Fault-tolerant Radar Signal Processing using Selective Observation Windows and Peak Detection

Michael Beyer, Andre Guntoro
Bosch Corporate Research
Robert Bosch GmbH
Renningen, Germany
{michael.beyer2, andre.guntoro}@de.bosch.com

Holger Blume
Institute of Microelectronic Systems
Leibniz University Hannover
Hannover, Germany
blume@ims.uni-hannover.de

***Abstract*—Soft errors, such as bit flips, pose a serious threat to the functional safety of systems. Thus, ensuring the correct operation even in case of errors is particularly relevant for safety-critical applications. In this paper, we present a novel error detection and mitigation method for parallel FFTs in radar signal processing. We systematically define small observation windows in the 2D spectrum to detect peaks caused by soft errors. This enables protecting FFTs with several orders of magnitude lower computational overhead compared to related work. We conduct fault injection experiments to validate our method. Our experiments show that targets can be reliably detected even at higher error rates where more than 500 bit flips are present.**

***Index Terms*—fault tolerance, soft error, FFT, radar, signal processing, automotive**

I. Introduction

With increasing demand for computational performance and power efficiency, transistor structure sizes have been continuously shrinking in the past decades. Consequently, integrated circuits are becoming more susceptible to transient faults [1]. Such faults are caused by fluctuating temperature or voltage, electromagnetic interference, or highly energetic particles, and they can manifest as bit flips when sensitive regions in an integrated circuit are affected [2]. If such errors remain undetected, they can lead to silent data corruption, which is particularly concerning for safety-critical applications. For autonomous driving, neural networks (NN) are used to reason on the car's surroundings based on various sensor data such as radar, lidar, and video. The susceptibility of NNs to soft errors, i.e., non-permanent and non-destructive errors, and methods for improving their robustness has been investigated previously [3], [4]. However, besides ensuring the correct computation of the NN itself, the integrity of its input data is also highly relevant. Previous work investigated protecting fast Fourier transforms (FFTs), which play a key role in radar signal processing. Classic approaches like triple modular redundancy are not desirable in this case because of the high overhead they entail. An attractive solution to reduce the computational overhead is leveraging algorithmic properties to introduce checksums that cover the computations [5], [6]. During radar signal processing, several FFTs are performed in parallel. Therefore, instead of protecting each FFT separately, multiple FFTs can be protected simultaneously to further reduce the overhead [7]–[9]. However, the aforementioned methods focus on single errors and are not capable of protecting the computations in case multiple errors occur. In this paper, we investigate the protection of parallel FFTs in radar signal processing. Specifically, we combine the FFT's behavior when affected by faults with characteristics of the radar signal processing chain to achieve efficient error detection. To summarize, our contributions are:

- We investigate the impact of soft errors on the radar signal processing chain. Specifically, we investigate the error propagation of 2D-FFTs and compare the robustness of floating-point to fixed-point processing.
- We present and evaluate a new error detection and mitigation method, which requires less computational overhead compared to related methods (see Section IV-D).

The remainder of this work is structured as follows. First, we provide background information on radar signal processing and present the effect of soft errors in Fourier transforms. In Section III we introduce our error detection method. Afterward, the experimental setup is outlined, and the results are presented and discussed. Section V concludes our paper.

II. Background

A. Radar Signal Processing

The basic working principle of a radar sensor is based on transmitting electromagnetic waves and extracting information about objects from the received echo signal. For frequency-modulated continuous wave (FMCW) radars, the transmitted signal consists of several sinusoidal chirps with increasing frequency over time (see Fig. 1). To obtain the distance and relative velocity of an object to the radar sensor, a two-dimensional FFT of the received signal is computed. The first FFT, also called range FFT, is computed for each chirp. Afterward, the extracted frequency components can be converted to a distance to the sensor. The second FFT or velocity FFT, is computed along a different dimension using the range FFT's output (see Fig. 1). In the resulting *range-Doppler* (rD) spectrum, distinct peaks can be observed, each associated with a relative velocity and distance to the sensor. For target detection, a threshold is used to distinguish targets from noise in the rD spectrum. A fixed threshold is not desirable as external interference characteristics are not known in advance

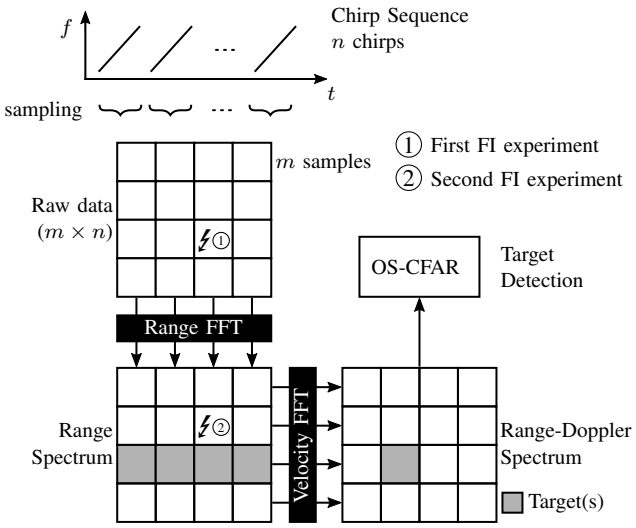

Fig. 1: Simplified overview of the radar signal processing chain (FI: fault injection, see Section IV-B).

and can change during system operation. In practice, constant false alarm rate (CFAR) detectors are employed to adjust the detection threshold online. Here, a sliding window is used to estimate the noise floor for each cell in the spectrum. One example for such a detector is the ordered statistic CFAR (OS-CFAR). OS-CFAR sorts the elements in the reference window and uses the k-th element as a noise estimate. For additional information on radar signal processing, we refer to [10].

B. Effect of Soft Errors in Fourier Transforms

Soft errors, such as bit flips, are single event effects caused by highly energetic particles. They are not permanent or destructive and can be reverted, e.g., by rewriting affected memory cells [2]. Computing an FFT over samples that are affected by bit flips, can be considered as a FFT over a signal with added random noise. The resulting spectrum exhibits an increased noise floor, depending on the magnitude and number of errors. In the worst case, the noise floor increases to such an extent that frequency components can no longer be extracted.

Considering the parallel FFTs that are performed during radar signal processing, errors in the input data create "error walls" in the two-dimensional spectrum. Depending on the dimension over which the FFT is computed, they manifest either vertically or horizontally across the spectrum as shown in Fig. 2. If such error walls are present after the range FFT, all subsequent velocity FFTs are affected, too. This distributes the error across the whole rD spectrum, resulting in an increased noise floor and thus fewer detected targets. Errors in the input of the velocity FFT that cause error walls in the rD spectrum, however, lead to false positive detections.

III. Method

To achieve efficient error detection, we combine the FFT's behavior when affected by faults with characteristics of the radar signal processing chain. As shown in Figure 2, errors in

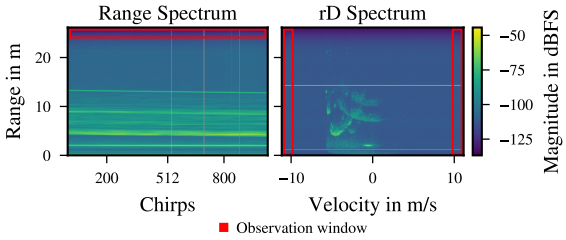

Fig. 2: Error walls in the range and range-Doppler spectrum and observation windows (red) for error detection.

the input data of a FFT can manifest as error walls in the two-dimensional spectrum. A straightforward solution for detecting such walls is examining the spectrum for vertical or horizontal lines after each FFT pass. If erroneous FFTs are detected, they can either be discarded or re-created in subsequent processing steps. Since errors in the input spread across a FFT's whole output, we can focus on smaller observation windows for error detection instead of analyzing the whole spectrum. For the range FFT, we select the top of the spectrum as the location for the observation window (see Figure 2 left). Its width is defined by the number of chirps in the chirp sequence, the height is a configurable parameter. As the loss in power of the received radar signal is proportional to r^{-4} [10, pp. 54–57], values in the last few bins located at the top of the spectrum are small. Therefore, this location is particularly advantageous for the observation window, as it allows detecting low error walls that would not be detected if the window is located at the bottom. For the velocity FFT we use two observation windows, one on each side of the rD spectrum (see Figure 2 right). Their height is defined by the number of samples for each ramp, the width is another configurable parameter. We combine detected errors of both windows and only report errors that appear in both. This ensures that possible targets located at the edges of the spectrum are not interpreted as errors. Nevertheless, there is a possibility of incorrectly classifying targets as errors, e.g., this corner case is triggered when two targets at the same distance to the sensor with same but opposing velocities are present. This effect can be mitigated by introducing an additional observation window. However, this entails an increased computational overhead.

The working principle of our method is as follows:

1) The squared magnitude for each element in the observation window is computed.
2) An error detection threshold is defined based on the estimated noise floor of the observation window. It is obtained by OS-CFAR processing.
3) The mean is computed along the height/width of the window depending on the FFT stage.
4) The obtained values are compared to the error detection threshold, resulting in a mask containing the location of any present error walls.

Using OS-CFAR for dynamic noise estimation enables parameter-less usage of our detection method. Furthermore, we assume that a dedicated unit for OS-CFAR is available in the

radar system, which is reused by our method. For final target detection with OS-CFAR, an accurate noise estimate increases detection performance by providing a better separation of targets and background noise. Our method focuses on detecting errors, which would otherwise affect the final target detection after OS-CFAR processing if not mitigated. Therefore, an accurate estimate of the noise floor is not necessary. A single noise estimate for the whole observation window is sufficient, which also reduces the required computational overhead. In practice, we suggest a similar false alarm rate and window size as it is used for the subsequent target detection with OS-CFAR. That way, possible undetected error walls after the velocity FFT will not result in false positive targets.

While our investigations focus on memory errors affecting the input data of a FFT, we analyze that our method is also capable of detecting computational errors to some degree. Considering the signal flow graph of a FFT, an error originating in a butterfly stage can be distributed over the entire output. Therefore, the detection mechanism may not be able to detect an error wall as it uses only a few bins of the output for detection. However, errors that occur during the computation of the range FFT and remain undetected can lead to error walls after the velocity FFT, which can be detected by our method (see results of the second experiment in Section IV-C). Errors during the computation of the velocity FFT may not be detected. This can be circumvented by adding additional observation windows and/or increasing their size such that more bins are included in the error detection.

Overall, our method enables efficient protection of parallel FFTs in the radar signal processing chain, since it operates only on a fraction of the spectrum to check whether errors occurred. Furthermore, apart from determining the error detection threshold, error wall detection and mitigation can be performed independently and is therefore highly parallelizable.

IV. Experiments

A. Measurement Setup

In our experiments we use data collected by an automotive radar sensor prototype similar to the one in [11]. The multiple-input multiple-output (MIMO) radar employs a frequency modulated signal and is mounted in the front of a test vehicle. It has two sending and 16 receiving antennas for a total of 32 virtual channels. The range and velocity FFTs are performed over 512 samples and 1024 chirps, respectively. After non-coherent integration of all virtual channels, the size of the rD spectrum is 256×1024. It is processed by an OS-CFAR detector with a false alarm rate of 10^{-6} for target detection.

B. Experimental Setup

We perform multiple fault injection (FI) experiments to evaluate our error detection mechanism. Furthermore, we use floating-point and fixed-point data representations and compare their robustness against soft errors when used in radar signal processing. For the fixed-point experiments, we use a word length of 32 bits to match that of the floating-point experiments. Although smaller word lengths are also used in practice, we explicitly use 32 bits to ensure a fair comparison between the two data representations, since bit flips are injected based on a bit error rate. Therefore, if one representation had fewer bits, fewer bit flips would be injected in total (more information below). For the fixed-point FFT, we use a radix-2 implementation with convergent rounding [12], where values are scaled after each stage.

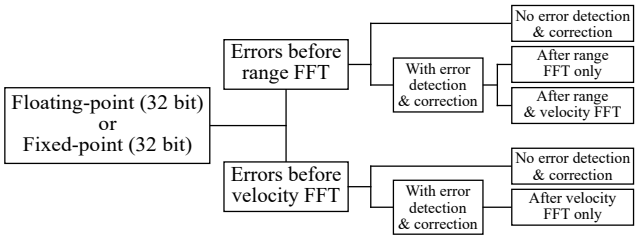

Fig. 3: Experimental setup.

In Fig. 3 an overview of the experimental setup is shown. For both data types, we conduct two sets of experiments where bit flips are injected into the input data of either the range or the velocity FFT (i.e., into the raw data or the range FFT's output. See Figure 1). Depending on where faults are injected, we test different configurations where either both, none, or just a single FFT pass of the processing chain is covered by our error detection method. In all cases, the spectrum of each virtual channel of our MIMO radar is protected separately, i.e., the error detection threshold and mask for detected errors are determined for each virtual channel. Based on empirical analysis, the reference and observation windows for error detection have a height/width of three bins and any detected error walls are removed by zeroing. All experiment configurations are investigated using five different radar measurements.

To simulate the effect of bit flips, we randomly generate an error mask that is applied to the input data of a FFT using the XOR operation. Each bit in the mask is set depending on a specified bit error rate (BER). In total, we investigate six different BERs corresponding to injecting just a single bit flip to more than $50\,000$ bit flips. For each BER, we perform 100 individual FI runs to ensure statistical confidence. In each run we pass the rD spectrum to the OS-CFAR detector for target detection. Finally, we compare the number of detected targets and their location in the spectrum (i.e., their distance and velocity) to the ground truth without faults. Our two performance indicators to evaluate our method are:

False negative rate (FNR). The proportion of cells in which no targets have been detected, but which are targets, out of all cells that actually contain targets. I.e., the proportion of missed targets to the total number of targets.

False positive rate (FPR). The proportion of cells in which targets have been detected, but which are not targets, out of all cells that actually contain no targets. I.e., the proportion of incorrectly detected targets out of all non-target cells.

C. Results

In this Section the results our FI experiments are presented. For all configurations listed in Fig. 3, we show the mean and

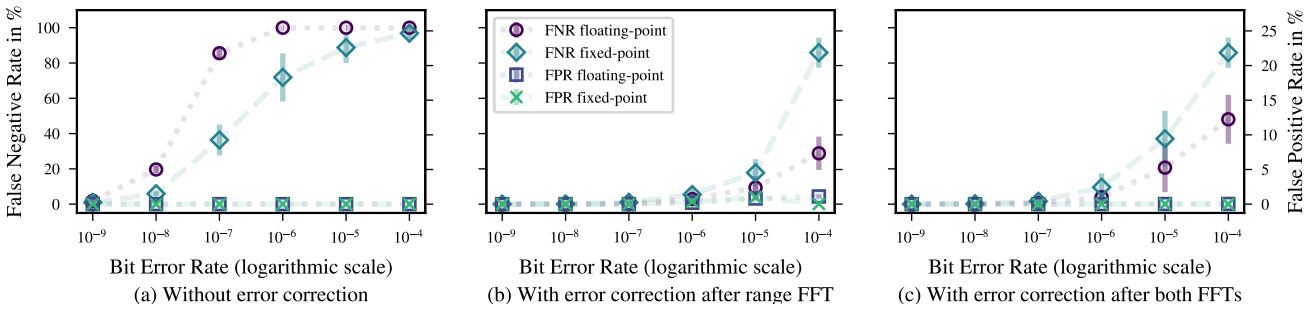

Fig. 4: Mean and std. dev. for FNR and FPR of the first set of experiments (errors before range FFT only). Connecting lines between markers are added for visual guidance.

standard deviation for FNR and FPR computed over all five radar measurements and the 100 runs for each BER.

Figure 4 shows the results of the first set of experiments where faults are injected into the raw input data. In case no error mitigation is used (Fig. 4a), less targets are detected with an increasing BER. The FPR remains zero for all tested BERs. This result is in line with the effects outlined in Section II-B. Faults in the raw input data lead to error walls in the range spectrum and since these error walls affect each individual velocity FFT, the error is spread across the whole rD spectrum, resulting in an increased noise floor. As a result, fewer targets are detected as the BER increases, since target responses are swallowed up by the noise floor. The FPR remains zero for all tested BERs since the noise floor rises uniformly throughout the rD spectrum. When using our error detection method after the range FFT only (Fig. 4b), targets can be reliably detected initially. However, as the BER increases, multiple effects that influence the error and target detection performance come into play: (i) At first, error walls are detected and removed reliably, resulting in low FNRs and FPRs. Removing error walls creates sidelobes that cause a wider target response in the rD spectrum. This influences the noise estimate during OS-CFAR processing, which can lead to masking of weaker targets and thus an increased FNR. Furthermore, they can be incorrectly interpreted as targets, increasing the FPR. (ii) With more and more error walls in the range spectrum, also the reference window used for determining the error detection threshold is affected. Consequently, the reference window's noise estimate is no longer correct which leads to a higher error detection threshold. As a result, fewer error walls are removed, which causes a higher noise floor in the rD spectrum and thus an increased FNR. Since fewer walls are removed, also fewer sidelobes are created. Therefore, the FPR decreases again for higher BERs. A FPR with a maximum of 1% at a BER of 10^{-5} may seem low, however, in this case it corresponds to a couple of thousand detected targets. By using our error detection method after the range and velocity FFT (Fig. 4c), false positive targets caused by sidelobes due to zeroing in the range spectrum can be mitigated. However, this entails a slight decrease in target detection performance at higher error rates. This is because removing sidelobes in the rD spectrum can also remove target responses, resulting in an increased FNR.

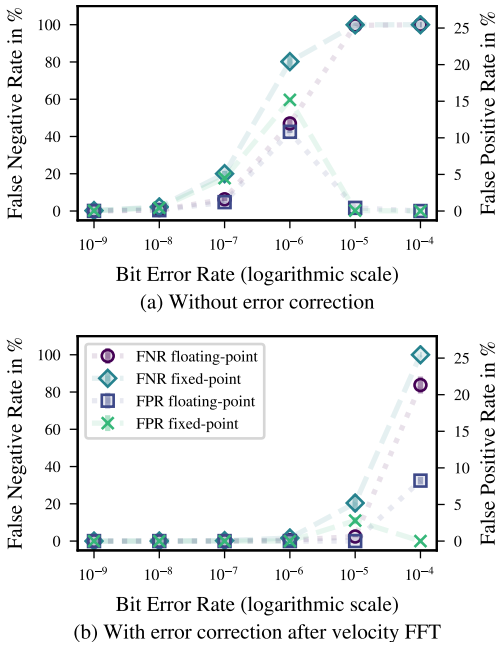

Fig. 5: Mean and std. dev. for FNR and FPR of the second set of experiments (errors before velocity FFT only).

Figure 5 shows the results for the second set of experiments where errors are injected into the velocity FFT's input data. Without error mitigation (Fig. 5a), error walls manifesting in the rD spectrum directly affect the target detection. When computing the noise estimate for the sliding window during OS-CFAR processing, error walls can increase the target detection threshold resulting in masking of targets with a weaker signal response. Consequently, with an increasing BER, more error walls cover the rD spectrum, which results in an increased FNR. At a certain point, the spectrum consists practically only of noise, so that fewer and fewer targets are detected until finally none are detected at all. Apart from affecting the detection of actual targets, error walls are additionally interpreted as targets, too. The FPR shown in Figure 5a initially increases with an increasing BER and reaches a peak at 10^{-6}, after which it decreases again. This effect is linked to how targets are detected with OS-CFAR. At first, the FPR increases with the increasing number of error walls that are present in the rD spectrum. However, similar to the FNR, at a certain point

the spectrum consists of just noise, leading to less detections overall. Figure 5b shows that up to an BER of 10^{-6}, both targets are reliably detected and the FPR remains low when our error detection method is used. In general, removing detected error walls through zeroing can also remove target responses, which should consequently increase the FNR. However, here this does not affect final target detection initially, since target responses can be preserved due to our MIMO radar setup with 32 virtual channels and its associated high data redundancy. Similar to previous cases, the error detection threshold is affected more and more at higher error rates, since more error walls are present in the spectrum. Therefore, with an increasing BER, less error walls are detected and removed, leading to an increasing FNR and FPR.

D. Discussion

Overall, errors before the range FFT have a more severe impact on target detection, since they spread throughout the whole rD spectrum. Errors before the velocity FFT on the other hand cause false positives. The robustness of the investigated data types w.r.t. reliable target detection is different in each experiment and whether error mitigation is used or not. In general, the behavior can be attributed to the value range and encoding for each data type and the signal processing stage in which errors occur. For floating-point, bit flips affecting the mantissa do not alter the value significantly, unlike bit flips in the exponent. A single error in the exponent before the range FFT can easily cause a high error wall that raises the whole rD spectrum's noise floor such that no targets are detected. The same error for the velocity FFT affects only part of the rD spectrum and is therefore less severe. This is different to fixed-point, where errors are limited to the same value range as the processed data. While in this case errors in the raw data can also raise the rD spectrum's noise floor, errors are not as drastic and increase the noise floor more gradually. This is also why the std. dev. is more variable for fixed-point compared to floating-point. When using our error detection method, floating-point is performing better at high BERs compared to fixed-point. The reason for this is that the magnitude of error walls for fixed-point is within the same value range as the processed data. This has a greater effect on the calculation of the noise estimate, and hence the error detection threshold, than errors walls that lie far outside the range of values in the processed data. Nevertheless, in practice fixed-point can be considered more robust, since the magnitude of erroneous values is limited, and smaller word lengths can be used. This significantly reduces the number of bits and thus the attack surface for highly energetic particles, since less data is stored. Overall, the results show that with our method, targets can be reliably detected for a BER up to 10^{-7} and 10^{-6} for the range and velocity FFT, respectively, which correspond to more than 100 and 500 bit flips in our setup. Although soft errors are rare and protecting FFTs with our method may therefore be more than required, our approach has less overhead than related work aimed at detecting single errors. In Fig. 6 the operations required for protecting the 2D FFT of our setup are compared to methods proposed in [9], which use error correction codes (ECC) and sum of squares (SOS) checks. As stated earlier, we assume that a dedicated unit for OS-CFAR is present in the system. Therefore, computing the noise estimate is not included in our overhead.

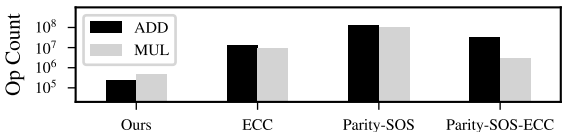

Fig. 6: Comparison of the computational overhead for protecting the 2D FFT in our setup with methods proposed in [9].

V. Conclusion

This paper presents an error detection and mitigation method for parallel FFTs in radar signal processing. Our method only needs to evaluate a fraction of the spectrum to determine if errors occurred, resulting in low computational overhead. Experiments show that target detection can be maintained even at high error rates where more than 500 bit flips are present. Furthermore, the method is highly parallelizable and therefore suitable for efficient hardware implementations.

Acknowledgments

This work was supported by the German federal ministry of education and research (BMBF), project ZuSE-KI-AVF under grant no. 16ME0062.

References

[1] R. Baumann, "Radiation-induced soft errors in advanced semiconductor technologies," *IEEE Trans. Device Mater. Rel.*, vol. 5, no. 3, pp. 305–316, Sep. 2005.

[2] V. Castano and I. Schagaev, *Resilient computer system design*. Cham: Springer International Publishing, 2015.

[3] E. Ozen and A. Orailoglu, "Sanity-Check: Boosting the Reliability of Safety-Critical Deep Neural Network Applications," in *2019 IEEE 28th Asian Test Symposium (ATS)*. IEEE, Dec. 2019, pp. 7–75.

[4] M. Beyer *et al.*, "Automated Hardening of Deep Neural Network Architectures," in *Proceedings of the ASME 2021 International Mechanical Engineering Congress and Exposition (IMECE) 2021*, Nov. 2021.

[5] K.-H. Huang and J. A. Abraham, "Algorithm-based fault tolerance for matrix operations," *IEEE Trans. Comput.*, vol. 100, no. 6, pp. 518–528, 1984.

[6] A. Jacobs, G. Cieslewski, C. Reardon, and A. D. George, "Multiparadigm Computing for Space-Based Synthetic Aperture Radar." in *ERSA*. Citeseer, 2008, pp. 146–152.

[7] P. Reviriego, S. Pontarelli, C. J. Bleakley, and J. A. Maestro, "Area efficient concurrent error detection and correction for parallel filters," *Electronics letters*, vol. 48, no. 20, pp. 1258–1260, 2012.

[8] P. Reviriego, C. J. Bleakley, and J. A. Maestro, "A novel concurrent error detection technique for the Fast Fourier Transform," in *IET Irish Signals and Systems Conference (ISSC 2012)*, 2012, pp. 1–5.

[9] Z. Gao *et al.*, "Fault Tolerant Parallel FFTs Using Error Correction Codes and Parseval Checks," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 24, no. 2, pp. 769–773, 2016.

[10] M. A. Richards, *Fundamentals of Radar Signal Processing*. US: McGraw-Hill Professional, 2005.

[11] F. Meinl, M. Stolz, M. Kunert, and H. Blume, "An experimental high performance radar system for highly automated driving," in *2017 IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*. IEEE, 2017, pp. 71–74.

[12] D. Menard, D. Novo, R. Rocher, F. Catthoor, and O. Sentieys, "Quantization mode opportunities in fixed-point system design," in *2010 18th European Signal Processing Conference*. IEEE, 2010, pp. 542–546.