A Framework for Integrity-Ensured URA Generation in Dedicated LEO PNT Systems

Dongchan Min and Sam Pullen
Stanford University

Jimin Do and Jiyun Lee
Korea Advanced Institute of Science and Technology

BIOGRAPHY

Dongchan Min is a postdoctoral scholar within the GNSS Laboratory at Stanford University. He holds M.S. and Ph.D. degrees in Aerospace Engineering from the Korea Advanced Institute of Science and Technology (KAIST).

Jimin Do is a candidate in the integrated M.S. and Ph.D. program at the Korea Advanced Institute of Science and Technology (KAIST). He holds a B.S. in Aerospace Engineering from KAIST.

Sam Pullen is a senior researcher within the GNSS Laboratory at Stanford University, where he received his Ph.D. in Aeronautics and Astronautics in 1996. He has supported the FAA and other service providers in developing system concepts, technical requirements, integrity algorithms, and performance models for GBAS, SBAS, and other GNSS services and applications. He has also performed GNSS system design, application development, risk assessment, and legal support through his consultancy, Sam Pullen Consulting. He was awarded the ION Early Achievement Award in 1999 and became an ION Fellow in 2017.

Jiyun Lee is an Endowed Chair Professor in the department of Aerospace Engineering at the Korea Advanced Institute of Science and Technology (KAIST). Her specialties lie in the area of GNSS augmentation systems, integrity architecture for multisensor integrated navigation systems, and GNSS remote sensing. She received her Ph.D. from Stanford University in aeronautics and astronautics. She served ION as secretary of Satellite division executive committee (2022-2024), technical representative of ION council (2019-2021), and pacific technical advisor (2016-2018, 2020-2022).

ABSTRACT

Dedicated Low Earth Orbit (LEO) constellations for Positioning, Navigation, and Timing (PNT) can augment Global Navigation Satellite Systems (GNSS). Such constellations exhibit fast orbital dynamics and comprise too many satellites for broadcast navigation messages to be uploaded from the ground at the rate that would be required. Many dedicated LEO PNT concepts therefore assume that each satellite generates its navigation message onboard from GNSS measurements. This message includes the User Range Accuracy (URA), a conservative estimate of the User Range Error (URE) standard deviation, which users employ to compute protection levels. Onboard message generation, however, introduces a fault mode absent from conventional GNSS: a fault in a one or more GNSS satellites can cascade through onboard orbit determination and appear as LEO ephemeris and clock errors. Because this affects multiple LEO satellites simultaneously, this cascaded GNSS fault is a common-cause fault and must be accounted for within the constellation fault probability. Among the fault modes of onboard generation, this paper focuses on the cascaded GNSS fault. We develop a framework that computes a URA guaranteeing integrity for a given allocation of the constellation fault probability to this fault. We extend the protection-level concept of civil-aviation user integrity monitoring to the URE domain. An onboard solution-separation monitor detects cascaded GNSS faults, and a URE-domain protection level is computed from the worst case among the faults that remain undetected. The URA broadcast by LEO satellites is derived from this URE-domain protection level such that the probability of an unalerted URE exceedance caused by cascaded GNSS faults does not exceed the allocated value. We evaluate this framework through simulations of representative high- and low-altitude LEO scenarios using synthetic GPS and Galileo observations.

I. INTRODUCTION

Low Earth Orbit (LEO) constellations designed for Positioning, Navigation, and Timing (PNT) can augment Global Navigation Satellite Systems (GNSS) by providing stronger received signal power, rapidly changing geometry, and an independent source of ranging signals (Reid et al., 2018). These features can improve position accuracy and filter convergence while adding system-level redundancy to Medium Earth Orbit (MEO) GNSS. The implementation of a dedicated LEO PNT system, however, departs from GNSS practice in how the broadcast navigation message is generated. A dedicated LEO PNT constellation includes

far more satellites than an MEO GNSS constellation, and the orbits of its satellites change rapidly under atmospheric drag and gravitational perturbations. Uploading updated parameters to every LEO satellite from the ground at the cadence these dynamics demand, as GNSS does, is generally not feasible. Therefore, many dedicated LEO PNT concepts rely on GNSS to generate the broadcast navigation message onboard rather than receive it from the ground control segment.

In GPS, the broadcast messages carry User Range Accuracy (URA) along with their orbit and clock parameters, so that users can compute protection levels. URA is a *conservative* standard deviation of the User Range Error (URE) caused by satellite orbit and clock errors. It is designed so that a Gaussian distribution with standard deviation URA overbounds the actual URE distribution to a specified probability level (DeCleene, 2000). The URE is considered nominal within the range where the URA Gaussian is guaranteed to overbound it. Beyond that range, an unalerted exceedance is treated as a Signal-in-Space (SIS) fault (U.S. Department of Defense, 2020).

GPS places this nominal/fault boundary at a Not-To-Exceed (NTE) threshold, defined as $\kappa \times \mathrm{URA}$, where κ is the value at which the two-sided standard Gaussian tail probability equals the target satellite fault probability P_{sat}. The legacy GPS service guarantees $P_{\mathrm{sat}} \leq 10^{-5}$, for which κ is 4.42, so its NTE is $4.42\,\mathrm{URA}$ (U.S. Department of Defense, 2020). GPS Block III, when a satellite's Integrity Status Flag is set, widens the NTE to $5.73\,\mathrm{URA}$ and guarantees $P_{\mathrm{sat}} \leq 10^{-8}$ (Kovach et al., 2008). This smaller P_{sat} guarantee does not follow from the larger κ alone, because the URA Gaussian that overbounds the true URE out to the 10^{-5} tail can no longer be assumed to do so out to 10^{-8}. GPS Block III instead adopts an early-warning strategy that monitors the root causes likely to introduce a large URE (Kovach et al., 2008). By lowering the probability of such an event, it keeps the URA Gaussian a valid overbound out to the 10^{-8} tail. Achieving a small P_{sat} without inflating the URA therefore depends on monitor effectiveness.

In a dedicated LEO PNT system, generating the broadcast message and its URA onboard introduces a fault mode absent from conventional GNSS. A fault in a GNSS satellite can propagate through onboard orbit determination (OD) and appear as a LEO ephemeris or clock error. We refer to this mode as a cascaded GNSS fault (CGF). LEO SIS fault modes include both faults that originate on the LEO satellite and cascaded GNSS faults that are not detected onboard. A cascaded GNSS fault may not be confined to one LEO satellite. Every LEO satellite that tracks the faulted GNSS satellite incorporates the same corrupted measurements into its OD, so a single GNSS fault can corrupt the broadcast ephemerides of several LEO satellites simultaneously. The probability of such faults, which affect two or more LEO satellites, should be allocated from the constellation fault probability P_{const}. The cascaded GNSS fault therefore belongs in the P_{const} of the LEO constellation. We denote the portion of P_{const} allocated to it by P_{CGF}.

URA generation for LEO PNT must account for cascaded GNSS faults explicitly. Prior work has addressed parts of this problem, but not the complete chain from onboard ephemeris generation through fault monitoring to URA. Wang et al. (2022b) introduced an orbit URA for predicted LEO orbits, using worst-location orbit URE computations that account for the Earth's ellipsoidal geometry, but focused on nominal ephemeris errors rather than fault modes. Wang et al. (2022a) proposed an integrity monitoring framework for real-time LEO OD under nominal conditions, but did not account for undetected faults. Gioia et al. (2025) developed a time receiver autonomous integrity monitoring scheme for LEO OD to detect cascaded GNSS faults, but did not extend the results to URA generation. A method is still needed that maps a target allocation P_{CGF} to a broadcast URA whose Gaussian overbounds the URE out to that probability even under the cascaded GNSS faults left undetected onboard.

This paper develops a framework for integrity-ensured URA generation given a target allocation P_{CGF} on dedicated LEO PNT satellites that generate their broadcast ephemeris onboard. The framework brings the civil-aviation user-integrity concept, in which a protection level bounds the user position error against undetected faults (Blanch et al., 2015), to the URE domain. An onboard solution separation (SS) monitor in the URE domain detects cascaded GNSS faults, and a URE-domain protection level for the target P_{CGF} is computed from the worst case over the faults the monitor leaves undetected. The broadcast URA follows from this protection level and keeps the probability of an unalerted URE exceedance caused by cascaded GNSS faults at or below the target. The integrity budget of a LEO satellite comprises many fault modes, of which the cascaded GNSS fault is one. This paper addresses that mode alone. The others, and their allocations within P_{sat} and the remainder of P_{const}, are outside its scope.

The remainder of this paper is organized as follows. Section II reviews how the broadcast navigation message is generated onboard and the error sources this introduces. Section III defines the URE and the URA, relates them to the satellite fault probability P_{sat} of GPS, and states the URA requirement that this paper meets for P_{CGF}. Section IV places the cascaded GNSS fault among the SIS fault modes of onboard generation and defines its allocation P_{CGF}. Section V formulates the URE-domain protection level under the onboard SS monitor and derives the broadcast URA from it. Section VI presents the simulation results, and Section VII concludes.

II. BACKGROUND: NAVIGATION MESSAGE GENERATION

The broadcast navigation message carries the satellite ephemeris and clock parameters from which users compute the position and clock correction of each ranging source. For conventional GNSS, these messages are determined and uploaded by the ground control segment. The GPS Control Segment, for instance, computes them and uploads new navigation data to each satellite at least once per day, with additional contingency uploads only when a satellite's range error grows faster than expected (U.S. Department of Defense, 2020). This suffices for conventional GNSS, whose satellites experience weaker perturbing forces and are fewer in number than those of a LEO PNT constellation. A dedicated LEO PNT system cannot follow this model, because its far larger constellation and rapidly evolving low orbits would require ground uploads at a cadence that the ground segment cannot sustain. Its satellites are therefore expected to generate the broadcast message onboard from GNSS measurements. The navigation message generation procedure for LEO satellites described by Kunzi et al. (2026) proceeds in three stages: real-time OD using GNSS measurements, forward prediction of the orbit and clock over the ephemeris validity interval, and least-squares fitting of the predicted states to a compact broadcast parameter set.

This work adopts the onboard generation procedure of Kunzi et al. (2026) as its baseline. In the OD stage, an Extended Kalman filter (EKF) combines ionosphere-free dual-frequency GNSS code and carrier-phase measurements with a reduced-dynamic orbit model. The filter estimates the satellite position and velocity together with empirical accelerations and the atmospheric-drag and solar-radiation-pressure (SRP) scale factors, allowing the dynamic model to absorb mismodeled forces. We assume that the satellite clock is synchronized with the onboard GNSS receiver clock. The filter therefore also estimates the receiver clock offset, which supplies the clock state needed for the broadcast message. Measurements from GPS and Galileo are used, optionally augmented by Galileo High Accuracy Service (HAS) corrections, yielding a real-time estimate of both the orbit and the clock.

To form a message that remains valid over a future interval, this real-time OD estimate is propagated forward with the same reduced-dynamic model over the ephemeris validity interval, on the order of twenty minutes, so that a single message covers a satellite pass. The predicted orbit and clock are then fit by least squares to a compact broadcast parameter set. Kunzi et al. (2026) adopt an extended-Keplerian model of roughly twenty parameters, augmenting the conventional GNSS broadcast elements with higher-order harmonics suited to the faster LEO dynamics, alongside the broadcast clock-correction coefficients. The accuracy of this broadcast message is limited by three error sources inherent to the procedure: an OD error in the real-time estimate, a propagation error from the predicted orbit and clock diverging from truth as the message ages, and a fitting error, because a finite parameter set cannot reproduce the predicted states exactly. The propagation error grows over the validity interval and dominates, whereas the fitting error remains comparatively small.

III. USER RANGE ACCURACY AND THE SATELLITE FAULT PROBABILITY

The broadcast message accuracy is characterized by the URE in the range domain. It is the satellite orbit error projected onto the user line of sight, combined with the satellite clock error (U.S. Department of Defense, 2020),

$$\mathrm{URE} = \mathbf{e}^{\top}\delta\mathbf{r} - c\,\delta t, \tag{1}$$

where $\delta\mathbf{r}$ is the broadcast orbit error, the satellite-position error of the broadcast ephemeris; $\mathbf{e}$ is the unit line-of-sight vector from the satellite to the user; c is the speed of light; and δt is the broadcast clock error. The orbit term depends on where the user lies within the satellite's coverage region. The URE arises from the OD, propagation, and fitting errors. Kunzi et al. (2026) report that the fitting error is roughly an order of magnitude smaller than the OD and propagation errors. The resulting URE stays below about 8 cm over a twenty-minute interval for a high-altitude satellite using GPS and Galileo, and grows to several decimeters at low altitude under high drag (Kunzi et al., 2026).

These error contributions vary with their underlying root cause, and the resulting URE is accordingly classified as either nominal or faulted. The nominal URE corresponds to fault-free operation, in which every error source behaves as modeled, whereas the faulted URE arises when a root cause, such as an erroneous ephemeris upload, a clock anomaly, or an unmodeled maneuver, causes one or more contributors to depart from their modeled behavior. Because the faulted URE is generally much larger than the nominal URE, it constitutes the dominant threat to user integrity. How the nominal and fault states are modeled, and the criterion by which they are separated, strongly influence the integrity assessment.

In the GPS integrity framework, the nominal URE is modeled as a Gaussian distribution, and URA is defined as a conservative one-sigma value whose Gaussian overbounds this distribution out to the NTE at the worst-case user location (DeCleene, 2000). URA is computed by the ground control segment and broadcast in the navigation message together with the orbit and clock parameters. The boundary between the nominal and fault states is set by the NTE threshold, a fixed multiple of URA. Let $\mathrm{Q}(\cdot)$ denote the upper-tail probability of the standard normal distribution and $\mathrm{Q}^{-1}(\cdot)$ its inverse. The NTE multiplier κ is $\mathrm{Q}^{-1}(P_{\mathrm{sat}}/2)$, a value at which the two-sided Gaussian tail probability equals the satellite fault probability P_{sat}. The GPS Standard Positioning Service (SPS) specifies $P_{\mathrm{sat}} \leq 10^{-5}$ and a corresponding NTE of $4.42\,\mathrm{URA}$, where $4.42 = \mathrm{Q}^{-1}(10^{-5}/2)$

(U.S. Department of Defense, 2020). A satellite is declared faulted when its instantaneous URE at the worst-case user location exceeds the NTE without a timely alert (U.S. Department of Defense, 2020). The broadcast URA thus overbounds the URE up to the 10^{-5} tail, beyond which the overbound is no longer guaranteed and the corresponding URE is treated as a fault.

GPS Block III tightens this commitment. When a satellite's Integrity Status Flag is set, the NTE is widened to $5.73\,\mathrm{URA}$ and the satellite is guaranteed $P_{\mathrm{sat}} \leq 10^{-8}$, where $5.73 = \mathrm{Q}^{-1}(10^{-8}/2)$ (Kovach et al., 2008). For a fixed URA, a wider NTE multiplier might appear to reduce P_{sat} by construction, but this inference does not hold. The nominal URA Gaussian cannot be assumed to remain valid in the 10^{-8} tail, and a probability of this magnitude is hard to validate empirically. An NTE exceedance accompanied by a timely alert is not a fault; only an unalerted exceedance is classified as such (U.S. Department of Defense, 2020). GPS Block III therefore meets the 10^{-8} commitment not by relaxing the NTE but by strengthening fault monitoring. Onboard monitors detect the root causes of a large URE, such as a clock run-off, and rapidly disable the affected signal, while the control segment validates each navigation-message upload before transmission (Kovach et al., 2008). Complemented on the ground by Independent URA Monitoring (Braff et al., 2012), this architecture lowers the probability that a large URE is broadcast without an alert and thereby allows the broadcast URA to support the 10^{-8} guarantee.

The framework proposed in this paper applies the same principle to the dedicated LEO satellite, focusing on the cascaded GNSS fault introduced in the Introduction. An onboard monitor lowers the probability that a large URE from this cause remains unalerted, and the broadcast URA is set so that its Gaussian overbounds the URE out to the tail probability allocated to this fault, P_{CGF}, whenever no alert is raised. Formally, the broadcast URA satisfies

$$P\big(|\mathrm{URE}| > \mathrm{Q}^{-1}(P_{\mathrm{CGF}}/2)\,\mathrm{URA} \;\big|\; \text{no alert}\big) \leq P_{\mathrm{CGF}}, \tag{2}$$

where $\mathrm{Q}^{-1}(P_{\mathrm{CGF}}/2)\,\mathrm{URA}$ is the NTE threshold introduced above, with P_{CGF} taking the place that P_{sat} holds in the GPS definition, and the probability is conditioned on *no alert*, the absence of an alert from the onboard monitor for cascaded GNSS faults. This conditional probability describes the URE to which users are exposed while the broadcast is in use, that is, while no alert is raised. It is also the stricter of the two possible forms. The joint probability of an exceedance and no alert equals the conditional probability multiplied by $P(\text{no alert})$, which is at most one, so a URA satisfying Eq. (2) also keeps the probability of an unalerted NTE exceedance, the joint form used in the GPS commitment, at or below P_{CGF}. The left-hand side of Eq. (2) decreases monotonically as the URA increases, so a URA satisfying the bound always exists. The true minimum URA cannot be found exactly, however, because the actual URE distribution is unknown. Using a conservative overbound of the URE instead, we adopt the smallest URA for which the bound is guaranteed. This value is no smaller than the true minimum, but keeping the gap small avoids an overly conservative integrity assessment and the attendant loss of availability.

The approach adapts the user-integrity concept of civil aviation, in which fault monitors detect the sources of large user position errors and a protection level meeting the integrity requirement is computed. Because the protection level is derived from the worst case over faults the monitors do not detect, the probability that the position error exceeds it is guaranteed to remain below the integrity requirement, on the order of 10^{-7} (Blanch et al., 2022). This paper applies the same framework in the URE domain. An onboard monitor detects cascaded GNSS faults, a protection level for the target P_{CGF} is computed in the URE domain from the worst case over the faults it leaves undetected, and the broadcast URA is defined from that protection level so that the target is guaranteed.

IV. CASCADED GNSS FAULTS IN ONBOARD LEO EPHEMERIS GENERATION

For a dedicated LEO satellite, the SIS fault modes are not those of conventional GNSS, because the broadcast ephemeris is generated onboard from GNSS measurements. Its integrity depends on those measurements and on the LEO platform itself. This dependence introduces fault modes that do not exist when the navigation message is produced and uploaded by a ground segment. As in the GPS integrity standards and the fault model of Advanced Receiver Autonomous Integrity Monitoring (ARAIM), these fault modes fall into two classes according to how many satellites they affect (U.S. Department of Defense, 2020; Walter et al., 2019). Faults that originate on a LEO satellite and affect only that satellite, such as an anomaly in its onboard GNSS receiver, an unexpected maneuver or clock event after the message is generated, or an error in fitting or encoding the broadcast parameters, are independent across the constellation and are budgeted by P_{sat}. Faults that affect several LEO satellites through a common cause, such as an atmospheric-drag surge during a geomagnetic storm, are budgeted by P_{const}. The cascaded GNSS fault belongs to this second class, and its allocation within P_{const} is P_{CGF}. Each class comprises many fault modes, and each mode requires a dedicated monitor and its own allocation. This paper addresses the cascaded GNSS fault alone, the mode through which the LEO SIS integrity depends on that of GNSS, and derives the URA required by P_{CGF}. The monitors and allocations of the other modes, within P_{sat} and the remainder of P_{const}, must be established separately, and the URA finally broadcast must satisfy those allocations as well.

A cascaded GNSS fault originates in one or more GNSS satellites as an orbit, clock, or signal fault that corrupts the measurements derived from those satellites. The corrupted measurements propagate through the OD and cascade into the broadcast LEO

ephemeris. Because every LEO satellite that tracks the faulted GNSS satellite incorporates the same corrupted measurements into its own OD, one GNSS fault can corrupt the broadcast ephemerides of several LEO satellites during the same interval, which is the defining property of a common-cause fault (U.S. Department of Defense, 2020). Within each onboard solution, the faulted satellites appear as inconsistent per-satellite contributions, but the reduced-dynamic filter absorbs part of the error and leaves them only partially revealed in the OD residuals, so a dedicated monitor is required. The prior probability of a cascaded GNSS fault is derived from the GNSS fault priors, the per-satellite and per-constellation fault probabilities to which the GNSS service commits, accumulated over the interval during which each GNSS satellite contributes to the onboard filter. Section V makes this dependence explicit through the fault-hypothesis priors of the protection-level equation.

The monitor for the cascaded GNSS fault operates in the URE domain, so that its decision applies directly to the broadcast ephemeris bound. It is realized fully onboard as a URE-domain SS monitor. The URE implied by the all-in-view ephemeris is compared with the URE from subsets that each exclude one or more GNSS satellites, and a large separation identifies the faulted satellites. Section V formalizes this monitor and the protection level it supports.

V. URE PROTECTION LEVEL FOR THE CASCADED GNSS FAULT

1. Definition of URE-Domain Protection Level and URA

Within the scope set in Section IV, the protection level in the URE domain bounds the effect of undetected cascaded GNSS faults on the broadcast orbit and clock error. As described in Section II, onboard message generation contributes two errors: an OD error and a propagation error that grows with the age of data and dominates over the validity interval. The fitting error of the broadcast parameter representation is an order of magnitude smaller and is neglected. The nominal one-sigma URE is obtained by projecting the predicted orbit and clock error covariance onto the line of sight,

$$\sigma^2_{\mathrm{URE}} = \boldsymbol{\ell}^\top \mathbf{P}_{\mathrm{Pred}}\, \boldsymbol{\ell}, \tag{3}$$

where $\boldsymbol{\ell}$ is the line-of-sight projection vector that combines the orbit line-of-sight direction and the clock term of Eq. (1), and $\mathbf{P}_{\mathrm{Pred}}$ is the covariance of the predicted orbit and clock states, obtained from the onboard EKF and accounting for both the OD and propagation errors. For protection-level computation, $\mathbf{P}_{\mathrm{Pred}}$ is formed from overbounding error models within the filter.

The protection level is defined through the integrity risk it must satisfy. Following the integrity-risk formulation of ARAIM (Blanch et al., 2022), but conditioning on the absence of an alert as in Eq. (2), the URE-domain integrity risk, also known as the probability of a hazardously misleading information (HMI) event, is the probability that the broadcast URE exceeds the protection level $\mathrm{PL}_{\mathrm{URE}}$ given that the cascaded GNSS fault monitor raises no alert,

$$\begin{aligned} \mathrm{IR}_{\mathrm{URE}} &= P(|\mathrm{URE}| > \mathrm{PL}_{\mathrm{URE}} \mid \text{no alert}) = \frac{P(|\mathrm{URE}| > \mathrm{PL}_{\mathrm{URE}},\ \text{no alert})}{P(\text{no alert})} \\ &\le \frac{1}{\eta}\left[\sum_{i=0}^{N_{\mathrm{modes}}} P(|\mathrm{URE}| > \mathrm{PL}_{\mathrm{URE}},\ \text{no alert} \mid H_i)\, P_{H_i} + P_{\mathrm{NM}}\right] \\ &\le P_{\mathrm{CGF}}, \end{aligned} \tag{4}$$

where H_i denotes the i-th fault hypothesis, with H_0 the fault-free case and $H_1, \ldots, H_{N_{\mathrm{modes}}}$ the monitored cascaded GNSS fault hypotheses, each a fault in one or more GNSS satellites. The term P_{H_i} is the prior probability of H_i, P_{NM} the combined prior of the GNSS fault combinations that are not monitored explicitly, and N_{modes} the number of monitored fault hypotheses. The numerator is bounded by the total probability theorem over the monitored hypotheses, with P_{NM} upper-bounding the unmonitored remainder, as in ARAIM. The denominator is bounded below by

$$P(\text{no alert}) \ge P(\text{no alert} \mid H_0)\, P_{H_0} \ge (1 - P_{\mathrm{FA}})\left(1 - \sum_{k=1}^{N_{\mathrm{modes}}} P_{H_k} - P_{\mathrm{NM}}\right) \equiv \eta, \tag{5}$$

because no alert is raised at least when no fault is present and the monitor does not alert falsely. Here P_{FA} is the false-alert probability of the monitor under H_0, the continuity allocation that its detection thresholds are set to meet in the next subsection, and the second factor is a lower bound on P_{H_0}, which excludes every fault, monitored or not. All probabilities refer to the message-generation epoch. The protection level is the smallest $\mathrm{PL}_{\mathrm{URE}}$ for which $\mathrm{IR}_{\mathrm{URE}} \le P_{\mathrm{CGF}}$, that is, for which the bracketed sum in Eq. (4) does not exceed $\eta\, P_{\mathrm{CGF}}$.

In Eq. (4), the fault hypotheses to be monitored are established following ARAIM practice (Blanch et al., 2015), with an additional consideration of the difference between the snapshot filter used in ARAIM and the recursive filter used for OD. A

recursive estimator departs from this snapshot count in two ways. First, its current estimate depends on every GNSS satellite incorporated over the filter's memory. Thus, the set that must be hypothesized as faulty extends beyond the satellites in view at the moment of message generation to include those that have recently set, whose earlier faulty measurements still bias the estimate. Second, the per-mode fault priors P_{H_i} are not constant, as in a snapshot test, but grow roughly in proportion to the interval over which each satellite has contributed to the filter. Both effects enlarge the set of fault hypotheses that must be explicitly monitored relative to the snapshot count, because the fixed reserve P_{NM} covers a smaller fraction of them.

Comparing Eqs. (2) and (4) shows that the definition of the URA and that of the HMI event in the URE domain have the same form. Both bound the probability that the URE exceeds a threshold given that the cascaded GNSS fault monitor raises no alert, and both are held to P_{CGF}. The two coincide when the NTE threshold of Eq. (2) equals the protection level of Eq. (4). We therefore define the URA from the protection level that satisfies Eq. (4),

$$\mathrm{URA} = \frac{\mathrm{PL}_{\mathrm{URE}}}{\mathrm{Q}^{-1}(P_{\mathrm{CGF}}/2)}. \tag{6}$$

The URA computed from Eq. (6) thus satisfies the bound in Eq. (2). The protection level itself depends on the monitor, which this paper realizes as a URE-domain SS monitor.

2. URE-Domain SS Monitor

The URE-domain SS monitor detects cascaded GNSS faults entirely onboard. The broadcast URE implied by the all-in-view OD solution is compared against a family of subset solutions, each formed by excluding one or more GNSS satellites from the onboard filter. A fault in one or more GNSS satellites displaces the all-in-view solution relative to the subset that excludes the faulted satellites, and a large separation between them indicates the fault. This is the SS principle of ARAIM (Joerger et al., 2014; Blanch et al., 2015), applied to the scalar broadcast URE rather than to a user position solution.

For each fault hypothesis H_k, the separation statistic is the difference between the subset and all-in-view predicted URE estimates,

$$u^0 = \boldsymbol{\ell}^\top \hat{\mathbf{x}}^0, \qquad u^k = \boldsymbol{\ell}^\top \hat{\mathbf{x}}^k, \qquad \Delta_k = u^k - u^0, \tag{7}$$

where $\hat{\mathbf{x}}^0$ and $\hat{\mathbf{x}}^k$ are the all-in-view and subset-k state estimates and u^0, u^k their line-of-sight URE projections. A fault is declared when $|\Delta_k|$ exceeds a detection threshold,

$$T_k = K_{\mathrm{FA}}\, \sigma^k_{\mathrm{ss}}, \qquad K_{\mathrm{FA}} = \mathrm{Q}^{-1}\!\left(\frac{P_{\mathrm{FA}}}{2\, N_{\mathrm{modes}}}\right), \tag{8}$$

in which σ^k_{ss} is the standard deviation of the separation statistic Δ_k and K_{FA} is the false-alert multiplier, set so that the continuity allocation P_{FA}, shared among the N_{modes} monitored hypotheses and the two sides of the test, is met.

Under this monitor, the protection level follows by expanding the bracketed sum of Eq. (4) over the fault hypotheses and bounding each fault-conditioned term with the triangle inequality, which gives

$$2\,\mathrm{Q}\!\left(\frac{\mathrm{PL}_{\mathrm{URE}}}{\sigma^0_{\mathrm{URE}}}\right) + \sum_{k=1}^{N_{\mathrm{modes}}} P_{H_k}\, \mathrm{Q}\!\left(\frac{\mathrm{PL}_{\mathrm{URE}} - T_k}{\sigma^k_{\mathrm{URE}}}\right) + P_{\mathrm{NM}} = \eta\, P_{\mathrm{CGF}}, \tag{9}$$

where the first term is the fault-free contribution and each summand is the residual HMI probability under fault hypothesis H_k, weighted by its prior P_{H_k}. The quantities σ^0_{URE} and σ^k_{URE} are the all-in-view and subset-k one-sigma UREs obtained by applying Eq. (3) to the respective filter covariances, T_k is the detection threshold of Eq. (8), and η is the no-alert lower bound of Eq. (5), evaluated with the same P_{FA} that fixes K_{FA}. This expression is the SS-RAIM form of Joerger et al. (2014), apart from the domain in which it is applied and the factor η, which converts the joint ARAIM bound into the conditional requirement of Eq. (2). Under the overbounded URE covariances and the SS monitor, the bracketed sum reduces to a sum of Gaussian tail probabilities, and the protection level is $\mathrm{PL}_{\mathrm{URE}}$ satisfying Eq. (9).

These standard deviations take the same form as in snapshot ARAIM and differ only in their source. In the onboard system they come from the reduced-dynamic OD filter and its subset filters, each excluding one or more GNSS satellites, rather than from an instantaneous measurement geometry (Joerger et al., 2014).

VI. SIMULATION RESULTS

1. Simulation Scenario

We evaluate the framework using simulated data: a high-fidelity truth orbit, synthetic dual-frequency GNSS observations, and the onboard reduced-dynamic filter that produces the broadcast orbit and clock. Two LEO scenarios bracket the dedicated PNT altitude range (Table 1): a high-altitude case, representative of a Sentinel-6A-like orbit (1336 km, 66°), and a low-altitude case, representative of a Swarm-C-like orbit (470 km, 87°) where atmospheric drag is strong. Figure 1 shows the orbital geometry and the resulting GPS and Galileo visibility over a six-hour arc. The higher inclination of Swarm-C extends its ground track to higher latitudes, while under the Earth-limb visibility model both scenarios continuously track on the order of several tens of GNSS satellites (on average about 38 and 31 for Sentinel-6A and Swarm-C, respectively).

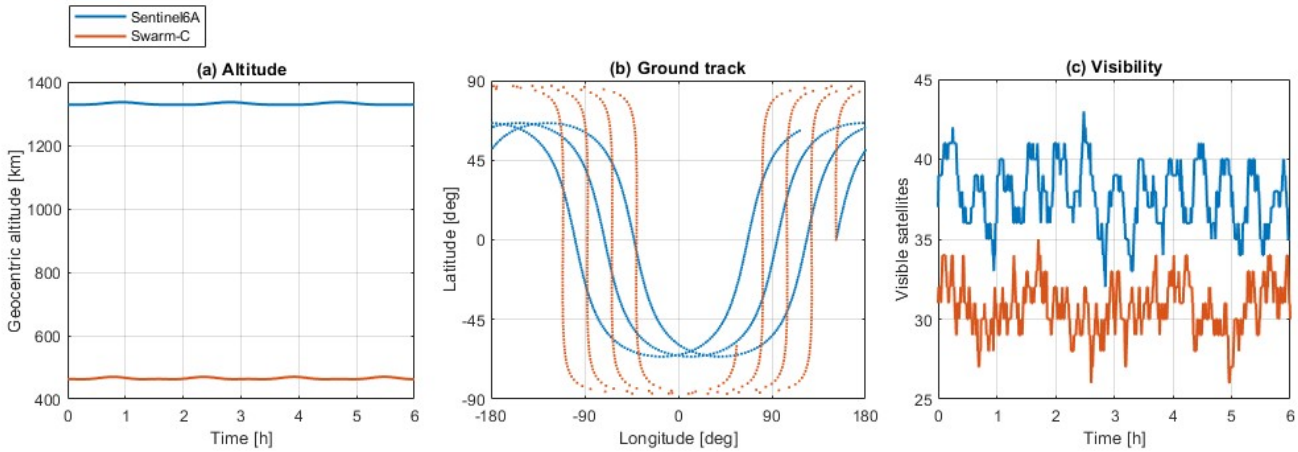

Figure 1: Geometry of the two simulated LEO PNT scenarios over a six-hour arc: (a) geocentric altitude, (b) ground track, and (c) number of visible GPS and Galileo satellites. The high-altitude *Sentinel-6A* scenario (1336 km, 66°) and the low-altitude *Swarm-C* scenario (470 km, 87°) are overlaid.

Table 1: Simulated LEO satellite scenarios.

Parameter	*Sentinel-6A*	*Swarm-C*
Altitude	1336 km	470 km
Inclination	66°	87°
Eccentricity	0.001	0.001
Mass	1180 kg	384 kg
Cross-sectional area	9.8 m^2	1.0 m^2
Drag coefficient C_D	2.2	2.2
SRP coefficient C_R	1.3	1.3

The truth orbit is integrated under a degree-120 GGM03S geopotential (Tapley et al., 2007) with Harris–Priester atmospheric drag and cannonball solar-radiation pressure (Montenbruck and Gill, 2000). The onboard filter deliberately adopts a reduced force model (degree-70 gravity and a lower assumed solar flux). Synthetic ionosphere-free code and carrier-phase observations are generated from precise GPS and Galileo products, and the onboard EKF estimates the satellite position and velocity, the receiver clock bias and drift, empirical accelerations, and atmospheric-drag and SRP scale factors, following the reduced-dynamic approach of Kunzi et al. (2026). The protection level is evaluated in the URE domain for a fixed ground user at Stanford. This single-user case demonstrates the framework. Extending the protection level to the worst-case user location over the satellite coverage region, as the broadcast URA definition in Section III requires, is left to future work. The remaining simulation, filter, and integrity parameters are summarized in Table 2.

Table 2: Simulation, filter, and integrity configuration (common to both scenarios).

Parameter		Value
Force model	Truth	GGM03S gravity (deg. 120), Harris–Priester drag ($F_{10.7} = 200$), cannonball SRP
	Filter	GGM03S gravity (deg. 70), Harris–Priester drag ($F_{10.7} = 150$), cannonball SRP (intentional mismatch with truth)
Estimated states		position, velocity, receiver clock (bias and drift), empirical acceleration (radial/along-/cross-track), C_D, C_R
Clock-drift process noise		1×10^{-4} or 1×10^{-3} m/s
GNSS measurements		GPS + Galileo; Ionosphere-free code and carrier measurements
Target P_{CGF}		1×10^{-7}
GNSS Fault prior		1×10^{-5} / 1×10^{-8} per hour (per-satellite / per-constellation)
Reference ground point for line of sight vector		Stanford University (37.43°N, 122.17°W, 30 m)

2. OD and Prediction Performance

The URE protection level of Section V is built on the predicted one-sigma URE of Eq. (3), which is valid only if the onboard filter covariance is both accurate and a conservative overbound of the true error. We verify these two properties for the orbit and clock states that form the URE of Eq. (1), first for the OD estimate and then as that estimate is propagated over the ephemeris validity interval.

Figures 2 and 3 show the real-time OD accuracy. After the filter converges, the three-dimensional orbit error is about 0.3 cm root-mean-square (RMS) for *Sentinel-6A* and 1.1 cm for *Swarm-C*, and the receiver clock-bias error remains at the centimeter level within its formal $\pm\sigma$ envelope. The formal standard deviation tracks the realized error in both the orbit and the clock channel, so the two contributors to the URE of Eq. (1) are estimated to centimeter accuracy.

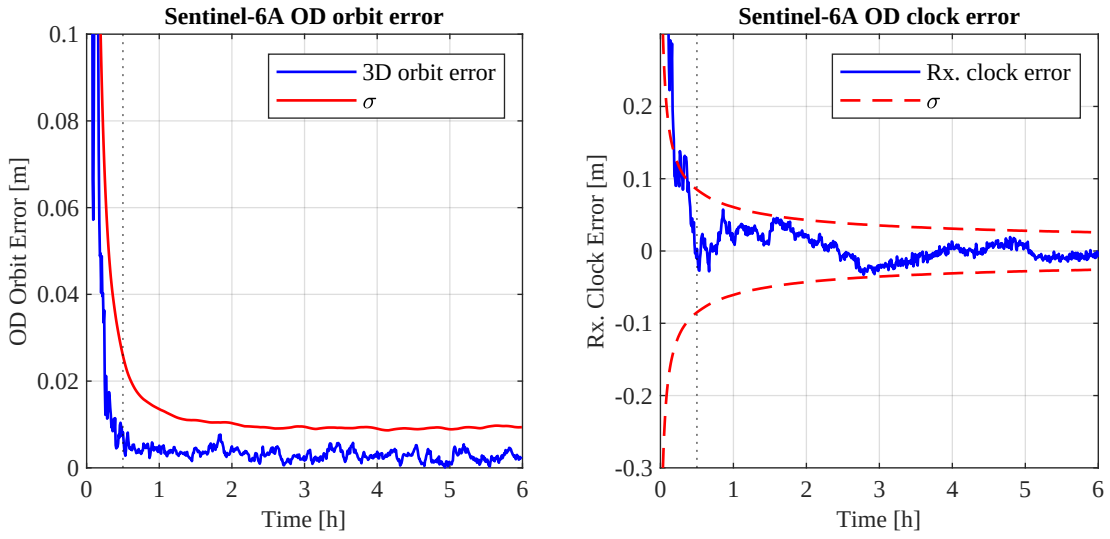

Figure 2: Real-time OD accuracy for the *Sentinel-6A* scenario: (left) three-dimensional orbit error, and (right) receiver clock-bias error.

Accuracy alone does not guarantee a valid protection level; the covariance must also overbound the error. Figure 4 reports the Normalized Estimation Error Squared (NEES) of the real-time estimate (Bar-Shalom et al., 2001), which tests the consistency of the full seven-state covariance against the realized error. For both scenarios the per-epoch NEES remains within the χ^2_7 confidence band, and the averaged NEES is 0.77 and 1.20, well below the expected value of seven. The covariance is therefore under-confident, that is, conservative. This NEES check does not by itself establish that the covariance is conservative enough to ensure integrity at the target P_{CGF}, since that would require validating the deep tails of the error distribution. With the limited simulation data available, it serves as a practical confirmation that the covariance is conservative.

Because the broadcast message is generated from *predicted* states, the covariance must remain a conservative overbound as

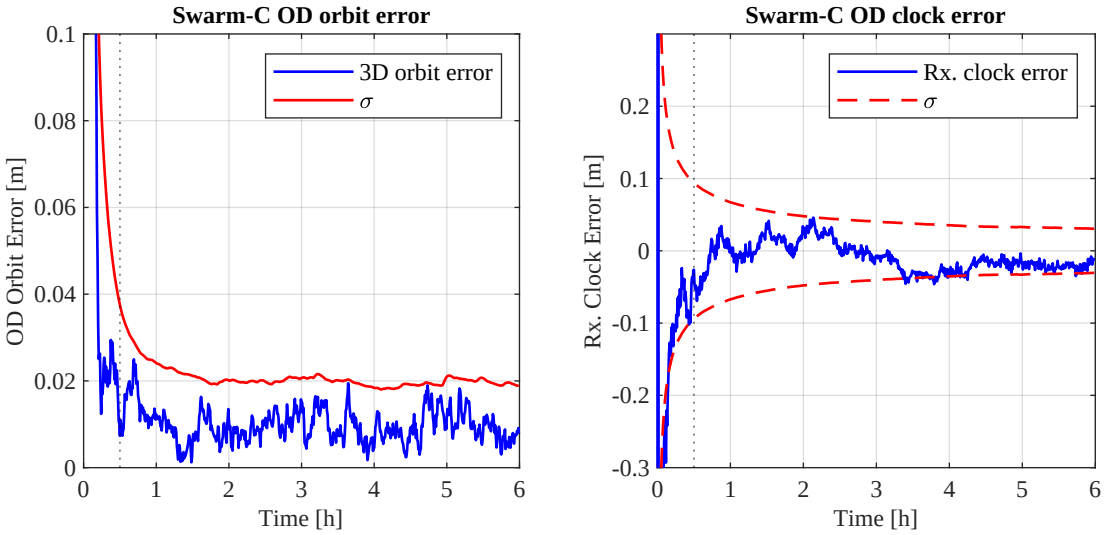

Figure 3: Real-time OD accuracy for the *Swarm-C* scenario, in the same format as Figure 2. The stronger drag at low altitude raises the converged orbit error to about 1.1 cm.

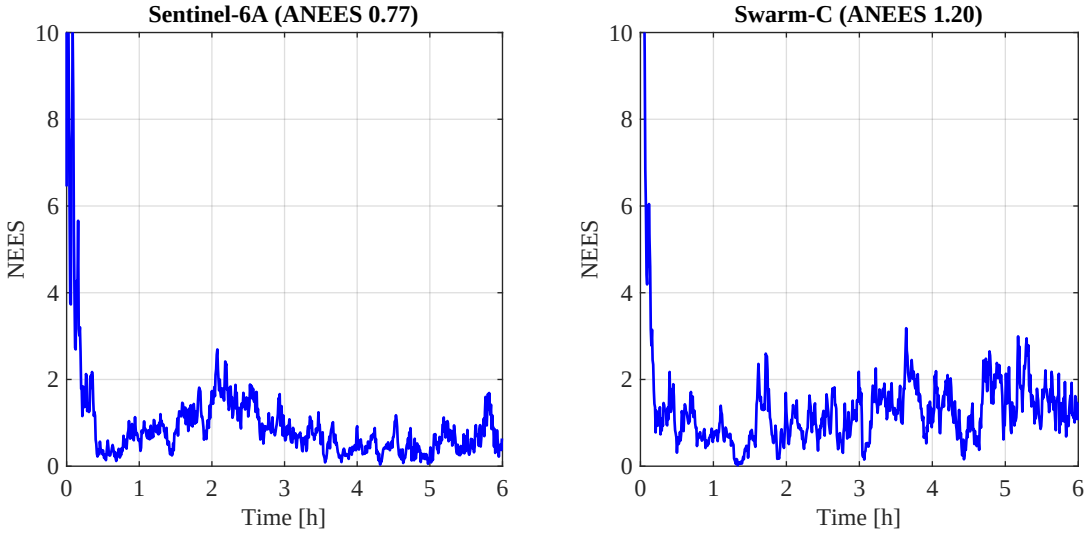

Figure 4: NEES (seven states) of the real-time estimate versus time, for (left) *Sentinel-6A* and (right) *Swarm-C*. The averaged NEES lies below the consistency value of seven, indicating a conservative covariance.

the message ages over the twenty-minute prediction period used in this paper. Figure 5 shows the NEES as a function of the prediction age. It is obtained by propagating the converged filter forward and comparing the predicted covariance against the realized prediction error over an ensemble of generation epochs. The NEES is largest at generation, about 0.8 for *Sentinel-6A* and 1.1 for *Swarm-C*, and stays near or below one at every age, far below the over-confidence bound, as the predicted covariance grows at least as fast as the prediction error. The predicted one-sigma URE of Eq. (3) thus remains a conservative input across the full twenty-minute validity interval.

3. Protection Level and URA

With the predicted URE covariance validated, we evaluate the URE protection level and the broadcast URA over the ephemeris validity interval, for the target $P_{\mathrm{CGF}} = 10^{-7}$ of Table 2 and a reserve $P_{\mathrm{NM}} = 9 \times 10^{-8}$ for the GNSS fault combinations not monitored explicitly.

Figure 6 shows the URE protection level of Eq. (9) under the SS monitor and the realized URE as a function of prediction age

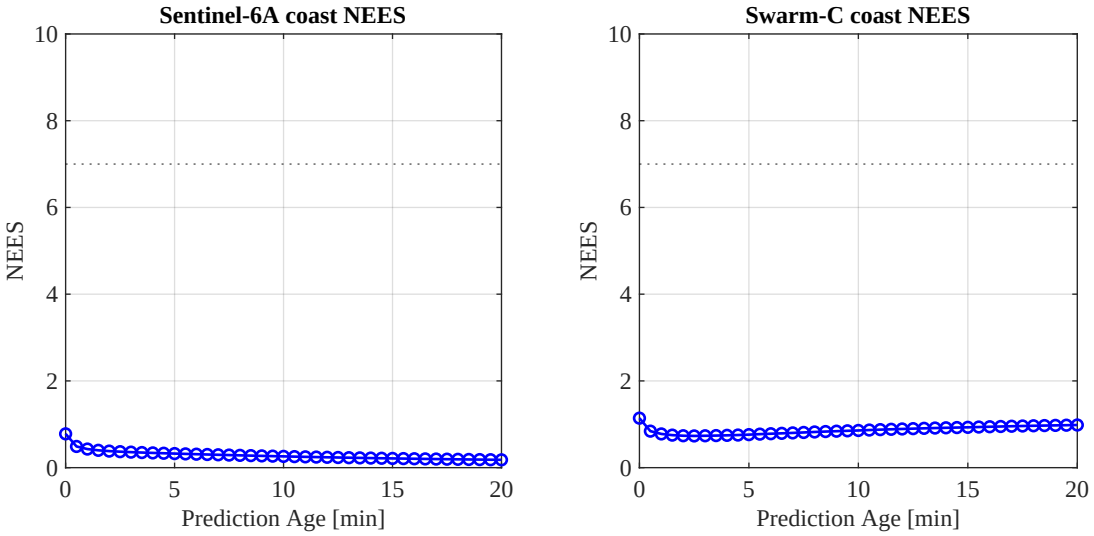

Figure 5: NEES versus prediction age, for (left) *Sentinel-6A* and (right) *Swarm-C*. The averaged NEES stays below the over-confidence bound at every age, confirming that the predicted covariance remains a conservative overbound over the validity interval.

for both scenarios. The protection level grows from about 0.5 m at generation to about 21 m at twenty minutes, following the growth of the predicted covariance. It bounds the realized URE, which stays below 1 m, by a large margin at every age.

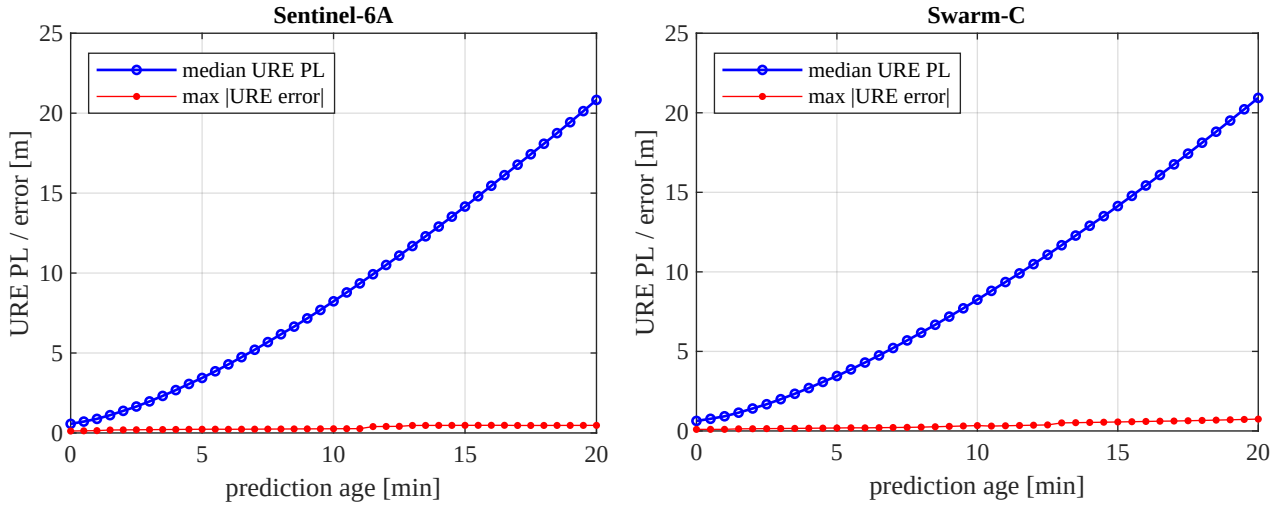

Figure 6: URE protection level under the onboard SS monitor and realized URE versus prediction age, for (left) *Sentinel-6A* and (right) *Swarm-C*. The protection level bounds the realized error by a large margin at every age.

Within Eq. (9), the fault-hypothesis terms contribute little to this protection level, which follows from the strong dynamical constraint on the onboard OD filter. Because the LEO orbit dynamics tightly constrain the estimate, excluding one or more GNSS satellites alters the solution only slightly, and the subset one-sigma UREs σ^k_{URE} differ little from the all-in-view σ^0_{URE}. The detection thresholds T_k of Eq. (8), set by the difference between the all-in-view and subset filters, are correspondingly small. The protection level is therefore set primarily by the fault-free term and the all-in-view one-sigma URE σ^0_{URE}.

The growth of the protection level reflects that of the predicted one-sigma URE. Figure 7a decomposes σ_{URE} of Eq. (3) into its orbit- and clock-induced contributions for the *Sentinel-6A* scenario. The two are comparable over the interval, each reaching about 2.3–2.5 m at twenty minutes with the clock term marginally larger, so the protection level is determined jointly by orbit- and clock-prediction uncertainty rather than by either alone.

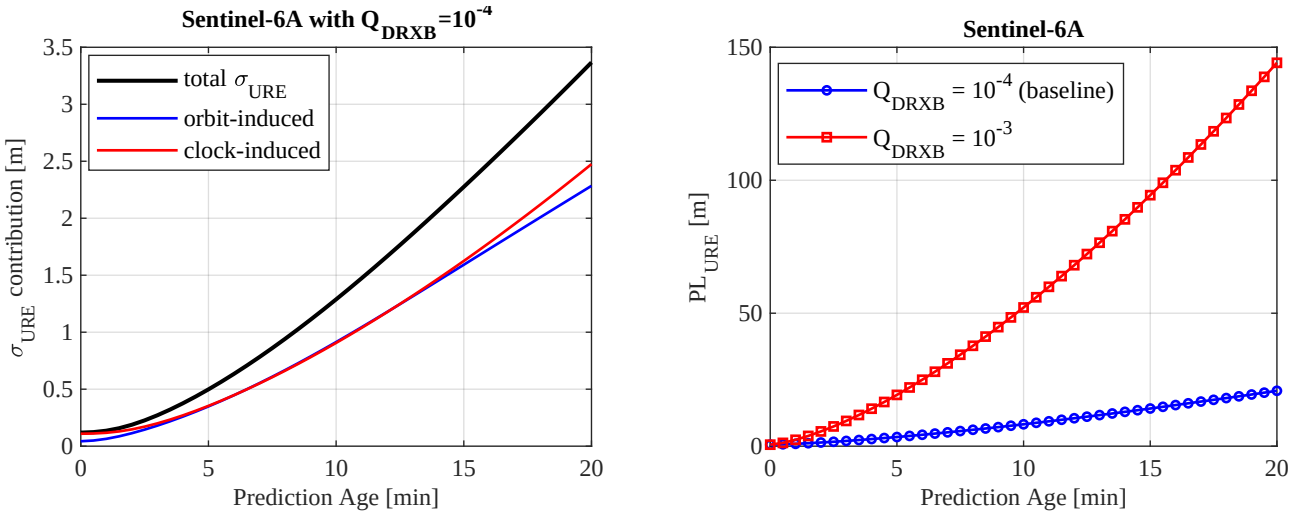

(a) **(b)**

Figure 7: Predicted URE behavior versus prediction age for the *Sentinel-6A* scenario: (a) decomposition of the predicted one-sigma URE into orbit- and clock-induced contributions, which are comparable and each reach about 2.3–2.5 m at twenty minutes; and (b) sensitivity of the URE protection level to the clock-drift process noise, reaching about 21 m at twenty minutes for the baseline 10^{-4} m/s versus about 144 m for 10^{-3} m/s.

Unlike the orbit error, whose growth the satellite dynamics limit over the validity interval, the predicted clock has no comparable dynamical constraint and depends directly on the stability of the onboard oscillator. The clock contribution to the URE is determined by the assumed clock-drift process noise, the dominant contributor to the predicted clock covariance over the validity interval. Figure 7b shows the sensitivity of the protection level to this parameter. The baseline value of 10^{-4} m/s is consistent with the manufacturer-specified short-term stability of high-stability space oscillators, such as ultra-stable oscillators (USOs) and oven-controlled crystal oscillators (OCXOs), on the order of 0.1 mm/s at one second (Mayank et al., 2025). Disclosed dedicated LEO PNT architectures indicate this clock class, disciplining a stable oscillator to GNSS time instead of relying on a high-stability onboard atomic standard (Kunzi and Montenbruck, 2022). At this value the protection level reaches about 21 m at twenty minutes. At 10^{-3} m/s, an order of magnitude less stable, the predicted clock covariance grows much faster and the protection level grows to about 144 m. The baseline value is required to keep the broadcast URA within an operationally useful range, and is the value adopted in Table 2.

Finally, the broadcast URA follows from the protection level through Eq. (6). Table 3 reports the protection level and the resulting URA at representative ages for both scenarios. The URA grows from about 0.1 m at generation to about 3.9 m at the end of the twenty-minute interval. These are the integrity-ensured URA values the framework produces for the target $P_{\mathrm{CGF}} = 10^{-7}$.

Table 3: URE protection level (PL) and broadcast URA versus ephemeris age of data under the onboard SS monitor ($P_{\mathrm{CGF}} = 10^{-7}$, $P_{\mathrm{NM}} = 9 \times 10^{-8}$).

Scenario	**Age [min]**	PL [m]	URA [m]
Sentinel-6A	0	0.57	0.11
	10	8.23	1.55
	20	20.83	3.91
Swarm-C	0	0.63	0.12
	10	8.25	1.55
	20	20.94	3.93

Despite the stronger atmospheric drag at low altitude, the two scenarios yield nearly identical protection levels and URAs

(Table 3). The realized prediction error is larger at low altitude, by about 26% at twenty minutes, consistent with its more demanding dynamics. The protection level, however, depends not on the realized error but on the formal predicted one-sigma URE, σ_{URE}, of Eq. (3). During propagation, the growth of σ_{URE} is dominated by the reduced-dynamic process noise, the empirical-acceleration and clock-drift terms, which is identical in the two scenarios. σ_{URE} therefore differs by less than 1% at twenty minutes, and the clock contribution, determined by the same clock-drift process noise and comparable in magnitude to the orbit contribution, narrows the difference further. The equal protection levels remain valid for both scenarios because the process noise is set conservatively to absorb dynamic mismodeling: σ_{URE} overbounds the realized error by a factor of roughly 2.3–2.9, so the larger realized error at low altitude still lies within this margin. The low-altitude scenario thus has a smaller protection margin, with the same protection level bounding a larger realized error, but its realized URE remains bounded.

VII. CONCLUSION

This paper addressed URA generation for dedicated LEO PNT satellites that generate their broadcast ephemeris onboard from GNSS measurements. Onboard generation makes the broadcast URE dependent on the GNSS measurements used in the onboard OD, introducing fault modes absent from conventional GNSS, where the navigation message is generated and validated on the ground. Among these, this paper treated the cascaded GNSS fault, in which a fault in a GNSS satellite propagates through the OD of every LEO satellite that tracks it. Because this fault affects several LEO satellites simultaneously, it belongs in the constellation fault probability P_{const} of the LEO constellation, and we assigned it the allocation P_{CGF}. We adapted the protection-level concept of civil-aviation user integrity to the URE domain. An onboard URE-domain SS monitor detects cascaded GNSS faults, and a URE-domain protection level is computed from the worst case over the faults the monitor leaves undetected. The broadcast URA follows from this protection level and keeps the probability of an unalerted URE exceedance from this cause at or below P_{CGF}, provided that the predicted URE covariance overbounds the error. The effect of undetected faults is thereby absorbed into the broadcast URA instead of being transferred to users as an elevated constellation fault probability. The framework was evaluated in simulation over representative high- and low-altitude LEO scenarios with synthetic GPS and Galileo observations. The predicted one-sigma URE covariance remained conservative relative to the realized error across the twenty-minute ephemeris validity interval, and the resulting protection level bounded the realized URE by a large margin. For a target $P_{\mathrm{CGF}} = 10^{-7}$, the URA grew from about 0.1 m at generation to about 3.9 m at twenty minutes.

The protection level here is evaluated for a single ground user, whereas the broadcast URA of Section III is defined at the worst-case user location over the satellite coverage region. Extending the evaluation to that worst case is required before the reported values can be broadcast as URA. The remaining fault modes of a LEO satellite, and their allocations within P_{sat} and the remainder of P_{const}, must be integrated with the cascaded GNSS fault before a complete broadcast URA can be stated. The framework should also be validated against real GNSS observations collected onboard LEO satellites in orbit. Finally, the sensitivity of the protection level to the quality and latency of the GNSS orbit and clock products used onboard will be analyzed.

ACKNOWLEDGEMENTS

This work was supported by the National Research Foundation of Korea (NRF) grants funded by the Korean government (MSIT) (Grant No. RS-2024-00354326, Project Title: Development of Integrity-Assured Navigation System for AAM Based on Hyper-Connected Network, 2025–2026, contribution 50%; Grant No. RS-2025-02213804, Project Title: In-Space Assembly & Manufacturing Research Centre (ISMRC), 2025–2026, contribution 50%).

REFERENCES

Bar-Shalom, Y., Li, X. R., and Kirubarajan, T. (2001). *Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software*. John Wiley & Sons, New York. https://doi.org/10.1002/0471221279.

Blanch, J., Walter, T., Enge, P., Lee, Y., Pervan, B., Rippl, M., Spletter, A., and Kropp, V. (2015). Baseline advanced RAIM user algorithm and possible improvements. *IEEE Transactions on Aerospace and Electronic Systems*, 51(1):713–732. https://doi.org/10.1109/TAES.2014.130739.

Blanch, J., Walter, T., Milner, C., Joerger, M., Pervan, B., and Bouvet, D. (2022). Baseline advanced RAIM user algorithm: Proposed updates. In *Proceedings of the 2022 International Technical Meeting of The Institute of Navigation (ION ITM 2022)*, pages 229–251. Institute of Navigation. https://doi.org/10.33012/2022.18254.

Braff, R., Bian, B., and Shively, C. (2012). Independent URA monitor for GPS IIIC. *NAVIGATION*, 59(1):37–49. https://doi.org/10.1002/navi.1.

DeCleene, B. (2000). Defining pseudorange integrity – overbounding. In *Proceedings of the 13th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2000)*, pages 1916–1924, Salt Lake City, UT. Institute of Navigation. https://www.ion.org/publications/abstract.cfm?articleID=1603.

Gioia, C., Menzione, F., Piccolo, A., Casotto, S., and Bardella, M. (2025). T-RAIM for precise orbit determination in LEO-PNT. *Sensors*, 25(23):7322. https://doi.org/10.3390/s25237322.

Joerger, M., Chan, F.-C., and Pervan, B. (2014). Solution separation versus residual-based RAIM. *NAVIGATION*, 61(4):273–291. https://doi.org/10.1002/navi.71.

Kovach, K., Dobyne, J., Crews, M., and Miles, C. (2008). GPS III integrity concept. In *Proceedings of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2008)*, pages 2250–2257, Savannah, GA. Institute of Navigation. https://www.ion.org/publications/abstract.cfm?articleID=8127.

Kunzi, F. and Montenbruck, O. (2022). Precise onboard time synchronization for LEO satellites. *NAVIGATION*, 69(3):navi.531. https://doi.org/10.33012/navi.531.

Kunzi, F., Montenbruck, O., and Perazzetti, F. (2026). GNSS-based real-time orbit determination and prediction for LEO-PNT ephemeris generation. *Advances in Space Research*, 77(1):923–938. https://doi.org/10.1016/j.asr.2025.11.082.

Mayank, Prol, F. S., Lundén, V., Saleem, Z., Foreman-Campins, G., Sharma, S., Lohan, E. S., Bhuiyan, M. Z. H., Kaasalainen, S., Kuusniemi, H., and Praks, J. (2025). LEO-PNT feasibility aspects: Satellite navigation payload size, weight, and power analysis. *IEEE Access*, 13:110069–110089. https://doi.org/10.1109/ACCESS.2025.3583080.

Montenbruck, O. and Gill, E. (2000). *Satellite Orbits: Models, Methods and Applications*. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-58351-3.

Reid, T. G. R., Neish, A. M., Walter, T., and Enge, P. K. (2018). Broadband LEO constellations for navigation. *NAVIGATION*, 65(2):205–220. https://doi.org/10.1002/navi.234.

Tapley, B. D., Ries, J. C., Bettadpur, S., Chambers, D., Cheng, M., Condi, F., and Poole, S. (2007). The GGM03 mean earth gravity model from GRACE. In *AGU Fall Meeting Abstracts*, pages G42A–03. Abstract G42A-03.

U.S. Department of Defense (2020). Global positioning system standard positioning service performance standard. 5th edition, U.S. Department of Defense, Washington, DC. https://www.gps.gov/technical/ps/2020-SPS-performance-standard.pdf.

Walter, T., Blanch, J., Gunning, K., Joerger, M., and Pervan, B. (2019). Determination of fault probabilities for ARAIM. *IEEE Transactions on Aerospace and Electronic Systems*, 55(6):3505–3516. https://doi.org/10.1109/TAES.2019.2909727.

Wang, K., El-Mowafy, A., and Rizos, C. (2022a). Integrity monitoring for precise orbit determination of LEO satellites. *GPS Solutions*, 26(1):32. https://doi.org/10.1007/s10291-021-01200-4.

Wang, K., El-Mowafy, A., and Yang, X. (2022b). URE and URA for predicted LEO satellite orbits at different altitudes. *Advances in Space Research*, 70(8):2412–2423. https://doi.org/10.1016/j.asr.2022.08.039.