Radio Science

RESEARCH ARTICLE
10.1029/2025RS008237

Special Collection:
Radio Frequency Interference (RFI) 2024

TranQuiL: Long Range Detection and Localization of Interference in Radio Quiet Zones

Atul Bansal[1], Mohamed Ibrahim[1,2], Kuang Yuan[1], Yiwen Song[1], Bob Iannucci[1], and Swarun Kumar[1]

[1]Carnegie Mellon University, Pittsburgh, PA, USA, [2]Hewlett Packard Labs, Milpitas, CA, USA

Key Points:

- We present TranQuiL, the first long-range detection and localization system to enforce RQZ regulations, achieving an accuracy of 13.2 m
- We achieve a range of 950 m for WiFi and 450 m range for Bluetooth devices by developing an enhanced beacon packet detection pipeline
- We perform a detailed experimental evaluation of TranQuiL at an actual Radio Quiet Zone (RQZ), the Green Bank Radio Observatory

Correspondence to:
A. Bansal,
atul.bansal49@gmail.com

Citation:
Bansal, A., Ibrahim, M., Yuan, K., Song, Y., Iannucci, B., & Kumar, S. (2026). TranQuiL: Long range detection and localization of interference in radio quiet zones. *Radio Science*, *61*, e2025RS008237. https://doi.org/10.1029/2025RS008237

Received 1 FEB 2025
Accepted 7 DEC 2025

Author Contributions:
Conceptualization: Atul Bansal, Mohamed Ibrahim, Kuang Yuan, Bob Iannucci, Swarun Kumar
Data curation: Atul Bansal
Formal analysis: Atul Bansal
Funding acquisition: Bob Iannucci, Swarun Kumar
Investigation: Atul Bansal, Swarun Kumar
Methodology: Atul Bansal, Mohamed Ibrahim, Kuang Yuan, Bob Iannucci, Swarun Kumar
Project administration: Atul Bansal, Swarun Kumar

Abstract Radio Quiet Zones (RQZs) have been established to prevent radio sources from causing harmful interference to sensitive radio telescopes, which study extremely faint cosmic radio waves. Even with strict regulations, such interference is growing due to the widespread use of consumer electronics, emitting in many different frequency bands, including Wifi, Bluetooth. Removal of interferers is often a matter of sending trucks with spectrum analyzers to perform localization, using signal power-based localization techniques, a human-intensive process. We present TranQuiL, a novel long-range detection and localization system that can detect and localize an interfering transmitter at large distances. Our key innovation is the development of an improved beacon packet detection pipeline, which enables significant range improvement. We implement and evaluate our system for an interfering WiFi and Bluetooth transmitter across two testbeds: (a) the Green Bank Observatory in West Virginia and (b) around a large manufacturing facility in a major U.S. city. We demonstrate a localization accuracy of 13.2 m in both test beds from 950 m away for WiFi transmitters and 450 m for Bluetooth transmitters, sufficient for building-scale identification of the interferer's location.

Plain Language Summary Just as optical telescopes provide information about cosmic events by capturing images of astronomical objects, radio telescopes capture RF signals radiated by astronomical objects. These signals are crucial in conducting research on the astrophysical processes taking place in space. Since these signals are transmitted from outer space, they have very low signal power and thus—require very sensitive receivers on Earth to effectively capture the underlying information. However, over the years, we have observed an increasing proliferation of various terrestrial wireless communication technologies around the world. These transmitters are much closer to the radio telescopes relative to the astrophysical sources and thus their transmissions can corrupt the low-energy signals radiated from astrophysical sources. To prevent this, radio spectrum regulators around the world have established regulations that limit and prohibit the operation of wireless transmitters near radio telescopes in areas spanning hundreds of kilometers in diameter. These areas are called Radio Quiet Zones (RQZs). However, enforcing these regulations with the ever-increasing proliferation of wireless technologies is quite difficult. In the presence of an existing interfering wireless transmitter in an RQZ, radio telescope operators need to send out moving vehicles that travel across the whole RQZ while scanning for high powered signals, so that they can pinpoint the location of the interfering transmitter. Due to the limited range of scanning (about 100 m), it takes a long time to move around an RQZ area spanning hundreds of kilometers. In this article, we present a novel scanning mechanism that increases the range of scanning across various technologies. This enables radio telescopes to significantly reduce the overall search time required to pinpoint the location of the transmitter. This scanning mechanism exploits specific signal properties such as periodicity to improve the range. We evaluate the effectiveness of this scanning algorithm with WiFi and Bluetooth technologies as interference at two different locations: (a) An actual radio astronomy telescope located in West Virginia, the Green Bank Observatory, and (b) around a large manufacturing facility in Pittsburgh. We observe that we can pinpoint the location of interference with an accuracy of 13.2 m at a range of 950 m for WiFi (9.5x improvement) and 450 m for Bluetooth (4.5x improvement).

1. Introduction

Governments around the world have established Radio Quiet Zones (RQZs), spanning areas a few square kilometers to ensure no or minimal radio operation in an area sur-rounding radio astronomy observatories. This is done to minimize the radio frequency interference (RFI) to sensitive radio equipment at these radio telescopes, enabling continuous scientific research on cosmic radiations from space. As the number of IoT devices continues to increase, radio frequency interference (RFI) is expected to become more severe in the coming years. Even

Resources: Atul Bansal, Mohamed Ibrahim, Kuang Yuan, Yiwen Song, Swarun Kumar
Software: Atul Bansal
Supervision: Bob Iannucci, Swarun Kumar
Validation: Atul Bansal, Mohamed Ibrahim, Kuang Yuan, Yiwen Song
Writing – original draft: Atul Bansal
Writing – review & editing: Atul Bansal, Mohamed Ibrahim, Kuang Yuan, Yiwen Song, Bob Iannucci, Swarun Kumar

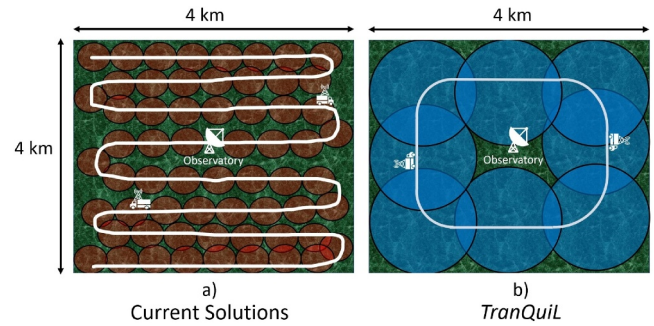

Figure 1. Area covered by the trucks for Detection and Localization of interfering transmitter: (a) Without TranQuiL, trucks have to move around for a longer time and (b) In the presence of TranQuiL, moving trucks can finish the localization with minimal movement.

though transmissions are power-limited or prohibited by law in RQZs (Cohen et al., 2003), they frequently experience stray interferers, such as WiFi access points, Bluetooth in cars, etc., often deployed due to lack of awareness or accidently (Bhattacharjee, 2010; Porko, 2011; RFI Scans, 2022).

This paper explores a scalable approach to detect and locate such errant interferers in RQZs. We focus primarily on WiFi and Bluetooth—two technologies operating on licensed spectrum that are highly common and challenging-to-address interferers to radio astronomy (Bhattacharjee, 2010; Porko, 2011). Indeed, while there is rich literature on WiFi and Bluetooth-based localization (Ayyalasomayajula et al., 2018; Bahl et al., 2000; Kotaru et al., 2015; Luo & Hsiao, 2019; Singh et al., 2018; Thompson et al., 2009; Wisanmongkol et al., 2019; Youssef & Agrawala, 2005), these are limited in localization range to 100 m (i.e., the traditional range limit of WiFi and Bluetooth Low Energy (BLE)). However, radio astronomy telescopes are known to be impacted by very low-powered interference, only 16 km away (Series, 2021) with astronomical signals having minuscule power levels in the order of 10^{-32} W/m^2. Thus, current approaches to locating such interferers in RQZs are fairly primitive. For example, the Green Bank Radio Observatory (Emberson, 1959) in West Virginia, one of the major radio observatories in the US, uses trucks with spectrum analyzers that drive around the facility in search of interferers based on RSSI values. This process is both time-consuming and inaccurate, especially over a large area (Mirror Article, 2022), primarily because the vehicle can only scan a short range around itself at any given time (see the left panel of Figure 1). Ideally, one would like an approach that further enhances the sensing range of these trucks that allows them to scan wide areas at each time and quickly pinpoint the interferer, saving unnecessary scouring of the environment (see Figure 1, right).

This chapter presents *TranQuiL* (TRANsmission QUIeting through Localization), a system that seeks to overcome the range limitations of traditional localization techniques with a focus on locating interferers of radio astronomy. To the best of our knowledge, *TranQuiL* is the first radio quiet zone enforcement system that performs long-range detection and localization of WiFi and Bluetooth sources, at the range of 950 m for WiFi and 450 m for Bluetooth. TranQuiL was evaluated in two wide-area testbeds: (a) The Green Bank Radio Observatory; (b) A large manufacturing facility in a major U.S. city. Across both testbeds, TranQuiL shows a median accuracy of 13.2 m (about building-scale).

At a high level, TranQuiL localizes interfering transmitters by measuring the difference in the time of arrival of its signal to a pair of receivers (each on different trucks)—that is, the Time Difference of Arrival (TDoA). Indeed, TDoA measurements can be tracked as the trucks move over time to assist in trilaterating the location of the interfering source. While TDoA is a well-studied technique (Jamali-Rad & Leus, 2013; Kaune, 2012; Yang et al., 2019), TranQuiL's main challenges stem from the significantly long range from which it needs to locate interfering transmitters. First, at such long ranges, packets from the WiFi and Bluetooth transmitter sources are simply undetectable and buried below the noise. Second, signal multipath becomes much more challenging to account for. We elaborate on these challenges below:

Packet Detection at Long Range: Prior to localization, one must first detect the transmitter packets. However, traditional packet detection algorithms require close proximity to the interferer, resulting in time-consuming searches. To this end, TranQuiL relies on transmissions that are very common across multiple wireless

technologies that are easier to detect at a longer range: advertising beacons. While beacon packets from each technology contain different and unique fields, most fields exhibit predictable patterns over time. This allows TranQuiL to essentially "brute-force" explore the received signals for these predictable fields, thereby increasing the preamble size. Furthermore, beacons are sent periodically—theoretically allowing for TranQuiL to improve beacon detection with time. However, due to clock imperfections, every beacon packet experiences random jitter around the expected time period of transmission, considerably degrading packet detection, especially when the received signals are below the noise floor. The key challenge, therefore, is to first align these beacons with one another within a reasonable timeframe and under the available computational constraints. To address this, TranQuiL builds on a signal processing technique used for time-series similarity applications called Dynamic Time Warping (DTW). If we divide the received signal into chunks each with the same length as the beacon period, the resultant groups should be very similar to each other (except for some jitter). TranQuiL then uses DTW to extract the most similar chunks among these signals and performs sample-level alignment of the chunks. Section 4 elaborates on how TranQuiL performs coherent combination of these aligned chunks to enhance beacon packet detection.

Long-Range Multipath Mitigation: Next, we consider a challenge common to localization systems—signal multipath. At long distances from an interfering transmitter source, multipath remains challenging, both due to the increased number of potential reflection sources and the signal attenuation of the direct path. TranQuiL addresses this challenge using two key properties—among the various signal paths as the receiving truck moves, the Line-of-Sight (LOS) path and reflections from dominant reflectors remains the most consistent and least impacted by changes in the environment. We leverage this consistency to render multipath sparse. Next, we use the terrain information at the receiver locations to further extract the LOS path from the sparse multipath present. Note that exploiting the consistency of LOS paths has been explored in the literature (Xiong & Jamieson, 2013; Wang & Katabi, 2013) to remove multipath from the AoA spectrum in indoor scenarios. In contrast, TranQuiL develops its own consistency metric for TDoA, tailored to the physical constraints of truck movement.

Evaluation at a Radio Quiet Zone: We have evaluated TranQuiL at Green Bank Observatory, a Radio Quiet Zone in West Virginia. Green Bank enforces a strict bar on deploying any transmitters around the radio observatory. To comply with this policy, we conducted TranQuiL's experiments during a limited 3-week window (its annual maintenance period) to evaluate TranQuiL, through careful coordination and with relevant permissions from the RQZ administrators. We implemented TranQuiL on a 2 sq. km testbed at the Green Bank Radio Observatory and a 1.33 sq. km testbed near a manufacturing facility in a major U.S. city. We used two USRP N210s spectrum analyzers and moved them around to emulate multiple distributed receivers, ensuring that they are GPS-synchronized in time and frequency. While our approach is general to any technology that uses beacon packets and can be extended to non-beacon packets, we primarily focus on WiFi and Bluetooth. We discuss how TranQuiL can be extended to non- beacon technologies as well. We implement TranQuiL with an ASUS AC2900 WiFi router as the interfering WiFi transmitter and the nRF5340 kit for Bluetooth. We show:

1. A median error in Time Difference of Arrival (TDoA) of 17.75 and 9.16 m respectively in the Green Bank observatory and manufacturing facility.
2. A median localization error of 17.4 and 9.1 m respectively in Green Bank and the manufacturing facility.
3. An increase of 9.5 × in the WiFi detection range and of 4.5 × in the Bluetooth detection range compared to state-of-the-art detection algorithms.

Contributions: The contributions of this paper are as follows:

1. To the best of our knowledge, our method implements the first long-range detection and localization system to enforce RQZ regulations that achieves an accuracy of 13.2 m.
2. Development of enhanced beacon packet detection at 950 m range for WiFi and 450 m for Bluetooth.
3. A detailed experimental evaluation at an actual Radio Quiet Zone (RQZ), the Green Bank Radio Observatory.

2. Primer on Radio Quiet Zones

From tracking the rotation of Earth to uncovering the secrets of Sagitarrius A (Remijan et al., 2002), the massive black hole at the center of our galaxy, radio observatories have been indispensable for scientific research. One such observatory in Green Bank, West Virginia is home to the world's largest 100-m fully steerable single-dish radio telescope (see Figure 2a).

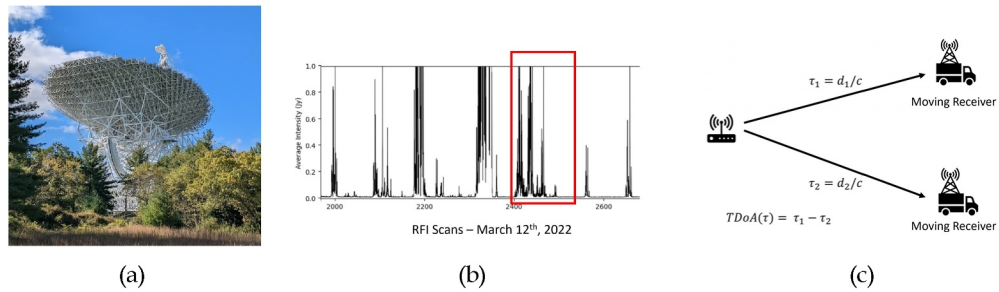

(a) (b) (c)

Figure 2. (a) Fully steerable single-dish radio telescope at Green Bank (b) RFI Scans in the S-band (Above 2 GHz) taken on 12 March 2022, publicly available on Green Bank website (c) Time Difference of Arrival (TDoA).

The need for Radio Quiet Zones (RQZs): A major challenge faced by radio observatories is Radio Frequency Interference (RFI) from WiFi, Bluetooth, etc. Astronomical radio waves are extremely weak at powers of 10^{-32} W/m^2 per Hz (RFI Scans, 2022). For reference, a typical WiFi router transmits power of about 10^{-9} W/m^2 at a 1 km distance, which is 23 orders of magnitude stronger than astronomical radio waves, rendering the latter unmeasurable. To alleviate these concerns, Radio Quiet Zones (RQZs) (at tens of kilometers radius) around observatories (Rfi Mitgation at Green Bank, 2023; Sizemore & Acree, 2002; West Virginia Radio Astronomy Zoning Act, 2022) highly restrict radio operations.

Why is Interference at the 2.4 GHz ISM band critical? Despite these regulations, enforcing them is quite challenging. As radio devices proliferate, radio observatories continually suffer from stray interferers. This is particularly true for ISM band technologies such as WiFi and Bluetooth, which are often deployed accidently or without awareness. Indeed, ISM interference is a bigger challenge compared to most other advanced radio communications (e.g., cellular) where accidental deployments are unlikely (see RFI Table:-(RFI Reports Table, 2023)). The 2.4 GHz ISM band and adjacent frequencies carry valuable data critical to key astronomical observations (Padovani, 2016; Vernstrom et al., 2016) that are irreparably damaged by terrestrial interference. Figure 2b demonstrates this with high interference at the 2.4 GHz ISM bands in RFI scans collected by the Green Bank Observatory. Indeed, several studies in the literature (Bhattacharjee, 2010; Porko, 2011) highlight how these technologies impede radio astronomy.

Current solutions are slow and inefficient: To address sporadic interference, radio observatories can pre-process the data to filter it out. However, persistent interference corrupts most or all data samples, necessitating a halt in data collection until the interfering source is located and mitigated. Typically, trucks equipped with spectrum analyzers and directional antennas scan the RQZ area by systematically traversing surrounding roads (see Figure 1) to detect signals from the interfering source. This process is time-consuming (taking hours) over a large area (Mirror Article, 2022), limiting the scanning range to about 100 m at any given time. Every minute of this slow search, during which the radio observatory remains shut down is a lost opportunity for multi-million- dollar scientific instruments. Thus, a solution is needed to extend the scanning range, allowing trucks to cover larger areas more quickly and to pinpoint the interfering source efficiently (see Figure 1).

3. System Overview

TranQuiL is a long-range interference detection and localization system, that addresses the range limitations of the observatory-owned spectrum scanning trucks. TranQuiL locates ISM interferers (e.g., WiFi access points, Bluetooth, etc.), saving valuable time for the radio observatory; the time needed for the advancement of scientific research. While applicable to any radio technology that broadcasts messages using beacons, TranQuiL primarily focuses on WiFi and Bluetooth. TranQuiL employs two mobile receivers and a dedicated wireless backhaul for inter-truck communication and performs detection and localization of the interfering source. To achieve this, TranQuiL leverages a well-established technique—Time Difference of Arrival (TDoA) (Gustafsson & Gunnarsson, 2003), which refers to the difference in the time of flights of a signal from the interfering source received by the two receiving trucks (see Figure 2c). Note that when the trucks perform their search using TranQuiL or using current solutions, the radio observatory must remain shut down. Thus, wireless backhauls between trucks do not affect the observatory data collection and are deactivated once the observatory resumes

operations. Multiple TDoA measurements, obtained as the trucks move to different locations, provide the spatial diversity necessary to trilaterate the location of the interfering source.

TDoA versus AoA: One may wonder why TranQuiL chose TDoA rather than Angle of Arrival (AoA) localization (e.g., antenna arrays or directional antennas). While AoA is well-suited for indoor localization, it incurs significant errors over long distances. A small AoA error of $\Delta\Theta = 0.05$ radians at a range $d = 1$ km results in a large localization error of $\Delta d = \Delta\Theta * d = 50$ m.

Assumptions: To ensure proper operation and compliance with RQZ regulations, TranQuiL is built upon several key assumptions that must hold true. First, TranQuiL requires the establishment of a wireless backhaul between the two mobile receivers to ensure accurate Time Difference of Arrival (TDoA) calculations. These backhauls must operate on frequencies that do not interfere with the radio observatory operations, thereby preventing potential disruptions. Additionally, these backhauls should be designed for rapid deployment and teardown, minimizing any delays in the localization of interference. Second, TranQuiL is designed to address unintentionally deployed interference; it is not equipped to handle intentionally adversarial interference scenarios. For instance, TranQuiL cannot mitigate cases in which an adversary has deliberately altered the beacon protocol specifications or modified the wireless signal modulation properties. Third, although TranQuiL is primarily designed for Non-Line-of-Sight (NLOS) scenarios, it assumes the presence of a Line-of-Sight (LOS) path in the received signal. This LOS path is extracted using TranQuiL's multipath mitigation algorithm to calculate TDoA. However, if the LOS path is obscured by noise, the accuracy of localization may be significantly reduced. Finally, TranQuiL utilizes USRP N210 devices in conjunction with an HP dual-core laptop with 16 GB of memory and a 1 Gbps Ethernet card for real-time collection of IQ samples. It is assumed that, in practice, IQ samples will be available to vehicle operators in real time, enabling them to execute TranQuiL's detection and localization algorithms to mitigate environmental interference.

Paper Outline: The rest of this paper explores two key challenges: (a) Section 4 studies the problem of overcoming the limited range of WiFi and Bluetooth—about 100 m. We specifically explore ways to leverage the structure of beacon packets of WiFi and Bluetooth radio technologies to enhance the detection range of WiFi to 950 m and of Bluetooth to 450 m; (b) Section 5 talks about how TranQuiL can be extended to non-beacon based technologies by exploiting the frequency domain properties of the signals (3) Section 6 explores a long-range localization system that processes detected packets to estimate the location of their sources. In doing so, it mitigates various challenges exacerbated at long range including signal multipath and synchronization. The paper concludes with system implementation (Section 7), evaluation (Section 8), and discussion (Section 9).

4. Beacon-Based Packet Detection

We present TranQuiL's enhanced packet detection technique that to improve range of typical localization systems.

4.1. Current Packet Detection Algorithms

Consider Figure 3 which represents the packet format of both WiFi and Bluetooth packets. At a high level, every packet in both technologies consists of a PHY layer preamble and payload that contains higher layer header fields and data bits. However, there are key differences in the structure of bits within both preamble and payload across both technologies. The preamble of WiFi contains repetitive patterns that are exploited by Schmidl and Cox algorithm (Schmidl & Cox, 1997). However, its performance degrades as the SNR decreases. Thus, there are cross-correlation-based approaches (Huang et al., 2022; Wang et al., 2020) which make packet detection more robust to noise. In Bluetooth, the preamble consists of alternating 1's and 0's, and existing approaches also use cross-correlation to perform packet detection. Figure 4 (see blue and red) shows the performance of existing WiFi packet detection algorithms. We observe that existing packet detection algorithms perform extremely poorly in detecting packets for long-range scenarios with SNR values as low as -10 dB.

Shortcomings: Why do current packet detection algorithms for both Bluetooth and WiFi fail at low SNR? Two reasons: First, both WiFi and Bluetooth were originally designed for indoor use typically operating at ranges of only a few tens of meters. Second, existing detection algorithms are intentionally simple, relying on only a few symbols to detect a packet. This is done aimed at minimizing computational complexity and power consumption in order to conserve battery life.

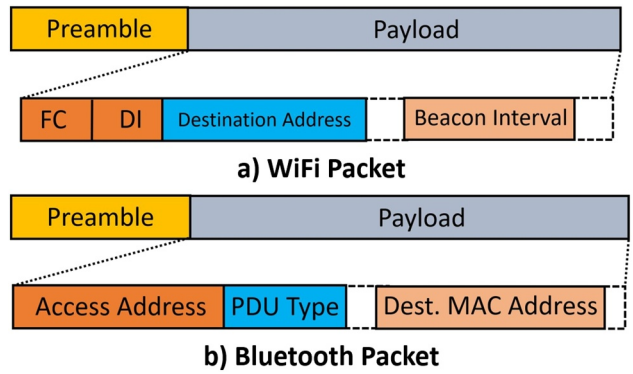

Figure 3. Packet Format of WiFi and Bluetooth beacons.

4.2. Exploiting Beacon Packets

We mentioned why existing detection algorithms fail in low SNR scenarios. This naturally leads to the question: "How do we significantly increase the detection range to detect an interfering transmitter source?." To understand our approach, we reflect on a key distinction between the state-of-the-art indoor WiFi/Bluetooth localization systems and our own.

Locating Access Points versus *Clients*: Much of the existing work focuses on the localization of WiFi/Bluetooth clients using uplink packets (Vasisht et al., 2016; Xiong & Jamieson, 2013; Zhao et al., 2018). In contrast, TranQuiL targets interferers in RQZs, where both the uplink and downlink packets are of interest. In fact, downlink packets with higher transmit power cause greater interference and thus facilitate easier localization of APs. Locating and turning off APs automatically mitigates interference from any/all client devices. Therefore, unlike conventional localization studies, TranQuiL prioritizes the identification and shutdown of APs.

Exploring Beacon packets for Localization: Instead of analyzing every packet, TranQuiL focuses on beacon packets periodically sent by APs to advertise their existence, as these packets are remarkably easier to detect owing to their low modulation rate and well-defined structure. To understand why, consider Figure 3. In both technologies, the "Payload" section can be decomposed into multiple fields that remain consistent over time with fields such as Frame Control, Duration ID, Destination MAC Address, Beacon Interval, etc. in WiFi and fields such as PDU Type, MAC addresses, Access Address, etc. in Bluetooth. This enables us to enlarge the preamble size for cross-correlation, thereby improving the detection range of the beacon packets. This can be observed in Figure 4 where we compare the detection rate of our extended preamble (see green) with state-of-the-art WiFi detection approaches (see red and blue). We observe an SNR gain of nearly 30 dB in packet detection in

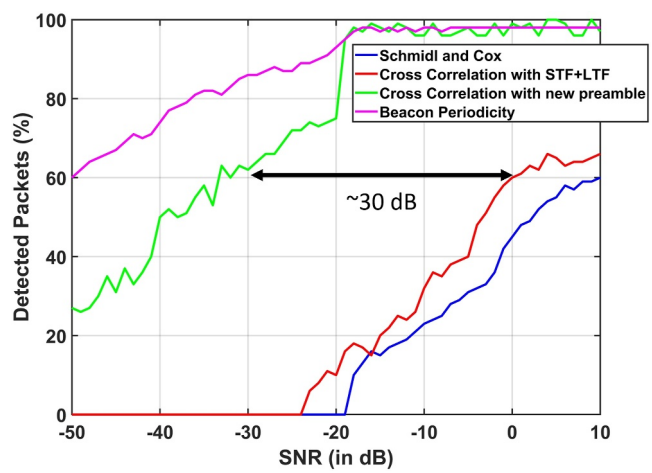

Figure 4. Comparison of packet detection performance of our approach as compared to the state-of-the-art packet detection algorithms in simulation.

simulation. To simplify exposition, we will henceforth primarily present techniques in relation to WiFi beacons. But, these techniques can be easily generalized to Bluetooth beacons as well.

Exploring beacon periodicity: However, a longer dynamic preamble is insufficient to perform detection at TranQuiL's target range. Therefore, we also exploit the beacon periodicity to enhance our detection range by looking for patterns of the beacon's *new preamble* repeated over time and add them together coherently. The major challenge here is that even though an interfering WiFi AP sends a beacon packet every 102.4 ms, the samples received are not exactly periodic to *high precision* due to clocking imperfections on WiFi APs leading to a chicken-or-egg problem that is, to perform detection, we require detection of the following beacons with high precision. A naıve solution would be to perform an optimization over the exact periods between consecutive packets and choose periods that maximize the detection metric. However, this approach suffers from exponential complexity, as each additional beacon packet introduces a new dimension to the search space, significantly increasing the computational burden for coherent summation.

To tackle this chicken-or-egg problem, we build on a well-known technique in signal processing to detect similarities between two temporal sequences—Dynamic Time Warping (DTW) (Muller, 2007). DTW is particularly effective at handling time warping and offsets, and is known for its robustness to noise. It also outputs the indices from two temporal sequences that are matched to one another. But, one might ask—"How does DTW, which requires two time-series as input, help us to resolve this dilemma?." We observe that because the received signal is nearly periodic, the period-sized chunks of this signal should be very similar to each other, assuming beacon packets are present in the signal. However, due to imperfect periodicity, there are some offsets between two chunks. This is where we use DTW to detect similarities between chunks and matched indices. DTW works well even for cases where beacon patterns are buried under noise, thus effectively mitigating the uncertainty introduced by timing perturbations between two consecutive beacon packets. Then, these matched chunks can be added coherently to increase the SNR of the beacon packet ensuring detection.

Mathematically, given the size n of signal, we first perform correlation with the long preamble. Then, we divide the result into multiple chunks equally, each with the length of a beacon period (p) (102.4 ms for WiFi beacons), leaving us with n/p chunks—a few of these chunks with beacon packets buried under noise and the rest of them with only noise. All the chunks with beacon packets should be very similar to each other and should have low DTW distance. Thus, to extract such chunks, we perform pairwise DTW among all chunks and sort them based on the DTW distance. This also outputs corresponding time sequences which are matched to one another. We then coherently add these sequences, until we observe the correlation peak exceed a certain threshold. Alg. 10 details the exact algorithm to exploit beacon periodicity to detect packet start sample. We observe that, in general, coherently adding 15−20 chunks results in detection, after which there are diminishing returns.

Algorithm 1. Beacon Periodicity Detection

Input: Correlated received signal—$C(t)$ with length n, peak threshold thresh and the beacon period p, $C_{final}(t) = 0$
Output: t_{start}

```
Divide C(t) into n/p equal chunks—C_i(t), i = 1.. n/p
for i = 1 to n/p do
  for j = i + 1 to n/p do
    dtwMatrix(i, j) = DTW(|C_i(t)|, |C_j(t)|)
(i_sorted, j_sorted) = Sort(Dist(dtwMatrix)) /* where Dist(dtwMatrix) captures pairwise time sequences' similarities */
for i, j in i_sorted, j_sorted do
  C_final(t) = C_final(t) + Match(dtwMatrix(i, j)) /* where Match(dtwMatrix(i, j)) matches the ith and jth sequence */
  if findPeaks(C_final(t) > thresh) then
    t_start = argMax findPeaks(C_final(t))
    break
```

Complexity: It is worth noting that the above algorithm does not explode exponentially in sharp contrast to the naıve optimization problem. For the following analysis, we assume that DTW has a linear runtime (e.g., FastDTW (Salvador & Chan, 2007)). Given the length n of the received signal and n/p chunks, the overall complexity for the algorithm would be:

$$O\left(\frac{\mathrm{n2}}{\mathrm{p2}}\right) * O(\mathrm{p}) + O\left(\frac{\mathrm{n2}}{\mathrm{p2}}\log\left(\frac{\mathrm{n2}}{\mathrm{p2}}\right)\right) + O\left(\frac{\mathrm{n2}}{\mathrm{p2}}\right) \approx O\left(\frac{\mathrm{n2}}{\mathrm{p2}}\right) \tag{1}$$

Which is polynomial in time complexity. Note that increasing the period p can further decrease the time complexity.

Effect of Frequency Offsets on WiFi beacons: To detect the start of the beacon packet, we mentioned that we simply added the matched DTW waveforms together to improve the correlation peak. However, due to the carrier frequency offset (CFO), the phase of beacon packets drifts over time, resulting in incoherent addition. Assuming that the CFO remains constant across all beacons, we mitigate this problem by iterating over a fixed set of possible CFOs and constructing a 2D correlation metric—one axis representing the sample index of a chunk and the other representing the set of CFO values. The maximum over this joint 2D metric provides us with the start sample of the beacon packets. Our assumption of CFO being constant across beacon packets holds fairly well since our packet detection algorithm only needs about 15–20 chunks to perform detection, which is about 2 s for WiFi beacons during which the CFO does not change more than ±30 Hz (Chen et al., 2019). Note that iterating over a set of CFOs only slightly increases the complexity of the algorithm and is much less than the exponential complexity of the naıve approach. Figure 4 shows the improvement we observe (see the magenta curve) by exploiting the WiFi beacon periodicity.

Effect of Frequency Offsets on Bluetooth beacons: Can the above approach be applied to eliminate CFO effects on Bluetooth beacons? Unfortunately, due to Bluetooth's frequency-hopping nature, each beacon experiences a random initial phase offset caused by phase-locked loop (PLL) locking, resulting in incoherent addition even after CFO correction. Thus, rather than performing detection individually at every receiver, we perform detection at relative signal measurements across two receivers (which are GPS synchronized). This relative channel, devoid of any frequency and phase offsets, completely eliminates the CFO correction step. The low bandwidth of Bluetooth ensures that packets across disparate GPS-synced receivers are received at the same/adjacent sample. This synchronization is not feasible for WiFi due to its wide bandwidth, however, WiFi transmissions are coherent and are transmitted at a single frequency.

What if the beacon period is not known? Thus far, we make a key assumption—the interfering source is accidently deployed and thus the beacon period of the transmitter source is well-known. This assumption is supported by the fact that about 60% of users of consumer-grade WiFi do not perform simple tasks such as updating firmware and changing passwords (Avast Threat Landscape Report, 2019). Hence, it is unlikely that most users would have the required expertise to change the beacon period.

But, how does TranQuiL tackle the remaining users or an adversary who deliberately changes the beacon period? To handle such cases, TranQuiL performs a window-based periodicity estimation before doing the actual detection. Using only the signal power, this algorithm gets a coarse estimate of the beacon period by leveraging the fact that noise power converges to a constant value as the sample size increases. In contrast, when a transmitter signal is embedded in the noise, the overall signal power exhibits non-converging behavior, revealing patterns corresponding to the signal's periodicity. TranQuiL exploits these patterns to estimate the period of the beacon signal.

Mathematically, TranQuiL iterates over multiple window sizes, calculates the average signal power in each window and observes its variation across all such windows. In cases when there is no signal present, the average signal power remains almost constant across all such windows and its deviation from the constant value decreases with increasing window size. This effect is captured in the standard deviation of average signal powers across all windows of a particular size, and is continuously decreasing. However, when there is a signal present, even when buried under noise, this standard deviation becomes convex and creates minima at the multiples of the period of the signal. This effect can be observed in the simulation study whose results are shown in Figure 5. The figure shows how with decrease in SNR, the standard deviation curve varies from a minima at the correct period value (∼50,000 samples) to become less convex at low SNRs, reaching its limit at −20 dB. At low SNRs, the period

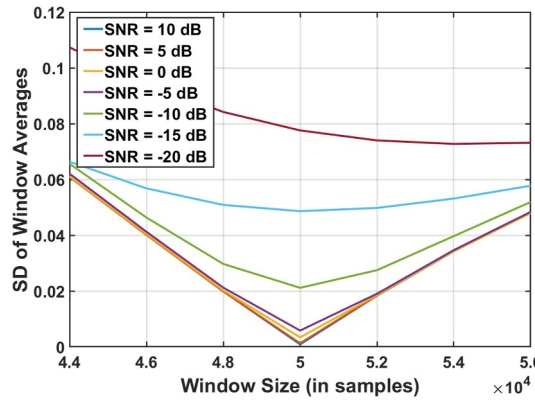

Figure 5. Simulation of our coarse period estimation algorithm across multiple SNRs from 10 to −20 dB.

estimation becomes too coarse and can lead to uncertain estimates. In such scenarios, the only option is to iterate over all the possible period configurations available to estimate the correct period value, albeit at a high computational cost.

Generalizing TranQuiL to other beacon-based technologies: TranQuiL primarily focuses on 2.4 GHz ISM band interferers such as WiFi and Bluetooth APs. However, the techniques mentioned in this chapter are generalizable to any technology that supports beacon transmissions. Table 1 shows examples of such technologies on unlicensed frequency bands that can support the operation of TranQuiL.

5. Generalizing TranQuiL to Non-Beacon Based Technologies

TranQuiL had made another assumption in its operation—beacon support. But there are many technologies that do not have support for beacons. For example, DVB-T lacks any beacon mechanism. Apart from this, even WiFi, LTE, 5G data packets do not satisfy the beacon requirements of packets being periodic.

So how does TranQuiL deal with interference from such technologies? Thus far, the detection range improvements achieved by TranQuiL have been due to the presence of key beacon characteristics — specifically, predictable field structures and periodic transmissions. In simple words, this translates to the temporal averaging of received signals to improve the overall SNR ratio. However, we have another degree of freedom that can help us achieve SNR gains—frequency domain information. Given we exploit frequency domain patterns of the transmitted signal at the receiver, we may be able to achieve better detected SNR at the receiver, thus resulting in enhanced detection performance.

Typically, these frequency domain patterns of a wireless signal are governed by the modulation scheme used by the transmitter. TranQuiL focuses on OFDM—one of the most common modulation schemes used by a variety of state-of-the-art wireless protocols due to its high spectrum efficiency and throughput. These include WiFi, LTE, 5G, Digital Video Broadcast—Terrestrial (DVB-T) transmissions. DVB-T protocols do not support any beacon packets in their transmissions, thus interference based on these protocols can't be detected based on the techniques we have mentioned until now. In this section, we explore how we can exploit frequency domain patterns of OFDM modulation to improve detection performance compared to the conventional cross-correlation approach.

Table 1
Beacon Periods of Different Technologies

Technologies	Frequency	Period
WiFi	2.4, 5 GHz	100.24 ms
Bluetooth	2.4 GHz	20 ms
LoRa-Class B	902–928 MHz	128 s
RFID	902–928 MHz	6 ms

5.1. Frequency Domain Properties of OFDM Signals

We explore two key properties that are very specific and unique to the frequency domain representation of OFDM signals. We discuss them below:-

1. *Frequency selectivity and orthogonality of wideband OFDM channel*:- For a wide bandwidth channel (tens of MHz of bandwidth), the channel response across frequencies changes significantly in power since the coherence bandwidth of wireless channels does not exceed a few MHz (Kristem et al., 2018). Thus, if we can detect this change across the channel response estimated from the received signal, we can infer that there was

some transmission. Conversely, if there is no change across the estimated channel response, we can infer that there is no transmission since noise power remains constant across a wide bandwidth under AWGN noise assumption. However, detecting this change in channel response is fairly difficult if the signal power is completely buried under noise. Thus, rather than estimating individual channel responses from each symbol, we observe the variation in Energy Spectral Density (ESD)—the channel response power across frequencies, averaged over all the symbols. Furthermore, this variation in average power can only be observed in OFDM signals, since the total symbol energy in each data subcarrier averaged over all OFDM symbols is constant. Mathematically, given that there is a signal present in the received signal:-

$$Y(k,f) = H(f)\,\mathrm{X}(k,f) + N(k,f) \tag{2}$$

where $f = -\mathrm{N}/2 + 1, \mathrm{N}/2$ and k = OFDM symbol number. The received signal power at kth OFDM symbol is given by

$$|Y(k,f)|^2 = (H(f)\,\mathrm{X}(k,f) + N(k,f))\,(H^*(f)\,X^*(k,f) + N^*(k,f)) \tag{3}$$

$$= |H(f)|^2\,|X(k,f)|^2 + |N(k,f)|^2 + H(f)\,\mathrm{X}(k,f)\,N^*(k,f) + H^*(f)\,X^*(k,f)\,N(k,f) \tag{4}$$

If we average this over K OFDM symbols where K is large, we can assume $E(N(k,f)) \sim 0$. Thus, the Energy Spectral Density (ESD) can be written as:-

$$E_k\left(|Y(k,f)|^2\right) = |H(f)|^2\,E_k\left(|X(k,f)|^2\right) + E_k\left(|N(k,f)|^2\right) \tag{5}$$

Assuming AWGN noise with noise power σ^2 and constant average signal power P_t at each subcarrier, the above expression can be written as:-

$$E_k\left(|Y(k,f)|^2\right) = |H(f)|^2\,P_t + \sigma^2 \tag{6}$$

If we plot $E_k(|Y(k,f)|^2)$ across f, we should see some variation owing to the range of f being more than coherence bandwidth. For the case, where no signal is present, the same expression can be written as which is constant.

$$E_k\left(|Y(k,f)|^2\right) = \sigma^2 \tag{7}$$

2. *Cyclostationarity*: OFDM signals have been shown to satisfy cyclostationarity properties (Heath & Giannakis, 1999). Cyclostationarity refers to a signal having statistical properties that vary cyclically with time. In an OFDM signal, few subcarriers only transmit fixed known symbols called pilot symbols. These pilot symbols are used to perform channel estimation to compensate for variations in channel response due to the Doppler shift with time. Furthermore, these pilot symbols are spread out uniformly across all subcarriers that is, periodic in the frequency domain. Figure 6 represents a simulation of how cyclostationarity property can be extracted from a WiFi 802.11b OFDM signal with no noise present. In this figure, we calculate the cyclic correlation (Lund'en et al., 2009) of an OFDM signal and we obtain peaks at the subcarrier numbers corresponding to the pilot subcarriers. The cyclic frequency axis represents the actual period and its multiples of the pilot subcarriers.

TranQuiL exploits both of these frequency domain properties to achieve improvement in the detection performance of wireless technologies that do not support beacon packets. However, each of these approaches has individual weaknesses. While the first approach can detect the presence of an OFDM transmission under low SNR, it can not provide the start sample of the packet, due to the received signal being averaged over a large number of samples. In contrast, the cyclic correlation approach can provide the start sample of the packet, but it fails to work reliably when the received signal is heavily dominated by a large number of noise samples. Thus, TranQuiL takes a hierarchical approach. First, it performs a coarse detection of OFDM packet by dividing the received signal into multiple windows and adopting the first approach to detect if a packet is present in a given window. Second, if a

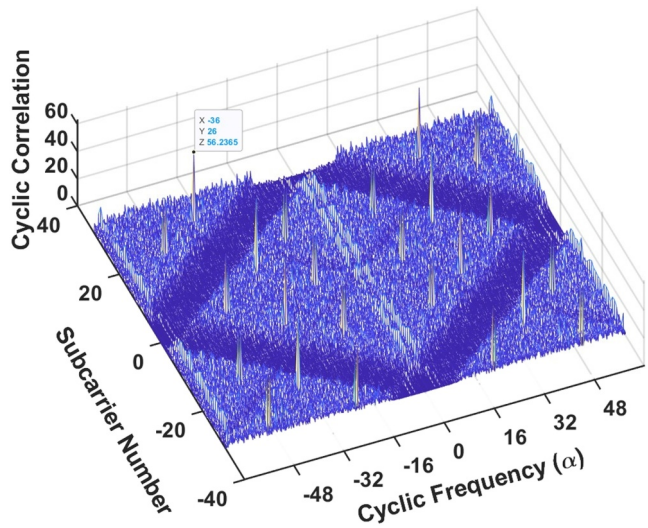

Figure 6. Cyclostationarity property observed through cyclic correlation.

packet is detected in a window, cyclic correlation is performed on the received signal within that window to detect the start sample of the packet.

6. Long Range Localization

In this section, we discuss how we calculate the Time Difference of Arrival (TDoA) across two receivers and then use multiple TDoA values across spatially diverse locations to find the interfering transmitter.

6.1. Base Station Synchronization

At long ranges, both our receivers must be synchronized in time and frequency to accurately measure TDoA. Hence, we use GPS clocks at both receivers to synchronize them in frequency and time with the clocks of GPS satellites. Thus, all the samples collected by both receivers have the same reference ($t = 0$) point in time and we obtain sample-level accurate I/Q measurements.

6.2. Mitigating Multipath

Consider two moving trucks and an interfering transmitter at a location of about a kilometer away. To perform accurate localization, TranQuiL needs the transmitter's Line of Sight (LOS) paths to both receivers, whose length difference corresponds to TDoA. However, achieving this is challenging because of dense multipath at such large ranges.

TDoA versus ToF? One may ask—Why do we not use Time of Flight (ToF) using a single receiver and move around to perform localization? To calculate the ToF, the time-of-departure of the signal at the AP needs to be known. However, since the transmitter is non-cooperative and doesn't provide this information, we use TDoA, which eliminates this requirement.

Why is multipath at long range challenging? At such long ranges, mitigating multipath is very difficult due to more opportunities for signal reflections, which lead to large variations in candidate TDoA values across all paths —ranging from a few meters to hundreds of meters. Almost all of the state-of-the-art WiFi localization systems (Gong & Liu, 2018; Kotaru et al., 2015; Xiong et al., 2015) leverage a sparsity assumption on multipath to extract the LOS path. However, this assumption breaks down at large ranges, where multipath is so dense and state-of-the-art techniques fail when directly applied on the received signals. To address this challenge, TranQuiL adopts a hierarchical approach:- First, eliminate multipath arising from fleeting reflectors and induce sparsity in multipath, containing only the effects of dominant reflectors. Then, we exploit the terrain information to extract the LOS path from sparse multipath.

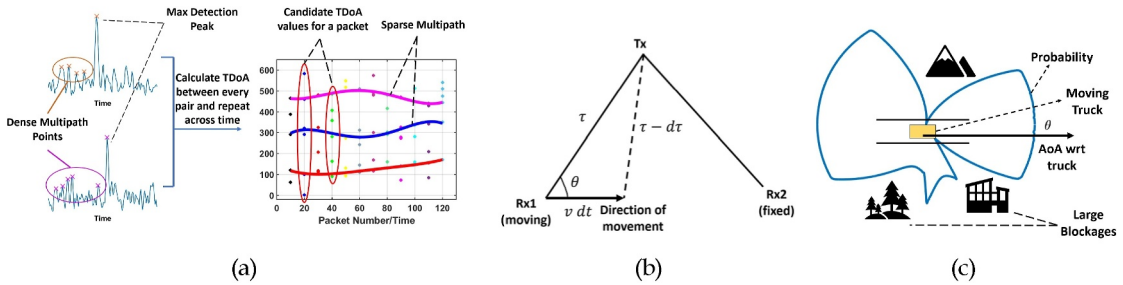

Figure 7. Multipath Mitigation: (a) Sparsify Multipath: We plot TDoA variation across time. Consistent TDoAs are assumed as dominant and sparse multipath (b) Angle of Arrival (θ) estimation (c) Maximum Likelihood Metric using Terrain information.

Rendering multipath sparse: To make multipath sparse, we make a key observation—Due to the movement of trucks, there is inherent temporal diversity in the packets received at different times. For all these packets, the LOS path along with paths from large reflectors such as mountains, buildings, trees etc. should remain *almost* consistent with time, with the variations bounded by the speed of the truck (e.g., 25 mph × 100 ms ≈ 1.2 m). However, other multipath peaks from fleeting reflectors are inconsistent and sporadically appear and disappear. We exploit this observation by plotting all the candidate TDoA values (see Figure 7a) with time and then performing multiple least square fits with the constraint that the slopes do not exceed the speed of the vehicle. TDoA values lying on this curve represent the TDoAs including the LOS path and reflections from large reflectors. Note that exploiting the consistency of LOS paths to mitigate multipath has been explored in the literature (Xiong & Jamieson, 2013; Wang & Katabi, 2013) for AoA-based systems. In contrast, TranQuiL develops its own consistency metric for TDoA by fitting a curve with a slope bounded by the speed of the vehicle, thus tying this metric to the physical constraints of the vehicle.

Exploiting terrain information: After obtaining multiple TDoA values, a natural question arises—"How does TranQuiL differentiate between the LOS path and the dominant multipath reflections that remain consistent with time?" To answer this, we first need to look at a key property of the TDoA curves we obtained. Consider Figure 7b where truck Rx1 is moving with a speed of v and Rx2 is static with τ as the distance between the source Tx and Rx1. As the truck moves, the distance between Rx1 and Tx after the time dt is $\tau - d\tau$. Assuming the original angle of the Tx with respect $\vec{v}$ is θ (the Angle of Arrival), we can write:

$$(\tau - d\tau)^2 = \tau^2 + (v.dt)^2 - 2.v.dt.\tau.\cos(\theta) \tag{8}$$

We then simplify the above (assuming dt and dτ are small): $\cos(\theta) = 1\ d\tau/v\ dt$—where dτ/dt is the slope of the TDoA curves. Thus, for every source Tx, either virtual or actual, we have an AoA estimate of the received paths. This is repeated with the second truck to get second receiver's AoA estimate.

We now create a maximum likelihood metric across 360° based on the terrain information at each individual receiver location. Consider Figure 7c, where we have a truck moving on a road surrounded by multiple blockages. We design a simple probability metric based on two key factors—the blockage probability of a reflector and the aperture angle of the reflector in the FoV of the receiver. Now, given the (AoA_1, AoA_2) pair for all the dominant paths across 2 receivers, we assign a likelihood to each of the dominant paths and take the TDoA of the (AoA_1, AoA_2) pair which has the highest likelihood as the LOS path.

TDoA for Bluetooth transmitters: Most of the concepts above apply to the TDoA calculation of Bluetooth as well with the primary difference being in their bandwidth. While WiFi uses wideband OFDM, Bluetooth employs narrowband GFSK modulation leading to a low resolution of TDoA. Fortunately, Bluetooth employs frequency hopping, and channels across frequencies can be stitched together to improve resolution. TranQuiL exploits this to emulate a wideband channel and improve the TDoA resolution.

6.3. Localization

Final Localization: After computing the TDoA values across multiple receiver pairs, localization is performed by tracing locus of constant TDoA, with the two receivers serving as the foci of a hyperbola. The intersection of multiple such hyperbolas provides us with the location of the transmitter. (see Figure 8a). We note that these

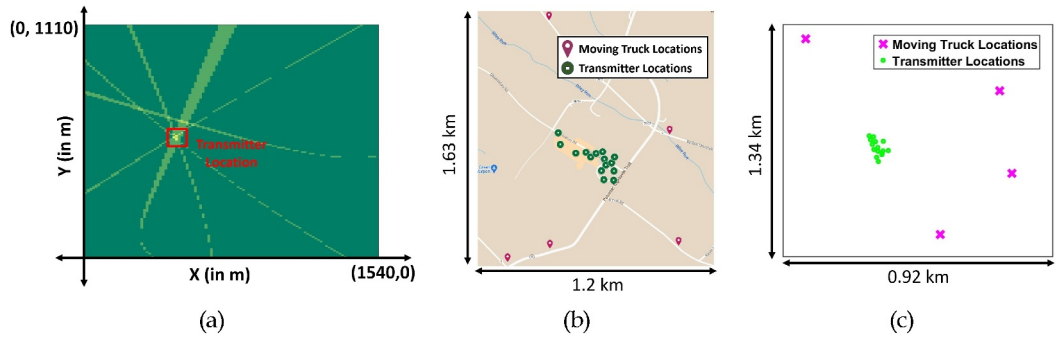

Figure 8. (a) Final Localization of the transmitter by the intersection of multiple hyperbolas at the Green Bank observatory testbed (b) Our testbed in the Green bank radio astronomy observatory (c) Our Manufacturing facility testbed in a major city (terrain removed for anonymity reasons).

techniques have been extensively studied in the literature (Ho et al., 2007). Modeling these hyperbolas require GPS estimates of the truck locations, and errors in these estimates influence our measurements. Various filtering strategies (e.g., Kalman filtering (Ullah et al., 2019)) can further reduce this error—although we present system accuracy without such filtering in our results to provide the readers a better sense of TranQuiL's raw error.

Impact of Dilution of Precision: The location of the base stations is an important factor that influences localization accuracy because of Geometric Dilution of Precision (GDOP) (Langley, 1999). A poor relative location of trucks (e.g., trucks very close to each other) can lead to hyperbolas whose intersection leads to a large error in the transmitters' location estimate. Thus, these trucks need to ensure proper GDOP with respect to the transmitter. Hence, we adopt the following truck movement route: As soon as a single truck detects the WiFi packet, one truck stops while the other truck moves around to ensure a large range of GDOP values.

7. Implementation and Evaluation

We implement *TranQuiL* on two Ettus USRP N210s emulating on-vehicle spectrum analyzers. Each USRP is kept in a car and moved around the testbed area to collect measurements. We use BG7TBL GPS Disciplined Oscillator (GPSDO) clocks to synchronize our 2 receivers in time and frequency using GPS signals. We also equip our receivers with omnidirectional WiFi 2.4 GHz antennas. We use an ASUS AC2900 WiFi router as the WiFi interferer and Nordic Semiconductor's nRF5340 Development Kit as the Bluetooth interferer to the nearby quiet zone. The collected data is processed to perform localization in MATLAB on the cloud.

Green Bank Observatory Testbed: Our first testbed (Figure 8b), is an actual Radio Quiet Zone (RQZ) in the Green Bank Observatory in West Virginia. As shown in Figure 8b, we moved our interferer (ASUS AC2900 WiFi router) to 15 different transmitter locations marked by green. For every transmitter location, we deploy 2 receivers on moving cars around a 2 km^2 area and the data is collected at all the possible pairs created by the 5 locations marked by red.

Careful Compliance at an RQZ: To prevent interference with the scientific data collected at the radio telescope from the deployed transmitter in our evaluation, we took relevant permissions from RQZ administrators and coordinated with them while evaluating. We performed our experiments during its annual maintenance window (3 weeks) when the observatory is shut down. We also set up a custom communication infrastructure to coordinate all of our experiments at Green Bank (since cellular is unavailable) that was taken down post the 3-week window.

Manufacturing Testbed: Our second testbed (Figure 8c) is in a manufacturing facility in a major US city. We deployed our receivers in a 1.33 km^2 area and the data is collected for all possible pairs with 4 different receiver locations marked by magenta for each of 15 different transmitter locations.

Ground Truth and Baseline: For ground truth, we use an off-the-shelf U-blox EVK- 7P GPS receiver (<1 m precision in an open sky setup). We note that the error in the GPS location estimates of the moving receivers affects our measurements. We also place an Ettus USRP N210 within the 100 m the interferer as a ground truth spectrum analyzer, which can detect the interferer packets using standard packet detection. To evaluate our systems' localization performance, we use two baselines:- state-of-the-art WiFi localization system SpotFi

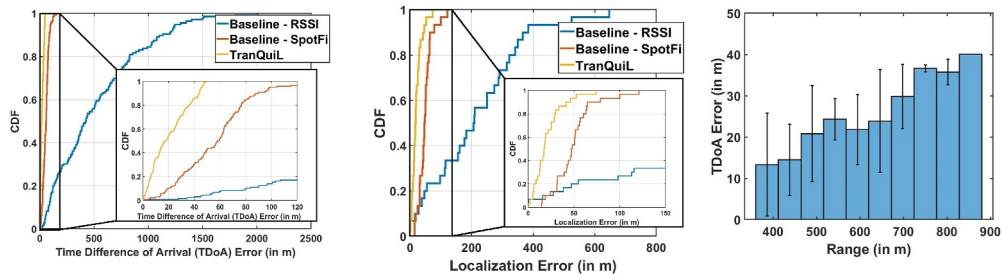

Figure 9. Green Bank Observatory testbed: (a) CDF of the TDoA Error of TranQuiL versus SpotFi and RSSI baseline (b) CDF of the Localization Error of TranQuiL versus SpotFi and RSSI baseline (c) Variation of TDoA with Range of TranQuiL.

(Kotaru et al., 2015) and a localization system purely based on RSSI. Both the baseline systems need prior information about the start of the packet to perform localization which we provide using the techniques mentioned in Section 4.

8. Experimental Results

8.1. Time Difference of Arrival Accuracy

Method: We evaluate our system with 30 different WiFi interferer locations in both testbeds. Note that because of our wide area testbeds, all the locations are in non-line of sight (NLOS). Therefore, we implement our multipath mitigation algorithm to extract the line of sight (LOS) path from both the receivers and estimate the TDoA.

Results(a): **Green Bank Testbed** Figure 9a depicts a 17.75 m median error and 99th percentile error of 48 m in estimating the TDoA of interfering WiFi source. We significantly outperform both SpotFi and RSSI baselines with median errors of 58 and 402.6 m respectively. It is worth noting that the CDF tail is significantly longer for SpotFi due to its poor multipath resolution capability in the presence of dense multipath at long ranges.

Results(b): **Manufacturing Facility Testbed** Figure 10a depicts a 9.16 m median error and 99th percentile error of 46.7 m in estimating the TDoA of interfering WiFi source. The median error is better than at Green Bank due to the relatively wider area of the Green Bank testbed. We significantly outperform the 2 baseline systems—SpotFi and RSSI, which provide median accuracies of 53.2 and 376.9 m respectively.

8.2. Localization Accuracy

Method: We evaluate TranQuiL across the two testbeds, with the WiFi interferers located at over 30 locations in both the testbeds. After estimating the TDoA across multiple receiver pairs, we perform localization of the interfering WiFi source.

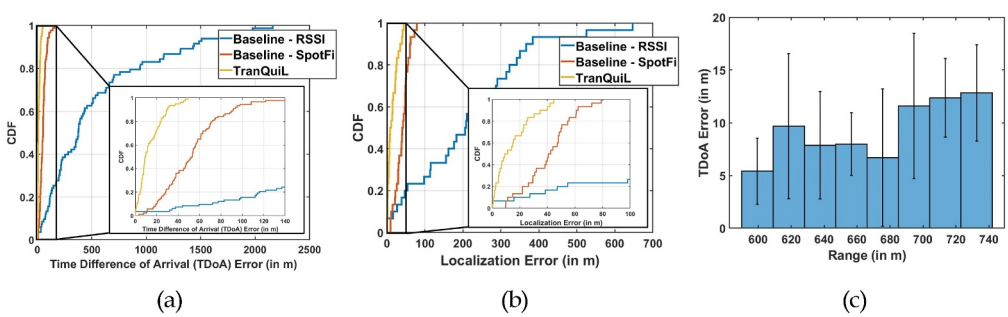

Figure 10. Manufacturing Facility testbed: (a) CDF of the Time Difference of Arrival Error of TranQuiL versus SpotFi and RSSI baseline (b) CDF of the Localization Error of TranQuiL versus SpotFi and RSSI baseline (c) Variation of Time Difference of Arrival Error with Range of TranQuiL.

Results(a): **Green Bank Testbed** Figure 9b depicts 17.38 m median error for TranQuiL in localization of the interfering WiFi source across 5 receiver locations with 99th percentile error of 52.5 m. We note that TranQuiL significantly outperforms the localization error of both baseline systems which had a median error of 47 m (about 3×) and 348.2 m (about 20×) respectively under identical settings.

Results(b): **Manufacturing Facility Testbed** Figure 10b depicts a 9.11 m median error and a 99th percentile error of 41.9 m for TranQuiL in the localization accuracy. Consistent with our previous observations, the localization accuracy of our manufacturing testbed is better than the Green Bank testbed because of the relatively wider area testbed at Green Bank with dense natural obstructions. However, we still significantly outperform both the baseline localization systems with a median error of 40.60 m (about 4×) and 205.92 m (about 22×) respectively. Note that the reported accuracy for both testbeds suffices to locate the building or a house where the interfering WiFi source lies, and thus is sufficient for enforcing radio quietness in the RQZ.

8.3. TDoA Error Variation With Range

Method: We calculate TDoA from all the receiver pairs at 30 different interferer locations in both testbeds. We calculate the range of all pairs by taking the minimum distance of both receivers with the transmitter. Also, we have removed the TDoA outliers with very high errors (>80th percentile) to understand the trend with range more clearly since most outliers have their LOS paths completely buried under noise.

Results(a): **Green Bank Testbed** The general trend in the TDoA error is increasing with the range (see Figure 9c). It varies from 13.3 m TDoA error at 400 m range to 40.07 m TDoA error at 850 m. This is due to the low SNR of signal received at larger distances, worsening our detection algorithm (Sec. 4), thus leading to errors in TDoA calculation. We observe that around the 600 m range, the TDoA error is lower than in other ranges. We attribute this to the high SNR of the signal obtained at that distance which may have happened due to constructive fading.

Results(b): **Manufacturing Facility Testbed** Similar to the Green Bank testbed, the TDoA error (see Figure 10c) is increasing with range because of decrease in SNR. It varies from 5.4 m TDoA error at 600 m range to a 13 m TDoA error at 730 m. As above, there are certain irregularities in the TDoA error at 620 m range and at 680 m range because of SNR variations obtained resulting in destructive and constructive fading respectively. In contrast with Green Bank experiment, the TDoA error is lower at the manufacturing testbed even at longer ranges because of the difference in terrain and major obstructions.

8.4. Detected Packet SNR Versus Range

Method: We calculate the received signal SNR at all receivers by taking the ratio of the peaks obtained **after performing detection** using 20 chunks and the noise power level. Therefore, we only detect packet if the detection peak is above noise (∼0 dB SNR). We obtain the SNR across all the receivers in both testbeds and plot its variation with the range.

Results: In Figures 11a, 11c and 11b, the SNR of both WiFi and Bluetooth signals decreases with range. For WiFi, it varies from 7.8 dB median SNR at 380 m to 2.46 dB median SNR at 930 m across both testbeds with no packet detected at distances >950 m. For Bluetooth, it varies from 13.7 dB median SNR at 170 m to 0.12 dB median SNR at 450 m. For WiFi transmitters in indoor environments, the median SNR varies from 3.8 dB at 50 m to 0.1 dB median SNR at 750 m. Thus, we conclude that 950 and 750 m are the maximum ranges of our system for WiFi outdoors and indoors respectively. For Bluetooth, the maximum range is 450 m.

8.5. Coarse Period Detection Versus Range

Method: To evaluate the coarse period detection algorithm explained in Section 4.2, we observe the error in detecting the correct beacon period across all signals. We plot the percent error (wrt the beacon period) across all ranges—that is, the distance of the receiver from the transmitter.

Results: In Figure 11d, the error in detecting the beacon period value increases from <5% to >10% after the 775 m range. With more than 10% uncertainty, iterating over the DTW-based detection algorithm across all candidate period values becomes computationally expensive. Thus, even without knowing the beacon period value, TranQuiL operates up to 775 m range—at the expense of some computation and reduction in the maximum 950 m range.

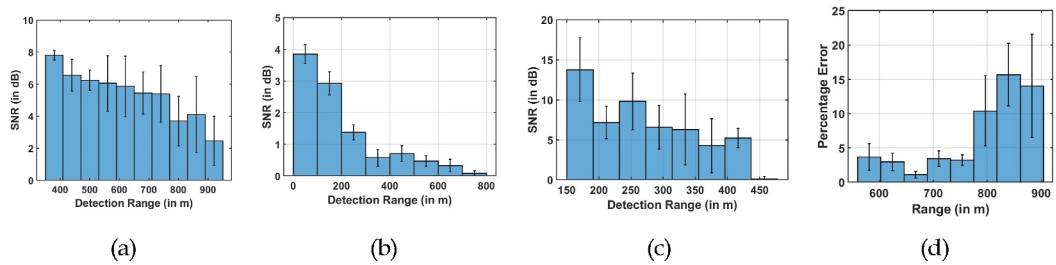

Figure 11. Microbenchmarks: (a) Variation of SNR (in dB) of the received WiFi signals in outdoor and (b) indoor environments and (c) Bluetooth signals respectively after coherent combining beacons across the distance from Tx to Rx (d) % Error in detecting the Period across the distance from Tx to Rx.

8.6. Percent of Detected Packets Versus SNR—Non-Beacon Based TechnoloGies

Method: To evaluate the effectiveness of frequency domain characteristics in performing the detection of non-beacon-based OFDM technologies, we perform a simulation study to plot the percentage of detected WiFi packets with cross correlation as baseline and frequency domain characteristics along with cross-correlation.

Results: Figure 13 represents our simulation results. We observe that we obtain about/sim8 dB of improvement in detection performance by adding the frequency domain characteristics of OFDM signals into the detection pipeline. The improvement is not as good as using beacon periodicity, due to the limited number of subcarriers (64 to be exact for a 20 MHz WiFi 802.11b signal)—thus limiting the frequency domain gains we can achieve.

8.7. Localization Error Versus No. of Receivers

Method: We evaluate the localization performance with the total number of receiver pairs used to calculate TDoA. We vary the number of receiver pairs from 2 to 5 to compute TDoA values and use them to perform localization.

Results: In Figure 12c, the localization error decreases with the increase in base station pairs. This is due to the increasing spatial diversity, thus enhancing our localization performance. Adding another pair only has a limited incremental effect at high spatial diversity, which can be observed in Figure 12c, where error doesn't improve after 3 pairs.

8.8. Detection Latency With Range

Method: To evaluate the latency gains from TranQuiL's enhanced detection, we perform a simulation study using the actual road network around Green Bank (see Figure 12a). First, we choose 15 transmitter locations at random in an 8 km × 5 km area around the observatory. For each transmitter, we sample 20 different routes which a truck can take to perform detection, ensuring that the truck starts moving from the observatory. We measure detection times for each route and get an average detection time at a transmitter location. We perform this for 2 cases:- (a) A moving truck with a 100 m range and (b) with a 1 km range.

Results: Figure 12b shows the variation of detection time with the transmitter distance from the observatory. The detection time significantly improves using TranQuiL's detection range of 1 km with 6 × latency improvement at

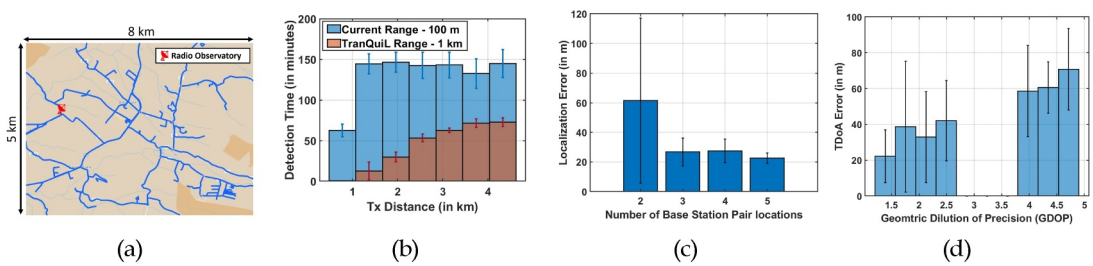

Figure 12. (a) Truck routes around the Radio Observatory to perform detection (b) Detection time using conventional methods with 100 m range (in blue) and TranQuiL's approach with 1 km range (in red) (c) Variation of Localization Error with the no. of base station pairs (d) Variation of TDoA Error with GDOP (No data points between 2.7 and 3.8).

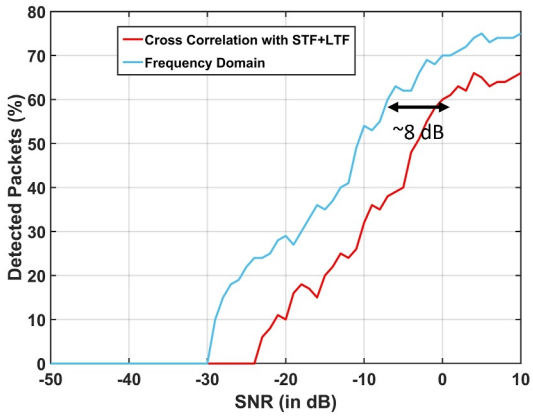

Figure 13. Comparison of packet detection performance of frequency domain properties of OFDM signals as compared to simple cross-correlation in simulation.

2 km range. All the transmitter locations are chosen close to actual roads so that the baseline can detect them. For other locations, the baseline would fail to perform detection that is, infinite detection time. Note that the baseline detection time almost remains consistent with transmitter distance. The primary reason is that a truck with a 100 m detection range, needs to be very close to the transmitter, and with randomized routes, the average detection time for every transmitter location is almost the same. We observe similar saturating trends in TranQuiL's curve as well at large transmitter distances.

8.9. Geometric Dilution of Precision

Method: As discussed in Section 6, the Geometric Dilution of Precision (GDOP) plays a very important role in the accuracy of the localization. In our scenario, we calculate GDOP across all the pairs of receivers and plot it against their respective TDoA Errors. Note that the lower the GDOP value is, the better the accuracy.

Results: Figure 12d represents the TDoA error variations with respect to GDOP. As expected, the TDoA error increases with the increase in GDOP. Unfortunately, for all the receiver locations we collected data from, we do not have any data samples with GDOP values between 2.7 and 3.8 due to topographic constraints. Thus, we observe a hole with no data points in this region.

9. Discussion and Limitations

Extending to technologies other than WiFi and Bluetooth: TranQuiL is designed to primarily tackle technologies whose signals demonstrate some form of repetition or redundancy in time—for instance, beacons repeated periodically. We can observe this property in various other technologies as well apart from WiFi and Bluetooth. We mention such technologies in Table 1 along with the beacon periods they support. The concepts and techniques used in TranQuiL to detect WiFi and Bluetooth interferers can be extended to these technologies, provided that the preamble is known up front as well as its periodicity. Along with beacon-based technologies, we also generalize TranQuiL to non-beacon based technologies as well, as shown in Section 5 by exploiting frequency domain patterns. Such frequency domain patterns are common in technologies that primarily use OFDM signals.

In short, TranQuiL generalizes to technologies that either have: (a) periodic beacons; *or* (b) frequency domain cyclostationarity (e.g., in OFDM). Put together, this allows TranQuiL to generalize in detecting and localizing a wide variety of wireless technologies that act as interferers in an RQZ.

However, it should be noted that the incorporation of known information on the interference signal enhances the range of detection. For example,— if a known RFI repeats itself consistently every few hours, TranQuiL can be extended to use this periodicity as the beacon period for its detection. The greater the amount of prior knowledge available about a signal, the greater the detection performance and range. Although not used in TranQuiL, we also envision support from the radio observatory in providing some of the prior information to the trucks going out in the environment, owing to their highly sensitive radios that are capable of detecting very low-powered signals.

We believe there is rich potential for future work that investigates ways to build databases of patterns that exist in RFI. For example, one can train machine-learning based approaches where using a large data set of classified RFI. Once trained, this model can be used to extract the time-domain and frequency domain patterns, that can be used by TranQuiL to perform detection and localization. This approach could potentially extend to RF interference whose general structure is less reliably known up front, such as from microwave ovens, wireless controllers, etc.

Multiple Interfering WiFi transmitters: We assume that there is only a single interfering source present in the environment. How would TranQuiL handle multiple interfering sources? TranQuiL can exploit the beacon periodicity information to differentiate multiple transmitters. It is very unlikely that two interfering sources start transmitting at the same time. Thus, the respective beacons transmitted would be separated by a certain offset in every chunk and can be differentiated easily.

Extreme Multipath/Occlusion: While TranQuiL does try to mitigate the effect of multipath, there can be cases where the LOS path is severely attenuated and cannot be measured due to deep occlusions. In such cases, our multipath mitigation algorithm would fail and would lead to large errors.

Transmitter Mobility: We have assumed that the interferer transmitter is static while the trucks are trying to localize it. This assumption may not hold true in cases such as WiFi hotspots, Bluetooth devices in moving cars, etc. However, even with moving transmitters, LOS and high reflection paths still remain consistent across time and TranQuiL's consistency metric can readily extract them.

Cost of deploying TranQuiL: It may seem that deploying two moving trucks to localize an interfering source is quite expensive. However, TranQuiL was designed to exploit the already existing resources at radio observatories with no extra cost. Typically, these observatories already have 2–3 trucks available (with mounted receivers) to detect interferers. Thus, we incur no extra cost to the observatory. We note that an alternative variant of TranQuiL may replace trucks with drone swarms carrying lightweight SDRs or arrays of narrowband radios, however, may be relatively limited in flight range, sensing range and/or payload weight/complexity.

Improvements in Range and Accuracy: The range of TranQuiL across various technologies is limited by the prior information we have while performing detection. In case of WiFi and Bluetooth, even though the period of Bluetooth beacons is shorter than WiFi beacons, the preamble size of Bluetooth is much smaller than WiFi, leading to small processing gain in the detection performance, thus leading to smaller detection range. To improve the range of Bluetooth interference detection, TranQuiL needs the presence of additional time domain or frequency domain patterns. The accuracy obtained by TranQuiL's localization is primarily dependent on mitigating multipath experienced by the interfering signal. To improve accuracy, we need more dimensions in which we can resolve multipath. To this effect, having more antennas on a single truck or having more trucks deployed concurrently to perform detection can improve the localization accuracy.

Uplink transmissions from Client to AP: One may wonder, how uplink transmissions from the clients to Access Points affect TranQuiL. The presence of client transmissions do not affect TranQuiL's performance due to the non-periodicity of client transmissions. Uplink client transmissions are generally governed by Carrier Sense protocols (CSMA) and are much weaker than beacon transmissions. For our manufacturing facility testbed, we had ambient WiFi traffic and TranQuiL still performed fairly well.

10. Summary

This chapter presents TranQuiL, a long-range interference detection and localization system for Radio Quiet Zones, which detects and locates WiFi packets at about a kilometer range and Bluetooth packets at about 450 m range. We develop an enhanced packet detection pipeline that leverages periodicity of beacons. Then, we perform localization by exploiting path consistency to mitigate the effects of multipath. We deploy TranQuiL on two testbeds—one at an actual RQZ (Green Bank observatory) spanning two sq. km. area and another at a manufacturing facility spanning 1.3 sq. km. area in a major city in the US. We observe an overall median error of 13.2 m across both testbeds. We believe that extending this system to detect and localize other non-beacon-based technologies remains a problem to be addressed in future work.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

The IQ data received in our evaluation of TranQuiL at Greenbank Radio Observatory are available at TranQuiL IQ Samples repository at Kaggle. These IQ samples are collected from various transmitter and five different receiver locations situated at distances ranging from 300 m to about a kilometer. We used our Beacon Detection pipeline on these IQ samples to detect packets at long ranges and perform localization at the end. The data set is available at https://www.kaggle.com/datasets/atulbansalcmu/tranquil-iq-samples/ data with a Creative Common license with Non-Commercial Sharelike (CC BY-NC-SA) and must be cited with (Bansal et al., 2025) if used in future research. Our code for evaluating TranQuiL can be found at www.github.com/atul-bansal/TranQuiL and can be cited with (Bansal, 2025).

Acknowledgments
We thank NSF (2030154, 2106921, 2007786, 1942902, 2111751, 2433903), ONR, DARPA-SOAP, Cylab-Enterprise and MFI for their support.

References

Avast threat landscape report. (2019). Avast threat landscape report. Retrieved from https://cdn2.hubspot.net/hubfs/486579/Avast_Threat_Landscape_Report_2019.pdf

Ayyalasomayajula, R., Vasisht, D., & Bharadia, D. (2018). Bloc: CSI-based accurate localization for BLE tags. In *Proceedings of the 14th International Conference on emerging networking experiments and technologies* (pp. 126–138).

Bahl, P., & Padmanabhan, V. (2000). Radar: An in-building Rf-based user location and tracking system. In *IEEE Infocom* (Vol. 2, pp. 775–784). https://doi.org/10.1109/infcom.2000.832252

Bansal, A. (2025). Tranquil. *Github*. Retrieved from https://github.com/atul-bansal/TranQuiL

Bansal, A., Ibrahim, M., Yuan, K., Song, Y., Kumar, S., & Iannucci, B. (2025). Tranquil IQ samples [Dataset]. *Kaggle*. https://doi.org/10.34740/KAGGLE/DSV/10628518

Bhattacharjee, Y. (2010). *Radio astronomers take arms against a sea of signals*. American Association for the Advancement of Science.

Chen, Y., Su, X., Hu, Y., & Zeng, B. (2019). Residual carrier frequency offset Estimation and compensation for commodity WiFi. *IEEE Transactions on Mobile Computing*, *19*(12), 2891–2902. https://doi.org/10.1109/tmc.2019.2934106

Cohen, R., Delgado, G., Hardy, E., Hasegawa, T., & Nyman, L. Å. (2003). Radio-quiet zones. *Light Pollution: Global Views*, 225–259. https://doi.org/10.1007/978-94-017-0125-9_21

Emberson, R. M. (1959). National radio astronomy observatory: The early history and development of the observatory at green bank, West Virginia, are reviewed. *Science*, *130*(3385), 1307–1318. https://doi.org/10.1126/science.130.3385.1307

Gong, W., & Liu, J. (2018). Roarray: Towards more robust indoor localization using sparse recovery with commodity WiFi. *IEEE TMC*, *18*(6), 1380–1392. https://doi.org/10.1109/tmc.2018.2860018

Gustafsson, F., & Gunnarsson, F. (2003). Positioning using time-difference of arrival measurements. In *ICASSP* (Vol. 6, p. VI–553).

Heath, R. W., & Giannakis, G. B. (1999). Exploiting input cyclostationarity for blind channel identification in OFDM systems. *IEEE Transactions on Signal Processing*, *47*(3), 848–856. https://doi.org/10.1109/78.747790

Ho, K., Lu, X., & Kovavisaruch, L.-O. (2007). Source localization using TDOA and FDOA measurements in the presence of receiver location errors: Analysis and solution. *IEEE Transactions on Signal Processing*, *55*(2), 684–696. https://doi.org/10.1109/tsp.2006.885744

Huang, Y., Yuan, L., & Gong, W. (2022). Research on IEEE 802.11 OFDM packet detection algorithms for household wireless sensor communication. *Applied Sciences*, *12*(14), 7232. https://doi.org/10.3390/app12147232

Jamali-Rad, H., & Leus, G. (2013). Sparsity-aware multi-source TDOA localization. *IEEE Transactions on Signal Processing*, *61*(19), 4874–4887. https://doi.org/10.1109/tsp.2013.2272288

Kaune, R. (2012). Accuracy studies for TDOA and TOA localization. In *IEEE fusion* (pp. 408–415).

Kotaru, M., Joshi, K., Bharadia, D., & Katti, S. (2015). Spotfi: Decimeter level localization using WiFi. In *ACM SIGCOMM*.

Kristem, V., Bas, C. U., Wang, R., & Molisch, A. F. (2018). Outdoor wideband channel measurements and modeling in the 3–18 GHZ band. *IEEE Transactions on Wireless Communications*, *17*(7), 4620–4633. https://doi.org/10.1109/twc.2018.2828001

Langley, R. B. (1999). Dilution of precision. *GPS World*, *10*(5), 52–59.

Lunden, J., Kassam, S. A., & Koivunen, V. (2009). Robust nonparametric cyclic correlation-based spectrum sensing for cognitive radio. *IEEE Transactions on Signal Processing*, *58*(1), 38–52.

Luo, R. C., & Hsiao, T.-J. (2019). Indoor localization system based on hybrid Wi-Fi/Ble and hierarchical topological fingerprinting approach. *IEEE Transactions on Vehicular Technology*, *68*(11), 10791–10806. https://doi.org/10.1109/tvt.2019.2938893

Mirror Article. (2022). Mirror article. Retrieved from https://www.mirror.co.uk/news/world-news/telescope-town-can-hear-within-6155014

Muller, M. (2007). Dynamic time warping. Information retrieval for music and motion (pp. 69–84).

Padovani, P. (2016). The faint radio sky: Radio astronomy becomes mainstream. *Astronomy and Astrophysics Review*, *24*(1), 13. https://doi.org/10.1007/s00159-016-0098-6

Porko, J. P. G. (2011). *Radio frequency interference in radio astronomy*. (Unpublished Master's Thesis). Aalto University.

Remijan, A., Snyder, L. E., Liu, S., Mehringer, D., & Kuan, Y. (2002). Acetic acid in the hot cores of Sagitarrius b2 (n) and w51. *The Astrophysical Journal*, *576*(1), 264–273. https://doi.org/10.1086/341627

Rfi mitgation at green bank. (2023). Rfi mitgation at green bank. Retrieved from https://greenbankobservatory.org/science/gbt-observers/rfi-mitigation/

RFI reports table. (2023). RFI reports table. Retrieved from https://safe.nrao.edu/wiki/bin/view/GB/Projects/RFIReportsTable

RFI Scans. (2022). RFI scans. Retrieved from https://greenbankobservatory.org/rfi-scans-and-known-sources/

Salvador, S., & Chan, P. (2007). Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, *11*(5), 561–580. https://doi.org/10.3233/ida-2007-11508

Schmidl, T. M., & Cox, D. C. (1997). Robust frequency and timing synchronization for OFDM. *IEEE Transactions on Communications*, *45*(12), 1613–1621. https://doi.org/10.1109/26.650240

Series, R. (2021). *Characteristics of radio quiet zones*. (Tech. Rep.). Technical Report, International Telecommunication Union (ITU).

Singh, H., Sarkar, S., Dimri, A., Bhaskara, A., Patwari, N., Kasera, S., et al. (2018). Privacy enabled crowdsourced transmitter localization using adjusted measurements. In *PAC* (pp. 95–106).

Sizemore, W., & Acree, J. (2002). *The national radio quiet zone*. Lewis and EmerSon. 217

Thompson, R.J., Balaei, A.T., & Dempster, A.G. (2009). *Outdoor localization of a WiFi source with unknown transmission power*. Gold coast, Australia.

Ullah, I., Shen, Y., Su, X., Esposito, C., & Choi, C. (2019). A localization based on unscented Kalman filter and particle filter localization algorithms. *IEEE Access*, *8*, 2233–2246. https://doi.org/10.1109/access.2019.2961740

Vasisht, D., Kumar, S., & Katabi, D. (2016). *Decimeter-level localization with a single WiFi access point*. Usenix NSDI.

Vernstrom, T., Scott, D., Wall, J., Condon, J., Cotton, W., Kellermann, K., & Perley, R. (2016). Deep 3-GHZ observations of the Lockman hole north with the very large array–II. Catalogue and μjy source properties. *Monthly Notices of the Royal Astronomical Society*, *462*(3), 2934–2949. https://doi.org/10.1093/mnras/stw1836

Wang, G., Guoyong, L., Zhang, F., Ye, J., & Guojia, X. (2020). Research and implementation of cross-correlation Symbol synchronization algorithm for long training sequence based on OFDM system. In *IOP Conference Series: Materials Science and Engineering*, (Vol. 782(3), p. 032115), https://doi.org/10.1088/1757-899x/782/3/032115

Wang, J., & Katabi, D. (2013). Dude, where's my card? Rfid positioning that works with multipath and non-line of sight. In *Proceedings of the ACM Sigcomm 2013 Conference on Sigcomm* (pp. 51–62).

West Virginia radio astronomy zoning act. (2022). West Virginia radio astronomy zoning act. Retrieved from http://www.wvlegislature.gov/wvcode/code.cfm?chap=37a&art=1

Wisanmongkol, J., Klinkusoom, L., Sanpechuda, T., Kovavisaruch, L.-O., & Kaemarungsi, K. (2019). Multipath mitigation for RSSI-based Bluetooth low energy localization. In *2019 19th International Symposium on Communications and Information Technologies (ISCIT)* (pp. 47–51).

Xiong, J., Sundaresan, K., & Jamieson, K. (2015). Tonetrack: Leveraging frequency-agile radios for time-based indoor wireless localization. In *ACM mobicom* (pp. 537–549).

Xiong, J., & Jamieson, K. (2013). Arraytrack: A fine-grained indoor location system. In *Usenix NSDI*.

Yang, M., Jackson, D. R., Chen, J., Xiong, Z., & Williams, J. T. (2019). A TDOA localization method for nonline-of-sight scenarios. *IEEE Transactions on Antennas and Propagation*, *67*(4), 2666–2676. https://doi.org/10.1109/tap.2019.2891403

Youssef, M., & Agrawala, A. (2005). The Horus WLAN location determination system. In *ACM mobisys* (pp. 205–218). https://doi.org/10.1145/1067170.1067193

Zhao, M., Li, T., Abu Alsheikh, M., Tian, Y., Zhao, H., Torralba, A., et al. (2018). Through-wall human pose estimation using radio signals. In *IEEE CVPR* (pp. 7356–7365).