

Fused Low-Earth-Orbit 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 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 pseudorange would cost less than 0.8% of downlink capacity for the largest of the new constellations, SpaceX's Starlink.

Index Terms—broadband; satellite; navigation.

I. INTRODUCTION

Anticipation is building for commercial broadband Internet services provided by mega-constellations of satellites in low Earth orbit (LEO). Such services' global reach, low latency, and wide bandwidth situate them to revolutionize broadband communications. This paper seeks to establish a less-obvious assertion: In addition to broadband service, these constellations could revolutionize satellite-based positioning, navigation, and timing (PNT). Their space vehicles (SVs) are far nearer and more numerous than those of traditional global navigation satellite systems (GNSS) in medium Earth orbit (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 those of traditional GNSS.

A. Background

All current GNSS have settled into a design in which the constellation consists of a few dozen SVs primarily in MEO whose signals are low-data-rate direct-sequence spread spectrum waveforms confined to L band and occupying less than a total of 125 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 SVs. L band was chosen because its 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 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 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. Although not originally intended for PNT, such 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 must be *economical*: that is, it must offer fundamental advantages over traditional (free) GNSS commensurate with the price tag, otherwise there will be

no demand; and must be sufficiently inexpensive to provide, otherwise there will be no supply. This paper explores 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 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 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 allocations.

D. Fused LEO GNSS

This paper introduces 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 pseudorange (on the order of 3 m), fused LEO GNSS can improve on this by more than an order of magnitude [13]. 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

Characteristic	Traditional GNSS	Hosted [11]	Fused (hemi 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
vert.	4.8 m	4.4 m	48 cm	25 cm
Anti-jam advantage	—	+30 dB	+25.3 dB	+56 dB

TABLE I: Contrasting traditional GNSS, previous hosted proposals, and fused LEO GNSS. Precise orbit determination (POD) here assumes onboard GNSS receivers in LEO (multi-tier GNSS). Positioning precision is 95th percentile in the horizontal and vertical directions. Anti-jam advantage is compared to an L-band choke-ring antenna [13].

as a source of user ranging error, even when the ranging signal used over the communications link adopts the same structure and spectral profile as the usual communications signals. Furthermore, because these signals have a much shorter wavelength than traditional GNSS, it is possible to build a (relatively) compact, highly-directional receiver phased array for an additional 30 dB of anti-jam performance.

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 a hemispherical antenna is preferred, such as in mobile devices, the fused SNR is not high enough to download ephemerides, so a back-up communications link such as cellular data service would be required. Note that certain design elements that give fused LEO GNSS its performance advantage could be incorporated into future hosted payload proposals. However, this paper will only compare against published proposals.

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 up to 1090 km² [14]) is roughly as costly, in terms of satellites’ resources spent providing PNT signals, as a single 5.9 Mbps downlink stream. Whereas broadband service expends constellation resources in proportion to the number and activity level of subscribers, GNSS service consumes resources in proportion to coverage area. For this reason, in dense urban centers where few broadband subscribers could be accommodated and alternatives for high-speed Internet service abound, fused LEO GNSS service may be particularly impactful and lucrative.

F. Contributions

This paper makes three primary contributions. First, it summarizes the features of modern broadband LEO system

design and operation relevant to dual-purposing such systems for PNT. Second, it presents a detailed concept of operations for fused LEO GNSS. Third, it provides an analysis of the opportunity cost to constellation providers for re-allocating resources to provide a fused PNT service. The paper is organized around these contributions.

A preliminary version of this paper was published in [13]. The current version narrows the focus to permit a more complete description of the fused LEO GNSS concept and a more detailed opportunity cost analysis. The reader is referred to [13] for analyses of achievable fused LEO GNSS positioning precision and anti-jam advantage compared to traditional GNSS; summary values from these analyses are provided in Table I for reference.

II. BROADBAND LEO CONCEPT OF OPERATIONS

This section presents a summary discussion, not of fused LEO GNSS, but of the type of modern broadband LEO system upon which fused LEO GNSS may be built. This digression is necessary for two reasons. First, fused LEO GNSS is an exercise in re-use. To re-arrange the building blocks of a broadband LEO system into a GNSS, one must first identify these building blocks and understand their operation. Second, a provider contemplating fused LEO GNSS faces costs arising from lost opportunities for profitable broadband service. Quantifying such opportunity costs, as will be attempted in a later section, requires an understanding of the resource constraints that will dominate the cost analysis. This section summarizes relevant portions of what is known, and lays out reasoned speculation about what is unknown, regarding the concept of operations of broadband LEO systems.

Plans for each of the leading broadband LEO projects, including SpaceX’s Starlink, Amazon’s Kuiper, and OneWeb, envision thousands to tens of thousands of space vehicles (SVs) in orbits ranging from 335 km to 1325 km altitude. These systems are proposed to provide broadband, worldwide Internet service to consumers via the K_u and K_a microwave bands between 10 GHz to 30 GHz. Compared with existing infrastructural wireless systems like LTE and Wi-Fi, broadband LEO systems offer greater availability in remote locations at a lesser infrastructural investment. In densely populated regions, the systems are not expected to compete with cable, fiber, and 5G to serve a large fraction of users, but would still offer competitive bandwidth and latency to a limited subset.

To the extent possible, this paper will remain agnostic to the details of any particular broadband LEO system. However, it will be helpful to refer to concrete examples at some points in the discussion. At such points, this paper will refer to information gleaned from public statements and filings with the U.S. Federal Communications Commission (FCC) regarding SpaceX’s Starlink constellation, currently the most mature and ambitious broadband LEO contender, with plans for up to 42,000 LEO spacecraft.

A. Architecture

A broadband LEO SV functions as a cross between a wireless router and a cellular base-station. It operates according to

a dynamic *schedule* which allocates resources of time, space, and frequency to route packets of data among gateways and users. *Users* are paying subscribers equipped with directional transceivers (“modems”), mounted either on static or vehicular platforms. *Gateways* are special ground sites with (potentially) superior antennas, unobstructed skylines, and high-bandwidth, low-latency connections to the Internet.

B. Initialization

When a user’s modem enters the network after power-on, it begins a search for information that will allow it to connect to an SV: orbital parameters, positioning, timing, and antenna orientation. Any parameters that it cannot obtain or recall, it must determine by a guess-and-check strategy. The user modem tunes its phased-array antenna to point in certain directions where an SV might be found, and then either listens for a signal on a particular downlink frequency, or transmits a connection request on a particular uplink frequency. Until it receives a signal or a response, it keeps trying with other directions and frequencies. Once a connection is established, the user modem and network exchange authentication tokens.

C. Steady-state operation

After initialization, operation likely proceeds as in Wi-Fi or cellular communications with central scheduling: the user modem may notify the SV, within certain windows of time and frequency, that it wishes to uplink a packet of e.g. Dynamic Host Configuration Protocol (DHCP) or Internet Protocol (IP) data. At its discretion, the SV grants the user modem an uplink time-slot. The modem listens for such a scheduling directive, waits for the appropriate moment, and then encodes and transmits the packet. Conversely, when a packet addressed to a user arrives at the SV from the Internet via a gateway, the SV queues the packet for downlink transmission. The user modem monitors the downlink channel for messages.

As with Wi-Fi or any other wireless Internet medium, the broadband LEO architecture need not make any special affordance for higher-level abstractions like data fragmentation and re-assembly, reliable and in-order delivery, Internet sockets, or applications like the Web: these features are provided by IP and communication layers built atop it, like the Transmission Control Protocol (TCP) and the Hypertext Transport Protocol (HTTP), in accordance with the so-called end-to-end principle.

D. Error correction and re-transmission

Each transmitted packet is encoded using error detection (e.g. CRC) and error correction (e.g. Turbo codes) mechanisms. If a packet is decoded successfully, the receiver (user modem or SV) sends back an acknowledgement (ACK). If a packet fails to decode, the receiver sends back a negative acknowledgement (NAK). A packet which has been NAK’d may be scheduled for another transmission attempt using automated repeat request (ARQ) or hybrid ARQ (H-ARQ).

Fused LEO GNSS will involve tasking the downlink transmitter to send additional PNT-specific signals. Because the SV knows the timing of these additional signals in advance,

it must schedule around them in provisioning data service: uplink packets and their re-transmissions must be scheduled so that an ACK need not be sent down while the SV transmitter is busy, and downlink packets and their re-transmissions must be scheduled so that no data transmission is required during transmission of a PNT signal. Section III-D offers further details on such constraints.

All broadband LEO systems will likely employ frequency-division duplexing (FDD), whereby the uplink and downlink operations are frequency-disjoint. The alternative, time-division duplexing (TDD), is wasteful for long-distance channels: avoiding simultaneous reception and transmission at both ends requires that the channel frequently sit idle. As with FDD LTE, the user modems will likely be frequency-division half-duplex (one frequency at a time) to keep modem cost low, whereas the SVs, like cellular stations, will support full FDD. This implies that an SV need not wait for an ACK to a downlink packet before it begins transmitting another packet. As will be discussed later, this property is significant for fused LEO GNSS.

E. Multiplexing

General wireless principles offer a number of options for partitioning a shared downlink channel among multiple users, or among distinct data streams, including code-, time-, frequency-, and space-division multiplexing (C/T/F/SDM). Typically, several techniques are employed to achieve high spectral efficiency at modest computational cost. K-band broadband LEO systems operate on a dynamic, wide-band, dispersive channel [15]. These are the same parameters that characterize the radio-layer design of existing mobile wireless networks like LTE and Wi-Fi, and hence one may expect similar solutions to apply.

LTE and Wi-Fi use orthogonal frequency-division multiplexing (OFDM) to eliminate dispersive inter-symbol interference due to, e.g., multipath scattering, and SDM in the form of beamforming to boost signal strength. This paper presumes that broadband LEO systems will also apply OFDM, acknowledging its preeminence in modern wireless technologies. OFDM’s spectrally-flat signals are not theoretically ideal for ranging, but when wide enough offer excellent precision [13]. Beamforming, discussed further in §II-F, is critical to close link budgets and to avoid causing interference. Apropos of fused LEO GNSS, beamforming implies that signals-of-opportunity from broadband LEO SVs cannot be expected to be strong enough for navigation unless a “spot” beam is directed towards the listener’s geographical vicinity.

The third form of multiplexing likely to be in use—and relevant for fused LEO GNSS—is time division. In FDD LTE, the channel is licensed to the mobile network operator, and is provisioned for a single transmitter in a given spatial region. Time division is used for sharing resources between users, with time being divided into frames, subframes, and OFDM symbols, and users being scheduled downlink opportunities in the time-frequency plane. In broadband LEO, one might expect the same to apply. However, continuous operation of the SV transmitter regardless of load would waste a significant

amount of power: wide-band OFDM signals have a high peak-to-average ratio (PAR) and require a high degree of linearity, and PAR tends to drive power requirements for all but the most sophisticated linear amplifiers. Power may be a significant constraint for the operations of broadband LEO systems (§II-K). This paper therefore assumes that broadband LEO transmitters operate in “burst mode”: that is, with intermittent transmissions. Such a mode aligns radio power consumption with load.

F. Beamforming

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 focusing an SV’s transmit power into a narrow beam targeted toward a relatively small ground service region. Each SV therefore must have independently-steerable directional antennas for transmit and for receive operations. Each Starlink craft, for instance, will support 15 downlink beams [14] for user data service.

Because each LEO SV is overhead only as viewed from a relatively small region at a time, a large number of SVs are needed to provide continuous global service. To mitigate launch costs, each SV must be (relatively) light and compact. This configuration strongly favors flat, electronically-steered two-dimensional phased array antennas. Such antennas require costly *phasing networks*. These are (analog or digital) circuits which form programmable combinations of one or more transmitted signals (“beams”) for each one of a large number of radiators (“array elements”). A schematic representation is shown in Fig. 1. The contribution of each beam to each array element is linear, and consists of a delay and/or scaling (from now on, “phasing”). Mathematically, this operation is a complex matrix multiply for each frequency, but such a cavalier description belies significant implementation challenges.




Fig. 1: An electronically steered phased array for transmission. Amplifiers not shown. Columns are beams; rows are array elements. Dotted lines carry coefficients; thick solid lines carry RF signals. For reception, the matrix of phasing elements and accumulators is transposed: the accumulators move from the right (antenna) lines to the top (modem) lines.

For this paper, what is relevant is the procedure for “steering” the array. The phasing elements potentially require many bits of configuration (“coefficients”). During steering, the CPU must compute (or retrieve from a look-up table) a new set of coefficients for the desired beam pattern. These coefficients must then be transferred to the phasing network. To free up the CPU’s I/O resources for other tasks, direct control over the phasing elements would typically be delegated to an external multiplexer, shown in Fig. 1 as a serial-to-parallel converter. (This example architecture is not the only way to design such a circuit; it serves only to facilitate the parameter definitions needed for this paper.) Over some interval $T_{\text{set-up}}$, the CPU loads coefficients into this register. Then, on a signal from CPU, the updated coefficients are simultaneously imposed upon the phasing network. Such an arrangement is favorable because it minimizes the “downtime” during steering, in which the phasing network is in an indeterminate or mixed state and so cannot be relied upon to produce a valid beam pattern. The downtime could perhaps be made as short as a few times the light-crossing time of the phased array: some few nanoseconds. Large or small, the downtime will be denoted by T_{switch} in what follows.

Note that $1/T_{\text{set-up}}$ is an upper limit on the rate at which the array may be re-steered. T_{switch} may be expected to be shorter than $T_{\text{set-up}}$, so while it also imposes an upper limit, it is a less stringent one.

For reasons of cost, it is reasonable to assume that user modems will only be capable of transmitting or receiving a single beam at a time.

G. Channels

The number of disjoint frequency channels available is a function of spectrum licensing and desired system bandwidth to a single user. Public filings for Starlink indicate that the K-band spectrum in the ranges 10.7 GHz–12.7 GHz, 17.8 GHz–18.6 GHz, 18.8 GHz–19.3 GHz, and 19.7 GHz–20.2 GHz will be broken into 76×50 MHz downlink channels. This paper does not assume that all channels are available on all beams, nor that all available channels may be transmitted simultaneously by any one beam. Instead, it assumes that there is some number N_{bc} of “beam-channels” that represents the greatest number of simultaneous transmissions from the SV. For Starlink, this number is at most 264, though it may be less. When concrete values are required in what follows, this paper will assume $N_{\text{beams}} = 15$, $N_{\text{channels}} = 76$, and $N_{\text{bc}} = 264$.

H. Flux

One consideration particular to first-generation broadband LEO proposals has been driven by the opening of spectrum in the K band. Operators can access this spectrum only by adhering to certain limits on the flux of RF energy at the surface of the Earth intended to prevent interference with terrestrial K-band services. For instance, Starlink has declared that the system will generate no more than -122.0 dB(W/m²/MHz) in the 10.7 GHz to 12.7 GHz band as observed at ground level at elevation angles above 25°.

It has been revealed in public filings for Starlink that a single beam from a single broadband LEO SV is powerful enough and directional enough to saturate this flux limit. While a shrewd operator might wish to increase system capacity in densely-populated urban regions by focusing several beams from several SVs onto one location, taking advantage of user antenna directionality to boost spatial re-use, this would not avail: each beam would have to operate with either reduced power or reduced duty cycle to avoid exceeding the flux limit.

Note, however, that the flux limits are per-MHz, so an operator might be permitted to focus multiple beams on the same location, provided these are on disjoint channels. This solution might cause other problems if it resulted in all available channels being occupied serving a dense region, leaving none available to serve a nearby sparse region due to the possibility of same-channel interference.

I. Cellular Provisioning

To avoid same-channel interference between users served by different SVs, broadband LEO operators will establish regions of exclusion within which a channel may not be re-used. The simplest way to do this would be to provision broadband LEO service into a grid of hexagonal cells on the surface of the Earth (Fig. 2). Rather than assigning individual users to individual SVs, users are assigned to a local cell, and the cell is assigned to an SV. Provided that SV antennas are suitably directional, a channel used in one cell may be re-used in other cells so long as these cells are not immediate neighbors. Provided that signals to each individual user are suitably multiplexed in time and/or frequency, a channel may serve a large number of users within one cell, just as in LTE.

One refinement would be to rely on user antenna directionality to avoid same-channel interference in adjacent cells while permitting additional spatial re-use. This might, however, run afoul of the flux limits, and would have to be done with care to maintain a large angular separation between the assigned SVs. Even in the best case, the need for angular separation would tend to push the typical SV assignments for each cell away from the zenith, potentially reducing availability for users with less-than-totally-clear skies.

J. Cellular Scheduling

For reasons of latency and efficiency, it makes sense to delegate, insofar as possible, the global network coordination and packet scheduling problem to the SVs. Setting aside for the moment the question of how SVs know which cells to serve on which beams and frequencies, each SV may independently make millisecond-level decisions about the scheduling of uplink and downlink packets; these decisions have no impact on other cells. The alternative would be for SVs to consult with their gateways, or even with a central server, before a user’s request for data could be satisfied, which would needlessly incur several milliseconds of additional latency and tie up gateway bandwidth.

On the other hand, it would not be easy to fully decentralize the global network coordination problem. One could imagine designing a fixed, pre-arranged schedule of assignments of




Fig. 2: Schematic geometry for geographical provisioning of broadband service for a LEO system covering the range of $\pm 60^\circ$ latitude. Real-world cells would be smaller, with diameters D measuring in the tens of kilometers. The area of the service region is $4\pi R_{\text{earth}}^2 \sin 60^\circ$, and the area of an individual cell is $\frac{3\sqrt{3}}{8} D^2$, ignoring north-to-south variations. For Starlink, 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 $D = 45$ km: this gives a number on the order of 400,000.

SVs to service cells, but this scheme fails on several counts. (1) As massive constellations of broadband LEO SVs are gradually deployed, the schedule must be constantly adjusted to account for unfilled orbits. (2) Dense deployments like Starlink are expected to be spread over a number of “shells” at different orbital altitudes, making their global configuration potentially aperiodic. (3) A pre-arranged schedule fails to accommodate dynamic variations in the health, level of battery charge, and access to gateway bandwidth of individual SVs.

It would be straightforward to periodically re-compute, at a central location, the assignment of SV beams to cells and gateways to SVs, and for commands effecting this schedule to be uplinked to the SVs via the gateways. Whether operating at initial operational capability or full capability, tasking each SV beam to serve many cells or just one, such central coordination would give the operator much-needed flexibility.

How often must this computation take place? The orbital motion of LEO SVs across the user’s sky plays out over a period of several minutes, so that during the interval from when an SV becomes the “most suitable” to serve a given cell to the time that it is no longer the most suitable, hundreds of thousands of packets may be sent. Assignments of gateways to SVs will evolve over similar time-scales.

Downlink scheduling plays a critical role in costing out the fused LEO concept, which involves additional steering costs and additional transmissions. For systems like Starlink that use

frequency-division duplexing, with uplink channels disjoint from downlink channels, the uplink schedule is expected to be largely unaffected by fused LEO GNSS.

Will broadband LEO SVs be able to steer their (possibly paired) transmit and receive beams separately and use them simultaneously? Public filings do not make this clear. However, if the differences in wavelength between the uplink and downlink channels are significant, as is the case for Starlink, it seems reasonable to assume that a separate set of phasing coefficients would be required for each, if not completely separate antenna hardware, implying independently-steerable uplink and downlink antennas at the SV.

K. Power Management

Broadband LEO constellations can be expected to operate with highly efficient power management strategies. Their SVs’ solar arrays will be sized just large enough to meet some low multiple of the expected average energy demand per orbit at end of life, and their batteries will be sized just large enough to meet some fraction of peak regional demand, and to sustain operations during eclipse [16]. Regional or global scheduling algorithms will optimize energy usage across the constellation by tasking each active SV to nearly exhaust its collected energy per orbit.

This is not to say that SVs will operate continuously at or near their maximum load power. Rather, they will be commanded to enter a “deep sleep” minimum-power state during a portion of their orbit while nearby SVs carry the burden of providing uninterrupted service. Such duty cycling is rational because, just as for smartphones [17], an SV’s useful work is a nonlinear function of its expended energy. Greater efficiency obtains when a subset of SVs operates near maximum power and the complement operates at minimal power than when each SV is active but not fully tasked. One may expect that over oceans a majority of SVs will be idled whereas over some populated areas all will be fully tasked.

The constellation enters a scarce energy regime when the constellation-wide energy collected per orbit is inadequate to support the demanded communications operations. One might expect the constellation to be designed never to enter this regime, given that opportunities for profitable exchange of data are lost. But as with electricity provision [18] and terrestrial broadband provision [19], it is wasteful to design a network for peak demand if the peak-to-average demand ratio is high. LEO broadband providers are likely to accommodate congestion just as their terrestrial counterparts do, namely, by throttling data speeds or by implementing time-dependent pricing [19].

L. Visibility

For K-band operation in particular, the ITU Radio Regulations require broadband LEO operators to avoid SV-to-user lines-of-sight that pass too close either to geostationary orbit or to the horizon. SV beams cannot be assigned to serve cells for which the line-of-sight falls into one of these “exclusion masks.” The two types of masks reduce the number of available SV beams at low and high latitudes, respectively.

III. FUSED LEO GNSS CONCEPT OF OPERATIONS

To achieve latency and performance competitive with traditional GNSS, fused LEO GNSS will employ single-epoch code-phase pseudorange using bursts transmitted from four or more SVs. This will require changes to broadband LEO operators' coordination and scheduling algorithms: multiple SVs must provide time-multiplexed ranging signals to each cell to support pseudorange-based PNT. These broadcast ranging signals are an addition to the broadband LEO system, but they share the same modulation and encoding as the data service.

A. Obstacles

A comprehensive proposal for fused LEO GNSS must address certain key obstacles: Broadband service in a given cell is expected to be provided by a single SV for minutes at a time. Multi-SV-to-cell pseudorange will require changes to both the central scheduler that matches SV beams with cells, and to the onboard schedulers that must avoid collisions with cross-cell ranging signals (i.e., those from other SVs). (§III-D) The downlink transmitters are not designed to produce traditional ranging waveforms. (§III-C) The broadband LEO SVs' positions and clock offsets may not be known to high enough accuracy either to the broadband system or to user modems, so a precision orbit and clock determination capability is needed. (§III-E) The onboard clocks are likely not stable enough for nanosecond-accurate forecasting beyond a minute, so a "zero age-of-ephemeris" solution is required. (§III-C) Finally, the SVs have limited power and limited resources for downlink scheduling, array steering, and command-and-control signaling. Any resources diverted from the constellation's primary communications mission must be paid for by PNT customers. (§IV)

B. System Overview

The underlying broadband LEO system may be depicted as in Fig. 3 in terms of the relationships between SVs, beams, cells, and users. The fused LEO GNSS concept in Fig. 4 mirrors this structure with two alterations. First, SV beams may receive additional "secondary" cell assignments. Each beam will broadcast periodic ranging signals to all of its assigned cells, primary and secondary, but it will only provide broadband service to the cells for which it is primary. Second, whereas broadband connectivity is tied back to a gateway, PNT is tied back to traditional GNSS signals, as observed via GNSS receivers on the SVs. In this way, LEO provides a second "tier" of PNT service, with traditional GNSS serving as the first tier (§III-E). Note that, although dependent on traditional GNSS, such multi-tiered fused LEO GNSS is well-protected from terrestrial GNSS interference sources: L-band signal spreading loss to LEO is more than 145 dB [20].

The total number of assignments (primary or secondary) for each served cell is equivalent to the number of signals provided for pseudorange, denoted n . To fully constrain the three-plus-one dimensional PNT solution, $n \geq 4$. When a concrete baseline value is required, $n = 5$ will be assumed in what follows. For a customer willing to pay a high price for ultimate

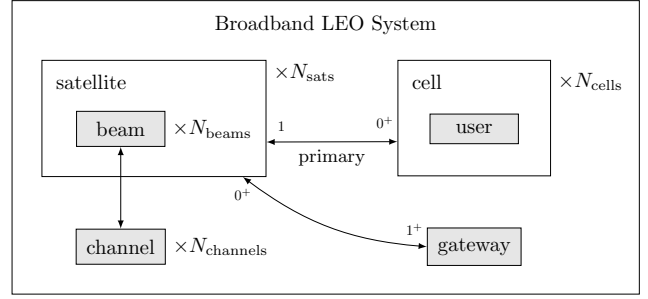


Fig. 3: Entities (boxes) and their relationships (arrows) in a broadband LEO system. Numbers by arrowheads indicate the cardinality of a relationship: for instance, the number of cells for which an SV is primary is zero or more, while the number of SVs which are primary for a cell is exactly one.

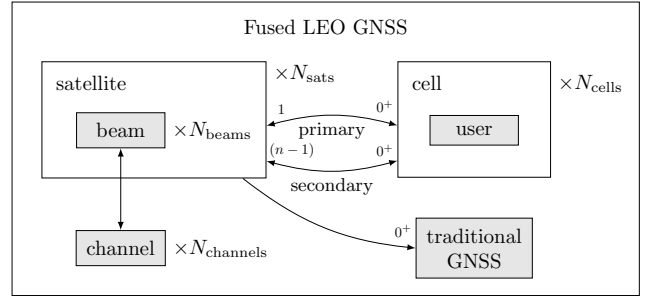


Fig. 4: As Fig. 3 but for fused LEO GNSS, with n denoting the number of ranging signals provided to users in each cell.

performance, n may be as large as the full number of SVs in view, circa 40 [13]. (While it may be less costly to send more signals from a few SVs, the benefits would be lesser due to more correlated user ranging errors.)

Two SVs in a ranging scenario are depicted in Fig. 5. Users in the service cell require a diversity of directions-of-arrival to obtain a robust positioning and timing solution. To maximize geometric diversity with $n = 5$, the operator could assign to a cell those SVs near the vertices of a square pyramid; e.g., the northernmost, easternmost, southernmost, westernmost, and zenith-most SVs in view among those that are not in an exclusion mask. (The pyramid need not align with the cardinal directions.)

C. Pseudorange Service

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 SV transmitter, 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 adopts 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 dominated by other sources of error.




Fig. 5: Two of the n SVs providing ranging signals to a given cell from different points in the sky. At least $n = 4$ signals from distinct directions are needed to solve for x, y, z , and time. The west-to-east axis of motion shown here is arbitrary, but consistent with Fig. 6. Angles θ_1 and θ_2 are the zenith angles corresponding to TX1 and TX2.

Service consists of a series of ranging bursts, each modulated just like ordinary broadband data, with three exceptions. First, the user modem must be able to regenerate the transmitted waveform and perform cross-correlation, so the data bits encoded and modulated to form the ranging burst must be largely known in advance to the modem. This cross-correlation results in a code-phase time-of-arrival measurement, which may then be compared with the nominal time-of-departure of the burst to form a pseudorange measurement. Second, the burst is very short. A duration of $T_{\text{burst}} = 500 \mu\text{s}$ is more than adequate: the contribution to ranging error due to finite burst duration is far less than other sources of error [13]. For this reason, it is also 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 ephemerides. This portion of the burst is ignored during correlation. Because the entire clock and orbit ephemeris fits into a small fraction of a single ranging burst—which might easily accommodate tens of kilobits of data given the reported bandwidth—user pseudoranges need never use stale or forecasted data. This “zero age-of-ephemeris” eliminates the need for atomic clocks on the LEO SVs. Third, the burst is not acknowledged by one ground receiver (i.e. unicast), but instead is broadcast to all receivers in the cell.

One significant challenge here is that, as noted in [15], the K-band channel is dispersive, with a worst-case coherence bandwidth of only 3 MHz. A naïve cross-correlation on a 50 MHz-wide signal would not produce a peak with width approaching, e.g., $1/(50 \text{ MHz})$; instead, a broad, incoherent peak would be expected with width on the order of $1/(3 \text{ MHz})$. The dispersion on the channel may be decomposed into factors: (1) the frequency response of the transmit filters, amplifiers, and phased array; (2) the frequency response of the atmospheric channel; (3) the aggregate effect of multipath scattering; and (4) the frequency response of the receive array,

filters, and amplifiers. Each of these can be managed in fused LEO GNSS: atmospheric dispersion at K-band is negligible for the bandwidths envisioned (Hobiger et al. [21] report sub-millimeter delay sensitivity to dry air pressure, water vapor, and surface air temperature for a 200 MHz-wide K_u -band signal), non-line-of-sight multipath effects will be suppressed by the directionality of the receive phased array, and the transmit and receive frequency responses may be estimated using the training preamble that is in any event required for OFDM.

D. Global Scheduling

This paper draws a distinction between two objects which might be termed “schedules.” *Global schedules* are computed centrally and in advance by the broadband LEO provider, and consist of assignments of SV beams to provide service to certain cells at certain times. *Local schedules* are computed on-the-fly on each SV as packets arrive and are dispatched. Global schedules may be valid for multiple minutes, and can be forecast due to orbital predictability. Local schedules cannot be forecast due to the unpredictable timing of packets on the Internet.

The global scheduler’s role in broadband LEO is to compute a conflict-free assignment of SV beams to cells: the “primary” assignments in the parlance of fused LEO GNSS. These assignments respect power, visibility, and exclusion mask constraints. In fused LEO GNSS, the global scheduler must additionally solve a system of constraints on the inter-departure and inter-arrival times of ranging bursts, particularly those “secondary” ranging bursts directed to a cell from a non-primary SV. Care is required in scheduling these cross-cell ranging bursts so as to avoid collisions on the ground. In Fig. 6, bursts TX1 and TX2 from two different SVs arrive in the same cell. The wavefronts of each burst sweep across the cell in different directions, consistent with the geometry of Fig. 5. If a user at the eastern edge of the cell is to decode both bursts (for instance, if the first burst is data addressed to this user and the second burst is a ranging broadcast), then an interval $T_{\text{switch}}^{\text{RX}}$ must be allowed between the end of the first burst and the arrival of the second for the user’s modem to re-tune its antenna.

1) *Feasibility*: A global schedule is *feasible* if and only if it satisfies all feasibility constraints at each transmitter, and all feasibility constraints at each receiver. (Recall that only the downlink is changed in fused LEO GNSS, so the transmitter here is one SV beam, and the receiver is one user modem.) Let *GNSS scheduler* refer to the subroutine of the global scheduler concerned with fused LEO GNSS operations, and let the *GNSS schedule* be its product. GNSS schedules are periodic. One may visualize their period T_{period} as one revolution of the hand of a clock (Figs. 7 & 8), with the schedule of ranging bursts repeating after this interval. The relative timing of events differs between transmit and receive clocks due to time-of-flight effects. Transmit and receive constraints under which the GNSS scheduler operates are shown in the two figures: bursts must not overlap at either the transmitter or the receiver, but two bursts to the same cell (or two bursts from the same




Fig. 6: Satellite-to-ground bursts shown in two axes of space (east-west and up-down) and one of time. The bursts begin at space-time points TX1 and TX2, and each continues for a duration of T_{burst} . Signals radiate outward in a light-cone, intersecting the ground at $z = 0$ in a segment of a hyperbola. The interval until the signal first reaches the ground somewhere in the service cell is T_{flight} . The interval from that time until the signal has reached the entire cell is T_{sweep} .

SV) may be back-to-back. However, bursts to different cells (or bursts from different SVs) must be separated by a suitable interval so that the transmitter and receiver can steer their antennas, and so that time skew across the cell (T_{sweep}) is taken into account. If some users do not need PNT, additional flexibility is possible; but this paper will make conservative scheduling assumptions:

- 1) Every cell receives n ranging bursts per T_{period} .
- 2) Bursts and/or TX switching intervals from one beam-channel of an SV do not overlap in time.
- 3) Bursts from one SV beam to different cells are separated by at least the TX switching interval T_{switch}^{TX} .
- 4) Bursts and/or RX switching intervals on one channel do not overlap in time from any viewpoint in the target cell.
- 5) Bursts to one channel in one cell from different SVs are separated by T_{switch}^{RX} from any viewpoint in the target cell.
- 6) TX switching events are separated by at least T_{set-up}^{TX} .
- 7) RX switching events are separated by at least T_{set-up}^{RX} .
- 8) Bursts to neighboring cells on the same channel are non-overlapping in time.

This paper assumes, based on the implementation of serial-to-parallel converters using, e.g., paged register files, that constraints 6 and 7 do not apply in the case of *switching back to the most recent coefficients*. That is, the array may be switched to new coefficients without erasing the old coefficients from its memory. Under this assumption, switching forward and then switching back requires only the short interval $2T_{switch}$, rather than the longer interval $2T_{switch} + T_{set-up}$.

A GNSS schedule consists of the following data:

- 1) T_{period}




Fig. 7: Scheduled events at each transmitter (i.e., SV beam) must satisfy feasibility constraints: bursts may not overlap, and bursts to different cells must be separated by at least T_{switch}^{TX} .

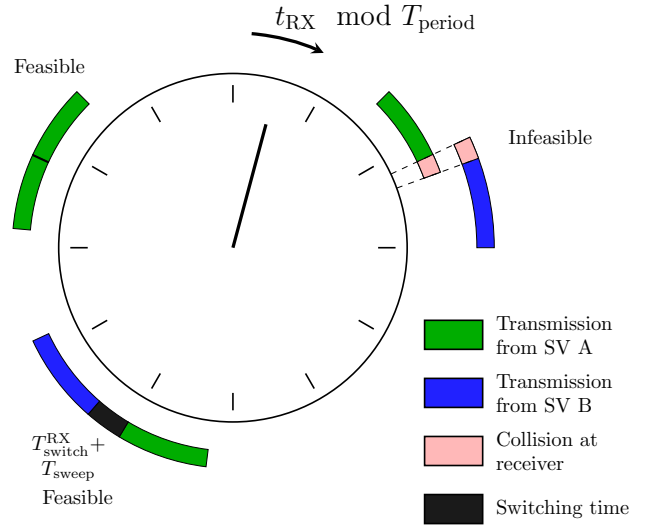


Fig. 8: Scheduled events at the receiver (i.e., user modem) must satisfy feasibility constraints: bursts may not overlap, and bursts from different SVs must be separated by at least $T_{switch}^{RX} + T_{sweep}$.

- 2) For each cell k , and for each ranging signal $s = 1 \dots n$, a tuple (SV, channel, time modulo T_{period} , T_{flight} , T_{sweep}) indicating that a ranging burst will be sent by this SV on this channel towards this cell at such times.

Timing information may be quantized to a resolution of $1 \mu s$ without any significant loss of scheduling flexibility. Each SV must be provided with all assignment tuples involving either itself or any of its primary cells. The time-of-flight data in secondary (cross-cell) assignment tuples allows each SV to determine when it is safe to schedule data bursts. Optionally, the timing and time-of-flight parameters for primary ranging bursts may be left out of the GNSS schedule, to be determined

by local scheduling on each individual SV.

2) *Optimization*: The GNSS scheduler must be an efficient algorithm for finding feasible global schedules. Among feasible schedules, it will also optimize: the GNSS scheduler will prefer assignments that minimize the impact on local data scheduling, maximize the geometric diversity of ranging signals in each cell, and minimize system power consumption. The relative importance of these objectives could depend on the prices paid by customers for broadband and PNT service, and on the number of customers of each type in each cell. Given the complexity and non-convexity of this problem, approximations or heuristics will likely be required.

3) *Complexity*: How complex is the global scheduling problem, and how over- or under-constrained is the feasibility problem? One way to grapple with this question is to consider a simple greedy baseline GNSS scheduler that runs after the global broadband scheduler has made its assignment of primary SV beams to cells over the next planning window. The GNSS scheduler allocates a “transmit cube” array to represent the ranging-specific status of SV i , beam-channel bc at time $t \bmod T_{\text{period}}$ (Fig. 9). Recalling that data service is scheduled locally, and hence does not appear in the global schedule, the enumerated status values stored in the TX cube are Idle, Burst, and Switch. The GNSS scheduler also allocates a “receive cube” array to represent the ranging-specific status of cell k , channel j at time $t \bmod T_{\text{period}}$ (Fig. 10). The enumerated status values stored in the RX cube are Idle, Burst, Switch, and Exclude. During construction of the TX and RX cubes, the GNSS scheduler consults an “availability” array, which indicates whether SV i is available to provide signals to cell k , or is excluded from doing so for any reason.

The greedy baseline GNSS scheduler iterates over cells k . For each cell, it iterates over ranging signals $s = 1 \dots n$. For each signal, it iterates over available SVs i . To maximize geometric diversity for $n = 5$, a “goal direction” is defined for each s , and available SVs are iterated in descending order of alignment of their line-of-sight vectors with the goal direction. The goal directions form a square pyramid of points in the sky [22]: for $s = 1$, the zenith; for $s = 2$, local north; for $s = 3, 4, 5$, local east, south, and west, respectively. The greedy scheduler proceeds from SV to SV in this order, attempting to find idle space in the corresponding planes of the TX and RX cubes to make an assignment. When it finds a suitable space, it adds the assignment tuple to the output GNSS schedule, updates the TX and RX cubes, and proceeds to the next signal $s + 1$. If all available SVs are exhausted and n signals have not been assigned for cell k , then the greedy scheduler fails. If all cells have been assigned n signals each, then the greedy scheduler succeeds.

Is this procedure likely to succeed? Does the answer depend upon whether cells are iterated in geographic order, or randomly? Or upon whether iteration is instead first over signals and then over cells? A key clue is that the TX and RX cubes remain overwhelmingly Idle in any event, as will be shown.

4) *Schedule Sparsity*: If one pre-supposes that a feasible schedule exists, it consists of $N_{\text{cells}} \cdot n$ assignment tuples, of which a fraction $(n - 1)/n$ are secondary, and the rest are primary. The fraction of the TX cube that is non-Idle is then

equal to the sum duration of these assignments, divided by the dimensions of the cube. This fraction will be called the *transmit reservation*. Similarly, the fraction of the RX cube that is non-Idle will be called the *receive reservation*.

Each primary assignment reserves a single downlink SV beam-channel for an interval T_{burst} . Each secondary assignment reserves the entire SV beam for $T_{\text{burst}} + 2T_{\text{switch}}^{\text{TX}}$. Summing up all the primary and secondary assignments and dividing by the TX cube volume $N_{\text{bc}} N_{\text{sats}} T_{\text{period}}$, one obtains TX reservation

$$\text{TX res.} \leq \frac{N_{\text{cells}} \cdot \left[T_{\text{burst}} + (n - 1)(T_{\text{burst}} + 2T_{\text{switch}}^{\text{TX}}) \frac{N_{\text{bc}}}{N_{\text{beams}}} \right]}{N_{\text{bc}} N_{\text{sats}} T_{\text{period}}}$$

where the inequality allows for possible over-counting of switching times, and non-uniform beam bandwidth is ignored.

Similarly, each primary or secondary assignment reserves one downlink channel in a cell for an interval of up to $T_{\text{burst}} + T_{\text{sweep}}$, and each secondary assignment additionally reserves the channel for $2T_{\text{switch}}^{\text{RX}}$. Finally, constraint 8 requires that each cell’s $N_{\text{neighbors}}$ adjacent cells be reserved for time $T_{\text{burst}} + T_{\text{sweep}}$ for each assignment. Summing over assignments and dividing by the RX cube volume $N_{\text{cells}} N_{\text{channels}} T_{\text{period}}$, one obtains

$$\text{RX res.} \leq \frac{n(N_{\text{neighbors}} + 1)(T_{\text{burst}} + T_{\text{sweep}}) + 2(n - 1)T_{\text{switch}}^{\text{RX}}}{N_{\text{channels}} T_{\text{period}}}$$

$$T_{\text{sweep}} \leq \frac{D}{c} \cos(\phi_0),$$

where the first inequality allows for possible over-counting of switching times and adjacent-cell exclusions, and ϕ_0 is the minimum elevation mask angle.

To obtain order-of-magnitude estimates for these reservations in practice, let $N_{\text{cells}} = 400\,000$, $N_{\text{sats}} = 10\,000$, $n = 5$, $N_{\text{channels}} = 76$, $N_{\text{beams}} = 15$, $T_{\text{burst}} = 500 \mu\text{s}$, $T_{\text{switch}} = 100 \mu\text{s}$, and $T_{\text{period}} = 1 \text{ s}$. For T_{sweep} , assume a minimum elevation angle $\phi_0 = 40^\circ$ and a cell diameter $D = 45 \text{ km}$, and let c be the speed of light. Let $N_{\text{neighbors}} = 6$. Then

$$\text{TX reservation} \leq 0.75\%$$

$$\text{RX reservation} \leq 0.03\%$$

That the TX and RX cubes are overwhelmingly Idle is significant because the scheduling problem within each cube is equivalent to the graph coloring problem [23] for which upper bounds exist on the number of required colors (analogous to the number of time slices of the cube) for various algorithms [24]. So long as the cube reservations are sparse, feasible schedules can be found with even very simple algorithms [23]. One may thus presume that the feasibility problem is greatly under-constrained, and the optimization problem will be of greater interest than the feasibility problem in future work.

The global TX and RX reservations also equal the mean reservations for any given SV beam-channel or cell, respectively. Thus, the impact on the local scheduling problem should be small as well. This paper accordingly does not elaborate on local scheduling.

One might question whether all, or only a subset of SVs in a broadband LEO system ought to be involved in fused LEO GNSS. Indeed, the low reservation numbers indicate that such sub-setting would be possible. However, due to switching




Fig. 9: The TX cube. The GNSS scheduler uses such a representation to track when each (beam, channel) pair of each SV has been reserved for ranging. Burst durations and cube occupancy are greatly exaggerated for clarity of visualization.




Fig. 10: The RX cube. The GNSS scheduler uses such a representation to track when each channel in each cell has been reserved for ranging. The true RX cube has a more complicated adjacency relationship than shown here, because service cells form a 2-D hexagonal grid rather than a 1-D line. Burst durations and cube occupancy are greatly exaggerated for clarity of visualization.

times, it is less costly for an SV to provide ranging bursts to its primary cells than otherwise. It is preferable, then, to task every SV having one or more primary cells to provide ranging service to those cells. Load shifting over regions with few subscribers could mean that some SVs are placed into power-saving modes, while the remainder are tasked close to 100% capacity. In this case, GNSS duties would also be concentrated as much as possible to minimize the number of active SVs.

E. Orbit and Clock Determination

Both hosted payload and fused approaches to LEO GNSS require continual, highly-accurate estimates of the SV orbital ephemerides and of the time offsets of the spaceborne clocks. Such estimation should ideally take advantage of constellation-to-ground ranging, intra-constellation ranging, and onboard GNSS receivers. In the near term, however, broadband LEO providers should not be expected to build out an observational network of ground stations extensive enough (e.g., covering ocean regions) or intra-constellation ranging accurate enough, to significantly improve on orbit and clock determination based on onboard GNSS receivers, which can be expected to constrain forecasting uncertainties for clocks and orbits to 14 cm RMS after one second, as estimated [13] based on the report of Montenbruck et al. [25].

Thus, in the near term, LEO GNSS will operate in what Reid et al. refer to as a multi-tier GNSS architecture [10], with each SV carrying, and dependent on, 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, but provides highly jam-resistant signals to users on the ground.

IV. COST MODEL

The opportunity cost of PNT provisioning for fused LEO GNSS is defined as the value of the best alternative use of the same resources. The resources of a broadband LEO system have two uses: broadband service, and PNT service. So the opportunity cost of fused LEO GNSS is the value of *foregone broadband service*. In order to determine what quantity and value of broadband service is foregone, one would need to know both the load on the system—since idle resources cannot be wasted by putting them to profitable use—and the billing structure for broadband data.

Neither is publically known. To accommodate the first issue, this paper will focus on *worst case* opportunity cost: the lost opportunities for profitable broadband service assuming that the broadband LEO system is *loaded at 100% at all times*, i.e., that there is at least one constellation resource that is 100% utilized, the so-called bottleneck resource. To address the second issue, rather than giving costs in monetary units, this paper will express costs as *reservations*: percentages of potential bottleneck resources.

Lost opportunities for broadband data transmission may be interpreted differently depending on the subscription model. For a given number of users with fixed load patterns, opportunity costs manifest as small penalties to latency and throughput. If the number of subscriptions is variable but subscriber

bandwidth and latency are held constant, then opportunity costs manifest as a reduction in the number of subscriptions that may be sold. The former model is more likely to apply for broadband LEO: guarantees in consumer Internet service tend to be aspirational because quality of service can be more nimbly modulated than number of subscribers.

This section will consider potentially-constraining resources one at a time, supposing that each, in turn, is the bottleneck resource. One presumes that a prescient broadband LEO system operator would provision the constellation such that each constraint is nearly binding in practice; otherwise, costs could have been cut in some sub-system(s). Discussion and interpretation of the results follows in §IV-E.

A. Downlink Capacity

If an SV’s local scheduling problem becomes over-constrained—that is, if no more downlink data packets can be scheduled, but packets are available from the gateway—then downlink capacity becomes a bottleneck resource. In this case, either data or ranging service is delayed, depending on what promises of latency, throughput, or reliability are in effect.

Consider one SV beam and one cell for which it is primary. In the worst-case scenario, the relevant TX and RX reservations are almost entirely non-overlapping: none of the secondary (cross-cell) ranging bursts originating from the SV beam are timed to coincide with any of the secondary ranging bursts directed towards this cell from other SVs. In this case, the portion of the beam’s local schedule rendered unavailable for data service due to ranging is union-bounded by the sum of the TX and RX reservations. By the linearity of expectation, this local reservation for a given beam and primary cell has a mean bounded by

$$\text{mean} \left(\frac{\text{downlink}}{\text{reservation}} \right) \leq \text{TX res.} + \text{RX res.} \leq 0.78\%$$

B. Phased Array Set-Up

Another constellation resource consumed by fused LEO GNSS is serial bandwidth for phased array set-up (§II-F). Broadband LEO SVs need to steer their arrays at angular rates up to 0.73°s^{-1} to compensate for orbital motion. This is somewhat smaller than the angular full width at half maximum (FWHM) of the beam, indicating that array steering must be adjusted on the order of once every few seconds. In the scenario where one SV beam provides broadband service to multiple cells, this increases to a few steering events per served cell per round-trip time (RTT), in order to avoid increased latency. For a target latency of 20 ms [26], downlink arrays will likely be designed for a small multiple of 50 steering events per second. To pick a concrete number, suppose that loading the array with coefficients requires $T_{\text{set-up}} = 5 \text{ ms}$.

Fused LEO PNT requires each array to be steered frequently to provide secondary ranging bursts. If the GNSS schedule is designed jointly with the broadband schedule, so that a beam’s broadband service assignments coincide with its ranging service assignments, then the need for additional steering events may be reduced. Otherwise, this resource limits opportunities for assigning many cells’ broadband service to the same SV.

Recall (§III-D1) the assumption that “switching back” has no set-up cost: 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.

Steering bandwidth is the least abundant resource considered so far. The GNSS schedule includes a total of $(n-1) N_{\text{cells}}$ secondary assignments, each requiring one SV beam’s phasing array to be busy loading coefficients for an interval $T_{\text{set-up}}$ once per T_{period} . Considered as a fraction of total steering set-up capacity in the same way as the TX reservation,

$$\text{mean}(\text{set-up reservation}) = \frac{(n-1) N_{\text{cells}} T_{\text{set-up}}}{N_{\text{beams}} N_{\text{sats}} T_{\text{period}}} = 5.3\%$$

C. Power

Available electrical power becomes a bottleneck resource when the constellation operates in a scarce energy regime (§II-K). Opportunity cost of fused LEO GNSS attains a maximum in this regime: a Joule expended on a ranging burst is a Joule that could have been spent on broadband service. Moreover, if the SV’s battery is sized so that peak power loads are amortized across the orbit, then the cumulative energy expended per orbit that drives the calculation of opportunity cost, in which case it makes no difference at what point in the orbit the ranging burst is transmitted, whether over a sparsely- or densely-populated area: the opportunity cost per burst remains the same (maximum) value for every burst.

As discussed in §II-K, SVs will be commanded to enter a “deep sleep” state during some portion of their orbit, and otherwise may be expected to operate near capacity. The burden of fused LEO GNSS in any given region is distributed among the active SVs. Let E_{burst} be the average energy expended by an SV to transmit a ranging burst via a single beam, let E_{orbit} be the average energy expended by an SV per orbit, and let T_{orbit} be the average orbital period. Then the mean energy reservation, or the mean energy allocated per orbit for fused LEO GNSS as a fraction of total-constellation mean energy expended, can be approximated as

$$\text{mean}(\text{energy reservation}) = \frac{n N_{\text{cells}} T_{\text{burst}} \text{PAR}}{N_{\text{bc}} N_{\text{sats}} T_{\text{period}}}$$

where $\text{PAR} = N_{\text{bc}} E_{\text{burst}} T_{\text{orbit}} / (T_{\text{burst}} E_{\text{orbit}})$ is the peak-to-average power ratio for each SV. Intuitively, the energy reservation increases with PAR because a larger PAR implies a greater fraction of total energy spent per burst.

One can speculate about the value of PAR by recognizing that it will approximately equal the peak-to-average population density ratio of subscribers to the broadband service who are visited by each SV over its orbit. (This assumes that the average power allocated to fused LEO GNSS is a small fraction of each SV’s total average power load.) But one must bear in mind that subscriber density will not be a linear function of actual population density because dense population centers will be better served by terrestrial broadband services.

Let p_{max} be the maximum population density that can be competitively served by broadband LEO systems. A value for p_{max} can be approximated for the developed world by

accumulating population density cells from the GPWv4 population model within the United States [27], in order from least to most populated, until the summed population reaches 42 million, which is the estimated number of U.S. citizens without options for purchasing terrestrial broadband Internet (circa 2020) [28]. The population of the densest cell thus accumulated is 63.2 people/km². This can be mapped to a global average value of $p_{\text{max}} = 92.7$ people/km² by noting that the ratio of broadband penetration in developed countries like the U.S. to the global average is 122/83 (circa 2019) [29].

Again using GPWv4, one can examine the peak-to-average population density ratio along a LEO SV’s ground track for all regions with density less than p_{max} . Population density cells are smoothed over the average availability footprint of the LEO SVs, and ground tracks are averaged over longitude of the orbital ascending node. For a LEO constellation at an altitude of 550 km and an orbital inclination of 53° (Starlink’s proposed B sub-constellation [14]), the peak-to-average population density ratio resulting from such calculations is 9.6 (see Supplement B), which may be taken as a proxy for PAR. Filling in the other quantities involved in the mean energy reservation with the values given in §III-D4, one finds

$$\text{mean}(\text{energy reservation}) = 0.36\%$$

D. Command-and-Control Bandwidth

Fused LEO GNSS will require only a negligible amount of command-and-control bandwidth. This paper has assumed that each SV independently performs precision orbit determination using an on-board GNSS receiver. This requires up-to-date communication of traditional first-tier GNSS orbits and clocks from the gateway, comprising a few kilobits per LEO SV every 15 min. The GNSS schedule must also be distributed, sending each SV its assignments and those secondary assignments which affect its primary cells (§III-D1). Overall, each primary assignment in the GNSS schedule will be sent to exactly one SV, and each secondary assignment will be sent to exactly two: the transmitter, and the SV that is primary for the target cell. The total bandwidth used distributing assignments is therefore $(2n-1) N_{\text{cells}}$ times the size in bits of one assignment. Examining the range and quantization of each parameter in an assignment tuple, one finds that the tuple may be encoded in 59 bits. The uplink bandwidth used by fused LEO GNSS is dominated by the cost of uplinking assignments. Summed over the entire constellation, this may be computed as

$$\text{assignment uplink cost} = (2n-1) N_{\text{cells}} \cdot 59 \text{ b} \lesssim 25 \text{ MiB}$$

The data portion of each primary ranging burst will include a copy of the secondary assignments. Users also need the line-of-sight direction to each secondary SV for steering. It is sufficient to provide orbits accurate to ~ 10 km for this purpose: once a secondary burst is decoded, the user will have high-accuracy ephemeris for that SV.

These assignments and ephemerides need to be updated on the SV no more than a few times per minute. On a per-SV basis, the total uplink cost is perhaps 1.7 kbps.

E. Discussion

Regarding the downlink reservation: If the broadband data load on an SV keeps its transmitters busy less than 99.2% of the time, then this reservation represents essentially zero opportunity cost. Recall that ranging bursts do not involve gateway or inter-satellite link traffic. It may be that gateway retransmissions or temporary link interruptions cause the SV's packet buffers to be empty more than 0.8% of the time, even under heavy offered broadband load. If buffers are kept deliberately small for low latency [26], link utilization might be bounded away from 100% due to the behavior of TCP. Like water poured into the interstitial spaces in a jar "full" of sand, the channel reservation for fused LEO GNSS ranging bursts may not displace any data traffic at all. In any event, a global schedule providing $n = 5$ ranging signals per second to every cell between $\pm 60^\circ$ latitude would tie up no more than 0.8% of system downlink capacity. In a rough estimation, this is comparable to adding one user to each cell consuming 5.9 Mbps of broadband service (see Supplement A).

The energy reservation, at 0.36%, may be optimistically low, as it assumes $N_{bc} = 264$, and hence that the SV's potential for peak power output is high. The code in Supplement A permits the reader to explore alternative scenarios.

The significant mean set-up reservation, at 5.3%, invites additional discussion. Does this make fused LEO GNSS impractical? A mitigating factor is that broadband LEO service does not fail catastrophically when beam steering is overtaxed. Instead, the link budget degrades slightly, as beams drift out of perfect alignment, and spill-over into adjacent cells increases slightly. Since beam profiles vary quadratically about the center of the main lobe, a 5% increase in time between steering corrections will increase maximum pointing loss by roughly $2\times$ this percentage: if pointing loss was 0.5 dB, it may now be 0.6 dB.

One should also bear in mind that the foregoing mean reservations have all been calculated assuming a global fused LEO GNSS service, without regard to the distribution of users. An alternative model would provide service only to areas where subscribers are located. For even greater flexibility in matching supply with demand, such targeted service could be paired with time-varying subscription rates. Under this model, inflexible customers demanding continuous high-accuracy LEO GNSS could obtain it, but at a significant cost during periods and within regions of peak broadband demand. Conversely, subscribers willing to accept opportunistic LEO GNSS service could obtain it cheaply when and where its provision presented a near-zero marginal cost to the broadband LEO system.

V. CONCLUSION

This paper presented a concept of operations for fused LEO GNSS, enabling the exploitation of powerful new broadband LEO constellations for global positioning, navigation, and timing (PNT). It laid out a summary and analysis of what is publically known and what may reasonably be inferred about broadband LEO systems, insofar as this information is needed to explore dual-purposing these systems for PNT. Finally, it

analyzed the opportunity cost to constellation providers for re-allocating resources to provide a fused LEO GNSS service. For a constellation such as SpaceX's Starlink to provide continuous service to 99.8% of the world's population would require reserving at most 0.8% of system downlink capacity, 0.36% of system energy capacity, and a trifling amount of uplink capacity. This provisioning scenario would also reserve 5.3% of the constellation's capacity for beam-steering, leading to at most a 0.1 dB increase in maximum pointing loss.

ACKNOWLEDGMENTS

The authors wish to thank Lakshay Narula and Daniel LaChapelle for valuable assistance. Research was sponsored by the Army Research Office under Cooperative Agreement W911NF-19-2-0333. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

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, *Position, Navigation, and Timing Technologies in the 21st Century*. Wiley-IEEE, 2020, vol. 1, ch. Civilian GNSS Spoofing, Detection, and Recovery.
- [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.
- [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, and A. Perkins, "Satellite navigation for the age of autonomy," in *Proceedings of the IEEE/ION PLANS Meeting*. IEEE, 2020.
- [13] P. A. Iannucci and T. E. Humphreys, "Economical fused LEO GNSS," in *Proceedings of the IEEE/ION PLANS Meeting*, 2020.
- [14] SpaceX, "SpaceX non-geostationary satellite system, Technical Parameters," https://licensing.fcc.gov/myibfs/download.do?attachment_key=1877844, Aug. 2019, SAT-MOD-20190830-00087.
- [15] E. L. Cid, M. G. Sanchez, and A. V. Alejos, "Wideband analysis of the satellite communication channel at Ku- and X-bands," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 4, pp. 2787–2790, 2015.

- [16] E. Mostacciolo, L. Iannelli, S. Sagnelli, F. Vasca, R. Luisi, and V. Stanzione, "Modeling and power management of a leo small satellite electrical power system," in *2018 European Control Conference (ECC)*. IEEE, 2018, pp. 2738–2743.
- [17] R. Pérez-Torres, C. Torres-Huitzil, and H. Galeana-Zapién, "Power management techniques in smartphone-based mobility sensing systems: A survey," *Pervasive and Mobile Computing*, vol. 31, pp. 1–21, 2016.
- [18] A. Faruqui, S. Sergici, and C. Warner, "Arcturus 2.0: A meta-analysis of time-varying rates for electricity," *The Electricity Journal*, vol. 30, no. 10, pp. 64–72, 2017.
- [19] S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, "Incentivizing time-shifting of data: A survey of time-dependent pricing for internet access," *IEEE Communications Magazine*, vol. 50, no. 11, pp. 91–99, 2012.
- [20] M. J. Murrian, L. Narula, P. A. Iannucci, S. Budzien, B. W. O'Hanlon, S. P. Powell, and T. E. Humphreys, "GNSS interference monitoring from low earth orbit," *Navigation, Journal of the Institute of Navigation*, 2020, submitted for review.
- [21] T. Hobiger, D. Piester, and P. Baron, "A correction model of dispersive troposphere delays for the aces microwave link," *Radio Science*, vol. 48, no. 2, pp. 131–142, 2013.
- [22] Y. Teng and J. Wang, "A closed-form formula to calculate geometric dilution of precision (gdop) for multi-gnss constellations," *GPS solutions*, vol. 20, no. 3, pp. 331–339, 2016.
- [23] D. J. Welsh and M. B. Powell, "An upper bound for the chromatic number of a graph and its application to timetabling problems," *The Computer Journal*, vol. 10, no. 1, pp. 85–86, 1967.
- [24] N. Zufferey, P. Amstutz, and P. Giaccari, "Graph colouring approaches for a satellite range scheduling problem," *Journal of Scheduling*, vol. 11, no. 4, pp. 263–277, 2008.
- [25] 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.
- [26] E. Musk, (@elonmusk), Twitter post: "Around 20ms. It's designed to run real-time, competitive video games. Version 2, which is at lower altitude could be as low as 8ms latency." [Online]. Available: <https://twitter.com/elonmusk/status/1272363466288820224>
- [27] C. for International Earth Science Information Network, "Gridded Population of the World, Version 4 (GPWv4): National Identifier Grid, Revision 11. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC)," Columbia University, Tech. Rep., 2018.
- [28] J. Busby and J. Tanberk, "FCC reports broadband unavailable to 21.3 million americans, BroadbandNow study indicates 42 million do not have access," BroadbandNow Research, Tech. Rep., 2020, <https://broadbandnow.com/research/fcc-underestimates-unserved-by-50-percent>.
- [29] I. T. Union, "Measuring digital development: Facts and figures," International Telecommunications Union, Tech. Rep., 2019, <https://www.itu.int/en/ITU-D/Statistics/Documents/facts/FactsFigures2019.pdf>.