Can't Hide Your Stride: Inferring Car Movement Patterns from Passive TPMS Measurements

Yago Lizarribar[§] *armasuisse*
Alessio Scalingi *Universidad Carlos III de Madrid*
Domenico Giustiniano *IMDEA Networks*
Pedro Miguel Sánchez Sánchez *University of Murcia*
Roberto Calvo-Palomino Universidad Rey Juan Carlos
Gérôme Bovet armasuisse
Vincent Lenders University of Luxembourg

***Abstract*—Tire Pressure Monitoring System (TPMS) transmissions of modern cars are sent over the air in clear text and entail a unique identifier that does not change over very long periods of time. In this work, we investigate the privacy implications for car owners of this design choice by collecting and analyzing TPMS transmissions from a network of low-cost spectrum receivers that we deploy along the road over a period of 10 weeks. Our measurement study comprises data from 12 verified cars, but malicious actors could easily scale their efforts to track several thousands of cars, given that we observed at least 20k cars during our measurements. Our results show that TPMS transmissions can be used to systematically infer potentially sensitive information such as the presence, type, weight, or driving pattern of the driver. The affordability of the equipment to cause these threats, as low as $100 per receiver, urges policymakers and car manufacturers to design a more secure and privacy-preserving TPMS for future cars.**

I. Introduction

Cars are increasingly being geo-tracked nowadays. For example, car manufacturers continuously monitor the position of vehicles sold on cellular networks for improved maintenance and optimization purposes. Navigation services such as Google Maps follow the position and velocity of car drivers' smartphones to improve their GPS navigation services with traffic jam updates. Road authorities rely on the identification of cars with cameras for road billing and fines.

Although continuous car tracking is becoming ubiquitous, it also comes with privacy risks. Car movement data entails a lot of private and sensitive information. Movement data provides insights into the everyday private lives of their owners. For example, in a recent data breach that affected four large European car brands[1], researchers were able to systematically observe the activities of police officers, military officials, or individuals visiting medical facilities, raising serious privacy concerns for those affected entities.

Given the importance of the security and privacy of car-related data, new regulations have been put in place. For example, in July 2022, 54 countries worldwide (including all EU countries and other OECD nations) approved new cybersecurity regulations, requiring traded cars to have a cybersecurity certificate [2], [3]. This certification implements the United Nations Regulation No. 155 for the vehicle cybersecurity management system [4]. Furthermore, the United Nations' regulation lists various cyberthreats to vehicle data, such as identity and location information.

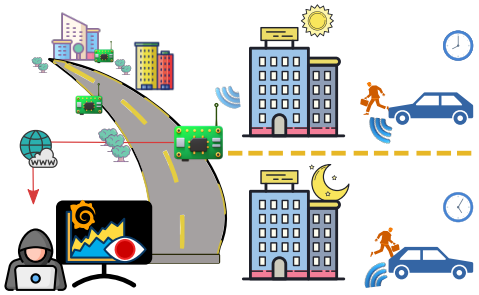

Fig. 1: This paper demonstrates how a network of low-cost spectrum receivers can capture and analyze the movement patterns of vehicles, inferring car owners' routines.

Perhaps surprinsingly, the Tire Pressure Monitoring System (TPMS) is not yet considered in this certification process. Originally designed for vehicle safety and maintenance, the TPMS system has now become mandatory in vehicles worldwide [5], [6], [7]. The system provides the owner of the vehicle with real-time information on tire pressure by having tires transmit regular wireless beacons to the electronic control unit of the vehicle. Although not providing the exact location of the tire or the car, researchers have discovered that most TPMS sensors transmit a unique identifier in clear text that never changes during the lifetime of the tire[5], [8]. This unencrypted wireless communication makes the signals susceptible to eavesdropping and potential tracking by any third party in proximity to the car.

In this paper, we investigate this risk and whether it is possible to infer car movement patterns from such passive TPMS measurements, as shown in the conceptual scenario presented in Figure 1. Previous research [5], [8] has highlighted that TPMS signals can be intercepted up to 40m from the car, but it is not clear whether they can enable extensive movement profiles due to their weak and sporadically transmission.

To address this question, we deploy five receivers along the road in our neighborhood over a period of 10 weeks receiving more than 6 million TPMS messages from more than 20,000 vehicles. We then develop algorithms to match the signals from different tires to the same car and to infer car movement from

[§]This work was carried out while the author worked at IMDEA Networks

ISBN: 978-3-903176-79-9

a subset of this data comprising verified cars. Our results show that TPMS transmissions can be used to systematically infer potentially sensitive information such as the presence, type, weight or driving pattern of the driver. The affordability of the equipment and the ease of deployment to cause these threats is as low as $100 per receiver, urging car manufacturers and policymakers to reconsider the use of unencrypted wireless transmissions with a fixed identifier for the transmission of TPMS messages in future cars.

In summary, our key contributions are as follows.

- **Feasibility**: We show the practical viability of collecting and analyzing TPMS data using a network of low-cost Software-Defined Radio (SDR) receivers deployed in a real setting.
- **Accurate vehicle identification**: We introduce techniques to reliably associate TPMS signals from individual tires with specific vehicles.
- **Uncovering Driver Patterns**: We employ pattern matching and data mining to reveal insights into driver routines and assess the potential for tracking and surveillance.
- **Novel data set**: We provide a novel TPMS dataset with 12 verified cars that span 10 weeks.

II. Primer on TPMS systems

When measuring tire pressure, TPMS systems are divided into 2 categories: Indirect TPMS (iTPMS) and Direct TPMS (dTPMS). The former measures wheel pressure through speed sensors (typically from the ABS system) and tracks the linear and rotational speed variation over time. When the difference in one of the wheels exceeds some predetermined thresholds, the system reports a pressure loss. There exist a variety of methods to determine the pressure [9], but they are regarded as low precision because they do not directly measure the pressure of the wheel. These systems are favored by companies in the Volkswagen Group (VW, Audi or Skoda).

On the other hand, dTPMS systems utilize sensors embedded in the rim of the wheel, or on the inner liner of the tire, to continuously directly measure its pressure and temperature [10]. dTPMS makes use of battery-powered pressure sensors to acquire the tire pressure of a vehicle. As there is no open dTPMS standard, dTPMS communication is based on proprietary protocols and simple modulation schemes (for instance, ASK or FSK at 315 MHz or 433 MHz). Each TPMS sensor transmits the information about tire pressure and identifier to the Electronic Control Unit (ECU), located inside the vehicle. Manufacturers like Toyota, Renault, Hyundai or Mercedes typically favor these systems over iTPMS.

The length of the message varies depending on the manufacturer, but the dTPMS messages usually are 100 bits in length with a symbol rate of 20 kbps, which means that a full transmission takes 5 ms. A TPMS message contains the following fields: **preamble**, with common hex structures like `0x55555556` or `0xaaaaaaa9`; textbfID, which is a 24 to 32 bit hexadecimal string with the dTPMS sensor identifier; **temperature**; **pressure**; **different flags** which may contain parameters like battery status; and **a checksum**.

Although the operation of these devices varies among different manufacturers, most of them transmit pressure information when sudden changes in tire pressure are detected or when the vehicle is moving. When the vehicle starts to move, motion sensors trigger the pressure sensor to start data transmission with a period of 30-120 seconds. Another way to trigger transmissions from a TPMS sensor is to send a pulse in the LF band (125 kHz), which is the operation principle of many TPMS monitoring tools used in repair shops, such as the one shown in [11].

As an example, [12] shows the modes of operation of a dTPMS sensor from Schrader Electronics. It defines up to nine modes of operation for these sensors, but the ones we are more interested in are *Drive*, *Stationary* and *ID Response* modes. The first sends up to eight messages in 60 seconds (1 message every 7.5 seconds on average) while the car is moving. The second mode sends messages every hour when the car is stopped. The last is a response to the trigger in the LF band that can be used to verify the sensor IDs of a car. As we will show in this work, this is not the case for all brands, as some transmit at high rates even when the car is stopped.

Because dTPMS send data wirelessly without encryption or obfuscation (other than by using proprietary, yet simple, communication protocols), we conjecture that they might pose several threats to driver privacy. Therefore, the focus of this paper is on these systems, and we will use the terms TPMS and dTPMS interchangeably.

III. Threats of a TPMS surveillance system

TPMS transmissions **are sent without any encryption or secure mechanisms and include a unique identifier**. This allows anyone with affordable equipment like a low-cost spectrum receiver and a standard off-the-shelf antenna to capture and track them throughout time and space. In this section, we explore how these transmissions could be exploited and what information an attacker could gain by gathering TPMS measurements.

For our analysis, we consider the following passive tracking system with low-cost spectrum receivers as a threat model. At a high level, an entity could set up an architecture like the one shown in Fig. 1, with receivers distributed over a geographical area, continuously capturing tire messages. These devices do not need to be expensive, as RF dongles such as the RTL-SDR attached to a Raspberry Pi have enough capability to collect messages continuously, as it has been shown in many other applications [13]. Apart from ID filtering, the rest of the components are open-source and easily available online, making this kind of architecture easy to deploy and affordable to anyone. Once the architecture is in place, the threats are manifold.

Malicious users could deploy passive receivers on large scales and track citizens without their knowledge. The advantage of such a system, over more traditional camera-based ones, is that no direct line-of-sight is needed with the TPMS sensors and spectrum receivers could be placed in covert or hidden locations, making them harder to spot by victims. A

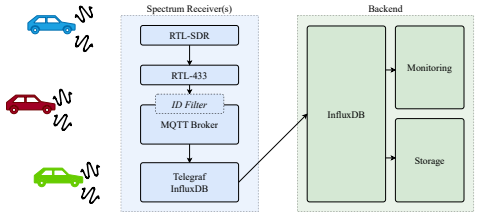

Fig. 2: The software architecture used in this work is composed of low-cost RTL-SDR platform publicly available TPMS decoder. The data pipeline and monitoring interface are also composed of open-source software.

data mining company could deploy receivers, gain insight on the types of traffic and routes taken, and then sell that data, all without the knowledge of the users (the drivers).

By establishing such a network of spectrum devices over a city, malicious users could track cars and infer behavioral patterns. In fact, another type of attack that results from passive surveillance could be for burglars in suburban residential areas. By tracking the vehicles of each household, they could infer the schedule and pattern of a particular household and take advantage of their absence. Finally, by combining passive monitoring with active spoofing, threat actors could track logistics trucks, spoof flat tire alerts to force stops, and then hijack the cargo.

This work does not explore active attacks (as in [5]), but the scenario presented in Fig. 1 could be extended to also include spectrum transmitters targeting specific TPMS sensors obtained through passive monitoring, for example with the objective of depleting the battery of TPMS sensors. However, this paper uncovers the significant warnings associated with such wireless transmissions, including potential threats in tracking vehicle patterns.

IV. System Deployment

In order to collect data for our research, we implemented a TPMS monitoring system consisting of five low-cost spectrum receivers.

A. Data collection setup

We position a total of 5 spectrum receivers along the perimeter of a workplace. The total area covered is the blue region of Fig. 3b and the surrounding area, which is approximately 100×50 m^2. The setup consists of an RTL-SDR connected to a Raspberry Pi, which is cost-effective while having a small form factor (see Fig. 3a). The equipment necessary to collect these data is as low as 100 dollars per spectrum receiver. Internet connectivity depends on the location of the deployment, with some places offering Ethernet or Wi-Fi, while mobile internet connections is used in more remote locations. In addition, if Internet connectivity is unavailable, data is recorded locally on the device and collected periodically.

In order to create a realistic deployment, we place the spectrum receivers indoors pointing at the parking and the nearby roads, and close to windows. The deployment is presented in Fig. 3b. The distances range from $10 - 50$ m, and we show later that it is sufficient to capture many transmissions.

We use the open source `rtl433` software for message decoding [14]. It includes around 250 different decoders operating at 433MHz, 28 of them related to the TPMS sensors of different manufacturers. `rtl433` provides various output formats such as JSON, MQTT, or InfluxDB. We opt for a combination of the MQTT and InfluxDB, so that we can reliably send data to our backend and for its flexibility in terms of managing discontinuous data, tools to deal with aggregate data and large data sets, and optimization for time series data, which makes it ideal in our scenario. In order to support data collection, we set up a centralized InfluxDB collecting data from all receivers, where we create a database per each. We show a general overview of the architecture in Fig. 2.

B. Maximum Distance Estimation

Although TPMS sensors are expected to transmit at low power, we are able to receive TPMS transmissions at distances greater than 50 m.

To test and validate that our receiver deployment is capable of obtaining enough tire pressure information, we measured Car 3 (`C3`) transmissions at different distances and conditions. The scenario is illustrated in Fig. 3c. We trigger the transmission of all the tires of the car and measure the number of messages and the Signal-to-Noise-Ratio (SNR) at the 7 positions of the map (green circles), with spectrum receivers deployed outdoors, with Line-of-Sight (LOS) conditions with respect to the car's tires. The farthest position is around 55 m. In this scenario, in all positions, we can receive all messages with a SNR of at least 10.0 dB.

For the second part of the test, we set our measurements in Non Line-of-Sight (NLOS) conditions (we place the spectrum receiver inside a building, without any window nearby), and we are able to capture and decode TPMS measurements in 3 out of the 4 selected locations, with a SNR of at least 8.0 dB. These results show that it is feasible to install receivers at relatively large distances from roads and still obtain many car tire measurements with low-cost equipment.

Therefore, we conclude that our deployment is suitable for measuring TPMS transmissions in the area shown in Fig. 3b as the distances involved are well within limits. Clearly, with a better antenna specifically designed for 433 MHz, it could be possible to reach even greater distances under both LOS or NLOS conditions.

C. Capturing TPMS Signals from Moving Vehicles

To evaluate signal reception from vehicles in motion, we conduct experiments using a car equipped with a dTPMS system. The experimental setup consists of a sensor positioned approximately 20m from the road, with a narrow field of view oriented toward the test route. The vehicle completes 10 laps around a 1km circuit encompassing the building block, maintaining speeds up to 50 km/h.

We are able to capture 32 TPMS messages across the 10 laps, with successful signal reception in 9 out of 10 laps as shown in Figure 4. This corresponds to an average of 3.5 transmissions per successful lap. These findings suggest that a

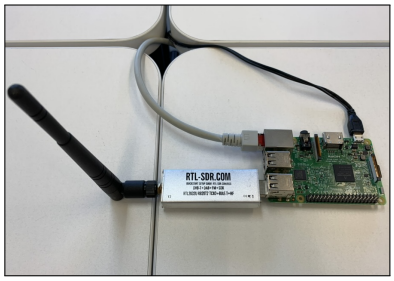

(a) Spectrum receiver setup.

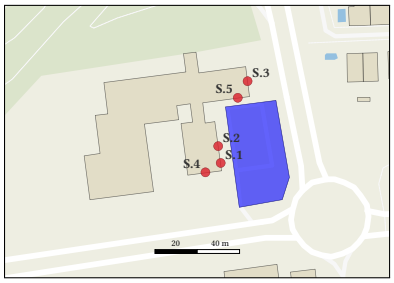

(b) Receiver deployment.

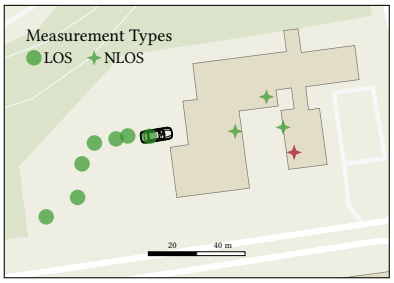

(c) Positions used to test TPMS range.

Fig. 3: Selected receiver and deployment site. We prepared five spectrum receivers around the selected area (Fig. 4b). The receivers (red) are placed indoors, and the violet area represents the parking lot where most measurements are focused. In Fig. 4c, we show the location of receiver positions used to test the range of TPMS messages. Marked circles indicate the position of spectrum receivers with LOS conditions. Marked stars indicate scenarios with NLOS conditions.

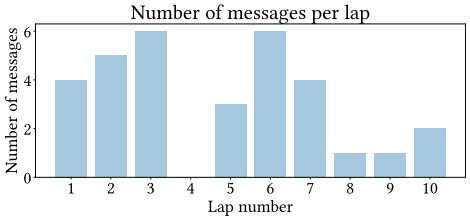

Fig. 4: Successfully received TPMS messages per lap from the experimental moving vehicle.

receivers network deployed throughout a city could potentially enable car tracking and surveillance infrastructure as cars will likely move around the speeds set in this experiment.

Our results show no significant bias in signal reception based on vehicle orientation, indicating that while directional determination may not be feasible, the relatively low signal attenuation by the vehicle body enables consistent message capture regardless of the transmitter's position. This omnidirectional characteristic enhances the probability of successful signal capture each time a vehicle passes.

V. Dataset

We collect data for our experiments from cars of 12 individuals. To verify the IDs of each TPMS sensor we use the Autel MaxiTPMS TS501 [11]. This device is able to transmit the required 125 Hz LF trigger and capture the response of the TPMS sensor back. At the same time, we also capture the response of each of the TPMS sensor with a laptop to verify that the IDs decoded by the software library `rtl433` coincide. Once we collect all TPMS IDs, we add them to a whitelist in each of our receivers so that only those are sent to the backend.

To protect the privacy of users, we will not show the real IDs of the cars used in this study. We provide the software used in this research in the following repository[1].

A. Receiver Sites

We conduct our measurements over a time span of up to 10 weeks. Overall, the first measurements start with 4 receivers, which are used to build the data collection framework. The network is then extended with 1 additional receiver starting after 12 days. A summary of the receiver sites, their availability, and the generated data is shown in Table I. The *Measured IDs* represent all the unique TPMS IDs each receiver has processed during the data collection campaign. However, since we are only interested in IDs of specific cars, we use the term *Recorded IDs* to represent the subset of data points from those authorized users. With this, we aim to provide insights into how many measurements it would be possible to collect had we gathered all data from the wild.

[1]https://github.com/yagoliz/tpms-analysis

B. Ethics and Privacy

Each of the identifiers, tied to the brand of the vehicle and the location of the authors' workplace, could lead to the connection of the daily patterns shown in this investigation with specific individuals. Our research is committed to complying with ethical standards and all authors involved in data measurement and collection have obtained the necessary Institutional Review Board (IRB) approval.

We asked participants to fill out an electronic form to allow us to collect the data from their cars. No other type of data is needed to start the capturing, which also highlights the low information amount needed to capture TPMS transmissions. We then used the TPMS measuring device to collect the 4 TPMS IDs belonging to their cars. When setting up the backend, we established a filter that only saved the IDs of the participants of our study, discarding the rest.

For the posterior analysis, we obfuscated the car brand and real IDs and substituted with the form `Ci-XX`, where `Ci` represents car i and `XX` represents one of the four wheels of the

TABLE I: Approximated receiver sites, their availability and generated data.

Site name	Recording Hours	Total Messages	Recorded Messages
Receiver 1	1.621	1.448.525	28.506
Receiver 2	1.367	930.931	31.202
Receiver 3	1.621	1.268.652	4.883
Receiver 4	1.329	1.169.537	13.170
Receiver 5	1.621	2.094.422	34.142

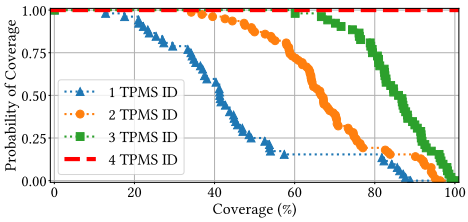

Fig. 5: Probability of achieving the same time coverage for a car when using 1,2 or 3 TPMS sensors instead of all 4.

car, namely `FL`, `FR`, `RL` and `RR` representing the Front-Left, Front-Right, Rear-Left and Rear-Right tires respectively.

VI. Car inference from tires

This Section details the data analysis approach for TPMS sensor correlation in order to infer cars and their patterns. The objective is to be able to reliably identify a car based on the transmissions of its four TPMS IDs (more than four when it is a truck or a bus).

A. Why Car Matching

TPMS transmissions are sparse and the likelihood of capturing all 4 sensors from a car passing by a road decreases with the vehicle's speed. In Fig. 5 we show the amount of TPMS transmissions that would be lost if we only used one, two or three IDs for tracking instead of all four. To analyze this, we aggregate all measurements of our dataset in one-minute windows –$\tilde{3}$0-60s is the standard transmission frequency– and count the number of times each TPMS ID appears compared to the number of times the vehicle to which it belongs appears. We call this coverage and it is represented on the x-axis of Fig. 5.

A coverage of 100% would mean that a specific ID or ID group is always captured whenever a transmission from the vehicle to which they belong is captured. The y-axis represents the probability of achieving the given coverage. When we use a single TPMS ID to track a vehicle, the median probability is slightly higher than 40%, meaning that for half of the vehicles in our dataset, we are losing at least 60% of the total messages. When using 2 or 3 IDs the coverage increases, but even when using 3 TPMS IDs from a vehicle, we would still be losing at least 20% of the transmissions for 25% of the vehicles.

These figures are relatively low, and we also must note that most of these measurements were with parked cars, so the percentage should be higher than in moving conditions, where these proportions might reduce significantly. Our findings motivates that by performing car matching, collection, and tracking improve significantly and gains might potentially be even higher in higher-mobility scenarios.

B. Car aggregation based on Jaccard index

We look at the Jaccard index as a similarity metric between two sets of correlated sensors (*A,B*). This idea is inspired by the preliminary work presented in [15], with the objective of overcoming some of its limitations, since it uses simulated and controlled data. Jaccard index is defined as follows:

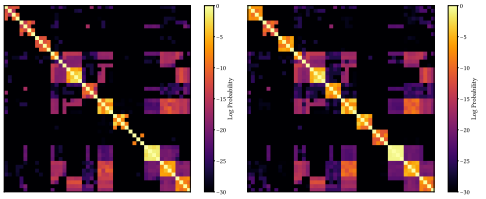

(a) One receiver. (b) All receivers.

Fig. 6: Jaccard matrix for all verified IDs. Agg: 30s.

$$Jaccard(A,B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|} \quad (1)$$

In our case, $|A|, |B|$ are the number of times sensors A, B appear respectively, and $|A \cap B|$ is the number of times both sensors are seen together in the same time window. From our dataset, we extract the tire sensor ID and the timestamp of each packet retrieved. Then, we measure the number of occurrences of each one of the identifiers with the rest of the dataset, grouping together the most frequent sensors. In particular, for each one of the transmissions registered from the sensor that is being processed, we generate a list of the sensors that appear in the same one minute window ($\pm$ 30 s). The amount of time required to compute Jaccard indices for all IDs is around 20 s with data coming from all sensors.

Table II shows the correlation of 4 different TPMS sensors belonging to 4 different cars and their corresponding Jaccard indices with the rest of the TPMS sensors in our dataset. We only show the first four IDs with the highest Jaccard index. We observe that in all cases all sensors belonging to the same car have the highest scores among each other. That is the case except for all except for one car in our dataset. For such cases, we will introduce an algorithm in Section VI-C that matches cars more robustly starting from the Jaccard index.

We choose a one-minute window, as this provides a clear pattern. In particular, Fig. 6 shows the distribution of the overlaps over all pairs considering the aforementioned one-minute window. Twelve clusters are visible (corresponding to the 12 verified cars), and many of them can be separated with relative ease. To obtain all the cars, we can take an iterative approach, where we select the most probable clusters and remove them from subsequent iterations. We discuss this approach in depth in the Section VI-C.

Apart from the fact that IDs belonging to the same car transmit at similar time windows, we also note that in most cases there are several bits in common. For our dataset, the number of common bytes ranges from 2-5, but we also observe that after the common bytes, the next octet is usually formed

TABLE II: Jaccard index example of Car 5 Front-Left Tire

Car 5, Front-Left

Target	**Other**	$\|A\|$	$\|B\|$	$\|A \cap B\|$	**Jaccard**(A,B)
`C5-FL`	`C5-FR`	991	1013	888	0.7957
`C5-FL`	`C5-RR`	991	1012	864	0.7586
`C5-FL`	`C5-FL`	991	1013	686	0.5205
`C5-FL`	`C12-FR`	991	302	8	0.0677

TABLE III: Comparison of Cross-Correlation and Jaccard index methods for different time aggregations and naive vs robust matching.

Method	Agg. (s)	Naive Match	Robust Match	Comp. Time (s)
Xcorr	5 sec.	5	9	2.0970
	10 sec.	6	10	1.5616
	30 sec.	8	11	1.1017
	1 min.	8	12	0.8708
	2 min.	10	12	0.6534
Jaccard	5 sec.	7	10	34.1570
	10 sec.	9	12	23.7702
	30 sec.	10	12	21.1729
	1 min.	11	12	19.1987
	2 min.	11	12	18.2871

by consecutive numbers. This insight, combined with the probabilistic observations provided by the Jaccard index, can provide a clear identification of individual cars.

C. Car Matching

To achieve robust tracking, our aim is to aggregate multiple TPMS IDs associated with a single vehicle (e.g., 4 for cars). A naive approach based on pairwise Jaccard index comparison suffers from limitations, such as potential mismatches due to stronger correlations between sensors of different cars and the inability to handle transitivity (e.g., if sensor A is correlated with B, and B with C, but A is not directly correlated with C).

To address these issues, we propose an algorithm which iteratively groups IDs based on a decreasing Jaccard index threshold. Starting with a high threshold, we identify strongly correlated groups (for example, the brightest clusters in Fig. 6). Smaller groups are temporarily stored and later merged in a transitive manner. This process continues with decreasing thresholds until all potential vehicles are matched. Cars which have not been fully matched (i.e. with only 2 or 3 IDs) could also be provided as an output of the algorithm.

We compared this approach with a classical time-series cross-correlation method (Table III). Although cross-correlation is generally faster, the Jaccard index method offers higher accuracy in identifying cars, especially in larger datasets with potentially ambiguous correlations. In scenarios with more cars, we predict that the difference between both approach would increase, as more vehicles to analyze might introduce more ambiguous correlation results.

This analysis demonstrates that vehicle identification can be achieved with relatively simple hardware and data analysis techniques. The choice between Jaccard index and cross-correlation methods depends on the specific application requirements: cross-correlation prioritizes speed, while the Jaccard index prioritizes accuracy.

VII. Privacy Analysis Results

A. Analysis of individual cars

Now, we look closer at the temporal dimension of individual cars that are observed in the dataset. In addition, this section analyzes the appearance patterns of the most common cars in the dataset, demonstrating how a malicious actor could infer private information about owner movements and behavior. We stress that since we do not have the ground truth (real user identity), the behaviors and patterns described here are hypothetical. The following use cases properly address the challenge mentioned in Sec.I, revealing the temporal insights obtainable from tire IDs, and the potential for attackers to harvest this data across time.

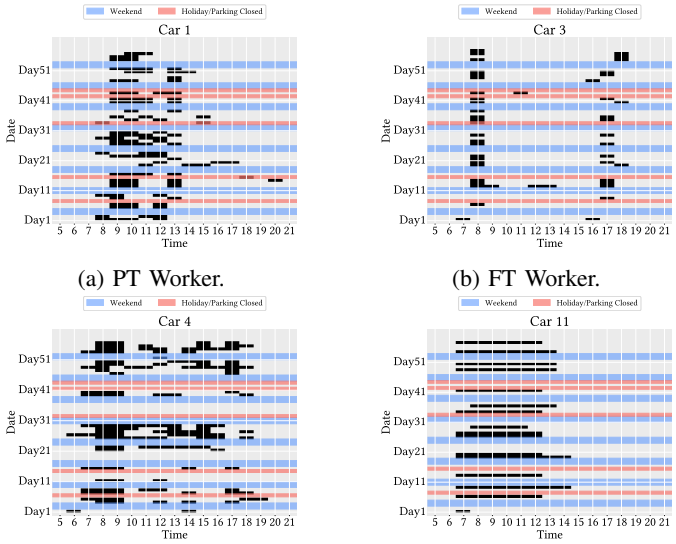

(a) PT Worker. (b) FT Worker.

(c) External Worker. (d) External PT Worker.

Fig. 7: Weekly patterns of different Worker profiles. X-axis represents the hour of the day, and the Y-Axis the Day. Each hour from each day that one of the 4 IDs of a car was seen, we paint that square black. Weekends and Holidays are shaded in blue and red respectively.

We study the profile of four workers. This car is systematically seen at a fixed time at 8:00 am (except for a day) during the week and always leaves at 5:00 pm. Due to the working hours, we can deduce that this profile belongs to a full-time worker of the company. We can also extract more information from these patterns. On Day 12, they went outside to have lunch and we measured their car's IDs at a time around 12:00 pm to 1:00 pm, as it took around 1 hour to have lunch. Another insight we can extract is that their car is never seen on Fridays, because they usually work from home on those days. Thus, from TPMS transmissions, we can also infer the remote/in-person work patterns of individuals.

Fig. 7a shows the observations of Car 1 over the course of the measurement campaign. The pattern we describe is clearly visible, since the worker usually arrives at 8:00 am and leaves at 12:00 pm. However, by overlaying weekends and local holidays, we can further explain the gaps in the data. This driving pattern corresponds to a part-time worker, which is actually the case.

In addition to regular working hours, we can also observe anomalies. First, on day 14 the driver left at their usual time, but appeared later in the evening. This can be explained because the driver attended a university course in the vicinity

and on their way home, they passed through a road near the workplace. The next day, we can also see that the same driver attended another class and decided to go back through the same route. The second day, the workplace was still closed for holidays, and yet we were still able to capture transmissions from the nearby road. This further enhances our intuition on the feasibility of deploying a car tracking infrastructure, as we are able to reliably capture transmissions from nearby roads and moving cars. These transmissions were captured by only one of the car IDs, which reinforces the car matching approach proposed in Section VI-C.

Fig. 7c has a similar pattern as Fig. 7b, but there are some key differences. First, we see more transmissions during the day (the brands are different), so we can capture their car transmissions more often. The working pattern is quite stable as well; however, there are some days where we see that the car did not appear for a few consecutive days. This pattern actually matches that of an external worker/collaborator at the building. The gaps during some days are explained by the fact that this person was working elsewhere or on a trip. In the last week of the measuring campaign, that worker went on a longer trip and was not visible throughout the time. Apart from tracking individuals' patterns, the study shows that companies could potentially use this information to track their employees without their consent or knowledge, as transmissions cannot be controlled by drivers.

The next case is of an external part-time worker, who comes a few days every week and does a shorter schedule than full-time workers. What is interesting is that we can capture TPMS transmissions every hour even when the vehicle was not moving. During our experiments we observe that each TPMS brand transmits with different strategies, as it can be seen in both Fig. 7b and Fig. 7. We observe that the TPMS sensors used by **Toyota** tend to transmit continuously, brands like **Ford** or **Nissan** do it less regularly, and brands like **Renault** only transmit when movement is detected on wheels.

B. Insights into Pressure

Additional information can also be inferred from the data provided by the open source decoders. For example, the pressure of the tire could indicate the size and weight of a car, as larger SUVs or trucks tend to require higher pressures. In addition, it could provide information on the weight of the load of a truck (the higher the load, the higher the pressure on the tires), which could be used in combination with the previously described analysis and allow more sophisticated attacks, by targeting the wheels that are in worse states and not raising any suspicion on the driver.

In Fig. 8, we show the pressure of the 4 tires for Car 1. From this we can infer that it is a utilitarian sized vehicle as the maximum pressure is around 230 kPa (note that due to temperature, we can expect some variation in the pressure). Another insight is that during the course of our measuring campaign, the driver noticed a pressure loss in the front wheels, and we can observe that, starting on the 13th day, they had inflated the tires.

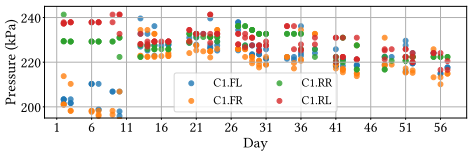

Fig. 8: Tire pressure of the selected car.

VIII. Related work

Cars are becoming increasingly advanced in terms of technology. However, its ever growing complexity has enlarged the attack surface and made them the target of numerous types of cyber-attacks. And in the near future, with more vehicles and road infrastructure interconnected, these attacks will become even more harming [16].

Works have focused in the past decade on attempting to breach the CAN-Bus [17], bluetooth stack [18] or the car's infotainment systems [19], as they control essential components of the vehicle as well as handle sensitive data, among which geolocation and routes of the vehicle are included. Attacks can also come through the vehicle connection to the cellular communications infrastructure [20]. All these types of attacks require more sophisticated access methods and they cannot be deployed passively. Furthermore, as these systems are being studied regularly, countermeasures appear as well which can be implemented in subsequent updates to the vehicle.

In this paper, we have focused on TPMS as it is often an overlooked component of the vehicle. And in fact, there is usually little control as to which TPMS sensor is installed in our vehicles. Although previous research has examined TPMS vulnerabilities [5], [15], [21], these studies were mainly based on simulations or limited real-world data. Furthermore, they often focused on active attacks (for example, spoofing) rather than passive tracking enabled by low-cost SDR networks. Unlike previous work, this paper presents a comprehenstive analysis based on a large-scale SDR dataset collected using a real-world deployment of low-cost receivers. This allows us to investigate long-term tracking and pattern analysis, revealing the practical privacy implications of widespread SDR use. Our findings highlight the need to revise existing cybersecurity regulations (e.g., UN Regulation No. 155 [2], [3]) to include TPMS vulnerabilities and ensure protection against transmission exploits.

Some studies have proposed adding encryption [22], [23], [24], [25] or even redesigning the sensor altogether [8], however, to the best of our knowledge, none of these solutions have been widely adopted by car manufacturers. Being able to capture such a large number of transmissions highlights the need to address the privacy risks that we discuss in this work.

Recently, a new TPMS system was proposed by Pirelli and Bosch called the Cyber Tyre [26]. This concept uses Bluetooth Low Energy (BLE) to send information to the car's ECU, which could entail its own risks as BLE devices can be eavesdropped [27]. Besides, this system is planned for higher-end road and race cars, so we do not expect it to be widely adopted in the near future by the majority of vehicles.

In summary, although previous work has studied possible privacy and security threats in TPMS, their approaches present significant limitations, mainly related to the lack of real data and analytical studies. This work seeks to solve previous limitations, by deploying low-cost spectrum receivers to collect real TPMS data, assess current protocol deployments, and analyze security and privacy implications for drivers. Our study is timely, given the aforementioned ongoing regulations on cyberthreats at the United Nations level, and in a large number of countries in the world.

IX. Conclusion

We have shown that TPMS signals, originally introduced to improve road safety, can be misused to track vehicles and therefore the movement pattern of their owners through a network of low-cost software-defined spectrum receivers. The lack of encryption and standardization is one of the root causes. This information can be used by attackers to infer the movement patterns of individuals or to track people around cities if multiple sites are available for recording. Furthermore, we argue that it is rather easy in practice to link TPMS sensors with a specific person of interest. A TPMS signal receiver can be linked to a camera or, if the person of interest's home address is known, a targeted recording at the person's home can reveal their unique TPMS sensor IDs. Once these IDs are known, the movements of that person can be tracked with inexpensive publicly available software-defined radios. Attackers can use this information to learn, predict, and exploit a person's movements, points of interest, and behavior patterns. Although there exist proposals to enhance privacy for TPMS sensors, we could not find any real deployment from manufacturers. Therefore, we urge legislators, policymakers, and manufacturers to take the necessary steps to improve the privacy and security of the pressure monitoring system.

Acknowledgments

The research conducted by IMDEA Networks was sponsored in part by armasuisse under the Cyber and Information Research Program, in part by the NATO Science for Peace and Security Programme under grant G5461 and in part by project PID2022-136769NB-I00 funded by MCIN/AEI /10.13039/501100011033 / FEDER, UE.

References

[1] M. Kreil, Flüpke, Wir wissen wo dein Auto steht (2024).
[2] E. United Nations, S. Council, ECE/TRANS/WP.29/2020/79: Proposal for a new UN Regulation on uniform provisions concerning the approval of vehicles with regards to cyber security and cyber security management system (2020).
[3] Unece, UN Regulations on Cybersecurity and Software Updates to pave the way for mass roll out of connected vehicles (2020).
[4] E. Union, Regulation (EU) 2019/2144 of the European Parliament and of the Council on type-approval requirements for motor vehicles and their trailers, and systems, components and separate technical units intended for such vehicles (2019).
[5] I. Rouf, R. Miller, H. Mustafa, T. Taylor, S. Oh, W. Xu, M. Gruteser, W. Trappe, I. Seskar, Security and privacy vulnerabilities of in-car wireless networks: A tire pressure monitoring system case study, in: 19th USENIX Security Symposium, USENIX Association, Washington, DC, USA, 2010, pp. 323–338.
[6] S. Velupillai, L. Guvenc, Tire pressure monitoring [applications of control], IEEE Control systems magazine 27 (6) (2007) 22–25.
[7] J. Caban, P. Droździel, D. Barta, Š. Liščák, Vehicle tire pressure monitoring systems, Diagnostyka 15 (3) (2014) 11–14.
[8] K. Daimi, M. Saed, Securing tire pressure monitoring system, in: Proceedings of the 2018 the 14th Advanced International Conference on Telecommunications Conference, Barcelona, Spain, IARIA, Barcelona, Spain, 2018, pp. 22–26.
[9] R. Suender, G. Prokop, T. Roscher, Comparative Analysis of Tire Evaluation Methods for an indirect Tire Pressure Monitoring System (iTPMS), SAE International Journal of Passenger Cars - Mechanical Systems 8 (2015-01-1519) (2015) 110–118.
[10] C. R. Bowen, M. H. Arafa, Energy Harvesting Technologies for Tire Pressure Monitoring Systems, Advanced Energy Materials 5 (7) (2015) 1401787.
[11] Autel, MaxiTPMS TS501 | Autel (2023).
[12] FCC, OET List Exhibits Report ID: MRXFG2R4MA (2011).
[13] R. W. Stewart, K. W. Barlee, D. S. W. Atkinson, Software Defined Radio Using MATLAB & Simulink and the RTL-SDR, Strathclyde Academic Media, Glasgow, GBR, 2015.
[14] B. Larsson, Rtl_433: Program to decode radio transmissions from devices on the ISM bands (and other frequencies) (2022).
[15] K. Hacker, S. Graham, S. Dunlap, Vehicle Identification and Route Reconstruction via TPMS Data Leakage, in: J. Staggs, S. Shenoi (Eds.), Critical Infrastructure Protection XIII, IFIP Advances in Information and Communication Technology, Springer International Publishing, Cham, 2019, pp. 123–136.
[16] X. Sun, F. R. Yu, P. Zhang, A Survey on Cyber-Security of Connected and Autonomous Vehicles (CAVs), IEEE Transactions on Intelligent Transportation Systems 23 (7) (2022) 6240–6259.
[17] M. Bozdal, M. Samie, I. Jennions, A Survey on CAN Bus Protocol: Attacks, Challenges, and Potential Solutions, in: 2018 International Conference on Computing, Electronics & Communications Engineering (iCCECE), IEEE, Southend, UK, 2018, pp. 201–205.
[18] D. Antonioli, M. Payer, On the Insecurity of Vehicles Against Protocol-Level Bluetooth Threats, in: 2022 IEEE Security and Privacy Workshops (SPW), IEEE, San Francisco, CA, USA, 2022, pp. 353–362.
[19] S. Jeong, M. Ryu, H. Kang, H. K. Kim, Infotainment System Matters: Understanding the Impact and Implications of In-Vehicle Infotainment System Hacking with Automotive Grade Linux, in: Proceedings of the Thirteenth ACM Conference on Data and Application Security and Privacy, CODASPY '23, Association for Computing Machinery, New York, NY, USA, 2023, pp. 201–212.
[20] Y. Li, Q. Luo, J. Liu, H. Guo, N. Kato, TSP Security in Intelligent and Connected Vehicles: Challenges and Solutions, IEEE Wireless Communications 26 (3) (2019) 125–131.
[21] K. T. Sterne, J. M. Ernst, D. K. Kilcoyne, A. J. Michaels, G. Moy, Tire pressure monitoring system sensor radio frequency measurements for privacy concerns, Transportation Research Record 2643 (1) (2017) 34–44.
[22] D. K. Kilcoyne, S. Bendelac, J. M. Ernst, A. J. Michaels, Tire Pressure Monitoring System encryption to improve vehicular security, in: MILCOM 2016 - 2016 IEEE Military Communications Conference, IEEE, Baltimore, MD, USA, 2016, pp. 1219–1224.
[23] A. Kolodgie, P. Berges, R. Burrow, M. Carman, J. Collins, S. Bair, G. D. Moy, J. M. Ernst, A. J. Michaels, Enhanced TPMS security through acceleration timed transmissions, in: MILCOM 2017 - 2017 IEEE Military Communications Conference (MILCOM), IEEE, Baltimore, MD, USA, 2017, pp. 35–39.
[24] C. Solomon, B. Groza, LiMon - Lightweight Authentication for Tire Pressure Monitoring Sensors, in: A. Bécue, N. Cuppens-Boulahia, F. Cuppens, S. Katsikas, C. Lambrinoudakis (Eds.), Security of Industrial Control Systems and Cyber Physical Systems, Lecture Notes in Computer Science, Springer International Publishing, Cham, 2016, pp. 95–111.
[25] M. Vaszary, A. Slovacek, Y. Zhuang, S.-Y. Chang, Securing Tire Pressure Monitoring System for Vehicular Privacy, in: 2021 IEEE 18th Annual Consumer Communications & Networking Conference (CCNC), IEEE, Las Vegas, NV, USA, 2021, pp. 1–6.
[26] Pirelli, Cyber™ Tyre | Pirelli (2024).
[27] M. Cäsar, T. Pawelke, J. Steffan, G. Terhorst, A survey on Bluetooth Low Energy security and privacy, Computer Networks 205 (2022) 108712.