Optimization and Performance Analysis of an FPGA-Based Structural Health Monitoring System for Aerospace Composite Structures

GRACIETH CAVALCANTI BATISTA,
LEANDRO BACHEGA TRUJILHO,
CARL-MIKAEL ZETTERLING, JOHNNY OBERG,
OSAMU SAOTOME, MALIN AKERMO and PER HALLANDER

ABSTRACT

Structural Health Monitoring (SHM) is a transformative technology in aerospace engineering, enabling real-time assessment of aircraft structural integrity to enhance safety, reliability, and maintenance efficiency. Traditional inspection methods, such as visual inspections and non-destructive testing (NDT), require manual labor and aircraft downtime. In contrast, a SHM system offers continuous and automated monitoring, thereby reducing operational costs and improving aircraft availability. This paper presents an optimization study of a previously developed FPGA-based SHM system for carbon fiber-reinforced polymer (CFRP) laminates. The system uses piezoelectric transducers (PZTs) to both excite and receive Lamb waves and employs signal processing and machine learning techniques for damage detection. In the original implementation, signal processing steps such as discrete wavelet transform (DWT) and feature extraction were executed in MATLAB. In this work, a fixed-point hardware implementation of the DWT was developed and integrated into the acquisition pipeline. The DWT stage, originally implemented in MATLAB, was migrated to hardware, resulting in a processing time reduction from 1.43 seconds to 5.68 microseconds, which led to over 250,000 times speedup. The classification stage, based on a support vector machine (SVM) enhanced with Mahalanobis distance, was preserved in hardware. A comparative analysis shows that the embedded version achieves the same classification accuracy as the software-based approach, with no loss of precision. The results demonstrate that feature extraction can be executed in real-time, advancing the system toward fully embedded SHM applications.

INTRODUCTION

Composite structures are widely used in aerospace engineering due to their high strength-to-weight ratio, corrosion resistance, and design flexibility. However, these materials are prone to complex failure modes such as delaminations, matrix cracking, and barely visible impact damage (BVID), which are difficult to detect with traditional

Gracieth Cavalcanti Batista, Postdoctoral researcher, Email: gracieth@kth.se. Lightweight structures group, Department of Engineering Mechanics at the School of Engineering Science (SCI), KTH Royal Institute of Technology, Stockholm, Sweden

inspection techniques [1, 2]. Maintenance strategies based on visual inspection and non-destructive testing (NDT) require scheduled downtime and manual intervention, increasing operational costs and potentially missing damage between inspections [3, 4].

Structural Health Monitoring (SHM) systems aim to overcome these limitations by enabling continuous or on-demand evaluation of structural integrity. These systems reduce the need for periodic disassembly and inspection, supporting condition-based maintenance strategies and increasing aircraft availability [5]. In this context, SHM based on guided ultrasonic waves, especially Lamb waves, has gained considerable attention. Lamb waves are sensitive to various defect types and can propagate over large areas, making them ideal for composite materials [6, 7].

Lamb wave-based SHM systems typically rely on piezoelectric transducers (PZT) to generate and receive signals, combined with digital signal processing algorithms for defect detection and classification. Although advanced signal processing and machine learning algorithms such as wavelet transforms, principal component analysis (PCA), and support vector machines (SVMs) can be implemented in software. These implementations often struggle to meet real-time constraints in embedded applications due to limitations in latency, parallelism, and energy efficiency [8, 9]. In contrast, hardware-based solutions using field programmable gate array devices (FPGAs) allow the integration of parallel signal processing and classification pipelines with deterministic execution and reduced power consumption, making them more suitable for in-situ SHM applications.

In previous work [10], we developed an FPGA-based SHM system for carbon fiber reinforced polymer (CFRP) laminates that integrates signal acquisition, preprocessing, feature extraction using discrete wavelet transform (DWT) and power spectral density (PSD) analysis, and classification using a SVM with Mahalanobis distance for outlier rejection. The system was validated experimentally using two CFRP specimens with quasi-isotropic layups: one containing embedded defects and the other surface damage. With one transmitter and four strategically placed receivers, the system achieved high accuracy in classifying both internal (96.25%) and external (97.5%) defects, while maintaining real-time performance.

This paper builds on that system by presenting a detailed optimization study targeting signal processing and hardware implementation. The focus is on reducing latency, optimizing the precision and representation of fixed-point data, and minimizing FPGA resource utilization. A new embedded architecture for the DWT was developed to replace the previous MATLAB-based implementation. A comparison between the software and hardware versions of the signal processing chain is presented. The results confirm that the hardware implementation maintains the same classification performance, highlighting the system's evolution toward a fully embedded and online-capable SHM solution.

SYSTEM'S OVERVIEW

The SHM system under analysis was specifically developed for real-time defect detection in CFRP laminates. It combines guided wave excitation, multichannel signal acquisition, and embedded machine learning classification, all implemented on a single FPGA platform. The goal of the system is to classify internal and external defects in composite structures using Lamb waves as a sensing modality. Figure 1 shows a picture of the system setup.

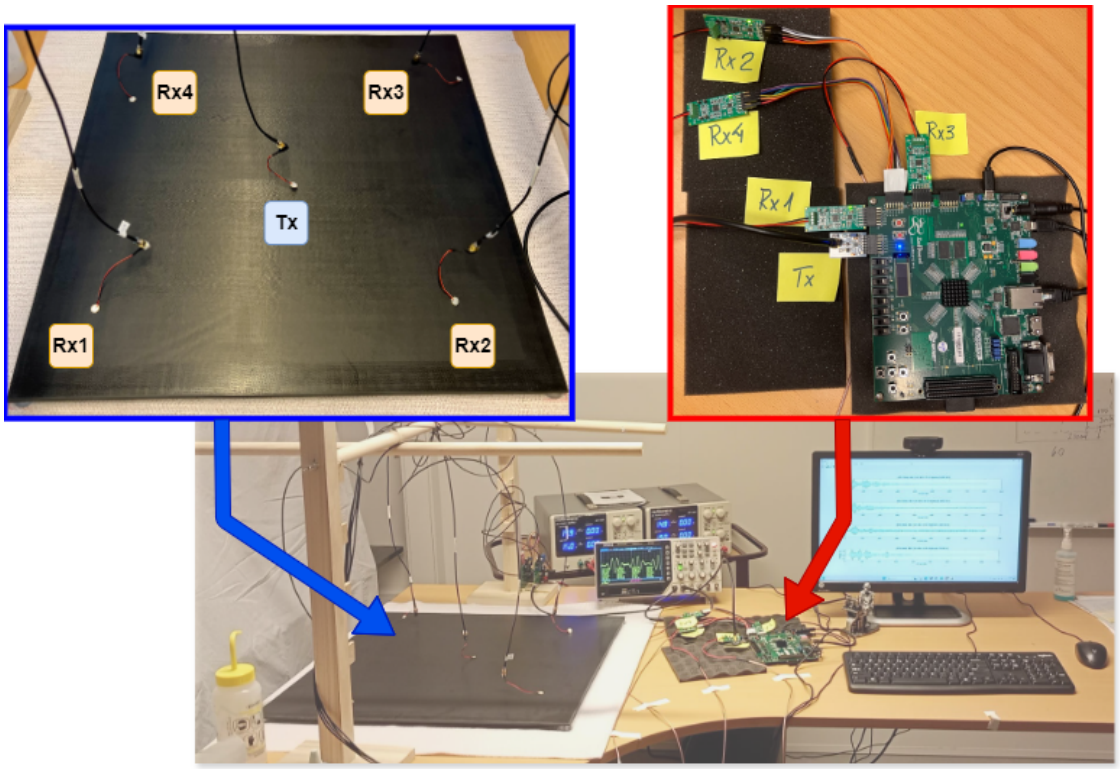

Figure 1. Setup of the SHM system described in [10].

The proposed system is implemented on a ZedBoard, which integrates a Xilinx Zynq-7000 System-on-Chip (SoC) composed of two main elements: a Processing System (PS), which includes a dual-core ARM Cortex-A9 processor, and a Programmable Logic (PL) fabric used for hardware acceleration. In this configuration, the PL is responsible for signal acquisition and hardware-based classification, while the PS handles system configuration, data export, and communication with external tools.

The experiments were carried out on two quasi-isotropic CFRP plates with dimensions of $600 \times 600 \times 5$ mm. The first specimen contains embedded defects introduced by inserting a fluorinated ethylene propylene (FEP) film between the plies during the lamination process. The second specimen exhibits surface damage, including circular cutouts and low-velocity impact marks. These two configurations emulate internal delamination and external damage typically found in composite aircraft components.

Each plate is instrumented with five PZT transducers: one central transmitter and four receivers placed symmetrically, one at each corner of the specimen. This layout ensures coverage of the propagation area while minimizing edge reflections. The excitation signal is a 15 kHz 3.5-cycle tone burst, selected based on finite element simulations to excite the A0 Lamb wave mode, which offers high sensitivity to delamination and surface cracks in thin composite structures.

The PL controls both signal generation and data acquisition. The tone burst is sent to a digital-to-analog converter (DAC) connected to an external amplifier. Reflected and transmitted waveforms are captured by four analog-to-digital converters (ADCs) interfaced to the FPGA through Pmod connectors. Each channel is sampled independently, and acquisition is synchronized with the excitation pulse to preserve timing accuracy

across all sensors. The raw signal data are stored in the block RAM (BRAM) and made available to the PS.

All signal processing steps, including DC offset removal, DWT with three decomposition levels, PSD estimation, and PCA, are performed offline in MATLAB. The output of the PCA projection is a reduced-dimension feature vector, which is exported to the FPGA for SVM classification on hardware.

The classification module is described in Verilog using fixed-point arithmetic. It applies a linear SVM and evaluates the Mahalanobis distance to improve decision reliability, especially in ambiguous cases. Inference is performed entirely in the PL, with deterministic latency.

SYSTEM OPTIMIZATION

To enable embedded processing of ultrasonic signals, a fixed-point implementation of the signal processing pipeline was developed using AMD/Xilinx Vitis HLS. The entire pipeline is written in C++ and encapsulated in a module named *dwt_pipeline_full*, composed of three functional blocks connected via *hls::stream* objects to support parallel execution with *#pragma HLS DATAFLOW*.

All data are represented in *ap_fixed<32,12>* format, providing 12 integer bits and 20 fractional bits. This configuration preserves the precision of the db4 wavelet coefficients across all decomposition levels while ensuring compatibility with the target hardware platform.

The *int_to_fixed* block receives raw samples from the ADCs in *uint16_t* format, applies a DC offset correction (subtracting 2048), and converts the data to fixed-point. The core transformation is performed by the *dwt_core_pipeline*, which executes a three-level DWT using Daubechies-4 wavelets. The final block, *fixed_to_int*, scales and converts the level-3 coefficients back to *int16_t*, making the output compatible with AXI-based data acquisition and communication systems.

This pipeline was synthesized as an independent IP block and integrated into the existing acquisition framework without altering the original finite-state machine (FSM) or FIFO-based buffering logic [11]. The updated design enables real-time pre-processing in hardware and supports the selective retrieval of raw or filtered data via MATLAB scripts communicating with the embedded processor.

The new pipeline was inserted without disrupting the original control flow. As shown in Figure 2, the updated design enables the selective retrieval of both raw ADC data and filtered DWT output via MATLAB scripts using TCP/IP communication with the embedded processor.

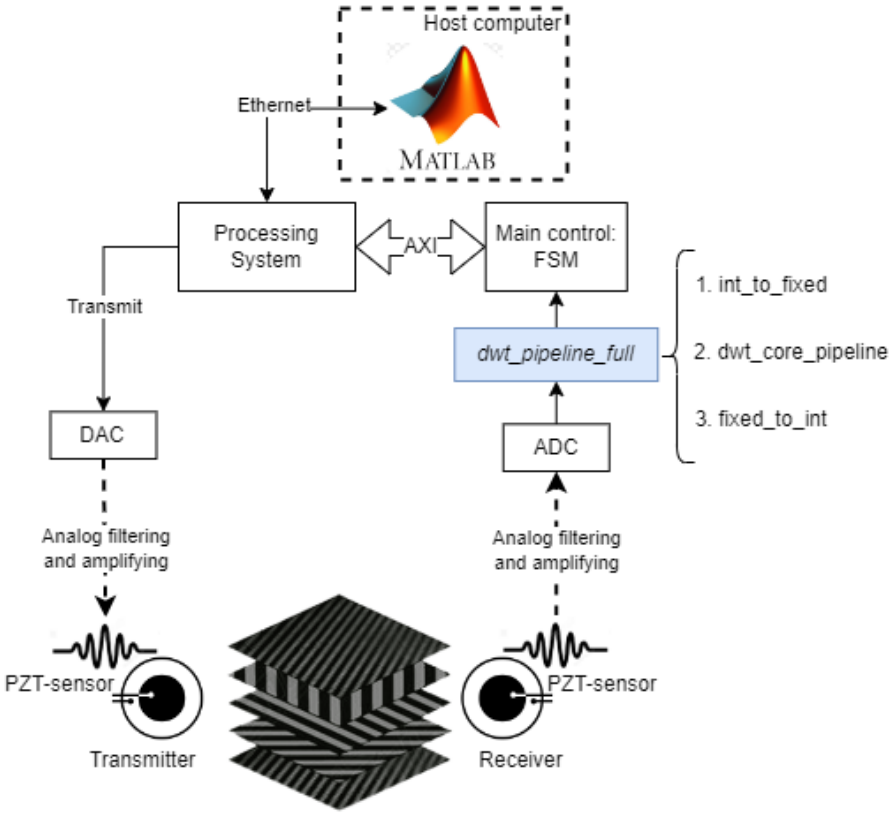

Figure 2. Detailed block diagram of the data acquisition phase of Figure 1.

The internal structure of the embedded DWT pipeline is shown in Figure 3. After digitization, the signal is converted from integer to fixed-point format by the *int_to_fixed* block. A dedicated stage then performs DC offset correction by subtracting 1.25 V from the input signal. The corrected signal is processed by a 3-level DWT using Daubechies-4 filters. Each level (D1, D2, and D3) applies filtering and downsampling to isolate features across different frequency bands. Finally, the *fixed_to_int* block converts the output to integer format for export to MATLAB.

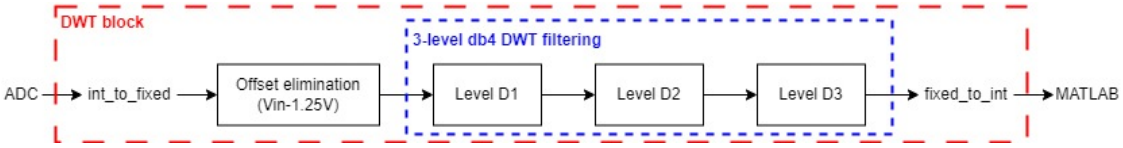

Figure 3. Detailed block diagram of the embedded DWT block.

RESULTS

The experimental validation was conducted on two CFRP specimens: one with embedded FEP film to simulate internal delaminations, and another with surface damage, including circular cutouts and impact marks. Both configurations were tested using the embedded excitation and acquisition system, with one transmitter and four receivers

placed symmetrically on each plate.

Each acquisition consisted of recording Lamb wave responses from the four receivers. A total of 34 signals were recorded per receiver, resulting in 136 waveforms per test case. All signals were initially acquired in *uint16_t* format and stored in local memory. The embedded processor retrieved the raw signals via BRAM and made them accessible to MATLAB over TCP/IP.

Figure 2 illustrates the ability to switch between unfiltered ADC data and filtered signals obtained after three-level wavelet decomposition. DWT filtering enhances time-localized features and suppresses high-frequency noise. The DWT outputs are exported in the *int16_t* format and plotted in MATLAB for visualization.

After implementing the *dwt_pipeline_full* block on the FPGA-Zedboard device with a clock period of 10 ns, it was possible to observe the results of Table I. This table summarizes the overall FPGA resource usage with and without the embedded DWT pipeline. The addition of signal processing blocks results in a noticeable increase in LUT, memory, and DSP utilization. The column *Delta* indicates the absolute difference in resource usage between both configurations. Even with these increases, the total usage remains well within the capacity of the Zynq-7000 device, confirming the feasibility of the hardware implementation.

TABLE I. FPGA resource utilization for the system with and without the embedded DWT pipeline.

Resource	**Available**	**Without DWT**	**With DWT**	**Delta**	**Overhead (%)**
LUT	53,200	5,275	8,936	+3,661	+69.4
LUTRAM	17,400	343	585	+242	+70.5
FF	106,400	8,283	13,089	+4,806	+58.0
BRAM	140	17.5	30	+12.5	+71.4
DSP	220	0	16	+16	–

The fixed-point processing pipeline was synthesized using AMD Vitis HLS 2023.2 for the Zynq-7000 device. The design includes three sequential blocks connected via *hls::stream* interfaces and parallelized using *#pragma HLS DATAFLOW*.

The total latency of the entire DWT block from input to output is 2038 clock cycles. Table II summarizes the utilization of resources per module.

TABLE II. Hardware resource usage and latency for each module in the DWT block.

Module	**Latency (cycles)**	**LUTs (%)**	**DSPs (%)**
`int_to_fixed`	8	2.1	0.5
`dwt_core_pipeline`	1780	7.8	8.4
`fixed_to_int`	250	2.3	1.1
Total	**2038**	**12.2**	**10.0**

The total latency of 1780 clock cycles comes from the three levels of the DWT, each processing a progressively smaller signal:

- Level D1: 1022 outputs $\rightarrow$ 1022 clock cycles
- Level D2: 508 outputs $\rightarrow$ 508 clock cycles
- Level D3: 250 outputs $\rightarrow$ 250 clock cycles

The classification accuracy obtained in this version of the system was identical to that achieved in the previous implementation, in which all signal processing, including the wavelet decomposition, was performed offline in MATLAB using floating-point arithmetic. This confirms that the fixed-point hardware implementation of the DWT does not degrade classification performance, preserving precision while enabling embedded execution.

DISCUSSION AND COMPARISON

The results obtained from the hardware implementation of the DWT pipeline show excellent agreement with the MATLAB reference model. A point-by-point comparison between the D3 outputs of the HLS pipeline and MATLAB revealed a mean absolute error of only 0.0003, with a maximum absolute error of less than 0.0014, corresponding to a minimum discrepancy of less than 0.001%. These results validate the reliability of the 32-bit fixed-point representation, where 12 bits are allocated to the integer part, and demonstrate that the sequential structure built with *hls::stream* objects and explicit parallelism using *#pragma HLS DATAFLOW* preserves the integrity of the signal at the three levels of decomposition. The precision was further confirmed by the accuracy of the data conversion stages between the integer and fixed-point domains, with reconversion errors of less than 0.002%. These results confirm that the pipeline is ready to replace the MATLAB-based processing stage with an efficient and deterministic hardware implementation on the programmable logic.

To better illustrate the performance trade-off introduced by embedding the DWT in hardware, Table III provides a direct timing comparison between the software and hardware implementations of the DWT stage. In MATLAB, the processing time for a single acquisition was measured as 1.43 s, while the embedded implementation requires only 5.68 µs, assuming a 100 MHz clock. This corresponds to a speedup factor of over 250,000×. These results confirm the feasibility of real-time feature extraction in hardware, while classification still depends on a pre-trained model and PCA transformation performed offline.

TABLE III. Execution time and processing characteristics of the DWT stage in software and hardware implementations.

Implementation	**Processing Time**	**Latency**	**Throughput**
MATLAB (software)	1.43 s	–	–
FPGA (hardware)	5.68 µs	568 cycles @100 MHz	≈100 MS/s

CONCLUSION

The goal of this work was not to propose a new SHM framework, but to improve the performance, scalability, and robustness of an existing real-time system, making it more suitable for deployment in realistic aerospace scenarios. By getting closer to end the loop between data acquisition and in-situ defect classification, the system represents a practical step towards onboard diagnostic capability for composite aircraft structures.

The main contribution of this work was the transition of the signal processing stage, previously implemented in MATLAB, to a fully embedded hardware version. A comparative evaluation demonstrated that the embedded implementation achieves the same classification accuracy as the software-based approach, confirming that signal processing can be executed in hardware without loss of precision.

In addition to preserving classification performance, the integration of the signal processing pipeline into hardware significantly improved the system's responsiveness. The reduction in execution time enables real-time feature extraction, which is essential for autonomous operation. The optimized architecture was incorporated into the existing acquisition system with minimal modifications, demonstrating the system's scalability and compatibility with embedded platforms. These improvements consolidate the processing chain and advance the system toward practical, low-latency SHM deployment.

REFERENCES

1. Rocha, H., C. Semprimoschnig, and J. P. Nunes. 2021. "Sensors for process and structural health monitoring of aerospace composites: A review," *Engineering Structures*, 237:112231, doi:10.1016/j.engstruct.2021.112231.
2. Ogunleye, A., R. D. Pollock, and A. E. Lyeo. 2024. "Advanced Sensors and Sensing Systems for Structural Health Monitoring in Aerospace," *Advanced Engineering Materials*, 26(4):2201041, doi:10.1002/adem.202201041.
3. Sause, M. G. R., E. Jasiūnienė, and R. Pullin. 2021. "Introduction to Structural Health Monitoring in Aerospace," in M. G. R. Sause and E. Jasiūnienė, eds., *Structural Health Monitoring Damage Detection Systems for Aerospace*, Springer, Springer Aerospace Technology, ISBN 978-3-030-72191-6, pp. 1–10, doi:10.1007/978-3-030-72192-3_1.
4. Pollock, R. D., A. J. F. Jones, and M. D. Gilchrist. 2021. "The Need for Aerospace Structural Health Monitoring: A Review of Failures," *Journal of Intelligent Material Systems and Structures*, 32(12):1395–1413, doi:10.1177/1045389X20984787.
5. Pant, S., Z. S. Khodaei, and M. G. Droubi. 2021. "Monitoring Tasks in Aerospace," in *Structural Health Monitoring Damage Detection Systems for Aerospace*, Springer, pp. 5–14, doi:10.1007/978-3-030-72192-3_2.
6. Ricci, F., E. Monaco, N. Boffa, L. Maio, and V. Memmolo. 2022. "Guided waves for structural health monitoring in composites: A review and implementation strategies," *Progress in Aerospace Sciences*, 129:100790, ISSN 0376-0421, doi: https://doi.org/10.1016/j.paerosci.2021.100790.
7. Du, F., S. Wu, Z. Tian, F. Qiu, C. Xu, and Z. Su. 2023. "An Improved Prototype Network and Data Augmentation Algorithm for Few-Shot Structural Health Monitoring Using Guided Waves," *IEEE Sensors Journal*, 23(8):8714–8726, doi:10.1109/JSEN.2023.3257366.
8. Azad, M. M., S. Kim, Y. B. Cheon, and H. S. Kim. 2024. "Intelligent structural health monitoring of composite structures using machine learning, deep learning, and transfer learning: a review," *Advanced Composite Materials*, 33(2):162–188, doi: 10.1080/09243046.2023.2215474.

9. Elahi, H., M. Eugeni, and M. Gaudenzi. 2021. “Energy-autonomous piezoelectric sensors for aircraft structural health monitoring: A review,” *Microsystem Technologies*, 27:3453–3465, doi:10.1007/s00542-020-05017-y.
10. Batista, G. C., C.-M. Zetterling, J. Öberg, and O. Saotome. 2024. “SHM System for Composite Material Based on Lamb Waves and Using Machine Learning on Hardware,” *Sensors*, 24(23), ISSN 1424-8220, doi:10.3390/s24237817.
11. Alhallak, A. 2024. *PZT Scanner System for a Structural Health Monitoring Application*, Master’s thesis, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, supervisors: Gracieth Cavalcanti Batista, Johnny Öberg; Examiner: Carl-Mikael Zetterling.