ECE 146A Lab: An Initial Exposure to mmWave Radar Sensing

Lab Objectives: The goal of this lab is to provide an initial exposure to the fundamentals of radar sensing, with parameters consistent with emerging mm wave sensing applications.

1 Background

The basic concept of radar is as follows. The transmitter sends out a signal $s(t)$. Assuming that the receiver is co-located with the transmitter, when the signal bounces back from a target, the "return" signal is received with a delay equal to the round-trip time between the transceiver and the target. The radar signal is designed so as to enable accurate estimation of this delay, which in turn provides an estimate of the distance between the transmitter and the target, termed the *range* R. Figure 1 shows a nominal geometry.

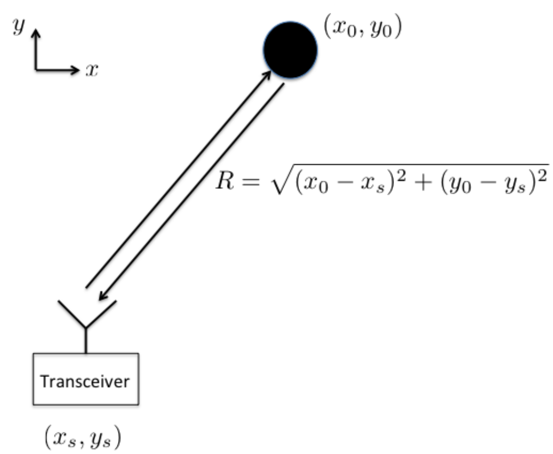

Figure 1: Basic radar geometry.

A particularly popular type of radar waveform is a linear chirp, where the instantaneous frequency is varied linearly with time. That is, the passband signal that is transmitted is given by

$$s_p(t) = \cos(2\pi f_c t + \pi S t^2) \ , \ 0 \leq t \leq T_d \tag{1}$$

where S is termed the *slew rate,* and T_d is the chirp duration. Thus, the instantaneous phase is given by

$$\theta(t) = 2\pi f_c t + \pi S t^2$$

and the instantaneous frequency by

$$f(t) = \frac{1}{2\pi} \frac{d\theta}{dt} = f_c + St$$

Running Example with mmWave chirp parameters: For example, mmWave radar sensors for the automotive market use the band from 77-81 GHz, so let us take our running example as $f_c = 80$GHz, $S = 100$MHz/μs, and $T_d = 10\mu$ s, so that the instantaneous frequency varies from 80 to 81 GHz. We will use these in our numerical experiments.

2 Lab Assignment

1) Suppose that a return signal $s_p(t-\tau)$ is mixed against the local copy $s_p(t)$ at the transceiver to obtain

$$y(t) = 2s_p(t-\tau)s_p(t)$$

Show that, after filtering out double frequency terms around $2f_c$, we get a sinusoid

$$y(t) = \cos(2\pi f_0 t + \phi_0)\ ,\ \tau \le t \le T_d \tag{2}$$

with

$$f_0 = S\tau\ ,\quad \phi_0 = -2\pi f_c \tau - \pi S \tau^2 \tag{3}$$

In practice, the phase ϕ_0 in a single chirp does not tell us much. Can you guess why?

2) The return signal from a target at distance R from the transceiver comes back with delay

$$\tau = \frac{2R}{c}$$

where $c = 3 \times 10^8$ m/s is the speed of light.
(a) Suppose we are interested in estimating targets at a maximum range of 30 m. What is the largest value of τ that we need to accommodate? How does this compare with the frame duration $10\mu s$?
(b) What is the largest frequency f_0 that we need to estimate? How does this compare with the 1 GHz variation in instantaneous frequency over the chirp duration?

3) Now, consider a generalization of (2) to K targets:

$$y(t) = \sum_{k=1}^{K} A_i \cos(2\pi f_k t + \phi_k) I_{[\tau_k, T_d]} + n(t) \approx \sum_{k=1}^{K} A_i \cos(2\pi f_k t + \phi_k) + n(t)\ ,\ 0 \le t \le T_d \tag{4}$$

ignoring the offsets at the beginning, since $\tau_k \ll T_d$ in our regime of interest. Assuming that we need to accommodate frequencies corresponding to a range of up to 30 m, how fast do we need to sample $y(t)$ so as to approximately preserve the information regarding the target ranges?

4) *(One sensor, multiple targets)* Build a discrete-time simulation model for the system in (5) by sampling it at the rate $\frac{1}{T_s}$ determined in part 3. Adding in noise, we obtain the model

$$y[n] = y(nT_s) = \sum_{i=1}^{K} A_k \cos(2\pi f_k n T_s + \phi_k) + W[n],\ ,\ n = 0, 1, ..., N-1 \tag{5}$$

where $N = \frac{T_d}{T_s}$, and $W[n] \sim N(0, \sigma^2)$ are i.i.d. noise samples. The integrated SNR for the kth target is defined as

$$SNR_k = \frac{N A_k^2}{\sigma^2}$$

(thus, you can set $\sigma^2 = 1$ in your model, and set A_k.
(a) Generate the received signal in (5) assuming that there are 3 targets at 3 m, 10 m, and 30 m, respectively. Keep the SNRs programmable, but for your numerical evaluation, set the SNR of the closest target to 10 dB, and set the others assuming the same radar cross section for each target, and $1/R^4$ loss (the radar return signal encounters free space loss scaling as $1/R^2$ in each direction). The radar cross section of a target is the ratio of the power of the reflected signal to the incident signal. It depends on a large number of factors (e.g., target size, geometry, material, direction from which radar signal is arriving), hence we are not modeling it explicitly,
(b) Since the range of each target maps to a frequency, let us use an FFT to estimate the dominant frequencies in $\{y[n]\}$. Compute the FFT and "look for peaks." You will have to come up with a scheme to do this–try to do it without using the prior knowledge that there are three targets. The size of the FFT should be a power of 2, and should be at least N (pad with zeroes as needed). You may also want to try increasing the FFT size to try and get better resolution.
(c) Estimate the dominant frequencies and map them to range estimates. Plot a histogram of the number of targets your scheme estimates. Define range error for a given target as the smallest error between one of the estimated ranges and the true range. This is well defined for each target, even if the

number of dominant frequencies is not equal to the number of targets. Plot a histogram of the range error for each target.
(d) Repeat, varying the SNR of the closest target (scaling the other SNRs with respect to it as before). At what SNR do you start doing "really well"? At what SNR do things break down?
(e) For a meaningful range of SNR as determined in (d), plot the standard deviation of the range error for each target versus SNR, averaging over multiple simulation runs for each SNR.
Remark: This is optional, and may be useful for getting smoother curves. You may wish to add a random dither to the range of each target on each run to avoid grid effects in range estimation. For example, for the nth run, you can set range $R_n = R_{nominal} + d_n$, where the dithers $d_n \sim Unif(-0.1m, 0.1m)$ are i.i.d. across runs. You would of course measure error relative to R_n in the nth run. You would dither the range of each of the three targets in this fashion, and use independent dither sequences for each target.

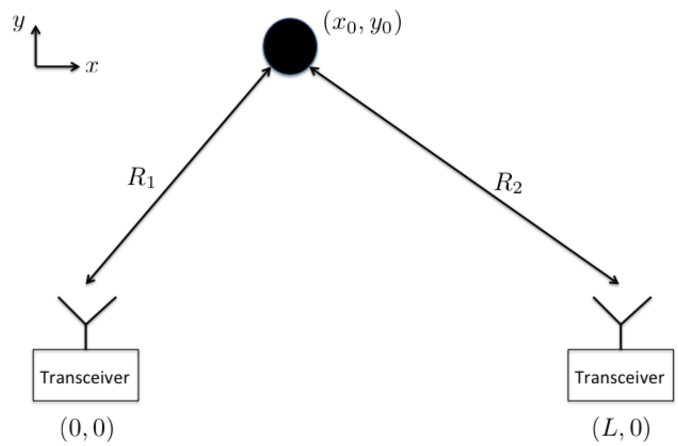

Figure 2: Estimating absolute location using two sensors.

5) *(One target, multiple sensors)* Now, consider two sensors separated by L, and a target as shown in Figure 2.
(a) Assuming you have ideal estimates of the ranges R_1 and R_2, find an estimate of the target's absolute location (x_0, y_0) via geometric calculations.
(b) Set $L = 4$ m, consistent with placing the sensors along the side of a car, and suppose that the target is at location $(1, 10)$ m. Set the SNR at each sensor to be the same value (ignoring the differences in path loss due to the different values of R_1 and R_2). Now, simulate the discrete time signal (5) at each sensor (the noises are independent at the different sensors). Estimate the ranges using the FFT as before, and then derive an estimate of the absolute target location. Using multiple simulation runs at SNR of 10 dB, plot a histogram of the errors in R_1, R_2, x_0 and y_0. Comment on any patterns that you see.
Remark: Just plugging in the range estimates $\hat{R}_1$ and $\hat{R}_2$ into the geometric calculations in (a) may not work well when the estimates are noisy. What we typically do is to find the estimate (x_0, y_0) such that the corresponding values of R_1, R_2 are close to the measured values $\hat{R}_1, \hat{R}_2$. Suppose the geometric relationships are given by $R_1 = f_1(x_0, y_0)$ and $R_2 = f_2(x_0, y_0)$. You can then get an estimate of (x_0, y_0) by solving the following problem:

$$(\hat{x}_0, \hat{y}_0) = \arg\ \min_{(x_0, y_0)} \ \left(f_1(x_0, y_0) - \hat{R}_1\right)^2 + \left(f_2(x_0, y_0) - \hat{R}_2\right)^2$$

Since the functions f_1, f_2 are nonlinear, the preceding is a *nonlinear least squares problem.* You can find off-the-shelf functions for doing this, or you can do a brute force grid search.
(c) After doing spot checks to figure out what range of SNR is relevant, average over multiple runs and report on the standard deviation of the errors in x_0 and y_0 as a function of SNR.

6) Let us now go back to the single sensor, single target scenario in Figure 1, and suppose that the target is moving, which means that the range R, and hence the delay τ, is a function of time. Going back to (3), note that the instantaneous phase is given by

$$\theta(t) = 2\pi f_0 t + \phi_0$$

which is a function of both t and τ.
(a) Now, if τ is a function of t, show that the instantaneous frequency is given by

$$f(t) = \frac{1}{2\pi}\frac{d\theta(t)}{dt} = S\tau + S\dot{\tau}t - f_c\dot{\tau} - S\tau\dot{\tau} \tag{6}$$

where the derivative of the delay, $\dot{\tau}$ is given by

$$\dot{\tau} = \frac{2\dot{R}}{c} = \frac{2v_R}{c} \tag{7}$$

where v_R denotes the radial velocity of the target with respect to the sensor, with positive sign if the target is moving away from the sensor.
(b) Keeping the most significant term involving $\dot{\tau}$, write the instantaneous frequency as

$$f(t) = S\tau - f_c\dot{\tau} = S\tau - f_D \tag{8}$$

where

$$f_D = \frac{2v_R}{c}\, f_c$$

is the Doppler shift due to the target's radial velocity. For a target at distance 10m, moving at a radial velocity of 3 m/s with respect to the sensor, what is the Doppler shift? What is the corresponding phase shift over a chirp of duration 10μs?
(c) How many consecutive chirps would you need to send for the net phase shift due to Doppler to be significant (e.g., a few multiples of 2π) for a target moving at 3 m/s? How much would the range of the target change over that duration?

The preceding background should help you understand the notion of 2D FFT for estimating range and Doppler discussed in the white paper from Texas Instruments posted along with the lab. The range can be estimated from a single chirp ("fast time"), but Doppler is estimated by comparing the phase shifts around a given range bin across multiple chirps ("slow time"). Taking a 2D FFT in fast and slow time therefore allows us to estimate range and Doppler.

7) For targets at range of up to 30 m, and radial velocity resolution down to 0.5 m/s, how many chirps of duration 10 μs do you need? What is a good choice of size (N_{fast}, N_{slow}) for the 2D FFT?

3 Lab Report

- Discuss the results you obtain, answer any specific questions that are asked, and print out the most useful plots to support your answers.
- Append your programs to the report. Make sure you comment them in enough detail so they are easy to understand.
- Write a paragraph about any questions or confusions that you may have experienced with this lab.