UAV-borne antenna and airborne radio measurement

Why a flying probe turns an antenna pattern into a timing problem, and what the receiver must guarantee before the geometry means anything.

ABSTRACT

An antenna's range pattern describes the antenna. The installed pattern, measured on the tower with its radome, its neighbors and its structure in place, describes the system. A rotary-wing platform reaches the elevation angles that matter, but it converts a static measurement into a moving one: every amplitude sample must be acquired at its true value before the probe leaves the angular cell it belongs to. This paper derives the far-field standoff, the angular sampling limit and the dwell budget that follows, and shows why a swept receiver cannot meet that budget on a duty-cycled emitter.

WHO THIS PAPER IS FOR

Antenna and site engineers who need an installed pattern without a crane. **Flight test and avionics engineers** characterizing an installed airborne radio. **Spectrum and interference engineers** who must prove where an emitter is and is not radiating.

CONTENTS

1 An antenna on a range is not the antenna you installed

Every antenna arrives with a pattern measured over absorber, on a rotator. It describes the antenna, not the installation. On a tower that aperture acquires a radome, tower legs, three neighbors and a mount fabricated on site, and the pattern that governs coverage is the installed one.

Taking it down costs a crane, an outage and a range booking, and answers a question about an antenna that is no longer installed. The cheaper substitute is a ground receiver at a surveyed bearing, and at its best it returns a genuine horizontal cut. But a downtilted panel, a stacked array and a cosecant-squared radar face all put their designed structure in **elevation**, and a ground receiver sits at one elevation angle fixed by the terrain. It reaches only the wrong cut.

A rotary-wing platform reaches the right cuts in an afternoon, and brings a problem the range never had. The probe is never stationary, so every sample must be acquired before it leaves the cell it belongs to. That constraint, not the geometry, decides whether the pattern is real.

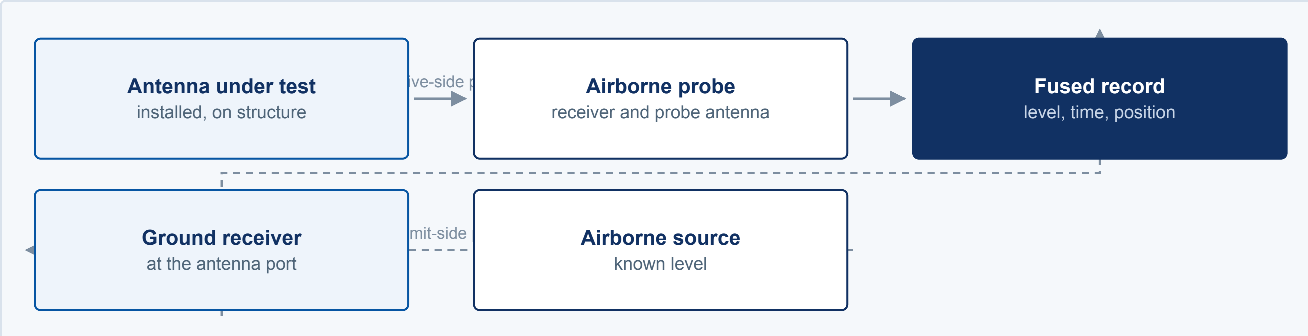

Figure 1. Reciprocity makes the two passes equivalent, so fly whichever the license holder will authorize: one needs the antenna radiating, the other needs it off the air. SCHEMATIC

Figure 1 sets out both topologies. Which one applies decides what flies, and everything after it in this paper follows from that choice.

2 What an in-situ pattern measurement demands

Three requirements fall out of the geometry before any instrument is chosen. The first is standoff: the far field begins where the path-length error across the aperture falls to a sixteenth of a wavelength.

$$R_{\mathrm{ff}} \geq 2D^2 / \lambda \quad (1)$$

where

R_{ff} is the minimum measurement range, in meters

D is the largest dimension of the antenna under test, in meters (5 m for an L-band radar face)

λ is the wavelength, in meters (0.231 m at 1.3 GHz)

One might read that as a boundary. It is a convention. A sixteenth of a wavelength is 22.5 degrees of residual phase error, tolerable in the main beam and not in a deep null: the criterion that costs 0.1 dB on boresight caps null depth near 20 dB. Null work is flown at three to five times Equation 1.

The second is angular sampling. A pattern is the transform of an aperture of finite extent, so it is bandlimited in angle and can be sampled without aliasing at:

$$\Delta\theta \leq \lambda / (2D) \quad (2)$$

where

$\Delta\theta$ is the maximum angular sample spacing, in radians; 0.0231 rad, or 1.32 degrees, for a 5 m face at 1.3 GHz

The beamwidth that spacing has to resolve follows from the same two quantities:

$$\theta_{3\mathrm{dB}} \approx 70^\circ \cdot \lambda / D \quad (3)$$

where

θ_{3dB} is the half-power beamwidth, in degrees; 3.23 degrees for the same face

The two are locked together, which is the useful part. Whatever the aperture and whatever the frequency, Nyquist sampling is a little over two samples per half-power beamwidth, and resolving a null wants four to eight. Once the platform moves, that sample count becomes a sample rate. A 5 m radar face at 1.3 GHz carries the arithmetic from here: far field from 216 m, a 3.2 degree beam, a sample every 1.32 degrees.

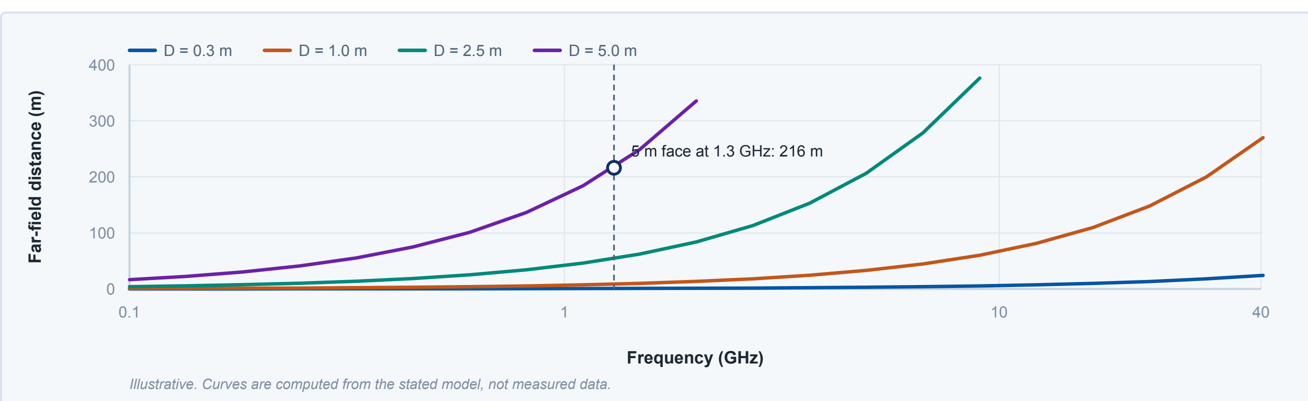

Figure 2. Standoff rises with frequency and with the square of the aperture, so the flyable region is bounded by aperture far more sharply than by band: a meter of aperture stays inside 300 m across the whole band. DERIVED

Figure 2 plots Equation 1 across the band, and the reading to take from it is that aperture, not frequency, is what puts a range out of reach.

3 Why the platform's motion breaks a swept measurement

Standoff and sample spacing convert angle into distance; ground speed converts it into a budget. The probe spends this long in one angular cell:

$$t_{\text{cell}} = R \cdot \Delta\theta \,/\, \upsilon = \lambda R \,/\, (2D\upsilon) \qquad (4)$$

where

R is the slant range flown, in meters (250 m here)

υ is the ground speed along the arc, in meters per second (8 m/s)

t_{cell} is the time available for one amplitude sample, in seconds

At 250 m and 8 m/s that is 250 × 0.0231 rad / 8, the 1.32 degree cell expressed in radians, or 0.72 seconds per sample, and a full arc is 273 samples in 196 seconds of flight. Slowing down buys time and costs endurance in equal measure.

Now fit a swept-tuned receiver inside 0.72 seconds. It drags a resolution filter across the span, and that filter must settle in every cell, giving a sweep time near $k \cdot S / \text{RBW}^2$ with k around 2. A 100 MHz span at 10 kHz resolution is 2.0 seconds: the cell is gone first.

Fixing the tuning helps and does not cure it, because the emitter is duty-cycled. A pulse far shorter than the reciprocal of the filter's impulse bandwidth reads low by roughly 20 $\log_{10}(\tau \cdot 1.5\ \text{RBW})$, 36.5 dB for 1 microsecond at 10 kHz. Widening the filter recovers 20 $\log_{10}$ of the bandwidth ratio while raising the noise floor by only 10 $\log_{10}$ of it, a net gain, which is why a pulsed pattern is measured at the coarse end of the resolution list rather than the fine end.

> **THE FAILURE THAT LOOKS LIKE A MEASUREMENT**
>
> Partial-dwell measurement of a pulsed emitter does not produce obvious garbage. It produces scatter, and scatter under max-hold produces a floor. Every cell reports something, the main beam is broadly right, and the nulls fill in to the depth of the scatter.

So the criterion installed here applies to any instrument from any vendor: **how many samples, each guaranteed correct in amplitude rather than merely detected, does the receiver deliver inside one angular cell?** Above one the pattern is sampled; below one it is interpolated.

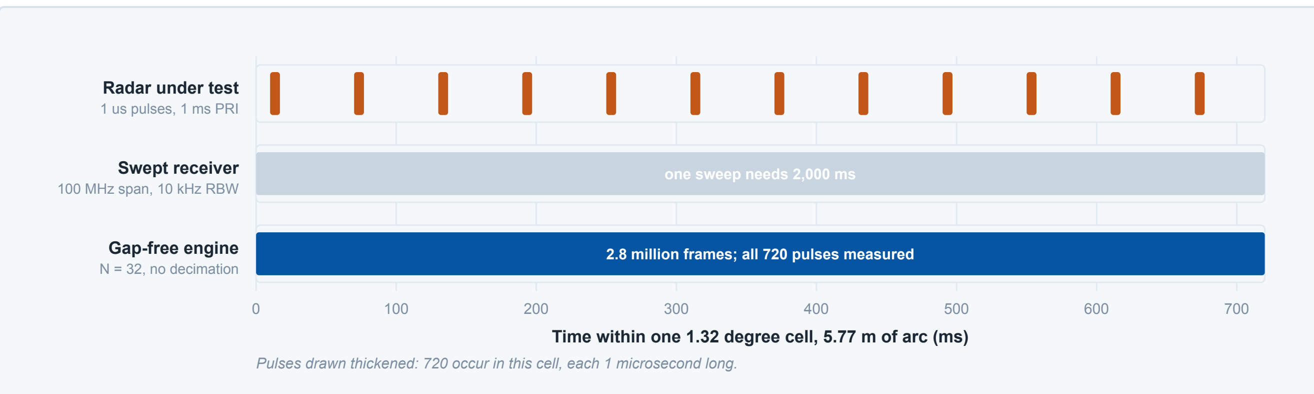

Figure 3. The angular cell, not the sweep, sets the measurement period. Inside one cell a 100 MHz swept acquisition at 10 kHz resolution completes 36 percent of one sweep; a real-time engine completes 2.8 million transforms. DERIVED

Figure 3 draws the cell budget against both architectures, and the gap between the two bars is the whole argument for a real-time engine on a moving platform.

4 Derived requirements

Table 1 is the hinge of this paper. Everything above it is physics; everything below it is instrumentation.

Table 1. Only two rows concern sensitivity or speed in the usual sense. The rest are about time and mass, which is what changes when the receiver leaves the ground.

Requirement	Value for the worked case	Why that number
Slant range	≥ 216 m, flown at 250 m; ≥ 650 m for a 40 dB null	$2D^2/\lambda$, 5 m face; 3× for nulls

Requirement	Value for the worked case	Why that number
Sample spacing	≤ 1.32°	λ/2*D*, two per beamwidth
Sample period	≤ 720 ms	arc per cell at 8 m/s
Amplitude guarantee	≤ 1 μs	the pulse width, not the repetition rate
Log time alignment	≤ 176 ms	0.1 beamwidth of smear, Eq. 3
Input handling	≥ +30 dBm	received level plus margin
Payload	≤ 1.2 kg, ≤ 16 W	a 7 kg-class multirotor

5 Meeting the requirements from a flying payload

A real-time analyzer transforms a continuous sample stream rather than tuning across it, so its amplitude guarantee is a boundary, not a probability. Berkeley Nucleonics publishes the relations for the ICX-FieldHawk family, so it can be computed:

$$\mathrm{POI} = 2 \times N \times M \times 8\ \mathrm{ns}, \quad f_{\mathrm{frame}} = 10^9\ \mathrm{ns} \,/\, (N \times M \times 8\ \mathrm{ns}) \qquad (5)$$

where

N is the transform size, in points (32 to 2048 as published)

M is the decimation factor (1 at full analysis bandwidth)

POI is the duration a signal must be present for its amplitude to be measured, not merely detected

The factor of two is the arrival-phase penalty: a burst inside one window is seen but reads low, because only part of the window held energy. The published points are 0.512 microseconds at a 32-point transform with 3,906,250 frames per second, and 32.768 microseconds at 2048 points with 61,035, at M = 1.

THE WORKED CASE, END TO END

The 1 microsecond pulse at a 1 kHz repetition rate puts 720 pulses in each cell. It is longer than the 0.512 microsecond guarantee at N = 32 and shorter than the 32.768 microsecond guarantee at N = 2048, so transform size is the choice between measuring the pulse and estimating it. At N = 32 the engine produces 3,906,250 × 0.72, or 2.81 million frames in one cell, and each of the 720 pulses lands in a frame that saw it whole. Averaging 720 correct amplitudes drives receiver noise well below the pattern error, which moves the limit on null depth back onto the geometry: at 250 m the standoff convention caps it near 20 dB, and a 40 dB null needs this same cell budget flown at 650 m or beyond, where the cell is 2.6 times longer and easier to meet.

The same path publishes resolution bandwidths from 14.73 MHz down to 3.59 kHz in thirteen Flat-top grades, and the coarse end is what a 1 microsecond pulse needs. Power detection is published at 8 ns with six detectors, turning each pulse into one number.

Figure 4. Power against time is the measurement that produces a pattern point: a pulse train at 1 GHz in a 62.5 MHz analysis bandwidth, a 16 ms record with a zoom on 6 ms of it, tops flat at −10 dBm and bases at −50. The pulses here are milliseconds apart rather than the microseconds of the worked case, so read it for the shape of the measurement rather than for its numbers. One cell yields hundreds of pulse amplitudes, and the statistic of that population is what carries an honest null depth. **MEASURED**

Figure 4 is that path on the instrument: the detector trace, not a swept trace, is where a pulsed amplitude is read.

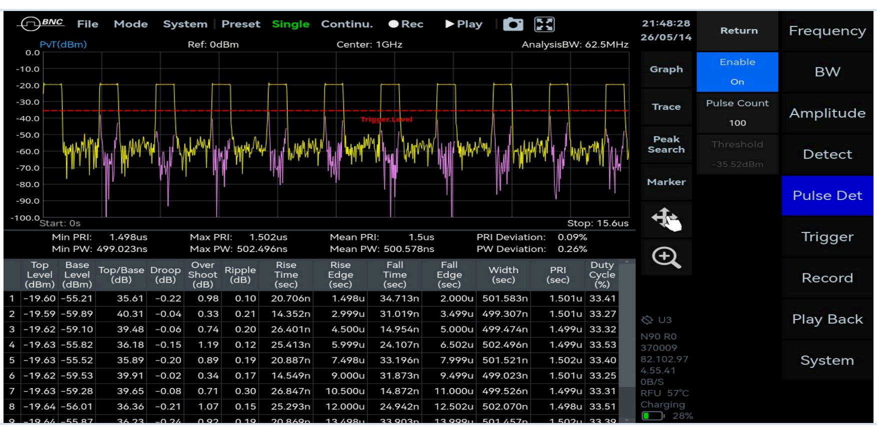

Figure 5. The same detector output turned into per-pulse numbers, here on a bench train of 500 ns pulses at a 1.5 microsecond interval, centered at 1 GHz in a 62.5 MHz analysis bandwidth. A hundred pulses become a hundred rows of top and base level, droop, overshoot, ripple, rise and fall times in nanoseconds, width and interval. The measured width deviation of 0.26 percent across that population is what an amplitude claim on a pattern point rests on: a pulse train that is not stable is not a calibrated illuminator, whatever the analyzer says. Pulse Detection is option 72. **MEASURED**

Figure 5 shows the parameters this produces. The bench train there is not the worked case above, and the point is the population statistics rather than the numbers: before a flight is scored, the illuminator itself should be measured this way, because a drifting pulse width appears in the pattern as a lobe that is not there.

Mass and power decide the project on a multirotor. Table 2 lists the published payload figures.

Table 2. Mass chooses the airframe. A 420 g module with a companion computer, probe antenna and mount lands near 1.2 kg, which a 7 kg-class multirotor carries without a redesign, and at 16 W it draws 6.7 Wh in 25 minutes.

Variant	Frequency	Analysis bandwidth	Mass	Power
ICX-090U, USB	9 kHz to 9.5 GHz	50 MHz std, 100 MHz opt	< 305 g	9 to 16 W
ICX-400U, USB	9 kHz to 40 GHz	100 MHz	< 420 g	9 to 16 W
ICX-400U, LAN	9 kHz to 40 GHz	100 MHz	< 665 g	9 to 16 W
ICX-400 handheld	9 kHz to 40 GHz	100 MHz	1.5 kg	25 W typical

Published on the ICX-FieldHawk USB and handheld datasheets. Antenna and mount mass are not.

Two published details matter more than their datasheet lines suggest. The stack runs on Linux on AArch64 as well as x64, with SCPI and Python standard, so the flight controller can be an ARM companion computer on the airframe, with SpectraCore on the ground for reduction. And option 02 adds a signal generator from 100 kHz to 6.3 GHz on USB models, which flies the transmit-side pass of Figure 1.

THE NUMBER THAT CATCHES PEOPLE OUT

A 25 kW L-band radar with 34 dBi of gain is +108 dBm EIRP. Free-space loss at 250 m and 1.3 GHz is 32.45 + 20 $\log_{10}(1300)$ + 20 $\log_{10}(0.25)$, or 82.7 dB. Into a 2 dBi probe that is 108 − 82.7 + 2 = **+27.3 dBm**, which is 4.3 dB above the published maximum continuous-wave input of +23 dBm at 50 MHz and above with the preamplifier off.

That inversion is the lesson. At 14.73 MHz of real-time resolution the noise floor sits near −88 dBm, from the published displayed average noise level of −159.9 dBm/Hz at 1 GHz, RBW 1 kHz, for the ICX-400 family, or 7.6 dB lower on the 090. A sidelobe 60 dB down is still 55 dB clear. **The first thing a flying receiver needs near a radar is an attenuator, not sensitivity.**

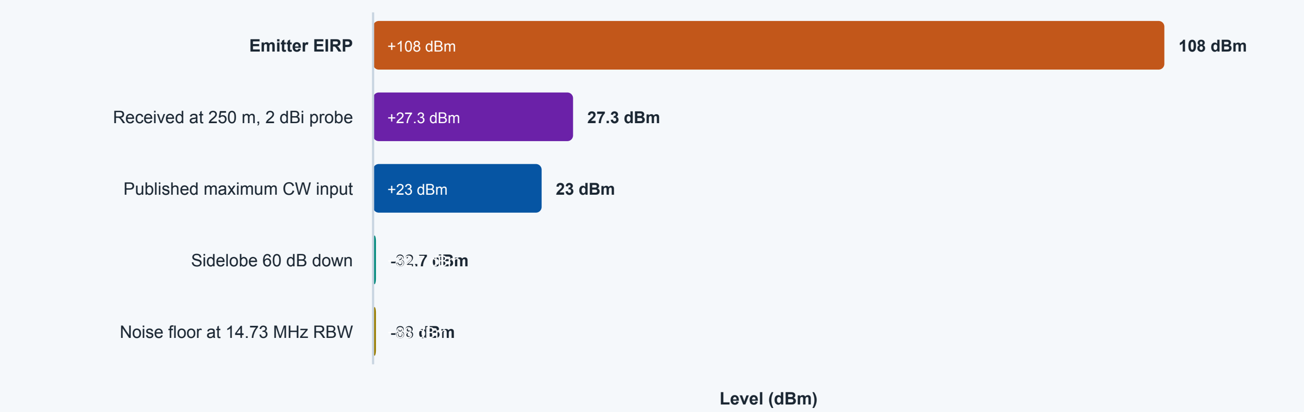

Free-space loss of 82.7 dB at 1.3 GHz over 250 m, less 2 dBi of probe gain, separates the first two bars.

Figure 6. The dynamic-range problem on this flight is at the top of the scale, not the bottom. The received main beam sits 4.3 dB above the published input maximum while the sidelobe of interest sits 55 dB above the noise floor, so a pad costs nothing that matters and its absence costs the measurement. **DERIVED**

Figure 6 puts the five levels on one axis. Read the gap between the second and third bars first: everything else on the flight is comfortable, and that one is not.

The honest limits belong here. Amplitude accuracy is published at ±2.0 dB from 9 kHz to 9.5 GHz and ±3.0 dB from 9.5 to 40 GHz, at 25 °C after 10 minutes of warm-up with spur reject on, and a pattern is a ratio to boresight, so the fixed part cancels. The probe antenna's own pattern, airframe scattering and the ground-reflected path do not cancel, and they usually dominate. A UAV pattern is not an anechoic chamber pattern, and it carries no phase.

6 Measuring the aircraft's own radio

Turn the geometry around and the same payload characterizes the airborne radio. A command link, a video downlink or a 1090 MHz transponder all have patterns shaped by the airframe, and the belly blade that looked omnidirectional on the bench is shadowed by a battery tray and a motor arm. Those emissions are bursts, and a 120 microsecond squitter is exactly what a partial dwell mishandles.

Run that case through the same arithmetic and it does not come out like the radar. At 120 microseconds the squitter is 3.7 times the 32.768 microsecond guarantee at N = 2048, so the amplitude of every message received is correct without any care being taken. The constraint moves somewhere else entirely. A transponder emits roughly two extended squitters a second, so the 0.72 second cell derived in section 3 contains about 1.4 of them, against the 720 pulses the radar put in the same cell. That clears the criterion of section 3, but only just, and one correct sample per cell leaves nothing to average: the difference between the radar and the transponder is not whether the pattern is sampled but how much of the scatter survives into it. **For a transponder the pattern point is set by message rate, not by the receiver**, and the honest options are to fly slower, widen the cell and accept the angular smear, or fly the orbit twice and interleave. No instrument specification changes that number.

Here the fusion of two logs, not the precision of either, sets the accuracy. Any unmodelled offset between the receiver's time base and the flight log smears it:

$$\delta\theta = v \cdot \Delta t \,/\, R \qquad (6)$$

where

$\delta\theta$ is the angular smear from log misalignment, in radians

Δt is the offset between the two logs, in seconds

Holding that smear to a tenth of the 3.23-degree beamwidth of Equation 3, at 8 m/s and 250 m, allows 0.00564 × 250 / 8, or **176 milliseconds** of skew. Published GNSS one-pulse-per-second timing is ±100 ns as standard: roughly 1.8 million times better than the requirement. **The instrument's clock is not the risk. The risk is an operating-system timestamp drifting against an autopilot log nobody disciplined.**

The installation is part of the measurement. Module, companion computer and battery all radiate and all scatter, so their positions relative to the probe must be fixed and recorded before the first sortie: every later correction is applied against that geometry.

7 Measurement practice

None of this depends on which analyzer is flown, and all of it decides whether the sortie was worth the battery.

1. **Fit the pad before the first flight, not after the first failure.** Size it from EIRP and slant range, add 10 dB for the beam you did not expect, and store the attenuator's measured loss as an amplitude offset.
2. **Choose the probe antenna for pattern flatness, not gain**, because probe ripple multiplies the error and does not average out. Apply its pattern as an antenna factor.
3. **Fly constant slant range, not constant altitude**, so amplitude change is pattern rather than geometry.
4. **Discipline both logs to the same one-pulse-per-second edge**, and fly one abrupt heading change at the start of each pass as a marker visible in both.
5. **Warm up on the ground**, since published specifications assume 10 minutes at 25 °C.

When not to do this: if the question is only whether an antenna points where it was aimed, and the emitter is a continuous carrier, a power meter and a surveyed position answer it for less.

Field evidence

This configuration flies commercially. Services companies mount the module on a UAV for airborne radio measurement across the full range and publish the resulting patterns, and their stated motivation is simpler than any argument in this paper: engineers should not have to climb high antenna structures to take a reading. Everything derived above is what it costs to make that substitution honest, and the standoff arithmetic of section 2 is the part that most often decides whether a given face can be flown at all.

8 Summary and selection

Table 3 lifts the six requirements whole into an evaluation matrix, so a vendor's published data can be scored against them line by line without any of it being taken on trust.

Table 3. Lift this whole. Every number here carries the condition it was measured under, and a competing figure quoted without its condition is not comparable to any of them.

Quantity	Relation or value	Condition
Far-field standoff	$2D^2/\lambda$	3 to 5× for null work

Quantity	Relation or value	Condition
Sample spacing, period	$\lambda/(2D)$, $\lambda R/(2Dv)$	constant-radius arc
100 percent POI	$2 \times N \times M \times 8$ ns	0.512 μs at $N = 32$; 32.768 μs at 2048
DANL	−159.9 dBm/Hz (400); −167.5 (090)	1 GHz, RBW 1 kHz, 25 °C, 10 min warm-up
Amplitude accuracy	±2.0 dB; ±3.0 dB	to 9.5 GHz; 9.5 to 40 GHz
Max CW input	+23 dBm	50 MHz and above, preamp off
GNSS 1PPS	±100 ns	standard

Table 4 maps the job to the configuration. Frequency decides the module and the payload budget decides the mount, in that order, because a mount that will not fly makes the frequency question academic.

Table 4. Frequency picks the module and the emitter picks the pad, in that order. Every measurement named in the rows is in the published SpectraCore set, so the configuration question is which probe and how much attenuation, not which license.

If the job is	Fly	With
Anything below 9.5 GHz, including radar faces	ICX-090U module	option 34 omni probe; a pad sized to the emitter
Links, radars and payloads above 9.5 GHz	ICX-400U module	option 35 directional probe; a pad sized to the emitter
Transmit-side pass, or cold sorties	either module	option 02 source, USB only; temperature class 40 or 41

Pulse Detection is option 72. Confirm any configuration against the current datasheet.

9 Definitions, symbols and further reading

Antenna factor
The ratio of incident field strength to the voltage delivered to the receiver, in dB/m.

Far field
The region beyond $2D^2/\lambda$ where the field's angular distribution no longer changes with distance.

Installed pattern
The pattern of an antenna in its mounting, with radome, structure and neighbors in place.

Probability of intercept
The minimum duration a signal must be present for a correct amplitude measurement.

Symbols used in this paper.

Symbol	Meaning	Units
D	Largest dimension of the antenna under test	m
λ	Free-space wavelength	m
R	Slant range to the probe	m
$\Delta\theta$, $\delta\theta$	Sample spacing, angular smear	rad
θ_{3dB}	Half-power beamwidth	deg
v	Ground speed along the arc	m/s
t_{cell}, Δt, τ	Cell dwell, log offset, pulse	s
N, M	Transform size, decimation	points, dimensionless

DANL Displayed average noise level
EIRP Equivalent isotropically radiated power
GNSS Global navigation satellite system
POI Probability of intercept
PRI Pulse repetition interval
RBW Resolution bandwidth
UAV Unmanned aerial vehicle

Further reading

- IEEE Std 149, *Recommended Practice for Antenna Measurements*, for far-field criteria and range error budgets; IEEE Std 1720 for near-field alternatives, which is the better answer for a large aperture that cannot be flown at three to five times its far-field distance.

- C. A. Balanis, *Antenna Theory: Analysis and Design*, Wiley, chapters 2 and 15.
- RTCA DO-160, section 21, for airborne emission limits; ITU-R SM.378 for field strength.
- Berkeley Nucleonics BNC-AN-102, BNC-AN-103 and the ICX-FieldHawk datasheets at berkeleynucleonics.com.

If you are planning a flight program and want to work this arithmetic against your own antenna and airframe, 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.

- Revision suffixes for the standards cited should be confirmed with their issuing bodies.
- Peak pulse input handling is not published; +23 dBm is the maximum continuous-wave input at 50 MHz and above, preamplifier off.
- Mass of the option 34 and 35 antennas is not published, and the rugged tablet's ingress rating carries an unverified specification row. The radar parameters used above are representative.
- 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.
- The two-per-second extended squitter rate used in section 6 is a representative 1090ES airborne-position rate rather than a specification, and should be confirmed against the installation it is applied to.

Berkeley Nucleonics Corp. | **BNC-AN-106 Rev A**

Test, Measurement and Nuclear Instrumentation since 1963 | August 2026

For details: info@berkeleynucleonics.com or 800-234-7858