

# Finding Phones Fast: Low-Latency and Scalable Monitoring of Cellular Communications in Sensitive Areas

Martin Kotuliak  
ETH Zurich  
Zurich, Switzerland

Simon Erni  
ETH Zurich  
Zurich, Switzerland

Jakub Polák  
Unaffiliated  
Zurich, Switzerland

Marc Röschlin  
Unaffiliated  
Zurich, Switzerland

Richard Baker  
University of Oxford  
Oxford, United Kingdom

Ivan Martinovic  
University of Oxford  
Oxford, United Kingdom

Srdjan Čapkun  
ETH Zurich  
Zurich, Switzerland

## Abstract

The widespread availability of cellular devices introduces new threat vectors that allow users or attackers to bypass security policies and physical barriers and bring unauthorized devices into sensitive areas. These threats can arise from user non-compliance or deliberate actions aimed at data exfiltration/infiltration via hidden devices, drones, etc. We identify a critical gap in this context: the absence of low-latency systems for high-quality and instantaneous monitoring of cellular transmissions. Such low-latency systems are crucial to allow for timely detection, decision (e.g., geofencing or localization), and disruption of unauthorized communication in sensitive areas. Operator-based monitoring systems, built for purposes such as people counting or tracking, lack real-time capability, require cooperation across multiple operators, and thus are hard to deploy. Operator-independent monitoring approaches proposed in the literature either lack low-latency capabilities or do not scale.

We propose WaveTag, the first low-latency, operator-independent and scalable system designed to monitor 5G and LTE connections across all operators prior to any user data transmission. WaveTag consists of several downlink sniffers and a distributed network of uplink sniffers that measure both downlink protocol information and uplink signal characteristics at multiple locations to gain a detailed spatial image of uplink signals. WaveTag then aggregates the recorded information, processes it, and provides a decision about the connection—all done prior to the complete connection establishment of a UE. To evaluate WaveTag, we deployed it in the context of geofencing, where WaveTag was able to determine whether the signals originate from inside or outside of an area within 2.3 ms of the initial base station-to-device message, therefore enabling prompt and targeted suppression of communication before any user data was transmitted. WaveTag achieved 99.66% geofencing classification accuracy, using only the information collected before connection establishment. Finally, we conduct a real-world uplink measurement evaluation on a commercial 5G SA network.

## 1 Introduction

In certain scenarios, cellular communication is undesirable or forbidden, demanding strict monitoring and rapid response. Examples include exam halls, where phones are used to cheat [9, 10, 15]; correctional facilities, where smuggled phones threaten security [12,

13, 16, 38]; protected airspace, where a drone controlled over a cellular connection [11, 21] can wreak havoc; or confidential voting processes, such as a jury deliberation room [1, 42] or even the Vatican conclave [7], where unauthorized communication could influence the outcomes. Consequently, there is a strong demand for monitoring solutions that can provide detailed real-time information about cellular devices within a sensitive area, fast enough to prevent unauthorized communication. Low latency is crucial, as targeted blocking of a malicious connection requires a decision from the monitoring system before the connection is fully established and user data is transmitted or received [19, 22, 23].

In addition to being low-latency, such monitoring solutions also need to be scalable (i.e., handle a large number of operators, cells, and UEs in parallel) and integrate with geofencing or localization. This can be challenging, particularly because UEs transmit only a limited number of messages during connection establishment and typically operate in complex, non-line-of-sight environments.

Cellular operators could provide such monitoring solutions [20, 34]. However, several reasons make this impractical: (1) Accuracy and Cost: Under the current LTE specification [3, 6], operators cannot monitor a user’s connection with the precision required in these situations [20, 33], especially in NLoS scenarios involving many indoor locations. Improving precision around sensitive areas would require deploying a much denser base-station infrastructure than exists today. (2) Latency: No 3GPP-specified solution provides a geofencing result based solely on initial uplink messages; localization happens later in the connection. Without a specification change, an operator aiming to prevent communication from sensitive areas would need to block all communication until geofencing results arrive. This would significantly degrade user experience for everyone near sensitive areas, who would have to wait before making calls or sending messages—an unacceptable impact in urban environments. (3) Feasibility: Every operator in the country would need to deploy this solution and collaborate with sensitive-area administrators. If even one operator refuses or is unable to cooperate, an adversary could bypass the system through that operator. Such nationwide coordination seems infeasible in practice.

On the other hand, state-of-the-art research [30, 31, 35] explores operator-independent monitoring but faces practical limitations. (1) None shows to be low-latency when deployed in a distributed

system. (2) They are restricted to a single cell per SDR, whereas realistic deployments involve multiple cells across different operators. (3) They combine downlink and uplink receivers into a single device, limiting their use to locations with good downlink visibility. These challenges make it difficult to collect high-quality signal metrics and integrate them with geofencing systems. (4) They do not support 5G SA.


In this paper, we introduce WaveTag, a low-latency 5G and LTE monitoring system independent of cellular operators that can provide information about a cellular connection no later than 2.3 ms after the initial Random Access Response message. WaveTag works in parallel on any frequency band across all operators and cells. The signal characteristics of every uplink message are measured by a distributed set of uplink receivers both inside and around the sensitive area, and the measurements are aggregated in a central unit, which can run further analysis steps. Due to its low latency, WaveTag provides enough time for other systems to react before any user data reaches the network.

To showcase the capabilities of WaveTag, we deployed it in the context of geofencing that combines metrics using a robust machine learning model. After aggregating the signal characteristics, the model classifies the communications as originating from inside or outside of a predefined restricted area. We deployed and tested our system in two different building complexes in urban settings where mobile phones are not allowed. In the two deployments, we installed six and ten uplink receivers, respectively. We also developed a platform with multiple UEs continuously re-connecting to the cellular network, in order to obtain a large dataset of connection attempts from inside and outside the building complex. Using WaveTag, we achieved 99.7% geofencing accuracy.

Finally, we evaluated WaveTag on a real-world 5G SA network. First, we validated that WaveTag can process multiple 5G cells on the same frequency, as well as measure signal characteristics from UEs connecting to a 100 MHz 5G TDD cell. Then, we evaluated WaveTag on a large open field against a commercial smartphone. This evaluation showed that existing LTE features are still valid in the context of 5G SA, and there is only marginal time overhead in the collection of 5G SA signal characteristics, compared to LTE.

In summary, we make the following contributions:

- We design and implement WaveTag, the first operator-independent, low-latency 5G and LTE monitoring system capable of making a decision about a connection within 2.3 milliseconds of the initial base station-to-device message, before any user data transmission occurs.
- We develop a scalable architecture consisting of one or more downlink receivers and a distributed set of uplink receivers, enabling precise spatial geofencing or localization even in complex, non-line-of-sight environments.
- We present an end-to-end prototype and full-system evaluation, validating the feasibility, scalability, and performance of WaveTag in real-world scenarios with more than 100 base stations across four operators and all frequency bands.
- We perform scalability, latency, and feature extraction tests in a commercial 5G SA network.
- We build a new data collection platform for collecting connection data based on COTS modems and collect a huge



**Figure 1: Diagram of a connection establishment procedure between a UE and an xNB.**

dataset of >170,000 connections with all the signal features for the initial UL messages. We make this dataset public.

- We integrate our system with a machine learning-based geofencing mechanism that achieves 99.7% classification accuracy using only uplink messages before a connection with the network core is established, enabling early and targeted response to unauthorized communication attempts.

## 2 Background

In 5G/LTE cellular networks, users communicate using user equipment (UE) over a wireless channel with base stations (gNodeB for 5G and eNodeB for LTE, we use the umbrella term xNB for both of them). Communication from the xNB towards the UE is the downlink (DL), and the opposite direction is the uplink (UL). Downlink and uplink are separated by frequency—Frequency Division Duplex (FDD)—or time—Time Division Duplex (TDD).

The International Telecommunication Union (ITU) defines allocations of radio spectrum for uplink and downlink communication across a number of frequency bands. In the country of our deployments, five FDD bands are allocated for 5G or LTE [14] and one extra TDD band for 5G. Operators tend to have frequencies allocated in most of the available bands and set up cells in all of their allocated frequencies. Therefore, any monitoring system must cover all frequency bands allocated to operators simultaneously.

### 2.1 Connection Establishment

Figure 1 shows the initial message exchange between a UE and an xNB during the connection establishment procedure. Initially, the UE scans the spectrum for available xNBs and acquires their configuration broadcast in system information blocks (SIBs). After selecting a suitable cell, e.g., depending on signal conditions and/or operator preferences, the radio connection is initiated with the UE transmitting a Random Access message on the Physical Random Access Channel (PRACH). Upon detecting this message,

the xNB responds with a Random Access Response (RAR) message on the Physical Downlink Shared Channel (PDSCH) containing the original PRACH preamble identifier, a unique connection identifier (RNTI), a Timing Advance (TA) command, and a resource allocation specifying time and frequency parameters for the next uplink message of the UE.

This next uplink message, transmitted on the Physical Uplink Shared Channel (PUSCH), is a Radio Resource Control (RRC) Connection Request (RRC Setup Request in 5G). In response to the request, the xNB sends an RRC Connection Setup (RRC Setup in 5G) message on the PDSCH containing dedicated connection configuration parameters for the UE. The UE acknowledges this message on the Physical Uplink Control Channel (PUCCH), and once the xNB sends another uplink resource allocation on the Physical Downlink Control Channel (PDCCH), the UE transmits the RRC Connection Setup Complete and Non-Access Stratum (NAS) Service/Attach Request messages on the PUSCH. Notably, the NAS Service/Attach Request is the first message sent from the UE to the network core, as all prior messages are exchanged solely between the UE and the xNB. We consider the connection to be re-established once the NAS Service Request is delivered, since the following messages are sent encrypted, integrity protected, and can contain user data.


## 2.2 Physical Uplink Channels

On the physical layer, a UE sends a message in a resource grid consisting of OFDM symbols with multiple subcarriers in each symbol. Subcarrier spacing is 15kHz for all the monitored FDD bands, and 30kHz for the 5G TDD band. For 15kHz subcarrier spacing, 14 OFDM symbols form a 1 millisecond long subframe, and 10 subframes form one frame. For 30kHz subcarrier spacing, one subframe consists of 28 OFDM symbols instead. One physical resource block (PRB) spans 12 subcarriers over 7 OFDM symbols.

In the uplink, there are three types of physical layer channels: the PUSCH used for transmitting user or control data, the PUCCH used for control information such as ACK/NACK or scheduling requests, and the PRACH used for the initial connection message. All uplink messages in these channels carry a known reference signal alongside the modulated data, which is used for channel estimation, or PRACH detection. Depending on the message, the uplink reference signal spans different OFDM symbols in the resource grid. The reference signal is a Zadoff-Chu sequence, a sequence known for its good autocorrelation properties and constant-amplitude signal. Figure 2 illustrates these three types of messages in the spectrum plot and highlights the reference symbols.

Depending on the configuration inside the SIB, a PRACH preamble consists of one or more 839- or 139-symbol-long Zadoff-Chu sequences spanning six PRBs. The UE randomly chooses one of the pre-configured Zadoff-Chu sequences and transmits the preamble in one of the pre-configured PRACH occasions. Upon detection, the xNB replies with a random access response, which contains the preamble index of the chosen sequence. The UE continues the connection establishment only if the preamble index in RAR matches the transmitted PRACH preamble index.

The UE sends a PUSCH message only after it receives an uplink allocation from the base station. The uplink allocation specifies the subframe of the uplink message, the frequency (subcarriers) of the



**Figure 2: Spectrogram of the LTE uplink for one subframe with all the possible uplink physical channels. The uplink reference signals (RS), which are used for uplink measurements (Section 5), are highlighted in blue. The X-axis represents the time, and in this figure it spans one subframe, lasting 1 millisecond. The Y-axis represents the frequency; in this case, the uplink channel is 20 MHz wide.**

uplink message, and the configuration parameters of the reference symbols within the uplink message, such as which symbols during the PUSCH transmission are reserved for the reference signals. In LTE, 2 out of 14 symbols are allocated to reference signals.

A PUCCH message is sent in pre-configured subframes by the UE or after a downlink transmission as an ACK/NACK. In LTE, PUCCH transmission occupies exactly one PRB on the edge of the cell's spectrum and lasts one subframe. Hopping is often configured for PUCCH, meaning that the first half of the PUCCH is transmitted on the lower edge of the spectrum, and the second half is on the upper edge. It also consists of 14 OFDM symbols; however, typically 6 of them are allocated to reference symbols. For 5G, the PUCCH is similar but with many more configurable parameters.

## 2.3 Physical Downlink Channels

Similar to the uplink, the physical layer of the downlink also uses OFDMA, with its resource grid split into channels. The PDCCH contains information about the resources for the current downlink transmissions and can also contain uplink allocations for upcoming user transmissions. In contrast, the PDSCH carries cell data, which can include unicast user data, broadcast paging messages, or configuration messages (e.g., SIBs).

## 3 System Design

WaveTag monitors the downlink and uplink of all cells (across multiple operators and frequency bands) visible from the sensitive area. For each UE connection originating within the sensitive area, WaveTag captures the required uplink signal features with sufficiently low latency to allow intervention *before* any communication with the operator's core network occurs.

Accurate spatial inference about a UE connection, such as precise localization or geofencing, typically requires multiple observations of uplink messages. Given the limited number of uplink messages available before connection establishment with the core network, WaveTag overcomes this constraint by deploying a distributed set of




Figure 3: In (A), we show a sketch of a deployment with two downlink receivers and six uplink receivers positioned around the sensitive area. The purple triangles on the uplink receivers show the azimuth of the two opposite facing antennas. In (B), we show how the system components interact. First, the downlink receivers detect a new connection on the downlink and inform the UL receivers about it. The UL receivers then measure the corresponding uplink signals. Finally, all measurements are aggregated in the central unit (CU), which makes a decision about the connection. In (C), we show how a concrete wall can be used for improved inside/outside separation of a single uplink receiver unit.

uplink receivers strategically positioned in and around the sensitive area.

In this section, we present the design of WaveTag, which addresses the three primary challenges that have remained unaddressed until now: scalable handling of multiple operators and frequency bands, adherence to stringent low-latency constraints, and the ability to support numerous distributed uplink receivers.

### 3.1 System Model and Assumptions

WaveTag consists of one or several downlink receivers and a larger number of geographically distributed uplink receivers. The downlink receivers are synchronized to all xNBs around the target area and monitor their downlink channels for new connections. Once a new connection is detected, the downlink receiver looks for uplink allocations that correspond to the connecting UE. Upon receiving an uplink allocation, the downlink receiver notifies all uplink receivers about the upcoming uplink transmission. All uplink receivers then capture the uplink message and measure its signal characteristics without the need to decode the data contained in the uplink messages. The measurements from all uplink receivers are aggregated in a central unit (CU), where they can be further processed. In our use case shown in Section 5, a model classifies the connection as originating from within or outside the sensitive area. This system model is depicted in Figure 3 in an example deployment.

Due to the requirement for UE detection and decision before full connection establishment with the core network, WaveTag collects metrics only for messages sent by the UE before full connection establishment. This includes at most one PRACH preamble, one RRC Connection Request on PUSCH, and one or two subsequent PUCCH messages. Subsequent PUSCH messages after the RRC Connection Request could already include messages such as a Service Request, reestablishing an existing security context and effectively communicating with the core network.

While we do not impose strict restrictions on the size or shape of the sensitive area that needs protection, our system is intended

for sensitive areas that can encompass multiple buildings and span thousands of square meters. We note that those areas are typically separated from the outside by a wall or fence, as this often applies to sensitive perimeters. These facilities can be located in dense urban environments and therefore the separation between inside and outside can be as low as 5 meters, requiring sufficiently precise spatial information for the classification mechanism to work. Furthermore, built-up areas exhibit rich radio frequency environments with several xNBs operated by different cellular network operators, leading to a predominance of non-line-of-sight (NLoS) signal paths between xNB and user equipment.

**3.1.1 Attacker Model.** An adversary aims to establish a cellular link from within a sensitive area without being detected. The adversary can use any commercial off-the-shelf (COTS) device that uses a specification-compliant cellular baseband chipset such as smartphones or cellular modules. Additionally, the adversary may use any software on such COTS devices, including running in higher privilege mode (e.g., a rooted phone). We further assume that the adversary is able to obtain any SIM card (e.g., roaming SIM card) in order to connect to any operator in the area. We also allow the adversary to interfere with the cell selection mechanism by, e.g., forcing the baseband chip to connect to any visible cell, even if further away or with a worse SNR. Finally, we assume the attacker does not use bulky analog radio equipment such as narrow-beam directional antennas, large poles or similar, as any equipment except COTS devices would be spotted during typical entrance inspections in the proposed use cases in the Section 1.

### 3.2 Downlink Receivers

The downlink receiver is an SDR-based platform that has the task of monitoring the downlink of cells visible on all bands from the sensitive area. To decrease inter-cell interference and enhance decoding rates, we use narrow beam-width directional antennas. To cover all the directions, we plan for four ideally positioned downlink receivers with a direct line of sight to the xNBs. In this way,

each xNB in the vicinity is covered by at least one downlink receiver. A downlink receiver with good downlink reception quality continuously remains synchronized with the xNBs and decodes all messages that are broadcast on the downlink.

After successfully synchronizing in time and frequency to a cell, the downlink receiver decodes the MIB and SIB and applies the cell-specific configurations. These configurations include cell-specific information required for new UE connection detection and uplink signal measurements. Each new UE connection includes a RAR sent by the xNB as a response to PRACH preamble (see Figure 1). By monitoring the RA-RNTIs defined in the SIB messages, we can decode all RARs, and thus detect all new connections on the xNB.

After the downlink receiver decodes the RAR, it immediately forwards the RAR to the uplink receivers, stores the UE-specific RNTI in the set of active RNTIs, and continues monitoring and decoding subsequent messages addressed for the active RNTIs. In case an RRC Connection Setup message is decoded, the downlink receiver applies the enclosed UE-specific configuration to the decoding process for this particular UE.

### 3.3 Uplink Receivers

Uplink receivers, operating also on an SDR-based platform, are connected to two wide-beam directional antennas facing away from each other, as pictured in Figure 3. This antenna orientation enables uplink receivers to estimate the direction of the signal (see Section 5 for more details). Each SDR streams IQ samples of the uplink of one or more cellular frequency bands to an uplink receiver, where it is buffered for at least 20 milliseconds. One of the main advantages of WaveTag over existing solutions is that uplink receivers are *decoupled* from downlink receivers. The uplink receivers do not need to monitor downlink frequencies, and therefore they do not need to acquire a sync to the xNBs or decode downlink messages. Instead, they detect and measure uplink signals using the correct uplink allocations and timings received from the downlink receiver.

A key requirement for WaveTag is that the downlink and uplink receivers share a common clock. This ensures that the uplink allocations received by the downlink receiver can be mapped to the correct time and frequency to be observed by the uplink receivers.

All relevant downlink control messages, including uplink allocations from RARs, are sent to all uplink receivers over a low-latency communication link. Each uplink receiver then independently looks up the corresponding recorded time and frequency resources for the uplink message in the buffered IQ stream. WaveTag then extracts the signal characteristics (see Section 5) based on the reference signals and finally sends them to the central unit.

Using this process, uplink receivers detect and measure all three types of uplink messages in LTE: PRACH, PUSCH, and PUCCH. In 5G, only PRACH and PUSCH feature extraction is currently implemented.

We position uplink receivers around the perimeter of the target area. We try to position them such that there is a line of sight to every location where a UE might be present. In case the UEs are inside a building and it is not possible to have a line of sight, we try to position the antennas such that they are on the clearest path out of the building, e.g., having only a window or a single wall in the path. We found this approach to work well when the target area is

relatively large and consists of multiple buildings separated from the outside world by a wall or empty space.

### 3.4 Central Unit

Measurements of individual uplink messages are sent over a low-latency communication link and aggregated at the central unit. Once we receive measurements from every uplink receiver, we can further process the signal metrics for geofencing (or localization) purposes.

Each uplink message sent by the UE is processed by the central unit as soon as all uplink receivers have submitted their signal characteristics measurements for that message. This can be a heavy-weight process, for example, running inference on a neural network. In a geofencing context, once the system is asked to make a final decision about the connection, the central unit aggregates the decisions from the individual messages into one result. This combination is a fast operation, as we can expect no more than four different message results (e.g., a downlink message might be split into multiple PDSCH messages requiring multiple PUCCH responses). We can therefore perform this operation in real time.

## 4 WaveTag Implementation

In this section, we explain how our system achieves both low latency and real-world scalability, including monitoring multiple cells per operator across all their allocated frequency bands. To this end, we show how the radio units acquire and process IQ samples and how the sub-components communicate with each other.

As an overview, in Figure 4, we show the messages exchanged between the UE and the xNB, the actions of our system's sub-components in response to these messages, and their associated mean latencies.

### 4.1 Downlink Monitoring


We note that while downlink decoding plays a role in this process, it is not the main contribution of this paper and will not be discussed in detail. There are many downlink receiver implementations already available [18, 22–24, 37], and we use XYZ<sup>1</sup> as a base for our downlink sniffer.

The main feature of this downlink receiver is the ability to decode multiple cells in parallel on a single SDR device. It continuously tracks multiple cells from the same operator or different operators in parallel. This is achieved by splitting the single IQ stream from the SDR into multiple parallel streams using our software channelizer implementation. Our implementation is an efficient implementation of overlap save filtering [17], using FFTW3 and AVX instructions.

### 4.2 Uplink Processing

To extract the metrics of the uplink messages reliably, the uplink and downlink receivers must be synchronized to the same time source with a low absolute pairwise offset. In our deployments, we have created a clock distribution tree that spans across all radio units. At the root, there is a high-precision clock synchronization device, generating highly stable 10MHz and 1PPS clock signals, disciplined

<sup>1</sup>To avoid de-anonymization, we will include a citation to our downlink receiver in the camera ready paper if the paper is accepted.



**Figure 4: Overview of a typical connection on one of the FDD bands with annotated mean processing latencies for end-to-end measurement of PRACH, PUSCH, and PUCCH transmissions. The downlink receiver decodes the RAR and PDSCH messages. The RAR message both informs about the PRACH message 3-13ms in the past and the PUSCH message 6ms in the future. Any PDSCH message triggers a corresponding PUCCH message 4ms later.**

by GPS. This signal is then distributed with the WhiteRabbit protocol [39] to all radio units.

Classically [27, 30], uplink monitoring of an LTE cell is achieved by tuning an SDR receiver to the center frequency of the uplink of the cell and collecting 30.72M IQ samples per second. However, in each of the frequency bands, there are three to four operators, each transmitting on different frequencies. Using the classical approach, we would require up to 24 SDRs in each uplink receiver to monitor the entire uplink spectrum. Such a system would be extremely expensive and difficult to maintain.

In our solution, we use SDRs with a sampling rate of at least 122.88Msps, such that they have enough bandwidth to receive a whole FDD frequency band. For the 5G TDD band with cells having as much as 100MHz bandwidth, multiple SDRs are needed. To ensure that the DC component does not disturb the uplink frequencies, we set the center frequency of our SDR between two operators allocated within a band, as close to the middle of the frequency band as possible. An uplink receiver continuously receives IQ samples from all SDRs and writes them into circular buffers. We also keep the timestamp of the last written sample. Using this timestamp, we can compute the absolute time of IQ samples in the buffer.

For each monitored cell, we spawn a new cell object, identified by its technology (LTE or 5G), center frequency, and Physical Cell Identity (PCI). When a cell object receives an uplink allocation from the downlink receiver, it is pushed into an event queue with the corresponding uplink transmission time. Each cell object continuously checks in the circular buffer if it has enough samples to process the earliest uplink allocation. Once samples for the uplink allocation are fetched from the circular buffer, our software extracts the uplink signal features.

During the feature extraction, WaveTag first shifts the copied IQ samples in frequency such that the cell's uplink band is centered in the middle of the received bandwidth. Afterwards, it computes an FFT of size  $\text{sampling rate} / \text{subcarrier spacing}$  for each OFDM symbol containing a known reference signal. This FFT operation gives us modulated values at each individual subcarrier. Then, only the subcarriers defined in the uplink allocation are further processed, i.e., we compute the circular correlation of the received reference signal with a known reference signal (see Section 5 for more details).

Given that the number of allocated subcarriers is generally very low for the initial messages, this circular correlation is quick.

This processing is executed independently in parallel for each of the two radio ports. In our real-world deployment on the frequency band with the most activity, such a measurement was done on average every 5 milliseconds across all cells. Once the measurement is completed for both ports, the uplink receiver transmits the signal characteristics to the central unit.

### 4.3 Message Exchange

The downlink receiver, the uplink receivers, and the central unit continuously communicate with each other over a low-latency link. They are connected with fiber links in a star architecture with a switch in the middle. Given the relatively small number of devices, the overhead of such a network topology is small. In Section 6.1, we evaluate how long a message exchange takes.

We use the ZMQ protocol [8] as it is a low-latency protocol with minimal overhead. We utilize its PUB/SUB pattern, where the transmitters are publishers and the receivers of messages subscribe to the publishers. Essentially, publishers broadcast messages and they do not have any information on the state of the subscribers. However, this also means that if a subscriber goes offline, the publisher does not know about it. In our system, we expect relatively low component downtime, and if any components shut down, the distributed system will be minimally affected. Therefore, a PUB/SUB pattern is a good choice for the minimal overhead it brings. We run all messaging services on a separate thread.

### 4.4 Central Unit

The measurements of the individual messages are sent over ZMQ to a central unit. When the first measurement for a message arrives, we start a timer with a timeout of 2 milliseconds. Once all uplink receivers have delivered the measurements, or when the timer runs out, we process the signal measurements for this uplink message as outlined in the next section. Therefore, even if an uplink receiver malfunctions, it does not have a big impact on the overall latency.


## 5 Low-latency Geofencing

The main challenge of precise geofencing is classifying a connection's origin—i.e., determining whether it is inside or outside a sensitive area—based on as few uplink messages as possible. To achieve this, WaveTag uses a combination of relative signal strength and signal quality features. These features are collected from multiple uplink receivers, each equipped with two directional antennas facing in opposite directions. Finally, a classification model running on the central unit infers whether the connection originates inside or outside the perimeter.

### 5.1 Uplink Signal Features

To measure the uplink signal features, we rely on the PRACH preamble or reference signals sent in PUSCH (and PUCCH in LTE). Since these signals are known from the cell configuration or user-specific configurations, our receivers can efficiently compute the circular cross-correlation of the received signal with the reference, without having to decode the actual messages. The OFDM symbols that contain the data are essentially ignored. Finally, the OFDM Cyclic Prefix allows the cross-correlation to be calculated even if the signal is received with a time offset. From the circular cross-correlation, we extract signal power (peak) and signal quality characteristics (peak to average power ratio). In Appendix A, we give more details on what exact features we collect.

The radio units of our uplink receivers are equipped with two receiving ports. Each port is connected to a separate directional antenna, oriented in opposite directions. Intuitively, if a signal is significantly stronger on one antenna, the signal must originate from the direction of that antenna's azimuth. On the other hand, if the signal power on both antennas is equal, the direction is perpendicular to the azimuth of both antennas.



**Figure 5: Measured signal power difference between two ports of a single uplink receiver in a real-world deployment. The color shows the measured signal power difference of an uplink signal transmitted from the corresponding point on the map.**

We can further adapt the direction-finding approach of the two ports to focus on classifying signals as either inside or outside the target area. The most straightforward method involves placing uplink receivers along the border of the target area, with one antenna positioned as an "inside" antenna and the other as an "outside" antenna. Often, the target area is separated from the outside by a boundary, such as a concrete wall or another physical feature. While a boundary such as a wall is not strictly required, it significantly enhances separation when available. If it is not feasible to position antennas on different sides of the boundary, both antennas can be placed on the same side. In such cases, one antenna is positioned above the wall and the other below, ensuring that at least one benefits from the wall's attenuation effects. This flexible approach enables the system to adapt to various boundary conditions. An example of this antenna setup is shown in Figure 3. To showcase this separation in the real world, we plot in Figure 5 the relative signal power difference between the two ports of one radio unit for uplink messages collected in one of our deployments.

### 5.2 Geofencing Model

Our classification model is a fully connected neural network that uses relative features as input. We first train the model offline with Python and pyTorch [4]. We use a simple architecture to prevent overfitting and ensure fast inference time. The architecture consists of BatchNorm, Dropout and two fully connected layers followed by an ReLU activation function. Finally, we pass the output neuron's result through a sigmoid function to produce the classification decision. We use binary cross-entropy as a loss function and optimize the model using an Adam optimizer. We then train the model in batches and stop training once the loss does not improve across several epochs. Finally, we perform a hyperparameter search for each deployment separately to ensure an optimal solution.

For all connections, we get between 2 and 4 messages before having to make a decision on whether the phone is inside or outside. Thus, we need to combine the decision results (a real number between 0 and 1) of the individual messages. To this end, we train a logistic regression model [2], which weights each classification result based on its message type. If we have not received a given message type, we replace its classification result with 0.5. If we have multiple classification decision results of the same message type, we take their mean value.

### 5.3 Data Collection

For any ML model, the most important factors are the quality of the training data and the number of data points. Generating uplink transmissions with a typical modern smartphone is easily possible, for example, by switching airplane mode on and off. However, there is normally no access to low-level base-band messages from the RRC and NAS layer. We therefore employ a set of cellular COTS modems. These modems have the same functionalities as modern smartphones in terms of cellular communication, and we can use these modems and smartphones interchangeably. However, they also provide low-level debugging features and logs. For example, we can force the modem to connect to a specific frequency band or even a particular cell. It also provides the connection identifier (RNTI), which we use to match the recorded signal features at the

| Operation             | Mean                         | StdDev                      | Count     |
|-----------------------|------------------------------|-----------------------------|-----------|
| Downlink Decoding     | 540 $\mu$ s                  | 155 $\mu$ s                 | 47        |
| Message Round-Trip    | 75 $\mu$ s                   | 2 $\mu$ s                   | 1000      |
| LTE PRACH Measurement | 587 $\mu$ s                  | 147 $\mu$ s                 | 2027      |
| LTE PUSCH Measurement | 505 $\mu$ s                  | 129 $\mu$ s                 | 2024      |
| LTE PUCCH Measurement | 423 $\mu$ s                  | 145 $\mu$ s                 | 3534      |
| Model Inference       | 254 $\mu$ s                  | 78 $\mu$ s                  | 147       |
| <b>LTE E2E PRACH</b>  | <b>2223<math>\mu</math>s</b> | <b>230<math>\mu</math>s</b> | <b>48</b> |
| <b>LTE E2E PUSCH</b>  | <b>877<math>\mu</math>s</b>  | <b>128<math>\mu</math>s</b> | <b>26</b> |
| <b>LTE E2E PUCCH</b>  | <b>1028<math>\mu</math>s</b> | <b>155<math>\mu</math>s</b> | <b>13</b> |
| 5G PRACH Measurement  | 479 $\mu$ s                  | 51 $\mu$ s                  | 140       |
| 5G PUSCH Measurement  | 602 $\mu$ s                  | 64 $\mu$ s                  | 140       |

**Table 1: Latency measurements of both individual parts of the system and the overall E2E latency from when the message is sent by the UE until a decision is available.**

central unit with the modem’s recorded location. This gives us a robust way to collect a large amount of labeled training data across multiple operators, bands, and base stations.

To improve the efficiency of the data collection, we fitted a backpack with multiple such modems that connect to the network almost every second. We can then walk a route and easily gather large amounts of uplink signal features from any location. Such an annotated dataset is crucial for training and testing of the model. We make this data available online to support further research.

## 6 Evaluation

To evaluate the performance of our system, we conducted tests in two different deployments with two different SDR platforms. The first deployment was used as a feasibility study with COTS SDRs and was limited to a single frequency band at a time. The second deployment is our full system, with hardware that is capable of monitoring all frequency bands from all operators in the country simultaneously. For both deployments, we had the necessary approvals from the relevant institutions to conduct the tests and we tested the geofencing performance against our own devices. At the time of these tests, only LTE and 5G-NSA networks were available at these locations. To show that our system is 5G SA capable, we performed a set of tests on a commercial 5G SA network. We confirm that key features—latency, feature extraction, and ability to track multiple cells—still hold for 5G SA.

### 6.1 Latency

In this section, we evaluate the latency of individual processing steps and also the end-to-end (E2E) latency of our system for various types of messages. Both individual steps and E2E latencies are shown in Figure 4. We use the Perfetto [25] and Tracy [41] libraries to measure the processing times of individual components during the runtime of the whole system. The exception is the measurement for the ZMQ message round-trip times, for which we built our own measurement tool.

In Table 1, we show the measurement results of the individual parts of the system, as well as the E2E latencies for various message types. For the E2E PRACH measurement, we measure the elapsed time from the reception of the RAR by the downlink receiver to the completion of the inference of the geofencing model. This time encapsulates the downlink decoding of the RAR, the message exchange between the downlink receiver and the uplink receiver, the PRACH measurement by the uplink receiver, the message exchange between the uplink receiver and the central unit, and finally the model inference. For the E2E PUSCH and E2E PUCCH measurements, we measure the time from the reception of the uplink message to the inference result. This measurement does not encapsulate downlink decoding and initial message exchange, as WaveTag knows about the uplink message already 4 subframes prior to the uplink transmission.

In summary, we need less than 2.3ms after the RAR is transmitted to know the result for the corresponding PRACH. For a PUSCH message, the classification is ready on average 0.88ms after the uplink signal is received. Finally, a PUCCH message is classified on average 1.0ms after the uplink signal is finished transmitting. The latency of WaveTag does not suffer when measuring features in 5G SA compared to LTE. From measurements done on the TDD band while monitoring a commercial 100MHz cell, we can see a slight improvement of 108 $\mu$ s for PRACH measurement, and a slight slowdown of 97 $\mu$ s for PUSCH measurement.

### 6.2 Scalability

In the second deployment, we observed that WaveTag was monitoring more than 100 cells at the same time. Given that we monitor 5 frequency bands, a single SDR can monitor the uplink of more than 20 cells. From our data set, we know that we observed our test phones connecting to 62 of those cells across all bands.


Overall, in the second deployment, we measured 1,172,772 uplink signal messages on 10 uplink receivers, with 620 (less than 0.01%) not received by the CU.

WaveTag is capable of monitoring multiple cells on the same frequency in parallel on a single SDR port for 5G SA. We evaluate this claim by setting up WaveTag on a COTS SDR platform and conducting a 1-hour-long monitoring campaign on two real-world cells in the vicinity. The two cells have a bandwidth of 15 MHz and are operating on the same frequency in an FDD band. Using this setup, we observed 9,106 connections in total, with 8,230 on the first cell and 876 on the second cell. For all connections, we successfully measured both the PRACH and the PUSCH messages.

Moreover, as discussed in Section 4.2, our system is capable of monitoring an entire 100MHz 5G SA cell. Similarly to before, we ran WaveTag on a 100 MHz real-world cell on a TDD band in the vicinity for 1 hour. In this time span, we detected 1,152 connections and we failed to measure only one PRACH for those connections. For PUSCH, we did not observe any failures.

### 6.3 Geofencing

**6.3.1 First Deployment.** The initial feasibility deployment, shown in Figure 6, was set in the middle of the urban area. The area is approximately 100  $\times$  200 meters and consists of a single multi-story building with many wings, and many outside areas/courtyards



**Figure 6: First deployment in an urban area showcasing six uplink receivers (the two triangles on each uplink receiver show the azimuth of their antennas) positioned around the sensitive area. Blue dots represent correct geofencing (inside/outside), with orange dots showing false positives and red dots representing false negatives.**

around the building. Part of the sensitive area is divided from the outside world by a concrete wall, but on the right side it is separated from buildings with just a road. The separation between the "inside" and "outside" classes is as little as 3 meters in some areas.

We equipped the site with one downlink receiver placed on the roof of the building and six uplink receiver units positioned around the sensitive area. As radio units, we used COTS SDRs: for the downlink receiver, we used a USRP B210 and for uplink receivers, we used  $2 \times$  USRP X310,  $2 \times$  USRP N310, and  $2 \times$  USRP X410. All SDRs were connected to a central unit using fiber optic cables.

We collected samples over 5 days using our custom data collection rig. We collected training data for 244,804 uplink messages, spanning 43,157 connections. In this set, we have 63.3% inside and 36.7% outside connections. To acquire the test set, we collected data on a separate day to ensure that we were not overfitting on specific environmental conditions. The test set consists of 29,809 uplink messages spanning 6,198 connections. 58% of those are inside connections and 42% outside connections. Finally, the geofencing model achieves an accuracy of 95.78% on individual messages from the test set.

|           |         | Actual |         |             |
|-----------|---------|--------|---------|-------------|
|           |         | Inside | Outside |             |
| Predicted | Inside  | 3503   | 68      | Acc = 97.42 |
|           | Outside | 92     | 2535    | FPR = 2.61% |
|           |         |        |         | FNR = 2.56% |

**Table 2: Geofencing accuracy results for test set from the first deployment.**

A logistic regression model, trained on the training data set, estimates the weight of each type of message in the final ensemble. The PRACH message has a weight of 0.36, PUSCH 0.37, and PUCCH 0.27. As reported in Table 2, when we run the logistic regression model on the test set, we achieve an accuracy of 97.42%. From

|           |         | Actual |         |              |
|-----------|---------|--------|---------|--------------|
|           |         | Inside | Outside |              |
| Predicted | Inside  | 10442  | 30      | Acc = 99.66% |
|           | Outside | 43     | 10796   | FPR = 0.28%  |
|           |         |        |         | FNR = 0.41%  |


**Table 3: Geofencing accuracy results for test set from the second deployment.**

Figure 6, we see that the model mostly misclassified messages close to the border of the sensitive area, which is expected.

**6.3.2 Second Deployment.** After the successful initial deployment, we deployed our system in a second, larger location. Here, we monitor an area of around  $250\text{m} \times 250\text{m}$ . The sensitive area consists of multiple buildings with outdoor locations around and in between the buildings. The sensitive area is located on the outskirts of an urban area with neighboring buildings on two sides. We deployed ten prototype uplink receivers based on the custom-built SDR.

We again collected measurements over multiple days using our custom data collection rig based on LTE COTS modems. For the training set, we collected 486,850 measurements for 103,564 connections over 6 days. The inside/outside split is 52.3%/47.7%. We collected the test data set on a separate day with 21,311 connections.

We trained a geofencing model as a fully connected neural network model, and furthermore, we trained a logistic regression that aggregates results over multiple messages. The estimated weights from the logistic regression are: 0.39 for PRACH, 0.38 for PUSCH and 0.23 for PUCCH, approximately matching the weights from the first deployment. The accuracy of the neural network model is 99.02% on individual messages. The ensemble model then increases the accuracy for the entire connection to 99.66%. In Table 3, we report more detailed results for the final geofencing model. Throughout the deployment, we also conducted regular spot-checks with a large number of commercial phones. We did not observe any degradation in performance for any particular model.




**Figure 7: Effect of uplink receivers being switched off.**

Furthermore, we performed robustness tests to determine how well the model behaves if fewer uplink receivers are available or

if fewer messages are available for classification. For the first test, we plot these results in Figure 7. In the second test, with just one message (PRACH), the model's accuracy is 99.17%, and with the first two messages (PRACH and PUSCH (RRC Connection Request)), the model's accuracy improves to 99.6%.

## 6.4 5G-SA Feasibility Study



**Figure 8: Test of the feature extraction from the 5G PHY layer performed on a commercial 5G SA network. Each point on the map represents a test location. Its color shows the relative signal power measured on two ports of the uplink receiver.**

To evaluate the 5G SA capabilities of WaveTag, we set up an uplink receiver running on a USRP X410 right in the middle of a large open field. We used a setup with two directional antennas pointing away from each other. We then conducted an extensive walk around the field with a smartphone (OnePlus 10 Pro 5G) connecting to the commercial 5G SA network. For each connection request, our system measured the features for the PRACH and the first PUSCH message and stored them in a database. Finally, to show our uplink receiver's opposite-facing antenna setup, we plot the relative signal power between the two ports on a map in Figure 8.

With this experiment, we confirm that feature extraction works in 5G SA. Moreover, at the farthest points from the uplink receiver (100m away), we still measured an SNR of more than 20dB. Furthermore, we can see a clear separation between the points on the left and right side of the field in Figure 8, meaning that the ratio of signal strengths measured on the individual ports correlates with the uplink signal direction. Given the similarities in the PHY layer of LTE and 5G, and the successful tests in the field, we conclude that if 5G SA were deployed at one of our sites, WaveTag would also correctly geofence 5G SA connections.

## 7 Related Work

Other projects that build operator-independent monitoring systems and rely on uplink signal measurements include [30, 31, 35]. They all choose to couple uplink and downlink processing on a single host using one (or multiple) SDRs. This approach comes with two disadvantages: (1) *Colocation of downlink and uplink sniffers*: This constraint requires downlink sniffers to be placed only where good downlink reception is available. As a result, uplink receivers cannot be positioned indoors or in complex non-line-of-sight environments. (2) *One-to-one requirement between uplink and downlink*

*receivers*: Each uplink receiver must have a dedicated downlink receiver, increasing the number of required devices and the system's processing load. In contrast, our system decouples the downlink and uplink receivers, allowing each to be placed optimally. This design also improves scalability, as a single downlink receiver can serve multiple uplink receivers.

In our deployment, we observed the activity of our phones on 62 different cells in all frequency bands. [30, 31, 35] lack the ability to continuously monitor multiple cells simultaneously on a single receiver. As a result, monitoring these 62 cells would require a separate pair of uplink and downlink receivers for each cell, multiplied by the number of distributed locations in a deployment. WaveTag processes the uplink signals of the cells in a single frequency band using just one SDR. To put it into perspective, our system in the second deployment could be implemented with as little as 38 COTS SDRs, compared to more than 620 SDRs required by [35].

However, the context in which these systems were implemented is also different. Mainly, they were designed to accurately localize users from multiple observations during a connection, e.g., looking for an uncooperative cellular device from a moving police car. Because to their use cases, it is enough to only observe a large number of messages on PUSCH, use rotational antennas to slowly but accurately infer the direct signal path, or potentially use active power-boosting attacks to increase the UE's power [35]. As a result, localization in [35] takes around 2 minutes, [30] measures ToA values of all messages from the whole connection, and [31] requires offline processing of recorded samples. Even though they achieve their respective tasks, these points show that they are unsuitable for low-latency geofencing. WaveTag compensates for the low number of observations with a higher number of uplink receivers, while also measuring PUCCH and PRACH messages.

Finally, all these works lack the capability to monitor 5G SA. The only existing 5G SA uplink sniffer, Sni5gect [32], is built to decode uplink (PUSCH) messages on just one xNB. However, it is not capable of extracting signal features crucial for geofencing and its monitoring capabilities were only evaluated in a 20m range on a lab base station. Even at this small distance, the decoding rate was as low as 55% for one of the phones. On the other hand, in 5G, our WaveTag extracts important signal metrics not only for PUSCH but also for PRACH. Our system and its low-latency signal extraction capabilities were tested on a commercial operator's network. Even at 100m distances, it consistently measured uplink signals with more than 20dB SNR. We also showed WaveTag's capability to monitor multiple base stations on a single SDR.

In this paper, we do not try to compare our geofencing accuracy with previous localization techniques. These techniques use TDoA, AoA, or RSSI principles [26, 28, 29, 36] to achieve a sufficiently high localization accuracy (localization being a harder problem to solve than geofencing) but also require expensive optimization algorithms. However, WaveTag is extensible and could integrate new localization algorithms for various use cases. To this end, we are open-sourcing our datasets and invite researchers to develop and test different algorithms in a real-world setting, which could then be ported to our platform.

## 8 Conclusion

In this paper, we presented WaveTag, a system designed to reliably monitor the uplink of all LTE and 5G connections within sensitive areas while addressing key challenges in latency and scalability. Unlike other systems, WaveTag relies only on the initial few messages to make a geofencing decision, enabling low-latency operation critical for rapid countermeasures. Through multiple evaluations on a commercial network, we show that WaveTag is the first system in its category capable of monitoring 5G SA. Our system achieves high geofencing accuracy, even in non-line-of-sight environments, as demonstrated through evaluations in two separate deployments using different hardware platforms. In the second, full-scale deployment, WaveTag achieved a geofencing accuracy of 99.66%.

## References

- [1] Ashley Benefield claims juror snuck phone into deliberations. URL: <https://www.courtstv.com/news/ashley-benefield-claims-juror-snuck-phone-into-deliberations/>.
- [2] LogisticRegression. URL: [https://scikit-learn/stable/modules/generated/sklearn.linear\\_model.LogisticRegression.html](https://scikit-learn/stable/modules/generated/sklearn.linear_model.LogisticRegression.html).
- [3] LTE Positioning Protocol (LPP) # 37.355. URL: <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationid=3710>.
- [4] PyTorch. URL: <https://pytorch.org/>.
- [5] resample — SciPy v1.14.1 Manual. URL: <https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.resample.html>.
- [6] Stage 2 functional specification of User Equipment (UE) positioning in E-UTRAN # 36.305. URL: [https://www.etsi.org/deliver/etsi\\_ts/136300\\_136399/136305/16.01.00\\_60/ts\\_136305v160100p.pdf](https://www.etsi.org/deliver/etsi_ts/136300_136399/136305/16.01.00_60/ts_136305v160100p.pdf).
- [7] Vatican deploys cellphone signal jammers for conclave. URL: <https://www.nbcnews.com/video/vatican-deploys-cellphone-signal-jammers-for-conclave-239076933776>.
- [8] ZeroMQ. URL: <https://zeromq.org/>.
- [9] Scanners catch JEE cheats, June 2015. URL: <https://www.telegraphindia.com/west-bengal/scanners-catch-jee-cheats/cid/1491086>.
- [10] Cheating in exams: Smartphones blamed for rise in incidents. March 2019. URL: <https://www.bbc.com/news/uk-northern-ireland-47406500>.
- [11] The era of the cellular-connected drone has arrived, September 2021. URL: <https://www.verizon.com/about/news/era-cellular-connected-drone-has-arrived>.
- [12] Neo-Nazi sold gun parts to extremists online while in prison using smuggled phones, DA Bragg says, June 2024. URL: <https://www.nbcnews.com/news/us-news/neo-nazi-sold-gun-parts-extremists-online-prison-using-smuggled-phones-rcna156547>.
- [13] French police launch prison hunt for Chinese-made miniature phones, May 2025. URL: <https://www.bbc.com/news/articles/cnv11qvq397o>.
- [14] 3GPP TS 36.101. 3GPP TS 36.101 V16.20.0, May 2024. URL: [https://www.etsi.org/deliver/etsi\\_ts/136100\\_136199/136101/16.20.00\\_60/ts\\_136101v162000p.pdf](https://www.etsi.org/deliver/etsi_ts/136100_136199/136101/16.20.00_60/ts_136101v162000p.pdf).
- [15] Richard Adams and Richard Adams Education editor. More than 1,000 pupils penalised for phones in GCSE and A-level exams. *The Guardian*, January 2018. URL: <https://www.theguardian.com/education/2018/jan/05/students-cheating-mobile-phones-gcse-exams-a-levels-2017>.
- [16] Donnie Anderson. CONTRABAND CELLULAR TELEPHONES, October 2019. URL: [https://oksenate.gov/sites/default/files/interim\\_materials/Interim-Studies/10-21-2019%20Contraband%20Cell%20Phone%20Presentation.pdf](https://oksenate.gov/sites/default/files/interim_materials/Interim-Studies/10-21-2019%20Contraband%20Cell%20Phone%20Presentation.pdf).
- [17] M. Borgerding. Turning overlap-save into a multiband mixing, downsampling filter bank. *IEEE Signal Processing Magazine*, 23(2):158–161, March 2006. Conference Name: IEEE Signal Processing Magazine. URL: <https://ieeexplore.ieee.org/abstract/document/1598092>, doi: 10.1109/MSP.2006.1598092.
- [18] Nicola Bui and Joerg Widmer. OWL: a reliable online watcher for LTE control channel measurements. In *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*, ATC '16, pages 25–30, New York, NY, USA, October 2016. Association for Computing Machinery. URL: <https://dl.acm.org/doi/10.1145/2980055.2980057>, doi: 10.1145/2980055.2980057.
- [19] Cristian Capotă, Mădălin Popescu, Eduard-Marian Bădulă, Simona Halunga, Octavian Fratu, and Mircea Popescu. Intelligent jammer on mobile network LTE technology: A study case in bucharest. 13(22):12286. URL: <https://www.mdpi.com/2076-3417/13/22/12286>, doi: 10.3390/app132212286.
- [20] Suma S. Cherian and Ashok N. Rudrapatna. LTE location technologies and delivery solutions. *Bell Labs Technical Journal*, 18(2):175–194, September 2013. Conference Name: Bell Labs Technical Journal. URL: <https://ieeexplore.ieee.org/document/6772141?arnumber=6772141>, doi: 10.1002/bltj.21612.
- [21] DJI Enterprise. Introducing 4G Connection for Drones with DJI Cellular Dongle, August 2024. URL: <https://enterprise-insights.dji.com/blog/introducing-4g-connection-for-drones-with-dji-cellular-dongle>.
- [22] Simon Erni, Martin Kotuliak, Richard Baker, Ivan Martinovic, and Srdjan Capkun. GLaDoS: Location-aware Denial-of-Service of Cellular Networks. pages 5307–5325, 2025. URL: <https://www.usenix.org/conference/usenixsecurity25/presentation/erni>.
- [23] Simon Erni, Martin Kotuliak, Patrick Leu, Marc Roeschlin, and Srdjan Capkun. AdaptOver: adaptive overshadowing attacks in cellular networks. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, MobiCom '22, pages 743–755, New York, NY, USA, October 2022. Association for Computing Machinery. doi: 10.1145/3495243.3560525.
- [24] Robert Falkenberg and Christian Wietfeld. FALCON: An Accurate Real-Time Monitor for Client-Based Mobile Network Data Analytics. In *2019 IEEE Global Communications Conference (GLOBECOM)*, pages 1–7, December 2019. ISSN: 2576-6813. URL: <https://ieeexplore.ieee.org/document/9014096>, doi: 10.1109/GLOBECOM38437.2019.9014096.
- [25] Google. google/perfetto. original-date: 2019-12-10T10:32:44Z. URL: <https://github.com/google/perfetto>.
- [26] Jiajun He and Hing Cheung So. A Hybrid TDOA-Fingerprinting-Based Localization System for LTE Network. *IEEE Sensors Journal*, 20(22):13653–13665,

- November 2020. Conference Name: IEEE Sensors Journal. URL: <https://ieeexplore.ieee.org/abstract/document/9122501>, doi: 10.1109/JSEN.2020.3004179.
- [27] Tuan Dinh Hoang, CheolJun Park, Mincheol Son, Taekkyung Oh, Sangwook Bae, Junho Ahn, BeomSeok Oh, and Yongdae Kim. LTESniffer: An Open-source LTE Downlink/Uplink Eavesdropper. In *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec '23*, pages 43–48, New York, NY, USA, June 2023. Association for Computing Machinery. doi: 10.1145/3558482.3590196.
- [28] Yanbin Hou, Xiaodong Yang, and Qammer H. Abbasi. Efficient AoA-Based Wireless Indoor Localization for Hospital Outpatients Using Mobile Devices. *Sensors*, 18(11):3698, November 2018. Number: 11 Publisher: Multidisciplinary Digital Publishing Institute. URL: <https://www.mdpi.com/1424-8220/18/11/3698>, doi: 10.3390/s18113698.
- [29] Jehn-Ruey Jiang, Chih-Ming Lin, Feng-Yi Lin, and Shing-Tsaan Huang. ALRD: AoA Localization with RSSI Differences of Directional Antennas for Wireless Sensor Networks. *International Journal of Distributed Sensor Networks*, 9(3):529489, March 2013. Publisher: SAGE Publications. doi: 10.1155/2013/529489.
- [30] Martin Kotuliak, Simon Erni, Patrick Leu, Marc Röschlin, and Srdjan Čapkun. {LTrack}: Stealthy Tracking of Mobile Phones in {LTE}. pages 1291–1306, 2022. URL: <https://www.usenix.org/conference/usenixsecurity22/presentation/kotuliak>.
- [31] Swarun Kumar, Ezzeldin Hamed, Dina Katabi, and Li Erran Li. LTE radio analytics made easy and accessible. In *Proceedings of the 2014 ACM conference on SIGCOMM, SIGCOMM '14*, pages 211–222, New York, NY, USA, August 2014. Association for Computing Machinery. URL: <https://dl.acm.org/doi/10.1145/2619239.2626320>, doi: 10.1145/2619239.2626320.
- [32] Shijie Luo, Matheus Garbelini, Sudipta Chattopadhyay, and Jianying Zhou. SNISGECT: A Practical Approach to Inject aNRchy into 5G NR. pages 5385–5404, 2025. URL: <https://www.usenix.org/conference/usenixsecurity25/presentation/luo-shijie>.
- [33] M. F. M. Mahyuddin, A. A. M. Isa, M. S. I. M. Zin, Afifah Maheran A. H., Z. Manap, and M. K. Ismail. Overview of Positioning Techniques for LTE Technology. *Journal of Telecommunication, Electronic and Computer Engineering (JTEC)*, 9(2-13):43–50, September 2017. Number: 2-13. URL: <https://jtec.utem.edu.my/jtec/article/view/2564>.
- [34] Robert Margolies, Richard Becker, Simon Byers, Supratim Deb, Rittwik Jana, Simon Urbanek, and Chris Volinsky. Can you find me now? Evaluation of network-based localization in a 4G LTE network. In *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, pages 1–9, May 2017. URL: <https://ieeexplore.ieee.org/document/8057097>, doi: 10.1109/INFOCOM.2017.8057097.
- [35] Taekkyung Oh, Sangwook Bae, Junho Ahn, Yonghwa Lee, Dinh-Tuan Hoang, Min Suk Kang, Nils Ole Tippenhauer, and Yongdae Kim. Enabling Physical Localization of Uncooperative Cellular Devices, March 2024. arXiv:2403.14963 [cs]. URL: <http://arxiv.org/abs/2403.14963>, doi: 10.48550/arXiv.2403.14963.
- [36] Kamari Radnosrati, Carsten Fritsche, Fredrik Gunnarsson, Fredrik Gustafsson, and Gustaf Hendeby. Localization in 3GPP LTE Based on One RTT and One TDOA Observation. *IEEE Transactions on Vehicular Technology*, 69(3):3399–3411, March 2020. Conference Name: IEEE Transactions on Vehicular Technology. URL: <https://ieeexplore.ieee.org/abstract/document/8966502>, doi: 10.1109/TVT.2020.2968118.
- [37] Alexander J. Ross and Bradley Reaves. Towards Simultaneous Attacks on Multiple Cellular Networks. In *2023 IEEE Security and Privacy Workshops (SPW)*, pages 394–405, May 2023. ISSN: 2770-8411. URL: <https://ieeexplore.ieee.org/document/10188622>, doi: 10.1109/SPW59333.2023.00040.
- [38] Joe Russo. Evaluating the Performance of Hand-held Cellphone Detectors in a Prison Setting. April 2016.
- [39] J Serrano, M Cattin, G Daniluk, E Gousiou, M M Lipinski, and M Lipiński. The white rabbit project.
- [40] Spacey. Is up-sampling prior to cross-correlation useless?, May 2012. URL: <https://dsp.stackexchange.com/q/2321>.
- [41] Bartosz Taudul. wolfpld/tracy, November 2025. original-date: 2020-03-17T19:43:04Z. URL: <https://github.com/wolfpld/tracy>.
- [42] Marcy Zora. THE REAL SOCIAL NETWORK: HOW JURORS' USE OF SOCIAL MEDIA AND SMART PHONES AFFECTS A DEFENDANT'S SIXTH AMENDMENT RIGHTS. *UNIVERSITY OF ILLINOIS LAW REVIEW*, 2012(2), 2012.

## A Detailed Features Description

As uplink signal features, we collect various signal power and signal quality features derived from the circular cross-correlation. The limiting factor in acquiring reference signals is the sampling rate of the signal, which translates into the same resolution of the cross-correlation. To increase this resolution, we upsample the resulting time domain cross-correlation and therefore infer information at

a higher sampling rate [40]. We use the Fourier method [5] for upsampling.


**A.0.1 Signal Power.** A naïve way to estimate the signal power is through  $\text{RMS}^2$  signal power measurement of the corresponding PRBs. However, it is influenced by unrelated signals on the same frequencies, such as uplink traffic from other UEs on other cells or other interference.

Instead, we use the interpolated correlation peak as an improved signal power measure. In Figure 9a, we ran 100 simulations for each SNR value to compare the interpolated correlation peak to the classical  $\text{RMS}^2$ . Especially in low SNR conditions, it outperforms naïve  $\text{RMS}^2$  signal power calculation as seen in Figure 9a. The correlation focuses on the target signal only and remains largely unaffected by noise, resulting in a more robust feature. In another simulation, we saw that the effect of a time delay on the signal power features is small and even with a time delay of  $5\mu\text{s}$ , the error magnitude is only around 0.5dB.

**A.0.2 Comparison of Uplink Message Types.** The measurement error of the features depends on the message type and its size. As shown in Figure 9b, the estimated signal power value based on correlation peaks is more accurate for PRACH and PUSCH with more PRBs than for PUCCH and PUSCH with only 1 PRB. This is because the PRACH includes the longest reference signal, allowing for more accurate correlation computations. In contrast, a PUSCH transmission with 1 PRB has only 2 reference signal sequences, each just 12 samples long, leading to poorer performance. A PUCCH transmission, despite spanning only 1 PRB, contains reference signals in 6 OFDM symbols, resulting in better performance than PUSCH with 1 PRB.

**A.0.3 Signal Quality.** To estimate signal quality, we use a more robust measurement technique compared to previous uplink sniffers [27, 30]. In their work, the circular cross-correlation is smoothed with a filter in the frequency domain to eliminate the impact of noise. The difference between the smoothed and non-smoothed version of the correlation is their resulting noise estimate. We call this technique *Smoothed Correlation SNR Estimation*. For our estimation technique, called *Correlation Peak-to-Average* measurement, we divide the peak value of the cross-correlation with the average value of the cross-correlation. The intuition is that the signal gets concentrated in and around the peak, while the noise is spread over the whole correlation. The comparison of the two features can be found in Figure 9c.

In 5G OFDM symbols allocated to the PUSCH reference signals, some subcarriers are left unoccupied. We can thus measure noise in the band from these empty subcarriers ( $N = \text{RMS}^2$  of empty subcarriers). On the other hand, the signal at the occupied subcarrier is the relevant signal plus the noise ( $P = S + N = \text{RMS}^2$  of occupied subcarriers). Then the SNR for these OFDM symbols can be calculated as  $P/N - 1$ .



(a) Estimation of the signal power of a 0dB signal using two different techniques.

(b) Estimation of the signal power of various messages sent with the same 0dB power with the Correlation Peak technique.

(c) Error of SNR estimation with two different techniques.

**Figure 9: Simulations of underlying features of our detection model. We varied SNR of the signal using AWGN channel. For each SNR, we ran the simulation 100 times.**