Analyzing 5G, Wi-Fi and Bluetooth in one instrument

Why multi-radio analysis is a continuity problem, and what a front end must do before any protocol decoder can help.

ABSTRACT

A device carrying an FR1 modem, an 802.11ax radio and a Bluetooth controller fails in ways no single-radio measurement reproduces. This paper derives why: a stepped receiver holds at most one Bluetooth packet in forty-nine intact, a composite error vector magnitude figure hides a rare severe symbol by construction, and a stitched carrier carries five independent amplitude, phase and timing references. The arithmetic becomes six numbered requirements a reader can score any receiver against, with a plain statement of where a 100 MHz gap-free window is not enough.

WHO THIS PAPER IS FOR

Test-equipment engineers specifying the front end inside a protocol analysis product. **Integration engineers** chasing a symptom that appears only when two radios are active. **Coexistence engineers** asked to prove simultaneity with instruments that measure one thing at a time.

CONTENTS

1 Three radios, one enclosure, and no way to see them together

The report never arrives from an RF engineer. It arrives from a systems engineer, and it says the Bluetooth link drops when the modem uplink is busy, or that Wi-Fi throughput falls forty percent when the audio stream starts. A telematics unit, a medical gateway and an industrial access point share an anatomy: separate radios, one antenna structure, one ground plane, one power rail. The measurements available are per-radio, taken at different times on different instruments. Coexistence is a question about simultaneity, and it is being investigated with non-simultaneous data.

The escape hatch a competent engineer reaches for first is a faster sweep. Section 3 derives why it cannot work: sweeping faster shortens the dwell in exactly the proportion it shortens the revisit interval. The second is a longer averaging interval, and that fails worse, because averaging is what hides the defect.

THE FRAMING WORTH INSTALLING FIRST

Wireless protocol analysis asks two things of an RF front end: enough instantaneous bandwidth to hold the signal whole, and IQ clean enough and continuous enough to demodulate. Everything above that line is software. The failures that cost the most are not sensitivity failures. They are continuity failures.

2 What the signals actually require

3GPP defines a scalable orthogonal frequency-division multiplexing (OFDM) numerology indexed by one integer. Fix it and every timing quantity in the frame follows:

$$\Delta f = 2^{\mu} \times 15\ \text{kHz} \quad \text{and} \quad T_{\text{slot}} = 1\ \text{ms} / 2^{\mu} \tag{1}$$

where

μ is the numerology index, 0 to 4 in Release 15 (1 is the FR1 workhorse)

Δf is the subcarrier spacing, in hertz

T_{slot} is the slot duration, fourteen symbols with normal prefix

Table 1 evaluates it. The normal cyclic prefix is 144 × 64 × $2^{-\mu}$ basic units of 0.5086263 ns, which at 30 kHz is 2.34375 µs.

Table 1. The cyclic prefix costs 6.57 percent of air time at every numerology, because 144/(2048+144) is a constant. NR trades prefix length against symbol length at fixed overhead, which is why one analysis bandwidth serves every row.

µ	Δf (kHz)	T_u (µs)	T_{CP} (µs)	Symbol (µs)	Slots / subframe	Used in
0	15	66.667	4.6875	71.354	1	FR1
1	30	33.333	2.3438	35.677	2	FR1, dominant
2	60	16.667	1.1719	17.839	4	FR1, FR2
3	120	8.333	0.5859	8.919	8	FR2
4	240	4.167	0.2930	4.460	16	SSB only

Occupied bandwidth is fixed by the resource-block allocation:

$$B_{\text{occ}} = N_{\text{RB}} \times 12 \times \Delta f \tag{2}$$

where

N_{RB} is the maximum transmission bandwidth configuration, TS 38.104 Table 5.3.2-1 (273 at 100 MHz and 30 kHz)

12 is subcarriers per resource block, fixed in NR

A 100 MHz FR1 carrier at 30 kHz spacing occupies 273 × 12 × 30 kHz = 98.28 MHz, leaving 0.86 MHz of guard per side. And 100 MHz is the largest single-carrier bandwidth 3GPP defines for FR1; above it, NR aggregates carriers rather than widening one. So a gap-free 100 MHz window is not merely enough bandwidth for a 5G signal: it is exactly one FR1 carrier, edge to edge, under one amplitude reference. The NR synchronization block sits on the channel raster rather than at carrier center, so an analyzer narrower than the carrier can be tuned to the carrier and still never see it *(verify)*.

The 2.4 GHz band lands the same way. The industrial, scientific and medical allocation is 83.5 MHz wide, so one record centered at 2441.75 MHz holds all 79 Bluetooth channels, all 40 Low Energy channels and Wi-Fi 1, 6 and 11 at once.

Reading a record before demodulating it

A cyclic prefix is a copy of the end of a symbol prepended to it, so two samples separated by the useful symbol length are identical whenever the earlier lies inside a prefix. That gives the metric of van de Beek, Sandell and Börjesson:

$$M(d, N) = |\ \Sigma\, r[n] \cdot r^*[n+N]\ |\ \ /\ \ \Sigma\, |r[n+N]|^2, \quad n = d \dots d+L-1 \tag{3}$$

where

$r[n]$ is complex baseband sample n of the record

N is the lag, the useful symbol length (4096 samples at 30 kHz, 122.88 MSa/s)

L is the window, the prefix length (288 samples at 30 kHz)

Run it as a hypothesis test over the four lags that occur in practice, those for 15, 30, 60 and 120 kHz spacing. Only the true lag yields a regular train of peaks. Symbol timing, the half-subframe grid and the frequency offset come free with it, at a coherent gain of 10 $\log_{10}(288)$ = 24.6 dB, for one pass over the record.

> **TAX ON THE IDEALIZATION JUST GRANTED**
>
> The prefix estimator resolves frequency offset only within half the subcarrier spacing, 15 kHz at 30 kHz. Beyond that it wraps, and the integer part needs the primary synchronization signal.

3 Why swept and segmented architectures fall short

A 100 MHz carrier stitched from five 20 MHz segments is five acquisitions, each with its own amplitude calibration, its own local-oscillator phase and its own trigger instant, so the composite carries five references where the signal has one. An error vector magnitude computed across a seam measures the instrument, not the transmitter. That is the second and better reason to forbid stitching in R1, and it is independent of bandwidth.

The antagonist is an architecture. Swept and stepped-tuned receivers, and the stitched measurement built on them, are excellent at characterizing stationary signals over wide spans at high dynamic range. Neither was designed for signals whose defining property is that they change.

Mechanism one, probability of intercept

Filter settling is not the binding constraint on Bluetooth; 83.5 MHz at 1 MHz resolution settles in roughly 209 µs. Dwell is. To capture a packet intact the receiver must stay on its channel for the packet's whole duration:

$$t_{\mathrm{dwell}} = T_{\mathrm{sweep}} \cdot (\mathrm{RBW} / S) \;\Rightarrow\; T_{\mathrm{sweep}} \geq T_{\mathrm{pkt}} \cdot S / \mathrm{RBW} \tag{4}$$

where

S is the span, 83.5 MHz for the 2.4 GHz band

RBW is the resolution bandwidth, 1 MHz to resolve one BR/EDR channel

T_{pkt} is the packet air time, at most 366 µs for a single-slot DH1 *(verify)*

A sweep long enough to hold one DH1 packet needs 366 µs × 83.5 = 30.56 ms, and in that interval Bluetooth executes 1600 × 0.03056 = 48.9 hops. One was caught intact and forty-eight were not, so the best-case intercept fraction is 1/48.9 = **2.05 percent**. Sweeping faster does not rescue it: at a 1 ms sweep the dwell is 12 µs, 3.3 percent of a packet, a fragment with no access code and no address. Both branches lose.

Now grant the stepped receiver everything and see whether it recovers. Take away the retune time entirely, let it settle instantly, and widen its instantaneous bandwidth. The dwell constraint becomes $T_{\mathrm{sweep}} = T_{\mathrm{pkt}} \cdot S/B$, and the fraction of packets caught intact collapses to a single ratio:

$$f = \min[1, B / (h \cdot T_{\mathrm{pkt}} \cdot S)] \tag{5}$$

where

B is the instantaneous bandwidth held at once, in hertz

h is the hop rate, 1600 per second for BR/EDR

S is the band to be covered, 83.5 MHz

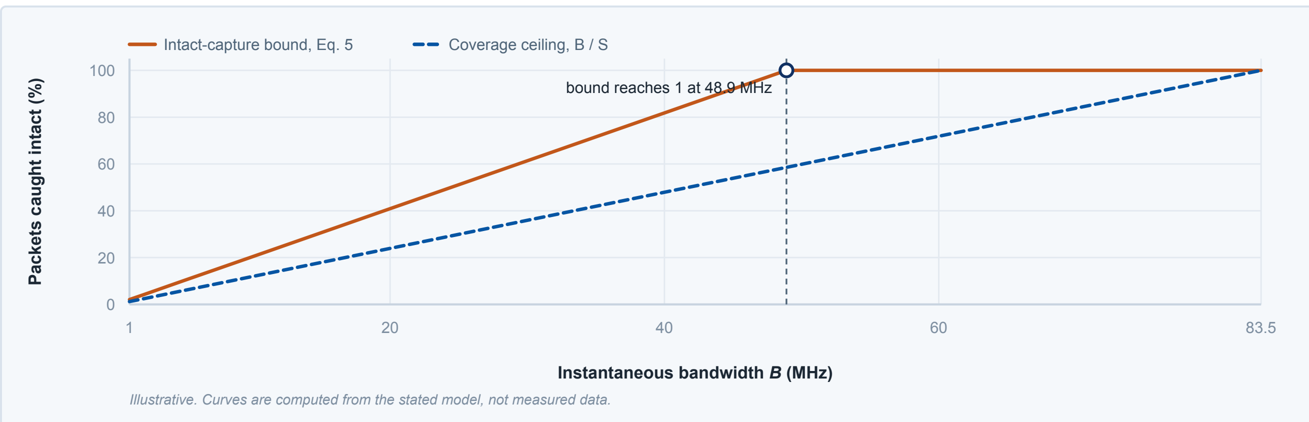

Figure 1. The bound of equation (5) against the coverage a receiver actually has. Coverage rises linearly in B the whole way; the bound rises linearly until it saturates at 48.9 MHz. Neither depends on sweep rate or retune time, which is the point. Below the knee the bound runs 1.71 times above coverage, so where it first reaches 1 the receiver is still blind to 41 percent of the band at every instant.

DERIVED

Figure 1 draws both. A stepped receiver can be argued up the orange line by engineering; it cannot be argued past the blue one by anything at all.

THE RESULT THAT CLOSES THE ARCHITECTURE QUESTION

At 1 MHz the idealized receiver catches at most 2.0 percent, at 20 MHz at most 40.9 percent, at 40 MHz at most 81.8 percent, and the bound reaches 1 at 48.9 MHz, where it stops constraining anything. Read it as a ceiling and not a promise: a receiver holding 48.9 MHz of an 83.5 MHz band is blind to the rest at every instant, so certainty needs $B = S$, and a receiver that holds the whole band at once is not a stepped receiver. The retune time never appeared. Speed never appeared. **The best a stepped receiver can do is set by how much of the band it holds at once**, and every other specification on its datasheet is an argument about how quickly it fails.

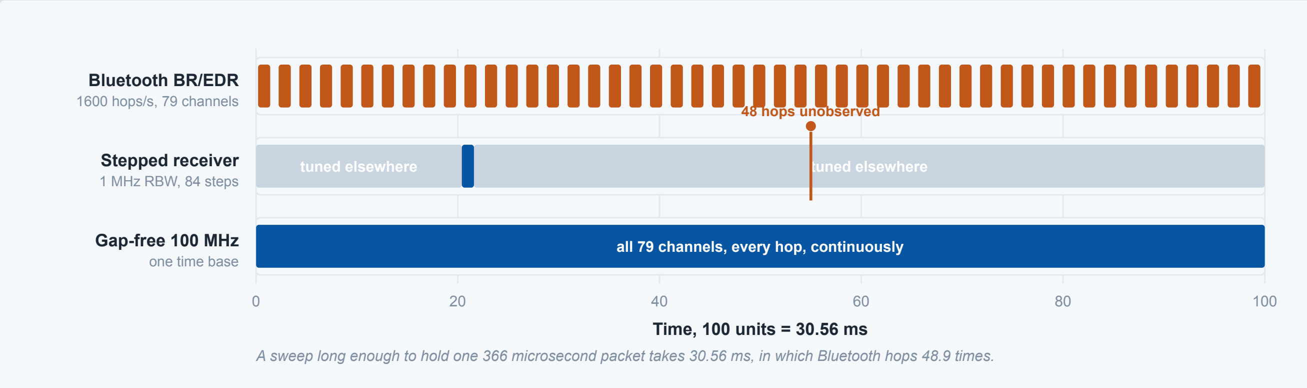

Figure 2. The best case for a stepped receiver on Bluetooth BR/EDR is one intact packet in 48.9, or 2.05 percent. What is lost is not power accuracy but the sequence: the hopping map, the retransmission pattern, the address in the access code. SCHEMATIC

Figure 2 draws that arithmetic, and the same physics governs any burst: an SS/PBCH block lasts 142.7 μs and repeats every 20 ms, so it is present 0.71 percent of the time.

Mechanism two, periodic blindness

Retune dead time is not random; it repeats at the revisit period. A duplex pattern switching every 2.5 ms, watched by an analyzer revisiting five slices on a 25 ms period, is sampled at the same phase every time. If the dead time overlaps the switch, the transient is never observed. Not rarely. Never. Running longer does not help.

Mechanism three, the composite that hides the defect

Error vector magnitude combines in root-sum-square, because device and instrument error vectors are uncorrelated:

$$\mathrm{EVM_{meas}}^2 = \mathrm{EVM_{DUT}}^2 + \mathrm{EVM_{inst}}^2 \quad (6)$$

where

EVM_{DUT} is the true device error vector magnitude, percent of the RMS reference

EVM_{inst} is the instrument residual, measured on a known-good loop at the frequency and bandwidth actually used

A transmitter sitting exactly at the 3GPP 256QAM limit of 3.5 percent reads sqrt(12.25 + 4.00) = 4.031 percent on an instrument with a 2.0 percent residual: a wrong verdict, not a less precise one.

The averaging is worse. Take a 70-symbol window in which 69 symbols sit at 2.5 percent and one, at a duplex switch, sits at 15 percent. The composite is sqrt[(69 × 6.25 + 225)/70] = sqrt(9.375) = **3.06 percent**, passing with margin. That symbol exceeds even the 16QAM limit of 12.5 percent, so it is undecodable and it fails the whole slot: one slot in five, a 20 percent block error rate, a throughput collapse no EVM report explains. It shows in one view only, EVM against symbol index over a contiguous record, which a segmented acquisition cannot produce.

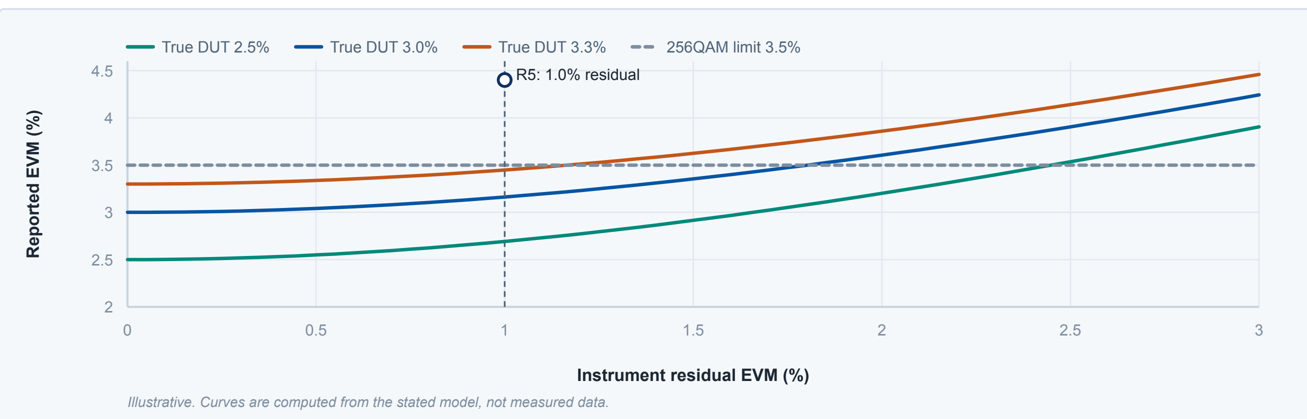

Figure 3. A transmitter with 14 percent true margin against the 256QAM limit is failed by an instrument residual of 1.8 percent and passed by one of 1.0 percent. At this modulation order the residual decides the verdict, it does not refine the number. DERIVED

Figure 3 plots the crossings. What it shows is that a receiver residual near the limit does not degrade the measurement gracefully; it moves the pass boundary.

4 Derived requirements

Table 2 turns those mechanisms into six numbers, each with the reason a looser value is wrong rather than merely worse. Score any receiver with it.

Table 2. Six requirements. R3, R4 and R6 are the ones most receivers miss, and section 7 says plainly that this one misses all three and cannot be scored on R5.

	Requirement	Number	Why not looser
R1	Gap-free, calibrated, flat analysis bandwidth, no stitching	≥ 100 MHz	80 MHz misses 18.28 MHz of the carrier; 98.28 MHz leaves the edges in the roll-off.
R2	Sample rate with an exact rational path to the native rate	≥ 122.88 MSa/s	122.88 MSa/s is exactly 3072/3125 of 125; any other ratio forces an interpolating resampler.
R3	Sustained, verified-lossless record throughput	500 MB/s for 60 s	8-bit halves the rate but yields 49.9 dB against the 85 dB R4 needs.
R4	Instantaneous usable dynamic range in the capture bandwidth	≥ 85 dB	75 dB from a −10 dBm Wi-Fi peak to a −85 dBm peripheral, plus 10 dB of crest.
R5	Instrument residual EVM at the modulation order under test	≤ 1.0 percent	At the 3.5 percent limit, 1.0 percent inflates it by 4.0 percent, 2.0 percent by 15.2.
R6	Frequency reference accuracy, lock state logged per record	≤ 0.005 ppm	The 0.05 ppm base-station limit is 175 Hz at 3.5 GHz; measuring it needs ten times that.

R3 is where gap-free stops being a word and becomes an engineering constraint:

$$\dot{R} = f_s \times 2 \times (b/8) \qquad (7)$$

where

f_s is the complex sample rate, in samples per second

b is the bits per component, 16 in a conventional format

At 125 MSa/s with 16-bit components that is 500 MB/s, or 4.0 Gbit/s of payload. Table 3 explains how an instrument can carry a 100 MHz converter and still not be gap-free.

Table 3. Two of these five paths cannot sustain 4.0 Gbit/s and a third fails once its cache is gone, which is why sustained throughput, not converter bandwidth, is the specification to interrogate.

Link	Payload ceiling	Sustained	Verdict at 500 MB/s
USB 3.2 Gen 1	4.0 Gbit/s, 8b/10b	about 400 MB/s	insufficient
USB 3.2 Gen 2	9.7 Gbit/s, 128b/132b	875 MB/s and up	sufficient
1 GbE	0.94 Gbit/s	118 MB/s	insufficient by four times
SATA III drive	4.8 Gbit/s, 8b/10b	500 MB/s in burst	fails after the cache
NVMe PCIe 3.0 ×4	about 3.9 GB/s	2 to 3.5 GB/s	sufficient

5 Meeting the requirements

A receiver satisfying R1 through R6 is a narrow class of instrument. The ICX-FieldHawk family meets R1 and R2 as published, and section 7 says where it does not and where it cannot be scored. Published, at 25 °C after ten minutes of warm-up with spur reject standard on: 100 MHz analysis bandwidth on the ICX-400, ICX-090R, ICX-400R and ICX-400U, 50 MHz standard with 100 MHz optional on the ICX-090U; IQ sample rate to 125 MSPS with decimation 1 to 4096; burst recording to 100 MHz into 128 Mbyte; continuous recording to 25 MHz.

R1 and R2 go together. A 100 MHz window holds an FR1 carrier plus both guard bands, and the whole ISM band. 125 MSPS clears the 122.88 MSa/s native rate, and 122.88/125 is exactly 3072/3125, so one polyphase stage reaches the standard's grid without interpolation. Ask any vendor for that numerator and denominator; resampler error arrives as EVM you will blame on the device.

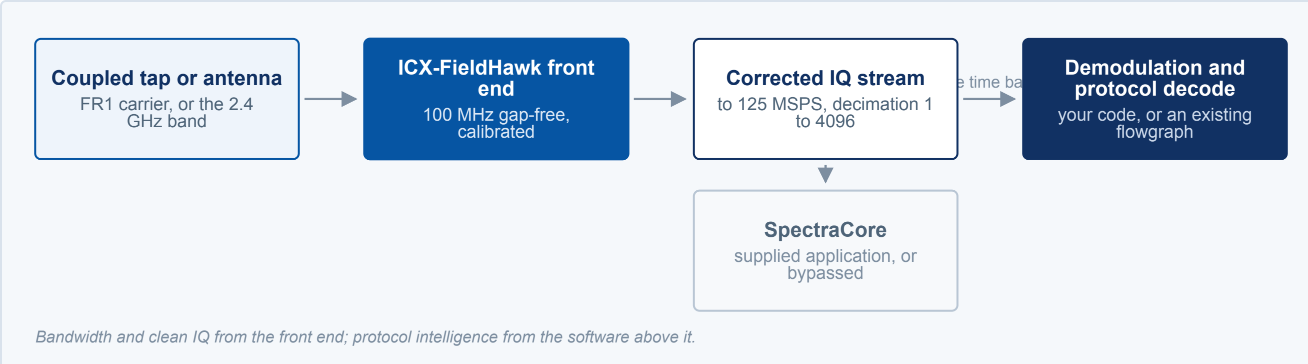

Figure 4. The division of labour that makes multi-radio analysis tractable. Corrected, continuous IQ on one time base lets the decoders above it be written once, in C, C++, Python, MATLAB or LabVIEW, against a receiver that does not drop samples. SCHEMATIC

Figure 4 draws the split. It matters most to anyone embedding a receiver in a product they sell: a dropped sample becomes a decode failure that carries their name, in code they cannot debug. The supplied application, SpectraCore, sits at the same boundary: it is where the corrected samples surface, and it is the layer a product design is free to replace without giving up the correction underneath it.

Now the hopping map. Treating the hops as sampling channels uniformly, the expected number of hops to touch all 79 at least once is 79 × (ln 79 + 0.5772 + 1/158) = 391, or **244 ms** at 1600 hops per second. The published 128 Mbyte capture memory, filled at 100 MHz with 16-bit components, holds 128 MB / 500 MB/s = **0.26 s**.

> **THE MEMORABLE RESULT**
>
> One fill of the built-in capture buffer at full bandwidth is 0.26 s. The expected time for Bluetooth to visit all 79 of its channels at least once is 0.244 s. A single burst capture is, to within seven percent, exactly one adaptive frequency hopping map. It also holds twenty-five 10 ms NR radio frames.

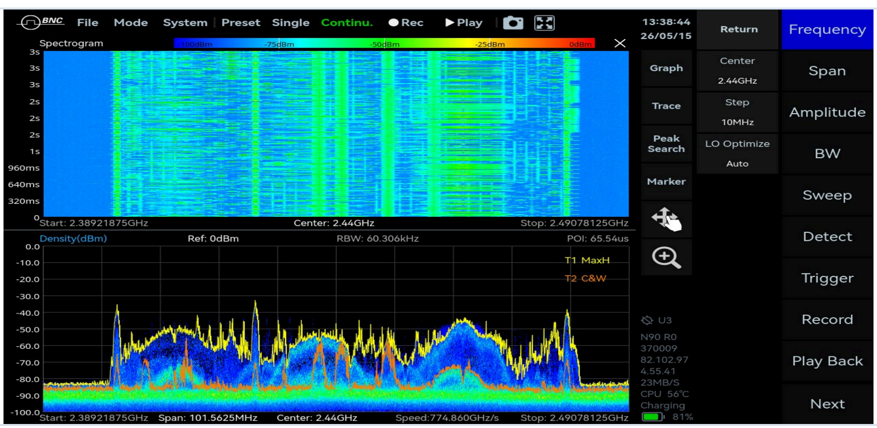

Figure 5. Real-time density on an ICX-FieldHawk, 2.389 to 2.491 GHz in a 101.56 MHz span at 60.3 kHz resolution, the display reporting a 65.54 microsecond probability of intercept. The density surface, which the real-time engine accumulates from every transform rather than from the peak of a sweep, means a hopping emitter resolves into a channel-occupancy statistic instead of a max-hold smear. MEASURED

Figure 5 follows from the real-time engine, whose published relation makes the guarantee computable. One hundred percent probability of intercept is 2 × N × D × 8 ns: 32.768 μs at 2048 points, 0.512 μs at 32. The single-slot DH1 packet worked in section 3 is eleven times the first figure and a 1 μs symbol twice the second, so neither sits near the boundary. Decimation is then the storage lever: at 64, one channel records at 7.8 MB/s, turning a minute from 30 GB into 469 MB.

Decoders that already exist

None of R1 through R6 asks for a software ecosystem, and that is the point. Once the IQ is gap-free, corrected and on one time base, the decoders above it are somebody else's finished work. The ICX-FieldHawk presents itself through SoapySDR, the vendor-neutral hardware abstraction layer, so it appears to GNU Radio, Gqrx or any other SoapySDR application as an ordinary software-defined radio device. The instrument calibration files install with the driver, which means the samples arriving in a flowgraph are corrected before the first block sees them.

For multi-radio work that matters more than it does anywhere else, because the protocol decoders for these three radios already exist. An 802.11a receiver, complete with packet detection, synchronization, transform, channel equalization and MAC decoding, is an existing GNU Radio flowgraph. Connecting it to a calibrated front end is a matter of changing the source block, not of writing a demodulator.

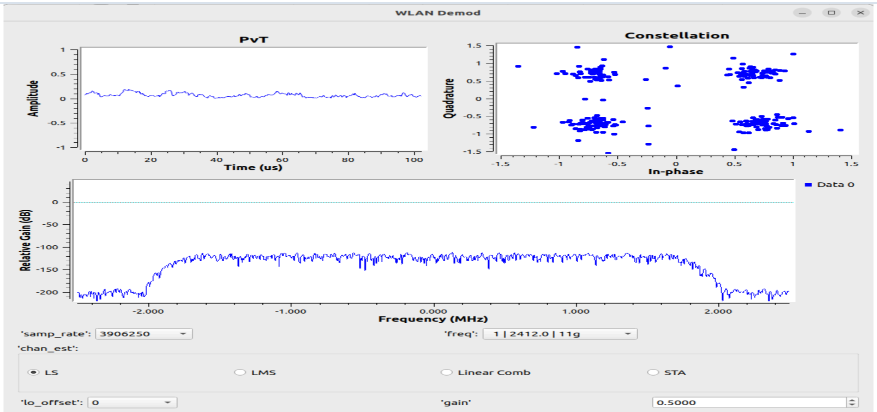

Figure 6. The 802.11a/g OFDM physical layer recovered at 2.412 GHz, with the flowgraph's own channel selector reading 11g and its least-squares channel estimator selected. The 12 Mb/s QPSK at rate 1/2 and the −40 dBm level are the integration guide's stated conditions rather than fields on this window. An existing flowgraph and a calibrated front end together produce a decoded frame with a traceable amplitude behind it. The frames leave as a PCAP file, so the coexistence argument continues in Wireshark rather than in a screenshot. MEASURED

Figure 6 is the result of exactly that substitution. The same path has been taken through 16-QAM demodulation, which the integration guide records as recovering a clean constellation from a −80 dBm, 500 kSym/s carrier at a root-raised-cosine roll-off of 0.35, and through ADS-B decoding to aircraft address and position. What makes it useful here is not the demodulation, which any receiver can be made to do, but that the amplitude underneath every decoded frame is still traceable. A coexistence dispute settled with a PCAP file and a calibrated power number is settled; one settled with a PCAP file alone is an opinion with better formatting.

HONEST LIMITATION

This path is demonstrated on x86_64 hosts running Ubuntu 22.04 or later with GNU Radio 3.9 or later over USB 3.0, and it does not yet appear on a datasheet. Treat it as demonstrated rather than specified, and ask before planning around a different environment.

6 Measurement practice

None of this depends on which analyzer is on the bench.

1. **Prove gap-free before trusting the record.** Confirm the file holds the sample rate times the duration, to the sample. Then feed a tone off center, unwrap the phase and fit a line: a dropped block of 4096 samples at 125 MSa/s is a 33 μs discontinuity. Repeat to storage for ten minutes.
2. **Set the reference level from the instantaneous-power distribution, not a peak marker.** Take the 0.01 percent point and add 3 dB. A peak marker never saw the OFDM peak, and clipping manufactures the interference you are investigating.
3. **Run the cyclic-prefix autocorrelation before demodulating anything.** One pass gives subcarrier spacing, symbol timing, the half-subframe grid and the frequency offset. If it disagrees with the configuration you were given, resolve that first.
4. **Capture the whole 2.4 GHz band once and filter in software**, and record the RF path and reference lock state with the samples. Retuning between aggressor and victim destroys the only property coexistence analysis is about, and a dBm is meaningless without the path behind it.
5. **Plot EVM against symbol index, not just composite RMS.** It costs nothing on a contiguous record and is the most diagnostic view in OFDM debug.

Field evidence

The boundary in section 5 is not a theoretical one. It is where a wireless test vendor drew it when they built a multi-radio analysis product on a receive module and wrote the protocol layer themselves, and the reason is the one in section 3: from the customer's side, a decode failure caused by a dropped sample is indistinguishable from a decode failure in the vendor's own code. The arrangement is not the evidence. The evidence is where the party with the strongest reason to own the front end decided the line was.

7 Honest limits

- **A 100 MHz window is not the whole problem.** 802.11ax at 160 MHz, 802.11be at 320 MHz and carrier-aggregated FR1 do not fit, so cross-component behavior goes unmeasured.
- **Burst and continuous recording are different specifications.** The ICX-FieldHawk publishes burst recording to 100 MHz into 128 Mbyte, about 0.26 s, and continuous recording to a host at 25 MHz. R3 asks for 500 MB/s sustained for 60 s at full bandwidth, which is not what this instrument publishes. If a measurement needs a minute of gap-free 100 MHz, this is not the instrument.
- **R5 is not answered here.** No residual error vector magnitude figure is published for this receiver at any modulation order, so the requirement Figure 3 makes the case for cannot be scored from a datasheet, ours or anyone else's. Measure it yourself, on the loop described under equation (6), and treat a residual quoted without its frequency and bandwidth as unusable.
- **The coexistence measurement is dynamic-range-limited.** R4 asks for more than a real converter delivers at these rates, and no dynamic-range figure is published. Use the proxies: display range to +20 dBm typical on the 40 GHz tier, third-order input intercept +40.3 dBm at 1 GHz with reference level +20 dBm, IF flatness ±2.0 dB. Beyond that, take two captures at different attenuation.
- **A conformance frequency verdict needs an external standard.** The published references, a TCXO below 1 ppm and an OCXO, option 01, below 0.15 ppm, give 525 Hz at 3.5 GHz against a 175 Hz limit: enough for hop tracking, not for adjudication.
- **If only one radio is ever active, a swept analyzer is the better tool.** Higher dynamic range, better close-in performance, lower cost.

8 Summary, definitions and further reading

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

Table 4. Published values with their conditions: 25 °C, ten minutes of warm-up, spur reject on.

Quantity	Relation or value	Condition
Occupied bandwidth	$N_{RB} \times 12 \times \Delta f$	98.28 MHz at 100 MHz, 30 kHz
Intact-capture sweep time	$T_{pkt} \cdot S$ / RBW	30.56 ms, DH1 across 83.5 MHz at 1 MHz
Best-case hop intercept	2.05 percent	one packet in 48.9, stepped receiver
Intact-capture ceiling, any stepped receiver	$B / (h \cdot T_{pkt} \cdot S)$	upper bound, not achievable; certainty needs $B = S$
Analysis bandwidth	100 MHz	ICX-400, 090R, 400R, 400U; 50 MHz on the 090U
Recording	burst 100 MHz into 128 Mbyte; continuous 25 MHz	about 0.26 s at 16-bit *(format not published, verify)*
100 percent POI	$2 \times N \times D \times 8$ ns	32.768 µs at N = 2048, D = 1
Frequency reference	TCXO < 1 ppm; OCXO option 01 < 0.15 ppm	after warm-up

Source: the ICX-FieldHawk handheld, rugged and USB datasheets.

Selection

Where a display matters, the ICX-400 handheld covers 9 kHz to 40 GHz at 100 MHz. Embedded, the ICX-400U gives the same bandwidth over USB or LAN; the ICX-090U covers 9 kHz to 9.5 GHz at 50 MHz standard, 100 MHz optional, enough for the 2.4 GHz band but not a full FR1 carrier. Option 71 adds digital demodulation, option 01 the OCXO.

Adaptive frequency hopping

The Bluetooth mechanism that drops congested channels from the hop sequence. A minimum number must stay in use, so a busy band forces some hops under interference.

Error vector magnitude

The root-mean-square distance between measured and ideal constellation points, normalized to the RMS reference.

Gap-free
Every sample reaches the record, with a verified count and no dropped block, for the stated duration. A claim about the path to storage, not the converter.

Probability of intercept
The minimum duration a signal must present for its amplitude to be measured correctly.

Symbols used in this paper.

Symbol	Meaning	Units
μ, Δ*f*	Numerology index, subcarrier spacing	dimensionless, Hz
T_u, T_{CP}	Useful symbol, cyclic prefix duration	s
N_{RB}, B_{occ}	Resource blocks, occupied bandwidth	Hz
f_s, *b*	Complex sample rate, bits per component	Sa/s, bits
S, RBW	Span, resolution bandwidth	Hz
f, *B*, *h*	Intact-capture ceiling; instantaneous bandwidth; hop rate	dimensionless, Hz, s^{-1}
T_{pkt}	Packet air time	s

BR/EDR Basic rate and enhanced data rate
EVM Error vector magnitude
FR1 Frequency range 1, the 3GPP sub-7.125 GHz range
ISM Industrial, scientific and medical
OFDM Orthogonal frequency-division multiplexing
POI Probability of intercept
SSB Synchronization signal and PBCH block

Further reading

- 3GPP TS 38.104, *NR; Base Station radio transmission and reception*: 5.3.2 for channel bandwidth and resource blocks, 6.5.1 for frequency error, 6.5.2 for EVM limits; TS 38.211 4.1 to 4.4 for frame structure and prefix lengths. IEEE Std 802.11-2020, 802.11ax-2021 and 802.15.2-2003, and the Bluetooth Core Specification, cover the rest.
- J.-J. van de Beek, M. Sandell and P. O. Börjesson, "ML estimation of time and frequency offset in OFDM systems," *IEEE Trans. Signal Processing*, vol. 45, no. 7, 1997, and A. Peled and A. Ruiz, *Proc. IEEE ICASSP*, 1980, which introduced the cyclic prefix.
- The SoapySDR project, for the abstraction layer in section 5. BNC-AN-101 sets out the calibrated open-instrument argument in full; BNC-AN-102 and BNC-AN-105 cover monitoring and record and playback.

If you are specifying a front end for a multi-radio product, or settling a coexistence argument with a measurement, our application engineers would be glad to work through these requirements against your own signals.

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.

- The Bluetooth minimum used-channel count and the 366 microsecond DH1 air time should be confirmed against the Core Specification version cited; the 2.05 percent result depends on it.
- The TS 38.104 EVM limits and the SS/PBCH raster offset in TS 38.101-1 should be confirmed against the releases cited.
- IQ sample formats are not published, so the 0.26 second figure assumes 16-bit I and Q. No dynamic-range figure, onboard AI compute, EMC function or VITA 49 is claimed here.
- The SoapySDR and GNU Radio path is documented in a platform integration guide and demonstrated with captured results, but does not yet appear on a datasheet. The distribution route for the driver package is not settled, so no package name, repository 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-108 Rev A
August 2026