ICX-FIELDHAWK TECHNICAL PAPER

BNC-AN-110 REV A

AUGUST 2026

Designing a spectrum analyzer into your own product

What an embedded receiver has to deliver, what it cannot fix, and the volume at which building one becomes the cheaper answer.

ABSTRACT

Monitoring, direction-finding and test products are commonly built around a receiver their maker did not design, and that decision is usually argued on radio-frequency performance. This paper derives the error budget of a distributed radio-location node and shows that the receiver's own time-delay estimation contributes under one hundredth of one percent of the position-error variance, while time transfer, inter-channel matching, enclosure temperature and site geometry contribute all of it. It then works the ten-year cost arithmetic for building a front end against buying one, gives the volume at which that answer reverses, and names the terms that decide whether an embedded module is an asset or a schedule risk.

WHO THIS PAPER IS FOR

System integrators and product architects embedding a receiver inside a branded product. **Program and supply-chain managers** who own the ten-year cost. **RF and mechanical engineers** sharing an enclosure, a thermal budget and a calibration interval.

CONTENTS

1 The problem an integrator actually has

A company that sells spectrum monitoring systems does not want to be in the receiver business. It wants to sell nodes, a detection engine, a database and a map, and underneath all of that it needs a calibrated receiver. Design one, buy one as a module, or buy an uncalibrated board and finish it yourself.

The instinct is to compare receivers, so datasheets go side by side and the argument is about noise floor and phase noise. Three years later, when the program is in trouble, the cause is almost never noise floor. It is a supplier who changed a front-end filter without calling the change calibration-affecting, so a hundred fielded nodes now disagree with the fifty built last quarter, and nobody kept as-found data to say by how much.

Take the escape route seriously, because most integrators try it. Buy a low-cost wideband board, characterize it against a reference source, correct it in software and keep the margin. It works on the bench and fails at the second unit. The radio design is not the expensive part; the calibration fixture, the traceability chain, the qualification and ten years of obsolescence engineering are, and none of those scale to fifteen units a year.

2 Where the boundary falls

An embedded receive module is a supply boundary before it is a component. What matters is what crosses the line: calibrated baseband samples, an absolute timestamp with a stated uncertainty, a calibration table with its traceability statement, a versioned interface, and a device identity carrying no vendor branding. Read Figure 1 as a list of deliverables. The timestamp uncertainty, the traceability statement and the interface deprecation period are the three most often left unspecified.

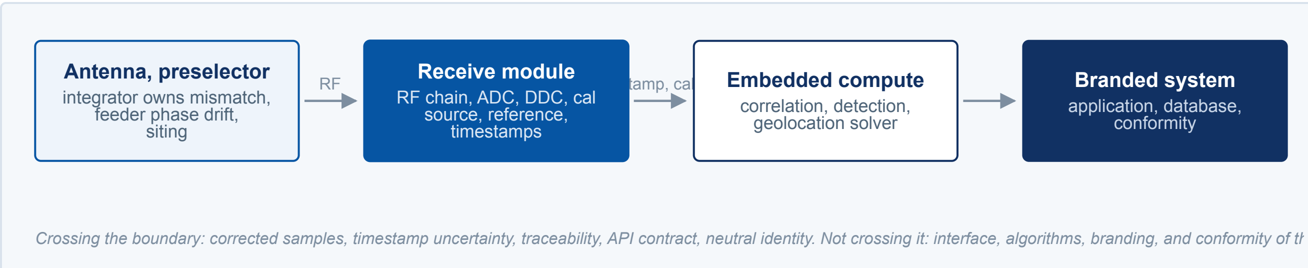

Figure 1. The module ends at calibrated, timestamped baseband. Every error term larger than the module's own, and every element a customer will recognize, falls on the integrator's side of the line. SCHEMATIC

3 What sets accuracy in a distributed system

Two measurements dominate these products: locating an emitter from the difference in arrival time between nodes, and taking a bearing on it from one node. Neither is limited by a receiver specification.

Time difference of arrival: a nanosecond is thirty centimeters

The conversion from timing to distance is the whole physics of the method:

$$\Delta r = c \cdot \Delta\tau \tag{1}$$

where

Δr is range difference between two nodes, in meters
c is the speed of light, so one nanosecond is 0.2998 m
$\Delta\tau$ is measured arrival-time difference, in seconds

One nanosecond is thirty centimeters. That conversion is worth more than any comparison of noise floors, because it says the answer is set by whatever assigns time to a sample. Geometry then multiplies the error:

$$\sigma_p = G \cdot c \cdot \sigma_\tau, \quad G = 1 / (2 \sin(\theta/2)) \tag{2}$$

where

σ_p is 1σ position error, in meters
G is geometric dilution of precision (2 to 4 for a workable layout)
σ_τ is 1σ arrival-time uncertainty, in seconds
θ is angle the baseline subtends at the emitter

Note what G is made of. It contains no receiver parameter at all. It is a site survey.

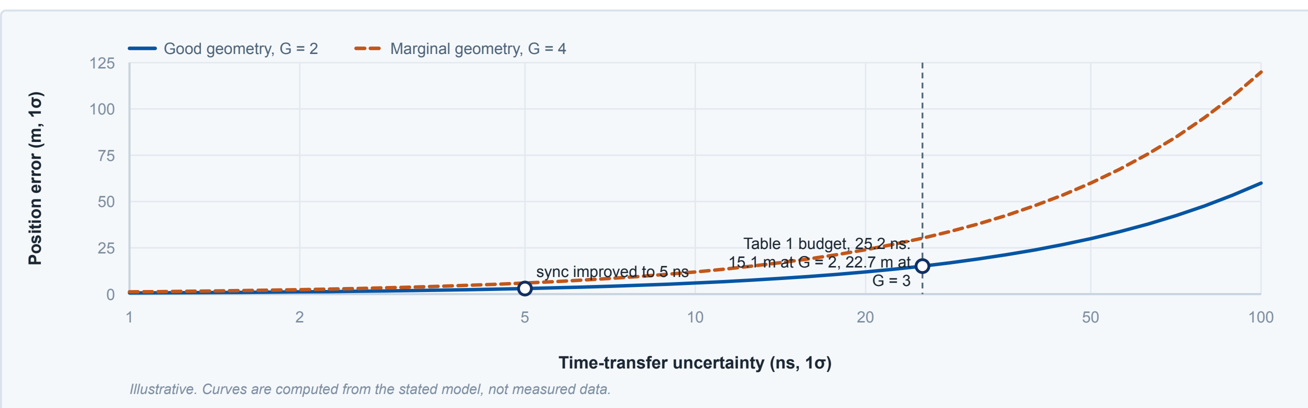

Figure 2. Position error is proportional to time-transfer uncertainty and to geometry, and independent of everything on a receiver datasheet. Moving one node to bring the dilution of precision from 4 to 2 is worth as much as halving the timing error, and it is usually cheaper. DERIVED

Figure 2 plots both effects on one axis. The logarithmic time axis is what bends the lines; the underlying relations are straight.

Direction finding: a decibel is three and a third

degrees

Amplitude-comparison direction finding takes the arctangent of the ratio of two channel outputs, so any gain error between them rotates the answer. For a source at 45 degrees, where the channels should read equal, a gain mismatch ΔG gives:

$$\varepsilon_\theta = \arctan(10^{\Delta G/20}) - 45^\circ \tag{3}$$

where

ε_θ is bearing error, in degrees

ΔG is inter-channel gain mismatch, in dB (0.3 dB is good practice, 1.0 dB is common in an uncontrolled set)

Work it for one decibel. $10^{1/20}$ is 1.1220, whose arctangent is 48.29 degrees, so the error is **3.29 degrees**, larger than most published direction-finding accuracy figures and produced by a mismatch smaller than most published amplitude accuracy figures.

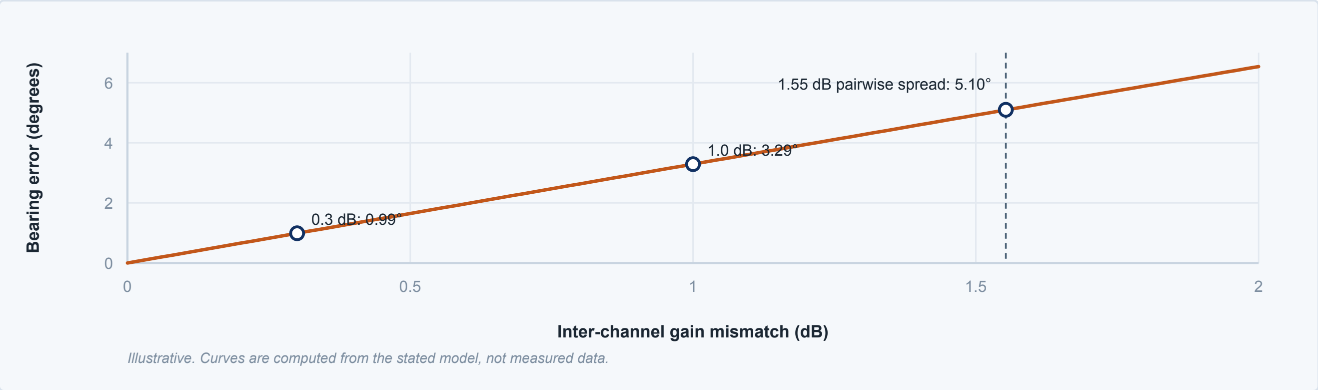

Figure 3. Matching is a pairwise property of a set of units, not an absolute property of each one. Two units individually certified to ±1 dB may differ by 1.55 dB at the 97.5th percentile, which is 5.1 degrees of bearing error, and no software recovers it. DERIVED

Figure 3 carries the commercial point. For two tolerances uniform on ±*A*, the difference has a 97.5th percentile of 1.553*A*, so a stack of perfect absolute certificates does not make a matched set.

4 Why the specification page is the wrong comparison

Arrival-time uncertainty at a node is the root sum of squares of independent terms:

$$\sigma_\tau = \sqrt{\sigma_{est}^2 + \sigma_{sync}^2 + \sigma_{cal}^2 + \sigma_q^2} \tag{4}$$

where

σ_{est} is correlator estimation error, about 0.8 / (B √(SNR · BT)), with SNR as a power ratio and T the integration time: 0.23 ns at B = 5 MHz, SNR = 10 and T = 10 ms

σ_{sync} is time-transfer error between nodes

σ_{cal} is residual antenna and feeder delay error

σ_q is timestamp quantization, all in seconds

Table 1 evaluates it for one realistic node, and the variance column is where the argument of this paper is won or lost.

Table 1. The synchronization term is 98.5 percent of the variance, and it is set by the GNSS option on the order form rather than by anything inside the receiver.

Term	1σ	Variance (ns²)	Share
Correlator, 5 MHz signal, 10 dB SNR, 10 ms integration [illustrative]	0.23 ns	0.05	0.008 %
Time transfer, from a ±50 ns 1PPS bound taken as 2σ	25 ns	625.0	98.5 %
Antenna and feeder delay residual [illustrative]	2 ns	4.00	0.6 %
Timestamp quantization at the 8 ns engine clock	2.31 ns	5.33	0.8 %
Root sum of squares	**25.19 ns**	**634.38**	**100 %**
Range difference, Eq. 1	7.55 m		
Position error at *G* = 3, Eq. 2	22.7 m		

The 1PPS figure is published, quoted with its option; the other terms are illustrative.

Check it. 0.05 plus 625.0 plus 4.00 plus 5.33 is 634.38 ns², whose square root is 25.19 ns; times *c* that is 7.55 m, and at a dilution of precision of 3 it is 22.7 m. Make the correlator ten times better and the first term falls to 0.0005 ns², the total to 634.33 ns², and the position answer improves by 0.9 mm.

THE LESSON WORTH REPEATING TO A COLLEAGUE

In a working time-difference geolocation network the receiver's own timing precision accounts for roughly one hundredth of one percent of the position-error variance. You are not buying a receiver. You are buying a timestamp, and the rest is packaging.

That is the buying criterion, and it applies to any supplier. Score a candidate on the uncertainty of the timestamp it hands you, the pairwise matching of a set of units, the temperature coefficient of amplitude, and the terms under which those may change. Sensitivity still has to be adequate. It is not the discriminator.

5 Derived requirements, and an embedded module against them

The ICX-FieldHawk family includes two variants for this role: the ICX-090U, 9 kHz to 9.5 GHz with 50 MHz of analysis bandwidth standard and 100 MHz optional, and the ICX-400U, 9 kHz to 40 GHz with 100 MHz. Both are USB or LAN devices with no display, and published values hold after 10 minutes of warm-up, at 25 °C ambient, with spur reject standard on. Table 2 sets them against requirements derived for a 25 m position and 2 degree bearing specification. It is the hinge: before it is physics, after it is instrumentation and contract.

Table 2. Six of the seven requirements are properties of timing, matching or thermal behavior. Note the third row: amplitude accuracy is absolute and says nothing about how two units agree, so pairwise matching is a separate deliverable that has to be asked for.

Requirement	Derived value	Published module value
Node time-transfer uncertainty	≤ 28 ns, 1σ, from Eq. 2 at G = 3	GNSS 1PPS ±100 ns standard; ±75 or ±50 ns with options
Timestamp granularity	≤ 10 ns, under the sync term	8 ns time resolution
Pairwise gain match across a set	≤ 0.6 dB, giving 1.97° from Eq. 3	Not specified; amplitude accuracy is ±2.0 dB to 9.5 GHz, ±3.0 dB above, and absolute
Reference stability through an outage	≤ 168 ns/hour of holdover drift, to hold 28 ns through a 10 minute outage, from Eq. 2	TCXO < 1 ppm; OCXO option 01 < 0.15 ppm
Continuous IQ to host	≥ 20 MHz sustained	25 MHz continuous; 100 MHz burst into 128 Mbyte
Dissipation and envelope	≤ 20 W	9 to 16 W; < 305 g (090U USB), < 420 g (400U USB)
Operating temperature	site ambient plus internal rise	0 to 50 °C; −20 to +65 °C option 40; −40 to +65 °C option 41

Published values from the ICX-FieldHawk USB and LAN module datasheet, conditions as stated above.

Read the first row again, because it is the one the paper turns on. The standard 1PPS does not meet the derived requirement. Taken as a 2σ figure, ±100 ns is a 50 ns synchronization term, which puts the node at 45.1 meters rather than 22.7; the ±75 ns option reaches 33.8 meters, and only the ±50 ns option reaches the 22.7 meters Table 1 reports. That option is the line most often left off a quotation, and it is the only one that meets the specification.

The fourth row cannot be scored from a datasheet at all. TCXO and OCXO figures are frequency accuracy, not holdover drift, and a part specified at 0.15 ppm reaches the 28 ns budget in 0.19 seconds of free run if that accuracy is read as drift. Ask for drift after the last correction, over the outage you intend to survive, and treat any answer quoted in ppm as an answer to a different question.

The fifth row is the one an integrator underestimates: burst depth and sustained rate read alike on a datasheet and behave nothing alike in a product.

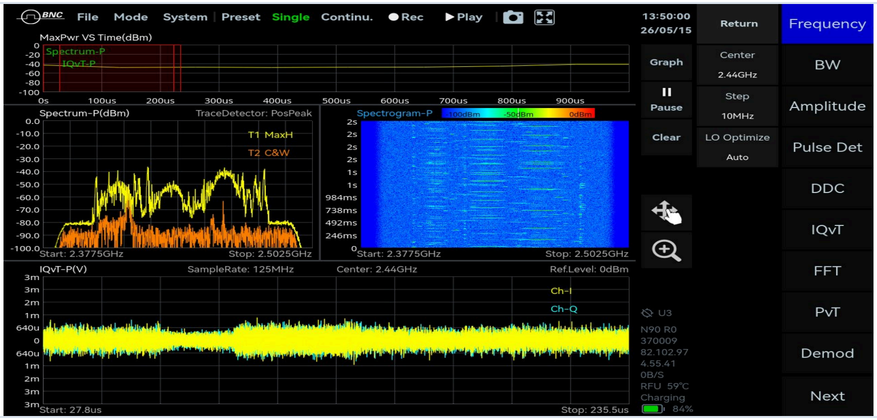

Figure 4. One acquisition at 2.44 GHz, 125 MHz of span at 125 MSPS, presented four ways at once: power against time, max-hold and clear-write spectra, a two-second spectrogram, and the I and Q traces across 208 microseconds. The span reported equals the sample rate while the usable filtered bandwidth is the published 100 MHz, and this is a single acquisition, so it is the fifth row's burst case and not its sustained one. That distinction is what an embedded design carries into its own throughput budget rather than inheriting from a display. MEASURED

Figure 4 is what the fifth row buys: every panel is a projection of one array of complex samples, so a host that cannot sustain the stream does not lose a panel, it loses the record.

The host software, SpectraCore, is a full desktop application, which for an embedded product is the wrong deliverable. What matters is that it can be bypassed: the modules are driven from C, C++, C#, Python, MATLAB, Qt and LabVIEW over SCPI, on Windows and Linux, x64 and AArch64. Bring a set up under SpectraCore on a bench, then ship an ARM node with none of it.

The thermal term, decided by somebody who was not at the RF review

A sealed enclosure reaches a steady rise above ambient set by its thermal resistance, on a time constant set by its mass:

$$\Delta T_\infty = P\,R_{\mathrm{th}}, \quad \tau_{\mathrm{th}} = m\,c_{\mathrm{p}}\,R_{\mathrm{th}}, \quad \Delta A = k_{\mathrm{T}}\,\Delta T \qquad (5)$$

where

P is dissipation, in watts (15 W for a module and host board)

R_{th} is enclosure thermal resistance, in K/W (1.67 for a sealed 0.055 m² box)

$m\,c_p$ is thermal mass, in J/K (1.2 kg of aluminium)

k_T is amplitude coefficient, in dB/K (0.02 typical)

Evaluate it. 15 W times 1.67 K/W is 25.1 K, and 1.2 times 900 times 1.67 is 1,804 s, so the enclosure is within 5 percent of final temperature after 90 minutes and not before, and at 0.02 dB/K that rise is 0.50 dB of amplitude walk. At 45 °C ambient it also puts internal air at 70 °C, outside even the −40 to +65 °C class of option 41. The fix is area, a conduction path or a lighter finish, and it is not a receiver purchase.

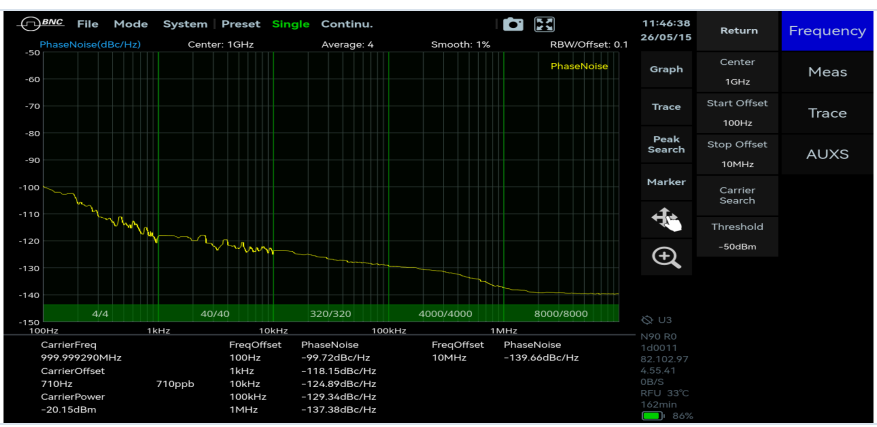

Figure 5. The measurement that characterizes a reference, and the reason it is not the answer here. Phase noise on a 1 GHz source from 100 Hz to 10 MHz offset: −99.7 dBc/Hz at 100 Hz, −124.9 at 10 kHz, −137.4 at 1 MHz, with the carrier itself 710 Hz off nominal, or 710 parts per billion. Parts of this trace sit below the analyzer's own published floor of −107.5 dBc/Hz at 1 GHz and 10 kHz offset, so they bound the instrument's margin rather than the source, exactly as BNC-AN-104 sets out. None of it reaches the fourth row of Table 2 in any case: the measurement starts at a 100 Hz offset, holdover across a ten minute outage lives near 1 mHz, five decades lower, and aging and temperature coefficient do not appear here at all. This is the wrong domain for the question, however good the numbers are. MEASURED

Figure 5 is worth reading against Table 1. The frequency domain is where a reference looks good and the time domain is where it has to perform, and the variance column shows which of the two the position answer depends on.

6 Build, buy, and the volume where the answer reverses

The crossover is where a fixed cost spread over a lifetime volume meets a flat unit price:

$$N^* = (C_{\text{build}} - C_{\text{buy}}) \,/\, [\, L\,(c_{\text{buy}} - c_{\text{build}})\,] \tag{6}$$

where

N^* is annual volume at which the paths cost the same, in units/year

C_{build}, C_{buy} is the fixed cost of each path: non-recurring plus ten-year sustaining for building, integration only for buying

L is program life, in years (10 for a fielded product)

c_{buy}, c_{build} is marginal cost per unit

Table 3. At fifteen units a year the built front end costs about $33,000 per unit against $9,000 bought, and arrives two years later. Eq. 6 puts the crossover at $4,125,000 over ten times $3,500: 118 units a year.

Line item [all illustrative]	Build	Buy
Design and layout: RF, digital, firmware, mechanical, three spins	$1,610,000	n/a
Calibration fixture development and infrastructure	$500,000	n/a
Qualification, documentation, first article	$265,000	n/a
Integration: enclosure, thermal path, host software, product qualification	included above	$450,000
Sustaining, metrology and obsolescence, 10 years	$2,200,000	included
Fixed total	**$4,575,000**	**$450,000**
Marginal cost per unit; time to first delivery	$2,500; 24 to 30 months	$6,000; lead time
Per unit at 15 units/year	**$33,000**	**$9,000**
Per unit at 500 units/year	**$3,415**	**$6,090**

Illustrative structure to substitute into, not a quotation.

Table 3 is the honest version, and the Buy column is still generous to itself. An integrator pays for the enclosure, the thermal path, the host software and the qualification of the finished product whichever way the decision goes; put $450,000 against those and the crossover falls from 131 units a year to 118. Buying still wins at fifteen units by nearly four to one and still loses at five hundred.

Lifecycle risk belongs in the same arithmetic. For M parts with no second source and end-of-life notices at rate h, expected events over a life t are $M(1 - e^{-ht})$: at 25 parts and 0.08 per year over ten years, 13.8 events, with a two-in-a-billion chance of none. Obsolescence is a workload, and the contract decides whose engineers absorb it.

Terms that make a module a component rather than a dependency

- **Two-tier change control.** A short notice period for changes touching nothing calibrated, a longer one for anything touching amplitude, phase, timing or the interface, and the classification fixed in the agreement rather than at the time of the change.
- **Continuity and calibration.** An end-of-life notice period, a last-time-buy window, a named migration target, and a serialized calibration certificate with its traceability chain, plus a pairwise matching record where bearings matter.
- **Interface stability and neutral identity.** A versioning policy, a deprecation notice period and firmware pinning. Neutral enclosure, model number and documentation, and also device descriptors, service names, server headers and log strings, any of which will show a supplier name to a customer with a packet capture.

7 Measurement practice

1. **Measure your own arrival-time uncertainty.** Feed two nodes from one source through a known splitter and cable set, log for an hour, and take the standard deviation of the difference. That is what goes into Eq. 4, not a datasheet value.
2. **Compute the dilution of precision before siting.** Evaluate Eq. 2 over the coverage area for each

candidate layout. Moving one node a few hundred meters routinely buys more than a tenfold receiver improvement.

3. **Wait out the thermal transient, and watch the feeder.** Amplitude means nothing until three time constants after power-on, roughly 90 minutes. Three meters of ordinary PTFE coax at 1 GHz, at 50 ppm/K and a velocity factor of 0.7, drifts 0.26 degrees per kelvin, so two nodes twenty kelvin apart are five degrees of phase apart before the receiver does anything. Phase-stable assembly at 10 ppm/K takes that to 0.05 degrees per kelvin, and it is the cheapest line in the node.
4. **Test bearing accuracy the way ITU-R SM.2060-0 prescribes.** A full azimuth traverse in 5 or 10 degree steps, across the band, reporting RMS error over all azimuths with the signal-to-noise ratio stated. Plot error against azimuth too: gain mismatch peaks at 45 degrees and nulls at 0 and 90, phase mismatch grows off broadside, and the shape says which defect you have.
5. **Keep as-found calibration data from the first unit.** A 200-unit fleet recalibrated annually at three weeks turnaround holds 11.5 units in transit and costs about $226,000 a year at $1,130 per unit [illustrative]. A 24-month interval halves both, and the only thing in the way is drift data nobody kept.

Field evidence

The pattern here is long-running rather than a single case. Several independent system houses have built their own monitoring and direction-finding products on receive modules from this family, write their own software and sell the result under their own name, and none is in the receiver business. Each made the section 6 judgement the same way, at volumes well below the crossover, where the arithmetic says buy and the instinct says build.

8 Where an embedded module is the wrong answer

- **Below about one nanosecond of time transfer, no receiver helps.** Meter-class geolocation needs length-calibrated fibre, GNSS common view with long averaging, or a physical timing network, and that infrastructure becomes the product.
- **For narrowband emitters, time difference of arrival is the wrong method.** A 12.5 kHz channel at 10 dB signal-to-noise ratio with a full second of integration reaches 181 ns of estimation error, 54 m before geometry or synchronization. Bandwidth decides the method; receiver quality does not.
- **A monitoring receiver is not a compliance measuring receiver.** Conformity measurement needs the CISPR detectors and bandwidths, a defined input impedance and overload factor, and an accredited uncertainty budget, and no CISPR function appears in the module specification. Nor is the module the limiting element for field strength: a ±0.3 dB receiver inside a ±2.5 dB antenna factor gives ±2.5 dB.
- **Above roughly 120 units a year, build.** Table 3 says so plainly. The argument for buying is an argument about this volume, not a universal claim.

9 Summary, selection and definitions

Table 4 collects the governing relations and published values with their conditions.

Table 4. Governing relations and published values with their conditions, in a form that lifts into an evaluation matrix.

Quantity	Relation or value	Condition
Range from time	$\Delta r = c\Delta\tau$; 1 ns = 0.2998 m	vacuum
Position error	$G\,c\sigma_\tau$	$G = 1/(2\sin(\theta/2))$, two nodes
Bearing error from mismatch	$\arctan(10^{\Delta G/20}) - 45°$; pairwise spread $1.553A$	source at 45°; two uniform $\pm A$
Build/buy crossover	118 units/year	illustrative Table 3 values, including buy-side integration cost
Module GNSS 1PPS; bandwidth; power	±100 ns std, ±75 or ±50 ns optional, and only the ±50 ns option meets the 25 m case; 50/100 MHz (090U), 100 MHz (400U); 9 to 16 W	10 min warm-up, 25 °C, spur reject standard on

Selecting a variant

Take the **ICX-090U** for nodes below 9.5 GHz, which covers most civil monitoring and direction finding, and the **ICX-400U** where coverage reaches microwave backhaul or radar bands, accepting ±3.0 dB above 9.5 GHz. Take the **±50 ns GNSS option** on every node in a locating network, because section 5 shows nothing below it meets a 25 m specification, **option 01** where time must also survive an outage, and **option 40 or 41** wherever site ambient plus the rise from Eq. 5 exceeds 0 to 50 °C.

Geometric dilution of precision
The factor by which node geometry multiplies measurement error into position error. A property of the layout, with no instrument parameter in it.

Pairwise matching
How closely two units of a set agree with each other, as distinct from how closely each agrees with an absolute standard.

Traceability
An unbroken chain of comparisons linking a calibration to a national standard, each with a stated uncertainty.

Symbols used in this paper.

Symbol	Meaning	Units
c	Speed of light	m/s
Δr, $\Delta\tau$	Range and arrival-time difference	m, s
G	Geometric dilution of precision	dimensionless
σ_τ, σ_p	Time and position uncertainty	s, m
ΔG, ε_θ	Gain mismatch, bearing error	dB, degrees
σ_{est}, σ_{sync}, σ_{cal}, σ_q	Estimation, time-transfer, residual delay and quantization uncertainty	s
τ_{th}, $m\,c_p$	Thermal time constant, thermal mass	s, J/K
θ, L, A	Subtended angle; program life; tolerance half-width	degrees, years, dB
P, R_{th}, k_T	Dissipation, thermal resistance, amplitude coefficient	W, K/W, dB/K
N^*, C_{build}, C_{buy}, M, h	Crossover volume, the two fixed costs, single-sourced parts, end-of-life rate	units/yr, currency, count, 1/yr

API Application programming interface
CISPR International Special Committee on Radio Interference
DDC Digital down-converter
GNSS Global navigation satellite system
IQ In-phase and quadrature baseband samples
OCXO Oven-controlled crystal oscillator
TCXO Temperature-compensated crystal oscillator
TDOA Time difference of arrival

Further reading

- IEEE Std 1588-2019, *Precision Clock Synchronization Protocol*, and IEC 62402, *Obsolescence management*.
- Recommendation ITU-R SM.2060-0, *Test procedure for measuring direction finder accuracy*, and ITU-R SM.854-3 on direction finding at monitoring stations.
- Berkeley Nucleonics BNC-AN-102, *Wide-area spectrum monitoring with automated classification*, and the ICX-FieldHawk USB and LAN module datasheet.

If you are sizing an embedded receiver into a product of your own, our application engineers would be glad to work the timing budget and the build-against-buy arithmetic against your volume and your emitters.

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.

- Every cost, thermal, calibration-price and obsolescence-rate value in Tables 1 and 3 and in section 7 is illustrative and labeled as such. The 25 ns 1σ synchronization term in Table 1 treats the published ±50 ns 1PPS bound as approximately 2σ; confirm it against the calibration report.
- SoapySDR and GNU Radio operation, including installation of device calibration files with the driver, is documented in a supplied integration guide rather than on a datasheet (verify). No export control classification is stated here; obtain it in writing, and confirm standard designations and editions against each issuing body.
- The field evidence is reported from account history rather than a published source, and no end user, vendor or country is named.
- A pairwise-matching capture across several units would evidence section 3 directly and does not yet exist; the argument there rests on the published amplitude accuracy and on Eq. 3.

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

BNC-AN-110 Rev A
August 2026