symmetry

MDPI

Article

The Hardware-Architecture Design of a Hybrid DWT–ADTF-Based Real-Time ECG-Denoising System: An Optimized FPGA Implementation

Hanane Elferdaoussi [1,*], Wissam Jenkal [1], Mostafa Laaboubi [1,2], Azzedine Dliou [1,2] and Rachid Latif [1]

[1] Laboratory of Systems Engineering and Information Technology (LISTI), National School of Applied Sciences, Ibn Zohr University, Agadir 80000, Morocco; w.jenkal@uiz.ac.ma (W.J.)
[2] Faculty of Applied Sciences, Ibn Zohr University, Ait Melloul 86150, Morocco
* Correspondence: hanane.elferdaoussi@edu.uiz.ac.ma

Abstract

Electrocardiogram (ECG) signals are highly susceptible to various types of noise, including power-line interference, electromyographic activity, motion artifacts, and baseline drift, all of which compromise diagnostic reliability. Although software-based denoising approaches have the potential for high performance, their computation cost and power requirements, and their nondeterministic execution are not suitable for the stringent real-time and power requirements in embedded and wearable systems. Such limitations restrict the practical applicability of computationally intensive denoising and highlight the need for hardware designs to support the required performance in long-duration monitoring systems. To address these limitations, this work introduces a hardware-efficient ECG-signal-denoising architecture based on the hybrid DWT–ADTF algorithm, which has been selected for its favorable trade-off between low algorithmic complexity and high denoising performance. The proposed design employs a fully pipelined structural-RTL architecture with dual-clock synchronization and memory-aware buffering to ensure deterministic real-time throughput under embedded constraints. The architecture is prototyped on an Intel Cyclone V SoC FPGA environment, which is used as a validation platform. Across multiple noise conditions, the proposed architecture consistently improves signal quality; for instance, using MIT-BIH records with additive white Gaussian noise at 5 dB, the system achieves an SNR improvement of 9.23 dB to 10.00 dB and a PRD of approximately 18%, with 1.09 ms system latency, 426 mW power consumption, and under 5% logic utilization. These results demonstrate the suitability of the proposed architecture for low-power, real-time biomedical monitoring devices.

Keywords: ECG; DWT; ADTF; FPGA; real-time signal processing; hardware design

Academic Editor: Jie Yang

Received: 9 June 2026
Revised: 6 July 2026
Accepted: 10 July 2026
Published: 14 July 2026

1. Introduction

Performing an ECG is one of the most common non-invasive approaches for monitoring and diagnosing cardiovascular disorders. Beyond clinical use, ECG signals support continuous cardiorespiratory monitoring, seizure detection through cardiac irregularities, biometric authentication, and AI-driven real-time rhythm analysis, all of which require accurate and reliable waveform acquisition [1–4]. However, the diagnostic quality of the electrocardiogram signal is often impaired by the presence of power-line interference, electromyogram activity, movement artifacts, and baseline drift. All of these affect the PQRST wave morphology and make it difficult to interpret ECG signals automatically or manually.

Therefore, denoising of ECG signals is an essential preprocessing condition that requires careful consideration, particularly in long-duration monitoring scenarios [5–7].

A number of denoising techniques have been proposed to improve ECG signals. These techniques range from classical linear techniques to more complex adaptive and transform-based techniques. Conventional infinite and finite impulse response (IIR/FIR) filters have widely been used to remove baseline wander and high-frequency interference due to their well-defined frequency responses and relatively low computational complexity [8–10]. Furthermore, the performance of adaptive filter algorithms such as least mean squares (LMS) and recursive least squares (RLS) has also improved the capabilities for suppressing noisy interferences through the dynamic modification of filter coefficients to address the time-varying interference, including motion artifacts and powerline disturbances [11,12]. In this regard, the performance of the class of signal processing algorithms based on discrete wavelet transforms (DWT) [13] has also been reported to be effective for the isolation and attenuation of the noisy components at various scales while preserving diagnostically relevant features of the ECG waveform. More recently, deep-learning-based approaches have also emerged as a major research direction for ECG denoising. Convolutional neural network (CNN) autoencoders and related hybrid architectures have demonstrated remarkable denoising performance by learning robust signal representations directly from noisy ECG recordings [14,15].

Despite their strong denoising performance, the implementation of ECG-denoising algorithms in embedded systems remains subject to several practical constraints. While techniques that make use of multi-scale transformations, iterative optimization schemes, floating-point arithmetic, and significant buffering can pose significant challenges in terms of the tight timing constraints that are inherent in the design of wearable and battery-powered devices, techniques that offer too much simplicity can fail to provide the required diagnostic information while maintaining the characteristics of the waveforms [16]. This represents a persistent trade-off between denoising efficiency and the practical constraints of embedded implementation [17].

In this context, hybrid techniques such as the DWT–ADTF method [18,19] offer a much more balanced compromise: the DWT stage performs a preliminary but efficient elimination of the noise-dominated detail coefficients, avoiding the need for the laborious and computationally expensive procedure involved in a full wavelet-thresholding scheme. The subsequent ADTF stage is a fixed-parameter local-thresholding technique that further suppresses the remaining noise at a reasonable computational cost. The peak correction step is a final operation that restores high-amplitude features like the R and S peaks that may have been slightly diminished during the processing. The above approach combines moderate algorithmic complexity with strong morphological preservation, thereby providing an even more favorable framework for embedded system realization compared to more computationally intensive alternatives. Despite its moderate computational demands, the DWT–ADTF pipeline derives significant advantage through architecture-level optimization in order to satisfy the rigorous energy and resource constraints of embedded platforms. Since the denoising operation represents an early stage of continuous activity in the processing chain, an effective hardware realization is necessary for preserving system-level margins and to support diagnostic-grade operation in constrained environments.

These system-level constraints are increasingly intensified by the growing shift toward continuous and autonomous cardiac monitoring. The challenge is to achieve strict energy limits, real-time determinism, and compact, integrated hardware while maintaining performance reliability in wearables and out-of-hospital environments. Long-duration wireless monitoring has also exacerbated such limitations, which has emphasized the need for architectures that meet continuous real-time requirements without any compromise

in terms of the quality of the measurements. At such times, traditional processors and software solutions have been associated with increased power consumption, irregular execution timing, and scalability issues. Conversely, hardware-oriented architectures provide a more appropriate basis, as this affords a deterministic dataflow execution, fine control of computation-based parallelism, and efficient resource management. These aspects will facilitate more sustained denoising capabilities in terms of embedded system constraints. The field-programmable gate array (FPGA) platforms offer a practical and reconfigurable environment for realizing hardware-oriented architectures, as well as assessing aspects of functional correctness, timing, and efficiency under representative hardware conditions.

Building on the established DWT–ADTF framework, this work develops an architecture-aware FPGA design based on embedded ECG-monitoring-implementation constraints. In contrast to prior FPGA realizations that primarily demonstrated functional feasibility, the proposed architecture emphasizes deterministic real-time dataflow, hardware integration of the morphological correction stage, memory-efficient streaming organization, and resource-aware RTL optimization. This design strategy preserves denoising performance while reducing hardware costs and providing sufficient architectural flexibility to support subsequent ECG-signal-processing stages.

The key contributions of this work are as follows.

1. A dedicated hardware architecture for the DWT–ADTF ECG denoising, explicitly tailored for the stringent latency, energy and resource constraints of wearable and continuous telemonitoring systems.
2. A fully pipelined, structural-RTL-design methodology, incorporating dual-clock-domain synchronization and memory-aware buffering to enable deterministic and continuous real-time dataflow.
3. An optimum hardware-level design, leveraging controlled parallelism, limited combinational depth, and effective FIFO/register allocation to achieve a strong balance between throughput, silicon area, and power consumption.
4. Comprehensive end-to-end FPGA-based experimental validation, demonstrating robust denoising performance (9.23–10.00 dB SNR improvement and 18% PRD at 5 dB AWGN) with low latency (1.09 ms), low power (426 mW), and minimal resource usage (<5% logic utilization).

The remainder of this paper is organized as follows: Section 2 reviews related works on ECG-denoising algorithms and FPGA implementations. Section 3 describes the proposed methodology, including the architectural design, dataflow organization, and signal processing pipeline. Section 4 reports the experimental results and hardware validation of the proposed system. Finally, Section 5 concludes the paper and proposes future directions.

2. Related Work

2.1. ECG-Denoising-Techniques Overview

Noise removal from ECG signals has been a focus of considerable research, with techniques varying from classical linear filtering to adaptive and transform-based methods. Empirical Mode Decomposition (EMD) has been widely employed for non-linear and non-stationary signals due to its ability to decompose signals into Intrinsic Mode Functions (IMFs) [20,21]. Nevertheless, EMD is computationally expensive, and suffers from mode mixing, which might compromise the accuracy of noise separation. This drawback has been addressed by employing an Ensemble Empirical Mode Decomposition (EEMD), in which white noise is introduced to enhance accuracy at the expense of processing time [5,22].

Adaptive filters such as the Least Mean Square (LMS) filter and the Normalized Least Mean Square (NLMS) filter have proven to be effective in the application of real-time denoising in the absence of a reference signal. Nonetheless, the stability and the rate of

convergence depend on the characteristics of the noise. Moreover, the iterative process in the application of the filters increases the computational complexity. Traditional filters such as the Finite Impulse Response Filter (FIR) and the Infinite Impulse Response Filter (IIR) remain dominant. These filters have the advantage of preserving the morphology of the waveforms due to the linear phase characteristics. They have the disadvantage of requiring a high-order filter to produce a sharp cutoff. On the other hand, the IIR filters have the advantage of producing a sharp cutoff using filters of lower orders, though they can introduce phase distortion [8–10].

Transform-based methods, in particular the Discrete Wavelet Transform (DWT), have been a source of considerable interest given their ability to localize features in terms of time–frequency. DWT-based filtering can efficiently suppress baseline drift and high-frequency noise while preserving clinically relevant information [13]. More recently, hybrid approaches have been studied, combining DWT with adaptive filtering or thresholding techniques to yield enhanced denoising capability. Nonetheless, these approaches are mostly evaluated in software-based environments with minimal concern for real-time performance under resource-constrained embedded environments. Adaptive Dual Threshold Filtering (ADTF), originally inspired by the double-threshold methods employed in image processing, has been especially adapted for ECG denoising. ADTF employs two local-signal statistic-based adaptive thresholds to eliminate the noise outliers without limiting waveform distortion [18]. Subsequently, an enhanced variant of DWT–ADTF has been developed in [19], incorporating additional adaptive processing stages to further improve denoising performance at the expense of increased computational complexity. Consequently, the standard DWT–ADTF formulation remains particularly well suited for embedded hardware implementation as it provides a favorable trade-off between denoising performance and computational complexity, whereas the enhanced variant introduces additional processing stages that increase hardware complexity.

Table 1 summarizes the inherent trade-offs across common ECG-denoising techniques [16]. FIR and IIR filters remain computationally effective but offer limited accuracy or may introduce distortion. Median and adaptive approaches provide improved robustness, though convergence problems and iterative adaptation increase the cost. Transform-based techniques such as DWT offer strong noise reduction with moderate complexity and reasonable QRS morphology preservation. Hybrid methods like DWT–ADTF take advantage of these strengths, giving improved accuracy while keeping the computational burden acceptable for embedded hardware.

Table 1. Comparative overview of representative ECG-denoising methods, highlighting their algorithmic complexity, reported quantitative metrics, and associated trade-offs.

Method	Complexity	Evaluation Parameters (Exact Values)	Trade-Off
FIR filter	Low: simple convolution	PRD = 26.63% (MIT-BIH 100, 60 Hz interference) [23]	Low cost; limited suppression of non-stationary noise.
IIR filter	Low: recursive filtering	PDR = 7.73% (MIT-BIH 100, 60 Hz interference) [23]	Efficient; possible phase distortion in recursive paths.
Adaptive LMS/NLMS	Medium: iterative adaptation	SNR_{imp} = 17.16 dB; RMSE = 0.0298 (MIT-BIH 100, 10 dB AWGN) [24]	Adaptive; convergence depends on step size and reference accuracy.
DWT (multi-level transform)	Medium: decomposition and reconstruction	$SNR_{imp} \geq 8$ dB; MSE = 0.000893; PRD = 7.75% (MIT-BIH 100, 10 dB AWGN) [13]	Multi-resolution; preserves ECG morphology with moderate complexity.
Hybrid (DWT + ADTF)	Medium–High: multi-resolution and adaptive thresholds	SNR_{imp} = 9.70 dB; MSE = 0.0044; PRD = 18.26% (MIT-BIH 100, 5 dB WGN) [18]	Combines robustness of DWT with adaptivity of ADTF; higher accuracy at moderate cost.

Table 1. *Cont.*

Method	Complexity	Evaluation Parameters (Exact Values)	Trade-Off
EMD	High: signal decomposition (IMFs)	SNR_{imp} = 9.30 dB; MSE = 0.02022; PRD = 34.32% (MIT-BIH 100, 0 dB AWGN) [25]	Effective for non-linear, non-stationary noise; mode mixing and high computational cost.
EEMD	Very High: multiple EMDs with added noise	SNR_{imp} = 10.08 dB; MSE $\approx$ 0.0015 (MIT-BIH 115, random noise) [26]	Improved decomposition accuracy; very costly due to repeated runs.

2.2. Implementation Strategies for Embedded ECG Denoising

Beyond algorithmic efficiency, the selection of implementation strategy is critical for enabling real-time ECG denoising under embedded constraints.

2.2.1. CPU-Based Implementations

Various works have addressed the implementation of CPU-based solutions for ECG-signal denoising on embedded-class processors using C/C++ or MATLAB environments. Performance evaluations include assessments of ADTF performance on multi-core ARM architectures with thread-level parallelization [27]. These platforms provide a practical environment for functional validation and are capable of achieving near real-time execution under moderate workloads. However, their performance and timing remain tied to the variability of operating system scheduling as well as memory accesses, which limits determinism. Furthermore, sustained multi-hour operation on general-purpose processors leads to non-negligible energy consumption, making such implementations unsuitable for long-duration, battery-powered ECG monitoring. CPU-based solutions, therefore, remain useful for algorithmic prototyping but do not fulfill the stringent real-time and low-power constraints required in wearable biomedical systems. GPU platforms have also been considered for computationally intensive denoising algorithms, especially deep-learning and transform-domain methods. Despite substantial parallel throughput, the high energy demands and lack of deterministic execution prevent such resources from being employed in portable, long-duration ECG systems.

2.2.2. Custom- and Reconfigurable-Hardware Implementations

ASIC/VLSI implementations assure high energy efficiency; nonetheless, because of their fixed-function performance and high development complexity, they are less adaptable to the hardware-oriented refinement and validation process adopted in this work. Reconfigurable logic platforms, such as CPLDs and more prominently FPGAs, have, therefore, been investigated for accelerating ECG-signal denoising based on their deterministic timing and support for parallel processing. CPLDs provide low-cost and low-power logic, but they lack the required arithmetic resources to perform multi-level transformations or adaptive filtering, restricting their use to control-oriented tasks. In contrast, FPGAs have a more complex architecture, with added support modules such as integrated digital signal processing (DSP) components, memory, and flexible routing, which are well-suited to pipelined designs to support real-time preprocessing for biomedical applications. Prior FPGA-based works have implemented finite impulse response/infinite impulse response (FIR/IIR) filters, wavelet thresholding, and adaptive strategies, which have achieved great speedups over CPU-based platforms. High-level design tools such as Vivado HLS and Simulink HDL Coder have also been utilized to synthesize MATLAB or C descriptions to FPGA circuits. Patel et al. [28] designed a multi-band FIR filter for denoising ECG signals on an FPGA with a partial-serial structure, ModelSim-simulated. Saha et al. [29] also synthesized a stable elliptic IIR structure on Zynq-7000 using XSG tools. These deployments resulted in adequate reconfigurability and real-time execution but generally at the expense of increased utilization

of resources and poor power efficiency due to auto-generated structure code. Low-level VHDL/Verilog designs provide more control over architecture, allowing dataflow, pipelining, and memory optimization. Sohal et al. [30] synthesized FPGA architectures, merging linear windowing with DWT, and comparing collateral and sequential architectures on Virtex, Kintex, and ZedBoard platforms. Closer to this work, Jenkal et al. [31] presented an FPGA-based ADTF implementation, while Mejhoudi et al. [32] integrated ADTF and DWT modules in a modular design. These studies demonstrated the feasibility of implementing DWT–ADTF-based ECG denoising on an FPGA and established an important basis for hardware-oriented signal processing. However, their primary objective was to demonstrate the functional feasibility of implementing the DWT–ADTF algorithm on an FPGA, rather than to redesign its architecture for the stringent resource, timing, and power constraints of embedded ECG systems. Several architecture-aware aspects were not explicitly addressed, including integrated morphological correction, resource-efficient streaming architecture, and hardware-optimization strategies intended to reduce resource utilization (with logic utilization occasionally reaching up to 60% on a Cyclone IV GX FPGA) and facilitate the integration of subsequent ECG-signal-processing stages.

Given these trade-offs, FPGA-prototyping solutions have been found to be an appropriate platform for experimenting with pipelined data flow, bit-width optimization, and resource-accuracy trade-offs, while maintaining the flexibility required during architectural-level optimization. These aspects make FPGAs well suited for validating hardware-oriented denoising strategies within the limit imposed by wearable, power-efficient, continuous ECG applications, thus facilitating the development of the architecture addressed in this paper.

Despite significant work on ECG-denoising algorithms and FPGA implementations, there remain several challenges yet to be addressed. Most studies are algorithm-centered, aiming at a specific filtering algorithm rather than leveraging complementary strengths via hybridization. Moreover, most works are mainly validated in the simulation or offline-test environment with lesser presentation of robust real-time implementation. Architectures generated through high-level synthesis (HLS) accelerate prototyping but generally do not consider optimizations such as pipelining, memory-aware buffering, and clock-domain synchronization. Furthermore, few designs inherently consider diagnostic integrity by preserving QRS morphology and ensuring signal-length consistency.

These limitations highlight the need for an FPGA architectural solution capable of efficiently deploying an established denoising algorithm while satisfying the stringent resource, latency, and power constraints of embedded ECG monitoring systems. Accordingly, the objective of this work is to present an architecture-aware FPGA design of the standard DWT–ADTF algorithm that preserves denoising quality while minimizing hardware resources and facilitating future integration with subsequent ECG-signal-processing stages.

3. Materials and Methods

This section describes the methodological approach adopted for the hardware design of a real-time ECG-denoising pipeline based on a hybrid DWT–ADTF method. The design is optimized for low-power and low-latency processing on a low-cost FPGA platform, and is fully implemented VHDL. The section addresses the algorithmic basis, the architectural design of each of the primary processing blocks, and the system-level integration and optimization techniques allowing real-time processing.

3.1. System Overview

The proposed system processes raw ECG signals through a three-stage hardware pipeline consisting of (1) multi-level Discrete Wavelet Transform (DWT) decomposition and reconstruction, (2) ADTF-based thresholding and filtering, and (3) correction and output

selection. These processing stages are integrated within a single pipelined system optimized for real-time performance under a unified clock domain. The system is targeted onto an Intel Cyclone V SoC DE10-Nano FPGA board. The overall processing flow, consisting of three interconnected phases, is illustrated in Figure 1.

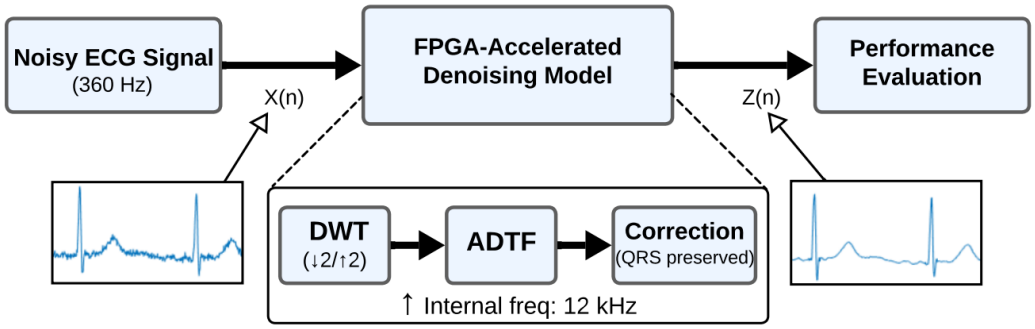

Figure 1. Overview of the proposed pipelined architecture.

3.2. Algorithmic Description

3.2.1. Multi-Level DWT

The Discrete Wavelet Transform (DWT) enables multi-resolution ECG-signal analysis through the decomposition of the signals into approximation and detail sub-bands. A two-level DWT decomposition and reconstruction is employed, with suppression of detail coefficients during reconstruction, as these sub-bands exhibit the highest concentration of noise energy [18]. One-level decomposition for a discrete input signal $x[n]$ is represented in Equation (1).

$$A[n] = \sum_k x[k] \cdot h[2n-k], \quad D[n] = \sum_k x[k] \cdot g[2n-k] \tag{1}$$

where $g[k]$ and $h[k]$ are low-pass and high-pass filter coefficients, respectively. Daubechies-6 (db6) wavelet is selected in this design for its effective time–frequency localization properties.

3.2.2. Adaptive Double-Threshold Filtering (ADTF)

The Adaptive Thresholding Filtering (ADTF) process enhances signal clarity by computing local statistical features (maximum, minimum, mean, and median) within a sliding window of 11 samples. These features are utilized to calculate adaptive thresholds defined in Equations (2) and (3).

$$Ht = mean + [(max - mean) \times \beta\,] \tag{2}$$

$$Lt = mean - [(mean - min) \times \beta\,] \tag{3}$$

where Ht represents the higher threshold of the selected window, Lt corresponds to the lower threshold of the selected window, the maximum and the minimum values of the window are denoted by max and min, respectively, and β is the thresholding coefficient. The coefficient β adjusts the thresholding process and ranges from 0 to 1. Lower values are required for significant noise levels, and higher values are recommended for increased tolerance. In the present implementation, β is fixed to 0.1, following the algorithmic evaluation reported in [18], in which several candidate values were assessed and $\beta = 0.1$ was identified as providing the highest SNR improvement.

3.2.3. Morphological Peak Correction

The correction stage ensures the preservation of diagnostically significant peaks, particularly QRS complexes, which may be compromised during the ADTF process. It

performs selective output fusion by dynamically selecting the most reliable denoised output according to an energy criterion [18].

An energy-based rule is applied over a 5-s window by estimating the maximum amplitude of the signal and using a single adaptive threshold, as given in Equation (4).

$$T = 0.7 * \max(|s(n)|) \tag{4}$$

where $|s(n)|$ is the intermediate signal absolute value after ADTF filtering The final output $y(n)$ is determined as described in Algorithm 1.

Algorithm 1: Threshold-based Output Selection

```
foreach sample s(n) do
    if |s(n)| < T then
        y(n) ← ADTF(s(n));
    else
        y(n) ← DWT(s(n)) in the local region;
```

A representative example of this effect is shown in Figure 2, which illustrates how the correction stage restores and conserves QRS morphology after DWT and ADTF processing.

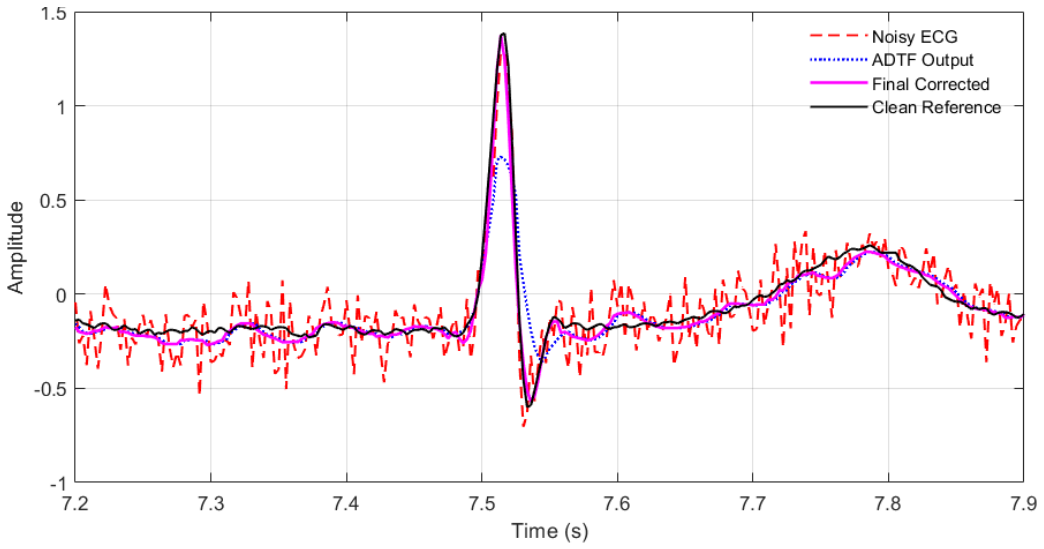

Figure 2. QRS morphology conservation across the successive processing stages.

3.3. Proposed Hardware Architecture

The designed hardware architecture implements DWT–ADTF-based denoising pipeline in a fully synchronous and deeply pipelined RTL design targeting the latency, throughput, and energy constraints of the embedded ECG monitoring. The architecture employs a streaming dataflow model in which each processing module is designed to handle one sample per clock cycle, thus supporting deterministic and hard real-time execution in the design with no overheads due to intermediate buffering. A uniform 16-bit fixed-point processing datapath has been adopted across all modules based on a bit-width optimization analysis for an effective trade-off between quantization accuracy, power consumption, and logic utilization while avoiding the high arithmetic cost of floating-point implementations. These processing blocks are interconnected via pipeline registers. As a result, there is a one-sample-per-clock transfer without the need for any intermediate buffering. This strategy eliminates memory overhead without compromising the precise cycle accurate-timing

alignment. The system is running with a clock of 12 kHz, selected to provide a substantially wide timing margin relative to the 360 Hz sampled input data while preventing unnecessary increases in dynamic power, thereby ensuring stable operation across all pipelined stages. Figure 3 provides an overview of the top-level hardware organization.

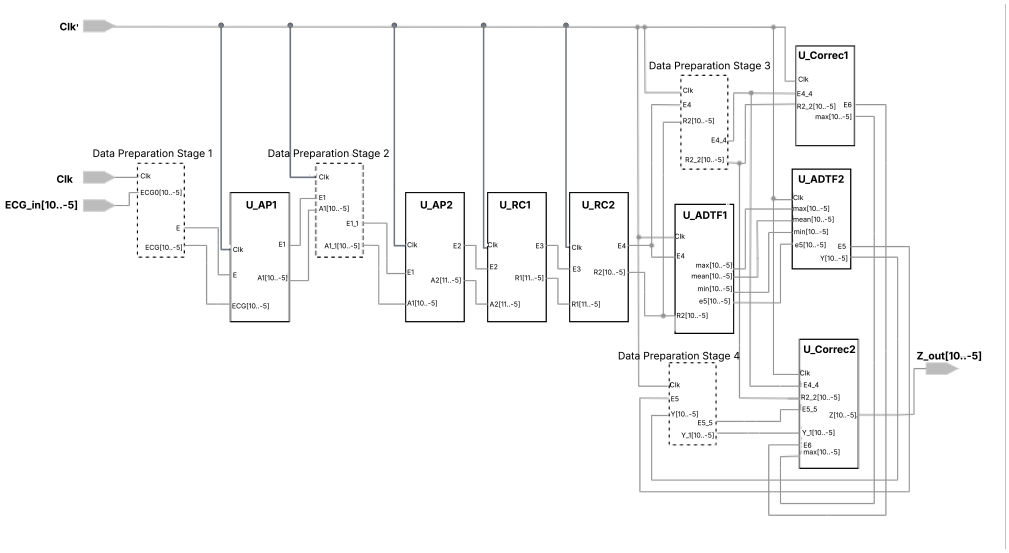

Figure 3. Top-level Register Transfer Level (RTL) schematic of the proposed FPGA-based DWT–ADTF ECG-denoising architecture, illustrating the organization and streaming dataflow of the principal processing modules from the input interface to the denoised ECG output.

3.3.1. DWT Entity

The architecture consists of two decomposition modules, namely U_Ap1 and U_Ap2, and their respective reconstruction modules, namely U_Rc1 and U_Rc2. This architecture, thus, offers an integrated and pipelined multilevel analysis/synthesis process in a streaming mode. Each module employs a 12-sample shift register, corresponding to the filter lengths of the db6 wavelet filters. This minimizes memory addressing overhead in convolution. The architecture supports a throughput rate of one sample per cycle while maintaining a limited combinational depth in each stage to meet timing requirements.

To maintain the length preservation and strict temporal correspondence with downstream stages, a duplication buffer is applied to the DWT input stage. Since the DWT requires 12 initial samples before generating its first valid output sample, a naïve implementation would incur a constant 12-sample delay and shorten the effective output window. The duplication buffer pre-appends replicas of the initial samples in order to reach its steady state without losing temporal synchronization or compromising the continuity required for real-time operation.

All DWT blocks operate in a deeply pipelined configuration, where each wavelet coefficient is mapped to a dedicated multiply accumulate stage. This organization sustains a one-sample-per-cycle rate with deterministic latency, ensuring real-time processing of the 360 Hz input rate with substantial timing margin. The modules operate under a shared clock domain of 12 kHz, thus eliminating the corresponding overhead due to clock-domain crossing, while balanced pipeline depths equalize combinational load and guarantee timing closure. In addition, lightweight synchronization logic ensures that the processing rate remains aligned with the incoming data stream, maintaining stable sample-to-sample continuity throughout the dataflow. The DWT analysis synthesis path is fully pipelined, allowing concurrent stage execution and sustaining a one-sample-per-cycle throughput after pipeline fill, as illustrated in Figure 4; the same steady-state pipelined streaming principle applies uniformly to all processing stages of the proposed architecture.

The hardware realization of the two-level decomposition and reconstruction process is illustrated in Figure 5.

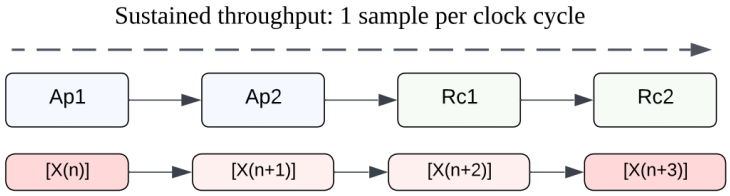

Figure 4. Steady-state streaming pipeline model of the DWT analysis–synthesis chain, illustrating the propagation of successive samples through the pipeline stages after the initial filling phase, enabling a sustained throughput of one-sample-per-clock cycle.

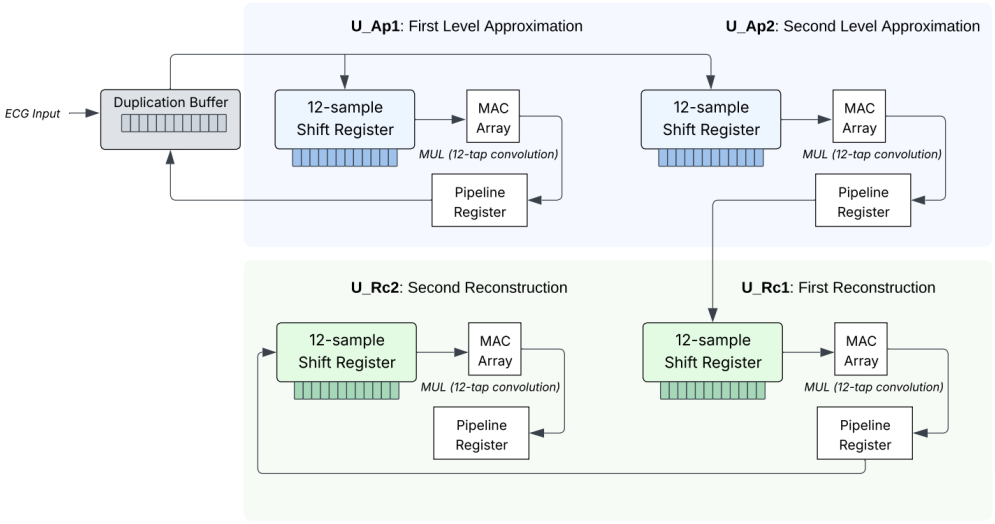

Figure 5. Internal structure of the DWT module.

3.3.2. ADTF Entity

The ADTF block receives the reconstructed signal from the second DWT stage and aims to further enhance the noise-removal operation. Its realization involves two pipelined modules, where the first component (U_adtf1) calculates the local statistical values required for thresholding using an 11-sample shift register structure that corresponds to the ADTF window size, and the second (U_adtf2) executes the double-thresholding operation through a fixed-point comparator chain. This implementation preserves a sample clock cycle, allowing the ADTF stage to integrate seamlessly into the overall streaming architecture.

Figure 6 illustrates the internal processing flow through the ADTF modules, from feature extraction to adaptive thresholding. The control signals synchronize both blocks under the unified 12 kHz clock. ADTF enhances the signal quality but may attenuate peaks located in the QRS regions, motivating the introduction of a correction stage to selectively restore them.

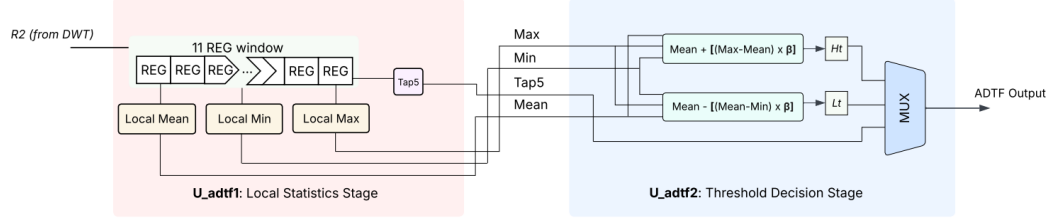

Figure 6. Internal structure of the ADTF subsystem.

3.3.3. Correction Entity: Output Selection and QRS Preservation

The correction block is the last stage in the processing pipeline. The hardware realization consists of two pipelined modules. The first module implements a fixed-point comparator chain to calculate the local maximum within a sliding window and provides the reference amplitude required to preserve the peaks. The second module selects the output using a thresholding-based approach. A segment-level FIFO is implemented for a 5-s window and is required to calculate the local maximum while maintaining the timing coherence between the ADTF output and DWT output during the selection phase. It is implemented as a single-cycle block while preserving the latency in accordance with the latency characteristics of the preceding stages. It maintains timing coherence in the unified clock domain operating at 12 kHz.

The system maintains a minimal memory footprint, being mainly dependent on register driven pipelining. No intermediate frame storage is required, and only a segment-level FIFO appears on the correction path to support maximum estimation and maintain temporal alignment between ADTF and DWT streams. This design choice avoids the use of large on-chip memories and supports the system's low-power operation. A lightweight control interface manages sample propagation across the three processing blocks, ensuring continuous streaming without stalling or backpressure.

Figure 3 presents the Register Transfer Level (RTL) schematic of the proposed ECG-denoising system. For clarity and readability, the illustration focuses on the core processing modules, namely the DWT, ADTF, and correction blocks, while omitting the lower-level interconnections and auxiliary control units of the complete design. This representation provides a concise yet representative overview of the dataflow and modular organization within the implemented architecture. The data-preparation stages in the global-RTL diagram represent functions related to pipeline alignment. These functions involve the duplication buffer at the input stage of the DWT and the segment-level FIFO in the correction path to preserve temporal consistency across the streaming datapath.

3.4. Timing Analysis and Optimization

Effective timing management is critical for real-time FPGA-based signal processing, particularly in biomedical applications that require minimal latency and high reliability. Motivated by earlier research that emphasized clock-domain synchronization and latency optimization [33,34]. Accordingly, the architecture's timing behavior has been established to support deterministic real-time operation despite stringent energy and resource constraints. To this end, the design combines clock-domain standardization, deep pipelining, and latency balancing, along with targeted multi-cycle constraints, in order to guarantee stable processing.

3.4.1. Clock-Domain Standardization and Operating Frequency

All processing stages operate within a common synchronous clock domain of 12 kHz. This frequency was selected in order to maintain a stable, real-time operation with a throughput of one sample per cycle, while preserving a sufficient timing margin for the longest arithmetic paths. This selection reflects the propagation of 16-bit fixed-point arithmetic on the target FPGA fabric and the register-to-register timing constraints associated with the longest combinational paths of the pipelined implementation.

Maintaining a single clock domain avoids cross-domain synchronization logic and allows uniform timing-constraints application over the entire design, which simplifies static timing analysis (STA) and contributes to robust and long-duration operation.

3.4.2. Pipeline Structuring and Critical-Path Mitigation

The system adopts a fine-grained pipeline design in which each stage of processing consists of a simple function of constant complexity. The design ensures a consistent register-to-register timing model throughout the datapath, eliminating the possibility of combinational delay in a single block of processing. The wavelet front-end processing is divided into a series of stages, each of which consists of a single 12-tap processing step. The timing of this wavelet processing is, thus, tightly bound. A similar technique is also employed in processing stages of ADTF and correction, in which a series of distinct levels of processing are employed for computing features, thresholds, and morphological details, thus maintaining a regular timing model. By enforcing a balanced stage granularity across the entire datapath, the architecture maintains a short, stable critical path. This provides a single-cycle deterministic progression of the data stream and enables robust timing closure under stringent embedded-system constraints.

3.4.3. Multi-Cycle Constraints

While the streaming datapath has to conform to strict single-cycle timing to preserve the one-sample-per-clock throughput, not all internal logic drives the per-cycle output. A subset of register-to-register paths related to auxiliary control and state-update logic exhibits more relaxed temporal requirements as their values are sampled in subsequent cycles rather than in the cycle of their computation. These paths, therefore, do not constitute throughput-critical timing constraints. This was implemented by applying multi-cycle timing exceptions to these non-critical paths. Allowing their logic to propagate over multiple clock cycles relaxes the combinational timing pressure, increasing slack margins while avoiding unnecessary operator duplication, all while preserving the one-sample-per-cycle throughput and the fixed-latency behavior of the streaming pipeline.

Static Timing Analysis verified that the multi-cycle paths meet their extended timing requirements and remain properly synchronized with the single-cycle sections of the DWT, ADTF, and correction modules.

4. Experimental Results and Hardware Validation

4.1. Experimental Setup

To evaluate the denoising quality of the developed FPGA-based system, various ECG records from the MIT-BIH Arrhythmia Database (360 Hz) were used to cover a variety of waveform morphologies, ensuring realistic and representative test conditions.

Each record was artificially corrupted to simulate real acquisition conditions. Table 2 lists the noise types along with the respective signal-to-noise ratio (SNR) values. Noisy ECG signals were forwarded directly to the denoising pipeline in order to make the evaluation effectively reflect the inherent ability of the proposed system for real-time embedded applications.

Table 2. Noise scenarios and corresponding SNR levels for ECG-denoising evaluation.

Noise Type	Source	SNR Levels (dB)
White Gaussian Noise (WGN)	AWGN function	0, 5, 10, 15
Baseline Wander (BW) Muscle Artifact (MA) Electrode Motion (EM) Mixed Noise (MN) (Combination of BW, MA, and EM)	MIT-BIH NSTDB	9, 12, 20

4.2. Evaluation Metrics

Performance was quantified using Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Percentage Root Mean Square Difference (PRD), and Signal-to-noise ratio improvement (SNR_{imp}). MSE and RMSE measure absolute reconstruction error, PRD expresses relative distortion, and SNR_{imp} captures noise suppression between the input and the output.

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(x_i - \hat{x}_i)^2 \tag{5}$$

$$RMSE = \sqrt{\text{MSE}} \tag{6}$$

$$PRD = 100 \times \frac{\sqrt{\sum_{i=1}^{N}(x_i - \hat{x}_i)^2}}{\sqrt{\sum_{i=1}^{N} x_i^2}} \tag{7}$$

$$SNR_{\text{in}} = 10 \cdot \log_{10}\left(\frac{\sum_{i=1}^{N}[x_i]^2}{\sum_{i=1}^{N}[x_i - y_i]^2}\right) \tag{8}$$

$$SNR_{\text{out}} = 10 \cdot \log_{10}\left(\frac{\sum_{i=1}^{N}[x_i]^2}{\sum_{i=1}^{N}[x_i - \hat{x}_i]^2}\right) \tag{9}$$

$$SNR_{\text{imp}} = SNR_{\text{out}} - SNR_{\text{in}} \tag{10}$$

where x_i is the clean ECG sample, $\hat{x}_i$ is the denoised output sample, y_i is the noisy input sample, and N is the number of samples.

4.3. Quantitative Results and Analysis

The denoising performance of the proposed DWT–ADTF design was evaluated on multiple ECG recordings of the MIT-BIH Arrhythmia Database under synthetic (AWGN) and realistic (NSTDB) noisy conditions. Table 3 summarizes the average values of the quantitative metrics for studied noise types and SNR levels. The proposed hardware architecture achieves robust performance for the suppression of the Gaussian and physiological artifacts (Baseline Wander, Muscle Activity, Electrode Motion, and Mixed Noise). The obtained results indicate that the denoising performance varies with the characteristics of the considered noise source and the input SNR level, while consistently preserving the clinically relevant ECG morphology across all evaluated scenarios. The stability of the MSE and RMSE across all test scenarios further validates the FPGA-based implementation for real-time biomedical signal processing.

The Figure 7 demonstrates the quantitative performance of the proposed architecture with different levels of AWGN. The FPGA implementation of this architecture is stable with different levels of noise, thus demonstrating its high level of generalization with regard to denoising. This robustness guarantees that the structure offers predictable behavior under different SNR environments, a critically important requirement in real-time clinical ECG recording.

A comparative analysis of the proposed implementation on the FPGA and its software reference model is illustrated in Figure 8. The two realizations exhibit nearly identical SNR_{imp} and PRD curves across all levels of noise input, thereby validating their algorithmic equivalence. Minor numerical differences are attributed to the inherent differences between fixed- and floating-point data types with regards to scaling, quantization, and saturation of arithmetic values. These differences are within acceptable tolerance ranges and reflect the expected behavior of hardware quantization rather than a deviation from the reference algorithm.

Overall, the quantitative findings demonstrate that the proposed FPGA-based design for DWT–ADTF provides a high level of denoising performance along with a low level of waveform distortion, thus achieving an optimum balance between accuracy, processing throughput, and hardware cost. Additionally, the level of robustness demonstrated for handling both synthetic and physiological noise sources provides a strong indication of the potential for the proposed design for deployment in low-power, real-time ECG-signal-monitoring systems.

Table 3. Average quantitative performance across noise conditions using the MIT-BIH Arrhythmia Database.

Noise Type	SNR_{in} (dB)	SNR_{out} (dB)	MSE	RMSE	PRD (%)
AWGN	0	9.8870	0.0201	0.1347	32.0652
	5	14.3643	0.0073	0.0809	19.1606
	10	18.1655	0.0031	0.0522	12.3984
	15	20.8665	0.0018	0.0388	9.2879
AWGN Average	–	15.8208	0.0081	0.0767	18.2280
Baseline Wander (BW)	9	8.8959	0.0676	0.2507	35.9153
	12	11.7874	0.0348	0.1798	25.7459
	20	18.7180	0.0070	0.0806	11.6405
BW Average	–	13.1338	0.03647	0.1704	24.4339
Muscle Artifact (MA)	9	11.7527	0.0349	0.1802	25.8541
	12	14.5919	0.0181	0.1005	18.6497
	20	21.2233	0.0041	0.0609	8.7819
			0.0190		
MA Average	–	15.8560	0.1139		
			17.7619		
Electrode Motion (EM)	9	9.2192	0.0622	0.2410	34.6105
	12	12.1078	0.0321	0.1729	24.8282
	20	18.9811	0.0072	0.0820	11.8850
EM Average	–	13.4360	0.0338	0.1653	23.7746
Mixed Noise (MN)	9	9.1088	0.0642	0.2445	35.0470
	12	11.9013	0.0336	0.1769	25.4166
	20	18.5025	0.0072	0.0821	11.9781
MN Average	–	13.1709	0.0350	0.1678	24.1472
Overall Average (Real Noise Conditions)	–	13.8992	0.0311	0.1544	22.5294

4.4. Qualitative Results and Visual Analysis

Beyond numerical analysis, a qualitative examination was conducted to visually assess the denoising behavior and morphological fidelity of the proposed FPGA-based DWT–ADTF architecture. Representative ECG signals from the MIT-BIH Arrhythmia Database were compared under synthetic (AWGN) and real-noise conditions. Figure 9 illustrates the system response for record 117 (5 dB AWGN) and Figure 10 illustrates the denoising process through each level of the processing chain—from the noisy input signal, to intermediate results after the DWT and ADTF modules, and finally to the corrected output generated by the decision module.

This graphical comparison illustrates the filtering action of the architecture to progressively eliminate additive and physiological noise and maintain the features of the diagnostic content of the ECG signal. Hardware-level validation was conducted using the SignalTap™ Logic Analyzer in order to verify real-time execution and functional correctness, as discussed in Section 4.5.

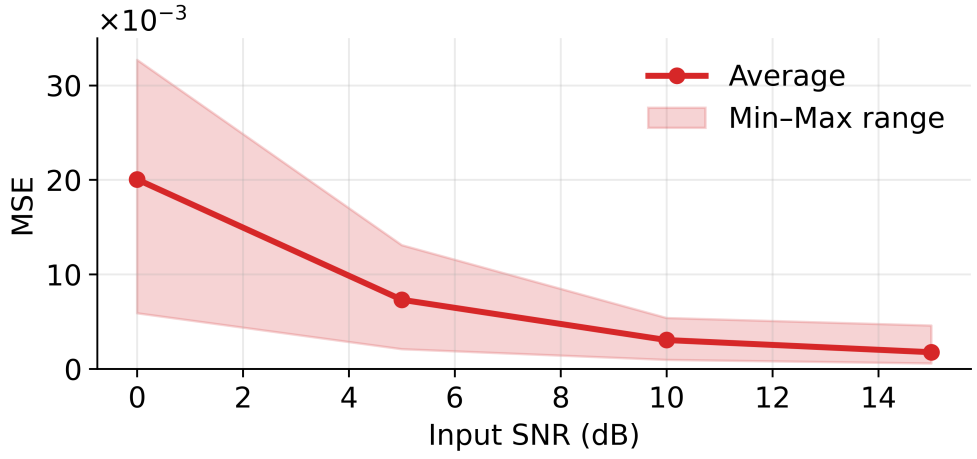

(**a**)

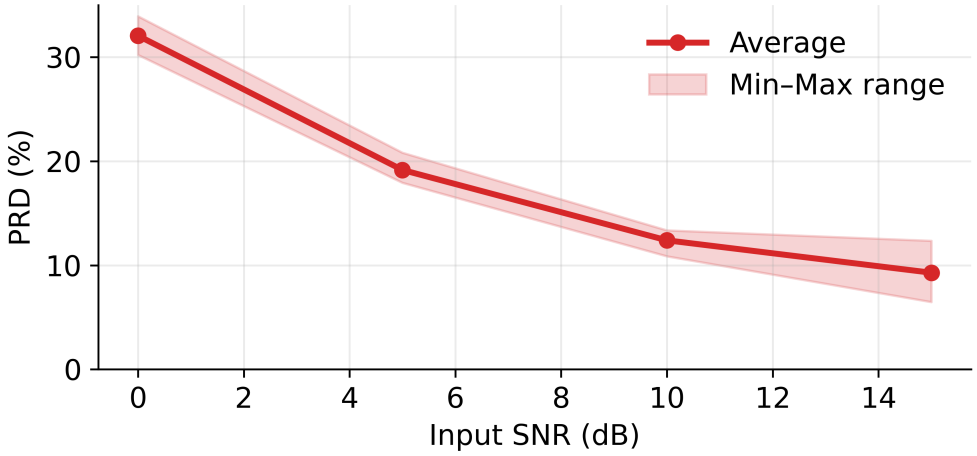

(**b**)

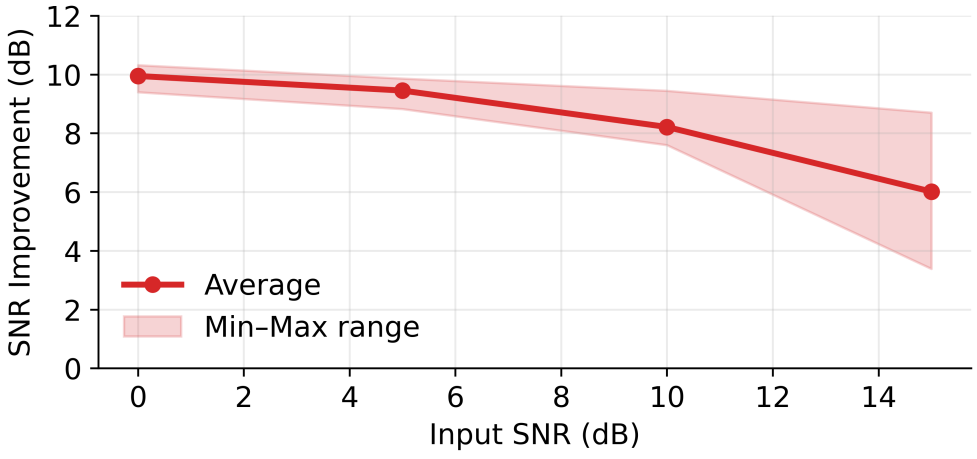

(**c**)

Figure 7. Performance evolution of the proposed architecture under AWGN conditions from 0 to 15 dB. (**a**) MSE variation; (**b**) PRD variation (%); and (**c**) SNR_{imp} variation (dB).

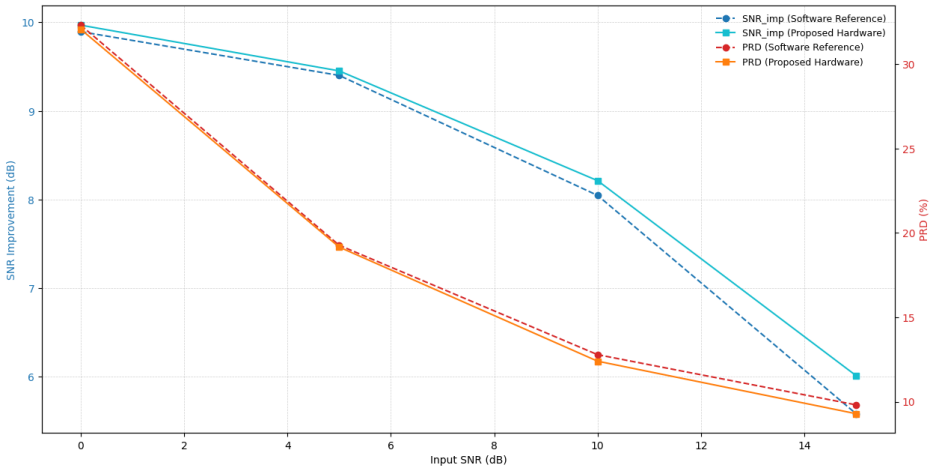

Figure 8. Comparison of SNR_{imp} and PRD between the software reference model and the proposed hardware design under AWGN (0 to 15 dB)).

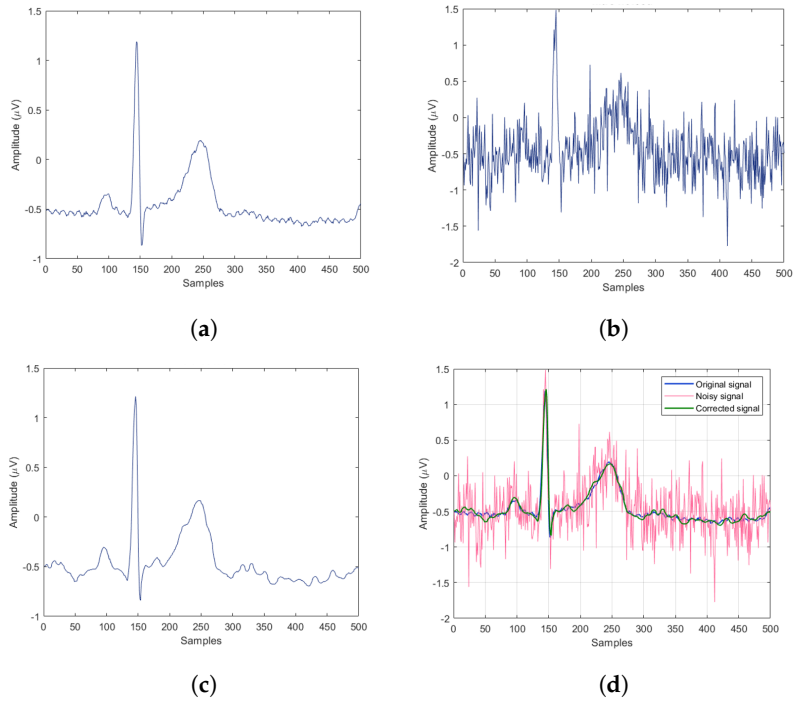

Figure 9. ECG signal denoising for the MIT-BIH record 117 with 5 dB added wgn noise. (**a**) Original ECG signal; (**b**) noisy ECG signal input; (**c**) denoised ECG signal output; and (**d**) compared signals.

4.5. SignalTap-Based Hardware Validation

The real-time behavior and functional correctness of the proposed DWT–ADTF ECG-denoising system were validated on the Cyclone V SoC FPGA (5CSEBA6U23I7) using the Intel SignalTap™ Logic Analyzer. This on-chip diagnostic tool provides real-time observation of waveforms and timing verifications, providing direct visibility into internal signal transitions without any interrupting system execution. Using SignalTap, dataflow synchronization and output transitions were captured through the main pipeline stages, confirming that each module operates correctly under real operating scenarios.

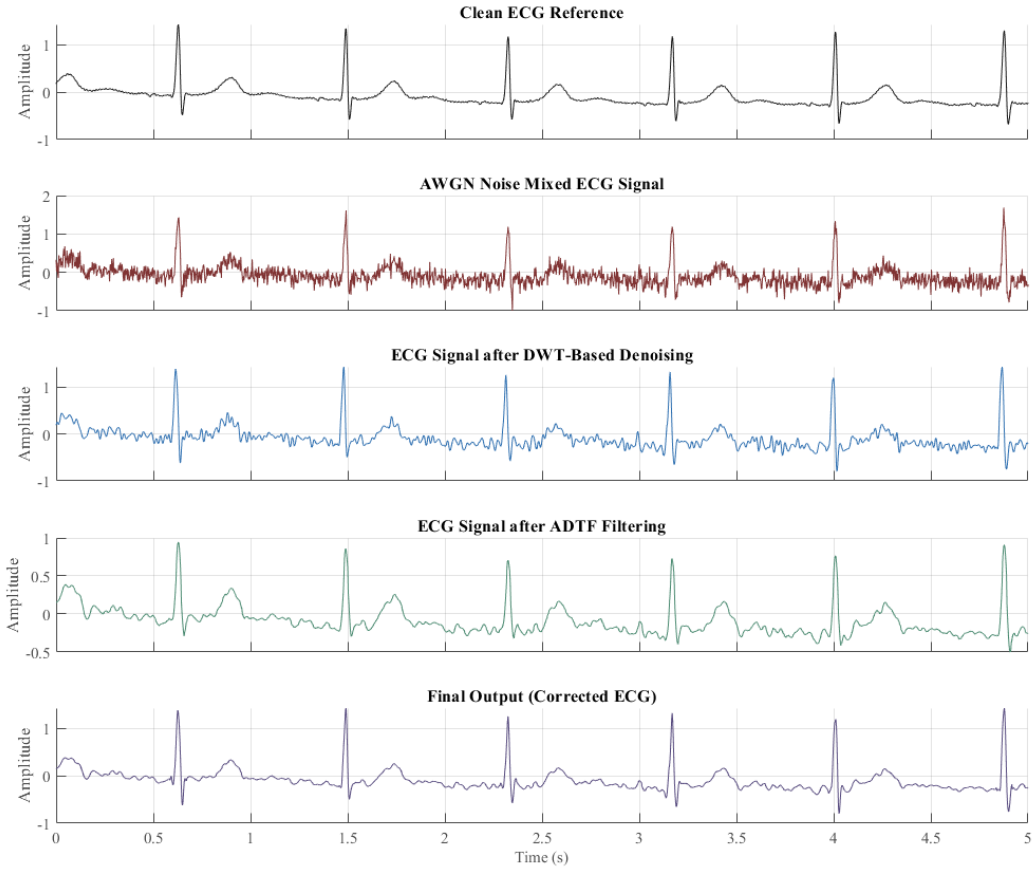

Figure 10. Application of the proposed method for ECG signal denoising under a 5 dB AWGN noise level affecting the MIT-BIH signal 103.

Figure 11 illustrates the experimental setup based on the DE10-Nano development board. The two clock domains in FPGA implementation are a low-frequency domain (360 Hz) for the MIT-BIH ECG sampling rate and a high-frequency domain (12 kHz) dedicated to the denoising pipeline. The synchronization interface between the two domains was verified through captured enable pulses and stable inter-module synchronization signals.

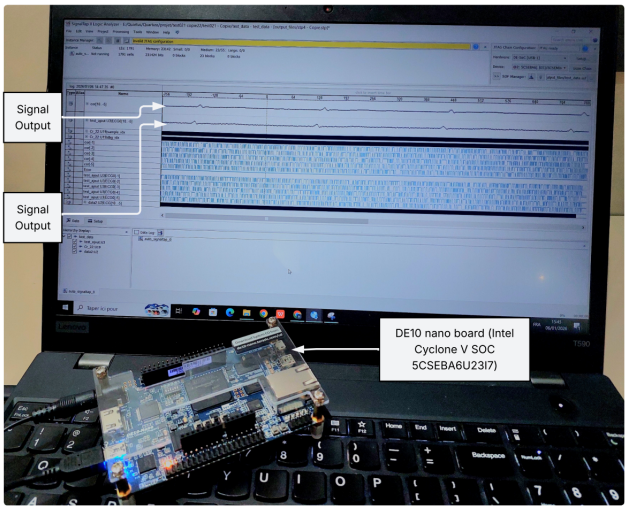

Figure 11. Experimental FPGA setup and on-chip monitoring environment using SignalTap™.

Representative SignalTap captures in Figure 12 demonstrate the temporal behavior of the architecture in real-time operation. The captured traces confirm continuous sample-by-sample processing, the enable signal remains constant, and latency becomes predictable once the pipeline is full. The lack of stalls and discontinuities in the output stream confirms that the alignment buffers and synchronization mechanisms are functioning correctly.

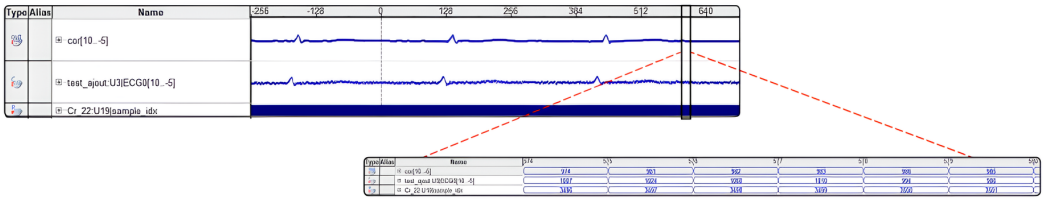

Figure 12. SignalTap on-chip capture for the MIT-BIH record 103 under 5 dB AWGN, showing synchronized data propagation and denoised-output stability.

Experimental results validate the FPGA's reliable operation under constant real-time data input, ensuring deterministic timing and stable throughput without off-chip latency or frame buffering. The SignalTap tool was utilized for verifying functional accuracy, timing, and streaming of the proposed architectural design under real operational scenarios. The hardware-verification process validates the designer's practical readiness for real-time embedded biomedical observation, considering the constraints of the system's timing and resources availability.

4.6. Resource Utilization and Power Analysis

The designed DWT–ADTF ECG-denoising model has been synthesized on various Intel FPGA series to evaluate the efficiency of the design. Table 4, providing the synthesis results of the proposed design, shows the efficiency of the proposed design across all the tested devices.

For the Cyclone V SoC FPGA device, the design under consideration requires only 5% of the logic elements, 6% of the registers, 1% of on-chip memory bits, and 44% of the DSP blocks in the SoC FPGA device. The total power dissipation of the design within the chip is found to be approximately 426 mW. This moderate level of power supports the design's compatibility with the requirements of the embedded system in the biomedical systems, for which energy efficiency and timing determinism is required.

The reported power values were obtained using the Intel Quartus Prime Power Analyzer following synthesis, fitting, and timing analysis. They, therefore, correspond to post-implementation power estimates. The estimated power consists of static and dynamic components. The relatively high static power is primarily associated with the intrinsic leakage characteristics of the Cyclone FPGA device, whereas the dynamic power is determined by the switching activity of the implemented architecture and is, therefore, directly influenced by the adopted hardware optimizations.

The compact resource footprint provides the flexibility to accommodate possible architectural extensions, such as adaptive filtering or on-chip feature extraction, without violating the real-time constraints or the device's capacity.

Figure 13 is a complement to Table 4 and provides a graphical representation of the hardware-resource utilization of the design across various FPGA-device families. The distribution clearly reflects the architectural balance achieved in the design, where the logical and memory requirements are low, while the DSP resources are optimally exploited to sustain real-time processing performance.

Table 4. FPGA-resource utilization and power dissipation across device families.

Parameter	Cyclone V	Cyclone IV GX	Cyclone IV E	Cyclone III LS	Arria II GX	MAX 10
Registers	1945 (6%)	1644 (3%)	1695 (3%)	2081 (3%)	1894	1695
Logic Elements (LEs)	2201 (5%)	5967 (20%)	12,353 (22%)	6563 (9%)	–	12,391 (31%)
Memory Bits	58,824 (1%)	57,952 (5%)	58,144 (2%)	60,482 (2%)	57,600 (2%)	58,144 (5%)
DSP Blocks	49 (44%)	–	–	–	134 (58%)	–
Static Power (mW)	412.63	118.93	80.31	99.67	319.48	85.19
Dynamic Power (mW)	3.14	8.27	3.70	6.83	3.28	4.86
I/O Power (mW)	10.52	22.69	51.12	37.47	13.77	7.26
Total Power (mW)	**426.01**	149.89	135.12	143.97	336.53	97.31

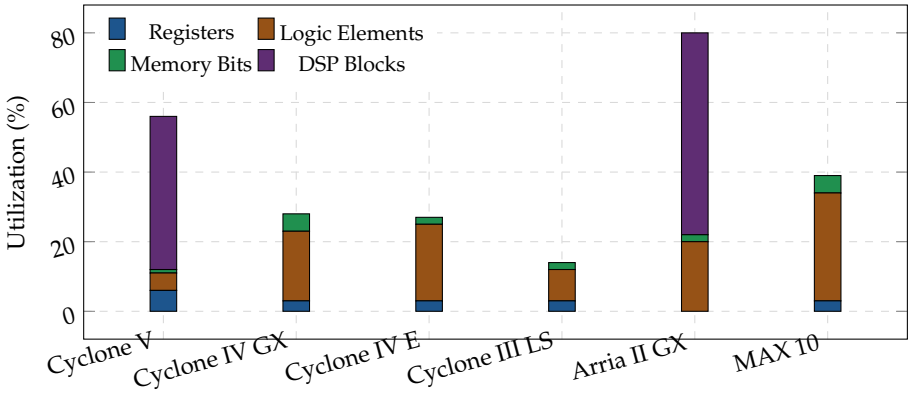

Figure 13. Graphical representation of the resources utilization in the proposed FPGA-based DWT–ADTF architecture across multiple Intel device families.

4.7. Comparative Performance and Discussion

In order to further evaluate the efficiency of the proposed architecture, Table 5 contrasts its hardware footprint as well as total power dissipation with representative state-of-the-art FPGA-based architectures for ECG denoising. Although numerical comparison should be interpreted cautiously, owing to variations in FPGA families, technology nodes, and synthesis settings, the table provides an informative insight into the compromises achieved by each system. Nevertheless, the comparison provides an overall architectural perspective on representative FPGA-based ECG-denoising implementations, highlighting the different trade-offs achieved in terms of hardware resources, performance, and power consumption.

Table 5. Comparison of total power consumption and resource usage with previous FPGA-based ECG-denoising architectures.

Denoising Technique	Reference	Device	TPC (W)	Resources Use	
				TR	TLE
DWT–ADTF(proposed)	–	Cyclone V	0.428	1945	5%
IIR elliptic filters	[29]	Xilinx Zynq-7000	0.321	1684	3.87%
FIR Filter	[35]	Xilinx Spartan-3E	0.186	–	–
FIR Filter (Multilevel Cascaded)	[36]	Xilinx viretex-6	3.926	8146	4%
Wavelet (Daubechies)	[30]	Zedboard (XC7Z007S)	0.312	1337	–
		Kintex Board (XC7K70T)	0.729	1337	–

TPC: total power consumption, TR: total registers, and TLE: total logic elements.

The design demonstrates an overall power dissipation of around 426 mW on the Cyclone V platform, which is slightly higher than that of simpler single-stage filter designs, but remains within the predicted range for multi-stage pipelined structures. This overhead can be attributed mainly to the hybridization of the DWT and ADTF cores along with the

synchronization logic and deeper pipelining strategy employed in the implementation to achieve real-time throughput. Part of the dissipation can be attributed to the static power consumption for the Cyclone V family's inherent static-power profile.

Thus, the architecture is able to provide low resource consumption, requiring only 5% LEs and 6% registers, while providing high accuracy in terms of denoising and morphological preservation, as shown in the quantitative evaluation. Compared with conventional FIR or IIR filters, which achieve reduced power at the expense of lower fidelity, the proposed architecture ensures a more balanced trade-off between precision, latency, and energy efficiency.

This trade-off validates the design's efficiency and scalability: it leverages parallel DSP utilization and on-chip synchronization with minimal logic or memory overhead, and, thus, represents a promising candidate for real-time, low-power biomedical systems.

Overall, the proposed architecture demonstrates that high-performance ECG denoising can be achieved on a mid-range FPGA device with moderate energy and resource consumption. This makes it a strong, scalable, and energy-efficient solution to be integrated in real-time biomedical and wearable ECG monitoring devices.

Figure 14 presents a comparative performance analysis of the proposed hardware-based DWT–ADTF denoising architecture compared to standard ECG-denoising algorithms, including empirical mode decomposition (EMD), finite impulse response (FIR) and infinite impulse response (IIR) filters, the extended Kalman filter (EKF), and the software-based DWT–ADTF approach [18,25,37,38]. The software implementation of the DWT–ADTF framework is used as the direct algorithmic reference, while FIR, IIR, EMD, and EKF were selected as representative ECG-denoising methods as they combine established denoising performance with relatively high computational complexity. This comparison, therefore, verifies that the proposed FPGA implementation preserves the denoising capability of the underlying DWT–ADTF framework despite the architectural optimizations introduced for efficient hardware realization. The evaluation was performed under various additive white Gaussian noise (AWGN) levels of 0 dB, 5 dB, 10 dB, and 15 dB.

The result shows that the proposed hardware architecture is consistently capable of achieving the optimal signal-to-noise ratio (SNR) improvement across all noise levels while maintaining the lowest percentage root mean square difference (PRD). This shows that the design ensures the morphological integrity of the ECG signal while also improving the robustness of the noise-removal process against both algorithmic and hardware alternatives. This improvement can be attributed to the hybridizing of the wavelet-based multi-resolution analysis and the thresholding for efficient suppression of broadband and transient noise components. Overall, the proposed framework offers the best compromise between the accuracy of the noise-removal process and the efficiency of the implementation.

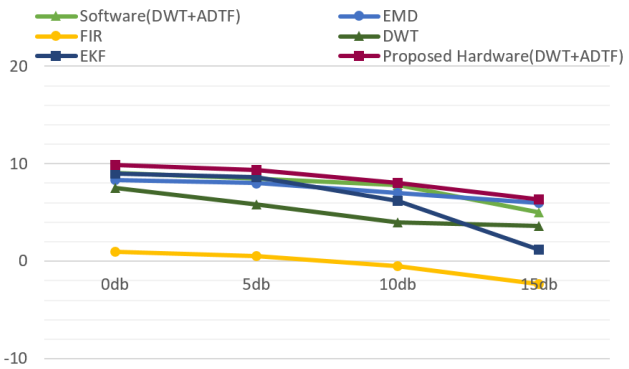

(**a**) SNR improvement across different AWGN levels.

Figure 14. *Cont.*

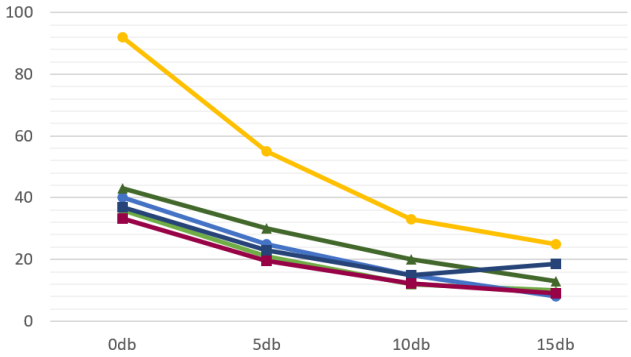

(**b**) PRD performance under identical noise conditions.

Figure 14. Performance comparison of the proposed FPGA-based DWT-ADTF denoising architecture with existing ECG-denoising techniques under varying levels of additive white Gaussian noise (AWGN). (**a**) Reports the signal to noise ratio (SNR) improvement, while (**b**) depicts the percentage root mean square difference (PRD).

5. Conclusions and Challenges

This work presented an efficient FPGA architecture of the hybrid DWT–ADTF ECG-denoising framework for deterministic real-time operation. The proposed architecture implements the established DWT–ADTF framework through a fully pipelined RTL architecture, dual-clock synchronization, hardware integration of the morphological correction stage, and memory-efficient buffering to ensure deterministic processing without compromising clinically relevant ECG morphology. The architecture was successfully synthesized on a Cyclone V SoC FPGA, requiring only 5% logic elements, 6% registers, and a moderate power consumption of 426 mW.

The real-time validation of the design using the SignalTap™ Logic Analyzer also ensured the functional correctness, the stability of the clock-domain cross-synchronization, and the deterministic nature of the pipelining of the dataflow. Furthermore, the comparative analysis of the existing FPGA-based methods for noise filtering also highlighted the best trade-off achieved by the proposed method.

The proposed architecture offers a robust foundation for the design of embedded biomedical monitoring systems, especially for low-power or wearable devices where precision of processing and power efficiency are critical. However, despite these results, there remain several inherent challenges associated with the implementation of hardware-based ECG denoising. First, achieving full adaptivity in the non-stationary noise condition may require parameter adjustments. This may lead to control-complexity issues and possible timing non-determinism if not properly managed. Additionally, applying such a validated architecture to ultra low-power or application-specific systems requires more research into technology-dependent issues such as voltage scaling, memory hierarchy, and long-term stability under sustained execution.

The above mentioned challenges need to be overcome in order to further achieve fully autonomous systems for long-duration ECG monitoring while retaining the balance between performance, power consumption, and architectural determinism achieved in this research work.

Author Contributions: Conceptualization, H.E.; methodology, H.E. and W.J.; software, H.E.; validation, H.E., W.J., M.L. and A.D.; formal analysis, H.E.; investigation, H.E., A.D. and R.L.; resources, H.E., W.J., M.L. and R.L.; writing—original draft, H.E.; writing—review and editing, H.E., W.J. and A.D.; visualization, A.D. and R.L.; supervision, W.J. and M.L.; and project administration, W.J. and M.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the doctoral scholarship (No. 17UIZ2023), from the National Centre for Scientific and Technical Research (CNRST), Morocco

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: We would like to thank the National Centre for Scientific and Technical Research (CNRST), Morocco, for the doctoral scholarship support received under scholarship number 17UIZ2023. The authors used AI for language editing, grammar correction, and improvement of the clarity and readability of the manuscript. All scientific content, results, analyses, and conclusions were developed and validated by the authors, who take full responsibility for the final content of the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bhaltadak, V.; Ghewade, B.; Yelne, S. A Comprehensive Review on Advancements in Wearable Technologies: Revolutionizing Cardiovascular Medicine. *Cureus* **2024**, *16*, e61312. [CrossRef] [PubMed]
2. Brüser, C.; Antink, C.H.; Wartzek, T.; Walter, M.; Leonhardt, S. Ambient and unobtrusive cardiorespiratory monitoring techniques. *IEEE Rev. Biomed. Eng.* **2015**, *8*, 30–43. [CrossRef] [PubMed]
3. Mporas, I.; Tsirka, V.; Zacharaki, E.I.; Koutroumanidis, M.; Richardson, M.; Megalooikonomou, V. Seizure detection using EEG and ECG signals for computer-based monitoring, analysis and management of epileptic patients. *Expert Syst. Appl.* **2015**, *42*, 3227–3233. [CrossRef]
4. Guan, Y.X.; An, Y.; Guo, F.Y.; Pan, W.B.; Wang, J.X. Intelligent Electrocardiogram Analysis in Medicine: Data, Methods, and Applications. *Chin. Med. Sci. J.* **2023**, *38*, 38–48. [CrossRef] [PubMed]
5. Chatterjee, S.; Thakur, R.S.; Yadav, R.N.; Gupta, L.; Raghuvanshi, D.K. Review of noise removal techniques in ECG signals. *IET Signal Process.* **2020**, *14*, 569–590. [CrossRef]
6. Jia, Y.; Pei, H.; Liang, J.; Zhou, Y.; Yang, Y.; Cui, Y.; Xiang, M. Preprocessing and Denoising Techniques for Electrocardiography and Magnetocardiography: A Review. *Bioengineering* **2024**, *11*, 1109. [CrossRef] [PubMed]
7. Satija, U.; Ramkumar, B.; Manikandan, M.S. A review of signal processing techniques for electrocardiogram signal quality assessment. *IEEE Rev. Biomed. Eng.* **2018**, *11*, 36–52. [CrossRef] [PubMed]
8. An, X.; Stylios, G.K. Comparison of motion artefact reduction methods and the implementation of adaptive motion artefact reduction in wearable electrocardiogram monitoring. *Sensors* **2020**, *20*, 1468. [CrossRef] [PubMed]
9. Saxena, C.; Sharma, A.; Srivastav, R.; Gupta, H. Denoising of ECG signals using FIR & IIR filter: A performance analysis. *Int. J. Eng. Technol.* **2018**, *7*, 1–5. [CrossRef]
10. Eminaga, Y.; Coskun, A.; Kale, I. IIR wavelet filter banks for ECG signal denoising. In *Proceedings of the 2018 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA)*; IEEE: New York, NY, USA, 2018; pp. 130–133.
11. Venkatesan, C.; Karthigaikumar, P.; Varatharajan, R. A novel LMS algorithm for ECG signal preprocessing and KNN classifier based abnormality detection. *Multimed. Tools Appl.* **2018**, *77*, 10365–10374. [CrossRef]
12. Mugdha, A.C.; Rawnaque, F.S.; Ahmed, M.U. A study of recursive least squares (RLS) adaptive filter algorithm in noise removal from ECG signals. In Proceedings of the 2015 International Conference on Informatics, Electronics & Vision (ICIEV), Fukuoka, Japan, 15–18 June 2015; pp. 1–6. [CrossRef]
13. Singh, P.; Pradhan, G. Denoising of ECG signal by non-local estimation of approximation coefficients in DWT. *Biocybern. Biomed. Eng.* **2017**, *37*, 599–610. [CrossRef]
14. Das, M.; Sahana, B.C. A Deep-learning-based auto Encoder-Decoder model for denoising electrocardiogram signals. *IETE J. Res.* **2025**, *71*, 326–340.
15. Singh, P.; Sharma, A. Attention-based convolutional denoising autoencoder for two-lead ECG denoising and arrhythmia classification. *IEEE Trans. Instrum. Meas.* **2022**, *71*, 1–10. [CrossRef]
16. Tripathi, P.M.; Kumar, A.; Komaragiri, R.; Kumar, M. A review on computational methods for denoising and detecting ECG signals to detect cardiovascular diseases. *Arch. Comput. Methods Eng.* **2022**, *29*, 1875–1914.
17. Chandrakasan, A.P.; Verma, N.; Daly, D.C. Ultralow-power electronics for biomedical applications. *Annu. Rev. Biomed. Eng.* **2008**, *10*, 247–274. [CrossRef] [PubMed]
18. Jenkal, W.; Latif, R.; Toumanari, A.; Dliou, A.; El B'charri, O.; Maoulainine, F.M. An efficient algorithm of ECG signal denoising using the adaptive dual threshold filter and the discrete wavelet transform. *Biocybern. Biomed. Eng.* **2016**, *36*, 499–508. [CrossRef]
19. Jenkal, W.; Latif, R.; Laaboubi, M. ECG Signal Denoising Using an Improved Hybrid DWT-ADTF Approach. *Cardiovasc. Eng. Technol.* **2024**, *15*, 77–94. [PubMed]

20. Choudhry, M.S.; Puri, A.; Kapoor, R. Removal of baseline wander from ECG signal using cascaded empirical mode decomposition and morphological functions. In *Proceedings of the 2016 3rd International Conference on Signal Processing and Integrated Networks (SPIN)*; IEEE: New York, NY, USA, 2016; pp. 769–774.
21. Han, G.; Lin, B.; Xu, Z. Electrocardiogram signal denoising based on empirical mode decomposition technique: An overview. *J. Instrum.* **2017**, *12*, P03010. [CrossRef]
22. Wang, T.; Zhang, M.; Yu, Q.; Zhang, H. Comparing the applications of EMD and EEMD on time–frequency analysis of seismic signal. *J. Appl. Geophys.* **2012**, *83*, 29–34. [CrossRef]
23. Betancourt, N.; Flores-Calero, M.; Almeida, C. ECG Denoising by using FIR and IIR Filtering Techniques: An Experimental Study. In Proceedings of the 2019 11th International Conference on Bioinformatics and Biomedical Technology, New York, NY, USA, 29–31 May 2019; ICBBT '19; p. 111–117. [CrossRef]
24. Wang, F.; Wang, Q.; Liu, F.; Chen, J.; Fu, L.; Zhao, F. Improved NLMS-based adaptive denoising method for ECG signals. *Technol. Health Care* **2021**, *29*, 305–316. [CrossRef] [PubMed]
25. Rakshit, M.; Das, S. An efficient ECG denoising methodology using empirical mode decomposition and adaptive switching mean filter. *Biomed. Signal Process. Control* **2018**, *40*, 140–148. [CrossRef]
26. Shweta, J.; Varun, B.; Anil, K. Effective de-noising of ECG by optimised adaptive thresholding on noisy modes. IET ence. *IET Sci. Meas. Technol.* **2018**, *12*, 640–644. [CrossRef]
27. Mejhoudi, S.; Latif, R.; Saddik, A.; Jenkal, W.; El Ouardi, A. Speeding up an adaptive filter based ECG signal pre-processing on embedded architectures. *Int. J. Adv. Comput. Sci. Appl. (IJACSA)* **2021**, *12*, 360–369. [CrossRef]
28. Patel, V.; Shah, A. Denoising of Electrocardiogram Signal using Multiband Filter and its Implementation on FPGA. *Serbian J. Electr. Eng.* **2022**, *19*, 115–128. [CrossRef]
29. Saha, S.; Mandal, S.B. FPGA implementation of IIR elliptic filters for de-noising ECG signal. *Biomed. Signal Process. Control* **2024**, *96*, 106544. [CrossRef]
30. Sohal, H.; Jain, S. FPGA implementation of collateral and sequence pre-processing modules for low power ECG denoising module. *Inform. Med. Unlocked* **2022**, *28*, 100838. [CrossRef]
31. Jenkal, W.; Latif, R.; Toumanari, A.; Elouardi, A.; Hatim, A.; El'Bcharri, O. Real-time hardware architecture of the adaptive dual threshold filter based ECG signal denoising. *J. Theor. Appl. Inf. Technol.* **2018**, *96*, 4649–4659.
32. Mejhoudi, S.; Jenkal, R.L.W.; Saddik, A.; Satie, A.E.O. Hardware Architecture for Adaptive Dual Threshold Filter and Discrete Wavelet Transform based ECG Signal Denoising. *Int. J. Adv. Comput. Sci. Appl.* **2021**, *12*, 0121106. . [CrossRef]
33. Methul, S.; Shedge, A.; Bartakke, P. CardioSync: Real-Time ECG Monitoring System with FPGA-Based Signal Processing and Analysis. In Proceedings of the 2024 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT), Bangalore, India, 12–14 July 2024; pp. 1–4. [CrossRef]
34. Desai, M.P.; Caffarena, G.; Jevtic, R.; Márquez, D.G.; Otero, A. A Low-Latency, Low-Power FPGA Implementation of ECG Signal Characterization Using Hermite Polynomials. *Electronics* **2021**, *10*, 2324. [CrossRef]
35. Janwadkar, S.; Dhavse, R. Power and area efficient FIR filter architecture in digital encephalography systems. *e-Prime Adv. Electr. Eng. Electron. Energy* **2023**, *4*, 100148. [CrossRef]
36. Bhaskar, P.C.; Uplane, M.D. FPGA based digital FIR multilevel filtering for ECG denoising. In Proceedings of the 2015 International Conference on Information Processing (ICIP), Pune, India, 16–19 December 2015; pp. 733–738.
37. Kabir, M.A.; Shahnaz, C. Denoising of ECG signals based on noise reduction algorithms in EMD and wavelet domains. *Biomed. Signal Process. Control* **2012**, *7*, 481–489. [CrossRef]
38. Christov, I.I.; Daskalov, I.K. Filtering of electromyogram artifacts from the electrocardiogram. *Med. Eng. Phys.* **1999**, *21*, 731–736. [CrossRef] [PubMed]