From Wobble to Wow: Real-Time Filtering Wizardry for Loadcells and Bridge Sensors

Author: Dr. Sanjeev Sarpal

Application note (ASN-AN031) **August 2025 (Rev 3)**

Synopsis

Loadcells and bridge sensors often suffer from noisy readings, slow settling times, and persistent oscillations, particularly in demanding industrial environments. When attempting to accurately measure load, strain, pressure, or torque in real-time bridge-sensor systems, engineers are often fighting three battles at once, namely:

1. Maintaining fast transitions to sudden changes.
2. Suppressing mechanical oscillations.
3. Reducing measurement noise without losing detail.

These requirements often contradict each other, as improving one usually worsens at least one of the others. Oscillations are typically non-stationary, i.e. their frequency and amplitude change over time, making them difficult to suppress without also degrading the filter's transient response. Conventional real-time filters, such as IIR Butterworth or Bessel, FIR equiripple, Savitzky–Golay, adaptive Gaussian, or even targeted notches, all demand multiple tuning parameters and are tricky to implement in changing environments.

This application note presents a robust, real-time digital signal processing (DSP) method that combines a Kolmogorov–Zurbenko (KZ) filter cascade with median filtering to deliver stable, responsive measurements with minimal tuning. The technique operates deterministically on three straightforward parameters, making it ideally suited for closed-loop control and Real-Time Edge Intelligence (RTEI) applications. Optimised for deployment on Arm Cortex-M based microcontrollers and other embedded targets, the method can be designed, evaluated, and implemented within minutes using the ASN Filter Designer and its integrated C SDK, ensuring that the deployed code matches the verified design.

1. Introduction

Precision loadcells and bridge sensors are indispensable in weighing systems, industrial automation, and high-accuracy measurement equipment. In real-time applications involving load, strain, pressure, or torque measurement, engineers are typically required to address three competing demands:

- Sharp step response for fast, accurate transitions.
- Oscillation suppression for stability.
- Low noise without sacrificing detail.

The difficulty lies in the fact that these requirements inherently conflict. Mechanical oscillations are often non-stationary, with frequency and amplitude varying over time. Conventional real-time filtering approaches — including IIR Butterworth or Bessel filters, FIR equiripple designs, Savitzky–Golay polynomial smoothers, adaptive Gaussian filters, and targeted notch filters — may reduce one problem but tend to exacerbate another. They require multiple parameters to tune, and maintaining stability under varying conditions can be challenging.

1.1. Measurement Challenges

Loadcells and bridge sensors convert small mechanical strains into electrical signals, typically through a resistive or capacitive Wheatstone bridge. In practice, these signals are subject to a variety of disturbance sources:

- Mechanical oscillations: introduced by supporting structures, mounting assemblies, or the measured load itself.
- Electrical noise: from nearby motors, variable-frequency drives, switch-mode power supplies, or radio-frequency sources.
- Sudden load changes: which demand rapid convergence of the measurement without excessive lag.
- Baseline drift: caused by temperature changes or slow environmental variations.

Designing a filter that balances the need for disturbance suppression with fast response time is a persistent challenge in such systems.

1.2. Limitations of Conventional Filtering Approaches

Over the years, many filtering methods have been applied to loadcell and bridge sensor applications. A summary of the most common approaches is given below:

- Moving average (MA) filter: excellent time-domain noise-cancellation performance, but poor frequency-domain performance, making it an inefficient low-pass filter. Although increasing the filter length improves cleaning performance, it also widens the transition width for step changes, leading to sluggish response, which is undesirable in most real-time applications.

- FIR filters: Gaussian and Savitzky–Golay (polynomial) filters can offer excellent results, but require extensive design work to achieve a practical balance between noise suppression, latency and responsiveness. Even then, their noise-reduction capability generally falls short of a moving average filter.

- IIR filters: Butterworth, Chebyshev, Elliptic, and Alpha–Beta filters can achieve effective noise reduction, but this typically necessitates sluggish responses to obtain sufficient attenuation. Furthermore, their non-linear group delay distorts waveform shape, which is problematic for precision measurement tasks.

- Notch filters (IIR and FIR): highly effective for fixed-frequency disturbances but of little use against varying resonances or broadband oscillations.

- Adaptive filtering: Adaptive Gaussian filtering is a promising technique that implements a lowpass filter whose bandwidth is dynamically adjusted by the characteristics of the signal. Although we have achieved good results with this method, the challenge is to accurately identify a ‘transition’ and ‘flat’ region, such that the filter’s sigma can be adjusted accordingly. This can be very tricky when operating conditions change, e.g., shifting noise profiles and new test setups that differ from the original design concept. Another method found in the literature is the Kalman filter. This method can only be effective when a model of the process is available and the noise is minimised. In the case of a loadcell application, unwanted resonances and slow mechanical sway greatly affect the performance, making it a poor choice for many applications.

- Machine-learning-based methods: typically non-deterministic in execution, requiring significant training datasets and computational resources during the design phase. The resulting ML model is inferenced based on data rather than scientific knowledge of the process being modelled, meaning that if any of the assumptions change, the model will perform badly. As with all ML solutions, latency is not deterministic, making them unsuitable for many mission-critical applications.

2. A hybrid filtering solution

As discussed in the previous section, there are many trade-offs with conventional approaches, especially under the varying dynamic and noise conditions typical of real-world loadcell and bridge sensor systems. In practice, engineers are often forced into lengthy cycles of trial-and-error tuning, only to end up with a compromise that still falls short in one or more performance areas — whether it be step response, noise suppression, or oscillation control.

To address these challenges, ASN developed a deterministic hybrid filter architecture that combines the strengths of median-based outlier suppression with the low-frequency handling capability of the Kolmogorov–Zurbenko (KZ) cascade. This approach achieves sharp transitions, stable output, and low noise without the complexity or unpredictability of adaptive or machine-learning methods.

The combined KZ-cascade and median filtering approach delivers a rare combination of stability, responsiveness, and robustness. In the steady state, the filter maintains exceptional stability, ensuring that readings remain consistent over long durations. When a step change occurs, it responds rapidly with minimal overshoot, enabling quick and accurate updates without introducing spurious artefacts.

It also provides strong suppression of both stationary and non-stationary mechanical oscillations, while the median filtering stage removes high-magnitude spikes and can attenuate certain low-frequency oscillations, such as slow mechanical sway or platform movement, when appropriately configured. Configuration is straightforward, with only three principal parameters to adjust, and the resulting performance is both predictable and deterministic — a critical attribute in control and safety-critical applications.

Key advantages include:

- Effective suppression of mechanical oscillations and electrical interference.
- Reliable spike rejection and reduction of certain low-frequency oscillations via median filtering.
- Minimal configuration effort with three primary parameters, L, R, and L_m
- Deterministic operation, making it suitable for closed-loop control systems — a key advantage over machine-learning-based approaches.
- Low computational cost, suitable for resource-constrained embedded platforms.

Taken together, these attributes make the method highly versatile, equally in laboratory instrumentation, industrial weighing systems, and embedded real-time control applications. By combining predictable real-time behaviour with excellent noise and oscillation suppression, it meets both the precision and timing requirements demanded by weighing and measurement control loops and by industrial automation systems.

This approach is also an example of Real-Time Edge Intelligence (RTEI) in practice. By embedding deep domain knowledge of loadcell and bridge sensor behaviour into a lightweight, deterministic DSP algorithm, the system achieves intelligent, context-aware processing without the complexity, uncertainty, or training burden of machine learning. Operating entirely at the edge, it delivers immediate, reliable results directly where the data is acquired.

The algorithm's efficiency and low memory footprint make it ideally suited to Arm Cortex-M based microcontrollers, such as STM32 devices, as well as other embedded processors used in IIoT/IoT applications. This enables engineers to implement high-performance filtering in compact, power-efficient hardware, avoiding the latency and connectivity overheads of offloading processing to remote systems. Moreover, the ASN Filter Designer C SDK can automatically generate fully optimised C code for direct integration, reducing development time and ensuring that the deployed implementation matches the verified design.

2.1. Kolmogorov-Zurbenko filter cascade

As discussed in section 1.2, the MA filter has excellent time-domain noise-cancellation performance, but poor frequency-domain performance, making it an inefficient lowpass filter. Although increasing the filter length improves cleaning performance, it also widens the transition width for step changes, leading to sluggish response, which is undesirable in most real-time applications.

A simple method of improving both the frequency domain performance, while retaining the excellent time domain performance, is by using a so-called Kolmogorov–Zurbenko (KZ) filter cascade, that in essence cascades R identical MA filters of length L into a single polynomial.

$$\mathrm{H(z)} = [\mathrm{A(z)}]^R = \left[\frac{1 + z^{-1} + z^{-2} + z^{-3} \;\ldots.. \; + z^{-(L-1)}}{L}\right]^R$$

A 3-section (R=3) Kolmogorov–Zurbenko cascade is shown below:

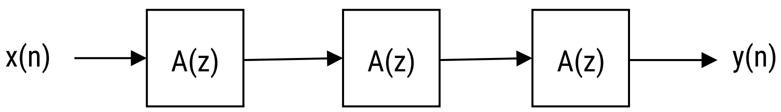

A key point here is that we can use a shorter MA filter with fast transition performance, and then improve its high-frequency noise cancellation performance by cascading it with copies of itself. Although the length of the resulting KZ filter will be equal to $L_{KZ} = R.L$, it will have very different properties to a single MA of length L_{KZ}, as shown below.

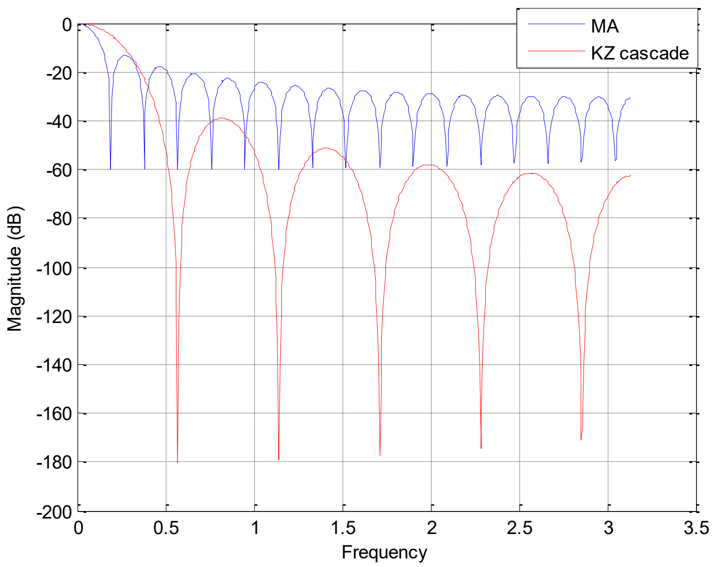

Figure 1 – Frequency response of Single MA (L=33) vs. KZ cascade (L=11, R=3)

Notice how the KZ filter has a much wider transition width (highly desirable for tracking fast transitions) and almost 55dB stopband attenuation, compared to the MA's narrow transition width and a modest 30dB attenuation. This unique property forms a large part of the 'wizardry' of the solution, achieving both excellent time and frequency domain filtering performance.

The KZ cascade is also extremely computationally efficient, i.e. each MA filter only requires one addition and one subtraction, rather than multiple MAC operations typically required with a traditional FIR filter.

As discussed in the previous section, although both the Moving Average (MA) and the Kolmogorov−Zurbenko (KZ) cascade are built from the same basic operation — averaging a set of consecutive samples within a window — they behave quite differently. The MA is mathematically optimal for reducing white-noise variance for a given window length, but it performs poorly in the frequency domain, allowing higher frequency artefacts to leak through its sidelobes. The KZ cascade on the other hand, sacrifices a small amount of this optimal white-noise reduction performance in exchange for much stronger stopband attenuation, making it far more effective at suppressing higher-frequency oscillatory disturbances and glitches in measurement data.

A KZ filter comprised of R sections of length M has an effective length of:

$$L_e = R(M-1)+1$$

Unlike the MA, the KZ's taps are not flat, so its noise reduction is slightly less than that of an MA with the same effective length. However, it offers much stronger stopband attenuation and is therefore better at suppressing oscillatory disturbances.

For practical engineering use, it is often sufficient to estimate the noise reduction factor directly from the KZ parameters, i.e. the number of cascaded sections, R and the per-stage window length, M. As such, a very accurate approximation (within $\pm 0.8\%$ for $2 \leq R \leq 6$ and $3 \leq M \leq 200$) is:

$$NR_{KZ} \approx 1.906 \left(\frac{R(M^2-1)}{12}\right)^{\frac{1}{4}}, \qquad \sigma_{out} \approx \frac{\sigma_{in}}{NR_{KZ}}$$

The equivalent moving-average length (the length of a flat MA achieving the same noise variance) can also be calculated as:

$$L_{eMA} \approx NR_{KZ}^2$$

The noise reduction performance of a moving average is well known to scale with the square root of its window length:

$$NR_{KZ} = \sqrt{L}\,.$$

This serves as the mathematical benchmark for white-noise suppression. By contrast, the KZ cascade distributes its smoothing power across several shorter moving averages. Because each stage works on the already-smoothed output of the previous one, the process becomes a *progressive convolution* rather than a single, longer convolution operation. In other words, the smoothing accumulates step by step, so the effective length is built up gradually rather than all at once. This cumulative structure means the KZ achieves slightly less noise reduction than an MA of the same effective length, but the staged convolution sharpens its frequency response, producing much stronger stopband attenuation. The result is a filter that is far more effective at suppressing higher-frequency oscillatory disturbances and glitches in measurement data.

For example, a KZ cascade with $R = 3$ and $M = 11$ has an effective length, $L_e = 31$ and achieves a white-noise reduction factor of $NR_{KZ} \approx 4.46$, compared to 5.3 for an MA of the same length. While its noise suppression is therefore only marginally lower, the KZ offers far superior stopband attenuation. In practice, this makes it much more effective at suppressing higher-frequency oscillations and other undesirable artefacts, and thus a stronger choice than the bread-and-butter MA for loadcell applications.

This numerical comparison highlights the trade-off between time-domain noise suppression and frequency-domain attenuation. To illustrate this, Figure 2 shows the Power Spectral Density (PSD) of white noise passed through successive stages of the KZ cascade. The plot illustrates how the progressive convolution sharpens the response: each stage adds further sidelobe suppression, improving stopband attenuation well beyond that of a single moving average. In practice, this enhanced attenuation is what makes the KZ cascade particularly effective at rejecting oscillatory disturbances in loadcell measurements.

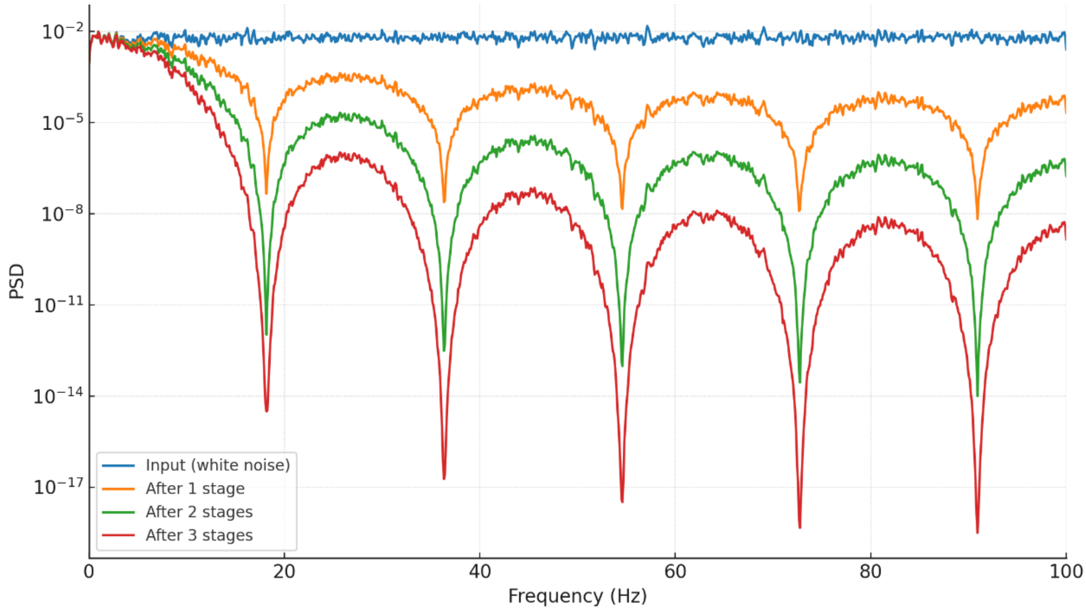

Figure 2 – PSD estimates using the Welch method of input versus outputs after 1, 2 and 3 stages.

Figure 2 highlights an important point: time-domain and frequency-domain performance measure different aspects of filter behaviour. A Moving Average is optimal for Gaussian white-noise suppression in the time domain, but the KZ cascade achieves nearly the same noise reduction while delivering much stronger attenuation of higher-frequency disturbances. For real-world sensor data, this trade-off makes the KZ cascade the more effective choice.

2.1.1.1. Beyond Gaussian White noise

In most textbooks, noise is assumed to be *white Gaussian noise*, i.e., normally distributed, and uniformly spread across all frequencies. However, in many industrial measurement systems, this assumption is unrealistic, as noise often exhibits other characteristics, such as:

- Coloured noise: Loadcells and bridge sensors frequently exhibit low-frequency drift, mains interference, or mechanical resonances. A MA filter is poorly suited to these conditions, as its weak stopband attenuation allows oscillatory components to pass through. By contrast, the KZ cascade provides substantially stronger suppression of such higher-frequency components.

- Impulsive and non-Gaussian noise. Many industrial environments exhibit noise closer to a Laplacian distribution, characterised by heavier tails than Gaussian. This reflects the prevalence of impulsive disturbances such as EMI spikes, switching transients, or quantisation glitches. Median filtering (as we will discuss in the next section) is particularly effective in such cases, since it is robust to outliers that would strongly bias a linear filter.

- Industrial signal conditions: Real-world signals typically contain a combination of drift, oscillations, and impulsive artefacts. No single filter can address all of these effects effectively. Hybrid configurations, such as median plus KZ filtering, are therefore the most practical and robust choice for industrial applications.

Thus, while *white Gaussian noise* provides a useful benchmark for analytical comparison, it does not adequately represent real operating environments. In practice, developers must consider the noise profile of their system: which frequency components dominate, whether impulsive artefacts occur, and how the noise evolves over time. The strength of the KZ cascade lies not only in reducing noise variance, but in reshaping the noise spectrum — thereby improving the stability and clarity of measurements under realistic industrial conditions.

2.2. Median Filtering

The median filter is a nonlinear filtering method that operates by sorting the data within its buffer and selecting the middle value as the output. Unlike linear filters (FIR, IIR) that perform a convolution, the median filter preserves the integrity of sharp transitions, making it particularly effective for removing impulsive noise and outliers without tainting valid signal edges. In the context of loadcell and bridge sensor applications, it excels at suppressing transient glitches, spurious spikes, and low-frequency oscillations such as those caused by slow mechanical sway, platform instability, or structural resonance.

More importantly, the median filter is especially effective in environments where noise is non-Gaussian. Many industrial noise sources follow distributions closer to the Laplacian model, characterised by heavy tails and a higher probability of large deviations than Gaussian noise. In such cases, conventional linear filters are strongly biased by outliers, whereas the median filter remains robust, providing reliable suppression of impulsive artefacts without distorting underlying signal features.

While the median filter alone is highly effective for eliminating impulsive noise and short-lived artefacts, many loadcell and bridge sensor systems are dominated by continuous broadband noise rather than isolated transients. Addressing this background noise efficiently calls for a complementary stage with strong low-pass characteristics, which is where the Kolmogorov–Zurbenko (KZ) cascade becomes invaluable.

The KZ cascade provides exceptional low-pass filtering, delivering strong attenuation of random noise and mid- to high-frequency disturbances while preserving the shape of valid signal transitions. However, because its passband extends down to DC, it will also pass very low-frequency oscillations, allowing slow disturbances from structural sway or quasi-static effects to creep through. In addition, occasional short-lived artefacts or high-amplitude transients can survive the KZ stage.

Therefore, cascading a median filter after the KZ output addresses these residual issues. Operating on a sorted buffer and selecting the middle value—rather than computing the arithmetic mean used in the KZ's moving-average process—the median stage rejects both slow, sustained deviations (such as structural sway), and any softened transient artefacts that slip through the KZ stage. This targeted suppression of very low-frequency oscillations and residual transients preserves legitimate step changes, making the median filter an ideal complement to the KZ cascade in precision measurement systems, where both low-latency and accurate feedback are essential.

The median filter window length is one of the three principal tuning parameters in the proposed hybrid design, alongside the KZ window size and the number of iterations. Together, these parameters define the balance between disturbance rejection and preservation of dynamic response. A shorter median window provides faster tracking of legitimate changes but less suppression of sustained oscillations, whereas a longer window enhances suppression at the cost of slightly increased delay. By adjusting these settings in ASN Filter Designer, engineers can visually optimise performance for their specific application before deploying the automatically generated C code directly to their target hardware.

2.3. Causality, Latency and Real-Time Filtering

In real-time systems, filters must be causal, meaning they can only operate on past and present samples. This requirement inevitably introduces group delay (latency), since each filter stage depends on a buffer of input samples before producing an output.

- MA/KZ (linear FIR case) – A MA filter is a special case of an FIR filter, as all coefficients are equal, making its impulse response perfectly symmetric. This symmetry guarantees linear phase, with a group delay of $(\mathrm{M}-1)/2$. A KZ filter cascade, being a cascade of MAs, preserves this symmetry. Its overall impulse response is still symmetric (binomial-shaped rather than flat), so it is also linear-phase, but with a group delay $\mathrm{R}(M-1)/2$.

- Median (nonlinear case) – The median filter does not perform convolution and has 'no phase' in the strictest sense. However, it still operates over a symmetric buffer, and to make it causal, the window must be shifted to use only past and present samples. This again produces a group delay of $(\mathrm{L}_m - 1)/2$ samples, even though the filtering operation is nonlinear.

For the hybrid filter described herein, comprised of a KZ cascade with R sections of length M followed by a median filter of length, L_m, the total group delay is:

$$\Delta = \frac{RM + \mathrm{L}_m - (\mathrm{R}+1)}{2}$$

For a representative configuration ($M = 31, R = 2, L_m = 99$) at a sampling rate of 1 kHz, the total group delay is 79 samples, corresponding to a latency of 79ms.

This latency is fully deterministic, making the proposed hybrid filter ideal for real-time control, where both speed and predictable timing are essential. By contrast, ML-based methods are inherently non-deterministic, as their inference introduces variable delays, making them unsuitable for mission-critical closed-loop systems.

This is precisely why the hybrid filter is such an excellent match for embedded sensing and control applications. It combines robust noise suppression with guaranteed timing, while the optimised C implementation via the ASN DSP library ensures a modest computational footprint. Although the median stage is inherently heavier than the KZ cascade, our efficient implementation keeps it lightweight enough for deployment on Cortex-M devices (see section 3.2 for a benchmark), making the solution both practical and powerful for a wide range of industrial control applications.

3. Designing the hybrid filter in the ASN Filter Designer

The hybrid filter cascade described herein can be very easily designed and customised in the ASN Filter Designer.

The KZ filter can be designed via the BLW tracker UI, that implements a KZ cascade and the median filter implemented via the H3 post filter.

1. Enable the BLW (baseline wander Tracker), and select **Tracking**

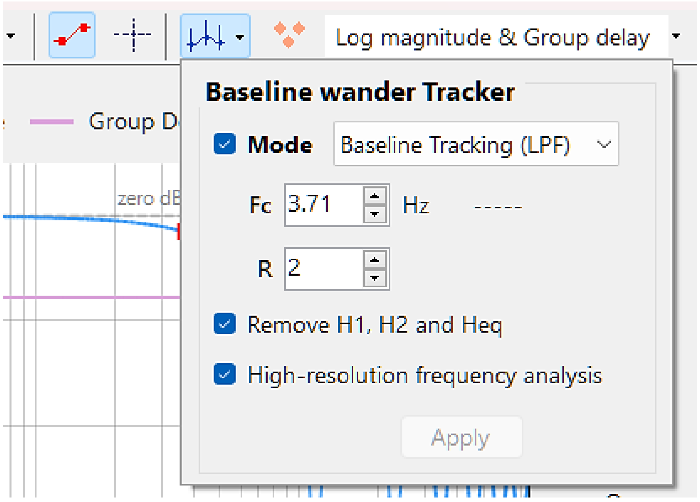

2. Fine-tune the cutoff frequency with the design marker.

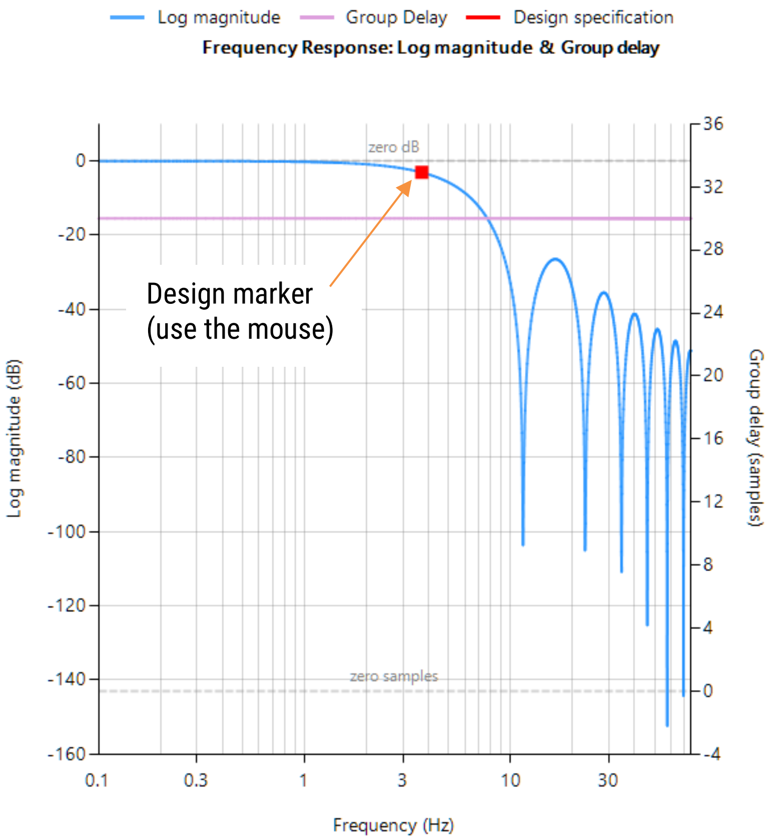

3. Open the signal analyser.

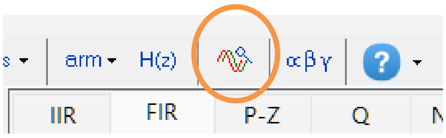

4. Enable the Median filter, setting $L_m = 99$ or so.

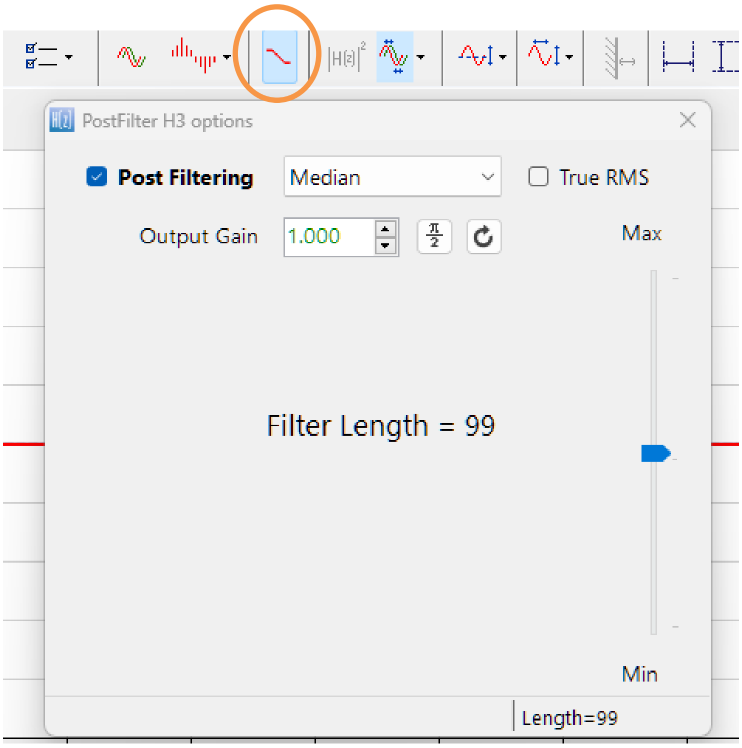

5. Enable Auto time alignment (to align the input and output data)

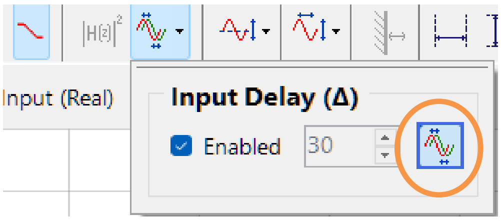

6. Drag and drop a CSV datafile onto the analyser canvas. If you do not have a CSV file, `loadcell_ex.csv` is available in the `\Datafiles` directory.

3.1. Filtering performance on real test data

The filtering performance of the hybrid filter ($M = 31, R = 2, L_m = 99$) is shown in the following plots, where the input and output data have been time-aligned for analysis. As expected, the filter has cleaned the signal and suppressed all oscillations, while maintaining sharp transitions.

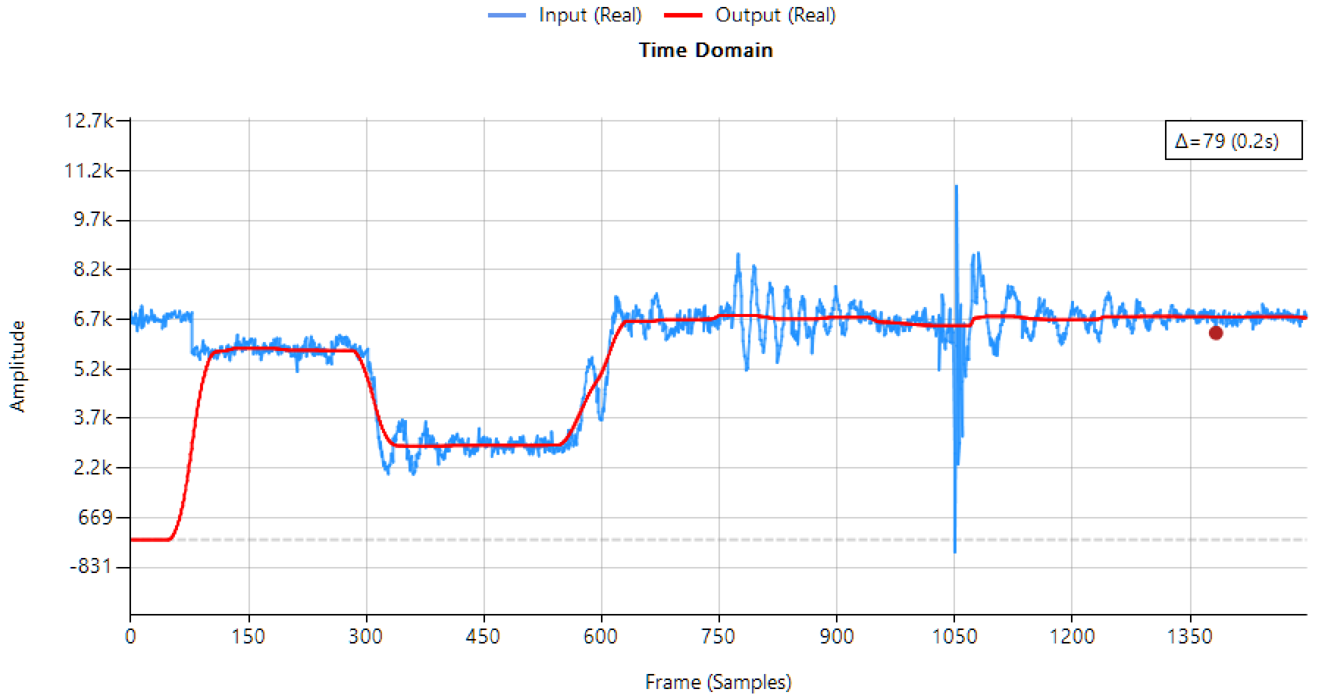

Figure 3 – filtering performance on the complete dataset.

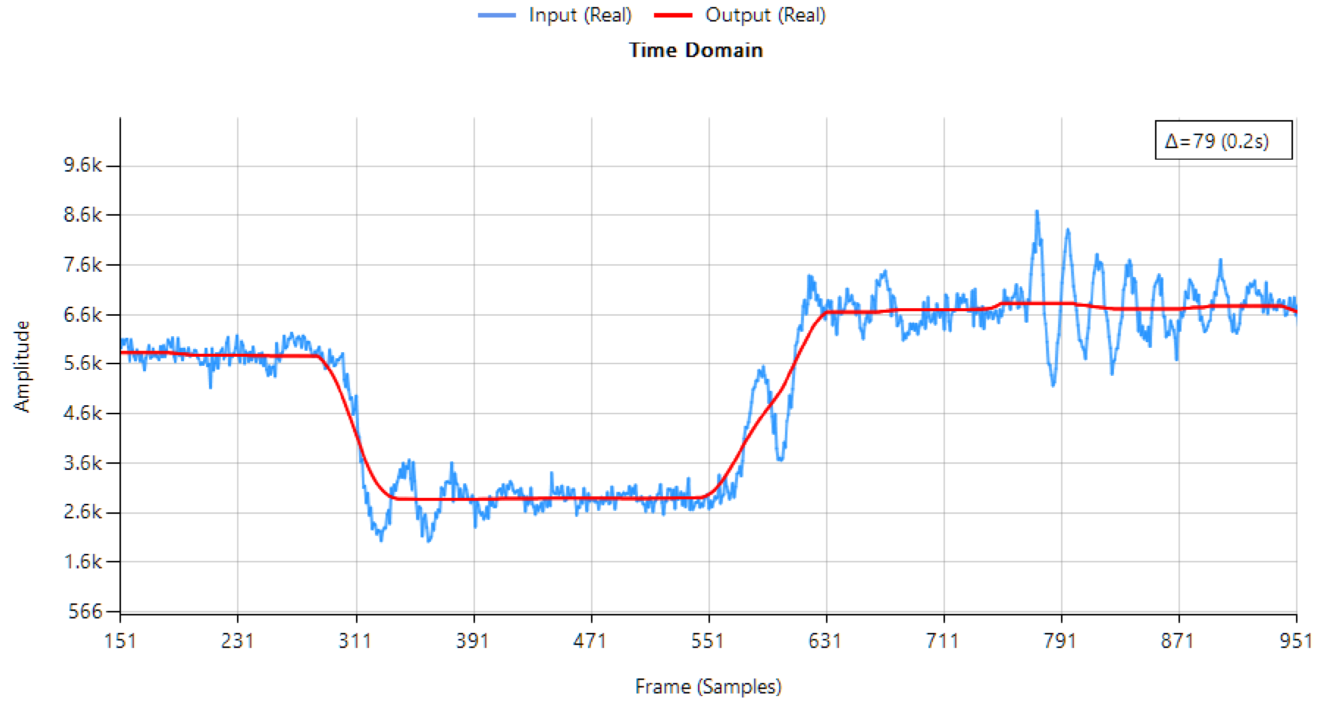

Figure 4 – zoomed view around the transitions. As seen, the tracking is fast, with no over/under-shoots, and all glitches and low-frequency oscillations are successfully suppressed.

3.2. Benchmark

Running the hybrid KZ–median cascade on an STM32F723 (Cortex-M7 core @ 216 MHz) with -O2 optimisation via the ASN DSP library using single-precision arithmetic, we achieved a processing rate of approximately **32 µs per sample**, equivalent to ~31.25 kS/s.

This performance comfortably meets the requirements of most practical applications. Typical loadcell and bridge sensor systems operate at sampling rates between 10 Hz and 1 kHz, leaving ample real-time headroom on an M7-class device. On more modest devices, such as the Cortex-M4F and Cortex-M33 (both with hardware floating-point support), throughput will be lower but still sufficient for these types of operations.

It should be noted that hardware floating-point support is not present on all Cortex-M4 devices; only the 'F' variants (e.g., Cortex-M4F) include this feature. Devices without such support, such as the Cortex-M3 and baseline M4, deliver reduced DSP performance, making them less attractive for DSP intensive filtering tasks. In practice, M4F and M7 cores are now the preferred platforms, offering the best balance of power efficiency and DSP capability. Thanks to advances in semiconductor technology, modern M4 and M7 devices achieve excellent low-power performance, in many cases matching or surpassing the power profiles of older M3-based designs.

More importantly, this benchmark demonstrates the efficiency of the ASN DSP library, showing that even advanced filtering schemes such as the hybrid KZ–median cascade can be deployed in real time on embedded Arm Cortex-M processors. This ensures not only high-quality filtering but also leaves ample compute budget for other control and signal-processing tasks.

4. Conclusion

The combination of a Kolmogorov–Zurbenko (KZ) cascade and a median filter offers a powerful yet computationally efficient solution for improving the quality of loadcell and bridge sensor measurements. The KZ stage delivers exceptional low-pass filtering, removing random noise and mid-to-high-frequency disturbances while preserving the shape of valid transitions. However, its passband extends to DC, meaning very low-frequency oscillations—such as those caused by structural sway or quasi-static disturbances—can still pass through, along with occasional high-amplitude transients.

Therefore, cascading a median filter after the KZ output addresses these residual issues. Operating on a sorted buffer and selecting the middle value—rather than computing the arithmetic mean used in the KZ's moving-average process—the median stage rejects both slow, sustained deviations (such as structural sway) that would otherwise be averaged into the signal, and any softened transient artefacts that slip through the KZ stage. This targeted suppression of very low-frequency oscillations and residual transients preserves legitimate step changes, making the median filter an ideal complement to the KZ cascade in precision measurement systems, where both low-latency and accurate feedback are essential.

Equally important, the hybrid design is controlled by just three intuitive parameters — the KZ window length, M, the number of sections, R and the median window length, L_m. This simplicity makes it straightforward for developers to tune performance, striking the right balance between noise suppression and dynamic response in real-world loadcell and bridge sensor applications.

By operating entirely at the edge, the hybrid solution exemplifies the principles of Real-Time Edge Intelligence (RTEI), delivering intelligent, physics-informed processing without reliance on machine learning. Its modest computational footprint make it ideally suited to Arm Cortex-M based microcontrollers and other embedded targets used in industrial and IoT/IIoT environments.

Using the ASN Filter Designer, engineers can design, verify, and export a fully optimised implementation in minutes. The integrated C SDK is highly efficient on Arm Cortex-M processors, such as the STM32, and ensures that the deployed code matches the verified design, streamlining the transition from development to production and enabling rapid, reliable deployment in both new and existing systems.

Author Bio

Dr. Sanjeev Sarpal is a RTEI (Real-Time Edge Intelligence) visionary and expert in signals and systems with a track record of successfully developing over 26 commercial products. He is a Distinguished Arm Ambassador and advises top international blue chip companies on their AIoT/RTEI solutions and strategies for I5.0, telemedicine, smart healthcare, smart grids and smart buildings.

Document Revision Status

Rev.	Description	Date
1	Document reviewed and released.	12/08/2025
2	Added noise reduction and benchmark sections	20/08/2025
3	Added latency and author bio sections	21/08/2025