What Low-Earth-Orbit Constellations Buy for Jam-Resilient Positioning: Geometry, Spectrum, Carrier Phase, and Antenna Arrays

Krishna R. Pattipati, *Life Fellow, IEEE*

***Abstract*—What positioning accuracy, recovery time, and resistance to jamming can a large constellation of low-Earth-orbit (LEO) navigation satellites deliver, and under what assumptions? Claims of "1 cm" and "unjammable" navigation motivate this question, but we treat them as hypotheses to examine, not conclusions to confirm. We develop a sequence of models that isolate the contributions of satellite geometry, signal strength, interference, antenna arrays, and carrier-phase estimation. LEO's principal advantages over medium-Earth-orbit (MEO) systems are stronger signals and faster geometry change; its snapshot geometry per visible satellite is not intrinsically better. Under the assumed altitude, transmit power, and receiver model, the jamming advantage reaches about 31 dB at 5,000 satellites and approaches a geometric ceiling near 33 dB. Modeling an antenna array by counting its nulls can underestimate protection when jammers outnumber its elements and overestimate protection when a jammer lies close in direction to a desired satellite. In carrier-phase simulations of a flying receiver, 5,000 LEO satellites reached formal 5 cm float precision about 25 s after a cold start and formal 2 cm fixed precision in about 100 s, compared with approximately 14 and 7 min, respectively, for the modeled MEO reference. Integer fixing and achieved accuracy depended strongly on residual errors: with orbit–clock residuals of about 2 cm, achieved 3-D RMS errors over 5–10 min were approximately 1–3 cm; with 5 cm residuals, they were approximately 3–5 cm. These are conditional simulation results, not demonstrated service levels. A companion paper (Part II) uses these results to evaluate navigation architectures and resilient routes for unmanned aircraft.**

***Index Terms*—Low Earth orbit, GNSS jamming, carrier-phase ambiguity resolution, antenna arrays, MVDR beamforming, constellation design, positioning navigation and timing.**

I. Introduction

A. Motivation

Global navigation satellite systems (GNSS), such as GPS, Galileo, GLONASS, and BeiDou, broadcast from medium Earth orbit (MEO), about 20,000 km up. In April 2026, GPS had 32 satellites in its operational constellation [1]; the four systems together have about 130 satellites, of which a receiver typically sees 30–40 [2]. Their signals reach the ground with about 10^{-16} W, below the receiver's own thermal noise [3], [4]. *Jamming*, transmitting noise in the navigation band, therefore denies an ordinary receiver over kilometers with a watt of power; *spoofing*, transmitting counterfeit signals, misleads it silently [5].

Constellations of thousands of satellites in low Earth orbit (LEO, 500–2,000 km) are now proposed as navigation sources [6], [7]; one commercial system has been authorized for 258 navigation satellites [8]. The largest LEO constellation, SpaceX's Starlink broadband system with about 11,000 satellites in orbit [9], broadcasts no navigation signal, but researchers have positioned receivers with its downlink signals to within several meters [10], [11]. Being 10–40 times closer, LEO satellites deliver signals hundreds of times stronger for the same transmitter power, and they cross the sky in minutes rather than hours [12]. Public statements about such systems speak of "1 cm accuracy" and "unjammable" navigation, usually without stating the operating conditions.

B. The question

This paper asks: *what positioning accuracy, recovery time, and interference resilience can a large LEO constellation deliver, and under what assumptions?* The question is deliberately open. A favorable answer, a disappointing one, and a trade-off are all useful, and the "1 cm" claim is a motivation, not a result this paper sets out to confirm. Three distinctions recur throughout. *Predicted precision* is what a filter's covariance says; *achieved accuracy* is the actual error; and *reliability* is how often the solution is badly wrong while appearing precise. A constellation's *size* (satellites in orbit) differs from the number a receiver *sees* (satellites in view): at $42°$N with a $10°$ elevation mask, about 2.4% of the three-inclination 550 km constellations studied here are in view (about 1.3–1.5% at the equator and at $70°$N, where the inclinations place fewer satellites), so the Walker patterns labeled 1,000 and 5,000 (988 and 4,976 satellites) give about 23 and 117 in view on average (Section III-D). A constellation's *signal strength* alone does not determine a receiver's *jamming resistance*, which also depends on the receiver's antenna.

C. Approach

We build the answer using five models, each adding one effect and each chosen to isolate one mechanism.

1) *Snapshot geometry* (Section III): How many satellites in view are needed for a desired accuracy, given the error of each range measurement? This separates the role of satellite count from that of measurement quality.

K. R. Pattipati is with Qualtech Systems, Inc., Rocky Hill, CT 06067 and the Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT 06269 USA (e-mail: krishna@teamqsi.com/krishna.pattipati@uconn.edu).

Working draft (revision 31, Part I of two), November 2026. The models and computations were developed with the assistance of an AI system; all tables are generated directly from the accompanying code and result files, and all results should be independently verified before publication.

2) *Mixing orbital layers and designing a constellation* (Section IV): When do satellites from several orbital layers improve each other's accuracy, and what constellation altitude or altitudes would be optimal if accuracy were the sole criterion?
3) *Jamming margin* (Section V): How much stronger do LEO signals make a receiver against jamming, and why does the gain stop growing? We distinguish a *geometric* limit (the strongest satellites cannot come closer than overhead) from a *spectral* one (the satellites' own signals, sharing the band, act as noise).
4) *Antenna arrays* (Section VI): How much does an adaptive array add, and is the common "count the nulls" model used in architecture studies right?
5) *Carrier-phase convergence* (Section VII): How fast does a flying receiver reach centimeter-level precision and accuracy after a cold start or an outage, and what limits it?

Each model section follows the same pattern: the question; the model and why it is appropriate; what to expect before seeing results; the results; and the mechanism behind them. Section II defines terms and summarizes the systems in service today, including their known weaknesses. Sections III–VII develop progressively more realistic models and the associated design questions. Section VIII collects what is achievable and under which assumptions. Part II [13] uses these results to design architectures and resilient flight policies for unmanned aircraft.

II. Background and Terminology

A. Terms

Positioning. A receiver measures the arrival time of each satellite's signal; when multiplied by the speed of light, this gives a *pseudorange*, the true distance plus the receiver's clock error expressed as a distance. Because the receiver's clock is cheap, its error is unknown and must be estimated together with the three position coordinates, so at least four satellites are needed. Accuracy depends on the error of each range and on the *geometry* of the satellites in the sky: satellites spread over the sky constrain position well, while satellites bunched together estimate the position poorly. An *elevation mask* ignores satellites near the horizon, whose signals are weak and distorted.

Orbits. LEO satellites circle Earth in about 95 min and each is visible from one place for several minutes; MEO (GNSS) satellites stay in view for hours; geostationary (GEO) satellites appear fixed. A *Walker constellation*[1] spreads satellites evenly over several orbital planes of one altitude and inclination, forming a *shell*.

Decibels. A power ratio R is $10\log_{10} R$ dB: 10 dB is a factor of 10, 30 dB a factor of 1,000 and so on. The parameter C/N_0 (dB-Hz) measures signal strength relative to noise. Unjammed GPS signals typically arrive at about 37–45 dB-Hz in a conventional receiver [16]; as C/N_0 falls, range measurements become noisier, and the receiver eventually loses lock. We take 30 dB-Hz as the threshold below which tracking is lost. A receiver's *jamming tolerance* is the largest jammer-to-signal power ratio it survives; "X dB of advantage" means it tolerates a jammer $10^{X/10}$ times stronger. The *equivalent* (also called effective) *isotropically radiated power* (EIRP) is the transmitter power times the antenna gain in the direction of interest.

Antenna arrays. An adaptive array combines the outputs of M antenna elements with complex weights. It can steer a *beam* toward a satellite and place deep *nulls* toward interferers. With M elements, it has $M-1$ spatial degrees of freedom: it can place exact nulls on at most $M-1$ separate directions while keeping its response toward the satellite.

Carrier phase. Precise positioning uses the phase of the received radio carrier, whose wavelength is approximately 19 cm at GPS L1.[2] Carrier-phase measurements have millimeter-level precision, but the receiver cannot directly determine the whole number of cycles between the satellite and receiver when tracking begins. This unknown cycle count is the *integer ambiguity*, which remains constant while tracking continues without a cycle slip. Estimating the ambiguities as real-valued parameters yields the *float solution*; resolving them to their correct integer values yields the more precise *fixed solution*. The LAMBDA method uses decorrelation to make the integer least-squares search efficient [17]; *partial ambiguity resolution* fixes only a subset of ambiguities that meets a specified reliability criterion [18]. Multiple frequencies help separate atmospheric effects from range: the first-order ionospheric contribution scales inversely with frequency squared and has opposite signs in code and carrier-phase measurements. Dual-frequency combinations can therefore remove this contribution, although the integer ambiguities must still be resolved. Orbit and clock errors are common across frequencies, whereas multipath depends on wavelength and the reflecting geometry.

B. Satellite navigation today

Table I summarizes the systems a receiver can use in late 2026; the notes below give the status behind each row, because

[1] Formalized by J. G. Walker of the UK Royal Aircraft Establishment in the 1970s [14], [15]. A Walker pattern, written $i{:}T/P/F$, places T satellites in P equally spaced orbital planes of the same altitude and inclination i, with T/P satellites evenly spaced in each plane; the phasing parameter F ($0 \le F < P$) offsets satellites in adjacent planes by $360°F/T$. The resulting highly symmetric pattern spreads coverage evenly; for small constellations Walker identified patterns that give continuous whole-Earth coverage with very few satellites. Galileo is a Walker $56°{:}24/3/1$ pattern; GPS, with six unevenly phased planes, is not.

[2] Satellite navigation signals occupy a few bands in the L band (1–2 GHz); the wavelength is $\lambda = c/f_c$ for carrier frequency f_c. GPS broadcasts on L1 (1575.42 MHz, $\lambda = 19.0$ cm), L2 (1227.60 MHz, 24.4 cm), and L5 (1176.45 MHz, 25.5 cm). Galileo uses E1 (1575.42 MHz, the same band as GPS L1), E5a (1176.45, the same as L5), E5b (1207.14), and E6 (1278.75); BeiDou uses B1I (1561.098), B1C (1575.42), B2a (1176.45), B2b (1207.14), and B3I (1268.52); GLONASS uses frequency-division signals near 1602 MHz (G1) and 1246 MHz (G2), and newer code-division signals at 1202.025 MHz (G3); QZSS uses L1, L2, L5, and L6 (1278.75 MHz); NavIC uses L5 and S band (2492.028 MHz), with L1 on its second-generation satellites. The systems share center frequencies by design (1575.42 MHz for L1, E1, and B1C; 1176.45 MHz for L5, E5a, and B2a), which lets one receiver front end track several systems but also lets one jammer cover them all, which is why jamming tolerance adds only across separate bands (Section V). Outside the GNSS bands, Iridium transmits at 1616–1626.5 MHz, a band of 5010–5030 MHz is allocated to satellite navigation in C band, and Starlink's broadband downlink is in Ku band (10.7–12.7 GHz).

several of the limitations bear directly on the questions of this paper and of Part II.

The four global systems. GPS had 32 satellites in its operational constellation in April 2026, of which 9 were the newest GPS III; the follow-on GPS IIIF is planned for up to 22 satellites. Its modernized civilian signals are still designated pre-operational: L2C is "use at risk" and L5 is set as unhealthy [1]. Galileo began its Open Service Navigation Message Authentication (OSNMA), which lets a receiver verify that navigation data come from the satellites, as an initial service in July 2025 [19], and broadcasts free High Accuracy Service corrections with a decimeter-level target accuracy [20]. BeiDou-3 has provided global service since 2020; its operator reported completing an in-orbit software upgrade of the constellation in July 2026 [21]. GLONASS is replacing its frequency-division signals with code-division ones on new satellites, a transition expected to take about a decade [22].

Regional systems. Japan's QZSS adds satellites that dwell near the zenith over Japan, valuable in urban canyons, and broadcasts centimeter-level corrections; its expansion to seven satellites, the number needed for standalone positioning, was set back when QZS-5 was lost in an H3 launch failure on 22 December 2025 [23].

India's NavIC (formerly IRNSS) illustrates how satellite failures and delayed replenishment can erode a constellation's operational redundancy. Its original design comprised seven satellites—three in geostationary orbit (GEO) and four in inclined geosynchronous orbit (IGSO)—providing a limited numerical margin above the four satellites required for standalone three-dimensional positioning and receiver-clock estimation. In 2016, all three rubidium clocks aboard IRNSS-1A failed; their supplier also provided clocks that experienced anomalies on Galileo [24]. Replenishment suffered further setbacks: IRNSS-1H was lost in August 2017 when the launcher's heat shield failed to separate, and NVS-02 remained stranded in transfer orbit following a propulsion failure in January 2025. In March 2026, IRNSS-1F's last functioning clock failed shortly after the satellite completed its ten-year design life, leaving three active navigation satellites. The government subsequently stated that standalone positioning would remain unavailable until a fourth satellite became operational, although timing services continued [25], [26]. On 30 October 2026, ISRO launched NVS-03, which, once raised from its transfer orbit to its operating slot and commissioned, would restore the four satellites that standalone positioning requires [27], [28]; the second-generation series is planned to five satellites, with NVS-04 and NVS-05 in preparation, to restore the full constellation [26], [29]. Four satellites would meet the minimum but not restore the margin the original design intended.

This history illustrates three architectural lessons: redundancy must accommodate losses while preserving usable geometry; shared components can introduce common-cause failures, so satellite count alone does not establish independence; and replenishment capacity must account for both on-orbit failures and unsuccessful launches. Multi-constellation receivers and independent navigation sources can mitigate such service degradation. Part II examines how source diversity and fallback capability contribute to mission resilience.

LEO navigation. Iridium's PNT service, broadcast from its communications constellation, is operational; independent tests by the European Commission's Joint Research Centre measured timing within about 100–150 ns of UTC and horizontal positioning of 24–27 m (95%), including indoors [30]. Dedicated navigation constellations are at an early stage: Xona launched its Pulsar-0 demonstrator in June 2025 and has authorization for 258 satellites [8].

Starlink. The largest constellation in orbit is not a navigation system. SpaceX's Starlink had about 11,000 broadband satellites in orbit in October 2026 [9]. Its downlink carries no navigation message, its precise orbits and satellite clocks are not published, and its signal structure is proprietary, so a receiver can use it only as a *signal of opportunity*: it measures the Doppler shift or timing of whatever the satellites transmit, takes their orbits from public two-line element sets, and estimates the unknown satellite clocks along with its own position. With these methods, researchers have positioned stationary receivers to within 6.5 m horizontally from six satellites [10] and 3.6 m horizontally (6.2 m in 3-D) with low-cost antennas and a nearby reference station [11], and have demonstrated navigation on ground vehicles, aircraft, a high-altitude balloon, and a ship [31]. The downlink also contains known synchronization sequences from which pseudoranges can be formed [32]. In 2025, SpaceX informed the U.S. Federal Communications Commission that Starlink can support positioning, navigation, and timing alongside broadband, without proposing a dedicated navigation spectrum [33]. Starlink is therefore both the strongest evidence that large LEO constellations can be built and a reminder that size alone does not make a navigation system: the missing pieces are a navigation signal, a navigation message, and precise orbit and clock products. This paper asks what a large constellation can deliver once it has them; Section VIII relates the results to Starlink.

The threat environment. Interference with satellite navigation now affects civil aviation routinely. International Air Transport Association (IATA) reported that jamming events in 2025 rose 67% from 2023 and spoofing incidents 193% [34], and European Union Aviation Safety Agency's (EASA's) safety bulletin on GNSS outages and alterations was in its fourth revision in 2026 [35]. This is the motivation for the jamming questions of Sections V–VI.

III. How Many Satellites Are Needed for a Desired Accuracy?

Question. If each satellite measures range with error σ, how many must be in view to reach a desired three-dimensional accuracy ε, and how much does their arrangement matter?

A. Model

A receiver at position $\mathbf{x}$ with clock offset b (in meters) observes pseudoranges $y_i = \|\mathbf{s}_i - \mathbf{x}\|_2 + b + v_i$ from satellites at $\mathbf{s}_i$, with independent errors $v_i \sim \mathcal{N}(0, \sigma^2)$. Linearizing about a

TABLE I
SATELLITE NAVIGATION SYSTEMS IN LATE 2026: STRENGTHS AND LIMITATIONS (STATUS FROM THE SOURCES CITED IN THE TEXT).

System	Orbits and status	Civil signals	Strengths	Limitations and known issues
GPS (US)	32 operational MEO satellites (Apr. 2026), 9 of them GPS III	L1 C/A operational; L2C and L5 "use at risk"	Longest record; universal receiver support; modernization under way (GPS IIIF planned)	Modernized civil signals not yet operational (L5 set unhealthy); no civil authentication.
GLONASS (Russia)	About 26 MEO; CDMA satellites (K, K2) replacing FDMA ones	FDMA L1/L2; CDMA on newer satellites	Higher inclination (64.8°) helps at high latitudes	FDMA inter-frequency biases complicate carrier-phase fixing; full CDMA expected only around 2035.
Galileo (EU)	MEO, Walker 56°:$24/3/1$ plus spares	E1, E5a/b, E6	Message authentication (OSNMA, initial service 2025); free High Accuracy Service (decimeter level)	Clock anomalies early in deployment; a week-long service outage in 2019 from a ground-segment failure.
BeiDou-3 (China)	MEO, inclined geosynchronous (IGSO), and GEO; global since 2020	B1I, B1C, B2a, B3I	Mixed orbits give many satellites over Asia; short-message service; PPP-B2b corrections; in-orbit upgrade 2026	Short-message and PPP-B2b services are regional; GEO satellites are low in the sky at high latitudes.
QZSS (Japan)	GEO and quasi-zenith IGSO; regional, 5 operational	L1C/A, L1C, L2C, L5, L6	High elevations over Japan help in cities; centimeter augmentation (CLAS, MADOCA-PPP)	Not standalone until 7 satellites; QZS-5 lost in a 2025 launch failure.
NavIC (India)	GEO and IGSO; regional (India and about 1,500 km beyond)	L5, S band; L1 on second-generation satellites	Independent regional system; GEO/IGSO geometry needs few satellites	Repeated clock and launch failures; 3 satellites active since Mar. 2026, timing only; NVS-03 launched Oct. 2026 to restore the minimum of 4 satellites.
Starlink (SpaceX)	About 11,000 broadband satellites, most lowered from about 550 to 480 km in 2026; no navigation signal	Ku/Ka-band broadband downlink; L-band direct-to-cell	Largest LEO constellation; strong signals; usable for positioning as signals of opportunity	No navigation message; public orbit data only, unknown satellite clocks; meter-level research results; one operator controls access.
LEO PNT (emerging)	Iridium PNT operational; Xona Pulsar demonstrator (2025), 258 authorized	Iridium band; Xona L and C bands	Signals far stronger than GNSS; fast geometry change; Iridium works indoors	Iridium positioning about 25 m (95%); large navigation constellations not yet deployed.

nominal position, a small position change $\delta\mathbf{x}$ and clock change δb change each pseudorange by

$$\delta y_i = -\mathbf{u}_i^\top \delta\mathbf{x} + \delta b + v_i, \tag{1}$$

where $\mathbf{u}_i = (\mathbf{s}_i - \mathbf{x})/\|\mathbf{s}_i - \mathbf{x}\|_2$ is the unit line-of-sight vector from the receiver toward satellite i. Each satellite therefore measures position only along its own line of sight, and every measurement contains the same unknown clock term. The clock matters because a common change in all ranges is indistinguishable from a clock change. Only the *differences* between satellite directions reveal position. The same structure arises in time-difference-of-arrival (TDOA) localization, where differencing arrival times removes the unknown emission time [36]. Formally, estimating the clock and eliminating it leaves the position information

$$J_x = \frac{1}{\sigma^2}\sum_{i=1}^{m}(\mathbf{u}_i - \bar{\mathbf{u}})(\mathbf{u}_i - \bar{\mathbf{u}})^\top, \qquad \bar{\mathbf{u}} = \frac{1}{m}\sum_i \mathbf{u}_i, \tag{2}$$

in which each direction enters only through its deviation from the mean direction $\bar{\mathbf{u}}$. The position covariance is $P_{\mathrm{pos}} = J_x^{-1}$, and we measure accuracy by $\sqrt{\operatorname{tr} P_{\mathrm{pos}}}$, the root-mean-square 3-D error.

To see how accuracy scales with the number of satellites, suppose their directions are scattered uniformly over the sky above the elevation mask e_0 (we write $c_0 = \sin e_0$). Then $J_x \approx (m/\sigma^2)C_u$, where C_u is the covariance of a random direction, and the error is $\sigma\sqrt{G(c_0)/m}$ with geometry factor

$$G(c_0) = \operatorname{tr} C_u^{-1} = \frac{12}{(1-c_0)(2+c_0)} + \frac{12}{(1-c_0)^2}, \tag{3}$$

which is 18 for the full hemisphere and 24.3 for a 10° mask (Appendix A). The second term, the vertical part, dominates: all satellites are above the receiver, so the vertical spread of directions is small, and the clock and the height are hard to separate. With few satellites, the actual geometry is a random sample, and the error is larger than this average suggests. A correction for the randomness of m directions (Appendix A) gives

$$\mathrm{RMS}_{3D} \approx \sigma\sqrt{G(c_0)/(m-5)}, \qquad m \geq 5 + G\sigma^2/\varepsilon^2. \tag{4}$$

B. What to expect

Equation (4) says the error falls only as $1/\sqrt{m}$: halving it takes four times as many satellites. The number of satellites needed therefore grows with the square of σ/ε. Improving each measurement tenfold is worth a hundredfold increase in the number of satellites.

C. Results

What it gives. We checked (4) by Monte Carlo over 200,000 random sky geometries per case. For $6 \leq m \leq 12$ it was 3–7% conservative, and within 2% for $m \geq 24$, at masks of 0°, 10°, and 20° (for example 4.93σ versus a simulated 4.68σ at $m = 6$, 10° mask). The formula's divergence at $m = 5$ is real for the expected squared error. The position is unobservable when all centered directions lie in one plane, that is, when the satellites lie on one circle of the celestial sphere (for example, all at the same elevation, which makes height and clock indistinguishable). Three directions always define such a plane; with five satellites the expectation includes the case

TABLE II
SATELLITES IN VIEW NEEDED FOR A DESIRED 3-D ACCURACY $\sqrt{\mathbb{E}[\operatorname{tr} P_{\text{pos}}]} \le \varepsilon$ (10° MASK, UNIFORM-SKY MODEL; COUNTS BELOW 30 FROM SIMULATION, LARGER ONES FROM (4)). SIX IS THE SMALLEST COUNT WITH A FINITE EXPECTED SQUARED ERROR; BY THE MEDIAN OF $\sqrt{\operatorname{tr} P_{\text{pos}}}$, FIVE SATELLITES WOULD MEET THE 1 M TARGET FOR $\sigma \le 10$ CM.

Range error σ (measurement)	$\varepsilon = 1$ m	$\varepsilon = 10$ cm	$\varepsilon = 1$ cm
1 m (code, single frequency)	29	2,431	243,000
10 cm (code, modern signals)	6	29	2,431
2 cm (carrier phase, LEO-grade orbits)	6	6	103
1 cm (carrier phase, precise orbits)	6	6	29

where the other two lie close to it at perpendicular distances d_4, d_5, and since $\operatorname{tr} J_x^{-1} \ge \sigma^2/(d_4^2 + d_5^2)$, the expectation contains $\iint (d_4^2 + d_5^2)^{-1}\,\mathrm{d}d_4\,\mathrm{d}d_5$ near zero, which diverges logarithmically. The divergence is too slow to show in a Monte Carlo average; with six or more satellites the integral converges. The typical geometry is nonetheless usable: with five satellites the median of $\sqrt{\operatorname{tr} P_{\text{pos}}}$ is 4.3σ, and with four it is 8.3σ. For few satellites, the spread across geometries is large; the 95th percentile is about 1.5 times the RMS at m = 6–8. Table II gives the counts needed for the RMS criterion, from simulation where the approximation is unreliable.

Insight. Measurement quality matters far more than satellite count. Centimeter accuracy from meter-level code ranges would need a quarter of a million satellites in view, more than any constellation provides; carrier-phase ranges good to 1 cm need about 30, which 1,000–2,000 LEO satellites provide. Arrangement matters too: bunching 12 satellites within 20° of zenith multiplies the error about twelvefold, almost all vertically. So any "1 cm" claim is a claim about carrier phase and its error budget (Section VII), not about the number of satellites.

D. From satellites in orbit to satellites in view

The count m in (4) and Table II is the number of satellites in *view*; a constellation is specified by the number N in *orbit*. A satellite at altitude h is in view when it is at least e_0 above the user's horizon, where e_0 is the elevation mask. On the sphere of radius $R_E + h$ on which the satellites move, these positions form a spherical cap (the region within a fixed angle of one point) centered on the point directly above the user, with Earth-central half-angle

$$\psi = \arccos\!\left(\frac{R_E \cos e_0}{R_E + h}\right) - e_0 \tag{5}$$

The cap holds the fraction $f = (1 - \cos\psi)/2$ of the orbital shell.[3] If the satellites were spread uniformly over the shell, Nf would be in view. Inclined orbits are not uniform: a plane of inclination i spends more time near the latitudes $\pm i$, where it turns, and none beyond them. With its nodes spread uniformly in longitude, its satellites have latitude distribution $P_i(l) = \frac{1}{2} + \arcsin(\sin l/\sin i)/\pi$ for $|l| \le i$, and the mean number in view at user latitude φ is

$$\bar{n} = Nf\,g(\varphi), \qquad g(\varphi) = \frac{1}{f}\sum_k \alpha_k \int q_\varphi(l)\,\mathrm{d}P_{i_k}(l), \tag{6}$$

where α_k is the share of satellites in planes of inclination i_k, and $q_\varphi(l) = \arccos\{(\cos\psi - \sin\varphi\sin l)/(\cos\varphi\cos l)\}/\pi$ is the fraction of the latitude circle l inside the user's cap. The latitude factor g is the satellite density over the cap relative to a uniform shell. For the 550 km constellations of this paper (53°, 70°, and 97.6° planes with 58%, 26%, and 16% of the satellites) and $e_0 = 10^\circ$, $\psi = 15.0^\circ$ and $f = 1.70\%$.

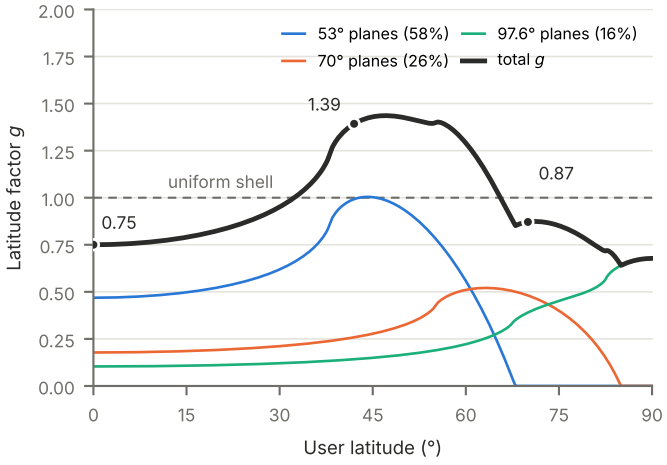


Fig. 1. Latitude factor g of (6) for the three-inclination 550 km constellations (10° mask), with the contribution of each set of planes; the mean number in view is Nfg with $f = 1.70\%$.

What it gives. Fig. 1 shows the latitude factor g as a function of user latitude. The latitude factor was 0.75 at the equator, rose above 1 near 32°, peaked at 1.44 near 47° (within 3% of the peak from about 42° to 56°), and fell to 0.87 at 70°N; the kinks occur where the edge of the user's cap reaches a set of planes' turning latitude, $i \pm \psi$. At 42°N, 2.36% of the constellation is in view, so 988 and 4,976 satellites (the Walker patterns labeled 1,000 and 5,000, with the plane shares those patterns actually realize) give 23.3 and 117.3 in view, the same as the time averages of the simulated constellations over one day (23.3 and 117.3; the instantaneous count varies between 15 and 30, and 111 and 125). Lowering the mask to 5° and 0° raises the share at 42°N to 3.4% and 4.9%.

Insight. Three consequences follow. First, the constellation size a given accuracy requires depends on where the user is: the 30 satellites in view that 1 cm carrier-phase ranges need (Table II) take about 1,300 satellites at 42°N but about 2,400 at the equator and 2,000 at 70°N. Second, the inclination mix is a choice of where to concentrate service: the mix studied here favors mid-latitudes, and a constellation meant for low latitudes would add low-inclination planes. Third, the elevation mask trades satellites in view against measurement quality. With equal, correctly weighted range errors, a lower mask

[3] By Archimedes' hat-box theorem, the area of a band cut from a sphere by two parallel planes equals the area of the same band on the circumscribed cylinder: it depends only on the distance between the planes, not on where the band lies. A cap of angular radius ψ has height $a(1 - \cos\psi)$ on a sphere of radius a, so it holds the fraction $(1 - \cos\psi)/2$ of the sphere's area. The same theorem is why a direction uniform in solid angle has a vertical component uniform on $[c_0, 1]$ (Appendix A).

can only help: at 42°N with the 988-satellite pattern, going from 10° to 5° and 0° raises the mean count from 23 to 33 and 48, lowers the geometry factor (3) from 24.3 to 20.7 and 18.0, and reduces the RMS error of (4) by 26% and 44%. But the lowest measurements are the poorest available. The slant range to a 550 km satellite grows from 550 km at zenith to about 1,800 km at 10° and 2,700 km at 0°, costing about 10 and 14 dB of received power; tropospheric delay grows roughly as $1/\sin e$; multipath and blockage concentrate near the horizon; a LEO satellite spends only seconds at the lowest elevations, which invites cycle slips; and ground jammers seen from the air lie just below the horizon, close in direction to low satellites, so an array protects those satellites the least (Section VI; Part II). The best mask therefore depends on the environment and the threat, which is why conventional values lie between 5° and 15°; the counts and margins in this paper use 10°. Finally, satellites in view are not necessarily satellites used: a receiver tracks a limited number of channels (48 in Section VII), and a broadband downlink serves only the satellites its scheduler assigns (Section VIII).

IV. Mixing Altitudes and Designing a Constellation

A. Several layers

Question. When a receiver uses satellites from several orbital layers (for example GNSS in MEO and a LEO layer), how do they combine, and does it matter whether each layer keeps its own time?

Equation (2) can be split by layer. Let satellite i have range error σ_i and weight $w_i = \sigma_i^{-2}$, and let $\mathcal{L}_k$ be the set of satellites in layer k. Define the total weight of layer k and of all layers,

$$W_k = \sum_{i \in \mathcal{L}_k} w_i, \qquad W = \sum_k W_k,$$

the weighted mean direction of layer k and of all satellites together,

$$\bar{\mathbf{u}}_k = \frac{1}{W_k} \sum_{i \in \mathcal{L}_k} w_i \mathbf{u}_i, \qquad \bar{\mathbf{u}} = \frac{1}{W} \sum_i w_i \mathbf{u}_i = \sum_k \frac{W_k}{W} \bar{\mathbf{u}}_k,$$

and the information within layer k, its directional spread about its own mean,

$$J_k = \sum_{i \in \mathcal{L}_k} w_i (\mathbf{u}_i - \bar{\mathbf{u}}_k)(\mathbf{u}_i - \bar{\mathbf{u}}_k)^\top.$$

With unequal errors, (2) becomes

$$J_x = \sum_i w_i (\mathbf{u}_i - \bar{\mathbf{u}})(\mathbf{u}_i - \bar{\mathbf{u}})^\top,$$

which reduces to (2) when all $\sigma_i = \sigma$. Writing $\mathbf{u}_i - \bar{\mathbf{u}} = (\mathbf{u}_i - \bar{\mathbf{u}}_k) + (\bar{\mathbf{u}}_k - \bar{\mathbf{u}})$ and noting that the cross terms vanish because $\sum_{i \in \mathcal{L}_k} w_i (\mathbf{u}_i - \bar{\mathbf{u}}_k) = 0$ splits it into within-layer and between-layer parts:

$$J_x^A = \sum_k J_k + \sum_k W_k (\bar{\mathbf{u}}_k - \bar{\mathbf{u}})(\bar{\mathbf{u}}_k - \bar{\mathbf{u}})^\top, \qquad J_x^B = \sum_k J_k. \tag{7}$$

Case A holds when all layers share one time scale, so one receiver clock offset applies to all of them. Case B holds when each layer has its own unknown time offset (for example two systems whose time scales are not tied together). Then, each layer's mean direction is absorbed by its own clock, exactly as in Section III, and only the directional spread within each layer remains.

The extra *between-layer* term in case A is information that comes from the layers looking, on average, in different directions. For two layers, it equals $(W_1 W_2 / W)\, \boldsymbol{\Delta}\boldsymbol{\Delta}^\top$ with $\boldsymbol{\Delta} = \bar{\mathbf{u}}_1 - \bar{\mathbf{u}}_2$, so in absolute terms it grows with the weight of the smaller layer, not as that layer shrinks. Its importance is relative. A separately estimated clock removes one degree of freedom from that layer's measurements: it absorbs the component common to all of the layer's ranges, and the position information lost with it is the between-layer term, of rank at most one, which vanishes when the layers' weighted mean directions coincide and grows with $W_1 W_2 / W$ and with $\|\boldsymbol{\Delta}\|$. It is not a fixed amount equal to one satellite's information, and it can vanish even for a sparse layer whose weighted mean direction matches the other layer's. A worked example at 42°N shows its size in two representative geometries. GNSS alone (29 satellites in view) gives an RMS error of 0.95σ. Adding a 100-satellite LEO layer (median 3.5 in view) improves this by 4% if the LEO layer has its own clock and by 8% if it shares the GNSS time scale; the shared time scale doubles the small layer's reduction of the RMS error (not its information matrix). With a 1,000-satellite layer (median 24.5 in view over the 40 sampled epochs), the improvements are 24% and 31%, respectively. In these examples, the recovered term was a large share of the sparse layer's contribution and a small share of the large layer's.

The number of satellites seen at each elevation also depends on altitude. Satellites per unit of sky at elevation e scale as $r(e)^2 / \cos\eta(e)$ (slant range r, angle η at the satellite), so low shells crowd toward the horizon, where they are least useful for the vertical.

What it gives. Per satellite in view, the geometry factor, G is 32.2 for a 550 km shell and 24.1 for MEO.

Insight. Per satellite, LEO geometry is slightly worse than MEO, because its satellites crowd toward the horizon. LEO's principal advantages are numbers, stronger signals, and faster geometry change; its snapshot geometry per visible satellite is not intrinsically better, although more satellites, and the freedom to choose among them, improve the usable geometry. A LEO layer with its own clock loses the information carried by the difference between its mean viewing direction and that of GNSS, a single direction, to estimating that clock. Tying its time to GNSS time, or knowing the offset precisely, recovers it. In the worked examples, the proportional increase in the RMS-error reduction was the largest for a LEO layer with few satellites in view (it doubled, from 4% to 8%, whereas the large layer's rose by more percentage points, from 24% to 31%), because the recovered term was then a large share of what the layer contributed; it need not be in every geometry.

B. Designing for accuracy

Question. Given N satellites, how should they be divided among altitudes and inclinations to give the best accuracy?

The quantity to minimize is the position covariance seen by users. For a candidate shell $c = (h, i)$ (altitude h, inclination i) and a user at latitude φ, let $\mathcal{J}_c(\varphi) \in \mathbb{R}^{4\times4}$ (calligraphic, to distinguish this per-satellite, orbit-averaged information from the summed matrices J_x and J_k) be the expected information that one satellite of that shell contributes to the four unknowns (three coordinates and the clock): the average over the satellite's orbital positions of $w\,\mathbf{h}\mathbf{h}^\top$ with $\mathbf{h} = [-\mathbf{u}^\top, 1]^\top$, zero when the satellite is below the mask, and w the inverse range-error variance. Information from independent satellites adds, so with $N_c \geq 0$ satellites in shell c the user's expected information is

$$M_\varphi(\mathbf{N}) = \sum_c N_c\, \mathcal{J}_c(\varphi) \in \mathbb{R}^{4\times4}. \tag{8}$$

The clock is eliminated only after the information of all shells has been added, because all satellites share the user's clock; eliminating it separately for each shell and adding the results would discard the between-layer term of (7). With $E = [I_3;\, \mathbf{0}^\top] \in \mathbb{R}^{4\times3}$ selecting the position block, the position covariance is $E^\top M_\varphi^{-1} E$, and the design problem is

$$\min_{\mathbf{N}} \max_{\varphi} \operatorname{tr}\big(E^\top M_\varphi(\mathbf{N})^{-1} E\big) \quad \text{s.t.} \quad \sum_c N_c = N,\ N_c \geq 0 \tag{9}$$

(or a population-weighted average over φ instead of the maximum). Because $\operatorname{tr}(E^\top M^{-1} E)$ is convex in M and M is linear in $\mathbf{N}$, this is a convex optimal-design problem [37], [38]: the optimum can be found reliably, and the optimal *fractions* N_c/N do not depend on N, because the objective scales as $1/N$. The counts are continuous in the optimization and are realized as Walker shells afterwards.

Altitude trades two effects. Higher satellites see more of the Earth, so each is in view of more users; but they are farther away, so their signals are weaker and their ranges noisier. We split each range error as $\sigma^2(r) = \sigma_{\text{sys}}^2 + \kappa r^2/\text{EIRP}$: a part independent of distance (orbit, clock, atmosphere) and a part from thermal noise that grows with the square of the slant range r. The two are equal at the crossover range $r_c = \sigma_{\text{sys}}\sqrt{\text{EIRP}/\kappa}$. Beyond r_c, signal weakness dominates, so we expect the optimal altitude to rise with r_c.

What it gives. Table III gives the optimal altitude. The optimum is flat: for each r_c, the best design confined to any single altitude within the band of the table was within 0.4% of the best mixed design, so the split among neighboring altitudes depends on the altitude grid and carries no meaning (for $r_c = 3{,}000$ km and $N = 1{,}000$, the optimizer put 58% of satellites at 2,400 km and 37% at 2,700 km; a coarser grid without 2,400 km put 95% at 2,700 km). Because the design uses time-averaged information, we checked the 2,700 km design by placing the satellites in Walker shells and evaluating the actual error at 150 distinct times: the worst-latitude error was 5% larger than predicted, because the inverse of an average information matrix is optimistic relative to the average of the inverses. Splitting the satellites equally between 2,000 and 3,000 km changed the time-averaged worst-latitude error by -1% and its 95th percentile by $+4\%$; wider splits were worse (1,000/5,000 km: $+15\%$, $+20\%$). Mixing inclinations improved worst-latitude accuracy by 6–14%.

TABLE III
ACCURACY-OPTIMAL ALTITUDE (KM) VERSUS CROSSOVER RANGE r_c (AVERAGED-INFORMATION DESIGN). FULL DESIGN: SATELLITE-WEIGHTED MEAN ALTITUDE OF THE OPTIMAL ALLOCATION; BELOW IT, THE RANGE OF TESTED SINGLE ALTITUDES WHOSE BEST DESIGN WAS WITHIN 0.4% OF THE OPTIMUM.

r_c (km)	1,000	3,000	10,000	20,000
Full design	1,150	2,540	6,080	10,540
within 0.4%	1,150–1,200	2,400–2,700	6,000–6,500	10,000–10,800
Efficiency only	1,020	2,390	6,010	10,070

Insight. Convexity guarantees a reliable optimum but not a single altitude: two shells that are informative in different directions are better mixed, as inclinations are. In this model, the optimum nonetheless concentrated in a narrow altitude band, set by the dominant error, within which the accuracy varied by less than 0.4%, and a design in that band survived the time-resolved check to within a few percent. Because r_c grows with the square root of transmitter power, cheap low-power satellites belong at low altitudes. Accuracy alone does not explain building in LEO at a few hundred kilometers; the reasons are jamming margin and convergence speed, which are examined next.

V. When Do More Satellites Stop Helping Against Jamming?

Question. As a LEO constellation grows, how much jamming margin does it add over GNSS, and what eventually limits it: geometry, or interference among the satellites' own signals?

A. Model

Let C_i be the received power (W) of LEO satellite i at the receiver's antenna output: its EIRP, reduced by the free-space loss over its slant range and multiplied by the receive-antenna gain at its elevation, for every satellite in view above the mask. A receiver needs at least $k = 4$ satellites to position itself, so its tolerance to a wideband jammer is set by the *k-th strongest* signal it receives, $C_{(k)}$. It keeps tracking that satellite while its carrier-to-noise density stays above a threshold $\gamma_{\min}$. A jammer of received power J spread over the signal bandwidth B_s, and the other satellites' signals sharing the band, act as additional noise density, so tracking requires

$$\frac{C_{(k)}}{N_0 + J/B_s + \kappa \sum_{i\neq(k)} C_i/B_s} \;\geq\; \gamma_{\min},$$

and the largest tolerable received jammer power is

$$J_{\max} = \frac{B_s}{\gamma_{\min}} C_{(k)} - N_0 B_s - \kappa \sum_{i\neq(k)} C_i. \tag{10}$$

The first term is the jamming budget the weakest needed signal can withstand ($\gamma_{\min} = 30$ dB-Hz is 1,000 Hz, so the term is a power); the second removes thermal noise (N_0); the third

removes interference from the other satellites' own signals, which share the band. Here, κ is their spectral coupling: $\kappa = 1$ for full overlap is an upper bound, and waveform design, Doppler separation, and spatial filtering reduce κ.

B. What to expect

Adding satellites helps in two ways and could hurt in one. More satellites make it more likely that four of them are in view and close to the receiver, so $C_{(4)}$ rises. But no satellite can be closer than directly overhead: at 550 km, a satellite at zenith is as strong as this constellation can deliver. Once the four strongest satellites are near zenith, adding more cannot raise $C_{(4)}$ further. Meanwhile the self-interference term grows with every satellite in view. Whichever limit is reached first decides whether the ceiling is geometric or spectral.

We evaluated (10) for three-inclination Walker constellations at 550 km (EIRP 24 dBW, $B_s = 20$ MHz, $\gamma_{\min} = 30$ dB-Hz) at 40 times of day and three latitudes, with and without the self-interference term. The reference is a 102-satellite GNSS constellation at its specified minimum received power (about -157 dBW for high satellites).

C. Results

What it gives. Table IV and Fig. 2 show the results. Each row of Table IV is a constellation of N satellites at 550 km; "In view" is the median number above the mask at $42°$N. The two advantage columns give how much stronger a jammer the LEO receiver can tolerate than a GNSS-only receiver can, in decibels. The first counts the other LEO satellites' signals as interference, assuming they fully overlap in frequency ("SI", self-interference); the second ignores them. The last column compares the fourth-strongest LEO signal, the one that sets the tolerance, with a satellite directly overhead; zero means the four strongest are effectively at zenith. The table shows three regimes. *Small constellations* (250–1,000 satellites; median 5–23.5 in view at $42°$N): each doubling adds about 3 dB, from 19.6 dB at 250 to 25.6 dB at 1,000. With few satellites in view, the fourth strongest is often low in the sky and far away (7–13 dB below zenith); doubling the constellation doubles the satellites in view and makes it likely that the fourth strongest is higher and closer. *Diminishing returns* (2,500–10,000; median 59–236.5 in view): the fourth strongest is already within 3 dB of zenith, and a doubling buys only 0.6–1.2 dB (30.1 dB at 2,500, 31.3 at 5,000, 31.9 at 10,000). *Saturation* (20,000 and beyond): the four strongest are within 0.4 dB of zenith, and the advantage stops near 32.9 dB; twenty-five times more satellites (20,000 to 500,000) gain about half a decibel. The two advantage columns agree to 0.1 dB up to 50,000 satellites; only with thousands in view does the sum of the other satellites' powers subtract noticeably from the jamming budget, and even at 500,000 satellites (median 11,805 in view) the loss is 0.3 dB. Fig. 2(a) shows the same pattern at the equator and at 70°N, with small early differences (at 1,000 satellites 23.8, 25.6, and 25.9 dB at 0°, 42°, and 70°), because the three-inclination pattern places more satellites over middle and high latitudes; by 20,000 satellites, all three are within 0.5 dB of the ceiling. Fig. 2(b) shows the mechanism directly: the fourth-strongest signal climbs toward zenith power as the constellation grows, at every latitude.

TABLE IV
LEO JAMMING ADVANTAGE OVER GNSS AT $42°$N (MEDIAN OVER THE DAY), WITH AND WITHOUT SELF-INTERFERENCE (SI, FULL SPECTRAL OVERLAP), AND THE 4TH-STRONGEST SIGNAL RELATIVE TO A SATELLITE AT ZENITH.

N	In view	Advantage (dB) with SI	Advantage (dB) without SI	4th strongest − zenith (dB)
250	5	19.6	19.6	−13.3
500	14	22.3	22.3	−10.6
1,000	24	25.6	25.6	−7.4
2,500	59	30.1	30.1	−2.7
5,000	117	31.3	31.3	−1.8
10,000	236	31.9	31.9	−1.0
20,000	473	32.4	32.4	−0.4
50,000	1,180	32.8	32.8	−0.2
100,000	2,362	32.8	32.9	−0.1
200,000	4,724	32.8	32.9	−0.0
500,000	11,805	32.6	32.9	−0.0

Insight. The margin is limited by geometry, not by interference among the satellites' own signals in the shared band. Its ceiling is simply the difference between a LEO satellite at zenith and a GNSS satellite at minimum power, about 33 dB for the assumed 550 km altitude, 24 dBW EIRP, signal, and receiver threshold; it is not a universal property of LEO. A lower orbit or higher power moves the ceiling up (every halving of altitude adds 6 dB). Self-interference is negligible over any plausible constellation size, because the received powers are spread (satellites low in the sky, near the horizon, are weak, since they are far away, about 1,800 km at $10°$ elevation against 550 km overhead, and the antenna has little gain toward the horizon) and the jamming budget $B_s C_{(k)}/\gamma_{\min}$ dwarfs the sum of satellite powers; a model that gives every satellite the power of the k-th strongest overstates it. Spectrum becomes the lever only through bandwidth: tolerances add across separate bands (Appendix A).

VI. Jamming With Adaptive Antenna Arrays

A. Denial radius

The signal advantage of Section V is easiest to interpret as a distance. A jammer of EIRP $P_J G_J$ denies a source with received power C_s and jamming tolerance $(J/S)_s$ within

$$d_s = \frac{\lambda_s}{4\pi}\sqrt{\frac{P_J G_J}{C_s (J/S)_s}}, \tag{11}$$

capped by the radio horizon (84 km for an aircraft at 300 m). Here C_s is the received signal power and the jammer is received through the same antenna gain, so the free-space factor $(\lambda_s/4\pi d)^2$ uses the source's carrier wavelength λ_s; any discrimination of the antenna or array against the jammer is counted in the protection, not in d_s. For a source of fixed EIRP whose jammer shares its band, the wavelength cancels: C_s itself scales as λ_s^2, so moving a signal from L to C band at the same EIRP and tolerance leaves d_s unchanged. Because power falls with the square of distance, every 20 dB of signal strength or receiver protection shrinks the denial radius tenfold

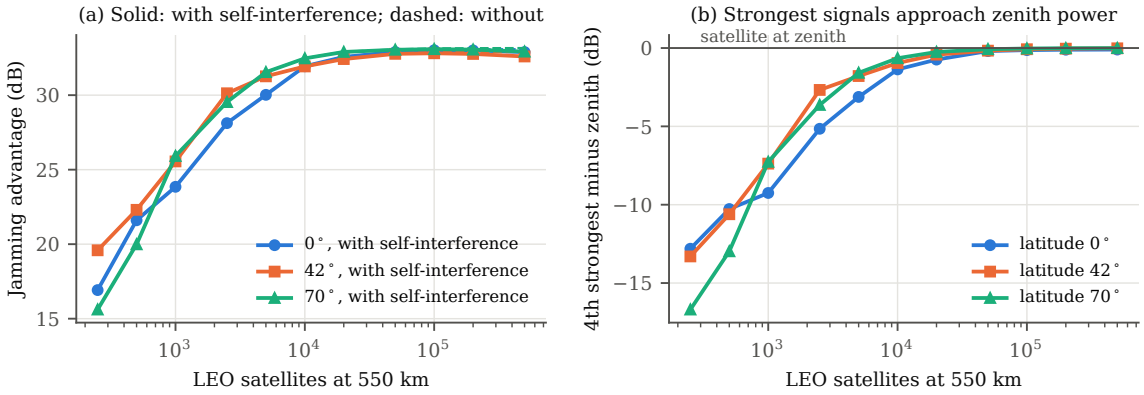


Fig. 2. (a) Jamming advantage over GNSS versus constellation size at three latitudes, with (solid) and without (dashed) self-interference; the curves coincide. (b) The 4th-strongest LEO signal relative to a satellite at zenith: the saturation is geometric.

TABLE V
DENIAL RADIUS (KM) OF ONE JAMMER VERSUS EIRP AND RECEIVER PROTECTION RELATIVE TO A SINGLE-ANTENNA GPS RECEIVER (AIRCRAFT AT 300 M).

Protection	10 W	1 kW	100 kW
0 dB (single-antenna GPS)	32	84 (horizon)	84 (horizon)
30 dB (GPS + array)	1.0	10	84 (horizon)
45 dB (weak LEO layer + array)	0.18	1.8	18
59 dB (LEO layer + array)	0.04	0.36	3.6
69 dB (LEO + 40 dB array)	0.01	0.11	1.1

(Table V). Receiver protection comes largely from the antenna, which is the subject of this section.

B. What an adaptive array optimizes

Question. Architecture studies often model an M-element array by counting nulls: the $M-1$ strongest jammers are suppressed by a fixed null depth, the rest receive no spatial suppression, and the satellite gains $10\log_{10} M$ dB of beam gain. How accurate is this, and when does it fail?

An adaptive array weights its element outputs by a vector $\mathbf{w}$. If R is the covariance of what the elements receive (jammers plus noise) and $\mathbf{a}_s$ the array's response to the satellite (its *steering vector*), the minimum-variance distortionless-response (MVDR) beamformer solves

$$\min_{\mathbf{w}} \ \mathbf{w}^{\mathsf{H}} R\,\mathbf{w} \quad \text{subject to} \quad \mathbf{w}^{\mathsf{H}}\mathbf{a}_s = 1, \quad \mathbf{w} = \frac{R^{-1}\mathbf{a}_s}{\mathbf{a}_s^{\mathsf{H}} R^{-1}\mathbf{a}_s}. \tag{12}$$

In words: keep the satellite's signal unchanged, and make the total output power, which is then interference plus noise, as small as possible [39]. The resulting signal-to-interference-plus-noise ratio is $\mathbf{a}_s^{\mathsf{H}} R^{-1}\mathbf{a}_s$ for unit signal power. Readers from statistics will recognize the form: MVDR maximizes the Rayleigh quotient $|\mathbf{w}^{\mathsf{H}}\mathbf{a}_s|^2/(\mathbf{w}^{\mathsf{H}} R\mathbf{w})$, as Fisher's linear discriminant maximizes the ratio of between-class to within-class scatter, with R in the role of the within-class scatter and the rank-one $\mathbf{a}_s\mathbf{a}_s^{\mathsf{H}}$ in the role of the between-class scatter; the solution $\mathbf{w} \propto R^{-1}\mathbf{a}_s$ matches Fisher's $S_W^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)$. With K satellites to protect, the signal term has rank at most K (exactly K when their steering vectors are linearly independent), as C classes give $C-1$ discriminant directions, and the beamformer that keeps unit gain toward all of them is the linearly constrained minimum-variance (LCMV) beamformer $\mathbf{w} = R^{-1}A_s(A_s^{\mathsf{H}} R^{-1} A_s)^{-1}\mathbf{g}$, with A_s the matrix of their steering vectors. Equivalently, with the snapshot model $\mathbf{y} = A_s\mathbf{s} + \mathbf{n}$, $\mathrm{Cov}(\mathbf{n}) = R$, the output is $\mathbf{w}^{\mathsf{H}}\mathbf{y} = \mathbf{g}^{\mathsf{H}}\hat{\mathbf{s}}$, where $\hat{\mathbf{s}} = (A_s^{\mathsf{H}} R^{-1} A_s)^{-1} A_s^{\mathsf{H}} R^{-1}\mathbf{y}$ is the weighted (generalized) least-squares estimate of the satellite amplitudes: the beamformer is the best linear unbiased estimator of the signals it protects. A receiver does not need that architecture to use several satellites: with digital beamforming, it can form a separate single-constraint MVDR beam for each satellite channel from the same element samples. One multi-constraint beamformer is needed only when a single set of weights must serve several satellites at once, and each added constraint uses a degree of freedom that could otherwise suppress interference. Nothing in (12) counts jammers. What matters is how the interference power is distributed over spatial patterns. Write R in terms of its eigenvalues and eigenvectors: each eigenvector is a pattern of amplitudes across the elements, and its eigenvalue is the interference power arriving in that pattern. The beamformer suppresses the strong patterns and leaves the weak ones. If many jammers arrive from similar directions, their steering vectors are similar, their power concentrates in a few strong patterns, and the array can suppress most of it with few degrees of freedom. Null counting cannot capture this: it treats every jammer beyond the $(M-1)$-th as unsuppressed, however close it is to the others. It also misses the opposite case. When a jammer lies close in angle to the satellite, the constraint $\mathbf{w}^{\mathsf{H}}\mathbf{a}_s = 1$ still passes the satellite's signal unchanged, but keeping that unit response while placing a null a small angle away requires large weights. Large weights amplify the receiver's own noise: the array's *white-noise gain* toward the satellite, $|\mathbf{w}^{\mathsf{H}}\mathbf{a}_s|^2/(\mathbf{w}^{\mathsf{H}}\mathbf{w}) = 1/(\mathbf{w}^{\mathsf{H}}\mathbf{w})$, which equals M for an undisturbed beam, collapses, and so does the output signal-to-interference-plus-noise ratio. The satellite is not attenuated; the array loses its gain toward it. (Actual attenuation of the satellite's signal is a separate mechanism, which occurs when the assumed steering vector is wrong.)

C. Model and what to expect

We compared the null-counting (screening) model with MVDR on planar arrays with half-wavelength spacing: a 7-element hexagon at L1 and 4×4 and 8×8 arrays at C-band. The steering vector toward azimuth ϕ and elevation e has elements $a_m = \exp\{j2\pi \cos e\,(x_m \cos\phi + y_m \sin\phi)\}$ for element positions (x_m, y_m) in wavelengths; all elements share a patch pattern ($+3$ dBi at zenith, about -7 dBi at the horizon, falling below it), applied to both models. Seen from an aircraft at 300 m, ground jammers at 1–80 km lie 0–17° *below* the horizon, uniform in azimuth; the satellite is at 15°–90° elevation; jammer powers span 30 dB; the total jammer-to-signal ratio at one element is 60 dB. Real nulls are not infinitely deep. Two effects limit them: channel decorrelation ζ, which adds ζ times the total jammer power as spatially white noise and limits nulls to about $-10\log_{10}\zeta$ dB, and estimating R from N_s snapshots. We use $\zeta = 10^{-3}$, $N_s = 500$ (about 30 dB nulls) and $\zeta = 10^{-4}$, $N_s = 2{,}000$ (about 40 dB). The MVDR weights are computed from the sample covariance and evaluated against the true covariance, so training and evaluation are independent. The screening model is given the same null depth, so the comparison isolates geometry. For each case, 300 random geometries compare the protection (output SINR relative to one element). We also record the eigenvalues of the jammer covariance and the beamformer's white-noise gain toward the satellite, which is M for an undisturbed beam and drops when the beam is distorted.

From the discussion above we expect: with fewer jammers than elements, the two models should agree in the typical case, with MVDR worse in the occasional geometry where a jammer lies close in direction to a satellite low in the sky (ground jammers appear just below the horizon, so a low-elevation satellite is near them in angle); with more jammers than elements, MVDR should do much better, because ground jammers all lie in a thin band just below the horizon.

D. Results

What it gives. Table VI and Fig. 3 show the results of the analysis. With fewer jammers than elements, the median difference was -4.4 to -0.1 dB, as expected. But with several jammers, the 5th percentile was 9–22 dB worse than the screening model; in these geometries the loss of MVDR white-noise gain toward the satellite (array gain below M) accounts for most of the shortfall. With more jammers than elements, MVDR was 9.5–18 dB *better* in the median with 30 dB nulls and 15–28 dB better with 40 dB nulls. The jammer covariance was not of low rank in the algebraic sense: with 128–192 jammers on the 64-element array, 34–35 of the 64 spatial patterns carried power above the noise-plus-decorrelation floor with 30 dB nulls (39–41 with 40 dB). But the power was concentrated: the *participation ratio* $(\sum\lambda)^2/\sum\lambda^2$, which counts the patterns that carry most of the power, was 13–14 there, and 3.2 with 21 jammers on 7 elements.

Insight. An adaptive array does not run out of nulls the way null counting assumes, because ground jammers seen from the air send most of their power through a few spatial patterns, and the array suppresses patterns, not jammers. The reason is geometric: seen from 300 m, ground jammers lie within 17° below the horizon, a band thinner than the array's elevation resolution, so their directions form a single ring of about $2\pi D/\lambda$ resolvable directions for an aperture D; for the 8×8 array ($D = 4\lambda$) that is about 25, and for a dense field of equal or range-weighted jammers 24–26 of the 64 spatial patterns carried 99% of the power. "More jammers than nulls defeats the array" is therefore not a reliable rule. The opposite error matters as well: null counting misses the geometries in which a jammer lies close in direction to a low-elevation satellite and the array's gain toward that satellite collapses. Architecture studies should use an element-level model, or at least the screening model with a tail correction. This experiment protects one satellite; positioning needs several. Part II therefore uses the MVDR model on its own jammer fields; there, requiring protection toward the fourth-best of eight satellites changed little, while requiring it toward a single satellite at 20° tripled the denied area for small arrays.

VII. Carrier-Phase Convergence and Ambiguity Resolution

Question. After a cold start or an outage, how long does a flying receiver take to reach centimeter-level precision with GNSS versus LEO, on one, two, or three frequencies; how accurate is it actually; and how reliable is the result?

A. How carrier-phase positioning converges

Four stages explain the results that follow.

1) *Phase is precise but ambiguous.* Each satellite's phase measures range to a few millimeters plus an unknown integer number of wavelengths. Until the integers are known, the position is no better than what the pseudoranges (good to decimeters) and the slowly accumulating phase history give.
2) *Changing geometry separates position from ambiguities.* An integer is constant while the satellite is tracked; the satellite's direction is not. As the geometry changes, the same position error produces different range errors at different times, while the ambiguities stay fixed, so the filter can tell them apart. Faster geometry change separates them sooner. This is LEO's structural advantage: its satellites cross the sky in minutes.
3) *Residual errors blur the separation.* Errors left after corrections, in orbit and clock products and in the ionosphere, also enter every range. If they are large, or if they change on the same time scale as the geometry, the filter cannot tell them from position and ambiguity changes, and convergence slows or stops.
4) *Fixing integers sharpens the solution, if it is right.* Once the float ambiguities are precise enough, recognizing them as integers removes most of the remaining uncertainty. A wrong integer, however, biases the position while the filter reports a small covariance. A small formal covariance is therefore not proof of a small actual error, and we report both.

TABLE VI

MVDR PROTECTION MINUS SCREENING-MODEL PROTECTION (DB) AT MATCHED NULL DEPTH, 300 RANDOM GEOMETRIES PER ROW; AND, FOR THE 30 DB CASE, THE MEDIAN NUMBER OF JAMMER-COVARIANCE EIGEN-DIRECTIONS ABOVE THE NOISE-PLUS-DECORRELATION FLOOR, THE PARTICIPATION RATIO $(\sum_i \lambda_i)^2 / \sum_i \lambda_i^2$ OF ITS EIGENVALUES λ_i (THE NUMBER OF SPATIAL PATTERNS THAT CARRY MOST OF THE INTERFERENCE POWER), AND THE MVDR WHITE-NOISE-GAIN LOSS TOWARD THE SATELLITE, $10 \log_{10} \left(1/(M\, \mathbf{w}^{\mathsf{H}}\mathbf{w})\right)$ WITH $\mathbf{w}^{\mathsf{H}}\mathbf{a}_s = 1$ (DB).

		30 dB nulls (matched)		40 dB nulls (matched)				
Array	K	median	5th pct.	median	5th pct.	Dir. above floor	Particip.	Beam loss
L1, 7 el.	1	−0.2	−3.5	−0.1	−2.4	1	1.0	−0.2
	3	−0.9	−11.2	−0.8	−8.7	3	1.4	−0.9
	6	−3.7	−18.5	−4.4	−21.8	5	1.9	−2.9
	7	+0.7	−14.9	+7.5	−11.6	6	2.0	−3.8
	14	+9.5	−6.3	+14.7	−2.9	6	2.7	−6.0
	21	+12.2	−2.4	+15.1	−1.1	6	3.2	−6.9
C, 16 el.	1	−0.2	−1.3	−0.1	−3.2	1	1.0	−0.2
	8	−0.9	−12.6	−1.0	−13.4	7	2.3	−0.7
	15	−2.6	−16.2	−3.6	−19.9	10	3.3	−2.0
	16	+0.0	−13.6	+5.1	−13.3	11	3.5	−1.9
	32	+11.5	−4.9	+19.9	+0.1	12	4.7	−3.3
	48	+14.3	−2.9	+23.6	−1.0	12	5.4	−5.1
C, 64 el.	1	−0.6	−0.8	−0.1	−0.2	1	1.0	−0.6
	32	−1.0	−9.8	−0.7	−11.7	24	6.8	−0.8
	63	−1.2	−15.2	−1.1	−19.4	30	9.6	−0.9
	64	−0.8	−14.4	+2.1	−14.3	30	9.7	−0.9
	128	+13.2	−3.7	+23.2	+2.3	34	12.6	−1.1
	192	+18.3	−0.1	+27.8	+3.7	35	14.3	−1.1

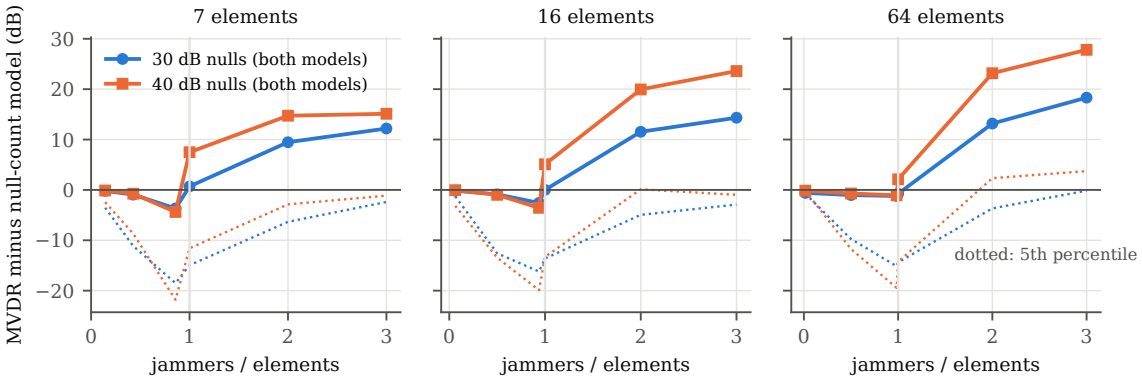


Fig. 3. MVDR minus null-count protection versus the ratio of jammers to elements, for two null depths (solid: median; dotted: 5th percentile).

Other ambiguity estimators. This paper uses the standard approach: LAMBDA decorrelation with partial fixing by integer bootstrapping. Alternatives exist and were not tested. Teunissen's best integer-equivariant (BIE) estimator [40] does not commit to one integer vector; it weights all integer candidates by their probability and is, within its class and under the assumed statistical model, the estimator with minimum mean-squared error; in mean-squared error it is therefore never worse than the float or the fixed solution, although this is neither a guarantee for every realization nor an integrity property. Mixture Kalman filters [41], [42] carry a bank of integer hypotheses through the recursion, so rising and setting satellites and cycle slips are handled in one filter; they have been evaluated with a few ambiguities, and their number of hypotheses grows quickly with the number of ambiguities, so the hundred or so ambiguities of a large LEO layer would need decorrelation or subset fixing first. Ambiguity resolution is also, structurally, the detection problem of multiuser and multiple-antenna (MIMO) communications: an integer vector observed through a known linear map in Gaussian noise. The correspondences are structural analogies, not demonstrated GNSS algorithms. LAMBDA's decorrelation plays the role of lattice reduction, and its integer search that of sphere decoding. Integer bootstrapping is the decision-feedback detector, which is the first solution a depth-first branch-and-bound search reaches; any-time branch-and-bound [43] offers bounded-effort searches between bootstrapping and full integer least squares. BIE corresponds to soft, minimum-mean-squared-error detection, which probabilistic data association (PDA) approximates at cubic cost in its communications formulation [44], [45]; whether that cost and performance carry over to the full GNSS ambiguity problem, with its strongly correlated float ambiguities and changing satellite set, has not been shown. Soft and any-time estimators of this kind may narrow the gap between formal precision and achieved accuracy, and may suit the frequent satellite turnover of LEO; testing them

is future work.

B. Model

For frequency j (wavelength λ_j, ionospheric factor $q_j = (f_1/f_j)^2$), code and phase observations of satellite i at epoch k are

$$\begin{aligned} P_{ijk} &= \rho_{ik} + b_k + m_{ik}T_z + q_j I_{ik} + o_{ik} + \epsilon^P_{ijk}, \\ L_{ijk} &= \rho_{ik} + b_k + m_{ik}T_z - q_j I_{ik} + o_{ik} + \lambda_j N_{ij} + \epsilon^L_{ijk}. \end{aligned} \tag{13}$$

Here, ρ is the geometric range; b_k the receiver clock (new each epoch); T_z the zenith troposphere (a slow random walk), mapped to each satellite by m_{ik}; I_{ik} the slant ionosphere left after regional corrections, which enters code and phase with opposite signs and scales with frequency (Gauss–Markov, standard deviation σ_I, 30-min correlation); o_{ik} the orbit-and-clock residual left after precise products, the same on all frequencies (Gauss–Markov, standard deviation σ_o, correlation time τ_o); N_{ij} the integer ambiguities; and elevation-dependent noise of 3 mm (phase) and 30 cm (code) at zenith. The ionosphere's frequency dependence is why a second frequency helps: comparing frequencies isolates I. The orbit-clock residual does not depend on frequency, so no number of frequencies removes it.

The receiver flies. Its motion between epochs is predicted by an inertial unit with the in-flight error model of Part II: industrial grade, with velocity random walk, a tilt random walk that acts through gravity as an acceleration random walk, accelerometer bias, and a gyroscope-bias-driven acceleration ramp, all inflated by vibration. Position, velocity, and the inertial error states are estimated jointly, and the inertial unit is independent of the satellite measurements. A Kalman filter estimates all states every second from up to 48 satellites, the highest in elevation (a simple rule, not an optimized selection; Section VIII). When a satellite rises, its states are added with their prior variances; when it sets, they are removed by marginalization, so all remaining correlations are kept. The true position starts at a random offset drawn from the filter's prior. At each epoch the float ambiguities are decorrelated by the LAMBDA reduction [17], and the largest subset whose integer-bootstrapping success probability is at least 0.999 is fixed [18]; combinations such as the wide lane are found automatically by the decorrelation. The fixed position follows by conditioning on the fixed subset. Satellites come from Walker shells at 42°N: GNSS (102 satellites in four MEO shells, all assumed to transmit on L1, L2, and L5), and LEO layers of 1,000 and 5,000 satellites at 550 km. For each case 50 start times (independent draws of time of day) are simulated over 30 min.

C. What we measure

We separate three questions. *Predicted precision*: the median time at which the formal error $\sigma_{3D} = \sqrt{\operatorname{tr} P_{\mathrm{pos}}}$ first reaches 5 cm (float) and at which a fixed solution with $\sigma_{3D} \leq 2$ cm first appears. If fewer than half the start times reach it within 30 min, the median is censored and we report the fraction that did. *Achieved accuracy*: the RMS of the actual 3-D error over 5–10 min, and the share of epochs after the first fix whose actual error is within 5 cm. *Reliability of the integer decisions*: over all formally precise fixed epochs, the share with any wrong integer, and the share whose actual error exceeds 5 cm (the *large-error frequency*). The last is not an integrity statement, which would require an alert limit, a risk allocation, and a monitoring rule. Bootstrap 95% intervals of the medians are in the result files; their half-widths are below 35% except for GNSS fix times with a 2 cm ionosphere residual (about ±65%), which vary widely with time of day. The two precision times are not ordered: the 5 cm time belongs to the float estimator's covariance, the 2 cm time to the fixed solution, whose covariance collapses once the integers are fixed. A fixed 2 cm solution can therefore appear before the float covariance reaches 5 cm, as it does for GNSS (about 7 against 14 min).

D. Results

What it gives. Table VII and Fig. 4 show the results of this analysis. With a 10 cm ionosphere residual, single-frequency receivers did not fix within 30 min in any case. With two frequencies, GNSS fixed in a median 400 s (bootstrap 95% interval 322–527 s), LEO 1k in 594 s (541–653 s), and LEO 5k in 104 s (97–106 s); a third frequency shortened these by 2–11%. With a 2 cm ionosphere residual, two frequencies gave 143, 217, and 71 s, and single-frequency LEO receivers also fixed (LEO 1k 550 s, LEO 5k 106 s), while single-frequency GNSS fixed in only 18% of start times. The float solution reached 5 cm in 15–78 s with LEO 5k (25 s, interval 25–26 s, with two frequencies and a 10 cm ionosphere residual), 32–260 s with LEO 1k, and 6–14 min with GNSS (never on one frequency with a 10 cm ionosphere residual). Achieved accuracy tracked formal precision: over 5–10 min, the actual RMS error with two or three frequencies was 1.2–1.4 cm for LEO 5k, 2.2–2.7 cm for LEO 1k, and 1.7–2.5 cm for GNSS. After the first fix the error stayed within 5 cm in 97–100% of epochs; wrong integers occurred in at most 0.4% of fixed epochs and large errors in at most 0.6% (Table VIII).

E. Three comparisons that explain the table

1,000 versus 5,000 LEO satellites (two frequencies, 2 cm orbit-clock residual, 10 cm ionosphere). The float solution reached 5 cm in 56 (50–66) versus 25 (25–26) s, the fix took 594 (541–653) versus 104 (97–106) s (bootstrap 95% intervals), and the achieved error was 2.7 versus 1.4 cm. Both layers change geometry at the same rate. The receiver tracks at most 48 satellites, so the difference is not 117 satellites in use against 23. With 1,000 satellites, it tracks all of the roughly 23 in view; with 5,000, it tracks 48, twice as many, chosen as the highest of about 117 in view, so they also have better elevations and lower noise. Twice as many tracked satellites give twice as many independent views of the same position, and their residual errors, which are independent, average down faster. The tracked set also turns over differently: with 5,000 satellites, a satellite leaves it when it drops out of the highest 48, not only when it sets. Constellation size buys convergence speed and a smaller final error, not a different mechanism.

TABLE VII

CARRIER-PHASE CONVERGENCE AT 42°N (1 HZ, FLYING RECEIVER WITH INDUSTRIAL IMU, COLD START): MEDIAN TIME (S) TO A FORMAL 5 CM FLOAT SOLUTION AND TO A FIXED SOLUTION WITH FORMAL ERROR AT MOST 2 CM, FOR 1 / 2 / 3 FREQUENCIES (L1; L1+L5; L1+L2+L5), OVER 50 START TIMES; IN PARENTHESES THE FRACTION OF START TIMES THAT REACHED IT WITHIN 30 MIN WHEN BELOW 1 ("—": NONE DID; ">1800": FEWER THAN HALF DID; "N/R": NOT RUN). ACHIEVED ERROR: RMS OF THE ACTUAL 3-D ERROR OVER 5–10 MIN WITH TWO FREQUENCIES. RESIDUAL CORRELATION TIME 600 S UNLESS STATED.

Constellation (orbit–clock residual)	Ionosphere residual	Float, 5 cm (s) 1 / 2 / 3 freq.	Fixed, 2 cm (s) 1 / 2 / 3 freq.	Achieved error 5–10 min, 2 freq. (cm)
GNSS MEO (1 cm)	2 cm	509 / 384 / 362	>1800 (0.18) / 143 / 133	2.2
	10 cm	— / 832 / 820	— / 400 (0.98) / 356 (0.98)	2.5
LEO 1k (2 cm)	2 cm	42 / 35 / 32	550 / 217 / 190	2.2
	10 cm	260 / 56 / 56	— / 594 (0.98) / 582	2.7
LEO 5k (2 cm)	2 cm	18 / 16 / 15	106 / 71 / 68	1.2
	10 cm	78 / 25 / 25	— / 104 / 100	1.4
Sensitivity (ionosphere residual 10 cm; 2 / 3 frequencies)				
LEO 1k, orbit–clock 5 cm	10 cm	200	—	5.1
LEO 5k, orbit–clock 5 cm	10 cm	62 / 62	— / —	2.9
GNSS, residual correlation 60 s	10 cm	788 / 781	83 / 73	1.6
LEO 5k, residual correlation 60 s	10 cm	68 / 67	— / —	2.4
LEO 1k, no motion knowledge	10 cm	56	622 (0.96)	2.5
LEO 1k, tactical IMU	10 cm	56	590 (0.98)	2.6
LEO 1k, 2 mm/$\sqrt{s}$ motion aid	10 cm	46	278	1.8
LEO 1k, 5 s epochs	10 cm	90	>1800 (0.10)	3.3
LEO 5k, 5 s epochs	10 cm	40	182	1.7
GNSS, 5 s epochs	10 cm	1595 (0.92)	718 (0.96)	4.7

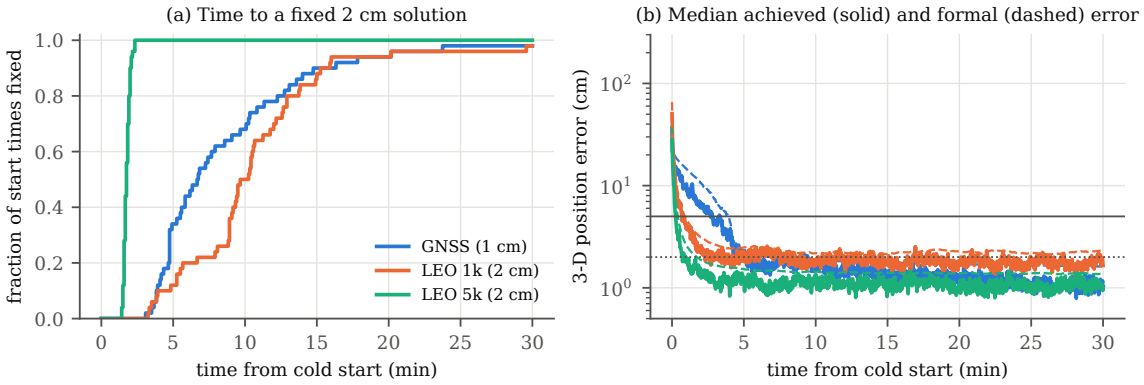


Fig. 4. Convergence across 50 start times (two frequencies, 10 cm ionosphere residual, 1 Hz, industrial IMU): (a) fraction of start times with a fixed 2 cm solution; (b) median actual (solid) and formal (dashed) 3-D error.

TABLE VIII

RELIABILITY OF FIXED SOLUTIONS (2 / 3 FREQUENCIES; 1 HZ, INDUSTRIAL IMU): SHARE OF FORMALLY PRECISE FIXED EPOCHS WITH A WRONG INTEGER; SHARE WITH ACTUAL ERROR ABOVE 5 CM (LARGE-ERROR FREQUENCY); AND SHARE OF EPOCHS AFTER THE FIRST FIX WITH ACTUAL ERROR WITHIN 5 CM.

Case	Wrong int.	Large err.	≤5 cm after fix
GNSS (1 cm), iono. 2 cm	0.00% / 0.00%	0.23% / 0.00%	0.997 / 1.000
GNSS (1 cm), iono. 10 cm	0.00% / 0.00%	0.12% / 0.03%	0.998 / 1.000
LEO 1k (2 cm), iono. 2 cm	0.12% / 0.01%	0.08% / 0.38%	0.992 / 0.986
LEO 1k (2 cm), iono. 10 cm	0.00% / 0.09%	0.60% / 0.23%	0.971 / 0.971
LEO 5k (2 cm), iono. 2 cm	0.26% / 0.01%	0.00% / 0.00%	1.000 / 1.000
LEO 5k (2 cm), iono. 10 cm	0.06% / 0.02%	0.01% / 0.00%	1.000 / 1.000

2 cm versus 5 cm orbit-clock residuals (same receivers). With 5 cm residuals, no LEO case reached a fixed 2 cm solution in 30 min, and the achieved error stayed at 2.9 cm (LEO 5k) and 5.1 cm (LEO 1k). The orbit-clock residual is the same on every frequency and, with a 10-minute correlation, almost constant over a pass, so it looks like an offset of the integer: a 5 cm residual is a quarter of a 19 cm wavelength, which keeps the probability of identifying the integer correctly below the 0.999 threshold. Within the tested grid, precise orbit and clock products were therefore the gate to fixing. Extra frequencies cannot separate a frequency-independent residual; more satellites can reduce its effect by averaging independent residuals, as the 2.9 cm achieved by 5,000 satellites versus 5.1 cm by 1,000 shows, but did not restore fixing. The float solution still delivers a few centimeters.

600 s versus 60 s residual correlation. Here, the two systems react in opposite directions. GNSS fixed faster (83 s instead of 400 s on two frequencies), while LEO 5k, which fixed in 104 s, never fixed. The magnitude of the error was the same; only its time structure changed. A plausible mechanism, not tested by a separate diagnostic, is the following. For GNSS, whose

geometry changes over hours, a residual that decorrelates in a minute averages out over many independent samples, leaving the constant integers easier to see. For LEO, the separation of ambiguities from position relies on geometry that changes over a few minutes; when the residual changes on that same time scale, its variation and the geometric signal are confounded, and the filter cannot attribute the phase history to integers. *We read this as a lesson about the time structure of the errors, which matters as much as their size: integer fixing with LEO, under the tested criterion, needs residuals that are stable over a pass. LEO's float benefit survives: with 60 s correlation, LEO 5k still reached 5 cm float precision in about a minute and an achieved error of 2.4 cm.*

Motion knowledge and measurement rate. With the in-flight inertial model, an industrial or a tactical inertial unit gave the same fix times as no motion knowledge at all (LEO 1k: 594, 590, and 622 s). In-flight tilt and velocity random walks leave meter-level uncertainty in the relative position over the minutes that geometry change needs, so the inertial prediction adds little to carrier phase. A 2 $\mathrm{mm}/\sqrt{\mathrm{s}}$ motion aid, as is sometimes assumed, corresponds to centimeter knowledge over a minute and halved LEO 1k's fix time (278 s); it is optimistic for an aircraft. The measurement rate matters more: with 5 s instead of 1 s epochs, LEO 1k fixed in only 10% of start times, and LEO 5k and GNSS took 182 and 718 s.

Recovery after an outage. The times above are for a cold start. After an outage the receiver is not cold: its troposphere and orbit-clock residual states are still known, with variances grown by their dynamics, and its position is known to within the inertial drift over the outage. Only the ambiguities are new, because lock is lost on every satellite. We simulated this warm restart: 900 s of normal tracking, an outage of 15, 60, or 180 s without measurements during which the filter propagates on the inertial model, and then reacquisition of every satellite with new integer ambiguities (25 start times each; two frequencies, 10 cm ionosphere and 2 cm orbit-clock residuals). The position uncertainty at reacquisition was 0.8, 16, and 235 m. With 5,000 satellites the float solution regained 5 cm in a median 15, 17, and 18 s and the fixed 2 cm solution in 49, 54, and 56 s, against 25 and 105 s for a cold start at the same start times. With 1,000 satellites the float solution took 42, 47, and 61 s (cold 55 s) and the fixed solution 470, 473, and 430 s (cold 566 s; the bootstrap intervals of the warm fix times, 300–570 s, overlap). No reacquired fix used a wrong integer. The gain comes from the states that survive the outage, so it shrinks as the inertial position prior degrades: after 180 s, the 1,000-satellite float solution was no faster than a cold start. Fix times depend little on the outage length because they are set by the new ambiguities, which must be resolved afresh either way. Recovery also depends on the state before the outage. With a 60 s outage after only 30, 120, or 300 s of tracking instead of 900 s, the float solution regained 5 cm in 20, 17, and 17 s with 5,000 satellites (56, 49, and 54 s with 1,000), and the fixed solution came in 94, 68, and 59 s (546, 529, and 407 s with 1,000, where one start time in 25 did not fix within 15 min). In these tests, an interruption before the float solution had converged recovered almost like a cold start, two minutes of continuous tracking before the outage already gave the full float benefit, and the fix benefit keeps growing with longer prior tracking, presumably because the retained troposphere and orbit-clock states are better estimated.

TABLE IX
SIMULATED PERFORMANCE OF A FLYING RECEIVER (TWO FREQUENCIES, 1 HZ, INDUSTRIAL IMU, 10 CM IONOSPHERE RESIDUAL, RESIDUAL CORRELATION 600 S); ACHIEVED ERROR IS THE RMS 3-D ERROR OVER 5–10 MIN AFTER A COLD START.

Configuration	To 5 cm (float)	Fixed	Achieved error
GNSS, 1 cm orbit–clock	832 s	400 s	2.5 cm
LEO 1k, 2 cm orbit–clock	56 s	594 s	2.7 cm
LEO 5k, 2 cm orbit–clock	25 s	104 s	1.4 cm
LEO 1k, 5 cm orbit–clock	200 s	none	5.1 cm
LEO 5k, 5 cm orbit–clock	62 s	none	2.9 cm

Insight. LEO's clearest advantage is recovery speed: with a large constellation, a formal 5 cm float solution in under a minute (about 14 min for GNSS) and a formal 2 cm fixed solution in under two (about 7 min for GNSS), which matters after every outage; after a short outage, a receiver that keeps its other states recovers faster still (about 15–18 s to float and 50–56 s to fix with 5,000 satellites). Whether and how fast a fixed solution is reached depends on both the error budget and the constellation. In the tested cases, fixing required orbit and clock residuals near 2 cm, residuals stable over a pass, and either two frequencies or good ionosphere corrections; within that budget, 5,000 satellites fixed in 104 s against 594 s for 1,000. Under those assumptions a 5,000-satellite layer achieved about 1.3 cm and a 1,000-satellite layer about 2.5 cm; with 5 cm residuals, they stayed at 3–5 cm. Realistic in-flight inertial aiding does not change these conclusions; the measurement rate does.

VIII. Discussion: What Is Achievable, and Under What Assumptions

Table IX collects the conditional answers to the question of Section I. They are simulation results under stated assumptions, not demonstrated service levels.

What LEO buys. Its principal advantages here are stronger signals and faster geometry change; its snapshot geometry per visible satellite is not intrinsically better (Section IV), although a larger constellation improves the usable geometry through more satellites and a larger choice among them. Signal strength gives about 25–33 dB of jamming margin, depending on constellation size, for the assumed altitude and power; the margin stops growing for geometric reasons; it is not limited by interference among the satellites' own signals in the shared band. Fast geometry change gives recovery to a few-centimeter precision in about a minute, and that benefit remains even where strict integer-fixing criteria are not met.

What this means for Starlink-class constellations. The 550 km shells simulated here are at the altitude where most Starlink satellites flew until SpaceX began lowering them to about 480 km in 2026 [46] (at 480 km a satellite overhead is about 1.2 dB stronger, and fewer satellites are in view for a given constellation size), and the 5,000- and 10,000-satellite cases are of comparable scale to Starlink, which is

somewhat larger. The results therefore describe approximately what a constellation of Starlink's size and altitude would buy *if* it broadcast a navigation signal with the assumed power, frequencies, and orbit and clock products good to about 2 cm: a jamming margin near the 33 dB ceiling, a float solution at 5 cm in under half a minute, and fixed solutions with about 1.3 cm achieved error. Today's signal-of-opportunity results are several meters horizontally. They were obtained with different receivers, signals, and measures than the 3-D RMS errors simulated here, so the two are not a matched comparison, and their ratio is not a measured performance gap; but the reasons for the difference lie mainly in the error budget, not in the geometry.

The first item in that budget is time. If all satellites transmitted on one common clock, differencing arrival times between satellites would remove the receiver clock, and Starlink would work as a time-difference-of-arrival (TDOA) system with exactly the clock-eliminated geometry of (2); it is the dual of emitter localization by TDOA, with many emitters on a common time and one receiver instead of one emitter and many receivers [36]. Measurements show that the premise does not hold today. Frame arrival times can be measured to about 2 ns, but the frame clocks are only loosely disciplined by GNSS and are not aligned to GPS time (each 15 s scheduling interval begins 2–4 ms before a GPS-time boundary), different beams of one satellite can run on separate frame clocks, older satellites adjust their clocks once per second by about 117 ns on average, and drifts reach several parts per million [47]. On these grounds the authors judged the downlink unsuitable for purely opportunistic pseudorange positioning; assuming aligned clocks gives a 95% horizontal error bound of about 54 m, and they regard even that model as too simple. They attribute the problem to software rather than hardware and name three remedies: aligning the frames to GPS time, broadcasting a clock model, or a ground reference network that measures the frame clocks and publishes corrections, the last being the analogue of differential GNSS and possible without Starlink's operator.

The second item is orbits. Two-line element sets locate a satellite only to roughly a kilometer; more accurate orbit data derived from the operator's own ephemerides are published [48], but their accuracy for navigation has not been established. Even a common clock and good orbits would leave a gap: Section VII found that 5 cm of orbit-clock residual already prevented integer fixing. A third limit is specific to a broadband system: the downlink is spot-beamed and scheduled, so a receiver hears the satellites currently serving its area rather than every satellite in view, and the number of usable satellites in (4) is set by the scheduler rather than by the constellation size. Two other points carry over. Starlink's broadband downlink is in Ku band, far from the GNSS bands, so a jammer aimed at GNSS does not affect it, and a receiver using both gains tolerance band by band (Appendix A). And the satellites' beamformed broadband transmissions are designed for communication, not for continuous, predictable navigation ranging; whether an operator would provide navigation signals, products, and continuity guarantees is a business and policy question that geometry cannot answer.

What "1 cm" requires. The studies here use a 2 cm formal fixed-solution criterion and measured achieved errors of about 1–3 cm; they do not demonstrate a 1 cm service. In the tested scenarios, fixed solutions needed orbit and clock residuals near 2 cm that are stable over a pass, and either two frequencies or good ionosphere corrections; achieved errors of 1–3 cm came with fixing, and float solutions without it stayed at about 2.4–5 cm. Questions to ask of any claim: float or fixed? Formal or achieved? RMS or 95%? Static or flying, and at what measurement rate? Which orbit and clock products, and how many frequencies?

Which satellites to track. With identical satellites and a filter matched to the error model, discarding measurements cannot improve the estimate: the filter already weights each measurement optimally. Choosing a subset matters for two other reasons. First, receivers have a limited number of channels: with 5,000 satellites about 117 are in view, while the simulated receiver tracks 48, chosen by highest elevation. Second, satellites can differ in the quality of their orbit and clock products, for example during eclipse or after a maneuver. Better selection rules might favor satellites with long remaining passes (fewer new ambiguities to resolve), satellites whose geometry changes fast relative to the correlation time of the residuals, or the subset that maximizes information over a horizon. For a fixed set of candidates whose measurements contribute additive information $J_i \succeq 0$ to a prior $J_0 \succ 0$, the information gain $F(S) = \log\det(J_0 + \sum_{i\in S} J_i) - \log\det J_0$ is normalized, monotone, and submodular, so choosing a fixed number of satellites greedily achieves at least $1 - 1/e$ of the optimal gain [49]; convex relaxations give bounds on the optimum [50]. The guarantee concerns that information gain only. For a single epoch with the clock as the only shared state, the 4×4 information matrix J of position and clock and the 3×3 position information J_x left after eliminating the clock (its Schur complement) satisfy $\det J = (\sum_i w_i)\det J_x$, so the log-determinant of the full information and that of the position information after clock elimination differ by $\log \sum_i w_i$, which is the same for every subset of a given size when all weights are equal; D-optimal selection is then unaffected by clock elimination. The position-error criterion used in this paper, $\mathrm{tr}\, P_{\mathrm{pos}}$, is A-optimal, and A-optimality is not submodular in general. The determinant is the product of the information eigenvalues, whereas $\mathrm{tr}\, P_{\mathrm{pos}}$ is the sum of their reciprocals, so a subset with a large determinant can still leave one axis poorly determined, typically the vertical when satellites bunch near zenith. It does not carry over automatically to the position information left after the clock and other shared states are eliminated, to the probability of fixing the integers, to convergence time, or to a selection repeated over time with satellite turnover and correlated errors. Whether such rules can restore fixing when the residuals decorrelate within a minute, the case in which LEO 5k did not fix, is an open question we leave to future work.

What "unjammable" requires. No constellation makes a single-antenna receiver unjammable: a 1 kW jammer still denies a receiver with 30 dB of advantage within 10 km (Table V). The user antenna contributes as much as the constellation, and its performance depends on jammer ge-

ometry in ways null counting misses: counting nulls is too pessimistic when jammers outnumber the array's elements, and too optimistic when a jammer lies close in direction to a satellite.

IX. Limitations

- Walker constellations, idealized EIRP, and a patch element pattern stand in for real designs; existing broadband constellations, such as Starlink, used as signals of opportunity, are discussed but not simulated; real LEO navigation signals may use different bands, waveforms, and powers.
- The array study is narrowband, with finite null depth modeled by channel decorrelation and sample size; mutual coupling, calibration drift, wideband dispersion, and front-end saturation by very close jammers are not modeled.
- Only hard partial fixing by integer bootstrapping was tested; soft (BIE, PDA) and multiple-hypothesis (mixture Kalman filter) estimators may behave differently. The carrier-phase filter matches the truth model (an optimistic assumption); the bootstrapped success-rate criterion is a model-based reliability measure and does not validate individual solutions (no ratio test); all GNSS satellites are assumed to transmit three frequencies; inter-frequency and phase biases are treated as corrected, and multipath enters only as white noise.
- Results are for one latitude (42°N) and a 30 min window; the 50 start times are independent draws of time of day, and bootstrap intervals of the medians have half-widths of 0–35% (up to 70% for GNSS fix times).

X. Conclusion

A large LEO constellation buys jamming margin and fast recovery; its snapshot geometry per visible satellite is not intrinsically better than that of GNSS. Its jamming margin saturates near 33 dB for the assumed altitude and power, for geometric reasons, not by interference among the satellites' own signals. In simulation, large LEO constellations accelerated float convergence (a formal 5 cm solution in about 25 s with 5,000 satellites, against about 14 min for GNSS), while fixing and actual accuracy remained sensitive to the residual-error assumptions: achieved errors of 1–3 cm where orbit and clock residuals were about 2 cm and stable over a pass, and 3–5 cm with 5 cm residuals. Antenna arrays must be modeled element by element. Future work includes soft and any-time ambiguity estimators, and rules for choosing which of the many satellites in view to track. Part II builds on these results to design jam-resilient architectures and flight policies for unmanned aircraft.

Appendix A
Derivations

Geometry factor (3). The derivation takes five steps.

Step 1: information as a covariance. By (2), $\sigma^2 J_x/m = \frac{1}{m}\sum_i (\mathbf{u}_i - \bar{\mathbf{u}})(\mathbf{u}_i - \bar{\mathbf{u}})^\top$ is the sample covariance of the m line-of-sight directions. If the directions are independent draws from a distribution with covariance C_u, the sample covariance tends to C_u, so $J_x \approx (m/\sigma^2)C_u$ and $\operatorname{tr} P_{\mathrm{pos}} \approx (\sigma^2/m)\operatorname{tr} C_u^{-1}$. The clock enters only through the centering: it removes the mean direction, exactly as differencing removes the emission time in TDOA localization [36].

Step 2: the direction distribution. Let the satellites be uniform in solid angle over the part of the sky above the elevation mask e_0. In east-north-up coordinates, a direction with elevation e and azimuth ϕ is

$$\mathbf{u} = \big(\sqrt{1-z^2}\cos\phi,\ \sqrt{1-z^2}\sin\phi,\ z\big), \qquad z = \sin e.$$

By Archimedes' hat-box theorem, the area of a spherical zone is proportional to its height, so uniform solid angle means z is uniform on $[c_0, 1]$, $c_0 = \sin e_0$, with ϕ uniform on $[0, 2\pi)$ and independent of z.

Step 3: means. Averaging over azimuth gives $\mathbb{E}[u_E] = \mathbb{E}[u_N] = 0$, and $\mathbb{E}[u_U] = \mathbb{E}[z] = (1 + c_0)/2$.

Step 4: variances and covariances. The vertical component is uniform on an interval of length $1 - c_0$, so $\operatorname{Var}(u_U) = (1 - c_0)^2/12$. For the horizontal components, $\mathbb{E}[\cos^2\phi] = \frac{1}{2}$ and $\mathbb{E}[z^2] = (1 + c_0 + c_0^2)/3$, so

$$\operatorname{Var}(u_E) = \operatorname{Var}(u_N) = \tfrac{1}{2}\big(1 - \mathbb{E}[z^2]\big) = \frac{(1-c_0)(2+c_0)}{6}.$$

All covariances vanish: those between u_E and u_N and between either and u_U average $\sin\phi\cos\phi$, $\cos\phi$, or $\sin\phi$ over a full turn. Hence

$$C_u = \operatorname{diag}\Big(\tfrac{(1-c_0)(2+c_0)}{6},\ \tfrac{(1-c_0)(2+c_0)}{6},\ \tfrac{(1-c_0)^2}{12}\Big).$$

Step 5: the trace of the inverse. Inverting the diagonal entries and adding gives

$$\operatorname{tr} C_u^{-1} = 2 \cdot \frac{6}{(1-c_0)(2+c_0)} + \frac{12}{(1-c_0)^2},$$

which is (3). For the full hemisphere ($c_0 = 0$) the two horizontal terms give $3 + 3$ and the vertical term 12, total 18; for a 10° mask ($c_0 = 0.174$) they give $3.5 + 3.5$ and 17.6, total 24.3. The vertical term dominates because the vertical components of all directions lie in the short interval $[c_0, 1]$: every satellite is above the receiver, so the vertical spread is small, and after the mean is removed (the clock) little vertical information remains.

From the average geometry to the expected error. Step 1 replaced the sample covariance by its mean. Because the matrix inverse is convex, $\mathbb{E}[J_x^{-1}] \succeq (\mathbb{E}J_x)^{-1}$, so the average-geometry value $\sigma^2 G/m$ is optimistic for a finite number of satellites; the correction below quantifies by how much.

Finite-m correction (heuristic). Treating the sample covariance of m directions as Wishart with $m-1$ degrees of freedom in dimension 3 gives $\mathbb{E}[\hat{C}_u^{-1}] \approx C_u^{-1}(m-1)/(m-5)$ and hence (4). The directions lie on a spherical cap and are not Gaussian, so this is an approximation; its accuracy for $m \geq 6$ is checked by simulation in Section III, and its pole at $m = 5$ matches the true divergence of the expected squared error shown there.

Additive power-density tolerances. An attacker splitting total power among bands ℓ with tolerated power densities Φ_ℓ denies all bands within $\pi d^2 = P_J G_J/(4\sum_\ell \Phi_\ell)$, so jamming tolerances add across separate bands.

APPENDIX B
EXPERIMENT SPECIFICATION

Constellations. All shells are circular Walker patterns $i{:}T/P/F$ with phasing $F = 1$, propagated as Keplerian orbits on a rotating Earth. The LEO layers use three shells at 550 km: 53° with 72 planes, 70° with 36 planes, and 97.6° with 20 planes, carrying about 58%, 26%, and 16% of the satellites. The layer labeled 1,000 has $576 + 252 + 160 = 988$ satellites, 5,000 has $2{,}880 + 1{,}296 + 800 = 4{,}976$, and 10,000 has $5{,}832 + 2{,}592 + 1{,}600 = 10{,}024$. The GNSS reference has four MEO shells: 20,180 km, $55^\circ{:}30/6/1$; 23,222 km, $56^\circ{:}24/3/1$; 21,528 km, $55^\circ{:}24/3/1$; and 19,130 km, $64.8^\circ{:}24/3/1$ (102 satellites).

Random draws. Jamming margin (Section V): 40 times of day drawn with seed 1. Finite-m geometry (Section III): 200,000 skies per case, seed 9. Array study (Section VI): 300 geometries per row, seed 17. Carrier phase (Section VII): start time r (the time of day, uniform over a sidereal day) drawn with seed $1000 + r$, $r = 0, \ldots, 49$; warm restarts use $r = 0, \ldots, 24$.

Code. The accompanying archive contains every script and a file `REPRODUCE.md` listing the commands, in order, that regenerate each table and figure from the stored results or from scratch.

REFERENCES

[1] B. Mediate, "GPS operations / modernization update," Civil GPS Service Interface Committee (CGSIC), https://www.gps.gov/sites/default/files/2026-04/Mediate%2020260421_CGSIC%20MD%2031_FINAL.pdf, Apr. 2026.

[2] Orbital Radar, "Navigation constellations," https://orbitalradar.com/navigation-constellations, counts as of Mar. 2026; accessed Oct. 2026.

[3] E. D. Kaplan and C. J. Hegarty, *Understanding GPS/GNSS: Principles and Applications*, 3rd ed. Norwood, MA: Artech House, 2017.

[4] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*, 2nd ed. Lincoln, MA: Ganga-Jamuna Press, 2006.

[5] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner, "Assessing the spoofing threat: Development of a portable GPS civilian spoofer," in *Proc. ION GNSS 2008*, 2008, pp. 2314–2325.

[6] T. G. R. Reid, A. M. Neish, T. Walter, and P. K. Enge, "Broadband LEO constellations for navigation," *NAVIGATION*, vol. 65, no. 2, pp. 205–220, 2018.

[7] F. S. Prol *et al.*, "Position, navigation, and timing (PNT) through low Earth orbit (LEO) satellites: A survey on current status, challenges, and opportunities," *IEEE Access*, vol. 10, pp. 83 971–84 002, 2022.

[8] DatacenterDynamics, "FCC approves Xona application for 258-satellite PNT constellation," https://www.datacenterdynamics.com/en/news/fcc-approves-xona-application-for-258-satellite-pnt-constellation/, Aug. 2026.

[9] HighSpeedInternet.com, "How many Starlink satellites are in orbit as of October 23?" https://www.highspeedinternet.com/resources/how-many-starlink-satellites-are-in-orbit-october-23-2026, Oct. 2026, estimate.

[10] M. Neinavaie and Z. M. Kassas, "Unveiling Starlink LEO satellite OFDM-like signal structure enabling precise positioning," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 60, no. 2, pp. 2486–2489, 2024.

[11] Y. Xu, Y. Liu, M. Lei *et al.*, "Joint pseudo-range and Doppler positioning method with LEO satellites' signals of opportunity," *Satellite Navigation*, vol. 6, p. 10, 2025.

[12] X. Li, F. Ma, X. Li, H. Lv, L. Bian, Z. Jiang, and X. Zhang, "LEO constellation-augmented multi-GNSS for rapid PPP convergence," *Journal of Geodesy*, vol. 93, no. 5, pp. 749–764, 2019.

[13] K. R. Pattipati, "Resilient navigation architectures and proactive routing for unmanned aircraft under uncertain interference (Part II)," 2026, working draft.

[14] J. G. Walker, "Continuous whole-Earth coverage by circular-orbit satellite patterns," Royal Aircraft Establishment, Tech. Rep. TR 77044, 1977.

[15] ——, "Satellite constellations," *Journal of the British Interplanetary Society*, vol. 37, pp. 559–572, 1984.

[16] A. Joseph, "Measuring GNSS signal strength," Inside GNSS, Nov./Dec. 2010, https://insidegnss.com/measuring-gnss-signal-strength/, 2010.

[17] P. J. G. Teunissen, "The least-squares ambiguity decorrelation adjustment: A method for fast GPS integer ambiguity estimation," *Journal of Geodesy*, vol. 70, pp. 65–82, 1995.

[18] ——, "Success probability of integer GPS ambiguity rounding and bootstrapping," *Journal of Geodesy*, vol. 72, pp. 606–612, 1998.

[19] EUSPA, "From testing to operations: Galileo OSNMA service now available to users," https://www.euspa.europa.eu/newsroom-events/news/testing-operations-galileo-osnma-service-now-available-users, 2025.

[20] ——, "Have you met HAS?" https://www.euspa.europa.eu/newsroom-events/news/have-you-met-has, 2023.

[21] Cislunar Space, "BeiDou in-orbit upgrade complete," https://cislunarspace.cn/en/space-news/2026/07/2026-07-31-beidou-in-orbit-upgrade-complete/, Jul. 2026.

[22] GPS World, "GLONASS: The decade of transition to CDMA signals," https://www.gpsworld.com/glonass-the-decade-of-transition-to-cdma-signals/, Dec. 2023.

[23] JAXA, "Launch result of H3 launch vehicle flight no. 8 with QZS-5 onboard," Press release, https://global.jaxa.jp/press/2025/12/20251222-2_e.html, Dec. 2025.

[24] GPS World, "3 atomic clocks fail on 1 Indian satellite; replacement prepped," https://www.gpsworld.com/3-atomic-clocks-fail-on-1-indian-satellite-replacement-prepped/, 2017.

[25] ETV Bharat, "India's navigation satellite system NavIC goes offline amid critical satellite loss," https://www.etvbharat.com/en/technology/indias-navigation-satellite-system-navic-goes-offline-amid-critical-satellite-loss-enn26031403113/, Mar. 2026.

[26] Business Today, "Only three satellites left active: Centre acknowledges NavIC standalone navigation gap," https://businesstoday.in/india/story/only-three-satellites-left-active-centre-acknowledges-navic-standalone-navigation-gap-546285-2026-07-30, Jul. 2026.

[27] Indian Space Research Organisation, "GSLV-F18/NVS-03 mission," https://www.isro.gov.in/Mission_GSLVF18.html, 2026, accessed 1 November 2026.

[28] ThePrint, "ISRO launches NVS-03 satellite to revive India's NavIC system," https://theprint.in/science/isro-nvs-03-satellite-launch-navic/3086653/, Oct. 2026.

[29] GPS World, "India to launch NVS-03 satellite to bolster NavIC," https://www.gpsworld.com/india-to-launch-nvs-03-satellite-to-bolster-navic/, Oct. 2026.

[30] Iridium, "Iridium PNT performance testing," Summary of tests by the European Commission Joint Research Centre, https://www.iridium.com/sites/default/files/2026-03/Iridium%20PNT%20Performance%20Testing_March%202026.pdf, Mar. 2026.

[31] Ohio State ElectroScience Laboratory, "Kassas and ASPIN team featured on GPS World cover," https://electroscience.osu.edu/news/2026/03/kassas-and-aspin-team-featured-gps-world-cover, Mar. 2026.

[32] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," arXiv:2210.11578, 2022.

[33] Inside GNSS, "SpaceX details Starlink's existing and potential PNT capabilities in response to FCC inquiry," https://insidegnss.com/spacex-details-starlinks-existing-and-potential-pnt-capabilities-in-response-to-fcc-inquiry/, May 2025, fCC WT Docket No. 25-110; corrected May 2026.

[34] Aviación al Día, "2025 IATA annual safety report: global accident rate improves, but fatalities and GNSS interference surge," https://aviacionaldia.com/en/2026/03/2025-iata-annual-safety-report-global-accident-rate-improves-but-fatalities-and-gnss-interference-surge.html, Mar. 2026.

[35] EASA, "Safety information bulletin 2022-02R4: Global navigation satellite system outages and alterations leading to communication / navigation / surveillance degradation," https://ad.easa.europa.eu/blob/EASA_SIB_2022_02R4.pdf/SIB_2022-02R4_1, 2026.

[36] R. J. Harvey, K. R. Pattipati, and P. K. Willett, "A smooth manifolds approach to TDOA-based localization under environmental uncertainty," 2026, IEEE Open Journal of Signal Processing, under revision.

[37] J. Kiefer and J. Wolfowitz, "The equivalence of two extremum problems," *Canadian Journal of Mathematics*, vol. 12, pp. 363–366, 1960.

[38] F. Pukelsheim, *Optimal Design of Experiments*. New York: Wiley, 1993.

[39] R. T. Compton, *Adaptive Antennas: Concepts and Performance*. Englewood Cliffs, NJ: Prentice Hall, 1988.

[40] P. J. G. Teunissen, "Theory of integer equivariant estimation with application to GNSS," *J. Geodesy*, vol. 77, no. 7–8, pp. 402–410, 2003.

[41] K. Berntorp, A. Weiss, and S. D. Cairano, "Integer ambiguity resolution by mixture Kalman filter for improved GNSS precision," *IEEE Trans. Aerosp. Electron. Syst.*, 2020.
[42] M. Greiff and K. Berntorp, "Optimal measurement projections with adaptive mixture Kalman filtering for GNSS positioning," in *Proc. American Control Conf.*, 2020, pp. 4435–4441.
[43] J. Luo, K. R. Pattipati, P. Willett, and G. M. Levchuk, "Fast optimal and suboptimal any-time algorithms for CDMA multiuser detection based on branch and bound," *IEEE Trans. Commun.*, vol. 52, no. 4, pp. 632–642, 2004.
[44] J. Luo, K. R. Pattipati, P. K. Willett, and F. Hasegawa, "Near-optimal multiuser detection in synchronous CDMA using probabilistic data association," *IEEE Commun. Lett.*, vol. 5, no. 9, pp. 361–363, 2001.
[45] D. Pham, K. R. Pattipati, P. K. Willett, and J. Luo, "A generalized probabilistic data association detector for multiple antenna systems," *IEEE Commun. Lett.*, vol. 8, no. 4, pp. 205–207, 2004.
[46] Interesting Engineering, "SpaceX to lower Starlink satellite orbits in 2026," https://interestingengineering.com/space/spacex-starlink-orbit-lowering-2026, Jan. 2026.
[47] W. Qin, A. M. Graff, Z. L. Clements, Z. M. Komodromos, and T. E. Humphreys, "Timing properties of the Starlink Ku-band downlink," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 62, pp. 727–744, 2026, arXiv:2501.05302.
[48] CelesTrak, "Supplemental GP element sets," https://celestrak.org/NORAD/elements/supplemental/, 2026, accessed Oct. 2026.
[49] M. Shamaiah, S. Banerjee, and H. Vikalo, "Greedy sensor selection: Leveraging submodularity," in *Proc. 49th IEEE Conf. Decision and Control*, 2010.
[50] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 451–462, 2009.