CMA Filter: Robust, IEC-Compliant, Ultra-Low Latency Feature Tracking for IoT Applications

Author: Dr. Sanjeev Sarpal

Application note (ASN-AN032) **September 2025 (Rev 2)**

Synopsis

Reliable real-time tracking of sinusoidal waveforms is essential across many application domains, from power quality monitoring to IoT communications. International standards such as IEC 61000-4-30 (dip–swell detection), IEC 61000-4-7 (harmonic measurement), IEEE 1547 (distributed energy resources), and IEC 61851 (EV charging) demand ultra-low latency detection and feature extraction, often within a half-cycle.

Meeting these requirements with conventional FIR and IIR filters is challenging, as achieving a narrow transition band with an FIR demands very high orders, which increases both latency and computational cost. By contrast, an IIR filter can offer lower latency, but its non-linear group delay means that latency changes with frequency, and sharper transition bands come at the cost of overshoot and more sluggish time-domain performance.

This article introduces the **Complex Moving Average** (CMA) filter — a simple but powerful analytic bandpass that provides instantaneous amplitude, phase, and frequency with minimal latency. Validated against IEC-compliant dip–swell test datasets, the CMA demonstrates robust tracking of voltage envelopes and frequency variations with only 9.5 samples of latency, outperforming conventional IIR approaches in real-time compliance scenarios.

Lightweight to implement (three multiplications and two additions per sample), and supported by the ASN Filter Designer with direct C SDK export, the CMA is deployable on resource-constrained Arm Cortex-M devices as well as safety-critical Arm Cortex-R platforms with lockstep execution, ensuring deterministic real-time performance with fault tolerance. These qualities make the CMA a practical building block for applications in smart grids, EV chargers, IoT predictive maintenance, gas sensing, and wireless communications such as LoRa, ZigBee, Bluetooth, Wi-Fi, and 5G.

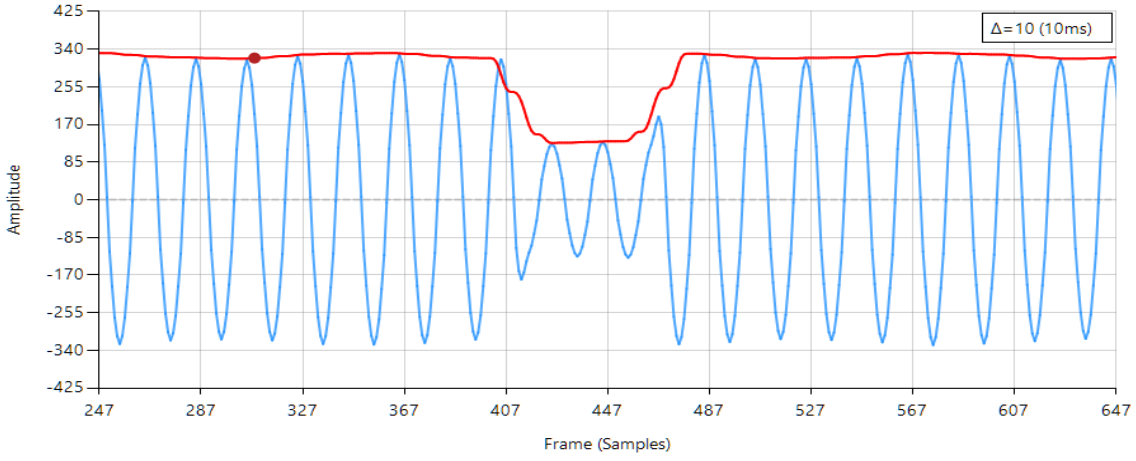

Figure 1 – Real-time envelope tracking on IEC 61000-4-7/-4-30 compliant powerline dataset with a dip event CMA filter output time aligned (10ms group delay) for easier visual comparison.

1. Introduction

In the real world, powerline (mains) waveforms are rarely ideal. Although nominally a 50Hz or 60Hz sine wave, they are continuously shaped by harmonics, switching events, and transient anomalies. Large disturbances can be triggered by everyday events — for example, a motor or industrial machine starting up and drawing a surge of current, which can momentarily pull down the grid voltage. Equipment faults, load imbalances, or sudden disconnections can equally cause the voltage to rise above nominal levels. The result is a voltage dip (a sudden reduction in RMS voltage) or a voltage swell (a sudden increase), both of which can trip sensitive equipment, disrupt processes, and, in the worst cases, destabilise the grid.

But dips and swells are not just a power system's problem. The same challenge arises anywhere a sinusoidal waveform must be monitored in real time. In wireless communications, the requirement is to track carrier envelopes and signal integrity across technologies such as LoRa, Zigbee, Bluetooth, Wi-Fi, and 5G, where reliable amplitude and frequency tracking is essential for maintaining link quality. In IoT measurement tasks such as predictive maintenance, vibration or even acoustic signals, must be tracked in real time to detect early signs of failure. Whether in smart grids, EV chargers, IoT sensors, or base-station antennas, the requirement is the same: extract and track amplitude, phase, and frequency reliably and with ultra-low latency.

International standards reflect this urgency. IEC 61000-4-30 requires Class A instruments to detect dips and swells in power systems within half a cycle (10ms at 50Hz). IEEE 1547 for distributed energy resources and IEC 61851 for EV chargers demand response within tens of milliseconds. In practice, measurement systems must therefore operate with ultra-low latency, extracting useful features from just a handful of samples.

Conventional filter approaches — Butterworth, Chebyshev, Gaussian FIR, and similar designs — can be used, but when the sampling rate is high and the passband is low, they quickly fall short. Both FIR and IIR band-pass filters still output a sinusoid, requiring rectification and smoothing that adds latency. FIRs demand high orders for sufficient narrowband selectivity, while IIRs, though more compact, suffer from non-linear group delay means that latency changes with frequency, and sharper transition bands come at the cost of overshoot and more sluggish time-domain performance. These drawbacks make them poorly suited to sub-cycle detection.

What is needed is a method that can deliver the essential properties of the fundamental waveform — amplitude, phase, and frequency — with minimal delay, while naturally rejecting DC offsets and harmonics. This article introduces the Complex Moving Average (CMA) filter as a suitable method: mathematically elegant, computationally efficient, and practical for deployment on embedded processors ranging from low-power Cortex-M devices to Cortex-R lockstep platforms for mission-critical grid and industrial applications.

The importance of robust and ultra-low latency feature tracking continues to grow as grids, vehicles, and industrial systems become more interconnected and standards-driven. Voltage dips and swells remain one of the most common causes of equipment malfunctions, from tripped EV chargers to disrupted industrial processes and unexpected downtime in automated plants. In smart grids and distributed energy systems, poor tracking can mask real disturbances or trigger false alarms, undermining both reliability and compliance.

International regulations reflect these concerns: IEC 61000-4-30 specifies strict performance for Class A instruments, IEC 61000-4-7 defines harmonic measurement protocols, and IEEE 1547 and IEC 61851 impose fast response requirements for grid interconnection and EV charging. These standards not only define how disturbances must be detected, but also set the benchmark for the design of measurement and control equipment worldwide. Against this backdrop, there is a pressing need for methods that combine mathematical simplicity with practical deployability, ensuring that real-time amplitude, phase, and frequency tracking can be achieved across diverse embedded platforms. The challenge is therefore not just theoretical accuracy, but delivering solutions that can be realised in practice on cost-sensitive devices. Addressing this challenge is the central focus of the solution presented herein.

2. The Moving Average Filter as a Prototype

The moving average (MA) filter is the simplest FIR low-pass prototype. For order N (where, $L = N + 1$), the impulse response is given by,

$$h[n] = \frac{1}{L}, \quad n = 0,1,2,3,\ldots L-1 \qquad (1)$$

with the transfer function:

$$H(z) = \frac{1}{L}\left(1 + z^{-1} + z^{-2} + \cdots + z^{-(L-1)}\right) \qquad (2)$$

Applying some lateral thinking to the computational challenge, we see that a much more computationally *efficient moving average filter* can be used in order to achieve the same result, namely:

$$H(z) = \frac{1}{L}\frac{1 - z^{-L}}{1 - z^{-1}} \qquad (3)$$

Upon initial inspection of the transfer function, it appears that the efficient Moving average filter is an IIR filter. However, analysing the pole-zero plot of the filter (shown on the right for L=8), we see that the pole at DC has been cancelled by a zero, and that the resulting filter is actually an FIR filter, with the same result as Eqn. (4).

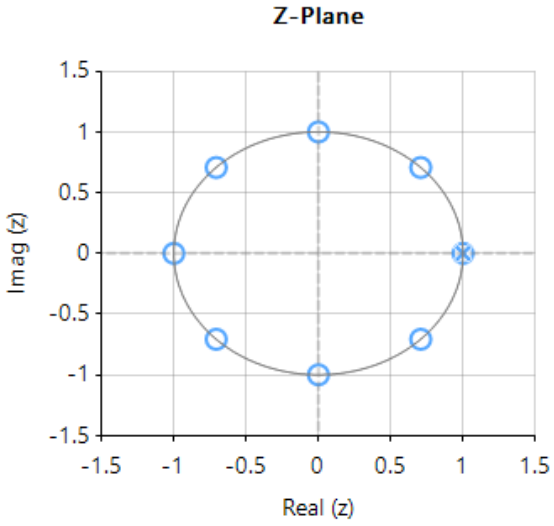

Analysing the transfer function, it can be seen that this prototype has regularly spaced notches in its frequency response at multiples of:

$$f = \frac{mf_s}{L}, \qquad m = 1,2,3.. \qquad (5)$$

making it a good candidate for supressing the harmonics of a fundamental.

2.1. Lowpass to Complex Bandpass Transformation

By applying the complex frequency shift transform, we can transform our low-pass prototype into a complex bandpass filter centred on f_0,

$$h_{bpf}[n] = \frac{2}{L} \cdot e^{j\omega_0 n} \qquad (6)$$

Where, $\omega_0 = 2\pi f_0/f_s$. Notice the explicit 2 gain factor in Eqn. (6). This normalisation is required because the complex shift produces a single-sided (analytic) response that suppresses the negative-frequency content; without it, the analytic output would be half the correct amplitude.

This can be readily seen via the two-term Fourier series of a real cosine of amplitude A, i.e.

$$A\cos(\omega_0 n) = \frac{A}{2}\, e^{j\omega_0 n} + \frac{A}{2} e^{-j\omega_0 n} \qquad (7)$$

In the ideal, infinite-length case, the spectrum consists of two impulses, each of amplitude $A/2$ located at $\pm\omega_o$

However, when viewed on a spectrum analyser with finite-length windows (see below), these impulses appear as narrow peaks with some spectral leakage. Nevertheless, the interpretation is the same: the total energy is split equally between positive and negative frequencies.

When forming the analytic signal, the negative-frequency term is suppressed, leaving only the positive-frequency exponential of amplitude $A/2$. To restore the true amplitude, we need to multiply by 2:

$$x_a[n] = Ae^{j\omega_0 n} \qquad (8)$$

The magnitude of this analytic signal is exactly A, i.e. the correct envelope of the original cosine wave.

NB. This is the essence of the analytic signal construction — halve the spectrum by discarding the negative frequencies, then double the positive frequencies so that amplitude and phase are preserved.

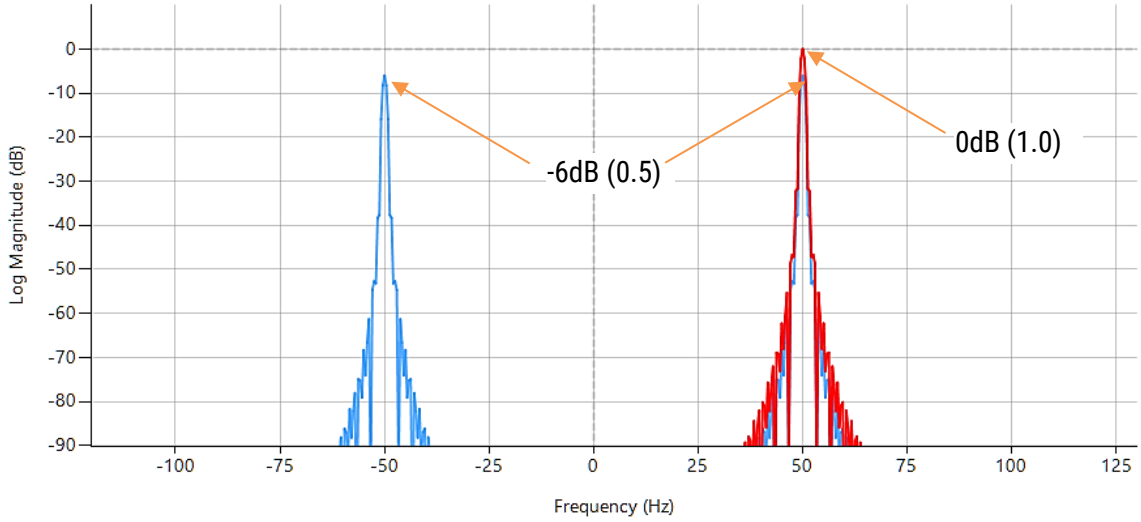

Figure 2 - Spectrum of a sinusoid before and after analytic filtering: The blue trace (input) shows a real sinusoid ($f_o = 50Hz, A = 1$. Its spectrum has equal positive- and negative-frequency lobes (at $\pm f_o$), each of amplitude 0.5 or -6dB. The red trace (output) shows the analytic signal: the negative-frequency lobe is suppressed, and the positive-frequency lobe is doubled. With the scaling factor of 2 applied, the spectrum now has a single lobe at amplitude 1, corresponding exactly to the true sinusoidal envelope.

2.2. Complex Moving Average (CMA) filter Definition

Combining the aforementioned concepts and simplifying, we can formally define the CMA as:

$$H_{bpf}(z) = \frac{2}{L} \cdot \frac{1 - e^{j\omega_0 L} z^{-L}}{1 - e^{j\omega_0} z^{-1}} \qquad (9)$$

Taking the inverse z-transform, we obtain the difference equation,

$$y[n] = e^{j\omega_0} y[n-1] + \frac{2}{L}\left(x[n] - e^{j\omega_0 L} x[n-L]\right) \qquad (10)$$

This may be further simplified for implementation, as $c_0 = e^{j\omega_0}$, $c_1 = \frac{2}{L}$, $c_2 = -\frac{2}{L}e^{j\omega_0 L}$

$$y[n] = c_0 y[n-1] + c_1 x[n] + c_2 x[n-L] \qquad (11)$$

What is striking about this result is its implementation simplicity. After all of the theory, $H_{bpf}\ (z)$ reduces to just 3 complex multiplications and 2 complex adds per sample, regardless of L – making it an exceptionally lightweight solution for real-time embedded deployment. The only downside is that the implementation requires Complex arithmetic support. While many modern C libraries (e.g. CMSIS-DSP) provide basic functionality, the ASN Filter Designer C SDK library offers full native support for complex filters, so developers can deploy $H_{bpf}\ (z)$ to their application very easily.

As an FIR of length L, the group delay is: $GD = \frac{L-1}{2}$. For $f_s = 1kHz, L = 20$ this corresponds to $GD = 9.5$ samples or 9.5ms – within the half-cycle (10ms) detection window for 50Hz mains tracking as stipulated by IEC 61000-4-30.

The PSD (Transfer function Estimation) of $H_{bpf}\ (z)$ is shown below. Analysing the plot, it can be seen that while the filter achieves strong attenuation at DC and at $\pm f_o$, its overall stopband rejection is relatively poor – consistent with the well-known limitations of a moving average as a low-pass prototype. Nevertheless, as discussed in previous articles, the moving average has excellent time-domain noise-cancellation properties. Therefore, in reality, as long as inter-harmonic noise is minimal (as is often the case in power and communications systems), the developed CMA should be more than sufficient for the applications considered herein.

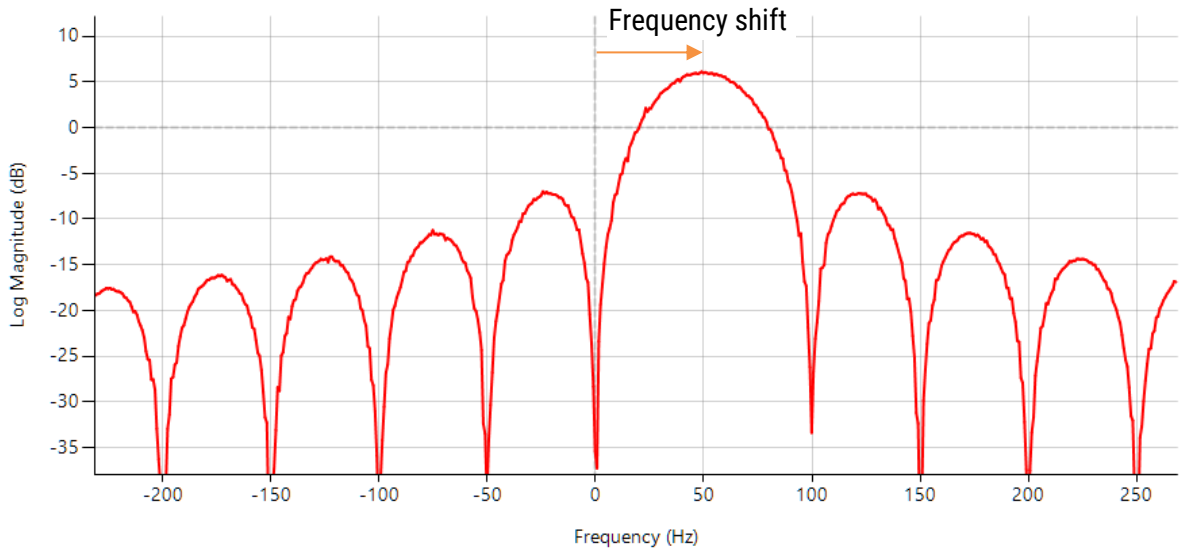

Figure 3 – PSD (Transfer function Estimation) of $H_{bpf}\ (z)$.

2.2.1. CMA Characteristics and extras

1. The CMA filter introduces zeros (notches) at DC and at harmonic frequencies $2f_0, 3f_0, 4f_{0,} \cdots$, giving strong suppression of unwanted components.
2. The passband peak is precisely centred at f_0, ensuring accurate tracking of the fundamental frequency.
3. 3 instantaneous cleaned features: amplitude, phase and frequency.
4. Linear phase and constant group delay – ideal for mission-critical applications that have strict latency requirements.
5. 3 complex multiplications and 2 complex adds per sample, regardless of L.
6. Despite using complex arithmetic, the workload is extremely lightweight, making this a very elegant and efficient solution for microcontrollers with floating-point support (e.g. STM32 Cortex-M4/M7/M33).
7. If better stopband attenuation is required, a Kolmogorov-Zurbenko cascade may be employed, which in essence cascades R identical MA filters of length L into a single polynomial. Although this will increase the overall latency, it's a very simple way of improving the MA's frequency domain performance, while retaining its excellent time-domain performance.

3. Designing the CMA filter in the ASN Filter Designer

The CMA filter can be very easily designed and customised in the [ASN Filter Designer] using the ASN FilterScript language. The filter can then be validated via the tool's signal analyser UI.

1. Set the Sample rate to 1000Hz.

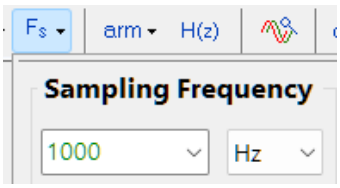

2. Open the ASN FilterScript IDE

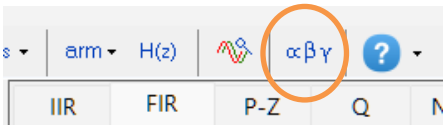

3. Paste the following script into the editor, and then Run it.

```
ClearH1;  // clear primary filter from cascade

Main()

fo=50;
L=round(fs/fo);
N=L-1;  // MA order (length L)
Hd=movaver(N,"numeric");
Hd=cplxfreqshift(Hd,fo,"symbolic");  // shift lowpass filter zeros
                                     // and make a bandpass
Num=getnum(Hd);       // get numerator coefficients
Den=1;                // get denominator
Gain=2*getgain(Hd);   // get gain
```

4. Open the signal analyser for testing.

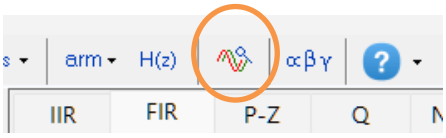

5. Drag and drop your test CSV datafile onto the canvas for analysis and click 'play' to start streaming the data.

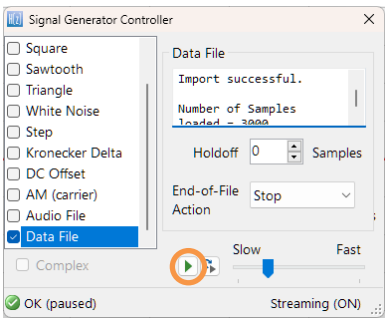

6. Tracking options
    a. Amplitude: set Output Math Function to `Abs`, alternatively, choose `RMS` if you want RMS amplitude.

    b. Frequency: set output math function to `Angle`.

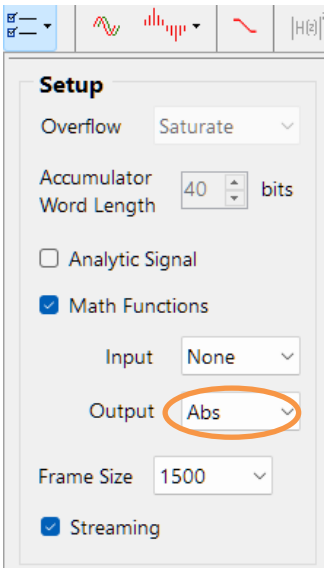

7. Setup time domain analytics
    a. Amplitude:

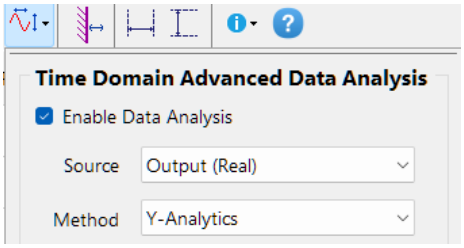

    b. Frequency:

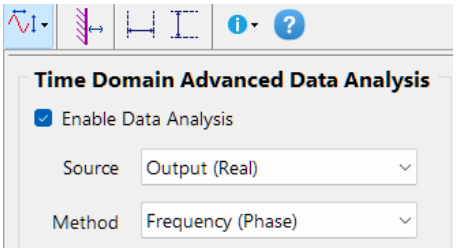

4. Results on a realistic dataset

An IEC 61000-4-7/-4-30 compliant powerline dataset was generated, consisting of a 50 Hz waveform sampled at 1 kHz with programmed dips, swells, and smooth edges, together with injected 3rd, 5th, 7th, and 9th harmonics at specified amplitudes. To better reflect real-world conditions, the dataset also included natural frequency sway, additive measurement noise, and a 2% flicker modulation.

This dataset was then processed with the proposed CMA filter. As shown below, the amplitude tracking follows the envelope very closely across both dips and swells, with only a small latency of 9.5 samples, demonstrating the suitability of the method for this type of application.

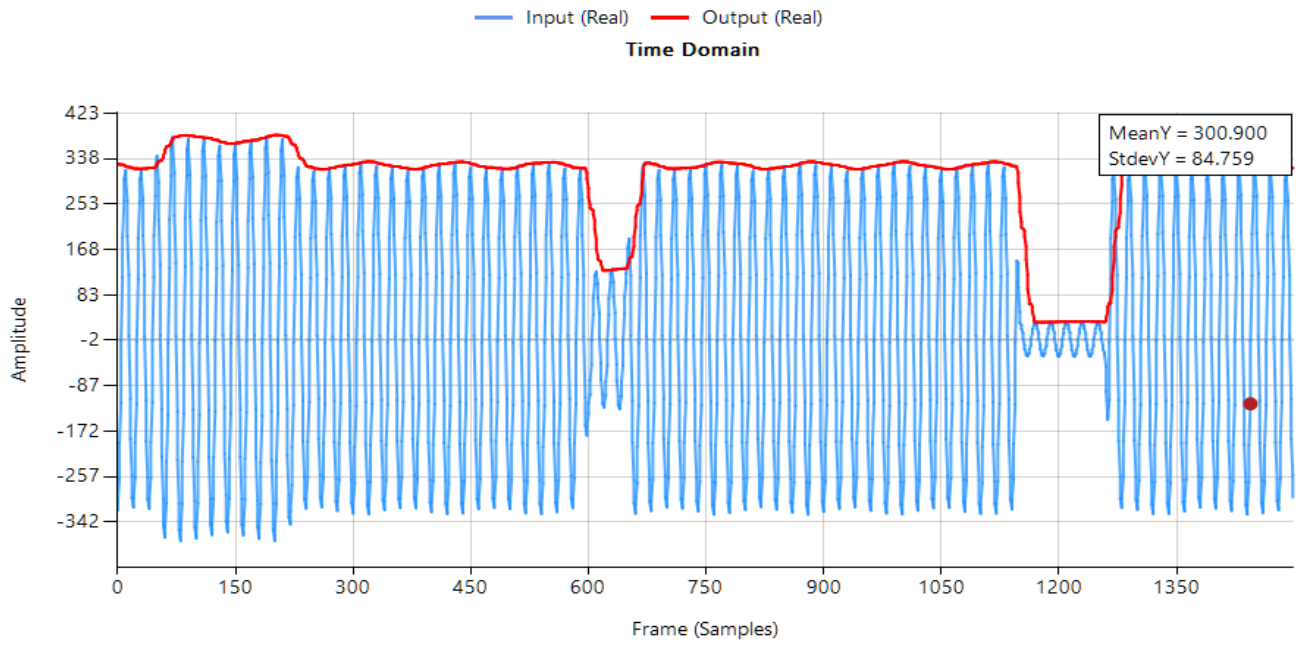

Figure 4: powerline envelope tracking

Likewise, frequency tracking is also very accurate, where a local, smoothed frequency estimate is obtained by calculating the gradient of short phase segments (within $\pm\pi/2$) and then scaling by $f_s/(2\pi)$.

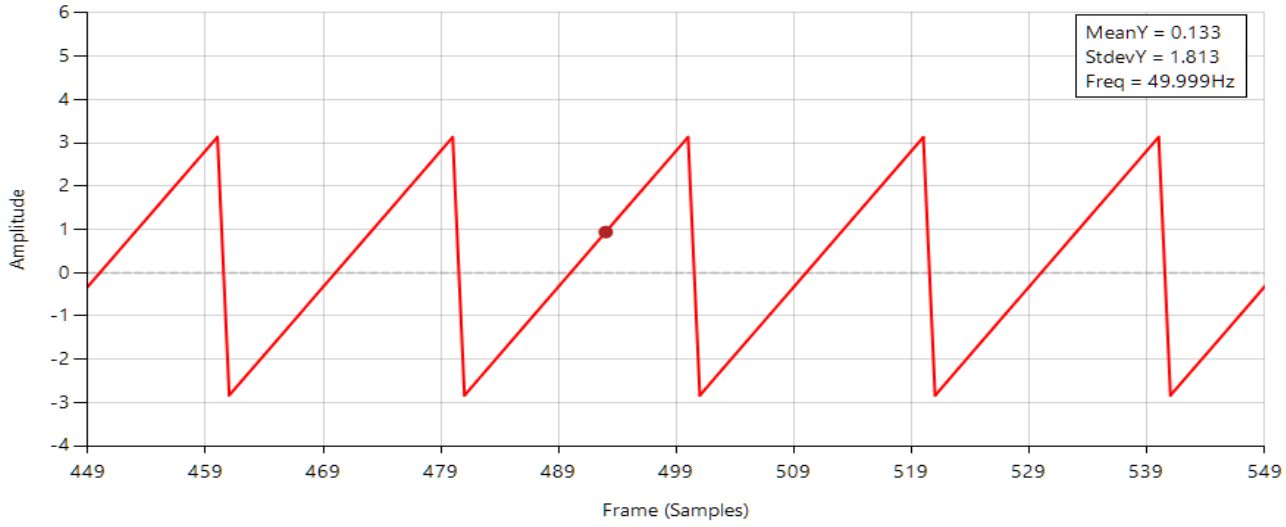

Figure 5: phase waveform of the 50Hz fundamental: Frequency calculated from scaled derivative phase data.

4.1. CMA vs. CIIR Butterworth filter

Generating a 10ms (half-cycle dip) at time step 300, we tested the performance of the CMA vs a complex 12th-order Butterworth IIR filter (CIIR).

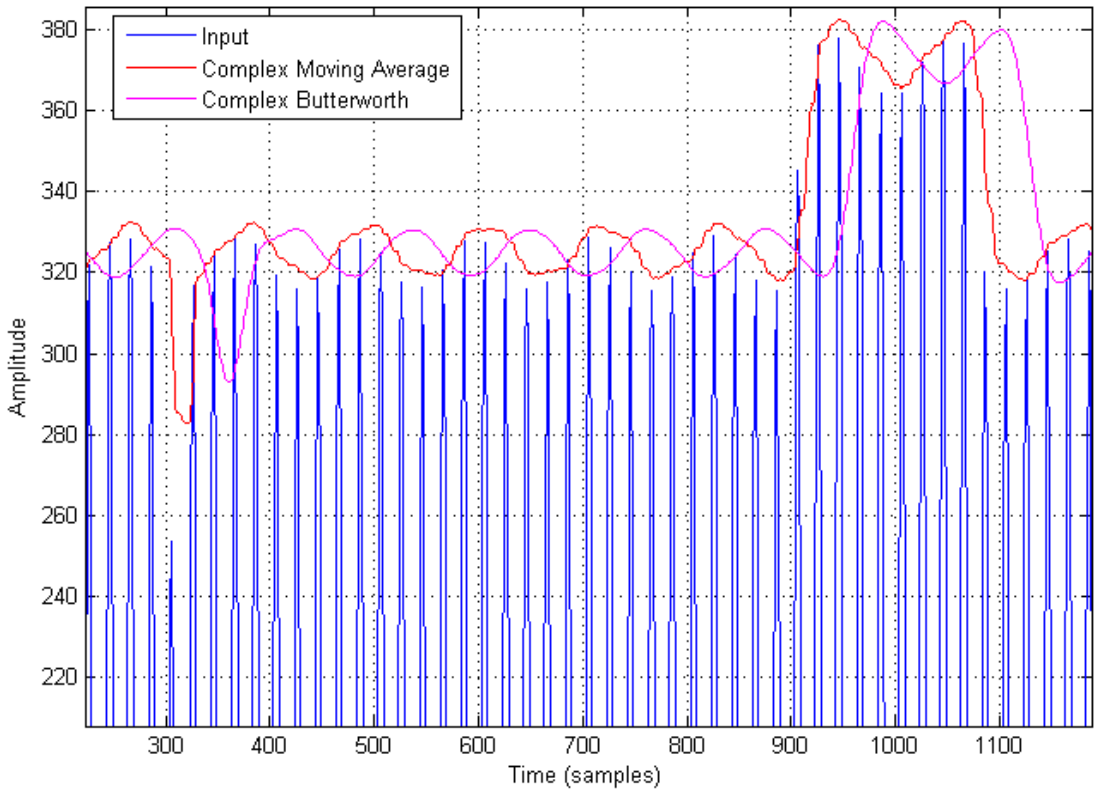

Figure 6: Butterworth CIIR vs. CMA

Analysing the plot, it can be seen that the IIR's high stopband attenuation delivers the strongest smoothing performance of the two filters. However, this advantage comes at a cost, as the CIIR exhibits higher latency and a more sluggish response than the CMA. It is also evident that the IIR fails to track the momentary dip around sample 300. These results demonstrate that a complex IIR is generally unsuitable for IEC-compliant dip–swell feature tracking.

4.2. Practical considerations

While the simulated dataset shown in Figure 4 illustrates the principle clearly, real-world IoT systems introduce a few practical considerations. The most important point is the sampling rate tolerance. Low-cost RC oscillators can drift by ±0.5–1%, which will shift the CMA's nulls away from their ideal positions and noticeably reduce harmonic suppression. This is why most manufacturers opt for a crystal oscillator with ±50 ppm tolerance (±0.005%), which offers a much more stable clock at low cost. However, in practice, the oscillator is divided down through prescalers or PLLs to generate the ADC sampling clock, so the effective sampling frequency may still be off by around ±0.01–0.1%.

A related issue arises when the design length $L = f_s/f_o$ is not an integer. In these cases, fractional values of L mean the filter no longer aligns perfectly with the fundamental, again slightly degrading tracking accuracy. These effects are usually minor but worth considering in system design. Typical mitigations include oversampling, designing around the expected clock tolerance, or applying resampling techniques when very high accuracy is needed.

Conclusion

This investigation has shown that the complex moving average (CMA) filter is a highly practical solution for ultra-low latency feature tracking in sinusoidal waveform applications. In wireless communications, the requirement is to track carrier envelopes and signal integrity across technologies such as LoRa, Zigbee, Bluetooth, Wi-Fi, and 5G, where reliable amplitude and frequency tracking is essential for maintaining link quality. In smart grids and EV chargers, it enables fast detection of dips, swells, and flicker for compliance and stability. The same idea extends to IoT measurement tasks such as predictive maintenance, where vibration or acoustic signals must be tracked in real time. It is equally applicable to gas sensing with infrared absorption waveforms, or to IEC sine-wave tests in biomedical instrumentation for verifying amplitude and frequency response. These examples highlight just a few of the areas where clean, real-time amplitude, phase, and frequency tracking is a core requirement.

The CMA achieves this by transforming a simple MA prototype into a complex analytic bandpass, yielding instantaneous amplitude, phase, and frequency with minimal delay — well within the IEC 61000-4-30 half-cycle detection requirement. Tests on IEC-compliant dip–swell datasets confirmed that the CMA can reliably follow voltage envelopes and frequency variations while maintaining robustness against measurement noise, flicker, and harmonic distortion.

Comparative analysis showed that although complex IIR (CIIR) filters (e.g. Butterworth, Chebyshev) offer higher stopband attenuation, their added latency and sluggish response make them unsuitable for sub-cycle dip–swell tracking. In contrast, the CMA achieves the right balance: computational simplicity, consistent linear-phase behaviour, and suitability for deployment on resource-constrained embedded processors. The implementation is extremely lightweight, requiring only three complex multiplications and two complex additions per sample, making it straightforward to realise in C on modern microcontrollers.

Crucially, the entire design, validation, and code generation flow can be carried out in the ASN Filter Designer, with direct export to the ASN C SDK. Since the SDK provides full floating-point and complex floating-point support in both single and double precision, developers can deploy the filter seamlessly on Arm Cortex-M devices for IoT and edge applications, or on Arm Cortex-R processors with lockstep execution where deterministic performance and fault tolerance are essential.

Overall, the results emphasise that the CMA is not just a theoretically elegant filter but a robust and developer-ready building block for IEC-compliant, real-time dip–swell detection in smart grids, EV charging, and distributed energy resources. In particular, it directly addresses the stringent requirements of IEC 61000-4-30 (Class A measurement within 10ms), IEC 61000-4-7 (harmonic measurement), and related grid-interoperability standards such as IEEE 1547 and IEC 61851. Beyond power systems, the same principles are equally relevant in IoT applications such as predictive maintenance, RF and telecoms for envelope tracking, and other sensor-based measurement tasks, such as gas concentration sensing. These examples highlight the versatility of the CMA as a general-purpose building block for real-time amplitude, phase, and frequency tracking across many domains.

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.