Proceedings of the 33rd European Safety and Reliability Conference (ESREL 2023)
Edited by Mário P. Brito, Terje Aven, Piero Baraldi, Marko Čepin and Enrico Zio

doi: 10.3850/978-981-18-8071-1_P124-cd

Towards Modelling Sensor Failures in Automotive Driving Simulators

Rhea C. Rinaldo

Institute for Quality and Reliability Management (IQZ GmbH), Hamburg, Germany.
E-mail: rhea.rinaldo@iqz-wuppertal.de

Aaron Blickle and Nico Müller

Institute for Quality and Reliability Management (IQZ GmbH), Wuppertal, Germany.
E-mail: blickle@iqz-wuppertal.de, nico.mueller@iqz-wuppertal.de

Testing autonomous vehicles is a costly and tedious task, but essential to ensure their safe operation. In order to cover diverse scenarios and stage emergency situations like accidents safely, many manufacturers rely on automotive driving simulators. These simulators are virtually mirroring real driving situations by simulating an ego vehicle and its sensors to generate perceptual output for training or testing purposes of the autonomous driving software. Currently, a "functional paradigm" is followed by creating a realistic representation of the environment and running an ideal simulation. However, in reality sensors do not generate ideal outputs, as they are victim to various environmental effects like interfering signals, reflections, blockage and naturally ageing and wear. Furthermore, the environment can be flawed, e.g., traffic signs may be damaged or obscured. Consequently, we suspect that in the future a shift to a "failure paradigm" will become necessary where simulators include faults and failures of the infrastructure, the ego vehicle and its components. With that, a more realistic representation of driving in the actual world could be achieved, producing valuable data for training and testing of fault detection mechanisms.
With this paper we aim to make first steps towards the direction of a failure paradigm by the example of automotive radar. We study existing literature regarding the degradation of the radar function and propose four generalized failure models. We present the partial implementation of these in the open-source simulator CARLA and discuss experimental results as well as future contributions to this topic.

Keywords: Autonomous Driving, Automotive, Driving Simulators, Safety, Security, Reliability, Resilience, Failure Behavior, Testing

1. Introduction

Every leading car manufacturer is currently working on autonomous driving solutions. Testing these is essential to ensure the vehicle's safe operation in any conceivable situation, though it is a costly and tedious task. A plethora of different test scenarios are required to train and validate the autonomous function. This is extremely challenging, as the amount and diversity of required test scenarios can hardly be covered by test drives alone. In order to increase the diversity of test scenarios and stage emergency situations like accidents safely, many manufacturers rely on automotive driving simulation. Automotive driving simulators extensively model some 3D environment and simulate an ego vehicle and its sensors. The usual output comprises perception data as well as ground truth and meta information, providing for a manifold usage: Actual autonomous driving software can be directly connected to the simulator for testing, or the perceived data can be collected for training and analyses purposes. This AI-based software relies on perception data of various sensors, prominently cameras, radars and lidars. To refine the scope, this paper focuses on automotive radar. The application of radar in the automotive area dates back over 50 years ago, e.g., Grimes and Jones (1974) discusses its application in the context of automatic breaking. Since then, its quality has improved significantly and it became a reliable perception component for obstacles and their velocity. Owing to that fact, automotive radar is a prime sensor of currently developed automated and autonomous solutions.

However, like every electronic component it can age and is not free of misperceptions due to environmental disturbances. These are, for instance, interfering signals, reflections, obscuration through mud, water, foliage etc. Yet, current automotive simulators, already challenged by creating a realistic representation of real driving scenarios, are following a "functional paradigm" where the modeled environment and the sensed perception is most ideal. We can see that in reality sensors do not generate an ideal output and likewise the vehicle's environment can be flawed, e.g., traffic signs may be damaged or obscured. Because of that, we suspect that in the future a shift to a "failure paradigm" will become necessary. Thereby, the idea is that simulators include faults and failures of the infrastructure, the ego vehicle and its components, with the goal of achieving a more realistic representation of driving in the actual world. The generated data is valuable to the training and testing of fault detection mechanisms. Whereas efforts for creating more realistic radar models up to synthetically altering generated data to match the real-world exist (see e.g. Ngo (2023)), best to our knowledge, there is no model that includes and combines effects of functional safety, safety of the intended function and cybersecurity to date.

To change this, we extend the open-source simulator CARLA with four parameterizable failure models for its existing radar sensor, based on a literature study on the various factors leading to a radar degradation. The paper is closed by showing first results and discussing future extensions to this topic.

2. Automotive Radar

Radar is a historically grown technology that has evolved to a reliable perception sensor for delivering object information. In the big picture, an autonomous vehicle (AV) fuses the data of multiple sensors to gain a most realistic representation of the environment. Thereby, the radar sensor, if used, prominently contributes to the calculation of velocities and directions of detected objects. However, even though it is a comparably robust technology, it can be influenced and flawed by diverse factors. Subsequently we provide a brief overview on its functional specification, and on behalf of that, discuss how various effects can lead to a decreased signal or information quality.

2.1. *Functionality*

Due to the versatile requirements of the AV a plethora of different automotive radar sensors, distinguishable by used frequencies, covered ranges and price sections, exist. An abstract functional scheme can be defined as depicted in Figure 1. There, a radar sensor generally consists of a transmitter and a receiver, both equipped with one or more antennas, and a signal processing unit. The transmitter sends out radio waves in a range of 76 to 81 GHz (see Waldschmidt et al. (2021); Yeh et al. (2017)). Influenced by the mounting angle (see also Murad et al. (2012)), these waves form a cone directed to the target area that forms the sensor's field of view (FOV). A careful calibration of the FOV is essential for performing subsequent sensor data fusion and combining the perceived data to an accurate representation of the real world. Transmitted radio waves are reflected and adapted by any hit object surfaces in that cone such as pedestrians and vehicles, but also other reflecting surfaces including the road and traffic signs. However, not all of these waves hit reflective surfaces, and thus only a portion of them (that potentially mark hit objects) are reflected back to the receiving antenna(s). Received waves are collected and forwarded to the signal processing unit where information on the hit objects is determined. This information comprises most importantly (see also Wolff (2009))

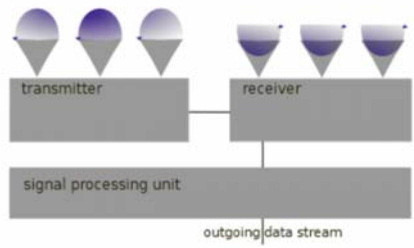

Fig. 1. Radar Structure

- the angle and direction to the object,
- the distance to the object based on the elapsed time between sending and re-

ceiving,

- and the velocity relative to the sensor based on the Doppler-Effect[a].

Further auxiliary information such as the radar cross section (RCS), describing how detectable an object is by radar based on its reflection properties, can be determined. With that, RCS provides valuable information for object classification (see also Hasirlioglu et al. (2016)). High resolution radar sensors allow for the determination of object shapes or even images. Consequently, radar sensors provide viable information for applications like automated cruise control, collision warning and mitigation systems, as well as the object detection, tracking and trajectory planning of the AV.

2.2. *Failures, Faults and Incidents*

There are various factors increasing the risk for failures of automotive radars ranging from minor errors in the object detection to the entire loss of the sensor or its provided information. We can categorize these classically as (i) hazards to the sensor's functional safety, leading to failures of the sensor module itself, (ii) hazards to the safety of its intended function, leading to functional insufficiencies where the sensor module is operating flawlessly, yet due to misuse or disturbances the output signal deviates from the expected and (iii) cybersecurity incidents, where the sensor's function is actively manipulated by an attacker.

The first category concerns any kinds of failures that are related to the sensor's hard- and software, ranging from production and design errors to natural effects of ageing and wear. As any electronic part, radar sensors can potentially degrade over their deployment time due to the ageing phenomena on transistor level. This can lead to manifold hardware faults (for more information see Halak (2019)), but also more commonly understood problems such as loose contacts, which may lead to data loss by disturbances of the existing data connections. The ageing effect is accelerated by environmental conditions such as extreme temperatures and temperature changes induced by the predominant climate and weather condition of the vehicle's operating area, or dispensing heat of the motor engine. Further, sensors are naturally exposed to nonlinear movements like vibrations, caused by the vehicle's engine, dynamic driving (braking and accelerating), as well as different road surfaces or due to bad sensor integration (see Hau et al. (2017, 2020)).

Vibrations can affect the radar in several ways. In the most straightforward case, we can observe a direct influence on the signal quality. Hau et al. (2017) present a radar signal model and test the effects of vibration on it. Thereby, the authors were able to show that the disturbance of the radar signal's phase information leads to a shift in the Doppler spectrum. Recalling, the Doppler-Effect is used to measure the velocity of hit objects by analyzing the frequency shift. Consequently, resulting velocity measurements can be flawed, and as stated by the authors, lead to a faulty distinction between stationary and moving targets. In this case, the failure lies somewhere in between category (i) and (ii), because we can assume that an automotive radar should be capable of dealing with expected vibrations, e.g., coming from the engine. Yet, heavy vibrations due to unexpected road surfaces, perhaps caused by driving offroads, can be considered beyond the scope of the intended usage and thus are not a failure of the sensor functionality. Besides that, vibrations can enhance the ageing phenomena. Commonly this phenomena is dealt with by predicting the component's degradation for the envisaged deployment time in regard of its functional requirements and the assumed environmental conditions. Based on that a maintenance plan is derived. For example, Tinga et al. (2017) model the ageing effect of a radar system's printed circuit board (PCB) due to *thermal fatigue* and *mechanical fatigue* as a consequence of vibrations for the purpose of predictive maintenance. This demonstrates that vibrations and temperature changes can accelerate the degradation of a radar system. Though, it must be mentioned that the authors regard the radar system

[a]The Doppler-Effect is the apparent change of a wave's frequency in relation to an observer that is moving relative to its source. Radar technology makes use of this effect by analyzing the Doppler frequency shift of the received signal and measuring the velocity of the moving object.

of a navy ship, which is much different from the radar sensor deployed in vehicles and further, the deployment time as well as the environmental conditions differ heavily between the automotive and the maritime area. This may be a reason why the effects of ageing in automotive radar is only superficially discussed in the literature.

Additionally to that, the sensor's mounting can be affected by vibration. A flawed mounting can potentially lead to shifts of the sensor's position and angle. As mentioned earlier, the position and angle define its conic FOV and thus are crucial for fulfilling the intended task. Depending on the degree, such a shift can remain unnoticed, but also lead to an entire loss of information. The latter is the case if the sensor is, e.g., pointing to the sky, the ground, or the FOV being blocked by the vehicle's covering. In addition to the natural ageing effect, this phenomena could also be caused by minor collisions with other vehicles or objects, for instance, a minor parking accident that does not seem to afford the vehicle to go to the repair shop. We can categorize these failures as (ii), because the functionality of the sensor itself remains unchanged, but the shifted FOV leads to unexpected or ineffective signals, qualifying as the loss of sensor. Additionally to the blockage of the sensor by an unwanted shift in position, blockage can also be caused by obscuration through mud, foliage, but also water and ice as stated in more detail in Murad et al. (2012). This kind of blockage can lead to the identification of very close detection points that may lead to an interpretation of very close objects and thus (partially) prevent the correct detection and interpretation of the surroundings (see also Fetterman and Carlsen (2016)).

With the radar's basic function being dependent on the reflection of signals, the reflection properties of surrounding objects make a critical factor. For instance, some object surfaces are more absorbing than others and thus only lead to a weaker detection. We do not consider this effect as a failure, because it is simply a result of the sensor's base mechanism. However, a commonly discussed problem based on reflections is the signal's backscatter on rainy weather. Raindrops, depending on their size and distribution within the antenna beam, reflect and absorb sent out radar signals (see Hasirlioglu et al. (2016)) potentially leading to a degraded perception. To test the influence of these effects, Hasirlioglu et al. (2016) measure the radar cross section (RCS) under different levels of rain. They come to the conclusion that rain can decrease the measured RCS by up to 67%. For object classification methods that rely on the reflection characteristics of a target this can lead to false identifications and potentially false decision making. Another reflection related problem is that the reflected signal takes an indirect path back to receiver, where it first bounces to a different target (a highly reflective surface like a traffic sign) creating a multi-path reflection (see Chamseddine et al. (2021)). This indirect reflection leads to the detection at a wrong position with different range and/or azimuth, often creating a so-called ghost target; a mirror-like reflection of the actual target at a different position.

Another vividly discussed phenomena is interference. With the increasing usage of automotive radar as a source for autonomous functions and automated safety functions (see Waldschmidt et al. (2021)), it is anticipated that vehicle-to-vehicle radar interference will become a significant problem (according to, e.g., Al-Hourani et al. (2017); Grimes and Jones (1974)). Thereby the victim radar is receiving another radar's signals. The consequences are noise polluted radar frequencies, and similar to reflections, the identification of ghost-targets is possible. Current research is focusing on mitigation techniques to decrease the consequence of this effect by, among others, various signal processing methodologies. While this effect classifies as category (ii), it is conceivable that interference are caused intentionally, classifying as category (iii). Regarding cybersecurity attacks on radar, Yeh et al. (2017) separate between jamming by saturating the receiver with noise, spoofing as the replication and retransmission of valid signals to provide false information and interference as the modification or disruption of the radar signal due to unwanted signals. In this way, an attacker could provoke the identification of ghost-targets, the denial of service due to jamming, or even the spreading of false velocity and distance information through sophisticated spoofing attacks. While these attacks are possible in theory, it must be mentioned that several difficulties remain such as the vehicle being a moving target that would need to be followed and the attacker requiring to sent on exactly the same frequencies.

3. Radar Failure Models

In the previous section we presented effects that can degrade or avert the radar perception. In reality their appearance is very unique to the situation. For example, the kinds of vibration working on a radar depend on the vehicle, the sensor itself and the road properties. Moreover, these effects change whenever the vehicle is braking, accelerating or the environment varies. Additionally to that, we could also see that some effects despite having a different nature actually share common results like interfering signals of other sensors or spoofing and intentional interference. Consequently, in this section we generalize the discussed effects by proposing four distinct failure models:

(1) Data transfer error
(2) Shifted FOV
(3) Signal disturbance
(4) Blockage

Further on, we are describing these failure models in more detail and discuss partially how their behavior is parametrized to achieve a more realistic staging of the effect. In the current state of our research we neglect considerations that are specific to the radar's environment. For that reason, we do not provide a model that covers reflection based failures, as it would afford the identification of potentially reflective surfaces in the simulation and thus presents a more sophisticated modeling and manipulation. In accordance with that, we are treating environmental aspects like road surfaces and the weather only superficially in the current state of the subsequently described models.

3.1. *Data Transfer Error*

The data transfer error characterizes by the (progressing) loss of single detection points up to the entire information produced by the sensor. Thereby, we aim to cover failures based on loose contacts, e.g., induced by the ageing phenomena, and partial jamming. Due to the progressing nature of loose contacts, as well as the timely behavior of jamming, we decided to make this model parametrizable by:

- Start time of the error
- Duration of an error
- Interval the error reappears
- Progression rate

Because we cannot assume that a failure is already active on start of the simulation, we make the first appearance of the failure a user specific choice in terms of a starting time. Further, especially regarding jamming, it is conceivable that the failure only persists for a given interval and then vanishes, possibly reappearing again. As loose contacts can progress over time, we also consider a progression rate to what the failure increases in its effect. By this we can model errors which occur for a short time period but with a high frequency, as well as errors that occur less often but for a longer time period.

3.2. *Shifted FOV*

This model represents shifts of the sensor's FOV due to collisions or ageing enhanced vibrations that work on the sensor's mounting. To cover these effects, the model is made adjustable by the following parameters:

- Shift of the position
- Shift of the rotation
- Start time / event
- Progression rate

The shift of the position (in meters) and rotation (in Euler angle) is crucial to define the direction and angle the radar can move to in case of a collision or a degradation of its mounting. These can be, among others, given by the properties of the mounting or simply for the desired simulation effect (e.g., considering a specific collision shall be staged). For the degradation of the mounting a start time has to be defined. This can be simply the start of the simulation, however, sometimes we want the failure to start after several driving hours have been performed. Once this start time is reached, the sensor starts to shift its rotation and/or position by the provided parameter. Optionally in that scenario, a rate describing the occurrence of the next shift can be determined, used in simulations where we want to express a progression of the failure, i.e., simulating a steadily progressing shift. For the collision based shift we require to define a triggering event rather than a starting time, such as *the first collision the ego vehicle experiences*.

3.3. *Signal Disturbance*

With this model we aim to generalize data falsification on signal basis through interference, vibration and possibly spoofing. Thereby this model covers two failure kinds: (1) the generation of detection points of non-existent objects, subsequently referred to as ghost detection points, and (2) partial velocity and distance falsification of valid detection points. The first failure kind abstracts the effect of interfering signals and maliciously placed signals and the second kind represents maliciously replayed signals as well as signal disturbances through vibration. In order to make a single model capable of covering all these effects, we require to make it adjustable in at least the following parameters:

- Start and end time of the error
- Ghost detection points
- Distance falsification
- Velocity falsification

Similar to previous models, the occurrence time of the failure can be defined. This is useful for various reasons. In reality the effect of vibration may always be present, though the effect of interfering signals and attacks is bound to other vehicles and potential attackers being within the range of the victim radar. Since we are currently not modelling the environment in this extent, we make an occurrence time definable to reflect that the effect may not persist throughout the entire simulation. Which kinds of failures are active is a user specific choice. The ghost detection points are modeled as a cluster or group of detection points that occur randomly within the FOV the sensor for an arbitrarily long time interval between milliseconds up to a few seconds. While distance and velocity falsifications are not required to be active at the same time, we model the failure in the same manner: A group of valid detection points (randomly chosen), that are assumably reflecting the perception of an actual object, are falsified. Similar to the ghost detection points, the effect lasts up to a few seconds.

3.4. *Blockage*

A heavy shift of the sensor's FOV or its obscuration through mud, foliage, water and ice can lead to a blockage of the sensor. The results are the perception of very close detection points without any relative velocity to the radar, because the object leading to the blockage and the radar are moving at the same speed. With this failure model we aim to represent blockage that is configurable by the following parameters to cover the properties of different sensor types and blockage kinds:

- Start time / event
- Blockage degree

In order to reflect that the blockage appears after a while of driving, e.g., by accumulating mud or snow, a start time for the effect must be determined. The effect on the signal depends on the kind of blockage. For instance, dust may only lead to partial blockage, while mud or a shift in the sensor's mounting may lead to entire blockage. Given that, it is essential to make the degree of the blockage (in percent) specifiable. In case the blockage is a result of a FOV shift due to a collision, instead of the start time a triggering event can be defined, similarly to the failure model concerning the FOV shift. With that we can create a dynamic transition between the shift in the FOV and sensor's blockage. Since blockage can affect the entire sensor or only certain, particularly exposed areas, a conceivable refinement would be to make a blockage area definable. However, this is highly depending on the radar's properties and thus affords further research to be considered accurately.

4. Application in CARLA

In the following we show first efforts regarding the implementation of the presented failure models, focusing on the data transfer error (1) and the shifted FOV (2). Therefore, we make use of the CARLA[b] open-source automotive driving simulator which is based on the Unreal Engine 4 and provides several flexible ways of manipulating the simulation and the data output. As pictured in Figure 2, CARLA is separated by a server-side that simulates the world including traffic participants, the ego vehicle and its sensors, and a client-side that manipulates the simulated objects via dedicated commands. Most of the relevant

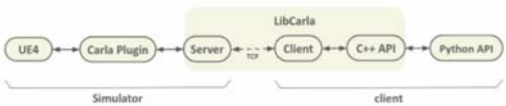

Fig. 2. CARLA Pipeline [d]

specifications, e.g., vehicle type and sensor setup, can already be made on the client-side. However, for manipulating CARLA internals an adjustment of the server's C++ base becomes necessary. While it is possible to extend the server by self-written models, typical automotive sensors are freely available. Based on the instantiated sensors, the server delivers sensor outputs and world information to the client. We base the implementation of our failure models on the native CARLA radar model. Since this model makes use of ray-casting rather than acting like real physical representation of a radar, we are required to implement signal based errors as a manipulation of the information that is sent to the client. Given that, the data transfer error is implemented by systematically leaving out transmissions of actual radar detection points to the client in the user defined duration. By making use of the start time, duration and reappearance interval parameters, we can simulate a steadily appearing data loss, for instance, every 5 seconds the transmission is left out for 100 milliseconds. In the current implementation, the progression rate is simply a defined increase of that duration.

The shift of the FOV is implemented in accordance with the presented model, though a progression rate is currently not considered. The user must define the degree of the shift in its position and/or rotation. Then the shift occurs by the given start time or is triggered by the first collision event. On that trigger, the sensor specification is dynamically rewritten by the adjustment of its position and/or rotation. For the collision based failure we listen to CARLA's internal collision sensor and request a position/rotation update on this event.

Figure 3 shows a simulation scene with a shifted FOV. In the top picture we can see the original radar placement, centered on our ego vehicle indicated by centered detection points (marked white). In the bottom picture we can see that the FOV is shifted to the right, as the detection points cover more of the building on the right, however, cease to cover the buildings on the left.

Fig. 3. Shifted FOV

[b] https //carla org/ accessed 2023-03-27

[d] https //carla readthedocs io/en/0 9 7/dev/how_to_add_a_new_sensor/ accessed 2023-03-30

5. Conclusion

The present paper provides a literature-based overview on diverse radar signal degradations. An abstraction of these effects to four general failure models was performed, covering failures of the categories of functional safety, safety of the intended function and cybersecurity. First efforts and results of an implementation in an automotive simulator have been presented that give an impression of further possibilities and data usage opportunities. In the future we want to implement the missing failure models and perform some refinements. We aim to include considerations of the environment, e.g., the sensor blockage being triggered by driving over muddy roads, snowy weather etc., or the loose contact progressing in relation to driving over rough road surfaces. At the same instance, we want to add more detail to the simulated effects and enable a smooth transition between instantiated failure models. As an example, instead of the blockage simply appearing at a defined time or event, we envisage to model a slowly progressing and increasing obscuration that eventually leads to a total blockage. Furthermore, we want to review degradation effects of other automotive sensors, like lidar and camera, and derive failure models similarly.

Acknowledgement

The present work is partially funded by the German BMWE joint project GAIA-X 4PLC-AAD.

References

Al-Hourani, A., R. J. Evans, S. Kandeepan, B. Moran, and H. Eltom (2017). Stochastic geometry methods for modeling automotive radar interference. *IEEE Transactions on Intelligent Transportation Systems 19*(2), 333–344.

Chamseddine, M., J. Rambach, D. Stricker, and O. Wasenmuller (2021). Ghost target detection in 3d radar data using point cloud based deep neural network. In *25th International Conference on Pattern Recognition (ICPR)*, pp. 10398–10403.

Fetterman, M. R. and A. Carlsen (2016). Blockage detection algorithm for automotive radars. In *International Conference on Microwaves for Intelligent Mobility (ICMIM)*, pp. 1–4.

Grimes, D. and T. Jones (1974). Automotive radar: A brief review. *Proceedings of the IEEE 62*(6), 804–822.

Halak, B. (2019). *Ageing of integrated circuits: causes, effects and mitigation techniques*. Springer Nature Switzerland AG 2020.

Hasirlioglu, S., A. Kamann, I. Doric, and T. Brandmeier (2016). Test methodology for rain influence on automotive surround sensors. In *19th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2242–2247. IEEE.

Hau, F., F. Baumgärtner, and M. Vossiek (2020). The degradation of automotive radar sensor signals caused by vehicle vibrations and other nonlinear movements. *Sensors 20*(21), 6195.

Hau, F., F. Baumgärtner, and M. Vossiek (2017). Influence of vibrations on the signals of automotive integrated radar sensors. In *International Conference on Microwaves for Intelligent Mobility (ICMIM)*, pp. 159–162.

Murad, M., J. Nickolaou, G. Raz, J. S. Colburn, and K. Geary (2012). Next generation short range radar (srr) for automotive applications. In *2012 IEEE Radar Conference*, pp. 0214–0219.

Ngo, A. (2023). *A Methodology for Validation of a Radar Simulation for Virtual Testing of Autonomous Driving*. Ph. D. thesis, Universität Stuttgart.

Tinga, T., W. Tiddens, F. Amoiralis, and M. Politis (2017). Predictive maintenance of maritime systems: models and challenges. In *European Safety and Reliability Conference (ESREL)*, Volume 40.

Waldschmidt, C., J. Hasch, and W. Menzel (2021). Automotive radar — from first efforts to future systems. *IEEE Journal of Microwaves 1*(1), 135–148.

Wolff, C. (2009). Book 1 – radar basics. `https://www.radartutorial.eu/druck/Book1.pdf`. accessed: 2023-03-23.

Yeh, E., J. Choi, N. Prelcic, C. Bhat, and R. W. Heath Jr (2017, 05). Security in automotive radar and vehicular networks. *Microwave Journal 60*, 148–164.