ICX-FIELDHAWK TECHNICAL PAPER
BNC-AN-104 REV A
AUGUST 2026

Commissioning and maintaining a satellite earth station

The recorded baseline is the document that settles the interference dispute you have not had yet.

ABSTRACT

Earth stations are built once and defended for years. This paper works through the measurements a field engineer actually makes, from the pre-build site survey to the interference call three years later: the link budget and where each decibel is measurable, figure of merit by the sun method, cross-polarization at line-up, spectral regrowth against back-off, and cable verification by scalar sweep. It then derives why a short swept survey is close to worthless against an intermittent interferer, and what the archived baseline is actually for.

WHO THIS PAPER IS FOR

Earth station and teleport engineers commissioning or maintaining stations. **Ground segment managers** specifying acceptance procedures. **Interference investigators** building a case that has to survive scrutiny.

CONTENTS

1 The chain, and where each fault lives

The transmit path runs from the modem, along the inter-facility link, into the block upconverter whose local oscillator is locked to the station's 10 MHz reference, through waveguide to the orthomode transducer, the feed and the reflector. The receive path mirrors it through the low-noise block. Faults do not wander; each class has an address.

The reasonable first instinct is to skip all of this and let the modem tell you. It reports carrier-to-noise, it logs errors, and it is already installed. The trouble is that the modem reports one number for every fault on the chain: a wet connector, a drifting reference, a misaligned feed and an adjacent-satellite carrier all arrive as degraded carrier-to-noise, and the modem cannot separate them because it only sees what survived. It is an excellent detector of trouble and it localizes nothing. Everything below exists to turn one alarm into an address.

PHOTOGRAPH NEEDED

Placeholder for a photograph of an earth station showing feed height and waveguide runs

Figure 1. What the block diagram below cannot show: scale, exposure and access. A photograph of a ground segment is wanted here because the geometry is the argument. Feeds sit several meters up on structures that need a lift to reach, the waveguide runs to them are long enough for their own thermal expansion to matter, and every flange along them is a mechanical joint in weather. The measurements in this paper are designed around that geometry rather than around the diagram below, because a diagram implies that reaching any point on the chain costs the same, and it does not. SCHEMATIC

Figure 1 is the half of the station that weather reaches and a technician cannot. The block diagram below it draws the line every fault sits on, and the two together are the map the rest of this paper indexes: one gives the topology, the other gives the cost of getting to any point on it.

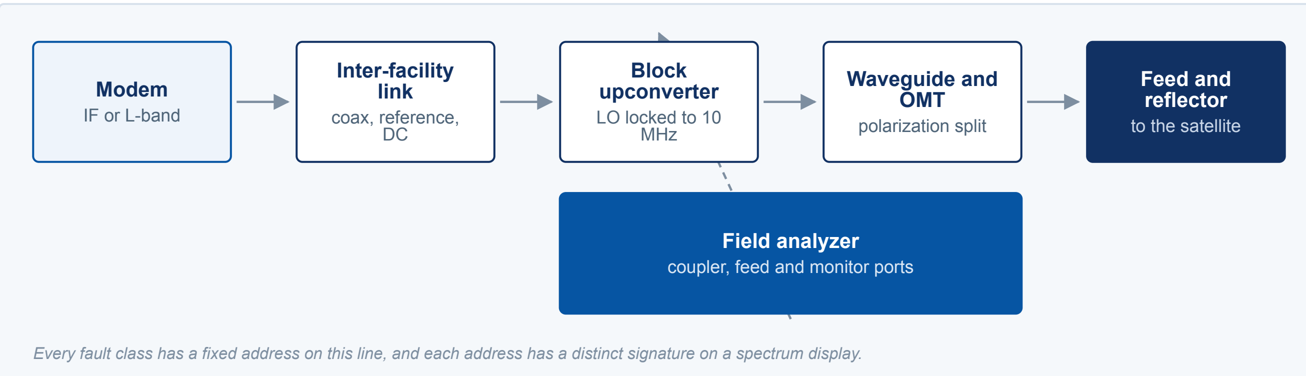

Figure 2. The transmit chain, with the points a field analyzer is actually connected to. Diagnosis is a routing problem: identify the signature, and the signature identifies the element. SCHEMATIC

Table 1 maps Figure 2 onto the signatures each element produces when it fails.

Table 1. Fault classes by address in the chain. The last row is the one that catches people out: an out-of-band blocker compresses the low-noise block and every carrier drops together, with nothing to see inside the band a filtered analyzer is looking at.

Element	Fault class	What it looks like on a display
Inter-facility cable	loss slope, water ingress, intermittent contact	level slope with frequency, standing-wave ripple, carrier jumping with vibration
10 MHz reference	phase noise, harmonics	widened carrier skirts after multiplication, spurs at multiples of 10 MHz
Block upconverter	LO leakage, compression, spectral regrowth	discrete spur at the LO, shoulders on the carrier
Waveguide run	moisture, arcing, flange misalignment	resonant suck-outs in a swept trace, broadband noise under transmit power
Feed and orthomode transducer	depolarization, wet radome	cross-polarization discrimination degraded
Low-noise block	gain drift, LO unlock, saturation by an out-of-band blocker	whole-band level shift, or wideband desensitization with nothing visible in band

The frequency plan matters because interference follows allocation. C band pairs a 5.925 to 6.425 GHz uplink with a 3.7 to 4.2 GHz downlink, Ku pairs 14.0 to 14.5 GHz with 10.95 to 12.75 GHz, and Ka pairs 27.5 to 31.0 GHz with 17.7 to 21.2 GHz. Two of those neighborhoods have become crowded: in the United States the 3.7 to 3.98 GHz segment was cleared for flexible terrestrial use and incumbent downlinks repacked into 4.0 to 4.2 GHz, and the 13.75 to 14.0 GHz uplink segment is shared with radiolocation.

2 The link budget, and which decibels you can measure

The satellite link equation is short, and the interesting thing about it is how few of its terms a field engineer owns:

$$C/N_0 = \mathrm{EIRP} - L_{fs} - L_{atm} + G/T + 228.6 \qquad (1)$$

where

$EIRP$ is the effective isotropic radiated power, in dBW; +48 dBW is representative of a Ku satellite toward a site, and +50 dBW of an earth-station uplink

L_{fs} is the free-space path loss, in dB; 205.7 at 12 GHz to geostationary orbit

L_{atm} is gaseous, cloud and rain attenuation, in dB; under 1 dB clear-sky at Ku, tens of dB in heavy rain

G/T is the receive figure of merit, in dB per kelvin; +23.0 for the 1.8 m station worked below and +37.1 for a 9 m one

228.6 is the negative of ten times the log of Boltzmann's constant, in dBW per kelvin per hertz

At a representative geostationary slant range of 38,500 kilometers, free-space loss is 196.2 dB at 4 GHz, 205.7 dB at 12 GHz and 210.2 dB at 20 GHz. Ka band pays fourteen decibels more than C band over identical geometry, and the obvious conclusion from that is the wrong one. A fixed aperture recovers all of it and more: dish gain rises as the square of frequency, so the same reflector buys fourteen decibels at each end of the link, twenty-eight in total, against the fourteen the path took. What Ka cannot recover is rain, which moves by tens of decibels in minutes, and that, not the static budget, is why Ka systems run adaptive coding rather than fixed margin.

Table 2 works a downlink budget through line by line.

Table 2. A worked Ku-band downlink budget. Every decibel in it is fixed by geometry or supplied by the operator except three, and those three, figure of merit, radiated power and interference, are exactly the commissioning measurements.

Line	Value	Where it comes from
Satellite EIRP toward the site	48.0 dBW	the operator's link budget
Free-space loss at 12 GHz	−205.7 dB	geometry
Clear-sky atmosphere	−0.5 dB	gaseous plus pointing allowance
Earth-station G/T, 1.8 m at 150 K	+23.0 dB/K	**measurable on site**
Boltzmann term	+228.6	physics
C/N_0	**93.4 dBHz**	
Carrier at 27.5 Msym/s	−74.4 dB	10 log of the symbol rate
Clear-sky C/N	**19.0 dB**	supports high-order modulation with margin

Figure of merit, measured with the sun

Figure of merit is receive antenna gain divided by system noise temperature, and it is measured in the field by pointing first at cold sky and then at the sun, at the same elevation, and taking the ratio of the two noise powers. That ratio, conventionally called Y, gives:

$$G/T = 8\pi k(Y - 1)L \,/\, (\lambda^2 F) \qquad (2)$$

where

Y is the ratio of sun-pointed to cold-sky noise power, linear; 33.6 for the 1.8 m station and 263 for the 9 m one

k is Boltzmann's constant, 1.380649×10^{-23} J/K

λ is the wavelength at the measurement frequency, in meters; 24.98 mm at 12 GHz

F is the solar flux density at that frequency, in watts per square meter per hertz, from a solar observatory on the day; 1,000 flux units at 12 GHz is 10^{-19} W m^{-2} Hz^{-1}

L is the beam-size correction for the sun's half-degree angular diameter; 1.10 for a 1.8 m dish and 3.52 for a 9 m one

Work it once and it stops being algebra, and the beam correction is where the intuition breaks. A 9 m Ku station measured at 12 GHz has a wavelength of 24.98 mm and a half-power beamwidth near $70\lambda/D$ = 0.194 degrees. The sun subtends about half a degree, so it *overfills* that beam and the correction is large: $L = 1 + 0.38(0.5 / 0.194)^2 = 3.52$. With a measured Y of 263, a 24.2 dB rise between cold sky and the sun, and a solar flux of 1,000 flux units at 12 GHz, which is 10^{-19} W per square meter per hertz, $8\pi k(262)(3.52) / (6.24 \times 10^{-4} \times 10^{-19}) = 5{,}120$, or **37.1 dB/K**.

Run the same measurement on the 1.8 m station in Table 2 and the numbers move the other way. Its beam is 0.97 degrees, so the sun now *underfills* it, the correction falls to 1.10, and a Y of 33.6 returns 199, or **23.0 dB/K**, which is the figure that budget assumed. The two stations differ by

14 dB of figure of merit and by a factor of three in their beam corrections, and it is the correction, not the ratio, that field measurements most often get wrong. Measuring rather than assuming is a two-hour job with no test equipment beyond the analyzer already on the site.

Two practical notes carry more weight than the algebra. Flux interpolation and the beam correction dominate the uncertainty, so a realistic field claim is a few tenths of a decibel rather than hundredths. And the measurement should be made above about twenty degrees of elevation, so that the atmosphere is comparable between the hot and cold looks.

3 Why a short survey proves nothing

Most interference is continuous, and continuous interference is easy: adjacent-satellite carriers, cross-polarization, a terrestrial link in the band, the station's own intermodulation products. Point an analyzer at it and it is there. The cases that consume months are the intermittent ones, and against those a short swept survey is close to worthless.

A swept analyzer covering span S at resolution bandwidth B listens to any one frequency for the fraction B over S of each sweep. For an interferer of burst length tau and period T_p, the probability of catching it on one sweep is the overlap of the dwell and the burst, and detection over an observation window follows the usual geometric law:

$$P_{det}(T) = 1 - (1 - P_1)^{T/T_r}, \quad P_1 \approx (\tau + t_d) / T_p \tag{3}$$

where

t_d is the dwell in one resolution cell, T_r × B/S; 90 μs in the worked case

T_r is the sweep revisit period, in seconds; 0.15 s here

τ is the interferer's burst length, in seconds; 10 ms for the radar, 40 ms for the arcing connector

T_p is the interferer's repetition period, in seconds; 3 s and 30 s respectively

Take a rotating maritime radar in the shared Ku uplink segment: twenty revolutions per minute, so a three second period, and a 1.2 degree beam, so ten milliseconds of illumination each turn. Against a 500 MHz survey at 300 kHz resolution sweeping every 150 milliseconds, the dwell in the affected cell is 90 microseconds and the single-sweep probability works out at three and four tenths parts in a thousand.

> **THE RESULT THAT SHOULD CHANGE THE PROCEDURE**
>
> A one minute survey catches that radar three times in four, and ninety-five percent confidence arrives at two and a quarter minutes. The radar is the easy case: it fires for ten milliseconds in every three seconds. An arcing connector firing for forty milliseconds every thirty seconds needs five and a half minutes for the same ninety-five percent, and a five millisecond burst arriving once a minute needs an hour and a half. Between those three the duty cycle moves by a factor of forty and the survey duration moves with it, so booking a site visit by the hour and hoping is not a procedure.

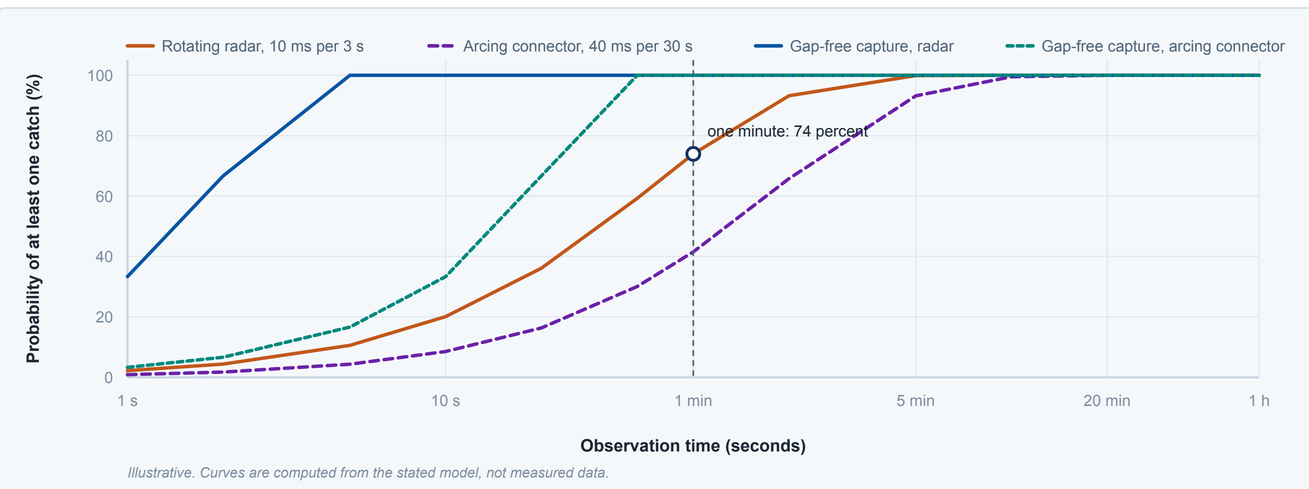

Figure 3. Detection confidence rises exponentially with observation time, at a rate set by the interferer's duty cycle, so the time to any chosen confidence scales inversely with that duty cycle. A gap-free receiver does not escape this. It removes the receiver's own blindness, leaving only the emitter's repetition interval to wait out, which is why its two curves rise on the emitter's own timescale rather than jumping to certainty. DERIVED

Figure 3 makes the case for engineering the survey rather than scheduling it. Two honesty clauses belong with it. The derivation assumes independence between sweeps, and if the revisit happens to be commensurate with the interferer's period the sweep can alias and systematically never see it. And max-hold across many sweeps does raise the catch probability exactly as the formula says, but it destroys the timing information: max-hold tells you that something happened, never when, and cannot separate two intermittent sources.

Sections 1 to 3 are a criterion. Table 3 states it as four requirements a reader can score any instrument against, and they are deliberately unglamorous: none of them is a sensitivity figure, because sensitivity is not what fails a ground-segment measurement.

Table 3. Four requirements, and the second is the one that decides the purchase. Ku and Ka both sit above 9.5 GHz, so a station working either carries the looser amplitude figure and its commissioning file has to say so. R3 is the row to read twice: the requirement is on the station's reference, and what the instrument supplies is a timebase good enough to judge it. R1 and R4 are procedure rather than hardware, which is why they are the two most often skipped.

#	Requirement	Derived value	Published
R1	Continuous observation matched to the suspected duty cycle	5.6 min for a 40 ms burst every 30 s at 95 percent, from Eq. 3	Max hold and real-time density; no time limit published
R2	Bounded amplitude accuracy across the band actually in use	±3 dB or better through Ka	±2.0 dB to 9.5 GHz; ±3.0 dB above, which is where Ku and Ka sit
R3	An analyzer timebase better than the reference it is asked to judge	20 log(1,305) = 62.3 dB is added between a 10 MHz reference and a 13.05 GHz LO, derived in §4	−107.5 dBc/Hz at 1 GHz, 10 kHz offset; TCXO < 1 ppm, OCXO option < 0.15 ppm; external reference input
R4	A recorded baseline that outlives the engineer who took it	state, trace and conditions together, not a screenshot	IQ record and playback, amplitude correction, offset tables

4 Line-up, and the measurements that go in the file

Commissioning is a sequence, and each step produces a record that has to survive being read by somebody hostile three years later.

1. **Pre-build site survey.** Sweep the full receive band and the adjacent segments at the planned look angles and around the horizon, with a calibrated antenna. Log azimuth-resolved spectra, and run a long monitored capture. A clean ten minute sweep proves nothing about a radar with a three second rotation and a shipping lane twenty kilometers away.
2. **Alignment.** Peak in both axes and measure the peaking curve. A flat-topped or asymmetric response indicates feed defocus or reflector distortion. Confirm you are on the right satellite by matching the observed carrier plan against the operator's, not by signal strength: the two-degree-off error is the origin of a large share of adjacent-satellite cases.
3. **Cross-polarization nulling.** Under the operator's direction, transmit a clean carrier at low power while they measure co-polarized and cross-polarized levels, and rotate the feed to null. Operators typically require at least thirty decibels of discrimination. Re-check after final mechanical torque-down.
4. **Line-up.** Step to the assigned radiated power under supervision, check linearity, then bring modulation up and check spectral shape, occupied bandwidth and shoulder levels against the operator's mask.
5. **Archive.** Export the SpectraCore records: full-band receive spectrum on both polarizations, per-carrier levels, the transmit spectrum at the coupler, the cross-polarization result, the figure of merit, the cable sweeps, the pointing readouts, and the instrument model, serial and calibration date, timestamped.

WHY THE BASELINE IS THE DELIVERABLE

Interference claims arrive months later as an assertion that your station is degrading a transponder. The commissioning file is the only timestamped, instrument-traceable evidence of the station's compliant state. Measured discrimination at line-up rebuts a depolarization accusation. The archived transmit spectrum against the mask rebuts a regrowth accusation. A site survey showing a pre-existing emitter turns your problem into the neighborhood's. Without a baseline the argument is the operator's monitoring data against nothing, and **nothing loses every argument it is ever in.**

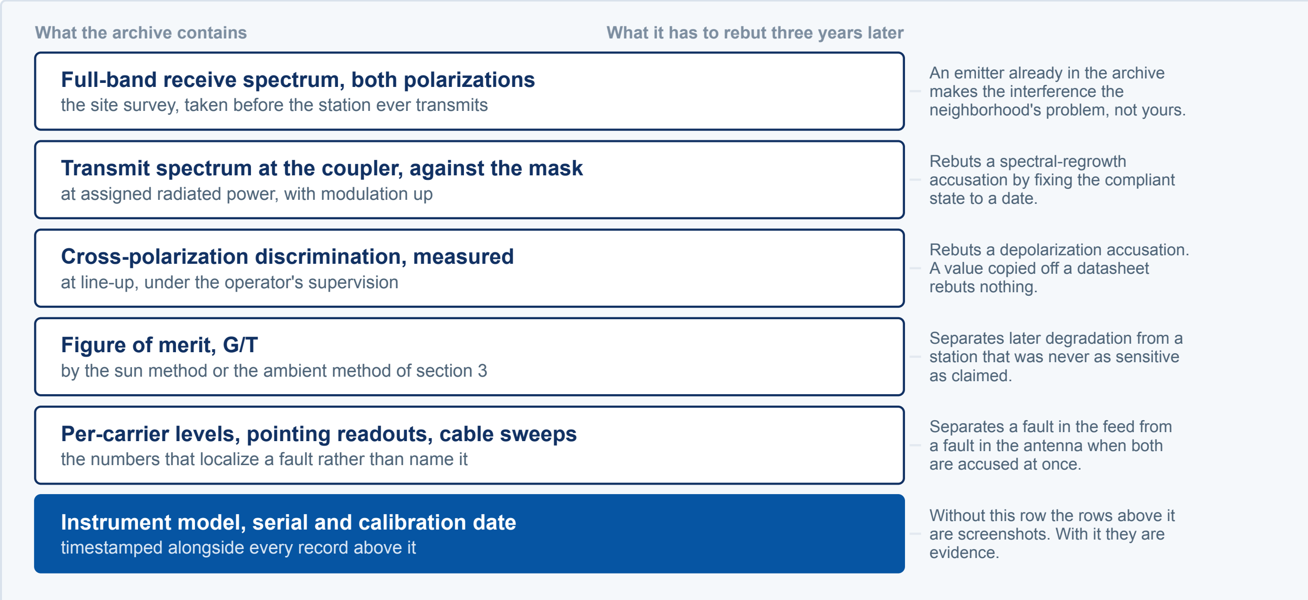

Figure 4. The deliverable is not the visit, it is the file the visit produces, and every row here is a measurement somebody will later claim was never made. Left, what the commissioning export contains; right, the specific accusation each element answers once the van has gone. The bottom row is the one that decides whether any of the others count: traces without the instrument's model, serial and calibration date beside them are screenshots, and a screenshot loses the argument on the day it is produced. SCHEMATIC

Figure 4 is the paper's argument in one picture, and it is worth reading in the right-hand column first. Each accusation in that column arrives as an assertion from a party with its own monitoring data and no obligation to share it. The only reply that changes the outcome is a measurement that already existed before the dispute did, which is why the instrument on site has to be one whose numbers survive being quoted back years later.

Two measurements worth doing properly

Phase noise multiplies. A Ku block upconverter locked to the station's 10 MHz reference multiplies it by 1,305 to reach a 13.05 GHz local oscillator, and phase noise rises by twenty times the log of that ratio, which is 62.3 dB. A mediocre reference that looks harmless at 10 MHz is fatal at 14 GHz, and the symptom is a modem that locks at low order modulation and refuses at high order while every level on the analyzer looks correct.

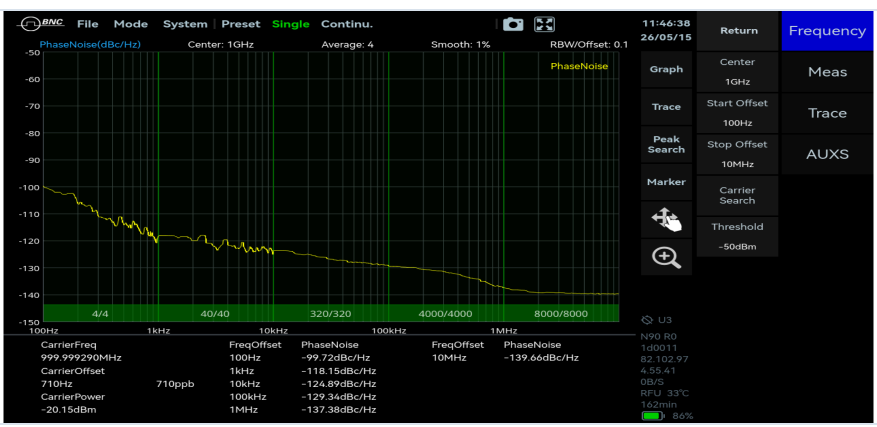

Figure 5. Phase noise of a 1 GHz source measured directly, from 100 Hz to 10 MHz offset: −99.7 dBc/Hz at 100 Hz, −118.2 at 1 kHz, −124.9 at 10 kHz, −129.3 at 100 kHz and −137.4 at 1 MHz, with the carrier itself 710 Hz off nominal, or 710 parts per billion. This is a measurement of the source, not of the analyzer: the instrument's own phase noise is a separate published figure and it is what bounds how good a source this mode can characterize. The published figure for this family at 1 GHz and 10 kHz offset is −107.5 dBc/Hz, which is a limit rather than a typical value, so a unit reading a source below that limit is measuring its own margin as much as the source: any phase-noise claim within about 10 dB of the analyzer's own floor wants a second method. MEASURED

Figure 5 is the measurement to take before the block upconverter is blamed, and it also shows how the penalty scales with where the chain starts. Carried from 1 GHz to 13.05 GHz the multiplication is only 13.05 times, or 22.3 dB, so this source would land at −102.6 dBc/Hz at 10 kHz offset, which a high-order modem will just about live with.

Start instead at the station's 10 MHz reference and the multiplication is 1,305 times, or 62.3 dB, so a reference sitting at −120 dBc/Hz at 10 kHz arrives at −57.7 dBc/Hz and nothing high-order will hold lock. Where the chain starts decides how much of it survives. The 710 parts per billion carrier offset is the second thing to read: at 13.05 GHz that same fractional error is 9.3 kHz of absolute frequency error, and a receiver's acquisition range is a specification with a number in it.

Back-off is the cheapest fix in the chain, and it is not free. Every decibel of back-off is a decibel off the link's carrier-to-noise, so it is cheap only where the margin exists to spend, and one decibel can be the one that drops the modulation order. Third-order regrowth on a nonlinear amplifier rises roughly three decibels for every decibel of drive, so it falls three decibels for every decibel of output back-off. When the shoulders fail the operator's mask, one decibel of back-off usually costs less than anything else on the site.

Figure 6. The spectral emission mask function in SpectraCore: a stored mask template over a 70 MHz span at 1 kHz resolution bandwidth, 10 Hz video bandwidth and peak detector. This is a bench illustration of the function, not a Ku-band earth-station capture. Shoulder levels against the operator's mask are the compliance record a regrowth accusation is answered with, and they are worth capturing at more than one back-off setting so the trade is documented rather than argued. MEASURED

Figure 6 is what the operator's engineer will be shown three years later, so the sweep behind it wants its resolution bandwidth, video bandwidth, detector and reference level recorded alongside it. A mask plot with unstated settings proves nothing about the shoulders it appears to clear.

Cable and waveguide, by scalar sweep

A tracking generator turns an analyzer into a scalar transmission set, which is enough for most of what goes wrong outdoors. Insertion loss should rise smoothly with the square root of frequency; a wet connector adds disproportionate loss at the high end and, more usefully, puts periodic ripple on the trace. That ripple is a tape measure: the distance to the reflection follows from the spacing between successive ripple peaks.

$$\ell = v_f \, c \,/\, (2\,\Delta f) \qquad (4)$$

where

ℓ is the distance to the dominant reflection, in meters

v_f is the cable velocity factor (0.83 for foam dielectric coax)

Δf is the frequency spacing between adjacent ripple peaks on the swept trace, in hertz; a 30 m foam-coax run puts it near 4.1 MHz

Table 4 converts between the quantities a field engineer reads and the loss they imply.

Table 4. Return loss, reflection coefficient, standing wave ratio and the resulting mismatch loss. Twenty decibels or better is a healthy field joint, fourteen is marginal, and below ten something is mechanically wrong.

Return loss	Reflection coefficient	VSWR	Mismatch loss
26 dB	0.050	1.11	0.011 dB
20 dB	0.100	1.22	0.044 dB
14 dB	0.200	1.50	0.18 dB
9.5 dB	0.335	2.01	0.52 dB

One caution before this goes into a procedure. Tracking generator availability varies by model and is not listed on the current handheld or rugged datasheets, so confirm it for the unit you have. Where the analyzer cannot supply the stimulus, an external source stepped in synchronism does the same job with more setup and the same arithmetic.

Field evidence

One pattern is worth reporting from stations already in service. Instruments of this class get specified against the maintenance workload rather than the commissioning one, which inverts what the datasheet comparison usually optimizes for. Commissioning is a scheduled day with the operator on the line and the good reference in the rack; maintenance is an unscheduled afternoon in weather with whatever is in the van, and it is where the recorded baseline of section 4 either exists or does not. The stations that resolve interference quickly are not the ones with the best instrument. They are the ones that recorded what good looked like.

5 Honest limits

- **Scalar is not vector.** A tracking generator gives magnitude only: no phase, no group delay, no complex impedance, no true vector distance to fault. Waveguide qualification and any dispute headed for a formal report want a vector network analyzer.
- **An analyzer does not demodulate the link.** Energy per bit, error rate and adaptive modulation behavior belong to the modem or a link analyzer. The analyzer proves the RF chain; the modem proves the service. Both records belong in the file.
- **Off-site sources can be characterized, not resolved.** A field instrument at the victim site establishes an interferer's frequency, bandwidth, timing signature, polarization and bearing, which is enough to classify it and open a case. Finding a distant or on-orbit source is multi-site geolocation work with the operator and the regulator. The instrument makes the escalation credible; it does not end it.

6 Summary, definitions and further reading

Table 5 collects the governing relations and published values with their conditions.

Table 5. Governing relations and published values with their conditions. The instrument rows are the ICX-FieldHawk published figures, and they are the only two lines in this table that depend on which analyzer is carried up the tower.

Quantity	Relation or value	Condition
Link equation	$EIRP - L_{fs} - L_{atm} + G/T + 228.6$	C/N_0 in dBHz
Free-space loss to GEO	196.2 / 205.7 / 210.2 dB	4 / 12 / 20 GHz at 38,500 km
Figure of merit	$8\pi k(Y - 1)L / \lambda^2 F$	sun method, above 20 degrees elevation
Reference multiplication	$20 \log_{10}(N)$	+62.3 dB for a 13.05 GHz LO from 10 MHz
Regrowth against back-off	about 3 dB per dB	third-order products
Distance to reflection	$v_f c / 2\Delta f$	from swept-trace ripple spacing
Amplitude accuracy	±2.0 dB to 9.5 GHz; ±3.0 dB above	25 °C, 10 min warm-up

Instrument values from the ICX-FieldHawk datasheets. Link and budget figures are worked examples, not a specific operator's numbers.

Quantity	Relation or value	Condition
Phase noise, 1 GHz at 10 kHz	−107.5 dBc/Hz	ICX-400 family

Instrument values from the ICX-FieldHawk datasheets. Link and budget figures are worked examples, not a specific operator's numbers.

For selection, the frequency plan decides it, and the boundary falls lower than most people expect. A C-band-only station is covered by the 9.5 GHz ICX-FieldHawk tier, because its uplink stops at 6.425 GHz. A Ku station does not: its downlink starts at 10.95 GHz and its uplink runs to 14.5 GHz, so Ku and Ka both need the 40 GHz tier. Amplitude accuracy loosens from plus or minus 2.0 dB to plus or minus 3.0 dB above 9.5 GHz, which is inside Ku, so a Ku or Ka commissioning file carries the looser figure and the report should say so.

Cross-polarization discrimination
The ratio of co-polarized to cross-polarized received field from the same transmission. Operators typically require at least 30 dB at line-up.

Figure of merit
Receive antenna gain divided by system noise temperature, in decibels per kelvin. The single number that describes a receive station.

Spectral regrowth
Shoulders raised beside a modulated carrier by amplifier nonlinearity, controlled by output back-off.

Symbols used in this paper.

Symbol	Meaning	Units
Y	Sun to cold-sky noise power ratio	linear
F	Solar flux density	W m^{-2} Hz^{-1}
L_{fs}	Free-space path loss	dB
τ, T_p	Interferer burst length, period	s
Δf	Frequency spacing between adjacent ripple peaks	Hz

Further reading

- Recommendation ITU-R S.524, on maximum permissible off-axis radiated power density, and S.728 for very small aperture terminals.
- Recommendations ITU-R S.465 and S.580, the coordination and design-objective antenna patterns.
- Recommendation ITU-R S.733, determination of the figure of merit for earth stations.
- ETSI EN 301 428 and EN 301 443, harmonized standards for Ku and C band terminals, and 47 CFR Part 25 sections 25.209 and 25.218.
- Recommendations ITU-R P.618 and P.838, for rain attenuation prediction.

If you are writing an acceptance procedure or preparing an interference case and want the survey duration engineered against the interferer you actually suspect, our application engineers would be glad to work through it with you.

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.

- Exact envelope segment values in ITU-R S.524 and S.728, the current revision of S.733, and the C and Ka band envelopes in 47 CFR 25.218 should be confirmed against the source documents before publication.
- Tracking generator availability and frequency coverage on each ICX model requires bench verification; no tracking generator appears on the current handheld or rugged datasheets.
- The spectral emission mask screenshot is a bench illustration: its displayed span and its segment table come from different acquisitions. A Ku-band capture with the carrier and its shoulders visible should replace it before release.
- The maintenance-versus-commissioning pattern in the field evidence section is drawn 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-104 Rev A
August 2026