

# GPU-accelerated Direct Geolocation of GNSS Interference

Jacob S. Clements<sup>\*†</sup>, and Zachary L. Clements<sup>†</sup>

<sup>\*</sup>*Department of Computer Science, University of Maryland, College Park*

<sup>†</sup>*Department of Aerospace Engineering and Engineering Mechanics, University of Texas at Austin*

**Abstract**—In recent years, there has been a sharp increase in Global Navigation Satellite Systems (GNSS) interference, which has proven to be problematic in GNSS-dependent civilian applications. Many currently deployed GNSS receivers lack the proper countermeasures to defend themselves against interference, prompting the need for alternative defenses. Satellites in Low Earth Orbit (LEO) provide an opportunity for GNSS interference detection, classification, and localization. The direct geolocation approach has been shown to be well-suited for low SNR regimes and in cases limited to short captures—exactly what is expected for receivers in LEO. Direct geolocation is a single-step search over a geographical grid that enables estimation of the transmitter location directly from correlating raw observed signals. However, a key limitation to this approach is the computational requirements. This computational burden is compounded for LEO-based receivers as the geographic search space is extensive. This paper alleviates the computational burden of direct geolocation by exploiting the independence of position-domain correlation across candidate points and time steps: nearly all computation can be accomplished in parallel on a graphics processing unit (GPU). This paper presents and evaluates the performance of GPU-accelerated direct geolocation compared to traditional CPU processing.

**Index Terms**—GPU acceleration; emitter geolocation; interference localization; spectrum monitoring.

## I. INTRODUCTION

Global Navigation Satellite Systems (GNSS) such as the Global Positioning System (GPS) offer precise positioning and timing across the globe. Besides their broad coverage, GNSS are robust to weather conditions and are available to users for free. Due to these useful properties, there are a myriad of technologies that rely on GNSS. For example, positioning information obtained from GNSS is used for navigation in automobiles, ships, and aircraft. Additionally, the timing information provided by GNSS is used to synchronize stock market exchanges and control power grid switches.

Despite their uses, GNSS signals are susceptible to interference due to being extremely weak. For users on Earth, GNSS signals have no more flux density than a 50 W light bulb shining at a distance of 2000 km away [1]. Because of this, GNSS receivers are highly susceptible to deliberate GNSS interference such as jamming or spoofing.

Although they are illegal for individuals to use, there are GNSS interference devices on the public market and can be purchased for as little as \$10. One of such devices is

a so-called “personal privacy device” which jams GNSS-based tracking devices so that the tracking device’s position cannot be relayed. Because the authentic GNSS signals are so weak, these personal privacy devices can potentially impact GNSS receivers multiple kilometers away [2]. In 2013, a man in New Jersey was fined \$32,000 when his GPS jammer unintentionally disrupted operations at Newark Airport [3].

Wide-area GNSS interference has been observed near ongoing conflict zones. The civilian maritime and airline industries are encountering GNSS jamming and spoofing at an alarming rate [4]–[8]. Widespread jamming and spoofing is indicated by monitoring anomalous positioning information broadcast by ships in Automatic Identification System (AIS) messages, and airplanes in Automatic Dependent Surveillance-Broadcast (ADS-B) messages. All of the civilian GNSS receivers falling victim to GNSS interference are likely unintended targets caught in the electronic warfare crossfire near ongoing conflict zones.

In general, GNSS interference can be separated into two categories: (1) jamming and (2) spoofing. Jamming is a blunt attack where the jammer transmits strong interference signals at the same frequency as GNSS, with the goal of denial of service. On the other hand, spoofing is aimed at deception, trying to trick a GPS receiver into inferring a false navigation solution. In the simplest case, jammers can emit a pure tone at the GNSS center frequencies. Chirp jammers can deny a broader range of the GNSS spectrum by varying the frequency of the emitted signal with time to occupy a larger bandwidth. Such chirp jammers appear to be popular in commercially available civil GPS jammers [2], [9]. More sophisticated attacks such as spoofing can trick GNSS users into computing a false navigation solution by feeding them false GNSS signals [10].

Given the vulnerabilities in GNSS, there has been significant research into navigation robustness and security within the past two decades [1], [11]–[16]. Another significant research thrust has been developing alternate forms of positioning, navigation, and timing (PNT). Mega-Low Earth Orbit (LEO) constellations can be opportunistically exploited for navigation [17]–[24]. Alternate terrestrial radio-based positioning [25], [26], alternate sensors such as cameras [27], [28], and different receiver architectures [29], [30] have been explored. Although such progress has been inspiring, there

is still a need to help defend the currently deployed receivers that have not incorporated such defenses.

For GNSS users to avoid or account for GNSS interference, it can be helpful to geolocate these interference sources. A solution that achieves this is receivers in LEO, which have been shown to be able to geolocate GNSS interference [31]–[37]. Satellites in LEO offer worldwide and consistent coverage. LEO satellites are also typically far enough away from GNSS interference sources to acquire and track authentic GNSS signals. The timing information obtained from authentic GNSS signals allows multiple receivers to synchronize samples in their data captures, which allows for precise interference emitter geolocation. With multiple satellites, time difference of arrival (TDOA) and frequency difference of arrival (FDOA) measurements can be exploited to estimate the location of GNSS interference sources.

The traditional approach to estimating an emitter's location is a two-step approach using time-frequency difference of arrival (T/FDOA) measurements [38], [39]. In this two-step approach, T/FDOA measurements are first made using the complex ambiguity function (CAF). A CAF is calculated at different snapshots to create a time series of T/FDOA measurements. Then, nonlinear estimation is applied to this time series data to estimate the emitter's location.

This two-step approach has some drawbacks. In the first step, the CAF may have properties that make it hard to extract T/FDOA measurements. Peaks in the CAF are used to determine T/FDOA measurements, but there may be multiple peaks in the CAF if there are multiple emitters or if an emitter is cyclostationary. Because of these ambiguities, a tracking algorithm is usually used to track and associate peaks across CAFs. This tracking and association algorithm struggles in crowded signal environments, especially ones with strong emitters with cyclostationary behavior.

An alternate approach to geolocation using multiple receivers is direct geolocation [33], [40]–[42]. In direct geolocation, the TDOA and FDOA are directly parameterized for a single geographical point, given knowledge of the receivers' position and velocity, and then the raw signals correlated in accordance with the T/FDOA. This is repeated for a set of geographical candidate locations. In the direct approach, cyclostationary behavior of a single emitter does not interfere with the geolocation of other emitters. This makes the direct approach ideal for scenarios with multiple emitters that possibly exhibit cyclostationary behavior and for scenarios that have both strong and weak emitters.

A key challenge these T/FDOA-based geolocation approaches face is computation time. As in the case of the direct approach, data captures that are potentially millions of samples long must be time-aligned according to the calculated TDOA and then FDOA aligned according to the calculated FDOA. Additionally, the direct approach requires this process to be repeated for each candidate point. This computational burden is compounded for LEO-based receivers as the geo-

graphic search space is extensive. For example, for a 10 degree by 10 degree grid defined in Latitude-Longitude-Altitude (LLA) coordinates with 0.01 degree spacing, a million position-domain correlation values must be computed.

This paper attempts to alleviate the computational constraints of direct emitter geolocation by using a graphics processing unit (GPU) to accelerate the computation. The independence of the position-domain correlation values between individual candidate locations is exploited by computing their position-domain correlation values in parallel on a GPU. To account for the vast amount of data being processed on a GPU with limited memory, batch processing is used and then optimized to minimize computational time. A simulation engine and synthetic dataset are developed to show a large speedup in the time required to run direct geolocation compared to traditional CPU processing.

## II. SIGNAL MODELS

### A. Received Signal Model

Consider a stationary emitter on the surface of the Earth transmitting a baseband signal  $x(t)$  modulating a carrier. When this emitted signal reaches a moving receiver, the observed signal includes both a time delay  $\tau$  due to the propagation time and a Doppler shift  $f_D$  due to the relative motion. Both  $\tau$  and  $f_D$  are accurately modeled as constant over a snapshot (short intervals less than 50 ms) as the relative motion over this short interval is negligible. After downconversion, the baseband signal received by the  $i$ th receiver is given by

$$y_i(t) = x(t - \tau_i) \exp(j2\pi f_i t) \quad (1)$$

where  $\tau_i$  and  $f_i$  are the time delay and Doppler shift for the  $i$ th receiver.

Let  $N_s$  represent the number of samples,  $f_s$  be the sampling rate, and  $T_s$  be the time between samples. Let  $y_i[k]$  represent the discretized sampled representation of the received signal  $y_i(t)$ .

$$y_i[k] = y_i(kT_s) \quad (2)$$

where  $k$  goes from 0 to  $N_s - 1$ . Let the time  $t_k$  be the time at the  $k$ th sample, starting from 0 and going to  $T_s(N_s - 1)$ .

Consider a stationary emitter with position vector  $\mathbf{p}_e$  specified in the Earth-centered-Earth-fixed (ECEF) coordinate frame. At a given snapshot, for the  $i$ th receiver, let  $\mathbf{p}_i$  be its position vector and  $\mathbf{v}_i$  be its velocity vector, both in ECEF. Define the range vector between the emitter and the  $i$ th receiver  $\mathbf{r}_i$  as

$$\mathbf{r}_i = \mathbf{p}_i - \mathbf{p}_e \quad (3)$$

The range between the emitter and the  $i$ th receiver  $\rho_i$  is then

$$c\tau_i = \rho_i = \sqrt{\mathbf{r}_i^T \mathbf{r}_i} \quad (4)$$

where  $c$  is the speed of light. The unit range vector  $\hat{\mathbf{r}}_i$  can then be defined as

$$\hat{\mathbf{r}}_i = \frac{\mathbf{r}_i}{\rho_i} \quad (5)$$

The observed Doppler shift at the  $i$ th receiver  $f_i$  is given by

$$f_i = -\frac{1}{\lambda} \hat{\mathbf{r}}_i^T \mathbf{v}_i \quad (6)$$

where  $\lambda$  is the wavelength of the signal.

Different receivers will receive the emitter's signal at different time delays. For two time synchronized receivers, the time difference of arrival between the  $i$ th receiver and the  $j$ th receiver is defined as

$$\Delta\tau_{ij} = \tau_j - \tau_i \quad (7)$$

The frequency difference of arrival of the  $i$ th and  $j$ th receivers is then defined as

$$\Delta f_{ij} = f_j - f_i \quad (8)$$

These physical received characteristics of a transmitted signal can be exploited to estimate an emitter's position.

These models assume the relative clock bias and clock drift between the receivers can be compensated for. Typically, this can be achieved by using the onboard GNSS navigation solution. Each receiver can use their onboard navigation solution to synchronize with GPS time, and in effect, synchronize to the other receivers.

### III. DIRECT GEOLOCATION

Direct geolocation is a single-step grid-search method that solves directly for the emitter position without the need for intermediate T/FDOA measurements. The CAF is maximized directly by parameterizing the delay and Doppler time histories in terms of the emitter's position, as well as the known receivers' position and velocity. The cross-correlation between received samples is performed in the position domain instead of the delay-Doppler domain. For a large number of measurements the two-step approach is equivalent to the direct approach [41]. However, the direct approach outperforms the two-step approach in low SNR regimes and in cases limited to short captures. This estimator is proven to be maximum likelihood in [41].

The set of candidate points defined by the user is typically a grid of points defined in LLA coordinates. The user only needs to define these points in latitude and longitude, since the altitude for each point can be constrained to the surface of the Earth, or be retrieved from a digital terrain model. For the previously stated measurement model to be valid, these coordinates are transformed into an ECEF coordinate system using the standard transformations. Since this coordinate system is rectangular, the measurement model applies if the receiver positions are also expressed in ECEF.

Let  $\mathbf{s}_i = [\mathbf{p}_i^T, \mathbf{v}_i^T]^T$  contain the position and velocity of the  $i$ th receiver and  $\mathbf{p}_c$  denote the candidate point. At a

snapshot with a fixed receiver geometry, the expected TDOA and FDOA for each  $\mathbf{p}_c$  can be calculated using (7) and (8). For  $\mathbf{p}_c$ , let the expected TDOA  $\Delta\bar{\tau}_{ij}$  and FDOA  $\Delta\bar{f}_{ij}$  be

$$\Delta\bar{\tau}_{ij} = \text{round}((\Delta\tau_{ij}(\mathbf{p}_c, \mathbf{s}_i, \mathbf{s}_j))/f_s) \quad (9)$$

$$\Delta\bar{f}_{ij} = \Delta f_{ij}(\mathbf{p}_c, \mathbf{s}_i, \mathbf{s}_j) \quad (10)$$

Then the position-domain correlation value at  $\mathbf{p}_c$  is as

$$S(\mathbf{p}_c) = \left| \sum_{k=0}^{N_s-1} y_1[k] y_2^*[k + \Delta\bar{\tau}_{ij}] \exp(j2\pi\Delta\bar{f}_{ij}t_k) \right| \quad (11)$$

The maximum value of  $S$  is the maximum likelihood estimate of the emitter's position [41]. Direct geolocation is summarized in Algorithm 1. Note, the shape and sharpness of the peak is dependent on: (1) receiver geometry, (2) transmitted waveform, (3) coherent integration time, and (4) received signal power. As will be shown later, a major advantage of the direct approach is the ability to noncoherently accumulate the position-domain correlation values across snapshots. Due to the transmitted waveform, structures can arise in a position-domain correlation grid that can make detecting peaks in an individual snapshot difficult. When position-domain correlation grids are noncoherently accumulated, those undesirable structures will end up being reduced to noise.

A position-domain correlation grid for a GPS L1 C/A spoofer is shown in Fig. 1. As expected, there is a sharp peak that corresponds to the GPS L1 C/A spoofer's position.




Fig. 1: Example position-domain correlation grid from direct geolocation for a single GPS L1 C/A spoofer.

Shown in Fig. 2 is an example position-domain correlation grid shown for a sawtooth jammer. A sawtooth jammer

exhibits cyclostationary properties and has a highly structured auto-ambiguity function, creating several sidelobes, even in the position domain. As will be shown later, this is not problematic when there are multiple snapshots with differing geometries.




Fig. 2: Example position-domain correlation grid from direct geolocation for a sawtooth jammer, which exhibits cyclostationary properties. There are several undesirable sidelobes due to the waveform.

---

**Algorithm 1:** Direct Geolocation

---

**Input :** set of candidate points; raw samples, positions, and velocities from the receivers

**Output:** Emitter position estimates

---

```

1 Convert all candidate positions to ECEF
2 for each snapshot do
3   Set  $\mathbf{s}_i = [\mathbf{p}_i^T, \mathbf{v}_i^T]^T$ 
4   Set  $\mathbf{s}_j = [\mathbf{p}_j^T, \mathbf{v}_j^T]^T$ 
5   for each candidate point do
6     Calculate TDOA  $\Delta\tau(\mathbf{p}_c, \mathbf{s}_i, \mathbf{s}_j)$ 
7     Calculate FDOA  $\Delta f(\mathbf{p}_c, \mathbf{s}_i, \mathbf{s}_j)$ 
8     Compute position-domain correlation value  $S$ 
9   end
10 end
11 Sum the position-domain correlation grids
12 Apply a threshold to detect emitters

```

---

#### IV. SIMULATION AND GEOLOCATION RESULTS

This section presents the developed simulation engine and an example capture scenario. A high-fidelity simulation engine is essential for testing and algorithm development.

##### A. Simulation Engine

A high-fidelity simulation engine was developed to explore direct geolocation of GNSS interference sources. Within this simulation engine framework, the user can specify any arbitrary number of receivers with any geometry. The user can also specify any number of GNSS interference sources at any location. Each GNSS interference source transmits one of the four of the most common GNSS interference waveforms: (1) GPS L1 C/A spoofing, (2) pure tone, (3) chirp, and (4) sawtooth. These transmitted waveforms are generated at baseband and then the received waveform at each receiver is generated.

The GPS L1 C/A waveform is a spread spectrum signal where the carrier is modulated by a 50 bits per second navigation data sequence, which is further modulated with the C/A code [43]. Both the navigation data and C/A employ binary phase shift keying (BPSK). The C/A code is a pseudorandom noise (PRN) sequence that is different for each navigation satellite. Each C/A code is a Gold code with maximal length linear shift register sequence of length 1023, that is upsampled to a chipping rate of 1.023 Mbps and lasts 1 millisecond long. Every millisecond, the upsampled 1023 length code repeats. All of the Gold sequences have good auto-correlation functions, which is important for ranging precision, and uniformly low cross-correlation properties. GNSS spoofers transmit spoofing signals that mimic the authentic signals, so a would-be GPS L1 C/A spoofer would transmit GPS L1 C/A signals. Additionally, matched-code and matched-spectrum interference will take on a similar construction.

A pure tone jammer is by far the simplest form of jamming, as it is simply a tone at the carrier frequency. Chirp and sawtooth jamming waveforms are different, as they can deny a broader range of the GNSS spectrum by varying the frequency of the emitted signal with time to occupy a larger bandwidth. Let  $B_{\text{chirp}}$  be the bandwidth of a chirp, and  $T_{\text{chirp}}$  be the chirp period. The baseband representation of a chirp  $x_{\text{chirp}}(t)$  is defined as

$$x_{\text{chirp}}(t) = \exp\left(j2\pi\left(\frac{B_{\text{chirp}}}{2T_{\text{chirp}}}t^2 - \frac{B_{\text{chirp}}}{2}t\right)\right) \quad (12)$$

A chirp jammer transmits  $x_{\text{chirp}}(t)$  repeatedly. A sawtooth jammer is slightly different, as it transmits  $x_{\text{chirp}}(t)$  followed by  $\text{conj}(x_{\text{chirp}}(t))$  for a period  $T_{\text{sawtooth}} = 2T_{\text{chirp}}$ .

The power spectral density and spectrogram of the GPS L1 C/A, pure tone, chirp, and sawtooth waveforms are shown in Fig. 3.



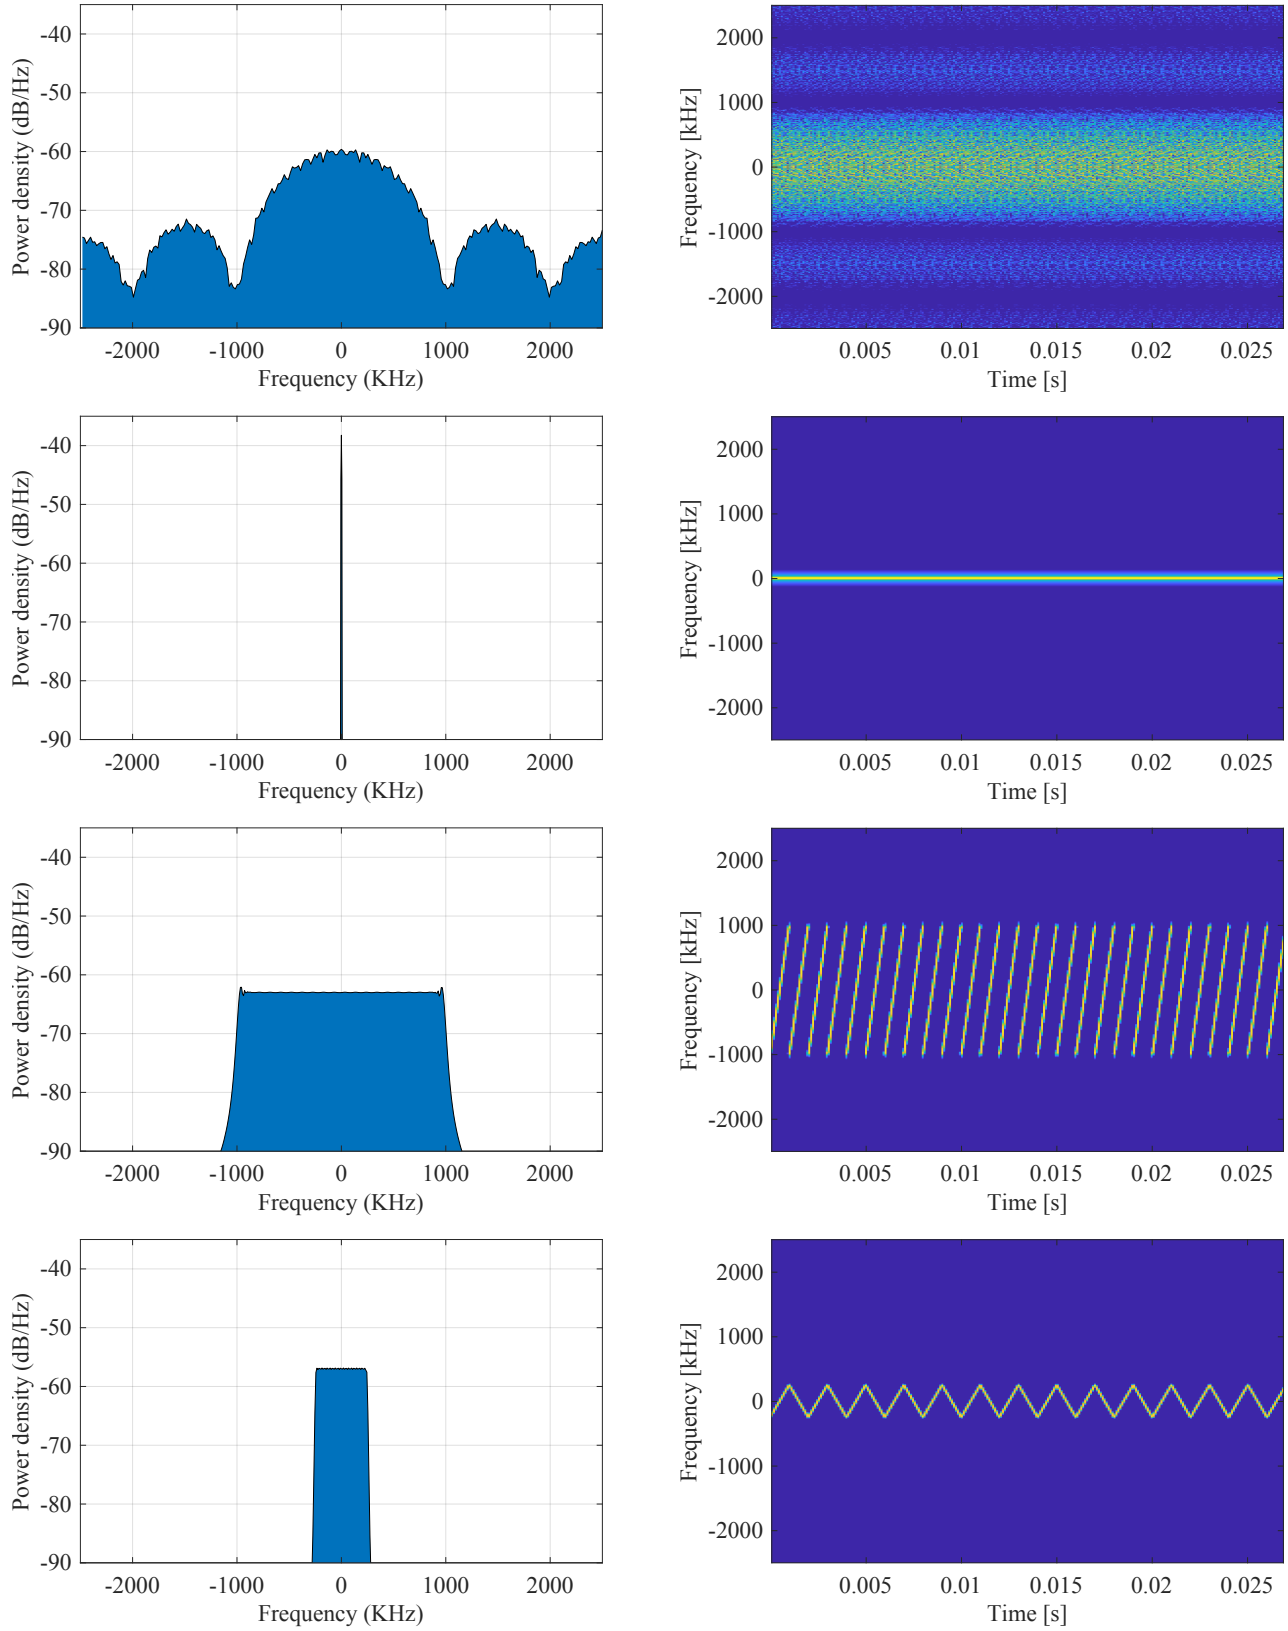


Fig. 3: Left: Power spectral densities of spoofing, pure tone, chirp, and sawtooth interference waveforms. Right: Spectrograms of spoofing, pure tone, chirp, and sawtooth interference waveforms.

### B. Simulated Capture Scenario and Geolocation Results

Fig. 4 shows the relevant physical portions of the simulated capture scenario. The ground tracks of the two LEO-based receivers over 58 seconds are shown. A GPS L1 C/A spoofer, pure tone jammer, chirp jammer, and sawtooth jammer were simulated to be in a rectangular formation on the surface of the Earth. The chirp jammer was set to have a bandwidth of 2 MHz and a chirp duration of  $20 \mu\text{s}$ . The sawtooth jammer was set to have a bandwidth of 200 kHz and a duration of 5 ms. All interference sources were set to be transmitting centered at the GPS L1 frequency (1575.42 MHz) at the same power. The receivers performed 58 raw captures, each spaced by a second. Each raw capture was 50 ms long and the sampling rate was set to 5 MHz (complex) centered at the GPS L1 frequency (1575.42 MHz).

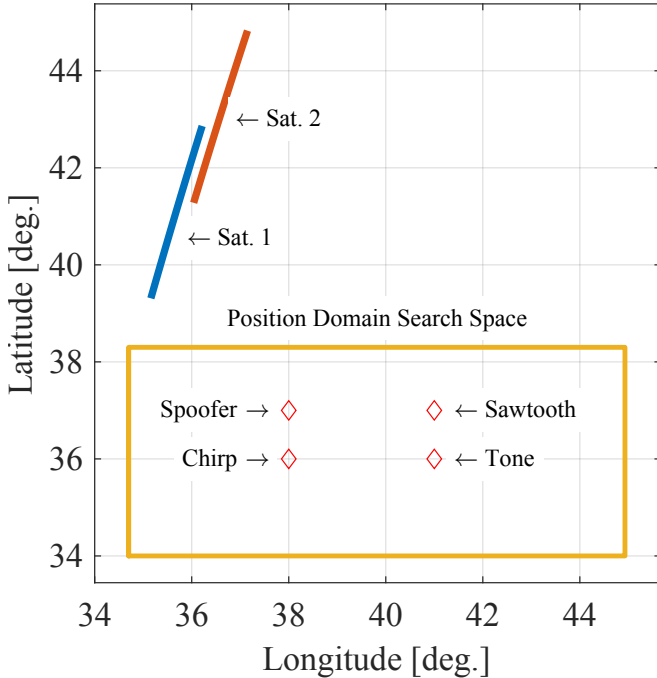


Fig. 4: Ground track and search space

The position-domain search space shown in Fig. 4 spans approximately  $10^\circ$  in longitude and approximately  $4^\circ$  in latitude. The position-domain correlation was computed at each of the 58 snapshots. The position-domain correlation grids for the first snapshot and the last snapshot are shown in Fig. 5. One can immediately see the different correlation structures for the different waveforms. The GPS L1 C/A spoofer creates a single sharp peak, the tone creates a large streak, the chirp jammer creates multiple equally prominent peaks, and the sawtooth jammer creates an “X” shape. At a single snapshot, the pure tone and chirp jammers’ positions are ambiguous. But, as seen between the first and last snapshot, the structure of each GNSS interference source changes. This is due to the changing receiver geometry. However, the peaks at the true

emitter positions do not change from the start of the capture compared to the end of the capture.




Fig. 5: Top: position-domain correlation at the first snapshot. Bottom: position-domain correlation at the final snapshot.

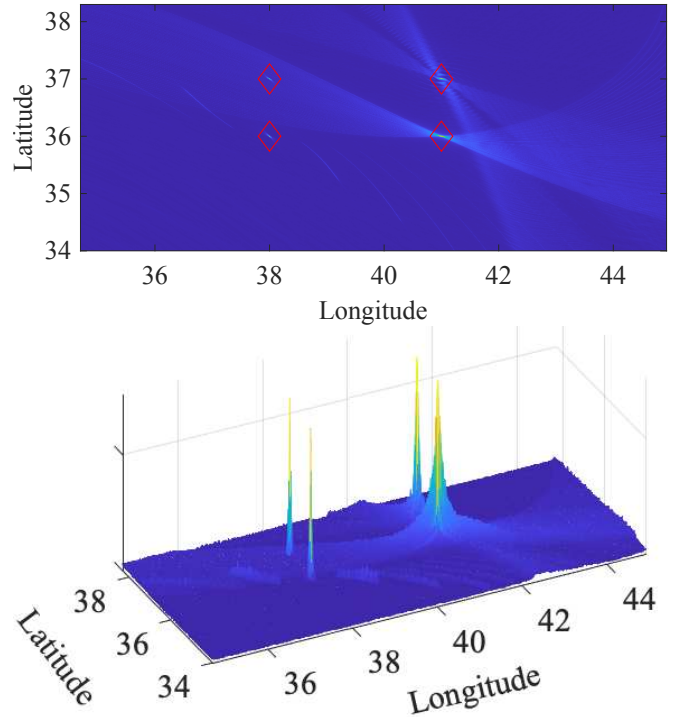


Fig. 6: position-domain correlation after accumulating 58 snapshots. The four distinct sharp peaks represent the MLE of each GNSS interference source’s position.

Shown in Fig. 6 is the position-domain correlation grid when the 58 position-domain grids at each snapshot are non-coherent accumulated. One can immediately appreciate four sharp peaks centered at each of the four GNSS interference sources' true position. As the receiver geometry changes over the 58 second interval, the positions of the sidelobes and other undesirable structures due to the waveform change. When position-domain grids are noncoherently accumulated, the undesirable structures are reduced to the noise floor, while the true peaks persist. This shows how powerful direct geolocation can be in crowded signal environments with waveforms that have undesirable structures.

## V. GPU-ACCELERATED ANALYSIS

### A. Hardware and Software Setup

The hardware setup consists of a computer equipped with an Intel 20-core, 28-thread i7-14700K processor and an Nvidia GTX 4060 Ti GPU. In terms of memory, the computer had 32 gigabytes of RAM and the GPU had 16 gigabytes of memory. To utilize the GPU for parallel processing, initial tests were done using Matlab's Parallel Computing Toolbox, PyTorch, and CuPy, with PyTorch being used for the final implementation. PyTorch was also used for the CPU implementation comparisons, and it is worth noting that PyTorch automatically parallelizes many CPU computations under the hood. Unless otherwise noted, all numbers related to performance were averaged over 10 runs.

### B. GPU-Accelerated Implementation

The loops shown in Algorithm 1 are embarrassingly parallelizable since the body of each iteration is independent of all the other iterations for both of the loops. The GPU implementation vectorizes the inner loop to compute the TDOA, FDOA, and position-domain correlation values in parallel on a GPU. Since the memory inside the GPU is limited, this process is batched up. Effectively, the inner for loop becomes a loop over batches of candidate points, and the contents of the for loop (TDOA, FDOA, and position-domain correlation) are all parallelized on the GPU.

Processing each batch of candidate points consists of three steps: (1) load the candidate points and relevant data onto the GPU, (2) perform the computation, and (3) offload the results to CPU memory. To efficiently utilize the limited GPU memory, the final steps of the algorithm, namely the summing over snapshots and thresholding, were left on the CPU. This is done because if the position-domain correlation values are left on the GPU, the amount of free memory on the GPU would deteriorate as the algorithm progresses from one snapshot to the next. Additionally, the summing and thresholding steps were already quite fast and very little performance would be gained by transferring the data back to the GPU to do the computation.

### C. GPU-Accelerated Analysis

To optimize the batch size, the time required to compute the TDOA, FDOA, and position-domain correlation value for about half a million candidate points is measured. In the coarse search shown in Fig. 7, there is an immediate drop going from a batch size of 1 to a batch size of 2. After reaching a minimum, the curve rises before plateauing at a batch size of around 500. This plateau happens because the GPU memory is full and has to borrow from CPU memory. After performing a fine search over the region around the minimum in the coarse search, a batch size of 8 was found to be the fastest. Interestingly, batch sizes of 32 and 64 appeared as local minimums, which corresponds with the conventional wisdom that powers of 2 are the most efficient choices for batch sizes.




Fig. 7: Top: Coarse search of computation time for different batch sizes. The plateau occurs when the GPU memory is full and has to borrow from CPU memory. Bottom: Fine search of computation time for different batch sizes.

In general, there were massive improvements in the computational time required to compute the same number of points on a GPU vs on the CPU. Fig. 8 illustrates the CPU vs GPU speedup multiplier for a varying number of input candidate points. The speedup remains relatively stable, and Fig. 8 shows that the GPU implementation provides a speedup in the range of about 26 times faster to about 28 times faster compared to the CPU implementation.




Fig. 8: Speedup multiplier when going from CPU to GPU.

## VI. CONCLUSION

GPU-accelerated direct geolocation paired with data from LEO-based receivers is a powerful combination that helps maintain GNSS security. Direct geolocation is a single-step search over a geographical grid that enables estimation of the transmitter location directly from correlating raw observed signals. However, a key limitation of this approach is the computational time required, since it involves correlating large amounts of data. This computational burden is compounded for LEO-based receivers as the geographic position-domain search space is extensive. This paper addressed this issue by performing the direct geolocation algorithm in parallel on a GPU. Specifically, the TDOA, FDOA, and position-domain correlation values are each computed in parallel. The results show that the GPU-implementation saw speedups of up to around 28 times when compared to a CPU-based implementation.

## REFERENCES

- [1] T. E. Humphreys, "Interference," in *Springer Handbook of Global Navigation Satellite Systems*, pp. 469–503, Springer International Publishing, 2017.
- [2] R. Mitch, R. Dougherty, M. Psiaki, S. Powell, B. O'Hanlon, J. Bhatti, and T. Humphreys, "Signal characteristics of civil GPS jammers," in *Proceedings of the ION GNSS Meeting*, 2011.
- [3] CBS, "N.J. man in a jam, after illegal GPS device interferes with newark liberty operations," *CBS News*, Aug. 2013.
- [4] C4ADS, "Above us only stars: Exposing GPS spoofing in Russia and Syria," April 2019. <https://c4ads.org/reports>.
- [5] S. Gebrekidan, "Electronic warfare confounds civilian pilots, far from any battlefield," *The New York Times*, Nov. 2023.
- [6] J. Arraf, "Israel fakes GPS locations to deter attacks, but it also throws off planes and ships," *NPR*, April 2024.
- [7] A. Tangel and D. FitzGerald, "Electronic warfare spooks airlines, pilots and air-safety officials," *Wall Street Journal*, Sept. 2024.
- [8] N. Bulos, "Israel's 'GPS spoofing' tricks missiles, but also commercial airplanes, dating apps in mideast," *Los Angeles Times*, July 2024.
- [9] R. Mitch, R. Dougherty, M. Psiaki, S. Powell, B. O'Hanlon, J. Bhatti, and T. E. Humphreys, "Know your enemy: Signal characteristics of civil GPS jammers," *GPS World*, Jan. 2012.
- [10] M. L. Psiaki and T. E. Humphreys, "GNSS spoofing and detection," *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258–1270, 2016.
- [11] M. L. Psiaki and T. E. Humphreys, "GNSS spoofing and detection," *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258–1270, 2016.
- [12] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, "Review article: GPS vulnerability to spoofing threats and review of antispooing techniques," *International Journal of Navigation and Observation*, pp. 1–16, 2012.
- [13] M. L. Psiaki and T. E. Humphreys, *Position, Navigation, and Timing Technologies in the 21st Century: Integrated Satellite Navigation, Sensor Systems, and Civil Applications*, vol. 1, ch. Civilian GNSS Spoofing, Detection, and Recovery, pp. 655–680. Wiley-IEEE, 2020.
- [14] Z. Clements, J. E. Yoder, and T. E. Humphreys, "Carrier-phase and IMU based GNSS spoofing detection for ground vehicles," in *Proceedings of the ION International Technical Meeting*, (Long Beach, CA), pp. 83–95, 2022.
- [15] Z. Clements, J. E. Yoder, and T. E. Humphreys, "GNSS spoofing detection: An approach for ground vehicles using carrier-phase and inertial measurement data," *GPS World*, vol. 34, no. 2, pp. 36–41, 2023.
- [16] C. Tanil, S. Khanafseh, M. Joerger, and B. Pervan, "An INS monitor to detect GNSS spoofers capable of tracking vehicle position," *taes*, vol. 54, pp. 131–143, Feb. 2018.
- [17] Z. M. Komodromos, W. Qin, and T. E. Humphreys, "Signal simulator for Starlink Ku-Band downlink," in *Proceedings of the ION GNSS+ Meeting*, pp. 2798–2812, 2023.
- [18] Z. M. Komodromos, W. Qin, and T. E. Humphreys, "Weak signal acquisition and tracking of the Starlink Ku-Band downlink to enable global PNT," in *Proceedings of the ION Joint Navigation Conference (JNC)*, 2024.
- [19] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–16, 2023.
- [20] W. Qin, Z. M. Komodromos, and T. E. Humphreys, "An analysis of the short-term time stability of the Starlink Ku-band downlink frame clock," in *Proceedings of the IEEE International Conference on Wireless for Space and Extreme Environments (WISEE 2024)*, 2024. Submitted for review.
- [21] W. Qin, A. M. Graff, Z. L. Clements, Z. M. Komodromos, and T. E. Humphreys, "Timing properties of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, 2024. In preparation.
- [22] W. Qin, Z. M. Komodromos, and T. E. Humphreys, "An agile, portable antenna system for LEO megaconstellation-based PNT," in *Proceedings of the ION GNSS+ Meeting*, 2023.
- [23] S. Morgan and T. E. Humphreys, "HOOC-EM: Fast beam sweeping for LEO mega-constellation customer terminals," in *Proceedings of the ION GNSS+ Meeting*, 2024.
- [24] T. E. Humphreys, Z. L. Clements, and W. Qin, "GNSS interference: Situational awareness and LEO backup," *Center for Automated Vehicles Research with Multimodal Assured Navigation (CARMEN+) Tier-1 University Transportation Center*, 2024.
- [25] R. Tenny and T. E. Humphreys, "Robust navigation for urban air mobility via tight coupling of GNSS with terrestrial radionavigation and inertial sensing," in *Proceedings of the ION GNSS+ Meeting*, (Denver, CO), 2022.
- [26] R. Tenny and T. E. Humphreys, "Robust navigation for urban air mobility," *Inside GNSS*, vol. 18, no. 2, pp. 47–47, 2023.



- [27] R. M. Tenny, L. C. Sun, A. Duru, and T. E. Humphreys, "Robust absolute headset tracking for extended reality," in *Proceedings of the IEEE/ION PLANS Meeting*, (Monterey, CA), 2023.
- [28] S. G. Catalan, B. A. Jones, and J. S. McCabe, "Implementation of machine learning methods for crater-based navigation," in *Astrodynamics Specialist Conference*, 2021.
- [29] Z. Clements, P. A. Iannucci, T. E. Humphreys, and T. Pany, "Optimized bit-packing for bit-wise software-defined GNSS radio," in *Proceedings of the ION GNSS+ Meeting*, (St. Louis, MO), pp. 3749–3771, 2021.
- [30] W. J. Morrison, T. E. Humphreys, and D. A. Ton-Nu, "Robust beam-pattern synthesis for UAV-swarm-based distributed beamforming," in *Proceedings of the ION GNSS+ Meeting*, 2024.
- [31] D. M. LaChapelle, L. Narula, and T. E. Humphreys, "Orbital war driving: Assessing transient GPS interference from LEO," in *Proceedings of the ION GNSS+ Meeting*, (St. Louis, MO), 2021.
- [32] M. J. Murrian, L. Narula, P. A. Iannucci, S. Budzien, B. W. O'Hanlon, S. P. Powell, and T. E. Humphreys, "First results from three years of GNSS interference monitoring from low Earth orbit," *NAVIGATION, Journal of the Institute of Navigation*, vol. 68, no. 4, pp. 673–685, 2021.
- [33] Z. Clements, P. Ellis, and T. E. Humphreys, "Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit," in *Proceedings of the IEEE/ION PLANS Meeting*, (Monterey, CA), pp. 458–469, 2023.
- [34] Z. Clements, P. B. Ellis, and T. E. Humphreys, "Pinpointing GNSS interference from low earth orbit," *Inside GNSS*, vol. 18, no. 5, pp. 42–55, 2023.
- [35] Z. Clements, P. Ellis, M. L. Psiaki, and T. E. Humphreys, "Geolocation of terrestrial GNSS spoofing signals from low Earth orbit," in *Proceedings of the ION GNSS+ Meeting*, (Denver, CO), pp. 3418–3431, 2022.
- [36] Z. Clements, I. Goodridge, P. Ellis, M. J. Murrian, and T. E. Humphreys, "Demonstration of single-satellite GNSS spoofer geolocation," in *Proceedings of the ION International Technical Meeting*, (Long Beach, CA), pp. 361–373, 2024.
- [37] Z. L. Clements, P. Ellis, I. Goodridge, M. J. Murrian, M. L. Psiaki, and T. E. Humphreys, "Analysis and demonstration of single-satellite geolocation of broadcast GNSS spoofers from low Earth orbit," *NAVIGATION, Journal of the Institute of Navigation*, 2024. In preparation.
- [38] A. Sidi and A. Weiss, "Delay and Doppler induced direct tracking by particle filter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, pp. 559–572, January 2014.
- [39] K. Ho and Y. Chan, "Geolocation of a known altitude object from TDOA and FDOA measurements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, pp. 770–783, July 1997.
- [40] A. M. Reuven and A. J. Weiss, "Direct position determination of cyclostationary signals," *Signal Processing*, vol. 89, no. 12, pp. 2448–2464, 2009.
- [41] A. Weiss, "Direct geolocation of wideband emitters based on delay and Doppler," *Signal Processing, IEEE Transactions on*, vol. 59, pp. 2513–2521, June 2011.
- [42] J. Bhatti, *Sensor Deception Detection and Radio-Frequency Emitter Localization*. PhD thesis, The University of Texas at Austin, Aug. 2015.
- [43] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*. Lincoln, Massachusetts: Ganga-Jumana Press, revised second ed., 2012.