Locating Payloads Using Doppler Shift in the Signal of a Backup CW Beacon

Richard Eason[a,1] **and Joseph Patton**[a]

This manuscript was compiled on May 10, 2019

There are a number of devices for locating high altitude payloads on their return to earth, including APRS, Iridium modems, SPOT® and custom radio transmitters. Such units can transmit the exact location, which is ideal. Having an accurate location ahead of time will often influence the direction of approach for recovery, especially if it is near obstacles such as a marsh, river or stream that is difficult to cross. Unfortunately, every device has some probability of failure due to a variety of causes, everything from component failure, loss of satellite lock or radio reception to human errors such as forgetting to charge or change batteries. Including a second independent location device as a backup can greatly increase the odds of locating the payloads, but it can also double the required weight. A simple, lightweight, nearly foolproof, continuous wave (CW) beacon can serve as an attractive backup. These units transmit a simple beep heard on the radio. The transmitter is typically located using a directional antenna by comparing signal strength in different directions from different locations. This technique can be tricky in practice, as real-world conditions create reflections making it difficult to identify the actual source direction. Furthermore, the signal works line-of-sight, so a payload on the ground might only be detected up to a half mile away in some cases, requiring a good rough guess of location to even get started. If one has access to an aircraft, the signal can usually be received up to 20 or more miles away from the air, but narrowing the location from an airplane can still be tricky using a directional antenna and signal strength.

This paper describes a technique we have developed for quickly and accurately locating a CW beacon using the Doppler shift heard in the received signal. Using a continuous wave (CW) beacon of 433 MHz, an aircraft traveling directly toward the source at 100 knots will experience a shift in the signal to 433,000,074 Hz. This 74 Hz change in the radio frequency will result in a 74 Hz shift in the audio frequency heard in a CW receiver. The audio signal will therefore vary between +/- 74Hz of a center frequency depending on the direction of travel relative to the beacon. Our experience has shown that we can narrow down the location of the transmitter much more quickly and accurately by listening to the Doppler shift in the received signal than we can using a directional antenna and signal strength. To aid in this process, we have developed a system based on a Raspberry Pi® that will listen to the Doppler-shifted beacon combined with the GPS track to automatically compute the location fairly precisely. While we have currently experimented in the use of this system in an aircraft, the technique also has some feasibility using an automobile should the beacon land in proximity of suitable roads, and we are also developing it for deployment on a quad-copter.

Payload recovery | CW beacon | Direction finding | Doppler shift

Introduction

High altitude payloads typically carry locating devices for tracking the balloon in flight, not only as a component of the data collected, but also for aiding in the recovery of the payloads after their return to earth. A variety of devices are used ranging from amateur radio's APRS, to satellite-based systems such as Iridium modems, SPOT®, Garmin inReach® and other personal locator beacons to custom radio transmitters. These units can transmit the exact location on landing, which is of great help in recovery, especially when the terrain is difficult. Knowing the exact location can influence the direction of approach for recovery, in particular if it is near obstacles such as a marsh, river or stream that is difficult to traverse.

Unfortunately, every location device has some probability of failure due to a variety of causes ranging from the physical to the human. This includes everything from component failure to loss of satellite lock or radio reception to human errors such as forgetting to charge or change batteries. Including a second independent location device as a backup can greatly increase the odds of locating the payloads, but it can also double the required weight. It can be helpful if the backup unit has different failure modes. One of our balloons accidently turned into a floater, and the APRS batteries were long dead before it came down, so having a backup unit with relatively short battery life wouldn't have helped any.

A simple, lightweight, nearly foolproof, continuous wave (CW) beacon such as the Beeline® transmitter from Big Red Bee can serve as an attractive backup. These units transmit a pulse of carrier frequency that is heard as a simple beep on the radio. With an external battery they can transmit for weeks.

CW beacons are typically located using a directional antenna by comparing signal strength in different directions from different locations. This technique can be tricky in practice, as real-world conditions create reflections making it difficult to identify the actual source direction. Furthermore, the signal works line-of-sight, so a payload on the ground might only be detected up to a half mile away in some cases, requiring a good rough guess of the location to even get started.

Our earliest balloon flights generally used two GPS-based location devices as well as a CW beacon for further backup. Naturally, as the primary author is a pilot and we have access to aircraft, we began experimenting with locating a CW beacon from the air. In our typical conditions, we can pick up the signal from perhaps 20 miles away and navigate to the general area, but we quickly discovered the difficulty in getting a very precise indication of location using the directional antenna and signal strength. As with terrestrial-based direction finding, real world conditions such as signal reflections can make it difficult to identify the actual source direction at any given time.

Early on, it was noticed that we seemed to get a Doppler shift in the audio frequency coming from the radio as we circled near the payloads. This shift was ignored at first as the received signal is a radio wave and we were obviously traveling at far slower than the speed of light, but then it was realized that a very small percentage shift in radio frequency received would translate to an equal number of cycles per second shift in the audio heard. For example, when traveling directly towards our 433.2 MHz beacon at 100 knots (185 km/h), the 433,200,000 Hz signal would shift 74 Hz to 433,200,074 Hz. In a CW receiver this shift would translate to a +/- 74 Hz shift in audio frequency heard as we fly towards, then away from, the beacon, which is quite noticeable. We quickly found that once we arrived at the general location we could narrow down the location of the transmitter much more quickly and accurately by listening to this Doppler shift than we could using a directional antenna and focusing on the strength of the received signal. Given our access to aircraft, the ease and relative accuracy in locating the payloads using this technique, we soon began including only a single GPS-based location system on most flights to reduce weight and simplify our base system, while always including the CW beacon as a backup. With nearly 100 launches so far and most landing in the forests where they are not likely to be stumbled upon for some time, we've yet to lose a payload. The CW beacon has proved its worth at least a half dozen times.

Using Doppler shift to locate a CW beacon

The CW beacon we use is the Beeline® transmitter from Big Red Bee (1). This beacon operates in the 70 cm amateur radio band and an amateur radio license is required to use it. Most parameters of this device are programmable. We have it set to use CW mode with an output power of +12dBm, and transmit a100 ms "beep" with a 900 ms delay between beeps, or one beep per second. While shorter pulses at a lower pulse rate and lower power would yield better battery life, with these settings the beeps are long enough and come often enough for us to get a good reading of the frequency. We found that with these settings the included lithium polymer battery will last less than two days, so we purchased transmitters without a battery and instead connect a "2CR5" battery which allows the unit to operate continuously for nearly two weeks, giving us a sufficient window of time to find conditions suitable for flying an aerial mission should it be required. The radio receiver we use is the Yeasu® FT-817 with the YF-122C Collins mechanical CW filter installed.

In the following, we will describe how we use the Doppler shift in the signal to locate the CW beacon by flying overhead in an aircraft, beginning with the technique of just listening to the audio from the radio followed by a computational solution yielding fairly high accuracy. Before beginning, however, we feel inclined to preface this with a brief safety message.

As every good pilot should know, safety is of primary concern. Maneuvering flight at low altitudes requires focus on the business of flying the airplane and the pilot should stick with maintaining aircraft control and paying attention to (aircraft) traffic, airspace, and obstacles such as towers and terrain. The pilot's distractions should be minimized with a passenger operating the CW radio and directing the desired search pattern. Every pilot should know minimum legal altitudes for the area flown, and those might be modified upward depending on conditions and the pilot's experience level. With that said, the required operations are no more difficult or dangerous than the ground reference maneuvers every private pilot must practice in their training, and they are in some sense a variation of what is done in the traffic pattern flown on every landing.

Beacon location by means of "fly by ear"

CW beacon location is begun as the aircraft approaches the general area where the payloads are suspected to have landed. Using a 7-element Arrow II® hand-held portable Yagi antenna (2) pointed in the direction of flight, we find we can begin to pick up the signal from around 20 miles (32 km) away. By flying an "S-pattern" as in Figure 1, we can listen to the signal intensity, determine the approximate direction of the source and fly to the general area. We know we've arrived when the signal becomes louder and the direction of greatest signal strength becomes more ambiguous. The altitude flown during this approach phase is not critical and can be at any altitude normally flown on a short-range cross country flight. Tip: the audio of the CW radio can be better heard in the noisy cockpit using an ear bud or by simply holding the radio's speaker near the mic input on a headset.

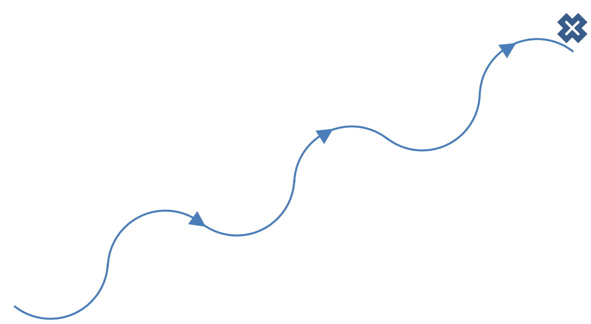

Figure 1. Flight path on initial approach to general area

After arriving at the general area, the signal is strong enough that we switch to a lower gain antenna.

A typical flight path once we arrive in the vicinity of the beacon is shown in Figure 2. The pattern involves two main components: circling while offset from the target and overflying the best guess of the target. During circling, approximate directions of maximum and minimum frequency can be determined to get an idea of the target direction. Then an overflight in that general direction can be made, noting the approximate location of where the frequency heard has seemed to have shifted halfway to from high to low while passing over the target. These midway positions can be marked as waypoints on a GPS and after several passes an average can be determined as the best guess of beacon location.

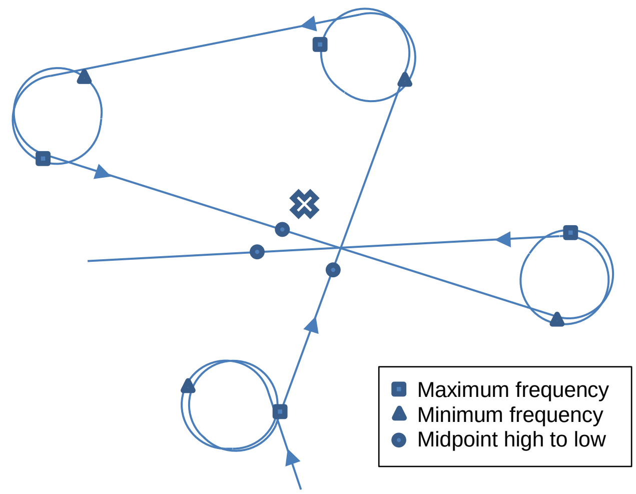

Figure 2. Flight path used to narrow down the location

Best results for this phase of the search are achieved by flying at a minimum safe altitude. Our target location accuracy seems to be within a few hundred meters when flown from an altitude of 1000'. However, on one recovery where we could get no lower than around 3000' above the target due to terrain, we found that our best guess was off by maybe a half mile. The main difficulty in narrowing down the location using this technique seems to be in determining the midpoint of the high to low frequency change while passing overhead at high altitude.

Computing beacon location from ping and GPS data

After a few aerial beacon hunting trips we started looking into developing a system that would automatically locate the target by listening to the frequency of the pings while recording the position and velocity from GPS information. The overall mathematical approach turned out to be relatively simple. The details are presented in the following.

The Doppler equation (3) is

$$f = \left(\frac{c + v_r}{c + v_s} \right) f_0$$

Where f is the frequency detected at the receiver, f_0 is the frequency of the source, c is the speed of light, v_r is the scalar velocity of the receiver relative to the source (positive towards the source) and v_s is the velocity of the source in the medium (zero in this case).

Since the source velocity is zero, the above equation can be written as

$$f = f_0 + \frac{v_r}{c} f_0$$

In our application we are interested in the Doppler shift in frequency, so we can rewrite this as

$$f - f_0 = \frac{f_0}{c} v_r$$

Note that the shift is therefore proportional to the relative velocity of the receiver, with the constant of proportionality given by the ratio of two constants, f_0 which is the radio frequency of the CW transmitter, 433.2 MHz, and c which is the speed of light. As described earlier, this change in frequency is the same for the audio frequency as it is for the radio frequency, that is, a 50 Hz change in radio frequency results in a 50 Hz change in audio frequency. We will therefore rewrite this equation using capital "F" for audio frequencies as:

$$F - F_{beep} = \frac{f_0}{c} v_r$$

where F is the audio frequency heard in the radio by a possibly moving observer and F_{beep} is the audio frequency that would be heard on the radio when stationary. Note that we will talk about these frequencies as if they were real, absolute quantities, but they are actually a function of the radio and its current tuning, and the radio should not be retuned as we are taking data. Also note that while the frequency shift for the remainder of our discussion is audio, f_0 is still the radio frequency of the beacon.

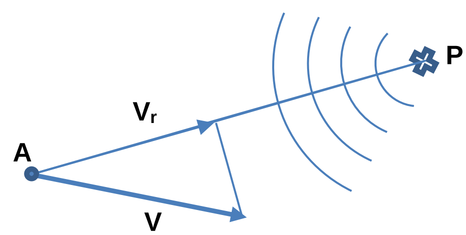

Figure 3. Relative velocity for determining Doppler shift

The scalar velocity of the receiver relative to the source, v_r, is easily found from Figure 3 to be the magnitude of the vector component of the aircraft velocity in the direction of the source. This is given by the dot product relationship

$$v_r = \frac{\boldsymbol{V} \cdot (\boldsymbol{P} - \boldsymbol{A})}{|\boldsymbol{P} - \boldsymbol{A}|}$$

where $\boldsymbol{P}$ is the position of the source, $\boldsymbol{A}$ is the position of the aircraft, $\boldsymbol{V}$ is the velocity of the aircraft, and $|\boldsymbol{P} - \boldsymbol{A}|$ represents the magnitude of that vector. Combining this with the frequency-shift equation above yields

$$F - F_{beep} = \frac{f_0}{c} \frac{\boldsymbol{V} \cdot (\boldsymbol{P} - \boldsymbol{A})}{|\boldsymbol{P} - \boldsymbol{A}|}$$

Now that we have this relationship, our problem can be stated as a problem of determining the unknown source position, $\boldsymbol{P}$ from a series of frequency observations, F_i, along with the aircraft positions, $\boldsymbol{A_i}$, and velocities, $\boldsymbol{V_i}$, at the time each was observed, keeping in mind that the (virtual) audio source frequency, F_{beep} is also unknown.

For a given potential beacon source location, $\boldsymbol{P}$ and the data from a given ping, we can compute the apparent audio source frequency based on that one observation as

$$F_{beep} = F_i - \frac{f_0}{c} \frac{\boldsymbol{V_i} \cdot (\boldsymbol{P} - \boldsymbol{A_i})}{|\boldsymbol{P} - \boldsymbol{A_i}|}$$

If that point in space happens to be the correct location of the beacon, we would expect this calculation for F_{beep} to be the same for all pings, although it would vary somewhat due to inaccuracies in the readings. At points in space other than the correct location we would expect this calculation for F_{beep} to vary among the different ping readings. This suggests the use of a least square error approach for computing the beacon location, assigning the beacon location to the position in space which minimizes the variance (least square error about the mean) in this source frequency calculation.

We can express our approach as follows. For any potential beacon location $\boldsymbol{P}$ and observation i we can use $B_i(\boldsymbol{P})$ to represent this computed source frequency

$$B_i(\boldsymbol{P}) = F_i - \frac{f_0}{c} \frac{\boldsymbol{V_i} \cdot (\boldsymbol{P} - \boldsymbol{A_i})}{|\boldsymbol{P} - \boldsymbol{A_i}|}$$

For a given prospective source position, $\boldsymbol{P}$ we determine the computed source frequency among all pings and compute the mean, $B(\boldsymbol{P})$, as

$$B(\boldsymbol{P}) = \frac{\Sigma\, B_i\,(\boldsymbol{P})}{N}$$

summed over all samples, where N is the number of samples. The mean square variation of all computed source frequencies for a given potential source point is given by

$$E(\boldsymbol{P}) = \frac{\Sigma\,(B_i\,(\boldsymbol{P}) - B(\boldsymbol{P}))^2}{N}$$

which can be shown to be equivalent to

$$E(\boldsymbol{P}) = \frac{\Sigma\, B_i^2(\boldsymbol{P})}{N} - B(\boldsymbol{P})^2$$

Our best guess of beacon position is then given by the position in space, P, for which this mean square variation is minimum. Note that we may need to guard against finding a local minima that is not the global minima.

A generic optimization function can be used to find this minimum. In our implementation, we use an adaptation of the “”amoeba” function given in *Numerical Recipes in C* (4). This method has the advantage of not having to supply any derivative information, while still operating rather quickly

Implementation

Figure 4 shows an implementation of our system using a Raspberry Pi® 3 Model B+ with a display, an Arduino and GPS module along with the Yaesu FT-817 radio.

Figure 4. Implementation using an Arduino Nano and Yaesu FT-817 radio

The Arduino® Nano is placed in-line between the radio’s audio output jack and a USB port on the Raspberry Pi. Its only job is to sample the audio output from the radio, detect when the beeps are occurring, then count the zero crossings during the beep period, and finally compute and report the beep frequency to the Raspberry Pi over a serial connection.

The GPS module is a GlobalSat® BU-353-S4 receiver which simply streams GPS data serially through USB.

On the Raspberry Pi, we’ve created a multi-threaded application that makes use of its multiple CPU cores, separating the different processes. One process collects GPS data, one collects ping data, one repeatedly computes the best guess of location from the data collected so far and one processes keyboard input.

We are also working on an alternate implementation, depicted in Figure 5, which uses a Software Defined Radio (SDR) in place of the Arduino and Yaesu radio resulting in a smaller, simpler setup.

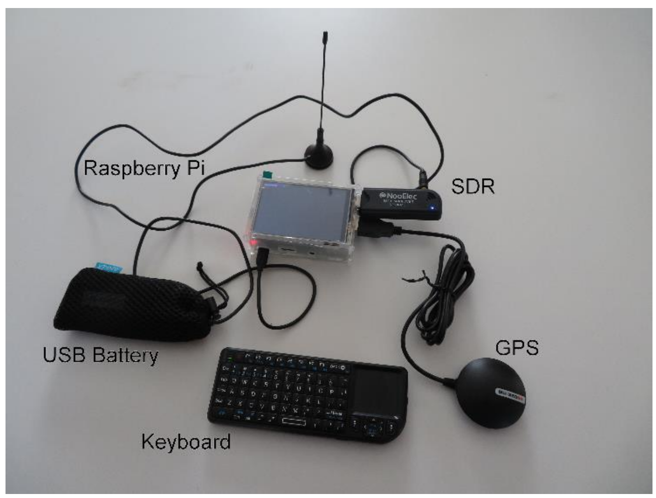

Figure 5 Alternate implementation using an SDR.

Here, the SDR performs the same function as did the Arduino and Yaesu radio, and it functions in nearly the same way, only it is looking at one component (I) of the IQ data rather than sampling the audio as the Arduino does.

Results

Figure 6 shows the aerial track flown on one of our test flights. In this case, we did some random maneuvering at about 1000’ (300 meters) above ground level in the general area of the beacon. The axis markings show meters from an arbitrary reference point. The “X” marks the actual location of the CW beacon.

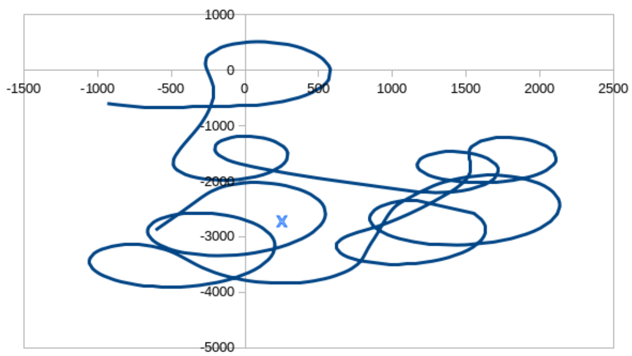

Figure 6 Track flown while collecting ping data

The frequency data shown on this run in given in Figure 7. It can be seen that the Arduino miscalculated the frequency on a number of pings. To correct this we created a filter that eliminated pings in which the forward differences from one ping to the next exceeded certain thresholds. The determination of ping frequency from the noisy signal from the radio seems to be one of our greatest challenges and we are still working on it.

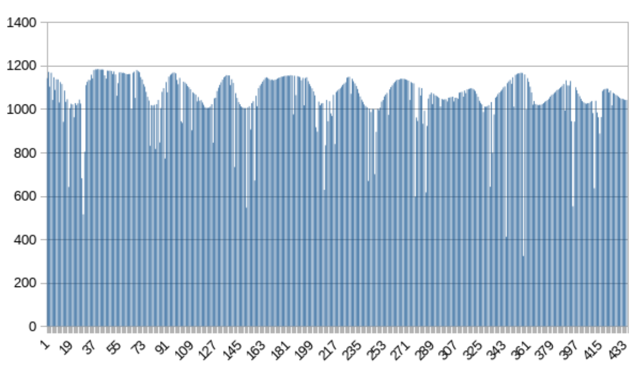

Figure 7 Frequency of the pings collected

Figure 8 shows the variance from all the (valid) samples taken at points within a kilometer in each direction and at the beacon altitude. Of course, lower values are in the center. You can see from this figure that the evaluation function is well behaved. In this case the optimization algorithm was able to locate the CW beacon horizontally to within 40 meters. We feel that we can improve on this accuracy by fine-tuning our frequency detection with the noisy signal.

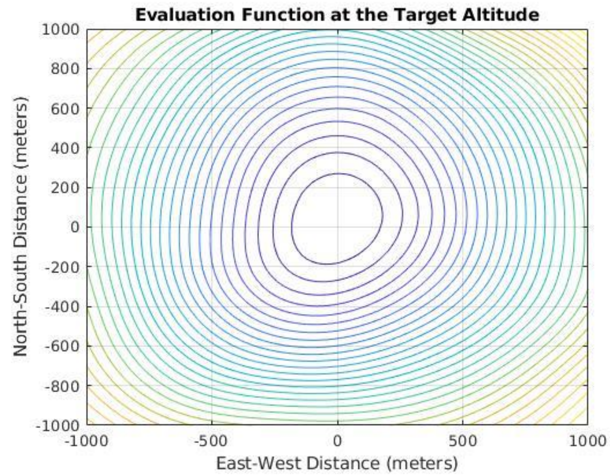

Figure 8 Evaluation Function at Target Altitude

Conclusions and Future Work

Our ability to locate a CW beacon from an airplane increased significantly when we started using the Doppler shift technique. Simply listening to the beacon’s frequency shift in the radio’s audio and using that to direct the flight path can greatly help to quickly get the beacon’s position with fair accuracy, while automating the process can do so to a much greater degree.

We are working on fine-tuning our system, so it can better detect the frequency of each ping with the noisy and sometimes weak signal that we have to work with. Improving on this would certainly do the most in improving the overall accuracy. The SDR is very convenient and does seem to give good frequency resolution, but we need to work on improving its sensitivity in our setup, as it seems to require a stronger signal than does the radio with Arduino. Also, cheaper SDR’s seem to suffer from drift in the detected signal and we are working on ways to account for this.

If we can get the system working on an SDR, one goal is to carry it aboard a drone that could circle overhead and get a better idea of the direction to the beacon. If we are not faced with the line-of-site operation limitation, the drone could potentially automatically locate the beacon with fairly high accuracy, as it can, and must legally, fly at lower altitudes than a manned aircraft safely can.

We have tried our setup using an automobile driving down nearby roads and that shows some promise. Automobile speeds, even at 45 mph are still 1/3 the aircraft speeds. The Doppler shifts are proportionally smaller, but that should still be enough to help in beacon location under the right conditions. We are confident we’ll get better results as we improve our frequency detection.

Finally, we are also considering an app for a tablet that simply listens to the audio output from the radio and, using its internal GPS, can do the work of our present system.

Acknowledgment

This work was made possible through the generous support we have received from the Maine Space Grant Consortium for our high altitude ballooning program.

References

1 BeeLine Transmitter by Big Red Bee, http://www.bigredbee.com/BeeLine.htm.

2 Arrow II Handheld Portable Antenna, http://www.arrowantennas.com/arrowii/440-7ii.html

3 Doppler effect, *Wikipedia*, https://en.wikipedia.org/wiki/Doppler_effect

4 William H. Press, Teukolsky, S.A, Vetterling, W.T and Blannery, B.P., *Numerical Recipes in C, 2nd Edition, Section 10.4, 1992*

---

[a] University of Maine.
[1] To whom correspondence should be addressed. E-mail: reason@maine.edu