

»GREEN ICT @ FMD« – COMPETENCE CENTER FOR ECOLOGICALLY SUSTAINABLE ICT

Resource-Efficient Radar Signal Processing on Microcontrollers

A Whitepaper by "HUB 1 –Sensor-Edge-Cloud"

Michel Sonntag (Fraunhofer FHR), Michael Gräf (Fraunhofer FHR), Sabine Gütgemann (Fraunhofer FHR),
Christian Krebs (Fraunhofer FHR)



The work presented is part of the »Green ICT @ FMD« project, your competence center for ecologically sustainable information and communication technology. The project is established by the Research Fab Microelectronics Germany and funded by the German Federal Ministry of Education and Research.

Kompetenzzentrum »Green ICT @ FMD«

c/o Forschungsfabrik Mikroelektronik Deutschland FMD
Anna-Louisa-Karsch-Str. 2
10178 Berlin, Germany

Main contact
Sabine Gütgemann, sabine.guetgemann@fhr.fraunhofer.de

www.greenict.de
www.forschungsfabrik-mikroelektronik.de

Date of publication
08.12.2025

1. Introduction

Resource-efficient information and communication technology (Green ICT) requires novel system designs that reduce energy and material expenditure over the entire life cycle without compromising functional performance. Radar systems pose a particular challenge: they typically rely on FFT-centered signal processing chains with high demands on compute and memory resources, which often suggests FPGAs or GPUs. The resulting power consumption is, however, a critical factor in mobile, distributed, or scaled sensor networks and also affects environmental footprints. Against this backdrop, this work addresses shifting essential processing steps to microcontroller (μ C) platforms and designing distributed sensor-edge architectures.

In the area of spectral analysis, it is known that the nominal frequency resolution of the FFT scales with block size and that windowing, leakage, and estimator variance limit practical accuracy [1]. Larger FFTs improve resolution but increase computational load and thus energy consumption. In contrast, high-resolution estimators based on peak interpolation (parabolic) allow sub-bin accuracy with smaller block sizes, provided SNR, windowing, and model assumptions are consistent [2]. In parallel, time-interleaved ADCs offer a way to increase the effective sampling rate; however, they require calibration of gain, offset, and timing skews, as well as low-jitter clock and trigger design to ensure phase coherence and SNR [5].

This work presents a radar-centered demonstrator that

1. provides a μ C-based backend for up to four phase-coherent channels (80 GHz),
2. can switch at runtime between independent (1–4 channels) and interleaved operation (1–2 channels, increased sampling rate),
3. integrates FFT-adjacent peak interpolation to reduce FFT block size while maintaining measurement quality, and
4. demonstrates distributed task sharing between sensor and edge.

The hardware is modular (frontend, filter/conversion, backend) and supports selective power gating to efficiently cover partial load and event-driven operation. Critical aspects such as phase-coherent acquisition across multiple ADCs/controllers, external block triggering, and clock/jitter management are addressed.

1.1 Problem Statement

How can central radar signal processing be migrated from FPGA/GPU platforms to μ C architectures such that (a) the overall system's energy consumption measurably decreases, (b) measurement quality (frequency/distance estimation) is maintained, and (c) the architecture is scalable for distributed sensor-edge scenarios? The following constraints are decisive:

- Limited compute/memory resources of the μ C platform compared to FPGA/GPU.
- Requirements for phase coherence across multi-channel and multi-controller structures (trigger/clock/jitter/skew).
- Flexibility in operating modes and bandwidth (independent vs. interleaved) with runtime-changeable ADC parameters.
- Life-cycle and environmental footprint perspective.

1.2 Scope and Delimitation

The focus is on FMCW short-range applications with requirements for distance estimation (target range, adjustable bandwidth, and update rate). Topics such as MIMO angle measurement, complex multi-target tracking algorithms, or safety-certified implementations are outside the immediate scope but are addressed in the outlook. The edge component serves visualization/feature aggregation. Cloud backend structures are not considered here.

2. System Architecture

This chapter describes the architecture of the radar-based demonstrator along the complete signal chain: from the frontend (80 GHz) through filter/conversion and the microcontroller backend in a master/slave configuration to the edge-cloud components (Xilinx Versal, Nvidia Orin). The focus is on resource-efficient, modular hardware, phase-coherent multi-channel acquisition, configurable ADC operating modes (independent/interleaved), and firmware that implements FFT-based signal processing with peak interpolation energy-efficiently on a μC platform. Design decisions target energy efficiency, measurement quality, and adaptability for different application scenarios.

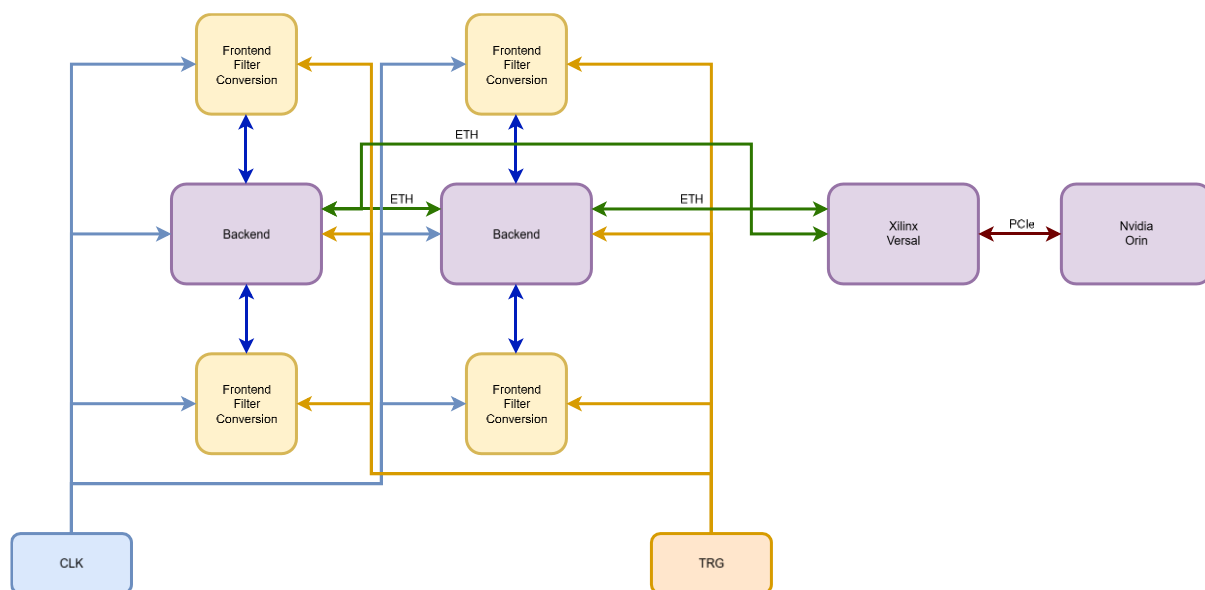


Figure 1 Simplified schematic of the FHR concept

2.1 Overview and Design Goals

The demonstrator aims to replace an FPGA-based reference solution with a μC backend while significantly lowering system power without compromising measurement quality. Core principles are:

- **Modularity:** Frontend, filter/conversion, backend, and adapters are implemented as standalone, pluggable modules. This simplifies switching modules on/off, isolated measurements of individual paths, and reuse in future systems.
- **Energy efficiency:** Selective power gating (including the slave μC and peripherals) and partial-load-optimized operating strategies (duty cycling, feature-first streaming) reduce average power consumption.
- **Phase coherence:** External, block-wise triggering and coherent clock distribution ensure identical phase across multi-channel and multi-controller operation [6].
- **Flexibility:** ADC modes are switchable at runtime (independent 1–4 channels; interleaved 1–2 channels with increased effective sampling rate). Resolution and sampling rate are adapted per scenario.
- **Edge integration:** Integration into a system with different components to intelligently distribute computational load.

2.2 Hardware Architecture

The backend consists of two identical microcontroller boards (e.g., STM32H735, Cortex-M7 with FPU and DSP instructions) operated in a master/slave topology. Both boards run the same firmware; role assignment is determined deterministically during the boot process (boot-strapping). The master can hard power off the slave controller and selected peripherals via power switches or place them in defined sleep states. This power gating is a central lever for partial and event-driven operation. The platform provides four ADC channels. In independent mode, 1–4 channels run in parallel at standard sampling rate. In interleaved mode, 1–2 channels achieve doubled effective sampling rate by two ADC cores sampling the same input signal alternately. Single-ended input is required for interleaved operation; differential-to-single-ended conversion is integrated on the filter board. The target configuration covers 16-bit resolution and worst-case sampling rates up to 7.2 MSps/channel; other rates/bit depths are runtime configurable.

To ensure phase coherence, both μC ADCs are externally and block-wise triggered (master GPIO). The clock path is coherently designed; lengths, impedances, and distribution elements are arranged so that channel and controller skews remain reproducibly small. Custom evaluation boards support measurement and optimization of trigger jitter and channel skews. A block trigger (instead of sample trigger) reduces interrupt load and increases robustness against jitter.

Table 1 System properties

Aspect	Status
<i>Radar principle</i>	FMCW
<i>Frontend frequency</i>	80 GHz
<i>Bandwidth</i>	25 GHz
<i>Chirp duration</i>	102,4 μs
<i>Channels</i>	1 – 4
<i>ADC modes</i>	Independent (1-4 channels), Interleaved (1-2 channels)
<i>ADC resolution</i>	10 – 16 Bit
<i>ADC sampling rate</i>	3,6 – 11 MSps
<i>Triggering</i>	Internal or external
<i>Power management</i>	Selective power-gating
<i>Edge integration</i>	Compute load distributable via Xilinx Versal or Nvidia Orin

3. Methods and Implementation

This chapter describes the concrete implementation of resource-efficient radar signal processing on a microcontroller platform. The focus is on (i) an FFT-adjacent frequency estimation with peak interpolation to reduce computational load, (ii) phase-coherent multi-channel acquisition across two controllers, (iii) flexible ADC operating modes (independent/interleaved) including runtime switching, and (iv) firmware and energy-management strategies to lower average power. Integration into edge components for visualization and application-near demos is also described.

3.1 Signal Processing Approach: FFT-Compatible Peak Interpolation

Radar sensors typically use FFT-based chains for distance and velocity estimation. Nominal frequency resolution increases with FFT length, driving compute and memory needs. To achieve equivalent measurement quality on μC structures with significantly reduced computational load, we use a two-stage approach:

- Preprocessing and windowing: Raw data are processed with proven windows (e.g., Hamming) to control leakage and reduce estimator variance [1]. Window selection remains configurable to reflect SNR and scenario requirements.
- Peak interpolation in the frequency domain: Instead of increasing FFT length, the strongest spectral component is estimated sub-bin precisely using parabolic interpolation (PI). Simulations with representative SNRs show nearly continuous frequency estimation with smaller FFT sizes (Figure 2). This lowers compute and memory load without compromising accuracy relevant for distance estimation [2, 3].

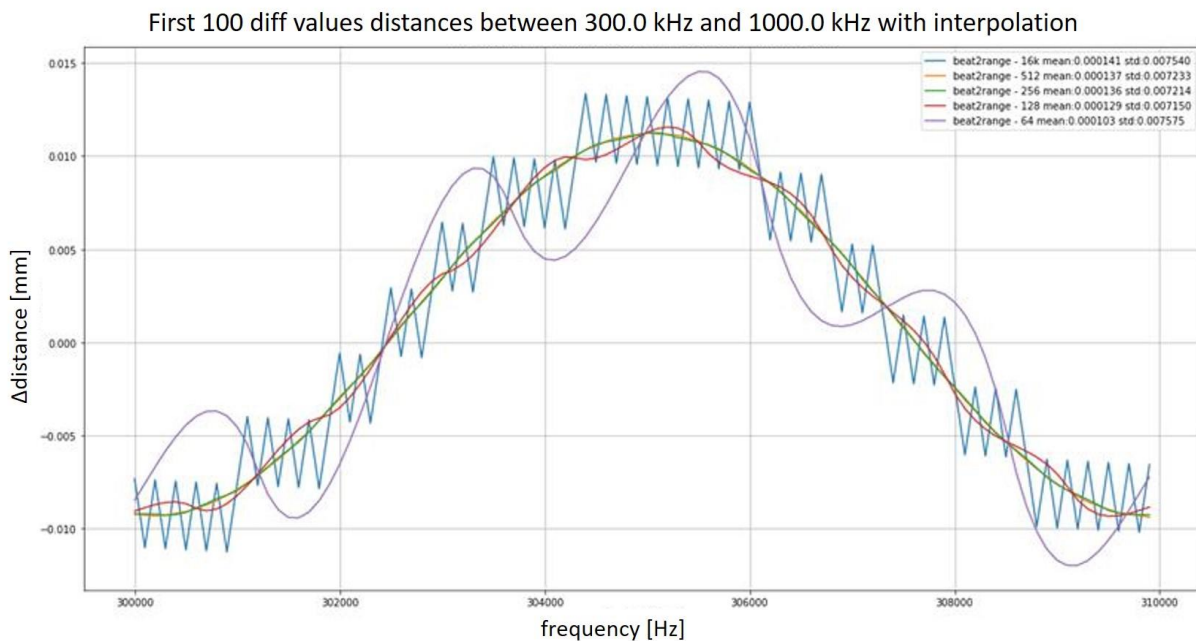


Figure 2 Simulation results of parabolic interpolation versus 16k FFT (FPGA); (blue) 16k FFT, (orange) 512 FFT with PI, (green) 256 FFT with PI, (red) 128 FFT with PI, (purple) 64 FFT with PI

The implementation uses CMSIS-DSP and μC -specific DSP instructions (Cortex-M7, FPU) [7]. Data paths and buffer sizes are chosen to maximize cache hit rates, minimize copy operations, and leverage DMA transfers efficiently. The entire chain remains modular to integrate later extensions (e.g., additional features like phase differences) with limited adjustments.

3.2 Phase-Coherent Acquisition and Trigger/Clock Concept

For multi-channel and multi-controller operation, phase coherence is essential so that master and slave deliver equivalent results for the same scene. The concept includes:

- External block triggering: A master GPIO block-triggers both μC ADCs. Block instead of sample-trigger reduces interrupt load and jitter sensitivity.
- Coherent clock distribution: Clock paths, line lengths, and impedances are symmetrically designed. Evaluation boards support measurement and optimization of trigger jitter and channel skews.
- Reproducibility: Trigger and clock paths are characterized; calibration routines (offset/gain/timing) are provided for interleaved operation.

3.3 ADC Operating Modes and Runtime Adaptation

To cover different measurement requirements and achieve a favorable energy-accuracy ratio, the ADCs support two modes:

- Independent mode: 1–4 channels at standard sampling rates; differential input is possible.
- Interleaved mode: 1–2 channels with doubled effective sampling rate via alternating sampling by two ADC cores. The required single-ended input is provided on the filter board via differential-to-single-ended conversion.

Resolution (e.g., 16-bit) and sampling rates (worst-case up to 7.2 MSps/channel) are runtime configurable. Switching is realized via a unified firmware API, allowing reaction to scenario changes (e.g., higher bandwidth, faster updates) without restart.

3.4 Firmware Architecture and Energy Management

Both μC boards (e.g., STM32H735, Cortex-M7) run identical firmware. Role assignment (master/slave) is determined deterministically at boot. Core elements:

- Power gating and duty cycling: The master can hard power off the slave μC and selected peripherals (e.g., frontends, filter/conversion) or place them in defined sleep states. Event-driven and partial-load operation significantly reduce average power consumption.
- Feature-first streaming: Only relevant features (e.g., peaks, distances, phases) are prioritized for transmission; raw data streaming remains optional for diagnostics and validation.
- Modular software: The pipeline is structured into reusable C modules (acquisition, preprocessing, FFT, interpolation, feature extraction, protocols). This accelerates testing, facilitates porting, and enables targeted runtime measurements of individual stages.

3.5 Edge Integration and Visualization

Sensor-near preprocessing occurs on the μC . A Raspberry Pi handles 2D/3D visualization (phase, magnitude, distance) for demos and customer discussions. For more complex workloads (e.g., tracking, higher-order filters), optional edge acceleration (e.g., Xilinx Versal, NVIDIA Orin) is prepared. The goal is a balanced sensor-edge partitioning: reduce network load, control latency, and lower energy system-wide.

4. Results

4.1 Compute and Measurement Quality

Simulations show: Smaller FFTs combined with parabolic peak interpolation deliver nearly continuous frequency estimation with consistent windowing and sufficient SNR [2, 3, 4]. In practice, this means that distance estimation results from the μ C chain can be equivalent to the FPGA reference while compute and memory access decrease. The approach is robust to moderate jitter/skew variations but requires clean clock/trigger design.

4.2 Power Consumption and Savings Potential

Measurements under full load (all components active; 16-bit, 7.2 MSps) show the following power consumption:

Table 2 Power consumption per module

Module	Quantity	Power consumption per module[W]	Total power consumption [W]
Backend	2	1,944	3,888
Filter/Conversion	4	1,022	4,088
Frontend adapter	4	0,564	2,256
Frontend	4	0,804	3,216
System			13,448

Compared to a 4-channel FPGA reference system, a worst-case power reduction of **24.577%** is achieved. In realistic partial-load scenarios (sensor-side preprocessing, reduced streaming, duty cycling), further reduction in average energy consumption is expected, as power gating and block triggering have stronger effect.

5. Environmental Life-Cycle Assessment

This chapter aims at a robust, life-cycle-based assessment of the climate impact of the μ C-based radar demonstrator. The manufacturing phase is based on BOM-derived results according to ISO 14040 framework guidelines [9]. The use phase is modeled with a concrete, plausible load profile and compared with an FPGA reference system. End-of-life effects are considered qualitatively since verified datasets are not yet available.

5.1 Manufacturing Phase: Results and Hotspots

BOM analysis shows that PCB manufacturing and integrated logic account for the largest share of manufacturing impacts. At system level (4× frontend, 4× adapter, 4× filter/conversion, 2× backend) the following finding emerges:

Table 3 CO₂e per module

Module	Quantity	Manufacturing [kg CO ₂ e]	Share of system [%]
Backend	2	13,381	32,6
Filter/Conversion	4	15,624	38,1
Frontend adapter	4	8,925	21,8
Frontend	4	3,068	7,5
System		40,994	100

Across all modules, the largest share is PCB processes and logic/processors (together around 85%). Aggregated by component class: PCB ≈ 18.43 kg CO₂e ($\approx 45\%$), CMOS logic ≈ 16.62 kg ($\approx 41\%$), MPU ≈ 4.59 kg ($\approx 11\%$); transistors and assembly are secondary. The leverage for further reductions lies primarily in PCB design and fabrication (area, layer count, materials, utilization) as well as functional integration and consolidation of logic/MPU.

5.2 Use Phase: Concrete Load Profile and Average Power

For operation, a practice-oriented, event-driven profile was defined that deliberately exploits partial load and standby phases. Modeling is based on ISO 14044 requirements [10]. The measured full load of 13.448 W (all modules active, 16-bit, 7.2 MSps) serves as reference. The four states and their relative power levels are:

- Full operation (30% of time): 100% of reference power.
- Partial load A – reduced rate/resolution-optimized (40%): 60% of reference power.
- Partial load B – “idle-armed” (20%): 30% of reference power.
- Standby/sleep (10%): 10% of reference power.

From these assumptions, the μ C system’s average power is 8.20 W. The FPGA reference system is set with 24.577% higher power based on measurements; its average power is thus 10.22 W.

5.3 Annual Consumption, Climate Impact, and Comparison to FPGA

For industrial use with 6,000 operating hours per year, the following energy amounts result. Conversion to greenhouse gas emissions uses an electricity emission factor of 0.40 kg CO₂e/kWh, based on reports from the German Environment Agency [11].

Table 4 Comparison μ C demonstrator vs. reference system

System	Average power [W]	Annual energy [kWh/a]	Use-phase GWP [kg CO ₂ e/a]
μ C demonstrator	8,20	49,22	19,69
FPGA reference	10,22	61,32	24,53
Difference	2,02	12,10	4,84

Thus, the μ C system avoids about 12.1 kWh and 4.84 kg CO₂e per year compared to the FPGA reference under the stated use profile. With higher utilization (more active time), savings grow proportionally; with a renewable electricity mix, absolute climate impact decreases, but the ratio between the two systems remains the same.

5.4 Life-Cycle Assessment, “Carbon Payback,” and Interpretation

The μ C system’s manufacturing impact is 40.99 kg CO₂e. Using the annual use-phase savings of 4.84 kg CO₂e (versus FPGA), the CO₂ “payback” is about 8.5 years if the reference system’s manufacturing is excluded. The assessment is based on ISO 14040 and ISO 14044 for life-cycle climate impact [9, 10], supplemented by emission data from UBA reports [11]. If the electricity mix is higher (e.g., 0.60 kg CO₂e/kWh) or operating time is 8,000 h/a, payback shortens to about 5.5–6.5 years; with very low-CO₂ electricity (0.25 kg CO₂e/kWh) it lengthens accordingly. Over 10 years of operation, use-phase savings in the baseline scenario sum to around 48 kg CO₂e; the net balance thus favors the μ C system. A full comparison of manufacturing impacts of both alternatives (μ C vs. FPGA) is planned once reliable PCF data of the reference are available; experience suggests the μ C approach’s lower logic/peripheral density offers additional manufacturing advantages.

5.5 Conclusions and Measures

The life-cycle analysis shows three key statements: First, hotspots clearly shift toward PCB and integrated logic, making design decisions at these levels particularly effective. Second, the μ C system achieves significant use-phase savings versus the FPGA reference under realistic partial-load operation; these are the key drivers of life-cycle advantages. Third, strict energy management (duty cycling, runtime switching of ADC modes, selective shutdown of entire modules) measurably accelerates “carbon payback.” For the next expansion stage, PCB optimizations (area/layer count), functional integration into logic/MPU, and inclusion of site- and customer-specific electricity mix scenarios in the assessment are recommended to make project-specific statements about amortization time more precise.

6. Discussion

The demonstrator shows that central steps of radar signal processing can be shifted from FPGA platforms to a microcontroller architecture without degrading measurement quality. The core is the combination of consistent windowing and an FFT-adjacent peak interpolation that delivers sub-bin precise frequency estimation with smaller FFT block sizes. Simulations and initial evaluations show that distance estimation can be equivalent to the FPGA reference if SNR, window choice, and clock/trigger design are executed cleanly. In practice, phase-coherent block triggering and coherent clock distribution have proven effective for handling channel skews and jitter. For interleaved operation, careful calibration of gain, offset, and timing skew remains essential.

From a system perspective, the technical strategy adds up to a measurable energy effect: Compared to a functionally equivalent FPGA solution, the μ C system's power consumption is 24.577% lower even in the worst case. Under a realistic, event-driven use profile (30% full operation, 40% partial load with reduced rate, 20% idle-armed, 10% standby), the μ C system's average power is 8.20 W versus 10.22 W for the reference. Over 6,000 operating hours per year, this corresponds to 49.22 kWh versus 61.32 kWh and an annual climate impact of 19.69 kg CO₂e versus 24.53 kg CO₂e (electricity EF 0.40 kg/kWh). The savings of around 12.1 kWh or 4.84 kg CO₂e per year are substantial and scale with higher utilization; with lower-CO₂ electricity mixes, absolute emissions decrease while the relationship between the systems remains.

From a life-cycle perspective, the μ C system's manufacturing impact (40.99 kg CO₂e) is clearly dominated by PCB manufacturing and integrated logic/MPU; together they cause about 85% of manufacturing emissions. The modular design (frontend, adapter, filter/conversion, backend) is a strategic advantage: It increases repairability and reusability and enables upgrades of individual modules instead of replacing entire systems. Under the described use profile, "carbon payback" versus FPGA use-phase emissions is about 8.5 years; with longer operating times or more CO₂-intensive electricity mixes it shortens; with very low-CO₂ electricity it lengthens. A full comparison of manufacturing impacts of both system alternatives is planned next; experience suggests the μ C approach's lower silicon and peripheral needs provide further benefits.

Technical limits remain: Interpolation is sensitive to SNR and leakage; window choice and data preparation are critical. Interleaved sampling increases effective rate but requires calibration discipline. Memory and cache sizes on the μ C platform demand efficient DMA and buffering strategies. For high-channel-count MIMO angle estimation and complex multi-target tracking, a hybrid approach is sensible: Sensor-side feature extraction on the μ C, compute-intensive steps on edge accelerators (e.g., Versal, Orin). This distribution aligns with Green ICT, as it balances network load, latency, and energy system-wide.

In sum, the demonstrator yields three robust insights: First, the μ C architecture with peak interpolation is technically viable and energy-efficient. Second, ecological benefits arise both in operation and, prospectively, in manufacturing through modularity and functional integration. Third, real-world impact depends strongly on operating profile and energy management—consistent power gating, runtime switching of ADC parameters, and feature-first streaming are central levers.

7. Conclusion and Outlook

This work demonstrates that a μ C-based backend can deliver central radar signal processing with significantly lower energy consumption without compromising measurement quality compared to an FPGA reference. This is enabled by FFT-compatible peak interpolation, a phase-coherent multi-channel concept, and strict energy and runtime management tailored to partial-load and event-driven operation. From a life-cycle perspective, the system shows clear use-phase advantages under realistic profiles; further design measures in PCB and logic/MPU can additionally reduce manufacturing impact.

In the medium term, a hybrid sensor-edge architecture is recommended for more demanding functions: The sensor efficiently extracts essential features, while edge accelerators take on complex processing steps. This enables stepwise access to MIMO angle measurement and robust multi-target tracking without undermining energy efficiency principles. Suppliers' product carbon footprints (PCFs) should be systematically requested and PCB fabrication options with lower environmental impact used to further optimize manufacturing.

In the long term, the demonstrator represents a transferable building block for resource-efficient sensing: One- to two-channel systems can often be moved directly to μ C, while high-channel, latency-critical applications benefit from sensor-edge partitioning. For exploitation, application-near demos (e.g., positioning, volume measurement), identification of validation partners, and a clear offering are key. The course is thus set: Greener signal processing becomes a designable property—technically robust, ecologically effective, and industrially compatible.

8. Acknowledgments

This whitepaper was created within the joint project “Green ICT @ FMD – Competence Center for Ecologically Sustainable ICT.” We thank the Federal Ministry for Research, Technology and Spaceflight (BMFTR) for funding. The project is funded under grant number 16ME0494 and coordinated by the Research Fab Microelectronics Germany (FMD). Our special thanks go to the project partners for their constructive collaboration and to Fraunhofer IZM for their support with life-cycle assessment. Responsibility for the content of this publication lies with the authors.

References

- [1] F. J. Harris, On the use of windows for harmonic analysis with FFT, IEEE, 66(1), 1978, pp. 51–83.
- [2] M. Gasior, J. L. Gonzalez, Improving FFT Frequency Measurement by Parabolic Interpolation, CERN, BIW 2004
- [3] J. I. Brown; D. C. Rife, R. R. Boorstyn, Tone Estimation, IEEE Trans. Info Theory, 20(5), 1974.
- [4] W. Grandke, Interpolation Algorithms for FFT Spectra, IEEE Trans. Measurement, 32(2), 1983.
- [5] TI, Interleaved Data Defense; ADC-Skew Resolution (Application Note SLAA105).
- [6] NI Whitepaper on Timing/Jitter Management.
- [7] Arm, CMSIS-DSP Library (CMSIS-DSP Repo).
- [8] STMicroelectronics, STM32H735 Datasheet: Cortex-M7 (Floating DSP).
- [9] ISO 14040, Environmental Management, Life Cycle Assessment Principles.
- [10] ISO 14044, Life Cycle Guidelines and Reporting Frameworks.
- [11] UBA, Electricity Emission Values Germany (Alliance CO2 Index 0.4).