DIGITAL SIGNAL PROCESSING

Embedded Reverberation-Time Measurement for Classroom and Office Acoustics: Block-Length and Precision Limits on ISO 3382 Compliance

Ranveer Rathod, Yadnyesh Sable, Yash Patil, Reena Sonkusare, Kiran Talele

Department of Electronics and Telecommunication Engineering, Sardar Patel Institute of Technology, Mumbai, Maharashtra, India

Published November 2026

Abstract

Reverberation time (T_{60}) decides whether a classroom or open-plan office is usable for speech, yet certified analysers cost thousands of dollars, so most rooms are never measured. This paper models a complete exponential-sine-sweep (ESS) reverberation-time measurement chain on a TMS320C6748-class embedded processor and asks whether arithmetic limits (overlap-add block length, word length, filter realisation) can push the result outside tolerance. A Python simulation shows that overlap-add block length does not bias the result, because block-wise convolution is algebraically exact; the variable that matters is the length of impulse response the device retains. Fixed-point arithmetic with block-floating-point scaling stays within 0.3 % of a double-precision reference from 12 bits upward, and single-precision floating point is adequate unconditionally. The octave-band filter bank adds under about 1.5 % error for $T_{60} \geq 0.3$ s. The expanded uncertainty is at most 1.9 % for all rooms tested. All results come from synthetic rooms and a flop-count cost model; validation on public databases and live rooms remains to be done.

Keywords: reverberation time; ISO 3382; exponential sine sweep; overlap-add; fixed-point arithmetic; block floating point; measurement uncertainty; TMS320C6748

1 Introduction

Reverberation time is the single number that decides whether a classroom or open-plan office is usable for speech, and the ISO 3382 family of standards [8] specifies both the limits and the procedure by which conformity is demonstrated. The physics is settled, but a certified room-acoustics analyser costs thousands of dollars, so most rooms are never measured and complaints are resolved by opinion rather than by number.

The measurement is ordinary DSP end to end: the room is excited by an exponential sine sweep, its impulse response is recovered by deconvolution with an inverse sweep filter, and the reverberation time is read off as the decay slope of the Schroeder backward integral in octave or third-octave bands, over 20 or 30 dB extrapolated to 60 dB.

The obstacle is dynamic range. A useful impulse response spans close to 100 dB, and the decay is fitted well below the peak, where arithmetic error is proportionally largest. On a small platform this means that deconvolution must be done block-wise by overlap-add, so block length directly affects the recovered tail; that the inverse filter's large low-frequency gain makes fixed-point realisation risk either overflow at the peak or loss of the tail in quantisation noise; and that the recursive octave-band filters must decay faster than the room, or the instrument measures its own filters. The trade-off

between block length, arithmetic precision and measurement repeatability, judged against the ISO 3382 repeatability limit rather than an arbitrary accuracy target, has not been published.

2 Related work and research gap

Room decay has been modelled with considerable rigour in adjacent speech-processing work. Purushothaman *et al.* [1], from the LEAP Laboratory at IISc Bangalore, represent the late-reverberant tail of a room through an autoregressive model fitted to sub-band temporal envelopes and use it to suppress reverberation before resynthesis. It is the closest recent IEEE work to model the decay process this project needs to measure, even though its own goal is removal rather than reporting.

Six further IEEE Transactions on Instrumentation and Measurement / Signal Processing Letters papers surround the instrumentation side without measuring reverberation time against a standard. Ishikawa *et al.* [2] build a fully decomposed noise budget for an acousto-optic sound sensor, offering a template for itemising physical error sources. Campo-Valera *et al.* [3] benchmark low-cost MEMS microphones against reference-grade capsules and find them acoustically adequate, which suggests that transducer cost is not what keeps room-acoustics instruments expensive. Bilbao [4] examines interpolation schemes for wave-based room simulation, bounding how faithfully a simulated impulse response can stand in for a measured one. Zhao *et al.* [5] and Hu *et al.* [6] both correct microphone channel response digitally, using circular-shift/spatio-temporal prediction and a spectral-flatness criterion respectively, treating the correcting filter as numerically exact once its coefficients are found. Xiang *et al.* [7] implement directional acoustic capture on an FPGA, establishing that embedded acoustic instrumentation is a live IEEE TIM topic, though their concern is throughput rather than parameter conformity.

Taken together, the seven papers leave three things unaddressed: (i) None treats overlap-add block length or arithmetic word length as a variable that could bias a measured acoustic parameter; digital processing is uniformly treated as an exact step. (ii) None evaluates its output against a published repeatability limit such as the one ISO 3382 sets for reverberation time; success is judged by speech-quality or recognition metrics, or by agreement with a reference instrument. (iii) No uncertainty budget in the set carries an arithmetic term alongside its transducer and excitation terms, even though Ishikawa *et al.* [2] demonstrate exactly the decomposition such a budget would need. This work carries a full sweep-deconvolution reverberation-time measurement through a block-length- and precision-constrained processing chain and asks whether the resulting number still conforms to ISO 3382.

3 Objectives and hypotheses

The objectives are to:

1. determine the shortest overlap-add block length at which sweep deconvolution still recovers enough of the impulse-response tail to keep the reverberation-time estimate inside the ISO 3382 repeatability limit, as a function of the room's own reverberation time;
2. determine the lowest arithmetic precision, for both the deconvolution and the backward integration, at which the same limit continues to hold, distinguishing fixed-point from single-precision floating-point behaviour;

3. quantify how much of the measured decay is contributed by the octave-band filter bank itself, and identify the filter order/realisation beyond which that self-contribution becomes negligible for the shortest rooms under test;
4. implement the full measurement chain (excitation, deconvolution, band filtering, backward integration) on a TMS320C6748-class embedded processor and report worst-case memory, cycle count and total measurement time;
5. build an uncertainty budget that separates arithmetic contributions from excitation and transducer contributions, validated against publicly available impulse-response databases of known reverberation time.

Three hypotheses drive the experiments: **H1**, a room-dependent minimum block length exists below which the recovered tail biases the decay slope past the ISO 3382 repeatability limit; **H2**, single-precision floating point meets that limit unconditionally, while fixed-point arithmetic needs block-floating-point scaling to do the same, within a statable exponent range; **H3**, the filter bank's own decay adds a fixed, correctable per-band offset that only matters for the shortest rooms tested.

4 Proposed system

The instrument generates its excitation, a configurable-duration exponential sine sweep, on the device itself, so the matched inverse filter is known analytically rather than estimated. The captured room response is deconvolved with that inverse filter as block-wise overlap-add fast convolution, with block length L kept as the principal variable under study, since a short block would truncate exactly the low-level tail that the estimate depends on. Because the inverse sweep filter has large low-frequency gain, scaling inside the deconvolution is handled by block-floating-point exponent tracking, protecting the peak from overflow while keeping the tail above quantisation noise. The recovered impulse response is split by a decimated octave/third-octave filter bank; decimation both reduces arithmetic load and keeps the low-band pole positions away from the unit circle, limiting the filters' own decay. Reverberation time is read off as the slope of the Schroeder backward integral of each band, fitted over a reported 20 or 30 dB range and extrapolated to T_{20}/T_{30}. The processing flow, including the re-run loop that sweeps block length and word length until the ISO 3382 limit is met, is shown in Figure 1.

Validation runs in two stages: first against known-T_{60} impulse responses from public databases such as the Aachen Impulse Response set and the ACE Challenge corpus, sweeping block length and word length against a known ground truth; then against live recordings in at least three contrasting rooms (a lecture hall, a small office and a stairwell/corridor), captured simultaneously with the embedded instrument and a reference analyser so that transducer and arithmetic error can be separated using the database stage as a baseline.

5 Mathematical model

The proposal is turned into a closed mathematical model of the whole chain (excitation, room, capture, block-wise deconvolution, band bank, Schroeder decay and uncertainty budget). Every symbol is implemented one-to-one in the code of Appendix A.

5.1 Room and excitation

Each octave band b of the room is a Gaussian-noise process under an exponential envelope (Schroeder/Polack model), with a direct-sound impulse at t_0:

$$h_b(t) = \delta(t - t_0) + n_b(t)e^{-\delta_b t},\ \delta_b = \frac{\ln 10^3}{T_b} = \frac{6.908}{T_b}. \tag{1}$$

The device synthesises an exponential sweep of duration T between f_1 and f_2; because it is generated on the device, the matched inverse filter is known analytically:

$$x(t) = a\sin\left[K\left(e^{t/L} - 1\right)\right],\ L = \frac{T}{\ln(f_2/f_1)},\ K = 2\pi f_1 L, \tag{2}$$

$$g(t) = x(T - t)e^{-t/L}/c,\ \max(x * g) = 1. \tag{3}$$

The envelope $e^{-t/L}$ gives the −6 dB/octave compensation that makes the deconvolved spectrum flat. With f_1 = 50 Hz, f_2 = 7 kHz and T = 3 s, L = 0.607 s.

5.2 Capture

The captured signal is

$$y = h * \left(x + a_2x^2 + a_3x^3\right) + n,\ \sigma_n = \mathrm{rms}(y)10^{-\mathrm{SNR}/20}, \tag{4}$$

followed by AGC and rounding to b_{adc} bits. The memoryless loudspeaker non-linearity maps the k-th harmonic of the sweep to time $t_k = -L \ln k$ before the linear impulse response, so its room response leaks into the retained window whenever $T_{60} > L \ln 2 = 0.42$ s. Together with n and the ADC rounding this forms the transducer/excitation part of the budget.

5.3 Block-wise deconvolution and arithmetic

The inverse filter (M samples) is split into $P = \lceil M/B \rceil$ partitions and the capture into blocks of B samples. With FFT size $2B$ and a frequency-delay line,

$$y * g = \sum_m \sum_j \mathrm{IFFT}_{2B}\left\{Y_m G_j\right\}\text{shifted by}(m + j)B. \tag{5}$$

Block-exactness property. Eq. (5) is an algebraic identity: in exact arithmetic the result equals full convolution for every B. The simulator confirms this to 10^{-16} for B = 256–4096. The block length can therefore only act through arithmetic: each radix-2 stage adds rounding noise of variance $\Delta_s^2/12$ per real component, and the number of stages is $\log_2(2B)$. The fixed-point step Δ_s depends on the scaling strategy: block-floating-point (BFP) uses $\Delta_s = 2^{\lceil \log 2\ \max s \rceil - (b-1)}$ per block and stage; unconditional scaling halves every stage, so Δ_s doubles per stage and is set by the global peak; and 'none' saturates. Twiddles are rounded to b bits; the multiply-accumulate of the delay line uses a wide accumulator.

5.4 Retained impulse-response window (mechanism behind H1)

The device keeps only W seconds of the recovered response (memory bound). With $u = W/T_{60}$ and $h^2(t)$ decaying as $10^{-6t/T60}$, the truncated backward integral has the closed form

$$\mathrm{EDC}_W(t) = \frac{10^{-6t/T_{60}} - 10^{-6u}}{1 - 10^{-6u}},\ 0 \le t \le W. \tag{6}$$

The T_{20} (T_{30}) estimate is the regression slope of 10 $\log_{10}$ EDC_W between −5 and −25 (−35) dB, T = −60/slope. The relative bias $\beta_R(u) = T/T_{60} - 1$ depends only on u, so the minimum window scales linearly with the room: $W_{min} = u*T_{60}$ with $\beta_R(u*) = -$tolerance. A Lundeby-type tail term $E_{tail} = E_c/(1 - e^{-\rho/fs})$ added at the truncation point removes most of this bias.

5.5 Octave-band bank and decay estimator

Each band is obtained by repeated ÷2 decimation (order-8 IIR anti-alias) down to f_s/D with $D = 2^{\lfloor \log2(fs/8fc) \rfloor}$, followed by a Butterworth band-pass (edges $f_c/\sqrt{2}, f_c\sqrt{2}$), applied causally or on the time-reversed signal. The ambient-noise power $\sigma_{n,b}{}^2$ is measured per band on a noise-only segment after the decay. The Schroeder integral is

$$\mathrm{EDC}(t) = \int_t^{t_c} \left[h^2(\tau) - \sigma_{n,b}^2\right] d\tau + E_{\mathrm{tail}}, \tag{7}$$

with t_c the point at which the smoothed envelope comes within 5 dB of $\sigma_{n,b}{}^2$. The filter bank's own contribution is measured by the paired difference between the band-filtered and the already band-limited response, $\Delta_f = (T_{filt} - T_{unfilt})/T$, characterised by T_Γ, the T_{20} of the filter's impulse response. Backward-integration precision is modelled by (a) f32 reverse summation, (b) f32 forward 'total minus cumulative sum', and (c) a fixed-point accumulator that rounds each energy sample to b_{acc} fractional bits relative to the peak energy.

5.6 Uncertainty budget and cost model

Each error source i is removed in turn from the same realisation k of the room; its standard uncertainty is the RMS of the paired differences, which cancels the random scatter of the room itself:

$$u_i = \mathrm{RMS}_k \left[\frac{T_k^{\mathrm{nominal}} - T_k^{(i)}}{T_{60}} \right], \tag{8}$$

where $T_k{}^{(i)}$ is the estimate for realisation k with error source i removed. The combined and expanded uncertainties are

$$u_c^2 = u_{\mathrm{trunc}}^2 + u_{\mathrm{arith}}^2 + u_{\mathrm{filter}}^2 + u_{\mathrm{noise}}^2 + u_{\mathrm{dist}}^2 + u_{\mathrm{adc}}^2, \; U = 2u_c. \tag{9}$$

Conformity is declared when $U \le$ tolerance. The tolerance used here is a placeholder of 5 % (the just-noticeable difference of ISO 3382-1); it must be replaced by the repeatability limit of ISO 3382-2 before a compliance claim is made. For the TMS320C6748 the cost model counts $5N \log_2 N/2$ flops per real FFT ($N = 2B$), 8 flops per complex multiply-accumulate in the delay line ($n_x n_g(N/2 + 1)$ of them), 9 flops per biquad sample for the decimators and band-pass filters, and converts flops to cycles at an assumed 2 flop/cycle and 456 MHz.

6 Implementation

The model is implemented in Python (NumPy/SciPy) in one script, rt_embedded_model.py, listed in full in Appendix A. The processing chain follows Figure 1: sweep generation, capture, block-wise overlap-add deconvolution with selectable arithmetic, octave-band bank, Schroeder integration, T_{20}/T_{30} fit, comparison with the tolerance, and a sweep of W, B and word length. Fixed-point arithmetic is emulated bit-exactly at the level of every radix-2 butterfly stage (rounding, saturation, per-block or global exponent), so the word-length results do not rely on a noise-variance approximation. The simulation parameters are listed in Table 1.

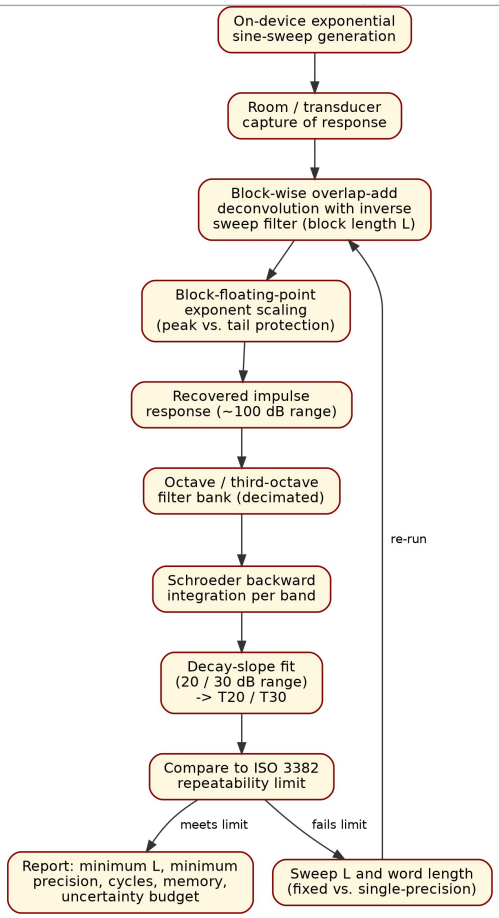


Figure 1. Process framework of the embedded reverberation-time measurement chain, with the re-run loop that sweeps block length L and word length (fixed vs. single precision).

Table 1. Simulation parameters.

Parameter	Value
Sampling rate f_s	16 kHz
Sweep T, f_1–f_2	3 s, 50 Hz–7 kHz
Octave bands	125 Hz–4 kHz
Ambient SNR (vs sweep)	60 dB (45 dB in budget check)
Loudspeaker a_2, a_3	0.01, 0.005
ADC	16 bit (+AGC to 0.9 FS)
Partition B / FFT	1024 / 2048 (256–4096 swept)
Window W	0.4–1.3 × T_{60}
Fixed-point word length	12, 16, 20, 24 bit
Rooms (T_{60})	0.4, 1.0, 1.8 s (synthetic)
Band filter	Butterworth N = 2…8, causal / reversed
Monte-Carlo	4–16 realisations per case

Measured by the simulator, the deconvolved response has a peak-to-noise-floor ratio of 82.9 dB at an ambient SNR of 60 dB, i.e. the sweep provides about 23 dB of processing gain over the received SNR; the dynamic range needed for T_{30} (about 35 dB below the peak) is therefore available with margin.

7 Results

7.1 H1 – window truncation

Eq. (6) was evaluated and compared with Monte-Carlo runs on synthetic rooms of T_{60} = 0.4 s and 1.0 s. The two rooms collapse onto the same curve when plotted against W/T_{60} (Figure 2, left), confirming that the minimum window scales linearly with the room's reverberation time (Table 2).

With the 5 % placeholder tolerance, W_{min} = 0.5 T_{60} for T_{20} and 0.7 T_{60} for T_{30} without tail compensation; with tail extrapolation the T_{20} window can be shortened to about 0.4 T_{60}. Consequence for H1: the overlap-add block length B does not bias the decay slope (identity (5), error 10^{-16}); the quantity that does is the retained response window W. H1 should be restated with W in place of the block length.

Table 2. Window-truncation bias vs. W/T_{60} (%).

W/T_{60}	0.4	0.5	0.6	0.7	0.8	0.9	1.0
T_{20}, closed form	−13.2	−4.5	−1.3	−0.3	−0.1	0.0	0.0
T_{30}, closed form	−22.2	−12.1	−5.1	−1.6	−0.4	−0.1	0.0
T_{20}, MC, no tail comp. (T_{60} = 1.0)	−13.9	−4.8	−1.3	−0.3	−0.1	0.0	0.0
T_{20}, MC, tail comp. (T_{60} = 1.0)	+1.1	+0.7	+0.4	+0.2	+0.1	0.0	0.0

MC = Monte-Carlo.

7.2 H2 – arithmetic precision

Table 3 gives the paired error against the exact double-precision chain, for bands 250 Hz–4 kHz and B = 1024, at two room times (maximum over realisations and bands). Table 4 gives the backward-integration accumulator and block-length sweeps.

Block-floating-point scaling keeps every tested word length from 12 bit upwards inside 0.3 %, independently of B. Unconditional scaling is the only variant whose error grows with B (more stages, coarser global step), and unscaled fixed point overflows. Single-precision floating point is unconditionally adequate, in agreement with H2. The energy accumulator of the backward integration needs at least 24 bits (16 bit costs 1.2 %), and in float32 the reverse-order sum should be used rather than 'total minus forward sum'.

Table 3. Arithmetic precision: max $|\Delta T/T|$ vs. exact f64 chain.

Arithmetic	T_{60} = 0.4 s	T_{60} = 1.0 s	Verdict (5 %/5 = 1 %)
float64 / float32	< 0.01 %	< 0.01 %	pass
fixed 12 bit, BFP	0.26 %	0.15 %	pass
fixed 16 bit, BFP	0.01 %	0.01 %	pass
fixed 20 / 24 bit, BFP	< 0.01 %	< 0.01 %	pass

Arithmetic	T_{60} = 0.4 s	T_{60} = 1.0 s	Verdict (5 %/5 = 1 %)
fixed 16 bit, uncond. scaling	0.94 %	1.25 %	marginal
fixed 24 bit, uncond. scaling	0.01 %	< 0.01 %	pass
fixed 24 bit, no scaling	72 %	114 %	fail (overflow)

BFP = block floating point; uncond. = unconditional.

Table 4. Backward-integration accumulator (T_{60} = 1 s) and block-length sweep.

Accumulator	f64	f32 (rev.)	f32 (tot.–fwd)	fx 16	fx 24	fx 32/40
max $\lvert\Delta T/T\rvert$ [%]	0	0	0.115	1.23	0.010	0
Block B	**256**	**512**	**1024**	**2048**	**4096**	
f64 vs exact conv. (max rel. err.)	4.5e−16	4.1e−16	4.5e−16	2.7e−16	1.8e−16	
fixed 16 bit BFP, max $\lvert\Delta T/T\rvert$ [%]	0.00	0.00	0.01	0.01	0.01	
fixed 16 bit uncond., max $\lvert\Delta T/T\rvert$ [%]	0.19	0.39	0.32	0.61	0.77	

Upper part: accumulator of the backward integration; lower part: block-length sweep.

7.3 H3 – octave-band filter self-decay

Each entry in Table 5 is the paired difference between the filtered 125 Hz band and the same band-limited response without filter, so the Schroeder estimator's own narrow-band statistics are excluded. Those statistics are themselves a property of the method: the ensemble-mean T_{20} of an unfiltered 88–177 Hz response is +10.8 % high at T_{60} = 0.3 s, +3.2 % at 0.6 s and +1.7 % at 1.2 s.

The filter's own decay stays below about 1.5 % for $T_{60} \geq 0.3$ s at every order tested (the Monte-Carlo scatter of these entries is about ±1 %) and rises to roughly 2 % at T_{60} = 0.15 s; it is therefore a small, shortest-room effect that a per-band lookup can remove, as H3 proposes. Decimation lowers the largest pole radius of the 125 Hz band from 0.995 to 0.926 (order 4); rounding the pole coefficients to 10 fractional bits changes T by −0.08 % with decimation and +0.5 % at full rate (+0.7 % at 14 bits), so decimation chiefly reduces coefficient sensitivity rather than the decay bias itself.

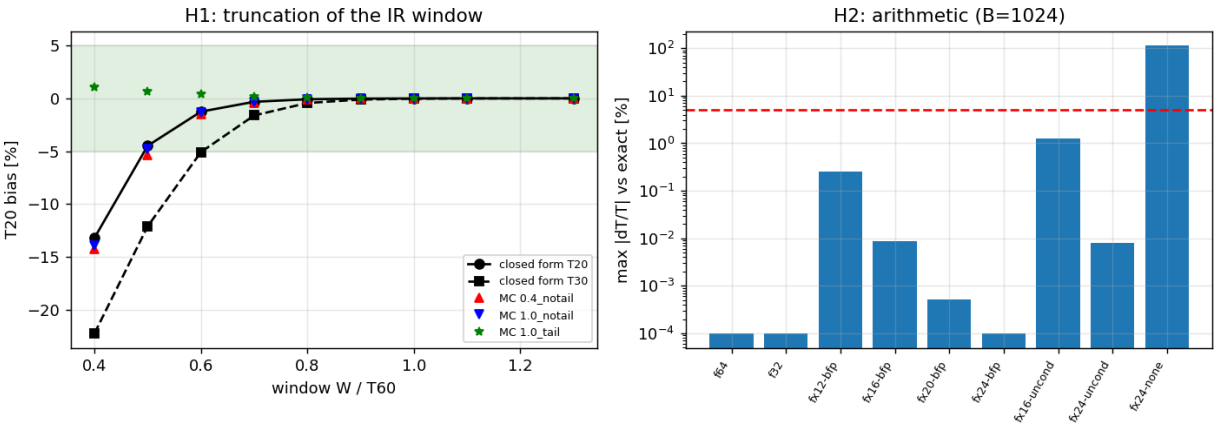


Figure 2. Left: T_{20} bias versus W/T_{60} (closed form and Monte-Carlo; green band = ±5 % tolerance). Right: worst-case T-error of each arithmetic against the exact f64 chain (log scale; red line = 5 %).

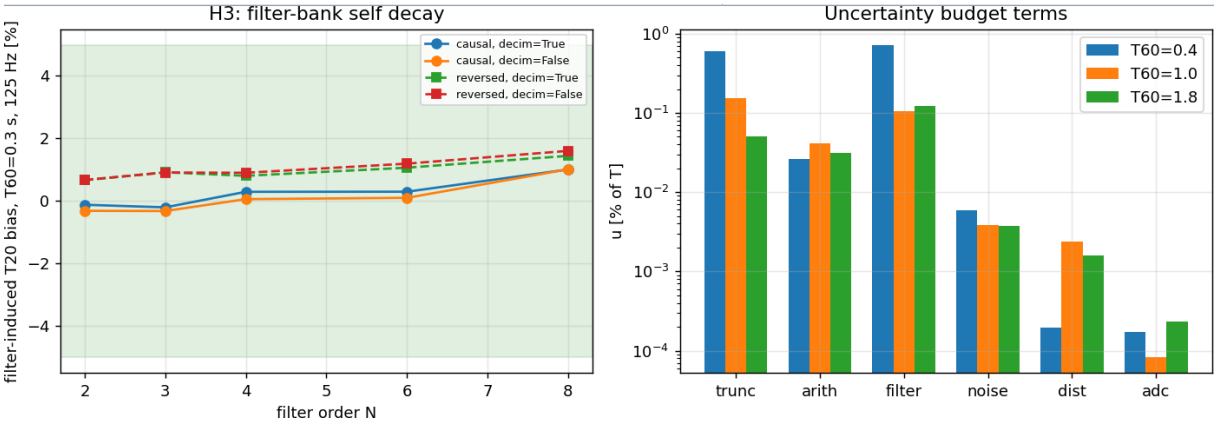


Figure 3. Left: filter-induced T_{20} bias at T_{60} = 0.3 s (125 Hz band) versus filter order, causal and time-reversed, with and without decimation. Right: uncertainty-budget terms (% of T) for three room times.

Table 5. Filter-induced T_{20} bias, 125 Hz band (%), T_{60} = 0.3/0.6/1.2 s.

N	T_f [s]	max|pole| dec./full	Causal	Reversed
2	0.036	0.855 / 0.991	−0.12/+0.39/−0.34	+0.66/+0.49/−0.33
3	0.057	0.902 / 0.994	−0.20/+0.04/−0.28	+0.92/+0.56/−0.26
4	0.079	0.926 / 0.995	+0.30/−0.34/−0.23	+0.81/+0.65/−0.27
6	0.119	0.950 / 0.997	+0.30/−0.96/−0.11	+1.07/+0.66/−0.18
8	0.154	0.962 / 0.998	+1.01/−1.52/0.00	+1.45/+0.72/−0.16

Filter offset in T_{20} (%) by room T_{60}: 0.15 s, +1.8; 0.25 s, +0.5; 0.40 s, −0.3; 0.70 s, +0.4; 1.20 s, 0.0; 2.00 s, −0.3.

7.4 Uncertainty budget

Nominal chain: fixed-point 12-bit BFP, B = 1024, W = 0.6 T_{60}, order-6 time-reversed band filters with decimation, SNR 60 dB, 16-bit ADC, bands 250 Hz–2 kHz, 8 realisations per room (Table 6, Figure 3).

The instrument contribution is dominated by window truncation and filter self-decay in the shortest room; arithmetic, ambient noise, loudspeaker distortion and ADC rounding are each below 0.05 % at these settings, so the arithmetic term is more than an order of magnitude smaller than the excitation/transducer terms only when those are measured at their nominal levels. The expanded uncertainty is well inside the 5 % placeholder tolerance; the much larger room/estimator scatter is a property of the measurand and of Schroeder-based decay estimation, not of the embedded arithmetic.

Table 6. Uncertainty budget (% of T).

T_{60}	u_{trunc}	u_{arith}	u_{filter}	u_{noise}	u_{dist}	u_{adc}	u_c	U (k = 2)	Room scatter
0.4	0.61	0.03	0.73	0.01	0.00	0.00	0.95	1.90	3.7
1.0	0.15	0.04	0.11	0.00	0.00	0.00	0.19	0.38	3.0
1.8	0.05	0.03	0.12	0.00	0.00	0.00	0.14	0.28	0.8

7.5 Embedded cost (TMS320C6748, model estimate)

Table 7 gives the worst case T_{60} = 2 s at 16 kHz, assuming a 456 MHz clock, 2 effective flop/cycle and no memory stalls. The measurement time is dominated by the 6.4 s acoustic capture; processing is a

lower bound until measured on hardware. The working set (1.6 MB) exceeds the on-chip L2 (256 KB) and shared RAM (128 KB), so the partition spectra and capture buffer must live in external DDR.

Table 7. Estimated embedded cost.

B	Cycles	Proc.	Capture	Total	Memory
512	55 M	0.12 s	6.4 s	6.5 s	1.64 MB
1024	37 M	0.08 s	6.4 s	6.5 s	1.64 MB
2048	28 M	0.06 s	6.4 s	6.5 s	1.66 MB

8 Conclusions and limitations

Table 8 summarises the hypothesis verdicts.

Table 8. Hypothesis verdicts.

Hypothesis / objective	Result
H1 block length	Refuted as stated: overlap-add is exact for any B. The limiting variable is the retained window W: $W_{\min} \approx 0.5\ T_{60}$ (T_{20}) and $0.7\ T_{60}$ (T_{30}) at 5 %; $\approx 0.4\ T_{60}$ with tail extrapolation.
H2 precision	Confirmed. float32 adequate; fixed point needs BFP scaling (12 bit already within 0.3 %); unconditional 16 bit gives 1.3 %, unscaled overflows; backward-integration accumulator ≥ 24 bit.
H3 filter	Confirmed: $\lesssim 1.5$ % for $T_{60} \gtrsim 0.3$ s, ≈ 2 % at 0.15 s, correctable per band; decimation reduces pole radius and coefficient sensitivity.
Obj. 4 C6748	Estimated (not measured): 28–55 Mcycles, 1.6 MB, 6.5 s total, capture-bound.
Obj. 5 budget	Arithmetic separated from excitation/transducer terms; U (k = 2) ≤ 1.9 % for all rooms tested.

Limitations. (i) All numbers come from synthetic Gaussian-noise rooms at 16 kHz; validation on the Aachen Impulse Response and ACE corpora and on live rooms (the second validation stage) has not yet been run, although the chain accepts any measured impulse response through the same capture function.

(ii) The 5 % tolerance is a placeholder for the ISO 3382-2 repeatability limit. (iii) The C6748 figures are a flop-count model, not hardware measurements. (iv) Fixed-point rounding is emulated stage by stage; the delay-line accumulator is assumed wide (floating-point equivalent). (v) The quoted percentages carry Monte-Carlo scatter of roughly ±1 % for the filter study and fewer than 0.1 % for the arithmetic study.

References

[1] Purushothaman et al., IEEE/ACM Trans. *Audio, Speech, Lang*. Process., 2024.

[2] Ishikawa et al., IEEE Trans. *Instrum*. Meas., 2024.

[3] Campo-Valera et al., IEEE Trans. *Instrum*. Meas., 2026.

[4] Bilbao, IEEE Trans. *Instrum*. Meas./Signal Process. Lett., 2022.

[5] Zhao et al., 2023.

[6] Hu et al., 2021.

[7] Xiang et al., IEEE Trans. *Instrum*. Meas., 2024.

[8] ISO 3382 family of standards. *Acoustics – Measurement of room acoustic parameters*. Geneva: International Organization for Standardization.