Principles of Echolocation
in Bats and Machines

Note 9

Range or Doppler? Choosing the Call

Ambiguity functions of sonar calls and bat calls

Jan Steckel
Cosys-Lab, University of Antwerp, and Flanders Make, Belgium jan.steckel@uantwerpen.be

Version 2026-11-18 16:29

1 Why the Call Matters

A sonar measures range by timing an echo, and it times an echo by correlating it with the call it sent. This correlation is called the matched filter. How sharp the correlation peak is, and what happens to it when the reflector moves, is decided by the call alone. A long tone, a short sweep and a sweep with a particular curvature each make a different trade-off between measuring range and measuring velocity. In this note, we compare five calls with one tool, the ambiguity function. Three of them are common in machines, one is inspired by horseshoe bats, and the last one, a multisine, lets several emitters send at the same time.

2 The Matched Filter and the Ambiguity Function

Let $s(t)$ be the emitted call. A reflector that approaches the sonar with radial velocity v does not just delay the echo. Because the reflector moves while the call is reflected, the echo is compressed in time by the Doppler factor

$$\eta = \frac{c+v}{c-v} \approx 1 + \frac{2v}{c}, \tag{1}$$

with $c = 343$ m/s the speed of sound. At 5 m/s, $\eta = 1.0296$, so every frequency in the call rises by 3%. For a reflector at range R, which returns its echo after $\tau_0 = 2R/c$, the echo is

$$e(t) = \sqrt{\eta}\, s\big(\eta\,(t-\tau_0)\big). \tag{2}$$

The argument $\eta\,(t-\tau_0)$ says that the echo is the call, delayed by τ_0 and played back η times faster. The factor $\sqrt{\eta}$ in front is only a normalization. Playing the call faster makes it shorter by a factor η, and with $\sqrt{\eta}$ the echo keeps the energy of the call, so that the matched filter of a moving reflector can be compared directly with that of a stationary one. Without it, the echo at 5 m/s would be 0.13 dB weaker, and the shapes of the curves that follow would not change at all. The matched filter correlates the echo with the call. Measuring the delay τ from τ_0, it computes

$$\chi(\tau, v) = \sqrt{\eta} \int s(\eta t)\, s^*(t-\tau)\, dt. \tag{3}$$

Its magnitude, over delay τ and velocity v, is the (wideband) ambiguity function. Along the delay axis at $v = 0$, it shows how well two reflectors at different ranges are separated. Along the velocity axis, it shows how quickly the matched filter loses the echo of a moving reflector.

In radar, the compression is usually approximated by a frequency shift, the Doppler shift $f_D = 2vf/c$, which is the same for all frequencies in the narrow band of a radar pulse. For in-air sonar, this narrowband approximation breaks down. Our calls span a factor of four in frequency, so the Doppler shift is four times larger at the top of the call than at the bottom, and a call of a few milliseconds is long enough for the compression to matter. A practical criterion is that the approximation holds while $(2v/c)\,BT \ll 1$, with B the bandwidth and T the duration. For a sweep of 60 kHz over 3 ms, $BT = 180$, and the criterion is already violated at 1 m/s. We therefore compute the wideband ambiguity function directly, by compressing the call (Section 4).

3 Five Calls

Figure 1 shows the five calls in time and frequency.

- The constant-frequency (CW) tone lasts 3 ms at 50 kHz.
- The linear FM (LFM) sweep falls from 80 to 20 kHz in 3 ms, at a constant rate.
- The hyperbolic FM (HFM) sweep covers the same band in the same time, but its period $1/f$ grows at a constant rate, so it falls fast at high frequencies and slowly at low ones. This is the type of call that the eRTIS sends, also known as linear-period modulation.
- The CF-FM call is modeled on horseshoe bats. A long constant-frequency part (10 ms at 80 kHz) is followed by a short FM tail (2 ms, 80 to 65 kHz).
- The multisine (MS) sends 181 tones from 20 to 80 kHz at the same time, for 3 ms. The tones are 333 Hz apart, the inverse of the duration, and each has a random phase.

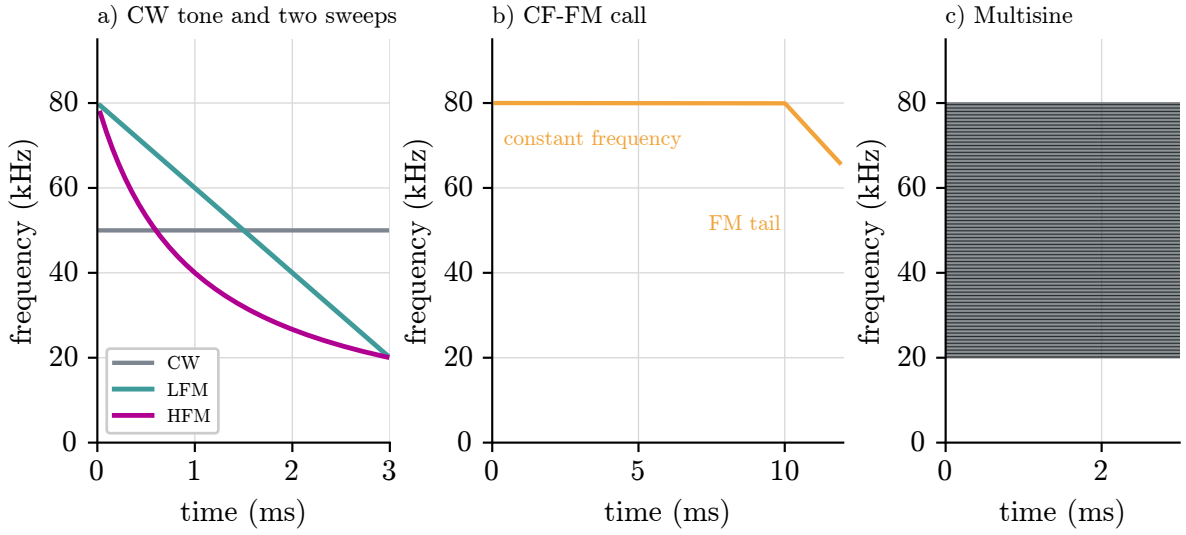


Figure 1: The five calls in time and frequency. a, b) The instantaneous frequency of the tone, the two sweeps and the CF-FM call. c) The multisine has no single instantaneous frequency. All its 181 tones sound during the whole call, so they fill the band.

4 What the Matched Filter Sees

Each call gets its own figure below (Figs. 2 to 6). On the left is its wideband ambiguity function, the matched-filter output over range offset ($c\tau/2$) and velocity (positive when the reflector approaches), in dB relative to the peak for a stationary reflector. On the right are three horizontal

cuts through it, for a reflector at rest, one approaching at 1 m/s and one at 5 m/s. Delays refer to the position of the reflector at the middle of the call. All maps use the same velocity axis and colour scale, but note the different range axes.

4.1 The CW Tone

The CW tone gives a correlation peak as long as the call itself (Fig. 2). Its range resolution, the width of the peak at −3 dB, is 294 mm. In return, it is very sensitive to motion. At 1 m/s, the 50 kHz tone is shifted by 292 Hz, almost the inverse of its duration, and the peak has already dropped by 8.2 dB. A tone is therefore a good call for measuring velocity and a poor one for measuring range.

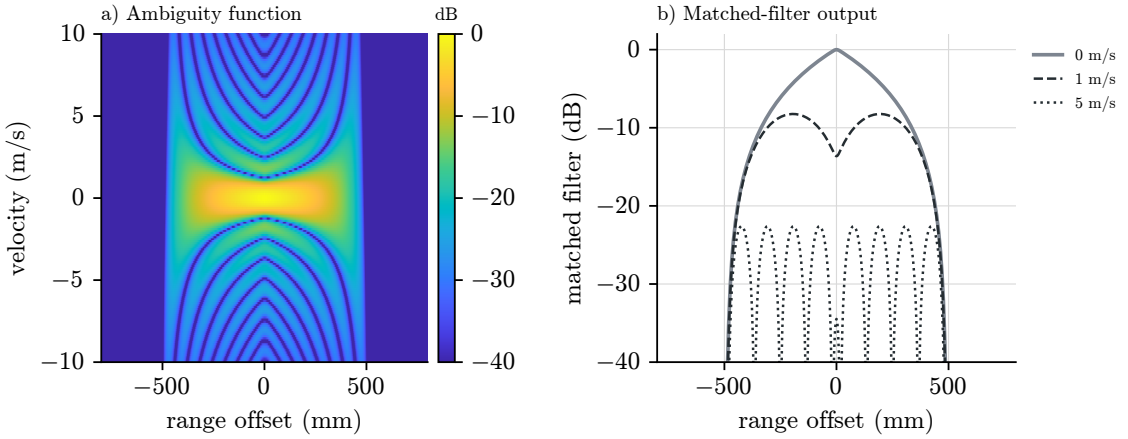


Figure 2: The CW tone (3 ms, 50 kHz). a) Ambiguity function. b) Matched-filter output for a reflector at rest and for one approaching at 1 and 5 m/s.

4.2 The Linear FM Sweep

The linear sweep spreads 60 kHz of bandwidth over the call, and its peak shrinks to 2.4 mm (Fig. 3). For a stationary reflector, that is the best of the five calls. Under motion, its ambiguity function shows the well-known tilted ridge. At low velocities, the peak shifts in range by about 2.5 mm per m/s, the range-Doppler coupling that the narrowband theory predicts. The ridge is short, though. Above about 2.0 m/s, the compressed sweep no longer fits the emitted one at all frequencies at once, and the peak drops and widens. At 5 m/s, it is 7.8 dB lower and 20 mm wide.

4.3 The Hyperbolic FM Sweep

The hyperbolic sweep has almost the same resolution (3.1 mm), and it behaves very differently under motion (Fig. 4). Its peak stays as high and as narrow as for a reflector at rest. Even at 10 m/s, it loses only 0.3 dB. A hyperbolic sweep has the special property that compressing it in time is almost the same as shifting it in time, so the matched filter finds the same peak at a slightly different delay (Section 5). That delay is a range error of 2.5 mm per m/s, or 12.3 mm at 5 m/s. It is predictable, so it can be corrected once the velocity is known, and for a mapping robot it is usually negligible. A call whose peak survives motion like this is called Doppler tolerant.

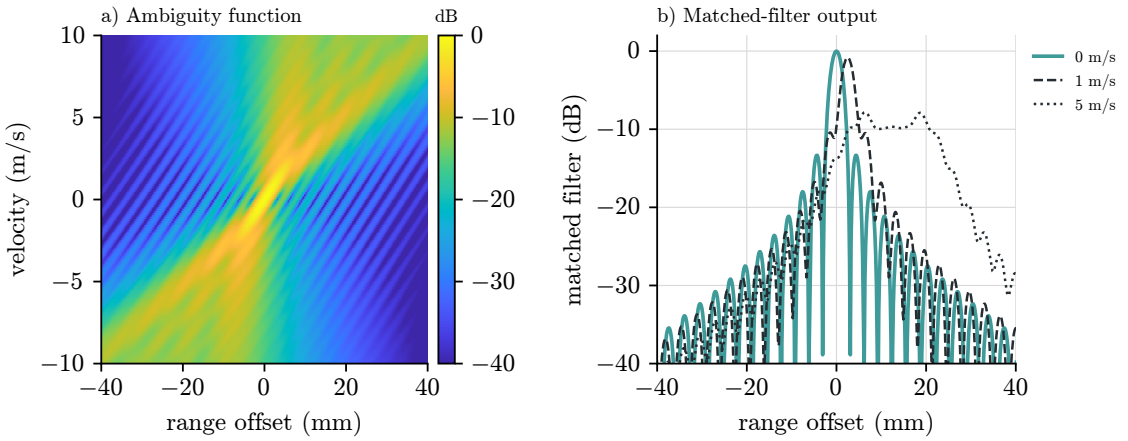


Figure 3: The linear FM sweep (3 ms, 80 to 20 kHz). a) Ambiguity function. b) Matched-filter output for a reflector at rest and for one approaching at 1 and 5 m/s.

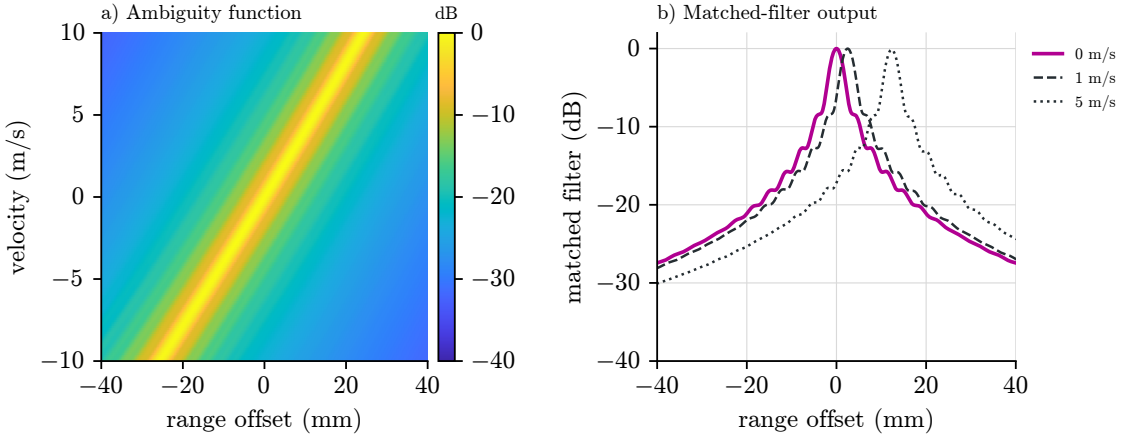


Figure 4: The hyperbolic FM sweep (3 ms, 80 to 20 kHz). a) Ambiguity function. b) Matched-filter output for a reflector at rest and for one approaching at 1 and 5 m/s.

4.4 The CF-FM Call

The CF-FM call has a long tone and a short sweep in one call (Fig. 5). The long CF part gives a broad correlation peak, 592 mm wide, but a very fine sensitivity to velocity. A 10 ms tone at 80 kHz is shifted by 466 Hz per m/s, and its inverse duration corresponds to a velocity of only 0.21 m/s. Its peak drops by 3 dB at 0.10 m/s. The short FM tail adds a narrow peak on top (on its own, 11 mm wide) for ranging. A single matched filter for the whole call mixes these two very different peaks, and Section 7 shows how to receive the two parts separately.

4.5 The Multisine

The multisine covers the same band as the two sweeps, so its peak is just as narrow, 2.4 mm (Fig. 6). Because the phases of its tones are random, the call looks like a short burst of noise, and the matched filter has random sidelobes around the peak, the highest 13 dB below it. Under motion, the multisine behaves like a long tone. Every tone shifts by $2vf/c$, which at 80 kHz is 466 Hz per m/s, more than the spacing of 333 Hz between the tones. At 1 m/s, the tones of the echo therefore fall between the tones of the call, and the matched filter no longer finds a match. The peak drops by 3 dB at 0.55 m/s and by 10.7 dB at 1 m/s, into the sidelobes.

The multisine thus has the narrow range peak of a sweep and the velocity sensitivity of a long

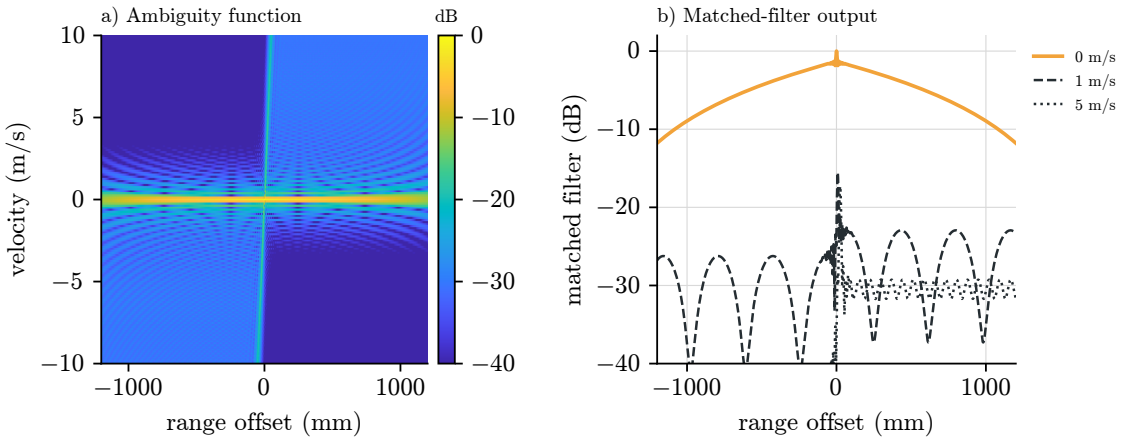


Figure 5: The CF-FM call (10 ms at 80 kHz, then 2 ms from 80 to 65 kHz). a) Ambiguity function. b) Matched-filter output for a reflector at rest and for one approaching at 1 and 5 m/s.

tone at the same time. Its ambiguity function is a single narrow spike, often called a thumbtack. There are two drawbacks. First, a matched filter tuned to the wrong velocity misses the echo, so a receiver needs a bank of filters, one for every velocity (Section 8). Second, the random tones add up to a peaky signal. Its crest factor (peak over rms) is 10.7 dB, against 3.0 dB for a sweep, so an emitter with a given peak amplitude puts less energy into a multisine than into a sweep.

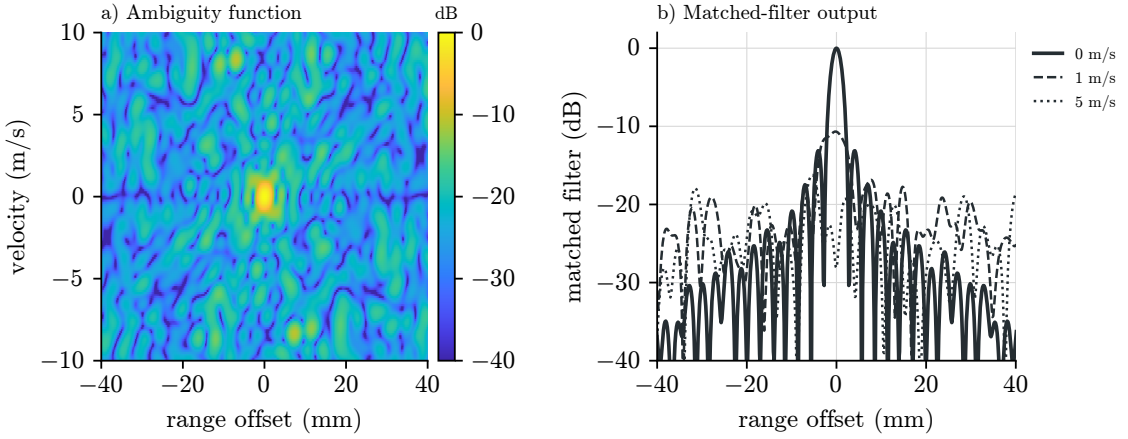


Figure 6: The multisine (3 ms, 181 tones from 20 to 80 kHz with random phases). a) Ambiguity function. b) Matched-filter output for a reflector at rest and for one approaching at 1 and 5 m/s.

4.6 The Five Calls Side by Side

Figure 7a shows how fast the peak of each call drops with velocity. The hyperbolic sweep is the only call that keeps its peak over the whole range of velocities. The linear sweep holds on up to about 2.0 m/s. The tone, the CF-FM call and the multisine lose their peak within a fraction of a metre per second, which is what makes them useful for measuring velocity. Figure 7b shows the range shift of the two sweeps. Both follow the narrowband prediction closely, the linear sweep only as long as its peak holds.

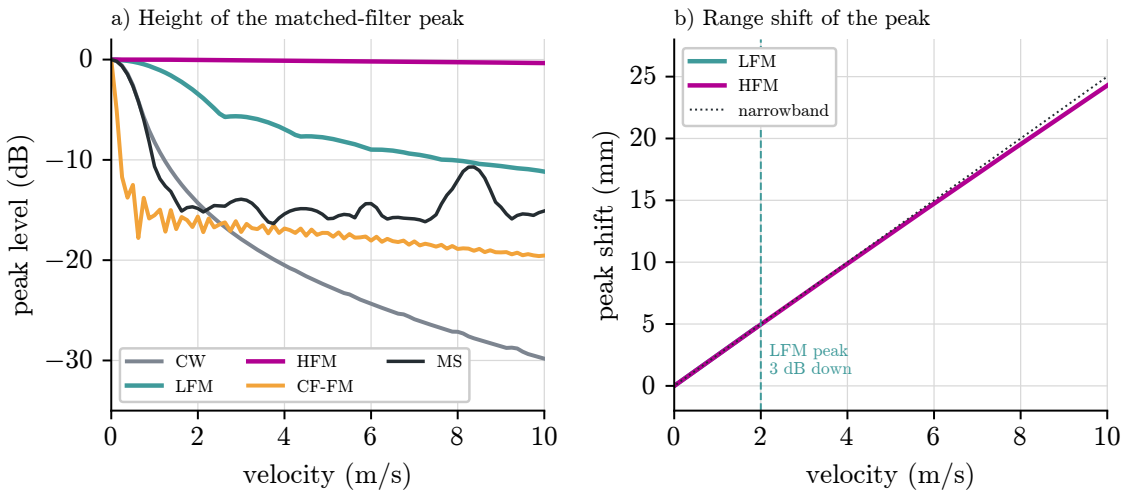


Figure 7: a) The height of the matched-filter peak against velocity, for the five calls. b) The range shift of the peak for the two sweeps (the LFM only as long as its peak stays within 3 dB), with the narrowband prediction (dotted).

Key idea

Range resolution needs bandwidth. Velocity resolution needs a long call whose echo stops matching the call as soon as the reflector moves, like a long tone or a multisine. A sweep has the bandwidth, and when the reflector moves, its peak shifts in range instead of disappearing. A hyperbolic sweep, like the eRTIS call, does so with almost no loss, which makes it the natural call for a moving sonar that measures range. A call that is sharp in range and velocity at once, like the multisine, needs a bank of matched filters, one for every velocity.

5 Doppler Tolerance of the Hyperbolic Sweep

The difference between the two sweeps is easiest to see in time and frequency. Let a sweep start at f_0 and pass frequency F at time $t_s(F)$. The echo of a moving reflector, $s(\eta t)$, contains frequency F where the call contained F/η, but η times earlier, at $t_e(F) = t_s(F/\eta)/\eta$. For the hyperbolic sweep, the instantaneous frequency is $f(t) = f_0/(1+bf_0t)$, with $b = (1/f_1 - 1/f_0)/T$, so

$$t_s(F) = \frac{1}{bF} - \frac{1}{bf_0}, \qquad t_e(F) = \frac{1}{bF} - \frac{1}{\eta\, bf_0}, \qquad t_e(F) - t_s(F) = \frac{1}{bf_0}\left(1 - \frac{1}{\eta}\right). \tag{4}$$

The difference does not depend on F. Every frequency of the echo arrives by the same amount later, so the compressed call is the same call, shifted in time (Fig. 8a). The matched filter therefore finds the full peak, only displaced. For the linear sweep, $t_s(F) = (F - f_0)/k$, and the same calculation gives a difference that grows with F. Across the band, the frequencies of the echo are displaced by different amounts, spread over about $2(\eta - 1)T$ (Fig. 8b). Once that spread exceeds the duration of the correlation peak, the peak drops and widens.

6 A Worked Example with the eRTIS Call

The eRTIS sends a hyperbolic sweep of 2.5 ms from about 75 to 20 kHz. Consider two cases, a robot that drives at 1 m/s towards a wall, and a bat that flies at 5 m/s. Table 1 compares the eRTIS call with a linear sweep over the same band and duration.

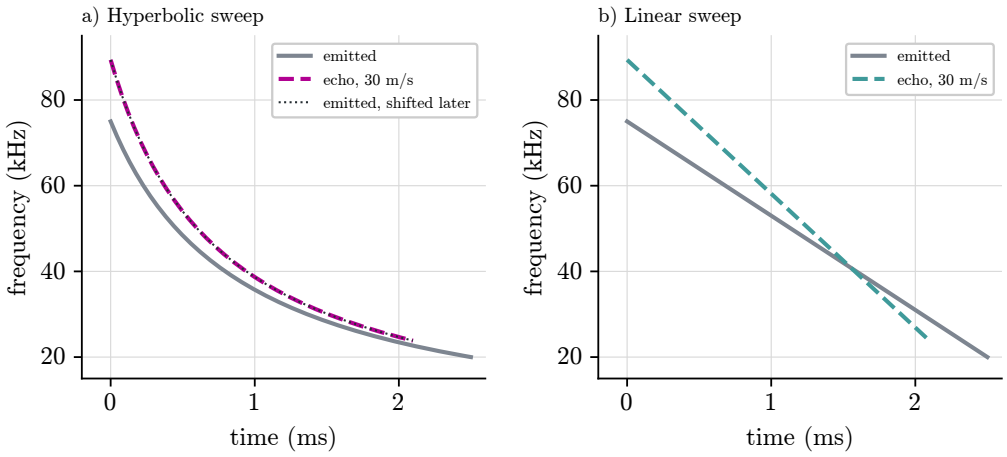


Figure 8: The instantaneous frequency of an eRTIS-like call (2.5 ms, 75 to 20 kHz) and of the echo of a reflector approaching at 30 m/s (exaggerated, so that the effect is visible). a) Hyperbolic sweep. The echo lies exactly on the emitted sweep, shifted later (dotted, hidden under the echo). b) Linear sweep. The echo crosses the emitted sweep, and no single shift fits all frequencies.

	robot, 1 m/s	bat, 5 m/s
Doppler factor η	1.0058	1.0296
frequency shift at 75 / 20 kHz (Hz)	439 / 117	2219 / 592
hyperbolic: time shift of every frequency (μs)	5.3	26.1
hyperbolic: peak loss (dB)	0.0	0.1
hyperbolic: range shift of the peak (mm)	2.2	10.6
linear: spread of the frequencies (μs)	29	148
linear: peak loss (dB)	0.5	6.5

Table 1: The eRTIS call (hyperbolic, 2.5 ms, 75 to 20 kHz) and a linear sweep over the same band, for a reflector approaching at 1 and 5 m/s. Range resolution: 3.4 mm (hyperbolic) and 2.7 mm (linear). Range shifts refer to the position of the reflector at the middle of the call.

For the robot, both calls work. The linear sweep spreads its frequencies over 29 μs, about 5 mm of range, which is twice its resolution, and loses 0.5 dB. For the bat at 5 m/s, the spread grows to 148 μs or 25 mm, ten times the resolution, and the linear sweep loses 6.5 dB. The hyperbolic sweep loses only 0.1 dB. Its peak moves by 10.6 mm, of which 4.5 mm come from the time shift of the hyperbola and 6.3 mm from the motion of the reflector during the first half of the call. For a robot or a bat that maps a static scene while moving, such a shift is the same for every reflector at the same radial velocity, and it is small compared with the distances that matter.

7 Two Receivers for One Call

The CF-FM call is best received with two different receivers (Fig. 9). The CF part goes to a bank of narrow filters, here 100 Hz wide, the inverse of its 10 ms duration. A reflector approaching at 1 m/s shifts the 80 kHz echo by 466 Hz, into a channel four or five steps up, and the channel with the most energy tells the velocity to within about 0.21 m/s. The FM tail goes to a matched filter. It is short and spans only 15 kHz, so its peak is 11 mm wide, and it keeps that peak under motion (1.6 dB loss at 5 m/s). But its shallow sweep couples range and velocity strongly, and at 5 m/s, its peak moves by 48 mm. The two receivers therefore complement each other. The velocity measured by the filter bank gives the size of that shift, so that the range can be corrected.

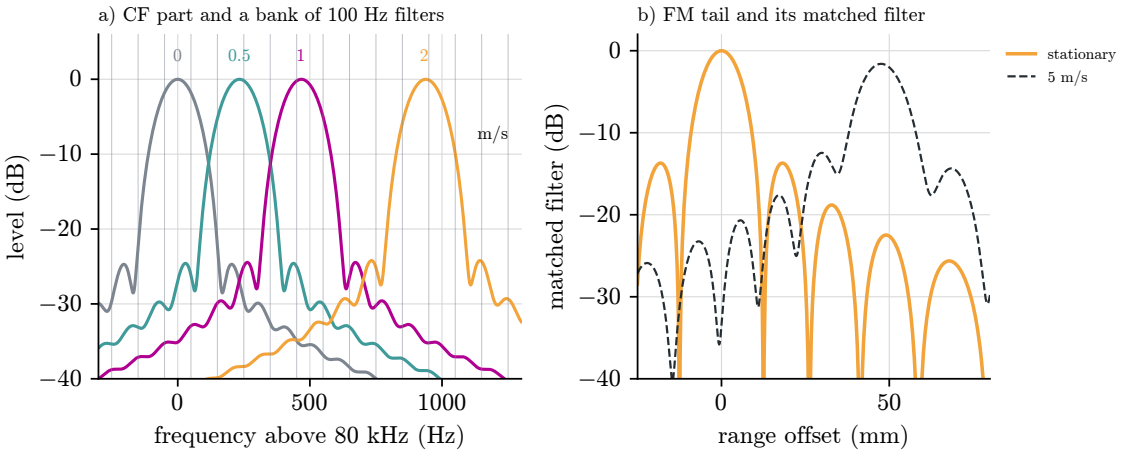


Figure 9: The two receivers of a CF-FM call. a) The spectrum of the echo of the 10 ms CF part at 80 kHz (Hann window), for reflectors approaching at 0, 0.5, 1 and 2 m/s, with the channels of a 100 Hz filter bank (gray lines). b) The matched-filter output of the 2 ms FM tail (80 to 65 kHz), stationary and at 5 m/s.

8 Several Emitters at the Same Time

Some sonars use more than one emitter. In ST-STAB (simultaneous transmit, synthetic aperture beamforming), an array of emitters sends at the same time, each emitter with its own code, and every receiver separates the emitters again by correlating its signal with each code (Steckel, 2015). The codes must be orthogonal, which means that the matched filter of one emitter should not respond to the call of another. Multisines make this easy. Split the 181 tones of the multisine at random between two emitters, A with 90 tones and B with 91 tones. Two tones 333 Hz apart are orthogonal over the 3 ms call, so the matched filter of A responds to the echo of B only at the level of its sidelobes, 22 dB below its own peak (Fig. 10b, at 0 m/s).

Motion breaks this separation. At 1 m/s, a tone at 65 kHz moves up by 379 Hz, more than the spacing between the tones, so the tones of A land between and next to tones of B (Fig. 10a). A matched filter of A tuned to a reflector at rest then loses 12 dB on the echo of A, while the echo of B reaches 17 dB below the original peak. The two emitters can hardly be told apart any more (Fig. 10b).

The remedy is the same compression that caused the problem. For every velocity v_k in a set, the receiver compresses the code of A by $\eta(v_k)$ and uses it as a matched filter. Figure 10c shows such a bank, with filters 0.25 m/s apart, for an echo of A from a reflector approaching at 2 m/s. The filter for 2 m/s finds the full peak, its neighbours lose 0.7 dB, and the filters further away see only sidelobes. In the filter for 2 m/s, the echo of B is 22 dB down again. The bank thus separates the emitters again and measures the velocity of the reflector at the same time. This is the step from ST-STAB to vSTAB, the work of Kerstens and Steckel, in which the velocity of each reflector is estimated together with its range and direction. The cost is computation. Covering ±5 m/s in steps of 0.25 m/s takes 41 matched filters per emitter, for every receiver channel.

9 What Bats Do

Bats use both kinds of call. Many species, such as the big brown bat *Eptesicus fuscus*, emit short, broadband FM calls, which suit precise ranging. Altes and Titlebaum pointed out that the FM calls of several bats resemble hyperbolic sweeps, and argued that these calls are Doppler tolerant for this reason. Horseshoe bats instead emit long CF calls with a short FM tail. Their

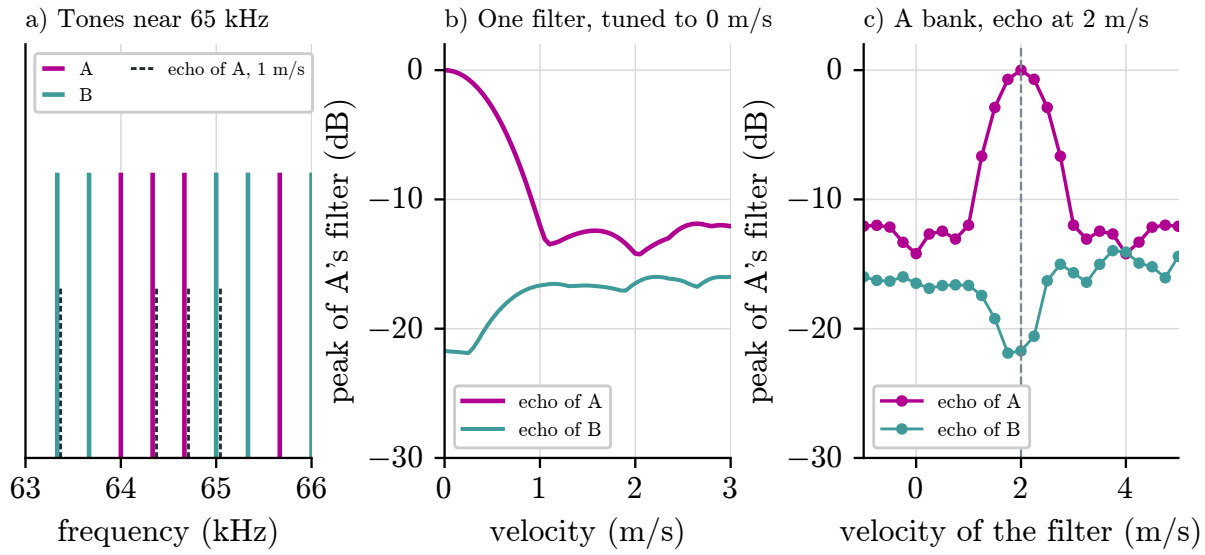


Figure 10: Two emitters with multisine codes, A (90 tones) and B (91 tones), drawn at random from the 181 tones of Fig. 6. a) Some tones of A and B near 65 kHz, and the tones of A in the echo of a reflector approaching at 1 m/s. b) The peak of the matched filter of A, tuned to a reflector at rest, for the echo of A and of B, against the velocity of the reflector. c) The peak of a bank of matched filters of A, each tuned to another velocity, for echoes of A and of B from a reflector approaching at 2 m/s.

CF part is sensitive to small changes in frequency, which allows them to detect the rhythmic glints of a fluttering insect's wings against a static background. Their hearing is sharply tuned to a narrow band just above their emitted frequency, and they lower their emitted frequency in flight so that the echoes return in that band, a behavior known as Doppler shift compensation (Schnitzler, 1968). The bat thus uses its CF part as a Doppler sensor and its FM tail for ranging, much like the two receivers of Section efsec:twor.

Exercise

1. A bat flies at 5 m/s towards a wall and emits a CF call at 83 kHz. At which frequency does the echo return? By how much must it lower its call to hear the echo at 83 kHz?

2. Why is the range shift of the hyperbolic sweep harmless for a robot that builds a map of a static scene, but not for one that estimates the range rate of a moving object?

3. For which call duration and bandwidth would the narrowband approximation be valid at 5 m/s? Is such a call still useful for an in-air sonar?

4. The bank of Doppler filters in Section 8 has its filters 0.25 m/s apart. Use the velocity at which the peak of the multisine has dropped by 3 dB to estimate the largest spacing for which an echo never loses more than 3 dB. How many filters does that save for ±5 m/s?

Further Reading

- J. Kroszczyński, "Pulse compression by means of linear-period modulation," Proc. IEEE, 57(7), 1260–1266, 1969.

- R. A. Altes and E. L. Titlebaum, “Bat signals as optimally Doppler tolerant waveforms,” J. Acoust. Soc. Am., 48(4B), 1014–1020, 1970.
- H.-U. Schnitzler, “Die Ultraschall-Ortungslaute der Hufeisen-Fledermäuse (Chiroptera-Rhinolophidae) in verschiedenen Orientierungssituationen,” Z. vergl. Physiol., 57, 376–408, 1968.
- J. Steckel, “Sonar system combining an emitter array with a sparse receiver array for air-coupled applications,” IEEE Sensors J., 15(6), 2015.
- N. Levanon and E. Mozeson, *Radar Signals*, Wiley, 2004.