# Detecting unmanned aircraft, and finding the operator

**Why the emission that matters most is the one that is hardest to hear.**

ABSTRACT

Radio detection of small unmanned aircraft is usually described as a sensitivity problem. It is not. It is a probability-of-intercept problem, and the arithmetic is unforgiving: against a burst lasting 645 microseconds once every 600 milliseconds, a scanning receiver needs about ninety seconds to reach ninety percent confidence while a gap-free receiver needs the emitter's own repetition interval. This paper derives that result, works the link budget for both ends of the link, sets out what each direction-finding method can actually deliver, and states plainly where radio detection is blind.

WHO THIS PAPER IS FOR

**Counter-UAS system integrators** designing or specifying an RF sensing layer. **Security and site-protection engineers** evaluating what an RF sensor can and cannot promise. **Spectrum enforcement staff** working unauthorized flight cases.

CONTENTS



## 1 Two emitters, two geometries

Almost every small unmanned aircraft in service is a pair of radios. The controller sends a narrow, low-duty-cycle uplink. The aircraft returns a wide, high-duty-cycle video downlink. They live in the same bands, they belong to the same system, and for the purpose of finding the person flying, they could hardly be less alike.

The instinct at this point is to reach for a cheap scanning detector, the kind sold as a drone-band monitor, and see what it finds. It will find something, which is the problem. Section 2 shows that such a receiver hears a hopping link about one percent of the time, so what it reports is a sparse and unrepeatable sample of the traffic, and an operator cannot tell a quiet sky from a missed one. Buying a second and a third does not fix it either: three receivers each blind ninety-nine percent of the time are still blind, and now the false alarms triple. The constraint is not sensitivity or channel count. It is coincidence, and coincidence is arithmetic.

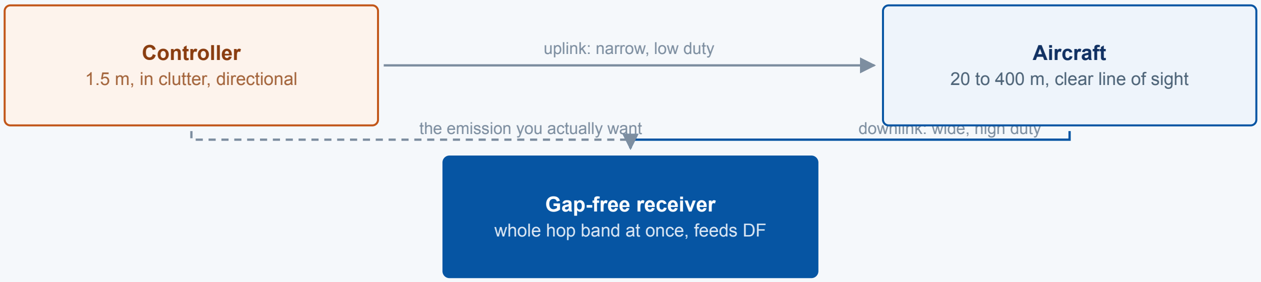


*The downlink is loud, wide and unobstructed. The uplink is quiet, narrow and buried in ground clutter, and only the uplink leads to the operator.*

**Figure 1.** Two emitters, two geometries, two problems. Detection is usually won on the downlink and the operator is only ever found on the uplink, which is the harder of the two by twenty to forty decibels of geometry alone. SCHEMATIC

Figure 1 is the whole difficulty in one picture. The aircraft sits in clear air with a quasi-omnidirectional antenna, transmitting tens of megahertz at high duty cycle. The controller sits a meter and a half above the ground, behind whatever the ground has on it, with an antenna pointed at the aircraft rather than at you, transmitting a short slot.

There is a third emission, and it is the easiest of all: the mandated Remote ID broadcast. Under ASTM F3411 and 14 CFR Part 89 a compliant aircraft transmits its identity, its position and its control station position in the clear over Bluetooth or Wi-Fi, at least once per second. The European requirement under Regulation (EU) 2019/945 is equivalent and explicitly names the remote pilot's position. That broadcast is designed to be received by anyone. It is also, by construction, only present on aircraft whose operators intended to comply.

## 2 Why detection is a timing problem, not a sensitivity problem

Control links hop. Regulation requires it: a frequency-hopping device in the 2.4 GHz band must use at least fifteen channels and may not occupy any one of them for more than 0.4 seconds within a defined period. The practical result is a set of emitters that are on a given frequency for a millisecond or less at a time.

Table 1 gives representative figures for the link classes a sensor actually meets.

**Table 1.** Representative hopping behavior for control and identification links. Figures are drawn from open-source implementations and published reverse engineering rather than vendor datasheets, and should be treated as representative rather than exact.

| Link class | Band | Channels | RF on per hop | Hop rate |
|---|---|---|---|---|
| Consumer RC, 16-channel | 2.4 GHz | 47 | about 3 ms | about 111 /s |
| Long-range RC, 500 Hz mode | 2.4 GHz | about 80 | about 0.25 ms | 500 /s |
| Long-range RC, sub-GHz | 868 / 915 MHz | about 50 | about 1.5 ms | 150 /s |
| Proprietary video link identifier | 2.4 / 5.8 GHz | wide OFDM | 645 us burst | about 1.7 /s |

Set a scanning receiver against that. If it covers a hop band of width $B_h$ in steps of its instantaneous bandwidth $B_i$, dwelling d at each step with a retune time $t_r$, its revisit period is the product of the number of steps and the time per step. The probability of catching any one burst is the fraction of the revisit period during which the dwell and the burst overlap by at least the detector's minimum observation:

$$P_1 = \min[1, (d + \tau - 2t_{min}) / T], \quad T = N_s(d + t_r) \qquad (1)$$

where

$d$ is the dwell at each tuning step, in seconds; 1 ms is representative of a fast scanner

$\tau$ is the burst duration, in seconds; 645 μs for the identification burst worked below

$t_{min}$ is the shortest contiguous observation the detector needs; 50 μs here

$N_s$ is the number of tuning steps across the hop band; 84 at 1 MHz across 83.5 MHz

$t_r$ is the retune and settling time, in seconds; 0.2 ms

Detection over many bursts follows a geometric law, so the bursts needed for ninety percent confidence is the log of one tenth over the log of one minus the single-burst probability. Working that through for a 645 microsecond identification burst repeating every 600 milliseconds gives the numbers in Table 2.

**Table 2.** Time to ninety percent confidence against a 645 microsecond burst repeating every 600 milliseconds, computed from equation 1 across an 83.5 MHz hop band in 84 steps, with a 1 ms dwell, a 0.2 ms retune and a 50 microsecond minimum observation. Change any of those five and the column changes with it. The gap-free figure is set by the emitter's repetition interval, not by the receiver.

| Receiver | Instantaneous BW | Revisit | P per burst | Time to 90 percent |
|---|---:|---:|---:|---:|
| Narrow scanner | 1 MHz | 100.8 ms | 0.015 | 89 seconds |
| Mid scanner | 20 MHz | 6.0 ms | 0.26 | 4.6 seconds |
| Gap-free capture | at least 100 MHz | continuous | 1.00 | 0.6 seconds |

> **THE NUMBER WORTH REPEATING**
>
> Eighty-nine seconds against six tenths of a second. An aircraft traveling at fifteen meters per second covers 1.3 kilometers while the narrow scanner is still accumulating confidence. That is the difference between a warning and a report written afterward.

There is a second, quieter penalty. A scanner observes any given channel for only the fraction of wall-clock time given by dwell over revisit, so it forfeits ten times the logarithm of that ratio in sensitivity. For the narrow scanner in Table 2 that is 20 dB, which at a free-space path-loss exponent is a factor of ten in range. The scanning receiver is not merely slower. It is also deaf by twenty decibels while it is being slow.

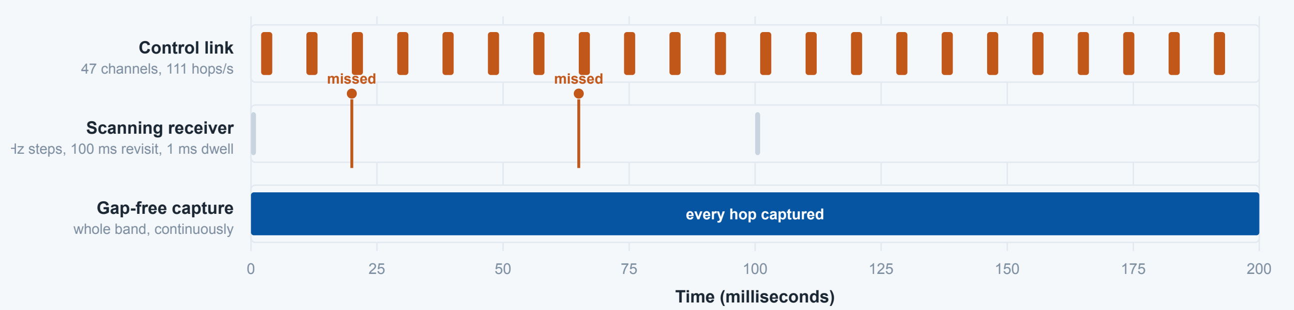


**Figure 2.** A hopping link fills the time-frequency plane with short marks while a scanner traces a single line through it. What determines detection is coincidence, not coverage. SCHEMATIC

Figure 2 shows where the twenty decibels go. Both the scanner's marks and the link's marks are sparse, and the receiver hears the link only where the two coincide, which for a narrow scanner is about one percent of wall-clock time.

Sections 1 and 2 are a criterion, not yet a specification. Table 3 states it as four requirements a reader can score any receiver against, and the reason they are stated here rather than at the end is that a reader who stops at this table should still be able to buy correctly.

**Table 3.** Four requirements, and the third is the one to ask about. R1 and R2 are published and answer sections 1 and 2 directly. R4 is published, but read it carefully: the instrument is not what limits a time-difference fix. R3 is the requirement a direction-finding system is actually built on and the one no datasheet in this class states, which is why section 4 spends its length there.

| # | Requirement | Derived value | Published |
|---|---|---|---|
| R1 | Instantaneous bandwidth covering the whole hop band at once | ≥ 83.5 MHz for the 2.4 GHz band, from §2 | 100 MHz; 50 MHz standard on the 9.5 GHz module |
| R2 | Probability of intercept published as a guarantee with its transform size | ≤ 645 μs to catch an identification burst whole | 0.512 μs at $N$ = 32; 32.768 μs at $N$ = 2048 |
| R3 | Coherent channels on one time base for a bearing | ≥ 2 for an interferometer, ≥ 4 for an array | Not published as a matched multi-channel figure *(verify)* |
| R4 | Time tagging fine enough that the instrument is not the limit | ≤ 10 ns, since 1 ns is 30 cm | 8 ns power detection resolution |


**Figure 3.** The 2.4 GHz control band held whole: 2.39 to 2.49 GHz in one 100 MHz window at 30 kHz resolution, with max hold (T1) and a single sweep (T2) overlaid in the center pane, a spectrogram above, and a 31.7 MHz zoom below on one video downlink. The gap between the two overlaid traces is what a receiver that visits this band rather than holding it would report as empty, and the spectrogram is what a hopping link looks like before anything has been identified. **MEASURED**

Figure 3 is the band this paper is about. Read the spectrogram rather than the spectrum: the marks are short, they move, and none of them is where it was a moment ago, which is the whole of the detection problem in one picture. Note what the figure cannot prove. A max-hold trace and a spectrogram look much the same behind a fast swept receiver as behind a gap-free one, so R1 is not settled by a screenshot. It is settled by the published frame rate in R2, which is why that row carries a transform size.

## 3 The link budget, and the asymmetry it creates

Free-space path loss sets the floor for both directions:

$$L_{fs} = 32.44 + 20\log_{10}(f_{MHz}) + 20\log_{10}(d_{km}) \qquad (2)$$

where

$f_{MHz}$ is the frequency, in megahertz; 2440 for the control band and 5800 for the video downlink

$d_{km}$ is the path length, in kilometers

At 2.44 GHz that is 100.2 dB at one kilometer and 109.7 dB at three. At 5.8 GHz it is 7.5 dB worse at every range, which is one reason long-range control links migrate down to 900 MHz, where the same geometry costs 8.5 dB less than 2.4 GHz.

The asymmetry does not come from transmit power, which is similar at both ends and capped by the same rules. It comes from geometry. The aircraft path is essentially free space. The controller path leaves a transmitter a meter and a half above the ground, so beyond a crossover of roughly 1.5 kilometers the loss follows a fourth-power law rather than a square law, and urban clutter adds 20 to 35 dB on top. For a sensor budget allowing 127 dB of path loss, the aircraft is detectable to around twenty kilometers in clear air while the controller is detectable to about 1.4 kilometers through urban clutter.

**HONEST LIMITATION**

That fourteen-to-one range ratio reverses more often than counter-UAS marketing admits. A split link with a 900 MHz uplink and a 5.8 GHz downlink hands the uplink a 16 dB path-loss advantage plus better clutter penetration. Photogrammetry and waypoint missions often fly with video off entirely. And an aircraft directly overhead sits in its own antenna null. Detection range is a distribution, not a number.

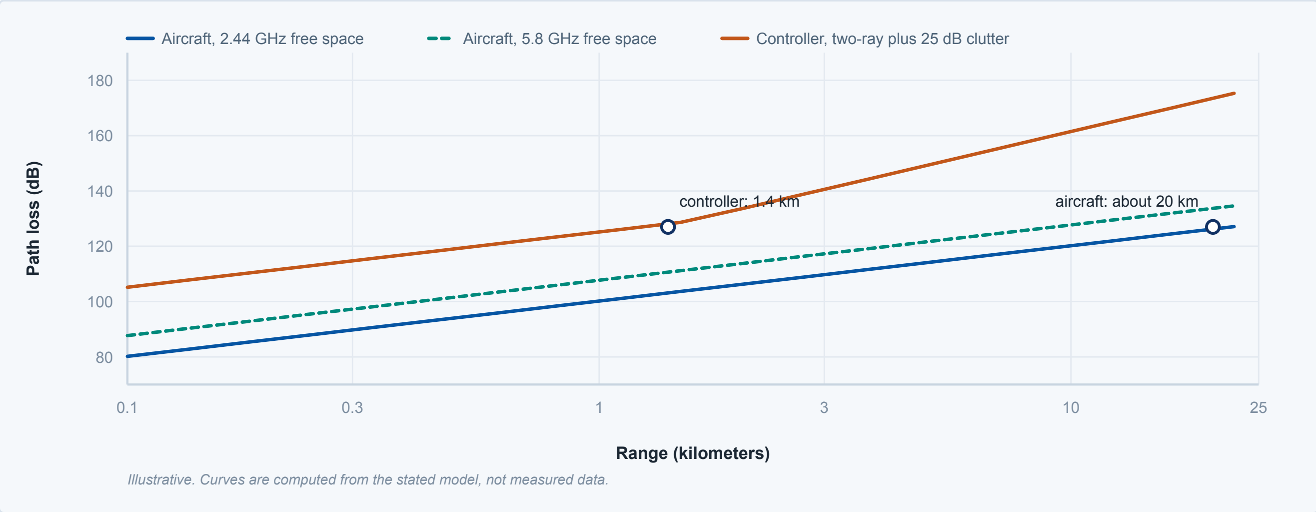


**Figure 4.** At identical transmit power the aircraft is detectable roughly fourteen times further than its controller. The difference is not power. It is ground reflection and clutter standing between the sensor and the person you are trying to find. DERIVED

Figure 4 plots both against a 127 dB budget. The crossing points are the whole planning problem: a sensor sited for the aircraft is not sited for the operator.

## 4 Finding the operator

A bearing is not a position, and received power is not a range. The methods available differ by an order of magnitude in what they deliver, and two of the classical ones are structurally unsuited to hopping links.

**Table 4.** Direction-finding methods and what each actually delivers. Note that pseudo-Doppler needs milliseconds of continuous carrier, which a 250 microsecond hop does not provide, and that interferometry and time-difference methods both require the same burst captured coherently, which a scanning front end cannot supply.

| Method | Achievable accuracy | What it needs | Where it fails |
|---|---|---|---|
| Received signal strength | range to plus or minus 30 to 60 percent | amplitude only | unknown emitter power is indistinguishable from range |
| Amplitude comparison, 4 to 8 sectors | 5 to 15 degrees | one burst, amplitude only | multipath ripple |
| Pseudo-Doppler | 2 to 10 degrees | a continuous carrier for several milliseconds | short bursts and fast hopping, structurally |
| Multi-baseline interferometry | 1 to 3 degrees clear, 5 to 15 urban | coherent, gap-free, multi-channel capture | coherent multipath |
| Time difference of arrival | 10 to 50 m inside the array | the same burst at three or more synchronized sites | emitter outside the array |

For a phase interferometer the geometry sets both the accuracy and the ambiguity. Bearing follows from the phase difference across a baseline, and unambiguous coverage of the hemisphere requires the baseline to be no more than half a wavelength:

$$\sigma_\theta = \lambda\, \sigma_\varphi / (2\pi D \cos\theta), \quad \text{unambiguous if } D \le \lambda/2 \tag{3}$$

where

$\lambda$ is the wavelength, in meters; 0.1229 m at 2.44 GHz

$D$ is the baseline length, in meters; 0.5 m in the worked case

$\sigma_\varphi$ is the phase measurement error, in radians, roughly the reciprocal square root of linear SNR; 0.1 rad at 20 dB

$\theta$ is the angle from broadside, in degrees; the relation degrades as 1/cos θ and is written for boresight

The tension is immediate. A half-meter baseline at 2.44 GHz gives about 0.2 degrees of bearing accuracy at 20 dB signal to noise, and eight ambiguous bearings across the hemisphere, one for every half wavelength of baseline. Resolving them takes a ladder of baselines, the shortest at half a wavelength and the longest to refine, which is why serious direction finding uses arrays rather than pairs.

Whatever produces the bearing, two bearings produce a fix, and the error follows from range and cut angle. Cross-range error is simply range times bearing error: three degrees at two kilometers is 105 meters. Taking circular error probable as 0.5887 times the sum of the two semi-axes, two such bearings crossing at right angles give about 124 meters. The same two bearings crossing at twenty-five degrees stretch the ellipse to 343 by 76 meters and give 247, with no degradation in the sensors at all. Geometry does that, not equipment.

Time-difference methods trade the bearing problem for a timing problem, and the conversion is brutal in its simplicity: one nanosecond is thirty centimeters. The estimator floor for a 645 microsecond burst across 10 MHz at 10 dB signal to noise is about 0.15 nanoseconds, which is four and a half centimeters, and no field system has ever achieved it. Real accuracy is set by inter-site clock synchronization, which for a satellite-disciplined oscillator is 10 to 50 nanoseconds, and by multipath leading-edge bias in urban settings, which is a bias rather than noise and therefore does not average away.


**Figure 5.** Marker signal track holding a drifting carrier at 999.90 MHz in SpectraCore: 1 MHz to 2 GHz span, 300 kHz resolution bandwidth, peak detector. This is the limit of what one channel can prove about one emitter. A bearing needs the same burst on several channels at once, and no single-channel display can show that. **MEASURED**

Figure 5 is the front-end requirement rather than the answer: every method in Table 4 below the first two rows needs the same burst, captured coherently and timestamped.

## 5 The legal boundary, which is also a design boundary

Radio detection of unmanned aircraft is lawful. Demodulating the communications those aircraft carry generally is not, and the distinction is sharper than most system specifications acknowledge.

In United States law the Wiretap Act turns on **contents**, meaning information concerning the substance, purport or meaning of a communication. The existence of an emission, its center frequency, its bandwidth, its timing, its hop pattern, its amplitude, its phase and its bearing are none of those things. A further provision withholds the protection of ready public accessibility from communications sent using modulation whose essential parameters have been withheld from the public, which means a proprietary drone link can be legally protected even though it carries no encryption. Unencrypted is not the same as lawful to decode.

**DESIGN RULE**

A lawful radio counter-UAS sensor detects, characterizes and direction-finds an emission. It does not demodulate the communication that emission carries. The constraint costs less capability than it appears to, because bearing, timing and hop structure locate an operator and payload content does not.

There is one clean exception and it is worth designing around. The mandated Remote ID broadcast is required by regulation to be open, documented and receivable by the general public. Parsing it is not interception in any meaningful sense. Its second use is less obvious and more valuable: correlating Remote ID against the non-cooperative detections lets a system suppress compliant traffic from the alert stream, so the one aircraft that did not announce itself is not buried under the dozens that did.

## 6 Measurement practice

1. **Detect on the downlink, then cue onto the uplink.** The downlink gives early warning cheaply. Finding the operator needs a longer dwell and more antenna gain pointed at a quieter, lower emitter.
2. **Do not set a detection threshold on a max-hold trace**, in SpectraCore or anywhere else. Max-hold destroys the underlying statistics, so a constant-false-alarm threshold computed against it has no characterisable false-alarm probability. It is a common and invisible field error.
3. **Use an order-statistic false-alarm threshold, not a cell average.** The 2.4 GHz reference window is full of interferers, and cell averaging collapses in their presence.
4. **Confirm before declaring.** A per-burst false-classification rate of one in a thousand, confirmed four times in ten opportunities, gives a track false-alarm rate around two in ten billion. Confirmation buys seven orders of magnitude for a few hundred milliseconds of latency.
5. **Survey the site before you trust a bearing.** Every reflector within a few hundred meters is a competing arrival, and interferometric bearing error in dense clutter is five to twenty degrees regardless of the array.

### Field evidence

Two things show up in deployment that the arithmetic above does not predict. The first is siting. Sensors placed for the aircraft sit high and clear, and the same mast is the worst place from which to hear a controller a meter and a half off the ground behind a building line, so installations that detect well and geolocate badly are the common failure rather than the rare one. The second is clutter, and it lands at the top of Table 4's urban range rather than the middle: sites with metal-clad buildings inside two hundred meters sit at the fifteen-degree end, which at two kilometers is the difference between a street and a district. Both are geometry problems, and both are cheaper to fix before the mast goes up than after.

## 7 Where radio detection is blind

This is the section most counter-UAS material omits, and omitting it is why buyers lose confidence in the category.

Table 5 sets the modalities against the threat modes.

**Table 5.** Radio sensing is necessary and not sufficient. It is the only modality that offers protocol-level identification and a path to the operator, and it has a complete blind spot against three operationally significant threat modes.

| Threat mode | RF | Radar | Acoustic | Electro-optical |
|---|---|---|---|---|
| Standard aircraft, link up | best: warning, identity, operator | good | weak | confirmation |
| Fibre-tethered aircraft | **blind** | best | fair | confirmation |
| Pre-programmed, link off | **blind** | best | fair | confirmation |
| Cellular-linked aircraft | **blind** | best | fair | confirmation |
| Locating the operator | **only sensor that can** | no | no | no |
| Payload and intent | no | no | no | **only sensor that can** |

Fibre-tethered aircraft deserve their own sentence. Spools of ten to twenty kilometers are now fielded at scale, and the control and video links run on glass. There is no emission to detect, at any sensitivity, at any bandwidth. Any radio sensor sold as a complete answer is being oversold, and a serious system pairs radio for cue and operator geolocation with radar for track and electro-optics for identification.

### Cutting the false-alarm rate with a decoder you already have

The failure mode that erodes operator trust fastest is not a missed drone, it is an alarm on a helicopter. Cooperative aircraft announce themselves continuously on 1090 MHz, and the same receiver that watches the control bands can decode that broadcast to an aircraft address, callsign, altitude and position. A track correlating with a decoded 1090ES report is cooperative traffic and can be suppressed before an operator sees it; a track with no such report stays a candidate. Watching 1090 MHz costs tuner time, and a second channel if it has to run while the control bands are being watched. No other hardware is needed, because the ICX-FieldHawk presents itself through SoapySDR, the vendor-neutral hardware abstraction layer, so an existing GNU Radio decoder runs against it with the instrument calibration already applied to the samples. The capability is demonstrated rather than specified; see the verification note.

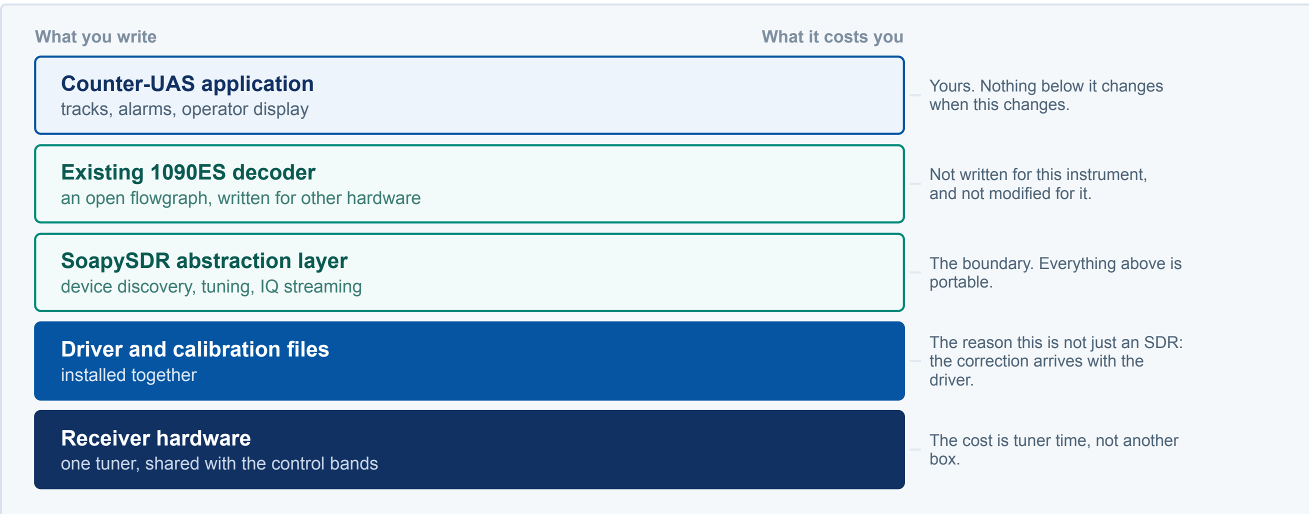


**Figure 6.** Why the suppression argument needs no second instrument. The decoder is an existing open flowgraph that was never written for this instrument, and it attaches at a standard boundary with the amplitude correction already applied beneath it. What a bearing solution consumes from the layer below is not a spectrum but a timestamped, corrected sample stream, which is exactly what section 4 required of a bearing and what a scanning front end cannot deliver. SCHEMATIC

Figure 6 draws the attachment point. The layer worth staring at is the fourth: the calibration files install with the driver, so the flowgraph above is handed a measurement rather than samples, and that is the whole difference between this and a development board.

## 8 Summary, definitions and further reading

Table 6 collects the governing relations and published values with their conditions.

**Table 6.** Governing relations and published values with their conditions.

| Quantity | Relation or value | Condition |
|---|---|---|
| Single-burst intercept | $(d + \tau - 2t_{min}) / T$ | unsynchronized timing |
| Bursts to 90 percent | $\ln(0.1) / \ln(1 - P_1)$ | independent looks |
| Free-space path loss | $32.44 + 20 \log f_{MHz} + 20 \log d_{km}$ | |
| Interferometer ambiguity | unambiguous if $D \le \lambda/2$ | full hemisphere |
| Timing to distance | 1 ns = 0.30 m | c = 299,792,458 m/s |
| Cross-range error | range × bearing error | 3 degrees at 2 km is 105 m |
| Gap-free bandwidth | 100 MHz | 50 MHz standard on the 9.5 GHz module |
| Published POI | 0.512 us at N = 32 | full amplitude accuracy |

Instrument values are from the ICX-FieldHawk datasheets. Link parameters are representative rather than exact and should be confirmed per target.

**Circular error probable**
The radius of the circle, centered on the estimate, containing the true position half the time.

**Remote ID**
A mandated broadcast identifying an unmanned aircraft, its position and its control station, sent in the clear.

**Symbols used in this paper.**

| Symbol | Meaning | Units |
|---|---|---|
| $P_1$ | Single-burst intercept probability | dimensionless |
| $d$, $t_r$ | Dwell, retune time | s |
| $T$ | Revisit period | s |
| $\tau$ | Burst duration | s |
| $D$, $\lambda$ | Baseline, wavelength | m |
| $\sigma_\theta$, $\sigma_\varphi$ | Bearing error, phase error | rad |

## Further reading


- ASTM F3411-22a, *Remote ID and Tracking*, and 14 CFR Part 89 §§89.305 and 89.315.
- Commission Delegated Regulation (EU) 2019/945, Annex Parts 2 to 4 and 6.
- FAA, DOJ, FCC and DHS, *Advisory on the Application of Federal Laws to the Acquisition and Use of Technology to Detect and Mitigate Unmanned Aircraft Systems*, August 2020.
- D. J. Torrieri, "Statistical theory of passive location systems," *IEEE Trans. Aerospace and Electronic Systems*, vol. AES-20, no. 2, 1984.


If you are specifying a sensing layer and want the intercept and link-budget arithmetic run against your own site geometry, our application engineers would be glad to help.

VERIFICATION NOTE

The following values in this paper are not yet confirmed against a published Berkeley Nucleonics datasheet and are marked verify in the text. They must be confirmed before this paper is released.

- Hop rates, per-hop transmit durations and burst timings in Table 1 come from open-source implementations and published reverse engineering rather than vendor datasheets.
- The current published revision of ASTM F3411 should be confirmed.
- Directional antenna coverage and gain should be confirmed against the current ICX-FieldHawk datasheet before quoting.
- The SoapySDR and GNU Radio path behind the 1090ES argument is documented in a platform integration guide and demonstrated, but is not on a datasheet, and the distribution route is unsettled, so no package or driver string is printed here.
- The deployment described in the field evidence section is reported from account history rather than a published source, and no end user, vendor or country is named.

**Berkeley Nucleonics Corp.**
Test, Measurement and Nuclear Instrumentation since 1963
For details: info@berkeleynucleonics.com or 800-234-7858

**BNC-AN-103 Rev A**
August 2026