

Detection of On-Road Vehicles Emanating GPS Interference

Gorkem Kar[†], Hossen Mustafa[‡], Yan Wang[‡]
Yingying Chen[‡], Wenyuan Xu[‡], Marco Gruteser[†], Tam Vu^{*}

[†]Winlab, Rutgers University, North Brunswick, NJ 08902, USA

[†]{gkar87,gruteser}@winlab.rutgers.edu

[‡]Stevens Institute of Technology, Hoboken, NJ 07030, USA

[‡]{ywang48,yingying.chen}@stevens.edu

[‡]University of South Carolina, Columbia, SC 29208, USA

[‡]{mustafah, wyxu}@cse.sc.edu

^{*}University of Colorado Denver, Denver, CO 80204, USA

^{*}tam.vu@ucdenver.edu

Abstract

The Global Positioning System (GPS) is widely used in critical infrastructures but is vulnerable to radio frequency (RF) interference. A common source of interference are commercial drivers that use GPS jammers to circumvent vehicle tracking systems. Existing mechanisms to detect and identify such interference emitting vehicles on roadways require a large number of specialized detectors or a manual observation process. In this paper, we design a practical, automated system to facilitate enforcement actions. Our system combines information from roadside monitoring points at key locations along the roadway as well as mobile detectors (e.g., smartphones and other mobile GPS systems). Rather than attempting precise localization at a given time, the system exploits the inherent variation in driving speeds and the resulting diverging trajectories of vehicles to uniquely identify the interfering vehicle. Through our experiments on a local highway with a vehicle transmitting interference in the 900MHz ISM band, we found that the vehicle identification rate of our mechanism is 65% for a single-point setup and 100% for a two-point setup. We performed 200 hours of passive monitoring of GPS L1 band on roadways and found two episodes of real interference. We also demonstrate that our mobile detector-based profiles are sufficiently consistent in time and space to enable reliable interference detection.

Categories and Subject Descriptors

C.5.0 [Computer System Implementation]: [General]

Keywords

GPS; jamming; vehicular

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS'14, November 3–7, 2014, Scottsdale, Arizona, USA.

Copyright 2014 ACM 978-1-4503-2957-6/14/11 ...\$15.00.

<http://dx.doi.org/10.1145/2660267.2660336>.

1. INTRODUCTION

Many networked systems, particularly in the transportation domain rely on the Global Positioning System (GPS) for location tracking. For example, automatic vehicle locator system are installed in truck fleets to monitor delivery progress and optimize routes [19], in taxicabs to facilitate cab-passenger scheduling [4], and also in aeronautical applications such as landing systems. If current plans for connected vehicles with Vehicle-2-Vehicle Dedicated Short Range Communications [28] are any guide, this reliance on GPS will only increase. In this effort, vehicles are expected to broadcast their current GPS reading and recent trajectory multiple times per second, so that each vehicle can be aware of nearby vehicle movements and use this information to improve efficiency and safety.

GPS Interference and Jamming. This reliance on GPS has also exposed its susceptibility to interference. Due to the 20,200km distance to the transmitting satellites, the received signal strength is extremely low at about -125dBm. This is much weaker than in most other wireless communication systems and below the noise floor. These signals can therefore easily be masked by other transmissions, either from jammers or malfunctioning electronic equipment. A particular source of concern are in-vehicle GPS jammers. These appear to be used mainly by drivers to gain location privacy, that is to prevent employer-installed GPS vehicle tracking systems from recording their trajectories. GPS jammers are easy to use, some can be simply plugged into the vehicles cigarette lighter for power, and require no knowledge about the vehicle tracking system or the location of the GPS receiver. They can be removed as easily as they are installed and leave no evidence of tampering. To be effective regardless of the type and mounting position of the GPS antenna and across vehicle types from large commercial trucks to small passenger vehicles, the output power of these devices will have to be relatively high. Such generously proportioned output powers, however, increase the likelihood of interfering with other nearby receivers. For this reason, the United States and many other countries have significant regulatory restrictions in place for such devices.

The Detection Challenge. Still, both anecdotal evidence such as incidents at Newark Airport and a study in the UK [7] indicate that use of jamming devices is common




Figure 1: Overview of jamming detection system consisting of monitoring points, mobile detectors, and central office for camera-based vehicle identification.

enough so that it can interfere with current critical infrastructure and could be a concern in any future connected vehicle deployment. Moving jammers might appear to cause only brief interference and be less of a concern. However, some GPS systems, particularly stationary ones, may not be designed to expect frequent interference and even a short interference spike could cause malfunction (as in the initial Newark airport system). For mobile GPS receivers passing at high speed, this may be less of an issue, but receivers in vehicles driving along with the interferer could still experience outages over several minutes. Presumably regulatory restrictions can reduce their use, but enforcing these restrictions, particularly at the individual vehicle level, is very labor intensive. The accuracy of standard wireless localization techniques [14, 16, 23] is usually not sufficient to pinpoint a particular vehicle in dense traffic. Current methods require manually pointing a handheld monitoring device with a directional antenna receiver at individual vehicles to scan for emissions. Expending this level of effort¹ is usually only practical when frequent interference has been reported at a fixed location.

Automated Monitoring and Detection. This paper seeks to overcome these challenges with an automated interference detection and vehicle identification system, with includes roadside interference monitors, mobile interference detectors and camera-based vehicle identification, as shown in Figure 1. The roadside component of such a system could remain in continuous operation near critical infrastructure and thereby quickly identify interference in such areas. The mobile detectors that leverage smartphones or other mobile devices opportunistically to identify interference can achieve wide coverage at low overhead. It requires a special profiling approach, however, to distinguish interference from benign fading and attenuation based on the limited GPS perfor-

¹As an example, consider a case at Newark airport where repeated interference at GPS reference receivers from the landing system was observed [15]. It took many months for FAA/FCC investigators to locate one jammer. During this time frame, FAA/FCC teams convened with contractors twice for three day investigation sessions until it was determined that the interference source is mobile and moving on the I-95 highway next to the airport. In an additional session, the FAA/FCC team with contractors used a spectrum analyzer on an overpass to identify a set of candidate vehicles that might carry a jammer. These vehicles were then pursued with patrol vehicles carrying mobile interference detectors or spectrum analyzers until a single vehicle could be isolated (apparently based on interference signal strength and proximity to the patrol car). This vehicle was then stopped to confirm that a jamming device is present inside the vehicle.




Figure 2: (a) shows one of the commonly available jammer and (b) shows the spectrogram of a 8 MHz wide jamming signal generated using Matlab function.

mance indicators available on widely used mobile devices. Our approach combines data from multiple detections to identify vehicles. The inherent variation of driving speeds means that a group of cars that passes a first monitoring point together likely has dispersed at a later monitoring point, which allows unique identification of the vehicle emitting such signals. The salient contributions of this work are summarized below.

- We propose an automated two-pronged interference monitoring framework with roadside monitoring stations and mobile detectors that can achieve wide coverage at low cost.
- We design profile-based algorithms that detect GPS signal anomalies based on widely available satellite signal-to-noise ratio measurements (e.g., accessible from Android apps). Even though GPS signals are highly variable due to fading, our method of constructing signal profiles is sufficiently robust to allow interference detection without any specialized chipset functions.
- Instead of precise triangulation with multiple receivers to locate the interference source, our system can work with just a single detector (assuming repeated commuting by the jammer) or at most two detectors to identify the interference source.
- We evaluate through multiple days of roadway experiments on local highways the likelihood that a monitoring station can uniquely identify an interference emitting vehicle and show that unique identification is possible with two appropriately spaced monitoring stations in most cases.
- We use our system to investigate the prevalence of GPS interference on US public roadways. During 200 hours of passive monitoring in two US cities, we captured two instances of interference.

2. SYSTEM OVERVIEW

Our proposed detection system uses unmanned roadside monitoring stations in critical locations (e.g., airports) and crowdsourced GPS signal-to-noise ratio data from mobile phones or vehicles to achieve broad coverage. It is designed to significantly reduce the effort required to detect and track down interference in the civilian use L1 GPS carrier at 1.57542 GHz. We envision that such a system can be used for identifying malfunctioning equipment that generates interference or for enforcement actions against GPS jammers.




Figure 3: Overview of our GPS jamming detection methodology. Each monitoring point uses a USRP for capturing wireless data and each mobile detector uses phone’s GPS APIs to collect SNR values. The centralized camera-based vehicle identification system uses traffic camera images for accurate vehicle identification.

2.1 Attack Model: Personal Privacy GPS Jammers

The attackers are mainly drivers that want to gain location privacy by circumventing employer-installed GPS vehicle tracking systems with GPS jammers. We target relatively widespread low-cost jammers that are used as personal privacy devices rather than sophisticated defense-related equipment. Such personal privacy devices are often small jammers that can be powered by an automotive power outlet.

These GPS jammers simply overpower GPS signals with interference. Although some advanced jammers can block the L1, L2, L3, L4 and L5 bands, the most commonly used ones block only the L1 band. Such GPS jammers can be categorized into two categories in terms of signal characteristics [1]: (a) Continuous wave jammers simply transmit an unmodulated continuous wave signal with a bandwidth less than 100 kHz [24], and (b) Chirp jammers transmit different types of narrowband and wideband chirp signals where the chirp signal can be a saw-tooth function, a multiple saw-tooth function, or a frequency bursts.

The jamming bandwidth ranges from 0.92 kHz to 44.9 MHz and the output power varies from -9.5 dBm to -30.8 dBm [2]. Chirp signal jammers are most common [2]—we show an example in Figure 2(a). For our experiments and analysis, we generated a chirp signals similar to the characteristics reported by Bauernfeind et al [2]. An 8 MHz chirp signal is depicted in Figure 2(b).

2.2 Design Requirements

Towards this goal, we identify the following design requirements for identifying vehicles emanating GPS interference.

- **Automated.** The detection and identification system should be able to operate around the clock with limited human intervention. While manual examination of detection results may still be required, continuous human presence at a monitoring location should not be necessary.
- **Low-cost.** To facilitate widespread usage, infrastructure costs should be minimal. Some infrastructure is appropriate for priority locations around critical GPS infrastructure (e.g., airports).
- **High accuracy.** Identifying the source of the interference source should be accurate enough to support enforcement actions. False positive can lead to inconveniencing innocent drivers and false negative represents missed detections.

2.3 Challenges

Identifying GPS jammers carried by vehicles in roadways is challenging for several reasons.

Availability of GPS performance metrics. Widely used GPS receivers only provide high-level GPS performance metrics that are not obvious indicators of interference. Identification of interference in the GPS L1 band is not particularly difficult with specialized equipment or high performance GPS receivers that monitor interference conditions. Such equipment, however, is costly or at least not in widespread use. Many GPS receivers only allow monitoring satellite signal-to-noise ratios (SNR) and outages (i.e., a loss of the position fix). Both metrics, however, are also affected by the varying signal conditions under normal operation.

Vehicle identification in complex environment. Roadways present both a complex radio propagation and traffic environment. The radio propagation on a multi-lane highway is highly irregular because of the strong attenuation and fading caused by the metal bodies of vehicles, the Doppler shift introduced by high traveling speeds, and unpredictable radio propagation paths due to the dynamic traffic patterns. All the irregularity make it challenging to accurately identify the vehicle that carries a GPS jammer.

Several projects have proposed to localize GPS jammers, but none of them can automatically identify vehicles with GPS jammers accurately and economically. For instance, most of the latest research on GPS jammer localization are based on the time difference of arrival (TDOA) technique [14], [16], [23]. However, TDOA-based approach requires either high accurate synchronization between multiple observing stations or the prior knowledge of GPS jamming signals. In addition, these localization mechanisms are not accurate enough to identify a single vehicle. While some of them can localize the source of jamming signals with an accuracy of 5 meters in certain cases, this is still not sufficient to distinguish two vehicles driving on adjacent lanes. Thus, a new vehicle identification solution is in need.

2.4 Approach

We chose a two-pronged approach with specialized roadside detectors and mobile detectors because the former promise higher system accuracy around critical infrastructure and the latter offer low-cost, broad coverage of an entire transportation network. As shown in Figure 3, both types of detectors will report signal anomalies to a centralized system. Rather than attempting precise localization at a single point in time, the system uses multiple reports to identify the source vehicle. To do so, it must have access to vehicle identification information such as electronic toll system identifiers or camera footage that reveals license plates. We focus here on the latter as an example.

For specialized roadside equipment, there exist established detection methods. In the interest of low complexity, we rely on energy-based detection, wherein a **Wireless Receiver** continuously monitors the GPS L1 frequency band and the **Interference Detector** monitors whether energy exceeds the normal noise levels at this location. Suspicious signals can further be validated against established jamming patterns by the **Signal Analyzer** and are reported to the centralized system together with their location.

Interference detection is most challenging when seeking to collect interference data from widely deployed, mobile GPS receivers. We chose to focus on smartphone GPS receivers accessible through Android apps, because this is arguably

the most common GPS platform. For example, such mobile detectors could be carried in government vehicle fleets or with appropriate energy optimizations offered as an Android background service to the public. The Android platform makes available carrier-to-noise density ratio—loosely called signal-to-noise (SNR)—measurements of each satellite in addition to the GPS position data, but does not provide any information that would directly indicate interference. Interference, however, can be expected to manifest itself as additional noise at the receiver and therefore decrease the carrier to noise density ratio (or, a loss of fix, in the extreme). The challenges lie in distinguishing this effect from signal fading or attenuation, which is prevalent in wireless propagation environments and also leads to carrier to noise density ratio decreases or outages.

Our mobile detection approach therefore profiles the environment to establish reference signal levels for a certain receiver. Signals are then classified as anomalies, when they deviate significantly from this profile. They can be further validated by detecting abnormal patterns in carrier-to-noise-density ratio variations. These three steps make up the GPS Data Collector, Interference Detector, and Jamming Validator components. We discuss detection in more detail in Section 3.

The centralized camera-based vehicle identification component identifies a vehicle using images from roadside units or traffic cameras. For roadside detection, it is most convenient to have cameras collocated with the interference detector. Still, in many cases one can expect multiple vehicles to be present in the image and it will be difficult to identify which one is the interference source. We address this challenge by relying on multiple detections in time or space. Since groups of vehicles will eventually disperse, multiple detections (of the same jammer) will eventually show only one common vehicle. When mobile detections are reported, there may not be any camera available at this location. In this case, the system will have to interpolate the vehicles path to a camera location, and work with a larger set of candidate vehicles to account for the path estimation uncertainties. Still, with a sufficient number of detections, unique identification can be expected. We discuss this component in Section 4.

3. JAMMING DETECTION

While many sophisticated jamming detection solution exist, we focus here on techniques that can be implemented with existing widely available GPS receiver hardware or low cost software radios. In particular, we develop a carrier-to-noise profiling approach which only relies on the carrier-to-noise metrics (sometimes imprecisely referred to as signal-to-noise ratio) reported by most GPS receivers. This metric is even accessible for apps through the Android location API. We also consider energy detection which can easily be implemented on dedicated monitoring devices and is also supported in more sophisticated GPS receivers. While not nearly as widely available as the carrier-to-noise metric, energy detection allows a more direct measure of interference since it does not need to account for variations in satellite signal propagation.

3.1 Carrier-to-Noise Profiling

To achieve broad coverage with this monitoring system, we design a solution that can take advantage of existing GPS chips and devices such as smartphones or in-car GPS

Algorithm 1 Mobile Detector

```

1: function PROFILESNR( $n, d$ )
    $\triangleright$  record time series data for  $n$  satellites
2:    $< time, SNR_{sat}, lat, lon > = record\_data(n)$ 
3:    $p = calculate\_path\_distance(lat, lon)$ 
4:   for each path segment  $p_i$  with length  $d$  do
5:      $sat_{snr}(p_i) = get\_sat\_snr(SNR_{sat})$ 
6:      $m_{snr}(p_i) = average(sat_{snr}(p_i))$ 
7:      $\tau_{snr}(p_i) = mode(sat_{snr}(p_i))$ 
8:      $\sigma_{snr}(p_i) = std\_dev(sat_{snr}(p_i))$ 
9:      $save\_profile(p_i, m_{snr}, \tau_{snr}, \sigma_{snr})$ 
10:  end for
11: end function
12: function DETECTSNRANOMALY( $n, k, p_i$ )
13:    $sat_{snr} = read\_sat\_snr(n)$ 
14:    $m_{snr} = average(sat_{snr})$ 
15:    $\tau_{snr} = mode(sat_{snr})$ 
16:    $\sigma_{snr} = std\_dev(sat_{snr})$ 
17:    $< m_{pr}, \tau_{pr}, \sigma_{pr} > = get\_profile(p_i)$ 
18:    $\delta = m_{pr} - k * \sigma_{pr}$ 
19:   if  $m_{snr} + (\tau_{pr} - \tau_{snr}) \leq \delta$  then
20:      $jamming\_status = verify\_all\_sat()$ 
21:     if  $jamming\_status = true$  then
22:        $[lat, lon] = get\_location()$ 
23:        $send\_report(lat, lon, m_{snr})$ 
24:     end if
25:   end if
26: end function

```

receivers connected to telematics units. The carrier-to-noise density ratio (C/N_0) is the most fine-grained GPS performance measure that conveys information about interference and is provided by nearly every GPS chip as well as made accessible through the NMEA standard or through the Android API. It is the ratio of the satellite signal power to noise density and is often somewhat imprecisely referred to as signal-to-noise ratio (SNR) in many GPS interface documentations.²

The C/N_0 is primarily affected by signal variations which are due to fading, both from obstructions (e.g., skyscrapers, tall billboards), satellite movement, and multipath propagation. Under interference conditions, however, the interfering signal acts similar to noise on the satellite signal reception process. The reported C/N_0 is therefore usually a signal-to-interference-plus-noise density ratio, which will decrease when either interference increases or signal fading occurs.

The challenge is then to distinguish fading from interference. We propose to detect abnormally low SNR values by comparing them with SNR profiles that characterize typical fading variations at a location. We rely here on three key observations: (i) in-vehicle jammers and interferers are mobile and rarely dwell for an extended amount of time in the same location (indeed, they are usually deactivated when a vehicle parks); (ii) vehicles pathways are highly constrained by roadways and tend to repeatedly pass through those; (iii) most significant signal fading is caused by stationary structures such as tall buildings, underpasses, tunnels, etc. The first observation leads us to the heuristic that a temporarily low SNR at a particular location is indicative of interference,

²In wireless communications, SNR typically includes noise over the entire receiver bandwidth, while C/N_0 uses noise density, that is noise over 1 Hz. We use the terms interchangeably here to refer to the carrier-to-noise ratio density.

while a more permanently low SNR at a particular location is more likely due to obstructions in the built environment. The second and third observations leads us to the idea to construct a profile of typical signal levels at each location. If detectors are mobile, each small interval of length d along a roadway is treated as a separate location, which will eventually lead to a signal map. By comparing current SNR readings with this profile, rather than a global threshold, we can then perform outlier detection to determine whether the readings are unusually low for this particular location. This should significantly reduce false positives. Thus, our algorithm contains two stages: profiling, and jamming detection. The generalized mobile version is shown in Algorithm 1 and we describe the individual stages next.

Profiling. In this stage, we first collect C/N_0 values along a route from the n strongest satellites for each of the position readings along the route. Note that as satellites orbit the earth, the identity of these satellites will change but we found C/N_0 readings to be sufficiently consistent to not warrant further distinction. For each interval p_i along a roadway, we calculate and store the mean m , the mode τ , and the standard deviation σ of the C/N_0 values from all satellites and all position updates within this interval. By calculating a mean over multiple C/N_0 readings (from different satellites but also different locations within the interval or different trips), we will average out multipath effects and the approach becomes more robust to small scale fading. The mode is provided as a statistic that is more robust to outliers. We chose 20m as the interval length d , because it is sufficiently fine-grained to capture signal variations caused by smaller structures, and large enough so that usually at least one sample from a trip falls into this interval. These values comprise the profile and can be continuously updated as additional C/N_0 data from the same location arrives (i.e., the vehicle travels over the same road again).

Anomaly Detection. In this stage, we use the parameter recorded in the profiling stage to determine whether the SNR at a particular location is abnormally low. First, the current reported latitude and longitude are mapped to a particular road interval p_i . Second, we will again calculate the mean SNR m_{snr} over all satellites and all location updates from the current trip that fall into this interval. This is reasonable even if we expect interference, because (i) interference will affect all satellites equally since they are transmitting on the same L1 carrier and (ii) in most cases interference will be present long enough for the vehicle to travel a short distance d (e.g., 20m). Thus, the averaging is unlikely to remove the effect of interference.

The mean of the profile and the sigma of the profile are then used to derive a threshold for outlier detection. An abnormal signal report is generated if $m_{snr} < m_{pr} - k * \sigma_{pr}$. Since the normal distribution is often successfully used to model the distribution of SNR readings one might expect at a given distance, it can serve also as a guide for determining the k parameter of the threshold. While this is not necessary for built-in vehicle GPS receivers with external antennas, the threshold for portable GPS receivers can be further calibrated to account for placement differences inside the vehicle from one trip to the next. For this purpose, the algorithm will also track the mode of the SNR τ_{snr} over an extended distance d_{ext} , which is chosen longer than the longest expected interference duration. The algorithm then estimates the profile mode over the same extended interval based on the m_{pr} that make up this interval. The difference




Figure 4: Ladder-shaped detection zone on the road.

between this estimated profile mode and the current measured mode is a calibration parameter that will be added to m_{snr} before the comparison.

When the SNR falls below this threshold, a signal anomaly report can be transmitted to a central office. This report will contain the difference between the threshold and the mean SNR value as well as the current or the last known GPS location, if the current one is unavailable. In the latter case, the first new location fix may also be provided subsequently, to allow a better estimation of the location of the low signal. Complete outages, where no satellites can be detected, are treated as the lowest possible SNR value in the algorithm. The low signal report can further include indicators of detected signal characteristics which imply a higher confidence that this signal level is due to jamming, such as the high variance described earlier. The central office can take the magnitude of the SNR difference and such indicators into account, when determining the reliability of such reports.

3.2 Comparison with Energy-based Detection

Roadside units designed for interference monitoring can rely on many custom techniques to detect interference. Energy detection is a simple yet effective technique in the GPS band, since the GPS signal level is typically below the thermal noise floor (confirmed in our experiments). Therefore, any energy-level above the noise floor is unwanted interference. This approach greatly simplifies the choice of detection threshold since no signal variations have to be taken into account but it requires hardware that supports such metrics.

Based on these considerations, we chose the carrier-to-noise-profiling approach in our mobile smartphone detectors and implemented energy detection in our dedicated roadside detector using a Universal Software Radio Peripheral (USRP) N210 [8] in conjunction with a WBX daughterboard and with GNURADIO [11] scripts. The roadside unit first automatically sets the threshold value just above the ambient noise level. For monitoring, the **wireless receiver** block (Figure 3) continuously captures 10 milliseconds of data and passes it to **interference detector** block. When the energy becomes larger than the threshold ($-75dBm$ to $-83dBm$ in our experiments), the interference detector records a suspicious signal and reports it together with timestamp information. A **signal analyzer** can optionally conduct temporal and spectral analysis to gain additional confidence in the report.

4. SOURCE VEHICLE IDENTIFICATION

In this section, we describe how to correlate the detection with license plates from traffic camera footage to allow identification of a set of candidate vehicles and how this set can be narrowed down to a unique vehicle after multiple detections. While we discuss this in the context of traffic cameras, similar methods could be applied to other vehicle identifi-




Figure 5: Positions of transmitting vehicle passing by the detector at time t_m when the maximum energy value is identified in multiple experimental runs in real-road scenarios. cation methods such as IDs from electronic toll collection systems.

4.1 Identification Using Roadside Detector

To facilitate vehicle identification, the system needs to map an energy detection to a portion of a video obtained from a traffic camera that the monitoring point is collocated with. It is therefore convenient to focus the wireless detections on a smaller area that is in view of the camera. This can be achieved by using a directional antenna at the monitoring point and results in a trapezoidal area of the roadway from which signals can be detected. Such a detection zone is illustrated in Figure 4 (marked as *Zone 1*). Let us first consider a naive approach to identify a set of candidate vehicles. Given a time interval $[t_s, t_e]$ during which interference was detected, consider every vehicle that was sighted in the monitoring zone at least at one point of this time interval. In practice, however it is challenging to determine this zone since it not only depends on wireless propagation effects but also the transmission power and antenna of a jammer.

Deriving Peak Detection Zone. To address this we use the concept of a peak detection zone. As an interfering vehicle approaches and leaves, the received signals at the monitoring point can be expected to rise and fall, respectively. By considering only the moment t_m at which the energy of interference signals peaks, we can define a smaller peak detection zone.

Theoretically the energy of interference signals reaches the maximum value when the transmitting vehicle is at the closest position to the detector. These closest positions are along a line perpendicular to the road and passing through the detector (red dotted line in Figure 4). Due to multipath such as reflections of interference signals by other vehicle bodies, there is however some uncertainty regarding the exact position of the vehicle when energy peaks. We therefore use a *peak detection zone* as the area that includes all those possible positions (i.e., the red trapezoid in Figure 4).

This discussion implicitly assumed omnidirectional transmitters. For a transmitter that is using a directional antenna, the peak may in theory not occur when the transmitting vehicle is closest to the detector but most commercially available jammer in the market have omnidirectional antennas and in practice, the tolerance zone can compensate many discrepancies.

Obtaining the boundary of the peak detection zone theoretically is complicated and thus we identify the zone empirically. We use the following steps: 1) For each experiment, we identify the position (p_m) of the transmitting vehicle at the moment t_m by examining the recorded video; 2) draw a ladder-shaped area to cover all the marked vehicle positions (p_m) across all the lanes. Figure 5 shows the identified

vehicle positions at the moment t_m for the 20 rounds of real-road experiments. Empirically, we determine that the peak detection zone for our receiver configuration can be enclosed by an trapezoid with a width of about 15 meters. Overall, the results are encouraging as it indicates it is possible to derive a peak detection zone that minimizes the number of candidate vehicles to uniquely identify the source vehicle.

4.2 Identification Using Mobile Detector

When multiple mobile devices (e.g., smartphones) on the road detecting the presence of the GPS jamming/interference, our system can track and predict the movement of the jammer on the road and coordinate with the roadside cameras to capture the candidate interference vehicle.

Determining the Moving Direction of Jamming /Interference Source. When the GPS jamming /interference is detected by a mobile detector (e.g., a smartphone), it is possible to determine the moving direction of the vehicle emanating interference by utilizing the *detection period*, which is the time period that the mobile detector finds low GPS SNR (i.e., below the detection threshold). The rationale behind this is if the mobile detector moves in the opposite direction as the jammer, the detection period will be very short, typically 1-2s. Whereas if the mobile detector moves towards the same direction as the jammer, the detection period will be longer than few seconds. Thus, by utilizing a time threshold we can determine the relative direction of the GPS jamming source.

Intuitively, the mobile detector moving in the opposite direction as the source vehicle is the most common case, since a vehicle has much more chances to encounter different vehicles in the opposite direction than in the same direction. Therefore, we focus on the cases when the moving direction of the source vehicle is opposite to that of the mobile detector. Since the source vehicle is moving along the road, it is reasonable that there are multiple mobile detectors sequentially detecting the presence of jamming at several different locations on the road. It is thus possible to determine the speeds of the source vehicle and its estimated locations, and further perform the path interpolation to estimate the time when the source vehicle passes the closest installed roadside camera.

To estimate the number of mobile detectors needed to guarantee at least one detector-jammer encounter, we built an emulation to analyze an actual dataset, called Next Generation SIMulation (NGSIM) [12]. The dataset contains actual trajectories of more than 2100 vehicles recorded by video cameras mounted on the side of an urban highway in North Hollywood, CA during a 15min rush hour period. The emulation randomly selects one vehicle to be a jammer and M vehicles to be mobile detectors. For each selection, the emulation looks at the vehicles' traveling timestamps and trajectories to detect if any mobile detector encounters the jammer. The emulator repeats 10,000 times for each value of M ranging between 1 and 10. The results shows that the jammer can be encountered 92% of the time with 10 mobile detectors. The number reduces to 91%, 83%, 74%, 67%, and 51% when the number of monitoring points reduces to 9, 5, 3, 2 and 1 respectively.

Estimating Recording Time Zone.

Next, we first present how to estimate the speed of the source vehicle based on two GPS signal SNRs observations from a single mobile detector at two time points. Then we discuss how to determine the recording time zone by

leveraging the estimated speed and location of the source vehicle.

Assuming that at time t_1 the mobile detector detects the presence of GPS jamming on a road with the opposite direction, the detector reports its speed v_m , the maximum SNRs of the GPS signal before the detection S^0 and after the detection S^1 to a centralized server. Based on S^0 and S^1 , the central server can calculate the noise power after the detection as $N_1 = S^0 - S^1 + N_0$. Then it can obtain the jamming signal's amplitude by subtracting the thermal noise from the noise after the detection, and further estimate the jamming signal power by taking the square of the jamming signal's amplitude using the equation:

$$\hat{J}_1 = 10 \log_{10} (\sqrt{10^{N_1/10}} - \sqrt{10^{N_0/10}})^2, \quad (1)$$

where N_0 can be calculated with room temperature and signal bandwidth [26]. Note that the $\hat{J}_1$ in Equation (1) is converted to decibel unit. Since most civilian GPS jammers have similar jamming power [2], the central server can estimate the relative distance d_1 between the jammer and the detector with the estimated jamming signal power $\hat{J}_1$ based on the free-space path loss model[22]. Similarly, the central server can estimate the relative distance $d_2 = d_1 - \Delta d$ at time $t_2 = t_1 + \Delta t$ by using the S^0 and another maximum SNR S^2 reported by the detector at time t_2 (within the detection period). The estimated speed of the source vehicle $\hat{v}_J$ can be calculated as following:

$$\hat{v}_J = (\Delta d)/(\Delta t) + v_m, \quad (2)$$

The estimated speed of the source vehicle keeps updating whenever there are new smartphones reporting the detection of the jamming/interference. Assuming the individual vehicle speeds follow a invariant normal distribution [18], the central server can fit the $\hat{v}_J$ in Equation (2) to a normal distribution $\mathcal{N}(\mu_v, \delta_v)$, where μ_v and δ_v are the mean and standard deviation of the speed, respectively. Since the central server can also estimate/track the locations of the source vehicle by utilizing the last-fixed locations reported by mobile detectors on the road, when the central server finds that the source vehicle will pass a roadside camera that is at a distance of d_c away, it can estimate the traveling time of the vehicle to reach the camera as:

$$\hat{t}_c = (d_c)/(\mu_v), \quad (3)$$

and start recording on the camera around $\hat{t}_c$ to capture the candidate vehicles. We define the *recording time zone* as $\{\hat{t}_c - \theta_1, \hat{t}_c + \theta_2\}$, where θ_1 and θ_2 are adjustable parameters capturing the estimation errors, which can be calculated as following in our work:

$$\begin{aligned} \theta_1 &= (d_c \cdot 3\delta_v)/(\mu_v(\mu_v + 3\delta_v)), \\ \theta_2 &= (d_c \cdot 3\delta_v)/(\mu_v(\mu_v - 3\delta_v)). \end{aligned} \quad (4)$$

To get an understanding of the relationship between θ and the distance from the interference observation, we analyze the difference in travel time of vehicles on a common road segment using the NGSIM dataset mentioned above. The emulation randomly selects 2000 pairs of intervals that are 100 to 500m long on this road segment. It then calculates the amount of time it takes for each vehicle to travel through this interval. Figures 13 shows the distribution of travel time for each interval length, in 100 meter increments. The spread in travel times is directly related to theta. For example, the results show vehicles took between 25 to 55s to travel the 500 meters intervals. This means that if the camera is 500 meters away from the mobile detection, an appropriate size for θ_1 and θ_2 would be 15s, which together results in the 30s spread in travel times observed.




Figure 6: The wired experiment circuit.

4.3 Monitoring Point Vehicle Identification

The formulation of the peak detection zone and recording time zone equip our system to derive a new approach that can automatically identify the transmitting vehicle by exploiting the camera visualization. When a single monitoring point with a roadside detector and a monitoring camera is available, the moment t_m when the detector finds the interference signals having the maximum value is utilized to check the video captured by the monitoring camera. And the peak detection zone is applied to capture all the candidate vehicles. We show in the next Section that we find over 95% of our experiments in real-road scenarios capture only one candidate vehicle inside the peak detection zone.

In low traffic conditions, using single monitoring point can identify transmitting vehicles with high confidence, because in most of the cases there is only one transmitting vehicle in both the peak detection zone and the recording time zone. However, if the traffic is heavy, multiple candidate vehicles could be captured by either the peak detection zone or the recording time zone. By using our low-cost solution, our system can employ multiple monitoring points separated along the major roads having high traffic volume to still uniquely identify the transmitting vehicle. The intuition behind this is that with two or more monitoring points along the same direction of the major roads, it is less likely for a transmitting vehicle to have the exact same neighboring vehicles inside of either the peak detection zone or the recording time zone when crossing different monitoring points.

Thus, the multiple monitoring points identification algorithm is separated into 2 steps:

Step 1: Collect Images of Candidate Vehicles from Monitoring Points. Assume the transmitting vehicle passing K monitoring points on the road, the cameras at these monitoring points are all triggered to record the videos when the transmitting vehicle passing by. Therefore, each monitoring point generates the images of candidate vehicles (including the transmitting vehicle) from the recorded videos, denoted as $R_k = \{r_1^k, \dots, r_m^k\}$, where $k = 1, \dots, K$ is the index of monitoring points, r_m^k is the images of m^{th} candidate vehicles captured at the k^{th} monitoring point. All R_k are transmitted to a central server. We note that this step is also valid if the transmitting vehicle passing the same monitoring point for K times in one day or different days, similarly, R_k from the same monitoring point will be transmitted to the central server.

Step 2: Cross Validate Candidate Vehicles. Once the central server receives the sets of vehicle's image R_k from all monitoring point cameras, it finds the common vehicle C that appears in all R_k by calculating the $C = R_1 \cap R_2 \cap \dots \cap R_K$.

5. PERFORMANCE EVALUATION

In this section, we evaluate the two most important components in the designed system: Carrier-to-Noise (C/N_0)




Figure 7: Result of cable experiments where (a) shows C/N_0 levels of satellites for different interference power levels, and (b) shows C/N_0 levels of satellites when GPS receiver lost location fix.

profiling based detection in both static and mobile cases, and source vehicle identification leveraging energy detection through roadside monitoring. Such performance evaluation provides fundamental insights to a complete system implementation for GPS interference detection.

5.1 Stationary Carrier-to-Noise Profiling and Interference Tests

To gain a deeper understanding of the effect of an jamming signal on SNR levels of GPS receivers, we conducted the following experiments. To comply with federal regulations and to minimize possible interference to GPS systems, we created a wired testbed depicted in Figure 6. In this setup, a jamming signal created by a USRP N210 software defined radio is combined with GPS satellite signals from an external antenna and fed into a u-blox EVK-7N GPS receiver [27]. We used an isolator to ensure that the jamming signal is not transmitted by the GPS antenna. In addition, we utilized a high gain (40dB) GPS antenna with attenuators to adjust the SNR of the received GPS signals. Using this setup, we first analyze the SNR patterns of different satellites with a 10 hour experiment. In this test, we did not use any attenuation to be able to observe the natural SNR behaviour of satellites. The result of our detection algorithm is depicted in Table 1 as the stationary case. Then by using this setup, we analyze the impacts of interference power on different satellites. As shown in Figure 7(a), increasing interference power results in a linear decrease of SNR with increased noise. This result verifies the expected effect of interference on the SNR readings.

We also observe that when a GPS receiver loses a location fix due to high-energy interference it sometimes still can receive data from some satellites, at least for part of the time. We show the SNR level of several satellites over time at an interference level of -45dBm in Figure 7(b) where the GPS receiver has lost the location fix. We can observe high variance in several satellites under this high level of interference, an effect that we have not observed under any of our non-interference tests—both in this wired setup and the vehicle experiments described later. This pattern can therefore be a further indicator for interference.

5.2 Mobile Detectors for Carrier-to-Noise Profiling

We study the performance of mobile detectors leveraging mobile devices (e.g., smartphone and GPS receiver) placed inside of vehicles. Specifically, we examine the detection rate and false positive rate under various interference powers and the behavior of SNR changes across different mobile devices over repeated road trips.




Figure 8: C/N_0 profile construction: (a) shows average C/N_0 for different trips for a route and (b) shows the constructed C/N_0 profile the route.

5.2.1 Experimental Setup

Hardware and Signals. We conduct our experiments in two routes in two US cities. In the first city, we use 1 HTC Evo 4G phone, 1 LG Nexus 4 phone, and 1 u-blox EVK-7N GPS receiver. We use 1 LG Nexus 4 phone in the second city. During the experiments, smartphones are located on the front panel of the car and GPS antenna for u-blox receiver is placed on top of the car.

Real Road Scenarios. We chose a suburban highway for route 1 with a distance of 12 km which includes one underpass and a city road for route 2 with a distance of 6 km. We drove our vehicle in regular traffic at different times of a day for 20 different days over 4 weeks and use the smartphones and the GPS receiver to record the GPS signal for all trips. We split the distance into 20 meters long road intervals which is relatively small and is long enough to get several readings from receivers.

Metrics. In this work, we utilize the mean value of 4 largest SNRs from all satellites, since GPS chips utilize the 4 strongest satellite signals to determine 3-D positions. In the rest of the paper, SNR level corresponds to such a mean value. We consider the case, where the SNR level reported by a receiver is less than the threshold, as a detection when the GPS jamming/interference is present. Based on this, we evaluate the performance of the detection algorithms in terms of following metrics: 1) detection rate (DR) is the percentage of distance intervals in which the detection is successful. For instance, in our route 1 experiment, there are 610 different road intervals (i.e $610 * 20 \text{ meters} = 12.2 \text{ km}$), and if the detection is successful in x of them, then the detection rate becomes $(x * 100 / 610) \%$. 2) False positive is the case when our algorithm determines there is a suspicious GPS interference activity although it is not coming from a jammer. We note that the detectable interference power is defined as the level where detection rate exceeds 95%.

5.2.2 Experimental Results

Feasibility Study of Profile Construction. Figure 8 shows the average C/N_0 over 4 days (each trip represents one day) versus the relative distances from the starting points of the trips in two cities. We observe that the C/N_0 levels are consistent at same locations across different days. For instance, there is always a big drop in C/N_0 at about 6 km in Figure 8(a) which is caused by an underpass during the road trip. We note that the C/N_0 level in smartphones fluctuate more than that in the GPS receiver.

To further illustrate the basic idea, a C/N_0 profile could be constructed by combining the C/N_0 values from all trips across different days into each corresponding distance interval. The constructed C/N_0 profile of the testing route is shown in Figure 8(b).




Figure 9: Estimated detection rate of our detection algorithm for different interference powers.

Detection Performance. As we discussed in Section 3, we can determine whether there are interference signals or not by comparing the newly reported SNR to a fixed threshold. We present the detection rate of our algorithm under different interference power by varying the threshold levels in Figure 9. In our profiling algorithm as presented in Section 3, the threshold is defined as $\delta = m_{snr} - k * \sigma_{snr}$. We observe the detection rate increases with the increase of received interference power. And the best detection rate is achieved by using $k=2$. We further present the false positive rate of our algorithm and detectable interference power(DIP) under different thresholds in Table 1. We find that low false positive rates are achieved under all cases, indicating the effectiveness of our approach.

The detection range is dependent on signal propagation conditions. We therefore first evaluate detection rate based on the actual received interference power. We further study how to decide on the distance of the jammer. From detectable interference power results in Table 1, we know that if the received power is -109 dBm, with 95 percentile possibility we can detect the jammer. By using the free space path loss(FSPL) model with two propagation exponents 3 and 3.5, we observe the jamming activity could be detected under the distance of 77 and 14 meters respectively.

5.3 Roadside Monitoring for Source Vehicle Identification

We next study the accuracy of source vehicle identification through roadside monitoring. To perform a representative case study, in place of an actual jammer, we use a transmitter with similar characteristics in the closest ISM band at 900MHz. This is to minimize the chance of causing any GPS interference. The shift to the 900MHz band primarily means somewhat better material penetration and higher antenna effectiveness. This means that the range in the L1 band will be somewhat shorter, but we believe these 900MHz results are sufficiently close to be useful guidance.

5.3.1 Experimental Setup

Hardware and Signals. As described in section 3, we conducted our experiments using USRP N210 at 915MHz band to implement the energy-based interference detector at each monitoring point. The USRPs were connected to laptops streaming data to or from the host processors. A

k	DIP u-blox stationary	DIP u-blox mobile	DIP phone mobile	False Positive
2	-108 dBm	-103 dBm	-109 dBm	5%
3	-105 dBm	-86 dBm	-91 dBm	1%
4	-101 dBm	-55 dBm	-69 dBm	0.1%

Table 1: Detectable Interference Power and False Positive Rate for different thresholds




Figure 10: Spectrogram of a jamming signal captured by the jammer detector from the simulated jammer transmitting jamming signal at 915 MHz.

video camera was setup with the USRP at the monitoring point to record videos during the experiments with its time synchronized with the laptop. Specifically, the camera is on an adjacent overpass at each monitoring point, and angled to capture the video of traffic on all lanes within the transmission range of the detector.

Real Road Scenarios. We chose a suburban highway as a representative route, and placed two monitoring points 3 miles apart. The locations were chosen to reflect different lane/road configurations and based on the availability of overpasses for camera placement as well as the availability of safe roadside positions (a parking lot and bus stop). We drove the vehicle carrying the transmitter on the right-most and left-most lanes, and passed by the monitoring point for 20 times, respectively. The experiments were across different days with regular traffic.

Metrics. We consider the number of candidate vehicles that are in the detection area(s) at the time of detection. If the number is one, and this is our transmitter vehicle, we consider a detection successful. Based on this, we evaluate the performance of our detection and identification approach in terms of the following metrics: 1) detection rate (DR): the percentage of passes during which our transmitting vehicle was correctly identified by our system, and 2) false detection: when a vehicle that likely does not emit interference signals is implicated by our system.

5.3.2 Experimental Results

Validation of the Emulated GPS Jammer To validate this detection approach and ensure a sufficient detection range as well as robustness to high speeds we conducted the following experiment. In addition to the roadside detector we used a second USRP with a Matlab generated (Figure 2) interference signal to emulate a jammer. Both were equipped with one omni-directional antenna [9] (824-960 MHz).

Given that the simulated GPS jammer jams at the power of -15 dBm, we were able to detect it 50 meters away, indicating a good chance of detecting GPS jammers from a roadside location. The spectrograms of the received signal at different speeds (Figure 10) shows that the chirp shape of the jamming signal remains visible even if the speed is 50 MPH and the received signal strength remains similar at different speeds.

Single Monitoring Point Case. We first consider a single monitoring zone. Provided with the video images from the time of signal detection, we manually identified the number of vehicles inside the detection zone and whether our transmitter vehicle is among them. As shown in Fig. 11, at monitoring point 1 there is only a single vehicle in the monitoring zone in 13 out of the 20 passes and 2 to 3 vehicles




Figure 11: Number of candidate cars at a monitoring point when jamming activity is detected.

are present in the remaining 7 trials. Our transmitter vehicle was present during all 20 trials, this indicates our system did not have any false positives and were able to uniquely identify the transmitting vehicle in 65% of the cases. When only considering monitoring zone 2, a single vehicle was present in 7 out of the 20 passes with 2 to 4 vehicles in the remaining ones, yielding no false positives and a lower detection rate of 35%. This is likely due to the larger number of lanes at this location.

Multi-Monitoring Point Case. The detection rate can be significantly improved by combining the information collected from both zones. In particular, in the 16 cases where more than one car was inside the detection zone at either monitoring point 1 or 2, we determined the intersection of the candidate vehicle set from zone 1 and 2, as per the cross validation method when multiple detection zones are present. This resulted in only a single detected vehicle and a correct identification of our transmitter vehicle in all cases.

To investigate these aspects further, we determined the time difference that has accumulated between cars that passed monitoring zone 1 together when they arrived at monitoring zone 2. Figure 12 (a) presents a histogram of this time difference. Positive timing means it takes more time for a regular vehicle to travel from monitoring point 1 to 2 than the transmitting vehicle and vice versa. This time difference allows us to judge how large detection zones can become before the detection rate would fall. Our current detection areas could be traversed in about 1-2 seconds at typical speeds. The histogram shows that only a detection zone taking more than 4 seconds to traverse (approximately 100m at 25m/s) would start to yield failed detections.

Detection Zone Size and Detection Rate. The relationship between detection zone size and detection rate is further examined in Figure 12 (b). Here the detection time interval represents a detection zone in the time domain (for example, a 50m zone is approximately equivalent to 2s at 25m/s). When this interval is too short, the detection zone may not include the transmitting vehicle and detection rate is low. If it is too long, many other vehicles are also in the detection area, which means unique vehicle identification is not possible and detection rate declines again. In our roadway experiments, detection rate was optimized with a 3-4 second interval, achieving a detection rate of 100%.

Thus, to determine the detection zone size, we drove the transmitting vehicle and passed by each monitoring point 20 times as mentioned in section 4. In each pass, when the detector is triggered by the transmitting signal, our system recorded the position of the transmitting car and marked that position in a screenshot obtained from the camera as can be seen in Figure 14. These positions let us empirically create a detection zone which covers all of these positions.




Figure 12: (a) Time difference between the jamming vehicle and other suspicious vehicles that are in the detection areas at both monitoring points, (b) The change of detection rate with the detection zone size.

Evolution of the Signal Through One Pass. When we were passing by the monitoring point, the detector detects several activities (10 ms long). In Figure 15(a), we present how the normalized amplitude of the signal is changing over time for three passes as an illustration. Finally, Figure 15(b) presents the received power for each detected signal with the duration of that signal together with a detection threshold at monitoring point 1. It is clear that in the false detection cases, the durations of the incident and the received power are much less than those in the correct detection cases. Similar observations are obtained from monitoring point 2. Our detection algorithm gives perfect results (precision 1) for both monitoring points.

5.3.3 Discussion

In our experiments, the distance between the monitoring points is about 3 miles and there is no traffic light in between. If the distance between the monitoring points is shorter, then we expect to see more cars to travel together and they would be seen in the monitoring points together. If there are some traffic lights between the monitoring points, it would force the drivers to move in the same groups and the number of common cars in both monitoring points is expected to be more than those in our case. Also driving styles can effect the results.

In the evaluation, we have not compared with competing solutions, because we are not aware of an approach that can isolate mobile GPS jammers with the level of resources that were at our disposal (we did not have access to the FAA/FCC teams and contractors described in Newark Airport incident).

5.3.4 Passive Roadside Monitoring

We deployed our detection system in several locations in two US cities for monitoring GPS L1 frequency band (1575.42 MHz) passively on the roadside. We mainly use a interference/jamming detector equipped with a multi-band GPS antennas [17] to continuously monitor the GPS L1 frequency (1575.42 MHz) band. The detector is tuned to log any suspicious signal at the L1 band when its power exceeds the pre-set threshold, which was -60dBm. The suspicious signals are studied through offline analysis.

To perform extensive roadside monitoring, we want to identify the interfering GPS signals on public roadways. With this purpose, we have selected several locations on three kinds of roads: a major highway, one of the busiest toll road, and an urban road for passive monitoring of GPS frequency band. In total we monitored the roadways passively for approximately 200 hours in the 5 locations in NJ and South Carolina.




Figure 13: (a) Spectrogram of no jamming activity at location 1, (b) suspicious interference activities detected at location 2, (c) Travel time of vehicles for different interval length, (d) HTC Evo phone for route 1.




Figure 14: Positions of the transmitter vehicle at the time of maximum detection in monitoring point 1.

In the 200 hours monitoring at five monitoring points, we detected two suspicious interference incident at the location 1 (on a busy toll road) and the location 2 (on an interstate highway) respectively. As shown in Figure 13, there was high energy at the L1 frequency band in both of the two incidents.

We also emphasize that although, it is not clear if the signals detected in passive monitoring are coming from a real jammer, these results show that significant interference events happen in the L1 frequency band and even non-jamming interference can cause problems.

6. DISCUSSION

System limitations for future jammers. Our goal is to design a low-cost detection mechanism that can be mass deployed in roadways, and thus our jamming detection mechanisms are designed to identify existing off-the-shelf GPS jammers that continuously emit interference. Future jammers that may emit the interference for a short period of time could still be caught by the mobile detectors (e.g smart phones) since they could still be reporting any outages or SNR anomalies even the strength of the jamming signals change over time. With such reports, the path of the jammer car can be interpolated to cameras for identification. The probability that jammer would be detected by a stationary monitoring point on a single pass is smaller but after a sufficient number of passes (say, for a commuter), detection would still be likely.

Opportunistic Crowdsourcing. While not every individual user will be motivated to install our mobile detector app for jammer detection, we envision that this mechanism can still be deployed in police vehicles, public transit systems, delivery truck fleets, etc. It could also become part of future connected vehicle standards or be integrated in smartphone platforms with careful energy management.

GPS Spoofing. Apart from GPS jamming, another security issue is GPS spoofing attacks, whereby attackers transmit fake GPS signals to fool the GPS receivers[20]. Note that GPS spoofing is outside the scope of this paper and we focus on designing a practical solution for detecting GPS jammers. Nevertheless, our method is complementary to GPS spoofing detection strategies.




Figure 15: (a) shows the change in received power over time as the transmitting vehicle passing by the monitoring point 1 and (b) shows correct detections and false detections at the same monitoring point.

7. RELATED WORK

There are quite a few works analyzing the characteristics of the GPS jammers available in the market and utilizing multiple static stations to detect and localize the jammer. Bauernfeind et al. [2] proposes to mitigate the in-car jammers by applying an inverse signal transformation [21]. Mitch et al. [16] propose a time difference of arrival (TDOA) based method that can localize the GPS jammer with maximum error of 15 meters, which however requires the prior knowledge of the jamming signal. Lindstrom et al. [14] builds a network of low cost application-specific integrated circuit (ASIC) front-end modules providing automatic gain control (AGC) values to detect and localize the interference sources in GNSS L1/E1 band with an error of 21 meters on average. Similarly, Poncelet et al. [23, 3] can detect and localize wide-band GNSS jammers based on TDOA measurements with the accuracy of 20 meters.

Vehicular ad hoc networks are also being employed for identifying and localizing GPS jammers. Bauernfeind et al. [1] uses a Differences-of-Received-Signal-Strength (DRSS)-based localization algorithm which requires at least signal measurements from four vehicles for localization. Fontanella et al.[10] further experimentally confirms the potential for using VENETs to locate in-car jammers. However, the assumption of having vehicles all around the jammer is not generally true and the estimation error is approximately 40 meters which is too high for automatic vehicle identification. In addition, Kramer et al. [13] recently proposes an Android based approach which relies on relative position change between a smartphone user and a static GPS jammer to perform localization. Scott et al. [25] proposes a crowd-sourcing approach to detect and localize the GPS jammer by exploiting the AGC values from the GPS receiver, which requires to incorporate the GPS jam-to-noise (J/N) ratio detector in cell phones.

Chronos Technology has two commercial devices [5][6] for detecting and localizing GPS jammers in the L1 frequency band. However, these devices can only detect jammer in static scenarios and can only be used manually for a short period of time. On the contrary, our solution is automatic, can be deployed at roadsides and can detect jammer in dynamic mobile scenarios.

8. CONCLUSION

Several incidents caused by in-vehicle jamming devices have illustrated their serious impact on the availability of the navigation services in critical infrastructure. To monitor and reduce the use of in-vehicle GPS jammers, we presented a low-cost jammer identification system that can be mass deployed in roadways to automatically detect and identify the vehicles with GPS jammers. Instead of covering all types of possible jammers, our jamming detection mechanisms focuses on achieving wide geographical coverage against the most prevalent off-the-shelf GPS jammers, which continuously emit interference. The dedicated detectors, however, could be extended for more sophisticated jamming signals. The key components of the system are monitoring stations (which are equipped with directional antennas and cameras) and mobile detectors (e.g., smartphones). Using an off-the-shelf software-defined radio (USRP) to emulate GPS jamming signals, we conducted a case study evaluation of our system with multiple trial drives on local highways in 2 US cities and found the monitoring stations effective.

We also demonstrated that by constructing location-based profiles of expected signals, our mobile detector component can detect interfering signals based on measurements that are readily available in most GPS receivers, including the ones in smartphones. Thus, it is possible to detect jammers via crowdsourcing. While not every individual user will be motivated to install our mobile detector app, we envision that this mechanism can be deployed in law enforcement vehicles, public transit systems, etc. End users may also volunteer to report data while using GPS navigation services, and thus provide opportunistic crowdsourced data. The positive results that are obtained when validating the most important components of our solution, i.e., the Energy-based monitoring point and the SNR-based mobile detector, are promising. They indicate the validity of a full-fledged system that consists of a centralized vehicle identification component integrated with monitoring points and mobile detectors.

9. ACKNOWLEDGEMENT

This material is based in part upon work supported by the National Science Foundation under Grant Nos. CNS-0845896, CNS-1329939, CNS-0954020 and by the Army Research Office under Grant No. W911NF-13-1-0288.

10. REFERENCES

- [1] R. Bauernfeind et al. In-car jammer interference detection in automotive gnss receivers and localization by means of vehicular communication. In *FISTS, 2011 IEEE Forum on*, pages 376–381. IEEE, 2011.
- [2] R. Bauernfeind et al. *Analysis, Detection and Mitigation of InCar GNSS Jammer Interference in Intelligent Transport Systems*. Deutsche Gesellschaft für Luft-und Raumfahrt-Lilienthal-Oberth eV, 2013.
- [3] J. Bhatti et al. Development and demonstration of a tdoa-based gnss interference signal localization system. In *PLANS, 2012 IEEE/ION*, pages 455–469. IEEE, 2012.
- [4] A. Chhabra. *Testimony of Ashwini Chhabra, Deputy Commissioner, Policy & Planning*. NYC Taxi and Limousine Commission, 2012.
- [5] Chronos Technology. *CTL3510 GPS Jammer Detector*. http://www.navtechgps.com/assets/1/7/CTL3510_DS.pdf.
- [6] Chronos Technology. *CTL3520 GPS Jammer Detector & Locator*. <https://www.ettus.com>.
- [7] Chronos Technology Ltd. *SENTINEL: GNSS Services Needing Trust In Navigation, Electronics, Location*. <http://www.gps-world.biz/index.php/sentinel-1157>.
- [8] Ettus Research. *USRP N210*. <https://www.ettus.com>.
- [9] Ettus Research. *VERT900*. <https://www.ettus.com>.
- [10] D. Fontanella et al. In-car gnss jammer localization using vehicular ad-hoc networks. 5/6 2013.
- [11] GNU Radio. <http://www.gnuradio.org>.
- [12] <http://ngsim-community.org/>. *Next Generation SIMulation Community*, 2013.
- [13] I. Kramer et al. Android gps jammer localizer application based on c/n0 measurements and pedestrian dead reckoning. In *ION-GNSS*, 2012.
- [14] J. Lindstrom et al. Gnss interference detection and localization using a network of low-cost front-end modules. In *ION GNSS*, 2007.
- [15] Merrill, John. *Patriot Watch Vigilance Safeguarding America*. <http://www.gps.gov/multimedia/presentations/2012/03/WSTS/merrill.pdf>.
- [16] R. Mitch et al. Civilian gps jammer signal tracking and geolocation. In *ION GNSS*, 2012.
- [17] MobileMark Antenna Solutions. *GPS Multiband Antenna*.
- [18] M. A. Mutaz et al. Leveraging platoon dispersion for sybil detection in vehicular networks. In *PST*. IEEE, 2013.
- [19] National Transportation Library. *Transit-Management Systems*. <http://ntl.bts.gov/lib/jpodocs/edldocs1/13480/ch6.pdf>.
- [20] T. Nighswander et al. Gps software attacks. In *CCS*. ACM, 2012.
- [21] X. Ouyang et al. Short-time fourier transform receiver for nonstationary interference excision in direct sequence spread spectrum communications. *Signal Processing, IEEE Transactions on*, 49(4), 2001.
- [22] J. D. Parsons and P. J. D. Parsons. *The mobile radio propagation channel*. J. Wiley, 2000.
- [23] J.-P. Poncelet et al. A low-cost monitoring station for detection & localization of interference in gps l1 band. In *NAVITEC*. IEEE, 2012.
- [24] G. D. Rash. Gps jamming in a laboratory environment. *NAWCWPNS*, pages 1–20, 1997.
- [25] L. Scott. J 911: Fast jammer detection and location using cell-phone crowd-sourcings. *GPS World*, 2010.
- [26] E. N. Skomal et al. *Measuring the radio frequency environment*. Van Nostrand Reinhold Company, 1985.
- [27] U-Blox. *EVK-7N*. <https://www.u-blox.com/en/evaluation-tools-a-software/gps-evaluation-kits.html>.
- [28] C. VSCC. Vehicle safety communications project: task 3 final report: identify intelligent vehicle safety applications enabled by dscc. 2005.