Satellites of Opportunity as Calibration Sources: Peak-Timing Pointing and Thermal-Node LO Drift

Greg Passmore

Abstract

Low Earth orbit constellations put a predictable Ku-band beacon through a fixed dish beam several times an hour, and that makes them a calibration source available to almost any station. I use those crossings to calibrate a 4.5 m Earth Station dish, with no dedicated beacon, radio star or boresight tower. Each scheduled stare yields one power-transient peak time. For a circular Gaussian beam, each peak time gives one linear constraint, $v_k^\top c + |w_k|, \tau = v_k^\top p_k + |w_k|, t_k$, on the beam center c and a common clock offset τ. The track geometry comes from SGP4 and public CelesTrak supplemental GP elements. From 138 Starlink and OneWeb stares over 35 hours, the least-squares solution places the beam at azimuth 341.608° and elevation 70.309°, with leave-one-out jackknife deviations of 0.035° cross-elevation and 0.021° in elevation, and $\tau = 0.007 \pm 0.026$ s. The same transits also give a 0.589° beamwidth, a Starlink identity test from the OFDM cyclic prefix, and a frequency reference that exposes a sun-driven LNB drift of about −0.14 ppm/°C. The method needs only a fixed dish, a wideband receiver that records time-stamped power, and published ephemerides. It returns pointing, a clock check and oscillator drift from routine operation rather than a dedicated calibration campaign. In my view, it suits any narrow-beam ground station that has to know where it is looking, before it can trust what it is recording.

Keywords: antenna pointing calibration; satellites of opportunity; peak-timing estimation; low Earth orbit; Starlink; OneWeb; Ku band; SGP4; supplemental GP ephemerides; tangent-plane projection; OFDM cyclic prefix; local oscillator drift; LNB thermal model

Introduction

Our lab operates a Ku-band Earth Station built around a 4.5 m dish. At installation, I measured the dish tilt and azimuth, which put the beam at a nominal azimuth of 342.0° and elevation of 70.0°. The dish has a narrow beam, 0.337° at half power for a uniformly illuminated aperture at 11.55 GHz, so an installation reading good to a few tenths of a degree cannot tell whether a recorded satellite transit crossed the beam axis or only its shoulder. Recording transits therefore required a precise dish pose, with an uncertainty attached. To do this, I use the satellites themselves as the calibration source.

Low Earth orbit gives me a dense supply of moving beacons. Starlink, OneWeb and Kuiper satellites cross the beam about six times an hour, each on a predictable path, and each leaves a power transient as it passes. I schedule a stare on each predicted crossing and invert the set of transient peak times for the beam center. Between 22:49 CDT on October 27 and 09:47 CDT on October 29, the station recorded 226 scheduled stares. From 138 of them, the peak-timing estimator places the beam center at azimuth 341.608° and elevation 70.309°, with leave-one-out jackknife standard deviations of 0.035° cross-elevation and 0.021° in elevation, and a common clock offset of 0.007 ± 0.026 s. A quadratic-form beam model halves the along-track residual and moves the center to 341.770° and 70.285°, so the elevation offset of 0.09° to 0.12° above the prior holds under both beam models, while the cross-elevation component depends on the model at the 0.06° level. The same captures show the LNB reference drifting with insolation at an apparent 0.27 ppm per degree of air temperature.

I start with the measurement chain and a single stare, then derive the peak-timing equation that turns one peak time into one linear constraint on the beam center. After that come the least-squares solution and its uncertainty, the residual fan

and the quadratic-form test that explains part of it, and the beam width the profiles imply. The last two sections use raw IQ: first as an identity test on the Starlink waveform, then as a frequency reference that exposes a sun-driven thermal drift in the LNB. Six people stand behind the methods. Todd Humphreys, who leads the Radionavigation Laboratory at the University of Texas at Austin, and his colleagues blindly identified the Starlink Ku-band downlink, including the OFDM cyclic prefix I use as a signature Humphreys and colleagues, 2023. T.S. Kelso runs CelesTrak, whose supplemental GP element sets are fitted to operator-supplied ephemerides rather than to radar and optical tracking Kelso, CelesTrak. David Vallado, then at the Center for Space Standards and Innovation, led the revision of the SGP4 propagator that every one of those element sets assumes Vallado and colleagues, 2006. Bradley Efron, the Stanford statistician, gave the jackknife method its modern treatment alongside the bootstrap Efron, 1982. Gideon Schwarz, at the Hebrew University, derived the large-sample model-selection criterion I use to compare beam and drift models Schwarz, 1978.

The measurement chain

The dish feeds a Ku universal LNB with a 9.75 GHz local oscillator. The stare frequency is 11.55 GHz sky, which lands at an IF of 1.80 GHz on an SDRplay RSPdx sampling 8 MS/s complex at gain setting 40, with no software frequency correction applied. My receiver manager code computes 8192-point power spectra and stamps each frame with the host read time in UTC. The site location comes from a u-blox ZED-F9P antenna, as the median of three daily standalone fixes: with a derived latitude and longitude, and an ellipsoidal height of 299.85 m. A lever arm of $\ell = 30$ m between that antenna and the feed moves the line of sight to a satellite at range ρ by

$$\delta\theta \le \frac{\ell}{\rho} = \frac{30\ \mathrm{m}}{504\ \mathrm{km}} \cdot \frac{180^\circ}{\pi} = 0.0034^\circ,$$

at the median stare range of 504 km. That is a sixth of the smallest uncertainties I report, so the unsurveyed lever arm does not limit the fit.

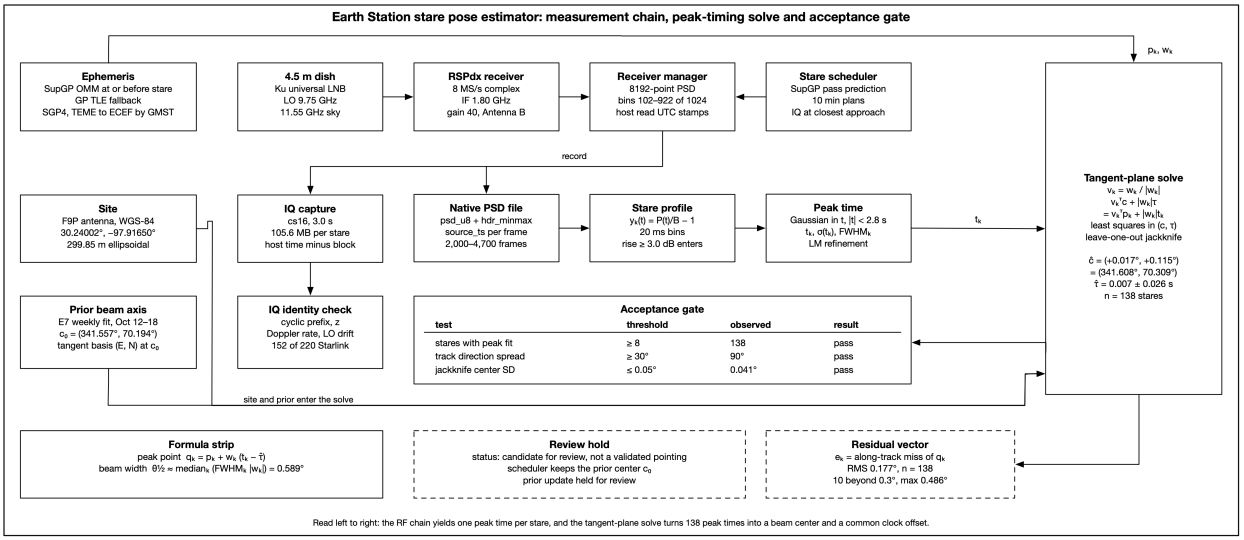


Figure 1. Measurement chain, peak-timing solve and acceptance gate. Imported operators enter on the left, the RF chain produces one peak time per stare, and the tangent-plane solve returns the beam center and a common clock offset.

The stare scheduler propagates the newest SupGP element set at or before each pass with SGP4, plans in 10-minute blocks, and holds the receiver on the stare frequency for about ±5 s around predicted closest approach. Each stare writes a native PSD file of 2,000 to 4,700 frames and, when the plan allows, a 3.0 s IQ capture of 105.6 MB at closest

approach. The prior beam axis c_0 is the E7 weekly fit from October 12 to 18, at azimuth 341.557° and elevation 70.194°.

One stare, one peak time

Figure 2 shows the strongest stare in the set, STARLINK-31853. Power rises across the whole 8 MHz window for about a second, centered 0.19 s after predicted closest approach. White columns are 50 ms bins that hold no frame.

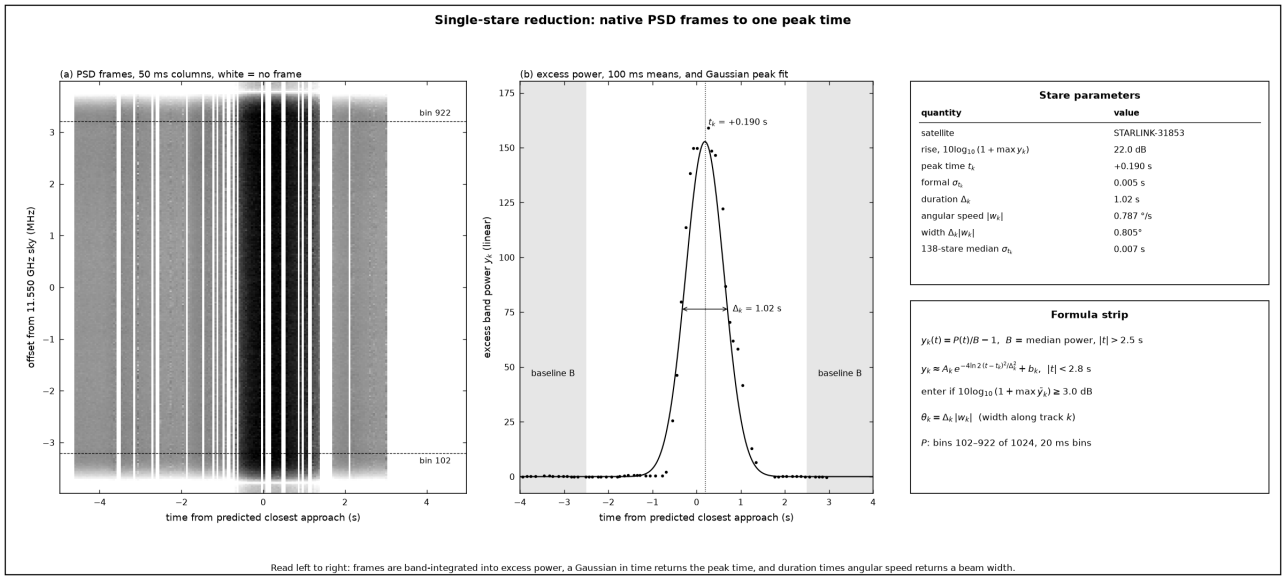


Figure 2. Single-stare reduction. (a) Native PSD frames for STARLINK-31853, with the integration band marked. (b) Excess band power, the Gaussian peak fit, and the duration Δ_k; shaded regions set the baseline.

With $P(t)$ the linear power averaged over bins 102 to 922 in 20 ms bins, and B the median power more than 2.5 s from predicted closest approach, the excess-power profile of stare k is

$$y_k(t) = \frac{P(t)}{B} - 1.$$

A stare enters the pose fit when its smoothed peak rise, $10\log_{10}(1+\max \bar{y}_k)$, reaches 3.0 dB. I fit a Gaussian in time over $|t| < 2.8$ s,

$$y_k(t) \approx A_k \exp\left[-4\ln 2\,\frac{(t-t_k)^2}{\Delta_k^2}\right] + b_k,$$

and keep the peak time t_k, its formal standard error σ_{t_k} and the duration Δ_k. For STARLINK-31853, the rise is 22.0 dB, $t_k = +0.190$ s, $\sigma_{t_k} = 0.005$ s and $\Delta_k = 1.02$ s. The median σ_{t_k} across the 138 fitted stares is 0.007 s, so profile noise does not limit the solution.

Peak timing

A satellite crossing a circular beam is brightest when it is closest to the axis, regardless of how far off axis the track runs. The miss distance scales the peak amplitude and leaves the peak time alone; so the peak time carries the along-track position of the axis.

Formally, I work in the tangent plane at c_0, with E along azimuth and N along elevation, in degrees. Each track is locally straight, $q_k(t) = p_k + w_k t$, with p_k the predicted offset at closest approach, w_k the angular velocity and $v_k = w_k/|w_k|$.

If the station clock runs late by τ, the recorded peak is $t_k = \tau + v_k^{\top}(c - p_k)/|w_k|$, one linear equation per stare,

$$v_k^{\top} c + |w_k|\,\tau = v_k^{\top} p_k + |w_k|\,t_k .$$

The equation assumes a circular beam, and the residuals below show where that assumption struggles.

The geometry behind p_k and w_k is a gnomonic projection. With b_0 the unit vector of c_0 in local ENU, $\hat{e} = b_0 \times \hat{z}/\|b_0 \times \hat{z}\|$ and $\hat{n} = \hat{e} \times b_0$, a line-of-sight unit vector u maps to

$$\Pi(u) = \frac{180^{\circ}}{\pi}\,\frac{\left(u^{\top}\hat{e},\; u^{\top}\hat{n}\right)}{u^{\top} b_0}, \qquad p_k = \Pi(u_k(t_c))\,, \qquad w_k = \Pi\big(u_k(t_c + \tfrac{1}{2})\big) - \Pi\big(u_k(t_c - \tfrac{1}{2})\big)\,,$$

with $u_k(t)$ from SGP4 and t_c the predicted closest approach. The peak-time relation follows from the beam pattern itself. For a circular Gaussian of half-power width $\theta_{1/2}$,

$$G(q) \propto \exp\!\left[-4\ln 2\,\frac{\|q - c\|^2}{\theta_{1/2}^2}\right],$$

$$\frac{d}{dt}\,\|p_k + w_k(t - \tau) - c\|^2 = 2\,w_k^{\top}\big(p_k + w_k(t - \tau) - c\big) = 0,$$

which gives $t_k - \tau = w_k^{\top}(c - p_k)/\|w_k\|^2$, and dividing by $|w_k|$ returns the master equation in degrees.

Each stare constrains the center along its own track direction only, so the solution needs crossing tracks. Stacking the 138 equations (one per transit that entered the fit) gives $A x = b$ with $x = (c_E, c_N, \tau)^{\top}$, rows $a_k = (v_{k,E}, v_{k,N}, |w_k|)$ and $b_k = v_k^{\top} p_k + |w_k| t_k$,

$$\hat{x} = \left(A^{\top} A\right)^{-1} A^{\top} b, \qquad e_k = b_k - a_k^{\top}\hat{x},$$

where e_k is the along-track residual in degrees. At the median angular speed of 0.812°/s, a clock error of 0.1 s moves every peak 0.081° along its track, which is why τ is fitted rather than assumed. The information the tracks carry about the center is set by the direction scatter matrix,

$$M = \frac{1}{n}\sum_{k=1}^{n} v_k v_k^{\top}, \qquad \operatorname{eig}(M) = (0.455,\; 0.545),$$

which is close to the isotropic $(\tfrac{1}{2}, \tfrac{1}{2})$, even though the track directions span only 50° to 167° from east toward north.

Solution and uncertainty

The estimator accepts a fit for review when at least 8 stares carry a peak fit, the track directions span at least 30°, and the jackknife center deviation, $\sqrt{\sigma_E^2 + \sigma_N^2}$, is at most 0.05°. The October 29 run used 138 stares, 125 Starlink and 13 OneWeb. Of the 88 that did not enter, 85 fell below the 3.0 dB rise, 2 native files were not stares, and 1 file was missing. The center deviation is 0.041°.

I estimate uncertainty with the leave-one-out jackknife. With $\hat{x}_{(i)}$ the solution without stare i and $\bar{x}$ the mean of the n such solutions,

$$\hat{\Sigma}_{\mathrm{jk}} = \frac{n-1}{n} \sum_{i=1}^{n} \left(\hat{x}_{(i)} - \bar{x}\right)\left(\hat{x}_{(i)} - \bar{x}\right)^{\top}.$$

The jackknife makes no assumption about the residual distribution, which matters here, because the residuals in Figure 4 are not identically distributed. As a cross-check, the homoscedastic formal covariance $\hat{\sigma}_e^2 (A^{\top}A)^{-1}$ with $\hat{\sigma}_e = 0.179^{\circ}$ gives 0.042°, 0.024° and 0.037 s, against jackknife values of 0.035°, 0.021° and 0.026 s. The two agree to within 30 percent, and the jackknife is the smaller of the pair. The tangent-plane offset maps to azimuth through the Jacobian of the projection at c_0,

$$\Delta az = \frac{c_E}{\cos el_0} = \frac{0.0173^{\circ}}{0.3388} = 0.051^{\circ}, \qquad \sigma_{az} = \frac{0.035^{\circ}}{0.3388} = 0.102^{\circ},$$

so a cross-elevation uncertainty of 0.035° reads as a tenth of a degree in azimuth at 70° elevation.

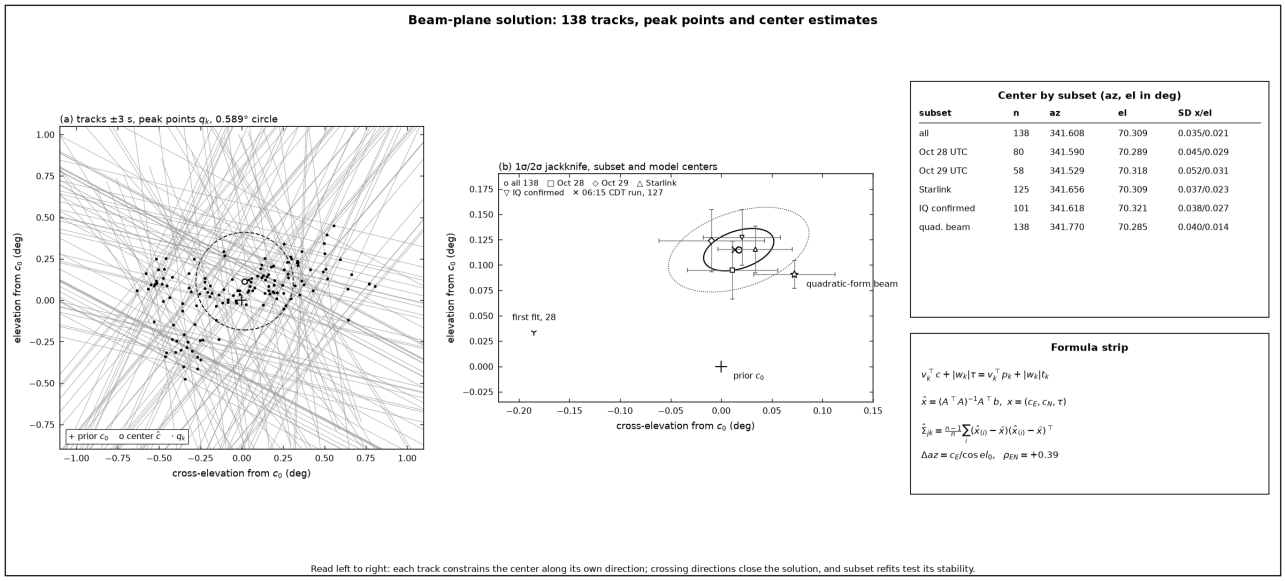


Figure 3. Beam-plane solution. (a) The 138 stare tracks over ±3 s in the tangent plane at the prior, their peak points $q_k = p_k + w_k(t_k - \hat{\tau})$, and the 0.589° half-power circle. (b) Center estimates with 1σ and 2σ jackknife ellipses, four subset refits, the quadratic-form center, the two earlier runs and the prior.

Quantity	Circular beam	Quadratic-form beam
Azimuth	341.608°	341.770°
Elevation	70.309°	70.285°
Offset from prior, cross-elevation	+0.017°	+0.072°
Offset from prior, elevation	+0.115°	+0.091°
Jackknife SD, cross-el / el	0.035° / 0.021°	0.040° / 0.014° (shape fixed)
Clock offset τ	$+0.007 \pm 0.026$ s	$+0.017 \pm 0.022$ s

Along-track RMS	0.177°	0.085°
BIC	−463	−656

Under the circular model, the elevation offset of 0.115° sits more than five standard deviations above the prior, while the cross-elevation offset of 0.017° is half a standard deviation. In beam terms, the prior sits a fifth of a half-power width below the measured axis.

I refit four subsets to check that no slice of the data carries the answer. The two UTC days are independent samples of geometry, ephemeris age and weather. Their centers, 341.590°, 70.289° from 80 stares and 341.529°, 70.318° from 58, differ by 0.021° cross-elevation and 0.029° in elevation, against combined standard deviations of 0.069° and 0.042°. Starlink alone (125 stares) gives 341.656°, 70.309°. The 101 stares whose IQ carries the Starlink signature give 341.618°, 70.321°, a shift of 0.003° cross-elevation and 0.012° in elevation, so mis-associated transients do not drive the result. Every subset puts the elevation between 70.289° and 70.321°.

The residual fan

The along-track residual RMS of 0.177° corresponds to 0.22 s at 0.812°/s, thirty times the median formal peak-time error. Figure 4(a) shows where it comes from. The residual RMS grows with miss distance, from 0.093° for 36 tracks within 0.15° of the axis, through 0.140° and 0.198°, to 0.242° for 38 tracks beyond 0.45°, and its sign splits by track direction. A circular beam predicts no such growth.

A beam whose contours are not circles moves the brightest point of an off-axis track along the track. I model the pattern as a quadratic form,

$$G(q) \propto \exp\!\left[-4\ln 2\,(q-c)^{\top} Q\,(q-c)\right], \qquad Q = R_{\phi}\begin{pmatrix} 1/a^{2} & 0 \\ 0 & 1/b^{2} \end{pmatrix} R_{\phi}^{\top},$$

with R_{ϕ} a rotation through the major-axis angle ϕ, measured from east toward north. Setting the time derivative of the exponent to zero along $q_k(t) = p_k + w_k (t - \tau)$ gives the generalized peak-time equation,

$$(Q w_k)^{\top} c + \left(w_k^{\top} Q w_k\right)\tau = (Q w_k)^{\top} p_k + \left(w_k^{\top} Q w_k\right) t_k,$$

which reduces to the circular case for $Q \propto I$ and stays linear in (c, τ) for fixed shape. Writing $p_k - c = s_k v_k - d_k n_k$ with $n_k \perp v_k$, the peak moves away from the foot of the perpendicular by

$$s_k^{*} = -\,d_k\,\frac{v_k^{\top} Q\, n_k}{v_k^{\top} Q\, v_k},$$

which is proportional to the miss distance d_k and changes sign with the track direction relative to the axes of Q, the same fan Figure 4(a) shows. I scan the axis ratio b/a from 0.50 to 1.00 in steps of 0.02 and ϕ in 3° steps, solve for (c, τ) at each point, and normalize every row to along-track degrees. The minimum lies at $b/a = 0.58$, with the major axis at $15^{\circ} \pm 6^{\circ}$, and the along-track RMS falls from 0.177° to 0.085°. With $n = 138$ and the Schwarz criterion $\mathrm{BIC} = n \ln(\mathrm{RSS}/n) + p \ln n$, two extra parameters lower the BIC from −463 to −656, and the nested F statistic is 222 on (2, 133) degrees of freedom.

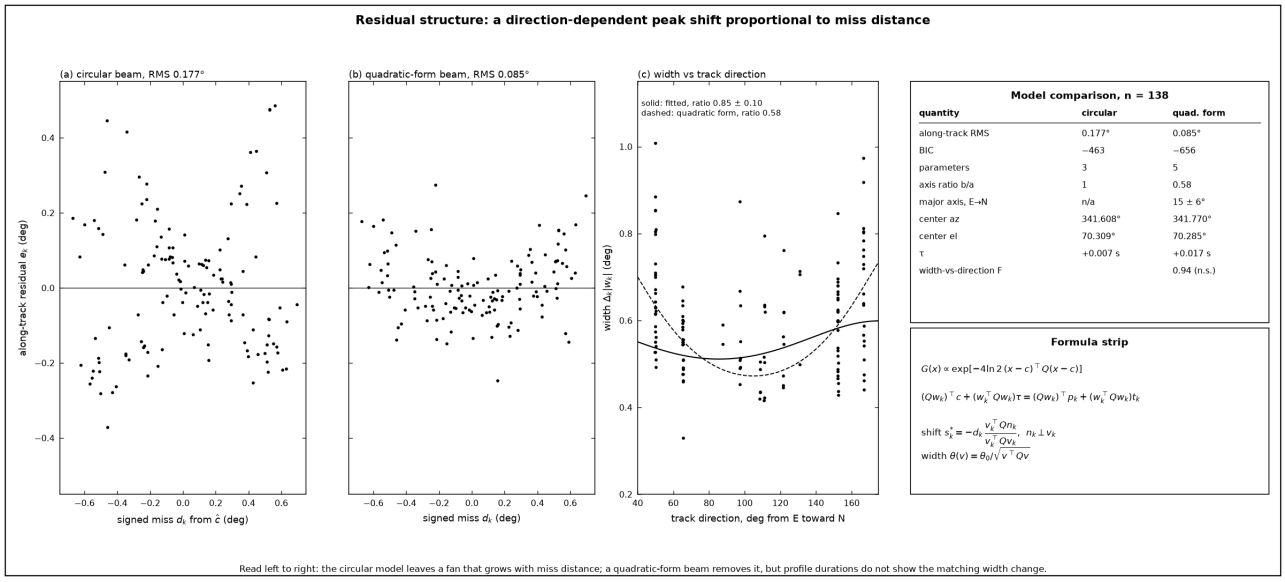


Figure 4. Residual structure. (a) Along-track residual against signed miss distance under the circular beam. (b) The same under the quadratic-form beam. (c) Profile width against track direction, with the fitted second harmonic and the width variation the quadratic form would require.

The profile durations test that shape independently, because they never enter the timing fit. Along direction v, a quadratic-form beam has a half-power width

$$\theta(v) = \frac{\theta_0}{\sqrt{v^\top Q\, v}}, \qquad \frac{1}{\theta_k^2} = \alpha_0 + \alpha_1 \cos 2\psi_k + \alpha_2 \sin 2\psi_k,$$

with ψ_k the track direction, so a ratio of 0.58 should make the widths swing by a factor of 1.72 across the observed directions. A least-squares fit of the second harmonic to the 138 widths gives a ratio of 0.85 ± 0.10, with its narrowest direction at $86^\circ \pm 22^\circ$, against the 105° minor axis the timing fit implies. The harmonic explains 1.4 percent of the width variance ($F = 0.94$ on 2 and 135 degrees of freedom). The widths therefore show no direction dependence at the level the quadratic form needs. In my view, the timing fan comes from an asymmetric, coma-like pattern, whose peak slides along an off-axis cut without the main lobe being elongated, and the quadratic form is the simplest operator that reproduces that slide. Either way, the beam model moves the center by 0.055° in cross-elevation and 0.024° in elevation, and the elevation offset survives it.

Beam width and detection

The duration of a profile, scaled by angular speed, gives a beam width that does not rely on miss distance for a circular Gaussian, because a straight cut through a circular Gaussian is still a Gaussian of the same width. The median over 138 stares is $\theta_{1/2} = \mathrm{median}_k\, \Delta_k|w_k| = 0.589^\circ$, with an interquartile range of 0.506° to 0.667°. With $\lambda = c/f = 25.96$ mm at 11.55 GHz, the beamwidth constant k in $\theta_{1/2} = k\lambda/D$ inverts to an illuminated diameter

$$D_{\mathrm{eff}} = \frac{k\,\lambda}{\theta_{1/2}} = \frac{k \times 25.96\ \mathrm{mm}}{0.589^\circ \times \pi/180^\circ},$$

which gives 2.57 m for uniform illumination ($k = 58.4$) and 3.08 m for a typical edge taper ($k \approx 70$). Either way, it falls well short of 4.5 m, and it seems the universal LNB under-illuminates this dish.

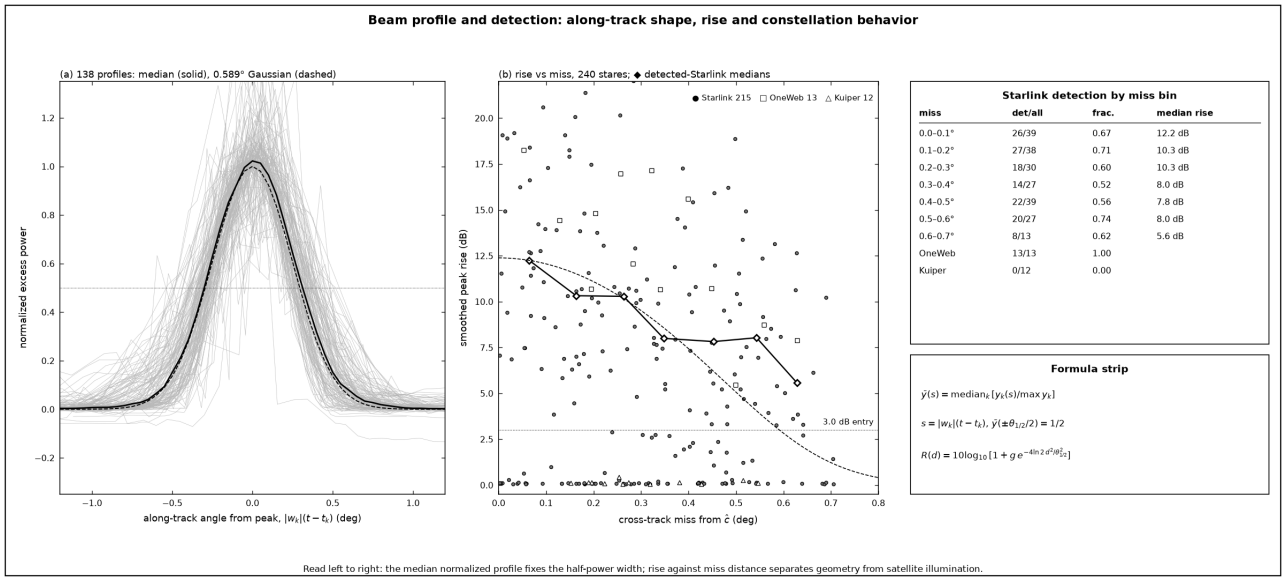


Figure 5. Beam profile and detection. (a) The 138 profiles, normalized to their peaks and placed on the along-track angle from peak, with their median and a 0.589° Gaussian. (b) Smoothed peak rise against cross-track miss distance for all 240 profiled stares, with binned medians of the detected Starlink stares.

Figure 5(a) collapses the profiles onto $s = |w_k|(t - t_k)$. Figure 5(b) separates geometry from satellite behavior. Among the 240 profiled stares, OneWeb produced an excess every time (13 of 13), Kuiper never (0 of 12), and Starlink 135 of 215. The Starlink detection fraction stays between 0.52 and 0.74 in every 0.1° bin out to 0.7° miss distance, so whether a Starlink pass shows up depends mostly on whether the satellite is illuminating this cell, not on where its track crosses the beam. The median rise of detected Starlink stares falls from 12.2 dB within 0.1° to 5.6 dB at 0.6° to 0.7°.

IQ identity

The power transient says something crossed the beam, and the IQ capture says what. Humphreys and colleagues showed that the Starlink Ku-band downlink is OFDM with a subcarrier spacing of $\Delta f = 234.375$ kHz, a useful symbol of $T_u = 1/\Delta f = 4.267$ μs, and a 4.4 µs symbol that includes the cyclic prefix [Humphreys and colleagues, 2023]. The prefix repeats the tail of each symbol, so the lag product

$$\Gamma = \sum_n x[n]\, x^*[n + N_u], \qquad N_u = T_u f_s,$$

carries a coherent term for Starlink and averages toward zero for other waveforms. The analyzer forms a statistic z from Γ over each 3.0 s capture and flags the feature at $z > 6$.

Across 248 captures through 18:13 CDT on October 29, the statistic separates cleanly. All 152 Starlink detections have $z \geq 6.21$, with a median of 31.6. The 68 Starlink non-detections stay at or below 2.97, and all 28 OneWeb and Kuiper captures stay at or below 3.16. Nothing falls between 3.16 and 6.21, so the signature is a clean identity test for the association the scheduler makes from ephemeris alone.

The Doppler rate is a weaker check. With $f_D = -(f_c/c)\,\dot{r}$, the rate is $\dot{f}_D = -(f_c/c)\,\ddot{r}$, and every crossing of a beam at 70° elevation has nearly the same $\ddot{r}$, so the SupGP rates span only −3.96 to −2.98 kHz/s. The per-capture standard error is 1.53 kHz/s over a median in-beam span of 0.93 s. The median residual is +150 Hz/s, the median absolute deviation 1.16 kHz/s, and 26 of 151 detections fall beyond ±4 kHz/s. The rates agree with SupGP on average, and a single pass constrains nothing.

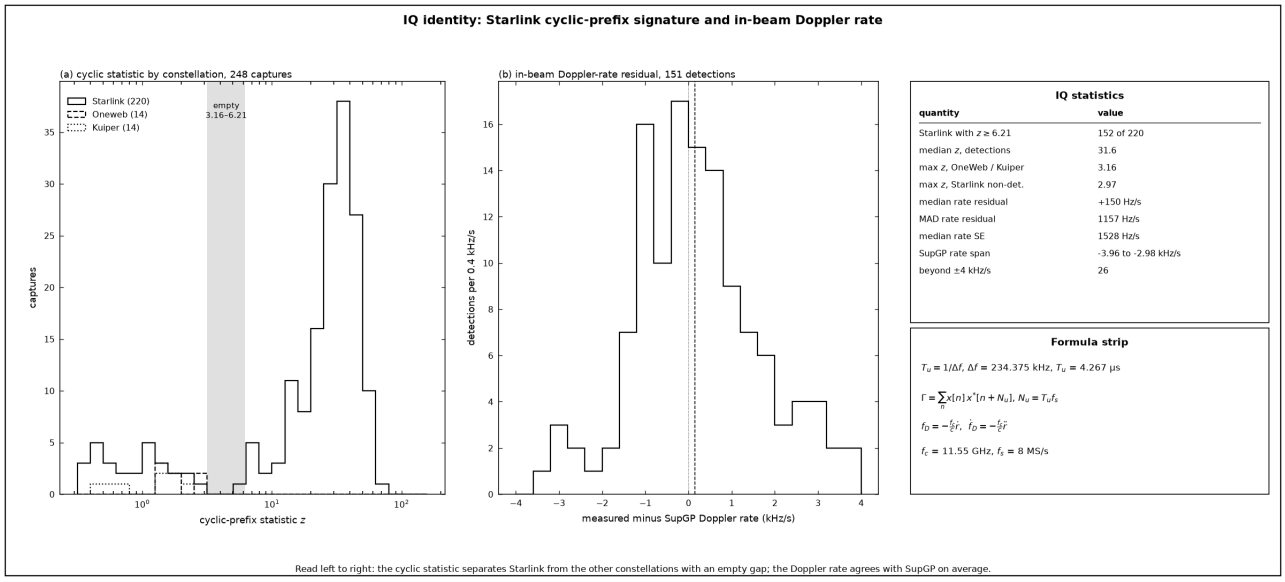


Figure 6. IQ identity. (a) Cyclic-prefix statistic by constellation across 248 captures. (b) Measured in-beam Doppler rate minus the SupGP prediction for 151 detections.

The LNB runs warm

After the predicted Doppler at closest approach is removed, the carrier offset should be a constant. On October 28, a clear day with insolation peaking at 783 W/m², it climbs from about −62 kHz in the early morning to about −34 kHz by 14:30 CDT. On overcast October 29 it stays between −63 and −53 kHz.

Formally, the offset carries both frequency references, $c_0 = -\epsilon_L f_{LO} - \epsilon_R f_{IF} \pmod{\Delta f}$, with fractional errors ϵ_L and ϵ_R, f_{LO} = 9.75 GHz and f_{IF} = 1.80 GHz. For equal fractional drift, the LNB term is 5.4 times larger.

I model the LNB reference as one thermal node heated by air temperature T_a and insolation S,

$$\tau_{th}\,\dot{T}_L = T_a + \kappa S - T_L, \qquad c_0 = a + g\,T_L.$$

The modulus leaves a ambiguous by multiples of Δf, and one node cannot separate housing from oscillator lag.

On October 28, air ran from 17.2 °C to 28.7 °C; on October 29, from 18.5 °C to 21.3 °C, under a 209 W/m² peak. Any transmit offset common to the satellites folds into a.

The offsets come from the 151 cyclic detections, unwrapped about −58 kHz, with 7 points removed that sit more than eight robust standard deviations from a 13-point running median (running MAD 515 Hz), which leaves 144. The drivers come from the WS90 station at one-minute resolution over 47.9 hours, without gaps. I discretize the node on that grid as $T_L[i] = T_L[i-1] + (1 - e^{-\Delta t/\tau_{th}})\,(T_a[i] + \kappa S[i] - T_L[i-1])$, sample it at each capture time, and solve for $(a, g, g\kappa)$ by least squares on a grid of τ_{th}.

The lag correlation in Figure 7(b) shows what air temperature alone cannot. With the driver sampled at $t - L$, the offset correlates best with air temperature at $L = -50$ min ($r = 0.947$), so the offset leads the air by about 50 minutes, and best with low-passed insolation at $L = +30$ min ($r = 0.934$), so it lags the sun. A reference that responds to air alone cannot lead the air, and the ordering sun, then oscillator, then air is what a sun-heated housing produces.

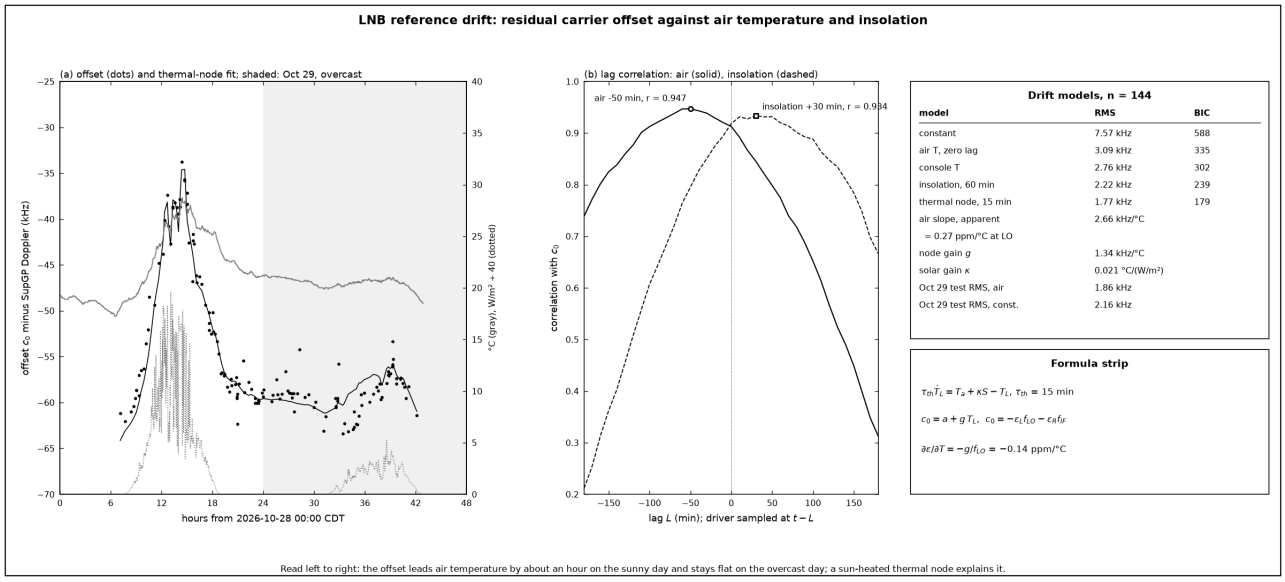


Figure 7. LNB reference drift. (a) Residual carrier offset, the thermal-node fit, air temperature and insolation over two days, with the overcast day shaded. (b) Correlation of the offset with lagged air temperature and lagged insolation.

Model	Parameters	RMS (kHz)	BIC
Constant	1	7.57	588
Air temperature, zero lag	2	3.09	335
Station console temperature	2	2.76	302
Insolation, 60 min low-pass	2	2.22	239
Thermal node, τ_{th} = 15 min	4	1.77	179

The thermal node lowers the BIC by 156 against air temperature alone. Its fit gives $g = 1.34$ kHz/°C and $\kappa = 0.021$ °C per W/m², so 800 W/m² of sun raises the housing about 17 °C above the air. The decomposition between g and κ trades against τ_{th}: from 10 to 20 minutes, the RMS stays within 0.04 kHz of its minimum while g runs from 1.48 to 1.22 kHz/°C. Referred to the LO,

$$\frac{\partial \epsilon_L}{\partial T_L} = -\frac{g}{f_{LO}} = -\frac{1.34 \times 10^3 \ \mathrm{Hz/^\circ C}}{9.75 \times 10^9 \ \mathrm{Hz}} = -0.14 \ \mathrm{ppm/^\circ C},$$

with a range of −0.13 to −0.15 ppm/°C over that band of τ_{th}. The apparent coefficient against air temperature, 2.66 kHz/°C or 0.27 ppm/°C, is twice the node gain, because air-only regression assigns the solar heating to the air. That makes it the right number for predicting the offset from a thermometer reading, and the wrong number for describing the oscillator.

The overcast day is the out-of-sample test. Fitting the air-only model on the 82 captures before 01:00 CDT on October 29 and predicting the 62 after it gives a test RMS of 1.86 kHz with a bias of +0.28 kHz, against 2.16 kHz for a constant. The node model, fitted the same way, gives 2.00 kHz with a bias of −1.37 kHz. October 29 spans only 2.8 °C of air and 209 W/m² of sun, so the test confirms the sign and size of the drift, and it does not discriminate between the drift

models. The node residuals show no dependence on predicted Doppler, Doppler rate or predicted separation, with $|r| \le 0.12$ in every case.

Convergence

The first peak-timing fit, at 07:10 CDT on October 28, used 28 stares and returned azimuth 341.011°, elevation 70.228° and a clock offset of 0.234 s, with a 0.18° jackknife center deviation. In the tangent plane, that point lies 0.22° from the current center, which is 1.2 times its own deviation. The scheduled run at 06:15 CDT on October 29 used 127 stares and gave 341.597°, 70.309°. The manual run at 09:47 CDT used 138 and moved the center by 0.011° in azimuth, or 0.004° cross-elevation, and 0.0003° in elevation. The clock offset fell from 0.234 s to 0.017 s and then 0.007 s as track directions filled in. In my view, the early 0.234 s was the solver absorbing a weakly constrained center component into τ, which is the degeneracy the direction-spread gate exists to catch.

Conclusions

The 4.5 m dish points at azimuth 341.608° and elevation 70.309° under the circular-beam estimator, from 138 Ku-band satellite crossings over 35 hours, with jackknife standard deviations of 0.035° cross-elevation and 0.021° in elevation. The elevation sits 0.09° to 0.12° above the E7 prior under both beam models. The station clock agrees with the ephemeris to within 0.026 s. The beam is 0.589° wide at half power, 1.75 times the uniform-aperture value, and its along-track residuals carry a peak shift proportional to miss distance that a quadratic-form beam halves, while the profile widths show no matching elongation. The Starlink cyclic prefix separates Starlink from OneWeb and Kuiper with an empty gap from 3.16 to 6.21. The residual carrier offset follows a sun-heated thermal node with a 15-minute time constant and a gain of −0.14 ppm/°C at the LO, and it leads the air temperature by 50 minutes. The next estimator carries a shape term into the gated fit, and the next drift test needs a clear day after an overcast one.

The jackknife deviations are statistical. They exclude errors shared by every stare, such as a bias in the SupGP ephemerides, an error in the site coordinates, or the selection imposed by the 3.0 dB entry threshold, which keeps off-axis passes only when they are bright. The same truncation lifts the binned rises in Figure 5(b) at large miss distances, so those medians are not a cross-track beam measurement. The quadratic-form center rests on a shape the duration test does not confirm, so I report it as a bound on model dependence, not as a second pointing. The thermal result rests on one clear day and one overcast day, one weather station and an inferred housing temperature, so g and κ are a fitted pair, not separate measurements. The center is a candidate for review, not a validated pointing, and until it is accepted the scheduler keeps aiming at the prior.

Relevance

For a reader with a dish and a receiver, the useful result is that a precise pointing calibration no longer needs a calibration campaign. Constellations of thousands of satellites now cross a fixed beam several times an hour on paths that public ephemerides predict to within a fraction of a beamwidth. Every crossing a station records during normal operation can therefore double as a pointing measurement. The method asks only for time-stamped power and a least-squares solve, and it reports its own uncertainty through the jackknife, so a station can tell when its pointing estimate is good enough to trust. Recording a few more seconds of IQ at each crossing makes the same transits check the clock to a few hundredths of a second and track the local oscillator against the weather. In my view, that combination matters most for small and independent stations, which rarely have a boresight tower, a strong beacon or time on a radio-astronomy schedule. A narrow beam that is a fifth of a beamwidth off target loses signal on every pass and misattributes what it does see. Ironically, the same satellites that cause those errors are also the ones that can correct them.

References

1. Humphreys, T. E., Iannucci, P. A., Komodromos, Z. M., and Graff, A. M. (2023). Signal structure of the Starlink Ku-band downlink. IEEE Transactions on Aerospace and Electronic Systems, 59(5), 6016–6030. https://arxiv.org/abs/2210.11578
2. Kelso, T. S. CelesTrak Supplemental GP Element Sets. https://celestrak.org/NORAD/elements/supplemental/
3. Vallado, D. A., Crawford, P., Hujsak, R., and Kelso, T. S. (2006). Revisiting Spacetrack Report #3. AIAA/AAS Astrodynamics Specialist Conference, AIAA 2006-6753. https://celestrak.org/publications/AIAA/2006-6753/
4. Efron, B. (1982). The Jackknife, the Bootstrap and Other Resampling Plans. CBMS-NSF Regional Conference Series in Applied Mathematics 38, SIAM. https://epubs.siam.org/doi/book/10.1137/1.9781611970319
5. Schwarz, G. (1978). Estimating the dimension of a model. The Annals of Statistics, 6(2), 461–464. https://doi.org/10.1214/aos/1176344136