed close-by is through the master/slave mode in TI radar devices. In this scheme, one device is designated the master. This device generates triggers to the slave devices when it transmits a frame. The slaves can then delay, using this signal to trigger their own frames after a precisely defined delay.

3.3 Sensing and Avoidance

In the absence of any synchronization, users can still perform ‘sense and avoid’. In this scheme, before a device begins transmission, it senses the spectrum. This is achieved by keeping the receivers active and the transmitter switched off.

If there have been no transmissions by other radars, the spectrum is silent. The ADC data should only show the thermal noise floor and the noise figure. If, on the other hand, there are transmissions from another radar device, expect spikes in the ADC data corresponding to the points where the crossing occurs.

XWR devices can generate fast chirps of the order of 250 megahertz per microsecond, allowing for fast scans.

A max hold of ADC data cross chirps shows the interferer clearly. In Figure 3-3, the interferer is chirping between frequencies f_1 and f_2. Thus, the ADC output shows the energies between f_1 and f_2. If the period of scanning is long enough to cover multiple frames, estimate the number of aggressor radars using the number of discrete bands of frequencies used. The user can also estimate frame periodicity bandwidth occupied by chirps. Most importantly, the user can find free spectra, or time slots, where interference-free transmission is possible.

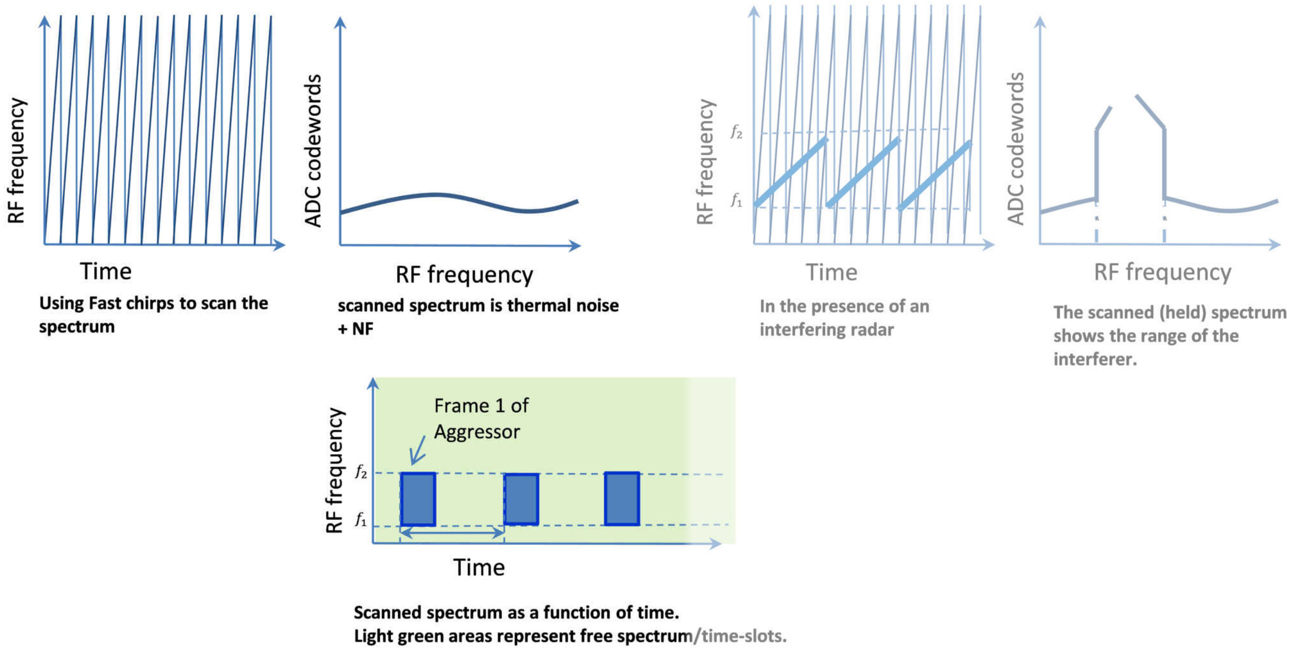

Figure 3-3. Sense and Avoid

When the interferers have been identified, the radar can start transmission in regions where the interferer is not active.

3.4 Antenna Polarization

The final method involves specific polarization of antennas. This method uses, for example, horizontal polarization for a certain set of antennas, and vertical authorization for another set. If an aggressor uses horizontally polarized antennas for transmission and the victim uses vertical polarized antennas for its receiver, then the signal from an aggressor is attenuated (by ~ 10 dB) at the antenna of the victim. This method requires expertise in antenna design. This is a useful method, but only two different options, vertical and horizontal polarization, are available. Also, this method involves increased complexity in antenna design.

4 Localization and Interference Mitigation

In many cases, it is impossible to avoid the crossing interferers. As shown in Section 2.3.1, crossing interference without any treatment can increase the noise floor and blind the weaker targets. This section describes the methods to reduce the system degradation by localization followed by mitigation. Localization refers to the process of finding which samples in a chirp are affected by interference.

4.1 Localization

Localization can be accomplished in one of two ways.

- First, find outliers in ADC data. Strong crossing interferers look like large glitches in ADC data. For example, if you were to take the energy of each sample in a chirp and plot it as a function of time, at the point where the interferer crosses, there is a large increase in the first sample energy of the chirp. A suitable threshold can be found and set, and samples that cross this threshold in energy can be marked as having been affected by interference. Figure 4-1 shows a sample ADC data (top) – When it its absolute value is plotted the glitch is clearly visible (middle), though not very distinct from the signal due to the large low frequency signals present. However, if these low frequency signals are suppressed using a simple difference filter, the resulting signal lets the glitch stand out even more.

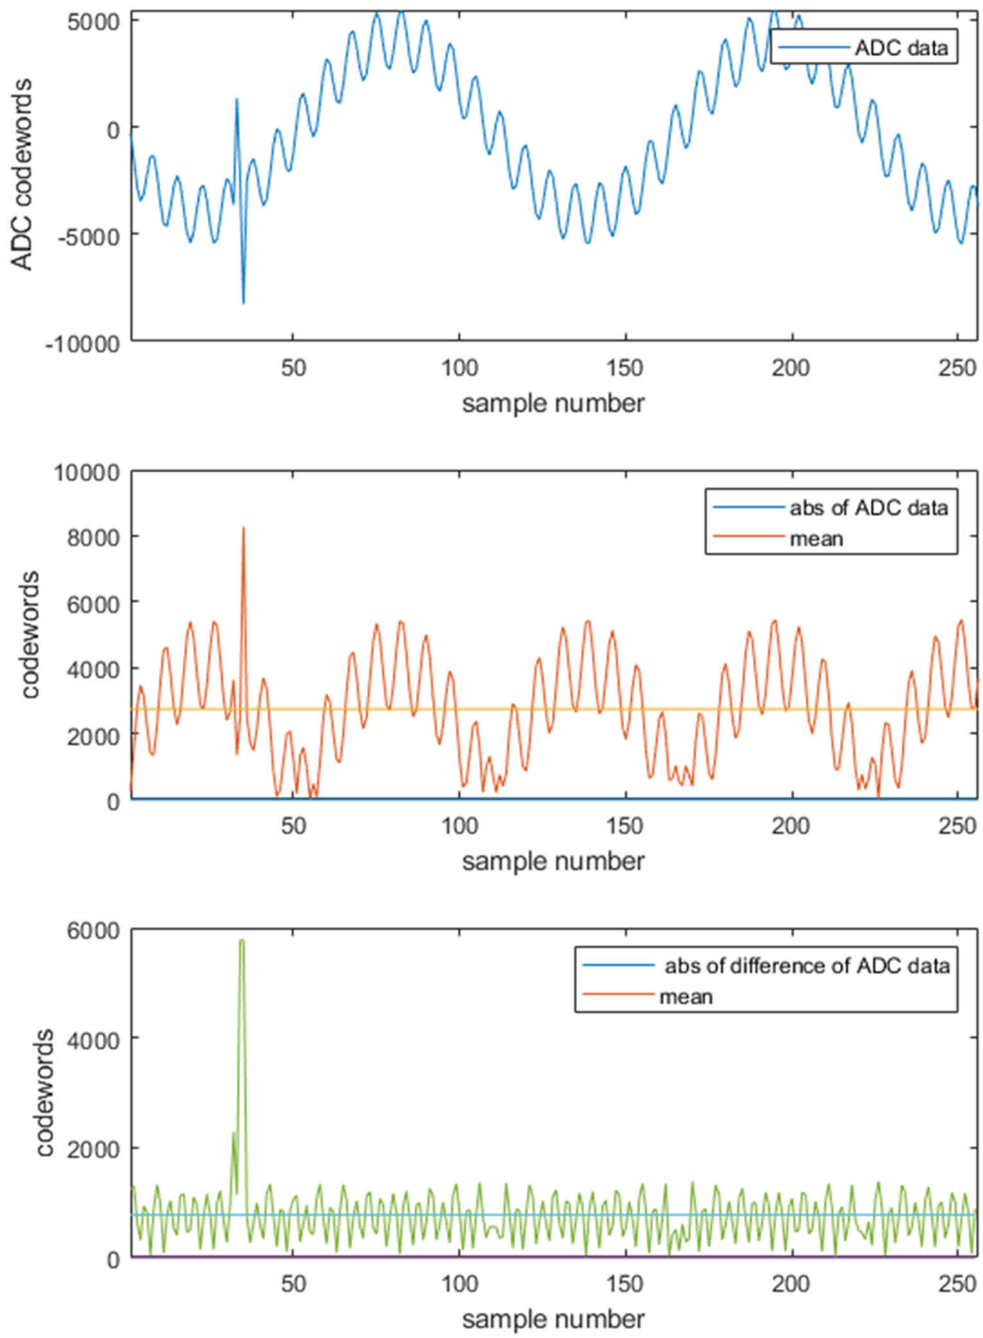

Figure 4-1. Detecting Outliers Using abs and abs-diff Filters

- Chirp quality metrics are additional metrics that are optionally attached to each chirp, and use some of the advanced features of the XWR devices to provide information about interference. The interference detection training video provide details to enable chirp quality metrics. As XWR devices have complex base-bands, they can discriminate between positive frequencies and negative frequencies. As stated in Section 2.3.1, delays can never be negative. Thus, if the slope is positive, all valid objects have positive frequencies (that is, they exist in the signal band). Any signal in the negative frequency (image band) is likely to be due to interference. The signal and image band monitor monitors these two bands separately. As can be seen in Figure 4-2, the signal band (blue line) is stronger than the image band (red line). However, when crossing interference appears, the image band has an sudden rise in energy. This indicator is used to locate weak interferers.

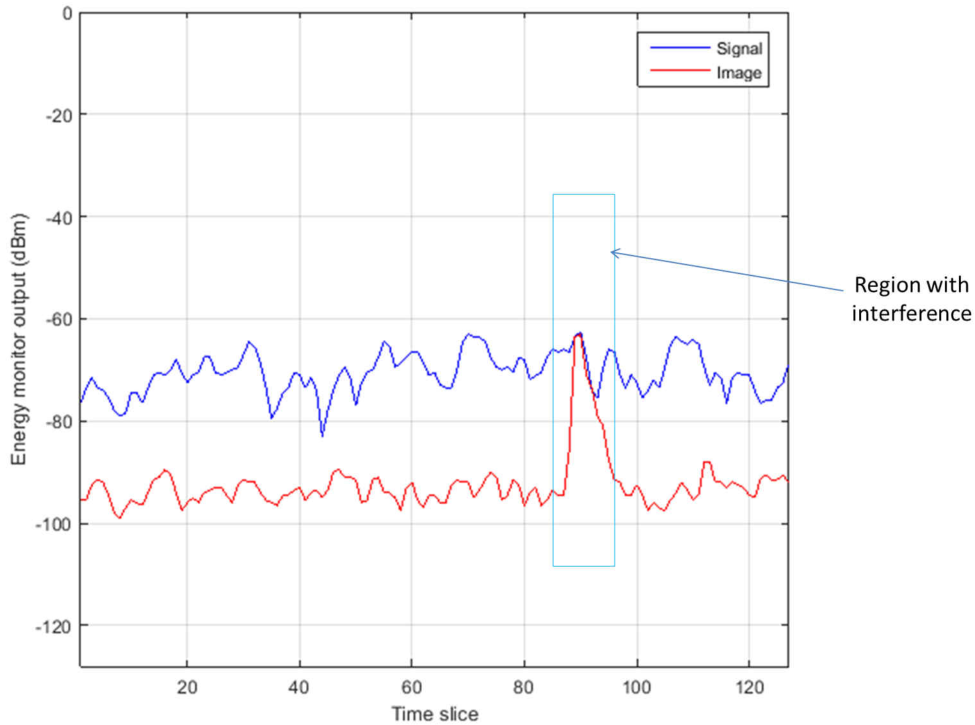

Figure 4-2. Using the Signal and Image Band Monitor to Locate Interferers

4.2 Mitigation

Having found interferer locations, we now would like to mitigate them. Mitigation here refers to the process by which the region of interference is healed.

The simplest method of mitigation is to replace the region of interference with zeros (Figure 4-3 - top). However, that has the side effect of creating large sidelobes that might bury weak targets. A better approach is to blank out with a window. A smoothing window is used to zero the samples that have been affected by interference. This leads to lower sidelobes and better detectability of weak targets (Figure 4-3 - middle).

A better approach is to perform linear interpolation in the blank region, using the last good ADC sample before interference and the first good ADC sample after the period of interference. Because the strongest reflectors are likely to be closer to the radar and thus have lower frequencies, this approach works well in many cases (Figure 4-3 - bottom).

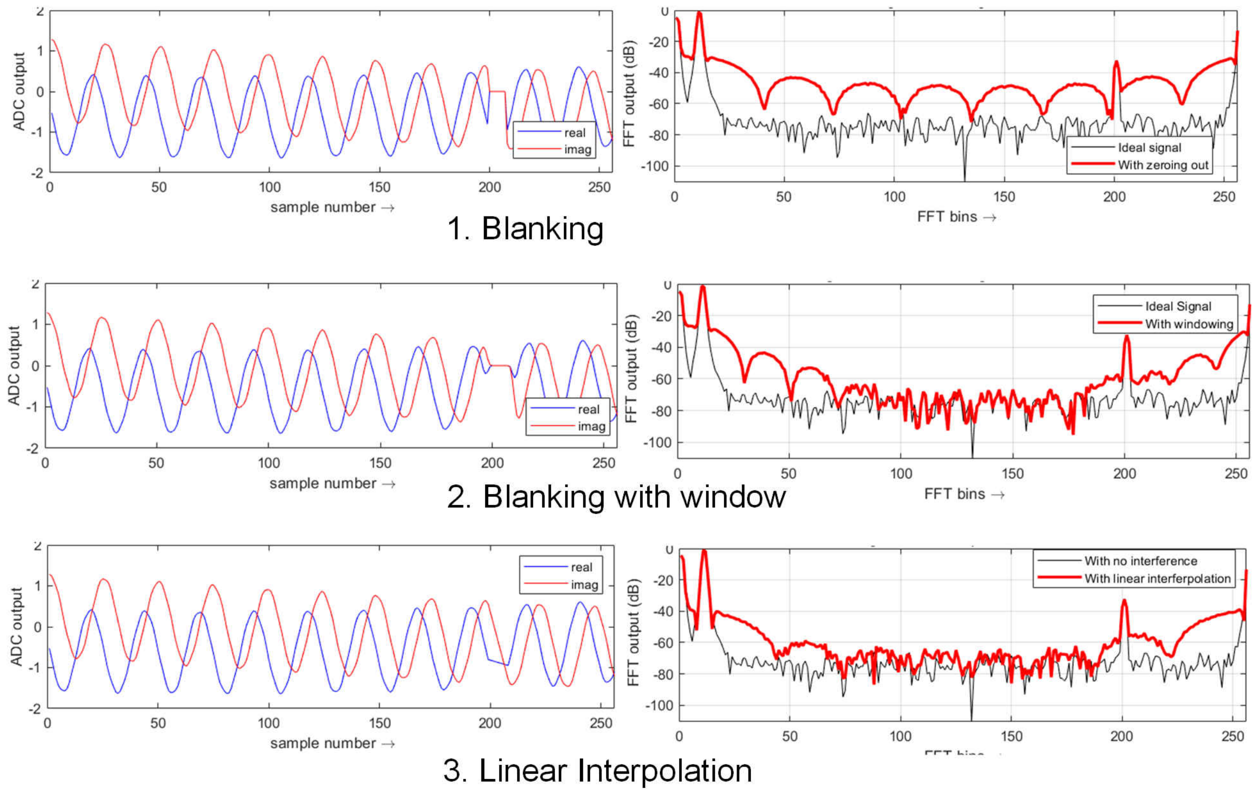

Figure 4-3. Different Mitigation Approaches

Mitigation is an active area of research, and more complex mitigation schemes possible than the three described here. However, as the mitigation scheme becomes more complex, one has to weigh the amount of MIPs consumed against the benefit gained by the more complex scheme.

5 Dithering and Randomization

When parallel interference occurs, the whole chirp (or most of the chirp) is damaged, and it is difficult to correct it. Therefore, localization and mitigation is not very useful for parallel interference.

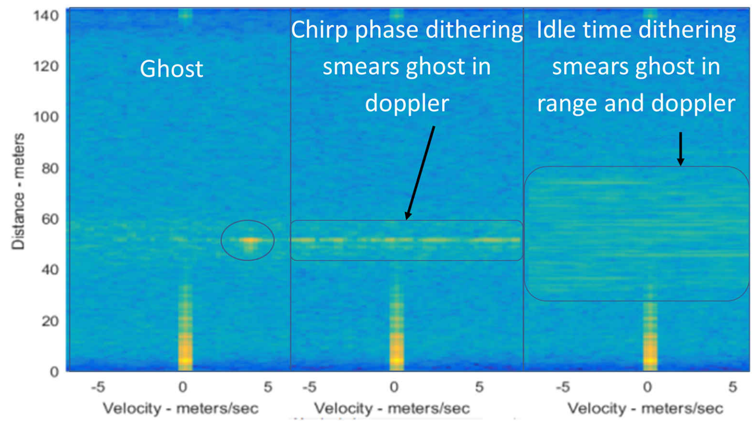

Figure 5-1. A Ghost from a Parallel Interferer is Spread with Randomized Chirp Phase, and Diffused with Idle-Time Dither

Parallel interferers can be weakened by a process called chirp dithering (or chirp randomization). In this process, certain parameters of a chirp are randomized on a per-chirp basis. For example, the phase of the chirp can be made random. Because the aggressor has no knowledge of the victim's randomization scheme, a parallel interferer is spread during the doppler processing. The chirp starting phase can be randomized using the per-chirp phase shifter API, or the binary phase shifter API. There are multiple other parameters of a chirp that can be randomized; chirp slope, chirp start frequency, and the chirp idle time can all be randomized using the chirp config API. This chirp config API *rlSetChirpConfig* is described in the interface control document.

Figure 5-1 shows how a ghost target due to a parallel interferer is spread in Doppler by random binary phase modulation (that is, chirp phase dithering), and further spread using chirp idle time dithering. If no randomization options are used, parallel interferers appear as ghost objects. If randomization is applied, the peak of the interferer is destroyed by dithering. Randomization works by damaging aggressor coherence across different chirps, thereby reducing their effect during 2D processing. The reduction is approximately 10*log*10 of the number of chirps in one frame. When the aggressor coherence is destroyed, CFAR algorithms can then be used to remove interference-related effects.

Dithering schemes introduce more complexity during Doppler processing, due to the fact that some correction must be applied. For example, chirp phase dithering can be corrected by applying an opposite phase shift to the chirp before Doppler processing. Certain dithering schemes, such as idle time dithering, can introduce high noise-floor in Doppler.

Because oscillators for radar devices vary, without clock synchronization, the chirp starting time moves relatively slowly between the radar devices. For example, two radar devices, programmed to the same chirp configuration, have a 200-ppm difference in local oscillator frequency. Assuming that the chirp must start within 1 µs to see the interference and frame rate is 0.1 s, then after one frame, the relative chirp starting time moves approximately 20 µs. Thus, it takes approximately 1/20 frames to move away from the interference zone. It takes approximately 8 minutes to shift one whole frame to ensure it gets back to the interference zone. Users will see interference in 1/20 frames before it moves away.

When the LO variation is smaller, such as a 1-ppm difference between the two radars, then it takes approximately 27.7 hours to shift back to the interference zone; users will see interference for 10 frames before it moves away. This introduces the possibility of frame start randomization. Each frame starts with a random time offset. In such cases, the parallel interferer may only affect one frame. In this way, the worst case system performance can be improved.

6 Conclusion

Radar-radar interference is a stumbling block to the wide (and dense) deployment of radars. If it is not accounted for, it can lead to detection failures, ghost objects, and reduced radar range. Using different schemes, such as randomization, dithering, frequency planning, and localization and mitigation, it is possible to manage interference and provide robust performance.

7 References

1. Sriram Murali, Karthik Subburaj, Brian Ginsburg and Karthik Ramasubramanian, *Interference Detection in FMCW Radar Using A Complex Baseband Oversampled Receiver*, https://ieeexplore.ieee.org/document/8378800

8 Revision History

Changes from Revision * (January 2020) to Revision A (September 2022)	Page
• Updated the numbering format for tables, figures, and cross-references throughout the document	1