StarAngle: User Orientation Sensing with Beacon Phase Measurements of Multiple Starlink Satellites

Raghav Rathi
Dept. of Computer Science, Florida State University
rrathi@fsu.edu

Zhenghao Zhang
Dept. of Computer Science, Florida State University
zzhang@cs.fsu.edu

ABSTRACT

Low Earth Orbit (LEO) satellite networks have been growing very rapidly in recent years. In this paper, we propose a novel method, called StarAngle, which estimates user orientation with the beacon signals of Starlink satellites. StarAngle measures the beacon phase difference between two receiving antennas because the phase difference is a function of the user orientation. The phase measurements are compared with mathematical calculations based on known orbital parameters of Starlink satellites and the value that leads to the best agreement is used as the estimation. We overcome challenges due to asynchronous clocks in our commodity antennas by subtracting the phase measurements of one satellite by another which cancels the biases caused by clock mismatch. We experimentally test StarAngle in 10 locations under challenging weather conditions and our results show that the median estimation error is 7.5 degrees. Our results also confirm that the phase information of Starlink satellites can be measured reliably and may be used to support other applications in addition to user orientation estimation.

CCS CONCEPTS

• **Networks** → **Location based services**.

KEYWORDS

User orientation estimation, LEO satellite, Phase difference

ACM Reference Format:
Raghav Rathi and Zhenghao Zhang. 2024. StarAngle: User Orientation Sensing with Beacon Phase Measurements of Multiple Starlink Satellites. In *The 22nd ACM Conference on Embedded Networked Sensor Systems (SENSYS '24), November 4–7, 2024, Hangzhou, China.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3666025.3699367

1 INTRODUCTION

User orientation estimation is an important function in guidance systems and can be very helpful in many scenarios. For example, it can be very frustrating when we need to drive out of a parking lot but the GPS is confused about the orientation of the car. Many orientation estimation methods have been proposed for indoor and outdoor scenarios. The indoor methods exploit the signals from RFID tags [32], Bluetooth signals [13], light signals [2, 33], Reconfigurable Intelligent Surfaces (RIS) [7], or multiple sensors attached to the user body or embedded in smartphones [6, 24]. The outdoor methods include using compass, gyroscope, or analyzing Radio Frequency (RF) signals with algorithms including MUSIC [31], ESPRIT [30], and SAGE [8]. In this paper, we focus on outdoor orientation estimation to which the indoor methods usually cannot be extended. For example, many methods need to measure signals at selected locations which is possible indoors but not feasible for very large outdoor areas. The existing outdoor methods also have their own strengths and weaknesses. For example, compass is inexpensive but must be calibrated to cope with non-uniform Earth magnetic field and can be inaccurate with stray magnetic fields or ferromagnetic material in the vicinity. Gyroscope readings may drift with low-cost units although the drift can be reduced with space-grade units but such units are also more expensive. RF signals in the outdoors, such LoRa signals [4], offer another possibility for orientation estimation; however, without long-range RF transmitters installed over the entire globe, global coverage cannot be achieved.

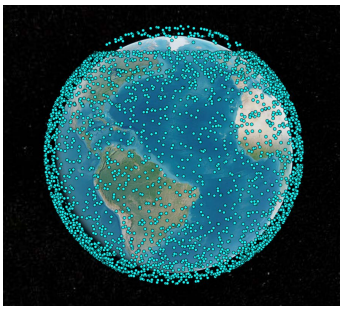

Figure 1: Starlink satellites.

In this paper, we propose *StarAngle*, which is an outdoor orientation estimation method based on the beacon signals of Starlink Low Earth Orbit (LEO) satellites received by commodity antennas. Compared to existing methods also based on RF signals, the key advantage of StarAngle is its free global coverage because Starlink satellites transmit beacon signals by default and cover the entire globe even in the ocean and very rural areas. By 2024, Starlink has launched over 6,000 satellites, the constellations of which are shown in Fig. 1. LEO satellite networks have seen very fast growing in the last few years and are expected to continue to grow. For example, Starlink plans to eventually maintain over 42,000 LEO satellites [28], after which there should be multiple satellites over any area with detectable signal at any time.

StarAngle finds the user orientation with two receiving antennas based on the phase difference of the same Starlink beacon received by the antennas because the phase difference is a function of the user orientation. StarAngle is novel because it is the first method

SENSYS '24, November 4–7, 2024, Hangzhou, China

ACM ISBN 979-8-4007-0697-4/24/11.
https://doi.org/10.1145/3666025.3699367

to extract the phase information of Starlink beacon signals using only commodity antennas. The main challenge is to cope with very weak signals under very heavy Doppler shift. To elaborate, due to the long distance from the satellites to Earth which is usually about 500 km, the Signal to Noise Ratio (SNR) of the signal picked up by our antennas is typically below -30 dB, i.e., the signal is more than 1000 times weaker than the noise. Also, as satellites travel at about 7,000 m/s, the *Doppler rate*, the amount frequency change of the RF signal, can be over 4,000 Hz/s. Our main contribution of this paper is to overcome these challenges by designing robust signal processing algorithms to detect the satellites and compensate for the Doppler rate. Another practical challenge is the asynchronous clock in the antennas, i.e., our commodity antennas are the Low-Noise Block downconverter (LNB) with individual builtin clocks not synchronized with each other. As a result, the clock difference generates large phase differences which completely masks the useful phase information. We overcome this challenge by exploiting the signals from multiple satellites received simultaneously. That is, suppose there are two satellites and let $\theta_1(t)$ and $\theta_2(t)$ be the phase difference measurements of satellite 1 and 2 at time t, respectively. We use $\theta_1(t) - \theta_2(t)$ to estimate the user orientation because the difference of the antenna clocks exists both in $\theta_1(t)$ and $\theta_2(t)$ and is canceled after the subtraction.

We conducted tests in 10 locations, including rooftops in universities, suburban residential areas, apartment complexes, and parks, where the sky condition can be clear or cloudy. The results show that the median estimation error is 7.5 degrees using only inexpensive commodity antennas and a Software Defined Radio (SDR). Our results show that the phase measurement, i.e., $\theta_1(t) - \theta_2(t)$, matches very well with the mathematical calculation according to published satellite orbital parameters, i.e., the measured data points usually land exactly on the calculated lines, an example of which has been shown in Fig. 13. The reason of this very high level of agreement is the line-of-sight path from the satellite to the receiver which is not contaminated by reflections and preserves the original geometry even from hundreds of kilometers away in the space.

The rest of the paper is organized as follows. Section 2 discusses related work. Section 3 gives a short description of the background. Section 4 gives an overview of StarAngle. Section 5 describes the detailed design of StarAngle. Section 6 describes the evaluation. Section 7 gives discussions and explains future directions. Section 8 concludes the paper.

2 RELATED WORK

Orientation estimation has attracted much interests both from the academia and industry. The outdoor estimation methods include 1) using compass, 2) using gyroscope, or 3) analyzing RF signals. As mentioned earlier, compass and gyroscope may suffer errors and biases in certain cases. StarAngle belongs to the third category but is unique because it leverages the existing Starlink infrastructure and provides global coverage without the need to install any base stations or gateways. StarAngle also solves unique problems because the satellite signal is very weak and suffer high Doppler rate, while existing signal processing algorithms and methods [4, 8, 26, 30, 31, 39] are not designed for such challenging cases. One of the key advantages of Starlink signals and satellite signals in general is that the signal is typically received from the Line-Of-Sight (LOS) path without strong multipath components, making it possible to infer information reliably from the phase measurement because it agrees very well with geometry. In addition, signals from different satellites are on distinct frequencies and can be easily separated in the frequency domain. As a result, the main challenge is to extract very weak signals and coping with high Doppler rate, rather than jointly estimating signals from difference sources or from multiple paths.

Many indoor orientation estimation methods have been proposed based on Received Signal Strength (RSS) measurements of RFID tags [32], light signals [2], Wi-Fi mmWave [15], etc., where the typical approach is to collect measurements at selected locations as fingerprints. Other approaches include using Reconfigurable Intelligent Surfaces (RIS) [7] or multiple sensors attached to the user body or embedded in smartphones [6, 24]. Bluetooth uses the the Angle of Arrival (AoA) or Angle of Departure (AoD) in multiple antenna systems for its Direction Finding service [13]. As mentioned earlier, the indoor methods typically do not extend to the outdoor scenarios. For example, the fingerprinting method is not feasible in outdoor large areas; the light and mmWave signals are designed for short range; the RIS is installed only at selected locations; the body-mounted sensors approach still needs a small number of control points with known coordinates; and the Bluetooth Direction Finding service are designed for Bluetooth communications and does not cope with very weak signals and high Doppler rate in outdoor satellite networks.

The Global Positioning System (GPS) uses signals from 4 or more satellites to calculate the user location and has been extremely successful and transforming. As the satellite orbital parameters are known, the possibility of using Starlink signals as a potential GPS alternative has been proposed [14, 17, 18, 25, 41] and generated media interest [10, 22]. The most recent result has achieved an accuracy of 4.3 m for user localization [18]. StarAngle is inspired by the earlier work but is unique because the information embedded in the phase difference between multiple antennas has not been measured and exploited. We show in this paper that such information can be reliably measured and used for purposes such as user orientation estimation.

There have also been extensive research on device localization which is related but different from orientation estimation. For example, fingerprinting methods use the signal strength measurement [3, 6, 47, 48] or the Channel State Information (CSI) [9, 29, 37, 40, 42] to find the device location without estimating characteristic of the signal propagation path. Other methods estimate the AoA [16, 23, 35, 49, 50], the phase difference [11, 27, 44] or Time Difference of Arrival (TDoA) [12, 34, 36, 38, 43] to locate the device. StarAngle is different because Starlink signal is dominated by the LOS path; as a result, StarAngle does not need to cope with multipath. As mentioned earlier, the main challenge of StarAngle is to cope with the very low SNR and high Doppler rate.

Satellite networks have attracted much interest recently. StarAngle focuses on extracting and exploiting Starlink beacon phase information and is different from studies on other problems such as understanding the downlink data signal structure [14], satellite constellation [19], downlink beamforming [20, 46], or secrecy [21].

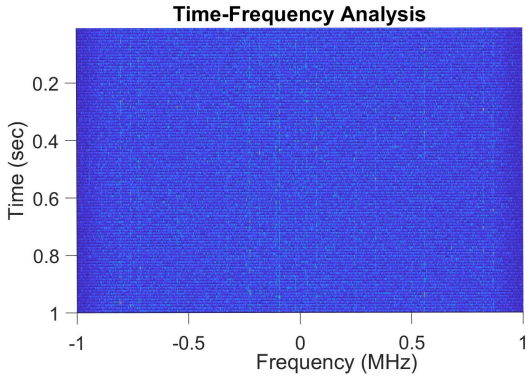

Figure 2: Beacons from Starlink satellites.

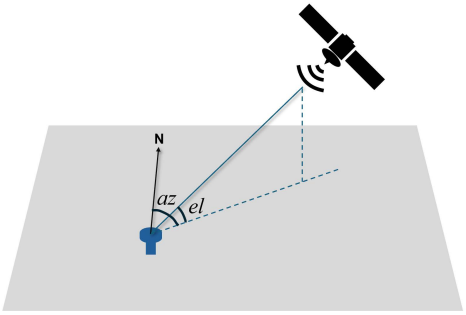

Figure 3: Azimuth and elevation angles.

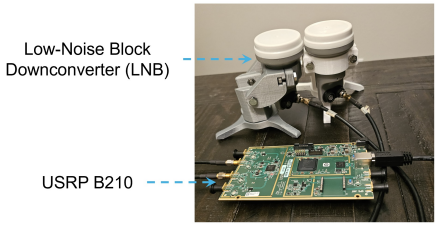

Figure 4: Devices used in our experiments.

3 BACKGROUND

Starlink satellites can be used for finding user location and orientation because they transmit beacons at known frequencies, such as 11.95 GHz, where each beacon appears to be a pure sine wave. Each satellite transmits multiple beacons separated by $\mathcal{F}$ Hz where $\mathcal{F}$ = 43955. It has been reported that the number of beacons is 9 [41]. However, satellites launched more recently, e.g., those with Starlink IDs over 30000, apparently transmit much more than 9 beacons, e.g., over 30 in our observations. Therefore, we do not assume the knowledge of the number of beacons by each satellite. Based on our observations and earlier work [41], Starlink satellites do not transmit signals on all beacons and do not transmit beacons at the same power. The purpose of transmitting such a large number of beacons and the beacon power allocation scheme are both unclear. We suspect some kind of coding scheme is adopted, which is out of the scope of this paper because we take measurements only with the strongest beacon.

Fig. 2 is the time-frequency plot of a typical 1-second signal where each horizontal line is the Fast Fourier Transform (FFT) of samples collected in 10 ms. As the beacons are pure sine waves, each strong beacon produces a high point in the FFT. The high points of the same beacon in consecutive FFTs are at close locations, which form a straight vertical line in the plot. There are multiple such lines in the plot, suggesting that there are many beacons in the signal. It can also be seen that the lines are actually not completely vertical, which is because the frequency of a beacon changes over time due to the relative movement between the satellite the ground observer. The *Doppler rate*, i.e., the amount of frequency change during one second, can be within -2000 to -5000 Hz/sec in our experiments.

The orbital parameters of all Starlink satellites are maintained in a database at [5]. Each satellite occupies three lines in the file, which are its ID, such as STARLINK-30787, followed by two lines known as the Two-Line-Elements (TLE). The TLE includes key parameters such as the inclination angle, the Right Ascension of the Ascending Node (RAAN), etc. It is possible to compute features of a satellite with respect to the ground receiver at a given time instant based only on the TLE. In our implementation, we use the MATLAB `dopplershift` function to compute the the Doppler rate and the `aer` function to compute the azimuth and elevation angles, which are the only three features needed in our computation. The azimuth and elevation angles are parameters to locate the satellite in the sky at a ground location. Roughly speaking, as shown in Fig. 3, they are how much the ground observer should turn from the true north and then look up in order to see the satellite, respectively. The TLE is not 100% accurate and must be updated every few hours. Therefore, before every experiment, we download the most recent TLE database and use it only for the particular experiment.

Our hardware is shown in Fig. 4, which consists of commodity LNBs and a USRP B210. The LNB is specifically designed to receive satellite signals, such as those from DirectTV networks. The LNB has an internal clock which can be configured at 9.75 GHz, and is therefore capable of down-converting the beacon signal at 11.95 GHz to 2.2 GHz to be received by the USRP. The USRP basically takes samples of the signal which are written to files to be processed. As we do not track any particular satellite, we point the LNBs up vertically. The sampling rate, denoted as Fs, is 2,000,000 measured in samples per second because beacons can be detected within ±1 MHz of the center frequency. The physical distance between the centers of the LNBs is denoted as $\mathcal{H}$ measured in meters where $\mathcal{H}$ = 0.062, which is the minimum we could achieve. Smaller distance is preferred because the wavelength of the beacon is only 0.025 m.

4 OVERVIEW OF STARANGLE

StarAngle finds the user orientation by solving three main problems, namely, 1) detecting satellites in the received signal, 2) measuring the phase differences between each pair of satellites, and 3) estimating the user orientation based on the phase measurements and satellite azimuth and elevation angles.

As shown in Fig. 5, satellites are detected by scanning for beacon signals which produce high peaks in the frequency domain because beacons are pure sine waves with energy concentrated within a very small frequency range. To use the satellites as reference points,

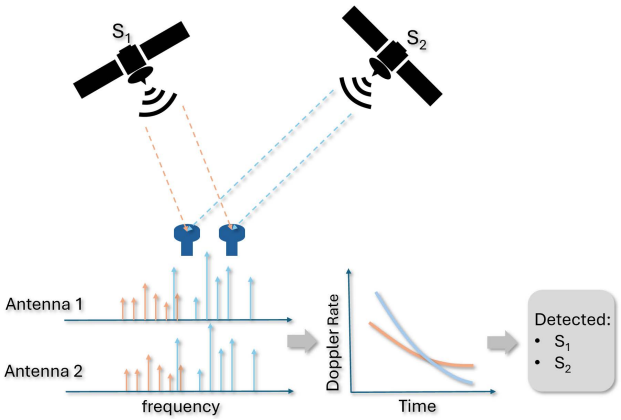

Figure 5: Satellite detection.

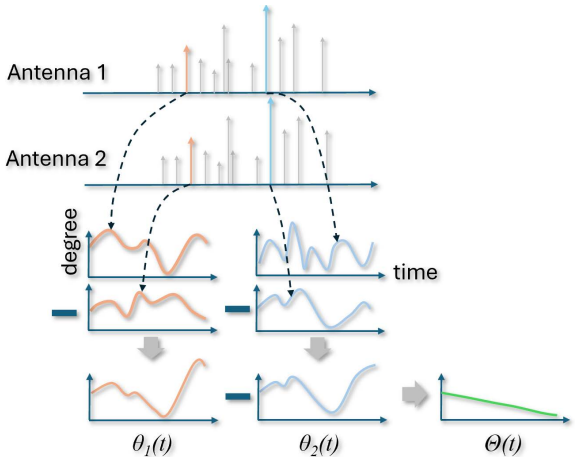

Figure 6: Taking phase measurements.

the IDs of the satellites must be determined, which however cannot be decoded from the beacon signal. Fortunately, satellites usually exhibit unique Doppler rates due to the uniqueness of their trajectories. Therefore, the measured Doppler rates of detected satellites can be compared with those calculated from the TLE database to match the detected satellites with the known satellites in the database, after which the satellite IDs can be determined. The main challenges at this step are the very low SNR such as -30 dB encountered in our experiments and the Doppler rate which can be -4000 Hz/sec. We design efficient algorithms to address these challenges by canceling the frequency shift and adding signals over a long period to accumulate more energy.

For each detected pair of satellites, phase measurements are obtained. As mentioned earlier, the measurement is $\theta_1(t) - \theta_2(t)$, where $\theta_i(t)$ is the phase difference of the beacon from satellite i between the two receiving antennas for $i = 1, 2$. The phase measurement process is shown in Fig. 6, where $\theta_1(t) - \theta_2(t)$ is defined as $\Theta(t)$. For each satellite, the strongest beacon is used. The time-domain signal of the selected beacon is reconstructed by applying a narrow-band filter that passes only the signal of this beacon but rejects beacons on other frequencies and most of the noise. $\theta_1(t)$ and $\theta_2(t)$ can then be obtained by subtracting the phase of the beacon received by one antenna by the other. $\Theta(t)$ can then be obtained by subtracting $\theta_1(t)$ by $\theta_2(t)$.

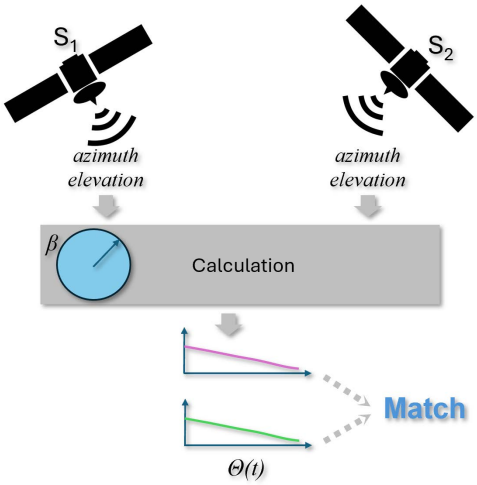

Figure 7: Estimating user orientation.

Table 1: List of Main Notations

β^*	user orientation
S_i	satellite i
$\theta_i(t)$	phase difference of S_i at time t
$\Theta(t)$	$\theta_1(t) - \theta_2(t)$
$\mathcal{F}$	beacon frequency separation
$\mathcal{H}$	antenna separation
Fs	sampling rate
α	Doppler rate
$B_k^i(t)$	beacon signal from S_i at antenna k
$b_k^i(t)$	phase of $B_k^i(t)$
$\delta_i(t)$	path lengths difference of S_i
$az_i(t)$	azimuth angle of S_i
$el_i(t)$	elevation angle of S_i

The actual user orientation, denoted as β^*, is defined as the azimuth angle of the line linking the center of the two LNBs. The estimation is based on finding the value that minimizes the difference between the phase measurements and phase calculation. To be more specific, after the satellites have been detected, their azimuth and elevation angles can be found according to their TLEs. Let β be a possible value of β^*. As shown in Fig. 7, the satellite azimuth and elevation angles along with β can be used as input to a phase calculation module to produce the calculated phase. When β is close to β^*, the calculated and measured phases should match.

5 DESIGN

In this section, the design of StarAngle is explained, including satellite detection, phase measurement, and user orientation estimation. Table 1 is a list of main notations.

5.1 Satellite Detection

Satellites are detected by their beacons, which is carried out in three steps. First, the received signal is processed in consecutive *segments* to find beacons from satellites in each segment, where the length of a segment is one second. This is because key satellite parameters, such as the Doppler rate, typically do not change significantly within one second so that they can be approximated as constants to simplify the detection algorithm. Second, clusters of beacons detected in consecutive segments are linked if they are believed to be from the same satellite. Third, the IDs of the satellites are found by comparing estimated Doppler rates with those calculated according to the TLE database.

5.1.1 Beacon Detection. Beacon detection is conceptually simple because a beacon is a pure sine wave on a certain frequency and should generate a detectable high point in the frequency domain. This is however complicated by two factors, namely, the *Doppler rate*, which can be -3000 Hz/sec, i.e., the frequency of the beacon changes over 3000 Hz in one second, and the very low SNR, which in our experiment can be less than -30 dB. Note that when the Doppler rate is high, the beacon frequency does not stay the same and does not generate a single high point in the frequency domain. For example, when the Doppler rate is -3000 Hz/sec, energy is spread among 3000 points in the FFT of 1 second. Therefore, the FFT should be performed over a shorter time period. However, when the SNR is very low, the high point generated by the beacon may not be high enough to be detected.

Our solutions to these problems are explained in the following. First, we divide a segment into *mini-segments* where each mini-segment is 10 ms. The initial time-frequency analysis, which is shown in Fig. 2, is the FFT of each mini-segment. When the Doppler rates of Starlink satellites are within ±5000 Hz/sec, which we found empirically to be true, a beacon generates a single high point in the FFT because within a mini-segment, the frequency change is within ±50 Hz while one FFT point corresponds to 100 Hz. Clearly, as explained earlier, when a mini-segment is as short as 10 ms, a beacon may not generate a detectable high point. That is, although the beacon usually generates a point of a descent height, points of similar height can also be generated by noise, making it impossible to set a good threshold for beacon detection. Our idea is to add the FFTs of all mini-segments together because a real beacon should consistently generate reasonable high points while high points due to noise should be random. The challenge, however, is that the locations of high points generated by the same beacon may change in different mini-segments due to the frequency change. Fortunately, the Doppler rate does not change significantly within 1 second; as a result, the frequency change can be assumed to be linear, which has been shown in Fig. 2 where the beacon frequency typically follows a straight line with a certain slope. Therefore, as shown in Fig. 8, if we shift the FFTs of the mini-segments with the negative of the slope, the effect of frequency change can be canceled, so that the high points generated by the same beacon in different mini-segments are aligned and can be added together to create a much higher point which can be detected reliably by algorithms such as `peakfinder` [45]. Fig. 9 shows an actual case of taking summations. The top figure is the summation without any shift which has a single high peak in the middle due not to any actual beacon but to interference. The bottom figure is the summation with the correct shift in which many beacons start to emerge.

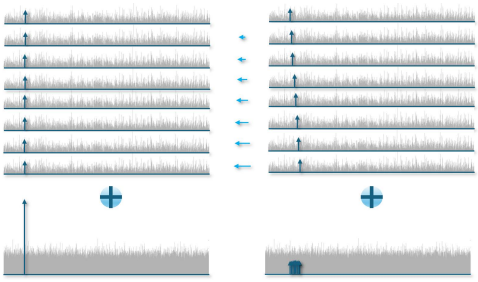

Figure 8: Taking summation after shifting the vectors.

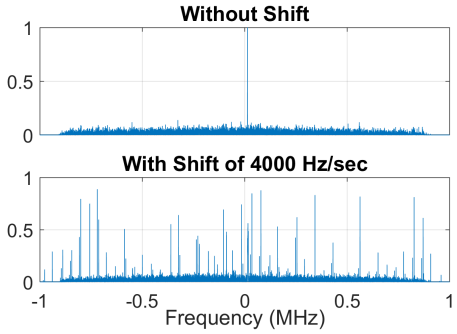

Figure 9: Example of summation vectors with and without the correct shift.

The beacon detection algorithm is shown in Algorithm 1. In the algorithm, "⌈⌋" denotes rounding a number to the nearest integer. The algorithm mainly consists of a linear scan of possible Doppler rate values at a step of 1000 Hz/sec, which corresponds to 0.1 point of shift between consecutive mini-segments. The summation vector for shift d is denoted as Γ_d and is passed to `peakfinder`. The last step of the algorithm is to merge common peaks which refers to peaks found in different summation vectors at close locations because such peaks are clearly generated by the same beacon that happen to be capable of generating peaks with different shift values.

Algorithm 1 Beacon Detection

1: Partition the signal into 10-ms mini-segments. Let mini-segment u be g_u where $u \in [1, U]$.
2: Let $G_u = FFT(g_u)$ for $u \in [1, U]$.
3: $W \leftarrow [-0.5, -0.4, ..., 0.5]$.
4: **for** $d = 1$ to $|W|$ **do**
5: $\quad G_u^d \leftarrow G_u$ shifted by $\lceil W(d)u \rfloor$ for $u \in [1, U]$.
6: $\quad \Gamma^d \leftarrow \sum_u^U G_u^d$
7: $\quad P^d \leftarrow$ the set of peaks found in Γ^d.
8: **end for**
9: Merge common peaks in $P^1, P^2, \ldots, P^{|W|}$.

Lastly, beacons separated by values close to multiples of $\mathcal{F}$ Hz are grouped together and are called a *cluster* because they are likely

from the same satellite. The Doppler rate of the cluster, denoted as α, is estimated based on the strongest beacon in the cluster. That is, first, an initial estimate of α is obtained by tracking the highest points generated by this beacon in G_u for $u \in [1, U]$, because the slope of the line that best fits the locations of these points is close to α. With this approach, the initial estimate is usually within ±100 Hz/sec of α. Second, the fine estimate of α within ±1 Hz/sec is obtained by applying different offset to the initial estimate and picking the one that produces the highest peak in the frequency domain. To simplify the search, first, offsets within ±100 Hz/sec are tested at a step of 10 Hz/sec; then, if the highest peak occurs with x Hz/sec, offsets within ±5+x Hz/sec are tested at a step of 1 Hz/sec. This method is used because it performs well even for very weak signals.

5.1.2 Linking Clusters. The cluster linking module links clusters of beacons found in consecutive segments so that signals from the same satellite can be traced throughout the experiment. The problem is not completely trivial because the signal may contain clusters from multiple satellites as well as those falsely detected by the beacon detection module.

Cluster linking is based on two features of beacons. First, from one segment to the next, the Doppler rate of a beacon should be similar in some manner. Second, from one segment to the next, the beacon frequency should have changed by an amount close to the estimated Doppler rate. One factor that slightly complicates the matter is that the Doppler rate can no longer be assumed to be a constant for the duration of the experiment, which is 40 sec, although it is close to a constant during a segment, which is 1 sec. Fortunately, the Doppler rate in most cases follows a smooth curve. Therefore, for each satellite, we can calculate the expected Doppler rate of the next segment based on its past Doppler rates by fitting the Doppler rate history with a polynomial.

To be exact, let cluster r found in the q^{th} segment be C_r^q. Let the Doppler rate of C_r^q be α_r^q measured in Hz/sec. Let the frequency of the strongest beacon of C_r^q be f_r^q Hz, which is within $[-Fs/2, Fs/2]$ Hz. Let a detected satellite, denoted as S_i, be a list of clusters attributed to the satellite so far: $\{C_{r_1}^{q_1}, C_{r_2}^{q_2}, \ldots, C_{r_l}^{q_l}\}$. Suppose we need to determine if $C_{r'}^{q'}$ should be linked to this list where $q' > q_l$. We do not consider $C_{r'}^{q'}$ if $q' - q_l > 5$ because the gap is too large and the estimation becomes inaccurate. This usually not does not cause problems because the gap exists only when the beacons are very weak and cannot be detected in multiple consecutive segments. As mentioned earlier, we estimate the change of Doppler rate by fitting the observed Doppler rate history with a polynomial. Let the polynomial evaluated for segment q' be $F_i(q')$. $C_{r'}^{q'}$ is linked if the following two conditions are both true:

$$|F_i(q') - \alpha_{r'}^{q'}| < 100, \tag{1}$$

and

$$X < 1000 \text{ or } X > \mathcal{F} - 1000, \tag{2}$$

where

$$X = \mod\left(f_{r_l}^{q_l} + (q' - q_l)\alpha_{r_l}^{q_l} - f_{r'}^{q'}, \mathcal{F}\right). \tag{3}$$

Eq. 1 checks if $\alpha_{r'}^{q'}$ is close to the expected Doppler rate of S_i in segment q'. Eq. 2 checks if the difference between $f_{r'}^{q'}$ and the expected frequency of S_i in segment q' is close to a multiple of $\mathcal{F}$ Hz. Both conditions are needed for better robustness because, for example, beacons from different satellites may have very similar Doppler rates at a certain time and can be distinguished only based on frequency values.

Lastly, let V_i be the *history vector* that can be used for curve fitting to obtain $F_i(q')$. In our implementation, V_i contains all the Doppler rate measurements of S_i if $l \leq 6$, otherwise the most recent 6 measurements. Let $P_{V_i}^a()$ be a polynomial of degree a that minimizes the fitting error to V_i. $F_i(q')$ is calculated according to $F_i(q') = P_{V_i}^a(q')$ where $a = \min\{2, l-1\}$. The maximum degree is 2 because the Doppler rate curve within the duration of an experiment can almost always be fitted well with degree 2.

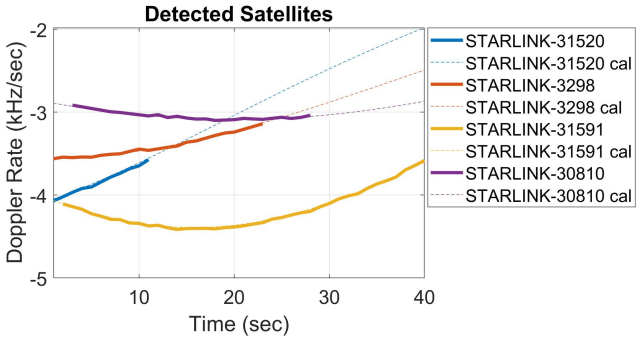

Figure 10: Example of 4 detected satellites.

5.1.3 Satellite Matching. Matching satellites refers to associating the detected satellites with those calculated from the TLE database, which is a necessary step because the IDs of the satellite cannot be decoded from the beacons. This is possible because the number of *candidate satellites*, i.e., satellites flying over the user location, is not very large and is usually no more than 30, while each candidate satellite has a unique trajectory, leading to unique Doppler rates which can be used for identification. Therefore, we compare the measured Doppler rates of each detected satellite with those of the candidate satellites, and associate a detected satellite with the candidate satellite with the minimum Euclidean distance. For example, Fig. 10 shows 4 detected satellites in one experiment. The solid and dashed lines represent the estimated and calculated Doppler rates of the same satellite, respectively, which match very well. The solid lines exist only in part of the experiment which is when the beacon signals can be detected.

The set of candidate satellites are obtained with a simple method that examines all Starlink satellites and finds those can be accessed theoretically according to the TLE database. A satellite is considered a candidate if its elevation angle is 30° or higher in the middle of the experiment.

5.2 Phase Measurement

The phase measurement module examines every pair of detected satellites to extract the phase information. In the following, it is described for two satellites denoted as S_1 and S_2 because the same process is applied to every pair of satellites.

As mentioned earlier in Section 4, the phase measurement at time t is $\Theta(t) = \theta_1(t) - \theta_2(t)$. The measurement process is performed for each segment of the received signal individually for the same reason the satellite detection process is performed for each segment. Therefore, in the following, it is described for one segment. The basic idea is to reconstruct the time-domain signals of the beacon from which the phase difference can be directly measured. That is, let B_k^i be the baseband signal of the beacon from S_i received by antenna k during this segment. Let $B_k^i(t)$ be the beacon signal at time t and let $b_k^i(t)$ be the phase of $B_k^i(t)$. By definition,

$$\theta_i(t) = b_1^i(t) - b_2^i(t). \tag{4}$$

It is possible to produce B_k^i even when the beacon is very weak and when there exist many beacons because the beacon is on a constant frequency and can be extracted with a narrow-band filter that keeps only the signal of the beacon but rejects most of the noise and beacons on other frequencies.

To be more specific, let Ω_k be the vector of the received baseband samples during this segment by antenna k and let $\omega_k(t)$ be the phase of the sample received at time t. As mentioned earlier, the strongest beacon from S_i is used for the measurement, the frequency of which is denoted as f_k^i for antenna k. Note that $f_1^i \neq f_2^i$ because the antennas have own builtin clocks with non-zero frequency offset, leading to different measured frequency values of the same beacon. The first step of the signal reconstruction step is to apply the negative of the Doppler rate of the beacon, denoted as α_i for S_i measured in Hz/sec, which has been estimated earlier during satellite detection. This operation is also called *dechirping* which cancels the change of frequency and returns the beacon to a sine wave on a constant frequency. Mathematically speaking, dechirping is to adjust $\omega_k^i(t)$ according to

$$\omega_k^i(t) \leftarrow \omega_k^i(t) - \frac{\alpha_i t^2}{2} \tag{5}$$

because when the Doppler rate is α_i Hz/sec, at time t, the accumulated phase change is

$$\int_{\tau=0}^{t} \alpha_i \tau \, d\tau = \frac{\alpha_i t^2}{2}. \tag{6}$$

Note that dechirping is different and more effective than adding shifted FFTs of the mini-segments described earlier in Algorithm 1 because all samples are added coherently. Dechirping cannot be applied in Algorithm 1 because the beacons have not been detected and the Doppler rates are not known.

Let Ω'_{k,α_i} represent the baseband signal vector after being dechirped by α_i Hz/sec. Let $\Gamma'_{k,\alpha_i} = FFT(\Omega'_{k,\alpha_i})$ and rearrange it so that the indices of the elements are in $[-Fs/2, Fs/2-1]$. As the length of the segment is one second, the beacon should generate a peak at f_k^i in Γ'_{k,α_i}. In practice, some energy may have spread to neighboring points but the energy is still mostly concentrated around location f_k^i. As $f_1^i \neq f_2^i$, it is necessary to convert the beacons received by both antennas to a common frequency because $\theta_i(t)$ will otherwise be dominated by the phase change due to the frequency difference. Therefore, Γ'_{k,α_i} is shifted by f_k^i points, denoted as $\Gamma'_{k,\alpha_i,f_k^i}$, which is equivalent to down-converting the beacon to 0 Hz.Then, a low-pass filter with 100 Hz bandwidth is applied to obtain B_k^i, which is currently implemented by taking the 100 points around 0 Hz in $\Gamma'_{k,\alpha_i,f_k^i}$ and performing an Inverse FFT.

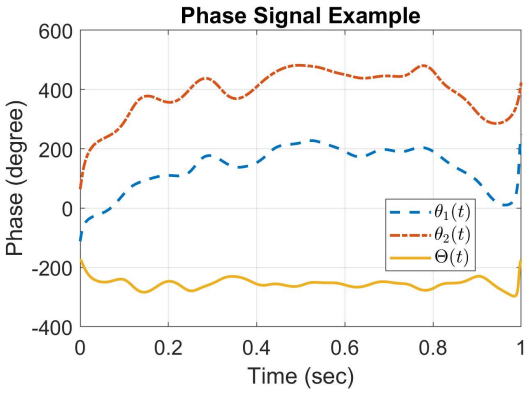

Figure 11: Phase measurement for STARLINK-31250 and STARLINK-3298.

Fig. 11 shows the phase measurement of one segment for the satellite pair STARLINK-31250 and STARLINK-3298 detected earlier in Fig. 10. It can be seen that $\theta_1(t)$ and $\theta_2(t)$ look similar and both have some large drifting and fluctuations caused by the clock mismatch. $\Theta(t)$, however, is mostly a constant with no large fluctuation, which is the expected behavior of the phase because the satellite angles with respect to a ground observer do not change significantly during one second. There do exist some small ripples in $\Theta(t)$ caused likely by overfiting, which do not cause problems because we eventually use the average of $\Theta(t)$ over one segment as one data point and the ripples are canceled during the averaging.

5.3 User Orientation Estimation

β^* is estimated by finding the minimum of a cost function that is determined by the difference between the measured phase values and those calculated according to the TLE database.

5.3.1 Mathematical Foundation. StarAngle is based on the fact that $\Theta(t)$ contains only the useful phase information and is not contaminated by the clock differences of the satellites and the antennas. To see this, let $\phi_k(t)$ be the carrier phase of antenna k at time t. Let $\psi_k^i(t)$ be the phase of the beacon from S_i received by antenna k at time t. Let λ be the wavelength of the beacon. Let $\delta_i(t)$ be the difference of path lengths from S_i to antenna 1 and antenna 2 at time t. Note that

$$\psi_2^i(t) = \psi_1^i(t) + \frac{2\pi\delta_i(t)}{\lambda}. \tag{7}$$

Therefore,

$$b_1^i(t) = \psi_1^i(t) - \phi_1(t) \tag{8}$$

while

$$b_2^i(t) = \psi_1^i(t) + \frac{2\pi\delta_i(t)}{\lambda} - \phi_2(t) \tag{9}$$

which leads to

$$\theta_i(t) = b_1^i(t) - b_2^i(t) = \frac{2\pi\delta_i(t)}{\lambda} - \big(\phi_1(t) - \phi_2(t)\big), \tag{10}$$

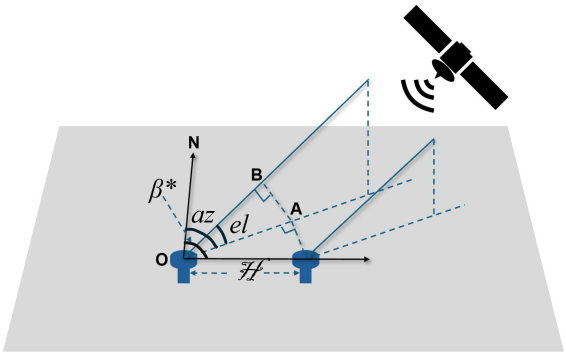

Figure 12: Phase calculation with satellite angles and β^*.

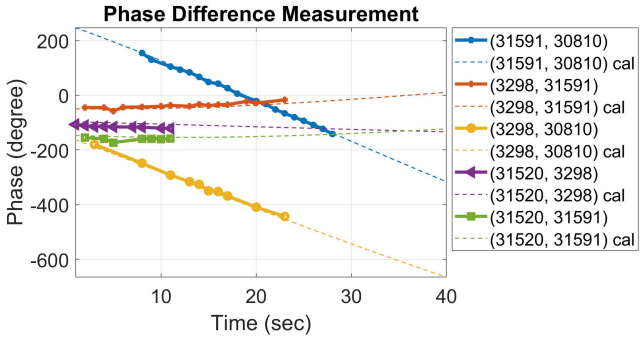

Figure 13: Example of phase measurements.

because $\psi_1^i(t)$ is canceled after the subtraction. Then,

$$\Theta(t) = \theta_1(t) - \theta_2(t) = \frac{2\pi[\delta_1(t) - \delta_2(t)]}{\lambda}, \tag{11}$$

because $\phi_1(t) - \phi_2(t)$ is canceled after the subtraction. $\Theta(t)$ is therefore irrelevant to the clock characteristics of the satellites and antennas and is only a function of $\delta_1(t)$ and $\delta_2(t)$.

$\delta_i(t)$ is determined by the location of S_i and the user orientation. To elaborate, let the azimuth and elevation angles of S_i observed at the user location at time t be $az_i(t)$ and $el_i(t)$, respectively. As mentioned earlier, $az_i(t)$ and $el_i(t)$ can be calculated based on the TLE database of the detected satellites, such as by the `aer` function in MATLAB. The situation is illustrated in Fig. 12 where $az_i(t)$ and $el_i(t)$ are shown as az and el for simplicity. In the figure, the length of OA, denoted as $|OA|$, is given by

$$|OA| = \mathcal{H}\cos(\beta^* - az), \tag{12}$$

where $\mathcal{H}$ is the antenna separation. The length of OB, which is the additional distance the signal needs to travel to reach the antenna on the left, is given by

$$|OB| = |OA|\cos(el). \tag{13}$$

Therefore,

$$\delta_i(t) = \mathcal{H}\cos\left(\beta^* - az_i(t)\right)\cos\left(el_i(t)\right). \tag{14}$$

5.3.2 The Estimation Method. The best estimate of β^* should be the value that leads to the best agreement between the calculated and measured versions of $\Theta(t)$. Fig. 13 shows the measured and calculated $\Theta(t)$ based on the best estimate of β^* for 5 pairs of satellites detected earlier in Fig. 10 for the entire duration of the experiment, where the solid and dashed lines represent the measurements and calculation, respectively. It can be seen that the measurements and calculation match very well.

We explain the details of the estimation method in the following, starting with some notations and definitions. Let $\hat{\Delta}(q)$ be the average of $\Theta(t)$ during the q^{th} segment and refer to it as a *phase measurement*. In Fig. 10, each marker is a valid phase measurement. $\Theta(t)$ is averaged over one segment because it does not change significantly during a short period such as one second. An additional step is taken to eliminate likely bad phase measurements, i.e., $\hat{\Delta}(q)$ is discarded if the standard deviation of $\Theta(t)$ during the q^{th} segment is too high. Let $Az_i(q)$ and $El_i(q)$ be the average azimuth and elevation angles during the q^{th} segment for S_i, respectively. Note that $Az_i(q)$ and $El_i(q)$ are known from the TLE database and are treated as constants. Let $\Delta(q;\beta)$ be the theoretical value of $\Theta(t)$ calculated according to $Az_i(q)$, $El_i(q)$, and β, where β is a possible value of β^*. Let $A_i(q)$ be the signal power of S_i during the q^{th} segment which should also be taken into consideration because measurements with higher signal power should have higher credibility. Therefore, $A_1(q)A_2(q)$ is called the *weight* of $\hat{\Delta}(q)$.

Without loss of generality, suppose there are N valid phase measurements with indices from 1 to N. Let the cost function be $\Xi(\beta)$ which reflects the likelihood for β to be β^*, the higher the cost the less likely. Based on the above discussions, $\Xi(\beta)$ can be defined as

$$\Xi(\beta) = \sum_{q=1}^{N}\left(\hat{\Delta}(q) - \Delta(q;\beta)\right)^2 A_1(q)A_2(q). \tag{15}$$

The estimation of β^*, denoted as $\hat{\beta}$, is clearly

$$\hat{\beta} = \arg\min_{\beta}\Xi(\beta). \tag{16}$$

In practice, there may be gaps in phase measurements. Eq. 15 can be easily adapted to this case by considering only segments with valid measurements.

5.3.3 Fake Minimums. A value β' is called a *fake minimum* if β' is far from $\hat{\beta}$ but $\Xi(\beta')$ is close to $\Xi(\hat{\beta})$. The condition for fake minimums to occur is discussed in the following. Usually, $\Delta(q;\hat{\beta})$ is close to $\hat{\Delta}(q)$ for every q. Therefore, in order for $\Xi(\beta')$ to be close to $\Xi(\hat{\beta})$, $\Delta(q;\beta')$ and $\Delta(q;\hat{\beta})$ must be close for almost all q. $\Delta(q;\beta') - \Delta(q;\hat{\beta})$ is the product of $\frac{2\pi\mathcal{H}}{\lambda}$ and the following term:

$$\left(\cos\left(\beta' - Az_1(q)\right) - \cos\left(\hat{\beta} - Az_1(q)\right)\right)\cos\left(El_1(q)\right) - \left(\cos\left(\beta' - Az_2(q)\right) - \cos\left(\hat{\beta} - Az_2(q)\right)\right)\cos\left(El_2(q)\right). \tag{17}$$

Therefore, $\Delta(q;\beta')$ is close to $\Delta(q;\hat{\beta})$ implies that

$$\frac{\cos\left(\beta' - Az_1(q)\right) - \cos\left(\hat{\beta} - Az_1(q)\right)}{\cos\left(\beta' - Az_2(q)\right) - \cos\left(\hat{\beta} - Az_2(q)\right)} \approx \frac{\cos\left(El_2(q)\right)}{\cos\left(El_1(q)\right)} \tag{18}$$

for every q, which is unlikely in most cases.

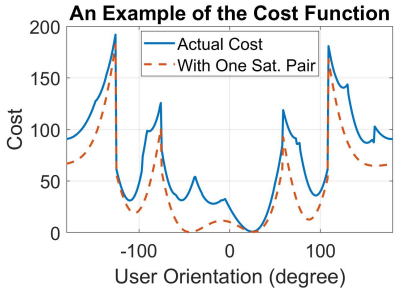

Figure 14: Example of the cost function.

5.3.4 Multiple Satellite Pairs. There could be more than one pair of satellites during the experiment, which can be of great help by eliminating fake minimums because the values of fake minimums are different for different satellite pairs. One straightforward extension is to calculate a cost function for each individual pair of satellites according to Eq. 15 then use the summation as the overall cost. One issue with this straightforward extension is information reuse because some measurements are not independent. For example, suppose there are 3 satellites, namely, S_1, S_2, and S_3. The phase measurements of the q^{th} segment can be denoted as $\hat{\Delta}_{1,2}(q)$, $\hat{\Delta}_{1,3}(q)$, and $\hat{\Delta}_{2,3}(q)$, respectively, among which only two are independent because the third is a linear combination of the other two. We therefore use a simple method to select the measurements of each segment. That is, let the cost function of the q^{th} segment be $\Xi_q(\beta)$, which is initially 0. Then, in each round, the measurement with the highest weight is selected. If the selected measurement is independent of the measurements that have been considered, its cost is added to $\Xi_q(\beta)$; otherwise, it is skipped. The process is repeated until no independent measurement can be added.

Fig. 14 shows an example of the cost function with the phase measurements of 5 pairs of satellites shown earlier in Fig. 13. It can be seen that the minimum cost is very low, implying a very good matching between the measured phases and those calculated according to $\hat{\beta}$. Also shown in Fig. 14 is the cost function by using only the first satellite pair, which clearly has two low points, one overlapping with the low point of the actual cost function, the other at a completely different location and is a potential fake minimum. This example therefore illustrates the benefit of having more satellite pairs.

6 EVALUATION

StarAngle has been experimentally evaluated, which is explained in this section.

6.1 Experiment Setup

We test StarAngle in 10 locations, including rooftops of university buildings, lawn between university buildings, suburban residential areas, space between apartment buildings, and public parks. Experiments 1, 7, 8, 9, 10 and experiments 2, 3, 4, 5, 6 are in two different cities separated by over 100 miles (161 km). The surroundings of experiment locations are shown in Fig. 15. It can be seen that our receiver is at the ground level in most cases and is often close to some building. There are cloud covers in some of the experiments and one experiment is under a tree. In many experiments, the weather conditions are very harsh because the devices are directly under the sun. The intense heat sometimes even melts the 3D-printed stand of the LNBs. We use the information at [1] to estimate if there are multiple satellites flying over our area to start a signal capture. In some experiments, we also use a third antenna and connect it to a separate SDR to display the signal in real time to helps us visualize the signal. The signal from the third antenna is not processed by StarAngle and in some cases we unplug it to reduce the load of our laptop computer and avoid overheating. Table 2 lists the detailed information about the experiments.

Table 2: Experiment information

ID	Location	Sky	Temperature	Humidity
1	backyard	part. cloudy	81°F (27°C)	82%
2	rooftop	cloudy	85°F (29°C)	57%
3	rooftop	clear	93°F (34°C)	58%
4	univ. lawn	part. cloudy	84°F (29°C)	74%
5	front yard	cloudy	85°F (29°C)	77%
6	apartment	cloudy	86°F (29°C)	72%
7	apartment	part. cloudy	83°F (28°C)	88%
8	park	part. cloudy	85°F (29°C)	76%
9	backyard	cloudy	78°F (25°C)	96%
10	under tree	part. cloudy	86°F (19°C)	88%

In all except 3 cases, at each location, 2 angles are randomly chosen, according to which the LNBs are orientated. For each angle, multiple *runs* of signal collection are performed, where each run is always 40 sec. For each angle, the results of the first 3 good runs are reported in this paper. A run is good if 1) two satellites appear simultaneously in the signal for at least 10 seconds, and 2) good match can be found between the detected satellites and those calculated from the TLE database. The second condition is sometimes not true likely due to outdated information in the TLE database. The 3 exceptions are explained in the following. First, Experiments 8 and 10 have data for only one angle due to weather constraints. Second, both angles of Experiment 2 have only 2 good runs because the results appeared to be good during the experiments but later turned out to be not good.

6.2 Results

Fig. 16 is the scatter plot of the estimation results from all experiments, where a marker is at location (x, y) if the actual user orientation is x degrees and the estimated is y degrees. It can be seen that the estimation results are mostly accurate with markers very close to the diagonal line. There are 2 cases in which the estimations are very far from the actual values, which are called *failures*. The failures are caused by fake minimums, i.e., the cost function happens to be slightly lower at another point than at the actual user orientation. As mentioned earlier, fake minimum can be eliminated with more satellite pairs. Indeed, in both failure cases, there is only one satellite pair. Fig. 17 is the Cumulative Density Function (CDF) of the estimation error, the median of which is 7.5 degrees. The failure cases, which constitute a very small percentage, are not shown to better exhibit the details of other cases.

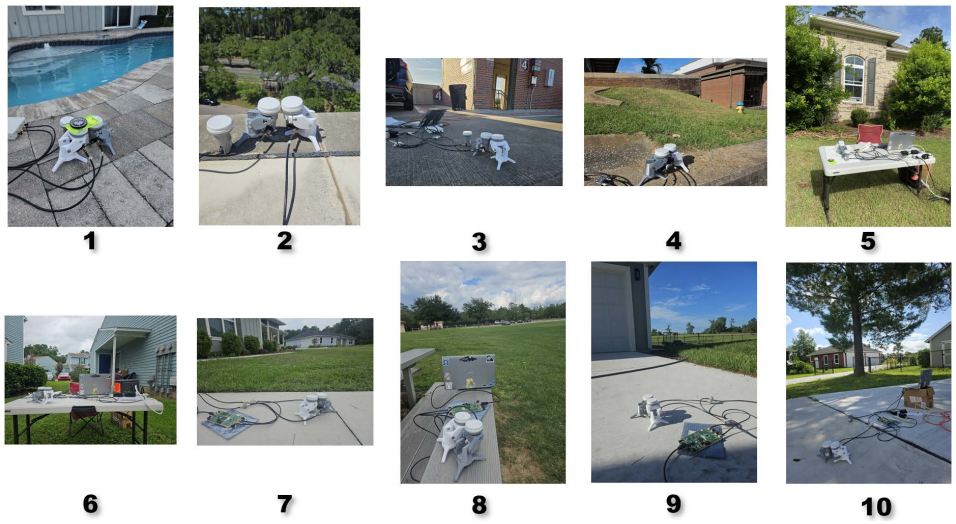

Figure 15: Surroundings of the experiment locations.

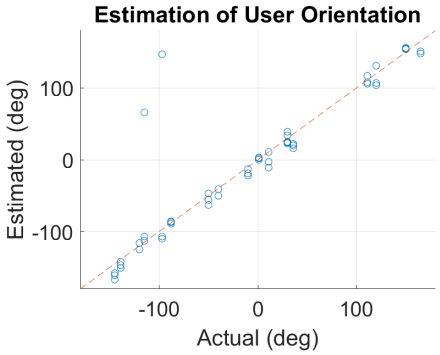

Figure 16: Scatter plot of estimation results.

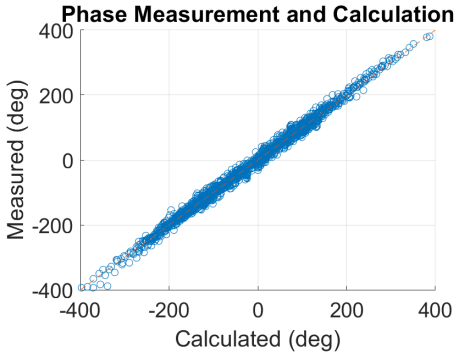

Figure 18: Scatter plot of phase measurement and calculation.

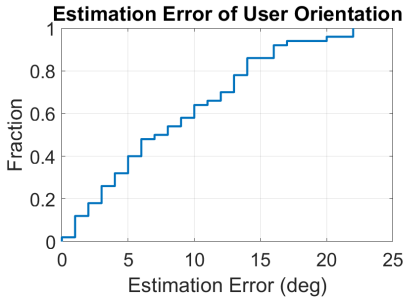

Figure 17: CDF of estimation errors.

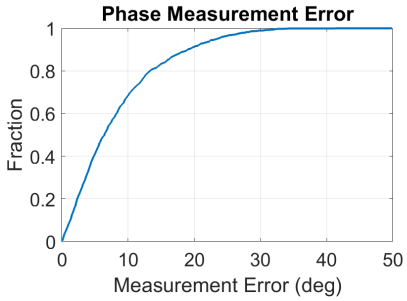

Figure 19: CDF of phase measurement errors.

Fig. 18 is the scatter plot of the phase measurements and the calculation based on the best estimate of the user orientation in all experiments. It can be seen that the measurements agree very well with the calculation, which has been previously hinted in Fig. 13 for a single run. Fig. 19 is the CDF of the difference between phase measurements and calculations, the median of which is 6.4 degrees. These results therefore suggest that the signal phase of Starlink satellites can likely be a useful source of information because it can be mathematically understood.

Fig. 20 reveals more details about the experiments. Fig. 20(a) is the CDF of the estimated SNR of each run, which is very low. StarAngle still produces reliable phase measurements, suggesting that it is capable of coping with very weak signals. Fig. 20(b) is the CDF of the number of satellite pairs, where it can be seen that

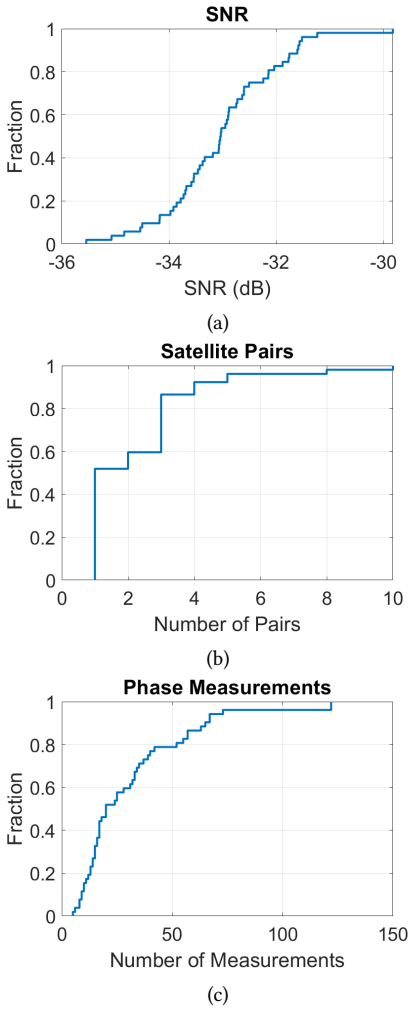

Figure 20: Various aspects of the experiment.

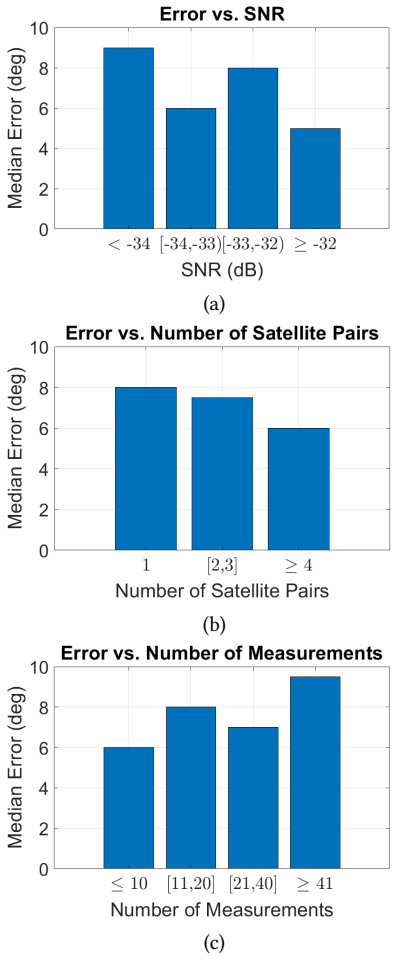

Figure 21: User orientation estimation error as functions of various aspects of the experiment.

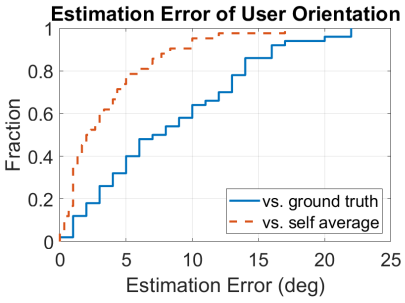

Figure 22: Consistency of estimation.

over half of the runs have only one useful pair, although there are some runs with many pairs. Fig. 20(c) is the CDF of the number of phase measurements, where each measurement is a basically a marker shown in Fig. 13. It can be seen that there are cases with only about 10 useful data points. We expect that, as the number of Starlink satellites keeps increasing, more satellite pairs will become available in the future which will produce more data points.

The estimation error is further studied in Fig. 21 with respect to various aspects of the experiments. Fig. 21(a) shows the medians of estimation errors in different ranges of SNR. The trend, as expected, is that the error is lower when the SNR is higher. This suggests that StarAngle should perform better with better hardware or under more benign weather conditions. Fig. 21(b) shows the medians of estimation errors with different number of satellite pairs. Also as expected, the error is lower with more satellite pairs. Fig. 21(c) shows the medians of estimation errors with different number of phase measurements, the trend in which is not as expected because the error appears to be larger with more phase measurements. The reason is due to large measurement errors in some cases as well as some bias which will be discussed shortly.

We found that the estimations produced by StarAngle, although sometimes may deviate from the ground truth, are usually consistent with each other for the same location and same angle. Fig. 22 compares the estimation errors when defined against the ground truth, which is the same as that shown in Fig. 17, and those defined

against the *self average*, i.e., the error of a particular run is the difference between the estimation of this run and the average of all runs for the same angle at the same location. It can be seen that the estimation error against self average is much lower, suggesting that StarAngle gives consistent estimations.

6.3 Summary

As a summary, our experiments show that: 1) StarAngle achieves good estimation accuracy of user orientation with median error of 7.5 degrees; 2) the phase measurements agree well with the mathematical calculation; and 3) the estimation accuracy can be improved with higher SNR, more satellite pairs, and better hardware setup.

7 DISCUSSIONS AND FUTURE DIRECTIONS

In this section, we discuss several issues as well as future directions.

7.1 Further Improving the Performance

The SNR in our experiment is very low because our current LNBs are less than $20 and do not collect very strong signals. If better antennas can be used, the performance of StarAngle can be improved, although the signal is still likely very weak unless the antennas have large reception areas commodity satellite dishes. A more sturdy and better leveled stand will also likely reduce measurement biases caused by tilting antennas in some cases. StarAngle can also benefit from more satellites because estimation failures are eliminated as long as there are more than one pair of satellites in our experiments. We could have insisted running the experiment until at least two pairs of satellites are observed but it would have needed more time and is not feasible under the weather conditions.

7.2 Measurement Opportunities

We refer to a *measurement opportunity* as an occasion in which StarAngle can obtain phase measurements. Each Starlink satellite revolves around Earth about 15 times a day and creates measurement opportunities at various locations as it travels. More satellites in the sky also creates more measurement opportunities. Currently, Starlink has over 6,000 satellites. It has been reported that Starlink eventually will maintain 42,000 satellites [28] which will create much more measurement opportunities.

7.3 StarAngle with Synchronized Antenna Clocks

Our prototype is built with commodity LNBs with asynchronous clocks, which forces us to subtract the signal of one satellite by another in order to cancel the clock differences. If two or more LNBs are driven by a common clock, the phase information can be measured with the signal from one satellite by subtracting the signal from one antenna by another. This will create more measurement opportunities because there is no longer the need to wait until more than one satellite is observed. In other words, StarAngle will produce estimates faster and likely more accurately. Most of the signal processing algorithms in StarAngle, such as the satellite detection algorithm in Section 5.1 and the phase measurement algorithm in Section 5.2, need not be changed because they are designed to detect and measure the information of individual satellites. Some minor change needs to be made for the user orientation estimation algorithm in Section 5.3 mainly by modifying the cost function in Eq. 15 for a single satellite.

7.4 Beyond Starlink

StarAngle currently processes beacon signals of Starlink satellites but is not restricted to Starlink because it only needs to receive a pure sinewave transmitted by a satellite. In other words, StarAngle is applicable to any satellite as long as the satellite transmits a beacon signal in the form of a pure sinewave.

7.5 Future Directions

Our future work includes broadening the applications of StarAngle. Our results confirm that the phase information of Starlink satellites can be extracted reliably. As the phase measurement does not contain biases caused by clock mismatch, it is very robust and can potentially be very useful for more applications in addition to user orientation estimation. For example, one promising direction is to use the phase information measured by StarAngle to make corrections to the satellite orbital parameters because the phase values are very sensitive to the satellite orbit. The orbital parameters are published and maintained in public databases such as [5] but are updated only once every few hours and some of them may be outdated which has lead to some errors in our reported results. We find that, by making small adjustments to the orbital parameters, the calculated and measured phase values sometimes match better, which indicates that adjusted adjustments may be closer to the reality.

8 CONCLUSIONS

In this paper, we propose StarAngle, which is the first method that uses the phase difference of Starlink satellite beacons for user orientation estimation. We overcome the challenges of very weak signals and high Doppler rate by designing algorithms to detect the satellites and compensate for the Doppler rate. We subtracting the phase difference measurement of one satellite by another to cancel the clock mismatch in our commodity antennas and obtain measurements that agree very well with the theoretical calculation. We test StarAngle in 10 locations and the results show that the median estimation error is 7.5 degrees. StarAngle exploits the LOS path between the satellite and the ground receiver that allows for useful phase measurements and could enable more applications in addition to user orientation estimation. The performance of StarAngle will further improve in the future with more Starlink satellites in the space offering more useful phase observations.

ACKNOWLEDGMENTS

We greatly appreciate our shepherd and the reviewers of our paper for their valuable feedback and guidance. This research work was supported by the US National Science Foundation under Grant 2312113.

REFERENCES

[1] 2024. https://satellitemap.space/.

[2] Mohamed Amine Arfaoui, Mohammad Dehghani Soltani, Iman Tavakkolnia, Ali Ghrayeb, Chadi M. Assi, Majid Safari, and Harald Haas. 2021. Invoking Deep Learning for Joint Estimation of Indoor LiFi User Position and Orientation. *IEEE Journal on Selected Areas in Communications* 39, 9 (2021), 2890–2905. https://doi.org/10.1109/JSAC.2021.3064637

[3] P. Bahl and V.N. Padmanabhan. 2000. RADAR: an in-building RF-based user location and tracking system. In *Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064)*, Vol. 2. 775–784 vol.2. https://doi.org/10.1109/INFCOM.2000.832252

[4] Noori BniLam, Dennis Joosens, Michiel Aernouts, Jan Steckel, and Maarten Weyn. 2021. LoRay: AoA Estimation System for Long Range Communication Networks. *IEEE Transactions on Wireless Communications* 20, 3 (2021), 2005–2018. https://doi.org/10.1109/TWC.2020.3038565

[5] celestrak. 2024. https://celestrak.org/NORAD/elements/supplemental/sup-gp.php?FILE=starlink&FORMAT=tle.

[6] Junyoung Choi, Gyujin Lee, Sunghyun Choi, and Saewoong Bahk. 2022. Smartphone Based Indoor Path Estimation and Localization Without Human Intervention. *IEEE Transactions on Mobile Computing* 21, 2 (2022), 681–695. https://doi.org/10.1109/TMC.2020.3013113

[7] Ahmed Elzanaty, Anna Guerra, Francesco Guidi, and Mohamed-Slim Alouini. 2021. Reconfigurable Intelligent Surfaces for Localization: Position and Orientation Error Bounds. *IEEE Transactions on Signal Processing* 69 (2021), 5386–5402. https://doi.org/10.1109/TSP.2021.3101644

[8] B.H. Fleury, M. Tschudin, R. Heddergott, D. Dahlhaus, and K. Ingeman Pedersen. 1999. Channel parameter estimation in mobile radio environments using the SAGE algorithm. *IEEE Journal on Selected Areas in Communications* 17, 3 (1999), 434–450. https://doi.org/10.1109/49.753729

[9] Zhihui Gao, Yunfan Gao, Sulei Wang, Dan Li, and Yuedong Xu. 2021. CRISLoc: Reconstructable CSI Fingerprinting for Indoor Smartphone Localization. *IEEE Internet of Things Journal* 8, 5 (2021), 3422–3437. https://doi.org/10.1109/JIOT.2020.3022573

[10] Mark Harris. 2022. Starlink signals can be reverse-engineered to work like GPS—whether SpaceX likes it or not. *MIT Technology Review* (October 2022).

[11] Cory Hekimian-Williams, Brandon Grant, Xiuwen Liu, Zhenghao Zhang, and Piyush Kumar. 2010. Accurate localization of RFID tags using phase difference. In *2010 IEEE International Conference on RFID (IEEE RFID 2010)*. 89–96. https://doi.org/10.1109/RFID.2010.5467268

[12] K.C. Ho and Y.T. Chan. 1993. Solution and performance analysis of geolocation by TDOA. *IEEE Trans. Aerospace Electron. Systems* 29, 4 (1993), 1311–1322. https://doi.org/10.1109/7.259534

[13] Dave Hollander. March 27, 2019. How AoA & AoD Changed the Direction of Bluetooth Location Services. https://www.bluetooth.com/blog/new-aoa-aod-bluetooth-capabilities/.

[14] Todd E. Humphreys, Peter A. Iannucci, Zacharias M. Komodromos, and Andrew M. Graff. 2023. Signal Structure of the Starlink Ku-Band Downlink. *IEEE Trans. Aerospace Electron. Systems* 59, 5 (2023), 6016–6030. https://doi.org/10.1109/TAES.2023.3268610

[15] Toshiaki Koike-Akino, Pu Wang, Milutin Pajovic, Haijian Sun, and Philip V. Orlik. 2020. Fingerprinting-Based Indoor Localization With Commercial MMWave WiFi: A Deep Learning Approach. *IEEE Access* 8 (2020), 84879–84892. https://doi.org/10.1109/ACCESS.2020.2991129

[16] Manikanta Kotaru, Kiran Raj Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM 2015, London, United Kingdom, August 17-21, 2015*, Steve Uhlig, Olaf Maennel, Brad Karp, and Jitendra Padhye (Eds.). ACM, 269–282. https://doi.org/10.1145/2785956.2787487

[17] Sharbel Kozhaya, Haitham Kanj, and Zaher M. Kassas. 2023. Multi-Constellation Blind Beacon Estimation, Doppler Tracking, and Opportunistic Positioning with OneWeb, Starlink, Iridium NEXT, and Orbcomm LEO Satellites. In *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. 1184–1195. https://doi.org/10.1109/PLANS53410.2023.10139969

[18] Sharbel E. Kozhaya and Zaher M. Kassas. 2023. Positioning with Starlink LEO Satellites: A Blind Doppler Spectral Approach. In *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*. 1–5. https://doi.org/10.1109/VTC2023-Spring57618.2023.10199264

[19] Yuanjie Li, Hewu Li, Wei Liu, Lixin Liu, Wei Zhao, Yimei Chen, Jianping Wu, Qian Wu, Jun Liu, Zeqi Lai, and Han Qiu. 2023. *A Networking Perspective on Starlink's Self-Driving LEO Mega-Constellation*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3570361.3592519

[20] Zhi Lin, Kang An, Hehao Niu, Yihua Hu, Symeon Chatzinotas, Gan Zheng, and Jiangzhou Wang. 2023. SLNR-Based Secure Energy Efficient Beamforming in Multibeam Satellite Systems. *IEEE Trans. Aerospace Electron. Systems* 59, 2 (2023), 2085–2088. https://doi.org/10.1109/TAES.2022.3190238

[21] Zhi Lin, Min Lin, Benoit Champagne, Wei-Ping Zhu, and Naofal Al-Dhahir. 2021. Secrecy-Energy Efficient Hybrid Beamforming for Satellite-Terrestrial Integrated Networks. *IEEE Transactions on Communications* 69, 9 (2021), 6345–6360. https://doi.org/10.1109/TCOMM.2021.3088898

[22] Andrew Liszewski. 2022. The University of Texas Hacked Starlink's Signal So It Can Be Used as a GPS Alternative. *Gizmodo* (October 2022).

[23] Jun Liu, Jiayao Gao, Sanjay K. Jha, and Wen Hu. 2021. Seirios: leveraging multiple channels for LoRaWAN indoor and outdoor localization. In *ACM MobiCom '21: The 27th Annual International Conference on Mobile Computing and Networking, New Orleans, Louisiana, USA, October 25-29, 2021*. ACM, 656–669. https://doi.org/10.1145/3447993.3483256

[24] Adi Manos, Tamir Hazan, and Itzik Klein. 2022. Walking Direction Estimation Using Smartphone Sensors: A Deep Network-Based Framework. *IEEE Transactions on Instrumentation and Measurement* 71 (2022), 1–12. https://doi.org/10.1109/TIM.2022.3144225

[25] Mohammad Neinavaie and Zaher M. Kassas. 2023. Signal Mode Transition Detection in Starlink LEO Satellite Downlink Signals. In *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. 360–364. https://doi.org/10.1109/PLANS53410.2023.10139993

[26] Ousmane Abdoulaye Oumar, Ming Fei Siyau, and Tariq P. Sattar. 2012. Comparison between MUSIC and ESPRIT direction of arrival estimation algorithms for wireless communication systems. In *The First International Conference on Future Generation Communication Technologies*. 99–103. https://doi.org/10.1109/FGCT.2012.6476563

[27] Ales Povalac and Jiri Sebesta. 2011. Phase difference of arrival distance estimation for RFID tags in frequency domain. In *2011 IEEE International Conference on RFID-Technologies and Applications*. 188–193. https://doi.org/10.1109/RFID-TA.2011.6068636

[28] Tereza Pultarova and Elizabeth Howell. May 30, 2024. Starlink satellites: Facts, tracking and impact on astronomy. https://www.space.com/spacex-starlink-satellites.html (May 30, 2024).

[29] Xinping Rao, Zhenzhen Luo, Yong Luo, Yugen Yi, Gang Lei, and Yuanlong Cao. 2024. MFFALoc: CSI-Based Multifeatures Fusion Adaptive Device-Free Passive Indoor Fingerprinting Localization. *IEEE Internet of Things Journal* 11, 8 (2024), 14100–14114. https://doi.org/10.1109/JIOT.2023.3339797

[30] R. Roy and T. Kailath. 1989. ESPRIT-estimation of signal parameters via rotational invariance techniques. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 37, 7 (1989), 984–995. https://doi.org/10.1109/29.32276

[31] R. Schmidt. 1986. Multiple emitter location and signal parameter estimation. *IEEE Transactions on Antennas and Propagation* 34, 3 (1986), 276–280. https://doi.org/10.1109/TAP.1986.1143830

[32] Fernando Seco, Antonio R. Jiménez, and Francisco Zampella. 2013. Joint estimation of indoor position and orientation from RF signal strength measurements. In *International Conference on Indoor Positioning and Indoor Navigation*. 1–8. https://doi.org/10.1109/IPIN.2013.6817871

[33] Shengqiang Shen, Shiyin Li, and Heidi Steendam. 2022. Hybrid Position and Orientation Estimation for Visible Light Systems in the Presence of Prior Information on the Orientation. *IEEE Transactions on Wireless Communications* 21, 8 (2022), 6271–6284. https://doi.org/10.1109/TWC.2022.3148169

[34] Xiufang Shi, Brian. D. O. Anderson, Guoqiang Mao, Zaiyue Yang, Jiming Chen, and Zihuai Lin. 2016. Robust Localization Using Time Difference of Arrivals. *IEEE Signal Processing Letters* 23, 10 (2016), 1320–1324. https://doi.org/10.1109/LSP.2016.2569666

[35] Yimiao Sun, Yuan He, Jiacheng Zhang, Xin Na, Yande Chen, Weiguo Wang, and Xiuzhen Guo. 2023. BIFROST: Reinventing WiFi Signals Based on Dispersion Effect for Accurate Indoor Localization. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems, SenSys 2023, Istanbul, Turkiye, November 12-17, 2023*, M. Rasit Eskicioglu, Polly Huang, and Neal Patwari (Eds.). ACM, 376–389. https://doi.org/10.1145/3625687.3625786

[36] Xinyu Tong, Han Wang, Xiulong Liu, and Wenyu Qu. 2023. MapFi: Autonomous Mapping of Wi-Fi Infrastructure for Indoor Localization. *IEEE Transactions on Mobile Computing* 22, 3 (2023), 1566–1580. https://doi.org/10.1109/TMC.2021.3108155

[37] Xuyu Wang, Lingjun Gao, Shiwen Mao, and Santosh Pandey. 2017. CSI-Based Fingerprinting for Indoor Localization: A Deep Learning Approach. *IEEE Transactions on Vehicular Technology* 66, 1 (2017), 763–776. https://doi.org/10.1109/TVT.2016.2545523

[38] Yue Wang and K. C. Ho. 2017. TDOA Positioning Irrespective of Source Range. *IEEE Transactions on Signal Processing* 65, 6 (2017), 1447–1460. https://doi.org/10.1109/TSP.2016.2630030

[39] Kainam Thomas Wong and M.D. Zoltowski. 2000. Self-initiating MUSIC-based direction finding and polarization estimation in spatio-polarizational beamspace. *IEEE Transactions on Antennas and Propagation* 48, 8 (2000), 1235–1245. https://doi.org/10.1109/8.884492

[40] Kaishun Wu, Jiang Xiao, Youwen Yi, Dihu Chen, Xiaonan Luo, and Lionel M. Ni. 2013. CSI-Based Indoor Localization. *IEEE Transactions on Parallel and Distributed Systems* 24, 7 (2013), 1300–1309. https://doi.org/10.1109/TPDS.2012.214

[41] Chun Yang and Andrey Soloviev. 2023. Starlink Doppler and Doppler Rate Estimation via Coherent Combining of Multiple Tones for Opportunistic Positioning. In

2023 IEEE/ION Position, Location and Navigation Symposium (PLANS). 1143–1153. https://doi.org/10.1109/PLANS53410.2023.10140055

[42] Kang Yang, Yuning Chen, and Wan Du. 2024. OrchLoc: In-Orchard Localization via a Single LoRa Gateway and Generative Diffusion Model-based Fingerprinting. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services, MOBISYS 2024, Minato-ku, Tokyo, Japan, June 3-7, 2024*, Tadashi Okoshi, JeongGil Ko, and Robert LiKamWa (Eds.). ACM, 304–317. https://doi.org/10.1145/3643832.3661876

[43] Kehu Yang, Gang Wang, and Zhi-Quan Luo. 2009. Efficient Convex Relaxation Methods for Robust Target Localization by a Sensor Network Using Time Differences of Arrivals. *IEEE Transactions on Signal Processing* 57, 7 (2009), 2775–2784. https://doi.org/10.1109/TSP.2009.2016891

[44] Xiheng Chen Yanhan Zeng, Yuxing Liao and Hong zhou Tan. 2023. UHF RFID Indoor Localization Based on Phase Difference. *IETE Journal of Research* 69, 6 (2023), 3511–3517. https://doi.org/10.1080/03772063.2021.1912649

[45] Nathanael Yoder. Retrieved April 23, 2024. peakfinder(x0, sel, thresh, extrema, includeEndpoints, interpolate). *MATLAB Central File Exchange* https://www.mathworks.com/matlabcentral/fileexchange/25500-peakfinder-x0-sel-thresh-extrema-includeendpoints-interpolate (Retrieved April 23, 2024).

[46] Li You, Xiaoyu Qiang, Ke-Xin Li, Christos G. Tsinos, Wenjin Wang, Xiqi Gao, and Björn Ottersten. 2022. Massive MIMO Hybrid Precoding for LEO Satellite Communications With Twin-Resolution Phase Shifters and Nonlinear Power Amplifiers. *IEEE Transactions on Communications* 70, 8 (2022), 5543–5557. https://doi.org/10.1109/TCOMM.2022.3182757

[47] Dian Zhang, Jian Ma, Quanbin Chen, and Lionel M. Ni. 2007. An RF-Based System for Tracking Transceiver-Free Objects. In *Fifth Annual IEEE International Conference on Pervasive Computing and Communications (PerCom'07)*. 135–144. https://doi.org/10.1109/PERCOM.2007.8

[48] Dian Zhang and Lionel M. Ni. 2009. Dynamic clustering for tracking multiple transceiver-free objects. In *2009 IEEE International Conference on Pervasive Computing and Communications*. 1–8. https://doi.org/10.1109/PERCOM.2009.4912777

[49] Tianyu Zhang, Dongheng Zhang, Guanzhong Wang, Yadong Li, Yang Hu, Qibin Sun, and Yan Chen. 2023. RLoc: Towards Robust Indoor Localization by Quantifying Uncertainty. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 7, 4 (2023), 200:1–200:28. https://doi.org/10.1145/3631437

[50] Minghui Zhao, Tyler Chang, Aditya Arun, Roshan Sai Ayyalasomayajula, Chi Zhang, and Dinesh Bharadia. 2021. ULoc: Low-Power, Scalable and cm-Accurate UWB-Tag Localization and Tracking for Indoor Applications. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 3 (2021), 140:1–140:31. https://doi.org/10.1145/3478124

9 ARTIFACT APPENDIX

9.1 Abstract

The artifact is the source code of StarAngle written in MATLAB, which analyzes the beacon signals of Starlink Low Earth Orbit (LEO) satellites received by two antennas and produces an estimate of the azimuth angle of the line connecting the center of the two antennas. Accompanying the source code are all the trace files of the experiments in the paper.

9.2 Artifact check-list (meta-information)

- **Algorithm: LEO Satellite Signal Detection; Doppler Rate Estimation; Orientation Estimation**
- **Data set: Starlink Beacon Signal Trace; USRP**
- **Run-time environment: MATLAB**
- **Hardware: Low-Noise Block Downconverter (LNB), USRP**
- **Metrics: Orientation Estimation Accuracy**
- **Output: Orientation Angle**
- **Experiments: Outdoor**
- **How much disk space required (approximately)?: Code less than 1 MB; Trace 35 GB in total**
- **How much time is needed to prepare workflow (approximately)?: Less than 1 hour to set up the code to process the uploaded traces**
- **How much time is needed to complete experiments (approximately)?: ; About 2 min to process a trace and about 90 min to process all traces**
- **Publicly available?: Yes**
- **Code licenses (if publicly available)?: Free**
- **Data licenses (if publicly available)?: Free**
- **Workflow automation framework used?: No**

9.3 Description

9.3.1 How to access. The source code can be found at https://github.com/raghavrathi10/StarAngle
The traces files can be downloaded at https://zenodo.org/records/13881520

9.3.2 Hardware dependencies. StarAngle is written in MATLAB, so any hardware that supports MATLAB can be used. We use a computer with 11th Gen Intel(R) Core(TM) i7-11700 @ 2.50GHz, 8 cores, 16 threads, 56 GB ram, NVIDIA GeForce GT 730 2GB memory.

9.3.3 Software dependencies. MATLAB R2023b or above along with the following toolboxes: Aerospace toolbox, Communications toolbox, Curve Fitting toolbox, Fixed-Point Designer, and Satellite Communications toolbox.

9.3.4 Data sets. The experiments discussed in the paper are conducted in 10 locations. For each location, we test 2 angles except two locations. For every angle, we collect 3 good runs except 2 cases for which we have 2 good runs. In total, we have 52 traces. The Starlink satellite beacons are on 11.95 GHz. We connect two LNBs facing vertically above to a USRP B210 and take samples at 2 Msps, where each sample consists of a real part and an imaginary part both as 16-bit integers. Each run include five files. For example, for a run we refer to as `sig1`, the files are:

- `sig1.tle`: TLE of satellites flying over during the experiment,
- `sig1.cfg`: information about the experiment, such as the start time, longitude, latitude, the ground truth of the orientation, etc.;
- `sig1_ant1`: samples received from antenna 1;
- `sig1_ant2`: samples received from antenna 2;
- `sig1_foundbeacons.mat`: saved results for comparison.

In addition, a file named `All_sat.tle` can be found for runs with the same angle at the same location, which is the TLE of all Starlink satellites downloaded before experiments.

9.4 Installation

The user will need to install MATLAB R2023b or above first and then install the required Toolboxes in MATLAB mentioned earlier in "software dependencies." After downloading the source file, there should a directory called `StarAngle`. Under `StarAngle`, there should be a directory called `stana`, which is the MATLAB source code directory. Also under `StarAngle` should be a directory called `GNURadio` which stores `2ant.grc` which is the file we use for signal capture. The trace data should be downloaded to a directory under `StarAngle` called `expdata`.

9.5 Experiment workflow

The first step is to start MATLAB, go to the `stana` directory, and run `stana_init` in the command window. There are three ways to run the code.

9.5.1 Processing All Traces. Simply run `stana_wrapper`.

9.5.2 Processing A Specific Trace. To process a specific trace, simply change the trace file name in `stana_process.m`, such as
`thisfilename = '../expdata/06-16/sig1_ant1';"`
and run `stana_process`.

9.5.3 Processing Own Trace. To collect and process own trace data, say, `sig2`, the following steps need to be followed:

- Obtain two LNBs and a USRP such as B210.
- Select an experiment location and place the LNBs on sturdy stands facing upwards.
- Create a directory to store files for this experiment.
- Download the latest Starlink satellite TLE database by running `updateTLE`, which produces `All_sat.tle`, and move it to the experiment directory.
- Write down information about the experiment and save it in `sig2.cfg` in the experiment directory. Lines in `sig2.cfg` should be:
  - Start time of the experiment in UTC, such as `16-Jun-2024 12:22:06`
  - Duration of the experiment measured in seconds.
  - Latitude of the experiment measured in degrees.
  - Longitude of the experiment measured in degrees.
  - Azimuth angle of the line connecting the center of the LNBs measured in degrees.
  - Distance between the center of the LNBs measured in meters.
- It is suggested to use `https://satellitemap.space/` to monitor the Starlink satellites flying over the experiment area and start the experiment when there are multiple satellites present.
- Run experiment with the GNU Radio Companion using `2ant.grc`. Name the files as `sig2_ant1` and `sig2_ant2`, and store them in the experiment directory.
- Modify the following lines in `stana_process.m`:
  `thisfilename = 'your directory/sig2_ant1';`
  `STANA_DO_GEN_EXP_TLE_FILE = 1;`
  `usrp_sig_data_type = 'float';`
  because `sig2.tle` need to be generated and the data type from `2ant.grc` is `float`.
- Run `stana_process`.

9.6 Evaluation and expected results

The final output of the program, for example, is:
`best matching user az angle: 24.00 (deg); cost: 0.44; actual angle: 30.00 (deg); error: -6.00 (deg)`