EasyChair Preprint

№ 870

The Fourier Tracing Approach for Modeling Automotive Radar Sensors

Martin Friedrich Holder, Clemens Linnhoff, Philipp Rosenberger and Hermann Winner

March 31, 2019

The Fourier Tracing Approach for Modeling Automotive Radar Sensors

Martin Holder, Clemens Linnhoff, Philipp Rosenberger, Hermann Winner

Institute of Automotive Engineering (FZD)
Technische Universität Darmstadt, 64287 Darmstadt, Germany
E-mail: holder@fzd.tu-darmstadt.de

Abstract: *The certification of autonomous driving by means of virtual testing methods is limited by the accuracy of simulated perception sensors, especially radar. This work presents a proof-of-concept implementation of Fourier tracing, a novel approach for synthesizing radar measurements of a virtual scene. Fourier tracing adapts ray tracing techniques from image synthesis to directly simulate a 3D periodogram of the range, range rate, and azimuth domains. Unlike previous approaches, Fourier tracing does not require computationally expensive FFT calculations and is specifically designed to capture the main characteristics of radar measurements, including interference and multi-path reflections. Our findings show that Fourier tracing captures key radar sensor performance parameters.*

1. Introduction

Large joint initiatives such as the PEGASUS and ENABLE-S3 projects actively pursue research in boosting the capabilities of virtual testing methods towards a level appropriate for the validation of self-driving cars. Sensor models are particularly crucial in the safety validation toolchain, due to their direct impact on the automated vehicle's perception system. No unrivaled standard has arisen for modeling perception sensors, and accurate modeling approaches are especially lacking for radar. In this work, we present *Fourier tracing*, a novel technique for modeling automotive radar sensors. Fourier tracing builds upon the mathematical foundations of ray tracing, synthesizing the received power of a radar sensor in a virtual 3D scene in the Fourier domain. radar sensor models utilizing ray tracing reported in prior work either do not discuss multiple targets [1], or do not provide comparisons to real-world measurements [2]. In contrast to previously reported approaches to radar sensor modeling, which attempt to directly produce object lists [3], Fourier tracing produces a periodogram and thus excels in capturing important radar characteristics while maintaining generalization across radar sensors.

The remainder of this paper is organized as follows: First, the key principle of Fourier tracing for synthesizing the *radar cube* from a virtual scene is introduced. This data structure contains a periodogram: the received power for each range, range rate, and azimuth angle bin. We then explain how key radar performance parameters are captured by the sensor model and outline an optimized ray tracing scheme for identifying reflection paths. We next detail how we implemented Fourier tracing for an automotive simulation framework and conclude by evaluating our modeling approach by comparing against measurements from static and dynamic scenarios.

2. Radar Sensor Modeling

Today's automotive radar sensors often use chirp sequence frequency modulated waveforms and MIMO antenna arrays. In a typical radar signal processing chain, the received signal at the antennas is digitalized and processed by a 3D FFT regime, where spectral peaks reveal range, range rate, and azimuth to targets present in the scene. More precisely, these Fourier spectra that record a power level for each frequency bin may be understood as unprocessed radar data. The achievable sensor performance is a compromise between hardware limitations and cost.

The most essential radar sensor characteristics are resolution and measurement range. These characteristics are principally determined by the chirp bandwidth Δf_{eff}, chirp sequence duration $T_{\text{M}} = I_{\text{v}} \cdot \Delta t_{\text{ch}}$ of I_{v} chirps, each comprised of I_{r} samples, and the separation distance Δy of I_{ϕ} evenly-spaced antennas. These parameters are summarized in table 1, where c and λ denote the speed of light and the wavelength, respectively. Note that aliasing can cause the detection of measurements with ranges exceeding the unambiguous interval bounds.

Table 1: Radar resolution parameters and unambiguous measurement intervals

	Range r	Range rate $\dot{\text{r}}$	Azimuth ϕ
Resolution cell	$\Delta r = \frac{\text{c}}{2\Delta f_{\text{eff}}}$	$\Delta \dot{r} = \frac{\lambda}{2T_{\text{M}}}$	$\Delta \phi = \arcsin\frac{\lambda}{I_{\phi}\Delta y}$
Unambiguous Interval	$\left(0, \frac{\text{c}I_{\text{r}}}{2\Delta f_{\text{eff}}}\right)$	$\left(\frac{-\lambda}{4T_{\text{M}}}, \frac{\lambda}{4T_{\text{M}}}\right)$	$\left(\arcsin\frac{-\lambda}{2\Delta y}, \arcsin\frac{\lambda}{2\Delta y}\right)$

Despite requiring several complex signal processing steps involving 3D FFT calculations and several estimation steps, the set of parameters describing core sensor performance is quite small. These parameters are often given in sensor data sheets, and the model inputs are readily available to the sensor model as part of the 3D virtual environment. For instance, the range and azimuth angle from the sensor to a target can be obtained from a ray tracer to a very high degree of precision.

2.1. Radar Modeling with Fourier Tracing

Ray tracing is a well-established technique for calculating propagation paths from a sender to a receiver and abstracts light propagation according to geometrical optics. This is a reasonable approximation for automotive radar sensors operating at frequencies near $77\,\text{GHz}$ by disregarding of diffraction. Fourier tracing brings these aspects together under the key insight that radar measurements essentially take place in the spectral domain. The core idea behind Fourier tracing is to realistically model the smearing of spectral peaks caused by limited measurement time. Fourier tracing makes use of window functions to capture the achievable measurement resolution and takes advantage of a ray tracer's ability to identify valid propagation paths between transmitters and receivers in a virtual scene and to measure ranges, range rates, and azimuth angles to the target with high precision. Preserving phase information enables consideration of interference effects.

The following section outlines the Fourier tracing workflow. First, a ray tracing algorithm determines the set of rays that return to the sensor. Each ray is initially understood as an infinitely sharp δ-peak in the respective periodogram, located at the particular frequency that is determined by range, range rate, and azimuth angle reported from the hit point of the received ray. This can be represented by the Kronecker-δ: $|\hat{P}_{\iota_\mathrm{h}}|\delta\left(\omega-\omega_\mathrm{ext}\right)$, where the ι_h-th hit point of the ray tracer reports a power $|\hat{P}_{\iota_\mathrm{h}}|$ and is represented as a single peak in the three periodograms, denoted as ω_ext. In a real FFT with limited measurement time, peaks usually do not appear at a single frequency, but affect neighboring bins through spectral leakage. This smearing of peaks is captured using a convolution of $\hat{U}_{\iota_\mathrm{h}}\delta\left(\omega-\omega_\mathrm{ext}\right)$ with a window function in the frequency domain, denoted $\mathcal{F}\left\{W_{(r,v,\phi)}\right\}$:

$$F_{\left(\iota_\mathrm{r},\iota_\mathrm{v},\iota_\phi\right)}\left(\iota_\mathrm{h}\right)=\hat{U}_{\iota_\mathrm{h}}\delta\left(\omega-\omega_\mathrm{ext}\right)*\mathcal{F}\left\{W_{(r,v,\phi)}\right\}. \tag{1}$$

Here, $F_{\left(\iota_\mathrm{r},\iota_\mathrm{v},\iota_\phi\right)}$ gives complex signal strength $\hat{U}_{\iota_\mathrm{h}}=\sqrt{|\hat{P}_{\iota_\mathrm{h}}|}e^{-j\varphi_{\iota_\mathrm{h}}}$ for each hit point $1\leq\iota_\mathrm{h}\leq I_\mathrm{h}$ at all range, range rate, and azimuth bin locations in the vicinity of the target. These signal values contain the power and the phase information. The phase is particularly important in order to account for wave interference effects and can be inferred from the length of the ι_h-th received ray, denoted $\ell_{\mathrm{ray},\iota_\mathrm{h}}$, with respect to λ.

$$\varphi\left(\ell_{\mathrm{ray},\iota_\mathrm{h}}\right)= \mod\frac{2\pi\ell_{\mathrm{ray},\iota_\mathrm{h}}}{\lambda},\quad\varphi\in[0,2\pi) \tag{2}$$

A change in power occurs whenever a ray's reflects off a surfaces. One must consider all reflections as a ray propagates through a scene. The power $|\hat{P}_{\iota_\mathrm{h}}|$ at the receiver is:

$$|\hat{P}_{\iota_\mathrm{h}}|=\frac{\mathcal{G}\cdot\prod_{l=1}^{I_\mathrm{h}}c_{\mathrm{R},l}}{r^4}\cdot\frac{\tilde{P}_\mathrm{Tx}}{I_\ell}. \tag{3}$$

Each reflection alters the power, which was set to unity when emitted from the sensor, by scaling it by the reflection coefficient $c_{\mathrm{R},l}$ of the reflecting surface. The Friis transmission formula [4] is applied using the MIMO antenna gain $\mathcal{G}$ and the range r to the target as the ray returns to the sensor. The fraction of the total radiated power $\tilde{P}_\mathrm{Tx}$ for each ray is normalized by the total number of launched rays I_ℓ.

The synthetic radar cube is populated by the absolute power values of the complex-valued $F_{\left(\iota_\mathrm{r},\iota_\mathrm{v},\iota_\phi\right)}$. The summation over all hit points realizes a phase-coherent superposition with respect to the ray's individual phase shift. Note that multipath propagation to any identified target in the scene is already considered by the ray tracing algorithm. The simulated power for given bin, denoted $\tilde{P}_{\left(\iota_\mathrm{r},\iota_\mathrm{v},\iota_\phi\right)}$, reads:

$$\tilde{P}_{\left(\iota_\mathrm{r},\iota_\mathrm{v},\iota_\phi\right)}=\left|\sum_{\iota_\mathrm{h}=1}^{I_\mathrm{h}}F_{\left(\iota_\mathrm{r},\iota_\mathrm{v},\iota_\phi\right)}\left(\iota_\mathrm{h}\right)\right|^2 \tag{4}$$

As can be seen from the derivations above, applying the Fourier tracing method on a 3D scene will only report the received power for each bin in each respective measurement dimension. Signal attenuation due to interference and range, as well as multipath propagation and antenna

gains are considered. Fourier tracing can therefore be seen as a method for synthesizing unprocessed radar measurements where no further signal processing, such as peak detection, spectral interpolation, or object tracking is initially considered. We argue that such processing steps are not part of the actual radar measurement process, which is modeled by Fourier tracing.

2.2. Radar Reflection Modeling with modified Ray Tracing

Using conventional image rendering methods to model the reflection behavior of electromagnetic mm-waves will fail as bidirectional reflectance distribution functions (BRDFs) are specific to a certain wavelength and therefore cannot be transfered from the visible light wavelengths of the image domain to mm-waves. For our proof-of-concept implementation of Fourier tracing, we assume the following radar-specific material reflection properties. We consider surfaces made from metal to be ideal reflectors and model edges on metallic surfaces as diffuse reflectors to account for their scattering behavior that is akin to antennas. Pavement reflectivity is modeled using the method proposed by Schneider [5]. All other material (e.g. tire rubber) is considered absorbent

Deploying ray tracing for modeling electromagnetic waves is challenging, as waves are a continuous front rather than comprised of individual discrete rays with zero thickness. This discretization causes surface meshes, which should result in a reflection back to the sensor, to be easily missed during ray tracing. To address this issue, for every hit point of a ray, we deploy an optimization algorithm to the area on the hit surface corresponding to the ray. This algorithm, inspired by the pattern search algorithm by Hooke and Jeeves, searches for further rays that reflect back to the sensor. The objective is to find new rays that minimize the distance between the sensor and potential reflection rays of examined points in the vicinity of the inbound ray. The original hit point of the inbound ray is locally shifted to optimize reflections back to the sensor. An edge detector is used to identify edges in the examined area, as metallic edges are modeled as diffuse reflectors.

3. Proof-of-Concept Implementation

We have integrated Fourier tracing in Vires Virtual Test Drive, an automotive simulation environment that provides an API for the Nvidia OptiX ray tracing engine. The only radar hardware parameters that are required to parametrize a sensor model with Fourier tracing are measurement ranges for range, range rate, and azimuth angle; bin resolutions and number of bins in each domain; the antenna pattern (incl. antenna and system gain); and the window functions for the FFTs. The simulation also requires the ray angle discretization and the sensor's detection range, which define the ray density and the total number of rays per time step. The model outputs periodogram values comprising spectral power for each range, range rate, and azimuth bin. This data can be arranged in a radar cube structure at every simulation time step. Sparse data structures are used to maintain efficiency. The only populated bins are those with significant spectral power, and their adjacent neighbors.

4. Verification Experiments and Results

Real-world data was collected with a series $77\,\mathrm{GHz}$ automotive radar sensor that outputs the received power in each bin. Individual reflection points are also produced through peak identification and spectral interpolation by a proprietary algorithm. Scenarios were re-created in a virtual environment using the recorded ground-truth positions of all agents. Two evaluation scenarios are considered. First, a static scenario is used to examine the range-azimuth plane of the Radar cube. Second, a dynamic scenario is used to examine the Doppler dimension of the cube as well as multi-path properties in the range-Doppler plane.

4.1. Static Scenario

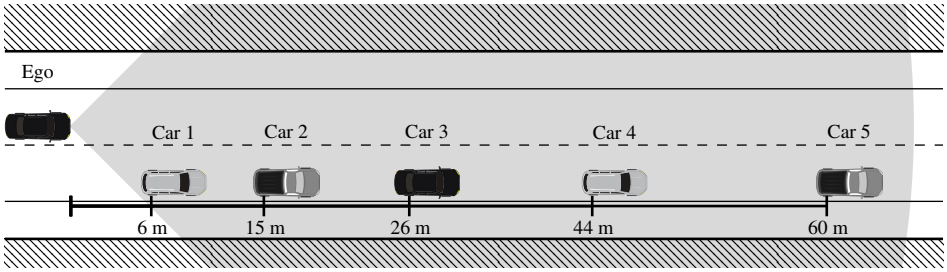

Figure 1: Arrangement of five stationary vehicles in front of the ego vehicle, which is equipped with a radar sensor. All cars are of different sizes and are located within the Radar's detection range.

This static scenario was constructed to study the reflective properties of extended targets at range. The static nature of the scenario allows one to neglect the Doppler dimension of the radar cube and only the maximum power value for each azimuth beam is considered. Five cars were placed along a straight road such that they were all visible at varying aspect angles to the radar sensor, see Fig. 1. We have chosen to connect our sensor model to a typical radar signal processing pipeline. This pipeline contains a detection threshold found using the OS-CFAR method [6] and the exact locations of the (range)-peaks are further refined by the QIFFT scheme [7] that interpolates sub-bin locations of peaks. Interpolated peaks are in the following referred to as *reflection points*. Both algorithms are applied to simulated and measured data and share the same parameters, such as false alarm probability, window size, order, and number of guard cells. This ensures a consistent false alarm rate throughout the comparison of the detected peaks. The scope of the experiment is to quantify differences occurring after applying a CFAR threshold and spectral interpolation on both measured and synthetic data.

Fig. 2 shows that reflections are successfully obtained from both the simulated and the measured power when one applies both an OS-CFAR threshold with a constant false alarm rate of 10^{-6} and spectral interpolation to the peaks in the periodogram. The reflections obtained from the simulated and measured power returns settle at similar ranges. Compared to the reflection points reported by the sensor, their cardinality is, however, smaller. This reduction is likely a result of the simplification to only one beam; the (unknown) signal processing in the sensor is able to derive multiple reflections for each car, which is not the case when processing the measured signal with

our pipeline. Nevertheless, our pipeline is still able to derive reflection points for every car and the power levels for reflections from simulated and measured power are similar to the reflections reported by the sensor. The noise levels of the simulation and the true measurement differ, and the CFAR threshold settles in a conservative manner. Areas in between the vehicles are noise and clutter regions, and are not taken into consideration.

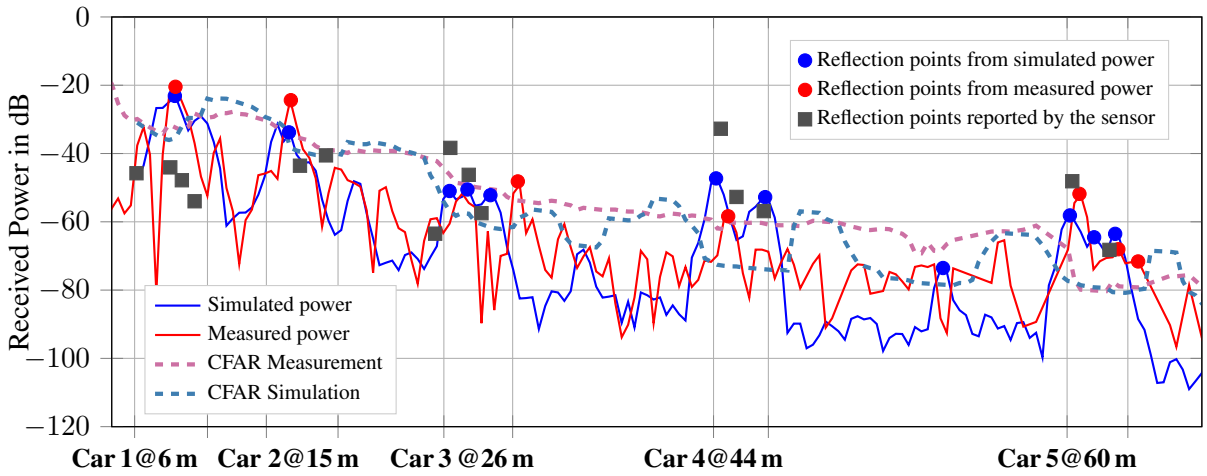

Figure 2: The received power (simulated and measured) along with a CFAR threshold. Radar reflection points reported from the radar and estimated from the power signals are also shown.

4.2. Dynamic Scenario

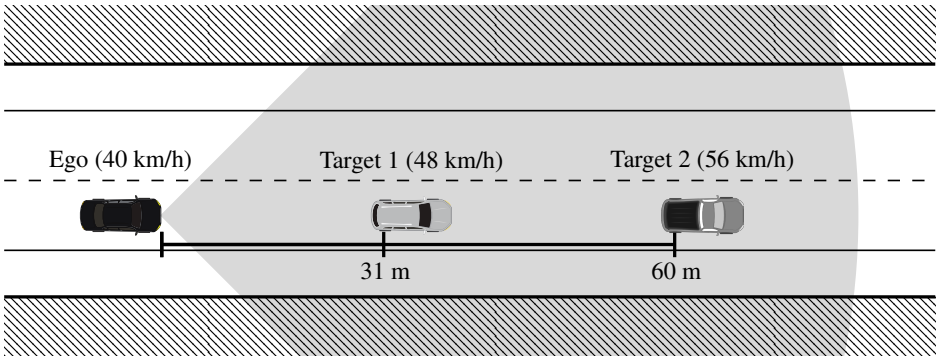

Figure 3: Overview of the dynamic scenario with two target cars. All cars move at different speeds. Target 2 is occluded by target 1.

Multi-path propagation and the correct mapping of occupied Doppler cells are studied in a dynamic scenario with two target cars, one of which is occluded. No additional post-processing steps are considered, as the scenario's main purpose is to demonstrate the model's ability to correctly collect Doppler and range information from the virtual environment under multi-path propagation. The target cars do not exhibit lateral displacement with respect to the ego car, allowing for a reduction of the 3D radar cube to a 2D range-Doppler map for the central azimuth beam. The dynamic scenario is depicted in figure 3. An instantaneous time-step is evaluated

where the distance from the ego car to the targets is $31\,\mathrm{m}$ and $60\,\mathrm{m}$ with relative speeds of $8\,\mathrm{km/h}$ and $16\,\mathrm{km/h}$, respectively. The measured and simulated range-Doppler maps are displayed in figure 4. It can be seen that the relative movement of both target cars is captured correctly in the simulation. However, the simulation results show a higher number of occupied bins due to two obvious reasons. First, the number of rays that return to the virtual sensor is too high and more cells were occupied in simulation. This issue can be addressed by tuning the material reflection properties. Second, the window function (Chebyshev) can be replaced by a sharper window from the Blackman-Harris or Hamming family.

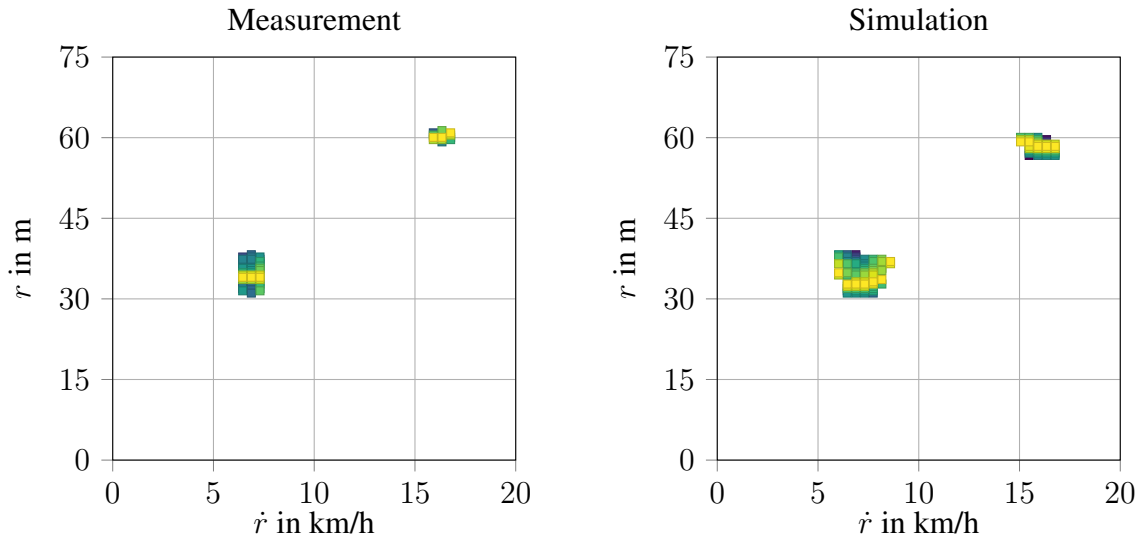

Figure 4: Measured (left) and simulated (right) Range-Doppler maps, cropped to the relevant portions of the dynamic scenario.

5. Conclusion

A proof-of-concept implementation of Fourier tracing, a novel method for modeling automotive radar sensors, is presented. Fourier tracing builds upon ray tracing and synthesizes the 3D radar cube from a virtual 3D scene, while bypassing the need for computationally expensive FFT calculations. An analysis of a static scenario found that reflection points where successfully interpolated from the power spectra. A dynamic scenario showed that Fourier tracing handles multi-path propagation for detecting occluded vehicles, and that cells in the range-Doppler domains where correctly populated. These findings stress the fact that Fourier tracing appears to be capable of rendering synthetic radar data in a manner comparable to real-world radar measurements. Future work will focus on incorporating more realistic material scattering properties, and improve sensor noise and clutter behavior. We will also evaluate whether relaxing the assumption of infinitely thin rays for capturing beam divergence is beneficial with regard to simulation accuracy.

6. Acknowledgment

This research was partly funded by the PEGASUS research initiative, promoted by the Federal Ministry for Economic Affairs and Energy (BMWi).

References

[1] N. Hirsenkorn, P. Subkowski, T. Hanke, A. Schaermann, A. Rauch, R. Rasshofer, and E. Biebl, "A ray launching approach for modeling an FMCW radar system," *2017 18th International Radar Symposium (IRS)*, 2017.

[2] B. Schick, G. Herz, R. Hettel, and H. Meinel, "Sophisticated Sensor Model Framework Providing Realistic Radar Sensor Behaviour in Virtual Environments," in *8. Fachtagung Fahrerassistenz*, 2017.

[3] S. Bernsteiner, Z. Magosi, D. Lindvai-Soos, and A. Eichberger, "Radar sensor model for the virtual development process," *ATZelektronik worldwide*, vol. 10, pp. 46–52, apr 2015.

[4] C. Levis, *Radiowave Propagation: Physics and Applications*. Hoboken, N.J: Wiley, 2010.

[5] R. Schneider, *Modellierung der Wellenausbreitung für ein bildgebendes Kfz-Radar*. Dissertation, Universität Karlsruhe, 1998.

[6] H. Rohling, "Radar CFAR Thresholding in Clutter and Multiple Target Situations," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 19, pp. 608–621, jul 2011.

[7] M. Abe and J. O. Smith III, "Design criteria for simple sinusoidal parameter estimation based on quadratic interpolation of FFT magnitude peaks," in *Audio Engineering Society Convention 117*, Audio Engineering Society, 2004.