Timing and Signal-Fidelity Trade-offs in a Low-Cost Wireless ECG Edge Pipeline: A Measurement Study

Tomasz Kot
Faculty of Computer Science and Mathematics
Cracow University of Technology
Krakow, Poland
tomasz.kot13@gmail.com

Filip Krużel
Faculty of Computer Science and Mathematics
Cracow University of Technology
Krakow, Poland
filip.kruzel@pk.edu.pl

Daniel Grzonka
Faculty of Computer Science and Mathematics
Cracow University of Technology
Krakow, Poland
daniel.grzonka@pk.edu.pl

Andrzej Mycek
Faculty of Computer Science and Mathematics
Cracow University of Technology
Krakow, Poland
andrzej.mycek@pk.edu.pl

***Abstract*—Low-cost wireless medical-sensing prototypes must deliver samples within tight timing budgets while preserving signal morphology needed by downstream tasks. This paper presents a measurement study that combines a Node–Hub electrocardiography (ECG) prototype with a retained denoising benchmark for a wireless edge-sensing pipeline. The prototype combines an AD8232 single-lead front-end, a microcontroller node, Bluetooth/RFCOMM transport, and a Raspberry Pi hub for buffering, filtering, storage, and visualization. Retained artifacts support an aggregate inter-sample interval of about 5 ms against an 8 ms engineering target, but not stronger claims about jitter, packet loss, or per-stage latency. Offline benchmark results show that lightweight smoothing can reduce mean absolute error (MAE) by up to about 80% for additive white Gaussian noise (AWGN), but gives weak or negative gains under mixed artifacts and a real low-cost AD8232 trace. The main systems lesson is that lightweight filtering should be a signal-quality-aware runtime option rather than a default preprocessing step.**

***Index Terms*—wireless sensing, ECG, edge computing, performance evaluation, body sensor networks, signal quality, Bluetooth**

I. Introduction

Wireless and mobile sensing systems often combine a low-cost sensor, embedded processing, short-range wireless transport, and hub-side software into one measurement pipeline. For wearable and bedside electrocardiography (ECG), this pipeline is both a communication system and a measurement instrument: samples must arrive on time, but digital processing must not smooth away diagnostically relevant morphology such as the QRS complex and R peaks.

This work is motivated by two retained project artifacts: a prototype-oriented wireless ECG monitor and a later ECG denoising benchmark [1], [2]. We reframe those artifacts as a systems measurement study rather than as a new clinical method or a new denoising algorithm. The paper is therefore intentionally conservative: it reports what the retained evidence supports and explicitly identifies the missing instrumentation required for stronger wireless-performance claims.

Two research questions guide the study. RQ1 asks whether the retained Node–Hub prototype supports the engineering target of one sample every 8 ms. RQ2 asks whether simple filters that are attractive for edge or hub-side deployment remain useful when noise changes from AWGN-like disturbances to baseline wander (BW), electrode motion (EM), muscle artifacts (MA), power-line interference (PLI), and mixed low-cost acquisition failures. We also make all major acronyms explicit: exponential moving average (EMA), MAE, mean absolute percentage error (MAPE), and symmetric MAPE (sMAPE) are used with their standard meanings.

TABLE I
Acronyms used in the paper.

Term	Meaning	Term	Meaning
ECG	electrocardiography	ADC	analog-to-digital converter
AWGN	additive white Gaussian noise	BW	baseline wander
EM	electrode motion	MA	muscle artifact
PLI	power-line interference	EMA	exponential moving average
QRS	ventricular depolarization complex	MAE	mean absolute error
MAPE	mean absolute percentage error	sMAPE	symmetric MAPE

The contributions are: (i) a compact formal model of a low-cost wireless ECG edge pipeline, (ii) a provenance-aware separation between supported timing evidence and missing measurements, and (iii) a condensed timing/fidelity evaluation showing why reconstruction error alone is insufficient for runtime filtering decisions.

II. Wireless Edge Pipeline and Evidence Model

The evaluated system follows a Node–Hub architecture. The Node uses an AD8232-class single-lead ECG front-end, a microcontroller, an analog-to-digital converter (ADC), and Bluetooth/RFCOMM transmission. The Hub receives packets, buffers them in a first-in-first-out (FIFO) queue, applies lightweight processing, stores samples, and updates a plot [1]. Each record contains a timestamp, lead-off information, and the ECG sample value.

For a sample $s_k = (\tau_k, l_k, x_k)$, where τ_k is the node timestamp, l_k is lead-off status, and x_k is the ADC-derived sample, the end-to-end delay can be decomposed as

$$T_{e2e} = T_{acq} + T_{tx} + T_{rx} + T_q + T_f + T_{io} + T_p, \quad (1)$$

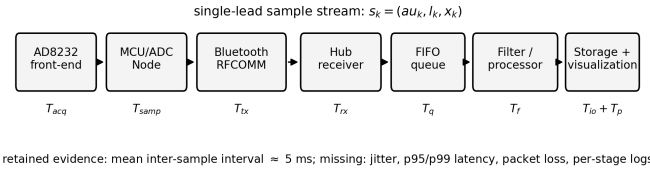


Fig. 1. Compact Node–Hub wireless ECG edge pipeline. The diagram separates acquisition, wireless transport, hub reception, buffering, filtering, storage, and visualization.

TABLE II
RETAINED EVIDENCE AND UNSUPPORTED PERFORMANCE CLAIMS.

Aspect	Supported by artifacts	Not supported
Timing	Mean inter-sample interval ≈5 ms; target 8 ms; plot refresh ≈3 s	Jitter, p95/p99 latency, packet loss, per-stage timestamps
Pipeline	AD8232–MCU–Bluetooth–Hub design; FIFO buffering; filtering/storage/plotting stages	Queue occupancy, reconnect latency distribution, CPU/energy per stage
Signal	MIT-BIH/NSTDB sweeps; five lightweight filters; one 10 s AD8232 trace	Multi-subject live validation, calibrated PLI, downstream R-peak F1 or heart-rate error

where the terms denote acquisition, wireless transmission, hub reception, queueing, filtering, storage, and plotting. The hub inter-arrival spacing is $\Delta_k = t_k^H - t_{k-1}^H$.

The prototype requirement can therefore be expressed as an average-rate condition $\overline{\Delta} < 8$ ms for the retained evidence. A stronger wireless-systems claim would additionally require the distribution of Δ_k, the probability of deadline misses $P(\Delta_k > 8\ \mathrm{ms})$, packet sequence gaps, reconnect intervals, and stage-specific timestamps. This distinction is central to the paper: the artifacts show that the design was fast enough on average, not that it provided bounded latency or reliable delivery under all channel conditions. The retained artifacts provide only an aggregate observation for Δ_k; they do not contain raw logs for jitter, tail latency, packet loss, or per-stage bottleneck attribution.

The 8 ms target originates from the prototype's engineering analysis: a faster heart rate shortens the time available to capture ECG wave components, so the prototype required sampling and transmission at least once per 8 ms [1]. This is not a clinical sampling recommendation; it is a design budget for a low-cost single-lead prototype. The reported average spacing of about 5 ms indicates aggregate feasibility, but Table II explains why it cannot be read as a complete wireless-performance characterization. In practical terms, the same average could arise from a stable stream, short bursts separated by stalls, or occasional lost samples masked by buffering. The current paper therefore treats the Node–Hub prototype as a useful feasibility artifact and uses the formal model to specify what a follow-up measurement campaign must log before stronger performance claims are made.

III. Evaluation Methodology

The offline benchmark uses MIT-BIH Arrhythmia record 118 as the clean ECG reference and the MIT-BIH Noise Stress Test Database (NSTDB) for ambulatory noise [3]–[5]. Additional synthetic noise includes AWGN, BW, and EM components. Signal-to-noise ratio (SNR) levels are 24 dB, 18 dB, 12 dB, 6 dB, 0 dB, and -6 dB. The retained live trace is a 10 s AD8232 recording acquired from one adult volunteer under non-clinical engineering conditions; it is used only as a qualitative acquisition-failure example, contains no directly identifying information, and is not used for clinical inference [6].

Five lightweight methods are compared using the retained sweep artifacts [2]: bin smoothing, EMA, an implemented overlap-window ensemble, running-line smoothing, and Savitzky–Golay filtering [7]. The ensemble method is not a clinically validated beat-aligned ensemble average; it is the implemented overlap-window procedure, which is why its failure under mixed artifacts is a useful systems result.

The primary numerical metric is

$$\mathrm{MAE} = \frac{1}{n}\sum_{t=1}^{n} |y_t - x_t|, \tag{2}$$

where x_t is the clean reference and y_t is the noisy or filtered signal. Relative improvement is $100(\mathrm{MAE}_{noisy} - \mathrm{MAE}_{filtered})/\mathrm{MAE}_{noisy}$. MAPE and sMAPE were not used because ECG samples can be zero or near zero. Because MAE is a reconstruction metric rather than an application metric, every conclusion is interpreted together with a QRS/R-peak morphology check. The retained artifacts do not permit reporting R-peak precision/recall, heart-rate error, or QRS duration error; those are explicitly listed as required future measurements rather than inferred from MAE.

The morphology check used in this retrospective paper is deliberately modest. We mark a filtered trace as problematic when the high-slope QRS region becomes flattened, when R-peak-like maxima shift or disappear visually, when low-frequency drift remains dominant, or when filtering introduces block-like steps and amplitude scaling artifacts. This check is not a substitute for detector-level evaluation, but it prevents the paper from equating a lower reconstruction error with a more useful ECG signal. It also connects the fidelity results to runtime policy: a hub should not activate a smoothing filter solely because it improves an isolated-noise MAE score.

IV. Results

A. Aggregate timing feasibility

The retained Node–Hub prototype met the aggregate timing target: the reported mean interval between successive ECG data points was about 5 ms, compared with the 8 ms requirement, and visualization refreshed in about 3 s, compared with a 5 s non-functional requirement [1]. This supports the statement that the basic acquisition, wireless transport, hub buffering, and display path was feasible under the tested conditions. It does not support stronger claims about latency tails, jitter, packet losses, or where the pipeline bottleneck occurs. Table III provides the complementary mixed-noise fidelity results.

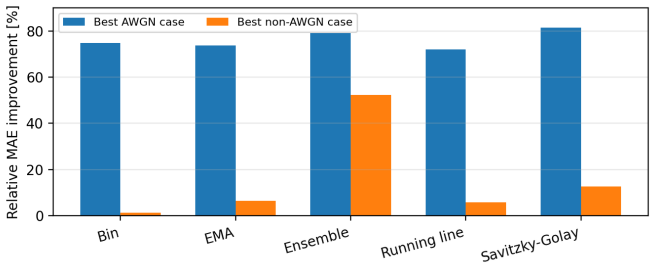


Fig. 2. Best observed relative MAE improvements on MIT-BIH record 118. AWGN cases look strong, but best non-AWGN improvements do not transfer to realistic mixed artifacts.

TABLE III
MIXED-NOISE MAE IMPROVEMENT ON MIT-BIH RECORD 118. NEGATIVE VALUES MEAN THAT FILTERING INCREASED MAE.

Method	0 dB	6 dB	12 dB
Bin smoothing	+1.17%	+0.39%	-1.56%
EMA	+1.82%	+0.46%	-3.10%
Ensemble	-206.64%	-272.95%	-473.36%
Running line	+2.33%	-3.71%	-19.68%
Savitzky–Golay	+2.90%	+0.93%	-3.80%

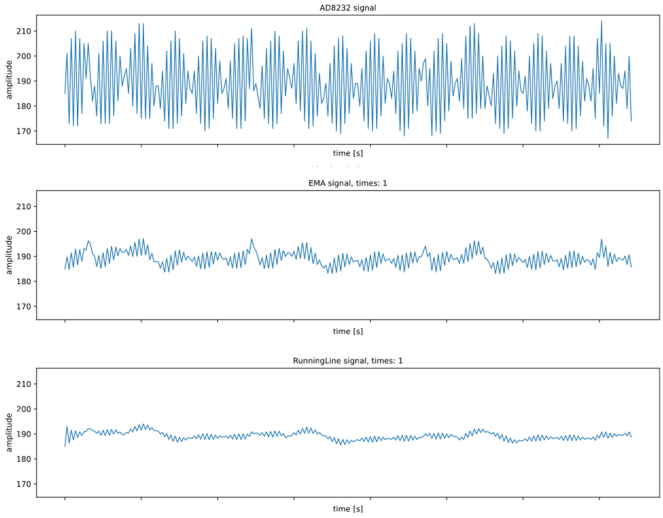


Fig. 3. Retained live AD8232 failure case. The plot is used as qualitative engineering evidence only: it illustrates periodic interference and the tendency of simple smoothing to flatten ECG morphology.

B. Signal fidelity under isolated and mixed noise

For isolated AWGN-like disturbances, all five lightweight methods can reduce MAE when tuned. The strongest best-case improvement on MIT-BIH record 118 is obtained by Savitzky–Golay filtering under −6 dB AWGN, with 81.5% relative MAE improvement. The implemented ensemble reaches 79.1%, bin smoothing 74.7%, EMA 73.7%, and running-line smoothing 72.1% in their best AWGN cases. Non-AWGN improvements are substantially less robust, as summarized in Fig. 2.

Mixed-noise results are the key negative result. Table III reports the selected follow-up parameters at 0 dB, 6 dB, and 12 dB. Bin smoothing, EMA, running line, and Savitzky–Golay produce at most small gains and can worsen MAE when the input is moderately readable. The implemented ensemble becomes unstable and amplifies/distorts the waveform. Therefore, a low reconstruction error in an isolated-noise sweep does not imply a robust runtime policy.

The live AD8232 trace reinforces this result. It contains strong periodic interference, and none of the simple methods produced a trace suitable for full ECG interpretation. EMA preserved the clearest R-peak-like spikes and may support approximate heart-rate estimation, but bin smoothing, running-line smoothing, and Savitzky–Golay removed or flattened important morphology, while the ensemble produced an amplitude-scaled near-constant output [2]. Fig. 3 illustrates the qualitative nature of this evidence: the upper trace is dominated by periodic interference, EMA leaves some peak-like structure, and running-line smoothing removes much of the sharp morphology. Thus, MAE and visual smoothness are not enough: a wireless ECG edge pipeline needs task-aware signal-quality logic.

V. DISCUSSION AND LIMITATIONS

The result most relevant to wireless and mobile systems is not that one filter is best. Rather, the pipeline exposes a coupling between sample-delivery feasibility and signal-fidelity risk. A throughput-only view would mark the prototype as successful because the mean inter-sample interval is below the engineering target. A signal-quality view shows that unconditional smoothing can destroy or distort the waveform even when it reduces scalar error in isolated cases.

Fig. 1 and the formal decomposition make the contribution less dependent on this exact ECG prototype. Any low-cost wireless sensing pipeline with acquisition, wireless transport, queueing, processing, storage, and visualization can use the same evidence hierarchy: average-rate feasibility is a first-layer claim, while jitter, p95/p99 latency, packet loss, reconnect behavior, queue occupancy, and per-stage energy require explicit instrumentation. In a follow-up prototype, each stage should export timestamps so that inter-arrival cumulative distribution functions (CDFs), tail latencies, packet losses, and per-filter runtime/energy can be audited instead of inferred.

A practical runtime policy would use three states rather than one fixed denoising configuration. In a normal state, the hub would apply only minimal processing and prioritize faithful delivery. In a random-noise state, detected from short-term variance or high-frequency residuals, a cheap smoother such as EMA or Savitzky–Golay could be enabled with conservative parameters. In a structured-artifact state, indicated by baseline drift, periodic interference, or electrode-lead warnings, the hub should either switch to a class-specific filter, request acquisition correction, or label the segment as low confidence. This policy view is the main systems implication of the denoising results.

The main signal-processing limitation is the lack of a retained downstream metric. Although related wearable ECG work increasingly evaluates R-peak detection and heart-rate utility [8], [9], the current artifact set supports only MAE and qualitative morphology inspection. The single MIT-BIH record

TABLE IV
MINIMUM LOGS NEEDED FOR A STRONGER FOLLOW-UP SYSTEM STUDY.

Log field	Purpose
Node sequence and timestamp	deadline misses, missing samples, clock drift
Hub receive timestamp	inter-arrival CDF, jitter, p95/p99 latency
Queue length and dequeue time	buffering pressure and back-pressure diagnosis
Filter start/stop timestamp	per-method runtime and deadline impact
Reconnect and lead-off events	wireless reliability and acquisition-quality context
Energy or CPU counters	cost of runtime filtering policies

and one 10 s live trace also limit generalization. We therefore position the paper as a measurement study and instrumentation agenda, not as clinical validation or a general ECG denoising benchmark.

VI. CONCLUSION

This paper presented a low-cost wireless ECG Node–Hub prototype and retained denoising benchmark as a wireless edge-sensing measurement study. The prototype met its 8 ms sample-delivery target at the aggregate level, with a reported mean spacing of about 5 ms, but the artifacts do not support stronger wireless-performance claims. Lightweight filters perform well for AWGN-like noise but are unreliable under mixed artifacts and real AD8232 interference. The practical implication is that filtering in wireless ECG edge systems should be conditional on signal quality and evaluated jointly through timing, morphology preservation, and downstream task utility. Future work should therefore combine stage-level wireless logs, packet-loss accounting, queue measurements, per-filter runtime/energy counters, and R-peak or heart-rate validation on a broader set of ECG records and live acquisitions.

REFERENCES

[1] T. Kot, "Bezprzewodowy monitor EKG," B.Sc. thesis, Cracow University of Technology, Faculty of Computer Science and Telecommunications, 2022, unpublished thesis; source of the Node–Hub prototype architecture and aggregate timing artifacts.

[2] T. Kot, "Analysis of algorithms for ECG signal processing and heart disease detection based on data from a high-precision analog-to-digital converter," M.Sc. thesis, Cracow University of Technology, Faculty of Computer Science and Telecommunications, 2023, unpublished thesis; source of the parameter sweeps, mixed-noise tests, and AD8232 trace.

[3] G. B. Moody and R. G. Mark, "The impact of the MIT-BIH arrhythmia database," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 3, pp. 45–50, 2001.

[4] G. B. Moody, W. K. Muldrow, and R. G. Mark, "A noise stress test for arrhythmia detectors," in *Computers in Cardiology*, vol. 11, 1984, pp. 381–384.

[5] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.

[6] World Medical Association, "Declaration of Helsinki: Ethical principles for medical research involving human participants," 2024.

[7] A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.

[8] H. J. Davies, G. Hammour, M. Zylinski, A. Nassibi, L. Stankovic, and D. P. Mandic, "The deep-match framework: R-peak detection in ear-ECG," *IEEE Transactions on Biomedical Engineering*, vol. 71, no. 7, pp. 2014–2021, 2024.

[9] E. Occhipinti, M. Zylinski, H. J. Davies, A. Nassibi, M. Bermond, P. Bachtiger, N. S. Peters, and D. P. Mandic, "In-ear ECG signal enhancement with denoising convolutional autoencoders," arXiv preprint arXiv:2409.05891, 2024. [Online]. Available: https://arxiv.org/abs/2409.05891