

Osprey: A mmWave Approach to Tire Wear Sensing

Akarsh Prabhakara
Carnegie Mellon University
aprabhak@andrew.cmu.edu

Swarun Kumar
Carnegie Mellon University
swarun@cmu.edu

Vaibhav Singh
Carnegie Mellon University
vaibhav3@andrew.cmu.edu

Anthony Rowe
Carnegie Mellon University
agr@ece.cmu.edu

ABSTRACT

Tire wear is a leading cause of automobile accidents globally. Beyond safety, tire wear affects performance and is an important metric that decides tire replacement, one of the biggest maintenance expense of the global trucking industry. We believe that it is important to measure and monitor tire wear in all automobiles. The current approach to measure tire wear is manual and extremely tedious. Embedding sensor electronics in tires to measure tire wear is challenging, given the inhospitable temperature, pressure, and dynamics of the tire. Further, off-tire sensors placed in the well such as laser range-finders are vulnerable to road debris that may settle in tire grooves.

This paper presents Osprey, the first on-automobile, mmWave sensing system that can measure accurate tire wear continuously and is robust to road debris. Osprey's key innovation is to leverage existing, high-volume, automobile mmWave radar, place it in the tire well of automobiles, and observe reflections of the radar's signal from the tire surface and grooves to measure tire wear, even in the presence of debris. We achieve this through a super-resolution Inverse Synthetic Aperture Radar algorithm that exploits the natural rotation of the tire and improves range resolution to sub-mm. We show how our system can eliminate debris by attaching specialized metallic structures in the grooves that behave as spatial codes and offer a unique signature, when coupled with the rotation of the tire. In addition to tire wear sensing, we demonstrate the ability to detect and locate unsafe, metallic foreign objects such as nails lodged in the tire.

We evaluate Osprey on commercial tires mounted on a mechanical, tire-rotation rig and a passenger car. We test Osprey at different speeds, in the presence of different types of debris, different levels of debris, on different terrains, and different levels of automobile vibration. We achieve a median absolute tire wear error of 0.68 mm across all our experiments. Osprey also locates foreign objects lodged in the tire with an error of 1.7 cm and detects metallic foreign objects with an accuracy of 92%.

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

MobiSys '20, June 15–19, 2020, Toronto, ON, Canada

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-7954-0/20/06...\$15.00
<https://doi.org/10.1145/3386901.3389031>

CCS CONCEPTS

• **Hardware** → **Sensor applications and deployments; Sensors and actuators; Wireless devices; Signal processing systems**; • **Computer systems organization** → **Embedded and cyber-physical systems**; • **Applied computing** → Computers in other domains.

KEYWORDS

Millimeter Wave, Wireless Sensing, Tire Wear, Tread Depth, Radar, Automotive, 77 GHz, FMCW, Super Resolution, Inverse Synthetic Aperture Radar Imaging, Debris, Spatial Coding, Orthogonal Codes, Free of Electronics, Foreign Object

ACM Reference Format:

Akarsh Prabhakara, Vaibhav Singh, Swarun Kumar, and Anthony Rowe. 2020. Osprey: A mmWave Approach to Tire Wear Sensing. In *The 18th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '20)*, June 15–19, 2020, Toronto, ON, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3386901.3389031>

1 INTRODUCTION

Tires are a pivotal part of automobiles and directly affect both their safety and performance. According to a 2012 report from the National Highway Traffic Safety Administration, a whopping 194,000 (5%) out of all vehicles involved in crashes between 2005-2007 encountered problems with their tires [9]. In spite of introducing Tire Pressure Monitoring Systems (TPMS) in all vehicles in 2007 (as per FMVSS No. 138 [47]), tire-related accidents are still prevalent even today. This is partly because a TPMS is effective only in sensing issues related to tire pressure. Other important issues such as tire wear / tread degradation, cited as one of the leading causes of tire-related crashes [9], go unnoticed. Beyond safety, tread degradation is an important metric that decides tire replacement, the second biggest maintenance expense (after fuel) of the US\$58 billion global trucking industry [16], which presently relies on extremely coarse manufacturer-provided heuristics to model when to replace or re-tread tires [2]. It is thus imperative to measure and monitor tire wear / tread depth in all automobiles, just as we monitor pressure using TPMS today.

Past efforts to design electronic sensors that automatically and always sense tread depth from within the tire [5] are expensive to manufacture and maintain over a tire's life. Such a sensor must necessarily survive automotive grade temperatures ranging from -35°C to 85°C and cope with pressures of 2.75 bar. Making the environment more challenging is the extreme friction at the surface, mobility of the vehicle, various road terrains, and varying




Figure 1: Osprey: A mmWave radar is placed in the tire well to image the tire and estimate the tire wear while filtering out the effects of debris.

weather conditions. For this reason, much of the state-of-the-art estimates tread depth indirectly using TPMS readings [11] or other sensors mounted on the inner lining [37], (see Fig. 2) – all of which are far away from the actual tread and therefore prone to large errors. Even off-tire solutions such as laser rangefinders in the tire well experience errors due to debris accumulation and varying tire pressure [42]. This motivates the need for continuous sensing of accurate tire tread depth without embedding electronics within the tire.

We propose Osprey¹, the first debris-resilient mmWave sensor system design to measure tread depth without embedding any electronics within the tire. This paper is a joint work with Bridgestone, a leading global tire manufacturing company whose expertise we utilized in developing and evaluating our system. We build our design on top of a commercial, automotive, mmWave radar (TI AWR1642BOOST) which is becoming more pervasive in automobiles for the purposes of collision detection and avoidance. We place this radar in the well of the tire, as shown in Fig. 1. Our approach then draws inspiration from traditional radar imaging and relies on extracting precise range locations of the surface of the tire and that of thin strips of metal placed in the grooves of the tire. With precise range locations of these two surfaces, tread depth can be estimated as the difference between the two. The rest of this paper describes our solutions to two key challenges in making our design practical:

(1) Insufficient Range Resolution: With a 4 GHz bandwidth, the range resolution that state-of-the-art automotive mmWave radars provide is 3.75 cm ($c/2B$). However, tread depths vary at significantly finer granularities between 2 - 20 mm. As a result, the reflections from both the surface and the groove of the tire are below the best resolution of mmWave radars and thus indistinguishable. To mitigate this challenge, we pose this range resolution problem as a parametric estimation problem to resolve the range bins by building a super-resolution algorithm that builds on Inverse Synthetic Aperture Radar (ISAR). Our approach effectively exploits the rotation of the tire to view the same segment of the groove and tire surface from different perspectives. Unlike traditional SAR in the far-field terrain imaging context, we model the rotating nature of tire where segments of the tire appear, disappear, and then re-appear into the view. In addition, we overcome unique challenges owing to the proximity of the tire, misalignment, uneven rotation, and vibrations.

¹Osprey is a fish-eating bird that can accurately gauge the depth of prey underwater, overcoming the effects of refraction.

Sec. 5 describes how our approach resolves millimeter differences in tread depth.

(2) Resilience to Debris: While sensing to recover the range of the tread and groove might seem enough to solve the problem of tread depth, small pieces of debris embedded in the groove of the tire can impede these measurements. This might lead to incorrect characterization of tread depth. To deal with this challenge, we carefully craft the layout of the metallic strips in the groove along specific coded patterns. We choose metal to design these patterns since they provide strong reflections capable of penetrating debris. Further, we note that today’s tire manufacturing processes are already tuned to add metallic strips to tires to disperse static charge. Our radar’s received signal from the groove therefore includes both the specially designed metallic spatial codes as well as debris settled within. We then correlate the received signal with the spatial code to isolate the signal from the groove’s surface and eliminate any effect of debris. Our specific design is inspired by optical orthogonal codes, used in optics to allow for correlating patterns on ISAR images. Sec. 6 describes how to design spatial codes that allow for efficient elimination of the impact of debris within the small surface area available within a tire’s groove.

Sensing Foreign Objects: Beyond tread depth, our approach is designed to be robust and to detect and locate foreign objects in the tire. Foreign objects like pieces of wood, metal, stone, etc. lodged on the surface or in the groove of the tire are a major cause of tire damage that can drastically reduce the tire’s lifetime. We present a foreign object localization system that can accurately determine the location of the object and notify the user in real time. Our approach processes the output of the Inverse Synthetic Aperture Radar algorithm, which appears akin to X-ray images showing components lodged within the tire. We receive multiple such images as the tire rotates over time. We then stitch these images together to generate a continuous image of the tire and use the known pattern of metal strips in the grooves to determine the exact location of the foreign object. We then use different machine learning algorithms to distinguish between harmful (metallic) and non-harmful objects lodged in the tire, chosen among common sources of debris dangerous to a tire. Sec. 7 describes our method for characterizing the kind of object present in the groove.

Limitations: We note the following important limitations of our system (see Sec. 10): (1) Our evaluation shows Osprey operating at tire speeds up to 5.45 kmph. We note that tire abrasion occurs at relatively slow timescales, meaning that periodic measurements when vehicles slow down are already very informative. (2) Our system makes use of the tire geometry and needs to be notified when tires are changed to download expected tire geometry from the manufacturer. (3) The cost of AWR1642 radar IC, one per per tire well, is <\$40. While this is not insignificant compared to cost of passenger car tires, we note that this cost is insignificant for expensive tires – buses, trucks, agricultural, off the road vehicles, and for trucking companies – where tire replacement alerts outvalue system cost [16].

We implement Osprey on a 77 GHz mmWave automotive radar. We evaluate our system on Falken Eurowinter HS449 and Bridgestone Blizzak LM001 tires, which differ in terms of tread patterns. We evaluate our system on moving a passenger car as well as on a




Figure 2: Depicts different components of a tire – distinct tread patterns, grooves running along the circumference, steel belts, nylon belts, etc. [50]

mechanical tire rotation rig. We rotate tires at speeds of up to 5.45 kmph. Our evaluation demonstrates:

- A median tread depth estimation error of 0.68 millimeters across all experiments.
- A median error of 1.7 cm in the location of foreign objects.
- An accuracy of 92% in the type of foreign object classified between metal and non-metal.

Contributions: Our main contributions include:

- A super-resolution Inverse Synthetic Aperture Radar algorithm that exploits the natural rotation of the tire to 3D image its surface at sub-millimeter accuracy, amidst tire dynamics.
- A sub-millimeter accurate tread-depth sensing solution in the presence of debris using spatial codes composed of metallic strips in the groove.
- A solution to detect the location and type of foreign object lodged in the tire.

2 A PRIMER ON TIRES

This section provides an overview of tire fundamentals needed to understand our problem domain and define tread depth.

What is Tread Depth?: Fig. 2 shows the cross section of a typical tire. The prominent patterns etched on the rubber provide traction. The shallow portion in these patterns are grooves, and rubber portions which touch the road are treads. Tread depth, a measure of tire wear, is simply the distance between tread’s surface and groove. In general, greater tread depth provides greater traction. This is why heavy-load truck and agricultural tires have greater tread depth than car tires.

Why does Tread Wear matter?: During the lifetime of a tire, tread wears down naturally and tread depth decreases. Natural wear is a slow process. Tires are rated to last for 20-25 thousand kilometers. For trucks which drive 100s of kilometers each day, it is useful for the truck fleet management to track changes in tread depth continuously and plan the maintenance costs of retreading. It also opens up opportunities like creating an economically viable leasing model for trucks. For other vehicles like passenger cars which travel around 20 km per day, measuring fine-grained changes due to natural wear continuously seems unnecessary. However, there are other factors which lead to sudden or faster wear – misalignment of wheels, improper pressure, varying load on the vehicle, characteristic driving traits of drivers, and many more. Insufficient tread depth leads to the tire not being able to grip the

	Accurate	Robust to debris	Real-time	Free of extra electronics within tire
Manual - Ruler / 1 cent coin ..	✓	✓	✗	✓
Mapping pressure and temperature to tread depth	✗	✓	✓	✓
In-tread sensors	✓	✓	✓	✗
Laser range finder in real time	✓	✗	✓	✓
Osprey	✓	✓	✓	✓

Figure 3: Why mmWave: A design space exploration.

road surface properly, and the driver can lose control of the vehicle due to loss of braking efficiency or hydroplaning. As a safety precaution, the Federal Motor Carrier Safety Administration (FMSCA) in the US defines a legal minimum tread depth requirement of 3.2 mm (4/32”) for steer tires of trucks and buses and 1.6 mm (2/32”) for other tires and vehicles. It would therefore be convenient to monitor tread depth continuously so that sudden or faster wear can be tracked and accidents can be avoided. To help put the minimum tread depth numbers in perspective, steer tires of trucks start with about 17.5 mm (22/32”) of tread depth, and passenger tires have an initial tread depth of about 9.5 mm (12/32”). This means that a tread depth measuring system will have to track changes at the level of a few millimeters! On top of this, a continuous tread depth measuring system should be designed with at least the two minimum requirements for on-road robustness - be immune to debris accumulating in the groove, and be able to measure as tires are rotating.

Tire Sensors: Ordinary sensors such as TPMS are mounted on the rim of the wheel / inner lining of tires. RFID tags used for inventory management of tires are embedded in the rubber along the sidewall, which makes them easy to scan. Embedding sensors in the tread is challenging, owing to high temperatures (about 85°C), pressure (about 2.75 bar), and extreme dynamics. In-tread sensors which can withstand these conditions (see Sec. 10) are expensive to embed as they require changes to existing manufacturing lines.

A tire is not all rubber: A tire is a highly engineered system, mostly but not entirely made of rubber. Just below the groove, layered steel belts (or sometimes nylon) run throughout the circumference of the tire. Their intended purpose is to reinforce the structural strength of the tread. In the absence of this structure, a tire would heavily deflect under load and inflate uncontrollably akin to a balloon when air pressure is applied. We use the metal-compatibility of tires to our advantage in our system design, as we show in Sec. 6.

3 WHY MMWAVE?

Our choice of mmWave radars for tire depth sensing stemmed from a deep design space exploration of different choices available, with close consultation with our tire manufacturing industry collaborator.

The Design Space: Fig. 3 depicts the design space of feasible approaches to tire tread depth sensing. Broadly, the problem of tread depth sensing has three approaches: (1) Human effort, e.g. measurement with a penny; (2) In-tire and indirect sensors, which lead to manufacturing and robustness problems; or (3) Off-tire sensors, typically vulnerable to debris. Sec. 11 elaborates on prior

work in each of these spaces and their known weaknesses. Among the three classes, we choose the off-tire sensor approach given that we seek a robust solution that requires minimal human intervention and no change to the tire-manufacturing process. Current solutions in the off-tire sensor space use visible light (e.g. LIDAR [42]) which is vulnerable to debris or X-ray imaging in factories [18] that are unsafe to use on the road.

Our Choice: We therefore chose an RF-based solution given that radio signals are known to penetrate through debris. Compared with commonly used RF solutions such as RFID and WiFi, mmWave frequencies provide one of the widest bandwidths among all radio systems. We specifically chose mmWave RF frequencies due to this high available bandwidth, keeping in mind the sub-millimeter resolution requirement of tread depth sensing. Further, mmWave radars are commonplace in the automotive industry as collision radars and are already optimized for vehicular mounting. Yet, the high-bandwidth of mmWave radars – about 4 GHz – while high relative to sub-6 GHz bands, still offers only a distance resolution of 3.75 cm. While this is sufficient for collision sensing, it is well over an order of magnitude poorer than our desired tread-depth resolution. The rest of this paper discusses the various challenges and opportunities in circumventing this resolution limit.

4 SYSTEM OVERVIEW

In this section, we present a high level overview of Osprey and present the organization of the rest of this paper. Fig. 4 depicts the high level architecture of our system. We mount a commodity automotive mmWave radar in the well of a tire with its antenna facing the tire. Our objective is to create a 3D depth image of the surface of the tire and grooves, while filtering out any impact of debris. We then measure tread depth simply as the difference in distance between the tire surface and groove. We further classify the filtered-out debris among different categories and map out their specific locations on the groove, particularly for debris of importance (e.g. metal lodged in groove). The rest of this paper describes three important challenges in achieving this design (Sec. 5-7), a discussion of limitations (Sec. 10), as well as a comprehensive system evaluation (Sec. 8-9):

(1) Super-Resolution Tire Surface Imaging: Sec. 5 describes how one can obtain the tire surface image as shown in Fig. 4 to study wear. Our approach at a high level uses Inverse Synthetic Aperture Radar (ISAR) that effectively exploits movement of the tire to improve spatial resolution. We further develop super-resolution algorithms that are specifically optimized along the depth axis. We further tackle important challenges from tire deflections and vibrations which can impact the performance of ISAR.

(2) Tread Depth Amidst Debris: Next, Sec. 6 describes how one can filter out debris that may create spurious peaks in Fig. 4, leading to misleading depth of the tread groove. To this end, Osprey places metallic strips in the groove along specific patterns – effectively, spatial codes. While the groove’s peaks will match this code, debris (on average) will not. We describe our choice of design of these spatial codes as well as how they remain robust to different kinds of debris.

(3) Sensing Foreign Objects: Finally, Sec. 7 discusses locating and identifying the nature of tire debris. Our approach to locating

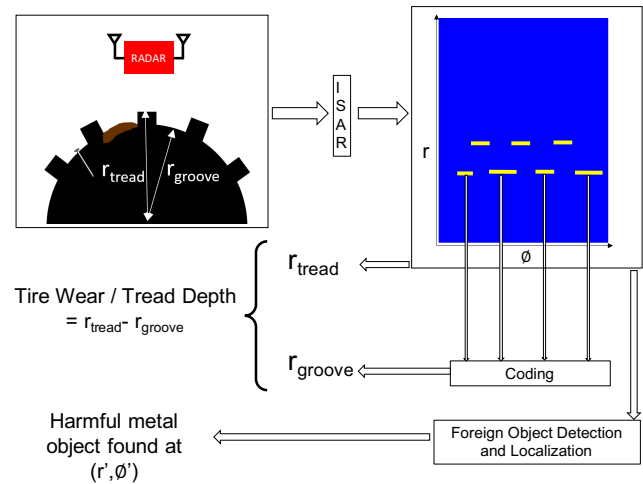


Figure 4: Osprey’s Architecture: (1) Generates a super-resolution ISAR tire image (2) Filters out debris using coding to obtain tread depth (3) Detects and locates foreign objects.

debris relies on how close their peaks are relative to the underlying metallic spatial codes on the groove which act as reference. We further devise features from the debris that help us uniquely map them to different types of debris.

5 TIRE SURFACE IMAGING

In this section, we aim to measure the depth of different points along the surface of the tire from the mmWave radar mounted along the tire well (see Fig. 1). Recall that the difference between this quantity and the depth of grooves along the tire results in the tread depth – our main quantity of interest. Note that the spacing between the tire surface and radar can change for two reasons: (1) Wear and tear of the tire; (2) Due to tire suspension, e.g. weight of load or dynamics of the road such as speed bumps. Hence, it is important to measure both the distance of the tire’s surface *and* groove relative to the radar – i.e., measuring the former alone does not suffice. This section specifically focuses on capturing the depth of the exterior surface alone, and we will describe the unique challenges in determining the depth of the grooves later in Sec. 6.

Osprey’s Approach: To address this challenge, Osprey develops a super-resolution algorithm that measures the depth of the tire surface at sub-mm accuracy. We do this by exploiting the mobility of the tire – particularly its rotation, to our advantage. Our specific approach is to build an Inverse Synthetic Aperture Radar (ISAR) solution for the tire context. In traditional radar systems used for military and scientific applications, ISAR integrates multiple received signals from a target moving at a known speed (e.g. an aircraft) to a static object (e.g. ground radar) to localize it at fine precision, significantly higher than what the bandwidth of the radar would allow. In an analogous fashion, our approach integrates signal reflections from the same point as the tire rotates to improve our precision of its depth relative to the radar. The rest of this section describes our solution to two key challenges unique to the tire context: (1) First, we need mechanisms to model the rotation of the tire, including dynamics of the tire due to vibration and




Figure 5: Osprey’s ISAR uses a cylindrical coordinate system centered at the center of tire. Any point is represented as (r, ϕ, z) - r is radius, ϕ is azimuth, z is z-coordinate. mmWave radar is at $(D, 0, 0)$. The difference between the radius of surface and groove is essentially the tread depth.

misalignment; (2) Second, one would need to isolate measurements from the tire’s surface from all other sources of reflection, such as the well of the tire.

5.1 Inverse Synthetic Aperture on a Tire

Osprey’s ISAR algorithm leverages tire rotation with respect to the mmWave radar placed in the tire well. For now, we assume that the precise angle of rotation of the tire is known at any time. Sec. 6.1 details our approach to estimate this by attaching known tire-compatible artifacts to the groove. However, we note that tire orientation can also be gleaned out-of-band through an encoder [36], an inertial measurement unit [34] placed on the axle, or from motor feedback [8]. We also assume that the prior tread pattern and tire model are known up front, for instance, specified at tire installation. As a result, the only unknown that needs to be captured is the depth of the tire’s surface relative to the radar.

Challenges: A unique challenge in formulating Osprey’s ISAR optimization is to model the effect of the rotation of the tire on the received signals at the radar. Specifically, note that the received signal at the radar is the sum of reflected signals that impinge on multiple points along the surface of the tire. As the tire rotates, these points rotate as well at a rate dictated by the tire’s speed. Further, some points progressively disappear from view as they move beyond the field-of-view of the radar, while others appear into view at the other end of the tire. Our formulation of ISAR therefore needs to model tire surface trajectories to isolate signals received from across points on the surface.

ISAR Formulation: Our mathematical formulation of ISAR models the journey of an imaginary point on the surface of the tire, to ascertain its depth – illustrated in Fig. 5. For mathematical convenience given that the tire is a rotating body, we choose cylindrical coordinates to express this point, with the origin defined as the center of the tire. Let (r, ϕ, z) represent the cylindrical coordinates of a point, X , in 3D space on the surface of a tire. Let ϕ denote the azimuthal angle that changes with the rotation of the tire about the z -axis. We assume that the mmWave radar is located at a distance D along the positive x axis (i.e. at $(D, 0, 0)$). Then by definition, r directly relates to tread depth – any wear of the tread automatically results in an equal reduction in r . Our objective is to therefore estimate r .

Next, we trace the journey of our imaginary point that traverses $(r, \phi(t), z)$ over time t , where $\phi(t)$ denotes the changing azimuth as the wheel spins. Let $d(t)$ denote the distance between the points $(r, \phi(t), z)$ and $(D, 0, 0)$. Then it is clear that the wireless channel contribution at any time over the trajectory of the point X , $h_X(t)$ due to the reflection of the signal from the radar off the point is [44]:

$$h_X(t) = \frac{1}{2d(t)} e^{-j4\pi d(t)/\lambda}$$

Where λ denotes the wavelength. Our algorithm to isolate the signal along any point located at $(r, \phi(0), z)$ at $t = 0$ therefore actively projects the received channel along $h_X(t)$. We specifically run a modified Bartlett algorithm for Inverse Synthetic Aperture Radar, akin to an inverse spatial Fourier transform, that accounts for the tire’s rotation. Specifically, we write the power of the received signal reflected off the point $(r, \phi(0), z)$ on the tire as:

$$P(r, \phi(0), z) = \sum_t h(t) e^{+j4\pi d(t)/\lambda} \quad (1)$$

Where $h(t)$ is the wireless channel read at time t .

Note that mmWave radars often have multiple antenna elements, e.g. eight in the radar used in our experiments. In such cases, the above process needs to be co-optimized across antenna elements. We lay the multiple antennas along the z -axis. Similar to the above process, we create expected wireless channels of reflection from a point for each antenna and then sum over the projection across all antennas in addition to summing across time.

Our choice of Bartlett, as opposed to other antenna algorithms such as MUSIC [52] or ESPRIT [32], stems from the non-uniformity of the rotation of the tire, where tires often rotate at uneven speeds or packet samples from the mmWave radar are obtained at unequal times. This may cause spurious peaks in algorithms such as MUSIC or ESPRIT that are best suited to uniform arrays [39], whereas Bartlett remains relatively more robust.

5.2 Resilience to Tire Dynamics

The above formulation assumes perfect awareness of the rotational dynamics of a tire over time. However, several dynamics of the tire may make its movements irregular and often noisy or unpredictable. Our approach needs to be resilient to tire dynamics. We detail how our approach handles three classes of common tire dynamics:

Suspension: Vehicles are equipped with suspension to provide resilience to changes in road topography, such as speed bumps, gravel, potholes, etc. When a vehicle drives over, say, a speedbump, the wheel traces the contours of the speedbump, while the well of the tire may follow the contour significantly less tightly due to suspension. For Osprey, this means that the distance between the radar and tire surface can change due to *both* the dynamics of suspension and the rotation of the tire. Neglecting tire suspension can introduce a fixed error Δ in our estimate of the distance of the tire surface relative to the radar – corresponding to the state of vehicle suspension. Fortunately, this fixed error does not meaningfully impact our effective estimate of tread depth. This is because our estimate of the depth of the groove also experiences the exact same error Δ owing to suspension. And, tread depth which is the relative distance between tread and groove does not change with suspension effects. Hence, the difference in distance between the




Figure 6: An ISAR image reveals the tire tread pattern.

tire surface and groove is independent of Δ , that is, Osprey’s tread depth estimate remains unaffected by suspension.

Vibrations: Vehicles routinely experience vibrations due to the movement of the motor. Unfortunately, these vibrations are noisy and therefore extremely challenging to predict and directly introduce noise to our surface depth estimates. Worse, different parts of the vehicle may vibrate differently – for instance, the (metallic) well of the tire typically vibrates at a higher amplitude compared to the (mostly rubber) tire itself. Osprey limits the impact of vibrations in two ways: (1) First, we place the radar itself at the centroid of the well lodged firmly with the body of the vehicle. This limits unnecessary extra mechanical vibrations that loosely fitted fixtures experience; (2) Second, we average measurements across several packets and drop outliers ($\sim 1\%$ of measurements) to discount spurious readings due to vibrations. Our evaluation in Sec. 9 measures Osprey’s accuracy on the well of a vehicle to study the impact of vibrations.

Misalignment: While our model assumes that z (the z -coordinate) of any point on the tire’s surface is fixed as the tire rotates about the z -axis, misalignment of the tire can cause (mild) variations in z as the tire rotates. If this is not accounted for, the tire geometry sensed experiences unnatural spatial distortions, skewing the boundary of the tire surface in accordance with the tire misalignment. Osprey accounts for tire misalignment by actively leveraging this distortion. Specifically, it models known features on the tire surface, such as its boundaries, and measures their skew along the z axis. It then performs a cubic-spline interpolation of this skew to estimate the corresponding offset in z as a function of time. Osprey then reruns its algorithm with the new estimates of z over time, until the skew along z in its ISAR output disappears.

Accentuating all the above dynamics are tire speed and signal multipath from the vehicle itself and its surroundings. We deal with these challenges explicitly in the next section.

5.3 Isolating Tire Surface

In this section, we explore ways to isolate signals from the surface of the tire from other sources of reflected signals – such as signals from the well of the tire, metallic parts of the car, or even objects along the street. Such reflectors can cause spurious peaks to appear within the measured ISAR image obtained from Eqn. 1. Osprey’s

approach to de-cluttering the surface from other objects relies on two solutions: tire pattern recognition and background subtraction.

Tire Pattern Recognition: Our approach to focus on the surface of the tire relies on the fact that we know the precise tread pattern of the tire based on its manufacturing specifications. This pattern inevitably manifests itself in the ISAR image, where grooves and bumps on the tire’s surface produce variations along the radial and azimuthal axes (r and ϕ). Our experimental results show that these patterns manifest across a wide-range of tire patterns. Fig. 6 depicts a candidate tire tread pattern and its corresponding ISAR image – a surface plot depicting $P(r, \phi, z)$ (from Eqn. 1), where ϕ and r denote the x and y axis respectively and pixel intensity denotes the value of $P(r, \phi, z)$. We note that the tire patterns closely align with corresponding ISAR images. As a result, Osprey can effectively identify points on the surface of the tire by correlating the ISAR image with this known pattern.

Impact of Tire Speed: An important effect that needs to be accounted for in designing our approach to correlation is the impact of tire speed. While our ISAR formulation explicitly accounts for tire speed in the evolution of $\phi(t)$, small errors can produce distortions in the ISAR image. In particular, two effects are prominent: (1) The images of the tread may be marginally stretched or squeezed based on whether the tire speed was over- or under-estimated; (2) Sharp edges along the tread could appear unduly smooth owing to vibrations and tire dynamics. Osprey explicitly accounts for these effects by applying a spatial smoothing Gaussian function (whose width is determined by the resolution of the image) on the known tread pattern. Further, rather than applying a standard matched filter, Osprey applies a 2D version of Dynamic Time Warping [31] used in speech and image processing to correct for minute spatial stretches and squeezes of the signal received from the tire. In effect, these allow Osprey to spatially map the precise locations of the surface of the tire between the grooves. Osprey then averages the depth information (defined by r) obtained at these locations across rotations of the tires to report r_{tread} , the location of tread surface.

Background Subtraction: An additional signal processing tool that Osprey employs to combat signal multipath from spurious objects around the tire is background subtraction [55]. Osprey subtracts out received signals along two different time windows to preserve dynamic artifacts (e.g. the tire) while canceling out static objects (e.g. the well of the tire). This effectively removes much of the static (relative to the car) objects surrounding the tire such as the well of the tire and the surface of the car itself – modulo vibrations, as we discussed in Sec. 5.1. Background subtraction when applied to two adjacent ISAR images across time can also reveal another effect – spurious objects that appear on the surface of the tire, such as debris (e.g. mud picked up by the tire) that soon after dissipates due to abrasion. Osprey can effectively be resilient to such distortions to surface depth measurements by identifying and rejecting these outliers.

An important point to discuss is the difference in the effect of debris on the tire surface versus the groove. While debris on the tread surface inevitably is worn away due to abrasion resulting in (at worst) local and short-term uncertainty, debris on the groove of the tire can settle in and create long-term errors in measurement. The next section explicitly discusses our solution to this problem.

6 DEBRIS-RESILIENT TREAD SENSING

This section studies Osprey’s approach to sensing the precise depth of the grooves of a tire. Coupled with the estimate of the depth of the tire surface in Sec. 5, this would allow Osprey to fully estimate tread depth. The primary challenge is the presence of debris (e.g. mud, stones, soil, etc.) within the grooves, which may cause spurious reflections of the mmWave radar signal that completely mask the true reflected signal from the tire grooves.

Osprey addresses this challenge by developing specialized mmWave compatible metallic codes embedded in the tire groove. We leverage the fact that mmWave reflects off metallic surfaces with a significantly higher power compared to other surfaces (e.g. mud, rubber, etc.), a property validated by prior work [15, 24] and our own results (see Fig. 11 in Sec. 9.1). We place thin and tiny metallic strips of Aluminum that are inexpensive and highly reflective for mmWave in the tire’s groove (see Fig. 7). We place these strips in 2 dimensions resembling barcodes with a pre-designed and well-known layout akin to a spatial code. These codes serve twin purposes: (1) First, we can modify our algorithm to explicitly look for reflections from a specific code in a specific groove within the reflected signal from the tire. This would help isolate the signal from the groove of the tire from all other reflections (e.g. debris within the groove), both due to high reflectivity of the metal and the coding gain of the spatial code. (2) Second, we can reuse these metallic tire codes placed at known groove locations as encoders to accurately measure (with mm-accuracy), the current azimuth $\phi(t)$ (i.e. rotational angle) of the tire. In other words, metallic codes can free us from needing out-of-band encoders to find the rotational angle of the tire, needed for Osprey’s ISAR algorithm in Sec. 5.1.

The rest of this section explores the design space of metallic codes. First, we describe our choice of tire code that accounts for the limited space available within the groove, resilience to error, and interference. Second, we present our algorithm to efficiently decode the metallic codes and estimate tread depth.

6.1 Design of Tire Codes

In designing metallic codes in the tire groove, our objective is to determine the depth of the groove with high accuracy, despite the limited area of the groove itself and the potential presence of debris. In exploring the design space of spatial codes, we pose three important requirements unique to the tire context: (1) First, we need a solution that is resilient to debris by ensuring that the code reflects mmWave radiation strongly; (2) Second, we need to be resilient to errors that stem from foreign objects or debris lodged in the groove; (3) Third, we need to be able to decode and disambiguate signals from codes along adjacent grooves, the reflected signals from which may interfere at the radar. We describe how our design meets each of these requirements below.

Debris-Resilient Code Modulation: Osprey’s spatial code needs a modulation that maps zeros and ones to metallic structures that maximizes resilience to debris. Intuitively, we need to select modulations that maximize the surface area of thin metallic strips needed to ensure best resilience to debris. Our solution is inspired by traditional 1D barcodes that use pulse-width modulation [27] for similar reasons: to maximize the amount of ink clearly observable and resolvable by a camera. Osprey uses pulse width modulation, mapping zeros and ones to metallic strips of different widths. We choose




Figure 7: Osprey combats debris by laying Aluminium strips in the groove, which emulate a spatial code. Different grooves have different coded bit patterns encoded using pulse width modulation.

widths of 4 and 2 cm dictated by the resolution of the mmWave radar. We orient these tags along the azimuth axis, which is our axis of maximum spatial resolution owing to ISAR.

We pack as many bits as possible within the available surface area of the groove. Our design largely uses one-dimensional codes, given that most grooves on tires are relatively narrow. However, should larger grooves be available (for e.g. on certain large truck or bus tires), we can lay out the patterns along two dimensions. However, these patterns need to be spaced based on the resolution along the azimuthal angle (about 3 mm in our implementation). Hence, our current implementation in Sec. 8 uses 1D codes.

We choose Aluminum foils with a thickness of 274 microns, to ensure minimal impact on tire dynamics while still providing strong surface reflections to the mmWave radar. We note that today’s tire manufacturing processes already insert metallic structures within the tire for structural stability. Our design of laying Aluminium foils in the groove introduces minimal disruption to existing tire manufacturing pipelines.

Resilience to Bit Errors and Code Collisions: Next, our system needs to be resilient to bit errors as well as collisions between codes from adjacent grooves. At first blush, one may consider using a family of orthogonal codes from traditional CDMA codes (e.g. Walsh codes [43], Gold codes [45], etc.) that are known to be resilient to collisions and have excellent resilience to bit flips. However, these codes assume that bits are encoded into polar modulations where bits are encoded as -1 and 1 . In the case of spatial codes, however, bits are encoded by different widths of the metallic strip for pulse width modulation. In other words, bits are encoded in the relative *amplitude* of the signals reflected off the code rather than the phase or sign. Said differently, Osprey needs orthogonal CDMA codes that allow for only positive (unipolar) modulations where bits are modulated as $(0, 1)$ instead of $(-1, 1)$. To illustrate why this difference in polarity of codes is important, note that $(-1, 1)$ and $(1, 1)$ are orthogonal polar codes of length two (i.e. their vector dot product is zero), while their unipolar twins $(0, 1)$ and $(1, 1)$ are not.

To address this challenge, Osprey borrows from optics where, similar to our context, signals add up in amplitude. Osprey specifically relies on Optical Orthogonal codes (OOC) [10], a family of coded bits of different lengths that are designed to have high auto-correlation and poor cross-correlation. Osprey specifically chooses OOC to maximize the amount of metal and remain within the space constraints of common grooves. We further note that our approach




Figure 8: Osprey was evaluated on a tire rotation rig. The tire, with one bit of Aluminium strip embedded in the central groove, is mounted on a turntable controlled by stepper motors. mmWave radar is mounted at a distance on a table, similar to the tire well, and streams I/Q samples to computer.

is compatible with pulse width modulation, given that it effectively results in marked differences in total amplitude reflected from zero and one bits. We note that our choice of codes simultaneously achieves two purposes: (1) OOC have poor cross-correlation and therefore are highly resilient to collisions when codes across adjacent grooves need to be disambiguated; (2) OOCs are also inherently resilient to bit flips showing high robustness to erroneous bits as demonstrated in [10].

6.2 Decoding of Tire Codes

Our approach to decoding tire codes correlates the codes with different possible known code sequences in order to detect the presence of a specific groove. This directly serves as an encoder as well. The main challenge, however, is to effectively measure the precise depth of the groove, given the known code pattern that is present. Further, this code pattern could experience distortions (e.g. smoothing) owing to the limited resolution of the mmWave radar and the dynamics of tire rotation.

Osprey explicitly accounts for this calibration by developing a model $M_{(r,\phi(t),z)}(C)$ that captures the expected wireless channels from the code C accounting for the expected distortion, when moving along the trajectory $(r, \phi(t), z)$. We then pose a maximum likelihood problem that determines the true depth of the groove given by r by correlating this model with the received channels. Specifically, we estimate the groove’s coordinates as:

$$r_{\text{groove}} = \underset{r}{\operatorname{argmax}} \operatorname{corr}(M_{(r,\phi(t),z)}(C), h(t))$$

We then subtract r_{groove} from our prior measurement of r_{tread} at the tread surface from Sec. 5.3 to compute tread depth.



Figure 9: Osprey was also evaluated on a passenger car. Osprey’s hardware is shown to be mounted in the tire well of a 2019 Honda Odyssey.

7 FOREIGN OBJECT SENSING

This section describes our approach to locating and detecting foreign objects lodged in the tire. Our primary approach is to detect and locate anomalies in the ISAR image that appear due to the presence of foreign objects. We then study the shape, intensity, and phase corresponding to these anomalies to classify the type of foreign object.

7.1 Locating Foreign Objects

To locate foreign objects on the surface, Osprey relies on background subtraction to constantly monitor any new reflectors that appear in the mmWave ISAR image. It then locates the (r, ϕ, z) location of objects that appear in the ISAR image and remain persistent when averaged across multiple frames. Given its high spatial resolution, ISAR image captures as small as 3 mm on the tire’s surface, which captures the vast majority of dangerous objects that may impede a tire’s structural integrity.

An important challenge in detecting foreign objects deep in the groove is that they may not reflect strongly under mmWave. Here, Osprey relies on the presence of known metallic coded patterns in the groove. Dangerous foreign objects made of metal will inevitably interact with mmWave and make their presence known, even in the groove by disrupting the shape of the code. However, even weaker reflectors of mmWave (e.g. sharp stones or glass) that penetrate into the tire via the groove are likely to wear out the metallic codes in the process of damaging the tire. This manifests as bit errors in the code when processed by Osprey’s algorithm in Sec. 6. Osprey therefore treats bit errors as potential foreign objects and passes on the received signals at the specific location within the groove to the classification algorithm below to detect and classify the object type.

7.2 Classifying Foreign Object Type

Osprey classifies object type by relying on both the magnitude and the phase of the received signal at a specific (r, ϕ, z) location as obtained from Eqn. 1. We rely on three specific properties of the foreign object’s impact on radar signals: (1) The amplitude of reflection (stronger for metallic objects); (2) The phase which

captures object reflectivity; (3) The shape and size of the object as it appears on the ISAR image. We choose a simple linear binary class classifier using Gaussian Mixture Models, which achieved the highest accuracy computed through cross-validation on the dataset among different models compared. Sec. 9.5 presents the results from an evaluation of our system.

8 IMPLEMENTATION AND EVALUATION

Hardware: We implement Osprey on commodity 77 GHz automotive radar developed by Texas Instruments – the TI AWR1642. The chip uses an FMCW chirp modulated between 77 GHz and 81 GHz with maximum allowable bandwidth of 4 GHz. To provide maximum flexibility, we use the AWR1642BOOST evaluation board that allows for quick integration with a processing pipeline and enables easy debugging. We interface this board with the DCA1000EVM FPGA board to collect and process raw I/Q samples from the radar board in real-time. The maximum speed our system can support is limited by the minimum allowed chirp periodicity time. Our system is hardware limited to tire speeds of about 12 kmph, for a tire outer radius of 33.5 cm. We are also limited in our ability to evaluate by the maximum reliable rotation speed of the motor - 5.45 kmph. However, with relaxed sampling rates of > 1 mm, higher speeds can be supported (see Sec. 10).

Software: Our implementation of Osprey’s algorithm is written in MATLAB. We assume that the tread pattern and the placement of metallic codes on the surface, developed using Aluminum strips as described in Sec. 6 (see Fig. 7), are known. The system outputs the following: (1) The tread depth of the tire at each groove where the corresponding codes are attached; (2) Detection and location of foreign objects.

Evaluation: We evaluate Osprey on two distinct test-beds (see Fig. 8, 9): (1) [Sec. 9.2-9.3 and 9.5] A mechanical rotational rig that mimics rigs used to test tire performance. The rig uses a Nema 23 stepper motor with a microstepping driver that allows fine-grained control of rotational angle and reliable speeds up to 5.45 kmph of the tire. We use the rig for the bulk of our experiments primarily due to safety reasons when experimenting with extreme tread wear and to ensure fine-grained control of variables (e.g. speed, debris depth, etc.). (2) [Sec. 9.4] We also present system performance on a large passenger car (a 2019 Honda Odyssey). Across experiments, we use the following models of tires: Falken Eurowinter HS449 and Bridgestone Blizzak LM001 – with diverse tread patterns to show system generality. Further, across experiments, we consider tread depths between 1.7 and 8.75 mm, ensuring diversity by adding thin layers of rubber in the groove. We evaluate the system under different types and amounts of debris, foreign objects, and tire dynamics. Note that unless specified otherwise, error bars denote first standard deviation.

Pipeline: Fig. 10 shows the entire pipeline for testing on rotation rig. Based on the desired speed of rotation, the radar chirp configurations and motor signals are created and sent to DCA1000EVM and motor controller respectively. AWR1642BOOST samples the reflections, and DCA1000EVM streams the packets over Ethernet to the host computer which runs the tread depth estimation algorithm and foreign object detection and localization. Similarly, for testing on car, we follow everything discussed above except the




Figure 10: Top view of Fig. 8. Block diagram showing the signal flow pipeline for tests performed on the rotation rig.

motor controller. Here, the driver is instructed to maintain a desired, sufficiently uniform, slow speed. Streaming from DCA1000EVM provides the option of connecting to the car’s computer and running algorithms.

Ground Truth and Baseline: To obtain ground truth, we first clear up the tire groove of any debris and use a digital Vernier caliper. We reintroduce any debris prior to running Osprey. For comparison in certain experiments (those involving debris), we also present accuracy numbers of a LIDAR-based Bosch GLM40 sensor as a baseline to demonstrate that the debris introduced impedes light-based measurement.

9 RESULTS

This section evaluates Osprey under debris, tire dynamics, and foreign objects. Note that besides Sec. 9.4, which evaluates vehicle dynamics on a passenger car, all other sections are evaluated on the mechanical rig for safety and repeatability.

9.1 Microbenchmark – Code Material

Method: In this microbenchmark experiment, we motivate our choice of Aluminum foil for the design of Osprey spatial codes for debris-resilient tread depth sensing. We place thin strips (1 mm thin) of six different materials: Aluminum foil, paper, rubber, cloth, plastic, and cardboard at distances of 29 cm directly below and facing the mmWave radar. We then measure the average power of the reflected signal per material across experiments.

Results: Fig. 11 plots the mean and standard deviation of the reflected power across materials. Aluminum has a clear advantage over other materials (a mean gain of ≈ 3 dB or equivalently, $2\times$ higher power). We therefore choose Aluminum strips for Osprey’s design.

9.2 Tread Depth with Debris

Method: In this experiment, we evaluate the performance of Osprey when debris is placed in tire groove. We first vary the amount of debris ranging from 3 mm to 8 mm thickness (34%-91% of tread




Figure 11: mmWave reflectivity of various materials




Figure 12: Tread Depth Error vs. Debris Level




Figure 13: Tread Depth Error vs. Debris Type




Figure 14: Tread Depth Error vs. Tire Speed




Figure 15: Tread Depth Error vs. Vehicle Vibrations




Figure 16: Tread Depth Error vs. Terrain Type

depth covered). We add sawdust as debris in the groove. We note readings from diverse trials in which codes are uncovered, partially covered, and completely covered with sawdust. To validate ground truth, we use a digital Vernier caliper. We additionally measure the error of a Bosch GLM40 laser rangefinder. Next, we evaluate Osprey with different types of commonly found debris: sawdust, mud, snow, and tiny stones that can be lodged in the groove of the tire.

Results: Fig. 12 measures Osprey’s error in tread depth for different levels of sawdust. We note that Osprey achieves a maximum error of 1.53 mm for the maximum thickness of debris (8 mm), remaining highly accurate with expected but small levels of degradation as debris is added. We also note that the baseline LIDAR system performs poorly under debris as expected with errors increasing from 3.8 mm to 8 mm as more debris is attached, essentially detecting debris level rather than true tread depth.

Fig. 13 measures the performance of Osprey in terms of error in tread depth for different types of debris. Observe that Osprey achieves a maximum tread depth error of 0.85 mm, remaining highly accurate with minimal degradation across different types of debris. Among the different types of debris, we see that sawdust and mud achieves the lowest performance (which is still 0.85 mm accurate). While it is easy to see that a high density object like stone can block the mmWave significantly, moist sawdust also attenuates mmWave due to its moisture content and particle size [12].

9.3 Tread Depth vs. Tire Speed

Method: In this experiment, we measure the accuracy of tread depth at different tire wear levels across different tire speeds. We perform this experiment on the mechanical rig to ensure repeatability and fine-grained control of the speed. We consider speeds of 0.62, 2.32 and 5.45 kmph of the tire. Strictly, the speed of interest is the rotational speed, but for convenience we report in kmph as well. We vary the tread depths by adding rubber in the groove of a Falken Eurowinter HS449 tire. The unaltered tread depth is 8.75 mm. We vary it to two other levels - one which is around the legal limit (1.7 mm) and another above the legal limit for passenger cars (5 mm).

Results: Fig. 14 shows that Osprey’s performance in terms of tread depth has a maximum error of 1.5 mm across tire speeds. On the other hand, as the speed increases, we see a general increase in error averaged across true tread depths. While our radar adapts the chirp periodicity parameter such that it always samples the tire surface at a fixed spatial sampling rate of 1 mm to ensure similar Inverse SAR image quality is maintained across different speeds (see Sec. 10), small sampling offsets result in this increasing error trend. Despite this trend, we note that the error is small irrespective of the true tread depth.

9.4 Impact of Vehicle Dynamics on Osprey

Method: In this section, we evaluate the performance of Osprey on a passenger car with the mmWave radar attached to the tire well of the car as shown in Fig. 9. We once again measure error




Figure 17: Foreign Object Localization Error

in tread depth by varying two distinct parameters: (1) First, we consider different levels of vibrations of the car by driving it on three different bumpy surfaces. These correspondingly induce low, medium, and high levels of vibrations. (2) Second, we consider the performance of Osprey when picking up natural debris on the road by driving the Osprey-enabled vehicle on a paved flat surface with no debris, a stony surface, a wet surface mimicking road after rain, and a slushy surface created by mixing snow and mud mimicking road after snowfall.

Results: Fig. 15 shows that the tread depth estimation in Osprey has average error of 0.21, 0.68, and 0.73 mm respectively for low, medium, and high levels of vibration. This demonstrates Osprey’s robustness to road-related vibrations and changes in vehicle suspension.

Next, Fig. 16 shows that average tread depth estimation error in Osprey is 0.21, 0.68, 0.52, and 0.19 mm respectively for paved flat, stony, wet and slushy surfaces. We note that the maximum error is 0.68 mm, showing that our system is robust to different kinds of terrain. We note that interestingly, our experiments also resulted in the mmWave radar itself being splattered with debris as the vehicle moved. However, our system remains robust to this effect. In contrast, a LIDAR system placed in the tire well reported error messages and refused to output range readings due to debris that obstructed the LIDAR’s light sensor itself when driven on all the above types of terrains.

9.5 Sensing Foreign Objects

Method: In this section, we measure Osprey’s accuracy in detecting and locating foreign objects placed in the groove of a tire. Note that we assume that grooves have Osprey spatial codes attached to them according to a known specification. We consider the following types of objects – nail, stone, sawdust, and ice cubes as foreign objects embedded in the groove of the tire. The sizes of the foreign object range from 5 mm to 25 mm. We run 50 trials with these objects in the tire at constant tire speed. Across these trials the object type and its location are varied.

Results: Fig. 17 depicts the CDF of the absolute error in estimate of the foreign object’s reported location from its true location. We note that this error in the location is 17 mm, comparable to the

size of the foreign objects themselves and more than sufficiently accurate to visually inspect the tire and find the objects.

Next, we trained multiple classifiers to classify the foreign object into 2 categories: harmful (metallic objects) and non-harmful (the rest) using the features described in Sec. 7. We choose the Gaussian Mixture Model which achieves the highest classification accuracy of 91.67% between metallic and non-metallic foreign objects.

10 DISCUSSION

Speed: We show that our system can operate at tire speeds up to 5.45 kmph owing to the maximum reliable speed supported by our motor. Moreover, as explained in Sec. 8, we adapt the chirp periodicity for different speeds. As speed increases, we lower the chirp periodicity. The minimum chirp periodicity supported by the radar hardware limits maximum speed supported by our system to 12 kmph. In order to support higher speeds, we either need a hardware which can support lower chirp periodicity or we can set the chirp periodicity to be the minimum supported by the system and trade off on spatial sampling. This means that instead of spatially sampling the surface of tire at 1 mm (as mentioned in Sec. 8), we relax it to higher values and sample coarsely. Higher spatial sampling values can lead to errors and potentially aliasing effects. But, we note that natural tire wear occurs at relatively slow timescales. Our system, although operating at slow speeds, provides a design which enables obtaining periodic measurements continuously as and when vehicles move slowly. This continues to be very informative and is reliable and convenient compared to manual sensing approaches.

For the purposes of this work, we assume we know the speed of rotation. We use this speed information to choose a chirp periodicity. The speed information can be obtained from the computer in the vehicle or from other sensors that can be mounted on the tire. Almost similar to the principle of an optical encoder, we can reuse the spatial codes in the grooves as an encoder to infer the speed through radar processing.

Another assumption we make is that the speed is sufficiently uniform. This is important because we do not have the ability to configure the chirp periodicity at very fine time intervals as and when the measured speed changes. Hardware with such capabilities can potentially deal with non-uniform speeds.

Debris-Resilient Metallic Structures: In our system, we lay the metallic structures in the grooves. Note that grooves don’t directly come in contact with the terrain. Unlike in-tread sensors, which have sensors embedded in the tread during manufacturing (requiring changes to assembly lines), we envision these structures to be placed in the grooves post-manufacturing.

Because our metallic structures are thin paper-like Aluminium foils, if they come in contact with a stone or a sharp object, they can get peeled off. This is possible to detect by monitoring the spatial code and observing changes. Once peeled, the spatial code has to be replaced. Adding metallic structures could also affect some functionalities of the groove such as ability to channel water.

To address these deficiencies, we believe it is possible to reuse existing steel belts in tires and their pattern as a spatial code instead. These belts are located just below the grooves (see Fig. 2). Tread depths calculated with respect to these will have a small and fixed

offset which can be accounted for. We believe this is an important topic for future work to explore.

Impact of Tire Geometry: We evaluate our system on two different passenger car tires. Our system uses information about the tread patterns and layout of spatial codes. Such information can be obtained from 3D models of tires that can be provided by manufacturers during purchase. In addition, our system needs to be re-calibrated when tires are replaced during service.

While both tires had different tread patterns, they had a groove running through the circumference of the tire. This allowed us to lay out our codes along the azimuth. These type of grooves are commonplace as they are essential for channeling water [51]. We leave it for future work to address other types of tread and groove structures.

Impact of Debris on Radar: While we focus on road debris getting accumulated in the grooves, debris accumulation on radar could affect measurements. However, mmWave is shown to propagate through dust and other adverse atmospheric conditions [6]. We also notice that our system remains robust when we observed debris splattered on the radar in Sec. 9.4.

Cost of System: Our system is based on TI AWR1642 mmWave sensor. The cost of this IC is \approx \$40 [41]. We note that this cost-overhead is insignificant for trucking companies, where the need for tire replacement alerts far outvalue system cost [16].

11 RELATED WORK

Past solutions explored for tire sensing are very diverse and use different sensing modalities, broadly grouped in the following categories.

Manual and Indirect Tread Sensing: The most common approach to tread depth measurement is manual – using a coin or ruler to measure tread depth. However, this approach requires human effort, doesn't provide continuous measurements, and is prone to error. Our approach instead seeks to provide a continuous and accurate mechanism to measure tread depth. [4, 11] uses tire pressure from TPMS and [28] uses in-tread temperature and pressure sensors to map indirectly to tread depth. While the methods have the advantage of reusing preexisting sensors, they are known to be crude and only good enough to alert that a tire change is needed [11].

In-Tread and Off-Tread Sensors: There has been a lot of activity to develop in-tread sensors exclusively for tread depth based on RFID [19, 33], 2.4GHz [46], and Surface Acoustic Waves [40]. Embedding these sensors would require manufacturing lines to leave a tread lug empty, manufacture the rest of the tire, embed the sensor in a lug outside manufacturing lines, and then cure the lug with the rest of the tire. This could decrease the efficiency of manufacturing lines, increase the time cost, and possibly negatively affect the economics of making a tire.

In cases where tread depth is measured occasionally but highly accurately over a specific region on tread, vehicles drive over laser scanners on the ground [25, 48]. There are other solutions in which laser scanners are installed in manholes on roads and scan for tread depth on all tires which go by [38]. Unfortunately, such solutions require the tire to be free of debris in the groove as they can corrupt measurements. Radio based off-tread solutions [3] lay an active

transmitter and a receiver sensor array over which the tire is driven. Received signal changes with any change in electric medium (capacitance) due to wear. Similar to other off-tread solutions, this too does not provide an infrastructure which allows to obtain tire wear in a periodic continuous fashion - which is important from the perspective of safety and performance. Our objective is to obtain continuous, robust to debris, accurate tread depth measurement.

mmWave Localization and Sensing: There has been rich prior work on using mmWave radios for sensing. Commercial applications of mmWave sensing include full-body scanning [22] and detecting automobile collisions [23, 24]. There has also been recent past work on using commodity mmWave radars for varied applications, such as: gesture-based interfaces [26, 49], object tracking [17], scene imaging [14], and beyond. Several past solutions have also sought to extend mmWave communication radios for sensing, e.g. for accurate localization [1, 7, 13] and sensing [35, 53, 54]. However, in our context, insufficient range resolution and debris offer challenges. Osprey complements this literature and develops a super-resolution algorithm specifically for the tire wear sensing context and achieves sub-mm accuracy in tread depth. Osprey is most closely related to recent work on mmWave barcodes [20, 21] used for product identification and more broadly with literature on chipless RFIDs and backscatter [29, 30] which require no electronics. The objectives of these – no electronics and tag identification in non line of sight conditions – are similar to our requirements. Osprey builds on and complements this literature by designing spatial codes and a decoding algorithm for the unique problem of tire sensing amidst debris and tire dynamics.

12 CONCLUSION

This paper presents Osprey, the first mmWave tire wear sensing system that measures accurate tread wear continuously even in the presence of debris. Osprey achieves this by imaging the tire's surface using a mmWave radar system at a resolution of sub-mm, an order of magnitude over its advertised resolution. We achieve this through a super-resolution Inverse Synthetic Aperture Radar algorithm that exploits the natural rotation of the tire. We further devise mechanisms to eliminate the impact of debris lodged in the tire when measuring tread depth. We demonstrate how our system detects and locates unsafe foreign objects lodged in the tire. A detailed evaluation of Osprey on commercial car and tire rotation rig reveals 0.68 mm error in tread depth estimation, 1.7 cm error in locating and 92% accuracy in detecting harmful foreign objects. While our current approach requires adding metallic structures in the groove, future work will explore leveraging the diversity of materials already within the tire to filter out debris. We will further explore abrasion of other surfaces beyond a tire, such as paint on aircraft or anti-fouling systems in ships.

ACKNOWLEDGMENTS

This research was supported in part by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA, the National Science Foundation (1823235, 1942902), Kavčić-Moura grant and Bridgestone. We thank our shepherd, Guoliang Xing and other reviewers for their insightful feedback which helped improve this paper. We thank all members of WiSE Lab and WiTech Lab at CMU for their inputs.

REFERENCES

- [1] Zohair Abu-Shaban, Xiangyun Zhou, Thushara Abhayapala, Gonzalo Seco-Granados, and Henk Wymeersch. 2018. Error bounds for uplink and downlink 3D localization in 5G millimeter wave systems. *IEEE Transactions on Wireless Communications* 17, 8 (2018), 4939–4954.
- [2] Saman Hassanzadeh Amin, Guoqing Zhang, and Pervaiz Akhtar. 2017. Effects of uncertainty on a tire closed-loop supply chain network. *Expert Systems with Applications* 73 (2017), 82–91.
- [3] Joseph B Andrews, Peter Ballentine, Jorge A Cardenas, Chin Jie Lim, Nicholas X Williams, James B Summers, Michael A Stangler, David A Koester, Steven A Cummer, and Aaron D Franklin. 2019. Printed Electronic Sensor Array for Mapping Tire Tread Thickness Profiles. *IEEE Sensors Journal* 19, 19 (2019), 8913–8919.
- [4] Continental Automotive. 2020. Electronic-Tire Information Systems (eTIS). Retrieved April 30, 2020 from <https://www.continental-automotive.com/en-gl/Passenger-Cars/Vehicle-Networking/Comfort-Security/Tire-Information-Systems/Electronic-Tire-Information-System>
- [5] Thomas A Brey. 2007. Tire tread wear sensor system. US Patent 7,180,409.
- [6] Graham Brooker, Ross Hennessey, Craig Lobsey, Mark Bishop, and Eleonora Widzyk-Capehart. 2007. Seeing through dust and water vapor: Millimeter wave radar sensors for mining applications. *Journal of Field Robotics* 24, 7 (2007), 527–557.
- [7] Joe Chen, Daniel Steinmetzer, Jiska Classen, Edward Knightly, and Matthias Hollick. 2017. Pseudo lateration: Millimeter-wave localization using a single RF chain. In *2017 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 1–6.
- [8] Long Chen, Mingyuan Bian, Yugong Luo, Zhaobo Qin, and Keqiang Li. 2016. Tire-road friction coefficient estimation based on the resonance frequency of in-wheel motor drive system. *Vehicle System Dynamics* 54, 1 (2016), 1–19.
- [9] E-H. Choi. (2012, April). Tire-Related Factors in the Pre-Crash Phase. (Report No. DOT HS 811 617). Washington, DC: National Highway Traffic Safety Administration.
- [10] Fan RK Chung, Jawad A Salehi, and Victor K Wei. 1989. Optical orthogonal codes: design, analysis and applications. *IEEE Transactions on Information theory* 35, 3 (1989), 595–604.
- [11] Continental. 2014. Future tire pressure sensors read pressure, load and tread depth. Retrieved April 30, 2020 from <https://www.continental.com/en/press/press-releases/2014-05-07-tpms-profile-104542>
- [12] Katia Cosentino, Amerigo Beneduci, Alfonsina Ramundo-Orlando, and Giuseppe Chidichimo. 2013. The influence of millimeter waves on the physical properties of large and giant unilamellar vesicles. *Journal of biological physics* 39, 3 (2013), 395–410.
- [13] Hua Deng and Akbar Sayeed. 2014. Mm-wave MIMO channel modeling and user localization using sparse beamspace signatures. In *2014 IEEE 15th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. IEEE, 130–134.
- [14] Shiwei Fang and Shahriar Nirjon. 2018. AI-Enhanced 3D RF Representation Using Low-Cost mmWave Radar. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*. ACM, 414–415.
- [15] Reinhard Feger, Clemens Pfeffer, Werner Scheibhofer, Christian M Schmid, Markus J Lang, and Andreas Stelzer. 2012. A 77-GHz cooperative radar system based on multi-channel FMCW stations for local positioning applications. *IEEE Transactions on Microwave Theory and Techniques* 61, 1 (2012), 676–684.
- [16] Fred Ignatz-Hoover, Mark Arigo, Dominica Hiu Ching Wong, and Leandro Forciniti. 2018. How market drivers fuel tire additive innovations. *Rubber and Plastics News* (2018).
- [17] Muhammad Z Ikram and Murtaza Ali. 2013. 3-D object tracking in millimeter-wave radar for advanced driver assistance systems. In *2013 IEEE Global Conference on Signal and Information Processing*. IEEE, 723–726.
- [18] Karine Le Gorju Jago. 2012. X-ray computed microtomography of rubber. *Rubber Chemistry and Technology* 85, 3 (2012), 387–407.
- [19] Kati Körbe Kaare, Kristjan Kuhl, and Ott Koppel. 2012. Tire and pavement wear interaction monitoring for road performance indicators. *Estonian Journal of Engineering* 18, 4 (2012).
- [20] Nimai Chandra Karmakar and Chow Kay Pern. 2011. mm-wave chipless RFID tag for low-cost item tagging. In *Asia-Pacific Microwave Conference 2011*. IEEE, 1462–1465.
- [21] Zhengxiang Li, Baicheng Chen, Zhuolin Yang, Huining Li, Chenhan Xu, Xingyu Chen, Kun Wang, and Wenyao Xu. 2019. FerroTag: A Paper-based mmWave-scannable Tagging Infrastructure. In *Proceedings of the 17th Conference on Embedded Networked Sensor Systems (New York, New York) (SenSys '19)*. ACM, New York, NY, USA, 324–337.
- [22] Benedetta Mencagli, R Vincenti Gatti, Luca Marcaccioli, and Roberto Sorrentino. 2005. Design of large mm-wave beam-scanning reflectarrays. In *The European Conference on Wireless Technology, 2005*. IEEE, 475–478.
- [23] Wolfgang Menzel, Dietmar Pilz, and Ralf Leberer. 1999. A 77 GHz FM/CW radar frontend with a low-profile, low-loss printed antenna. In *1999 IEEE MTT-S International Microwave Symposium Digest (Cat. No. 99CH36282)*, Vol. 4. IEEE, 1485–1488.
- [24] Toshiya Mitomo, Naoko Ono, Hiroaki Hoshino, Yoshiaki Yoshihara, Osamu Watanabe, and Ichiro Seto. 2010. A 77 GHz 90 nm CMOS transceiver for FMCW radar applications. *IEEE journal of solid-state circuits* 45, 4 (2010), 928–937.
- [25] Andrew Nevin and Eric Daoud. 2014. *Evaluation of Advanced Machine-Vision Sensors for Measuring the Tread Depth of a Tire*. Technical Report. SAE Technical Paper.
- [26] Avishek Patra, Philipp Geuer, Andrea Munari, and Petri Mähönen. 2018. mm-Wave Radar Based Gesture Recognition: Development and Evaluation of a Low-Power, Low-Complexity System. In *Proceedings of the 2nd ACM Workshop on Millimeter Wave Networks and Sensing Systems*. ACM, 51–56.
- [27] Songwen Pei, Guobo Li, Gang Chen, and Baifeng Wu. 2008. Prototyping system of codec for novel 2-D continuous barcode. In *2008 12th International Conference on Computer Supported Cooperative Work in Design*. IEEE, 1128–1132.
- [28] Pirelli. 2020. Meet the Cyber Car. Retrieved April 30, 2020 from <https://www.pirelli.com/global/en-ww/road/meet-the-cyber-car>
- [29] Stevan Preradovic, Isaac Balbin, Nimai Chandra Karmakar, and Gerhard F Swiegers. 2009. Multiresonator-based chipless RFID system for low-cost item tracking. *IEEE Transactions on Microwave Theory and Techniques* 57, 5 (2009), 1411–1419.
- [30] Stevan Preradovic and Nimai Chandra Karmakar. 2010. Chipless RFID: Bar code of the future. *IEEE microwave magazine* 11, 7 (2010), 87–97.
- [31] Toni M Rath and Raghavan Mammatha. 2003. Word image matching using dynamic time warping. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, Vol. 2. IEEE, II–II.
- [32] Richard Roy and Thomas Kailath. 1989. ESPRIT-estimation of signal parameters via rotational invariance techniques. *IEEE Transactions on acoustics, speech, and signal processing* 37, 7 (1989), 984–995.
- [33] Navtej S Saini, Shuai Shao, Asimina Kiourti, Robert J Burkholder, and John L Volakis. 2016. RFID tags for in-situ tire monitoring. In *2016 URSI International Symposium on Electromagnetic Theory (EMTS)*. IEEE, 575–578.
- [34] Sergio M Savaresi, Mara Tanelli, Peter Langthaler, and Luigi Del Re. 2008. New regressors for the direct identification of tire deformation in road vehicles via “in-tire” accelerometers. *IEEE Transactions on Control Systems Technology* 16, 4 (2008), 769–780.
- [35] Stefano Savazzi, Stephan Sigg, Federico Vicentini, Sanaz Kianoush, and Rainhard Finding. 2019. On the Use of Stray Wireless Signals for Sensing: A Look Beyond 5G for the Next Generation of Industry. *Computer* 52, 7 (2019), 25–36.
- [36] Kabir Siddiqui. 1999. Z-encoder mechanism. US Patent 5,912,661.
- [37] Kanwar Bharat Singh and Saied Taheri. 2019. Accelerometer Based Method for Tire Load and Slip Angle Estimation. *Vibration* 2, 2 (2019), 174–186.
- [38] smallformfactors.mil embedded.com. 2020. Tiny modules measure tire treads. Retrieved April 30, 2020 from <http://smallformfactors.mil-embedded.com/articles/tiny-measure-tire-treads/>
- [39] Gerd Sommerkom, Dirk Hampicke, Ralf Klukas, Andreas Richter, Axel Schneider, and Reiner Thoma. 1999. Reduction of DoA estimation errors caused by antenna array imperfections. In *1999 29th European Microwave Conference*, Vol. 2. IEEE, 287–290.
- [40] R Steindi, C Hausleitner, Hans Hauser, and W Bulst. 2000. Wireless magnetic field sensor employing SAW-transponder. In *ISAF 2000. Proceedings of the 2000 12th IEEE International Symposium on Applications of Ferroelectrics (IEEE Cat. No. 00CH37076)*, Vol. 2. IEEE, 855–858.
- [41] ti.com. 2020. AWR1642. Retrieved April 30, 2020 from <https://www.ti.com/product/AWR1642#order-quality>
- [42] Roger H Tracy, Edwin H Reeves, Nicholas J Radclyffe, and Robert Mark Longden. 2004. Hand held probe for measuring tire tread wear. US Patent 6,789,416.
- [43] Shiauhe Tsai, Farideh Khaleghi, Seong-Jun Oh, and Vieri Vanghi. 2001. Allocation of Walsh codes and quasi-orthogonal functions in cdma2000 forward link. In *IEEE 54th Vehicular Technology Conference. VTC Fall 2001. Proceedings (Cat. No. 01CH37211)*, Vol. 2. IEEE, 747–751.
- [44] David Tse and Pramod Viswanath. 2005. *Fundamentals of wireless communication*. Cambridge university press.
- [45] Adel MD Turkmani and US Goni. 1993. Performance evaluation of maximal-length, Gold and Kasami codes as spreading sequences in CDMA systems. In *Proceedings of 2nd IEEE International Conference on Universal Personal Communications*, Vol. 2. IEEE, 970–974.
- [46] Roesler Tyres. 2020. Roesler Tyres INTREAD SENSOR. Retrieved April 30, 2020 from <https://www.roesler-tyres.com/en/otr-tyre-sensor/>
- [47] Sankaranarayanan Velupillai and Levent Guvenç. 2007. Tire pressure monitoring [applications of control]. *IEEE Control systems magazine* 27, 6 (2007), 22–25.
- [48] Xi-Bo Wang, Ai-Juan Li, Qin-Peng Ci, Meng Shi, Tian-Long Jing, and Wan-Zhong Zhao. 2019. The study on tire tread depth measurement method based on machine vision. *Advances in Mechanical Engineering* 11, 4 (2019), 1687814019837828.
- [49] Teng Wei and Xinyu Zhang. 2015. MTrack: High-Precision Passive Tracking Using Millimeter Wave Radios. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (Paris, France) (MobiCom '15)*. Association for Computing Machinery, New York, NY, USA, 117–129. <https://>

- [//doi.org/10.1145/2789168.2790113](https://doi.org/10.1145/2789168.2790113)
- [50] [wikimedia.org](https://commons.wikimedia.org/wiki/File:Bridgestone_tire_cross_section.png). 2017. Bridgestone tire cross section. Retrieved April 30, 2020 from https://commons.wikimedia.org/wiki/File:Bridgestone_tire_cross_section.png
- [51] [wikipedia.org](https://en.wikipedia.org/wiki/Tire). 2020. Tire. Retrieved April 30, 2020 from <https://en.wikipedia.org/wiki/Tire>
- [52] Jie Xiong and Kyle Jamieson. 2013. Arraytrack: A fine-grained indoor location system. In *Presented as part of the 10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13)*. 71–84.
- [53] Chenhan Xu, Zhengxiong Li, Hanbin Zhang, Aditya Singh Rathore, Huining Li, Chen Song, Kun Wang, and Wenyao Xu. 2019. WaveEar: Exploring a mmWave-based Noise-resistant Speech Sensing for Voice-User Interface. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 14–26.
- [54] Zhicheng Yang, Parth H. Pathak, Mo Sha, Tingting Zhu, Junai Gan, Pengfei Hu, and Prasant Mohapatra. 2019. On the Feasibility of Estimating Soluble Sugar Content Using Millimeter-Wave. In *Proceedings of the International Conference on Internet of Things Design and Implementation (Montreal, Quebec, Canada) (IoTDI '19)*. Association for Computing Machinery, New York, NY, USA, 13–24. <https://doi.org/10.1145/3302505.3310065>
- [55] Rudolf Zetik, Stephen Crabbe, Jozef Krajnak, Peter Peyerl, Jürgen Sachs, and Reiner Thomä. 2006. Detection and localization of persons behind obstacles using M-sequence through-the-wall radar. In *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense V*, Vol. 6201. International Society for Optics and Photonics, 62010I.