Economical Fused LEO GNSS

Peter A. Iannucci and Todd E. Humphreys
Radionavigation Laboratory
The University of Texas at Austin
Austin, TX, USA
peter.iannucci@austin.utexas.edu, todd.humphreys@utexas.edu

***Abstract*—In addition to Internet service, new commercial broadband low-Earth-orbiting (LEO) satellites could provide a positioning, navigation, and timing (PNT) service far more robust to interference than traditional Global Navigation Satellite Systems (GNSS). Previous proposals for LEO PNT require dedicated spectrum and hardware: a transmitter, antenna, and atomic clock on board every broadband satellite. This paper proposes a high-performance, low-cost alternative which fuses the requirements of PNT service into the existing capabilities of the broadband satellite. A concept of operations for so-called fused LEO GNSS is presented and analyzed both in terms of positioning performance and in terms of the economy of its use of constellation resources of transmitters, bandwidth, and time. This paper shows that continuous assured PNT service over $\pm 60^\circ$ latitude (covering 99.8% of the world's population) with positioning performance exceeding traditional GNSS pseudoranging would cost less than 2% of system capacity for the largest new constellations, such as SpaceX's Starlink or Amazon's Project Kuiper.**

***Index Terms*—broadband; satellite; navigation.**

I. Introduction

Anticipation is building for positioning, navigation, and timing (PNT) technologies based on commercial broadband low Earth orbiting (LEO) satellite constellations. These space vehicles are far nearer and more numerous than traditional global navigation satellite systems (GNSS) in medium (MEO) or geostationary orbit (GEO), and their communications transponders have both exceedingly high gain and access to a vast allocation of spectrum. Potential commercial LEO PNT signals are thus more precise, powerful, survivable, and jam-resistant than traditional GNSS.

A. Background

All current GNSS have settled into a design in which the constellation consists of a few dozen satellites primarily at MEO whose signals are low-data-rate direct-sequence spread spectrum waveforms confined to L band and occupying less than a total of $125\,\mathrm{MHz}$.

There are good reasons behind these architectural choices. The first GNSS, TRANSIT, was a LEO system that employed Doppler positioning using narrowband signals in UHF. But the trade studies from which GPS was later conceived revealed that a MEO system with wideband signals would be more resistant to jamming and would be capable of satellite-redundant instantaneous positioning with only a few dozen satellites. L band was chosen because it wavelengths are short enough for ionospheric transparency, yet long enough to avoid attenuation due to rainfall and water vapor [2]–[4].

For all its advantages, the architecture into which existing GNSS have settled also has significant drawbacks. Non-GNSS uses of the congested space-to-Earth spectrum in L band have prevented allocation of much greater bandwidth for GNSS in that band. Constellation survivability is limited by the small number of satellites (also called space vehicles, or SVs), which make attractive targets for anti-satellite warfare [5]; jamming immunity is limited by the weakness of the signals, which, being diffused over an entire hemisphere, are easily overwhelmed [6], [7]; and positioning precision is limited by both signal weakness and modest bandwidth, which place information-theoretic lower bounds on ranging uncertainties [2].

Despite a pressing need for much greater GNSS robustness in contested environments, GNSS evolution over the past few decades has been more incremental than radical. New signals have been introduced at different frequencies, with binary offset carrier (BOC) waveforms that more efficiently allocate signal power. But GNSS remains principally MEO, L-band, and confined to less than $125\,\mathrm{MHz}$. Given tight budgets and enormous design inertia due to the need for backward compatibility, radical changes in traditional GNSS over the next 30 years are unlikely. Spot beams, the most promising feature of the GPS III program for improved jamming immunity [8], have been abandoned. Calls to introduce new GNSS signals in C band (e.g., [9]) have not gained traction. In short, as regards robustness to interference for both civil and military users, traditional GNSS is inadequate and has stagnated.

B. LEO GNSS

LEO GNSS based on commercial broadband mega-constellations may be a way forward. Such constellations are already being developed and deployed to provide global broadband Internet service. They also represent a promising platform for LEO GNSS: Although not originally intended for PNT, these satellite constellations are designed for rapid technological refresh via software or hardware, and so may be adaptable for PNT.

Unlike traditional GNSS, in which costs are borne by nation-states and service is free-of-charge, commercial GNSS providers would have to pass costs on to users. For such a scheme to be viable, it would have to be *economical*: that is, it would have to offer fundamental advantages over traditional (free) GNSS commensurate with the price tag, or else there would be no demand; and the service would have to be

 Preprint of the 2020 IEEE/ION Position Location and Navigation Symposium St. Louis, MO, September 23–25, 2020

sufficiently inexpensive to provide, or else there would be no supply. This paper will explore both facets of this problem.

C. Hosted Payload LEO GNSS

In their groundbreaking work, Reid et al. [10], [11] analyzed the performance of potential LEO GNSS implemented using *hosted payloads*: dedicated PNT hardware onboard each satellite. There are excellent reasons to explore a hosted payload solution. Such payloads are modular, independent of the satellite's primary communications mission, and may be iterated and upgraded for future launches. More recent work envisions a dedicated LEO GNSS constellation [12] with similar advantages. As laid out by Reid et al., hosted PNT signals provide continuous omnidirectional global coverage and may be incorporated into user pseudorange navigation equipment nearly as readily as traditional GNSS signals. Reid et al. estimate that the system would enjoy a $30\,\mathrm{dB}$ improvement in signal-to-noise ratio, and thus resistance to jamming, over traditional GNSS.

A hosted payload approach along those lines is not radically dissimilar to traditional GNSS. No theoretical obstacle bars the way. However, space hardware development is costly and challenging as a practical matter. And a hosted payload would be costly: besides the cost of each payload, there are costs associated with renting space and hookups on the host satellite, costs for running necessary radiofrequency interference (RFI) and compatibility testing, and both costs and risks of delay in securing the necessary frequency allocation.

D. Fused LEO GNSS

This paper will introduce an alternative *fused* approach to LEO GNSS, whereby the hardware already designed and the spectrum already allocated for the satellites' primary broadband mission is dual-purposed for PNT. Fused LEO GNSS sacrifices nothing in performance while eliminating the costs of special-purpose on-orbit hardware. In fact, where previous proposals targeted positioning precision on-par with traditional GNSS pseudoranging (on the order of $3\,\mathrm{m}$), this paper will show that fused LEO GNSS can improve on this by more than an order of magnitude. Moreover, it offers a significant anti-jam advantage over L-band hosted-payload solutions in terms of signal-to-interference ratio, thus making it attractive as a means for delivering assured PNT (A-PNT).

These strengths are made possible by a confluence of factors.

First, one may exploit the plentiful data bandwidth present in each broadband satellite transmission burst for up-to-the-instant orbit and clock products. If such zero-age-of-ephemeris products are available, then the need for atomic clocks in LEO may be eliminated. Second, one may obtain these orbit and clock products by performing precision orbit determination (POD) on-orbit using traditional GNSS in a multi-tier architecture [10]. This eliminates the need for an extensive ground segment to observe satellite orbits. Third, unlike traditional L-band services, commercial broadband signals in K-band and V-band will have both high signal-to-noise-ratio (SNR) and large bandwidth. This greatly reduces receiver noise and multipath as a source of user ranging error, even though no special ranging code is used over the communications link. Furthermore, because these signals have a much shorter wavelength, it is possible to build a (relatively) compact, highly-directional receiver phased array for an additional $30\,\mathrm{dB}$ of anti-jam performance.

Characteristic	Traditional GNSS	Hosted [11]	Fused (omni RX)	Fused (array RX)
Single-epoch PNT	✓	✓	✓	✓
Low Earth Orbit		✓	✓	✓
Mega-constellation		✓	✓	✓
On-orbit POD		✓	✓	✓
Non-atomic clocks			✓	✓
Zero age-of-ephemeris			✓	✓
Time multiplexed			✓	✓
Excess bandwidth			✓	✓
Highly directional				✓
Localized power boost	\$\$\$	\$	\$	\$
Precision horz.	3.0 m	3.0 m	37 cm	19 cm
Precision vert.	4.8 m	4.4 m	48 cm	25 cm
Anti-jam advantage	—	+30 dB	+26 dB	+56 dB

TABLE I: Contrasting traditional GNSS, previous hosted proposals, and fused LEO GNSS. Precise orbit determination (POD) here uses either inter-satellite ranging (i.e. autonav capability) or onboard GNSS receivers in LEO (multi-tier GNSS). Positioning precision is 95^{th} percentile in the horizontal and vertical directions.

PNT precision, anti-jam performance, and other constellation characteristics are compared in Table I for traditional GNSS, hosted-payload LEO GNSS, and fused LEO GNSS. In the case where an omnidirectional antenna is preferred, such as in mobile devices, the SNR is not high enough to download ephemeris, and a back-up comms link such as cellular data service would be required. It is also worth noting that certain design elements that give fused LEO GNSS its performance could be incorporated into hosted payload proposals. However, this paper will only compare against prior proposals as they are currently understood.

E. Economical Operation

To be viable, a fused LEO GNSS service must be cost-effective for providers. This is the case: this paper will show that providing PNT service to every user in one cell (e.g., for the Starlink constellation, a hexagon of $1315\,\mathrm{km}^2$) is roughly as costly, in terms of satellites' time spent providing PNT signals and re-orienting their antennas, as a single 1 MB/s downlink stream. This paper will further show how the system can flexibly accommodate users willing to spend more to buy improved accuracy or resistance to jamming.

F. Contributions

This paper makes three primary contributions. First, it introduces the concept of fused LEO GNSS, an architecture in which existing clocks, modems, antennas, and spectrum of broadband satellite mega-constellations are dual-purposed for PNT. Second, it provides an analysis of the opportunity cost to constellation providers for re-allocating resources to provide a

fused PNT service. Third, it lays out an analysis of achievable positioning precision for the fused architecture. The remainder of the paper is organized around these contributions.

II. Concept of Operations

There is tension in trying to achieve a fused architecture. The clocks onboard satellites in broadband satellite mega-constellations are not likely to be stable enough to allow for nanosecond-accurate clock forecasting beyond a minute. The modems are not designed to produce traditional ranging waveforms. The highly-directional antennas are scheduled to be steered so that only one satellite serves a given cell at any given time, complicating single-epoch pseudorange positioning. The K-band spectrum set aside for commercial LEO broadband is shared with terrestrial and geostationary services, and suffers greater atmospheric attenuation than L band. Moreover, the satellites have resource constraints in (A) downlink scheduling, (B) antenna array re-phasing, (C) electrical power, and (D) command-and-control bandwidth that must be respected. Any resources diverted from the constellation's primary communications mission must be paid for by PNT service users.

Section III will consider these resource costs in detail.

There are three points worth re-iterating. Fused LEO GNSS requires the PNT user to adapt to the characteristics of the broadband system: burst-type, rather than continuous emissions; directivity at the satellite, so that each burst is only detectable within a few kilometers of its intended target; and, optionally, directivity at the receiver, so that the receiver must decide in advance to listen for a particular satellite.

A. Orbit and Clock Determination

In both the hosted payload and the fused approaches to LEO PNT, it is necessary to obtain highly-accurate continual measurements of the orbital ephemerides and the phases of the spaceborne clocks. For this reason, each satellite will perform onboard precision orbit determination using, for the time being, onboard GNSS receivers. This may be expected to constrain forecasting uncertainties for clocks and orbits to $14\,\mathrm{cm}$ RMS after one second.

One has essentially three tools: constellation-to-ground ranging, intra-constellation ranging, and onboard GNSS receivers. Ideally, one should use all available tools.

The functions of such a network could likely be shouldered by the LEO broadband provider's own ground station network, but wide oceans and relatively small service footprints below each satellite make constellation-to-ground ranging an insufficient tool in isolation. Rapid measurements and rapid ephemeris updates are required (Fig. 7), unless the LEO broadband provider makes the unlikely decision to use full-scale atomic clocks.

In the near term, the nature and capabilities of the intra-constellation links are in flux. Regarding Starlink in particular, recent public statements by SpaceX officers [13], [14] indicate that their first-generation satellites lack the optical inter-satellite links slated for later generations.

Without intra-constellation ranging to supplement constellation-to-ground ranging, it would not be possible to build an autonomous (i.e. GNSS-agnostic) LEO PNT system today. The difficult engineering problems of designing an autonomous orbit determination system for a massive LEO constellation are, therefore, left to future work.

Thus, in the near term, LEO PNT will operate in what Reid et al. refer to as a multi-tier GNSS architecture: each satellite will carry a GNSS receiver. This architecture is intermediate in assurance between traditional GNSS and fully autonomous LEO A-PNT: it is reliant on GNSS being available in orbit, and provides highly jam-resistant signals to users on the ground. Jamming or spoofing the GNSS receiver on a satellite in LEO is a far more challenging task than jamming a ground- or air-based receiver, and the high orbital speed of the space vehicles will rapidly carry them out of range of any fixed jammer. The results of the present analysis are for the case of multi-tier GNSS.

B. Modems and Ranging

Service consists of a series of ranging bursts. These bursts are modulated just like ordinary broadband data, with three exceptions: first, the contents of the burst are (largely) known in advance to the ground receiver, so that it is possible to cross-correlate the received burst and the expected waveform to obtain a code-phase pseudo-range measurement; second, the burst is very short, on the order of $500\,\mu\mathrm{s}$; and third, the burst is not acknowledged by one ground receiver (i.e. unicast), but instead is broadcast to all receivers in the cell.

It is well known that broadband data signals ought to be spectrally flat in order to maximize information capacity within a given power and bandwidth budget. For this and other reasons, broadband transmitters may be anticipated to use a spectrally-flat modulation such as OFDM, and to use a spectrally-flat synchronization preamble.

It is also well known that ranging signals ought to be two-peaked in spectrum, concentrating power towards the edges of the channel, in order to maximize ranging precision within a given power and bandwidth budget (Appendix B).

Ideally, one uses a ranging signal for ranging, and a communications signal for communications. For fused LEO GNSS, however, it is greatly disfavored to change the function of the modem, both because modulations tend to be fixed in hardware, and because a high-bandwidth side-channel for orbit and clock data embedded in the ranging burst is highly valuable. Fused LEO GNSS uses the unmodified communications waveform for both data and ranging. There is a small degradation in possible ranging precision arising from this compromise, but it is far less than other sources of error.

A marked difference with traditional GNSS is that even the shortest bursts possible on a broadband LEO data modem are expected to achieve pseudo-ranging uncertainties far smaller than other sources of error in the system: the so-called receiver noise and multipath error is effectively negligible (§A-F). For this reason, it is not a problem to set aside a portion of the ranging burst to contain data not known in advance to the receiver, such as up-to-date clock and orbit ephemeris. This portion of the burst is ignored during correlation.

In order to perform single-epoch pseudorange positioning, it is necessary for multiple satellites to transmit ranging bursts to the same cell e.g. once per second. A receiver collects a pseudorange observation from each satellite by comparing the time-of-arrival of the ranging burst to the clock and orbit ephemeris data modulated into the burst. The receiver combines these observations with its prior beliefs using standard non-linear pseudorange positioning and/or extended Kalman filtering, and provides posterior position and time estimates to the user.

C. Antenna Directionality

LEO signals are fundamentally visible to a much smaller portion of the Earth's surface than are traditional GNSS signals. Furthermore, the requirement of a broadband system for simultaneous high SNR and bandwidth is possible only by steering the majority of a satellite's transmit power into a narrow beam focused on a relatively small ground service region ("cell"). For broadband service, a given cell need only be served by one satellite; but for pseudorange PNT, the satellites must be reprogrammed to provide cross-cell ranging bursts.

By way of example, consider SpaceX's Starlink constellation, eventually slated to comprise between ten and forty thousand satellites. The number of cells is approximately equal to the surface area of the Earth between $\pm 60^\circ$ latitude divided by the area of a hexagon of diameter $45\,\mathrm{km}$: this gives a number on the order of 400,000.

Each satellite has a set of independently-steerable phased array antennas for transmit and for receive operation. According to relevant public FCC filings [15], [16], each Starlink craft will support eight downlink beams for user data service, each with a distinct associated $250\,\mathrm{MHz}$ channel at the low end of the $\mathrm{K_u}$ band. At any given time, each beam may be focused towards a single $45\,\mathrm{km}$ cell. Once a beam is focused, bursts of modulated data may be transmitted to users in the cell. Users in distant cells will not be able to detect these transmissions. A complementary set of uplink phased arrays are also present, operating on a disjoint set of channels in the $\mathrm{K_u}$ and $\mathrm{K_a}$ bands. In this paper, it is assumed that the steering operation of the transmit and receive beams is independent.

D. Spectrum

Due to spatial re-use of spectrum in the K-band, non-geostationary broadband services like Starlink are required to employ careful masking, avoiding interference with higher-precedence terrestrial and geostationary services (§V). Atmospheric attenuation is compensated by antenna gain at the satellite and on the ground: but if a compact or an omnidirectional receiver antenna is required, as in the case of mobile devices, the SNR is reduced by as much as $30\,\mathrm{dB}$ relative to the performance of a phased array. Under these conditions, the signal is still suitable for correlation-based ranging, but not for data transmission. Orbit and clock products may be transmitted via an alternative (e.g. terrestrial) wireless communications network.

III. Cost Model

A. Channel Reservations

Again taking SpaceX's Starlink constallation as an example, based on the same public FCC filings, each cell with one or more paying broadband users is expected to be assigned to one satellite at any given time. There is little value in assigning multiple satellites to a single cell, because one satellite is quite capable of saturating the regulatory power flux limit at ground level. Pointing additional beams at the same cell would thus require each beam to be attenuated in power, undermining the objective of improved spatial re-use.

However, as noted above, it is a requirement for pseudorange PNT that ranging bursts from multiple satellites are provided to a single cell. Let the satellite assigned to a cell be the "primary" satellite, and let all the others be "secondary". The primary must avoid transmitting either broadband or PNT bursts at any time when these bursts might arrive at the ground simultaneously with a secondary PNT burst. Because secondary PNT bursts arrive from a different point in the sky than primary bursts, the time difference of arrival of the two bursts will vary by some tens to hundreds of microseconds across a service region. For this reason, the primary downlink beam must observe a moment of silence somewhat longer than the burst itself to accommodate each secondary PNT burst. (It is not necessary for the primary uplink to be disrupted during this time.)

This disruption in the primary downlink represents a lost opportunity to carry valuable commercial broadband traffic. Similarly, to transmit a secondary PNT burst, one of the secondary satellite's downlink beams must forgo broadband traffic while it steers away from its own cell, transmits the burst, and steers back. (As before, it is not necessary for the secondary uplink to be disrupted during this time.)

Thus, one cost of fused LEO PNT service is the fraction of downlink time which must be set aside by satellites to transmit PNT bursts. An additional cost is the time for which a channel must be set aside at ground level to receive PNT bursts. These quantities may be referred to as the transmit reservation and the receive reservation. The worst-case opportunity cost of providing PNT service is determined by the larger of the two.

Let each PNT burst be transmitted once per $\tau \approx 1\,\mathrm{s}$ (the epoch length). Let the burst length be $t \approx 500\,\mu\mathrm{s}$, the number of satellites serving each cell be $n = 5$, the number of cells $C = 400\,000$, the number of satellites $N = 10\,000$, and the number of beams (channels) per satellite $b = 8$. Let the time-difference-of-arrival penalty for secondary PNT bursts be T, a function of the minimum elevation angle $\phi_0 = 40^\circ$. Let the cell diameter be $D \approx 45\,\mathrm{km}$ and the speed of light c. Then

$$T = \frac{2D\cos(\phi_0)}{c} \approx 230\,\mu\mathrm{s}$$

$$\begin{aligned}\mathrm{mean}(\text{TX reservation}) &= \frac{nCt}{bN\tau} \\ &\approx 1.25\% \end{aligned} \tag{1}$$

$$\begin{aligned}\mathrm{mean}(\text{RX reservation}) &= \frac{(n-1)(t+T)+t}{b\tau} \\ &\approx 0.04\% \end{aligned} \tag{2}$$

B. Phased Arrays

An additional cost which may not be negligible is phased array re-configuration bandwidth. For the primary communications mission, LEO broadband providers have had little reason to optimize the rate at which new phasing coefficients can be loaded into the phased array hardware. If a beam is serving only one cell, the time between phase updates needed to keep the pattern pointed to within tolerance may be measured in seconds. For fused LEO PNT, however, it may be necessary for each array to be steered many times per second.

In this work, it is assumed to be possible to steer an array no more than 200 times per second. This paper assumes that coefficient pre-loading takes $5\,\mathrm{ms}$ but actual array switching is glitch-free and (relatively) instantaneous. This paper also assumes that "switching back" is free: during switching, the banks of memory holding old and new coefficients are swapped without erasing, so that switching back does not require re-loading coefficients. Otherwise, the steering bandwidth utilization must be doubled.

Perhaps unexpectedly, this resource turns out to be the least abundant:

$$\begin{aligned}\mathrm{mean}(\text{steering bandwidth utilization}) &= \frac{1}{200}\frac{nC}{bN} \\ &\approx 12.5\%\end{aligned} \tag{3}$$

C. Electricity

A third quantifiable cost is electrical power. This depends on transmit duty cycle, and potentially also on lost opportunities to place satellite hardware into a "deep sleep" state with the radio powered off. To minimize electrical costs due to loss of sleep, it is better to schedule PNT work to a small number of satellites so that 100% of whichever-is-their-bottleneck-resource is utilized, and put the rest to sleep. With phased array re-configuation bandwidth as the bottleneck, only one out of every eight satellites loses sleep.

D. Command-and-Control Spectrum

Use of command-and-control spectrum is increased by the secondary PNT mission. In this work, it is assumed that each satellite independently performs precision orbit determination (POD) using an on-board GNSS receiver. This requires very little communication, except to provide up-to-date GNSS orbits and clocks. Scheduling of secondary PNT bursts, however, requires coordination between satellites.

For the primary broadband mission, the ground control segment is assumed to continuously update the assignment of cells to satellites, even as it monitors constellation health and satellite power levels. This functionality may be expanded to control PNT operation. It would be sufficient to direct each satellite to send regular PNT bursts to certain cells at certain times, and to avoid transmitting at other times that might result in collisions with secondary PNT bursts. For instance, satellite 35, beam 4 might be directed to transmit one PNT burst to cell 87654 every time the microsecond hand of the GPS clock reads 0.123456. The full set of such assignments consists of

$$\begin{aligned}\text{PNT C\&C data} &= nC\log_2(bN\cdot 10^6)\ \text{bits} \\ &\approx 8.7\,\mathrm{MiB}\end{aligned} \tag{4}$$

of data, and needs to be uploaded no more than a few times per minute.

E. Alternatives

If more satellite signals are provided to a cell (i.e. $n > 5$), geometric factors affecting positioning accuracy improve. Holding the burst and interval durations constant, horizontal accuracy reaches an optimum of $0.191\,\mathrm{m}$ when all available satellites are used (mean number of satellites visible over the United States $= 43.5$). The constellation is not expected to supply this level of service to every cell at all times: the density of paying customers and their preferences should be used to enable high-accuracy service on-demand.

Elsewhere in this work, only the all-signals case is considered. This represents the accuracy of the system when, for instance, the cost for covering an entire geography is borne by one large, cost-insensitive customer with many receivers in that area. It is straightforward to convert accuracy estimates between the five-signals and the all-signals cases: five-signals accuracies are worse than all-signals accuracies by a factor of $2.7\times$ in the horizontal and $2.3\times$ in the vertical directions. (These factors may be optimized further with a more careful selection of five satellites: the numbers given here simply assume the northernmost, southernmost, easternmost, westernmost, and zenith-most visible and non-masked satellites are chosen.)

IV. Precision Analysis

The third primary contribution of this paper is an analysis of the achievable positioning precision for the fused architecture, extending that of Reid et al. [10] (for hosted LEO GNSS).

While the precision equation (5) is unchanged from Reid, the extension to fused LEO GNSS requires each contribution to be revisited. There are contributions from a number of fundamental effects: orbit and clock prediction errors (uncertainty in exactly where and when a signal is emitted), orbit and clock quantization errors (uncertainty due to limited digital bandwidth), altitude-dependent weight factors (relating orbit and clock errors to average ranging errors), ionospheric and tropospheric errors, receiver noise and multipath errors, and dilution of precision. Each effect has been the subject of extensive scholarship, conclusions from which are gathered in the Appendix.

In the fused scenario, ranging observations are intermittent, with duty cycles below 1×10^{-3}; space-borne clocks fall short of atomic quality by a factor of 30 or more at long ephemeris update intervals; and downlink datarates for ephemeris distribution exceed those of GNSS by a factor of 4×10^6 or more. This paper focuses on the part of the trade-space which achieves sub-meter positioning precision, in order to demonstrate that the cost of such a performance goal is modest. It appears in retrospect that Reid et al. may have been overly cautious in targeting merely precision parity with traditional GNSS.

Improvements to the analysis of Reid et al. include refined coverage maps incorporating up-to-date orbits from public regulatory filings and more accurate service masks (§V) and

Quantity	Description	Estimate
τ	Ephemeris Update Interval	$1\,\mathrm{s}$
T	Clock RMS error, worst case	$0.022\,\mathrm{m}$
R	Orbit RMS error radial, worst case	$0.059\,\mathrm{m}$
A	Orbit RMS error along-track, worst case	$0.093\,\mathrm{m}$
C	Orbit RMS error cross-track, worst case	$0.083\,\mathrm{m}$
w_R	Geometric weight factor, radial	0.774
w_A	Geometric weight factor, along-track	0.448
w_C	Geometric weight factor, cross-track	0.448
σ_{SISURE}	Satellite timing and positioning errors	$0.088\,\mathrm{m}$
σ_{IONO}	Ionospheric delay error	$0.028\,\mathrm{m}$
σ_{TROPO}	Tropospheric delay error	$0.050\,\mathrm{m}$
σ_{RNM}	Receiver noise and multipath errors	$0.005\,\mathrm{m}$
σ_{URE}	User ranging error per satellite	$0.105\,\mathrm{m}$
$\Delta x_{95,H}$	95% horizontal error	$0.191\,\mathrm{m}$
$\Delta x_{95,V}$	95% vertical error	$0.246\,\mathrm{m}$
Δx_{95}	95% total error	$0.413\,\mathrm{m}$

TABLE II: Predicted performance statistics for fused LEO GNSS.

much more aggressive performance targets for precise on-orbit orbit and clock determination (§II-A) and prediction (§A-D) based on estimates of steady-state filter covariances.

The final results of the analysis in Appendix A are a 95% horizontal error of $0.191\,\mathrm{m}$ and a 95% vertical error of $0.246\,\mathrm{m}$ for the case of fused LEO GNSS with oven-controlled crystal oscillators and spaceborne single-frequency GNSS receivers (Table II).

V. Orbits and Masks

Hosted and fused solutions naturally have identical geographical service regions. This section provides revised orbital visibility and Dilution of Precision (DoP) maps relative to those produced by Reid et al. [11]. These maps incorporate new information about Starlink satellite orbital parameters from FCC filings, and an improved understanding of avoidance strategies for K-band interference with geostationary orbit (GSO) services.

A. GSO Proximity

In particular, in order to meet the FCC's implementation of the ITU Radio Regulations' limits on K-band non-GSO interference (Figure 6), each Starlink satellite will avoid providing spot-beam service to any geographical location from whose perspective the satellite appears in the sky within 12° of the GSO arc.

The most obvious way to compute this angle is to construct a line of longitude on the celestial sphere through the satellite, and measure the angle between the satellite and the celestial equator: that is, the declination angle. However, due to the finite distance of GSO from the Earth, the true figure of GSO as seen in perspective wanders from the celestial equator by $\pm 9°$. Approximating the true figure of GSO by a circle of constant celestial latitude improves the approximation to $\pm 2°$; using the true figure further improves the approximation to ± 26 seconds of arc. (Each approximation requires somewhat more computation, imposing costs on simulating, and, indeed, operating, the constellation. The improved approximation costs 9 additions, 12 multiplies, 1 division, and 2 square roots per satellite-visibility-test, of which billions may be required per simulated minute.)

The GSO proximity question may also be solved exactly. Construct a circular cone with vertex at the ground receiver's location and axis pointing along the ray from the receiver to the satellite. Let its opening angle be 24°. Represent the GSO orbit as a circle in the Earth's equatorial plane. If and only if the cone intersects the circle, then this satellite cannot serve this receiver at this time. The problem may be restricted to two dimensions in the equatorial plane, since any intersection, if one exists, must lie in this plane. The cone and circle thus become two planar conic sections. The numerical intersection of planar conics has a well-known solution using eigenvector techniques in homogeneous coordinates. This approach is computationally costly, and a number of edge cases caused by numerical round-off must be carefully considered to produce a robust usability test.

B. Revised Visibility and Performance Maps

The geometric masks by elevation angle (e.g. 35°) and GSO proximity affect positioning accuracy by limiting which satellites in a given receiver's sky are available to provide PNT service. The fewer satellites, or the more unfavorable their geometric distribution across the sky, the worse the receiver's positioning accuracy.

Results for evaluating the entire 11,927-satellite Starlink constellation with accurate geometric masks are shown in Figure 1. For a single epoch in time, under a zero-mean-error assumption, these maps show the constant of proportionality between ranging accuracy and each of a number of geometrical summaries of receiver PNT solution accuracy. PDOP relates ranging accuracy to RMS position error, VDOP and HDOP to vertical and horizontal RMS error, and TDOP to clock RMS error. These results are for the extreme case of a user who wishes to buy signals from every visible satellite.

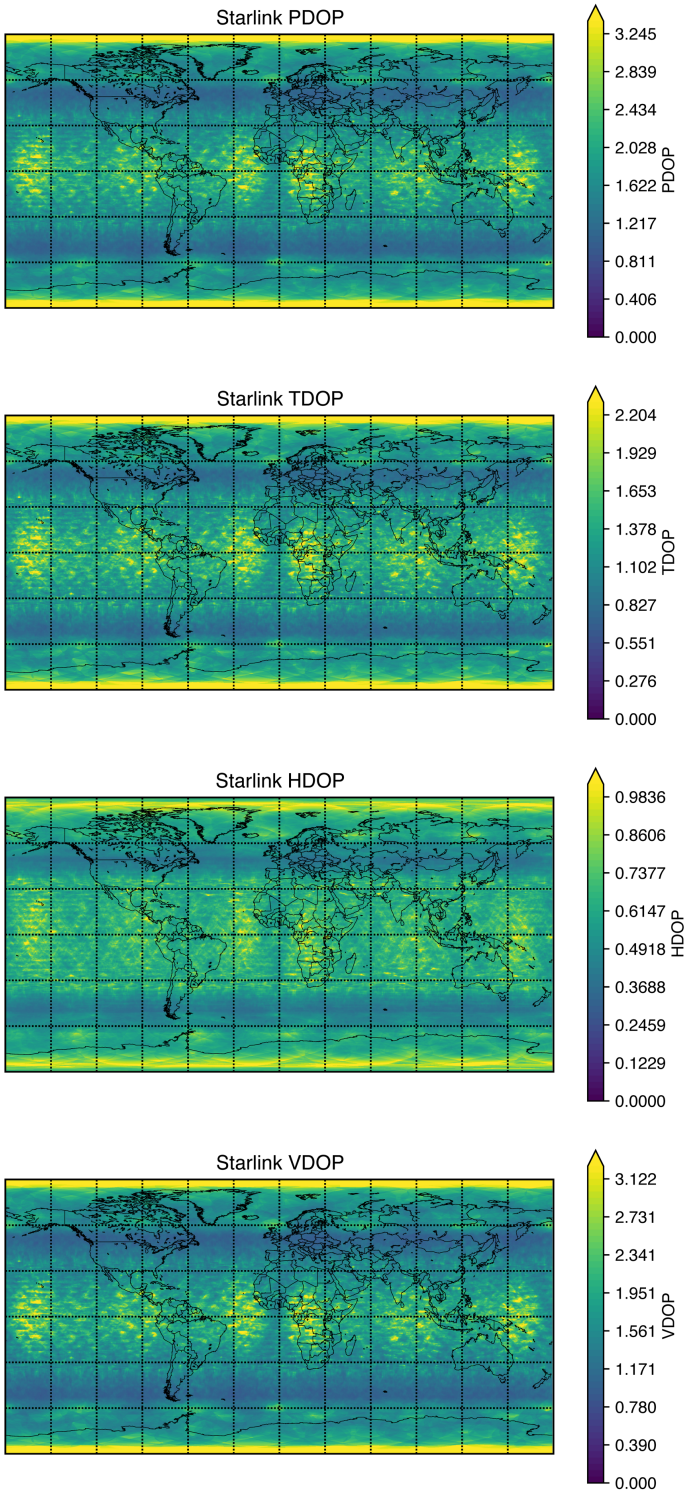

Fig. 1: Constant of proportionality, or Dilution of Precision, between receiver ranging accuracy and receiver PNT solution accuracy for the full (LEO + VLEO) Starlink constellation, based on public FCC filings.

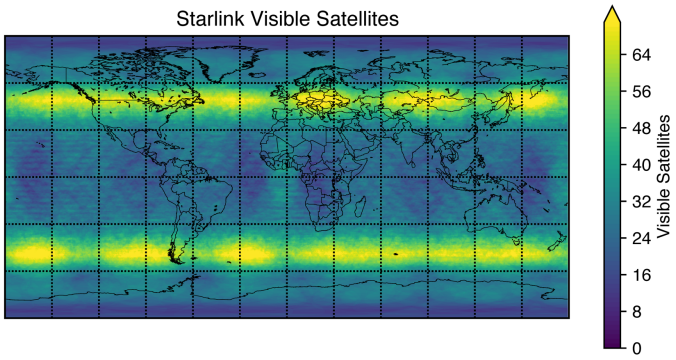

Fig. 2: Number of satellites visible above elevation mask and outside of geostationary arc exclusion zone.

VI. Intermittency and Clock Quality

This section presents and analyzes a simple model for the short-time dynamics of pseudorange positioning using discrete ranging bursts.

Because traditional GNSS signals are transmitted continuously, they are amenable to a wide variety of carrier-tracking and other phase locking techniques. Some receivers may track only intermittently to save power, but this is not necessitated by any part of the system architecture. In fused LEO PNT, however, intermittent observations (on a potentially irregular schedule) are fundamental to the concept of operations. A stochastic system simulator tool (Appendix C) was created to assist in quantifying the time-varying estimator uncertainties that might be observed during the operation of a fused LEO PNT system. Rather than generating individual trajectories that might correspond to a particular realization of the non-stationary Gaussian processes that make up the model, the simulator constructs closed-form expressions for the covariance of all possible combinations of dynamical variables using automatic differentiation. It can then analyze these covariances using either causal or non-causal (i.e. filtering or smoothing) assumptions to produce posterior estimates of the receiver's positioning performance.

Two examples of simulated positioning service are shown in Figures 3 and 4; see figure captions for details.

A. *Qualitative Insight from 1-D Model*

One goal of this work has been to gain insight into the effect of intermittent observation on PNT performance in the high-SNR regime. The filtered instantaneous state uncertainty of an intermittently-observed continuous-time system may exhibit fluctuations (Fig. 8). The Dilution of Precision (DOP) matrix encodes the leading-order relationship between ranging uncertainty and PNT state uncertainty arising from the single-epoch nonlinear least squares pseudo-ranging problem. In the non-single-epoch case, that is, the case of filtering, the DOP matrix may be generalized to a dynamical quantity whose inverse encodes the amount of lingering information remaining from previous ranging epochs. This dynamical DOP grows during the time between ranging bursts, and shrinks rapidly during a ranging burst.

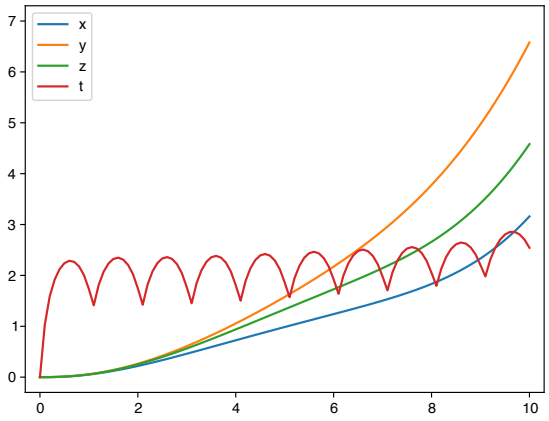

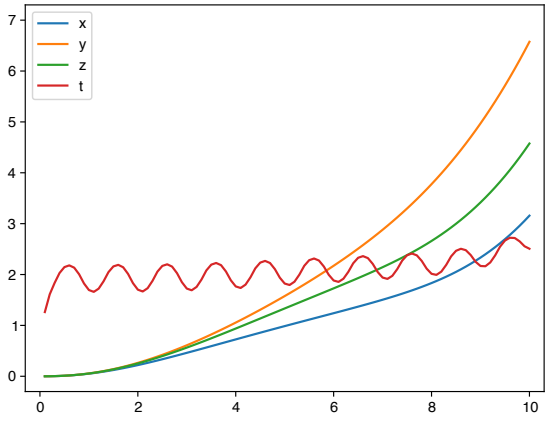

Fig. 3: Simulator outputs for a toy example: positioning uncertainty over time as duty cycle is varied from very low (top) to very high (bottom) while keeping ranging uncertainty constant. A multi-component stochastic dynamical model for the clock on board each of four satellites is simulated, along with a similar model for the ground-based receiver. In addition, the stochastic model includes dynamics for the translational degrees of freedom of all five vehicles.

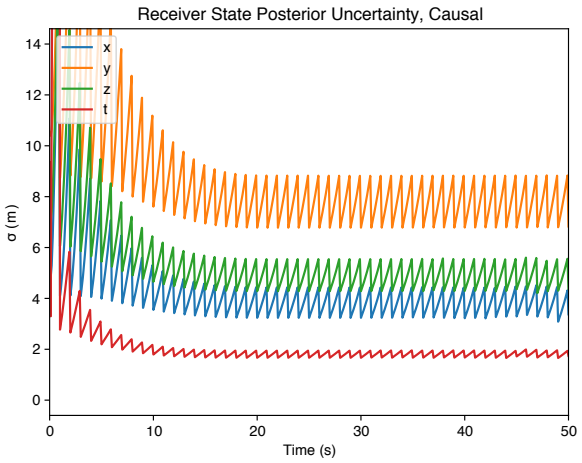

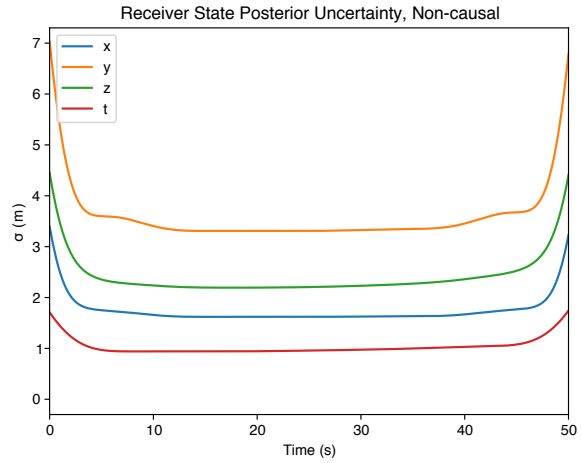

Fig. 4: Simulator outputs for a more realistic example. The satellites are parameterized as though equipped with chip-scale atomic clocks, and the receiver is parameterized as though equipped with a low-end temperature-compensated crystal oscillator. An initially large position uncertainty is reduced by the action of successive pseudorange observations to the four satellites. In the causal case (top), inferred position estimate uncertainties depend only on the observations that would be available up to the time of the position being estimated. In the non-causal case (bottom), the entire position time history is inferred after all observations have been collected, in a batch processing mode.

In one dimension, it is possible to build and solve a closed-form model for the dynamical DOP (see Appendix C-D for derivation and definitions). The result is plotted against duty cycle in Fig. 5 for the case of $Q = 1\,\mathrm{m}^2/\mathrm{s}$ of user position random walk and σ_{URE}, $t_{\mathrm{integrate}}$ as in Appendix A.

Fig. 5 shows the mean and the envelope of the dynamical DOP fluctuations. While this is a simplified model, its qualitative behavior reflects underlying mechanics that are unchanged from the complete 4-D model. Note especially the three distinct regions of behavior. For extremely low duty cycles d, mean and instantaneous DOP together run like $\sqrt{d}$; for intermediate duty cycles, mean-DOP plateaus, while best-case instantaneous DOP (from now on, "min-DOP") continues to improve, even faster than $\sqrt{d}$, towards an asymptotic value; and for high duty cycles, mean-DOP saturates, quadratically approaching min-DOP, while the worst-case instantaneous DOP (max-DOP) stubbornly remains high.

One intuition to take away from this analysis is that if an application only requires periodic position estimates at the same rate as the satellite sends ranging bursts, then the navigation system need only report estimates featuring min-DOP. In this case, duty cycles in the middle (according to log) of the plateau region are ideal: the knee of the curve for min-DOP is located here, and additional ranging measurements provide little benefit.

Note that if an application requires more frequent filter outputs, it is greatly preferable to negotiate a higher rate of ranging measurements from the satellites. This would allow

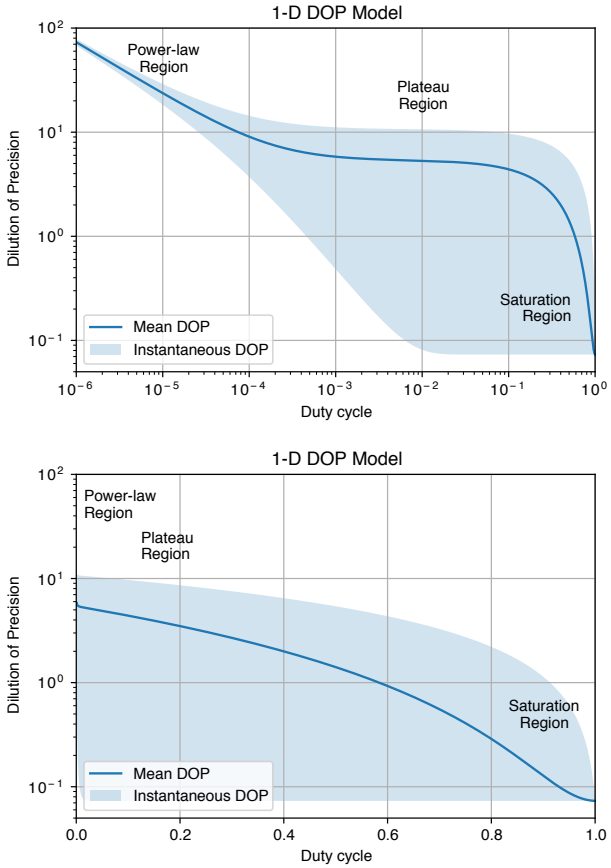

Fig. 5: Results showing min, mean, and max dilution-of-precision over each second of observation for the one-dimensional simplified model. The qualitative behavior shown here is believed to be representative of the full model. See the text for a discussion of the three regions of operation. Top: dilution-of-precision versus log duty cycle. Bottom: dilution-of-precison versus linear duty cycle.

the filter to consistently achieve the min-DOP. The alternative is to propagate the filter forward and report estimates intermediate between ranging epochs; but in this case, the application will occasionally suffer unfavorable max-DOP estimates.

Note also that for fused PNT using certain LEO broadband providers, burst lengths shorter than $500\,\mu$s may not be possible. In the simplified model with ranging epoch $1\,$s, this corresponds to duty cycles below 5×10^{-4}. Lower duty cycles are possible with longer epochs.

VII. Anti-Jam

Anti-jam performance is a key performance indicator for providing assured PNT from LEO. This section briefly analyzes differences in anti-jam performance between traditional and fused LEO GNSS.

Consider the comparative effectiveness of a 100 watt average power jammer against traditional, continuous GNSS signals, versus that same jammer against a fused LEO PNT signal with discrete downlink bursts. The receiver's greater antenna directivity in the K-band scenario gives it some advantage over

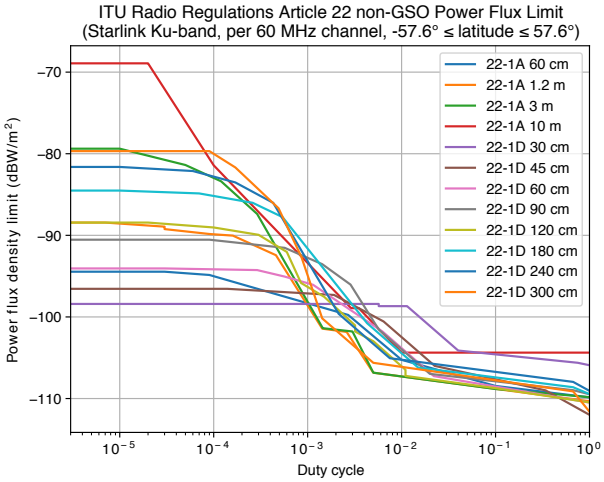

Fig. 6: Power flux density limits imposed under various tabulated scenarios from the ITU Radio Regulations [19].

the L-band scenario. (This advantage is possible because the ratio of typical receiver dimensions to wavelength is an order of magnitude larger in the K band.)

Note that a steerable phased-array antenna designed for comms is not necessarily capable of null-forming, particularly if it is not full-digitizing (more detail below). Even if it is capable, null-forming may be mooted if the jammer consists of many distributed low-power transmitters rather than one centralized high-power transmitter, so that the number of signals to be nulled is large.

On the other hand, the K-band user terminals designed for the comms mission are subject to stringent ITU limitations on low-elevation radiated power, so as not to interfere with terrestrial K-band comms. In particular, the SpaceX Blanket CPE-1 filings for Starlink Earth stations describes a transmit antenna capable of $33.8\,$dB of selectivity between the satellite and the horizon. By Helmholtz reciprocity, this antenna array may be expected to be equally selective between incoming satellite signals and incoming low-elevation jamming signals, as would be typical for over-the-horizon jamming signals, or for ground-based jammers with line-of-sight to distant receivers. This analysis therefore supposes that the K-band receiver has $33.8\,$dB of anti-jam selectivity. An L-band choke ring with the same elevation mask would give only $8.5\,$dB of selectivity between the satellite and the horizon [17]: an anti-jam advantage of $25.3\,$dB for K-band fused LEO PNT.

The K-band fused LEO PNT advantage is further extended by the greater field strength. A typical power flux density for GNSS signals is on the order of $-135\,\mathrm{dBW/m^2}$ over a $20\,$MHz-wide band. The power flux density for K-band LEO broadband is generally set according to limits imposed by ITU Radio Regulations, but based on FCC filings [18], it may be as high as $-104.2\,\mathrm{dBW/m^2}$ over a $60\,$MHz-wide band. This extends the K-band fused LEO PNT anti-jam advantage to $56.1\,$dB.

A. Jamming and Ranging Bursts

One possible concern is that if the jammer is limited in its average power, rather than its instantaneous power, its effectiveness may be substantially increased by generating pulsed output. If the jammer can perfectly anticipate downlink bursts for the fused LEO PNT service, and if those bursts have a duty cycle of, say, $10 \times 500\,\mu\mathrm{s}$ / $10\,\mathrm{s}$, then the receiver would lose $33\,\mathrm{dB}$ of jamming immunity.

In response to this concern, this work offers two minor and two major notes. The first minor note is that a pulsed microwave transmitter may be considerably more expensive, bulky, and power-hungry than a continuous transmitter of similar average power output, particularly under the additional requirement that the jamming signal should be modulated in a non-trivial way and confined to a target band. These attributes make the pulsed jammer more difficult to construct, more costly to operate, and more of a target.

The second minor note is that the jammer cannot perfectly anticipate downlink bursts, even with access to the downlink schedule, because the time-of-arrival of each downlink burst at each receiver depends on that receiver's location. If the jammer wishes to blanket a target area, they must increase their duty cycle accordingly. The light-crossing time of a single $7\,\mathrm{km}$ spot beam is $23\,\mu\mathrm{s}$, compared with an anticipated burst length of $500\,\mu\mathrm{s}$. The jammer must therefore produce pulses 9% longer than those from the satellites in order to ensure full overlap at every receiver. For a jammer with longer range than the size of one spot beam, the jamming pulses must be proportionally longer, costing the jammer $3\,\mathrm{dB}$ of power at a $75\,\mathrm{km}$ radius, and an additional $10\,\mathrm{dB}$ of power for each additional order of magnitude in radius (over and above the $20\,\mathrm{dB}$ of power required to meet spreading losses).

The first major note in response to the concern of pulsed jammers is to focus in on the matter of perfectly antipating downlink bursts. If downlink burst scheduling is shared across many PNT customers, then it must be, to some extent, knowable in advance. Otherwise, a receiver has no way to pre-emptively beam-form towards the satellite that is transmitting at any given time. The alternative to pre-emptive beamforming is to construct a full-digitizing phased array, with one amplifier, mixer, and analog-to-digital converter for each antenna—of which there may be dozens to hundreds. This would be prohibitively expensive for the vast majority of potential PNT customers, particularly compared with traditional GNSS.

However, the low duty cycle of the fused LEO PNT service opens up another possibility. Suppose that the satellites (or their central scheduler) have a cryptographically secure pseudo-random number generator and a well-known public-private key pair. Then for each downlink epoch, for each satellite, the scheduler announces that the downlink burst will take place within a certain interval of time—say, $500\,\mathrm{ms}$. The scheduler picks a time for the burst uniformly at random within this interval. It then composes a message indicating the satellite, epoch, and exact transmission time, adds a digital signature using its private key to attest to the validity of the message, and uploads this message to the satellite to be included in the downlink burst.

The legitimate PNT customer steers her antenna to point at the satellite and waits patiently. When the burst arrives, she compares the time-of-arrival against the signed transmission time, and rejects the burst if the signature is invalid or if she detects a timing discrepancy (that is, a replay attack). She then steers her antenna towards the next satellite from which she expects to receive a burst, and waits once again. If signals are transmitted from each satellite with an average interval of $10\,\mathrm{s}$, then she can track up to 20 satellites in this way.

The jammer does not know in advance exactly when the burst will be scheduled; it only knows what everyone else does: that the burst will occur within some broad window. Within any given $500\,\mathrm{ms}$ window, then, the jammer must transmit continuously until the downlink burst arrives. On average, the jammer's duty cycle is 50%; it is no longer pulsed, and the $33\,\mathrm{dB}$ of advantage it gained by pulsing is reduced to $3\,\mathrm{dB}$.

A drawback of this solution is that the legitimate PNT customer cannot track more than a certain number of satellites. One can boost this number somewhat in an opportunistic fashion: on average, the PNT customer spends $250\,\mathrm{ms}$ waiting for any given satellite. The scheduler may, at its discretion, schedule overlapping downlink windows from multiple satellites. Suppose that satellites 1 and 2 have both been scheduled for the same interval, $[0\mathrm{ms}, 500\mathrm{ms}]$. Suppose also that satellites 3 and 4 have both been scheduled for the interval $[500\mathrm{ms}, 1000\mathrm{ms}]$, and so forth. At $t = 0\mathrm{ms}$, the legitimate PNT customer steers her antenna towards satellite 1. When she receives the downlink burst from satellite 1, she immediately steers her antenna towards satellite 2. If she is lucky, she will receive a second downlink burst some time before $t = 500\mathrm{ms}$. At $t = 500\mathrm{ms}$, she steers her antenna towards satellite 3.

In each $500\,\mathrm{ms}$ interval, then, she certainly receives one burst, and with probability 0.5 she receives a second burst. On average, she receives 3 bursts per second. This procedure can be extended, with decreasing efficiency, to a limiting value of $e - 1 \approx 1.718$ bursts per interval, or 3.44 bursts per second.

The second major response is that an A-PNT customer who desires improved anti-jam performance may choose to pay for a second allocation of downlink bursts, separate from those shared among other customers. These new downlink bursts may be scheduled in a way that is totally unknown to the jammer, negating the advantage of bursts. The A-PNT customer may also purchase higher duty cycle, or additional ranging burst frequency, as desired to combat current jamming conditions in a pay-as-you-go fashion.

VIII. Conclusion

The present report has demonstrated that sub-meter positioning, navigation, and timing from LEO is both economical and attractive: a fused LEO GNSS system with a 95% horizontal error of $0.191\,\mathrm{m}$ and a 95% vertical error of $0.246\,\mathrm{m}$ is feasible using low-cost oven-controlled crystal oscillators and spaceborne single-frequency GNSS receivers. Using only five satellites to serve each geography sacrifices a factor of $2.7\times$ in horizontal and $2.3\times$ in vertical accuracy, and yields

a low opportunity cost to provide this service: even serving all geographical locations below $\pm 60^\circ$ latitude simultaneously, the portion of transmit capacity reserved for PNT use is only 1.25%.

ACKNOWLEDGMENTS

The authors wish to thank Lakshay Narula and Daniel LaChapelle for invaluable assistance.

APPENDIX A
ERROR BUDGET

A. Receiver Positioning Error

The receiver performs two operations: time-of-arrival (ToA) estimation, and nonlinear least squares (single-epoch) or Kalman filtering (multi-epoch) to find the best state estimate(s). The effects of these two operations on receiver positioning error can be encoded in two factors: the per-satellite User Range Error (σ_{URE}), and the Dilution of Precision matrix (DOP), respectively. Various traces over the DOP matrix extract horizontal (HDOP) and vertical (VDOP) components. Assuming an uncorrelated Gaussian error model leads to the expressions below, which are directly analogous to those in Reid et al. [11].

$$\left.\begin{aligned}
\begin{pmatrix}\text{95\% horiz.}\\ \text{error}\end{pmatrix} &= \sigma_{\mathrm{H}} \cdot \sqrt{F^{-1}_{\chi^2_2}(0.95)}, \text{ where}\\
\sigma^2_{\mathrm{H}} &= \mathrm{HDOP} \cdot \sigma^2_{\mathrm{URE}}\\
\begin{pmatrix}\text{95\% vert.}\\ \text{error}\end{pmatrix} &= \sigma_{\mathrm{V}} \cdot \sqrt{F^{-1}_{\chi^2_1}(0.95)}, \text{ where}\\
\sigma^2_{\mathrm{V}} &= \mathrm{VDOP} \cdot \sigma^2_{\mathrm{URE}}\\
\sigma^2_{\mathrm{URE}} &= \sigma^2_{\mathrm{SISURE}} + \sigma^2_{\mathrm{IONO}} + \sigma^2_{\mathrm{TROPO}}\\
&\quad + \sigma^2_{\mathrm{RNM}}\\
\sigma_{\mathrm{SISURE}} &= \text{Signal-in-Space User Range Error}\\
\sigma_{\mathrm{IONO}} &= \text{Ionospheric delay uncertainty}\\
\sigma_{\mathrm{TROPO}} &= \text{Tropospheric delay uncertainty}\\
\sigma_{\mathrm{RNM}} &= \text{Receiver noise and multi-path}
\end{aligned}\right\} \quad (5)$$

B. Signal-in-Space User Range Error

This component represents the portion of the one-satellite ranging error which is characteristic of the signal itself, and would be present even with no noise, no atmosphere, and an ideal receiver. If this ideal receiver under ideal conditions measures the ToA of the signal, it suffers a ranging error when the signal is emitted from a different position or at a different time than is indicated by the orbit and clock ephemeris.

To quantify this, one defines σ_{SISURE} to be the ideal receiver's RMS ranging error due to ephemeris inaccuracy, averaged over a single satellite's ground service area. The radius of this circular service area is determined by the satellite's altitude and the receiver's elevation mask. The average can be expressed analytically in terms of the RMS orbit and clock errors.

$$\begin{aligned}
\sigma^2_{\mathrm{SISURE}} &\le (w_R R + T)^2 + w_A^2 A^2 + w_C^2 C^2\\
T &= \text{RMS clock ephemeris error}\\
R &= \text{RMS orbit ephemeris error, radial}\\
A &= \text{RMS orbit ephemeris error, along-track}\\
C &= \text{RMS orbit ephemeris error, cross-track}\\
w_A^2 &= \text{ Geometric weight factor for along-track errors to contribute to average ranging error}\\
&= \frac{a^2 + a(u+1) + 1}{8a^2} + \frac{(a^2-1)^2 \log\left(\frac{(a-1)^2}{1-2au+a^2}\right)}{16a^3(1-u)}\\
w_C^2 &= \text{ Geometric weight factor for cross-track errors to contribute to average ranging error}\\
&= w_A^2\\
w_R^2 &= \text{ Geometric weight factor for radial errors to contribute to average ranging error}\\
&= 1 - 2w_A^2\\
u &= \frac{\cos^2 E + \sqrt{a^2 - \cos^2 E}\sin E}{a}\\
a &= 1 + \text{altitude}/r_{\text{Earth}}\\
E &= \text{elevation mask} = 35^\circ\\
c &= \text{speed of light}
\end{aligned}$$

To generate ephemerides, the ground segment must continually estimate the state of each satellite's clock and orbit. It must then extrapolate these states forward to cover the prediction interval of the ephemeris. Thus, the uncertainty in orbits and clocks begins each ephemeris epoch at some small value, and then grows steadily up until the beginning of the next ephemeris epoch, as the extrapolation becomes gradually less reliable.

C. Clock Prediction Error

The well-known two-parameter clock model, given in terms of Allan deviation parameters h_{-2}, h_0, can be expressed as a linear Itô stochastic differential equation with Gaussian priors:

$$\begin{aligned}
\mathrm{d}\begin{bmatrix} f_t \\ \phi_t \end{bmatrix} &= \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}\begin{bmatrix} f_t \\ \phi_t \end{bmatrix}\mathrm{d}t\\
&\quad + \begin{bmatrix} \pi\sqrt{h_{-2}} & 0 \\ 0 & \sqrt{h_0}/2 \end{bmatrix}\begin{bmatrix} \mathrm{d}B_{1,t} \\ \mathrm{d}B_{2,t} \end{bmatrix}\\
\begin{bmatrix} f_0 \\ \phi_0 \end{bmatrix} &\sim \mathcal{N}(0, \Sigma_0)
\end{aligned}$$

This may be solved to give

$$\mathrm{Var}\{\phi_t\} = \tfrac{2\pi^2}{3}h_{-2}t^3 + \tfrac{1}{2}h_0 t + \begin{bmatrix} t & 1 \end{bmatrix}\Sigma_0\begin{bmatrix} t \\ 1 \end{bmatrix}$$

giving, at the end of an ephemeris epoch of duration τ, the RMS ranging error component

$$T = c \cdot \sqrt{\tfrac{2\pi^2}{3}h_{-2}\tau^3 + \tfrac{1}{2}h_0\tau + \begin{bmatrix} t & 1 \end{bmatrix}\Sigma_0\begin{bmatrix} t \\ 1 \end{bmatrix}}$$

In this analysis, the clock parameters will be set to model an oven-controlled crystal oscillator (OCXO):

$$h_{-2} \approx 6 \times 10^{-25}\,\mathrm{s}^{-1} \qquad h_0 \approx 2 \times 10^{-25}\,\mathrm{s}$$

1) Clock Estimation Error: In order to use this model, one must supply the initial clock state uncertainty. This analysis will rely on the carrier-phase residuals of $2\,\mathrm{cm}$ reported by Montenbruck et al. [20] as a baseline for the clock state uncertainty. Those authors suggest applying a factor-of-two correction to estimate the residuals under real-time operation, as in the present analysis; their analysis was for batch-mode operation.

The clock state is a two-element vector of frequency and phase. The uncertainty at the ephemeris epoch is therefore a 2×2 covariance matrix Σ_0. Only the $(\Sigma_0)_{22}$ element in this matrix, the phase variance, is known at this point: it has been fixed to the value $(2\,\mathrm{cm}/c)^2$. How can one obtain the other matrix elements?

The solution is to impose a steady-state condition: the information flowing in from clock phase observations (by the reference network or ground control segment) is in balance with the information lost to process noise. That is, one may assume that the satellite clock has been continuously measured for all $t < 0$. Provided that the clock dynamics are observable from phase measurements, and for any given rate of information gain from observations, the ground stations' estimates for the satellite clock state will reach a steady-state covariance. This steady-state may be computed using a standard Continuous-time Algebraic Riccati Equation (CARE) solver.

The CARE solver provides the steady-state covariance for a given observation information rate (i.e. inverse covariance per unit time). This information rate parameter may be determined by bisection (trial and error) from the known phase variance, which ought to appear in the CARE solver output as the $(\Sigma_0)_{22}$ element. This procedure gives the complete matrix-valued initial clock state uncertainty.

$$\Sigma_0 = \begin{bmatrix} (0.0031\,\mathrm{m})^2 & (0.0066\,\mathrm{m})^2 \\ (0.0066\,\mathrm{m})^2 & (0.02\,\mathrm{m})^2 \end{bmatrix} / c^2$$

Jitter in the timing of the downlink modem components (i.e. unpredictable differences between scheduled and actual burst timings) can be modeled as white phase noise in the satellite clock. For the purposes of this analysis, which considers only a single ranging measurement, jitter may be added into the at-epoch clock uncertainty $(\Sigma_0)_{22}$. A more detailed model would account for the fact that clock errors are correlated across ranging measurements to the same satellite, while jitter differs from burst to burst.

2) Ephemeris Interval and Data Rate: Fig. 7 shows how clock prediction uncertainty grows with time for various types of clocks. At $t = 0$, all of the clocks are initialized to the same uncertainty. (Note the log scale on the horizontal axis; $t = 0$ is infinitely far to the left of the frame).

For the remainder of this analysis, the interval between ephemeris epochs will be fixed as

$$\tau = 1\,\mathrm{s}.$$

This imposes a modest overhead on the broadband communications system for broadcasting new ephemerides. Suppose that a satellite broadcasts ephemerides in a burst of bandwidth $60\,\mathrm{MHz}$ and duration $500\,\mu\mathrm{s}$. Suppose further that the system uses an error-correction scheme that achieves 75% of the Shannon limit for the additive white Gaussian noise (AWGN) channel. Then the data content of the burst may exceed $50\,\mathrm{kb}$: enough to spend 1000 bits on orbit and clock ephemeris for each of several dozen visible satellites.

In Reid et al. [10], great finesse is applied to packing ephemerides into the same 302-bit message size as is used for GPS. The clock correction parameters occupy an additional 56 bits. For the satellite network considered herein, that level of economy is plainly not necessary.

D. Orbit Prediction Error

Montenbruck et al. [20] demonstrate $4\,\mathrm{cm}$ RMS orbit determination for the GRACE-B spacecraft (altitude $500\,\mathrm{km}$) using dual-frequency GPS observations. Their errors are reported for forward-backward smoothed (i.e. non-causal) estimation. LEO PNT requires real-time orbit estimation–and even prediction–and hence this analysis applies a factor-of-two correction, as suggested by Montenbruck et al., to estimate the errors under real-time (i.e. causal) operation.

At an ephemeris epoch, then, the RMS orbital estimation errors are assumed to be

$$R_0 = 0.059\,\mathrm{m} \qquad A_0 = 0.093\,\mathrm{m} \qquad C_0 = 0.083\,\mathrm{m}$$

It is not sufficient to estimate the orbits at the ephemeris epoch; one must predict them out to the next epoch, so that receivers may use them for real-time positioning. To determine the worst-case RMS orbit prediction errors, one must posit a dynamical model for the residual (i.e. unmodeled) accelerations suffered by the spacecraft.

Montenbruck et al. summarize their dynamical model in Table 1 of that work [20]. In the present error analysis, the only component that needs be carried over is the residual (i.e. unmodeled) acceleration: all the other components of the dynamical model are deterministic.

Montenbruck et al. model the radial, along-track, and cross-track components of the residual acceleration of the spacecraft as a set of three uncoupled first-order Gauss-Markov processes, each characterized by a steady-state variance σ^2 and a correlation time-scale τ. The parameters are not given, but their empirical values are roughly

$$\sigma_{a,R}^2 = (100 \times 10^{-9}\,\mathrm{m/s^2})^2 \qquad \tau_{a,R} = 2400\,\mathrm{s}$$
$$\sigma_{a,A}^2 = (100 \times 10^{-9}\,\mathrm{m/s^2})^2 \qquad \tau_{a,A} = 2400\,\mathrm{s}$$
$$\sigma_{a,C}^2 = (20 \times 10^{-9}\,\mathrm{m/s^2})^2 \qquad \tau_{a,C} = 2400\,\mathrm{s}$$

As with the clocks, the full 3×3 covariance matrix of the initial position, velocity, and acceleration in the R, A, and C directions may be obtained from the RMS position uncertainties R_0, A_0, and C_0 by by applying a steady-state condition. Once again, one may use bisection to find the rate of information flow from ranging observations needed to balance the rate of information loss to process noise under

the constraint that the steady-state covariance should yield the pre-specified position uncertainties.

$$\mathrm{d}\begin{bmatrix} p_t \\ v_t \\ a_t \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & -1/\tau_a \end{bmatrix} \begin{bmatrix} p_t \\ v_t \\ a_t \end{bmatrix} \mathrm{d}t + \begin{bmatrix} 0 \\ 0 \\ \sigma_a/\sqrt{\tau_a} \end{bmatrix} \mathrm{d}B_t$$

$$\begin{bmatrix} p_0 \\ v_0 \\ a_0 \end{bmatrix} \sim \mathcal{N}(0, \Sigma_0)$$

By solving this stochastic differential equation, the ephemeris errors after t seconds of prediction may be expressed as

$$R, A, C = \sqrt{\frac{1}{3}\sigma_a^2\tau_a\Big(2(t-\tau_a)^3 - 12e^{-t/\tau_a}t\tau_a^2 + 5\tau_a^3 - 3e^{-2t/\tau_a}\tau_a^3\Big) + v^T\Sigma_0 v},$$

$$\text{with } v = \begin{bmatrix} 1 \\ t \\ \tau_a(t + \tau_a(e^{-t/\tau_a} - 1)) \end{bmatrix}$$

The solutions are depicted in Fig. 7: for ephemeris intervals longer than a few seconds, for all clocks of lesser quality than cesium or rubidium, orbital errors may be expected to be a sub-dominant component of the overall prediction error.

E. Atmospheric Errors

These components are much the same as in traditional GNSS. The tropospheric component is largely insensitive to frequency, as long as one operates well below the water resonance at $22.24\,\mathrm{GHz}$. The value below assumes that surface weather data is available at the receiver [21]. If numerical weather model data is available, then the RMS error is reduced to $2.8\,\mathrm{cm}$ [22].

The ionospheric component runs as the inverse square of the frequency. Note that the quantities here are not atmospheric delays, but atmospheric delay *uncertainties*. The ionospheric errors assumed here are for the IGS Rapid global ionospheric model [23], [24]. This parameter is not critical to the final analysis, since ionospheric errors form a small portion of the total error budget. However, if desired, data from a dense reference network (i.e. one reference node per $20\,\mathrm{km}$) can reduce this uncertainty to $0.03\,\mathrm{TECU}$ [25].

$$\begin{aligned} \sigma^2_{\mathrm{IONO}} &= \left(\frac{40.3 \times 10^{16}\,\mathrm{m/s^2/TECU} \cdot \sigma_{\mathrm{STEC}}}{f^2}\right)^2 \\ &\approx (0.028\,\mathrm{m})^2 \\ \sigma^2_{\mathrm{TROPO}} &= \sigma_{\text{TROPO (GPS)}} \\ &\approx (0.050\,\mathrm{m})^2 \\ f &= \text{center frequency} \\ &\approx 12\,\mathrm{GHz} \\ \sigma_{\mathrm{STEC}} &= \text{Uncertainty in Slant Total Electron Content} \\ &\approx 10\,\mathrm{TECU} \end{aligned}$$

Note that the typical thin-shell ionospheric model is less appropriate for satellites in very low Earth orbit. The true ionospheric error is bounded above by these figures.

F. Receiver Noise and Multipath

Specializing the expression of Appendix B to the case of a spectrally-flat signal within a $60\,\mathrm{MHz}$ passband,

$$\begin{aligned} \sigma^2_{\mathrm{RNM}} &\geq \frac{3c^2 k_{\mathrm{B}} T_{\mathrm{RX}}}{2\pi^2 W^2 P\, t_{\text{integrate}}} \\ k_{\mathrm{B}} &= \text{Boltzmann's constant} \\ T_{\mathrm{RX}} &= \text{noise temperature} \\ &= 273\,\mathrm{K} \cdot 10^{0.1 \cdot \text{noise figure}} \\ W &= \text{bandwidth} = 60\,\mathrm{MHz} \\ P &= \text{received power} \\ t_{\text{integrate}} &= \text{ranging burst duration} = 500\,\mu\mathrm{s} \\ \text{noise figure} &= 6\,\mathrm{dB} \end{aligned}$$

Using the receiver antenna gain of $33.2\,\mathrm{dBi}$ and a power flux density of $-104.2\,\mathrm{dBW/m^2}$, one obtains a best-case received signal power of $-84.0\,\mathrm{dBm}$.

1) Multipath: Compared to traditional GNSS, LEO PNT receivers will be far less sensitive to multipath and non-line-of-sight errors. This improvement comes from a combination of factors: rapid decorrelation, greater bandwidth, and receiver antenna selectivity.

The rapid decorrelation of multipath components for LEO PNT signals arises from their shorter wavelength (so that the relative phases of individual multipath components are $7\times$ more sensitive to path-length differences/excess delays) and their lower orbital altitude (so that angles of incidence between the satellite signal and reflectors in the environment vary $100\times$ more rapidly). These effects multiply: with the same reflector geometry, multipath components "average out" in the integrated correlator outputs $700\times$ more rapidly for LEO PNT. Reid et al. estimate a multipath decorrelation time-scale of $600\,\mathrm{s}$ for traditional MEO GNSS and $60\,\mathrm{s}$ for Iridium. Their ratio of $10\times$ is surprisingly low. This may be due, in part, to Iridium's operation in L-band at a higher altitude; however, even these assumptions would favor a ratio of at least $50\times$. Supposing the greater number of $700\times$, the multipath decorrelation time-scale will be less than $1\,\mathrm{s}$.

The greater bandwidth of LEO PNT signals compared to e.g. the GPS L1 C/A signal potentially reduces the time-of-arrival ambiguity associated with multipath by $30\times$, so that multipath components with path-length differences larger than $5\,\mathrm{m}$ will not appear in the output of a narrow correlator. This reduces the volume of space within which potential reflectors must reside (if they are toV generate a problematic multipath component) by an even greater factor.

Receiver antenna directivity further suppresses multipath components arriving far from boresight. In traditional GNSS, one advantage of circular polarization is that the receive antenna can be selective against multipath, since singly-reflected signals have flipped helicity. However, if these flipped-helicity signals arrive far from boresight, even high-end antennas

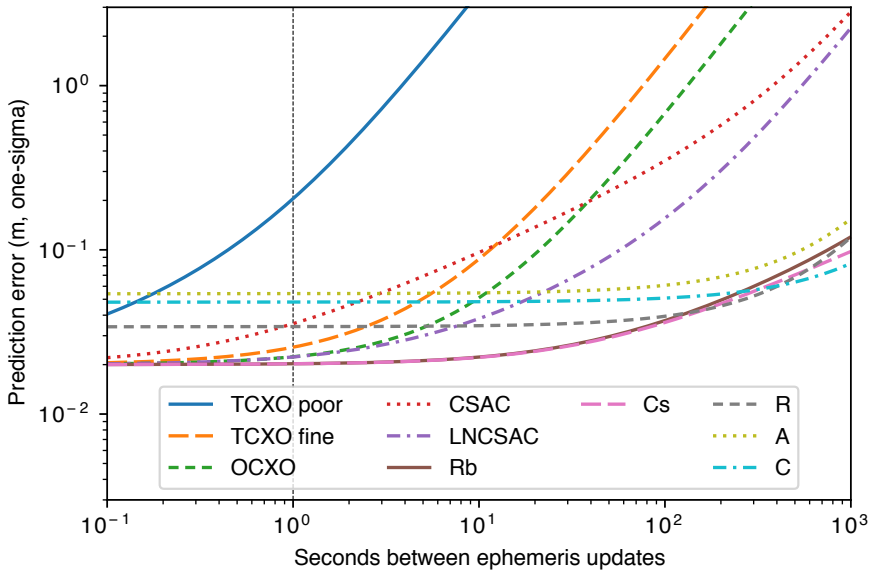

Fig. 7: Prediction errors, and hence ephemeris errors, as a function of time since the last ephemeris epoch. Three curves, labeled R, A, and C, correspond to orbit ephemeris errors which may be expected using the reduced dynamics model of Montenbruck et al. [20]. The remaining seven curves correspond to clock ephemeris errors which may be expected using COTS clocks of varying performance. This analysis assumes an OCXO-quality oscillator on each satellite (green curve). The selected epoch interval of $1\,\mathrm{s}$ is shown by the dashed vertical line.

cannot offer a very good axial ratio. For LEO PNT, since multipath components arriving far from boresight are excluded by antenna directivity, the axial ratio can be more favorable.

For these reasons, the present analysis neglects multipath errors.

G. Dilution of Precision

For single-epoch positioning, the dilution-of-precision factors HDOP and VDOP may be read from Fig. 1. This analysis uses values averaged over the land area of the United States, where the mean number of visible satellites is 43.5:

$$\mathrm{HDOP} = 0.55 \qquad \mathrm{VDOP} = 1.43$$

For filtered (i.e. multi-epoch) positioning, the dilution-of-precision factors may be determined by the dynamical model of Appendix C.

H. Summary

Based on this analysis, an estimate for the 95$^{\text{th}}$ percentile user ranging accuracy of a Starlink-based A-PNT system using OCXO clocks, $1\,\mathrm{s}$ ephemeris updates, and PPP-based GNSS orbit determination is $0.191\,\mathrm{m}$ horizontal, $0.246\,\mathrm{m}$ vertical. See Table II for more details.

Appendix B
Power, Bandwidth, and Signal Structure

For any given signal structure $s(t)$, under additive white Gaussian noise of power spectral density S_{nn}, the Cramér–Rao lower bound (CRLB) requires that the uncertainty in a single ranging measurement, expressed in units of time, should be at least

$$\sigma_{\tau\tau}^2 \gtrapprox \frac{S_{nn}}{2\int dt |s'(t)|^2},$$

where τ is the duration of the ranging measurement. The least lower bound among signals of bandwidth W, power P, and duration $T \gg 1/W$ is

$$\sigma_{\tau\tau}^2 \gtrapprox \frac{S_{nn}}{2\pi^2 W^2 P T},$$

where k_{B} is Boltzmann's constant and T is the noise temperature of the receiver. Compared with traditional GNSS, a fused LEO PNT system based on current public information about Starlink would deliver $30.8\,\mathrm{dB}$ more power to the receiver across $3\times$ to $30\times$ more bandwidth than a traditional GNSS signal. This gives code-phase fused LEO PNT a factor of $10^4\times$ to $10^6\times$ better ranging variance than traditional GNSS over a similar averaging period. (Furthermore, the ionospheric delay bugbear of L-band ranging is suppressed by a factor of 100 at the Ku band, bringing un-modeled terms to the sub-meter level). In practice, this advantage will be traded off for lower duty cycle operation: if fused LEO PNT targets sub-meter errors, compared with typical $5\,\mathrm{m}$ RMS errors for single-epoch positioning using GPS L1 C/A, then duty cycles as low as 6×10^{-6} are feasible.

An excellent analysis of signal structure principles beyond the CRLB is given by Nanzer et al. [26]. For broadband LEO, the signal-to-noise ratios are expected to be extremely high by design, even before correlation. This provides justification

for the use of the CRLB. Unlike the scenario considered by Nanzer et al., a fused PNT system would have a relatively fixed modulation structure designed to be spectrally flat for high spectral efficiency. The appropriate bound, then, is the CRLB for a spectrally-flat signal:

$$\sigma_{\tau\tau}^2 \gtrapprox \frac{3S_{nn}}{2\pi^2 W^2 PT}$$

For GPS L1 C/A with $\tau = 1\,\mathrm{s}/1\,023\,000$ and a receiver bandwidth of W, the CRLB is

$$\sigma_{\tau\tau}^2 \gtrapprox \frac{S_{nn}}{4PTW/\tau}$$

Appendix C
Stochastic System Simulator Tool

A. Background and Definitions

Some of the most widely used models for dynamical, non-deterministic orbit, clock, and vehicle states are non-linear Gaussian stochastic differential equations (SDEs). These models are appropriate when the underlying non-determinism has the Markov property, and the system state varies continuously over time. The Markov property states that the future of the system is conditionally independent of the past, given full information about the present; it encodes the idea that the system does not have memory outside what is encoded in the state vector. Such a system may be written as

$$\mathrm{d}X_t = f(t, X_t)\mathrm{d}t + g(t, X_t)\mathrm{d}W_t,$$

where X_t is the state vector at time t, f is the deterministic part of the state evolution, g is the non-deterministic part, and W_t is a vector of Wiener processes (that is, $\mathrm{d}W_t$ may be thought of as a vector of infinitesimal i.i.d. Gaussian increments at every time t). All SDEs in this work are of the Itô type.

In many cases of interest for positioning, timing, and navigation, weak non-linearities may be neglected by assuming that all uncertainties in the problem are much smaller than e.g. the distance between the satellite and the receiver. In this case, one may break the state vector into a "large" deterministic part and a "small" non-deterministic part.

$$\begin{aligned}
X_t &= X_t^{(0)} + X_t^{(1)} \\
\mathrm{d}X_t &= \mathrm{d}X_t^{(0)} + \mathrm{d}X_t^{(1)} \\
&= f(t, X_t^{(0)})\mathrm{d}t \\
&\quad + (\nabla_X \otimes f)(t, X_t^{(0)}) \cdot X_t^{(1)}\mathrm{d}t + g(t, X_t^{(0)})\mathrm{d}W_t, \text{ or} \\
\mathrm{d}X_t^{(0)} &= f(t, X_t^{(0)})\mathrm{d}t \\
\mathrm{d}X_t^{(1)} &= (\nabla_X \otimes f)(t, X_t^{(0)}) \cdot X_t^{(1)}\mathrm{d}t + g(t, X_t^{(0)})\mathrm{d}W_t
\end{aligned}$$

The large deterministic part $X^{(0)}$ plays no role whatsoever in the present analysis. It is the fixed background, the railroad-track, the stage atop which the small non-deterministic part $X^{(1)}$ plays its role: and it is this small non-deterministic part that is of interest when one sets out to quantify accuracy—that is, uncertainty, or non-determinism. The large part $X^{(0)}$ is therefore neglected throughout, and the superscript on the small part $X^{(1)}$ is dropped.

Over time-scales short compared to e.g. the overflight time of one satellite as viewed from one geographic location, f and g may be assumed to vary slowly (for instance, piecewise constant). One may now make some definitions to bring the notation more in line with standard practice.

$$\begin{aligned}
A &= (\nabla_X \otimes f)(t, X_t^{(0)}) \\
\sqrt{Q/2} &= g(t, X_t^{(0)}) \\
\mathrm{d}X_t &= AX_t\mathrm{d}t + \sqrt{Q/2}\,\mathrm{d}W_t
\end{aligned}$$

The matrix A is the generator of deterministic time-evolution, and the symmetric positive semi-definite matrix Q is the process noise per unit time.

It is not of interest in the present work to "solve" this equation in the sense of generating trajectories. Rather, it is of interest to solve this equation in the aggregate, determining the probability distribution over possible trajectories. The first step is to apply the Fokker-Planck equation to get a partial differential equation for the probability distribution function of X_t as a function of time. Next, since the equation is linear and the increments are Gaussian, the trajectories X_t are Gaussian processes. Provided that one uses a Gaussian prior distribution for X_0, the instantaneous probability distribution function is Gaussian for all time. One may therefore substitute the ansatz, or functional form, of a Gaussian into the Fokker-Planck equation to reduce it to a system of ordinary differential equation for a mean μ and a covariance Σ:

$$\begin{aligned}
\partial_t \mu_t &= A\mu_t + b \\
\partial_t \Sigma_t &= A\Sigma_t + \Sigma_t A^T + Q
\end{aligned}$$

where ∂_t denotes the partial derivative with respect to time. These ODEs may be solved in terms of the matrix exponential:

$$\begin{aligned}
\mu_t &= e^{A(t-t_0)}\mu_0 \\
\Sigma_t &= e^{A(t-t_0)}\Sigma_0 e^{A^T(t-t_0)} + \int_{t_0}^{t} d\tau\, e^{A(t-\tau)} Q e^{A^T(t-\tau)}
\end{aligned}$$

This "instantaneous" solution alone is not sufficient. One also requires a tool for computing the covariance between X_{t_1} and X_{t_2} for distinct times t_1, t_2. This can be achieved using the Markov property: the future dynamics of the system after t_1 may be viewed as the convolution of the instantaneous distribution of the state at t_1 with the conditional distribution of the state at t_2 given the state at t_1. This finally gives, for $t_2 \geq t_1$,

$$\begin{aligned}
\mathrm{Cov}\{X_{t_1}, X_{t_2}\} &= e^{A(t_1-t_0)}\Sigma_0 e^{A^T(t_2-t_0)} \\
&\quad + \int_{t_0}^{t_1} \mathrm{d}\tau\, e^{A(t_1-\tau)} Q e^{A^T(t_2-\tau)}
\end{aligned}$$

The case where $t_2 < t_1$ may be found by matrix transposition. Since a joint Gaussian distribution is fully characterized by its means and covariances, the analysis is complete. The stochastic simulator tool evaluates these matrix exponentials and definite integrals analytically in certain cases of interest; otherwise it evaluates them numerically.

B. Tooling

Beyond implementing these core mathematics, the simulation tool developed for this work provides additional tools: nonlinear observables with automatic differentiation, means, and covariances; labeled states and sub-systems; polynomial approximations for "large" dynamics; and assembly and solution of normal equations for posterior covariance after observations, with support for both causal and non-causal processing.

Using the labeled state and sub-system tools, one may assemble a vehicle dynamics model from, e.g., three copies of a position-velocity-acceleration (PVA) sub-system, and one copy of a clock drift sub-system. The location of each component in the resulting matrices is tracked, so that observables may be extracted either by index or by name.

A nonlinear observable (NLO) is an algebraic object tied to a particular SDE system. It is stored as a constant plus a linear combination of terms, each of which points to a particular scalar state variable at a particular time. That is, not all terms in a single NLO need be evaluated at the same time. It is therefore possible to construct a pseudorange non-linear observable in terms of the difference between the satellite clock phase measured at the time of transmission, and the receiver clock phase at the time of arrival. Pseudorange observations by correlation over a finite time may be represented approximately by the average of a number of pseudorange observations shifted in time by small increments.

The NLO object overloads various arithmetic operations so that it can be used in complicated expressions (for instance, the Euclidean distance formula). In each operation, the NLO automatically computes the Taylor series of the result to first order in the "small" variables, plus a new constant term. The result is a new NLO object.

Because the NLO is represented as a linear combination of state variables, means and covariances of NLOs or pairs of NLOs are straightforward to compute.

To draw conclusions about position estimation accuracy as a function of time, one selects a set of discrete estimation times, a set of estimation variables, a set of observation times, and a set of observables for each time. For instance, the observables might be satellite-to-receiver pseudoranges for each visible satellite, or perhaps all-to-all pseudoranges among the receiver and the visible satellites; the observation times might be once every second; the estimation times might be ten times per second; and the estimation variables might be the receiver position and clock offset.

Given these four pieces of information, the simulator tool constructs a covariance matrix among all the estimated and observed quantities at all relevant times. This Gaussian distribution can then be marginalized over e.g. only observations at time t and causal observations up to time t, and then conditioned on these observations using the Schur complement. The final result is a covariance matrix over receiver states: the positioning and timing uncertainty.

Because the analysis has been linearized in terms of "small" variables in the vicinity given by the "large" variables, it may be "non-dimensionalized": that is, one may divide the initial covariance Σ_0, the ranging variances, and the positioning variance Σ_t through by σ^2_{URE}, and still have a valid solution to the differential equation. In this case, one may interpret the solution variable $\Sigma_t/\sigma^2_{\text{URE}}$ as the time-varying dilution-of-precision matrix arising from the specified observation geometry.

C. Intuition

For the simplified case of a position-random-walk vehicle dynamics model (i.e. no velocity or acceleration state variables, no clock frequency drift), the time-varying dilution-of-precision has a particularly straightforward ordinary differential equation form, obtained by solving the filter (i.e. causal) covariance update equations and taking the appropriate limits:

$$\partial_t \text{DOP} = \frac{Q}{\sigma^2_{\text{URE}}} - \text{DOP} \cdot \left(\sum_i^{\text{\# sats}} d_i(t) \frac{h_i h_i^T}{t_{\text{integrate}}} \right) \cdot \text{DOP}, \text{ where}$$

$$h_i = \begin{bmatrix} \frac{x_{\text{sat,i}} - x_{\text{recv}}}{|x_{\text{sat,i}} - x_{\text{recv}}|} \\ 1 \end{bmatrix}$$

$$d_i(t) = \begin{cases} 1 & \text{satellite } i \text{ is sending a burst at time } t \\ 0 & \text{otherwise} \end{cases}$$

$$Q = \text{position process noise } [m^2/s]$$

$$t_{\text{integrate}} = \text{ duration of ranging burst; same as in Appendix A-F: Receiver Noise and Multipath}$$

This form makes explicit the intuition that information flows in (reducing dilution) from observations along the directions to the satellites, and information flows out (increasing dilution) according to the non-deterministic vehicle dynamics.

D. 1-D Problem

Consider the one-dimensional case

$$\text{DOP} = x(t)$$

$$\partial_t x = a^2 - b^2 d(t) x^2, \text{ with piecewise solutions}$$

$$x = \begin{cases} x_0 + a^2(t - t_0) & d = 0 \\ \frac{a}{b} \tanh\left(ab(t - t_0) - C\right) & d = 1 \wedge x_0 < \frac{a}{b} \\ \frac{a}{b} \coth\left(ab(t - t_0) - C\right) & d = 1 \wedge x_0 > \frac{a}{b} \end{cases}$$

$$C = \frac{1}{2} \ln \left| \frac{a - b x_0}{a + b x_0} \right|$$

$$a^2 = \frac{Q}{\sigma^2_{\text{URE}}}$$

$$b^2 = \frac{1}{t_{\text{integrate}}}$$

In order to understand the "steady-state" behavior of this system under periodic observations, one must identify its *limit cycle*: that is, the stable periodic orbit towards which its oscillations converge, regardless of initial conditions (Fig. 8). Suppose that ranging bursts are transmitted continuously for the first τ seconds of each period of length T, forever. Then the limit cycle of this system is given by

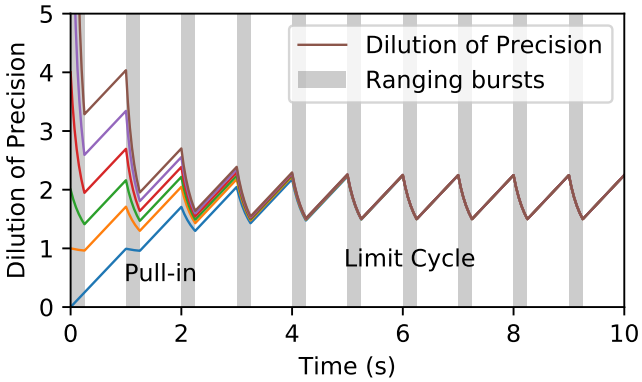

Fig. 8: Simulated dilution-of-precision over the course of ten cycles for the 1-D simplified model. Duty cycle exaggerated for visibility. Regardless of initial value of DOP, the trajectory asymptotically approaches the limit cycle.

$$
\begin{aligned}
x &= \frac{a}{b} \cdot \begin{cases} \frac{1+u_0 w}{w+u_0} & (t \bmod T) \le \tau \\ u_1 + ab((t-\tau) \bmod T) & \text{otherwise} \end{cases} \\
u_0 &= \sqrt{z^2/4 + z\coth(ab\tau) + 1} + z/2 \\
u_1 &= \sqrt{z^2/4 + z\coth(ab\tau) + 1} - z/2 \\
z &= ab(T-\tau) \\
w &= \coth(ab(t \bmod T))
\end{aligned}
$$

Averaging over a period of the limit cycle, the mean dilution-of-precision is

$$
\begin{aligned}
\text{mean DOP} = &\frac{\ln(\cosh(ab\tau) + u_0 \sinh(ab\tau))}{b^2 T} \\
&+ \frac{a}{b}(1-\tau/T)(u_1 + z/2)
\end{aligned}
$$

Results explored in §VI-A.

References

[1] Center for International Earth Science Information Network - Columbia University, "Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11," Palisades, NY, 2018. [Online]. Available: https://doi.org/10.7927/H4JW8BX5

[2] J. J. Spilker, Jr., *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 3: GPS Signal Structure and Theoretical Performance, pp. 57–119.

[3] ——, *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996, ch. 20: Interference Effects and Mitigation Techniques, pp. 717–771.

[4] P. Teunissen and O. Montenbruck, Eds., *Springer handbook of global navigation satellite systems*. Springer, 2017.

[5] E. C. Dolman, "New frontiers, old realities," *Strategic Studies Quarterly*, vol. 6, no. 1, pp. 78–96, 2012.

[6] T. E. Humphreys, *Springer Handbook of Global Navigation Satellite Systems*. Springer, 2017, ch. Interference, pp. 469–504.

[7] M. L. Psiaki and T. E. Humphreys, *Global Navigation Satellite Systems: Theory and Applications*, 2020, ch. Civilian GNSS Spoofing, Detection, and Recovery, pp. 469–504.

[8] O. Luba, L. Boyd, A. Gower, and J. Crum, "GPS III system operations concepts," *IEEE Aerospace and Electronic Systems Magazine*, vol. 20, no. 1, pp. 10–18, 2005.

[9] T. E. Humphreys, L. Young, and T. Pany, "Considerations for future IGS receivers," in *Position Paper of the 2008 IGS Workshop*, 2008, http://www.ngs.noaa.gov/IGSWorkshop2008/docs/recDev-positionpaper.pdf.

[10] T. G. R. Reid, A. M. Neish, T. Walter, and P. K. Enge, "Leveraging commercial broadband LEO constellations for navigation," in *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, Portland, Oregon, Sept. 2016, pp. 2300–2314.

[11] T. G. Reid, A. M. Neish, T. Walter, and P. K. Enge, "Broadband LEO constellations for navigation," *Navigation*, vol. 65, no. 2, pp. 205–220, 2018. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/navi.234

[12] T. G. R. Reid, B. C. Chan, A. Goel, K. Gunning, B. Manning, J. Martin, A. Neish, A. Perkins, and P. Tarantino, "Satellite navigation for the age of autonomy," in *2020 IEEE/ION Position, Location and Navigation Symposium (PLANSx)*. IEEE, 2020.

[13] E. Musk, (@elonmusk), Twitter post: "Direct links aren't needed to offer service." [Online]. Available: https://twitter.com/elonmusk/status/1228604628750626816

[14] J. Wattles, "Here's what you need to know about SpaceX's Starlink internet service," *CNN Business*, Oct. 2019.

[15] SpaceX, "SpaceX non-geostationary satellite system, Schedule S," https://licensing.fcc.gov/myibfs/download.do?attachment_key=1380782, July 2017, SAT-LOA-20170726-00110.

[16] ——, "SpaceX non-geostationary satellite system, Attachment A," https://licensing.fcc.gov/myibfs/download.do?attachment_key=1252848, July 2017, SAT-LOA-20170726-00110.

[17] M. Maqsood, S. Gao, T. Brown, and M. Unwin, "Effects of ground plane on the performance of multipath mitigating antennas for GNSS," in *2010 Loughborough Antennas Propagation Conference*, Nov 2010, pp. 241–244.

[18] SpaceX, "SpaceX non-geostationary satellite system, Attachment A," https://licensing.fcc.gov/myibfs/download.do?attachment_key=1569860, Nov. 2018, SAT-MOD-20181108-00083.

[19] International Telecommunication Union, "Radio regulations," 2016, http://www.itu.int/pub/R-REG-RR.

[20] O. Montenbruck, T. Van Helleputte, R. Kroes, and E. Gill, "Reduced dynamic orbit determination using GPS code and carrier measurements," *Aerospace Science and Technology*, vol. 9, no. 3, pp. 261–271, 2005.

[21] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*, revised second ed. Lincoln, Massachusetts: Ganga-Jumana Press, 2012.

[22] J. Boehm, B. Werl, and H. Schuh, "Troposphere mapping functions for GPS and very long baseline interferometry from European Centre for Medium-Range Weather Forecasts operational analysis data," *Journal of Geophysical Research: Solid Earth*, vol. 111, no. B2, 2006.

[23] M. Hernández-Pajares, J. M. Juan, J. Sanz, R. Orus, A. Garcia-Rigo, J. Feltens, A. Komjathy, S. C. Schaer, and A. Krankowski, "The igs vtec maps: a reliable source of ionospheric information since 1998," *Journal of Geodesy*, vol. 83, no. 3, pp. 263–275, Mar 2009. [Online]. Available: https://doi.org/10.1007/s00190-008-0266-1

[24] A. Rovira-Garcia, J. Juan, J. Sanz, G. González-Casado, and D. Ibáñez, "Accuracy of ionospheric models used in GNSS and SBAS: methodology and analysis," *Journal of geodesy*, vol. 90, no. 3, pp. 229–240, 2016.

[25] M. Murrian, C. Gonzalez, T. E. Humphreys, and T. D. Novlan, "A dense reference network for mass-market centimeter-accurate positioning," in *Proceedings of the IEEE/ION PLANS Meeting*, Savannah, GA, 2016.

[26] J. A. Nanzer, M. D. Sharp, and D. Richard Brown, "Bandpass signal design for passive time delay estimation," in *2016 50th Asilomar Conference on Signals, Systems and Computers*, Nov 2016, pp. 1086–1091.